Muon Ionisation Cooling Experiment

MICE Note ???
IC-Pre/10-??
RAL-TR-???
Sheffield-Pre/10-??

September 1, 2011

The design, construction and performance of the MICE target

Abstract

The pion-production target that serves the MICE Muon Beam consists of a titanium cylinder that is dipped into the halo of the ISIS proton beam. The design and construction of the MICE target system is described along with the quality-assurance procedures, electromagnetic drive and control systems, the readout electronics, and the data-acquisition system. The performance of the target is presented together with the particle rates delivered to the MICE Muon Beam. Finally, the beam loss in ISIS generated by the operation of the MICE is evaluated as a function of the particle rate and the operating parameters of the MICE target are derived.

M. Ellis, P.R. Hobson, P. Kyberd, J.J. Nebrensky
Brunel University, Uxbridge, Middlesex, UB8 3PH, UK

G. Barber, D. Clark, I. Clark, R. Hare, A. Jamdagni, J. Leaver, K.R. Long
Department of Physics, Blackett Laboratory, Imperial College London, Exhibition Road, London SW7 2AZ, UK

S. Griffiths, I. Mullacrane
STFC Daresbury Laboratory, Daresbury, Cheshire

D. Adams, T. Bradshaw, E. Capocci, C. MacWaters, E. McCarron, A. Nichols, J. Tarrant, J. Treadgold
STFC Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire, OX11 0QX, UK

C. Booth, P. Hodgson, R. Nicholson, P. Smith
Department of Physics, University of Sheffield, Sheffield, UK

Contents

1 Introduction

Muon storage rings have been proposed for use as sources of intense high-energy neutrino beams in a Neutrino Factory [1] and as the basis for multi-TeV lepton-antilepton colliding-beam facilities [2]. To optimise the performance of such facilities requires the phase-space compression (cooling) of the muon beam prior to acceleration and storage. The short muon-lifetime makes it impossible to employ traditional techniques to cool the beam while maintaining the muon-beam intensity. Ionisation cooling, a process in which the muon beam is passed through a series of liquid-hydrogen absorbers interspersed with accelerating RF cavities, is the technique proposed to cool the muon beam. The international Muon Ionisation Cooling Experiment (MICE) will provide an engineering demonstration of the ionisation-cooling technique and will allow the factors affecting the performance of ionisation-cooling channels to be investigated in detail [3]. Muon beams of momenta between 140 MeV/c and 240 MeV/c, with normalised emittances between 2 πmm and 10 πmm, will be provided by a purpose-built beam line on the 800 MeV proton synchrotron, ISIS [4], at the Rutherford Appleton Laboratory [5].

MICE is a single-particle experiment in which the position and momentum of each muon is measured before it enters the MICE cooling channel and once again after it has left (see figure 1) [6]. The MICE cooling channel, which is based on one lattice cell of the cooling channel described in [7], comprises three 20 l volumes of liquid hydrogen and two sets of four 201 MHz accelerating cavities. Beam transport is achieved by means of a series of superconducting solenoids. A particle-identification (PID) system (scintillator time-of-flight hodoscopes TOF0 and TOF1 and threshold Cherenkov counters CKOVa and CKOVb) upstream of the cooling channel allows a pure muon beam to be selected. Downstream of the cooling channel, a final hodoscope (TOF2) and a calorimeter system allow muon decays to be identified. The calorimeter is composed of a KLOE-like lead-scintillator section (KL) followed by a fully active scintillator detector (the electron-muon ranger, EMR) in which the muons are brought to rest. For a full description of the experiment see [6].

A schematic diagram of the MICE Muon Beam is shown in figure **??**. A cylindrical target is dipped into the edge of the circulating proton beam. The depth at which the target is dipped into the proton beam is characterised by the 'beam centre distance' (BCD) which is defined to be the distance from the tip of the target to the proton-beam axis at the target's maximum excursion into the beam. Pions produced in the target are captured by a quadrupole triplet and transported to a dipole magnet by which the pion momentum is selected. A 5 T superconducting 'decay' solenoid follows the dipole. The additional pion path-length in the decay solenoid increases the muon-production efficiency. Following the solenoid, a second dipole is used to select the muon momentum and the beam is transported to MICE using a pair of large-aperture quadrupole triplets.

This paper is organised as follows. The requirements for the target system and an overview of its design are presented in section 2. Section 3 describes the design of the linear motor. The mechanical design of the target mechanism and the mechanical interface to the ISIS accelerator infrastructure is presented in section 4. Section 5 describes the optical position-measurement

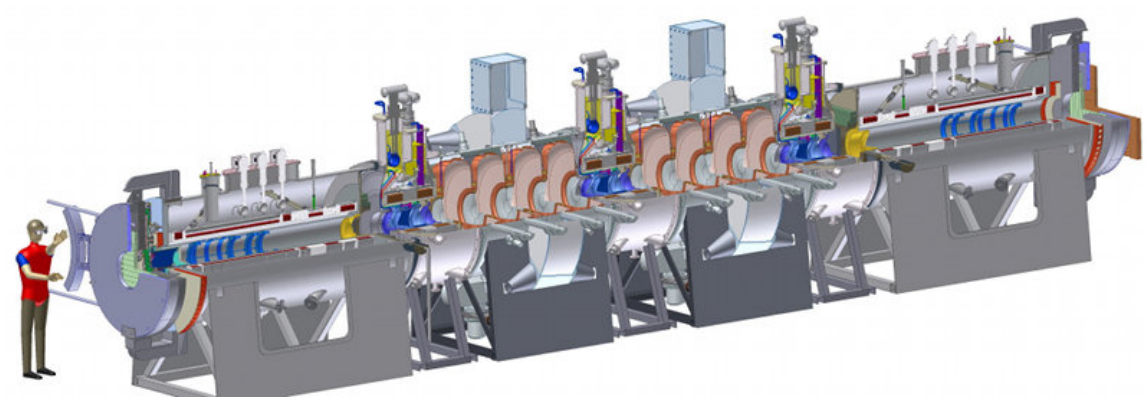

Figure 1: Cutaway 3D rendering of the international Muon Ionisation Cooling Experiment (MICE). The muon beam enters from the bottom left of the figure. The upstream PID instrumentation (not shown) is composed of two time-of-flight hodoscopes (TOF0 and TOF1) and two threshold Cherenkov counters (CKOVa and CKOVb). The upstream spectrometer is followed by the MICE cooling channel, which is composed of three 20 l volumes of liquid hydrogen and two sets of four 201 MHz accelerating cavities embedded in a solenoidal transport channel. This in turn is followed by the downstream spectrometer, a third time-of-flight hodoscope (TOF2), and a calorimeter system (KL and EMR).

system, the control system is described in section 7 and the power electronics used to drive the linear motor is described in section 6. The performance of the system, the particle rate delivered for the MICE Muon Beam and the beam losses induced in ISIS are presented in section 8. Finally, a summary is presented in section 9.

2 Requirements and Overview

Lead author: Chris Booth

The ISIS synchrotron [4] operates on a basic cycle of 50 Hz. Protons are injected with a kinetic energy of 70 MeV and accelerated to 800 MeV over a period of 10 ms prior to extraction. In the following 10 ms, the currents in the focusing and bending magnets are reduced to their initial values, ready for the next pulse of protons to be injected and accelerated.

MICE operation is parasitic to the functioning of ISIS and must cause minimal disruption to its main operation of neutron production. On selected pulses, a small target is caused to dip into the outer low-density halo of the accelerated proton beam just before extraction. Pions produced in the target emerge through a thin window in the ISIS vacuum system towards the muon beam line. On injection, the proton beam effectively fills the beam pipe. At the location of the MICE target, the beam has a vertical radius of 67 mm. During acceleration, the beam shrinks to a radius of about 48 mm. To produce the required muon flux, the target must enter the beam by at least 5 mm, so a minimum travel of 24 mm is needed. (In practice, the exact position of the edge of the beam and the intensity of the halo show long-term variations. The position of the target at maximum insertion must therefore be controllable.) The target must be outside the beam envelope for the first 8 ms of the machine cycle, only entering the beam for

the last 1 to 2 ms before extraction, when the protons are close to their maximum energy. The exact time of insertion must also be controllable.

In order to meet the demands described above, a linear electromagnetic drive was implemented to move the target vertically into the beam from above. The technical challenges are considerable. The mechanism must be extremely reliable, to avoid disrupting normal accelerator operation. It must provide an acceleration approaching 1000 ms^{-2} in order that the target overtakes the shrinking beam envelope and is removed before the next injection. Operation must be precise and reproducible, both in position and timing relative to the beam cycle. The mechanism must operate within a high radiation environment, and all moving parts must utilise materials compatible with the stringent constraints of the accelerator's high vacuum system. In case of any failure of the target mechanism, it must be possible to separate it both mechanically and in terms of vacuum from the rest of the synchrotron.

The complete target mechanism, described in detail in this paper, consists of a number of sub-assemblies.

- The linear electromagnetic drive assembly, which contains:
  - The shaft, forming the target at its lower end and carrying a set of fixed permanent magnets, an optical readout vane at its upper end and a stop to prevent the magnets falling from inside the coils in the absence of power;
  - A pair of bearings to support, guide and align the shaft;
  - The stator, consisting of stationary coils with a water cooling system;
  - A central steel tube forming a vacuum barrier between the target shaft and the stationary coil unit;
  - An optical readout enclosure, with sapphire windows.
- The mechanical support assembly, which contains:
  - Flanges to provide accurate location of the stator and bearings;
  - Conflat seals, to ensure the integrity of the vacuum system.
- The mechanical and vacuum isolation system, to allow the unpowered target to be raised out of the beam, and the target vacuum to be separated from the accelerator vacuum. This contains:
  - A structural frame, carrying the weight of the target assembly;
  - A jacking unit and support structure with motorised screw-jack allowing a vertical travel of approximately 150 mm;
  - Centralising units and guide rods that guarantee the target returns to its predefined position when lowered into its operating position;
  - Edge-welded bellows, to allow relative movement of the components under vacuum;
  - A vacuum gate valve to isolate the vacuum systems.

