



Muon Ionisation Cooling Experiment

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

April 5, 2012

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 are 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 target is evaluated as a function of the particle rate, and the operating parameters of the MICE target are derived.

G. Barber, K.R. Long

Department of Physics, Blackett Laboratory, Imperial College London, Exhibition Road, London SW7 2AZ, UK

D. Adams, E. Capocci, E. McCarron, J. Tarrant,

STFC Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire, OX11 0QX, UK

C.N. Booth, P. Hodgson, L. Howlett, R. Nicholson, E. Overton, P.J. Smith

Department of Physics and Astronomy, University of Sheffield, Sheffield, UK

Contents

1	Introduction	1
2	Requirements and overview	3
3	Linear motor	4
3.1	Electromagnetic design	5
3.2	Stator	5
3.3	Permanent magnets	9
4	Mechanical design and construction	10
4.1	Target Shaft	11
4.2	Target Bearings	19
4.3	Stator	22
4.4	Mechanical Integration	26
4.5	Quality Assurance and Quality Control	29
5	Optical position-measurement system	32
5.1	Optical Vane	32
5.2	Laser Source	33
5.3	Optical Fibres	33
5.4	Collimators, Lenses and Mechanical Mount	34
5.5	Optical Sensors	35
6	Stator Operation and Power Electronics	35
6.1	Introduction	36
6.2	Coil Current Switching Sequence	37
6.3	Magnetic Assembly and Modelling	37
6.4	Zero-Force Points	39
6.5	Actuation	41

6.6	Coil Switching and Current Control	43
6.7	The Target Power Supply	46
6.8	System Placement in ISIS	48
6.9	Fibre Optic Links	48
7	Target Control, Electronics and the Data Acquisition System	49
7.1	Target Controller Overview	49
7.2	Control: Park, Hold and Actuate Enable Modes	50
7.3	Control: Target Actuation and Capture	51
7.4	The ISIS Trigger Signal	53
7.5	DAQ	54
8	Performance	55
9	Summary	55

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 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 2π mm and 10π mm, 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 2 [8]. 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 is presented in section 4. Section 5 describes the optical position-measurement system, the power




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).




Figure 2: The MICE Muon beam-line

electronics used to drive the linear motor is described in section 6 and the control system in section 7. 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

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 principal function as a spallation neutron source. On selected pulses, the MICE target is caused to dip into the outer low-density halo of the proton beam just before extraction. Pions produced in the target emerge through a thin window in the ISIS vacuum system in the direction of the MICE 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 controlled.) 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 800 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 use 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 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

A linear actuator was chosen as the most appropriate mechanism to drive the target into the beam. 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, to keep it levitated out of the beam. At the appropriate time, a large accelerating force is required over a short period 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 large travel of the motor rule out the placing of coils on the moving parts. The integration of permanent magnets into 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 Joule heating of the coils.

The initial design of the linear motor was based on studies performed by an electrical engineer specialising in motor design, and outlined in [9]. 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: Nigel Schofield / Paul Smith

3.2 Stator

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, the temperature of which 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 resin. This yields a high temperature winding wire that is rated to 200°C operation [10]. These coils are wound on an 18.1 mm diameter former, each coil having 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. The clearance of ~ 0.15 mm between the coils is

used to insert thin copper shims which act as heat sinks. After winding, the outer diameter of the coils is 30 mm. Each coil is double dipped into a varnish that seals the windings and provides additional electrical insulation [11]. Each coil is tested to 1 kV before being built into a stator. A photograph of a finished coil and some of the inter-coil copper shims is shown in figure 3. The copper shims provide a thermal path between the coils and the cooling jacket. The inner diameter of the shims is 19 mm and the outer diameter 36 mm. The shims are split to reduce eddy-current losses.




Figure 3: A complete stator coil and some copper heat-sinking shims.

3.2.2 Stator Bore

A thin walled non-magnetic stainless steel tube passing through the centre of the coils forms the stator bore. It provides isolation between the stator body and the ISIS vacuum and ensures the mechanical alignment of the coil stack. The nominal wall thickness of 0.5 mm is reduced to 0.3 mm where it passes through the coils. The reduction in magnetic field strength within the bore caused by the stainless steel tube was estimated to be 1%. The stainless steel tube is insulated from the coil stack using three layers of self-adhesive kapton tape.

3.2.3 Cooling Jacket

The cooling of the stator 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. Every time the target mechanism is actuated an additional heat load of 400 J is deposited in 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 $\sim 5^{\circ}\text{C}$ with every actuation. Therefore, if this heat is not removed quickly, repeated actuation of the stator will 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 [12]. Running the stator for extended periods has demonstrated that coil temperatures of $80\text{--}90^{\circ}\text{C}$ does not lead to demagnetisation.

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 4.




Figure 4: The copper cooling jacket is over-fitted with the water cooling pipes. The cooling pipes are soldered onto the jacket.

The inner diameter of the cooling jacket is slightly smaller than the outer diameter of the copper shims. When the jacket is slid over the coil stack this has the effect of bending the copper shims over, thus ensuring a good thermal contact between them and the jacket. A photograph of the jacket placed over the coil stack is shown in figure 5. The cooling pipe has a narrow bore and so the water flow rate is quite low. Typically, at ~ 4 bar, a flow rate of ~ 1 litre min^{-1} is achieved. This 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, well within the working temperature range of all the components.




Figure 5: The cooling jacket, placed over the coil stack.

3.2.4 Stator Assembly

The coil stack is assembled over a former that has the same diameter as the insulated bore tube. The stack starts with a spacer followed by 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 by selection of coils of appropriate thickness. An assembled coil stack is shown in figure 6.

The split in the copper shims aids in keeping the coils parallel during construction. 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 allows thermocouples to be inserted into the coil stack.

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.



Figure 6: 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.

3.3 Permanent magnets

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 7. 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.

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 8. The peak radial component of the surface field of the magnets is approximately 0.6-0.9 T.

A FLUKA [14] simulation has been performed to ascertain the radiation levels expected around the target. Assuming 24 hour operation at 1 Hz for 180 days in a nominal year, the expected radiation dose to the magnets is estimated to be $\sim 1 \times 10^4$ Gy. To date, no evidence of degradation due to radiation has been observed.




Figure 7: Drawing showing the structure of the magnet assembly




Figure 8: Diagram showing the flux lines from the permanent magnets as simulated by MAXWELL 2D Student Version [13].

