

UNIVERSITÀ
DI PAVIA

Physics Department

Master's Degree in Physical Sciences

Optimisation and characterisation of scintillation detectors at RAL:

muon spectroscopy for material analysis at the ISIS Muon Source

Thesis by:
Lisa Malliolio

Supervisor:
Prof. Alessandro Menegolli

Co-supervisors:
Dr. Riccardo Rossini
Dr. Daniel Pooley

Academic year 2025 – 2026



UNIVERSITÀ
DI PAVIA

Dipartimento di Fisica

Corso di laurea magistrale in Scienze fisiche

Ottimizzazione e caratterizzazione di rivelatori a scintillazione a RAL:
spettroscopia muonica per l'analisi dei materiali alla ISIS Muon Source

Tesi per la laurea di:
Lisa Malliolio

Relatore:
Prof. Alessandro Menegolli

Correlatori:
Dr. Riccardo Rossini
Dr. Daniel Pooley

Anno accademico 2025 – 2026

Contents

Introduction **1**

1 Muons **3**
1.1 Artificial muons . . . 6
1.1.1 Proton acceleration . . . 7
1.1.2 Target stations . . . 9
1.1.3 Muon production . . . 11
1.2 Positive muons . . . 16
1.2.1 μSR . . . 16
1.3 Negative muons . . . 19
1.3.1 Muonic atoms . . . 20
1.3.2 μ-XES . . . 22
1.3.3 Muon laser spectroscopy . . . 29

2 Scintillation detectors **35**
2.1 Scintillators . . . 35
2.1.1 Organic scintillators . . . 36
2.1.2 Inorganic scintillators . . . 39
2.2 Photodetectors . . . 41
2.2.1 Photo-Multiplier Tubes . . . 42
2.2.2 Silicon Photo-Multipliers . . . 44
2.3 Detectors . . . 47
2.3.1 Geometry . . . 47
2.3.2 Coupling . . . 47

3 Towards the next generation of μSR spectrometers **49**
3.1 The current instruments . . . 49
3.1.1 MuSR . . . 50
3.1.2 EMU . . . 50
3.1.3 HiFi . . . 52
3.2 Super-MuSR . . . 53
3.2.1 Modules and stave . . . 54
3.2.2 Testing and optimisation . . . 58

3.3 Correlated pulses 64

4 CHNet-MAXI **73**
4.1 The experiment 73
4.1.1 Introduction 73
4.1.2 Experimental setup 74
4.2 Beam monitor 78
4.2.1 Data acquisition 79
4.2.2 Equalisation 82
4.2.3 Absolute calibration 83
4.2.4 Spectra normalisation 92

Conclusion **97**

Bibliography **103**

Abstract [Italiano]

La tesi tratta del processo di ottimizzazione e caratterizzazione di scintillatori per l'esperimento CHNet-MAXI, e per gli spettrometri per analisi di muon spin relaxation: Super-MuSR, HiFi ed EMU. Tutti gli esperimenti sono utilizzati per la spettroscopia muonica per l'analisi dei materiali alla ISIS Muon Source, presso il Rutherford Appleton Laboratory a Didcot, UK.
Gli spettrometri della EC facility a ISIS sono utilizzati per la muon spin relaxation e rotation, tecniche che permettono di sondare le proprietà magnetiche dei materiali, utilizzate in chimica, biologia, industria e ricerca nel campo della materia condensata, superconduttori e semiconduttori.
Ho partecipato, grazie al tirocinio ERASMUS+ Traineeship, all'ottimizzazione del prototipo dello spettrometro Super-MuSR, in previsione dell'installazione nel 2027, che consiste in un array di scintillatori plastici letti da coppie di SiPM per il conteggio dei positroni. Lo sviluppo comprende test sulla performance delle fibre e dei SiPM e sull'ottimizzazione sia con fascio che con sorgente, in condizioni sia tipiche che anomale di campo magnetico. Ho inoltre potuto caratterizzare dei segnali anomali negli spettrometri EMU e HiFi, che erano stati notati dopo un intervento al sistema di acquisizione. Il lavoro ha confermato la presenza di questi segnali prima dell'intervento e ha identificato la causa in infiltrazioni di elio all'interno dei fotorivelatori (PMT).
Infine, ho partecipato all'esperimento CHNet-MAXI: uno studio di fattibilità per l'analisi isotopica del piombo in beni culturali che sfrutta la tecnica della μ-XES. L'analisi isotopica non invasiva di artefatti storici permette di localizzare geograficamente i reperti senza compromettere il loro stato. L'esperimento comprende diversi detector, tra cui rivelatori di raggi X (HPGe), scintillatori inorganici e un monitor di fascio a fibre di scintillatore organico. Ho caratterizzato il monitor di fascio, essenziale per la normalizzazione degli spettri dei rivelatori di raggi X e per la quantificazione dell'intensità del fascio: dai segnali è stato infatti possibile estrarre un valore per il flusso di muoni, monitorare lo stato del fascio e i cambiamenti di flusso dovuti a cambi nei parametri del fascio, ed estrarre valori di carica accumulata per la normalizzazione.

Abstract [English]

The thesis describes the process of optimisation and characterisation of scintillating detectors for the CHNet-MAXI experiment and for the Super-MuSR, EMU and HiFi spectrometers for muon spin rotation. The experiments are used for muon spectroscopy for material analysis at the ISIS Muon Source, at Rutherford Appleton Laboratory, Didcot, UK.
The muon spectrometers at the EC facility at ISIS are used for muon spin rotation and relaxation, techniques used to probe magnetic properties of materials: it is used in chemistry, biology, industry and research in condensed matter, superconductors and semiconductors.
Thanks to the ERASMUS+ Traineeship I participated in the optimisation of the Super-MuSR prototype, to be installed in 2027. The prototype was a sector of the spectrometer, which consists of plastic scintillating positron counters read by SiPM pairs. The development involved tests on fibre and SiPM performance, on beam and with a radioactive source, in typical and anomalous magnetic field conditions. I was also tasked with the characterisation of anomalous signals from the EMU and HiFi instruments, noticed after a change to the DAQ system. I confirmed that the anomaly was present before the change, and the problem was identified as helium infiltration in the photodetectors (PMTs).
Finally, I participated in the CHNet-MAXI experiment, which is a feasibility study for the lead isotopic analysis using the μ-XES technique. The non-invasive isotopic analysis of historical artefacts enables the identification of the geographical origin of the objects without damage. The experiment uses several detectors: X-ray detectors (HPGe), inorganic scintillators, and a beam monitor based on organic scintillating fibres. I characterised the beam monitor, which is essential for the X-ray spectra normalisation and the measurement of the beam intensity: from the signals, both the muon intensity and the deposited charge were extracted, and it proved able to monitor the beam status and changes to the intensity correlated to changes of the beam parameters.

Introduction

Since their discovery, muons have been used to advance the knowledge in many areas of physics and science: from their uses in fundamental physics to understand leptonic interactions to the study of cosmic rays, materials, and reaction dynamics, muons have proved a great probe to understand nature.
Muon techniques for material analysis have been in development for over 50 years, relying on both natural and artificial muons: the most prolific are Muon Spin Rotation, Relaxation and Resonance (μSR), a group of techniques based on the Larmor precession of the muon for extraction of the local magnetic properties of materials. More novel is the application of negative muons for Muonic Atoms X-ray Emission Spectroscopy (μ-XES), used for elemental analysis of the bulk of small objects in both historical and industrial applications through the analysis of the characteristic X-rays emitted during and after the formation of a muonic atom.
Both these techniques are in constant evolution: new muon spectrometers are being developed, able to handle increased rate and higher magnetic field, allowing the exploration of many more areas of physics. Improvements in resolution will enable the use of μ-XES for isotopic analysis in addition to elemental characterisation, hugely helpful for the analysis of historical artefacts to pinpoint their geographical origin.
In both areas, scintillation detectors are widely used, for their durability, characteristics and adaptability to different situations.
I had the opportunity of working on the optimisation of the prototype of the next generation of muon spectrometers at the ISIS Muon Source, based on scintillating counters read by SiPMs, while also working on characterisation of anomalous pulses in the current spectrometers based on PMTs. I also followed the muon beam monitor for the CHNet-MAXI experiment from beginning, in my bachelor thesis working on the DAQ, to end with the characterisation of the response during data taking, for my master's thesis.

In this thesis I will highlight the work on scintillation detectors used for material analysis at the ISIS Muon Source, following the ERASMUS+ Traineeship in the Detector System Group I was able to take part in from September to December 2025, with the addition of the work carried out during and after the data taking runs of CHNet-MAXI in December 2025.
In Chapter 1 there will be an overview of muons, muon beam production and the different techniques that utilise these particles for fundamental, industrial and historical

applications, focusing on Muon Spin Relaxation and Rotation, and Muonic Atom X-ray Emission Spectroscopy.
The technical aspects of scintillation detectors will be illustrated in Chapter 2, focusing on scintillating materials and the photodetectors used in the previously-mentioned applications.
Chapter 3 focuses on the upgrade of muon spectrometers at the ISIS EC facility, with the optimisation of the new SiPM-based spectrometer, Super-MuSR, and the investigation of anomalous correlated pulses, later identified ad ionic afterpulses due to gas infiltration in the PMTs of the two other spectrometers, EMU and HiFi. These tasks were important for the ongoing upgrade to the facility, that will wrap up next year during the ISIS long shutdown.
Finally, Chapter 4 will feature the work carried out in the characterisation of the response from the muon beam monitor for the CHNet-MAXI experiment, a cultural heritage experiment for the determination of lead isotopic ratios in historical artefacts using muons. The beam monitor is used for normalisation of the X-ray spectra, while also being able to measure the intensity of a muon beam thanks to the absolute calibration performed.

Chapter 1

Muons

Muons are heavy fundamental particles. Their discovery, in the 1930s, came from the study of cosmic rays in cloud chambers, which revealed a particle with a mass-to-charge ratio 200 times the electron and 1/9 of the proton: a new particle was discovered, with both positive and negative charge [1].
In the following years, with further studies and the development of the Standard Model, the properties of the muons were discovered: they are second generation leptons, with a mass of 105.7 MeV/c^2 and a lifetime of 2.2 μs [2]. The properties of the muon can be referenced in Table 1.1, while the fundamental particles in the Standard Model are illustrated in Figure 1.1.
The electric charge of the muon is -1, while its antiparticle has a charge of +1: because muons are not in ordinary matter, and they are produced in similar quantities in symmetric processes, they will be referred to as *negative muons* and *positive muons* respectively, as charge is the defining factor for their behaviour in matter.
These particles are unstable and decay following the scheme:

$$\mu^- \rightarrow e^- \, \overline{\nu_e} \, \nu_\mu \qquad\qquad \mu^+ \rightarrow e^+ \, \nu_e \, \overline{\nu_\mu} \tag{1.1}$$

Table 1.1: Muon properties [2, 3]

Charge	-1
Spin	1/2
Mass	105.7 MeV/c^2(0.11 m_p or 207 m_e)
Magnetic moment	3.2 μ_p
Lifetime	2.2 μs (free) or less

Muon interactions are studied in many areas of physics, and many experimental techniques utilise them. Depending on charge, muons can either bond with atoms to form muonic atoms, substituting an electron, they can implant into an object and decay carrying information on the local magnetic dipole distribution, or they can pass through matter and provide information on the density distribution. Their uses are varied, in

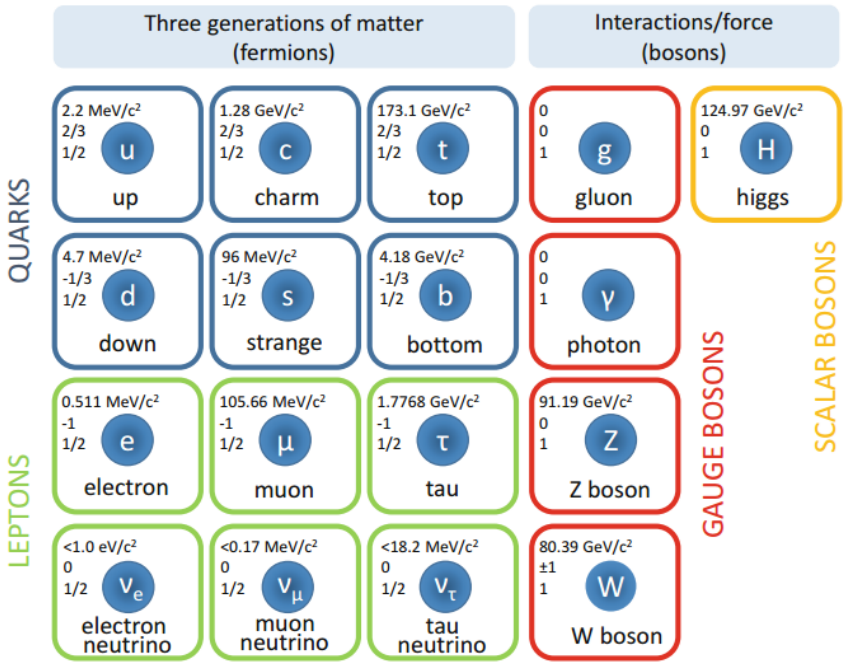


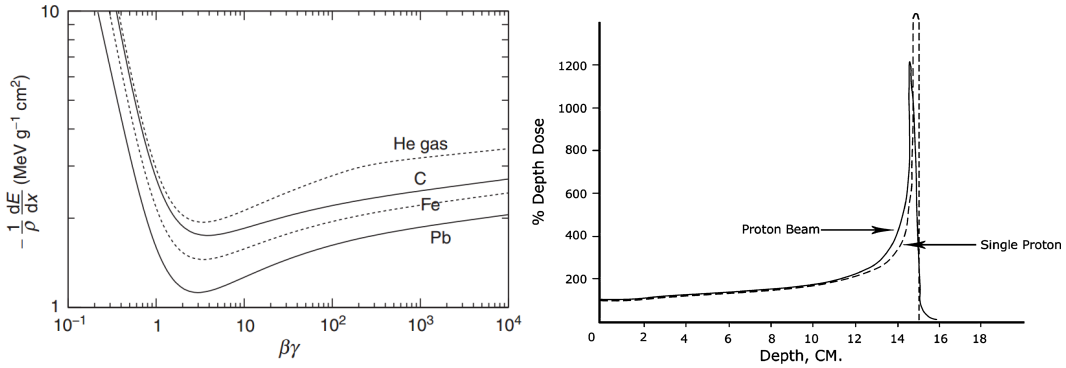
Figure 1.1: Subatomic particles in the Standard Model. The muon is a second generation lepton [3]

fundamental research, as well as in industrial applications and material analysis, with muons coming from both natural and artificial sources.

Muons in matter One of the strong advantages of using muons comes from their nature as heavy charged particles.
Relativistic charged particles, in their passage through matter, lose energy continuously along their track because of ionisation of atoms in the material, through the Coulomb force [4, 5]. The formula that describes the *stopping power*, or the energy loss per unit length of track, is the Bethe-Bloch:

$$-\frac{dE}{dx} \approx 4\pi\hbar^2 c^2 \alpha^2 \frac{nZ}{m_e v^2}\left\{\ln\left[\frac{2\beta^2\gamma^2 c^2 m_e}{I}\right]-\beta^2\right\}\approx$$
$$\approx \rho\frac{4\pi\hbar^2 c^2\alpha^2}{m_e v^2 m_u}\frac{Z}{A}\left\{\ln\left[\frac{2\beta^2\gamma^2 c^2 m_e}{I}\right]-\beta^2\right\} \quad (1.2)$$

In the formula, v is the velocity of the particle and Z its charge in units of e, $\alpha = e^2/(4\pi\varepsilon c\hbar)$ is the fine structure constant. m_e is the mass of the electron, while n is the number density of atoms which can be rewritten as $n = \rho/(Am_u)$, where $A = N + Z$ is the mass number of the target and $m_u = 1.66 \cdot 10^{-27}$ kg is the atomic mass unit. Finally, I represents the ionisation potential and the average excitation, experimentally derived



(a) Energy loss curve for a particle in different materials [4]

(b) Ionisation (dose) as a function of depth in water for proton beams at different energies [6]

Figure 1.2: Bethe-Bloch and Bragg Peak profiles

parameter for each material, and $\beta = v/c$ is the velocity in units of c.

The behaviour of the stopping power can be understood as a function of the velocity:

- $v \ll c$: non-relativistic particles, the stopping power inversely depends on the energy of the particle ($\frac{dE}{dx} \propto \frac{1}{v^2}$). This effect is the source of the *Bragg Peak*;
- at energies more than several hundred MeVs, the stopping power reaches a near-constant minimum of circa 2 (MeV·cm^2)/g (for light materials): particles in this region are already relativistic and they are referred to as *minimum ionising particles*;
- $v \sim c$: relativistic particles, the logarithmic term dominates. The stopping power depends logarithmically on the velocity ($\frac{dE}{dx} \propto \ln \beta\gamma^2$). This is called the *relativistic rise*.

As seen in Figure 1.2a, the profile for the loss of energy for the same particle in different materials is similar, if dependence on density is removed. The difference between charged particles is also minimal, depending on mass only: the only exception are electrons, which are much lighter and have a different energy loss profile.

The dependence of the stopping power to v^{-2} means that the energy loss is the greatest for slow particles: the maximum ionisation happens at the end of range, where all particles stop. This was discovered by Bragg in 1904, and it results in the fact that particles have a well-defined range inside the material, where they lose most of the energy, as seen in Figure 1.2b. The range, dependent on incident momentum, is also called the *implantation* or *penetration depth*, and it must be taken into consideration when performing any measurements with muons (in Figure 1.3 simulations of the penetration depth of muons as a function of mass) [6, 7].

Muon implantation means that muons can be used as local probes for the properties of the materials to be analysed, allowing the analysis of any depth in the bulk without interference from the surface.

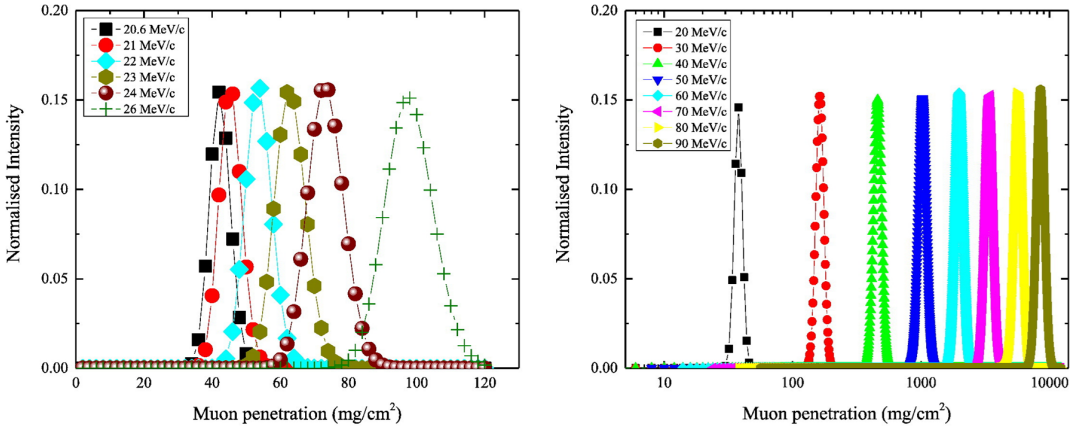


Figure 1.3: Simulations of penetration depth for muons at different incident momentum [7]

1.1 Artificial muons

Muons from both natural and artificial sources have been successfully used for analysis in multiple areas of physics, the choice depending on the type of study and the requirements in flux, energy and direction.

Natural muons are produced in the high atmosphere by the interaction of primary cosmic rays (protons or light ions accelerated by different sources in the galaxy and beyond) with the air: when these energetic particles interact, they produce hadronic showers, where both neutral and charged pions are produced. While neutral pions decay into photons, initiating electromagnetic showers, charged pions most likely decay into muons, which can reach the ground thanks to their relativistic speed despite the limited lifetime. This lead to the discovery of the muon, about one century ago.
Natural cosmic muon rate is very low ($\sim$ 170 Hz/m^2), but it is still useful for specific types of studies in several areas of research, such as muon absorption radiography, muon tomography, and muon metrology, where a uniform flux is necessary. The advantages are the use of a natural, continuous source with no need for dedicated production facilities, no radiological risks, but the low flux extends the time of measurement, rendering the use of natural muons only applicable in quasi-static conditions [8].

For other types of analysis, such as material analysis, laser spectroscopy or spin relaxation studies, the use of natural muons is not advised or possible: a much higher intensity is needed, with momentum selection and collimation. In particular, momentum selection is very important in muon spectroscopy measurements, where the great advantage of the technique is the possibility of depth profiling the object in analysis. In these cases, the use of beams is essential. In the world, only a few selected research facilities produce muon beams: SμS at PSI (Switzerland), TRIUMF (Canada) and MuSIC at RCNP (Japan),

with continuous beams, and ISIS at RAL (United Kingdom) and MUSE at J-PARC (Japan) with pulsed beams [1].
All experiments illustrated in this thesis take place at the ISIS Neutron and Muon Source, so a more thorough description of its methods of production will be explored.

ISIS Neutron and Muon Source is a STFC research facility located at the Rutherford-Appleton Laboratory in Harwell (Didcot, Oxfordshire, United Kingdom). The laboratory, focused on research in physical, life, and engineering science, hosts a 26 m radius proton synchrotron that accelerates protons up to 800 MeV kinetic energy, which are then used to produce both neutrons, for both scientific research and industrial applications, and muons [9]. The facility hosts more than 30 instruments, specialised in different techniques used to probe objects, reactions, or to test electronic ageing and stability in multiple different areas of science, thanks to the use of techniques that can study the atomic structure, dynamics and composition.
A schematic of the facility is illustrated in Figure 1.4.

1.1.1 Proton acceleration

To produce secondary particles, protons need to be extracted and accelerated in a synchrotron [9, 10]. The accelerator system is depicted in Figure 1.5.

ISIS protons start out as hydrogen negative ions, sourced from hydrogen plasma in the **Pulsed Ion Source**. Ionised hydrogen gas produces both H^+ and H^-, in equal quantity: to skew production toward negative ions, the extraction of additional electrons from the cathode metal surface is enhanced by coating the cathode with a thin layer of Group-1 elements, such as caesium used by ISIS. The additional electrons can bind with the hydrogen atoms producing more negative ions. These are extracted and selected using electrodes and a 90° bending magnet, to remove H^{2-}, H^+ and residual electrons. The source works with DC current at -35 kV, meaning extracted ions have 35 keV energy.

After extraction, the ions are transported to the RFQ accelerator via the **Low Energy Beam Transport line** (LEBT), where focusing magnets prevent the beam from expanding and align the beam with the RFQ entrance. The **Radio-Frequency Quadrupole** uses radio-frequency electric fields to accelerate the ions, up to 665 keV, and to create bunches ready for injection in a linear accelerator.
The next stage of acceleration is the **Alvarez-type linear accelerator** (linac), $\sim$ 50 m long, with four successive hollow copper drift tubes (tanks); it is operated at 202.5 MHz, and less than 1°C to eliminate effects of thermal dilation. It can accelerate the ions up to 70 MeV, with pulses lasting $\sim$ 200μs. The beam is delivered to the synchrotron via the **High Energy Drift Space** (HEDS), a 50 m beam transport line which includes 0.3 μm thick carbon stripping foils, used to remove the electrons from the negative ions to obtain a bare proton beam.

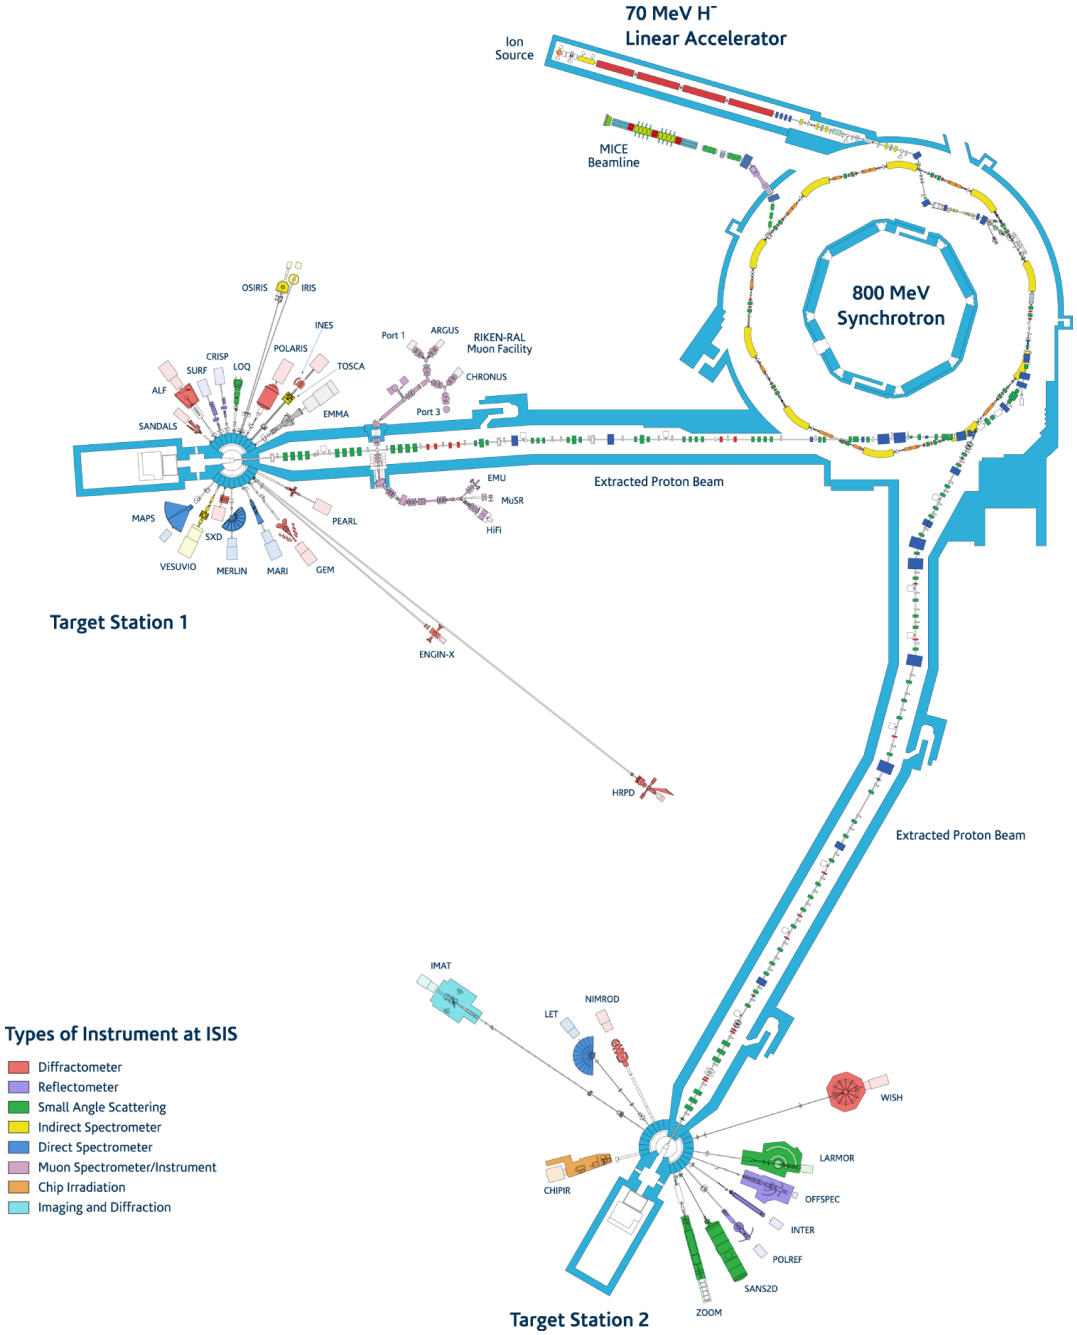


Figure 1.4: Schematic of the ISIS accelerator. At the top the ion source and initial accelerators, which feed the synchrotron. On the left, Target Station 1 with the muon facilities, at the bottom Target Station 2 [10]

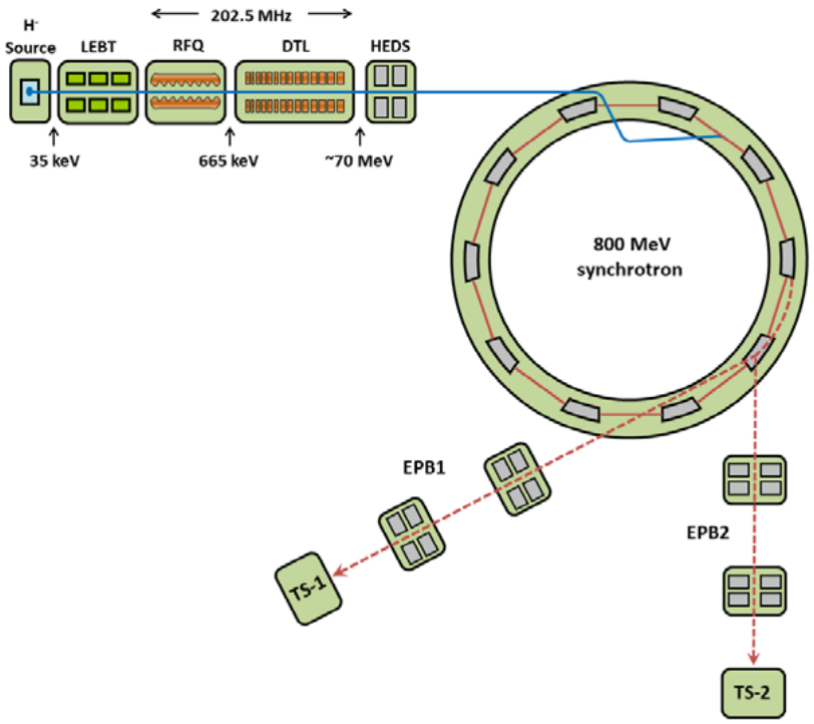


Figure 1.5: Close up of the accelerator system in ISIS [9]

The 70 MeV protons are injected into the **synchrotron**, a 26 m radius (163 m circumference) accelerator which uses 10 sections of RF cavities, quadrupoles, and dipoles to accelerate protons to 800 MeV kinetic energy. The synchrotron works with two co-rotating bunches, accelerated at the same time on opposite sides of the circumference. Each rotation in the synchrotron, protons pick up ~ 0.1 MeV, so each accelerating cycle is 8000 rounds, in the end reaching 84% of the speed of light. The synchrotron works at a 50 Hz alternating current, biased sinusoid, meaning an acceleration cycle must be 20 ms long: the beam is extracted from the synchrotron via the **Extracted Proton Beam lines**: EPB1 and EPB2.
These two beamlines carry the protons respectively to Target Station 1 and 2, where particle production can happen.

1.1.2 Target stations

The facility is split into two target stations, Target Station 1 (TS1) and Target Station 2 (TS2). The first is the original facility, which hosts the muon target and one of the neutron targets. The main focus is historical sample analysis and material analysis. TS2, on the other hand, is newer and more efficient, with several neutron experiments, employing a wide variety of techniques.

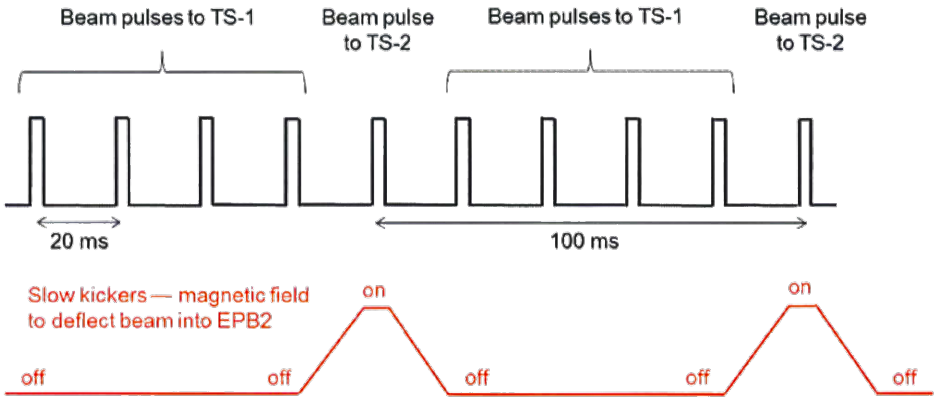


Figure 1.6: Beam pulses scheme: one out of every 5 pulses is directed towards TS2, making the effective rate 40 Hz for TS1 and 10 Hz for TS2 [9]

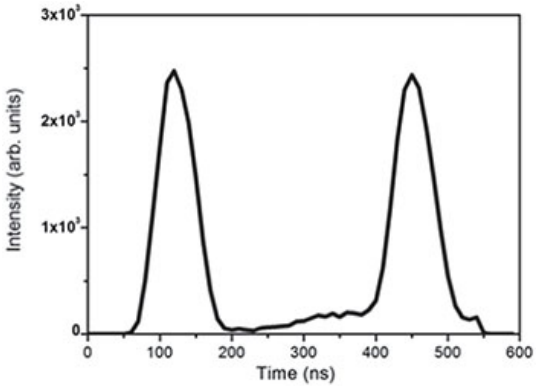


Figure 1.7: Temporal double bunch structure of a spill [11]

The beam is directed to the target stations via the EPB lines: one every 5 pulses (or spills) is directed to TS2 by using kicker magnets, as seen in Figure 1.6 [9].

Each spill has a charge of ~ 4.5 μC, or 2.8×10^{13} protons, corresponding to a current of 192 μA to TS1 and 45 μA to TS2. Each spill consists of two bunches, as seen in Figure 1.7, corresponding to the two co-rotating proton bunches in the synchrotron. At 84% of the speed of light, a full rotation of the 163 m synchrotron takes about 640 ns, meaning that the bunches are separated by 320 ns, with a FWHM of 70ns. The spills are 20 ms apart due to the 50 Hz acceleration cycle.
Interactions in the target involve circa 5% of the protons, leaving 95% of the beam for neutron production downstream.
For each spill, a digital signal is provided for synchronised triggering from Cherenkov detectors placed inside each beamline.

Only Target Station 1 is equipped with a muon producing target, feeding two facilities: the European Commission (EC) facility and the RIKEN-RAL facility, both named after the collaborations that lead to their creation.

1.1.3 Muon production

To get from protons to muons, we leverage the same process that produces natural muons, which is the decay of pions produced from energetic cosmic rays colliding with atoms in the high atmosphere. In a laboratory, pions are produced via the collision of energetic protons with light elements. In particular, at ISIS the target is a graphite sheet, 10 mm thick, set at 45° with respect to the beam direction [9].

