UNIVERSITÀ DI PISA

Dipartimento di Fisica “Enrico Fermi”
Graduate Course in Physics

Ph.D. Thesis
XXXII Cycle

The new trigger and data acquisition system for LFV searches in the MEG II experiment

Candidate:
Marco Francesconi

Supervisors:
Dr. **Alessandro Baldini**
Prof. **Donato Nicolò**

April, 2020

Contents

List of Figures

List of Tables

Abstract

This PhD. thesis has been developed in the frame of the MEG II experiment, which is currently in commissioning phase at Paul Scherrer Institut (CH), with the goal of probing charge lepton flavour violation (CLFV) to an unprecedented level. The experiment was developed on the basis of the knowledge coming from operating the former MEG experiment[1]; it aims principally at exploring the muon decay $\mu^+ \rightarrow \mathrm{e}^+ + \gamma$ with an order of magnitude increase in sensitivity, down to branching ratios of $\mathrm{BR} < 6 \cdot 10^{-14}$ at 90% confidence level[2]. The compelling theoretical motivations for such an effort is the extreme sensitivity of CLFV channels to physics beyond Standard Model, while having a negligible background. Requiring very performant detectors, MEG II can also probe other exotic channels, such as the search of a possible $\mu^+ \rightarrow \mathrm{e}^+ X$ decay with X being a neutral long-living boson, often referred as Majoron or Familon.

The first concept of a new Trigger and Data Acquisition system[3] for the MEG II experiment dates to 2012, when it was realised that the former system could not scale to fit the requirements of the new apparatus. A complete picture of the system characteristics, detailing the interplay of the various components, was finalised in late 2014 with first prototypes available by the end of 2015.

Starting from the 2016, with the beginning of this Ph.D. work, all the major detectors of the experiment were installed; the TDAQ system gradually grew from 256 to 1536 channels and correspondingly complexity of operation increased. Currently the TDAQ system is being qualified, some design flaws were identified and will be fixed in the mass production foreseen for next year.

The online event selection capabilities grew together with the experiment being assembled, from individual channel pulse height discrimination to more refined algorithms based on signal amplitudes and sub-nanosecond time-stamping, all being processed in half a microsecond.

MEG II will have a small physics trigger rate, ranging between 6 Hz and 24 Hz depending on the final reconstruction performances. However it will feature a raw event data size of just below 14 MB resulting in an overall throughput around 250 MB/s. Custom readout boards paired with commercial networking equipment assemble the data before being zero suppressed and written to disk.

During my Ph.D., I spent roughly two years out at PSI to contribute in loco to the commissioning of the experiment. I have been directly involved not only in the commissioning of the Trigger subsystem but also on the Data Acquisition, which is strictly integrated with the fist one. As the local TDAQ expert, I strictly collaborated with colleagues from each detector during the commissioning phase.

Therefore my contributions to the MEG II experiment can be summarised in the development of the firmware providing the trigger logic and the ancillary functionality needed for its operation, in the development from scratch of the multithreaded readout code being used since 2019, in the reconstruction of trigger-related information at the offline processing level and, more in general, in the commissioning and daily maintenance of the actual TDAQ hardware. To provide a physics case for this thesis beyond the $\mu^+ \rightarrow e^+ + \gamma$ golden channel, I also evaluated the sensitivity and the feasibility of a $\mu^+ \rightarrow \mathrm{e} + X$ search at MEG II.

In the first chapter of this thesis, a review of the physics being involved in the search for $\mu^+ \rightarrow \mathrm{e}^+ + \gamma$ and $\mu^+ \rightarrow \mathrm{e}^+ X$ is given, also including a short description of some theoretical models that may contribute to the processes.

The second chapter describes the structure of the MEG II detector and how it will improve with respect to the former setup and in the third chapter, I describe the new Trigger and Data Acquisition system for the experiment, which is known as WaveDAQ.

The fourth and the fifth chapters will cover respectively the first performances of the TDAQ in the 2018 engineering data-taking and an evaluation of the trigger performances and the final data acquisition throughput.

Finally, in the last chapter, a feasibility study on a search of $\mu^+ \rightarrow \mathrm{e}^+ + X$ with the MEG I dataset is reported to highlight possible others CLFV searches within the reach of the upgraded MEG II detector.

1 CLFV Introduction

The Standard Model of particle physics (SM)[4, 5], representing our best knowledge of subatomic scale processes, is constructed on an enormous set of measurements performed at a wide range of energies, from atomic spectroscopy measurements at the electronvolt scale, to the energy frontier large scale experiments at collider, probing TeV processes.

All the current efforts in searching discrepancies between experiments and SM, both increasing the accelerators centre of mass energy, looking directly for unexpected new particles, or probing with detail known reaction rates, that may be boosted or decreased by unforeseen or anomalous interactions, proved unsuccessful.

While in the first case (centre of mass energy increasing) an assumption of the underlying physics is needed to define the experimental signature, the indirect probing method, by its nature, utilises an already defined process and the experimenter's ability lies in choosing the most sensitive channel to a given new physics search.

To qualify for indirect searches of new physics, a channel should have a clear Standard Model prediction and be sensitive to the largest possible number of theoretical models. Possibly the chosen channel should also couple to sectors of Standard Model that already show tensions with theoretical predictions.

The Charged Lepton Flavour Violation (CLFV) search, that includes the $\mu \rightarrow \mathrm{e} + \gamma$ or other exotic muon decay channels, fits this requirements by having a null value in the minimal SM and a negligibly small value assuming the model extensions including massive neutrinos[6]. The underlying physics is also strictly connected to the muon magnetic dipole moment measurement $g-2$ that is puzzling the particle physics community since a decade. CLFV allows to probe new physics up to a limit given only by the experimental detector resolutions, having at the same time a high sensitivity to beyond Standard Model physics, especially in the framework of Super-Symmetric and Grand Unification extensions[7].

1.1 The Flavour Structure in Standard Model of Particle Physics

Particles inside the Standard Model are grouped, according to their spin, into Fermions, having spin 1/2, and Bosons being either scalar particles with spin 0, such as the Higgs boson, or vector bosons with spin 1:the gluon, the charged $\mathrm{W}^{\pm}$, the neutral Z^0 and the photon. Excluding the octet of gluon fields that is associated with the $\mathrm{SU}(3)_C$ symmetry of Quantum Cromo Dynamics, all the other vector boson quantum fields are part of the ElectroWeak theory and are generated from the mixing following the spontaneous breaking, induced by the Higgs field, of the $\mathrm{SU}(2)_T \otimes \mathrm{U}(1)_Y$ weak isospin and hypercharge symmetry.

Fermions can be split into two groups: quarks, interacting with all the mentioned fields, and leptons, that have a null colour charge and therefore do not couple with the nuclear strong force mediated by gluons. The latter ones can be divided into three families each composed of a charged massive field paired with a neutral field; each family is tagged with a Lepton Flavour quantic number named by the corresponding charged particle being electron (e), muon (μ) and tau (τ) particle.

In the Standard Model there are indeed two, a priori distinct, bases to be used when describing the three families of fermion fields:

- The previously described base is known as Flavour base and defines the behaviour of particles according to their coupling with the weak force mediated by $W^{\pm}$ or, in other words, by their grouping to the same $SU(2)_T$ isospin doublet.
- The Mass base is the one that diagonalise the Yukawa coupling of fermions with the Higgs vacuum value, and are the ones describing the non interacting particle equation and, subsequently, a particle mass.

The common assumption of the neutrino fields being massless has a very huge impact on the SM Lagrangian: in terms of coupling, it means that they do not couple to Higgs field, so neutrino behaviour is only described by the Flavour, causing also the charged leptons mass to be diagonal in the flavour base and hence the impossibility of direct coupling of fields from different families. All interactions between families happen in SM through the exchange of a vector boson respecting the Flavour base and ultimately there are three independently conserved quantities: the Lepton Flavour numbers.

1.2 Lepton flavour violating processes in Standard Model with massive neutrinos

Neutrinos were observed to mix while propagating in vacuum with a now well understood mechanism[8] that, in a simplified two-flavour scheme, prescribes an oscillation probability given by:

$$P(\nu_\ell \rightarrow \nu_{\ell'}) = \sin^2(2\theta) \cdot \sin^2(1.27\frac{\Delta m^2[\mathrm{eV}^2]}{E[\mathrm{MeV}]}L[\mathrm{m}]) \tag{1.1}$$

The probability of equation 1.1 is composed of two independent factors, the first one is only dependent on the mixing angle θ that defines the transformation between the Flavour and the Mass bases. The second factor involves the neutrino energy E and the distance L between production and detection but also the difference of the two neutrino mass values $\Delta m^2 = (m_2)^2 - (m_1)^2$.

In this picture, the observation of a non null oscillation probability automatically selects a non trivial mixing angle θ and at least one of the two neutrinos to being massive ($(m_1)^2 \neq (m_2)^2$).

Following the first observation of a solar neutrino deficit in solar flux by the Homestake experiment[9], later confirmed by several independent experiments[1], the evidence of an oscillation in neutrino flavour was obtained by Super-Kamikande[11] and Sudbury Neutrino Observatory[12] experiments and the oscillation pattern was established.

A plethora of experiments studied the flavour-dependent boost or deficit in neutrino fluxes at different E/L ratios and provided[13] a measurement of the Δm^2_{ij} that are present in the three flavour extension of equation 1.1:

$$\Delta m^2_{21} = 7.39^{+0.21}_{-0.20} \cdot 10^{-5}\ \mathrm{eV}^2 \tag{1.2}$$

$$|\Delta m^2_{32}| = 2.525^{+0.033}_{-0.032} \cdot 10^{-3}\ \mathrm{eV}^2 \tag{1.3}$$

and the values for the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix that describes the three flavour transformation:

$$\begin{pmatrix} \nu_\mathrm{e} \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{\mathrm{e}1} & U_{\mathrm{e}2} & U_{\mathrm{e}3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \cdot \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.4}$$

$$= \begin{pmatrix} 0.797 \leftrightarrow 0.842 & 0.518 \leftrightarrow 0.585 & 0.143 \leftrightarrow 0.156 \\ 0.233 \leftrightarrow 0.495 & 0.448 \leftrightarrow 0.679 & 0.639 \leftrightarrow 0.783 \\ 0.287 \leftrightarrow 0.532 & 0.486 \leftrightarrow 0.706 & 0.604 \leftrightarrow 0.754 \end{pmatrix} \cdot \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.5}$$

where the values are the current best estimates at the 3σ level.

In the three flavours neutrino model the PNMS matrix, assuming that there are not other lepton families, is an unitary matrix which can be parametrised in terms of the three Euler rotation angles and a complex phase that is not yet bounded by experimental measurements.

There are different possibilities for the neutrino mass terms[8] and for the corresponding mixing matrix, to be added to the SM; the easiest one being the addition of three right handed neutrino fields that can then produce mass terms by Yukawa coupling to the Higgs vacuum expectation value, similarly to what happens for the other fermions.

The direct effect of such modification is the inclusion in the SM of vertexes like $\ell - \nu_i - W^\pm$ were $\ell = e, \mu, \tau$ and $i = 1, 2, 3$ whose coupling is given by the elements $U_{\ell\alpha}$ of the matrix in equation 1.5, in analogy to what happens on the quark sector with the Cabibbo-Kobayashi-Maskawa (CKM) matrix.

In such an extended model[6] the lepton flavour violating process $\mu \to \mathrm{e} + \gamma$ receives non null contribution to its amplitude by 1-loop level Feynman graphs like the one shown in figure 1.1.

This results in a decay width of

$$\Gamma(\mu \to \mathrm{e} + \gamma) = \frac{G_F^2 m_\mu^5}{192\pi^3} \cdot \frac{3\alpha}{32\pi} | \sum_{i=1}^{3} U_{\mu i} U^*_{\mathrm{e}i} \frac{m_i^2}{m_W^2} |^2 \simeq \Gamma\mu \to \mathrm{e}\nu\bar{\nu} \cdot 10^{-55 \leftrightarrow -54} \tag{1.6}$$

---

[1] A deficit of solar neutrinos was confirmed by Gallex/GNO, Kamiokande, Super-K, Borexino and SNO experiments while other effects were observed in atmospheric neutrinos by Super Kamiokande, MACRO and Soudan experiments[10]

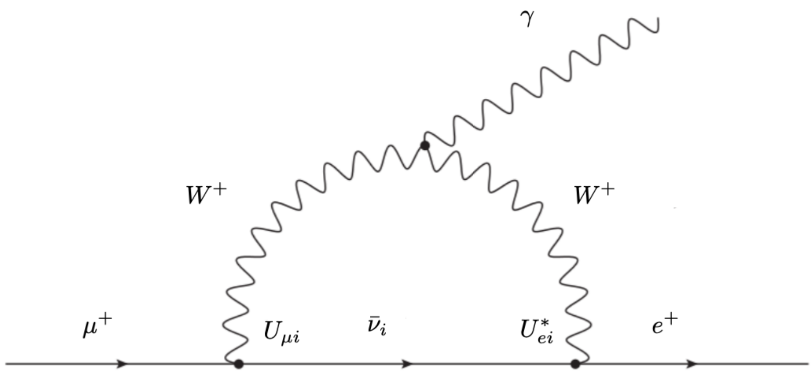

Figure 1.1: Feynman graphs contributing to $\mu \to \mathrm{e} + \gamma$ in massive neutrino extensions of SM; other diagrams are given by the photon line being attached to other charged particles.

where the first factor describes the SM decay width of muons according to $\mu \to \mathrm{e}+\nu_\mu+\nu_\mathrm{e}$ [2] and the other factor, equal to the $\mu^\pm \to \mathrm{e}\gamma$ branching ratio, can be evaluated by plugging the values from neutrino measurements.

This value, despite being non null, is well beyond the experimental reach (as reported later the sensitivity range of current searches is $\sim 10^{-14}$). Such small probability comes from the smallness of neutrino masses compared with W boson; the mass ratio suppresses the amplitude even if the values of the relevant elements of the PNMS matrix are significantly different from zero.

The smallness of the $\mu \to \mathrm{e}\gamma$ decay width in SM extensions with right handed neutrinos is directly correlated to an unsolved theoretical issue in this model: the required neutrino Yukawa coupling value to the Higgs vacuum expectation value needs to be fine tuned to values at the order of 10^{-13}, to push the 246 GeV vacuum value down to hundredths of eV of the expected neutrino mass range.

Other theoretical models[14] were developed to provide more appealing alternatives and explain the enormous gap between neutrinos masses and the electroweak scale.

An option is to generate neutrino mass terms not through Yukawa couplings with right handed fields but indirectly, by having additional physics at a much higher energy scale that weakly couples with the SM Lagrangian producing a suppressed neutrino mass term.

A family of such models is known as See-Saw theories. In such a framework a set of massive right handed neutrino fields interacts only with leptons through the Higgs field. If the right handed neutrino masses are pushed far beyond the electroweak scale, decoupling them from the SM Lagrangian, the conventional neutrino field turns out to

---

[2] G_F is Fermi constant which can be expressed in terms of the Standard Model constants like W boson mass m_W and electroweak coupling strength g resulting in: $G_F/\sqrt{2} = g^2/8m_\mathrm{W}^2$

have a very small mass term.

A side effect of the presence of massive right handed neutrinos is that the lepton-neutrino-W vertex is not only given by the PNMS matrix 1.5 but receives an additional contribution that may boost the $\mu \to \mathrm{e}\gamma$ reaction rate at observable levels.

1.3 Flavour Violation beyond Standard Model

The Standard Model, regardless of the many successes, still presents some shortcomings[15] on top of the already mentioned problem of neutrino masses ed oscillations.

- Dark matter and energy: cosmological measurements showed that usual matter (contained in SM) accounts only for 5% of the energy content of the universe, the remaining 95% is split into 20% of unknown matter (Dark Matter) and 75% of energy (Dark Energy).
- Matter-Antimatter asymmetry: while CP violation is a well established feature of weak interaction through the complex phase of the CKM matrix, the observed amount is not enough to explain the abundance of matter in the observable universe.
- Mass spectrum and force unification: SM does not explain the source of different Yukawa couplings and the origin of the gravitational, electroweak and nuclear strong force.

In the last years SuperSymmetric (SUSY) theories[15] guided both the theoretical and experimental searches of new physics, providing, when inserted in the bigger picture of grand unification theories (GUT), some interesting solutions to the mentioned shortcomings while at the same time also solving other technical theoretical problems such as the stability of the Higgs mass under radiative corrections. Both direct searches at LHC[16] or indirect ones like the Super-Kamiokande limit on proton life time[17] showed no evidence of SUSY-induced effects; however the amount of different SUSY theories and their huge space of parameters can be exploited to consider models that can explain the smallness of the effects searched for.

The main idea originating SUSY theories is that each standard model particle has an associated supersymmetric partner following the opposite statistics: a SUSY vector boson is associated to each SM fermion and an additional SUSY fermionic field to each boson, as described in table 1.1.

The CLFV effects are generated because there is no reason for an exact one to one correspondence between SM fields and SUSY counterparts; each SM particle can correspond to a single one but to a mixture of the various SUSY particles. This clearly gives contribution to $\mu \to \mathrm{e}\gamma$ and other lepton flavour violating processes, if SUSY mixing occurs, through diagrams like the ones in figure 1.2, whose contribution varies according to the masses of the sneutrino/chargino or slepton/neutralino of each particular SUSY theory. In general, nevertheless, there is consensus that, if a SUSY theory exists, lepton flavour violation is boosted by the processes of figure 1.2.

SM particle	**SUSY associated**	**SUSY spin**
charged leptons (e, μ, τ)	sleptons ($\tilde{\ell}_i$)	1
neutrinos (ν_{e}, ν_μ, ν_τ)	sneutrinos ($\tilde{\nu}_i$)	1
quarks	squark	1
neutral gauge bosons (γ, Z^0, h)	neutralinos ($\tilde{\chi}_i^0$)	1/2
charged gauge bosons ($\mathrm{W}^{\pm}$)	charginos ($\tilde{\chi}_i^{\pm}$)	1/2

Table 1.1: Fields being added in supersymmetric theories with their corresponding spin.

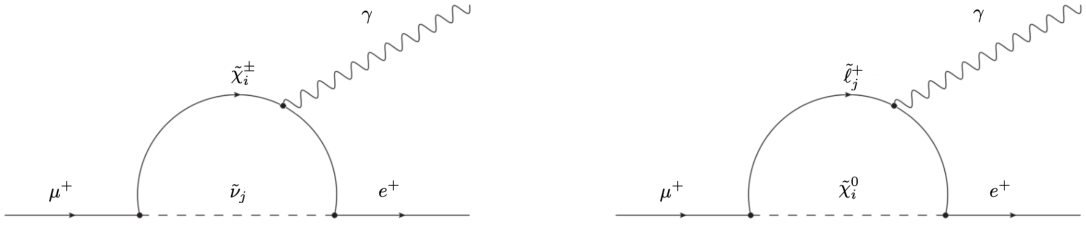

(a) chargino-sneutrino

(b) neutralino-slepton

Figure 1.2: Feynman graphs contributing to $\mu \to \mathrm{e} + \gamma$ in SUSY theories.

In models with sufficiently large mixing[18, 19], the existing bound on $BR(\mu^+ \to \mathrm{e}^+\gamma) < 4.2 \cdot 10^{-13}$[1] can be even more effective setting limits on SUSY particle mass spectrum than direct searches.