4 Mechanical design and construction

The MICE Target unit contains mechanical components and sub-assemblies that are designed to perform several functions, in some cases in conjunction with the electronics, electrical power and optical readouts:

- Accurately and stably combine all the mechanical, electronic, electrical power and optical

readout functions into the MICE Target unit;

- Provide a vacuum tight volume connected hermetically 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 thus 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 that could contaminate the RF cavities.
- Enable an operational lifetime that is economical with respect to the number of new or refurbished target changes required over a 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

The MICE Target shaft comprises of two sections of titanium (Ti) alloy grade 5 (6% Al, 4%V). The sections are a solid upper section and a lower tubular section with various functional features on each. The two halves are joined with a shrink-fitted plug and socket arrangement then electron beam welded together. The shaft is coated with a hard diamond-like carbon (DLC) coating that acts as a bearing surface. The shaft is finally fitted with a permanent magnet, slotted graticule for position read-out that works in conjunction with a laser head sitting on the unit and their associated fixings and fasteners.

The shaft is 528 mm long and when fully assembled with the magnet, slotted graticule and fixings it has a weight of around 51 g. A radially segmented permanent magnet assembly is rigidly fixed to the shaft by bonding and a mechanical clamp. The permanent magnet sits inside a stator; the stator is a stack of electromagnetic coils that are powered sequentially to drive the permanent magnet, thus the attached shaft, linearly up and down. The stator accelerates and




Figure 9: MICE Target assembly, (left) with bellows removed, (right) a cutaway view (Descriptions and information on the components in the assemblies shown here follows.)




Figure 10: MICE Target shaft assembly (shown horizontally oriented, operates vertically with tip down).

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 before dwelling and then dipping again; with the dwell the shaft dips at a frequency of just under 1 Hz. The tip of the shaft has a cylindrical cross section of $\sim 11.5 \text{ mm}^2$ (5.95 mm outside diameter and 4.55 mm inside diameter); as it is an integral part of the tubular lower shaft it is grade 5 titanium alloy. It is an adjustable predetermined length of this tip that momentarily grazes the halo of the ISIS beam during the dip to produce pions.

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; this tight fit accurately constrains 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 along its length and its geometry and the material from which it is made do not give the shaft perfect stiffness. High speed video photography had been used to film the path of travel of an earlier shaft design with a '+' cross section lower shaft, rather than the tubular 'O' cross section lower shaft shown above. The video showed the shaft followed a single axis of travel on the down stroke; however on the return the shaft wandered off centre and then waggled a number of times as it was captured in its powered park position. It is suspected that the off-axis movement of the shaft was due to it 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 few 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:

- The time periods for large changes in velocity and direction take place at around 15 ms intervals, i.e. twice over the 30 ms cycle. Excitations induced would be at around 33 Hz and 67 Hz. If the natural frequency of the shaft meets either of these frequencies, some distortion would be induced.
- 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 as the sequential coil switching takes place during dipping. 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




Figure 11: Output plot of FE modal analysis of shortened shaft with ‘O’ shaped lower section.

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 depends on the force seen by the shaft when switching between coils; this is related to the quality of the coils and their axial location relative to the permanent magnet. With a 0.5 mm offset in coaxiality between the permanent magnet and the coil’s magnetic centre the lateral forces have been calculated to be up to 2.5 N and the torque up to 32.5 mN m.

- 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 the unit and the bearings. There is significant vibration felt on the test frame after each dip.

As mentioned the first design of shaft had a ‘+’ shaped cross section. This ‘+’ shape fitted a bearing with a matching ‘+’ shaped hole, thus preventing rotation of the shaft and the attached flat graticule used for laser alignment. 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. The ‘+’ shaped shaft’s first two resonant modes at 27.6 Hz are very close to the potential excitation frequency of 33 Hz.

To stiffen the shaft the ‘+’ lower section was replaced by a tubular ‘O’ section. The ‘O’ section had comparable material to the ‘+’ section to maintain the low mass but due to the geometric change it had over twice the stiffness. 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 geometric change on the lower ‘O’ shaped shaft has significantly raised the first two resonant modes; the 3rd and 4th modes are related to the stiffness of the upper section. The upper section is the same as the previous shaft with the exception that it 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.




Figure 12: Plug and socket detail, upper shaft on left, lower shaft on right

The laser readout relies on the orientation of the flat graticule remaining within a few degrees of perpendicularity to the lasers. The anti-rotation ability of a ‘+’ shaped shaft fitting a ‘+’ bearing hole is lost with the ‘O’ shaped cross section; with the new shaft a pair of flats on the upper section engage with an anti-rotation component on the upper bearing, see section 4.2. Due to complexity of the ‘+’ shape to manufacture (on both the shaft and in the bearing) tight tolerances cannot be easily achieved, also form along the length of the shaft is an issue due to the many manufacturing operations required to achieve the shape. The simpler ‘O’ shaped tubular shaft can be formed by turning or grinding processes designed specifically for producing cylindrical objects. Unlike the ‘+’ shaped shaft it has only one critical dimension for fit with the bearings, the outside diameter (OD) and its form along the length can be constrained by a single geometric tolerance of cylindricity; this accounts for straightness and roundness.

The tubular lower section is created by rough machining a bar to leave a slightly oversized external profile, including an integrated stop. This is then heat treated to remove residual stresses. The bore is then wire eroded to 4.55 mm ID over the length of the lower shaft (about 320 mm); if internal stresses were present the shaft may deform along its length as this material is removed. After wire erosion the external profile, including the stop, is finished to final size referenced to the bore to maintain concentricity and thus material balance. Finally the socket for the shrink-fit joint is added.

The upper shaft stock material is stress relief heat treated before manufacture. It is then ground; this includes the flats which are very finely ground to ensure internal stresses are minimised. A larger diameter is left at one end onto which the socket for the shrink-fit joint is machined. Also added are the undercut for the magnet clamp and the fixing holes for the laser readout graticule. The very fine slot for the graticule is then cut by wire erosion.



Figure 13: Failed weld and displacement of shaft halves after tensile testing

To alleviate the risk of undertaking so many operations on a long single piece shaft it is produced in two sections. The additional benefit of this is that they can be produced concurrently with benefits to the schedule. As mentioned the lower shaft has an accurately machined socket and the upper shaft has 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 cold weld. This is unlikely to be uniform and so the joint would skew causing a significant form error between the upper and lower shaft sections; it may even seize before it is fully pressed together. To achieve a clearance as they are assembled (i.e. prevent cold welding or galling) the socket is expanded by heating to 250°C and the plug is shrunk by cooling to -190°C in nitrogen. A precision jig is used to press them together to ensure correct alignment. The interference fit is made as the joint comes back to room temperature, contracting the socket and expanding the plug.