Pions are the lightest hadrons: they are mesons, bound states of two first-generation quarks. They exist in three different charges:

$$\pi^+ = u\bar{d} \qquad \pi^0 = \tfrac{1}{\sqrt{2}}\left(u\bar{u} - d\bar{d}\right) \qquad \pi^- = \bar{u}d \tag{1.3}$$

Pions are produced via the following processes [1], either alone (Equation 1.4), for $E_p > 280$ MeV or in pairs (Equation 1.5, for $E_p > 600$ MeV), depending on the energy of the colliding proton:

$$p^+\ p^+ \to p^+\ n\ \pi^+ \qquad p^+\ n \to p^+\ p^+\ \pi^- \tag{1.4}$$

$$p^+\ p^+ \to p^+\ p^+\ \pi^+\ \pi^- \qquad p^+\ n \to p^+\ n\ \pi^+\ \pi^- \tag{1.5}$$

The pions can be both positive and negative (ignoring the neutral pions because they decay overwhelmingly into two photons, which is irrelevant in the matter of muon production).
Charged pions are most likely to decay into a leptonic channel. The decay is:

$$\pi^+ \to \ell^+\ \nu_\ell \qquad \pi^- \to \ell^-\ \overline{\nu_\ell} \tag{1.6}$$

The decay in the electronic channel is highly suppressed compared to the muonic channel, by a factor of 10^4 $\left(\frac{\Gamma_\mu}{\Gamma_\mu+\Gamma_e} \simeq 99.998\%\right)$: this is due to helicity suppression. Weak interactions can only couple with LH particles and RH antiparticles. Neutrinos can be approximated as massless, so their helicity and chirality states are the same. The pion is a spin-0 particle, so the lepton-neutrino system must be produced in a spin-0 singlet state: the only way this can happen is with neutrino and lepton spin in opposite directions. Antineutrinos must be RH[1], so for conservation of angular momentum the charged lepton must have the same helicity. If the charged lepton were massless, the decay would be impossible because chiral states and helicity states would be equivalent and a particle could not be in a RH state. This is not the case, and we can decompose a RH helicity spinor into its chiral components:

$$u_\uparrow = \frac{1}{2}\left(1 + \frac{p}{E+m}\right) u_R + \frac{1}{2}\left(1 - \frac{p}{E+m}\right) u_L \tag{1.7}$$

[1]This is considering the decay of a π^-, for the decay of a π^+ change RH for LH

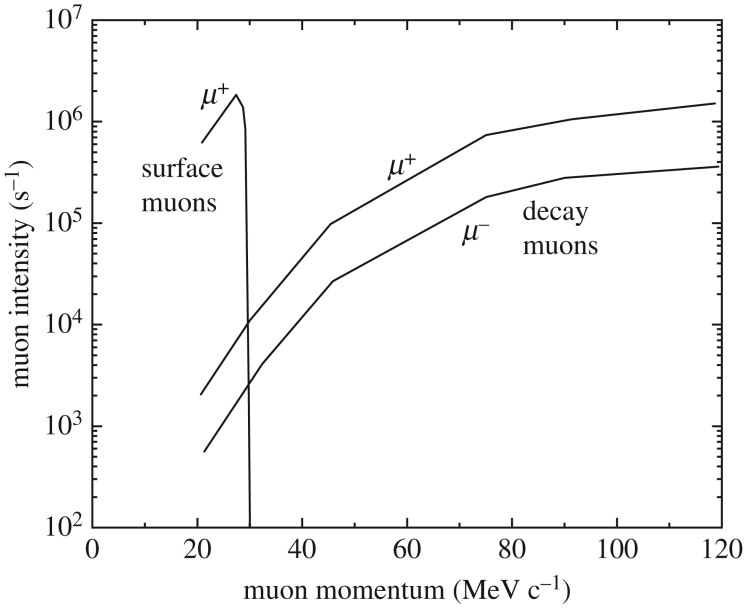


Figure 1.8: Simulated instantaneous intensity for decay and surface muon beams produced at ISIS [10]

Only LH components can interact in a weak vertex, giving an amplitude that is dependent on the mass of the lepton:

$$\mathcal{M} \sim \frac{1}{2}\left(1 - \frac{p}{E+m}\right) u_L \sim \frac{m_\ell}{m_\pi + m_\ell} \tag{1.8}$$

Due to the greater mass of the muons, the electronic channel is suppressed by a factor that is $\sim \left(m_e/m_\mu\right)^2$ [4].

The daughter particles are emitted back-to-back, in the Centre-of-Mass rest frame the muon momentum is 29.65 MeV/c [12].
The interaction of protons with the target can generate two different types of beamlines, differentiated by the behaviour of the pions in the target itself [1]:

- **Surface muons**: pion production happens at rest close to the surface of the target. It can only happen for positive muons, as negative muons are captured in the target before decay. The advantage of using a surface muon beam is the high flux and almost complete spin polarisation ($\sim$ 95%) due to the pion decay at rest. The momentum of the beam is fixed at $p = 29.65$ MeV/c.
- **Decay muons**: pions are collected and transported away from the interaction point. The decay happens in flight, and muons can be selected either emitted along the pion momentum or against (forwards and backwards). The maximum momentum is limited by the proton accelerator, while the minimum momentum is 29.65 MeV/c for forward emitted muons and almost 0 MeV/c for backward emitted muons, albeit with limited intensity. The selection happens in a long solenoid,

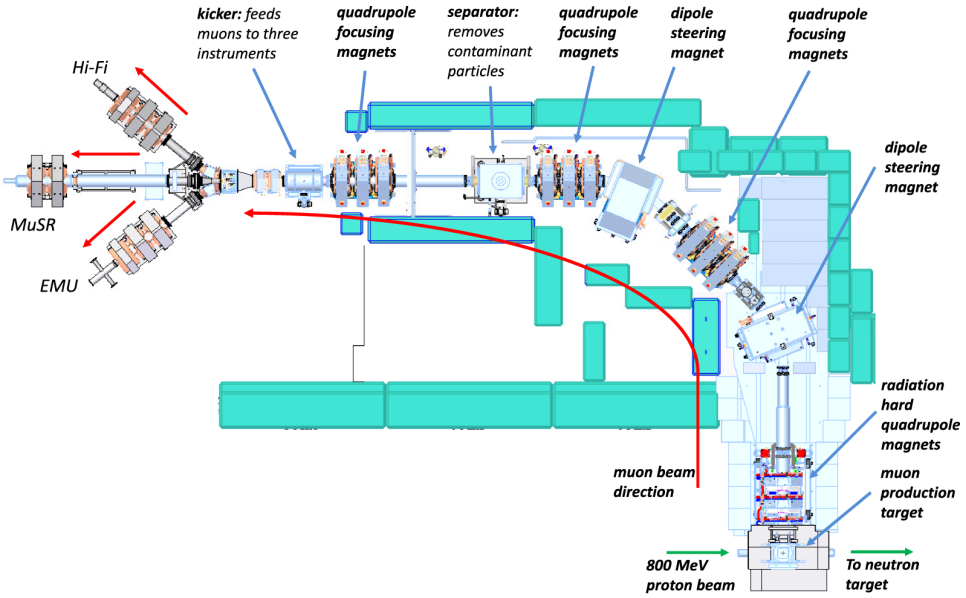


Figure 1.9: Schematic of the EC facility at ISIS [10]

which selects the momentum of pions which maximise the flux for the selected muon characteristics. Muons in decay beamlines can be both positive and negative, with a wide range of momentum (at ISIS, the maximum momentum is ~ 120 MeV/c).

The choice between the different type of beamline comes down to the choice between charge and momentum selection, but with higher costs and lower intensity, or fixed momentum and charge but higher intensity and less complicated beamlines. At ISIS, both are present: the surface muons feed the EC facility, with three dedicated μSR experiments, while the decay muons are used in the RIKEN-RAL facility, with four multipurpose ports. The simulated flux of the two beamlines as a function of momentum can be seen in Figure 1.8.

EC The EC beamline, commissioned by the European Commission in 1987, collects the positive surface muons, with fixed momentum (29.65 MeV/c, or energy of 3.8 MeV). A schematic of the beamline can be seen in Figure 1.9.
The collection of muons from the target happens via three quadrupole magnets, which select the muons from other particles produced in the interaction. The beam is then directed to the experimental rooms, through a 90° bend via two dipole magnets. Between the magnets there are an additional pair of quadrupoles and a set of momentum slits. The slits are used to reduce the momentum spread of the beam, removing muons with momentum greater or smaller than a 2% spread, with the additional perk of reducing the intensity of the beam at the instruments if needed. After the bend, a focusing quadrupole triplet is used.
The final piece of the beamline consists of a separator and a kicker. The first removes

contamination from other particles (positrons that have decayed in flight, protons form the beam...) by having perpendicular electric and magnetic field set to only let muons through (Wien filter). It leverages the Lorentz force ($\boldsymbol{F} = q\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}$): the fields can be tuned to exactly balance the muon velocity, leaving their path undisturbed. The kicker is instead used to split the beam into the three experimental rooms, allowing the simultaneous use of all three instruments. It is an electrostatic kicker, consisting of three plates: the external two are grounded, while the middle one is set at 37 kV. This way, muons passing on the left are directed towards EMU, and on the right towards HiFi. After the first bunch has passed, the voltage is dropped and the second bunch goes straight to MuSR [10].
By using a kicker, all experiments can work simultaneously with a reduced flux (10^6 muon/s for MuSR and 5×10^5 muons/s for EMU and HiFi), which is only possible thanks to the high flux provided by the surface muons. Separating the bunches is also essential for the technique explored at the EC experiments, which focus on muon spin relaxation (μSR), a technique used to probe the magnetic properties of materials in both industry and science. The two muon bunches are too close together to allow the natural muon decay, rendering the entire technique unusable, so a single bunch can be used per instrument [1]. The μSR technique will be explored in Section 1.2.1 and the EC instruments in Chapter 3.

RIKEN-RAL The RIKEN-RAL facility was born from the collaboration between RAL and RIKEN, started in 1990 and culminated in the first beam delivery in 1994, building off of the EC target installed a few years prior.
The RIKEN-RAL facility (in Figure 1.10) hosts 4 different ports, where a high intensity beam of positive or negative muons, from 17 MeV/c to 120 MeV/c with 4% momentum spread, can be directed. It employs a similar layout as facilities at KEK, PSI and TRIUMF, with the decay happening in a long solenoid and the muons then collimated and directed towards ports [3, 10, 13].

Pions extracted from the graphite target are focused with quadrupoles, then transported to the superconducting solenoid magnet via the pion injector, a dipole bending magnet. Inside the 5.5 m solenoid, the pions decay in flight into muons, which stay confined into a beam thanks to the magnet. The solenoid is tunable, up to 5.1 T, allowing for the selection of both positive and negative muons, maximising the intensity for the selected momentum.
After another set of quadrupoles, there is a kicker, analogous to the one in the EC beamline but with only a single separation, splitting the two bunches to different ports (only usable up to 65 MeV/c). Finally, other sets of bending and focusing magnets direct the beam to the selected port.
The peculiarity of the facility is the presence of 4 ports, dedicated to different activities:

- *Port 1*: originally designed for muon catalysed fusion, it now houses the FAMU experiment (see Section 1.3.3);
- *Port 2*: μSR spectrometer for condensed matter research (ARGUS), which can use

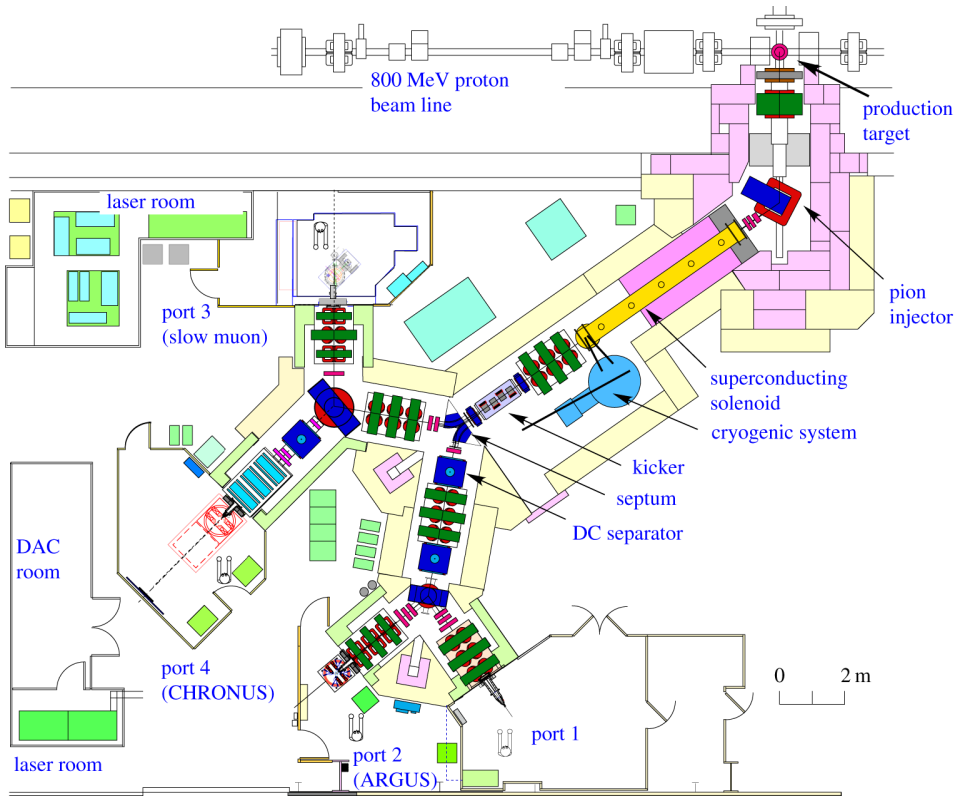


Figure 1.10: Schematic of the RIKEN-RAL facility at ISIS [10]

both positive and negative muons with tunable momentum, allowing analysis at different depths;

- *Port 3*: multipurpose, originally for the development of a slow muon facility;
- *Port 4*: dedicated to elemental analysis using μ-XES (MuX) and irradiation of electronic components.

The kicker allows the ports to either receive a double pulse or a single pulse, the other (with the same momentum) directed to another port. This is especially needed for μSR.

At each port is a lead beam nose, with a mylar window to separate the vacuum of the beam from the air. The nose, 4.5 cm of inner diameter, is placed right after a set of quadrupoles that focus the beam onto the sample position. In addition, a set of lead collimators can be used to reduce the size of the beam spot, with the perk of reducing the amount of muons hitting the target, helping in some cases with saturation of slower detectors such as HPGe. To avoid detector saturation, one of the bending magnets can also be adjusted, defocusing the beam and reducing the flux: this process is handled via a *detune* parameter in the beam setup process.

1.2 Positive muons

At ISIS, the main use of positive muons is for μSR, using the surface beam at the EC facility and decay beam in the RIKEN-RAL ARGUS port.

1.2.1 μSR

μSR refers to the technique of muon spin relaxation/rotation, where a positive muon[2] implants inside the sample, then decays carrying information about the magnetic properties of the sample itself [1].
Surface muon beams are always polarised, thanks to pion decay at rest, and when implanted, if a magnetic field is present (either external or the one generated by the local atoms), the muon spin will start precessing at a frequency dependent on the field:

$$\omega = \gamma_\mu B \tag{1.9}$$

This phenomenon is called **Larmor precession**, and it can be leveraged in such experiments to extract the information on the intensity of the field. γ_μ is the muon gyromagnetic ratio, measuring at $2\pi \times 135.5$ MHz T^{-1}
Another important note is that positrons produced during the muon decay don't have isotropic emission, but they are preferentially emitted along the direction of the muon spin, as noted in Figure 1.11a [1].

Leveraging this asymmetry in positron decay we can measure the average polarisation of the muons.

Considering a muon beam with spin polarisation along the z axis, and a magnet applying a B field along the x axis, the Larmor precession will happen in the $y-z$ plane, with frequency given by Equation 1.9. The polarisation along z can be measured as a function of time:

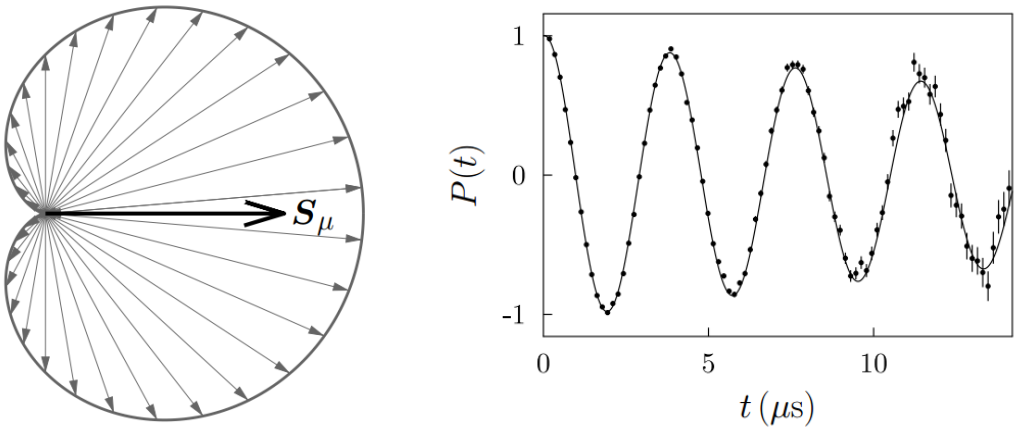
$$P(t) \propto \frac{N_F(t) - N_B(t)}{N_F(t) + N_B(t)} \tag{1.10}$$

This is done by measuring the asymmetry in the decay of positrons, in practice by placing identical detectors in front and behind the sample, perpendicularly to the B external field. In Figure 1.11b the polarisation oscillation can be appreciated.

Multiple geometries are possible, comparing the direction of the beam and the external field, as seen in Figure 1.12 [1]:

- **Longitudinal field (LF-μSR)**: the magnetic field is applied parallel to the muon spin polarisation;
- **Transverse field (TF-μSR)**: the magnetic field is applied transverse to the muon spin polarisation;

[2]Positive muons are most often chosen to avoid the formation of muonic atoms and to leverage surface muon beams, but μ^-SR is also possible



(a) The probability of emission of the positron as a function of direction, indicated by the length of the arrows [1]

(b) Muon precession in a 2 mT external magnetic field, as visible from the positron counts asymmetry in the front and back direction [1]

Figure 1.11: Fundamental properties that allow the use of μSR as a technique

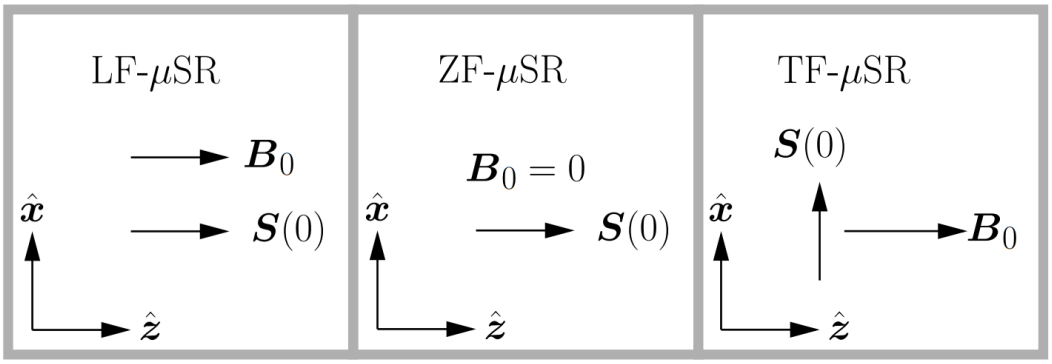
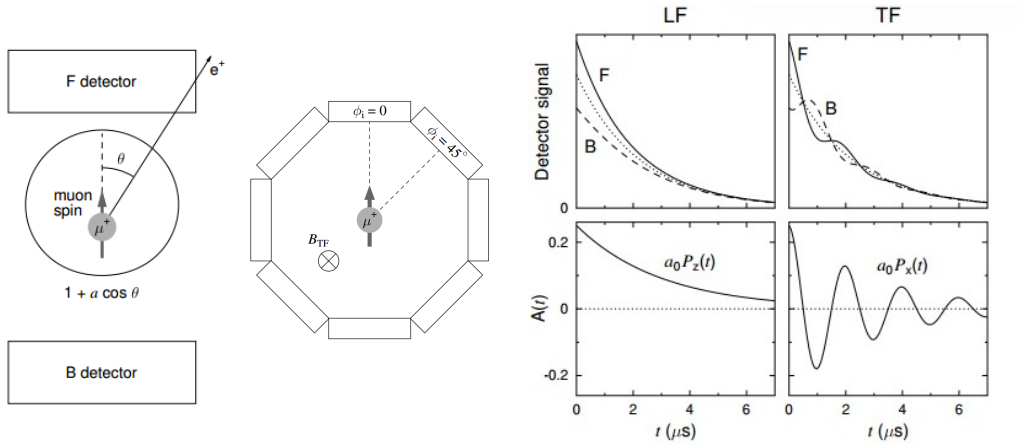


Figure 1.12: Spin and field directions for the various experimental configurations [1]

- **Zero field (ZF-μSR)**: no field is present, using the external magnet to counteract the Earth's magnetic field.

The types of experiments that can be done with this setup are multiple, but not all give the best results: for example, resonance experiments (such as NMR or EPR) are not usually done with muons, although possible, because muons have a limited lifetime, reducing the potential for spin manipulation via RF and because muons are not ordinary matter, and thus introduce perturbations in the object.
The most common technique is the analysis of the *relaxation* (analysing time dependence of polarisation/relaxation functions), and the analysis of the *rotation* (analysing



(a) Setup for a LF-μSR experiment [1]

(b) Setup for a TF-μSR experiment [1]

(c) Asymmetry (polarisation) curves for LF and TF configuration [1]

Figure 1.13: A simplified experiment, with results for different configurations, analysing fields that are longitudinal (LF) or transverse (TF) to the initial spin direction

frequency shifts) [1].

In a simplified experiment, as seen in Figure 1.13a, there are two counting detectors, placed upstream and downstream of the muon implantation, along the spin direction. In a longitudinal configuration, the counts in the detectors are:

$$N_F = N_F^0 e^{-t/T_\mu}[1 + a_0 P_z(t)] \tag{1.11}$$

$$N_B = N_B^0 e^{-t/T_\mu}[1 - a_0 P_z(t)] \tag{1.12}$$

P_z is the longitudinal polarisation function, which is equivalent to the relaxation function in LF configuration. To extract the function, the forward-backward asymmetry is calculated:

$$A(t) = \frac{N_F - N_B}{N_F + N_B} = \frac{(N_F^0 - N_B^0) + (N_F^0 + N_B^0)a_0 P_z(t)}{(N_F^0 + N_B^0) + (N_F^0 - N_B^0)a_0 P_z(t)} = a_0 P_z(t) \tag{1.13}$$

The last equality holds true if the detectors are the same in terms of geometry and sensitivity ($N_F^0 \equiv N_B^0$), or if they can be balanced by introducing a balancing factor ($\alpha = N_F^0 N_B^0$) [1].

In the transverse field configuration, a different polarisation function is probed, not equivalent to the relaxation function (G). In this case, we define the x axis as the spin direction, and the B field is directed along the z axis. We find the polarisation function as before, from the asymmetry, but the relation between polarisation and relaxation differs.

$$P_x(t) = \cos(\gamma_\mu B_{TF}\ t) G_x(t) \tag{1.14}$$

In this case, sometimes a setup such as the one in Figure 1.13b is used. In that case, for each detector the polarisation function contains a phase offset:

$$P_i = \cos(\gamma_\mu B_{TF}\ t + \phi_i)G_x(t) \qquad N_i(t) = N_i^0 e^{-t/\tau_\mu}[1 + A_i] \tag{1.15}$$

Front and back, in this case, are defined by the i-th detector and the one placed exactly opposite [1].
The two different results for the polarisation calculation are depicted in Figure 1.13c.

A more in-depth description of muon spectrometers will be in Chapter 3, with the description of the HiFi, EMU and MuSR instruments.

Applications

The applications of μSR rely on the implantation of the muon, which makes the muon a local probe of the magnetic properties of the material.
Properties that can be derived from a μSR experiment are the detection and observation of magnetic order, spin fluctuations, local electronic environment in semiconductors...
This is possible because the polarisation and relaxation functions can provide important information on both static and dynamic local magnetic fields in magnetically ordered materials, with contributions from different sources such as dipoles and hyperfine fields [1].
One of the industrial applications is for the analysis of superconductors, identifying key magnetic characteristics (such as London penetration depth) but also thermal and resistive characteristics as a function of magnetic field.
The study of impurities in semiconductors is also a common use of the technique, which extracts information on the hyperfine couplings and the dynamics of impurities.
Another industrial application is batteries, where with μSR it is possible to probe diffusing ions, especially using depth profiling to analyse both the bulk and the surface in cells.
μSR is also used in chemistry to extract reaction rate in catalysed reactions by integrating the muon in free radicals, and in biology to study electron transfer properties in biomolecules.

1.3 Negative muons

Negative muon beams are exclusively available at decay muons beamlines, coming from the decay of pions produced in the collision of protons with the graphite target. Their distribution follows the line in Figure 1.8, allowing momentum selection. At decay beamlines positive muons are also available.

Multiple techniques rely on negative muons, such as μ^-SR (muon-spin relaxation spectroscopy), very much similar to the positive counterpart highlighted in Section 1.2.1.
Another completely different technique relies on the properties of negative muons, and their ability to form muonic atoms: it is a relatively novel techique, with many names. It

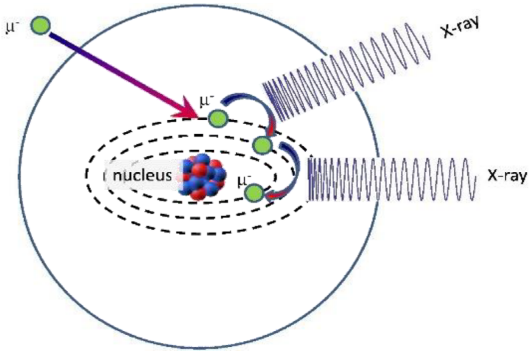


Figure 1.14: Muon capture and cascade [14]

can be referred as Muonic atoms X-ray Emission Spectroscopy (μ-XES), Muon Induced X-ray Emission (MIXE[3]), Muonic Atoms X-Ray Spectroscopy (MAXRS) and possibly more. In this thesis we will refer to it as **μ−XES**.

1.3.1 Muonic atoms

The uses of negative muons are mostly centred around the production and analysis of muonic atoms, which are produced in the interaction between negative muons with local atoms in the target.
The great advantage of the use of muons is that, being massive particles, they implant in the material at a well-defined penetration depth, dependent on momentum and predictable via the use of MC simulations. The muons at the end of range stop, and either decay as free particles via the usual decay scheme (Equation 1.1) or they can be captured by atoms, substituting an electron and forming a *muonic atom*[4]. The process is illustrated in Figure 1.14.

The electron is captured by the nucleus Coulomb potential at a level with high principal quantum number, $n \gtrsim 14$, and it transitions to the innermost shell because of the greater binding energy of the muon to the nucleus compared to the electron [14]. This is due to an increase in the reduced mass of the system, because the muon is ~ 200 times more massive than the electron. The energy levels of atoms are correlated to the reduced mass of the system: in the simplest approximation, the Bohr model, the energy levels are only related to the principal quantum number, without considering angular momentum, and the energy is inversely related to the Bohr radius, defined as $a_0 = \frac{4\pi\varepsilon_0\hbar^2}{Ze^2\mu}$.
Considering the hydrogen atom, we can calculate the difference between the reduced

[3]With the same naming standards as PIXE and PIGE, techniques for the analysis of cultural heritage with particle beams

[4]Muonic elements are designated with a μ in front: element X becomes μX

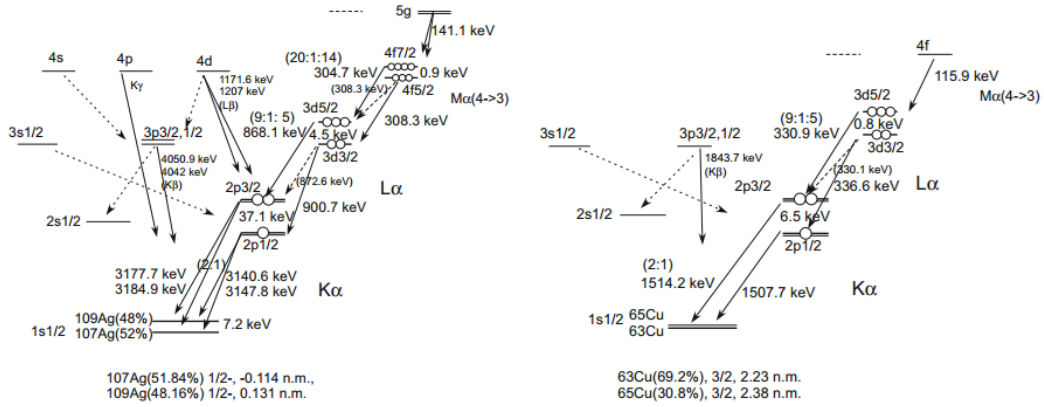


Figure 1.15: Typical muonic atom cascade for copper and silver [3]

masses:

$$\mu_{\mu H} = \frac{m_p \cdot m_\mu}{m_p + m_\mu} = 207\frac{m_p \cdot m_e}{m_p + 207m_e} = 207\frac{m_p \cdot m_e}{m_p + 207m_e}\frac{m_p + m_e}{m_p \cdot m_e}\mu_H = \\ = 207\frac{m_p + m_e}{m_p + 207m_e}\mu_H \approx 186\mu_H \quad (1.16)$$

This change in reduced mass, and thus energy gaps, is an asset in μ-XES because it allows better resolution than classic spectroscopy.

As the system evolves, the muon cascades down towards the innermost shell, first transferring its energy to bound electrons which are then emitted via Auger decay. At lower quantum numbers, $n \sim 5$, the muon loses energy via radiative decay, emitting energetic X and γ rays. Finally the cascade stops when the muon reaches the K shell, 10^{-13} s after capture [15]. The cascade is the source of prompt prompt X-ray emission, caused by the de-excitation of more energetic electrons that can jump levels to fill the lower shells.
The process terminates with either the decay of the bound muon into an electron and neutrinos, or the K-capture of the muon by the nucleus. The latter is more likely the heavier the atom is, and it is not dissimilar to electron capture, where a proton is converted to a neutron with the emission of a neutrino and the absorption of a muon:

$$\mu^- \; p \; \rightarrow \; n \; \nu_\mu \quad (1.17)$$

The capture promts the nucleus to an excited state, which then de-excites with the emission of X and gamma rays and/or neutrons.
A typical cascade can be seen in Figure 1.15.

The increase in likelyhood of muon capture by heavy nuceli is at the basis of another technique for elemental analysis: by measuring the change of the lifetime of muons in

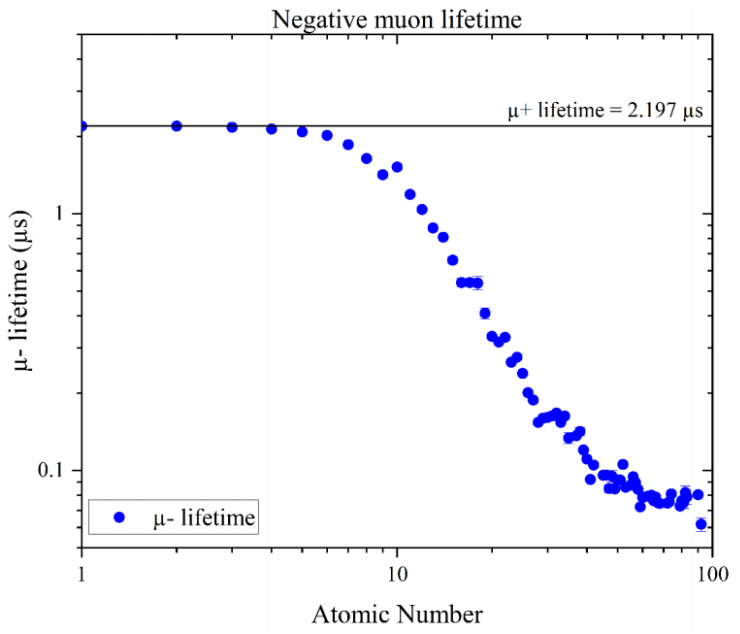


Figure 1.16: Muon lifetime in matter as a function of Z: the higher the atomic number and the more likely K caputure is to happen [12]

matter, as seen in Figure 1.16, we can extract information on the elemental composition of the material in analysis.

Other techniques, such as μ-XES, which we will explore in Section 1.3.2, analyse the characteristic X-rays emitted in the cascade: these photons are much more energetc than the corresponding electronic X-rays, with usual range from 10 keV to 10 MeV, for both light and heavy atoms. This change in energy compared to ordinary matter increases the transmittance of the light through objects, making detection of photons originating from the bulk of the material possible, thus allowing depth probing in non invasive, non destructive way [14].

Muonic atoms are also used in high precision nuclear experiments, decreasing the uncertainty thanks to the higher energies of the emitted X-rays and the bigger influence on the nucleus given the closeness of the muon to the nucleus thanks to the increased reduced mass. This will be explored in Section 1.3.3.

1.3.2 μ-XES

One of the use of negative muons in applied physics is the application of muonic atom X-ray emission for the identification of the elemental composition of an object. This technique has been used in several areas, the main drive coming from heritage science for depth analysis and profiling of artefacts, but it can also be used in battery characterisa-

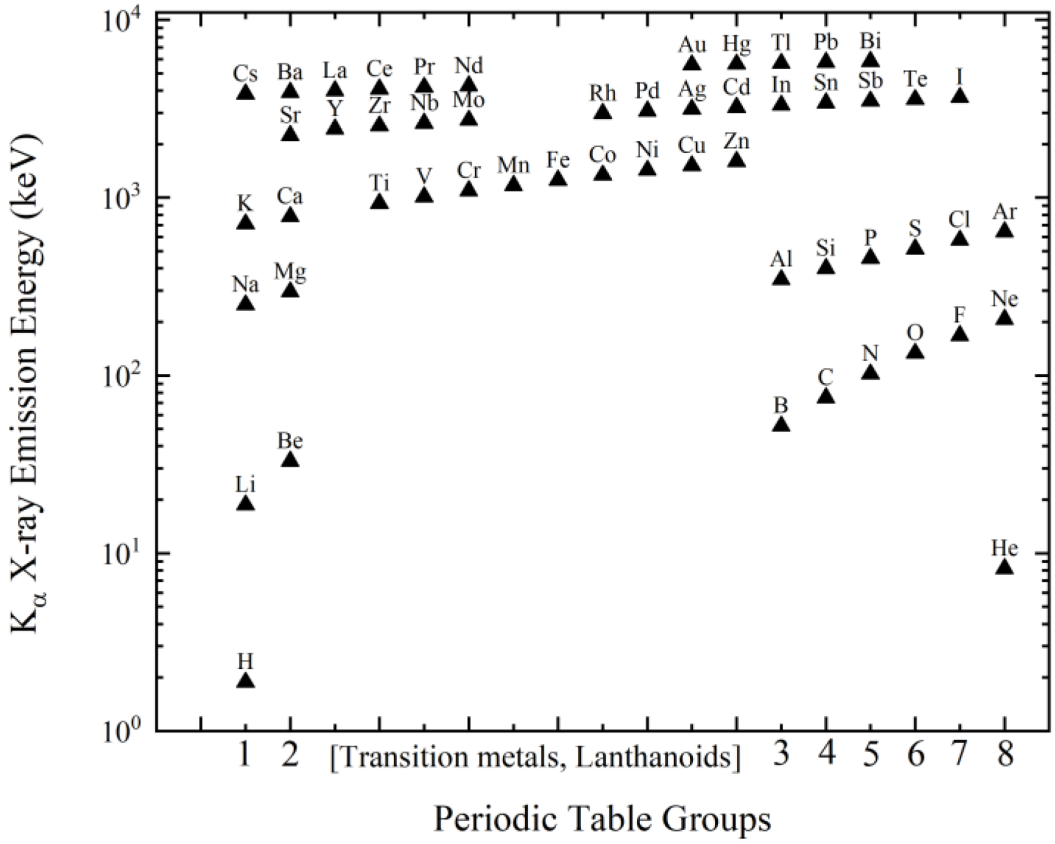


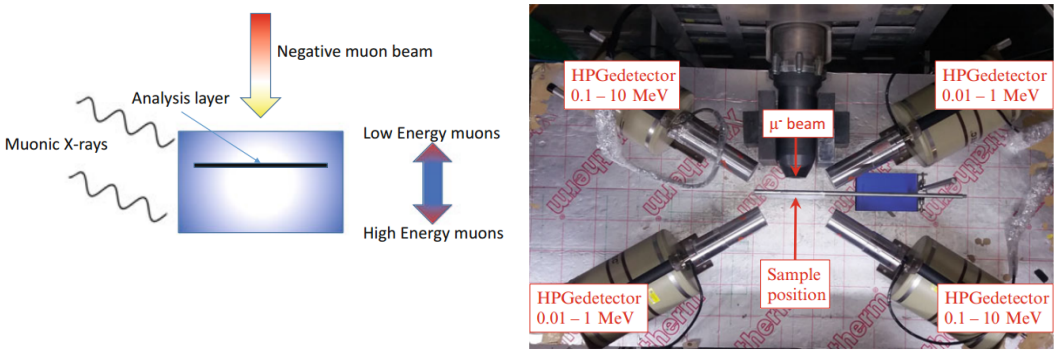
Figure 1.17: K_α X-ray energy for the elements in the periodic table [14]

tion [15].