3 Linear motor

Lead author: Chris Booth, Paul Smith

A linear actuator was chosen as the most appropriate mechanism to drive the target into the beam for a short period every second. This implementation does not require any moving parts to cross the vacuum chamber walls, and can be realised without the need for lubricated bearings. For most of the duty cycle, the actuator is only required to exert a small force on the target, in order to keep it levitated out of the beam. At the appropriate time, a large accelerating force is required over a short period in order to accelerate the target into and out of the beam, and then bring it to rest at its levitated holding position. For this short period, high currents can be employed.

The motor must be of a permanent magnet brush-less design, as the high acceleration and travel of the motor rule out the placing of coils on the moving parts. The integration of permanent magnets onto the moving assembly removes the need for electrical contacts between the stator and the moving parts, simplifying the interface between the motor and the ISIS vacuum. The magnets on the moving components interact with the field produced by a set of stationary coils in the stator body. These coils are outside the ISIS vacuum, directly wired to the driving electronics. Positioning outside the vacuum also allows the use of a water cooling circuit, to remove the energy deposited by ohmic heating of the coils.

The initial design of the linear motor was based on studies performed by an electrical engineer specialising in motor deign, and outlined in [**?**]. The important constraints were to maximise the accelerating force while minimising the mass of the moving components. The mass of the magnetic materials thus form a significant fraction of the total mass. Different magnet and coil topologies were investigated, and the resulting design is documented in the following sections.

3.1 Electromagnetic design

Lead author: Paul Smith

3.2 Stator

Lead author: Paul Smith, Chris Booth

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 six 25 μm copper shims are sandwiched between each pair of coils to facilitate heat conduction out of the coil stack. 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 to the stator body.

3.2.1 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 beam.

Each coil is composed of thirty six turns of 0.56 mm polyester-imide enamelled copper wire, over-coated with a polyamide-imide. This is a high temperature motor winding wire that is rated to 200°C operation [**?**]. These coils are wound on an 18.1 mm diameter former and each coil is manufactured to a depth of 2.85±0.1 mm. Coils outside this tolerance are rejected due to the limited space between coils and the required pitch of 3 mm. These dimensions provide ∼0.15 mm clearance between adjacent coils, which is used to insert thin copper shims which act as heat sinks. 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 [**?**]. A photograph of a finished coil and some of the inter-coil copper shims is shown in figure 2.

Figure 2: A complete stator coil and some copper heat-sinking shims.

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 after construction due an invisible mechanical defect in the outer insulation of the coil.

3.2.2 Shims

As already noted there are copper shims filling the 0.15 mm gap between the coils. These shims are designed to act as a thermal path between the coils and the cooling jacket. The copper shims are a split annulus with an inner diameter 19 mm and an outer diameter of 36 mm. They are split to reduce any potential eddy-current losses in the copper (although the effect of such losses has not been established).

3.2.3 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 alignment of the coil stack.

A thin walled non-magnetic stainless steel tube passes through the centre of the coils. It has a nominal wall thickness of just 0.5 mm, reduced to 0.3 mm where it is surrounded by the coil stack. (The effect of the metallic tube, compared with a thicker ceramic tube, was investigated. It was shown that the reduction in magnetic field led to a decrease in force on the permanent magnets of approximately 1%, a result subsequently confirmed in simulations.)

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 where it is reduced in wall thickness 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 between the inner surface of the coil stack and the earthed steel tube.

3.2.4 Cooling Jacket

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 mechanism can be actuated. Typically when the stator is levitating the shaft out of the beam the stator's power consumption is ∼30 W although every time the target mechanism is actuated (and the 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 heat-sink the coils would rise in temperature by $\sim 5^{\circ}$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 demagnetisation. The exact maximum safe operating temperature is hard to determine, as the Curie temperature is field dependent. There is also some evidence that the risk of demagnetisation at elevated temperatures is accentuated when running permanent magnets in a radioactive environment [**?**]. Experience with running the actuator for extended periods of time has demonstrated that with coil temperatures of 80 - 90°C no problems with demagnetisation have occurred.

A cooling circuit is required to remove heat from the coil stack. This consists of an external, water-cooled, split cylindrical copper jacket. The jacket has a thin bore copper cooling tube soldered onto its outer surface through which a flow of water can be maintained. This is illustrated in figure 3.

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

Figure 3: The copper cooling jacket is over-fitted with the water cooling pipes. The cooling pipes are soldered onto the jacket.

Figure 4: The cooling jacket, placed over the coil stack.

copper shims over ensuring a good thermal contact between them and the jacket. A photograph of the jacket placed over the coil stack is shown in figure 4. The cooling pipe has a narrow bore and so the water flow rate is quite low. Typically at $\sim$4 bar a flow rate of $\sim$1 litre min^{-1} is achieved. This water flow rate has proved to be sufficient to remove the heat from the 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.

3.2.5 Stator Assembly

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, then four copper shims. A coil, lightly coated on both sides with a thermal paste to aid with heat-sinking, is added. Six copper shims follow. The second coil is then added to the coil stack and the process is repeated for all twenty-four coils. As each coil has a nominal thickness of 2.85 ± 0.1 mm the cumulative error is tracked and minimised during the coil stack construction process by selection of appropriate thickness coils. An assembled coil stack is illustrated in figure 5.

The split in the copper shims aids in keeping the coils parallel to each other during construction of the coil stack. This is achieved by placing the split in the shim 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: 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.

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. The former on which the coil stack was formed is now removed and the bore tube inserted through the bore of the stator. A final electrical insulation check is then performed, to ensure that all coils remain isolated from metal parts including the bore tube and cooling jacket. The installation of the stator body into the core of the target drive assembly is described in section 4.

3.3 Permanent magnets

Lead author: Paul Smith, Chris Booth

The permanent magnet assembly interacts with the field of the stator coils to produce the force on a central shaft which accelerates the target into and out of the proton beam. To achieve the maximum magnetic field, the assembly is constructed from sintered neodymium-iron-boron (NdFeB) magnets. Twenty-four segments are arranged in 3 rings, with 8 magnets per ring, as shown in figure 6. They are glued to a mild steel core to produce a cylinder that is 18 mm long with an outer diameter of 15 mm and an internal diameter of 4 mm. Thin ceramic washers separate the three rings. The central ring is 7.8 mm long, twice the length of the two outer rings. The whole assembly has a mass of about 25 g.

Figure 6: Drawing showing the structure of the magnet assembly

The individual sectors are manufactured by wire erosion from un-magnetised NdFeB material. The sectors are then magnetised radially and assembled in a jig before finally being glued in position using a strong two-component epoxy. Once the glue is cured, the magnet unit is lightly machined to the precise outer radius required. The middle ring is magnetised so that the outer surface is a North pole, while the outer rings have the opposite polarity. As a result of the relative polarisation of the rings, roughly circular flux lines are produced, as shown in figure 7. The peak radial component of the surface field of the magnets is approximately 0.6-0.9 T.

Figure 7: Diagram showing the flux lines from the permanent magnets as simulated by MAXWELL 2D Student Version [?].

The magnet material currently in use, NdFeB, is the material which shows the highest remanence field of all permanent magnets, allowing the highest force to be achieved for a given coil current. However, concerns about the temperature dependence and radiation intolerance of these magnets could mean that samarium-cobalt (SmCo) magnets will have to be considered in the future. Since SmCo has a lower field strength than NdFeB, a higher coil current would be required to achieve the same acceleration. A FLUKA [**?**] simulation has been performed to ascertain the radiation levels expected around the target, looking specifically at the dose received by the magnets. The expected radiation levels for 1 year (assumed to be 180 24-hour days) of running at 1 Hz and the maximum beam intensity required for MICE have been estimated. The dose received by the magnets is found to be $\sim 1 \times 10^4$ Gy. The literature on radiation damage of NdFeB magnets [**?**] shows a large dependence on a variety of factors: temperature and exact composition of magnets, radiation type and orientation of the magnetic field with respect to the direction of the particles' travel. However, it seems likely that the degree of degradation in one year's running should be acceptable.

Figure 8: MICE Target assembly, (left) with bellows removed, (right) a cutaway view (Descriptions and information on the components in the assemblies shown here follows.)

4 Mechanical design and construction

Lead authors: Jason Tarrant, Eddie McCarron (James Treadgold?)

The MICE Target mechanical components, sub-assemblies and assembly (in some cases with the electronics, electrical power and optical readouts) are designed to perform several functions:

- Accurately, rigidly and stably combine all the mechanical functions, electronics, electrical power and optical readouts into the MICE Target unit;
- Provide a vacuum tight volume to connect to the ISIS beam-line in which the MICE Target shaft operates;
- Enable the target operation to balance particle production with ISIS beam loss by:
  - Providing controlled drive of the target's shaft into and out of the beam;
  - Enabling the velocity and location of the shaft to be determined;
  - Constraining the shaft along its path of travel preventing significant off-axis movement;
  - Providing stiffness in the shaft that resists significant deformation and vibration during operation.
- Eliminate the possibility of a failure that prevents continued operation of the ISIS beam such as:
  - The shaft breaking and falling into an irretrievable position in the path of the beam;
  - Welds or seals, that form the vacuum containment, failing causing a leak into the ISIS beam line;
  - Contamination in the beam line, particularly by conductive dust in the RF cavities.
- Enable an operational lifetime that is economical with respect to the number of target changes required per number of user runs.

The following sections describe the construction of the MICE target assembly starting with the target shaft at the very heart of the unit.

4.1 Target Shaft

Lead authors: Jason Tarrant

The MICE Target shaft is a titanium shaft consisting of 2 parts, a solid upper section and a lower tubular section. The 2 are joined with a shrink-fitted plug and socket arrangement then

Figure 9: MICE Target shaft assembly (shown horizontally oriented, operates vertically with tip down).

electron beam welded together. The shaft is coated with a hard bearing coating then fitted with a permanent magnet, readout vane and their associated fixings and fasteners.

The shaft is 538 mm long and when fully assembled with the magnet, readout vane and fixings has a weight of around 51 g. The permanent magnet is rigidly fixed to the shaft and is driven by a stator that sits around it in the assembly. The stator drives the permanent magnet up and down, which in turn drives the shaft into and out of the edge of the ISIS beam to produce pions. The stator accelerates and then decelerates the shaft in both the down stroke and up stroke with a total travel in each direction of around 48 mm. The shaft achieves this 96 mm of reciprocating movement in about 30 ms. After it dips and returns to its start position it dwells for just under 1 s before dipping again to give a dip frequency of nearly 1 Hz. The tip of the shaft that catches the edge of the ISIS beam has a cylindrical section with a 5.95 mm outside diameter and 4.55 mm inside diameter.