The two halves are further secured by welding at 8 points along the seam between the plug & socket using electron beam (EB) welding. A full circumferential weld was originally tried, however under-run or over-run at the end of the weld always made the two halves skew too much. Tensile testing was used to determine the strength of the shrink-fit & point welded joint. Tensile testing broke the spot welds at 4 kN which coincided with an audible crack as the welds broke. After the welds broke the interference fit of the plug and socket prevented the shaft breaking into two halves; instead the tensile load began to climb in finite steps against displacement to 6 kN before the test was stopped (as was shown on a load displacement plot). This was due to a cold welding effect; this is beneficial as it significantly reduces the chances of the shaft breaking into two should high impact load occur during a fault.

4.1.2 Material

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 needed to be low mass to enable rapid and efficient acceleration by the stator. The shaft also needed to be stiff to minimise elastic deformation and / or vibration during its operation; titanium and its alloys have a good stiffness to mass ratio.

The initial redesign of the lower shaft led to the use of a tube with a mechanical stop welded



Figure 14: A bulge can be seen in the shaft under the stop. This occurred during impact testing.

on; tube in unalloyed grade 2 is widely available and was initially chosen for the shaft material. Impact testing was undertaken on this welded stop design; this closely represented a shaft's mechanical end stop impacting the lower bearing at a maximum speed of 9.3 m s^{-1} during a fault condition, including a $1.5\times$ safety factor. A mass was dropped from a set height onto the stop on the lower shaft. This stop failed at around 4 impacts. In addition a bulge started to form in the wall of the tube due to the impact, see figure 14. Though the EB welding is a semi-automated process there is potential for variability in the strength of the weld (via contamination, set-up etc), so the welded on stop was not an entirely risk free choice. In addition the welding of the stop caused some distortion to the tube that had to be corrected by mechanical manipulation. It was decided to pursue a more costly manufacturing route to manufacture a stop integrated into the lower section from a single bar, a much stronger design. This overcame the restriction of obtaining only grade 2 tube and allowed consideration of titanium alloys.

Grade 5 (Ti-6AL-4V) titanium alloy was eventually chosen. There is a slight improvement in stiffness with the grade 5 titanium alloy over the unalloyed grade 2 titanium, 114 GPa and 103 GPa respectively. The additional strength of the grade 5, 895 MPa minimum over the 395 MPa minimum of grade 2, is of no significant benefit as the loads on the shaft are minimal during normal operation and the integrated stop design is much stronger. The biggest benefit of grade 5 over grade 2 is in the hardness - Rockwell C 36 and 21 respectively in an annealed state. The harder alloyed grade 5 is easier to grind, allowing better finishes and tighter tolerances to be achieved. This is particularly useful for the upper shaft where the diameters and flats are ground.

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.

Grinding is used to produce the main geometry of the upper shaft. The depth of cut in the grinding process is reduced 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 via 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 is undertaken on raw materials before processing as well as at certain stages throughout the processing. This removes internal stresses induced by manufacturing processes. Unrestrained heat treatment to around 670°C with a dwell period of 40 – 50 minutes then a slow cool relaxes the stresses through strain deforming the components slightly. This deformation is removed during subsequent processing to achieve the final component's size with minimum internal stresses. The temperature and heat treatment cycle information was taken from DIN 65084:1990 Aerospace; Heat treatment of wrought titanium and titanium alloys.

A stress relief anneal is also applied to the assembled shaft prior to the final DLC coating process to straighten it. The shaft is held firmly and accurately in a jig, see figure 15. The jig is made from mild steel which is itself stress relieved (at 700°C) before an accurate 'V' is finely ground into it that supports the shaft. Mild steel has a similar coefficient of expansion to titanium alloy, 11 ppm/°C and 9 ppm/°C respectively; this means during expansion and contraction upon heating and cooling there will be minimal force induced by the jig on the shaft. Mild steel is also a dissimilar material to the titanium alloy so there is less chance of the jig and the shaft bonding at the elevated temperatures during heat treatment. The latest shafts have all undergone this process and have been straightened from >0.5 mm run-out post manufacture to around 0.1 mm run-out post treatment.

Conventional machining, such as milling and turning, is used on many features including the profile of the lower shaft, the fixing holes for the graticule and the plug and socket features used to join the upper and lower shaft sections.

Wire erosion is used to create a 90 mm long by 0.22 mm slot in the 4 mm diameter upper shaft to support the graticule. The 4.55 mm inside diameter by 320 mm long bore through the lower shaft is also made by wire erosion.

EB welding is used to join the two halves of the shaft using eight point welds around the periphery where the plug and socket meet. EB welding was chosen for the joining of the shaft halves as the beam is very small and can be finely controlled. The fine beam minimises heating of the parts around the weld joint and so minimises distortions and internal stresses that can occur. EB welding is also a semi-automated controlled process, which allows a high degree of repeatability between batches once the program is written and proved.

Figure missing!

Figure 15: Shaft in straightening jig

The target shaft sub-components and the final assembled shaft are inspected at many stages during manufacture to determine if they meet the specification on the technical drawings. Further information on the inspection of the shaft can be found in section 4.5.

Diamond like carbon (DLC) coating is applied to the bearing surfaces of the target shaft. The bearing surfaces are polished prior to DLC coating to ensure the final finish is smooth. The tip of the shaft that dips into the beam is not coated. The tip is left uncoated to prevent any thermal heating, from contact with the ISIS beam, causing failure of the coating through high temperature ($>400^{\circ}\text{C}$) or through thermal strain differentials.

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 two 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 incorporating an anti-rotation feature to prevent significant rotation about the longitudinal axis of the shaft. Excessive rotation will cause a misalignment between the slotted graticule and the readout lasers. The upper bearing has the anti rotation component; it is finely adjustable via screws to ensure an accurate location with respect to the anti-rotation flats on the upper shaft. The bearings have plastic (Vespel® SCP5000) inserts and anti-rotation component, the



Figure 16: DLC coated shaft, dark grey/black coating extending over and beyond the bearing regions of shaft.




Figure 17: (left) Shaft and bearings, (middle) Upper bearing with anti-rotation, (right) Lower bearing

dark components in figure 17. The plastic parts are assembled into stainless steel (Nitronic® 60) stainless steel carriers. Nitronic® was chosen as it minimises cold welding when the fine location taper on the outside of the bearing carriers are lightly pressed into the tapered bearing seat in the 1.4404 / 316L stainless steel body.

