

**The Development of the Target Mechanism for the Muon  
Ionisation Cooling Experiment (MICE)**

**Paul Jason Smith**

**Department of Physics and Astronomy**

**University of Sheffield**

**February 2010**

**Thesis Submitted for the Degree of**

**Doctor of Philosophy**

**in the Subject of Physics**



# Abstract

The Muon Ionisation Cooling Experiment (MICE) has been designed to practically demonstrate the principle of ‘Ionisation Cooling’ a theoretically sound but technologically unproven method of reducing the emittance of a muon beam. Ionisation Cooling is one technology that will be required to build a next generation high intensity neutrino source, the Neutrino Factory.

The source of these muons for the MICE experiment will come from the MICE ‘target mechanism’ an electro-mechanical device that operates parasitically on the ISIS accelerator. This is an 800 MeV proton accelerator that forms part of a neutron spallation source, situated at the Rutherford Appleton Laboratory in the UK, home of the MICE experiment. This target mechanism, operating at a maximum frequency of 1 Hz, inserts a small titanium target into the ISIS proton beam on demand. The target remains outside the beam envelope during acceleration and then overtakes the shrinking beam envelope to enter the proton beam during the last 2 ms before beam extraction. The target interacts with the ISIS beam halo during these 2 ms to produce pions for the MICE experiment. This is a demanding application, the target must accelerate at over 80  $g$ , whilst the components of the system must remain compatible with the ultra high vacuum of the ISIS system and be able to operate in a radioactive environment.

The first fully operational target mechanism was installed onto ISIS in January 2008. This thesis details both the mechanical development of the target mechanism and the electronic control systems that operate it. Finally, an analysis of the target’s performance on the ISIS beamline will be given.



*To John,*

*This thesis is dedicated to my late uncle and good friend John Alan Smith. John gave me both the original inspiration and the opportunity to study physics. It was with deep regret that John was not able to see the completion of my PhD.*



## Acknowledgements

It goes without saying that this thesis could not have been written in isolation and I am indebted to many people. It is not possible for me to thank them all, Chris has quite strictly told me that my thesis is already long enough! However there are a few people who undoubtedly deserve mentioning.

Firstly thanks to Chris for giving me the opportunity to study for a PhD and for his tireless effort in proof-reading this thesis. Thanks to Paul for his help, political education and for convincing me to write my thesis in  $\text{\LaTeX}$  as opposed to Microsoft Word. Neither can I forget Richard, the man who can make anything even if the drawing is scribbled on the back of a piece of tissue paper; but don't forget to double his time estimate and add a week!

Closer to home I have to thank Danni for her patience whilst I have been writing my thesis. I would also like to thank my children, Jamie and Ellie, both for their patience with me and for 'generally' being quiet whilst 'dad' has been spending all of their summer holiday on the computer. Last, but certainly not least, a thank you to my mum for providing me with a quiet room and regular cups of tea during the periods when the children were not so 'generally' being patient and quiet.



## Preface and Author's Declaration

When I first started this PhD in the summer of 2004, MICE was an established collaboration that was beginning the transition from the design to the build of the experiment. The design of several major components of the MICE system was close to being finalised and a significant amount of R&D, both practical and simulated, was taking place throughout the collaboration.

Design work for the target mechanism was underway, the operational requirements for the target mechanism were defined and the use of a linear motor as the mechanism of choice had been decided upon. The first stator, which would be significantly different to the final stator installed on ISIS, had been received at Sheffield University. Some ideas were being actively explored, for example the development of a pair of diaphragm springs to support the shaft and the use of a magnetic quadrature encoder to monitor its position. Further details of this early developmental work and the justification for using a linear motor can be found in chapter four.

My first task upon starting this PhD was to build the circuitry that was necessary to drive this prototype six coil stator. The first circuits that I built really set the stage for me during the rest of my PhD and beyond. The experience I gained by building the necessary functionality into these circuits was both continuously utilised and added to, enabling the target mechanism to work in an integrated and consistently reliable manner.

Much of the development of the target mechanism was only achieved because of the advice, experience and often practical help from a significant number of people. Design decisions were generally made as a team, which is of course the correct way to proceed, even if the advice of the expert heavily influenced the outcome! With this in mind there are several pieces of work that are described in this thesis where I feel that I made a 'significant contribution', and for clarity I will outline this work here.

Primarily I designed all of the control electronics and algorithms for the target system and by the time the first system was installed on ISIS I had written all of the firmware that resided on the target control microprocessors. I also did all of the simulation work that determined that the stator could be operated in two distinct modes giving a mechanism by which the target could be passively levitated - this is detailed in chapter seven. I designed and built the initial 10 A power supply for running the stator. Many features of its design were subsequently incorporated into the high current power supply that was designed by Steve Griffith's team at Daresbury Laboratories, including the suggestion that the power supply should work from a capacitor bank system - and the even later suggestion to redesign the capacitor bank so that the stator could run from a split supply. I also designed the temperature monitoring and interlock systems that we use at Sheffield, and designed the way that the control electronics would interface to the DAQ PC. I had no involvement with the subsequent software that was written for the DAQ PC though!

I also designed the optics and the optical quadrature system that is described in chapter six. I developed the initial idea for the prototype system that is described in this chapter, and proved that the system was workable. The final details of the optical vane layout were undertaken by Chris Booth and the method of manufacture of this component was suggested by Richard Nicholson. The final mounting system for the optical components, based upon the prototype, was designed by Martin Baldwin, a mechanical engineer at RAL.

Nigel Schofield should be credited for the original design of the target stators (both the prototype and the final) who at the time of writing is based in Electrical Engineering at Manchester University. Since this design was handed over to Sheffield, and along with Richard Nicholson, I have made a significant contribution to the further development of the coil stack, principally with respect to the design of the cooling system that used a cooling jacket and shims. The success of this heat exchanger proved instrumental in enabling the stator to run at 1 Hz at high current.

Throughout the design I have had a significant input into many of the decisions concerning the installation of the target mechanism at RAL, particularly with regards to the electrical and optical systems. However it would be remiss of me not to mention the efforts of Steve Griffith's team at Daresbury who have been instrumental in sorting out the details and providing the manpower to install these particular systems.

The original analysis of the November 2006 run in chapter nine was largely completed by Lara Howlett. The subsequent analysis of the data from the target mechanism that was installed on ISIS was done by myself although it would be unfair of me not to mention the large contribution made by Paul Hodgson in writing the DAQ software, the MICE Analysis Program (MAP) and the help and time that he gave showing me how to use ROOT. Finally the test rig for taking the stator field measurements in chapter ten was set up largely by Craig MacWaters and myself although in the end Craig took the actual data sets. The subsequent FEA was all my own work.

# Contents

<b>1</b>	<b>Neutrino Physics</b>	<b>1</b>
1.1	Introduction . . . . .	1
1.2	The Standard Model . . . . .	2
1.3	Neutrinos . . . . .	5
1.4	The Solar Neutrino Experiment . . . . .	6
1.5	Neutrino Oscillations . . . . .	7
1.5.1	Three Neutrino Oscillation . . . . .	8
1.5.2	Neutrino Oscillations in Matter . . . . .	11
1.5.3	Two Neutrino Approximation . . . . .	12
1.6	Measuring Neutrino Parameters . . . . .	14
1.7	Neutrino Experiments . . . . .	14
1.7.1	Solar Neutrino Experiments . . . . .	15
1.7.2	Current values of Solar Oscillation Parameters . . . . .	17
1.7.3	Atmospheric Neutrino Experiments . . . . .	19
1.7.4	Reactor Neutrino Experiments . . . . .	19
1.7.5	Accelerator Neutrino Experiments . . . . .	22
1.8	Conclusions . . . . .	25
<b>2</b>	<b>The Neutrino Factory</b>	<b>27</b>

2.1	Creating Neutrino Beams . . . . .	27
2.1.1	Super Beams . . . . .	28
2.1.2	Beta Beams . . . . .	28
2.2	Neutrino Factory . . . . .	29
2.3	Neutrino Factory Design . . . . .	30
2.3.1	Proton Source and Driver . . . . .	30
2.3.2	Target . . . . .	30
2.3.3	Muon Phase Rotation and Ionisation Cooling . . . . .	30
2.3.4	Acceleration . . . . .	31
2.3.5	Storage . . . . .	32
2.4	Physics Motivation . . . . .	32
2.5	Conclusion . . . . .	33
<b>3</b>	<b>The MICE Experiment</b>	<b>35</b>
3.1	Introduction . . . . .	35
3.2	Emittance . . . . .	36
3.2.1	Geometrical Interpretation of Emittance . . . . .	39
3.3	Emittance Reduction With Ionisation Cooling . . . . .	42
3.4	The ISIS accelerator . . . . .	43
3.5	The MICE Cooling Channel . . . . .	45
3.5.1	Diffuser . . . . .	45
3.5.2	Emittance Measurement and Particle ID in MICE . . . . .	47
3.5.3	The MICE magnets . . . . .	54
3.5.4	The MICE Beamline . . . . .	54
3.6	Conclusion . . . . .	55

<b>4</b>	<b>An Introduction to the MICE Target Mechanism</b>	<b>57</b>
4.1	Introduction . . . . .	57
4.2	Background . . . . .	57
4.3	MICE cooling Channel Requirements . . . . .	59
4.4	Target Operational Requirements . . . . .	60
4.5	Possible Designs for the Target Mechanism . . . . .	60
4.5.1	Resonant Flexure Mechanisms . . . . .	61
4.5.2	Rotary Mechanisms . . . . .	62
4.5.3	Linear Motor . . . . .	67
4.6	Choice of Design . . . . .	70
4.7	Earnshaw's Theorem and its Implications . . . . .	70
4.8	Bearings . . . . .	71
4.8.1	Linear Bearings . . . . .	72
4.8.2	Magnetic Bearings . . . . .	73
4.8.3	Sprung Bearings . . . . .	73
4.8.4	Sliding Bearings . . . . .	74
4.9	Motor Development . . . . .	74
4.10	Test Stand for Development . . . . .	76
4.11	Design Requirements of Sliding Bearings . . . . .	79
4.12	Conclusion . . . . .	81
<b>5</b>	<b>Mechanical Design</b>	<b>83</b>
5.1	Target Mechanism Design . . . . .	83
5.1.1	Construction . . . . .	83
5.1.2	Function . . . . .	86
5.2	Stator Body . . . . .	87

5.2.1	Stator Bore . . . . .	87
5.2.2	Coils . . . . .	88
5.2.3	Shims . . . . .	88
5.2.4	Assembly of the Coil Stack . . . . .	89
5.3	Cooling . . . . .	90
5.3.1	Temperature Monitoring . . . . .	92
5.3.2	Cooling Jacket Design . . . . .	92
5.4	Flanges . . . . .	94
5.5	The Target Shaft . . . . .	94
5.5.1	Mechanical Considerations . . . . .	95
5.5.2	Description of the Shaft Design . . . . .	95
5.5.3	Target Heating . . . . .	97
5.5.4	Diamond Like Carbon . . . . .	98
5.5.5	Magnet Assembly . . . . .	99
5.5.6	Optical Vane . . . . .	100
5.6	Bearings . . . . .	102
5.6.1	Bottom Bearing . . . . .	102
5.6.2	Top Bearing . . . . .	103
5.7	Mechanical Support and the Beam-Pipe Interface . . . . .	103
5.7.1	A-Frame . . . . .	104
5.7.2	Beam-pipe . . . . .	105
5.7.3	Bellows and the Gate Valve . . . . .	106
5.8	Demonstrator System . . . . .	107
5.9	Conclusion . . . . .	107

<b>6</b>	<b>The Optical Positioning System</b>	<b>109</b>
6.1	Position Monitoring . . . . .	109
6.2	Description of a Quadrature System. . . . .	110
6.3	Overview of the Quadrature System . . . . .	112
6.3.1	Theory . . . . .	113
6.4	Details of the Optical Design . . . . .	114
6.4.1	Laser Source . . . . .	115
6.4.2	Optical Fibres . . . . .	115
6.4.3	Collimators and Lenses . . . . .	116
6.4.4	Optical Vane . . . . .	117
6.4.5	Mechanical Mount - The Optics Block . . . . .	118
6.4.6	First Prototype Mount . . . . .	118
6.4.7	Final Mount . . . . .	120
6.4.8	Optical Sensors . . . . .	121
6.5	Fibre Installation Issues . . . . .	122
6.6	Conclusions and Future Developments . . . . .	123
<b>7</b>	<b>Stator Operation and the Power Electronics</b>	<b>125</b>
7.1	Phased Motors . . . . .	125
7.2	Three Phase ‘Star’ Wiring . . . . .	126
7.3	Coil Current Switching Sequence . . . . .	129
7.4	Magnetic Assembly . . . . .	129
7.5	Stator Simulation . . . . .	131
7.6	Zero-Force Points . . . . .	132
7.6.1	Moving the Zero-Force Points . . . . .	133
7.7	Position Monitoring . . . . .	134

7.8	Actuation . . . . .	135
7.9	Coil Switching and Current Control . . . . .	139
7.9.1	Switching Mechanism . . . . .	139
7.9.2	Current Control . . . . .	141
7.9.3	Controlling the Current with Pulse Width Modulation . . . . .	142
7.10	System Placement in ISIS . . . . .	144
7.11	Fibre Optic Links . . . . .	146
7.12	10 Ampere Hex Bridge Circuit . . . . .	146
7.12.1	Performance of the 10 Ampere circuit . . . . .	147
7.12.2	Correlation between Current and Acceleration . . . . .	148
7.13	100 Ampere Power Supply . . . . .	149
7.13.1	Description of the 100 Ampere Supply . . . . .	150
7.13.2	Performance of the 100 Amp Supply . . . . .	150
7.14	Over Temperature Protection . . . . .	152
7.15	Conclusion . . . . .	154
<b>8</b>	<b>Target Control, Electronics and the DAQ</b>	<b>155</b>
8.1	Target Controller Overview . . . . .	155
8.2	Control: Park, Hold and Actuate Enable Modes . . . . .	157
8.3	Control: Target Actuation and Capture . . . . .	159
8.3.1	Issues . . . . .	161
8.4	The ISIS Trigger Signal . . . . .	163
8.5	The Electronic Control System . . . . .	164
8.5.1	PIC Micro-Processors . . . . .	166
8.5.2	Power Supply . . . . .	167
8.5.3	Temperature Monitoring Module . . . . .	167

8.5.4	Quadrature Optical Amplifier . . . . .	167
8.5.5	Quadrature Counter . . . . .	168
8.5.6	24 Bit Timer . . . . .	169
8.5.7	Switch Panel . . . . .	170
8.5.8	Controller . . . . .	171
8.6	DAQ and the Software Monitoring . . . . .	172
8.6.1	Hardware . . . . .	173
8.6.2	Software . . . . .	174
8.6.3	Data Acquisition Code . . . . .	175
8.6.4	Plotting and Reduction of the DAQ data . . . . .	175
8.6.5	Conclusion and The Future of the Target Control System . . . . .	177
<b>9</b>	<b>Operational Performance and Analysis</b>	<b>179</b>
9.1	ISIS . . . . .	179
9.2	Experimental Setup for Nov 2006 . . . . .	180
9.2.1	Target . . . . .	180
9.2.2	Detectors . . . . .	181
9.2.3	DAQ . . . . .	181
9.3	Experimental Run 2006 . . . . .	181
9.4	Analysis of 2006 Data . . . . .	182
9.4.1	Sheffield Analysis . . . . .	183
9.4.2	Glasgow Analysis . . . . .	185
9.5	Experimental Setup for 2008 . . . . .	188
9.6	Running of Primary System at RAL . . . . .	189
9.7	Analysis . . . . .	190
9.7.1	Mechanical Performance . . . . .	191

9.7.2	Mechanical Stability . . . . .	192
9.7.3	Actuation Depth . . . . .	193
9.7.4	Acceleration . . . . .	194
9.7.5	Time Jitter . . . . .	199
9.7.6	Capture Reliability . . . . .	202
9.7.7	Operational Temperature of the Stator . . . . .	202
9.8	Particle Production . . . . .	204
9.8.1	ISIS Beam Loss and Actuation Depth . . . . .	205
9.8.2	Establishing Parasitic Operation . . . . .	213
9.8.3	Pion Production for MICE . . . . .	213
9.9	Conclusion . . . . .	215
<b>10</b>	<b>Further Developments</b>	<b>217</b>
10.1	Introduction . . . . .	217
10.2	Coil Insulation . . . . .	217
10.3	Bearing Development . . . . .	222
10.4	Bearing Wear . . . . .	223
10.5	Analysis of the Target Mechanism Failure . . . . .	224
10.5.1	DLC Coating . . . . .	226
10.5.2	Magnetic Survey and FEA . . . . .	227
10.5.3	Mechanical Survey . . . . .	232
10.6	Target Mechanism Modifications . . . . .	235
10.6.1	Target Shaft . . . . .	235
10.6.2	Target Bearings . . . . .	236
10.6.3	Stator Body . . . . .	237
10.7	Conclusion . . . . .	237

*CONTENTS*

ix

**11 Conclusion** **239**

11.1 Outlook . . . . . 240



# List of Figures

1.1	Diagram showing the correlation between the flavour and the mass bases [18]. . . . .	9
1.2	The two possible neutrino mass hierarchies. The $\nu_e$ fraction of each mass eigenstate is crosshatched, the $\nu_\mu$ fraction is indicated by right-leaning hatching, and the $\nu_\tau$ fraction by left-leaning hatching ( $\Delta m_{12}^2 = \Delta m_{\odot}^2$ & $\Delta m_{13}^2 \approx \Delta m_{23}^2 = \Delta m_{atm}^2$ ) [12]. . . . .	12
1.3	The solar neutrino spectrum [22]. . . . .	15
1.4	Neutrino-oscillation contours. (a) SNO only: $D_2O$ & salt day and night spectra, neutral current detection (NCD) phase fluxes. The best-fit point is $\Delta m^2 = 4.57 \times 10^{-5} \text{ eV}^2$ , with $\tan^2 \theta = 0.447$ , (b) Solar Global: SNO, SK, Cl, Ga, Borexino. (c) Solar Global + KamLAND. The best-fit point is $\Delta m^2 = 7.94 \times 10^{-5} \text{ eV}^2$ , $\tan^2 \theta = 0.447$ [27]. . . . .	18
1.5	The zenith angle distribution for fully-contained 1-ring events in the Super-K detector. The points show the data, box histograms show the non-oscillated Monte Carlo events and the lines show the best-fit expectations for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with $\sin^2 2\theta = 1.00$ and $\Delta m^2 = 2.1 \times 10^{-3} \text{ eV}^2$ [30]. . . . .	20
1.6	The upper bound on $\sin^2 \theta_{13}$ as a function of $\Delta m_{13}^2$ from the Chooz experiment. . . . .	21
1.7	The ratio of measured to expected $\bar{\nu}_e$ flux from reactor experiments. The solid dot is the KamLAND data point plotted at a flux-weighted average distance (the dot size is indicative of the spread in reactor distances). The shaded region indicates the range of flux predictions corresponding to the 95 % C.L. LMA region found in a global analysis of the solar neutrino data. The dotted curve corresponds to $\sin^2 2\theta = 0.833$ and $\Delta m^2 = 5.5 \times 10^{-5} \text{ eV}^2$ [36] and shows the best-fit LMA predictions while the dashed curve illustrates the no oscillation case [35]. . . . .	23

1.8	Allowed regions at 68 %, 90 %, 99 % Confidence Level in the $\Delta m^2, \sin^2 2\theta$ plane from a fit to the far detector reconstructed energy spectrum for the MINOS experiment. The best-fit point, which occurs at $\sin^2 2\theta = 1$ and $\Delta m^2 = 2.74 \times 10^3 \text{ eV}^2/c^4$ is represented by the star [38]. . . . .	24
1.9	Oscillation Parameter Plot showing the current experimental limits on the neutrino oscillation parameters [47]. . . . .	26
2.1	Current conceptual layout of the neutrino factory [56]. . . . .	29
2.2	A pictorial representation of how phase rotation will alter the beam dynamics by the application of suitable RF fields. . . . .	31
3.1	Particle trajectories $x(s)$ within the envelope $E(s)$ of the beam. The upper figure shows a single particle trajectory whilst the lower figure shows multiple trajectories. The beam envelope is defined from the outermost trajectory $E(s)$ at each point along the $s$ axis [64]. . . . .	37
3.2	The phase space ellipse of particle motion in the 2D $x - x'$ plane [64]. . . . .	38
3.3	An overview of the ISIS accelerator facility [66]. . . . .	44
3.4	A 3D drawing of the MICE cooling channel showing the position of the magnets, trackers, absorbers and RF cavities. The diffuser and TOF1 are at the front of the MICE machine whilst the calorimeter and TOF2 are at the rear. The white arrows represent the location of the coupling/focusing coils [70]. . . . .	46
3.5	The MICE diffuser. A carousel system installs lead discs of different thicknesses into the MICE beamline causing an increase in the beam's emittance. This allows MICE to measure the amount of emittance reduction that can be achieved on beams with different initial emittances [70]. . . . .	47
3.6	Diagram of one of the 8 bars per plane TOF stations including some of the mechanical supports [65]. . . . .	48
3.7	Cross-section through one of the two Cherenkov detectors. Note the aerogel box on the front of the detector [65]. . . . .	49
3.8	the whole Cherenkov1 (CKOV1) device is constructed by putting two near identical detection units one after the other along the beam line upstream of MICE. Each unit contains aerogel with a different refractive index [65]. . . . .	49

3.9	Particle Momentum vs Photoelectron density for the two Cherenkov detectors with different aerogel. The different triggering behaviour for the Cherenkov detectors is indicated by the regions I, II and III [65]. . . . .	50
3.10	A schematic view of one of the MICE scintillating fibre trackers. The trackers are inserted into the long solenoidal magnets with a field of $\approx 4$ T [65]. . . . .	52
3.11	A cross-sectional view of the MICE absorber focus coil module illustrating the liquid absorber [73]. . . . .	53
3.12	Four RF cavities form one RF assembly. There will be two such assemblies in MICE sandwiched between three absorbers [65]. . . . .	53
3.13	The MICE Beamline . . . . .	55
3.14	The six stages of the MICE experiment [65]. . . . .	56
4.1	Schematic diagram of the old HEP target. The target was suspended from two arms, the first arm was oscillated at 50 Hz and the second arm was oscillated at 150 Hz. . . . .	58
4.2	Simulated target insertion. Superimposed 50 Hz and 150 Hz oscillations allow the target to dip into the ISIS beam at 50 Hz . . . . .	58
4.3	The trajectory of the target mechanism (blue) must be capable of accelerating past the shrinking beam (red) so that it only intercepts the beam during the last 2 ms of the acceleration cycle. The target must be clear of the beam envelope by the start of the next ISIS cycle and must remain clear of the beam until the next actuation $\approx 1$ s later. . . . .	61
4.4	Diagram illustrating the use of a continuously rotating arm to intercept the ISIS beam. The case for intercepting the beam with a rotation that is perpendicular to the beam's direction (b) is much better than for the case where the arm rotates in the direction parallel to the beam (a). Both cases result in a long arm due to the small value of the angle $\theta$ and creates a moving production point for particles. . . . .	62
4.5	The minimum arm radius for a constant speed rotary target drive as a function of dip depth for three different insertion times. The motion of the arm is parallel to the beam pipe. One can see that this leads to a totally impractical solution. . . . .	63
4.6	If $r \gg R$ then the arm will to a first approximation travel in a straight line through the beam. The path is indicated by $x$ . . . . .	64

4.7	The minimum arm radius ( $r$ ) for a constant speed rotary target drive as a function of dip depth ( $d$ ) for three different insertion times. The radius of the ISIS beam for these plots ( $R$ ) is 40 mm. The motion of the arm is transverse to the beam. The solution here looks much better than for the previous case; for a 2 ms insertion time the required radius of the arm is of the order of $\approx 1.5$ m at the estimated required insertion depth. . . . .	65
4.8	The required rotation angle $\theta$ changes as a function of radius for a given insertion depth ( $d$ ) for a transverse rotating arm. Note that the formula is only valid if $r \gg R$ therefore this is only used for $r \geq 0.4$ m. The angle $\theta$ should be compared to a $\theta$ of 0.012 rad in the case of a constant rotational velocity of 1 Hz for an insertion time of 2 ms. . . . .	67
4.9	The estimated acceleration required to move the target through a given actuation depth $d$ in time $t_A$ using a series of quadratics defined by a constant magnitude driving force. . . . .	69
4.10	Example of FEA results for the two shuttle design topologies considered showing coil current excitation and the direction of magnetisation. The radial field topology was shown to be superior. Diagram taken from Schofield's paper[90]. . . . .	75
4.11	The diaphragm springs were designed to provide a mechanism of supporting the shaft for a travel of $\pm 20$ mm. . . . .	76
4.12	The completed test assembly. This test rig was produced for testing the diaphragm springs and the quadrature read head. . . . .	78
4.13	A simulation of how the magnetic field changes radially for a current of 7.5 Amps at a coaxial point that corresponded to the peak of the magnetic field. (The magnetic field changes sinusoidally coaxially) The vertical lines represent the position of the stator bore tube. It is clear from this diagram that if the magnets are not positioned coaxially to the bore tube, and therefore the coil stack, that there will be a resultant lateral force on the magnetic assembly. . . . .	80
5.1	A close up drawing of the stator itself. Here the details of the coil stack within the stator can be seen. . . . .	84
5.2	A cross sectional view of the target mechanism including the full length of the target shaft and the target. The black plane protruding from the stator block on the right hand side of the diagram represents the wires protruding from the stator coils. . . . .	85

5.3	A complete stator coil and some copper heatsinking shims. Two kapton shims and four copper shims are placed between each pair of coils to provide electrical insulation and heatsinking capability. . . . .	89
5.4	Bottom spacer, copper and Kapton shim and the first coil. . . . .	90
5.5	All 24 coils assembled on a former. The copper shims can be clearly seen protruding from the coil stack where they will later make contact with the cooling jacket. . . . .	91
5.6	The copper cooling jacket is over-fitted with the water cooling pipes. The cooling pipes are soldered onto the jacket. . . . .	93
5.7	The cooling jacket has been placed over the coil stack. . . . .	93
5.8	Drawing of the target shaft . . . . .	96
5.9	The temperature of the target (first 35 mm) as a function of distance up the target shaft for normal (bottom line) and failure (top line) conditions. The graph shows that the target temperature doesn't become excessively hot under both normal and failure conditions and that there is minimal conduction of heat up to the bearings and the magnets[96]. . . . .	98
5.10	Drawing showing the structure of the magnet assembly . . . . .	99
5.11	Diagram showing the flux lines from the permanent magnets as simulated by MAXWELL 2D Student Version [98]. . . . .	100
5.12	Diagram showing the BH curve for the magnetic material that is used in the magnet assembly. The blue curve represents the magnetisation, whereas the red curve represents the flux density at the surface of the magnets. Note the temperature dependence of the maximum demagnetisation field that can be tolerated[99]. . . . .	101
5.13	Drawing of the Optical Vane. One side of the vane has 157 teeth, these are used to provide the quadrature signal, whilst the other side has just a single tooth that is used to provide a fixed reference point. Dimensions in mm. . . . .	102
5.14	The bottom bearing is constructed in quadrants. Note the cruciform cross section that corresponds to the cross section of the shaft and prevents the shaft from rotating. . . . .	102
5.15	The top bearing is constructed in two pieces. . . . .	103
5.16	A drawing of the A-frame, small frame, target, bellows and the gate valve. The small frame sits within the A-frame and can be easily removed with a crane. RAL drawing [100]. . . . .	104

5.17	A close-up of the ‘Small Frame’ that sits inside the A- Frame. The Small Frame makes the attachment to the target’s lower flange as illustrated [100]. . . . .	105
5.18	Drawing of the beam-pipe section to which the target is attached[100]. .	106
6.1	How a quadrature system can be utilised to determine position. The top figure illustrates how the quadrature system responds to movement in the positive direction. The middle figure shows how the system responds to movement in the negative direction. The bottom figure shows how the quadrature system can be utilised to keep track of the position when the direction of motion changes. . . . .	111
6.2	A cartoon illustrating the key components to the optical quadrature system used to ascertain the position of the target shaft. . . . .	112
6.3	Diagram showing how the collimator and lens system works magnifying the optical fibre source to create a finite spot size on the optical vane. .	113
6.4	The attenuation of the SM600 single mode fibre as a function of frequency [101]. . . . .	116
6.5	The construction of the multimode fibre. The silica core gives more radiation resistance than polymer cored fibres [102]. . . . .	117
6.6	The attenuation of the BFH37/200 fibre as a function of frequency [102].	117
6.7	Diagram illustrating a cross section through a collimator and a lens cap. The collimators are individually aligned at the factory and are precision pieces. . . . .	118
6.8	Left - test bench for a single channel. Right - prototype optical block before addition of the retaining grub-screws. Pairs of collimators are inserted into the horizontal bores, one either side of the vertical bore. .	119
6.9	The glass tube on top of the stator can be seen in this photograph with a prototype optics block loaded with collimators on the left. . . . .	120
6.10	An exploded view of the optical mount. The optical vane sits inside the main body where the collimators and lens caps are able to focus the laser light onto the vane. . . . .	121
6.11	A photo showing the laser focused onto the optical vane through one of the sapphire windows. It is clear from this diagram how small the spot size is. (The distance between fingers on the comb is 300 $\mu\text{m}$ .) . . . . .	122

- 7.1 Diagram illustrating the phase relationship of the coils to their physical position within the stator. Coils marked with a prime are wired so that the current flows through them in the opposite direction to those coils that are unprimed. . . . . 127
- 7.2 This diagram represents how the coils are wired in the stator, illustrating the star point connection and the three external connections from the stator. Each coil in the diagram represents 4 coils i.e. The A coil represents the A coil from banks 1 to 4. . . . . 128
- 7.3 By switching the coils through the circular sequence given by States 1 to 6 this gives an apparent motion in the coil switching upwards through the stator as is illustrated here. Reversing the sequence gives an apparent downwards motion. The pattern repeats itself every six coils. The permanent magnets on the target shaft can be made to interact with the fields produced by passing sequentially through these states to produce motion of the target shaft. . . . . 130
- 7.4 Graph illustrating a simulation of how the force on the magnetic assembly varies as a function of its position within the stator when the stator is left in one of its six states with a coil current of 58 A. A good sinusoidal fit can be made to these points as illustrated. . . . . 132
- 7.5 If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 5’ then the magnets would see a restoring force (black arrow) that would move the shaft to the new zero force position (red arrow) 3 mm further up the stator. By repeatedly incrementing or decrementing the states the shaft can be moved up and down the stator body. This graph shows the force for a 3 A coil current typical of that used to levitate the shaft. . . . . 134
- 7.6 If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 6’(top) or ‘State 2’ (bottom) then the magnets would see a restoring force (black arrow) that would want to move the shaft to the new zero force position. During the first 3 mm of movement the shaft would see an average force that is  $\approx 96\%$  of the peak force. . . . . 136
- 7.7 The shaft moves from levitation (3 A coil current) to acceleration (58 A coil current). By tracking the position of the magnets the states can be switched to ensure the shaft continually accelerates using the maximum available force. This results in a hard acceleration of the target shaft. . . . . 137
- 7.8 An actuation. The shaft is accelerated downwards, decelerated, accelerated upwards, decelerated and then captured. The position of the shaft must be tracked accurately so the controller knows into which state to put the coils to maximise the accelerating/decelerating forces. . . . . 138

- 7.9 The Hex Bridge. This circuit allows bidirectional current control through any two out of the three phases of the stator by application of the appropriate control signals to the six inputs that switch the transistors on. . . . . 140
- 7.10 Top: Current path for ‘State 1’. If the PWM signal is imposed onto both the high-side and low-side transistors this causes the induction current to pass through the power supply (middle). This can stress the power supply and creates a lot of EMI. If the PWM signal is only supplied to the low-side transistor then this improves the situation as the induction current is now localised within the driving circuit (bottom). . . . . 145
- 7.11 A target actuation at 10 A using the 10 A Hex-Bridge. At this time capture was done based upon the target position. This proved to be unreliable, and was later changed so that capture occurred at the point where the target trajectory changed direction at the top of the pulse. Some damped motion of the target is expected at capture as mode change does not necessarily happen at a zero force point and the target may have some residual velocity. . . . . 148
- 7.12 Plots of some of the first actuations at higher current. This plot represents the complete data set. At this time no attempt was made to capture the target. This was done to obtain some indication as to whether the required acceleration would be achieved with the 100 A power supply. . . . . 152
- 7.13 A plot showing the acceleration achieved using the trajectories shown in figure 7.12. A straight line can be fitted to this confirming the that the accelerating force is proportional to the current as expected. When an extrapolation is done to the lighter cross shaped target then this shows that a coil current of about 50 A to 60 A would be required to accelerate the target to 80 *g*. This was confirmed experimentally several months later in the lab. . . . . 153
- 8.1 This figure illustrates the major components of the target mechanism installation at ISIS and how these components relate to each other. The power supply rack and the PC that controls it were built at Daresbury Laboratory. The target assembly, the DAQ PC and the control electronics were built at Sheffield University. The target frame was designed and built at the Rutherford Appleton Laboratory. . . . . 156
- 8.2 Diagram illustrating the four states of the target system. . . . . 157

8.3	The actuation trajectory is split into four quadrants from the controllers point of view. Switch Point 1 reverses the coil currents and so determines the actuation depth. Switch Point 2 determines the capture point. These two points are not at the same position due to the decay on the capacitor bank. . . . .	160
8.4	Diagram illustrating the target triggering system. At present the MICE synchronisation is not utilised and the ‘target ready’ output is looped back into the ‘MICE ready’ input. . . . .	165
8.5	The laser light from the optical system is converted into a square wave digital signal. There are three of these amplifiers, one for each channel of the optical system. The outputs from these boards are sent to the quadrature counter crate where the signals are counted to give the absolute position of the target shaft. . . . .	168
8.6	The quadrature counter crate takes the quadrature signals and converts them into an absolute position. This position, a 9-bit number, is then output to the controller and the DAQ system. A digital to analogue converter also gives an analogue output that is used to display the target position on an oscilloscope. . . . .	170
8.7	The ‘24 Bit Timer’ crate provides timing information on every position change that occurs during actuation. This information is synchronised with the quadrature counter and fed into the DAQ. . . . .	171
8.8	The controller is the heart of the system as it directly interfaces to and controls the motion of the target shaft within the stator. . . . .	172
8.9	A screenshot from the MAP program showing target data from a single actuation being plotted. The red line represents the cleaned target trajectory as recorded in the data stream. The green line is a third order polynomial fit. The cyan line shows where there was beam in the synchrotron. (The value goes high.) The pink line shows the recorded total beam loss signal. The light blue line is a fit to the beam loss data. Computed parameters are shown displayed on the right hand side of the plot area. The data rate is proportional to the target velocity. This produces blind spots in the data. . . . .	176
9.1	Evolution of momentum (red line) and energy (blue line) during an ISIS acceleration cycle [56]. . . . .	180
9.2	Target Test Setup in November 2006 . . . . .	182

9.3	An averaged beam loss plot as a function of time for the actuations starting 14.5 ms before injection. This time setting corresponds to the deepest penetration achieved into the ISIS beam during the 2006 tests [109]. . . . .	184
9.4	A plot showing the beam loss as a function of target position for an individual time bin, along with the exponential fit. This plots shows the time corresponding to 1 ms after injection. A sequence of these plots was created giving the time evolution of the beam loss for different target trajectories [109]. . . . .	185
9.5	Total beam loss as a function of position and time as calculated using exponential fits. This graph gives an indication of the position of the ISIS beam edge with respect to time. The horizontal black line represent the maximum depth of the stationary target trajectory and the curved black line shows the earliest actuated target trajectory used to obtain the data. This makes it clear what is interpolation and what is extrapolation. The cut-off at 0.4 V represents the maximum recorded beam loss reading in the data [109]. . . . .	186
9.6	A figure showing how a 30 ms 80 g quadratic target trajectory can be made to fit into the ISIS spill structure (50 Hz) to intercept the beam during the last 2 ms of a spill. The target must not touch the beam during the next injection [109]. . . . .	187
9.7	Diagram illustrating the position of counters GVA1 and GVA2 in the MICE beamline [112]. . . . .	190
9.8	A histogram of the actuation depth for the stator running in the laboratory over five days at 1 Hz. . . . .	195
9.9	A histogram of the actuation depth for the stator running on ISIS. This plot shows $\approx 15$ hours running at 0.4 Hz. . . . .	195
9.10	The time evolution of the actuation depth for the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points). . . . .	196
9.11	The time evolution of the actuation depth for the stator running on ISIS. This plot shows $\approx 15$ hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points). . . . .	196
9.12	A histogram of the acceleration achieved for the stator running in the laboratory over five days at 1 Hz. . . . .	197
9.13	A histogram of the acceleration achieved for the stator running on ISIS. This plot shows $\approx 15$ hours running at 0.4 Hz. . . . .	197

- 9.14 The time evolution in the acceleration for the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points). . . . . 198
- 9.15 The time evolution of the acceleration for the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points). . . . . 198
- 9.16 A histogram of the time it takes for the shaft to reach its minimum position. This data was taken with the stator running in the laboratory over five days at 1 Hz. . . . . 200
- 9.17 A histogram of the time it takes for the shaft to reach its minimum position. This data was taken with the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. . . . . 200
- 9.18 The time evolution of the time to minimum dip depth. This plot shows data from the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points). . . . . 201
- 9.19 The time evolution of the time to minimum dip depth. This data was taken with the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points). . . . . 201
- 9.20 This temperature profile shows the coil temperature of over 1 hours running from power up to power down. Note how the temperature monitored directly on the coil (blue - Coil 7) does not differ significantly from the measured on the other side of the kapton shim (pink - Kapton 7). The microstructure seen as noise during actuation is caused by the stator cooling between actuations. . . . . 203
- 9.21 The plot shows how the acceleration of the target shaft changes as the stator warms up during the first couple of minutes of operation. The current through the coils is linear with temperature and the acceleration is directly proportional to the current through the stator coils. The heating of the stator will be exponential. This means that the acceleration should change exponentially with time and is confirmed by the fit. The graph is for a stator with kapton insulated coils running at 1 Hz. . . . . 205
- 9.22 This plot shows the relative number of actuations done at each actuation depth during one continuous run of data taking. The actuation depth in this plot is the distance that the target actuates to measured from the centre of the beam-pipe. The individual populations were then individually extracted to estimate the error. This represents about half the data used to produce the final beamloss/actuation depth relationship. . . . . 207

9.23	The first four individual entries in the previous plot (figure 9.22) have been extracted to allow an estimation of the positional error. . . . .	208
9.24	The last four individual entries in the previous plot (figure 9.22) have been extracted to allow an estimation of the positional error. . . . .	209
9.25	Histogram of the beam loss distribution for the first four actuation depths (with respect to the centre line of the beam-pipe) for the data set shown in figure 9.22. . . . .	210
9.26	Histogram of the beam loss distribution for the last four actuation depths (with respect to the centre line of the beam-pipe) for the data set shown in figure 9.22. . . . .	211
9.27	This plot shows the relationship between the beamloss and target's actuation depth. In this plot the second data set has also been added. Once again the distance scale represents the distance from the centre of the beam-pipe. An exponential fit through the data points has been shown. . . . .	212
9.28	Rates of proton counts in scintillation counters GVA1 (circles) and GVA2 (triangles) placed in the MICE beamline, plotted as a function of target delay, $\Delta t$ , for a constant dip depth after the optimum insertion time. The one standard deviation background rate bands are shown for GVA1 (upper band), and GVA2 (lower band). . . . .	214
10.1	A photograph of the coil stack clearly showing the location of the coil failure. . . . .	218
10.2	The current surge had caused the coil to weld itself to the heatsinking shims. Approximately half a turn of coil was vaporised from the current pulse. . . . .	219
10.3	On the second electrical failure the coil insulation failed to the centre bore tube causing considerable damage; the current pulse had embedded copper from the coil into the bore tube. . . . .	220
10.4	The new power-supply design [117] nearly halves the voltage to earth helping to reduce the risk of a coil failure to earth. The highest operational voltage that the coils should see to earth when the new power supply is used is $\approx 125$ V, plus a 30 V spike on switching. All the coils are tested to 1 kV before being used in a stator, which gives a significant safety margin. . . . .	222
10.5	Upon opening up the mechanism there was significant particulate debris on the top bearing and the DLC coating had been worn through on the shaft. . . . .	224

- 10.6 There is considerable damage to the bearing faces on the upper bearing. Although it isn't clear in the photograph, a lot of the DLC has been removed in the damaged areas. . . . . 225
- 10.7 Much of the DLC has become detached from the surface of the cruciform cross section where it had passed through the bottom bearing. . . . . 225
- 10.8 Bottom bearing quadrants: Of the four bearing faces that can be clearly seen in this photograph, three of them are badly damaged. The other four faces (not possible to see clearly due to the angle of the photo) showed similar damage. . . . . 226
- 10.9 The voltage output from the Hall probe changed as a function of axial position within the stator bore in an azimuthal progression of  $30^\circ$  over a full  $360^\circ$ . The coils were wired in 'State 1' and the coil current was 7200 mA. 1 mV is equivalent to 10 mT. . . . . 228
- 10.10 The voltage output from the Hall probe changed as a function of axial position within the stator bore in an azimuthal progression of  $30^\circ$  over a full  $360^\circ$ . The coils were wired in 'State 2' and the coil current was 7200 mA. 1 mV is equivalent to 10 mT. . . . . 229
- 10.11 The peaks in figure 10.9 are numbered from 1 to 9 (P1 - P9) and the peak value plotted as a function of azimuthal angle. In the case of a perfect field with no measurement error one would expect to see that each set of points would lie on a straight horizontal line. The average value is taken from points P2 - P8 for each azimuthal angle. . . . . 230
- 10.12 The peaks in figure 10.10 are numbered from 1 to 10 (P1 - P10) and the peak value plotted as a function of azimuthal angle. The step in the data is suggestive of a systematic error in the data taking. . . . . 230
- 10.13 The measured data corresponds very well to the modelled data at a radial distance of 6 - 7 mm from the central axis of the stator bore. This distance corresponds to the radial distance that the Hall probe was at when the axial data was taken. Note in the legend that 'AB data' means the 'State 1' data, the other two data sets are produced from FEA. . . . 231
- 10.14 A simulation of the change in radial field strength across the bore of the stator at a position that corresponds to a peak in the axial field measurements based on experimental data. The two vertical lines represent the position of the inner radius of the stator bore. . . . . 232

- 10.15 The standard deviation of the azimuthal data sets for ‘State 1’ is plotted as a function of axial displacement. The absolute of the average value of the azimuthal data sets is also plotted (offset corrected) but note that it has been scaled by a factor of 1/20 to better fit the graph. The correlation indicates that the standard deviation is giving a measure of either the real field asymmetry and/or there is a positional error in the measurement. . . . . 233
- 10.16 If it is assumed that the magnets are fixed and that there is sufficient clearance in the bearings then the lower part of the shaft will have a resonance at 27.6 Hz[120]. . . . . 234
- 10.17 The new shaft design is composed from two separate pieces of titanium. The anti-rotation device is now part of the upper bearing. The bottom section of the lower tube (RHS) acts as the target[121]. . . . . 235
- 10.18 Left: The top bearing is built from a single piece with the addition of a wedge and a rotation stop. The wedge allows the inner surfaces to be coated with DLC as the process must have line of sight. Right: The bottom bearing is similar to the top bearing although the bearing aperture is larger and there is no rotation stop[121]. . . . . 236
- 10.19 The new body has been designed to strengthen and simplify the current design and will allow the bearings to be accurately located[121]. . . . . 238

# List of Tables

1.1	The fermions . . . . .	3
1.2	The fundamental forces and their force carriers . . . . .	3
1.3	Predicted solar neutrino fluxes from seven solar models. The table presents the predicted fluxes, in units of $10^{10}$ ( $pp$ ), $10^9$ ( ${}^7Be$ ), $10^8$ ( $pep$ , ${}^{13}N$ , ${}^{15}O$ ), $10^6$ ( ${}^8B$ , ${}^{17}F$ ), and $10^3$ ( $hep$ ) $\text{cm}^{-2}\text{s}^{-1}$ [22]. . . . .	17
2.1	Possible neutrino oscillation channels available to a neutrino factory muon storage ring. . . . .	32
7.1	The six states that create an apparent upward motion in the fields generated within the stator. The states are circular. Moving through the states in the opposite direction reverses the apparent motion - See figure 7.3 . . . . .	129
7.2	This table shows how the six stator states can be switched using the Hex Bridge. . . . .	141
9.1	This table shows the number of data sets collected with the target held in a stationary position. The distance is given with respect to the target's nominal holding position which was 78 mm from the centre of the beamline. The 0 mm data-sets were used to obtain background beam loss information [109]. . . . .	183
9.2	This table shows the number of data sets collected with the target actuating. The time is given with respect to the injection of protons into ISIS (ms) [109]. . . . .	183
9.3	Un-normalised rates for both shielded and unshielded detectors using MARS for $1 \times 10^7$ protons on target [56]. . . . .	187

9.4	The calculated number of non-neutron singles into the 1 cm <sup>2</sup> detector acceptance compared to the observed data from the November 2006 target tests [56]. . . . .	188
9.5	Simulation parameters used in initial modeling of the target ISIS beam. The proton interception rate was determined by the need to obtain 600 good muons through MICE per spill [65]. . . . .	188
9.6	Table showing target operation on the ISIS beamline from installation in Jan 2008 to July 2008. Generally the runs have been short as the MICE target has been operated cautiously. This cautious approach will continue until a second ‘demonstrator’ target system has been installed and operated in building R78. . . . .	191
10.1	A summary of the results of using different bearing materials . . . . .	223
10.2	A summary of the correlation between the measured standard deviation in the azimuthal data, the rate of change of magnetic field as a function of radius at an axial field peak and the resultant estimated field asymmetry in microns. . . . .	234

# Chapter 1

## Neutrino Physics

### 1.1 Introduction

There has been an increasing interest over the past couple of decades in neutrino physics; the discovery that neutrinos have mass and their ability to oscillate from one flavour to another has turned them from the rather uninteresting particles that they were once perceived to be into objects of much attention within the particle physics community. The prospect of finding physics beyond the standard model is exciting and this has opened up a plethora of experiments probing the properties of neutrinos. The background to neutrino oscillations and a description of some of the experiments that have probed the neutrino oscillation phenomena are discussed in this very chapter.

As will be illustrated these experiments have shown conclusively that neutrinos do oscillate from one flavour to another as they traverse through both space and matter. However the mass of the neutrinos and the parameters that determine the rate at which the neutrinos oscillate from one flavour to another are still not known to an accuracy that permits a good understanding of these sub-atomic particles. Better experimental data is clearly needed and in order to obtain this a new kind of neutrino experiment, on a much larger scale, is going to be required.

One possible solution to this problem is the creation of a ‘neutrino factory’, a machine that would be capable of producing the largest accelerator based flux of neutrinos yet created by man. Accelerator based neutrino sources offer many advantages to the experimentalist, primarily because of the well defined state and energy of neutrinos at the source. A large flux of neutrinos permits better statistics in the detectors and this ultimately improves the accuracy of the oscillation parameters that are being measured. Chapter two of this thesis looks at the possibility of building a neutrino factory and gives a glimpse of some of the technologies that will be required to build such a machine.

One of the technologies required by a neutrino factory will be ‘muon cooling’. Muons provide a good source of neutrinos but their rapid manipulation within an accelerator

will be necessary if they are to prove to be a suitable source of accelerator neutrinos. More than this, to create a large flux of neutrinos requires that the accelerator accepts many more muons into the beamline than is currently possible with conventional accelerator technology. One way of improving the flux of muons accepted into the accelerator is to first cool the muons by reducing their emittance. A more formal definition of emittance will be given in chapter three but essentially muon cooling reduces the muons emittance by making the muon beam more nearly parallel, i.e. collimating the muon beam, thus allowing the muon accelerator to accept more of the muons into its beamline.

The theory of muon cooling is not new but it has never been experimentally proven. MICE, the Muon Ionisation Cooling Experiment, has been set the task of implementing muon cooling in a short section of cooling channel. MICE will be based at the Rutherford Appleton Laboratory in the UK, home of ISIS, an 800 MeV proton synchrotron accelerator and Diamond, a 3 GeV electron synchrotron. An overview of the MICE experiment and how it reduces the emittance of the muons that pass through the cooling channel will also be explained in chapter three.

The MICE experiment requires a muon source and this is provided by the MICE target mechanism. The MICE target mechanism, operating parasitically on the ISIS synchrotron, dips a small titanium paddle into the halo of the ISIS beam when the protons have been accelerated to their maximum energy. The pions produced by the interaction of the target with the proton beam are then collected by the front end of the MICE beamline. Many of these pions will decay to muons within the MICE beamline; it is these muons that will ultimately be transported by this beamline for use in the MICE cooling channel.

The specification of the target mechanism is demanding, requiring that the target is accelerated at greater than 80  $g$  into the ISIS beam and operating in an environment that severely constrains the way in which the device has been designed and built. It is the design, build and operation of this device that is the focus of this thesis and details of this can be found from chapters four to eight.

Finally, chapter nine of this thesis will illustrate how the target mechanism has performed both from a electro-mechanical point of view, and in terms of its ability to produce particles for the MICE experiment.

## 1.2 The Standard Model

The standard model is currently the best working theory that offers an explanation of how matter interacts with radiation. It is able to model three of the four fundamental forces in a way that is consistent with both relativity and quantum mechanics in extremely good agreement with experimental data. The standard model falls short of being a complete theory of fundamental interactions because it does not incorporate

the fourth fundamental force, gravity.

The standard model explains fundamental interactions by consideration of how matter, fermions, interacts with the fundamental forces that are also known as force carrying particles, bosons. Electromagnetism and the weak force are collectively known as the electroweak force.

The fermions describe matter particles and can be defined by their spin; all fermions are spin  $\frac{1}{2}$  particles. Fermions are split into two sub-categories, the leptons and the quarks. The leptons only experience the electroweak force whereas the quarks feel both the electroweak and strong forces. Both the leptons and the quarks each contain three generations of particles where each generation is described by a pair of particles as shown in table 1.1.

Fermions			
Quarks		Leptons	
up $u$	down $d$	electron $e$	electron neutrino $\nu_e$
charmed $c$	strange $s$	muon $\mu$	muon neutrino $\nu_\mu$
top $t$	bottom $b$	tau $\tau$	tau neutrino $\nu_\tau$

Table 1.1: The fermions

The existence of anti-matter particles, particles of the same mass but opposite charge, serves to double the number of fermions so that there are a total of twenty-four fundamental particles. These fundamental particles interact with each other via the exchange of force bosons; the photon, weak gauge bosons, gluon and the graviton. All of the force bosons, except the graviton, have a spin value of 1 and are shown in table 1.2. The graviton is not considered in the standard model but is included in the table for completeness.

Force Carrier	Particle (Boson)
Electromagnetism	Photon
Weak Force	$W, Z^0$ Bosons
Strong Force	Gluon
Gravity	Graviton

Table 1.2: The fundamental forces and their force carriers

**Photons.** These are massless particles and mediate the electromagnetic force between electrically charged particles. Being massless the photon is able to carry its force over long range. The way in which photons interact with matter is described by the theory of quantum electrodynamics.

**Weak gauge bosons.** These mediate the weak force which is responsible for weak interactions. There are three particles responsible for mediating this force, namely the  $Z^0$ ,  $W^+$  and  $W^-$  particles. The  $Z^0$  is the most massive at 91.2 GeV. The  $W^+$  and the  $W^-$  have a mass of 80.4 GeV. The large mass of these particles means that the force is only felt over a very short range. These particles act on both the quarks and leptons and the behaviour of these gauge bosons is described by the weak interaction. The most familiar physical phenomenon that requires the weak force for its explanation is beta decay.

The standard model shows that electromagnetism and the weak force are in fact unified at high energy and become the ‘electroweak’ force. These two different facets of the ‘electroweak’ force come about due to a process known as symmetry breaking. The electroweak model predicts a particle called the ‘Higgs’ Boson. This particle has yet to be observed but its existence would provide a mechanism for giving particles mass. The observation of the Higgs particle at the Large Hadron Collider (LHC) is expected by a significant number of particle physicists over the next few years.

**Gluons.** These mediate the strong force and are responsible for the binding of quarks together. The strong force is extremely short range and is described by the theory of ‘Quantum Chromo-Dynamics’ or QCD.

The standard model has proven to be very successful at describing the fundamental interactions and has matched up very well to most experimental data. However it cannot be considered as a ‘complete’ theory as it does not include the force of gravity and it relies upon the insertion of eighteen ‘free’ parameters. All of these parameters have to be measured experimentally and these cannot be predicted from the theory itself. The free parameters of the standard model are:

- Six quark masses
- Three charged lepton masses
- Four parameters of the CKM matrix: 3 mixing angles and a CP violating phase
- Four free parameters in the electroweak model
- The strong coupling constant

Seventeen of these parameters have been measured and the only one that has not yet been observed and measured is the Higgs boson mass from the electroweak model. Despite the successes of the standard model it is clear that it is not a complete theory and that there are several physical phenomena that are not accurately described within the framework of the model. Some examples of these phenomena include:

**CP Violation and matter/antimatter discrepancy.** No known method of CP violation in the standard model can account for the observed matter/antimatter discrepancy in the universe today. This problem is related to the ‘Strong CP’ violation problem.

**Dark Matter.** The observed gravitational evolution of the universe cannot be accurately described if consideration is only given to baryonic matter. This is suggestive that a significant proportion of the mass in the universe is derived from non-baryonic matter. The standard model does not include this dark matter.

**Gravity.** As already noted gravity is not included at all in the standard model.

**Neutrinos.** In the standard model neutrinos are modelled as massless particles. Recent experimental data has confirmed that neutrinos oscillate from flavour to flavour, consequently neutrinos must have mass. The mechanism for describing massive neutrinos is not part of the standard model.

Whilst each of these topics is clearly very important, the subject of neutrinos is of particular relevance to this thesis. For this reason the rest of this chapter will focus on neutrino physics.

### 1.3 Neutrinos

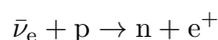
As briefly touched on above, in the standard model neutrinos are massless neutral particles of spin  $\frac{1}{2}$  that do not change flavour. They are modelled as the neutral partners to the charged leptons such that each charged lepton partners with a neutrino of the same flavour. These are often shown in a hierarchy of three doublets:

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix} \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \quad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}$$

Experiments have shown that all the observed neutrinos have left-handed helicity (spin is antiparallel to the neutrino's momentum) whereas all the antineutrinos are shown to have right handed helicity (spin is parallel to the anti-neutrino's momentum). For standard model neutrinos this is absolute, however for massive neutrinos this helicity is relative as it is clear that a change of reference frame can change the helicity of a neutrino by changing its momentum vector.

The neutrino was first postulated by Pauli in 1930 [1] to explain how the spectrum of the  $\beta$  particles emitted from radioactive beta decay could be continuous. Fermi shortly after worked out the theory of  $\beta$  decay and a cross section for neutrino interaction was then calculated by Bethe and Peierls [2].

Fred Reines and Clyde Cowan were the first to discover the electron (anti-) neutrino [3] by using the huge flux of neutrinos from the Savannah River nuclear reactor to look for inverse  $\beta$  decay which is given by the reaction:



The muon neutrino was first observed in 1962 by Lederman, Schwartz and Steinberger [4] at the then new Alternating Gradient Synchrotron at the Brookhaven facility in the USA. This landmark experiment not only established the existence of the muon neutrino but also demonstrated that the electron neutrino and the muon neutrino were distinct particles. When the tau lepton was discovered in 1975 at the Stanford Linear Collider [5] it was anticipated that this too would have its own neutrino partner. However the first observation of the tau neutrino had to wait another 25 years; observation of the tau neutrino was first announced by the Direct Observation of Nu-Tau (DONUT) collaboration based at the Tevatron Linear Collider at Fermilab in 2000 [6].

The elegance of the symmetry between the number of charged leptons and the number of different types of neutrinos was further confirmed from  $Z^0$  production in the electron-positron collider at LEP in the late 1990's. A number of measurements of the decay width of the  $Z^0$  resonance (L3 collaboration of LEP [7], ALEPH [8], OPAL [9] and DELPHI [10]) showed that the number of light neutrino flavours must equal three. These results are valid where the mass of the neutrino species is less than  $M_{\nu_x} < M_{Z^0}/2$ . It is worth noting that this result does not preclude the possibility of there being massive neutrinos (mass greater than  $M_{Z^0}/2$ ) or sterile neutrinos. Sterile neutrinos are a hypothetical neutrino species that do not take part in electroweak reactions. Both of these options are considered to be quite exotic within the framework of the current understanding of neutrino physics.

## 1.4 The Solar Neutrino Experiment

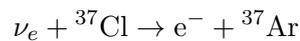
The standard solar model (SSM) is a model of the sun that aims to accurately describe its internal workings by incorporating all of the important physical processes that take place within it. Current knowledge of the physical processes that take place within the sun are constrained by both the instrumentation used to make direct measurements on the sun and by earth based experiments that aim to replicate the nuclear processes that occur within it. As knowledge of these processes has improved the SSM has evolved in both its detail and accuracy.

Validating the nuclear processes taking place within the sun is a difficult process as there is no easy 'probe' that can be used to directly measure the conditions within the core of the sun. Certainly before the discovery of the neutrino there was no known method of directly probing the nuclear reactions taking place at the sun's core. The discovery of the neutrino led to the realisation that the neutrino could possibly provide the necessary 'probe' that would be able to provide validation of the details of the nuclear reactions taking place within the sun's core.

The sun should be producing an enormous quantity of neutrinos within its core whilst burning its fuel. The neutrino's small cross section for interaction would make it an ideal candidate for obtaining valuable information about the energy spectra of the reactions in the core of the sun as many of the neutrinos should escape from the sun's

core unimpeded. Of course, the same small cross section for the neutrino makes them equally difficult to detect at the surface of the earth, therefore such measurements are extremely difficult to obtain.

The first experiment to try and detect solar neutrinos was undertaken by Ray Davis et al at the Homestake Mine in South Dakota USA in 1968 [11]. Their detection method was based upon an electron neutrino capture on a  $^{37}\text{Cl}$  atom which was provided by a 100,000 gallon tank of trichloroethylene  $\text{C}_2\text{Cl}_3\text{H}$ . The solar neutrinos would then undergo the following reaction with the chlorine atoms at a threshold energy of 814 keV [12]:



The result of this painstaking experiment was that Ray Davis only observed one third of the number of electron neutrinos that were predicted by the standard solar model. The difference between this experimental result and what the SSM predicted became known as the ‘Solar Neutrino Problem’. This discrepancy is now understood to have arisen from the inadequate but accepted understanding of neutrinos at the time of the experiment. The current understanding of this deficit can be explained by the phenomenon of neutrino oscillations.

## 1.5 Neutrino Oscillations

Neutrino oscillation is a theory about how one neutrino flavour, e.g. an electron neutrino, can turn into another neutrino flavour, e.g. a muon neutrino, as it passes through space and/or matter. Neutrino oscillations are a quantum mechanical effect. It is interesting to note that the theory of neutrino oscillation started to be independently put forward at about the same time that Ray Davis’ experiment was taking place at the Homestake mine.

In 1957, several years before Ray Davis published his results on his solar neutrino experiment a Russian physicist by the name of Bruno Pontecorvo first postulated that if neutrinos had mass then there was the possibility that they could oscillate from one flavour to another as they traveled through space [13]. In 1962 Maki, Nakagawa and Sakata [14] extended the work of Pontecorvo by introducing flavour mixing to show how neutrinos could oscillate from one flavour to another.

A quantitative theory of neutrino oscillations was first put together by Pontecorvo in 1967 and a paper illustrating this was first published in 1969 entitled ‘Neutrino Astronomy and Lepton Charge’ [15]. Further theoretical work by Mikheyev and Smirnov [16], based upon earlier work by Wolfenstein [17] showed how neutrino oscillations could be modified as neutrinos passed through matter to give an amplification of the oscillation effect. This theory has become important in understanding solar neutrino oscillation at

different neutrino energies and is known as the MSW effect; an outline of this process will be detailed in a following section of this chapter.

### 1.5.1 Three Neutrino Oscillation

Neutrinos are born in one of three flavour states  $\nu_e$ ,  $\nu_\mu$  and  $\nu_\tau$ . In neutrino oscillation theory it is assumed that each of these flavour states is comprised of an admixture of neutrino mass states called  $\nu_1$ ,  $\nu_2$  and  $\nu_3$ . The neutrinos propagate through space as independent mass states but can only be observed as flavour states.

More formally:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle \quad (1.1)$$

$$|\nu_i\rangle = \sum_\alpha U_{\alpha i} |\nu_\alpha\rangle \quad (1.2)$$

where the neutrino flavour state  $|\nu_\alpha\rangle$  is a summation over the mass eigenstates  $|\nu_i\rangle$ . Similarly the mass eigenstates are a summation over the flavour eigenstates.  $U$ , the unitary mixing matrix, can be written as:

$$U = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{bmatrix} \quad (1.3)$$

The above matrix components can be considered as a multiplication of three matrices where each matrix represents a rotation about one of the orthogonal base axes, the complex phase  $i\delta$  representing a CP violating phase. This can be written out fully as

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1.4)$$

where  $s_{xy} = \sin\theta_{xy}$  and  $c_{xy} = \cos\theta_{xy}$ . The matrix is often seen in the form where these rotations are multiplied out giving:

$$U = \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \quad (1.5)$$

The mixing matrix can be considered as a basis change between the orthogonal flavour states and the orthogonal mass states. This can be visualised as a rotation of one of the basis states onto the other where the rotation is given by the matrix. CP violation can only occur for a three neutrino matrix and only if none of the mixing angles equals zero. Figure 1.1 illustrates how the mass states and flavour states are correlated.

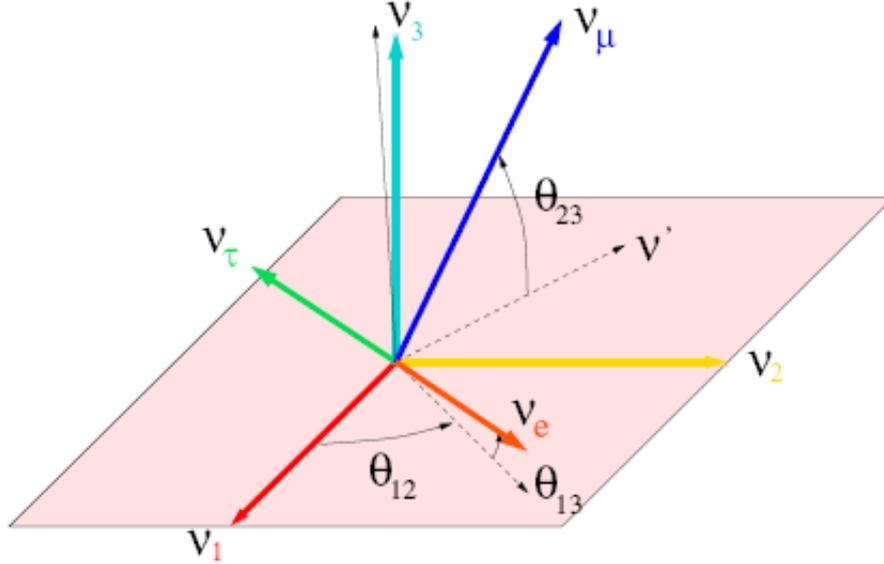


Figure 1.1: Diagram showing the correlation between the flavour and the mass bases [18].

To understand how neutrinos change from one flavour to another it is necessary to see how the neutrino evolves over time. This is done by applying the Schrodinger equation to a neutrino mass state which yields:

$$|\nu_i(\tau_i)\rangle = e^{-im_i\tau_i} |\nu_i(0)\rangle \quad (1.6)$$

where  $\tau_i$  is represent the time from the moment that the mass state  $m_i$  was created. This equation describes the neutrino in its own frame of reference. The space time translator  $e^{-i(Ht-px)}$  modifies equation 1.6 into a Lorentz invariant quantity [19]. Moving the result to the lab frame yields the result that:

$$e^{-im_i\tau_i} = e^{-i(E_it-p_iL)} \quad (1.7)$$

Now it is reasonable to assume that in most circumstances the neutrino will be ex-

tremely relativistic as its mass is so small, therefore  $t \approx L$  and this modifies equation 1.7:

$$|\nu_i(\tau_i)\rangle = e^{-im_i\tau_i} |\nu_i(0)\rangle = e^{-i(E_i-p_i)L} |\nu_i(0)\rangle \quad (1.8)$$

Given that the neutrino mass is much smaller than its energy then a further approximation can be made:

$$E_i = \sqrt{p^2 + m_i^2} \approx p + \frac{m_i^2}{2p} \quad (1.9)$$

and substituting this into equation 1.8 yields:

$$|\nu_i(\tau_i)\rangle \approx e^{(-im_i^2/2p)L} |\nu_i(0)\rangle \approx e^{(-im_i^2/2E)L} |\nu_i(0)\rangle \quad (1.10)$$

The mixing matrix gives the mass eigenstates in terms of the flavour eigenstates as shown in equation 1.1 so a simple substitution of equation 1.10 can be made that gives the time evolution of the flavour eigenstates in terms of the mass eigenstates.

$$|\nu_\alpha(L)\rangle \approx \sum_i U_{\alpha i}^* e^{(-im_i^2/2E)L} |\nu_i\rangle \quad (1.11)$$

Of course it is now a simple substitution to show the flavour evolution in terms of the original flavour eigenstates as defined in equation 1.2.

$$|\nu_\alpha(L)\rangle \approx \sum_\beta \left( \sum_i U_{\alpha i}^* e^{(-im_i^2/2E)L} U_{\beta i} \right) |\nu_\beta\rangle \quad (1.12)$$

This equation shows that as a flavour neutrino travels a distance  $L$  its state vector is rotated, or oscillated into a superposition of all the other flavours. The probability that a flavour has oscillated from one to another is given by the square of the amplitude for the process i.e.

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta | \nu_\alpha(L) \rangle|^2 \quad (1.13)$$

and this yields the general form for neutrino oscillations in the vacuum [20]:

$$\begin{aligned}
P(\nu_\alpha \rightarrow \nu_\beta) = & \delta_{\alpha\beta} - 4 \sum_{i>j} \text{Re}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left( \Delta m_{ij}^2 \frac{L}{4E} \right) \\
& + 4 \sum_{i>j} \text{Im}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin \left( \Delta m_{ij}^2 \frac{L}{4E} \right) \cos \left( \Delta m_{ij}^2 \frac{L}{4E} \right) \quad (1.14)
\end{aligned}$$

Note that the neutrino mass difference  $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$  and that when the units of  $c$  and  $\hbar$  are added then:

$$\Delta m_{ij}^2(L/4E) \equiv 1.27 \Delta m_{ij}(\text{eV}^2) \frac{L(\text{km})}{E(\text{GeV})} \quad (1.15)$$

In the case of no CP violation the probability for oscillation of a neutrino is the same as that for an anti-neutrino. Now for an anti-neutrino the complex conjugate of equation 1.14 is taken reversing the sign of the last term. Therefore the probability of oscillation can only remain the same if the imaginary term in equation 1.14 equals zero [20].

This equation shows that the probability of oscillation depends on several parameters; the mixing angles between the two basis' as contained within the mixing matrix, the squared mass differences of the mass eigenstates, the distance between the neutrino source and the point at which it is detected as well as the energy of the neutrino. The three neutrino equation yields three mass splittings,  $\Delta m_{12}^2, \Delta m_{23}^2, \Delta m_{13}^2$ , so whilst neutrino oscillations can yield the mass difference between the neutrinos it does not give the actual neutrino masses. The other subtlety is that it is not possible from this to establish the hierarchy of neutrino masses. This leaves two possibilities where  $\Delta m_{12}^2 < \Delta m_{23}^2 < \Delta m_{13}^2$  which is the normal hierarchy and  $\Delta m_{12}^2 < \Delta m_{13}^2 < \Delta m_{23}^2$  which is called the inverted hierarchy. Figure 1.2 shows this relationship diagrammatically.

### 1.5.2 Neutrino Oscillations in Matter

The last section has dealt with neutrino oscillations within a vacuum. However, as previously touched upon, Mikheyev and Smirnov [16] were to show that when neutrinos pass through matter this can modify the neutrino oscillation behaviour.

This effect is due to the fact that normal matter contains electrons that will only couple to the electron neutrino. As is well understood matter is composed largely of both nucleons and electrons. Now the nucleons will interact with the all the flavours of neutrinos identically, causing a certain amount of scattering via the neutral current interaction,  $Z^0$  exchange. As this process affects all of the neutrinos identically there is no net effect on the relative phases of the three mass states and so the oscillation parameters of the neutrinos are unaltered by their passage through baryons.

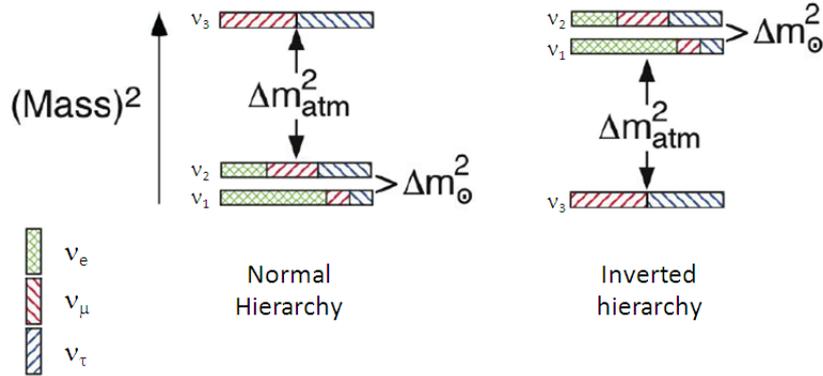


Figure 1.2: The two possible neutrino mass hierarchies. The  $\nu_e$  fraction of each mass eigenstate is crosshatched, the  $\nu_{\mu}$  fraction is indicated by right-leaning hatching, and the  $\nu_{\tau}$  fraction by left-leaning hatching ( $\Delta m_{12}^2 = \Delta m_{\odot}^2$  &  $\Delta m_{13}^2 \approx \Delta m_{23}^2 = \Delta m_{atm}^2$ ) [12].

However the electrons in the matter will also interact with the electron component of the neutrino state via a charged current exchange  $W^+$ . This asymmetry in how the electron component of a neutrino interacts with matter has a significant impact upon the way in which the neutrinos oscillate. The effect is most pronounced where there is a high electron density, such as in the sun or, to a lesser extent, within the earth and it is also dependent upon the energy of the neutrino.

For higher energy solar neutrinos the MSW effect is important and leads to the result that there is a significant increase in the probability of an electron neutrino transforming to a muon neutrino. The MSW effect is now recognised as being the primary cause of the solar neutrino problem. For lower energy neutrinos the MSW effect is less relevant and often the neutrino oscillation in vacuum formula can be used instead.

A formal derivation of the MSW effect and how it alters the neutrino oscillation is quite complex; an example can be found in the book ‘Physics of Massive Neutrinos’ by Felix Boehm and Petr Vogel [21].

### 1.5.3 Two Neutrino Approximation

If it is assumed that mixing to one of the neutrino families is negligible then the three neutrino case can be simplified to a two neutrino approximation. This works based upon the assumption that the mixing angle  $\theta_{13}$  is very small.

Starting with the neutrino mixing matrix 1.4 it is immediately obvious that gives us:

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} \quad (1.16)$$

letting  $\sin \theta_{13} = 0$  yields:

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} \quad (1.17)$$

The three by three matrices are rotation matrices and so an inverse matrix to these matrices can easily be found. Performing this manipulation yields:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & -s_{23} \\ 0 & s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} \quad (1.18)$$

From this point it is possible to go one of two ways depending on what the experiment is looking for. The first case is where the experiment is concerned with looking for oscillations where the electron neutrinos are detected and the experiment cannot discern between mu and tau neutrinos. Taking the left hand side this can be simplified and redefined so:

$$\begin{bmatrix} \nu'_\mu \\ \nu'_\tau \end{bmatrix} = \begin{bmatrix} c_{23} & -s_{23} \\ s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} \nu_\mu \\ \nu_\tau \end{bmatrix} \quad (1.19)$$

$$\nu'_e = \nu_e \quad (1.20)$$

Note that  $\nu'_\mu$  is an admixture of both the  $\nu_\mu$  and the  $\nu_\tau$ . Equating this with the right hand side of equation 1.18 then gives a two neutrino oscillation matrix ( $\theta = \theta_{12}$ ):

$$\begin{bmatrix} \nu_e \\ \nu'_\mu \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} \quad (1.21)$$

Of course  $\nu'_\tau = \nu_3$  but as it was assumed that  $\sin\theta_{13} = 0$  this component will be negligible. For the case where the electron neutrino plays almost no role a similar argument yields ( $\theta = \theta_{23}$ ):

$$\begin{bmatrix} \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} \nu'_2 \\ \nu'_3 \end{bmatrix} \quad (1.22)$$

In both cases the mixing matrix is the same. Consideration of the mixing probability leads to a single mixing angle and a single mass splitting. To avoid confusion the mass splitting and the angle are differentiated for solar and atmospheric values using a suffix after the quoted result.

## 1.6 Measuring Neutrino Parameters

The implication of neutrino mass and neutrino oscillations has led to an active research program to ascertain the key parameters of this new area of physics. Particularly, the neutrino oscillation matrix requires the determination of the mass splittings, the mixing angles and as to whether there is any evidence of CP violation. Experiments designed to observe neutrinos are categorised by the source of neutrinos. As a number of sources has been available over differing base lines a range of measurements has been possible over a wide energy range. Broadly speaking the experiments are classified as either:

- Solar Neutrinos
- Atmospheric Neutrinos
- Reactor Neutrinos
- Accelerator Neutrinos

The following section will summarise the current status of these measurements.

## 1.7 Neutrino Experiments

Due to the significant increase in the interest of neutrino physics there is now a considerable number of neutrino experiments that have either taken place or are in the process of running. As there is a significant amount of information available on these experiments it would not be realistic to cover all of these experiments in detail in this thesis. What is presented here is a summary of some of these experiments and the key results that have been obtained from them. For further details on the experiments that

are detailed here the interested reader is referred to the references given in the text. As the neutrino experiments are largely broken up into four categories that describe the various sources of neutrinos the experiments will be considered in terms of these.

### 1.7.1 Solar Neutrino Experiments

As already discussed the sun is an enormous source of neutrinos and its distance from earth gives a long baseline for neutrino experiments. Solar neutrinos have a broad energy spectrum with neutrino energies up to about 20 MeV - see figure 1.3 [22]. For those reactions that give neutrino energies greater than 5 MeV the sun's mass means that the MSW effect becomes important.

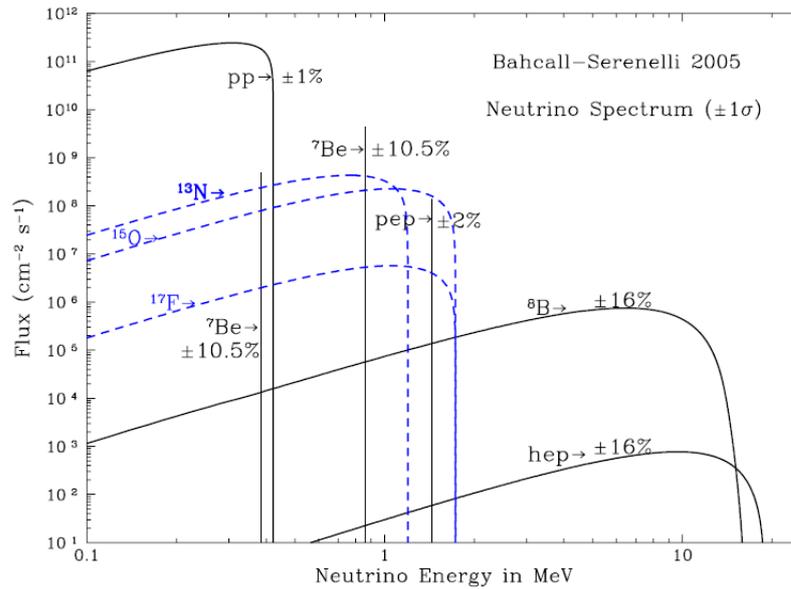


Figure 1.3: The solar neutrino spectrum [22].

Solar Neutrinos have to date been detected by one of two techniques. Firstly there is the radio chemical neutrino capture method. Here a neutrino is captured by the nucleus transmutating the element via the reaction:



Due to the small cross section of this type of reaction it is necessary to use a large volume of very pure material to capture the neutrinos in. Extraction and/or counting of the transmuted element is also difficult. The transmuted element needs to be radioactive with a fairly short half life so that its decay can be detected. One of the

key deficiencies of this method is that it does not give any detail as to the direction of the neutrino that took part in the reaction. Consequently it is not possible to be sure that any given interaction had its neutrino source come from the sun. This method is only sensitive to electron type neutrinos.

The other method is via the detection of Cherenkov light when an electron is scattered by the interaction with a neutrino. Typically this is achieved by using a large volume of liquid surrounded with many photo multiplier tubes to detect the Cherenkov light. The advantage of this method is it can give an indication of both the initial direction and the energy of the incident neutrino.

### Radio Chemical Experiments

The Homestake mine experiment, as introduced earlier in this chapter, by Ray Davis [11] was a radio chemical method. This experiment ran for over 30 years and consistently measured about 1/3 of the expected flux of electron neutrinos.

SAGE [23] and GALLEX [24] were two experiments that used the same radio chemical technique but this time on Gallium. Here neutrino capture would turn the Gallium into Germanium:



Defining one Solar Neutrino Unit (SNU) as  $10^{-36}$  captures per second per absorber nucleus then the expected standard solar model rate is  $128_{-7}^{+9}$  SNU for this reaction at an energy threshold of 233 keV [12]. Both of these experiments showed a consistent deviation from the expected flux of electron neutrinos. The SAGE experiment from January 1990 through December 2001 and showed a neutrino flux of:

$$70.8_{-5.2}^{+5.3}(\text{stat})_{-3.2}^{+3.7}(\text{sys}) \text{ SNU}$$

which was consistent with the GALLEX result of:

$$77.5 \pm 6.2(\text{stat})_{-4.7}^{+4.3}(\text{sys}) \text{ SNU}$$

### Cherenkov Detection

The Sudbury Neutrino Observatory (SNO)[25] that uses 1,000 tonnes of heavy water ( $\text{D}_2\text{O}$ ) as the detection medium for observing neutrinos. This heavy water detector, placed 2 km underground is surrounded by numerous photomultiplier tubes that detect

the cherenkov light from neutrino interactions. SNO detects neutrinos from three principle reactions channels; elastic scattering (ES), charged current (CC) interactions and neutral current (NC) interactions. Only electron neutrinos produce charged-current interactions, while the neutral-current and elastic scattering reactions have sensitivity to all neutrino flavours  $\nu_x$ .

The experiment has been run in two phases. The first phase ran for several months using pure ( $D_2O$ ) in the fiducial detection volume. The experiment was able to reject the no oscillation scenario of  $^8B$  neutrinos at  $5.3 \sigma$  [26] In the second phase of the experiment two tonnes of NaCl were added to the heavy water to enhance the neutron detection efficiency, this allowed for a better statistical separation of the NC and CC reactions. This improves the measured fluxes for the  $\nu_e$  and the  $\nu_x$  reactions. The results of this is that the total flux of solar  $^8B$  neutrinos measured from the sun was found to be [27]:

$$5.54_{-0.31}^{+0.36}(\text{stat})_{-0.34}^{+0.36}(\text{sys}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

This measurement is in agreement with the standard solar models as is shown in table 1.3 [22].

