Methods for muon beam diagnostics

J.S. Lord

October 12, 2011

1 Introduction

In a muon spin relaxation experiment, spin-polarised muons are implanted in a sample of interest. The muons are produced when high energy protons hit a target, initially giving pions which decay to muons. The decay of the pion leads to the muon's spin polarisation which is antiparallel to the muon momentum in the pion's rest frame. The purpose of a muon beamline is to select and transport muons from the target to the sample.

Muons are available with either positive or negative charges, produced from the corresponding charge of pion (although there is a larger flux of the positive species due to the positive charges of the beam and target nuclei). They are implanted in a sample, and having come to rest they interact with the material before decaying. Positive muons generally act as an interstitial hydrogen-like atom while negative muons usually cascade into a low orbit round a nucleus. The decay positron is preferentially emitted in the spin direction — detecting these allows the time dependent polarisation to be calculated.

Muons are charged particles, so a beamline can use magnetic or electrostatic elements to steer and focus the particles. Any muon beamline must deliver muons to a well defined, usually small, sample region. The momentum range of the beam must be well defined so that muons do not stop in the sample cell window or pass through the sample. Other particles, particularly positrons produced from muon decays in the target, must be excluded.

A "surface" muon beamline uses the muons resulting from those pions which come to rest in the production target. As the pion decay is a two-body process the muons have a well defined initial momentum of 29.7 MeV/c. This results in a momentum spectrum outside the target with a cutoff at this value, but extending to lower momentum for those muons produced below the target surface. The polarisation of a surface muon beam is close to 100%, limited only by the angular range accepted and any background of muons formed from pions in flight between the target and the first beamline elements. The advantage is a high intensity beam, though with a limited momentum range, and only positive muons are obtained since negative pions coming to rest are rapidly re-absorbed by the target nuclei.

The alternative design, the "decay" muon beam, collects pions leaving the target and allows them to decay in-flight before the sample. In order to give a

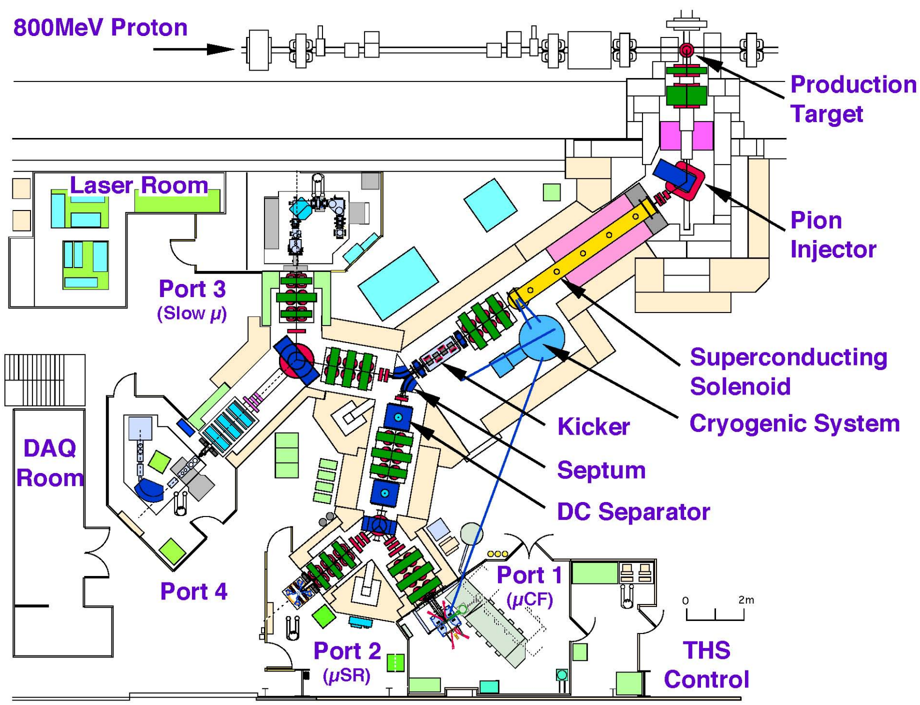

Figure 1: Layout of the RIKEN-RAL beamlines at ISIS, which can produce either decay or surface muons.

high polarisation, the momentum ranges of the pions and muons must be separately selected so that only “backward” (or occasionally “forward”) decaying muons reach the sample. Backward decays are usually preferred since the momentum difference between them and any residual pions (or other particles of the same initial momentum) is greater, giving a cleaner beam. The advantage of such a beam is the much higher momentum available, for example to penetrate the window of a high pressure cell. Positive or negative muons can be selected by changing the polarity of the bending magnets.

A diagnostic system will enable the performance of a muon beamline to be checked, to confirm that a new beamline is performing as expected from simulations, and identify faults if they occur. These may include mechanical misalignment of beamline elements, magnets or other elements not producing the expected field, stray magnetic fields deflecting the beam, or a misaligned primary proton beam or target.