1.4 The $\mu \to \mathrm{e}\gamma$ channel

As already pointed out muon decays provide a very interesting probe to new unknown physics by relying on well known Quantum Electro Dynamics and Electro Weak processes. Moreover the limited energy of such processes makes the experimental setup much smaller in size than the ones used for high energy searches. On the other hand, the smaller energies involved make these particles more subject to material effects while being measured so that very low X^0 detectors must be used. Furthermore exceptional timing and spatial resolutions are needed to remove the accidental background and reconstruct event pile-up.

1.4.1 Phenomenology of muon decays

The muon is one of the three charged leptons of the Standard Model; it carries an electric charge and is part of a weak isospin doublet with its corresponding neutrino ν_μ. Because of this, it is subject to both the electromagnetic and the weak nuclear forces with the corresponding vertexes to photon, $\mathrm{W}^{\pm}$ and Z^0.

The muon mass value[20] of $105.6583745(24)\ \mathrm{MeV/c}^2$ is very precisely known thanks to muonic atom spectroscopy measurements; the mass term arises from the muon coupling to the Higgs field with, as a side effect, the generation of an additional vertex involving muons and Higgs boson that resulted fundamental in the discovery of Higgs boson in 2012.

Due to the weak interaction the muon decays into an electron plus two neutrinos $\mu^+ \to \mathrm{e}^+\bar{\nu}_\mu\nu_\mathrm{e}$ with a life time of 2.1969811(22) μs, a very large value which is due to how small muon-electron mass difference is when compared to the electroweak scale[3].

In nature muons are abundantly produced from cosmogenic proton interactions on top of the atmosphere where nuclear interactions generate charged pions that subsequently decay into muons. Thanks to their long decay time and the boost they get at the production, atmospheric muons can reach the see level where they were discovered in 1927[21]. Following a similar chain, muons can be produced at accelerators by proton interactions in solid targets and delivered from the highly active production target by means of magnetic fields to experiments with fluencies as high as $10^8\mu$/s.

1.4.1.1 Michel Decay and Radiative Michel Decay

Within the Standard Model and according to all experimental evidences so far[20], the full list of muon decays is reported in table 1.2 with the corresponding branching ratios (BR).

Process	Name	BR	Notes
$\mu^+ \to \mathrm{e}^+\ \nu_\mathrm{e}\ \bar{\nu_\mu}$	Michel Decay	$\sim 100\%$	
$\mu^+ \to \mathrm{e}^+\ \nu_\mathrm{e}\ \bar{\nu_\mu}\ \gamma$	Radiative Michel Decay (RMD)	$6.0(5) \cdot 10^{-8}$	[22] $E_\mathrm{e} > 45$ MeV and $E_\gamma > 40$ MeV
		$3.3 \pm 1.3 \cdot 10^{-3}$	[23] $E_\gamma > 20$ MeV
$\mu^+ \to \mathrm{e}^+\ \nu_\mathrm{e}\ \bar{\nu_\mu}\ \mathrm{e}^+\ \mathrm{e}^-$		$3.4(4) \cdot 10^{-5}$	[24] $E_\mathrm{e} > 17$ MeV

Table 1.2: Muon SM Branching Ratios for positive muons, CP conjugate of the process takes place for negative muons.

The second and third decay channels in the table 1.2 originate from radiative corrections of the main decay mode $\mu \to \mathrm{e}\ \nu\ \bar{\nu}$, also known as "muon Michel decay" from the french physicist Louis Michel who firstly performed a systematic theoretical study of the features of this decay.

In case of positively charged leptons, the positron spectrum expected for a tree level calculation of the Michel decay is shown in figure 1.3a for fully polarised muons. The difference with negative leptons is given by the angular distribution: the charged lepton emission angle θ_e with respect to the muon polarisation P is $\propto\ 1 + P\cos(\theta_\mathrm{e})$ for positive muons and $\propto\ 1 - P\cos(\theta_\mathrm{e})$ for negative ones.

---

[3] The muon decay width in SM is indeed given at tree level by $\Gamma = \frac{g^2 m_\mu^5}{6144\pi^3 m_\mathrm{W}^4} \simeq \frac{m_\mu^5}{m_\mathrm{W}^4}$, neglecting small contributions due to electron mass and W boson propagator; the same does not apply to charged pion decays even though their mass is similar to muon one, in fact the spin suppression in $\pi^\pm$ decays favours the muon decay mode over the electron one and thus reduces their live time by two orders of magnitude, down to 10^{-8} s.

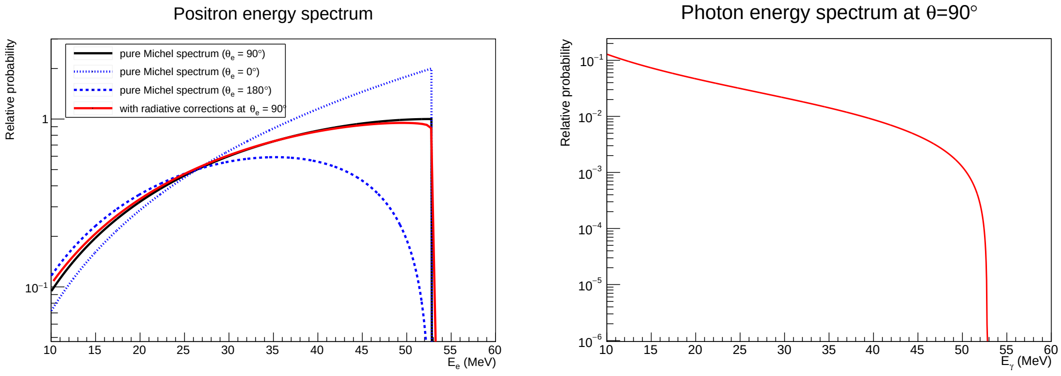

(a) (b)

Figure 1.3: Michel energy spectrum characteristic of muon decays (in the plot $\mu^+ \rightarrow \mathrm{e}^+ \, \nu_\mathrm{e} \, \bar{\nu}_\mu$), evaluated from equations (30, 51) in reference [6], with and without first order radiative corrections (a) and photon spectrum emitted from radiative processes (b).

Such dependency can be understood by spin considerations of the involved particles: in particular because the two neutrinos have fixed helicities (one being left and one right handed); the emission of the three particles needs therefore to be aligned with the muon spin direction because of the underlying W boson exchange and, for the same reason, the positron is preferentially emitted in the muon spin direction with the two neutrinos aligned in the opposite direction. In such kinematic configuration, the positron has a momentum at the kinematic endpoint value of $m_\mu/2 \simeq 52.8$ MeV.

Radiative corrections to 1.3a automatically produce a $\mu \rightarrow \mathrm{e} \, \nu \, \bar{\nu} \, \gamma$ “Radiative Michel decay” (RMD) channel with the photon spectrum shown in figure 1.3b. The radiative nature of the photon spectrum is clearly shown by the infrared divergency of such distribution for small photon energies. The effect on the Michel spectrum (at $\theta = 90°$) is shown in red in figure 1.3a, with an observable reduction of events near the endpoint due to part of the energy being subtracted by the soft photon emission.

If photon coming from a radiative Michel converts into an electron-positron pair the decay $\mu \rightarrow \mathrm{e} \, \nu \, \bar{\nu} \, \mathrm{e}^+ \, \mathrm{e}^-$ spectrum is observed. In this case the positron spectrum ranges from 0 to the kinematic endpoint due to muon mass, but with an average total kinematic energy of the three leptons peaked at $m_\mu/2$.

1.4.1.2 Lepton flavour violating decays

Theoretical efforts[25] have been made to provide a consistent model-independent description of CLFV processes: in recent years a framework was established by explicitly adding all allowed non renormalizable operators $O^{(i)}$, with dimension i, to the SM lagrangian:

$$\mathcal{L} = \mathcal{L}_{\mathrm{SM}} + \frac{1}{\Lambda} \sum_\alpha C^{(5)}_\alpha O^{(5)}_\alpha + \frac{1}{\Lambda^2} \sum_\alpha C^{(6)}_\alpha O^{(6)}_\alpha + \ldots \tag{1.7}$$

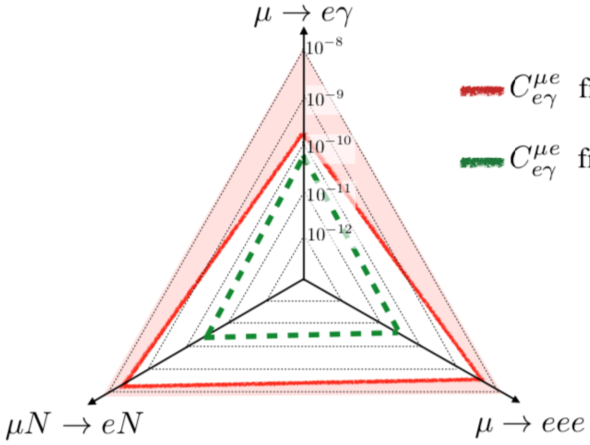

Figure 1.4: Current and future limits on the $C^{\mu\mathrm{e}}_{\mathrm{e}\gamma}$ from various charged lepton number violation searches, from [26].

In this sum the cut off energy Λ represents the energy limit of the resulting effective field theory. It parametrises the scale of energy where new physics is expected to give sensible differences with respect to the introduced point-like high dimension operators, just like the Fermi theory of beta decays with the dimensioned G_F parameter is a suitable description of the processes for energy much smaller than m_W.

In equation 1.7 the weight of the various operators O_α is described by the corresponding adimensional constant C_α. In this framework each experimental result can be converted either in a lower limit on the new physics scale Λ, by assuming the coupling of the corresponding operator equal to 1, or in an upper limit on the coupling strength C_α by fixing the energy scale. In this effective field theory the prediction for the $\mu \to \mathrm{e}\gamma$ process is given at tree level by:

$$\Gamma(\mu \to \mathrm{e}\gamma) = \frac{m_\mu^3 \nu^2}{8\pi\Lambda^4}(|C^{\mathrm{e}\mu}_{\mathrm{e}\gamma}|^2 + |C^{\mu\mathrm{e}}_{\mathrm{e}\gamma}|^2) \tag{1.8}$$

with ν being the Higgs vacuum expectation value and the $C^{\mathrm{e}\mu}_{\mathrm{e}\gamma}$ and $C^{\mu\mathrm{e}}_{\mathrm{e}\gamma}$ the constants associated with the operators directly containing the $\mu - \mathrm{e} - \gamma$ vertex.

In addition to $\mu \to \mathrm{e}\gamma$, which is mediated by $\mu - \mathrm{e} - \gamma$ operators at tree level, other processes such as $\mu \to \mathrm{eee}$ decay and muon conversion in nuclear field $\mu N \to \mathrm{e}N$ receive loop-level contributions involving the same vertex; all the channels can therefore probe the couplings of $\mu - \mathrm{e} - \gamma$ operators, although with a different sensitivity.

In figure 1.4 the limit on the $C^{\mu\mathrm{e}}_{\mathrm{e}\gamma}$ extracted from the $\mu \to \mathrm{e}\gamma$ is compared, at fixed energy $\Lambda = 1$ TeV, with the corresponding limits from the other processes. In order to disentangle the contribution of $C^{\mu\mathrm{e}}_{\mathrm{e}\gamma}$, the other coupling strengths are fixed at zero while producing the figure 1.4; therefore highlighting the different sensitivities of each channel if the effect of any new physics is limited to that operator.

With the recent limit[1] by the MEG experiment $BR(\mu^+ \to \mathrm{e}^+\gamma) < 4.2 \cdot 10^{-13}$, the limit on this particular coupling is heavily constrained by the $\mu \to \mathrm{e}\gamma$ process (at

$\sim 2 \cdot 10^{-10}$ level); planned next generation experiments will have a similar sensitivity, an order of magnitude below the current one. To reach such goal $\mu \to$ eee and $\mu N \to \mathrm{e}N$ limits need to be two orders of magnitude more stringent than $\mu \to \mathrm{e}\gamma$: they both contain an additional electromagnetic vertex which introduces a fine structure constant $\alpha \sim 1/137$ and suppresses the branching ratio; the QED vertex accounts, respectively, for the conversion of the photon into a pair of leptons (in $\mu \to$ eee) or to have it absorbed by the nuclear field (for $\mu N \to \mathrm{e}N$).

This should not indeed be considered as a direct competition among different experiments; the other two channels are more sensitive to other operators, like the 4-fermion point-like operators $\mu - \mathrm{e} - f - f$ (where f is an electron in one case and a quark in the other) that only contributes at loop level to $\mu \to \mathrm{e}\gamma$. Therefore only the combination of various measurements can have some discrimination power between new physics models in case of any positive CLFV measurement.

If compared with $\tau \to \ell\gamma$ decays the sensitivity to new physics scales as $(m_\tau/m_\mu)^3 \simeq 5 \cdot 10^3$ because of the dependency of equation 1.8 on the lepton masses, neglecting any theory dependent scaling of the $\tau - \ell - \gamma$ vertexes with respect to the $\mu - \mathrm{e} - \gamma$ one. However, experimental τ searches in CLFV processes, with current BR limits in the order of 10^{-8}, are more complex given the difficulty to collect the required statistics with particles having a lifetime of $2.9 \cdot 10^{-13}$ s. For this reason τ beams cannot be created and $\tau - \bar{\tau}$ pairs need to be produced by proton or electron beam collisions in experiments covering a wider panorama of searches in addition to CLFV ones.

1.4.2 History, signature and backgrounds in $\mu \to \mathrm{e}\gamma$ search

Ideally an experiment looking for lepton flavour violation is a counting experiment, with a window in all the observables defined only as a function of the experimental resolutions, considering that the SM induced background is negligibly small. However the smearing introduced by the response function of actual detectors and the corresponding finite dimensions of the counting window causes some events nearby the kinematic endpoint to enter the count. To improve the sensitivity modern experiments employ a likelihood analysis to fit a signal shape over an experiment-dependent background to blinded data, ensuring an unbiased result.

Looking for decays of a given particle automatically fixes the best experimental observation frame which is the particle rest frame but this ultimately limits the available energy for the produced particles, in particular for muon decays where the signal energy is below 100 MeV. Having such small energies requires indeed a careful design of the detectors to achieve the goal resolutions, with the limit evolving according to the development of detector and acceleration technologies.

1.4.2.1 Evolution of CLFV limits

Following the discovery of the muon lepton in 1937 by Neddermeyer and Anderson in cosmic ray events[21], several experiments were carried out to try understanding the characteristics of such particle. After 10 years[27], it became evident that it was a particle

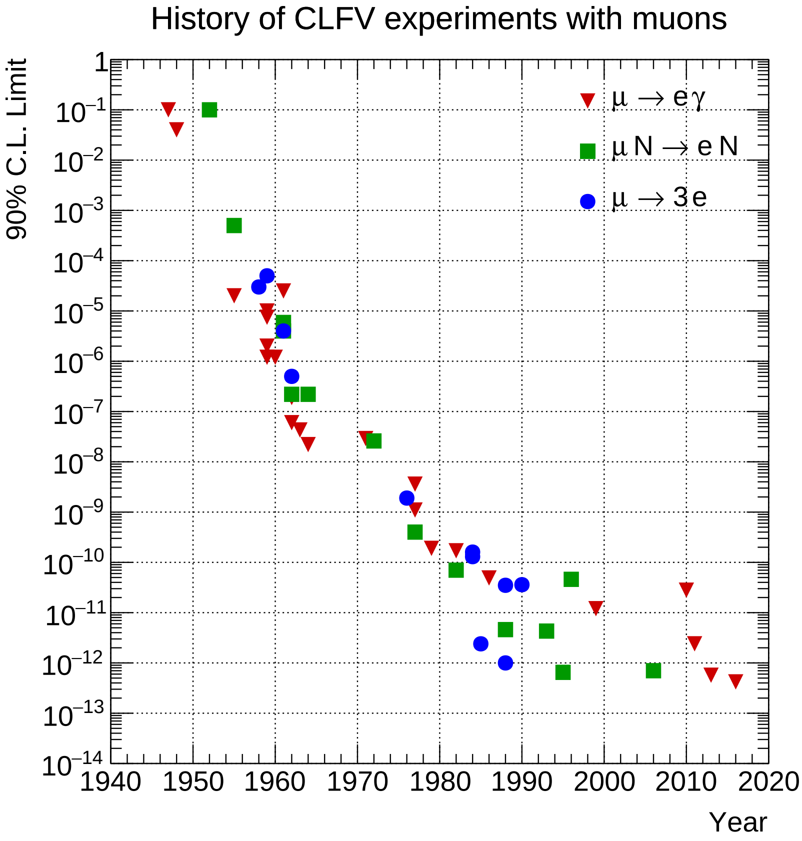

Figure 1.5: Historic evolution of limits in three lepton flavour violating channels.

with the same charge of electron but with mass 200 times larger; it was also understood that the particles mediating the strong interaction, in Yukawa theory of strong forces, were the pions π, discovered in 1947 by Powell[28], and not the muons as it was initially believed. The particle physics community began speculating if the muon could be an excited state of the electron, but early measurements of charged particle spectrum in muon decays already showed that it was not compatible with a two-body decay and that at least two neutral particles were generated in the process. Such measurement by Pontecorvo and Hincks results in an initial limit on $\mu \to e\gamma$ branching ratio at few percent[29].

The real step forward needed a much larger statistic than that possible with cosmic rays which became possible with developments in accelerator technologies. The first direct search of the $\mu \to \mathrm{e}\gamma$ process (and $\pi \to \mathrm{e}\gamma$) was carried out at Columbia University cyclotron in 1955[30] and set an upper limit on the branching ratio to $2{\cdot}10^{-5}$. As shown in figure 1.5, it was then followed by other experiments[31] all exploiting the same principle: pion beams from proton accelerators were stopped in a target and a delayed coincidence was opened to allow for muon decays to be selected.

The modern theoretical framework was being established[32] at the same time with the idea of the two neutrinos coming from muon decays not being particle and antiparticle, as it was believed until then, but instead two distinguishable particles carrying some additional quantum numbers: the lepton flavour. Such hypothesis was experimentally confirmed in 1962 with the famous experiment by J. Steinberger, M. Schwartz and others at the Brookhaven AGS[33] accelerator. In the article conclusions the experimentalists clearly state:

> However, the most plausible explanation for the absence of the electron showers, and the only one which preserves universality, is then that $\nu_\mu \neq \nu_\mathrm{e}$; i.e., that there are at least two types of neutrinos. This also resolves the problem raised by the forbiddenness of the $\mu^+ \to \mathrm{e}^+ + \gamma$ decay.

In the very same period similar studies were carried out in the quark sector observing a very small, but not null, probability for flavour changing neutral current processes such as $K_L \to \mu^+\mu^-$ and $K^+ \to \pi^+\mu^+\mu^-$, which paved the way for the Glashow-Iliopoulos-Maiani mechanism[34] and, after the observation of CP violation in the neutral kaons (1964)[35], to the CKM matrix that foresaw a third family of fermions.

Such prediction proved to be astonishingly correct by the experimental discovery of J/Ψ and Υ quarkonium resonances, the τ lepton and, more recently, the top quark and the τ neutrino.