Model	pp	pep	hep	$^7Be$	$^8B$	$^{13}N$	$^{15}O$	$^{17}F$
BP04(Yale)	5.94	1.40	7.88	4.86	5.79	5.71	5.03	5.91
BP04(Garching)	5.94	1.41	7.88	4.84	5.74	5.70	4.98	5.87
BS04	5.94	1.40	7.86	4.88	5.87	5.62	4.90	6.01
BS05( $^{14}N$ )	5.99	1.42	7.91	4.89	5.83	3.11	2.38	5.97
BS05(OP)	5.99	1.42	7.93	4.84	5.69	3.07	2.33	5.84
BS05(AGS, OP)	6.06	1.45	8.25	4.34	4.51	2.01	1.45	3.25
BS05(AGS, OPAL)	6.05	1.45	8.23	4.38	4.59	2.03	1.47	3.31

Table 1.3: Predicted solar neutrino fluxes from seven solar models. The table presents the predicted fluxes, in units of  $10^{10}$  ( $pp$ ),  $10^9$  ( $^7Be$ ),  $10^8$  ( $pep$ ,  $^{13}N$ ,  $^{15}O$ ),  $10^6$  ( $^8B$ ,  $^{17}F$ ), and  $10^3$  ( $hep$ )  $\text{cm}^{-2}\text{s}^{-1}$  [22].

### 1.7.2 Current values of Solar Oscillation Parameters

SNO often publishes a combined analysis of all solar neutrino experiments at the conclusion of their papers. As of 2008 a combined analysis of all solar and reactor neutrino experiments yields the following results with respect to the solar mass splitting and mixing angle assuming a two species oscillation [27]:

$$\Delta m_{solar}^2 = 7.94_{-0.26}^{+0.42} \times 10^{-5} \text{ eV}^2$$

$$\theta_{\text{solar}} = (33.8^\circ)_{-1.3}^{+1.4}$$

Also see figure 1.4.

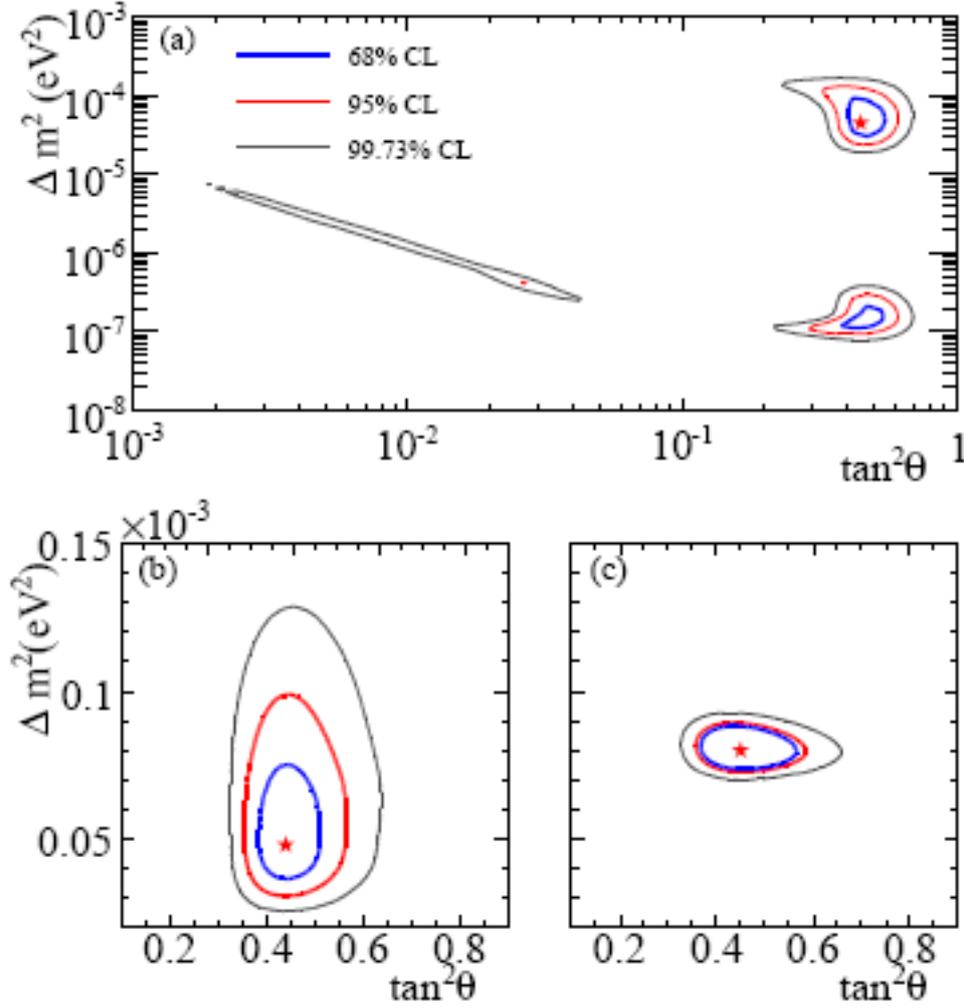


Figure 1.4: Neutrino-oscillation contours. (a) SNO only:  $D_2O$  & salt day and night spectra, neutral current detection (NCD) phase fluxes. The best-fit point is  $\Delta m^2 = 4.57 \times 10^{-5}$  eV<sup>2</sup>, with  $\tan^2 \theta = 0.447$ , (b) Solar Global: SNO, SK, Cl, Ga, Borexino. (c) Solar Global + KamLAND. The best-fit point is  $\Delta m^2 = 7.94 \times 10^{-5}$  eV<sup>2</sup>,  $\tan^2 \theta = 0.447$  [27].

### 1.7.3 Atmospheric Neutrino Experiments

Atmospheric muons are created indirectly when cosmic rays create pions in the upper atmosphere. The pions quickly decay to muons.

$$\pi^\pm \rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu)$$

Atmospheric neutrinos are also created by muon decay via the reaction

$$\mu^\pm \rightarrow e^\pm + (\bar{\nu}_\mu)\nu_\mu + (\nu_e)\bar{\nu}_e$$

Muon neutrinos created in the atmosphere typically have an energy from 100 MeV to several hundred GeV and above.

Super Kamiokande [28] (Super-K) was an upgrade to the large water Cherenkov detector Kamiokande that had been built for the purpose of detecting neutrinos from multiple sources. Kamiokande had been operated for several years and had observed several neutrino events but the results were inconclusive with respect to neutrino oscillation [29]. However, from the viewpoint of atmospheric neutrino observations the Super-K experiment has been the most significant.

Super-K is a 50,000 ton ultra pure water detector that is placed 1,000 m underground in a mine in Kamioka, Japan. Atmospheric neutrinos are typically detected via the leptonic reaction  $\nu_{e,\mu} + N \rightarrow l + X$  although other channels are also possible. In the case of the charged current reactions the ultra-relativistic lepton is detected via the Cherenkov radiation it emits when travelling through the water. The Cherenkov light is detected via the >11,000 photo-multiplier tubes (PMTs) that surround the detector.

Super-K has good angular, energy and time resolution. The directional sensitivity of this detector was important as there is clearly a correlation between the origin of the neutrino and the amount of matter that the neutrino has travelled through. With respect to atmospheric neutrinos, the neutrinos produced directly above the detector would only have to travel 15 km to get to the detector; however those produced in the atmosphere at the other side of the planet would have to travel through 12,600 km of rock before passing through the detector. The directional sensitivity of the detector was able to conclusively show a  $\nu_\mu$  to  $\nu_\tau$  oscillation [30] as is illustrated in figure 1.5.

### 1.7.4 Reactor Neutrino Experiments

Neutrinos are a byproduct of the nuclear reactions that take place within reactors. Nuclear reactors can produce a substantial flux of neutrinos. For example the Chooz B nuclear power station in France is a 4.2 GW power station that emits of the order of

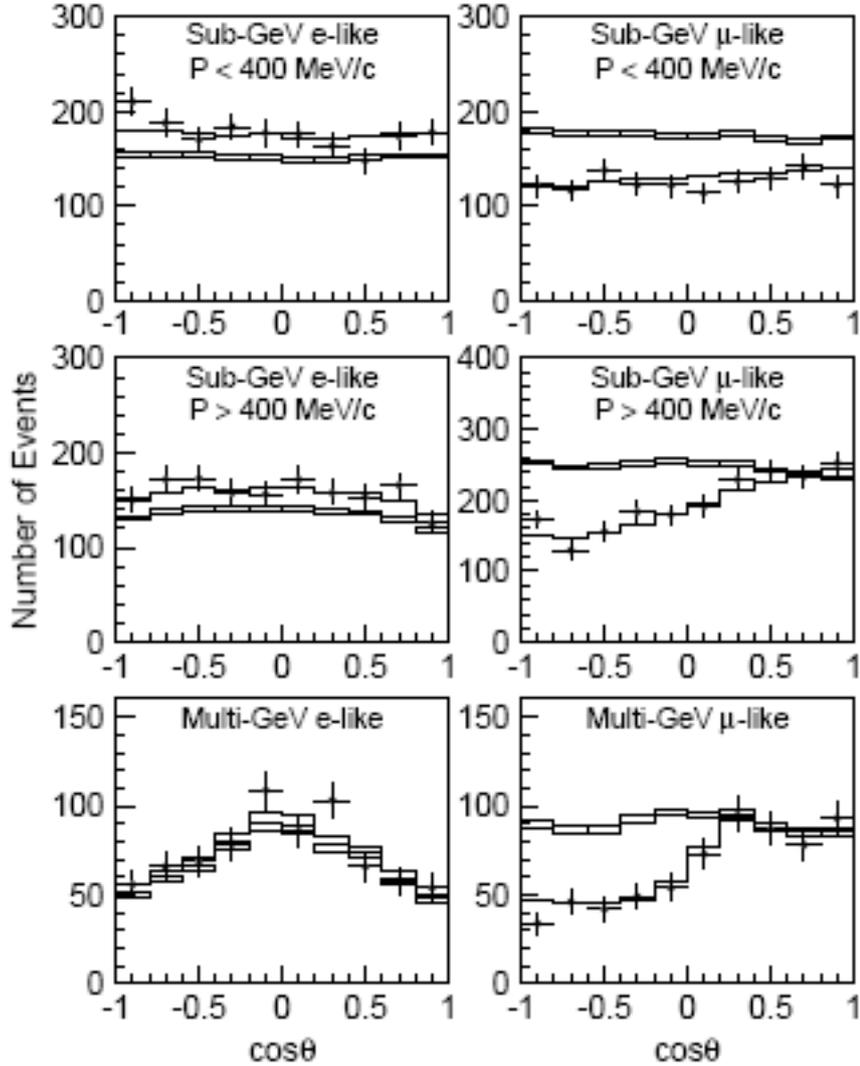


Figure 1.5: The zenith angle distribution for fully-contained 1-ring events in the Super-K detector. The points show the data, box histograms show the non-oscillated Monte Carlo events and the lines show the best-fit expectations for  $\nu_\mu \leftrightarrow \nu_\tau$  oscillations with  $\sin^2 2\theta = 1.00$  and  $\Delta m^2 = 2.1 \times 10^{-3} \text{eV}^2$  [30].

$1 \times 10^{21}$  neutrinos per second [31]. This extremely high flux of neutrinos makes doing experiments with these neutrino sources feasible. Another advantage of using nuclear reactors as a source for experiments is that both the energy spectrum and the neutrino flux is well known.

Increased sensitivity to low mass neutrino oscillations requires a large value of  $L/E$ . This is because the smaller the mass difference then the rate of change of any neutrino oscillation with respect to distance is reduced. As the experiment needs to see an

observable reduction in the number of neutrinos which is often statistically limited this places a limit on the observable mass difference. Reactor experiments typically offer a large value of  $L/E$  as the neutrino energy is so low and so provide good sensitivity to the mass difference over reasonable baseline distances. The short baseline clearly offer the added incentive of an increased flux of neutrinos at the detector.

As a result of the advantages that reactor neutrinos have to offer several neutrino experiments using nuclear reactors as their neutrino source have been run. The Chooz experiment produced an important result as it placed an upper limit on the value of  $\theta_{13}$  as a function of  $\Delta m_{13}^2$  is shown in figure 1.6 [32]. Genuine three flavour oscillation effects, such as CP violation, are only going to be observed if  $\sin^2 \theta_{13} > 0$ . Establishing the value of  $\sin^2 \theta_{13}$  is clearly of paramount importance to the neutrino community. A proposal to extend Chooz into Double Chooz has been put forward and this aims to either measure the mixing angle  $\sin^2 \theta_{13}$ , or observe a limit on the value of  $\sin^2(2\theta_{13})$  to less than 0.0025 [33].

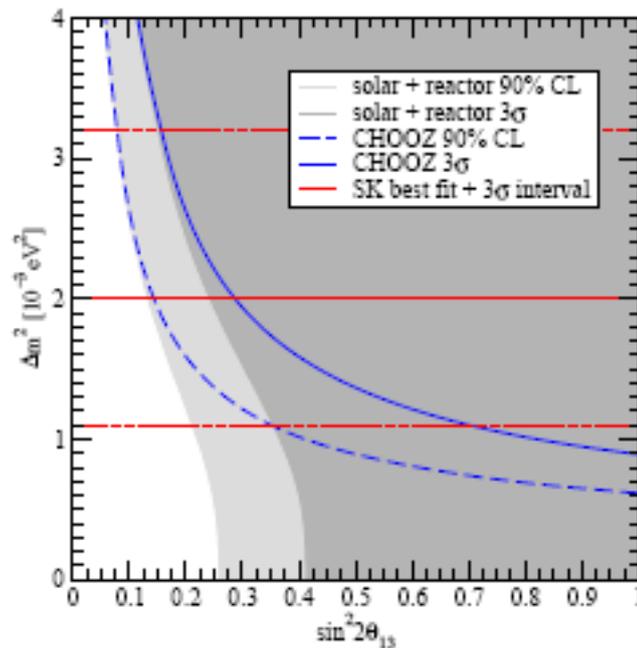


Figure 1.6: The upper bound on  $\sin^2 \theta_{13}$  as a function of  $\Delta m_{13}^2$  from the Chooz experiment.

The Kamioka Liquid scintillator Anti-Neutrino Detector (KamLAND)[34] is a neutrino oscillation experiment that uses reactor neutrinos to look at the solar neutrino problem. The experiment studies the disappearance of electron antineutrinos  $\bar{\nu}_e$  from multiple nuclear reactors across Japan. Although the anti-neutrino flux is sourced from many nuclear reactors at a variety of distances,  $\approx 80\%$  of the flux originates from 26 reactors that are within 138 km - 214 km of the KamLAND detector.

The KamLAND detector itself is constructed from 1 kiloton of ultra-pure Liquid Scintillator contained in a 13 m diameter spherical balloon. An array of 1879 photomultiplier tubes (PMTs) are mounted on the inner surface of the detector and this comprises the inner detector system. The Liquid Scintillator containment vessel is also surrounded by a 3.2 kton water-Cherenkov detector with 225 PMTs and this comprises the outer detector system. This outer detector provides a tag for atmospheric muons and absorbs  $\gamma$  rays and neutrons from the surrounding rock thereby reducing backgrounds. In an exposure of 145.1 days the ratio of the number of  $\bar{\nu}_e$  events above 3.4 MeV that the KamLAND detector observed compared to the expected number without oscillation is:

$$\frac{\bar{\nu}_e \text{ (observed)}}{\bar{\nu}_e \text{ (expected)}} = 0.611 \pm 0.085(\text{stat}) \pm 0.041(\text{syst})$$

This figure is consistent with non-oscillation being excluded at a 99.95 % confidence level [35].

Kamland was well placed to probe the large mixing angle solution for solar neutrino oscillations. When the Kamland results are placed with the results from other solar neutrino experiments a best-fit prediction can be made with this data. Doing so favours a large mixing angle (LMA) prediction of  $\sin^2 2\theta = 0.833$  and a mass splitting of  $\Delta m^2 = 5.5 \times 10^{-5} \text{ eV}^2$  [36]. This fit is shown in figure 1.7.

### 1.7.5 Accelerator Neutrino Experiments

Particle accelerators have the advantage that they can produce extremely pure beams of neutrinos due to their monochromaticity. Neutrino energies of up to a few GeV are currently possible. Accelerator complexes usually generate  $\nu_\mu$  beams and these are useful for exploring the following channels:

$$\begin{aligned} \nu_\mu &\rightarrow \nu_e \\ \nu_\mu &\rightarrow \nu_\tau \\ \nu_\mu &\rightarrow x \end{aligned}$$

where  $x$  can be any neutrino species. The first two reactions are useful for appearance experiments where the neutrino species on the right hand side is searched for and compared to the known flux of  $\nu_\mu$ . The third reaction is a disappearance reaction, so the experiment does not care what the decay products are. In this case two detectors are used at two different distances from the source. The normalised  $\nu_\mu$  rates will differ in the two detectors if neutrino oscillations have occurred.

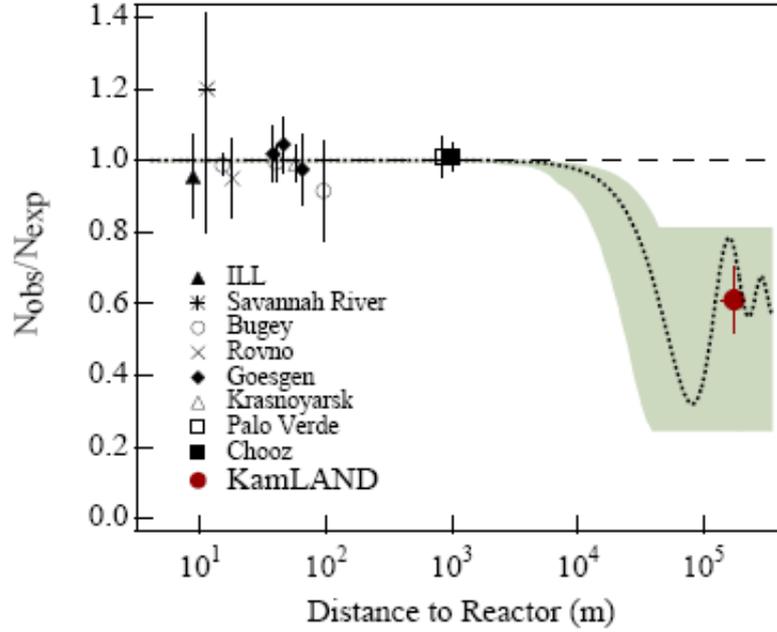


Figure 1.7: The ratio of measured to expected  $\bar{\nu}_e$  flux from reactor experiments. The solid dot is the KamLAND data point plotted at a flux-weighted average distance (the dot size is indicative of the spread in reactor distances). The shaded region indicates the range of flux predictions corresponding to the 95 % C.L. LMA region found in a global analysis of the solar neutrino data. The dotted curve corresponds to  $\sin^2 2\theta = 0.833$  and  $\Delta m^2 = 5.5 \times 10^{-5} \text{ eV}^2$  [36] and shows the best-fit LMA predictions while the dashed curve illustrates the no oscillation case [35].

To date many accelerator experiments have been run or are running and examples include the the Main Injector Neutrino Oscillation Search Experiment (MINOS) and KEK to Kamioka - the long baseline neutrino oscillation experiment (K2K). T2K and NOVA are two further experiments that are in the process of being built. The CERN Neutrinos to Gran Sasso (CNGB) is another accelerator neutrino experiment that is being commissioned at the time of writing this thesis.

The MINOS experiment is based in North America and uses two detectors; one of these is located at Fermilab and the other is located at the Soudan Underground Mine State Park in northern Minnesota a further 450 miles away [37]. The MINOS beam is a muon neutrino beam but it does contain a small amount of electron neutrinos. MINOS measures the neutrino beam composition in both the near and far detectors. This technique of using near identical detectors has the advantage that the uncertainties in neutrino flux and detector efficiency become less important and therefore reduce the systematic errors in any measurement. The normalised  $\nu_\mu$  rates in the two detectors are expected to be different if neutrino oscillations have occurred. However it is not

a pure disappearance experiment as the far detector is also capable of looking for the appearance of electron neutrinos. This technique allows for the possibility of making a precision measurement on the  $\sin^2\theta_{12}$  mixing angle and  $\Delta m_{12}$  mass-squared term.

In late 2007 the MINOS collaboration was able to report that they had measured a deficit of  $\nu_\mu$  events in the far detector and that this was consistent with two flavour neutrino oscillation with:

$$\Delta m_{atm}^2 = 2.74 \pm_{0.26}^{0.44} \times 10^{-3} \text{eV}^2$$

and a mixing angle of  $\sin^2 2\theta_{atm} > 0.87$  at  $1\sigma$  confidence; see figure 1.8 [38].

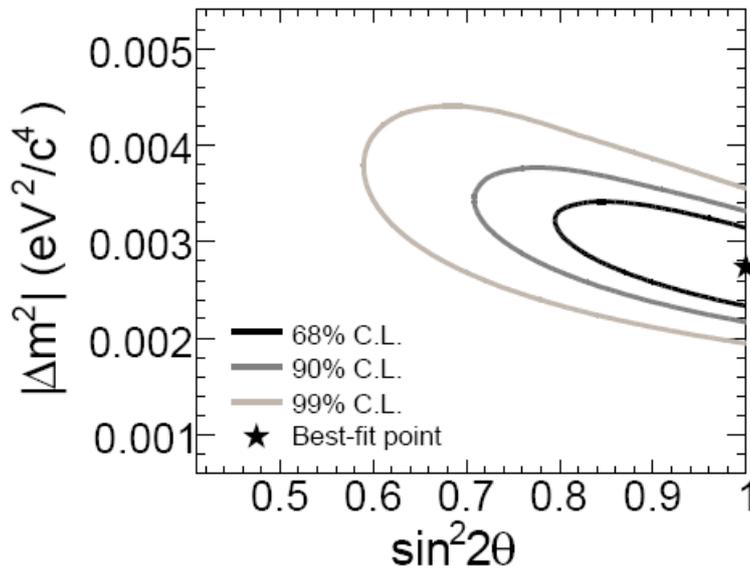


Figure 1.8: Allowed regions at 68 %, 90 %, 99 % Confidence Level in the  $\Delta m^2, \sin^2 2\theta$  plane from a fit to the far detector reconstructed energy spectrum for the MINOS experiment. The best-fit point, which occurs at  $\sin^2 2\theta = 1$  and  $\Delta m^2 = 2.74 \times 10^3 \text{ eV}^2/c^4$  is represented by the star [38].

The K2K experiment [39] is another  $\nu_\mu$  to  $\nu_\tau$  experiment that uses an accelerator produced  $\nu_\mu$  beam to explore the same  $\Delta m^2$  region as produced by atmospheric neutrinos. The neutrinos are measured by a complementary pair of detectors; the first detector is a 1 kiloton water Cherenkov detector that is 300 m away from the muon source and the second detector is the Super Kamiokande detector, a 50 kiloton water Cherenkov water detector which is 250 km away. In December 2006 the collaboration reported  $\Delta m^2$  is between  $1.9$  and  $3.5 \times 10^{-3} \text{ eV}^2$  when  $\sin^2 2\theta = 1$  at the 90 % confidence level with a best fit value of  $2.8 \times 10^{-3} \text{ eV}^2$  [40].

The Liquid Scintillator Neutrino Detector (LSND) experiment is a well known experiment that also looked for  $\nu_\mu$  to  $\nu_e$  events. Its results suggested a mass splitting for  $\Delta m_{12}^2$  of  $0.2 \text{ eV}^2$ - $10 \text{ eV}^2$  [41]. This result is inconsistent with the data produced from other solar and atmospheric experiments.

The MiniBooNE experiment is another accelerator based neutrino experiment built to test the LSND result. The miniBooNE detector is based at Fermilab and is filled with 800 tons of mineral oil and lined with 1,280 photomultiplier tubes. The detector is placed only 500 m from the particle production point. This experiment recently published results that refuted the LSND experimental result [42].

## 1.8 Conclusions

To date, the combined results from the neutrino experiments lead to the following known values for the neutrino oscillation parameters:

$$\Delta m_{12}^2 = 7.59 \pm 0.21 \times 10^{-5} \text{ eV}^2$$

$$\tan^2 \theta_{12} = 0.47_{-0.05}^{+0.06} \text{ [43][44]}$$

$$1.9 \times 10^{-3} \text{ eV}^2 < \Delta m_{23}^2 < 3 \times 10^{-3} \text{ eV}^2$$

$$\sin^2 2\theta_{23} > 0.9 \text{ [45] [46]}$$

$$\Delta m_{13}^2 \approx \Delta m_{23}^2$$

$$\sin^2(\theta_{13}) < 0.032 \text{ [45] [46]}$$

The CP violating phase  $\delta$  is unknown. These experimental limits are also summarised in figure 1.9.

The results are illuminating but important measurements are missing and the neutrino parameters are not currently known accurately enough to answer some of the important questions with respect to neutrino phenomenology. Primarily these questions are:

- How small is  $\theta_{13}$ ? Is it zero?
- If  $\theta_{13}$  is not zero what is the value of the CP violating phase,  $\delta$ ?
- Is  $\theta_{23}$  maximal?
- What are the absolute masses of the neutrinos?
- Are the neutrino masses normal or inverted?

To make the measurements that will answer these questions will require a new accelerator facility that produces a significantly higher neutrino flux than is currently available. The current ideas behind how these next generation neutrino sources can be provided will be discussed in the next chapter.

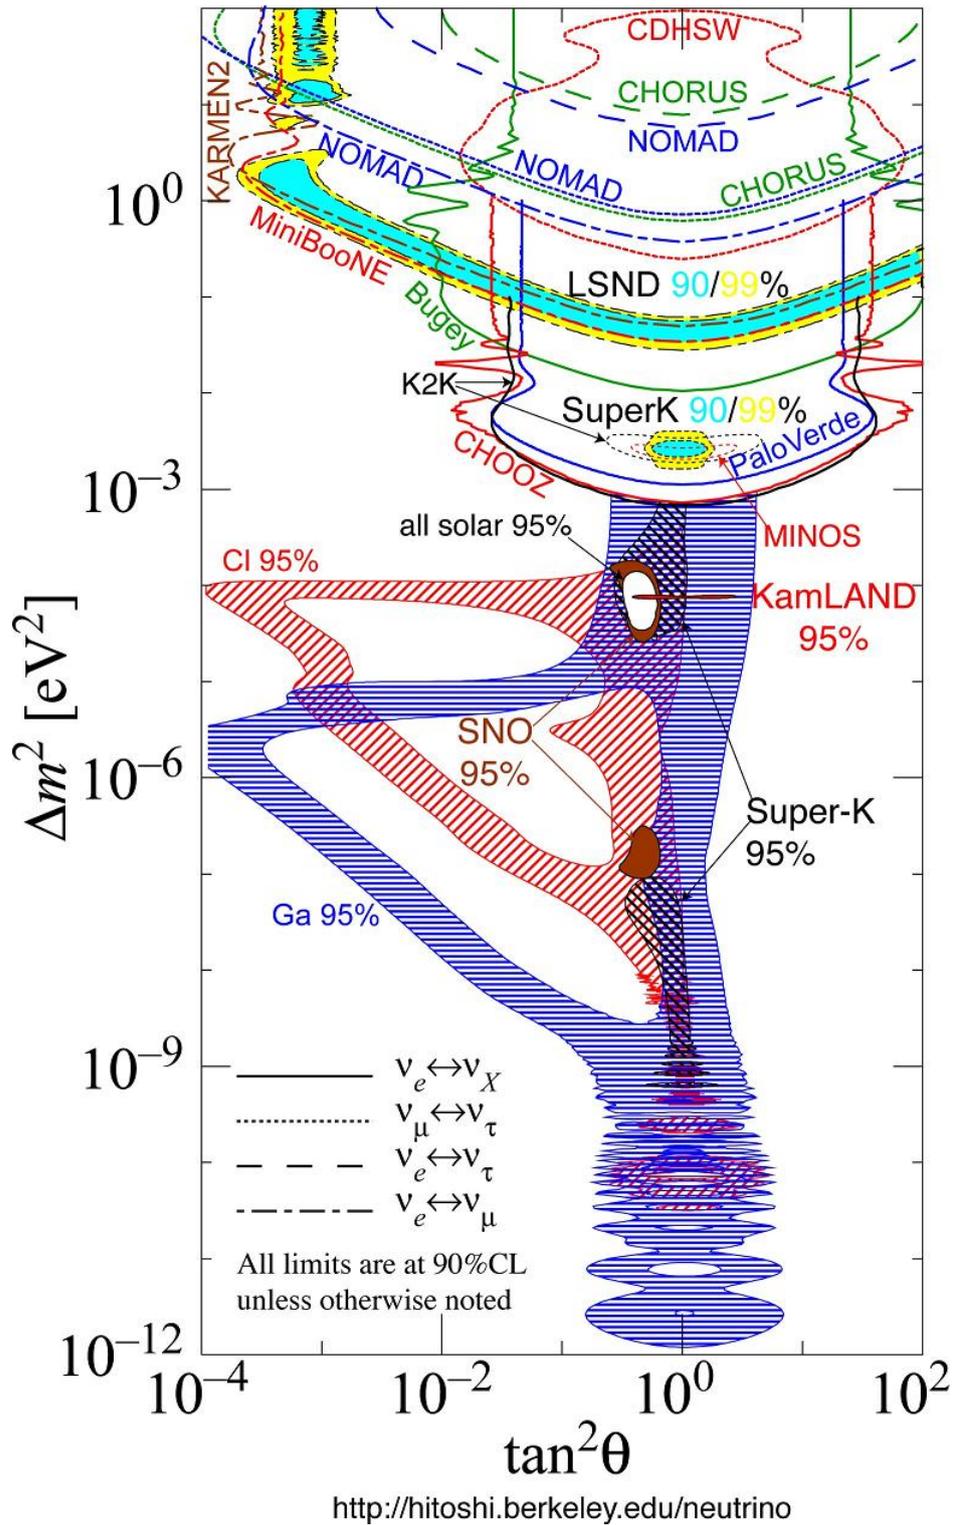


Figure 1.9: Oscillation Parameter Plot showing the current experimental limits on the neutrino oscillation parameters [47].

## Chapter 2

# The Neutrino Factory

### 2.1 Creating Neutrino Beams

It is clear that an improvement in the understanding of neutrinos is only going to be obtained by the production of high intensity neutrino beams via particle accelerators. Neutrino beams derived from particle accelerators have the advantage of being a well understood and characterised source as the neutrinos have a well known energy spectrum and flux. Future neutrino facilities have the potential to provide a much higher neutrino flux than any other known source that could be utilised.

All of these requirements are necessary to be able to collect enough experimental statistics with a high enough precision to answer the outstanding questions in neutrino phenomenology as described in the previous chapter. The high flux requirement for a future neutrino source will be principally dependent upon one of the three following technologies.

- Super Beams
- Beta Beams
- Neutrino Factory

Several independent designs utilising these technologies have been put forward in conceptual design reports and their various merits considered [48][49][20][50][51]. The advantages and disadvantages of these conceptual designs have now been examined and combined into a single baseline design that was reported in a scoping study comprising of three separate documents [52][53][54]. (These documents can also be found together on the neutrino factory international scoping study website [55].) The first of these reports puts together the physics case for an extensive experimental programme that would involve building a facility to enable a detailed study of the neutrino. The

three types of facility are compared based upon their physics potential, whether it is technologically feasible to provide such a facility and the cost. It is clear that the neutrino factory will offer the richest physics potential as it has the capability of measuring neutrino parameters to a greater accuracy than either a super beam or a beta beam source. It will also offer the potential for much muon physics as its neutrino production is based around the decay of an intense muon source.

Given that the target mechanism which is the subject of this thesis is being built for the MICE experiment, and in turn that this experiment is part of the research program for the development of a neutrino factory, this chapter will primarily concern itself with describing the design and physics potential of a neutrino factory. First, a very brief description will be given of what super beam and beta beam technology is.

### 2.1.1 Super Beams

Several experiments have used, or plan to use, pion decay from conventional proton beams to produce relatively clean muon neutrino beams. Super beams are an extension of the conventional way of producing neutrino beams through pion production and decay by increasing the intensity of the proton beam onto the target.

$$p + p(n) \rightarrow p + n(p) + \pi^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}(\bar{\nu}_{\mu})$$

A higher intensity proton beam produces more pions, these clearly create a larger neutrino flux from the initial pion decay. A super beam is usually taken as a conventional neutrino beam but with a proton source that has a power in excess of 2 MW. This is a significant amount of beam power and is in excess of what is typically produced from proton drivers today although much work is currently underway to increase this power up to the megawatt range [52].

### 2.1.2 Beta Beams

Beta beams use unstable radioactive ions that beta decay. By first accelerating these ions and then keeping them in a storage ring with long straight sections an intense beam of neutrinos can be created. Beta beams offer several advantages over other sources of neutrino beams. Firstly a beta beam has a very well characterised beam spectrum as the energy of the decay products is well known. Secondly, the nuclides can be chosen to have a half life that allows them to be accelerated relatively slowly to the required energy (compared to muon accelerators); this permits the use of conventional accelerator technology.

## 2.2 Neutrino Factory

As already alluded to the neutrino factory creates a neutrino beam from a muon beam via pion decay. Because muons have a finite lifetime before decaying this requires that the muon beam be accelerated up to its required energy quickly. The need for a high neutrino flux necessitates that a large muon flux be captured into the accelerator. This can only be achieved if the muon beam is effectively ‘collimated’ by having its emittance reduced. The emittance reduction has to happen quickly and the technology required to do this, ionisation cooling, is going to be demonstrated with the MICE experiment. The mechanism and technology to do this will be presented in the next chapter.

The use of a muon beam is challenging but if successful it will offer several advantages over the other two technologies; primarily it would significantly improve the ability to provide the most accurate measurements of the neutrino parameter space. A design for the neutrino factory will inherently provide both muon and anti muon beams, thus allowing the facility to also pursue an active physics program in making high precision studies of the muon. It has also been shown that a neutrino factory would provide an order of magnitude better sensitivity to  $\sin \theta_{13}$  than either super beams or beta beams. This will clearly be an important consideration if  $\sin \theta_{13}$  turns out to be small and non-zero. It also provides significantly greater sensitivity to other neutrino oscillation parameters than the other facilities can offer. The following section gives an overview of the principle features in an expected neutrino factory design.

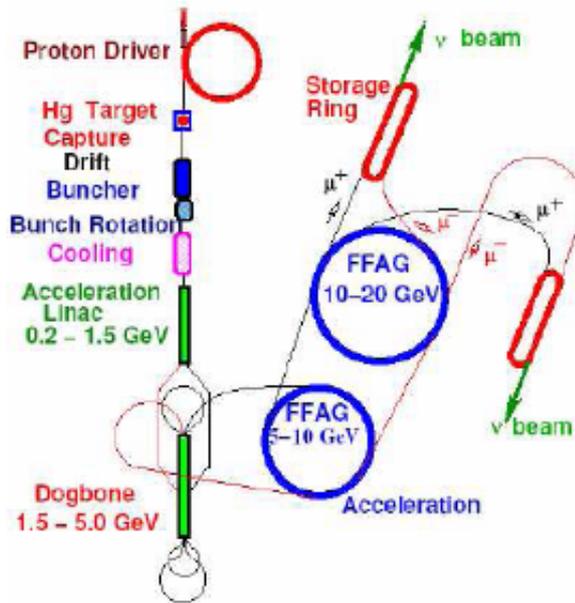


Figure 2.1: Current conceptual layout of the neutrino factory [56].

## 2.3 Neutrino Factory Design

This section provides an overview of some of the design considerations for a neutrino factory complex. However this review is by no means complete or comprehensive. For further information the reader is referred to the neutrino collaboration's literature. A good place to start is the NuFact website [57].

### 2.3.1 Proton Source and Driver

Any neutrino factory design requires a source of pions produced through a proton beam incident upon a target. The requirements for the proton driver are similar to that required to make a super beam. The difference between a super beam and the neutrino factory is that the latter takes the muons as the source for the neutrinos, whereas the super beam source would use the neutrinos produced from the pion decay. The high intensity of muons required at the front end of the neutrino factory places significant demands upon the proton source and target. A high power proton driver will be essential in any design for a neutrino factory. Initial studies indicate that a beam power of up to 4 MW will be required [52].

### 2.3.2 Target

The high powered proton beam will need to be directed to a target to create the initial pion flux. It is clear that the target will be an enormous technical challenge as the extreme energy density of the proton beam will destroy any fixed solid target.

Two solutions to this problem are currently being investigated. The first idea is to use a rotating solid cylindrical metal target. The rotation of the target will both reduce the power density at any given point and give the active volume of the target opportunity to cool between proton pulses. The second possible solution is to use a jet of liquid mercury as the target. This idea is currently being pursued by the MERIT experiment; preliminary results have recently been released demonstrating the validity of the principle [58].

### 2.3.3 Muon Phase Rotation and Ionisation Cooling

The neutrino factory target will produce pions in short bunch lengths over a wide momentum range. The subsequent decay of these pions leaves the bunched muons with the same bunch structure and a similar momentum spread. If the bunches of muons are allowed to drift over a distance of  $\approx 100$  m, an energy time correlation will develop. A high power radio frequency (RF) field can then be applied to the muon bunches splitting them into several micro bunches. These micro bunches can then be

accelerated by a second RF system. This RF system being phased with respect to the muon bunch so that the slower muons are accelerated, and the faster muons are decelerated towards the central momentum of the initial macro-bunch. This is done by timing the RF system so that the later muons see a positive RF phase and the earlier muons see a negative RF phase. This process is called phase rotation and an illustration of the principle is shown in figure 2.2.

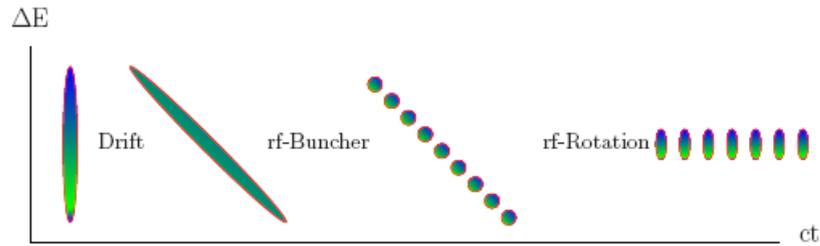


Figure 2.2: A pictorial representation of how phase rotation will alter the beam dynamics by the application of suitable RF fields.

After the phase rotation the emittance of the beam must be reduced, i.e. the beam must be collimated, so that more of the muons can be accepted into the accelerator. Fast emittance reduction can be achieved by a process called ‘Ionisation Cooling’. Ionisation Cooling will reduce the 6D phase space volume (reduces the momentum spread of the particles in all 3D) of the muon beam quickly with minimal particle losses. This technology will be proven by the Muon Ionisation Cooling Experiment (MICE) and a more detailed discussion of both emittance and how muon ionisation cooling reduces the emittance of the beam can be found in chapter three. This is no trivial effect; Ionisation Cooling could increase the beam’s intensity in a neutrino factory by as much as a factor of sixteen [59].

### 2.3.4 Acceleration

Once ionisation cooling of the muons has taken place the acceleration of the muons must be done promptly so that they can be placed in storage rings before they decay. Conventional synchrotron designs would be unable to accelerate the muons quickly enough so a novel acceleration techniques is required. The three most popular alternatives are the use of Rapid Cycling Synchrotrons (RCS’s) also known as a rapid pulsed synchrotrons [60], Recirculation Linear Accelerators (RLA’s) [20] and Fixed Field Alternating Gradient Accelerators (FFAG’s) [50].

### 2.3.5 Storage

The accelerated muons will then be stored in a suitable storage ring. The design of the storage ring must maximise the neutrino flux on the far detectors. This is accomplished by maximising the length of the straight sections and minimising the length of the bends. Two designs have been proposed for the storage ring; one is based upon a racetrack design and would be suitable for a single neutrino detector, and the second is based upon an isosceles triangle design. In the latter case two of the straight sections could be utilised and aimed at two separate neutrino detectors.

## 2.4 Physics Motivation

The construction of a neutrino factory would represent a significant investment in both labour and cost and so the physics motivation behind the project must be convincing. The evidence for neutrino oscillation is now accepted and these oscillations can be parameterised by three neutrino oscillation mixing angles, and two mass splittings and a CP violating phase as described in chapter one of this thesis. The existence of these parameters quite naturally leads to the physics motivation of measuring them as accurately as possible.

As the design of the neutrino factory utilises muons as its source of neutrinos and because the storage ring can store both muon charges this offers the possibility of twelve different oscillation channels that can be used to explore neutrino oscillations (The muon provides a source of both muon and electron type neutrinos/anti-neutrinos). These oscillation channels are shown in table 2.1 [61].

Channel	
$\nu_e(\bar{\nu}_e) \rightarrow \nu_e(\bar{\nu}_e)$	Disappearance Channel
$\nu_e(\bar{\nu}_e) \rightarrow \nu_\mu(\bar{\nu}_\mu)$	Golden Appearance Channel
$\nu_e(\bar{\nu}_e) \rightarrow \nu_\tau(\bar{\nu}_\tau)$	Silver Appearance Channel
$\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$	Platinum Appearance Channel
$\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_\mu(\bar{\nu}_\mu)$	Disappearance Channel
$\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_\tau(\bar{\nu}_\tau)$	Appearance Channel

Table 2.1: Possible neutrino oscillation channels available to a neutrino factory muon storage ring.

Studies have shown that a neutrino factory will give a significantly greater sensitivity to the value of  $\sin\theta_{13}$ , sensitivity to the hierarchy of the mass splittings and more sensitivity to the CP violating factor  $\delta$  when compared to any other neutrino facility.

Optimisation studies of the neutrino factory for detectors at several baselines have indicated that the neutrino factory shows superior performance in parameterising the

neutrino [62][63]. Primarily the following key measurement could, in principle, be achieved;

- The angle  $\theta_{13}$  can be measured to a precision of up to tenths of a degree, dependent upon its magnitude. The sensitivity of the experiment to  $\sin \theta_{13}$  is effectively improved by at least an order of magnitude above what is presently achievable.
- If the solar deficit can be explained using the large mixing angle (LMA) MSW effect then the neutrino factory will be sensitive to CP violation.
- The sign of the atmospheric mass splitting  $|\Delta m_{23}^2|$  can be determined (dependent upon the value of  $\sin \theta_{13}$ ) thereby allowing the neutrino mass hierarchy to be determined.

## 2.5 Conclusion

The neutrino factory would undoubtedly be the best of the proposed facilities for both improving the accuracy of, and making new measurements on, the neutrino parameter space. The physics from such a facility would permit a breakthrough in the understanding of the neutrino and perhaps cast new insights on leptonic CP violation. Much research and development is currently in progress to make such a facility a reality. MERIT is one example that has been mentioned in this chapter but there are many others active areas of research including, but not limited to, FFAG's, ion sources and targetry.

One area of active research is towards realising muon ionisation cooling. The Muon Ionisation Cooling Experiment (MICE) has been charged with the responsibility of building and successfully demonstrating the Ionisation Cooling Technology that would be one of the principle components required to make the neutrino facility a reality. The next chapter will discuss the principle of emittance reduction and the MICE experiment.



## Chapter 3

# The MICE Experiment

### 3.1 Introduction

The Muon Ionisation Cooling Experiment (MICE) [18] has been charged with two principal roles in its experimental program. Firstly there is the technical aspect of actually building and running a section of cooling channel. Technically implementing muon cooling will be challenging and the proof that this can be done will be a significant step towards the development of a neutrino factory. It is worth noting that this issue is quite distinct from the testing of the theory of ionisation cooling, the principle itself being considered sound. Secondly there will be testing of the cooling channel with a muon beam such that the performance of such a channel can be ascertained over different operating conditions both with respect to changing machine parameters and using muon beams with differing momenta.

MICE will demonstrate cooling for muons of momentum within the range of 140 MeV/c to 240 MeV/c. It is estimated from simulation that the MICE cooling channel will show a reduction in beam emittance of approximately 10 % although the actual reduction in emittance will be dependent upon the muon beam's input emittance. A more thorough discussion of beam emittance can be found in the next section.

The MICE experiment will measure the emittance of the beam by individually tracking the muons that go through MICE on a particle by particle basis. The reason that this is being done is two-fold. Firstly a particle by particle measurement will allow a much more accurate measurement of the reduction in emittance. It is hoped that the beam emittance will be measured to better than 0.1 % [18]. The measurement of the emittance to this level of accuracy has demanded a significant amount of research and simulation to ensure that this requirement can be met. A brief overview of the detectors that will be used to make this measurement can be found later on in this chapter.

Secondly, given that several sets of data will be collected over several input emittances

and momentum ranges, then providing enough statistics have been collected, it will be possible to simulate the cooling effect that the cooling channel would achieve for any arbitrary input beam covered by the input parameterisations. This will prove to be a powerful tool in ascertaining the advantage that a cooling channel will give on a yet unknown beam input emittance; i.e. for the neutrino factory.

MICE must be large enough that it will produce measurable cooling and will allow a variety of configurations to be tried yet economic considerations will limit its length. Financial constraints have led to delays in the construction of MICE although at the time of the writing of this thesis the MICE beamline had been constructed. The planned stages in the construction of the the MICE beamline and cooling channel are discussed in the conclusion after consideration of the complete MICE experiment.

## 3.2 Emittance

This section will give a qualitative description of emittance and its geometrical interpretation. In reality emittance is a complex subject and only an overview of the subject is given here.

In essence emittance gives a measure of how well collimated the particle beam is. Much more than this, an understanding of the beam's emittance is necessary to understand how that beam will both be accepted into and propagate along a section of beam-pipe. A knowledge of the beamline's beta function along with the beam's emittance allows the physical dimensions of the beam to be determined at any point along the beamline.

In fact this last statement serves as a useful tool to introduce emittance. This can be understood further by consideration of the diagram in figure 3.1 [64]. This diagram illustrates a section of hypothetical beamline and a pair of quadrupole magnets are inferred in the diagram where the beam focus is changing.

As is illustrated in the figure the physical size of the beam is given by:

$$E(s) = \sqrt{\epsilon\beta(s)} \quad (3.1)$$

where  $E(s)$  is the beam envelope at position  $s$  along the beam,  $\epsilon$  is the emittance of the beam and  $\beta(s)$  is the beta function. The beta function effectively gives the shape of the beam as a function of position ( $s$ ). A derivation of the beta function will not be considered here although the interested reader is referred to 'The Physics of Particle Accelerators' by Klaus Wille [64].

Emittance is measured by consideration of the beam's phase space coordinates. In 3D this is represented with  $(x, x')$ ,  $((y, y')$  and  $(E, t)$ . The coordinates  $x, y$  represent the transverse position of the particle with respect to some beam axis, and  $x', y'$  represent

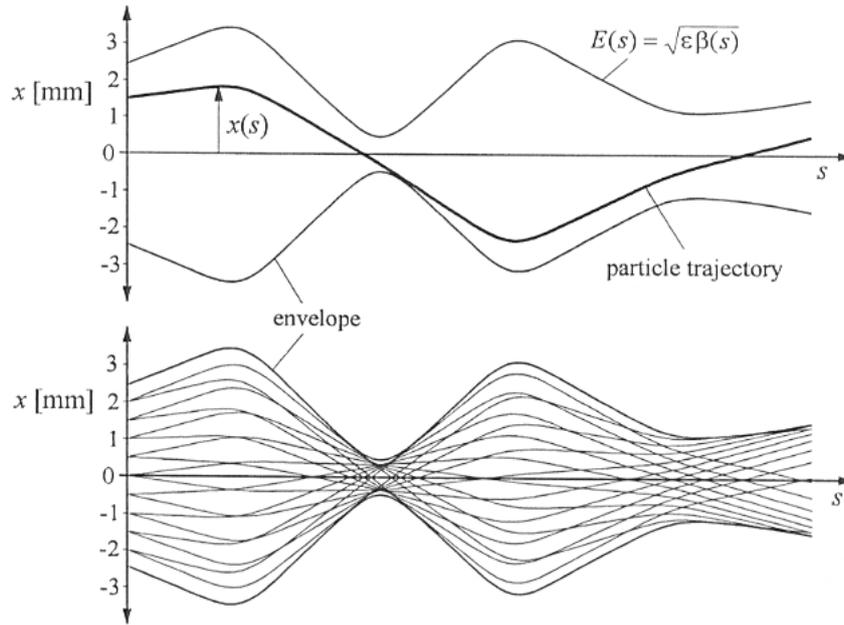


Figure 3.1: Particle trajectories  $x(s)$  within the envelope  $E(s)$  of the beam. The upper figure shows a single particle trajectory whilst the lower figure shows multiple trajectories. The beam envelope is defined from the outermost trajectory  $E(s)$  at each point along the  $s$  axis [64].

$\frac{dx}{dt}$  and  $\frac{dy}{dt}$  respectively. Usually these co-ordinates are considered in 3 pairs and when these are all considered together this leads to consideration of emittance in terms of a 6D phase space. Theoretically the emittance is a constant for a given beam transport system although in real beamlines several effects tend to increase the beam's emittance over time.

Figure 3.2 [64] shows an example of the 2D phase space occupancy for a particle population in terms of  $x$  and  $x'$  at some hypothetical point in an accelerator. This diagram shows an ellipse and this ellipse represents the Full Width Half Maximum (FWHM) distribution of the particles at some hypothetical point within an accelerator beam-pipe. The fact that the particles describe an ellipse is not coincidental, however as will be shown shortly, this geometrical symmetry is useful.

As the particles traverse a beamline the shape of the ellipse changes. In the case where the three pairs of co-ordinates are un-correlated the area of the 2D ellipse will not change. Where correlations do exist between the pairs then either the 4D or the 6D hyper-ellipsoid needs to be considered. With respect to these latter cases conservation of the volume of the hyper-ellipsoid applies as the particles propagate through the beamline.

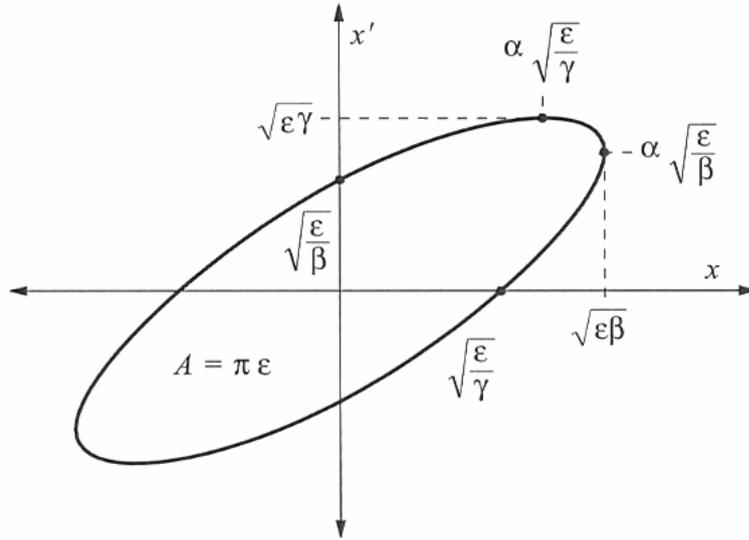


Figure 3.2: The phase space ellipse of particle motion in the 2D  $x - x'$  plane [64].

The conservation in area and/or volume is a consequence of Liouville's theorem which states that the volume occupied in phase space by these particles is constant with respect to time. A direct measure of this constant of area or volume is called the emittance. Now Liouville's theorem is only valid when the system is under the influence of Hamiltonian forces. Within an idealised accelerator this is true as the particles interact conservatively with the magnetic fields, and so the emittance is conserved. In the case where there are non-hamiltonian forces, Liouville's theorem is not conserved. These forces can come from many sources including stochastic processes, inelastic processes and dissipative forces. Ionisation is an inelastic process and as will be shown later in this chapter, in special circumstances, this can be used to reduce the emittance of a particle beam.

The correlations between pairs of coordinates often determines which measure of emittance is both useful and to be used. Where correlations do not exist between pairs of co-ordinates then either the 2D, 4D or 6D emittances can be considered as appropriate. Where correlations do exist then it important that the correct dimensional emittance be measured. As an example of this the tracking solenoids in MICE correlate both the  $x - x'$  and the  $y - y'$  co-ordinates and so the 4D transverse emittance is often considered in MICE when dealing with beam transport. Of course as the aim of MICE is emittance reduction both the transverse and longitudinal emittance are important and so the full 6D emittance measurement is also of critical value in MICE.

### 3.2.1 Geometrical Interpretation of Emittance

The statistical definition of rms emittance is related to the common geometrical interpretation. This is of interest because the way in which emittance is defined for the MICE experiment is elucidated in this discussion. For a real particle beam composed of  $N$  particles the rms emittance is defined from the second order moments of its phase space distribution and this gives its statistical definition:

$$\epsilon_{rms} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \quad (3.2)$$

$$\langle x^2 \rangle = \frac{\sum_{n=1}^N x_n^2}{N} \quad (3.3)$$

$$\langle x'^2 \rangle = \frac{\sum_{n=1}^N x'_n{}^2}{N} \quad (3.4)$$

$$\langle xx' \rangle = \frac{\sum_{n=1}^N x_n x'_n}{N} \quad (3.5)$$

A real particle beam does not have a clearly defined edge but can often be reasonably approximated by a Gaussian distribution. The distribution is defined so that the particles spaced one standard deviation ( $1\sigma$ ) away from the peak intensity define the edge of the beam ( $\epsilon_{rms}$ ). The definition of the emittance as indicated in equation 3.2 can be correlated to an ellipse in  $x - x'$  with an area of  $\pi\epsilon_{rms}$ .

$$\epsilon = A/\pi \quad (3.6)$$

In the 2D phase space  $x, x'$  the general case of an ellipse can be represented using a second order polynomial such that:

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0 \quad (3.7)$$

In the case where the axes coincide with the centre of the ellipse this can be simplified so that:

$$Ax^2 + Bxy + Cy^2 = K \quad (3.8)$$

It can also be shown from consideration of the geometry of the ellipse that the constant on the right hand side of equation is equal to the area of the ellipse divided by  $\pi$  so that from equation 3.6

$$Ax^2 + Bxy + Cy^2 = \epsilon \quad (3.9)$$

In consideration of the 2D  $x - x'$  emittance plot then these coefficients can also be represented in matrix notation; at this point the coefficient notation is also changed and this gives:

$$\beta x'^2 + 2\alpha xx' + \gamma x^2 = \epsilon \quad (3.10)$$

$$(3.11)$$

$$\begin{bmatrix} x' & x \end{bmatrix} \begin{bmatrix} \beta & \alpha \\ \alpha & \gamma \end{bmatrix} \begin{bmatrix} x' \\ x \end{bmatrix} = \epsilon \quad (3.12)$$

There are some key measurements on the ellipse in terms of  $\epsilon$ ,  $\beta$  and  $\gamma$  as shown in figure 3.2. By defining  $X$  as the furthest right hand point on the ellipse it can be shown that:

$$X = \sqrt{\epsilon\beta} \quad (3.13)$$

and by defining  $Y$  as the highest point on the ellipse, similarly:

$$Y = \sqrt{\epsilon\gamma} \quad (3.14)$$

then  $\beta$  and  $\gamma$  can be put in terms of  $X$  and  $Y$ . It can be easily shown from similar considerations that:

$$\alpha = XY/\epsilon \quad (3.15)$$

These can be rearranged and substituted into the matrix from equation 3.12 so that:

$$\begin{bmatrix} x' & x \end{bmatrix} \begin{bmatrix} \frac{X^2}{\epsilon} & \frac{XY}{\epsilon} \\ \frac{XY}{\epsilon} & \frac{Y^2}{\epsilon} \end{bmatrix} \begin{bmatrix} x' \\ x \end{bmatrix} = \epsilon \quad (3.16)$$

which clearly gives:

$$\begin{bmatrix} x' & x \end{bmatrix} \begin{bmatrix} X^2 & XY \\ XY & Y^2 \end{bmatrix} \begin{bmatrix} x' \\ x \end{bmatrix} = \epsilon^2 \quad (3.17)$$

By comparing equation 3.17 to equation 3.2 it can be seen that the statistical definition of emittance can be correlated to the ellipse as the determinant of the matrix defining the ellipse:

$$\epsilon_{rms} = \sqrt{|\det U|} \quad (3.18)$$

where  $U$  is the  $2 \times 2$  matrix in equation 3.17.

Under acceleration the emittance as defined so far isn't conserved. It can be shown that the quantity:

$$\epsilon_{nrms} = \beta_r \gamma_r \epsilon_{rms} \quad (3.19)$$

is conserved under acceleration where  $\beta_r$  and  $\gamma_r$  have their normal relativistic meanings. This is called the normalised emittance.

In MICE the helical path traced out by the muons in the trackers due to the solenoidal field means that  $x - x'$  and  $y - y'$  are correlated. As a result for MICE the transverse emittance is defined in 4D and is given by:

$$\epsilon_{nrms} = \frac{1}{m_\mu c} \sqrt[4]{|\det V|} \quad (3.20)$$

where  $V$  is the 4D covariance matrix of  $x$ ,  $y$ ,  $p_x$  and  $p_y$ . The 6D emittance can be similarly defined for  $x$ ,  $y$ ,  $p_x$ ,  $p_y$ ,  $E$  and  $t$ .

The transformation from  $x'$  to  $p_x$  gives the factor of  $(m_\mu c)^{-1}$  as:

$$x' = \frac{p_x}{\beta_r \gamma_r (m_\mu c)} \quad (3.21)$$

### 3.3 Emittance Reduction With Ionisation Cooling

As can be qualitatively seen from the last section the emittance of the beam is an important consideration in an accelerator design. The transverse emittance along with the beta function determines the physical size of the beam at any point along the beamline. The longitudinal emittance works in a similar manner to give the temporal spread in a bunch of particles that is dependent on the spread in energies of the particles. Longitudinal emittance reduction reduces the emittance by making the beam more nearly mono-energetic.

Both of these measurements are important as the acceptance of an accelerator is both constrained in terms of the maximum physical transverse size of the beam and its longitudinal bunch structure. Particles that do not fall within the machine's acceptance specification will be lost. For a neutrino factory this is a vital consideration as a loss of particles means a loss of neutrino flux. The small cross section of neutrinos means that any method that can increase the acceptance of the muons into an accelerator has a significant effect upon the statistics gained in any detectors that are making measurements on neutrinos produced by this source.

Muon Ionisation Cooling is a method that reduces the emittance of the incoming particle beam such that the flux of muons that is accepted into the accelerator is increased. This increase is not insignificant; as was pointed out in the last chapter it has been calculated that ionisation cooling could be able to increase the flux of neutrinos coming out of the neutrino factory by a factor of sixteen [59].

Ionisation cooling reduces the emittance of the particles in a two step process. Firstly the particles pass through an absorber where each particle loses momentum in all three spatial dimensions. Then the particles are accelerated longitudinally to replace the energy lost in the  $z$  direction. The net result is that the particles show emittance reduction in all three dimensions.

Before describing this in more detail it is instructive to consider the equation that defines cooling as it gives some useful insight. For MICE the rate of change in normalised emittance is given by [65]:

$$\frac{d\epsilon_n}{dz} = -\frac{\epsilon_n}{\beta^2 E} \left\langle \frac{dE}{dx} \right\rangle + \frac{\beta_t (0.014 \text{ GeV})^2}{2\beta^3 E m_\mu X_0} \quad (3.22)$$

here  $\left\langle \frac{dE}{dx} \right\rangle$  is the energy loss and  $X_0$  is the radiation length of the material.  $\beta_t$  is the transverse beta function and  $\beta$  is the normal relativistic factor.

The first term in the equation relates to the energy lost through ionisation cooling whereas the second term is due to heating through multiple coulomb scattering within the absorber. This heating term clearly limits the amount of cooling that can be

accomplished using this method. Equilibrium emittance is achieved when the first term equals the second and so for maximum cooling this implies that the following conditions would be required.

1. A small transverse beta function.
2. An absorber with a long radiation length.

A small transverse beta function is achieved through strong focussing of the beam through the absorbers. MICE achieves this by using super-conducting focusing coils in its design.

The best absorber for ionisation cooling is hydrogen due to its low  $Z$  and therefore its long radiation length  $X_0$ . Liquid hydrogen will be used in MICE as gaseous hydrogen would not provide the necessary  $\frac{dE}{dx}$  given its low density. Provision will be made to allow the absorbers in MICE to be changed for other materials to enable their emittance reducing properties to be experimentally tested. It is not expected that any other absorber material will perform as well as liquid hydrogen.

Ionisation of the particles in the absorber leads to a loss of momentum in all 3D for the incident muon. This is the desired effect. To demonstrate ionisation cooling the momentum in the  $z$  direction will then need to be replaced. High gradient RF cavities capable of operating in large magnetic fields will be necessary to replace the muons' lost momentum in the  $z$  direction. This loss of momentum in all 3D by the absorber followed by the replacement of the momentum in the  $z$  direction by the RF cavities is fundamentally how the emittance of the particle is reduced. Effectively this process makes the beam parallel, collimating the muon beam and therefore reducing its emittance.

### 3.4 The ISIS accelerator

The Rutherford Appleton Laboratory (RAL) is home to ISIS, a 800 MeV (1.4 GeV/c momentum) proton synchrotron accelerator that is one of the world's most intense muon and neutron sources. The low energy of the proton beam precludes much high energy particle physics work and the accelerator is mainly used for material science research such as crystallography. Figure 3.3 [66] shows an overview of the ISIS accelerator facility.

The proton beam starts at an ion source that generates  $H^-$  ions for insertion into a Linac. The Linac accelerates these ions from 0.7 MeV to 70 MeV in four stages. Once the  $H^-$  ions have been accelerated to 70 MeV by the linac they are then passed through a stripper as they inserted into the main synchrotron. The stripper removes the two electrons from a  $H^-$  ion leaving a proton.

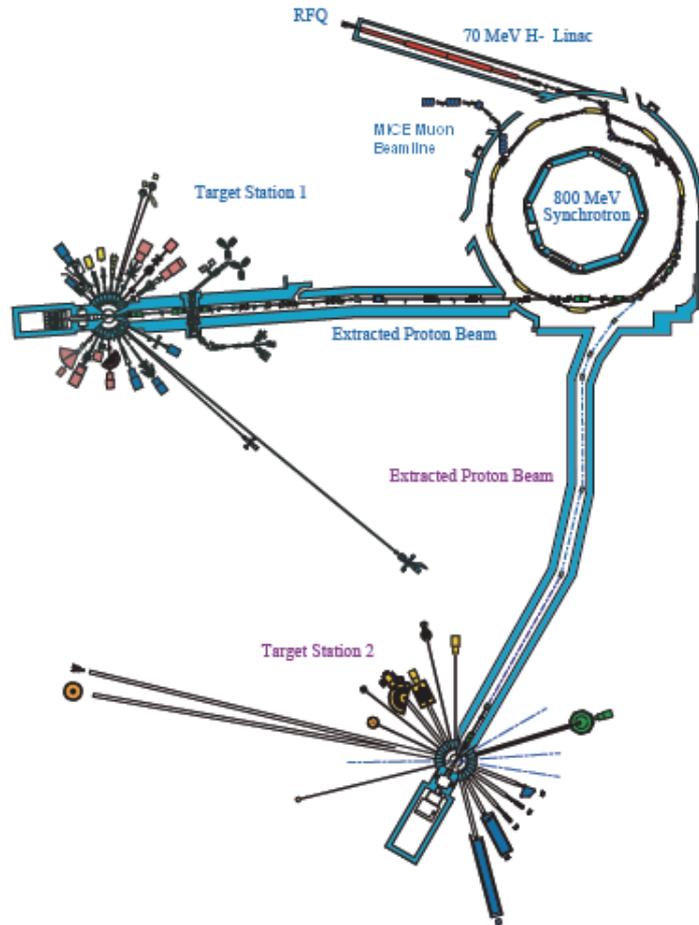


Figure 3.3: An overview of the ISIS accelerator facility [66].

The protons are accelerated from 70 MeV to 800 MeV as they make approximately 15,000 circuits of the synchrotron. The acceleration is done by six Radio Frequency (RF) cavities placed at various locations on the circular orbit around the synchrotron. As the protons are accelerated the currents through the dipoles and quadrupoles are ramped up increasing the magnetic field strength to compensate for the higher energy of the protons. A powerful kicker magnet then deflects the protons out of the synchrotron and towards one of the two target areas. Neutrons are produced by spallation from a water-cooled tungsten target when the proton beam is directed onto it.

ISIS operates at a frequency of 50 Hz. A spill of particles is accelerated and dumped in a 10 ms cycle, allowing an additional 10 ms between spills for the synchrotron magnets to be ramped down to their injection currents. As the beam is accelerated the beam

cross-section shrinks in size. The amount of physical shrinkage observed clearly depends upon the beta function of the beam at the point where the measurement is taken and so varies from point to point around the ring.

The cross section of the ISIS beam can be measured using gas ionisation monitors that are located around the ring at ISIS [67][68]. These gas ionisation monitors provide beam profile information by measuring the positive ion current from the proton beam's residual gas interactions. These monitors are placed at fixed positions around the ring because they need line of sight access to the beam itself; unfortunately this does mean that they cannot be easily moved to other parts of the synchrotron ring. The gas ionisation mechanism allows a profile of the core of the beam to be made but due to the mechanism that the profile is built up gas ionisation monitors tend to be slow and can take up to several minutes to build up a profile. This means that small pulse to pulse variations in the beam are averaged out. The fixed location of these monitors prevents them from being used to obtain information on the beam cross-section at the point where the target is located. The low resolution of these devices means that there is little reliable information on the proton density in the halo around the beam.

The core of the ISIS beam shows a gaussian profile. The beam centre is known to show long term drift due to changes in the linac and ion source conditions. This is monitored by the ISIS control crew through the use of beam position monitors.[69]

## 3.5 The MICE Cooling Channel

Figure 3.4 illustrates the major components of the MICE cooling channel. The rest of this chapter concentrates on describing these components and their purpose within the context of the experiment. The MICE beamline, although not shown in this illustration, is clearly a critical component of the experiment and will also be discussed.

### 3.5.1 Diffuser

The diffuser has been placed at the end of the MICE beamline to diffuse the incoming muon beam before it enters the cooling channel. This diffusion process has the effect of increasing both the input beam's emittance and its cross-sectional size. The amount of emittance increase for a given diffuser material is a function of its thickness and its density. Lead, with its high atomic weight is ideal for this purpose and has been chosen as the material to be used in the MICE diffuser. As MICE requires muon beams of different input emittances a mechanical design was required that would not only support the diffuser but would also allow it to be interchanged.

Remote interchangeability of the diffuser has been achieved by using an electromechanical carousel system. This system allows for one of four different diffusers to be independently selected and inserted into the MICE beamline at any one time. It is also

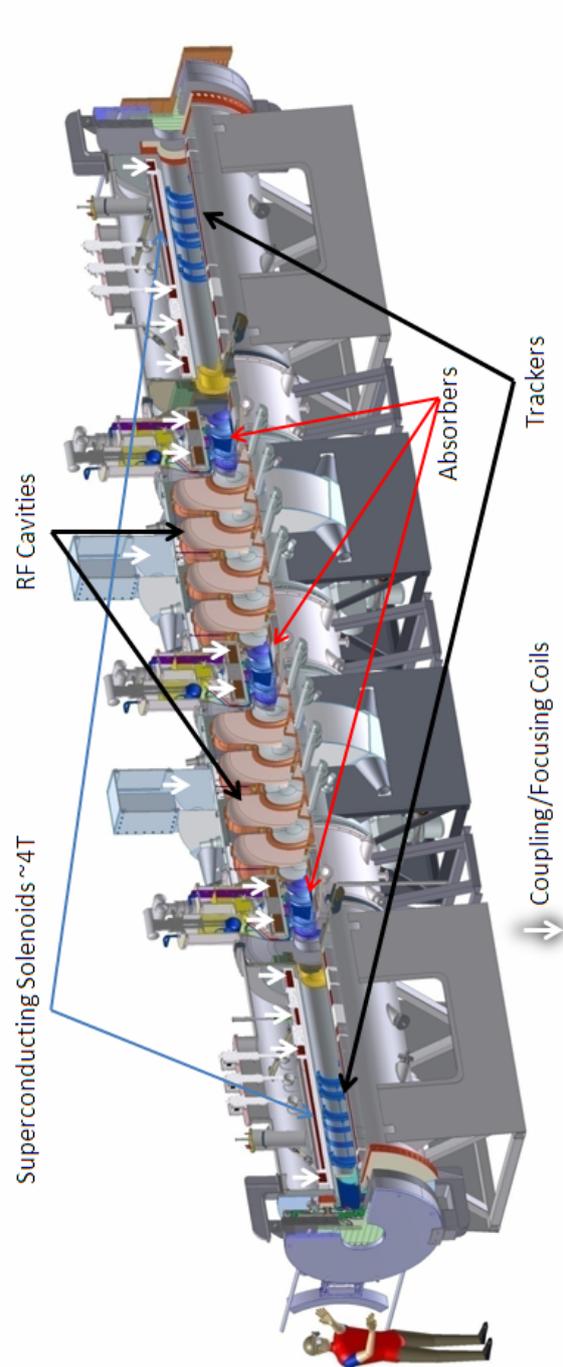


Figure 3.4: A 3D drawing of the MICE cooling channel showing the position of the magnets, trackers, absorbers and RF cavities. The diffuser and TOF1 are at the front of the MICE machine whilst the calorimeter and TOF2 are at the rear. The white arrows represent the location of the coupling/focusing coils [70].

possible to have no diffuser installed thereby giving an input beam with the minimum possible emittance for any given beamline settings. An illustration of the carousel diffuser system that has been designed for the MICE experiment is shown in figure 3.5 [70].

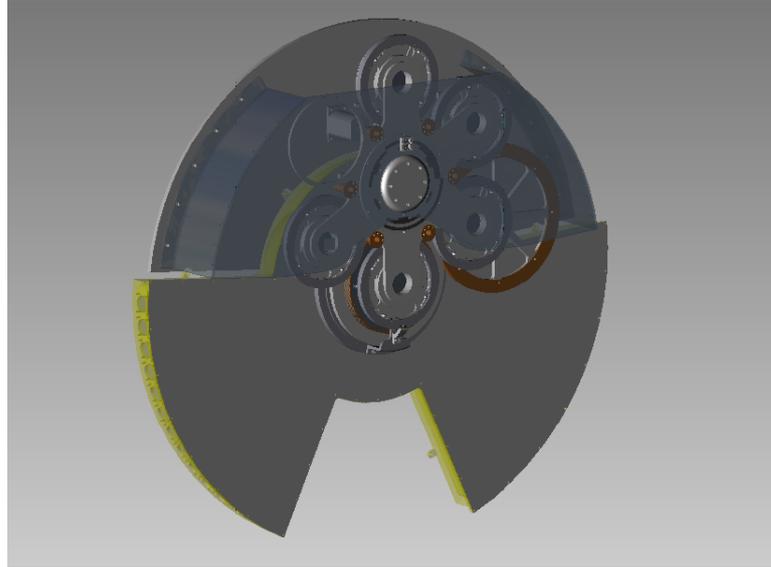


Figure 3.5: The MICE diffuser. A carousel system installs lead discs of different thicknesses into the MICE beamline causing an increase in the beam's emittance. This allows MICE to measure the amount of emittance reduction that can be achieved on beams with different initial emittances [70].

### 3.5.2 Emittance Measurement and Particle Identification in MICE

Precise measurement of the emittance requires both accurate particle identification and detectors that are capable of allowing the particles to be tracked accurately. MICE will take particle by particle measurements in order to build up a full picture of the 6D emittance. The detectors need to be able to identify the muons from any other particles that travel through the cooling channel with greater than 99.9 % success. Only at this level will the emittance measurement be possible to the required degree of accuracy. A much more comprehensive description of the detectors for the MICE experiment can be found in the MICE 'Technical Reference Document' [65]. The following sections are included to provide the reader with the necessary overview.

#### Time Of Flight Detectors TOF

MICE will be provided with three time of flight detectors called TOF0, TOF1 and TOF2 respectively. The TOF0 and TOF1 detectors will be placed upstream of the

cooling channel whilst TOF2 will be placed after the cooling channel. All three TOF detectors are designed to assist in providing muon identification. TOF0 and TOF1 will also provide the necessary trigger for MICE whilst TOF2 will help to eliminate muons that have decayed whilst traveling through the cooling channel from the emittance measurement statistics.

Each TOF counter is constructed from two planes of scintillating bars lying perpendicular to each other. In TOF0 there are 12 bars per plane whereas in TOF1 and TOF2 there are 8 bars per plane. Each bar is connected to a photo-multiplier tube via a light guide. A coincidence in two planes of the TOF counter will give a time resolution in the TOF counters of approximately 70 ps.

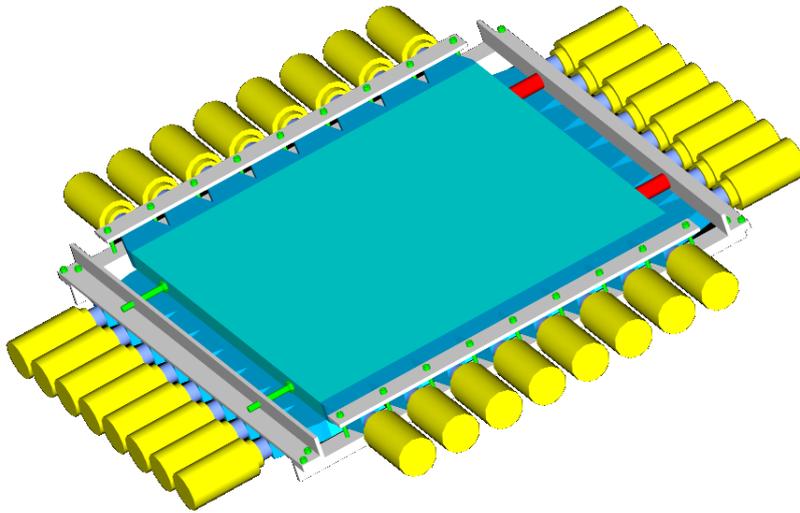


Figure 3.6: Diagram of one of the 8 bars per plane TOF stations including some of the mechanical supports [65].

### Cherenkov Detectors for Particle ID

Given the importance of obtaining high accuracy identification of particles that enter into the cooling channel it is necessary to provide some redundancy in the particle detection mechanism. This mechanism is provided for in MICE by a pair of Cherenkov detectors. Alongside the TOF detectors these detectors will help to discriminate muons, pions and electrons.

The Cherenkov detectors are placed one after the other and are situated upstream in the MICE beamline between the the two sets of triplet quads Q4-Q6 and Q7-Q9. These detectors use the property of Cherenkov radiation to discriminate between the pions and muons of a given momentum that are entering into the cooling channel.

Cherenkov light is produced as a particle passes through a medium at a speed greater than light. By choosing two mediums that have different refractive indices it is possible to discriminate between particles because their different velocities within these respective mediums at a given momentum gives a different light yield. The Cherenkov light is detected by four photomultiplier tubes that surround each detector. The problem is that the momentum range covered by MICE makes it impossible for any single Cherenkov detector to reliably discriminate between muons and pions. In MICE this problem is solved by using two detectors.

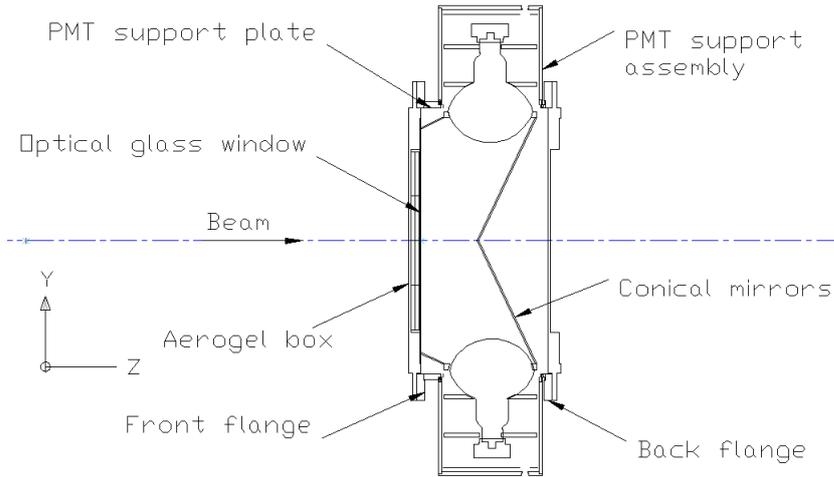


Figure 3.7: Cross-section through one of the two Cherenkov detectors. Note the aerogel box on the front of the detector [65].

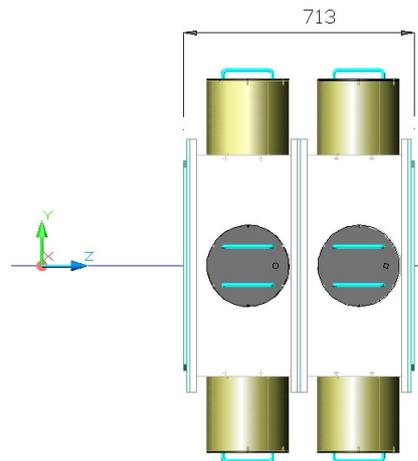


Figure 3.8: the whole Cherenkov1 (CKOV1) device is constructed by putting two near identical detection units one after the other along the beam line upstream of MICE. Each unit contains aerogel with a different refractive index [65].

The way that this works is illustrated in figure 3.9 [65]. By using aerogel in the detectors with refractive indices of 1.07 and 1.12 respectively this defines three momentum regions as illustrated in the figure. For incident muons in the first momentum range of 207 MeV/c to 276 MeV/c only unit one will fire when the muons pass through both detectors. In the second momentum range of 276 MeV/c to 365 MeV/c both units will generate a trigger for muons only.

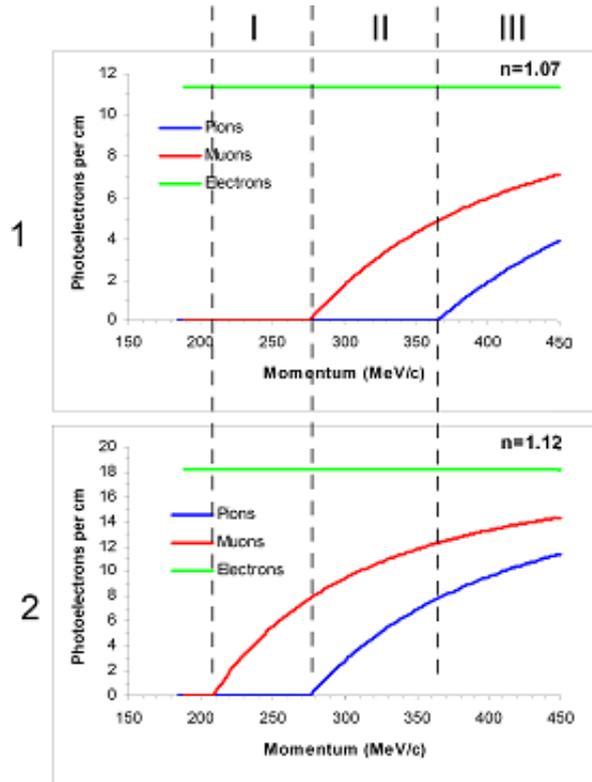


Figure 3.9: Particle Momentum vs Photoelectron density for the two Cherenkov detectors with different aerogel. The different triggering behaviour for the Cherenkov detectors is indicated by the regions I, II and III [65].

At momentum above 365 MeV/c both muons and pions will generate a trigger in both units and so will be indistinguishable. However as MICE is only interested in muons up to an energy of 350 MeV/c this is not considered an issue as the tracker will be able to eliminate these particles. Both of the Cherenkov detectors are sensitive to electrons at all momentum ranges of interest, however the TOF, tracker and calorimeter will be able to eliminate these from the beamline analysis.

