Politecnico di Torino

Dipartimento di Ingegneria Meccanica e Aerospaziale

Design and characterisation of the CsI Crystal Calorimeter, Cosmic Ray Tagger and SiPM front-end electronics for the Mu2e experiment at Fermilab

Relatori
Prof. Eliodoro Chiavazzo
Dott.ssa Raffaella Donghia
Prof. Matteo Martini
Prof. Fabio Subba

Candidato
Daniele Paesani

Corso di Laurea Magistrale in Ingegneria Meccanica

Anno Accademico 2019-2020

♪

...and all this science I don't understand

it's just my job five days a week...

♪

Introduction

The Mu2e experiment [1] at Fermi National Accelerator Laboratory will investigate Charged Lepton Flavour Violation by searching for coherent muon to electron conversion in the Coulomb field of an Al nucleus. Although allowed due to neutrino oscillations, this conversion is suppressed by the Standard Model with an expected BR $< 10^{-54}$. Therefore, the observation of this process would be a clear evidence of New Physics beyond the Standard Model. Over three years of run time, more than $6 \cdot 10^{17}$ muons will be stopped on the Mu2e target, allowing a single event sensitivity of approximately $2.5 \cdot 10^{-17}$, probing 4 orders of magnitude beyond the current best experimental limit of $7 \cdot 10^{-13}$ @ 90% CL set by SINDRUM II.

Mu2e will produce a high intensity pulsed negative muon beam, generated via the interaction of 8 GeV proton bunches on a W target. By means of a dedicated 28 metres long Nb-Ti superconducting solenoid system, low-momentum negative muons are selected and transported to the Detector Solenoid, where they are stopped on an Al target and undergo nuclear capture. The 104.96 MeV signature of monoenergetic conversion electrons will be identified by a complementary measurement carried out by a high-resolution straw-tube tracker and the electromagnetic calorimeter (EMC).

The calorimeter has a high granularity, high energy resolution ($\sigma_E/E < 10\%$) and fast timing ($\sigma_t < 500$ ps). It is composed of 1348 undoped CsI crystals, each read by two custom UV-extended SiPMs, arranged in two annular disks. The EMC will need to maintain extremely high levels of stability and in the harsh Mu2e operating environment, for radiation exposure up to 15 krad/y and for neutron fluency up to $10^{11}\mathrm{n}_{1\mathrm{MeV}}/\mathrm{cm}^2$. The calorimeter design has been validated through an electron beam test on a large-scale 51-crystals prototype at Laboratori Nazionali di Frascati (LNF) Beam Test Facility (BTF). An extensive test campaign has been carried out in order to characterise and verify the performance of crystals, photodetectors and of the front-end electronics, including hardware stress tests and irradiation campaigns using both neutrons and gamma.

A Cosmic Ray Tagger (CRT) system has also been developed at LNF, to allow each disk to be individually calibrated using cosmic rays prior to installation. The CRT system will provide 3D MIP track reconstruction capabilities, in order to allow the equalisation and calibration of the energy response of all channels to a level below 1% against the expected 21 MeV MIP peak deposit. The CRT consists of two sub-modules, featuring a single layer of 8 parallel 1600 mm long bars made of EJ-200 plastic scintillator with a dual-sided Mu2e-SiPM readout. The scintillating elements and the relative readout system have been characterised at the LNF laboratories, showing a timing resolution better than 100 ps. A template fit algorithm was used for timing reconstruction. Time-of-flight difference was used for the reconstruction of the interaction point, with a spatial resolution relative to the axial coordinate better than 6 mm.

The final verification steps of the front-end system, including the complete characterisation of the optical and electronic performance of the hardware are reported, with focus on the most recent irradiation tests carried out with an intense ^{60}Co γ source at ENEA Calliope facility. Furthermore, the work relative to the design and characterisation of the CRT system is included, starting from the first R&D stages, up to the characterisation steps of the instrument with cosmic rays. The work is organised as follows.

Chapter 1 contains an introduction about the New Physics implications of the Mu2e experiment, along with an abridged description of the muon-to-electron coherent conversion process. The analysis of the expected signals and backgrounds are also included.

Chapter 2 contains the description of the Mu2e experimental setup, the proton and muon beam systems. The description of the Production, Transport and Detector Solenoid is included, along with the CRV and the peripheral systems.

Chapter 3 describes in detail the requirements and the design of the Mu2e crystal calorimeter, along with all its peripheral and calibration subsystems. Monte Carlo studies of the calorimeter performance are reported, along with the characterisation steps carried out on a large-scale prototype, including cosmic ray studies and test beams. The relative data analysis algorithms are discussed in detail.

Chapter 4 contains a detailed description of the Mu2e SiPM design and of their characterisation, along with a complete description of the front-end electronic design and of the signal chain, from SiPMs to digitisation.

Chapter 5 reports the radiation hardness characterisation of the calorimeter hardware. It contains an extensive report on the irradiation studies of crystals, SiPMs and electronics with neutrons and ionising dose.

Chapter 6 contains a detailed discussion of the R&D stages relative to the design and characterisation of the Cosmic Ray Tagger system.

1 Mu2e and CLFV

1.1 CLFV

The Standard Model (SM) is currently the best theoretical description of our understanding of Particle Physics. It is a quantum field theory that describes matter's basic constituents and the interactions among them [2]. Fermions (leptons and quarks) constitute matter's building blocks and are organised in three flavour generations. Transitions between generations have been observed both in the quark sector and in the neutral lepton sector, but not for charged leptons. The SM is based on the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge group, where C, L and Y indicate respectively the colour charge, the left-handed fields (which are the only ones subject to weak interactions) and the hypercharge. Even though it is currently the best model describing the subatomic world and no obvious contradictions have been observed in its application, the SM is not a complete theory of the universe, it incorporates only three of the four fundamental forces. In fact, the SM is currently lacking for an explanation of gravity, flavour structure, neutrino masses, dark matter/energy and the observed baryon asymmetry. Thus, there is evidence of Physics Beyond the Standard Model (BSM) to be discovered.

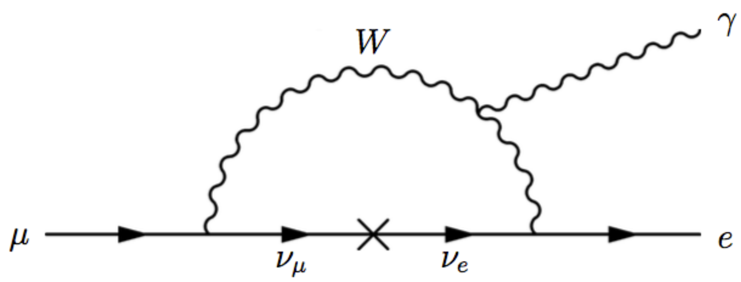

Figure 1.1 Feynman diagram of the CLFV violating process $\mu \to e\gamma$ in the SM.

The conservation of lepton flavour is accidental in the SM, as it is not related to the gauge structure of the model, but it arises from its particle content and in particular from the absence of right-handed neutrinos. Nevertheless, the experimental measurements of the neutrino mixing parameters and the discovery of their oscillations demonstrated the lepton family numbers are not conserved

[4][5]. As a consequence, an extension of the SM to include neutrino mass terms was necessary. In principle, the neutrino Pontecorvo-Maki-Nagasaki-Sakata (U_{li}) mixing matrix [6], can give rise to CLFV (Charged Lepton Flavour Violation). This violation can only occur through loop diagrams involving neutrinos and W bosons. For example, Figure 1.1 shows a Feynman diagram relative to the $\mu \to e\gamma$ decay in the SM. This process is strongly suppressed in the SM, as its the branching ratio (BR) is given by the following relation [7]:

$$BR(\mu \to e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_i U^*_{\mu i} \, U_{ei} \, \frac{m^2_{\nu i}}{M^2_W} \right| < 10^{-54} \tag{1.1}$$

where α is the electromagnetic coupling constant, U is the PMNS matrix, $m_{\nu i}$ are the neutrino masses and M_W is the W boson mass. On the other hand, many New Physics (NP) models predict significant enhancements to CLFV rates [3].

1.2 Negative muon conversion

The $\mu^- \to e^-$ conversion in a nucleus field represent a very attractive target in the context of the CLFV experimental searches in the muon sector: it is a coherent process with a signature given by a mono-energetic electron which recoils off the nucleus (according to a two-body interaction process). Its signature is associated to a number of appealing experimental advantages, including the quasi absence of accidental hits and the suppression of background electrons near the conversion energy from muon decays.

From the experimental point of view, the $\mu^- \to e^-$ signal corresponds to a mono-energetic peak to be separated from the fast falling spectrum of the decay-in-orbit events (cf. 1.3.1.1). The $\mu^- \to e^-$ process is also well suited for exploiting the new high intensity accelerators. In fact, this conversion is in principle non-rate limited, as the associated backgrounds scale at most linearly with the beam intensity. In contrast, accidentals related background (of great relevance, for example, in the MEG experiment) has a quadratic dependence on the event rate [8]. Furthermore, because the virtual particle exchanged with the nucleus can be either a photon or a heavy neutral boson, this reaction is sensitive to a broad swath of Beyond Standard Model Physics parameter space [3].

The best limit in the search of neutrino-less conversions is currently set by the SINDRUM-II experiment [9]:

$$BR(\mu \; + Au \rightarrow \; e^- + \; Au) < 7 \cdot 10^{-13} \text{ @ } 90\% \; CL \tag{1.2}$$

The Mu2e experiment aims to probe four orders of magnitude beyond this limit, by using one of the most intense muon beams ever created (cf. Chapter 2).

1.2.1 Muon capture

Mu2e will search for coherent muon-to-electron conversion in the Coulomb field of an Al nucleus. Muons fired at a thin aluminium target are stopped in the material and captured by the nucleus coulomb field to form exotic atoms, in which they replace electrons. The thus captured muons rapidly fall down to the *1s* state, through a cascade process associated with the emission of X-rays [10].

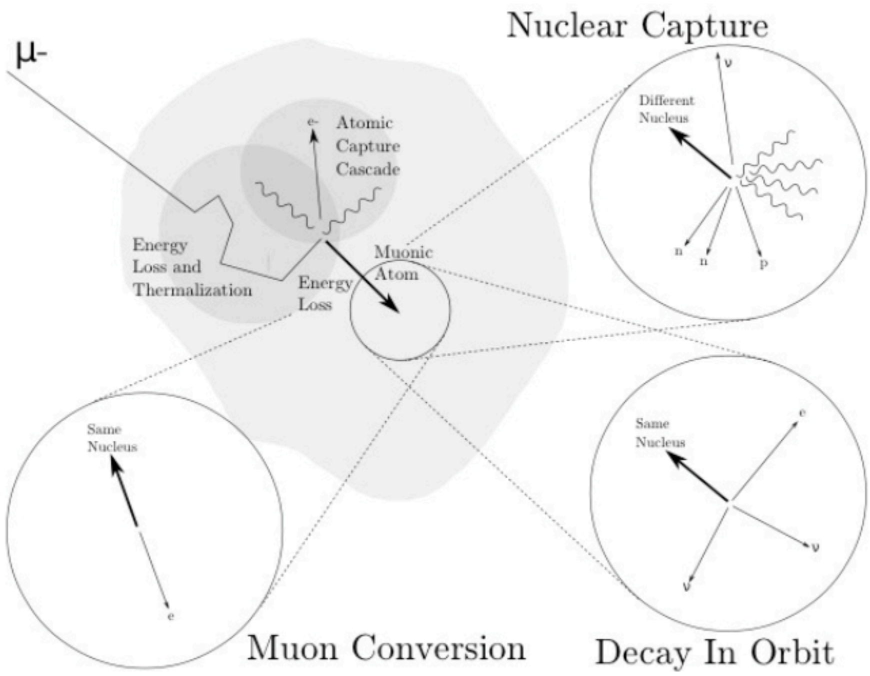

Figure 1.2 Main decay paths of a captured muon.

When a muon reaches the *1s* state, it can either decay or be captured on a bound proton. The muonic orbit will be much closer to the nucleus in comparison to an electronic one: in the *1s* state the Bohr radius will be comparable to the nuclear charge distribution radius, and, for heavier elements, muonic orbits can even be internal to the nucleus. In these conditions, a probability exists that the muon can undergo a nuclear capture process, described by the following semi-leptonic reaction (whose likelihood increases with Z):

$$\mu^- + N(A,Z) \rightarrow \nu_\mu + N(A, Z-1)^* \tag{1.3}$$

Muonic capture phenomena show a mean lifetime that has been measured in various materials and ranges from less than 100 ns (for high-Z nuclei) to over 2 μs (for low-Z nuclei). Another important mode is represented by the radiative capture (which is 10^4 times less likely if compared to ordinary capture), in which the energy released during capture is emitted via a photon:

$$\mu^- + p \rightarrow \nu_\mu + n + \gamma \tag{1.4}$$

After nuclear capture and once in the ground state, muons can undergo one of the following processes (Figure 1.2):

1. decay in orbit (DIO) via $\mu^- \rightarrow e^- + \nu_\mu + \overline{\nu_e}$ with a 39% BR;
2. nuclear capture with a 61% BR;
3. interaction with the whole nucleus through coherent conversion.

1.2.2 **Signature**

In the coherent $\mu^- \rightarrow e^-$ conversion the electron recoils off the entire Al nucleus, and thus the kinematic of a two-body decay process can be applied. As the nuclear mass is large if compared to the electron's one, the recoil terms are small and the conversion electron has a mono-energetic signature and carries most of the muon rest mass:

$$E_e = m_\mu - B_\mu - E_{rec} \approx m_\mu - B_\mu = 104.96 \text{ MeV} \tag{1.5}$$

where E_{rec} is the nuclear recoil energy, $m_\mu =$ 105.6 MeV and $B_\mu \approx Z^2\alpha^2 m_\mu$ is the muon binding energy for an atom with atomic number Z. For a recoiling nucleolus of mass m_N, nuclear-recoil energy can be approximated as $E_{rec} \approx$

$\frac{(m_\mu - B)^2}{2m_N}$. Since B_μ is different for various nuclei, the peak energy of the conversion electron signal changes accordingly.

In the Mu2e experiment, muons are stopped in a thin Al target and will cascade to the *1s* state in ~ 10^{-10} seconds. The associated muonic lifetime in aluminium is $\tau_{Al}^{1s} = 864$ ns. The coherent conversion in Al has an associated energy of 104.96 MeV.

1.2.3 Conversion rate and SES

Mu2e will measure the conversion rate, which is defined and currently estimated as the following ratio:

$$R_{\mu e} = \frac{\mu^- + N(A,Z) \rightarrow e^- + N(A,Z)}{\mu^- + N(A,Z) \rightarrow \nu_\mu + N(A,Z-1)} < 8.4 \cdot 10^{-17} \tag{1.6}$$

This quantity represents the number of coherent conversion events in the field of a nucleus, normalised to the total number of muonic captures (Figure 1.3). This normalisation scheme is used in order to eliminate various details of the nuclear wave-function.

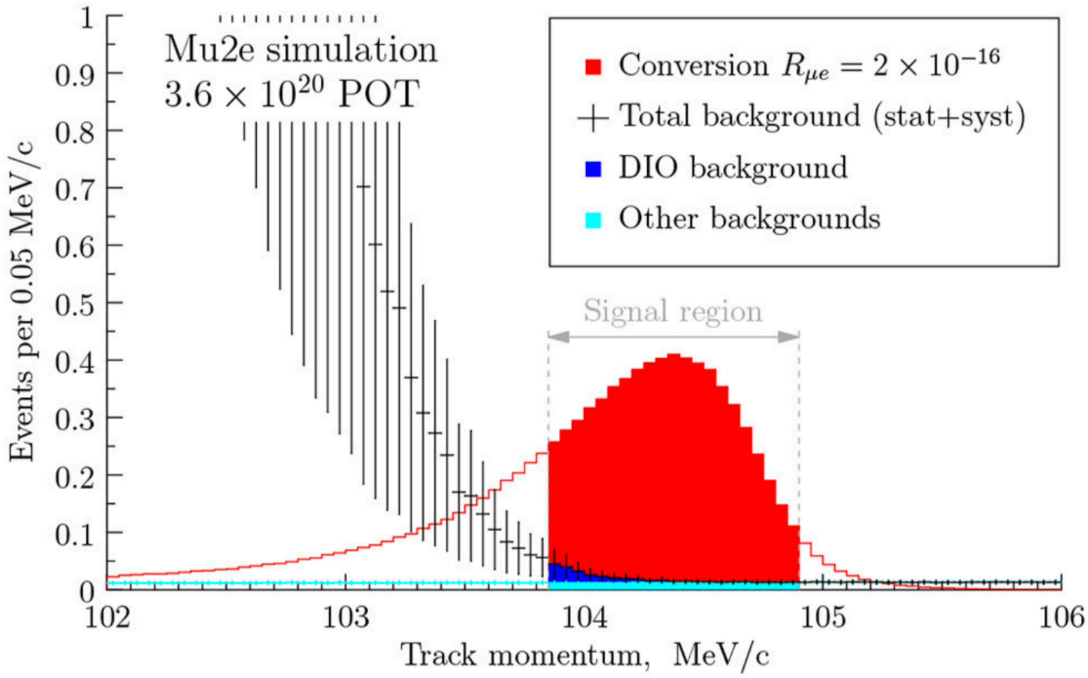

Figure 1.3 Mu2e discovery sensitivity for $R_{\mu e} = 2 \cdot 10^{-16}$.

The Single Event Sensitivity (SES) parameter is instead defined as the conversion rate associated to a number of measured CE events equal to one. The Mu2e experiment aims to reach a SES of ~2.5 $\cdot 10^{-17}$. For this target value, the background events rate needs to be smaller suppressed below a level of 0.5.

If the $\mu^- \rightarrow e^-$ process will not be observed, Mu2e will set an upper limit for the conversion rate of $R_{\mu e} < 8.4 \cdot 10^{-17}$ (four order of magnitude above the current world's best limit set by the SINDRUM-II experiment [9]). To provide this $R_{\mu e}$ limit with a 90% confidence level, approximately 10^{18} stopped muons are needed over 3 years of runtime.

1.2.4 Stopping target material

Given the Z-dependence of the CLFV operator, the choice of Al results from a trade-off between a number of conflicting requirements and competing factors. In fact, the number of muon conversions scales as Z^5 (the interaction scales with Z^2 and the probability of the wave-functions overlapping as Z^3). The BR is normalised to the rate of muon capture which is proportional to Z^4. As a result, the muon conversion rate scales linearly with Z (it peaks for tungsten and antimony, and then drops for higher Z values [11], as shown in Figure 1.4), thus suggesting the choice of a high-Z material.

Conversely, the long muonic atom lifetime of low-Z materials is preferred in order to allow data collection in a delayed acquisition window (cf. 2.1.1). Indeed, the muon lifetime in the nucleus field should ideally be greater than the transit time of prompt backgrounds by a few ns, as well as being compatible with the pulsed proton beam period. The muonic Al lifetime $\tau_{Al}^{1s} = 864$ ns is indeed approximately one half of the proton pulse period used in the experiment (cf. 2.1.1).

Another element supporting the choice of a low Z target relies in its less severe interaction with the outgoing electron, whose energy trace should be corrupted as little as possible. Moreover, high chemical purity and stability are needed in the material so that signals from muons stopped in the impurities can be neglected.

In the event of a successful identification of a CE signal, it will be possible to investigate the underlying New Physics mechanisms by comparing the rates relative to different material choices for the target.

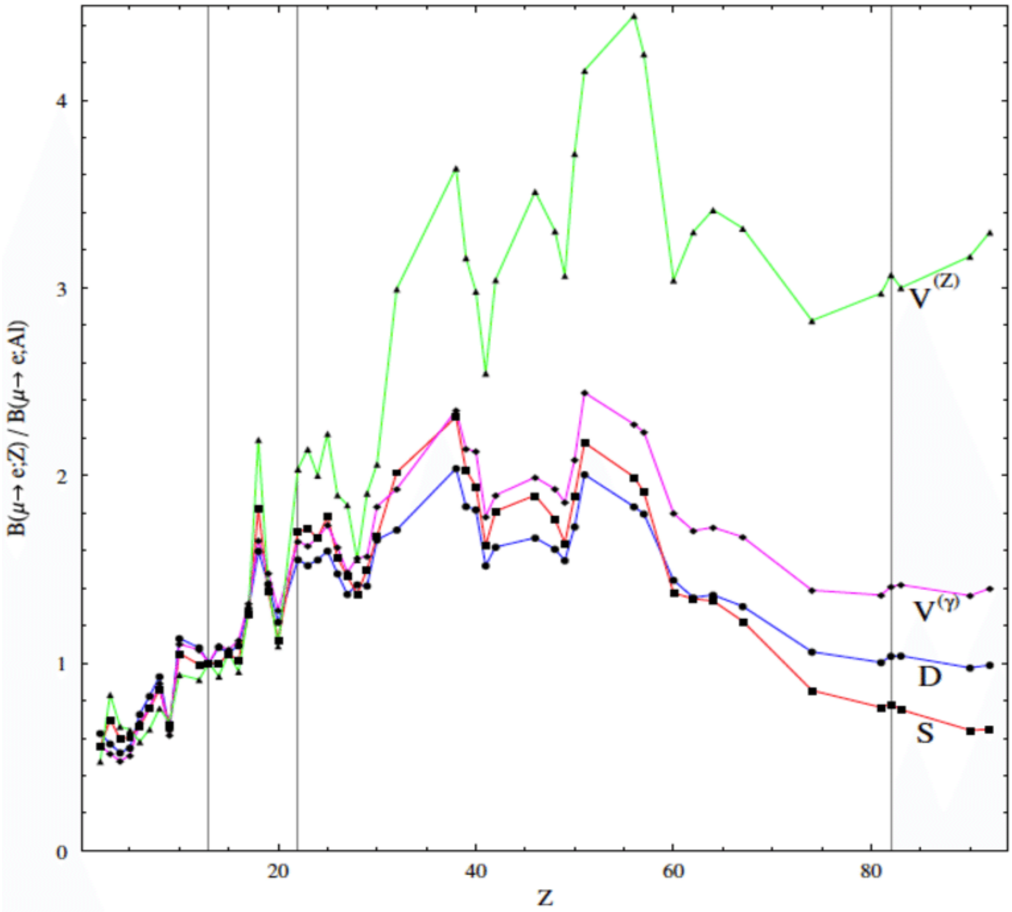

Figure 1.4 Conversion rate as a function of the stopping target material Z and for several NP operators [11].

1.3 Physics background

In the Mu2e experiment, several types of events can mimic the conversion electron signal. These processes can be classified in the following categories.

1.3.1 Intrinsic background processes

Intrinsic processes which scale with the beam intensity are muon decay in orbit (DIO) and radiative muon capture (RMC). A high-resolution detector is needed in order to mitigate the effects of this type of background, which cannot be reduced by the delayed acquisition gate scheme.

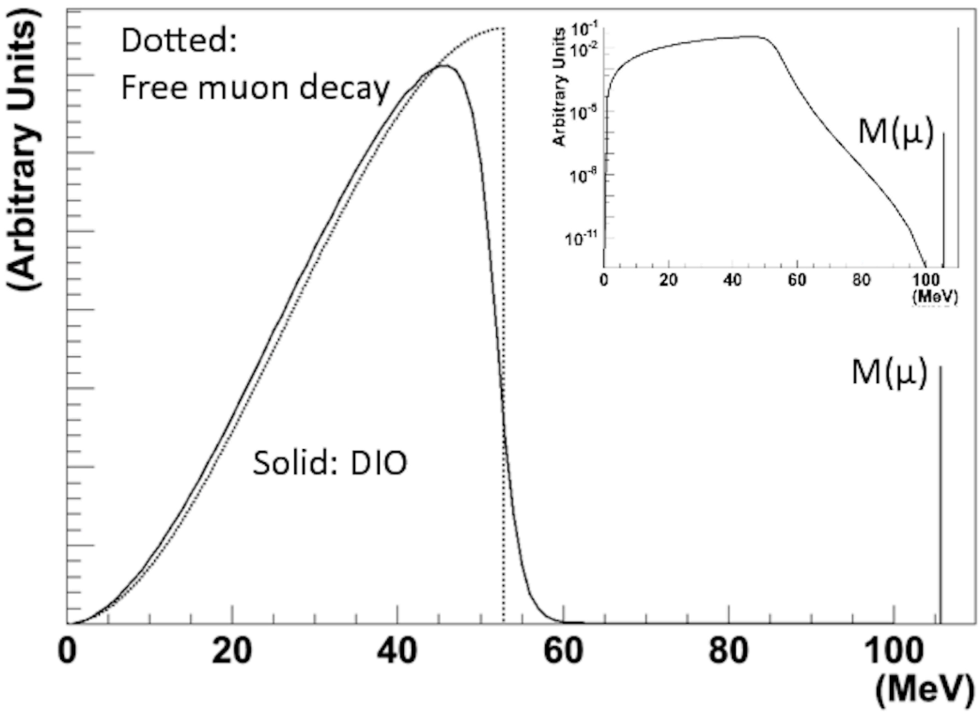

Figure 1.5 Michel and DIO spectra.

1.3.1.1 Decay in orbit

Electrons from muon DIO are one of the biggest contributions to the intrinsic conversion electron background. The dominant muon decay mode (for a free muon) is represented by the Michel decay ($\mu^- \to\ e^- + \nu_\mu +\ \overline{\nu_e}$). In this case, the electron presents a kinematic endpoint - associated to the electron and the two neutrinos emitted in opposite directions - which cannot exceed the value of 52.8 MeV. During a DIO process, instead, the outgoing electron can exchange momentum with the atomic nucleus: this nuclear recoil slightly distorts the Michel peak [13], creating a tail after the value of 52.8 MeV which approaches zero as $1/\mu^5$. This tail (Figure 1.5) extends up to the CE energy region and can thus results in a small but still significant probability for the outgoing electron to mimic a CE. To date, no measurements of this spectrum are available near the CE energy region. A recent theoretical calculation [13] gives an uncertainty near the endpoint which is better than 20%.

1.3.1.2 **Radiative muon capture**

During a radiative muon capture event, a muon is absorbed by an Al nucleus, with the emission of a photon:

$$\mu^- + N(A,Z) \rightarrow \gamma + \nu_\mu + N(A,Z-1) \tag{1.7}$$

These high energy photons can produce e^+e^- pairs in the secondary target or the surrounding materials. This phenomenon can be mitigated by choosing a target material suited to keep the photon energy below the CE threshold. For the Mu2e Al target, the photon energy endpoint is 101.9 MeV, located 3.1 MeV below the CE region, as the minimum Mg mass is slightly above the rest mass of Al.

1.3.2 Beam related background

1.3.2.1 **Prompt**

Beam-related background is associated with the production and transport of muons to the stopping target. Contaminations in the muon beam, indeed, can be associated with the production of electrons close to the conversion signal energy[1]. The main sources of contamination are hereby reported.

Pions that have not decayed and are able to reach the stopping target can undergo radiative pion captures (RPC) according to

$$\pi + N \rightarrow \gamma + N^* \tag{1.8}$$

with an endpoint spectrum near the π rest energy and a distribution peaking at 110 MeV. The mean decay time for pions in this process is $\tau_{\pi}^{Al} \approx 26$ ns. The asymmetric $e^{+}e^{-}$ conversion of such photons can yield an outgoing electron near the CE signal region.

A small fraction ($BR = 1.23 \cdot 10^{-4}$) of pions can decay in flight into high energy electrons via $\pi \rightarrow e^{-} + \overline{\nu_e}$. Finally, muons which decay in flight and have a momentum greater than 77 MeV can produce an electron with an energy in the CE region.

In these prompt processes the electron detection can be coincident in time with the arrival of a beam particle to the stopping target. These beam related backgrounds can be tamed by using a pulsed beam structure, along with a delayed observation window, which allows the associated prompt events to settle before acquisition. However, these schemes need a primary beam designed with an exceptionally good extinction factor (cf. 2.1.1 and 2.2.1).

1.3.2.2 **Delayed**

Beam-related backgrounds arise from the main proton pulse interaction with the production target: examples are represented by antiprotons and neutrons, which can be delayed by the transport lines. When produced on the primary target, antiprotons arrive late to the detector surface, as a result of their low speed. These types of background are reduced by implementing a sufficiently long muon transport line and, again, by a delayed acquisition window. Furthermore, antiprotons background is reduced to a negligible contribution by means of a specific absorber in the Mu2e Transport Solenoid described in section 2.3.2.

1.3.3 **Other**

1.3.3.1 **Cosmic induced**

Cosmic ray (CR) initiated events are expected to occur at a rate of approximately 1/day and will be tracked down by the CRV systems described in section 2.4.6. This asynchronous background is of great significance during the observation window and scales with its time length. Cosmic rays by themselves can be misidentified as CEs when their paths resemble trajectories originating in the vicinity of the target. Moreover, CR decay or interactions could produce secondary particles able to mimic the energy and the trajectory of a CE emitted from the stopping target (Figure 1.6). CR can indeed produce ~105 MeV/c electrons (or positrons) from secondary and delta-ray emission trough their interaction with materials, as well as via muon decay-in-flight.

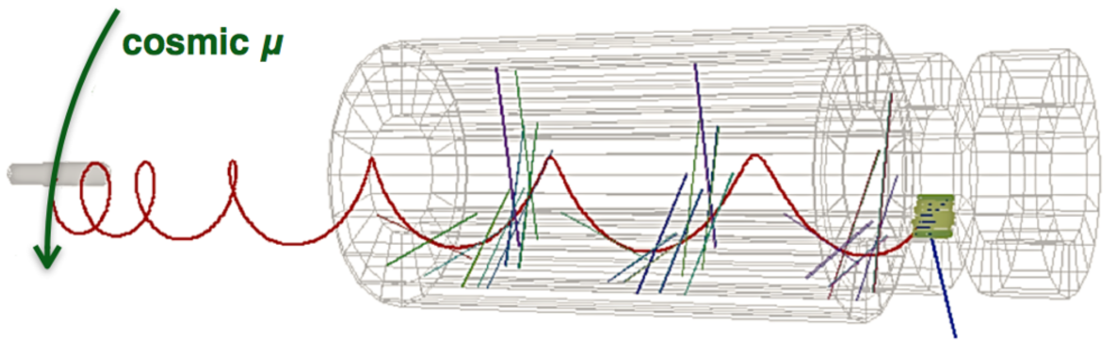

Figure 1.6 Cosmic induced events can mimic CE signals.

1.3.3.2 **Antiproton induced**

Antiprotons are generated by the proton beam interaction with the primary target or analogously by CR and can either mimic both the timing and the energy signature of CEs. Antiprotons can moreover be a source of additional background due to their interaction with matter. Low momentum antiprotons travel slowly through the transport line and their rates are not mitigated by the pulsed beam and delayed acquisition scheme. To prevent antiprotons from reaching the stopping target a thin absorber will be installed in the transport line (cf. 2.3.2).

1.3.3.3 **Mis-reconstruction errors**

Track mis-reconstructions can occur in conjunction with the presence of spurious signals in the detector, which can push low energy DIO electrons into the CE region. Thus, spurious events due to pile-up of low energy hits, as well as the environmental background, need to be taken into account trough the systematic uncertainties associated with the detector measurement to account for unwanted tails in its energy resolution.

Moreover, low energy protons ejected from the nucleus after muon capture are very ionising and can thus blind the tracker from detecting low energy electrons, potentially leading to reconstruction errors. Ejected neutrons, instead, once captured on the surrounding materials can produce low-energy photons, whose hits are very difficult to remove from measurements.

Finally, processes such as Compton scattering, photo-electric effect, pair-production and secondary electrons can all produce signals which are both in-time and with an energy spectrum close to the CE region.

As a consequence, high resolution detectors are needed in the experiment for the identification and suppression of all these background contributions.

2 Mu2e setup

2.1 Layout overview

The Mu2e apparatus is composed of three superconducting magnet sections with an S-shaped layout (Figure 2.1 and Figure 2.3). The magnets (Figure 2.2) rely on Al-stabilised Nb-Ti cables, protected from radiation damage by means of a bronze shield. The functions of each section are hereby reported.

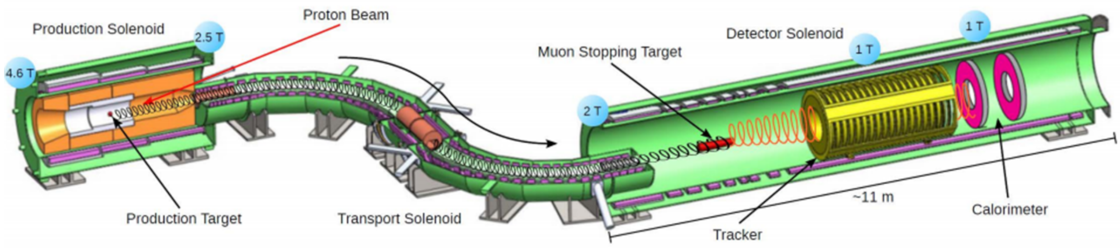

Figure 2.1 Mu2e setup and architecture.

- In the Production Solenoid (PS) the 8 GeV pulsed proton beam from the FNAL Accelerator System impinges on the tungsten primary target producing mostly pions, which are collected and focused by means of a magnetic lens into the Transport Solenoid section.

- The Transport Solenoid (TS) captures charged pions, selects the negative secondary muons resulting from pion decays and transports them to the stopping target. Its characteristic s-shaped geometry and its length are designed in order to allow the decay of all hadrons and suppress neutral line-of-sight particles. The TS selects low momentum muons in order to ensure that at least 40% of the transported ones can be brought to rest within the thin Al target housed in the Detector Solenoid.

- The Detector Solenoid (DS) houses the Al muon stopping target, where the transported muons will impinge at ~ 10 GHz, along with the detection instrumentation (tracker and calorimeter, cf. 2.4) needed for the efficient identification of the conversion electrons.

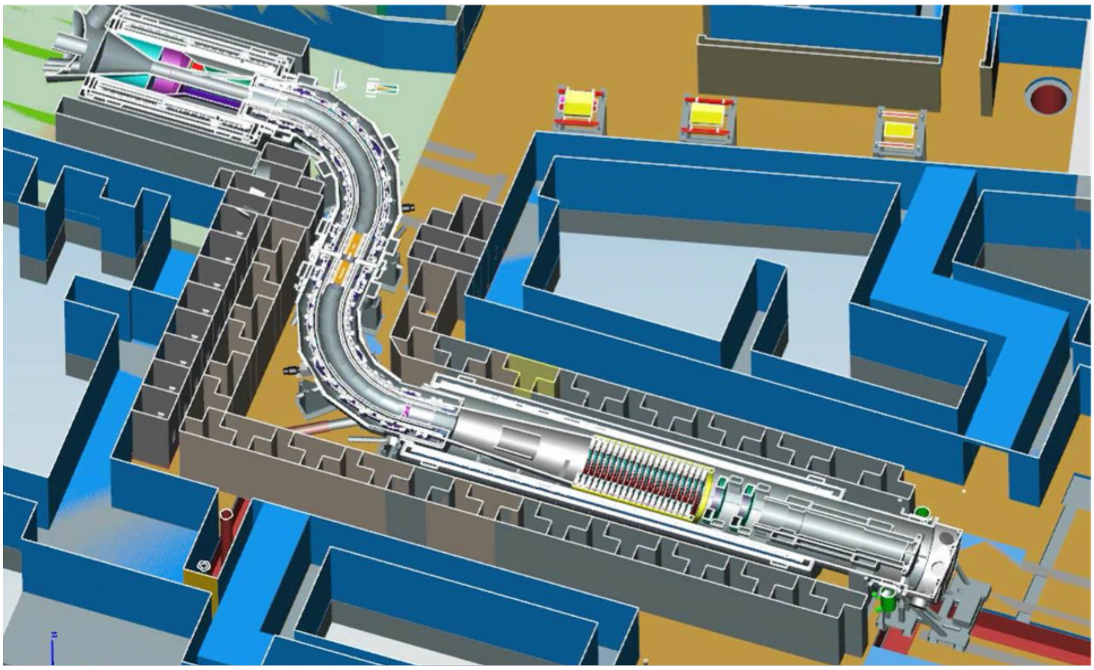

Figure 2.2 The Mu2e experimental hall.

In order to keep under control multiple scattering and muon stopping phenomena, as well as preventing arch discharges in the HV electronics and cabling, the entire muon beamline, as well as the totality of the DS, will be evacuated to down 10^{-4} Torr.

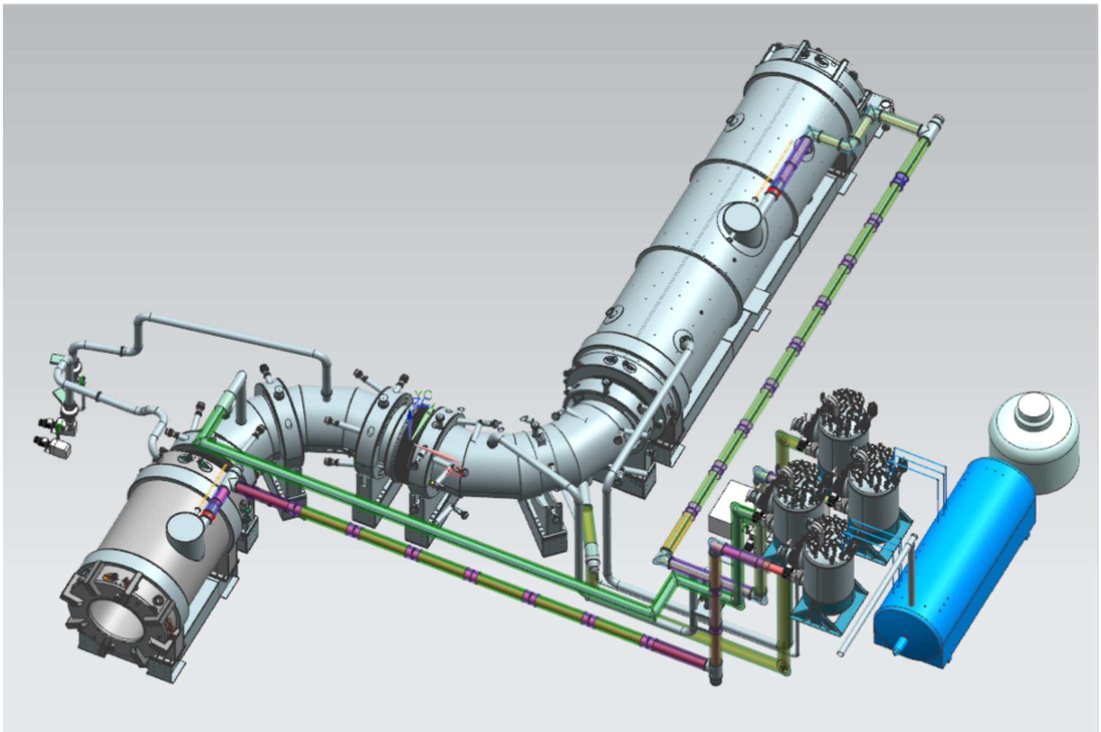

Figure 2.3 Mu2e cryostats, cryogenic distribution systems and feedthroughs.

2.1.1 Mu2e timing

Figure 2.4 reports the Mu2e timing scheme. The proton stopping process should ideally be distributed in a narrow time burst, the Proton-On-Target (POT) pulse, with a full width of less than 200 ns. It takes around 250 ns for the secondary muon bunches to reach the stopping target. Each bunch will consist of $\sim 2 \cdot 10^4$ muons, yielding a rate of approximately 10^{10} muons per second. The muon stopping process should ideally be distributed in a time window narrower than 250 ns.

The data acquisition window is gated 670 ns after the proton injection burst, in order to allow prompt background rejection (especially RPC, with $\tau_\pi^{Al} \approx 26$ ns). The sampling window is then closed after 925 ns, just before the next proton pulse, which follows a 1.7 µs spacing. The lifetime of muons stopped in Al is of around 864 ns (roughly one half the proton pulse scheme period): this timing scheme thus maximises the total number of muons on target and the signal level in the sampling window, minimising the prompt brackground level.

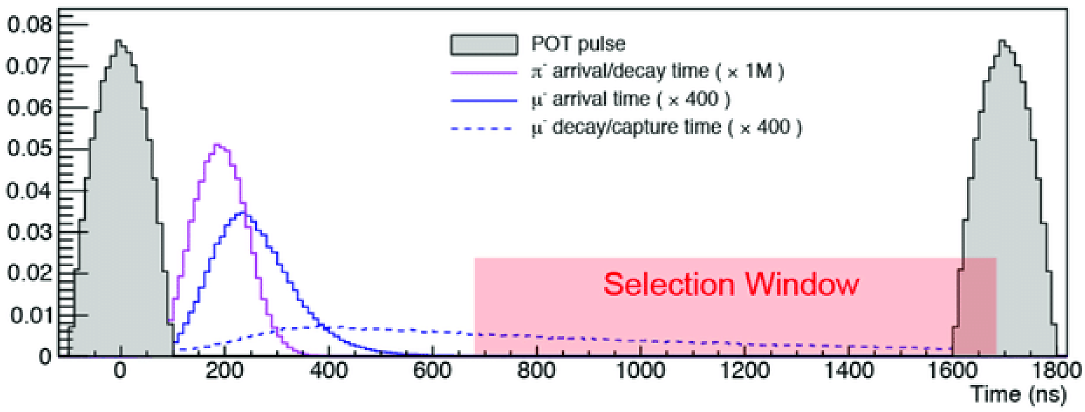

Figure 2.4 Mu2e timing scheme.

2.2 Proton beam and accelerator system

Mu2e is expected to receive more than 10^{18} protons during its run time [1]. The proton beam consists of a 200 ns pulse train with a 1.7 μs spacing. Figure 2.5 reports the layout of the accelerator complex [14] and its relative infrastructure.

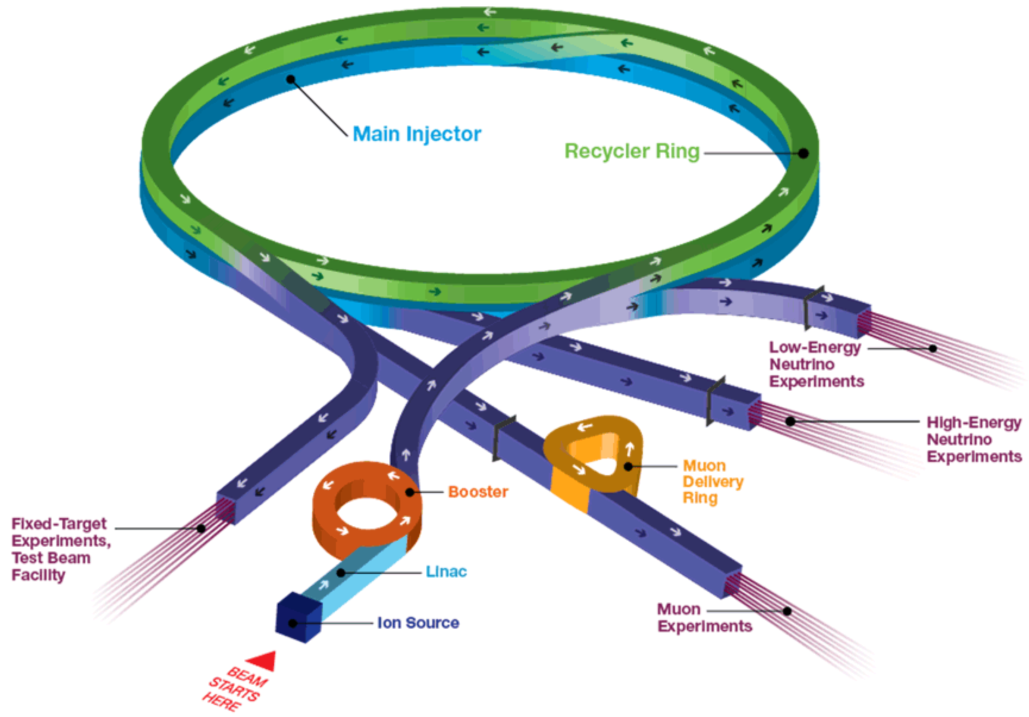

Figure 2.5 Fermilab accelerator complex architecture.

Protons are spilled for Mu2e according to the following stages:

1. 8 GeV bunches of $4 \cdot 10^{12}$ protons each from the Booster Ring are extracted into the Main Injector beamline and injected into the Recycler Ring;

2. by means of an RF manipulation sequence, the Recycler Ring performs a re-bunching at a rate of 15 Hz into 14 protons batches, two of which are reserved for Mu2e, while the other 12 are delivered to the NOνA [15] experiment (Figure 2.6);

3. bunches are slow-extracted and synchronously transferred to the Muon Delivery Ring;

4. a resonant extraction system injects a micro-bunch of $\sim 3 \cdot 10^7$ protons into the Mu2e beam line, once for every revolution in the Delivery Ring, which has a period of 1.7 µs; after the resonant extraction sequence completion, a clean-up abort kicker is fired to remove any remaining beam.

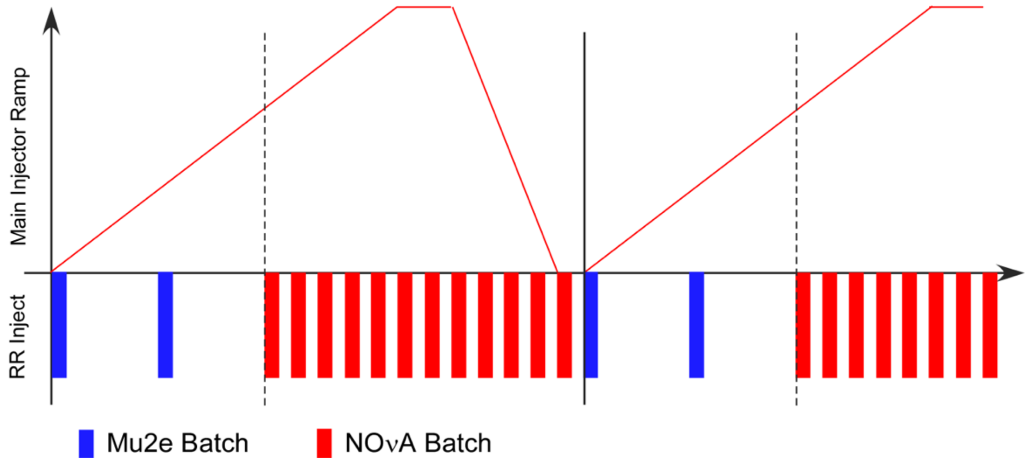

Figure 2.6 Mu2e spill scheme.

2.2.1 **Proton beam extinction and monitor**

In order to allow the delayed acquisition scheme to perform according to the requirements, it is mandatory to strongly suppress late arrivals, which are mainly associated with RPC background, not suppressible via the delayed gate. This dictates for a beam extinction factor (EF) better than 10^{10}, calculated as the ratio of in-time and out-of-time protons, with respect to the 200 ns full-width pulsed sequence shown in Figure 2.7. This roughly translates into a single out-of-time proton allowed every 300 bunches. An extinction factor of 10^5 is already provided by the micro-bunch resonant extractor and a further factor of 10^7 is provided by a system of resonant magnets and collimators in the beam line, to ensure a safety factor $> 10^2$ with respect to the extinction requirement.

An extinction monitor system is located above the Production Target (Figure 2.7) and measures scattered protons during the run-time to provide a direct indication of the residual beam. This instrument embeds a filter magnet, which selects 4.2 GeV/c momenta of scattered protons, a set of two collimators aimed at the target, and a spectrometer magnet. The momentum measurement is entrusted to two series of Silicon pixel detectors, triggered by scintillating counters. Finally, a muon range stack carries out muon identification.

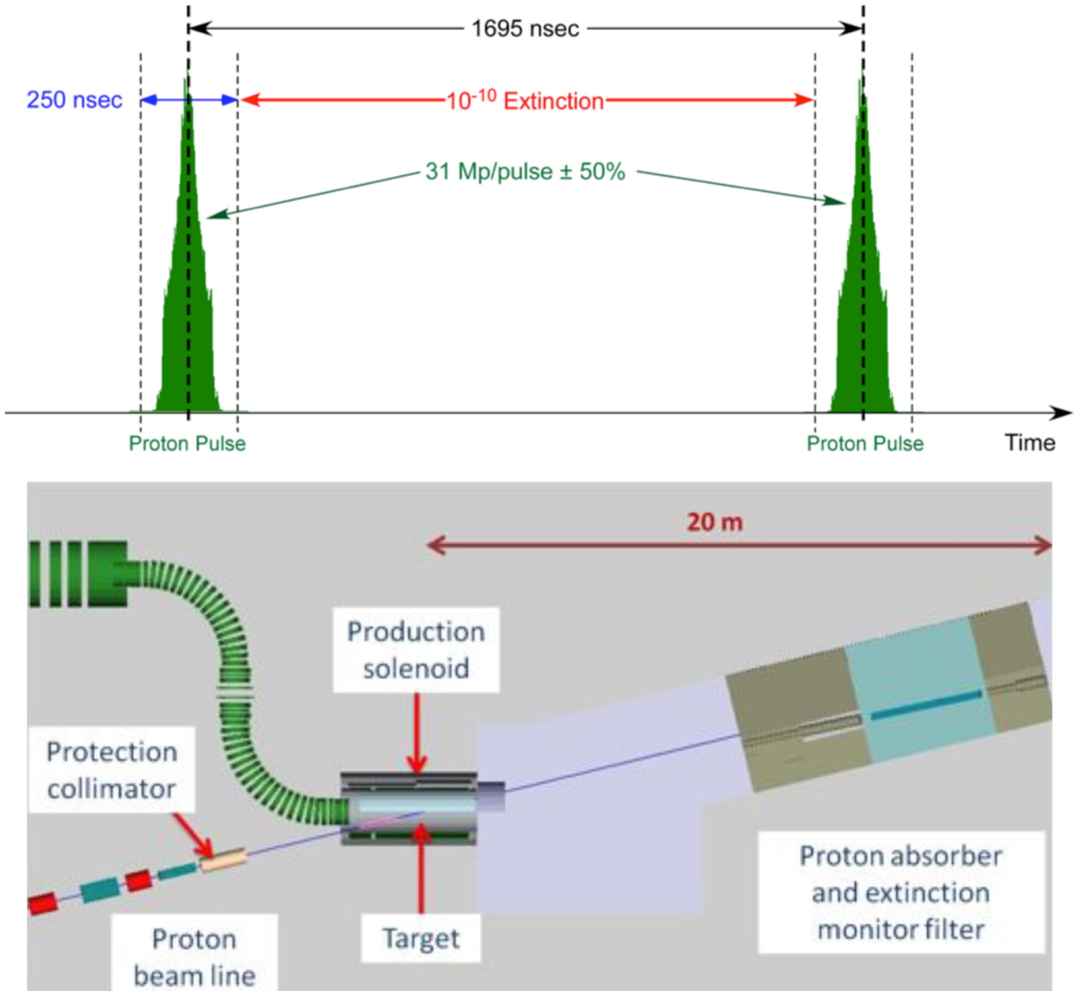

Figure 2.7 Top: proton pulse and extinction time diagram. Bottom: extinction monitor architecture.

2.3 Muon beam

The muon beam must be efficiently produced and transported to the stopping target in order to achieve the required detection sensitivity: intensity, time structure and purity represent the main parameters to be optimised in the secondary beam. A very high muon rate is essential for the required SES, and the currently proposed value is $4.21 \cdot 10^{10}$ μ^-/s. The delayed acquisition window requires the muon beam to have a pulsed structure with the already mentioned 1.7 μs period (cf. section 2.1.1), in order to suppress the prompt background. The muon stopping process should ideally be distributed in a time window narrower than 250 ns. Finally, inter-burst extinction is again fundamental to suppress non prompt background phenomena.