In the field of heritage science, the technique has many advantages over techniques such as XRF and mass spectrometry used for elemental analysis:

- it is both non-invasive (it doesn't require sample extraction) and non-destructive, allowing the analysis of the bulk of even the most precious artefacts without damage:
- it allows the analysis of both the surface and the bulk of the material with the same apparatus, just changing momentum, reducing handling and the need for other instruments;
- it has a high specificity, allowing the detection of elemental traces.

Compared to PIXE and PIGE, the emitted photons are greater in energy and superior in transmittance, given by the fact that they were generated not in ordinary matter but in muonic atoms, so the corresponding energy levels don't exist in the surrounding material and the probability of absorption is reduced.



(a) Experimental scheme for a μ-XES experiment [3]

(b) Detector setup at RIKEN port 4 [3]

Figure 1.18: Example of experimental setup

As displayed in Figure 1.17, even light elements (excluding helium and hydrogen) have characteristic X-rays above 10 keV, which is often the energy limit for detection with HPGe detectors. This means that the technique is suitable for both organic and inorganic materials, although it is often applied to metals because it is quite efficient in the identification of alloys [3].

A typical μ-XES experiment involves a small sample, placed in the muon beam (Figure 1.18a). A set of photodetectors are placed close to it, usually High Purity Germanium (HPGe) detectors. HPGe detectors are semiconductor diode detectors: the interaction of a photon with the germanium crystal produces electron-hole pairs in the depletion region. The charge carriers are then collected by an electric field and converted to an electric pulse, which is proportional to the number of carriers and the energy of the photon. HPGe detectors, thanks to the increased size of the depletion region and the small energy required to create the e-h pair (~ 3 eV) provide a very high resolution and are capable of detecting X-rays for tens of keV to tens of MeV energy range [5].
Their main drawback is that they are rate-limited due to long detection time and long dead-time. Faster silicon-based detectors and some types of scintillators can also be used: for example, Si(Li) detectors can be useful because they have higher resolution at lower energies, but unfortunately they can only reach ~100 keV, while inorganic scintillators can be made very small[5] and they can be placed closer to the sample, proving very useful despite the limited resolution. Specialised, low energy detectors (10 keV to 1 MeV, either Si(Li) or HPGe optimised for low energy), are sometimes used in conjunction with high energy detectors (100 keV to 10 MeV, usually HPGe). This setup can be seen in Figure

[5]Contrary to germanium photodetectors, which need to be kept cooled with liquid nitrogen and thus require large tanks to keep the temperature stable, scintillators can be made very small, especially if paired with SiPMs

Table 1.2: Range for 30 MeV energy ($\sim$ 85 MeV/c momentum) muons stopping in typical materials in artefacts [16]

Material	Range
Copper (Cu)	$\sim$ 0.79 cm
Silver (Ag)	$\sim$ 0.75 cm
Iron (Fe)	$\sim$ 0.87 cm
Lead (Pb)	$\sim$ 0.85 cm
Carbon (C, compact)	$\sim$ 2.35 cm
Carbon (C, graphite)	$\sim$ 3.13 cm
Water (H_2O, liquid)	$\sim$ 4.76 cm

1.18b, which depicts the typical setup at RIKEN port 4 for a μ-XES experiment
The reason for using a small sample is that muonic X-rays are few and isotropically generated, the detectors must be very close to the origin to obtain a signal over the background, and in a big sample the photons would need to pass through more matter, increasing the likelihood of absorption of the muonic X-rays. The rate must be optimised to reduce counting times without saturating the slower detectors, so it is usually adjustable [3].

Implantation depth can be controlled via momentum selection, with either Monte Carlo simulations or an empiric formula:

$$R = ap^{3.5} \tag{1.18}$$

a is proportional to the inverse of the stopping power, only slightly dependent on element and dependent on density, while p is the momentum.
Depth can be controlled within tens of microns.
The differences in range of typically analysed materials for the same momentum is summarised in Table 1.2.

In order to obtain quantitative measurements of the elemental ratio, we need to consider the likelihood of capture in the different atoms: this is, in first approximation, proportional to the number of electrons.

$$A\left(\frac{Z_1}{Z_2}\right) = \frac{n\ Z_1}{m\ Z_2} \tag{1.19}$$

Where Z is the atomic number and n/m is the ratio between the two elemental abundances.
This is not exact, as the presence of loosely bound electrons can introduce inconsistencies. The experimental formula uses a correction term of the form:

$$\sigma = \left(1 - \frac{E_B}{E_0}\right) \tag{1.20}$$

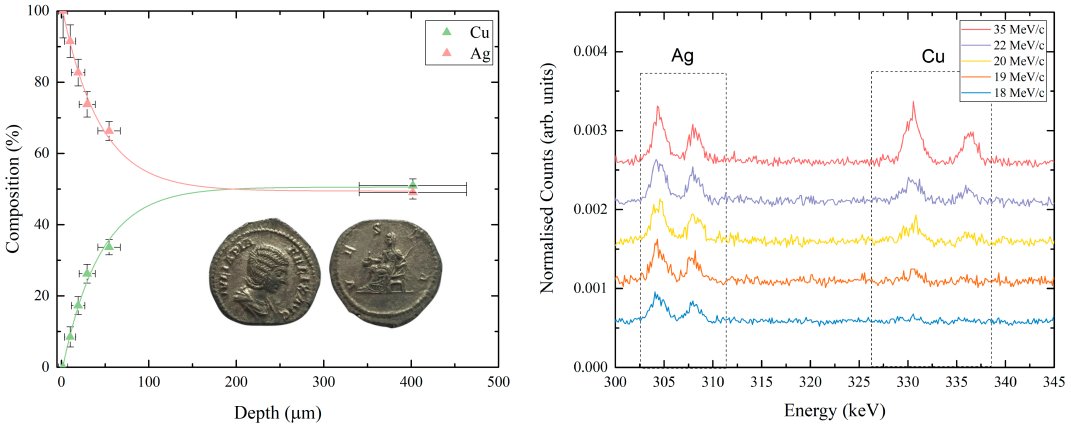


Figure 1.19: μ-XES depth profiling of a Julia Domna Roman silver coin from 211-217 CE [17]

In addition, we have a correction for the valence electrons. In the end we obtain:

$$A\left(\frac{Z_1}{Z_2}\right) = 0.6\rho(1+\alpha\rho)\left(\frac{Z_2}{Z_1}\right)^{\frac{1}{8}}(1+5.53V^{5.45}10^{-5}) \tag{1.21}$$

ρ is the density, V is the valency and α depends on the type of material (oxides, chlorides, ...). To avoid the use of this formula, especially if composition of the sample is not exactly known, it is also possible to measure standards with known, similar composition in addition to the sample, and the probability of capture can be deduced by comparing them [3].

Uses in Cultural Heritage

The applications in cultural heritage are multiple: several examples that demonstrate the perks of the technique will be highlighted [3, 14].

Ancient coins The analysis of ancient coins can shed light on the fiscal health of the nation that produced them: in cases of crisis, the coins are often plated or debased, reducing the value of the metal inside without changing the perceived value. The use non-destructive depth probing allows the detection of such cases without the spoilage of the coin itself: this allows the identification of the alloy ratios at different depths.

A notable example can be observed in Roman coins from the end of the Empire [17]. As seen in Figure 1.19, the coins were found to be made of a silver-copper alloy with silver enrichment of the surface, with only silver peaks visible at the lowest momentum. By looking at the depth profile (a smooth curve, not a step), the debasement process could also be deduced: by using an organic acid on a silver-copper minted coin, the copper on the surface was stripped away, leaving only silver.

Thanks to μ-XES, it was observed that gold coins were much less affected from debasement than the less precious silver [18]. The use of muons on ancient coins can also be for identification of counterfeit, as it was done on Japanese Tempo-Tsuho coins, where to apparently similar coins were found to have different composition in the bulk [19].

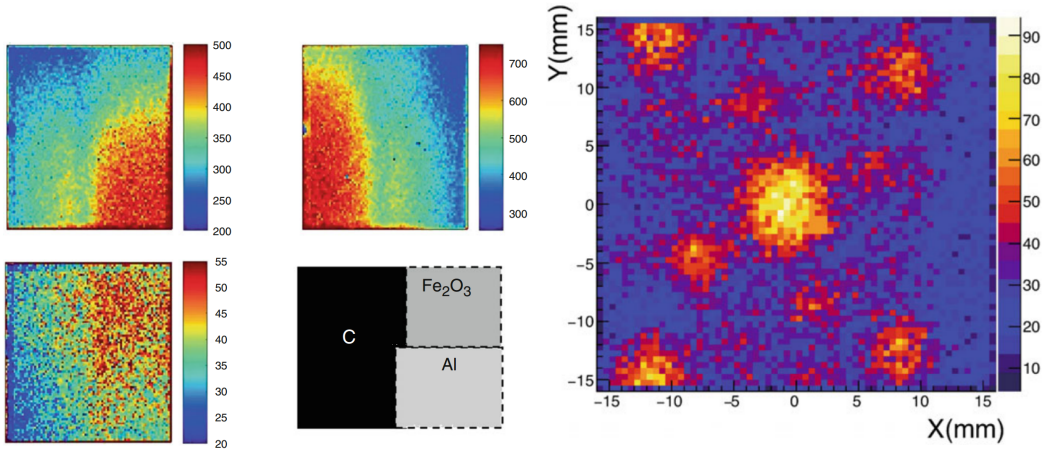
Bronze artefacts The study of bronze artefacts with negative muons can help with identifying means of production, sequence, and the specific alloy recipe looking through corrosion and external layers.
The analysis of a first century BCE bronze mirror from China showed the composition of bronze using a ternary alloy (71.8% copper, 22.1% tin and 6.1% lead) with contaminations [19].
In the analysis of Sardinian Nuragic votive ships, from the CHNet-TANDEM experiment, showed different recipes used in casting different parts of the artefacts, hinting at the adjustment of the ratios (increasing the lead) to improve on molten metal flow while casting the more detailed parts of the ships [20].

Meteorites μ-XES is particulary useful for the analysis of extraterrestrial samples, because of its sensitivity to light elements without being destructive. Moreover, the depth penetration can be used to analyse the fragments without extracting them from the protecting container, in order to avoid contamination with terrestrial material that would hinder future studies.
A notable study is the analysis of the the Murchinson and Allende meteorites: the first was found to contain significant amounts of organic material, carbon, magnesium, silicon, iron, calcium, and sulphur, while the other, much simpler, only contained magnesium, silicon, iron, potassium, and calcium [21].

Organic compounds and biomaterials The technique is suitable to identify not only metals but low Z atoms as well: it was used for the analysis of biological tissues, where it was able to identify components up to 1500 ppm (0.15%) [22].
The use in cultural heritage comes from the analysis of glass-encased biomaterials or organic compounds (such as preserved remains or medicines). The latter was explored with the analysis of the OGATA Koan medicine chests [23], where an unknown medicine was sealed in a small glass bottle. The μ-XES analysis identified both the glass that composed the bottle and the medicine, which was chalomel, composed of mercury and chlorine.

Imaging

Imaging by using negative muons could help identify distribution of materials inside an enclosed object, especially for low Z materials. Feasibility studies have been conducted for both downstream and pinhole detector setup, as seen in Figure 1.20.
The downstream setup consists of using a highly pixelated CdTe detector, placed behind the sample to capture the produced X-ray retaining spatial information.



(a) Downstream setup [3] (b) Pinhole setup [3]

Figure 1.20: Results from the feasibility studies

The pinhole setup is instead based on CdTe strip detectors, with a pinhole collimator placed 90° from the beam direction. It was successful in identifying the positions of lithium fluoride inclusions in the sample.

Isotopic analysis

Given the dependence of the energy levels to the reduced mass of the nucleus-lepton system, the identification of isotopes is theoretically possible: the K_α peak is shifted higher in energy the larger the mass of the isotope nucleus, as seen in Figure 1.21. This shift, only visible at 5 MeV, is minimal and only resolvable with a very high resolution detectors, but it is also possible to quantify the ratio by measuring the characteristic X-rays after K-capture: these, visible in Figure 1.22, are in the hundreds of keV, more easily measured and differentiated even with scintillators with limited resolution.
The isotopic analysis of specific materials, especially lead (Pb), can be used for provenance studies, as lead is the endpoint of most of the natural decay chains of uranium and thorium: small differences in the initial conditions of the cave where lead is extracted are reflected in the isotopic ratios, which are effectively a fingerprint of the geographical origin. The current applications of the technique unfortunately use an invasive approach, not usable on most archaeological findings [3, 14].

The CHNet-MAXI experiment, which will be the focus of Chapter 4, is a feasibility study to understand whether the current setup and technology allows the determination of the isotopic ratios of lead with the required precision ($< 1\%$), via the quantification of the characteristic X-rays.

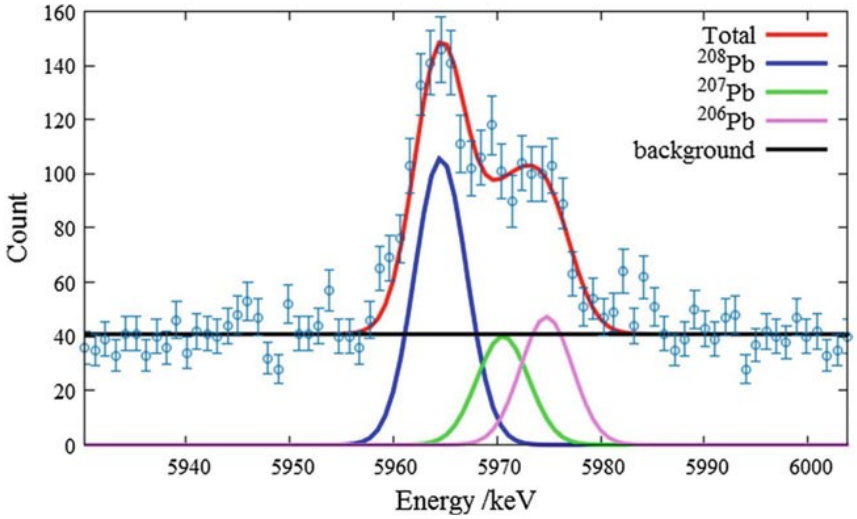


Figure 1.21: K_α line shift for muonic lead, in isotopes 206, 207 and 208 [3]

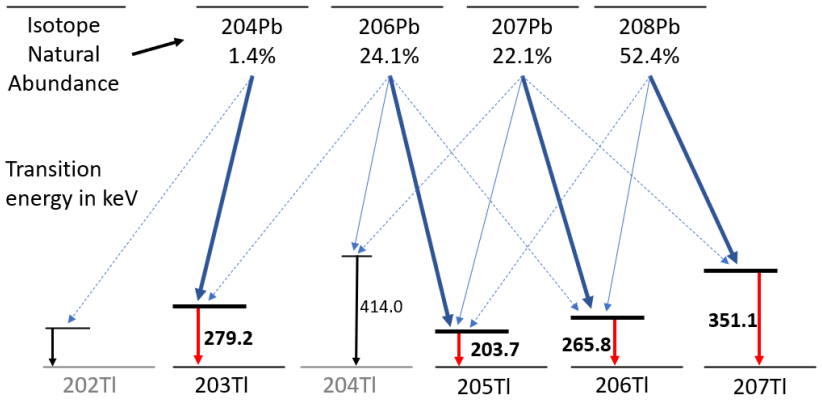


Figure 1.22: Most likely μPb transitions and their characteristic X-rays (courtesy of M. Cataldo)

1.3.3 Muon laser spectroscopy

The increase in resolution is of great use in experiments looking for very small effects with high resolution requirements, such as the ones looking to quantify with extreme precision the hydrogen hyperfine splitting, related to the Zemach radius of the proton.

The effect of the proton charge distribution can be neglected in low energy experiments, but with the increase in energy and probing depth, the experimental results deviate from the point-like proton theoretical cross-section for the elastic collision, even when accounting for the electric and anomalous magnetic moments [24]. To correct for the finite size of the proton, *form factors* are introduced: the first approximation was the

electric form factor, a quantity that is defined as the Fourier's transform of the charge distribution.

$$\left(\frac{d\sigma}{d\Omega}\right)_{\rho} = \frac{Z^2\alpha^2}{4E^2\sin^4(\theta/2)}\cos^2\frac{\theta}{2}|F(\vec{q}^2)|^2 \qquad F(\vec{q}^2) = \int d^3\vec{x}\; e^{-i\vec{q}\cdot\vec{x}}\rho(\vec{x}) \tag{1.22}$$

Considering the magnetic effects in a point-like proton (Z=1), with recoil and spin-spin interactions, we obtain the following cross-section:

$$\left(\frac{d\sigma}{d\Omega}\right)_{Dirac} = \frac{\alpha^2}{4E^2\sin^4(\theta/2)}\frac{E_3}{E_1}\left[\cos^2\frac{\theta}{2} + \frac{Q^2}{2M_p^2}\sin^2\frac{\theta}{2}\right] \tag{1.23}$$

A better approximation accommodates for the magnetic and electric distribution in the proton, using the Dirac and Pauli form factors:

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{Rosenbluth}} = \frac{\alpha^2}{4E^2\sin^4(\theta/2)}\frac{E_3}{E_1}\left[A\cos^2\frac{\theta}{2} + B\frac{Q^2}{2M_p^2}\sin^2\frac{\theta}{2}\right] \tag{1.24}$$

This is the Rosenbluth formula, with the A and B functions defined from the Dirac and Pauli form factors:

$$A(Q^2) \equiv F_1^2(Q^2) + \frac{Q^2}{4M_p^2}F_2^2(Q^2) \qquad B(Q^2) \equiv \left(F_1^2(Q^2) + F_2^2(Q^2)\right)^2 \tag{1.25}$$

The more meaningful magnetic and electric form factors, G_M and G_E, can be used instead:

$$\begin{cases} G_{E(p,n)} \equiv F_{1(p,n)} - \tau F_{2(p,n)} \\ G_{M(p,n)} \equiv F_{1(p,n)} + \tau F_{2(p,n)} \end{cases} \qquad \tau = \frac{Q^2}{4M_p^2} \tag{1.26}$$

Which lead to the Rosenbluth elastic differential cross-section with the Sachs form factors:

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{Rosenbluth}} = \frac{\alpha^2}{4E^2\sin^4(\theta/2)}\frac{E_3}{E_1}\left[\frac{G_E^2 + \tau G_M^2}{1+\tau}\cos^2\frac{\theta}{2} + 2\tau G_M^2\sin^2\frac{\theta}{2}\right] \tag{1.27}$$

These can be extracted from cross-section measurement in Deep Inelastic Scattering, using electron beams on hydrogen:

$$\frac{d\sigma}{d\Omega} \propto G_M(Q^2) \qquad \text{for big } Q^2 \text{ or } \theta \tag{1.28}$$

$$\frac{d\sigma}{d\Omega} \propto G_E^2(Q^2) + \tau G_M^2(Q^2) \qquad \text{for small } Q^2 \text{ or } \theta \tag{1.29}$$

By measuring the cross-section at small and big angles G_E and G_M can be extracted, and from that, the charge and magnetic radii:

$$r_E^2 := -6\frac{dG_E}{dQ^2}\bigg|_{Q^2\to 0} \qquad r_M^2 := -6\frac{dG_M}{dQ^2}\bigg|_{Q^2\to 0} \tag{1.30}$$

These results for the radii correspond to the square root of the second moment of either the electric charge density or the magnetic dipole density respectively. The k^{th} moment is defined as:

$$\langle r_X^k \rangle := 4\pi \int_0^{+\infty} dr\ r^{k+2} \rho_X(r) \tag{1.31}$$

In the calculation for the hyperfine splitting of hydrogen, the contributions at first order of the magnetic radius cancel out: the first non-zero contribution is a function of the Zemach radius.

$$r_Z := -\frac{4}{\pi} \int_0^{+\infty} \frac{dQ}{Q^2} \left[\frac{G_E(Q^2) G_M(Q^2)}{1+\kappa_N} - 1 \right] = \int r\ d^3r \int d^3r'\ \rho_E(\vec{r} - \vec{r'}) \rho_M(\vec{r}) \tag{1.32}$$

κ_N is the anomalous magnetic moment of the proton. The precise evaluation of the hyperfine splitting would thus allow the extraction of the proton Zemach radius with a novel approach [12].

Compared to measurements taken with electron-proton scattering and hydrogen spectroscopy, muonic laser spectroscopy is expected to obtain an uncertainty one order of magntitude smaller, thanks to the increased energy of the characteristic X-rays of muonic atoms and the increased interaction between the muon and the nucleus spins due to the reduced radius: the hydrogen hyperfine transition line is at 21 cm, while the same line in muonic hydrogen is at 7 μm. The current state of the measurement is summarised in Figure 1.23.

There are three experiments that are currently measuring the hyperfine splitting using muonic hydrogen: at PSI (Switzerland, CREMA), at J-PARC (Japan), and at ISIS (UK, FAMU), with different techniques.

FAMU (Fisica degli Atomi MUonici, Muonic Atom Physics) is an INFN (Istituto Nazionale di Fisica Nucleare) experiment for the measurement of the Zemach radius by means of measuring the muonic hydrogen hyperfine splitting. It will be in operation at RIKEN-RAL port 1 until the end of 2026, and is currently analysing data from the 2023-2025 data taking campaign.
It aims to achieve the resolution goal of $< 1\%$ on the Zemach radius (10^{-5} on the hyperfine splitting), one order of magnitude better than current measurements.

It is a muonic laser spectroscopy experiment, the general scheme pictured in Figure 1.24. In such experiment, muonic atoms of the species in analysis are formed in an enclosed target by stopping a negative muon beam inside. A moment later, a laser at the selected wavelength is injected in the cavity, exciting the $1S - hfs$ (hyperfine splitting of the $1S$ level) and inducing spin-flip. The resonance is correlated to an excess in the experiment observable, which in FAMU is delayed μO X-rays: the muon transfer rate

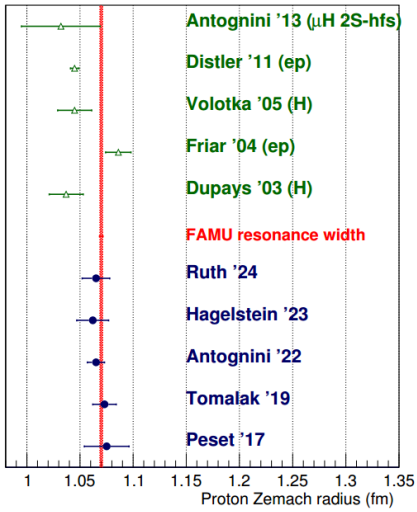


Figure 1.23: Experimental results (green) from ep scattering and hydrogen spectroscopy compared to the theoretical predictions (blue): in red, the expected resonance width (80 pm in hyperfine splitting, 0.003 fm in radius) [12]

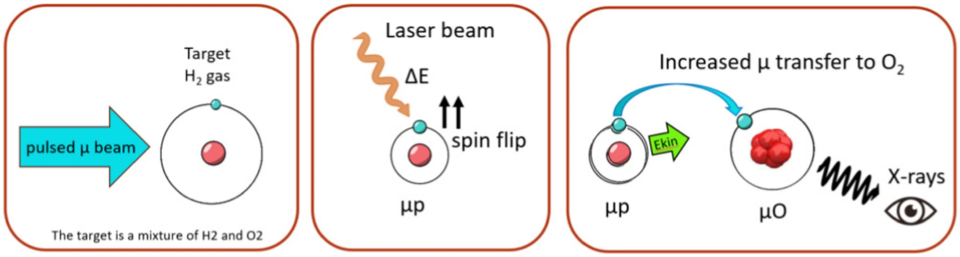


Figure 1.24: General scheme for muonic laser spectroscopy [25]

from hydrogen to oxygen is increased in the case of resonance, increasing the amount of muonic oxygen present in the target. A certain amount is present from the incident beam, producing *prompt* X-rays, but, as after laser injection the transfer is not instantaneous: the transfer increases the amount of *delayed* X-rays, after the injection of the laser.

Theoretical predictions for the hyperfine splitting define the explored spectral range, 2σ from the weighted average as in Figure 1.25, from 6788.040 nm to 6789.140 nm. To achieve the expected resolution, FAMU takes data every 25 pm, in order to differentiate between true resonance and fluctuations, because the expected resonance width is $\sim$ 80 pm: we would expect an increase of the X-rays in three or more adjacent measurements.

The expected delayed X-rays are $K_\alpha \simeq 133$ keV, $K_\beta \simeq 158$ keV, and $K_\gamma \simeq 167$ keV detected, via an array of detectors:

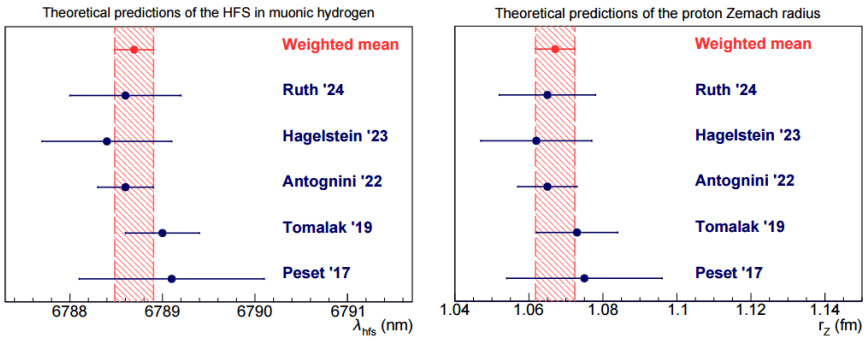


Figure 1.25: Theoretical predictions for the Zemach radius and the hyperfine splitting [12]

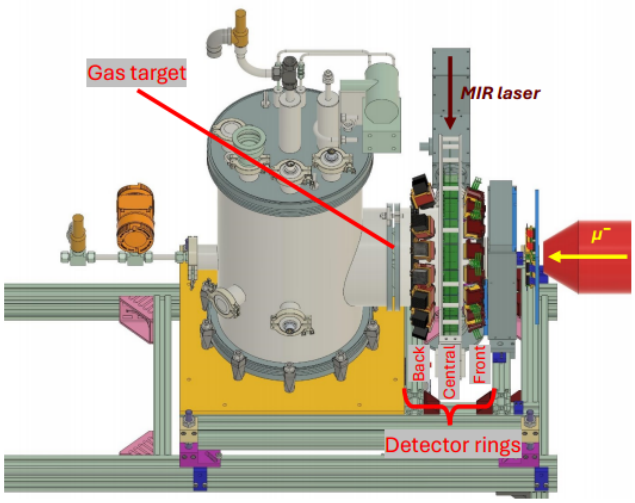


Figure 1.26: FAMU detector setup: from the right the muon beam is injected into the gas target, around it the three rings [12]

- $LaBr_3$(Ce) scintillators, in sizes:
  - 6 × 1" diameter cylinders, 1" thick, read by PMTs;
  - 28 × 1" diameter cylinders, 0.5" thick, read by SiPMs ;
- one HPGe detector (with length 20.8 mm, diameter 30.4 mm and a 0.9 mm carbon entrance window) placed downstream.

The detectors are placed in three rings, upstream, central, and downstream of the target, as visible in Figure 1.26. The HPGe detector was not depicted.

In front of the beam nose, a beam monitor is placed to measure the muon flux and normalise batches by the number of muons hitting the target. The design of the beam monitor will be discussed in detail in Section 4.2.

Chapter 2

Scintillation detectors

The use of scintillating detectors is ubiquitous in particle physics: they are varied as their applications, and can be used in all areas of physics and all types of measurements. Trackers, counters, and calorimeters can all be made with scintillators, their characteristics fine tuned to serve the intended purpose.

The basic function of a particle detector is to detect the presence of a particle, gathering additional information to serve the experiment where it is employed. Energy, position, flux, direction, nature of the particles, and more are all deduced by interactions with the detectors.
The detection process is possible by leveraging the interaction of a particle in matter, which releases energy in the active medium of the detector.
The energy of massive charged particles is released following the Bethe-Bloch (Equation 1.2) energy loss function, while photons can undergo either (coherent or incoherent) scattering, pair production, or photoelectric effect, which then transfer energy to a charged particle that can be easily measured. Even neutrons and neutrinos can be detected by using the interaction with charged particles.
The detection strategies are differentiated by the *method* of energy measurement: scintillators convert the energy into light. Different basic detectors can then be employed in geometrically significant ways in order to obtain additional information.

A scintillation detector is composed of two main parts: a scintillating medium, which produces light as a consequence of the release of energy in the material, and a photodetector, which quantifies the light and converts the information into an electric signal.

2.1 Scintillators

Scintillation is one of the oldest techniques for particle detection, but modern improvements keep it evergreen.
An ideal scintillator is a uniform material that has specific characteristics [5]:

- *scintillation efficiency*: the amount of detectable light produced from the deposit of energy is high and in a optimal wavelength for detection;
- *light yield*: the amount of light produced is linearly proportional to the deposited energy (photons/MeV);
- *transparency*: it should have high transmittance for the emission wavelength;
- *time response*: the signal produced should be fast, to accommodate high particle rate;
- *radiation hardness*: capability of maintaining original characteristics after prolonged use in highly radioactive conditions;
- *quenching factor*: fraction of energy converted into non-radiative processes, must be minimised.

Other characteristics, such as cost, ease of coupling to photodetectors, and ease of manufacturing should also be taken into consideration when choosing the detector for a specific application.
The compromise between the different characteristics is inevitable, as no single material can be optimal in all. The two main groups are organic and inorganic scintillators, the main difference being the light yield and time response.

The light emission mechanism is that of fluorescence: the light is emitted after excitation of some kind, immediately after the passage of the particle. Several other phenomena are also possible, such as delayed fluorescence and phosphorescence, but with much longer emission time. These longer lived emissions are undesirable, and in an ideal scintillator should be minimised.

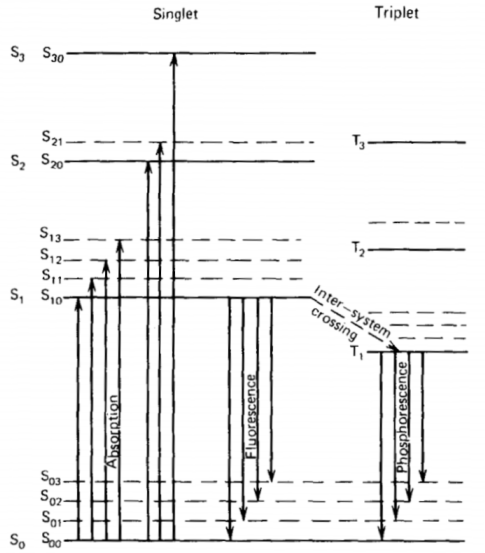
2.1.1 Organic scintillators

Organic scintillators are characterised by the presence of organic molecules (carbon, oxygen, hydrogen and other low Z elements). They can be liquid or solid, in this case they are also referred to as *plastic* scintillators, and they are made with the the main component of a solvent or polymer, which acts as absorber and transmits energy to the scintillating additives dissolved into the medium, for the conversion of energy into light.

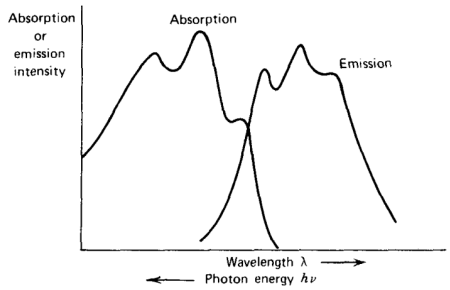
Mechanism of scinitllation

For light production, the mechanism that is leveraged is the excitation of singlet electronic states in molecules with a π-electron structure, often arising in organic molecules with particular symmetric properties [5].

As visible in Figure 2.1a, the electronic level structure of these molecules has singlet and triplet states, with circa 3-4 eV of separation between levels, each then further separated into vibrational states (15 meV separation between these levels).
The energy deposited by the charged particles is absorbed by the electrons, which, at normal temperature, mostly occupy the ground state. The electrons promoted at higher



(a) Electronic states. Energy absorption is indicated with an up arrow, emission via a down arrow [5]



(b) Absorption and emission spectra, highlighting the Stokes shift [5]

Figure 2.1: Characteristics of light emission from organic scintillators

levels, in picoseconds, are de-excited to the S_1 level via internal conversion without emission of photons. This produces an excess of electrons in the S_1 excited state, which then transition back to the ground state or one of its vibrational states, emitting photons. This is the *prompt component* of the scintillation light, with a typical time profile that is:

$$I = I_0 e^{-t/\tau} \qquad \tau \sim \text{nanoseconds} \tag{2.1}$$

Via intersystem crossing, some of the excited electrons can also decay into the T_1 excited triplet state, which has a much longer lifetime (up to milliseconds): this produces the *delayed component*, containing both phosphorescence (direct decay from T_1 to S_0, with a longer wavelength) and delayed fluorescence (thermal excitation from T_1 to S_1, then decay into the ground state, with the same wavelength as the prompt component).

The fact that most of the light emission decays to a vibrational state of the ground state means that the light has lower energy than the one needed to excite electrons from S_0 to S_1, thus increasing transparency. This is the Stokes shift, responsible for the transmittance of the scintillating light. In Figure 2.1b it is possible to observe the shift in the absorption and emission spectra.

Characteristics

Organic scintillators have good transparency and excellent response time, although phosphorescence and delayed fluorescence can be limiting factors.