### The Calorimeter

A calorimeter is to be placed at the end of the MICE beamline. The principle aim of the calorimeter is to provide electron-muon identification. Most of the electrons will be able to be separated out from the beamline using the other detectors however for the muons that decay in flight this calorimeter will provide a useful tool to accurately detect these decays.

The calorimeter is based upon a design that was used for the KLOE [71] experiment. It is constructed from 180 lead and scintillating fibre layers. The calorimeter is built from four layers, each layer being composed of 18 ‘cells’ that measure  $4\text{ cm} \times 4\text{ cm}$  wide. This creates a detector that is 16 cm thick and has an active area of  $72\text{ cm} \times 72\text{ cm}$ . The scintillating fibres are bundled and led out of the detector via light guides to photomultiplier tubes. The response of the calorimeter to both electrons and photons in the energy range of 20 MeV to 300 MeV is linear. Typically the signal deposited by a minimum ionising particle in a single cell is the same as that deposited by a 30 MeV electron.

It is expected that electrons will deposit most of their energy within the first two cells whilst the muons will pass through the calorimeter. In combination with the other detectors in MICE the calorimeter should be able to provide an electron rejection of 99.9 % allowing accurate identification of muons from electrons.

### Trackers

The trackers in MICE are used to measure the emittance of the individual muons that pass through the cooling channel. MICE will contain two trackers, each one sat in a 4 T solenoidal field, one before and one after the absorbers and RF cavities. The pair of trackers will work in a complementary fashion in order to measure the muon emittance before and after the muon has been cooled. Clearly these detectors will be essential in ascertaining the amount of muon cooling achieved.

The trackers are based on scintillating fibre technology. Each tracker contains five sets of fibres called stations and each station is separated by several cm. Each station is composed of three layers of  $350\text{ }\mu\text{m}$  fibres orientated at 120 degrees to each other. The fibres are bundled together in groups of seven and fed out to via light guides to Visible Light Photon Counters (VLPCs). The VLPCs are cryogenically cooled and have the high quantum efficiency needed to detect the small light output obtained from the fibres created by a passing muon.

As the muons pass through the solenoidal channel they describe a helix. Points on this helix are registered by the tracker stations as the particle traverses them. These tracks are then reconstructed using a piece of software called a Kalman filter. This program aims to minimise the mean squared difference between the measured hits in the tracker stations and a fitted helical path. Many simulation studies have been completed on

the trackers using a software package developed for MICE called G4MICE. This has shown that the tracker is capable of giving a good enough resolution to allow the muon emittance to be measured to within 0.1 % over the range of momenta of interest to MICE [72].

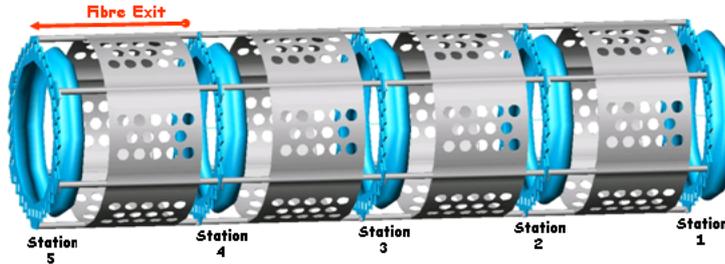


Figure 3.10: A schematic view of one of the MICE scintillating fibre trackers. The trackers are inserted into the long solenoidal magnets with a field of  $\approx 4$  T [65].

### Absorbers

The three MICE absorbers are effectively removable containers for holding 21 litres of liquid hydrogen. As discussed in section 3.3, liquid hydrogen provides the best absorbing material for muon cooling due to its minimal multiple coulomb scattering effect although the removable design of these modules does mean that other absorber materials, such as Beryllium, can be tried by MICE.

Careful consideration has been given to the design of the absorbers as the use of liquid hydrogen in the proximity of RF cavities is clearly a safety concern. The absorbers have specially strengthened windows and are surrounded by a large volume vacuum system so that any potential leaks can be safely contained. A cross-sectional view of the MICE absorber focus coil module showing the liquid absorber is illustrated in figure 3.11 [73].

### RF Cavities

The MICE RF cavities are to provide the necessary acceleration to replace the longitudinal energy lost by the muons in the absorber modules. Each cavity is made from copper and has a radius of 610 mm and a length of approximately 430 mm. The cavities are sealed using beryllium windows; these help both the electromagnetic properties of the cavities and keeps the cavities isolated from the hydrogen system of the absorbers, which are in close proximity. The RF cavities will operate at a frequency of 201 MHz and will initially have a peak field gradient of  $8 \text{ MVm}^{-1}$ .

MICE will contain eight cavities set in two blocks of four. A block of four cavities can be seen in figure 3.12. In the final MICE setup two sets of cavities will be situated between three absorbers; this setup is shown in figure 3.4.

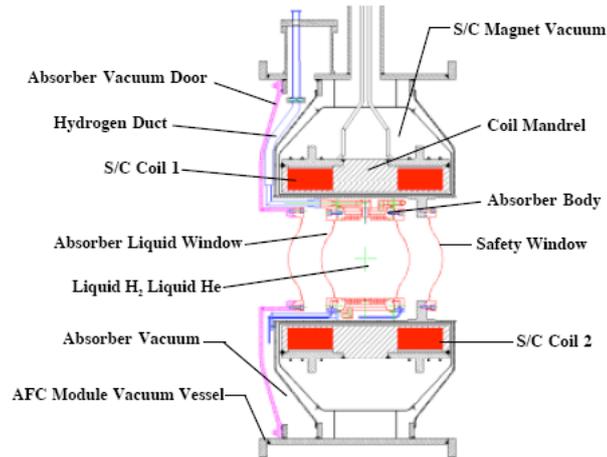


Figure 3.11: A cross-sectional view of the MICE absorber focus coil module illustrating the liquid absorber [73].

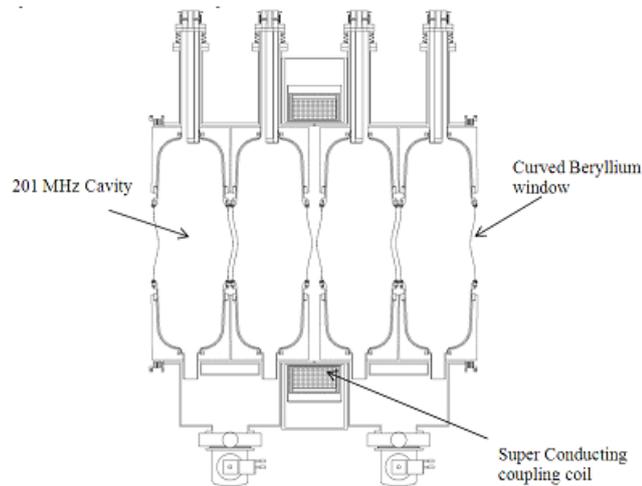


Figure 3.12: Four RF cavities form one RF assembly. There will be two such assemblies in MICE sandwiched between three absorbers [65].

As MICE RF cavities will be working in close proximity to super-conducting coils they will be working in a high magnetic field gradient ( $B$ ). A high  $B$  field gradient is known to limit the maximum electric field ( $E$ ) gradient that can be achieved within a cavity. As this is not a well understood phenomenon MICE have been testing a similar RF cavity at Fermilab to try and understand the effect of large  $B$  fields on cavity operation. To date these tests have shown that operating a cavity in a  $B$  field similar to that expected in MICE reduces the maximum  $E$  field gradient achievable compared to no  $B$  field by about half [74].

### 3.5.3 The MICE magnets

The MICE cooling channel contains 18 solenoidal super-conducting magnets. The purpose of 16 of these magnets is to control beam transport through the cooling channel. The other 2 magnets surround the trackers and give the passing muons a helical path so their emittance can be measured by the tracker - See figure 3.4 [70]. Broadly speaking the 18 magnets can be categorised into three sections; ten tracker solenoid coils, two RF coupling coils and six absorber (focusing) coils.

Each tracker solenoid is surrounded by four super-conducting magnets. Two of these coils provide field correction at the end of the solenoid to ensure that the field uniformity throughout the tracking channel is better than 1 %. The other two coils are there to match the beam optics either into or out of the tracker itself. Due to the different functions that these magnets provide each magnet requires its own power supply so that the current through it can be individually regulated.

The RF coupling coils consists of two coils, one that surrounds each section of four RF cavities. These are large coils with an inner diameter of over 1.3 m and the coils are to provide focussing of the beam as it is transported through the RF cavities. Both coupling coils will be operated in series from a single 10 V 300 amp power supply.

Each set of absorber coils is constructed from a pair of super-conducting coils separated by 240 mm and surrounding each absorber module. The aim of these focusing coils is to provide strong focusing to minimise the beta function of the beam within the absorber modules. The close proximity of these magnets means that they exert huge forces, up to 2.3 MN on each other. These forces must be carried by the frame supporting these magnets.

Quench protection of these magnets is passive and is done through a series of diodes and resistors. The quench protection is important because if one of the coils goes normal the close proximity of the other magnets tied with the inductive coupling is likely to quench the other magnets.

### 3.5.4 The MICE Beamline

For MICE it is necessary to have a dedicated muon beamline for the experiment. The principal role of the beamline is to provide a matched muon beam into the MICE experiment where the beam has well understood properties and minimum contamination from non-muonic particles. The beamline is illustrated in figure 3.13 and can be considered as two separate parts. The first part of the beamline is responsible for pion production, capture and decay. The second part of the beamline matches the muons into the MICE experiment.

The beamline starts with the target mechanism. The target mechanism is inserted into the ISIS beam timed to intercept the last couple of ms of an ISIS spill. The target

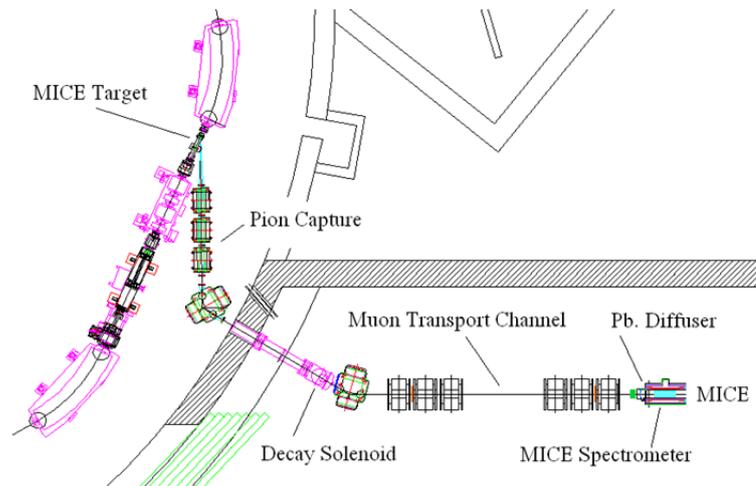


Figure 3.13: The MICE Beamline

does not intercept every spill but as a baseline it is expected to intercept one in fifty ISIS spills, i.e. it should operate at 1 Hz. This requirement was initially imposed due to the frequency at which the RF cavities within MICE could be operated although as will be described in later sections, the demanding acceleration required from the target also limits its operation to approximately 1 Hz thereby constraining the rate at which particles can be inserted into the MICE beamline.

Pions produced by the target are collected via a quadrupole triplet. These particles are then passed through the first beamline dipole that serves to both steer the beam towards the MICE hall and momentum select the particles so that a fairly clean beam of pions enters the pion decay solenoid. A second dipole is then placed after the decay solenoid to momentum select muons from the remaining pions and protons thereby giving a clean muon beam. The muons then pass through another two pairs of triplet quadrupoles, their purpose is to match the muon beam into MICE.

### 3.6 Conclusion

MICE is a complex experiment that will be constructed in a modular fashion in several stages over a period of several years. The main components of each of these stages are illustrated in figure 3.14. The modular construction of the experiment offers the advantage that many of the detectors can be progressively characterised and understood before the complexity of the emittance reduction and RF acceleration stages are added. This will be necessary if an accurate and precise emittance measurement is to be obtained over several different input beam momenta and emittances.

Stage one represents the MICE beamline and installation of this was completed in January of 2008. Much of the first part of 2008 used the newly installed target mechanism

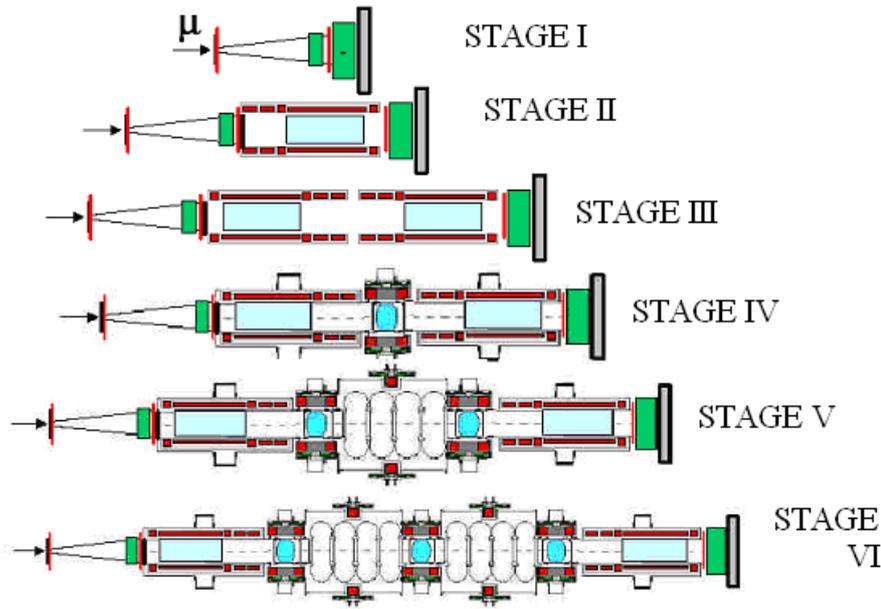


Figure 3.14: The six stages of the MICE experiment [65].

to commission the beamline. Stage two commenced during the summer of 2008. As of the spring of 2009 the Cherenkov detectors have been installed alongside TOF1, however the tracker is still awaiting the manufacture of the tracker coil before it too can be installed. Once the installation of stage one has been completed this will allow the first measurements of the beam input emittance. The addition of the tracker along with the calorimeter and the TOFs will give MICE the ability to perform some particle identification. Stage three will see the addition of the second tracker. By step four the addition of an absorber will allow emittance reduction to be demonstrated. However the ability to re-accelerate the muons will have to wait until step five. Only with the addition of three absorbers and the two sets of RF cavities in step six will MICE be able to demonstrate full muon cooling with an emittance reduction of up to 10 % and an emittance measurement accuracy to within 1 %.

## Chapter 4

# An Introduction to the MICE Target Mechanism

### 4.1 Introduction

The MICE target mechanism has been built to provide a source of muons for the MICE experiment by operating parasitically on the ISIS accelerator at the Rutherford Appleton Laboratories (RAL) in the UK. The design for the target mechanism is based upon an electromagnetic linear motor that is capable of accelerating a small titanium target into the ISIS beam halo during the last 2 ms of every fiftieth ISIS beam cycle. The timing of the insertion of the target into the proton beam is crucial for producing pions with the correct energy spectrum for the MICE experiment.

This chapter considers the motivation for the design of the target mechanism. After considering the background on the use of a parasitic target at ISIS the specification for a new target drive is given. From this specification the use of rotary and linear mechanisms are critically analysed and the justification for selecting a linear motor is given. Key aspects of the drive's design, the bearings and the motor, are discussed. Tests of the preliminary choices for bearings that led to the selection of sliding bearings are then described. Subsequent further evolution of these bearings are described in Chapter ten.

### 4.2 Background

A target mechanism running parasitically on the ISIS synchrotron is not a new idea. Prior to the development of the MICE target mechanism an oscillating target was run on the synchrotron producing particles for the then named HEP PPD beamline. This old target mechanism was situated on the same section of the synchrotron as the MICE

target is now placed, super-period seven, and the old HEP beamline ran into what is now the MICE hall.

The HEP PPD mechanism[75] was composed of two oscillating levers where the second lever was physically attached to the end of the first. The whole arrangement is illustrated in fig 4.1. One of the levers was oscillated at 50 Hz and the second at 150 Hz. The superposition of the motion of the two levers created a motion as illustrated in figure 4.2. This motion was able to insert the target into the ISIS beam once a spill, i.e. at 50 Hz.

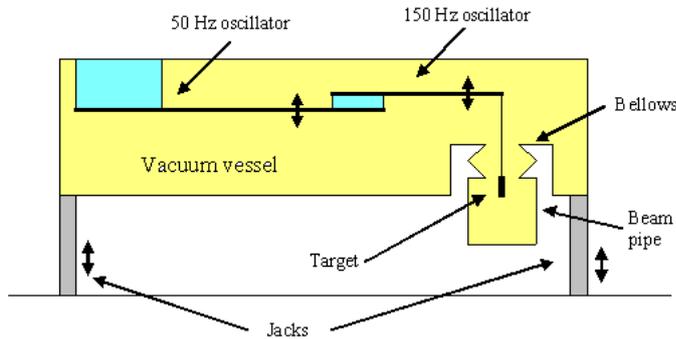


Figure 4.1: Schematic diagram of the old HEP target. The target was suspended from two arms, the first arm was oscillated at 50 Hz and the second arm was oscillated at 150 Hz.

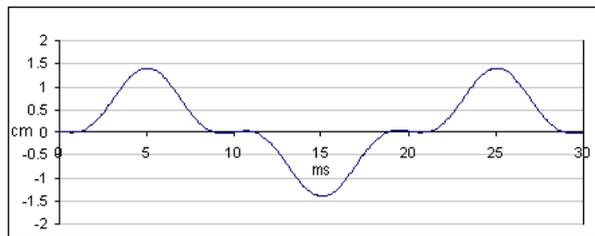


Figure 4.2: Simulated target insertion. Superimposed 50 Hz and 150 Hz oscillations allow the target to dip into the ISIS beam at 50 Hz

The target was operated for several years but suffered from several problems. In particular the target system took a significant amount of time to stabilise and it was easy for the system to cause a beam trip by dipping into the ISIS beam too far (variations in dip amplitude). When ISIS was upgraded to improve its beam intensity it became more difficult to adjust the target and the machine became much more sensitive to beam variations. This behaviour was attributed to a decreasing halo around the beam. The target was designed to dip at 7-8 ms into the acceleration cycle and was aided by a beam bump. (This was not believed to be possible at 9-10 ms, the time when the MICE target will be inserted, as it would interfere with extraction[76].) It was clear the

MICE experiment needed both a new target mechanism and a beamline to match, one that was designed to specifically provide the MICE experiment with the muon beam that it requires.

### 4.3 MICE cooling Channel Requirements

Before outlining the kind of device that would be required to produce muons for the MICE experiment it is worth briefly considering the MICE cooling channel requirements. The MICE beamline will be required to provide at least nine different matched muon beams into the MICE experiment. These nine conditions will be nominally met by providing three different momentum beams each at three different emittances. The baseline is to provide muons with momentum of 140 MeV/c, 200 MeV/c and 240 MeV/c with emittances of  $1\pi$ ,  $6\pi$  and  $10\pi$  mm mrad.

In order to maximise the number of good muons of the correct momenta going through the MICE cooling channel it is necessary for the target to intercept the ISIS beam when the protons have been accelerated towards their maximum momentum of 1,400 MeV/c. To do this the target must intercept the ISIS beam during the last 2 ms of the ISIS cycle (Also see figure 9.1 in chapter nine). Many simulations have been done by the MICE beamline group to ascertain the muon rate into MICE using several independent codes.

The target itself is a cuboid of titanium measuring 1 mm wide by 10 mm long by 35 mm tall. These dimensions are partly historical and partly from the results of simulation. The old HEP beam target was nearly cuboid and measured approximately 1 mm wide by 20 mm long by 8 mm tall. It was initially envisaged that the MICE target would need to fit into the same slot in the beampipe as the HEP PPD target, and this was only 2 mm wide. As it is the volume of material in the beam that dictates the particle production rate, the change in dimensions results in the target particle production rate being less sensitive to slight changes in dip depth but this does come at the cost that a deeper penetration into the beam is required for higher particle production rates. ISIS accelerator produces  $2.5 \times 10^{13}$  protons per spill at a rate of 50 Hz during normal operation. Based upon these assumptions simulations were performed in 2004/2005 to estimate the good muon rate in the MICE cooling channel. It was estimated that it would require  $1.4 \times 10^{12}$  protons to interact with the 10 mm length of the target in order to generate 600 good muons per intercepted spill for the MICE experiment[77][78]. (A good muon is a muon that makes it through both the MICE cooling channel and the detectors so that an emittance measurement can be made on it.)

#### 4.4 Operational Requirements and the Physical Constraints with respect to Operating the Target Mechanism

The target mechanism baseline design was determined in 2003 after a series of meetings with ISIS staff[76][79]. The design requirements were:

- Enter last 1-2 ms of beam.
- Not disturb next injection.
- Variable dip depth.
  - Required entry into beam is unknown but  $\approx 5$  mm.
  - Beam Shrinkage 17.8 mm.
  - Travel  $\geq 25$  mm.
  - Reproducibility  $\approx 0.2$  mm.
- Accurately synchronised to ISIS, drift/jitter  $\leq 0.2$  ms.
- Baseline frequency of 1 Hz.
- Target titanium 1 mm across beam, 10 mm along.
- Slit across beam pipe 2 mm wide max.
- Radiation hard, UHV materials.
- Resonant system disfavoured. On demand actuation preferred.

Figure 4.3 illustrates the motion required of a target tip based upon what was known of the ISIS beam profile from tests in 2006. The key point to be taken from this diagram is that the motion of the target tip has to be accelerated sufficiently quickly to overtake the shrinking beam so that it can be inserted into the ISIS beam (by a variable amount) during the last 2 ms of the beam cycle. Subsequent to its acceleration into the beam the target would need to be accelerated away from the volume required by the beam envelope by the start of the next cycle. This diagram was drawn from the perspective of a linear drive, therefore the detail of the trajectory may vary somewhat with an alternative mechanism; however the general operational constraint remains the same.

#### 4.5 Possible Designs for the Target Mechanism

Various mechanisms are considered with respect to the design requirements reproduced in the last section. All of the mechanisms considered utilise electromechanical drives. The use of conventional kicker magnets to steer the ISIS beam onto a stationary target

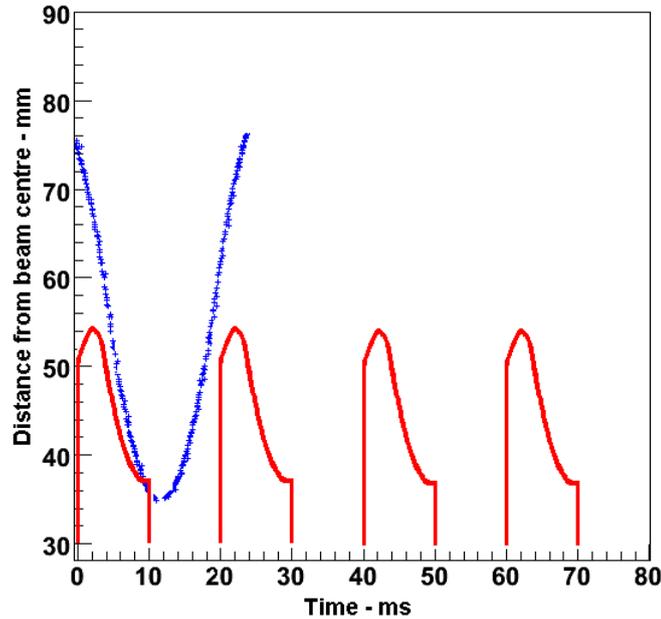


Figure 4.3: The trajectory of the target mechanism (blue) must be capable of accelerating past the shrinking beam (red) so that it only intercepts the beam during the last 2 ms of the acceleration cycle. The target must be clear of the beam envelope by the start of the next ISIS cycle and must remain clear of the beam until the next actuation  $\approx 1$  s later.

had been dismissed as the original design brief specified that the target had to be developed as an isolated stand-alone system that could be bolted to and operated parasitically on ISIS so that MICE could use the old PPD beam hall. Additionally, attempting to electromagnetically steer the beam onto a stationary target so close to beam extraction was rejected as would likely have an adverse effect on the ISIS beam[76].

#### 4.5.1 Resonant Flexure Mechanisms

A re-design of a vibrating sprung mechanism such as the one that had been used on the old HEP PPD line had also been discounted for two reasons. Firstly this was a resonant system and so it could not be operated on demand and its frequency of operation was fixed. Secondly the resonance system had its own problems with variation in amplitude of dip and temporal drift. For these reasons, and as previously discussed, it was specified that a new and novel non-resonant mechanism should be sought.

## 4.5.2 Rotary Mechanisms

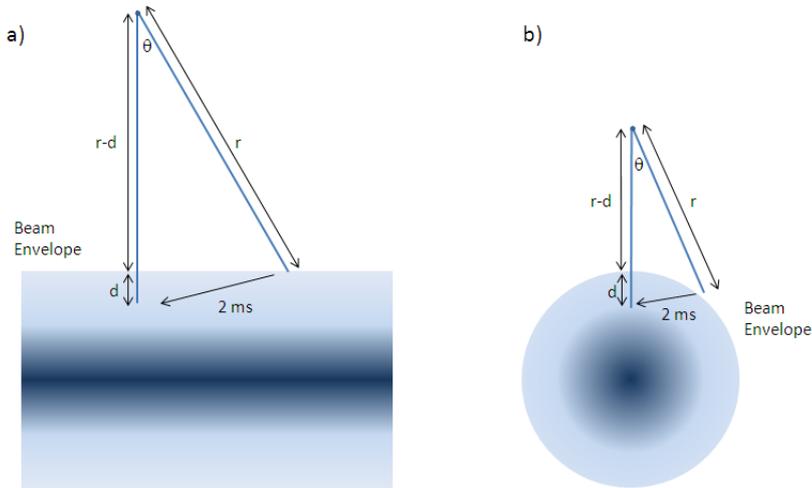


Figure 4.4: Diagram illustrating the use of a continuously rotating arm to intercept the ISIS beam. The case for intercepting the beam with a rotation that is perpendicular to the beam's direction (b) is much better than for the case where the arm rotates in the direction parallel to the beam (a). Both cases result in a long arm due to the small value of the angle  $\theta$  and creates a moving production point for particles.

It is perhaps instructive to first consider the possibility of using a rotary mechanism to achieve the design aims and to ask what the implications of using such a mechanism are. If a rotary mechanism is a viable idea then this would be encouraging as this type of mechanism is well understood and in principle should be easier to implement.

There are at least a couple of examples of rotary mechanisms being used to intercept a particle beam. The 1 MW continuous proton beam from the 590 MeV PSI cyclotron is aimed at two rotary targets that are placed in series in the beamline[80]. The first target is a solid disk and the second is effectively a solid annulus although in the latter case there are slits in the annulus to permit relief from thermal stresses due to the higher beam loading that this target receives. Both of the PSI targets rotate at 1 Hz, the purpose of which is to reduce both the thermal loading and the radiation damage to the devices. Because these devices are subject to continuous beam the purpose of the rotation is significantly different to that which would be required for the MICE target, namely ensuring that the ISIS beam is only intercepted at the appropriate time. As will be demonstrated in this section utilising a rotary mechanism to provide this functionality would be a radically different problem and as such consideration of the PSI design is inappropriate for the MICE target. A rotary target was also used in the Fermilab Anti-proton Source[81]; like the PSI targets the purpose of the rotation was to reduce heat damage and extend the lifetime of the target. A rotary beam chopper has at least been suggested as a source of positrons for a future ILC [82], however the higher frequency at which the beam would be chopped would give a significant advantage for

a rotary solution in this case.

It is possible to imagine such a system where a counterweighted arm is rotated through an opening in the beam-pipe and parallel to it in order to intercept the beam at a frequency of 1 Hz. Therefore if the angular velocity of the rotor is a constant  $2\pi \text{ rad s}^{-1}$  and that the insertion time is  $t$  ms, it can be shown that the radius of the arm would have to be equal to  $r = d/(1 - \cos \theta)$  where  $\theta$  is the angle subtended in time  $t$  ms, i.e.  $2\pi t$  and  $d$  is the final insertion depth into the beam - see figure 4.4. This assumes no beam shrinkage during the last couple of ms which is a reasonable approximation (see figure 9.5 in chapter nine). If  $r$  is plotted against the insertion depth  $d$  one can immediately see that this solution is impractical for all dip depths because the arm is unreasonably long - see figure 4.5.

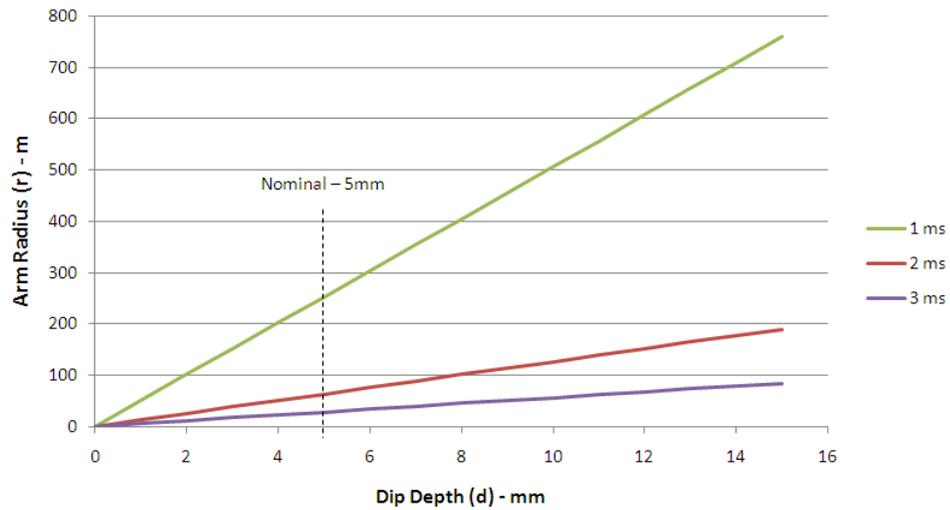


Figure 4.5: The minimum arm radius for a constant speed rotary target drive as a function of dip depth for three different insertion times. The motion of the arm is parallel to the beam pipe. One can see that this leads to a totally impractical solution.

The situation appears to improve dramatically if one attempts to cut perpendicularly across the beam, because the curvature of the beam envelope shortens the distance that the arm needs to be in the beam, allowing a much smaller arm radius. Firstly if it is assumed that  $r \gg R$  (where  $r$  is the arm length and  $R$  is the beam radius) then to a first approximation the motion of the arm across the beam can be approximated by a straight line at a constant distance of  $R - d$  from the beam centre - see figure 4.6. This approximation can be justified because of the small value of  $\theta$

This gives the distance  $x$  as

$$x = \sqrt{(2Rd - d^2)} \quad (4.1)$$

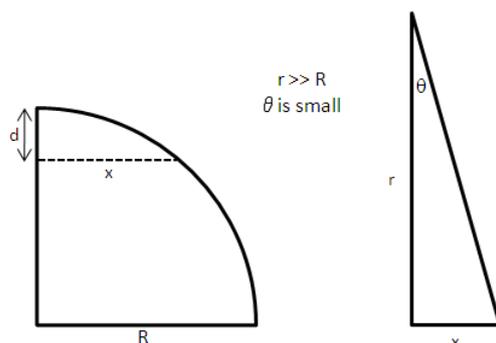


Figure 4.6: If  $r \gg R$  then the arm will to a first approximation travel in a straight line through the beam. The path is indicated by  $x$ .

As previously, the rotation rate is 1 Hz so

$$\theta = 2\pi t \quad (4.2)$$

This leads to the approximation that:

$$r \approx \frac{\sqrt{(2Rd - d^2)}}{\tan(2\pi t)} \quad (4.3)$$

One can see from figure 4.7 that using an arm that cuts transversely across the beam gives an arm radius that is much more reasonable than for an arm that cuts parallel to it, although the arm length is still on the large size for a practical mechanism particularly as the required dip-depth was not known at this time and so the arm length would have had to be chosen pessimistically.

The use of this mechanism would require making a significant cut into the beampipe to permit the transverse motion. Whilst such a change to the beampipe may be possible it is clear that this would be undesirable and would require a renegotiation of the specification and a study of the likely effect on the beampipe currents.

### Other Considerations with using a Rotary Mechanism

The use of a rotary mechanism would necessitate the use of some form of drive and possibly a gearbox to allow the mechanism to rotate at 60 rpm. Obtaining a commercial motor to accomplish this would be trivial if one worked on the assumption that the

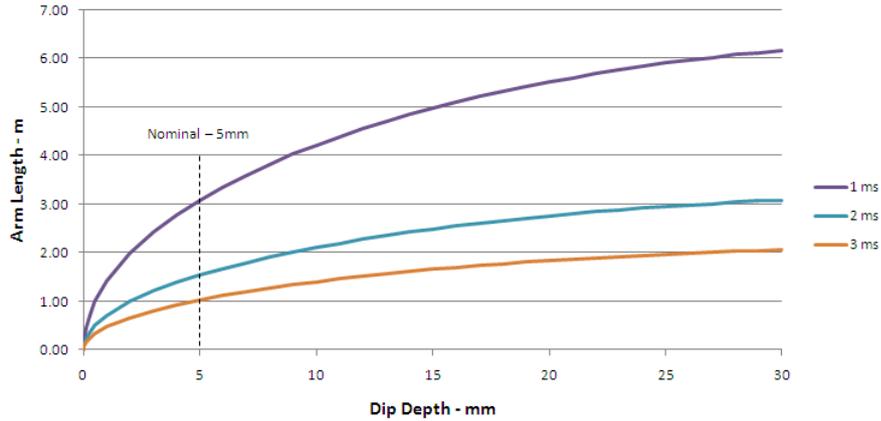


Figure 4.7: The minimum arm radius ( $r$ ) for a constant speed rotary target drive as a function of dip depth ( $d$ ) for three different insertion times. The radius of the ISIS beam for these plots ( $R$ ) is 40 mm. The motion of the arm is transverse to the beam. The solution here looks much better than for the previous case; for a 2 ms insertion time the required radius of the arm is of the order of  $\approx 1.5$  m at the estimated required insertion depth.

drive and any gearing could remain external to the vacuum. UHV rotary feedthroughs are commercially available and models can be obtained from several different manufacturers. For example ‘Rikagu’ [83] do a UHV rotary feedthrough where the seal is created with the use of a magnetic fluid. ‘UHV Design’ [84] also do rotary feedthroughs that are based on the use of a magnetic clutch. The amount of torque that can be handled by the magnetic clutch is model dependent. These products demonstrate that in principle providing a mechanism to transfer a rotary drive through to a UHV environment does not look insurmountable.

However, there is then the additional problem of how to manage changes in the insertion depth. Assuming a fixed length arm an additional mechanism would be required to raise or lower the axis of the machine in small controlled increments. This could in itself be another major challenge with a rotary design because the large rotating components would need to remain enclosed within the vacuum of ISIS.

Working on the assumption that the raising and lowering mechanism would be external to the vacuum then a set of extendable bellows could provide the means to allow vertical motion of the entire mechanism. If the top of bellows were at a distance  $z$  from the tip of the target then the minimum width  $w$  of bellows required to accommodate the rotational motion would be:

$$w = 2\sqrt{2zr - z^2} \quad (4.4)$$

For example if  $r = 1\text{ m}$  and  $z = 100\text{ mm}$  (20 mm of beam-shrinkage, 10 mm of clearance and 70 mm for compressed thickness of bellows plus end flanges) this would give an internal bellow width of 87 cm. Elliptical bellows are available from a number of companies and sizes of up to 70 cm are easily available off the shelf. Therefore it would not be unreasonable to assume that UHV bellows of this dimension could be manufactured.

Most of this discussion has been based upon the assumption that the driving mechanisms would remain external to the UHV system in which the rotary drive would be placed, and given the availability of both vacuum rotary feedthroughs and bellows this assumption seems entirely appropriate. This would almost certainly be preferable because if it proved necessary to place the driving mechanism within the vacuum itself, then this would significantly increase the difficulty of implementing such a design due to the limitations of what materials can be used within the ISIS vacuum.

One important drawback with using a rotary system with arm lengths as long as those suggested here is that the production point of the pions would move considerably during the 2 ms period that the arm was moving through the beam. In both cases discussed so far this movement would create a moving production point that, for the majority of the time, would fall outside of the acceptance angle of the the MICE quadrupoles. Typically the axis of the quadrupoles is physically aligned to within a few mm of the production point for optimum particle collection.

### Variable Acceleration Rotary Mechanism

So far only the case of using an arm rotating at a constant rotational velocity has been considered. It can be demonstrated that it is possible to use a shorter arm length if the arm could be suitably accelerated and decelerated during each rotational period. The ability to accelerate the arm would allow the angle  $\theta$  subtended during the insertion period to be much larger and therefore the radius of the arm could be much shorter. In principle, if the accelerating torque was high enough the rotating arm could be made an arbitrary length determined by the position of the rotor mechanism, the torque that can be achieved and the required depth of penetration into the ISIS beam.

Clearly the value of these variables would be design dependent and so only a few generalisations can be made here. If one starts with the formula given for a rotary mechanism earlier then by a simple rearrangement:

$$\theta = \tan^{-1} \left( \frac{\sqrt{(2Rd - d^2)}}{r} \right) \quad (4.5)$$

A plot of  $\theta$  as a function of  $r$  for  $r \gg R$  for two insertion depths (5 mm and 10 mm) is shown in figure 4.8. The angle  $\theta$  is related to the equivalent rotational frequency  $f$

by

$$f = \frac{\theta}{2\pi t} \quad (4.6)$$

where  $t$  is the required insertion time. If  $t$  is assumed to be 2 ms then for  $d=5$  mm the equivalent rotational frequencies in plot 4.8 range from 1.0 - 3.8 Hz and for  $d=10$  mm it ranges from 1.4 - 5.3 Hz. A detailed design would be required to further quantify the amount of torque required to achieve this rotational acceleration. Whilst this solution is clearly much more attractive than using a fixed rotational speed due to possibility of using a shorter arm length it is quite clear that an additional external mechanism would still be required to alter the insertion depth and there would remain the problem of a moving particle production point.

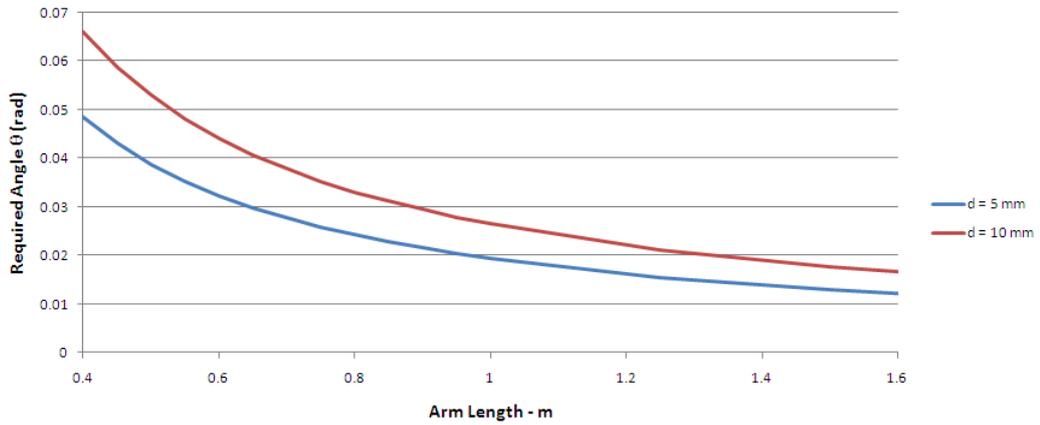


Figure 4.8: The required rotation angle  $\theta$  changes as a function of radius for a given insertion depth ( $d$ ) for a transverse rotating arm. Note that the formula is only valid if  $r \gg R$  therefore this is only used for  $r \geq 0.4$  m. The angle  $\theta$  should be compared to a  $\theta$  of 0.012 rad in the case of a constant rotational velocity of 1 Hz for an insertion time of 2 ms.

### 4.5.3 Linear Motor

The fact that two separate mechanisms would be required to operate a rotary machine suggested that a linear motor may be a favorable solution as it would permit the target to have a variable penetration depth into the ISIS beam in a single mechanism. A linear motor should fit into a much smaller footprint as the space required to accommodate a large rotating arm would not be required; the smaller footprint would also be consistent with the physical space available at ISIS for the mechanism. Additionally a linear device would suffer much less with the problem of a moving source of particle production as its

velocity would be low at the point where particle production occurs. On the other hand the required acceleration to make a linear device work was high, originally estimated at 100  $g$ , and achieving and controlling this acceleration would require development of a suitable feedback, power and control system; there was also the decision of how the moving components would be best supported in the vacuum of ISIS.

### Motion of a Linear Motor

The motion of a target being accelerated into the ISIS beam would go through four distinct stages.

- Acceleration towards the beam
- Deceleration towards the beam
- Acceleration away from the beam
- Deceleration away from the beam to bring it back to its start position

If one assumed that the target was accelerated by a constant force (and as will be shown in later chapters this is a reasonable and justifiable assumption) then the target would follow a set of trajectories that are described by a sequence of quadratic equations. In fact for a constant force the trajectory for the each stage could be described by translating the quadratic for the first stage with a suitable reflection in either the  $z$  or  $t$  axis. The point where the acceleration needs to be reversed to produce a deceleration is half way through the required travel. The motion of the target can therefore be described by the textbook mechanics equation:

$$z = ut + \frac{at^2}{2} \quad (4.7)$$

If this equation was used to describe the first stage alone then it is clear that  $z$  represents half the required actuation distance  $d$  and  $t$  represent  $t_A/4$  where  $t_A$  is the required actuation time. This allows an estimation of the required acceleration for the linear device to be made.

$$a = \frac{16d}{t_A^2} \quad (4.8)$$

ISIS operates on a 20 ms duty cycle, 10 ms of acceleration followed by a period of 10 ms where there is no beam. Using equation 4.8 it is possible to plot the required

acceleration against actuation depth to get a feel for the magnitude of acceleration that would need to be achieved by the device. The original estimation[65] of the required travel was approximately 25 mm and as can be seen from figure 4.9 at this travel the actuator would need to accelerate the target at 100  $g$  to meet the specification.

Pre-empting the fact that a design based on a linear motor was chosen it is worth adding at this point that the actual actuation depth was significantly underestimated during the early developmental stage because the beam dimensions were not fully understood and the requirement that the target should sit out of the beam-pipe when not being actuated was not specified until later. The maximum actuation depth required would be about  $\approx 50$  mm; this doubles the acceleration requirement in a 20 ms window. However, as will be shown in Chapter 9, if both beam shrinkage and the fact that the target can start accelerating a set distance away from the beam envelope is accounted for, then the value of  $t_A$  can actually be relaxed. Noting that the acceleration is dependent upon  $t_A^{-2}$  then relaxing  $t_A$  has a significant impact on the required acceleration even if the actuation depth is increased. Figure 9.6 in Chapter 9 shows a trajectory fitted to beam profile data taken at the target location in 2006. It can be seen from the plot that a  $t_A$  of 30 ms can be used and so this value is also included in figure 4.9. One can see that this indicates that the acceleration requirement is  $\approx 90 g$ .

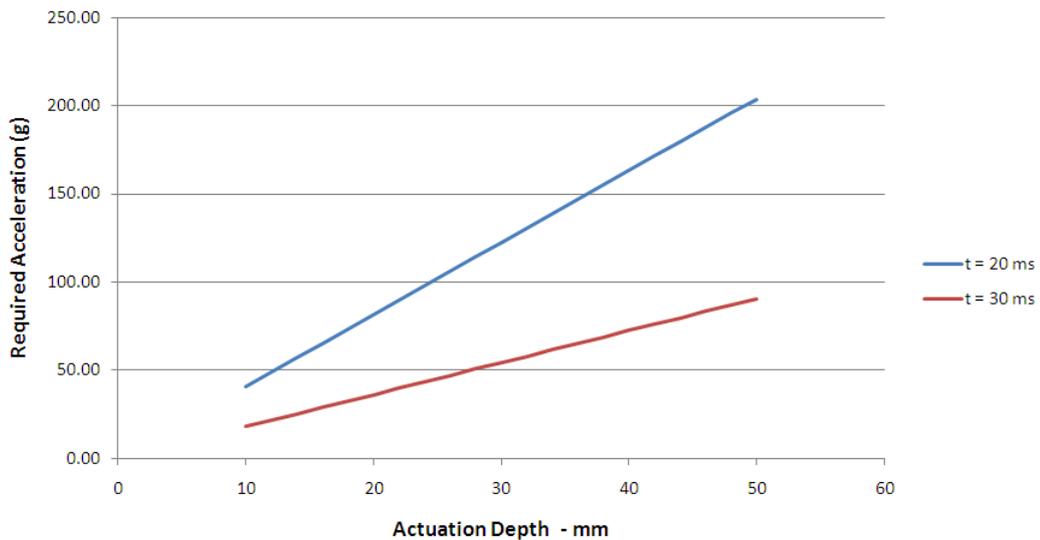


Figure 4.9: The estimated acceleration required to move the target through a given actuation depth  $d$  in time  $t_A$  using a series of quadratics defined by a constant magnitude driving force.

So far only the acceleration requirements of a linear motor design has been considered. When the linear motor is not actuating then the target would require levitating out of the ISIS beam. The necessity for this is because the mechanism must be installed above the ISIS beampipe due to the lack of space beneath it. The target would need levitating against gravity when it is not in its active actuation state. Nominally this

would be for  $\approx 970$  ms out of its 1 s baseline duty cycle.

When this problem was originally recognised it was thought that longitudinal levitation of the shaft would need to be done actively via a feedback loop on the position monitoring system. However it was later realised by the author that a much simpler system could be used to levitate the mechanism that didn't rely on active feedback but just used the symmetry of the driving field of the stator to advantage. It should be noted that a passive system is distinctly desirable because it gives a stable starting position at the start of an actuation. However as the purpose of this chapter is to justify the use of the linear motor design a description of this system is left until Chapter 7. The point of referring to the development of the levitation system here is that the inferred stability in the  $z$ -direction due to symmetries within the drive structure of a linear motor has important implications for the form of the expected radial forces acting upon any shaft running through the drive.

## 4.6 Choice of Design

Of the two possible target insertion mechanisms considered in detail, the linear motor is best able to meet the design specification. Rotating designs of a reasonable size are incapable of meeting the 1-2 ms insertion time requirement. Any rotary mechanism would have meant making a large opening in the ISIS beampipe which was not favoured at the time. Furthermore, the linear motor design provides an efficient mechanism by which the insertion depth can be easily altered on an insertion by insertion basis. For these reasons, the MICE target team elected to develop a linear motor.

## 4.7 Earnshaw's Theorem and its Implications

Earnshaw's theorem simply states that there is no static configuration of permanent magnets that can achieve stable levitation. Note that a permanent magnet can be a coil with a stable current passing through it.

The force acting on a magnetic dipole is given by

$$\mathbf{F} = \nabla(\mathbf{m} \cdot \mathbf{B}) \quad (4.9)$$

The divergence of the force is related to the energy through the relationship

$$\nabla \cdot \mathbf{F} = \nabla \cdot (-\nabla U) = -\nabla^2 U \quad (4.10)$$

Stable levitation can only be achieved if each term in the energy Laplacian has a value greater than zero.

$$\nabla^2 U = \frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2} \quad (4.11)$$

The proof is involved but it can be shown that for all static configurations of magnets that  $\nabla^2 U = 0$ . Principally the result is obtained because the divergence of a magnetic field,  $\nabla \cdot \mathbf{B}$ , is always equal to zero.

$$\nabla^2 U = \frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2} = 0 \quad (4.12)$$

So this leaves three possibilities.

- All three terms must equal zero.
- Two terms are positive and a third term is negative.
- One term is positive and the other two terms are negative.

This is relevant to the design of a linear motor for this particular application because as stability in the  $z$  direction is desirable for a fixed current to allow levitation to occur then this means that in either the  $x$  and/or the  $y$  direction the magnetic field must be magnetically unstable. In fact, pre-assuming rotational symmetry of the coil structure then stability in the  $z$  direction automatically imposes radial instability in the magnetic structure.

In fact the inferred radial instability due to the thrusting force is not really a choice. Even in a general sense if a drive is designed such that it uses a set of coils in a configuration that maximises the force on a set of magnets by utilising multiple poles (particularly desirable for a high force device as it helps to close the magnetic loop thereby reducing the reluctance of the magnetic circuit) then the symmetry of the coil/magnet construction that closes the loops will inevitably provide a magnetic well structure in the direction of the intended motion. Therefore instability in the other dimensions is forced upon the design and will need to be countered by either mechanical or electromagnetic means.

## 4.8 Bearings

The previous section justifies the need for radial support of the target shaft in a linear mechanism. There are four types of bearing that can be considered, namely linear

bearings, magnetic bearings, sprung bearings and sliding bearings.

### 4.8.1 Linear Bearings

Linear bearings refer to bearings that contain several sets of small ball bearings to carry the load. The bearings can either be in a fixed location, restrained by a retainer or placed within a racetrack. In the latter case the movement of the ball bearings around these racetracks provide the bearing mechanism. Considering the moment of inertia of a solid sphere one can easily calculate the required force to accelerate a ball bearing about a static axis at  $\approx 90 g$ .

$$F = \frac{8\pi r^3 \rho a}{15} \quad (4.13)$$

In the case where the ball bearing also has translational motion

$$F = \frac{18\pi r^3 \rho a}{15} \quad (4.14)$$

Here  $r$  is the bearing radius,  $\rho$  is the material density and  $a$  is the acceleration of the bearing on its outer radius. For a steel ball bearing with a 1 mm radius this gives a required force on a static axis of  $\approx 0.01$  N and  $\approx 0.02$  N for translational motion. Although the forces required to accelerate a suitable target have yet to be discussed it is possible to estimate the likely order of magnitude. Based upon the assumption that the moving mechanism would have a mass of at least several tens of grammes, a thrust force of several tens of Newtons will be required to obtain the necessary acceleration. Therefore the force required to accelerate many ball-bearings would not be a major constraint.

For the acceleration to occur without slippage would require a normal contact force that is of the approximately the same order of magnitude as that of the accelerating force if one estimates it using the coulomb force  $F_{T_{max}} = \mu F_N$  and assumes that  $\mu \approx 1$ . A more formal and accurate calculation would require calculating the Hertzian contact stress, this defines the contact area and is a function of the preloading, i.e. the normal contact force. An increase in the preload increases the coulomb friction reducing the risk of slippage[85].

The suggests that the required contact forces are small but there are problems with providing these forces. It will be demonstrated later on in the thesis that for the design of linear motor that was used the radial forces are theoretically zero when the shaft is on axis. However as was explained in the previous section, this is an unstable minimum with a negative second order differential in the energy function. In reality there will

be contact forces between any bearing and shaft but they will not be well defined due to the tolerances of construction and because of the dynamic forces when the shaft is accelerated towards the beam, this would almost inevitably lead to slippage between the shaft and bearings.

Slippage would be bad for any ball bearing contact as it would quickly wear flats on the small contact area. This would render the bearing useless as any flats would prevent the bearings from rotating properly thereby further exacerbating the problem. For this reason linear bearings are not considered appropriate for this application and the idea has never been pursued further.

### 4.8.2 Magnetic Bearings

Stable levitation is not possible with a passive array of magnets/electromagnets. However if the position of the levitated object is tracked and active feedback is used to alter coil current then this invalidates the assumption of a static field from which Earnshaw's theorem is derived and 'stable' levitation can be achieved. Stable magnetic levitation, aka magnetic bearings, are a speciality in their own right and it would not be possible to do the subject justice here. Their main advantage is frictionless motion, therefore they do not wear out (in a mechanical sense), do not require lubrication and, in the case of rotary bearings, high rotational speeds are possible.

A search of the literature reveals that a large part of their application is for rotational bearings[86][87]. Magnetic bearings have been used to support a shaft undergoing longitudinal motion; a paper produced by Feeley et al [88] for NASA modelled the use of magnetic bearings in a long life Sterling cooler for a space application.

It is clear from these papers that the use of magnetic bearings in any given application requires a good understanding of the dynamics of the system, else the feedback mechanism may not be able to properly respond to dynamic changes in the driving forces. The dynamic nature of the driving forces for a yet to be designed linear motor were completely unknown when a suitable bearing mechanism was being considered. Additionally, one could soundly argue that reliable data of the dynamic forces can only be obtained from measurements of a prototype build. Given that only limited manpower was available for the project and that much of this man-power would be required to construct and design the necessary control system for a high acceleration stator it was realised that further complicating a linear design by adding magnetic bearings was seen as both a significant risk and possibly unnecessary.

### 4.8.3 Sprung Bearings

Sprung bearings were at first considered to be an attractive proposition for a linear motor because they allow motion without any parts having to slide or roll over each other. Providing that the sprung mechanism has been properly designed they can also

be operated for significant periods without fatigue, therefore this type of bearing should never wear out or break during its operational lifetime. A spring is clearly suited to working in UHV.

Diaphragm springs were considered to be the most appropriate method of supporting a linearly moving shaft because they should constrain radial motion whilst allowing longitudinal motion. Once again pre-empting the conclusion that a linear motor was preferable over a rotary mechanism a prototype actuator was built using diaphragm springs; more details of this will be given later on in the chapter.

#### **4.8.4 Sliding Bearings**

Sliding bearings are the simplest type of bearing and have the advantage of being extremely stiff. Usually sliding bearings are lubricated to help improve their lifetime, but lubrication is not always necessary and this makes them suitable for use in UHV. The main problem with sliding bearings is wear; two unlubricated surfaces rubbing together will wear, usually the less hard surface being the one that wears. This can create problems with dust production, seizing and short life if the wear is significant.

Once it had become clear that sprung bearings were unsuitable for the target mechanism, sliding bearings were the logical choice. Finding a suitable bearing material that works reliably in UHV for an extended period of time has proven to be a difficult task and has required the assistance of tribologists.

### **4.9 Motor Development**

As no commercial product could be found the expertise of an electrical engineer, N. Schofield[89], who specialised in motor design, was sought from the outset. The motor would have to be of a permanent magnet brushless design; this is because the high acceleration and the travel of the motor ruled out the possibility of placing coils on the moving components. The integration of the permanent magnets onto the moving assembly would remove the need for electrical contacts between the stator and the moving parts thereby simplifying the interface between the motor and the ISIS vacuum. The magnets on the moving components would interact with a set of stationary coils in the stator body. These coils can then be kept outside of the ISIS vacuum and directly wired to the driving electronics.

A paper produced by N.Schofield in 2005[90] gives details of the magnetic design of the linear actuator. The position monitoring and control of the stator became the responsibility of the Sheffield target group, primarily that of the author. As the electromagnetic design of the stator is central to the operation of the target mechanism a summary of the main results of Dr Schofield's paper[90] is reproduced here.

The main requirement of the actuator is to produce a large accelerating force for a relatively small duty cycle, therefore large instantaneous coil currents are achievable. There is a strong correlation between electromagnetic design and mechanical implementation; the design needs to maximise the force whilst minimising the mass of the moving components. The target will be of low mass therefore the magnetic components will constitute a significant fraction of the mass of the design.

For an actuator with permanent magnets on the moving components the force on the magnets is given by:

$$F = 2\pi r_a B_{ave} Q \quad (4.15)$$

where  $r_a$  is the magnetic assembly (shuttle) radius,  $B_{ave}$  is the magnetic loading due to the shuttle's permanent magnets and  $Q$  is the loading of the stator's coils in A/m. Maximum magnetic loading is achieved by using NdFeB permanent magnets. An air-cored design is used because manufacturing tolerances and losses from using an iron core would otherwise compromise the maximum achievable force.

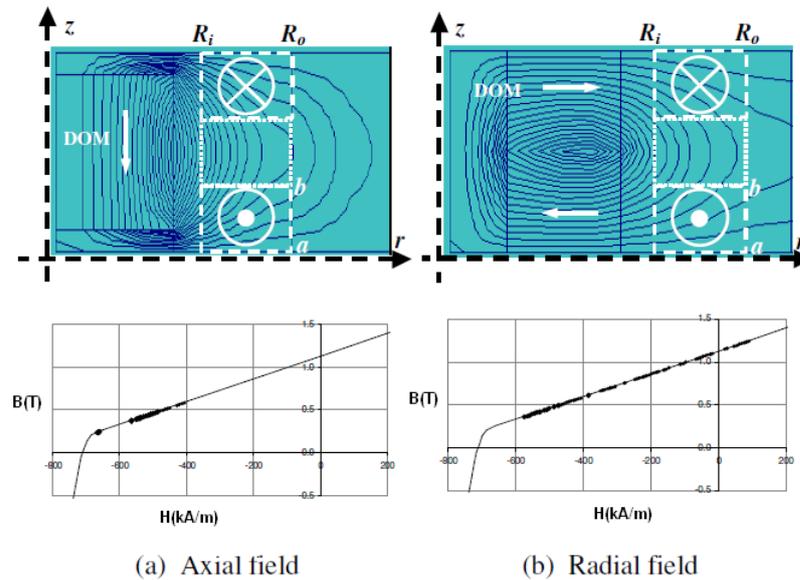


Figure 4.10: Example of FEA results for the two shuttle design topologies considered showing coil current excitation and the direction of magnetisation. The radial field topology was shown to be superior. Diagram taken from Schofield's paper[90].

Two stator topologies were considered and modelled using Finite Element Analysis (FEA). The first design used multi-pole axially magnetised disc magnets separated by soft iron pole pieces whereas the second design considered radially magnetised magnets

attached to a soft iron core. Cross sections of these topologies are shown in figure 4.10. FE analysis indicated that the radial topology was preferred because it showed a 30% larger force than for the axial field topology. The radial field topology was also more resistant to demagnetisation, which was a particularly important point as the exact mass of the moving components was unknown at the time of the electromagnetic design and so the stator coil currents could only be estimated. Once the topology had been decided FEA work was then used to optimise the component geometries. Mechanical/dimensional details of the design are shown in the next chapter.

The design was optimised to produce a maximum force on a low mass target shaft but the actual acceleration that could be achieved with the design could not be determined until the combined masses of the mechanical components were known and the feedback, power and control systems were built. Chapters six, seven and eight describe the design building and testing of the feedback, power and control components necessary to operate the stator, and this forms the main part of the work in this thesis.

## 4.10 Test Stand for Development

In order to fully ascertain the viability of using a linear motor for the target mechanism some early developmental work was done to understand some of the important issues involved. This section highlights some of the important work done here and how it shaped the mechanical design that is described in Chapter 5. As was indicated earlier in this chapter a sprung method of supporting the target shaft was initially considered to be advantageous and so a test stand was set up to ascertain whether diaphragm springs were suitable for this purpose.

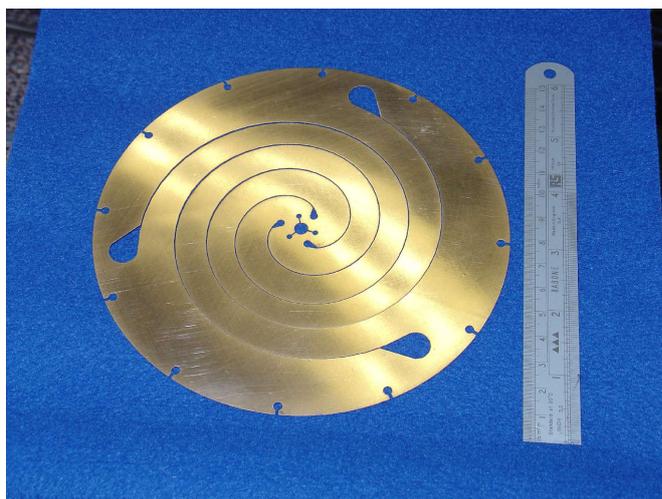


Figure 4.11: The diaphragm springs were designed to provide a mechanism of supporting the shaft for a travel of  $\pm 20$  mm.

A set of diaphragm springs were designed at Sheffield University using finite element software based upon a design provided by Dr Bradshaw[91] a member of the MICE collaboration. The finite element package was used to ensure that the springs could withstand the repeated stress of being cycled many millions of times through  $\pm 20$  mm. The final springs were approximately 150 mm in diameter and were manufactured from 250  $\mu\text{m}$  and 300  $\mu\text{m}$  thick beryllium-copper using a wire erosion technique - See figure 4.11. In order to test these springs a basic linear motor was built that would permit 36 mm of travel. This motor was not in any way optimised electromagnetically and it was not intended to demonstrate the high acceleration required for the target mechanism but was built to allow various peripheral ideas to be demonstrated, including the use of diaphragm springs. The motor comprised of just six coils at a pitch of 15 mm. These coils were individually wired and embedded in an epoxy resin. The magnet assembly was bulky and comprised of two permanent magnets sandwiched between three pole pieces. The magnets and poles had a mass of  $\approx 150$  g. The final assembly is shown in figure 4.12. The driving electronics was relatively simple, its purpose was to just move the shaft in a continuous oscillatory motion through a distance of 30 - 40 mm at a maximum frequency of a few Hz. (The maximum frequency achieved with this system was 6.4 Hz, this is equivalent to an acceleration of  $\approx 2$   $g$ . However at this frequency the longitudinal motion was unstable due to the lack of any feedback!) No positional feedback mechanism was used when powering the coils, which were fired in a timed sequence. This particular test assembly was operated in air as opposed to vacuum because the maximum momentary velocity of the device at  $< 1$   $\text{ms}^{-1}$  was sufficiently low enough that any air resistance would have been negligible.

This experimental rig also provided an opportunity to test a magnetic quadrature encoder. Full details of how a quadrature encoder works is presented in chapter 6, however it is sufficient for now to state that it permitted the travel of the shaft to be measured accurately. Position monitoring of the shaft would be a necessary requirement for both control and data-acquisition. The quadrature encoder used for this purpose was a commercial unit and comprised of two components, a read head that contained the encoder's sensors and electronics and a magnetic strip that passed underneath the read head. As the strip passed underneath the read head the distance travelled between the two components was measured to an accuracy of 20  $\mu\text{m}$ . The magnetic strip was attached to the shaft and the read head was mounted on the top of the stator as indicated in figure 4.12.

As the diaphragm springs extended it was found that they lost their lateral stiffness, therefore they were not effective at keeping the shaft centrally aligned once their displacement had exceeded a few mm. Additionally, it is shown below that the inertial mass of the diaphragm springs is problematic and would have prevented the necessary acceleration from being achieved even if the lateral stiffness of the springs could be improved.

The diaphragm spring will need accelerating along with the mass of the shaft and the shuttle. The centre of the spring accelerates at the same rate as the shuttle but this acceleration tapers off linearly to zero at the edge of the active area of the diaphragm

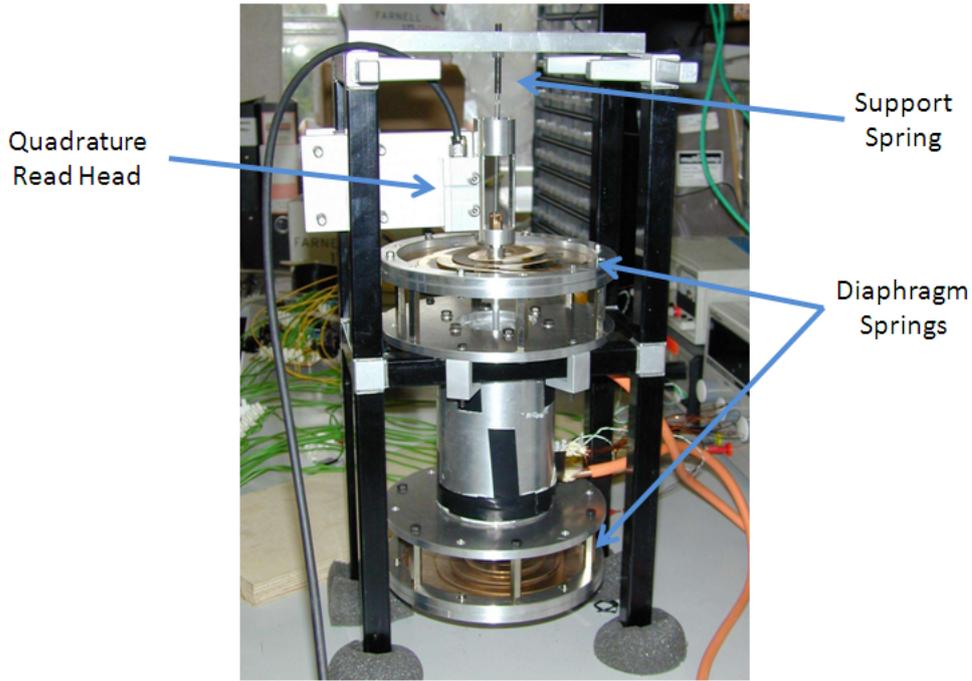


Figure 4.12: The completed test assembly. This test rig was produced for testing the diaphragm springs and the quadrature read head.

spring. It can be shown that the required force to accelerate the mass of the spring is

$$F = 2\pi\rho t a_m \int_0^{r_a} r \left(1 - \frac{r}{r_a}\right) dr$$

$$F = 2\pi\rho t a_m \frac{r_a^2}{6} \quad (4.16)$$

where  $\rho$  is the density of the metal  $\approx 9000 \text{ kg m}^{-3}$ ,  $t$  is the thickness of the spring which in the prototype was  $250 \text{ }\mu\text{m}$ ,  $r_a$  is the active radius of the spring which is  $0.057 \text{ m}$  and  $a_m$  is the maximum acceleration which was estimated at  $90 \text{ g}$ . This gives an inertial force of  $6.7 \text{ N}$  per spring, therefore the effective inertial mass of the diaphragm spring is  $7.6 \text{ g}$ . There were two springs in the prototype setup, which is the minimum number that can be used. This force does not include the masses of any bosses that connect the spring to the shaft and does not account for the spring constant. Therefore it would be reasonable to estimate that a minimum force of  $\approx 15 \text{ N}$  (equivalent inertial mass of  $\approx 15 \text{ g}$  + weight of bosses etc) would be required to suitably accelerate a pair of springs alone.

At about the same time, the magnetic studies for the motor design were coming to completion. The final motor was designed for a very low mass load and given that it was later demonstrated that the motor design could only produce a thrust force of  $\approx 45$  N at 1 Hz without overheating, it was presumed that the additional mass of the springs would have prevented the necessary acceleration from being achieved.

The lack of stiffness of the springs constitutes the main argument for not pursuing this type of support mechanism for the target shaft; however it can be seen from the above calculations that the inertial mass of the springs would have been a problematic issue in itself. Therefore the decision was made not to use diaphragm springs in the design<sup>1</sup>. Mechanical details of the sliding bearings that are used in the design to support the shaft can be found in the next chapter and chapter 10.

The quadrature counter was successful at measuring the position of the target shaft in the experimental rig. However it was clear that these commercial components could not be used in any final design. Firstly the read-head contained much of the decoding electronics and, from discussions with the manufacturer, it was clear that the electronics could not be separated out from the read-head. An extensive search of other manufacturers showed there to be no alternative commercial products that offered a quadrature system where all active electronics could be separated from the target mechanism by a distance greater than a few metres. Secondly in the case of this particular product the magnetic attraction of the read-head to the magnetic strip created another source of frictional force that would have reduced the net achievable acceleration. For this reason a novel optical quadrature measurement system was designed in house. The full development of this piece of equipment is described in Chapter 6.

## 4.11 Design Requirements of Sliding Bearings

As previously noted the actuator has to run in a Ultra High Vacuum (UHV) in a radiation environment. These two constraints and the experimental work just described determined that the target shaft had to be supported on a pair of sliding bearings.

The bearings must keep the target shaft coaxial down the centre of the bore of the stator. The radial magnetic field across the stator bore is symmetric but unstable, Earnshaw's theorem states that this is to be expected, and in the case of the target stator this has been confirmed both by measuring the field inside the stator and by Finite Element Analysis. From consideration of figure 4.13 it can be seen that if the magnet assembly deviates laterally from the central axis then this will create a resultant lateral force on the magnets due an asymmetry in the integrated B field on either side of the centre of the magnet assembly. Currently, at the time of writing the thesis,

---

<sup>1</sup>A thrust force of 47 N was eventually achieved on a target mass of  $\approx 52$  g giving a maximum of 90 *g* of acceleration. With a conservative estimated additional inertial mass of at least 15 g from the diaphragm springs this would then reduce the maximum acceleration to 71.5 *g*. This acceleration would not have been sufficient to meet the design specification.

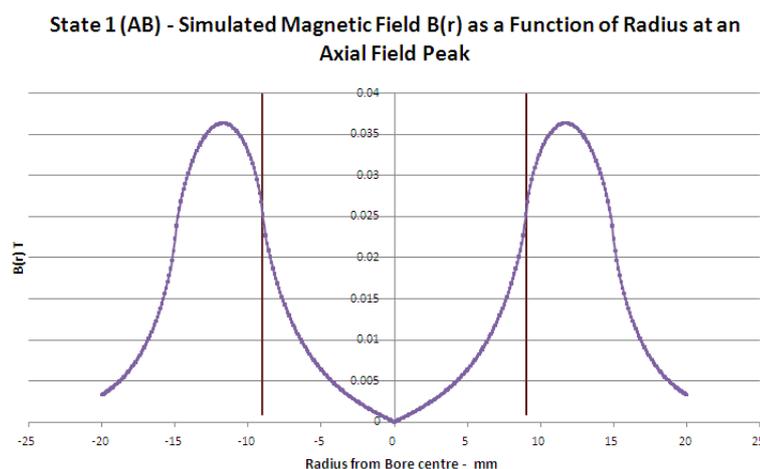


Figure 4.13: A simulation of how the magnetic field changes radially for a current of 7.5 Amps at a coaxial point that corresponded to the peak of the magnetic field. (The magnetic field changes sinusoidally coaxially) The vertical lines represent the position of the stator bore tube. It is clear from this diagram that if the magnets are not positioned coaxially to the bore tube, and therefore the coil stack, that there will be a resultant lateral force on the magnetic assembly.

the magnitude of this force cannot be calculated due to inconsistencies in the FEA model, however it is clear from the magnetic model that such a force would exist. The increased coulomb friction from any lateral forces would then impede shaft movement and increase the rate of wear. Carefully aligning the magnets coaxially within the stator bore will help to alleviate any problems with respect to this issue. Minimal clearance in the bearing apertures assists with keeping the shaft coaxial and will help to prevent excessive ‘rattle’ during the shaft travel.

As will be illustrated in the next chapter it was natural to use the bearings as a mechanism to prevent the shaft from rotating, necessary for the optical position system that will be described in Chapter six. The bearings must also exert minimal static and dynamic frictional force upon the shaft. This is aided by ensuring that either the bearing or its coating has a low coefficient of friction. The consequences of not minimising the frictional forces are significant; it reduces the net acceleration and increases wear. Wear also increases the clearances in the bearings which in turn increases the lateral movement available to the shaft during its travel exacerbating the problem. Once bearing wear is established failure of any sliding bearing would be quickly expected.

Several different bearing materials and configurations were tried in a prototype actuator based upon the magnetic design described in this chapter over a period of several months before a reliable solution was found. It was found that the best solution was to use steel bearings with 5 mm deep flat bearing surfaces coated with Diamond Like Carbon. (Diamond Like Carbon is discussed further in the next chapter.) These bearings have demonstrated an operational life of up to 2 million actuations before macroscopic wear

became evident. At the time of installing the first target mechanism DLC was proven to be by far the most promising material to date and was approved for use on the target mechanism by ISIS. Details of the mechanical construction of the bearings can be found in the next Chapter whereas further details of the history of the bearing development can be found in Chapter 10.

## 4.12 Conclusion

This chapter has discussed the development of the baseline design for the MICE target mechanism. The choice of a linear motor approach has been justified based on design requirements, the operational environment and constraints imposed by the operators of the ISIS accelerator. Crucial to the selection of the linear motor technology is a suitable design for the bearings. A test stand for the target mechanism development was described, and this test stand was used to conduct tests of diaphragm spring bearings and a quadrature position sensor. Based on these tests, diaphragm spring bearings were rejected. As the development of magnetic bearings was considered to be both complex and risky, sliding bearings were shown to be the only remaining option likely to satisfy the design requirements. Tests of the quadrature position sensor were successful, and prompted the development of an optical quadrature sensor that is now implemented on the current MICE target.

The next four chapters consider the mechanical build, the optical position measuring system, the power electronics and stator operation and then the control electronics and data acquisition (DAQ). The combination of these systems provides a description of the first target system that was installed on ISIS.



## Chapter 5

# The Mechanical Design of the MICE Target Mechanism

The last chapter gave a background to the required performance of the target mechanism and gave consideration to the choice of using a linear motor to provide the necessary solution. Some early development work that was instrumental in shaping the design of the target mechanism was also described.

The purpose of this chapter is to describe the design and construction of the mechanical components of the stator body in more detail. Further details of the operation of the stator in terms of how the permanent magnets on the shaft actually interact with the stator coils to produce an accelerating force, is covered separately in Chapter 7. Where appropriate, technical drawings and photographs have been included. High resolution drawings of the stator components can be found on the Sheffield University target website[92].

### 5.1 An Overview of the Target Mechanism Design

This first section has been written to provide the reader with a brief overview of the assembled target mechanism and how it is operated. Further details of the mechanical components of the stator will follow in subsequent sections, however the design and operation of the optical positional system is discussed in the next chapter.

#### 5.1.1 Construction

The linear motor that drives the target into and out of the ISIS beam consists of a moving magnet assembly on a long shaft (target shaft) carrying the target inside a

series of coils (the stator). Figure 5.2 shows a cross sectional diagram of the whole of the target mechanism whilst figure 5.1 shows a close up drawing of the stator body.

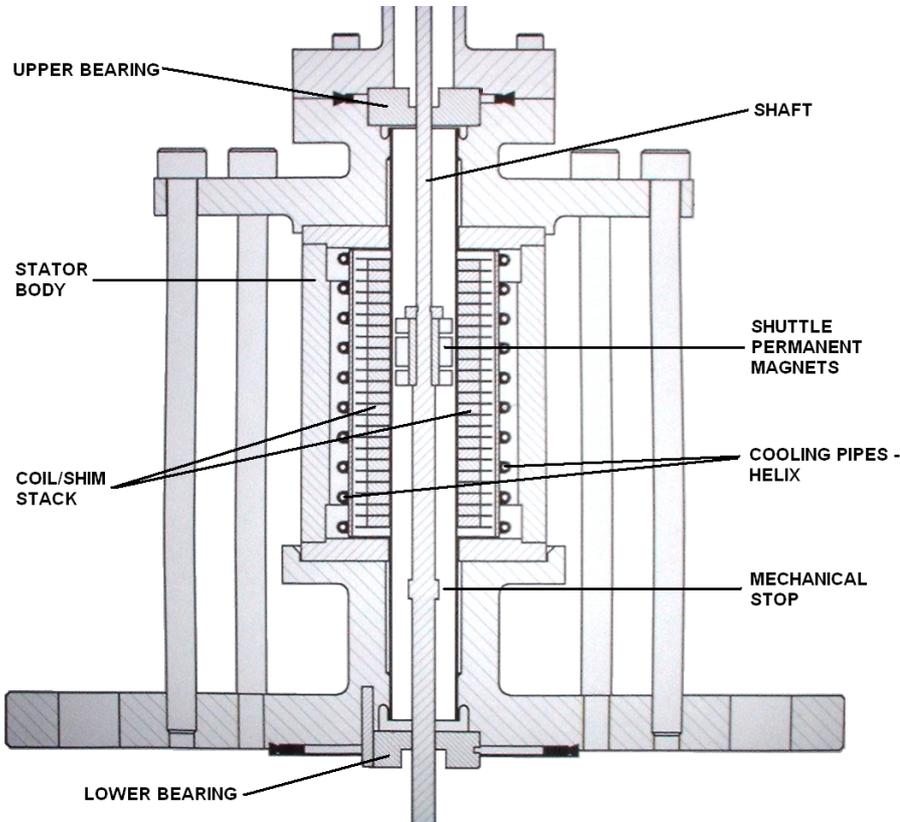


Figure 5.1: A close up drawing of the stator itself. Here the details of the coil stack within the stator can be seen.

The stator, which is cylindrical in shape, contains 24 flat coils mounted around a thin walled steel tube. Individual coils, with an inner diameter of 18.3 mm, consist of 36 turns of copper wire and have an axial thickness of 2.85 mm. After winding, each coil is impregnated with insulating varnish to form a stable compact unit. During assembly four 25  $\mu\text{m}$  copper shims are sandwiched between two 25  $\mu\text{m}$  kapton shims and these are placed between each pair of coils to facilitate heat conduction out of the coil stack whilst maintaining electrical isolation. The addition of the copper shims gives a coil pitch of 3 mm. Connecting leads from the coils are led radially outwards. Three thermocouples are inserted between three pairs of coils to enable the temperature of the coil stack to be monitored. A coiled copper tube soldered onto a solid copper jacket is placed around the coils and is in contact with the copper shims. This carries the cooling water, and the temperature of the water is monitored at either end with two more thermocouples. The entire assembly is inserted into an aluminium outer cylinder, the stator body, with the insulated copper wires and the cooling pipes emerging through a slit in the side. The individual coils are wired up at terminal blocks placed external

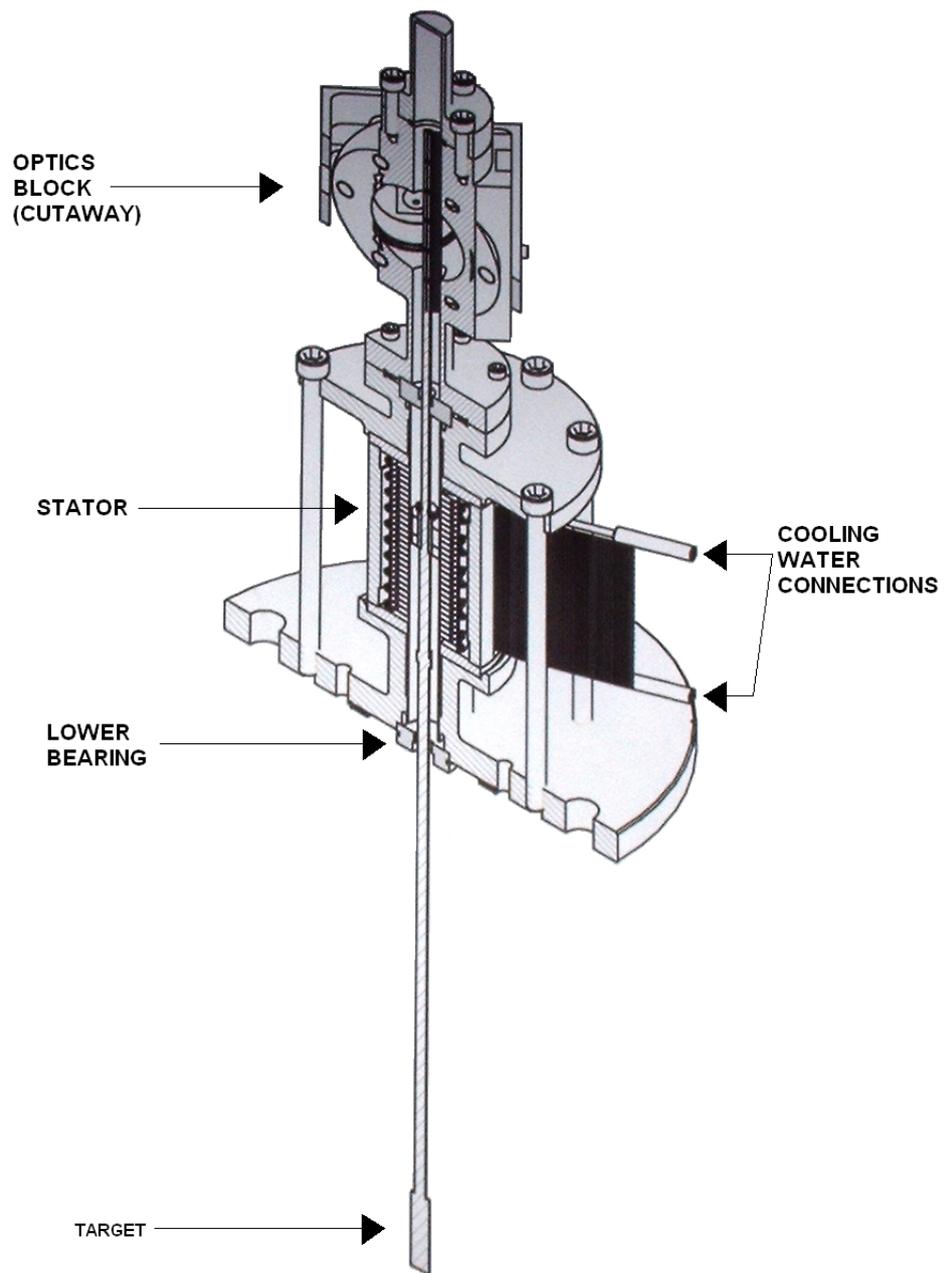


Figure 5.2: A cross sectional view of the target mechanism including the full length of the target shaft and the target. The black plane protruding from the stator block on the right hand side of the diagram represents the wires protruding from the stator coils.

to the stator body.

