Imperial College
London

Particle Rate and Host Accelerator Beam Loss on the MICE Experiment

Adam James Dobbs

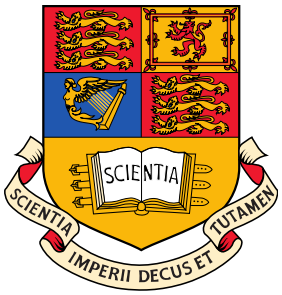

A Thesis submitted in partial fullfilment of the requirements of the degree of Doctor of Philosophy in the Department of Physics at Imperial College London

June 2011

Abstract

A study is presented of particle rates in the MICE Muon Beamline and their relationship to beam loss produced in ISIS. A brief overview of neutrino physics is presented, together with a discussion on the Neutrino Factory as a motivation for MICE. An overview of MICE itself is then presented, highlighting the need for a systematic understanding of the relationship between the MICE target parameters, ISIS beam loss, and MICE particle rate.

The variation of beam loss with target depth is examined and observed to be non-linear. The variation of beam loss with respect to the target dip time in the ISIS cycle is examined and observed to be approximately linear for dip times between 11.1 ms and 12.6 ms after ISIS injection, before tailing at earlier dip times. The variation of beam loss with particle rate is also observed to follow an approximately linear relationship from 0.05 V.ms to 4.7 V.ms beam loss, with a further strong indication that this continues up to 7.1 V.ms. Particle identification using time-of-flight data is used to give an insight into the relative abundances of each particle species present in the MICE beam. Estimates of muon rate are then produced as a function of beam loss. At a level of 2 V.ms beam loss $\sim$ 10.9 muons per spill for a 3.2 ms spill with negative $\pi \rightarrow \mu$ optics, and $\sim$ 31.1 muons per 1 ms spill with positive $\pi \rightarrow \mu$ optics are observed. Simulations using the ORBIT particle tracking code of the beam loss distributions around the ISIS ring, caused by the MICE target, are also presented and the implications for MICE running discussed.

Acknowledgements

This thesis could not have been produced in isolation and there are many people who deserve thanks. In particular I would like to express my heartfelt thanks to:

Ken Long, for giving me the chance to do a PhD in the first place, and for all the support and advice over the last three and a half years.

Linda Coney, for going above and beyond as a colleague, whether it be paddling in the Pacific or still managing to be passionate about physics even at 3 a.m. in the local control room.

Marco Apollonio, for all the support and always being willing to help.

All of my Imperial office mates, especially Dave Lee for being a good person to share a cup of tea with, and Andri Alekou for making things fun and for being someone who reminds me that banging your head on the desk is a normal part of the PhD process.

My family, Mam, Dad, Gran, Bill and Rhys, for far more than can be mentioned.

As a Christian I thank God for giving me a beautiful and complex world to poke, prod and measure. *"And God said to them, 'Be fruitful and multiply and fill the earth and subdue it...' "* - Genesis 1:28 (ESV).

Computing resources provided by STFC's e-Science facility.

The work presented here is my own unless otherwise stated, and is produced within the context of the international Muon Ionisation Cooling Experiment (MICE).

Contents

List of Figures

List of Tables

Chapter 1

Neutrino Physics

I have hit upon a desperate remedy.

Wolfgang Pauli

1.1 A Short History of the Neutrino

1.1.1 Discovery of the Neutrino

The neutrino, the name of which means "small, neutral one", was first proposed in 1930 by Wolfgang Pauli, in order to account for missing energy and momentum observed in the β-decay spectra of atomic nuclei. His hypothesis was not confirmed until 26 years later when the (anti-)neutrino was first detected by Frederick Reines and Clyde Cowan[1], who observed anti-neutrinos from a nuclear reactor using scintillators immersed in water and cadmium chloride. The anti-neutrinos were observed via their interaction with hydrogen nuclei according to:

$$\bar{\nu}_e + p \rightarrow n + e^+ \tag{1.1}$$

The emitted positron quickly annihilates with atomic electrons emitting two

γ-rays. The neutron is captured by a cadmium nucleus which subsequently emits another γ. These three γ-rays, detected by the photomultiplier tubes (PMTs), provided a characteristic signal for an anti-neutrino interaction. This signal was observed by Reines and Cowan and the existence of the (anti-)neutrino confirmed (however the discovery was not rewarded with a Nobel Prize until much later, in 1995).

Later experiments at Brookhaven[2] and CERN[3] in the 1960s determined that the neutrino produced in association with the muon (itself discovered in the 1930s in cosmic rays) was distinct from that produced in association with the electron, leading to the discovery of the muon neutrino as a distinct entity. The tau particle was then later discovered in the mid 1970s at SLAC[4] and its decay observed to have similar missing energy and momentum to that of β-decay, leading to the postulation of the tau neutrino. This was finally discovered in 2000 by the DONUT collaboration at Fermilab[5], the last particle in the Standard Model to be discovered to date (the search for the Higgs, the last proposed Standard Model particle, continues however). Thus the picture of three neutrinos to partner the three charged leptons, and six leptons to partner the six quarks was completed, as illustrated in Table 1.1. The question as to whether there exist still further generations of neutrinos, in particular "sterile" neutrinos which do not participate in the weak interaction but do mix with the known flavours, remains open (see for example [6]).

1.1.2 Helicity

A particularly unusual property of the neutrino is its *helicity*. Helicity is a property of particles that is defined to have a value of +1 if the particle's momentum vector is aligned with the spin vector (known as right-handed),

u $+\frac{2}{3}$ 2.4MeV	c $+\frac{2}{3}$ 1.27GeV	t $+\frac{2}{3}$ 171.2GeV
d $-\frac{1}{3}$ 4.8MeV	s $-\frac{1}{3}$ 104MeV	b $-\frac{1}{3}$ 4.2GeV

e -1 0.511MeV	μ -1 105.7MeV	τ -1 1.777GeV
ν_e 0 < 2.2eV	ν_μ 0 < 0.17MeV	ν_τ 0 < 15.5MeV

Table 1.1: The Standard Model fermions with their electric charge and masses.

and -1 if the momentum vector points in the opposite direction to the spin vector (known as left-handed). It is given by the expression:

$$\mathcal{H} = \frac{\boldsymbol{\sigma} \cdot \boldsymbol{p}}{|\boldsymbol{p}|} \tag{1.2}$$

where $\boldsymbol{\sigma}$ is the particle's spin vector, and $\boldsymbol{p}$ the particle's momentum vector. The first direct measurement of the helicity of the neutrino was performed by Goldhauber *et al* in the 1950s, using the electron capture of ^{152}Eu. The interaction proceeds according to:

$$^{152}\mathrm{Eu} + e^- \rightarrow \nu_e + ^{152}\mathrm{Sm}^* \rightarrow ^{152}\mathrm{Sm} + \gamma \tag{1.3}$$

An ingenious experimental set-up was used, exploiting the properties of both the europium and samarium, to measure only those photons emitted by the excited samarium whose helicity matched that of the emitted electron neutrino (see the original paper[7], or the account in [8]). The helicity of the photon is simply its circular polarization which was measured by Compton

scattering. It was found that $67\% \pm 15\%$ of the photons had a helicity of -1, while the expectation for the case of 100% left-handed neutrinos was a little under $\sim 75\%$. This result was quite general, and in all subsequent neutrino experiments only left-handed neutrinos, ν_L, and right-handed anti-neutrinos, $\bar{\nu}_R$, have been observed, never right-handed neutrinos, ν_R, or left-handed anti-neutrinos, $\bar{\nu}_L$. This is in contrast to both the charged leptons and the quarks, where particles of either handedness are observed.

Helicity however can only be a good quantum number if a particle is constrained to travel at the speed of light. If this is not so, then it is always possible to perform a Lorentz boost to a different frame of reference in which the momentum vector no longer aligns with the spin vector, so that the helicity is no longer Lorentz invariant. If neutrinos have a non-zero mass, and are so constrained to travel at less than the speed of light, then it always possible to find a frame in which the helicity of a left-handed neutrino flips, so that a right-handed neutrino is observed (and similarly for anti-neutrinos). Thus a non-zero neutrino mass implies the existence of the as-yet unobserved right-handed neutrinos and left-handed anti-neutrinos. These matters will become important for the discussion of the implications of neutrino mass in Section 1.3. For a more detailed treatment of helicity see for example [8].

1.1.3 Discovery of Oscillations

The discovery of the tau neutrino did not mark the end of neutrino physics, nor is helicity the neutrino's only unusual characteristic. One of the greatest discoveries in physics during the last decade or so has been that neutrinos possess a non-zero mass and undergo mixing between the different flavours. The Standard Model had, until relatively recently, assumed a zero neutrino

rest mass, which was implied by neutrinos seeming to have a definite helicity. The first evidence against this assumption of a zero mass, however, appeared in the Homestake Mine experiment of Ray Davis and John Bahcall, which has been measuring the flux of solar neutrinos since 1968[9, 10, 11]. The measured flux was found to differ significantly from the levels predicted by the Standard Solar Model. This became known as the "Solar Neutrino Problem".

Further evidence that some new physics was in play came later from experiments such as Super Kamiokande[12], a 50 kton water Cherenkov detector, which observed a deficit in the ratio of the muon neutrino to electron neutrino rate for neutrinos produced in the atmosphere by cosmic rays ("atmospheric neutrinos"), compared with that expected from theory. Further, the atmospheric muon neutrino rate was shown to be dependent on the azimuthal angle from which the neutrinos arrived. As muon neutrinos produced directly overhead of Super Kamiokande had a shorter distance to travel than those created on the other side of the globe, this again indicated that the rate depends on the path length between the source and the detector. These observations led to a pivotal announcement of evidence in favour of neutrino oscillations at the 18th International Conference on Neutrino Physics and Astrophysics (Neutrino98) in Takayama, Japan[13].

The Sudbury Neutrino Observatory (SNO)[14], another water Cherenkov detector based in Sudbury, Canada, also provided a milestone on the road towards the acceptance of neutrino oscillations, this time through the study of solar neutrinos. By using heavy water (D_2O) the SNO is able to observe neutral current (NC) neutrino interactions of the form:

$$\nu + d \rightarrow \nu + n + p \tag{1.4}$$

whereby a neutrino inelastically scatters from a deuterium nucleus (d), causing it to split into its constituent proton and neutron. The neutron may then be detected after an interaction with another deuterium nucleus, leading to tritium and a high energy γ-ray which is detected by PMTs. The NC interaction in Eqn. 1.4 occurs for neutrinos of all flavours, allowing the SNO to measure the total neutrino rate from all flavours, in addition to the rate from electron neutrinos only caused by the usual charge current (CC) interaction:

$$\nu_e + d \rightarrow e^- + p + p \tag{1.5}$$

and elastic scattering interaction:

$$\nu_e + e^- \rightarrow \nu_e + e^- \tag{1.6}$$

Using this technique the SNO collaboration demonstrated that the total flux from all neutrino flavours observed in the NC interactions was in excellent agreement with the Standard Solar Model, but still different from the electron neutrino rate observed in the CC interaction. As there is no known mechanism to produce muon or tau neutrinos in the Sun, and yet a muon / tau neutrino flux is observed, this demonstrated that neutrino flavour changes are occurring between the stellar core and the detector, solving the Solar Neutrino Problem[15].

Lastly the KamLAND experiment[16], a liquid scintillator experiment based at Kamioka, Japan, demonstrated conclusively that flavour oscillations occur as neutrinos travel through spacetime. Electron anti-neutrinos emitted by various Japanese nuclear reactors were used to provide signals along baselines of differing lengths (the average being about 180 km). The anti-neutrinos interact in the scintillator via inverse β-decay:

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{1.7}$$

The positron quickly annihilates dumping its energy in the scintillator. This then produces light that is detected by an array of PMTs. The neutron goes on to be captured on a proton, emitting a 2.2 MeV γ-ray. This again produces light that is detected later, giving a characteristic signal for the anti-neutrino. The anti-neutrino fluxes from the reactors were known and so the survival probabilities for the anti-neutrinos could be observed over the different baseline lengths. In 2005 results were published[17], demonstrating that the survival probabilities of electron neutrinos oscillated with baseline length over neutrino energy.

Taken together these experiments, amongst others, led to the formulation and acceptance by the particle physics community of the idea of neutrino flavour oscillations. In Section 1.2 a brief description is provided of this formulation based on quantum mechanics, followed by a discussion in Section 1.3 on the repercussions of neutrino oscillations and mass for particle physics and cosmology in general. An overview is then given in Section 1.4 of the current experimental status of neutrino physics, followed by a discussion of potential next generation oscillation experiments in Section 1.5. This introduces the idea of the Neutrino Factory, which is then discussed in more detail in Section 1.6, in turn introducing the concept and requirement of ionisation cooling. The MICE experiment, designed to demonstrate ionisation cooling, is then described in Chapter 2, providing the context for the work presented in this thesis.

1.2 Formalism of Neutrino Oscillations and Mass

1.2.1 General Formalism

Neutrino flavour oscillation is a quantum mechanical phenomenon whereby a neutrino that is created in a particular weak flavour eigenstate may subsequently be detected at a later time in another, different flavour eigenstate. This is possible because the flavour eigenstate, the wavefunction representing the neutrino immediately after it has undergone a weak interaction, is not the same as the mass eigenstate, the wavefunction which describes how the position of the neutrino evolves with time. They are, however, closely related, the one being a coherent sum of the other connected via a unitary matrix. Following the account given in [8]:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \tag{1.8}$$

$$|\nu_i\rangle = \sum_\alpha U^\dagger_{i\alpha} |\nu_\alpha\rangle = \sum_\alpha U^*_{\alpha i} |\nu_\alpha\rangle \tag{1.9}$$

with

$$UU^\dagger = 1 \tag{1.10}$$

Here $|\nu_i\rangle$ are the mass eigenstates, that is neutrinos with a definite mass ($i = 1, 2, 3$), and $|\nu_\alpha\rangle$ are the flavour eigenstates, that is, neutrinos with a definite flavour ($\alpha = e, \mu, \tau$). $U^\dagger$ is the Hermitian conjugate of U, and U^* the complex conjugate of U. For anti-neutrinos the states are related by:

$$|\bar{\nu}_\alpha\rangle = \sum_i U^*_{\alpha i} |\bar{\nu}_i\rangle \tag{1.11}$$

The time dependence of the mass eigenstates is given by:

$$|\nu_i(x,t)\rangle = e^{-iE_i t} |\nu_i(x,0)\rangle \tag{1.12}$$

The spatial dependence in the 1-dimensional case for a neutrino of momentum p is then given by:

$$|\nu_i(x,0)\rangle = e^{ipx} |\nu_i(0,0)\rangle \tag{1.13}$$

As neutrinos are produced (and detected) by weak interactions they are created in pure flavour eigenstates. By substituting Eqn. 1.12 and Eqn. 1.13 into Eqn. 1.8 the equation describing the spacetime evolution of a neutrino created as a flavour eigenstate $|\nu_\alpha\rangle$ may be arrived at:

$$\begin{aligned} |\nu_\alpha(x,t)\rangle &= \sum_i U_{\alpha i} |\nu_i(x,t)\rangle \\ &= \sum_i U_{\alpha i} e^{-iE_i t} |\nu_i(x,0)\rangle \\ &= \sum_i U_{\alpha i} e^{ipx} e^{-iE_i t} |\nu_i(0,0)\rangle \\ &= \sum_i \sum_\beta U_{\alpha i} U^*_{\beta i} e^{i(px - E_i t)} |\nu_\beta\rangle \end{aligned} \tag{1.14}$$

If the neutrino masses m_i are different it implies the phase factors $e^{i(px-E_i t)}$ in Eqn. 1.14 are also different. Therefore, due to the (x,t) dependence of the phase factors, the relative phases of the mass eigenstates will change as the neutrino propagates, leading to flavour oscillations. This may be thought of as each mass eigenstate state propagating with a different speed, compared with the case of a zero neutrino mass where all the states would

travel at the speed of light, c, and have a constant relative phase. Hence the observation of neutrino oscillations also implies the existence of non-zero neutrino masses. The transition amplitude to go from a flavour $|\nu_\alpha\rangle$ to a flavour $|\nu_\beta\rangle$ is then given by:

$$A(\alpha \to \beta) = \langle \nu_\beta | \nu_\alpha(x,t) \rangle = \sum_i U^*_{\beta i} U_{\alpha i} e^{i(px - E_i t)} \tag{1.15}$$

Considering the relativistic limit where $p \gg m_i$ and $E \approx p$, E_i may be rewritten as:

$$E_i = \sqrt{m_i^2 + p_i^2} \approx p_i + \frac{m_i^2}{2p_i} \approx E + \frac{m_i^2}{2E} \tag{1.16}$$

Setting $x = L$, the distance travelled by the neutrino, this leads to an expression for the transition amplitude of:

$$A(\alpha \to \beta) = \langle \nu_\beta | \nu_\alpha(x,t) \rangle = \sum_i U^*_{\beta i} U_{\alpha i} \exp\left(-i \frac{m_i^2}{2} \frac{L}{E}\right) \tag{1.17}$$

The probability of a transition $|\nu_\alpha\rangle \to |\nu_\beta\rangle$ is then given by the squared modulus of Eqn. 1.17.

1.2.2 Three Flavour Case

In the case of three neutrino flavours the transition probability derived from Eqn. 1.17 is given by[8]:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} &- 4 \sum_{i>j=1}^{3} \mathfrak{Re}\left(K_{\alpha\beta,ij}\right) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) \\ &+ 4 \sum_{i>j=1}^{3} \mathfrak{Im}\left(K_{\alpha\beta,ij}\right) \sin\left(\frac{\Delta m^2_{ij} L}{4E}\right) \cos\left(\frac{\Delta m^2_{ij} L}{4E}\right) \end{aligned} \tag{1.18}$$

where $K_{\alpha\beta,ij} = U_{\alpha i} U^*_{\beta i} U^*_{\alpha j} U_{\beta j}$ and $\Delta m^2_{ij} = m^2_i - m^2_j$.

The unitary matrix of Eqn, 1.18 is now the well known Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix[8], giving a formulation of neutrino mixing analogous to the mixing found in the quark sector:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{\text{PMNS}} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.19}$$

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}s_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}$$

$$\times \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha} & 0 \\ 0 & 0 & e^{i\beta} \end{pmatrix} \tag{1.20}$$

where $s_{ij} = \sin\theta_{ij}$, $c_{ij} = \cos\theta_{ij}$, and α, β, δ are CP violating phases. Oscillation experiments are insensitive to α and β, which are only non-zero in the case that neutrinos are Majorana particles (that is, identical to their anti-particles, see Section 1.3 below).

1.3 Neutrino Oscillations and New Physics

The implications for particle physics of neutrino oscillations and non-zero neutrino masses are profound. The observed patterns of fermion masses have no explanation within the Standard Model, in particular why the neutrino masses should be so much less than that of the charged fermions, as shown in

Table 1.1. The origin of neutrino mass is a subject of considerable theoretical effort, with various proposals having been put forward. One such theory proposes the use of the Higgs mechanism to generate mass, with neutrinos being Dirac particles and the as-yet unobserved right-handed neutrinos, ν_R (introduced in Section 1.1.2), existing but having extremely weak interactions. An alternative view with greater popular support is that neutrinos are Majorana particles, with ν_R now pictured as being extremely massive and short lived. Normal left-handed neutrinos, ν_L, are then thought to have a mass similar to the other standard model particles, but due to brief transformations into the large ν_R through interactions with the Higgs field, their mass is effectively shrunk by the a factor of 1 over the ν_R mass, the so-called "see-saw" mechanism. Such theories also offer a potential route to explaining the matter - anti-matter asymmetry of the Universe, by postulating that the decay of the massive Majorana neutrinos in the early Universe violates CP symmetry with a bias towards matter, so called "leptogenesis" [8, 18, 19].

The question of whether a neutrino is a Dirac particle, that is, possesses a distinct antiparticle, or a Majorana particle, that is, the neutrino being its own antiparticle, is also in and of itself of fundamental interest. Should the neutrino be Majorana in nature it will be the first case of a fermion being so, permissible as the neutrino unlike the other fermions possess no electric charge. This has led to a number of experiments searching for a process known as neutrino-less double β-decay, $0\nu\beta\beta$, which can only take place if the neutrino is Majorana (see Section 1.4.1).

In cosmology too, neutrino mass has a large role to play. It is thought that relic neutrinos from the Big Bang form a background similar to the cosmic microwave background of photons. These relic neutrinos outnumber relic baryons by 10^9:1 and are thought to act as relativistic or "hot" dark

matter, with cosmological models of structure formation depending heavily on the relative amounts of hot and cold dark matter present in the early Universe [8, 18].

1.4 Experimental Programme

The central efforts in neutrino physics at present consist of the measurement of the parameters of the PMNS mixing matrix given in Eqn. 1.20, measuring the absolute neutrino mass scale and determining whether neutrinos are Dirac or Majorana in nature. In the case of the oscillation parameters the principal goals are the precise determination of the three mixing angles, with an emphasis on θ_{13} which remains mostly unknown; the accurate measurement of the mass squared differences Δm^2_{31} and Δm^2_{21}; the determination of the sign of Δm^2_{31} (and so the form of the mass hierarchy); and the measurement of the CP violating phase δ. An accurate knowledge of these parameters is vital if the opportunities for new physics outlined in the previous section are to be realised.

In Section 1.4.1 a short account is presented of the current efforts towards finding the absolute neutrino mass scale and the determination as to whether the neutrino is Dirac or Majorana, the two being linked by the search for neutrino-less double β decay. This is followed by an overview of the current limits placed upon the oscillation parameters in Section 1.4.2, while in Section 1.5 future neutrino oscillation experiments are discussed.

1.4.1 Absolute Mass Scale and Neutrino-less Double Beta Decay Searches

Measurements of the absolute mass scale of the neutrino are possible by looking at the highest energy region of the spectrum of electrons emitted in

β-decay, typically using tritium, or by searching for neutrino-less double beta decay. The former is being pursued by the KATRIN collaboration[18]. In the case of a finite neutrino mass, instead of the β electron energy spectrum ending in a smooth tail it will curve downwards and end sharply, as no electrons can receive the full energy of the decay as some of this energy must go into making up the neutrino mass. Further, the spectrum end point region will also feature "kinks", resulting from the observed spectrum being a superposition of different spectra with different end points corresponding to the different neutrino mass eigenstates. By observing the spectral end point region with extremely high precision it should thus be possible to perform an absolute measurement of the neutrino mass.

The other class of experiment which is sensitive to the absolute scale of the neutrino mass are searches for neutrino-less double β-decay ($0\nu\beta\beta$). Standard double β-decay involves a nucleus emitting two electrons and two anti-neutrinos, with a corresponding conversion of two neutrons to protons, at the same time. This process, while very rare, has been experimentally observed. In $0\nu\beta\beta$ decay the same process occurs but with no anti-neutrino emission, as the outgoing anti-neutrino from one interaction is re-absorbed as an incoming neutrino at the second interaction. For this to occur it requires that neutrinos be Majorana rather than Dirac particles (as described previously in Section 1.3). If this is so, then the rate at which these interactions take place will be sensitively dependent on the absolute neutrino mass scale. Various experiments, such as NEMO3[20], have searched for $0\nu\beta\beta$ decay, with no signal confirmed to date, though the area remains highly active with a variety of experiments planned with increased sensitivity[19].

1.4.2 Current Limits on the Oscillation Parameters

The account here mainly follows that given in [21]. A calculation of the current limits on the neutrino mass and mixing parameters from global data fits are shown in Table 1.2. Oscillation experiments are not sensitive to the absolute neutrino mass scale, only to the mass squared differences, $\Delta m^2_{21} = m^2_2 - m^2_1$, $\Delta m^2_{31} = m^2_3 - m^2_1$ and $\Delta m^2_{32} = m^2_3 - m^2_2$. Empirically it is known that:

$$|\Delta m^2_{31}| \approx |\Delta m^2_{32}| \gg \Delta m^2_{21} \tag{1.21}$$

Δm^2_{21} is constrained principally by KamLAND and Δm^2_{32} like θ_{23} is constrained by atmospheric data from Super Kamiokande. The bounds on $|\Delta m^2_{31}|$ come from MINOS[22]; note the sign ambiguity, which leads to two possibilities for the neutrino mass hierarchy, known as normal ordering $m_1 < m_2 < m_3$ and inverted ordering $m_3 < m_1 < m_2$.

Turning to the mixing angles, the bounds on θ_{12} are provided by solar experiments, in particular at the SNO, and also from KamLAND. θ_{23} is close to being maximal and the bounds on its value arise from atmospheric neutrino data from Super Kamiokande, with additional support from the long baseline accelerator experiments, K2K[23] (which used Super Kamiokande as a far detector) and MINOS. θ_{13} is small, but is expected to be non-zero, the current limits being provided from a number of data sets, in particular data from the CHOOZ reactor experiment[24], combined with the present determination of $|\Delta m^2_{31}|$.

Parameter	**Best Fit** $\pm 1\sigma$	3σ **Interval**
$\Delta m^2_{21}[10^{-5}eV^2]$	$7.59^{+0.23}_{-0.28}$	7.03 - 8.27
$\Delta m^2_{31}[10^{-3}eV^2]$	$\pm 2.40^{+0.12}_{-0.11}$	$\pm$(2.07 - 2.75)
$\sin^2\theta_{12}$	$0.318^{+0.019}_{-0.016}$	0.27 - 0.38
$\sin^2\theta_{23}$	$0.50^{+0.07}_{-0.06}$	0.36 - 0.67
$\sin^2\theta_{13}$	$0.013^{+0.013}_{-0.09}$	≤ 0.053

Table 1.2: Limits on neutrino mass and oscillation parameters from global data fits. Taken from [21, 25].

1.5 Near Future and Next Generation Oscillation Experiments

To better measure the oscillation parameters described in Section 1.4.2 a variety of new experiments are in production or have recently begun. The sensitivities of these current and near future neutrino oscillation experiments to $\sin^2 2\theta_{13}$ are shown in Fig. 1.1. The experiments listed are the Super Beam project T2K[26], which began taking data in 2010; the NOνA experiment[27], another proposed Super Beam based at Fermilab; and the Double CHOOZ[28], Daya Bay[29] and RENO[30] reactor experiments. If θ_{13} is large, then its value should be within the reach of these experiments. If however θ_{13} is small it will be necessary to look still further ahead, beyond the current and near future generation of experiments to more advanced "Next Generation" facilities in order to measure θ_{13}. Such a facility will also be neccessary, even in the case of a large θ_{13}, in order to measure the sign of Δm^2_{31} (though for large θ_{13} NOνA may give a hint) and in order measure the CP violating phase δ[31].

Three types of facility have been proposed for such next generation experiments: advanced Super Beams, β Beams, and the Neutrino Factory. A

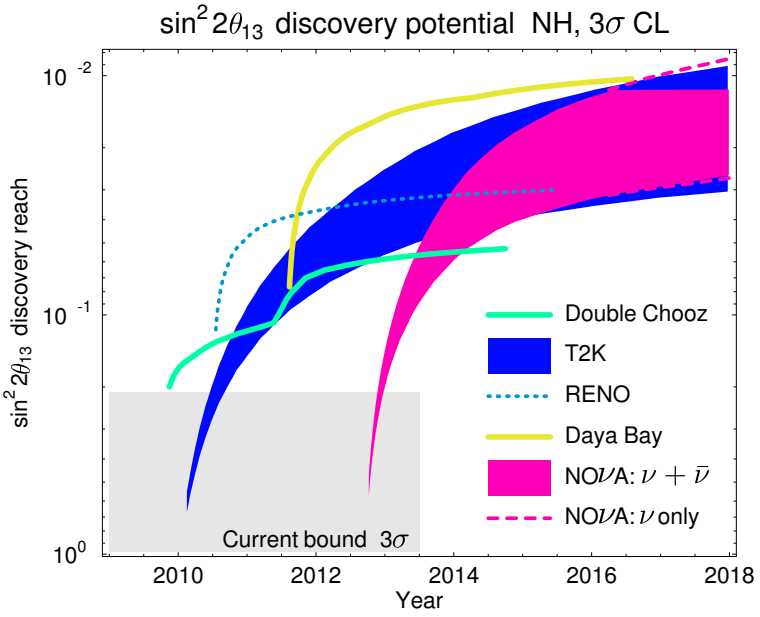

Figure 1.1: Discovery potential for $\sin^2 2\theta_{13}$, at a 3σ confidence level (CL), of recent and near future experiments, for the case of a normal neutrino mass heirarchy (NH). Taken from[31].

description and detailed comparison of the relative merits of each facility have been performed in [19] and [31], which shall be briefly summarised here.

The first of the proposed new facilities, the Super Beam[31], represents an evolution of the present generation of accelerator-driven neutrino experiments, such as the T2K project. A high power proton beam is collided with a target in order to generate pions which are then focused by a magnetic horn and decay generating the neutrino beam. Two main proposals for such a facility are under way, one based on the CERN Super Proton Linac (SPL)[32], the other in the context of the US based Long Baseline Neutrino Experiment (LBNE)[33]. Both projects envisage a MW range proton beam, of 4 GeV energy in the case of SPL and 120 GeV for LBNE[31].

The second type of new facility, the β Beam[31], is based on the decay

of accelerated radioactive isotopes which undergo β-decay to generate the neutrino beam. Isotopes with half lives in the region of 1 second are favoured to allow sufficient time for acceleration, but also to produce a reasonable neutrino flux. Relativistic γ factors, achieved through the acceleration, are planned to be in the region of $\gamma \simeq 50 - 500$. The maximum energy that can be given to the decay neutrinos is given by:

$$E_{Max} = \gamma Q_{\beta} \tag{1.22}$$

where Q_{β} is the maximum energy available to the neutrino from the β-decay of the parent nucleus. Q_{β} has a value in the region of 1 - 10 MeV, implying that β Beams are capable of producing maximum neutrino energies of 0.1 - 5 GeV. A benefit of β Beams over Super Beams is that they produce a pure ν_e beam, while Super Beams produce a ν_μ beam with a small ν_e component, leading to increased backgrounds due to uncertainties in the beam content.

The last next generation facility is the Neutrino Factory, which is based on the principal of creating a neutrino beam using the decay of muons in a storage ring. The muons themselves are generated from pion decay, with pions being produced from a high power proton beam intersecting a target, as in the case of the Super Beam. Like the β Beam, the beam content is well understood; for a μ^- stored beam the resultant neutrino beam is a 50% mix of $\bar{\nu}_e$, ν_μ, or for the μ^+ case a beam with a 50% mix of ν_e, $\bar{\nu}_\mu$ is produced. Whilst being challenging technically, the Neutrino Factory offers the best coverage of most of the oscillation parameter space, as illustrated in Fig. 1.2. In particular, the Neutrino Factory has the greatest sensitivity to a small θ_{13} and to δ. The design of a future Neutrino Factory, including the necessity of ionisation cooling, is discussed in Section 1.6 below.

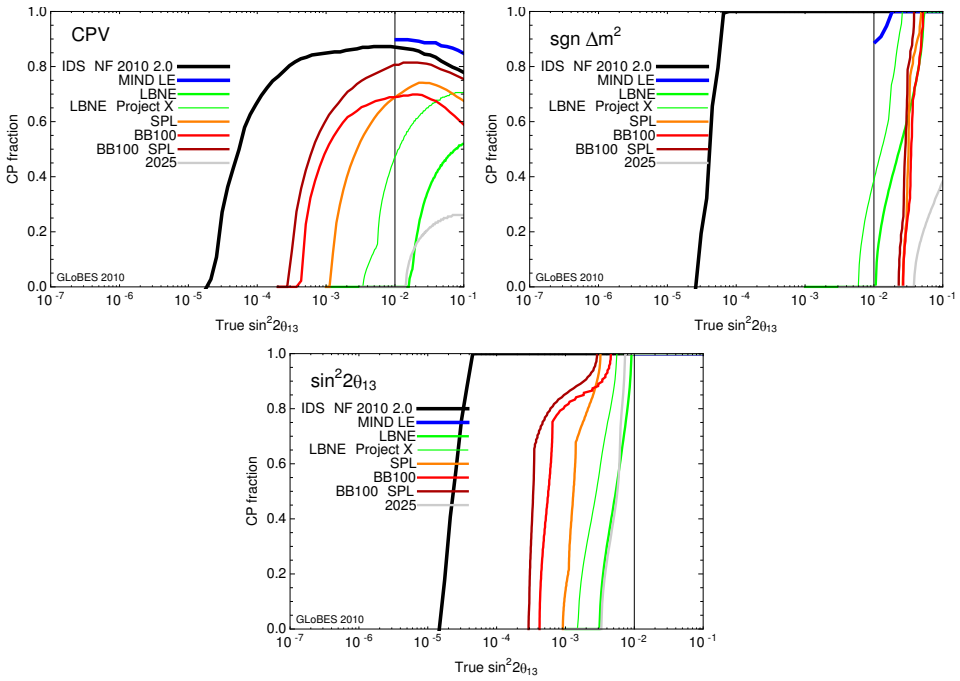

Figure 1.2: The discovery reach for the CP violating phase δ (top left), the mass hierarchy (top right) and $\sin^2 2\theta_{13}$ (bottom middle) of the proposed next generation neutrino experiments. The area to the right of each band shows the discovery potential at 3σ for each facility. SPL is a CERN-based Super Beam proposal while LBNE is the Long Baseline Neutrino Experiment, another proposed Super Beam experiment based in the US. BB refers to a β Beam. IDS NF refers to the International Design Study for the Neutrino Factory (see Section 1.6), while MIND LE refers to a single baseline low energy Neutrino Factory optimised for large $\sin^2 2\theta_{13}$. Project X is a proposed high intensity proton accelerator upgrade for Fermilab in the US (taken from[31]). Further comparisons are available in [19].

1.6 The Neutrino Factory

1.6.1 Introduction

The Neutrino Factory is a proposed next generation neutrino source, producing a high energy, high intensity, well understood beam from the decay of muons in a storage ring. As described in the previous section the Neutrino Factory offers the best coverage of the neutrino oscillation parameter space when compared with second generation Super Beam and β Beam facilities. It consists of a high power proton driver, pion production target, pion decay channel, muon cooling system, muon acceleration system and storage ring. The Neutrino Factory baseline, as defined by the International Design Study for the Neutrino Factory (IDS-NF), is described in this section, following the account given by the IDS-NF itself in [31].

A Neutrino Factory holds many components in common with another proposed facility known as the Muon Collider, which offers a potential route to multi TeV lepton - anti-lepton collisions. In particular both facilities require a muon cooling system (see Section 1.6.5). The Muon Collider will not be discussed further here, but for an overview see for example [34].

A schematic of the current Neutrino Factory baseline, as defined by the IDS-NF, is shown in Fig. 1.3[35]. A proton beam of high power is collided with a fixed target leading to hadronic interactions which produce a pion shower. These pions are then captured and allowed to decay to muons, which are then bunched and rotated in longitudinal phase space such that the large initial energy spread from the pion decay is reduced, leading to a corresponding increase in the time spread. The beam is then cooled to fit within the acceptance of the acceleration system. Once accelerated to the desired energy, currently specified to be 25 GeV, the muons are transferred

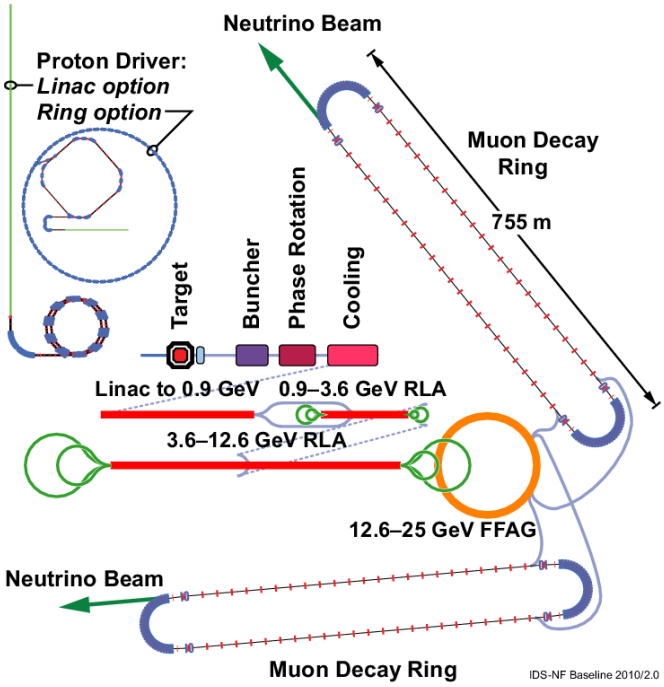

Figure 1.3: The IDS-NF baseline design for a Neutrino Factory, see[36].

to a storage ring where they decay to neutrinos. The storage rings are built with long straight sections, creating neutrino beams in the directions of the far detectors. Two source - detector baselines are defined, one of length 3000 - 5000 km, the other 7000 - 8000 km, each with a magnetised iron detector.

1.6.2 Proton Driver

In order to generate the required muon rate and energy, the proton driver for the Neutrino Factory is required to deliver a 4 MW power beam, at 5 - 15 GeV in energy, to the target. The proton beam must be bunched, with each bunch being 1 - 3 ns long, and with a 120 μs separation between each bunch. Three bunches then form a bunch train, of length 240 μs, the trains themselves being spaced every 20 ms (a 50 Hz repetition rate). Each individual proton bunch will go on to form a pion bunch downstream[31, 37].

Both linacs and circular machines are being considered to serve as a proton driver. In the case of a linac two possible implementations are being considered, either the SPL at CERN[38] or the Project-X linac at Fermilab[39]. In the case of a circular machine this could either be a Rapid Cycling Synchrotron (RCS) or a non-scaling Fixed Field Alternating Gradient (FFAG) system, potential sites being the new J-PARC complex in Japan or an upgraded version of the ISIS accelerator at RAL in the UK[40].

1.6.3 Target and Pion Capture

Due to the high power of the proton driver the Neutrino Factory will require a particularly robust target design. The current target baseline is a free liquid mercury jet, which presents a high atomic number material to the beam leading to the generation of a large pion flux. The baseline prescribes a jet 8 mm in diameter with a flow rate of 20 ms^{-1} in order to overcome the downward curvature caused by gravity. The jet must also be free, that is not contained within a pipe, as any such pipe would be damaged by the intense pressure waves induced in the mercury by the proton beam. Once the mercury has passed through the interaction region it enters a pool where it serves as the beam dump.

Both the target and the beam dump are contained within a superconducting solenoid, tapered from 20 T to 1.5 T, used to capture the produced pions and transport them to the muon front end (see Section 1.6.4). Placing the target in the solenoidal channel also has the benefit of greatly improving the stability of the mercury jet, as demonstrated by the MERIT experiment[31, 41]. This allows the jet to regenerate between the arrival of proton bunch trains up to the desired repetition rate of 50 Hz.

In addition to the baseline mercury jet, alternative designs are also being

considered. These include a solid target composed of tungsten rods which are interchanged between beam pulses[42], and a tungsten powder jet propelled by gas[43].

1.6.4 Muon Bunching and Phase Rotation

The region between the target and muon acceleration system in the Neutrino Factory is known as the *front end*, where pion capture, pion decay, bunching, phase rotation and beam cooling take place. The purpose of these last three is to manipulate the beam phase space such that it fits within the acceptance of the downstream acceleration system as closely as possible.

The pions are captured in the region of the target by the 20 T solenoidal field as mentioned in the previous section. This field then tapers down over a distance of 15 m until it reaches a value of 1.5 T. Over the same over length there is an increase in the beam pipe radius from 0.075 m to 0.3 m, which together with the tapering magnetic field forms the magnetic focusing horn.

Following the magnetic horn there is an empty magnetic lattice for longitudinal drift. The allows a correlation to develop between particle energy and longitudinal position. Here also a large fraction of the pions undergo decay to the desired muons, according to:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{1.23}$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \tag{1.24}$$

As the pion decay is a three body decay the muons are created with a large energy spread.