2 Technologies

Muon beams have very low flux compared to the primary proton beam (fractions of pA compared to μA or mA of protons). Therefore electromagnetic techniques such as current transformers, beam position monitors picking up induced charge in the beam pipe walls, or wire grid profile monitors, cannot be used as the signals would be too small. We must use methods which measure the kinetic energy deposited by individual particles in a scintillator, ionisation chamber, microchannel plate or similar detector. Such techniques are usually "destructive" as the muons are then removed from the beam, and the detector would be withdrawn for normal operation — an exception would be detecting particles that would otherwise have stopped at slits.

The time structure of the primary proton beam, and the resulting muon beam, may be either continuous (on timescales of the muon lifetime or shorter) or pulsed (with pulses of length less than the muon lifetime spaced by many muon lifetimes). In addition to the different experiments that can be performed, the diagnostics will have to be designed to take advantage of this. Continuous beams can more easily use particle counting while pulsed beams may have to use analogue detection of a large number of particles hitting a detector nearly simultaneously - though the arrival time should be well known and the time of flight may allow muons, pions and positrons to be distinguished.

For a low flux beam, sufficient particles must be detected to give reasonable statistics. However the timescale of the measurement should be as short as reasonable, for example to distinguish a beam whose position is unstable from a badly-focused beam.

Any measurement made upstream of momentum-separating bending magnets, crossed field separators or slits may include particles of the wrong type or which would not have reached the sample anyway, so simply maximising the gross rate or minimising the beam envelope at this point may be meaningless. However a slit or aperture could be scanned across the beam and the rate further downstream then measured. The collimating slits in the ISIS beamline can be used for this, see figure 2.

3 Rate measurement

The simplest diagnostic value is the rate of muons arriving at the sample position. On a time-differential instrument on a continuous beam, this is already available from the muon (start) counter. Alternatively the positron decay counts can be summed, allowing for detection efficiency and variation of average polarisation which may have varied during an experiment.

A variation is the rate stopping within a defined (small) sample area within a larger beam spot. Again this may be obtained from muon start and active veto counters, or from the positron counts and P(t) on a test sample. One such is haematite (Fe_2O_3) with a silver mask, and a small applied transverse field: the muons in the silver precess at a known frequency while those in the

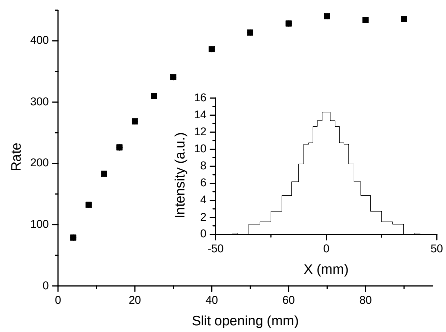

Figure 2: Measured rate in the MuSR instrument as a function of the collimation slit opening width. Inset: reconstructed beam profile at the slit position, assumed symmetric.

Fe_2O_3 see a distribution of much higher fields from the internal magnetism. The simple "background fraction" from this measurement relates directly to the experimental data on a similar sized sample.

To obtain a profile from the test sample, either the sample itself can be moved (if space permits) or the beam can be steered across it. Care must be taken not to clip the beam on any apertures: if the steering magnets are located before the last focusing quadrupole magnets, the steering deflection within the quadrupoles may be larger than that at the sample.

Rate measurements should also be possible at other points in the beamline by inserting a large scintillator to block the beam, taking care which particles are present at that point.

4 Profile measurements

In this case a narrow detector is scanned across the beam aperture, or a row of detectors are spaced across the aperture. Usually there will be a pair of such arrays in x and y directions. An example would be a scintillating fibre, connected to a photodetector. Such a detector assembly has been developed and is in use at PSI [1], using APDs. A similar system with multi-anode photomultipliers has been developed by RIKEN and is in use at the J-PARC facility [2].

The ISIS remote collimating slits have a fibre fixed to each edge which is scanned as the slit is opened, so can also be used to plot the profile. As the two sides are linked and have a common photomultiplier this limits the flexibility — also they only open in the horizontal direction and there is no corresponding vertical slit.

The signal from the profile monitor can be fitted to give an integrated intensity, mean position and width/height, to help with tuning or steering the