In the meanwhile searches for $\mu^+ \to \mathrm{e}^+ + \gamma$ and other lepton flavour violating processes were improved with the creation of meson farms where high intensity beams of muons (or pions) are used to collect the required statistics for rare searches in a reasonable time frame. In the next future very high intensity muon beams will be available which could allows, if coupled with refined detector technologies, to push down the CLFV limits by order of magnitude as will be discussed later.

1.4.2.2 $\mu \to \mathrm{e}\gamma$ detection strategy

The signature of the $\mu \to \mathrm{e}\gamma$ process is very clean in the muon rest frame: regardless of the underlying physics the energy of the resulting particles are monochromatic and equal to 52.8 MeV, neglecting the small contribution of electron mass. Furthermore the photon and the electron are emitted back to back so that the total momentum is conserved and the process is instantaneous; the time coincidence can be therefore used to identify the decay.

In order to exploit such peculiarity, muons are stopped in some target so that they can decay at rest in the experiment frame. Positive muons μ^+ are best suited because, differently from μ^-, they cannot orbit around the nuclei and do not require correction of the decay spectrum due to atomic binding energy.

There are two possibile sources of events that can mimic, close to the endpoint, the signal topology:

Radiative Michel decays can mimic the signal, as described in 1.4.1.1, if the energy subtracted by neutrinos is small. It constitutes the only physic background for a $\mu \to \mathrm{e}\gamma$ search and can be only suppressed by carefully measuring the energies and the relative angle of the generated particles.

Accidental coincidences are generated by a random superimposition of a positron from a Michel decay with a photon coming either from an uncorrelated Radiative Michel decay or from the annihilation in flight or Bremmstrahlung of any Michel positron. Given the accidental nature of the event, the emission time of the positron and photon are not correlated.

A proper comparison of the abundance of the two categories of background events is possible[6], in the approximation of a counting experiment, once the signal region is identified in the kinematic space constituted by E_e and E_γ energies and $\Theta_{\mathrm{e}\gamma}$ relative angle. In order to give an estimate of the relative contributions of there two backgrounds one can think of establishing the signal region at $E_\gamma^\mathrm{sig} \simeq E_\mathrm{e}^\mathrm{sig} \simeq 52.8$ MeV and $\Theta_{\mathrm{e}\gamma} = \pi$, with a size defined by a two sigma window so to be 95% efficient on signal. The following adimensional parameters can be defined:

$$\delta x = 2\frac{\sigma(E_\mathrm{e})}{E_\mathrm{e}^\mathrm{sig}} \tag{1.9}$$

$$\delta y = 2\frac{\sigma(E_\gamma)}{E_\gamma^\mathrm{sig}} \tag{1.10}$$

$$\delta z = 2\sigma(\Theta_{e\gamma}) \tag{1.11}$$

$$\tag{1.12}$$

The radiative branching ratio can be integrated in the signal region resulting in functions of δx, δy and δz. A full calculation is beyond the scope of this thesis and is reported in [36].

In the MEG and MEG II case, the angular resolution is good enough for δz to be relevant in the integration of the RMD branching ratio[4]. The expected branching ratio in the signal region, assuming a combined angular resolution of 8 mrad, is given in the plot of figure 1.6.

The accidental background can also be parametrised[36] by means of the same adimensioned variables with the addition of δt, the time coincidence resolution, and R_μ, the muon stop rate; under the assumption of the photon spectrum being dominated by RMD:

$$B_\mathrm{acc} = 2R_\mu \delta t \cdot 2\delta x \cdot \frac{\alpha}{2\pi}(\delta y)^2[\ln(\delta y) + 7.33] \cdot \frac{(\delta z)^2}{4} \tag{1.13}$$

where B_acc is the equivalent branching ratio of accidental events.

In the case of equation 1.13 the branching ratio can be easily factorised in its various components:

- the first term accounts for the probability of a second muon decaying within the time window δt;

---

[4] In the decay E_e, E_γ and $\Theta_{\mathrm{e}\gamma}$ are correlated, so a kinematic constraint can be set on the physically allowed $\Theta_{\mathrm{e}\gamma}$ range only by measuring the two energies. If the angular resolution is poor $\delta z > 2\sqrt{\delta x \cdot \delta y}$, the kinematic constraint defines the signal range to integrate the branching ratio.

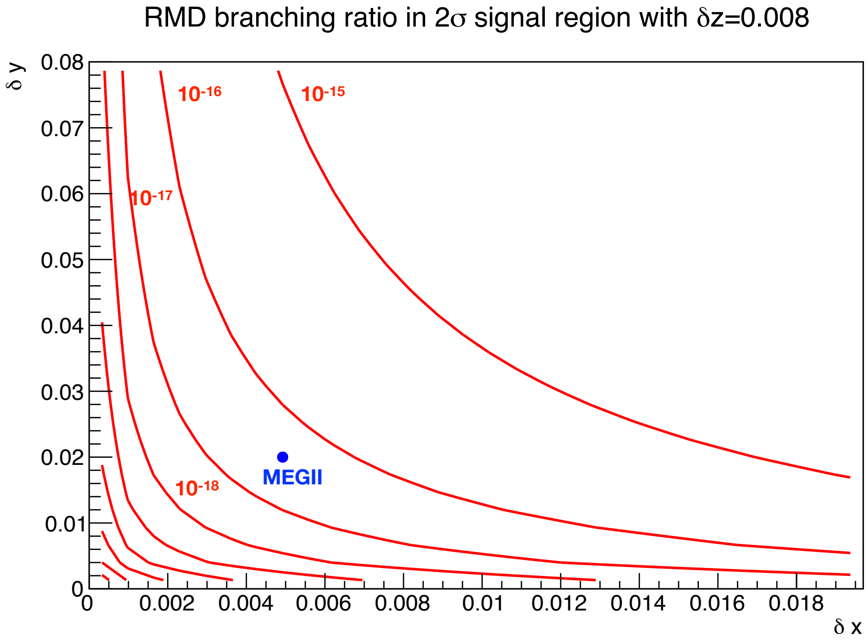

Figure 1.6: Radiative Muon decay branching ratio in signal region, MEG II case ($\sigma(E_{\mathrm{e}}) =$ 130 keV, $\sigma(E_\gamma)/52.8 = 1\%$, $\Theta_{e\gamma} = 5$ mrad) the value is $2 \cdot 10^{-16}$.

- the second term is the cut of positron energies on a flat background from muon decays at the endpoint;
- the third term is the requirement on the RMD photon energy and it scales roughly quadratically $(\delta y)^2$. This different dependence with respect to δx is induced by the different shape of the two spectra at the endpoint;
- the last term is the relative angle cut on an isotropic distribution, scaling as $(\delta z)^2$, due to the first Taylor term of $cos(\Theta_{\mathrm{e}\gamma})$ term for small angles.

Considering all these terms, a muon stopping rate of $R_\mu = 7 \cdot 10^7$ μ^+/s and a time coincidence window $\delta t = 80$ ps, the accidental branching ratio in the signal region is $1 \cdot 10^{-15}$, dominating on the radiative Michel background by a factor 10 and thus constituting the largest background source for the MEG II experiment. For this reason, excellent timing performances are crucial for MEG II: the request of time coincidence provides a way to reduce the overwhelming accidental background of modern and future searches using high intensity muon beams.

The branching ratio of accidentals, if compared to the expected sensitivity of $6 \cdot 10^{-14}$ (reported in section 2.5), indicates that the background of the MEG II experiment (as well as for the former MEG I) is non null, which motivates the use of likelihood analysis tools to exploit all possibilities to separate the signal from the background.

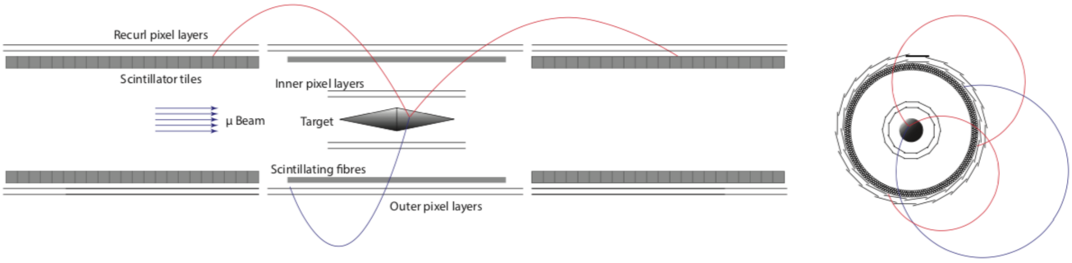

Figure 1.7: Sketch of the Mu3e experiment, in the configuration with a sensitivity of 10^{-15}.

1.5 Other CLFV serches with muons

As previously described, the $\mu \to e\gamma$ decay is not the only process being looked for in the CLFV panorama. Several searches have been made not only with muons but also with τ (in similar channels such as $\tau \to \mu\gamma$), in the kaon sector (the most famous one being $K_L^0 \to \mu e$), as well as in the B decays. The recent upgrade of the Belle experiment[37] is expected to push the limits on both B and τ decays by almost two orders of magnitude.

In the muon sector, in addition to $\mu \to e\gamma$, the most studied channels are $\mu \to eee$ and $\mu \to e$ conversion in a nuclear field.

1.5.1 $\mu \to eee$, Mu3e experiment

The $\mu^+ \to e^+e^-e^+$ decay is extremely challenging especially because of the reduced kinematical constraints on signal topology coming from a three body decay, the most stringent cut comes from the request of all particle momenta lying in the same plane and from their simultaneous emission. The variability on the energy and emission angle requires a very precise and fast spectrometer with a minimal radiation length and (possibly) a low momentum threshold.

The goal of the Mu3e experiment[38] is to probe the $\mu \to eee$ decay down to 10^{-16} in three separate phases, each with increased sensitivity until reaching the four orders of magnitude below the current limit of $1 \cdot 10^{-12}$ held by the SINDRUM experiment[39]. This improvement in sensitivity is obtained with the extensive use of ultra-thin silicon pixel tracking modules[38] inside a solenoidal spectrometer field: these detectors will be placed in four layers around a double-cone Mylar target, resulting in a full 2π azimuthal acceptance (as shown in figure 1.7).

The main contribution to the sensitivity comes from the three steps in momentum resolution that will be achieved, ranging from 730 keV/c RMS in the first phase to 280 keV/c in the last one.

A set of scintillating fibres and tiles will improve the timing performances, from the second phase onward, allowing the experiment to scale in muon stopping rate up to $2 \cdot 10^9$ μ^+/s.

The Mu3e experiment is the most direct competitor of the MEG II experiment, not just because the two underlying physics share similarities as described in section 1.4.1.2, but mainly because Mu3e will be hosted at the same facility that MEG II is currently using. Mu3e is still in an early construction phase, and while some complex part like the magnet are being already assembled, the final pixel hybrid module is still being developed.

1.5.2 $\mu A \to \mathrm{e}A$: Mu2e and COMET experiments

The muon to electron conversion in nuclear field has the advantage of a higher signal energy, respect to $\mu \to \mathrm{e}\gamma$: a beam of negative muons will be stopped in an aluminium target where muonic atoms are formed. When a muon decays in the nuclear electric field according the usual Michel decay, a recoil of nucleus can provide the conservation of four-momentum even if the neutrinos are emitted with very small energies. Given the huge mass difference between nucleus and electron the resulting kinematic endpoint is much higher ($\sim m_\mu$).

The conversion signal is given by a monochromatic electron exactly at muon mass energy minus the nuclear binding and recoil energy; the same value can be reached, at the endpoint, by the mentioned in-orbit Michel muon decay. However the single particle nature of the signal exposes the experiment to any type of background process that could produce an electron near the signal energy.

The structure of the two experiments Mu2e (shown in figure 1.8a from [40]) and COMET[41] (figure 1.8b), both searching for μ-e conversion, is very similar and dictated by the necessities of the measurement: protons from a pulsed beam are stopped in a production target and muons from pion decays are transported by means of a curved solenoid (S-shaped for Mu2e and COMET phase II and C-shaped for COMET phase I) to the final aluminium muon target[5]. The signal electron is measured by a spectrometer using straw tubes followed by a crystal calorimeter (CsI for Mu2e and LYSO for COMET).

Several sources of backgrounds have been identified, the most important being the beam-induced one coming from pion contamination of muon beam when moved between the production and the measurement targets. Such events, and others induced by proton beam, can be suppressed by vetoing the electrons detected in coincidence with the spill of particles impinging on the measurement target. In order for this to work a peculiar time structure of the proton beam is required: Mu2e and COMET will use beams with pulsed structure (each pulse ~ 100 ns wide spaced by $\sim 1-2$ μs) so that in the inter-spill a detection window can be opened[6].

Current best limits on μ-e conversion came from the SINDRUM II collaboration, using a more conventional solenoid spectrometer with drift chambers for tracking. The limit

---

[5] The choice of target material is crucial for the delayed window approach described later: muonic atom live time decreases with the nuclear charge and is also susceptible to the nucleus internal charge distribution. Aluminium have one of the longest decay time of 865 ns[42], for comparison gold shows a muonic decay time of only 74 ns.

[6] An additional challenge is posed by the suppression of protons out of time with respect to the main bunch; experiments define a beam “extinction factor”, being the fraction of out-of-bunch protons, which have to be kept under 10^{-10} in order to provide the required background rejection.

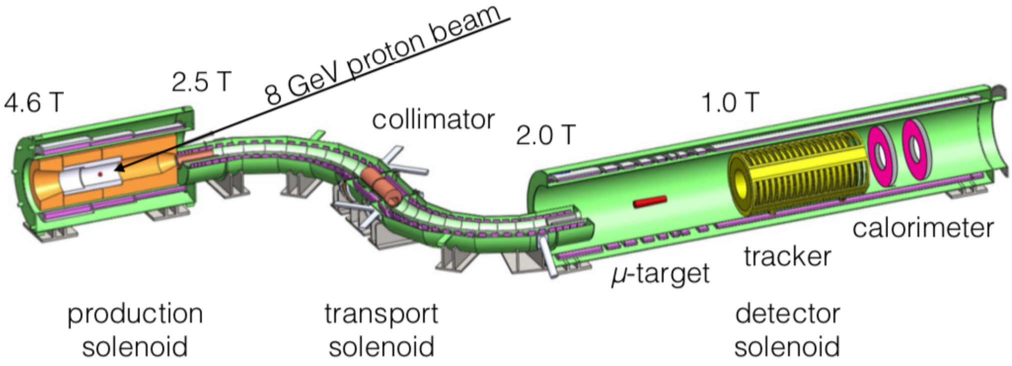

(a) Mu2e experiment

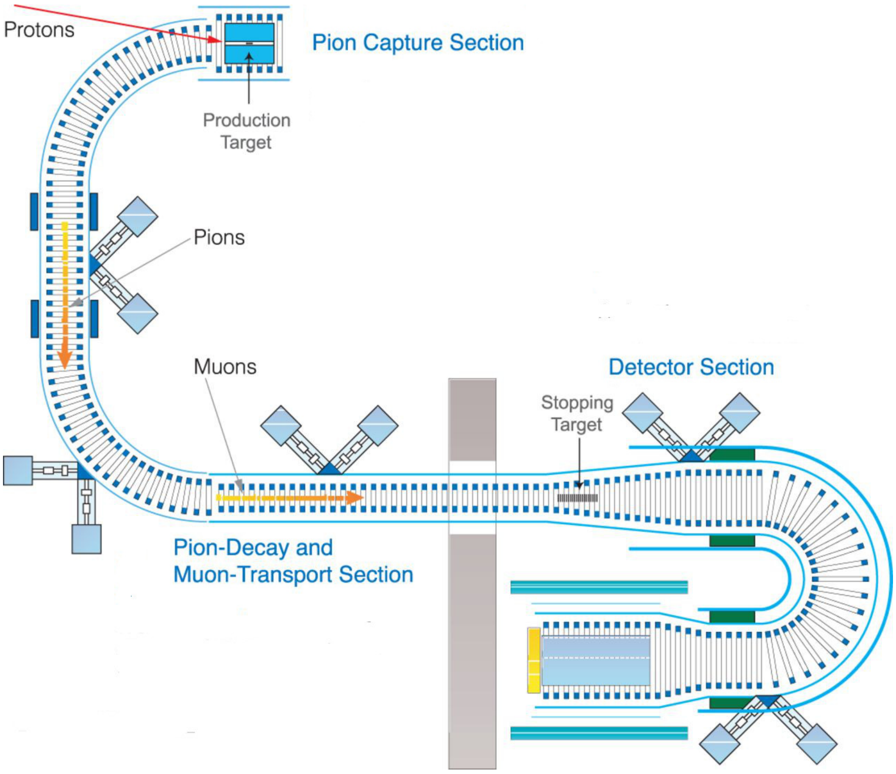

(b) COMET experiment

Figure 1.8: Sketch of future $\mu - \mathrm{e}$ conversion experiments.

on muon conversion probability is usually quoted as ratio to muon captures by nuclei so to reduce the theoretical uncertainties introduced by the nucleus internal structure. The SINDRUM II conversion ratio limits are $4.3 \cdot 10^{-12}$, using a titanium target[43], and $7 \cdot 10^{-13}$, using a gold one[44].

1.6 Exotic CLFV searches with muons

Beyond the three golden processes for CLFV searches with muons, which were reported previously, there are others involving more exotic searches and usually prescribing the presence of new particles. In this section few channels looking for direct beyond Standard Model physics with muons are presented, all involving the search for a neutral scalar field X that may be produced in interactions violating the lepton flavour conservation.

1.6.1 Physics motivations

There are several theoretical models[14, 45, 46] that include a new field X that may couple to fermions belonging to different families; the most interesting group of these theories aims at explaining the three-family structure of Standard Model by introducing additional symmetries to originate the fermion families. Such symmetries must be broken at a large energy scale since quarks and leptons show experimental mass values that span several orders of magnitude among the different families. If the flavour symmetry is a continuous global symmetry, Nambu-Goldstone bosons known as “familons” are generated in the symmetry breaking process.

A subset of familons can include the breaking of the global lepton number symmetry with the goal of generating Majorana neutrino mass terms; these were the first familons introduced by theoreticians and were historically named Majorons[47].

Supersymmetric models violating the R-parity[7] can also generate familons.

Familons are commonly assumed to be pseudo-scalar fields and, for generality, the couplings of their current to the fermionic ones happen through two distinct matrices in the flavour space: one describes the coupling to the vector current and the other to the axial one. This results in a possible difference in the behaviour of the matrix element that is generated by the underlying broken theory. Pure V-A interaction is the most theoretically interesting for its possible analogy with the weak interactions.

An interesting class of theories[48] identifies the familon pseudo-scalar field X with the one required by Peccei-Quinn theory of strong interaction[49]: among all allowed terms of QCD Lagrangian, only one describes a CP-violating effect and the experimental bound on its strength θ was proven to be incredibly small by the precise measurement of the electric dipole moment of neutron. The required fine-tuning of θ is commonly referred to as the “strong CP problem” or “θ-term puzzle”. The Peccei-Quinn solution was to promote the coupling constant θ to a field, later named “axion”.

---

[7] R-parity is an internal symmetry of some supersymmetric theories explaining why the couplings of vertexes violating baryon number B and lepton one L are proven experimentally below observation[15]. Having the R-parity charge defined as $(-1)^{3(B-L)+2s}$, where s is the particle spin, the R-parity symmetry is also known as “B minus L” symmetry.