Following on from the drift length there comes the buncher, where Radio

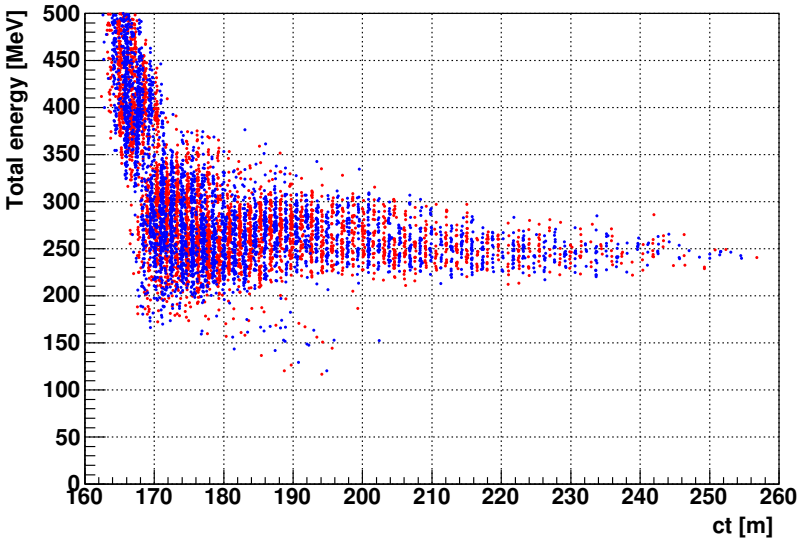

Figure 1.4: An example plot showing the longitudinal phase space of the beam at the end of the rotator. Red indicates μ^+, blue μ^-. The vertical bands represent the bunch structure. Taken from [31].

Frequency (RF) electric fields are used to impose a bunch structure on the muon beam, on top of the existing bunch train structure of 240 μs produced by the proton driver. The bunched beam then proceeds to the rotator, which exploits the position - energy correlation built up in the drift length in order to shift the large energy spread of the beam into a spread in longitudinal position. Using RF fields, those bunches which have an energy greater than a central reference bunch are decelerated, while those with less energy are accelerated, so that upon leaving the rotator all bunches have approximately the same energy. An example plot showing the longitudinal phase space of a bunched, rotated beam is given in Fig. 1.4[31].

1.6.5 Cooling

The final section of the muon front end is the cooling channel. The term cooling refers to reducing the phase space volume (or *emittance*) of the beam, either just the transverse (4D cooling) or the longitudinal as well (6D cooling). A modest degree of beam cooling is necessary in a Neutrino Factory in order for the muon beam to fit well into the acceptance of the downstream accelerating components (in comparison with a Muon Collider which requires more stringent beam cooling, see [44]).

Due to the short muon lifetime of 2.2 μs standard beam cooling techniques, such as stochastic cooling, cannot be used as the beam decays too quickly. This led to the idea of *ionisation cooling*. In ionisation cooling a beam of particles is passed through some absorber medium leading to an energy / momentum reduction in *all* directions. The beam is then subsequently re-accelerated using RF cavities in the longitudinal direction only, leading to a reduction in beam emittance. This method of cooling is fast enough to be applied to muons but, although proposed many years ago (see [45]), has yet to be demonstrated. The MICE experiment is designed to demonstrate this technique for the first time, being a prototype of a single cell of a full Neutrino Factory cooling channel, based on the Neutrino Factory Feasibility Study IIA[46] (the precursor to the current baseline). The MICE experiment is described in Chapter 2, where in Section 2.2 beam emittance is introduced and given a mathematical basis, followed by a more detailed treatment of ionisation cooling in Section 2.3.

The baseline Neutrino Factory cooling channel consists of 50 cooling cells, each 1.5 m in length, consisting of two RF cavities with a lithium hydride absorber disc placed before and after each cavity (4 per cell). Each cell also possesses two superconducting solenoidal coils of opposite polarity,

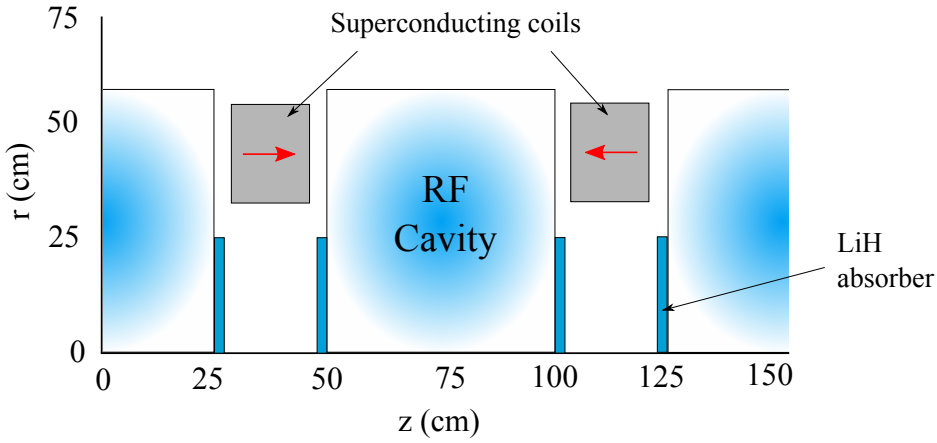

Figure 1.5: Schematic of a single Neutrino Factory cooling cell. LiH absorber discs (four per cell) remove particle momentum in all directions, while longitudinal momentum only is replaced via the RF cavities (two per cell), leading to reduction in beam emittance. Two superconducting coils of opposite polarity are used to provide transverse focusing. Taken from [31].

producing a field which varies approximately sinusoidally from +2.8 T to −2.8 T over the length of a cell, and is used to provide transverse focusing. A schematic of a single cell is shown in Fig. 1.5.

Simulations reported in the IDS-NF Interim Design Report (IDR, [31]) estimates that the current baseline cooling channel reduces an input normalised transverse rms beam emittance of $\epsilon_N = 0.018$ m to $\epsilon_N = 0.0075$ m after the cooling. Taking this together with the effect of the rest of the muon front end, the total muon rate that falls within the acceptance of the acceleration components is estimated by the IDR to rise by a factor of 10.

1.6.6 Acceleration and Storage

The short muon lifetime and large emittance, even post-cooling, set constraints on the Neutrino Factory acceleration system. The short lifetime leads to the necessity of high gradient accelerating fields, while the large

emittance requires a high longitudinal and transverse acceptance. The Neutrino Factory baseline envisages a complex acceleration system based on a linac, two "dog-bone" recirculating linacs (RLAs) linked via chicanes, and finally a non-scaling Fixed Field Alternating Gradient (FFAG) ring.

The layout of the acceleration system appears in Fig. 1.3. The beam arrives from the front end with a momentum of 230 MeV/c and is accelerated by the linac system to an energy of 12.6 GeV. After the linacs the muons are further accelerated by the FFAG ring, taking the beam to its final energy of 25 GeV. As the name implies FFAGs do not require magnets to be ramped, in contrast to the magnet system of a synchrotron, so they are able to perform acceleration very quickly, and also naturally have large acceptances, making them well suited to the task of muon acceleration. For details see [21, 31].

Once accelerated to the desired energy the muons must be held in a storage ring while they decay. These rings are designed with long straight sections, of greater length than the corner arcs, so that the neutrinos are produced preferentially in the directions of these straight sections to create a beam. The baseline design calls for a "race track" topology, consisting of two long straight sections and two arcs, as was illustrated in Fig. 1.3. The total length of each ring is ∼1600 m, each straight section being ∼600 m, with a central momentum of 25 GeV/c.

Of the two straights only one is used for neutrino production; the other, known as the return straight, is used for beam collimation, RF cavities and controlling the tune (that is, the number of betatron oscillations made by the beam over a full circuit). Two such rings are planned, one pointing towards the intermediate baseline detector at 3000 - 5000 km, the other to the far detector at 7000 - 8000 km (see Section 1.6.7 below). Due to the curvature

of the Earth the production straights must also be tilted downwards, by ∼18° for the intermediate detector and ∼36° for the far detector[21, 31].

1.6.7 Detectors

Beams and Signals

Muons and anti-muons in the Neutrino Factory storage ring decay according to Eqn. 1.25 and 1.26, with the corresponding Feynman diagrams shown in Fig. 1.6:

$$\mu^- \rightarrow \ e^- + \bar{\nu}_e + \nu_\mu \tag{1.25}$$

$$\mu^+ \rightarrow \ e^+ + \nu_e + \bar{\nu}_\mu \tag{1.26}$$

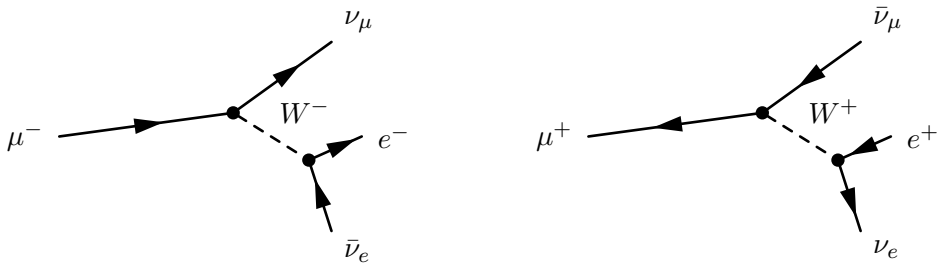

Figure 1.6: Feynman diagrams of muon and anti-muon decay.

Hence μ^- produce a beam composed of 50% $\bar{\nu}_e$ and 50% ν_μ, while μ^+ produce a beam composed of 50% ν_e and 50% $\bar{\nu}_\mu$, with which to study oscillations.

The principal signal to be searched for in the Neutrino Factory is known as the "Golden channel", based on the detection of $\nu_e \rightarrow \nu_\mu$ oscillations. This channel is favoured as it is sensitive to $\sin^2 2\theta_{13}$ and to δ, even for small values of $\sin^2 2\theta_{13}$[19]. Further, it requires looking for "wrong-sign" muons in the detector, that is, muons which have an opposite charge to that of

those held in the storage ring, allowing the signal to be better observed above the dominant background from "right-sign" muons. Taking the case of anti-muons in the storage ring, a neutrino beam is produced according to Eqn. 1.26. The ν_e in the beam then oscillate over the length of the baseline and may interact in the intermediate and far detectors as ν_μ, as shown in Eqn. 1.27. Similarily the interaction for the case of a muon beam in the storage ring is given by Eqn. 1.28, with the corresponding Feynman diagrams shown in Fig. 1.7:

$$\nu_\mu + n \rightarrow \mu^- + p \tag{1.27}$$

$$\bar{\nu}_\mu + p \rightarrow \mu^+ + n \tag{1.28}$$

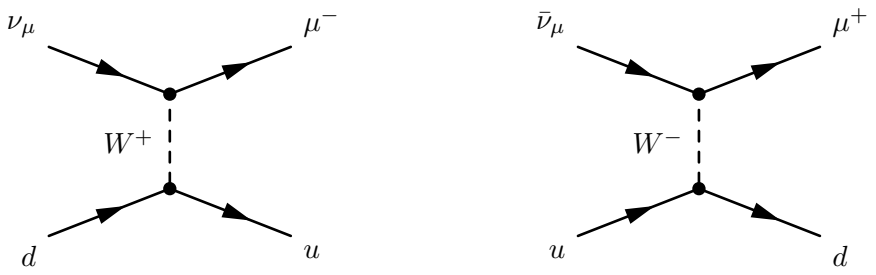

Figure 1.7: Feynman diagrams of the Golden channel.

By measuring the flux and content of the neutrino beam first at the near detector and then at the intermediate or far detector and observing the change, oscillation probabilities may be deduced. These may then be used to infer the oscillation parameters (see Section 1.2).

Besides the Golden channel, numerous other oscillation channels are available to the Neutrino Factory (12 in total). Of particular importance is the "Silver channel", another appearance mode, this time looking for $\nu_e \rightarrow \nu_\tau$ oscillations. This again produces wrong-sign muons in the intermediate and far detectors, this time resulting from tau decay. The chain

proceeds as in Eqn. 1.26, with the resultant ν_e oscillating to ν_τ which then interacts in the intermediate or far detector producing a tau according to Eqn. 1.29. The resultant tau then decays according to Eqn. 1.30 producing the wrong-sign muon (the corresponding Feynamn diagrams are shown in Fig. 1.8):

$$\nu_\tau + n \rightarrow \tau^- + p \tag{1.29}$$

$$\tau^- \rightarrow \nu_\tau + \mu^- + \bar{\nu}_\mu \tag{1.30}$$

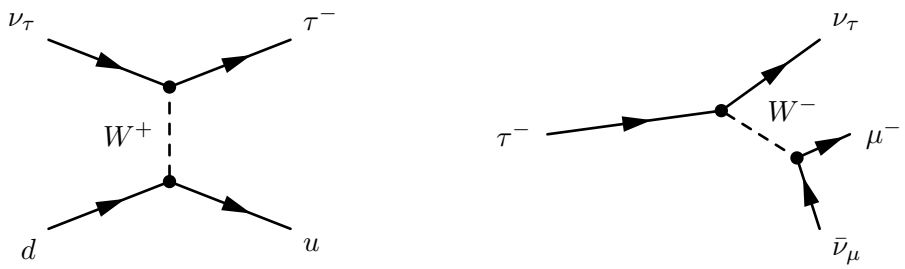

Figure 1.8: Feynman diagrams for the Silver channel.

Baselines

The favoured Neutrino Factory design involves two baselines, one to an intermediate detector at 3000 - 5000 km, and the other to a far detector at 7000 - 8000 km. The latter is known as the "magic" baseline, as matter effects cancel most of the effect of the CP violating phase δ, allowing a detector at this baseline to make a clean measurement of $\sin^2 2\theta_{13}$ and the sign of Δm^2_{13}. It does however have the obvious disadvantage of then not being able to measure δ, as well as having relatively low statistics due to the long baseline. Hence a second baseline is proposed at 3000 - 5000 km, optimised for measuring δ, and in order to generate higher statistics than

are possible with the magic baseline. For details see[47].

Detector Designs

The far detector for a high energy Neutrino Factory (25 GeV muons) as defined by the Neutrino Factory baseline is known as the Magnetised Iron Neutrino Detector (MIND). It consists of layers of iron and plastic scintillator, following the same basic design as the MINOS detector. The baseline envisages two such detectors, a 100 kton version at the intermediate baseline, and a 50 kton version at the magic baseline, primarily looking for the Golden channel signature. The magnetisation is necessary as both muons and anti-muons are produced by the neutrino beam, making charge separation a requirement. The MIND detectors may be supplemented by a further Magnetised Emulsion Cloud Chamber detector (MECC) looking for the Silver channel. MECC is of particular importance in looking for physics beyond the Standard Neutrino Model[35].

Alternatives to MIND being consider by the Neutrino Factory community include a Total Active Scintillating Detector (TASD), which would be suitable for a low energy Neutrino Factory (5 GeV muons), or a liquid argon based detector. See[48] for details.

Chapter 2

MICE

...to boldly go where no man has gone before.

James T. Kirk

2.1 Introduction

The international Muon Ionisation Cooling Experiment (MICE) is designed to investigate systematically the feasibility and performance of ionisation cooling, for application to a future Neutrino Factory or Muon Collider. In particular MICE, when complete, is to represent a prototype of a single lattice cell of the Neutrino Factory cooling channel (based on the Neutrino Factory Feasibility Study IIA[46]), with an input muon beam of variable momentum between 140 MeV/c and 240 MeV/c. The cooling channel is designed to cause a decrease in beam emittance of $\sim$10% and possess the ability to measure emittance to one part in 10^3 before and after the cooling section by means of high precision scintillating fibre trackers[49].

MICE is hosted at the Rutherford Appleton Laboratory (RAL), UK and is a *staged* experiment, that is, it is built and run in discrete periods. The MICE staging plan is shown Fig. 2.1 (currently under review), together with

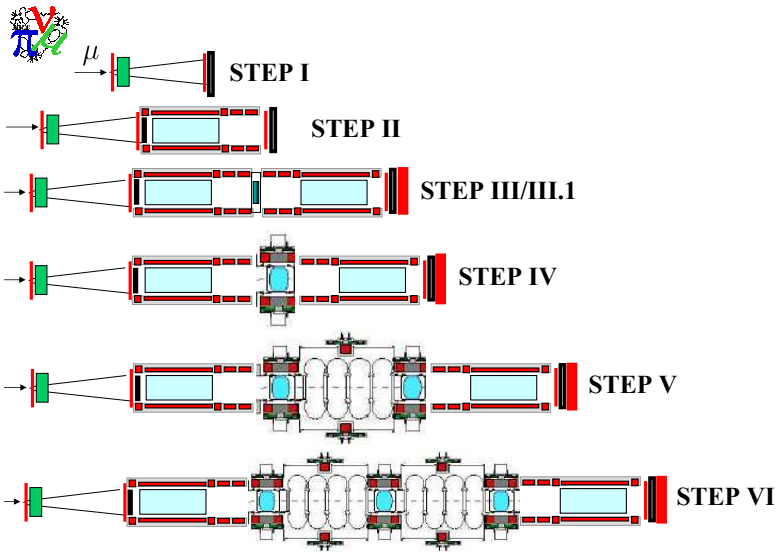

Figure 2.1: The MICE stages as of March 2010 (A. Blondel, University of Geneva). Step I represents a functional muon beamline with some particle identification. Step II introduces the first tracker, followed by the second tracker in Step III. In Step IV the first absorber module arrives, followed by the second absorber and first RF cavity in Step V. Finally in Step VI the third absorber and second RF cavity are added.

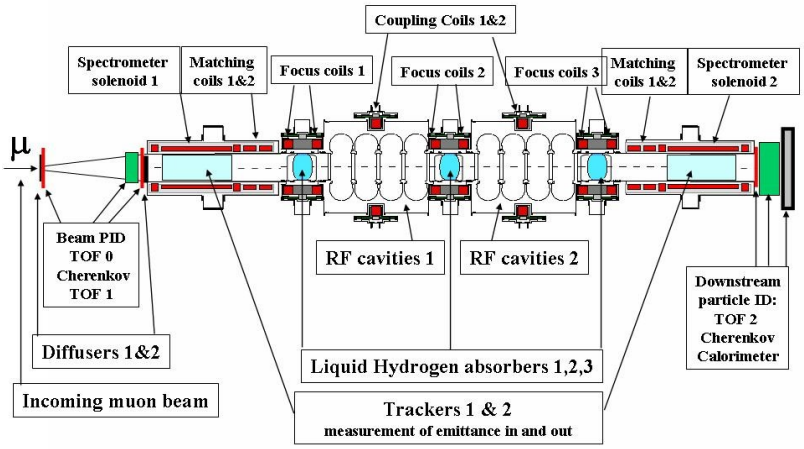

Figure 2.2: The key to the schedule representing MICE Stage VI (A. Blondel, University of Geneva).

a key showing the final, completed MICE in Fig. 2.2. MICE is currently at Step I, with the absorber modules, RF cavities and the solenoids for the trackers currently still in production.

In the following sections, the concept of emittance is described (section 2.2), how in the case of muons ionisation cooling may be used to reduce it (section 2.3), and how this is to be implemented in the MICE cooling channel (section 2.4). The beamline used to supply muons to the cooling channel is then described (section 2.5), followed by the proton source (section 2.6), pion production target (section 2.7) and the detectors and data acquisition systems (section 2.8). Lastly the issue of ISIS beam loss and particle rate in the MICE beamline, representing the central question addressed by this thesis, is introduced (section 2.9).

2.2 Emittance

2.2.1 Definition

As discussed in Section 1.6.5, beam cooling is required by a Neutrino Factory in order for the muon beam produced at the front end to be sufficiently "small" to fall within the acceptance of the downstream acceleration system. In accelerator physics, the idea of beam size and divergence is encapsulated by the concept of *emittance*, and it is this quantity which must be minimised in order to fit the beam efficiently into the acceptance of the downstream acceleration components. The following provides an introductory description of emittance and its conservation, roughly following the accounts given in [50] and [51] (the latter provides a more detailed treatment of emittance for the interested reader).

Each particle within a beam may be described by a set of six coordinates

known as the *phase space*. This space consists of the two transverse spatial coordinates (x, y), the corresponding two transverse momenta (p_x, p_y), and the longitudinal coordinates (ϕ, E) corresponding to the phase with respect to the synchronous particle and the energy respectively (ϕ may be replaced with the longitudinal distance coordinate z, or by the time, t). It is useful also to define a slightly different space known as the *trace space* which may be obtained from the phase space by transforming the momenta according to:

$$x' = \frac{p_x}{p_z} = \frac{dx}{dz} \tag{2.1}$$

$$y' = \frac{p_y}{p_z} = \frac{dy}{dz} \tag{2.2}$$

The trace space is then given by (x, x', y, y', ϕ, E) where x' and y' represent the divergence of the particle trajectory away from the longitudinal axis, z. The emittance, ϵ, of a beam in a given direction may then be defined as the volume within the trace space occupied by the beam, divided by a factor π to give units of mm.rad:

$$\epsilon = \frac{1}{\pi} \int x' \, \mathrm{d}x. = \frac{A}{\pi} \tag{2.3}$$

The beam distribution within trace space usually takes the form of an ellipse, as illustrated in Fig. 2.3. The equation describing this ellipse, sometimes called the Courant and Snyder invariant, is then given by:

$$\gamma(z)x^2 + 2\alpha(z)xx' + \beta(z)x'^2 = \epsilon \tag{2.4}$$

where α, β and γ are the parameters of the ellipse, known as the Twiss parameters. These three are not independent but, using the geometry of ellipses, are related by:

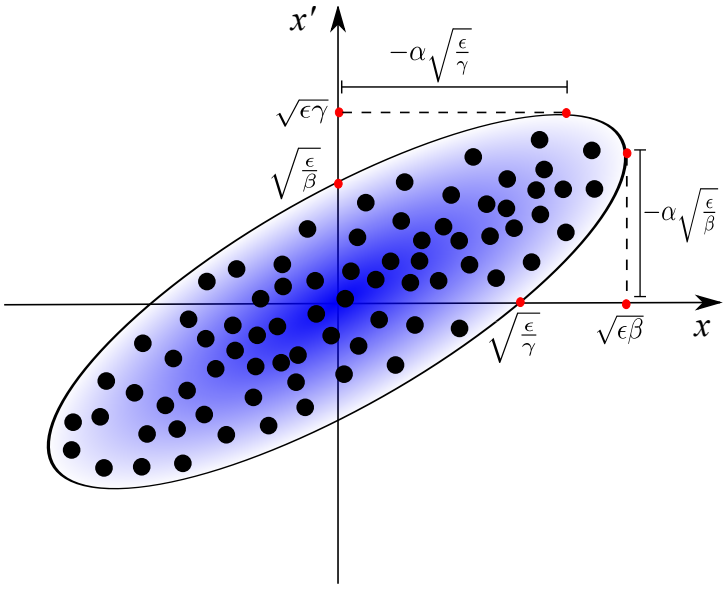

Figure 2.3: A trace space plot illustrating the concept of emittance. ϵ is the emittance, defined as the area of the ellipse, while α, β and γ are the Twiss parameters.

$$\beta\gamma - \alpha^2 = 1 \tag{2.5}$$

More generally, the equation describing an n-dimensional ellipse can written as:

$$u^T V^{-1} u = 1 \tag{2.6}$$

u is a vector of the trace space coordinates and V is known as the *covariance matrix*. The volume of the hyper-ellipsoid representing the beam in this space is then given by:

$$\mathrm{Vol}_n = \frac{\pi^{\frac{n}{2}}}{\Gamma(1+\frac{n}{2})}\sqrt{\det V} \tag{2.7}$$

where Γ is the gamma function. In the 2D case we may compare Eqn. 2.6 with Eqn. 2.4 to get:

$$\begin{pmatrix} x & x' \end{pmatrix} \frac{1}{\epsilon} \begin{pmatrix} \gamma & \alpha \\ \alpha & \beta \end{pmatrix} \begin{pmatrix} x \\ x' \end{pmatrix} = 1 \tag{2.8}$$

$$\implies V^{-1} = \frac{1}{\epsilon} \begin{pmatrix} \gamma & \alpha \\ \alpha & \beta \end{pmatrix} \tag{2.9}$$

$$\implies V = \epsilon \begin{pmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{pmatrix} \tag{2.10}$$

Substituting Eqn. 2.10 into Eqn. 2.7 and using the Twiss parameter relation in Eqn. 2.5 we recover the earlier definition of emittance as the phase space volume of the beam ellipse divided by π:

$$\text{Vol}_2 = \frac{\pi}{\Gamma(2)} \sqrt{\det V} = \pi \sqrt{(\epsilon\beta)(\epsilon\gamma) - (\epsilon\alpha)^2} = \pi\epsilon \tag{2.11}$$

For the 2D case to be valid the two transverse directions must not be coupled to each other, or to the longitudinal direction. If there is a correlation between the x and y components then the larger 4D space must considered, with the emittance then representing the volume occupied by the beam in this hyperspace. If the longitudinal phase space is also coupled to the transverse then the full 6D space and emittance must be considered. In the case of the MICE solenoidal beamline it is sufficient to first order to consider the 4D transverse emittance with the 2D longitudinal emittance decoupled (see Section 2.2.5 below).

2.2.2 Conservation: Liouville's Theorem

A conservation law applies to Eqn. 2.4 known as Liouville's theorem, which states that ϵ is conserved if the beam is subject to only conservative forces,

such as magnetic focusing and bending forces. Examples of non-conservative forces include synchrotron radiation and space charge. When Liouville's theorem holds, the ellipse shown in Fig. 2.3 may take different forms as it moves along the beamline (the Twiss parameters are functions of z) but the area must remain constant. As the maximum excursion of the beam in real space is given by $\sqrt{\epsilon_{x,y}\beta_{x,y}}$ then provided ϵ and $\beta(z)$ are known, the beam envelope may be calculated for the whole beamline.

2.2.3 Normalised emittance

Using the definition of emittance in Eqn. 2.3 implies that Liouville's theorem will not hold for an accelerated beam, limiting its usefulness. If however the concept of emittance is refined to use a relativistic definition of momentum, after the manner of a conjugate momentum in Hamiltonian mechanics, the invariance of the emittance may be asserted even in the presence of acceleration. The appropriate conjugate momentum is given by multiplying the classical expression for momentum by the relativistic γ_r factor:

$$p = \gamma_r m_0 v = m_0 c \gamma_r \beta_r \tag{2.12}$$

where v is the particle velocity and β_r and γ_r have their usual relativistic meaning (differentiated from the twiss parameters by use of the subscript r):

$$\beta_r = \frac{v}{c} \tag{2.13}$$

$$\gamma_r = \frac{1}{\sqrt{1-\beta_r^2}} \tag{2.14}$$

This leads to a definition of the *normalised emittance* which remains invariant even under acceleration:

$$\epsilon_n = \gamma_r \beta_r \epsilon \tag{2.15}$$

2.2.4 RMS emittance

In practice beams do not possess sharply defined edges, but may often be approximated as being a Gaussian distribution in both x and y. The ellipse in Fig. 2.3, and so the emittance, may then be defined statistically. In the case of electrons a quantity known as the RMS emittance, ϵ_{rms}, is defined as the volume enclosed by those particles 1σ distant from the mean of the distribution in both x and x' (and similarly for y), σ being the standard deviation of the Gaussian distribution. In the case of a proton beamline, ϵ_{rms} is usually given an alternative definition as the area enclosed by 2σ about the mean of the distribution; here any further references to ϵ_{rms} will refer to the electron definition unless stated otherwise. This then leads to the definition of ϵ_{rms} in terms of the second-order moments of the beam distribution (see [51] for a derivation):

$$\epsilon_{rms} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle x x' \rangle} \tag{2.16}$$

where

$$\langle x^2 \rangle = \frac{\sum_{i=1}^{N} x_i^2}{N} \tag{2.17}$$

$$\langle x'^2 \rangle = \frac{\sum_{i=1}^{N} x_i'^2}{N} \tag{2.18}$$

$$\langle x x' \rangle = \frac{\sum_{i=1}^{N} x_i x_i'}{N} \tag{2.19}$$

N being the number of particles in the beam. ϵ_{rms} may then be normalised as in Eqn. 2.15 to give the normalised RMS emittance:

$$\epsilon_{rms,n} = \gamma_r \beta_r \epsilon_{rms} \tag{2.20}$$

2.2.5 MICE emittance

Having now defined our terms, the expression for the MICE 4D transverse normalised emittance is given (without derivation) by [52, 53]:

$$\epsilon_n^{trans} = \frac{1}{m_0 c} \sqrt[4]{|V_{4\mathrm{D}}(x, p_x, y, p_y)|} \tag{2.21}$$

While the 2D longitudinal normalised emittance is given by:

$$\epsilon_n^{long} = \frac{1}{m_0 c} \sqrt[2]{|V_{2\mathrm{D}}(t, E)|} \tag{2.22}$$

Eqn. 2.21 represents that quantity which MICE is to reduce and measure in order to demonstrate ionisation cooling.

2.2.6 Acceptance

Having described emittance we may now go on and describe the other half of the problem, the *acceptance*. If emittance represents the size of the peg, acceptance refers to the hole into which it must fit, being the maximum emittance able to pass through a transport line or an acceleration structure. In the transverse case this is simply the size of the beam pipe in phase space, while in the longitudinal case it refers to the maximum phase difference from the synchronous particle which allows a particle to remain within the RF bucket (for more detail, see for example [51]).

2.3 Ionisation Cooling

Beam cooling refers to the use of non-conservative processes to reduce the emittance of a particle beam. Due to the short muon lifetime of 2.2 μs at rest, traditional beam cooling techniques cannot be used to reduced the emittance. This led to the proposal in the 1980's of *ionisation cooling*[45], whereby a beam is passed through an absorber, reducing particle momenta in *all* directions through ionisation energy loss, followed by re-acceleration using RF cavities in the longitudinal direction only. This technique is then sufficiently fast to overcome the difficulty of the short muon lifetime.

At the energies relevant to MICE the two most significant processes that affect muons traversing a material are the desired ionisation energy losses and also multiple scattering. Scattering is undesirable as it causes the beam emittance to increase, known as beam *heating*, which acts in opposition to the cooling effect. Stochastic effects in the ionisation energy loss can also lead to heating[54].

The two competing effects of heating and cooling may be related to give an overall expression for the change in beam normalised transverse emittance per unit length. The following expression was originally put forward by Neuffer[55] and is used in the MICE Technical Design Report[53]:

$$\frac{d\epsilon_n}{dz} = \left(\frac{\epsilon_n}{E\beta_r^2} \left\langle \frac{dE}{dz} \right\rangle \right)_{cooling} + \left(\frac{\beta_\perp 13.6^2}{2m_\mu \beta_r^3 X_0 E} \right)_{heating} \tag{2.23}$$

where $\beta_\perp$ is the transverse beta function, E is the beam energy, m_μ is the muon rest mass and X_0 is the radiation length of the material. Thus the rate of cooling is proportional to the ease with which muons ionise the absorber material, while heating is proportional to the beam size ($\beta_\perp$ dependence) and inversely proportional to the radiation length of the absorber material.

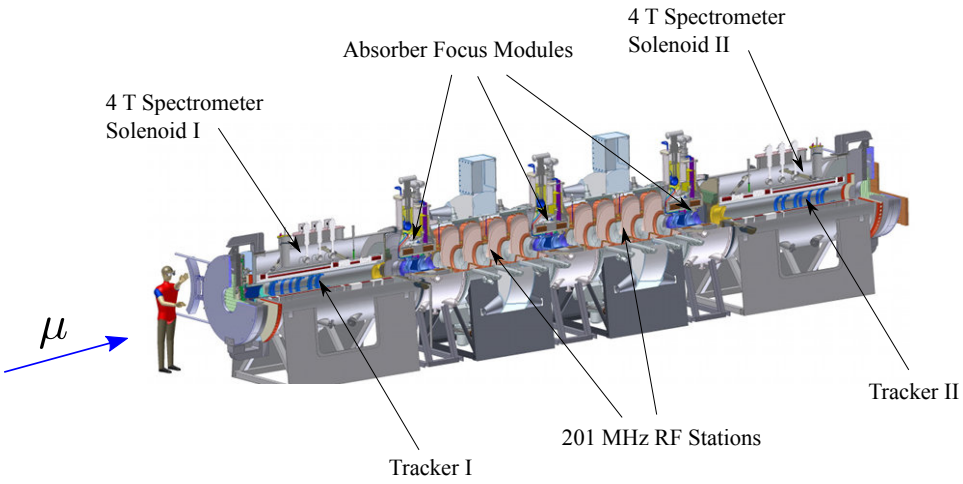

Figure 2.4: The MICE cooling channel, comprising upstream and downstream scintillating fibre trackers with 4 T solenoids, three liquid hydrogen absorber focus modules and two RF stations with four 201 MHz cavities per station[49, 56].

2.4 The MICE Cooling Channel

The cooling channel represents the heart of the MICE experiment, taking the muon beam supplied by the upstream beamline and reducing the momentum in all directions by passing the beam through a series of absorbers, with re-acceleration in the longitudinal direction via RF cavities. The channel is contained within solenoidal magnets to provide beam focusing. The layout of the cooling channel is shown in Fig. 2.4.

2.4.1 Magnetic Channel

Superconducting solenoidal magnets are used to focus and match the MICE muon beam over the whole length of the cooling channel and trackers. The lattice consists of seven magnet assemblies, with eighteen superconducting coils in total. Each tracker uses an assembly of three coils to create a uniform 4 T field in the detector region, together with a further 2 coils each to match

the beam into and out of the cooling channel. Each absorber is enclosed by an assembly of two coils in order to focus the beam and reduce the beta function, so as to maximise the cooling efficiency (a minimal beta function implies less multiple scattering which causes heating, see Eqn. 2.23). Lastly, positioned centrally along the length of each RF station there sits a coupling coil.

2.4.2 Absorber Focus Modules

The full MICE Step VI cooling channel contains three absorber modules, interspersed by the RF cavities. The absorber material chosen for MICE is liquid hydrogen (LH_2). LH_2 was chosen as it has a high dE/dz energy loss rate to provide cooling, while causing only a small amount of multiple scattering (which causes heating), as described by Eqn. 2.23.

A schematic of an absorber module is shown in Fig. 2.5. Each module measures 35 cm in length and 30 cm in diameter, containing 21 litres of LH_2. The windows to allow entry and exit of the beam sit at either end of each absorber module, being cylindrically symmetric and composed of aluminium. Each window is curved to be able to withstand higher pressure and enable thinner construction (minimising scattering of the beam). A second outer safety window is also present for each inner window, to further mitigate the risk of a LH_2 leak.

The absorber modules are also designed to be interchangeable, allowing for materials in addition to LH_2 to be tested. In particular solid absorbers composed of lithium-hydride (LiH) represent an interesting possibility which, while providing a slightly poorer cooling performance, offer practical advantages over the more volatile LH_2.

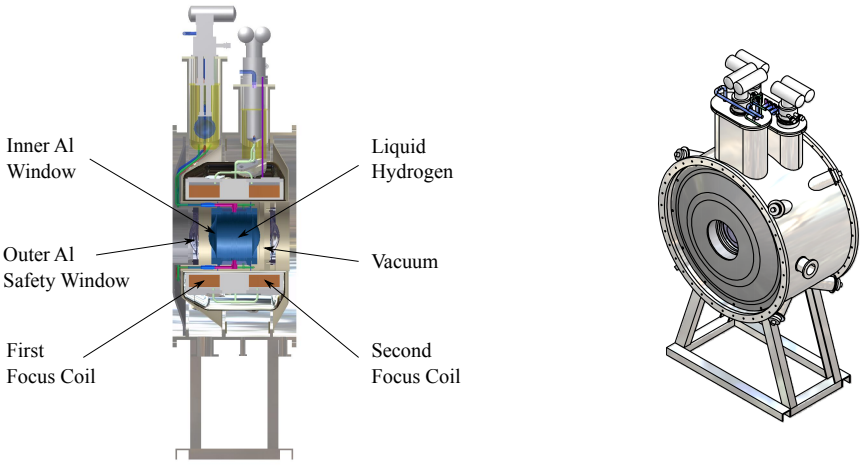

Figure 2.5: Absorber focus coil schematic.

2.4.3 RF Cavities

MICE uses two RF linac stations placed between the absorbers to replace the energy lost by the muon beam in the absorbers, re-accelerating the particles in the longitudinal direction. Each station consists of four normal-conducting copper cavities in the shape of "squashed" cylinders, each cavity of length 43 cm and radius 61 cm. A schematic of a station is shown in Fig. 2.6. The cavities are closed to improve the electromagnetic properties (see [57]) using beryllium windows, which should be almost invisible to the incident muons, to allow passage of the beam.

The cavities operate at 201 MHz and with an accelerating gradient initially of 8 MV/m. The Kilpatrick limit of the cavities, the point at which RF breakdown occurs, is higher at 15 MV/m, but financial constraints on the power supplies available prevent immediate exploitation of this. Further, the high magnetic field present from the focusing solenoids causes a reduction in the gradient at which breakdown occurs. To understand this phenomenon and its implications for muon cooling channels better, a test

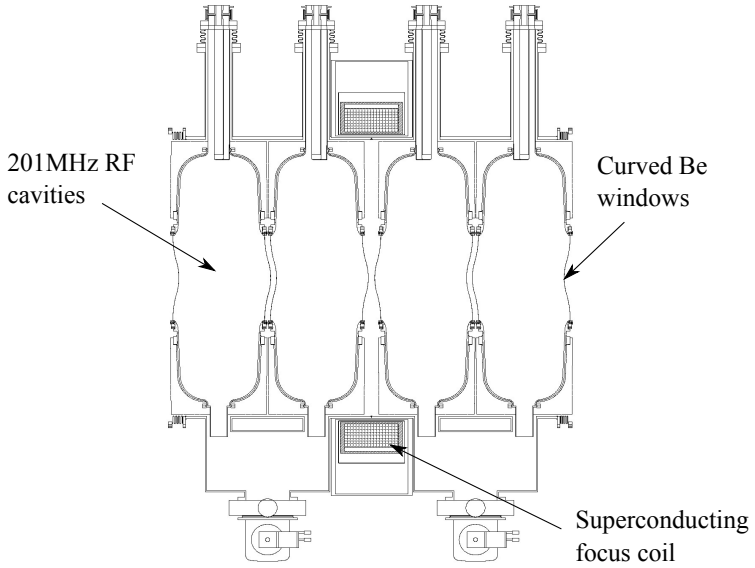

Figure 2.6: An RF station consisting of four cavities[53].

facility has been set up with 805 MHz and 201 MHz cavities at Fermilab[58].