2.3.1 Production solenoid

The Production Solenoid (PS) is approximately 4 m long and has an inner bore of about 1.5 m (Figure 2.8). A bronze shielding is inserted in between the inner bore and the PS superconductor coil to mitigate radiation damage.

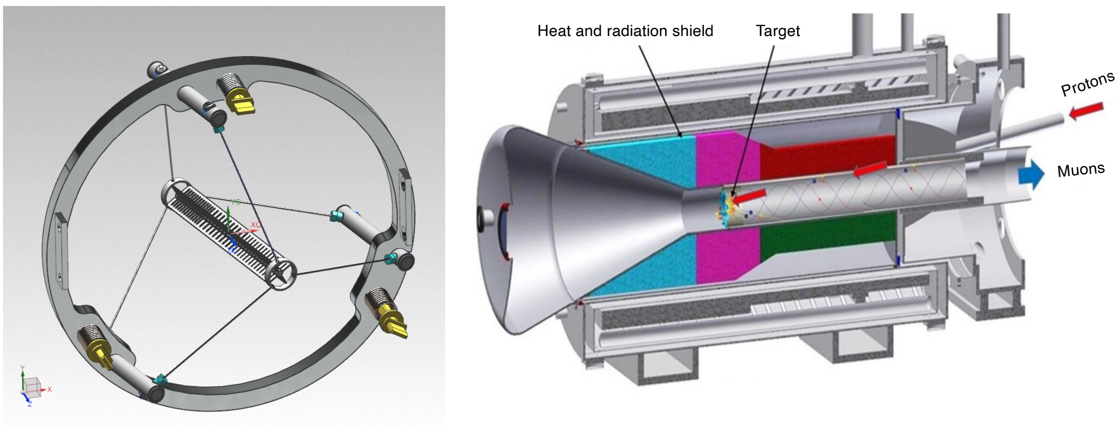

Figure 2.8 Production Solenoid architecture (right) and detail of the production target (left).

The 8 GeV proton beam enters the PS end and strikes a radiatively cooled tungsten target, consisting of a 160 mm long, 6.3 mm diameter segmented tungsten rod (Figure 2.9), suspended in a titanium support frame (Figure 2.8, left). The particle primary flux comes from beam flash, the early burst of particles hitting the production target. The primary proton beam enters the PS from the opposite direction of the TS exit point, so as to reduce the flux of secondary neutral particles.

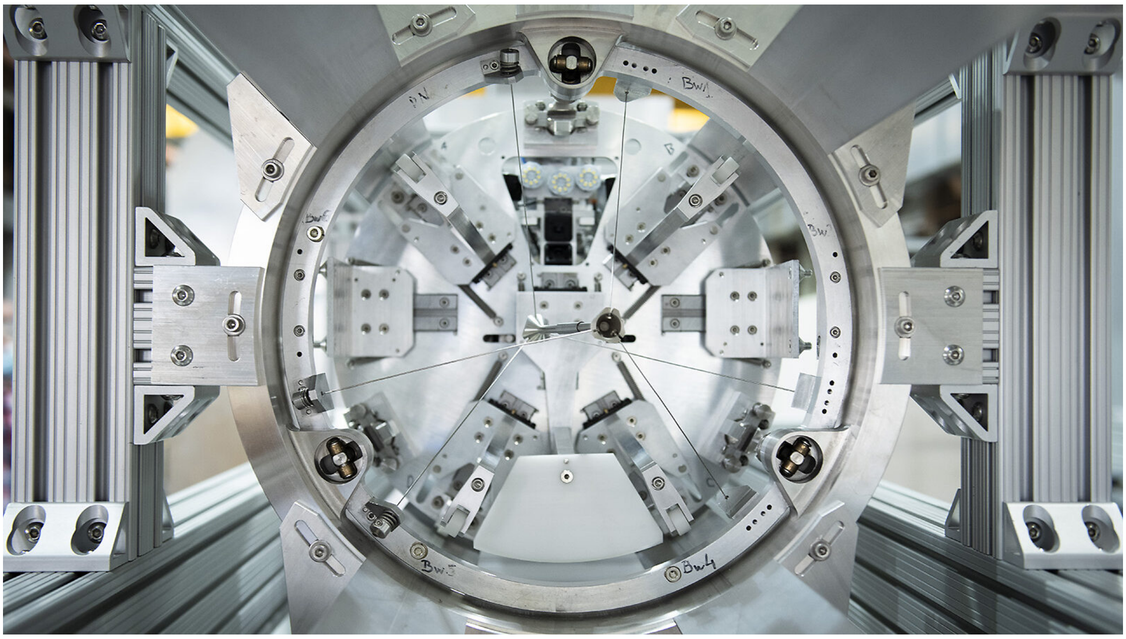

Figure 2.9 Tungsten production target installed in the support frame.

A high Z material such as tungsten is needed for the production target in order to maximize pion production and ensure an intense muon beam. The target geometry is designed to minimize pion reabsorption. The target supports are designed with a small physical profile to minimize scattering and absorption of pions and muons and the diameter of the warm bore of the Production Solenoid is large enough to allow pions and muons within the acceptance of the Transport Solenoid to pass through unobstructed. Furthermore, tungsten has been chosen as a function of its high melting point and its low thermal expansion coefficient.

The production solenoid generates a high intensity axially graded B-field varying from 4.6 T to 2.5 T, which serves a magnetic mirror to recover the backwards emitted charged particles and reflect them towards the low B-field region and the TS. The ratio between the transverse momentum (with respect to the axial coordinate) and the total momentum of a particle decreases as the magnetic field decreases, thus guiding the particle movement along the decreasing gradient.

2.3.2 Transport solenoid

The TS is composed of 27 superconducting sections: it allows the selection of low momentum negatively charged muons and their transport from the PS to the DS, thanks to a series of collimators and absorbers. Only a population of low momentum muons (< 80 MeV/c) are selected in order to ensure that at least 40% of the transported particles can be brought to rest within the thin Al targets housed in the DS.

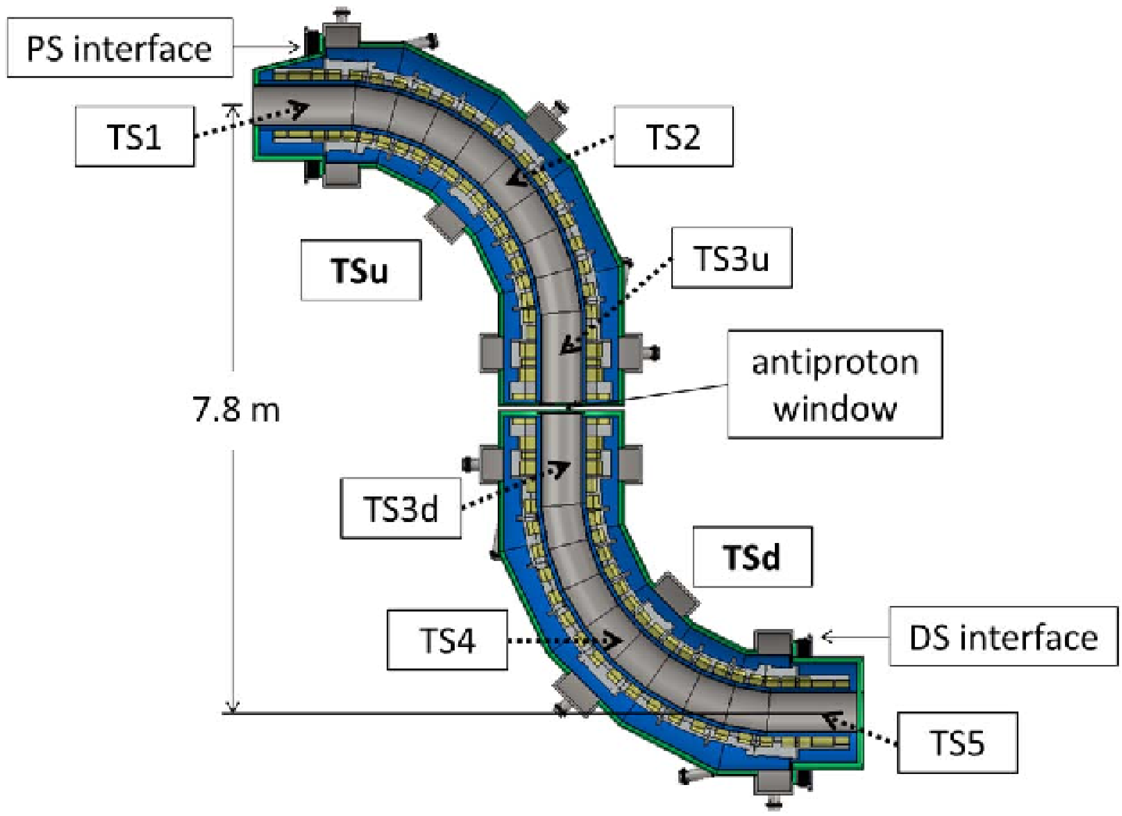

Figure 2.10 Transport solenoid architecture.

The S-shape of the TS and its length are designed in order to allow the decay of all hadrons and suppress neutral line-of-sight particles. The magnetic field in the TS has a negative gradient, from 2.5 T to 2 T, in order to minimise the transport of particles spending a long time in the magnetic system. This technique eliminates possible magnetic traps (in which particles are stuck bouncing between two local field maxima) and allows to accelerate particles trough-out their route to the DS.

The Transport Solenoid is composed of five distinct regions hereby described (Figure 2.10).

1. the 1 m long straight section TS1 links the PS to the TS and houses a collimator for the selection of particles with momentum inferior to the value of 100 MeV/c.

2. the TS2 prevents neutral particles from propagating from the PS to the DS, thanks to its quarter toroid shape.

3. the TS3 is a 2 m straight section which contains two collimators, which allows to filter particles by their sign and their momentum. The two sections which house the collimators are separated by a beryllium window which stops antiprotons produced in the PS.

4. similar to the TS2, the TS4 is another quarter of toroid, which prevents neutral particles from the beam interactions in the TS3 to reach the DS.

5. the TS5 interfaces with the DS entrance and is equipped with another collimator for momentum selection.

2.4 Detector solenoid

The Detector Solenoid DS is a ~11 m long, 2 m bore straight section which houses the muon target and the detection system composed by the tracker and the calorimeter (Figure 2.11 and Figure 2.12).

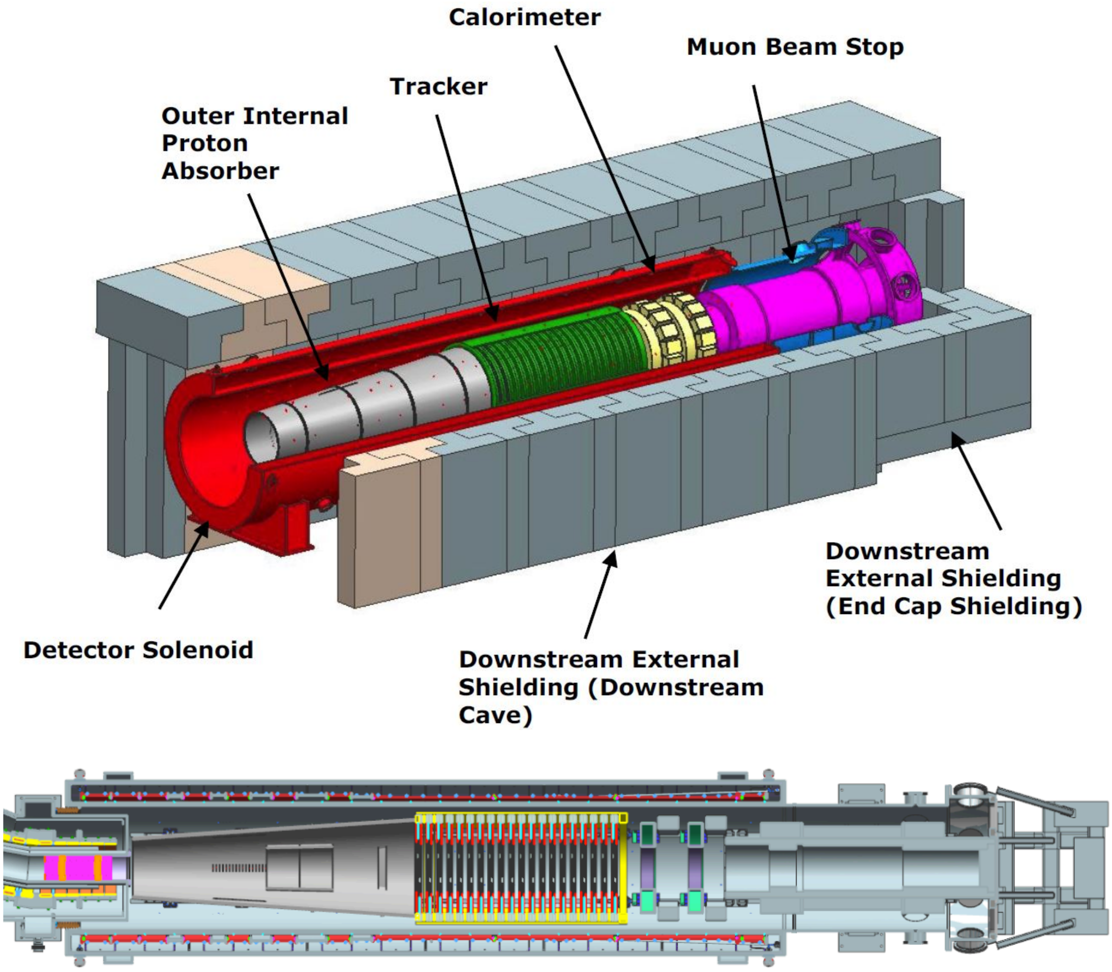

Figure 2.11 Detector Solenoid configuration (top) and cross section (bottom).

The magnetic field in which the stopping target is positioned is graded axially from 1.9 T to 1 T. This magnetic mirror scheme, once again allows to retrieve backwards emitted electrons and to reflect them towards the detector, thus increasing the acceptance for CE electrons. A significant portion of these reflected electrons will be scattered by or loose energy in the nearby materials, thus failing in the analysis selection and not appearing in the final data sample. The graded field serves the additional purpose of reducing the background from high energy electrons, which are accelerated towards the detectors with a resulting helix angle which does not match the CE trajectory and is thus rejected by the tracker system. In the region occupied by the detectors (after the stopping target and up until the end of the DS), the B-field remains uniform at 1 T.

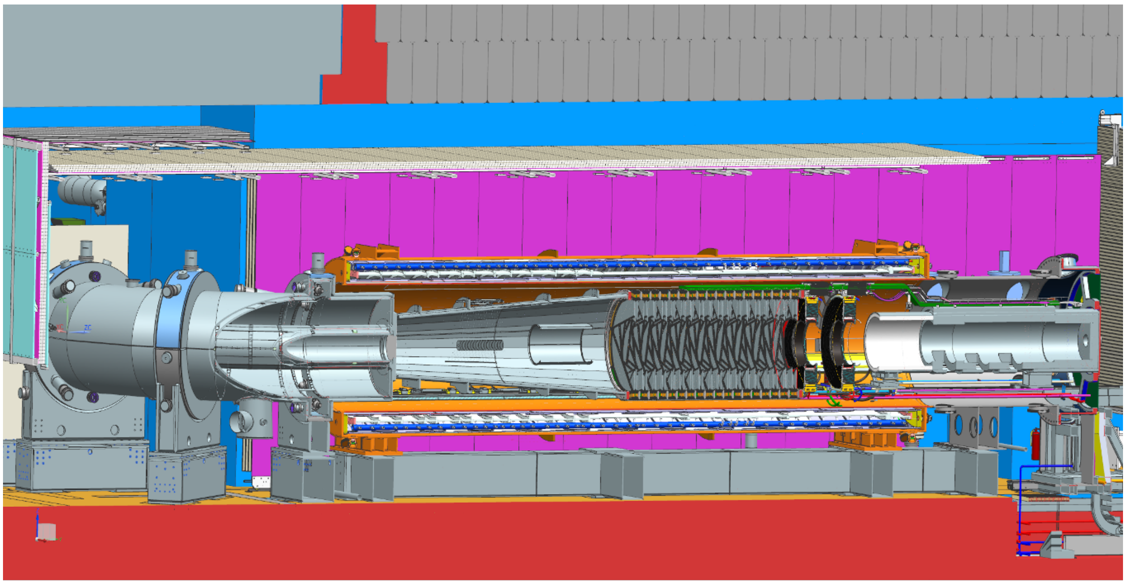

Figure 2.12 Detector solenoid design.

2.4.1 **Stopping target**

The muon stopping target design has to satisfy a series of conflicting requirements, such as stopping efficiency and reduced interference with the CE measurements:

- the target has to be sufficiently thick in the axial direction in order to allow the capture of large fraction (> 40%) of the impinging muons;
- the target geometry should not corrupt the momentum measurements of CEs by providing a controlled radiation length X_0;
- thickness and geometry are designed in order to tame background from Bremsstrahlung, secondary electrons, CR interactions, scattering of the incoming beam et similia.

The target design (Figure 2.13) is composed of 37 100 µm thick aluminium disks, axially stacked with 22.2 mm disk-to-disk spacing. The target assembly is supported by thin tungsten wires. As stated in section 1.2.4, the choice of Al results from a trade-off between a number of conflicting requirements and competing factors.

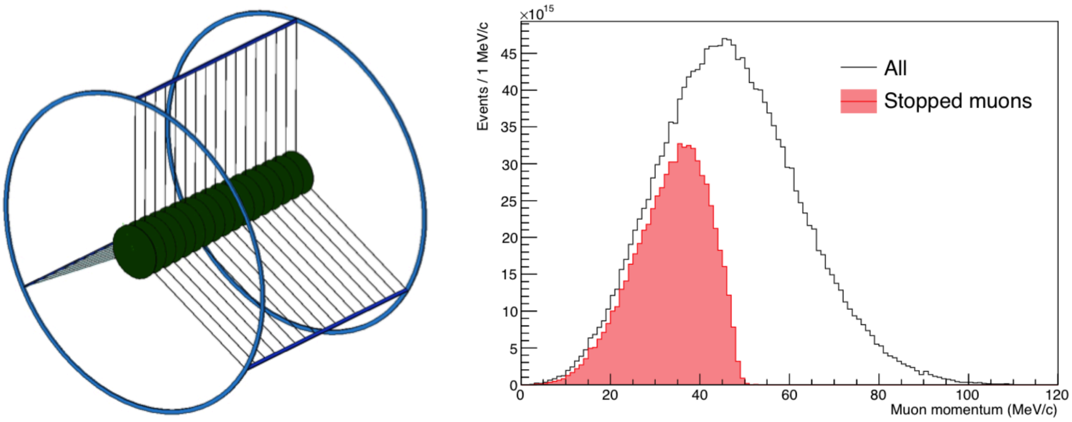

Figure 2.13 Left: schematic layout of the Mu2e stopping target and its mechanical support. Right: momentum distribution of muons delivered to the stopping target, with the superimposed distribution of stopped muons in the Al foils (red).

2.4.2 Stopping target monitor

The Stopping Target Monitor (STM) fulfils the task of measuring the total number of stopped muons in the target (with a relative accuracy of 10%), thus providing an important parameter for the conversion rate normalisation.

The STM is able to detect x-rays emitted during the capture cascade process of the muon toward the 1S state (cf. 1.2.1). The $2p \to 1s$ radiative transition is associated with the highest yield of X-rays. There is also a number of other processes, such as the $3p \to 1s$ and the $4p \to 1s$ transitions, which a have significant x-ray yield. Moreover, it is possible to identify muons captured by impurity atoms by their distinctive X-ray emission spectra

2.4.3 Muon beam stop

The Muon Beam Stop (MBS) is placed at the downstream end of the DS (Figure 2.11), in order to absorb the beam particles (mostly muons) while minimising the associated background levels (mostly due to muon decays) within the bore of the DS and the CRV. It consists of alternating cylindrical structures of stainless steel and HDPE.

2.4.4 Straw-tubes tracker

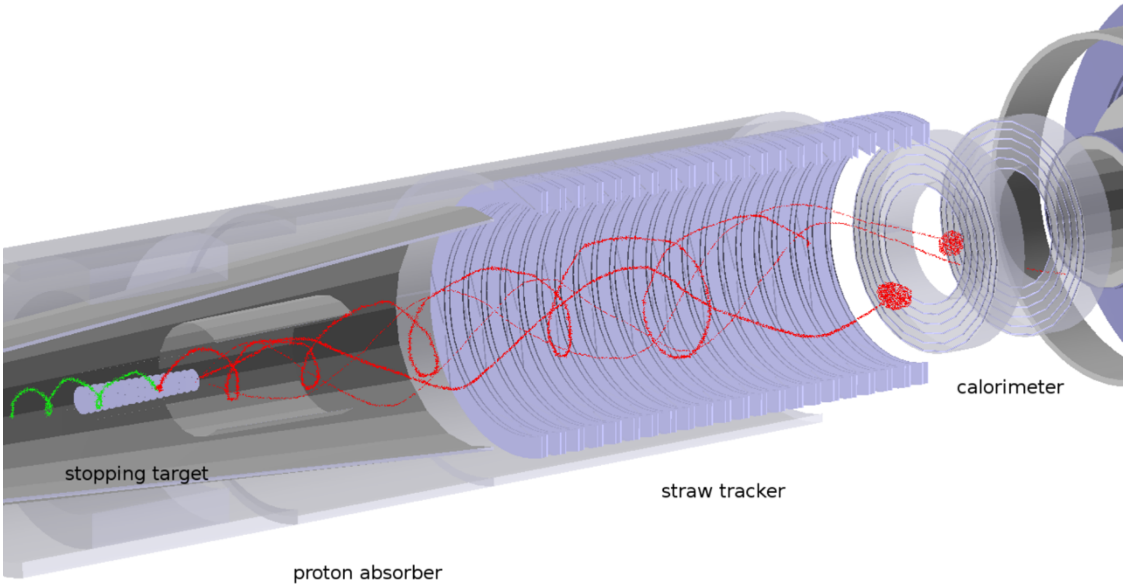

Figure 2.14 CE tracks inside the DS.

The trajectory of ~ 105 MeV/c CEs (Figure 2.14) and the associate momentum reconstruction is entrusted to the tracker. A low-mass and highly segmented detector is needed to suppress multiple scattering (the main cause of reconstruction errors) and cope with the high experimental rates, as spurious low energy hits can combine and mimic the passage of a CE.

Each element of the tracker is based on a 5 mm diameter drift tube, or straw, made of two layers of 6 µm spiral wound Mylar (a 3 µm layer of adhesive is present between the Mylar layers), for a cumulative wall thickness of 15 µm (Figure 2.16). The inner surface has a 500 Å aluminium coating and a 200 Å gold plating, which acts as the cathode layer. The outer surface has a 500 Å aluminium plating which acts as an additional electrostatic shielding, while also reducing the leak rate. In each straw a 25 µm gold plated tungsten sense wire is suspended. The chosen drift gas mixture is Ar:CO2 80:20 with a maximum operating voltage of 1.5 kV. The straws have active lengths varying from 334 mm to 1174 mm and are supported at their ends. A 1.25 mm gap between the straws allows for circumferential expansion due to gas pressure and manufacturing tolerances. The straw design is able to withstand changes in differential pressure ranging from 0 to 1 bar during operation in vacuum.

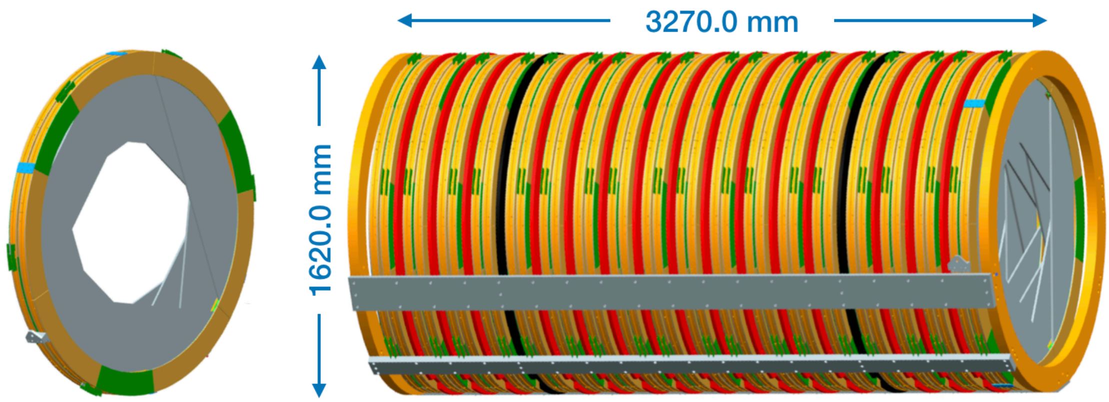

Figure 2.15 Straw tubes tracker architecture.

The straw tubes are oriented transversely with respect to the DS axis. Groups of 96 staggered straws are then organised in subunits, or panels, composed by two layers of straws each, to reduce the right-left ambiguity (Figure 2.15). The combination of six panels, shifted by 60° each, defines a plane. Finally, two planes combined together form a station (Figure 2.15, left). Each station is separated by 46 mm and is shifted by 180°. The 20736 straw elements which compose the tracker are divided in 18 measurement stations, distributed over a length of approximately 3 metres and assembled together in an annular disk, so that low momentum electrons (< 56 MeV/c) will pass through the inner void section without being detected, thus improving the measurement purity (Figure 2.17, left).

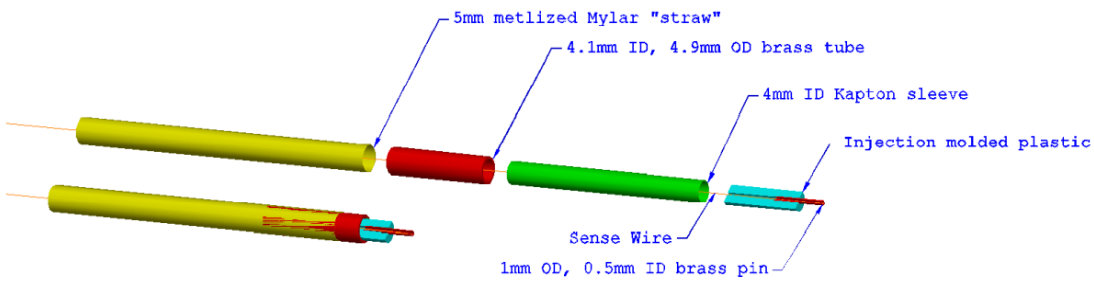

Figure 2.16 Detail of a straw tube.

Each straw end is pre-amplified and then read-out by two TDCs for timing reconstruction and one ADC for dE/dx capability. Signal digitisation is performed in the vicinity of the tracker and transmitted via optical fibres links, in order to minimise the number of feedthroughs. A liquid cooling system removes the heat dissipated by the electronic sections.

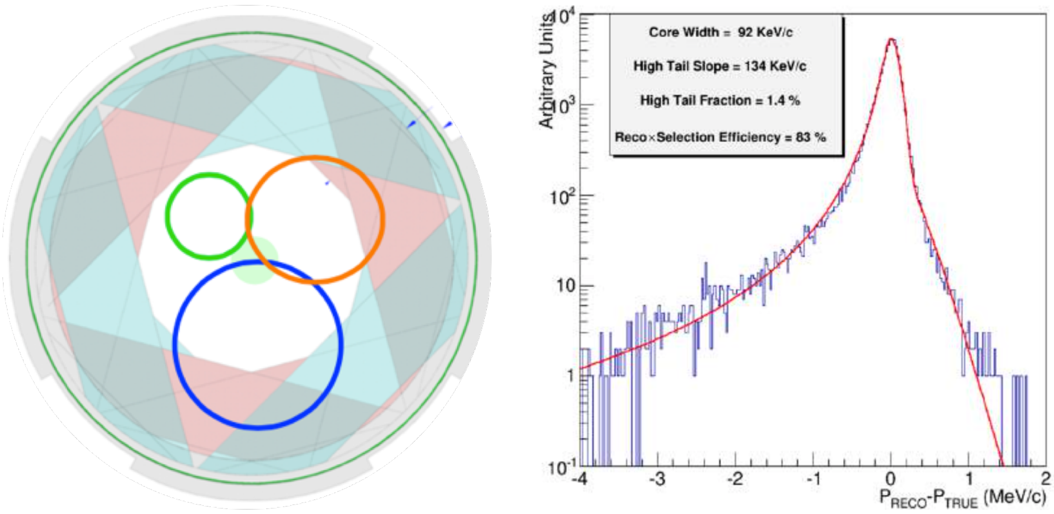

Figure 2.17 Left: low momentum particles will not produce hits in the tracker (green) and spurious tracks (orange) will be rejected with the selection cut. Right: tracker momentum resolution simulation for CEs. Full background overlay and pattern recognition are included. The data are fitted to a split normal distribution, with standard track fit quality cuts).

The position along the straw is reconstructed by weighing time information with charge deposition. The required longitudinal resolution on the drift radius is 100 μm. Tracker momentum resolution is crucial to positively identify the levels of several critical backgrounds: the instrument is required to have a high-side resolution $<$ 200 keV, as shown in Figure 2.17 (right). Simulations report that the tracker net resolution is smaller than the estimated deterioration due to the energy loss in the upstream material.

A small panel prototype was built at Lawrence Berkeley Laboratories to validate the design using cosmic rays and radioactive sources [16].

2.4.5 **Crystal calorimeter**

The primary function of the calorimeter is to provide a second set of complementary measurements and to validate the tracker reconstruction, by providing precise information about particle energy, timing and position. Moreover, it allows the identification of backgrounds and their rejection to a level consistent with the SES target level. A calorimeter with high acceptance for CE, good energy and time resolution is thus needed (requirements are stated at 105 MeV):

- $\sigma(E)/\langle E\rangle < 10\%$
- $\sigma(t) < 500$ ps
- $\sigma_{xy} < 10$ mm

In addition, the calorimeter will have to satisfy the following additional requirements:

- PID capability with e/μ separation with a μ rejection factor >200;
- cluster seeding to support the tracker in efficient pattern recognition and reconstruction, thus combining the extrapolated helix trajectory with a precise measurement of energy deposit, timing and position;
- standalone on-line trigger capability based on the deposited cluster energy;
- high level of background rejection for non-vetoed CR initiated events.

The more detailed description and characterisation of the Mu2e calorimeter is reported in Chapter 3.

2.4.6 CRV

Cosmic ray induced background suppression is of uttermost importance for the target Mu2e SES, as previously discussed (cf. 1.3.3.1) and it is expected to produce conversion-like events with a rate of 1/day.

The CRV (Figure 2.18) incorporates an active veto system, which encapsulates the totality of the DS and the final part of the TS (for a total surface of 336 m^2). The veto signals are generated by the coincidence of adjacent counters in different layers. In the region of the muon stopping target, the CRV has a muon veto efficiency of 99.99%, in order to tame the CR induced background rate down to 0.10 events over the entire data-taking period.

The active veto system relies on 86 modules. Each module is composed of 4 layers of long extruded polystyrene plastic scintillator strips, with the interposition of absorber layers, which reduce electron punch-trough (Figure 2.19). Each scintillator counter is equipped with a TiO2 coating, and embeds wavelength shifting fibres, which couple light to the SiPM in a double-end read-out scheme. The required light yield is of at least 15 pe/cm/SiPM, along with a signal rise time inferior to 5 ns.

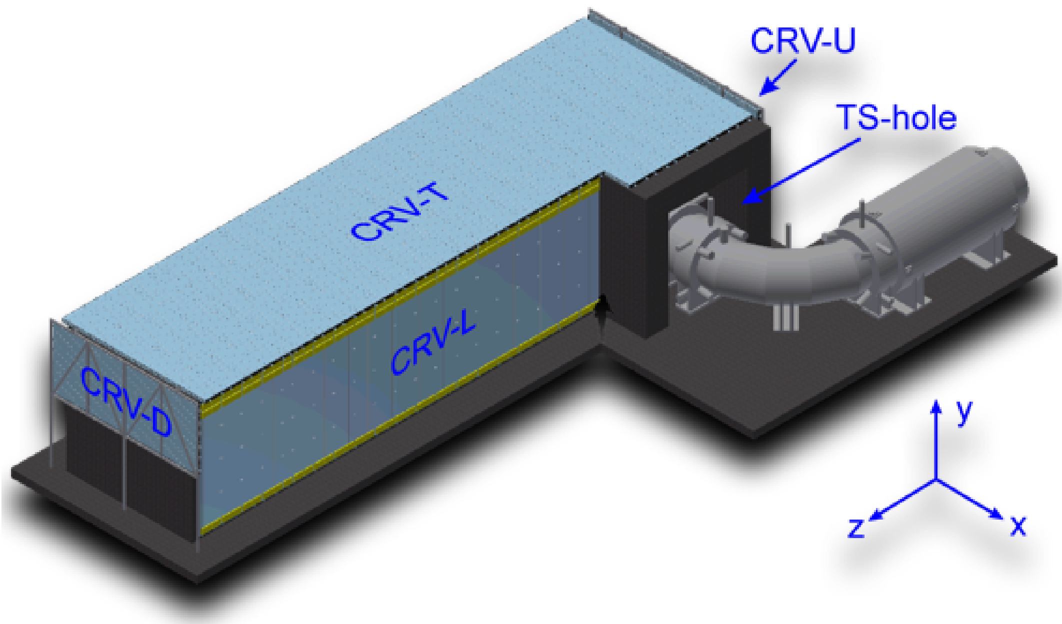

Figure 2.18 Cosmic Ray Veto layout.

A passive shielding between the CRV and the solenoids has been designed in order to minimise the CRV dead time associated with the intense neutron flux arising primarily from the muon stopping target. Most of the neutrons have kinetic energies below 10 MeV, with the most probable energy value at 1 MeV. Simulations [1] show that gamma emission associated with their capture on hydrogen is the primary source of counts for the active veto detectors. The passive shielding moderates and captures most of the neutrons. The magnitude and pattern of energy deposition in multiple layers of scintillator is expected to be different for neutrons and muons. Therefore, false veto signals from neutrons can be reduced to an acceptable level. Moreover, a passive shielding consisting in a thick layer of concrete surrounds the DS and allows for the suppression of background sources other than penetrating muons.

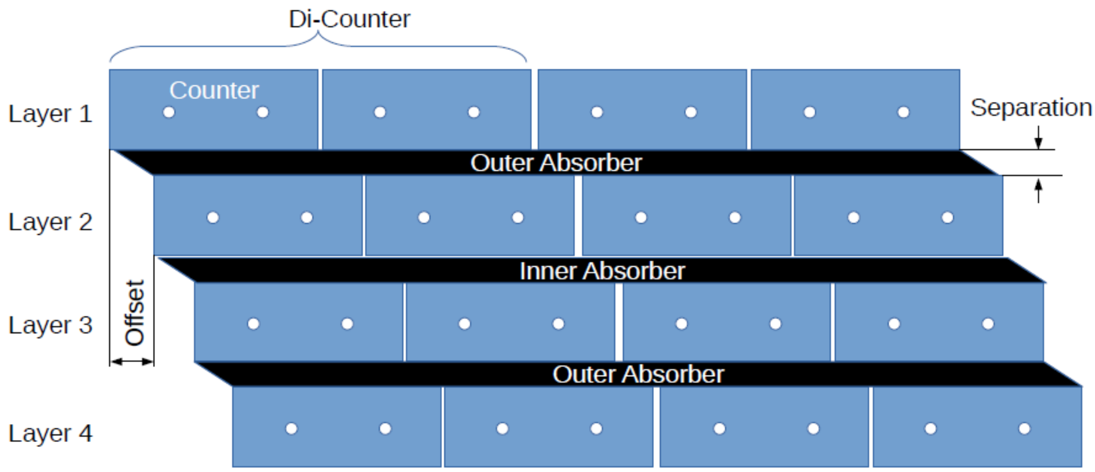

Figure 2.19 CRV active veto system architecture.

3 Calorimeter design and characterisation

3.1 General requirements

The primary purpose of the calorimeter is to provide the tracker data with a complementary set of measurements of particle energy, timing and position. This allows to validate and optimise the particle identification and track reconstruction processes, as well as allowing backgrounds rejection to a level consistent with the SES level (1.2.3).

A calorimeter with high acceptance for CEs and with the following resolution requirements at 105 MeV is thus needed:

- $\sigma(E)/\langle E\rangle < 10\%$
- $\sigma(t) < 500$ ps
- $\sigma_{xy} < 10$ mm

The energy resolution requirement is needed in order to allow the validation of the momentum measurement carried out by the tracker, which is much more precise ($\sigma < 200$ keV/c). As a consequence, temperature and gain stability are needed within ± 0.5% with respect to specifications. The position and time resolution requirements allow the position and timing of the calorimeter clusters to be matched with the extrapolated trajectory of the reconstructed tracks.

The other general functional requirements are hereby reported.

- PID capability with a minimum e/μ rejection factor of 200, in order to discriminate and reject the ~40 MeV clusters due to 105 MeV/c muons.
- Cluster seeding to support the tracker in efficient pattern recognition and reconstruction. Tracker hits are selected across a ± 50 ns window with respect to the calorimeter cluster timing and on the basis of their

compatibility with the cluster azimuthal angle. Fitting a helix trajectory to this selected subset of tracker hits will increase the tracking efficiency.

- Standalone on-line trigger capability based on the deposited energy cluster.
- High level of background rejection for non-vetoed CR initiated events.
- Fast response time ($\tau < 40$ ns) to cope with the high expected experimental rates and prevent pile-up.

The Mu2e calorimeter is required to achieve extremely high levels of stability in the measurements by maintaining the previously stated performances throughout the entire three-year expected run time and in the harsh Mu2e operating environment. This demand can be summed up in the following key points:

- keep an efficient operation in the high-radiation Mu2e environment and for radiation exposures up to ~15 krad/year (in the hottest region) and for equivalent neutron fluencies up to 10^{11} $n_{1\mathrm{MeV}}/\mathrm{cm}^2$;
- operate inside an environment evacuated to 10^{-4} Torr;
- operate in the 1 T B-field of the DS.

3.1.1 Trigger system

The Mu2e trigger will filter the events of interest with the highest possible efficiency, in order to keep the total rate below 2 kHz and the data rate to disk inferior to 0.7 GB/s. The events of interest also include calibration and control samples, as well as data for background evaluation. The trigger system must comply with the following requirements: an efficiency greater than 90%, a background rejection factor better than 100 and a processing time inferior to 3.6 ms per event.

3.2 Calorimeter design

A total absorption calorimeter employing a homogeneous continuous medium is required to meet the Mu2e precision requirements. After a long R&D phase the detector design was concluded in 2015 with the release of the final technical choice [17]. The calorimeter has a geometry optimised for maximum CE acceptance and is composed by two annular disks with an inner (outer) radius of 35 cm (66 cm) and a relative distance of 70 cm, value which corresponds to approximately half the expected pitch of the helical CE trajectory (Figure 3.1). Each disk has a high granularity and houses 674 parallelepipedal $3.4 \times 3.4 \times 20$ cm^3 crystals.

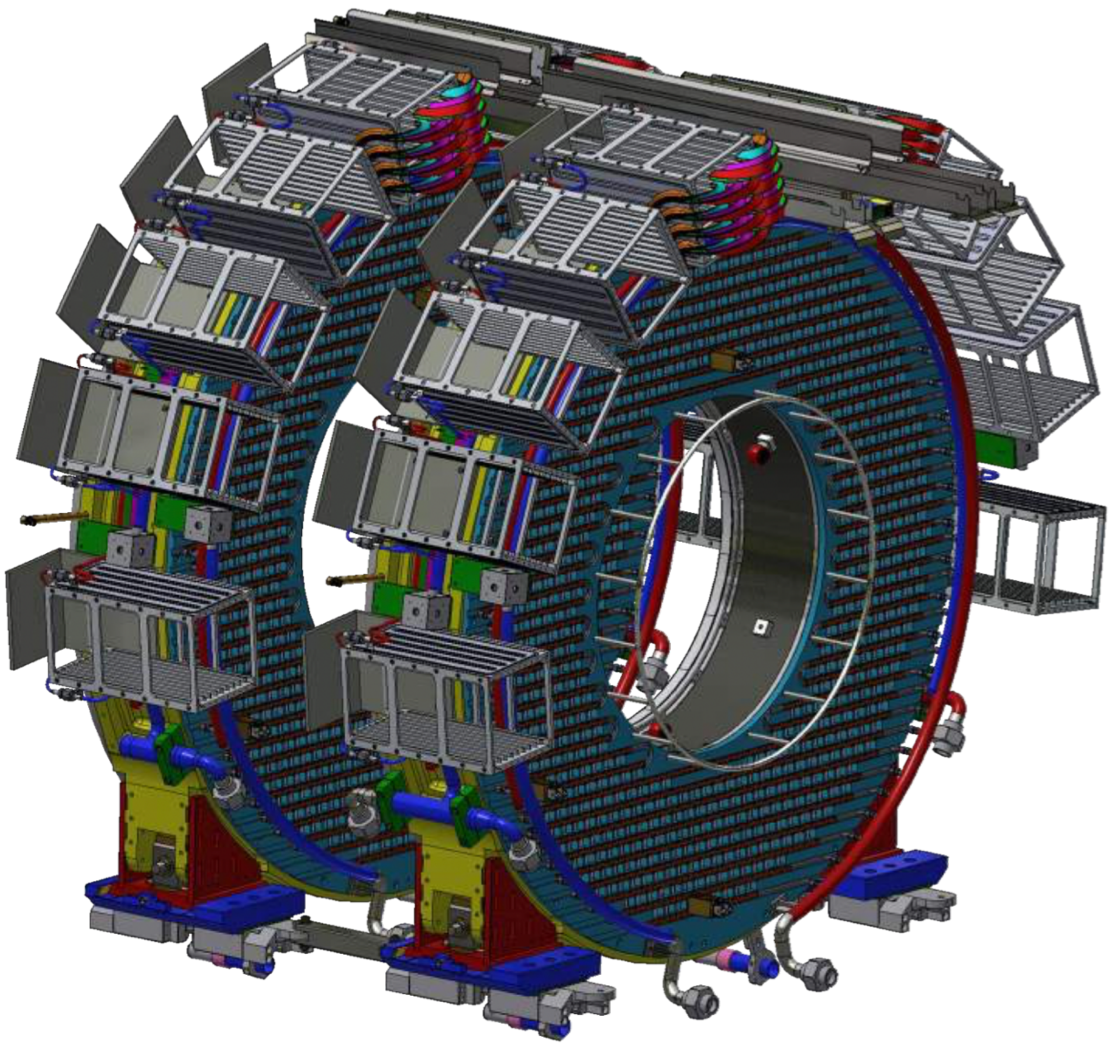

Figure 3.1 Rendering of the Mu2e calorimeter, along with the cooling lines and electronic crates.

3.2.1 Crystals

The choice of the crystal properties is of paramount importance for a successful development of the Mu2e calorimeter. During the R&D phase, four crystals were considered: Lutetium-Yttrium Oxyorthosilicate (LYSO), Lead Tungstate ($PbWO_4$), Barium Fluoride (BaF_2) and pure Caesium Iodide (CsI). LYSO was the first option considered, but due to the increase in Lutetium salt prices, it was eventually discarded as the associated costs would have been prohibitive. Moreover, $PbWO_4$ was also excluded due to its poor light output (LO). The two choices left, BaF_2 and CsI, have similar performance in terms of LO and X_0 values. Eventually, the use of BaF_2 was discouraged due to its scintillation properties: apart from its principal fast decaying component at 218 nm (0.9 ns decay time), this material presents a significant slow component (650 ns decay time) peaking above 280 nm. This element calls for a photodetector with a spectral responsivity able to suppress the slow decaying light and prevent pile-up. The material of choice for the crystal is thus non-doped CsI, which represents an excellent match with new generation photosensors.

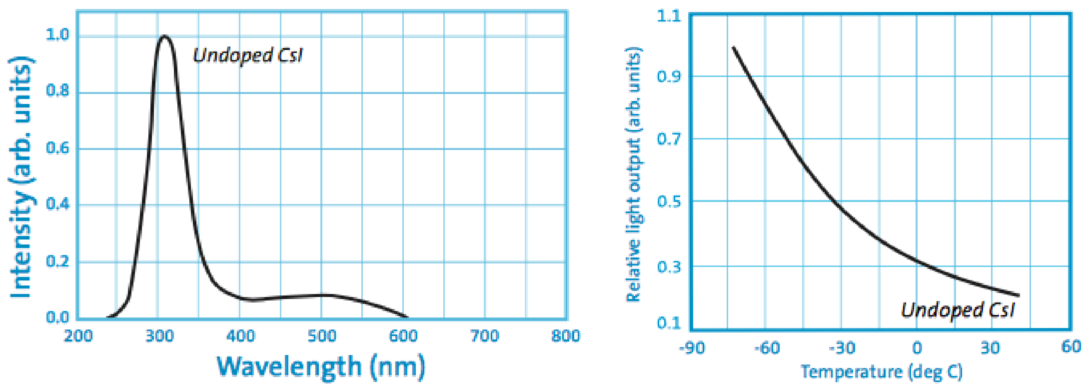

Figure 3.2 Left: CsI emission spectrum. Right: CsI LO versus temperature.

The emission properties of pure CsI (Figure 3.2) are heavily affected by the presence of trace impurities even below the ppm level and by the growth process itself. Generally, CsI has a fast decaying emission component ($\tau \approx 20$ ns) which accounts for more than 80% of the LO, with a spectral peak at ~ 315 nm. This emission peak is due to the radiative decay of self-trapped excitons and is associated with the luminescence of the pure material. A much slower (and smaller) component with a ~ 500 nm peak and a decay time of ~ 1000 ns is also present. This component is instead associated to the presence of I^- vacancies in the lattice, due to trace impurities in the crystal: as a result, its intensity decreases with crystal purity. Figure 3.3 reports the optical pulse of a CsI sample crystal tested with a ^{22}Na source and an UV-extended PMT readout, fitted with a triple exponential curve: the fitted fast decay time is of approximately 30 ns, and the slow component accounts for a very small fraction of the total detected light (< 1%).

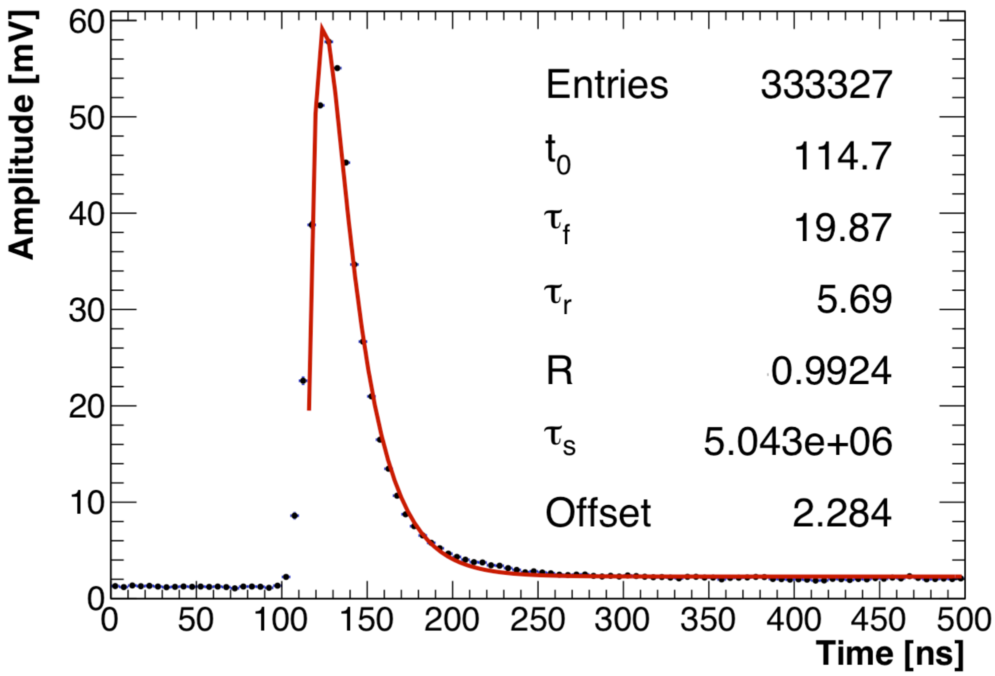

Figure 3.3 CsI light output waveform with a superimposed triple exponential fit (with PMT readout).

The crystals length has been set to 200 mm, which corresponds to a nominal value of $10 \cdot X_0$. If the average incidence angle of 50° - which corresponds to the average CE helix inclination - is taken into account, the length becomes effectively 300 mm, resulting in an acceptable shower containment of $15 \cdot X_0$.

Each crystal will be individually wrapped with a 150 µm Tyvek foil, for which dedicated studies report a ~ 90% reflection efficiency at 315 nm. In the final assembly, a 2 mm gap will be present between the crystal and the SiPM, in order to reduce thermal coupling between the two, both for mechanical and optical reasons. Indeed, the CsI light output is strongly quenched for increasing temperatures and even at room temperature: cooling the crystals to -77 °C would indeed increase the LO by an order of magnitude (Figure 3.2). For these reasons, the presence of any thermal gradients along the crystal would be significantly harmful for its optical response performance.

The technical specification relative to the quality control for the production crystals are verified using a 150um Tyvek wrapping and an UV-extended PMT (such as the Hamamatsu R2059 [18]), coupled trough an air gap. The following requirements must be satisfied when testing the crystals with 511 keV annihilation photons:

- a light yield LY > 100 pe/MeV;
- an energy resolution < 20%;
- a light response uniformity LRU < 5%;
- a fast-to-total ratio F/T greater than 75%.

The LRU is evaluated as the LY RMS relative for 8 measuring spots along the crystal axis.

As far as the mechanical specifications are concerned, it is required for the crystals to have a flatness below 100 µm for all faces, along with dimensional tolerances of 100 and 200 µm respectively for the transversal and longitudinal dimensions.

3.2.2 **Readout system**

The choice of Silicon Photomultipliers is driven by the magnetic environment of the DS. As hands-on maintenance on the calorimeter is allowed only once per year [1], the photosensors must be designed and installed with good reliability. Each crystal is readout by two identical photosensors, so as to improve reliability through redundancy. In order to achieve the desired spectral responsivity at 315 nm, custom UV-extended SiPMs from Hamamatsu are used, with an operation characterised by low bias voltage, high speed, high detection efficiency and gain (as discussed in section 4.2). Figure 3.4 shows the Mu2e SiPMs, along with the relative custom FEE boards and SiPM holders.

The photosensors signal are processed by custom front-end electronics (FEE) boards, mounted directly behind the SiPMs, which incorporate a fast pulse shaper and allow for independent control over the bias voltages, as well as providing individual current monitoring capability for each sensor. The systems responsible for the HV and power supply distribution, the slow control and the digitisation are housed around each disk in custom crates (Figure 3.5). The Data Acquisition (DAQ) system samples each channel at 250 Msps, allowing timing reconstruction and pileup deconvolution capabilities. The detailed description of the electronics and the signal chain of the Mu2e calorimeter is reported in Chapter 4.