4.2.2 Material

For the bearings the material and coating combinations at the bearing interface with the shaft were very important for function and longevity. A number of material & coating combinations have been trialled for suitability for use on the target. From the early trials of different coatings, a diamond like carbon (DLC) coating on both the shaft and the bearing was found to be suitable. Despite some early successes there appeared to be too much variability in the quality of the coating. This may have been due to the geometry of the parts being coated; the small bearing hole, even if produced in several segments, produced a “hollow cathode effect” in the RF coating chamber thus variability in the thickness and adhesion of the DLC coating. This variability led to coatings on some shafts and bearings wearing away easily in later trials.

The alternative to DLC came from cryocooler technology as 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. DuPont™ Technical were contacted to determine if an unfilled alternative may be suitable. From the following requirements it was determined that Vespel® SCP5000 would meet the needs of the Target when operating on the ISIS beam-line.

Bearing surfaces have a higher coefficient of friction in the ISIS vacuum 10^{-7} mbar [?], i.e. without the presence of moisture or lubricating films. A high coefficient of friction between bearing surfaces will increase the power required to drive the shaft, it will cause accelerated wear of the bearing surfaces and it will raise the temperature at the bearing interface through frictional heating. SCP-5000 was recommended for use as it has a relatively low coefficient of friction, even in a vacuum environment, of 0.26 or better [?].

Bearing materials 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. Vespel® is a polyimide which will tolerate a significant total accumulated ionising radiation dose before a loss in mechanical properties.

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. In addition particles that are conductive should not be released from the unit as they may damage the RF cavities adjacent to the MICE target in ISIS. Dust is produced when using the Vespel® bearings, however with a finely polished shaft this has been minimised. The Target unit has been fitted with an extended housing with dust traps to prevent the dust produced from the Vespel® escaping outside of the unit, see figure 18. The bearing surfaces need to withstand a significant operational life and be able to be manufactured with fine surface finishes and tight tolerances to minimise the rate of wear; Vespel® is a relatively hard plastic so suits this.




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

The material must not significantly out-gas in the ISIS vacuum. Outgassing tests were undertaken on vacuum baked samples of Vespel® SCP5000 (80°C for 72 hours) and were acceptable for use on the ISIS beam-line provided the total volume used was kept below 3 cm³.

Solutions without the use of sliding bearing were investigated. The most promising used diaphragm springs to support the shaft. The added mass hence inertia from the use of these springs however was problematic 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 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 finely polished parallel flats that run against a flat faced Vespel® anti-rotation component on the bearing, see figure 19. The flat faces of the anti-rotation component are finely adjusted so that the flats on the upper shaft only contact them if the shaft rotates. In trials using an anti-rotation component with two completely flat faces, it was found that after running for some time the corners of the flats on the upper shaft would dig-in to the plastic and cause the shaft to temporarily lock; the semicircular reliefs shown in figure 19 were added to prevent this.

4.3 Stator

Lead author: Jason tarrant, Geoff Barber, Eamonn Capocci

Figure missing!

Figure 19: Upper bearing with Vespel® bearing insert and anti-rotation component




Figure 20: (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).

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 the next section) to coaxially align the shaft assembly and permanent magnet with the stator coils.




Figure 21: (left) Core body assembly with welded vacuum tube, (right) core tube body with upper flange assembled on top with fasteners and taper dowels.

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 through 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 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

Lead author: Jason Tarrant

The core structure of the MICE Target is provided by 2 robust stainless steel (1.4404 / 316L) components, the core tube body and the upper flange. These contain the stator unit and bearings which in turn support the shaft. The other components of the target assembly, such as the electrical wiring and covers, are also fastened to these core components.

The core is required to accurately and stably align the other components, in particular the accurate coaxial alignment of the permanent magnets on the shaft and the stator coils. This produces the most efficient and accurate drive. Careful consideration of the interfaces between

the permanent magnet and the stator coils was required. The order of the interfaces relating to this critical alignment of the permanent magnet and stator coils are represented by ‘→’, and shown below:

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

This number of interfaces requires each component to be manufactured to tight tolerances then accurately assembled with respect to each other. Coaxial alignment between the permanent magnet and the magnetic field of the stator 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 build up to 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 or high power loading on the stator. 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 two components and three 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 in achieving tight tolerances, the two 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 bobbin that supports the stator has a central bore where the permanent magnet runs up and down. The bobbin is welded into the core assembly so that it forms a hermetic seal from the flange on the upper flange to the flange on the core tube body. On these flanges are CF type seal features which use knife edges to compress copper seals. A hermetically sealed internal vacuum volume is created when the optical housing and bellows are fitted to the core assembly and the bellows are then fitted to the vessel in the ISIS ring or the test vessel on the test rig.

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




Figure 22: Target assembly suspended by rods from motorised stage above ISIS beam line, connected to beam-line with bellows.

building (for offline running and trials), as described in the next section. 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, see figure 22 . 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 assem-




Figure 23: Motorised platform.

bly 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?)

A heavy steel plate rests on the main support frame. This carries a motorised screw jack driven by a stepper motor, shown in figure 23. 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.




Figure 24: (left) Bellows shown assembled under core tube body; (right) Compressed and uncompressed bellows states.

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, illustrated in figure 24. 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 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 design, development & prototyping, procurement, processing and manufacture & 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

All materials used to produce the target assembly parts were purchased with appropriate materials certification. The certificates shows the actual characteristics and properties of the supplied material rather than generalised ones. Reference is made to the certificate to ensure the characteristics and properties are within acceptable limits; this ensures the material supplied is of a suitable standard for use in the target assembly.

Process cycle data and copies of equipment calibration certification were requested for processing for critical parts. A heating cycle is plotted during the stress relief annealing of the target's shaft; this is supplied with the finished parts along with a copy of the calibration certificate for the vacuum oven.

Welding procedures were developed through prototyping for welding the two halves of the MICE target's shafts. The procedure defines a set-up for the EB welding machine which is then identical for each batch of shafts; this guarantees a high degree of consistency between batches.

Manufacturing of components and assembly to specified tolerance limits assures a level of performance for the MICE target. The tolerances are based on typical variations of size and geometry as measured in successful prototype units. Sizes & geometry are specified with tolerances on the engineering drawings. Finished parts and assemblies are inspected to these ensure they are dimensionally and geometrically compliant. This compliance ensures that the MICE target components are able to be assembled together without need for costly remake or adjustment to fit. Size and shape variations within tolerable limits between different batches of components also allows interchangeability of parts to refurbish worn targets thus saving time and money by recycling the majority of the target unit for another operational run.