In the last years there was a huge theoretical effort in merging the familon and the axion field and this sparkled the experimental interest in the search of processes involving a flavoured axion field. The rare decays of K, D and B mesons is the primary channel used for the search with channels such as $K^+ \to \pi^+ X$ or $B^\pm \to K^\pm X$ as well as the neutral ones. The most stringent limits on the mentioned kaon meson decays come from the E949+E787 experiments[50] for the charged decay with a branching ratio $< 7.3 \cdot 10^{-11}$ and from the KOTO experiment[51], being $< 5 \cdot 10^{-8}$, for the neutral one. The CLEO collaboration[52] holds the most stringent limits on the charged B with a branching ratio of decays of $4.9 \cdot 10^{-5}$. Further improvements are expected from new generation flavour experiments such as NA62[53] and Belle II[37].

Familons in the lepton sector are searched by means of the possible $\mu \to \mathrm{e}X$ decay and its radiative counterpart $\mu \to \mathrm{e}X\gamma$; though existing and future limits are not competitive in flavoured axion sensitivity with the ones from meson decays, the measurements of these additional channels provides further handles for discriminating among models because they rely on vector and axial coefficients which are different from those of the main channels.

A viable option, being currently pursued with the MEG I dataset to look for familon-like particles, is to search for $\mu \to \mathrm{e}X$ events with the X decaying in two photons which are also detected by the Liquid Xenon detector. Such process is indeed less inclusive than the direct $\mu \to \mathrm{e}X$ but the experimental signature is much cleaner. On the other hand, the angular acceptance for both photons is very reduced and trigger biases must be carefully taken into account. Furthermore displaced vertexes for the X decay (up to 1 cm)[54] are not yet bounded, causing the photon acceptance to depend on the X vertex displacement hypothesis.

Currently the best limits comes[54] from beam dump searches and from $\mathrm{e}^+\mathrm{e}^- \to \gamma\gamma$.

1.6.2 $\mu \to \mathrm{e}X$, the TWIST experiment

A direct search for the X boson is possible by looking at modifications of the positron spectrum: the two-body topology produces, similarly to $\mu \to \mathrm{e}\gamma$, a monochromatic signal positron. The X mass might be significantly different from zero and thus the positron from $\mu \to \mathrm{e}X$ will have a well defined energy of:

$$E(\mathrm{e}^+) = \frac{m_\mu^2 - m_X^2 + m_\mathrm{e}^2}{2m_\mu} \tag{1.14}$$

that, for massless X, reduces to $\sim m_\mu/2$, in analogy with $\mu \to \mathrm{e}\gamma$.

The angular distribution is usually parametrised as $\propto\ 1 + AP\cos(\theta_\mathrm{e})$ with $A = -1$ being the same of standard Michel decay, $A = 0$ being an isotropic decay and $A = +1$ the decay preferentially opposite to the Michel one.

There is no a priori reason to prefer a particular value of A, however A=-1 is the most experimentally challenging because no angular information can be used to discriminate it from the Michel decay with the only observable being the positron energy.

The Michel background provides a huge background for the signal peak, especially near the endpoint were modelling of the Michel spectrum is fundamental for these searches.

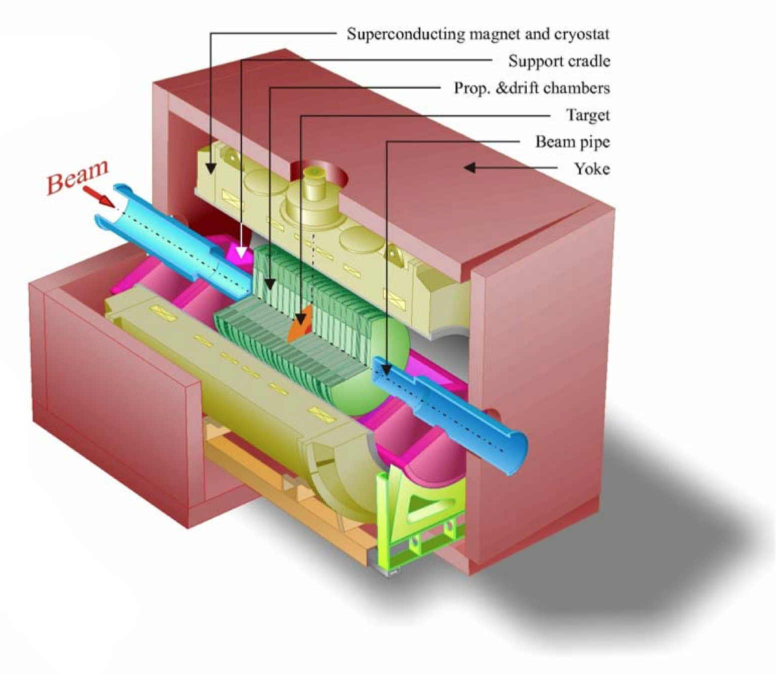

Figure 1.9: Drawing of the TWIST experiment at TRIUMF.

The combination of massless X and $A = -1$ angular distribution is a very interesting, but very difficult to explore, scenario.

1.6.2.1 TWIST experiment

The TWIST experiment (TRIUMF Weak Interaction Symmetry Test, shown in figure 1.9) was an experiment at the TRIUMF Canadian national laboratories designed to carefully measure the spectrum parameters of Michel decays[55]. A beam of surface muons was stopped in a 75 μm aluminium or 30 μm silver plane target placed perpendicular to the beam direction, at the centre of the experiment. The complete experiment was immersed in a 2 Tesla solenoidal magnetic field which caused both the incoming muon and the generated positron to move in helix trajectories. Both the incoming and the outgoing tracks were detected by an array of planar drift or proportional chambers also perpendicular to the beam direction, for an overall resolution of (58 keV/c)/$|\sin\theta_\mathrm{e}|$ at 52.8 MeV (with θ_e being the angle relative to the muon spin and, for construction, to the magnetic field axis).

The TWIST experiment collected a dataset of $1.1 \cdot 10^{10}$ muon decays with a minimum bias trigger on a muon entering the apparatus. After quality cuts and the selection of positrons in the fiducial acceptance $5.8 \cdot 10^{8}$ events were selected for the $\mu \to \mathrm{e}X$ analysis.

The limit on $\mu \to \mathrm{e}X$ by the TWIST experiment is divided in two distinct mass regions. For $m_X > 13$ MeV/c^2 the momentum of the signal positron is < 52 MeV/c, this simplifies the search because the energy scale and spectrometer performances can

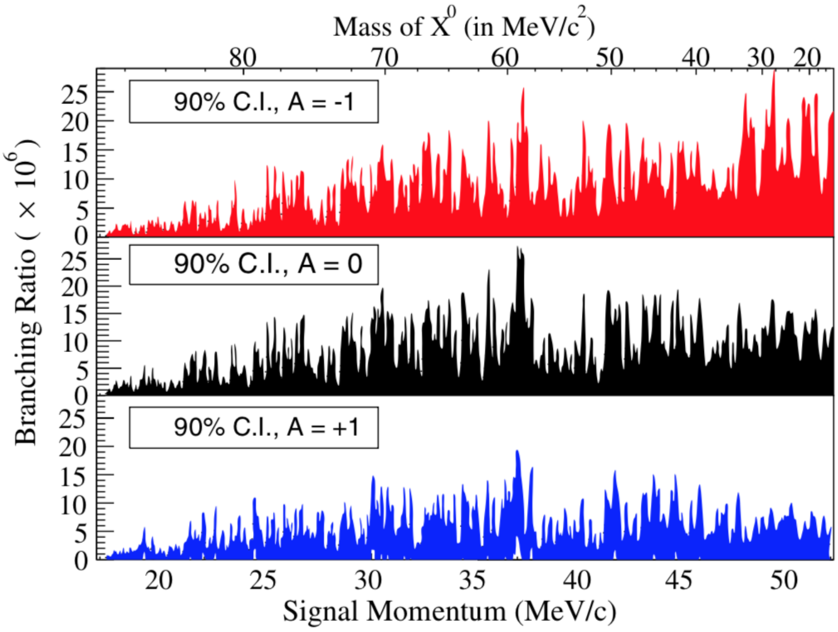

Figure 1.10: Best limit on $\mu^+ \to e^+ X$ with massive X from TWIST experiment[56], including statistical and systematic uncertainties. The A represent the angular shape parameter, $A = 0$ is the isotropic and $A = -1$ is the same as Michel decay.

be fitted at the endpoint that is far enough not to alter the fit. Figure 1.10 shows the limit for three angular distribution hypothesis; on average the limit on the branching ratio is $1.5 \cdot 10^{-5}$ for the small mass range[56].

For light or massless X ($m_X < 13$ MeV/c^2) the position of the peak is near the endpoint of the Michel decay; the analysis in this case needs a careful evaluation of the energy scale contribution. The TWIST collaboration carried out a careful a priori study of various systematics entering the definition of the spectrum absolute scale and extracted an a priori correction of the energy scale being -4.3 ± 6.1 keV. The obtained upper limits for massless X are shown in table 1.3 together with the propagated error on the limit resulting from the a priori estimate of the energy scale.

Angular distribution	upper limit $m_X = 0$ (at 90%)	upper limit uncertainty
A=-1	$5.8 \cdot 10^{-5}$	$10.73 \cdot 10^{-5}$
A=0	$2.1 \cdot 10^{-5}$	$3.64 \cdot 10^{-5}$
A=+1	$1.0 \cdot 10^{-5}$	$1.26 \cdot 10^{-5}$

Table 1.3: Best limit on $\mu^+ \to e^+ X$ with massless X from TWIST experiment[56] and corresponding 1-sigma error.

In their paper, the TWIST collaboration does not clearly state how the momentum

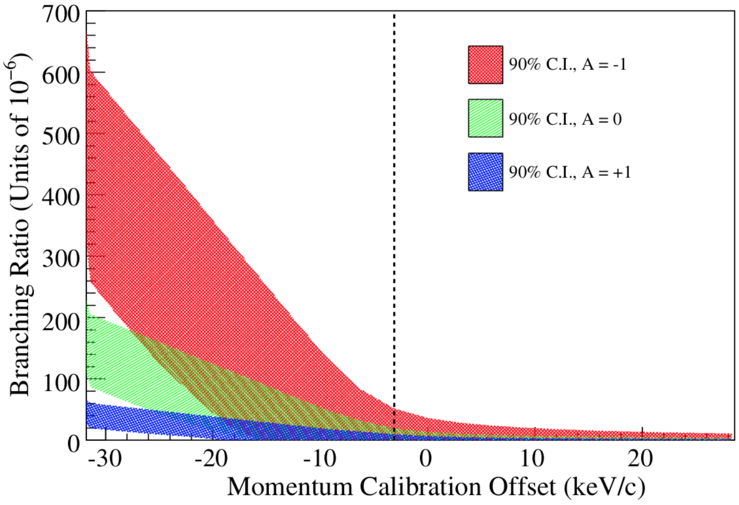

Figure 1.11: Feldman-Cousins intervals with respect to the spectrometer momentum scale for $\mu^+ \to e^+ X$ with massless X from TWIST experiment[56]. Dashed line is the best a priori estimate of the error scale.

scale uncertainties of table 1.3 should be interpreted. It is pretty unconventional to quote an upper limit uncertainty, considering that, in principle, the upper limit is an exact number and any uncertainty should be incorporated in the limit itself. At first glance an upper limit of $(5.8 \pm 10.73) \cdot 10^{-5}$ makes little sense because the uncertainty is much bigger than the obtained value; the computed branching ratio would be therefore compatible with 0 as well as with negative values, which have unphysical meaning.

In their paper, TWIST authors shows, for $m_X < 13$ MeV/c^2, "the Feldman-Cousins confidence intervals determined at the endpoint as a function of the momentum calibration offset", which are reported in figure 1.11 for the three possible angular distributions. The dashed vertical line shows the central value of their a priori momentum offset estimate and the intercept with the bands gives the nominal TWIST limits of table 1.3. From the article it is not clear how the uncertainty in the momentum scale affects the plot of figure 1.11: it might be included in the calculation of the bands but this is not clearly stated.

For this thesis the nominal values of TWIST upper limit will be used; a careful check on the analysis procedure by TWIST should however be pursued, in the future, if their value has to be combined with measurements by other experiments.

An alternative way, which will be used later in this thesis, to obtain an upper limit in this situation is to evaluate a 90% contour in the branching ratio-momentum scale plane. In this case the upper limit is given by the projection of such contour on the branching ratio axis. This procedure indeed results in an overcoverage of points in the 2D space but provides the worst possible branching ratio value which is still acceptable at the 90% level.

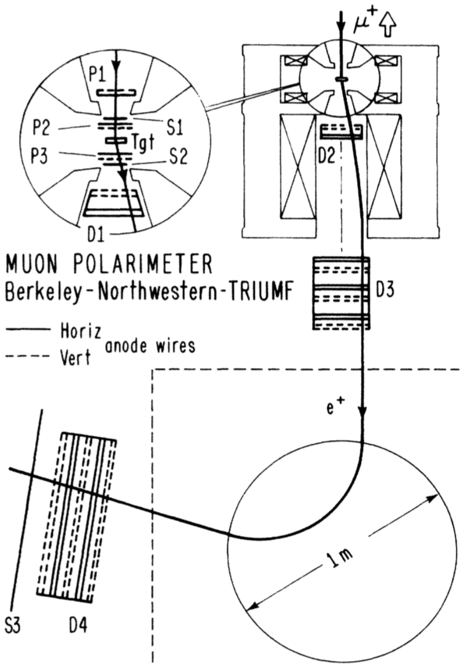

Figure 1.12: A. Jodidio spectrometer sketch with closeup of target are where the spin resonator is placed. PX are proportional chambers, DX are drift chambers and SX are plastic scintillators.

1.6.2.2 Other $\mu \to eX$ searches

The most sensitive search for $\mu \to eX$ is presently from an experiment by A. Jodidio also at TRIUMF[57]. It set a branching ratio upper limit of $2.6 \cdot 10^{-6}$, but it was carried out in the specific case of A=0, isotropic distribution of a massless X boson.

The setup, in figure 1.12, is made of a single arm spectrometer made of proportional/drift chambers and plastic scintillators with a 1.1 Tesla magnetic field in a 1 m region. The resulting momentum resolution was gaussian with a sigma of 68 keV/c for 52.8 MeV/c positrons.

To study the angular dependence of the positron spectrum, muons were stopped inside a spin rotator; such device employs magnetic fields either to hold the spin direction, when a longitudinal field is applied, or to rotate it, if a perpendicular magnetic field is applied.

The collected data can be then categorised in two datasets: the first one had the muon spin being held with its original negative polarisation and the spectrometer, which is placed downstream from the target, recorded a vanishing Michel spectrum due to the angular positron emission distribution. The second dataset have the spin rotator in operation collecting the abundant Michel decays in the spin direction.

If the signal angular distribution is not $A = -1$, the signal efficiency will not reduce in the spin held sample, resulting in better signal to noise ratio with respect to the spin precessed data. At the same time, that second dataset provides the spectrometer absolute

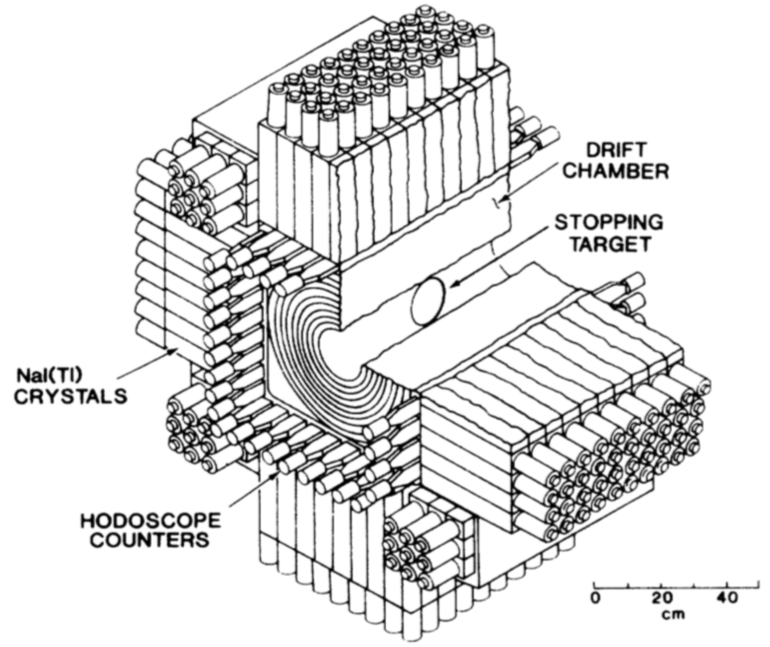

Figure 1.13: Schema of Crystal Box experiment at Los Alamos Meson Physics Facility (LAMPF).

scale measurement that is still found to be the biggest systematics of the analysis.

1.6.3 $\mu \to \mathrm{e}X\gamma$, the Crystal Box experiment

Being strictly connected to the $\mu \to \mathrm{e}X$ decay, the $\mu \to \mathrm{e}X\gamma$ process contains two observable particles that can be used to reduce the Standard Model background which is, in this case, the radiative Michel decay $\mu \to \mathrm{e}\nu\bar{\nu}\gamma$.

The current best upper limit on $\mu \to \mathrm{e}X\gamma$ branching ratio is $1.1 \cdot 10^{-9}$ and it comes from the Crystal Box experiment[58], a multipurpose muon apparatus that also set limits on $\mu \to \mathrm{e}\gamma$, $\mu \to \mathrm{eee}$, $\mu \to \mathrm{e}\gamma\gamma$ in 1988. The first two limits were superseded by later experiments but the $\mu \to \mathrm{e}\gamma\gamma$ and $\mu \to \mathrm{e}X\gamma$ still hold to date.

The Crystal Box experiment, shown in figure 1.13, is designed to detect positrons and photons with an almost 2π acceptance: muons decay at rest in a 350 μm slanted plastic target, the produced positrons are tracked in a cylindrical drift chamber before being stopped in a segmented NaI calorimeter that is also used to measure photons. Some plastic scintillators were sandwiched in between the drift chamber and the calorimeter to provide charged particle identification for trigger purposes and to improve the timing. Having no magnetic field there was no charge sign identification and the drift chamber was used only for vertexing and not for momentum measurement.

The $\mu \to \mathrm{e}X\gamma$ analysis relies on measuring the invariant mass of the X particle that can be computed from the positron track and the reconstructed photon cluster assuming that they are produced from the same muon decay. The invariant mass calculation requires a very precise measurement of the photon-positron relative angle as well as a good measurement of the particle energies.

observable	**Crystal Box**	**MEG II**
tracking resolution	130 μm	110 μm
positron momentum/energy	3.7 MeV	130 keV/c
photon position	1.8 cm @ R=23 cm	2 mm @ R=58 cm
photon energy	7% @ 55 MeV	1% @55 MeV
relative time	488 ps	84 ps

Table 1.4: Comparison of Crystal Box and MEG II performances.

The Crystal Box apparatus had very good positron tracking performances however the energy and time measurement are dwarfed from those of MEG II, as shown in table 1.4. There is indeed the possibility for MEG I/MEG II to explore the $\mu \rightarrow \mathrm{e}X\gamma$ decay but, having a detector acceptance much smaller than Crystal Box, an extrapolation of the reachable sensitivity is not straightforward.