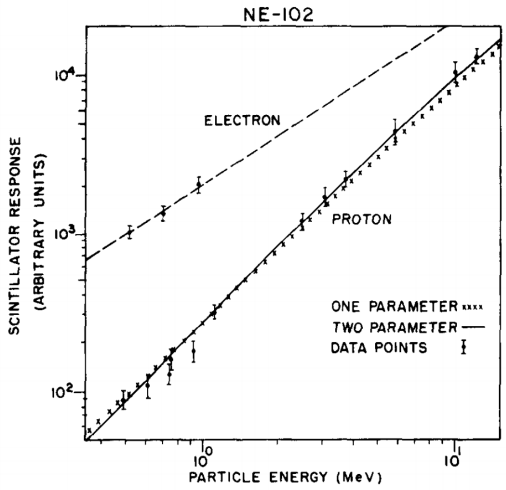


Figure 2.2: Light yield for a common organic scintillator [5]

The scintillation efficiency is limited by quenching, caused by impurities: these degrade the light output by providing additional non-radiative decay modes, such as dispersion of energy via lattice vibrations and heat.
Organic scintillators are also minimally dense, which makes them good choices for trackers and monitors: the particles can pass through the detector undisturbed, keeping their trajectory and most of their energy. They are also often used for detection in beta spectroscopy and fast neutron experiments, the latter because of the large percentage of hydrogen which is a good neutron absorber [5].

For most of the scintillators available, the light yield is linear for energies above 125 keV for electrons and at much higher energies for heavier particles, as observable in the plot in Figure 2.2.
Heavy particles also have a lower light yield: this is due to saturation along the track, and is empirically modelled by Birks' Law:

$$\frac{dL}{dx} = \frac{S\frac{dE}{dX}}{1 + kB\frac{dE}{dx}} \tag{2.2}$$

The fraction $kB\frac{dE}{dx}$ represents the amount of *spent* molecules along the track, which cannot produce scintillation light. For electrons, saturation has virtually no effect, while for alpha particles the energy loss per unit path is so large that the saturation term dominates:

$$\begin{cases} \frac{dL}{dx}\Big|_e = S\frac{dE}{dX} \\ \frac{dL}{dx}\Big|_\alpha = \frac{S}{kB} \end{cases} \implies kB = \frac{\frac{dL}{dE}\Big|_e}{\frac{dL}{dx}\Big|_\alpha} \tag{2.3}$$

The alpha-to-beta ratio is often used to compare scintillators, and it is always less than 1 in organic scintillators because the energy loss for alpha particles is always larger [5].

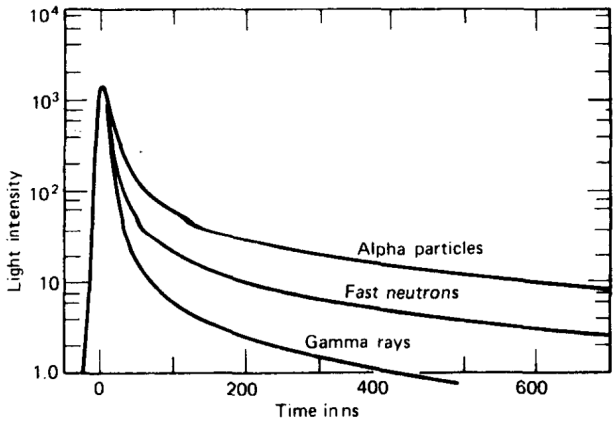


Figure 2.3: Time response comparison for neutrons and alpha and gamma particles [5]

The light emission of a scintillator can be modelled as a sum of two emissions, the prompt and the delayed components. While most of the light is produced in the initial fluorescence, the second is not insignificant and, most importantly, is different for different particles.

$$I(t) = Ae^{-t/\tau_p} + Be^{-t/\tau_d} \tag{2.4}$$

B is a factor that depends on the particle, larger for alpha and smallest for gammas, as visible in Figure 2.3. This effect, caused by the excitation of triplet states along the track of highly ionising particles, enables the use of Pulse Shape Discrimination to differentiate the particles passing through the scintillator. PSD is often used to remove undesired background [5].

2.1.2 Inorganic scintillators

The category of inorganic scintillators contains mostly crystalline materials, typically without carbon, oxygen and hydrogen.
The most diffused are alkali halides crystals with small amounts of activator impurities: most commonly used is caesium iodide with thallium impurities (CsI(Tl)), but also with other activators and other alkali (potassium, and lithium). Outside alkali, options are abundant, but less common and all with very specific use-cases (GAGG, CLLB, YAG, BGO, ...) [26].

Mechanism of scintillation

Scintillating crystals are insulators, meaning that the electronic energy levels, depicted in Figure 2.4, are described as bands with a large energy gap:

- valence band: the "ground state", representing electrons bound to the lattice;

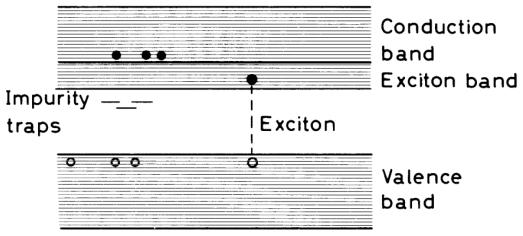


Figure 2.4: Energy band structure of an inorganic crystal [26]

- conduction band: represents the free electrons in a crystal;
- forbidden band: between the valence and conduction band, the electrons **cannot** stay in these states.

The energy gap between the valence band and the conduction band, or the minimum energy to be absorbed by the electron for it to be promoted to the conduction band, for insulators is too high to utilise the unaltered crystal for scintillation: in this case, the energy required to promote the electron to the conduction band would be too high, making the process unlikely. The decay of the electron directly to the valence band would also produce photons that are too energetic, reducing the efficiency of photodetectors.
To make the use of crystals viable, there needs to be a modification of the band structure: this is done via activators. Deliberately added impurities add energy states in the forbidden gap, called *luminescence or recombination centres*, reducing the band gap and shifting the scintillation photons to the visible range [5].

The passage of an ionising particle in the crystal releases energy, promoting electrons from the valence band to the conduction band and creating electron-hole pairs: the hole quickly travels to the position of an activator ground site, while the electrons, free to travel through the crystal, will do so until reaching an impurity, with an activator state. The electron will then decay into the ground activator state, emitting light.
The lifetime of these states is the limiting factor to time response, which is usually in the order of 100 ns or more.
Competing, longer lived processes are also possible, such as the excitation of the electron to an excited activation state where the direct decay to the ground state is forbidden. These are much longer lived, and produce phosphorescence. Quenching is also possible. Another option is the formation of an exciton, where electron and hole remain loosely associated, but free to drift. The decay is similar to free electrons and holes, with recombination via activation sites [5, 26].

The crystals, thanks to the small amount of activators, are transparent to scintillation light. The emission spectrum, in Figure 2.5, depends on both the crystal and activator, giving way to a wide range of possible spectra. It is also possible to have multiple decay modes, resulting in bimodal or multimodal spectra. The light yield is linear with respect to the energy deposit, with minimal non-linearity introduced by quenching (al-

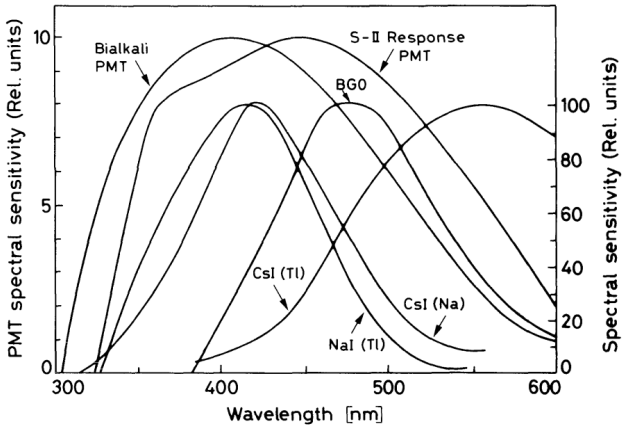


Figure 2.5: Spectra for common inorganic scintillators and the response curves for two PMTs [26]

beit in much lower impact than in organic scintillators). A similar effect as Birk's Law is observed, where the alpha-to-beta ratio is almost always less than 1, but much closer to 1 than in organic scintillation [5].

Characteristics

The main use is in detection of gamma rays and high energy electrons/positrons, due to their high density and thus high stopping power. The light yield is also much higher than organic scintillators ($\sim 10^4$ photons/MeV), increasing the resolution of the detector.
The main drawbacks of inorganic scintillators are that they are 2-3 orders of magnitude slower than organic scintillators, mostly due to the high amount of phosphorescence. Moreover, most scintillating crystals are hygroscopic, meaning that extra precautions with air humidity (or encapsulation) is needed [26].

2.2 Photodetectors

The light produced in scintillators must be collected and measured. This is often done with photo-multipliers, which are devices that collect light and produce a current dependent on the intensity of the input.

The choice of photodetector depends on both scintillation wavelength and experimental requirements, such as size, magnetic fields, longevity, rate...
The two most commonly used devices are silicon photo-multipliers and photo-multiplier tubes. Both are able to *count* single photons, measuring even the faintest light and producing a measurable current pulse.

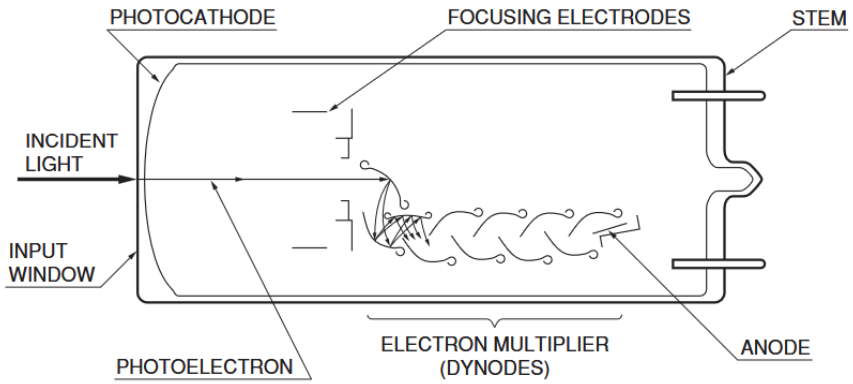


Figure 2.6: Typical structure for a classic PMT [27]

2.2.1 Photo-Multiplier Tubes

One of the first photodetectors ever used, and still one of the most widely used, Photo-Multiplier Tubes (PMTs) are simple and reliable. The uses are varied, same as their size, as PMTs can go from a few centimetres to half a metre in diameter.
Their function is rather simple: a photon interacts with the photocathode, extracting an electron, which is then accelerated and multiplied in its path to the anode. The structure, in the simplest case, is depicted in Figure 2.6: it is composed by a vacuum tube with a glass endcap on one side (the photocathode), several electrodes, called *dynodes*, and a final collection anode. The dynodes compose the accelerating and multiplying structure of the PMT [5].

The photocathode is an absorber coated with photoemittive material. A photon deposits $\sim$ 3 eV of energy into an electron in the absorber, which then migrates to the surface of the photoemitter. If after migration the electron is able to overcome the potential, it escapes: the photocathode is chosen with a potential barrier of 1.5-2 eV, usually made of semiconductors.
The properties of the photocathode define the absorption characteristics of the PMT, defining the wavelength cutoff and the efficiency: the first depends on the potential barrier, the second on the thickness of the glass [5].
The photodetection efficiency is described by the *quantum efficiency*, defined as:

$$QE(\lambda) = \frac{\#\text{ of photoelectron emitted}}{\#\text{ of incident photons}} \tag{2.5}$$

The quantum efficiency for a typical photocathode is around 20-30%, greatly dependent on wavelength [5].

The energy resolution of a PMT is instead linked with the electron multiplication: the process happens between the dynodes, where a positive high voltage ($\sim$ 100 V) is ap-

plied. An electron exiting the photocathode will be directed to the dynode, increasing it energy from a few electronvolts to several hundreds. In the dynode, the incident electron can excite several other electrons from the valence band: if all energy is deposited, and dynodes have a band gap of ~ 3 eV, then ~ 30 electrons are excited to the conduction band, a few of them will be freed from the material ($\delta \sim 3-5$). The multiplication is repeated for each dynode, exponentially increasing the number of electrons:

$$G = \alpha\delta^{\#\text{ of dynodes}} \tag{2.6}$$

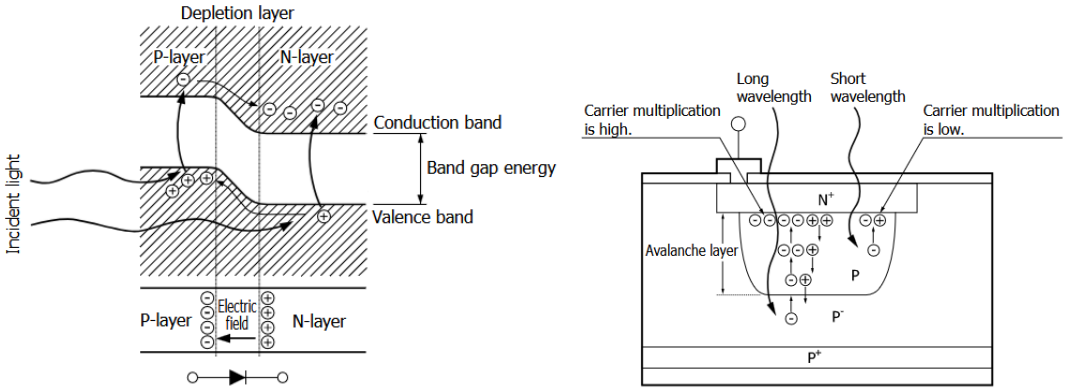
α describes the fraction of electrons surviving the initial collection stage. For typical PMTs, the gain can easily reach values of $10^6 - 10^7$ [5].
The energy resolution depends on both δ, α, and QE.

Time resolution instead depends mostly on the structure of the photocathode, where electrons produced far apart can have much different collection time, worsening the spread in transit time. It also minimally depend on the electron transit time in the dynodes ($\sim 20-80$ ns).
The limiting factor on resolution, both time and energy, is the number of initial electrons, directly related to the light yield of the scintillator [26].

The simple structure is able to maintain linearity over a large difference in number of photoelectrons, from single to several thousands, with non-linearity arising with very large pulses [5].

Any electron or light-emitting process in a PMT can cause noise. The most important source is thermoionic emissions from the cathode, proportional to the temperature and to the area of the detector, which generates the *dark current*. Natural radioactivity present in the tube can also produce spurious electrons that can be accelerated.
Afterpulsing can happen for multiple reasons, the most prominent being light emission from later stages of multiplication: if the photon reaches the photocathode, a photoelectron may be emitted. The resulting afterpulse is correlated to the first but much smaller. Another source of afterpulsing is the presence of gas infiltrations in the PMT vacuum: these atoms are ionised by the passage of electrons through the gas, resulting in positive ions. These ions are then accelerated towards the photocathode, releasing thousands of photoelectrons upon impact. These afterpulses are much larger than the photon-initiated pulses, sometimes even larger than the initial pulse, and usually happen later than with photon afterpulses, due to the slower acceleration of ions (depending on the size of the PMT, the delay can vary from hundreds of nanoseconds to microseconds) [5].

One of the drawbacks of PMTs is the operation voltage, which needs to be of several hundred volts per dynode and can easily reach 1-2 kV, which needs to be kept very stable as δ greatly depends on the inter-dynode voltage. Of note is the fact that PMTs, to be used in environment with high, variable, magnetic fields need to be either shielded or produced with continuous dynodes, because of the distortion to the electron path caused



(a) Electronic band structure for a photodiode [29] (b) Avalanche formation in a SPAD [29]

Figure 2.7: Photodiode detection principle

by the Lorentz force [26].

2.2.2 Silicon Photo-Multipliers

A silicon photomultiplier (SiPM) is a device composed of simple photodiodes operated above the breakdown voltage, also called a Single-Photon Avalanche Diode (SPAD), arranged in an array and readout in parallel.
Silicon photodiodes have been introduced to improve on the limitations of PMTs: they require lower voltage and are insensitive to magnetic fields. They can also be made much smaller and have a higher quantum efficiency [5, 28].

In semiconductor materials (silicon, germanium...), incident light can generate an electron-hole pair. A photodiode is a p-n junction (in Figure 2.7) where the generation of an electron-hole pair in the depletion region can be detected via amplification. Depending on the reverse voltage applied to the diode, the electric field in the depletion region changes, introducing different multiplication behaviour for the free charges, as see in Figure 2.8 [30]:

- $V < V_{ADP}$ or photodiode mode: no amplification;
- $V_{ADP} < V < V_{BD}$ or avalanche photodiode mode: amplification proportional to the number of electron-hole pairs, thanks to the avalanche sustained by electrons only. The avalanche is self-quenched;
- $V > V_{BD}$ or Geiger mode: the amplification is no longer proportional to the input, the avalanche involves electrons and holes, triggering a self-sustaining avalanche process that must be quenched by external circuitry.

The avalanche happens when electrons (or holes) gain enough energy to generate another pair. In the Geiger mode (or SPAD mode), the photodiode essentially becomes a photon counter, irrespective of initial energy.

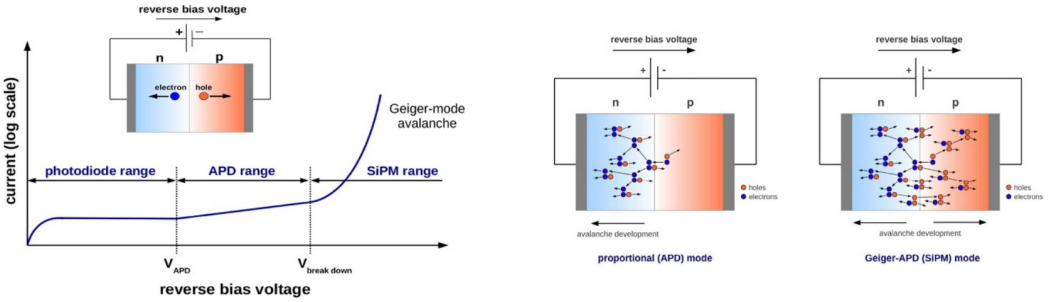


Figure 2.8: Photodiode working modes. On the right the different voltage regions, on the left the comparison between avalanches in APD and SPAD [30]

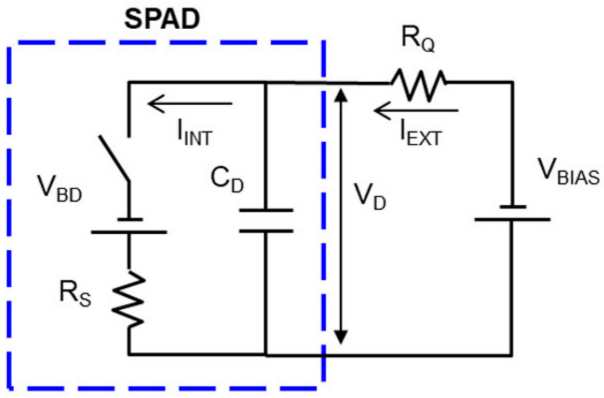


Figure 2.9: SPAD equivalent circuit [28]

The equivalent electrical circuit of a SPAD (Figure 2.9) can illustrate how to derive the important characteristics of detectors, such as gain and timing.
In ready condition, the diode capacitance (C_D) is charged and no current is flowing. The electron-hole pair generation is akin to closing the switch, which in turn allows the discharge of the capacitance from V_{BIAS} to V_{BD}. The internal current flows through the resistance R_S, while the external current flows through R_Q: both tend to the steady state value of $(V_{BIAS} - V_{BD})/R_Q$ with $C_D R_S$ time constant. The avalanche quenches thanks to the quenching resistance: the switch opens, the internal current goes to 0 and the external current decreases exponentially, recharging the capacitance.
The important parameters in a SPAD are linked to the behaviour of the switch, which closes with a probability P_T, defining the quantum efficiency, and opens with a probability P_Q, which defines the time duration of the signal [28].

The gain is defined as:

$$G = \frac{\text{avalanche charge}}{\text{electron charge}} = \frac{(V_{BIAS} - V_{BD}) \cdot C_D}{q} = \frac{V_{OV} \cdot C_D}{q} = \tag{2.7}$$

The gain is linear with respect to the overvoltage, defined as the difference between operating voltage and breakdown voltage. As breakdown voltage depends on temperature, the gain does as well [30].
The breakdown voltage can be extracted in two ways:

- I-V curve: the point of maximum derivative (dI/dV) is the breakdown voltage;
- G-V line: the first photoelectron peak position represents the gain, it can be measured across a wide range of voltage, and then extracting the breakdown voltage from the intercept with the $G = 0$ axis.

SiPMs are arrays of SPADs connected in parallel, their current summed to obtain the final signal. The dimension, density and fill factor (active area/total area) that define the array are important to the definition of the Photon Detection Efficiency of the photodetector [30]:

$$PDE(V_{OV}, \lambda) = QE(\lambda) \cdot P_T(V_{OV}, \lambda) \cdot FF \tag{2.8}$$

Three sources of noise are most prevalent:

- dark count rate: random firing in no-light condition due to thermal excitation;
- optical crosstalk: caused by photon emission from adjacent cells, correlated in time to the initial pulse;
- afterpulsing: similar to crosstalk but from photons originating in the same SPAD. Due to the fact that afterpulsing happens before the full recovery of the cell, the signal will be smaller.

The last two are correlated to the initial pulse: the first is indistinguishable from real signals, while the second will always be smaller. The recovery time of the SPAD depends on the capacitance C_D, $\tau = C_D R_Q$, meaning that excessive noise could cause too much deadtime for the detector to function [30].

One of the main drawbacks of SiPMs is saturation and limited linearity. The factors that are most important for the definition of the linear range are the number of cells and the amount of noise (ENF, excessive noise factor).

$$N_{fired} = N_{cells}\left[1 - e^{-\frac{N_\gamma \cdot PDE \cdot ENF}{N_{cells}}}\right] \tag{2.9}$$

This is caused by the finite number of cells and the probability of two or more electrons hitting the same cell: the SPAD response depends only on activation, not the number of incident photons [30].

Time resolution depends on fluctuations in the avalanche formation process, but it is usually around 20-40 ps FWHM. The energy resolution is instead dependent on gain fluctuations, influenced by temperature and noise [28].

2.3 Detectors

From the building blocks of scintillators and photodetectors, many different detectors can be built: from thin trackers to calorimeters, the properties of the final device will depend on the nature of the scintillators, the effectiveness of the coupling, and the geometry of the device.

2.3.1 Geometry

The geometry of scintillators greatly affect both the light output and their uses.
Light produced in scintillators undergoes two attenuation processes [26]:

1. bulk attenuation: $L(x) = L_0 e^{-x/\ell}$, where ℓ it the attenuation length;
2. boundary attenuation: each interaction with a boundary causes a part of the light to be reflected and part to be transmitted.

ℓ is often large, so only very big detectors are affected by bulk attenuation, but transmission to the photodetector and internal reflection are both cases where the light is attenuated. To maximise transmission to the photodetectors, see the next Subsection for a more in-depth discussion.

The maximisation of internal reflection is instead very important for scintillating fibres [5, 26], which are often used to construct position sensitive detectors. In this case, the most important factors are:

- refractive index: the total internal reflection can only happen for angles smaller than the critical angle for TIR, dependent on the ratio of the internal and external refractive index ($\theta_c = \arcsin(n_{out}/n_{scint})$). The most important light transport mode in fibre scintillators is in fact total internal reflection, which allows the transmission over long distances and even over minimum bends without loss of signal;
- coating: to improve the reflection off of the walls, a reflective coating can be used (often aluminium or titanium dioxide). Sometimes a protective cladding is added to reduce impurities infiltrations and optimise the critical angle;
- shape: minimisation of the travelled path, by either changing the shape of the scintillator or using multiple collection points, can reduce the amount of wall reflections, improving the collection.

From thin fibres, position of the particle can be extracted by arranging fibres in orthogonal planes, the position inferred from the coincidence of signals in the different planes. Plastic scintillators are often used, to minimise disturbance to the particle, while the strict timing requirements and the small size often require SiPMs as photodetectors.

2.3.2 Coupling

Coupling of a photodetector with a scintillator ensures the maximum efficiency, by both matching the wavelength of emission and absorption and maximising transmittance.
The first challenge is the selection of the refractive index of the scintillators, which should be the closest possible to the photodetector, in order to avoid reflections. The second is

the use of optical grease to bridge the air gap between the components, reducing the amount of interfaces with different refractive index (for glass, plastic and liquid scintillators, and grease, $n \simeq 1.5$, while it can be higher for crystals).

Coupling doesn't only refer to the physical connection, but to the matching of the emission and sensitivity windows: the peak of emission should be well inside the sensitive window of the photodetector. If the wavelength is not optimal for detection, the solution is the use some strategies for shifting the wavelength.

For organic scintillator, a tertiary component can be added to the initial mixtures, which absorbs the primary scintillation light and re-emits it with longer wavelength. The addition is either in solution or a piece of extraneous material which can be embedded in the plastic or be placed alongside. The embedded wavelength shifting fibres can also act as light guides (similar to optical fibres, with higher refractive index than the scintillator), collecting the scintillation light and funnelling it to the photodetector.
In plastic scintillators the most used flour is the one converting blue to yellow-green wavelengths, which are more easily detected by both PMTs and SiPMs. The quantum efficiency for these type of additives varies from the typical 70-80% up to some cases where it reaches 90%, with a time response that is slightly slower than the primary flour [5].

Chapter 3

Towards the next generation of μSR spectrometers

The EC facility at the ISIS Neutron and Muon Source hosts three muon spectrometers, specialised in different applications of the muon spin relaxation and rotation technique.

The facility was commissioned in the early 1980s after the development of the μSR technique. The initial facility only contained a single instrument, what is now MuSR, but was upgraded after increase in beamtime demand with the installation of the kicker and two new experiments, EMU and DEVA, specialised in RF experiments. DEVA was later decommissioned, leaving space for the new, high field spectrometer: HiFi. The upgrades are still ongoing: MuSR is going to be replaced by Super-MuSR, with improvements in performance and rate capability, and the entire facility is in the process of integration into the novel data acquisition architecture, Digital Muons [31, 32]. The Digital Muons project will be discussed in detail in Section 3.3, while Super-MuSR will be the focus of Section 3.2.

3.1 The current instruments

The three instruments currently operating on the decay muon beamline are, in order of commissioning, MuSR, EMU and HiFi. They all follow the basic design of a muon spectrometer for μSR measurements, but with slight differences to cater to differnt experimental needs.

The basic muon spectrometer consists of two detector arrays, usually symmetric rings, positioned upstream and downstream from the sample position, with an electromagnet to generate the magnetic field.

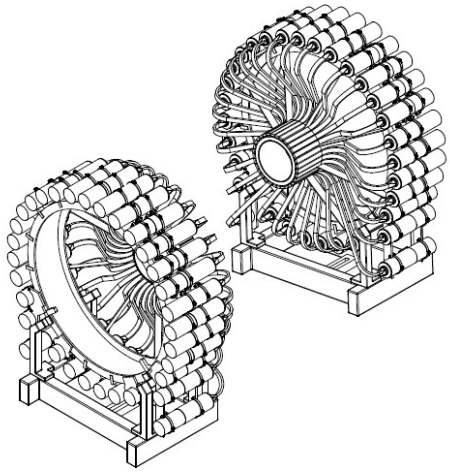


Figure 3.1: The MuSR spectrometer [33]

3.1.1 MuSR

The first of the spectrometers that was commissioned is MuSR: it was first constructed in 1987, with 32 scintillators (16 per side), improved in 2004 doubling this number. The spectrometer is the only rotating one in the facility, and it allows all three modes of measurements for μSR [33].
In particular, thanks to the rotation, it allows transverse field measurements with a magnetic field of up to ~ 80 mT, higher than any other facility [34].

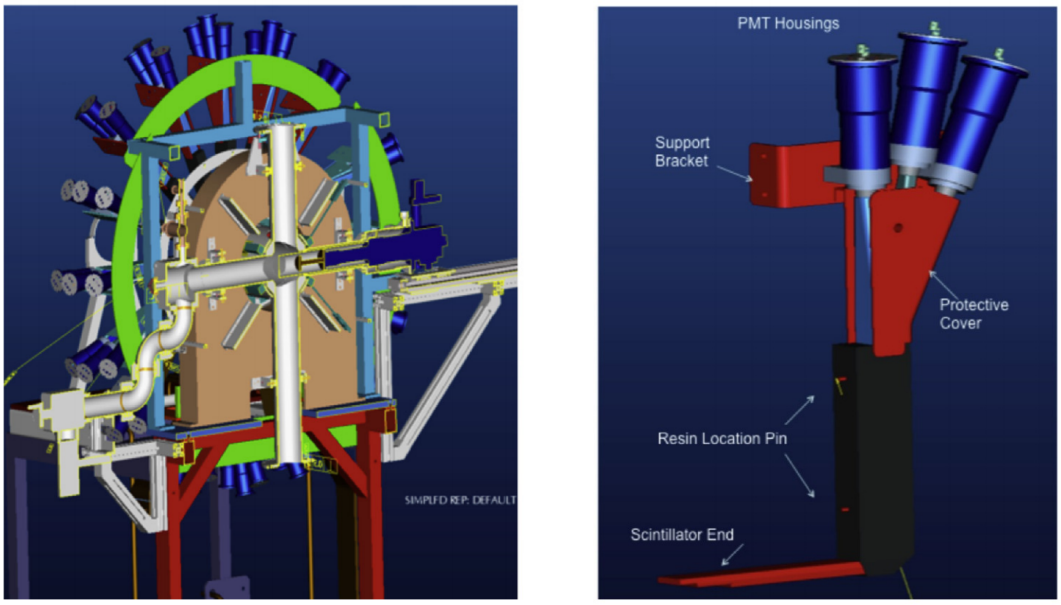
The current detectors, visible in Figure 3.1 are cylindrical scintillators, set in a ring formation. They are connected to Hamamatsu 2-inch R1166 PMTs via perspex light-guides. A set of cylindrical brass and aluminium rings is used as a degrader, to remove low-energy positron signals.

The MuSR instrument is planned for decommissioning during ISIS long shutdown in 2027, when it will be substituted by Super-MuSR(see Section 3.2).

3.1.2 EMU

EMU was the second spectrometer installed at the facility, commissioned in 1993 following the high demand of the user community. The instrument was subject to a major upgrade in 2014, with new detector segmentation and data acquisition electronics, to make full use of the increase in flux of the beam [35].

The number of detectors was chosen to be able to handle up to 500 events per 32 μs acquisition frame. The muon spectrometer, in Figure 3.2a, consists of 96 scintillating counters, in two rings. Each ring has three steps, each with 16 scintillators (see Figure 3.2b). The scintillators are connected to Hamamatsu R5505 1-inch PMTs via light guides,



(a) Cross section of the EMU instrument. In the centre, the sample holder [35]

(b) Single ring section, the scintillator is composed of three segments, each connected to a PMT [35]

Figure 3.2: EMU instrument scheme and detail of the detectors

mounted perpendicular to the field to minimise distortion to the path of the electron: the PMTs are unshielded, subject to 0.05 T stray fields.
The electromagnet system is able to provide up to 0.5 T longitudinal field and 15 mT transverse field in the sample position. Three orthogonal pairs of coils enable the Zero-Field condition, compensating for the Earth's own magnetic field (0.2 mT). The sample is placed in a vacuum chamber, where it can be measured at cryogenic and high temperatures, from 0.03 K to 1400 K.

The data acquisition is handled using a leading edge discriminator, with a threshold of 75 mV which is easily able to distinguish positron signals from noise. The signal is then passed through a Time-To-Digital converter to timestamp the arrival of the positron in the frame, excluding signals outside a 32 μs window.

Thanks to the 2014 upgrade to the TDC, the instrument is able to perform radio frequency μSR: the method removes the frequency measurements limitation given by the pulsed source, making it comparable to a continuous source. The resonant radio pulse is applied after the accumulation of the entire bunch ($\sim$ 80 ns, during which the spin is kept from precessing by using a static field), rotating the spin by 90° and then allowed to precess and decay.

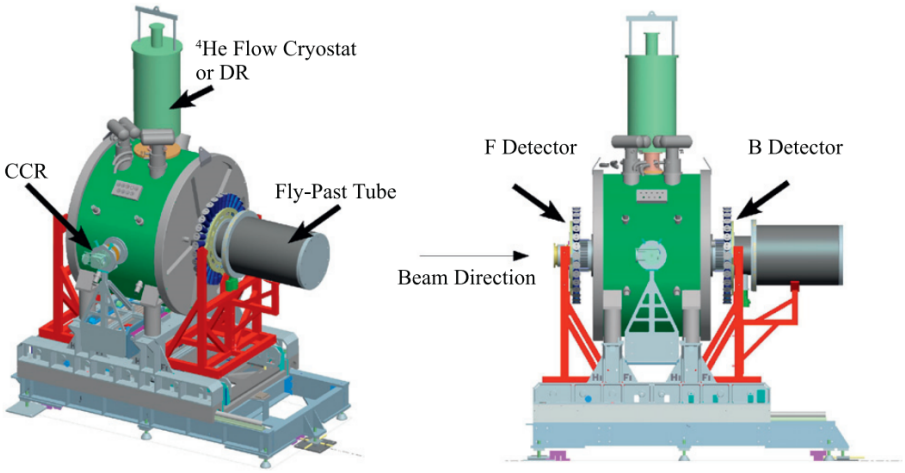


Figure 3.3: The HiFi spectrometer [36]

3.1.3 HiFi

The youngest instrument, HiFi, substituted DEVA in 2013, with an upgrade in 2017 for the addition of a laser. The main difference between the other two and HiFi is the magnetic field, which can reach up to 5 T. Such high fields allow the study of wider range of spectroscopy for condensed matter and molecular dynamics, and RF-μSR [36].

The requirements for HiFi were for the spectrometer to be comparable to measurements at continuous sources, with high rate and optimised for LF-μSR.
The most important piece is the superconducting magnet, with no rotation. The beam enters longitudinal to the field, interacts with the samples in the vacuum sample holder, equipped with cryostat and furnace to perform temperature studies (30 mK to 1500 K using four different devices), similarly to EMU but with much higher fields. The magnet provides a very stable and homogeneous field over a large area, and it is self-shielded with extra coils to avoid stray fields that could affect the rest of the facility[1]. The sample is inserted perpendicularly to the beam: this design, in Figure 3.3, allows the use of a fly-past tube to collect the muons passing through and around the sample, to reduce background by removing them from the sensitive area of the detector.
The magnetic field is so large that it affects the beam, and additional steering magnets were installed to reduce the beam size for the efficient analysis of small samples.
The positron detectors are placed inside the magnet, arranged in two rings: each ring is made of plastic scintillators, segmented into 32 sectors, as in Figure 3.4. They are arranged in two rings with a single step, contrary to EMU, as the high field removes the need for concentric segmentation: the step closest to the beam would receive all signals. The scintillators are coupled to R5505 PMTs via light guides, which cannot be completely

[1]The field is 0.2 T at 3 m from the magnet

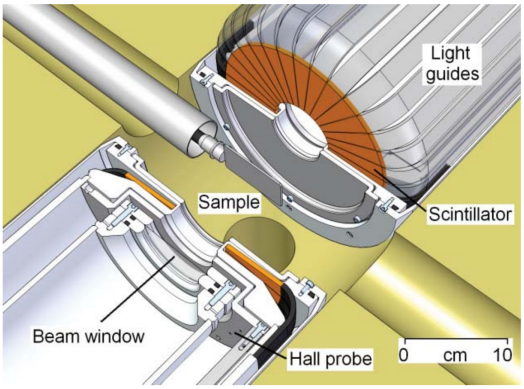


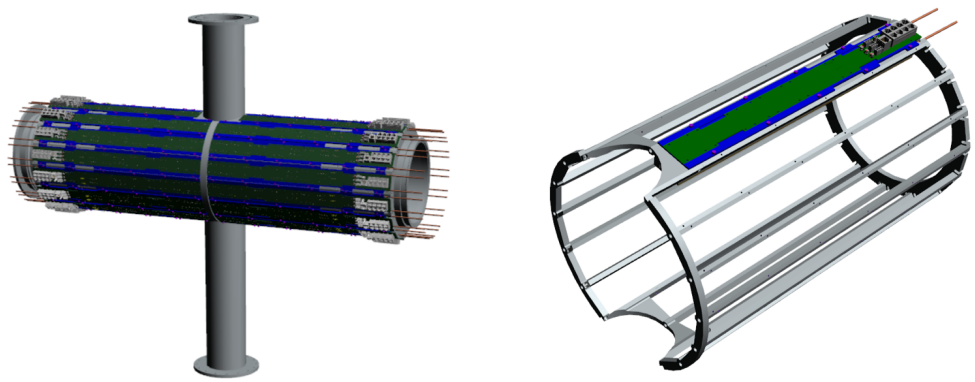
Figure 3.4: The HiFi scintillating detectors [37]

shielded via mu-metal but can tolerate a stray field of 0.15 T if placed radially.
Similarly to EMU and MuSR, the data acquisition is handled with a discriminator and a TDC [37].