2.5 The Muon Beamline

The MICE Muon Beamline, which supplies the muon beam for the cooling channel proper, took first beam in the spring of 2008. A schematic of the current beamline is shown in Fig. 2.7. The ISIS proton accelerator is used together with a titanium target in order to generate the initial pion shower. These pions are captured by the first quadrupole triplet (Q1-3), and transported to the first dipole (D1), which deflects the beam through a hole in the side of the accelerator vault wall into the MICE Hall, performing the first momentum selection on the beam. Positioned in the wall is a 5 T superconducting decay solenoid (DS) used to increase the pion path length and so the muon content in the downstream beam. The beam then enters a section of the MICE Hall partitioned off from the rest with a concrete

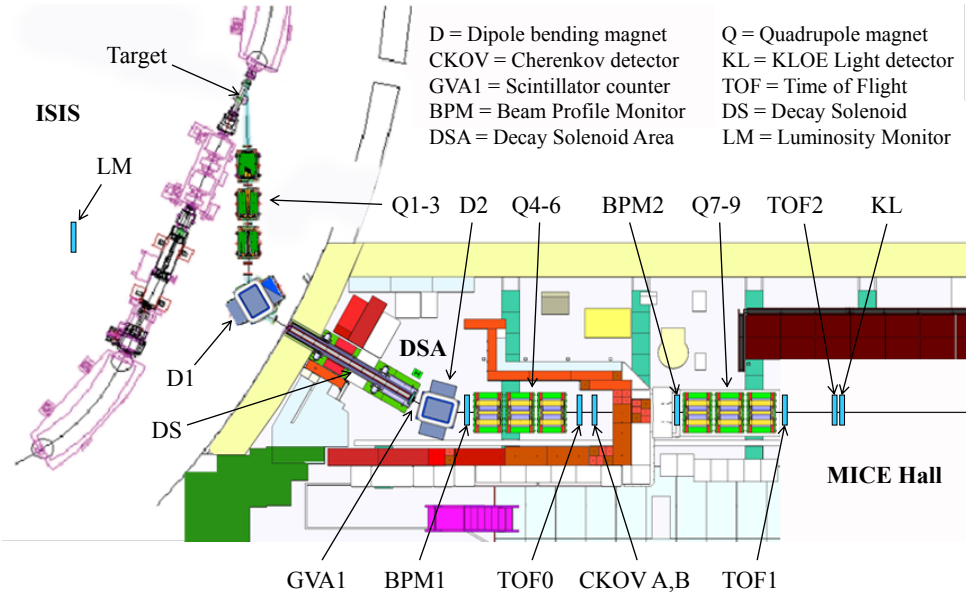

Figure 2.7: A schematic of the MICE Beamline, showing ISIS, the pion production target, quadrupole, dipole and decay solenoid magnets, Luminosity Monitor and beam diagnostic detectors.

wall, known as the Decay Solenoid Area (DSA). Here a second dipole (D2) corrects the beam trajectory and performs the second momentum selection, before the beam enters a second quadrupole triplet (Q4-6) prior to passing out of the DSA into the outer MICE Hall via a sealable aperture in the wall. The beam then passes through a final quadrupole triplet (Q7-9), before some final particle identification (PID) stations and the beam dump. The trackers, absorbers and cavities are to be positioned between the last quadrupole triplet and the final PID stations. Positioned along the length of the beamline are various detectors used for characterising the beam, including a Luminosity Monitor (LM), a rate counter (GVA1), beam profile monitors (BPM1 and 2), Cherenkov detectors (CKOVA and B), Time-of-Flight (TOF) stations (TOF0, 1 and 2) and a "KLOE Light" calorimeter (KL). The host accelerator ISIS, pion production target, the detectors and

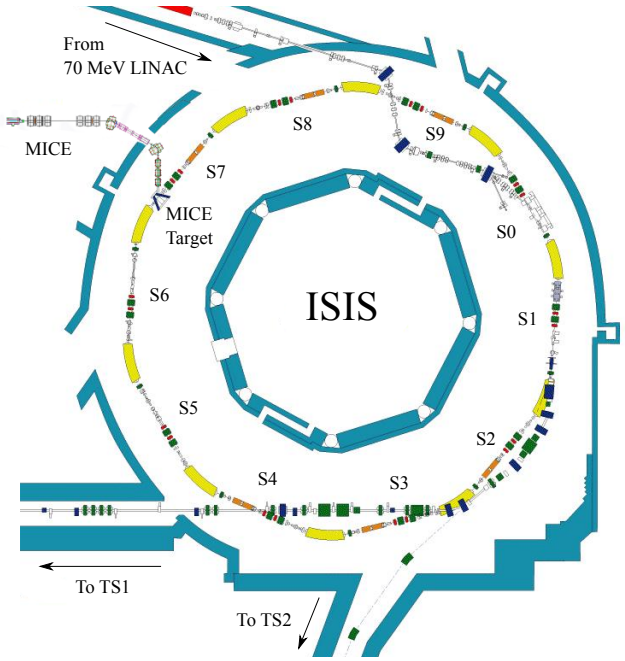

Figure 2.8: A schematic of the ISIS accelerator. Each sector is labelled, with the MICE target sitting at the start of sector 7. TS1 refers to Target Station 1, and TS2 to Target Station 2.

the data acquisition systems are described in more detail in the following sections.

2.6 ISIS

2.6.1 Introduction

The MICE host accelerator is the ISIS 800 MeV proton synchrotron. A schematic of the accelerator is shown in Fig. 2.8. ISIS is a high intensity machine, the primary function of which is as a neutron spallation source, for which its has two extraction lines each with a fixed target (Target Station 1 and Target Station 2). A third target used to generate muons is situated in the extraction line prior to Target Station 1. The injector is a 70 MeV linac

with an H^- ion source. The main ring is of 26 m radius and split into 10 "super periods" or "sectors" (SP0 - SP9) each of length 16.3 m. The ring contains 6 fundamental RF cavities which accelerate the proton beam from 70 to 800 MeV over $\sim$ 12,000 turns in 10 ms. During standard ISIS running each injection-extraction cycle is repeated at a rate of 50 Hz (20 ms between each injection).

2.6.2 Beam Loss Monitoring

The ISIS synchrotron is also equipped with various diagnostic detectors, including an intensity monitor and 39 ionisation chamber beam loss monitors (BLMs), four to a sector, positioned around the inner radius of the main ring. An individual BLM may be referred to by its number, e.g. BLM1 or its position within a sector; hence R8BLM1 refers to the first BLM in sector 8. "Beam loss" refers to particles that are lost from the circulating beam and travel outwards to interact with the accelerator components and surroundings. Excessive beam loss can cause radiation damage to the machine components and cause areas to become too radioactive to permit safe hands-on maintenance.

Time (ms)	**Beam Energy (MeV)**	**Sensitivity (V.ms/proton)**
0	70	2.2×10^{-13}
3	172	2.6×10^{-13}
5	374	4.3×10^{-12}
7	617	1.6×10^{-11}
9	780	3.5×10^{-11}
10	800	3.8×10^{-11}

Table 2.1: Beam loss monitor sensitivity as a function of beam proton energy and time in the ISIS injection-extraction cycle. Taken from[59].

The BLMs in the main ring are positioned $\sim$2 m away from the beam

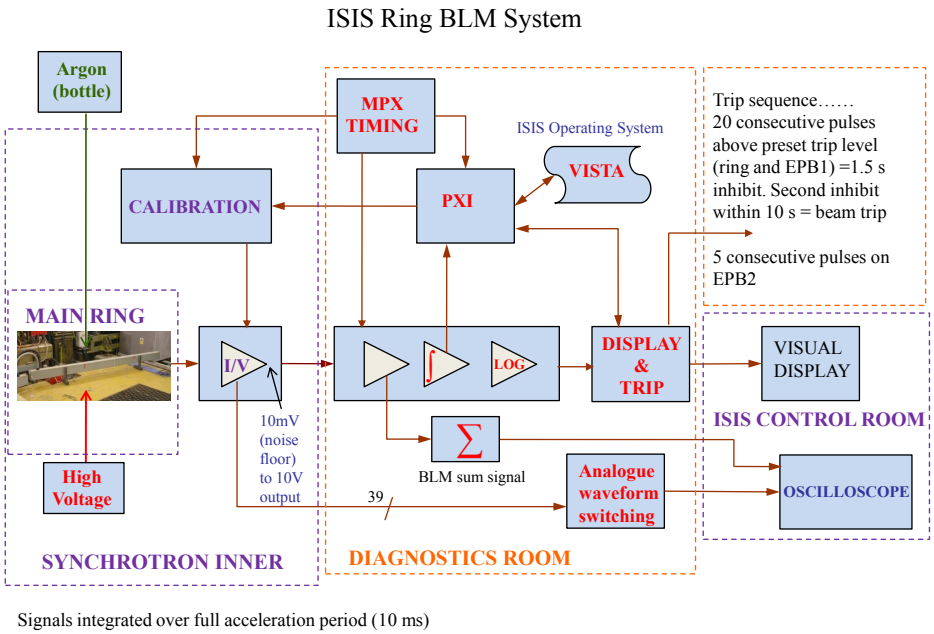

Figure 2.9: Diagram showing the generation and processing of the ISIS beam loss monitor signals.

pipe and consist of 3 m long, 16 mm diameter tubes filled with argon gas with a wire held at an electrical potential running through them[60, 61]. Protons which are lost from the beam interact with atomic nuclei in the surroundings, such as the vacuum pipe, causing neutrons to be emitted. Some of these neutrons then hit the ionisation chamber walls and undergo another nuclear interaction , releasing secondary protons. A fraction of these protons then cause the creation of electrons and argon ions in the gas, which register as a change in voltage[61]. The amount of signal generated is dependent on the energy of the incident protons, and hence the signal response per proton is dependent on time in the ISIS spill cycle. The sensitivity per proton of the monitors as a function of spill time and beam energy is given in Table 2.1.

A diagram illustrating how the signals generated by the BLMs are pro-

cessed is shown in Fig. 2.9[1]. The signal from each monitor is read out by an ISIS DAQ, while three signals, the summed signal for all the sector 7 monitors, the summed signal for all the sector 8 monitors and summed signal for all 39 monitors, are available to the MICE target DAQ system. Two screen displays are also available to MICE showing the output of each individual monitor; one displaying a live feed of each monitor, the other an average value over multiple cycles sent over a network link, an example of which is shown in Fig. 2.10. These however are mainly used as guides when operating the MICE target and are not digitised and recorded by the target DAQ.

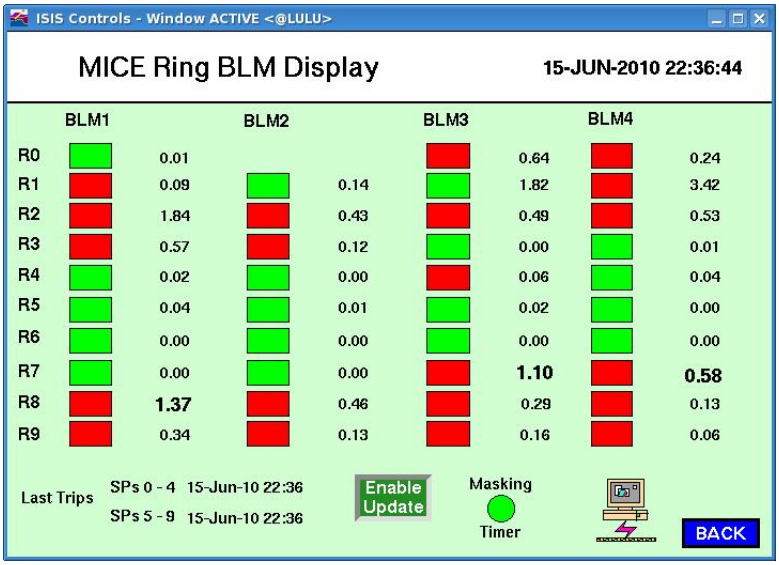

Figure 2.10: A screenshot of the averaged beam loss monitor display available via a network link in the MICE local control room (green indicates a normal value, red a trip condition). This particular shot is taken from MICE run 1989.

[1]Supplied by Di Wright(ISIS, STFC)

2.7 Pion Production Target

2.7.1 Target Design

The MICE pion production target consists of a titanium cylinder of inner radius 2.275 mm and outer radius 2.975 mm. It is situated in sector 7 of ISIS (near the MICE hall), where it can be pulsed into the beam on demand. This pulsing is achieved by attaching magnets to the target shaft, which is then placed inside a 24-coil stator (a magnetic drive), a schematic of which is shown in Fig. 2.11. Upper and lower steel bearings are used to hold the shaft in the correct vertical position, both the shaft and the bearings being coated in Diamond Like Carbon (DLC) to minimise friction and provide a hard-wearing surface (upgraded designs now involve using DLC on Vespel, a hard wearing plastic, to minimise dust generation; see [62]). Attached to the top of the target shaft is a steel optical vane with 157 slits of 0.3 mm width, the centres of the gaps being separated by 0.6 mm. This is read out by a three channel laser optical system in quadrature to give the target position to an accuracy of 150 μm (for details see [63]).

The shaft and stator are mounted in a frame which may be raised away from the beam using a linear stepper motor to make the target safe, or lowered when ready to pulse. In order to intercept the highest energy protons the target is required to enter the beam in the last 2 ms of an ISIS cycle before extraction, overcoming the beam contraction caused by the acceleration, and then withdraw fully from the beam before the next injection occurs. This leads to a target acceleration requirement of order 900 ms^{-2}[62].

The depth the target penetrates into the ISIS beam may be controlled from the MICE local control room by altering the duration a drive current is sent to the stator. The depth that the target penetrates into the beam is

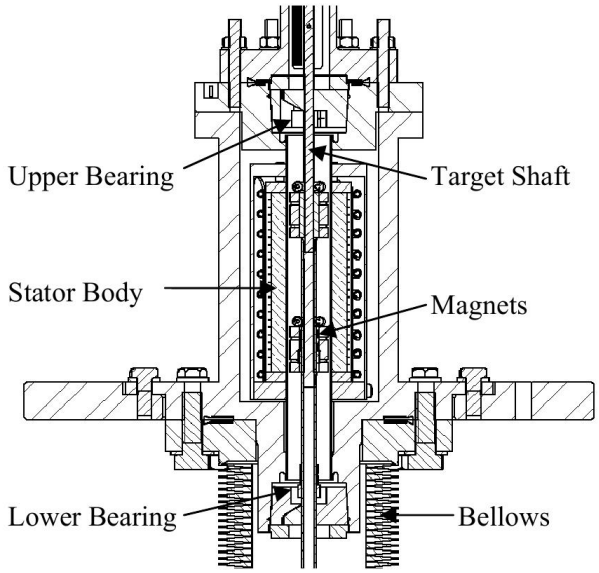

Figure 2.11: A schematic of the MICE target shaft and stator drive, taken from [62]. A 24 coil stator is used to drive a shuttle, consisting of a titanium shaft upon which are mounted permanent magnets to couple to the field produced by the stator. The lower end of the shaft takes the form of a hollow cylinder, which is pulsed into the ISIS beam by the stator. Upper and lower bearings are used to maintain the transverse position of the shaft.

characterised by a parameter known as the "beam centre distance" (BCD), defined as the distance from the target tip at the point of maximum excursion into the beam to the nominal beam centre (normally quoted in mm). Thus deeper excursions into the beam correspond to *smaller* values of the BCD, which cause increased ISIS beam loss levels.

2.7.2 Target Timing

When set to actuate, the target samples only a small subset of the total number of ISIS cycles. The relationships between the ISIS spill cycle and MICE target actuations are shown shown graphically in Fig. 2.12. Each actuation is triggered by the arrival of a signal known as Machine Start

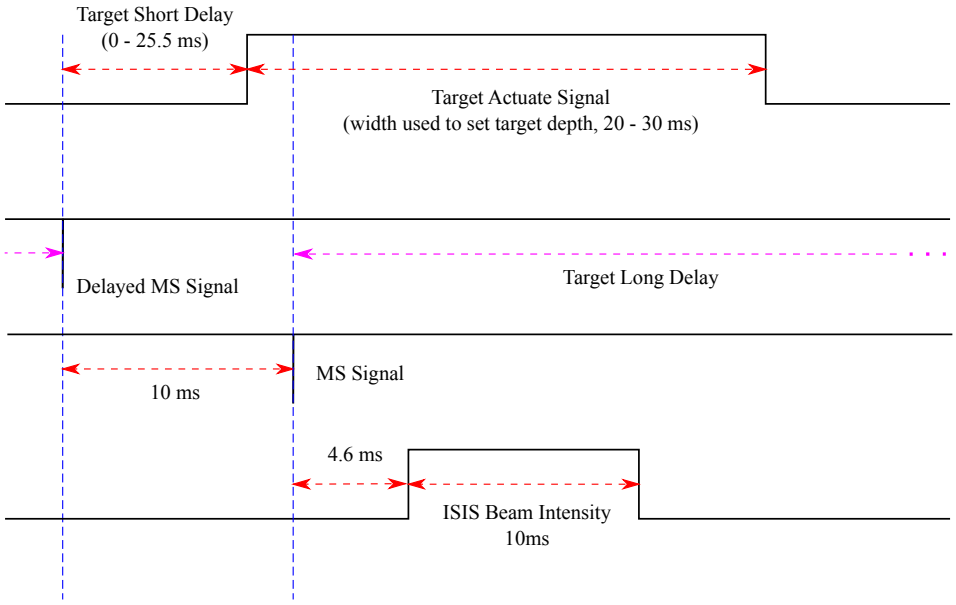

Figure 2.12: ISIS and MICE target timing diagram, based on [64]. "MS" refers to the machine start signal, which identifies on which spills the MICE target may enter the ISIS beam (the *legal* spills) and when to dip. As the MS signal occurs too close to the spill immediately following it (beam extraction occurs 14.6 ms after the MS signal) for the target to enter and leave in time, a delayed MS signal is created to occur before the *next* legal spill, providing an additional 10 ms for the target to enter and leave the beam successfully (hence in the diagram the delayed MS signal is from the MS signal of the preceding legal spill, and so occurs before the MS signal of the spill shown).

(MS) from ISIS. The ISIS cycle immediately following a MS signal is known as a "legal spill" which the MICE target may intercept; it is important that the target only intercepts legal spills, as the ISIS beam loss trip thresholds are altered to accommodate increased losses from MICE on these spills only.

The time between the MS signal arriving and ISIS beam injection is 4.6 ms, with a further 10 ms until beam extraction. This gives ∼14.6 ms for the target to enter the beam at extraction which, even with the large target acceleration, is insufficient. To avoid this, the target actuation is delayed by just under the period of arrival of the MS signal, so that the target can

successfully intercept the *next* legal spill. This is known as the "long delay", with a length given by the MS signal period minus 10 ms, giving the target ∼24.6 ms to enter the beam, which is sufficient. A second delay, known as the "short delay", can then be added to this to set precisely when during this following legal spill the MICE target begins to actuate. Early dips with respect to beam injection tend to produce higher beam losses with lower energy particles. During the period of running covered by this study (2009 - 2010) the MS signal, and so the target actuation, was set to have a period of either 3.2 s or 2.56 s, depending on ISIS running conditions. For more details of the target timing see[64].

2.8 Detectors and Data Acquisition

2.8.1 Luminosity Monitor

In January 2010 a Luminosity Monitor was installed in the ISIS vault. Its position is shown in Fig. 2.13. It is mounted on a stand raising it to the height of the ISIS beam pipe, and set at a distance of 10 m from the target such that a line from the target to the monitor forms an angle of 25° from a tangent to the ISIS ring at the point of the target. This is (almost) the same angle as the MICE beamline, but on the opposite side (the inner side) of the ISIS ring. It provides a gauge of the particle flux produced by the target, independent of the beam loss monitors, allowing for comparison with the rates recorded by other detectors further down the beamline.

A schematic of the Luminosity Monitor is shown in Fig. 2.14. It consists of two sets of scintillators, each set being composed of two blocks. The light from both sets of scintillators are read out by a pair of photomultiplier tubes (PMTs) each. One scintillator / PMT set is shielded by 150mm of

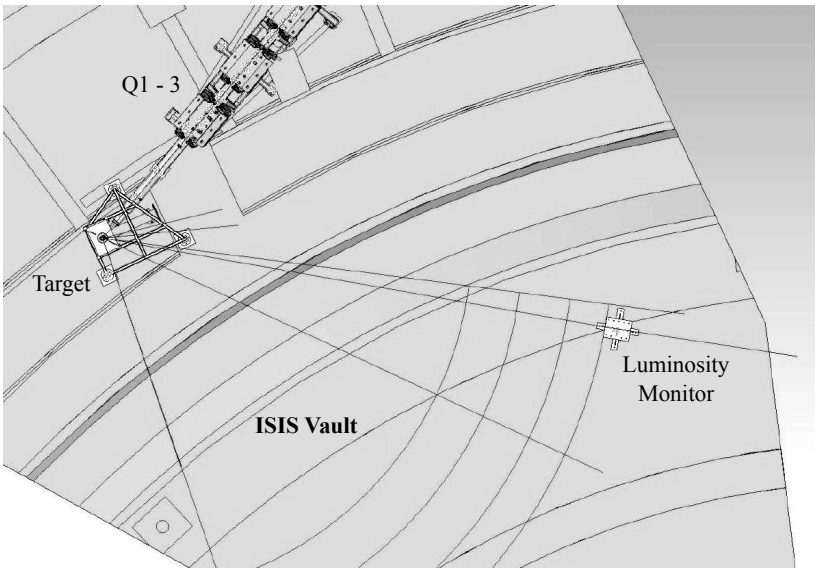

Figure 2.13: The position of the Luminosity Monitor in the ISIS vault. The three magnets at the top of the diagram are Q1-3, the target is shown on the left and the Luminosity Monitor itself on the right. The ISIS beam curves in from left to right, with the centre of the ISIS ring being located off the diagram towards the bottom right. Adapted from [65].

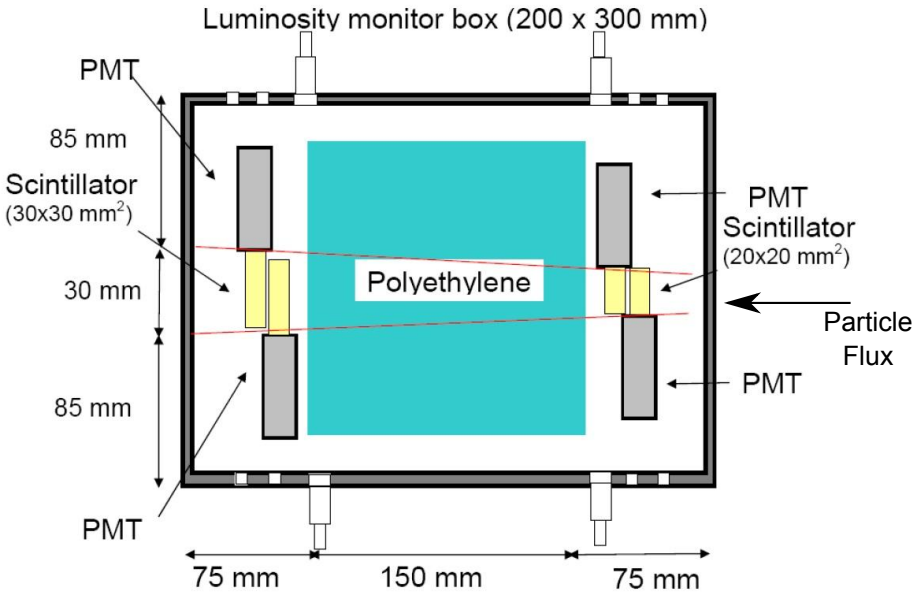

Figure 2.14: A schematic of the Luminosity Monitor, adapted from [65].

polyethylene plastic, and has a cross-sectional area of 20×20 mm^2, while the unshielded scintillator has an area of 30×30 mm^2[65].

2.8.2 GVA1

GVA1 is a simple rate counter; it consists of a $200 \times 200 \times 10$ mm^3 block of scintillator, read out by a single PMT. It was one of the first detectors to be added to beamline, being originally intended only as a temporary measure. In spite of (or because of) its simplicity, it has proved consistently reliable and useful since MICE took first beam.

2.8.3 Beam Profile Monitors

There are two beam profile monitors, known as BPM1 and BPM2, situated in the MICE beamline. Both consist of two layers of Kuraray scintillating fibres, one vertical, one horizontal, read out by one multi-anode PMT per direction (two per BPM). The upstream detector, BPM1, has dimensions of 200×200 mm^2, with 192 fibres per layer, all ganged (that is, read out as a single unit) in groups of 3 (giving 64 channels). BPM2 is larger having dimensions of 450×450 mm^2, with 432 fibres per layer, ganged in groups of 3 in the inner region (16 channels), and in groups of 8 in the outer region (48 channels).

2.8.4 TOF Detectors

There are three time-of-flight stations, TOF0, TOF1 and TOF2, in the MICE beamline, which are used for extracting the time coordinate in emittance measurements, for PID and for beam profile measurements. In this study only TOF0 and TOF1 are used. A schematic of TOF0 is shown in Fig. 2.15, with the other stations following a similar design. Each station

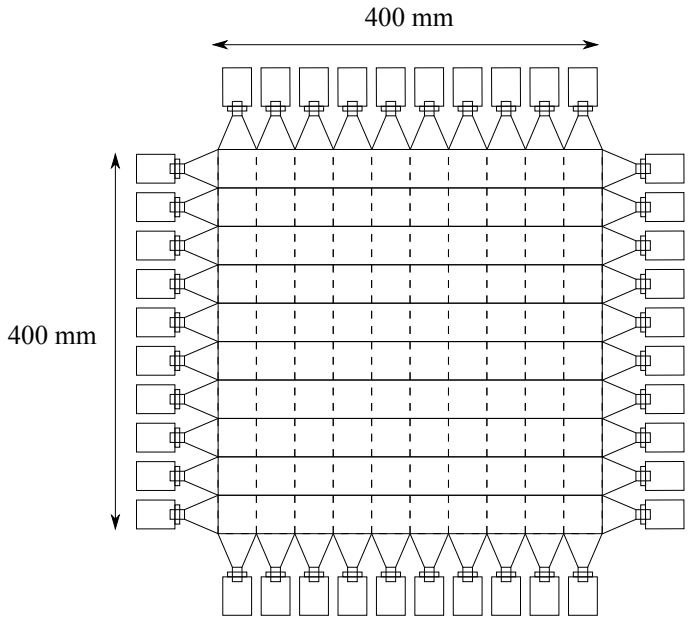

Figure 2.15: A schematic of TOF0. The station consists of two layers of scintillating bars perpendicular to each other. Each bar measures 40 mm in width with 10 bars to a layer. The intersection of two bars each from a different layer is used to form a pixel, allowing a coarse measure of the beam profile to be produced.

consists of two layers of scintillating bars, every bar having a PMT on each end, the bars of the two layers being perpendicular to each other. For each layer there are 10 bars of 40 mm width in TOF0, 7 bars of 60 mm width in TOF1, and 10 bars of 60 mm width in TOF2. Each station has a timing resolution of 50 - 60 ps[66, 67].

The TOFs perform PID by measuring the time taken for a particle to travel between two of the stations. As the distance is known, this gives the velocity, β. The particle momentum is also approximately known, as the beam momentum is set by the field strength of the dipoles. With both the momentum and velocity known the mass may be calculated, identifying the particle. The beam profile may also be measured with a limited resolution,

the intersection of two slabs from different layers on the same station forming pixels from which the profile may be determined.

2.8.5 Cherenkov Detectors

MICE is also equipped with two aerogel Cherenkov counters, used for pion / muon separation, in concert with the TOFs. No suitable single medium was found which would only radiate Cherenkov light in the case of muons and not pions over the whole required momentum range for MICE, leading to the choice of two forms of aerogel with different refractive indices in the two different counters. The refractive indices and muon threshold momenta are $n = 1.07$, $p_\mu = 278\text{MeV/c}$ and $n = 1.12$, $p_\mu = 220\text{MeV/c}$. Each counter is then read out by a set of four EMI9356KA PMTs[66, 68]. Unfortunately the Cherenkovs remain in the commissioning phase and were not available for use during these studies.

2.8.6 KLOE Light Detector

The KL forms half of the downstream electromagnetic calorimeter, which is to be used to distinguish muons from decay electrons, the other half being an electron-muon ranger (EMR) which is still under construction. The KL consists of grooved lead inlaid with 1 mm blue light scintillating fibres[66, 68], following the design used in the KLOE experiment[69]. The KL is not used in these studies, as it was not available for use, and was also not necessary.

2.8.7 The Trackers

The target MICE emittance change is 10%, which is required to be measured with an accuracy of 1%, leading to the requirement of measuring the

beam emittance to within ±0.1%. This is to be achieved by the use of two scintillating fibre trackers, positioned before and after the cooling channel.

Each tracker consists of ∼22,000 scintillating fibres, organised into 5 stations, all surrounded by a 4 T solenoid. The scintillating fibres are of 350 μm diameter, kept small to reduce multiple scattering, and are read out by Visible Photon Light Counters (VLPCs). Each station comprises a carbon fibre frame supporting 3 doublet layers of scintillating fibres, each layer being offset from the other by 120°, as shown in Fig. 2.16 (right). The doublet layers in turn are formed from two rows of scintillating fibres, packed so to remove any dead space, as shown in Fig. 2.16 (left). Detailed accounts of the trackers can be found in [56, 70, 71].

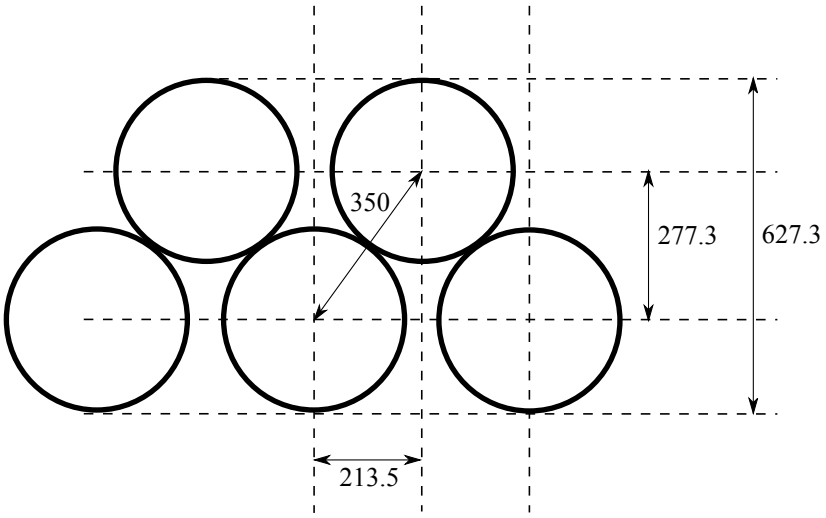

Figure 2.16: Tracker fibre layer schematic. The figure on the left illustrates the layout of the individual fibres in a doublet layer (all lengths are given in μm). Adapted from [70].

2.8.8 Data Acquisition Systems

Two DAQs are used on MICE. The main MICE DAQ handles all the data from the various beamline detectors, while the target DAQ stores the target

trajectory for each pulse, the ISIS beam intensity, and the three summed beam loss signals.

Main DAQ

The main DAQ records data in a time window of variable width, known as the spill gate, for every target dip. The window is set using the ISIS MS signal such that its end always coincides with ISIS beam extraction. The spill gate start may then be altered to change the gate width, usually being set to 3.2 ms. The Luminosity Monitor is an exception to this, having a different spill gate typically set to around 10 ns in order to deal with the very high particle fluxes in the ISIS synchrotron. Signals received from detector PMTs during the spill gate are sent to various digitisation and logic boards before being sent to PCs for recording to disc and tape.

In the case of the TOFs, a current pulse from an individual PMT is sent to a patch panel then to splitter and shaper boards. This signal is then sent to a flash analogue-to-digital converter (fADC) and also to a discriminator. From the discriminator the signal is sent to a time-to-digital converter (TDC) and to a set of logic units used to form "scaler hits". The fADC, TDC and scaler signals are then sent to acquisition PCs, before finally being written to disc on the central data server. In the case of GVA1, the BPMs and the LM only scaler hits are recorded (profile data from the BPMs is currently recorded using the EPICS-based[2] MICE controls and monitoring system).

Scaler hits provide a simple measure of the particle flux through a detector in a given spill. For GVA1 a scaler hit simply requires the PMT to produce a voltage above the discriminator threshold. In the case of the

[2]See http://www.aps.anl.gov/epics/

BPMs, a scaler hit is formed when any anode from the first plane of fibres (the logical OR of the anodes) fires in coincidence (within 20 ns) with any of the anodes in the second plane of fibres (the logical AND of the two planes). For the TOFs, a scaler hit occurs when two PMTs attached to the same slab fire in coincidence (within 100ns), in logical OR over all the bars, then in logical OR over both planes (any coincident hit in any bar causes a scaler hit).

The fADC and TDC signals can be used for more complex operations and require the system to be triggered in order to be recorded, in addition to falling within the spill gate. The trigger station is normally set to be TOF0 or TOF1 (TOF1 is used exclusively here), a trigger being formed by the same logic as a scaler hit. Any fADC or TDC signal within 1.28 μs of the trigger is recorded. A dead-time window, during which no triggers will be accepted, occurs during approximately the same period.

Data extraction and analysis from the main DAQ is performed by applications within the standard MICE analysis and simulation framework G4MICE[72], and is described in Section 3.2.3.

Target DAQ

The target DAQ is used to digitise signals from the target position readout, the ISIS beam intensity, the summed signal from all four beam loss monitors in sector 7, all four monitors in sector 8, and all 39 monitors over the entire ISIS ring.

The digitisation is performed using a PC equipped with a National Instruments PCI-6254 card. The data is sampled at a rate of 100 kHz in a 50 ms window, giving 5000 samples per pulse. When ISIS is running at its standard repetition rate of 50 Hz, this leads to the capture of several ISIS

spills only the first of which contains the MICE target pulse (unless the target is still scraping the beam at injection of the following cycle). The data is stored as one gzipped ascii file per pulse, organised by timestamp (the target DAQ does not have a "run" structure). Extraction and analysis of the data is not performed using G4MICE, but rather the more specialised *MICESoftware* package[3]. The analysis of the target DAQ data is described in Section 3.2.1.

2.9 MICE Particle Rate and ISIS Beam Loss

Particle rate in the MICE beamline may be increased by raising the number of protons-on-target per spill. This can be done by increasing the exposure time of the target to the beam, achieved by decreasing the target short delay so that the target intercepts the beam earlier in the ISIS spill, or by increasing the volume of material presented to the beam, achieved by increasing the maximum depth which the target is dipped into the beam (decreasing the BCD). In practice the former is normally held constant and the latter used to vary the particle rate.

An increased number of protons-on-target however also causes an increase in ISIS beam loss levels which, as explained in Section 2.6.2, is undesirable. This then motivates a study of the effect of the MICE target parameters (the BCD and short delay) on the induced ISIS beam losses and MICE particle rate, and in particular the muon rate present in the MICE beamline as function of induced beam loss. This forms the central subject of this thesis. In Chapter 3 the methodologies of the various studies undertaken to pursue this are described, in Chapter 4 the results are presented and discussed, in Chapter 5 simulations of the ISIS beam in the presence of

[3]By James Leaver (Imperial College London).

the MICE target are presented, and in Chapter 6 the study is concluded, summarising results and open issues, and outlining future directions.

Chapter 3

Methodology

If we knew what it was we were doing, it would not be called research, would it?

Albert Einstein

3.1 Experimental Methodology

Five studies are reported here, taken on the 6th November 2009, the 15th and 16th June 2010 and the 14th August 2010. Four studies vary the target depth in order to alter the induced beam loss, while the fifth, consisting of the last runs taken on the 14th August, uses the target short delay to alter the induced losses. Each study consists of a set of runs where all the parameters of the experiment, such as the optics, DAQ spill gate, trigger, *etc.* are held constant, with the exception of the target depth for the first four studies or the short delay for the last. The target depth / short delay is held constant for each individual run but varied between runs; thus every run represents a single data point in a given study. A run itself is a set of data for many concurrent target actuations stored together by the DAQ; the run numbers used for each study can be found in Appendix A.

3.1.1 Depth Studies

In the case of the depth studies the usual procedure is to set up MICE as for any other data taking shift and move the target until the highest beam loss levels permitted are achieved. The target short delay is then set to ensure the spill produced by the target is aligned with the DAQ spill gate. This is necessary as the timing of the target maximum point of excursion within the ISIS beam gets later with respect to the injection - extraction cycle as the target maximum depth is increased. As these studies represent the deepest excursions of the target to date, it is particularly necessary to ensure the spill generated matches the gate. The gate width must also be set such that the particle rate does not overload the DAQ.

During the current phase of MICE the usual DAQ spill gate width is 3.2 ms, but the high rates present in some of these studies have necessitated the use of shorter gates as well, of 1 ms and even 0.5 ms. In addition, a 1 ms gate also has the virtue of being the gate width that is expected to be used for the later stages of MICE when the cooling channel is present. This then allows for an easier extrapolation of particle rates from the present data to the future cooling channel.

Once the correct target delay and gate settings have been determined a run consisting of a several hundred pulses is taken. Once completed the target depth is altered to achieve the next beam loss setting and another run taken, and so on until data for all of the desired beam loss settings has been achieved.

Due to concerns for ISIS at very high loss levels, it has not always been possible to start each depth study from the highest beam loss position, meaning that best guesses must be used for the short delay and DAQ spill gate based on previous studies. This was necessary for the August 2010

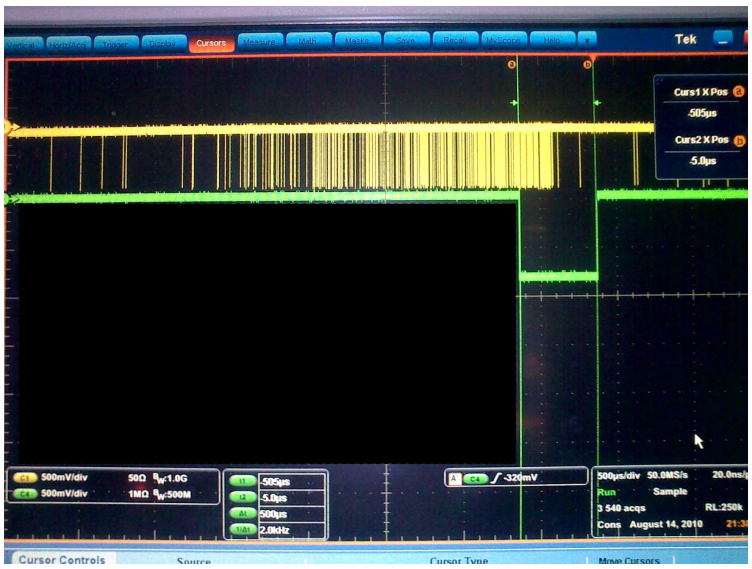

Figure 3.1: Picture of the DAQ oscilloscope taken during one of the highest beam loss runs of the August 2010 study. The yellow lines indicate trigger requests and the green line the main DAQ spill gate. The sharp cut-off of the triggers indicates ISIS beam extraction. This is not aligned with the end of the spill gate, illustrating the mis-alignment of the gate for this study.

study, which achieved the highest beam loss levels to date, leading to the requirement that the study begin at the lowest beam loss levels and each subsequent study point be taken in succession up to the maximum. This contributed to an error being made in the setting of the DAQ spill gate for this study, giving a partial misalignment of the spill gate, illustrated in Fig. 3.1. The hard cut off of the triggers indicates that the end of the spill gate was not aligned with ISIS beam extraction as it should be, leading to only part of the spill being captured. The gate was already set to a short width of 0.5 ms to deal with the high expected particle fluxes, and this mis-alignment served to decrease the gate artificially by a further $\sim 50\%$. In addition the target short delay was set very early to counter the effect of the very deep dip depth causing the target dip to move later in the ISIS cycle.

This leads to complications in the August depth data compared with the other studies.

While the experiment parameters are held constant throughout each particular study, the parameters between studies do vary, including MICE beam optics and available detectors. The various conditions present for each study are given in Table 3.1. The different conditions should be borne in mind when comparing data from different study dates.

3.1.2 Delay Study

In the case of the delay study, a suitable target depth was chosen (30.3 mm BCD), and a long spill gate of 10 ms set. The rest of the experiment conditions are then the same as that of the 14th August 2010 study, as shown in Table 3.1 (the delay study following on directly from the depth study of that date). Runs were then taken at various values of the short delay, the delay values themselves being used to define a run. This is in contrast to the depth studies where the beam loss values where used to define where to take a run, with the target depth merely being used to induce these desired losses.

Parameter	**6 Nov 09**	**15 June 10**	**16 June 10**	**14 Aug 10**
Max Beam Loss (V.ms)	4.7	2.8	3.4	6.0
Target Short Delay (binary)	0010001011	0010000011	0010000011	0010000000
Target Short Delay (ms)	13.9	13.1	13.1	12.8
Optics	$+\ \pi$	$-\ \pi \rightarrow \mu$, no Q3	$+\ \pi \rightarrow \mu$, no Q3	$+\ \pi \rightarrow \mu$
Proton Absorber	No	No	No	83 mm
Detectors	GVA1 BPM1 BPM2 TOF0 TOF1 x	GVA1 x BPM2 TOF0 TOF1 LM	x x x TOF0 TOF1 LM	GVA1 x BPM2 TOF0 TOF1 LM
DAQ Gate Width (ms)	0.5	3.2	1	0.5
DAQ Trigger	TOF1	TOF1	TOF1	TOF1
Approximate pulses per run	200	400	400	200

Table 3.1: Study conditions. "+" indicates a beamline set to transport positively charged particles, "-" negatively charged particles. The November 09 study took place before the arrival of the Luminosity Monitor, BPM1 failed during the June and August 10 studies, and BPM2 failed and GVA1 behaved erratically for the 16th June study. The magnet Q3 was unavailable for the June studies requiring a modification of the $\pi \rightarrow \mu$ optics (see Appendix B for the actual magnet settings). The delay study was performed with the conditions listed for the August 10 depth study (it having been performed immediately after the depth study), but with a 10 ms spill gate, a 30.3 mm target BCD and a varying target short delay.

3.2 Analysis Methodology

A flow diagram, illustrating the analysis procedure used to go from raw beam loss, particle rate and TOF data to final plots is shown in Fig. 3.2. The procedure may be split into five sections: beam loss, target delay, particle rate, combined beam loss and particle rate, and PID, which are discussed in the sections which follow.