2 The MEG II Experiment

As described in section 1.4.2.1, the study of rare muon decays moved from bench-top setups, employing cosmic rays, to more complex experiments using accelerator-produced muons as the only way to collect the required amount of statistics. The MEG experiment completed its three years long data-taking at Paul Scherrer Institut (PSI) in Switzerland ad set the most stringent limit on $\mu^+ \to e^+ + \gamma$ branching ratio to $4.2 \cdot 10^{-13}$ at 90% confidence level[1] with an involvement of ~ 60 researchers from Italy, Japan, Switzerland, United States of America and Russian Federation.

The MEG II[2] was built on top of the know-how of the previous experiment, maintaining the same detection scheme but upgrading or replacing all the detectors to achieve an additional order of magnitude sensitivity.

2.1 Description of the MEG II Experiment

The MEG experiment exploits the fixed kinematic of the two-body decay by having the muons being stopped in a thin foil target at the centre of the experiment and there decaying at rest in the experiment reference system, the usage of positive muons μ^+ is required to exploit the constraints of a possible two body decay at rest and to avoid the formation of bound states that could provide additional decay channels.

The reference system is defined with the Z axis along the beamline and the Y axis pointing vertically as shown in figure 2.1, in such a frame the photon is detected by a Liquid Xenon detector located at negative X and covering 11% of the whole solid angle seen from the target. The signal positron is emitted opposite to the photon and bends in a non uniform magnetic cylindrical spectrometer surrounding the target and composed by a low mass drift chamber followed by a fast plastic scintillator time of flight detector to improve the timing of the track.

2.1.1 $\pi E5$ PSI Muon Beamline at $S\mu S$ facility

Muon beam production at particle accelerators is typically achieved by bombarding carbon targets with protons to produce charged pions that then decay into muons. In the case of the Swiss Muon Source $S\mu S$ facility at PSI, sketched in figure 2.2, the main proton beam is produced by a Cockroft-Walton source paired with an injector stage and is then accelerated to the final 590 MeV energy by a ring cyclotron; this energy is chosen to maximise the pion production while being under threshold for kaon production(1.58 GeV) and capable to drive a huge beam current, as high as 2200 μA ($1.4 \cdot 10^{16}$ particle/s). Extracted protons hit a first thin target named “Target M” before reaching the thicker rotating “Target E”, in this way a total of 7 experimental

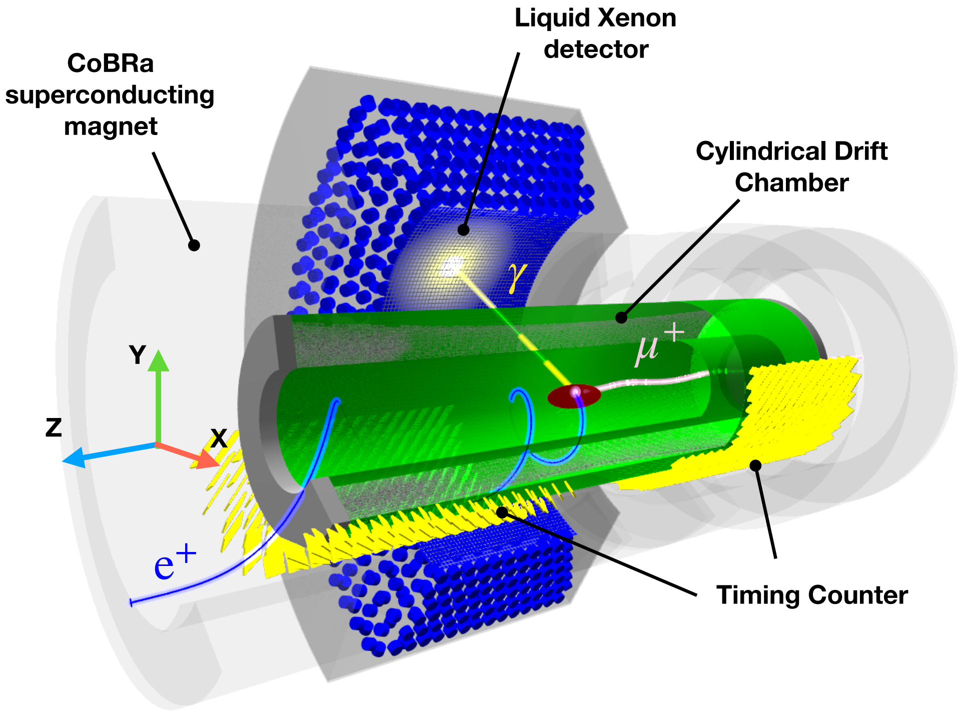

Figure 2.1: Sketch of the MEGII detectors.

beamlines are created, branching from the production targets, each with a peculiar beam component and energy spectrum. For completeness, the same accelerator complex can also provide protons to the new n2EDM experiment, aiming to measure the neutron electric dipole moment. Moreover, the proton beam emerging from the Target E is recycled by the SINQ neutron spallation source.

The $\pi E5$ Muon Beamline is the one used by the MEG experiment: it collects particles emitted backward in the Target E, with magnets being set to collect “Surface muons”; the beamline muons are indeed generated by pions that decayed at rest on the surface of the same production target. This results in an almost monochromatic beam of $8 \cdot 10^8$ μ^+/s at 28MeV which is extracted from the production region and is subsequently cleaned from positron contamination by an electrostatic separator.

At this point the muon flux does not show any structure of the original 40 MHz radio-frequency used in the proton cyclotron: the decay time of charged pions is 26 ns that closely match the 25 ns radio-frequency period, resulting in a continuous “DC-current” beam. In this configuration no accelerator-related timing of events is possible but the experiment does not suffer from the duty cycle originating from events being bunched with gaps having no events in between, and the required statistics can be integrated in

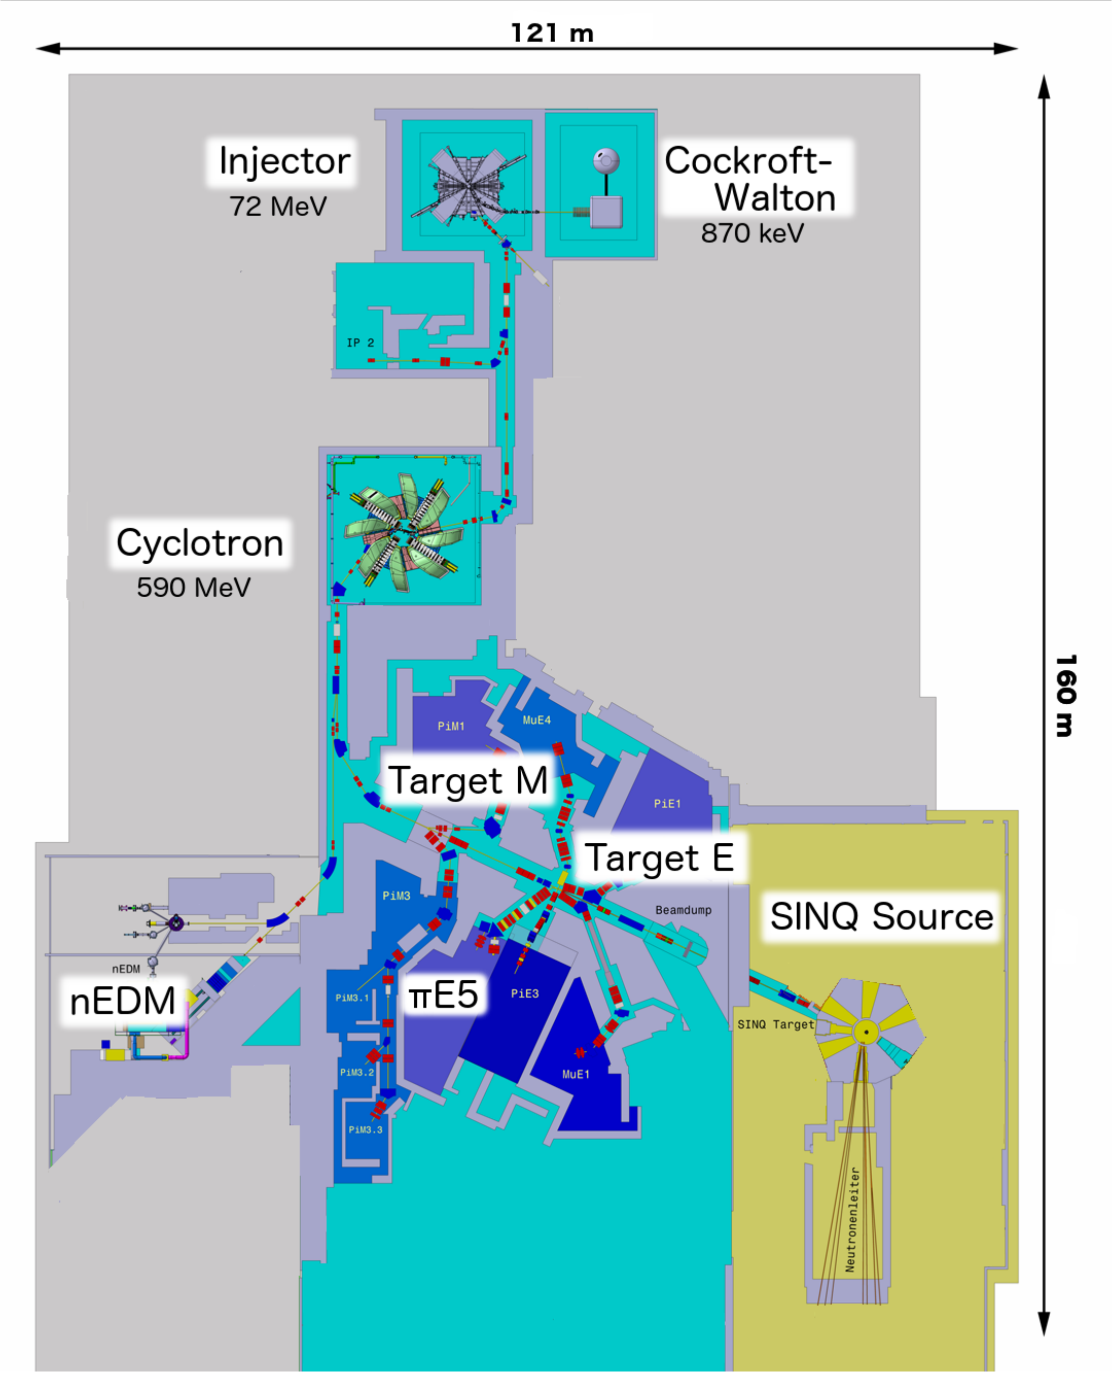

Figure 2.2: PSI $S\mu S$ facility as part of the High Intensity Proton Accelerator.

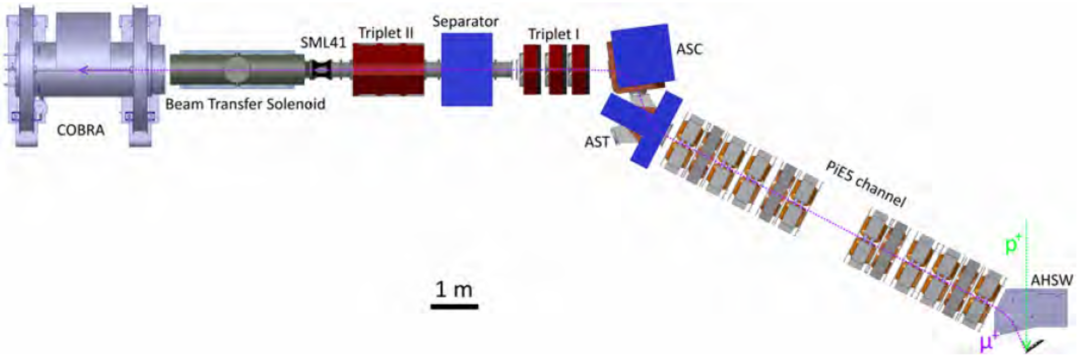

Figure 2.3: Magnet arrangement in $\pi E5$ beamline. Proton are stopped in Target E at the right and muons are collected by the $AHSW$ large-acceptance dipole. The PiE5 channels is filled by alternated Quadrupole and Sextuplole magnets with slits to extract the beam from the target area. Muons enter the experimental area by means of two dipoles (AST and ASC) and then cross two quadrupole triplets and an electrostatic separator to obtain the final focus. Finally the Beam Transport Solenoid feeds the beam into the experiment CoBRa magnet.

a shorter amount of time.

The complete structure of the $\pi E5$ beamline is shown in figure 2.3: starting from the proton stopping on target E, the muons are collected from a large acceptance dipole and extracted from the target area by means of the "PiE5 channel". This part of the beamline shapes the muon beam using alternated quadrupole and sextupole magnets and ends with a pair of dipoles selecting the desired momentum. Once entered the experimental area, the muon beam passes two triplets of quadrupoles with the electrostatic separator in between for particle selection. Then it enters the Beam Transport Solenoid (BTS), a 0.54 T superconducting magnet that provides the final focusing and contains a 300 μm degrader to reduce beam energy and separate accelerator vacuum from the helium atmosphere in the following part of the beamline. The degrader thickness has been chosen to tune the range of the muons in the MEG target and achieve the maximum number of stopped muons. At production, all muons coming from pion decays are fully polarised and, at the end of the beamline, a residual polarisation of -0.86 was measured[59], with the biggest unpolarised contribution coming for pions decay not at rest but nearby the surface of the production target.

For calibration purposes, the same beamline can be tuned to select a 70 MeV negative pion beam or a 52 MeV positron beam.

2.1.2 MEG II Target and beam monitor devices

Muons propagate in the helium volume up to the centre of the MEG II experiment where they are stopped in a 130 μm thick sheet of Polyvinyltoluene-based scintillating

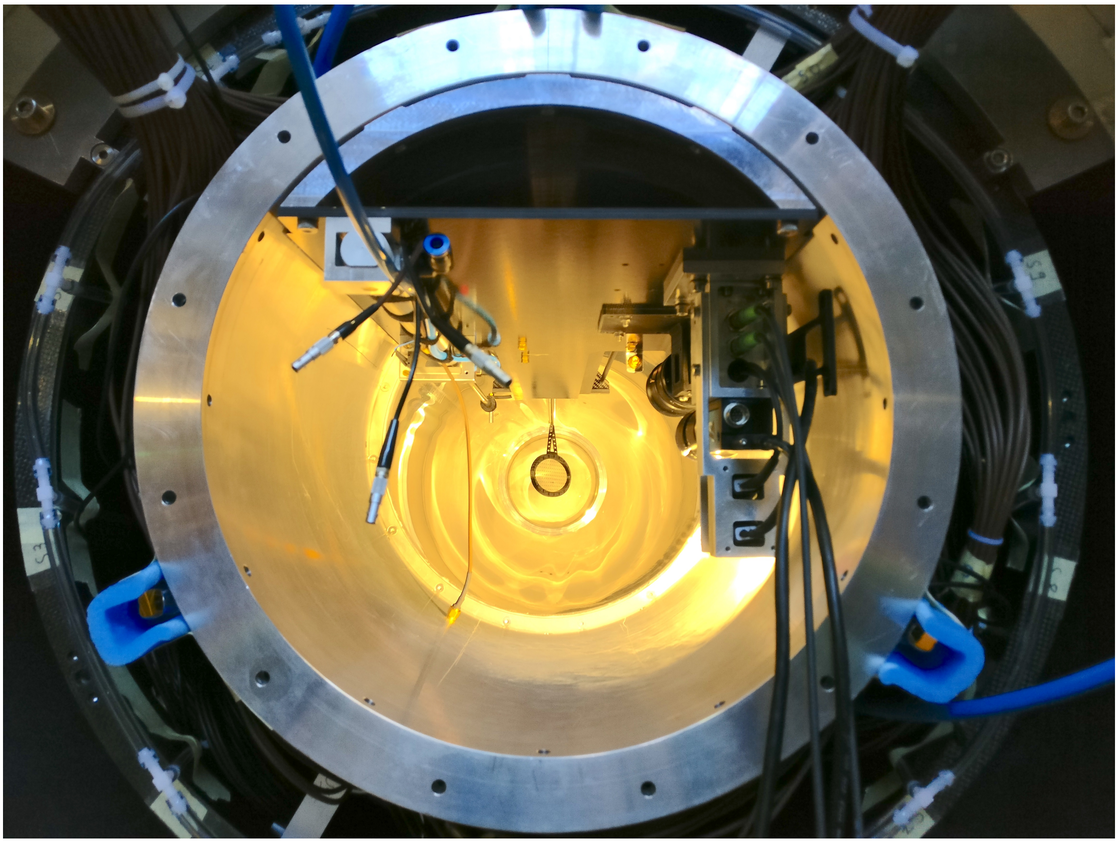

Figure 2.4: First prototype of MEG II scintillating target as seen from the beam upstream, the final support structure will include three corner cubes for laser tracker position measurement. Target Monitoring CCD cameras are shown on the right.

material (BC400B from Saint Gobain). The target has two opposite requirements: a high stopping efficiency for muons, requiring a thick amount of material, and a small scattering angle for the produced particles, requiring a small path inside its material. Those two requirements can be met by having the thin target normal being tilted with respect to the beam propagation direction by approximately 20°, increasing the path for the incoming particles by the inverse of the cosine of the tilt angle (see figure 2.4).

The target can be inserted and extracted from the beam by means of a pneumatic system but the reproducibility of such operation is of the order of one millimetre which, together with an overtime loss of planarity, constituted one of the biggest systematics of the former MEG experiment. Once every year the target position will be surveyed with respect of the experiment coordinate system by means of an external laser tracker measuring three reflectors that are mounted on the target frame. An array of commercial CCD cameras, visible on the right in figure 2.4, will then continuously image MEG II target and compare the position of some reference dots printed on its surface against those of a reference image to detect short term displacements and planarity alterations.

The choice of a scintillating material for the target, even if the scintillating properties will age pretty fast, allows for quick imaging, by a dedicated CCD camera placed downstream, of the beam position on the target. Some external factors such as the proton hit position on the target E can actually vary the outside the direct control of the experiment and produce a shift of the muon beam on the target.

Images from CCD camera are paired by two auxiliary detectors: a scintillating fibre sampling device (SciFi) can be inserted into the beam between BTS and the experiment to provide an almost non destructive monitoring of the upstream beam position and size together with a quantitative measurement of the muon flux. The other detector is a matrix of 9 × 9 small LYSO crystals that can be inserted in place of the target without dismantling any part of the experiment, and provides a second, more precise, measurement of the beam properties at its final destination.

2.2 The Liquid Xenon detector

2.2.1 Liquid Xenon as scintillating material

The reason for the choice of liquid xenon as detection material can be understood by the comparison (table 2.1) with other scintillating materials used in particle physics, in particular it shows a very good light yield, just 25% smaller than the brightest NaI(Tl) with a similar radiation lengths, fixing the required amount of material to be almost the same to fully contain the shower.

	Light yield ph/keV	**Radiation length** cm	**Molière radius** cm	**Decay time** ns
Liquid Xenon	26	2.9	5.2	4.3 - 27
Liquid Argon	-	14	9.0	5.0 - 1600
NaI(Tl)	38	2.6	4.1	250
BGO	9	1.1	2.3	300
LYSO	27.6	1.2	2.1	45

Table 2.1: Liquid Xenon properties compared with other common scintillators. Radiation length and Molière radius from [60], Light yield and decay time from [61, 62].