The spectrometer was later equipped with a high power pulsed laser system, with the goal of performing high field photoexcitation μSR studies, very useful to condensed matter research and semiconductors. The laser system is constructed in order to produce short, high-powered pulses in a wide range of wavelengths: a Nd:YAG and an optical parametric oscillator were chosen for the task, jointly providing the required characteristics [38].

3.2 Super-MuSR

The Super-MuSR spectrometer will substitute the MuSR instrument, with planned installation during ISIS long shutdown in 2027. The main difference between the older and the new spectrometer will be the use of SiPMs instead of PMTs: this will allow an increase in the magnetic field available, as SiPMs are insensitive to magnetic fields, and an increase in time resolution and rate capability, thanks to the new digital data pipeline. The upgrade also involves the use of a beam chopper to reduce the time width of the beam pulse from 70 ns to 40 ns.
The current instrument covers a $\sim$ 40% solid angle and utilises brass degraders to remove the low energy positron: only 20% of the positrons are actually counted, and only 10% of the beam flux can be used: great improvement in counting rate can be achieved with an instrument upgrade to allow the use of the beam at full intensity despite a small change in the solid angle ($\sim$ 42%). The new instrument must be able to work both in transverse and in longitudinal field, with high performance in both to be able to inherit the current workload of MuSR [34].



(a) Detectors and sample holder, to be placed in the bore of the magnet

(b) Schematic of one side of the barrel

Figure 3.5: Schematic of the Super-MuSR detector array (courtesy of RAL-Technology Project Engineering Group 2023)

The improved rate capability (10^9 counts·hr^{-1}) and time resolution will allow better in-operando measurements and extend the range of materials that can be analysed. The magnet will stay the same (a Helmholtz coil set) but the use of SiPMs will enable the use of fields from 0 to 0.650 T.
To achieve these goals the choice of geometry is different from the currently used ones: instead of small scintillators close to the sample, the Super-MuSR array will be composed of 960 scintillators, allowing for an increased rate capability. The geometry will be barrel-like, with two identical cylinders (forward and backward detector), 57 cm in length and 35 cm in diameter, with 480 scintillators each. A schematic of the barrel and the detector layout is in Figure 3.5.
Each cylinder will be segmented in 16 sections, each equipped with a stave, running lengthwise along the beam path. The stave (Figure 3.6) contains four scintillating modules, each composed of 8 scintillating tiles arranged in 4 columns and two rows. Each tile is read by two SiPMs, coupled to a WLS fibre running along and inside the plastic scintillating counter. The barrel is arranged in 8 segmented scintillating rings of different widths [39].

3.2.1 Modules and stave

Each ring width is optimised to cover the same solid angle, in order to all receive a comparable amount of positrons. A schematic of the ring sizes can be referenced in Figure 3.7.

Each stave, an independent sector of the barrel, contains 32 scintillating tiles, in 4 columns lengthwise. Each column is split lengthwise in the 8 rings, and 60 adjacent columns create the sides of the detector.

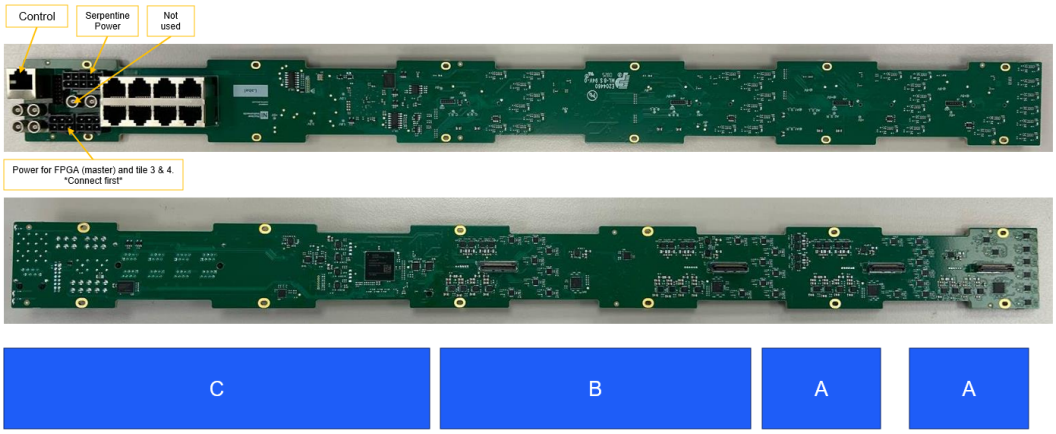


Figure 3.6: The Super-MuSR stave. In order from the top: the outside, the inside, the place where modules attach and the module's layout (courtesy of D. Pooley)

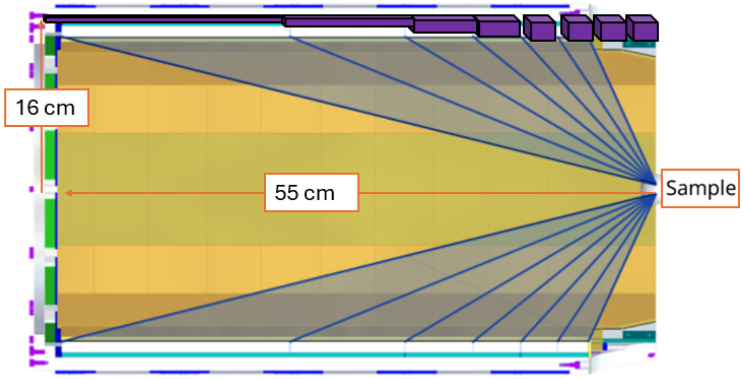
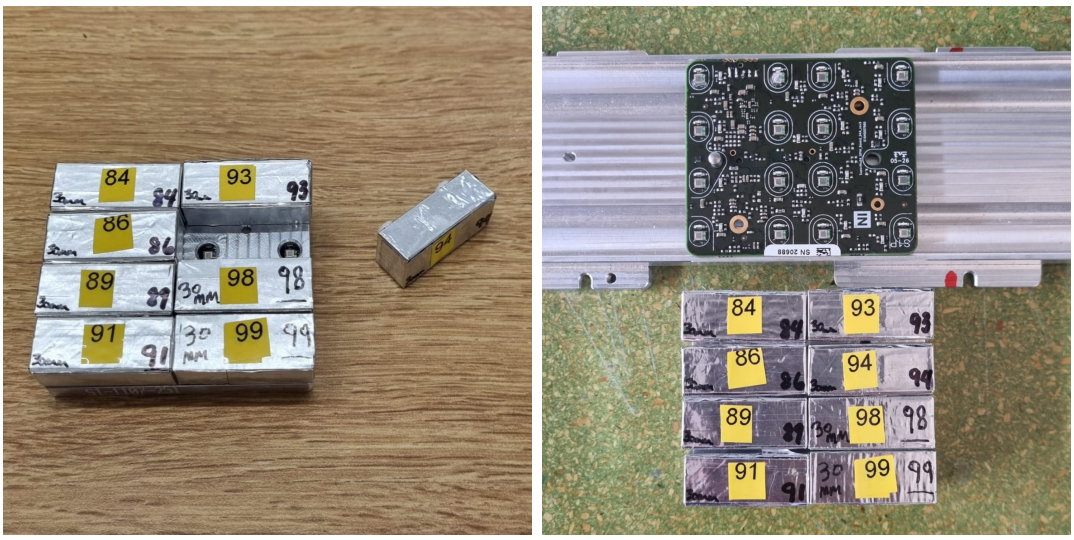


Figure 3.7: Schematic of one of the barrels showing the difference in length of the tiles. The length is chosen to obtain a similar solid angle for all scintillators (courtesy of D. Pooley)

A module is set of 8 scintillators, grouping together two rings and spanning 4 columns. The modules are differentiated by the size of the tiles that compose them:

- Module A: composed of 8 identical tiles, 30 mm long. It is repeated twice and composes the first 4 rings of the barrel;
- Module B: composed of 4 tiles, 43 mm long and 4 tiles, 63 mm long, composing ring 5 and 6;
- Module C: composed of 4 tiles, 105 mm long and 4 tiles, 210 mm long, composing ring 7 and 8.

Each module is composed of the 8 tiles mounted over a machined aluminium cooling



(a) Layout of the scintillating tiles. The holes correspond to the position of the studs and SiPMs

(b) PCB for the module, with the SiPM readout electronics

Figure 3.8: Module A: tiles and SiPMs (courtesy of B. Thompson)



Figure 3.9: The DAQ121 module [40]

plate (Figure 3.8a). The PCB with the SiPMs (Figure 3.8b) is mounted on the underside, with holes where the studs of the tiles couple to the SiPMs via optical grease. The PCB module then connects to the stave via a pin connector.
The stave is a custom-built circuit with integrated HV and cooling. It provides power to the SiPMs that read the scintillating tiles, their signal then amplified and directed to the DAQ121 module. The DAQ121 (Figure 3.9) is both the controller and the digitizer: it can digitise up to 32 channels per module, and one is required for each stave. The data acquisition software is custom-built by Nuclear Instruments and allows data acquisition, analysis, calibration, and control of the operational parameters.

Tiles

The tiles are made of plastic scintillators with an embedded WLS fibre, which shifts the wavelength of the scintillation light to be more closely matched with the SiPM sensitive range. The tiles are produced and then tested in batches using a PMT rig, to ensure

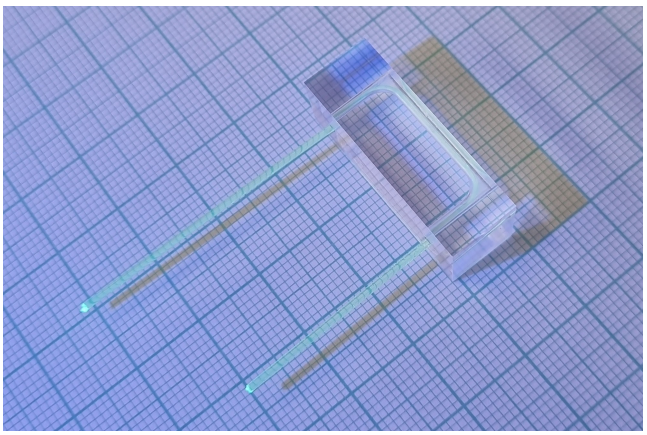


Figure 3.10: A 30 mm scintillator tile with the WLS fibre (courtesy of B. Thompson)

uniformity and high light yield: the testing is done before assembly on the SiPMs.

The size of the tiles depends on which module they are going to be mounted on: the shorter tiles are thicker than the longer ones, steadily diminishing in thickness the further from the sample. This is done to even out the path length of the positrons in the counters, ensuring a similar response for all the scintillators. In the closest tiles, the positrons originating from the sample will have tracks that are almost perpendicular to the face of the tile, while the further out a tile is and the shallower the angle between the track and the surface becomes, meaning a longer path in the tile if they were made all with the same thickness. The number of photons reaching the SiPM per positron track is circa 30.

Each tile has two studs at each end, providing a mounting point for the SiPMs (Figure 3.10). Because of the simultaneous readout of the tiles at the ends, the light collection increases, improving the resolution of the fibre. Moreover, because of the reduced path that the scintillation light takes, the timing resolution is also improved. The main drawback is the need to gain-match the two SiPMs, a topic that will be discussed later in the chapter.

SiPMs

The SiPMs are Hamamatsu S13360-1325PE, 1.3 mm in size and with a 25 μm pitch [41]. The breakdown voltage is around 50 V.
Compared to the previous PMTs, it improves the timing resolution and it is insensitive to magnetic field, allowing an increase in the available field. Their small size means that no light guides are needed, and the Super-MuSR barrel can be inserted in the bore of the magnet, much closer to the sample. This reduces timing uncertainty introduced by the transport in the light guides.

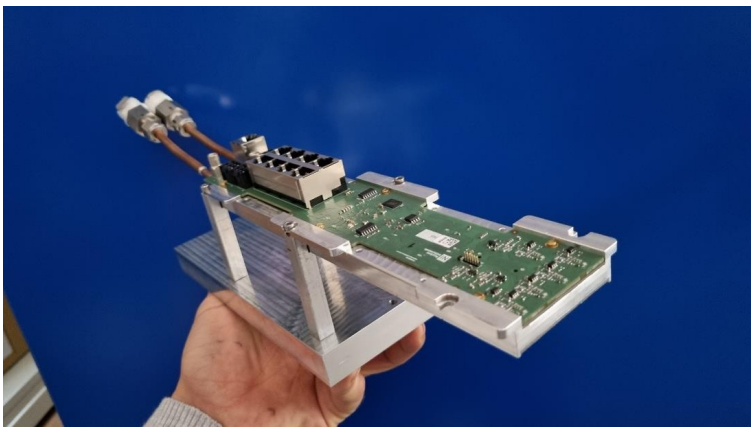


Figure 3.11: Photo of the single-module protostave (courtesy of D. Pooley)

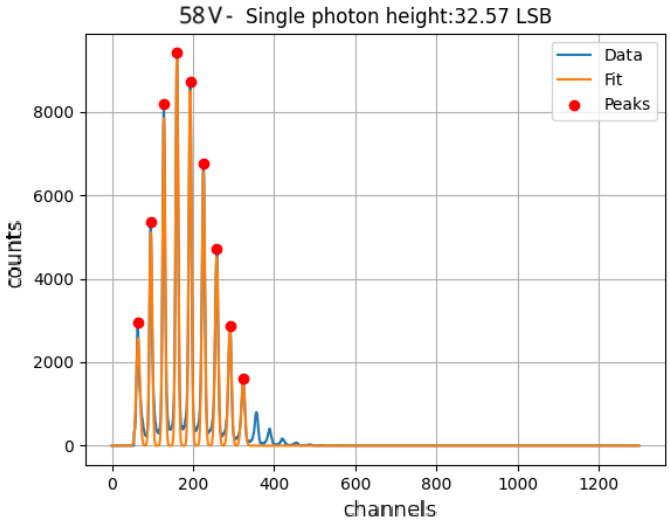


Figure 3.12: PHS for one of the module B SiPMs at 58 V

3.2.2 Testing and optimisation

During the ERASMUS+ Traineeship I had the opportunity of working on the first steps in optimisation of the detector.

My main contribution came from the testing and analysis of the SiPMs and the optimisation of the gain for counting purposes.
The tests were performed by using the *protostave* (Figure 3.11), a simplified, single-module version of the full stave, connected to a DAQ121 with a firmware for the extraction of the pulse height spectrum (PHS) and the calibration.

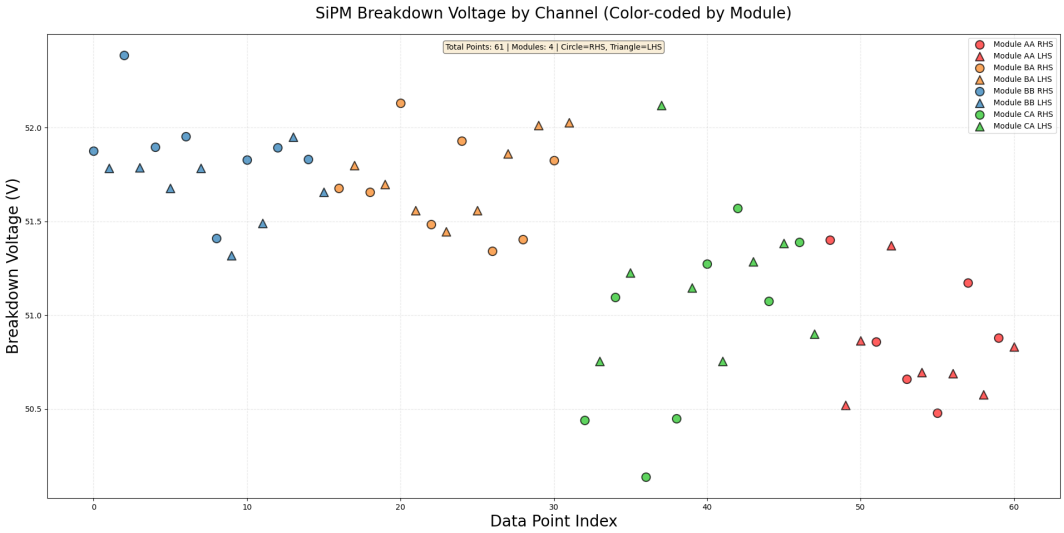


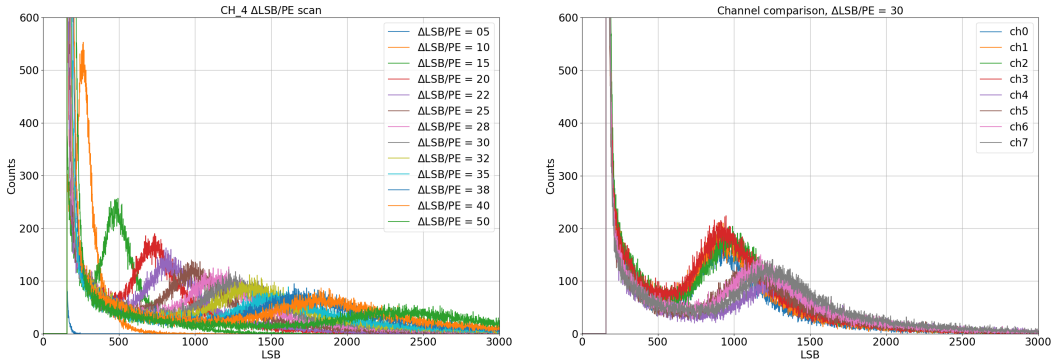
Figure 3.13: Summary of the breakdown voltage for the modules available (courtesy of B. Thompson)

Before any tests on beam, the tiles and SiPMs needed to be characterised: by using the control capability of the DAQ121, which not only provides the acquisition but also the power to the SiPMs, we acquired the PHS of the SiPMs under pulsed light conditions, for each module as a function of voltage applied. This was done by leveraging the use of the integrated LED pulser system, which will be used in the full detector for online calibration.
The PHS for this type of measurement showcases a typical peak structure, each peak corresponding to multiples of a single firing cell: the distance between the peaks is proportional to the gain of the SiPM, and by extrapolating the gain to 0 we obtain the breakdown voltage for the SiPMs. One of the spectra acquired can be seen in Figure 3.12, while the results of the breakdown voltage calculation are highlighted in Figure 3.13. The extrapolation was done via a python script, first the extraction of the peak positions (`scipy.signal.find_peaks`) and then the fit of the straight line through the points using the function `scipy.optimise.curve_fit` for the fit of the straight line.

The same fit line was also used for the matching of the SiPMs: given that two SiPMs need to measure the same tile, matching their gain is essential to them working as if they were one single detector. We define the gain in units of LSB[2] per photoelectron (that is the distance between two peaks in the pulsed laser PHS), denoted as PE.

The rest of testing was performed on beam, parasitically to other user measurements. The main task was the optimisation of the gain for the discrimination of the positrons

[2]Least significant bit, equivalent to ADC channel



(a) PHS as a function of PE (b) Difference in the tile response for PE 30

Figure 3.14: Module B test on beam

from the background. Several gain scans were performed over the usable range, from 5 PE to 50 PE. The results, one of which is in Figure 3.14a, show the position of the peak, but also the spread, increasing as the voltage increases. By looking at the difference in module B channels (channels 1-4 are connected to the shorter, 43 mm tiles, while 5-8 are connected to the 63 mm tiles), we can notice that the intrinsic properties of light transport through a longer fibre have a big effect on the spread of the peak (Figure 3.14b).

To quantify the quality of a PHS for counting experiments, we can look at some quantities, defined in Figure 3.15, that can help to define a Figure of Merit (FoM):

- peak height: the height of the main peak, measured in counts;
- valley height: minimum of the counts between the shoulder and the peak, measured in counts;
- peak-to-valley ratio (ptv): peak height divided by valley height;
- right hand half width half maximum (rHWHM): corresponds to the width of the main peak, measured in LSB;
- peak-to-noise separation (width): separation between the peak and the noise shoulder, measured in LSB at the height of the peak;

Several figures of merit can be defined. In general, the idea is to define a quantity that is maximised when the PHS exhibits the sought-after characteristics: maximisation of peak-to-valley ratio, minimisation of the width of the peak and maximisation of the separation between the peak and the noise. Low PE PHS have a very sharp peak with a very good peak-to-valley ratio, but the separation between noise shoulder and peak is very small. High PE PHS instead have very good separation but wider peaks with a lower peak-to-valley ratio.

Two FoM were proposed for the task:

$$\text{FoM} = \frac{\text{PtV} \cdot \text{width}}{\text{rHWHM}} \qquad \text{FoM} = \text{PtV} \cdot (\text{width} - \text{rHWHM}) \tag{3.1}$$

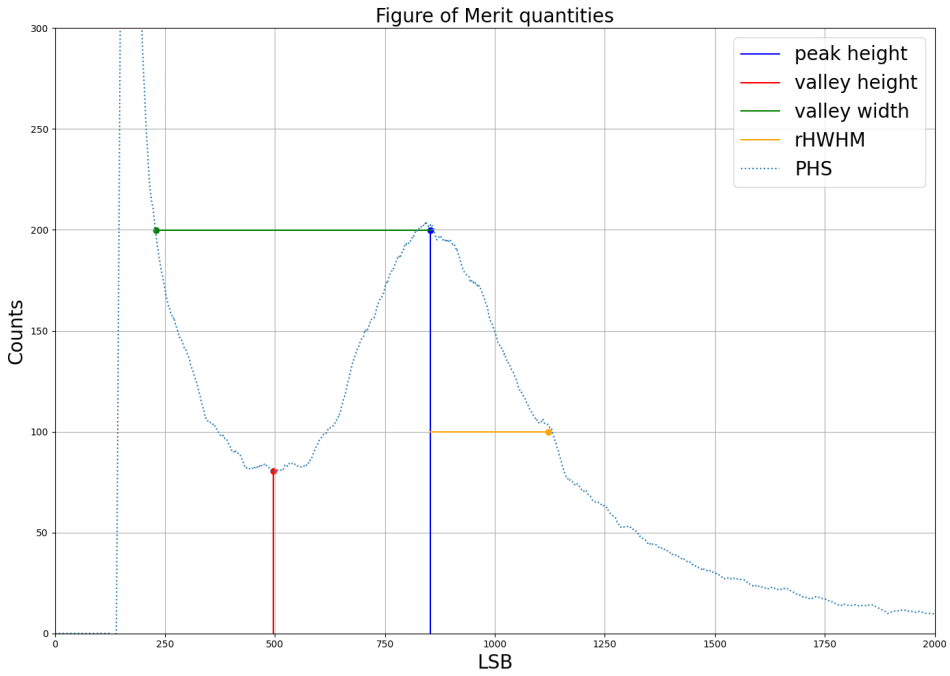


Figure 3.15: Definition of the parameters extracted from the PHS for the definition of the Figure of Merit

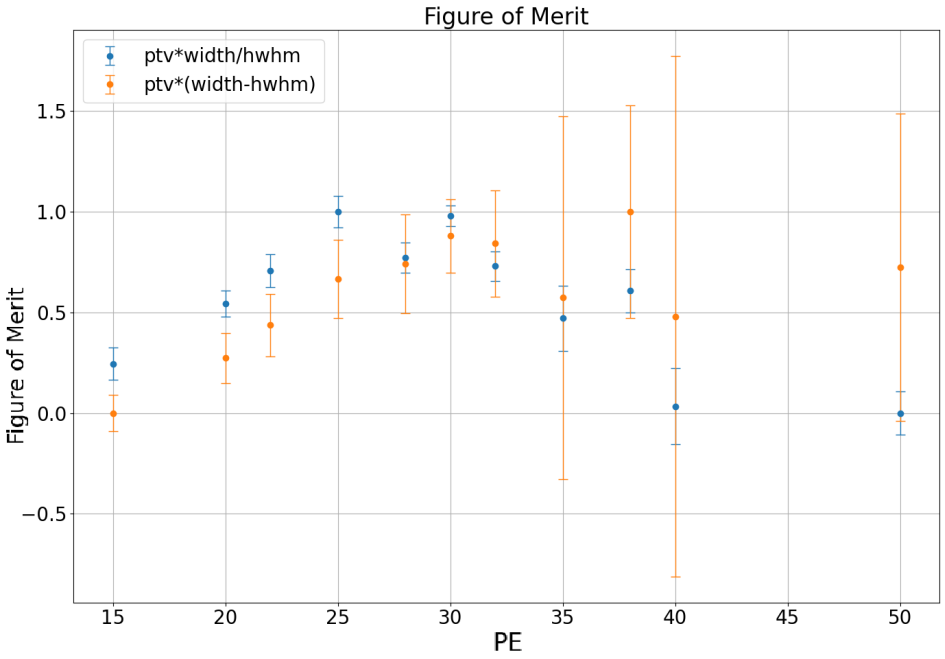
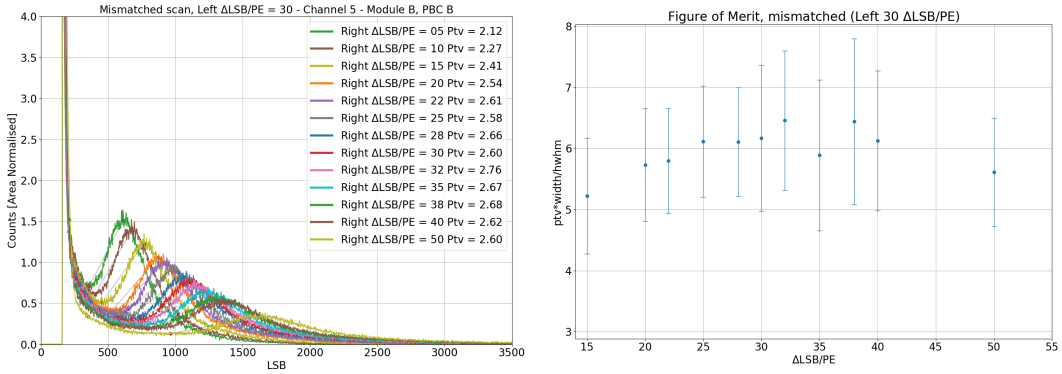


Figure 3.16: Figure of Merit comparison for module B



(a) PHS of the mismatched PE scan

(b) Figure of Merit

Figure 3.17: Mismatched run on module B, with the left SiPM fixed to a 30 PE gain

As in Figure 3.16, both indicate a preference for PE$\simeq 25-30$.

The PE scans were also used to test the effect of a gain mismatch between the two SiPMs: they were performed by keeping one SiPM fixed to 30 PE and changing the other's gain. The effect was appreciable but not so impactful to require gain-matching to a very high precision: the effect of a mismatch can be seen in Figure 3.17. The distortion can be seen in the calculation of the peak-to-valley ratio and in the figure of merit, which is higher for runs with matched gain. This result highlights that the effect of a small mismatch can be negligible.

The testing of the different modules in the expected geometry of the final barrel, matching the angle to the one in the final geometry, highlighted the difference in resolution and output. In Figure 3.18 the smoothed typical PHS for 25 PE SiPMs shows that the shorter fibres have a much more defined peak. The inner ring data was taken after a modification to module A to increase amplification and cannot be compared.

Finally, a test with high magnetic field was performed to ensure the response of the SiPMs and the electronics in the Super-MuSR working conditions would not be distorted. We were mostly concerned with alterations to the pulse shape due to magnetised components in the electronic circuit, as it was determined that a few components, such as the connector shields, were made of ferromagnetic alloy.
The analysis was performed at 1.3 T, during the ramp up, ramp down and at 0 T in the HiFi spectrometer, which is double the expected 0.65 T available to the final Super-MuSR spectrometer. The protostave equipped with module C was placed close to, but outside, the sample cavity. The SiPMs were powered to 56 V during ramp up and then changed to 54 V during ramp down, to determine if changes in the gain would also change the response to the magnetic field. The average shape of the pulses was analysed (Figure

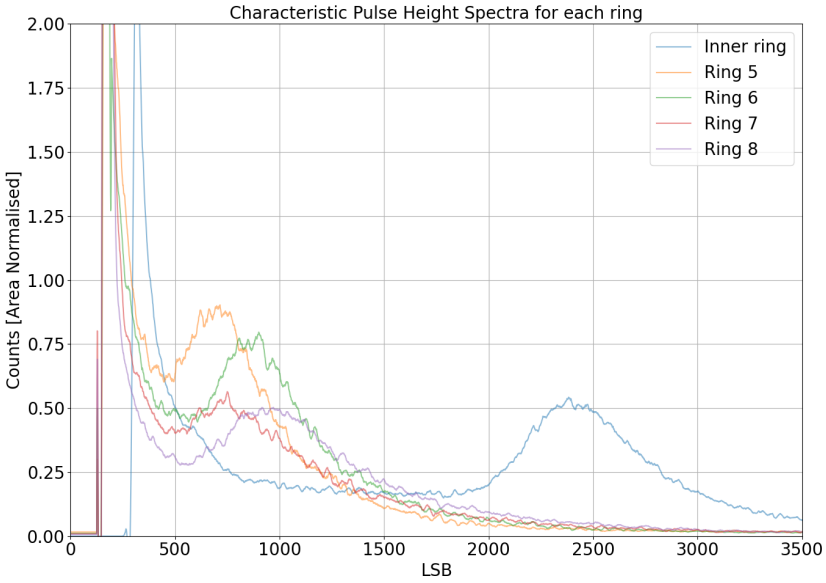


Figure 3.18: PHS for each ring, smoothed

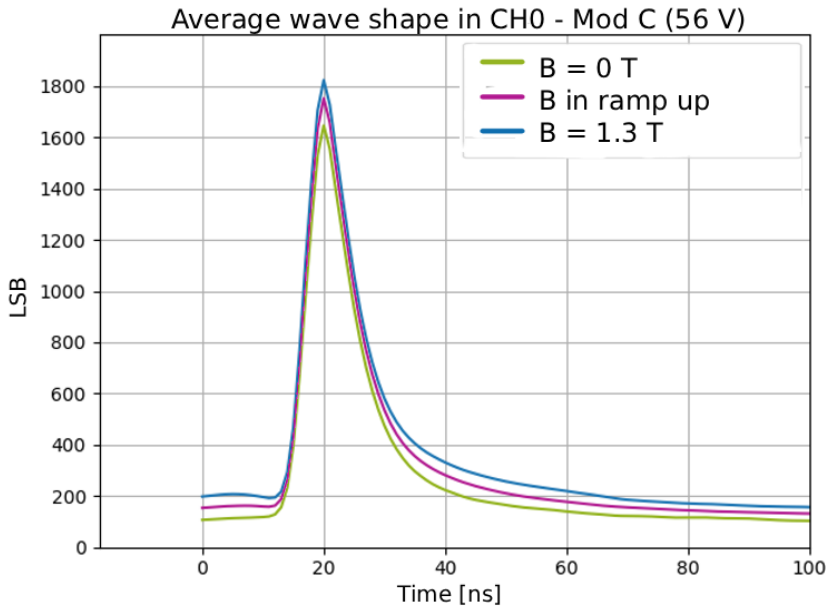


Figure 3.19: Results of average pulse shape in different magnetic field conditions

3.19), no effect correlated to the changes in magnetic field was found: the apparent baseline shift is due to changes in the density of events, thanks to the positrons being curved towards the detector.

3.3 Correlated pulses

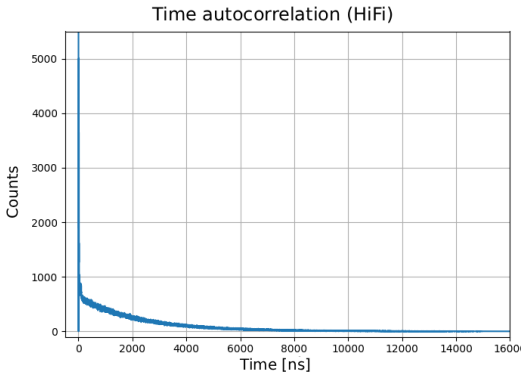
The Digital Muons project will improve the data acquisition of the muon facilities, which will provide better timing resolution and increased rate capability by switching to a digital, event-based architecture [32].

The upgrade will not only involve the new spectrometer, but also EMU, HiFi, and the RIKEN-RAL facility which will be able to make full use of the recent upgrade to the source and accelerators, resulting in the increase of the muon beam intensity.
The old data acquisition pipeline could not keep up with the increased rate, and is now working at $\sim 1/10$ of the maximum intensity. This prompted the need to invest in the improvement of the data acquisition architecture, which would increase rate capability and improve data processing: the idea is to implement a Digital Data Pipeline (DPP), based on custom FPGA-based modules (Nuclear Instruments DAQ121). The pipeline will acquire and record the waveforms for each detector within a frame (32 μs, extendable) with high fidelity, and process them with peak-finding algorithms based on the signal derivative. The process must take less than 20 ms, to avoid signal pileup with 50 Hz pulses [42].

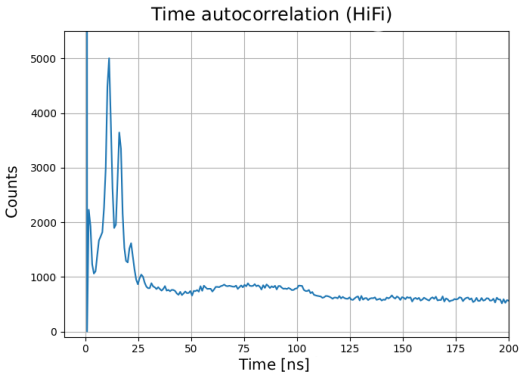
The pipeline was first developed for the Super-MuSR upgrade, but the DAQ was installed on EMU and HiFi in parallel to the current DAQ architecture, to test the setup and provide feedback before the long shutdown and installation of Super-MuSR. To do so, splitter cards were inserted in the original data pipeline to extract a copy of the signal for the DAQ121.