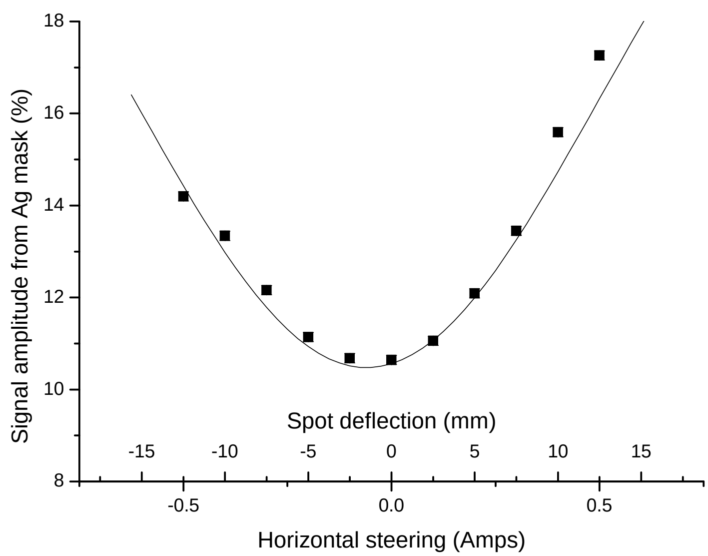

Figure 3: Typical horizontal steering curve from the MuSR instrument, with the spot (slits=8) moved across a 20mm Fe2O3 sample. Runs 23861–71.

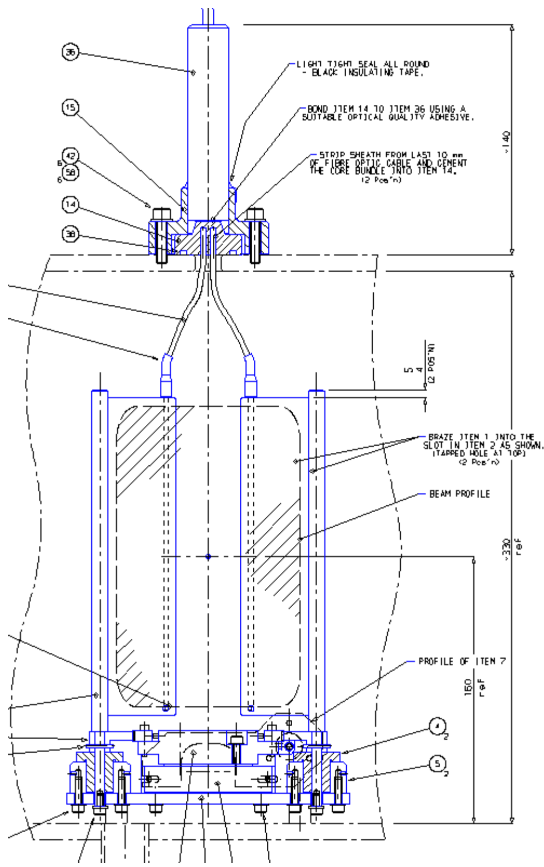

Figure 4: Drawing of the ISIS EC slit assembly, with scintillating fibres fixed to the edges of the slits.

beam.

Profile monitors would usually be in pairs (x and y) to give a full picture of the beam envelope, as with the PSI design. Correct interpretation of the results also relies on the x and y profiles being independent, which should be the case on a beamline with only horizontal and/or vertical bends and "upright" quadrupoles focusing in horizontal or vertical planes. Any focusing by longitudinal magnetic fields (solenoids), rotated quadrupoles or bends, some higher order aberrations, or even clipping by non-rectangular apertures, will couple x and y and break this assumption.

While not strictly a muon diagnostic measurement, the ISIS proton beam includes several profile monitors constructed of wire grids, which are inserted remotely when taking measurements and withdrawn to reduce beam losses when not in use. One of these is positioned immediately upstream of the muon production target to confirm the beam focus and alignment at this point, and gives a guide to the size of the source spot for the muon beam.

5 Full imaging

This can be done with a scintillating screen and camera (or traditionally a sheet of photographic film inserted directly in the beam path). The result will be an analogue integral of the beam profile over the exposure time. The ISIS beam cameras (both the original Eaton version and the more recent HiFi one) work in this manner. With a fast camera the individual flashes from muons hitting the screen can be distinguished and could be used to give a "particle-counting" image.

The optimum thickness for such a screen should just cover the end of the muon's range, where energy deposition is greatest. This could be as little as 0.5mm thick plastic scintillator, and degrading foil can be placed in front, even for surface muon beams. Decay positrons also deposit some energy in the scintillator and if it is too thick, their tracks could extend some distance away from the stopping site and give an apparently larger spot. A thin scintillator which does not stop the muons might give a cleaner, if lower intensity, image but the eventual stopping site of the muons may have to be considered.

Alternatively, if signals from X and Y profile monitors are combined and coincidences recorded, a 2-D histogram image can be built up. An ionisation chamber with x and y wire grids can also operate in this mode, as can a microchannel plate detector with a resistive anode, using charge division. The instantaneous rate must be low enough to avoid too many double hits, so this is easiest to perform on a continuous beam.

6 Beam loss measurement

If we can determine that muons are stopping inside the beamline, when they should have been delivered to the sample, this could be a useful indication of a problem. Beam rates might be measured at several points and subtracted to get beam loss, but this could not be done simultaneously and the measurement uncertainty would usually be too large.