4.5.2 Shaft Measurement

The shape of the shaft is very important. The shaft cannot be manufactured perfectly straight due to its thin long shape; however as long as the form is controlled within tolerable limits operation will be acceptable. During manufacture care is taken to produce the shaft sections

with the best possible straightness thus minimising run-out between the bearing regions as well as finely balancing material about the axis of the shaft. Manufacturing includes the use of jigs to hold the shape during processing, heat treatment stages to relieve stresses and a final constrained heat treatment to achieve the desired shape. Shafts can now be produced with a run-out between bearing regions of less than 0.12 mm; through prototyping it was proved that shafts with around 0.12 mm of run-out between bearing regions would operate acceptably in the target unit.

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's shaft a Taylor Hobson Talleyrand is used to measure the run-out between the lower bearing region and the other parts of the shaft. The Talleyrand has a rotating table onto which annular parts are clamped; a probe on an extending arm then measures the exterior of the annular component as the table rotates. A jig was designed and produced for shaft measurement that allows the Talleyrand to access relevant sections of the shaft for measurement. It uses an upright post to support a collet assembly; by lightly press fitting the shaft into the collets then the collets into a holder the shaft is secured and aligned. The larger diameter lower bearing region is chosen to be the datum to which all other surfaces are measured. The Talleyrand is programmed to locate this surface by measuring over 360° circumferential over a number of positions along its axis. The Talleyrand then aligns the lower bearing region with the axis of rotation; it does this by automatically adjusting the plane of the table. Next it is programmed to measure the other sections of the shaft. Again the Talleyrand measures circumferential over a number of places along the axis of the region to produce a profile of a 2D slice. These 2D slices can be stacked to produce an extrapolated 3D model of the shape of the region as well as being referenced to the lower bearing region datum. The probe has an extremely light touch (1.5 mN), however this is enough to elastically bend the shaft when measuring further away from the collet supports. A fine sprung arm counterbalance is placed to oppose the probe arm to counteract any forces from the probe during measurements.

4.5.3 Magnetic Inspection

Lead author: Ben Shepherd

To characterise the magnetic performance of the target mechanism, magnetic measurements are carried out on the stator and on the permanent magnet assembly. The equipment used is a three-axis Hall probe mounted on the end of a travelling arm. The position of the probe can be set to a resolution of 1 μm in three axes.

Stator

For the stator, measurements are focused on finding the position of the magnetic centre of the coils, relative to the mechanical centre. If the coils are offset, the small radial field generated by the solenoidal coils may have the effect of moving the shaft laterally as it moves through the stator. Accurate measurements of the magnetic centre are therefore important to minimise wear




Figure 25: Talleyrand with jig, jig being aligned vertically for shaft measurement using gauge rod.

on the bearings during operation.

The coils are powered using a constant current of 7 A; this is lower than the peak current when the shaft is dipped into the beam due to heating concerns - a constant current of 60 A would overheat the coils. A current of 7 A heats the stator to around 50°C when water cooling is used.

The stator is set up and aligned to the measurement bench. A cross-hair on the end of the Hall probe is viewed through a telescope as the probe is moved up and down the bench; this allows us to set the telescope axis parallel to the bench axis. Cylindrical inserts with a precisely machined central hole are placed in either end of the stator bore in turn, to allow the stator axis to be aligned parallel to the bench axis. The estimated accuracy for this is around 100 μm at each end, which gives an overall alignment accuracy of 0.8 mrad over the 180 mm length of the stator.

An initial scan of the stator along the z (longitudinal) axis shows a series of peaks corresponding to the individual coils. To find the centre of each coil, a two-dimensional scan at each peak position is carried out. The magnetic centre is the point at which the longitudinal field is a minimum. The longitudinal field is fitted to the function:

$$B_z = ax^2 + bx + cy^2 + dy + exy + f$$

The centre point can then easily be found from the fit to the data.

Permanent Magnets

For the permanent magnets, measurements are carried out to determine azimuthal variability of the radial field (i.e. B_r vs θ). Due to variations in the PM blocks, there is a certain amount of

variation in the peak (and trough) values of the radial field. An assembly with smaller differences between the peaks would be preferable, as it may reduce lateral movement and vibration of the target shaft during operation.

The PM assembly is attached to a rotational stage in order to map the radial field as a function of angle and longitudinal position. In this case, the Hall probe is attached to a spring-loaded head to keep it the same distance from the PMs (about 1 mm) as the assembly is rotated. This was found to be necessary since the PM assembly was not parallel to the Hall probe axis. The effect is small - a few tens of microns as the PM assembly is rotated - but enough to give a significant difference in the measured field, as the field drops off rapidly with radius.

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° 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.

5.1 Optical Vane

A diagram of the vane is shown in figure 26. 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




Figure 26: Two optical vanes. One side has 157 apertures to provide a quadrature signal. The other has a single slot to give a fixed reference point.

provides reliable position information, the manufacturing tolerances of the apertures were kept to less than 5% of the $150\ \mu\text{m}$ 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. Commercial 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 5.5) 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, $200\ \mu\text{m}$ 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 27. 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 MgF₂ coated to minimise reflections and maximise the light transmitted into the fibres.




Figure 27: 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 28. 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 28: 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 before it is converted into a digital signal.

6 Stator Operation and Power Electronics

Lead authors: Paul Smith

6.1 Introduction

The target stator is a three phase brushless permanent magnet linear 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.

The 24 coils in the stator 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, as shown in figure 29. All of the A and A' coils are wired together in series, and similarly for the B and B' coils and the C and C' coils. The difference between the unprimed and primed coils is the sense in which they are wired, 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.

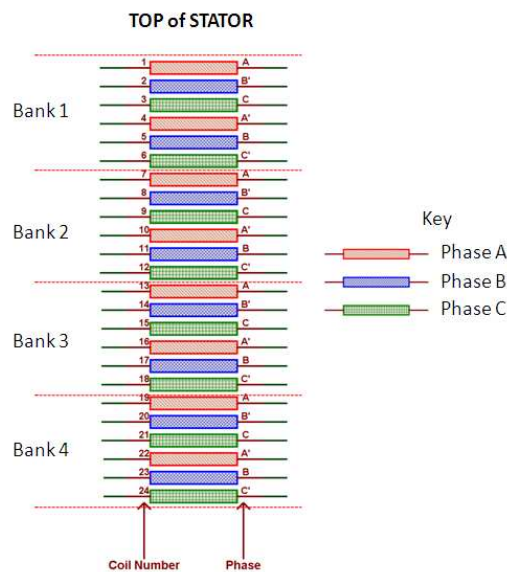


Figure 29: The relationship of phase of the coils to their physical position within the stator. Coils marked with a prime are wired so that current flows through them in the opposite direction to those that are unprimed.