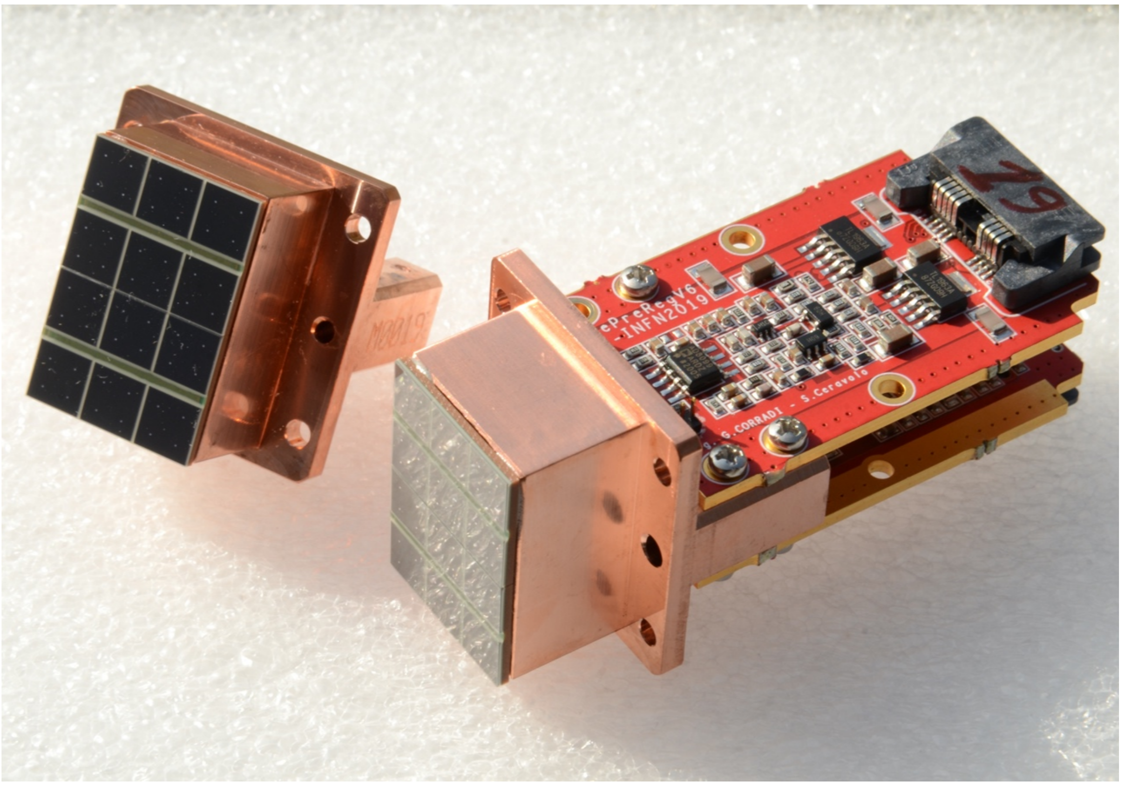

Figure 3.4 Mu2e SiPMs with the respective FEE boards and holders.

3.2.3 Mechanical structure

The calorimeter mechanical design (Figure 3.5) has to fulfil a series of tasks which are of uttermost importance for the calorimeter success and stability, while also being able to withstand the Mu2e harsh environment. It has to provide sufficient structural rigidity, alignment of all crystals and readout systems, as well as allowing fine regulations of the calorimeter alignment in the DS.

Moreover, the structure will have to integrate the distribution lines for the cooling systems and the radioactive source fluidic circuits (cf. 3.3.1.1 and 3.2.5). Each stack of 674 crystals is supported by two coaxial cylinders. The inner one has to be as light as possible, so as to minimise the amount of passive material which could spoil particle energies: the material of choice is indeed carbon fibre. The outer one, instead, carries the quasi totality of the load given by the crystal weight and is realised in a single CNC-milled stabilised Al block. This material choice is due to its mechanical and radiation hardness properties, as well as its machinability. This structure will be monolithic and able to integrate the two support bases (Figure 3.5).

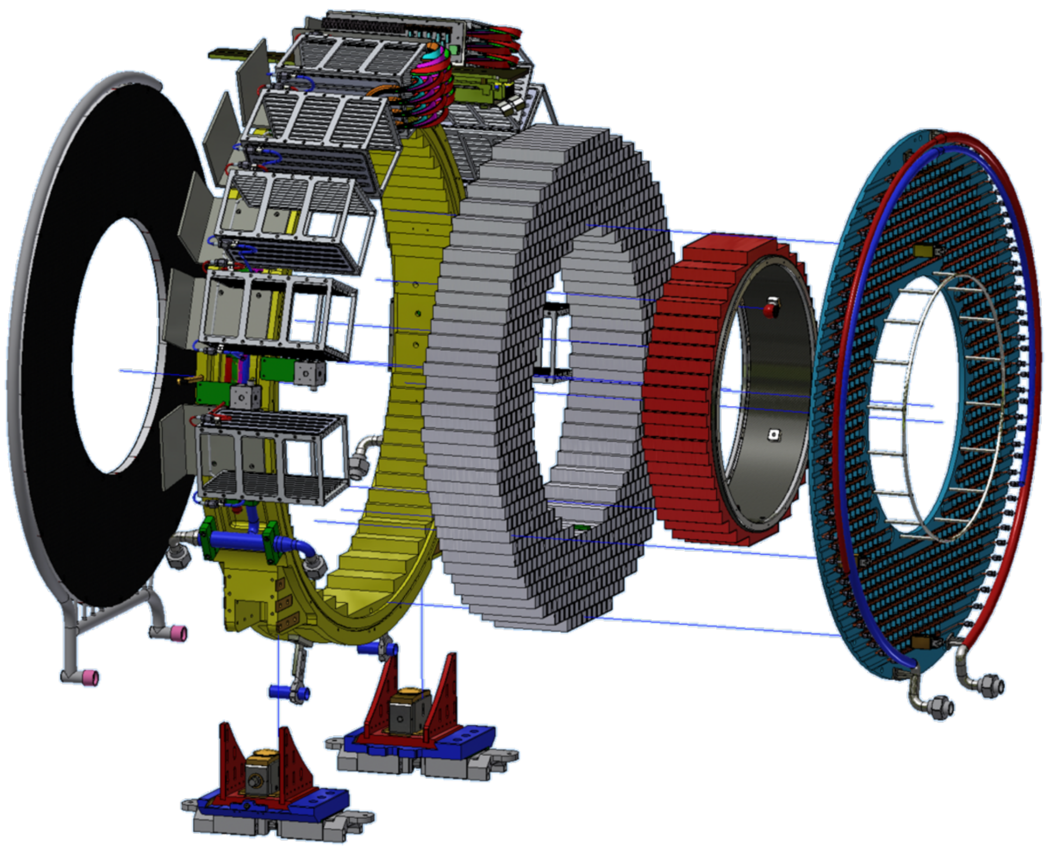

Figure 3.5 Exploded view of the calorimeter. From left to right: front plate, outer ring with 10 crates installed, crystal stack, inner ring and backplate.

Two cap plates complete the assembly. The one facing the beam is designed in carbon fibre, as to minimise the degradation of particle energies. A piping system for the transportation of a liquid radioactive calibration source will be embedded in this structure (cf. 3.3.1.1). The back one is instead made of PEEK. It will support the SiPM holders along with the associated FEE, while providing a sufficient level of electrical insulation. It also has to embed the cooling blocks and coolant distribution lines for the FEE system.

Inside each disk of the calorimeter, the 674 crystals are stacked vertically in staggered rows, resulting in a self-supporting array weighing approximately 700 kg, which are distributed over the outer annular chassis (Figure 3.5). This arrangement reduces the load on the inner cylinder by providing a catenary arch-like upper structure. The inner cylinder is constrained to the external cap plates, to which it transfers the load due to the crystal. Its design is optimised to minimise deformation, so as to prevent any misalignment of the crystals. The crystals are registered with respect to any radial or longitudinal movement by means of ASA additively manufactured frames which acts as shims and allow handling the calorimeter disks without any harm to the crystal alignment.

The whole disk assembly is able to slide on linear rails which allow its installation inside the DS. The calorimeter anchoring mechanism embeds a position fine-tuning system with a 10 mm travel per axis, relying on screw/wedges trimming systems and able to compensate for any misalignment with great accuracy.

3.2.4 **Vacuum considerations**

As stated in Chapter 2, the DS will be evacuated down to 10^{-4} Torr. This vacuum level imposes a design requirement which limits the outgassing rate of the whole calorimeter to a value of $8 \cdot 10^{-3}$ Torr $\cdot$ litre/s. This value is set to represent a negligible contribution if compared to the allowed tracker outgassing, which is one order of magnitude larger and is dominated by gas leaks in the straw tubes. The outgassing performance of each individual calorimeter component made of not-standard materials has been verified at the LNF vacuum facility by kinetic outgassing tests.

3.2.5 **Cooling systems**

The calorimeter will have a dedicated cooling system to cope with the heat generated by the front-end and back-end electronic systems. As far as the latter is concerned, the control and digitisation electronics in each crate will dissipate a total of 230 W (considering a 25% safety factor). This heat is removed by means of an aluminium block which is in contact, through a thermal pad, with the electronic PCBs and allows conduction towards the cooling lines (Figure 3.6, right). The digitisation PCBs have a maximum allowed operating temperature of 60 °C. Cooling power is derived from the circulation of a 35% glycol-water mixture, which enters the network at -20 °C.

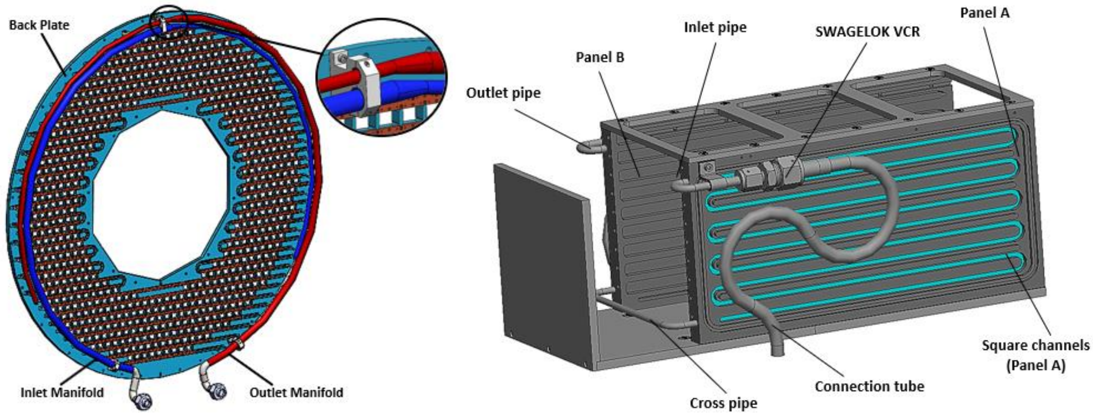

Figure 3.6 Left: FEE cooling manifolds and distribution lines. Right: back-end electronics crate cooling system.

As far as SiPMs and FEE boards are concerned, they are associated more strict requirements in terms of cooling and temperature homogeneity. The dark current of the photosensors will increase dramatically after irradiation, up to a design limit value of 2 mA (cf. Chapter 4 and Section 5.3). A target operating temperature of 0 °C has been chosen for the photosensors, in order to minimise their dark current. In addition, the cooling system should allow to decrease the operating temperature down to -10 °C, so as to preserve the photosensor performance after radiation exposure.

Moreover, although a 2 mm void gap is present between each crystal and its relative photosensor to reduce thermal coupling, the cooling system is still of crucial importance with respect to the CsI stability: as previously stated, excessive radiative heat transfer from the SiPM can result in thermal gradients of significant entity along the crystal axis, able to degrade its optical response and performance.

The FEE cooling system (Figure 3.7) circulates the 35% glycol-water mixture through pipes vacuum brazed to copper plates installed on the PEEK backplane. Each SiPM holder is also made of copper and is directly bolted on the cooling plates. The cover faraday cages are also in contact with the cooling plates. The cooling system is required to be free of fault and maintenance-free for a minimum operating time of one year.

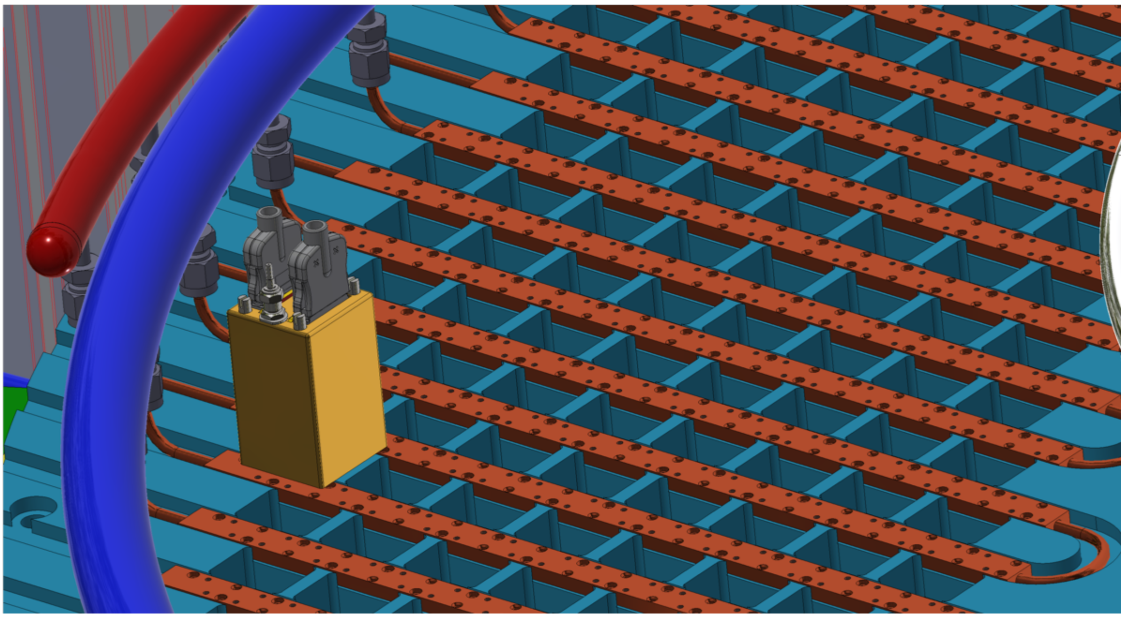

Figure 3.7 Detail of the FEE heat exchangers.

3.3 Calibration systems

The SES target level requires the Mu2e detection systems to be precisely and frequently calibrated in order to prevent the onset of systematic errors and to keep the operational performance extremely stable throughout the whole run-time period. Calibration is needed both for energy and timing measurements, and it relies on the systems and methods described in the following sections.

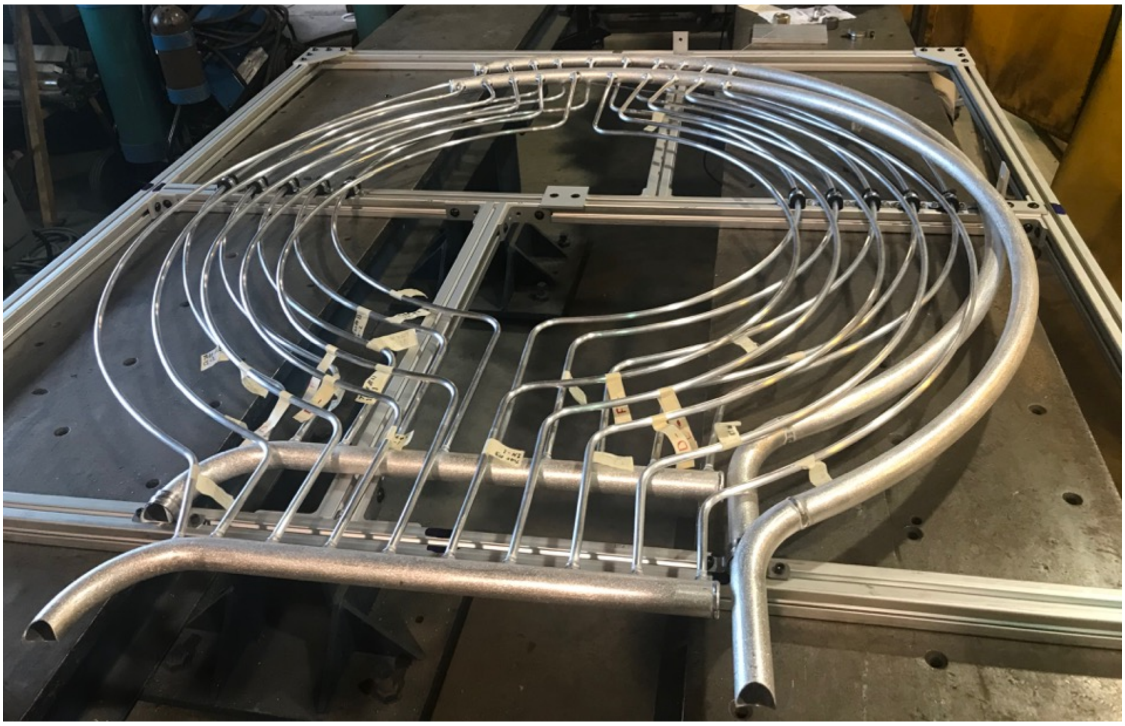

Figure 3.8 Liquid radioactive source circulation pipes.

3.3.1 Embedded calibration systems

3.3.1.1 Liquid radioactive source

The verification of the absolute energy scale of the calorimeter and a fast response equalisation between crystals is carried out by means of a liquid radioactive source, which is circulated through aluminium pipes installed on each disk cover plate. The Fluorinert (FC-77) based source fluid is activated via the production of the ^{16}N isotope, by means of irradiation with 14.2 MeV neutrons produced by an external deuteron-tritium generator with a flux of 10^8 n/s).

Calibration is carried out using the 6.13 MeV photons, produced via the hereby reported decay chain:

$$
\begin{aligned}
{}^{19}\mathrm{F} + \mathrm{n} &\rightarrow {}^{16}\mathrm{N} + \alpha \\
{}^{16}\mathrm{N} &\rightarrow {}^{16}\mathrm{O}^{*} + \beta^{-} + \overline{\nu_{e}} \; (\tau_{1/2} \approx 7\ \mathrm{s}) \\
{}^{16}\mathrm{O}^{*} &\rightarrow {}^{16}\mathrm{O} + \gamma(6.13\ \mathrm{MeV})
\end{aligned}
$$

The activated fluid is then pumped with a flowrate of 3.5 l/s in twelve 3/8 inch ID Al pipes (Figure 3.8). Their geometry allows a uniform illumination of the disk, with an intensity variation < 5% [19]. The 6.13 MeV photons deposits are reconstructed on a crystal as shown in Figure 3.9, which includes the full absorption 6.13 MeV peak and the two escape ones (at 5.62 MeV and at 5.11 MeV respectively), superimposed with the Compton spectrum. A rate of $\sim 10^4$ photons per 10 minutes is expected for a single crystal. An 1.5 % equalisation can be reached. A 10 minutes long calibration will be performed weekly.

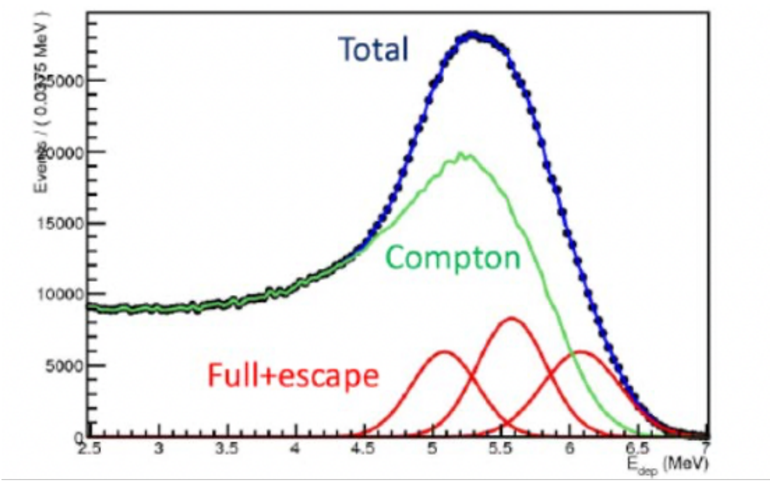

Figure 3.9 Spectrum for the energy deposited in a single crystal irradiated during the calibration cycle.

3.3.1.2 **Laser monitor**

A monitor system is implemented in order to verify gain, charge and timing resolutions of each photosensors, as well as to perform timing alignment between their response. It is based on a 530 nm laser source: this wavelength has been selected as it is not interested in the CsI emission spectrum and, most importantly, because the optical transmittance of CsI is not affected by irradiation at this frequency.

The laser is pulsed and focused by standard collimation optics into an optical 1-to-8 splitter; through the use of 8 1 mm ID, 60 m long quartz fibres, light is carried inside the DS. For each disk, light is injected into 8 Ulbricht integrating spheres, each of which is equipped with a monitoring photodiode and two output fibres bundles. Each bundle has 110 Si fibres, with a 200 μm core and an F-doped Si cladding. The total is of 880 fibres/disk. Each fibre is then inserted - via a custom needle - inside each SiPM holder, so that it can illuminate the crystals and the SiPMs.

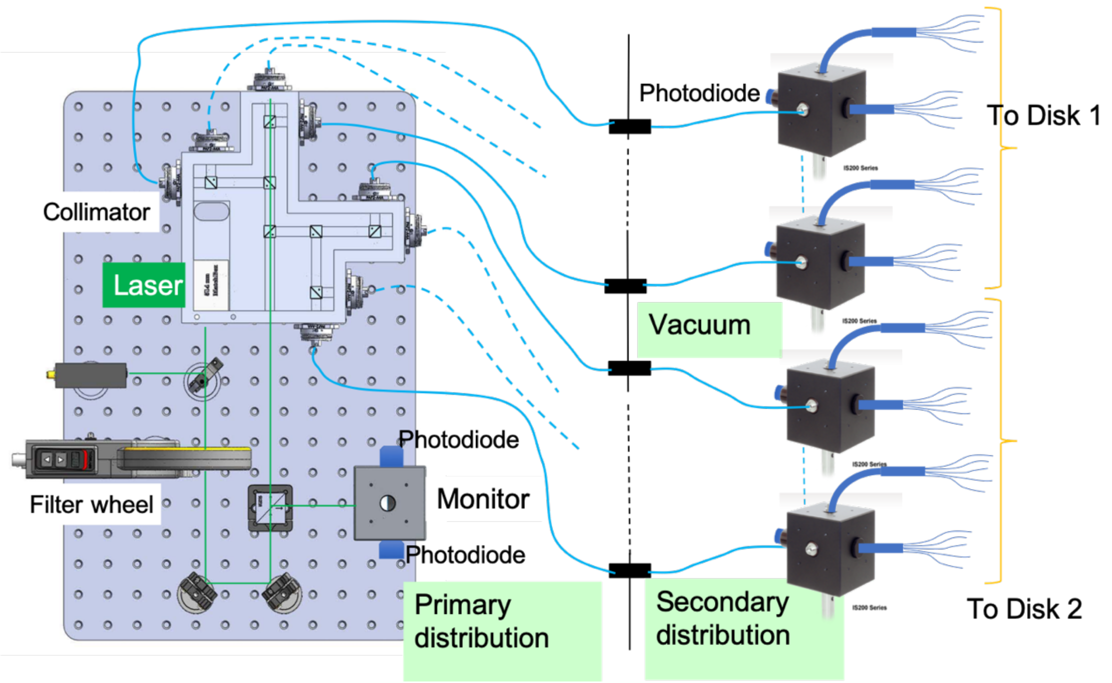

Figure 3.10 Block scheme of the laser calibration system.

A laser pulse equivalent to a 50 MeV energy deposition is generated during the spill-off of each injection period (~ 1.4 s). This laser calibration system is expected to reach a 0.5% gain equalisation level and a time-offset alignment within 25 minutes of operation.

3.3.2 Other calibration methods

3.3.2.1 DIO muons and π decays

Energy calibration below 100 MeV can be achieved by means of muons and pions decays, according to these two most promising channels:

- $\pi^+ \to e^+ + \nu_e$ which produces a monochromatic positron around 69.8 MeV;
- muon decays in orbit (DIO) and $\mu^- \to e^- + \nu_\mu + \overline{\nu_e}$ with a 52.8 MeV spectrum edge.

The tracker can reconstruct the momentum of such electrons and positrons with an accuracy better than 0.5%. Thus, the calorimeter can be calibrated with similar accuracy by matching its cluster with the respective track.

At the nominal 1 T B-field, these decays will illuminate only a reduced section of the calorimeter: dedicated runs at reduced B-field will be carried out to extend this calibration to the totality of the calorimeter crystals. During such runs, the detector hits rate increases, and the tracker resolution deteriorates, making it necessary to scale down beam intensity.

In the case of pion decays, some adjustments to the Mu2e beam line are required. The B-field must be reduced to 0.7 T and a rotation of the TS collimator is needed for selecting positive particles, as well as a modified acquisition window. A degrader has to be installed before the stopping target to increase the number of stopped muons, in order to cope with the low BRs associated of this decay.

The calibration with DIO electrons, instead, can be performed without any modification to beam line and can be carried out in shorter times. The only requirement is a reduction of the B-field to 0.5 T and of the beam intensity to half its nominal value. These calibration and equalisation data at reduced B-fields will be extrapolated to 1 T.

3.3.2.2 **Cosmic rays**

CR events are acquired using a dedicated trigger during the normal experimental conditions. The CR calibration can be carried out through the reconstruction of path length and energy deposit of minimum ionising particles. Figure 3.11 shows the deposited energy spectrum of cosmic muons in a single CsI crystal: the MIP peak can be identified at ~ 20 MeV. The expected achievable equalisation level is of ~ 1% for such energies. In addition, MIP transit time can be used to align the relative time offsets of all channels to a maximum expected deviation of ~ 50 ps and without the need for an external reference (cf. Section 6.1)

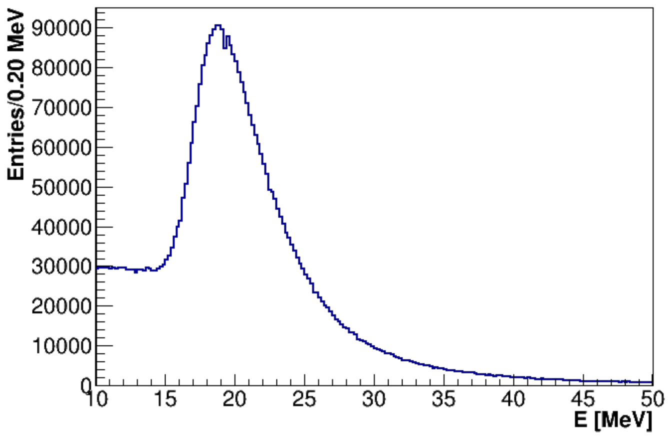

Figure 3.11 Spectrum of the deposits of cosmic muons in a crystal.

3.3.3 **CRT pre-installation calibration at SiDet**

Before the final installation in the Mu2e facility, each calorimeter disk will be individually tested and calibrated using MIPs. Calibration runs will be carried out by means of the Cosmic Ray Tagger described in Chapter 6. This calibration will provide the reference performance values for the instrument and will include the following tasks:

- calibration of the absolute energy response over the expected 21 MeV MIP deposit;
- equalisation of the channel responses below the 1% level;
- T_{0_i} alignment at a level below 10 ps (RMS);
- LRU check at 21 MeV;
- study of timing resolution dependence on crystals z-axis.

The CRT system and the calibration procedure are discussed in detail in Chapter 6.

3.4 Monte Carlo study

The hereby reported simulated performances were obtained via Monte Carlo methods using Geant4 [21]. The deposited energy is converted into optical photons, based on a yield of 30 pe/MeV and an LRU of a few percent points. SiPM response was simulated including 150 keV equivalent gaussian electronic noise as well as a Poissonian photostatistic fluctuation for each simulated energy deposit. A fully digitised waveform is extracted from the number of photoelectrons produced, also simulating pile-up phenomena.

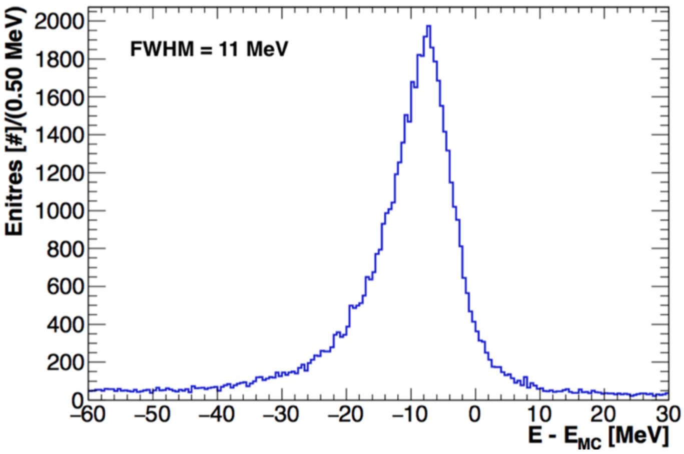

Figure 3.12 Energy resolution residuals.

Clusters are determined with an algorithm seeded by the location of the crystal with the largest reconstructed energy. Pixels are then individually added to the cluster, based on the following rules: they should share a common side with previously included crystals; their timing should be within $\pm$ 10 ns from the seed hit; their reconstructed energy should be significantly larger than the expected electronic noise.

3.4.1 Energy, position and time resolution

Energy resolution has been determined by simulating CEs together with the expected background. The distribution of the reconstructed energy residuals is plotted in Figure 3.12, where the presence of a left tail due to leakage can be noted. No significant contributions due to pileup are present. The energy resolution has been evaluated as $FWHM/2.35 = (3.8 \pm 0.1)$ MeV, or 5 MeV

RMS. A dedicated study (Figure 3.13) demonstrated that an LRU $< 5\%$ and a LY > 30 pe/MeV are sufficient to achieve good resolution.

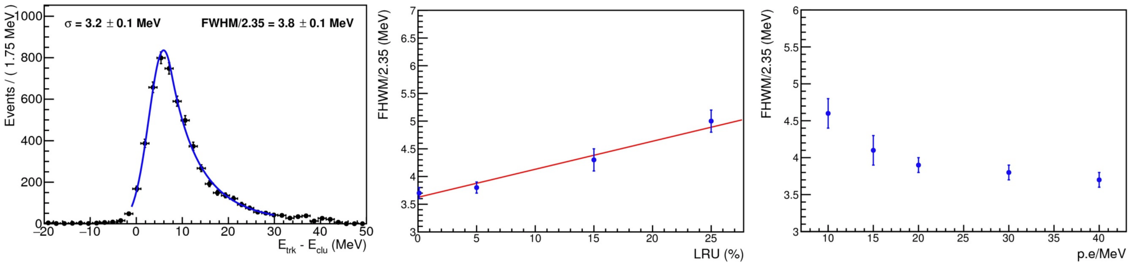

Figure 3.13 Left: double-sided Crystal Ball function fit of the energy resolution. Middle: FHWM/2.35 as a function of the LRU. Right: FHWM/2.35 as a function of LY.

As previously stated, cluster position is essential to match the reconstructed tracks with the calorimeter hits. Using a linear energy weighted mean, hit positions were extracted from the simulations: the distribution of the residuals of the predicted and effective position of the track showed that an xy resolution of ~ 6 mm is possible [1].

Timing was extracted by means of a fit with an analytical function and a constant fraction discrimination (CFD), assuming a constant pulse shape. The cluster time was then calculated as the linear energy weighted over the entire cluster. A Gaussian fit applied to the residuals showed an expected value for the time resolution of approximately 110 ps at 105 MeV (Figure 3.14).

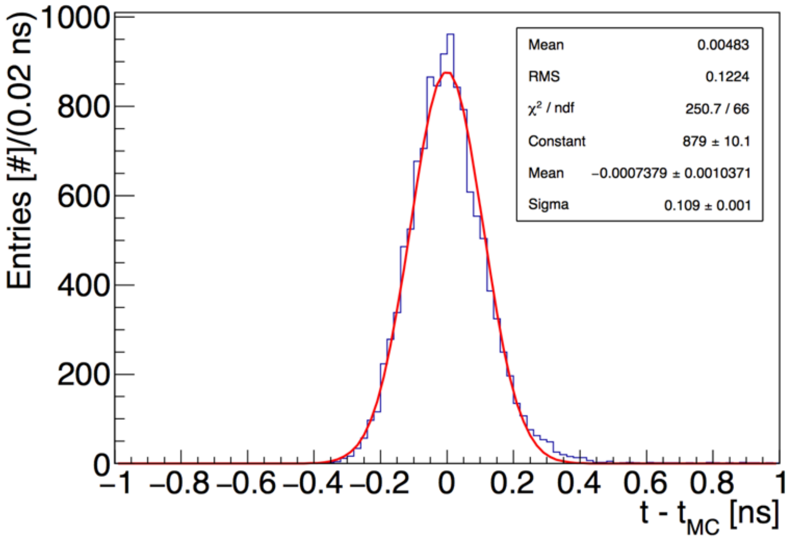

Figure 3.14 Residuals distribution relative to the timing reconstruction.

3.4.2 Particle identification

Extrapolated trajectories from the tracker and calorimeter hits are correlated in time in order to perform particle identification (PID) and optimise the pattern reconstruction algorithm. The cluster reconstruction algorithm makes use of the timing information provided by the calorimeter for connecting/rejecting cells and for cluster merging. Cluster seeding - carried out by selecting tracker hits which are within ± 50 ns with respect to the calorimeter hits and are compatible with the cluster azimuthal angle - increases the tracking efficiency and speed (Figure 3.15). The impact timing extracted from the helix extrapolated from the tracker measurement should be within 0.5 ns from the actual calorimeter hit: this requirement drives the demand for a good timing resolution of the latter.

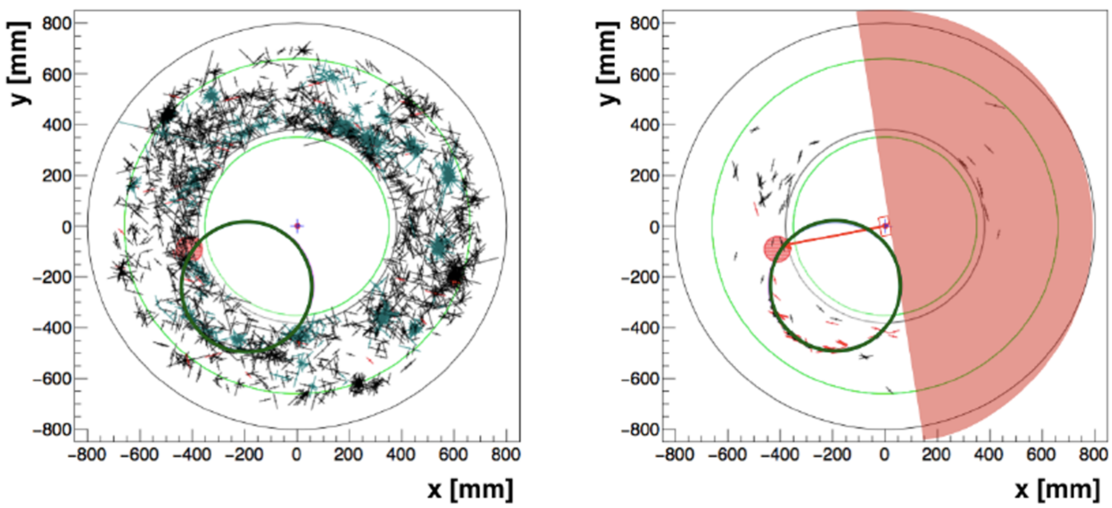

Figure 3.15 Tracker hits before (left) and after (right) the application of the calorimeter selection cuts.

Timing correlation is also needed for efficient PID, which is necessary for the discrimination between electrons and muons, either beam or cosmic related. The latter can either be trapped in the DS B-field or interact with the materials producing electrons. Assuming an inefficiency of about 10^{-4} for the CRV system, simulations showed that during a 3-year run-time the expected number of cosmic muons with $103.5 < p < 105$ MeV/c which are able to enter the DS undetected by the CRV and to survive the analysis cuts is around 3. Thus, to keep this background at a level inferior to 0.01 events, it is required to have a muon rejection factor better than 200. Figure 3.16 illustrates how the timing and the E/p ratio distributions can be used to discriminate 105 MeV/c electrons from muons [1].

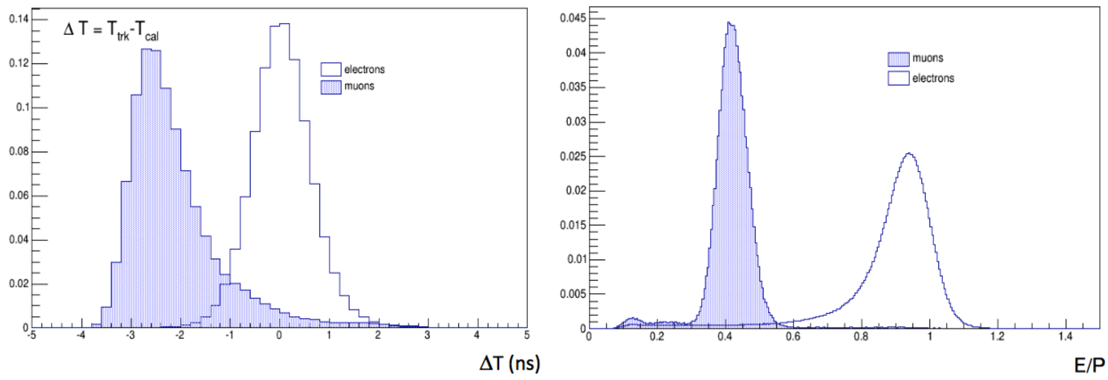

Figure 3.16 Left: time differences between the tracks and the clusters. Right E/P for CEs and muons.

The acceptance of the calorimeter has been optimised to detect $(99.4 \pm 0.1)\%$ of the CEs generated on the stopping target, by selecting only clusters which have an energy > 10 MeV and which are correlated to a track compatible with the necessary quality cuts [21].

3.5 Test beam on a large-scale prototype

The hereby reported tests have been carried out on a large-scale prototype, referred to as the Module-0, which was developed at LNF-INFN (Figure 3.17). It is composed by a matrix of 51 staggered crystals, wrapped in 150 µm Tyvek and assembled according to the final Mu2e calorimeter design (Figure 3.17). Each crystal is readout by two custom Mu2e SiPMs with prototypal FEE boards. The signals of the Module-0 were acquired using two CAEN V1742 [22], while an 12-bit ADC was used for the digital conversion. The digitiser was operating with a $0 \div 1$ V range and at 1 Gsps over an ~ 1 µs acquisition window.

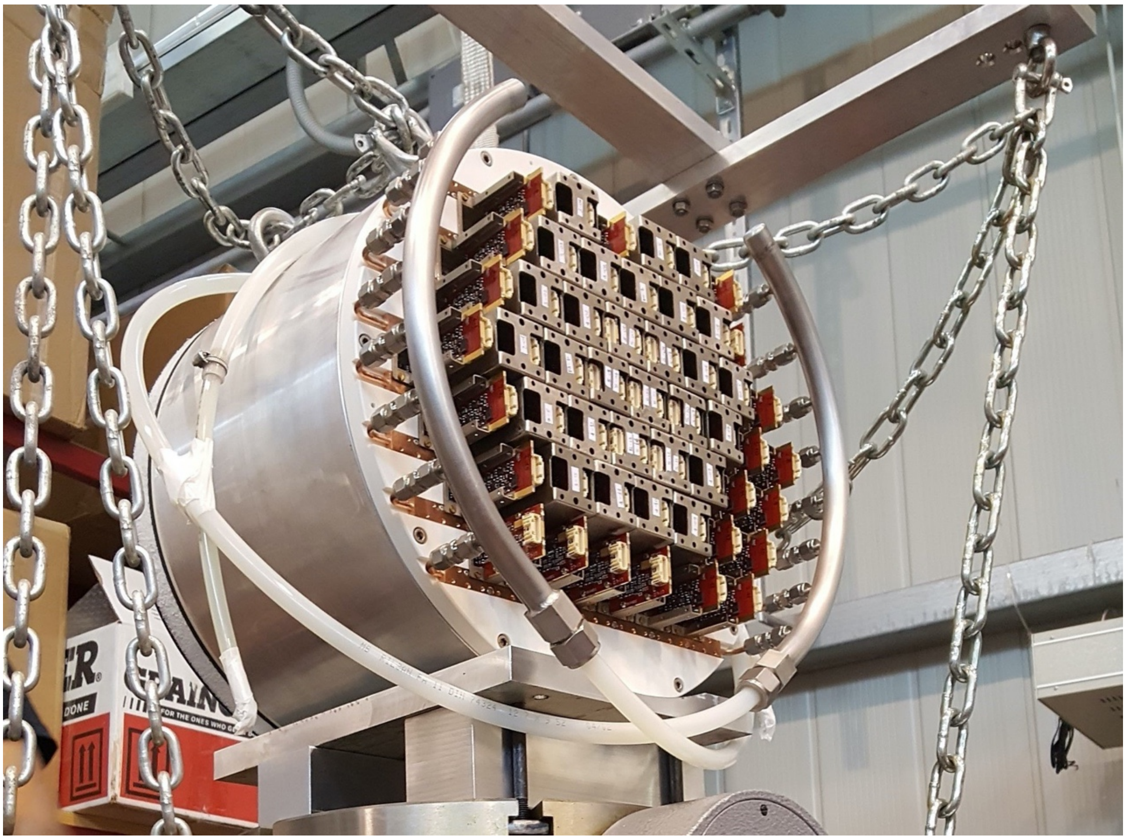

Figure 3.17 Module-0 inside the BTF hall at LNF.

3.5.1 Setup

The Mu2e calorimeter design was validated by performing a test beam on the Module-0 at the LNF Beam Test Facility [23] using an electron beam with a $60 \div 120$ MeV energy. The beam energy spread was evaluated to be O(2%) at 100 MeV.

Two small plastic scintillators crossed at 90° were positioned on the beam axis before the Module-0, thus providing a trigger system also able to count particle multiplicity. By tuning the beam intensity into single-particle configuration, ~ 0.7 particles/bunch were obtained with a bunch rate of 50 Hz. In addition, a large plastic scintillator was placed above the calorimeter and readout on two sides via PMTs, so as to track CR events.

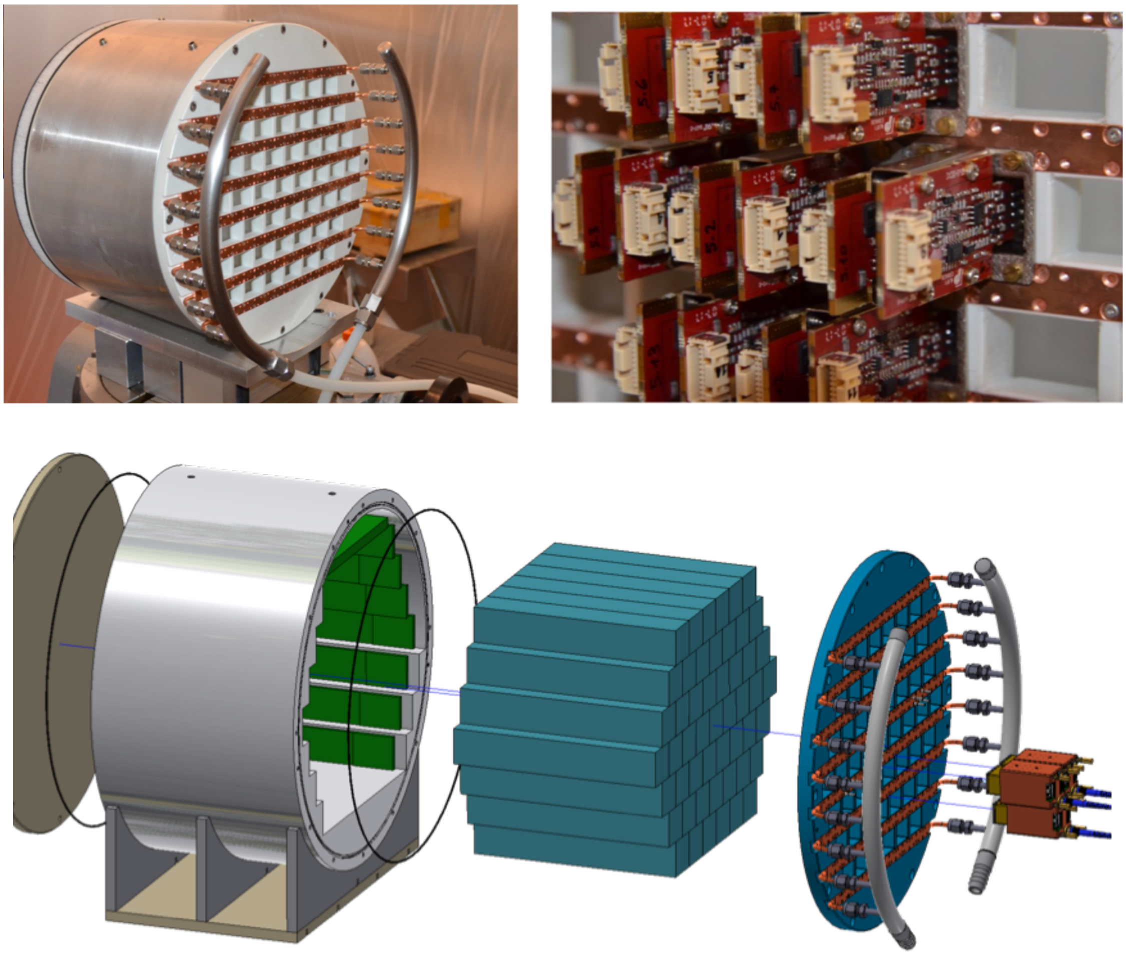

Figure 3.18 Top: Module-0 cooling system (left) and detail of the FEE boards (right). Bottom: exploded view of the Module-0. From left to right: cap plate, chassis, crystals stack, backplate with cooling lines and SiPM/FEE modules.

SiPMs and FEE boards were thermalized using the Module-0 integrated cooling lines (Figure 3.18) by means of an external chiller. A laser system (cf. section 3.3.1.2) was used to hourly monitor the SiPM response stability during the run-time. The maximum observed absolute deviation in the response to the laser pulses was lower than 3%.

Data were collected over the course of one week using a trigger scheme which included the following signals in OR:

- beam trigger (BT) produced by the coincidence of the discriminated beam counters;
- trigger (BTF) provided by the BTF itself;
- cosmic ray trigger (CRT) generated by the coincidence of the top plate scintillator PMTs;
- a syncing signal from the Laser system (LT) for monitoring purposes.

Tests were carried out with two geometrical configurations, with the electron beam impinging on the Module-0 centre at 0° and 50° angles. The latter configuration reproduces the expected CE incidence angle.

3.5.2 Charge and LY

Numerical integration was performed over the waveforms using a 250 ns gate in order to perform charge reconstruction. A sample waveform obtained with a laser pulsed in the centre crystal was integrated over different time windows, to assess whether any change to the integration gate width could degrade the performance of the matrix. Figure 3.19 reports the resulting resolutions, which were calculated as the relative standard deviations for a gaussian fit on charge distributions.

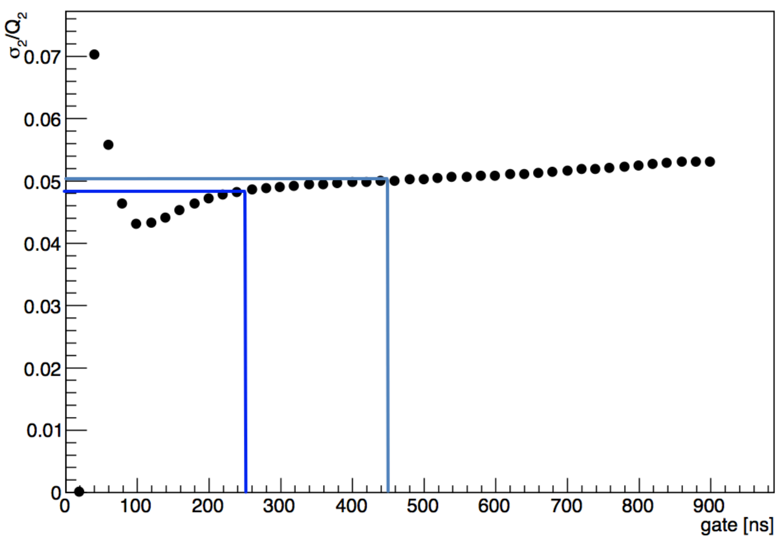

Figure 3.19 Charge resolution versus integration gate time from laser runs.

The number of photoelectrons emitted per MeV was evaluated with 100 MeV electrons and using the chosen 250 ns integration window, resulting in a value of approximately 22 pe/MeV.

3.5.2.1 **Single electron event selection**

In order to correctly perform time and energy measurements, a more accurate offline selection of single particle events is needed, as the electrons from the BTF could still have a multiplicity greater than one, even after proper intensity tuning and collimation. Events were rejected if coincident with the CR trigger or a laser calibration pulse and if a single particle event was not reported by the beam counters.

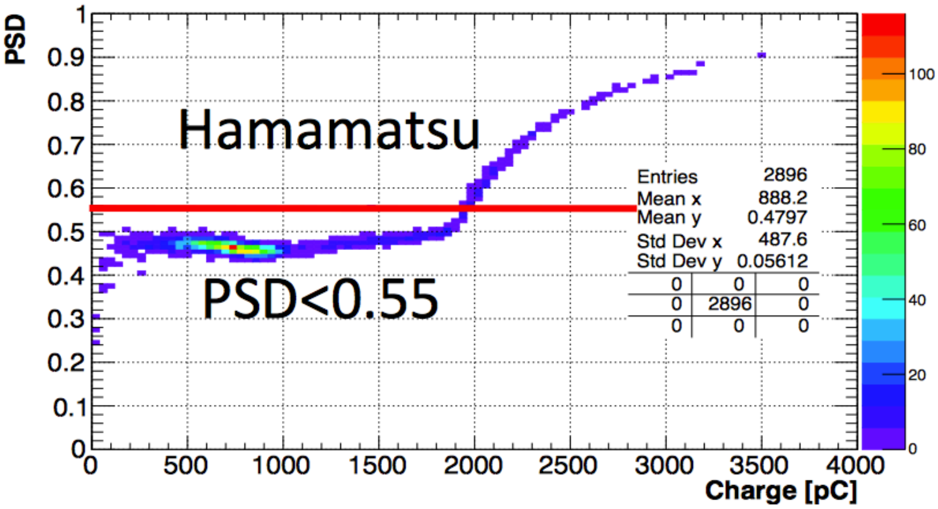

Figure 3.20 PSD distribution versus reconstructed charge for Mu2e SiPMs.

Moreover, a pulse-shape discrimination (PSD) was applied on each crystal waveform $w(t)$ in order to discard channels which presented saturation due to pileup:

$$PSD = \frac{\int_{t_1}^{t_2} w(t)\, dt}{Q_T} \tag{3.1}$$

where t_1 and t_2 are the time positions which correspond to the 10% and 90% of the maximum pulse height, respectively for the rising and falling edges. The total waveform charge is instead represented by Q_T.

Data were collected for a 100 MeV beam in the orthogonal configuration. The application of all these cuts (Figure 3.20) reduced the sample to ~ 1600 events, or ~ 2% of the original dataset.

3.5.3 Energy scale and equalisation

The offline response equalisation between channels was carried out using energy depositions from the TB and CR events (Figure 3.21). In the first case, a 100 MeV electron beam was fired orthogonally at each crystal and an asymmetric log-n was fitted to the charge distribution of each channel: the charge peak value was used for the evaluation of the calibration factor of each crystal with respect to the central one. The statistical error of this procedure was below 1.5% for all channels. The cosmic ray equalisation was carried out selecting MIP events and the relative energy deposit. The calibration factors were again calculated with respect to the central crystal, with an associated statistical error below 0.5% for all channels. The presence of any systematic effect was verified by plotting the distribution of the ratio between the calibration factors relative to the two calibration methods: the average is well centred around 1, with a sigma of around 3% (Figure 3.21, left).