The target shaft assembly consists of a magnet assembly mounted on the long target shaft, the target blade at the bottom and a readout vane at the top. To prevent the magnets from falling out of the range of the coils in the absence of power, a larger diameter section of the target shaft acts as a stop that can rest on the lower bearing.

The target blade, shaft and stop are machined together from a single piece of titanium. The target, at the lower end of the shaft, consists of a blade of titanium 1 mm thick, 10 mm wide and 35 mm high. The shaft, for most of its length of 530 mm, has a cross-shaped cross-section, with material thickness of 1 mm and a total width of 6 mm. The cross-shaped form passes through a similarly shaped aperture in the lower bearing, thereby maintaining the orientation of target and readout vane. The upper third of the shaft is circular in cross-section, of diameter 4 mm. The magnet assembly slides onto the shaft from above, resting on the top of the cross-shaped section. It is held in place with a stop that is clamped to the shaft. The final 94 mm of the shaft has a slot cut vertically into it to carry the readout vane. This is used with an optical encoder to monitor the position of the target shaft during operation. The sections of the shaft that are in contact with the bearings are coated with Diamond Like Carbon (DLC) to minimise friction and to give a hard wearing surface.

The magnet assembly consists of three sets of radial magnets produced from sintered Neodymium-Iron-Boron (NdFeB). Each of the radial magnets is produced in octants. The upper and lower radial magnets are magnetised so that the outer face of each sector is a South pole and the middle radial magnets are magnetised with the outer face as a North pole. The unit is then attached to the target shaft, as previously described. The shaft passes through two steel bearings, one above and one below the stator assembly. These bearings maintain the magnet unit on the axis of the stator while allowing longitudinal (vertical) movement with minimal friction. The bearings, like the shaft, are coated with diamond like carbon to give a hard wear resistant surface. Position sensing is performed using a quadrature system viewing an optical vane mounted in the slot at the top of the shaft. The vane is a wire-eroded double-sided ‘comb’ of 0.2 mm thick steel, having 157 teeth 0.3 mm wide (with 0.3 mm gaps) and 3 mm long on one side of a 6 mm wide spine, and a single similar tooth two-thirds of the way down the vane on the other side. More details of how this optical position monitoring system was developed and works will be presented in chapter 6.

### 5.1.2 Function

There are a number of modes of operation of the target drive. These include a reversible movement from a powered off ‘park’ to a raised ‘hold’ position that levitates the target outside the beam envelope. There is then an ‘actuate enabled’ mode, when the electronics is waiting for a trigger and ‘actuating’, the triggered rapid insertion of the target into the beam. All require the appropriately phased application of currents through the stator coils, and are under microprocessor control. The three-phase, bi-directional

supply to the coils is switched through six Integrated Gate Bipolar Transistors (IGBTs) powered by a capacitor bank. The movement between power-off and the shuttle's holding position is done passively without feedback from the optical system. These movements and levitation of the shuttle at its holding point can be done at a relatively low coil current of approximately 3 A.

Target insertion is synchronised to the ISIS machine start signal. The ISIS machine start signal is correlated to the injection of protons into the synchrotron by a pre-defined delay. After a second programmed delay, the current through the coils is increased to 60 A to drive the shaft through its trajectory at high acceleration. Feedback from the position sensing ensures that the correct coils are powered in sequence maintaining the maximum force on the shaft magnets. When the target is halfway through its descent, the controller reverses the currents so that the shaft experiences a decelerating (upward) force. This decelerates the target until the target reaches its maximum insertion depth and then re-accelerates the shaft and target back up the actuator. At a second preset point the currents are reversed again, decelerating the shaft so that it comes to a halt at its intended holding position. At this point the microprocessor changes the mode to keep the target shaft levitated at its hold point until another actuation signal is received. Further details of the control system is given in Chapter 8. The remainder of this chapter will consider the mechanical components of the target system in more detail.

## 5.2 Stator Body

### 5.2.1 Stator Bore

The stator bore runs through the centre of the stator and provides isolation between the stator body and the ISIS vacuum as well as providing mechanical support between the upper and lower flanges. Early prototypes of the stator contained a machinable ceramic tube of 1 mm thickness and 18 mm outer diameter and this was sealed to the end flanges to provide a UHV seal using a compression fitting with Indium metal. Unfortunately, Indium seals creep under the pressure of the join and this eventually leads to the failure of the seal over time. The use of this type of seal was unacceptable to ISIS and so an alternative joining mechanism needed to be found.

A thin walled non-magnetic stainless steel tube was swapped into a working stator, replacing the ceramic tube. This tube had a nominal wall thickness of just 0.5 mm and was reduced to 0.3 mm where it came into contact with the coil stack. Subsequent analysis of trajectory data showed that the use of a thin walled steel tube gave no discernable difference in stator performance when compared to the ceramic tube. Given the ease of creating a vacuum tight steel-to-steel join through a simple weld this thin walled steel tube was adopted as the baseline material for the bore. Subsequent simulations confirmed the practical result showing that the use of a thin walled steel tube would reduce the flux such that the magnetic force would degrade by approximately

1%, consistent with what was observed.

This steel tube would come into direct contact with the inner radius of the coils if it were not suitably insulated. The slight recess in the tube that is created by machining the wall thickness down to 0.3 mm allows three layers of self adhesive kapton tape to be applied to the outside of the tube. This provides protective insulation by preventing the inner surface of the coil stack from coming into contact with the steel tube.

### 5.2.2 Coils

The stator contains twenty four identical coils that are stacked vertically and numbered one to twenty four starting from the top of the stator. The stator coils are responsible for interacting with the permanent magnets on the shaft both to levitate the target shaft when the target is being held out of the ISIS beam and to produce the accelerating force when the target needs to be inserted into the ISIS beam. The interaction between the stator coils and the magnet assembly is quite complex and a discussion of this has been left until chapter 7. This section primarily deals with the mechanical construction of the coils.

Each coil is composed of thirty six turns of 0.56 mm polyester-imide enameled copper wire, over-coated with a polyamide-imide. This is a high temperature motor winding wire that is rated to 200 °C operation[93]. These coils are wound on an 18.1 mm diameter former and each coil is manufactured to a depth of 2.85 mm  $\pm$ 0.1 mm. Coils outside of this tolerance are rejected due to the limited space between coils which is fixed to a pitch of 3 mm. Clearly these dimensions provides  $\approx$ 0.15 mm clearance between adjacent coils, this space is used to insert thin kapton and copper shims providing both the necessary electrical isolation between coils and the ability to heat sink them. After winding, the outer diameter of the coils is 30 mm. The coils are then double dipped into a varnish that seals the windings and provides additional electrical insulation[94]. A photograph of a finished coil and some of the intercoil copper shims are shown in figure 5.3.

Each coil is individually tested on a boss to 1 kV before it is built into a stator. These tests are part of the stator build quality assurance procedure and are designed to ensure that the stator will not prematurely fail shortly after construction due an invisible mechanical defect in the outer insulation of the coil. Turn-to-turn insulation tests are not performed as the voltage seen across an individual coil is very low during normal operation ( $<20$  V) and so internal coil breakdown is considered unlikely.

### 5.2.3 Shims

As already noted there are two types of shim filling the 0.15 mm gap between the coils. These shims serve two purposes. Firstly the kapton shims electrically isolate the coils from the copper shims as these shims form part of the stator earthing. Secondly, the

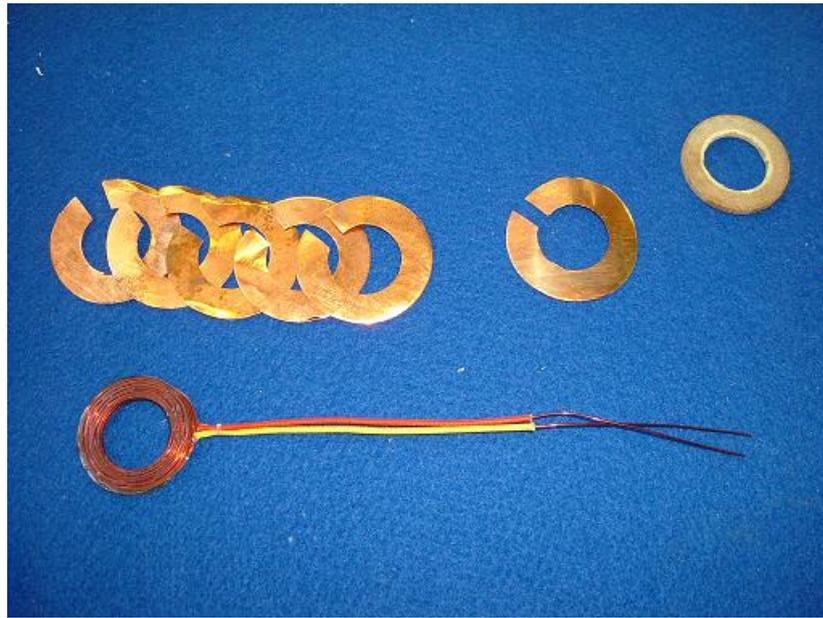


Figure 5.3: A complete stator coil and some copper heatsinking shims. Two kapton shims and four copper shims are placed between each pair of coils to provide electrical insulation and heatsinking capability.

copper shims are designed to pull the heat out of the coils and to transfer it to the cooling jacket.

The kapton shims have an inner diameter of 17 mm; this is smaller than the coil's inner diameter of 18 mm and slightly smaller than the outer diameter of the bore tube at 17.5 mm. This serves to hold the bore tube away from the coils, helping to prevent mechanical damage on insertion and removal of the bore. The kapton seals onto the tube centering the stack and effectively 'pocketing' each coil. The kapton shims are produced with two different outer diameters of 32 mm and 36 mm. The larger diameter of these two shims is equal to the outer diameter of the copper shims.

The copper shims are a split annulus with an inner diameter 19 mm and an outer diameter of 36 mm. They are split to help reduce any potential eddy-current losses in the copper although the effect of any eddy current losses has not yet been established.

#### 5.2.4 Assembly of the Coil Stack

The coil stack is assembled over a former that is of an identical diameter to an insulated stator bore tube. The stack starts off with a spacer, four copper shims and a kapton shim. A coil, lightly coated on both sides with a thermal paste to aid with heatsinking, is added followed by a second kapton shim. Four copper shims are added followed by another kapton shim. The second coil is then added to the coils stack and the

process is repeated for all twenty-four coils. As each coil has a nominal thickness of  $2.85 \pm 0.1$  mm the accumulative error is tracked and minimised during the coil stack construction process. The coils stack assembly process is illustrated in figures 5.4 - 5.5.

The split in the copper shims aids in keeping the coils parallel to each other during construction of the coil stack. This is because the split in the shim is placed where the wires to the coil exit (There is sometimes a small bump at this point on the coil due to the wire exiting from the centre of the coil back over the top of the other windings). The gap created by the split in the shims also serves as a useful point to insert the thermocouples into the coil stack.

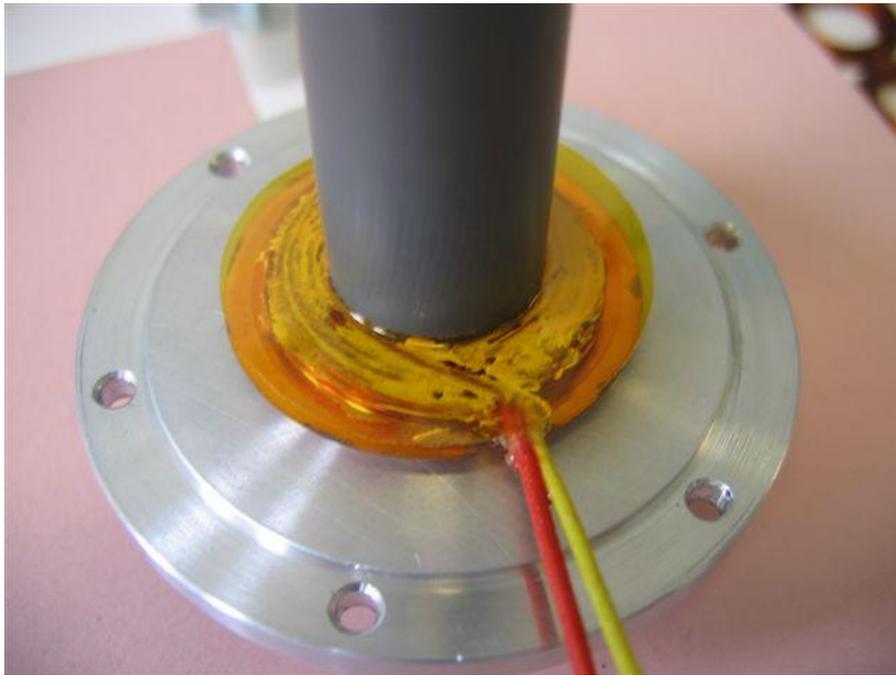


Figure 5.4: Bottom spacer, copper and Kapton shim and the first coil.

At this stage a cooling jacket is slid over the coil stack. The stator body is then completed by adding a split outer jacket and two end-caps. The split in the outer jacket allows the wires and the cooling pipes to protrude for external connection. The end caps provide light compression on the coil stack, keeping it in place. At this stage the tube that the coil stack was formed upon is removed and the bore tube that is welded to the flanges is pushed through the bore of the stator.

### 5.3 Cooling

The cooling of the actuator is extremely important as the rate of heat transfer from the coils to the cooling water ultimately limits the maximum rate at which the target



Figure 5.5: All 24 coils assembled on a former. The copper shims can be clearly seen protruding from the coil stack where they will later make contact with the cooling jacket.

mechanism can be actuated. Typically when the stator is levitating the shaft out of the beam the stator's power consumption is  $\approx 30$  W although every time the target mechanism is actuated (target inserted into the ISIS beam) this adds a heat load of 400 J into the stator coils. The coils are small and therefore the heat capacity of the coil stack is correspondingly low. Without any heatsink the coils would rise in temperature by  $\approx 5$  °C with every actuation. Therefore, if this heat was not removed quickly, repeated actuation of the stator would rapidly result in the temperature of the coils rising above their maximum rated working temperature of 200°C .

Unfortunately, the permanent magnets that are attached to the shaft will not operate up to this temperature without there being a serious risk of demagnetising them. To date it has not yet been possible to model the maximum temperature and current that the actuator will work to without demagnetisation occurring. This is not trivial as the Curie temperature is field dependent. Experience with running the actuator for extended periods of time has demonstrated that with coil temperatures of 80 - 90°C no problems with demagnetisation has been seen to occur. There is some evidence that the risk of demagnetisation at elevated temperatures is accentuated when running permanent magnets in a radioactive environment[95].

### 5.3.1 Temperature Monitoring

As was previously noted, space between the coils is limited and so after consideration of several different temperature sensors the K-type thermocouple was chosen for monitoring the temperature of the stator as it is the smallest and the easiest sensor to interface to. Being just a metal junction, as opposed to a semi-conductor material, means that this sensor should prove to be radiation hard. Typically five sensors are used per stator, three are distributed evenly along the length of the coil stack and the other two monitor the temperature of the water flowing into and out of the cooling jacket.

For the three thermocouples that are located on the coil stack the junction is placed between the copper shims rather than upon the coils themselves; this is for electrical safety reasons. The copper shims are earthed and so by placing the thermocouples onto the shims it is not possible for the thermocouples to become live in the case of a faulty coil. Unfortunately this does mean that the thermocouple does not monitor the temperature of the coils themselves as there exists a small temperature gradient across the kapton shim (see chapter 9). This effect was measured using a specially built coil stack and the effect was shown to be of only a few degrees at nominal working temperatures.

### 5.3.2 Cooling Jacket Design

This early design had a single helical copper pipe wrapped around the outside of the coil stack. The coil stack and cooling pipe was then potted together with epoxy resin as a single unit. No heatsinking mechanism was placed between the coils and the cooling pipe and paper spacers were used to fill the inter-coil gaps. This structure meant that heat from the coils would need to flow out radially from the edges of the coils in the coil stack and then through a few millimetres of epoxy before it could be transferred to the cooling pipe and carried away by the cooling water. It was clear that there was no mechanism for the heat to be effectively transferred out of the coils and into the cooling pipe with this design.

The cooling circuit was redesigned from the coils outwards. Firstly four copper shims and two kapton shims were placed between each coil so that the heat could be conducted directly out of the faces of the coils and into the copper shims as already described. Secondly, a split cylindrical copper jacket was also designed and constructed. This jacket has a thin bore copper cooling tube soldered onto the outside of it so that the jacket can be water cooled by pumping water through the tubing. A drawing illustrating the jacket and the cooling tube is shown in figure 5.6.

The inner diameter of the cooling jacket is slightly smaller than the outer diameter of the copper shims. When the jacket is slid over the coil stack this has the effect of bending the copper shims over ensuring a good thermal contact between the shims and the jacket. A photograph of the jacket placed over the coil stack is shown in

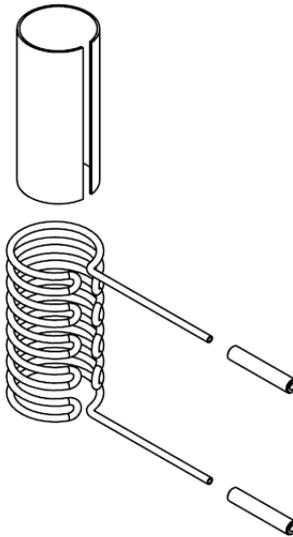


Figure 5.6: The copper cooling jacket is over-fitted with the water cooling pipes. The cooling pipes are soldered onto the jacket.

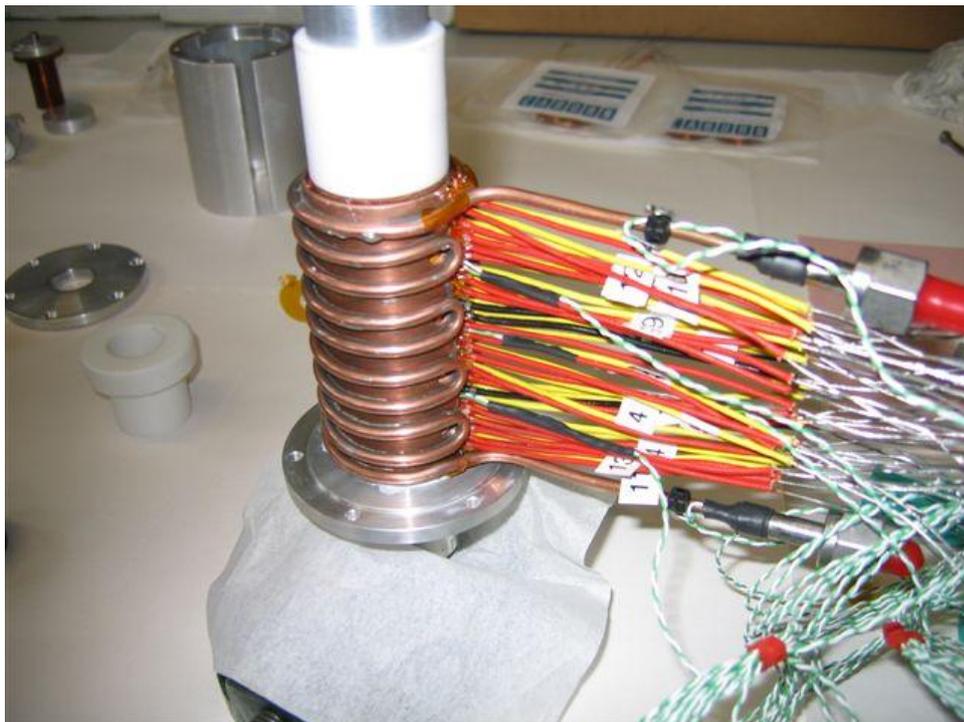


Figure 5.7: The cooling jacket has been placed over the coil stack.

figure 5.7. The cooling pipe has a narrow bore and so the water flow rate is quite low. Typically at  $\approx 4$  bar a flow rate of  $\approx 1$  litre  $\text{min}^{-1}$  is achieved. This water flow rate has proven to be sufficient at removing the heat from both the cooling jacket and coil stack during normal operation. A nominal stator operating temperature of  $80$  °C has been maintained; this is well within the working temperature range of all the components.

## 5.4 Flanges

The stator as described in the previous sections is supported between a pair of flanges; these were illustrated earlier, see figure 5.1. These flanges serve several purposes:

1. They provide firm mechanical support to the stator and the bottom flange provides the interface to the external framework that holds the stator in situ above the ISIS beamline.
2. The flanges contain recesses that hold the bearings; these bearings support and guide the target shaft.
3. The top flange also provides the attachment point for the optical block and the bottom flange attaches to the bellows. The bellows interface the target mechanism to the ISIS beampipe underneath.

The stator is trapped between these two flanges using eight long bolts; great care is taken when this is done to ensure that the flanges are both parallel and coaxial. The central stainless steel bore tube is then welded at both ends to the top and bottom flanges. This weld then provides a vacuum tight seal. These welds fix the position of the flanges but the thin walled bore tube is not strong enough to mechanically support the two flanges; additional support is provided by the eight long bolts. The bottom flange is attached to ISIS through a set of bellows, and metal to metal seals are used to maintain UHV. The bellows and other external supporting components are discussed later on in this chapter.

## 5.5 The Target Shaft

The target shaft must be light, strong and able to withstand high temperatures and large temperature gradients during normal operation. It also needs to provide a suitable contact surface for a sliding bearing and its design must prevent the shaft from rotating during use. The amount of loading on the shaft at this stage of the design had been underestimated and this led to a design with a small second moment of area where the anti-rotation feature had been added. This led to some subtle problems with the stiffness of the shaft. The design of this shaft is presented here, however further analysis of this issue and the corrective measures taken are presented in Chapter 10.

### 5.5.1 Mechanical Considerations

The target shaft and all of the attached components need to experience a high acceleration to meet the timing requirements for insertion into the ISIS beam. For a given force the achieved acceleration is clearly inversely proportional to the sum of the component masses being accelerated, therefore it is essential that the component masses be minimised to achieve the required acceleration of  $\approx 80\text{-}90\text{ g}$ . The mass of the other major component to be accelerated, the magnets, cannot easily be reduced as the material properties of the magnets are clearly defined by their magnetic properties. However it was clear that the design of the shaft would have a significant impact upon the total mass.

### 5.5.2 Description of the Shaft Design

The design of the shaft was undertaken by members of the MICE collaboration at Oxford University[70]. The target shaft itself measures 565 mm long and is made out of a single piece of titanium. The mass of the shaft, with no magnets or optical readout vane added, is  $\approx 27\text{ g}$ . From consideration of figure 5.8 it is possible to observe that the shaft consists of several functional sections. The length of the shaft is determined primarily by the geometry of the stator, optical block and the distance between the target mechanism and the ISIS beam-pipe. The latter distance is primarily determined by the length of the bellows and the gate valve. The function of these two items are discussed later.

The target end of the shaft is a blade measuring  $10\text{ mm} \times 35\text{ mm} \times 1\text{ mm}$ . The last few millimetres of this blade is the section that will intercept the proton beam and the geometry here was initially determined by the required particle production rate for MICE. Above the target, the shaft is cruciform in cross section. This section of the shaft is in contact with the bottom bearing and provides the shaft with mechanical support whilst also preventing it from being able to rotate in the bearing. Towards the top of the cruciform cross section the cruciform widens as illustrated in figure 5.8 of the shaft. The wider cross section acts as a stop preventing the shaft from falling through the bottom bearing. When the stator is in its upright position and the system is un-powered the bottom of this stop naturally rests upon the upper surface of the bottom bearing.

A little further up the shaft is the magnet stop. The permanent magnets are fitted onto the shaft by first inserting the top of the shaft down the centre bore of the magnet assembly and then pushing them up to this stop. A clip is fitted on top of the magnets and this is secured by tightening the clip up into a groove that is cut into the shaft above the magnets. This mechanical clip has to be tight up against the magnets to prevent any movement of the magnets when the shaft is accelerated. A mechanical clip was chosen, as this reduces the amount of glue in the ISIS beam-pipe and allows the magnets to be easily recovered when the shaft reaches the end of its operational life.

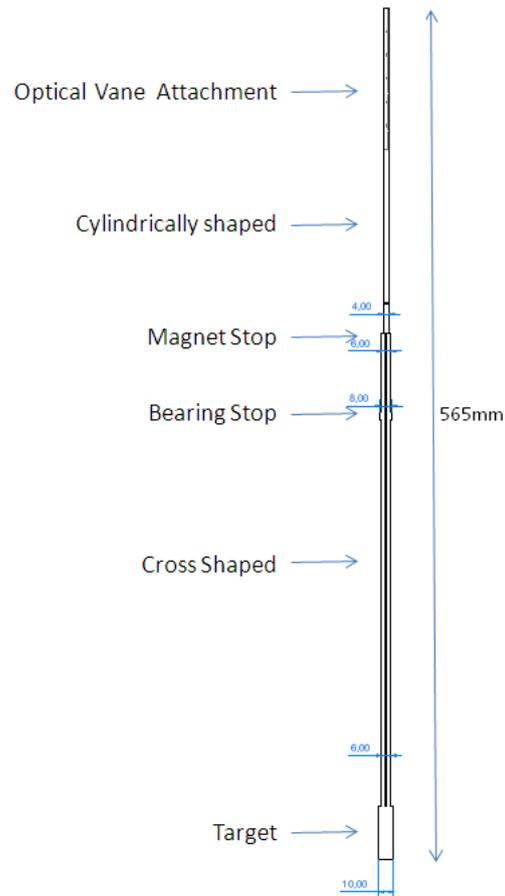


Figure 5.8: Drawing of the target shaft

Above the magnet section the shaft turns cylindrical and this is where the shaft moves against the upper bearing; this bearing is seated between the upper flange and the optics block. The top of the cylindrical section of shaft, where it never passes the top bearing, is split; this is designed to hold the vane which is used for the optical readout.

The optical vane, which will be discussed in more detail later, is a delicate component that has to be manufactured separately from the shaft. This is placed in the slit cut into the top of the cylindrical section of the shaft and is then mechanically secured with five screws that pass through the shaft, the optical vane and then the shaft again. Once again the method of attachment gives physical security, allows the component to be recovered and eliminates the worry of radiation damage and outgassing if an epoxy was to be used.

After manufacture the shaft itself is coated with DLC (Diamond Like Carbon) on the sections that will be in contact with the bearings. This material is coated on to the shaft at these points to reduce shaft and bearing wear therefore extending the operational

life of the shaft and bearing components. The bearings themselves are also coated.

The shaft is very long in comparison to its cross-sectional dimensions, and this means that the shaft will flex along its length if a lateral force is applied. Unfortunately the cruciform cross-section has a small ‘second moment of area’ and so presents little resistance to bending. The amount of loading on the shaft at this stage of the design had been underestimated and, as will be examined in chapter 10, this issue created a subtle problem that has resulted in the shaft having to undergo a significant redesign to increase its stiffness.

### 5.5.3 Target Heating

Before discussing the magnetic assembly it is worth considering another effect that is pertinent to the shaft’s design, namely target heating. There are two distinct issues with the heating of the target mechanism. The first source of target heating is the ohmic heating due to the passing of an electrical current through the stator coils. This is primarily an issue for the stator, although heat from the stator is radiated onto the magnets and so the stator temperature needs to be kept as low as possible. Further consideration of this is discussed later on in the thesis.

However, of interest at this point is beam heating, the passage of the protons through the part of the target that intercepts the ISIS beam. This is an important consideration for two reasons. Firstly target heating by this mechanism could have important consequences with respect to the mechanical properties of the material as it is important that the target tip does not melt under both normal and abnormal operating conditions. Secondly there was the concern that conduction of the heat through the target and up into the motor may cause additional heat loading on the motor components. A study of the effect of beam loading was undertaken by Dr Booth and the simulation showed that under normal operating conditions (actuating the target at 1 Hz into the ISIS beam) then there should be no more than 1.9 W of particle heating in the titanium target [96]. This was shown to correspond to the target running at an equilibrium temperature of 360 °C at the tip of the target.

The study also modelled the heating of the target under a worse case fault condition; this is when the target is dropped into the ISIS beam and remains there until ISIS is turned off. Under such circumstances the ISIS accelerator would detect large amounts of beam loss and it would automatically turn the machine off within three spills. In such a case the target would reach a maximum temperature of 815 °C. It can be seen from these results that the temperature of the target does not come anywhere near the melting temperature of titanium either during normal operation or during a fault condition.

Titanium has the property of being a poor conductor of heat; its thermal conductivity is  $21 \text{ W}^{-1} \text{ m}^{-1} \text{ K}^{-1}$ . The poor heat conduction of titanium turns out to be a useful characteristic of the metal in this application because it helps to reduce the rate of

heat transfer from the tip of the target thereby preventing the magnets from being excessively heated from conducted heat. The estimated temperature profile of the target shaft (for a design that is shortly to be discussed) after insertion into the ISIS beam as a function of distance from the target has been simulated is shown in figure 5.9[96].

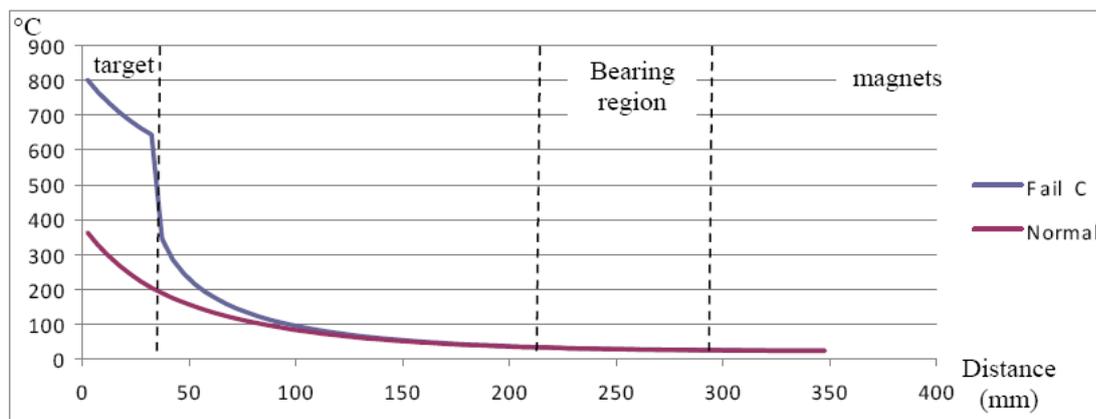


Figure 5.9: The temperature of the target (first 35 mm) as a function of distance up the target shaft for normal (bottom line) and failure (top line) conditions. The graph shows that the target temperature doesn't become excessively hot under both normal and failure conditions and that there is minimal conduction of heat up to the bearings and the magnets[96].

#### 5.5.4 Diamond Like Carbon

Diamond Like Carbon (DLC) is a coating applied to materials that replicates some of the properties of diamond. The main advantage of using such a coating is that DLC makes the surface of the coated material much harder, thus significantly increasing the wear resistance of the material. DLC is applied to a material by a vapour deposition process and this coats the underlying material with a layer of DLC that is typically 3-5  $\mu\text{m}$  thick. Unusually, it was recommended that DLC be used on both contacting surfaces[97]. (More usually two dissimilar materials are used and the less hard one wears.)

The properties of this material make it ideal for coating the moving parts of the target mechanism given the harsh environment in which it has to operate. A sample of DLC was sent to RAL for vacuum testing and was approved for use on the target mechanism. DLC is now used on the parts of the target shaft that contact the bearings and the bearings themselves to prolong these components' lifetimes.

DLC can denature at temperatures greater than 400°C. The heating studies described

in the last subsection demonstrated that the only component that could possibly reach a temperature in excess of  $400^{\circ}\text{C}$  is the end of the shaft that is the target blade. It is possible to prevent this part of the target shaft being coated in the vapour deposition process by covering it with a suitable mask.

### 5.5.5 Magnet Assembly

The permanent magnet assembly is responsible for interacting with the stator coils to produce the force on the shaft that accelerates the target into and out of the ISIS beam. The details of the interaction of the magnets with the coils will be considered in chapter 7. The purpose of this section is to describe the mechanical construction and the magnetisation.

As has already been indicated the magnet assembly is composed of 24 segments of Neodymium-Iron-Boron (NdFeB) magnets arranged into 3 rings, 8 magnets per ring, glued onto a steel core to produce a cylinder that is 18 mm long with a outer diameter of 15 mm and an internal diameter of 4 mm. Thin ceramic washers separate the three rings. A drawing of the construction of the magnet assembly is shown in figure 5.10. As can be seen in the diagram the central magnetic ring at 7.8 mm deep is about twice as deep as the outer rings at 3.9 mm. The whole assembly has a mass of  $\approx 25$  g.

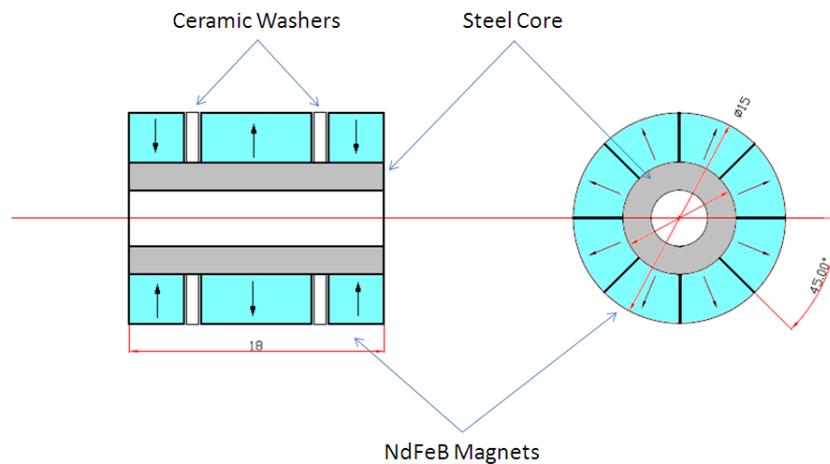


Figure 5.10: Drawing showing the structure of the magnet assembly

The individual sectors are manufactured by wire eroding them from the un-magnetised magnetic material. The sectors are then magnetised diametrically and then assembled in a jig before finally being glued into position using a strong two component epoxy. The middle sector is magnetised so the outer surface is a North pole and the upper and lower sectors are magnetised with the outer face as a South pole. As the middle magnetic sectors are magnetised in the opposite sense to the upper and lower rings

circular flux lines are produced, see figures 5.11. The peak surface field of the magnets  $B_r$  is approximately 0.6-0.9 Tesla, although the value of the field is position dependent.

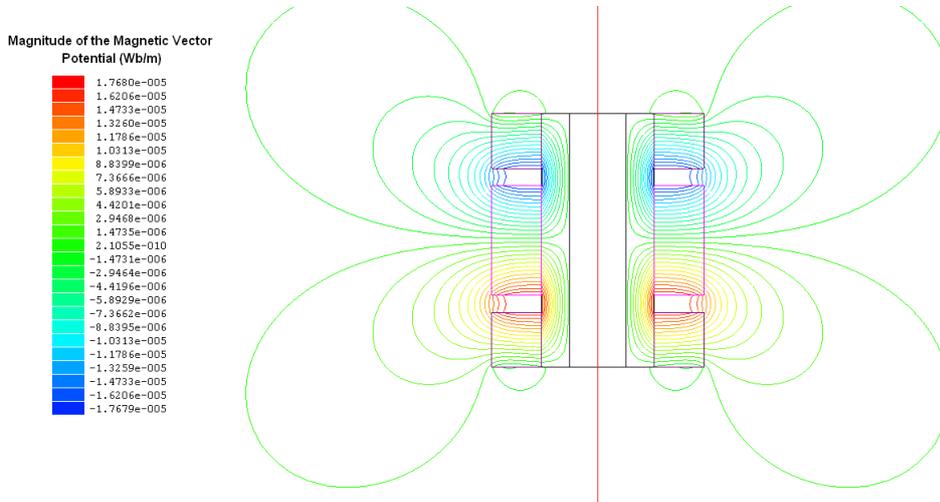


Figure 5.11: Diagram showing the flux lines from the permanent magnets as simulated by MAXWELL 2D Student Version [98].

To date the magnet material that has been used is NdFeB as this magnetic material shows the highest remanence field of all permanent magnets, allowing the highest force to be achieved on the shuttle for a given coil current. However, concerns about the temperature dependance of these magnets and their radiation intolerance [95] could mean that Samarium Cobalt (SmCo) magnets will need to be considered in later designs. Unfortunately SmCo has a lower field strength than NdFeB so a higher coil current would be required to achieve the same acceleration. This is an area in which further work will be necessary; the determination of any radiation induced demagnetisation will require magnets that have been exposed to the ISIS beam for a considerable period of time. The BH curve for the magnetic material used in the NdFeB magnets is shown in figure 5.12[99].

### 5.5.6 Optical Vane

The optical vane and the optical block are two mechanical components in the optical position monitoring system. Position sensing is performed using a quadrature system viewing an optical vane that is mounted on the shaft. This system gives the position of the shaft to an accuracy of  $150 \mu\text{m}$ . The position of the shaft must be known accurately in order that the trajectory of the target can be both monitored and controlled. A full description of the optical system and how this vane and the optical block work together along with the other optical components to provide positional information on the target shaft is given in chapter 6.

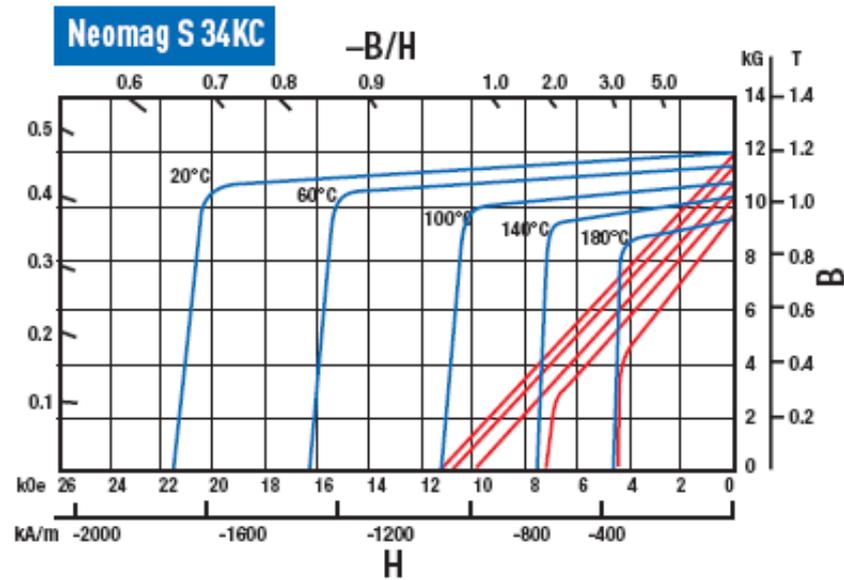


Figure 5.12: Diagram showing the BH curve for the magnetic material that is used in the magnet assembly. The blue curve represents the magnetisation, whereas the red curve represents the flux density at the surface of the magnets. Note the temperature dependence of the maximum demagnetisation field that can be tolerated[99].

A diagram of the vane is shown in figure 5.13. The vane is a wire-eroded double-sided ‘comb’ manufactured from 0.2 mm thick steel, having 157 teeth 0.3 mm wide and 3 mm long on one side of a 6 mm wide spine. The whole vane is 93.9 mm long. The gaps between the teeth are also 0.3 mm. There is a single tooth two-thirds of the way down the vane on the other side that is also 3 mm long and 0.3 mm wide. The vane has a mass of  $\approx 1$  g.

The fine detail of the vane required that it was manufactured using the technique of wire erosion. It is important that the manufacturing tolerances of the teeth were kept minimally small along its length; less than 5 % of  $150\ \mu\text{m}$  deviation on each tooth, and that the tracking error along its full length was also less than this percentage.

The optical block is a mechanical assembly that holds three laser channels that allow the vane to be read by the position monitoring system. Two of these laser channels are focused onto the side of the vane with many teeth and these two channels provide relative positional information. The third channel is focussed on to the other side of the vane. The single tooth provides an index position to which the positional information provided by the other two channels can be referenced, giving an absolute position. The optical block is mounted on the top flange of the stator as can be seen in figure 5.2 from chapter 4. A more detailed discussion of this component and how it fits into the position monitoring system is left until chapter 6.

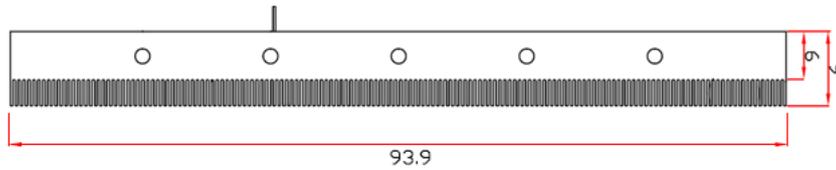


Figure 5.13: Drawing of the Optical Vane. One side of the vane has 157 teeth, these are used to provide the quadrature signal, whilst the other side has just a single tooth that is used to provide a fixed reference point. Dimensions in mm.

## 5.6 Bearings

The following subsections give an overview of the mechanical features of the two sliding bearings that were used in the first target mechanism that was installed on ISIS.

### 5.6.1 Bottom Bearing

The bottom bearing has an aperture that is cruciform in cross section. The aperture is slightly wider than the cruciform cross section of the shaft to allow the target blade to pass through the bearing on assembly of the target mechanism. The cruciform cross section successfully prevents both lateral and rotational movement of the shaft whilst allowing the shaft to move freely longitudinally.

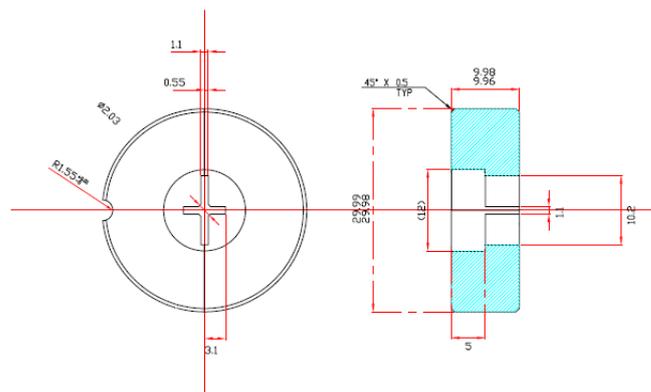


Figure 5.14: The bottom bearing is constructed in quadrants. Note the cruciform cross section that corresponds to the cross section of the shaft and prevents the shaft from rotating.

The bearing, see figure 5.14, is manufactured from four quadrants of steel and DLC

coated. The quadrant structure serves several purposes; it eases construction of the bearing and it permits coating of the inner bearing surface with the DLC. The coating would not be possible if the bearing was constructed from a single piece as the DLC would not penetrate into the aperture. Before the need for the DLC coating was recognised the bottom bearing had been manufactured both as a split bearing and as a single piece. (Although in the latter case this was with a prototype shaft that had a square cross section as opposed to a cruciform cross section.) The bearing is prevented from rotating in its housing by the use of a pin that passes through the edges of two of the quadrants and into the bottom flange.

### 5.6.2 Top Bearing

The upper bearing has a circular aperture that matches the circular cross section of the upper part of the shaft. This provides the shaft with the necessary support but clearly it does not provide any assistance in preventing the shaft from rotating. This bearing is manufactured from two pieces of steel. The split construction, like the bottom bearing, permits coating by the DLC. Just as important, this split design aids assembly as this bearing cannot be put into place until after the shaft is located inside the stator. If this bearing was not split it would not be possible to locate the bearing around the shaft. The circular symmetry of the top bearing has meant that there has been no need to pin it into the bearing housing.

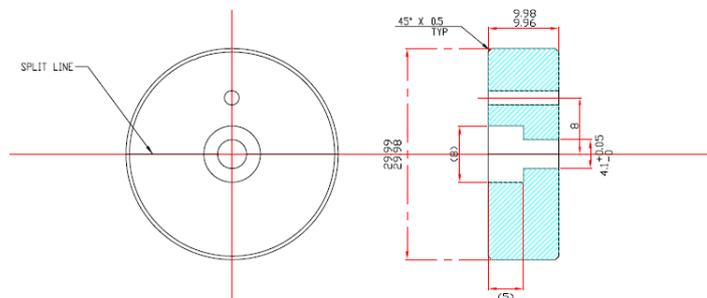


Figure 5.15: The top bearing is constructed in two pieces.

## 5.7 Mechanical Support and the Beam-Pipe Interface

This section gives some details of the surrounding mechanical support for the actuator and how the actuator interfaces to the ISIS beam-pipe. Much of the design and construction of the work surrounding mechanical support has been done at RAL so

the interface between the target mechanism and ISIS was defined quite early on in the design process.

### 5.7.1 A-Frame

The main structural support for the target mechanism is a three legged stand that straddles the ISIS beam pipe and is called the A-frame. The A-frame is supported upon three concrete pillars and once the frame has been mounted on the concrete pillars it stands about 2 m high. See figure 5.16.

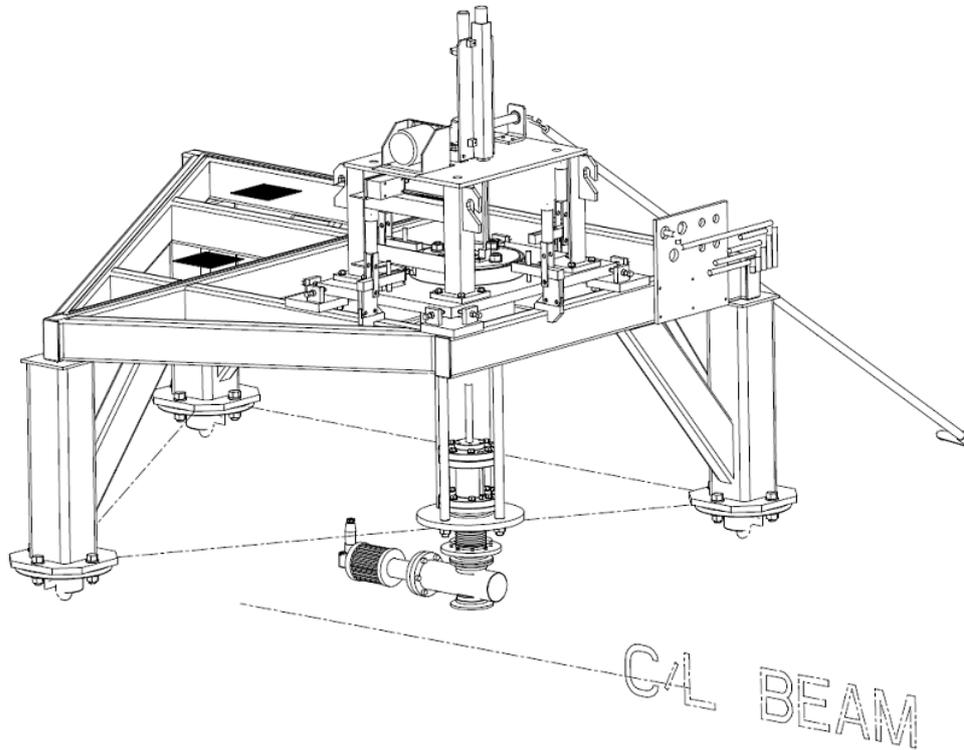


Figure 5.16: A drawing of the A-frame, small frame, target, bellows and the gate valve. The small frame sits within the A-frame and can be easily removed with a crane. RAL drawing [100].

The A-Frame provides a receptacle for accurately positioning and retaining a smaller frame. This smaller frame contains all the infrastructure for holding and moving the target mechanism itself. As can be seen from diagram 5.17, four pillars descend from a plate from within the smaller frame. The bottom flange of the target mechanism is attached to the bottom of these pillars.

The vertical movement of the upper plate, and hence the pillars, is controlled by a stepper motor that is housed on the top of the frame. Therefore raising and lowering

this plate has the effect of raising and lowering the target mechanism. The system is designed to operate in one of two positions; fully raised and fully lowered. The lowered position corresponds to the position in which the target will sit when operating. If the target was left un-powered in the lower position the target would be left in the beam path and ISIS could not run. The raised position is the position that the target will be in when it not in operation. This means that the target can be left in the upper position un-powered and ISIS can still run. In the case of a fault with the target system the target can be raised into this position until it is possible to gain access to service it. The interface between the target mechanism and the beam-pipe will be described shortly.

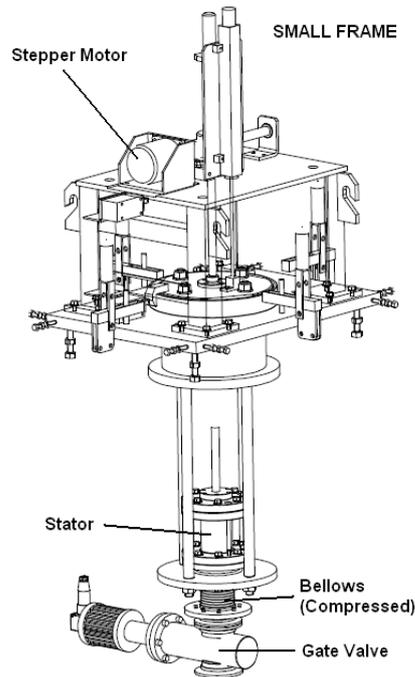


Figure 5.17: A close-up of the ‘Small Frame’ that sits inside the A- Frame. The Small Frame makes the attachment to the target’s lower flange as illustrated [100].

### 5.7.2 Beam-pipe

The beam-pipe that the target attaches to forms part of the ISIS ring in super period seven. The beam-pipe itself is a tapering rectangular box section, flanged at both ends, that is big enough to accommodate the ISIS beam. Additionally a ‘window’ has been added to the box section. This ‘window’ has been placed in the path between where the target intercepts the ISIS beam and the first three quadrupole magnets in the MICE beamline. The ‘window’ is a section of steel that is thinner than in the rest of the box section. This reduces both scattering and the  $dE/dx$  losses of the particles that

traverse it.

A flanged port has also been added to the top of the beam pipe that allows the target mechanism to be inserted into the ISIS beam. The target mechanism is not attached directly to this point as some means of allowing vertical movement of the whole mechanism must be added between the two. Below the beam-pipe is a second port that acts as a collection pot beneath the target mechanism. In the unlikely event of any debris being created by the target mechanism this should be captured in this pot. This may help to prevent obstructing the ISIS beam which would otherwise result in a costly intervention. There are some other ports on this section of beam-pipe that are for use by ISIS but these are not relevant to the target mechanism.

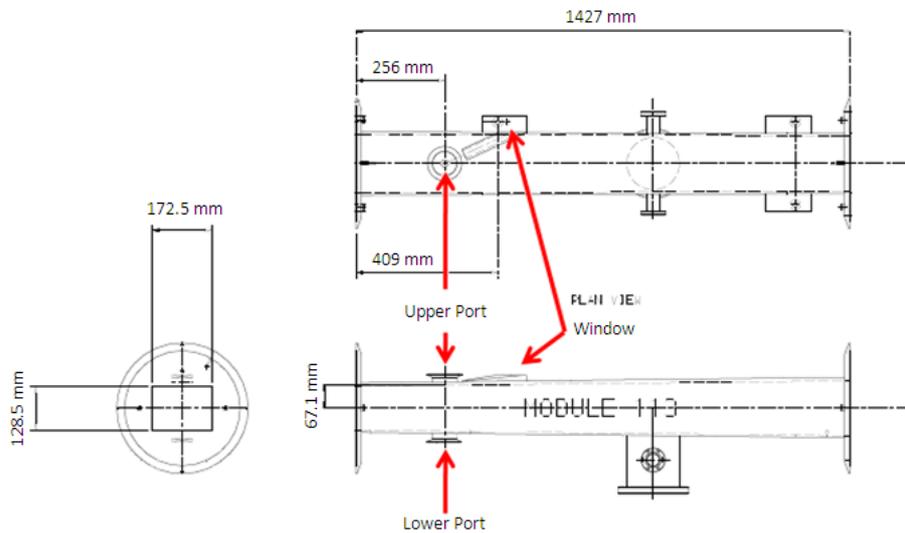


Figure 5.18: Drawing of the beam-pipe section to which the target is attached[100].

### 5.7.3 Bellows and the Gate Valve

The target does not attach directly to the beam pipe, the attachment is made through a set of bellows and a gate valve. The set of bellows provide a vacuum tight mechanism that freely compresses and decompresses along its longitudinal axis. This allows the whole target mechanism to be raised and lowered using the stepper motor on the smaller frame without compromising the ISIS vacuum.

The gate valve acts to isolate the ISIS vacuum from the target vacuum thereby creating a physical barrier between the two systems. The lower flange on the gate valve is attached directly to the upper port on the beam-pipe and the upper flange on the gate valve attaches to the lower flange on the bellows. If a problem occurs with the target

or the target vacuum then the target system can be completely isolated from ISIS by means of this gate valve until a repair can be effected. The gate valve is interlocked with the target frame so that it cannot be closed when the target frame is in its lowered position. Closing the gate valve when in this position would damage both the gate valve and the target as it attempted to close onto the target shaft.

## 5.8 Demonstrator System

Due to the expected finite lifetime of the bearings it was decided that it would be necessary to have two separate and complete target systems built and installed at RAL. One of the systems would be installed in ISIS whilst the second system would operate in an assembly hall, building R78. The second system, the demonstrator system, would be accessible 24 hours a day. The demonstrator system should be identical to the system mounted in ISIS and it was to be run three times as many actuations ahead of the system installed on ISIS. This way the demonstrator system could act as an early warning system for any potential gradual failure, with particular emphasis on acting as an indicator of the failure of the bearings.

Because of an unforeseen stator failure due to an electrical problem the second system was not installed in R78 when the primary system was installed on ISIS. Without a second system the primary stator installed on the ISIS beam-pipe has been run cautiously. A more thorough review of the problems experienced with the stators and the remedial action that has been taken will be discussed in Chapter 10.

## 5.9 Conclusion

This chapter has looked primarily at the mechanical construction of the target mechanism and how it interfaces to ISIS. The next chapter will look at the development of the optical position monitoring system, a vital subsystem for enabling accurate control and monitoring of the movement of the target shaft.



## Chapter 6

# The Optical Positioning System

It was apparent early on in the design that a positional monitoring system would be required to track the position of the target shaft at all times. This chapter describes the development of the positional monitoring system and how it works.

### 6.1 Position Monitoring

The movement of the shaft is tightly constrained by the bearings and for all intents and purposes its motion can be considered as one dimensional. The position of the shaft with respect to the stator needs to be carefully and reliably monitored as this positional information is essential for the correct operation of the target mechanism. Namely

1. The positional information is used in the controller as feedback to ensure that the correct coils are powered to either hold or actuate the target.
2. The position of the target can be monitored and recorded by the Data Acquisition (DAQ) System.
3. Long term reliability and stability can be accurately monitored.
4. It provides part of a fault detection system. If the target isn't where it is supposed to be then the controller can flag an error or initiate a corrective measure.

Motion along a single axis lends itself very well to using a single axis quadrature system to monitor the target's position. Quadrature is a method that is used extensively in many motion devices as it provides a reliable way of ascertaining positional information for a relatively small footprint in terms of space and mechanical complexity. For this reason a quadrature system was considered a good way of obtaining the positional information for the target shaft.

## 6.2 Description of a Quadrature System.

A quadrature encoder is an electro-mechanical device that converts either linear or rotary movement into an electrical pulsed signal. A quadrature system is usually composed of a total of three transducers and an electronic counter that work together to give both the direction of motion and the absolute position of the moving components.

In such a system two of these transducers ascertain both the relative movement and the direction of movement of the system. These transducers are referred to as channel A (CHA) and Channel B (CHB). A third transducer is then often used to give a zero reference point or an index marker and in this thesis this channel is called the index channel. All three of these channels give a digital output so each channel can be in one of two states at a given instance; high or low. The distance that the system has to move for a change of state to be observed on either of the transducers CHA or CHB defines the resolution of the system.

The reason that this system allows position and direction to be monitored is because the transducers CHA and CHB are arranged so that there is a phase difference between the outputs of  $90^\circ$ . Motion of the system under measurement in one direction will give an output signal where the output of CHA will lead CHB; motion in the other direction will give an output where the output of CHB will lead CHA. The two outputs from these channels can be interpreted as a two bit gray code giving four discrete states. Each change of state represents a given measurable movement of the system. The direction in which these states change gives the direction. By counting these state changes (taking care to account for the direction) this gives the total distance travelled by the system. Figure 6.1 shows how this works diagrammatically.

The third channel, the index marker, acts a reference point that can reset the counter. If the reset point is at a known location on the system then multiplying the count by the resolution of the quadrature system gives the total distance moved. The index also acts to remove erroneous counts that could otherwise build up and create an offset error.

Any quadrature encoder requires an electronic interface at some point in the system to convert the signals produced from the quadrature transducers to digital signals for summation in an electronic counter. The high levels of radiation that are likely to be observed at the actuator [95] would makes the long term use of electronic circuitry at the actuator difficult to achieve. Clearly, any system utilised on the target mechanism would require that the electronics be well separated from the actual sensing mechanism. The sensing mechanism itself needs to be radiation hard and also able to send the quadrature signal to any external electronics without degradation of the signal. Noise could be a problem as any errors in the positional information could lead to the control of the target being lost.

Some time was spent exploring the possibility of using a commercially available quadrature encoder that could be adapted to measure the position of the target shaft. A



commercially available solution would have been ideal because the manufacturer would have spent considerable time developing and debugging their system; this would have allowed the acquisition of a quadrature counter for the target system that would have minimised development time. Unfortunately there was no commercially available solution that fulfilled the necessary requirements because all of the quadrature encoders had the necessary electronics embedded into the device at or close to the read head. The result of this was that an optical quadrature system had to be developed specially for the target mechanism.

### 6.3 An Overview of the Quadrature Optical Position Measuring System

A cartoon of the optical position measuring system designed for the target mechanism is shown in figure 6.2. Three 635 nm lasers are used to provide a light source and this light is passed to three collimators via single mode optical fibres. Once the light has been collimated it is then focused onto the vane that is attached to the top of the target shaft. A mechanical description of the vane was given in section 5.5.6. The mount that holds the collimators and the lens system is called the ‘Optics Block’.

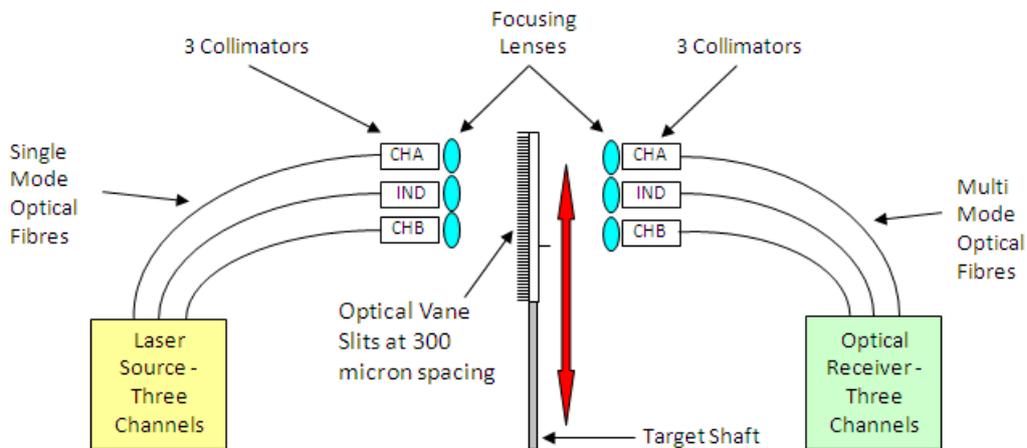


Figure 6.2: A cartoon illustrating the key components to the optical quadrature system used to ascertain the position of the target shaft.

Two of the channels, Channel A and Channel B, are focused onto the comb side of the vane whilst the Index Channel is focused onto the side of the vane with the single tooth. The focussing is tight; the spot size of the laser light is smaller than the half pitch of the teeth on the vane. This means that a single tooth, when suitably positioned, can completely prevent the light from a given channel from passing through to the receiving collimator. As the target shaft moves vertically up and down inside the stator the vane

moves with the shaft. Consequently the focused points of laser light are alternately blocked and passed through the comb. Channel A and Channel B are focussed so the relative positions of their focal points are  $90^\circ$  out of phase with respect to the teeth on the vane. This phase shift in the relative position of these two collimators results in a ‘optical’ quadrature output where the quadrature signal can be interpreted from whether light is being passed (or not) through the vane on these two channels. The single finger on the index channel is positioned two-thirds of the way down the vane and this gives an absolute position reference.

The light that is passed by the vane is picked up by a separate lens and collimator for each channel and is then passed down multi-mode optical cables to three individual optical sensors. The optical sensors convert the light into small electrical signals which are then amplified and conditioned to produce the three quadrature outputs that are required for passing into a digital counter. The counter is a separate integrated circuit that is specifically designed for counting quadrature signals. The quadrature chip when interfaced with a suitable microprocessor outputs a binary number that can be interpreted as the position of the shaft.

### 6.3.1 Theory

For this system to work the spot size at the focal point must be smaller than the pitch of the graticule. Ideally, to obtain a nice clean signal, the focal point should be less than half the pitch of the graticule. This then allows for a clean pass/no pass signal to be detected as the graticule moves through the plane of the focal point. For the current system this entails having a focal point that is  $< 150 \mu\text{m}$  in diameter. Ideally a focal point of less than this is preferred as some slight rotational and lateral movement of the graticule in and out of the focal plane does occur due to the tolerances in manufacture and construction of the target bearings.

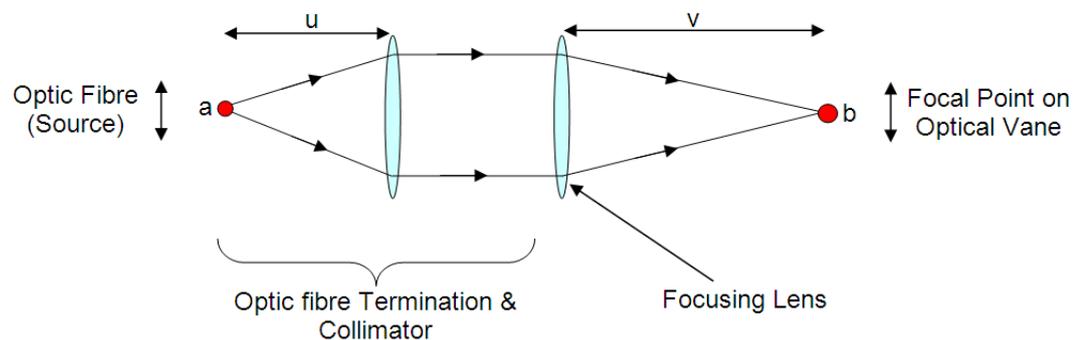


Figure 6.3: Diagram showing how the collimator and lens system works magnifying the optical fibre source to create a finite spot size on the optical vane.

The simplest optical way to achieve a focal point is to focus a parallel beam of light using a convex lens. The collimator creates a parallel beam from a diverging point source of light, this source being the end of the optical fibre that is attached to the collimator. The focusing lens and the collimator effectively work together as a two lens system as illustrated in figure 6.3. This optical setup will magnify the source point by a factor of:

$$b = \frac{v}{u}a$$

where  $b$  is the image size,  $a$  is the object size,  $u$  is the object focal length and  $v$  is the image focal length.

The object focal length is set by the collimator and for the collimators that are used their focal length is 15.3 mm. The image focal length is determined by physical constraints as there is a limit as to how close to the graticule these components can be placed. The vane must be inside the ISIS vacuum and the collimators and lenses must be out of the vacuum on the other side of an optical window that separates them. Initially a distance of 36 mm was used for a prototype optics block; later this needed revising out to 45 mm because of mechanical considerations with the design of the optics block. Using the latter figure this clearly gives a magnification factor of three and so the object source must be smaller than 100  $\mu\text{m}$  in diameter to satisfy the criteria of having an image size < 300  $\mu\text{m}$ .

Creating an object source with a diameter of less than 100  $\mu\text{m}$  diameter whilst still having sufficient light intensity to allow a detectable signal to pass through the system would be difficult to achieve with a conventional light source. The use of a laser coupled to a single mode fibre solves this problem as the end of the fibre then acts as a very small source. As collimators are readily available to attach directly to single mode fibres this allowed the optical system to be constructed as described earlier.

A measurement on the achieved spot size at the focal point on the vane using 36 mm focal length lenses showed it to have a FWHM of  $\approx 75 \mu\text{m}$ . This translates to the source of light at the end of the fibre being less than  $\approx 25 \mu\text{m}$  diameter.

## 6.4 Details of the Optical Design

The use of optical fibres confers an enormous advantage to this position measurement system. The very low attenuation of the light transmitted within the fibres allows a significant distance to be placed between the source, receiving electronics and the optics block. In fact, this is utilised to keep the source and the electronics out of the synchrotron altogether. This mitigates any concerns about radiation damage to the electronic components and keeps them accessible at all times. In fact, both the laser source and the receiving electronics can be kept in the MICE local control room along

with the target control electronics, a distance of 100 m away from the target mechanism itself.

#### 6.4.1 Laser Source

In order to assist the setting up and calibration of the optical system a visible laser was chosen. This also has the added safety implication that it is easy to see when the beam is on. Laser sources at 635 nm wavelength with a controlled output power from 0 to 2.5 mW were purchased. A relatively high light output is required because the optical losses in the system are quite high; there are a significant number of optical interfaces, including fibre to fibre connectors, and each interface loses light. Another reason for maximising the light intensity through the system was because the photodiodes that convert the light to an electrical signal do not generate a large output voltage, typically it is of the order of a few mV at best. Maximising the light on the sensor maximises the devices output voltage, this in turn reduces the signal to noise ratio and reduces the amount of amplification needed; this significantly eases the design of the electronics.

In operation these laser units do not need to be operated at their maximum power. Typically the light power is set to about 1 mW and the gain of the electronics is then set to obtain the best required response from the circuits. This setup leaves overhead in both the laser light-source and the electronics in case any future adjustments need to be made. This facility could be important as it expected that the fibres close to the the synchrotron could suffer some radiation damage over time attenuating the light signal.

#### 6.4.2 Optical Fibres

Two types of optical fibre were used in the system. On the transmitting side single mode fibres were used as they give the required small spot size at the collimator focal point ensuring a small spot size at the vane as was illustrated earlier. On the receiving side 200  $\mu\text{m}$  multimode fibres were used.

The single mode fibres that were used are SM600 with FC/FC connectors. The SM600 fibre is constructed from NuFern 630 nm fibre [101] and has a pure silica core. Silica is considered a radiation hard material and so these fibres are ideal for the location into which they will be going. Unfortunately no information is available on the proton/neutron dose that these fibres can withstand before the the fibres suffer attenuation losses due to radiation damage.

If single mode fibres were to be used on the receiving side the collimators would have had to be aligned to a much higher precision as the tolerance on alignment increases with a decreasing fibre size. Using single mode fibres could have made a system that was difficult to set up and that could have easily gone out of alignment. Multi-mode fibres make the alignment process much easier, although it is still a time consuming

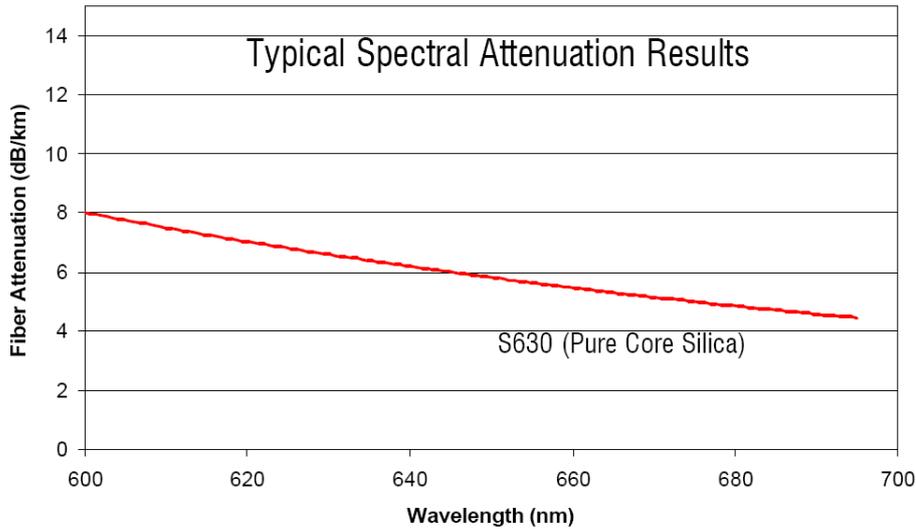


Figure 6.4: The attenuation of the SM600 single mode fibre as a function of frequency [101].

process. 200  $\mu\text{m}$  fibres were chosen because these minimised the cost per metre of fibre whilst still allowing good light collection and making it relatively easy to align the system.

The multimode fibres that are used are BFH37/200 with SMA to SMA connectors. These too are constructed from a 200  $\mu\text{m}$  silica core with a fluorine silica cladding. Like the single mode fibers the silica construction helps to ensure that the fibres have a higher tolerance to radiation damage. Figures 6.5 and 6.6 show the construction of the fibre and the attenuation as a function of wavelength respectively [102]. It can be seen from the latter that the fibre is particularly suited to operation at 635 nm.

### 6.4.3 Collimators and Lenses

Figure 6.7 shows a diagram of a typical collimator. The transmitting collimators use an FC/FC connector and the collecting collimators use SMA connectors as this eases identification. The ability to connect these collimators directly to optical fibres and their compactness has significantly simplified the design of the system. A collimator with a focal length of 15.3 mm was chosen for this application; collimators with shorter focal lengths are available. The longer focal length minimises the magnification of the focal spot reducing the final spot size. The trade-off for this is that they do decrease the numerical aperture of the collimator, theoretically making alignment a little harder. Experience has shown that this has not proven to be a problem with the current mounting system.

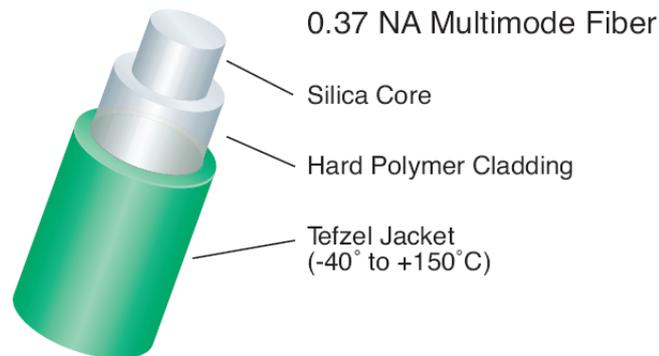


Figure 6.5: The construction of the multimode fibre. The silica core gives more radiation resistance than polymer cored fibres [102].

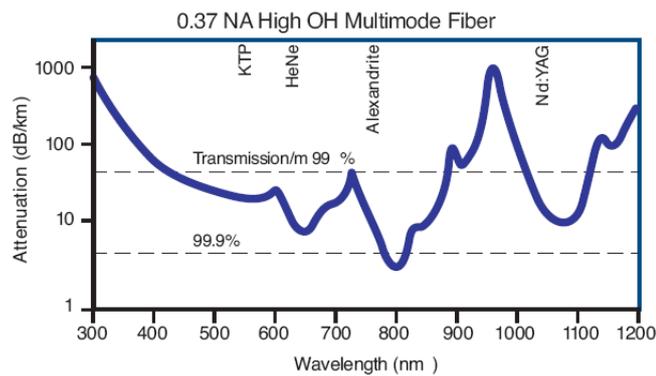


Figure 6.6: The attenuation of the BFH37/200 fibre as a function of frequency [102].

The focusing lenses are attached to the front of the collimators by inserting the lens into a lens holder which is then screwed onto the front of the collimator. On the transmitting side aspheric double achromatic lenses are used to obtain the minimum spot size. On the collecting side, where the focusing is less critical, double convex lenses are used to re-collimate the beam. The lenses are  $\text{MgF}_2$  coated to minimise reflections and maximise the light transmitted into the fibres.

#### 6.4.4 Optical Vane

The optical vane is the mechanical device that blocks and passes the laser light as it moves transversely through the focal points of the laser channels. This component was

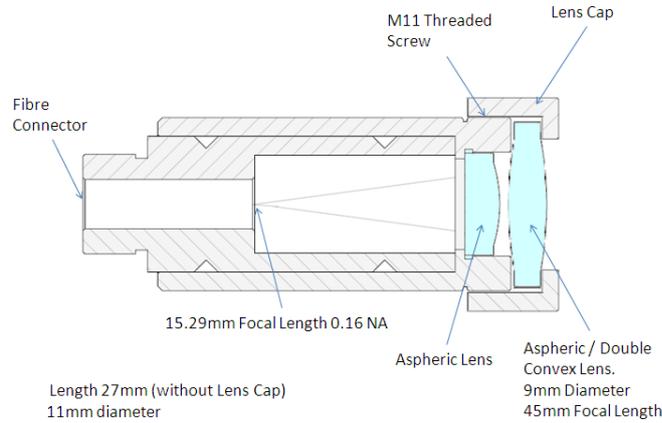


Figure 6.7: Diagram illustrating a cross section through a collimator and a lens cap. The collimators are individually aligned at the factory and are precision pieces.

discussed in detail in the previous chapter in section 5.5.6.

#### 6.4.5 Mechanical Mount - The Optics Block

Although the optical system was good in theory the constraints both in terms of the physical space required to implement the system and the requirements for a system that sat outside of the ISIS vacuum meant that the mechanical mounting system, the optics block, went through a significant evolution from the first prototype to the final system that is in use today. A significant amount was learnt from the prototype so details of this are also included here.

#### 6.4.6 First Prototype Mount

After some initial bench tests on a single channel to ensure that the principle worked, a prototype mount was constructed from a solid block of plastic. Three 12 mm diameter bore holes were machined horizontally parallel to the long axis of the block to hold the six collimators in three opposing pairs. These holes were tapped into from the outside so that each collimator would rest on 6 grub-screws. Enough clearance was given around each collimator so that these grub-screws would provide the mechanism by which the collimators could be aligned. The focal length of the focusing lenses was 36 mm. A 15 mm diameter hole was drilled vertically down through the block from the top so that it intercepted all three of the other bores holes, perpendicularly to them. The target shaft ran vertically through this last hole so that the vane passed through the focal point of all three collimator channels. Initially the target was run in

air although later the top of the shaft was placed inside a glass tube. This tube fitted into this vertical bore hole so that vacuum condition could be maintained on the target shaft during operation.

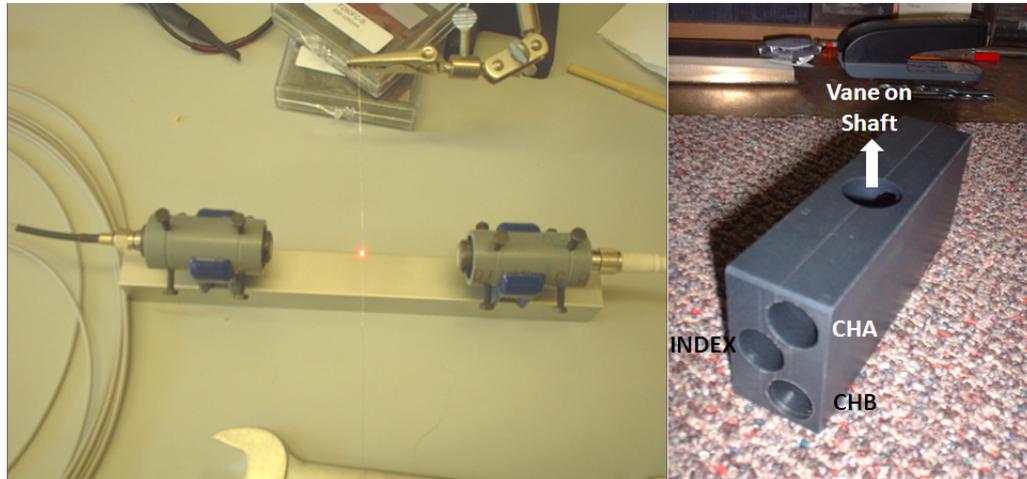


Figure 6.8: Left - test bench for a single channel. Right - prototype optical block before addition of the retaining grub-screws. Pairs of collimators are inserted into the horizontal bores, one either side of the vertical bore.

The amount of freedom that the alignment grub screws allowed made setting up of the three channels very difficult. It could easily take a day to align them. Once aligned the prototype worked well with the shaft running in air, although frequent tweaking of the alignment was necessary as the grub screws that accurately positioned the collimators crept in the plastic housing over time.

Inserting the glass tube into the mount caused several problems. In principle it is possible to focus the laser light onto the surface of the graticule through a curved glass window, but in practice this was very difficult to do. Imperfections in the glass tube meant that it could be very difficult to focus and align all three channels simultaneously. Rotating the glass tube could solve the problem (with a realignment procedure) but the lack of reproducibility was not good for a system that needs to be robust.

Alignment of the system is not possible at RAL due to local laser safety rules and so the alignment has to be performed at Sheffield. As this mount was not robust enough to keep the alignment over time and so it was not deemed suitable for long term use. The amount of light collected from each channel was also dependent upon the position of the glass tube.

Despite these problems the prototype was very useful in proving that the optical system worked but there was clearly several flaws in this particular design that needed addressing in order to create a robust reliable system.



Figure 6.9: The glass tube on top of the stator can be seen in this photograph with a prototype optics block loaded with collimators on the left.

#### 6.4.7 Final Mount

The prototype mount was redesigned by an engineer at RAL, to give a much more solid and reliable mount - see figure ???. The construction was all metal, with the exception of a pair of flat sapphire windows that replaced the curved glass tube of the stator. These windows are bonded into a metal flange and are typical of those used for observation in UHV environments and are rated for such use. The flat windows were much better as they removed the optical distortions produced by the imperfections in a curved glass surface and solved the problem of creating a vacuum tight seal that would be approved for use in ISIS.

Each pair of collimators for each channel is held in an arm that wraps around the mount. Due to constraints on space the collimator focussing had to be changed from 36 mm to 45 mm; this had the effect of increasing the magnification of the fibre source by 50 % but this was not problematic. The 'Index' and 'Channel A' arms are adjustable in the vertical direction. This allows some freedom as to where the index is set and allows the quadrature phase to be obtained. Originally the new mount was designed so that the collimators just slotted into the arms, however a small modification had to be made with the addition of some small grub screws to permit accurate alignment.