The Target's shaft is the most critical of all the mechanical components of the MICE Target in terms of production quality. The shaft is manufactured within very tight dimensional and geometric form tolerances. The tight dimensional tolerances ensure a close fit of the shaft with the bearings which accurately constrain the shaft's vertical motion. The shaft requires a high degree of straightness and good roundness. Tight tolerances of form (cylindricity and run-out) are applied to the manufacture of the shaft, good form minimising deformation and vibration of the shaft during the rapid acceleration and deceleration.

4.1.1 Design

The rapid acceleration and deceleration of the shaft may cause it to deform and vibrate during operation. This is because it is not perfectly straight, it is not perfectly constrained vertically and the structure and material do not give the shaft perfect stiffness. The design of shaft, the design of bearings, the manufacturing tolerances and the choice of materials were all carefully chosen to minimise deformation or vibration during operation. High speed video photography had been undertaken on an earlier shaft design with a '+' cross section lower shaft, rather than the 'O' cross section lower shaft shown above. This inspection was undertaken to determine what deformation or vibrations were occurring during shaft travel. The video showed the shaft stayed close to the ideal axis of travel on the down stroke; on the return however the shaft wandered off centre and then waggled a number of times as it was captured in the park position. It is suspected the shaft was passing through one or more resonant modes as it was being accelerated and decelerated by the stator.

The stiffness of the shaft is very important. Increasing the stiffness of the shaft, without increasing mass, pushes up the frequencies of resonant modes. Increasing the frequencies, especially those of the first resonant modes, reduces the number of resonant modes the shaft passes through as it is accelerated and decelerated. Vibrational excitation of the target's shaft may come from a number of sources, example of these are:

Figure 10: Output plot of FE modal analysis of shortened shaft with 'O' shaped lower section.

- The time periods for large changes in velocity and direction take place at around 15 ms intervals (twice) over the 30 ms cycle, i.e. from the initial firing of the shaft at the beginning of its dip cycle to the change in direction and eventual capture at the end of its cycle. Excitations induced would be at around 33 Hz and 67 Hz. Longitudinal forces in the unit imposed by the shaft's change in acceleration and deceleration have been calculated to be just over 50 N [**?**]. *Is this correct? Surely force is caused by - or causes - acceleration, not change in acceleration? CNB*

- The switching of power in the stator from one set of coils to the next may introduce an excitation force. The longitudinal driving force (up to 50 N) may be accompanied by off-axis torque or lateral forces. The frequency and severity of the excitation would be directly linked to the timing and magnitude of these forces. The timing of the switching is variable as the shaft is accelerated and decelerated in both directions, so the frequency of excitation is variable too. This may cause the shaft to pass through several excitation frequencies that correspond to the shaft's resonant frequencies. The magnitude of the excitation force depends on the quality of the switching between coils and in the quality of the coils and their location to the permanent magnet. If longitudinal switching is not smooth then there will be large steps in the longitudinal force to 50 N. If there is any asymmetry in the coil windings, or the stator coils are not mechanically coaxial to the permanent magnet, then torque and lateral off axis forces may be significant. With a 0.5 mm offset in coaxiality between the permanent magnet and the magnetic centre the lateral forces have been calculated to be up to 2.5 N and the torque up to 32.5 mNm [?].

- The unit or surrounding support frames may be excited by the dipping shaft. There is potential for this vibration to transfer back from the support frame to the shaft through unit then the bearings. There is significant vibration felt on the test frame after each dip.

As mentioned the first design of shaft (as investigated with the high speed camera) had a '+' shaped cross section. A finite element (FE) model of the shaft was built and the resonant frequencies determined through a modal analysis. The analysis showed that with this '+' cross section and with the shaft being fixed about the magnet, the 1st and 2nd resonant modes appeared when excited at 27.6 Hz, the 3rd and 4th at 80 Hz [**?**].

To stiffen the shaft the '+' lower section was replaced by a cylindrical 'O' section. The 5.95 mm outer diameter and 4.55 mm inner diameter of the cylindrical section give a cross sectional area of 11.5 mm^2, marginally greater cross sectional area than the '+' section but with very little increase in mass. This material displacement raises the second moment of area to 40.5 mm^4, resulting in a stiffness more than two times greater than in the previous design. With further changes shortening the shaft and changing the bearing locations, a new FE modal analysis showed the 1st and 2nd resonant modes at 50 Hz, but the 3rd and 4th at 68 Hz [**?**].

The frequency of the first modes with the stiffer 'O' section are as expected higher than those for the '+' cross section, however the next modes are marginally lower. In both cases the 3rd

Figure 11: Plug and socket detail, upper shaft on left, lower shaft on right

and 4th modes are related to stiffness of upper section, the design of which has changed slightly. The upper section now has a pair of flats to constrain rotational movement. These flats reduce the material and in particular the second moment of area in line with the flats. This is the cause of the decrease in frequency of the 3rd and 4th resonances. As mentioned above there is a potential for significant excitation at around 33 Hz between the start and finish of travel for the shaft. The '+' shaped shaft's first two resonant modes at 27.6 Hz are very close to this. The stiffer 'O' shaped lower section of the new shaft raises the first two modes to 50 Hz which is clear of this potentially significant source of excitation.

The cylindrical shaft can be formed by turning or grinding processes designed specifically for producing cylindrical objects. It has only 1 critical dimension (the OD) and the geometry of the whole cylinder can be constrained within one geometric tolerance, i.e. cylindricity, which accounts for straightness and roundness. The laser readout relies on the vane orientation remaining within a few degrees of perpendicularity to the lasers, hence the need for an anti-rotation device. In order to resist rotation of the shaft, a pair of flats on the upper shaft engage with an anti-rotation component on the upper bearing, see section 4.2.

The obvious problem with the 'O' shape if attempting a 1-piece manufacture would be creating the 4.55 mm bore. The bore is nearly 320 mm long and has to maintain a close concentricity with the 5.95 mm outside diameter to ensure the material in the shaft is well balanced about its axis. Creating the bore is not impossible as it might be achieved by either using deep drilling techniques with correction for errors in concentricity or electro-erosion processes. Though it is not impossible the manufacture of the single item becomes risky as it involves managing the full 538 mm length during each process; also the 1-piece shaft is becoming more and more valuable at each stage so errors in later manufacture are more costly. To alleviate the risk the shaft is produced in 2 sections, the upper and lower sections, which are produced separately then joined together. The additional benefit of separately manufacturing the 2 parts is that they can be produced concurrently with benefits to the schedule.

The lower shaft has an accurately machined socket and the upper shaft a mating plug. The plug and socket are manufactured with a nominal 0.01 mm interference press fit. If these were pressed together at ambient temperature the material would gall, i.e. some of the material at the interface of the plug and the socket would cold weld together. Galling is unlikely to be uniform around the whole circumference of the joint so the 2 halves would skew as they are pressed together. There is also a chance the joint could seize before it is fully pressed together. To mate the plug and socket without galling a clearance fit is required. To achieve this, differences in temperature were used to shrink and expand the plug and socket respectively to fit them together. The socket was expanded by heating to 250°C and the plug was shrunk by cooling to −190°C ; on the 4.75 mm size this gave a clearance fit of just under 0.001 mm on diameter. The 2 halves were pressed together using a precision jig to ensure correct alignment. The interference fit is made as the joint comes back to room temperature and the socket contracts and the plug expands. The circumferential joint between the 2 halves is then spot welded in 8 places using

Figure 12: Failed weld and displacement of shaft halves after tensile testing

Figure 13: A bulge can be seen in the shaft under the stop. This occurred during impact testing.

electron beam (EB) welding. A full circumferential weld was originally tried, however even using the very fine controllable EB welding process the distortion in the weld was unacceptable.

The strength of the spot welded joint was determined by tensile testing [**?**]. The tensile test broke the spot welds at 4 kN. This coincided with an audible crack as the welds broke. As the 2 halves are also joined by the interference fit of the plug and socket the shaft did not break apart, instead the tensile load began to climb in steps against displacement as shown on the load displacement curve plotted during the tensile test. It is suspected that the material at the interface between the plug and socket was galling at this point, so cold welding the two halves of the shaft together. This cold welding appears to be the cause of the steps in the load displacement curve, in that with each increase in load the shaft will pull apart slightly then further cold welding will cause greater adhesion between the plug and socket, thus requiring an increase in load to pull the halves further apart as shown in subsequent steps.

4.1.2 Material

Grade 5 titanium alloy (Ti-6Al-4V) was chosen for the shaft. Titanium is a good target material with respect to particle production, it is widely available, it can be easily worked and can be joined by welding. The shaft needs to be low mass to enable rapid and efficient acceleration by the stator. Also the shaft needs to be stiff to minimise elastic deformation and / or vibration during its operation; titanium and its alloys have a good stiffness to mass ratio. Impact testing was undertaken which closely represented a shaft impacting the end stop [**?**]. Using a 1.5x safety factor, a mass was produced to represent the impact the shaft would see if driven into the end stop at a full speed of 9.3 ms^{-1} [**?**]. The mass was dropped from a set height onto the stop on the lower shaft to prove the design of the stops. However it also showed that the shaft may plastically deform under such loadings. There is a slight improvement in stiffness with the grade 5 titanium alloy over the unalloyed titanium, 114 GPa and 103 GPa respectively, which will help to reduce the deformation under such failure conditions.

Though there is only a modest stiffness improvement between grade 5 and the unalloyed titanium there is a much larger strength and hardness difference. Hardness comparisons between grade 5 titanium alloy and the unalloyed titanium using the Rockwell C scale are 36 and 21 respectively and the strengths are 895 MPa min. and 395 MPa min. respectively (both properties pertain to the materials in their annealed state). The additional strength and hardness are beneficial for some of the manufacturing processes, e.g. grinding, where it is easier to work than the softer unalloyed material.h is allows tighter to be achieved.