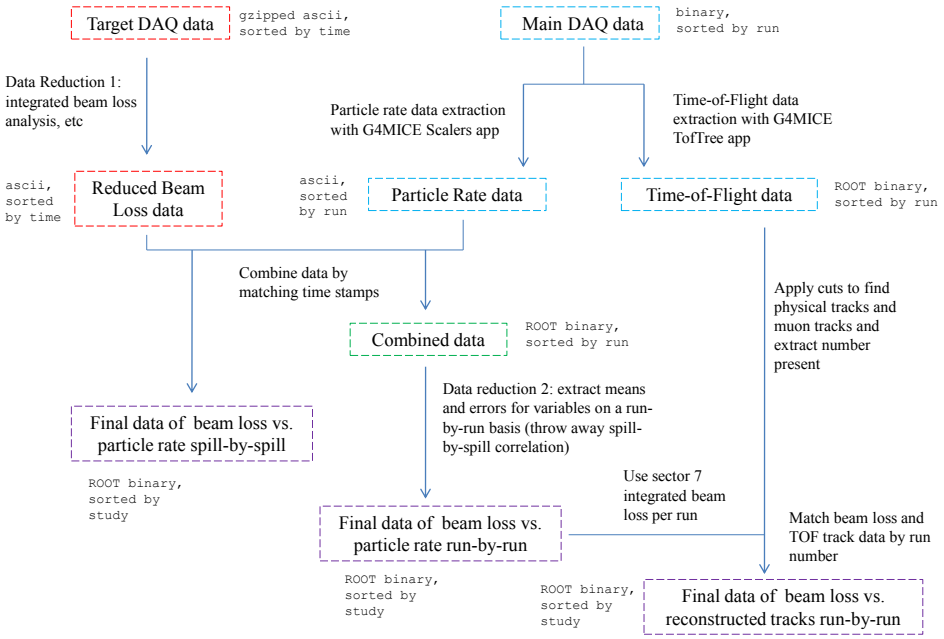

Figure 3.2: Flow diagram of the analysis procedure.

3.2.1 Beam Loss Analysis

As discussed in section 2.8.8 three beam loss signals are available to the MICE target DAQ for digitisation and analysis: the summed voltage from all four monitors in sector 7, the summed voltage from all four monitors in sector 8, and the summed voltage from all 39 monitors for the whole ISIS

ring. The data from the target DAQ are stored on the target control PC on the micenet subnet. These are ascii and human-readable, but compressed and stored as gzipped files, with one file per target actuation. The output for one cycle is shown in Fig. 3.3. From these data, "figures of merit" are defined for the beam loss produced by the action of the MICE target for each spill. The most important of these is found by integrating the summed signal of the 4 sector 7 BLMs. The integration is performed by a piece of code known as BeamLossAnalysisPart2, which operates within the MICETrackerSoftware package, which also serves the target DAQ system itself. As such both the online display and offline analysis use the same beam loss analysis classes.

For each data file (corresponding to one target dip), the beginning and end of each ISIS cycle are found by looking at the ISIS beam intensity column. This has a sharp turn on at injection and a sharp cut off at extraction, leading each cycle to create a "top hat" function in the beam intensity. Each file normally contains either one or four ISIS cycles, depending on the ISIS running mode. The average duration of a cycle is found by taking the mean over all the cycles in the current file, and this average together with the first injection edge found, is used to define the first cycle. Once found, the sector 7 beam loss is integrated, entry by entry, from the beginning to the end of the first cycle[1], to produce the integrated sector 7 value. Note that there are two integrations here, the summing of the voltage signals from the four sector 7 monitors, and the integrating of this signal over the ISIS injection - extraction cycle. Similar procedures are performed to produce other figures of merit at the same time, including the integrated losses in sector 8, the integrated losses over the whole ring, and peak losses in sector 7, sector 8 and over the whole ring.

[1]The MICE target should only intercept the first ISIS cycle recorded in the target DAQ. If this is not so it indicates a probable error in the target timing.

Each file in a particular dataset is looped over to create a new "reduced" data file, with one entry now holding the beam loss figures of merit for a particular target dip. Target data are also recorded, including the target BCD (defined in section 2.7.1).

ISIS tends to use another value as its main figure of merit when considering beam loss due to MICE, that of the integrated voltage observed in the single monitor R8BLM1. As this is not available in the MICE target DAQ at present, only via a screen display in the MICE local control room (shown in Fig. 2.10), the integrated sector 7 value will be used here.

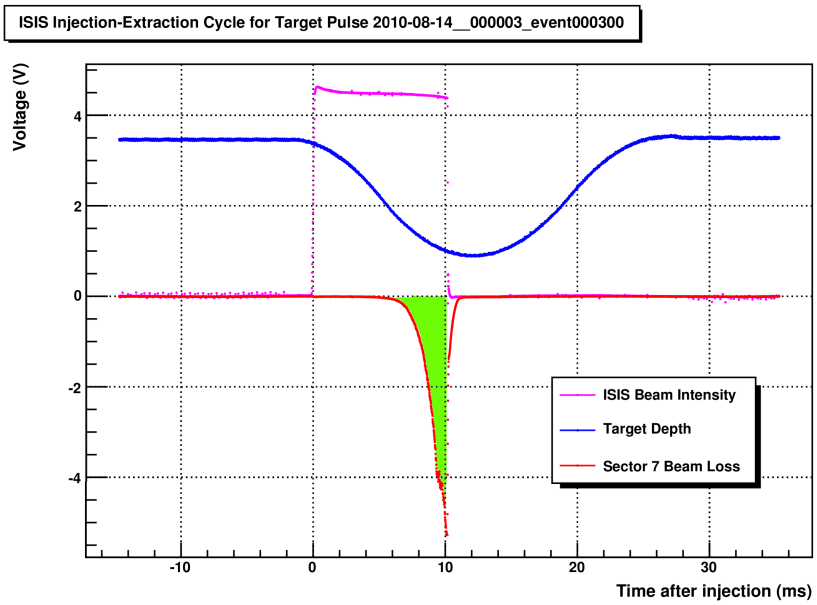

Figure 3.3: A high beam loss event as recorded by the MICE target DAQ system. The area shaded in green represents the beam loss of the summed signal for all four monitors in sector 7, integrated over a whole ISIS cycle. This value, here referred to as the integrated sector 7 beam loss, is one of the standard measures of beam loss per spill produced by the action of the MICE target. The area after extraction, still under the beam loss curve but not shaded, is an electronics artefact and not physical.

3.2.2 Target Data Analysis

In addition to the beam loss signals, the target DAQ also processes the target position signal. The readout is as a drive voltage, but may be calibrated to beam centre distance using the following:

$$\text{BCD (mm)} = \text{Gradient} \times \text{Drive Voltage (V)} + \text{Offset} \tag{3.1}$$

where Gradient = 15 V^{-1} and Offset = 12.35 mm.

The time with respect to ISIS beam injection at which the target reaches its maximum excursion into the beam, which shall here be referred to as the *target dip time*, may also be calculated for use in the target delay study. The target dip time is both a function of the short delay and the maximum BCD (the target parabola "swings out" further from injection as it requires more time to achieve deeper depths). It is calculated for each dip by finding the DAQ sample numbers for injection and for the minimum of the target position curve. The difference is then taken and calibrated to milliseconds using the DAQ sampling rate (1 MHz).

3.2.3 Particle Rate Analysis

The particle rate analysis falls into two parts: extracting the relevant scaler channel data from the main DAQ output files and extracting the time-digital-conversion (TDC) channel data for the reconstruction of TOF tracks. The former is used to give the total particle flux through the beamline, while the latter is used for particle identification (PID). As mentioned in section 2.8.8, an important difference between the two data types is that the scaler channel data record every hit created in a detector during the DAQ spill gate, while the TDC data only record those hits which coincide with

(occur within ~ 1.28 μs of) a particle trigger. This will become important later on when comparing scaler data and reconstructed TOF tracks (see section 4.5).

The scaler channel data are extracted from the main DAQ output files by the G4MICE application Scalers[2]. This outputs the number of hits per spill for every rate counter, together with a time stamp and the spill gate width. A spill is here defined as the particle shower resulting from one target dip intersecting an ISIS cycle (also referred to as a spill), not to be confused with a particle event (when the particle trigger condition is met, of which there are many per spill), or a run (a series of many spills taken consecutively, here with the same experiment parameters).

A similar procedure is then performed for PID. The TDC data must be extracted from the main DAQ output files, as in the case of the scaler channels, and the tracks of the particles between each TOF station reconstructed. Once done, the time of flight between two stations for each particle can be determined. As the distance between the stations is known, the velocity, β, can be calculated which, taken with the momentum (selected by the dipoles), can be used to determine the mass and so which particles are present. The data extraction and reconstruction of the particle tracks is performed by the TofTree and DataQualityCheck applications[3], using the TofRec class[4], again part of the G4MICE framework. The data are extracted as values per run, in contrast to the scaler data which are available spill-by-spill. This value per run is turned into an average value per spill using the number of spills present in the run. Getting the number of spills per run is not entirely trivial; the number of target pulses per run is recorded in the Run Condi-

[2]Created by Vassil Verguilov (University of Geneva). The application ScalersAnalysis has also been used but is now antiquated.

[3]Created by Mark Rayner (Oxford University).

[4]Created by Yordan Karadzhov (University of Sofia).

tions spreadsheet, but by hand not automatically. Further the target may pulse before the DAQ has been set to start the run, leading to a variance between the number of pulses and the number of spills. Here the number of spills is taken to be equal to the number of spills extracted from the scaler channels that are successfully combined with the beam loss data.

3.2.4 Combined Analysis

The particle rate data must be combined with the beam loss data, due to the two data streams being captured by separate DAQ systems. The is done by ensuring that both DAQ PCs have their system clocks properly synchronised and then combining the data on a spill-by-spill basis by matching the time stamps associated to the data for each target dip.

There is a sufficient variation in the beam loss, primarily due to fluctuations in the ISIS beam, even when the target depth is held constant, that correlations should be seen even by looking at data for single runs where all the experiment parameters have been held constant. If no spill-by-spill correlation for such a run is observable it points to a flaw either in the combination algorithm or in one of the time stamp readouts (in practice the beam loss time stamp readout is very reliable and trusted whereas the particle rate time stamp readout has in the past been problematic). These spill-by-spill correlations may be looked at for single runs, or over whole study sets.

Even when a spill-by-spill correlation is not possible, it is in general still possible to do a run-by-run correlation, whereby the data for each channel is averaged over each run, giving a single rate-per-spill value for each run. These can then be correlated even when originating from different DAQs providing the start and end times of each run are accurately known. The error on each run-averaged data point is taken as the error on the mean for

the data from the whole run, using the standard formula:

$$\epsilon = \frac{\sigma}{\sqrt{n}} \tag{3.2}$$

where ϵ is the error, σ the standard deviation, and n the number target pulses in the run.

In the case of the June and August 2010 studies it has been possible to correlate the data on a spill-by-spill basis, and for these studies plots of both spill-by-spill and run-by-run averaged data are shown (see section 4.3). For the November 2009 study however, it has not been possible to perform a spill-by-spill correlation. This was due to difficulties in reliably extracting the time stamp from the main DAQ data files. It was still possible however, to do a run-by-run correlation as, although the time stamp for each individual spill is not trusted in the study, the start and end time recorded for each run is believed to be good to within a spill or so. In addition, these times extracted from the main DAQ data stream can be further checked by looking at the beam loss data from the target DAQ, where the start and end of runs are more apparent. Hence averages for particle rate data and beam loss data from each run may still be calculated and plotted against one another.

Chapter 4

MICE Particle Rate and ISIS Beam Loss

Audaces fortuna iuvat.
(Fortune favours the daring.)
Virgil, Aeneid 10,284

The results of the ISIS beam loss and MICE particle rate studies are here presented. The four studies where the target depth is used to induce beam loss are presented collectively, organised by the relationship being investigated. The variation of induced beam loss with target depth is discussed in Section 4.1; the relationship of sector 7 beam loss with sector 8 beam loss in Section 4.2; particle rate in the MICE beam line and its relationship to ISIS beam loss in terms of the total rates observed in the detectors in Section 4.3; the rate for each particle species, in particular muons, in Section 4.4; and the relationship between scaler hits in the detectors and subsequent reconstructed TOF tracks in Section 4.5. Lastly the study into effect of the target short delay on induced beam loss is presented in Section 4.6.

4.1 Beam Loss and Target Depth

Histograms of the beam loss produced for each depth study as a whole (containing all the runs) are shown in Figs. 4.1 - 4.4. Individual peaks correspond to one run at a particular beam loss / target BCD setting. The width of each peak gives a rough measure of the variation in beam loss observed even for a nominally fixed target BCD. The variation is thought to be due primarily to variations in the ISIS beam, rather than in the target position. This spread in beam loss allows the various spill-by-spill plots of beam loss against particle rate to show trends to higher beam losses compared with those which appear in the run-by-run averaged plots.

In the case of the November 2009 study, shown in Fig. 4.1, the peaks from each run are well seperated, and become broader with increasing beam loss levels. For the June and August 2010 studies the peaks become less well separated, and again acquire a large width at higher losses. This confirms what had already been known operationally, that the precision with which the target BCD can control induced beam loss decreases significantly as the beam loss levels increase. This is probably a result of the ISIS beam becoming more dense closer to the beam centre, something which may be further probed by looking at the induced beam loss levels as a function of target BCD.

The results showing this variation of beam loss with target depth (in terms of BCD) on a run-by-run basis for all studies are shown in Fig. 4.5, while spill-by-spill plots are shown in Figs. 4.6 - 4.9. At high BCDs the curve is clearly non-linear, while for smaller BCDs the curve flattens and it becomes difficult to tell whether the relationship is tending towards linearity or not. In particular the August study, shown in Figs. 4.5 and 4.9, still exhibits non-linear behaviour even at smaller BCDs. It should however

be noted that for the highest beam loss points in this study, ISIS was experiencing problems with its ion source, leading to large fluctuations in the beam current.

The non-linearity in the observed beam loss increase may arise from a number of factors. Firstly it may indicate a non-linear increase in the density of the ISIS proton beam with depth (as would naturally be expected). In addition, as the target dips dynamically during the ISIS cycle, rather than simply being at a fixed depth throughout, smaller BCDs may be expected to lead to a non-linear increase in the total of amount of material shown to the beam, integrated over the time the target is present in the beam (the ISIS cycles for a subset of the August depth study runs are shown later in Fig. 4.38 and help to illustrate this point).

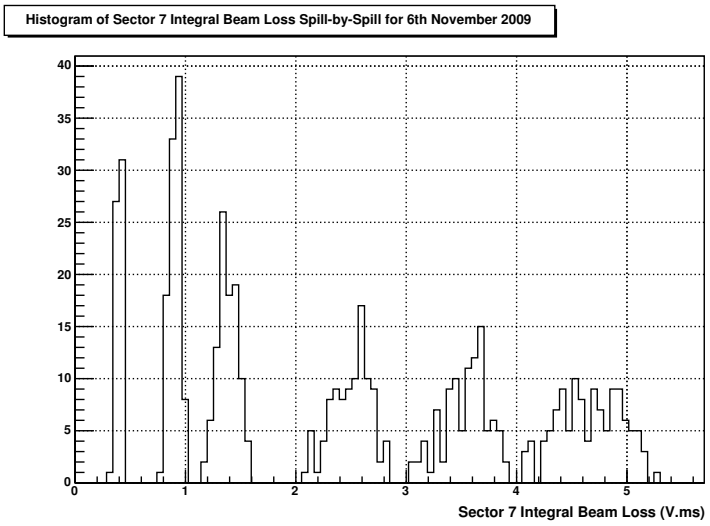

Figure 4.1: Spill-by-spill histogram of the integrated sector 7 beam loss for the whole 6th November 2009 study. The peaks correspond to individual runs, so that each peak represents a fixed target BCD. The width of each peak gives a rough measure of the variation observed in the beam loss levels even for constant target BCD. Most of this variation is thought to be caused by fluctuations in the ISIS beam and increases with smaller target BCDs.

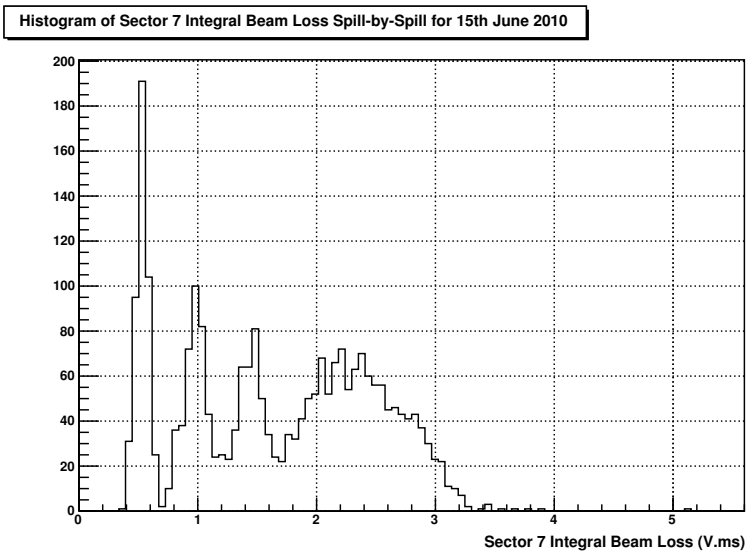

Figure 4.2: Spill-by-spill histogram of the integrated sector 7 beam loss for the whole 15th June 2010 study. Individual peaks corresponding to particular runs at particular target BCDs are discernible at lower beam losses, but merge at higher losses as the peaks broaden.

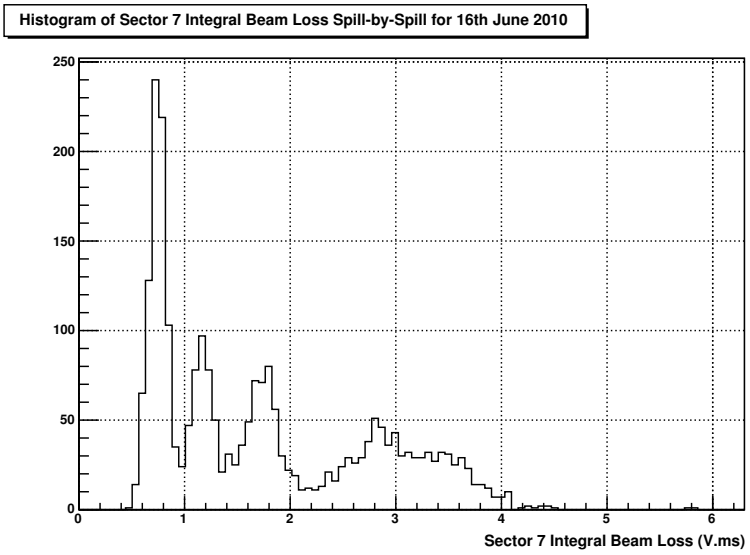

Figure 4.3: Spill-by-spill histogram of the integrated sector 7 beam loss for the whole 16th June 2010 study.

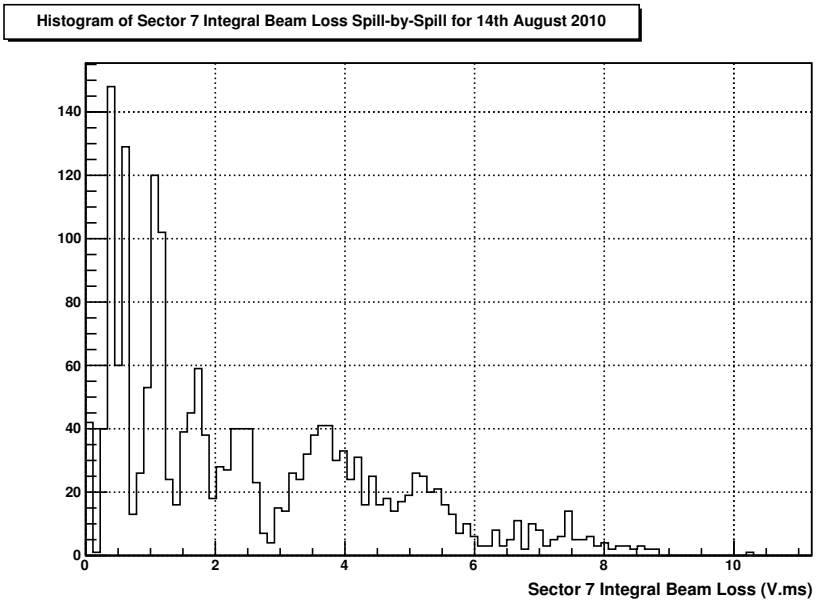

Figure 4.4: Spill-by-spill histogram of the integrated sector 7 beam loss for the whole 14th August 2010 study. The data in the tail on the far right represents the highest beam loss levels achieved to date. The broad width there indicates the large beam loss variation for a given target BCD at these high losses, an effect which may be being further enhanced by ISIS ion source difficulties leading to a fluctuating beam current for the highest loss levels of this study.

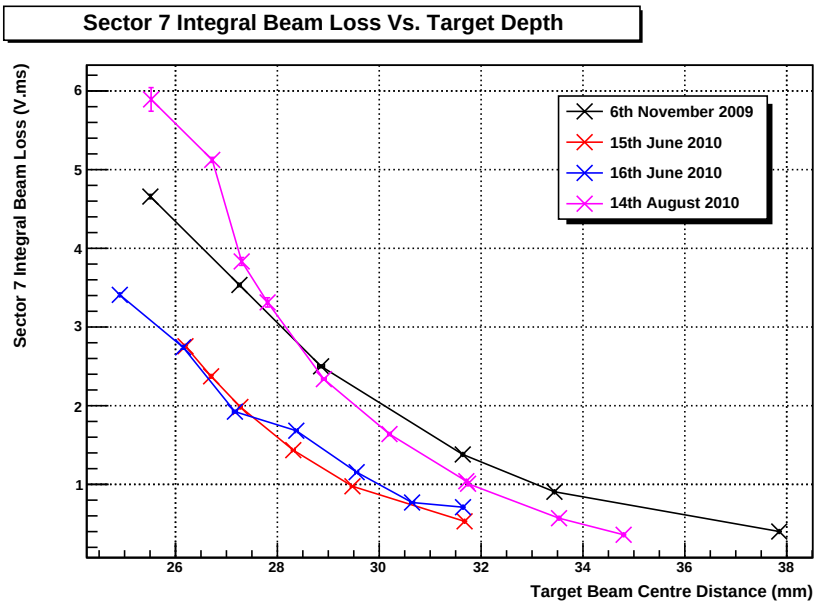

Figure 4.5: Beam loss as a function of MICE target depth in terms of BCD (smaller BCDs on the left of the graph imply deeper excursions into the ISIS beam) for all four studies. Simple lines between points are shown rather than fits. The variation appears overall to be non-linear, possibly the result of a non-linear increase in integrated beam - material exposure with increasing depth, or a non-linear increase of beam particle density with depth.

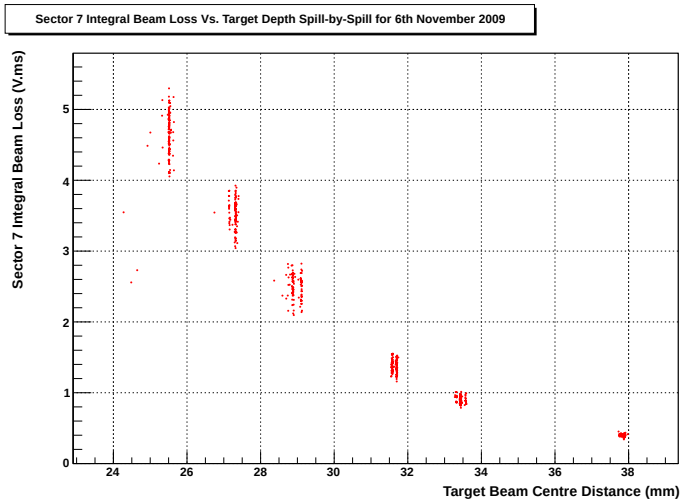

Figure 4.6: Spill-by-spill plot of beam loss as a function of MICE target BCD for the 6th November 2009 study. The clumps of points represent individual runs where the target BCD is held constant. The increasing vertical width of the clumps at deeper depths represents the greater variation of beam loss produced at higher beam loss levels. This is in contrast to the target BCD variation which remains small throughout.

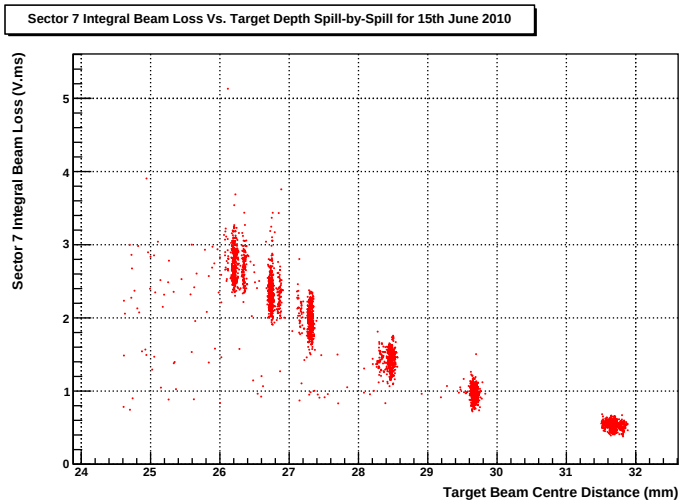

Figure 4.7: Spill-by-spill plot of beam loss as a function of MICE target BCD for the 15th June 2010 study.

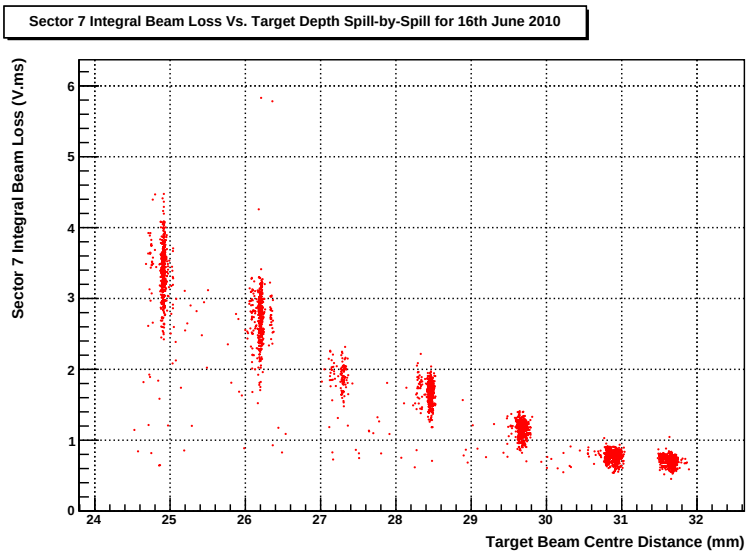

Figure 4.8: Spill-by-spill plot of beam loss as a function of MICE target BCD for the 16th June 2010 study.

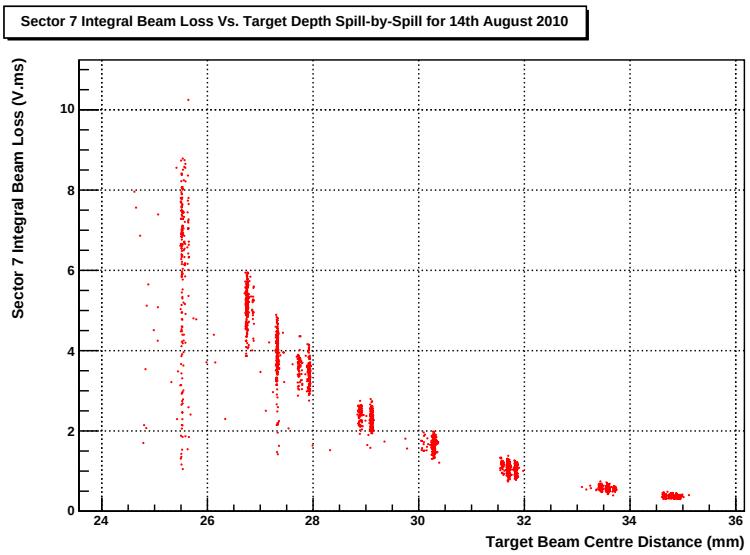

Figure 4.9: Spill-by-spill plot of beam loss as a function of MICE target BCD for the 14th August 2010 study. The large spread in beam loss at the highest losses is enhanced in this study by a problem with ISIS ion source which lead to large fluctuations in the beam current, when the largest beam loss data points were being taken.

4.2 Sector 7 and Sector 8 Beam Loss

In this study the main figure of merit used to evaluate the effect of the MICE target on ISIS beam loss are the integrated losses over a whole ISIS cycle for the summed signal of the four monitors in sector 7. The same integrated losses in sector 8 are also a good gauge. The relationship between the two, for all studies, is shown in Fig 4.10. Run-by-run plots are used as this more clearly illustrates the shape of the curve. A linear relationship is clearly visible to a high degree of accuracy. The absolute magnitudes are also similar, though slightly higher in sector 8. For the August study, a slight deviation from linearity may appear at the very highest losses, but again this may be related to the ion source problems at the highest data points.

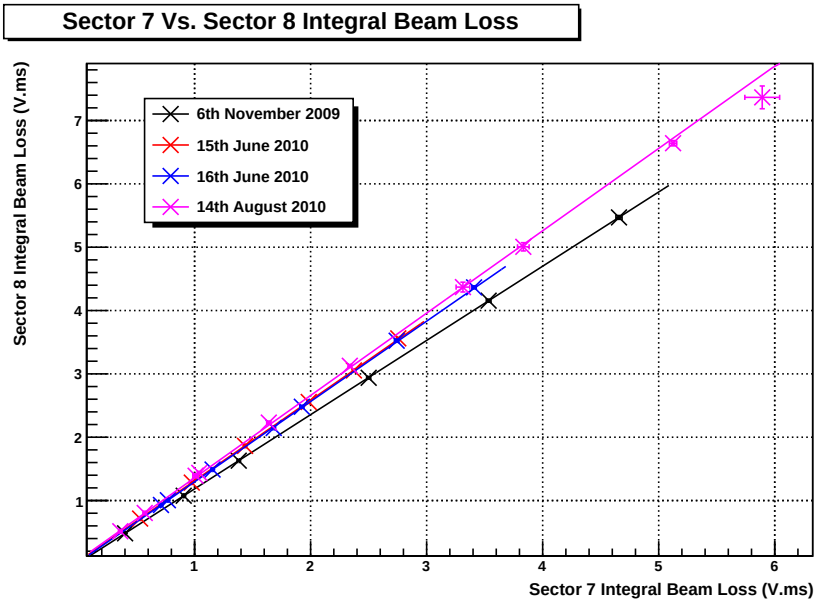

Figure 4.10: Sector 8 integrated beam loss as a function of sector 7 integrated beam loss for all four studies. Linear fits are also shown.

4.3 Beam Loss and Total Particle Rates

As described in Section 2.8, the total particle rate in the MICE beamline can be measured using scaler hits produced by the GVA1, BPM1, BPM2, TOF0 and TOF1 detectors, while scaler hits in the Luminosity Monitor (where available) give a measure of the number of particles produced or scattered by the target. Results for these scaler rates shall now be presented for each individual study.

4.3.1 6th November 2009 Study

A graph showing the run averaged rate data as a function of induced sector 7 beam loss for the November 09 study is shown in Fig. 4.11. For all detectors and across the whole beam loss range up to the highest losses achieved at 4.7 V.ms, a linear relationship can be clearly observed between rate and beam loss.

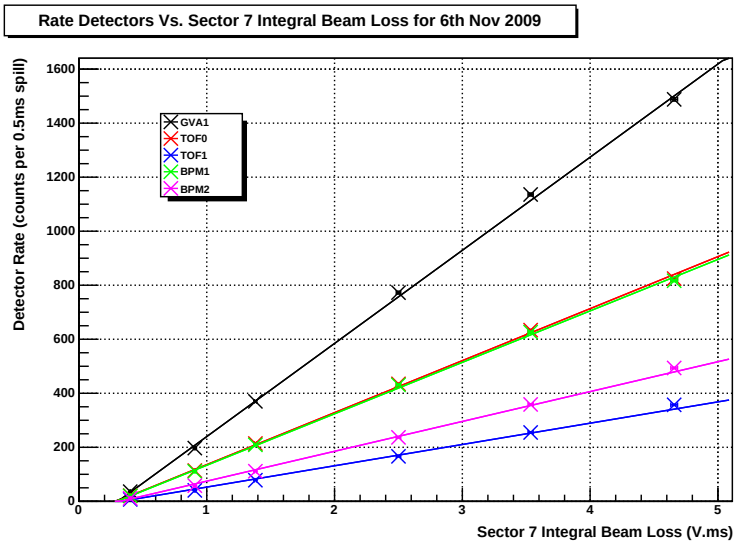

Figure 4.11: Total particle rate as recorded in GVA1, BPM1, TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 6th November 2009 study. Linear fits are also shown giving good agreement with data.

4.3.2 15th June 2010 Study

Plots of the beamline particle rate as a function of induced beam loss for the 15th June 2010 study are shown in spill-by-spill form in Fig. 4.12, and run averaged form in Fig. 4.13. As GVA1 dominates the scale the same plots are also shown in Figs. 4.14 and 4.15 with GVA1 removed. Again linearity can be seen to hold across the entire beam loss range.

Luminosity Monitor rate as a function of induced beam loss is shown in Figs. 4.16 and 4.17, again showing a linear relationship. Lastly beamline particle rates as a function of Luminosity Monitor rate are shown in Figs. 4.18 - 4.21, also being linear.

4.3.3 16th June 2010 Study

Plots of the beamline particle rate as a function of induced beam loss for the 16th June 2010 study are shown in spill-by-spill form in Fig. 4.22, and in run averaged form in Fig. 4.23. Again linearity can be seen to hold across the entire beam loss range, though with slightly more variation than was present for the 15th June study.

Luminosity Monitor rate as a function of induced beam loss is shown in Figs. 4.24 and 4.25. Again an approximately linear relationship is visble, though curiously the spill-by-spill plot hints that the slopes may be different for different runs, possibly being the result of saturation begining to affect the Luminosity Monitor at the higher losses (an effect which becomes obvious in the August data). Lastly beamline particle rates as a function of Luminosity Monitor rate are shown in Figs. 4.26 - 4.27, again being approximately linear.

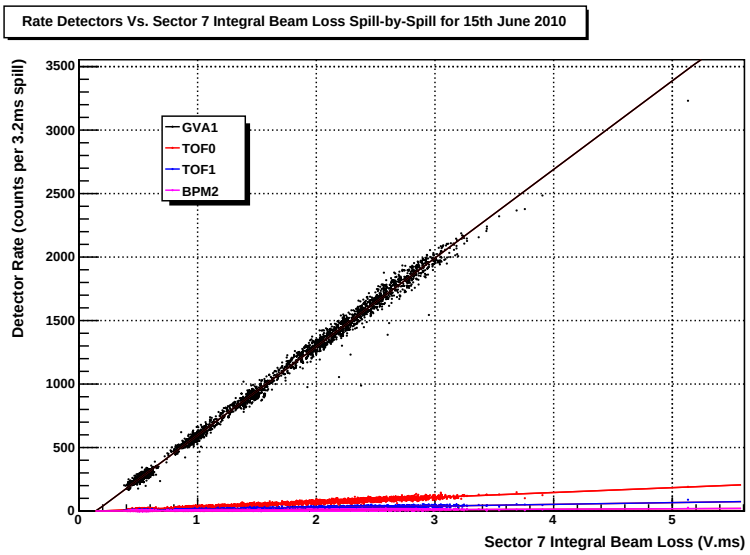

Figure 4.12: Spill-by-spill plot of total particle rate as recorded in GVA1, TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 15th June 2010 study. Linear fits are also shown. GVA1 dominates the scale, see Fig. 4.14 below for just the lower rate detectors.

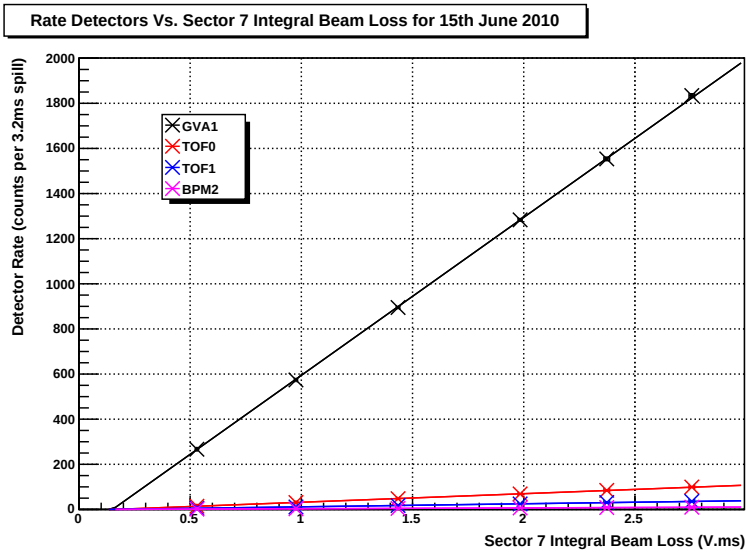

Figure 4.13: Total particle rate as recorded in GVA1, TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 15th June 2010 study. Linear fits are also shown. GVA1 dominates the scale, see Fig. 4.15 below for just the lower rate detectors.

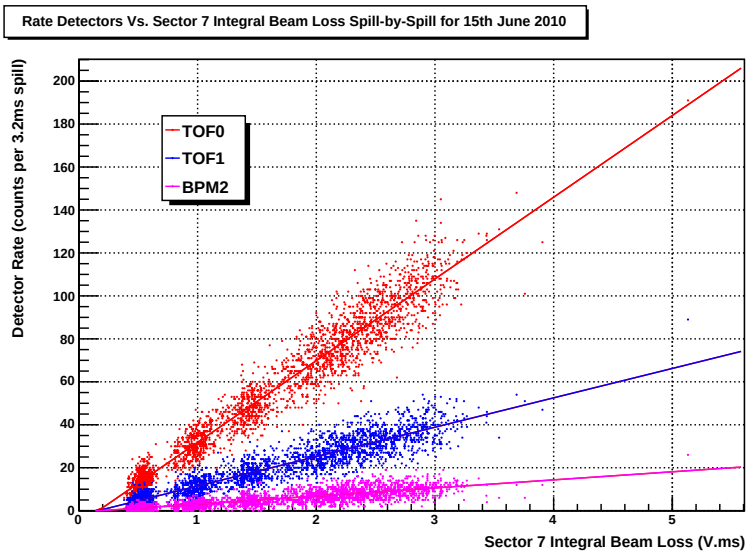

Figure 4.14: Spill-by-spill plot of total particle rate as recorded in TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 15th June 2010 study. Linear fits are also shown.

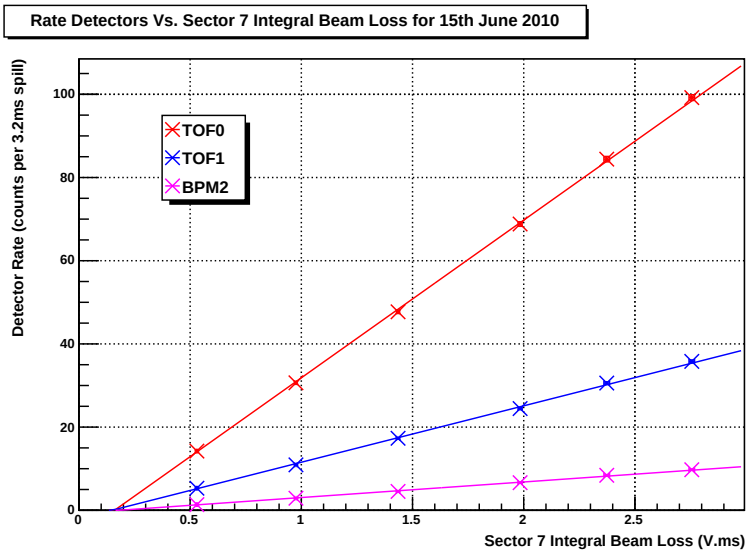

Figure 4.15: Total particle rate as recorded in TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 15th June 2010 study. Linear fits are also shown. Previously published in [73].