The coils in each phase are wired in series, leaving 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 to form the ‘star-point’. The three at the top of the stack are connected to the stator power supply. The stator wiring is shown in figure 30. 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.

External Connections to the Stator




Figure 30: The wiring of the coils in the stator, illustrating the star point and the three external connections from the stator. Each coil in the diagram represents 4 coils i.e. the A coil represents the A coil from banks 1 to 4.

6.2 Coil Current Switching Sequence

For the target mechanism, only two 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. With three phases, current can therefore be switched through the stator in six different ways. Using the external connection labels A , B' and C as shown in figure 30, 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 states in a predetermined order, the coils can be switched to create a 'ripple' motion in the magnetic fields generated within the stator. This pattern can be visualised by considering the way that the coil switching sequence moves up or down the stator body, as illustrated graphically in figure 31. The direction of motion is determined by the direction in which these 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.

6.3 Magnetic Assembly and Modelling

As described in section 3.3, the magnet assembly is composed of 3 radially magnetised rings, with a total length of 18 mm, matching the depth of 6 coils. 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.

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 31

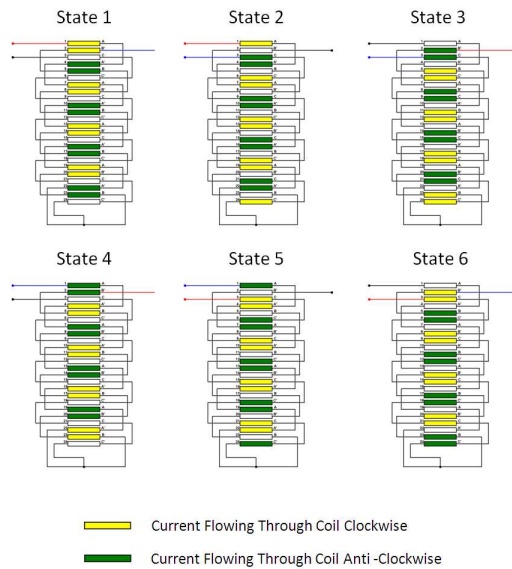


Figure 31: Switching the coils through the circular sequence given by States 1 to 6 gives an apparent motion in the coil switching upwards through the stator as illustrated. Reversing the sequence gives an apparent downwards motion. The pattern repeats itself every six coils. The permanent magnets on the target shaft can be made to interact with the fields produced by passing sequentially through these states to produce motion of the target shaft.

A simulation of the interaction between the magnetic assembly and the coils was produced using a magnetostatic model in Opera 3D [?]. This 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, corresponding to a plane that sits on top of the 13th coil. For this model a current of 58 A was passed through the coils, approximately equal to the peak current that passes through the coils during operation.

The simulation ran 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 are plotted in figure 32.

Simulated Force on the Magnetic Assembly as a Function of Displacement for a Fixed State.




Figure 32: Simulation of the force on the magnetic assembly as a function of its position within the stator when the stator is in one of its six states with a coil current of 58 A. A good sinusoidal fit can be made to these points as illustrated.

Although not illustrated here the sinusoidal pattern of the force curve repeats every 18 mm, the length of one bank of coils. Figure 32 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. 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.

6.4 Zero-Force Points

From figure 32 it can be seen that 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 results in a force that pushes the magnetic assembly further away. However the other zero force point is stable, and 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. The stable zero-force points provide a mechanism by which it is possible to passively levitate the target shaft. 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 than that required to achieve the high accelerations needed to insert the target into the ISIS beam.

The zero-force points in figure 32 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 33. The shaft will undergo damped harmonic motion as it settles at this new point.




Figure 33: If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 5’ then the magnets would see a restoring force (black arrow) that would move the shaft to the new zero force position (red arrow) 3 mm further up the stator. By repeatedly incrementing or decrementing the states the shaft can be moved up and down the stator body. This graph shows the force for a 3 A coil current typical of that used to levitate the shaft.

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’.

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 gives 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.

6.5 Actuation

From figure 32 it can be seen that the zero force point is 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 32 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. 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, as illustrated in figure 34.




Figure 34: If the shaft was sitting at the zero-force point that corresponds to ‘State 4’ and the stator was switched to ‘State 6’(top) or ‘State 2’ (bottom) then the magnets would see a restoring force (black arrow) that would move the shaft towards the new zero force position. During the first 3 mm of movement the shaft would see an average force that is ≈ 96 % of the peak force.

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 would execute damped harmonic motion about this 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 35. 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 32 and comparing this to figure 31 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.

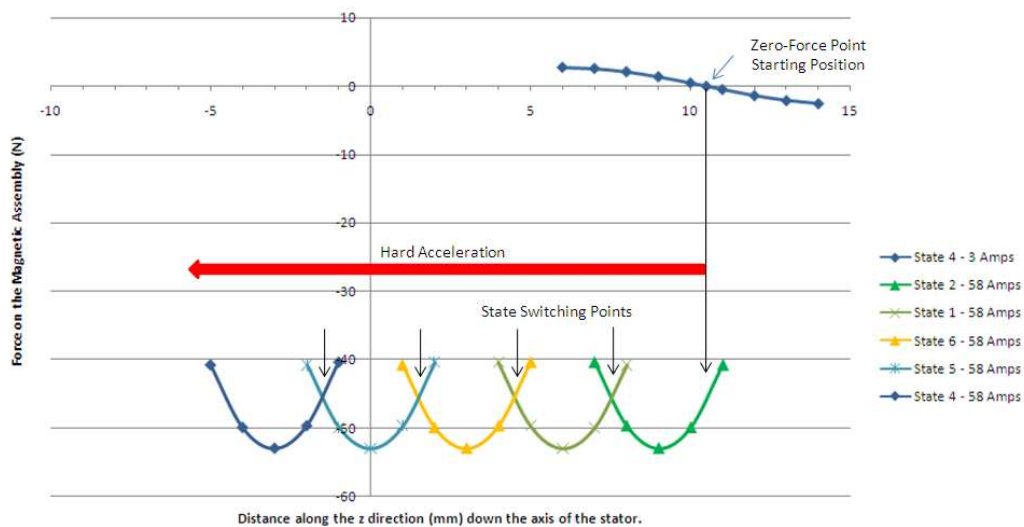


Figure 35: The shaft moves from levitation (3 A coil current) to acceleration (58 A coil current). By tracking the position of the magnets the states can be switched to ensure the shaft continually accelerates using the maximum available force. This results in a hard acceleration of the target shaft.