4.1.3 Manufacture

The most processed part of the MICE Target unit is the shaft. It currently takes around 6 weeks from receipt of the raw materials to produce a batch of shafts. Processing includes many different production, inspection and treatment steps. Some of these steps are outlined below:

- Grinding: This process is used to produce the overall shape of the shaft. The depth of cut in the grinding process is stepped down towards the final cut to minimise the amount of strain induced into the surface of the component, this in turn minimises internal stresses in the final shaft. Internal stresses may relax through strain over time causing deformation in the shaft. Grinding also gives the shaft a fine surface finish so that polishing removes only a minimal amount of material in achieving the final surface roughness of 0.05 Ra.

- Stress relief annealing: After rough sizing of the target's shaft sub-components a stress relieving anneal is undertaken. The initial rough machining / grinding will induce strains that leave the components with internal stresses. Heat treatment to around $0.3T_m$ (where T_m is the melting temperature) for a period between 0.5 and 1 hour causes the relaxation of stresses through strain and can deform the components slightly. This deformation is removed during subsequent fine sizing of the components. The heat and cycle information for heat treatment is taken from DIN 65084:1990 Aerospace; Heat treatment of wrought titanium and titanium alloys.

- Conventional machining: There are many features of the shafts that are milled, drilled or turned on, for example the fixing holes for the vane or the plug and socket features used to join the upper and lower shaft sections.

- Wire erosion: A 0.22 mm slot by around 90 mm long is required to support the vane. This is produced by a wire erosion process.

- EB welding: The 2 halves of the shaft are permanently fixed by EB welding after they are joined together by the plug and socket features. EB welding was chosen for the joining of the shaft halves as the beam is very small and can be finely controlled. This minimises heating of the parts around the weld joint and so minimises distortions that can occur. It is also a processor controlled process, a feature that allows procedural control for the EB weld program and so a high degree of repeatability between batches.

- Inspection: The target shaft sub-components and the final assembled shaft are inspected at many stages during manufacture to determine if it meets the specification on the technical drawings. Further information on the inspection of the shaft can be found in section 4.5. Some corrective work can be undertaken on the shaft to achieve the required geometric shape, however these adjustments have to be undertaken carefully to prevent damage to the shaft, in particular the welded joints.

- DLC coating: This process applies diamond like carbon (DLC) coating to the bearing surfaces of the target shaft.

Figure 14: DLC coated shaft, dark grey/black coating extending over and beyond the bearing regions of shaft.

Figure 15: (left) Shaft and bearings, (middle) Upper bearing with anti-rotation, (right) Lower bearing

- Assembly: After final inspection the shaft is fitted with the permanent magnet (see section 3.3), readout vane (see section 5) and associated fasteners and fixings. The magnet assembly is bonded and mechanically clamped onto the main shaft. This strong rigid connection is required to resist the high acceleration and deceleration forces that may cause the magnet to become loose or break away from the shaft. Bonding with adhesive also centres the magnet on the shaft to balance the mass. It does this by evenly filling the fine clearance gap between the iron core and the shaft. In addition to the adhesive, the shoulder of the main shaft and a 2 part clamp provide strong mechanical restraints. Assembly (as with the target units) is undertaken in clean-room conditions (ISO 14644-1 ISO 5 or better). The majority of tooling used is dedicated to assembly of the MICE Target, especially the torque wrenches and torque analysers which are calibrated for use.

4.2 Target Bearings

Lead authors: Jason Tarrant

4.2.1 Description

The target bearings are plain bearings that support and constrain bearing regions on the shaft. There are 2 bearings, one to restrain the upper section of shaft and one to restrain the lower section of shaft. They are required to constrain the vertical motion of the shaft as it dips as well as preventing significant rotation about the longitudinal axis of the shaft. Excessive rotation will cause a misalignment between the readout vane and the readout lasers; the upper bearing has the anti rotation component.

The bearings have plastic (Vespel® SCP5000) inserts and anti-rotation component, the dark components in figure 15. The plastic parts are assembled into stainless steel (Nitronic® 60) carriers. The carriers are 30 mm and have tapered sides to accurately locate them in the MICE target core structure. The anti-rotation component is finely adjustable to bring the contact face into light contact with the flats on the target's shaft.

4.2.2 Design

There are a number of requirements related to the running of bearings in ISIS:

- Bearing surfaces will generally have a higher coefficient of friction in vacuum conditions, i.e. without the presence of moisture or wet lubricating films. A high coefficient of friction

between bearing surfaces is likely to increase the power required to drive the shaft, it will cause accelerated wear of the bearing surfaces, and raise the temperature at the bearing interface through frictional heating. The bearing system to be used has to have an acceptable coefficient of friction when used in the ISIS vacuum (10^{-7}) *units?* [**?**].

- Bearing materials, coatings, etc are required to be tolerant to the effects of the nuclear radiation generated in ISIS, i.e. their mechanical properties must not change significantly during the operational life of the unit.
- The bearing surfaces need to withstand a significant operational life, preferably 10^7 cycles.
- The bearing materials, coatings, etc should not become significantly activated by the ISIS radiation, especially if there is a possibility they produce fine particles of dust during wear.
- Particles created during wear should be kept to a minimum. If they are created they must be managed and contained within the unit. This is particularly important for particles that have become activated, for health and safety reasons, also for those particles that are conductive as they may damage the RF cavities adjacent to the MICE target in ISIS.
- Material, coatings, etc must not significantly out-gas in the ISIS vacuum.

In addition to these requirements materials, coatings, etc used for the bearing surfaces should be able to be manufactured to tight tolerances. This requirement will enable minimum radial clearances between the shaft and bearings to enable tightly constrained vertical motion. The radial clearance between the inside diameter of the bearings and outside diameter of the shaft is critical. If the shaft is too loose in the bearings it will be able to move significantly off axis. This lateral movement affects the coaxiality of the permanent magnet to the stator's magnetic field. If the permanent magnet is not coaxially aligned with the magnetic field then non-axial forces caused by the magnetic field become more pronounced.

If the radial clearances are too tight the shaft may be gripped too tightly; in this situation the stator has to work harder to accelerate the shaft, it might move erratically, or at the extreme case it might not move at all. Even with a radial clearance between the upper bearing and upper shaft and similarly lower bearing and lower shaft, there might be no allowance in radial clearance for other factors. Such factors are the inevitable misshapenness of the shaft and the coaxiality of the bearing centres. Though these factors are relatively small, a lack of additional radial clearance may cause the shaft to be constantly in contact with the bearings with some degree of lateral force. This lateral force may cause erratic motion, accelerated wear of the bearings or may cause the shaft to flex during operation. This constant flexing may cause fatigue cracks to occur in the shaft, in particular between the bearing regions. If this flexing is significant and if the shaft is operated for long enough a deep fatigue crack may form which leads to a complete failure of the shaft.

Bearing free solutions were investigated. The most promising used diaphragm springs to support the shaft. The added mass hence inertia from the use of these springs was problematic

Figure 16: Tests of an early MICE Target unit using diaphragm springs and the diaphragm spring.

for control of the trajectory. In addition to the mass issues, diaphragm springs lose more lateral stiffness the further they operate; the large travel of the target's shaft was thus a problem.

4.2.3 Material

A number of material combinations have been tried. From the early trials a diamond like carbon (DLC) coating on both the shaft and the bearing was found to be the most successful. Despite some early successes with the DLC coating there appeared to be too much variability in the quality of the coating. This appeared to be due to the geometry of the parts being coated. This variability lad to coatings on some shafts and bearings wearing away easily.

The alternative to DLC came from cryocooler technology designed and built by STFC's Cryogenics and Magnetics Group. The cryocoolers have demonstrated in excess of 10^{10} cycles in a dry helium atmosphere. The material combination used in these cryocoolers is titanium (both commercially pure and 6% Al, 4% V) against Vespel® SP-3 grade. The SP-3 grade of Vespel® is a polyimide that contains molybdenum-disulphide (MoS_2). MoS_2 is a dry lubricant which lowers the coefficient of friction. However it was recommended by ISIS health physics to avoid the use of MoS_2 in the high radiation environment. Communication with DuPont^{TM}, the manufacturers of Vespel®, led to the selection of the SCP-5000 grade, an unfilled polyimide. SCP-5000 was recommended for use as it has a relatively low coefficient of friction, even in a vacuum environment, 0.26 or better [**?**], also polyimide has an inherently good tolerance to the effects of radiation. More familiar bearing plastics such as those based on PTFE are not tolerant to the levels of radiation found in ISIS and would soon degrade mechanically. As with most plastics Vespel® is hygroscopic. The absorbed moisture will out-gas into the ISIS vacuum if not removed. To remove the moisture after production and before they are to be assembled into the target unit, the bearings including their stainless steel carriers are vacuum baked at 80°C for 72 hours. As the bearing hole is small there is no concern regarding the moisture strain (shrinkage) that will occur as the Vespel® is dried; due to the small size there will be an insignificant reduction in the radial clearance around the shaft in the bearing. After vacuum baking the bearings are stored for a short period in a vacuum desiccators, this is until they are required for assembly and subsequent integration onto ISIS or onto the test vessel. After integration, a vacuum is again applied which will prevent further absorption of moisture from the environment, as well as removing the small amount of moisture uptake that would have occurred during assembly and integration.

4.2.4 Rotational constraint

As mentioned above the bearings not only constrain vertical motion but also rotation. The upper section of the shaft has a pair of flats. The anti-rotation feature of the latest bearing

Figure 17: Upper bearing with Vespel® bearing insert and anti-rotation component

Figure 18: Dust 'can' assembly, that captures the dust below the lower bearing.

with Vespel® inserts has a Vespel® component that contacts both flats. Figure 17 shows these plastic bearings.

4.2.5 Running with plastic bearings

Bearing trials have shown that wear will occur and that some wear particles will be produced. Minimising and managing these wear particles is important in preventing contamination of the ISIS beam-line. There are a number of changes that have been implemented on the latest target assembly to minimise and control these wear particles, see 4.5.3 for more details:

- DLC coated bearing surfaces polished to better than 0.05 Ra surface roughness with a lay parallel to the direction of travel.
- Plastic bearing surfaces will be burnished or run-in.
- The capture or management of dust outside of the target unit.

4.3 Stator

Lead author: Jason Tarrant

4.3.1 Description

The stator assembly is the next component in the Target unit. The stator assembly surrounds the middle portion of the shaft assembly, in particular the permanent magnet which is driven by the stator. The bearings and stator assembly are fixed in the core structure (covered in next section) to coaxially align the shaft assembly and permanent magnet with the stator coils.