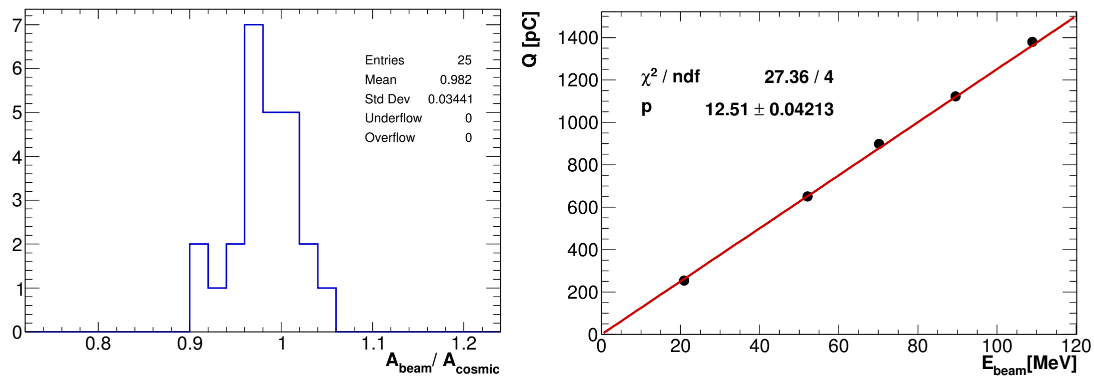

Figure 3.21 Left: ratio of the calibration factors resulting from the TB to the ones relative to CRs. Right: reconstructed charge as a function of the expected energy deposit for different beam energies.

The equalised total charge was evaluated by summing up the contributions of all channels presenting a charge greater than 3 times the noise floor and using the calibration factors evaluated from CR data. The response was linear for both beam configurations and the energy scale factor of (12.51 ± 0.04) pC/MeV was evaluated for the orthogonal run (Figure 3.21, right).

3.5.4 Energy resolution

The energy resolution $\sigma(E)/\langle E \rangle$ was calculated via the log-n distribution fitted over the energy distributions. An optimisation algorithm was developed to select the best fit range by minimising the χ^2.

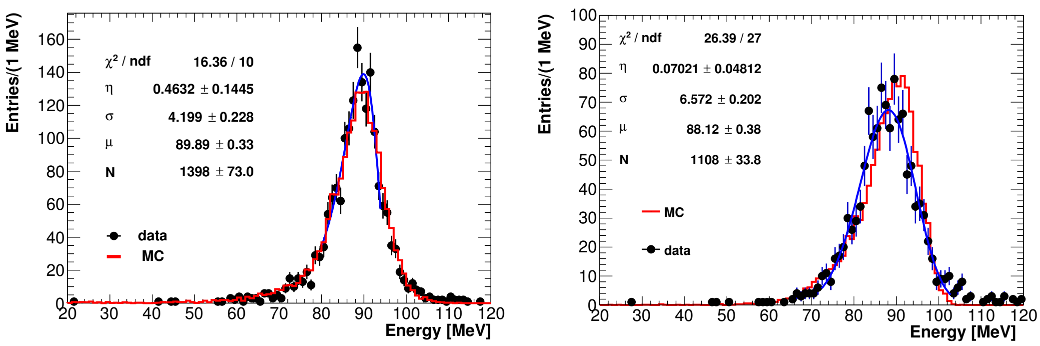

Figure 3.22 Energy reconstructed on the calorimeter for the 100 MeV electron beam impinging at 0° (left) and 50° (right). A good data-MC agreement was found.

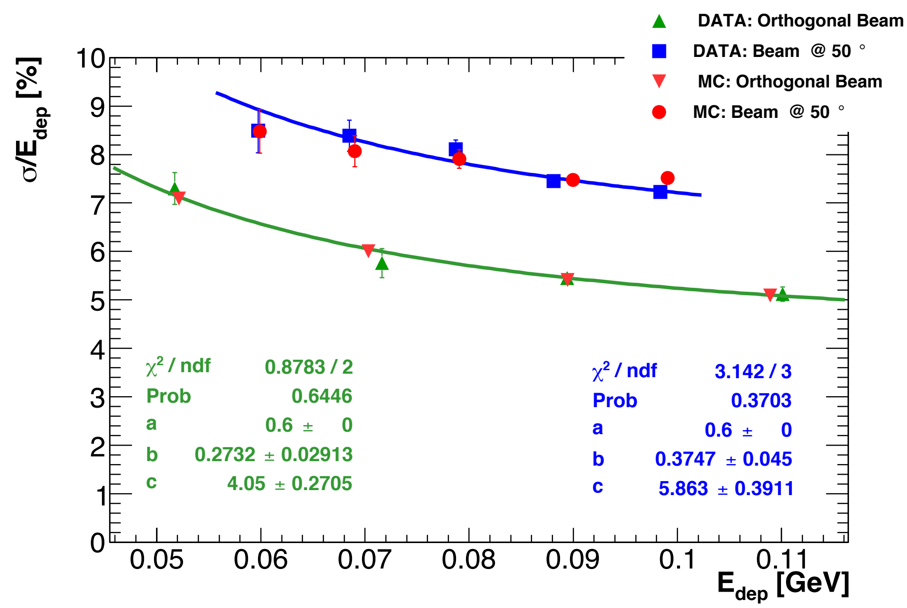

Figure 3.23 Energy resolution as a function of the energy deposited in the Modue-0 for the orthogonal (green) and tilted (blue) runs.

The values obtained for the energy resolution are 5% and 7%, respectively for the 0° and 50° beam angle runs (Figure 3.22). The observed non-gaussianity is due to leakage - especially for the orthogonal run, due to the inferior shower containment - and energy loss in passive materials, as well as in the beam scintillators. As expected, the majority of the energy deposition happened in the first two innermost crystal rings.

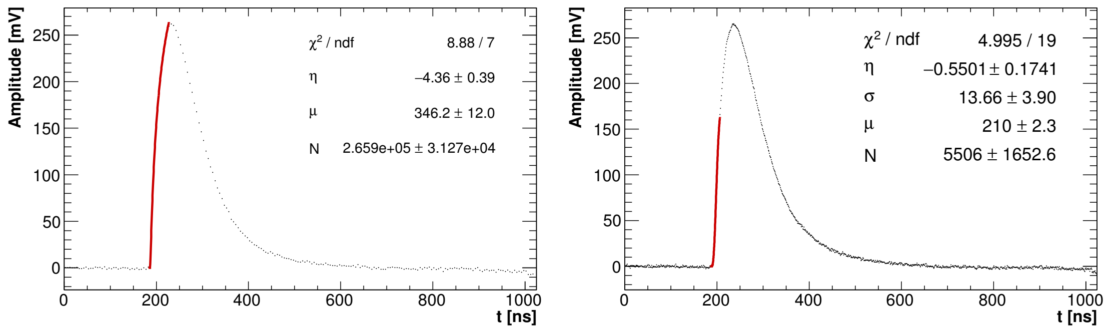

Figure 3.24 Optimised fit applied to waveforms sampled at 1 GHz (right) and offline resampled at 200 MHz (left).

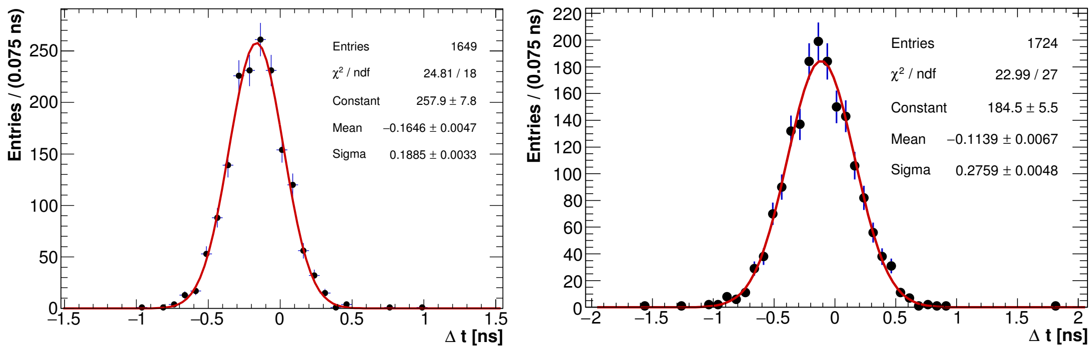

Figure 3.25 Time resolution of the central crystal. Left: 1 GHz sampling rate. Right: offline resampling at 200 MHz.

3.5.4.1 **Energy resolution deterioration**

The energy resolution dependence over different beam energies shown in Figure 3.23 has been modelled with the following relation:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E[\mathrm{GeV}]}} \oplus \frac{b}{E[GeV]} \oplus c \tag{3.2}$$

where the coefficients a, b and c account respectively for the stochastic, noise and constant terms [24]. The stochastic term contribution is due to the fluctuations in the shower development. The noise term is related to the signal chain and is due to electronic noise. The constant term accounts for deteriorations not dependant on the particle energy, such as instrumental effects associated with non-uniformities in the calorimeter response.

The fit reported in Figure 3.23 was rather insensitive to the stochastic term (which is generally really small in homogenous calorimeters), so that it was fixed to 0.7%, as provided by the 20 pe/MeV light yield. In the case of tilted runs, a worsening of the fluctuations associated to lateral leakage term occurs.

3.5.5 Timing reconstruction

As the pulse shape is preserved for different amplitudes, the timing information was retrieved by fitting an asymmetric log-n function on the rising edge (Figure 3.24) and by subsequently performing a CFD. The procedure was optimised by trimming the fit range and the CF threshold, which was set to 5% of the pulse-height. No significant correlation was found between the reconstructed time and the associated bin position, thus confirming the absence of systematic effects and a uniform distribution over the bin interval. Waveforms were resampled offline at 5 ns to verify the performance at the Mu2e 200 Msps acquisition rate (Figure 3.24, right).

3.5.5.1 Time resolution and deterioration

Time resolution was evaluated from the time difference between the signals relative to the two SiPMs reading the same crystal (Figure 3.25). Both CR and TB data were used for the analysis. From the distribution plotted in Figure 3.25, the σ_t was deduced dividing the standard deviation from a gaussian fit by $\sqrt{2}$, so as to take into account the contribution of the two individual photosensors. In the case of orthogonal beam at 100 MeV, the resulting values were 132 ps and 195 ps, respectively for the 1 GHz acquisition frequency and for the 200 MHz resampling.

For a single channel, the time resolution dependence over the deposited energy is modelled by the following relation:

$$\sigma_T = \frac{a}{E[\text{GeV}]} \oplus b \tag{3.3}$$

The parameter a is proportional to the optical emission time constant of pure CsI, while b represents additional contributions due to jitter in the readout electronics. Figure 3.26 reports the resolution as a function of the highest energy deposit in a single crystal, for both sampling frequencies. A deterioration of less than 30% is extrapolated after resampling.

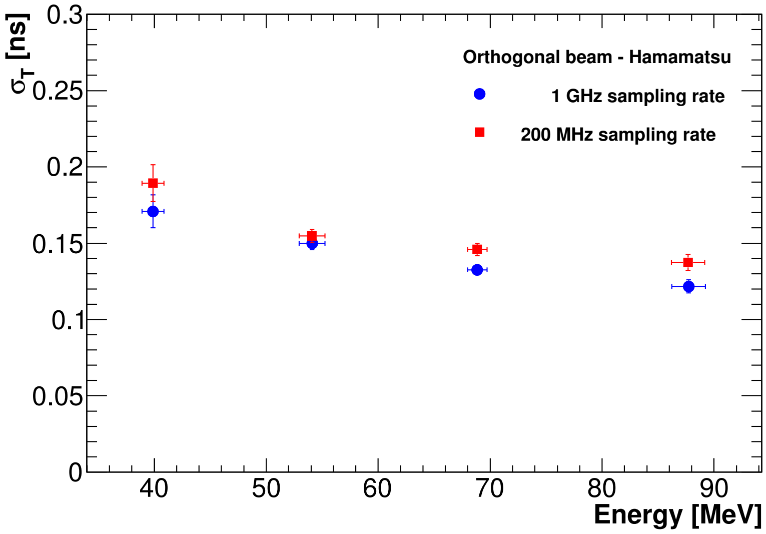

Figure 3.26 Time resolution as a function of the highest energy deposit in a single crystal before and after offline resampling.

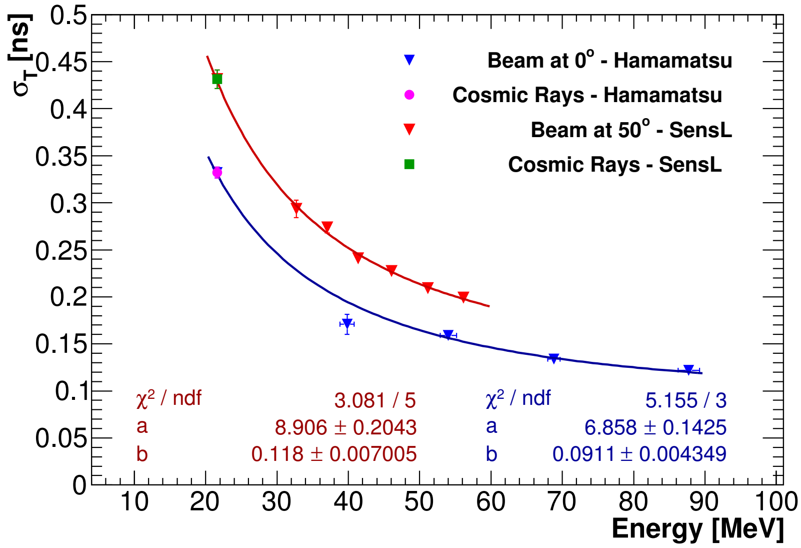

Figure 3.27 Time resolution as a function of the highest energy deposit in a single crystal. Both CR and TB events are shown.

The results of the TB demonstrated that the CsI option, along with the custom SiPM readout completely satisfied the Mu2e requirements.

3.6 Template fit

By using a large area scintillator on top of the Module-0 as a trigger source, the seven innermost channels were used to acquire MIP events to be processed using a waveform template fit algorithm for the reconstruction of the peak time position, as a way to overcome the limited 5 ns sampling interval of the Mu2e DAQ system, so that better precision and faster processing times can be obtained in the TOF analysis. The same algorithm will also be used in section 6.3.5. For this test, CR data were digitised at 250 Msps.

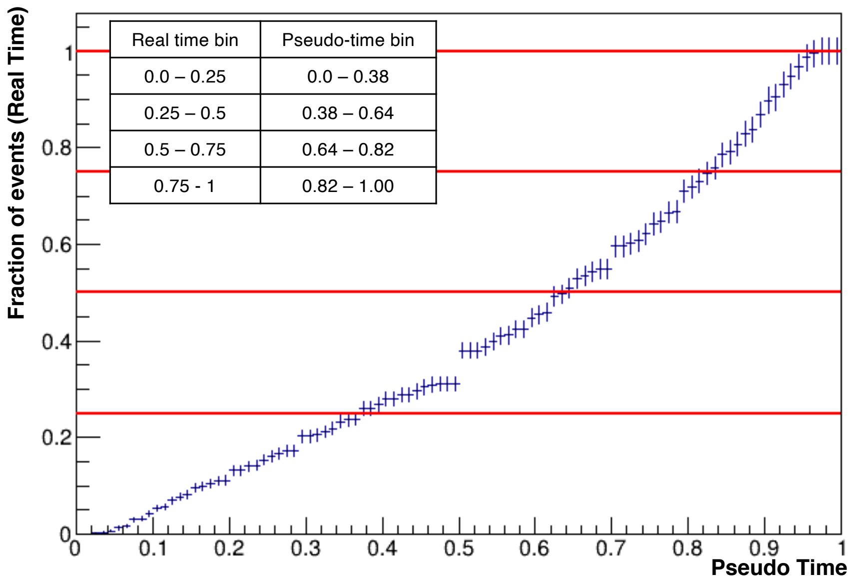

Figure 3.28 Mapping of the pseudo-time and real time cumulative distribution.

3.6.1 Template construction

A waveform template was constructed based on the cosmic ray data for each channel. At first, the evaluation of the pseudo-time variable t_p was performed using the following relation:

$$t_p = \frac{2}{\pi} arctan \left(\frac{s_0 - s_-}{s_0 - s_+} \right) \tag{6.7}$$

The parameter s_0 represents the signal value associated to the bin containing the sampled waveform peak, while s_- and s_+ represent the signals contained in the adjacent bins, respectively in the previous and in the following one. The pseudo-time is not linearly correlated to the actual real peak position t_r, the latter defined as the relative position of the true waveform peak relative to the bin in which it is contained and normalised over the bin width. If the MIP hits are (reasonably) assumed to have a uniform distribution in time, the pseudo-time t_p should vary smoothly from 0 to 1 along with the cumulative distribution of the real times t_r (which is assumed to grow linearly over the bin width). The events can thus be mapped with respect to t_p, so that any arbitrary fraction of events F having $0 < t_p < t_{p0}$ contains events which must have a real time between 0 and $t_{r0} = t_f$ (where t_{p0} is an arbitrary value between 0 and 1). This means that a one-to-one correspondence exists between t_r and t_p. The results of this process are reported in Figure 3.28.

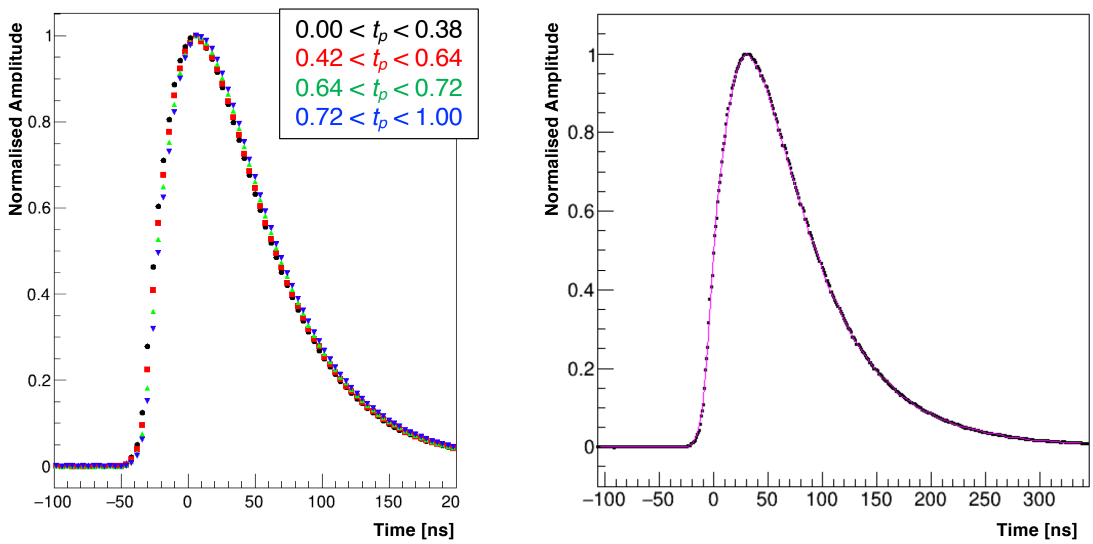

Figure 3.29 Left: template construction using the four averaged waveforms used for the template generation. Right: template constructed using the TSpline3 ROOT class.

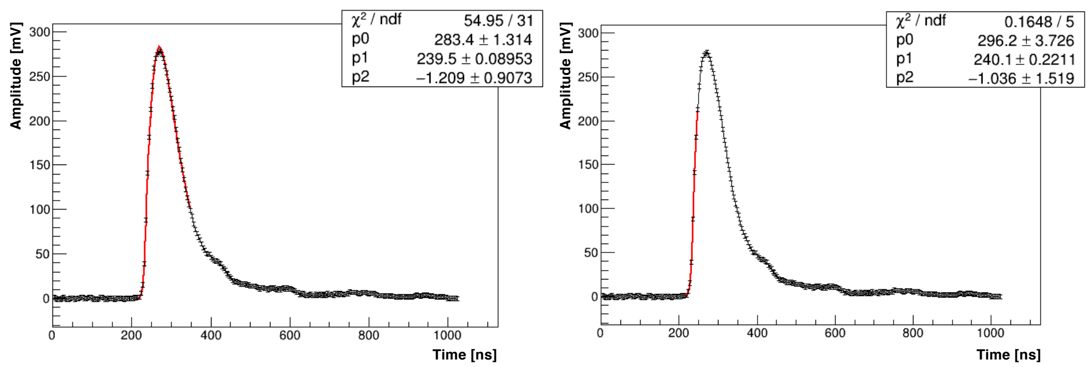

Figure 3.30 Fit examples for two different fit ranges.

This distribution was then equally split into 4 bands with respect to the real time variable t_r, and events were split accordingly into an equivalent number of sub-sets. The ensuing mapping between t_p and t_r with respect to the thus obtained four bins is also reported in Figure 3.28. By summing and normalising the waveforms contained in each one of this bins, four averaged waveforms were obtained. These waves were progressively shifted of ¼ of a bin's width each and combined together. The resulting 1-ns binned waveform (Figure 3.30, left) was then interpolated using a TSpline3 ROOT object [54], resulting in the final waveform template which embeds the superposition of the 4 waveform types associated to the different locations of the SiPM pulses with respect to the sub-binning of the main sampling interval (Figure 3.29, right). The choice of 4 of sub-divisions for the 4 ns sampling time was chosen after the careful analysis of the timing reconstruction accuracy versus the number of sub-bins used in the construction of the waveform template: this choice represents an optimal trade-off between resolution and processing times. In fact, it was verified that, for this setup, a further increase in the number of sub-bins would not yield a significant improvement in the timing resolution.

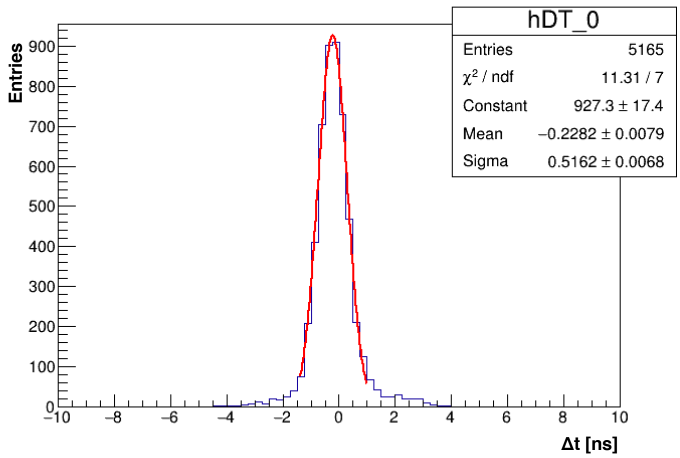

Figure 3.31 Time difference distribution example.

3.6.2 Template fit performance

Fit examples carried out on the rising edge of a MIP pulse are reported in Figure 3.30. The template fit procedure over the sampled waveforms can thus be carried out through the optimisation of 3 parameters: peak position (p0), the pedestal (or offset) value (p1) and a normalisation factor (p2). The fit range was determined by an optimisation procedure carried out in order to minimise the time resolution, evaluated as the RMS spread of the time difference distribution relative to two adjacent SiPM channels which read the same crystal (Figure 3.31).The distribution shown in Figure 3.33 was built to map the correlation of the upper and lower bounds of the fit range with the resulting resolution.

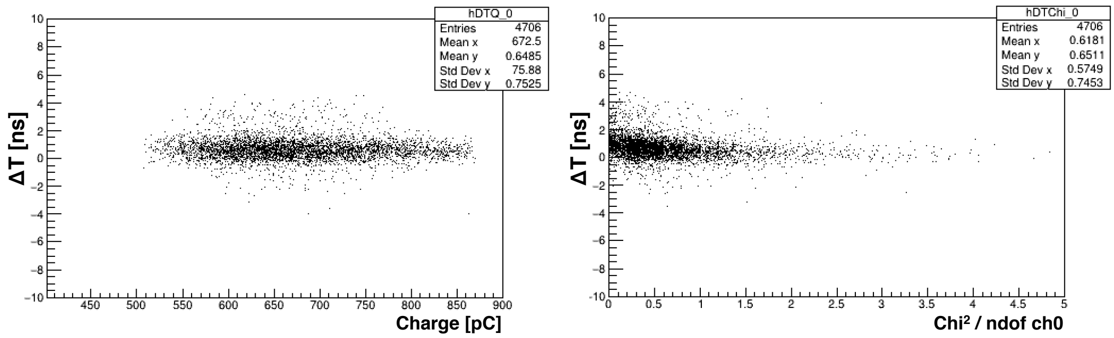

Figure 3.32 Time difference distribution versus reconstructed charge and fit goodness.

For the cosmic ray dataset, no slewing was observed in the behaviour of the time difference distribution with respect to both the reconstructed charge and the χ^2 (Figure 3.32). Finally, the resolutions obtained with the template fit algorithm are in good agreement with the Module-0 TB results reported in the previous sections.

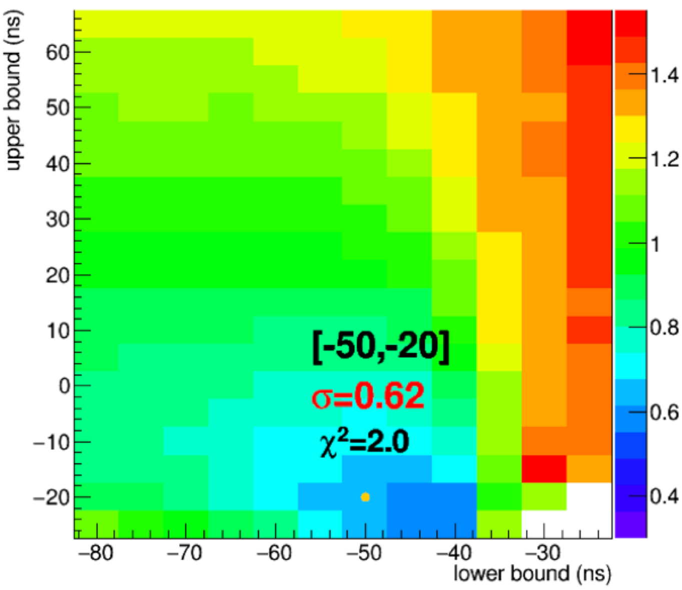

Figure 3.33 Map of the obtained resolution for different fit range bounds (with respect to the peak position). All values are in ns. The yellow dot represents the chosen optimal parameters.

4 Electronics and signal chain

4.1 Silicon photomultipliers

Solid-state photomultipliers - equivalently known as Silicon Photomultipliers (SiPM) or Multi-pixel Photon counters (MPPC) - are photodetectors capable of sensing, timing and quantifying light pulses down to the single-photon level. MPPCs rely on arrays of microcells (pixels), composed by single photon avalanche diodes (SPADs), as shown in Figure 4.1.

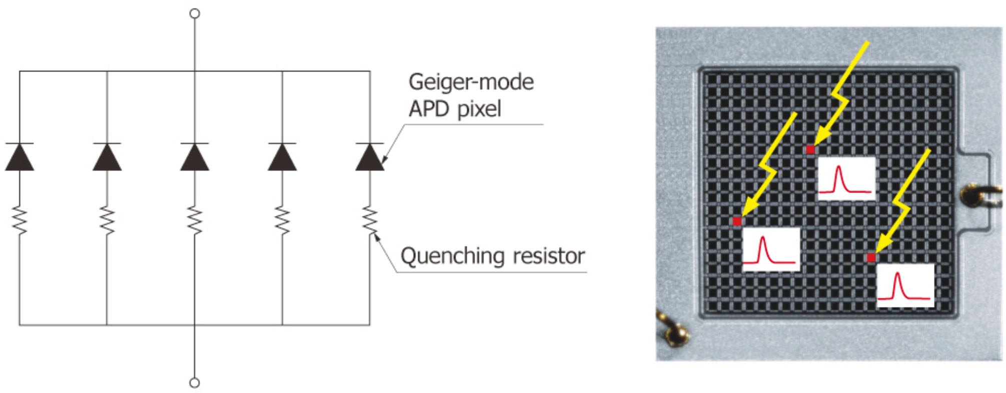

Figure 4.1 Schematic representation (left) and photo (right) of the individual SPAD microcells which compose an MPPC sensor

4.1.1 SPAD microcell

Single photon avalanche diodes (SPADs) are p-n junctions specifically designed to operate with a reverse bias above the breakdown voltage: in the presence of a sufficiently intense E-field (generally greater than $5 \cdot 10^5$ V/m), a single charge carrier injected or generated in the depletion layer will gain sufficient kinetic energy to initiate a self-sustaining avalanche process in which secondary charge pairs are generated through successive impact ionisations, resulting in an ionisation cascade able to spread through the silicon volume interested by the E-field (Figure 4.2, right). The ensuing junction breakdown will thus amplify the initiating event - such as the absorption of a single photon - into a macroscopic photocurrent flow, through a process known as Geiger discharge [39]. The discharge is generally quenched through a current-limiting series resistor (passive quenching), whose voltage drop during the discharge will restore a reverse voltage close to the breakdown one across the SPAD, thus halting the avalanche process. The diode is then recharged to the operating voltage trough the same resistor and is ready to detect subsequent photons.

This breakdown, avalanche, quenching and recharge cycle allows Geiger-mode APDs to operate as photon-triggered switches, which produce output pulses independent of the initial number of simultaneously absorbed photons (and thus not proportional to the instantaneous photon flux magnitude). This lack of proportionality is overcome by integrating parallel single SPAD elements with embedded quenching resistors into a dense array (100 to 1000 cells per square mm) of independent identical microcells or pixels (Figure 4.1, right). Each pixel detects photons identically and independently, while allowing the containment of the Geiger discharge in the single microcell it was initiated in. The sum of each individual photocurrent combines into a pseudo-analogue output, capable of providing information relative to the instantaneous photon flux magnitude.

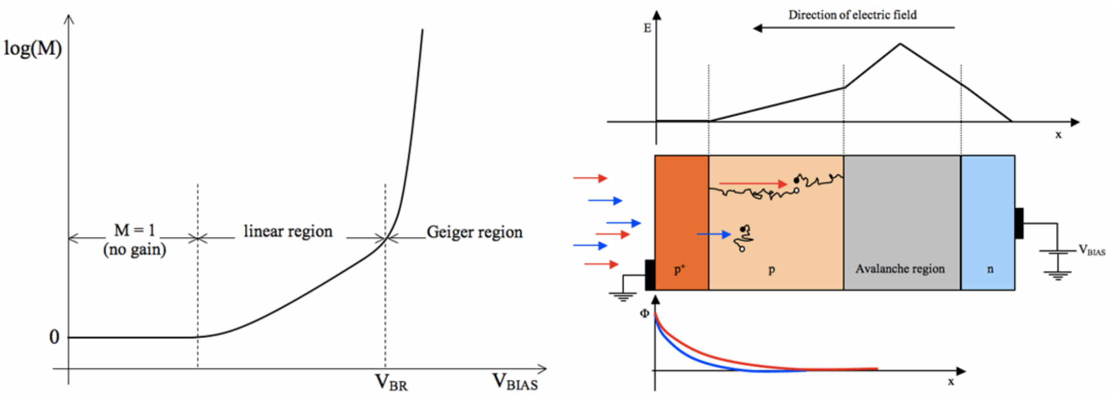

Figure 4.2 Left: multiplication factor behaviour for the different operating regimes of an ideal SPAD. Right: schematic representation of a p^{+}-p-n SPAD.

4.1.2 Avalanche and breakdown voltage

In Figure 4.2 (left) the relationship between the bias voltage and the multiplication factor of an ideal SPAD is illustrated. The threshold between the region of unitary multiplication factor (in which the APD is operated as a photodiode) and the linear region demarks the point for which the electric field in the avalanche region is able to provide charge carriers with minimum energy required for impact ionisation of the lattice atoms and thus initiate avalanche multiplication. The breakdown voltage V_{br} is defined as the minimum bias required for the initiation of a Geiger discharge and it is identified as the voltage for which a divergence in the multiplication factor is observed (or equivalently in the current–voltage characteristic, cf. Section 4.2.1.1), which demarks the two regions of linear multiplication and self-sustained avalanche. The breakdown voltage value is dependent on the junction structure and doping, as well as on temperature. SiPMs are generally operated with a reverse bias typically $10 \div 25\%$ higher than the breakdown voltage. The difference between the bias point and the breakdown voltage is referred to as overvoltage V_{ov}.

According to the drift-diffusion model [40], the number of electron-hole pairs generated per unit distance travelled by a single charge carrier is identified by the impact ionisation rates α and β, respectively relative to the motion of electrons and holes (in the case of silicon, $\alpha > \beta$ for any value of the E-field). These coefficients can be considered exponentially dependent on the E-field.

Non-uniformities in the E-field profile and edge-effects can alter the response uniformity of SPAD devices. In fact, the curvature effect of the junction creates a higher electric field at the edges, resulting in locally differing multiplication factors: at the breakdown voltage, only a reduced peripheral portion of the SPAD is working in Geiger mode. For example, to suppress this edge-breakdown in an n-on-p SPAD, perimetral p-wells (or p-enriched implants) can act as guard rings (or virtual ones) by reducing the E-field at the edges and preventing breakdown non-uniformities. This happens (especially in small area devices) at the expenses of the effective active area, reduced both in the transition from the central high-field region, and due to carriers at the border not drifting vertically, but rather laterally towards the lower-field region.

As temperature increases, scattering of charge carriers from the crystal lattice becomes more likely and mean free paths are shortened, in turn reducing the average energy of the accelerated carriers, so that an increase in the critical E-field value is required in order to compensate for the impact ionisation rate decrease with temperature. This translates into a positive dependence of the breakdown voltage with respect to temperature. Furthermore, the breakdown voltage is heavily affected by the depletion region geometry: a wider depleted region allows for a higher number of possible collisions and thus for a lower critical E-field, which is useful in order to suppress field-enhanced noise generation. Conversely, as the depletion region is wider, the breakdown voltage is also larger. Wider depletion regions will exhibit larger temperature related variations in the breakdown voltage value.

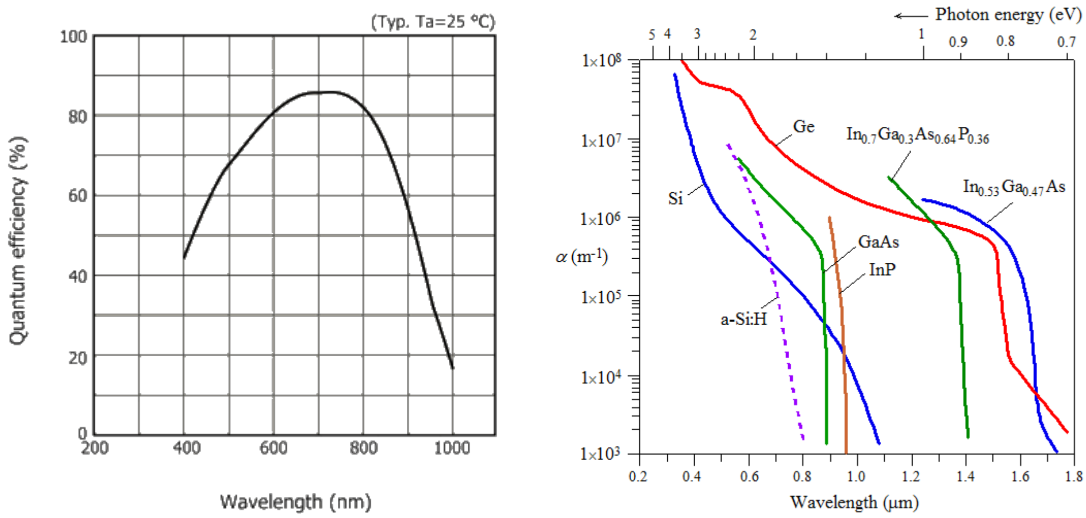

Figure 4.3 Left: typical QE behaviour for SIPMs. Right: photon absorption coefficients for several semiconductors.

4.1.3 **PDE and junction structure**

The photon detection efficiency (PDE) characterises the SiPM spectral response and represents the probability - at a given wavelength - for a single photon impinging on the detector surface to initiate an avalanche.

$$PDE(V_{ov},\lambda) = QE(\lambda) \cdot ATP(V_{ov},\lambda) \cdot FF_e(V_{ov},\lambda) \tag{4.1}$$

The avalanche triggering probability ATP is dependent on the E-field and on the generated carrier location. Moreover, it differs significantly based on the carrier type. The fill factor (FF) represents the ratio of the light sensitive surface of the photodetector to its total area, thus accounting for the mechanical separation between cells (due to optical and electrical isolation reasons), as well as for the area required to host quenching resistors and tracks. The effective fill factor FF_e also takes into account the E-field non-uniformity over the cell and the edge-effects described in the previous section. The silicon quantum efficiency QE (Figure 4.3, left) expresses the probability for a photon to enter the detector without being reflected and to be absorbed either in the depletion layer or in a zone of the detector sufficiently close to it, in order for the generated carriers to be able to escape recombination and reach the active region trough random diffusion (active volume).

The PDE differs for p-on-n and n-on-p junction types. If a junction's width of a few hundreds of nanometres taken into account, the majority of photons in the red or green part of the visible spectrum are absorbed at a greater depth, beneath the junction, as opposed to photons in the blue region which are absorbed above the junction (Figure 4.3, right). As electrons are associated to a higher probability of initiating an avalanche if compared to holes, a greater PDE can be achieved by absorbing the incoming photon in the p-layer. Thus, p-on-n (n-on-p) junctions are more efficient at shorter (longer) wavelengths. Shorter (blue) or longer (NIR) wavelengths will result in shallow or deep absorption of the photon: in this case, only one type of charge carrier (per doping type) is able to drift toward the high-field region to initiate an avalanche (single carrier multiplication). For green and red wavelengths, instead, light is absorbed directly in (or in the vicinity of) the junction and both carrier types are involved in the avalanche initiation. These phenomena contribute to the different spectral responsivity shapes for these photodetectors.

4.1.4 **Gain**

SiPM gain G is defined as the current pulse (or charge) generated for a single detected photon: each pixel generates a defined and quantised charge in each avalanche event (Figure 4.4). The gain of a microcell is represented by normalising this deposited charge over the elementary charge of the electron:

$$G = C_{pix} \cdot V_{ov}/q_e \tag{4.2}$$

In this relation, C_{pix} represents the total microcell capacitance, equal to the sum of the junction capacitance and the parasitic one due to the presence of the quenching circuitry. Although the effect is negligible for high bias voltages, the depletion capacitance also depends on the bias voltage as $1/\sqrt{V_{bias}}$.

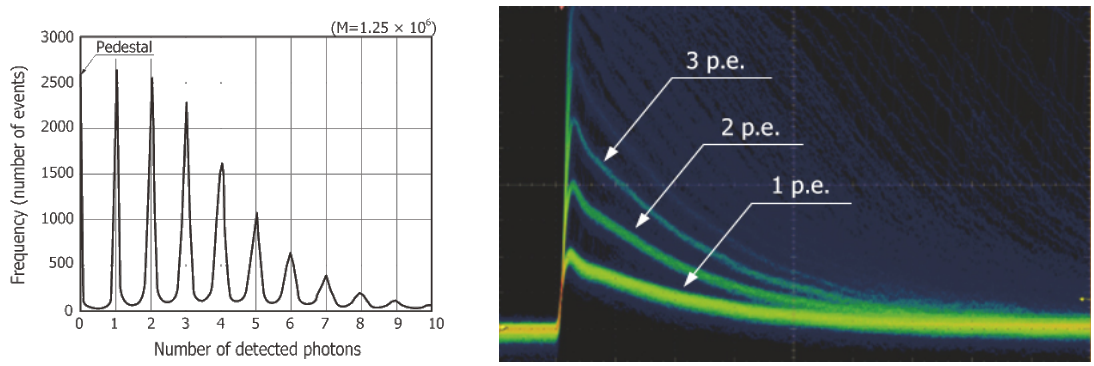

Figure 4.4 Poissonian charge spectrum (left) and quantised pulses (right) of an MPPC

Intrinsic gain fluctuations are very small in SiPMs and can be expressed by the following relation:

$$\frac{\sigma(G)}{\langle G\rangle} = \frac{\sigma(V_{br})}{\langle V_b\rangle} \oplus \frac{\sigma(C_{pix})}{\langle C_{pix}\rangle} \tag{4.3}$$

Variations relative to the cell capacitance are controlled by assuring cell to cell uniformity down to the percent level trough the manufacturing process. Doping densities can contribute to deviations greater than 0.3 V in the breakdown voltage value [43], while the manufacturing tolerances have a smaller impact on these fluctuations (inferior to 0.1 V). Moreover, variations in the quenching time can also lead to non-uniformities in the microcell's gains.

4.1.5 **Primary noise sources and dark rate**

Primary dark noise sources in SiPMs are related to (field-enhanced) thermally generated carriers, as well as tunnelling in the high-field region (dominant at low temperatures). At room temperature, the dominant mechanism is represented by the Shockley–Read–Hall generation and recombination, due to deep level traps associated to lattice defects. The generation rate depends on the epitaxial layer properties and quality, as well as on the cell active volume. A higher presence of defects will result in a shorter equivalent lifetime and thus in a more frequent generation of avalanches: this phenomenon is either measured via the dark count rate (DCR) - generally expressed in kHz - or integrated into a dark current value. The total leakage current I_D in a SiPM is given by the sum of the surface leakage current I_{ds} (which does not undergo multiplication) and the current I_{dg} generated inside the substrate:

$$I_D = I_{ds} + G \cdot I_{dg} \approx C_{pix} \cdot V_{ov} \cdot DCR \tag{4.4}$$

Overvoltage, temperature, semiconductor quality and active area will all affect the dark rate performance. Furthermore, due to the high E-field present in the SPAD element, phenomena such as Pole-Frenkel or thermal-enhanced trap assisted tunnelling can also increase the dark rate.

4.1.6 **Correlated noise: after-pulsing and crosstalk**

In SPADs, multiplication (or gain) noise, is accounted for by the excess noise factor (ENF), which represents a correction applied to statistical (Poisson) noise in order to account for the noise associated with the stochastic avalanche multiplication process. In SiPMs, the contribution of the intrinsic gain fluctuation term to the ENF is negligible, while correlated noise (after-pulsing and crosstalk) and saturation effects became more important. As the presence of correlated noise implies, the average avalanche charge presents an increase in its mean and variance values (quantified in the excess charge factor) and it is thus not merely proportional to the number of detected photons (plus noise) multiplied by the average gain.

Afterpulsing is related to trapping of carriers flowing in an avalanche, which can occur in deep and shallow bandgap levels in the active volume of the photosensor, which lead to secondary avalanches in a pixel that is recovering from a primary avalanche process.

$$P_{ap} = \frac{P_c \cdot P_t}{\tau_s} exp(-t/\tau_s) \tag{4.5}$$

The releasing of these carriers typically happens with an exponentially decaying time distribution, associated with the possibility of secondary spurious avalanches. The afterpulsing probability P_{ap} depends on the relation between the release time constant τ_s (i.e. the trap lifetime) and the microcell characteristic recharge time (4.5). The latter can be tailored to have the majority of releases happening while the recharge process is still ongoing. For long τ_s exceeding the pixel recovery time, the afterpulse waveform becomes indistinguishable from the one relative to a single photoelectron, thus leading to an erroneous increase in the photon detection efficiency and in a degradation of the recovery time of the fired pixel. Afterpulsing is promoted by an increase in the overvoltage, which, apart from enhancing the avalanche mechanism, eases the de-trapping of charge carriers. A linear dependence on the overvoltage is present in either the avalanche triggering probability P_t or the trap capture probability P_c.

Another mechanism associated to after-pulsing is represented by optical induction: secondary photons generated during avalanches can be reabsorbed in the inactive material around the active volume, in turn photo-generating carriers able to diffuse back to the active region and trigger a spurious secondary avalanche. This process is dominated by the carrier lifetime in the substrate, which can range from a few to hundreds of nanoseconds. These secondary photons are generally emitted in the near-infrared band and can travel quite far trough the silicon substrate. It is estimated that in a single avalanche, typically two NIR photons are emitted for every 10^5 electrons crossing the junction. These secondary photons are emitted isotropically and can also induce avalanches in the neighbouring cells (crosstalk), directly reaching adjacent microcells or being reflected at the bottom of the silicon substrate or at the interface with the optical windows. These processes can be direct (prompt) or delayed up to several microseconds due to carrier diffusion. Prompt crosstalk can occasionally mimic the signal of a multi-photon detection.

4.1.7 **Linearity and dynamic range**

SiPMs dynamic range upper limit is represented by the totality of the pixel fired altogether within the microcell dead-time. Increasing the number n of incoming photons results in multiple hits on single pixels, thus decreasing output linearity. The number fired microcells is described by the following relationship:

$$N_{fired} = N_{total} \cdot \left(1 - exp\left(-\frac{n \cdot PDE}{N_{total}}\right)\right) \tag{4.6}$$

4.1.8 Equivalent model and pulse shape

The linearised small signal model of a SiPM is reported in Figure 4.5. A single microcell can be modelled via the series of a SPAD and a quenching section, the latter including the integrated quenching resistor in parallel with a quenching capacitance C_q (which can be either parasitic or designed to enhance fast pulse output). The SPAD section model is instead composed by the inner depletion layer capacitance C_d (which includes the area capacitance along with the sidewall one) in parallel with the series connection of the junction resistance R_d, a voltage source representing V_{br} and a normally open switch which models the impulsive avalanche initiation. The application of this model is valid under the assumption that all time constants present in the circuit are much larger than the characteristic time of the avalanche phenomenon and that only a small number N_f of microcells are fired simultaneously with respect to the total number of pixels of the device (N_{tot}).

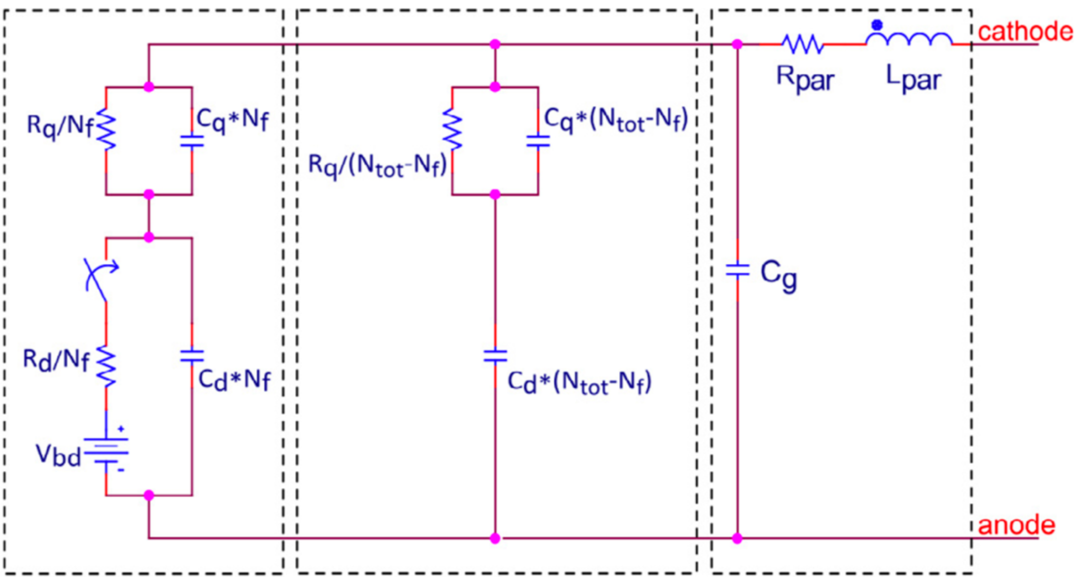

Figure 4.5 Equivalent SiPM model. The active, passive and parasitic sections of the circuit are highlighted from left to right.

Before the avalanche, C_d is charged at the bias voltage. When an avalanche is initiated (or equivalently the switch is closed), C_d is discharged trough R_d, with an elevated inrush current given by V_{ov}/R_d. The voltage drop on R_d translates into a voltage change across C_q, which in turn starts charging, thus producing a current spike available at the device terminals, with a theoretical rise time constant of $(C_d + C_q) \cdot (R_d // R_q)$, which can range down to tens of picoseconds. The discharge of C_d is halted as the voltage across the SPAD section reaches the threshold value which allows the avalanche process to self-sustain. This event is simulated by the opening of the switch at the pre-set threshold voltage, which can generally be assumed to be very close (or even equal) to V_{br}. After quenching, the cell is recharged with a characteristic recovery time $R_q \cdot (C_d + C_q)$ and C_q is discharged via R_q.

The output waveform of a passively quenched microcell is given by the superposition of a slow and a fast component (Figure 4.6). The sub-nanosecond overall rising edge is due to the fast path through C_q for the current at the beginning of the avalanche, with a characteristic time constant $\tau_{rise} \approx R_d \cdot (C_q + C_d)$. By indicating with $R_s \ll R_q$ the shunt resistance relative to an ideal transimpedance amplifier section connected to the SiPM, and with C_t the total capacitance seen by the said amplifier, the following approximate relations can be derived [41][42]:

$$V(t) \approx \frac{Q \cdot R_s}{C_q + C_d} \left(\frac{C_q}{\tau_f} exp(-t/\tau_f) + \frac{C_d}{\tau_s} exp(-t/\tau_s) \right) \tag{4.7}$$

where the fast and slow fall times are respectively $\tau_f \approx R_s \cdot C_t$ and $\tau_s \approx R_q \cdot (C_q + C_d)$ and $C_t \approx C_g + C_q // C_d$. The avalanche charges (total, prompt and slow) are expressed by the following relationships:

$$Q = V_{ov} \cdot (C_d + C_q) = (V_{bias} - V_{br}) \cdot (C_d + C_q) \tag{4.8}$$

$$Q_f = Q \cdot \frac{C_q}{(C_q + C_d)} \approx Q_s \frac{C_q}{C_d}$$

The following relationships can finally be written for the voltage peaks of the two components:

$$\frac{V_s^{max}}{V_f^{max}} \approx \frac{C_q C_t R_s}{C_q^2 R_q} \tag{4.9}$$

Increasing the ratios C_q/C_d or R_q/R_s thus yields a better response in terms of timing and pulse height. In this simple analysis the parasitic grid capacitance C_g (which is due to the sensor substrate and should include fringe effects), as well as the load effects of the other microcells present between the photodetector terminals are not included. Moreover, the parasitic inductance and resistance due to the interconnections can also be included in the model. These elements are associated with undesirable low pass filtering and attenuation of the pulse.

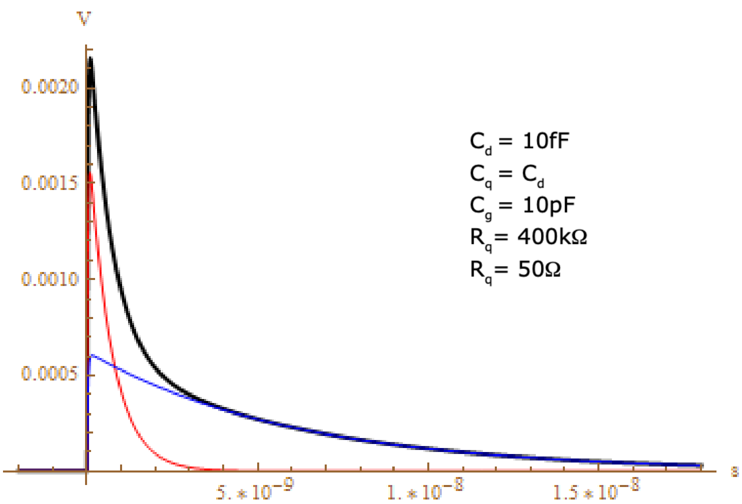

Figure 4.6 Simulated response of a single SiPM microcell (black). The slow and fast components are respectively shown in red and blue.