This new mount has proven to be easy to set up and extremely robust. Once aligned the mount has shown to remain in alignment indefinitely, even with a reasonable amount of handling. The change from 36 mm to 45 mm focal length had no detrimental effect on the device; this was expected as the FWHM of the laser source at 75  $\mu\text{m}$  is small

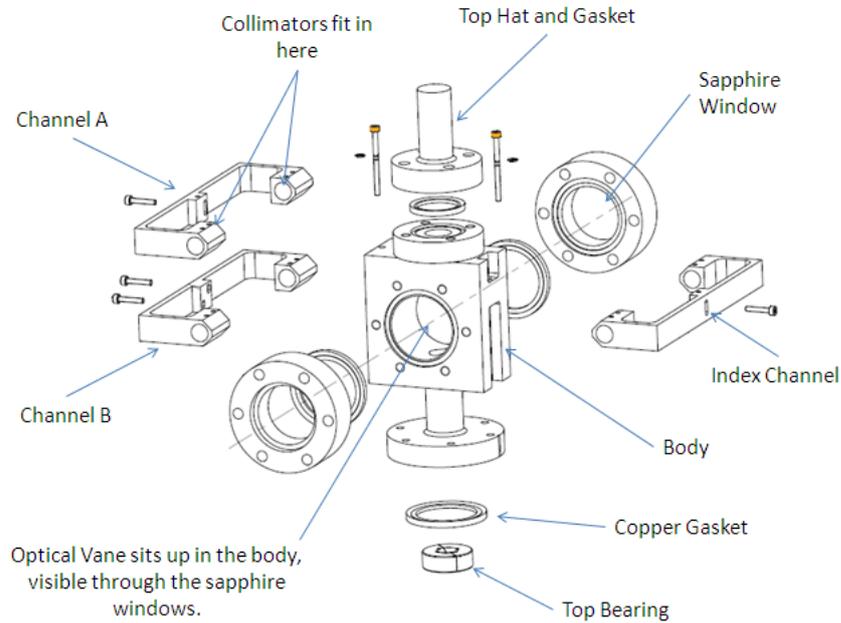


Figure 6.10: An exploded view of the optical mount. The optical vane sits inside the main body where the collimators and lens caps are able to focus the laser light onto the vane.

enough that this change of magnification was not a concern. The mount also solved the problem of sealing the vacuum as the all metal construction of the body meant that it could be sealed to the upper flange on the stator body using a standard copper gasket, ensuring reliable UHV operation.

#### 6.4.8 Optical Sensors

The multimode fibres are returned to a SMA photodiode (H3R880IR). This is a broad spectrum photodiode (400 nm to 1100 nm) [103]. Given that the maximum velocity of the target is no more than  $10 \text{ ms}^{-1}$  then this gives a maximum data rate per channel of 33 kHz, well within the response capability of these devices. The output from the photodiode is amplified and conditioned. More details on the electronic circuitry that accomplishes this is given in chapter 8.

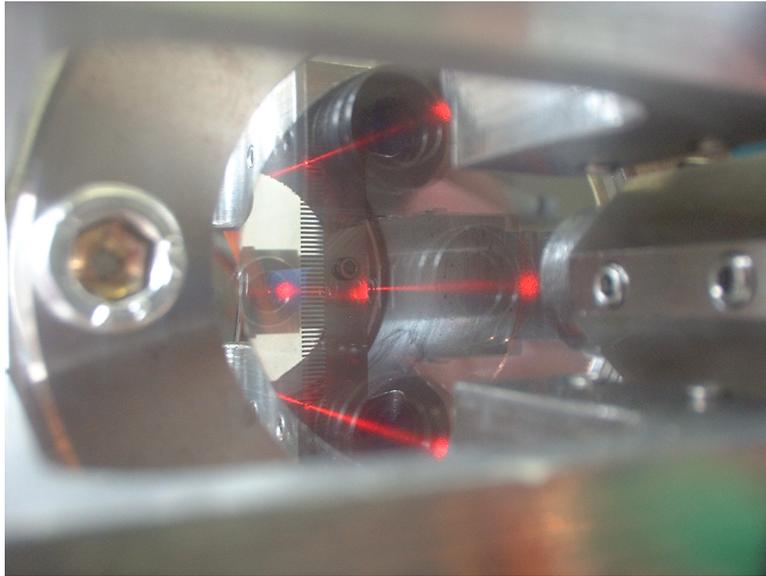


Figure 6.11: A photo showing the laser focused onto the optical vane through one of the sapphire windows. It is clear from this diagram how small the spot size is. (The distance between fingers on the comb is  $300\ \mu\text{m}$ .)

## 6.5 Fibre Installation Issues

The most vulnerable part of the fibre connections is at the ends. The optical fibres are delicate and can soon be damaged through bad handling. The long runs of fibre (over 600 m of fibre installed into RAL for the optics system alone) constitutes a significant expense and so much attention has been paid to minimising the damage to the fibre runs. At RAL each run of fibre to and from the optics block is split into three sections. At either end of each fibre run are two short patch cables, one in the synchrotron and one in the MICE local control room. These patch cables are connected together via a long section of fibre; this latter section is protected along its length by being placed within a metallic duct.

It is expected that the fibres that will experience the largest radiation dose will be those positioned closest to the actuator. By ensuring that the fibres at this point are composed of short replaceable patch leads this will help prevent having to change out the entire length of fibre should this position prove vulnerable to radiation damage. Of course these length can only be changed when ISIS is not operating so, should experience show that the fibres age quickly, this may not prevent down time.

During installation it was specified that the route the fibres took from the electronics to the actuator minimised close proximity to the beam line as much as possible. Unfortunately, the fibres were incorrectly installed, running along the underside of the ISIS beamline itself. This has created an excess of fibre that has been curled up underneath

the beam-pipe close to the target. These two mistakes may provide an acid test as to how well the fibres cope with irradiation.

## 6.6 Conclusions and Future Developments

The final system has performed very well and mechanically it is remarkably robust. Once the optical system has been aligned and calibrated it has worked for periods of several months without requiring any attention.

The system has not yet operated for long enough on ISIS for it to be known whether the same is true for the target operating in situ, the main concern being the possibility of having to regularly compensate for losses in the fibres due to radiation damage by adjusting the amplifier gain or the laser power. At the time of writing this thesis minimal running of the target system at low beam loss had taken place. From the limited operational time experienced so far it is too early to know whether fibre degradation will prove to be a problem.

The small amount of mechanical rotation of the target shaft that is inevitable due to the tolerances of the components moves the graticule into and out of the plane of focus of the laser system. This has the effect of slightly changing the mark/space ratio of the quadrature signals. Although the effect has been noted no problem with counting errors has yet been observed with a correctly set up system. It is also noted that when the mark/space ratio is not 50:50 this does introduce an error into the positioning system. The maximum error on the position is to within half the resolution of the counting system; i.e.  $\pm 75 \mu\text{m}$ .

Despite the success of the system there are a couple of possible mechanical improvements that could be considered for the future. Two possible upgrades are outlined here.

### **Changing the Index Marker to a flag**

By changing the index marker to a flag it is possible to utilise this to give an additional cross check as to whether the target is in its raised or lowered position. This simple modification could be useful to provide an additional interlock system to prevent ISIS operating if the target was ever in an 'In-Beam' position when it was not actuating.

### **Increasing the resolution of the system**

Increasing the resolution of the system may confer several advantages to the target system; better knowledge of the target position could help to improve the control over the target dip depth. The current vane gives 360 individual positions based upon the normal range of allowed target movement (54 mm) on a 0.15 mm resolution graticule. The current graticule is close to the limit of what can be accurately manufactured using wire erosion, so any increase in positional resolution must be matched by a machine that

can cut the graticule to an increased tolerance. There are several possibilities that could be considered including photo-etching and laser cutting. Only with more experience running on ISIS can the decision be made as to whether improving the resolution will confer any advantages. However, any increase in resolution will make the system more sensitive to unwanted mechanical movement due to the tolerances of the mechanical components. This will ultimately limit the maximum achievable resolution.

## Chapter 7

# Stator Operation and the Power Electronics

As the title suggest this chapter is split into two parts. The first part describes how the stator coils interact with the permanent magnets on the target shaft to either levitate or accelerate the shaft. The details of how this interaction is used to control the trajectory of the target is covered in the next chapter on the target electronics.

Later on in this chapter attention is turned to the development of the power supply that supplies the stator coils with current. The development of the power supply is integral to the way in which the stator coils are both wired and operated and for this reason the synergy between the stator and its power supply dictated that both of these items were covered in the same chapter.

### 7.1 Phased Motors

Modern industrial high power rotary motors are often designed to utilise the three phases of the AC power supply to operate. The availability of a three phase power supply lends itself very well to running a three phase motor; the one to one relationship greatly simplifies the design. A secondary effect is that it can also give a more even load distribution on the available power supply. It is possible for a motor to have as many phases as the designer wishes or requires but this can come at the cost of the added complexity of designing a specialist power supply; the application ultimately determines the design.

A three phase motor design does not necessarily have to run on three separate power supply phases nor does it need to run off an AC supply. Clearly the design of the motor dictates what the power supply requirements are, however it is increasingly common to find smaller three phase motors running off a single phase power supply or even off

a DC power supply. This is possible because modern semiconductors can switch large currents quickly and this makes it possible to emulate a three phase power supply by an appropriate switching mechanism, even if only a single phased supply is available. The use of microprocessors to control these semiconductors gives a significant amount of control to the designer. The general theory of utilising multi phase motors is well understood and documented. Continued improvements in microprocessors have allowed significant improvements in the way that this type of motor can be operated.

The target stator was designed as a three phase linear motor and could theoretically be run on either an AC or a DC supply. The decision was made to run the stator on a DC supply for a number of reasons.

- Running on a DC supply reduces the peak voltages that need to be applied to the stator to obtain an average thrust force. Minimising the voltage was important as it reduces the risk of breakdown and allows more of the inter-coil space to be used for heatsinking as opposed to insulation.
- If an AC supply is used then the voltage applied is positionally dependent, not time dependent. This means that there would be positions throughout the actuation cycle where the net force on the target would be zero. Having a varying thrust force as a function of position would make varying the actuation depth in a controlled manner difficult.
- A DC supply is much easier to construct and implement. There is no purpose in making the power-supply more complicated than is necessary. Using an AC supply would mean rating the power supply for both a higher voltage and current.

## 7.2 Three Phase ‘Star’ Wiring

Before considering how the force on the magnetic assembly changes with respect to its position within a powered stator it is necessary to understand how the coils in the stator are wired and how they are switched. There are 24 coils in the stator and these are split into three phases with eight coils in each phase. To understand how the phases are wired a convention is used where the coils are lettered in a cyclic sequence from the top of the stator down, namely A B' C A' B C'. A block of six coils labeled thus is called a ‘bank’. This sequence is repeated four times so there are four ‘banks’ of coils in the target stator - See figure 7.1.

All of the A and A' coils are wired together in series as are all of the B and B' coils and the C and C' coils. The difference between the unprimed and primed coils is that the primed coils are wired in the opposite direction to the unprimed coils so that when a current passes through an A coil in a clockwise direction then the same current passes through an A' in an anticlockwise direction. This has the effect that for a given phase current the magnetic field direction is opposite for primed and unprimed coils. The

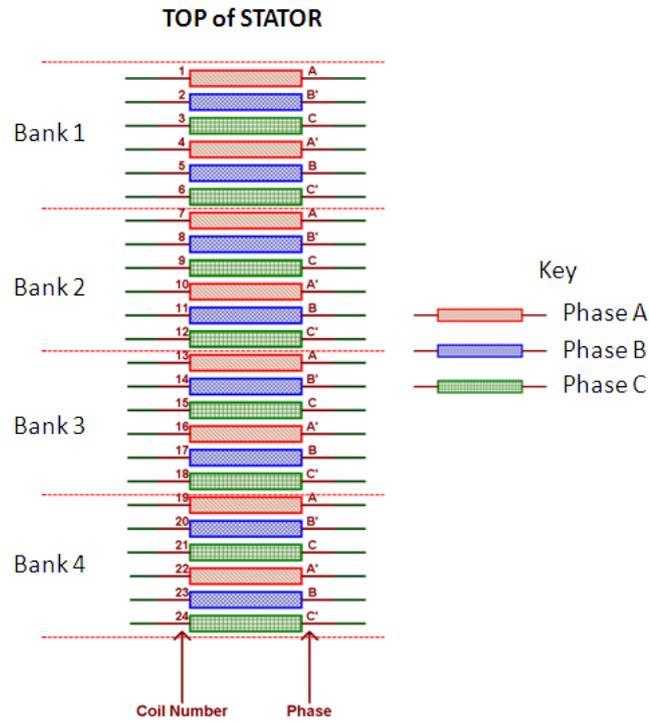


Figure 7.1: Diagram illustrating the phase relationship of the coils to their physical position within the stator. Coils marked with a prime are wired so that the current flows through them in the opposite direction to those coils that are unprimed.

same applies to the B and B' coils and the C and C' coils. This wiring gives the three different phases for the stator stack.

Once all of the coils in each phase have been wired in series this leaves two connections for each phase, one at the top of the stack and one at the bottom. The three connections at the bottom of the stack are wired together and this connection is called the star point. The three connections at the top of the stack serve as the connections to the stator power supply. A diagram representing the stator wiring is shown in figure 7.2. Note that all four of the banks are wired in series as this gives the simplest setup. It is possible to wire the stator as two sets of two banks or to wire each bank individually. Doing this would complicate the switching of the power into the stator but could offer the advantage of reducing the power consumption, allowing the stator to run at a lower temperature or for a given operational temperature this would allow an increase in the actuation rate. It will be noted from this wiring setup that there is no direct return path from the star point. This is deliberate and simplifies the design albeit at the cost of losing the ability to control each phase independently.

If power is fed in through one of the phases then clearly it must return through one

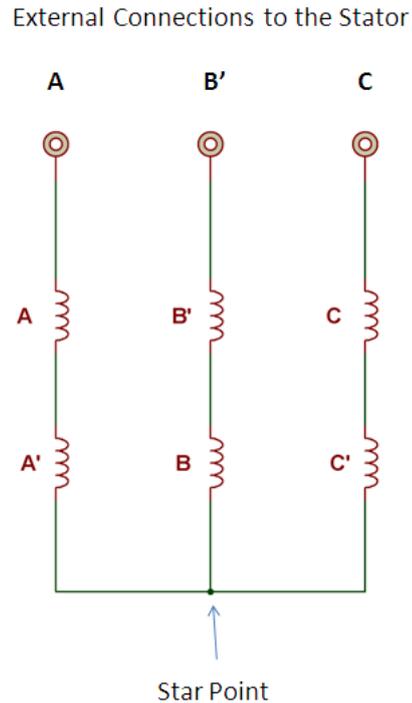


Figure 7.2: This diagram represents how the coils are wired in the stator, illustrating the star point connection and the three external connections from the stator. Each coil in the diagram represents 4 coils i.e. The A coil represents the A coil from banks 1 to 4.

(or both) of the other two phases. Therefore, at a minimum, two phases have current passing through them at any one time. This type of operation where the current flow in one phase mirrors that in one of the other two phases forms part of its operational design and, as will be seen shortly, the geometry of the coil structure is such that this mirrored current adds to the force applied to the magnetic assembly. How this works will be illustrated later on in this chapter.

It is not possible with this setup to independently control the current through each phase, however the design does not preclude wiring the stator up in a manner which allows independent control of each phase. The possibility of doing this and the benefits that it would realise will almost certainly be considered as an upgrade at some time in the future, however such an upgrade would require some further development work first as it would not be possible to realise such operation with the current control electronics.

### 7.3 Coil Current Switching Sequence

The three phase star wiring allows current to be switched through the stator in six different ways if it is assumed that only two out of three phases are ever powered. Using the external connection labels A, B' and C as shown in figure 7.2 to describe the current direction through the stator, these six states are clearly;  $A \rightarrow B'$ ,  $A \rightarrow C$ ,  $B' \rightarrow A$ ,  $B' \rightarrow C$ ,  $C \rightarrow A$  and  $C \rightarrow B'$ .

By stepping through these six different states in a predefined order then the coils can be switched so as to create an apparent 'ripple' motion in the magnetic fields generated within the stator. At this stage this pattern is best visualised by considering the way that the coil switching sequence moves up or down the stator body when passing through this sequence; the visual pattern is repeated for each bank of six coils. The direction of this motion is determined by the direction in which these states are stepped through; reversing this order the apparent motion moves towards the bottom of the stator. The required order of the states to observe this apparent motion upwards through the stator is shown in table 7.1. This is best illustrated graphically and is shown in figure 7.3.

State	Current Flow
1	$A \rightarrow B'$
2	$A \rightarrow C$
3	$B' \rightarrow C$
4	$B' \rightarrow A$
5	$C \rightarrow A$
6	$C \rightarrow B'$

Table 7.1: The six states that create an apparent upward motion in the fields generated within the stator. The states are circular. Moving through the states in the opposite direction reverses the apparent motion - See figure 7.3

The importance of this particular switching sequence will be made clear once the relationship between the fields produced by the coils and how they interact with the magnetic assembly is understood. From the information given in the last few paragraphs it is perhaps already possible for the reader to intuitively understand how by tracking the position of the magnets (by monitoring the position of the target shaft) and by correlating this to the fields produced by passing the stator sequentially through these six states, that target motion can be achieved.

### 7.4 Magnetic Assembly

The magnetic assembly has been discussed at some length in section 5.5.5 but given the importance of its geometry with respect to how it interacts with the stator coils a quick

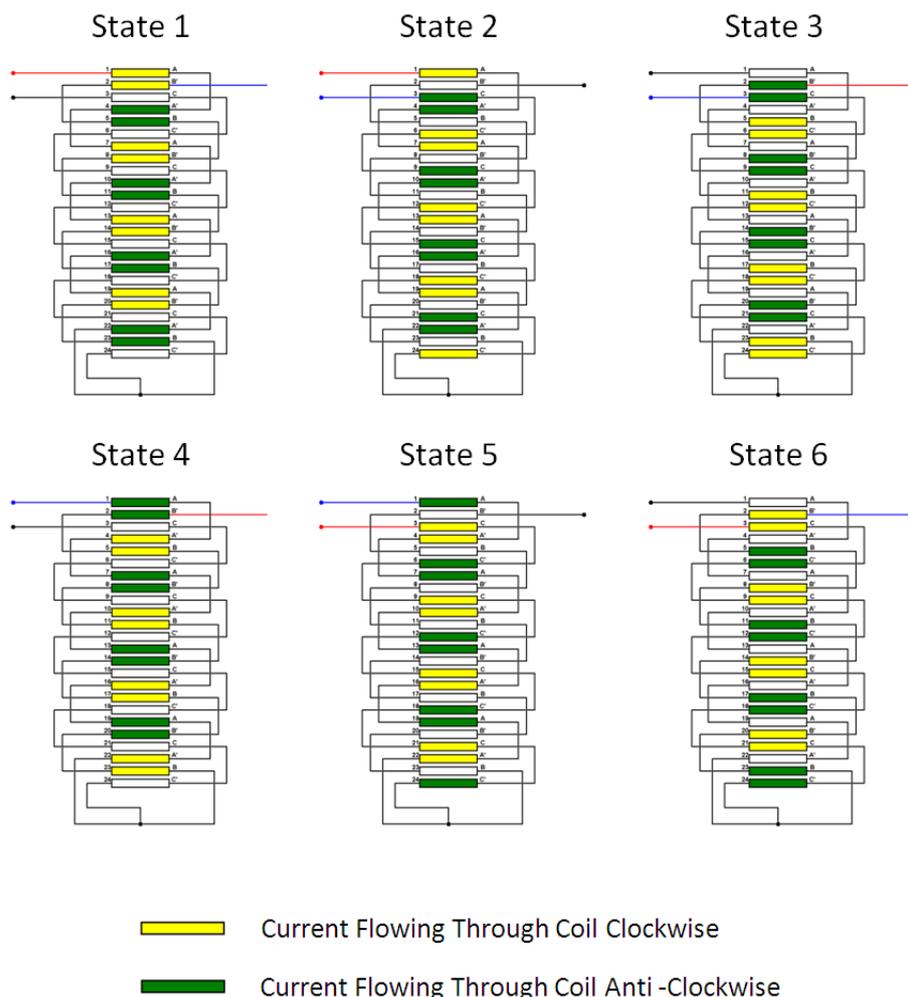


Figure 7.3: By switching the coils through the circular sequence given by States 1 to 6 this gives an apparent motion in the coil switching upwards through the stator as is illustrated here. Reversing the sequence gives an apparent downwards motion. The pattern repeats itself every six coils. The permanent magnets on the target shaft can be made to interact with the fields produced by passing sequentially through these states to produce motion of the target shaft.

recap will be given. The magnetic assembly is composed of three distinct rings where each ring is made up from eight sector pieces. The outer sector pieces have a depth of 3.9 mm whilst the inner sector pieces have a depth of 7.8 mm. Once assembled with the ceramic spacers the assembly has a length of 18 mm and a diameter of 15 mm.

The depth of the magnetic assembly at 18 mm is not coincidental but has been designed to match the depth of six coils. This corresponds to the symmetry of the coil fields

and alignment of the magnets with these fields gives the ability, with suitable feedback, to either hold the magnetic assembly in a predetermined place or to maximise the accelerating force on the permanent magnets.

In order to understand more fully the interaction between the magnetic assembly and the coils a simulation of the hardware setup was produced using Maxwell SV (Student Version) [98]. This software allows magnetostatic simulations to be constructed and solves the result using an iterative finite element process. Once a problem is solved the software is able to show the flux lines, magnetic field lines and the resultant force on a magnetic assembly. Unfortunately the Student Version does not allow parametric analysis but by altering the geometry of the model information can be extracted that allows additional analysis to be done externally.

## 7.5 Simulation of the Magnetic Assembly Moving through the Stator

The purpose of this section is to illustrate how the force varies on the magnetic assembly as it is moved through a powered stator that is kept in a single predetermined state. A model was set up that simulates the stator in ‘State 4’. This is arbitrary as it does not matter which state the stator is in. The model is axis-symmetric in  $R$ ,  $Z$  and the origin is placed at the centre of the stator 0.15 mm below the plane that defines the bottom of the 12th coil. A current of 58 A was passed through the coils. This value was determined by the fact that 58 A corresponds to the peak current that passes through the coils during operation.

The simulation was run over several iterations, with each iteration the magnetic assembly being moved in either 0.5 mm or 1 mm increments vertically through the centre of the stator. This was done over a total distance of 20 mm. The maximum force on the magnetic assembly was calculated for each of these positions and the results of this are plotted in figure 7.4. There are a couple of points to note with respect to this plot. Firstly the maximum force as indicated in the figure is higher than the average force over an actuation, a calculation to show this will be given later. Even when the averaging has been done this graph overestimates the force by  $\approx 10\%$ . It is not uncommon in a basic FE analysis for the force calculation to be inaccurate, as there are many assumptions made in the model. In this particular case the 3D stator has been simplified as a 2D model and so any 3D effects (such as the azimuthal variation of B field over the face of the magnets) cannot be accounted for. A more accurate force calculation would require an improved 3D model. However the form of the solution is correct and many insights can be gained from it.

Due to the symmetry of the powered stator it was not necessary to simulate the magnets moving through the entire stator body. Each state shows symmetry over a distance of six coils therefore symmetry dictates that, if fringe effects are neglected at the ends of the stator, the force curve will repeat itself over the distance of one bank of coils,

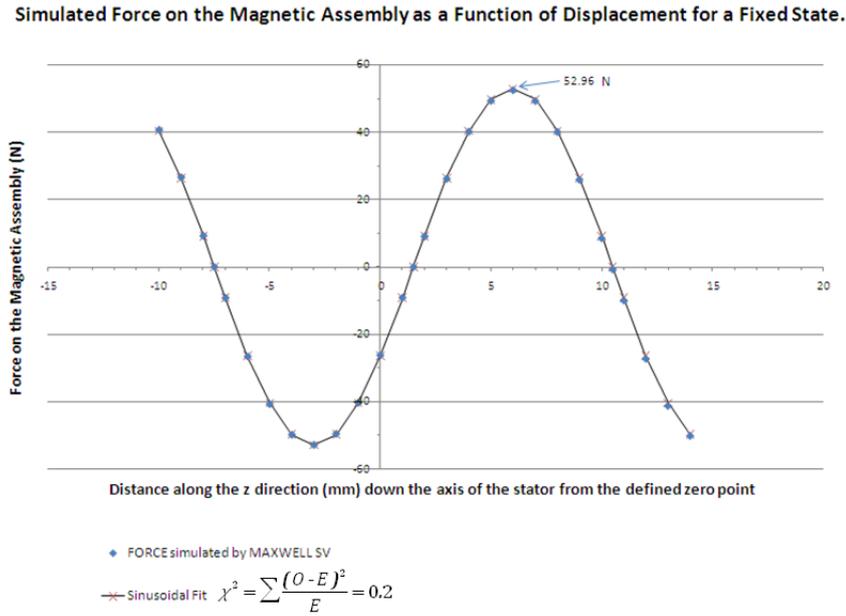


Figure 7.4: Graph illustrating a simulation of how the force on the magnetic assembly varies as a function of its position within the stator when the stator is left in one of its six states with a coil current of 58 A. A good sinusoidal fit can be made to these points as illustrated.

a distance of 18 mm. The graph shows a clear sinusoidal pattern for the resultant force on the magnetic assembly and a sine wave can be fitted cleanly to the simulation results.

## 7.6 Zero-Force Points

Figure 7.4 shows several interesting features that need to be explained if the operation of the stator is to be understood. Firstly there are three points on this graph where the force on the magnetic assembly is zero, although clearly only two points occur through a distance of six coils that define the symmetry.

One of these zero-force points is unstable, any movement of the magnetic assembly away from this point will result in a force that will push the magnetic assembly further away. The other two zero-force points are stable, any movement of the magnets away from these zero-force point results in a restoring force that pushes them back towards the zero-force point position; effectively this point is a magnetic well. The unstable force point can be seen at 1.5 mm whereas the two stable zero-force points are illustrated at -7.5 mm and 10.5 mm (18 mm apart, reflecting the symmetry) respectively on figure 7.4.

These stable zero-force points are of particular interest to the operation of the stator as they give a mechanism by which it is possible to passively levitate the target shaft. Providing that the force of gravity is less than the peak force that can be exerted on the combined weight of the shaft, magnets and vane for a given coil current it can be seen that levitation of the shaft is possible. The shaft will sit at an equilibrium position where the force of gravity is counteracted by the restoring force exerted by the magnetic field. Any small deviation from its equilibrium position would result in the shaft undergoing damped harmonic motion where the damping would primarily come from the frictional forces between the shaft and the bearings.

The stator is designed to be a high acceleration device, accelerating the target in excess of 80  $g$ . Stable levitation should occur when the electromagnetic force on the shaft is equal to 1  $g$ . To a good approximation the force scales linearly to the coil current so stable levitation of the target shaft in close proximity to these zero-force points can be achieved at currents far below those required to achieve the high accelerations needed to put the target shaft into the ISIS beam. Given that in operation the real zero force point (no electromagnetic force) and the apparent zero force point (electromagnetic force in equilibrium with the gravity) are in close proximity to each other these two points will be considered as the same point and called the ‘zero force point’. This assumption will be justified in section 7.9.2.

### 7.6.1 Moving the Zero-Force Points

These zero-force points in figure 7.4 are shown for the stator in ‘State 4’. If the state sequence is progressed then the position of these zero-force points clearly moves with the state sequence. Moving forward through the sequence moves the zero force points up through the stator in 3 mm increments, moving backwards through the sequence moves the force point down the stator in 3 mm decrements. If the zero-force point is utilised to levitate the target shaft then the shaft will passively track the movement of these zero-force points as the state sequence is progressed. When the state changes the shaft will move to the new zero force point, as a restoring force will push the shaft to the new equilibrium position. This is illustrated in figure 7.5. The shaft will undergo damped harmonic motion as it settles to this new point.

The ability to move the shaft in this way is clearly very useful and it is utilised to control the target shaft position when the target is not actuating. By switching on the appropriate state when the target system is powered up, the target shaft is picked up from its resting position, also known as its ‘parked’ position. It is then moved to its holding position, out of the ISIS beamline, by progressing cyclically through these states. This final state, the ‘hold mode’ then holds the shaft indefinitely until it is ready for actuation. Progressing back through these states allows the shaft to be lowered back down to its ‘park-mode’.

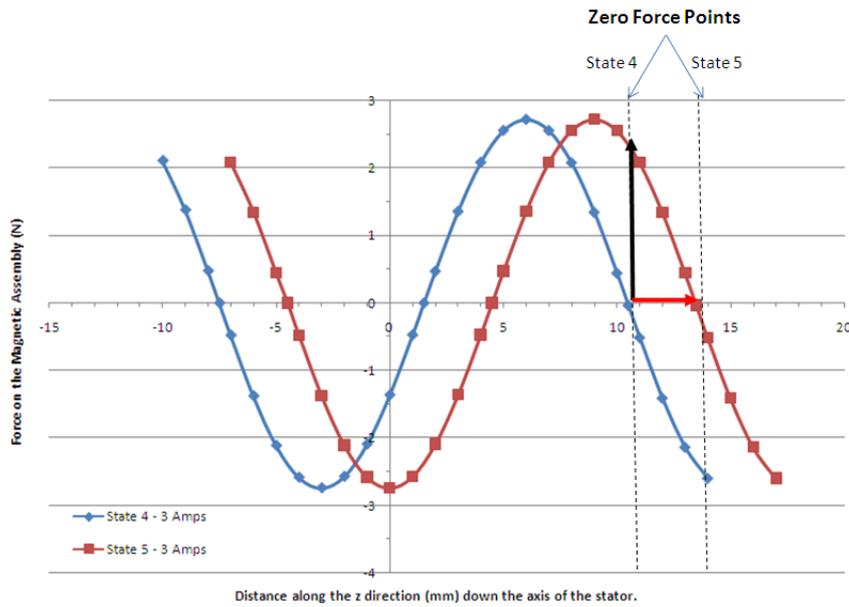


Figure 7.5: If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 5’ then the magnets would see a restoring force (black arrow) that would move the shaft to the new zero force position (red arrow) 3 mm further up the stator. By repeatedly incrementing or decrementing the states the shaft can be moved up and down the stator body. This graph shows the force for a 3 A coil current typical of that used to levitate the shaft.

## 7.7 Position Monitoring

Using the stator in this way to control the target shaft does limit the positioning of the shaft to one of a set of predefined points that are 3 mm apart. This means that the hold position is not entirely arbitrary but the 3 mm step size does give enough freedom to ensure that the target shaft is held out of the beam-pipe in ‘hold-mode’.

This system of moving the target shaft by allowing the magnets to track the position of the zero-force points is entirely passive and does not require any positional feedback system to operate. Of course positional feedback is useful as it provides a secondary check that the target shaft is actually tracking the zero-force point and that the shaft is where it should be. When moving the target shaft by using the zero-force point system position monitoring is provided by the optical position measuring system.

As discussed in chapter 6 the optical position measurement system requires an index signal to ensure that the position read by the system is referenced to a point defined by the position of the index marker. When the target system is initially powered up the absolute position as given by the quadrature counter is in an indeterminate state,

therefore the position determined by the quadrature counter will not be reliable until it has received an index reset. This is accomplished by moving the target shaft past the index point so that the system receives an index reset signal.

Because the zero-point system is passive and does not require positional feedback it provides the mechanism on power-on to 'blindly' move the shaft past the index reset point up to the hold point. This first index reset then ensures that the quadrature readout is reading correctly. This is necessary as it provides a useful check that the quadrature system is functioning correctly. Actuating the target requires the quadrature counter to provide the control system with accurate positional information.

## 7.8 Actuation

The zero force points sit half way between the maximum force points; these are the positions within the stator where the magnets on the target shaft would feel the strongest repulsive force. One of these maximum force points will push the magnets in one direction whilst the other maximum force point will push the magnets in the opposite direction. Consideration of figure 7.4 shows that these maximum force points are always positioned  $\pm 4.5$  mm away from the zero force points; this of course is true for any one of the six states that the stator can be in.

By looking at the force peak it can be seen that the force changes very little  $\pm 1.5$  mm either side of the peak force. In fact, integration of the fitted sine wave shows that the average force over this distance is 95.5 % of the peak force. Now if the shuttle was sitting at a zero force point and the stator coil state was to increment or decrement by two states this would put the shuttle magnets 1.5 mm on the far side of the peak force as the zero force point would have been moved by 6 mm. At this point the shaft would accelerate back towards the shifted zero point. In the first 3 mm of movement the shaft would see an average accelerating force that is 95.5 % of the value of the peak force. See figure 7.6 for clarification of how this would work.

Of course if nothing further was done the shuttle would continue accelerating towards the zero force point and it would then execute harmonic motion about the zero-force point until it came to rest 6 mm further up or down the stator. However if the position of the shaft is tracked using the optical quadrature system then another state change can be made when the target has travelled 3 mm, placing the magnets 1.5 mm on the far side of the next force peak, maintaining the accelerating force on the shaft. This process can of course be continued down the entire length of the stator as illustrated in figure 7.7.

Whilst in this accelerating mode, deceleration can be achieved by suddenly switching the coil state by three positions. This has the effect of placing the maximum opposite force on the permanent magnets. Once again by referring to figure 7.4 and comparing this to figure 7.3 it can be seen that a switch of three states changes the direction of

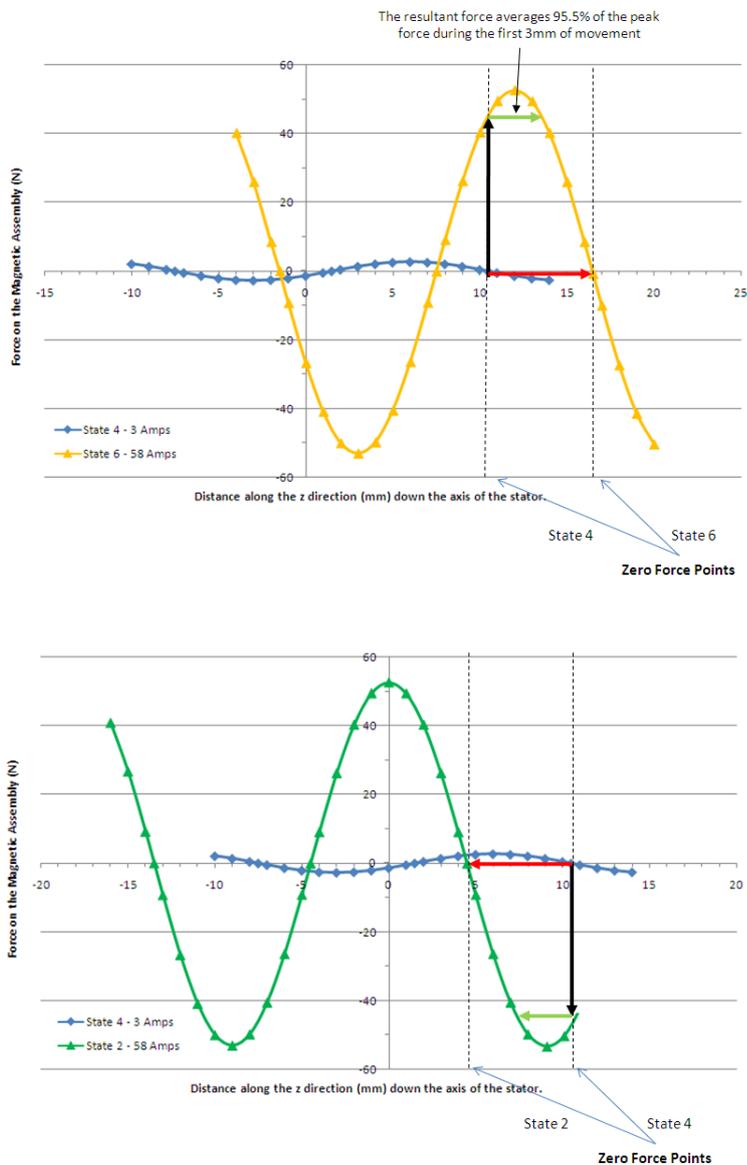


Figure 7.6: If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 6’(top) or ‘State 2’ (bottom) then the magnets would see a restoring force (black arrow) that would want to move the shaft to the new zero force position. During the first 3 mm of movement the shaft would see an average force that is  $\approx 96\%$  of the peak force.

the peak force because it reverses the current flow through the coils.

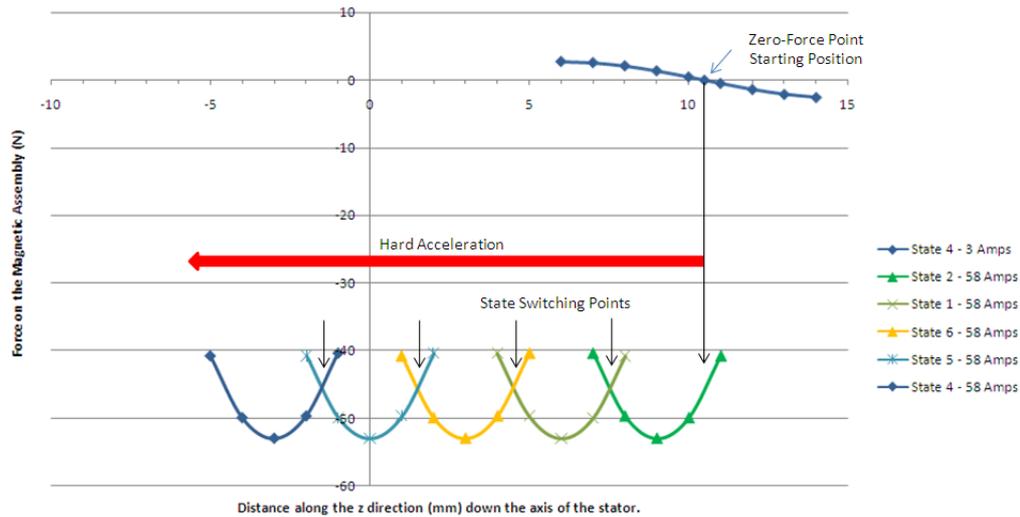


Figure 7.7: The shaft moves from levitation (3 A coil current) to acceleration (58 A coil current). By tracking the position of the magnets the states can be switched to ensure the shaft continually accelerates using the maximum available force. This results in a hard acceleration of the target shaft.

If a previously accelerated shaft is decelerated, and it is decelerated for a long enough period of time, then by simple mechanics there will be a point where the shaft will change its direction of motion. In chapter 6 it was pointed out that by monitoring the direction in which the position reading is changing, the quadrature system is able to determine the direction of motion of the shaft. Ideally the change in direction of the shaft will be observed by the quadrature system when the shaft has near zero velocity, although the shaft's actual velocity when the change of direction is detected will depend on the accelerating force and the resolution of the position monitoring system. This is because the quadrature system detects a change in direction by motion in the opposite direction so the shaft clearly cannot have zero velocity at the point where this change in direction is detected.

This change in direction of the shaft can be used as a trigger to capture the shaft at the nearest zero-force point. This is done by switching the stator into one of the six states that places the closest possible zero force point to the shaft's current position. Clearly good positional knowledge of the shaft is required for this to be done. The coil pitch dictates that the shaft will never be more than 1.5 mm away from a zero-force point and so, providing the shaft does not have a high velocity, capture of the shaft at the zero-force point is inevitable.

It will no doubt be clear to the reader that the processes just described provides the necessary mechanism such that with a suitable controller it would be possible to accel-

erate the target into and out of the ISIS beam whilst enabling its capture again at the end of the cycle. As already alluded to, controlling the shaft's motion in such a manner is called 'actuating'.

The minimum positional resolution required to be able to control the target shaft in the way described is 1.5 mm whereas the optical system provides the shaft's position to within 150  $\mu\text{m}$ . As will be shown in the next chapter the high resolution of the optical system allows better control of the actuation depth of the trajectory of the target shaft than the minimum resolution requirement would allow. The target positional information also gets passed to the target data acquisition system (DAQ). High resolution position monitoring is required by the target DAQ so that the relationship between the position of the target with respect to the ISIS beam and the resulting pion production for MICE can be understood. A full description of how the target electronics controls the trajectory of the shaft using the mechanisms described in this chapter will be discussed in chapter 8 although figure 7.8 shows how in principle it is done. The target DAQ is discussed in more detail in chapter 8 and the results of analysis obtained from the use of this system are discussed in chapter 9.

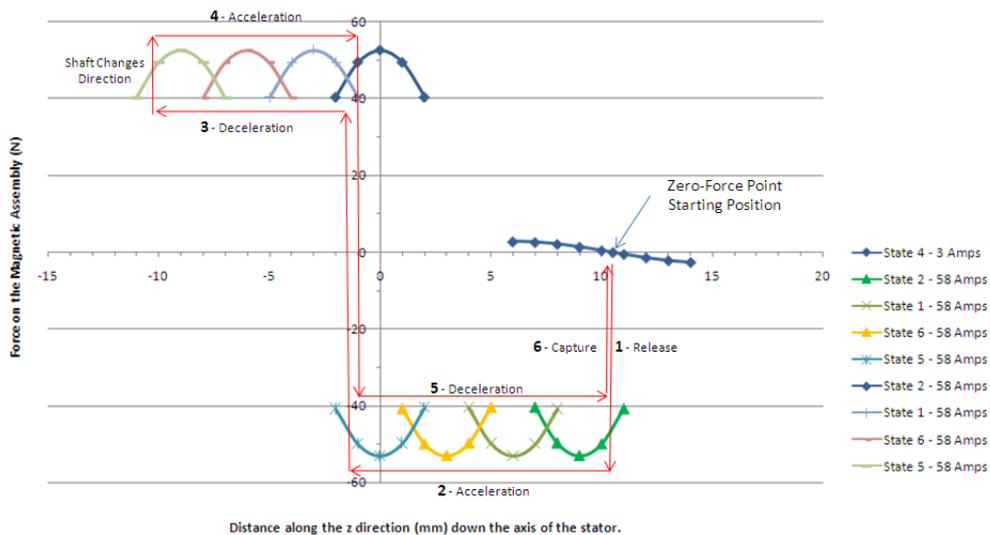


Figure 7.8: An actuation. The shaft is accelerated downwards, decelerated, accelerated upwards, decelerated and then captured. The position of the shaft must be tracked accurately so the controller knows into which state to put the coils to maximise the accelerating/decelerating forces.

If the accelerating force is assumed to be linear with current, and constant, then the trajectory of the target during actuation can, to a first approximation, be described as a set of intersecting quadratic curves. This approximation works well with the first power supply that was built, and this was proven by fitting calculated curves to the trajectory profiles recorded by a local PC. However a correction had to be added with an upgraded high current power supply as the current supplied from this PSU droops

exponentially during actuation due to the use of a capacitor bank to provide the high current. Both of these power supplies will be discussed shortly.

## 7.9 Implementing the Switching of the Coils and Current Control

From the first half of this chapter it has been shown that by wiring up the stator in three phases and applying the appropriate currents to those phases it is possible to control the movement of the target shaft. To be able to do this there are two principal design requirements. The first of these is the ability to switch current bi-directionally through any two of the three phases. The second requirement is the need to have the ability to control the amount of current that passes through the coils. The rest of this chapter concentrates on describing the development of the power supply that provides these three controlled phase currents to the stator.

### 7.9.1 Switching Mechanism

Switching current bi-directionally through the three phases of the stator can be accomplished using six transistors. The six transistors form three pairs where each pair of transistors is connected together serially between the power rails of the power supply. The mid-point of each pair of the three pairs of transistors then connects to one of the three phase wires of the stator. Figure 7.9 illustrates how these transistors are connected. This type of circuit is known as a ‘Hex Bridge’ or a ‘Three Phase Inverter’ and is a variant of the ‘H Bridge’ circuit that utilises four transistors to give bi-directional current control to an electronic or electrical device.

The three transistors across the top of the circuit are called the high side transistors as they are connected to the positive power rail. The other three transistors are therefore known as the low side transistors.

It is possible to see how this circuit can be utilised to switch the current bidirectionally through the phases by comparing table 7.1 with figure 7.9. The inputs are labelled in the figure Q1-Q6. This is often abbreviated to 1-6 or sometimes the inputs are labeled A-F. Note that the odd numbered inputs correspond to the high side transistors, whilst the even numbered inputs correspond to the low side transistors.

When using transistors to switch currents of more than a few amps it is usually necessary to use either Integrated Gate Bipolar Transistors (IGBTs) or Metal Oxide Semiconducting Field Effect Transistors (MOSFET’s) as the switching elements. These transistors will switch quickly and will saturate very close to the power supply rail voltage as they have a small internal resistance ( $R_{ON}$ ). Both of these characteristics are good for minimising the power dissipation within the transistors themselves and this makes them very efficient at transferring the power to the load. By utilising these

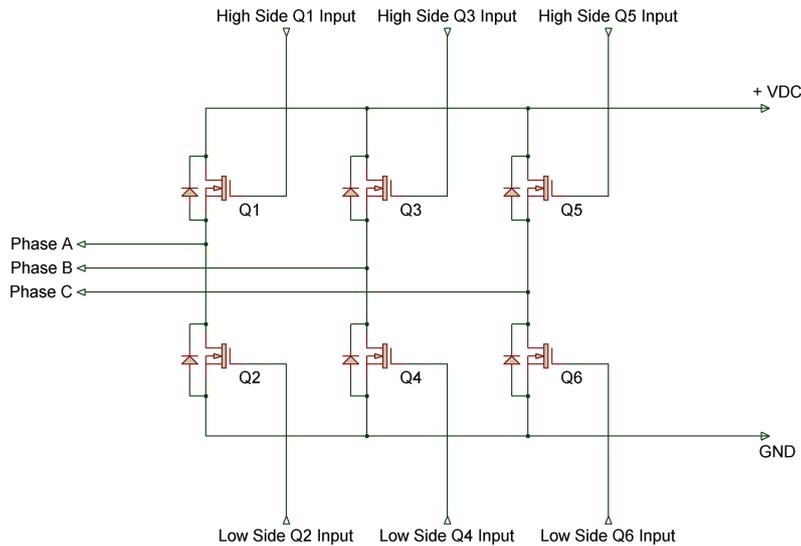


Figure 7.9: The Hex Bridge. This circuit allows bidirectional current control through any two out of the three phases of the stator by application of the appropriate control signals to the six inputs that switch the transistors on.

transistors a circuit that can switch a significant amount of power can be built with a minimum footprint. The low quiescent power consumption of these transistors improves the reliability of the switches and extends their lifetime.

Unfortunately both MOSFET's and IGBT's are very intolerant of over voltage. Fast switching of large currents generates voltage transients that can exceed the DC supply voltage, and possibly over-rate the device. Transients can also couple capacitively to the switching terminal, the gate, which is the most susceptible pin to over-voltage. This means that good care must be taken when designing the circuit to ensure that the devices are over-rated and that transients are suppressed, otherwise the devices will quickly fail.

One other risk with this circuit is that if the two transistors that control the current flow into and out of each phase are turned on simultaneously e.g. if inputs Q1 and Q2 are activated simultaneously then this creates a short circuit between the power supply rails. This almost invariably permanently damages the transistors due to excessive current flow. It is also possible to damage the power-supply itself if it is not current limited. This type of short circuit due to transistor switching errors is known as a 'shoot-through'. Great care has to be taken to avoid passing the wrong signals to the bridge circuit to prevent this situation from occurring. Protection circuitry can be added to prevent this from happening but this is difficult to do without compromising the performance of the switching.

State	Current Flow	Hex Bridge Inputs
1	$A \rightarrow B'$	1 and 4
2	$A \rightarrow C$	1 and 6
3	$B' \rightarrow C$	3 and 6
4	$B' \rightarrow A$	3 and 2
5	$C \rightarrow A$	5 and 2
6	$C \rightarrow B'$	5 and 4

Table 7.2: This table shows how the six stator states can be switched using the Hex Bridge.

The risk of shoot-through also comes from another subtlety with respect to the time that it takes to switch transistors on and off. When a given transistor is sent a signal to turn it off it takes a finite amount of time for the transistor's source/drain current to reach zero. Because of this it is important that when one transistor is turned off on a given phase that a certain amount of dead time is allowed before the other transistor can be turned on. If this extra time is not allowed then a 'shoot-through' can occur. The amount of time this takes is dependent upon the transistor but usually varies from tens of nanoseconds for small power transistors to a few tens of microseconds for the larger power transistors.

It is worth noting that 'shoot-through' is not normally a problem when stepping through the state sequence in a circular fashion as there is always at least one state on a given phase between a low-side transistor turning off and the complementary high-side transistor being turned on (or vice-versa). Given that the frequency of the state changes tends to be no more than a few kHz this ensures that the minimum dead-time is observed. However this is no longer true when the force direction is suddenly reversed. Here this creates a situation where the reversal of the current means that the two transistors on each powered phase change state. When this occurs a suitable dead time must be introduced by the controller to prevent 'shoot-through'.

### 7.9.2 Current Control

When the shaft is being levitated in its holding position the gravitational force will pull the shaft slightly away from the zero-force point such that there will be a small restoring force acting upon the shaft to exactly counter the gravitational force. The amount of counter force that the stator needs to supply to the shaft is small, reflecting the low mass of the target shaft.

Ideally this equilibrium position needs to be as close to the zero force point as possible. The true zero force point represents a reference position when setting up the stator and as previously stated the equilibrium position and the true zero-force point are assumed to be equal. Now it can be seen from the plots that  $\frac{dF}{dz}$  is maximal at the zero-force point so the amount of displacement required along the  $z$  axis to produce the

restoring force is minimal and approximately linear to the applied current over short displacements.

The amount of current used by the coils needs to be enough so that the levitation is stable but also power dissipation is minimised within the stator. A holding current of around 3 A has been found to satisfy both of these criteria. This can be seen by looking at figure 7.6; the shaft, magnets and vane has a mass of  $\approx 57$  g so the gravitational force is about 0.56 N. This force is counteracted by the electromagnetic force at a displacement of  $\approx 500$   $\mu\text{m}$  from the true zero force point.

On actuation the shaft will be accelerated at  $\approx 80$  g to intercept the ISIS beam at the correct time and to be out of the beam-pipe before the next ISIS beam injection. This implies that a large electrical current will need to pass through the coils. It has been found that a coil current of  $\approx 60$  A is needed to obtain the required acceleration. This current, if sustained, would quickly overheat and damage the coils and so care must be taken to ensure that the actuation current is only supplied to the stator for the required amount of time.

The ability to pass a large current through the coils for short periods of time during actuation and to pass a smaller current through the coils for longer periods of time when the shaft is in its holding position means that some mechanism is required to control the current to the stator coils.

### 7.9.3 Controlling the Current with Pulse Width Modulation

As alluded to earlier on the chapter Pulse Width Modulation (PWM) is a standard technique used to control current into motors, or any other device that requires a significant amount of power. PWM uses a square wave signal where the duty cycle is modulated to alter the average time that the signal is turned on for. This average time is usually representative of some other analogue signal. By altering the PWM duty cycle the amplitude of the continuous signal is represented by the duty cycle of the PWM signal.

By pulse width modulating the current through a given electrical component then the average current that flows through the device for a given duty cycle is, to a first approximation, linearly proportional to the maximum current that would flow through the device if the duty cycle is set to 100 %. Typically a PWM frequency of several kHz is used for controlling motors, the frequency being determined by the design of both the motor and power supply, although the exact frequency used is not usually critical. An example helps to elucidate the principle. If an induction-less motor pulls a current of 60 A then superimposing a switching signal with a duty cycle of 25 % would give an average current through the motor of 15 A. The momentum of the moving parts in the motor make the motor ‘blind’ to the high frequency switching effect and the motor runs as if the current was actually a steady DC 15 A.

Of course no electrical component is induction-less, this particularly being the case with motors. The induction of the motor reduces the sharp rise/falls in the current value, the effect being dependent upon the inductance of the motor and the frequency of the PWM signal being used. The effect of the rise time in current is offset by the fall time in current. For motors where the frequency of the PWM signal is much higher than the inductive time constant this leads to a continuous DC current passing through the motor with a superimposed AC ripple.

The induction of the motor has important consequences for the driving circuit. If the design of the circuit took no consideration of the inductive currents generated by the stator then the sharp turn off of the transistors would lead to a large back emf that would inevitably destroy the transistor that turned the current off to the motor. Most modern IGBTs and MOSFETs that are designed to switch power have an internal diode that connects back across the device to protect it against inductive surges. This diode will always be rated to at least the same maximum current as the transistor.

The path that the inductive current takes can also have an effect upon circuit operation. If both the high-side and the low-side transistors are fed with PWM signal then the inductive current will take a path through the power supply. If the power supply is not designed to withstand this current loading then it can be damaged. For capacitor banks the inductive current can under certain circumstances reverse polarise the capacitors causing them to fail prematurely. Taking the current path through the power supply can also significantly extend the length of the current path, particularly if the power supply is not situated close to the electronics. This increases the amount of electromagnetic noise produced by the circuit.

This problem can be partially mitigated by placing a reverse biased power diode across the power supply rails. However there is a neater solution. If only the low-side transistors are pulsed then the induction current path remains localised. Because the high side transistors stay on (for a given state) this provides a shorter return path for the induction current. This reduces the EMI and reduces the potential for damage to the power supply. This is illustrated in figure 7.10. For this reason it is usual to superimpose the PWM signal onto the low side transistors only and this was the way that both of the driver circuits for the MICE target were operated.

The assumption of a linear relationship between the duty-cycle and the current as seen through the motor is generally a very reasonable first order approximation. However the finite switching times of the transistors can introduce significant non-linearities as the duty cycle approaches either 0 % or 100 %. This effect is more noticeable for power transistors due to their longer switching times. If the duty cycle of the PWM signal is expressed as a percentage then this effect can become problematic if the transistor switching time  $T_{On}$  becomes greater than:

$$T_{On} \ll \frac{\text{Duty Cycle (\%)}}{\text{Frequency}} \quad (7.1)$$

The PWM is used to reduce the current required for levitating the target when moving to and from ‘park’ mode and when holding the shaft ready for actuation in ‘hold’ mode. The PWM is set at 100 % duty cycle when actuating, which is equivalent to a DC current. In the latter case the DC voltage on the power supply will set the actuation current, the current being primarily limited by the DC resistance of the stator coils, cables and the transistor’s ‘on’ resistance, in that order. The required holding current, at 3 A, is then scaled from the maximum delivered current during actuation to give the PWM duty cycle when the actuator is not actuating.

An early prototype Hex Bridge circuit that could deliver 10 A to the stator coils worked at a PWM frequency of 20 kHz. The 10 A design was later upgraded to a design that delivered 60 A to the stator at Daresbury laboratories. This PSU was initially restricted to 5 kHz operation, due to problems with the higher powered transistors’ longer  $T_{On}$  time. This problem was later resolved but operation of the stator now still continues at 5 kHz. A more detailed discussion of these power supplies can be found later in this chapter.

Switching large currents at high frequency can invariably lead to a significant amount of noise on the power supply which can lead to electromagnetic interference or EMI. Given that modern controller circuitry is complex and requires a clean power supply of only a few volts, voltage transients from the motor and its power supply can easily cause significant problems with the controller if the two circuits share a common power supply. For this reason the control circuitry almost always has to be placed on a completely separate power supply from that of the motor so that electrical isolation between the circuits is achieved. The switching signals that need to be passed from the controller circuit to the motor power supply are then usually passed through some kind of optical device so that complete isolation of the two circuits is accomplished.

Snubber circuitry is also required on the power supply to reduce the amplitude of and dampen the high frequency transients created by the pulsed switching of the transistors. This also helps to stabilise the current flow through the motor and serves to reduce the EMI emitted by the motor, its power supply cable and its driving circuitry.

## 7.10 System Placement in ISIS

It was clear early on in the project that the installed layout of the actuator system at ISIS would require there to be considerable distances between the various components in the system. The power supply cannot sit close to the target mechanism as the radiation produced by ISIS would likely cause premature failure of the electronics. Additionally, it also makes good sense to have the electronics accessible so that should any maintenance be required the power supply can be accessed without having to wait for the synchrotron to be opened up. The high currents required by the target mechanism mean that the distance to the power supply should be minimised to reduce the ohmic losses in the power cable connecting the two together. The stator coil resistances are fairly low so

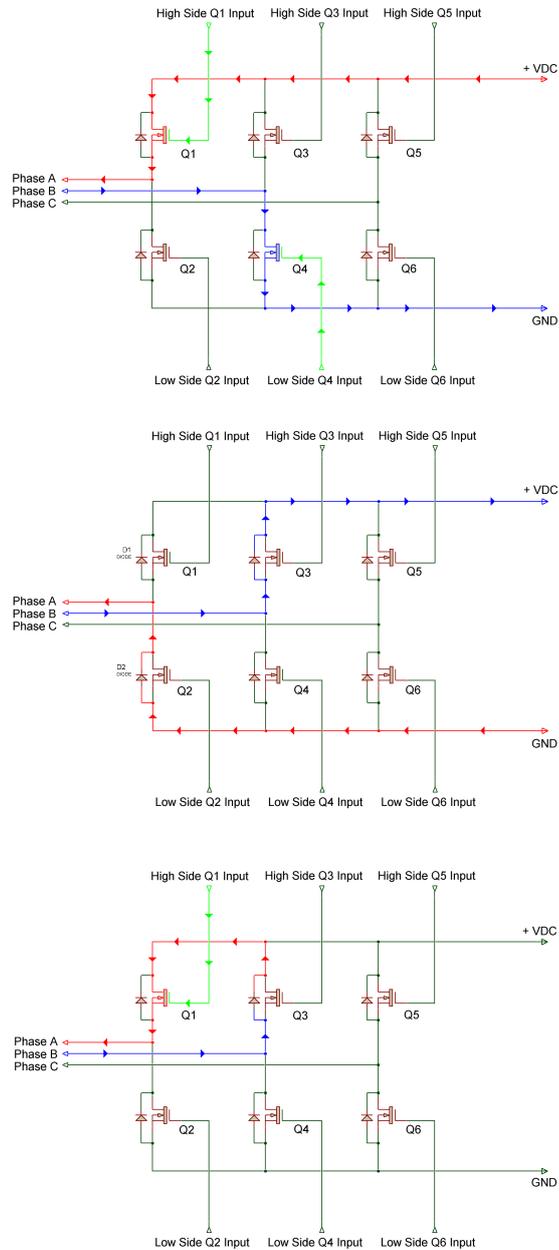


Figure 7.10: Top: Current path for 'State 1'. If the PWM signal is imposed onto both the high-side and low-side transistors this causes the induction current to pass through the power supply (middle). This can stress the power supply and creates a lot of EMI. If the PWM signal is only supplied to the low-side transistor then this improves the situation as the induction current is now localised within the driving circuit (bottom).

losses in the power cables that run from the target power supply to the target itself can soon mount.

For this reason the power supply to the target has been installed on the outside wall of the synchrotron at a location as close to the target area as possible. This has still necessitated a minimum power cable length of 25 m between the target mechanism and its power supply. The control electronics for the target system have been situated in the MICE control room and this is also a significant distance from both the target and the power electronics. Its installed position has meant that there is a 70 m cable run between the control electronics and the power electronics and nearly a 100 m run between the control electronics and the stator.

## 7.11 Fibre Optic Links

The fibre optic cables that are run between the control electronics and the stator's optical block have already been discussed in chapter 6. The six signals that control the switching of the transistors in the power supply have been sent from the controller to the power supply via fibre optic links. The use of optical fibres guarantees signal integrity, completely isolates the power supply from the controller and eliminates the risk of noise on the signal lines which could have caused additional problems with either the controller or the transistor drivers in the power supply.

This fibre optic link between the controller and the power supply has been designed by using a commercially available optical transmitter/receiver pair. The transmitter is a high powered Infra-Red LED and it has a bandwidth up to 5 MHz. This was deemed more than sufficient as a 1 % duty cycle resolution on a 20 kHz PWM signal requires a bandwidth of only 2 MHz. The receiver is a stand alone unit that gives a TTL compatible output for use in the hex bridge driving circuitry.

The optical fibres used were the same multi-mode fibres used to return the laser light in the quadrature optical counter, (BFH37-200). These fibres were well matched optically to the transmitters and experiments in the lab showed lower losses using these cables than by using the standard polymer cables that came with the transmitter/receiver pairs. The simplicity and reliability of the optical link has meant that this optical link circuitry has remained unchanged since its initial design in the laboratory.

## 7.12 10 Ampere Hex Bridge Circuit

The first Hex-Bridge circuit built to supply power to the stator was designed to deliver 10 A of current. At the time of its design there were many uncertainties in the design of the target system itself; the first prototype actuator had only recently been built, and a power supply was required to allow testing, hardware development and firmware

development. At this point the main issues were:

1. The final mass of the shaft and magnets was unknown as these components needed a lot of development.
2. The optical quadrature system had yet to be developed.
3. The required dip depth hence trajectory of the shaft was unknown and no information was available regarding stator position on the beamline.
4. The efficiency of the stator in terms of the correlation between the current and acceleration was unknown.
5. The details of the algorithms necessary to reliably control the shaft's trajectory had yet to be coded; the mechanism for levitating the shaft had yet to be discovered.

With all these unknowns the decision to build a 10 A supply was based primarily upon the fact that there was a power supply available that could deliver 10 A into the stator at the required voltage.

The Hex-Bridge was built using relatively small MOSFET's although they were suitably rated to handle 10 A. Isolation from the driving circuit was initially achieved using optical isolators although later this was upgraded to the fibre optic system as described earlier. MOSFET driver IC's were used to achieve minimum switching times on the MOSFET's although this was damped somewhat by an RC circuit that was placed to prevent coupled transients from the source/drain of the MOSFET's upon switching exceeding the maximum gate voltage.

Snubber circuitry was added across each MOSFET to reduce transient ringing on switching, an effect due to the combination of the LCR of the circuit [104]. Fast switching Schottky diodes protected all inputs of the driver circuit to protect the circuitry from transients fed back into the circuit through the gate.

### 7.12.1 Performance of the 10 Ampere circuit

After several revisions of the circuit, reliable operation was achieved and this circuit was used for over a year before a more powerful power supply became available. During this time the quadrature optical system was designed and built and integrated into the setup, the control algorithms were developed, and much of the mechanical system as described in chapter 5 was implemented. This power supply was used with a stator installed on ISIS for some beam tests in November 2006 to obtain information on the position of the ISIS beam and particle production for MICE. ISIS had to run in a mode that allowed single pulses on demand as the acceleration of the shaft at 10 A was not quick enough to allow operation at 50 Hz. More details of the results from these beam-tests are discussed in chapter 9.

### 7.12.2 Correlation between Current and Acceleration

Although the operation of the stator at 10 A only gives a single data point to extrapolate from it did give some indication of the coil current that would be required to operate the stator at ISIS. Operation of the stator showed that on average an acceleration of  $\approx 14 \text{ ms}^{-2}$  was achieved with 10 A, although this figure did change depending on the mechanical setup.

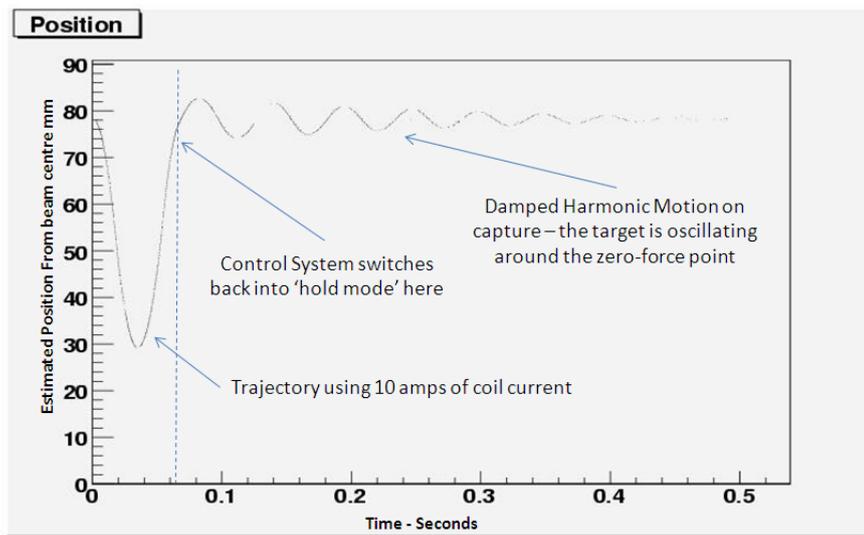


Figure 7.11: A target actuation at 10 A using the 10 A Hex-Bridge. At this time capture was done based upon the target position. This proved to be unreliable, and was later changed so that capture occurred at the point where the target trajectory changed direction at the top of the pulse. Some damped motion of the target is expected at capture as mode change does not necessarily happen at a zero force point and the target may have some residual velocity.

It also became clear from the measurements obtained during the November 2006 beam tests at ISIS that a target trajectory could be chosen that did not necessarily require a full 100 *g* of acceleration but could be possible with about  $\approx 80 \text{ g}$  yet would still meet the requirement of getting into the last 2 ms of one ISIS spill without interfering with the injection of the next spill 10 ms later. The advantage of the smaller acceleration was that the power supply would only need to provide  $\approx 60 \text{ A}$  assuming a linear extrapolation. As the power dissipated in the stator goes as  $I^2$  any reduction in the current used provides a significant advantage.

## 7.13 100 Ampere Power Supply

The expectation that the stator would require at least 60 A of current to realise the required acceleration meant that a new power supply needed to be built. The requirement for 60 A was an estimate based upon the performance at 10 A, a single data point. As a result a significant overhead was to be built in; the new supply was to be able to deliver up to 100 A to the stator. As the design of a new power supply was considered a major project the specification, design and build for this power supply was handed over to a group of electrical engineers at Daresbury Laboratory based upon input provided from the 10 A circuit built at Sheffield [105]. This way the design was kept compatible with the existing control electronics; the fibre optic receivers and the switching sequence were also maintained.

The design of a power supply that could theoretically deliver up to 100 A to the stator poses several unique problems. Having a linear power supply that could provide 100 A upon demand is possible but it would be both expensive and bulky. The stator only requires the high currents for relatively short periods of time and the actual average current drawn is significantly lower than the peak current.

In situations like this it is more economical and efficient to power the driving circuit from a capacitor bank so that the capacitors can provide the short high current pulses on demand when necessary. The capacitor bank is then charged by a second power supply that is specifically designed to supply the necessary charging current to top the capacitor bank up and provide the stator's holding current. This supply is called a capacitor charging unit or CCU. The CCU effectively provides the average current that the stator uses whilst the capacitor bank is there to provide the peak currents when required. The size of the capacitor bank attached to the Hex Bridge is 70 mF and is rated to 400 VDC

The use of a capacitor bank is inherently safer than using a linear power supply because should a fault occur that leaves the system in a state that demands high current then once the capacitor bank has discharged the only current that can be supplied to the stator is that provided by the CCU. The CCU used can only give a current of 4 A and into a 3.6  $\Omega$  load (The DC resistance across two phases of the stator at 20°C) this gives a power dissipation of only  $\approx 60$  W. This is significantly lower than the power dissipation during normal stator operation when it is actuating at 1 Hz and so this safely limits the energy deposition into the stator. At 70 mF the capacitor bank stores a significant amount of energy. This energy could be deposited into either the stator or the bridge circuit under a fault condition. Experience has shown that that this safety mechanism is very effective at preventing the stator from overheating but fault conditions that cause the capacitor bank to discharge into the stator can cause significant damage - see chapter 10.

### 7.13.1 Description of the 100 Ampere Supply

The power supply is very similar in its broad design to the 10 A power supply. The MOSFET's were changed to IGBT's as it was considered that these gave better performance at the required currents. The IGBT's that were chosen are power devices and can switch up to 250 A. The gate switching circuits were improved as the IC used to charge the gate in the 10 A circuit would not be able to provide the necessary gate current to switch the IGBT's at their maximum speed. The Bridge Circuit was connected to a capacitor bank that originally had 15 mF of capacitance but was later upgraded to 70 mF.

All six sections of the gate drivers (one for each IGBT) were taken to a floating DC power supply as this gives better individual protection to the circuits. (The 10 A circuit only had its 'High' MOSFET drivers on a floating DC supply as it is not absolutely necessary to have the 'Low' MOSFETs floating.)

The logic circuitry within that drives the gates on the IGBT was supplied with power from a separate power supply that is integral to the unit. As pointed out the capacitor bank is charged through a power supply that is specifically designed to charge capacitor banks. In this case a Xantrex 300-4 was used for this purpose. This unit can supply up to 4 A at up to 300 V i.e. a maximum output of 1200 W to both charge the capacitor bank and keep the stator levitated in its holding position.

### 7.13.2 Performance of the 100 Amp Supply

The 100 A supply was tested at Daresbury Laboratory providing up to 70 A into a dummy load that was impedance and inductance matched to the coil arrangement within the stator. It has subsequently been used to provide up to 60 A into the target stator.

The stator was initially run at 10 A with the Daresbury supply. At this point it was clear that the first capacitor bank at 15 mF was not large enough to provide the necessary current throughout the actuation cycle. As the capacitor bank discharges the voltage on the capacitor bank droops thereby limiting the maximum available current and extending the actuation time. With a 15 mF the capacitor bank was seeing a droop of about 50 % towards the end of the actuation cycle. This was significantly slowing the rate of deceleration and created problems with the target trajectory and capture. An upgrade to the capacitor bank increasing its capacitance to 70 mF was made at this time and although some droop is still observed it is no longer a problem; although changes to the control firmware had to be made to compensate for the droop.

Once reliable operation at 10 A current had been established the voltage was increased on the capacitor bank so that the actuation current increased in 10 A increments. Each voltage increase was followed by a several days continuous running at the new settings until reliable operation had been demonstrated. Often this involved making changes to

the control firmware. This step wise increment continued up to an actuation current of  $\approx 60$  A, corresponding to a working DC voltage on the stator of 210 V.

Due to the slower switching speed of the IGBT's the PWM frequency when the stator is in hold mode is determined by the maximum current flowing into the stator. At high maximum currents the duty cycle on the PWM for the hold current can become so short that it does not allow time for the IGBT to turn on as the inequality in equation 7.1 is no longer valid. This is frequency dependent, the higher the PWM frequency then for a fixed duty cycle the shorter is the time that the transistors are turned on for. When this occurs the only solution is to reduce the PWM frequency so the duty cycle time always exceeds the minimum switching time  $T_{On}$  of the IGBT's. For the first power supply the minimum  $T_{On}$  was 10  $\mu$ s which meant that the PWM frequency had to run at 5 kHz. As previously mentioned this problem was later resolved and the PSU can now operate at up to 20 kHz. As no problems were found with running at 5 kHz the firmware continues to use this PWM frequency.

The Daresbury power supply has now been operated for many millions of actuations and has proven to be reliable in operation both in the laboratory and on the installed system at ISIS. Since its original design several features have been added to the power supply that allows it to be controlled remotely from a PC terminal in the MICE local control room.

Figure 7.12 shows an overlay of several target trajectories at increasing currents. This was done during the development stage of the power supply to provide evidence that the relationship between current and force is linear (as was expected) and that a current of 60 A would provide the necessary accelerating force on the cross-shaped target. These plots represent single actuations and no attempt to capture the target at the end of its trajectory was made. (The power supply was only on loan for a single afternoon from Daresbury.) The control algorithms for catching the target at the higher accelerations had not been developed at this time that this plot was produced.

An analysis program was then used to calculate the acceleration for each of the trajectories in figure 7.12. Unfortunately there was no estimate for the errors on the points on this figure but it does show that the relation between the current and driving force appeared to be linear, as expected. It should be noted that the data for figure 7.12 was obtained using a target shaft that had a square-cross section; this was slightly heavier than the cross section shaft that was obtained shortly after creating this plot.

Using the linear correlation between force and current with a single data point for a lighter target shaft on a 10 A trajectory correctly suggested that the required acceleration would be obtained at a coil current of between 50 A and 60 A - see figure 7.13. An accurate answer was not really required or even expected from such a limited data set. Also there were still uncertainties with respect to the target mass as the design had not been finalised. However this was still useful as it did provide good indication that a coil current of 60 A was a reasonable estimate of what would be needed to drive the stator at the required accelerations.

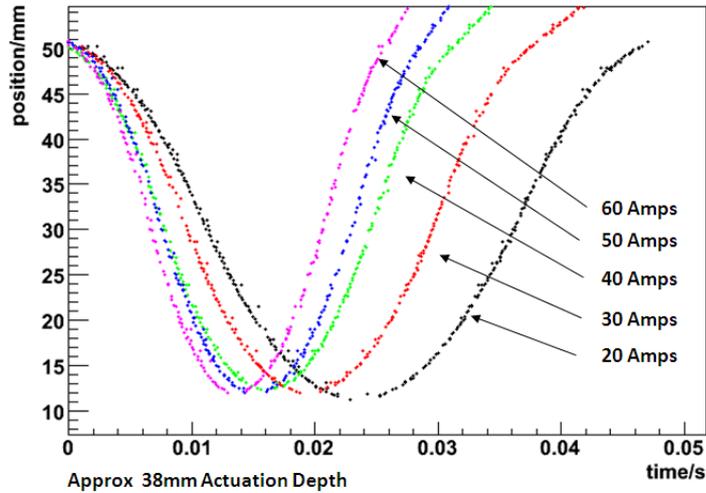


Figure 7.12: Plots of some of the first actuations at higher current. This plot represents the complete data set. At this time no attempt was made to capture the target. This was done to obtain some indication as to whether the required acceleration would be achieved with the 100 A power supply.

## 7.14 Over Temperature Protection

As discussed both the 10 A and 100 A power supplies are capable of depositing a significant amount of energy into the coils. In fact due to the use of the capacitor bank the 10 A circuit posed more of a risk than the 100 A circuit as it could deliver more current in a fault condition. Of bigger concern was if the stator water cooling circuit ever became blocked as this would almost invariably lead to a situation where the stator would overheat and possibly become a fire risk. As the stator was run for longer periods of time to build up reliability statistics it became clear that it would need to be left running unattended. This necessitated building a protection system that would cut the stator's power supply off in the case where the stator exceeded a preset temperature.

As was discussed in chapter 5 the stator has several thermocouples embedded into the coil stack so that the temperature of the coils can be monitored. One of these thermocouples is used for a over-temperature protection system by connecting it into a temperature monitoring unit. The temperature monitoring unit displays the thermocouple's temperature on a front panel and contains two relays that can be pre-programmed to switch state at any desired temperature. These programmable switches are called set-points. One set-point is programmed to put one of the relays in a non-latching mode should the temperature exceed this set-point. This has been wired to send a signal to the control electronics that then prevents further actuation if the set-point temperature is exceeded. Once the stator has cooled enough (this temperature also being determined

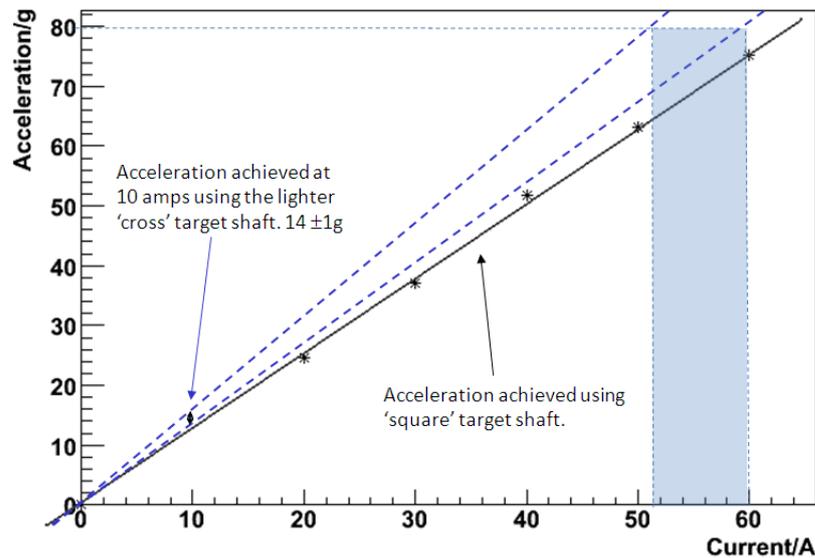


Figure 7.13: A plot showing the acceleration achieved using the trajectories shown in figure 7.12. A straight line can be fitted to this confirming that the accelerating force is proportional to the current as expected. When an extrapolation is done to the lighter cross shaped target then this shows that a coil current of about 50 A to 60 A would be required to accelerate the target to 80 *g*. This was confirmed experimentally several months later in the lab.

by the process monitor) the relay will reset without manual intervention allowing the actuator to resume operation. The temperature of set-point one is normally set a few degrees Centigrade above the normal operating temperature of the stator.

A second set-point is programmed to switch one of the process monitor's relays into a latching mode. The temperature at which this occurs is higher than that set for the first set-point relay. This relay cuts off the power to the power supply preventing any further power deposition into the stator. As this is a latching relay manual intervention is required to reset this. The second set point is a fail-safe for the first set-point. The first set-point assumes that the control firmware is functioning and that there isn't a fault anywhere in the system that is causing excessive power consumption. The second set-point kills all power to the stator preventing any further power deposition into the stator.

This temperature monitoring system as just described was used in the laboratory at Sheffield but the hardware has not been used on an installation in ISIS. This temperature protection system was replicated and improved by the control group at Daresbury Laboratory. This has given a system that is more flexible as it is now possible to monitor and log the temperature remotely via a PC terminal.

## **7.15 Conclusion**

This chapter has shown how the stator interacts with the permanent magnets attached to the shaft to create target motion through the appropriate switching of the stator coils through a set of six ‘states’. The method of utilising six transistors in a Hex-Bridge configuration to control the current through the three phases of the stator has also been illustrated. Of course, having knowledge and power is no good if there is no control. The next section will describe how the actuator is controlled by the target electronics and monitored with the target DAQ.

## Chapter 8

# Target Control, Electronics and the Data Acquisition System

This chapter shows how the target trajectory is controlled by the system electronics and describes the functioning of the individual electronic control systems that run the target mechanism. The interface between the electronics and the computer that records the performance of the target, the Data Acquisition (DAQ) system, is also discussed. Figure 8.1 gives an overview of the major components of the target system and how they relate to each other as installed at the Rutherford Appleton Laboratory.

### 8.1 Target Controller Overview

The actual implementation of the target control system is done with a system of integrated circuits (ICs) that include PIC micro-controllers, dedicated function ICs (i.e. amplifiers, digital to analogue convertors) and a significant number and variety of logic ICs. The system has been designed in functional modular sections, where each module has been constructed on circuit boards and then installed inside one of six crates. Each crate contains several interconnected circuit boards, the crates themselves communicate with each other through a system of multi-core cables on ‘D’ connectors. The total number of circuit boards in the target control system numbers over forty.

The single chip PIC microprocessors run programs from their own ‘on chip’ firmware. They interface to the electronics via their input/output (I/O) pins which are fully configurable in the firmware. Clearly, all of these microprocessors play an important part in the target system but there is one microprocessor that directly controls the stator switching and this one contains all the code and algorithms for directly running the target. This microprocessor is central to the target operation and is referred to as the ‘control PIC’ and is housed in the control crate. Compared to this microprocessor most of the other microprocessors perform minor, but important, peripheral functions.