Beam loss monitors using scintillator telescopes could detect the positrons emitted from muons stopped in the beam pipe walls. Positrons will easily escape through a few mm of aluminium or stainless steel pipe (but unfortunately would be blocked by the iron pole pieces of magnets). This monitor would give a "false positive" reading where particles are intentionally stopped, such as slits. A similar system is used on the ISIS proton synchrotron and extracted proton beam, in that case with ionisation chambers located parallel to the beam pipe.

It is not known if such a system is used on a muon beamline. For a much higher power proton beam, promptly detecting losses and automatically turning off the beam is important to prevent damage to the accelerator.

7 Momentum measurement

Electromagnetic elements such as quadrupoles and bending magnets will deflect muons through different angles depending on their momentum. A quadrupole doublet or triplet will therefore have a focal length dependent on the momentum giving chromatic aberration. To limit the resulting increase in spot size or losses, the momentum range transmitted must usually be restricted. A surface muon beam should normally select muons from just below the 29.7 MeV/c cutoff. A decay beam must also ensure that the correct "backward" muons are being selected in order to maintain high polarisation. Within the range of muon momenta transmitted, the flux should be maximised and there should be no dispersion at the sample position.

The central momentum of the beam should be determined only by the fields in the bending magnets, if the beam follows the central path. The momentum spread or "bite" is harder to check and depends on the correct focusing of the beam as well as the setting of any momentum slits.

The momentum range arriving at the sample also determines the penetration depth before the muons come to rest. Although there is a spread of depths even from a mono-energetic beam, due to the random nature of collisions, much of the range is usually due to the momentum spread in the beam. A "range curve" can be measured at the instrument by adding thin foils of one material over a sample of another with a different muon relaxation behaviour.Materials containing only light elements have a slightly lower spread of stopping ranges.

The momentum setting of a surface muon beam can be checked by performing a momentum scan, where all the magnetic elements are varied in proportion to the intended central momentum, in order to maintain equivalent focusing. Any electrostatic elements would need to be varied as the square of the momentum - though in the case of a crossed field separator it is possible to leave the voltage unchanged and vary the magnetic field inversely with momentum to maintain zero deflection for the selected species. The rate will show a gradual rise with momentum until the cutoff, where it will fall to leave only the "cloud"

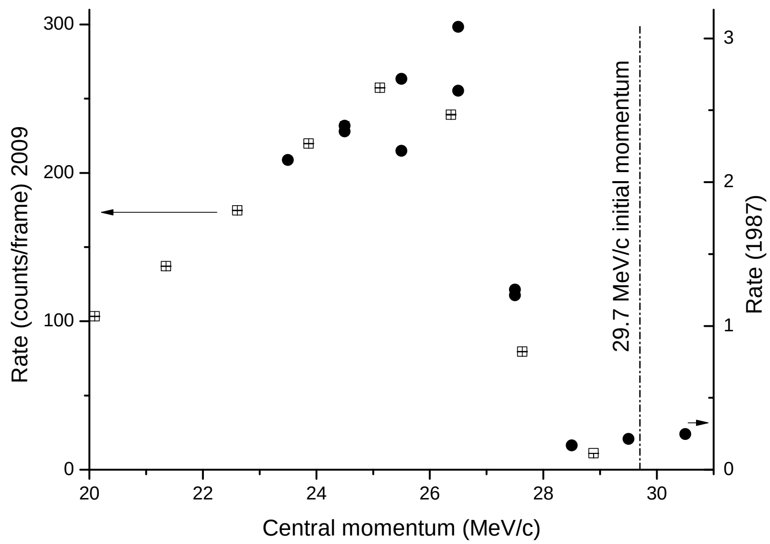

Figure 5: Momentum scan measured on the MuSR beamline at ISIS, by varying the fields in all magnets. Two data sets shown, taken in 1987 (when commissioned) and 2009. The cutoff is lower than the initial 29.7 MeV/c because of a thin Al window between proton and muon beam lines.

muons, accompanied by a drop in polarisation. The sharpness of the drop should be related to the momentum resolution, though this may be affected if there are windows in the beamline between the target and the momentum-selecting bend(s).

8 Spin rotation

A crossed-field separator is usually provided to remove positrons from the beam, since those with the same momentum as the muons, resulting from muon decays in and around the target, would otherwise be transmitted and they may scatter into the instrument's detectors and give a background signal. A spin rotator is a variation on this with higher fields acting over a longer beam path, rotating the muon spins through typically 40–90° relative to the momentum and allowing muon precession measurement in high magnetic fields (since low momentum "surface" muons would be deflected if implantation perpendicular to the field was attempted). On a decay beamline, due to the double momentum selection a separator may not be needed, and a spin rotator may be impractical at high momentum due to the high E-field needed to get a useful rotation — but the higher momentum muons can enter transvsersely to an applied field on the instrument.

Verifying the spin rotation angle can be done using the instrument, a sample showing little or no spin relaxation, and applying weak transverse magnetic fields parallel to that in the spin rotator itself. The observed polarisation will have a phase shift equal to the spin rotation angle.