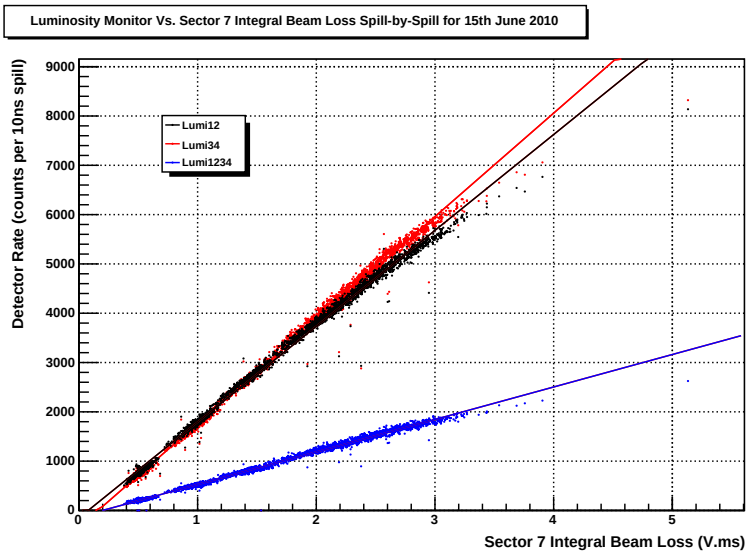

Figure 4.16: Spill-by-spill plot of Luminosity Monitor rate as a function of beam loss for the 15th June 2010 study. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs.

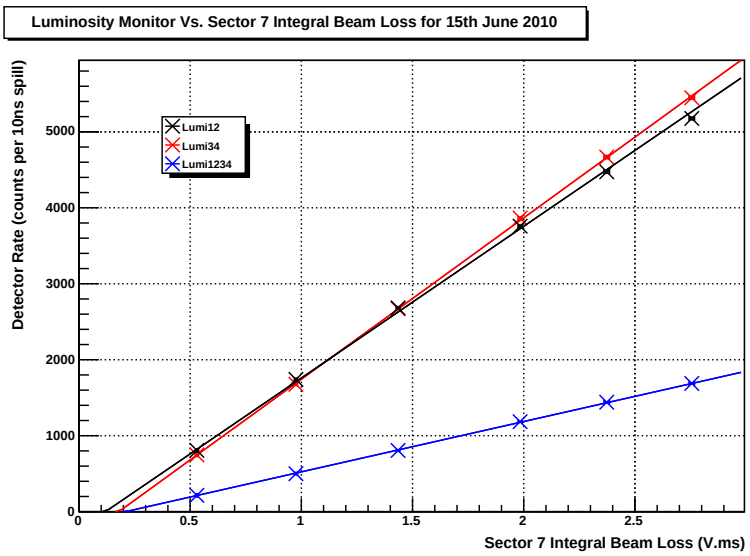

Figure 4.17: Luminosity Monitor rate as a function of sector 7 beam loss for the 15th June 2010 study. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs.

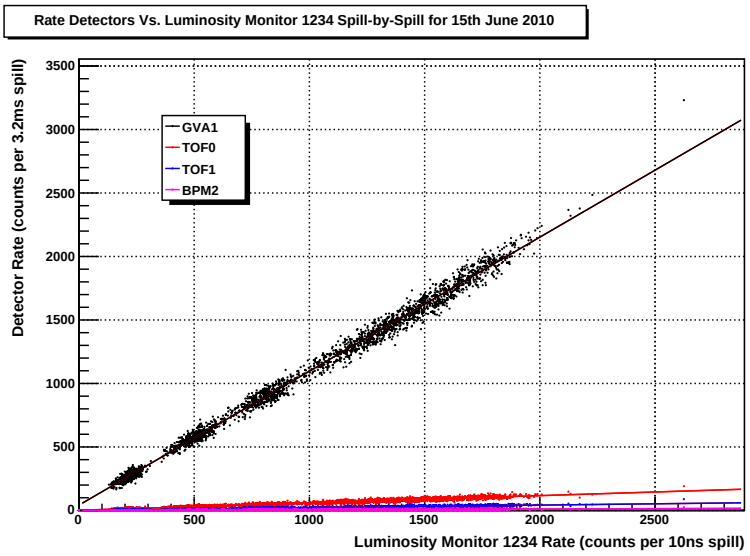

Figure 4.18: Spill-by-spill plot of total particle rate as recorded in GVA1, TOF0, BPM2 and TOF1, as a function of Lumi1234 rate, for the 15th June 2010 study. Linear fits are also shown. GVA1 dominates the scale, see Fig. 4.20 below for just the lower rate detectors.

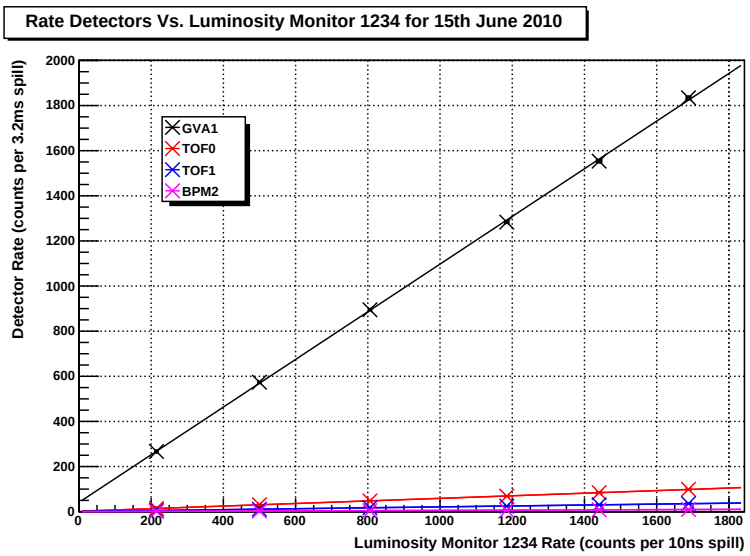

Figure 4.19: Total particle rate as recorded in GVA1, TOF0, BPM2 and TOF1, as a function of Lumi1234 rate, for the 15th June 2010 study. Linear fits are also shown. GVA1 dominates the scale, see Fig. 4.21 below for just the lower rate detectors.

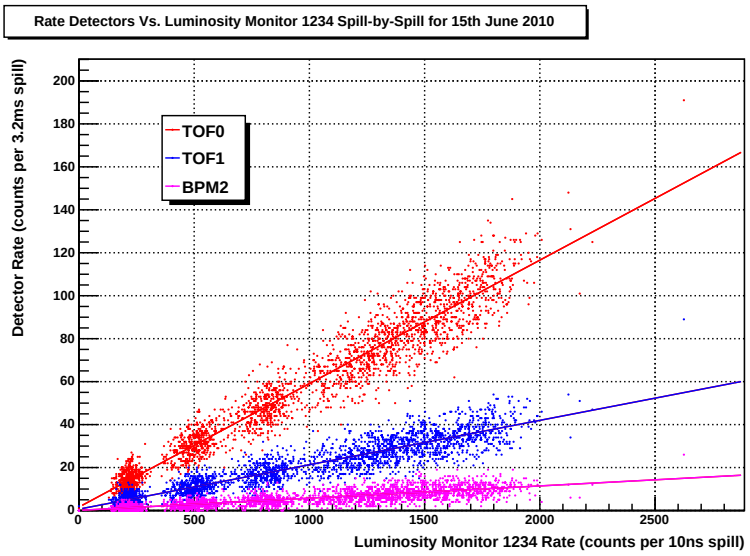

Figure 4.20: Spill-by-spill plot of total particle rate as recorded in TOF0, BPM2 and TOF1, as a function of Lumi1234 rate, for the 15th June 2010 study. Linear fits are also shown.

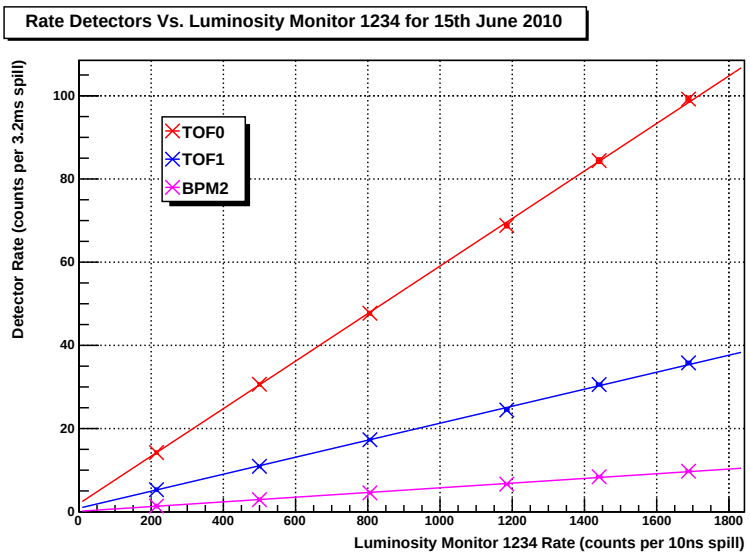

Figure 4.21: Total particle rate as recorded in TOF0, BPM2 and TOF1, as a function of Lumi1234 rate, for the 15th June 2010 study. Linear fits are also shown.

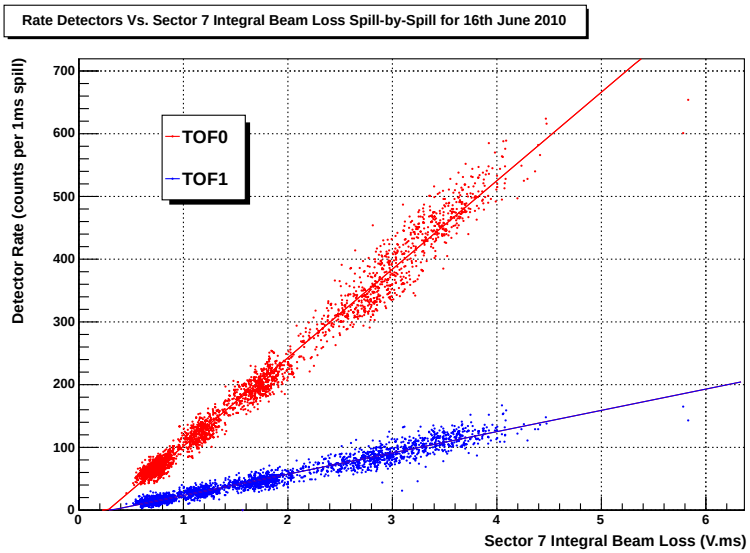

Figure 4.22: Spill-by-spill plot of total particle rate as recorded in TOF0 and TOF1, as a function of sector 7 beam loss, for the 16th June 2010 study. Linear fits are also shown.

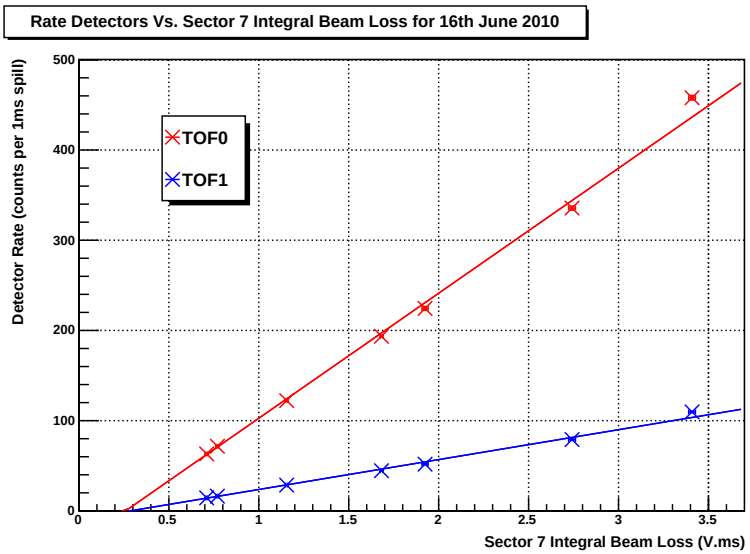

Figure 4.23: Total particle rate as recorded in TOF0 and TOF1, as a function of sector 7 beam loss, for the 16th June 2010 study. Linear fits are also shown. Previously published in [73].

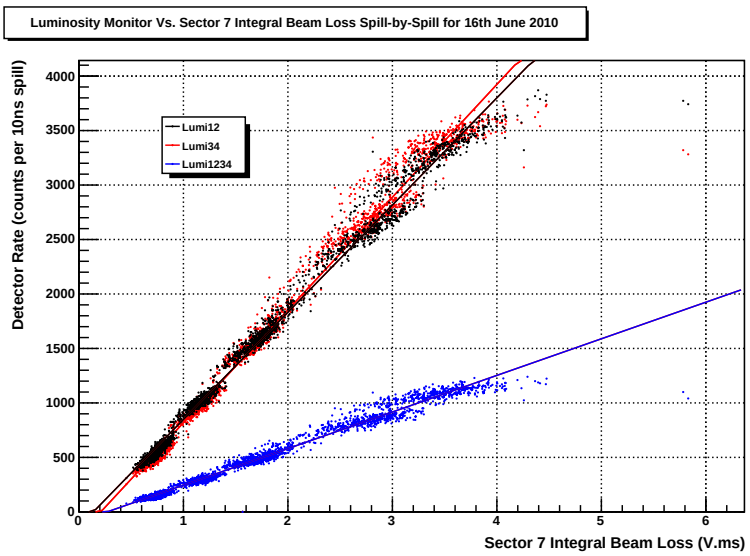

Figure 4.24: Spill-by-spill plot of Luminosity Monitor rate as a function of beam loss for the 16th June 2010. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs.

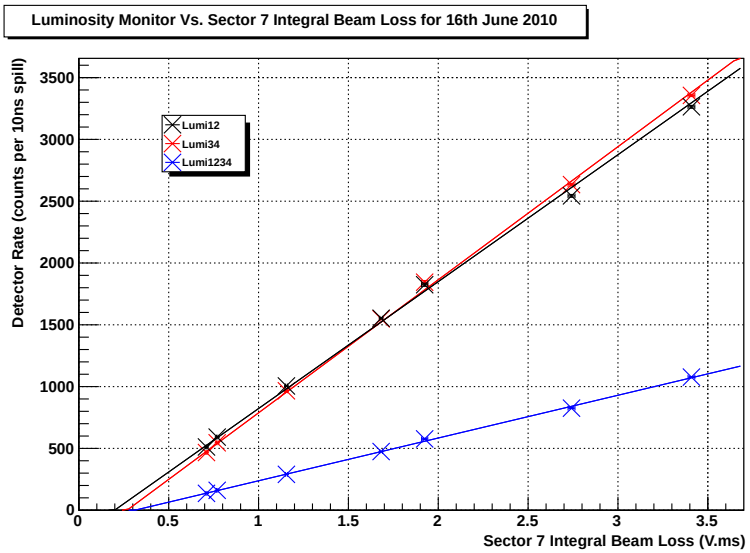

Figure 4.25: Luminosity Monitor rate as a function of sector 7 beam loss for the 16th June 2010 study. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs.

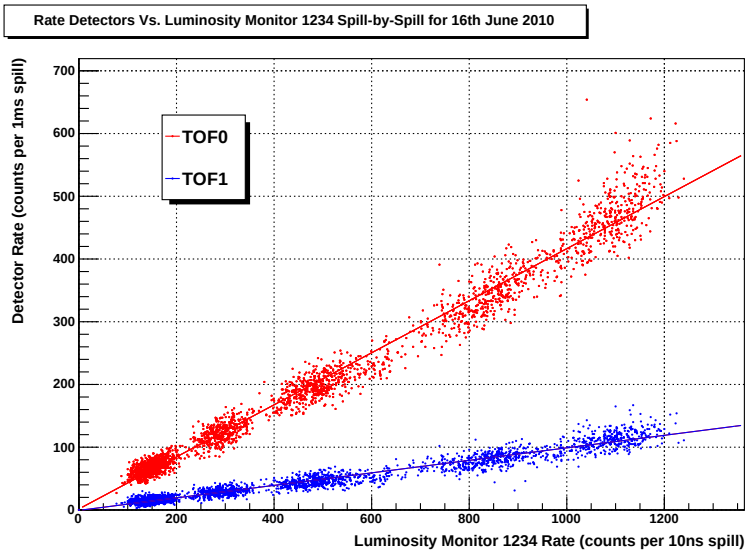

Figure 4.26: Spill-by-spill plot of total particle rate as recorded in TOF0 and TOF1, as a function of Lumi1234 rate, for the 16th June 2010 study. Linear fits are also shown.

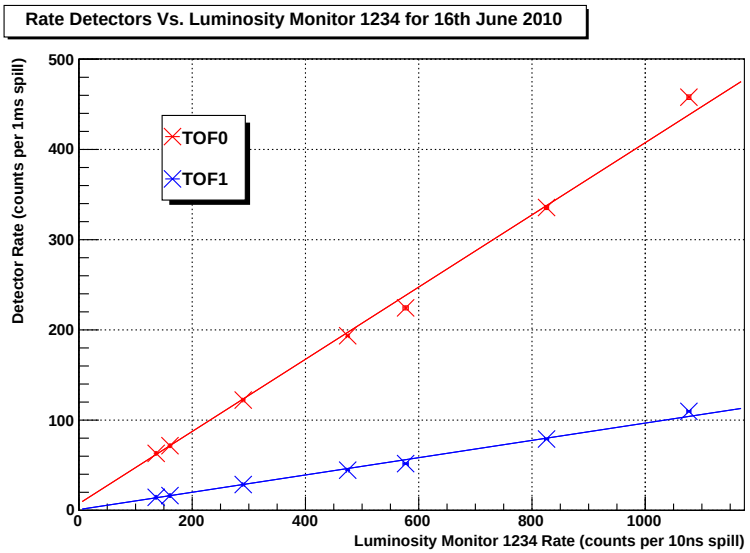

Figure 4.27: Total particle rate as recorded in TOF0 and TOF1, as a function of Lumi1234 rate, for the 16th June 2010 study. Linear fits are also shown.

4.3.4 14th August 2010 Study

Plots of beamline detector rate as a function of beam loss for the 14th August 2010 are shown spill-by-spill in Fig. 4.28 and in run averaged form in Fig. 4.29. GVA1 can be seen to reach saturation point at $\sim$ 3 - 4 V.ms. As GVA1 also dominates the scale, the plots are reproduced in Figs. 4.30 - 4.31 with just TOF1, TOF1 and BPM2 shown.

The highest beam loss point in the August 2010 study possesses a large spread in the beam loss distribution, shown in Fig. 4.32. In particular a large tail down to lower beam losses, again possibly a result of the ion source problems during the last few runs of this study, causes the average value of the last run averaged data point to lower considerably. Cuts applied to the beam loss distribution of 5.7 V.ms $<$ Beam Loss $<$ 9.0 V.ms help correct this by removing the tail. The resultant run averaged plot of particle rate as a function of beam loss is shown in Fig. 4.33. The beam loss range can be seen to be extended by the cut to higher losses by $\sim$ 1 V.ms, while actually leading to a reduction in the size of the error bar.

In both the cut and uncut plots a non-linear increase of rate with beam loss can be seen at lower target BCD values, tending back towards linearity as the target enters further into the beam. This is illustrated by the linear fits shown, which exclude the first four data points and then match data well for the remaining higher loss points. Fig. 4.34 shows a zoomed in view of the low beam loss region, from which it be seen that the particle rate is higher at low beam loss than that expected from the fits to the higher points, pointing to an excess in particle rate or conversely a lack of beam loss over the expectation.

The cause of the non-linearity may be elucidated by looking at the variation of beamline particle rates with Luminosity Monitor rate. Looking at

the variation of Luminosity Monitor rate with beam loss however, as shown in Figs. 4.35 and 4.36, it can be seen that like GVA1, the Luminosity Monitor saturates. While this complicates matters, as the saturation only begins to become pronounced at $\sim$ 4 - 5 V.ms, and the non-linearity in particle rate occurs only in the 0 - 1.5 V.ms region, the Luminosity Monitor may still be used if appropriate cuts are made to remove the high beam loss runs.

Fig. 4.37 shows the variation of the beamline particle rate with Luminosity Monitor rate, excluding the two highest beam loss points where the Luminosity Monitor is clearly saturated. Now it can be seen that linearity holds to a reasonable approximation across the low beam loss range, as expected from the previous studies. This indicates that the low beam loss non-linear behaviour may have its cause in some unexpected behaviour in the beam loss (or in how it is measured), rather than from any effects present in the MICE beamline. Fig. 4.38 shows the induced beam loss profile, target dip parabola and ISIS beam intensity for a selection of runs from this study, indicating how beam loss profile changes with increasing target BCD. While the beam loss profile is clearly broader at lower beam losses, there is still no obvious cause for the observed non-linear behaviour. This issue is discussed further in Section 6.3, which gives a recommendation for a further study.

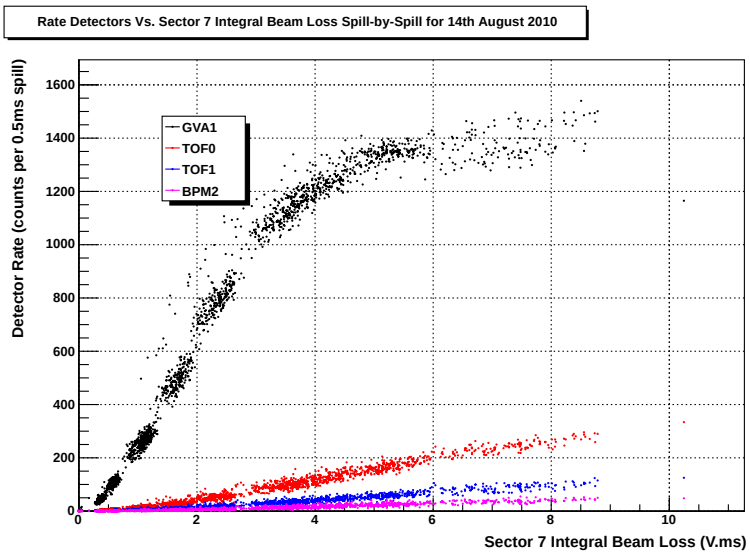

Figure 4.28: Spill-by-spill plot of total particle rate as recorded in GVA1, TOF0, BPM2 and TOF1, as a function of sector 7 beam loss, for the 14th August 2010 study. Saturation in GVA1 is clearly visible above $\sim$ 3 - 4V.ms.

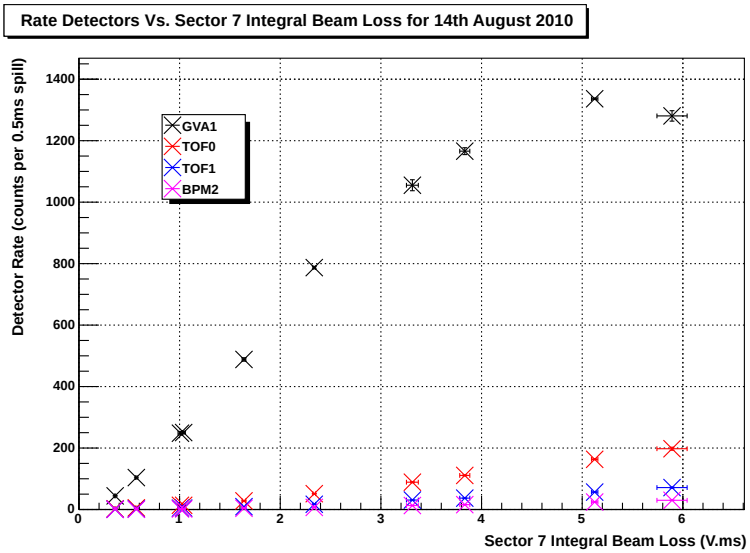

Figure 4.29: Total particle rate as recorded in GVA1, TOF0, TOF1 and BPM2, as a function of sector 7 beam loss, for the 14th August 2010 study. Saturation in GVA1 is clearly visible above $\sim$ 3 - 4V.ms.

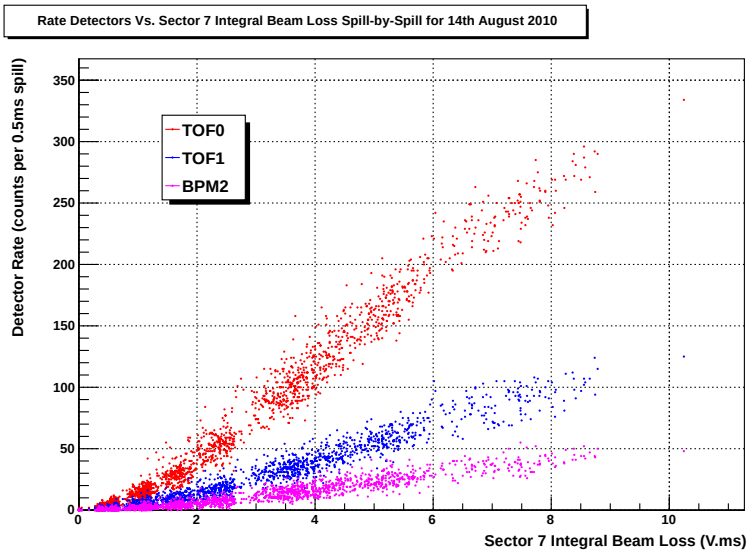

Figure 4.30: Spill-by-spill plot of total particle rate as recorded in TOF0, TOF1 and BPM2, as a function of sector 7 beam loss, for the 14th August 2010 study.

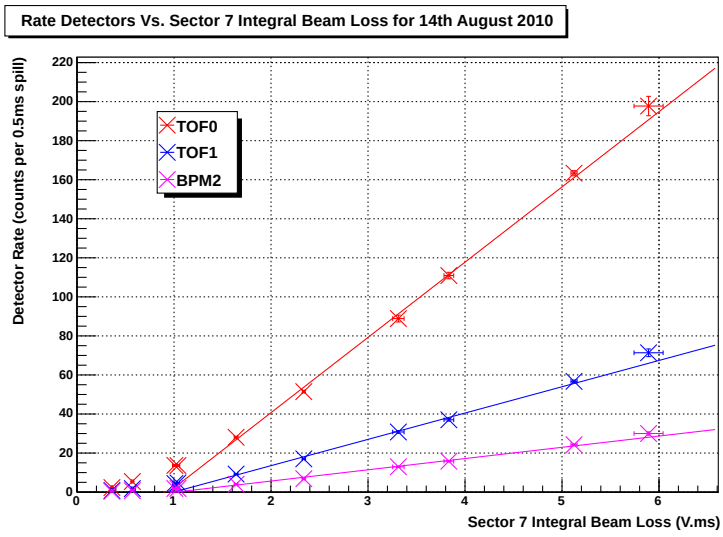

Figure 4.31: Total particle rate as recorded in TOF0, TOF1 and BPM2, as a function of sector 7 beam loss, for the 14th August 2010 study. Linear fits, fitted over the range 1.5 V.ms to 8 V.ms and then extrapolated back to 0 V.ms, are also shown. A significantly more pronounced deviation from linearity than in the other studies is visible at low beam losses, tending towards linearity again at higher beam losses.

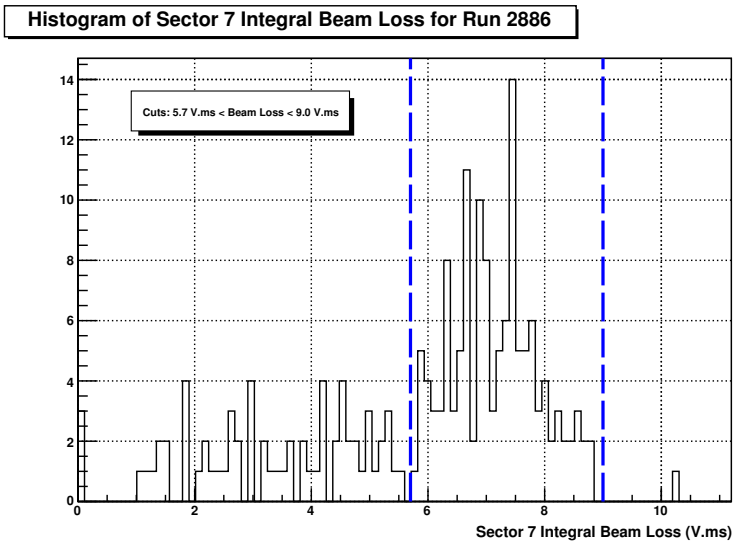

Figure 4.32: Histogram showing the beam loss distribtution for run 2886. Cuts placed at 5.7 V.ms < Beam Loss < 9.0 V.ms isolate the main peak and increase the value of the beam loss point used in the averaged data graph (shown in Fig. 4.33).

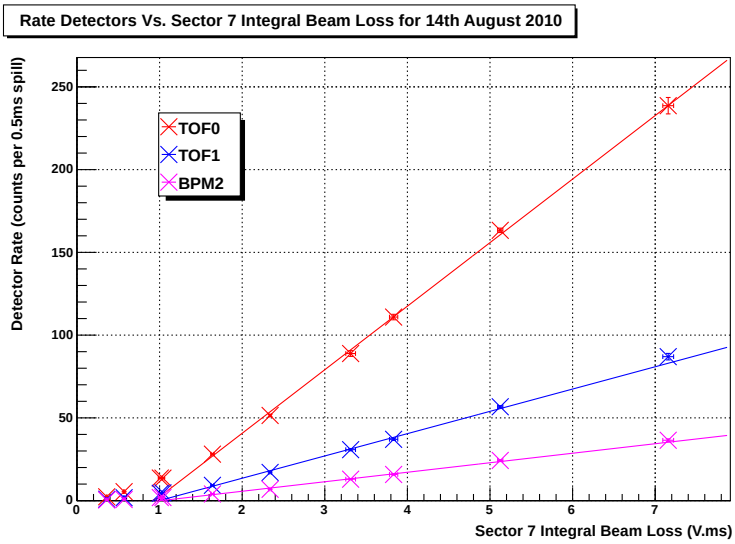

Figure 4.33: Total particle rate as recorded in TOF0, TOF1 and BPM2, as a function of sector 7 beam loss, for the 14th August 2010 study. Cuts of 5.7 V.ms < Beam Loss < 9.0 V.ms are placed on the distribution of the highest beam loss point (run 2886). Linear fits, fitted over the range 1.5 V.ms to 8 V.ms and then extrapolated back to 0 V.ms, are also shown.

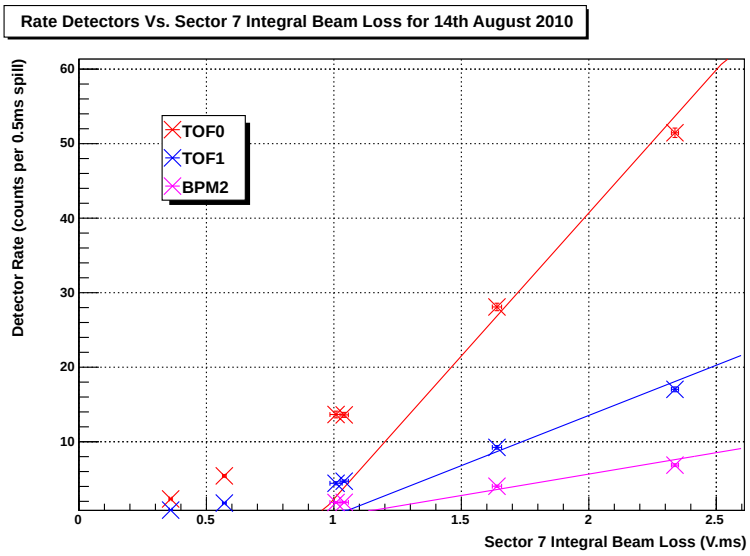

Figure 4.34: Total particle rate as recorded in TOF0, TOF1 and BPM2, as a function of sector 7 beam loss, for the 14th August 2010 study, zoomed to the low beam loss region.

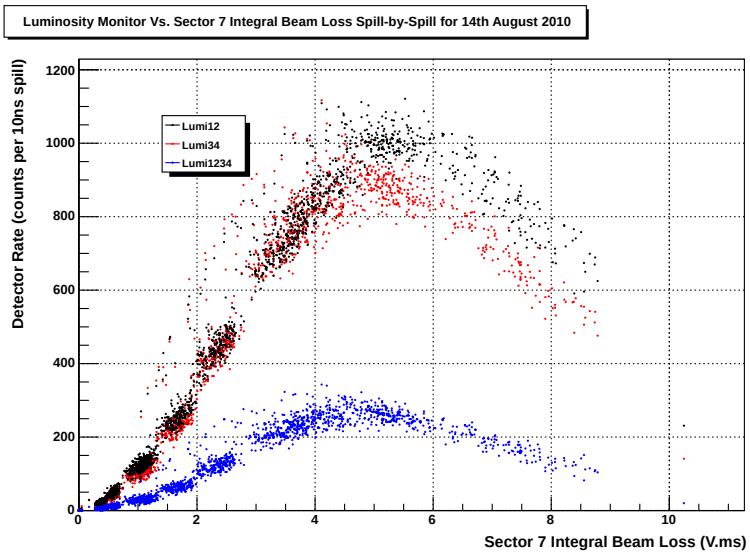

Figure 4.35: Spill-by-spill plot of Luminosity Monitor rate as a function of beam loss for the 14th August 2010. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs. Saturation of the monitor is visible at $\sim$ 4 - 5 V.ms.

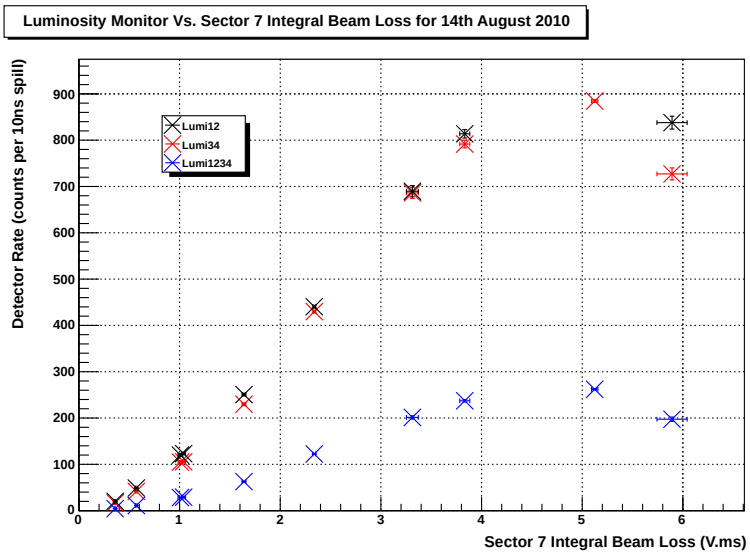

Figure 4.36: Luminosity Monitor rate as a function of sector 7 beam loss for the 14th August 2010 study. Lumi12 represents coincident hits in PMT1 and 2 of the monitor, Lumi34 coincident hits between PMT3 and 4, Lumi1234 coincident hits between all 4 PMTs. Saturation of the monitor is visible at $\sim$ 4 - 5 V.ms.

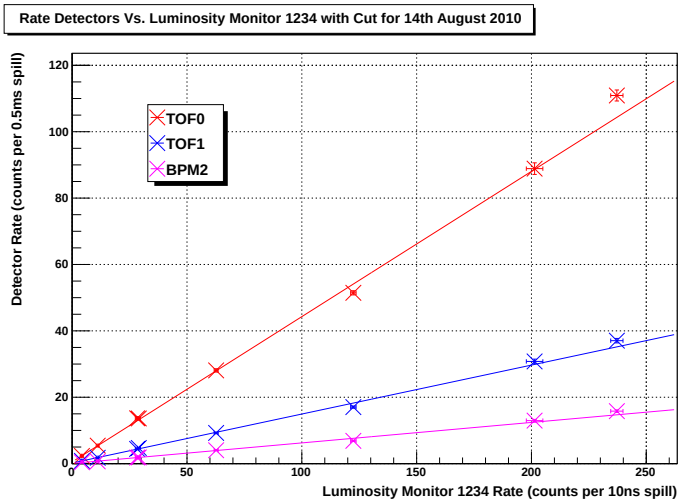

Figure 4.37: Beamline particle rate as a function of Lumi1234 rate for the 14th August 2010 study. The two highest beam loss points of the study are excluded due to the saturation of the Luminosity Monitor for those rates. Linear fits are also shown.

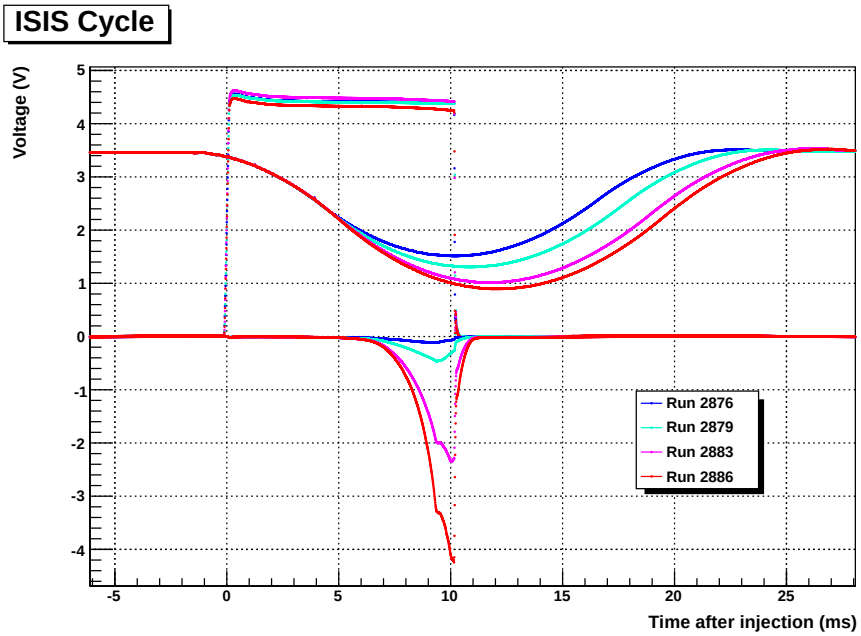

Figure 4.38: The ISIS cycle for a selection runs of the 14th Aug 2010 study. Increasing run number indicates increasing target BCD and so increasing induced beam loss. The "top-hat" like profile represents the ISIS beam intensity, the smooth curve represents the MICE target dip parabola, and the negative going profile the sector 7 beam loss.

4.4 Beam Loss and Particle Rates per Species

4.4.1 6th November 2009 Study

An example TOF distribution, between TOF0 and TOF1, for run 1231 from the November 2009 study is shown in Fig. 4.39. It was created using the TofMonitor G4MICE application.[1] For each run in this study, cuts were applied around the peaks, and the total number of entires within each peak counted to give the relative particle rates. The cuts used are:

- All real particles: 25 ns $< dt <$ 32 ns;
- Pions: 29 ns $< dt <$ 32 ns;
- Muons: 27.5 ns $< dt <$ 29 ns and
- Electrons: 25 ns $< dt <$ 27.5 ns.

The ratio of each beamline species using the above cuts is shown in Table 4.1. It can be seen that the most prevalent species is pions, followed by muons and lastly electrons, as could be expected for a pion optics beamline.

The reconstructed TOF track data is extracted by G4MICE per run without a spill structure, in contrast to the scaler hits data which is recorded for each individual spill. In order to produce plots of rate as a function of beam loss per spill, as done for scaler hits in Section 4.3, the approximate number of target pulses per run is used (as discussed in Section 3.2.3). This then provides an average value of the number of muon tracks per spill for each run in each study.

A plot showing the total number of particles per species per spill as a function of beam loss is given in Fig. 4.48. After ∼2 V.ms it can be

[1] Created by Mark Rayner (Oxford University). Now incorporated into the DataQualityCheck application.

Run	Electrons	Muons	Pions	Total Tracks
1231	1.00	1.51	2.34	4.85
1232	1.00	1.52	2.44	4.95
1233	1.00	1.51	2.47	4.97
1234	1.00	1.56	2.45	5.01
1235	1.00	1.62	2.44	5.06
1236	1.00	1.32	2.15	4.48
Average	1.00	1.51	2.38	4.89

Table 4.1: Ratio of muon, pion and total physical tracks per pulse normalised to electron tracks per pulse.

observed that the rate of each species stops increasing with higher beam loss, in contrast to what would be expected from the scaler data. This is most likely a consequence of dead-time in the DAQ and decreasing software reconstruction efficiency with increasing rate. The former is illustrated in Figs. 4.43 and 4.44, which show the number of accepted particle triggers as a function of the number of particle trigger requests. As descibed in Section 2.8.8 a request is denied if it occurs in the DAQ dead-time window, that is, if it is sufficiently close to a previously accepted trigger (within approximately 1.28 μs). A loss in the number of triggers of $\sim$50% can be observed at the highest losses. While this saturation is a highly undesirable effect, it only becomes a major concern at higher particle rates, as illustrated in the studies below.

4.4.2 June 2010 Studies

Example time-of-flight distributions for the June 2010 studies, created using the G4MICE TofTree application, are shown in Figs. 4.40 and 4.41, the highest beam loss data points of the 15th and 16th respectively. A central peak is observable, believed to consist primarily of muons, together with a small electron peak on the left, according with expectation for $\pi \rightarrow \mu$ optics.