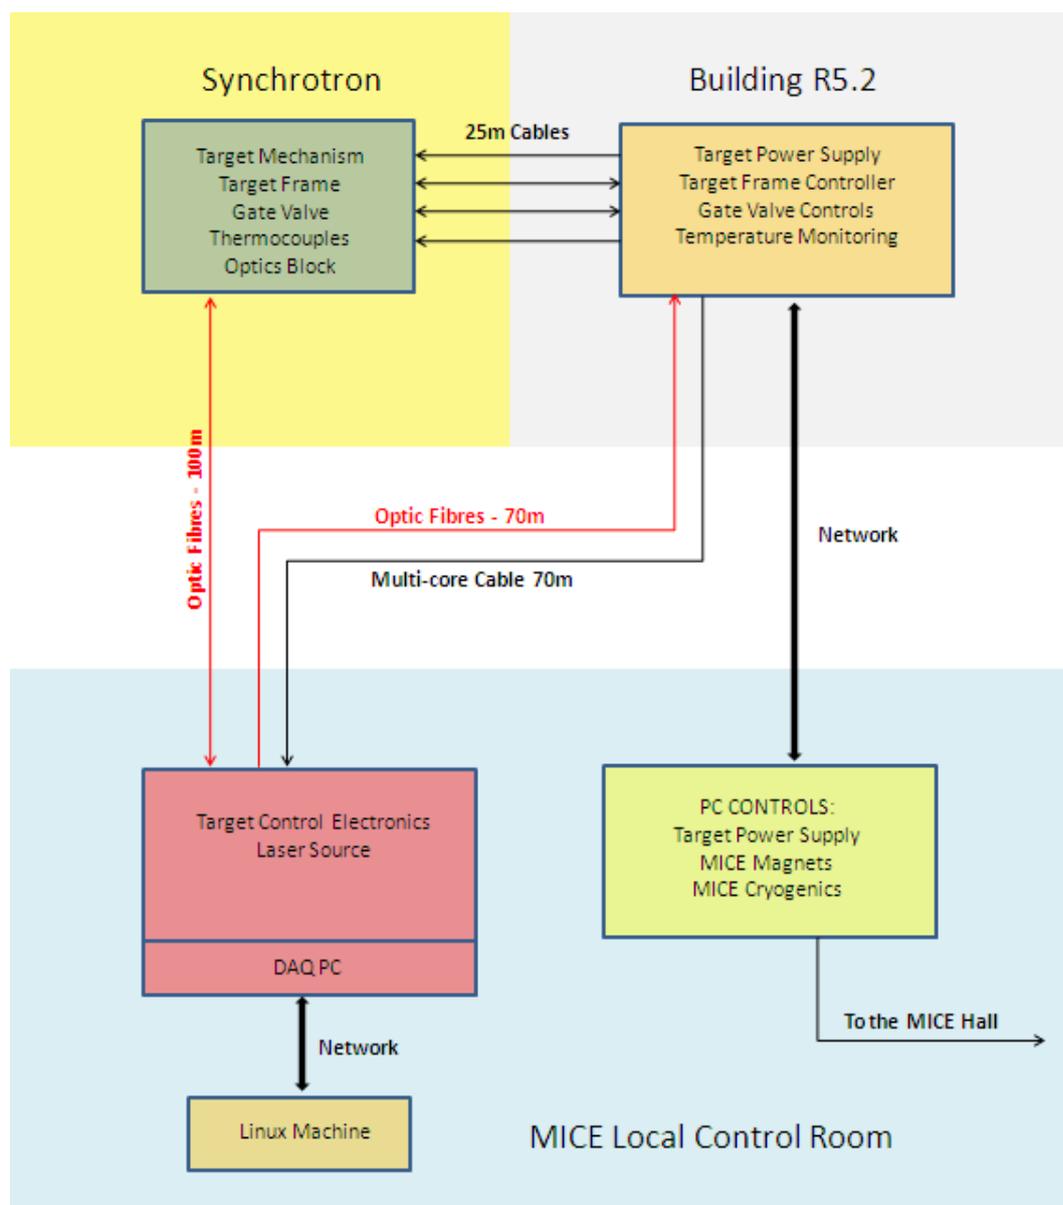


Figure 8.1: This figure illustrates the major components of the target mechanism installation at ISIS and how these components relate to each other. The power supply rack and the PC that controls it were built at Daresbury Laboratory. The target assembly, the DAQ PC and the control electronics were built at Sheffield University. The target frame was designed and built at the Rutherford Appleton Laboratory.

Figure 8.2 shows a block diagram of the four states that the target controller uses to control the movement of the target shaft and how they relate to each other. More details on the PIC processors and the individual electronic systems will be given later on in this chapter.

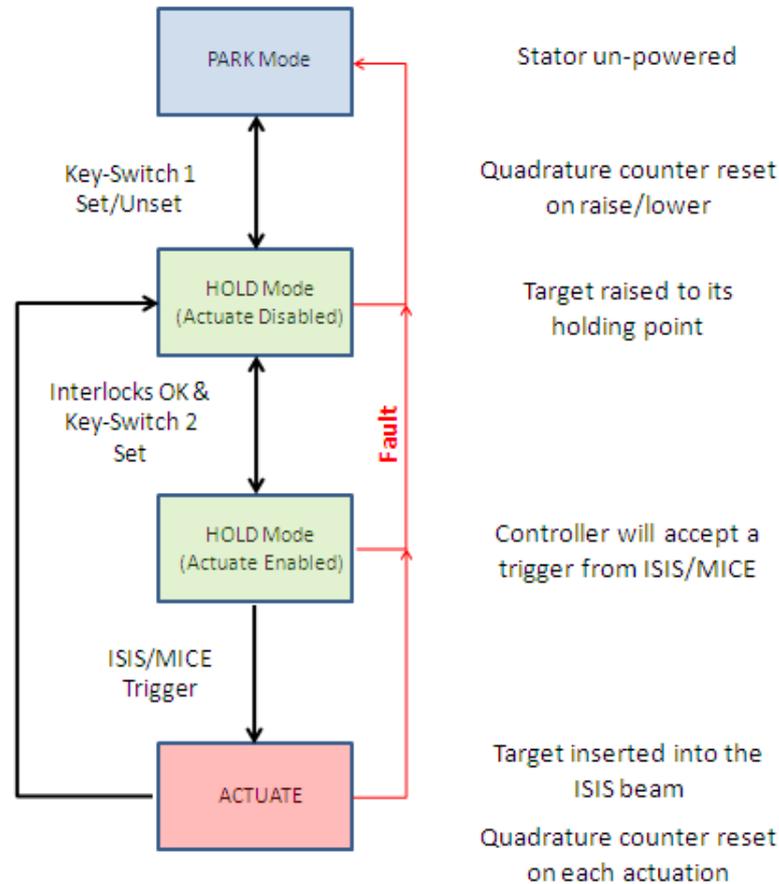


Figure 8.2: Diagram illustrating the four states of the target system.

## 8.2 Control: Park, Hold and Actuate Enable Modes

When the target control and power supply systems are first powered ‘on’ the stator itself is left un-powered. As noted in the last chapter, when the stator is un-powered the shaft stop is resting on the bottom bearing and the target is referred to as being in its ‘parked’ position. Before the target can be actuated the target must be moved into its ‘hold’ mode. This is a levitated position where the target is electromagnetically suspended at a zero-force point at a position inside the stator such that the target blade is held outside of the ISIS beam-pipe.

If the target frame was in its lowered position and the target was in park mode then this would leave the target sitting in the path of the ISIS beam. This dictates that when the target is in its un-powered state the target frame must be in its raised position so that the target does not intercept the ISIS beam. If the target frame is put into its

lowered position then the target must already be in hold mode or ISIS must be turned off.

A key-switch on the controller signals the controller to pick the target shaft up and raise it to its hold position. The actual hold position is determined by the firmware and is configurable by reprogramming the control processor. The holding position is achieved by advancing the zero-force point up the stator as described in Chapter seven. By picking up the target and moving it to its holding position this will take the optical vane, that is attached to the the shaft, through the index reset. This resets the quadrature counter so that an absolute position of the target shaft is obtained. The process of raising the target from park to hold is done blindly. However once the target shaft has gone through these steps the shaft should be in its holding position and this position should correspond to a predetermined position that should be indicated by the quadrature counter.

The controller checks the quadrature position to ensure that the shaft is where it expects it to be. If the shaft is not where it is supposed to be then the controller will flag an error. In the case where the position measured by the quadrature system is incorrect then this suggests either a problem with the quadrature counter system or that the target shaft has not been picked up. In both cases the actuator would be unable to actuate and so the system is prevented from going into the next state, 'actuate enable' mode.

The change of state from hold mode to actuate enable mode reflects a change of the internal state of the control system, the system will not enter actuate enable mode until certain internal and external conditions have been met. Internally, as already pointed out, the position counter must have been reset and be reading a value that the system expects. Externally all the interlocks must be good. At the time of writing this thesis these interlocks consist of a temperature interlock, and a target frame interlock. The temperature interlock prevents the target going into actuate enable mode if the stator temperature is above a preset value and is a non-latching interlock. The trip temperature is set on a temperature monitoring system and this system is integrated into the target power supply rack. This temperature system provides a go/no-go signal to the target electronics.

The target frame interlock is a latching interlock and works from the target frame limit switches. If the target frame is not in a well defined position; either fully raised or fully lowered, then the target will not actuate. This prevents actuation when the target frame is moving between states and there are two reasons for this. Firstly, if the target frame was moving and actuation was attempted the additional movement of the stator may cause a target drop. Secondly, if the system was actuating in the raised state and the target frame was accidentally lowered then this prevents the target from being actuated into the ISIS beam. Accidental actuation could cause ISIS to trip off if the target parameters are not set correctly.

Once the interlocks are 'ok' a key-switch needs to be turned in the front panel of the target electronics system to take the target system from hold mode into actuate enable

mode. This key-switch is also used to reset the latching interlock from the target frame.

### 8.3 Control: Target Actuation and Capture

‘Actuation’ refers to the state when the stator is actively accelerating the target into and out of the ISIS beam. During this state the target is accelerated so that it passes into the ISIS beam just before beam extraction and is then removed from the ISIS beam-pipe before the next ISIS spill begins. Actuation is completed when the target shaft had been electromagnetically recaptured at a zero-force point. Reliable capture of the target is essential; if the target is not captured reliably it will inevitably fall into the ISIS beam causing ISIS to trip off. Whilst dropping of the target shaft is unlikely to cause any damage to the target itself it could cause significant disruption to ISIS if it occurred frequently.

The target actuation and capture is a finely balanced process that took a significant amount of programming effort to ensure that the process worked reliably at high acceleration. An overview of how target actuation is achieved from the perspective of the coil/magnet relationship was given in chapter 7. As stated in Chapter seven the controller and power supply have been designed to apply a constant force during the actuation as this gives several distinct advantages. Foremost amongst these is the ability to accurately control the dip depth. This section provides a perspective of the actuation sequence from the point of view of the controller.

For the target to be actuated the position of the target must be accurately known. The current system uses the eight most significant bits of the nine bit position number to give a byte size position. Actuation is performed actively; the coils are switched so as to provide the maximum force on the permanent magnets at all times during the actuation process. For the maximum force to be maintained on the magnets the coil switching has to track the current position of the magnets. The controller is also provided with a signal that gives the direction of the target motion. Whilst this is easy to compute from the position it does require some computational overhead and so the function of determining the direction of motion has been given to a separate microprocessor that provides a signal to the controller.

From a control point of view the actuation process goes through four distinct stages and these stages have been named ‘quadrants’. The actual capture of the target is also considered as a separate process in its own right. These four quadrant states are shown diagrammatically with reference to a target trajectory in figure 8.3.

Upon receiving a trigger signal the control system enters ‘quadrant 1’ of the actuate sequence. Here the controller signals the stator coils to switch so that the shaft accelerates downwards. The position of the shaft is tracked and the coils are switched to keep the maximal force on the shaft until the shaft reaches position ‘switch-point 1’. Switch-point 1 is the position at which the coil currents are reversed and at this point

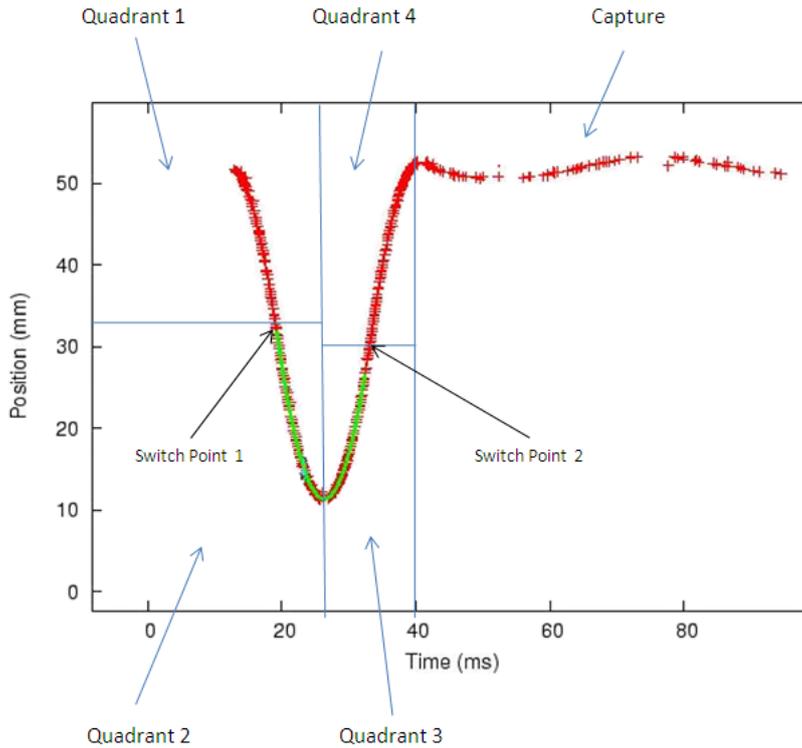


Figure 8.3: The actuation trajectory is split into four quadrants from the controllers point of view. Switch Point 1 reverses the coil currents and so determines the actuation depth. Switch Point 2 determines the capture point. These two points are not at the same position due to the decay on the capacitor bank.

the target begins decelerating and the controller enters into ‘quadrant 2’ of the actuate sequence. The target is decelerated until it reaches a position where the motion of the target changes direction as indicated by the quadrature counter. At this point there is no physical change to the system and the coil currents are kept the same; however this change in the target direction defines the point at which the controller state moves into ‘quadrant 3’. The target is now accelerated upwards until the target reaches the position defined by ‘switch-point 2’. At this point the coil currents are reversed again and the target begins decelerating; this is accompanied by the controller state moving into ‘quadrant 4’. The target continues to decelerate and at the point where the direction of the target’s motion changes again the controller moves to the capture state.

When the controller enters the capture state this point it calculates where the nearest zero force point is to its current position. (Recall that the zero force points is a discrete function with a resolution of 3 mm.) It then switches the coils to place this zero force point as close as possible to the target’s current position; the target is then left to be passively captured into this zero-force point. Ideally if ‘switch-point 2’ is set to the

correct value then capture position should be the same as the hold position. If the zero-force point does not correspond to the target's hold position then after capture the target will then be moved to the hold position ready for the next actuation.

Switch-point one defines the target's dip depth and to a first approximation it is half the way down the trajectory. It is possible to change its value via a switch on the front panel of the target electronics. The target electronics only sees the 8 most significant of the 9 bits that describe the target's position so it effectively sees the target with a resolution of  $300\ \mu\text{m}$  represented by an 8-bit number. Changing the value of switch-point 1 by a single digit therefore increases the target dip depth by approximately double this distance i.e.  $\approx 600\ \mu\text{m}$ .

Switch-point two is set lower than switch-point 1 because the decay on the capacitor bank means that the rate of acceleration in quadrant one and two is greater than the rate of acceleration in quadrants three and four. Changing the dip depth changes both the switch points by an equal amount. Multiple changes of the switch point values can result in the target either under or overshooting the hold point and being captured either at a higher or lower zero force point. If this happens persistently then the controller will automatically make a correction to switch-point 2 until the target is being captured at the hold point. This system of using 'quadrants' to define the actuation cycle allows the target to be accurately tracked by the controller as it passes through the trajectory ensuring reliable actuation.

The whole actuation process is position dependent; the trajectory is determined solely by the acceleration of the target and the switch points. This results in a small amount of jitter from pulse to pulse, both in terms of timing and actuation depth. The jitter is small and a quantitative measure of this will be shown in the next chapter. The deviations in the target's trajectory have a notable effect on capture as the amount of damped oscillation seen as the target moves into its zero-force point varies from pulse to pulse.

### 8.3.1 Issues

Moving between the 'park' and 'hold' modes makes use of the zero-force points that can be electromagnetically created within the stator as described in chapter 7. As alluded to in that chapter, the levitation of the shaft can be done passively without the requirement for positional feedback. To recap, this method of moving the target offers a couple of distinct advantages. Firstly, this mode is naturally stable and so it significantly reduces the complexity of the control electronics and the firmware algorithms. Secondly, it is useful in obtaining an index reset of the quadrature system when the electronics is first powered up as the shaft can be moved from park to hold blindly. Despite the advantages that this method offers in terms of the ease of moving the target shaft about within the stator body it does have a couple of problems.

The first problem is caused by the fact that the number of stable passive positions is

determined principally by the number of coils, and the distance between these stable positions is determined by the pitch of the coils, 3 mm. This presents a problem when the shaft is moving out of actuate mode and into hold mode because the position at which the target comes out of actuation and the position of the nearest zero-force point do not necessarily align. The restoring forces cause damped harmonic motion of the target shaft during the capture process because the shaft acquires a velocity moving between these two points. These damped oscillations can take several hundred milliseconds to settle. Trying to capture the target without first arresting the velocity of the target shaft can result in the amplitude of the oscillation being so large that the magnetic potential well cannot trap the target and the target is then dropped into the ISIS beam.

This issue became problematic when currents of greater than 20 A were used. At this current the target capture was not reliable and the target was dropped frequently by the controller as it entered into capture. The cause of this unreliability was traced to the delay in the quadrature system detecting the target changing direction from moving up to down. The target would have to move a minimum of 150  $\mu\text{m}$  in the downward direction before the quadrature system detected this change of direction in which time the target has received a downwards 'kick' from the accelerating force. This small downwards kick gave the target enough velocity to make the target capture unreliable.

To correct this a piece of code was added that gave the stator a very short current pulse at full actuate current providing a decelerating force in the opposite direction to the target's motion as the target enters into capture mode. The duration of this pulse is hard coded and typically lasts around 1 mS. The short pulse was called a 'kick-back' as it has the effect of removing the velocity acquired by the target due to the latency of the quadrature counter. The kickback significantly improved the reliability of the capture mechanism and allowed the actuate currents to be increased up to 60 A, the current required to obtain the necessary target acceleration.

The second problem is also caused by the damped harmonic motion. When picking the target up from park to hold or taking it from hold to park the movement from these two modes is done by moving the zero-force within the stator, as previously discussed. As these points are separated by 3 mm, a significant amount of damped oscillation occurs between each step. These oscillations must be given time to dampen out before the next step is taken else the timing of the steps can cause these oscillations to become resonant, this causes the target to be dropped. Typically a delay of 500 ms is required between steps. These delays mean that moving the target by this method is a relatively slow process and it takes several seconds to raise the target from its parked position to its hold position and vice versa. This is not normally an issue, however if the target were ever to be accidentally dropped into the ISIS beam it cannot be removed quickly, nevertheless ISIS would stop operating due to a beamloss trip.

## 8.4 The ISIS Trigger Signal

The triggering system for the target is fairly straightforward but has been designed with some flexibility in mind. In designing this system three assumptions were made:

1. The target needs to synchronise with ISIS, but the timing signal that is received from ISIS may not arrive when the target system needs to start actuating.
2. The exact time that the target intercepts the ISIS beam may need to be altered.
3. The target may also need to synchronise with another component of MICE.

Before considering these assumptions it is useful to consider some of the timing implications for running the target system.

The whole target actuation sequence from start to finish takes about 30 ms so it takes about 14-15 ms to get from the hold point to the apex of the actuate trajectory. The actual timing is dependent upon the actuation depth and the target accelerates harder when moving into the beam-pipe than when it is moving out because of gravity ( $\approx \pm 1\%$  each way) and the capacitor bank discharge ( $\approx 10\%$  deterioration in current through actuation.)

The ISIS beam is injected into the synchrotron just 10 ms before it is extracted, so the target system needs to start actuating before the beam is injected into ISIS if it is to catch the beam during the last 2 ms of its time in the synchrotron. The ISIS triggering system is flexible and can provide a timing signal from 5 ms before injection up to extraction. Unfortunately this makes it impossible for the target to trigger off a given ISIS pulse and to intercept that pulse 2 ms before extraction.

It was correctly foreseen that the timing of the target actuation would need tuning to optimise the particle production for MICE. This is for two reasons. Firstly, the target trajectory is dependent on the actuation depth. The deeper the target has to actuate then the earlier it must be triggered as it takes longer to reach the apex of its trajectory. The second reason is to allow the timing of the interception point of the target with the ISIS beam to be altered. Hitting the beam slightly earlier or later may give unforeseen advantages with respect to particle production for MICE. The acceleration of the target during actuation is effectively fixed so to enable an adjustable interception of the ISIS beam requires some mechanism to control exactly when the target enters actuation mode with respect to the injection of the ISIS beam into the synchrotron.

Both of these problems are solved by the use of programmable delays. To solve the first problem of synchronising the target with ISIS so that the target can be inserted into the last 2 ms of a given ISIS pulse it is necessary to use the ISIS MS signal from one spill to trigger the target to be inserted into a later ISIS spill. The ISIS control room wishes to monitor the beam loss produced by the target and they have a system that

is set up to specifically monitor the beam loss on those spills that coincide with the trigger signal sent to the target. If the beam loss regularly exceeds a given limit on the spills that coincide with the target trigger signal then ISIS will be stopped. Because of this setup the target cannot be inserted into the next ISIS pulse after receiving the trigger but must wait for a further  $2^n$  ISIS spills before being inserted into the beam. ( $n$  is determined by agreement with ISIS. Typically  $n = 8$ ) This means introducing a programmable delay of  $2^n \times 20 \text{ ms} - 15 \text{ ms}$ . The value of 15 ms is not critical, as a second delay is triggered from the output of the first delay and the output of this second delay is used to obtain the actual actuation start time. The second delay is set from the front panel of the target electronics and it has a  $100 \mu\text{s}$  resolution on an eight bit counter giving it a range of 25.5 ms.

As the target may be used in a mode where it synchronises to one ISIS pulse and actuates to intercept a later one there must be some way of communicating this to the rest of the MICE so that MICE can be synchronised to the target operation. The target system has outputs that are used to indicate that the target is actuating but additionally it also incorporates a ‘target ready’ output and a ‘MICE Ok’ input allowing it to handshake with an external MICE synchronisation system. This would prevent the target from actuating until the rest of MICE is ready; for example it will be useful for ensuring that the target only actuates when the MICE RF and the MICE DAQ are ready. Such a system has not yet been implemented but it is envisaged that it will be necessary as further components of MICE are installed [106]. When such a system is implemented the target electronics will send a signal out to the MICE synchronisation system to indicate that it is ready to actuate. The target system will then wait for a confirmation signal from the MICE synchronisation system; this will confirm that all the other systems on MICE are ready. The target system will then use the next valid pulse from the delay generator to trigger operation of the target. When this MICE synchronisation system is not required the ‘target ready’ output signal is just looped back into the ‘MICE Ok’ input signal. Figure 8.4 shows how the target triggering system has been designed.

## 8.5 The Electronic Control System

This section gives a more detailed overview of the electronic systems that control the actuator. In total, the control system is composed of several ‘crates’ where each crate provides independent functionality to the system. Every crate is isolated from all the other crates and where communication between different crates is necessary this is done via optically-isolated connectors. The isolation helps to reduce the possibility of earth loops and the fixed bandwidth of the opto-isolators attenuates any noise present on the signals. As expected the target electronics has ‘evolved’ considerably from the early prototype circuits that ran the prototype stators. As the system was developed more complex circuits were added to the target electronics to improve its operation. This evolution has led to a very modular system that is built upon many independent functioning units. The following is a list of the crates that comprises the target system,

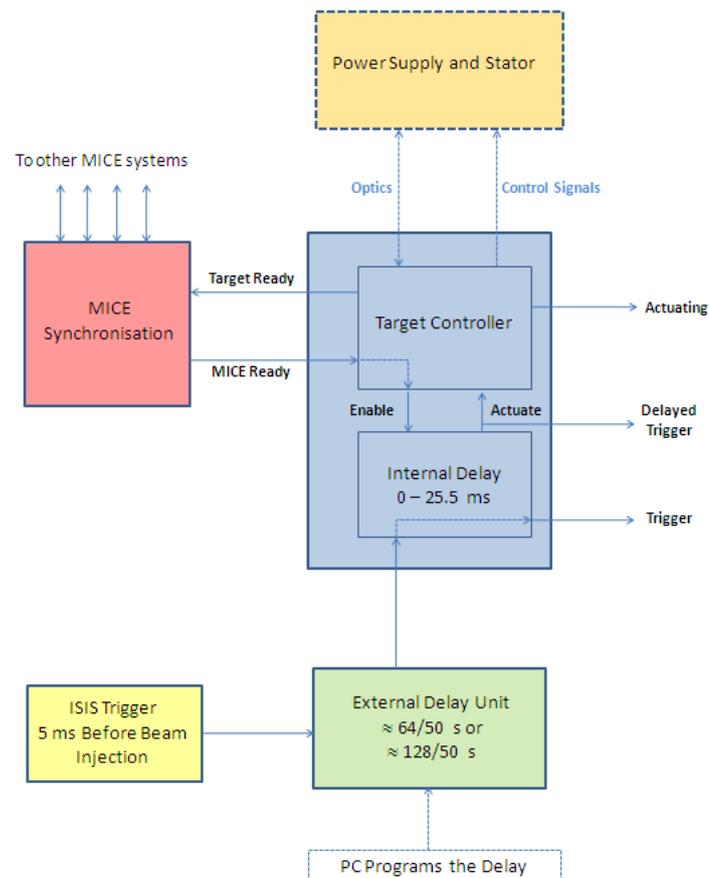


Figure 8.4: Diagram illustrating the target triggering system. At present the MICE synchronisation is not utilised and the ‘target ready’ output is looped back into the ‘MICE ready’ input.

the overall functionality of each crate will be discussed shortly.

- Power Supply
- Temperature Monitoring
- Quadrature Optical Amplifiers
- Quadrature Counter
- 24 Bit Timer
- Switch Panel
- Controller

Now that considerable experience has been gained understanding how to build a functioning actuator it is expected that much of the modularity in the control system will be re-integrated into a much smaller footprint. It is also expected that with continued development the control system will be turned from its present incarnation as an ‘expert system’ into something that is much more ‘turn-key’ and user friendly. The rest of this section describes the functionality of the control system in further detail.

### 8.5.1 PIC Micro-Processors

The target control system makes extensive use of PICs, a Harvard architecture micro-controller manufactured by Microchip. Microchip manufactures many different PICs with a variety of peripheral functions, I/O and processing capability. Many of these PIC’s can be flash programmed via a suitable programmer. Most of the flash PICs can be reprogrammed many times, this feature makes them ideal for development work. Principally three different PICs have been used in the target electronics and these come from two distinct families; the ‘16F’ and the ‘18F’ families. Both of these families are based on an 8-bit architecture.

The 16F PICs can run at 20 MHz and have a throughput of up to 5 million instructions per second (MIPS). They have a small instruction set comprising of 35 instructions. The PIC chosen as the controller is a PIC 16F871, a 40 pin PIC with 32 I/O lines, 2 KB of flash memory and a 10 bit resolution PWM module.

Most of the peripheral functions are performed by 18F PICS. These run at up to 40 MHz and have a throughput of up to 10 MIPS. They also have approximately double the instruction set and more built-in functionality. The two PICS used for peripheral functions are the PIC 18F1220 and the PIC 18F4520. The PIC 18F1220 has 20 pins with 4 KB of flash memory and is suited where a minimum of I/O is required. The PIC 18F4520, a 32 KB flash device, has 40 pins with up to 36 of these configurable as I/O.

It may strike the reader as odd that the peripheral functions use more powerful PICs than the main controller. The controller code started development before many of the peripheral functions were designed and at this time a PIC 16F IC seemed an ideal processor to start programming with. As confidence grew with using these devices a more complex device was chosen when processors were required elsewhere in the circuitry. Porting the 16F code over to an 18F chip was considered but it would have been time-consuming and would have meant a board redesign. As a result the main processor is currently the least powerful of all the microprocessors on the target system.

The firmware for the microprocessors is developed under a Windows application called the Integrated Development Environment or IDE. Once the assembly code has been compiled it can be downloaded to a programmer that flashes the firmware into the PIC’s memory. Further information on PICs, the IDE, associated hardware and application notes can be found on microchip’s website [107].

### 8.5.2 Power Supply

The power supply crate, as its name suggests, has the sole function of providing power to the rest of the control electronics crates with the exception of the temperature module. This crate runs directly from the mains supply at 240 V AC. All of the other electronics crates run off a single +5 V power supply. Each crate is protected with a 1 A fuse although typically each crate only consumes a couple of Watts of power.

The power supply crate contains one linear power supply that can provide 9 A at 5 V and two switch-mode power supplies. One of these provides 5 V at 1 A and the other provides a 24 V output at 1 A. The 9 A supply feeds six 5 V outputs on the back plane of the crate. These six outputs are then used as power feeds to five of the electronics crates, leaving one spare channel. The other isolated 5 V output is fed to a single output on the back of the power supply and is provided as an independent isolated 5 V supply. Currently this is utilised to provide power for an interlock loop on the over-temperature system in the laboratory. The 24 V power supply also goes to a single output on the back panel of the power supply. This 24 V output is used in the laboratory to power an external relay that switches power on to the stator power supply and forms part of the temperature interlock system.

### 8.5.3 Temperature Monitoring Module

This unit contains two commercially bought temperature process monitors that have dual relay outputs. This module is used to monitor the temperature of the actuator system when the system is running in the laboratory. This unit is not used for the actuator that runs on ISIS as an independent temperature monitoring system has been built by Daresbury Laboratory for this purpose.

The first of these units monitors the cooling water temperature on the output side of the stator. None of the output relays are used on this unit and it is used to monitor the water temperature only. The second unit monitors the temperature of a particular pair of coils in the stator body via the K-type thermocouples that are inserted into the coil stack during manufacture. The functionality of this unit has already been discussed in section 7.14.

### 8.5.4 Quadrature Optical Amplifier

The quadrature optical amplifier contains three identical circuit boards, one for each channel of the quadrature system. A figure of the main functions of one of these circuit boards is shown in fig 8.5.

A HB880IR photodiode converts the laser light that falls onto it from the optical fibre into an analogue signal. Typically the signal from the photodiode is a few tens of

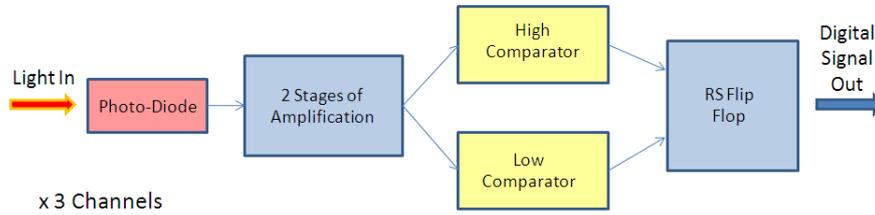


Figure 8.5: The laser light from the optical system is converted into a square wave digital signal. There are three of these amplifiers, one for each channel of the optical system. The outputs from these boards are sent to the quadrature counter crate where the signals are counted to give the absolute position of the target shaft.

millivolts and so it requires amplification to be useful. The primary amplification is done by an instrumentation amplifier as this helps noise rejection and therefore gives a cleaner signal. The second operational amplifier has a multi-turn variable resistor in its feedback path so that its gain can be accurately set. Typically this is set so that a 12 V ground to peak signal is observed. This analogue signal is then fed into a pair of comparators. One comparator is set to trigger at one third the ground to peak voltage and the other at two thirds the ground to peak voltage. These comparators feed the ‘Reset’ and ‘Set’ inputs of an RS flip flop. The output of the flip flop is taken as the output of the quadrature channel.

Two comparators and an RS flip flop were used to help prevent multiple transitions on the output as the signal passes through the comparator threshold. The high level of gain meant that multiple transitions of the comparator were sometimes observed on earlier prototypes due to the comparators causing a small voltage spike on the power rails at transition that was fed back into the input of the circuit. This fault could potentially lead to miscounts. The dual comparators also allows the transition levels to be independently set on both the rise and fall of the signal on each channel. This allows a degree of control over the setting of the duty cycle on each channel of the quadrature signal; ideally this needs to be as close to 50 % as possible.

### 8.5.5 Quadrature Counter

The quadrature counter crate contains several circuits performing different functions. The primary function of the quadrature counter crate is to take the quadrature signals from the quadrature optical amplifier crate and to convert them to an absolute position. This is done by a specialised quadrature counting IC, the HCTL-2032, interfaced to a PIC microprocessor.

A second set of boards inside the quadrature counter then adds a fixed programmable

offset to this computed position, this offset is called an ‘index offset’. The index channel normally resets the counter at just below the hold position. Without an index offset then the ‘raw position’ would roll-over during actuation because the position of the index causes the quadrature counter to reset to a count of zero. This behaviour is not particularly problematic but it is easier if the counter doesn’t roll over during the normal range of motion of the target shaft. It makes programming the look up tables in the control processor easier which in turn reduces the chance of any bugs being introduced into the code. The second big advantage of adding an ‘index offset’ to the ‘raw position’ value is that if the index marker moves due to small positional variations from vane to vane or if the absolute position of the index collimator on the optical block moves then it is easy to correct for this by altering the ‘index offset’ value that is added to the ‘raw position’ value.

As the mechanical setup has become more refined these issues have become much less of a problem but in earlier setups the position of the index could change enough from shaft to shaft to warrant making a change to the index offset value. Without this facility the lookup tables in the PIC that correlate the target’s position to the coil state would have needed recomputing and reprogramming every time the index changed position; a time consuming and error prone task.

The quadrature counter crate contains two digital to analogue convertors so that an analogue output of the target’s position can be viewed externally on an oscilloscope. Another microprocessor on board determines whether the quadrature signal is counting up or down and there are several signals produced that allow this crate to synchronise to the ‘24 bit timer’ crate. A diagram showing this crate’s functionality is given in figure 8.6.

### 8.5.6 24 Bit Timer

This crate was designed to provide timing information on position changes to the DAQ system. This crate contains a 24-bit binary counter that counts upwards at a rate of 5 MHz. Counting starts when a trigger signal is received and finishes 100 ms after the end of actuation. The system then automatically resets the counter ready for the next actuation. The counter value is latched into a buffer when the target system sees a change in the position of the target shaft; this is signalled from the quadrature counter crate. This means that for every positional change, the time that this took place is latched by the timer and then recorded by the DAQ.

The reason for doing this was that the DAQ could only support an input data rate of  $\approx 20$  kHz. Therefore the timing resolution of the DAQ is poor; it does not give the required accuracy to allow the acceleration of the target to be computed accurately. By taking this function away from the DAQ and giving it to an external system the time that each positional change occurred was recorded to 200 ns accuracy. Because the time was only latched when the target system registered a positional change this gave good timing information on the position without exceeding the bandwidth of the

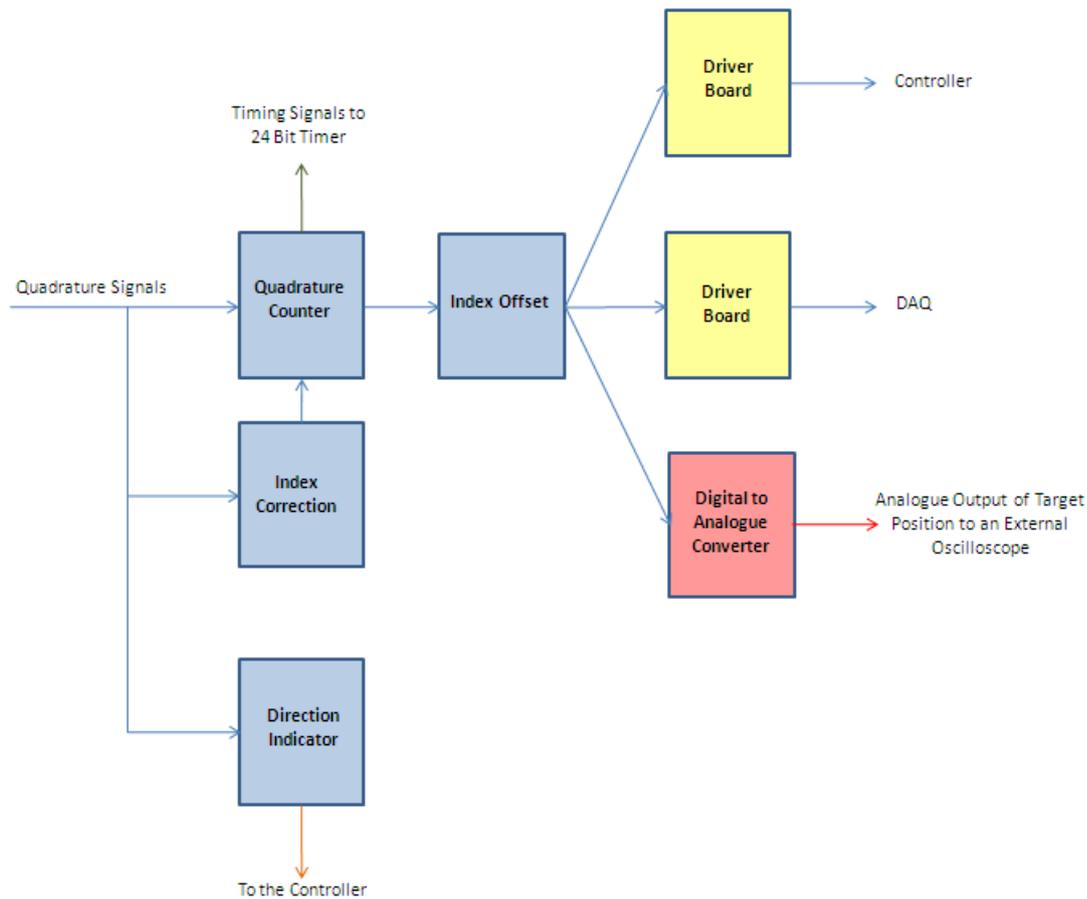


Figure 8.6: The quadrature counter crate takes the quadrature signals and converts them into an absolute position. This position, a 9-bit number, is then output to the controller and the DAQ system. A digital to analogue converter also gives an analogue output that is used to display the target position on an oscilloscope.

DAQ system. Further information on the DAQ system will be provided later on in this chapter. The timer circuit is conceptually very simple, although building a triple buffered 24 bit timer from logic IC's did require several circuit boards. A diagram of its operation is shown in figure 8.7.

### 8.5.7 Switch Panel

The switch panel crate was initially designed to provide a place to locate the key-locks and the raise/lower actuation depth switches to operate the target controller. However as there was some space left in the crate after construction it has also been used to

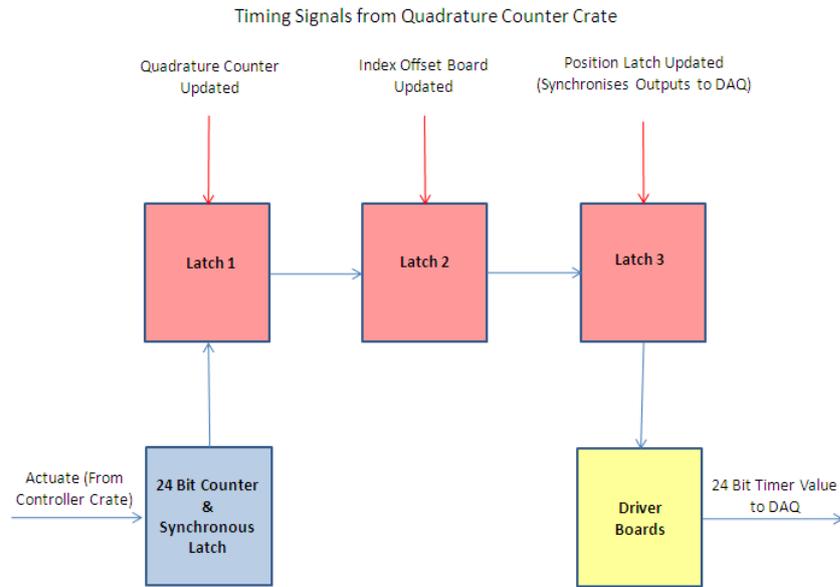


Figure 8.7: The ‘24 Bit Timer’ crate provides timing information on every position change that occurs during actuation. This information is synchronised with the quadrature counter and fed into the DAQ.

provide the target frame status and its interlock to the actuate enable key-switch. The crate also houses an analogue to digital converter that provides the DAQ with a 10-bit digitisation of the ISIS total beam loss signal. A one bit conversion of the beam intensity is also done; this effectively provides a beam/no beam signal into the DAQ.

Switching the stator from park into hold mode and switching the controller from hold mode into actuate enable mode are done with key-switches to prevent the target from being mistakenly actuated. The target frame is interlocked into the actuate enable mode so that if the frame is not in a well defined position, either fully up or fully down, then the controller will not allow the stator to actuate.

### 8.5.8 Controller

The controller crate contains the main microprocessor board that controls the currents through the stator coils. The method by which levitation and actuation is achieved has been discussed at length in chapter 7 and earlier in this chapter.

Also located in this crate are the trigger boards and the second digitally controlled delay. The value of this delay is set from the front panel using a rotary quadrature encoder switch. This crate also contains two circuit boards that drive high powered Infra-Red LEDs. These LEDs connect to the optical fibres that transmit the IGBT switching

signals to the Hex-Bridge located in the target power supply rack. One of these circuit boards switches the LEDs that signal the high side IGBTs whilst the other boards have their outputs logically AND'ed with a PWM signal generated by the controller; this drives the LEDs that are connected to the low side IGBTs. This crate also provides the triggering signals for the 24 bit timer and the DAQ. Several other signals giving information on the target timing and triggering are fed out to BNC connectors on the back of the crate, these are intended for providing synchronisation signals to external equipment

The controller crate contains an interlock board which allows up to four non-latching external interlocks to be connected to the controller preventing the target from actuating. These can be configured as active high or active low inputs. One of these interlocks is used as a temperature interlock when the actuator is being used in the laboratory. Another will shortly be utilised as part of the personnel protection system within ISIS and the MICE Hall; the stator will stop actuating if anyone enters the MICE hall whilst it is operating. A diagram of the main functions of the controller crate is shown in figure 8.8.

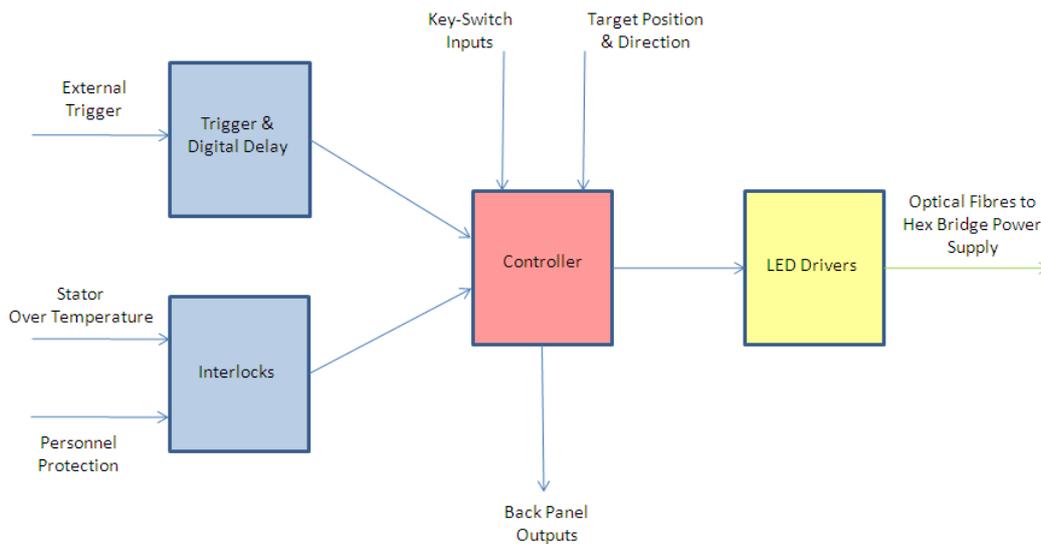


Figure 8.8: The controller is the heart of the system as it directly interfaces to and controls the motion of the target shaft within the stator.

## 8.6 DAQ and the Software Monitoring

Monitoring and recording of the target position is necessary on an actuation by actuation basis to allow analysis of the performance of the target mechanism. This is done

by the target data acquisition system or target DAQ. As the MICE experiment progresses it will also be a requirement that the data collected by the target DAQ can be integrated with the particle detection data received from instruments that are around and part of the MICE experiment. This will allow the performance of the target to be correlated to beam conditions for MICE.

This target DAQ encompasses both the hardware that is required to physically allow the target data to be recorded and the software that transfers this data from the hardware onto a PC. The software will also enable both online and offline analysis of the data once it has been recorded. The following sections describes the construction of the target DAQ system. Analysis of the data collected by operating the target system will be shown in the next chapter.

### 8.6.1 Hardware

The DAQ hardware starts with a local PC computer that contains a specialised PCI interface card; this PC sits in close proximity to the target electronics. Data that is sent from the electronics is wired to the interface card and data is then recorded by the computer. The PCI card used is a an isolated 64 channel digital input card manufactured by Amplicon. This card is double buffered so that one buffer can be read whilst the other continues to read data. The following inputs are wired into the card and recorded by the computer.

1. Target Position - 9 bits.
2. Target Time - A 5 MHz timer starts when the target actuates. The time is latched and synchronised to the target position outputs - 24 bits.
3. Target Actuate Status. This is used as a trigger to start the PC recording what is seen on the DAQ onto the computer's hard drive - 1 bit.
4. Beam Loss. This records a digitisation of the total beam loss signal produced by ISIS - 10 bits.
5. Beam On. This is a single bit that shows when there is beam in the ISIS synchrotron - 1 bit.
6. Target Frame Position - The state of the target frame is recorded on these bits. This tells the user whether the actuation data was obtained with target frame lowered, raised or somewhere between the two - 2 bits.

This gives a total of 47 bits recorded by the target DAQ. The 'Total Beam Loss' and 'Beam On' are signals that are generated by ISIS and then inserted into the target controller for digitisation before being sent into the target DAQ data stream.

The target DAQ is an event based system rather than a time based system. The difference between the two and the consequences of this require a little explanation. The DAQ card is continuously recording what it sees on its input pins into its buffers at a rate of approximately 20 kHz. These buffers are monitored by the PC and when it sees a trigger signal the PC creates a new file, and selected data in the buffers are recorded to the computer's hard drive. The computer only records data to the hard drive where it sees a change in the target's position. (The change in position is also accompanied by a change in the time data as this is latched by the 24 bit timer when the target changes position.) The other inputs to the DAQ card can change asynchronously with position and time but their values are only recorded if the PC sees that the position value has changed. Recording them at any other time doesn't give any more useful information because the time stamp produced by the 24 bit timer correlates to the change in positional data only. Once the actuate signal disappears, the PC continues to record the data for another 100 ms. Because the inputs are only recorded to file if the PC sees a change in the position of the target shaft the DAQ is event based rather than time based.

The target DAQ was initially built in the laboratory to assist with debugging the controller during development of the firmware. The advantage of using this event based system is that it gives accurate timing information on each positional change without swamping the DAQ with lots of data. The system worked well in the laboratory for the purpose for which it was built. The plan at the time of its construction was to replace this system with a time based DAQ before the target was installed at RAL. Unfortunately other developmental issues prevented this system being upgraded before the target was installed at ISIS and so the developmental DAQ system is now currently being used at RAL.

The use of this event based DAQ system has caused a subtle problem on the installed system. Because the DAQ channels only record data when the target shows displacement the data rate into the DAQ is proportional to the velocity of the target. This means that the total beam loss measurement loses resolution when the target is moving slowly, such as at the apex of its trajectory. It is unfortunate that it is at the point when the target enters deepest into the ISIS beam that the beam loss signal is recorded at its worst resolution as it is at this point that this signal is of most value. It is for this reason that the current event based target DAQ system is now less than ideal for purpose. The upgrading of the target DAQ system to a time based system so that this problem can be resolved was a high priority item at the time of writing this thesis. Despite this caveat, the target is still accelerating rapidly so despite the loss of resolution at the apex enough data is gained to be of use. Analysis on the target performance done using data gained from this DAQ system will be presented in the next chapter.

### 8.6.2 Software

The software for the target DAQ system can be considered in two parts. Firstly there is the code that talks directly to the DAQ card and places the data onto the computer's

hard drive as text files. Secondly there is the software that post-processes the raw data in these text files and produces plots, typically showing how the key parameters vary with respect to time. These plots are usually available for viewing whilst the target system is running giving a graphical interpretation of what is happening in almost real time.

There is of course a third type of software and this is the analysis code. The exact application of the analysis code is clearly dependent upon what is wanted to be achieved but typically it is written to look for long term trends and patterns in the target data and to correlate this with data from MICE, ISIS and the MICE beamline. This type of analysis software is not included in this discussion because this section describes the software used in the day to day running of the target DAQ. However code for analysing the target data has been written and is responsible for much of the analysis presented in the next chapter.

### 8.6.3 Data Acquisition Code

As indicated in the last section the primary data acquisition is performed by a short piece of software that polls the data acquisition in a continuous loop until it sees a trigger signal. At this point it opens a file on the computer, gives it an event number, and records the target data into this file. A new directory is created for every 10,000 pulses and the directory structure is based upon the date appended by a number. No compression is used on the files so they are readable with a text editor.

Each target actuation has its own independent dated file. This allows an accurate count of the number of actuations and also helps to place a time stamp on any analysis looking for long terms trends in the target behaviour. The time-stamping is also necessary to correlate target pulses with data obtained by other DAQ systems on MICE. There is a long term goal in MICE that all of these separate DAQ systems will have their data automatically integrated by an EPICS system into a single data stream. This will help to ease future analysis and prevent errors occurring due to mis-matched data sets from different DAQ sources.

### 8.6.4 Plotting and Reduction of the DAQ data

The data reduction and plotting software has undergone several revisions since its first inception. In its current format the plotting software, the MICE analysis program (MAP), can run concurrently with the DAQ software and provides near realtime graphical updates of key target parameters. MAP also provides a data reduction mechanism.

When plotting DAQ data MAP plots the target trajectory and computes a third order polynomial fit to the apex of the trajectory. It also plots the beam status (beam on/off) and the total beam loss. A fit is also done to the beam loss, and the integrated beam loss and maximum beam loss are calculated. Before plotting, this data is cleaned to

remove spikes in the analogue channels (real noise) and aberrations in the digital data due to sticky bits. From the polynomial fit to the trajectory the acceleration of the target is computed as well as its actuation depth.

Sticky bits are when the input bits on the DAQ card do not all change at exactly the same time. The result of this is that a single change to a value (such as position) looks like multiple changes in the value to the DAQ system within a very short space of time. This occurs most frequently when the binary representation of the value changes from  $2^{n-1}$  to greater than  $2^n$ . Most of these points can be removed by applying filtering algorithms to the data but because this problem occurs for both the position and time data points it is difficult to remove every single faulty point. Further work is needed in this area to understand the effect that these points have on the analysis of the target data. A screen-shot of the online DAQ monitoring software is shown in figure 8.9.

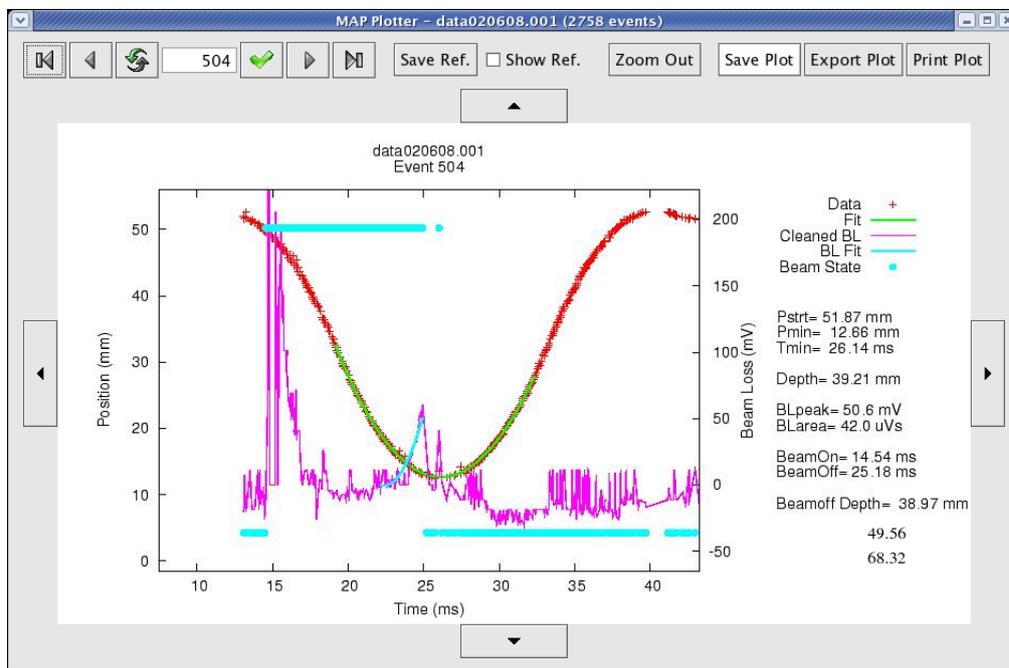


Figure 8.9: A screenshot from the MAP program showing target data from a single actuation being plotted. The red line represents the cleaned target trajectory as recorded in the data stream. The green line is a third order polynomial fit. The cyan line shows where there was beam in the synchrotron. (The value goes high.) The pink line shows the recorded total beam loss signal. The light blue line is a fit to the beam loss data. Computed parameters are shown displayed on the right hand side of the plot area. The data rate is proportional to the target velocity. This produces blind spots in the data.

MAP can reduce the 10,000 event data set from each directory by examining the text files and extracting key parameters from each target actuation. These key parameters from each one of the 10,000 actuations in each directory are then written to a single file, hence the name ‘reduction software’. The reduction software may be utilised in

the future for passing key parameters to the MICE DAQ system and for long term archive of target pulses. The key parameters that the reduction software pulls from the individual actuation files are:

- Target Actuation Depth
- The Time to Minimum Dip Depth
- Acceleration To Minimum Dip Depth
- Acceleration From Dip Depth To Capture
- Integrated Total Beam Loss
- Maximum Total Beam Loss

### 8.6.5 Conclusion and The Future of the Target Control System

A significant amount of time was spent both building and debugging the original electronics system. This has meant that there wasn't the time to upgrade the prototype electronic system to a full functioning 'turn key' solution before it was installed at RAL in January 2008. The electronics that has been installed at ISIS is functional and has been proven to work reliably, both in the lab and with the limited operation that it has so far seen at ISIS. Clearly these are the two most important considerations. Consequently there is a plan in place to upgrade the electronics over the forthcoming year by moving and rewriting the firmware to run on a 16 bit processor. The 16 bit processor will be both faster and more powerful than the current 8 bit microprocessors that are utilised. The improvement in processing power will allow several key improvements to be made to the system.

The use of 16 bit processors will force the circuitry onto printed circuit boards as there is no possible way that 16 bit processors can be interfaced to veroboard, which is the method currently used for creating all of the circuits. Clearly PCB's offer several advantages over veroboard including increased reliability, a smaller footprint for the circuits, improved signal to noise ratio and the possibility of increasing the operational frequency of the circuits.

It is intended that the whole system should be integrated so that all the microprocessors will communicate with each other over a system bus such as I<sup>2</sup>C. This will make the system more 'intelligent' and with the addition of a suitable interface it will allow the remote control and monitoring of many of the target system parameters from a PC terminal.

With more electromagnetic modeling of the stator and with some minor changes to the drive circuitry it should be possible to improve control over the target shaft movement reducing the damped oscillations and making the transition from park to hold quicker

and smoother. It is not clear at the moment whether implementing a time based feedback system will offer a significant advantage but if it does then time based feedback could in principle be introduced. The collection of additional data from operational use of the target mechanism and its subsequent analysis will assist in deciding whether time based feedback is worth pursuing.

The control system as described in this chapter has been pivotal in turning the target mechanism into a functioning piece of equipment. It has reliably actuated the target for tens of millions of times in a consistently reliable and repeatable manner that has met both the original timing and positioning specification that was defined at the start of the project. The next chapter will demonstrate this by quantifying the performance of the target mechanism.

## Chapter 9

# Operational Performance and Analysis

This chapter discusses the data produced by the target system both in the laboratory and on the system installed on ISIS at the Rutherford Appleton Laboratory. The first opportunity that the target team had to run a target system on ISIS was in November of 2006. This was an actuation limited run and the purpose was to obtain further information on the ISIS beam position and ascertain particle rates from ISIS.

Full running of a completed target system on ISIS did not commence until early 2008. By 2008 many changes had been made to the system since the target test in 2006; mechanical changes had been made to the stator, the electronics had been through significant revision, and the layout of the target system at ISIS had been changed.

### 9.1 ISIS

Some details of ISIS have already been given; a summary of the ISIS facility was given in section 3.4. The evolution of the beam momentum and energy in ISIS can be calculated from a knowledge of the magnetic fields of the steering dipoles within ISIS. A diagram of the result is shown in figure 9.1. This diagram and the derivation of the curves is given in K. Walaron's thesis [56]. It is well known that the radius of the core of the beam shrinks during the acceleration cycle. The amount of shrinkage of the beam halo radius at the point of the target was unknown before the experimental tests in 2006.

The ISIS beam loss monitoring system is provided by thirty-nine 3 m long gas ionisation chambers [108]. These chambers surround the inner circumference of ISIS and are wired to a set of amplifiers that give an output voltage that is proportional to the beam loss from the synchrotron. The output from the beam monitors is dependent upon the beam energy and this is reflected in the output voltage for a given loss. The relationship is

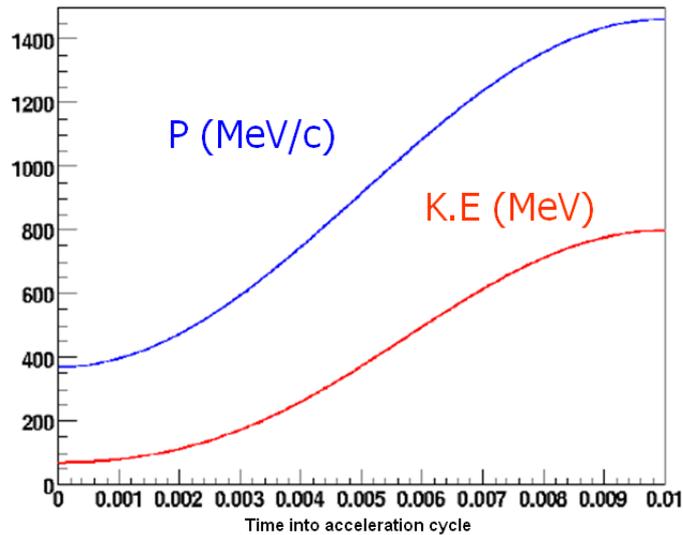


Figure 9.1: Evolution of momentum (red line) and energy (blue line) during an ISIS acceleration cycle [56].

not linear and the beam loss monitors are more sensitive at higher proton energies. It is possible to convert the output voltage from the beam loss monitors to an actual particle loss rate from ISIS after suitable calibration but there is a large error on this of  $\approx 50\%$  [56].

## 9.2 Experimental Setup for Nov 2006

Two shifts were granted in November 2006 during an ISIS maintenance period to actuate the prototype target up to a maximum of one hundred times into the ISIS beam. These shifts were done in collaboration with Glasgow University and colleagues at RAL. The Sheffield University team were concerned with gaining experience with installing the target on ISIS and obtaining experimental information on how the ISIS beam radius changes with respect to time. Glasgow University installed a pair of particle detectors as they were interested in looking at the particle rates coming from the target and how this would translate into the muon rate for MICE.

### 9.2.1 Target

A stator with ceramic bearings was installed on the ISIS beam pipe in sector seven of the synchrotron. At this time the power supply could provide a maximum of 10 A to the target and so the target acceleration was limited to about 14  $g$ . The target supply was placed in the synchrotron about 5 m away from the target. The target controls

were placed in the MICE hall and optical fibres were installed from the target controls to the target optical block and the target power supply in the synchrotron. A total beam loss signal, and a beam intensity signal was provided by ISIS as well as the local beam loss signal from super-period seven, the location of the target mechanism.

### 9.2.2 Detectors

Glasgow university provided and installed two particle detectors into sector seven of the ISIS synchrotron vault as they were primarily interested in looking at the particle rates produced by the MICE target and comparing this with the expected flux for MICE [56]. Both detectors contained a pair of photomultiplier tubes (PMT) coupled to a section of scintillating fibre. One detector was unshielded and the other detector was shielded homogenously with 5 cm of polyethylene. The double PMT design allows coincident signals to be detected thereby greatly reducing the background noise.

### 9.2.3 DAQ

At this point the only digital signals available to the target PC were the target position and time. An analogue target position output from the target electronics and the beam loss signals from ISIS were input to a Tektronix scope for screen capture. The outputs from Glasgow's detectors were connected to a high resolution LeCroy oscilloscope. Both scopes were triggered from the target electronics and then read out from the oscilloscope to separate files using a second PC running Linux. The digitised files from the target PC and the detector PC were then later reconstructed off-line. Figure 9.2 illustrates the experimental setup. The Tektronix oscilloscope sampled the data with a 10  $\mu$ s resolution over a period of 100 ms whilst the LeCroy oscilloscope sampled at a 10 ns resolution over a period of 1 ms.

## 9.3 Experimental Run 2006

The target acceleration was limited by the 10 A power supply so there was not enough acceleration to get the target cleanly into and out of the ISIS beam. To mitigate this ISIS was operated in an on-demand single pulse mode. With only single proton pulses in ISIS then it did not matter how long the target took to get out of the beam. Some data was taken with the target actuating, and some was taken with the target in a static position. In the latter case the target was levitated in a hold position that placed it in the path of the beam when ISIS was pulsed. Tables 9.1 and 9.2 shows the breakdown of the data.

By good fortune the target acceleration was quite closely matched to the beam shrinkage rate. This meant that as the target was accelerated it continually scrapped the beam

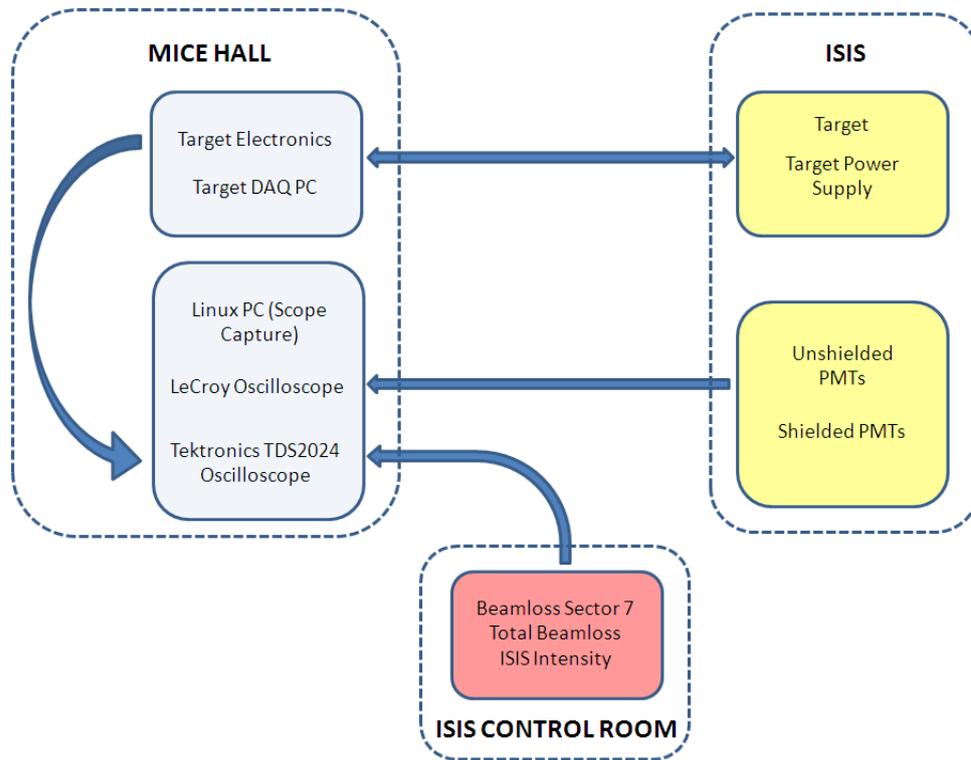


Figure 9.2: Target Test Setup in November 2006

as it shrank generating more data than could have otherwise been expected with so few actuations. The target dip depth was effectively altered by actuating the target earlier with respect to beam injection. This pulls the target trajectory further into the beam. The limited acceleration available to the target did mean that it could not penetrate deeply into the ISIS beam toward the end of an ISIS cycle, the point at which was of interest to MICE, however useful data was obtained.

## 9.4 Analysis of 2006 Data

Two separate analyses were performed with the data obtained from the November 2006 run, one by Glasgow on the particle rate monitors and the second in Sheffield on the beam loss data.

Stationary Position	Number of Data Sets
0 mm	4
-12 mm	3
-13.5 mm	3
-15 mm	3
-16.5 mm	3
-18 mm	1

Table 9.1: This table shows the number of data sets collected with the target held in a stationary position. The distance is given with respect to the target's nominal holding position which was 78 mm from the centre of the beamline. The 0 mm data-sets were used to obtain background beam loss information [109].

Time from ISIS Injection (ms)	Number of Data Sets
-0.045	3
-8.2	3
-12.5	5
-13.0	3
-13.5	16
-14.0	4
-14.5	5

Table 9.2: This table shows the number of data sets collected with the target actuating. The time is given with respect to the injection of protons into ISIS (ms) [109].

#### 9.4.1 Sheffield Analysis

The motivation behind the analysis of the beam loss data was to obtain a profile of the shrinking ISIS beam with respect to time. Both the cross sectional radius as a function of time and the physical position of the ISIS beam within the beam-pipe at the target mechanism's location was unknown before these tests. The information gained from these beam tests was crucial in confirming that the target design permitted a large enough actuation depth to reach the ISIS beam and allowed the target trajectory to be clarified. Much of the original analysis was done by L. Howlett and was presented in a MICE note titled 'Analysis of Target Beam Test Data' in February 2007 [109]. The following discussion is a summary of the analysis done in this report.

Firstly the files in the target PC were matched up to the files stored on the oscilloscope DAQ PC. The matching of the files was necessary as the target trajectory on the target PC gave accurate positional information which was not available on the oscilloscope trace. The target position trace on the oscilloscope was not calibrated.

The beam loss data acquired by the oscilloscopes was noisy so once the files had been matched the data set were then reduced by averaging over 10 consecutive values in the

data-sets. This reduction left the data in bins of  $10 \mu\text{s}$  width. A sinusoidal background was present in the beam loss data and this was removed by using the beam loss data collected with the target in its fully raised position; in this location no beam loss could have been caused by the target because it would be positioned outside of the beam-pipe. A reproduction of the plot of beam loss as a function of time for the target, starting  $14.5 \text{ ms}$  before injection, is shown in figure 9.3 [109].

From these reduced data sets plots of beam loss vs target depth were produced for time slices in  $100 \mu\text{s}$  bins and an exponential was fitted to each plot. The exponential fit shows integrated beam loss and therefore gives an indication of the ISIS beam location as a function of time. An example of such a fitted plot is reproduced in figure 9.4 [109].

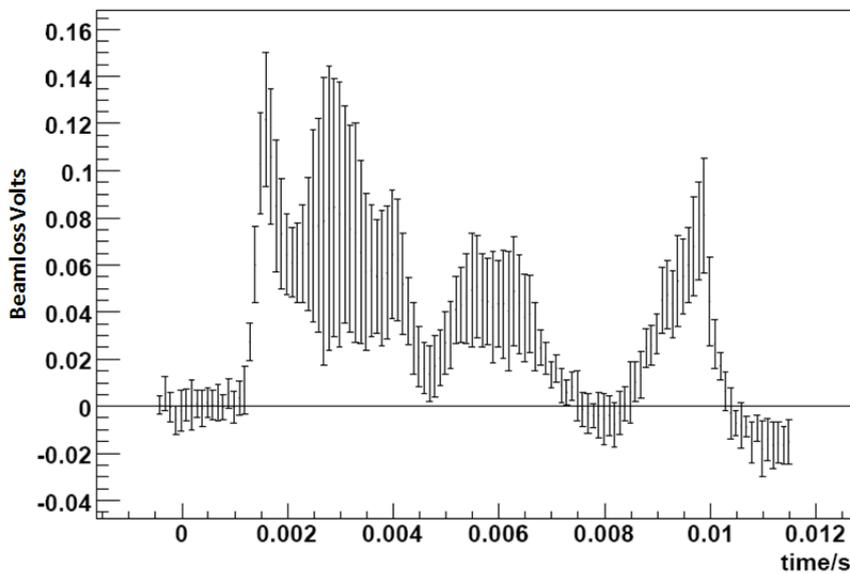


Figure 9.3: An averaged beam loss plot as a function of time for the actuations starting  $14.5 \text{ ms}$  before injection. This time setting corresponds to the deepest penetration achieved into the ISIS beam during the 2006 tests [109].

These fitted plots were finally placed together in a 2D colour plot to give an indication of the beam position with respect to time. This is illustrated in figure 9.5. This graph gives an indication of the position of the ISIS beam edge with respect to time. The horizontal black line represent the maximum depth of the stationary target trajectory and the curved black line shows the earliest actuated target trajectory used to obtain the data. This makes it clear what is interpolation and what is extrapolation. The cut-off at  $0.4 \text{ V}$  represents the maximum recorded beam loss reading in the data. It can be noted that the beam edge starts out with a radius of  $55 \text{ mm}$ , expands within the first couple of  $\text{ms}$  to  $60 \text{ mm}$  before shrinking to a radius of about  $40 \text{ mm}$ . There is a gap in the graph in the period between  $7.5\text{--}8.5 \text{ ms}$  where there is no beam loss as no data was recorded at this time during the limited data set. It is possible that the beam radius was shrinking at a faster rate than the target motion during this period.

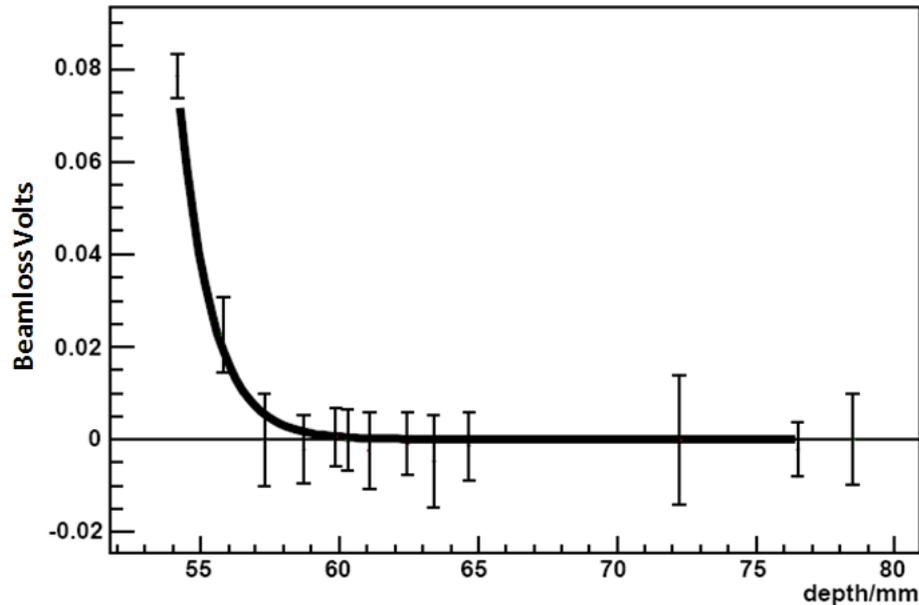


Figure 9.4: A plot showing the beam loss as a function of target position for an individual time bin, along with the exponential fit. This plots shows the time corresponding to 1 ms after injection. A sequence of these plots was created giving the time evolution of the beam loss for different target trajectories [109].

This profile proved to be extremely useful as it assisted in determining the kind of trajectories that are required to enable the target to dip into the beam in the last couple of ms and still be out of the beam before the next ISIS injection. Figure 9.6 shows some quadratic trajectories superimposed upon figure 9.5. These trajectories showed that it was possible to get the target into and out of the ISIS beam whilst meeting these requirements with an acceleration of  $80 g$ . This was well inside the original estimate of  $100 g$ ; this is fortunate as the power dissipation in the stator is significantly lower at these reduced accelerations. At lower accelerations the target takes longer to actuate but even accounting for this it can be shown that the energy deposited in the stator still scales at  $E \propto I^{3/2}$  if a quadratic trajectory is assumed.

#### 9.4.2 Glasgow Analysis

The Glasgow team used the particle production data from running the target in November 2006 to validate several simulations of the particle production rates from the target. This was an important step as these simulations calculate the expected particle rate into the ISIS beamline; this later reflects in the observed muon flux for MICE.

The work by the Glasgow team was covered in K. Walaron's Thesis [56], a conference

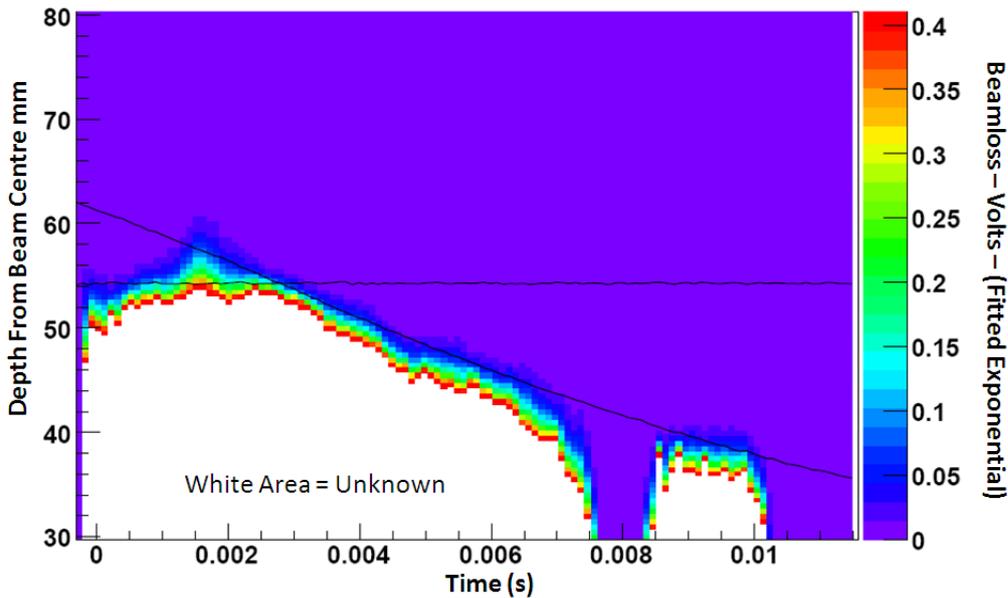


Figure 9.5: Total beam loss as a function of position and time as calculated using exponential fits. This graph gives an indication of the position of the ISIS beam edge with respect to time. The horizontal black line represent the maximum depth of the stationary target trajectory and the curved black line shows the earliest actuated target trajectory used to obtain the data. This makes it clear what is interpolation and what is extrapolation. The cut-off at 0.4 V represents the maximum recorded beam loss reading in the data [109].

paper submitted to NuFact07 [110] and a talk given to the MICE collaboration [111] and further details can be found from these sources. The following is a brief summary of their analysis.

Simulations of the shielded and unshielded detectors and their geometrical relationship to the ISIS beam and MICE target were modelled using GEANT4 and MARS; these are independent particle tracking codes. For every  $10^7$  protons incident upon the MICE target the number of particles incident upon the detectors was calculated. The efficiency of the detectors was also calculated given that a hit would only be registered if two PMT's fired in coincidence. Table 9.3 shows the result of the MARS particle simulation as given in K. Walaron's thesis. [56]

These codes were then compared to the data collected from the real detectors. First a linear relationship between the number of particles recorded in the detectors and the measured beam loss was established; this was then converted to particle rates using the relationship between the beam loss monitor output and particle loss from ISIS ( $3.5 \times 10^{-14} \text{ Vs}^{-1}\text{proton}^{-1}$  at a proton energy of 780 MeV [56]). When scaled this led to an agreement between the simulated data and the real data. The results of this

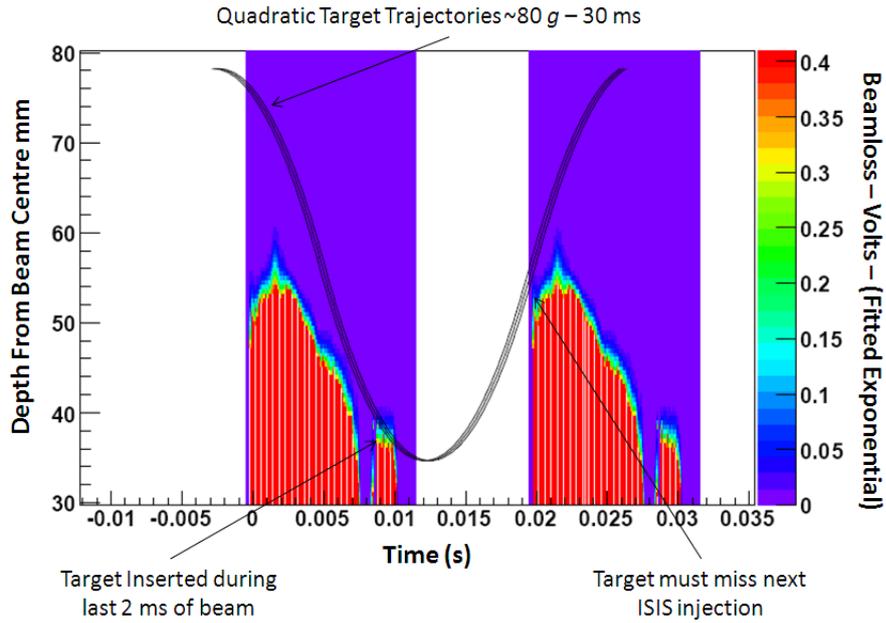


Figure 9.6: A figure showing how a 30 ms 80 g quadratic target trajectory can be made to fit into the ISIS spill structure (50 Hz) to intercept the beam during the last 2 ms of a spill. The target must not touch the beam during the next injection [109].

Particle Species	Unshielded Detector	Shielded Detector
Protons	251	226
Neutrons	446	405
$\pi^+$	6	7
$\pi^-$	2	2
$\mu^+$	4	4
$\mu^-$	1	1
$e^+$	9	10

Table 9.3: Un-normalised rates for both shielded and unshielded detectors using MARS for  $1 \times 10^7$  protons on target [56].

analysis are shown in table 9.4. This result allows an estimation of the produced muon rate compared to what is required for MICE ( $600 \mu^+$  per spill). Table 9.5 shows the original assumptions and calculations made for the target/ISIS interaction as laid out in the MICE technical reference document [65].

Further simulation studies by Walaron [56] showed that the number of  $\pi^+$  that would be accepted into the MICE beamline per  $1.7 \times 10^{12}$  incident protons is  $\approx 6 \times 10^5$  per ms. Earlier simulations by T. Roberts [65] showed that about 1 in 3,900 of these accepted pions will become ‘good’  $\mu^+$  (a muon that makes it through both the MICE cooling

	Singles/(p.o.t.) (Unshielded) ( $\times 10^{-8}$ )	Singles/(p.o.t.) (Shielded) ( $\times 10^{-8}$ )
MARS	1.70 ( $\pm 0.10$ )	1.52 ( $\pm 0.096$ )
GEANT4	2.47 ( $\pm 0.12$ )	1.61 ( $\pm 0.010$ )
DATA	1.59 $\pm 0.24$ (stat) $\pm 0.81$ (sys)	1.29 $\pm 0.22$ (stat) $\pm 0.65$ (sys)

Table 9.4: The calculated number of non-neutron singles into the 1 cm<sup>2</sup> detector acceptance compared to the observed data from the November 2006 target tests [56].