4.2 Mu2e SiPMs

The SiPMs used in the Mu2e calorimeter are custom components designed specifically for this application. Two adjacent identical photodetectors are used to read each Mu2e crystal, as a mean to increase redundancy (cf. section 3.2). Each photodetector consists of a large area 2×3 array of individual MPPCs (Figure 4.7). Each MPPC in the array counts 14400 50×50 μm^2 pixels, with an effective active area of 6×6 mm^2 each. The size of the pixels represents a good compromise between gain and photo-detection efficiency. The photodetectors are UV-extended, and high efficiency in matching the 315 nm emission peak of pure CsI is guaranteed by silicone optical windows, thanks to the use of silicone protection layer (SPL) technology. The Mu2e SiPM also feature trough-silicon vias (TSV) technology.

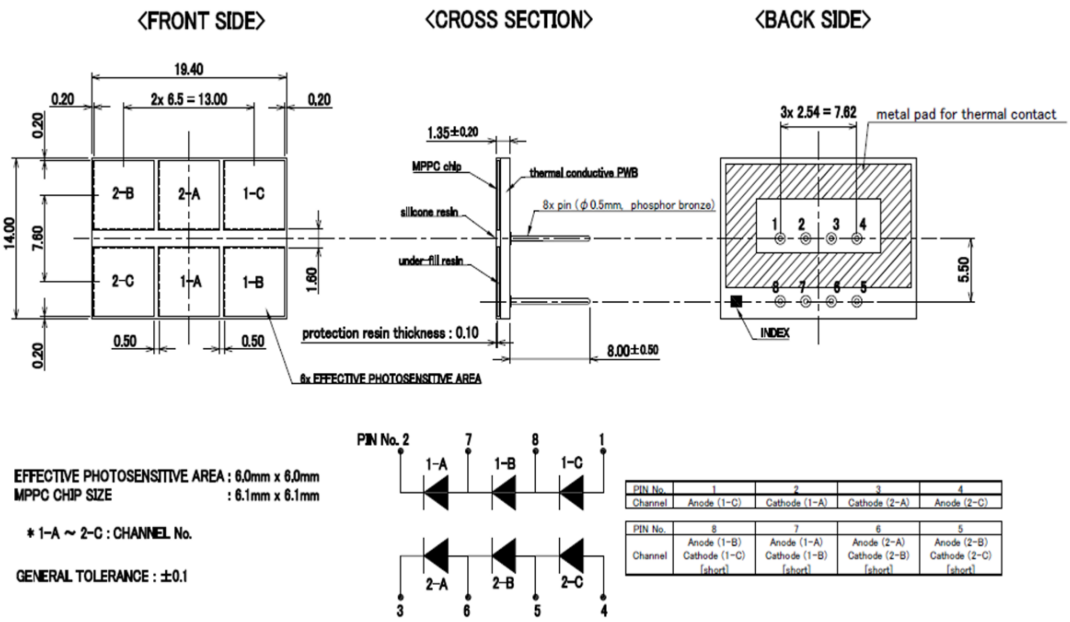

Figure 4.7 Mu2e SiPM drawing

In order to comply with the requested short (<40 ns) pulse width (as a mean to prevent pile-up), the 6 cells are connected as the parallel of two series made of 3 MPPCs each. This arrangement allows the total capacitance to be reduced to 2/3 of the value corresponding to the parallel of 6 cells. An important advantage of the series connection of MPPCs (as opposed to the parallel one) is associated to the self-aligning capacity of this configuration with respect to the individual over-voltages, regardless of the individual breakdown voltages of the cells. Under the series connection, the biasing voltage applied to each SiPM is determined by the common leakage current, so that the difference in the individual breakdown voltages is balanced and the over-voltages are aligned to some extent.

The chosen manufacturer for the Mu2e SiPMs is Hamamatsu [57]. A total of 4000 Mu2e SiPMs have been supplied to the experiment for the calorimeter production phase. The Quality Control (QC) procedure was performed as follows: 300 SiPMs/month were shipped from the manufacturer to the Fermilab Silicon Detector facility (SiDet) where they underwent a visual survey and a control of the mechanical specifications. The sensors passing this first check were fully characterised in a dedicated station. From each lot, randomised batches of 20 SiPMs were extracted, 15 of which were used for the MTTF evaluation (cf. 4.2.1.6), while the remaining 5 were irradiated at HZDR (cf. 5.3.3).

Each photodetector has to fulfil the following specifications:

- PDE > 20 % at 315 nm for each cell;
- gain $> 10^6$ for a nominal overvoltage $V_{ov} = 3$ V (for each cell);
- a recovery time for each cell inferior to 100 ns, relative to a load impedance of 15 Ω (or greater);
- matched bias voltages for the 6 cells on each sensor, with a spread smaller than ± 0.5% RMS;
- a dark current spread over the 6 cells inferior to ± 15% RMS under operating conditions;
- a dark current inferior to 10 mA and a gain drop smaller than a factor of 4 after exposures up to $3 \cdot 10^{11}$ $\mathrm{n_{1MeV}/cm^2}$ at 20 °C;
- a MTTF better than 10^6 hours at 0 °C;
- thermal resistance inferior to $7 \cdot 10^{-4}$ $\mathrm{m^2K/W}$.

4.2.1 Characterisation and quality control

4.2.1.1 Breakdown voltage and dark current

A characterisation of the breakdown voltage V_{br} and of the dark current I_D is carried out for each production SiPM, through a I-V scan performed individually on each one of the 6 cells. The breakdown threshold is evaluated as the voltage corresponding to a change in concavity of the I-V characteristic, calculated as the maximum of the logarithm derivative of the current fitted with a log-n function (Figure 4.8). The value of the quenching resistor can be deduced from the forward-bias I-V characteristic: by increasing the overvoltage, the I-V curve eventually becomes linear, as the resistance of the depletion region in the SPAD approaches zero, and the current is only limited by the quenching resistor.

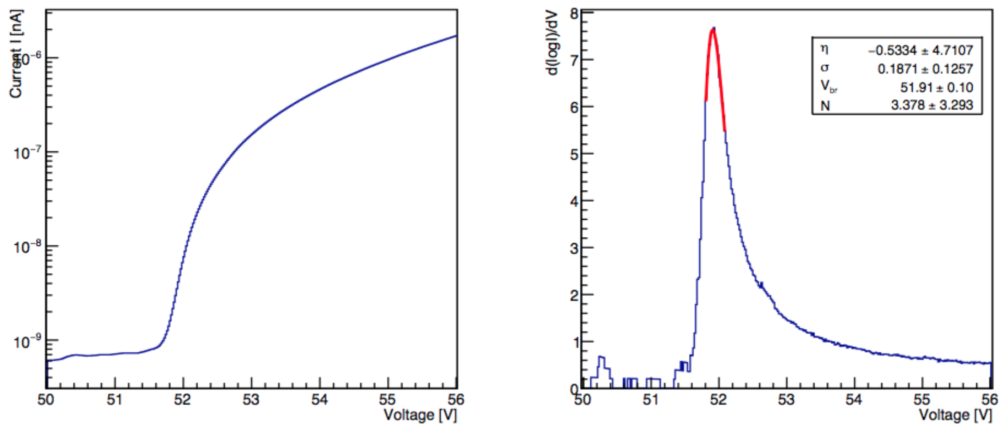

Figure 4.8 Left: I-V scan example (at 25 °C). Right: dark current logarithmic derivative versus bias voltage, with a superimposed log-n fit.

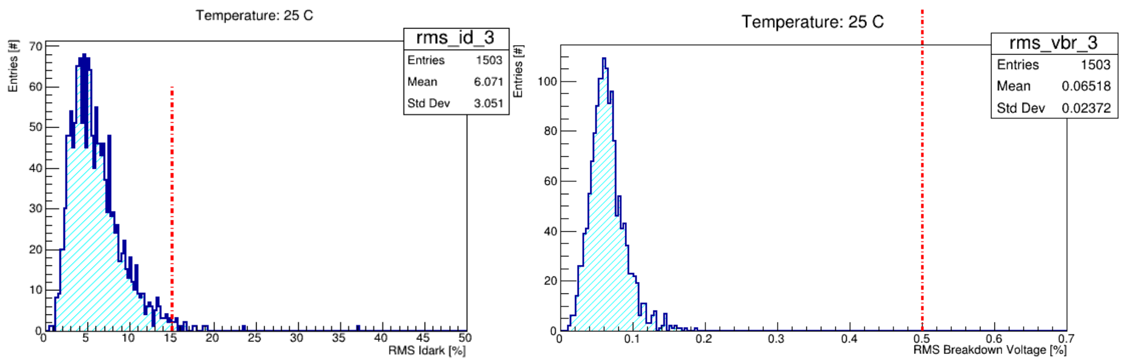

Figure 4.9 Distribution of the RMS spreads of V_{br} and I_D for all production SiPMs.

The RMS spread of V_{br} and I_D relative to the 6 cells of each SiPM were evaluated, as reported in Figure 4.9: red lines represent the quality control cuts.

Just 0.6% of the whole production was not in compliance with the dark current requirement and thus rejected. The breakdown voltage temperature dependence of the Mu2e SiPMs was also characterised with a positive temperature coefficient of 54 mV/°C (for a single cell).

4.2.1.2 **Gain**

Gain measurements were performed by illuminating the sensors with a 315 nm UV LED using 20 ns pulses at 100 kHz. Pulse amplitude was attenuated using progressive neutral density filters until only a few photons were able to reach the sensor surface. The SiPMs were biased with the nominal overvoltage of 3 V (per cell). A custom front-end amplifier with a fixed gain of 250 was used to process the pulses, which were then integrated over a 150 ns window into the charge spectrum shown in Figure 4.10 (left). This distribution follows a Poissonian statistics and the individual quantised peaks are clearly discernible, each of which corresponds to a different number of detected photons: the separation between each pair of adjacent peaks indicates the charge generated from a single pixel being fired. The first charge peak centred around zero is the noise pedestal and the underlying spectral noise is due to correlated noise and minor intrinsic gain fluctuations, as discussed previously (cf. Section 4.1.4). The gain of Mu2e SiPMs at the nominal overvoltage were in the range $(1.58 \div 1.74) \cdot 10^6$. Figure 4.10 (right) reports the measured gain of a single cell for various bias voltages (at 25 °C).

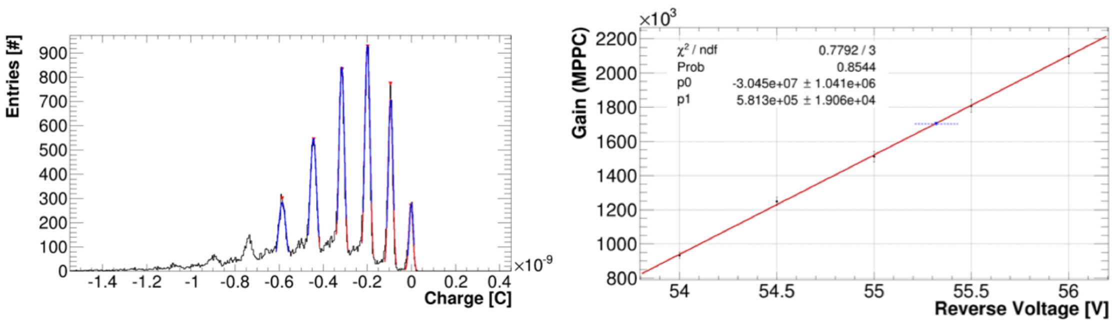

Figure 4.10 Left: charge distribution of a single SiPM cell illuminated by a LED light. Right: cell gain as a function of the bias voltage at 25 °C.

4.2.1.3 **PDE**

At a given wavelength, the PDE represents the average number of detected photoelectrons (n_{pe}) normalized to the average number of incident photons hitting the sensor active area (n_γ). Under the illumination of a stable flux of incident photons, the probability $P(n)$ for n detections to happen in a defined integration gate is subject to a Poissonian statistic:

$$P(n, n_{pe}, n_{dark}) = \frac{(n_{pe} + n_{dark}) \cdot e^{-(n_{pe}+n_{dark})}}{n!} \tag{4.9}$$

The 4.9 can be inverted as

$$n_{pe} = -ln(P(0, n_{pe}, n_{dark})) + ln(P(0{,}0, n_{dark})) \tag{4.10}$$

In the 4.10, $P(0, n_{pe}, n_{dark})$ is the probability of having zero counts under no illumination of the sensor. The same setup described in the previous section was used, and the pulse amplitude was further reduced to have an average of 1 photon hitting the SiPM area. The LED intensity was used as a reference indication of the emitted flux N_γ and its stability was monitored using a reference sensor. Acquisitions of 1 µs width were triggered on each pulse with a 500 ns pre-time and low pass filtered (via FFT) with a cut-off frequency of 150 MHz for noise suppression. During the acquisition pre-time, the dark pulses count (N_D) was evaluated over a 20 ns gate, and a second 20 ns window was used to count LED pulses containing at least one photoelectron (N_{nGE1}). Being N_t the total number of recorded events, the 4.10 can now be rewritten as follows:

$$n_{pe} = -ln\left(1 - \frac{N_{nGE1}}{N_t}\right) + ln\left(1 - \frac{N_D}{N_t}\right) \tag{4.11}$$

These measurements were compared to a reference sensor with a known PDE of 22% at 315 nm, as a way to determine N_γ:

$$\frac{PDE}{PDE_{REF}} = \frac{n_{pe}}{n_{pe\,REF}} \cdot \frac{N_{\gamma\,REF}}{N_\gamma} \tag{4.12}$$

where the $N_{\gamma\,REF}/N_\gamma$ is evaluated as the ratio of the respective currents measured by the reference light sensor. Following this procedure, the Mu2e photosensors showed a PDE of (28.0 ± 1.3)%.

4.2.1.4 **Gain and reconstructed charge**

A second characterisation run of the Mu2e SiPMs has been carried out at LNF, with the following setup: a ~ 20 cm diameter integrating sphere was fabricated through additive manufacturing in order to accommodate a complete SiPM holder with two SiPMs installed, and the sphere inner surface was coated with BC620 reflective paint (Figure 4.11). A 420 nm LED source was pulsed at 1 kHz with a 20 ns pulse width and attenuated using progressive neutral density filters mounted on an automatic wheel selector. The attenuated LED light was injected into the sphere via fibre optics. The FEE gain was set to 1 and the SiPMs signals were digitised at 1 Gsps. Throughout the acquisition runs, the whole setup was installed in a climatic chamber with a constant temperature of 30 °C.

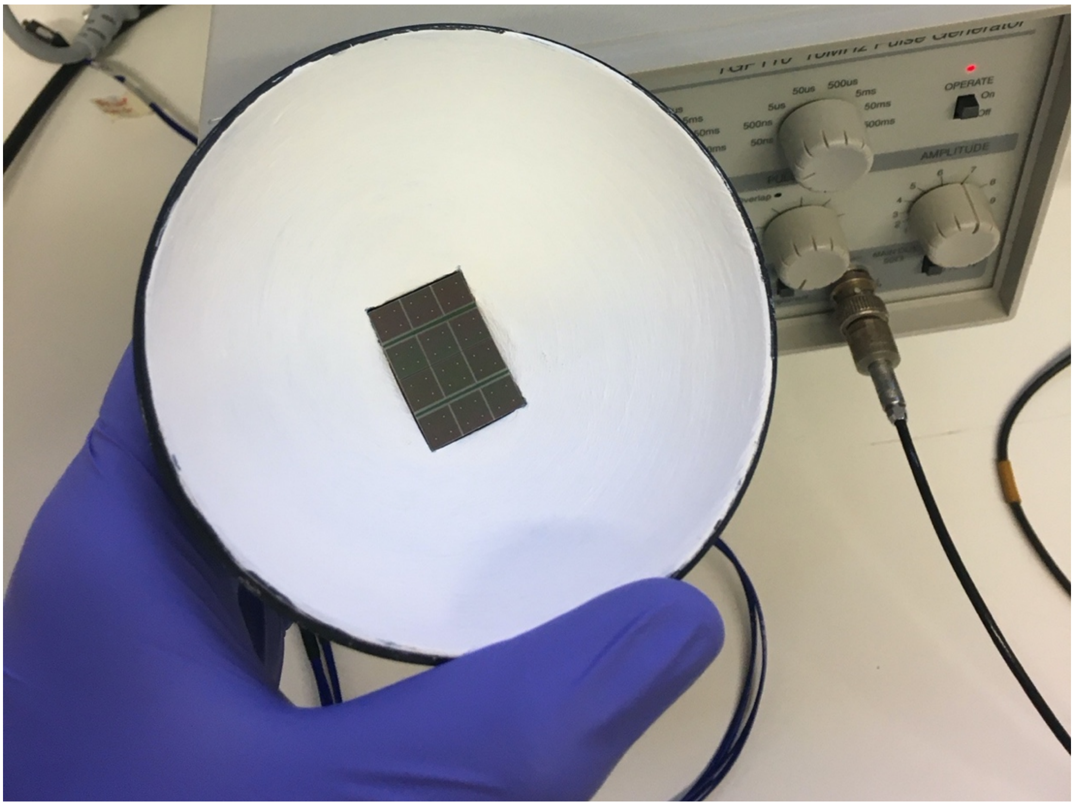

Figure 4.11 Mu2e SiPM holder installed inside the integrating sphere.

An automatic scan was performed using 9 progressive neutral density filters on the selector wheel: $2 \cdot 10^4$ waveforms were acquired for each filter and successively integrated over a 200 ns gate, resulting in the charge spectra reported in Figure 4.12. Each charge peak was fitted with a gaussian distribution in order to extract the respective mean and standard deviation. The obtained reconstructed charge resolutions were evaluated as the relative standard deviations σ_Q/Q. The thus obtained values for a 9 filter scan are plotted in Figure 4.13.

These resolution versus charge measurements (Figure 4.13) can be fitted with the following relationship:

$$\frac{\sigma_Q}{Q} = \sqrt{\frac{p_0}{Q}} \oplus \frac{p_1}{Q} \oplus p_2 \tag{4.13}$$

The parameter p_1 represents the contribution due to electronics noise, while the constant term (p_2) accounts for the fluctuations of the light source. Given the Poissonian photoelectrons statistics and considering only the contribution of the stochastic term p_0 (while also neglecting the Fano factor), it can be seen that $\sigma_Q/Q = \sigma_{N_{pe}}/N_{pe} = 1/\sqrt{N_{pe}}$. As a result, it is clear that $p_0/q_e = G$ represents the gain of the photodetector, expressed in pC/pe.

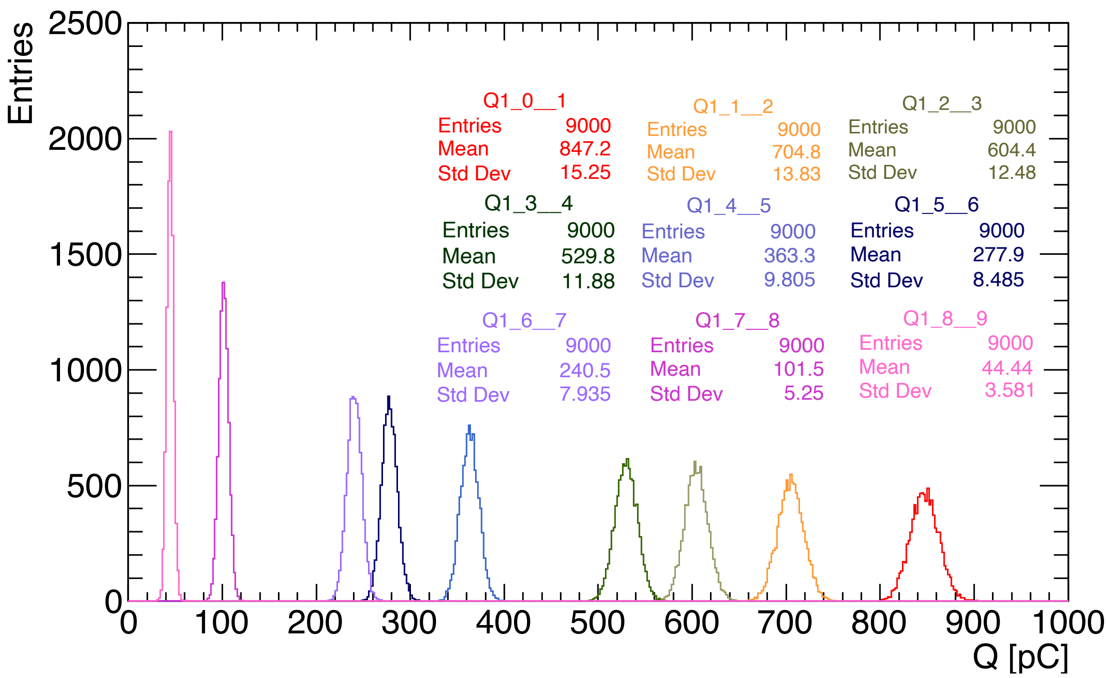

Figure 4.12 Reconstructed charge for 9 progressive neutral density filters.

The procedure was repeated 10 times to exclude any systematic effects: it was verified that $\sigma_{Gi}/\sqrt{N_{rep}} \approx \sigma_{\langle G\rangle}$ for each one of the N_{rep} repetitions, where σ_{Gi}

is the sigma relative to a single repetition and $\sigma_{\langle G \rangle}$ is the error on the calculated value if G. The measurement repeatability was $\sim 1\%$.

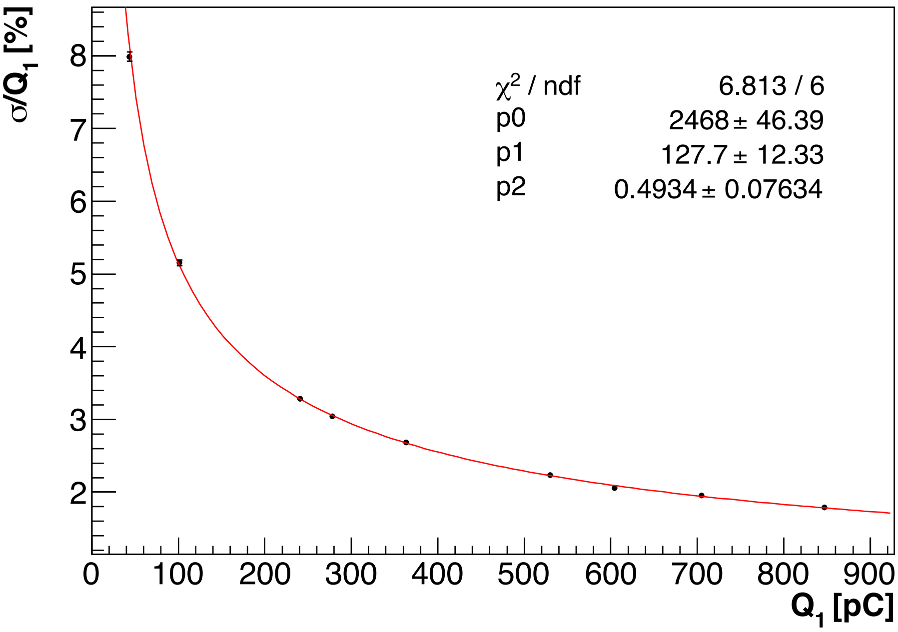

Figure 4.13 Charges and charge errors for a single 9 position scan.

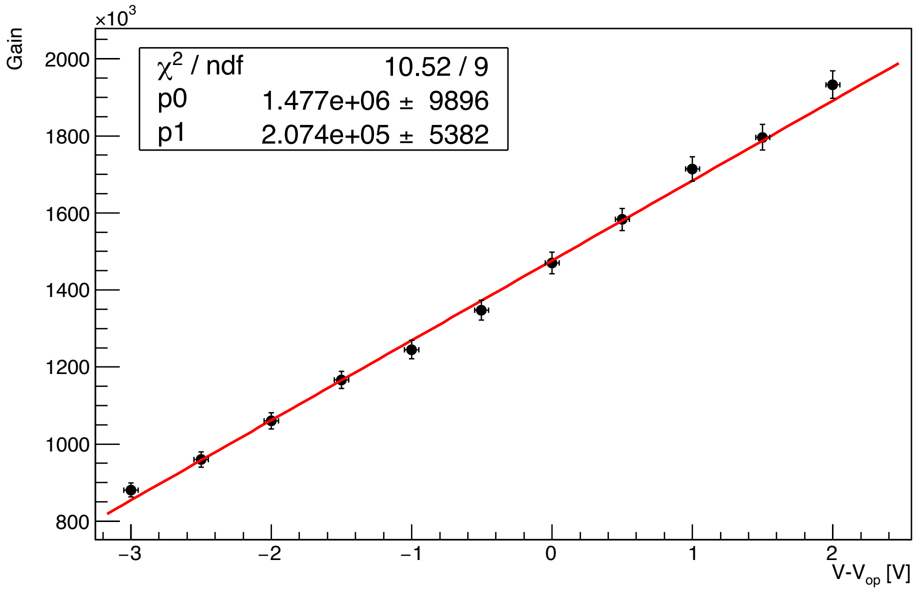

Figure 4.14 Gain as a function of $V - V_{op}$.

During this procedure SiPMs were biased at the nominal overvoltage and the gain values evaluated with this measurement were corrected for the 30 °C temperature of the climatic chamber, using the previously stated 54 mV/°C/cell temperature coefficient. The gain value at 25 °C (at nominal overvoltage) was successfully extrapolated as ~ $1.7 \cdot 10^6$, in agreement to what is reported in Section 4.2.1.2.

By relying on the same setup, a verification of the gain dependence on the overvoltage was carried out by performing a scan over the range $V - V_{op} = [-3;\ 2]$ V, where $V_{op} = V_{br} +\ 3$ V represents the nominal biasing voltage. The result of this procedure is shown in Figure 4.14. Gain variation with bias voltage is thus expressed as $p_1/p_0 \approx 13$ %/V. This translates to a temperature coefficient of ~ -2.2 %/°C.

4.2.1.5 **Pulse shape and recovery time**

As previously stated (cf. Section 4.1.8), the quenching time incorporates two decay constants: a fast component determined by the microcell recharging (which accounts for approximately 90% of the recovery process) and a slow one associated to the recharge of the SiPM bulk. A single cell has a capacitance of $C_t \approx 1280$ pF and a nominal value for the quenching resistor of 150 kΩ, yielding a nominal time constant $\tau_C = C_t \cdot (R_q/N_{pix} + R_s) \approx 80$ ns, where R_s is the 50 Ω load impedance and $N_{pix} = 14400$ is the total number of pixels.

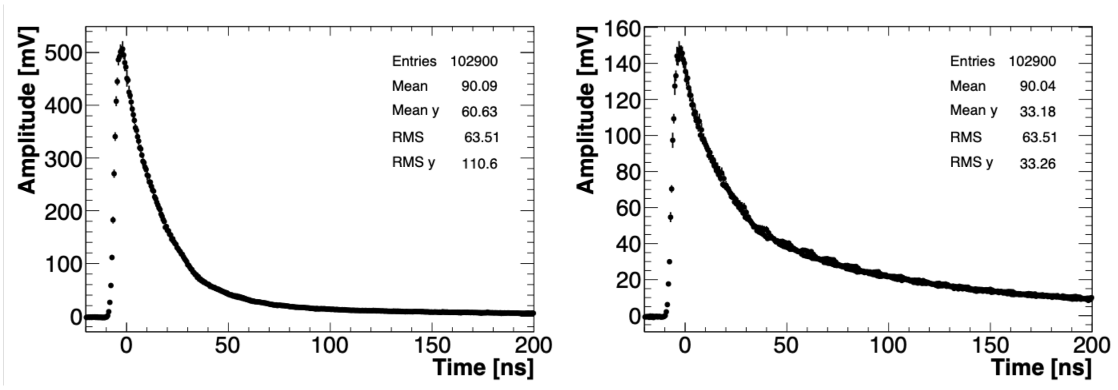

Figure 4.15 Right: example of a Mu2e SiPM pulse shape (right). Left: single cell waveform.

Photodetectors were thermalised at 20 °C and characterised by firing a fibre coupled picosecond pulsed blue laser at 1 kHz rate. This time, no preamplification was used as not to alter the pulse shape. A set of 5000 waveforms were acquired per photosensor using a 1 Gsps CAEN digitiser (DT5751). Waveforms (Figure 4.15) were then fitted using a triple exponential model (cf. section 4.1.8). All sensors successfully passed the characterisation step with a $\tau_C = 94.3$ ns for $R_S = 50\ \Omega$. This value was corrected for the 15 Ω load values listed in the requirements, yielding an extrapolated value of $\tau_C = 28.3$ ns for $R_S = 15\ \Omega$.

4.2.1.6 **MTTF test**

Mean time to failure (MTTF) was characterised in accelerated life tests using randomised batches of 15 SiPMs per production lot. The sensors were operated under illumination from pulsed blue LEDs and at a temperature of 65 ± 0.5 °C for a total of 342 hours (Figure 4.16).

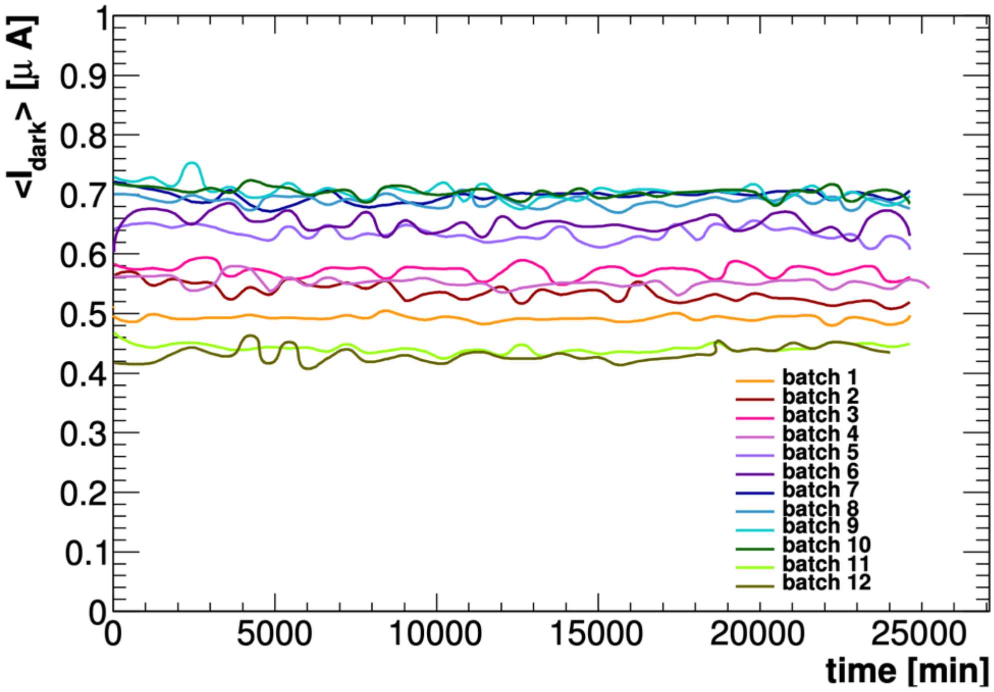

Figure 4.16 SiPM dark currents monitored during the MTTF test.

This setup translates to the following acceleration factor, derived from the Arrhenius equation:

$$AF = exp\left(\frac{E_a}{k}\left(\frac{1}{T_{op}} - \frac{1}{T_{stress}}\right)\right) \tag{4.13}$$

The Silicon activation energy is $E_a = 0.7\ eV$, and the operating and stress temperature are respectively 273 and 338 K. Under the aforementioned conditions, the MTTF can be evaluated as

$$AF = \frac{1}{2} \cdot N_{DUT} \cdot N_{hours} \cdot AF \tag{4.13}$$

The totality of the SiPMs under test ($N_{DUT} = 15$ devices) were able to survive to the test without any change in their charge and amplitude response, yielding a mean time to failure greater than 10^7 hours.

4.3 Front-end design

The design and characterisation of the front-end electronics (FEE) has been carried out by the SEA (Servizio Elettronica e Automazione) laboratory at LNF-INFN.

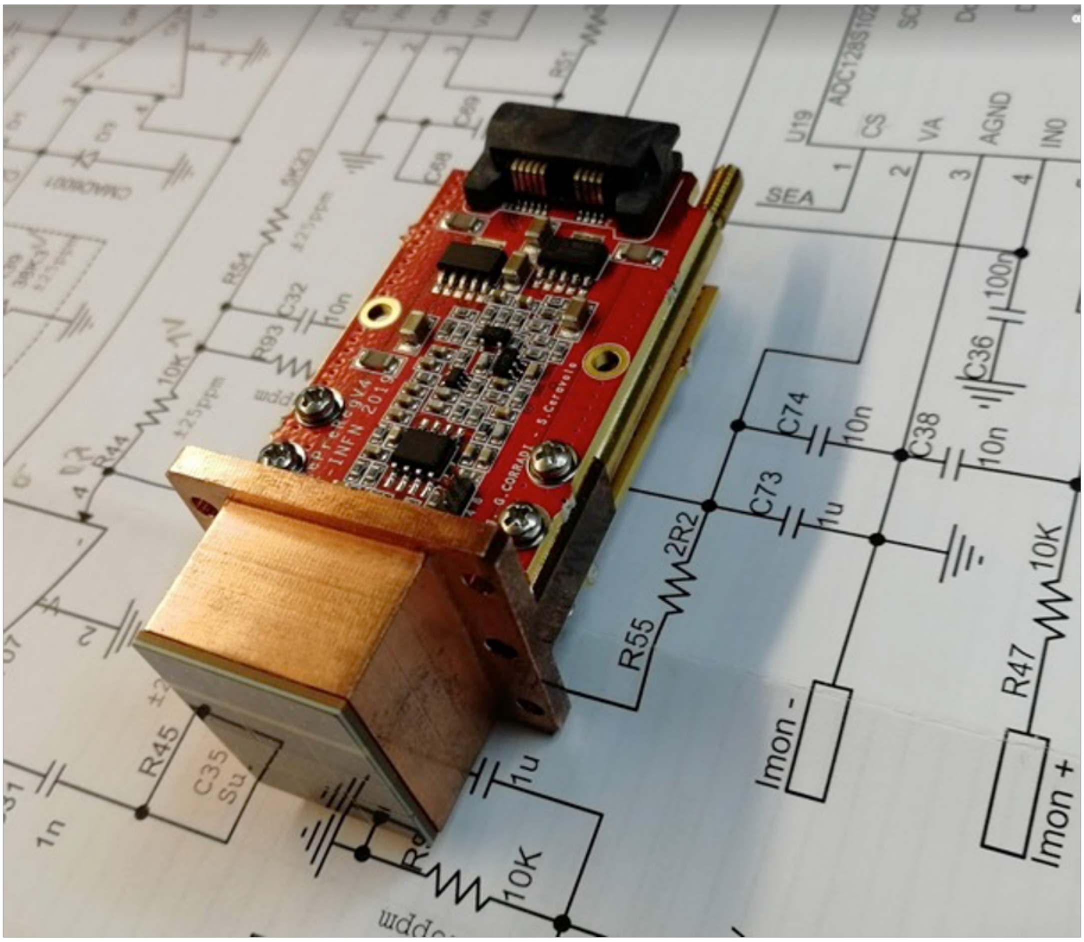

Figure 4.17 Mu2e SiPMs along with the respective FEE boards.

The general requirements which need to be satisfied by the Mu2e front-end electronics are hereby stated:

- selectable SiPM preamplifier gain at 1 or 2 (through settable jumper);
- signal rise time larger than 25 ns (5 times larger than the digitiser sampling interval to guarantee at least 5 points on the leading edge for appropriate timing reconstruction);
- short falling time to prevent pileup;
- low ripple, high precision and high stability voltage regulation ($\pm$ 50 mV);

- stability in the voltage regulation, signal amplification and shaping for event rates up to 1 MHz per channel;
- low power consumption;
- stable HV regulation after a radiation-induced dark current increase of up to 2 mA per SiPM:
- a radiation hardness level in compliance with the requirements stated in Section 3.1.

Two identical and independent FEE boards will handle the two SiPMs which read a single CsI crystal (Figure 4.17). Each board will provide individually programmable bias voltages for each photosensor, perform signal amplification and shaping, while monitoring the current and temperature of each SiPM.

4.3.1 **SiPM**

As previously stated, the 6 individual cells of the Mu2e SiPM are internally connected in two series of 3 elements each (Figure 4.7). As shown in Figure 4.18, on the FEE boards, the top cathodes of each series (K1 and K2) are connected to the regulated HV bias voltage, while the bottom anodes (A1 and A2) are connected through two separate 50 Ω shunt resistors (R53 and R54) to a common node, the latter shunted to the analogue ground via the 100 Ω current sense resistor R1. The pulses are separately extracted at the bottom anode of each series and capacitively coupled to the preamplifier stage (IN+A and IN+B). The regulator and preamplifier sections, respectively marked REG and PRE, can be identified in the schematic and will be discussed in the following sections.

4.3.2 **Linear regulator section**

The linear regulator section (Figure 4.19) of each FEE board (Figure 4.21) provides a programmable bias voltage to the SiPM with ± 50 mV stability and a maximum current sourcing capability of 2 mA. The regulator has a nominal 3mV peak-to-peak ripple and a 100 us settling time.

The regulated high voltage (HV) supply is distributed at 200 V to the individual FEE boards, which - by means of a linear regulator - provide the individual bias voltages to each SiPM.

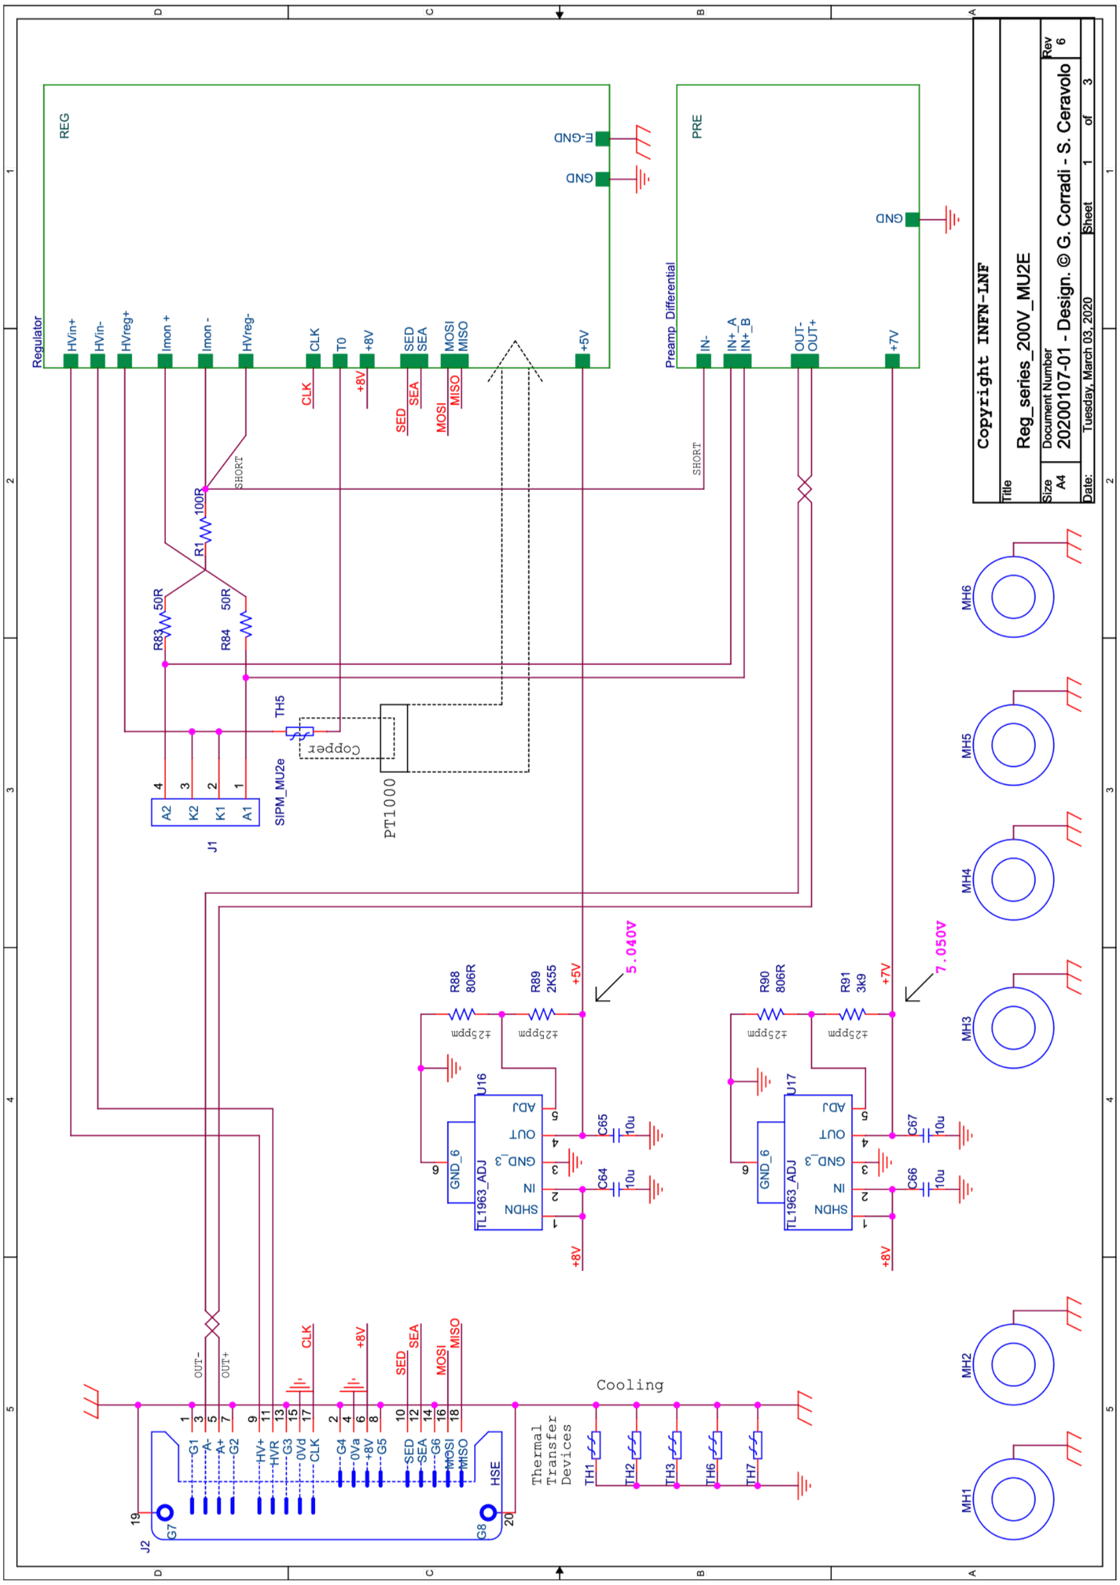

Figure 4.18 SiPM connection, linear regulators and thermal management of the front-end electronics board

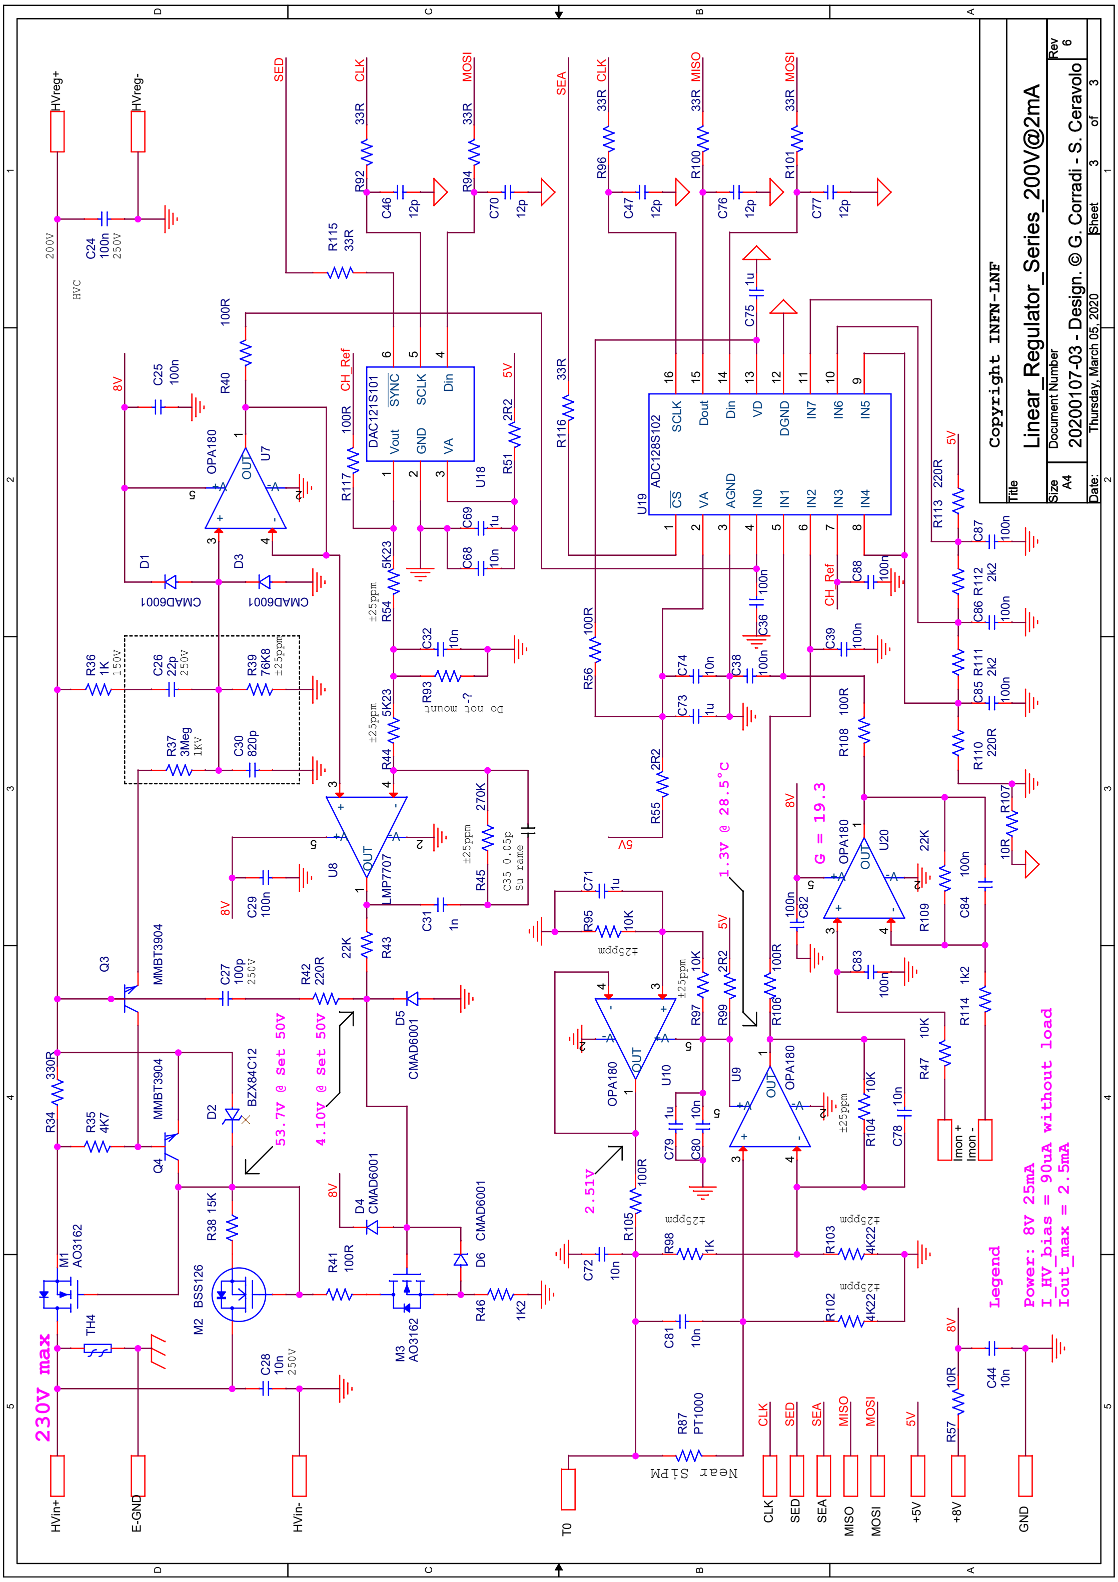

Figure 4.19 Linear regulator, temperature and current sensing section

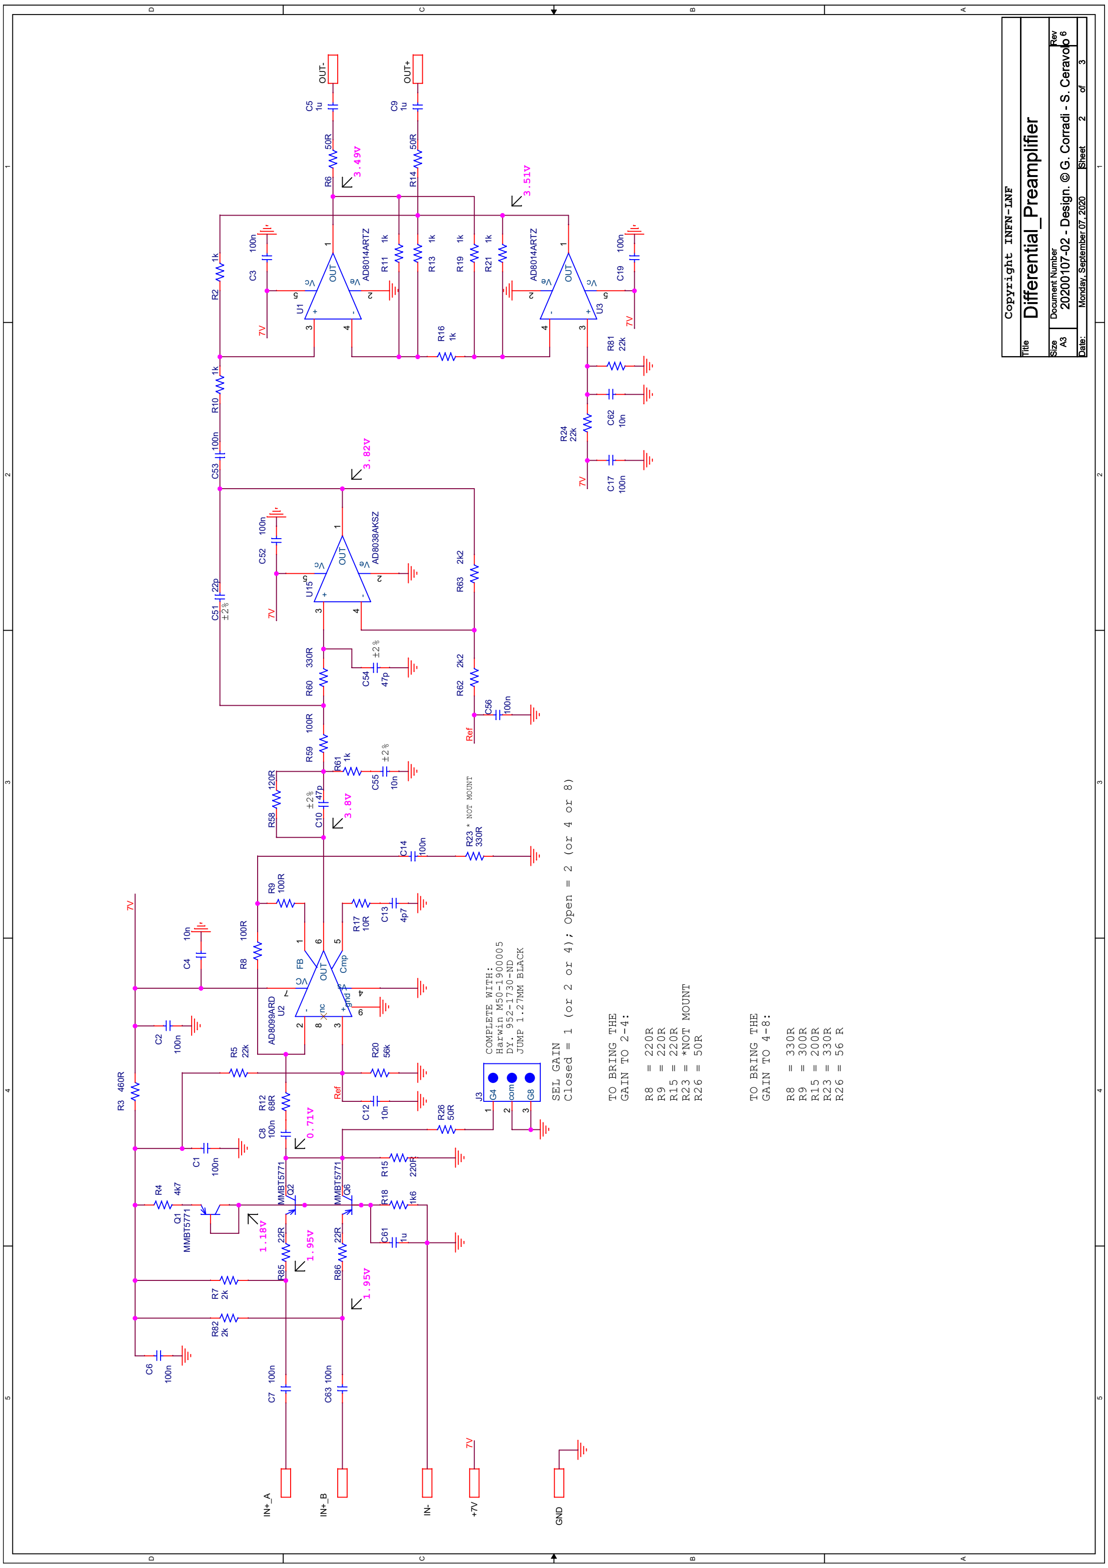

Figure 4.20 Schematic of the SiPM preamplifier, pulse shaper and differential line driver

With respect to the schematic shown in figure Figure 4.19, M1 and M3 are AO3162 n-mosfets, which have a low static drain-source on-resistance and are capable of fast operation, thanks to their low input and reverse transfer capacitance. M2 is a BSS126 depletion mode n-mosfet, which, along with R38, operates as constant current source providing ~ 130 μA of biasing trough R41. Mosfet M3 is driven by the error amplifier U8 and acts as a variable load, thus regulating the voltage at M1 gate. Resistor R46 provides negative feedback over the gate-source voltage of M3, while also increasing the bandwidth. The CMAD6001 diodes D4 and D6 act as a gate protection from reverse voltages during transients, while diode D6 is also useful to achieve a faster M3 turn-off time. The R43 and C27 network provides an additional compensation to the M3 dominant pole to provide additional phase boost.