After the addition to the DAQ architecture, the new digitizers were able to acquire the waveforms in high resolution, providing the ability to analyse in depth not only the counting histograms, as before, but also detector signals shape and timing properties. Thanks to this, it had been observed that, in the HiFi newly digitised data, some pulses exhibited peculiar properties: first, they seemed to be of a different shape, and secondly they seemed to be correlated in time to other pulses, coming 75-120 ns after.
The causes of such pulses can be varied, either coming from the changes to the data acquisition pipeline or intrinsic to the detectors. Problems with the DAQ could be both digital, in the processing of the waveforms, such as systematic errors in peak-finding or window veto, or hardware, such as signal reflections caused by impedance mismatch. PMT ageing and afterpulsing could also be causing them, either from photons or ions.

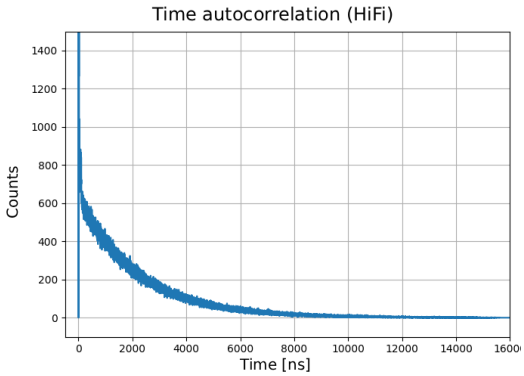
In order to identify the cause of the pulses, the first step was to pinpoint the start of the phenomenon: waveforms from the instruments were digitised in 2020, in several



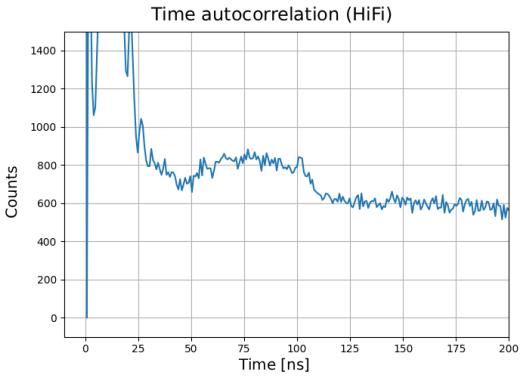
(a) Complete picture



(b) Zoomed to highlight the features close to zero



(c) Zoomed to highlight the exponential curve



(d) Zoomed to highlight the position of the anomalous signals

Figure 3.20: Autocorrelation plots for the HiFi 2020 dataset

beam and field conditions, prior to modifications to the hardware. The waveforms were digitised using a PicoScope Oscilloscope, model 6402D, with 1.25 GHz sampling rate (every 0.800 ns).
The analysis of the time correlation can then be done by measuring the time difference from a signal pulse to all the others in a frame: given that the decay of a muon is a stochastic event, the expected histogram of inter-arrival times is a negative exponential with the time constant linked to the probability of the stochastic process [5].

To locate the peaks, the `find-peaks` function from the `scipy.signal` python module was used. The function locates the position of peaks in a frame when the signal is above a pre-established threshold (15 mV, similar to a digital leading edge discriminator). The histogram of the inter-arrival times between a peak and all the others in the frame (the *autocorrelation plot*) is then produced, in Figure 3.20, excluding the negative times (the autocorrelation plot is symmetric around zero).
The process was first applied to a HiFi dataset with in the zero field condition and with

low flux (5.5 mm slits), to get the best case scenario of only a few peaks per frame: if the problem had been the electronics or digital reconstruction, the autocorrelation plot would show a simple exponential decay, with no features, but in the case of detector side problems, the expected result would be to see the same features in the 2025 and the 2020 datasets.
The result are illustrated in Figure 3.20: on the left (Figure 3.20a and 3.20c), the autocorrelation plot in a 16 μs window, on the right (Figure 3.20b and 3.20d) the zoom over the first 200 ns. The same correlation is visible in the original dataset, excluding the modifications as the cause and hinting to an intrinsic property of the photodetectors. The plots highlight an excess of events in window 50-110 ns after another pulse: this will be referred to as the *autocorrelation excess window*. The peak close at zero is due to noise summed over the waveforms, which can be reduced but not eliminated by changing the parameters of the peak-finding algorithm, such as introducing a prominence threshold. This would however remove pileup peaks, therefore we chose not to proceed with the prominence filter.

The second step of the investigation of the signals was to identify the source of the correlated pulses: in PMTs, the main culprits of correlated pulses is afterpulsing, caused either by photons emitted from late stage dynodes or by ionised gas due to infiltrations. The first thing to note is that HiFi PMTs are Hamamatsu R5505 PMTs, cylinders 25 mm in diameter and 40 mm in height [43]: photons travelling at light speed emitted from the last dynode would take around 0.1 ns to get to the photocathode to initiate a new pulse, and we could not explain the 50-110 ns delay in the correlation, even if only emitted at the end of the pulse.
The first hint at the cause being ionic afterpulses came from a study of the dependence of the autocorrelation excess window on the voltage applied to the PMTs. The results are in Figure 3.21: the study highlights that correlated pulses happen closer to the main pulses for higher voltage, consistent with the hypothesis of massive charged particles being accelerated in an electric field.

Gas infiltration in the PMT vacuum could be a production defect, which would be possible because the same problem is found in different PMTs of the same instrument, or it could be caused by gas infiltration during the years of use in the facility. The easiest way to discriminate between the two cases is to look at a different set of the same PMTs: in this case, the EMU PMTs are the same model and size, but a few decades older: if the cause is infiltration, we expect the same type of correlation with similar or greater intensity, as the two instruments are in the same environment, just metres apart, while if the cause is from the construction, we would not see the same in EMU. The analysis of a 2020 dataset acquired on the EMU spectrometer, with 20 mm slits and zero field, yields the result in Figure 3.22, where the same excess can be seen: this is a hint towards the infiltration hypothesis, with the most likely cause being helium. This is a known phenomenon in PMTs and other objects where vacuum is applied: helium infiltration is likely because the noble gas has a relatively high permeability through glass compared to other

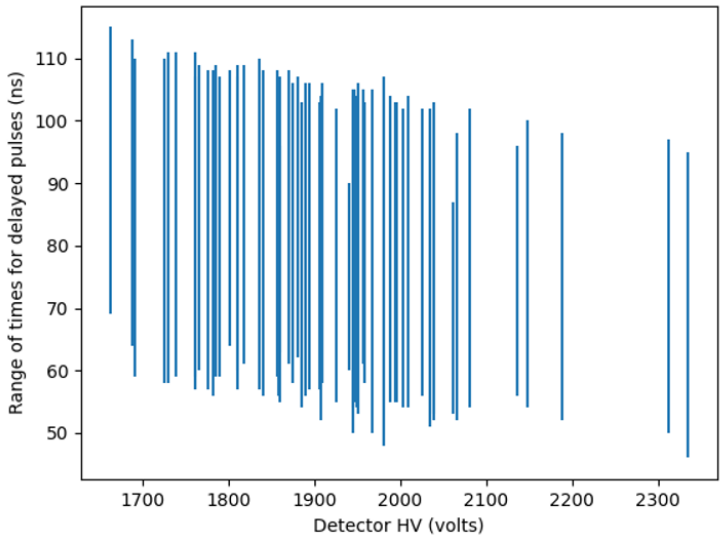


Figure 3.21: Position and size of the autocorrelation excess window as a function of the HV voltage applied to the PMTs (courtesy of J. Lord)

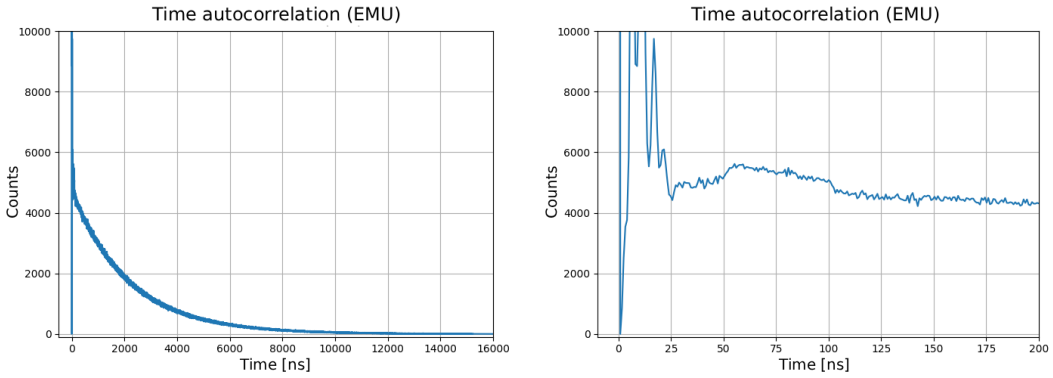
gasses in the atmosphere [44].

After identifying the culprit of the signals, it was noticed that such signals could affect the asymmetry calculation, and a method to discriminate the afterpulses from the real positron signals should be devised, if possible.
Because the first hint at a problem came from the fact that signals looked different, a quick way to differentiate them could lie in some type of pulse shape discrimination. The process entailed the identification of the peaks, fitting with a function representative of the PMT signal shape, the rejection of the poor fits and then the plotting of correlation plots between different function parameters.
After finding the peaks, the waveform is extracted (10 ns before the peak and 24 ns after) and fitted, using the python function `curve_fit` from the library `scipy.optimize`. The function of choice for fitting was the Back-To-Back exponential [45], which had already been used to fit ISIS PMT pulse shapes:

$$f(x) = I\frac{AB}{2(A+B)}\left[\exp\left(\frac{A[AS^2+2(x-X_0)]}{2}\right)\operatorname{erfc}\left(\frac{AS^2+(x-X_0)}{S\sqrt{2}}\right)+\right.$$
$$\left.+\exp\left(\frac{B[BS^2-2(x-X_0)]}{2}\right)\operatorname{erfc}\left(\frac{BS^2-(x-X_0)}{S\sqrt{2}}\right)\right] \quad (3.2)$$

The parameters extracted from the fitting routine were:

- A and B: respectively rise constant and fall constant of the exponential, measured in [1/ns]. The larger the value, the steeper the exponential curve;
- S: the width of the shape, measured in [ns];



(a) Zoomed to highlight the exponential curve (b) Zoomed to highlight the position of the anomalous signals

Figure 3.22: Autocorrelation plots for the EMU 2020 dataset

- X_0: the position of the peak in the fitting window, then translated to the position in the frame;
- I: the area of the pulse;
- PH: the pulse height, derived numerically from the final fit.

The resulting fits were tested to ascertain the goodness of the representation of the wave: the R^2 test was used. The coefficient of determination, R^2, is defined as:

$$R^2 = 1 - \frac{\sum_i^N (y_i - f_i)^2}{\sum_i^N (y_i - \bar{y})^2} \qquad \text{with } f_i = f(x_i),\ \bar{y} = \frac{\sum_i^N y_i}{N} \tag{3.3}$$

The R^2 parameter is between 0 and 1, the closest to 1 and the better the model fits the data. The fits with $R^2 < 0.985$ were excluded, with the threshold chosen conservatively by visual analysis. A small sample of the plots can be seen in Figure 3.23.

The correlation plot that exhibited the most significant features is the one comparing width and fall constant of the waves (Figure 3.24). The plot exhibits a peculiar distribution, where most of the pulses aggregate in a wide area with larger width and small fall constant, but small set of waves have a well-defined, smaller width and larger fall constant: these are the sharp waves noticed during HiFi use. Several different detectors, both backward and forwards and in different flux conditions, exhibited the same correlations. As expected, the EMU plot shows a much larger percentage of peaks outside the "normal" region, due to the much longer time for ion accumulation after production.

The final confirmation of the relation between the sharp pulses and the afterpulses is obtained by looking at the autocorrelation plot with the sharp peaks removed: this is done by defining a separation boundary (the line in Figure 3.24) and tagging the waves that fall to the right. The autocorrelation is then calculated excluding the tagged waves: the comparison, in Figure 3.25, clearly shows that by simply removing those pulses,

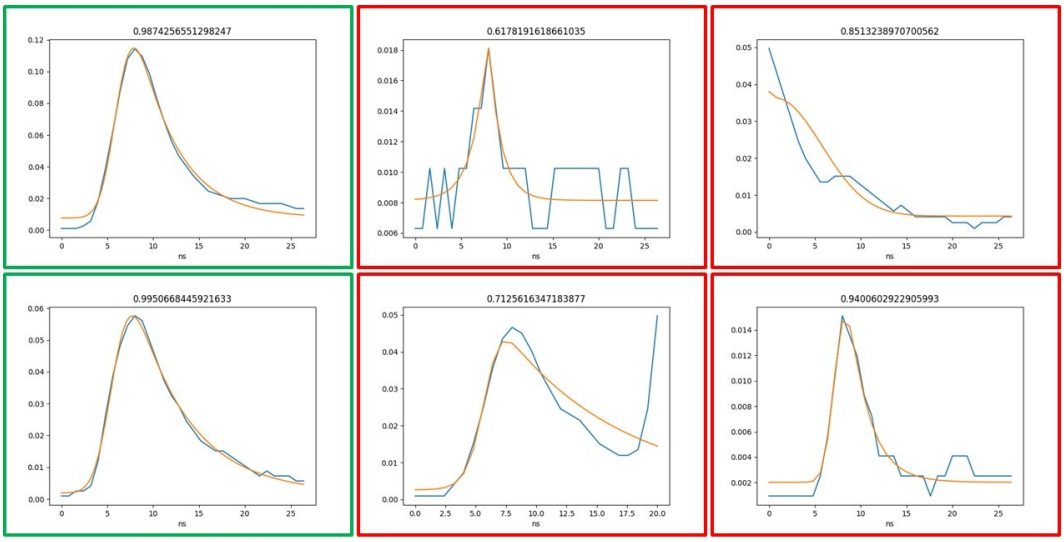
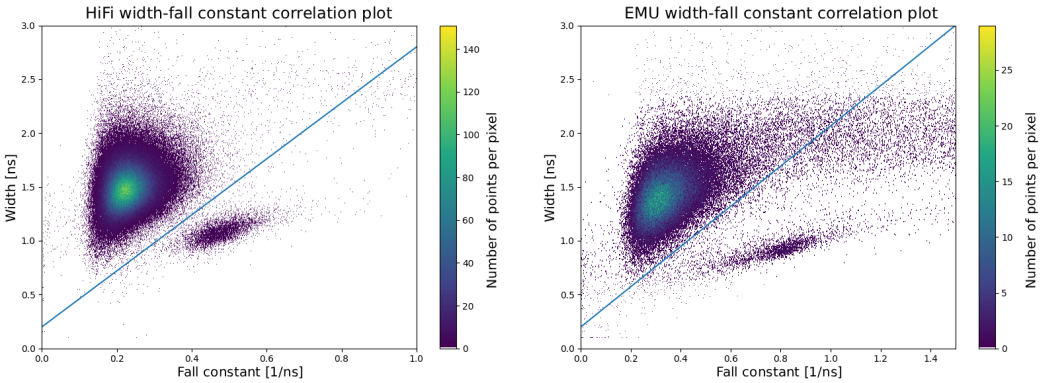
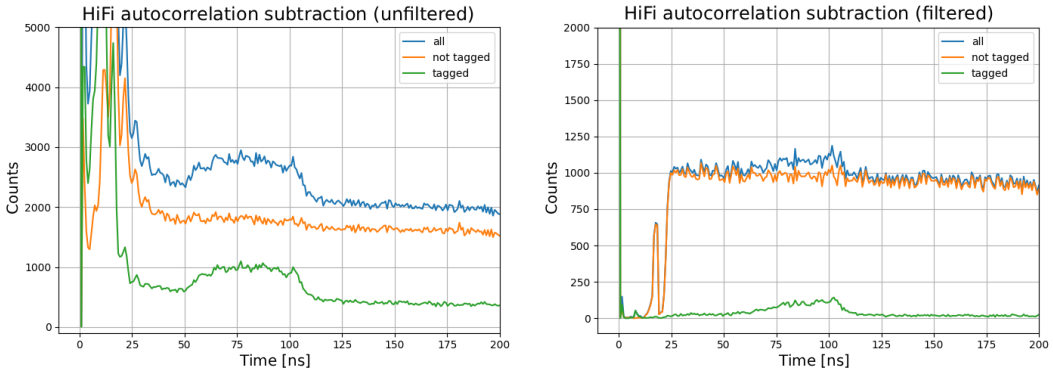


Figure 3.23: Sample of the fitted waves, using the function in Equation 3.3. In green, the fits that passed the R^2 test, while in red those that failed. The titles show the R^2 parameter



(a) HiFi (b) EMU

Figure 3.24: Correlation plots between width and fall constant of the waves. The line indicates a possible separating boundary between the real pulses and the afterpulses



(a) Unfiltered (b) Filtered

Figure 3.25: Autocorrelation plot comparing the calculation with all peaks to the one without sharp peaks (not tagged) and to the one with only sharp peaks (tagged)

the afterpulses are eliminated. The process can be done for both filtered and unfiltered pulsed: the R^2 test is unfortunately biased toward tall and isolated peaks, but by looking at the comparison between the filtered (Figure 3.25b) and the unfiltered (Figure 3.25a) plots, the same effect is visible.

The reason for the difference in shape is unknown, but one hypothesis lies in the mechanism of scintillation. The pulse height and width correlation did not exhibit any peculiar patterns (Figure 3.26), so the width and pulse height of the afterpulses could be uncorrelated: another strong hint for the helium theory, as ionic discharge can be of any height (different from photon-induced afterpulses). One of the reasons for this is that PH is defined by the gain and number of photons at the photocathode, while width is defined by the time difference of arrival of the first to the last photon at the photocathode. A normal signal is defined by the time spread caused by the scintillation phenomenon and the transit in the light guide, while electrons produced in the photocathode by the collision of a heavy ion are released all at the same time. The width of a real pulse would therefore be correlated to the scintillation phenomenon, and having a much sharper signal for the ion would be because the signal does not depend on scintillation.

The pulse height of the afterpulses may not be correlated to the width, but it is correlated to the delay: if an atom is ionised close to the photocathode, it has less space to travel so it will be closer to the main peak in time, but it will also accelerate less due to the limited potential differential between the first dynodes and the cathode. The opposite happens if the ion is produced close to the anode: the effect is visible in Figure 3.27, strong evidence for the cause of the pulses being accelerated massive particles.

The MuSR spectrometer was also analysed, despite it using a completely different

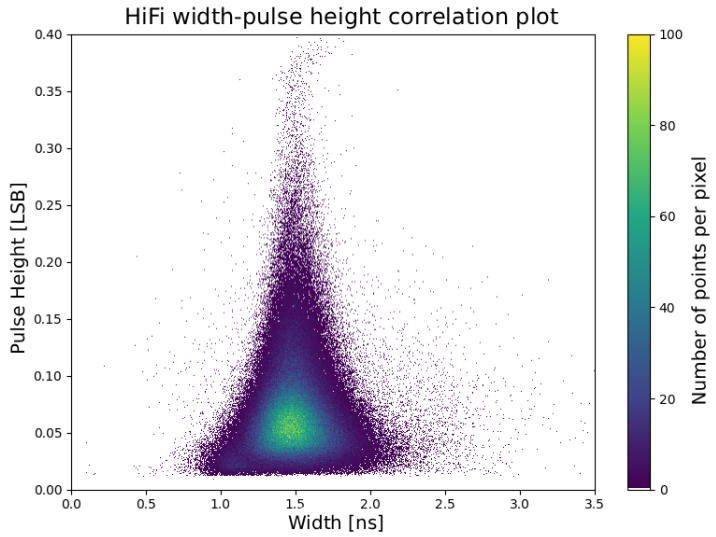


Figure 3.26: Pulse Height-width correlation plot. No peculiar distribution for the afterpulses is visible, meaning that those pulses can be of any height retaining a sharp shape

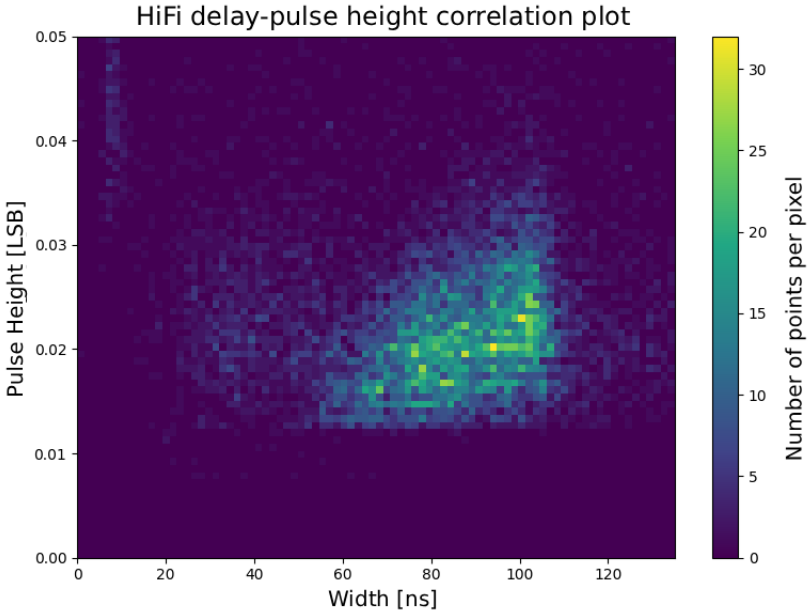
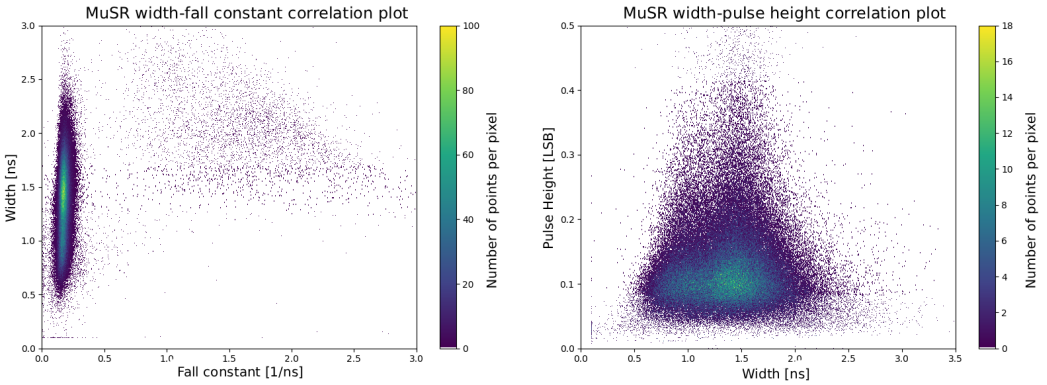


Figure 3.27: Pulse Height-delay correlation plot for the tagged waves, in the HiFi dataset



(a) Width-fall constant (b) Pulse height-width

Figure 3.28: MuSR correlation plot

PMT (2-inch Hamamatsu EMI-9954). The width-fall constant exhibited a completely different behaviour (Figure 3.28), which was not investigated further as the MuSR spectrometer is shortly going to be decommissioned. There seemed to be two classes of events in the pulse height-width correlation plot.

The ionic afterpulsing effect has been found to happen in several different PMTs in both HiFi and EMU. The fitting routine cannot be applied directly to the digital pipeline, due to the calculation taking more time than the 20 ms between muon pulses, which would cause pileup in the DAQ pipeline. A method to implement the discrimination will be devised, as the afterpulses could skew the asymmetry calculation. This is the great advantage of using a digital data pipeline with high precision digitizers: waveform analysis can be extremely useful, both for finding and addressing any undesired effect in the detectors, such as background, noise and afterpulsing.

Chapter 4

CHNet-MAXI

4.1 The experiment

CHNet-MAXI is an INFN (Istituto Nazionale di Fisica Nucleare, National Institute for Nuclear Physics) funded experiment for the feasibility study of isotopic analysis of lead in historical artefacts using μ-XES. It is part of the Cultural Heritage Network (CHNet) of INFN, a collection of applied physics experiments aimed towards the use of physics techniques for the innovative analysis of cultural heritage.

The project, a collaboration started in 2023 between the Universities and INFN divisions of Pavia and Milano-Bicocca, University of Sassari and INFN division of Cagliari, took data in December 2025. The experiment consisted in the measurements of several standards and an archaeological lead coin, aiming to the extraction of a quantitative measurement of the isotopic ratios and the definition of a benchmark for this type of analysis.

4.1.1 Introduction

The use of elemental lead analysis lays in the fact that lead isotopic ratios are a fingerprint of the geological area of excavation of the ore: they provide geographical information on the origin of the lead, a small but valuable piece of the puzzle for historians.

Isotopic ratios tracking is a known technique, common elements chosen being strontium and lead. With strontium, the measured ratio is Sr-87 to Sr-86, but accumulation can happen in different ways, and it is not always possible to distinguish the origin. Lead instead is present in traces in virtually all rocks, both igneous, metamorphic and sedimentary. It often comes from the decay of uranium and thorium (it is the endpoint of the 3 of the 4 alpha decay families), it is then leeched to surrounding rocks where the ratios stabilise and keep constant until present day. The phenomenon that creates the rocks, such as tectonic plates or volcanic eruptions, generates a signature that is mostly dependent on the geologic era when the phenomenon happened and the location. As

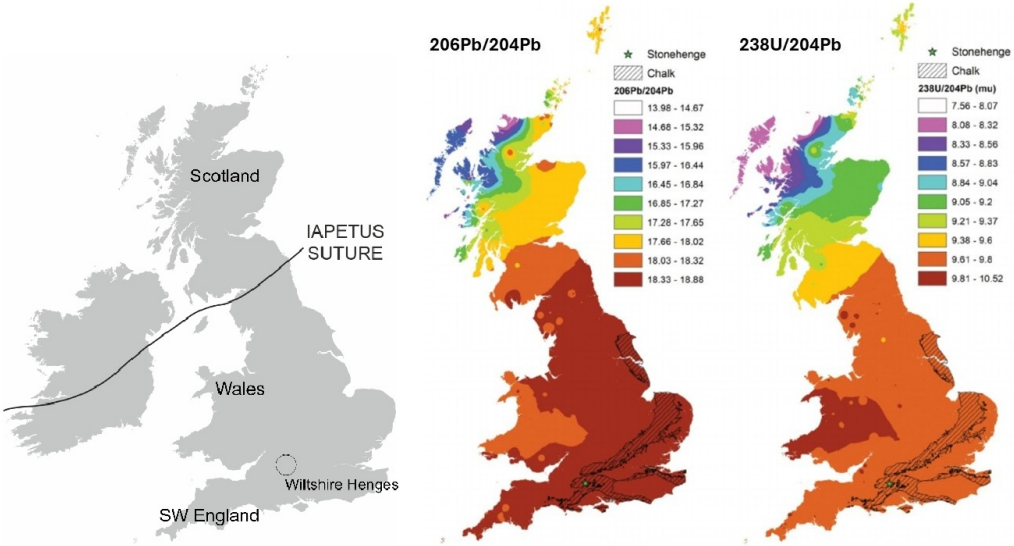


Figure 4.1: Map of Great Britain with the isotopic ratio of lead 205 to lead 204 in the centre, on the right the ratios of uranium 238 to lead 204. On the left the position of the suture line. The same map can be recreated via geological survey of any region of interest [46]

seen in Figure 4.1, the ratios both depend on the geography and the age of the rocks [46].

The choice of μ-XES as a technique came from the need of a non-invasive, non-destructive technique for the measurement of precious artefacts: the most common isotopic analysis tool is mass spectrometry, but it requires the removal of a sample of the object and its destruction. Moreover, the invasive analysis is much more destructive if we want to analyse the bulk, where surface contamination is minimal. Muons allow the flexibility of using the same technique for both surface and bulk, reducing handling and transportation, and without the need of any intervention on the object.

4.1.2 Experimental setup

The experiment was conducted at RIKEN-RAL, in port 4. The port can provide both positive and negative muons up to 120 MeV/c momentum, providing the flexibility of choosing the implantation depth.

The samples chosen were isotope standards of lead (204, 206, 207 and 208), with a thickness of 0.25 mm and an approximate size of $\sim$ 2 cm in diameter. The standards consist of certified samples of lead with a known isotopic composition, acquired from CK Isotopes. The samples were enriched with the specific isotope to more than 90%, the lowest being lead-207 at 92%.

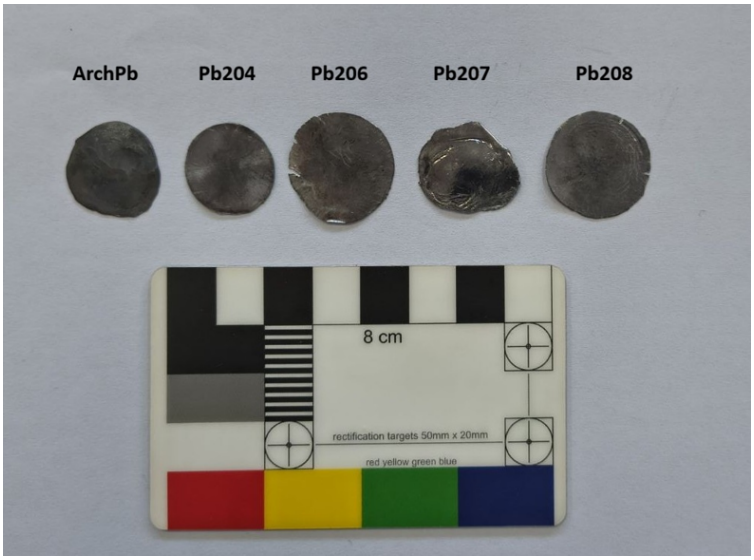


Figure 4.2: Samples enriched to around 100% in isotopes Pb-204, Pb-206, Pb-207, PB-208, and an archaeological sample from the INFN Roman lead collection, measured in the December 2025 CHNet-MAXI experiment

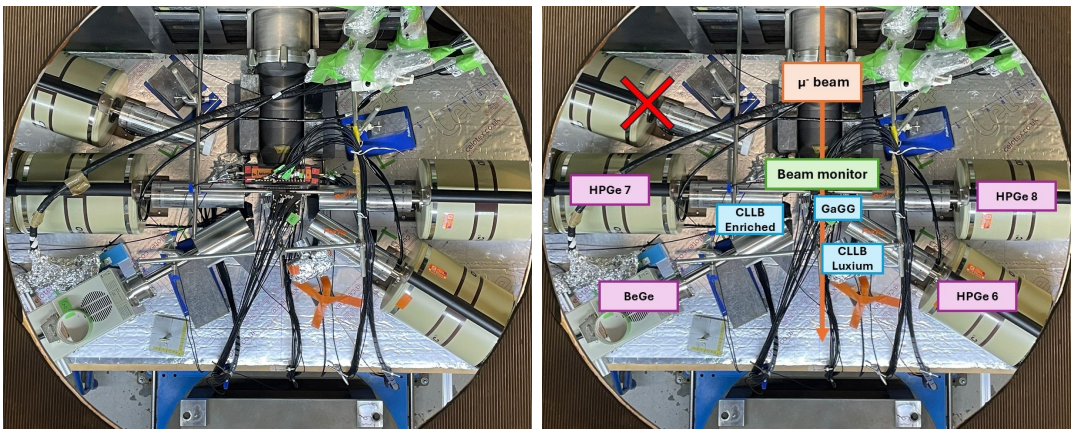
Table 4.1: Summary of measurement time and momentum for the samples

Sample	Measurement time (hr)	Momentum (MeV/c)
Pb-208	10.5	49.5
Pb-206	11	49.5
Pb-207	23.5	49.5
Pb-204	23	49.5
ArchPb (thin)	30.5	49.5
ArchPb (thick)	6	60

To confirm the accuracy of the results, an archaeological sample was also measured, from the INFN Roman lead collection[1], both a thinner one, similar to the standards, and one with a slightly bigger size (0.5 mm in thickness and $\sim$ 3 cm in diameter, not pictured). The samples can be seen in Figure 4.2. Each sample was measured for several hours, as summarised in Table 4.1[2]. The momentum chosen for the beam was 49.5 MeV/c, stopping the muons in the middle of the 0.25 mm discs. For the thicker archaeological sample, the momentum was increased to 60 MeV/c to better measure the bulk.

[1]The INFN Roman lead is the remnants of a shipment of lead from Sardinia to Rome, shipwrecked. It is poor in lead-210, being shielded from cosmic rays for two millennia in the Mediterranean Sea: Pb-210 is the main culprit of low energy background contamination from lead, and a lot poor in Pb210 means that shielding for low background experiments can be the most effective

[2]Lead-206 was measured in three shorter runs, summing to 11 hours



(a) Top view (b) With labels

Figure 4.3: Top view of the CHNet-MAXI setup at port 4. On the right the detectors are highlighted

The measurements were carried out in port 4, using both the ISIS High Purity Germanium (HPGe) detectors and custom detectors built for the task. The detectors in the CHNet-MAXI setup, in Figure 4.3, are:

- 4 high purity germanium detectors:
  - Ge6 (10.6 cm from the sample, 30° from the beam direction);
  - Ge7 (10.5 cm, 90°);
  - Ge8 (10.4 cm, 90°);
  - Ge9 or BeGe (12.5 cm, 30°);
- 2 CLLB scintillating crystals read by SiPMs (14 cm, 60° and 10 cm, 30°), specialised in neutron identification via PSD;
- 1 inorganic GAGG scintillator array (18 cm, vertically from the sample), with 16 different crystals;
- a 64-channel muon beam monitor, placed upstream of the sample, flush to the nose of the beam. The beam monitor consists of two orthogonal planes of 32 plastic scintillating fibres each.

The sample was suspended in front of the nose via the use of a aluminium foil pouch, placed 10 cm from the nose.

A calibrated, non normalised spectrum that shows the characteristic isotopic low energy X-rays can be seen in Figure 4.4.

The entire setup was developed to test the use of a portable setup, using small scintillators, the facility-provided HPGe detectors and simple setup for powering the detector and the data acquisition.

Port 4 provides two different acquisition modes: the analog acquisition via Ortec

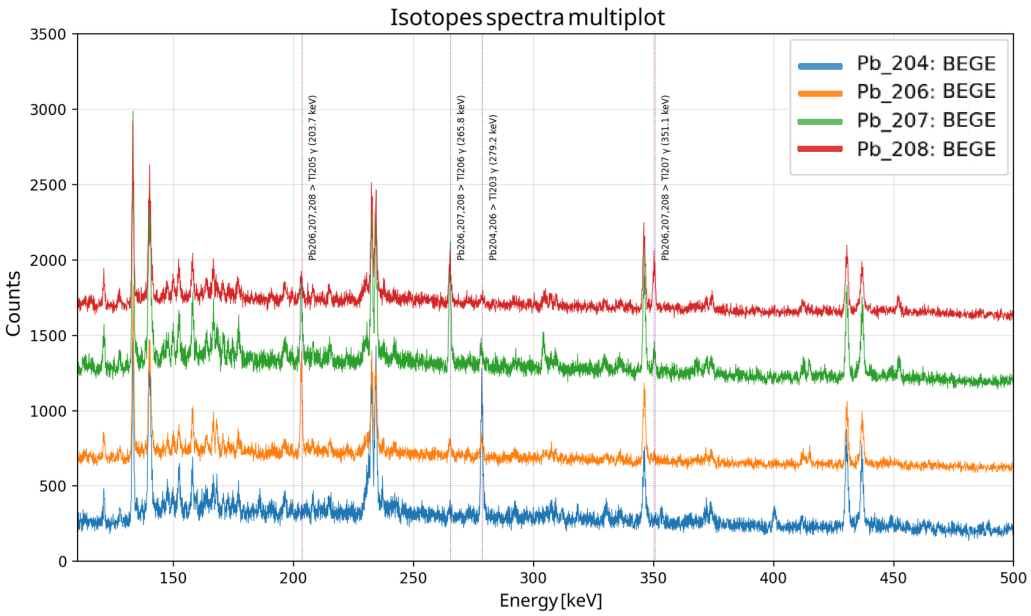


Figure 4.4: Spectrum of the measured X-rays in the 0 keV to 500 keV region, using the BeGe detector (courtesy of M. Cataldo)

modules, or the digital acquisition via Hexagon modules. Both have their drawbacks, such as the number of channels available and the inability of adding custom processing, such as coincidence.
The data acquisition of the experiment is instead based on a single rack module, the CAEN R5560SE digitizer. It has 128 Channel (divided into 4 independent blocks), digitising at 125 MS/s (8 ns) with 14-bit precision. It is a Open FPGA digitizer, meaning that using the SciCompiler program it is possible to set the FPGAs up with ease, without the use of the VHDL programming language [47]. The firmware for the digitizer, developed in conjunction with Nuclear instruments and INFN Milano Bicocca, and the graphical user interface (GUI), provide a complete package that allows the control and readout of all detectors. The germanium detectors are at the same time acquired using the R5560SE, split between high and low gain, and the Hexagon digitizer, used as backup because it was the first time testing the full setup.
Each of the blocks is programmed for a subset of detectors, in slightly different ways. The digitizer and the division can be referenced in Figure 4.5.