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. 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 have 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

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 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 36 shows the principle of how it is done.




Figure 36: An actuation. The shaft is accelerated downwards, decelerated, accelerated upwards, decelerated and then captured. The position of the shaft must be tracked accurately so the controller knows into which state to put the coils to maximise the accelerating/decelerating forces.

6.6 Coil Switching and Current Control

It has been shown that by wiring 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 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, and the second is 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.

Switching current bi-directionally through the three phases of the stator can be accomplished using six transistors 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 37 illustrates how these transistors are connected. This type of circuit is known as a ‘Hex Bridge’ or ‘Three Phase Inverter’. The three transistors across the top of the circuit are called the high side transistors as they are connected to the positive power rail, while the other three are the low side transistors.

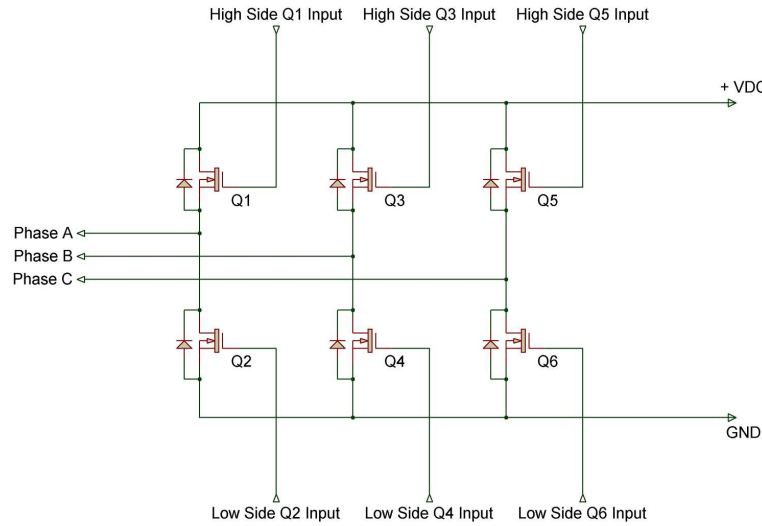


Figure 37: The Hex Bridge. This circuit allows bidirectional current control through any two of the three phases of the stator by application of appropriate control signals to the six inputs that switch the transistors on.

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 37. 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.

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: How the six stator states can be switched using the Hex Bridge.

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

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’. Shoot-through is also possible due 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, to prevent the risk of ‘shoot-through’ occurring. 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 imposed by the controller.

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 which exactly counters gravity. 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 passing through 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 from figure 34; 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 during actuation and to pass a smaller current through the coils for longer periods 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. Most modern power IGBTs have an internal diode across the device to protect against inductive surges.

6.7 The Target Power Supply

Based upon extrapolation from an early 10 A prototype power supply and through simulation, it was expected that the target mechanism would require approximately 60 A of current during actuation to achieve the required acceleration. The specification, design and build for this power supply was managed by 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 such a power supply poses several 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. It is therefore more economical and efficient to power the driving circuit from a capacitor bank which can provide the short high current pulses on demand when necessary. The capacitor bank is then charged by an external power supply that is designed to provide the necessary charging current to top the capacitor bank up and provide the stator's holding current. This capacitor charging unit or 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 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 current is limited to that provided by the CCU. The CCU used only provides a current of 4 A, and into a 3.6 Ω load (the DC resistance of 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 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 has shown that this safety mechanism is very effective at preventing the stator from overheating.

To switch the high currents necessary to drive the target power, Integrated Gate Bipolar Transistors (IGBTs) are 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) are taken to a floating DC power supply as this gives better individual protection to the circuits.

The logic circuitry that drives the gates on the IGBT's is supplied with power from a separate power supply that is integral to the unit. The capacitor charging unit utilises a Xantrex 300-4. 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 assumption of a linear relationship between the duty-cycle and the current through the motor is a reasonable first order approximation. However the finite switching times of the transistors can introduce significant non-linearities as the duty cycle approaches either 0% or 100%. This effect is more noticeable for power transistors due to their longer switching times. If the duty cycle of the PWM signal is expressed as a percentage then this effect can become problematic if the transistor switching time T_{ON} becomes greater than:

$$T_{ON} \ll \frac{\text{Duty Cycle (\%)}}{\text{Frequency}}$$

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.

During development, two stator failures occurred within a period of 18 months due to electrical shorts to ground from a coil. This led to an improvement in the quality assurance procedures in selecting coils for use in a stator. Furthermore 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 [?]. 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.

6.8 System Placement in ISIS

It was clear early on in the project that the installed layout of the target system at ISIS would require considerable distances between the various components in the system. The power supply cannot be close to the target mechanism as the radiation produced by ISIS would likely cause premature failure of the electronics. It is also advantageous to have the electronics accessible so that power supply maintenance can be performed without having to wait for the synchrotron to be opened. 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 be significant. 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.

6.9 Fibre Optic Links

The fibre optic cables that run between the control electronics and the stator's optical block have been discussed in the section on optics. Additionally fibre optic cables are 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 cause additional problems with either the controller or the transistor drivers in the power supply.

This fibre optic link utilises a commercially available optical transmitter/receiver pair. The transmitter is a high powered infra-red LED with a bandwidth up to 5 MHz. This is 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 are the same multi-mode fibres used to return the laser light in the quadrature optical counter, (BFH37-200). These fibres are well matched optically to the transmitters, and experiments in the lab showed lower losses using these cables than 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 gives 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 38 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.