BGO and LYSO crystals can better contain the shower localising the interaction region and thus providing more spacial resolution but only the LYSO has a short enough decay time to help disentangling pileup in a high flux environment. However the solid state of those materials makes difficult to deploy a large scale homogeneous calorimeter presenting losses on the connection of two crystals that impacts on the energy resolution.

A liquid scintillator is best suited for a large volume uniform detector and, among liquid noble gases, Xenon was the most promising for the already described light yield and for the fast response time. The decay time of liquid noble gases is largely affected by the energy deposit mechanism; in case of highly ionising particle in Xenon the excited atoms of the material form a Xe_2^* dimer before de-exciting to two separated atoms:

$$\mathrm{Xe} + \mathrm{energy} \rightarrow \mathrm{Xe}^* \tag{2.1}$$

$$\mathrm{Xe}^* + \mathrm{Xe} \rightarrow \mathrm{Xe}_2^* \rightarrow 2\ \mathrm{Xe} + \mathrm{photon} \tag{2.2}$$

The Xe_2^* dimer can be produced either in a triplet or in a singlet state: the singlet one being the most probable for highly ionising particle interactions. Both states decay but with different life times (reported in table 2.1) which may differ by several orders of magnitude as for the Argon case.

As described in Bethe theory[20], minimum ionising particles with relativistic factor $\beta\gamma > 3.5$ are fast enough with respect to atomic time scales to cause the energy loss being dominated by scattering on free electrons composing the material, so their energy deposit highly localised into material electron with few cases going so high to produce δ rays. The same is true, for other reasons, for photons of in the 10 MeV − 100 MeV energy range: they initially interact in liquid xenon by pair production followed by Bremmstrahlung, just like electromagnetic showers, then the high critical energy of Xenon (10 MeV) quickly quenches the shower with photons undergoing Compton scattering and electron depositing their energy by the already mentioned ionisation. The combination of the mentioned physics processes results in cases where the atomic electron energy to be so high to directly ionise the atom instead to excite it.

$$\mathrm{Xe} + \mathrm{energy} \rightarrow \mathrm{Xe}^+ + \mathrm{e}^- \tag{2.3}$$

$$\mathrm{Xe}^+ + \mathrm{Xe} \rightarrow \mathrm{Xe}_2^+ \tag{2.4}$$

$$\mathrm{Xe}_2^+ + \mathrm{e}^- \rightarrow \mathrm{Xe}_2^{**} \rightarrow \mathrm{Xe}_2^* + \mathrm{heat} \tag{2.5}$$

$$\mathrm{Xe}_2^* \rightarrow 2\ \mathrm{Xe} + \mathrm{photon} \tag{2.6}$$

Ionisation causes the formation of a Xe^+ ions and then a Xe_2^+ dimer that can only decay to the ground state if an electron is recombined to neutralise the charge. The neutral dimer Xe_2^{**} is then formed in a highly excited state that takes time to decay back to the same Xe_2^* but with a different mixture of triplet and singlet states.

The ultimate observable effect is a different waveform shape for highly ionising or minimum ionising particles and photons (figure 2.5) which proves useful for identification of events in a particular calibration of the MEG II Liquid Xenon detector (described in section 2.4.2.1).

2.2.2 The structure of Liquid Xenon detector

The structure of the Liquid Xenon detector is tailored to the needs of the MEG II experiment[63]: with 900 litres volume it can contains roughly 2.6 tonnes of liquid xenon, with a temperature between the 161.4 K melting and the 165 K freezing points. The

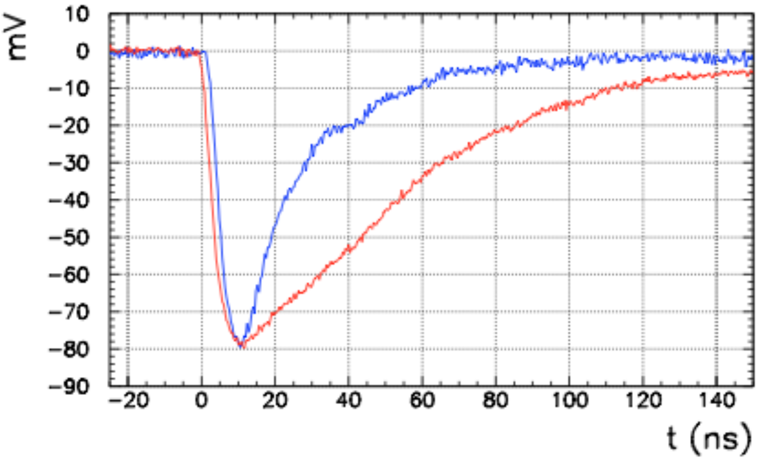

Figure 2.5: Example of a signal induced by α (blue) and γ (red) events, pulse scaled arbitrarily to have the same peak amplitude.

liquid is contained in a C-shaped cryostat (figure 2.6) surrounding 1/3 of the aximutal angle around the target for a total of 11% of the solid angle from the target, defining the signal photon acceptance and, opposite to that, the expected emission range for the corresponding signal positron.

The photon entrance window is composed of an external 0.4 mm thick steel plate enclosing the vacuum volume used for thermal insulation and another 0.4 mm steel inner wall reinforced with a carbon fibre/aluminium honeycomb, to keep the amount of material between the target and the internal sensible volume at a minimum value. The inner volume is filled with liquid xenon with the scintillation light detectors being submerged in it (figure 2.7a), 4092 Hamamatsu S10943-4372 MPPC “Multi Pixel Photon Counter”[64] (a silicon photomultiplier) cover in the inner face and 668 Hamamatsu R9869 photomultipliers are spread over the remaining five faces.

The devices used, showed in figure 2.7b, are sensitive to the VUV light (around 180 nm) emitted by xenon scintillation, MPPCs have several advantages over photomultipliers (PMTs): the inner face higher segmentation, passing from the 2” PMT diameter to a squared 15×15 mm^2 packaging, will increase the spatial resolution of the photon conversion point and help to identify γ pileup events. Furthermore, the squared shape allows for a tighter packing of the sensors, reducing the amount of dead material, a limiting factor for the MEG I resolution on shallow events. MPPCs also have no gain reduction due to the spectrometer magnetic field however, since they have an overall bigger capacitance, their waveform show a slower rising edge and a longer decay time; therefore the single channel timing might be reduced but can be compensated by averaging a bigger number of channels.

The Silicon photomultipliers (SiPMs) detects photons by an array of geiger-mode diode pixels (for MEG II Hamamatsu MPPCs, 50 μm size) each connected in parallel through a dedicated quenching resistor as shown in figure 2.8a, hence the large capacitance of

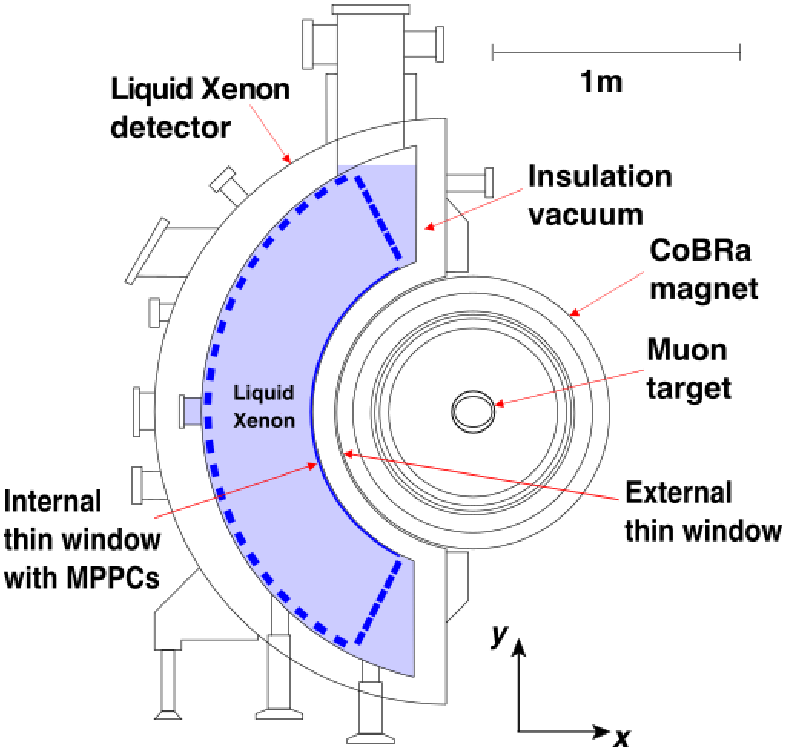

Figure 2.6: Liquid Xenon detector structure and positioning.

the device. To limit this issue the 12×12 mm^2 sensitive area is composed of four independent matrices with 6×6 mm^2 size, clearly visible in figure 2.7b. The four section composing the MPPC are connected to the readout electronic with a mesh of resistors and capacitors as shown in figure 2.8b, that have been chosen to limit the amount of bias and readout channels to one each MPPC: for the constant bias voltage each section is in parallel allowing a single ~ 60 V source to power the full MPPC. On the other hand, fast signals from photons are on multiple sections are summed up by the capacitors resulting in a similar pulse height to series connected sections.

2.2.3 Photon detection performances

The performances of the upgraded Liquid Xenon detector are evaluated using a Montecarlo simulation that takes into account the full γ shower generation, UV photon propagation to the sensor and waveform simulation.

The energy resolution is found to be around 1% and not to depend much on how deep the conversion point turns out to be. This was not the case of MEG I due to the different arrangement of photon sensors on the inner face, described in more detail in section 2.5.1.

The quoted resolution assumes an additional worsening of 0.7% from MC predictions which was needed, in the previous MEG I experiment, to account for some unknown

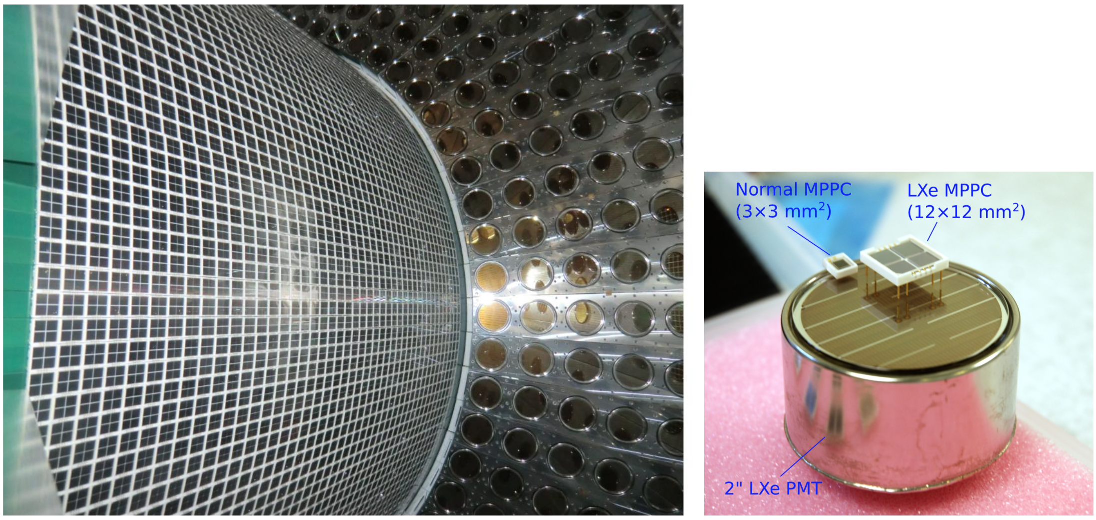

(a) Inner view of the Liquid Xenon detector. (b) MPPC and PMT

Figure 2.7: Light sensors in Liquid Xenon detector.

mismodelling of the signal response.

Timing resolution was found to be in the 50 to 70 ps range, assuming, in the waveform simulation, a gaussian noise with a sigma of 1 mV, which still needs to be checked in the actual hardware.

2.2.4 MPPC position measurement survey

The Liquid Xenon detector spatial resolutions are expected to be $\sim$ 1 mm, requiring the knowledge of MPPC positions at the order of $\sim 100\ \mu$m when the detector is cooled down to liquid xenon temperatures and then not accessible to any direct survey.

A dedicated setup was designed exploiting X-Rays from a Cobalt ^{57}Co source. 122 keV and 136 keV photons have a conversion depth of few millimetres in liquid xenon and produce a very localised energy deposit. A commercial point-like source of $3 \cdot 10^{10}$ Bq is placed on a movable stage at the centre of the CoBRa magnet with a very narrow collimator and can be moved along the beamline (experiment Z axis) and tilted to point to any point of the detector inner face. The collimator is such that the X-Rays interact only in a 0.5 mm × 33 mm area on the Liquid Xenon inner face and can be scanned across the surface, in two orientations, to identify both edges of the SiPMs.

Two data taking campaigns were carried out, the first one in 2017, involving only channel scalers: the source was moved to a known position and single channel rates were recorded measuring X-Rays over a continuous background mainly due to cosmic rays. This proved successful with an average X-Ray rate of $\sim$ 50 Hz on top of a $\sim$ 40 Hz background (figure 2.9).

The following year a background-free measurement was carried out with the goal of improving the resolutions and crosscheck previous year measurements after a second

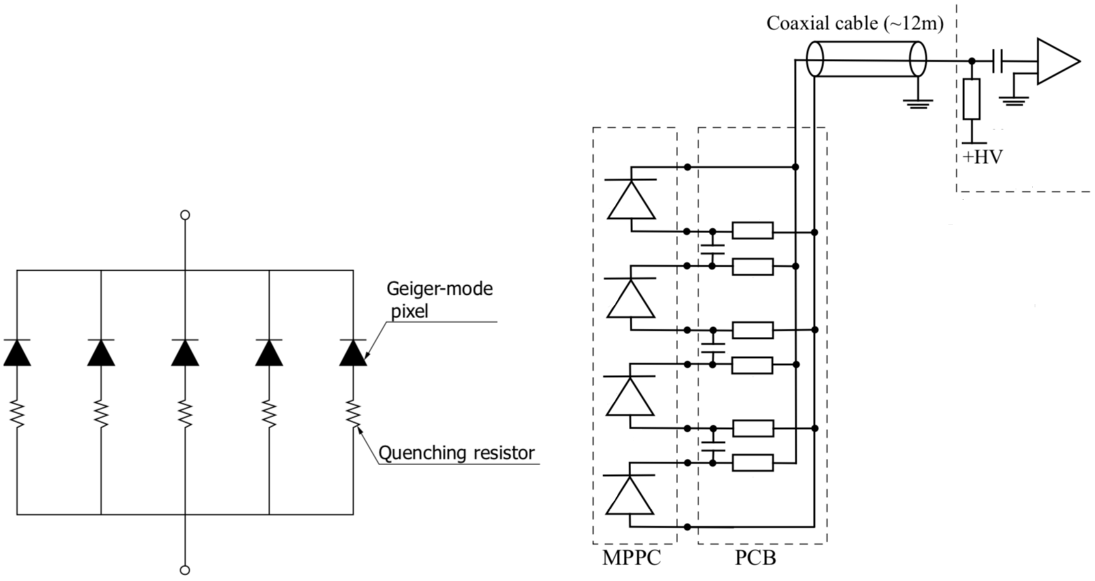

(a) Internal sipm schematic (adapted from [65]) (b) MPPC section connection to readout electronics

Figure 2.8: Internal MPPC schematic and connection schema for the Liquid Xenon readout.

thermal cycle of the cryostat. The removal of the background is possible in the offline by recording the waveforms of the MPPC nearby where the X-Ray beam was pointed and selecting only events with an energy deposit localised in few channels, that was paired with an online selection of events: the signals from the few channel of interest were summed up and subtracted by the same amount of channels placed somewhere else on the inner face, this provided a removal of any coherent component of the noise and, at the same time, the amplitude of the the waveform provided a value to reject of the background. The offline cut selected 60% of the online selected sample, the biggest online contamination source were very high energy cosmic rays that produced a similar energy deposit on the full inner face and the resulting signal was partially removed by the online noise subtraction.

A signal/noise ratio larger than 5 was achieved even if the same cobalt source was employed with an activity which had decreased by more than a factor 4 since the previous run.

The red dashed lines in figure 2.9 mark the edges of the eight carbon fibre plates, attached to the cryostat inner wall, where the MPPC are positioned. Figure 2.10 shows the difference in the measured and expected position, including the expected shrinking due to temperature effect. The different behaviour of the eight carbon fiber plates is evident: coherent displacements up to few millimetres indicate independent translation and rotation of the different parts. Only MPPCs having $|z| < 120$ cm could be scanned due to the presence of the spectrometer support structure.

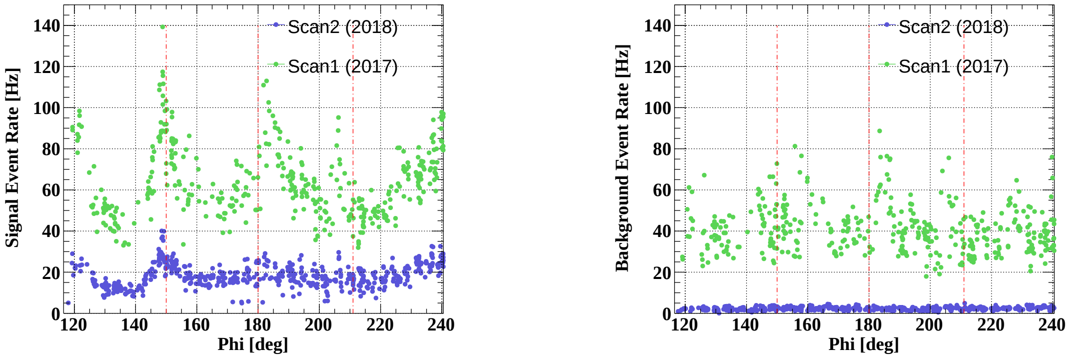

(a) Signal (b) Background

Figure 2.9: X-Rays signal and background rates in the two scans. Red lines mark boundaries of the MPPC support structure.

2.3 Positron Spectrometer

The signal positron, which would be emitted in the direction opposite to the Liquid Xenon detector, would be bent by the magnetic field of the Constant Bending Radius (CoBRa) magnet in such a way to have trajectories weekly dependent on the emission angle (see figure 2.11). The magnetic field gradient produces a further quick sweeping out of positrons emitted perpendicularly to the main field reducing the particle flux to $\sim$ 30 kHz/cm that is a manageable amount for the newly designed single volume Cylindrical Drift Chamber (CDCH).

The positron track timing is measured by the new pixelated Timing Counter (pTC), an array of plastic scintillators placed outside of the drift chamber volume, highly segmented to cope with the huge particle flux. The high number of independent drift cells and scintillator tiles of the pTC, helps disentangling the pileup and, at the same time, increases the resolution by providing multiple independent measurements of the positron position and timing.

2.3.1 The CoBRa Magnet

The CoBRa Magnet[66] is a thin wall superconducting magnet enclosing the positron spectrometer while allowing photons to pass through its 0.197 X_0 radiation lengths to reach the Liquid Xenon detector. To achieve the almost constant bending radius independent of the emission angle at the target, the magnetic field ranges from 1.27 T at the magnet centre down to 0.49 T at the end of the sensitive region, needing 5 different coils arranged at three different radii (as depicted in figure 2.11).