In particular, the division is as such:

- Analog 0: beam monitor, vertical fibres;
- Analog 1: beam monitor, horizontal fibres;
- Analog 2: ADC for scintillator coincidences (GAGG);
- Analog 3: simple ADC for all other spectra (germanium detectors, CLLB).

The beam monitor data acquisition will be explored in the next section.



Figure 4.5: CAEN R5560SE digitizer, with highlighted the function of each block [47]

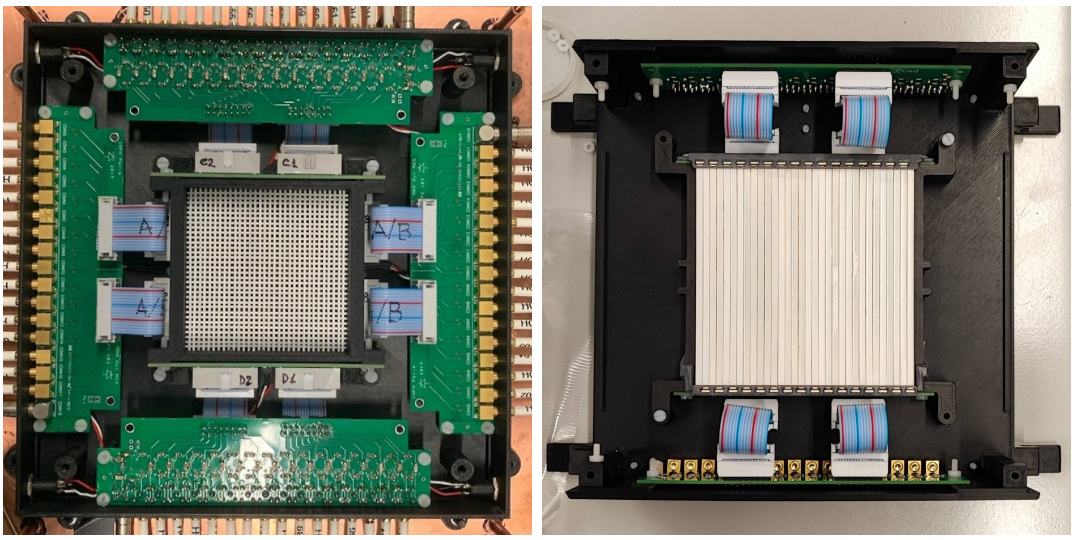
4.2 Beam monitor

A beam monitor is a type of detector dedicated to the analysis and monitoring of the beam flux, position and stability. The design can vary, depending on the particles in the beam and the needs of the experiment.

The muon beam monitor for the CHNet-MAXI experiment was developed inheriting the design and expertise from the FAMU beam monitor. The FAMU/CHNet-MAXI beam monitors rely on the use of perpendicular scintillating fibres for the identification of the muons position and flux: a hit in the front plane and one in the back plane signify the passage of a particle. The main difference between the two is the size of the fibres: FAMU relies on all muons stopping in a low density gas, so the beam monitor must be as thin as possible to minimise momentum straggling and reduce the energy absorbed in the fibres. CHNet-MAXI instead requires muons to stop in lead, much denser, which reduces the effect of straggling. The CHNet-MAXI beam monitor therefore uses a simpler design for ease of use, construction, and extraction of data. Both use polystyrene square fibres, FAMU uses 1 mm thick fibres with 1 mm interspacing, resulting in an area of 64 × 64 mm^2 (Figure 4.6a), while CHNet-MAXI uses 3 mm fibres with no interspacing, resulting in a 96 × 96 mm^2 area (Figure 4.6b). The area of the FAMU beam monitor is noticeably smaller despite the gaps between the fibres, due to a limited availability of channels and the need to use thin fibres.
The fibres are cut to the appropriate length (104.3 mm for CHNet-MAXI to accommodate the paint layer thickness) and then coated in a thin layer of TiO_2 diffusive paint, to improve the internal reflection and optimise light transport [12, 48, 49].

The fibres are read out at alternating ends by silicon photomultipliers (see Figure 4.7a), chosen for their small size and fast response. For the CHNet-MAXI beam monitor, the SiPMs are Hamamatsu S14160-3015PS, with a breakdown voltage in average of 39.2 V, the distribution seen in Figure 4.7b [49].

The beam monitor is housed in a custom, 3D printed case (Figure 4.8a), ensuring



(a) FAMU beam monitor [48] (b) CHNet-MAXI beam monitor [49]

Figure 4.6: Comparison of the two beam monitors

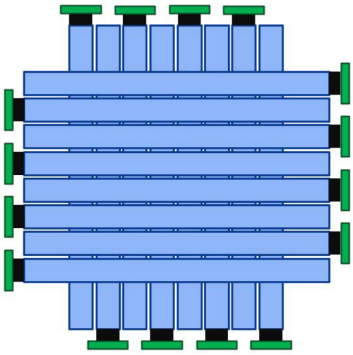
light tightness with a thin (0.5 mm) PLA front and back window, and an external case holding the custom boards for the SiPM readout and the distribution of the high voltage, designed by INFN Pavia. The HV is handled via the DT8032M HV desktop CAEN module. The output of the boards is read via MCX-MCX equal-length cables.
In addition to the case, the beam monitor is held upright via the use of aluminium mounting brackets, as in Figure 4.8b. These were reduced in size for the December 2025 run to make space for the scintillating detectors.

4.2.1 Data acquisition

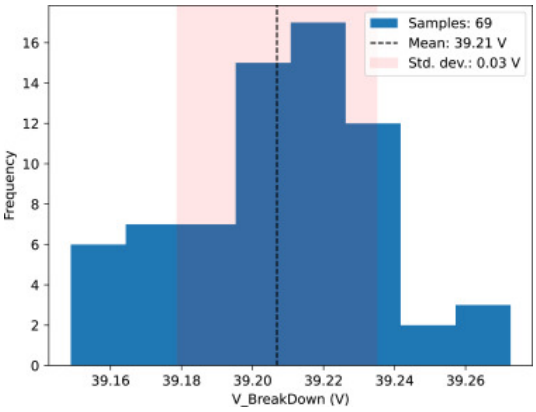
The data acquisition for the beam monitor is handled using a custom firmware and software from the collaboration between INFN Milano Bicocca, Nuclear Instruments and University of Pavia, which I worked on during my bachelor thesis.

The acquisition is based on the R5560SE Open FPGA digitizer, which can be used with the SciCompiler tool to produce FPGA-compatible scripts by using block-style programming, akin to connecting cables to default or custom modules in digital crates. The digitizer is then able to analyse the waveforms to extract meaningful, distilled information. The digitizer works at a 125 MS/s sampling rate, extracting the waveforms from the SiPMs. The charge for each channel is collected and integrated over the pulse to extract a current [47].

The beam monitor firmware allows multiple operating logics, the two most used being the internal and external trigger modes.

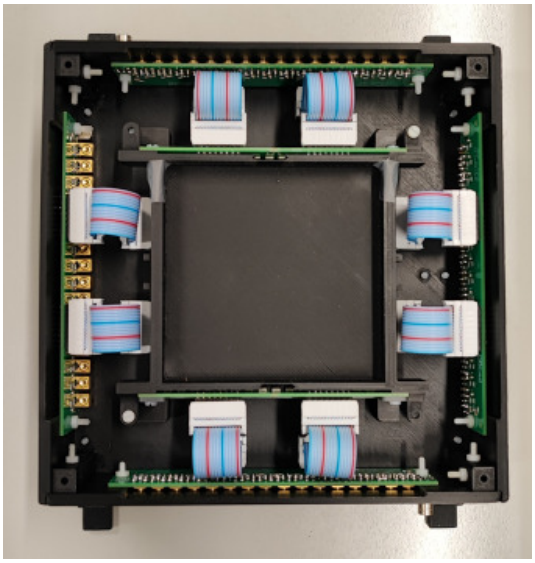


(a) Layout of the SiPM readout for the CHNet-MAXI scintillating fibres



(b) Distribution of the breakdown voltage of the SiPMs used [49]

Figure 4.7: SiPMs used for the CHNet-MAXI beam monitor: layout and properties



(a) Custom housing for the beam monitor [49]



(b) Beam monitor mounted upright for measurements

Figure 4.8: Housing and mounting for the beam monitor

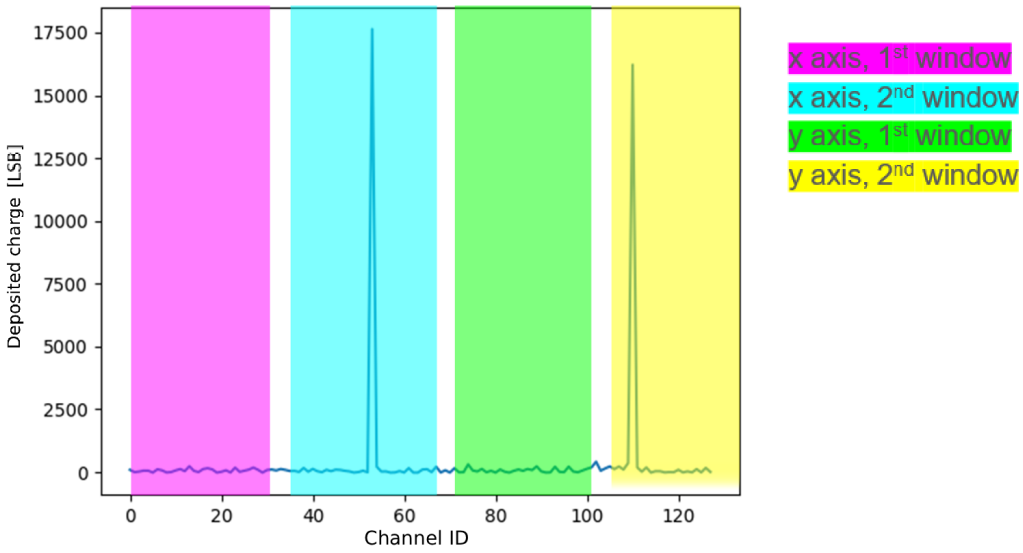
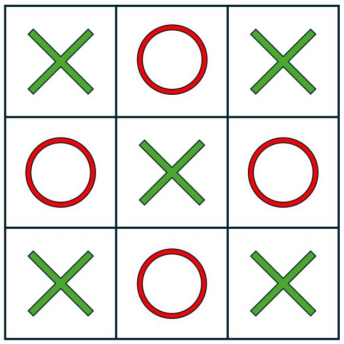


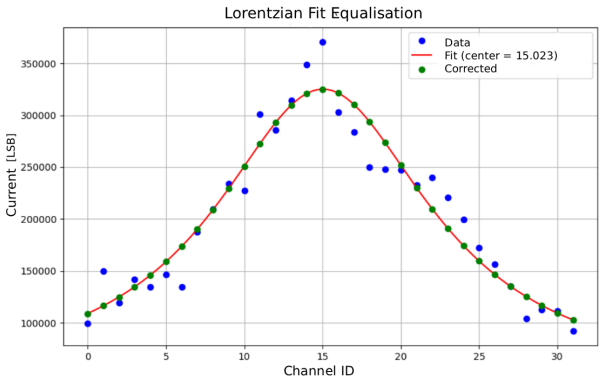
Figure 4.9: Example of a coincidence in the second temporal window (second spill), highlighting the arrangement of the data in the current output

Coincidence mode/internal trigger The waveforms are analysed using a discriminator set with a threshold, producing a digital signal for a `hit`. The signal is also integrated over a variable window depending on when the hit is detected: this is the `current` variable. The energy is timestamped and recorded if a coincidence is detected between a fibre in the X plane and one in the Y plane.
This mode is mostly used during tests using radioactive sources or if the trigger is unavailable.

External trigger The waveforms are acquired and integrated over the two bunch windows, thanks to the external trigger provided by the Cherenkov detector placed near the muon-producing target. The delay, separation, and size of the windows are all programmable parameters that need to be optimised to the beam, set thanks to registers, making the system adaptable to multiple beam scenarios.
The waveforms are simply integrated over the window (with baseline subtraction) and the `current` is recorded for each trigger. The currents are divided per fibre and per bunch, allowing the reconstruction of the front and back plane behaviour, in the first and second spill.
This mode is used during the actual data taking runs.
In case of external trigger, the data output for the current for each trigger is arranged in a 128-column array, as in Figure 4.9.



(a) Position of the source for the first equalisation (green crosses) and the second (red circles)



(b) Example of Lorentzian fit and equalisation calculation for the X profile, with the source in the central position

Figure 4.10: Equalisation procedure: positions and results of the fit using Equation 4.1

4.2.2 Equalisation

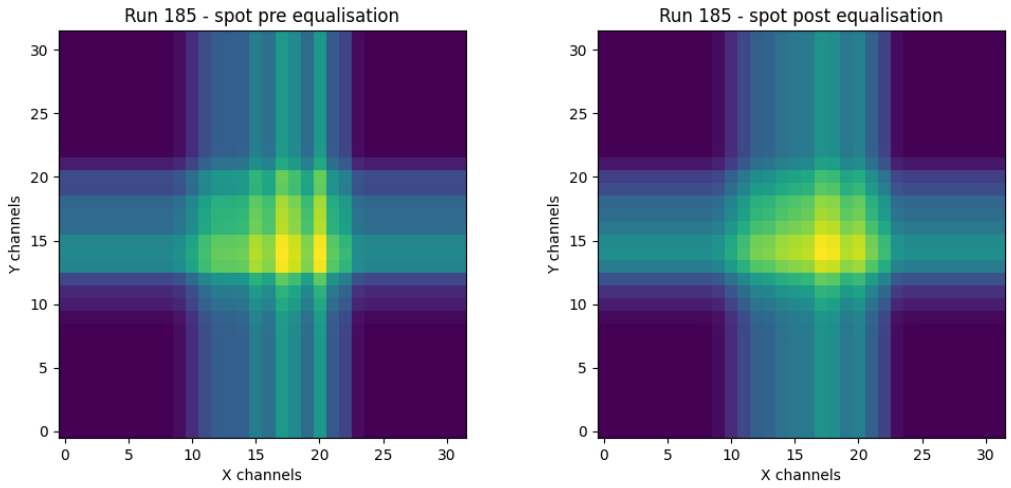
The initial calibration step, before starting the data taking, was the equalisation of the gain of the SiPMs. For ease of use, all SiPMs are powered to the same voltage, 46 V, the small differences in gain are then corrected by using a digital gain controller for each channel, which multiplies the analog input by the number selected.

To determine the digital gain factors, an equalisation procedure was developed. We took a radioactive source of Cobalt-60, emitting gamma rays that have a similar release of energy in the fibres as $\sim$ 50 MeV/c muons. The source was placed in strategic positions on the surface of the beam monitor (Figure 4.10a), flush with the external frame, 2 cm above the surface of the fibres. The profile was then acquired using the internal coincidence mode for 1 minute each.
The equalisation procedure was performed, calculating the variations, and applied. The process was repeated for 4 more points on the surface, and additional corrections were made.

For each plane (X and Y fibres), the extracted current profile was fitted using a Lorentzian curve: the theoretical form of the 2D projection of a spot due to a point-like source over fibres is the convolution of the Lorentzian and the Gaussian functions, also known as the Voight function, but in the case of thin fibres with a source placed far enough, the Lorentzian is a good approximation. The (normalised) Lorentzian function is:

$$f(x) = \frac{1}{\pi} \frac{\Gamma}{(x - x_0)^2 + \Gamma^2} \tag{4.1}$$

The main parameters are Γ (the half width at half maximum) and the position of the



(a) Before equalisation (b) After equalisation

Figure 4.11: Effect of equalisation on the beam spot for Run 185 (test run)

centre.
The points were fitted using the `curve_fit` routine from the `scipy.optimize` library in python, resulting in a fit such as Figure 4.10b. The correction factor was calculated as:

$$G_{ch} = 1 + \frac{q_{fit}}{q_{exp}} \tag{4.2}$$

For each of the positions the correction factor was calculated only for 6 fibres on each side of the centroid of the curve, providing 2-3 correction values for each fibre, which where then averaged. After applying the correction, the second equalisation fine-tuned the correction factors, resulting in the change visible in Figure 4.11.

4.2.3 Absolute calibration

One test run (run 185), with the X-ray detectors acquisition not optimised, was performed with lead-208 to confirm everything was working correctly and the acquisition ran smoothly. This was not the case, as the acquisition stopped at 5 hours, instead of the expected 9. The acquisition problems were resolved, and run 185 was chosen to estimate the absolute intensity of muons through the beam monitor for an absolute calibration, in order to compare the rate to other experiments and facilities.

What we expect to extract is the instantaneous muon intensity through the beam monitor [48]:

$$\varphi_\mu(t) = r \cdot \frac{1}{k_g \cdot Q_\mu} \cdot Q(t) \tag{4.3}$$

Where φ is the instantaneous intensity of the beam in muons per second, hypothesizing a fixed rate, r is the rate, Q_μ is the single muon charge and k_g is the simulated, momentum dependent, geometric factor, which considers the fraction of muons stopping in the active and dead layers.
This formula can be used to estimate the beam intensity in real time, from the extraction of the integrated charge per single trigger, $Q(t)$, which is the sum of the currents of the 64 SiPMs in both bunches.

Geometric factor simulation

The geometric factor was simulated via the use of a GEANT4 simulation. The factor is:

$$k_g = \eta_0 \cdot W_0 + \eta_{1f} \cdot W_{1f} + \eta_{1b} \cdot W_{1b} + \eta_2 \cdot W_2 \tag{4.4}$$

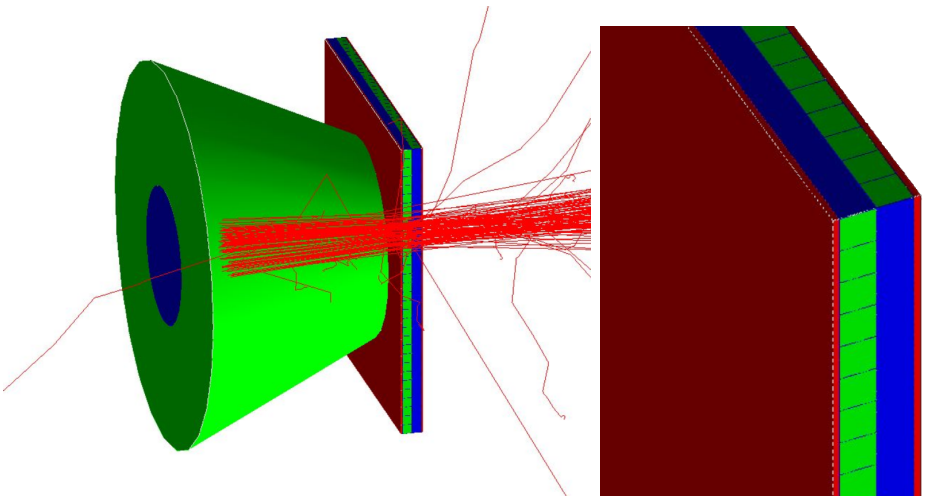
W represents the fraction of muons that interact with zero, one, or two fibres, and η is the fraction of the total detectable deposited energy for that interaction, compared to the single muon charge deposit in the two fibres, which is Q_μ. Therefore, $\eta_2 = 1$ and $\eta_0 = 0$.
Muons can interact with different materials in the beam monitor, but only if the interaction happens in the fibres the measurement can happen:

- if muons only interact with the PLA window, paint coating, or anywhere where the deposited energy is not detected (W_0);
- muons can also interact with a single front or back fibre (respectively W_{1f} and W_{1b}). This can happen if muons pass through the gaps between fibres in one layer and interact with the other, much more likely for an interspaced design such as the FAMU beam monitor;
- muons deposit energy in two fibres, different layers (W_2), resulting in a coincidence event (this is the optimal condition).

Table 4.2: Results for the beam monitor simulation

	W	Energy	η
2-hit	$(99.92 \pm 0.10_{\text{stat}})\%$	$(5.545 \pm 0.004_{\text{stat}})$ MeV	1
1-hit-front	$(0.061 \pm 0.002_{\text{stat}})\%$	$(2.34 \pm 0.04_{\text{stat}})$ MeV	$0.4220 \pm 0.008_{\text{stat}}$
1-hit-back	0	non def.	non def.
0-hit	$(0.0213 \pm 0.0015_{\text{stat}})\%$	0	0

The GEANT4 simulation was built only considering the active area of the beam monitor and the dead layers in front and in the back. Two 32-fibre planes were constructed using the ideal dimensions of the fibres, estimating the coating thickness to be 15 μm, and the front and back windows were estimated to be 0.5 mm thick. The beam, originating in the lead nose behind a mylar window that separates the beam vacuum from air, is a single muon particle gun with a 5% Gaussian energy spread. The geometry of the



(a) Geometry of the GEANT4 simulation with 100 simulated muons

(b) Close-up of the beam monitor fibre structure

Figure 4.12: Geometry used in the GEANT4 simulation

simulation can be seen in Figure 4.12.

1 million total events were simulated with a momentum of 49.5 MeV/c (as used in the runs), and it yielded the result summarised in Table 4.2. The result of the simulation confirm that most muons interact with two fibres, depositing ~ 5 MeV of energy in the layers. The hits in the front only are due to muons stopping in the first layer, while no muons hit only the second layer. The final figure for the geometric factor is:

$$k_g = 0.99946 \pm 0.00010_{\text{stat}} \simeq 1 \tag{4.5}$$

The instantaneous intensity can be simplified to be:

$$\varphi(t) = r \cdot \frac{Q(t)}{Q_\mu} \tag{4.6}$$

The geometric factor's role is quite important in interspaced fibres geometries, such as the FAMU beam monitor: in that case, the probability of hitting no fibres or a single fibre is not small, because the fibres are arranged with gaps, resulting in a much smaller geometric factor.

Single muon charge

The next step of the calculation is the extraction of the single muon charge deposit: this corresponds to the integrated current measured at each SiPM when a single muon hits both a front and a back fibre.
To obtain the data, measurements were taken in *low intensity condition*: the goal is to

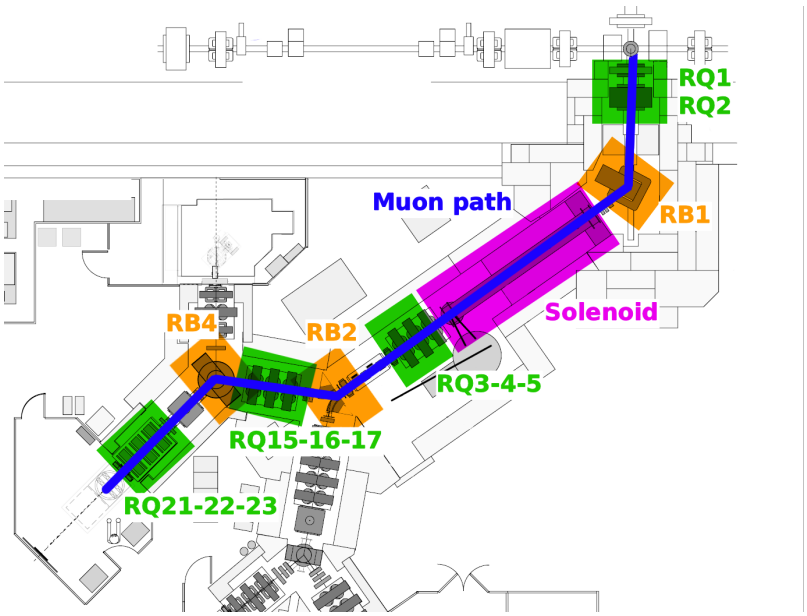
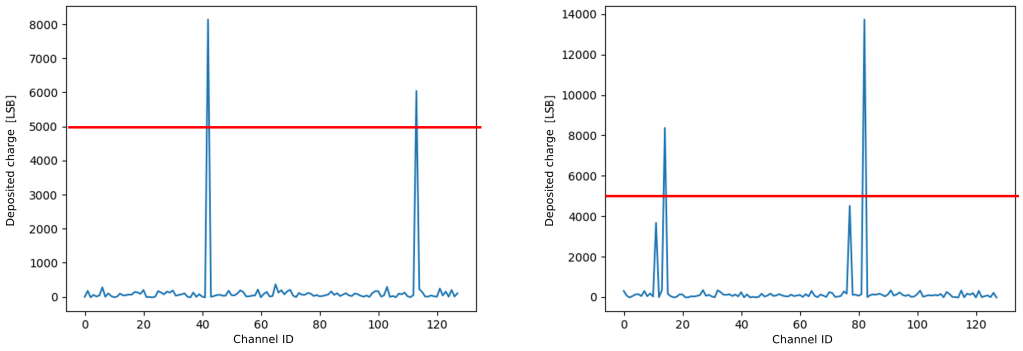


Figure 4.13: Schematic of the RIKEN-RAL facility highlighting the path of the muons to port 4 and the position of the magnets, adapted from [10]

keep the momentum the same but reducing the number of muons. Intensity reduction is generally achieved by de-focusing the initial pion beam by reducing proportionally to a factor *detune* the current of the two pion quadrupole magnets RQ1 and RQ2, highlighted in Figure 4.13. This increases the pion beam size, reducing the number of pions injected by bending magnet RB1 in the decay solenoid.
However, the conditions required by this activity are not yet met setting `detune` $= 0$, *i.e.* shutting down RQ1 and RQ2. For this reason, it is also necessary to change the current in the RB1 magnet: the magnet can be set to direct the centre of the beam out of the decay solenoid, transporting only the halo of the pion beam. This results in a much lower muon flux, while momentum selection is ensured by setting properly the decay solenoid and the next two bending magnets (RB2 and RB4) as in usual condition, to only transport the selected momentum. The muons transport path is highlighted in Figure 4.13.

By setting the current in the RB1 dipole magnet 30-50% larger or smaller than the nominal current, it is possible to obtain a low enough intensity to count single muons per spill in the beam monitor.
Four different, 5-minute runs were taken, at either extreme of the magnet settings: two of these are usable to collect the single charge data, run 223 and 226, respectively with a current setting of -265 A and -285 A (nominal is -190 A). Both provided a good amount of single coincidence events.

The selection of single coincidence event consists in an exclusive coincidence logic:



(a) Good event

(b) Noise and double hits

Figure 4.14: Examples of good and bad events for the same choice of threshold (5000 LSB)

we require a single hit (XOR) in each plane (AND), with the caveat that no hit must happen in the other time window. In a plot such as the one in Figure 4.9 we expect exactly two peaks, one in the X plane and one in the Y plane, separated by a number of channels between 32 and 64 to ensure the coincidence happening in the same spill.

The number of selected events depends on the choice of threshold for the identification of a hit: this is a difficult task, as two effects compound to give a very variable number of events and their charge.

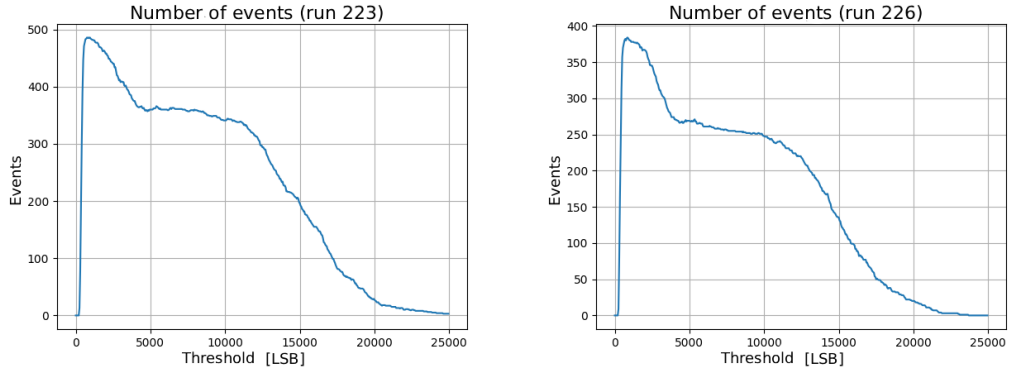
- a low threshold reduces the number of single coincidence events because the selection logic is exclusively a hit in only two channels, both in the same temporal window and with different planes: too low of a threshold and many events are excluded due to noise, but also several noise-only hits can be found, reducing the average charge deposit;
- increasing the threshold too much could include events such as in Figure 4.14b, where a double hit is not excluded. A high threshold also introduces a bias, where all selected events have a higher-than-average deposited charge.

A good selected event is represented in Figure 4.14a.

The solution was determined to be the use of a scan in threshold to determine a region of threshold where the number of events stays relatively stable, to be used for the determination of the error.
The threshold scan was performed every 50 LSB[3] and resulted in the plots in Figure 4.15, which show the number of exclusive single muon hits as a function of threhsold.

The number of selected events for the two runs show similar characteristics: the num-

[3]Least significant bit, equivalent to ADC channel



(a) Run 223 (b) Run 226

Figure 4.15: Number of single muon events as a function of threshold

ber starts at zero, where the threshold is too low and finds multiple peaks for each plot, then an initial rise due to the noise peaks being excluded with an increasing threshold. This reaches a maximum, with a subsequent decrease, a plateau and final decrease due to increasing threshold and exclusion of real, low energy hits.
The region we are interested in is right after the exclusion of the noise floor and before the high threshold region where the result would be biased. This is represented by the plateau.
To extract the region of the plateau, the easiest solution came from the shape of the derivative, calculated via the central derivative formula:

$$\frac{dN}{dx}(x_0) = \frac{N(x_0+h) - N(x_0-h)}{2h} \tag{4.7}$$

In this case, h is the threshold gap between two subsequent calculations of N, 50 LSB. The derivative of both runs, as in figures 4.16a and 4.16b, shows two depressions where events decrease, and a relatively flat region in the middle, the *plateau*.

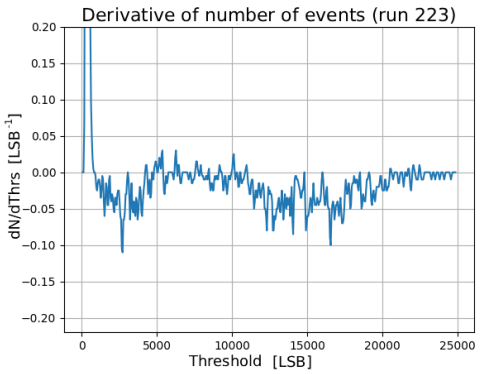
To extract the actual extension of the plateau, the two depressions were fitted using the Gaussian function, extracting the standard deviation (σ) and the position of the centres (X). To make the calculation easier, the negative of the derivative was used, and the results are in figures 4.16c and 4.16d. The plateau extension was chosen as:

$$[X_1 + 2\sigma_1, X_2 - 2\sigma_2] \tag{4.8}$$

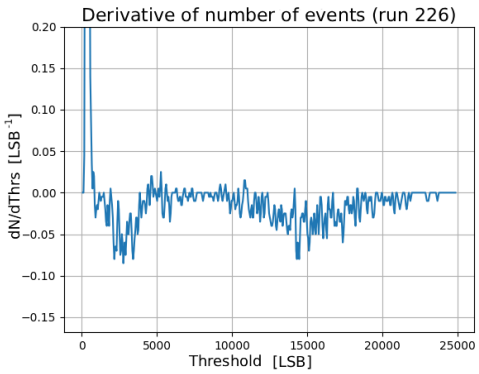
For the two runs, the extension of the plateaus was determined, resulting in:

$$\begin{cases} \text{Run 223} : [4900 \text{ LSB}, 8750 \text{ LSB}] \\ \text{Run 226} : [4350 \text{ LSB}, 9550 \text{ LSB}] \end{cases} \tag{4.9}$$

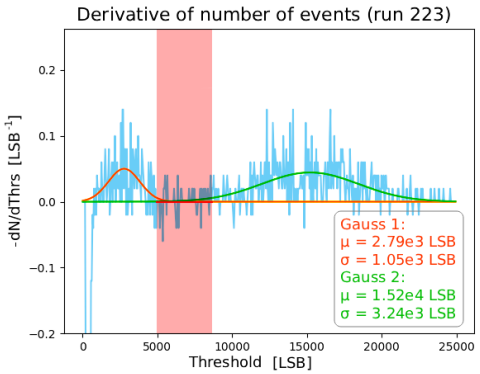
For each threshold in the selected range, the event extraction was performed: this consisted in the selection of single muon events via the single-window, double-layer logic. The charge of the entire selected event was then calculated, summing the charge for



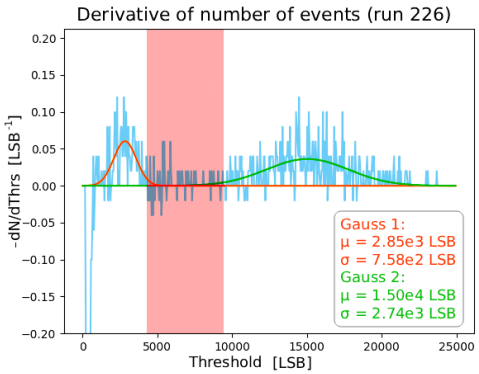
(a) Run 223



(b) Run 226



(c) Run 223 - fit



(d) Run 226 - fit

Figure 4.16: Derivative of the number of muons as a function of threshold, the bottom two plots show the fit to extract the plateau region

all channels and both spills. This resulted in a histogram of charges for each threshold, which could then be fitted to extract the maximum of the histogram for the determination of Q_μ.