The stator, described in section 3.2, is a self contained sub-assembly; it is assembled around, but not fixed to, the central vacuum tube. As it is self contained it can be tested independently of the rest of the MICE target structure. After assembly and test the stator unit is incorporated into the target unit by welding the ends of the vacuum tube into the 2 main core body components. The tube and welds provide an uninterrupted vacuum volume though the middle of the unit. Above and below this central vacuum tube, the components are fixed to the body with copper knife edge (CF) seals. When the stator drives the shaft there is a reactive force, which has been

Figure 19: (left) Stator assembled with shaft and disc spring plate on top, wiring not shown; (right) Tested stator unit ready for assembly (shown with temporary plastic clamps).

Figure 20: (left) Core body assembly with welded vacuum tube, (right) Core Tube Body with Upper Flange assembled on top with fasteners and taper dowels.

calculated at up to 50 N [**?**]. As the stator is not fixed to the vacuum tube then this load is not transferred to the vacuum tube welds. To prevent the coils within the stator moving or the stator moving as a whole, the stator unit is clamped down inside the core body using 3 stacks of disc springs. These springs exert a clamping force of around 100 N onto the stator unit when the Upper Flange of the core body is fully tightened onto the main Core Tube Body component.

4.3.2 Core structure

The core structure of the MICE Target is provided by 2 stainless steel (1.4404) components, the Core Tube Body and the Upper Flange. These contain the stator unit, bearings and support the shaft; the other components of the target assembly are fastened to these.

The core is required to accurately and stably align the other components, in particular the shaft assembly and the stator coils. Careful consideration of the interfaces between the permanent magnet and the stator coils was required. The order of the interfaces for the permanent magnet to the stator coils, i.e. the critical alignment in the target unit, represented by '→', are shown below:

Permanent magnet → Shaft → Bearings → Core structure → Vacuum tube → Stator coils

This number of interfaces requires tight tolerances on size and shape of the above components. Tight tolerances are also required on assembly positions between these components to accurately constrain their respective locations. Only by controlling these interfaces will the permanent magnet be accurately coaxially aligned with the magnetic centre of the stator stack. Coaxial alignment between the permanent magnet and the magnetic field of the stator ensures a high quality vertical drive. Coaxial alignment reduces or eliminates off-axis torque or lateral forces that the permanent magnet may be subjected to as it passes through the stator's changing magnetic field. Significant off-axis forces may lead to accelerated bearing wear and possible trajectory issues through high loading at bearing interfaces or by excitation of the shaft at its resonant frequencies. As well as the interfaces above, accurate alignment of the bearings is very important. The bearing alignment, bearing bore accuracy, straightness of the shaft and tolerances on outer diameter of the shaft all increase the minimum radial clearance tolerable between the bearings and the shaft. Keeping this radial clearance to a minimum ensures the shaft's travel is well constrained. If the radial clearances have no added compensation for manufacturing variations, the fit could be too tight which may lead to flexing of the shaft and possible fatigue failure. To achieve accurate bearing alignment the core structure and the way it supports and locates the stator assembly has been simplified to the minimum number of components and interfaces. There are only 2 components and 3 interfaces between the 2 bearings, bearing → core lower → core upper → bearing. Minimising the components and interfaces is a practical solution for machining as there are few tight tolerances to be achieved. To further aid the accuracy of alignment of the bearings and reduce the manufacturer's liability

Figure 21: Target assembly suspended by rods from motorised stage above ISIS beam-line, connected to beam-line with bellows.

in achieving tight tolerances, the 2 core parts are made with a slight clearance fit, they are then assembled, inspected, adjusted to achieve a tight coaxial alignment, then dowelled together. The Taylor Hobson Talyrond is used to undertake the inspection of the core body component alignment. The bearing seats are currently being aligned to each other within 20 μm of run-out; the repeatability of reassembly alignment with the taper dowels is around 10 μm.

The Core Tube Body has a small flange which is fitted with a larger flange extension. This flange extension is the main reference component for position of the assembly within ISIS. Rods are suspended from a frame, either above the ISIS beam-line or on its test vessel in a separate building (for offline running and trials). These rods fix to the larger flange extension and support the target unit for connection to the beam-line or vessel. Connection to the beam-line or vessel is made via a bellows component. The 2 part flange is a practical solution for manufacturing as it minimises the material cut away from stock billets to form the components. The framework that the rods attach to is motorised. This enables the Target unit to be lowered into position; the bellows compress when this happens. If there is a serious fault with the target unit, the motorised stage will drive the target unit up away from the ISIS beam-line past a shut off valve. This valve can then be closed to isolate the MICE Target from the ISIS beam-line.

4.4 Mechanical Integration

Lead authors: Chris Booth pp Jason Tarrant, Eddie McCarron (James Treadgold?)

The target mechanism must be rigidly supported in a position which is accurately aligned with a port in the ISIS beam-pipe, and in a manner which minimises the transfer of vibrations caused by the linear motor to the synchrotron itself. The target must also be reliably removed from ISIS when it is not in use, and complete isolation must be possible in the case of a fault developing in the mechanism, particularly one which could affect the ISIS vacuum. These objectives are met by suspending the target from a heavy, rigid steel frame which is itself bolted to the synchrotron floor. A motorised platform raises and lowers the target drive, and this is connected to the beam-pipe via a set of extensible bellows and a remotely operable valve which can separate the vacuum systems of the synchrotron and the target drive.

A duplicate section of beam-pipe, complete with target support frame, is situated in an assembly hall where all target mechanisms were commissioned. This allowed testing and alignment of the support before it was moved into the synchrotron.

4.4.1 Motorised Platform

Lead authors: Chris Booth pp Eddie McCarron (James Treadgold?)

Figure 22: Motorised platform.

A heavy steel plate rests on the main support frame. This carries a motorised screw jack driven by a stepper motor. The jack can raise and lower a steel ring, which carries the main target mechanism via three sturdy rods, which pass through holes in the plate. The jack has a travel of 200 mm(?), and when in its lowest position the ring fits into a locating socket. The motor is controlled remotely, and limit switches indicate when the mechanism is at its highest and lowest allowed positions. These are interlocked to the controls to prevent the equipment from being driven outside safe limits. An independent position switch is linked to the Personnel Protection System, and prevents access to the MICE hall when ISIS is operating and the target mechanism is lowered, so that the target could intercept the beam.

The plate is supported on levelling screws, and its lateral position is also adjustable with locking screws to ensure the target is centred over the port in the beam-pipe. Once positioned, the plate is clamped in place.

During normal target operation, the support frame is moved to its lowest position. The target is then outside the beam envelope only while it is magnetically levitated. When the target is not in use, the frame is raised to its highest position, and the target is well clear of the beam even if powered down.

4.4.2 Bellows and ISIS Vessel Interface

Lead authors: Chris Booth pp Eddie McCarron (James Treadgold?)

To allow vertical movement of the target drive while maintaining a vacuum seal, the lower flange of the drive is connected to the beam-pipe using edge-welded bellows. These have an internal diameter of X mm, a compressed length of 40 mm(?) between flanges and an extended length of 250 mm(?). The UHV seal between the bellows and target flange is maintained with a standard Conflat seal and copper gaskets.

As a fault in the target mechanism could compromise the ISIS vacuum, it is important that the former can be isolated from the accelerator. This is achieved with a gate valve, mounted between the flange on the upper face of the vacuum vessel and that at the lower end of the bellows. Seals are formed using aluminium gaskets compressed by V-band clamps. The valve is operated by a low pressure nitrogen line, controlled remotely, and has an aperture of X mm to allow the passage of the target shaft. The control system ensures that the target frame cannot be lowered while the valve is closed, and that the valve cannot be closed unless the frame is at its upper limit switch.

The ISIS vessel is adapted for the presence of the MICE target in several ways. As well as the port on the upper surface, to accept the target mechanism via the gate valve and bellows, it has a thin-walled steel particle window to allow the passage of pions, produced by interactions in the target, into the MICE secondary beam-line. In addition, it has a glass inspection window

Figure 23: (left) Bellows shown assembled under core tube body; (right) Compressed and uncompressed bellows states.

directly below the target. This allows visual observation of the target, and was incorporated when the production of dust from target bearings was a concern.

4.5 Quality Assurance and Quality Control

Lead author: Jason Tarrant

4.5.1 General QA and QC

Quality assurance is applied to the MICE target in all areas. This includes prototyping and development, purchasing, processing / manufacture and assembly. Quality assurance is used to ensure the fit of parts in the assembly, to ensure performance is as expected as determined through prototyping and development and that traceability is quick and effective should problems be encountered. The following are some examples of the quality assurance which is applied to the components and sub-assemblies of the MICE Target.

- Materials: All materials used to produce the Target assembly parts were purchased with certification. Certification shows the material meets a particular standard as well as giving information about certain characteristics, e.g. chemical composition. For the target parts the way materials react to heat treatment and their weldability are important characteristics that can be affected by the chemical composition.
- Processing: Cycle data and copies of in-date equipment calibration certification were requested for processing where applicable. For the target shaft the heating cycle print out for stress relief annealing is supplied with the finished parts, along with a copy of the calibration certificate for the vacuum oven. The welding procedures were developed through prototyping and tests, this led to the development and use of working procedures for EB welding of the shaft parts. In following these working procedures the EB welding machine set-up is identical for each batch of shafts guaranteeing a high degree of consistency between batches.
- Inspection: Machined parts produced for the target assembly are inspected to ensure they are dimensionally and geometrically correct. The measuring equipment is all calibrated to ensure the measurements are accurate. Compliance with drawing tolerances through measurement ensures that component parts of the MICE Target are able to be assembled together without need for remake or adjustment to fit. The size and shape of components not only affects their fit in current assemblies, it also affects their interchangeability with other components from batch to batch; interchangeability ensures that used targets can be refurbished with new parts in the future saving manufacturing time and money. Manufacturing of components and then assembly to specified tolerance limits also assures a level of performance for the MICE Target.

Figure 24: Talyrond with jig, jig being aligned vertically for shaft measurement using gauge rod.