A set of two additional resistive coils is needed to reduce the stray field of the superconducting magnet in the liquid xenon volume to below $5 \cdot 10^{-3}$ T allowing a stable operation of PMTs. If compensation coils are disabled the Liquid Xenon PMT gain is

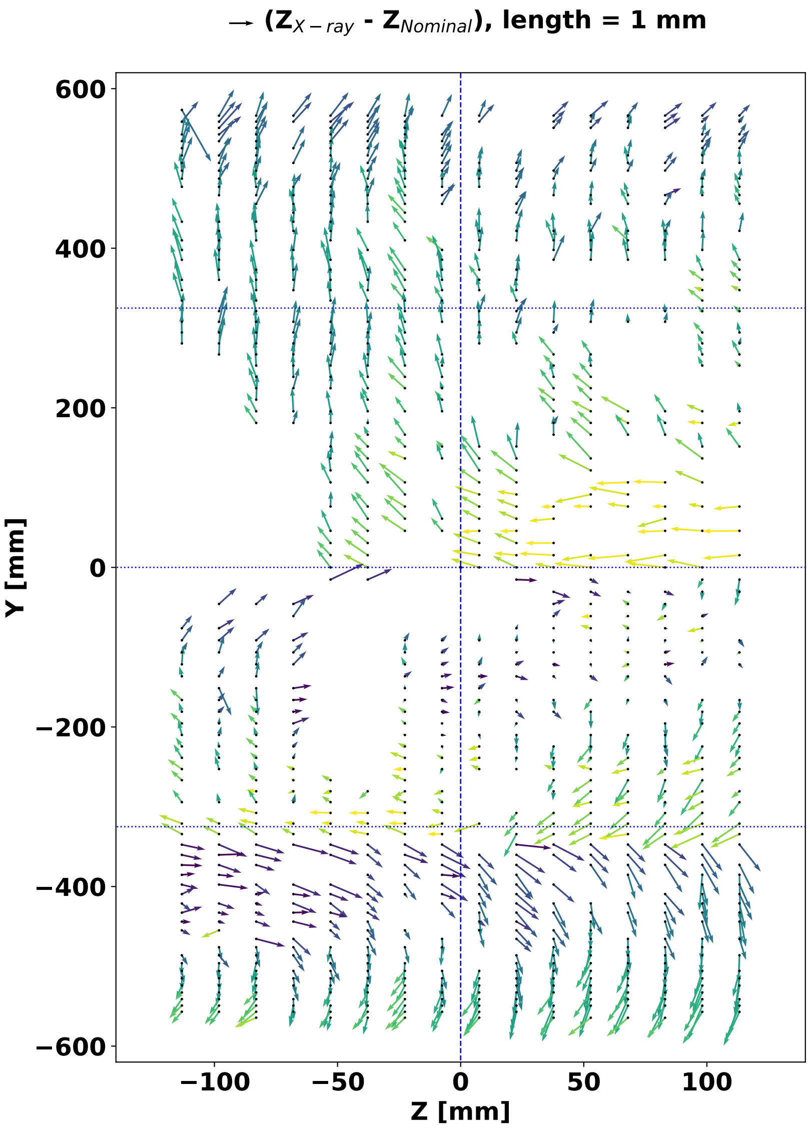

Figure 2.10: X-Ray reconstructed position offsets. The blue lines separate MPPCs attached to different support structures.

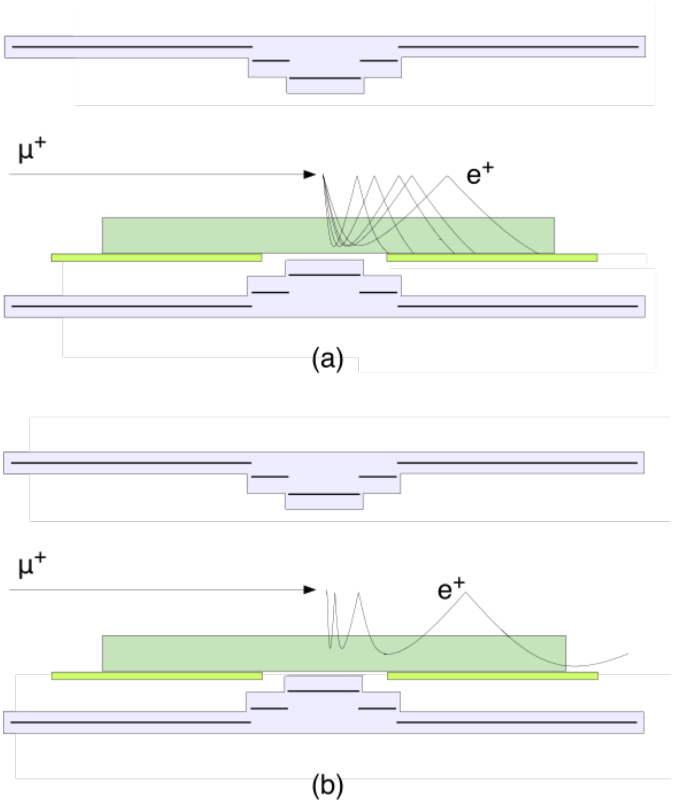

Figure 2.11: Effect of non-uniform CoBRa magnetic field on signal-like positron emitted in acceptance (a) and on smaller energy positron emitted perpendicularly to the beam.

reduced by factor 50 with respect to null magnetic field. Such a reduction becomes 5% when the compensation coils is switched on.

2.3.2 The Cylindrical Drift Chamber

The new Cylindrical Drift Chamber[67] is the tracking detector of the spectrometer, and covers the cylindrical volume from 17 cm to 30 cm radius for a 1.98 m length around the target. It measures positron trajectory with the associated momentum of the fitted track and, extrapolating the track to the target plane, the emission angles as well as the identification of muon decay point.

Positron tracking at very low momenta is made extremely difficult by the not negligible possibility of scattering scattering angle; if multiple scattering dominates, it can be described by a gaussian with σ given by:

$$\sigma = \frac{13.6\ \mathrm{MeV}}{\beta c p} z \sqrt{\frac{x}{X_0}} \tag{2.7}$$

where β c is particle velocity and p is particle momentum, z is the charge and x is the

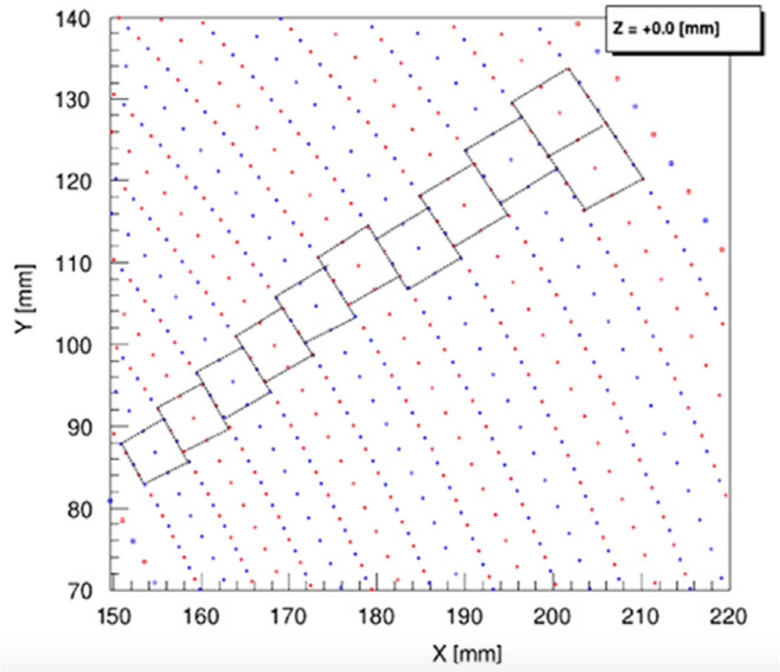

Figure 2.12: Position of CDCH wires at Z=0, red and blue marker highlight wires on different stereo angles.

distance travelled in the material having X_0 as radiation length. The dependency of 2.7 from the inverse of the momentum causes low momentum particles to be very susceptible to scattering limiting the reconstruction capabilities, and needing very transparent detectors, in terms of low amount of material having a very long X_0.

The very limited material budget is the guiding principle of the design of the MEG II drift chamber: the helium target volume is isolated by a 20 μm aluminised Kapton foil from the tracking volume filled with a very lightweight helium-isobutane gas mixture. The 90%helium 10% isobutane mixture has been chosen because of the very small X_0 of ~ 1300 m, with the associated drawback of reducing the primary ionisations to ~ 13 ionisation cluster/cm.

However equation 2.7 holds only in the limit of an average of small angle scattering, the CDCH radiation length for a complete positron track turn is $1.45 \cdot 10^{-3}$ X_0, and it makes unlikely, but still possible, that in correspondence of higher concentration of material, like the Mylar foil or the wires inside the chamber, a large angle scattering might occur, worsening the tracking and producing non gaussian tails in the quantities being reconstructed.

2.3.2.1 Wire structure

Electrons generated by ionising particles drift in cells with a geometry defined by the wire disposition in the active volume, the result are almost squared drift cells ranging from

6 to 8 mm size depending on the radial distance from the axis. Each cell is composed of a 20 μm gold plated tungsten anode, placed at the centre, surrounded by eight silver plated aluminium cathodes of 40 or 50 μm diameters.

Wires are placed in the CDCH in parallel sets of three layers starting with a layer of 40 μm cathodes followed by a layer of anodes/50 μm cathodes and another layer of 40 μm cathode closing the cell as visible in figure 2.12. Each set of layers is placed with alternating stereo angles of approximately $\pm 8^\circ$ with respect to the previous set to reconstruct the longitudinal position of the hits. Since no space is left between the two 40 μm layers of two consecutive sets, this structure results in meshes of crossing 40 μm cathodes, propagating in two stereo directions, between consecutive anode/50 μm cathodes layers. This mesh provides a very stable boundary condition for the electric field.

The full chamber is finally made of a stack of 9 cell sets for a total of 1728 signal wires. Two additional layers made of crossing 50 μm wires, at a reduced potential, are used to equalise the electric field of the innermost and outermost cells.

The elastic limit of the 40 μm wires is $\sim$ 8 mm with respect to the null tension length, so cathode wires are tensioned at 75% of their electrostatic limit to provide enough mechanical tension to balance the electrostatic attraction between anode and cathodes at the operational voltage of $\sim$ 1500 V. In this condition the total pull force on the supporting endplates, where wires are attached, is of the order of 300 kg to be added to the force applied by the internal kapton foil. This axial load is sustained by a 3 mm carbon fibre shell which separates the volume of the chamber from the exterior and keeps the endplates in position, providing also an external shielding from noise sources having a 50 μm aluminium foil attached to its inner surface. When the load is applied a shrinking of only 20 μm is observed on the full drift chamber length, well within requirements.

2.3.2.2 Wiring and Assembly procedures

Featuring a wire density at the endplate of $\sim$ 6 wires per cm^2, there is no possibility to use standard crimping procedures to attach the wires to the endplate, a PCB-based feedthrough was used[68]: the endplate, shown in figure 2.13a, was made of 12 aluminium spokes defining the same number of drift chamber sectors where 20 different kinds of PCBs could be piled up by means of plastic PEEK spacers.

Wires were soldered to two PCBs (figure 2.13b) outside of the chamber using the automatic system, it took care of the positioning, soldering and handling of the wired PCBs. Wire positioning was achieved by wounding the wire around a cylinder whose length matched the final CDCH length at the design tension. Subsequently a soldering laser fixed the wires into position by using a soldering tin and, finally, the wired PCBs were unrolled by means of automatic suction cups.

The stacking of PCBs (figure 2.14) into the chamber is made in order from the innermost layer to the outermost one by taking care of the stereo geometry that causes half of the PCBs, and the associated wires, to go from a given sector to two sectors clockwise and the other half to two sectors counterclockwise.

When wiring was completed, the inner foil and the external carbon fibre structure

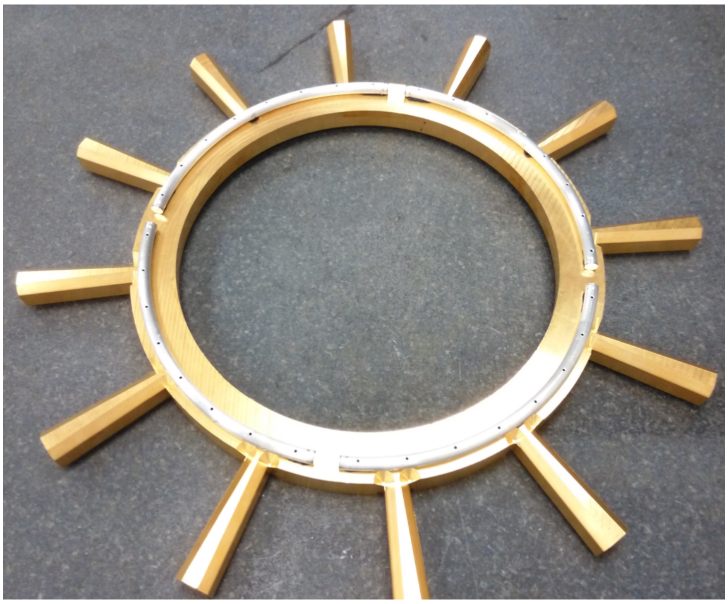

(a) Endplate with 12 spokes

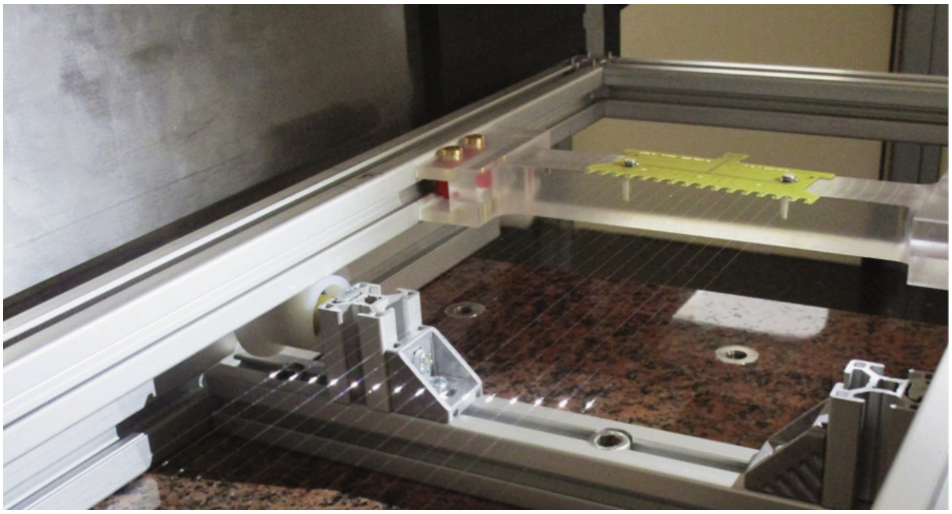

(b) Wires attached to PCBs

Figure 2.13: Wiring of CDCH.

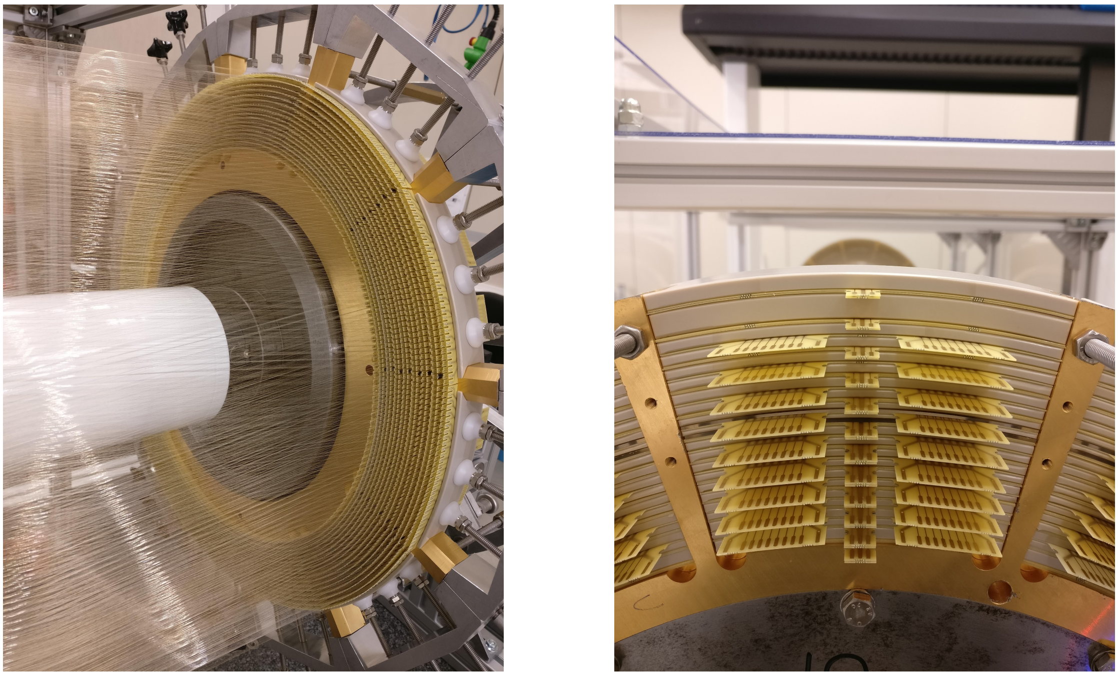

(a) PCBs pile during assembly

(b) Final PCB stack

Figure 2.14: CDCH assembly starting from wired PCBs.

were installed and the endplate stacks of PCBs were sealed by a vacuum-grade glue to assure air-tightness of the wire volume.

2.3.2.3 Signal readout path

Each anodic PCB is connected to two frontend circuits[69] by using a custom made connector (figure 2.15), so that the High Voltage can be applied and the signal can be amplified and sent to readout electronics. The 16 drift cells belonging to the same wire PCB are connected together by 10 MΩ resistors on the frontend circuits so that a single High Voltage channel from an ISEG EHS-f230p source[70] can supply the tension needed to generate the drift field.

Each anode signal is AC-coupled to an amplification circuit with a gain of 20 dB made by a two cascaded gain stages to achieve a bandwidth of $\sim$ 900 MHz after a 5 m long shielded parallel pair readout cable. Such a high bandwidth was chosen because of the very small amount of primary ionisation clusters of the gas mixture. Clusters are separated by roughly 80 μm biasing the impact parameter estimation by a similar amount if only the first cluster is used. The high bandwidth used gives the opportunity to reconstruct time, and therefore distance from the wire, for few different clusters and, in average, correct for the bias on the impact parameter[71].