Conditions	Scaler Hits	All Tracks	Muon Tracks	Muon / All	Efficiency
15th Deriv. 3.2ms	13.56	5.98	5.85	0.978	0.43
15th 1.3V.ms 3.2ms	15.62	6.95	6.82	0.981	0.44
15th 2V.ms 3.2ms	25.12	11.14	10.91	0.980	0.43
16th Deriv. 1ms	33.14	16.43	16.35	0.995	0.49
16th 1.3V.ms 1ms	33.64	19.72	19.66	0.997	0.58
16th 2V.ms 1ms	56.84	31.22	31.10	0.996	0.55

Table 4.2: Summary of results for the June 2010 studies. Scaler Hits are counted by TOF1. All Tracks and Muon Tracks represent TOF tracks reconstructed by software between TOF0 and TOF1, in the range 25 ns $< dt <$ 32 ns. Muon Tracks are distinguished by the cut $dt = 26.2$ ns. Deriv. refers to the derivative of the first order polynomial fitted to the data using ROOT. The rates and tracks at 1.3 V.ms and 2 V.ms are also evaluated using these fits. Muon / All represents the ratio of All Tracks to Muon Tracks. Efficiency refers to fraction of Scaler Hits that are converted to Muon Tracks. The expected particle rate in optimum conditions may be slightly higher than those observed here, due to the lack of Q3 in the first quadrupole triplet during these studies. Previously published in [73].

The exact content of the central peak is the subject of ongoing work. The successful operation of the Cherenkov detectors (not yet available) would aid pion / muon discrimination, and input from simulation (either G4BeamLine or G4MICE) could also help to determine the beam content with greater precision. Simulation would also be necessary to determine what fraction of these particles are "good" muons, that is muons which traverse the whole cooling channel and fall within the acceptance of TOF2 (see the discussion on future work in Section 6.3).

If the approximation is taken that the central peak consists purely of muons, then the muon track rate between TOF0 and TOF1 per unit beam loss can be calculated for these studies. The electron peak visible in the TOF spectrum is removed by means of a cut at $dt = 26.2$ ns, as is illustrated in Figs. 4.40 and 4.41. The TOF track data is again extracted for a whole run

and coverted to per spill values, as for the November 09 study.

The results of the analysis are shown as plots of muon track rate as a function of beam loss in Figs 4.50 and 4.51 for the 15th and 16th June studies respectively. A linear relationship can again be seen between muon track rate and the induced beam loss. This implies only a low incidence of DAQ dead-time, which is confirmed by looking again at the ratio of trigger requests to accepted triggers, shown in Fig. 4.43 for all studies and in Figs. 4.45 and 4.46 for the 15th and 16th June respectively.

Linear fits to the muon track rate against beam loss plots are provided, allowing the muon rate to be evaluated for arbitrary values of beam loss. Typical MICE operating values of 1.3 V.ms (which corresponds to ∼2 V.ms on the ISIS scale) and 2 V.ms are used, together with the derivative of the fit function. The results are shown in Table 4.2, together with the absolute scaler rates from Section 4.3 (it should be remembered that as the particle rate across the spill is non-linear the figures for the 15th and 16th cannot be compared directly, again see Section 6.3). This then provides an estimate of the number of muons currently available at TOF1 which may be used to demonstrate cooling. The implications of these rates for MICE are discussed in Section 6.2.

4.4.3 14th August 2010 Study

An example time-of-flight distribution for the August 2010 study is shown in Fig. 4.42. It can be seen to have a similar shape to that observed for the June 2010 studies, which is to be expected as both studies have optics set for $\pi \rightarrow \mu$ transport (with the small difference that Q3 was available for the August study).

Due to the small fraction of the spill captured in this study (as discussed

in Section 4.3) it is unwise to use the muon rates observed in this study as a guide to those expected in normal MICE running at different beam losses, as was done for the June studies in Section 4.4.2. Nevertheless, the total reconstructed TOF rate as a function of beam loss is shown in Fig. 4.52 and contains some interesting features. At the lowest beam losses a small non-linear rate increase can be seen, probably resulting from the non-linearity seen in the same region in the scaler hits data (see Section 4.3.4). The curve then becomes linear before again becoming non-linear at the highest beam losses, but now towards decreasing particle rate gains. This feature is not present in the scaler data which, taken with its appearance at high losses, implies that it arises either from DAQ dead-time or software reconstruction inefficiency, in a similar manner to the November track rates discussed in Section 4.4.1. The prescence of significant dead-time is confirmed by again looking at the accepted trigger rate as a function of the trigger request rate, shown in Fig. 4.47.

It is interesting to note that the high beam loss non-linearity is not present in the data from the 16th June, which achieved higher total track rates. This is however still consistent with the above explanation as, while the total rate per spill is not as high as that for the 16th June due to the small gate and early spill present in the August study, the instantaneous rate per unit time is extremely high due to the large beam losses. The issue of reconstruction efficiency when moving from scaler hits to TOF tracks is one of considerable importance and is discussed further in Section 4.5 below.

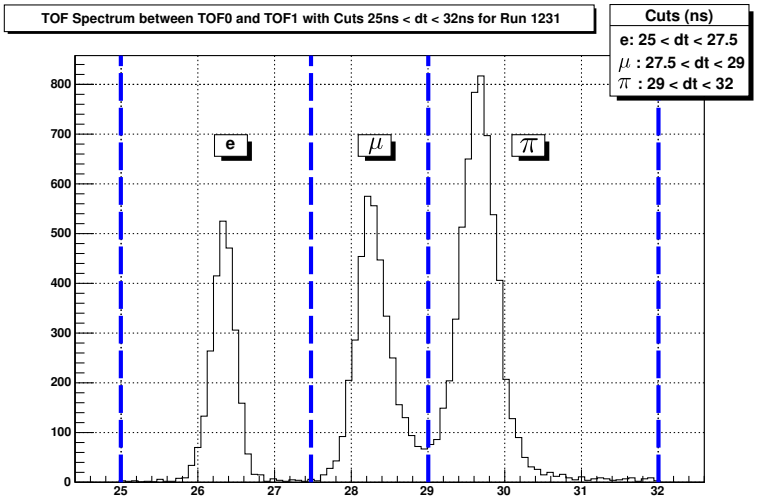

Figure 4.39: TOF Spectrum between TOF0 and TOF1 for run 1231 from 6th November 2009 study. Cuts applied to the reconstructed tracks are 25 ns $< dt <$ 32 ns. From the left the peaks represent positrons, then muons, then pions.

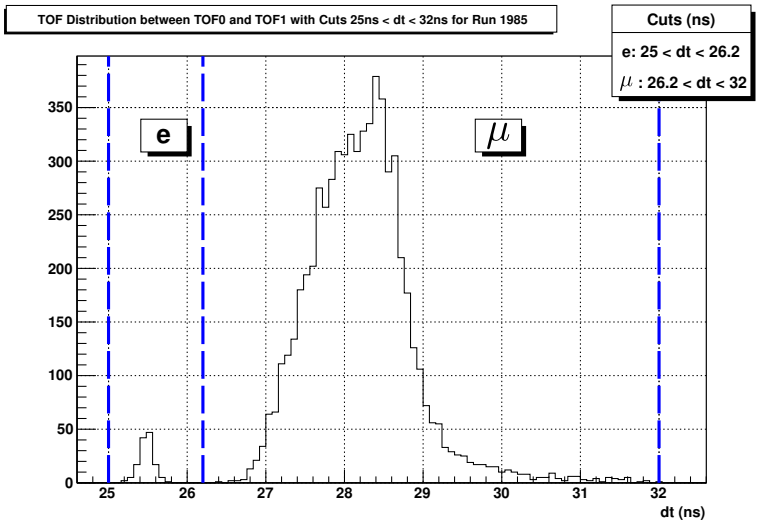

Figure 4.40: TOF Spectrum between TOF0 and TOF1 for run 1985 from the 15th June 2010 study. Cuts applied to the reconstructed tracks are 25 ns $< dt <$ 32 ns. A small electron peak is visible on the left, while the centre peak is predominately muons. A further cut of 26.2 ns $< dt$ is used to select the muons. Previously published in [73].

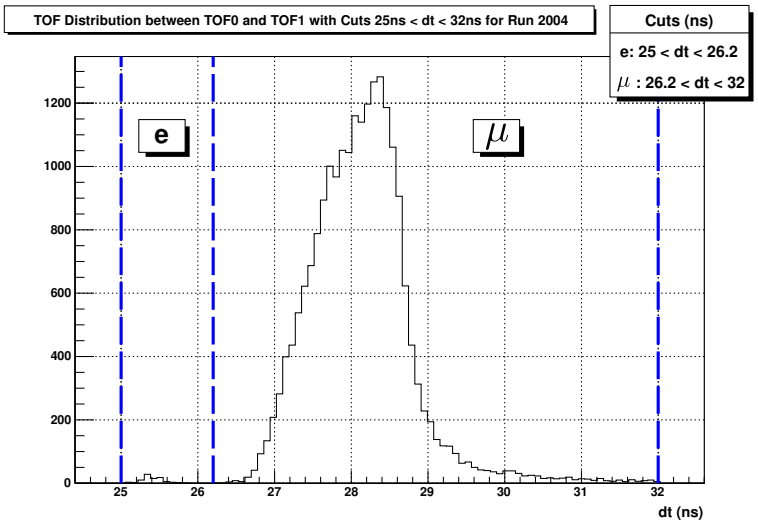

Figure 4.41: TOF Spectrum between TOF0 and TOF1 for run 2004 from the 16th June 2010 study. Cuts applied to the reconstructed tracks are 25 ns $< dt <$ 32 ns. The small positron peak on the left is now barely visible, the centre peak is predominately muons. A further cut of 26.2 ns $< dt$ is used to select the muons.

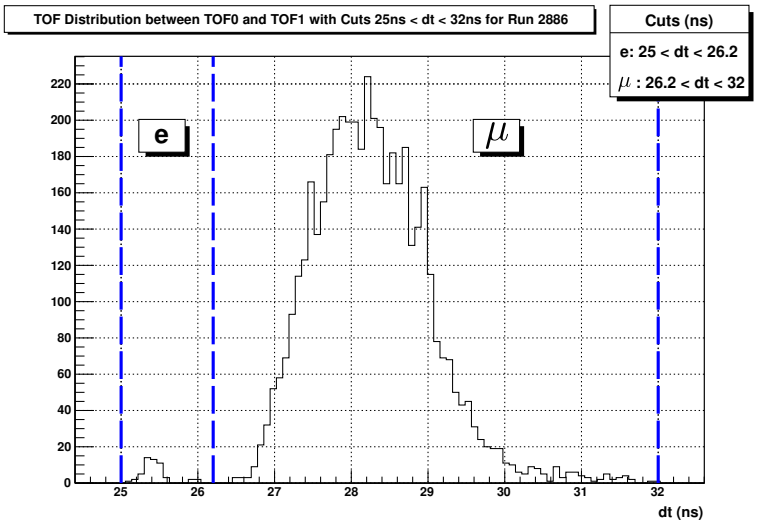

Figure 4.42: TOF Spectrum between TOF0 and TOF1 for run 2886 from the 14th August 2010 study. Cuts applied to the reconstructed tracks are 25 ns $< dt <$ 32 ns to give phsyical tracks. A small positron peak is visible on the left, while the centre peak is predominately muons. A further cut of 26.2 ns $< dt$ is used to select the muons.

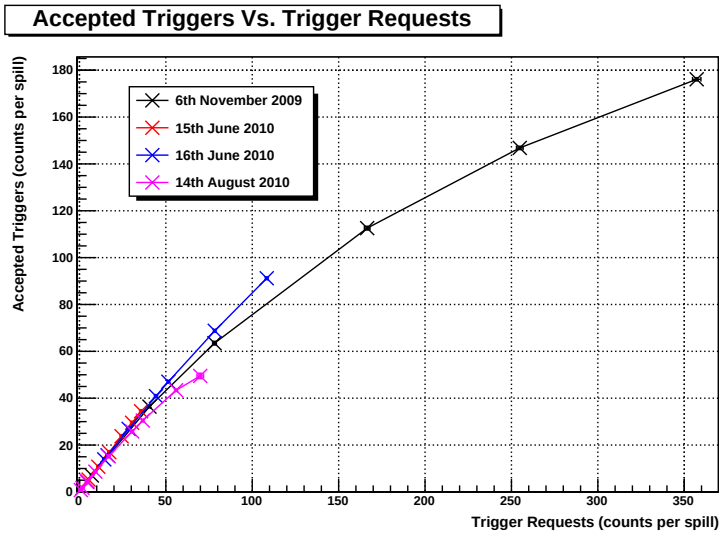

Figure 4.43: Accepted particle triggers as a function of particle trigger requests for all four studies. The scale is dominated by the higher trigger rate of the 2009 study.

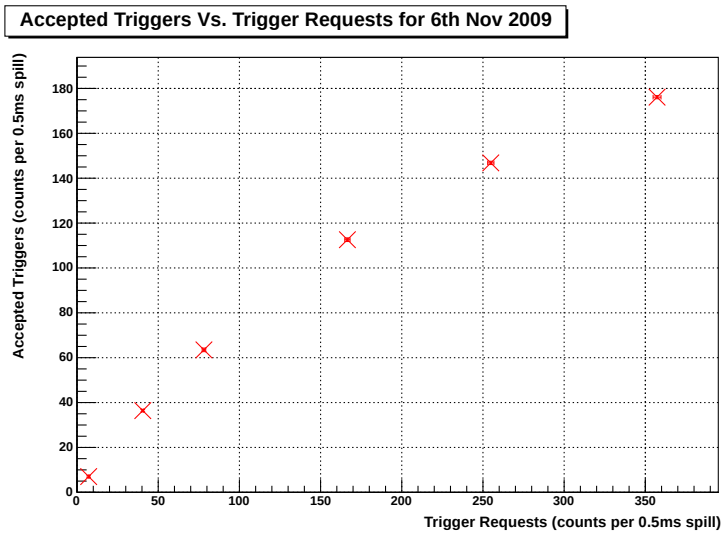

Figure 4.44: Accepted particle triggers as a function of particle trigger requests for the 6th November 2009 study. A curve is clearly visible showing the increasing effect of detector dead-time as rates increase.

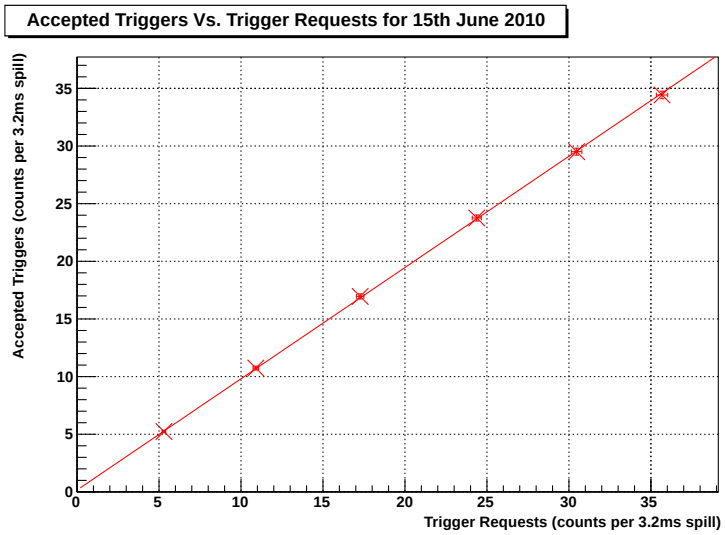

Figure 4.45: Accepted particle triggers as a function of particle trigger requests for the 15th June 2010 study. Here the curve is linear showing only a negligible effect from dead-time.

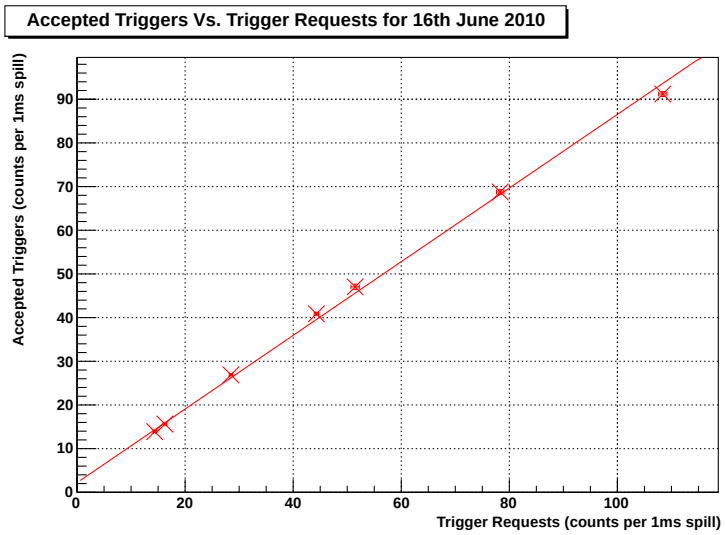

Figure 4.46: Accepted particle triggers as a function of particle trigger requests for the 16th June 2010 study. Here the curve is approximately linear showing only a negligible effect from dead-time.

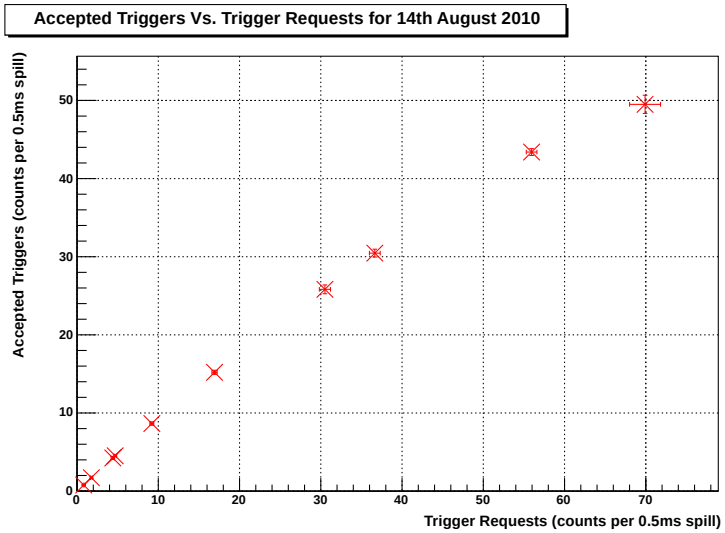

Figure 4.47: Accepted particle triggers as a function of particle trigger requests for the 14th August 2010 study. A curve is clearly visible showing the increasing effect of detector dead-time as rates increase.

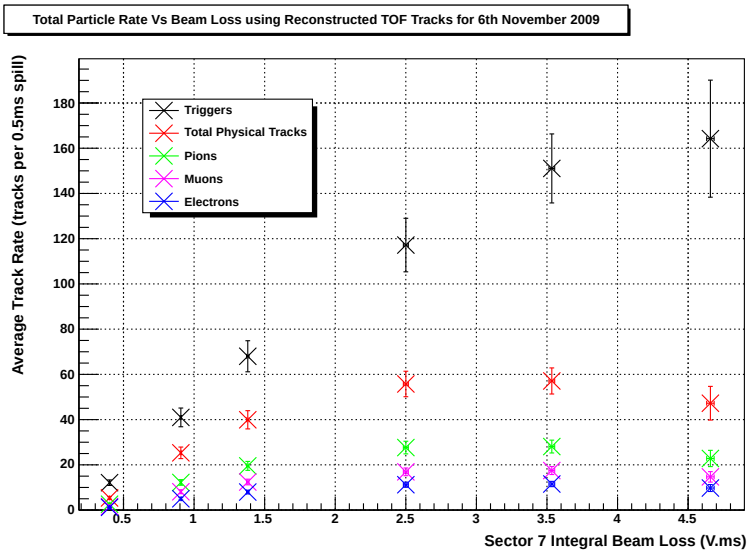

Figure 4.48: Average particles rates per species per spill as a function of beam loss for the 6th November 2009 study. The rates are determined by looking at reconstructed TOF tracks. The total number of MICEEvents in a spill is interpretted as the the number of particle triggers. A clear saturation is observable above ~ 2 V. This is most likely due to detector dead-time and software reconstruction inefficiencies when reconstructing the tracks.

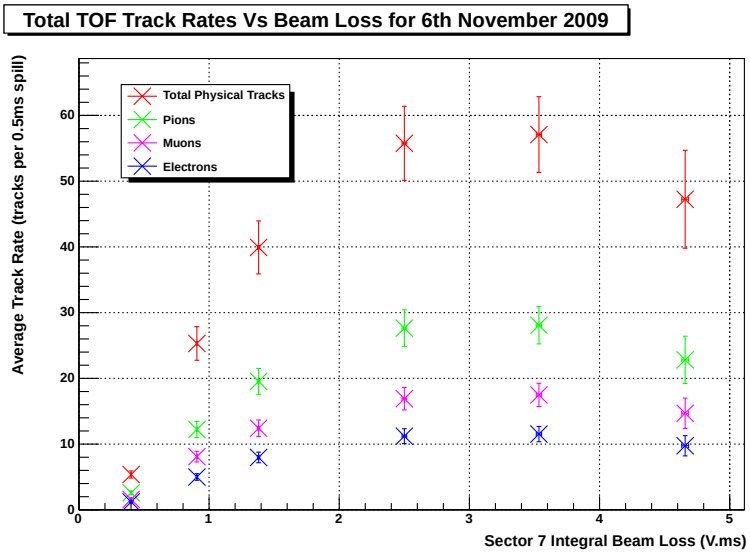

Figure 4.49: As Fig. 4.48 with only reconstructed tracks showing.

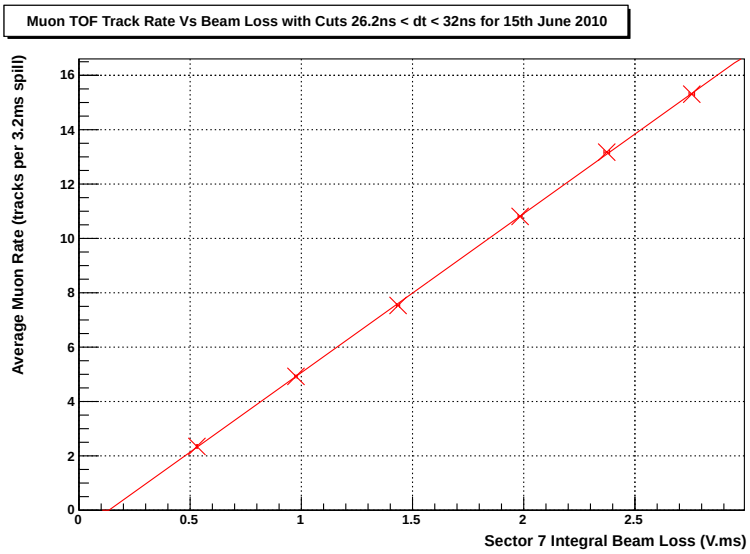

Figure 4.50: Average muon rate per spill from reconstructed TOF tracks as a function of beam loss for the 15th June 2010. The cuts applied to the TOF spectrum to isolate the physical tracks from muons are 26.2 ns $< dt <$ 32 ns.

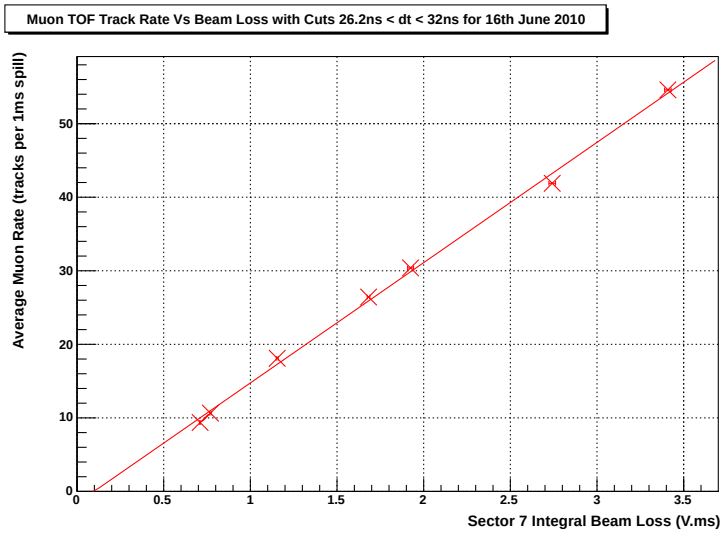

Figure 4.51: Average muon rate per spill from reconstructed TOF tracks as a function of beam loss for the 16th June 2010. The cuts applied to the TOF spectrum to isolate the physical tracks from muons are 26.2 ns $< dt <$ 32 ns.

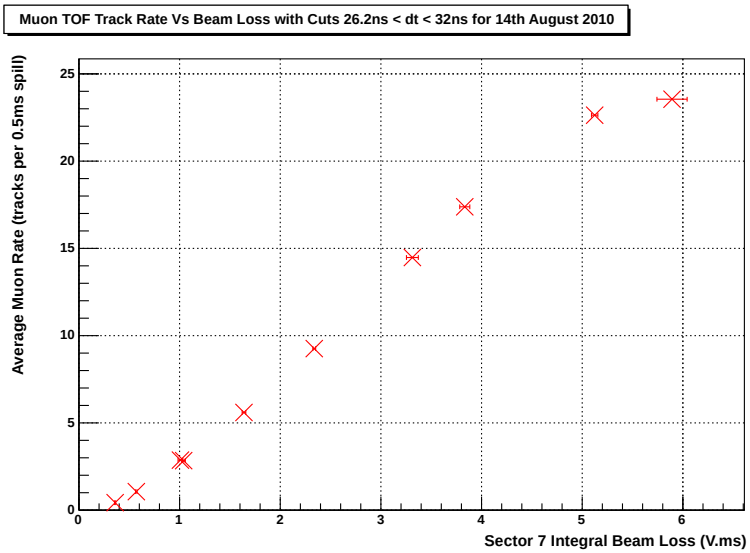

Figure 4.52: Average muon rate per spill from reconstructed TOF tracks as a function of beam loss for the 14th August 2010. The cuts applied to the TOF spectrum to isolate the physical tracks from muons are 26.2 ns $< dt <$ 32 ns. A small non-linear increase at lower beam losses can be observed, together with a non-linear decrease at higher beam losses.

4.5 Efficiency: Scaler Hits to TOF Tracks

Table 4.2 contains a summary of the results for the June studies in terms of both scaler hits and reconstructed TOF tracks. As can be seen, a substantial number of hits which register in the scaler channels are not subsequently reconstructed into TOF tracks. There are at least three contributing factors, touched on previously, which causes this difference: dead-time; software reconstruction effciency; and neutral particles.

As stated in Section 3.2.3 the scaler channels record every hit that occurs in each detector within the spill gate, while the TDC data used to create TOF tracks only records hits in coincidence with a particle trigger, again within the spill gate. In addition, any trigger request received in coincidence with another trigger is rejected, giving rise to dead-time, and so will not lead to a reconstructed TOF track, but will register as a scaler hit. The effect from DAQ dead-time is quantified by looking at the number of accepted particle triggers as a function of the number of particle trigger requests; these were shown for the June studies in Figs. 4.45 and 4.46. As can be seen the effect is negligible for the lower rate 15th study, and cause a loss of about 15% of triggers at the highest beam losses for the 16th. Dead-time plays an even larger role at higher beam losses, causing an approximately 50% loss of triggers at the highest beam losses in the November 2009 study (Fig. 4.43) and approximately 30% for the Aug 2010 study (Fig. 4.47).

In addition to dead-time the reconstruction efficiency of the TofRec class in G4MICE decreases with increasing particle rate, as multiple coincidences within a trigger window make it harder to disentangle which hits belong to which tracks. The effect has not yet been systematically studied, see Section 6.3 for more details.

Neutral particles can also contribute to the difference in rate between

scalers and TOF tracks. Neutrals are only capable of registering a hit in a TOF station by undergoing an interaction which removes them from the beam. As such it is not possible for a neutral to form a TOF track, barring a random coincidence of two neutrals, one at each TOF station; an effect which is thought to be extremely small. Neutrals may however contribute to scaler hits without this impediment, again leading to a reduction in the ratio of TOF tracks to scaler hits.

4.6 Beam Loss and Target Delay

The previous studies induced changes in beam loss by altering the target BCD while keeping the target short delay constant. The results presented in this section show the effect on particle rate of changing the beam loss using the target short delay while keeping the BCD constant. The results are given in terms of the target dip time (see Section 3.2.2) obtained from data, rather than in terms of the target short delay (see Section 2.7.2) used to define the run. Table 4.3 shows the target dip times obtained for given input target short delays (at 30.3 mm BCD), allowing for a conversion between the two systems.

Run	**Short Delay (bin)**	**Short Delay (ms)**	**Dip Time (ms)**
2890	0010001111	14.3	12.550
2889	0010001011	13.9	12.196
2891	0010001001	13.7	11.935
2888	0010000111	13.5	11.708
2893	0001111111	12.7	11.066
2894	0001111011	12.3	10.645
2895	0001111001	12.1	10.414

Table 4.3: Target short delay values and their equivalent target dip times for a BCD of 30.3 mm.

The effect of changing the dip time on the target parabola and the induced beam loss profile in the ISIS cycle is illustrated in Figs. 4.53 and 4.54. These show the latest (run 2890) and earliest (run 2895) target dip times of the study repectively. The earlier dip time can be seen to increase the induced beam loss and lead to substantial losses beginning to appear earlier in the cycle.

In a similar manner to the target depth study, plots of beam loss as a function of target dip time over the whole study may be produced. This is shown in run averaged form in Fig. 4.55, and Luminosity Monitor rate as a function of target dip time in Fig. 4.56. In both cases a roughly linear variation is visible for dip times from 12.6 ms to approximately 11 ms, before tailing off for earlier dip times (in contrast to the non-linear variation of target BCD with beam loss that was shown in Fig. 4.5).

A number of possible effects may be contributing to this trend away from linearity at early dip times. The ISIS beam envelope shrinks over the course of the acceleration cycle. For earlier dip times the target will be moving slower close to extraction (as illustrated in Figs. 4.53 and 4.54), allowing the beam to pull away from the target, increasing the effective target BCD. This effect may be enhanced by the beam orbit perturbation used for extraction. The perturbation begins $\sim$1.8 ms before extraction and causes a reduction in the beam position near the MICE target. Further, the sensivity per proton of the BLMs is lower at earlier times in the ISIS cycle (as described in Table 2.1), and earlier dips may produce a higher proportion of their losses at earlier times. Further study would be beneficial to pin down the matter in a definitive manner.

In addition to beam loss, the subsequent particle rate in the MICE beam line may also be plotted as a function of target dip time, as shown in

Fig. 4.57. Using the same procudure as before of applying cuts to the TOF distribtution, the reconstructed muon track rate may be plotted against target dip time, shown in Fig. 4.58. In both cases the relationship between dip time and induced rate is approximately linear above 11 ms, below which the curves tend towards plateaux. This may be attributed in part to the trend away from linearity at early dip times of beam loss with target dip time, shown in Fig. 4.55.

The trend away from linearity however appears slightly stronger for particle rates than for beam loss. This may be explained by particles with lower energy being lost during the momentum selection which occurs at the dipoles D1 and D2. This could lead to an additional reduction in rate beyond that due to the previous reduction in beam loss. This may be clarified by looking at plots of rate as a function of sector 7 beam loss (induced by the changing target dip time, in contrast to the plots of Section 4.3 and 4.4). Fig. 4.59 shows total detector rates as a function of beam loss, while muon track rate as a function of beam loss is shown in Fig. 4.60. Again, the trend is linear above 11 ms, tending towards plateaux below. This implies that there is an effect causing the particle rate to move away from linearity beyond that occuring due to a similar trend away from linearity in beam loss.

Further, if momentum selection, or a similar effect occuring within the MICE beamline, is reducing the particle rate in addition to the lower beam loss, then the variation of Luminosity Monitor rate with beam loss should remain approximately linear as the Luminosity Monitor is situated outside the beamline. This variation is shown in Fig. 4.61 and indeed, the trend is very close to linear. This supports the assumption that the plateau effect is caused by momentum selection at the dipoles or a similar, beamline-based effect, but more work is required to reach to a solid conclusion (see Chapter 6

for further discussion). Whatever the cause however it can be stated that the optimum target dip time for particle production in a 10 ms spill gate, together with the other conditions present on the 14th August 2010 (in particular at 30.3 mm BCD), occurs at approximately 11 ms, corresponding to a target short delay of 12.7 ms.

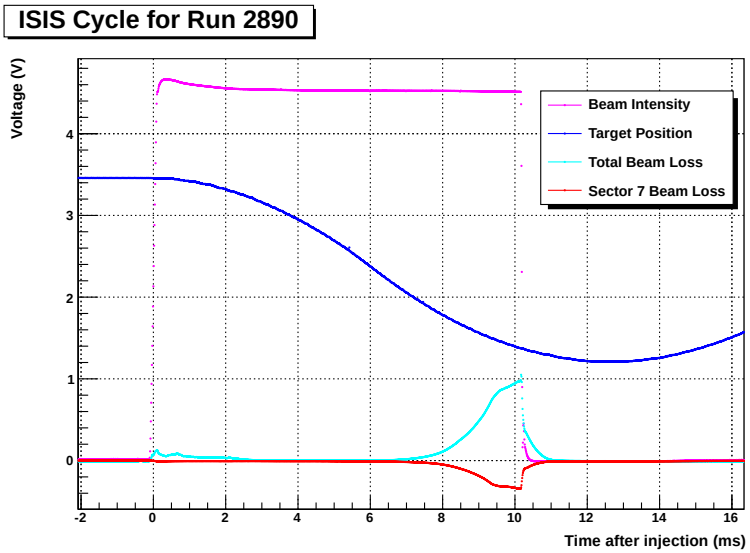

Figure 4.53: The averaged ISIS cycle for run 2890. The data for the curves has been avaraged over the whole run, using 199 target pulses. The target dip time is the latest in the delay study at 12.55 ms (equivalent to a short delay of 14.3 ms). The sector 7 integral beam loss is 0.539 V.ms

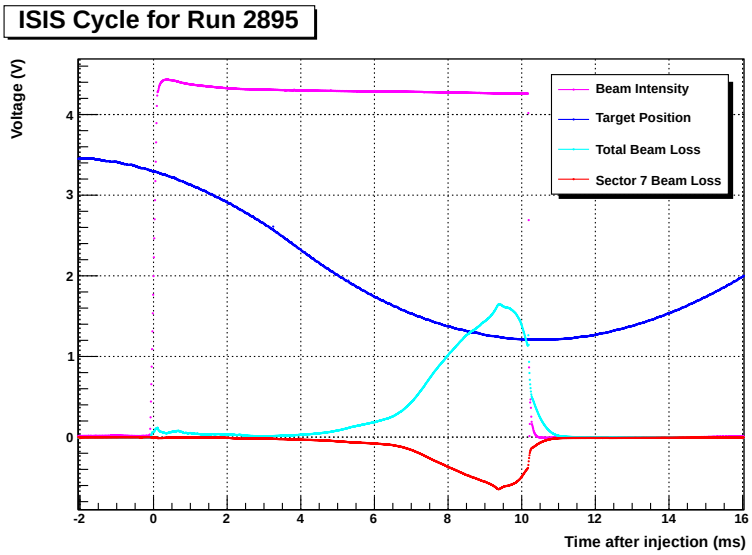

Figure 4.54: The averaged ISIS cycle for run 2895. The data for the curves has been avaraged over the whole run, using 197 target pulses. The target dip time is the earliest in the delay study at 10.41 ms (equivalent to a short delay of 12.1 ms). The sector 7 integral beam loss is 1.66 V.ms.

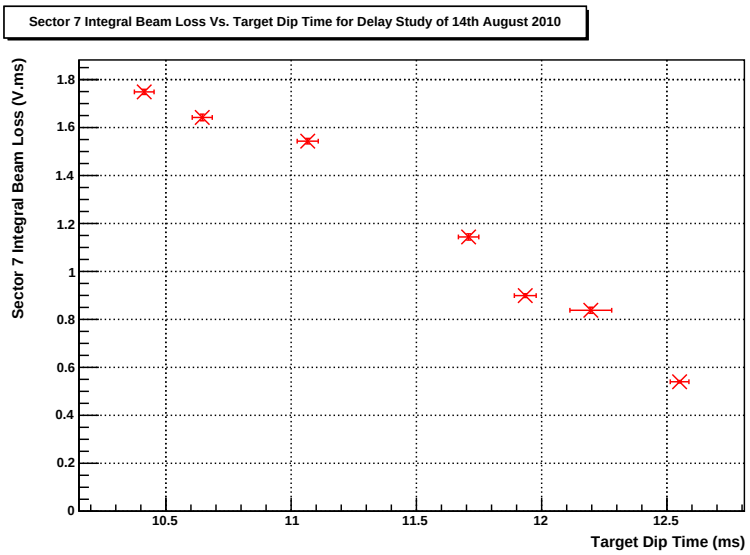

Figure 4.55: Integrated sector 7 beam loss as a function of target dip time.

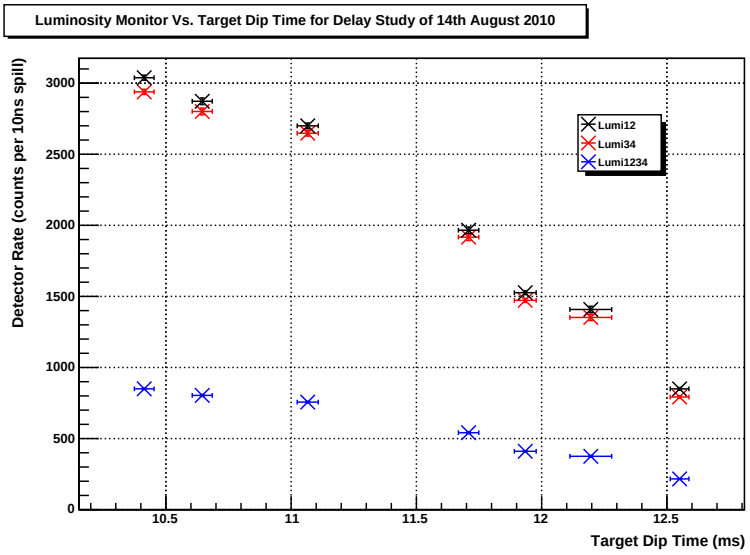

Figure 4.56: Luminosity Monitor rate as a function of target dip time

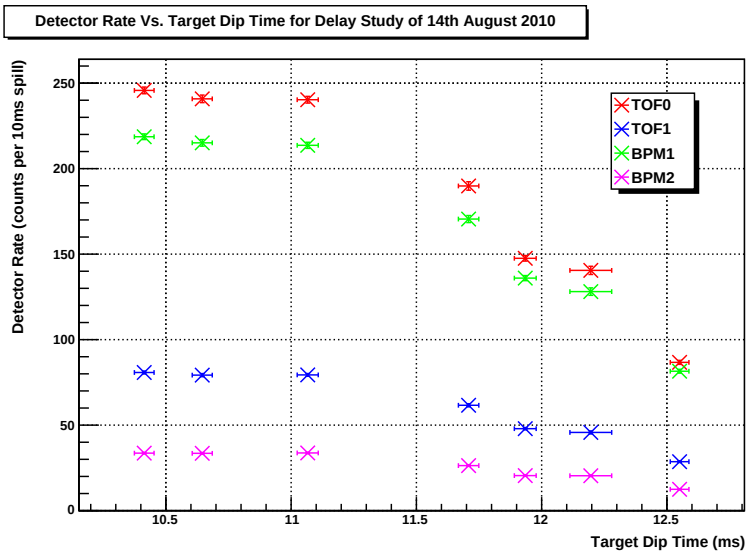

Figure 4.57: Total particle rates as a function of target dip time.

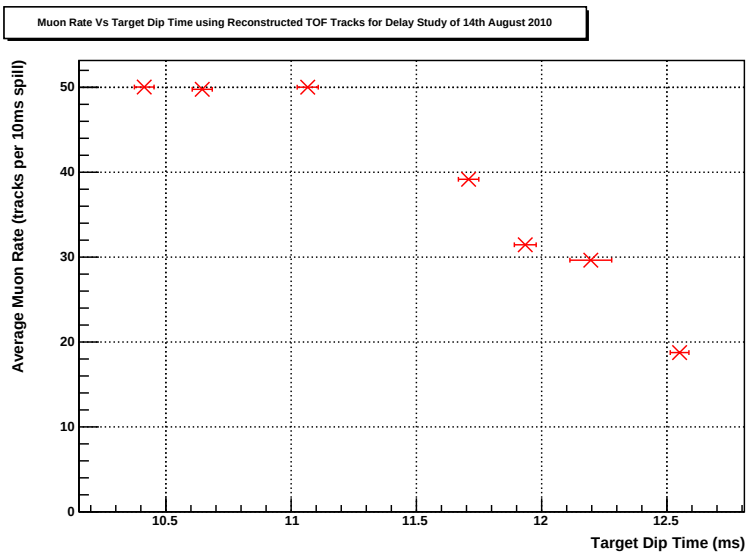

Figure 4.58: Reconstructed muon TOF track rate as a function of target dip time.