A separator with its electric and magnetic fields mis-matched will act like a bend, and cause either the position or arrival angle of the muons at the sample to vary, depending on the beamline tune. A reduction of flux is also likely. This can be checked with a position monitor or camera.

The effectiveness of a separator in removing positrons can be checked by scanning the matched electric and magnetic fields starting at zero and monitoring any of a variety of parameters: the total beam flux (including positrons), the pulse height spectra from a thin muon counter (positrons lose less energy than muons of the same momentum in the same path length, allowing them to be distinguished), the time structure on a pulsed beam (the earliest positrons will arrive somewhat ahead of the muons) or even the apparent muon polarisation with a small transverse field applied to a test sample.

9 Kickers and Timing

Many muon beam lines contain pulsed electrostatic or magnetic kickers. On a pulsed beam these may be used to divide a double pulse into two single pulses or reduce the pulse length [3, 4], both of which will improve the frequency response. A continuous beam may be switched away from an instrument once one muon is known to be implanted in the sample, to reduce measurement background and prevent pile-up — the MORE system [5]. In either case the kicker voltage (or magnetic field) waveform should be monitored, together with any beam monitors in the beamlines leading away from the kicker (such as muon counters on the instruments). Figure 6 shows the operation of the ISIS electrostatic kicker where two pulses of width 80ns and spacing 330ns are divided to feed different instruments. The noise on the counter traces at later times comprises both electrical noise picked up in the photomultipliers and positrons from muons implanted in the slits adjacent to the slit counter fibres, then decaying.

10 Logging

Any diagnostic measurements that can be made while the beamline is in operation ought to be logged, to allow investigation of when a fault developed. Instrument count rates are already available in run files. Slit counter amplitudes could also be measured.

Settings of slits and beamline magnets should also be logged. While unlikely to change unexpectedly they will be useful to confirm what beam tune was in use for a past experiment, and so should be included in run files. Diagnostic

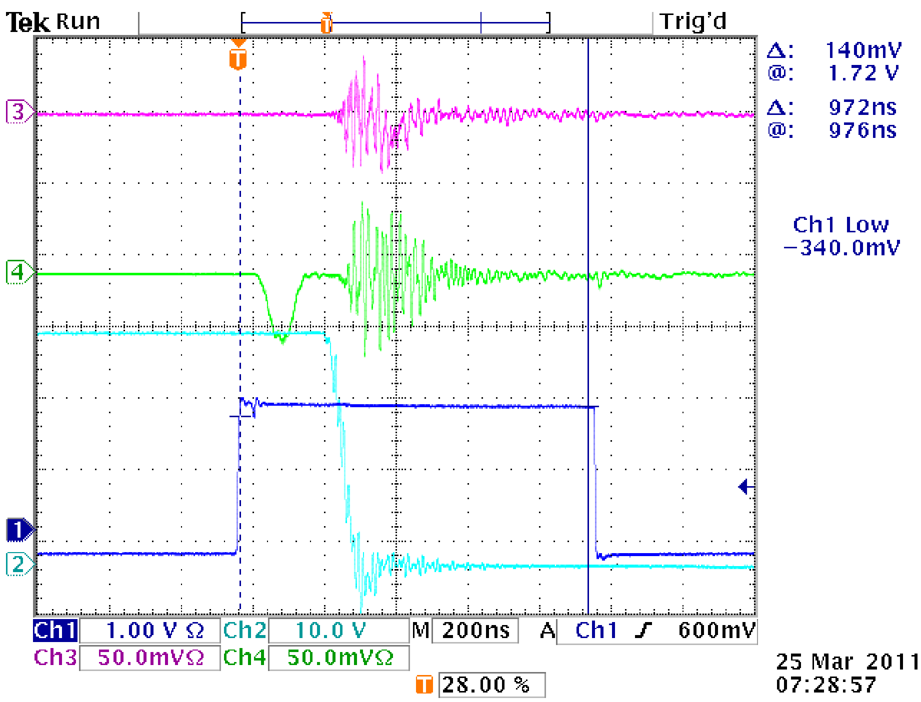

Figure 6: Monitoring signals from the ISIS kicker. From top, Ch3: Slit counter signal from the MuSR beam taking the second muon pulse (coincident with high frequency electrical noise). Ch4: Slit counter from EMU showing the 1st pulse. Ch2: Monitor voltage from the kicker electrode (scaled $\times 10^{-3}$). Ch1: Trigger pulse.

information from equipment, such as magnet voltage drops or kicker timing waveforms, would be useful — this may be easier to use if collected into a single log file for the beamline.

The magnetic fields in the beamline elements could be monitored by Hall probes mounted on the pole tips. These would confirm correct operation of the magnet independently of measuring the current in its coils. Hysteresis in the magnet yoke (although low, due to the use of soft magnetic material) and faults such as shorted turns may be detected.