For each threshold, histograms are constructed with a range of number of bins (from 40 to 200 with steps of 5). Each histogram is fitted using the `scipy.optimize.curve_fit` function in python, using as a model the Exponentially-Modified-Gaussian, which is the convolution of a Gaussian with the exponential function [48]:

$$F(x) = A\int_{-\infty}^{x} dt\ \ e^{-t/\tau}\ \ G_{\mu,\sigma}(x-t) = \\ = C + A\cdot\exp\left[-\frac{x-\mu}{\tau}+\frac{\sigma^2}{2\tau^2}\right]\left(1+\mathrm{erfc}\left(\frac{x-\mu-\sigma^2/\tau}{\sqrt{2}\sigma}\right)\right) \quad (4.10)$$

The fit was repeated on all thresholds in the ranges in Equation 4.9, and all binning (from

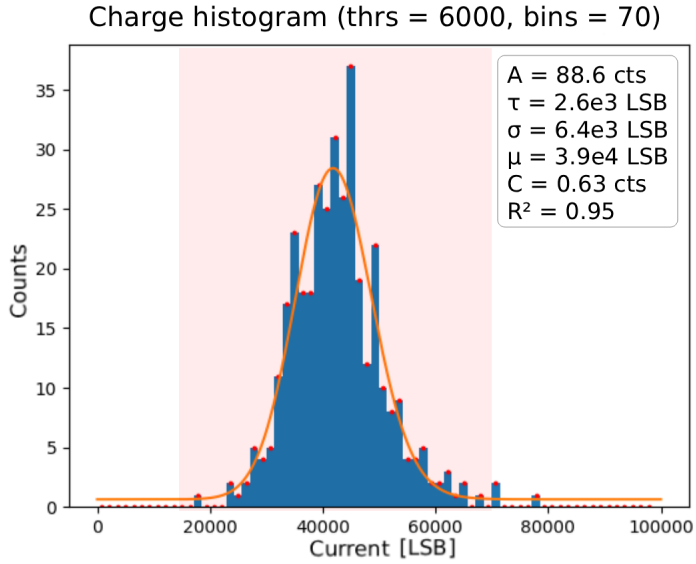


Figure 4.17: Fit using Equation 4.10 for one of the histograms used to extract Q_μ, the position of the maximum

40 to 200 bins per histogram), extracting numerically the position of the maximum, which corresponds to the charge deposited by a single muon coincidence event, as the EMG does not have an analytically-defined maximum. An example of a fit can be seen in Figure 4.17, with the red shaded area indicating the extension of the area where the function was fitted.

For each run, I obtained a set of values of Q_μ: the average was obtained by simple arithmetic mean, and the systematic error via the standard deviation.

$$\begin{cases} \text{Run 223}: \quad Q_\mu = 41750(1 \pm 0.896\%)\ \text{LSB} \\ \text{Run 226}: \quad Q_\mu = 41761(1 \pm 0.632\%)\ \text{LSB} \end{cases} \quad \Rightarrow \quad \overline{Q_\mu} = 41757(1 \pm 0.517\%)\ \text{LSB} \tag{4.11}$$

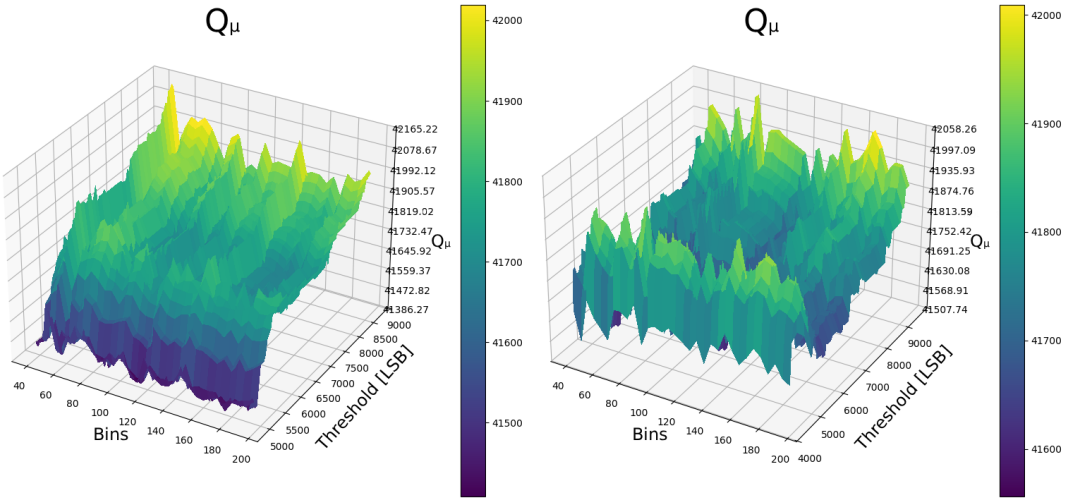
The distribution of Q_μ as a function of both threshold and bin number can be referenced in Figure 4.18.

Instantaneous intensity

The final result for $Q_\mu = (4.175 \pm 0.021) \times 10^4$ LSB, which determines the moment-dependent factor:

$$\varphi_\mu(t) = r \cdot \frac{1}{k_g \cdot Q_\mu} \cdot Q(t) = 9.584 \times 10^{-4} \cdot Q(t) \frac{N_\mu}{\text{s LSB}} \tag{4.12}$$

For the test run, run 185, the calculated muon intensity averaged over 5 hours is 1.45×10^3 muons/s, $\sim$ 18 muons per spill, which is much lower than the expected



(a) Run 223 (b) Run 226

Figure 4.18: Values of Q_μ as a function of threshold and bin number

intensity extracted from Figure 1.8, which would be 2.2×10^4 muons/s: this is partly because the measurements were made with reduced intensity to avoid saturating the HPGe detectors: the *detune* factor changes the settings of the magnets to reduce the intensity of the beam, but the reduction does not have a modelled dependence on the detune factor, and we cannot estimate it.

Another thing we can do is to look at the beam parameters to spot correlations between them and the current extracted form the beam monitor. In particular, the two most important parameters are the *synchrotron current* (number of primary protons per spill) and the *pulse width* (the time difference of arrival of the first and the last proton in a spill). The correlations for run 185 can be seen in Figure 4.19. This is done by using a different data file, which only records the total current: that part of the acquisition stopped later, at almost 7 hours, so the plot shows a longer time than was used in the average calculation. Only the first 5 hours were used, in the most uniform region, to extract the average intensity.

While the effect of a change in synchrotron current is easily interpreted (such as the dip in current at 1:45 AM, corresponding to stopping in the synchrotron), the effect of the pulse width is still debated. A correlation is noticeable, opposite to what was detected in the FAMU experiment.
The effects of the pulse width on the muon beam intensity are multiple:

- a decrease in pulse longitudinal width (time) means a more dense proton bunch, increasing the Coulomb repulsion between same charges: the horizontal profile of the bunch widens, reducing the number of protons hitting the target and the

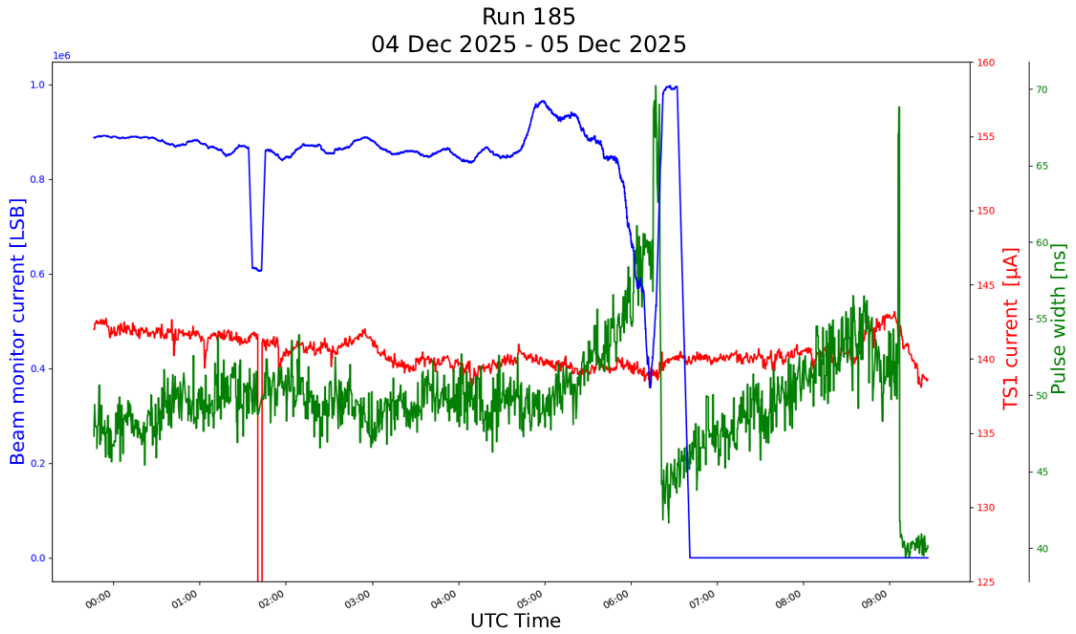


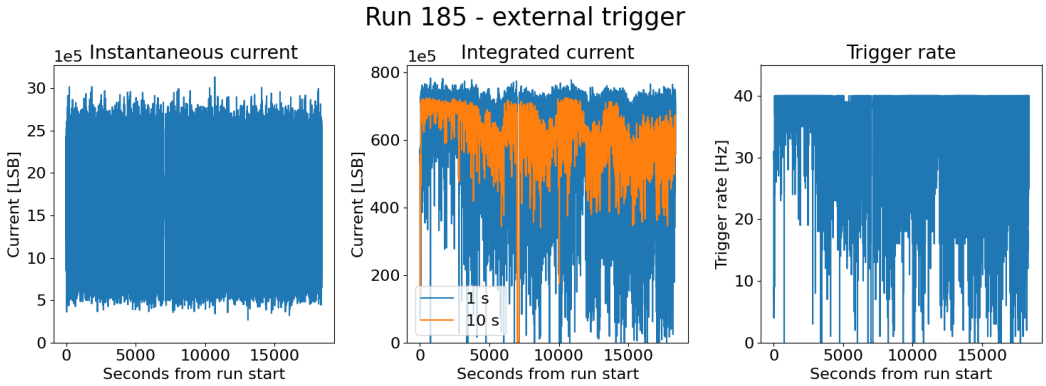
Figure 4.19: Beam monitor current (smoothed with a rolling average over 500 s) compared to the beam parameters

number of pions. This is a positive correlation, which was first noticed in the FAMU beam monitor;

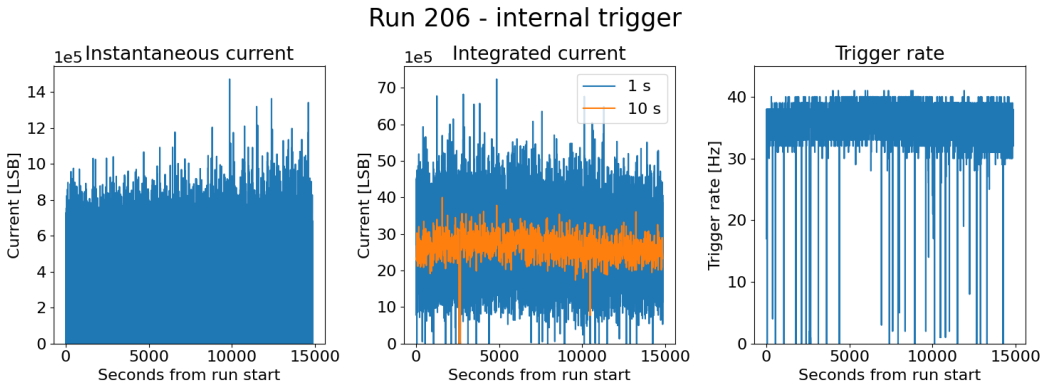
- an increase in pulse width means that the muon spill is more spread out in time: if the acquisition window is fixed (such as in CHNet-MAXI), we expect a reduction of detected muons. This is a negative correlation, in FAMU it is not relevant as it acquires data in a wide window that contains both bunches, while CHNet-MAXI integrates the data in narrow windows, separating the bunches.

4.2.4 Spectra normalisation

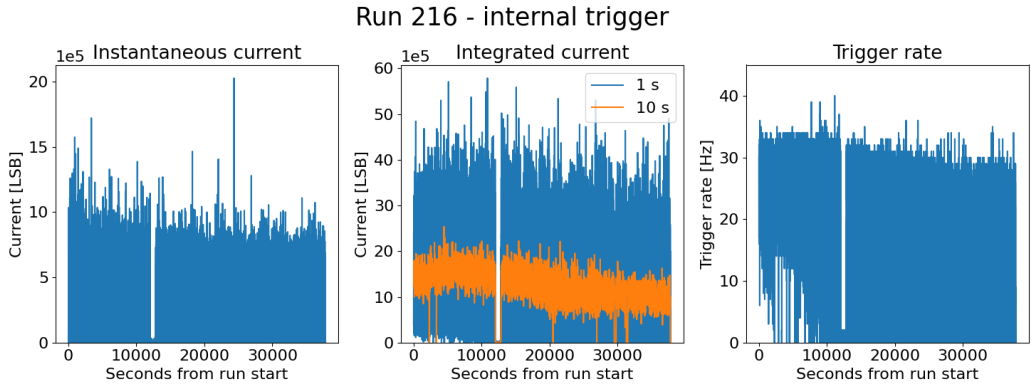
The main goal of the beam monitor was providing a normalisation factor for the X-ray spectra: knowing the number of muons that hit the target and generate X-rays removes the need of calculating the probability of capture, making the extraction of the ratios much easier.
Without a beam monitor, the spectra are normalised by counts: this ignores the difference in probability of capture for different isotopes, which must be taken into account when calculating ratios of the areas of the peaks. As discussed in Section 1.3.2, the formula used to derive the capture probability is empirical and cannot be easily applied to isotopes of the same element: the knowledge of the amount of muons that interacted with the material can be used to derive the probability of capture, and the emission of any of the characteristic X-rays can be calculated as a probability form the measurements of the standards without the use of any formula.



(a) Run 185 (test, external trigger)



(b) Run 206 (Pb-206, internal trigger)



(c) Run 216 (Pb-208, internal trigger)

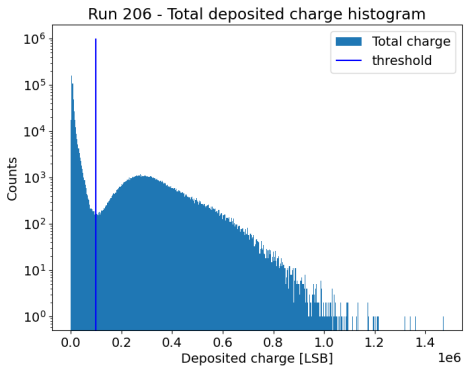
Figure 4.20: Plots showing the problems with the inconsistencies in the internal trigger acquisition

In normal operation, spectra normalisation is very easy, as it only consists of the sum of the charge deposited over the entire run in the beam monitor. If spectra must be compared to different beams, the absolute calibration can be used to derive the number of muons that hit the target.
Unfortunately, during setup, a value in the beam monitor data acquisition software was changed, switching from external trigger to coincidence mode. This resulted in all the experimental runs being acquired in non-optimised internal triggering, resulting in runs with many empty events and very variable rate of acquisition. The error is due to a variable in the setting files being changed, and it only involves the December 2025 data taking campaign: in all future uses of the same setup the error will be prevented by implementing strategies to highlight the triggering mode of acquisition and preventing accidental changes.
The main reason we cannot use the simple sum is that the acquisition rates of events (effectively the internal triggering rates) are much different in different runs: for example for run 206 (second run of Pb-206) the rate was 35.6 Hz but for run 216 (Pb-208) it was 17.5 Hz, less than half, as examples in Figure 4.20 show. Moreover, the amount of empty events changes between different runs, making the simple sum unreliable: a run with many empty events, even if the acquisition rate is close to the nominal 40 Hz, will underestimate the intensity of the beam.

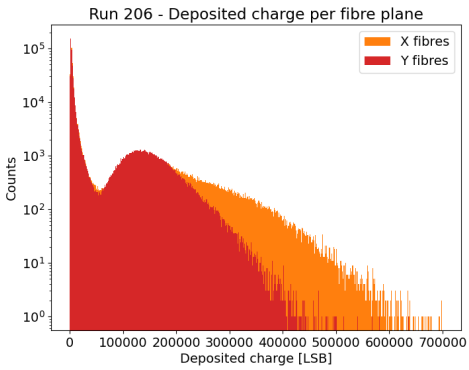
One of the important quantities that was recorded was the number of triggers acquired by the other detectors, which were running synchronously with the beam. This can be used, together with an estimation of the charge of non-empty events, to simulate a synchronous beam monitor by multiplying the number of triggers by the average charge deposited in the beam monitor (filtering only for non-empty events, or events with sum of the currents that is over a determined threshold).

The threshold, chosen at 10^5 LSB on the total charge deposited in the two planes, was determined by identifying the separation between the noise (empty events) and the real signals. The histogram of the deposited charge, in Figure 4.21, shows a very clear separation in logarithmic scale, making the choice of threshold easy. The position of the separation in the histograms of the charge per plane is at half the charge, as expected, but the histogram shows another effect: the second plane (X fibres) has a higher average charge deposited, due to the muons interacting with the first plane and losing energy, meaning that the deposit in the second layer is even greater due to the dependence of the Bethe Bloch to the inverse squared of the velocity (Equation 1.2).
This allowed the determination of the normalisation parameters with respect to the longest run, in Table 4.3. The normalisation can only be done with samples analysed at the same momentum, so the thicker archaeological coin cannot be compared to the standards.

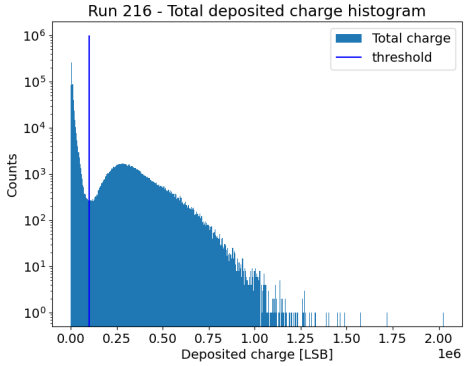
These results were used for the normalisation of the spectra: the same analysis is



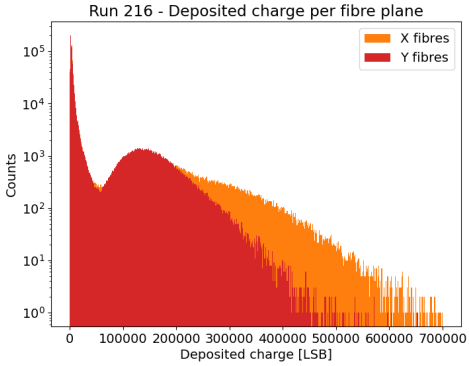
(a) Run 206 - total charge



(b) Run 206 - charge in fibres



(c) Run 216 - total charge



(d) Run 216 - charge in fibres

Figure 4.21: Histogram of the deposited charge. On the left, the total charge per event, with highlighted in blue the chosen threshold (10^5 LSB). On the right the charge deposited in the X and Y fibres, showing the effect of the Bethe Bloch: the second set of fibres encountered by the beam (X fibres) have a larger average energy, because they are closer to the end of range

Table 4.3: Summary of normalisation factors with respect to the archaeological coin and acquisition time

Sample	Measurement time (hr)	Normalisation factor wrt ArchPb
Pb-208	10.5	34.8%
Pb-206 (1)	3.5	12.6%
Pb-206 (2)	4	13.4%
Pb-206 (3)	3.5	11.4%
Pb-207	23.5	75.8%
Pb-204	23	75.9%
ArchPb (thin)	30.5	100%

also repeated by using the number of X-rays, to compare the two results.
More importantly, the discovery of the anomaly discussed in this Section prompted a complete redesign of the way variables are input in the DAQ software, reducing the risk of such problems arising in the future. The DAQ system was also changed to raise a warning flag in the case of different trigger modes being used for the acquisition of different detectors in the same run.

Conclusion

In conclusion, the work contained in this thesis is helping to develop upgrades to currently-used muon techniques for material analysis.
The work carried out with the Detector Systems Group at the ISIS Neutron and Muon Source is going to be used in the further development of the Super-MuSR spectrometer, which is in the final testing stages of the staves at the moment. I studied the response of the detector to different gain, to define a figure of merit for the optimisation of the Pulse Height Spectra as a function of gain. I also studied the effect of gain mismatch between the two SiPMs that read each scintillating tile out and analysed the wave shapes to ensure no systematic effects appeared in high magnetic fields. The final detector will be deployed next year during the ISIS long shutdown, improving the available rate capability, time resolution and magnetic field intensity compared to MuSR. The investigation I conducted of the anomalous pulses in the PMTs of HiFi and EMU instead identified them as ionic afterpulses due to helium infiltration, and will directly lead to changes in the data acquisition systems to reject the identified afterpulses, which affected the asymmetry calculations in the muon spectrometers. The Digital Muons upgrade project, aiming to improve the rate capability and resolution of the entire muon facility without changes to the detectors themselves through the use of a digital data pipeline for event processing, will enable the rejection of the afterpulses: a filter will be implemented into the data processing of PMT-based experiments acquisition, for now EMU and HiFi, thanks to the discovery of the shape parameters correlations that can easily be used to select them.
At the ISIS Muon Source, in RIKEN-RAL port 4, the advancements to μ-XES are led by the instrument scientists and the user community. The CHNet-MAXI experiment will add to the body of work, providing detailed methodology and a benchmark for isotopic analysis of lead. The work done on the beam monitor will be used to better normalise the X-ray spectra, to get a more precise evaluation of the isotopic ratios, and will be able to compare results of the same method in other facilities thanks to the absolute calibration. The calibration resulted in the extraction of the conversion factor to measure the instantaneous intensity:

$$\varphi_\mu(t) = 9.584 \times 10^{-4} \cdot Q(t) \frac{N_\mu}{\text{s LSB}}$$

This can be used to quantify the muons per second for any beam with momentum 49.5 MeV/c. The procedures for extraction of the geometric factor and the single muon response will need to be repeated for any experiment using a different beam momentum,

and the method here described could also be implemented if using the beam monitor in other applications or any other facility. The detailed study of the beam intensity also helped highlight the modulation to the muon beam intensity due to the proton pulse-width, first noticed during the beam monitor data analysis for the FAMU experiment.
I also developed a technique for the calculation of the normalisation parameters in internal coincidence triggering conditions, which will be used for the normalisation of the X-ray spectra for the extraction of the isotopic ratios despite problems in the acquisition method.

My master's thesis was spent working on different steps of the data taking, analysis and processing, and different moments of instruments' lives, from the prototype of Super-MuSR, to the first data taking run of the beam monitor and the effects of older PMTs on the HiFi and EMU spectrometers. This helped me gather experience in many areas of the experimental process, shaping a complete picture of a physics experiment, from inception to deployment and beyond.

Bibliography

[1] S. Blundell, R. De Renzi, T. Lancaster and F. L. Pratt, *Muon spectroscopy: an introduction* (Oxford University Press, 2022).

[2] F. Takahashi et al. (Particle Data Group), «Review of Particle Physics», Int. J. Mod. Phys. A **41**, 2630011 (2026) DOI: 10.1142/S0217751X26300115.

[3] S. D'Amico, V. Venuti et al., *Handbook of cultural heritage analysis* (Springer Cham, 2022), DOI: 10.1007/978-3-030-60016-7.

[4] M. Thomson, *Modern particle physics* (Cambridge University Press, 2013).

[5] G. F. Knoll, *Radiation detection and measurement* (John Wiley & Sons, 2010).

[6] A. Brown and H. Suit, «The centenary of the discovery of the Bragg peak», Radiotherapy and Oncology **73**, 265 (2004) DOI: 10.1016/j.radonc.2004.09.008.

[7] A. D. Hillier, D. M. Paul and K. Ishida, «Probing beneath the surface without a scratch—bulk non-destructive elemental analysis using negative muons», Microchemical Journal **125**, 203 (2016) DOI: 10.1016/j.microc.2015.11.031.

[8] G. Bonomi, P. Checchia, M. D'Errico, D. Pagano and G. Saracino, «Applications of cosmic-ray muons», Progress in Particle and Nuclear Physics **112**, 103768 (2020) DOI: 10.1016/j.ppnp.2020.103768.

[9] D. Findlay, J. Thomason, S. Fletcher, R. de Laune and E. Cooper, *A Practical Guide to the ISIS Neutron and Muon Source*, https://www.isis.stfc.ac.uk/wp-content/uploads/2025/12/A-Practical-Guide-to-the-ISIS-Neutron-and-Muon-Source.pdf, (accessed: 15.08.2026).

[10] A. D. Hillier, J. S. Lord, K. Ishida and C. Rogers, «Muons at ISIS», Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **377**, 20180064 (2018) DOI: 10.1098/rsta.2018.0064.

[11] C. Verona, G. Verona-Rinati, S. Platt, E. Schooneveld and A. Pietropaolo, «Comparison between silicon and single-crystal diamond photodiodes for localized pulsed fast neutron detection», Europhysics Letters **121**, 68002 (2018) DOI: 10.1209/0295-5075/121/68002.

[12] R. Rossini, «Exploring the proton structure with the FAMU experiment: detector performance and first physics results», (2025).

[13] T. Matsuzaki, K. Ishida, K. Nagamine, I. Watanabe, G. Eaton and W. Williams, «The RIKEN-RAL pulsed muon facility», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **465**, 365 (2001) DOI: 10.1016/S0168-9002(01)00694-5.

[14] M. Cataldo, M. Clemenza, K. Ishida and A. D. Hillier, «A novel non-destructive technique for cultural heritage: depth profiling and elemental analysis underneath the surface with negative muons», Applied Sciences **12**, 4237 (2022) DOI: 10.3390/app12094237.

[15] E. Quérel, S. Biswas, M. W. Heiss, L. Gerchow, Q. Wang, R. Asakura, G. Müller, D. Das, Z. Guguchia, F. Hotz et al., «Overcoming the probing-depth dilemma in spectroscopic analyses of batteries with muon-induced X-ray emission (MIXE)», Journal of Materials Chemistry A **13**, 2275 (2025) DOI: 10.1039/d4ta05112b.

[16] D. E. Groom, N. V. Mokhov and S. I. Striganov, «Muon stopping power and range tables 10 MeV-100 TeV», Atomic Data and Nuclear Data Tables **78**, 183 (2001) DOI: 10.1006/adnd.2001.0861.

[17] B. V. Hampshire, K. Butcher, K. Ishida, G. Green, D. M. Paul and A. D. Hillier, «Using negative muons as a probe for depth profiling silver roman coinage», Heritage **2**, 400 (2019) DOI: 10.3390/heritage2010028.

[18] G. A. Green, K. Ishida, B. V. Hampshire, K. Butcher, A. M. Pollard and A. D. Hillier, «Understanding Roman gold coinage inside out», Journal of Archaeological Science **134**, 105470 (2021) DOI: 10.1016/j.jas.2021.105470.

[19] K. Ninomiya, M. K. Kubo, P. Strasser, T. Nagatomo, Y. Kobayashi, K. Ishida, W. Higemoto, N. Kawamura, K. Shimomura, Y. Miyake et al., «Elemental analysis of bronze artifacts by muonic X-ray spectroscopy», in Proceedings of the 2nd international symposium on science at j-parc—unlocking the mysteries of life, matter and the universe— (2015), p. 033005, DOI: 10.7566/JPSCP.8.033005.

[20] M. Clemenza, G. Baldazzi, G. Ballerini, M. Bonesini, M. Carpinelli, O. Cremonesi, E. Di Stefano, E. Fiorini, G. Gorini, A. Hillier et al., «CHNET-TANDEM experiment: use of negative muons at RIKEN-RAL Port4 for elemental characterization of "Nuragic votive ship" samples», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **936**, 27 (2019) DOI: 10.1016/j.nima.2018.11.076.

[21] K. Terada, K. Ninomiya, T. Osawa, S. Tachibana, Y. Miyake, M. K. Kubo, N. Kawamura, W. Higemoto, A. Tsuchiyama, M. Ebihara et al., «A new x-ray fluorescence spectroscopy for extraterrestrial materials using a muon beam», Scientific reports **4**, 5072 (2014) DOI: 10.1038/srep05072.

[22] J. J. Reidy, R. L. Hutson, H. Daniel and K. Springer, «Use of muonic x-rays for nondestructive analysis of bulk samples for low z constituents», Analytical Chemistry **50**, 40 (1978) DOI: 10.1021/ac50023a015.

[23] K. Shimada-Takaura, K. Ninomiya, A. Sato, N. Ueda, M. Tampo, S. Takeshita, I. Umegaki, Y. Miyake and K. Takahashi, «A novel challenge of nondestructive analysis on OGATA Koan's sealed medicine by muonic X-ray analysis», Journal of natural medicines **75**, 532 (2021) DOI: 10.1007/s11418-021-01487-0.

[24] A. Rubbia, *Phenomenology of particle physics* (Cambridge University Press, 2022).

[25] A. Adamczak, D. Bakalov, G. Baldazzi, M. Baruzzo, R. Benocci, R. Bertoni, M. Bonesini, S. Capra, D. Cirrincione, M. Clemenza et al., «First operation of the FAMU experiment at the RIKEN-RAL high intensity muon beam facility», The European Physical Journal A **61**, 284 (2025) DOI: 10.1140/epja/s10050-025-01758-4.

[26] W. R. Leo, *Techniques for nuclear and particle physics experiments: a how-to approach* (Springer Science & Business Media, 1994).

[27] *PHOTOMULTIPLIER TUBES AND ASSEMBLIES*, https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/High_energy_PMT_TPMZ0003E.pdf, (accessed: 31.08.2026).

[28] C. Piemonte and A. Gola, «Overview on the main parameters and technology of modern Silicon Photomultipliers», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **926**, 2 (2019) DOI: 10.1016/j.nima.2018.11.119.

[29] K. S. A. Ghassemi and K. Kobayashi, *MPPC: Technical guide to silicon photomultipliers (MPPC)*, https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/mppc_kapd9005e.pdf, (accessed: 07.09.2026).

[30] S. Gundacker and A. Heering, «The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector», Physics in Medicine & Biology **65**, 17TR01 (2020) DOI: 10.1088/1361-6560/ab7b2d.

[31] P. J. King, R. De Renzi, S. P. Cottrell, A. D. Hillier and S. F. Cox, «ISIS muons for materials and molecular science studies», Physica Scripta **88**, 068502 (2013) DOI: 10.1088/0031-8949/88/06/068502.

[32] A. Hillier, J. Lord, S. Cottrell, P. Baker, A. Louat, R. Stewart, K. Yokoyama, J. Wilkinson, S. Biswas, D. Pooley et al., «Pulsed muons at ISIS-40 years and counting», in Journal of physics: conference series, Vol. 3222, 1 (IOP Publishing, 2026), p. 012018, DOI: 10.1088/1742-6596/3222/1/012018.

[33] R. da Silva Afonso, P. Baker, J. Lord and A. Hillier, «Monte Carlo simulations of the MuSR spectrometer at ISIS: current instrument and future designs», STFC ePubs website (2001).

[34] P. Baker, J. Lord, P. Biswas and A. Hillier, «Super-MuSR scientific design: Progress towards a step-change in muon capabilities at ISIS», in Journal of physics: conference series, Vol. 2462, 1 (IOP Publishing, 2023), p. 012032, DOI: 10.1088/1742-6596/2462/1/012032.

[35] S. Giblin, S. Cottrell, P. King, S. Tomlinson, S. Jago, L. Randall, M. Roberts, J. Norris, S. Howarth, Q. Mutamba et al., «Optimising a muon spectrometer for measurements at the isis pulsed muon source», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **751**, 70 (2014) DOI: 10.1016/j.nima.2014.03.010.

[36] Z. Salman, P. Baker, S. Blundell, S. Cottrell, S. Giblin, A. Hillier, B. Holsman, P. King, T. Lancaster, J. Lord et al., «HiFi—A new high field muon spectrometer at isis», Physica B: Condensed Matter **404**, 978 (2009) DOI: 10.1016/j.physb.2008.11.213.

[37] J. Lord, I. McKenzie, P. Baker, S. Blundell, S. Cottrell, S. Giblin, J. Good, A. Hillier, B. Holsman, P. King et al., «Design and commissioning of a high magnetic field muon spin relaxation spectrometer at the ISIS pulsed neutron and muon source», Review of Scientific Instruments **82** (2011) DOI: 10.1063/1.3608114.

[38] K. Yokoyama, J. Lord, P. Murahari, K. Wang, D. Dunstan, S. Waller, D. McPhail, A. Hillier, J. Henson, M. Harper et al., «The new high field photoexcitation muon spectrometer at the isis pulsed neutron and muon source», Review of Scientific Instruments **87** (2016) DOI: 10.1063/1.4972827.

[39] D. Pooley, A. Abba, F. Akeroyd, P. Baker, F. Caponio, Z. Clark, P. Coode, S. Cottrell, A. Cusimano, J. Day et al., «Development of Super-MuSR's 960 Channel SiPM Detector Barrel», in 2025 ieee nuclear science symposium (nss), medical imaging conference (mic) and room temperature semiconductor detector conference (rtsd) (IEEE, 2025), pp. 1–1, DOI: 10.1109/NSS/MIC/RTSD57106.2025.11287654.

[40] NuclearInstruments, *DAQ121*, https://www.nuclearinstruments.eu/products/daq121/, (accessed: 31.08.2026).

[41] *MPPC S13360-1325PE*, https://www.hamamatsu.com/us/en/product/optical-sensors/mppc/mppc_mppc-array/S13360-1325PE.html, (accessed: 29.08.2026).

[42] D. Pooley, A. Abba, F. Akeroyd, P. Baker, A. Boston, A. Brooks, S. Cottrell, F. Caponio, C. Evans, S. Franklin et al., «Development of a digitising data acquisition pipeline for next generation pulsed muon spectrometers», in Journal of physics: conference series, Vol. 2462, 1 (IOP Publishing, 2023), p. 012028, DOI: 10.1088/1742-6596/2462/1/012028.

[43] *Datasheet R5505 Datasheet (HAMAMATSU)*, https://manualmachine.com/datasheet/r5505/4287462-datasheet-hamamatsu/, (accessed: 28.08.2026).

[44] J. Incandela, S. Ahlen, J. Beatty, A. Ciocio, M. Felcini, D. Ficenec, E. Hazen, D. Levin, A. Marin, J. Stone et al., «The performance of photomultipliers exposed to helium», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **269**, 237 (1988) DOI: 10.1016/0168-9002(88)90885-6.

[45] MantidProject, *BackToBackExponential*, https://docs.mantidproject.org/nightly/fitting/fitfunctions/BackToBackExponential.html, (accessed: 28.08.2026).

[46] J. A. Evans, V. Pashley, K. Mee, D. Wagner, M. Parker Pearson, D. Fremondeau, U. Albarella and R. Madgwick, «Applying lead (pb) isotopes to explore mobility in humans and animals», PLoS One **17**, e0274831 (2022) DOI: 10.1371/journal.pone.0274831.

[47] CAEN, *R5560SE 128 Channel 14 bit 125 MS/s Open FPGA Digitizer*, https://caen.it/products/r5560se/, (accessed: 30.08.2026).

[48] R. Rossini, G. Baldazzi, S. Banfi, M. Baruzzo, R. Benocci, R. Bertoni, M. Bonesini, S. Carsi, D. Cirrincione, M. Clemenza et al., «The muon beam monitor for the FAMU experiment: design, simulation, test, and operation», Frontiers in Detector Science and Technology **2**, 1438902 (2024) DOI: 10.3389/fdest.2024.1438902.

[49] R. Rossini, M. Bonesini, M. Cataldo, G. Ceruti, A. De Bari, C. De Vecchi, L. Malliolio, A. Menegolli, P. Monza, P. Oliva et al., «A muon beam monitor for the CHNet-MAXI experiment at the RIKEN-RAL laboratory», Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1069**, 169824 (2024) DOI: 10.1016/j.nima.2024.169824.