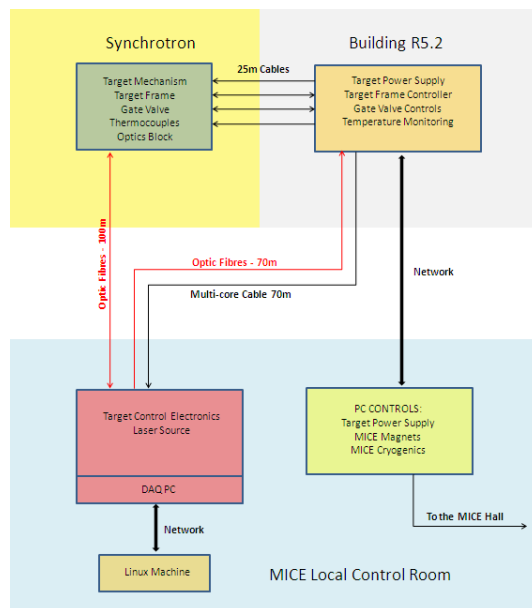


Figure 38: The major components of the target mechanism installation at ISIS, showing the relationship between the components.

7.1 Target Controller Overview

The target controller is implemented on a Field Programmable Gate Array (FPGA), which 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.

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!)

The four states defined for the target mechanism are Parked, Hold Mode, Actuate Enable and Actuating. These are illustrated in figure 39, with the paths which connect them. 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 being dependent upon input from both external signals and the user. The controller’s FSM has many tens of sub-states to ensure that the target operates correctly in a safe manner, though this level of detail is not described here.

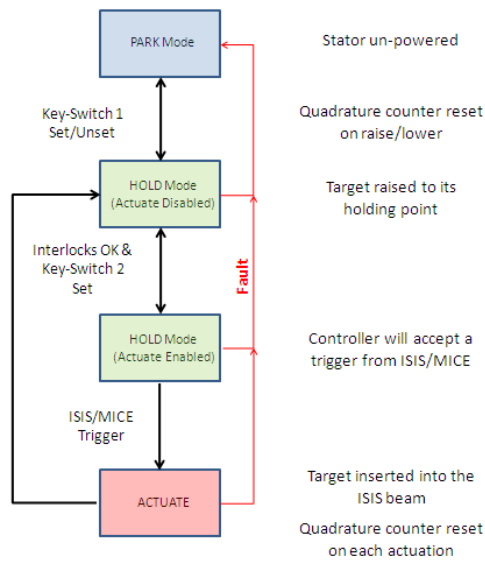


Figure 39: Diagram illustrating the four states of the target system.

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 rests on the bottom bearing and the target is referred to as being in its ‘parked’ position.

Before the target can be actuated it 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 the ISIS beam-pipe.

If the target frame was in its lowered position while the target was in park mode then this would leave the target 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, the controller can be commanded from the GUI to raise the target mechanism to its hold position. The latter 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, 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 performed 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, so that it enters the beam just before extraction and is then removed from the beam-pipe before the next 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 to ensure it does not fall into the ISIS beam causing the accelerator 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 known as ‘quadrants’. The actual capture of the target is considered a separate process in its own right. The four quadrant states are shown diagrammatically with reference to a target trajectory in figure 40.




Figure 40: The actuation trajectory is split into four quadrants from the controller's point of view. Switch Point 1 reverses the coil currents and so determines the actuation depth. Switch Point 2 determines the capture point. These two points are not at the same position due to the decay on the capacitor bank.

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 maintain the maximal force on the shaft until the shaft reaches position 'switch-point 1'. This 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 $150 \mu\text{m}$ this means that the target dip depth is controllable to within $300 \mu\text{m}$ (although there is some actuation to actuation jitter which depends on the amount of bearing wear.)

As the currents are reversed the target begins decelerating and the controller enters '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 it 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. (As described earlier, the zero force points are separated by 3 mm.) It then switches the coils to place this zero force point as close as possible to the target's current

location; 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 the 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.¹

Switch-point two is offset with respect to switch-point 1 because the decay on the capacitor bank means that the rate of acceleration in quadrant one and two is greater than that 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 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 the 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 must 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 for a 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 ($\approx 10\%$ reduction in current during actuation) and gravity ($\approx \pm 1\%$ 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 if it is to catch the beam during the last 2 ms of its time in the synchrotron. The ISIS triggering system has some flexibility and can

¹ 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.

provide a timing signal between ≈ 5 ms before injection up to extraction, but this earliest time is not sufficient to allow triggering from a given ISIS injection signal (ISIS MS) to intercept that same spill 2 ms before extraction.

The timing of the target actuation also needs tuning to optimise the particle production for MICE. This is firstly because the target trajectory is dependent on the actuation depth. The deeper the target has to actuate, the earlier it must be triggered as it takes longer to reach the apex of its trajectory. Secondly the timing of the interception point of the target with the ISIS beam must be adjustable. 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 allow synchronisation of the target with ISIS so that insertion occurs in 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 monitors the beam loss on the specific 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 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 and the output of this second delay is used to determine the actual actuation start time. All of this timing is handled by 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 require different values 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 the experiment 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 incorporates a ‘target ready’ output and a ‘MICE OK’ input allowing it to handshake with an external synchronisation system. This would prevent the target from actuating until the rest of MICE was 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 PC's 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, in close proximity to the target control electronics, that contains a specialised interface - a National Instruments 6254 PCI card with 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 sample rate for the target DAQ so currently only analogue data is recorded. Most of the signals that need to be recorded are naturally analogue, but this strategy 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 recorded are target position, beamloss from ISIS in sector 7, beamloss from in sector 8, and total beamloss. (The target is situated in sector 7 with sector 8 immediately downstream.)

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, 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 laboratory 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.
- [8] **MICE** Collaboration, “The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment,” [arXiv:1203.4089](#). Submitted to JINST (2012).
- [9] N. Schofield, C. Booth, and P. J. Smith, “A low mass, brushless permanent magnet linear actuator for the ISIS target accelerator,” in *50th Annual Conference on Magnetism and Magnetic Materials (MMM05)*. 2005. (Paper FF-12).
- [10] “P.A.R. Insulations and Wires Ltd. Technical data MAGNETEMP CA-200.”
<http://www.par.gb.com>.
- [11] “P.A.R. Insulations and Wires Ltd. Technical data DOLPHON-2102.”
<http://www.par.gb.com>.
- [12] L. Howlett, “Simulation of radiation levels in the MICE target magnets.” MICE Note 166, 2007. <http://mice.iit.edu/micenotes/public/pdf/MICE0166/MICE0166.pdf>.
- [13] Ansoft, “MAXWELL SV. The student version of MAXWELL 2D.”
<http://www.ansoft.com/maxwellsv/>.
- [14] A. Fasso *et al.*, “FLUKA: a multi-particle transport code,” 2005. CERN-2005-10, INFN/TC-05/11, SLAC-R-773.