11 Tuning

A beamline will usually have more adjustable parameters (magnet settings) than measurable results such as rate and spot size. Assuming the dispersion and aberrations are well controlled, the final spot focus can be adjusted using the last quadrupole doublet or triplet alone. For most normal experiments with solid or liquid samples, a well focused spot would never be a disadvantage so there is no need to refocus regularly unless there is a fault. A possible exception would be in high magnetic fields when the focusing effect of the field can be taken into account. With some special sample environment such as high pressure gas cells, the beam might need to be focused through a small diameter beam entry window some distance in front of the instrument centre, while the muons then stop in the centre of the instrument where an applied magnetic field is homogeneous and the detector coverage is optimised.

The beam flux may have to be adjusted in order to maximise the count rate while avoiding pile-up in the sample (on a continuous beam) or in the detectors (for a pulsed beam). It is often best to do this using collimating slits or apertures while keeping the beam tune unchanged.

Adjustments further upstream in the beamline may allow the effective magnification from source to sample to be varied, so that the rate and spot size at the sample vary while keeping an optimum focus. This may be of more use when optimising the beam for experiments on large or small samples. It may still be better to use one of a standard set of beam tunes, obtained with the use of cameras and all other diagnostics, rather than retune from scratch for each experiment.

Automatic tuning procedures are in use at the TRIUMF muon beamlines [6]. Two methods are in use: one scans each element in turn and maximises the rate (or a related quantity such as asymmetry squared times rate). A variation on this can scan coupled elements, for example varying the quadrupoles in a group to change the focal length in one plane while leaving the other unaffected. The second and more commonly used method uses simplex minimisation which copes well with strongly coupled parameters.

Where a beamline is dedicated to one experiment tuning can be done as required to optimise the beam for each setup or sample. Where the extracted muon beam is subdivided or distributed to more than one experiment, such as all the ISIS beams and the piE3 beam at PSI, a compromise tune may have to be chosen at the start of an operating period to get a reasonable rate into each instrument, both for experiments using the full beam on a large sample and those with a collimated beam on a small sample.

A The ISIS beam cameras

A.1 Eaton Mk. I camera

This is a self-contained unit with a camera and 60mm diameter scintillating screen enclosed in a light-tight can. It uses an image-intensified CCD camera, with a real-time "TV" image output. This can be viewed directly on a monitor for interactive tuning, where individial muons are seen as flashes on the screen and build up to give a spot. The camera can be mounted on any beamline where the muons can exit into air and there is sufficient space around and downstream of the sample position to install it (unfortunately it does not fit on the new HiFi instrument at ISIS).

To record images, the original method was to take a long-exposure photograph from the monitor, thus averaging over a large number of muons. More recently the video signal has been digitised on a computer. Real-time viewing is still available, and to measure the spot a video file is recorded and then processed to average its frames. The resulting 320×240 pixel image can then be fitted to a Gaussian profile and width, height, position and intensity extracted.

The electron image intensification stage in the camera is very sensitive to magnetic fields - for example a 100 Gauss longitudinal field accidentally left on will rotate the image by $45°$. It can therefore not be used to follow beam deflection in high fields. To ensure alignment of the image even in the presence of small stray fields, a graticule pattern is engraved on the scintillator sheet and can be observed on the image as a bright line with ticks, when examining an averaged image such as figure 7. This is caused by scintillation light initially captured in the sheet by total internal reflection then escaping at the line. An additional check on alignment is to mount a small "fly-past" sample holder, or set of cross-hairs, in the instrument then bring the camera up behind it. The sample holder (usually 1mm thick Ag) will leave a recognisable shadow on the image. It can then be removed for more accurate spot size measurement.

A.2 HiFi Mk. II camera

To help set up the new high magnetic field muon instrument HiFi, a new camera design was developed. It uses a cooled CCD sensor, as used for astronomical imaging. There is no electron image intensifying tube and the electronics inside the camera head are all insensitive to magnetic fields. To maximise image brightness a high numerical aperture lens is used. This camera has a relatively slow readout so instead of viewing at a 50 Hz frame rate, the image is built up in the CCD and read out after typically 20 seconds to 1 minute. The resolution is higher although for beam imaging this is not critical. A direct USB connection