- Tooling: Tightening torques for threaded fasteners are specified on all assemblies to ensure components are preloaded correctly and fasteners are not over tightened. Calibrated torque wrenches and torque analysers with calibration certification are used in the assembly of the MICE Target.

4.5.2 Shaft Measurement

The shape of the shaft is very important. As mentioned previously, the shaft will not be perfectly straight due to its thin long shape and the material stiffness. However as long as the form is controlled within tolerable limits operation will be acceptable. It has been shown that current manufacturing can produce a shaft where the maximum run-out from the lower bearing region to the upper bearing region or from the lower bearing region to the tip of the shaft can be controlled to within 0.12 mm. A run-out of 0.12 mm means that the eccentricity of the worst case feature of the upper bearing region with respect to the average centre axis of the lower bearing region is 0.06 mm. Target 1 has been successfully run with a shaft that has a total run-out of 0.114 mm.

The size of the shaft is very easy to measure using conventional micrometers. The shape however is more difficult and is undertaken in an inspection laboratory. For the target shafts a Taylor Hobson Talyrond is used to measure the run-out between the lower bearing region and the other parts of the shaft. A jig was designed and produced for shaft measurement. A bore on the jig holds the shaft vertically with 2 sects of collets. This collet bore is aligned vertically to the machine. The machines probe gauges the position of the bore then automatically adjusts the table of the machine. The shaft is clamped into the jig on the lower bearing region vertically; it is oriented either way up depending on the measurements to be taken. The probe then measures predetermined vertical positions. The surface profile is measured over a full 360° rotation and is output as a 2D slice. These 2D slices can be stacked to produce an extrapolated 3D model of the shape. For the upper shaft bearing region 5 vertical measurements are taken. For the tip, only the very end is measured. As well as the 2D and 3D plots, the Talyrond automatically outputs total run-out and eccentricity measurements for each slice which are used for quality control of the target shaft. The probe has an extremely light touch (1.5 mN). However this would be enough to bend the shaft when measuring away from the support point at the lower bearing region. A counterbalance has been used to counteract any forces from the probe during measurements. This counterbalance touches on the opposite side of the shaft to where the probe is measuring.

The quality control, in particular the reports and documentation produced when measuring the size and form of components, have been beneficial during post operation investigations. The mechanical size differences between Target 1 (T1) and Target 2 (T2) were very quickly identified after T2 failed to perform as well as T1. There were very slight differences in the shaft of T2 to that of T1; the shaft of T2 had a maximum total run-out of about 0.17 mm compared to

Figure 25: Wear in bore of lower bearing, polishing and longitudinal scratches.

0.114 mm of T1. Another set of shaft and bearings were selected and refitted, the shaft of which had a better straightness (0.103 mm) than that of T1. T2 was run again with these new set of shaft and bearings but T2 was still not performing as expected. By comparing all the inspection data it was clear that the problem was not related to the mechanics of the target assembly, in particular the shaft and bearings. This launched an investigation into other causes.

4.5.3 Plastic Bearing Investigation

Running with plastic bearings

Targets have been run for 1 or 2 million cycles, and then stripped down for examination. Initial tests showed the importance of having a well-polished titanium shaft in the bearing regions, prior to DLC coating. On inspection, some wear particles were found, in the vicinity of the bearings. Fine dust was noted, both on the inside of the unit and on the outside around the bearings. Both were relatively well adhered to the bearing and needed to be wiped off. The lower section of shaft in contact with the lower bearing was polished (0.03 to 0.05 Ra); although dust and wear had been produced it was not as great as the amount produced in the upper bearing or early devices, where polishing was not as good.

An optical microscope examination of both the upper and lower bearings showed that wear had occurred. General wear appeared as a polishing of the surface of the bearing and removal of the machined finish. On the anti-rotation component, there was clear evidence of relatively deep wear. In the lower bearing the wear was not uniform around the bore. This was checked against the orientation of certain features of the target unit, in particular the drive components, to see if there was a correlation between any features and the wear. There were no conclusive correlations.

Conclusions

The investigations showed that wear would produce wear particles, in this case a fine dust. The dust in large particles is, by its nature as an unfilled plastic, a dielectric material. However the dust produced is very fine, and there are concerns it might break down in the ISIS beam leading to conductive carbon contamination of the RF cavities in adjacent sections of the ISIS beam-line. In a review of the performance of the Vespel® bearings it was recommended that the dust levels be reduced. It was also recommended that the dust be managed to prevent it escaping the unit and contaminating the ISIS beam-line. A number of changes were implemented to achieve this:

- DLC coated bearing surfaces on the shaft will be polished to better than 0.05 Ra surface roughness. The lay of the finish will be parallel to the direction of travel.
- Plastic bearing surfaces will be burnished or run-in to remove machined surfaces and give greatest contact area from the start of operation.

Figure 26: Difference between run-in surface and rougher as machined surface on anti-rotation component (viewed under microscope with the component tilted over).

- The capture or management of dust within the target unit. This has been achieved by using a long 'can' assembly that is mounted just below the lower bearing. This 'can' envelops the volume just below lower bearing region on the shaft and has a number of internal features to trap dust. The shaft is also fitted with a skirted stop that works with the can to prevent dust migrating down the shaft and out of the unit.

There is also the possibility of trialling alternative plastic materials to determine if they are harder or more wear resistant than Vespel® SCP5000.

4.5.4 Magnetic Inspection

Lead author: Paul Smith/Geoff Barber

5 Optical position-measurement system

Lead authors: Paul Smith, Chris Booth

In order to switch the current through the stator coils at the correct time to drive the target shaft, it is necessary to sense the exact position of the shaft while it is moving. To avoid disrupting the motion of the low mass shaft by mechanical contact, and to remove the necessity for electrical feed-throughs traversing the vacuum wall, an optical method is adopted. Furthermore, since active electrical components would not survive in the high radiation environment near the target, optical fibres are used to convey the optical signals, with all sensitive electronics being placed outside the accelerator vault (in the experiment control room some 100 metres away).

A graticule or vane attached to the shaft and interrupting a light beam will generate a series of pulses, which could be used to determine the speed but not direction of movement of the shaft. If it is arranged that two beams are interrupted by the vane, with the pulse trains being 90 deg out of phase, then both speed and direction can be inferred. (Such quadrature encoders are used in many commercial movement detectors.) With a third beam producing a pulse at only one well-defined point, a zero of position can be defined. In this way, by counting pulses away from the zero, an absolute position measurement can be made.

In addition to providing feedback to the controller, position measurement allows the trajectory of the target to be monitored and recorded by the data acquisition system, allowing long-term reliability and stability to be monitored and fault conditions to be diagnosed.

Figure 27: The optical vane. One side has 157 apertures to provide a quadrature signal. The other has a single slot to give a fixed reference point.

5.1 Optical Vane

A diagram of the vane is shown in figure 27. The vane is a double-sided graticule manufactured from 0.2 mm thick steel, having 157 slots 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 spacing between slots is also 0.3 mm. There is a single similar aperture two-thirds of the way down the vane on the other side. To protect the fine features of the vane, it has a continuous edge. The vane was produced by photographic etching, and it is attached to the shaft by four M1.6 fasteners. In order to ensure that the readout system provides reliable position information, the manufacturing tolerances of the apertures were kept to less than 5% of the 150 μm resolution, and the tracking error along the length of the vane was also less than this percentage. The vane and fixings have a mass of about 1 g, which is balanced about the axis of the shaft.

5.2 Laser Source

Three fibre-coupled solid-state Lasers provide the light beams which intercept the vane. Commerical red (635 nm) Lasers, with a variable 0 to 2.5 mW output power are used. Visible light has advantages both for alignment and safety. The milli-Watt power level is required due to significant losses in the optical system, due to the number of interfaces such as at fibre connections. Maximising the light on the optical sensors (Section **??**) increases the signal to noise ratio and simplifies the design of the electronic amplifiers.

In practice, the Lasers are not operated at maximum power, but at about 1 mW. This leaves sufficient overhead (both in light source and amplification gain) to adjust for any degradation in the optical fibres, e.g. due to increased attenuation as a result of radiation damage in the vicinity of the synchrotron.

5.3 Optical Fibres

Two types of optical fibre are used in the light path. On the transmitting side, single mode fibres are required, in order to achieve the necessary small spot size at the collimator focal point and hence at the plane of the vane. On the receiving side, 200 μm multimode fibres are used.

The single mode fibres used are SM600 with FC/FC connectors [**?**]. This has a core of pure silica, which is a radiation-hard material and so ideal for the operating environment of the target. If single-mode fibres were used on the receiving side, the collimators would have to be aligned to an extremely high precision, and it would be hard to achieve an adequate light transmission. Multi-mode fibres have a higher acceptance, so make the alignment less critical. The fibres used are BFH37/200, with SMA to SMA connectors [?].

5.4 Collimators, Lenses and Mechanical Mount

Collimators and lenses are used to focus the light from the fibres to a spot in the plane of the optical vane (where it needs to be significantly smaller than the pitch of the vane), and to receive the transmitted light coupling it into the return fibre. The collimator is a commercial unit which produces parallel light from the diverging beam emitted by the fibre. Collimators with a focal length of 15.3 mm are used, as the longer focal length minimises the final spot size. The focusing lens, with focal length 45 mm, is attached to the front of the collimators by inserting it into a holder which is screwed onto the front of the collimator as shown in figure 28. On the transmitting side, aspheric double achromatic lenses are used to obtain the minimum spot size. On the collecting side, where focusing is less critical, double convex lenses are used to re-collimate the beam. The lenses are MgF_2 coated to minimise reflections and maximise the light transmitted into the fibres.

Figure 28: Cross section through a collimator and lens cap.

The optical vane moves inside the accelerator vacuum system, while all optical elements are kept outside the vacuum chamber to allow alignment adjustments. A mechanical mount was produced which has two circular sapphire windows, to allow entry and exit of the light, and provides a rigid assembly to support the optical collimator and lens units. This is illustrated in figure 29. The flat windows are bonded into metal flanges, and are rated for use in UHV environments. The pair of collimators for each optical channel is held in an arm that wraps around the mount. The exact alignment of each collimator is performed by adjusting pairs of opposed grub screws, with 4 pairs per collimator: horizontal and vertical at front and back of each collimator. The "Channel A" and "Index" arms are also adjustable in the vertical direction to allow the correct quadrature phase to be obtained and an appropriate index position to be set. The mount has proved to be easy to set up and extremely robust. Once aligned on the bench, no further optical adjustments have been necessary during periods of operation lasting several years.