Transistor Q4, along with R35 and R34 acts as an overcurrent cut-off, which intervenes when the voltage across R34 reaches the base saturation voltage, thus yielding a protection limit of ~ 2.2 mA on the output current. The regulated HV is sampled trough Q3, connected in an emitter-follower configuration: the HV (minus the 0.65 base-emitter drop) is thus applied from the emitter of Q3 to the divider formed by R37 and R39, and then filtered through C30. The C26 and R36 network act as damping branch. After being buffered by U7 (OPA180), the feedback voltage is applied to the non-inverting input of the error amplifier U8. After low pass filtering through R40 and C36, the feedback voltage is also measured by channel 0 of the 128S102 ADC (U19). The reference voltage of the regulator is set via the DAC U19 and applied to the inverting pin of the LMP7707 error amplifier.

The dynamic performance of the linear regulator must guarantee excellent stability and extremely fast regulation against the high-rate SiPM signals and all along the programmable 1-200 V regulated output range. The error amplifier is based on the LMP7707 fast precision op-amp which has a GBWP of 14 MHz at gain 10, along with ultra-low input bias current and input offset voltage. The type-2 compensation network formed by R44, R45, C31 and C35 provides one origin pole, one zero at ~ 590 Hz and one high-frequency pole at ~ 12 MHz. This configuration yields good phase response and flat gain in the mid-frequency range, with a substantial phase boost due to the wide separation between the zero and the finite pole. It should be noted that C35 is a 0.05 pF distributed element component embedded directly on the PCB.

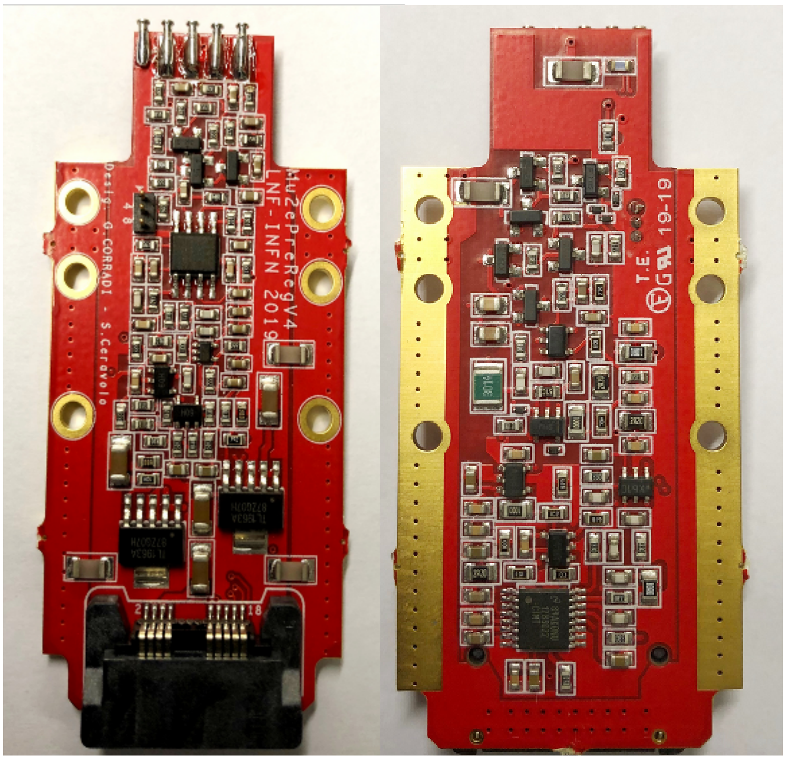

Figure 4.21 Photo of an assembled front-end PCB.

The 12-bit ADC and DAC allow a minimum step of 50 mV (corresponding to 1 LSB) relatively to the HV regulation and sensing. The resistor ladder composed by R110-R113 and R107 and the relative ADC channels are used for diagnostic purposes, as will be discussed in section 5.3.2.1.

4.3.3 Thermal management and current sensing

The worst-case maximum dissipated power on the linear regulator is 350 mW. The thermal management of the board is entrusted to electrically insulating aluminium nitride thermal bridges (Therma-Bridge), which are marked as TH1-TH7 in Figure 4.18. These bridges are used to remove the heat generated by the linear regulator circuitry from the analogue ground (AGND) towards the peripheral thermal ground plane on the PCB. This ground represents the calorimeter instrument ground (E-GND) and it is in electrical and thermal contact with the SiPM holders, with the calorimeter frame and with the cooling system described in Section 3.2.5.

The photosensors temperature is sensed via an SMD PT1000 (R87) which is mounted in close proximity to the SiPM holder and the thermal bridge TH5. The PT1000, read through the Wheatstone bridge composed by R87, R98, R102 and R103, is amplified via U9 and read by ADC channel 2. The op-amp U10 provides a buffered 2.51 V supply voltage for the Wheatstone bridge.

The SiPM current is instead sensed across the 100 Ω sensing resistor R1, at the nodes Imon+ and Imon- (Figure 4.18). As shown in Figure 4.19, U20 is configured as a current sense amplifier with a gain of approximately 20. Its output is read by ADC channel 1, with a 50 to 2000 μA useful range.

4.3.4 **Amplifier and shaper section**

SiPM pulses from the two 3-cell series arrays are coupled to the preamplifier section, respectively trough C7 and C73 100 nF capacitors. The first amplification stage is represented by Q2 and Q6 (MMBT5771) which independently amplify the signals relative to each one of the two 3-cell SiPM series on each photodetector.

The Q2 and Q6 PNP BJTs operate in forward-active mode and are configured as transimpedance amplifier in common base configuration. A 1.18 V biasing voltage is applied to both BJT bases via R4, R18 and Q1. The latter is connected as a diode in order to compensate Q2 and Q6 junctions against temperature fluctuations. The DC component is restored via R7 and R82 (and analogously R82 and R86), so that a 1.95 V voltage is present at the node between these two resistors and a quiescent current $I_e \approx 2$ mA flows into the emitters of the BJTs. Under small signal AC analysis, the BJT bases are effectively grounded trough C61, and thus the impedance seen by the SiPM pulses is effectively the 22 Ω of R85 (or R86) plus the $1/g_m$ emitter impedance. Considering a $V_T = \frac{k_T \cdot T}{q_e} \approx$ 25 mV and an emitter impedance $r_e = V_T/I_e \approx 12\ \Omega$, the input impedance is thus approximately 35 Ω.

The two signals relative to the two SiPM 3-cell series can be safely summed (at the R15 node) over the ~$1/h_{ob}$ output impedance of the BJT stage. The gain of the system section is set by the collector resistance and can be selected as either 1 or 2 by respectively closing or opening the settable jumper which puts R26 (50 Ω) in parallel with R15 (220 Ω). A second amplification stage composed by the AD8099 high speed op-amp follows the BJT stage and buffers the signals before their injection in the shaping section.

The lead-lag compensator (pole-zero) formed by R58, C10, C55 and R61 performs the first shaping of the signal by suppressing the slow decaying component of the SiPM pulse (cf. 4.1.8). The voltage feedback op-amp AD8038 is used in an Allen-Key architecture and configured as a Bessel filter, chosen for its maximally linear phase response, in order to preserve the SiPM pulse shape. The shaper is tailored in order to achieve a 25 ns minimum rise time for the SiPM pulses sent to the digitisation section. In fact, the 25 ns rise time is five times larger than the digitiser sampling period and has been chosen so as to guarantee that the timing reconstruction algorithms can work with at least 5 sampled points on the leading edge.

Finally, the signal is fed into a balanced differential driver (formed by U1, U3 and the relative passives) which feds the signal trough C5 and C9 to the twisted pair balanced line with an output impedance of 100 Ω. The differential to single ended conversion and the common mode adaptation happen on the digitiser board, before the acquisition ADC.

4.3.5 Mezzanine, DIRAC and DAQ

Each one of the calorimeter disks are subdivided into 34 similar pseudo-azimuthal sectors, each one counting 20 crystals. Dedicated mezzanine boards (MB) manage groups of 20 FEE boards each (Figure 4.22). The mezzanine boards handle the 200 V HV distribution and the regulated (low) voltages needed by the FEE circuitry, obtained through the use of synchronised switch mode step down converters (LTM8033). Communication with the ADCs and DACs on the FEE boards happens trough SPI protocol, allowing the reading/writing of individual bias voltages, along with the monitoring of temperatures and dark currents of each board. The step-down regulators and the SPI communications are handled by an ARM LPC1788 Cortex-M3 microprocessor. The onboard flash memory of the ARM microcontroller stores all the individual calibration parameters for each SiPM and FEE board, such as board and SiPM IDs, high-voltage 2-point calibration (slope and intercept) and SiPM bias setpoints.

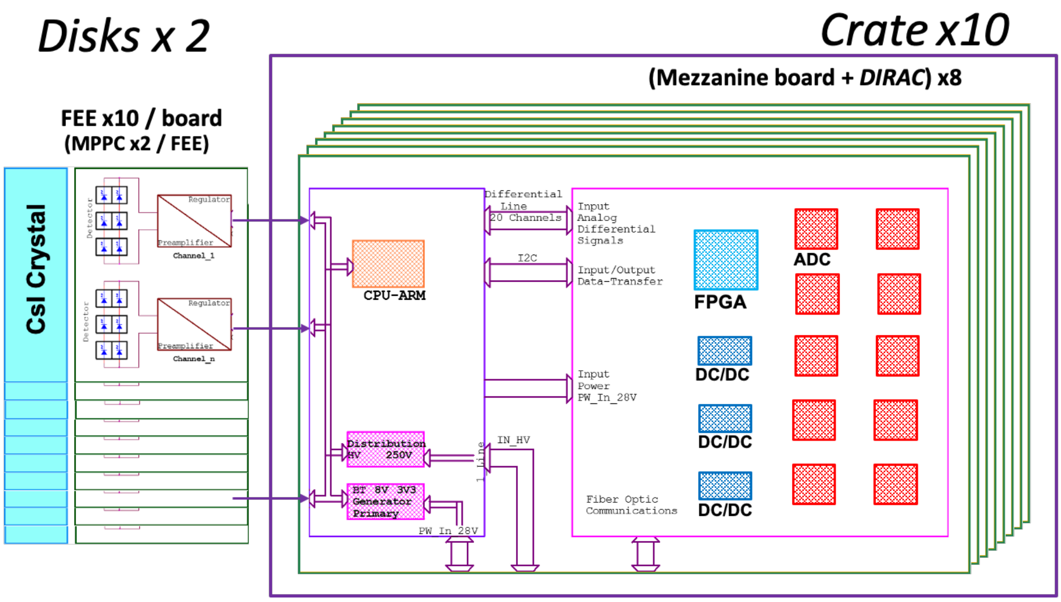

Figure 4.22 FEE, MB and DIRAC architecture.

The DAQ system has to digitise the 2696 fast analogue signals generated by the relative FEE boards and relayed trough the MBs (Figure 4.22). These signals feature SiPM pulses with a 150 ns maximum width, a minimum 25 ns rise time and a 2 V dynamic range. The DAQ system comprises a total of 140 DIRAC boards, mounted in 10 crates for each disk. By relying on simulations and on the analysis of test beam data, the 200 Msps sampling frequency and the 12-bit resolution of the digitisation ADCs were chosen as a good compromise between performance, power, costs and data through-put.

The DIRAC (Figure 4.24) architecture revolves around a MicroSemi SmartFusion2 M2S150T FPGA SoC (Figure 4.23), which controls 10 ultralow-power dual-channel 12-bit ADS4229 ADCs, which are operated at 200 Msps (instead of the maximum 250 Msps available sampling rate). A LMK04828 ultra-low noise clock jitter cleaner provides the board clocking tree (jitter needs to be reduced below 100 ps). The FPGA controls the ADCs protocol and timing, sparsifies and compresses the digitised data, and assembles the packets which are sent optically to the event builder trough a RJ-5G-SX-DPLX dual-fibre optical transceiver. Online operations such as zero suppression, baseline calculation, mean charge and timing reconstruction are also carried out by the FPGA. The FPGA also interfaces with the ARM controllers on the MB board, and the relative dataflows are forwarded to the Detector Control System via the optical link (Figure 4.23).

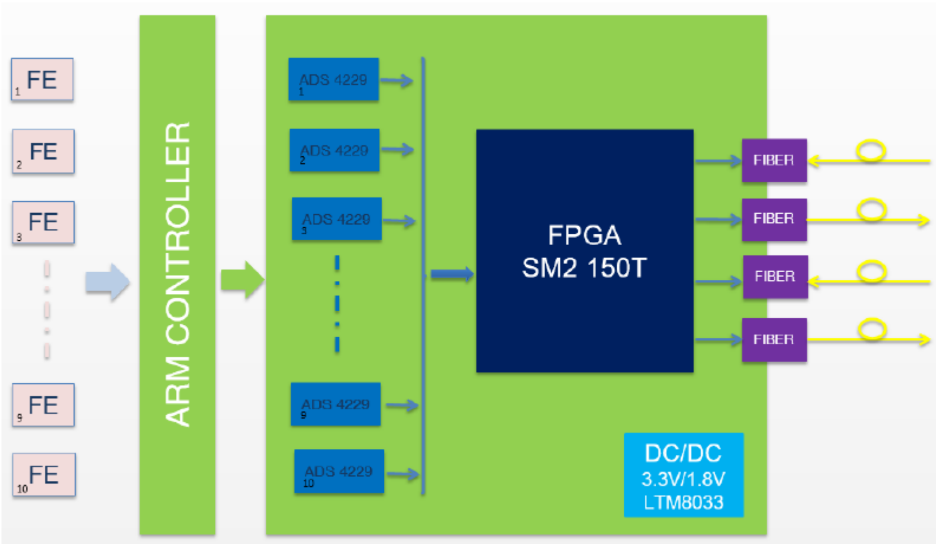

Figure 4.23 Dirac block scheme, along with the ARM controller of the MB.

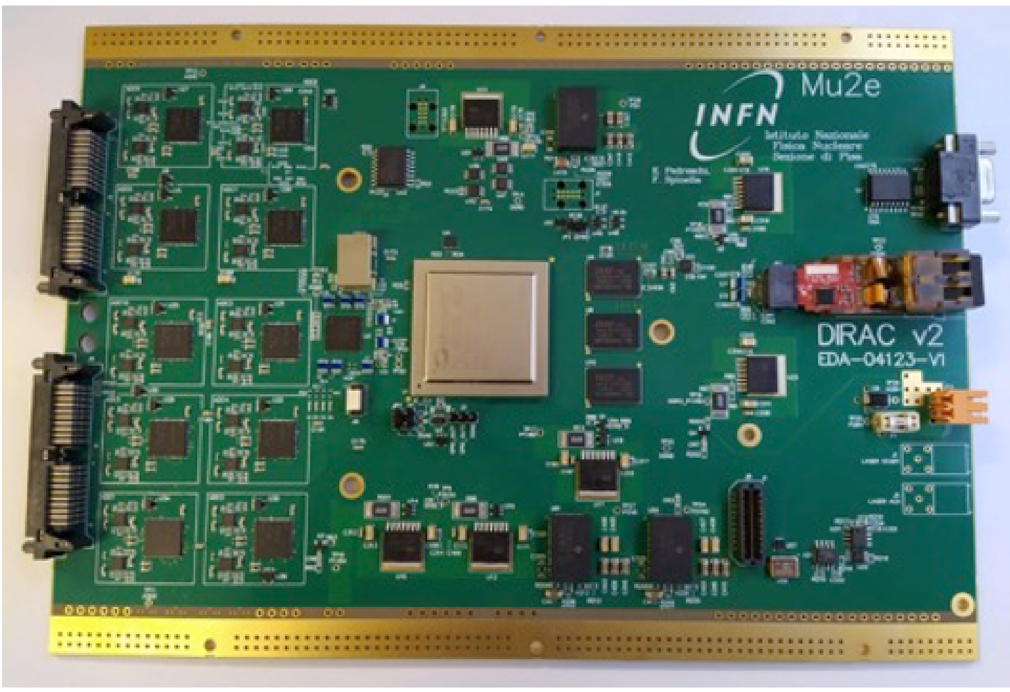

Figure 4.24 Photo of a Dirac board.

5 Radiation hardness

5.1 Radiation environment simulation

The study of the Mu2e radiation environment was carried out via a detailed Geant4 [20] simulation which included the following phenomena:

- beam flash (all particles within the beam other than a muon stopped either in the stopping target or in the beam dump);
- DIO electrons;
- neutrons, protons and photons due to nuclear capture in the stopping target;
- out of target (OOT) particles produced by muons stopped outside the Al stopping target.

5.1.1 Ionising dose

Figure 5.1 reports the expected dose distribution relative to the two calorimeter disks [25]. The averaged values for crystals for the front and back calorimeter disk are shown in Figure 5.2 as a function of their radial coordinate. The expected rate in the hottest region is ~ 1 rad/h. It is clearly visible that the dose values are dominated by the beam flash contribution.

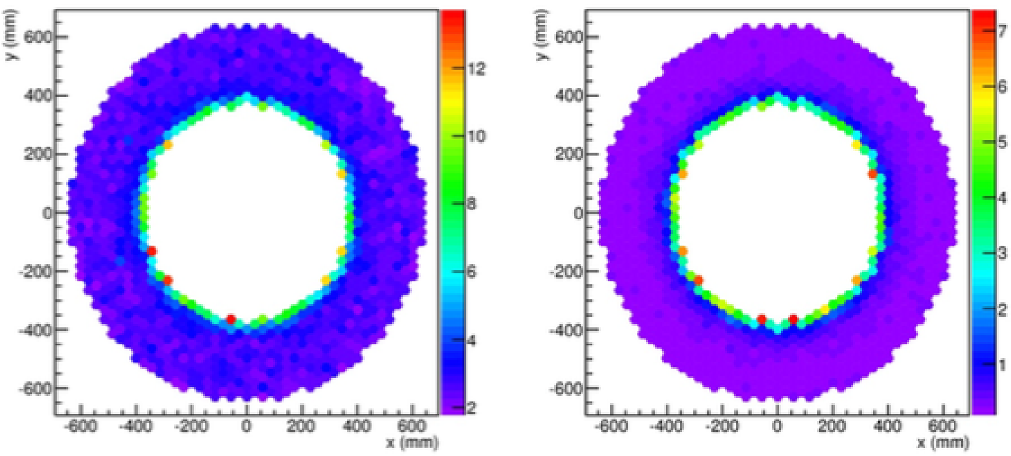

Figure 5.1 Averaged crystal dose map in krad/year over the calorimeter front (left) and back (right) disks.

Mock SiPMs are simulated as Si boxes of equivalent geometry. The deposited dose on SiPMs for varying radiuses is also reported in Figure 5.2. For the back disk SiPMs with a radial coordinate greater than 540 mm, contributions due to OOT muons are comparable in intensity to the beam flash contributions.

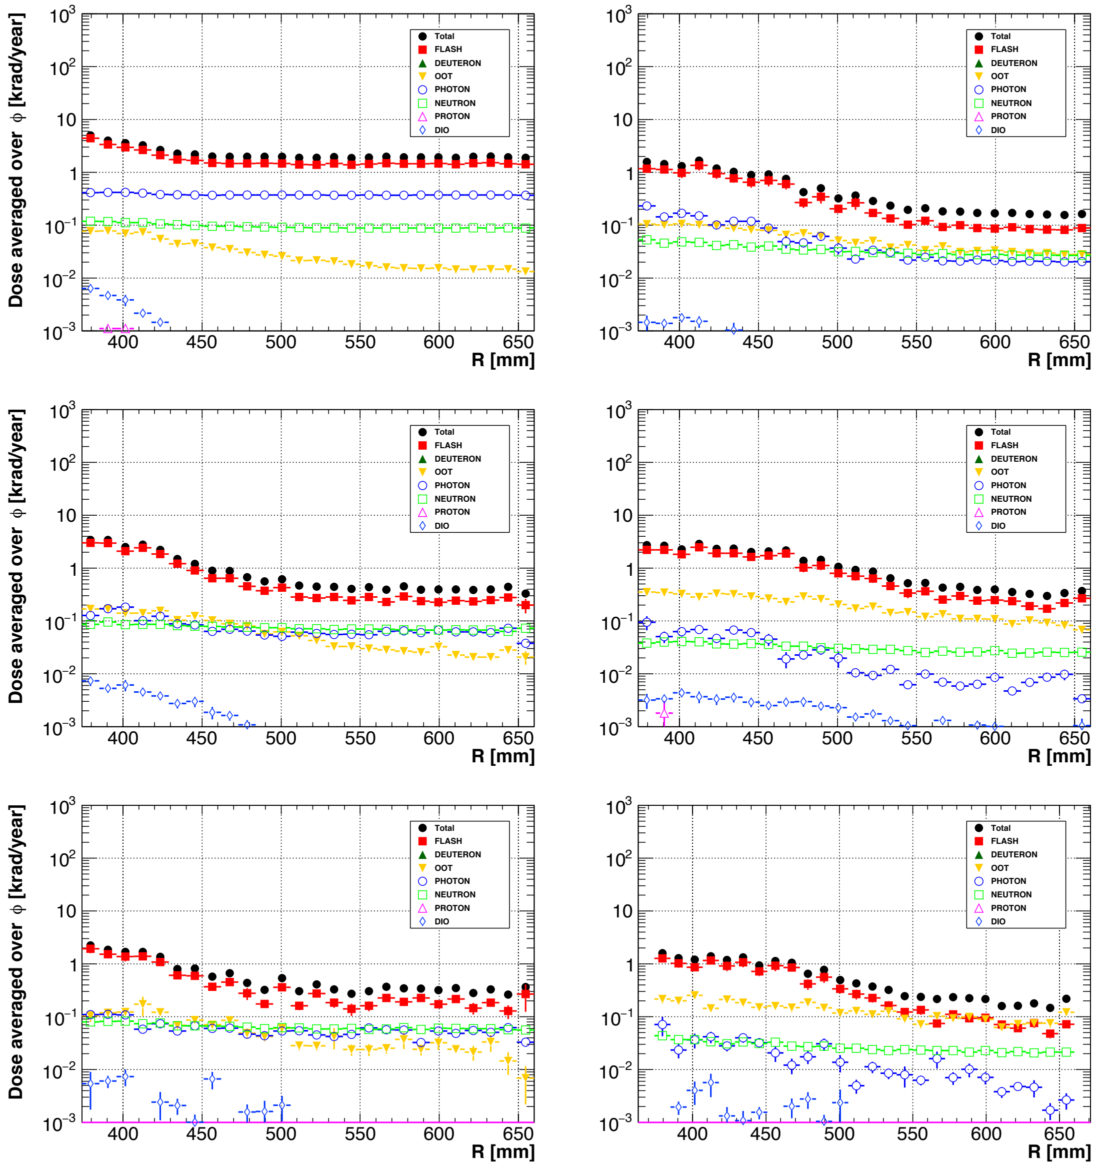

Figure 5.2 Simulated dose for the back (left) and front (right) calorimeter disk, relative to crystals (top), FEE boards (middle) and SiPMs (bottom).

The FEE boards are mocked by boxes of equivalent geometry, containing a compound of Si, Au, Al, Cu and PE. In this case, the averaged values are $3.5 \div 0.4$ and $2 \div 0.3$ krad/year, respectively for the first and second disk.

5.1.2 Neutrons

Neutron induced displacement damage is generally expressed by normalisation to the damage induced by 1 MeV neutrons (referred to as 1 MeV equivalent damage [26]). The expected flux incoming at the front calorimeter disk is around $2 \cdot 10^{11}$ $n_{1MeV}/cm^2/year$. For the second disk, the flux peaks at 10^{11} $n_{1MeV}/cm^2/$ year in the innermost region (not shielded by the first disk) and falls down to $4 \cdot 10^{10}$ $n_{1MeV}/cm^2/year$ in the external one. The maximum expected flux for crystals is inferior to 10^{11} $n_{1MeV}/cm^2/year$. Figure 5.3 shows the equivalent neutron flux expected for the SiPMs of the two disks, as a function of the radial coordinate.

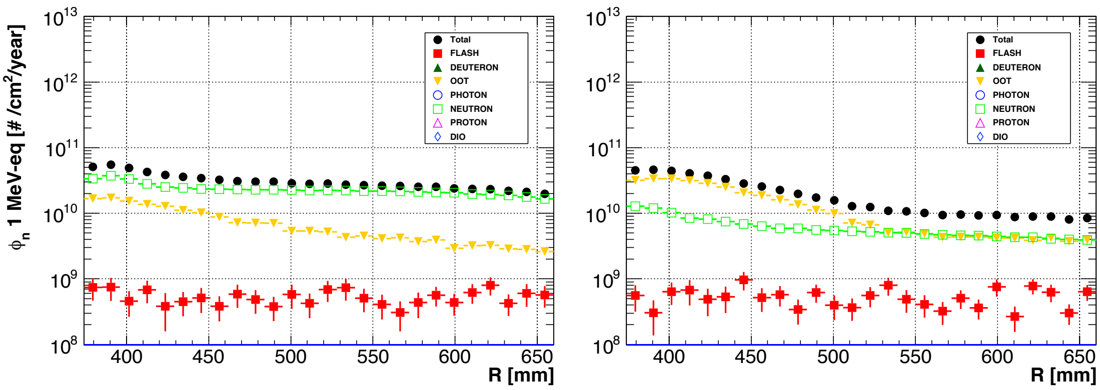

Figure 5.3 Neutron flux relative to the SiPMs of the front (left) and back (right) calorimeter disk expressed as $n_{1MeV}/cm^2/year$ and as a function of the radial coordinate.

5.2 Crystals

5.2.1 Radiation damage mechanisms

All known crystals are prone to radiation induced damage. The most common mechanism [28] is represented by the occurrence of radiation-induced absorption bands due to the formation of colour centres, resulting in the occupation of anion vacancies by unpaired electrons (Farbe centres), or of cation vacancies by holes (V centre). Interstitial anion atoms (H centre) or ions (I centre) also constitute colour centres. A reduction in the crystal attenuation length arises from these defects, along with a decreased light output (LO). Damages may recover even at room temperature, which renders the evolution of the colour centre density (at equilibrium) significantly rate-dependent. As opposed to other types of crystal, thermal annealing and optical bleaching were found ineffective for the elimination of colour centres in CsI [29]. Another element which can degrade the crystal performance is represented by radiation induced phosphorescence (afterglow), which can increase the optical noise in the readout systems. Additional radiation effects might include damage to the scintillation mechanism itself, leading again to a reduced LO and also a modification of the LRU. The LRU degradation mechanism can be induced in crystals which are subject to a non-uniform radiation profile. As opposed to the LO degradation, which can be inter-calibrated, the loss of energy resolution ensuing from an altered LRU cannot be compensated. Thus, to preserve the intrinsic energy resolution of crystals, this parameter must be strictly kept within specification (cf. Section 3.1).

5.2.2 Radiation requirements

As far as crystals are concerned, a maximum total ionising dose (TID) of 90 krad is expected during 5 years of run-time. The maximum 1 MeV-equivalent neutron total fluence expected in the same period of time is $3 \cdot 10^{12}$ n_{1MeV}/cm^2. Trough dedicated irradiation studies [30], it was found that the expected neutron fluence would induce a damage to the CsI properties - quantified in terms of LO and emission weighted longitudinal transmittance (EWLT) - of negligible entity if compared to the TID effects. A ~ 70% decrease in the LO was indeed observed after exposing a 20 cm CsI specimen to 1 Mrad TID. In addition, the radiation induced noise (RIN) due to ionising dose was again found to be more significant if compared to the one due to neutrons (cf. section 5.2.4). As a result, the evaluation and quality control of the Mu2e crystals is carried out solely with respect to a test campaign with ionising dose. The associated requirements include a safety factor of 3 on the doses derived from simulations and are hereby listed:

- a LY better than 85% (60%) of the initial value after exposure to a 10 krad (100 krad) TID;
- a RIN inferior to 0.6 MeV for a dose rate of 1.8 rad/h.

The production crystals will undergo a RIN control run, in order reject out-of-spec ones and identify the ones with the smallest fluorescence component to be installed in the hottest calorimeter areas.

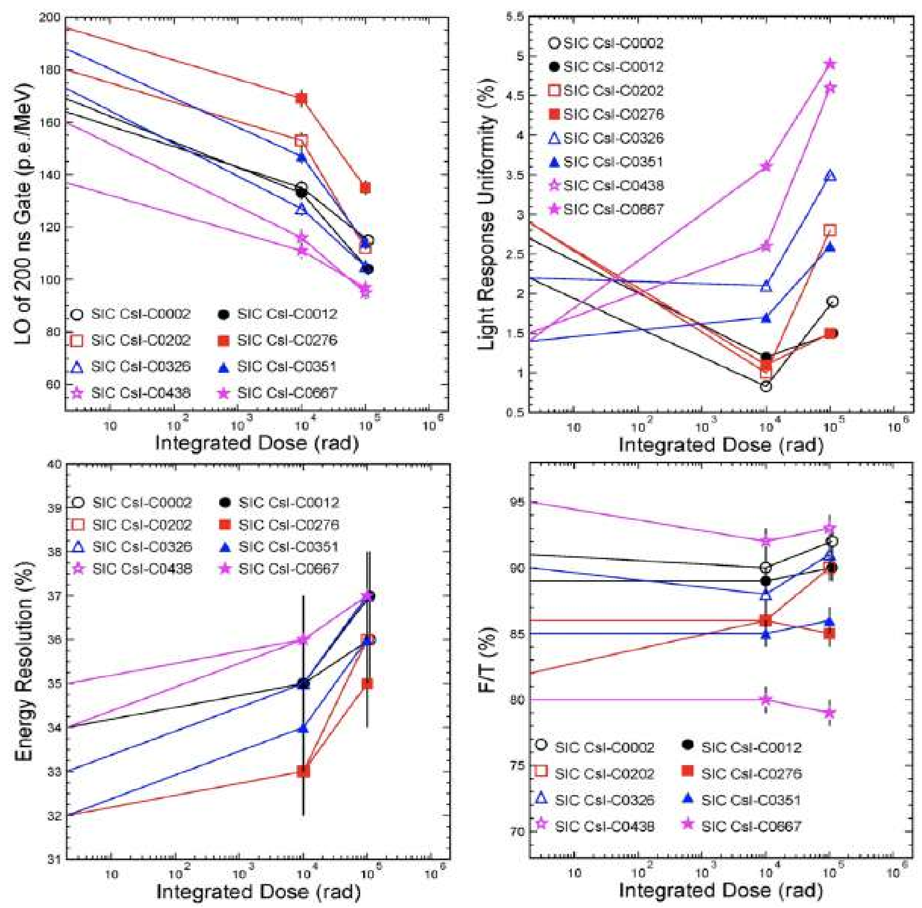

Figure 5.4 Summary of the crystal irradiation test at Caltech.

5.2.3 Irradiation tests

A characterisation run was carried out at Caltech using a two-step (10 and 100 krad) gamma irradiation test with a ^{137}Cs source to verify the optical properties of 16 production crystals. Figure 5.4 summarises the results relative to LO, energy resolution, LRU and F/T ratio after exposure. Only a single crystal out of 16 was unable to comply with the requirements after exposure. An additional irradiation test was performed at the ENEA Calliope facility by uniformly irradiating the crystal samples over their axes using an intense ^{60}Co gamma source. A negligible LY reduction was associated to doses up to 20 krad. The test was continued up to a value of approximately 90 krad, with an associated 20% LY reduction. The shape and the slope of the LRU curves was found to remain similar before and after irradiation.

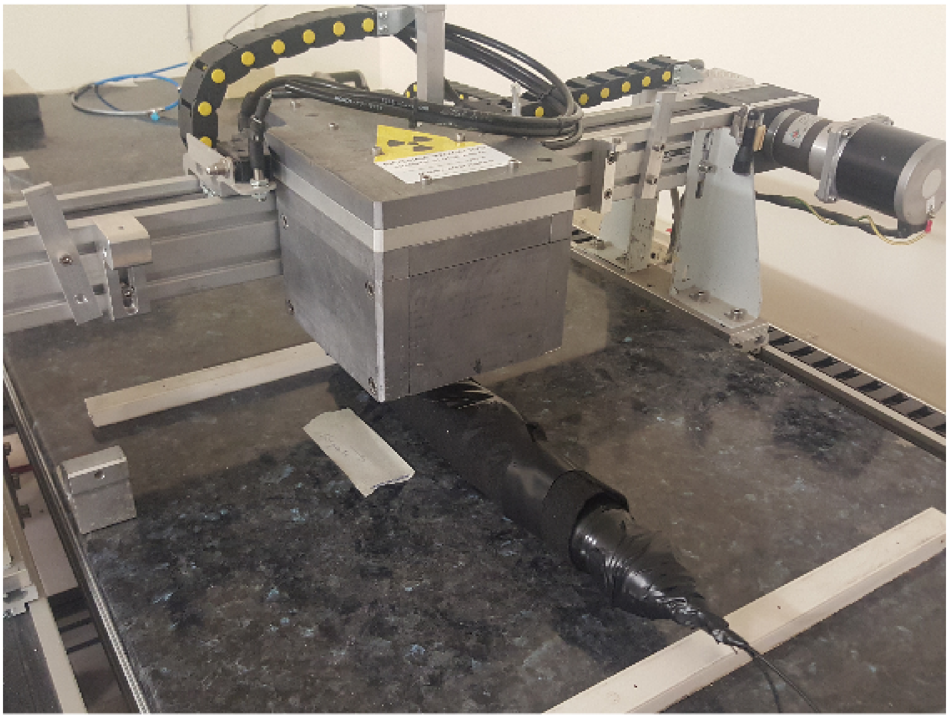

Figure 5.5 RIN test station with ^{137}Cs gamma source.

5.2.4 Gamma RIN test

The RIN was evaluated using 0.66 MeV photons from ^{137}Cs sources, swept over the crystals under test by linear actuator (Figure 5.5). A 1.5 mm diameter collimation aperture is used to reduce the dose rate and to confine the photon emission towards the crystal. Crystals were readout using two Mu2e SiPMs and the respective dark currents were recorded using a picoammeter. For each crystal, the dark current was sampled for 5 minutes without exposure to the source. This step was followed by a 15 minutes long measurement, under gamma irradiation, of the photocurrent I_{RIN}. Finally, an additional 15 minutes long acquisition of the photocurrent was performed to monitor the crystal decay time [31]. Figure 5.6 reports an example of the process.

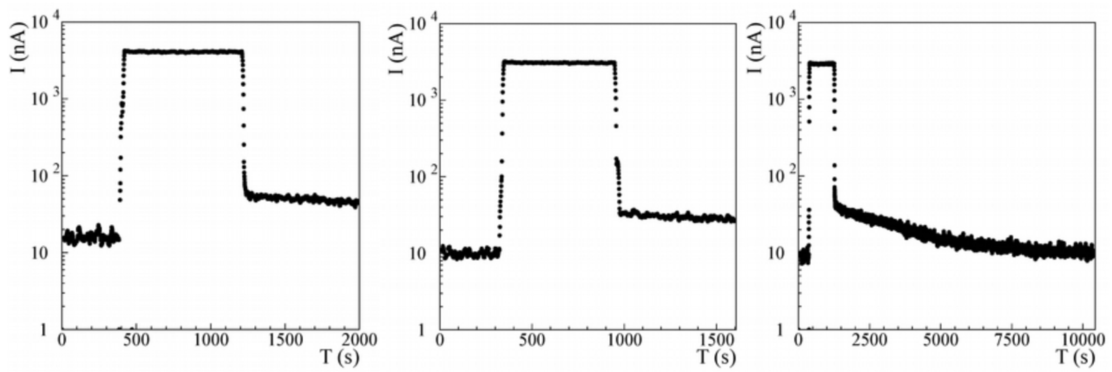

Figure 5.6 Typical photocurrent outputs relative to RIN tests.

The number of photoelectrons induced by radiation per unit dose rate is given by the following relation, normalised over the source rate $\varphi_s = 42$ mrad/h:

$$F = \frac{I_{RIN}}{G \cdot q_e \cdot \varphi_s} \tag{5.1}$$

The number of background photoelectrons due to RIN expected in the Mu2e operating environment $(\varphi_e \sim 1.8 \text{ rad/h})$ can be extrapolated by scaling and integrating over the typical 200 ns gate:

$$N_{pe_{RIN}} = F \cdot \varphi_e \cdot (200 \text{ ns}) \tag{5.2}$$

The equivalent value in MeV is then evaluated considering the statistical fluctuations of $N_{pe_{RIN}}$ and the typical LY:

$$RIN \;=\; \frac{\sqrt{N_{pe_{RIN}}}}{LY} \tag{5.3}$$

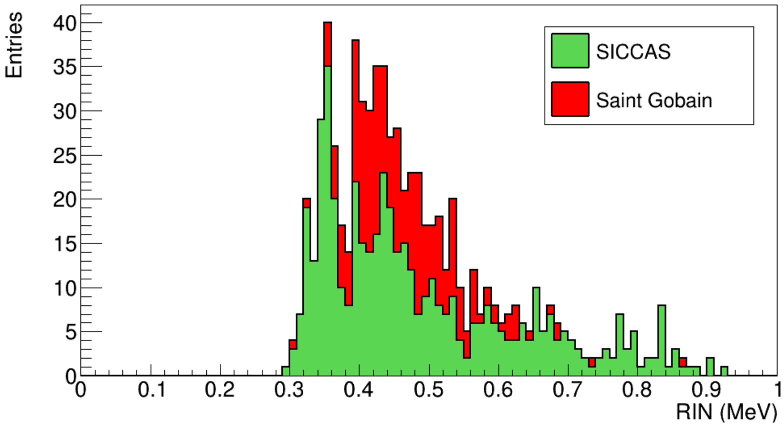

Figure 5.7 Summary of the RIN test results.

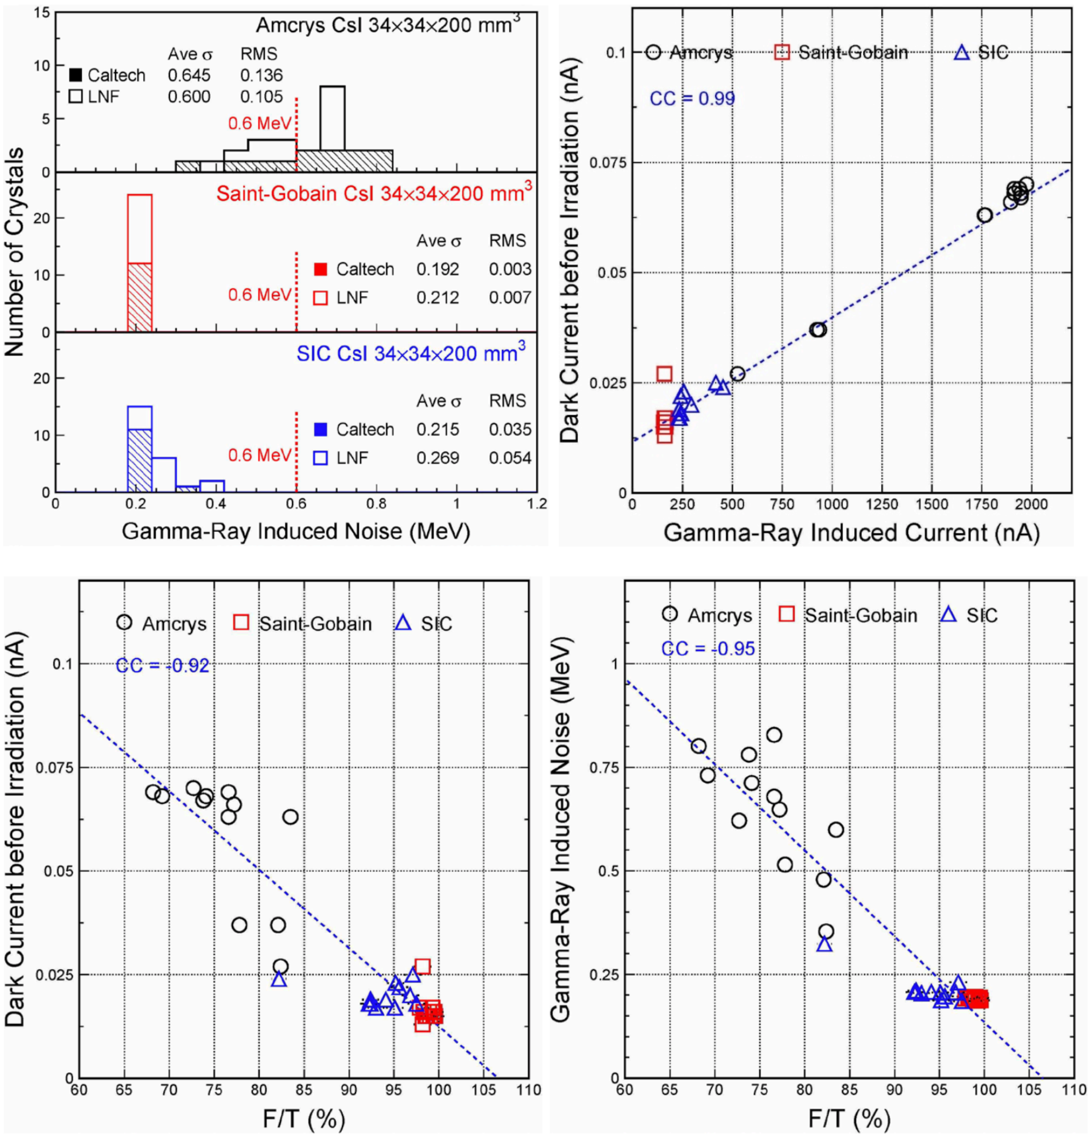

Figure 5.8 Caltech/LNF summary of the RIN test campaign.

A summary of the RIN test results relative to the whole crystal production is reported in Figure 5.7. Furthermore, a linear correlation was clearly observed between the initial dark current and the photocurrent measured during irradiation (Figure 5.8, top). In addition, the correlation with the F/T (Figure 5.8, bottom) proves that the suppression of the slow component results in a reduction of the RIN.

5.3 Front-end

5.3.1 Damage mechanisms in semiconductors

In semiconductors, radiation effects can be classified either as bulk or surface damage [32]. The first class of phenomena is associated to an alteration of the semiconductor structure via the introduction of deep levels defects acting as traps, leading to an increase in leakage current and noise, an alteration of the effective space charge concentration and finally to a decrease in the charge collection efficiency and in the carrier mobility and/or lifetime.

Bulk damage occurs mainly in the form of displacement damage (DD) trough particles interaction with the nuclei in the crystal structure [27]. For Si atoms, a minimum recoil energy of 25 eV is needed to achieve displacement of nuclei. This translates into a minimum energy of 260 keV for electrons and of only 190 eV for protons or neutrons. DD is associated to the non-ionising energy loss (NIEL), which is dependent on the mass and energy of the incident particle: this mandates for damage data to be always referred to the specific particle type and energy considered (DD tests are generally performed using protons or neutrons).

Surface damage is instead associated to the generation of defects in the overlaid dielectrics and at the Si/dielectric interfaces. It is responsible for the generation of oxide charges and interface traps, associated to alterations of break-down voltages and inter-electrode isolation and/or capacitance, potentially also impacting on the charge collection performance. Moreover, radiation effects can also include a permanent alteration of the materials resistivity.

5.3.1.1 Neutron and displacement damage

Neutron-matter interactions are primarily associated to non-ionising processes, such as DD, typically due to scattering: the dominant phenomena (from 50 keV to 14 MeV) are neutron-neutron elastic or inelastic scattering, as well as neutron-proton and neutron-alpha interactions. The NIEL for these processes can be associated to about half of the neutron energy transferred to the nucleus of the material [27]. As neutrons transfer momentum to the matter atoms, displacement of the atoms from their original lattice positions can occur, resulting in vacancies. Displaced atoms can eventually settle in interstitial positions. In a single interaction, several atoms can be displaced at once, which in turn can lead to the displacement of other atoms in a cascading process, resulting in the formation of a cluster of defects. Vacancies and interstitials will migrate and eventually either recombine (approximately 90% of the times) or form stable defects (Frenkel pairs) [32].

In semiconductors, these defects in the lattice periodicity are able to generate deep levels within the bandgap [33], thus producing a modification of their (electro-optical) properties, associated to following classes of phenomena:

- production of recombination-generation centres with an ensuing increase in leakage current and noise;
- alteration of the carrier's lifetime (for example by acting as recombination centres and thus degrading gain figures in BJTs);
- charge-transfer processes and charge collection efficiency thwarting;
- compensation of donors or acceptors (thus altering properties related to carrier concentration, such as bias voltages);
- promotion of electron tunnelling, which causes reverse currents in junctions;
- defects can act as scattering centres for charge carriers, thus decreasing their mobility.

5.3.1.2 **Charge trapping**

Particle impinging on a semiconductor can lose energy trough primary or secondary ionisation processes, leading to the generation of electron-hole pairs along the particle path. Strongly ionising particles are associated to dense charge columns and a higher recombination rate, while weakly ionising ones are able to generate isolated charge pairs which are less likely to recombine [34].

For example, in mosfet devices, if a bias E-field is present, the electrons will drift towards the gate and the holes will move towards the oxide-Si interface, partially recombining. This initial recombination has a strong dependence upon the bias field applied and the incident radiation. As far as the E-field is concerned, as its strength increases, the recombination rate decreases. The fraction of charge carriers able to escape recombination is known as charge yield. These carriers are eventually trapped in the oxide layers or can build-up at Si/oxide interfaces (such as circuit passivation layers or gate oxides).

Holes will travel (with a smaller mobility if compared to electrons) towards the gate/oxide (oxide/Si) interface when a negative (positive) bias is applied, causing a distortion of the oxide potential field: this increases the trap depth in a localised site (oxide trapping). In addition to oxide traps, positive, negative or neutral traps can occur within the silicon bandgap at the Si/oxide interface. Traps in the (upper) lower portion of the band gap can act as (acceptors) donors: for (n-mos) p-mos devices, (negatively) positively charge trapped at the interface is associated to a (positive) negative threshold voltage shift, as it produces a progressive (inhibition) activation of the gate. With increasing dose, n-mos (p-mos) devices will eventually fail in a permanently activated (deactivated) state. Moreover, in n-mos (p-mos) devices, charge trapped at the interface will compensate (worsen) the effects of the oxide trapped charge, which is positive for both cases.

Charge build-up in dielectrics are associated to cumulative effects related to increases in timing margins and/or modifications in the leakage or subthreshold currents. Annealing due to tunnel or thermal effects can lead to charge de-trapping, thus restoring device functionality.

5.3.2 **FEE TID test**

The calorimeter FEE boards are required to withstand a 100 krad TID. In January 2019, after an intense R&D campaign, the final choices for the rad-hard ADCs and DACs of the FEE boards (cf. Chapter 4) were tested up to 100 krad at the ENEA Calliope facility (Figure 5.9) using 1.25 MeV photons from a ^{60}Co source. The component choice was successfully confirmed. In November 2019, a stability test was carried out on the complete FEE boards with a TID of 105 krad at Calliope [35]. Test results showed stability values of 0.3% for the DAC and of 3% for the ADC. The 0.3% DAC variation alone is unacceptable, as it translates to a 480 mV change on the biasing applied to the SiPM, which far exceeds the gain stability specifications. In addition, mosfet M3 was reported to fail when the boards were irradiated with a 0 V setpoint on the regulated HV supply output.

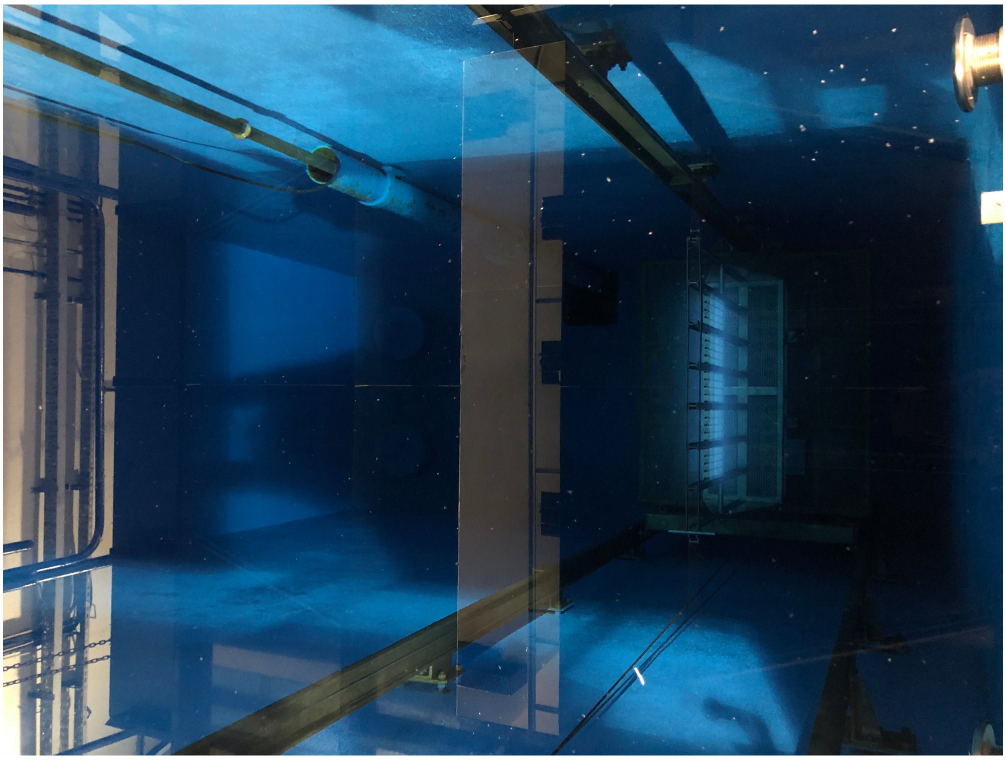

Figure 5.9 ENEA Calliope facility: ^{60}Co source pool. The source is arranged in 48 cylindrical bars. The activity of the source during the tests was $0.35 \cdot 10^{15}$Bq.