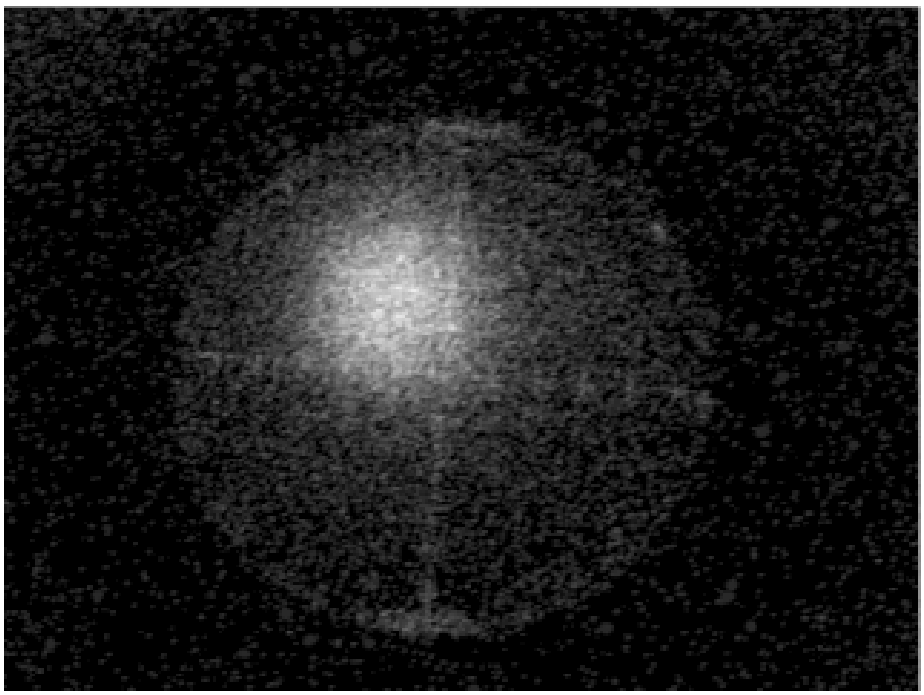

Figure 7: Averaged beam spot picture from the Eaton camera

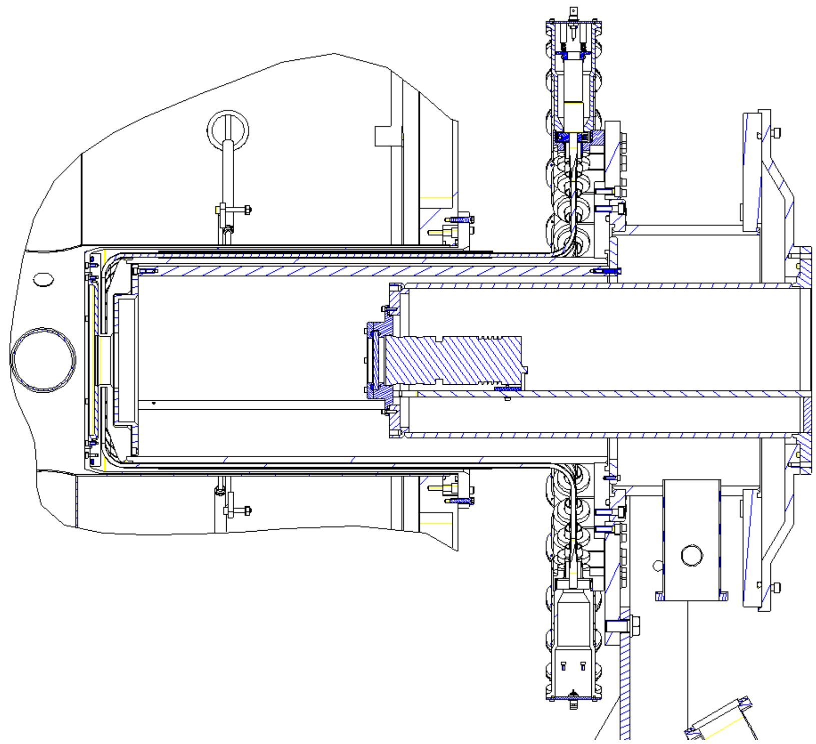

Figure 8: Drawing of the HiFi beam camera in position in the instrument. The CCR with scintillator is inserted through the magnet side port, shown end-on at the left.

makes the camera easier to set up — it is usually run from a laptop computer in the experimental area, which is remotely accessed when the area is locked up and the muon beam turned on.

The scintillator is now mounted directly on the sample stage — in the HiFi instrument we use the closed cycle refrigerator (CCR), with its radiation shield cover left off. The camera and lens are introduced from the downstream direction in a re-entrant vacuum vessel with a window. The design can be used on any beamline where the sample region is enclosed in a light-tight sample tank, solenoid or "cruciform". Our re-entrant vessel is designed to also fit the EMU instrument at ISIS. The front of the lens makes a light-tight join with the window leaving the camera's electrical connections and the focusing ring of the lens accessible. On HiFi the camera views the scintillator through the hole in the downstream detector bank, which cuts off the corners of the image. As the beamline window and upstream detector bank have the same sized hole (60mm diameter) no useful image area is lost.