Figure 29: An exploded view of the optical mount. The optical vane sits inside the main body, where the collimators and lenses are able to focus the light.

5.5 Optical Sensors

The multimode fibres are returned to a SMA photodiode (H3R880IR) [**?**]. This is a broad spectrum photodiode (400 to 1100 nm). As the maximum velocity of the target is less than 10 ms^{-1}, the maximum data rate per channel is only 33 kHz, well within the response capability of these devices. The output from the photodiodes are amplified and conditioned (as described somewhere else?????).

6 Stator Operation and Power Electronics

Lead authors: Paul Smith, Steve Griffiths (or delegate)

6.1 Introduction

The target stator is a three phase brushless permanent magnet linear motor and is run as a DC motor. 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 A, B and C, with eight coils in each phase. Starting from the top of the coil stack the coils are lettered in a cyclic sequence that follows the pattern, 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 **??**.

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 either an A, B or C coil in a clockwise direction then the same current passes through an A' B' or C' coil in an anticlockwise direction. This has the effect that for a given current flow through a phase the induced magnetic field direction for a primed coil is in the opposite direction to that of an unprimed coil.

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 separate connections at the bottom of the stack are wired together and this single connection is called the 'star- point'. The three connections that remain at the top of the stack serve as the connections to the stator power supply. A diagram representing the stator wiring is shown if figure **??**. Note that all four of the banks are wired in series as this gives the simplest setup.

If current is fed into one of the phases then it must return through one (or both) of the other phases. Therefore when current flows, at a minimum, two phases must have current passing through them at any given time.

6.2 Coil Current Switching Sequence

For the target mechanism only two out of three phases are ever powered at the same time, i.e. one phase is switched to provide the current source and another phase provides the current return path. This means that with three phases current can be switched through the stator in six different ways.

Using the external connection labels A, B' and C as shown in figure **??** to describe the current direction through the stator, these six states are; $A \rightarrow B'$, $A \rightarrow C$, $B' \rightarrow A$, $B' \rightarrow C$, $C \rightarrow A$ and $C \rightarrow B'$.

By stepping through these six different states in a predetermined order then the coils can be switched so that they create a 'ripple' motion in the magnetic fields generated within the stator. 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; this is best illustrated graphically and is shown in figure **??**. The direction of this motion is determined by the direction in which these predetermined states are stepped through; when reversing this order the ripple motion moves in the opposite direction. By tracking the position of the magnets and correlating this position to a given state, target motion can be achieved. The required order of the states to observe this apparent motion upwards through the stator is shown in table 1.

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 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 **??**

6.3 Magnetic Assembly

The magnetic assembly has been discussed at some length in an eariler section but given the importance of the magnet's geometry with respect to how it interacts with the stator coils a short 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 - the coils are placed on a pitch of 3 mm. This distance corresponds to the symmetry of the axial field 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 the interaction between the magnetic assembly and the coils a simulation of the target mechanism was produced using a magnetostatic model in Opera 3D [**?**].

6.4 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'. (The state is arbitrary due to the symmetry of the device.) The model is axis-symmetric in R, Z and the origin is placed at the centre of the stator; this corresponds to a plane that sits on top of the 13th coil. For this model a current of 58 A was passed through the coils. This value closely approximates the peak current that passes through the coils during operation.

The simulation was run over several iterations, and with each iteration the magnetic assembly was moved in either 0.5 mm or 1 mm increments axially along the centre line of the stator 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 **??**.

Although not illustrated here the sinusoidal pattern of the force curve will repeat itself over the distance of one bank of coils, a distance of 18 mm. Figure **??** 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. Although not illustrated, further simulations confirmed that as the magnets move towards the end of the stator bore there is no deviation from this sinusoidal pattern. This is due to the fact that the magnet assembly is always completely contained within the stator body so any end effects have minimal impact.

6.5 Zero-Force Points

Figure **??** shows several interesting features that need to be explained if the operation of the stator is to be understood. Firstly there are two points for each repetition of the sine wave where the force on the magnetic assembly is zero. 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. However the other zero force point is stable, any movement of the magnets away from this zero-force point results in a restoring force; effectively this point represent the centre of a magnetic well.

These stable zero-force points are of particular interest to the operation of the stator as they provide a mechanism by which it is possible to passively levitate the target shaft. Providing that the force of gravity is less than the electromagnetic force that can be exerted on the shaft's magnets when the magnets are in close proximity to the zero force point, then levitation of the target 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 force. 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 with a coil current that is much smaller that that required to achieve the high accelerations needed to put the target shaft into the ISIS beam.

6.6 Moving the Zero-Force Points

These zero-force points in figure **??** are shown for the stator in 'State 4'. If the state sequence is progressed then the position of these zero-force points move in step with the state sequence. Moving forward through the sequence moves the zero force points up through the stator in 3 mm increments, likewise moving backwards through the sequence moves the zero 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 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 point, as a restoring force will push the shaft to the new equilibrium position. This is illustrated in figure **??**. The shaft will undergo damped harmonic motion as it settles at this new point.

The ability to move and hold the shaft in this way is utilised to control the target shaft position when the target is not actuating. For example by switching the coils to 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, and then moved to its holding position by progressing cyclically through these states. This final state, the 'hold mode' then holds the shaft indefinitely until it is ready for actuation. Reversing the process allows the shaft to be lowered back down to its 'park-mode'.

6.7 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 way of the ISIS beam 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 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. When moving the target shaft by using the zero-force point system position monitoring is provided by the optical position measuring system.

6.8 Actuation

By consideration of figure **??** it can be seen that the location of a zero force point is positioned half way between two maximum force points; we define a maximum force point as a position

within the stator where the magnets on the target shaft would experience a maximal repulsive force. For each set of two maximum force points one of them will push the magnets in one direction whilst the other will push the magnets in the opposite direction. Figure **??** shows that these maximum force points are positioned ±4.5 mm away from the zero force points; this is of course true for any zero force point in any one of the six states that the stator can be in.

By looking at the position of the peak forces it can be seen that the forces change very little ±1.5 mm either side of the peak. Integration of the fitted sine wave ±1.5 mm either side of the peak shows that the average force is 95.5 % of the peak force. If the shaft was levitated at a zero force point and the stator coil state was to either increment or decrement by two states this would then put the shuttle magnets 1.5 mm on the far side of one of the peak forces as the zero force point would have been moved by 6 mm. At this point the shaft would accelerate back towards the shifted zero force point. In either case the resultant force would cause the shaft to pass through the location of the peak force during the first 3 mm of acceleration. See figure **??** for clarification of how this would work.

If after switching the coil two states no further switching was done the shaft would continue to accelerate towards the new zero force point 6 mm further up or down the stator and it would execute damped harmonic motion about this zero-force point until it came to rest. However because the position of the shaft can be 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. This has the effect of maintaining a maximal accelerating force on the shaft. This process can of course be continued down the entire length of the stator as illustrated in figure **??**.

If the shaft is accelerated via this mechanism, deceleration can be achieved by switching the coil state by three positions. This has the effect of placing a force of equal magnitude but opposite direction on the permanent magnets. Once again by referring to figure **??** and comparing this to figure **??** it can be seen that a switch of three states changes the direction of the accelerating force because it simply reverses the current flow through the coils. If a previously accelerated shaft is decelerated by this process then there will be a point where the shaft will change its direction of motion.

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.[1] This change in direction of the shaft can be used as a trigger to capture the shaft at the nearest zero-force point at the end of an actuation. This is done by switching the stator to a state that places a zero force point close to the magnets's current position. The coil pitch of 3 mm dictates that the shaft will never be more than 1.5 mm away from a possible zero-force point and so, providing the shaft does not a high velocity, capture of the shaft at the zero-force point is inevitable.

The processes just describe provide the necessary mechanism with a suitable controller to

[1] Note that the shaft's actual velocity when the change of direction is detected is dependent upon the accelerating force and the resolution of the position monitoring system. Therefore the shaft will have some small velocity as the change of direction is detected because the quadrature system detects the change after the event.

accelerate 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 μm. As will be described in the next section the high resolution of the optical system allows better control of the actuation depth of the trajectory of the target shaft than the minimum positional resolution requirement would allow. A full description of how the target electronics controls the trajectory of the shaft using the mechanisms just described will be discussed in a later section although figure **??** shows the principle of how it is done.

6.9 Implementing the Switching of the Coils and Current Control

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 effectively 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 section explains these requirements in more detail and describes the power supply that provides these three controlled phase currents to the target mechanism.

6.10 Switching Mechanism

Switching current bi-directionally through the three phases of the stator can be accomplished using six transistors. The six transistors are arranged in 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 transistors then connects to one of the three phase wires of the stator. Figure **??** illustrates how these transistors are connected. This type of circuit is known as a 'Hex Bridge' or a 'Three Phase Invertor' 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 2 with figure **??**. 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.

One risk that needs to be managed with this circuit is that if both transistors connected to a given 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 will almost invariably

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

permanently damage the transistors due to excessive current flow. This type of short circuit due to transistor switching errors is known as a 'shoot-through'.

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 necessary to ensure that a certain amount of dead time is allowed before the other paired transistor can be turned on. If this extra time is not allowed then a 'shoot-through' can occur.

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. In this case both sets of transistors that are on a live phase change state and so a suitable dead time must be introduced by the controller to prevent 'shoot-through'.

6.11 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 although a small offset can be accounted for. 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 large enough that the levitation is stable but small enough that power dissipation is minimised within the stator. A current of around

3 A has been found to satisfy both of these criteria. This can be seen by looking at figure **??**; the shaft, magnets and vane has a mass of ≈57 g so the gravitational force is about 0.56 N. This force is counteracted by the electromagnetic force at a displacement of ≈500 μm from the true zero force point.

On actuation the shaft will be accelerated at ≈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. It has been found that a coil current of ≈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.

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.

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 power IGBTs have an internal diode that connects back across the device to protect it against inductive surges.

6.12 The Target Power Supply

It was expected that the target mechanism would require approximately 60 A of current during actuation to achieve the required acceleration. This original figure was based upon extrapolating from an early 10 A prototype power supply and through simulation.