Target material	Titanium
Target height	10 mm
Target height intercepting ISIS proton beam	2 mm
Target width	1 mm
Target length parallel to beam	10 mm
ISIS beam area at target	4400 mm <sup>2</sup>
Target Insertions per second	1
Target duration in beam during good RF	1 ms
ISIS protons per bunch	$2.5 \times 10^{13}$
Bunch Rate at target	1.5 MHz
Protons intercepting the target per ms	$1.7 \times 10^{12}$
Fraction of protons taken on a spill	6.8 %
Fraction of protons taken on average (50 Hz)	0.14 %

Table 9.5: Simulation parameters used in initial modeling of the target ISIS beam. The proton interception rate was determined by the need to obtain 600 good muons through MICE per spill [65].

channel and the detectors so an emittance measurement can be made on it) for MICE giving a muon rate into MICE of 153  $\mu^+$  per spill for  $1.7 \times 10^{12}$  protons on target. Comparing this to the number of intercepted protons in table 9.4 the estimated good muon rate from the target at the level of beam loss recorded during the 2006 tests would be 1 muon per spill. The amount of beam loss produced in this test is inconsistent with the good muon rates required by MICE.

## 9.5 Experimental Setup for 2008

A stator was installed on the ISIS beam line during the synchrotron shutdown early in 2008. The target electronics were installed in the MICE local control room and the power supply electronics were installed on the outside of the synchrotron wall as close to the target as possible. See section 7.10 for more details. This stator had been run in the laboratory for several million pulses before being installed onto the ISIS beamline. When installed onto the beamline a brand new shaft and bearings were used.

A second complete stator system was also installed in a separate building. This system was nominally supposed to run three times as many pulses as the one in the synchrotron. Primarily this was done to monitor the wear rate on the DLC coated bearing; it also was to serve as an early warning system for other potential problems. The second stator system had been assembled quickly and it was not been possible to test this stator fully before it was sent to RAL. The second stator failed catastrophically within a few hundred actuations after installation at RAL due to coil insulation breakdown. The insulation problem has now been solved with improved quality assurance procedures on the coil manufacture and the through the addition of kapton insulation. The stator currently installed on the ISIS beamline does not have any kapton insulation. More details on this failure can be found in chapter 10.

Shortly after installation of the target the first six MICE quadrupoles were made functional as well as both of the dipoles. The decay solenoid has remained non-operational up to the time of writing this thesis although it is hoped that this will come online during the autumn of 2008. To allow monitoring of particle rates a pair of scintillation counters GVA1 and GVA2 were installed in the MICE beamline along with some checkerboard scintillation counters BC1 and BC2. The GVA scintillation counters give particle rates and can help determine particle type through time of flight which is useful for discriminating protons from pions; however they do not give any positional information. The BC counters can give low resolution positional information and were intended to assist beam alignment.

In July the Cherenkov counters were installed along with the TOF0 counters and the calorimeter. The TOF counter has been placed in front of the third quadrupole triplet and the calorimeter has been placed at the end of the beamline. As the calorimeter must be the last detector in the experiment its position will move further back as more equipment is added to MICE.

As the Cherenkov, TOF and KLOE detectors were not installed until July only information from the GVA counters have been used in the target performance analysis presented in this thesis. A diagram illustrating the position of the GVA counters in relation to the MICE beamline is shown in figure 9.7 [112].

## 9.6 Running of Primary System at RAL

Despite the failure of the secondary demonstrator target system the primary target system has been run cautiously at RAL during periods when ISIS has not been providing particles for other experiments. The beam loss caused by the target was also initially limited to 50 mV, a value consistent with that with which the target had been run during 2006. This level was set as a baseline so that the the areas surrounding the target in the synchrotron could be monitored for activation. The target has been operated at higher beam loss levels than this for short periods of time to aid the understanding of beam loss with particle production rates for MICE. The beam loss rate will need to be

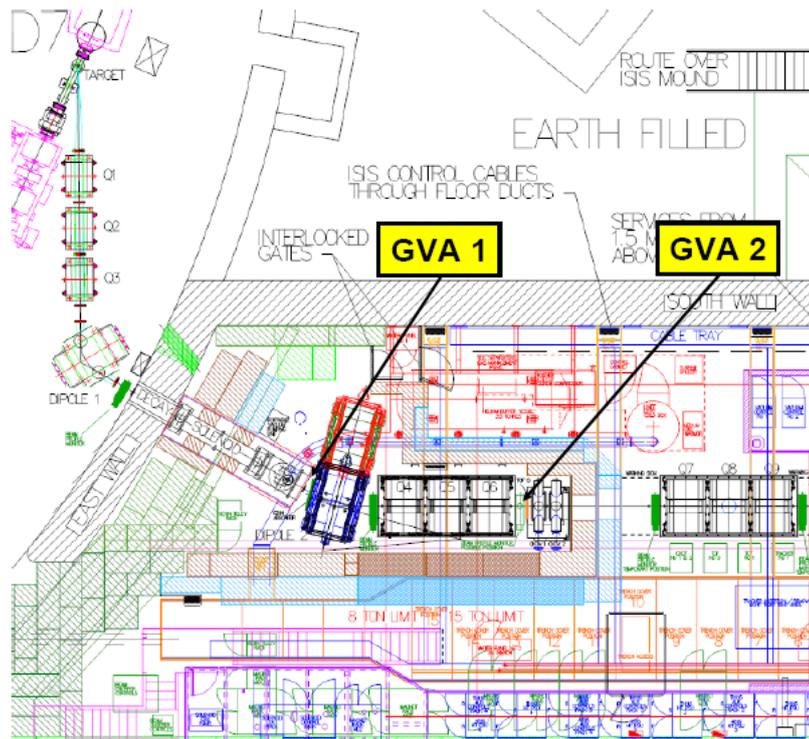


Figure 9.7: Diagram illustrating the position of counters GVA1 and GVA2 in the MICE beamline [112].

increased in due course to provide MICE with the necessary muon rates. However this is unlikely to happen until the activation effects of the target are fully understood and mitigated.

Table 9.6 shows the run dates from installation of the target up until the end of July 2008. There was no DAQ for the ISIS beam loss data until the run on the 29th March as the digitisation circuit board needed redesigning to enable full digitisation of the beam loss signal.

## 9.7 Analysis

This analysis section will be broken down into two sections. Firstly, analysis was undertaken to look at the mechanical performance of the actuator. For this, data from the target on the ISIS beam-pipe and data obtained from laboratory testing was used. Secondly an analysis of the actuator's performance in terms of beam loss on ISIS and particle production for the MICE beamline is considered.

Month	Day	Data-set	No. of Pulses
March	14	data140308.001	17
	15	data150308.002	8
		data150308.003	3654
	29	data290308.001	804
	30	data300308.001	1529
April	2	data020408.001	1997
	4	data040408.001	5464
	5	data050408.001	411
May	14	data140508.001	90
	15	data150508.001	3788
	16	data160508.001	4642
	20	data200508.001	804
June	2	data020608.001	3257
	3	data030608.001	2315
	8	data080608.001	3142
	20	data200608.001	3919
	21	data210608.001	10002
		data210608.002	10002
data210608.003	1872		
July	11	data110708.001	4089
	12	data120708.001	692
	14	data140708.001	5645
	15	data150708.001	3832
	18	data180708.001	9987
	19	data190708.001	10001
		data190708.002	7992
	20	data200708.001	10001
	20	data200708.002	3827
Total			113407

Table 9.6: Table showing target operation on the ISIS beamline from installation in Jan 2008 to July 2008. Generally the runs have been short as the MICE target has been operated cautiously. This cautious approach will continue until a second ‘demonstrator’ target system has been installed and operated in building R78.

### 9.7.1 Mechanical Performance

The mechanical performance of the stator is an important consideration as both long term stability and reliability of the stator will be necessary. Mechanical stability for the purposes of this thesis is defined as reproducibility of the target’s motion. The reliability of the target system can be considered from two perspectives. The first is component failure, the second is the reliability of the target capture at the end of an actuation. Target capture is as important as component reliability because failure to

catch the target results in it being dropped into the ISIS beam and the ISIS beam will trip off. If this occurs frequently then MICE may not be allowed to run due to the impact that this would have on the other experiments running off ISIS. Finally, the operational temperature of the stator is considered.

### 9.7.2 Mechanical Stability

The mechanical stability of the target actuator can be monitored by several means. The DAQ monitors the position of the target with respect to time and this allows several parameters to be monitored by the subsequent analysis. The target is usually set to a particular actuation depth and left to run unless more or less beam loss is required. In particular, the lab setup left the target running at a 44 mm dip depth. This dip depth is deeper than that typically used by the target at RAL - as will be shown the actuation depth used at RAL is typically around 35 mm, the actual depth being dependent upon the allowed beam loss.

The DAQ allows the following parameters to be monitored for long term stability:

- Actuation depth
- Time from the start of actuation to apex of trajectory
- Acceleration

Statistical variations on these parameters from pulse to pulse allows the repeatability of the target system to be quantified. This proves useful in terms of identifying any real changes in the target behaviour over time either due to component failure or a change in operating conditions. Quantifying the variation in the actuations from pulse to pulse can also help to isolate the cause of any drift in the system. For example, it is expected that the ISIS beam will move about over time and if this occurs it will cause changes in the ISIS beam losses, this will impact upon the particle production rate for MICE. Clearly it will be useful to be able to isolate and identify the movement of the ISIS beam separately from the natural variation in target behaviour. Obtaining statistics on the repeatability of the target motion will help to achieve this goal.

In this thesis an example of the target's repeatability of motion will be illustrated using two sets of data. The first set of data comes from a five day continuous run in the laboratory where the target actuated  $\approx 426,000$  times at 1 Hz. This data is used because no system changes were made during this run and so the large number of actuations gives a good statistical sample from which to create these plots. This data was taken from the 6th - 10th December 2007. The second set of data is taken from the same stator after it had been installed on the ISIS beamline. Because running of the target system on the beamline has been restricted the amount of data available for analysis is much smaller. This has been further impacted by the fact that the target has

been rarely run under conditions where the system has not had its parameters changed for a significant number of pulses. Time spent operating the target system has often involved changing target parameters such as dip depth and timing in an attempt to make measurements on the beam loss and the subsequent particle rates in the MICE beam line. The second data set therefore is restricted to a set of  $\approx 22,000$  actuations performed by the target on 21st and 22nd of June, 2008 at 0.4 Hz. This represents  $\approx 20\%$  of the total data available at the time of writing this thesis.

### 9.7.3 Actuation Depth

The actuation depth when running on ISIS is less than the target had been operating at in the laboratory. In the laboratory the target actuation depth was set at  $\approx 43$  mm. When operating on ISIS beam loss is initially observed at a target actuation depth of about 35 mm. The target cannot be taken more than a few mm deeper than this before the beam loss becomes so high that the ISIS machine protection trips ISIS off. Therefore it is difficult to make absolute comparisons between the two sets of data and only variation within the data sets are considered.

Figures 9.8 and 9.9 show histograms of the variation in actuation depth over the sample data. The first set shows the five day run in the laboratory and the second histogram shows the target run on ISIS. (All plots relating to this data are shown as pairs in this order). A Gaussian has been fitted to the first histogram in figure 9.8 where a better fit is expected. It is clear that there is a slight tail on figure 9.9 which is not understood. Both of these plots show a FWHM of  $\approx 0.3$  mm. The original specification was to have a target that could be repeatedly placed to within 0.5 mm (although there was no justification for this figure) so the device is performing better than originally specified. It is not clear yet as to whether the amount of jitter seen is problematic or not. More studies on the relationship between the actuation depth and the beam loss would be required.

Figures 9.10 and 9.11 shows the time evolution of the dip depth. In figure 9.10 each point represents the mean of one hour's data whereas in figure 9.11 each plotted point represents the mean of six minute's worth of data. (All the pairs of plots are shown with the same time division.) In both cases the dip depth is stable. There is some evidence of a fractional decrease in the dip depth during the 15 hours of running on ISIS on plot 9.11 but this is much smaller than the FWHM and is not inconsistent with the data shown in plot 9.10 which is of course plotted over a much longer timescale.

Figure 9.11 shows two distinct populations at 37.6 mm and 37.7 mm. Each point on this plot is an average over 140 actuations. The position system has a resolution of 0.15 mm and so there is quantisation of the actuation depth to this resolution. The dual population will be an artifact of this quantisation effect. This is not seen in the plot 9.10 as the larger number of samples per point (3,600) effectively washes this quantisation out.

### 9.7.4 Acceleration

Figures 9.12 and 9.13 shows histograms of the variation in acceleration over the sample data and figures 9.14 and 9.15 shows the time evolution. The acceleration is obtained by taking the second derivative of the fitted function to the acceleration curve. The first plot shows an acceleration of  $826 \text{ ms}^{-2}$  with a FWHM of  $\approx 15 \text{ ms}^{-2}$  whereas the second plot shows an acceleration of  $818 \text{ ms}^{-2}$  with a FWHM of  $\approx 12 \text{ ms}^{-2}$ . The slight difference in the magnitude of the acceleration can be attributed to a number of factors; the difference in cable lengths and operating voltages between the laboratory and in ISIS and the difference in operating frequency. In the laboratory the stator was run at 1 Hz but with fairly cold cooling water typically  $\approx 15^\circ\text{C}$ . On ISIS the actuator has been run at 0.4 Hz but the cooling water was nearer  $30^\circ\text{C}$ . As will be demonstrated later in this chapter the target acceleration is very sensitive to the temperature of the stator. The low impedance of the coils within the stator means that connection resistances and the cable resistance can have a measurable effect on the coil current for a given driving voltage. In the laboratory a cable length of 40 m ( $2.5 \text{ mm}^2$ ) was used with a PSU voltage of 240 V DC. The cable length in ISIS is 25 m ( $4.0 \text{ mm}^2$ ) and the PSU voltage is 220 V DC. A small discrepancy in the acceleration would be expected. The figure of  $818 \text{ ms}^{-2}$  represents an acceleration of  $83 g$ . This exceeds the estimated requirement of  $80 g$  and as will be shown later allows the target to operate parasitically on ISIS.

The small difference in the FWHM can also be seen in the time evolution plots (figures 9.14 and 9.15). The five day run in the laboratory shows more deviation than the fifteen hour run. The slightly larger FWHM in plot 9.12 could be due to larger temperature fluctuations seen in the laboratory which is likely to be observed with the target running over extended periods of time. Unfortunately there are no temperature logs available so any correlation between cooling water/laboratory temperature and acceleration cannot be established at present.

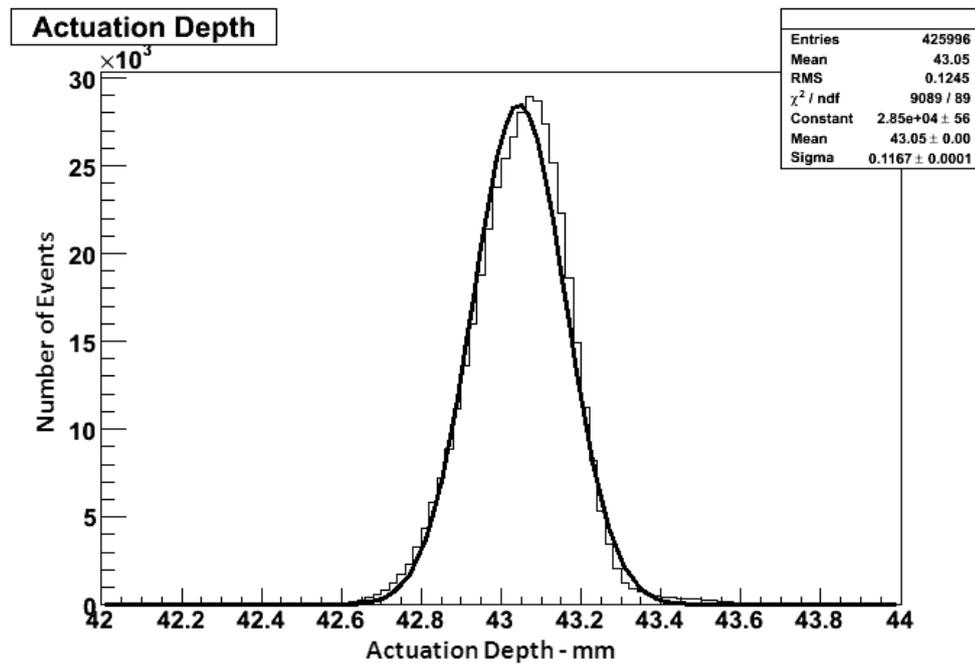


Figure 9.8: A histogram of the actuation depth for the stator running in the laboratory over five days at 1 Hz.

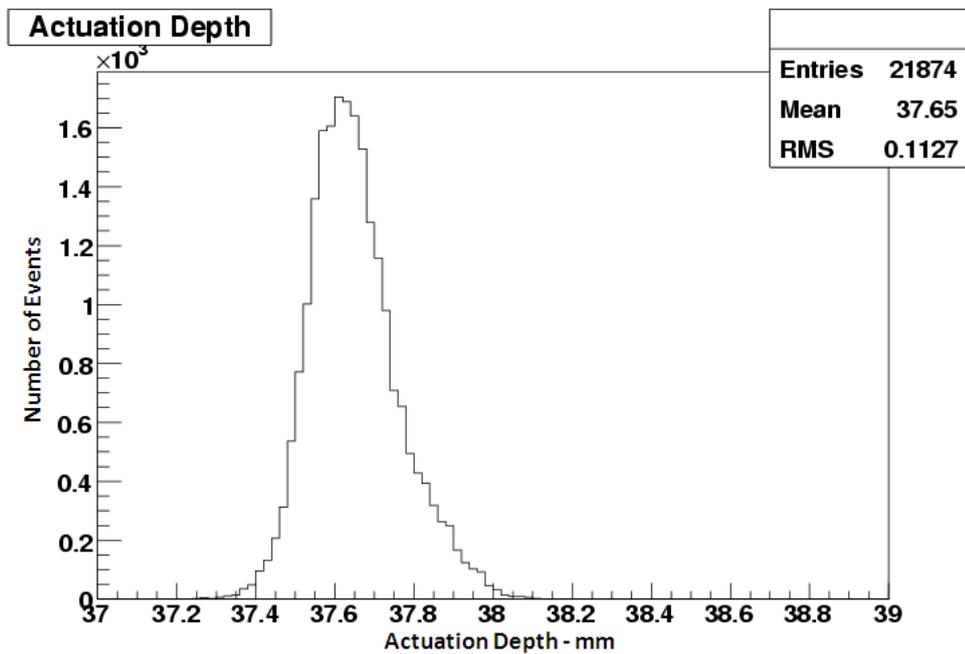


Figure 9.9: A histogram of the actuation depth for the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz.

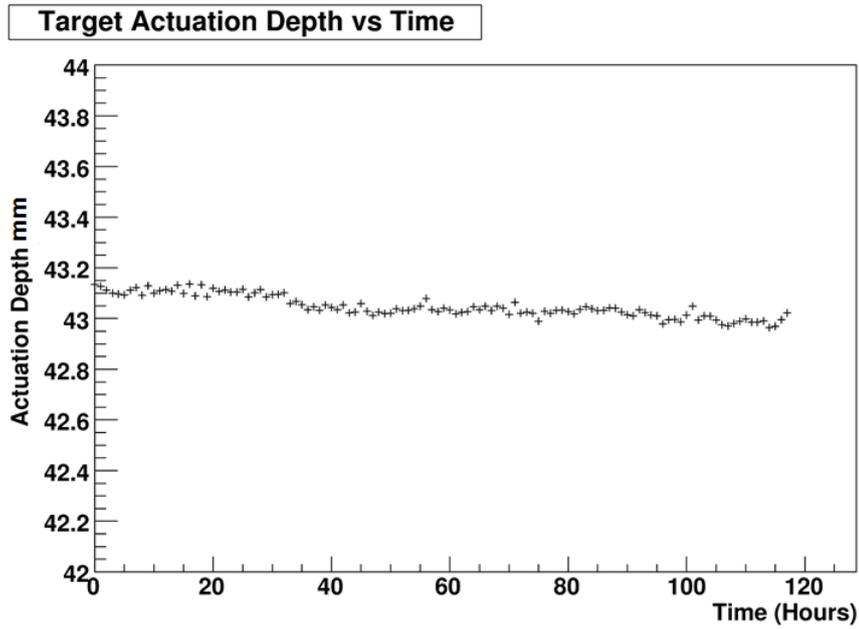


Figure 9.10: The time evolution of the actuation depth for the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points).

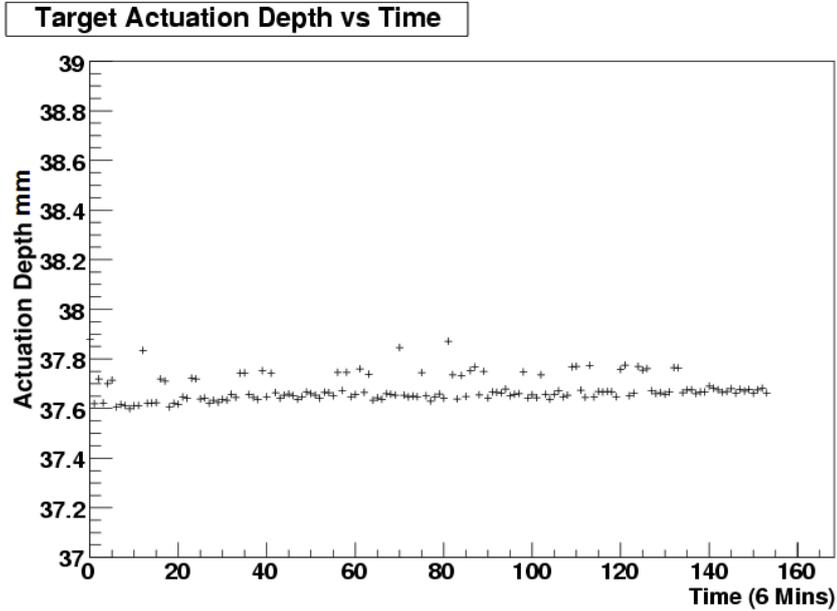


Figure 9.11: The time evolution of the actuation depth for the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points).

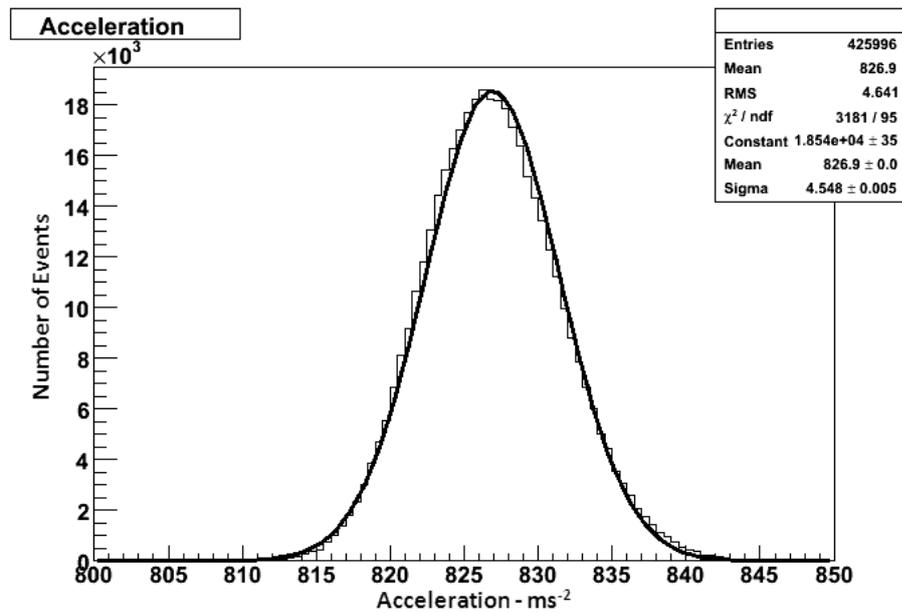


Figure 9.12: A histogram of the acceleration achieved for the stator running in the laboratory over five days at 1 Hz.

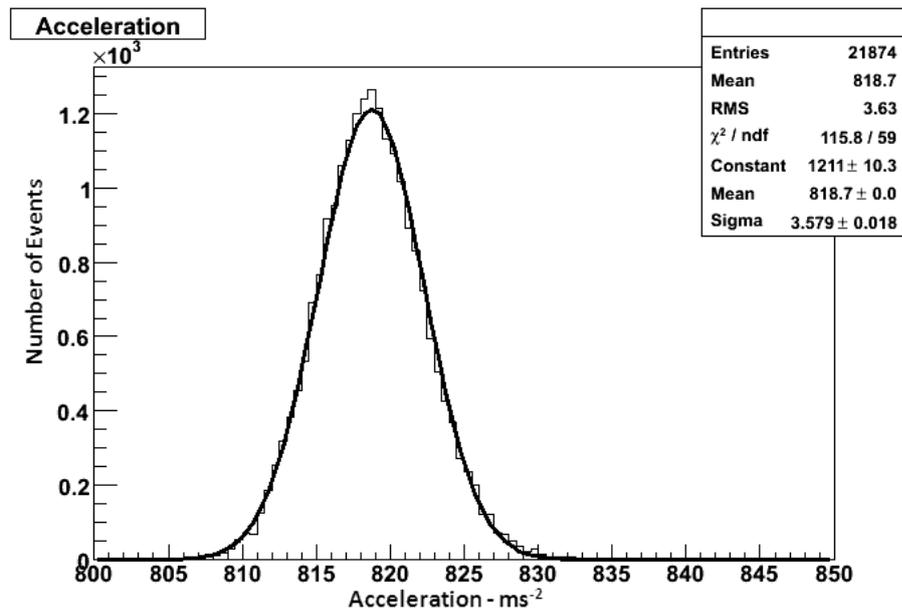


Figure 9.13: A histogram of the acceleration achieved for the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz.

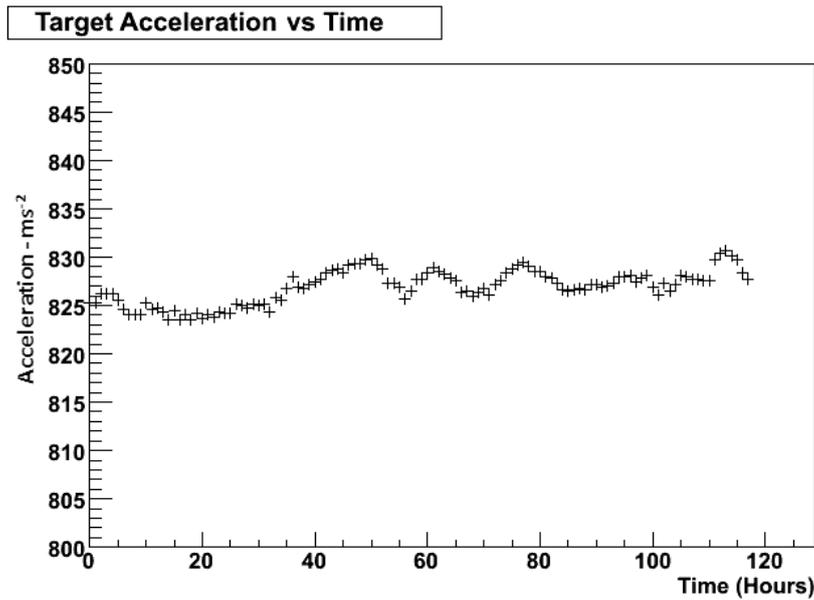


Figure 9.14: The time evolution in the acceleration for the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points).

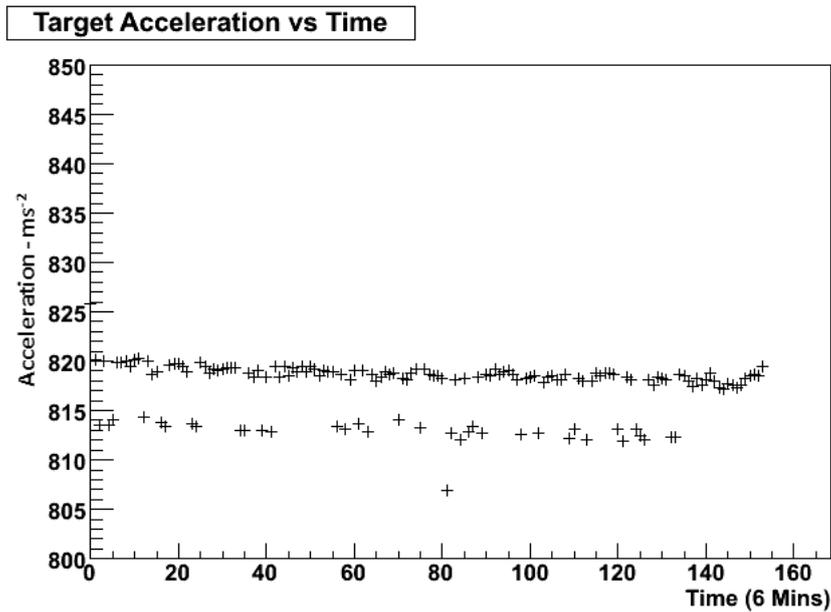


Figure 9.15: The time evolution of the acceleration for the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points).

### 9.7.5 Time Jitter

Figure 9.16 and 9.17 shows histograms of the variation in the time to minimum dip depth over the sample data. Figures 9.18 and 9.19 once again show the time evolution. The data in the plots from the laboratory show the time it takes from the target going into actuate mode until it reaches the minimum point of its trajectory. The plots shown from running the target on ISIS have to be treated a little more carefully as the time recorded is the time from when the DAQ started recording and not from the point where the target starts actuating. This is done because it allows an estimation of the delay setting from the trigger signal. Nonetheless the five day run histogram plots shows a FWHM time jitter of  $\approx 100 \mu\text{s}$  and the shorter run on ISIS shows a time jitter of  $\approx 50 \mu\text{s}$ . Looking at the time evolution graphs it is apparent that there was more drift in the five day run than in the shorter run on ISIS therefore explaining the wider FWHM on the longer run. It is not currently known what effect, if any, the target time jitter has on the beamloss produced from ISIS. The target intercepts the last 2 ms of the ISIS spill so this jitter represents  $\approx 5 \%$  of this time. The effect of this time jitter on the ISIS beamloss will require additional study when better statistics have been obtained.

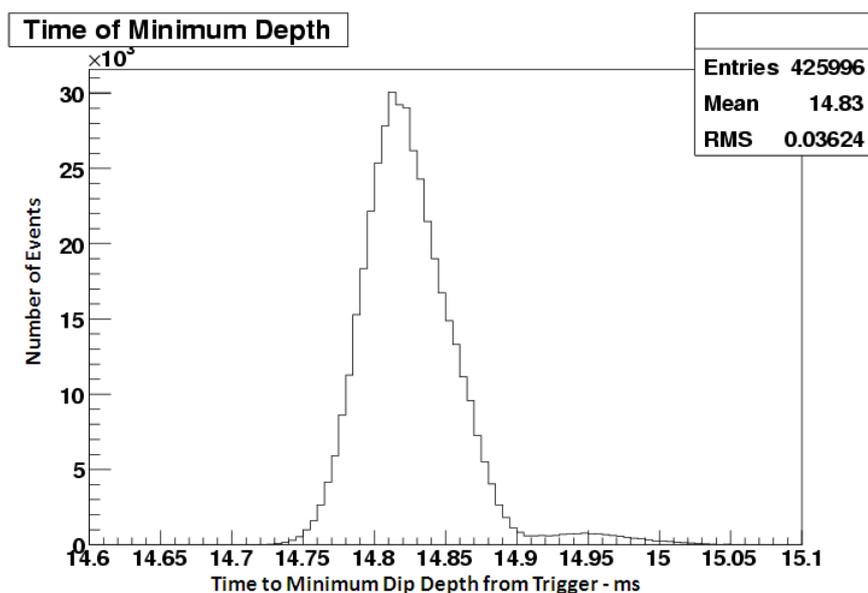


Figure 9.16: A histogram of the time it takes for the shaft to reach its minimum position. This data was taken with the stator running in the laboratory over five days at 1 Hz.

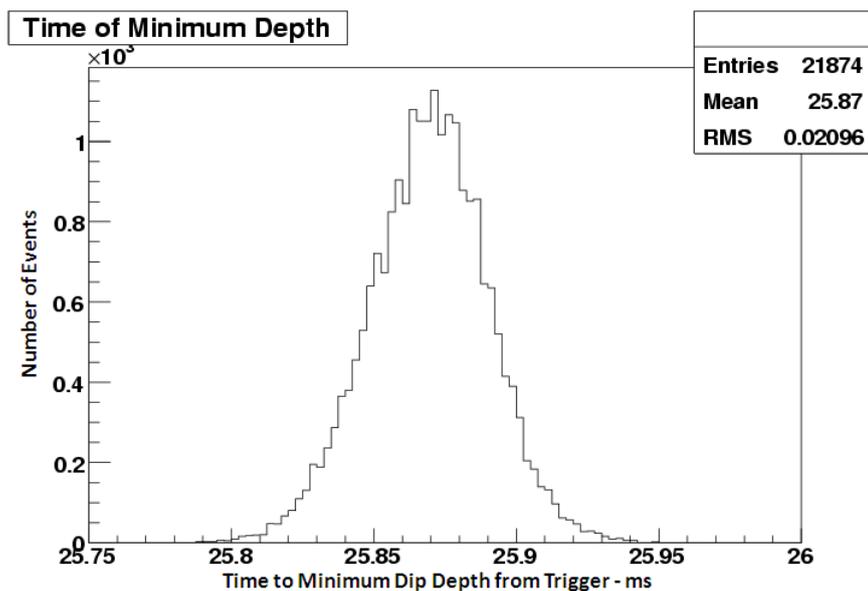


Figure 9.17: A histogram of the time it takes for the shaft to reach its minimum position. This data was taken with the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz.

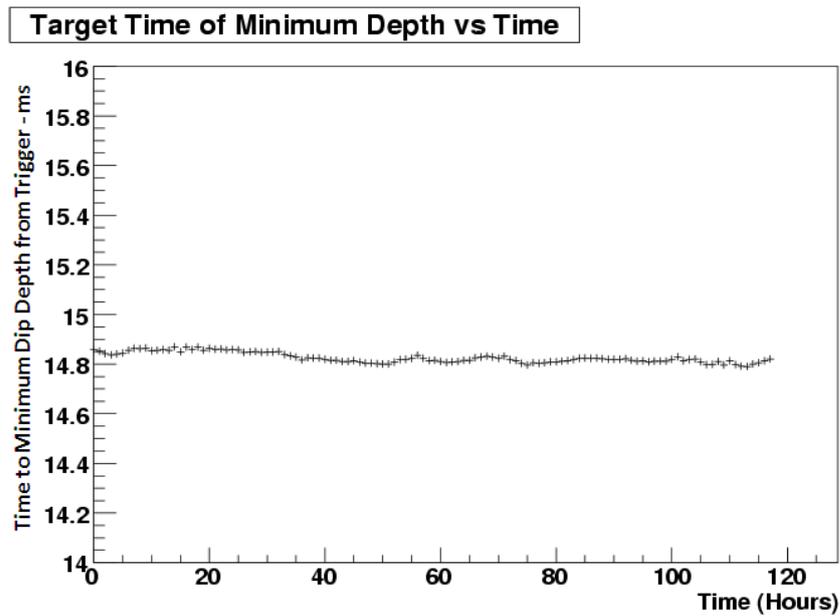


Figure 9.18: The time evolution of the time to minimum dip depth. This plot shows data from the stator running in the laboratory over five days at 1 Hz. Each data point represents the average value over 1 hour of running (3,600 points).

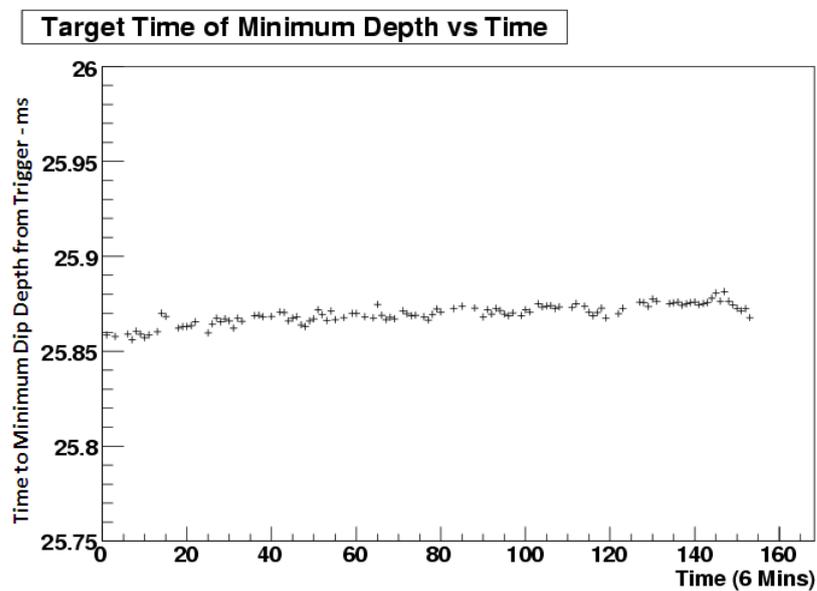


Figure 9.19: The time evolution of the time to minimum dip depth. This data was taken with the stator running on ISIS. This plot shows  $\approx 15$  hours running at 0.4 Hz. Each data point represents the mean value over 6 minutes of data (144 points).

### 9.7.6 Capture Reliability

A piece of software was written that monitors the time stamps added to each event created in the DAQ. This software was written to locate target drops that are a result of the controller failing to capture the target after an actuation. If the controller fails to capture the target the target shaft will drop on to the stop. The controller will then automatically try to pick the target shaft up and resume operation. The process of picking the target up is relatively slow as it typically takes 10 seconds and so any target drops can be recognised by a time gap in the time stamps in the events created by the DAQ stream.

During the development of the target system the controller algorithms have been improved to the point where target capture failures that are not due to an electrical or mechanical problem are extremely rare. Typically the target could be run continuously for up to a week at 1 Hz without a single target drop. By combining runs it is apparent that the capture failure is now less than 1 in  $10^6$ . In the data sets presented here there have been no capture failures in either the laboratory data set or the data sets from operation on ISIS.

Capture failures are most common when initially setting an actuator up. Incorrect alignment and calibration of the optics, connecting the coil phases incorrectly to the power supply and mechanical issues have been the greatest cause of capture errors.

### 9.7.7 Operational Temperature of the Stator

The internal temperature of the stator is important as it directly impacts upon the resistance of the coils and the operating temperature of the magnets. As the coils' resistance goes up the acceleration for a given voltage across the stator goes down. It is known that the permanent magnets on the shaft are adversely affected by their temperature. The magnetisation of the magnets is temperature dependent and for NdFeB magnets radiation damage is also accelerated at higher temperatures [95]. These last two problems are not currently understood and for this reason it is important that the operational temperature of the stator is both monitored and minimised.

The stator is a complex device and any accurate thermal modeling of it would be difficult. It is intended that the stator will be thermally modelled in the future as such a model may assist in implementing improvements to its thermal behaviour. However, at present the thermal data on the stator is obtained by monitoring the thermocouples that are located between the coils within the stator.

The addition of kapton insulation between the stator's coils meant that the thermocouples could no longer be placed in direct contact with the coils. Instead they now have to be placed in direct contact with the copper shims; this has to be done for electrical safety. Unfortunately this does leave a kapton shim between the thermocouple and the coil. During stator operation there is a temperature gradient across the kapton shims

which, with the new position of the thermocouples, means that they are no longer measuring the true coil temperature. Earlier tests, using a stator with thermocouples in contact with the shims and with additional thermocouples in contact with the coils show that the temperature difference across the kapton shims is minimal and that the temperature read by the thermocouples can be taken as a good indicator of the temperature of the coils within the stator. This is illustrated in figure 9.20. Of course if higher accuracy is needed an extrapolation between the measured temperature and the actual coil temperature could be made.

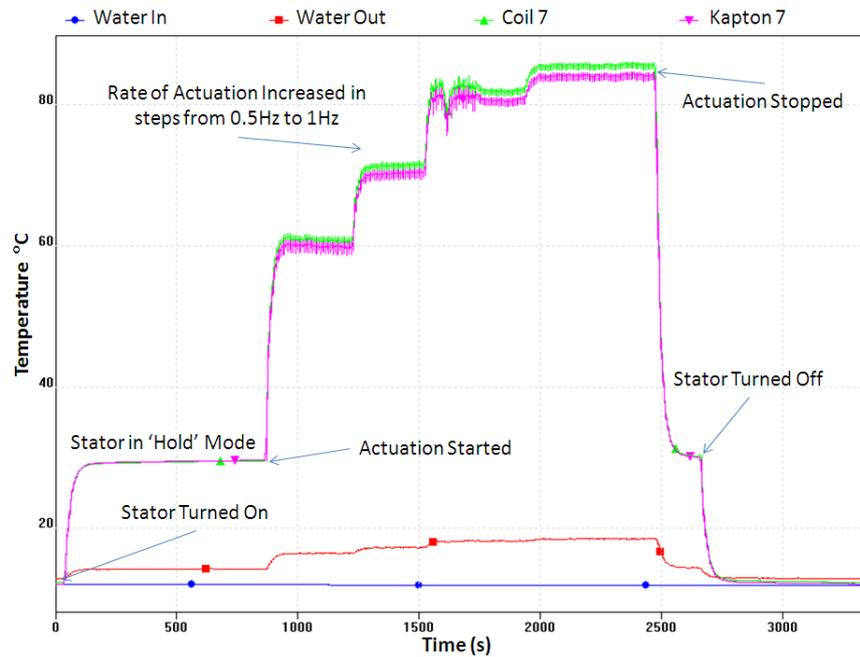


Figure 9.20: This temperature profile shows the coil temperature of over 1 hours running from power up to power down. Note how the temperature monitored directly on the coil (blue - Coil 7) does not differ significantly from the measured on the other side of the kapton shim (pink - Kapton 7). The microstructure seen as noise during actuation is caused by the stator cooling between actuations.

The change in the resistance of the coils as the stator warms up has been found to be an important effect as it notably affect the acceleration achieved for a given driving voltage. The resistivity of a metal is approximated by:

$$\rho = \rho_0[1 + \alpha(T - T_0)]$$

The temperature of the stator coils in hold mode is shown in figure 9.20 to be approximately 30°C. The temperature during operation rises to approximately 84°C, giving a temperature difference of 54°C. Now  $\alpha = 3.9 \times 10^{-3} \text{ K}^{-1}$  at a temperature

of 20°C [113]. Therefore the resistivity of copper will be increased by  $\approx 20\%$  for a 54°C temperature rise from 30°C.

To a first approximation the heat deposited into the stator can be considered as being constant when the stator is actuating and so the temperature rise will be exponential. The current through the coils will be inversely proportional to the temperature and the acceleration should be proportional to the current. This inverse relationship between the temperature and the acceleration implies that as the temperature of the stator rises then the acceleration should drop exponentially. Figure 9.21 shows a plot of the acceleration vs time as the stator warms up. The data used for obtaining the acceleration was taken from a stator built in June 2008. As can be seen an exponential fit can be placed through the data. Taking the raw figures from the plot the acceleration started at  $850 \text{ ms}^{-2}$  and reduced to  $710 \text{ ms}^{-2}$ . This represents a reduction in acceleration of  $\approx 17\%$  as expected. (An increase of 21% is equivalent to a reduction of 17%)

It may be noted that the acceleration achieved with this stator is less than that quoted for the stator that has been installed on the ISIS beamline. The biggest difference between these stators is that the stator installed on the ISIS beamline does not have any kapton insulation between its coils. As may be recalled, the kapton not only electrically insulates the stator it also thermally insulates the stator and so for a given power supply voltage the kapton insulated stator will run hotter and so the acceleration will be lower. Further analysis is required to determine whether the use of Kapton shims would be acceptable for a stator that is to be installed in ISIS.

## 9.8 Particle Production

As discussed earlier on in this chapter the particle losses from ISIS are monitored by the beam loss monitors that surround the inner circumference of ISIS. At present the quality of the beam loss signals that have been recorded by the target DAQ has been poor, the signal has been noisy and this has made it difficult to make accurate quantitative measurements. The signal recorded by the target DAQ has been the total beam loss, the summation of the losses from all 39 beam loss channels. The fact that this is a summation beam loss makes it difficult to extract the beam loss that is caused solely by the target and to establish where the beam loss is occurring. The total beam loss signal is predominantly shaped by the injection and extraction losses from ISIS.

There is currently a significant amount of work being undertaken to improve the beam loss monitoring on the target DAQ. This will include splitting the beam loss signal such that the beam loss monitors local to the target can have their signals individually recorded. Effort is also being made to improve the signal quality. When this is complete a more thorough analysis of these signals will be possible. When the new DAQ system has been installed the aim will be to understand the beam loss signals from the target area in terms of the actual rate of particle losses from ISIS. This should then correlate with the particle rates seen in MICE.

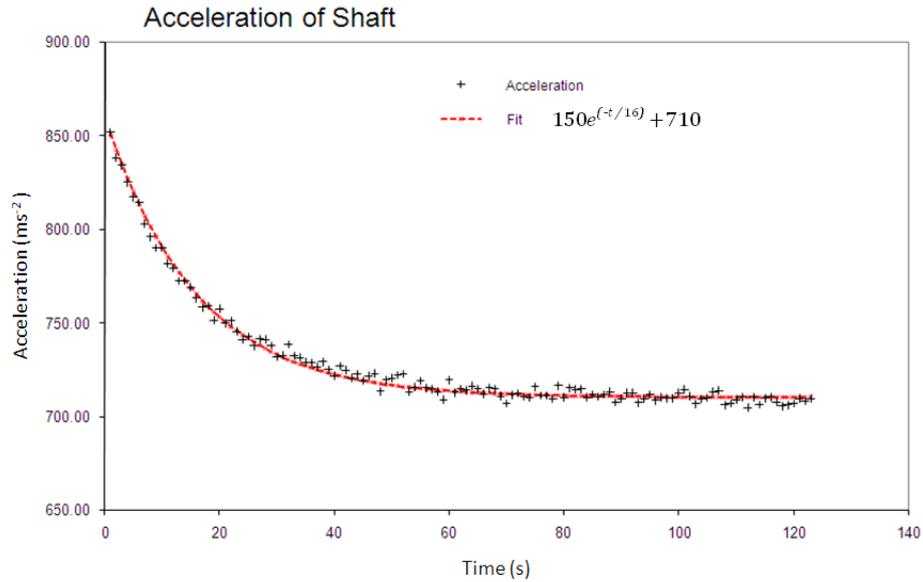


Figure 9.21: The plot shows how the acceleration of the target shaft changes as the stator warms up during the first couple of minutes of operation. The current through the coils is linear with temperature and the acceleration is directly proportional to the current through the stator coils. The heating of the stator will be exponential. This means that the acceleration should change exponentially with time and is confirmed by the fit. The graph is for a stator with kapton insulated coils running at 1 Hz.

Despite the problems with the beam loss signals, the beam loss monitors have proven useful for making relative comparisons of beam loss and with the aid of the GVA counters on MICE some quantitative data on the particles going through the MICE beamline has been obtained. The information gleaned from these measurements will be discussed in this section.

### 9.8.1 ISIS Beam Loss and Actuation Depth

This section will focus on demonstrating the relationship between the total beam loss signal received from ISIS and the target actuation depth. For the purposes of this section the actuation depth is measured vertically upwards from the centre of the beam-pipe, therefore the smaller the number the further into the beam the target has actuated. Details of the geometrical relationship between the internal target measurements and the beam-pipe can be found in the MICE note 212 'Position of MICE target number 1 in the ISIS beam pipe' [114]. Simply, the translation is achieved by adding an offset to the target position measurement recorded by the DAQ. The data for this plot was taken from two runs, the first on the 16th May 2008 and a second run on the 2nd June 2008. The analysis for the first set of data (that was taken in May) is

illustrated; this process was repeated for the second set of data and both sets of data are used in the fitting in the final plot that shows the beam loss vs actuation depth relationship. At the time of running the number of actuations permitted was severely limited, particularly at higher beam loss levels. This is reflected in the number of events in some of the populations being quite small, typically  $\approx 200$  actuations or less.

After the DAQ system has recorded the beam loss it is post-processed to obtain the beam loss signal at extraction. This is done by heavily cleaning and then by fitting a curve to the last couple of ms of beam loss data before ISIS beam extraction. The average value of this curve is then taken to be the beam loss signal and is measured in mV - see figure 8.9 in chapter 8 for an example of a beam loss fit. The output from the DAQ system is not well calibrated and the absolute value should only be taken as an indication of the beam loss. Obtaining an accurate measurement of the beam loss that is consistent with ISIS's own readings will be undertaken when the beam loss DAQ has been upgraded. The aim of this exercise was to map the relative increase in beam loss for a given actuation depth. ISIS limits the actuation depth of the target by setting a beam loss limit. This beam loss limit restricts the pion production rate for MICE.

Figure 9.22 shows the number of events as a function of distance from the beam pipe centre. This plot shows the populations of events that correspond to different actuation depths. These populations were individually extracted and are shown again in figures 9.23 and 9.24. It can be seen that the error in the position is consistent with the resolution of the position measurement system  $\pm 0.15$  mm and so this latter figure is used as the error on the position.

For each event in a population the beam loss was calculated in the DAQ post processing using the method described above. Figures 9.25 and 9.26 shows the histograms of the beam loss for a given distance from the centre of the beam-pipe and allows an estimate of the error for both the beam loss and actuation depth for each population to be obtained. These are then plotted together with the second set of data in figure 9.27 to show the relationship between beam loss and the depth to which the target actuates with respect to the centre of the beam-pipe.

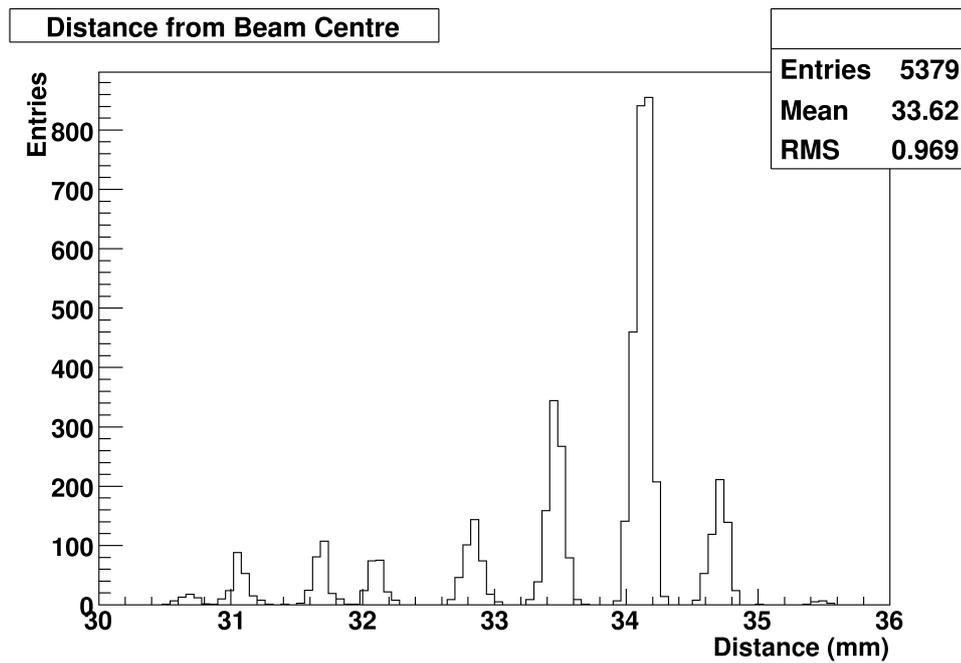


Figure 9.22: This plot shows the relative number of actuations done at each actuation depth during one continuous run of data taking. The actuation depth in this plot is the distance that the target actuates to measured from the centre of the beam-pipe. The individual populations were then individually extracted to estimate the error. This represents about half the data used to produce the final beamloss/actuation depth relationship.

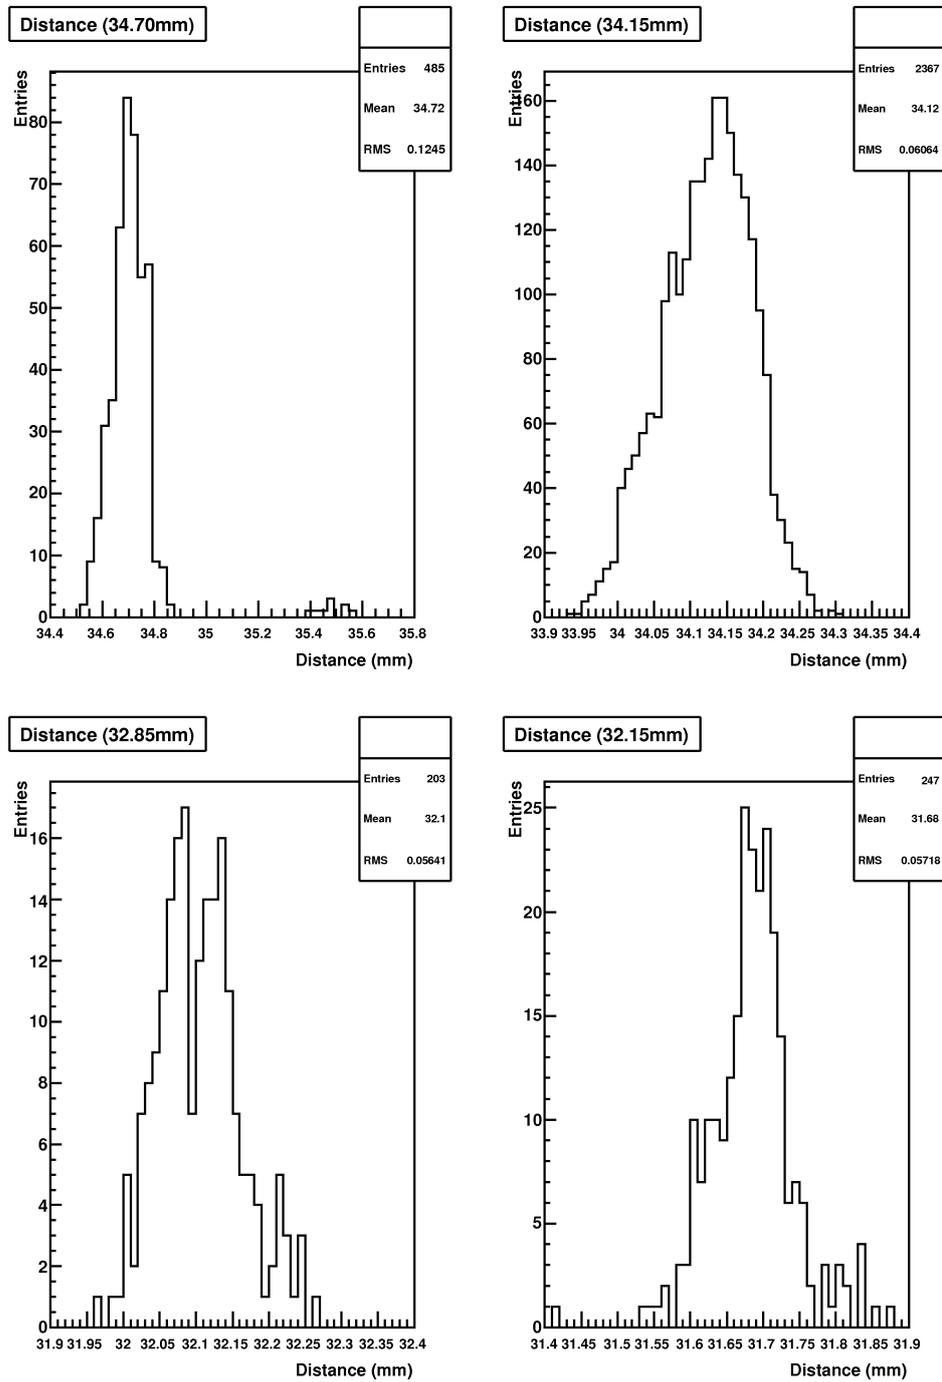


Figure 9.23: The first four individual entries in the previous plot (figure 9.22) have been extracted to allow an estimation of the positional error.

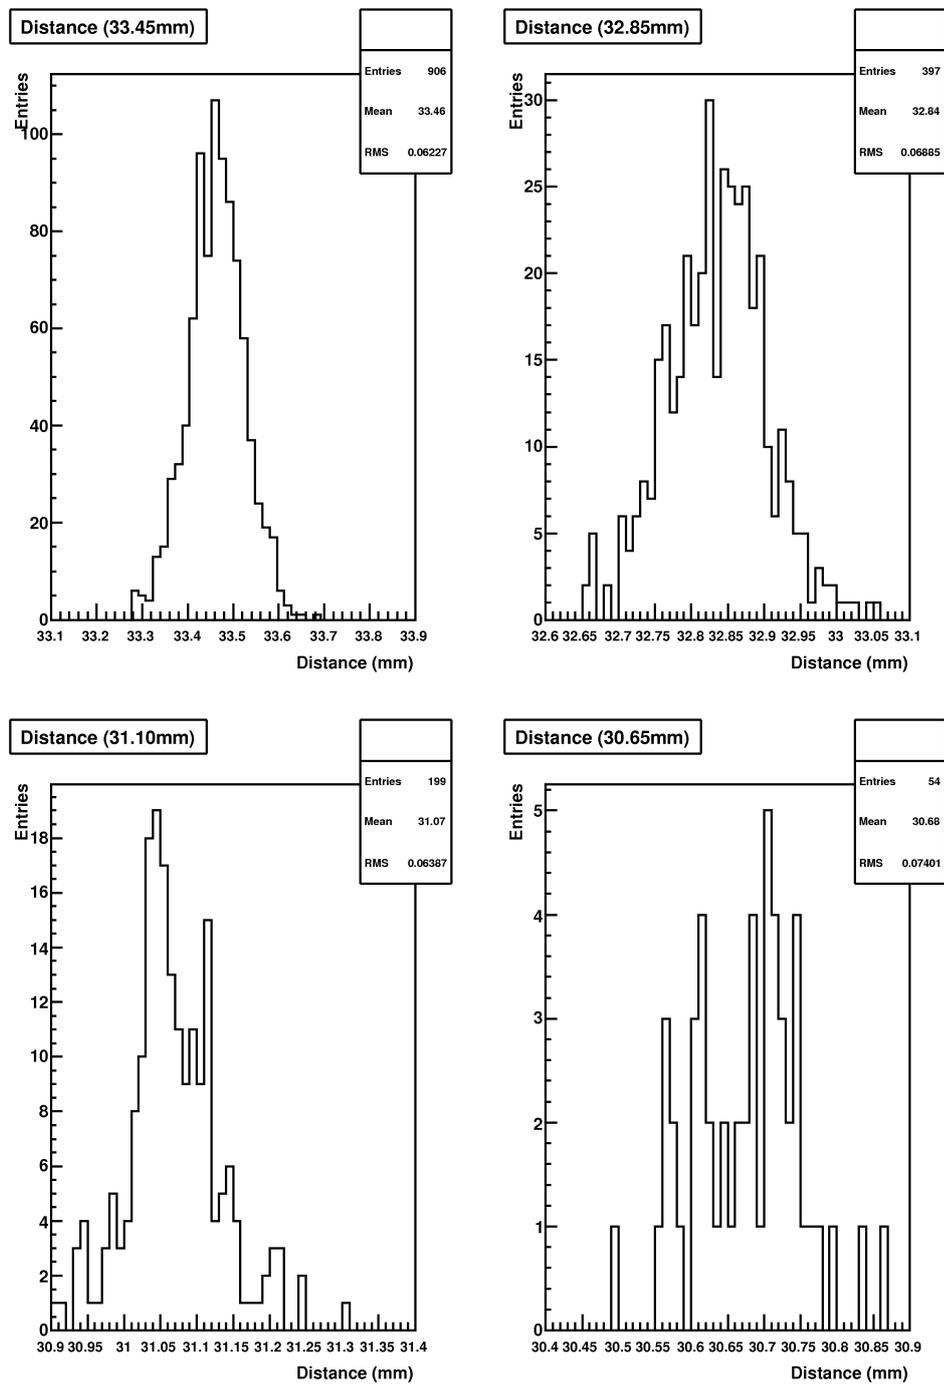


Figure 9.24: The last four individual entries in the previous plot (figure 9.22) have been extracted to allow an estimation of the positional error.

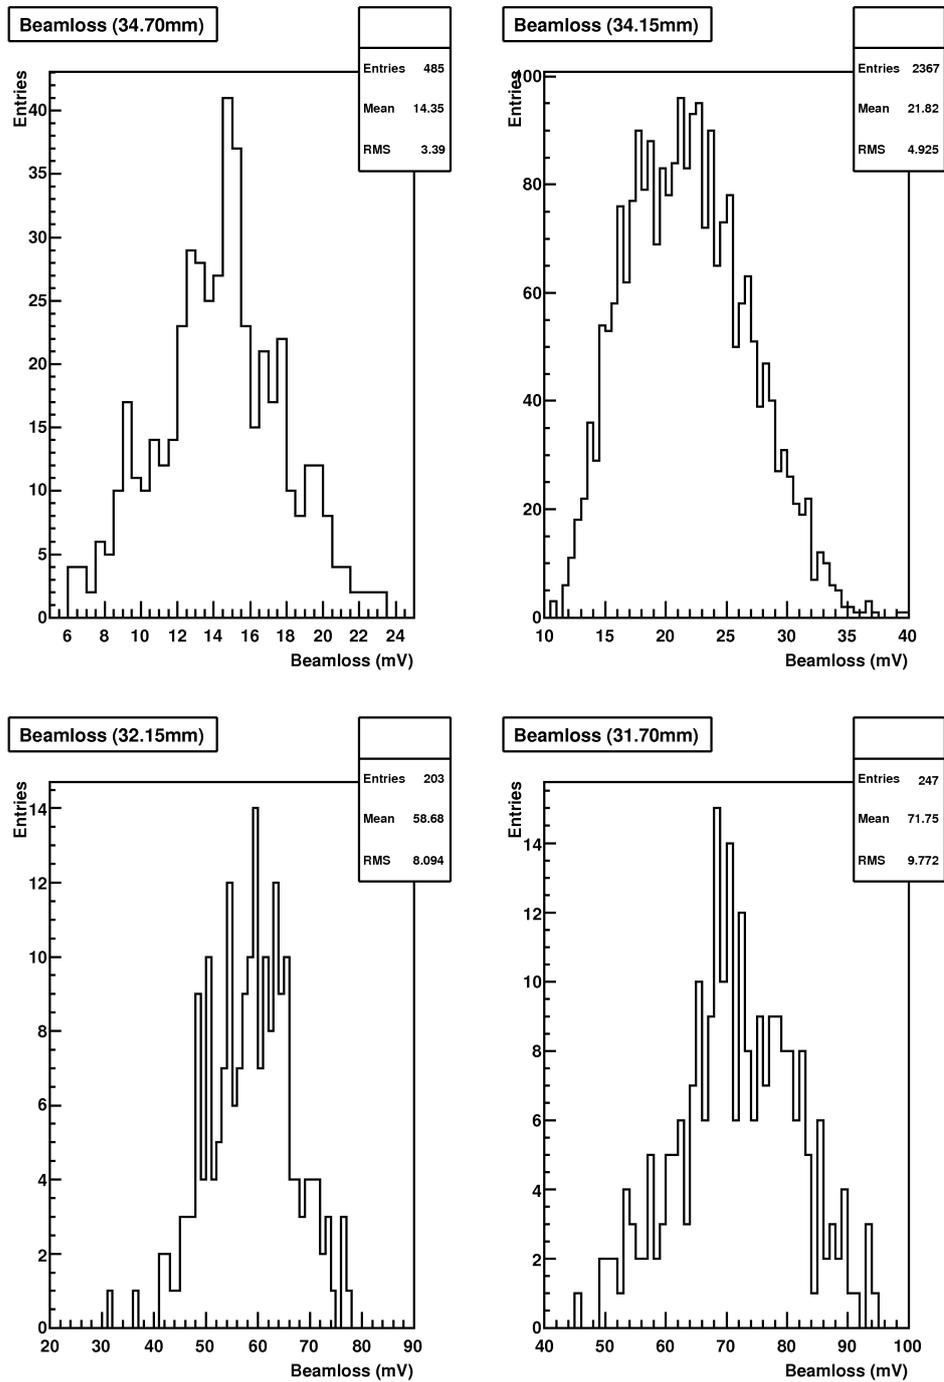


Figure 9.25: Histogram of the beam loss distribution for the first four actuation depths (with respect to the centre line of the beam-pipe) for the data set shown in figure 9.22.

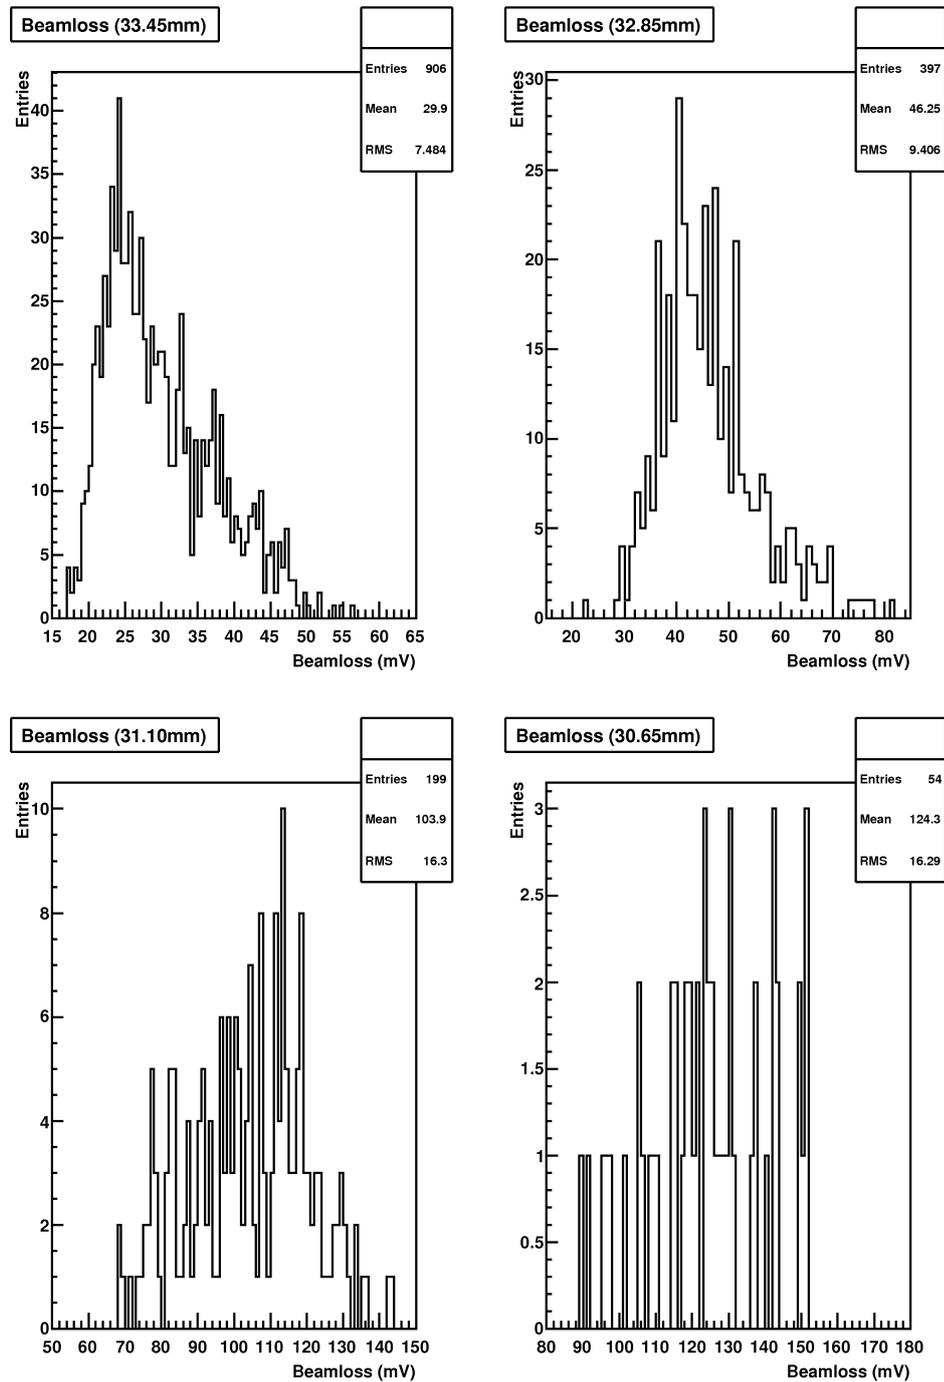


Figure 9.26: Histogram of the beam loss distribution for the last four actuation depths (with respect to the centre line of the beam-pipe) for the data set shown in figure 9.22.

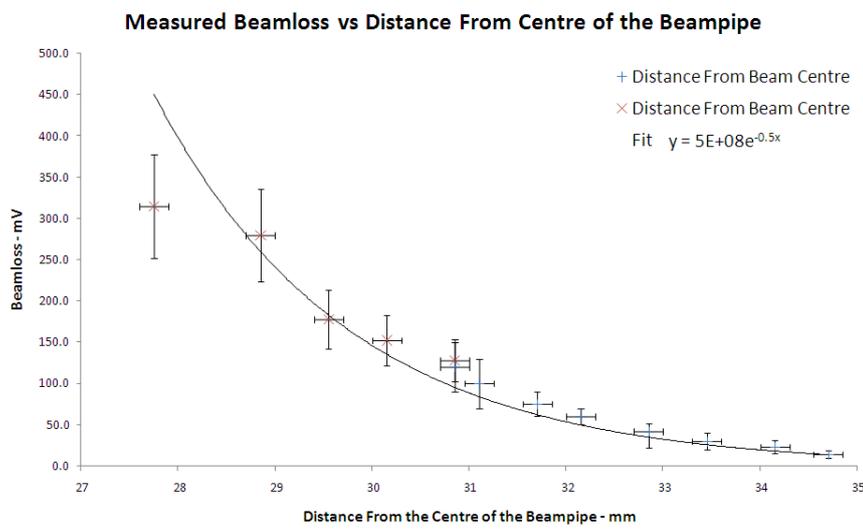


Figure 9.27: This plot shows the relationship between the beamloss and target's actuation depth. In this plot the second data set has also been added. Once again the distance scale represents the distance from the centre of the beam-pipe. An exponential fit through the data points has been shown.

### 9.8.2 Establishing Parasitic Operation

Once target operation had been established it was important to prove that the target could work parasitically on ISIS. Parasitic operation is defined as when ISIS is running at 50 Hz and the target is actuating on demand producing beam loss only during the last couple of milliseconds before beam extraction but is out of the way of the beam at all other times. Prior to this test ISIS had been run at 50/64 Hz to allow target operation to be established.

If the target is not able to accelerate quickly enough to get out of the way of the next beam injection an excess of beam loss would be seen on the beam loss monitors as the target intercepted the low energy protons. Excessive beam loss at injection causes ISIS to trip off. The target was actuated at a depth that gave 50 mV of beam loss during the last couple of ms before beam extraction. The target was then actuated at progressively later times in the cycle until the target intercepted the beam at the next ISIS injection. A measurement of the particle rate in the two scintillation counters positioned in the MICE beamline, GVA1 and GVA2, was performed. This data was obtained on the 15th May 2008.

This procedure showed that there was a 3 ms window between the optimum target insertion time and where injection losses on the next ISIS pulse were observed. Figure 9.28 illustrates how the beam rate in the MICE beam-line fell as a function of target delay during this 3 ms window and is taken from a MICE note entitled ‘Demonstration of parasitic running of the MICE Target’ [112]. It is clear from this that the target can manoeuvre significantly beyond its optimum position for MICE without causing unwanted disturbance to ISIS operation.

### 9.8.3 Pion Production for MICE

Much of the analysis in this chapter has been centred around demonstrating the capabilities of the target mechanism. However, ultimately the target mechanism has been built to serve MICE and its ability to provide MICE with pions is going to be critical to its success. Unfortunately for this thesis, delays with the target and the MICE experiment has meant that a good understanding has yet to be obtained of how the beam loss from ISIS corresponds to particle production for the MICE experiment.

What is clear so far from estimates correlating beam loss measurements to the proton flux in the MICE beamline is that the target is currently producing significantly less pions than will be required for MICE operation. A MICE note [115] was produced in April of 2008 that correlated the known beam loss and measured rates in the GVA1 and GVA2 counters to a simulation of the MICE beamline to estimate the implications for the muon flux in MICE. This note estimated that with 50 mV of beam loss the target would be producing about one good muon per ms per actuation. This correlates well with the results obtained by K. Walaron [56][111] discussed earlier in this chapter.

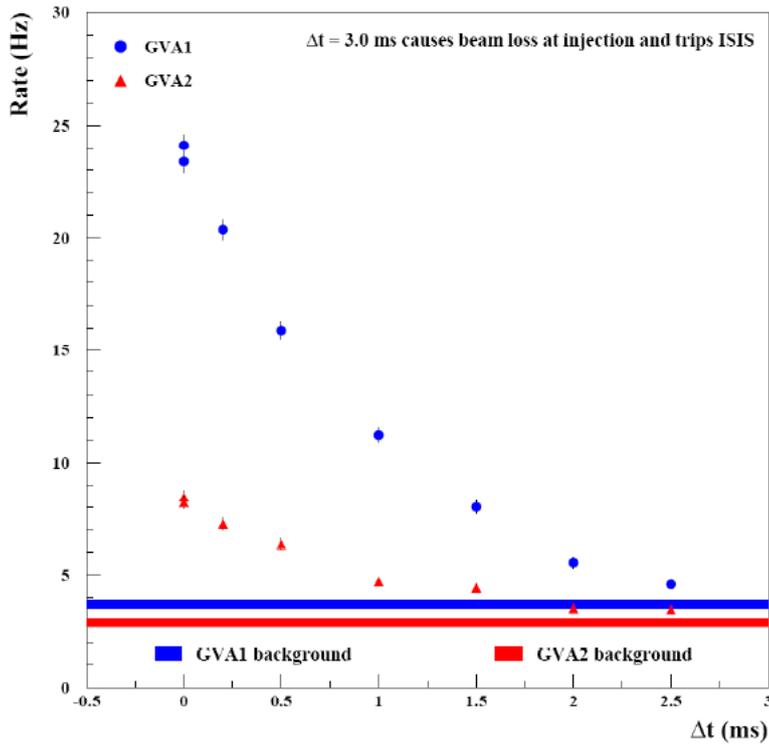


Figure 9.28: Rates of proton counts in scintillation counters GVA1 (circles) and GVA2 (triangles) placed in the MICE beamline, plotted as a function of target delay,  $\Delta t$ , for a constant dip depth after the optimum insertion time. The one standard deviation background rate bands are shown for GVA1 (upper band), and GVA2 (lower band).

However the crucial point is that this figure is currently 600 times lower than the muon rate required for MICE.

One important consideration with respect to this muon rate is that this has been obtained without the use of the pion decay solenoid. The pion decay solenoid is expected to increase the muon rate into MICE by about an order of magnitude once it has been switched on and commissioned. This is expected to happen in 2009. Further improvement to the muon rate can be anticipated if the target is allowed to produce more beam loss, however the amount of beam loss allowed will ultimately be constrained through radiological concerns as excessive beam loss from ISIS would likely result in the target area becoming a radiological hazard. The exact amount of beam loss that will eventually be allowed will be determined by the level of induced radioactivity caused through target operation. This limit has yet to be determined.

Simulations are currently being programmed to try and understand the interaction between the target and the ISIS beam. Optimisation of the target geometry may increase the yield of pions whilst minimising beam loss. Consideration is also being

given to the addition of scrapers around ISIS to assist containing the beam loss, hence the radiological hazard, to particular regions of the synchrotron. It is unlikely that any single solution is going to give the required muon flux. However by incremental improvements it is hoped that the muon flux will be increased to a value makes it practical to perform the MICE experiment.

## 9.9 Conclusion

The first part of this chapter showed the results of some early target tests undertaken in November 2006 that allowed the determination of several key measurements that were critical to the target design. Primarily the ISIS beam position was mapped; this established that the target geometry was correct (with respect to the target's reach into the beam pipe) and allowed the required target acceleration to be determined.

The installation of the target in 2008 and its subsequent operation has enabled the acquisition of a small amount of run time data. Some elementary analysis has been done on the data obtained during running on ISIS. This was compared to the data obtained from running in the laboratory. This data analysis quantified the performance of the target and verified the consistency of target operation after installation on ISIS. The addition of the beam loss data has also established the relationship between the target actuation depth and the beam loss produced from ISIS. Parasitic operation of the target has also been successfully demonstrated.

It is clear that the DAQ system needs upgrading so that the beam loss data can be better understood in terms of particle production for MICE. This upgrade is currently taking place. There are also concerns that the good muon rate for MICE is going to be significantly lower than expected due to the limitation on the beam loss that can be produced from the target. The MICE target mechanism has proven to be, and still continues to be, an extremely challenging project. Even in view of the successes highlighted in this chapter, there is still a significant amount of development work to be done. The next chapter will highlight the details of some of the more serious problems with the mechanism and how those problems have or are being resolved.



# Chapter 10

## Further Developments

### 10.1 Introduction

This chapter gives additional details on some of the practical problems that have been encountered with the target mechanism since its original installation onto ISIS. It should be pointed that this chapter has been written some ten months after the rest of the thesis was originally completed and so some of the information presented here benefits from greater hindsight and operational experience with the mechanism. Broadly speaking the issues to be discussed fall into two categories. Firstly there were repeated catastrophic failures of the stator coils. The investigation by the author and final solution to this problem is described here. Second, the target mechanism has had problems with bearing wear. For completeness the efforts of the collaboration to find a viable bearing solution are outlined here. Finally to address these mechanical wear issues a range of studies were carried out. One of these studies was to field map the stator and compare the results to a electromagnetic Finite Element Analysis (FEA) for which the author took responsibility. This study, along with other studies undertaken by the collaboration that are outlined in this chapter, have driven work on a redesign of the target shaft and bearing design. Again for completeness, the proposed redesign is described.

### 10.2 Electrical Insulation of the Stator Coil Windings

The stator is comprised of three phases where each phase contains eight coils. At any time during operation two phases are connected in series and so this means that the voltage that coil  $N$  (where  $N$  represents the number of coils counted with respect to the grounded connection) is being held at is

$$\frac{\text{PSU Voltage}}{16} \times N \tag{10.1}$$

and this voltage is the coil's floating voltage. The effect of the cooling jacket being earthed is to place this floating voltage between the coil's outer windings and the copper shims leaving the coil insulation to support its floating voltage plus any additional transient spikes. The close proximity of the earthing arrangement to the coils was not foreseen to be a problem because the coil wire is a 'grade 1' high temperature enamelled wire rated at 1.5 kV at room temperature and 750 V at working temperature. The highest voltage that would be seen at a coil is 240 V DC plus any transients. The transients were later shown to be approximately 60 V maximum.



Figure 10.1: A photograph of the coil stack clearly showing the location of the coil failure.

The first stator that was installed onto ISIS in January 2008 had been used in the laboratory at Sheffield University for many millions of actuations prior to its installation to test various bearing combinations. Therefore this particular stator had been well tried and tested before being placed onto ISIS. At the time the stators were operating at 220-240 V DC across the coil stack, this meant that each coil saw  $\approx 15$  V across its windings; at this voltage failure internal to the coil due to insulation breakdown is extremely unlikely and no evidence of this has ever been seen. During its operation at Sheffield and before installation on ISIS the stator suffered from a single coil failure. Insulation break-down resulted in a short circuit from a coil to the copper heat-sinking shims that sandwiched each coil. (It should be noted that the kapton shims detailed in Chapter five were not used in this stator build and only copper shims were placed between adjacent coils.) The faulty coil was replaced and further to this the stator had no further coil failures during its operational lifetime.