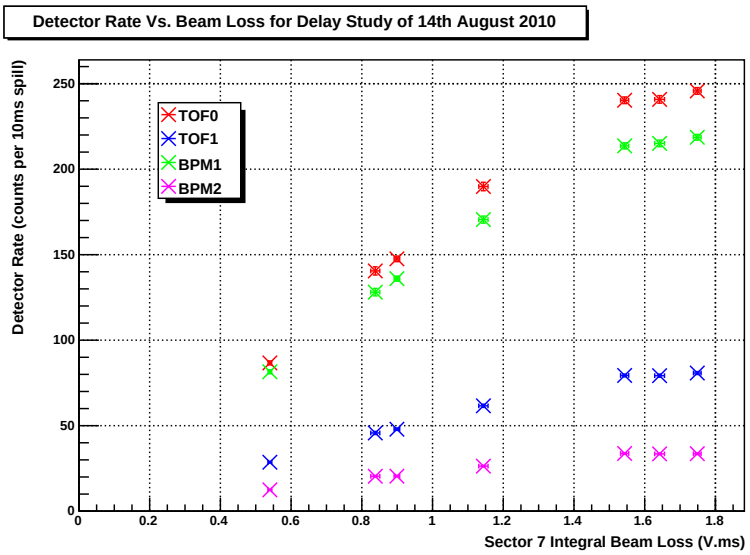

Figure 4.59: Total particle rates as a function of sector 7 integral beam loss, induced by changes in target dip time.

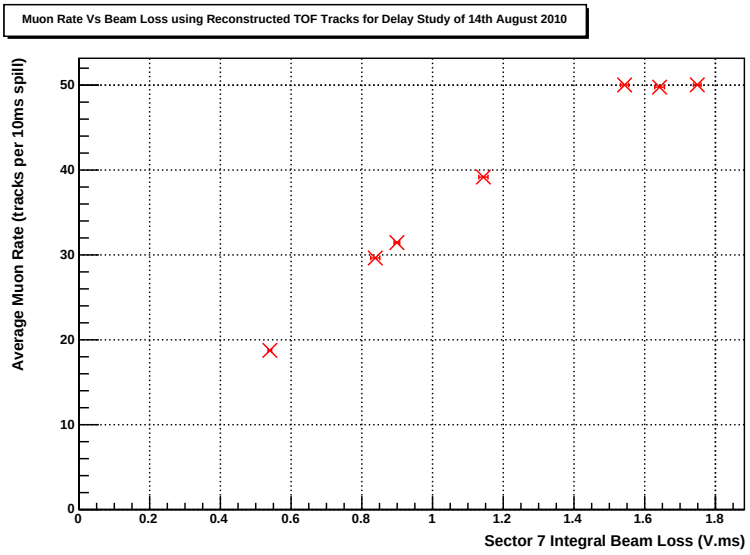

Figure 4.60: Reconstructed muon TOF track rate as a function of sector 7 integral beam loss, induced by changes in target dip time.

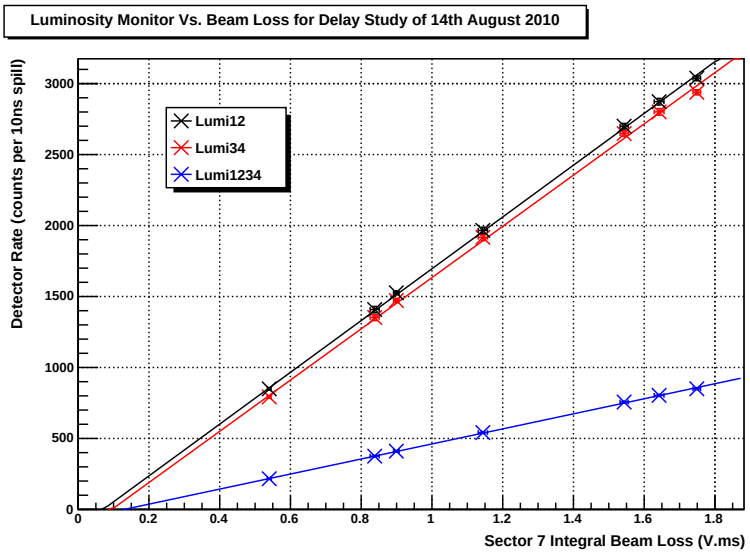

Figure 4.61: Luminosity Monitor rate as a function of sector 7 integral beam loss, induced by changes in target dip time.

Chapter 5

Simulations of ISIS Beam Loss

Never trust a computer you can't throw out a window.

Steve Wozniak

5.1 Motivation

In order to increase the particle rate in the MICE Muon Beamline to the desired levels it has already proven necessary to increase the beam loss levels in ISIS beyond any previously seen, and it may prove necessary to go still higher. It is important to understand the effects of these large losses on ISIS, in particular with regard to activation, and how they may be mitigated, for example by an upgrade of the collimator system. Simulating the effect of the MICE target on ISIS beam loss allows a greater understanding to be developed of where in the ISIS ring the extra beam losses are deposited, and how altering the target parameters, such as the material, affects this.

5.2 The ORBIT Particle Tracking Code

The Objective Ring Beam and Injection Tracking code (ORBIT) was developed in the late 1990s at the SNS, Oakridge, USA. It provides particle tracking for rings, with a particular emphasis on high intensity rings as it includes collective effects, such as space-charge, via "particle-in-cell" (PIC) methods. The tracking is symplectic, with the lattice functions being generated externally by the MAD programme[74].

Particles within ORBIT are represented by objects known as "macro particles", which are pushed around the accelerator ring in collections known as "herds" (in these studies only a single herd is used per simulation). The ring is represented by a series of elements known as nodes, which represent various accelerator structures, such as quadrupoles, dipoles, apertures and collimators, which act upon the macro particles as they are pushed around the ring. Diagnostic nodes may also be inserted as desired to read out the beam parameters at various points around the ring, while particles lost from the beam, such as from intersecting an aperture, are also recorded and may be read out.

The ORBIT code consists of a series of compiled C++ modules interfaced with a driver shell known as SuperCode (similar in design to the Python language). The compiled modules perform the general purpose, intensive physics calculations. The shell, via an interpreted script, is used to call these modules and control the flow of execution of the programme as a whole, including the layout of the ring, setting up loops for turns about the ring, and performing similar simulation dependent operations. For more details see [75, 76, 77].

5.3 Modelling ISIS and the MICE target

The ORBIT script used to described the ISIS ring, including the necessary MAD output, was provided by D. Adams (ISIS, STFC). A description of the MICE target, detailed in Section 5.3.1, was added to the ISIS script, together with a routine to calculate the target depth, described in Section 5.3.2.

A version of ORBIT customised for ISIS is used to perform the simulations, running on the STFC SCARF supercluster[78]. This may be run in one of two modes. In the first, which shall here be called the continuous mode, ORBIT is run as per its design, each turn following on immediately from the next with no break in the simulation. The second, called the turn-by-turn mode, writes out all the simulation data to disk after each turn, stops the simulation, then restarts the simulation for the next turn reading back into memory the data from disk, all controlled by a calling BASH script. The first method is far quicker and more straight forward, while the second is used in order to be able to alter the target depth turn-by-turn, as ORBIT does not permit the ring geometry to be modified whilst a simulation is in progress. As the target does not enter the beam until the last few milliseconds of the ISIS cycle, a continuous simulation is run up to 5000 turns while the target remains out of the beam, the output of which is fed into turn-by-turn simulations to cover the remaining $\sim$7000 turns.

Some example ORBIT output showing the transverse x, y coordinates of the particles in the ISIS beam after 5000 turns (at the start of the turn-by-turn simulations) is given in Fig. 5.1. The corresponding longitudinal phase space of the particles (ϕ and dE, the phase and energy difference with respect to the synchronous particle respectively) is also shown in Fig. 5.2.

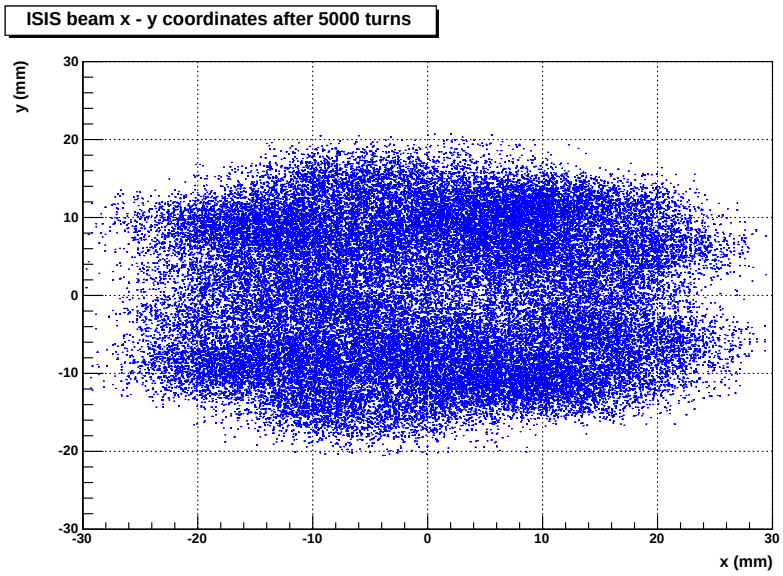

Figure 5.1: x, y coordinates of the particles in the ISIS beam after 5000 turns.

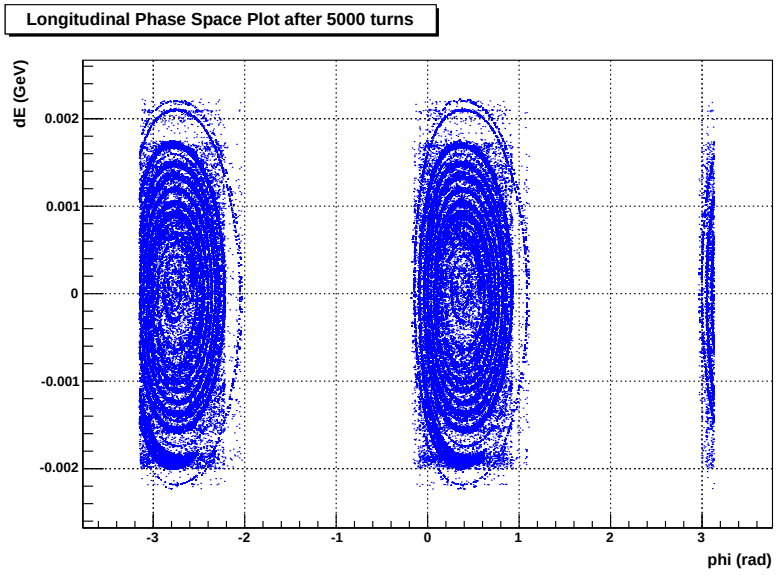

Figure 5.2: Longitudinal phase space coordinates of the particles in the ISIS beam after 5000 turns. The two buckets of the ISIS machine are clearly visible, while the distinct contours demonstrate the lack of space charge in the simulation (see Section 5.3.3).

5.3.1 Geometry and Material

The MICE target is modelled in ORBIT using the Collimator class. As titanium is not available in ORBIT the target is modelled, when in the standard configuration, as iron. Also, due to constraints on the shapes available for components within ORBIT, the target geometry is approximated by 10 rectangular collimator elements, as illustrated in Fig. 5.3. The length of each element in the longitudinal (z) direction and the transverse direction parallel to the floor (x) is determined to approximate a solid circle with the same area as the hollow circle of the target. As stated in Section 2.7.1 the target inner radius, r_{in}, is 2.275 mm and the outer radius, r_{out}, is 2.975 mm, so that the effective area, A_{eff}, is given by:

$$A_{eff} = A_{out} - A_{in} = \pi(r_{out}^2 - r_{in}^2) = 11.545 \text{ mm}^2 \tag{5.1}$$

The radius of a solid circle with the same effective area, r_{eff}, is then given by:

$$r_{eff} = \sqrt{\frac{A_{eff}}{\pi}} = \sqrt{\left(r_{out}^2 - r_{in}^2\right)} = 1.917 \text{ mm} \tag{5.2}$$

The length of each collimator in z is then given by r_{eff} divided by half the number of collimator elements used to model the target (half as the radius is being used, rather than the diameter). This gives a value of $1.917 \text{ mm} \div 5 = 0.3834$ mm. The transverse width in x of each collimator is then set so as to approximate a circle of radius r_{eff}. For the first 5 collimators in order of increasing z the following formulae are used:

$$x_{max}(i) = \sqrt{r_{eff}^2 - a^2} \qquad \text{for} \quad 1 \leq i \leq 5 \tag{5.3}$$

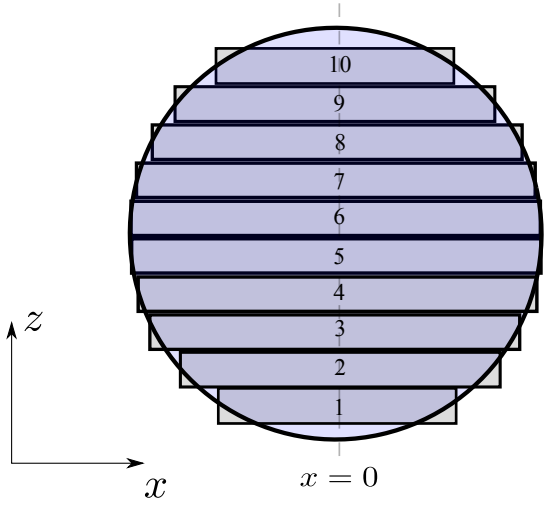

Figure 5.3: Top-down view of the ORBIT cylindrical target model, based on 10 collimator elements. Each element is 0.3834 mm in length in the z direction. The width in x is calculated using Eqns. 5.3 through 5.6. The ISIS beam travels in the direction of positive z.

$$a = \left(\frac{r_{eff}}{\frac{n}{2}}\right) \times \left(\left(\frac{n}{2} - i\right) + 0.5\right) \qquad \text{for} \quad 1 \leq i \leq 5 \tag{5.4}$$

$$x_{min}(i) = -x_{max}(i) \tag{5.5}$$

where i is the index of the collimators modelling the target, $i = 1$ having the lowest value of z and $i = 10$ the highest. n is the number of collimators modelling the target, having a value of 10. $x_{max}(i)$ is the coordinate of one end of the i^{th} collimator, $x_{min}(i)$ the other (with $x = 0$ being the middle of the beam, the target being centred around this). As the target cross section is a circle the remaining half circle for $6 \leq i \leq 10$ may be described as the

inverse of the first:

$$\begin{aligned} x_{max,min}(6) &= x_{max,min}(5) \\ x_{max,min}(7) &= x_{max,min}(4) \\ x_{max,min}(8) &= x_{max,min}(3) \\ x_{max,min}(9) &= x_{max,min}(2) \\ x_{max,min}(10) &= x_{max,min}(1) \end{aligned} \tag{5.6}$$

This then approximates a solid circle with the same effective area as the solid cross sectional area of the target cylinder. Output from ORBIT showing the simulated area is given in Table 5.1, showing that the simulated area matches the true to within ∼0.1 mm.

Collimator	$\mathbf{x_{min}}$	$\mathbf{x_{max}}$	**Area**
Mice1	-0.835601	0.835601	0.640749
Mice2	-1.36901	1.36901	1.049773
Mice3	-1.66017	1.66017	1.273038
Mice4	-1.8287	1.8287	1.402268
Mice5	-1.90739	1.90739	1.462609
Mice6	-1.90739	1.90739	1.462609
Mice7	-1.8287	1.8287	1.402268
Mice8	-1.66017	1.66017	1.273038
Mice9	-1.36901	1.36901	1.049773
Mice10	-0.835601	0.835601	0.640749
Total			**11.657**

Table 5.1: Collimator dimensions used for the Cylindrical MICE target model. All units are mm and mm^2. The beam is centred on $x = 0$.

Besides the standard cylindrical geometry, two alternatives are also used. The first, known as LongThin, consists of 10 collimator elements, each of length 1 mm. The width of each collimator is set so that the cross sectional

area presented to the beam is the same as in the real target:

$$w = \frac{A_{eff}}{10 \text{ mm}} = 1.1545 \text{ mm} \tag{5.7}$$

This then presents a long length target to the beam, with a narrow width.

In contrast the second alternative geometry, known as ShortFat, presents a wide but short target to the beam. 1 collimator is used, of length 1 mm (the other 9 collimators being set to be far above the beam pipe). The width is again set to mimic the cross sectional area of the real target:

$$w = \frac{A_{eff}}{1 \text{ mm}} = 11.545 \text{ mm} \tag{5.8}$$

For all geometries, the vertical y plane, corresponding to the direction along the length of the target shaft, is set with an upper coordinate far outside the beam pipe, and a lower coordinate determined by the target dip profile, discussed in Section 5.3.2 below.

5.3.2 Dip Profile

The target dip profile, that is the path of the target as it enters the ISIS beam, is modelled using target position data from the target DAQ. Profiles are created for a set of runs by averaging the position data over all the dips for each individual run, giving one profile per run. Each profile is then fitted between 0 to 10 ms of the ISIS cycle with a fourth order polynomial using ROOT, an example of which is shown in Fig. 5.4. The parameters of these fits are then used to create a function, accessible by ORBIT via a header file, which takes a run number and a time between 0 and 10 ms as arguments, and returns a value for the target depth.

During the initial continuous mode simulation the target depth is set to

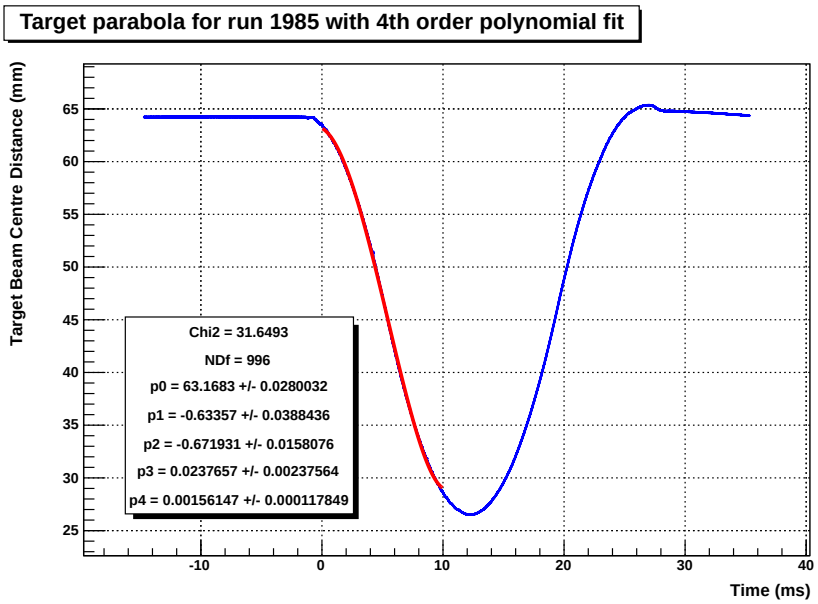

Figure 5.4: Averaged target dip profile for run 1985. The fourth order polynomial used to model the dip is shown in red, together with the parameters of the fit (p0 represents the value of the constant term, p1 the value of the coefficient of x^1, etc.)

a constant 99 mm BCD (safely out of the beam). For the following turn-by-turn mode simulations the target depth is calculated for each turn using the target depth function and the target collimators added accordingly. In practice it was found that this led to extremely low levels of lost particles, reflecting the fact that due to computing restrictions ORBIT simulations may only be conducted with a fraction of the true number of particles present in ISIS. In order to compensate for this the target profile was artificially lowered by a constant offset of 15 mm, in order to generate better lost particle statistics.

5.3.3 Simulation Parameters

In order to allow the simulations to be performed in a reasonable time, only a fraction of the total number of particles present in ISIS may be used. The number of particles injected into the ring in ORBIT is given by the number of injection pulses, set to 133, multiplied by the number of macro particles injected per turn, set to 250, giving a maximum herd size of 33250. This is far less than that required for reasonable space-charge calculations, which produces a further increase in CPU time, and is therefore also left out.

5.4 Analysis Methodology

The main quantity of interest in this study is the number of particles lost from the ISIS beam, corresponding to beam loss in the real machine. In particular it is of interest to determine where the losses occur around the ring and at what point during the ISIS acceleration cycle. The simulations are set to record all particles lost during transport into and around the ring, outputting the turn number, longitudinal position, z, and position in the 6D phase space $(x, x', y, y', \phi, \mathrm{d}E)$ when each particle was lost. The turn number may then be translated into time in the acceleration cycle with respect to injection via a lookup table produced from other ORBIT output. This then allows the production of histograms of lost particles as a function z, or of time, or of both.

It is useful to split the 2D space of z and time into 4 approximately equal quadrants, considering the lost particle levels in each. The first quadrant corresponds to injection losses located at the ISIS collimators, referred to as "IC" (Injection Collimators). As these studies are concerned with losses due to the action of the MICE target, the losses here are a constant, stemming

from a single continuous mode simulation (as was described in Section 5.3). It is defined by the cuts $time < 5.0$ ms and $z < 87.1$ m, where the origin of the z axis is the exit face of the dipole at the begining of sector 0. The second quadrant corresponds to injection losses at the MICE target, referred to as "IT" (Injection Target). It is defined by the cuts $time < 5.0$ ms and $z > 87.1$ m and is again constant here. The third quadrant represents the losses due to the MICE target located at the collimators, referred to as "EC" (Extraction Collimators). It is defined by the cuts $time > 5.0$ ms and $z < 87.1$ m. The last quadrant represents the losses due to the MICE target located just after the target itself, referred to as "ET" (Extraction Target). It is defined by the cuts $time > 5.0$ ms and $z > 87.1$ m. In particular it is of interest to compare the ratio ET to EC for different target configurations.

In addition to lost particles, it is also interesting to look at how this translates into beam loss profiles. This is done by using the beam loss monitor sensitivity data given in Table 2.1. Linear fits are performed between each data point in the table to give functions relating protons incident on the BLMs to beam loss covering the whole 10 ms range of the ISIS cycle, which may then be used to convert lost particles to beam loss. Note that this gives a shape calibration only; as the number of particles in the ORBIT simulations is far less than present in the actual beam and due to the requirement of artificially lowering the target (as described in Section 5.3), it does not represent a calibration to the same scale as the actual data. An attempt at such quantitative callibration using beam loss data is however given in the depth study section below.

5.5 Depth Study

The target depth study consists of six simulations using target profiles taken from the 15th June 2010 study presented in Chapter 4. The lost particle distribution as a function of time in the ISIS cycle is shown in Fig. 5.5, together with the corresponding beam loss distribution in Fig. 5.6. The losses due to injection are visible on the left-hand side of plots, together with the losses due to the MICE target on the right. The increasing depth of the dynamic target is clearly visible in the loss distributions.

The lost particles distribution in terms of longitudinal position is shown in Fig. 5.7, with the corresponding beam loss distribution in Fig. 5.8. The collimators are positioned at ∼22 m, corresponding with the peak that can be seen on the left of the plots. The MICE target is positioned at ∼115 m, immediately after which another series of peaks are visible. The first of these at 120 m is caused by the proximity of the target. The second at 140 m corresponds to a quadrupole doublet in sector 7 which constricts the beam pipe aperture in both transverse directions. The third at 150 m corresponds to another quadrupole doublet, now in sector 8. For the sake of brevity these last three loss positions shall all be referred to as losses at or in the vicinty of the target.

2-dimensional histograms of the full time - longitudinal position space are shown in Fig. 5.9 for lost particles and Fig. 5.10 for beam loss. A peak is present in quadrant IC (bottom lefthand corner) of both histograms, corresponding to the beam losses which are due to injection, and are located at the collimators. Quadrant IT (top lefthand corner) is empty, showing there are no losses caused by injection which occur in the vicinity of the MICE target. Quandrants EC (bottom right corner) and ET (top right corner), both possess distinct peaks, corresponding to losses which are located at

Run	BCD (mm)	LP All	BL All (V.ms)	LP ET/EC
1985	26.5	10525	1.70×10^{-7}	2.06
1986	27.0	9718	1.43×10^{-7}	2.01
1987	27.6	8924	1.16×10^{-7}	2.10
1988	28.7	7544	6.97×10^{-8}	2.00
1989	29.9	6521	3.48×10^{-8}	2.01
1991	31.9	5707	6.46×10^{-9}	1.74

Table 5.2: ORBIT depth study results. The target short delay is 13.1 ms for every run. "LP" refers to lost particles, "BL" to beam loss, "All" over the whole ring over the whole acceleration cycle, "ET" to particle losses caused by the MICE target in the vicinity of the target, "EC" to particle losses caused by the target in the vicinity of the ISIS collimators.

the collimators and due to the action of the MICE target, and losses which are located in the vicinity of the MICE target and due to the action of the MICE target, respectively. By comparing Fig. 5.9 and Fig. 5.10 it can be seen that in moving from lost particles to beam loss the emphasis shifts from losses due to injection, to losses due to the action of the MICE target (due to the energy - sensitivity curve of the BLMs, shown in Table 2.1).

The results for beam loss distributions produced around the ISIS ring for different target BCDs are summarised in Table 5.2. The run number used to created the target dip profile is given, together with the target BCD, the total number of lost particles produced, the equivalent beam loss and the ratio of losses in ET to EC. The target short delay has a constant value of 13.1 ms. Little variation is observed in the ratio of ET to EC as a function of target BCD (which may be compared with the results in Geometry and Materials studies below).

It is also interesting to look at how beam loss in ORBIT varies with target BCD and then compare this directly with the results from data. Fig. 5.11 shows the ORBIT results of ET beam loss as a function of target

BCD, while the results from data for the 15th June 2010 study, plotting the integrated sector 7 losses as function of target BCD, are shown directly beneath in Fig. 5.12. The shapes between data and simulation can be seen to match very well.

The procedure of matching ORBIT results to data may then be taken further by considering how to normalise ORBIT results quantitatively. Due to the ORBIT simulation running many times fewer particles than are present in ISIS, the need to lower artificially the target depth by a fixed offset, and the lack of space-charge effects, it is to be expected that such a normalisation is necessary, as indeed the beam loss scales in Fig. 5.11 and Fig. 5.12 illustrate. Fig. 5.13 shows the total beam loss levels observed in ORBIT as a function of total integral beam loss levels from data, correlated by run number, while Fig. 5.13 shows ORBIT ET losses related to sector 7 integral losses from data. In the case of the total losses the relation between the two is complicated, and no straightforward scaling can be applied. More work is necessary to understand this effect better. In the case of the sector 7 data, while again it is not sufficient to relate simulation to data with a constant multiplicative factor, a linear fit provides a very good match, thus allowing ORBIT to make actual beam loss level predictions beyond that available in data. It should be borne in mind however that as the relationship between beam loss and target depth varies between studies (as was illustrated in Fig. 4.5), any normalisations between simulation and data would also be subject to the conditions present when the data was taken. This implies that in order to estimate beam loss levels beyond data for a study, some beam loss data points would be necessary to give the normalisation.

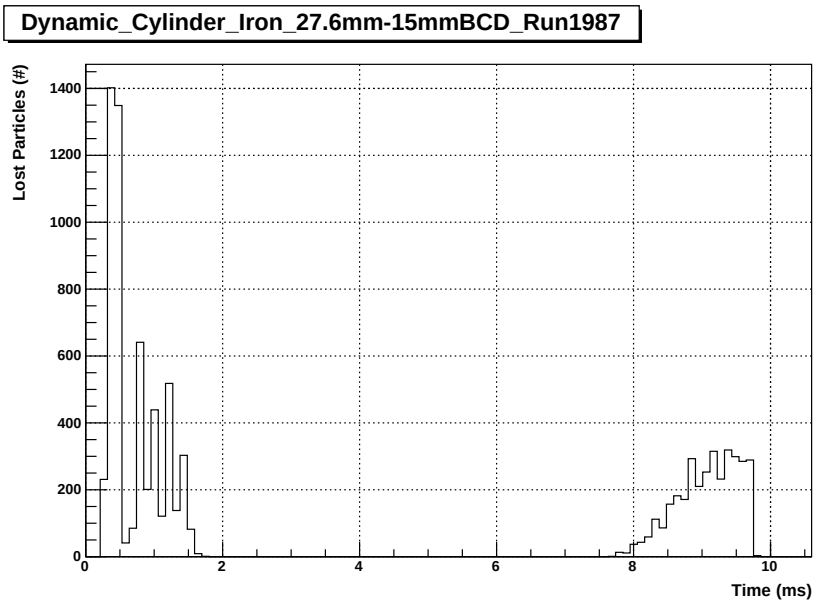

Figure 5.5: Histogram of lost particles as a function of time in the ISIS cycle for a MICE target profile approximating run 1987.

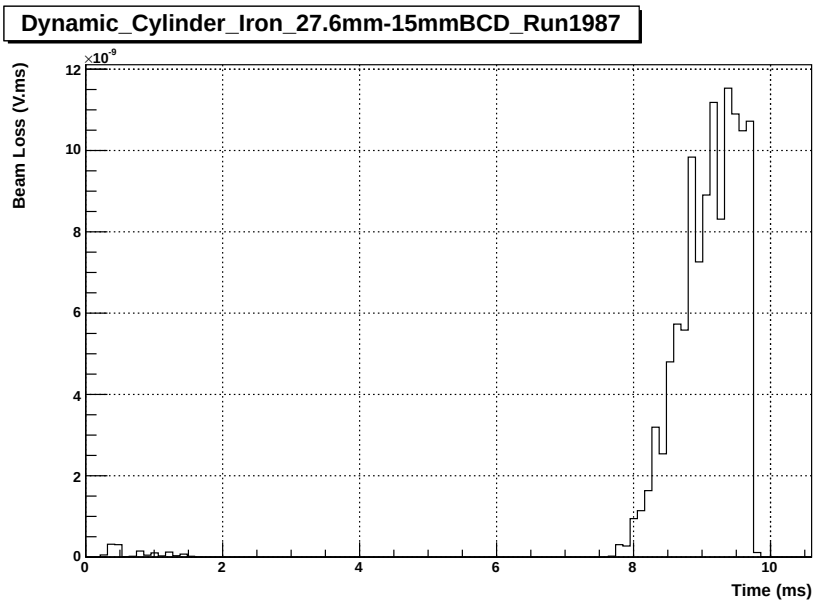

Figure 5.6: Histogram of beam loss as a function of time in the ISIS cycle for a MICE target profile approximating run 1987.

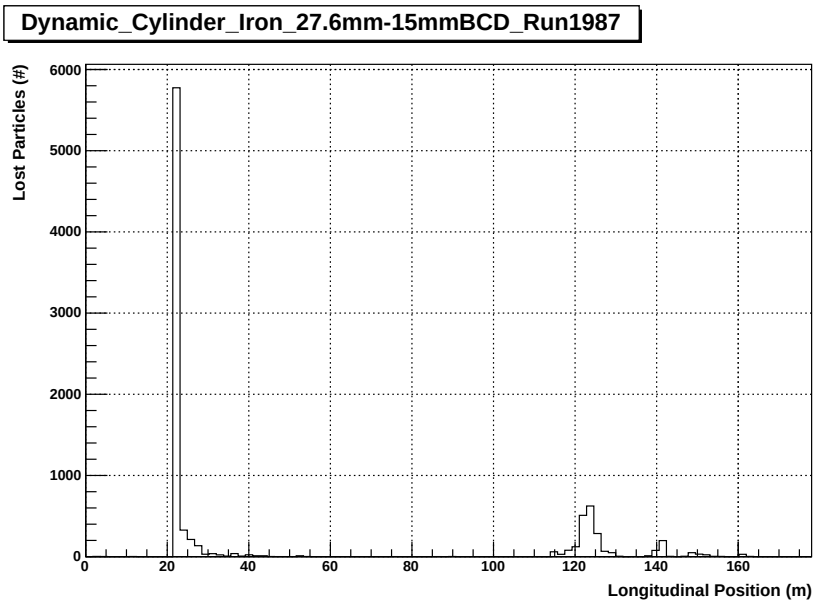

Figure 5.7: Histogram of lost particles as a function of position around the ISIS ring for a MICE target profile approximating run 1987.

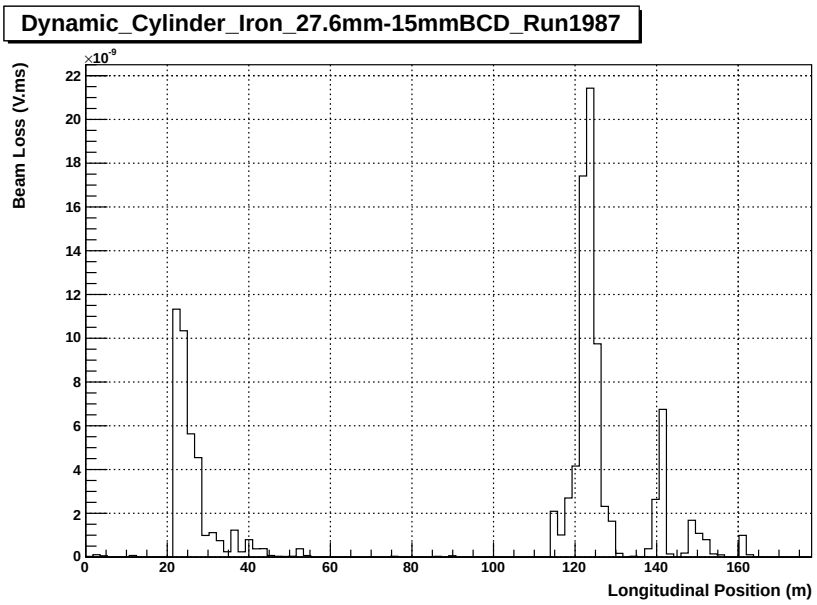

Figure 5.8: Histogram of beam loss as a function of position around the ISIS ring for a MICE target profile approximating run 1987.

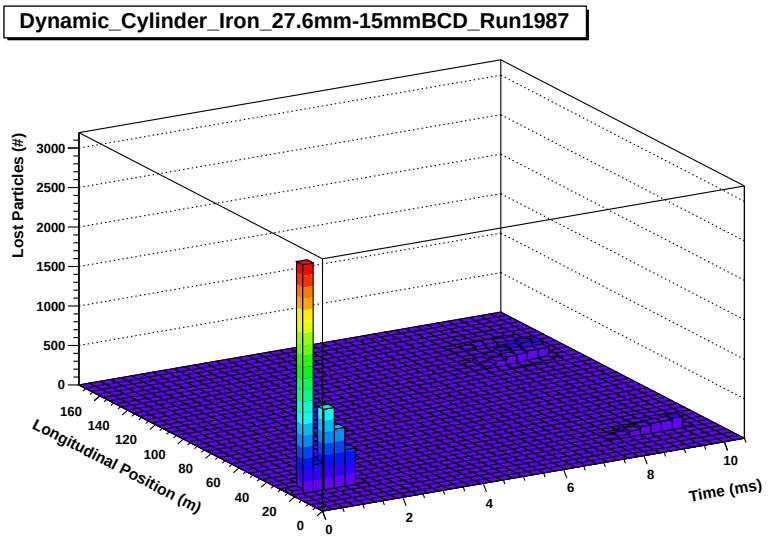

Figure 5.9: 2-dimensional histogram of lost particles as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987.

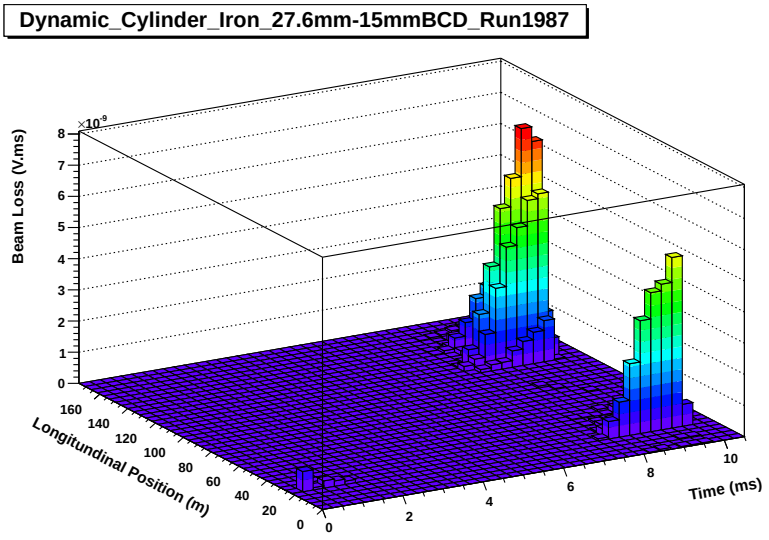

Figure 5.10: 2-dimensional histogram of beam loss as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987.

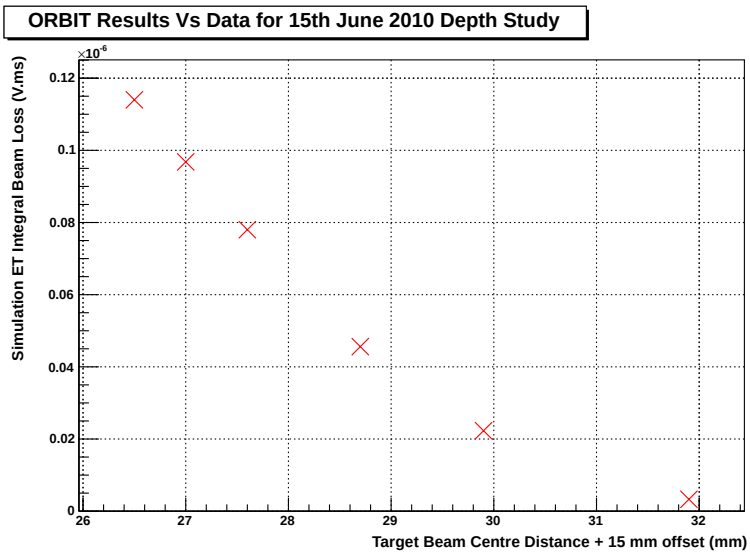

Figure 5.11: ORBIT ET beam loss as a function of target BCD, using target profiles generated from the 15th June 2010 study.

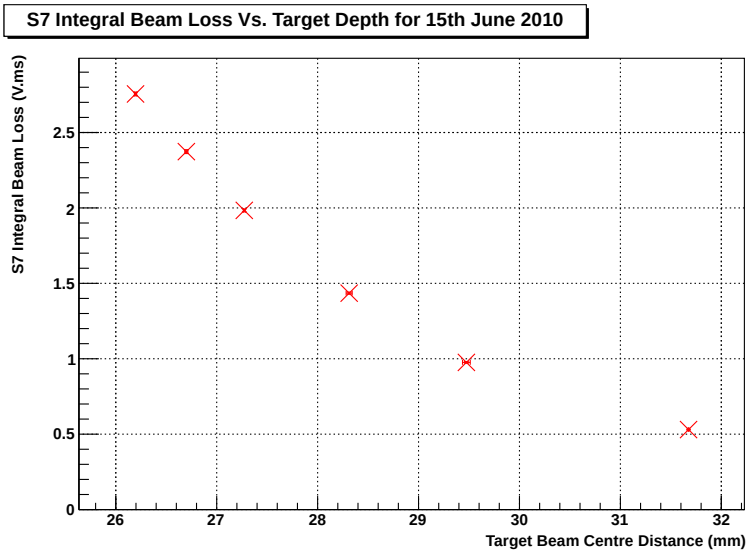

Figure 5.12: Sector 7 integral beam loss as a function of target BCD from data for the 15th June 2010 study.