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 STFC. To ensure that enough overhead was built in to permit any future changes or upgrades, the power supply was originally specified to operate at up to 300 V and to provide a current of up to 100 A.

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 Ω load (The DC resistance across two phases of the stator at 20°C) this gives a power dissipation of only ≈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 from the prototypes have shown that this safety mechanism is very effective at preventing the stator from overheating.

To quickly switch the high currents necessary to drive the target power Integrated Gate Bipolar Transistors (IGBTs) were used. These transistors switch current quickly and saturate very close to the power supply rail voltage due to a small internal resistance (R_{ON}). Both of these characteristics are good for minimising the power dissipation within the transistors and this makes them very efficient at transferring power to the load. By utilising these transistors a circuit that can switch a significant amount of power can be built with a minimum footprint.

The IGBT's that were chosen for the target power supply are power devices and can switch up to 250 A. 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 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.

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. This supply was then used for an extended period operating at 220 V providing up to 60 A into the target stator.

The assumption of a linear relationship between the duty-cycle and the current 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 the this effect can become problematic if the transistor switching time T_{ON} becomes greater than:

$$T_{ON} << \frac{\text{Duty Cycle (\%)}}{\text{Frequency}} \tag{1}$$

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 required holding current, at 3 A, is scaled from the maximum delivered current during actuation to give the PWM duty cycle when the actuator is not actuating; however a certain amount of trial and error is required to get the correct duty cycle as the power supply starts to show non-linear behaviour at around this current.

The PWM is set at 100 % duty cycle when actuating, which is equivalent to a DC current. In this case the DC voltage on the power supply will set the actuation current, the current being primarily limited by the discharge on the capacitor bank, the DC resistance of the stator coils, cables and the transistor's 'ON' resistance.

Snubber circuitry has been added to 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 EMI emitted by the motor, its power supply cable and its driving circuitry.

Two stator failures that occurred within a period of 18 months both due to an electrical short to ground from a coil led to an improvement in the quality assurance procedures in selecting coils for use in a stator. 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 %. 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 -115 to +115 V DC - see figure **??** [**?**]. 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 so to obtain a similar actuation time the capacitor banks have to be operated at slightly above half the original voltage.

6.13 System Placement in ISIS

It was clear early on in the project that the installed layout of the target 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 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 run between the control electronics and the power electronics and nearly a 100 m run between the control electronics and the stator.

6.14 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 the section on optics. Additionally fibre optic cables have been used to transmit the six signals between the controller and the power supply. 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 has utilised a commercially available optical transmitter/receiver pair. The transmitter is a high powered Infra-Red LED and it is 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 which 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.

7 Target Control, Electronics and the Data Acquisition System

Lead authors: Paul Smith, James Leaver

This section give a description of how the target is controlled by the system electronics. The interface between the electronics and the computer that records the performance of the target, the Data Acquisition (DAQ) system, is also discussed. Figure **??** 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.

7.1 Target Controller Overview

The target controller is implemented on a Field Programmable Gate Array (FPGA) IC. An FPGA consists of many thousands of logical units that can be programmably configured to create a bespoke digital circuit. For the target controller the logical units are configured to create a finite state machine (FSM) that both interfaces the controller to the outside world and controls the operation of the target mechanism.

More specifically the FPGA used for the target controller is a Xilinx XC3S1000 Spartan-3 FPGA. This device contains 1 million logic gates arranged as 17,280 logic cells. It has 120K of distributed RAM, 432K of block RAM, 391 I/O pins and a host of other features. **REFERENCE UG331** The FPGA has been bonded to a custom made PCB that permits access to most of the FPGA's IO pins and has a USB interface. The USB interface allows the FPGA to communicate bidirectionally with a PC permitting soft control of the underlying FSM based upon its current state through a GUI interface. This FPGA/pcb combination has been used in other experiments and has proven to be reliable. (REFERENCE - need to ask James!)

Broadly speaking control of the target mechanism is based around the target mechanism being in one of four states at any given time; Parked, Hold Mode, Actuate Enable and Actuating. The job of the FPGA controller is to control the movement of the target shaft, ensuring that it is in the correct place for its given state. The correct state will of course be dependent upon input from both external signals and the user. The reality is that the controller's FSM has many tens of sub-states to ensure that the target operates correctly in a safe manner, however this level of detail will not be described here. Figure **??** 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.

7.2 Control: Park, Hold and Actuate Enable Modes

When the target controller and power supply systems are first powered 'on' the stator itself is left un-powered. In this state 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 tip of the target shaft 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. To ensure that these conditions are met the target controller will not allow the target frame to be lowered until the target is in its hold position. Conversely it will not park the target while the target frame is in its lowered position.

An enable key-switch on the front controller unlocks the system and permits it to receive control signals over the USB connection. Providing that all the interlocks are good, then by using the GUI, the controller can be commanded to raise the target mechansim to its hold position. The actual hold position is determined by a register value that is loaded over the USB interface and so is configurable in software. The default value is set to raise the target shaft by 51 mm which is enough to hold the tip of the target out of the way of the ISIS beam when the target frame has been lowered. The holding position is achieved by advancing the zero-force point up the stator as described in an earlier section.

Once the target has been raised to its hold position a check is done by the controller to ensure that a correct count has been obtained by the quadrature counter and that an index signal has been received. An incorrect count or missing index signal could indicate a problem with the quadrature system and actuation would then not be possible.

The change of state from hold mode to actuate enable mode reflects a change of the internal state of the controller which is initiated by the user from the control PC. However 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, additionally there must be no other internal errors. All the external interlocks must be good.

7.3 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 has 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 disruption to the operation of ISIS if it occurred frequently.

As stated in a previous subsection the controller and power supply have been designed to apply a constant force during the actuation as this gives the advantage that it is possible to accurately control the dip depth. Actuation is performed actively; the coils are switched 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 position of

the target shaft and an accurate measure of the shaft's position is essential.

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 **??**.

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 therefore determines the actuation depth. The register on the FPGA that holds the value of switch point 1 has its value calculated by the GUI, dependent upon the requested dip depth. As the resolution of the quadrature system is 150 μm this means that the target dip depth is controllable to within 300 μm (although there is some actuation to actuation jitter on this which has some dependence upon the amount of bearing wear.)

As the currents are reversed 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 now accelerates upwards until the target reaches the position defined by 'switch-point 2'. Here 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 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. [2]

Switch-point two is offset wrt 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. If the value of switch point 2 is too high or too low this results in the target being

[2] Because the shaft will have some small velocity by the time the change of direction is detected, it is necessary to reverse the current flow again for a short period of time to arrest this residual velocity. Typically this is done for a few hundred microseconds. This process is called a 'kickback' and without it reliable target capture would not be possible.

captured at either a higher or lower zero force point. If under/over capture occurs persistently then the controller will automatically make a correction to switch-point 2 offset 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.

7.4 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 for an actuation with a 44 mm strike 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 of it because of capacitor bank discharge (≈10 % deterioration in current through actuation) and gravity (≈ ±1 % each way).

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, even with this much flexibility in the trigger system it is still impossible for the target to trigger off a given ISIS injection trigger (ISIS MS) and to intercept that same spill 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 $2^n \times 20$ ms $-$ 15 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. All of this timing is taken care of in the FPGA the only detail that the user needs to enter is the required second delay value and the mode that ISIS is running in. (ISIS can run in various modes that alters the required value of the first delay.)

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 [**?**].

7.5 DAQ

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 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 PCs hard disk drive. The software also enables both online and offline analysis of the data once it has been recorded.

The DAQ hardware starts with a local PC computer that contains a specialised PCI based interface card; this PC sits in close proximity to the target control electronics. The card is a National Instruments 6254 PCI card that has 48 digital inputs and 32 analogue inputs. Because of driver problems under linux the card cannot simultaneously record both digital and analogue data at a sufficiently high enough sample rate for the target DAQ so currently only analogue data is recorded. As most of the signals that need to be recorded by the DAQ are analogue

in nature this is not too problematic, however it has meant that the target position has to be converted to an analogue signal by a Digital to Analogue convertor (DAC) before the signal is interfaced to the card. The analogue signals that are recorded are target position, beamloss from ISIS in sector 7, beamloss from ISIS in sector 8, and total beamloss. The target is situtated in sector 7 and sector 8 is immediately downstream from there. The next section will give more information on what these signals mean and how the target affects them.

It is also worth noting that the target controller GUI also creates a data stream of target parameters on an actuation by actuation basis. This gives the status of the controller, as well as a record of controller settings and errors and other key actuation parameters. Currently the GUI data stream is independent of the data files generated by the DAQ, however there are plans to integrate these in the future.

8 Performance

Lead authors: Paul Hodgson, Adam Dobbs

9 Summary

Lead author: Chris Booth

Acknowledgements

We gratefully acknowledge the ISIS Division at the STFC Rutherford Appleton Laboratory for the warm spirit of collaboration and for providing access to laboartory space, facilities, and invaluable support. We are indebted to the MICE collaboration, which has provided the motivation for, and the context in which, the work reported here was carried out. We would like to acknowledge *any contractors, manufacturers, or suppliers whom we would like to acknowledge.* This work was supported by the Science and Technology Facilities Council under grant numbers PP/E003214/1, PP/E000479/1, PP/E000509/1, PP/E000444/1, and through SLAs with STFC-supported laboratories.

References

[1] S. Geer, "Neutrino beams from muon storage rings: Characteristics and physics potential," *Phys. Rev.* **D57** (1998) 6989–6997, hep-ph/9712290.

[2] D. Neuffer, "Multi-TeV muon colliders," *AIP Conference Proceedings* **156** (1987), no. 1, 201–208.

[3] **MICE** Collaboration, G. Gregoire *et al.*, "An International Muon Ionization Cooling Experiment (MICE)," *MICE Note* **167** (2003) http://mice.iit.edu/mnp/MICE0021.pdf.

[4] ISIS Pulsed Neutron & Muon Source. http://www.isis.rl.ac.uk/.

[5] Rutherford Appleton Laboratory (RAL). http://www.scitech.ac.uk/About/find/RAL/introduction.aspx/.

[6] **MICE** Collaboration, "MICE Technical Reference Document - Draft Version." http://www.isis.rl.ac.uk/accelerator/mice/TR/MICE_Tech_Ref.html, 2005.

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