Figure 10.2: The current surge had caused the coil to weld itself to the heatsinking shims. Approximately half a turn of coil was vaporised from the current pulse.

After installation of the primary stator on ISIS a second system was set up to run in building R78 at the Rutherford Appleton Laboratory. The purpose of the second system was to run ahead of the primary system on ISIS such that it was always to have completed three times as many actuations as the stator on ISIS. This demonstrator system should then act as an early warning system to indicate if excessive bearing wear was taking place.

The demonstrator stator was built shortly before its installation in R78. Time constraints had only permitted the stator to be actuated  $\approx 50,000$  times before it was transported from Sheffield to RAL for installation. Shortly after installation the demonstrator stator failed catastrophically due to a coil short circuiting to a copper heat-sinking shim. Photographs of this failure are shown in figures 10.1 and 10.2. This was the second coil failure witnessed on a stator over a period of about 18 months. Despite this failure there was still significant pressure to provide the target demonstrator system and so within just a few days another stator was rapidly assembled from spare components at Sheffield and taken to RAL untested. This second demonstrator stator also failed catastrophically after only a few hundred actuations due to a similar coil insulation failure to earth, although in this case the short circuit was from a coil to the centre bore tube - see figures 10.3.

It had become clear that a serious problem with the stator had been discovered and that this needed addressing before any further stators could be built. The lack of a demonstrator system led to the stator on ISIS being run cautiously over the forthcoming year with the total number of actuations being completed whilst installed on ISIS just exceeding 180,000.

Given the relatively low voltages at which the system operates these coil failures were

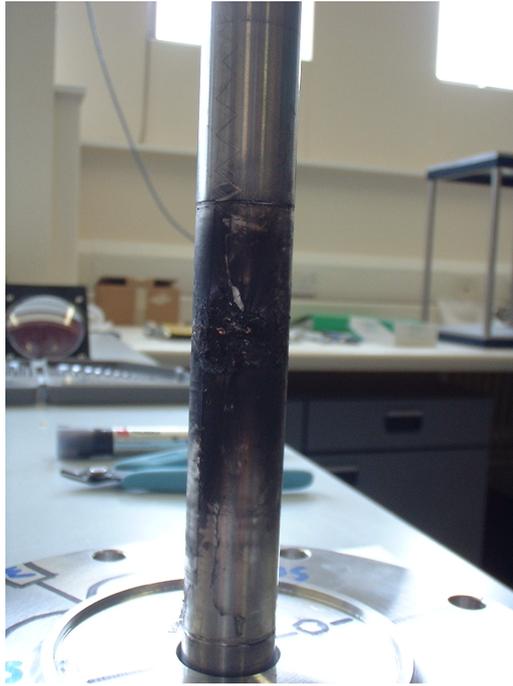


Figure 10.3: On the second electrical failure the coil insulation failed to the centre bore tube causing considerable damage; the current pulse had embedded copper from the coil into the bore tube.

indicative of imperfections in the insulation of the coil wire. Close inspection of the coils from the failed stators indicated that mechanical damage was occurring to the inner bore of the coil stack during assembly. A trial stack also demonstrated that damage to the coil faces could very easily occur if small foreign bodies were present within the coil stack during assembly. Given that these stacks had been put together in a workshop this possibility could not be overlooked. It is also worth noting that up until this point the integrity of the coils after receipt from the manufacturers had been assumed but this was now challenged. It was clear that each coil needed electrically testing before being installed into a coil stack. Several steps were taken to improve the quality assurance of the coils to prevent a repeat of this failure mode.

The stator bore tube was to have a soft plastic head placed upon it before insertion into the coil stack to prevent mechanical damage to the inner bore of the coils. A double wrapped layer of kapton tape was added to the bore tube, wrapped in such a way that there were no exposed seams down the length of the bore that sat within the coil stack, providing additional insulation between the bore tube and the inner face of the coils. This was also the stage where two of the six copper shims placed between adjacent coils were removed and replaced by a  $25 \mu\text{m}$  kapton film as detailed in Chapter 5. This film was tested in the lab and was shown to be able to reliably hold off 2 kV DC for a test period of one minute. This was well within the specification of the material which has a dielectric strength of  $\approx 300 \text{ kV mm}^{-1}$  [116].

A brand new batch of coils was manufactured and these coils were double dipped in a sealing varnish to help coat any imperfections in the insulation. All coils received were kept scrupulously clean before being used in a stator. A mechanical boss was constructed to permit insulation tests on individual coils. The boss mechanically clamps a coil so that the top, bottom and inner faces of a given coil are in contact with the boss. 1 kV is placed upon the coil leads for several seconds whilst the boss itself remains grounded. This coil testing mechanism revealed that approximately 10 % of the coils from the manufacturer's could not withstand 1 kV without breakdown. It is worth noting that the failures at 1 kV does not mean that all of the coils that failed the test would have necessarily failed in a stator because the stator's operational voltage is maximally only 30 % of the test voltage. However the large failure rate of the coils at 1 kV does give a clear indication of why coil failure had become problematic prior to the implementation of any quality assurance (QA) on the coils. After assembly the completed stator was then subsequently tested to 1 kV before being operated to ensure that no mechanical damage had occurred to the coils during assembly.

Two stators were built with new quality assured coils and kapton insulation was added between adjacent coils as described in chapter 5. Both stators were operated for 500,000 actuations at Sheffield. One of these stators was then installed in the demonstrator system where it operated for 340,000 actuations. The second stator was not utilised at RAL due to mechanical concerns. There were no coil failures.

The addition of the kapton shims to the stator coil stack had increased the operational temperature of the stator from approx 70 to 90 °C because of their thermally insulating effects. As pointed out in the previous chapter the consequent increase in the core temperature of the coil stack also increased the coil resistance. The stator works from a fixed voltage so the increased resistance had the effect of decreasing the coil current. The net effect is that the acceleration of the shaft fell from 80g to  $\approx 72g$ . This loss of acceleration was eventually determined to be unacceptable because it reduces temporal manoeuvrability of the target and increases unwanted beamloss in ISIS due to the longer actuation period. Therefore it became apparent that the kapton shims had to be removed and replaced with the copper shims to provide the necessary heatsinking.

Further to this a design change of the power supply was used to reduce the effective voltage seen by the coils to ground by nearly 50 %. Up until this point the power supply had been providing the stator with a voltage of between 0-220 V DC. By adding a second identical capacitor bank and a second capacitor charging unit to the system it then becomes possible to operate the stator from a split supply of -125 to +125 V DC - see figure 10.4 [117]. From the stator's point of view nothing has changed, it still sees the full voltage across both phases, however the maximum voltage to ground has now been reduced by  $\approx 50\%$ . The two identical capacitor banks are charged independently but because they appear in series to the stator the effective capacitance has been halved. This effective reduction in capacitance increases the rate of discharge of the banks and has to be offset by operating the capacitor banks at a slightly higher voltage. A voltage increase of about 15 V on each bank was found to be sufficient to counter the effects of the faster discharge.

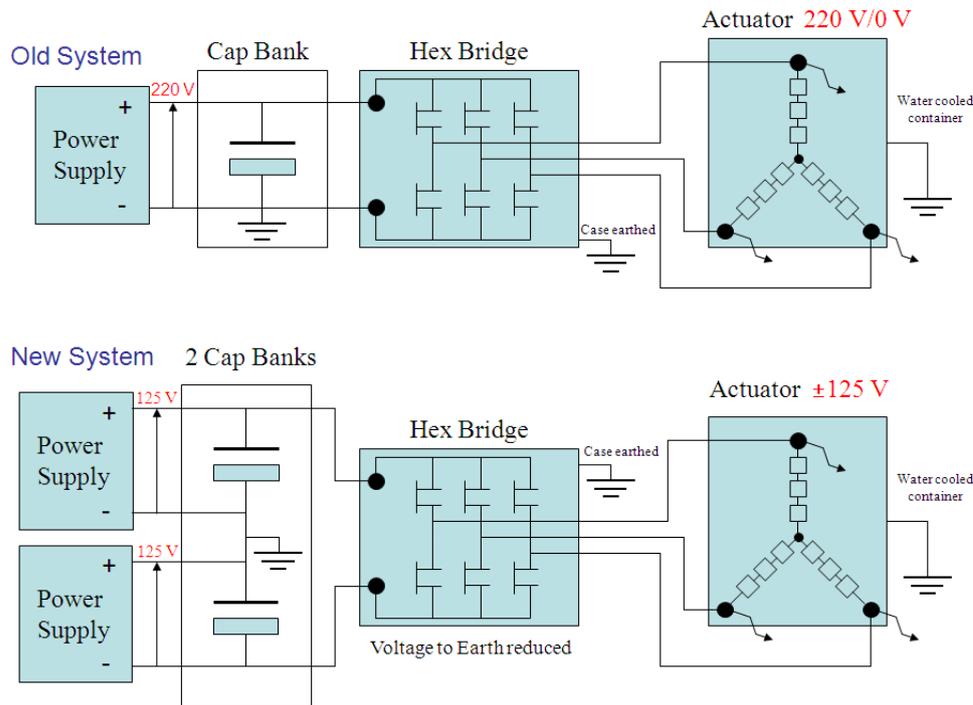


Figure 10.4: The new power-supply design [117] nearly halves the voltage to earth helping to reduce the risk of a coil failure to earth. The highest operational voltage that the coils should see to earth when the new power supply is used is  $\approx 125$  V, plus a 30 V spike on switching. All the coils are tested to 1 kV before being used in a stator, which gives a significant safety margin.

The split power supplies have now been installed on the target system in Sheffield and on both systems at RAL. Given the additional stringent QA criteria with the coils the problem of premature coil failure appears to have been solved.

### 10.3 Bearing Development

Since the early developmental work that was discussed in Chapter 4 led to the conclusion that the accelerated components of the target mechanism needed to be supported on sliding bearings if the required acceleration was to be achieved, there has been concern that the wear rate of these bearings would prove to be problematic. To recap; the target mechanism must work in the vacuum of ISIS and this means that these sliding bearings must remain unlubricated. Therefore, the wear rate of these bearings will be determined primarily by the magnitude of the lateral forces on, the hardness and the quality of the surface finish of the contact surfaces. It is perhaps useful at this stage to give the reader an overview of the work that went into development of the bearings before the first working target mechanism was installed onto ISIS. A summary of the

different bearing materials that were tried can be found in table 10.1.

Material	Actuations to Failure	Comment
Machinable Ceramic	100,000	These bearings failed quickly when operational acceleration was achieved. (The ceramic bearings showed little wear when operated at low acceleration.)
Brass	250,000	These bearings generated a lot of debris.
Ruby	50,000	Ruby Bearings failed quickly due to large contact forces. They also caused significant damage to the shaft.
Leaded Bronze	2 million +	These looked promising but the use of the material was not permitted within ISIS.
Diamond Like Carbon (DLC)	2 million +	The performance of this coating has been variable but up to 2 million actuations has been demonstrated. The coating was approved for use in ISIS and will be used until a better coating/material can be found.

Table 10.1: A summary of the results of using different bearing materials

## 10.4 Bearing Wear

Following the resolution of the electrical insulation issue two identical target mechanisms were built. The plan was to run the first of these two systems for  $\approx 500,000$  actuations and, if all was well, this target mechanism would then be installed onto ISIS, replacing the mechanism that was installed in January 2008. The second target mechanism would then run in the demonstrator system. Both of these target mechanisms had been run for 500,000 actuations in Sheffield using an older shaft and an older set of bearings prior to installation in R78 at RAL. Both mechanisms completed 500,000 actuations without incident and no noticeable additional wear was seen on either the shaft or bearings that were used during this trial period on the Sheffield installation. Each of the stators were to have a new set of bearings and a new shaft upon installation at RAL.

The first mechanism was installed and ran for 340,000 actuations before problems were encountered. Upon examination of the mechanism there was considerable damage to both the shaft and the bearings and it was clear that the DLC coating had failed

prematurely. There was a significant amount of dust but this had failed to penetrate through the bottom bearing and so had not been observed in the viewing port.

Figures 10.5 - 10.8 show photographs of the damage to the shaft and the bearings. The upper shaft and bearing showed the worst wear. The lower bearing surfaces also showed significant wear but it is not clear whether this was primary damage or damage induced by the introduction of particulate matter into this bearing as a result of the dust falling through from the upper bearing.

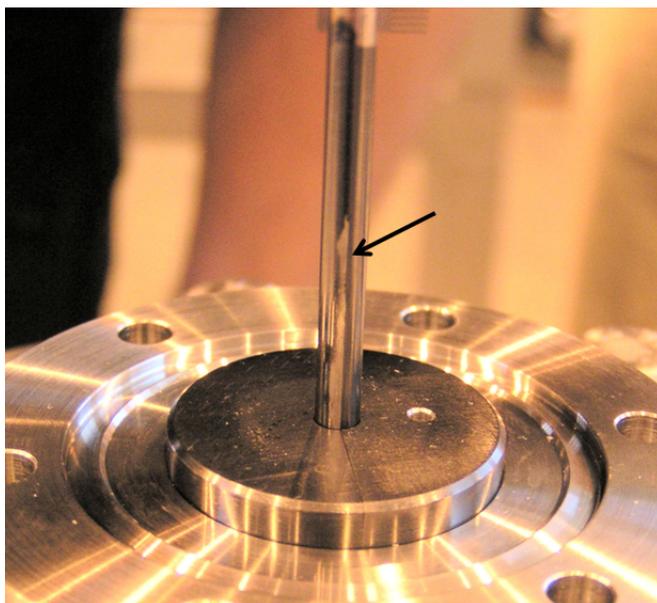


Figure 10.5: Upon opening up the mechanism there was significant particulate debris on the top bearing and the DLC coating had been worn through on the shaft.

The failure of this mechanism led to a comprehensive internal investigation as to why this particular target mechanism had failed so quickly. This was the first mechanism to show catastrophic failure of the DLC bearings; at least two other sets of bearings had demonstrated an operational lifetime in excess of 2 million actuations at Sheffield. The next two sections of this chapter details the results of this investigation and discusses the mechanical details of the subsequent redesign that resulted from it. It should be noted that as of writing this chapter in July 2009 the redesigned mechanism was still being being manufactured so no operational experience with the new design had been acquired.

## 10.5 Analysis of the Target Mechanism Failure

The subsequent analysis of the actuator failure followed three lines of enquiry and all of these will be discussed in this section. First, an investigation of the DLC coating

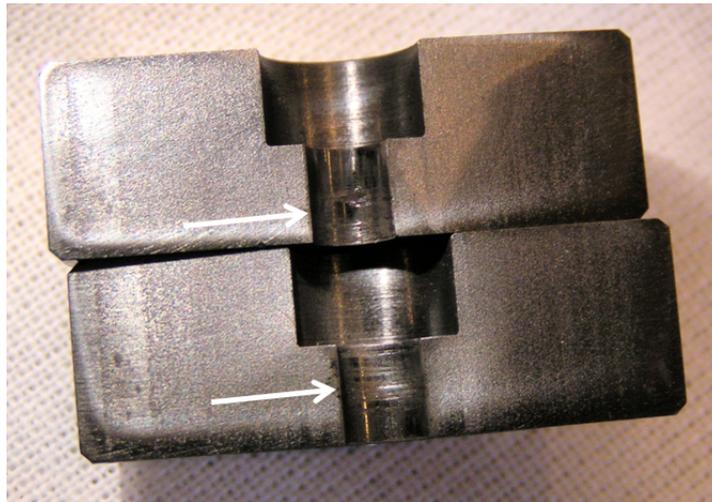


Figure 10.6: There is considerable damage to the bearing faces on the upper bearing. Although it isn't clear in the photograph, a lot of the DLC has been removed in the damaged areas.



Figure 10.7: Much of the DLC has become detached from the surface of the cruciform cross section where it had passed through the bottom bearing.

technique was undertaken to try and ascertain whether a process problem had contributed to the premature failure of the bearing. Second, a test rig was set up to measure the magnetic field inside the stator to ensure that no significant radial field asymmetries existed within the powered stator. These magnetic field measurements were complemented by an electromagnetic finite element analysis. Third, a mechanical survey was undertaken to try and ascertain whether there were any fundamental mechanical problems with the design of the shaft and bearings.

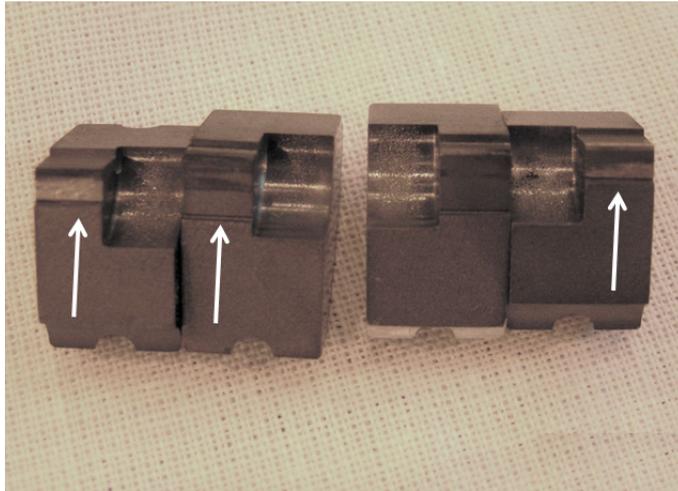


Figure 10.8: Bottom bearing quadrants: Of the four bearing faces that can be clearly seen in this photograph, three of them are badly damaged. The other four faces (not possible to see clearly due to the angle of the photo) showed similar damage.

### 10.5.1 DLC Coating

The rapid failure of the DLC coating after the previous success warranted investigation and initially advice was sought from the company that applied the DLC coating to the shaft and bearings. The company examined the components and concluded that the DLC had failed to adhere properly to the surface of the components and that this was the likely cause of the rapid failure of the coating[118]. The cause of the lack of adhesion was unclear but several possibilities were presented.

- Remnants of any polish used to pre-treat the surfaces before coating with DLC.
- Any silicon based chemicals left on the components from machining.
- The cruciform cross-section in the lower bearing was created using wire erosion. Wire-erosion creates an oxidation layer on the surface of the material that can reduce the adhesion of the DLC to that surface.

The first of these possibilities would have been the responsibility of the company as the pre-treatment of the components is done on site just before coating. However the other two items were clearly relevant to the manufacture of the bearings and in future the use of wire erosion or silicon based oils clearly needed to be avoided. It was also pointed out the DLC adheres preferentially to smoother surfaces and so the components would benefit from an improved  $R_a$ , the roughness amplitude in microns. Subsequent metrology showed that the  $R_a$  of the components that had gone for DLC coating was 1.6. Preferentially this needed reducing to an  $R_a$  of 0.1 - 0.02.

### 10.5.2 Magnetic Survey and FEA

It was speculated that the cause of the excessive wear could have been a significant lateral force on the magnets due to an asymmetry in the magnetic field inside the stator. Up until this time the symmetry of the field in the stator had always been assumed. However there were small but clear differences between the size and winding detail of different coils and the tolerances on the build of the coil stack could leave the coils positioned off centre by up to a couple of hundred microns. To ascertain whether this was having a significant impact upon the symmetry of the field inside the stator a magnetic field survey of the stator from the failed mechanism was done to determine if any asymmetry existed. The results from this survey were then compared to an FEA model of the stator coils and this model permitted an estimation of the field asymmetry within the stator bore.

The radial component of the stator field was measured by scanning a Hall probe down the inside of the bore of the stator that had been used in the demonstrator system. The data taking was done by using a motorised unit that allowed the probe to be moved in the  $\pm z$  direction to a resolution of  $200\mu\text{m}$ . The stator itself was mounted on a rotating stage centred on an aluminium boss that protruded a short distance into the end of the stator bore tube. This allowed the stator to be rotated through a full  $360^\circ$  so that an azimuthal angle in the stator bore could be selected. The Hall probe was mechanically constrained so that it rested against the inner edge of the bore tube that runs through the centre of the stator. The Hall probe tip contained two Hall sensors set at  $90^\circ$  to each other. Only the radial sensor was utilised permitting the radial field of the stator to be measured as a function of axial and azimuthal position. The geometry of the stator bore tube and the geometry of the sensor head meant that the active area of the sensor head would have been at a radial distance of 6.5 - 7 mm from the centre line of the stator bore.

Connected to the stator was a recirculating water chiller. The use of this permitted a maximum current of 7200 mA DC to be maintained through the coils indefinitely. Data was taken with the stator in 'State 1' and 'State 2', therefore all three phases of the stator would have their fields measured. Figure 10.9 shows the measured voltage from the Hall probe as a function of axial distance for the 'State 1' measurements and figure 10.10 shows the measured voltage for the 'State 2' measurements. Each graph comprises of several sets of data, each data set corresponding to an azimuthal progression of the stator of  $30^\circ$ .

There are several features to note from the data. (Where the data is shown in Volts the conversion to Tesla is  $1\text{ mV} = 10\text{ mT}$ ). Firstly, for perfect axis-symmetry then the data from the different azimuthal angles should coincide, so taking the standard deviation on the spread of the data for each set of points at a given axial depth gives some indication of either the field asymmetry or misalignment/error of the measuring rig. Secondly it would be expected that the field would drop to zero as the Hall probe is moved out of the stator. In both of these graphs it is apparent that there is an offset of approximately  $400\mu\text{V}$ . In each of the figures it can be seen that the graph is composed

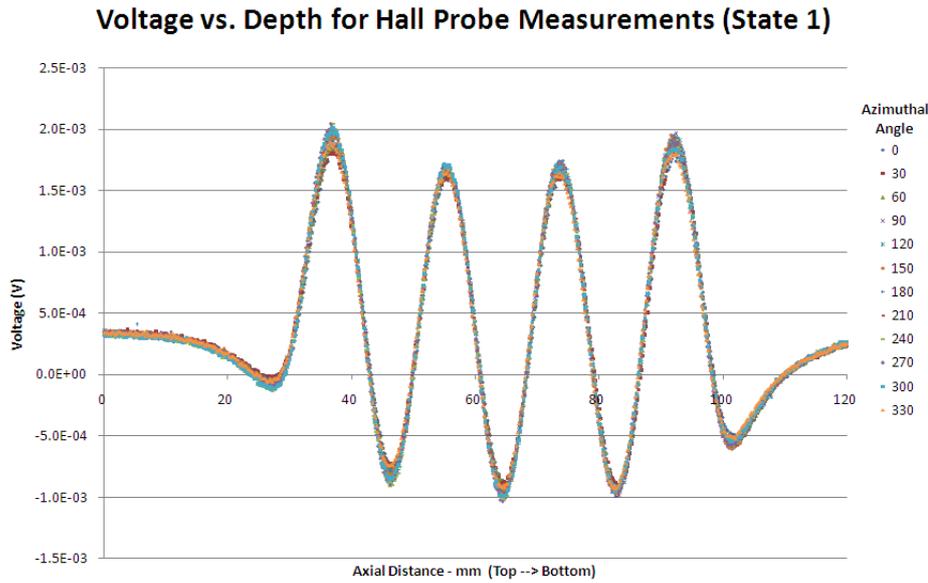


Figure 10.9: The voltage output from the Hall probe changed as a function of axial position within the stator bore in an azimuthal progression of  $30^\circ$  over a full  $360^\circ$ . The coils were wired in ‘State 1’ and the coil current was 7200 mA. 1 mV is equivalent to 10 mT.

of several peaks and troughs. If these peaks and troughs are labelled in numerical order from left to right and the field value plotted for each azimuthal measurement then a pattern emerges for both the sets of data as shown in figures 10.11 and 10.12. For the case of a perfect field and an accurate measurement a series of horizontal lines would be expected.

For the ‘State 1’ data there appears to be a general trend for the average field values to increase as a function of angle. This could be caused by the field not being quite coaxial, however it is not possible to tell if this effect is actually due to a field inhomogeneity or due to the precision to which the Hall probe can be placed. It would certainly be reasonable to assume that the axial repeatability of the placement of the Hall probe was accurate to within a couple of hundred microns. For the ‘State 2’ data the step in the data between  $210^\circ$  and  $330^\circ$  is suggestive that there was a systematic error during the data taking. This would be seen to be confirmed by the fact that the step occurred at a point where the stator had to be rotated  $330^\circ$  backwards on its stage instead of  $30^\circ$  forward to prevent the wires from tangling.

A Finite Element model of the stator was built using an electromagnetic FEA program called Opera 3D[119]. The axial field produced inside the bore of the coils was then calculated at different radii. As the radial position of the Hall probe sensor was reasonably well known, matching the simulated results to the measured data and cross checking the results provided a good verification on the general validity of the model

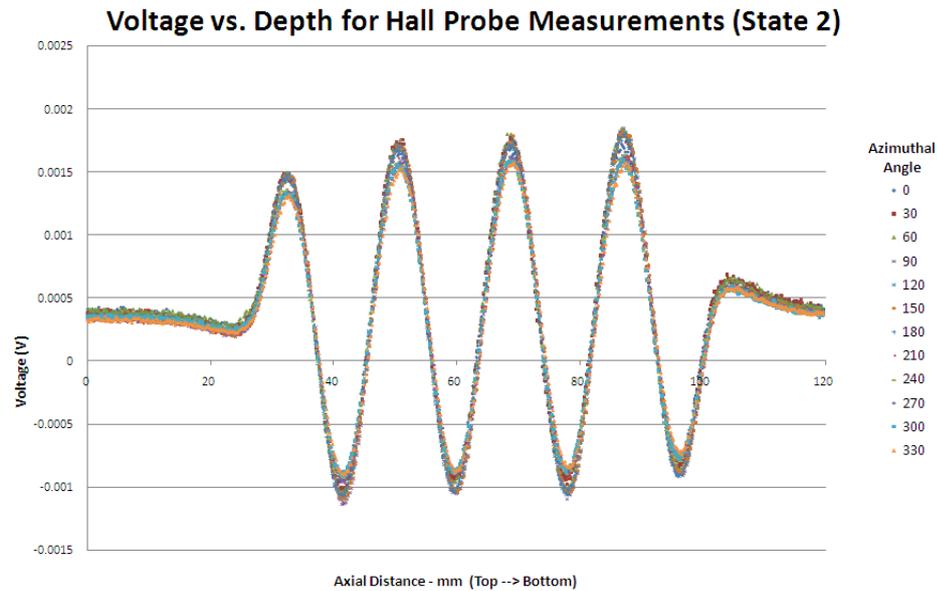


Figure 10.10: The voltage output from the Hall probe changed as a function of axial position within the stator bore in an azimuthal progression of  $30^\circ$  over a full  $360^\circ$ . The coils were wired in ‘State 2’ and the coil current was 7200 mA. 1 mV is equivalent to 10 mT.

which was then used to provide additional insight. Figure 10.13 shows the best fit of the model to the data for ‘State 1’. This shows two axial plots superimposed on the measured data at two different radial distances. A good correlation, neglecting end effects, is obtained at a radial distance of between 6 - 7 mm and this distance matches the position of the Hall probe when the measurements were taken. A correspondingly good match was also obtained for the ‘State 2’ data.

The model was then used to calculate the radial field strength as a function of radial distance at a fixed axial location. The cross-section is plotted at the point where the axial field strength is at a peak value - see figure 10.14 for the ‘State 1’ simulation. The plot for ‘State 2’ is almost identical and is not included here. In both cases the rate of change of field with respect to radius at the axial peaks is quite large. This gradient indicates that any deviation of the centre of the magnets from the central axis of the stator could have a significant effect on the lateral forces felt on the magnetic assembly, although it would require a more complicated model to quantify and confirm this as the permanent magnet assembly would need to be included in the model. Numerical differentiation of the plot at a radial distance of 7 mm reveals that  $dB/dR$  is  $\approx 4 \text{ mT mm}^{-1}$  for both ‘State 1’ and ‘State 2’.

The standard deviation in the spread of values over the angular measurements for each axial position within the measured data sets was taken. Figure 10.15 show the results of this for the ‘State 1’ data. There is some indication in these graphs that the

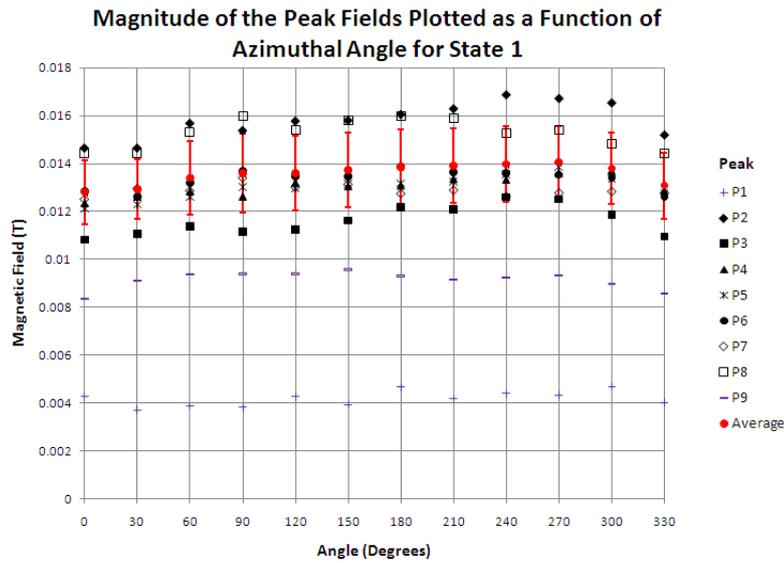


Figure 10.11: The peaks in figure 10.9 are numbered from 1 to 9 (P1 - P9) and the peak value plotted as a function of azimuthal angle. In the case of a perfect field with no measurement error one would expect to see that each set of points would lie on a straight horizontal line. The average value is taken from points P2 - P8 for each azimuthal angle.

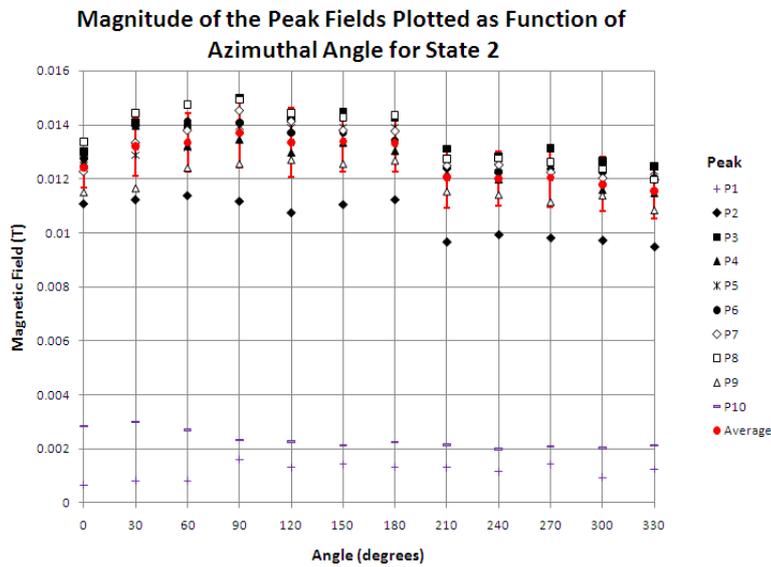


Figure 10.12: The peaks in figure 10.10 are numbered from 1 to 10 (P1 - P10) and the peak value plotted as a function of azimuthal angle. The step in the data is suggestive of a systematic error in the data taking.

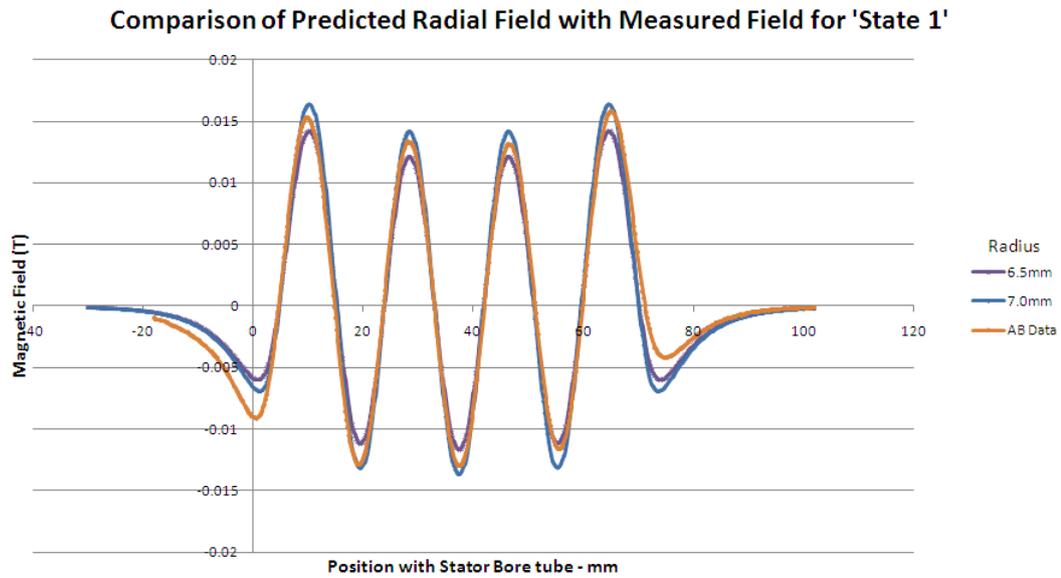


Figure 10.13: The measured data corresponds very well to the modelled data at a radial distance of 6 - 7 mm from the central axis of the stator bore. This distance corresponds to the radial distance that the Hall probe was at when the axial data was taken. Note in the legend that 'AB data' means the 'State 1' data, the other two data sets are produced from FEA.

standard deviation is weakly correlated to the measured field strength. As the value of the standard deviation falls significantly at the axial positions that correspond to being outside the stator volume this is indicative that the standard deviation is giving a true measure of the either the asymmetry in the field and/or the error in the measurement and is not due to sensor noise.

The maximum observed standard deviation from the measurements couples with the rate of change of field strength as calculated from the model to allow an upper estimate of the physical displacement of a perfectly axis-symmetric field that would be required to account for the spread in measured values. These results are tabulated in table 10.2. The conclusion from this initial analysis was that the maximum field displacement within the stator bore was  $\approx 300 \mu\text{m}$ . However the standard deviation in the measured field has two uncorrelated components, one due to the actual field asymmetries and the other due to errors in the positioning of the probe. The estimated effect of the Hall probe positioning errors is to contribute  $200 \mu\text{m}$  to the scatter of the field displacement, so the maximum field displacement is therefore somewhat less than  $300 \mu\text{m}$ . There is a plan to produce more accurate field maps of the stator's coil stack as a routine part of the QA of the stator manufacture in the future; however, creating an accurate field map due to the small feature size is non-trivial.

Alongside the field mapping it would also prove to be useful to add the shuttle's per-

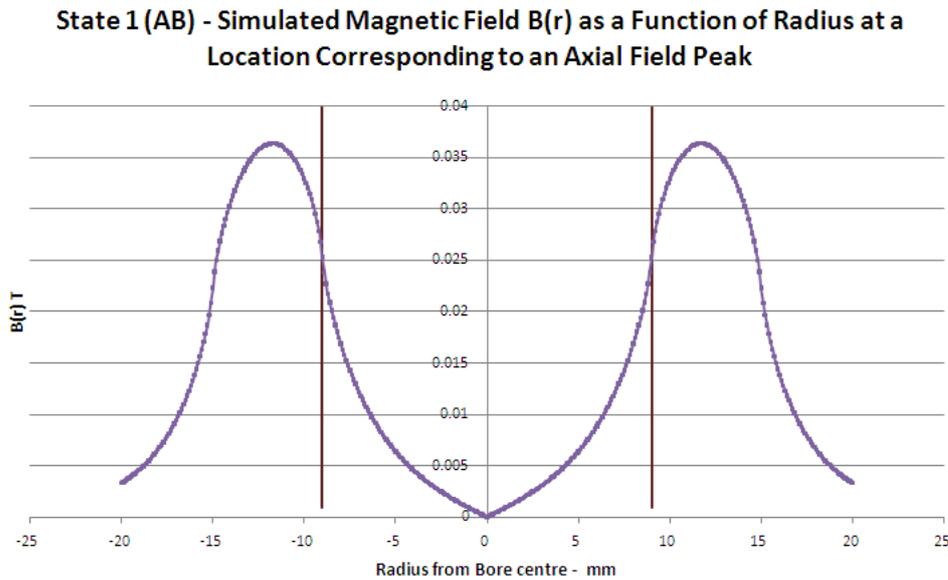


Figure 10.14: A simulation of the change in radial field strength across the bore of the stator at a position that corresponds to a peak in the axial field measurements based on experimental data. The two vertical lines represent the position of the inner radius of the stator bore.

manent magnets to the stator FEA model to enable a calculation of the resultant radial forces on the magnets as they pass through the bore tube. Once the results have been obtained for perfect symmetry it would then be useful to displace the shaft from its central axis by  $\approx 300 \mu\text{m}$  to allow an estimation of the maximum radial force that the magnets will see during actuation. As of writing this chapter these additional studies are being undertaken but are not yet complete.

### 10.5.3 Mechanical Survey

The mechanical survey started with a high speed camera recording of the target motion. The camera used was capable of recording several thousand frames per second. This allowed the motion of the target to be captured and analysed. The results from these recordings showed a definite oscillation of the target blade during its actuation that was indicative of a vibrational mode. These were then compared to the findings of an FEA model of the target's mechanical setup. This model was built to understand any possible vibrational modes of the target system[120].

This report concluded that the design of the shaft would likely see a natural oscillation of 27 Hz which was close enough to the 'equivalent' driving frequency of 33 Hz (based upon 1 cycle taking 30 ms) of the target that this mode of oscillation would be excited - see figure 10.16[120]. The report also analysed the effects of this mode of oscillation

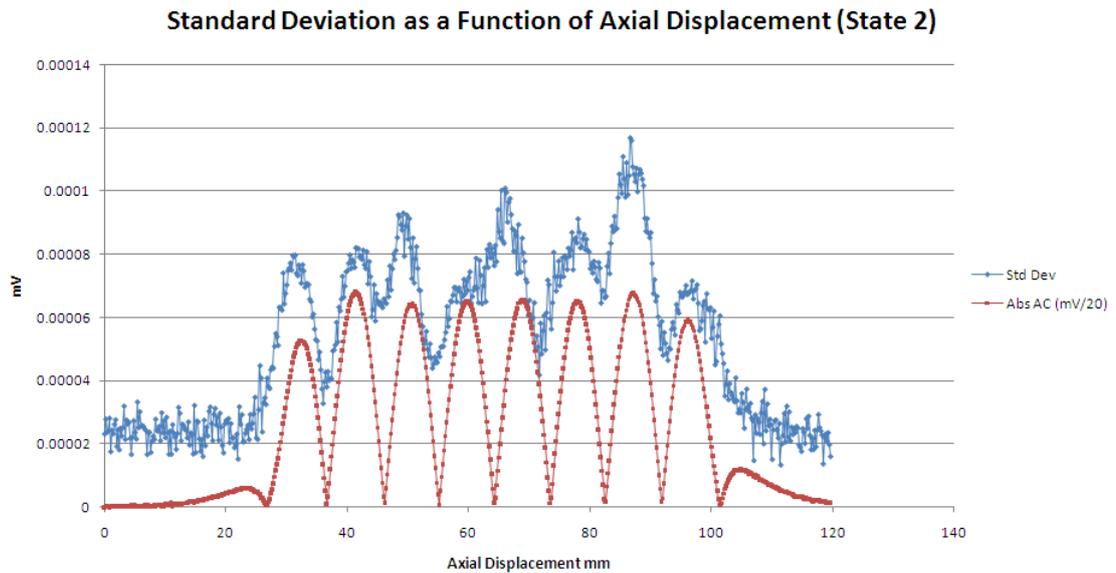


Figure 10.15: The standard deviation of the azimuthal data sets for ‘State 1’ is plotted as a function of axial displacement. The absolute of the average value of the azimuthal data sets is also plotted (offset corrected) but note that it has been scaled by a factor of 1/20 to better fit the graph. The correlation indicates that the standard deviation is giving a measure of either the real field asymmetry and/or there is a positional error in the measurement.

and some interesting conclusions were drawn. If this mode of oscillation was excited to the point where the extreme of the magnet deflection is allowed, given the approximate clearance in the bearings of  $100\ \mu\text{m}$  and the available clearance in the stator bore tube of  $1\ \text{mm}$ , it was estimated that this could momentarily place a lateral force of up to  $20\ \text{N}$  on the lower bearing during actuation. Additionally because the lower section of the shaft is now not running parallel to the bearing surface this load would be placed between the shaft and the top edge of the bearing, significantly increasing the pressure on the titanium shaft and the DLC coating.

Both the recordings from the high speed camera and the lack of any scoring inside the stator bore tube demonstrated that the extreme case presented in the report was, as deliberately expected, exaggerating any perceived effect. It should also be noted that the bearing clearances used in the FEA analysis are close to double the actual bearing clearances of  $\approx 50\ \mu\text{m}$ . However the high speed camera did indicate that a target shaft deflection of between one-third to one-fifth of the value indicated in the report was occurring during actuation. Given the small angle between the shaft and the bearings when the shaft is under deflection the lateral force on the bearing should approximately scale to the deflection and so it is reasonable to assume that the bottom shaft/bearing interface is seeing a force of several Newtons during actuation due to this resonant mode.

Phase	Hall Position (mm)	$\sigma$ Max (mT)	$dB/dR$ (mT mm <sup>-1</sup> )	Max Field Displacement ( $\mu\text{m}$ )
State 1	7	$7.5 \times 10^{-3}$	3.8	200
State 2	7	$1.2 \times 10^{-2}$	3.9	300

Table 10.2: A summary of the correlation between the measured standard deviation in the azimuthal data, the rate of change of magnetic field as a function of radius at an axial field peak and the resultant estimated field asymmetry in microns.



Figure 10.16: If it is assumed that the magnets are fixed and that there is sufficient clearance in the bearings then the lower part of the shaft will have a resonance at 27.6 Hz[120].

Although not discussed in the FEA report, the implications of it along with other relevant material that has largely been indicated in this chapter were discussed in subsequent target workshop meetings. It was clear that the target shaft would benefit from a redesign to make it stiffer, thereby pushing the resonant mode above the ‘equivalent’ operating frequency of the mechanism. The bearings also needed a redesign to remove the sharp edges so that they do not act as a knife-edge on the target shaft during actuation; there were also a couple of other subtle issues with the bearing design that needed addressing and these will be discussed in the next section. Given the requirement for a redesign of several key components of the target mechanism this provided an

opportunity for reconsideration of the overall mechanical design; this was subsequently undertaken by a mechanical engineer at RAL.

## 10.6 Target Mechanism Modifications

This section will give some details of the modifications to the target mechanism that are being undertaken. As of writing this chapter no operational experience has been gained with this improved design. Much of what is to be described in this section can be found in a presentation given to the MICE collaboration meeting[121], further technical drawings can be found on the web[122].

### 10.6.1 Target Shaft

The target shaft has been extensively revised and a diagram of the new design can be seen in figure 10.17[121]. The redesigned target shaft will be manufactured from two separate parts. The upper part of the shaft will be a solid titanium rod of 4 mm diameter and the lower part of the shaft will be a titanium tube with 5.95 mm outer diameter and a 4.55 mm inner diameter. The change from a cruciform cross section to a tubular cross-section increases the second moment of area from  $1.8 \times 10^{-11} \text{ m}^4$  to  $4 \times 10^{-11} \text{ m}^4$  doubling the stiffness of the shaft. This also changes its fundamental frequency from  $\approx 27 \text{ Hz}$  to  $\approx 50 \text{ Hz}$ [121].

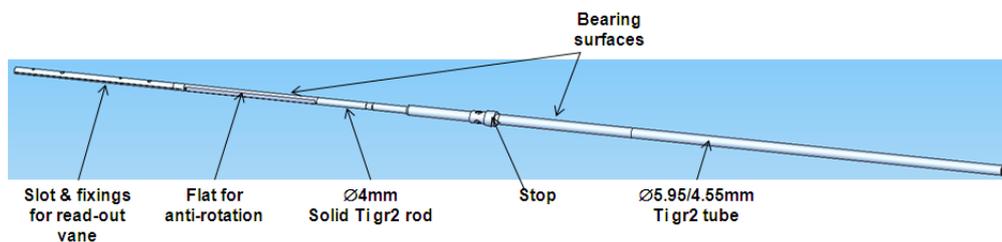


Figure 10.17: The new shaft design is composed from two separate pieces of titanium. The anti-rotation device is now part of the upper bearing. The bottom section of the lower tube (RHS) acts as the target[121].

The target blade seen on the cruciform cross-section has now gone. A study done on the shape of the target blade showed that for a given intersecting mass the shape of the target blade had little impact upon the beamloss distribution[123]. Consequently the bottom section of the tubular shaft will now intercept the ISIS beam; therefore significantly simplifying the design. Anti-rotation of the shaft is now provided by the use of a rotation stop that runs against one of two flat section on the upper shaft. The rotation stop is part of the upper bearing assembly and will be described in the bearing section. The main body assembly has been shortened by 41 mm, whereas the

shaft has been shortened by only 31 mm. This should allow the shaft to penetrate 10 mm further into the ISIS beam, increasing beamloss and improving the muon rates for MICE, whilst the shorter shaft contributes to the stiffness.

Both of the shaft sections will be polished to provide the necessary surface finish for good adhesion of the DLC. The two halves are push-fitted and then welded together. Great care will be taken to ensure that the two halves are well aligned so that the shaft remains straight along its length. Metrology of the components will be undertaken during all stages of manufacture to ensure that the shaft components are manufactured to within the tolerances stated in the design.

The vane design has also been altered to better protect the fragile fingers on the comb side. A solid section now joins the outer edges of the comb together. The change in topology means that the comb detail can no longer be manufactured using wire erosion and a chemical etching technique is used instead.

### 10.6.2 Target Bearings

The design of the target bearings can be seen in figure 10.18[121]. The main features are the circular cross section and that the anti-rotation feature had been moved to the upper bearing. The bearings are manufactured from a single piece of steel with a doweled removable wedge. The move towards a single piece was prompted by the requirement to ensure that the shaft and bearings were aligned within the stator to an accurate tolerance and to reduce the number of edges on the bearing/shaft interface. There was a concern that if the previous bearings were not accurately housed (to within microns) that the edges created by the split design of the bearings could be allowing an uneven load on the bearing surfaces therefore increasing the rate of bearing wear.

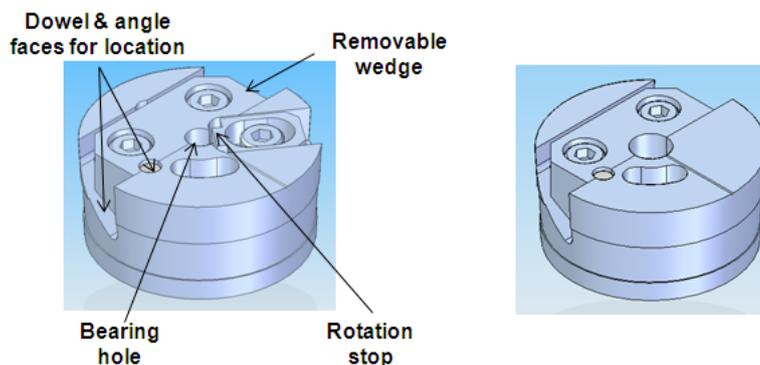


Figure 10.18: Left: The top bearing is built from a single piece with the addition of a wedge and a rotation stop. The wedge allows the inner surfaces to be coated with DLC as the process must have line of sight. Right: The bottom bearing is similar to the top bearing although the bearing aperture is larger and there is no rotation stop[121].

The wedge section in the new design allows the bearing surfaces to be properly coated with DLC as it opens the bearings surfaces up when disassembled. The bore of the bearing is machined with the wedge fixed in situ to ensure that the bearing features are accurately aligned; the bearing surface is 5 mm deep. Sharp edges are bevelled so that the shaft cannot scrape on the edge of the bearing if it is pulled out of alignment due to lateral forces on the shaft or magnet assembly. The anti-rotation feature is now an adjustable stop in the top bearing. This stop bears up to the flat surface on the upper part of the shaft thereby preventing its rotation.

The clearances in the bearings have been reduced to the maximum extent possible,  $\approx 10 \mu\text{m}$ , given the tolerances on the shaft components; this will help to reduce the rattle on the DLC surface. With the clearances reduced to this extent it is imperative that the bearing locations are accurately aligned with respect to one another and the centre of the coil stack. One of the main reasons why the stator body has undergone significant revision is to permit and maintain this accurate alignment of these bearings. These bearings could be easily be modified to allow a suitable boss insert to be used. This would allow other bearing materials to be tried in the future, e.g. radiation hard plastics, to see if they will have a longer life than the DLC.

### 10.6.3 Stator Body

The stator body as described in chapter 5 is composed of six interfaces and five components. To get the bearings accurately aligned it was seen as necessary to reduce the number of interfaces and redesign the body making the whole structure more compact and solid. A solid design was seen to be beneficial as it would ensure that the bearings would remain aligned after assembly taking into account handling and installation of the stator. The redesign also had to allow the coil stack and the cooling system to be built and tested independently; the design of these components has not been altered in any way. Figure 10.19 shows the main features of the redesigned stator body.

## 10.7 Conclusion

The target mechanism is a functioning piece of equipment that is being used to produce muons for the MICE experiment. It is a unique piece of equipment, and this uniqueness has meant that there have been various hurdles to overcome to improve the long term reliability of the mechanism. Some of this work is ongoing; this is demonstrated by the material presented in this chapter.

Electrically the target has performed flawlessly. As of writing there has not been another coils failure since both the quality assurance procedures on the coils were improved and the power supply altered to reduce the voltage stress. The controller has performed tens of millions of actuations in a consistently reliable manner. The controller will be upgraded to an FPGA based system, but primarily this is to improve

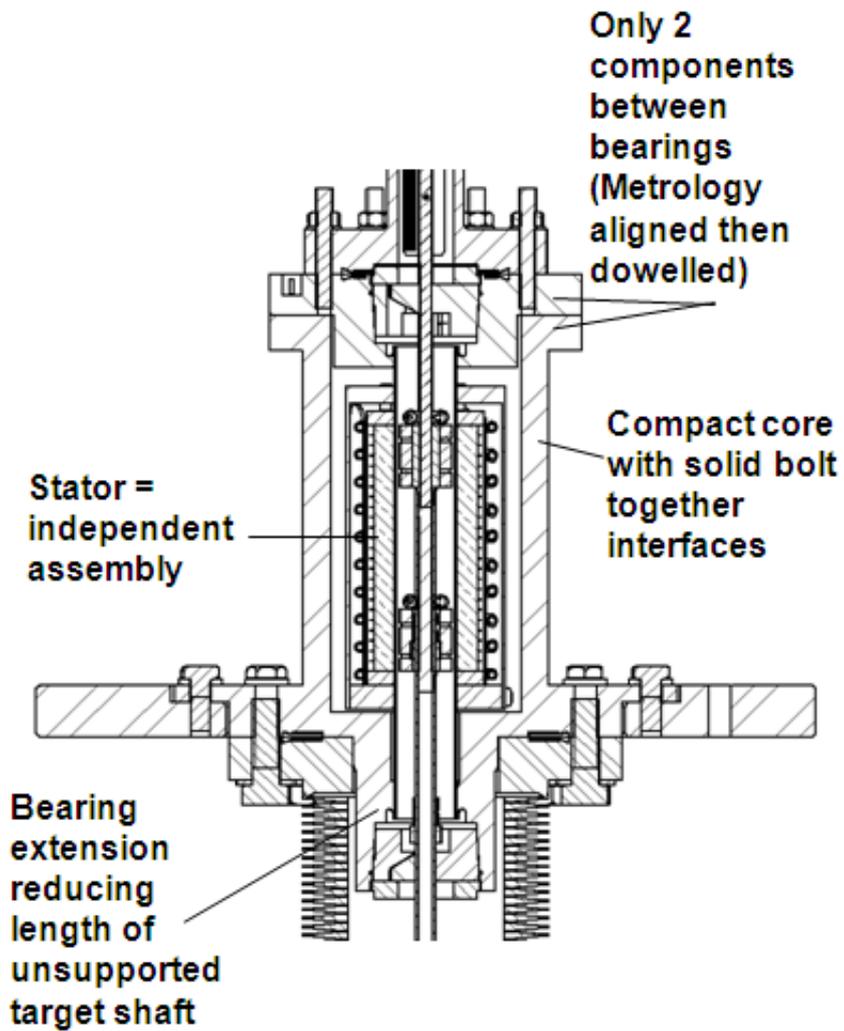


Figure 10.19: The new body has been designed to strengthen and simplify the current design and will allow the bearings to be accurately located[121].

the user interface and to add more safety interlocks, the underlying control logic will remain the same.

## Chapter 11

# Conclusion

This thesis has described the development and operation of the target mechanism, a vital piece of equipment for generating pions for the Muon Ionisation Cooling Experiment (MICE). The MICE experiment is part of a program to build a high intensity neutrino source called the neutrino factory. The large flux of neutrinos produced by such a machine would allow measurements of the neutrino mixing matrix to a precision never before achieved. This machine could also open the possibility of detecting CP violation and establishing the mass hierarchy of the neutrino mass states.

Chapter one of this thesis considered the background to neutrino physics, discussed several of the key experiments undertaken that have helped to parameterise the neutrino mixing matrix and finally illustrated the need for a high intensity neutrino source to further the understanding of these particles. The neutrino factory is considered by many to be the best proposed facility to enable precision measurement on the neutrino mixing matrix to be made and a discussion of this proposed facility was presented in chapter two. In order to achieve the required neutrino intensity the neutrino factory will require several new technologies, one of which is ionisation cooling. Ionisation cooling reduces the emittance of a muon beam which would enable a significant increase in the acceptance of muons into the neutrino factory muon accelerator thereby creating a higher flux of neutrinos. The MICE experiment has been set the task of proving that muon ionisation cooling is a practical technology. Both the principle of emittance reduction and an overview of the MICE experiment were considered in chapter three.

The rest of the thesis described the development of the target mechanism, a high acceleration electromagnetic linear motor that is situated on the ISIS synchrotron. This device operates parasitically on ISIS inserting a small titanium target into the halo of the ISIS beam for the last 2 ms before beam extraction. This creates pions that are captured by the MICE beamline and these subsequently decay into muons for use by the MICE experiment.

Chapter four starts by giving an overview of the requirements for the MICE target mechanism and describes why a linear motor was chosen to perform the task. Details

of some early developmental work that shaped the design of the target mechanism were also given. Chapter five then describes the details of the mechanical construction of the first target mechanism that was installed on ISIS in 2008. The design of an optical position monitoring system, essential for controlling the position of the target, was discussed in chapter six. The dynamics of the motor in terms of the coil-magnet relationship was discussed in chapter seven showing how the motor can be controlled to produce the necessary target motion. Control of the motor requires the switching of currents through its three-phase coil structure. As the motor power supply provides both the direction and magnitude of coil current to the stator the target power supply was also discussed in this chapter. The stator power supply is controlled via an electronic controller. The target motion requires positional feedback and chapter eight discussed the electronic circuitry that was built to control the operation of the target mechanism. The monitoring system, the target DAQ, was also discussed in this chapter.

Target tests were performed in November 2006 before the first proper target system was installed on the ISIS beam-pipe in January of 2008. Chapter nine gave details of what was learnt from the November 2006 tests and also detailed the operational performance of the stator both in the laboratory in Sheffield and in situ on the ISIS beam-pipe. As particle production for MICE is the primary aim of the target mechanism some initial data from the use of the target mechanism for this purpose was analysed and the conclusions from this data were considered.

Finally chapter ten details some of the more challenging problems that have been encountered from operation of the target mechanism at the Rutherford Appleton Laboratory. The two main problems with the mechanism have been electrical breakdown of coil insulation and excessive bearing wear. The former problem has apparently been resolved. It is hoped that shaft and bearing redesign resulting from studies by the author and the collaboration will resolve the wear issue.

## 11.1 Outlook

The target design has succeeded in providing a mechanism that can successfully place a target into the ISIS beam during the last 2 ms of the spill and withdraw it before the next spill is injected into the synchrotron. Prolonged operation at 1 Hz has also been established. Although this demonstrates that the primary purpose of the target mechanism has been realised there is still a significant amount of work that has yet to be accomplished. This work can broadly be broken into three categories, understanding the particle production data, continuing to improve the mechanical reliability of the target mechanism and upgrading the control electronics to allow more automated functionality. The DAQ system on the target will also be upgraded as part of the control upgrade; the aim of the DAQ upgrade will be to improve the data quality.

The target mechanism is a functioning piece of equipment that is currently being used

to produce muons for the MICE experiment. Electrically the target is performing flawlessly. Finally, as of writing, the target controller has reliably actuated the target for tens of millions of times in a consistently repeatable manner, fulfilling the control specification that was first described in 2003.



# Bibliography

- [1] W. Pauli. Letter proposing the neutrino (1930). *Physics Today*, September 1978.
- [2] H. A. Bethe and R. E. Peierls. The “neutrino”. *Nature*, 133:532, April 1934.
- [3] C. L. Cowan et al. Detection of the free neutrino: a confirmation. *Science*, pages 103–104, 1956.
- [4] G. Danby. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36, 1962.
- [5] M. L. Perl et al. Evidence for anomalous lepton production in  $e^+e^-$  annihilation. *Phys. Rev. Lett.*, 35:1489, 1975.
- [6] K. Kodoma et al. Observation of tau neutrino interactions. *Phys. Lett. B*, 504:218, 2001.
- [7] L3 Collaboration. Determination of the number of light neutrino species from single photon production at LEP. *Phys. Lett. B*, 431:199-208, 1998.
- [8] D. Buskulic et al. A direct measurement of the invisible width of the  $Z$  from single photon counting. *Phys. Lett.*, B313:520–534, 1993.
- [9] R. Akers et al. Measurement of single photon production in  $e^+e^-$  collisions near the  $Z^0$  resonance. *Z. Phys.*, C65:47–66, 1995.
- [10] P. Abreu, W. Adam, T. Adye, E. Agasi, and I. Ajinenko et al. Search for new phenomena using single photon events at LEP1. *Z. Phys.*, C74:577–586, 1997.
- [11] R. Davis, D. S. Harmer, and K. C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205-1209, 1968.
- [12] C. Amsler et al. Review of particle physics. *Phys. Lett.*, B667:1, 2008.
- [13] B. Pontecorvo. Mesonium and antimesonium. *Sov. Phys. JETP*, 6:429, 1957.
- [14] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870, 1962.
- [15] V. N. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Phys. Lett.*, B28:493, 1969.

- [16] S. P. Mikheev and A. Yu. Smirnov. Neutrino oscillations in an inhomogeneous medium: adiabatic regime. *Sov. Phys. JETP*, 65, 1987.
- [17] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D.*, 17:2369–2374, 1978.
- [18] MICE Collaboration. MICE, an international Muon Ionisation Cooling Experiment: proposal to the Rutherford Appleton Laboratory, submitted to CCLRC and PPARC on the 10th January 2003.
- [19] Jean-Michel Levy. On the derivation of the neutrino oscillation length formula. *hep-ph/0004221:arXiv*, 2000.
- [20] S. Ozaki Ed., R. Palmer, M. Zisman, and J. Gallardo. Feasibility study II of a muon-based neutrino source. *BNL-52623*, 2001.
- [21] F. Boehm and P. Vogel. *Physics of Massive Neutrinos*. Cambridge University Press, second edition, 1992.
- [22] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. New solar opacities, abundances, helioseismology, and neutrino fluxes. *Astrophys. J.*, 621:L85–L88, 2005.
- [23] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate by the Russian-American gallium solar neutrino experiment during one half of the 22-year cycle of solar activity. *J. Exp. Theor. Phys.*, 95:181–193, 2002.
- [24] GALLEX collaboration P. Anselmann et al. GALLEX solar neutrino observations: complete results for GALLEX II. *Phys. Lett. B*, 357:237, 1995.
- [25] Sudbury Neutrino Observatory SNO homepage.  
<http://www.sno.phy.queensu.ca/>.
- [26] SNO Collaboration. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301.
- [27] SNO Collaboration. An independent measurement of the total active  $^8\text{B}$  solar neutrino flux using an array of  $^3\text{He}$  proportional counters at the Sudbury Neutrino Observatory, 2008.  
<http://www.citebase.org/abstract?id=oai:arXiv.org:0806.0989>.
- [28] Super-Kamiokande (Super-K) homepage.  
<http://neutrino.phys.washington.edu/superk/index.html>.
- [29] S. Hatakeyama et al. Measurement of the flux and zenith angle distribution of upward through-going muons in Kamiokande II + III. *Phys. Rev. Lett.*, 81:2016–2019, 1998.
- [30] Y. Ashie et al. A measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I. *Phys. Rev.*, D71:112005, 2005.

- [31] The CHOOZ proposal. [http://duphy4.physics.drexel.edu/chooz\\_pub/](http://duphy4.physics.drexel.edu/chooz_pub/).
- [32] K. Anderson et al. White paper report on using nuclear reactors to search for a value of  $\theta_{13}$ , 2004.  
<http://www.neutrinooscillation.org/white.html>.
- [33] F. Ardellier et al. Double CHOOZ: A search for the neutrino mixing angle  $\theta_{13}$ . *hep-ex/0606025:arXiv*, 2006.
- [34] KamLAND home page.  
<http://kamland.lbl.gov/>.
- [35] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti- neutrino disappearance. *Phys. Rev. Lett.*, 90:021802, 2003.
- [36] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Getting the most from the statistical analysis of solar neutrino oscillations. *Phys. Rev. D*, 66(5):053010, Sep 2002.
- [37] MINOS homepage:  
<http://www-numi.fnal.gov/>.
- [38] P. Adamson et al. A study of muon neutrino disappearance using the Fermilab main injector neutrino beam. *Phys. Rev.*, D77:072002, 2008.
- [39] K2K homepage:  
<http://neutrino.kek.jp/>.
- [40] M. H. Ahn et al. Measurement of neutrino oscillation by the K2K experiment. *Phys. Rev.*, D74:072003, 2006.
- [41] A. Aguilar et al. Evidence for neutrino oscillations from the observation of  $\bar{\nu}_e$  appearance in a  $\bar{\nu}_\mu$  beam. *Phys. Rev.*, D64:112007, 2001.
- [42] A. A. Aguilar-Arevalo et al. A search for electron neutrino appearance at the  $\Delta m^2 \sim 1\text{eV}^2$  scale. *Phys. Rev. Lett.*, 98:231801, 2007.
- [43] W. M. Yao et al. Review of particle physics. *J. Phys.*, G33:1–1232, 2006.
- [44] S. Abe et al. Precision measurement of neutrino oscillation parameters with KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.
- [45] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillation. *Phys. Rev. Lett.*, 93:101801, 2004.
- [46] S. Eidelman et al. Review of particle physics. *Phys. Lett.*, B592:1, 2004.
- [47] Neutrino physics.  
<http://hitoshi.berkeley.edu/neutrino/>.
- [48] A. Baldini Ed. et al. Beams for European neutrino experiments (BENE) midterm scientific report. *CARE-2006-009 BENE CERN-2006-005 ECFA/06/242*, May 2006.

- [49] A feasibility study of a neutrino factory in Japan, 2001.  
<http://www-prism.kek.jp/nufactj/nufactj.pdf>.
- [50] S. Geer Ed. and M. Zisman. Neutrino factory and beta beam experiments and development (feasibility study IIa). *Technical Report BNL-72369-2004*, 2004.
- [51] J. S. Berg et al. A cost effective design for a neutrino factory (feasibility study IIb). *MU-COOL Note 327*, 2005.
- [52] A. Bandyopadhyay et al. Physics at a future neutrino factory and super-beam facility. *0710.4947:arXiv*, 2007.
- [53] The ISS Accelerator Working Group. Accelerator design concept for future neutrino facilities. *0802.4023:arXiv*, 2008.
- [54] T. Abe, H. Aihara, C. Andreopoulos, A. Ankowski, A. Badertscher, G. Battistoni, A. Blondel, J. Bouchez, A. Bross, A. Bueno, L. Camilleri, J. E. Campaigne, A. Cazes, A. Cervera-Villanueva, G. De Lellis, F. Di Capua, M. Ellis, A. Ereditato, L. S. Esposito, C. Fukushima, E. Gschwendtner, J. J. Gomez-Cadenas, M. Iwasaki, K. Kaneyuki, Y. Karadzhov, V. Kashikhin, Y. Kawai, M. Komatsu, E. Kozlovskaya, Y. Kudenko, A. Kusaka, H. Kyushima, A. Longhin, A. Marchionni, A. Marotta, C. McGrew, S. Menary, A. Meregaglia, M. Mezzeto, P. Migliozzi, N. K. Mondal, C. Montanari, T. Nakadaira, M. Nakamura, H. Nakumo, H. Nakayama, J. Nelson, J. Nowak, S. Ogawa, J. Peltoniemi, A. Pla-Dalmau, S. Ragazzi, A. Rubbia, F. Sanchez, J. Sarkamo, O. Sato, M. Selvi, H. Shibuya, M. Shozawa, J. Sobczyk, F. J. P. Soler, P. Strolin, M. Suyama, M. Tanak, F. Terranova, R. Tsenov, Y. Uchida, A. Weber, and A. Zlobin. Detectors and flux instrumentation for future neutrino facilities. *0712.4129:arXiv*, 2007.
- [55] International scoping study of a future neutrino factory and super-beam facility. Website can be found at:  
<http://www.hep.ph.ic.ac.uk/iss/>.
- [56] K. A. Walaron. *Neutrino Factory Targets and the MICE beam*. PhD thesis, University of Glasgow, 2007.
- [57] The neutrino factory and muon collider collaboration.  
<http://www.cap.bnl.gov/mumu/>.
- [58] I. Efthymiopoulos et al. The MERIT(nTOF-11) high intensity liquid mercury target experiment at the CERN PS. *EPAC 08 Conference Proceedings*, CERN EDMS Id 929446, 2008.
- [59] K. Hanke. Muon front-end without cooling, 2000. CERN-Nufact-Note-059.  
<http://slap.web.cern.ch/slap/NuFact/NuFact/NFNotes.html>.
- [60] D. J. Summers, A. A. Garren, J. S. Berg, and R. B. Palmer. A pulsed synchrotron for muon acceleration at a neutrino factory, 2003.  
<http://www.citebase.org/abstract?id=oai:arXiv.org:physics/0310028>.

- [61] Rajendran Raja et al. The program in muon and neutrino physics: Super beams, cold muon beams, neutrino factory and the muon collider. *hep-ex/0108041:arXiv*, 2001.
- [62] A. Cervera et al. Golden measurements at a neutrino factory. *Nucl. Phys.*, B579:17–55, 2000.
- [63] P. Huber, M. Lindner, M. Rolinec, and W. Winter. Optimization of a neutrino factory oscillation experiment. *Phys. Rev.*, D74:073003, 2006.
- [64] K. Wille. The physics of particle accelerators - an introduction. Oxford University Press. Second Edition, 2000. ISBN 0-19-850549-3.
- [65] MICE Collaboration. MICE technical reference document, 2004.  
[http://www.isis.rl.ac.uk/accelerator/MICE/TR/MICE\\_Tech\\_ref.html](http://www.isis.rl.ac.uk/accelerator/MICE/TR/MICE_Tech_ref.html).
- [66] J. Thomason. Upgrades to ISIS for the New Second Target Station. *Proc. 11th European Particle Accelerator Conference (EPAC '08), Genoa, Italy*, pages 2902 – 2906, 2008.
- [67] R. E. Williamson, B. G. Pine, S. J. Payne, and C. M. Warsop. Analysis of measurement errors in residual gas ionisation profile monitors in a high intensity proton beam. *Proc. 11th European Particle Accelerator Conference (EPAC '08), Genoa, Italy*, 2008.
- [68] S. J. Payne, P. G. Barnes, G. M. Cross, A. Pertica, and S. Whitehead. A self calibrating real time multi-channel profile monitor for the ISIS proton synchrotron. *Proc. 8th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators (DIPAC 2007), Venice, Italy*, 2007.
- [69] Private communication with Dean Adams at the ISIS accelerator division, Rutherford Appleton Laboratory, UK.
- [70] S. Yang W.Lau. Technical drawings of MICE and associated components. Oxford University.  
[http://www.physics.ox.ac.uk/design/MICE/mice\\_page.htm](http://www.physics.ox.ac.uk/design/MICE/mice_page.htm).
- [71] A. Aloisio et al. The KLOE detector: technical proposal. (Laboratori Nazionali di Frascati Report Number: LNF-93-002-IR. A preprint can be obtained from the CERN Document Server  
<http://cdsweb.cern.ch/>.
- [72] A. Khan et al. MICE scintillating-fibre tracker. MICE note 90.  
<http://mice.iit.edu/micenotes/public/pdf/MICE0090/MICE0090.pdf>.
- [73] D. E. Baynham, P. Bish, T. Bradshaw, M. A. Cummings, M. A. Green, S. Ishimoto, I. Ivaniouchenkov, W. Lau, S. Q. Yang, and M. S. Zisman. A liquid cryogen absorber for MICE. MICE note 128  
<http://mice.iit.edu/micenotes/public/pdf/MICE0128/MICE0128.pdf>.

- [74] A. Moretti, Z. Qian, J. Norem, Y. Torun, D. Li, and M. Zisman. Effects of high solenoidal magnetic fields on rf accelerating cavities. *Phys. Rev. ST Accel. Beams*, 8(7):072001, Jul 2005.
- [75] P.S. Flowers and R. Edgecock. Notes on the old PPD beam extraction target, July 2007. A collection of personal notes, diagrams and photos provided to P.Smith on various aspects of the old HEP target construction and running.
- [76] C. Booth. Private communication.
- [77] T. Roberts. MICE Target Source Calculations. Presented to the MICE Collaboration Meeting, 30th March 2004  
[http://www.mice.iit.edu/cm/cm8/cm8\\_roberts\\_target.ppt](http://www.mice.iit.edu/cm/cm8/cm8_roberts_target.ppt).
- [78] T. Roberts. MICE beamline analysis TRD Sept 2004. Presented to MICE Collaboration Meeting, 10th Feb 2005.  
[http://www.mice.iit.edu/cm/cm11/cm11\\_roberts\\_g4beamline.ppt](http://www.mice.iit.edu/cm/cm11/cm11_roberts_g4beamline.ppt).
- [79] C. Booth. Target specifications and history. Talk given to an internal MICE meeting at the University of Sheffield Dec 2009  
<http://mice.iit.edu/bl/Target/Meetings/2009-12-02/06.2-1-2009-12-02-Booth.ppt>.
- [80] G. Heidenreich. Carbon and Beryllium Targets at PSI. *American Institute of Physics Conference Series*, 642:122–124, December 2002.
- [81] J P Marriner M D Church. The Antiproton Sources: Design and Operation. *Annual Review of Nuclear and Particle Science*, 43:253–295, 1993.
- [82] I.R. Bailey et al. Development of a Positron Production Target for the ILC Positron Source. *Proc. 10th European Particle Accelerator Conference (EPAC '06), Edinburgh, UK*, pages 2484 – 2486, 2006.
- [83] Rigaku Vacuum Products  
<http://www.rigaku.com/vacuum/uhv.html>.
- [84] UHV Design  
<http://www.uhvdesign.com/>.
- [85] Private communication with Sven Soechting at the Rutherford Appleton Laboratory, UK.
- [86] G Schweitzer. Active Magnetic Bearings - Chances and Limitations. *Proc. 6th Internat. IFToMM Conf. on Rotor Dynamics, Sydney, 2002*.
- [87] Eric Maslen. Magnetic Bearings, 2000. Lecture Notes on Magnetic Bearings - Department of Mechanical, Aerospace, and Nuclear Engineering Charlottesville, Virginia.

- [88] J.J. Feeley and D.J. Ahlstrom. A New Eddy Current Model for Magnetic Bearing Control System Design. *The 1992 4th NASA SERC Symposium on VLSI Design 8 (SEE N94-21694 05-33)*, 1992. Provided by the SAO/NASA Astrophysics Data System.
- [89] N. Schofield. Power Conversion Group: A6C Sackville St Building, School of Electrical and Electronic Engineering, The University of Manchester.
- [90] N. Schofield, C. Booth, and P. J. Smith. A low mass, brushless permanent magnet linear actuator for the ISIS target accelerator. *50th Annual Conference on Magnetism and Magnetic Materials (MMM05)*, (Paper FF-12), October 2005.
- [91] Tom Bradshaw, Rutherford Appleton Laboratory, UK.
- [92] M. Mohammad and R. Nicholson. CAD drawings of the target components. All these drawings can be found on the Sheffield MICE webpage at:  
<http://www.mice.group.shef.ac.uk/drawings.php>.
- [93] P.A.R. Insulations and Wires Ltd. Technical data MAGNETEMP CA-200 enamelled copper winding wires. Datasheet on the properties of the wire used in the manufacture of the stator coils.  
<http://www.par.gb.com>.
- [94] P.A.R. Insulations and Wires Ltd. Technical data DOLPHON-2102 one pack low emissions polyester impregnating resin clear. Datasheet on the properties of the varnish used in the manufacture of the stator coils.  
<http://www.par.gb.com>.
- [95] L. Howlett. Simulation of radiation levels in the MICE target magnets, March 2007. MICE note 166  
<http://mice.iit.edu/micenotes/public/pdf/MICE0166/MICE0166.pdf>.
- [96] C. Booth. Beam heating of the MICE target. MICE note 206  
<http://mice.iit.edu/micenotes/public/pdf/MICE0206/MICE0206.pdf>.
- [97] R. Dwyer-Joyce. Tribology Department, University of Sheffield, UK., 2007. Meeting to discuss alternative bearing materials and techniques for the MICE target mechanism.
- [98] Ansoft. MAXWELL SV. The student version of MAXWELL 2D. A feature limited electromagnetic and electrostatic field simulation software package using finite element analysis.  
<http://www.ansoft.com/maxwellsv/>.
- [99] Information on the magnetic properties of NdFeB Magnets was obtained from the Energy Conversion Systems website.  
<http://www.kanemagnetics.de/products/pdf/magnets/NeomagS.pdf>.
- [100] E. McCarron et al. Technical drawings courtesy of Rutherford Appleton Laboratory.

- [101] NUFERN 630 nm select cut off single mode fibre. Graph taken from a datasheet provided by Thorlabs on request.
- [102] 0.37 NA hard polymer clad multimode fiber. Datasheet provided by Thorlabs.  
<http://www.thorlabs.com/Thorcat/12200/12255-S01.pdf>.
- [103] H3R880IR - FDH3 housing with FDR880IR pin photodiode. Datasheet obtained from Farnell.  
<http://www.farnell.com/datasheets/99561.pdf>.
- [104] P. C. Todd. Snubber circuits: Theory, design and application, 1993. Texas Instruments online archive.  
[focus.ti.com/lit/an/slup100/slup100.pdf](http://focus.ti.com/lit/an/slup100/slup100.pdf).
- [105] S. Griffiths, J. Theed, J. Cartledge, I. Mullacrane, G. Charnley, et al. Target power supply development and installation. Electrical Engineering Dept at Daresbury Laboratory, UK.
- [106] J. S. Graulich. MICE data acquisition terminology, 2006. MICE note 147.  
<http://mice.iit.edu/micenotes/public/pdf/MICE0147/MICE0147.pdf>.
- [107] A significant amount of documentation on the PIC microprocessors can be found on the Microchip website:  
<http://www.microchip.com>.
- [108] M. A. Clarke-Gayther, A. I. Borden, and G. M. Allen. Global beam loss monitoring using long ionisation chambers at ISIS. *Proc. European Particle Accelerator Conference (EPAC '94)*, pages 1634 – 1636, 1994.
- [109] L. Howlett. Analysis of target beam test data, February 2007. MICE note 165  
<http://mice.iit.edu/micenotes/public/pdf/MICE0165/MICE0165.pdf>.
- [110] F. J. P. Soler, K. Walaron, C. Booth, M. Carson, P. Hodgson, L. Howlett, P. Smith, D. Adams, R. Edgecock, W. Murray, K. Tilley, J. Cobb, M. Rayner, and T. Roberts. Measurement of particle production from the MICE target. *American Institute of Physics Conference Series*, 981:259–261, 2008.
- [111] K. Walaron. Measurement of particle production from the MICE target. Talk given to the MICE collaboration meeting at RAL June 2007  
[http://www.mice.iit.edu/cm/cm18/cm18\\_walaron\\_measurementofparticleproduction.ppt](http://www.mice.iit.edu/cm/cm18/cm18_walaron_measurementofparticleproduction.ppt).
- [112] D. Adams, M. Apollonio, A. Dobbs, P. Hodgson, C. MacWaters, J. Leaver, K. Long, H. Sakamoto, P. Smith, K. Tilley, V. Verguilov, K. Walaron, and D. Wright. Demonstration of parasitic running of the MICE target, May 2008. MICE note 205.  
<http://mice.iit.edu/micenotes/public/pdf/MICE0205/MICE0205.pdf>.
- [113] Resistivity - from Wikipedia, the free encyclopedia. Coefficient of resistivity for copper.  
<http://en.wikipedia.org/wiki/Resistivity>.

- [114] C. Booth, P. Hodgson, K. Long, E. McCarron, A. Nichols, R. Nicholson, and P. Smith. Position of MICE target number one in the ISIS beam pipe, July 2008. MICE note 212.  
<http://mice.iit.edu/micenotes/public/pdf/MICE0212/MICE0212.pdf>.
- [115] M. Apollonio and K. Tilley. Estimate of protons on target and implications for muon flux in MICE, April 2008. MICE note 204.  
<http://mice.iit.edu/micenotes/public/pdf/MICE0204/MICE0204.pdf>.
- [116] DuPont Kapton HN POLYIMIDE FILM. Technical Data Sheet  
[http://www2.dupont.com/Kapton/en\\_US/assets/downloads/pdf/HN\\_datasheet.pdf](http://www2.dupont.com/Kapton/en_US/assets/downloads/pdf/HN_datasheet.pdf).
- [117] S. Griffiths. Powerpoint slides presented to Target Workshop meeting on 8th January 2009 detailing changes to be made to the target electrical system.  
<http://mice.iit.edu/bl/Target/Meetings/2009-01-08>.
- [118] J. Tarrant. MICE Target Mechanical Design. Powerpoint slides presented to a Target Workshop meeting on 8th Jan 2009.  
<http://mice.iit.edu/bl/Target/Meetings/2009-01-08/>.
- [119] Opera 3D. Electromagnetic Finite Element Analysis Software.  
<http://www.vectorfields.com>.
- [120] Stephanie Yang. Interim Report on the Target Shaft Mechanical Analysis and Assessment – FEA. Initially presented to a Target Workshop meeting in October 2008.  
<http://mice.iit.edu/bl/Target/Notes/MICE>.
- [121] J. Tarrant. MICE Target Mechanical Design. Powerpoint slides presented to the MICE Collaboration Meeting 24 on the 1st June 2009.  
[http://www.mice.iit.edu/cm/cm24/agenda/cm24\\_tarrant\\_targetreadiness.ppt](http://www.mice.iit.edu/cm/cm24/agenda/cm24_tarrant_targetreadiness.ppt).
- [122] J. Tarrant. Technical drawings of the new stator, shaft and bearing designs.  
[ftp://ftp.eng.rl.ac.uk/pub/jason\\_tarrant/ExportForManufacture/](ftp://ftp.eng.rl.ac.uk/pub/jason_tarrant/ExportForManufacture/).
- [123] M. Apollonio and A. Dobbs. Target/Beam Interaction. Powerpoint slides presented to the Target Workshop at Imperial College on 27th November 2008.  
[http://www.mice.iit.edu/bl/Target/Meetings/2008-11-27/01-MICE-Beam-Tgt-Sim\\_MA-AD.ppt](http://www.mice.iit.edu/bl/Target/Meetings/2008-11-27/01-MICE-Beam-Tgt-Sim_MA-AD.ppt).