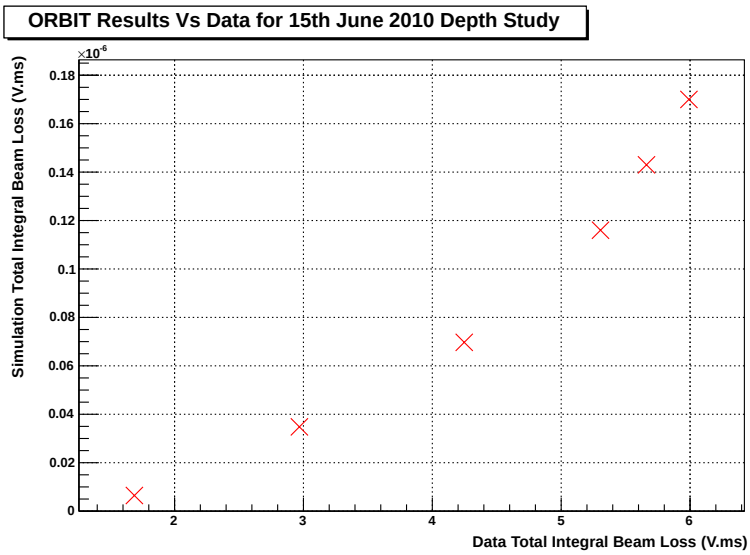

Figure 5.13: ORBIT total beam loss as a function of total integral beam loss for the 15th June 2010 study.

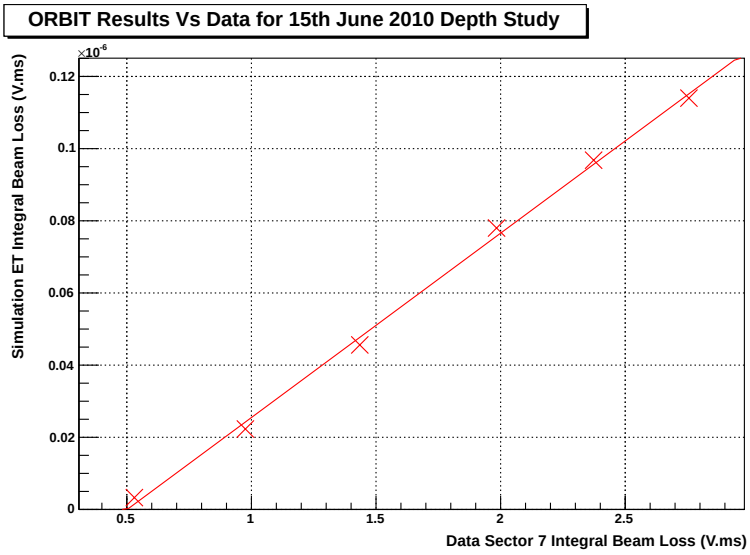

Figure 5.14: ORBIT ET beam loss as a function of sector 7 integral beam loss for the 15th June 2010 study. A linear fit is also shown, possessing a gradient of 5.11×10^{-8} and an intercept of -2.56×10^{-8} V.ms.

5.6 Material Study

The target materials study consists of three simulations, using a low atomic weight target composed of carbon (Z = 6), a medium atomic weight target approximating the real target, made from iron (Z = 26), and a high atomic weight target made from tungsten (Z = 74). The target short delay is again set to a constant value of 13.1 ms, using the target dip profile of run 1987, corresponding to a BCD of 27.5 mm.

2-dimensional histograms of lost particles and beam loss as a function of longitudinal position and time for a carbon target are shown in Figs. 5.15 and 5.16. The equivalent plots for tungsten are shown in Figs. 5.17 and 5.18. These may be compared with the plots for iron already shown in Figs. 5.9 and 5.10. A clear shift of losses away from the collimators and towards the MICE target is evident as the atomic number increases. This may be quantified by looking at the ET to EC ratio, as shown in Table 5.3.

Material	**LP All**	**BL All (V.ms)**	**LP ET/EC**
Carbon	8111	8.99×10^{-8}	0.26
Iron	8924	1.16×10^{-7}	2.10
Tungsten	9034	1.20×10^{-7}	8.92

Table 5.3: ORBIT material study results. The target short delay is 13.1 ms for every run. "LP" refers to lost particles, "BL" to beam loss, "All" over the whole ring over the whole acceleration cycle, "ET" to particle losses caused by the MICE target target in the vicinity of the target, "EC" to particle losses caused by the target in the vicinity of the ISIS collimators.

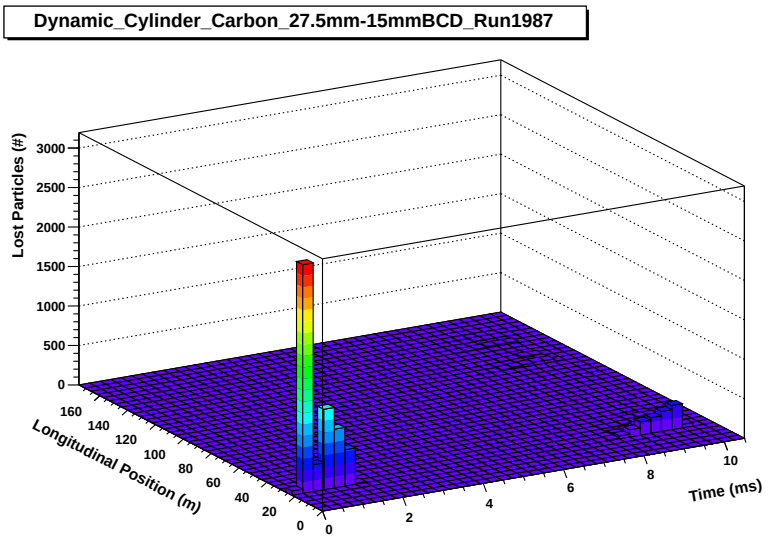

Figure 5.15: 2-dimensional histogram of lost particles as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a carbon target.

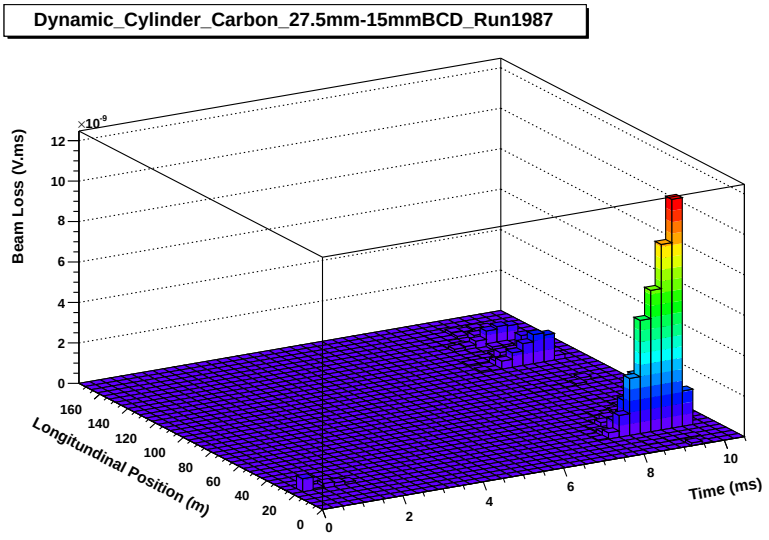

Figure 5.16: 2-dimensional histogram of beam loss as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a carbon target.

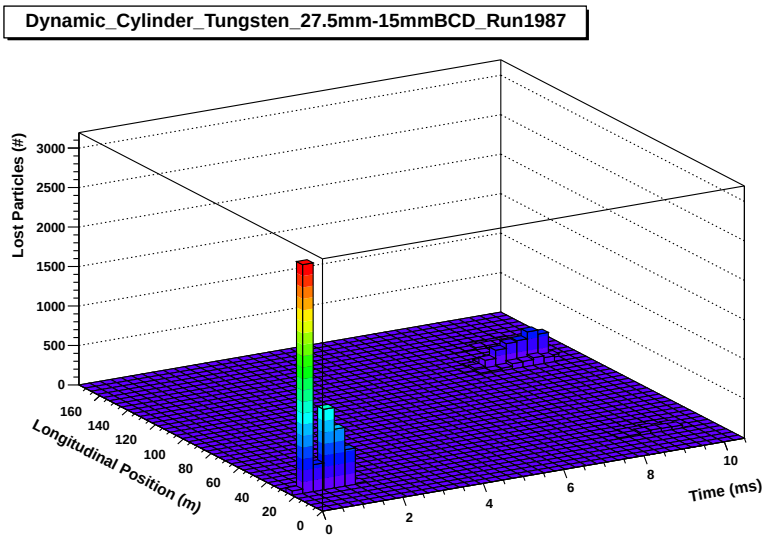

Figure 5.17: 2-dimensional histogram of lost particles as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a tungsten target.

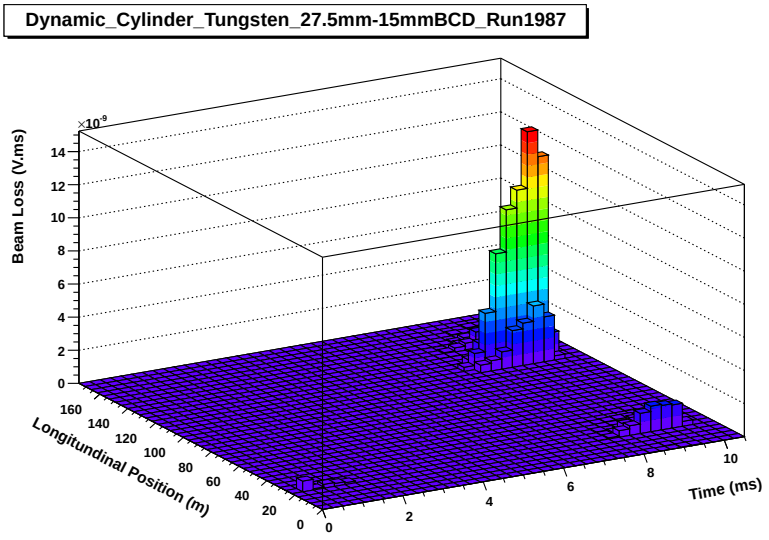

Figure 5.18: 2-dimensional histogram of beam loss as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a tungsten target.

5.7 Geometry Study

The target geometry study consists of three simulations, which use the Cylindrical, LongThin and ShortFat target geometries, as described in Section 5.3.1. For all the simulations the target material used is iron, the target short delay is set to a constant value of 13.1 ms, and the target dip profile modelled using run 1987, corresponding to a BCD of 27.5 mm.

The usual 2D histograms of lost particles and beam loss as a function of longitudinal position and time for the LongThin geometry are shown in Figs. 5.19 and 5.20. The equivalent plots for ShortFat are shown in Figs. 5.21 and 5.22. These may be compared with plots for the standard cylindrical geometry already shown in Figs. 5.9 and 5.10. For LongThin the losses are strongly biased to fall in the ET quadrant after the MICE target, compared with the standard cylindrical geometry which shows a smaller bias in favour of losses at the target. For ShortFat the pattern shifts still futher, with slightly more losses now being present in the vicinity of the collimators, than are present at the target. Again, this may be quantified by looking at the ET to EC ratio, as shown in Table 5.4.

Material	**LP All**	**BL All (V.ms)**	**LP ET/EC**
LongThin	8631	1.07×10^{-7}	5.39
Cylindrical	8924	1.16×10^{-7}	2.10
ShortFat	9036	1.20×10^{-7}	0.90

Table 5.4: ORBIT geometry study results. The target short delay is 13.1 ms for every run. "LP" refers to lost particles, "BL" to beam loss, "All" over the whole ring over the whole acceleration cycle, "ET" to particle losses caused by the MICE target target in the vicinity of the target, "EC" to particle losses caused by the target in the vicinity of the ISIS collimators.

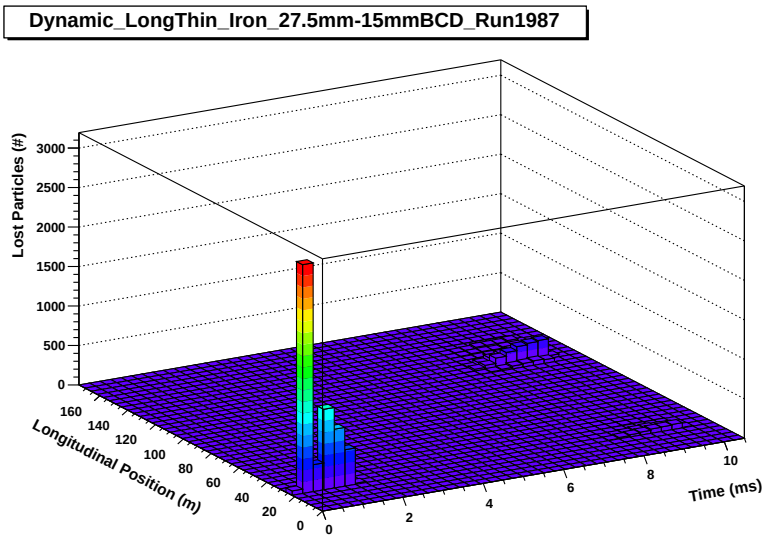

Figure 5.19: 2-dimensional histogram of lost particles as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a LongThin target geometry.

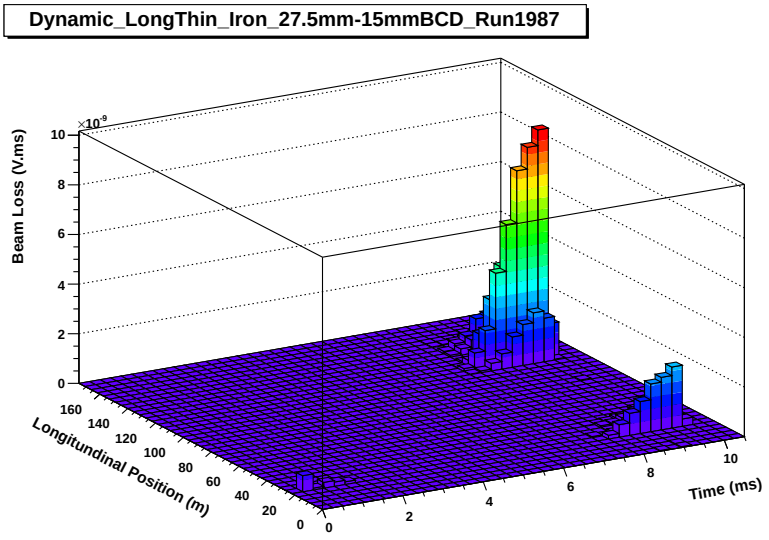

Figure 5.20: 2-dimensional histogram of beam loss as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a LongThin target geometry.

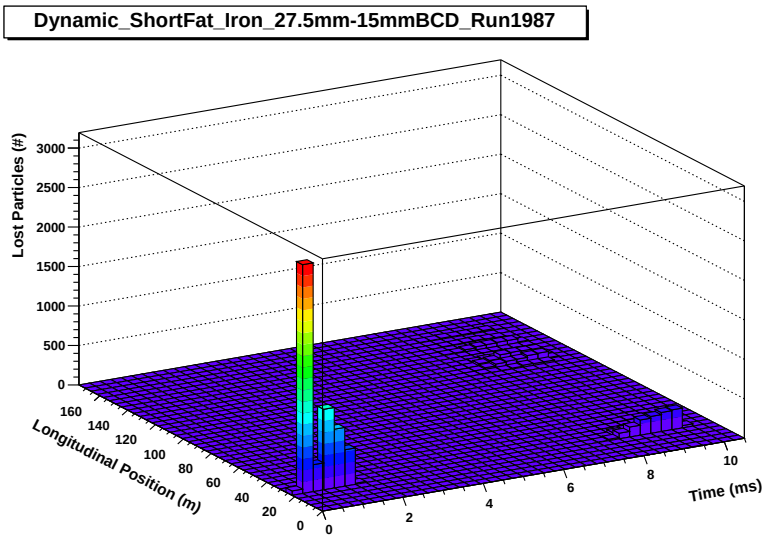

Figure 5.21: 2-dimensional histogram of lost particles as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a ShortFat target geometry.

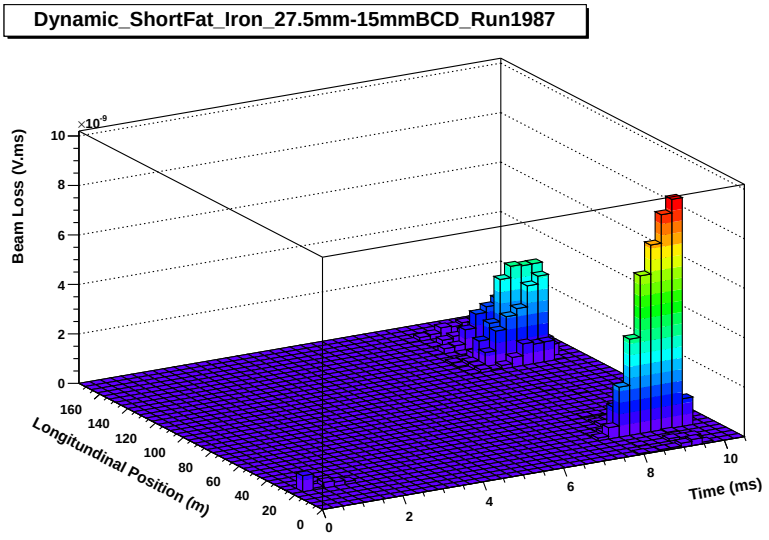

Figure 5.22: 2-dimensional histogram of beam loss as a function of position around the ISIS ring and time in the ISIS cycle for a MICE target profile approximating run 1987 with a ShortFat target geometry.

Chapter 6

Conclusion

Sosban fach yn berwi ar y tân,
Sosban fawr yn berwi ar y llawr,
A'r gath wedi sgrapo Joni bach.

Sosban Fach, Welsh Folk Song

6.1 Summary of Findings

The principal findings of the beam loss and particle rate studies may be summarised as follows:

- The particle rate in the MICE beam line increases linearly with host accelerator beam loss from 0 - 4.7 V.ms as measured by the integrated losses in sector 7, when induced by changing the target dip depth. There is a strong indication that this trend continues up to at least 7.2 V.ms (10 V.ms on the ISIS scale);
- The absolute particle rates passing through TOF1 are described in Table 4.2, valid for the optics described in Appendix B. For negative $\pi \rightarrow \mu$ optics this gives an absolute rate at TOF1 of 13.6 scaler hits

per V.ms beam loss per 3.2 ms spill gate. For positive $\pi \rightarrow \mu$ optics this gives an absolute rate at TOF1 of 33.1 scaler hits per V.ms beam loss per 1 ms spill gate;

- The reconstructed muon TOF track rate increases linearly with beam loss for low instantaneous rate values. At higher values values this becomes a curve tending towards decreasing rate gains per unit beam loss due to effects such as DAQ dead time. At the highest loss point for the August 2010 study the ratio of accepted triggers to triggers requests was 1:1.41, illustrating the large dead time effect at high rates;

- The reconstructed muon TOF track rates between TOF0 and TOF1 are described in Table 4.2, valid for the optics described in Appendix B. For negative $\pi \rightarrow \mu$ optics this gives a muon rate of 5.9 tracks per V.ms beam loss per 3.2 ms spill gate. For positive $\pi \rightarrow \mu$ optics this gives a muon rate of 16.4 tracks per V.ms beam loss per 1 ms spill gate;

- Induced beam loss varies non-linearly with increasing target dip depth, as shown in Fig. 4.5;

- Induced beam loss varies linearly with increasing target dip time between approximately 11 and 12.5 ms, for a target BCD of 30.5 mm. Between approximately 11 and 10.4 ms the increase in beam loss per unit dip time is reduced, as shown in Fig. 4.55;

- Particle rate in the MICE beamline also varies linearly with increasing target dip time between approximately 11 and 12.5 ms, for a target BCD of 30.5 mm, but then reaches a plateau between approximately 11 and 10.4 ms as shown in Fig. 4.57.

In addition to the studies based on real data, the results from ORBIT simulations allow the following conclusions to be drawn about the nature of the beam loss induced in ISIS by the MICE target (both for the true target design and for possible alternative designs):

- The induced losses are primarily concentrated around the ring shortly after the position of the MICE target and at the ISIS collimator system;
- The material used to construct the target strongly affects the pattern of induced losses. Low Z materials such as carbon create losses preferentially at the collimators, while high Z materials such as tungsten produce greater losses in the vicinity of the target. Medium Z materials such as iron produce results intermediate between low Z and high Z materials;
- The geometry of the target employed also affects the beam loss distribution. The ShortThin geometry produces losses roughly evenly distributed between the collimators and the target, while LongThin produces greater losses at the target. The cylindrical geometry produces a loss distribution intermediate between ShortFat and LongThin.

6.2 Implications for MICE

The MICE beamline has been shown to be operating successfully, delivering muons for use by the MICE cooling channel for the demonstration of ionisation cooling, whilst causing a tolerable increase in ISIS beam loss levels. The rate observed however, falls short of the favoured rate for MICE operation, that of 600 "good" muons per 1 ms spill (muons which traverse

the whole cooling channel and then fall within the acceptance of TOF2)[79]. Further, should the beamline be able to achieve such a rate the DAQ at present would be unable to make use of it due to the dead time associated with particle triggers.

Addressing the question of decreasing the DAQ dead time is an issue under investigation by the MICE Online Group and the amendments necessary are beyond the scope of this thesis. Much may be said however about how best to address the issue of the particle rate itself.

Assuming that there is no constraint from the DAQ, a rough estimate may be made of the beam loss levels required to produce 600 muons per 1 ms spill at TOF1, for a setup identical to that present for the June 2010 studies. Looking first at raw hit rate, the linear fit shown for TOF1 in Fig. 4.15 for a positive beamline may be extrapolated to higher losses. This then indicates a required beam loss level of ∼18.5 V.ms to produce 600 hits per 1 ms spill. Applying the same technique to reconstructed muon TOF tracks using the linear fit shown in Fig. 4.51, again for a positive beamline, the indicated beam loss level becomes ∼37 V.ms. Both these values are not feasible for standard ISIS running. In the case of a negative beamline it is more difficult to produce estimates as the gate used on the 15th June 2010 was 3.2 ms rather than 1 ms, and the rate does not vary linearly across the gate. A rough figure would be at least a factor of 5 times beyond the beam loss necessary in the positive case (a factor of ∼2.5 from the rate reduction between positives and negatives and a further factor of ∼2 from the spill gate length), which is again impractical from an ISIS standpoint.

This issue of low rate per spill may be addressed in several ways. The figure of 600 good muons per spill which may require such large beam losses is calculated to allow the demonstration of ionisation cooling in a timely

fashion. Lower rates are however quite acceptable and will simply require longer running of the experiment. Further, the rate at which spills themselves arrive, *i.e.* the target dip rate, may also be increased, going from the present value of 1 dip every 2.56 or 3.2 seconds to something closer to 1 dip per second, decreasing the time needed to acquire sufficient statistics.

Steps may also be taken to improve the rate per spill, to bring the present value closer to the desired rate. Alternative target geometries are one possibility, such as moving from a cylinder to a rectangle with a large cross-sectional area. Alternative target materials are also being considered, such as moving to a high atomic number material to increase the number of hadronic interactions. Both these options however have constraints from the mechanical considerations of the target system.

A particularly ingenious method of increasing the rate per unit beam loss has also been proposed, whereby a vertical offset or "bump" is introduced into the ISIS beam orbit at the location of the MICE target. In effect the beam rises up to meet the target in addition to the target dipping to meet the beam. This causes the target - beam interaction to be more localised to the end of the ISIS cycle, where the MICE DAQ spill gate is located, leading to a greater rate efficiency. Initial tests of this technique have proven promising and work on implementing it is ongoing[80].

Lastly, there exists the possibility of modifying the ISIS ring to be able to withstand higher beam loss levels safely, by enhancing the collimator system. This would allow the particle rate to be increased by simply running at higher beam loss levels than are presently permissible. The ORBIT simulation work presented in Chapter 5 indicates the mostly likely form such an enhancement could take would be an extension of the existing collimator system shortly after beam injection into the ring, or an entirely new

collimator set positioned just downstream of the MICE target. Such a system would be a solid way of providing increased particle rates in the MICE beamline, but would likely come with a high cost both financially and in terms of implementation time.

6.3 Open Issues and Future Work

The are several productive ways in which the studies presented here could be extended. The high beam loss study of August 2010 should be repeated with a correct spill gate, better matched target short delay and a stable ISIS beam, in order to investigate conclusively the linearity of beam loss and particle rate up to the highest beam losses presently available. This should then allow a quantitative estimate of particle rate per spill for comparison with the June 2010 studies, and help to resolve the issue of the non-linear low beam loss behaviour reported in Section 4.3.4.

The target short delay study would benefit from being repeated in order to understand more fully the relationship between beam loss and dip time. The existing study reported in Section 4.1 produced some interesting features at early dip times, for which plausible solutions were presented, but which would still benefit from a repeat study. In particular it may be helpful to do a study with a greater number of data points to give a better resolution for the beam loss - dip time plots.

The change in particle rate across the spill produced by the target has been stated to be non-linear (as illustrated in Fig. 3.1). This prevents a straightforward comparisson of data taken with different size spill gates. It would be interesting to perform a study of how the particle rate varies with the position of the spill, and if possible to quantify it, so that data from different spill gates can be sensibly compared. If this does not prove

possible, it would also be useful to retake data for negative $\pi \rightarrow \mu$ optics using only a 1 ms spill gate (as opposed to the 3.2 ms gate used for the 15th June 2010 study), to allow an estimate of the number of muons (as opposed to anti-muons) available per spill to the full cooling channel.

The neutral particle contamination of the MICE beam, an issue touched on in Section 4.5, should be investigated thoroughly in order to understand what contribution this makes to the loss in efficiency when moving from scaler hits to reconstructed TOF tracks. Additionally work should also be done to understand more firmly the levels of pion contamination present in a beam optimised for $\pi \rightarrow \mu$ transport. This is presently hard to measure from data due to the CKOV detectors remaining in the commissioning phase, but useful insights might be gained from simulation.

Lastly it would be useful to perform simulations of the MICE beamline using the Simulation application in G4MICE in order to estimate how many of the muons present at TOF1, which have been measured in these studies, go on to traverse the whole cooling channel successfully and fall within the acceptance of TOF2. This could then be used to infer how beam loss is related to good muons specifically, which may be used in the ionisation cooling measurement.

Appendix A

Run Numbers

Sometimes I lie awake at night, and I ask, "Where have I gone wrong?"
Then a voice says to me, "This is going to take more than one night."

Charlie Brown (Charles M. Schultz)

6th November 2009:	1231, 1232, 1233, 1234, 1235, 1236
15th June 2010:	1985, 1986, 1987, 1988, 1989, 1991
16th June 2010:	2004, 2008, 2011, 2012, 2013, 2014, 2016
14th August 2010:	
Depth study:	2876, 2877, 2878, 2879, 2880, 2881, 2882, 2883, 2884, 2886
Delay study:	2888, 2889, 2890, 2891, 2893, 2894, 2895

Appendix B

Magnet Optics

B.1 6th November 2009 Study

Positive π optics.

Momentum at target = 300.0, at D1 = 296.4 , at D2 = 291.2 MeV/c.

Magnet	**Current (A)**
Q1	74.94
Q2	93.6
Q3	65.11
D1	225.25
DS	488.88
D2	115.35
Q4	195.51
Q5	262.2
Q6	173.89
Q7	176.7
Q8	267.45
Q9	228.49

Table B.1: Magnet currents for the November 2009 study.

B.2 June 2010 Studies

15th: Negative $\pi \rightarrow \mu$ optics, no Q3.

Momentum at target = 408.6, at D1 = 405.3, at D2 = 238.0 MeV/c.

16th: Positive $\pi \rightarrow \mu$ optics, no Q3.

Momentum at target = 408.6, at D1 = 405.3, at D2 = 238.0 MeV/c.

Magnet	**Current (A)**
Q1	126.96
Q2	111.11
Q3	0
D1	323.15
DS	668.63
D2	94.15
Q4	158.1
Q5	212.02
Q6	140.57
Q7	138.67
Q8	209.82
Q9	179.18

Table B.2: Magnet currents for the June 2010 studies.

B.3 14th August 2010 Study

Positive $\pi \rightarrow \mu$ optics.

Momentum at target = 408, at D1 = 405, at D2 = 237 MeV/c.

Magnet	**Current (A)**
Q1	102.38
Q2	127.91
Q3	89
D1	323.15
DS	668.63
D2	94.15
Q4	158.1
Q5	212.02
Q6	140.57
Q7	138.67
Q8	209.82
Q9	179.18

Table B.3: Magnet currents for the August 2010 study.

Bibliography

[1] C. L. Cowan *et al.* Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.

[2] G. Danby *et al.* Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9(1):36–44, 1962.

[3] J. K. Bienlein *et al.* Spark chamber study of high-energy neutrino interactions. *Phys. Lett.*, 13(1):80–86, 1964.

[4] M. L. Perl *et al.* Evidence for anomalous lepton production in e+e– annihilation. *Phys. Rev. Lett.*, 35(22):1489–1492, 1975.

[5] DONUT Collaboration. Observation of tau neutrino interactions. *Phys. Lett. B*, 504:218–224, 2001.

[6] M. Sorel, J. M. Conrad, and M. H. Shaevitz. Combined analysis of short-baseline neutrino experiments in the $(3 + 1)$ and $(3 + 2)$ sterile neutrino oscillation hypotheses. *Phys. Rev. D*, 70(7):073004, 2004.

[7] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109(3):1015–1017, 1958.

[8] K. Zuber. *Neutrino Physics.* IOP, 2004.

[9] R. Davis. Solar neutrinos. ii. experimental. *Phys. Rev. Lett.*, 12(11):303–305, 1964.

[10] B. T. Cleveland *et al.* Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496(1):505, 1998.

[11] R. Davis. A review of the Homestake solar neutrino experiment. *Prog. Part. Nucl. Phys.*, 32:13–32, 1994.

[12] The Super-Kamiokande collaboration. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[13] T. Kajita. Discovery of neutrino oscillations. *Rep. Prog. Phys.*, 69:1607–1635, 2006.

[14] J. Boger *et al.* The Sudbury Neutrino Observatory. *Nucl. Instr. and Meth. A*, 449(1-2):172–207, 2000. arXiv:nucl-ex/9910016v2.

[15] SNO collaboration. Direct evidence for neutrino flavor tranformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89, 2002. No. 1, 011301.

[16] Andreas Piepke. KamLAND: A reactor neutrino experiment testing the solar neutrino anomaly. *Nucl. Phys. B (Proc. Suppl.)*, 91(1-3):99–104, 2001.

[17] The KamLAND collaboration. Measurement of neutrino oscillation with KamLAND: evidence of spectral distortion. *Phys. Rev. Lett.*, 94, 2005. 081801.

[18] KATRIN collaboration. KATRIN letter of intent. 2001. arXiv:hep-ex/0109033v1.

[19] International Scoping Study Physics Working Group. Physics at a future Neutrino Factory and Superbeam facility. *Rep. Prog. Phys.*, 72:106201, 2009. arXiv:hep-ph/0710.4947v2.

[20] The NEMO Collaboration. Study of 2b-decay of Mo-100 and Se-82 using the NEMO3 detector. *JETPLett*, 80:377–381, 2004. arXiv:hep-ex/0410021v1.

[21] The International Design Study for the Neutrino Factory Collaboration. First progress report. IDS-NF-017. See `https://www.ids-nf.org/wiki/FrontPage/`.

[22] P. Adamson *et al.* First observations of separated atmospheric ν_μ and $\bar{\nu}_\mu$ events in the MINOS detector. *Phys. Rev. D*, 73:072002, 2006.

[23] M. H. Ahn *et al.* Measurement of neutrino oscillation by the K2K experiment. *Phys. Rev. D*, 74:072003, 2006.

[24] M. Apollonio *et al.* Initial results from the CHOOZ long baseline reactor neutrino oscillation experiment. *Phys. Lett. B*, 420(3-4):397–404, 1998.

[25] T. Schwetz, M. A. Tortola, and J. W. F. Valle. Three-flavour neutrino oscillation update. *New Journal of Physics*, page 113011, 2008.

[26] Y. Itow *et al.* The JHF-Kamioka neutrino project. 2001. Prepared for NOON2001. arXiv:hep-ex/0106019v1.

[27] NOνA Collaboration. NOνA proposal. 2005. NOVA-doc-593-v1. See `http://nova-docdb.fnal.gov/cgi-bin/ShowDocument?docid=593`.

[28] F. Ardellier *et al.* Double CHOOZ: A search for the neutrino mixing angle θ_{13}. 2006. arXiv:hep-ex/0606025v4.

[29] Daya Bay Collaboration. A precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya Bay. 2006. Proposal to DOE. arXiv:hep-ex/0701029v1.

[30] RENO Collaboration. RENO: An experiment for neutrino oscillation parameter θ_{13} using reactor neutrinos at Yonggwang. 2010. arXiv:1003.1391v1.

[31] The IDS-NF collaboration. International design study for the neutrino factory: Interim design report. 2011. IDS-NF-020. See `https://www.ids-nf.org/wiki/FrontPage/Documentation`.

[32] R. Garoby, E. Benedetto, M. Aiba, and M. Meddahi. Linac-based proton driver for a neutrino factory. CERN-NuFact-Note-157. See `http://cdsweb.cern.ch/collection/Neutrino%20Factory%20Notes`, 2009.

[33] See `http://lbne.fnal.gov/`.

[34] S. Geer. Muon Colliders and Neutrino Factories. *Annl. Rev. Nucl. Part. Sci.*, 59:345–367, 2009.

[35] The International Design Study for the Neutrino Factory Collaboration. Neutrino Factory: Specication of baseline. IDS-NF-002. See `https://www.ids-nf.org/wiki/FrontPage/`.

[36] See `https://www.ids-nf.org/wiki/FrontPage`.

[37] International Scoping Study Accelerator Working Group. Accelerator design concept for future neutrino facilities. *JINST*, 4:P07001, 2009. arXiv:0802.4023v2.

[38] R Garoby, E Benedetto, M Aiba, and M Meddahi. Linac-based proton driver for a Neutrino Factory. oai:cds.cern.ch:1227854. Technical Report CERN-NUFACT-Note-157, CERN, 2009.

[39] M. Popovic. Project X and a muon facility at Fermilab. In *AIP Conference Proceedings*, volume 1222, pages 336–338, 2010. See `http://link.aip.org/link/?APC/1222/336/1`.

[40] C. R. Prior. The synchrotron option for a multi-megawatt proton driver. *Nucl. Phys. B (Proc. Suppl.)*, 155(1):312–314, 2006. Proceedings of NuFact05.

[41] J. Roger *et al.* Studies of a target system for a 4-MW, 24-GeV proton beam. Technical Report CERN-INTC-2004-016, CERN, 2004.

[42] G. P. Skoro *et al.* Solid target for a Neutrino Factory. In *Proceedings of PAC09*, page 2878, 2009.

[43] C. J. Densham *et al.* The potential of fluidised powder target technology in high power accelerator facilities. In *Proceedings of PAC09*, page 1833, 2009.

[44] R. Palmer *et al.* A complete scheme of ionization cooling for a muon collider. In *Proceedings of PAC07*, page 3193, 2007. arXiv:0711.4275v1.

[45] A. N. Skrinsky and V. V. Parkhomchuk. Cooling methods for beams of charged particles. *Sov. J. Part. Nucl.*, 12:223 – 247, 1981.

[46] S. Ozaki *et al.* Feasibility study-II of a muon-based neutrino source. BNL-52623, 2001.

[47] P. Huber and W. Winter. Neutrino factories and the 'magic' baseline. *Phys. Rev. D*, 68:037301, 2003. arXiv:hep-ph/0301257v2.

[48] International Scoping Study Detector Working Group. Detectors and flux instrumentation for future neutrino facilities. *JINST*, 4(05):T05001, 2007. arXiv:0712.4129v1.

[49] D. M. Kaplan and K. Long. MICE: The International Muon Ionization Cooling Experiment. 2007. Contributed to LP07. arXiv:0707.1915v1.

[50] E. Wilson. *An Introduction to Particle Accelerators*. Oxford University Press, 2001.

[51] H. Wiedemann. *Particle Accelerator Physics, 3rd edition*. Springer, 2007.

[52] G. Penn. Beam envelope equations in a solenoidal field. NFMCC-doc-71-v1. Muon Collider Note 71, revision 1, 2000.

[53] MICE Collaboration. MICE Technical Design Report. See `http://mice.iit.edu/trd/MICE_Tech_ref.html`.

[54] C. Rogers. *Beam Dynamics in an Ionisation Cooling Channel*. PhD thesis, Imperial College, University of London.

[55] D. Neuffer. Principles and applications of muon cooling. *Part. Accel.*, 14:75, 1983.

[56] M. Ellis *et al.* The design, construction and performance of the MICE scintillating fibre trackers. arXiv:1005.3491v2, 2010.

[57] D. Li *et al.* A 201 MHz RF cavity design with non-stressed pre-curved Be windows for muon cooling channels. In *Proceedings of PAC03*, volume 2, pages 1243 – 1245, 2003.

[58] D. Stratakis, J. Gallardo, and R. Palmer. Effects of external magnetic fields on the operation of high-gradient accelerating structures. *Nucl. Instr. and Meth. A*, 620:147 – 154, 2010.

[59] D. Adams, R. Edgecock, I. Gardner, and K. Long. Pion measurements with the HEP test beam and implications for MICE. 2008. MICE Note 217. MICE-NOTE-BEAM-0217. See `http://hep04.phys.iit.edu/mice/notes/notes.html`.

[60] S. J. Payne *et al.* Beam diagnostics at ISIS. In *Proceedings of HB2008*, page 466, 2008.

[61] M. A. Clarke-Gaythe *et al.* Global beam loss monitoring using long ionisation chambers at ISIS. In *Proceedings of EPAC94*, page 1634, 1994.

[62] P. Hodgson, C. N. Booth, P. J. Smith, and J. Tarrant. MICE target hardware. In *Proceedings of IPAC10*, page 3488, 2010.

[63] Paul Smith. *The Development of the Target Mechanism for the Muon Ionisation Cooling Experiment (MICE).* PhD thesis, Department of Physics and Astronomy, Univeristy of Sheffield, 2010.

[64] P. Smith. Instructions for setting the T560 digital delay generator for the target delay. See `http://mice.iit.edu/mico/manuals/`.

[65] P. Soler. Luminosity monitor documentation. 2010. See `http://micewww.pp.rl.ac.uk/documents/7`.

[66] M.A. Rayner and M. Bonesini. The MICE PID detector system. In *Proceedings of IPAC10*, page 1164, 2010.

[67] L. Coney, A. Dobbs, and Y. Karadzhov. Particle production in the MICE beamline. In *Proceedings of IPAC10*, page 3530, 2010.

[68] V. Verguilov and V. Palladino. Particle identification devices in MICE. In *Proceedings of DIPAC09*, page 321, 2009.

[69] KLOE Collaboration. The KLOE electromagnetic calorimeter. *Nucl. Instr. and Meth. A*, 494:326–331, 2002.

[70] A. Fish. *Construction and testing of the scintillating fibre trackers for MICE*. PhD thesis, Imperial College, University of London, 2009.

[71] H. Sakamoto. Status of the MICE tracker system. In *Proceedings of IPAC10*, page 996, 2010.

[72] C. Rogers. G4MICE tracking and cooling performance. 2005. MICE Note 93 version 1.0. See `http://hep04.phys.iit.edu/mice/notes/notes.html`.

[73] A. Dobbs, K. Long, and A. Alekou. The MICE muon beamline and induced host accelerator beam loss. In *Proceedings of PAC11*, 2011. Extracts Copyright IEEE.

[74] C. Iselin. The MAD program. In *Computing in Accelerator Design and Operation*, volume 215 of *Lecture Notes in Physics*, pages 146–151. Springer Berlin / Heidelberg, 1984.

[75] J. D. Galambos *et al.* ORBIT user manual, version 1.10. Technical Report 011, SNS/ORNL/AP, 1999. See `http://neutrons.ornl.gov/APGroup/Codes/orbit.htm`.

[76] J. Galambos *et al.* ORBIT - a ring injection code with space charge. In *Proceedings of PAC99*, page 3143, 1999.

[77] J. A. Holmes. Recent enhancements to the ORBIT code. In *Proceedings of IPAC10*, page 1901, 2010.

[78] See `http://sct.esc.rl.ac.uk/index.html`.

[79] P. Hodgson, C. N. Booth, and P. J. Smith. MICE target operation and monitoring. In *Proceedings of IPAC10*, page 3485, 2010.

[80] D. Adams and M. Popovic. Vertical orbit bump perturbations at the MICE target. 2010. MICE Note 284. MICE-NOTE-BEAM-0284. See `http://hep04.phys.iit.edu/mice/notes/notes.html`.