All other side ports on the instrument are closed off to exclude background light, which would increase the statistical noise level in the camera as well

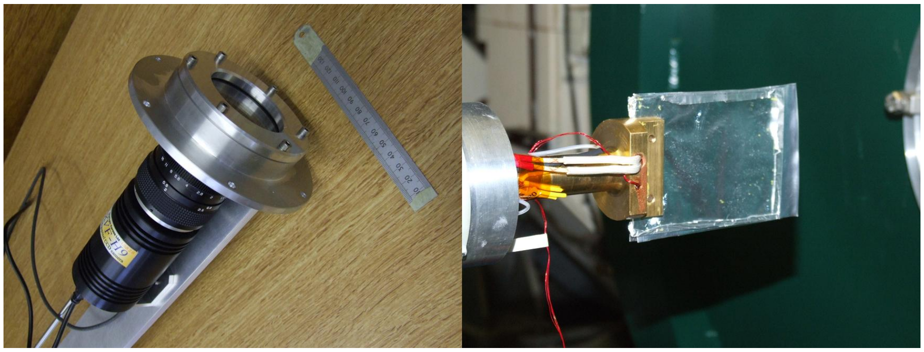

Figure 9: (a) The HiFi camera, lens and window. (b) The scintillator

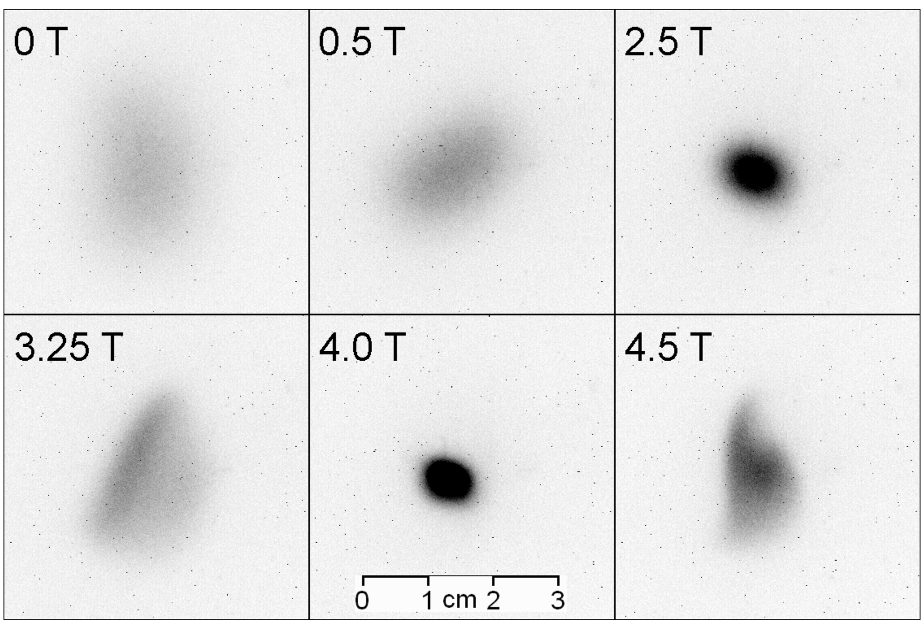

Figure 10: Typical beam pictures from HiFi at different magnetic fields

as possibly introducing reflections that may be confused with the spot. The beamline window is lightly aluminised Mylar and might admit some light from further up the beamline — in practice this has not been found to be a problem, though if necessary a piece of thin aluminium foil could be fixed over it. The degrader foil in front of the scintillator blocks any direct line of sight. When the camera is in use the instrument vacuum is usually pumped to reduce any scatter of the beam in air between window and focal point.

For focusing and alignment of the camera, the scintillator is replaced with a paper grid showing the sample centre, illuminated via another side port on the instrument.

References

[1] A.V. Stoykov, R. Scheuermann, T. Prokscha, C. Buehler, and Z.Y. Sadygov. A scintillating fiber detector for muon beam profile measurements in high magnetic fields. *Nucl. Inst. Meth. A*, 550:212, 2005.

[2] P. Strasser, K. Shimomura, A. Koda, N. Kawamura, H. Fujimori, S. Makimura, Y. Kobayashi, K. Nakahara, M. Kato, S. Takeshita, M. Hiraishi, M. Miyazaki, W. Higemoto, T. U. Ito, K. Ninomiya, K. Ishida, M. K. Kubo, R. Kadono, K. Nishiyama, and Y. Miyake. J-parc decay muon channel construction status. *J. Phys. Conf. Series*, 225:012050, 2010.

[3] G. H. Eaton, M. A. Clarke-Gayther, C. A. Scott, C. N. Uden, and W. G. Williams. *Nucl. Instrum. Methods A*, 342:319, 1994.

[4] A. D. Hillier, S. P. Cottrell, P. J. C. King, G. H. Eaton, and M. A. Clarke-Gayther. *Physica B*, 326:275–278, 2003.

[5] R. Abela, A. Amato, C. Baines, X. Donath, R. Erne, D. C. George, D. Herlach, G. Irminger, I. D. Reid, D. Renker, G. Solt, D. Suhi, M. Werner, and U. Zimmermann. *Hyperfine Interactions*, 120/121:575, 1999.

[6] D. Arseneau. Private communication, 2011.