To address all these issues, a final test was carried out in June 2020 at Calliope. The tested FEE boards presented minor revisions: unused ADC channels (6, 7 and 8) were connected to a resistor ladder in order to sample 0.05%, 50% and 95% of the supply voltage, respectively (cf. Chapter 4). These resistor ladders will be included in the final PCB version as a reference tool for on board diagnostics. Moreover, the value of the ADC input resistors (R56, R108 and R40 in Chapter 4) was lowered from 1 kΩ to 100 Ω, for reasons discussed in the following sections.

During the test, two lots of DACs (lots 1 and 2) and two lots of ADCs (lots A and B) were tested in 6 different configurations with 6 different FEE boards (Figure 5.10). Channels 1 and 2 had a 160 V setpoint on the HV regulator and a resistive load dissipating respectively 2 and 1 mA. The setup of channel 3 setup was identical to the one of channel 2, but with different components lots.

Channel 4 and 6 were configured with a HV setpoint of 20 and 10 V respectively, and no load connected. Channel 5 had no HV applied to the input of the regulator, but 1.5 mA, 200 ns current pulses were applied to the regulator output. During the test, the boards were exposed to a 103 krad TID, with an average 0.85 krad/h dose rate.

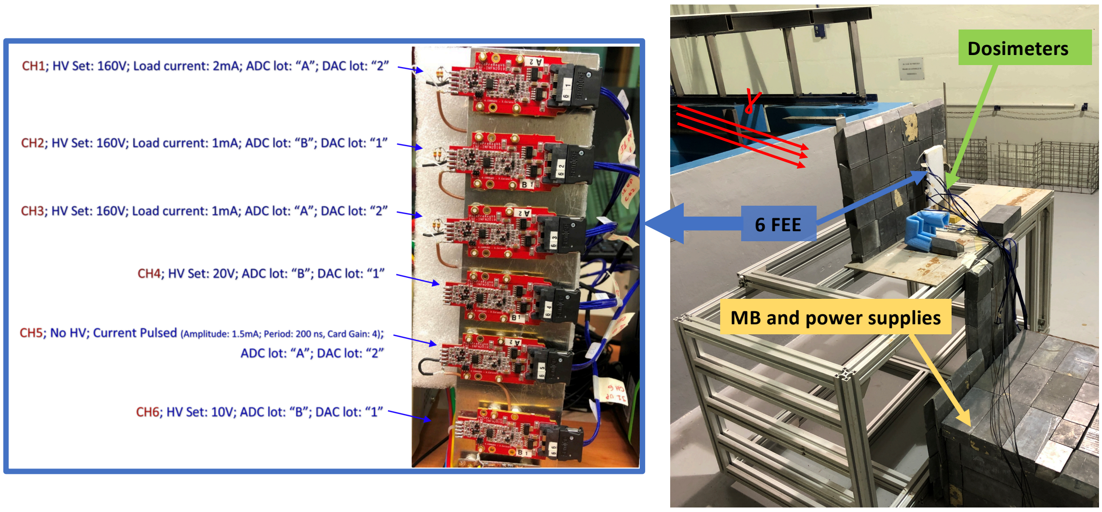

Figure 5.10 Experimental setup of the TID test at Calliope. Left: detail of the FEE boards under test. Right: Calliope experimental hall.

5.3.2.1 ADC and DAC performance

During exposure of the FEE boards, the average of 10 successive ADC readings was logged every 2 seconds. The observed standard deviation on oversampled and averaged data was $\leq 1\ LSB$ (with respect to the 4096 full scale counts). Thus, by assuming a white quantisation noise, the standard deviation on raw data is expected to be approximately $\sqrt{10}$ times larger.

During the test, the HV monitoring ADC channel showed excellent stability (Figure 5.11, top left), with a maximum absolute deviation of 2 bits (with respect to 4096 full scale counts) and a coefficient of variation better than 0.03% for all channels (calculated as the ratio between the standard deviation and the mean of the ADC readings). This improvement was due to the modification carried out on the ADC input resistors. In fact, the previously discussed drifts observed during the Calliope 2019 test were related to a dose-dependent increase of the input bias current of the programmable gain amplifier input stage on the ADCs.

By lowering the ADC input resistors (R56, R108 and R40) from 1 kΩ to 100 Ω, these detrimental effects were successfully tamed.

Good stabilities were also recorded with respect to the current and temperature monitor ADCs, the latter showing a synchronous trend for all channels which follows the day-night temperature fluctuations (Figure 5.11, top right). As previously states, ADC channels 6, 7 and 8 were connected to a resistor ladder in order to respectively sample 0.05%, 50% and 95% of the supply voltage, in order to provide an absolute indication of the ADC stability at different positions along its dynamic range: a maximum deviation of 4 bits was observed for all channels (Figure 5.12).

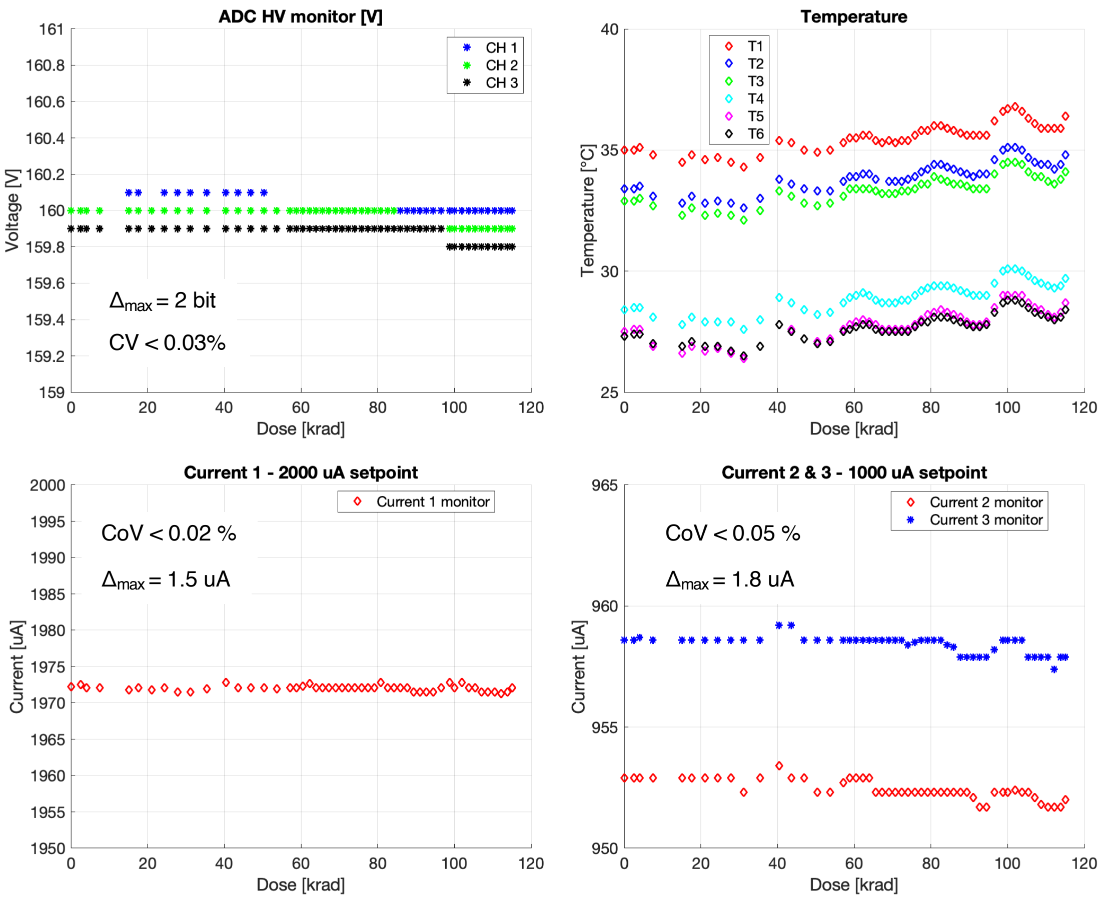

Figure 5.11 Stability of the HV (top left), temperature (top right) and current (bottom) monitoring ADCs during exposure at Calliope.

In conclusion, during irradiation the ADC showed a worst-case 0.1% drift, which is within the design parameters and the components tolerances. These variations are comparable to the stability of the non-irradiated boards.

It should also be noted that the choice of rad-hard ADCs and DACs for the final FEE design has forced the installation of 12-bit components, as opposed to the original 16-bit (non rad-hard) choice: 1 LSB corresponds to a minimum step of 50 mV on the HV regulation, in contrast with the original requirement of a 10 mV target precision. This element does not represent a significant issue for the design. In fact, a 50 mV variation on the Mu2e SiPM setpoint corresponds to a 17 mV variation for each cell, with an associated gain variation is of 1.7% (in linear gain-bias regime, cf. Section 4.2). From an operational point of view, this O(2%) channel-by-channel variation will have no impact on the physics, as calibrations are due weekly with an equalization level better than 0.5 %.

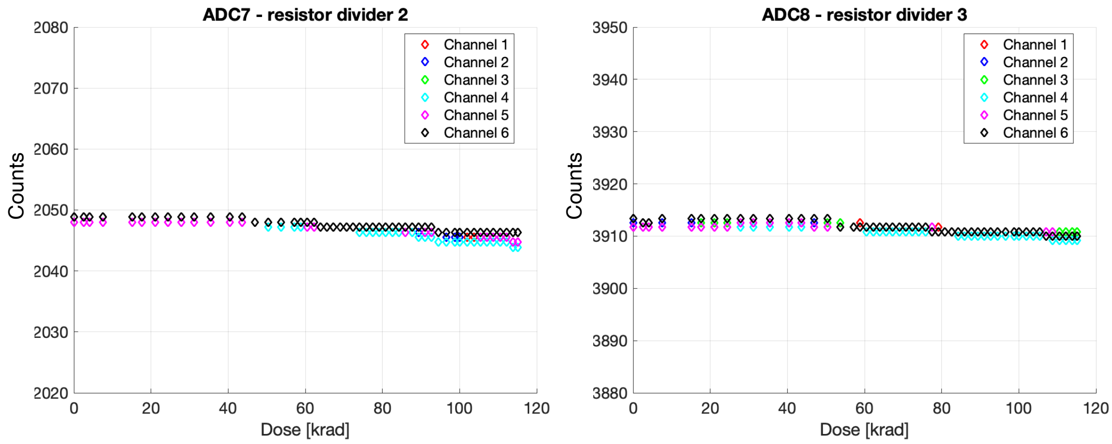

Figure 5.12 Reference resistor ladder sampled by ADC channels 7 and 8 during irradiation.

5.3.2.2 Mosfet charge trapping

During the test campaign at Calliope, the only operating condition in which the linear regulator section was found to be subject to failure under a 100 krad TID was related to the FEE being irradiated while powered up, with the HV applied at the linear regulator input, and with a 0 V setpoint on the regulated voltage. Even a set-point of 10 V on the regulator were sufficient in order to prevent failure of the boards. In fact, the June 2020 Calliope test demonstrated that, while the CH5 board failed as expected, even a 10 V setpoint on the HV linear regulator (CH4 and CH6) is sufficient to prevent failure (Figure 5.11). The culprit of this failure mode is the mosfet M3 which fails in an always-on state due to charge trapping. An annealing test carried out on an unpowered damaged board at 90 °C for 24 hours was able to restore M3 functionality, thus confirming the hypothesis.

A 0 V setpoint on the linear regulator translates into mosfet M3 being in saturation (cf. Section 4.3.2). The high E-field on the gate promotes charge trapping (as it increases the charge yield of carriers generated by the ionizing particles trails which escape the initial recombination). This produces a progressive gate activation of this n-mosfets (the threshold voltage is shifted negatively). The n-mosfet M3 will eventually fail in an always-on state, thus preventing the regulator from turning on and providing any regulated voltage output.

In any case, during the Mu2e experimental runs, the HV regulators on the FEE boards are expected to be always kept on, even when SiPMs are deactivated, so that the conditions which lead to M3 failure should never be met. Additionally, the expected dose rate will be several orders of magnitude lower if compared to the one provided in the irradiation tests at Calliope.

5.3.3 SiPMs

5.3.3.1 Total ionising dose tests

Irradiation tests were performed on SiPMs at ENEA Calliope using 1.25 MeV photons from a ^{60}Co source. A single SiPM was exposed over the arch of three days to a TID of 20 krad.

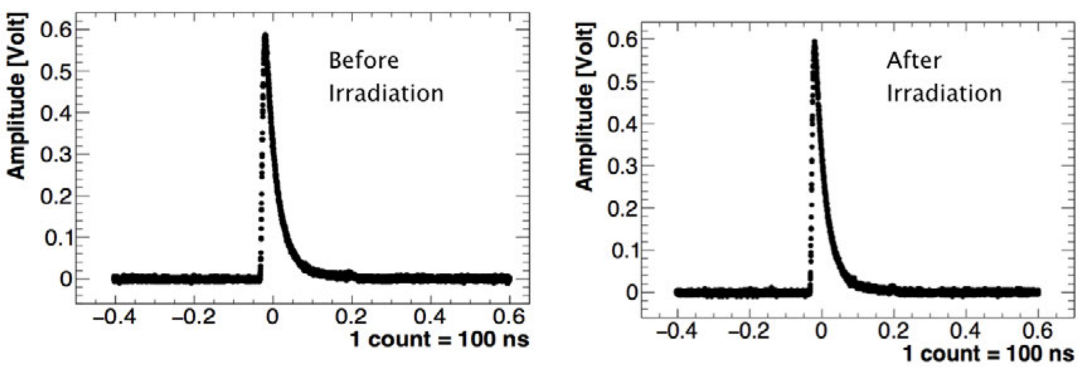

Figure 5.13 Mu2e SiPM pulse shape before and after irradiation at Calliope. A pulsed LED source was used.

An increase from 0.15 to 0.6 μA was observed in the dark current as soon as the irradiation started, due to Compton scattering on the active surface. During the exposure time, the dark current value showed of an additional 0.15 μA increment, which is of negligible entity with respect to the Mu2e requirements, also considering the much larger entity of the neutron induced detrimental effects on this parameter. Moreover, it was verified that the gain, as well as the signal amplitude and shape were preserved after irradiation (Figure 5.13). This test campaign's results confirmed the negligible effects of ionising dose on the SiPM performance in comparison to neutron damage, as discussed in the following section.

5.3.3.2 **Neutron damage test**

Mu2e SiPMs are required to withstand a TID of 45 krad and a total equivalent fluence of $6.5 \cdot 10^{11}$ n_{1MeV}/cm^2. Neutron irradiation tests were performed at the EPOS facility of the HDZR [36], where the interaction of a 30 MeV electron beam with a 10 mm Tungsten target is used to produce photons and neutrons. The SiPMs under test were placed on top of the shielding roof, in a zone in which the photon contribution to the dose is negligible by a factor of 10^{-4} (Figure 5.14). Three SiPMs were exposed to a total fluence of $8.5 \cdot 10^{11}$ n_{1MeV}/cm^2 over the course of 29 hours [37]. For each SiPM under test, 5 out of 6 cells were left unbiased: only one was held at the bias voltage and its dark current was continuously monitored.

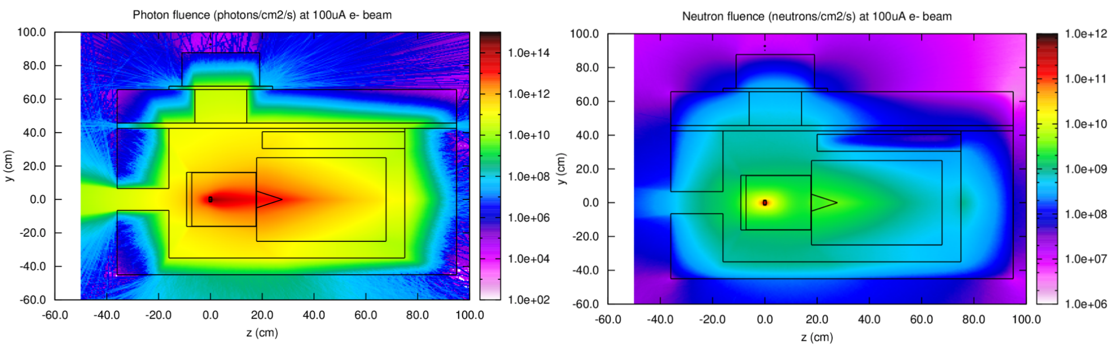

Figure 5.14 Map of the neutron and photon fluences for an electron beam current of 100 μA at EPOS-HZDR.

The dark current behaviour versus the total fluence is reported in Figure 5.15 (left). Up to a ~ $3.5 \cdot 10^{11}$ n_{1MeV}/cm^2 fluence, the correlation is well modelled by a linear function. After the end of the test, the unbiased cells presented a bias current slightly larger if compared to the biased ones, as a consequence of thermal de-trapping [38]. The bias current acquisition was continued for 22 hours after the test was ended, in order to monitor the self-annealing process (Figure 5.15, right). The bias current had decreased by approximately 50% in a follow-up measurement performed two months after the test

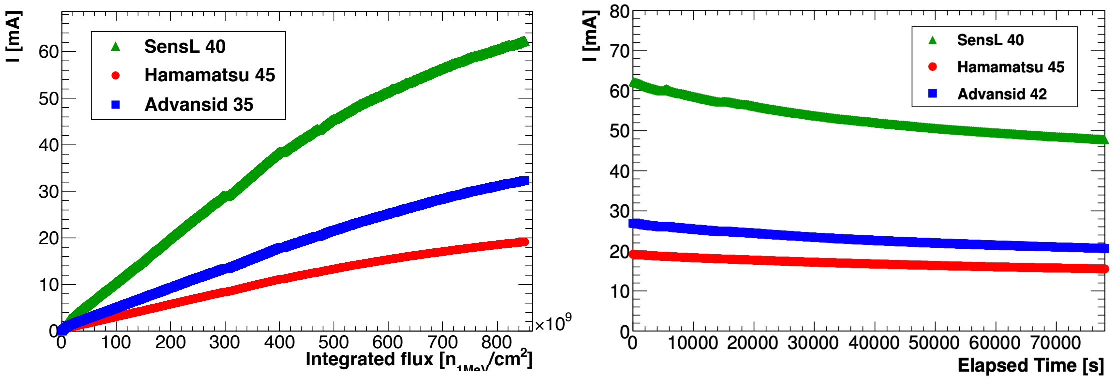

Figure 5.15 Left: dark currents of individual cells as a function of the equivalent fluence. Right: cell dark currents as a function of the annealing time after the irradiation test.

A characterisation of the dark current dependence on temperature and bias voltage was also carried out on the irradiated SiPMs. In a light-tight climatic chamber fluxed with nitrogen, temperature was ramped down from 20 to -5 °C, simultaneously decreasing the bias voltage of 0.1% per °C (according to specifications), in order to keep the operating point unchanged. This regulation was verified by checking the breakdown voltages at 20, 10 and 0 °C. The measurements are reported in Figure 5.16 (right) and show a halving of the bias current for a temperature decrease of 10 °C. The bias current dependence over the biasing point was characterised at 0 °C. Results are shown in Figure 5.16 (left) as a function of the overvoltage.

Trough dedicated measurements carried out on the Module-0 (cf. Section 3.5), the effect of the SiPMs dark current on the RIN was evaluated: it can be estimated that dark currents of 500 and 1500 μA correspond respectively to 1 and 1.73 MeV of equivalent noise.

As a result of this test campaign, an operational temperature of 0 °C was set for the Mu2e SiPMs. In order to cope with the effects of radiation exposure, a reduction of the nominal overvoltage and of the operating temperature of the photosensors after the first period of run-time have been taken into account as a viable strategy to achieve the 5 years run time target, or equivalently to keep the dark current below the 2 mA design limit of the FEE.

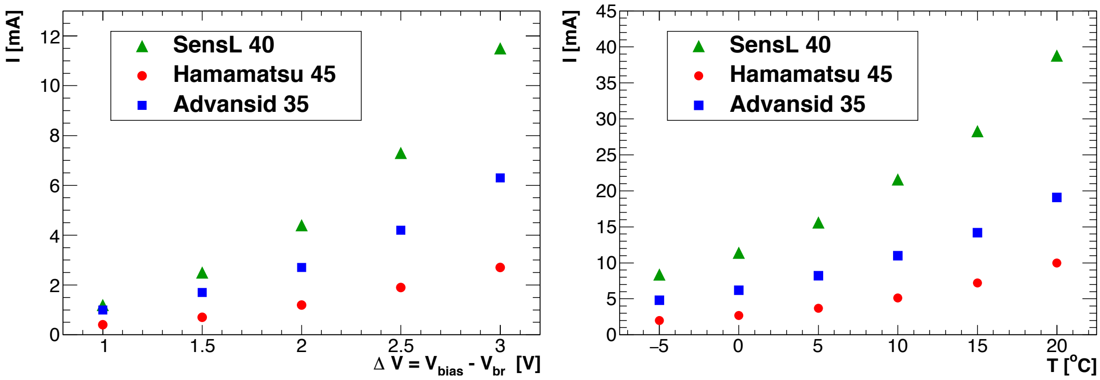

Figure 5.16 Dark current dependence on the overvoltage (left) and on temperature (right) after irradiation.

6 Cosmic rays calibration and Cosmic Ray Tagger

6.1 Cosmic muon calibration

Cosmic rays (CRs) represent a great tool to perform calibration runs of the Mu2e calorimeter, as the specific energy loss of MIPs is uniform throughout the whole calorimeter active area, thus allowing a precise energy scale equalisation of each channel. Thanks to their relativistic nature, MIPs are also used to compensate the relative timing offsets of all channels without the need for any external time reference. Moreover, the flux of cosmic ray particles is intense enough to enable the collection of a sufficient amount of calibration data in a short period.

6.1.1 MIP simulation and selection

Simulations have been carried out using a MC cosmic rays generator developed by the Mu2e collaboration, which implements the Gaisser model, along with corrections to account for muon decay rates and Earth curvature [46]. The simulation also embeds the effects due to the presence of the external neutron shield of the DS (2.4.6) and of the 1.8 metres thick concrete ceiling of the experimental hall. These corrections are implemented to precisely account for the conditions encountered during the online MIP calibration procedure carried out on a daily basis during the experimental run time (as described in section 3.3.2.2).

In the simulations, MIP hits are recorded whenever a minimum 3 MeV deposit is detected in at least one of the calorimeter crystals. This single digitisation threshold results in an upper limit for the detected CR rate of approximately 130 Hz, calculated as the total rate observed on a disk, corrected according to the fraction of total events satisfying the 3 MeV selection. After the calibration cuts reported in the following section, the rate decreases to a value of approximately 15 Hz, which translates to 1000 events per crystal collected during a 5 hours long run.

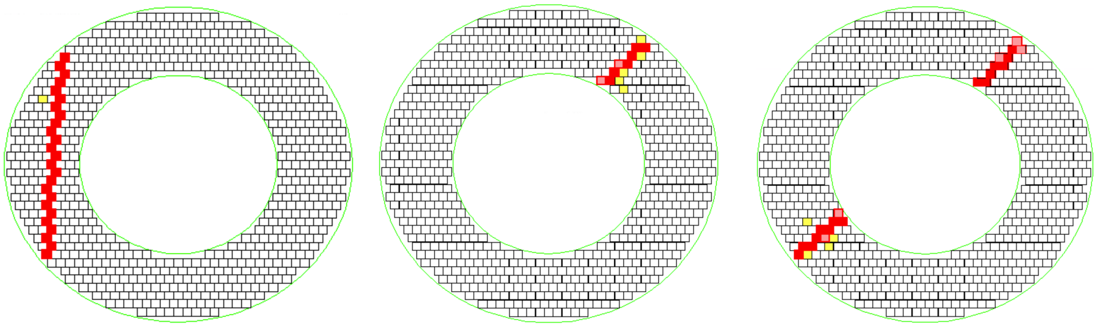

Figure 6.1 MIP clusters according to the three tracks selection categories (a), (b) and (c) from left to right.

For the MIP track analysis, crystal hits are clustered similarly to what is described in section 3.4, and a minimum 6 MeV deposit is required in at least 3 crystals of each cluster. Subsequently, MIP events are sorted in three different categories, according to their trace projection on the transverse plane of the calorimeter (Figure 6.1): tracks secant to the outer disk circumference but external to the inner hole (a), tracks which cross twice the calorimeter active area and are secant to the inner calorimeter circumference (c) and oblique tracks which cross the calorimeter torus once and escape through the inner hole of the disk (b).

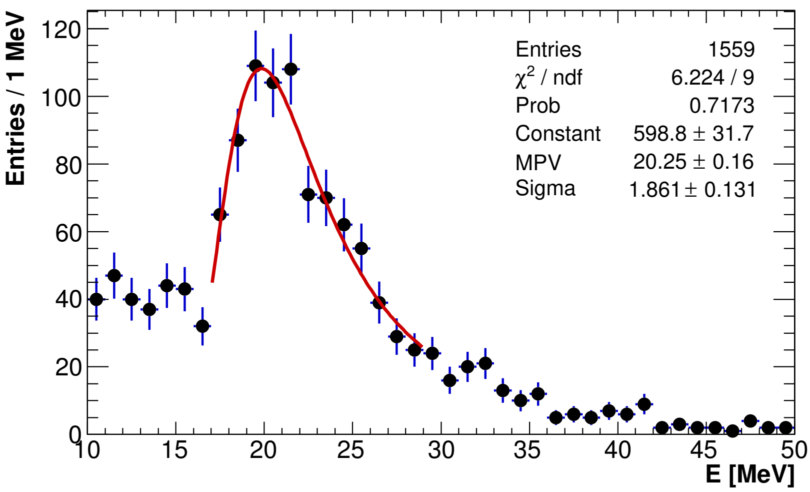

Figure 6.2 Distribution of the energy deposition of cosmic ray muons in a single crystal, as derived from simulations.

MIP tracks are reconstructed by applying a least square linear regression over the clusters. Figure 6.2 shows the distribution of the energy deposit in a single crystal: the peak due to the MIP is clearly visible and it was fit with a Landau function to evaluate the Most Probable Value (MPV). Only clusters having $0.6 < \chi^2/ndof < 1.4$ from the linear regression are included in the calibration dataset. Moreover, individual hits are included only if relative to a crystal distant at most 17 mm (half the crystal size) from the fitted trajectory. Finally, tracks with a total path length in the transverse plane inferior to the crystal lengths are also discarded.

6.1.2 Energy scale

To verify the uniformity of the energy response for MIPs events, the specific energy loss peak shown in Figure 6.3 (right) is evaluated as the MPV of a Landau fit for each channel. Figure 6.4 shows the calorimeter pixel map obtained by applying this procedure on the simulated data.

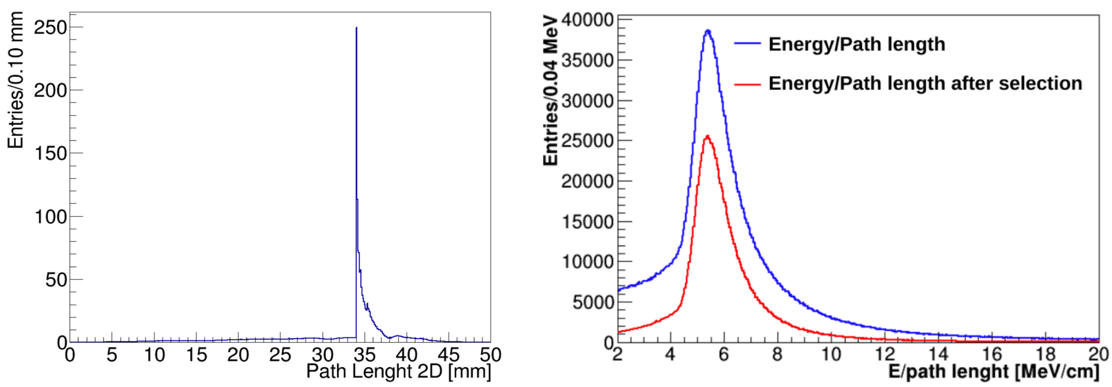

Figure 6.3 Left: bidimensional path length distribution. Right: specific energy loss before and after the application of the selection cuts.

Results among channels are homogenous, showing a $\sim 1\%$ RMS spread compatible with the statistical error of the fit procedure. Moreover, being the Mu2e calorimeter unable to provide any information with respect to the longitudinal Z position of the hits, the specific energy loss is calculated over the path length projection on the transverse plane (Figure 6.3, left). The same level of homogeneity is conserved with respect to the the 3D path case. This allows the calorimeter to be correctly equalised on-line, without the need for an external MIP tracking system. As in this case the path lengths are systematically underestimated, the bidimensional specific energy loss distribution is shifted upwards. From these estimations, the energy scale factor of each channel is calculated against the ~ 6 MeV/cm specific energy loss peak deriving from the simulation. The calibrated values of specific energy losses are distributed with $\sigma < 1\%$ (for both 2D and 3D path lengths).

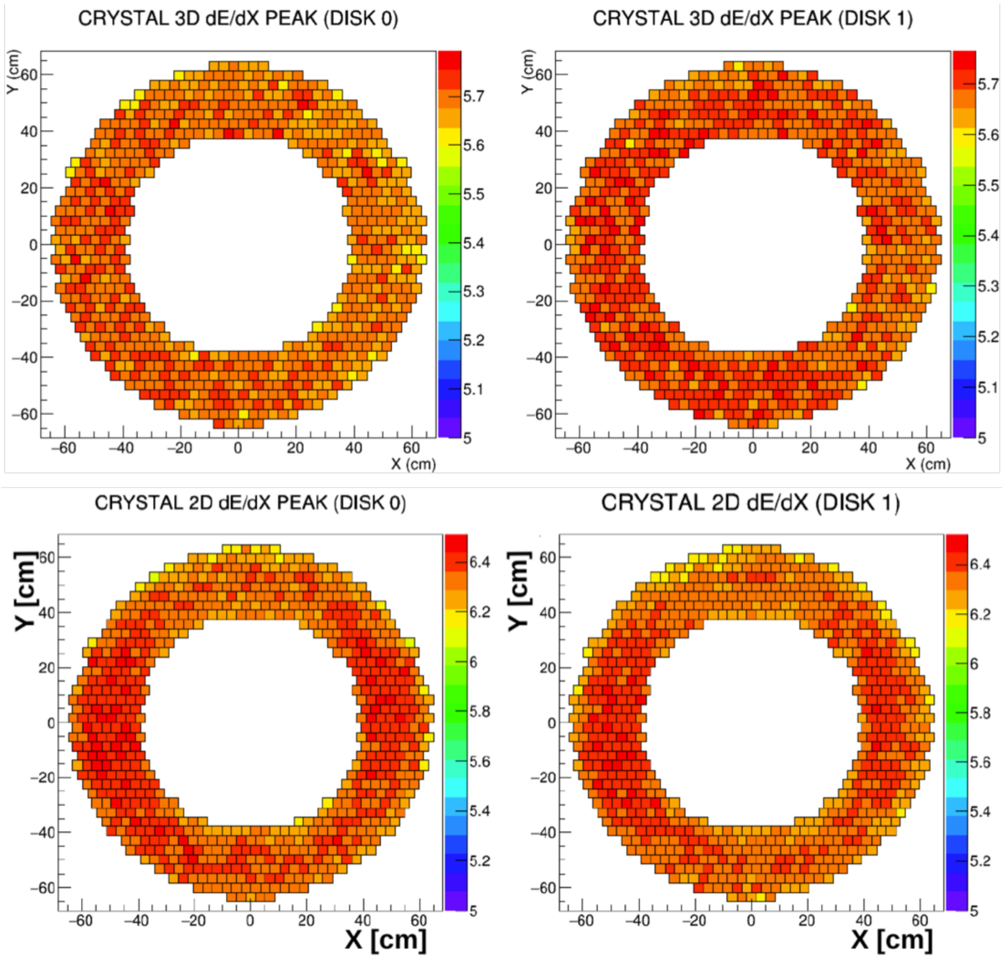

Figure 6.4 Map of the 3D (top) and 2D (bottom) specific energy loss in the crystals of the front (left) and back (right) calorimeter disks obtained through simulation.

6.1.3 Time offset alignment

By exploiting the relativistic nature of MIPs, the compensation of the individual time offsets (T_{0_i}) of each channel can be carried out in order to improve the calorimeter timing resolution. The following procedure is used:

1. hits clustering and 3D (or 2D in the case of on-line stand-alone calibration) MIP trajectory selection, as described in 6.1.1;
2. construction of the calibration dataset with the same cuts implemented in 6.1.1;

3. fit of the individual times of interaction t_i of each channel against the reconstructed MIP time-of-flight for each event in the dataset;
4. extraction of the parameter T_{0_i} from each fit;
5. calculation of the correction offset t_{co} as the mean value of a normal fit over the distribution of the T_{0_i} value for all the run events;
6. correction of all timing offsets T_{0_i} with the t_{co} value;
7. iteration of the points 3 to 6 until the mean of the T_{0_i} distribution converges to a value arbitrarily close to zero.

For a given MIP track, the (measured of simulated) time of interaction values t_i is reported with respect to their coordinates (x, y) in the transverse calorimeter plane and are referenced to the time of the first hit (the quantities relative to the latter are marked with the subscript 0). The function used to fit the time-of-flight data for each event is calculated as follows:

$$t_i = T_{0_i} + \frac{|(x,y)_i - (x_0,y_0)|}{c \cdot sin\ \theta} + \frac{z_i}{\langle v_p \rangle} \tag{6.1}$$

In this formula, θ represents the angle of the MIP track with respect to the calorimeter longitudinal axis, $\langle v_p \rangle$ is the propagation velocity of optical photons inside the crystals, and z_i is the distance of the hit from the photodetector surface along the i-th crystal axis. Rearranging the equation and inserting the crystal length parameter $L = 20\ cm$ leads to the following relation:

$$t_i - \frac{|(x,y)_i - (x_0,y_0)|}{c \cdot sin\ \theta} = T_{0_i} + \frac{\Delta z_i}{v_p} + \frac{z_0 - L/2}{v_p} + \frac{L}{2v_p} \tag{6.2}$$

It should be noted that in the case of the online calibration, the calorimeter is unable to provide any information with regards to the axial coordinate z_i of the hit; regardless, the contribution of the second and third term contained in the right side of the (6.2), in which $\Delta z_i = \frac{|(x,y)_i - (x_0,y_0)|}{c \cdot sin\ \theta} cot\theta$ becomes negligible for large datasets, getting smeared over all possible impact points along the crystal axis and CR impact angles.

The procedure has been validated by applying uniformly distributed random offsets within the 1 ns interval to all calorimeter channels in the MC simulation, along with a 350 ps gaussian time smearing on signal. After five iterations, the residuals followed a normal distribution with a mean of approximately 0.008 ps and a standard deviation of 343 ps. Furthermore, as reported in Figure 6.5 (left), the distribution relative to the $t_{co} - T_{0_i}$ deviation is gaussian and centred around zero, with a sigma which lowers with every iteration. After the fifth one, the mean value from the normal fit is 3 ± 1 ps and the relative standard deviation is 347 ± 11 ps. As shown in Figure 6.5 (right), this procedure allows a time offset correction down to a level below 90 ps (i.e. an RMS spread of ~ 35 ps).

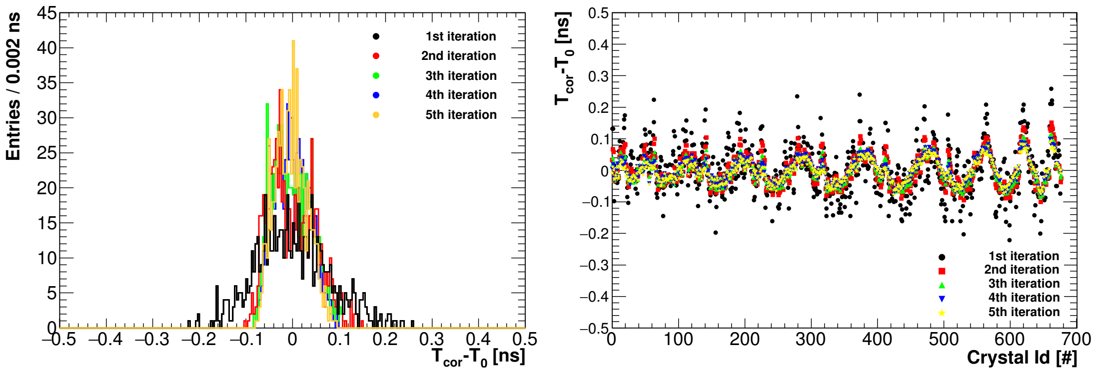

Figure 6.5 Left: single channel distribution example of the difference between the initial time T_{0_i} and the correction t_{co} found at each iteration. Right: same distribution as a function of the crystal number.

6.1.4 Calorimeter pre-calibration and CRT requirements

After the final assembly at Fermilab Silicon Detector facility (SiDet) and prior to the installation into the detector solenoid, each calorimeter disk will be individually tested. Performance validation and reference calibration of the detector will be carried out using MIPs traced by means of the Cosmic Ray Tagger (CRT) system described in the following sections. This system is needed to perform these procedures:

- Calibration of the absolute energy response over the expected 21 MeV MIP deposit;
- Equalisation of the channel responses below the 1% level;
- T_{0_i} alignment at a level below 30 ps (RMS);
- Crystals Longitudinal Response Uniformity (LRU) check at 21 MeV;
- Study of timing resolution dependence on crystals z-axis.

In order to perform these pre-calibration tasks in compliance with the requirements, a tagger capable of fast timing performance and high precision in the reconstruction of MIP tracks is needed. At the moment of writing, the instrument is under construction and characterisation at the LNF-INFN [56] in Italy. As previously stated, the Mu2e calorimeter is unable to provide any information with respect to the longitudinal crystal position of detected hits. The CRT system will instead allow a MIP 3D track reconstruction with excellent spatial resolution. Thus, apart from providing MIP trigger and track selection capabilities, the CRT system will allow to evaluate the LRU and the timing resolution performance of all channels with good precision, with respect to the specifications listed in Section 3.1.

6.2 CRT Design

6.2.1 Layout and mechanical design

The CRT system is composed by two sub-modules to be installed above and below the calorimeter disk under test, with the scintillator axes orthogonal to the crystal ones, in order to track the passage of MIPs throughout the entire calorimeter volume.

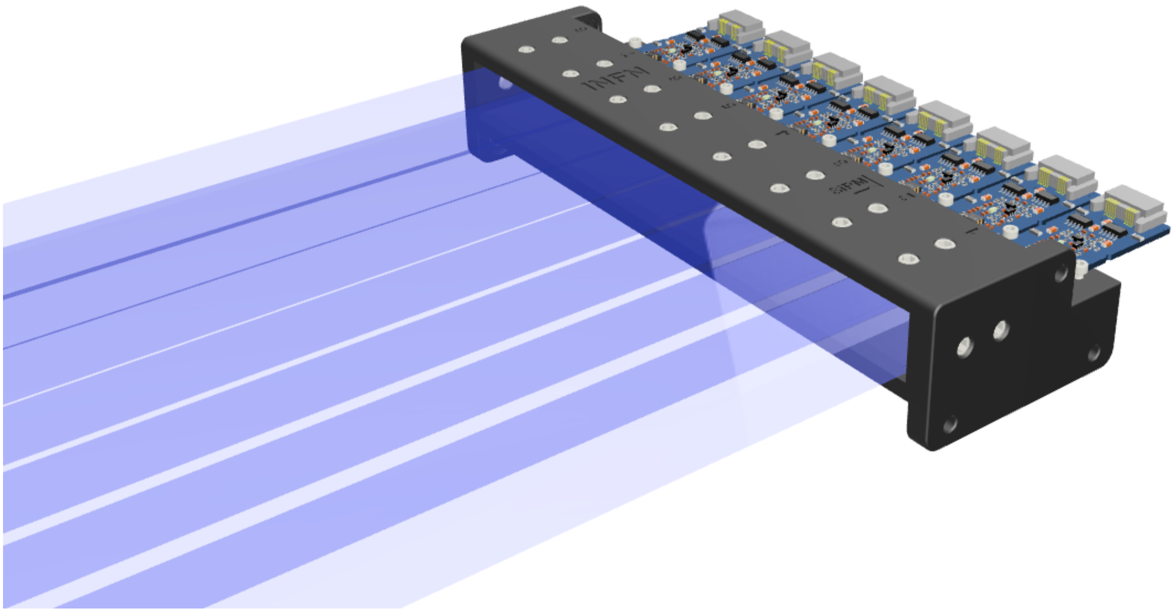

Figure 6.6 CAD rendering of a single module of the CRT system. The optical coupler block is shown for the readout system of single side.

Each sub-module is composed of a single layer of 8 parallel scintillator bars which have a 15×25 mm^2 square cross-section and measure 1600 mm in length. These elements have been properly sized in order to cover the whole crystal lengths and provide a 25 mm granularity with respect to the crystal longitudinal axis. Considering the two CRT plates installed above and below the calorimeter disk with a relative vertical separation of approximately 2500 mm, the 1600 mm length of the scintillators translates into an acceptance angle for cosmic muons on the calorimeter transverse plane of approximately $\pm$ 32 ° from the vertical.

The scintillating elements are individually wrapped with a reflective 150 µm Tyvek foil and closely packed together (with a relative separation of approximately 350 µm), after the interposition of a 50 µm Tedlar foil for the suppression of optical crosstalk. Each bar has a dual-side readout featuring a single Hamamatsu Mu2e SiPM plus its FEE board (cf. Chapter 4). Two integrated holders have been designed to be mounted at each end of the scintillator pack. The holder blocks (Figure 6.6) are additively manufactured via fused-deposition modelling on a Stratasys Fortus [48] in black acrylonitrile styrene acrylate (ASA) at LNF-INFN. The components embed recesses for the installation and alignment of the SiPM detectors, along with the mounting flanges for the FEE boards. Inside these holders, the scintillator pack is registered against a stopping and locating fixture which surrounds the detector recesses, leaving a 180 µm gap between the SiPM and scintillator faces in order to prevent mechanical coupling between the two components, so as to prevent excessive stresses on the SiPM windows and substrates due to dimensional or assembly tolerances.

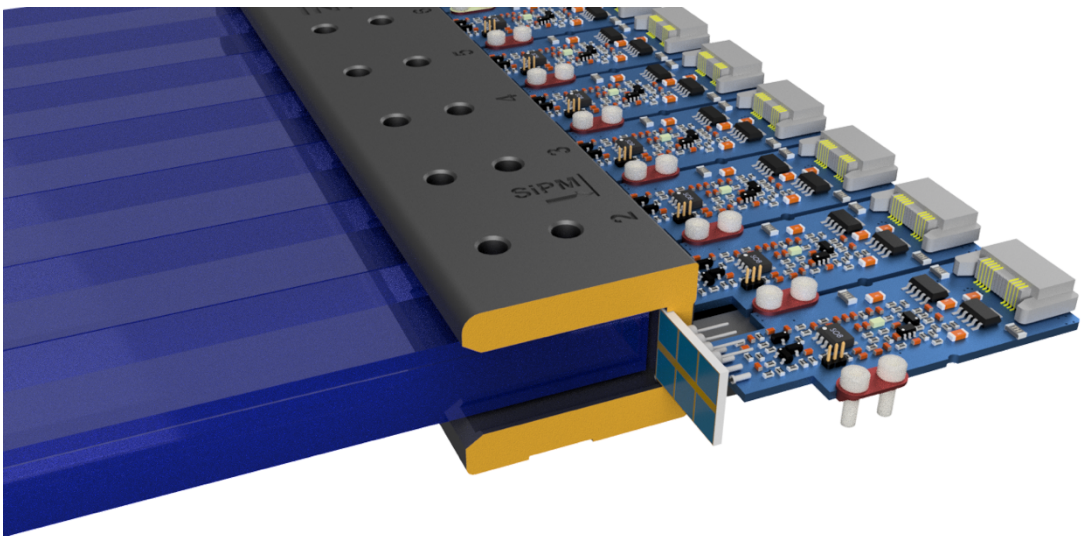

Figure 6.7 CAD rendering: detail of the SiPM coupling to the scintillator bars and of the front-end electronic boards.

The ratio between the bar cross-section and SiPM surface area is 0.58. The optical coupling is direct (Figure 6.7) and achieved through the use of optical grease. The integrated holder block also embeds lateral flanges for bolting to the chassis frame, which is realised in high precision aluminium alloy extrusions, to guarantee the parallelism of the assembly. Two additional flanges (also manufactured in ASA via fused deposition modelling) support the weight of the bars and provide additional flexural stiffness to ensure a small maximum deflection, in order to prevent inefficiencies in the light collection. The two completed sub-modules will be installed in light-tight PVC boxes.

6.2.2 Scintillating elements

Figure 6.8 reports the schematic representation of the cascaded process typical of organic scintillators. For high concentrations of the primary fluor (> 1% w/w), the latter and an excited base unit exchange energy predominantly via the Förster resonance energy transfer mechanism, due to their short (~ 100 Å) intermolecular distance: this strong coupling increases the light yield and shortens the decay time of the base material [52]. A secondary fluor is used as a wavelength shifter in order to obtain the desired emission spectrum and attenuation length (a large Stokes' shift is desirable in scintillators in order to reduce self-absorption).

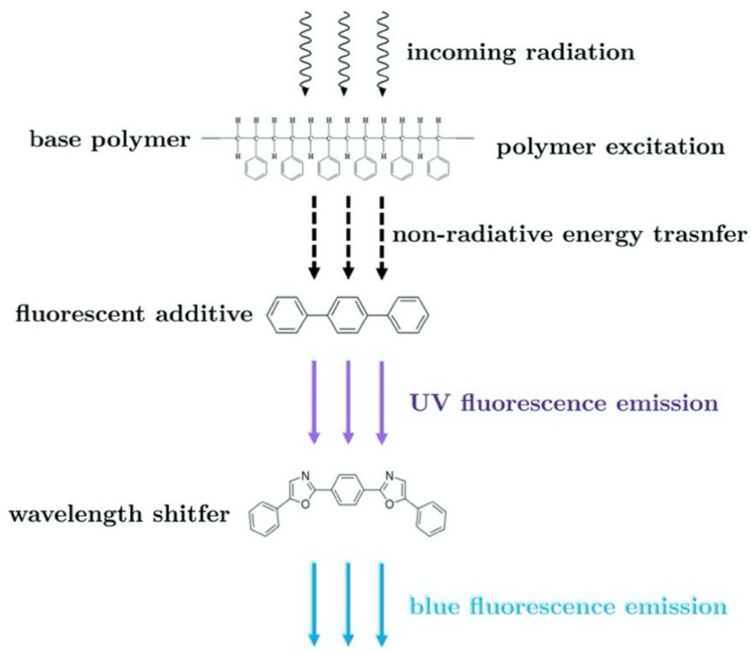

Figure 6.8 Schematic representation of the two-step cascaded scintillation mechanism of plastic scintillators with a primary and secondary fluor.

The EJ-200 polyvinyltoluene based plastic scintillator material (Eljen Technology) has been chosen for its emission and optical properties [49]. In fact, this material has a very good scintillation efficiency of 10^3 photons per deposited

MeV and an emission spectrum peaking at 425 nm, which conforms well to the spectral detection efficiency of the Mu2e SiPMs. The long bulk attenuation length of 3800 mm is well suited for the 1600 mm bar longitudinal dimension. Furthermore, in order to achieve the desired precision in timing reconstruction, the EJ-200 has been chosen for its fast rise time of 900 ps and decay time of 2.1 ns, with a slow-to-fast component intensity ratio of 0.27 (Figure 6.9).

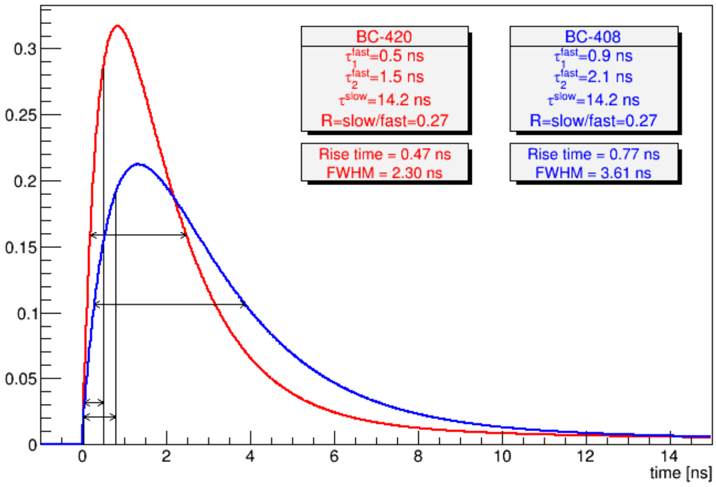

Figure 6.9 Emission time characteristics for Eljen EJ-200 (or Saint-Gobain BC-408, in red) and Eljen EJ-240 (or Saint-Gobain BC-420, in blue) plastic scintillators.

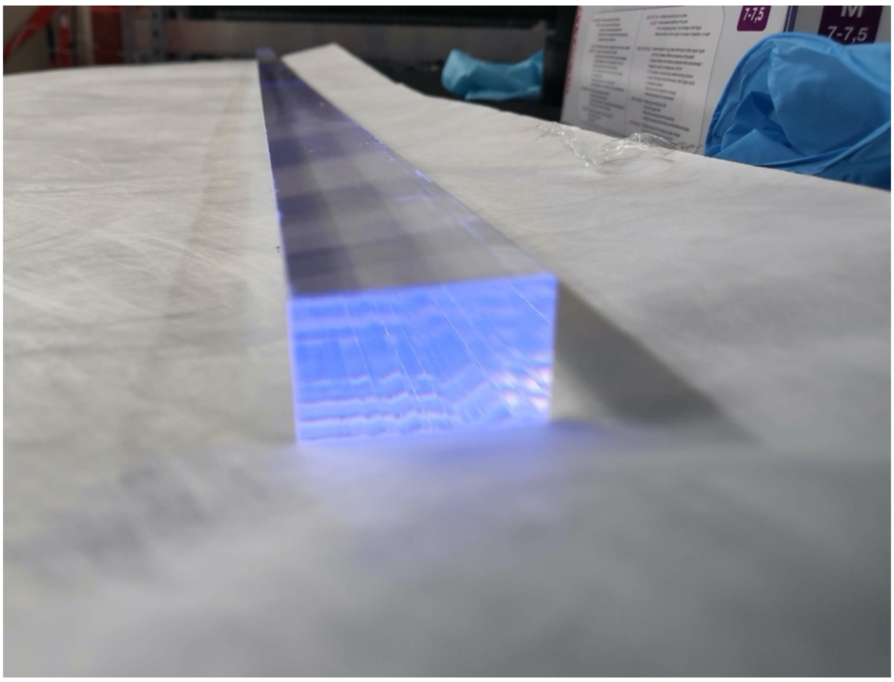

Figure 6.10 One of the EJ-200 scintillating elements about to receive its Tyvek reflective wrapping at LNF-INFN.

The $\langle dE/dx \rangle$ value for the EJ-200 base material (PVT) is around 2 MeV/cm (which corresponds to a minimum ionisation of 1.956 $\frac{\mathrm{MeV \cdot cm^2}}{\mathrm{g}}$). Considering a 15 mm path length and the 10^3 photons/MeV light output, a MIP deposit of 3 MeV should result in the generation of 30000 photons.

6.3 CRT characterisation

In order to perform the characterisation of a single scintillator bar, two SiPM optical couplers (Figure 6.11) were designed and additively manufactured in stereolithography (Formlabs Form 2 [50]) and installed on the bar under test, which was previously wrapped in Tyvek and in an outer darkening cladding.

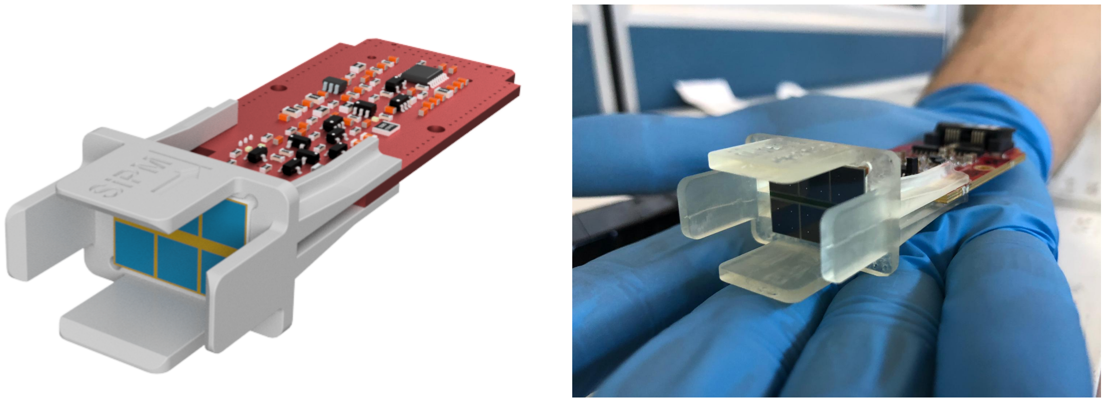

Figure 6.11 CAD rendering (left) and actual realisation (right) of the SiPM test coupler.

A double readout scheme was implemented using the standard Mu2e FEE with a gain setting of 2. A CAEN V1720 digitiser [51] was used to sample at 250 Msps both the SiPMs pulses from the EJ-200 and the signals from two small ($1 \times 1 \times 10$ cm^3) scintillating counters read by PMTs and used as a triggers.

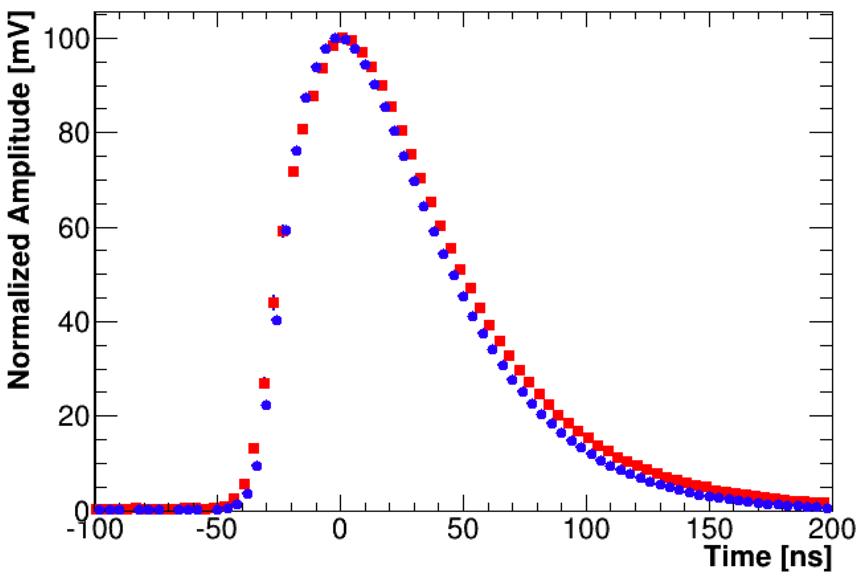

Figure 6.12 Comparison of a picosecond UV laser pulse fired directly at the SiPM (blue) and MIP (red) waveform on EJ-200 (after amplitude equalisation).

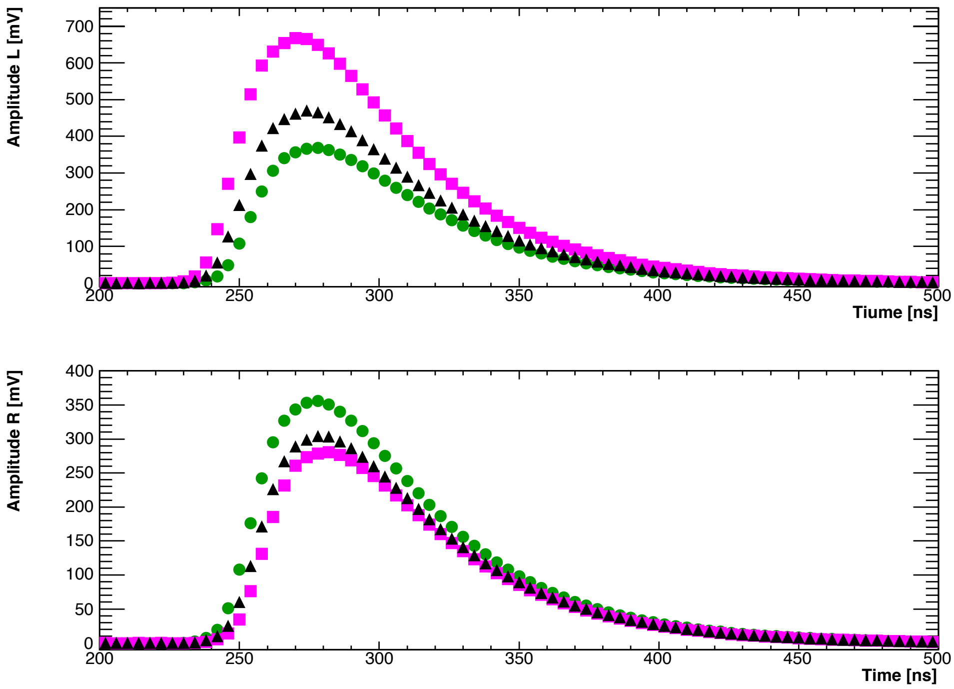

Figure 6.13 Sample waveforms for the left (top) and right SiPM (bottom) collected for three different positions of the trigger counters (relative to the left SiPM): 10 mm (purple), 400 mm (black) and 800 mm (green).

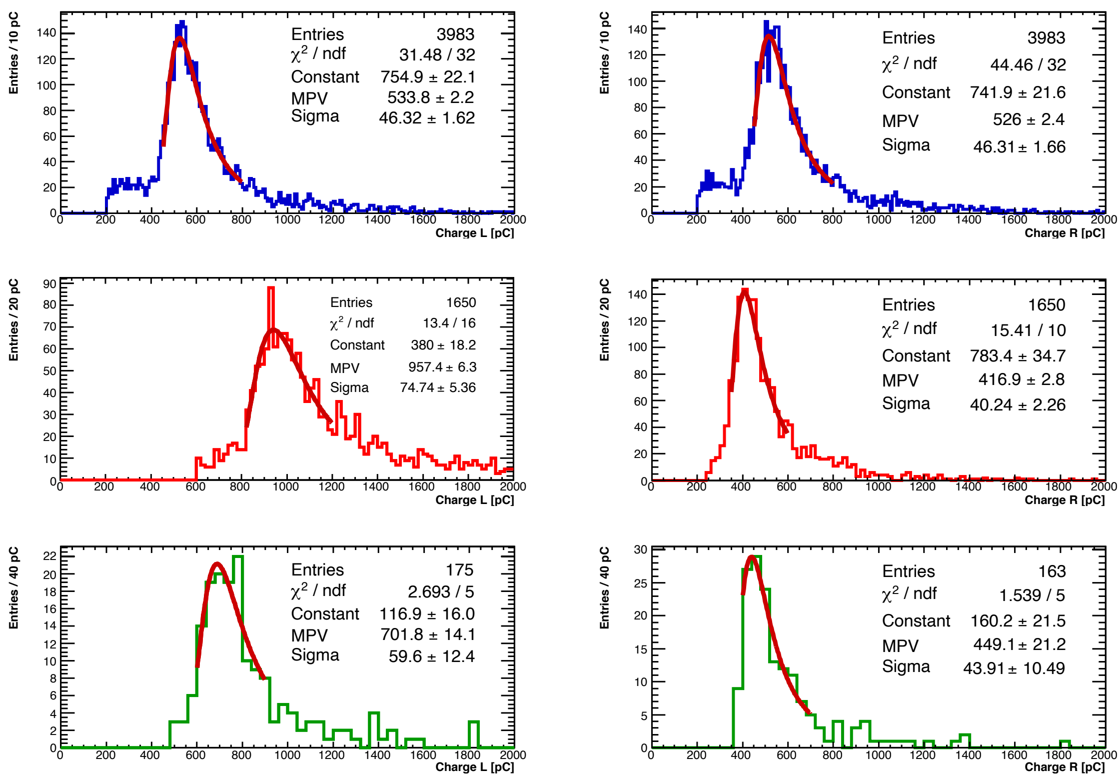

Figure 6.14 Charge spectra for the left and right SiPM collected for three different positions of the trigger counters (relative to the left SiPM): 10 mm (red), 400 mm (green) and 800 mm (blue). MIP peaks are clearly visible. A Landau fit is superimposed.

6.3.1 External trigger run

A first characterisation run was performed by triggering the acquisition on the coincidence of the discriminated signals relative to the two counters, placed parallelly above and below the bar under test, in order to have a precise spatial reference of the interaction point. The position of the trigger scintillators was scanned across the scintillator axis in 10 steps. Runs of 6000 MIP events each were carried out with an approximate trigger rate of 40 mHz. Sample waveforms are shown in Figure 6.13, relatively to the right and left SiPM for three different trigger positions. Figure 6.12 shows instead a MIP pulse in comparison with a waveform collected using a picosecond UV laser source pulse on the EJ-200.

The collected waveforms were integrated over a 250 ns gate into the charge spectra reported in Figure 6.14 for 3 different scan positions, along with a superimposed Landau fit. MIPs peaks are clearly visible, and a Landau fit is superimposed. The signals and charge spectra recorded with the trigger system in the middle of the bar show a good amplitude agreement between the two channels, confirming a good equalisation of the readout systems and of the optical coupling.

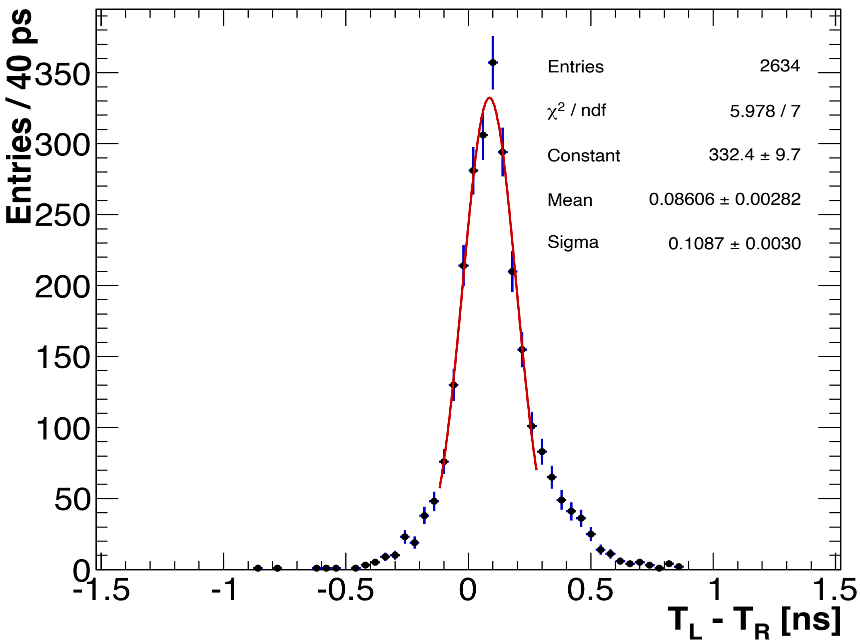

Figure 6.15 Time-of-flight difference distribution for the trigger counters positioned in the middle point of the EJ-200 bar.

6.3.2 Light propagation and time resolution study

A preliminary timing analysis was performed by fitting the MIP pulse leading edges with a log-normal function and applying a 20% constant fraction discrimination. For the (1 cm wide) trigger counters positioned in the middle of the scintillator, the time-of-flight (TOF) differences relative to the two readout channels (T_L-T_R) are distributed normally, as shown in Figure 6.15, with a standard deviation of ~100 ps.

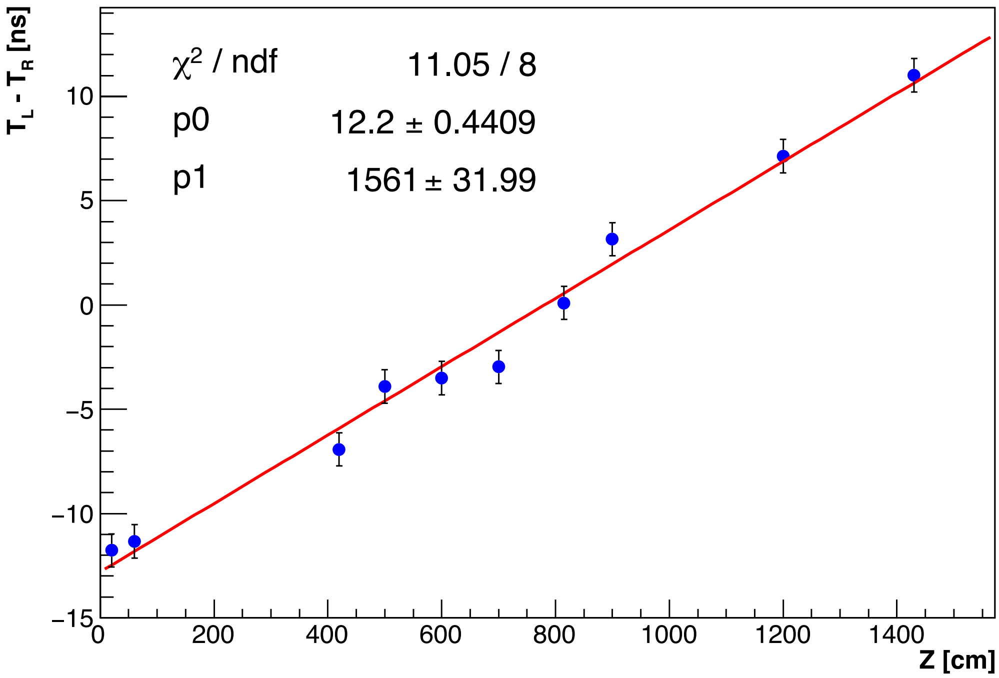

Figure 6.16 Time-of-flight difference T_L-T_R versus the scan position of the trigger counters. The average propagation velocity and angle are extrapolated using a linear regression.

For different positions of the trigger system along the scintillator axis, the mean values relative to distributions of the time differences T_L-T_R were evaluated using a normal fit (Figure 6.16). The obtained times as a function of the trigger position were fitted with the following relationship, in which z represent the scan position of the triggers. The variables L and $\langle v_p \rangle$ are the fit parameters and represent respectively the scintillator bar length and the average light propagation speed:

$$T_L - T_R = 2 \cdot z/\langle v_p \rangle - L/\langle v_p \rangle \tag{6.3}$$

This fit procedure reconstructed the scintillator length with a 2.5 % deviation and evaluated an average propagation speed of ~12.2 cm/ns for optical photons. From this value, an average light propagation angle $\langle\theta_p\rangle$ inside the bar of approximately 50° was evaluated using the following relation, where n represents the refractive index (which is equal to 1.58 for the EJ-200):

$$\langle\theta_p\rangle = \left(\tfrac{n\cdot\langle v_p\rangle}{c}\right) \approx 49.9° \tag{6.4}$$

A preliminary estimation of the timing performance can be obtained from these data. The time resolution is determined to be $\sigma/\sqrt{2} \approx 70$ ps. The spatial longitudinal resolution is estimated to be $\langle v_p\rangle \cdot \sigma/2 \approx 6$ mm (i.e. 0.4 % of the scintillator length).

6.3.3 Attenuation lengths

The main phenomena responsible for light attenuation are represented by self-absorption and Rayleigh scattering [53]. These effects are cumulatively taken into account in the effective attenuation length $\lambda_{BULK} = (1/\lambda_{abs} + 1/\lambda_{scatt})^{-1}$, which represents the number of downward e-foldings of the initial light intensity per unit length of material. The attenuation behaviour is wavelength dependent.

In plastic scintillators with reflective (or diffusive) coatings, light emitted with skew angles suffers from a stronger absorption due to the longer path lengths and to the greater number of reflection losses, if compared to meridionally emitted light (Figure 6.17). As a result, the transmission function should be considered dependant on the light emission axial angle. In practice, the light propagation inside the scintillator can conveniently be approximated with the following double exponential model (directly reported in terms of to the reconstructed charge):

$$Q(z) = a \cdot (e^{-z/\lambda_{BAL}} + \tfrac{I_r}{I_d} e^{-z/\lambda_{TAL}}) \tag{6.4}$$

The λ_{BAL} term is the bulk attenuation length associated with the bulk optical properties of the medium (specifications indicate a nominal value of 3800 mm for the EJ-200). This parameter is associated with the portion of light which propagates directly or with a single reflection from the interaction point to the photodetector.

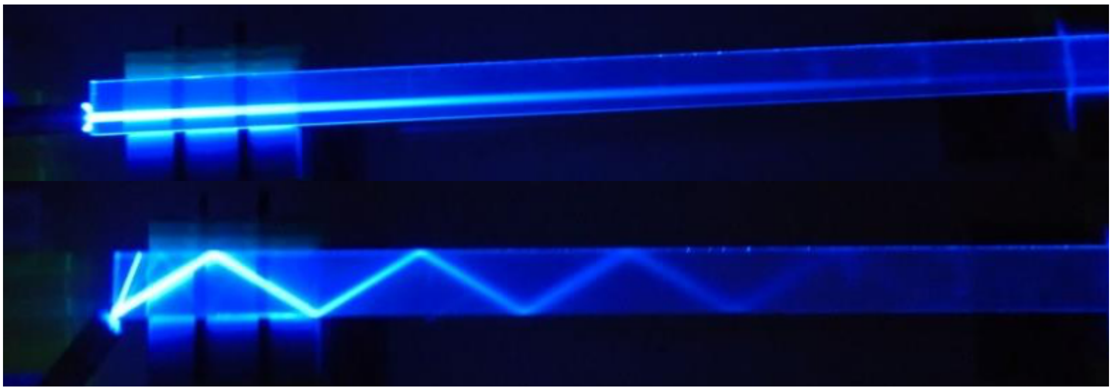

Figure 6.17 Direct and reflected laser light propagating inside a scintillator bar.

The technical attenuation length λ_{TAL} is instead significantly shorter and accounts for the light travelling indirectly to the sensor and which is subject to multiple reflections or Lambertian diffusion at the scintillator-reflective wrapping interface. The λ_{TAL} also depends on the optical configuration of the system and on the experimental layout in general. The $\frac{I_r}{I_d}$ factor represents the ratio between the relative intensities of these two components (i.e. the reflected to direct light intensity ratio). The intensities which appear in this ratio can be considered constant in the limited inner region of the scintillator bar, but they actually show a dependence on the hit position along the scintillator axis. In fact, non-linearities in the light collection - associated with the different solid angles seen by the detector - will appear at the extremal positions, where the hits happen very close to the SiPM surface. Finally, the term a is a normalisation factor.

Considering an m number of internal reflections for the indirect portion of the light, what follows can be considered as quick way of estimating the λ_{TAL} value:

$$\frac{h}{\lambda_{BAL}} + sin\langle\theta_p\rangle \cdot log(1/R) = \frac{cos\langle\theta_p\rangle \cdot h}{\lambda_{TAL}} \tag{6.6}$$

where $d \cdot cos\langle\theta_p\rangle = z$ and $d \cdot sin\langle\theta_p\rangle = m \cdot h$, being d the average light path between successive reflections and h the cross-sectional dimension of the scintillating element. For the EJ-200 bars, the latter has been approximated as $h = 2\frac{h_1 \cdot h_2}{h_1 + h_2}$, where h_1 and h_2 define the bar rectangular cross section. A reflection efficiency $R \approx 0.9 \div 0.96$ is reported for Tyvek wrapped scintillators for wavelengths close to the EJ-200 spectral peak. Considering the previously calculated average propagation speed and angle, a TAL of approximately 30 cm is expected.

6.3.4 Noise pedestal

The system noise has been evaluated as the RMS value of the reconstructed charge pedestal on noise runs acquired with a random trigger generator. All other running conditions (such as biases, readout and cabling) were left unchanged with respect to the triggered data sample, so as not to alter the load on the FEE. The noise pedestal as charge is reported in Figure 6.18. The charge noise values expressed in percentage with respect to the MIP peak is below 2.5% for both channels. This in turn translates to an energy value of approximately 60 keV.

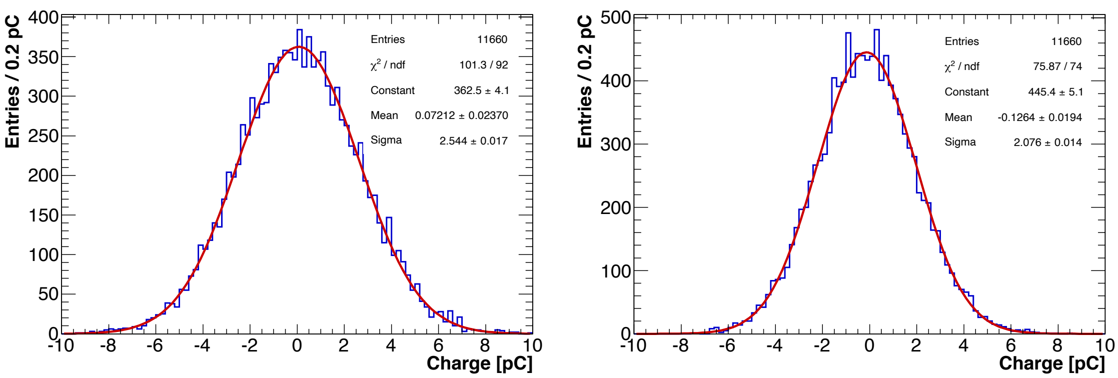

Figure 6.18 Noise pedestal for left and right channels expressed in charge.

6.3.5 Template fit, timing and calibration

Timing studies were performed by executing acquisition runs of 60000 events each based on the coincidence of the two readout SiPM. The left readout signal discriminator was adjusted in order to have a long output pulse, and a 32 ns transmission line delay was inserted after the right channel discriminator. The delay was chosen to be larger than $L/\langle v_p \rangle$, so that the right channel would always lag in the coincidence unit. The timing analysis has been performed using a template fit procedure [53] on the CRT data, in analogy to what is already described in section 3.6.

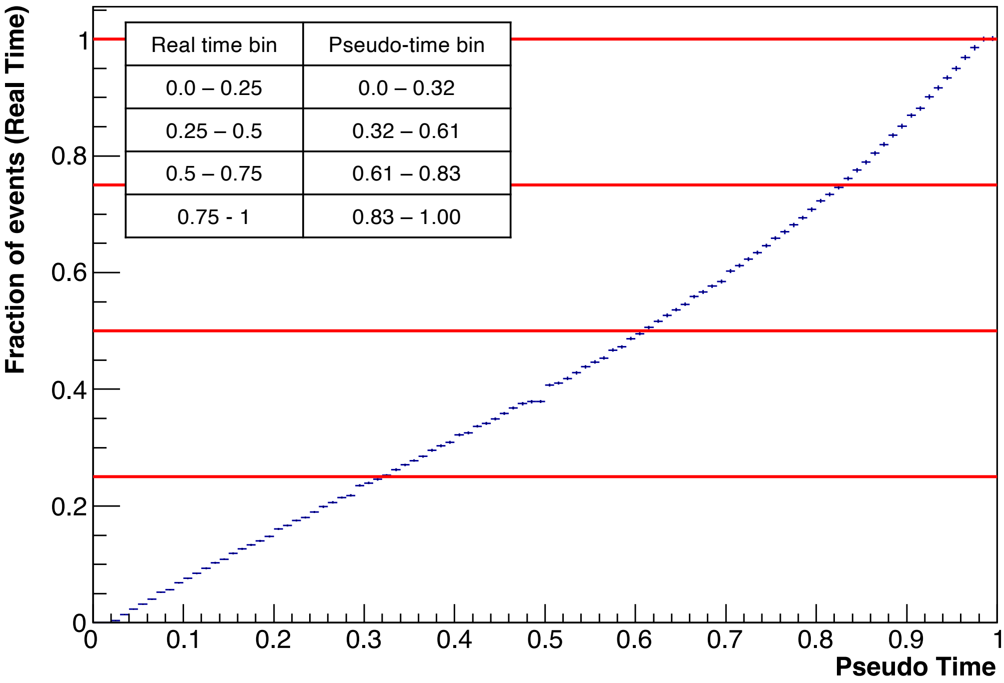

Figure 6.19 Mapping of the pseudo-time and real time cumulative distribution.

6.3.5.1 **Template construction and fit**

The template was constructed using cosmic ray data from a single channel (the left one), sampled at 250 Msps. At first, the pseudo-time variable t_p was evaluated as

$$t_p = \frac{2}{\pi} arctan\left(\frac{s_0 - s_-}{s_0 - s_+}\right) \tag{6.7}$$

The parameter s_0 represents the signal value associated to the bin containing the sampled waveform peak, while s_- and s_+ represent the signals contained in the adjacent bins. Although not necessarily linearly correlated to the actual peak position, the pseudo-time t_p should vary smoothly from 0 to 1 along with the cumulative distribution of the real times t_r, given the uniform distribution of MIP hits in time. The real time t_r (which also varies from 0 to 1) represents indeed the actual unknown position of the peak, relative to the bin in which it is contained. As a result, given an arbitrary value t_p for the pseudo-time, the events can be mapped with respect to t_p so that to any arbitrary fraction of events F having $0 < t_p < t_{p_0}$ contains waveforms which must have a real time between 0 and $t_{r_0} = F$. This translates into a one-to-one correspondence between t_p and t_r. The results of this procedure are shown in Figure 6.19 which reports the pseudo-time cumulative distribution.

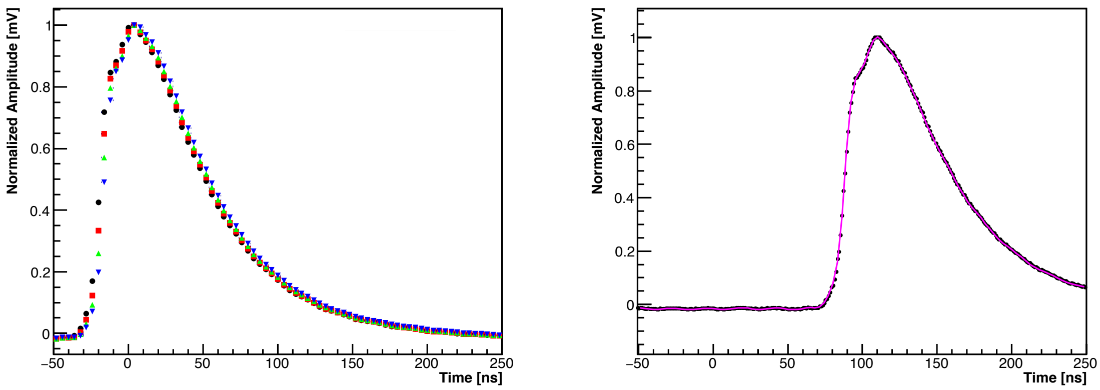

Figure 6.20 Left: template construction using the four averaged waveforms. Right: template constructed using the TSpline3 ROOT class.

The obtained cumulative distribution was subsequently split into 4 equal bins of the real time variable, and events were split into 4 sub-sets accordingly. The ensuing t_p and t_r fractions are thus mapped as reported in the same figure. Events in each subset were successively summed and normalised over the number of entries. The four obtained averaged waveforms were progressively shifted of ¼ of a bin each (i.e. 1 ns) and combined together (Figure 6.20, left).

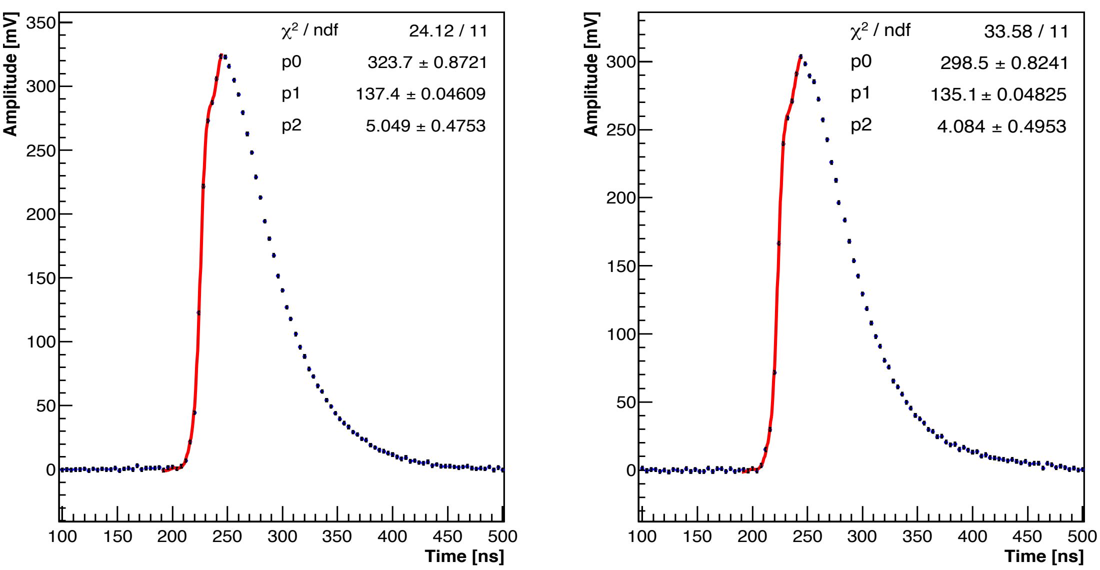

Figure 6.21 Examples of template fits for the waveforms for the left and right channel.

The template shown in Figure 6.20 (right) was finally generated using a TSpline3 ROOT object [54] over the combined 1 ns resampled waveforms. Templates are then fitted over signals using an optimisation of only 3 parameters: peak position (p_0), pedestal (p_1) and a normalisation factor (p_2). Two examples of the fit procedure carried out on the rising edge of a MIP pulse are reported in Figure 6.21 for both SiPMs. It should be noted the presence of a small dimple in the template shape: this feature is due to the response of the Mu2e FEE shaper - which is not optimised for the fast rising times of the EJ-200 - and correctly reproduces the actual signal shape, without any adverse effects on the goodness of the fitting procedure.

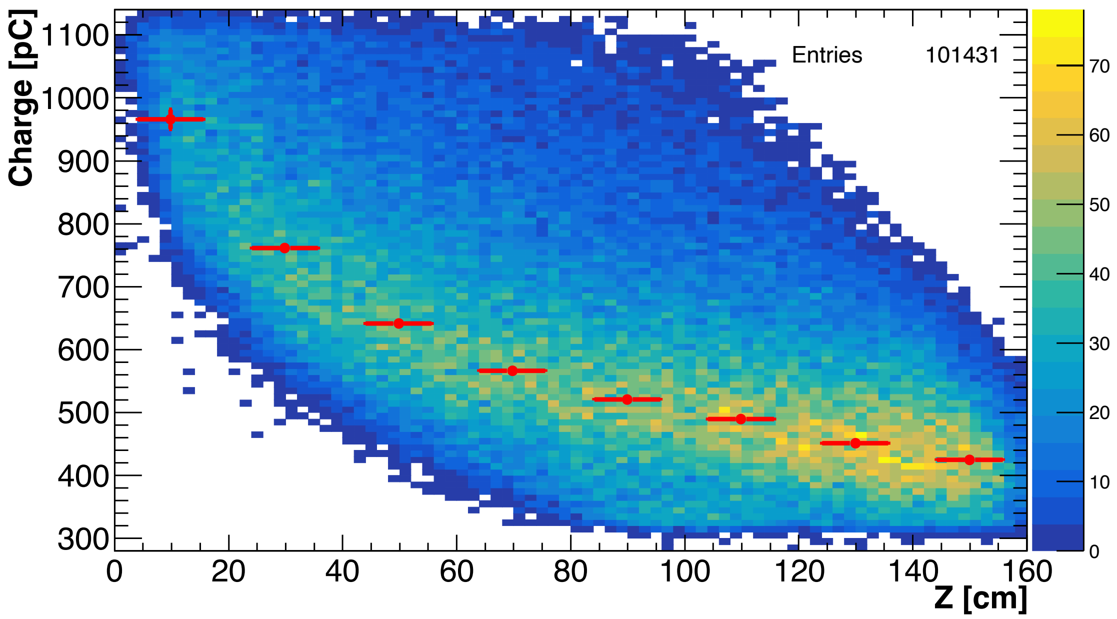

Figure 6.22 Reconstructed charge versus the hit position along the scintillator axis (left channel) with the superimposed fitted charge profile.

6.3.5.2 **Template fit, TOF and light propagation**

The events acquired by triggering on the coincidence of the two readout SiPMs are plotted in Figure 6.22, after the application of the proper cuts on signal amplitude, reconstructed charge and on the χ^2 of the fitted waveforms. This 2D distribution reports the reconstructed charge versus the hit position along the scintillator axis for the left channel. The template fit procedure described in the previous section was applied to reconstruct the time-of-arrival difference, and the interaction point was evaluated using the (6.3), along with the light propagation parameters derived in section 6.3.2. The distribution of the reconstructed hits was divided into 8 equal 20 cm slices with respect to the Z coordinate. The charge distributions relative to the events contained in each slice were fitted with a Landau function and the resulting MPVs were assembled into the charge profile which is overlaid in Figure 6.22. An example of the Landau fits carried out on four Z slices are reported in Figure 6.23, for the left channel. The errors relative to the charge are the errors resulting from the fit procedure, while the quantity $L/8\sqrt{12}$ was used as the error relative to the Z position, assuming a uniform distribution of hits on each slice of length L/8 = 20 cm.

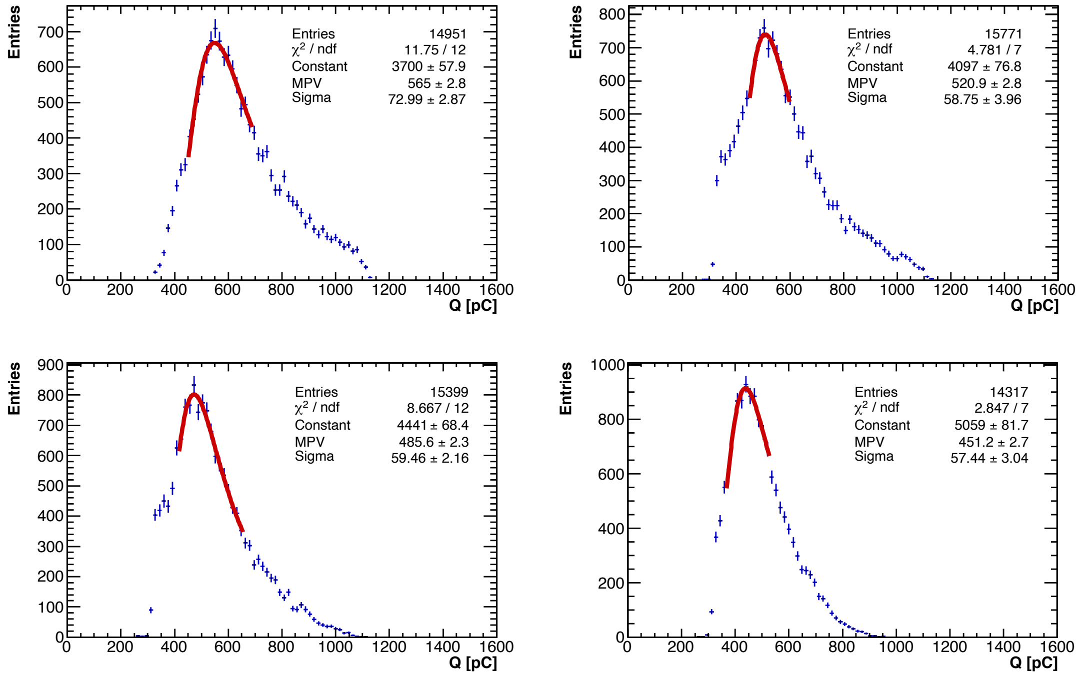

Figure 6.23 Examples of the Landau fits applied on the reconstructed charges relative to 4 of the 8 Z slices (left channel).

The obtained profiles were then fitted using the double exponential light propagation model for both channels, as shown in Figure 6.24 and in analogy to what is reported in section 6.3.3: the BAL was fixed to 380 cm and the parameters shown in Figure 6.24 are the normalisation factor (p_0), the technical attenuation length (p_1) and the indirect to direct light intensity ratio (p_2). The values are in good agreement between the two channels and in line with the estimations. Prior to this analysis, an equalisation factor of 1.103 was applied to the charge values of the left channel. This calibration factor was evaluated by comparing the charge response of the two SiPMs to MIPs events selected using a trigger scintillator positioned in the middle of the bar under test (cf. section

6.3.2). This simple correction allows to account for and equalise any unbalancing in the optical coupling of the SiPMs and in the response of the two readout systems.

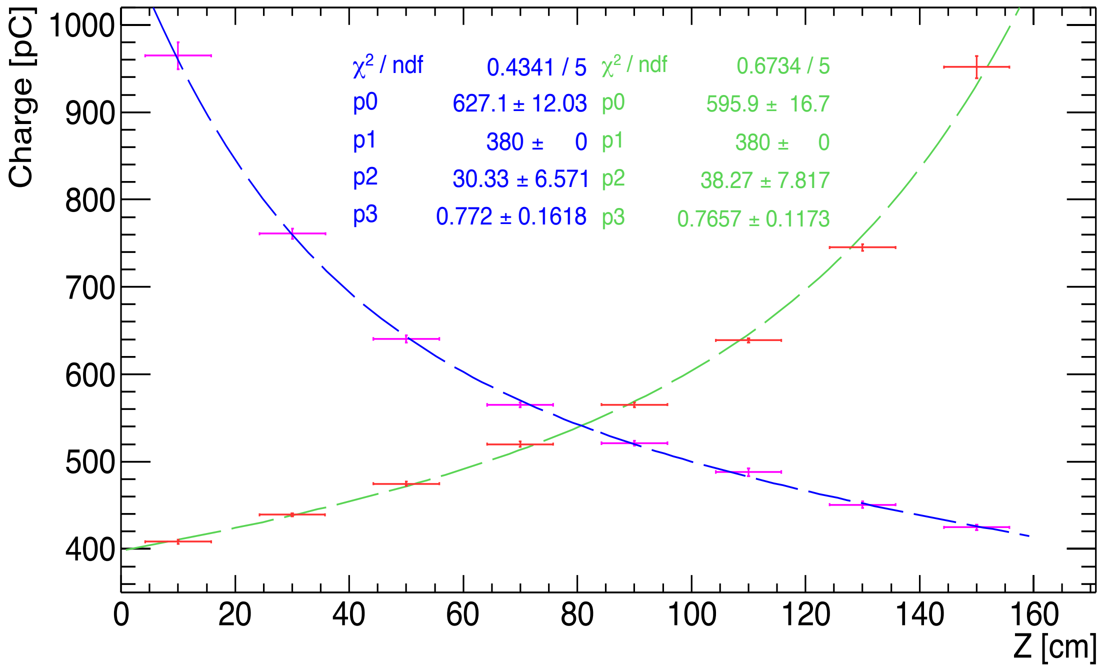

Figure 6.24 Charge profile of the Left and Right readout channels as a function of the reconstructed Z position for the CR data run. A double exponential light transmission model is fitted over the charge profiles for both channels.

The distribution along the scintillator axis of the interaction points collected in a single run is reported Figure 6.25. As expected, the spatial distribution of the MIP events is quite uniform and well reproduces the scintillator profile, apart from the extremal sections which suffer from leakage effects. Using a double sigmoid fit over the profile, the scintillator length can be correctly reconstructed. The function used for the fit procedure is the following one:

$$y = \frac{p_3}{1 + e^{-\frac{z-p_0}{p_1}}} - \frac{p_3}{1 + e^{-\frac{z-p_2}{p_1}}} \tag{6.7}$$

Where y represents the number of hits recorded as a function of the reconstructed position z along the scintillator axis.

For verification purposes, the position z_t of a 1 cm wide reference scintillation counter placed at a known position over the bar axis was correctly extrapolated as $2 \cdot z_t = \langle v_p \rangle \cdot (T_L - T_R - T_d) - L$, where T_d represents the 32 ns transmission line delay.

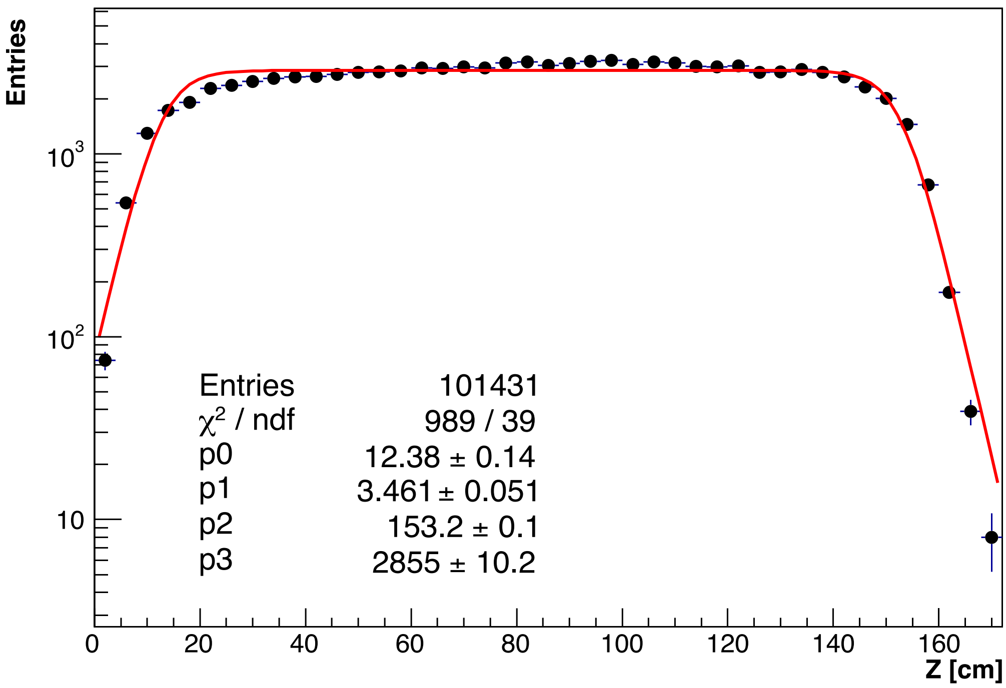

Figure 6.25 Distribution of the number of hits recorded as a function of the reconstructed Z position for the CR data run. A double sigmoid fit is superimposed on the scintillator profile.

Conclusions

The final design of the Mu2e crystal calorimeter ensued from a long R&D phase. The design was validated via test beams and cosmic ray studies carried out on the Module-0 large scale prototype. At the moment of writing the production phase of all SiPMs and crystal is concluded, and the calorimeter disks and support structures have been milled and are ready for assembly. The final assembly of the Mu2e calorimeter is scheduled in 2021.

The final assembly of the 2696 SiPMs on the respective holders have started in the LNF cleanroom. The design of the FEE has been finalised with the latest performance test of the linear regulator and preamplifier, with focus on the final irradiation tests carried on the electronics using an intense gamma source. At LNF, the final verification steps of the complete front-end system have been carried out, including a comprehensive characterisation of the optical and electronic performance of the hardware. At the moment of writing the production phase of the FEE boards has just started.

Recent simulations have been performed in order to refine the estimations of the RIN impact on the calorimeter. Moreover, comprehensive study of the digitisation and trigger systems are in progress and have reached a mature stage, along with the particle identification algorithms. A template fit algorithm has been developed in order to allow fast high resolution time reconstruction on the Mu2e crystals and on the CRT scintillators.

The work relative to the design and characterisation stages of the Cosmic Ray Tagger has successfully started and reached a mature stage. The complete characterisation of the scintillating elements and of the interaction point/time reconstruction algorithms have been successfully carried out, thus validating the design, performance and resolution of the instrument. The final assembly of the CRT is due by the end of 2020.

The construction of the Mu2e facility in the FNAL Muon Campus has been completed, and the experimental hall is ready to house the Mu2e instrumentation. The beginning of the data-taking period is currently scheduled in 2023.

References

[1] L. Bartoszek et al., *Mu2e Technical Design Report* [arXiv:1501.05241]

[2] G. Altarelli, *The Standard Model of Particle Physics,* 2005 [arXiv:hep-ph/0510281]

[3] L. Calibbi and G. Signorelli, *Charged Lepton Flavour Violation: an Experimental and Theoretical Introduction*, Riv. Nuovo Cim. 41 (2018) no.2, 1 (2018-02) [DOI:10.1393/ncr/i2018-10144-0, arXiv:1709.00294]

[4] Super-Kamiokande Collaboration, Y. Fukuda et al, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. 81, 1562 (1998) [arXiv:hep-ex/9807003]

[5] J.N. Bahcall, M.C. Gonzalez-Garcia and C. Peña-Garay, *Global analysis of solar neutrino oscillations including SNO CC measurement*, Journal of High Energy Physics, Volume 2001, JHEP08

[6] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elemen- tary particles*,Prog. Theor. Phys. 28 (1962) 870

[7] Marciano W. J., Mori T. and Roney J. M., *Charged Lepton Flavor Violation Experiments*, Annual Review of Nuclear and Particle Science 58, 315-341 (2008)

[8] https://meg.web.psi.ch

[9] SINDRUM II Collaboration (Bertl, Wilhelm H. et al.), *A Search for muon to electron conversion in muonic gold*, Eur. Phys. J. C47 337-346 - 2006.

[10] Physics Reports 354 (2001) 243–409

[11] V. Cirigliano, *et al.*, PRD 80, 013002 (2009)

[12] A. Czarnecki, M. Dowling et al., *Michel decay spectrum for a muon bound to a nucleus*, Phys. Rev. D 90, 093002 (arXiv:1406.3575) - 2014.

[13] A. Czarnecki, X. G. Tormo and W. J. Marciano, *Muon decay in orbit: spectrum of high-energy electrons*, Phys. Rev. D84 013006 - 2011

[14] P. Derwent, *Accelerators for Intensity Frontier Research*, Conf.Proc. C1205201 (2012) 4185-4189

[15] https://novaexperiment.fnal.gov

[16] FERMILAB-CONF-18-660-E

[17] O. Atanova et al., *Measurement of the energy and time resolution of a undoped CsI + MPPC array for the Mu2e experiment*, JINST 12 (2017) no.05, P05007 [DOI:10.1088/1748-0221/12/05/P05007, arXiv:1702.03720]

[18] https://www.hamamatsu.com/resources/pdf/etd/R1828-01_R2059_TPMH1259E.pdf

[19] G. Pezzullo et al., *Progress status for the Mu2e calorimeter system*, Journal of Physics: Conference Series 587, 012047 (2015)

[20] https://geant4.web.cern.ch

[21] D. Brown, *Tracker Selection Cut-set C* Mu2e-doc-3996-v3

[22] https://www.caen.it/products/v1742/

[23] http://w3.lnf.infn.it/acceleratori/btf/

[24] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments,* Springer, 1994

[25] G. Pezzullo et al., *Status report on dose estimates and future plan*, Mu2e-doc- 18417 - 2018

[26] ASTM-722–14 (2014)

[27] C. Leroy et Al., *Particle interaction and displacement damage in silicon devices operated in radiation environments*, Reports on Progress in Physics 70, (2007) 493

[28] R.Y. Zhu, *Radiation damage in scintillating crystals*, NIM A Vol. 413 (1998) Issues 2-3 [DOI:10.1016/S0168-9002(98)00498-7]

[29] Z.Y. Wei and R.Y. Zhu, *A study on undoped CsI crystals*, NIM A 326 (1993) 508

[30] S. Baccare et al., *Radiation hardness test of un-doped CsI crystals and Silicon Photomultipliers for the Mu2e calorimeter*, CALOR 2016, J.Phys Conf. Ser. 928, conference 1, (2016) 928 012041 [arxiv:1606.07255, DOI: 10.1088/1742- 6596/928/1/012041]

[31] Mu2e-DocDB 20131

[32] R. Garcia et al., *Introduction to radiation effects and modelling of radiation damage*, Reliability and Availability Workshop for the High Luminosity LHC, CERN (2014)

[33] J. R. Srour et al., *Review of displacement damage effects in silicon devices*, Doi: 10.1109/TNS.2003.813197

[34] J. Schwank, *Radiation Effects in MOS Oxides*, Sandia National Laboratories, SAND2007-3309C

[35] Mu2e DocDB 29910-v1

[36] https://www.hzdr.de/

[37] arXiv: 1710.06239

[38] arXiv: 1804.09792

[39] F. Acerbi at al., *Understanding and simulating SiPMs*, Nucl. Instrum. Methods Phys. Res., *A* 926 (2019) 16-35

[40] W. Maes et al., *Impact Ionization in Silicon: A Review and Update*, Solid-State Electronics, Vol.33, No.6:705–718, 1990

[41] F. Corsi et al., NIM A572 (2007) 416

[42] S. Seifert et al., IEEE TNS 56 (2009) 3726

[43] Shockley, Sol. State Ele. 2 (1961) 35

[44] H. Otono et al., *Study of the internal mechanisms of Pixelized Photon Detectors operated in Geiger-mode*, arXiv:0808.2541

[45] https://www.caen.it/products/v1742/

[46] M. Guan et al., *A parametrization of the cosmic-ray muon flux at sea-level,* Chinese Physics C (2015) [arXiv:1509.06176]

[47] P. Shukla and S. Sankrith, *Energy and angular distributions of atmospheric muons at the Earth* (2016) [arXiv:1606.06907]

[48] https://www.stratasys.com/

[49] https://eljentechnology.com/products/plastic-scintillators/

[50] https://formlabs.com/

[51] https://www.caen.it

[52] J. B. Birks, *The theory and practice of scintillation counting* (Pergamon Press, Oxford, 1964)

[53] C. P. Achenbach, *Active optical fibres in modern particle physics experiments* (2004), Phys. Rev. D 73, 072003

[54] https://root.cern.ch/doc/master/classTSpline3.html

[55] https://root.cern.ch/doc/master/classTProfile.html

[56] http://w3.lnf.infn.it

[57] https://www.hamamatsu.com