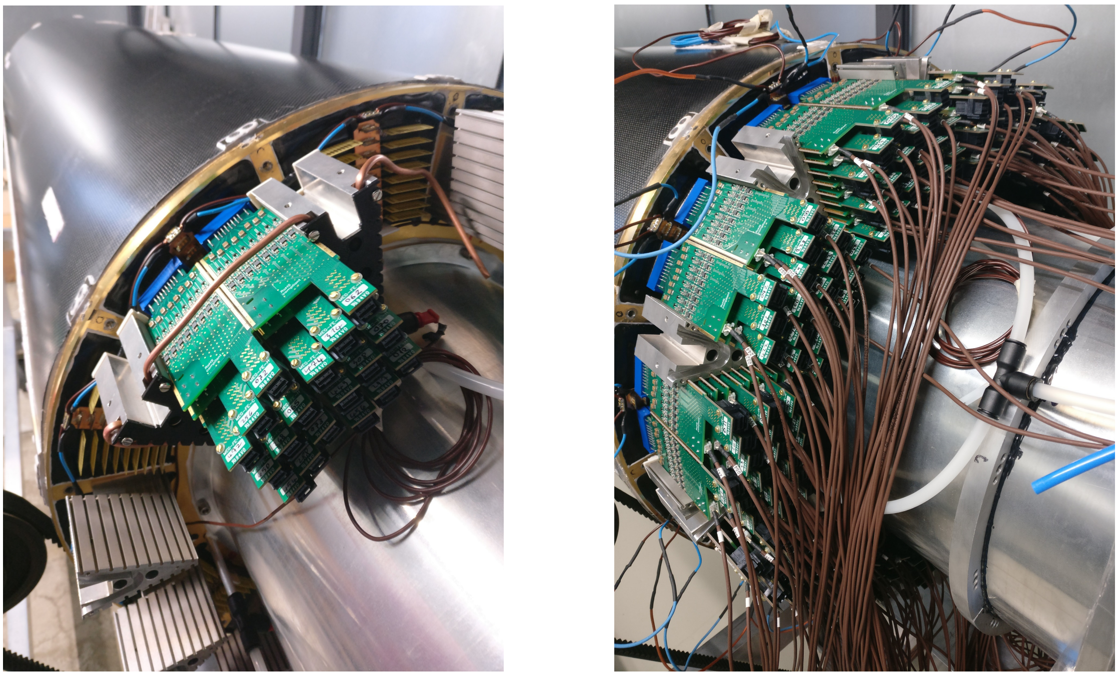
(a) Frontend circuit plugged into the CDCH

(b) Full set of Frontend circuits wit HV cables

Figure 2.15: CDCH Frontend circuits.

2.3.3 The Timing Counter

A set of 512 plastic scintillator tiles[72]is used to provide a very good timing of the track but also to identify the positrons for trigger purposes and to provide an external seeding for the CDCH track reconstruction. To achieve this Tile counters are placed in the region between the CDCH outer carbon fibre and the CoBRa magnet, far from the central region, where signal-like positrons emitted in the solid angle opposite to the Liquid Xenon are expected to end up.

Tiles showed in figure 2.16 are all 120 mm wide and 5 mm tick but they come in two heights of 40 mm and 50 mm to provide a compromise between bigger acceptance of the wider counter with a better resolution of the smaller ones.

Each Saint Gobain BC422 Tile is wrapped in Enhanced Specular Reflector (ESR) foil and then in a 25 μm Tedlar black sheet for light shielding; on the smaller faces six or seven AdvanSid ASD-NUV3S-P SiPMs, depending on the height of the Tile, collect scintillation photons and are connected in series so that a single MCX connector at the end of each tile can pass the signal to a common PCB backplane which is also the mechanical element holding the tile in position.

The full Timing Counter is made by two identical modules (whose drawing is reported in figure 2.17) with one located upstream and downstream from the target. The Tiles inside each of the two modules are closely packed and oriented to be perpendicular to the

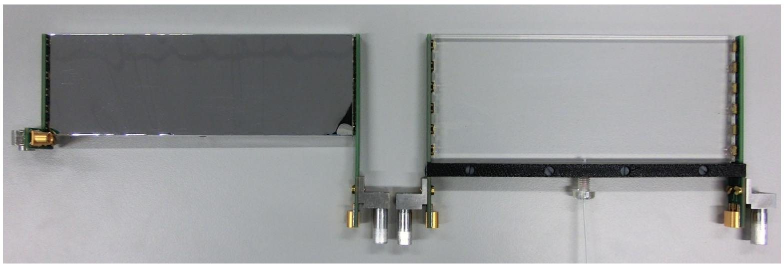

Figure 2.16: Two Timing Counter Tiles, 50 mm one on the on the right and 40 mm one on the left.

average 52 MeV positron track emitted in acceptance. Each module is supported by an aluminium structure which houses 33 PCBs with the MCX connectors, Innermost Tiles cannot connect directly with those PCBs because of the reduced space due to CoBRa inner coil, therefore these counters (an example of them is the left one in the figure 2.16) have a different shaped SiPM support PCB to anchor them directly onto the aluminium support structure.

The full aluminium block is cooled down to around 10 degrees and a constant flux of cooled dry air is applied so to fix the SiPM temperature and limit the variation of SiPM parameters that typically are highly susceptible to heating.

2.3.3.1 Laser system

The goal of monitoring a timing detector with 30 ps resolution requires a very precise and fast calibration system.

This is accomplished by a laser signal coming from a picosecond pulsed diode[73] source which is distributed (see figure 2.18a) to all Tiles through an optical switch and a set of optical splitter so that roughly a 300 photoelectrons signal is readout by any illuminated channel. The laser wavelength of 401 nm was chosen to match, as much as possible, the 370 nm emission of the scintillator and the peak detection efficiency of the SiPMs which is 420 nm. Each component of the optical path has been individually characterised both in optical length and in output power ratio so that, not only the system can be used for stability monitoring between two calibrations based on tracking, but also to measure the absolute time offset by subtracting the contributions of the Laser system.

A side product of the Laser system commissioning in 2017 data-taking was to monitor the stability (plot in figure 2.18b), not only of the Timing Counter but also of the full TDAQ system that, combined, gave an RMS fluctuation of ~ 8 ps for 60 counters over a 1 month period.

2.3.4 Expected performances of the spectrometer

Prototypes[74, 75] with the same drift cell size of the final CDCH were to validate the Geant4 simulation of the CDCH detector while the full timing counter was successfully

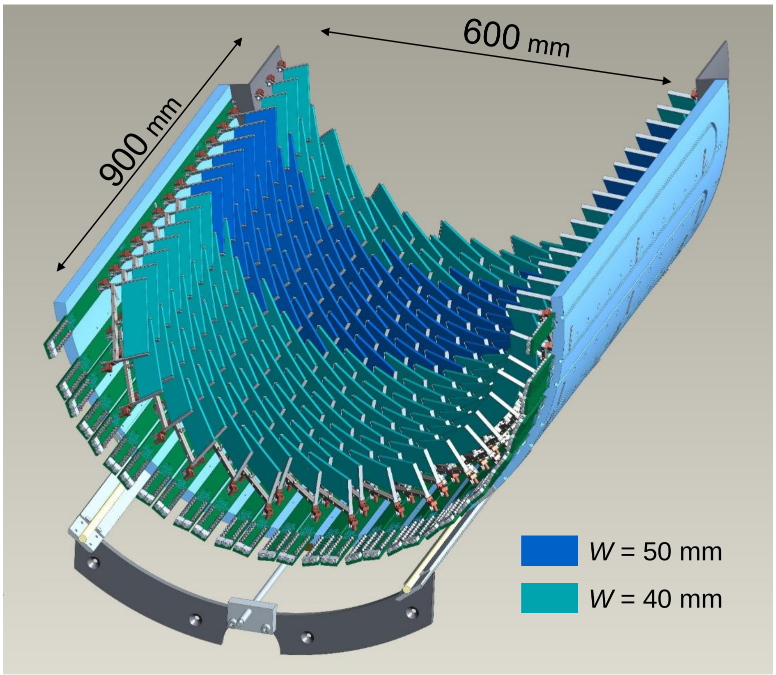

Figure 2.17: Drawing of a Timing Counter Module.

tested in its final position inside the MEG II apparatus achieving the goal of 30 ps[76]. The CDCH was installed on beam for the first time in 2018 run but could not operate due to the lack of readout channels and to a wire breaking inside the tracking volume that caused only a subsample of the drift cells to be operated.

The wire breaking turned out to be due to a corrosion problem and not to a mechanical one. The CDCH was subsequently by roughly 0.5 mm above the final length to remove wires that might be weakened by the corrosion problem. Several wires broke during this operation but no other wire failure was observed since August 2019.

The CDCH was reinstalled inside CoBRa and cosmic rays/positron signals were recorded during the 2019 data-taking with a small number of cells (96 out of 1728 cells) due to the limited number number of readout channels.

The combined performance of the spectrometer could only be studied by means of simulations. The core resolutions obtained are reported in table 2.2, MEG I performances are also reported for comparison.

	E_{e^+}	θ_{e^+}	ϕ_{e^+}
observed on MC	77 keV	5.7 mrad	6.0 mrad
expected based on prototypes	130 keV	5.3 mrad	3.7 mrad
MEG I	306 keV	9.4 mrad	8.7 mrad

Table 2.2: Spectrometer MC resolution vs expected by smearing of the MC Hit.

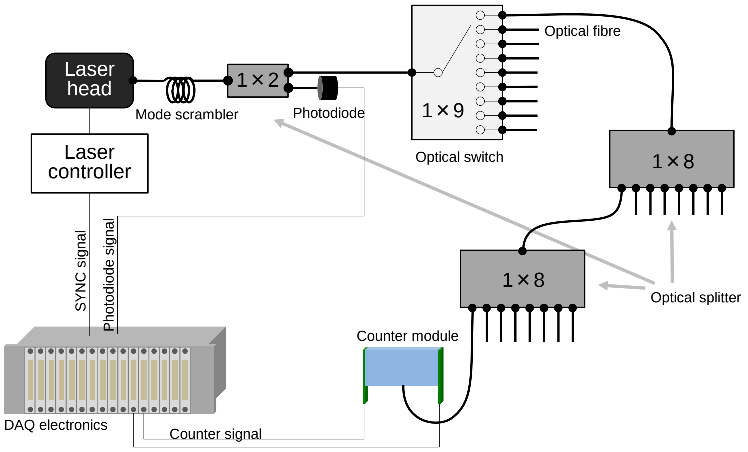

(a) Schematic drawing of the laser system.

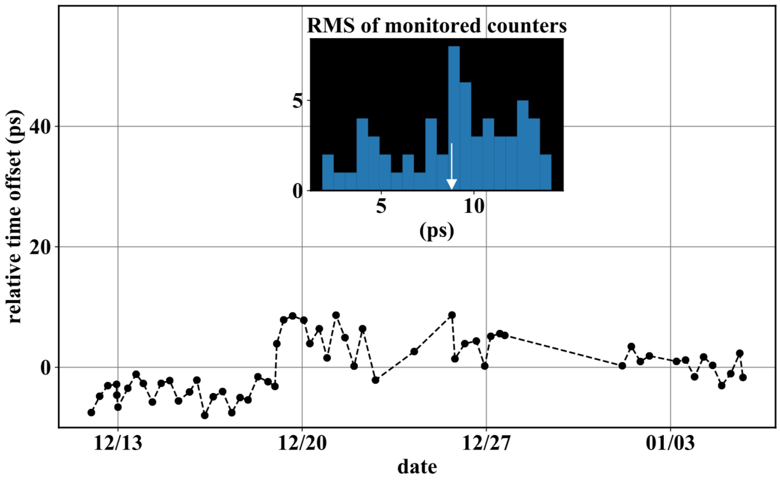

(b) TC+TDAQ Stability over 2017 run.

Figure 2.18: TC Laser system.

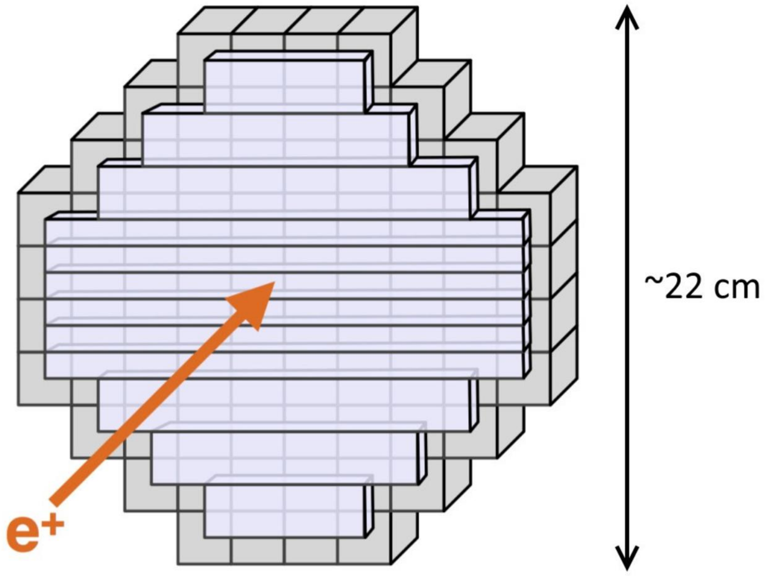

Figure 2.19: Sketch of the Radiative Decay Counter, the blue layer are the plastic detectors followed by the grey LYSO crystals.

2.4 Other Detectors and Calibration Items

In addition to the previously described detectors, a set of additional detectors and calibration items are installed to contribute the analysis, monitor the performances of the experiment or provide detector-specific tools. Most of these equipments are adopted from the MEG experiment while others, like the Radiative Decay Counter, are brand new.

2.4.1 The Radiative Decay Counter

There are two main sources of 52.8 MeV energy photons in the MEG experiment: radiative Michel decays can reach, at the endpoint, the signal energy, and photons from the in-flight annihilation or Bremsstrahlung of a Michel positron inside the CDCH volume. The latter one is very relevant because photons spectrum can extend above the signal photon energy due to the additional contribution coming from the positron mass.

Annihilation in flight and Bremmstrahlung events are reduced, with respect to MEG, by the much lighter spectrometer, while, on the other hand, the RMD contribution is dictated only by physics.

The Radiative Decay Counter (RDC), shown in figure 2.19, is placed along the beamline, downstream from the target, with the goal of tagging radiative Michel decays whose positron is emitted with small momentum (1-5 MeV) while the corresponding high energy photon (with energy > 48 MeV) enters the Liquid Xenon detector. The positron helix of low energy positrons in the CoBRa magnetic field is constrained inside the beam line volume and is not detected by the spectrometer. The RDC promises to tag such positrons with 40% efficiency providing an improvement of 15% on the final experimental sensitivity. Another RDC placed upstream would provide a higher tagging efficiency but the requirements of transparency and resistance to the muon beam passing through it poses a big challenge in term of detector design.

The downstream RDC can instead employ a ticker detector because the beam is not passing through it: positrons are detected by a layer of plastic scintillator followed by a calorimetric layer of $2 \times 2 \times 2$ cm^3 LYSO crystals, both layers being readout by SiPMs. Those two components are required to separate RMD events from standard Michel positrons that are emitted with a very small angle respect to the beam line. The energy deposit in LYSO is very different for the two events which can be therefore clearly separated, while the plastic scintillator provides a very precise timing for the coincidence with the photon in the Liquid Xenon detector.

2.4.2 Xenon Calibration sources

The 1% energy resolution goal for the Liquid Xenon detector is very challenging and requires to monitor its energy scale well below the 1% level. A set of calibration tools were developed for MEG I and adapted for the current experiment to provide an absolute definition of the energy and to monitor it continuously to correct for detector drifts.

2.4.2.1 LED and point-like α sources in Liquid Xenon detector

The gain of each Liquid Xenon PMT and MPPC is determined by using a group of LEDs[77] to illuminate each sensor with various, reproducible, light amplitudes and then, assuming that the photon detection follows poissonian statistic, for each sensor the mean μ and the variance σ at a given light amplitude are:

$$\mu = G \cdot n_{ph} \tag{2.8}$$

$$\sigma = G \cdot \sqrt{n_{ph}} \tag{2.9}$$

where n_{ph} is the number of detected photons hitting the sensor at such light level. So by a linear fitting σ^2 vs μ for various light amplitudes the gain can be estimated.

An array of 25 point-like ^{241}Am α sources is installed on five 100 μm wires[78] inside the liquid xenon volume (as shown in figure 2.20): having a range of only 40 μm in the dense liquid an α particle event is very well localised and, once the source that emitted the particle is identified, the amount of recorded photons collected by each photodetector can be compared with a Montecarlo estimates to extract the detection efficiency, knowing the gain from the LED measurement. Having multiple sources helps to illuminate all the channels equally and, by comparing the number of photoelectrons with the distance from the source, the xenon purity can be monitored by measuring its light attenuation coefficient.

The pulse shape of α particle events is also clearly recognisable by the shorter decay time, being produced by high ionising particle (see section 2.2.1), and the reconstructed position of the event is reconstructed in characteristic rings around the wire because the

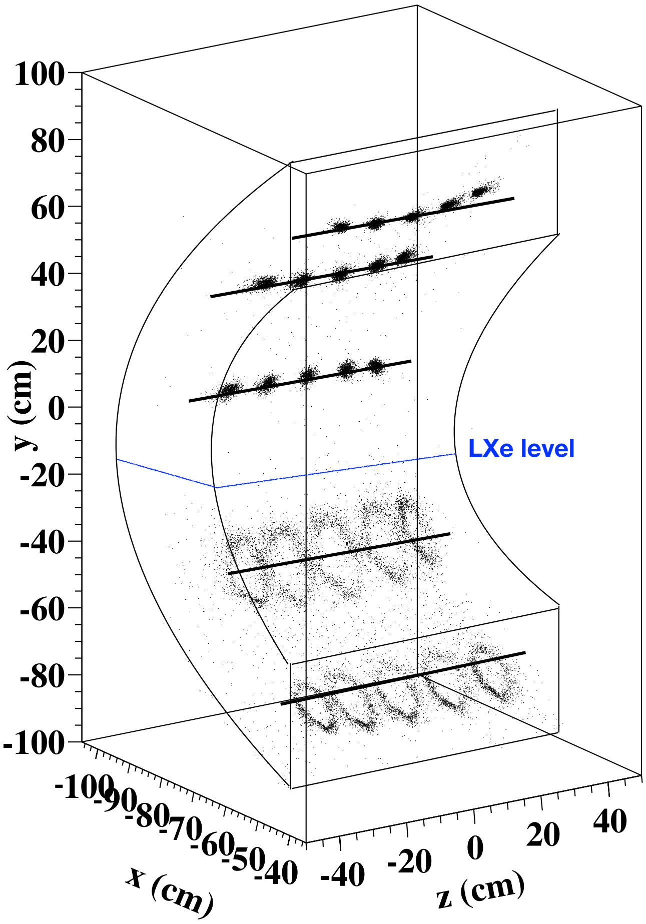

Figure 2.20: Alfa events topology in gaseous and liquid xenon inside a partially filled Liquid Xenon detector.

shadow projected by the same wire pushes the reconstructed position far from it. Alfa events in gaseous xenon have a much longer range ($\sim$ 8 mm) so the rings disappear.

2.4.2.2 AmBe and nickel gamma generators

Two low energy calibration points for the Liquid Xenon detector are obtained by using 50 kBq AmBe source providing 4.4 MeV γ rays and a 9.0 MeV photon line from neutron capture on a Nickel target, $^{58}\mathrm{Ni}(n, \gamma_{9.0})^{59}\mathrm{Ni}$. In the latter case the thermal neutrons producing the nickel reaction come from a pulsed deuterium-deuterium generator performing $\mathrm{d}(\mathrm{d},^3\mathrm{He})\mathrm{n}$ with an intensity of 2.5×10^4 n/pulse; such a generator can operate at a programmable rate (1-100 Hz) and provides a synchronisation signal. The synchronisation is sent to the trigger system set a $\sim$ 100 μs gate (the average capture time of thermal neutrons in nickel) after a delay of $\sim$ 15μs (needed for the neutron to moderate from 2.5 MeV to thermal energy). The gated operation mode enables the use of the

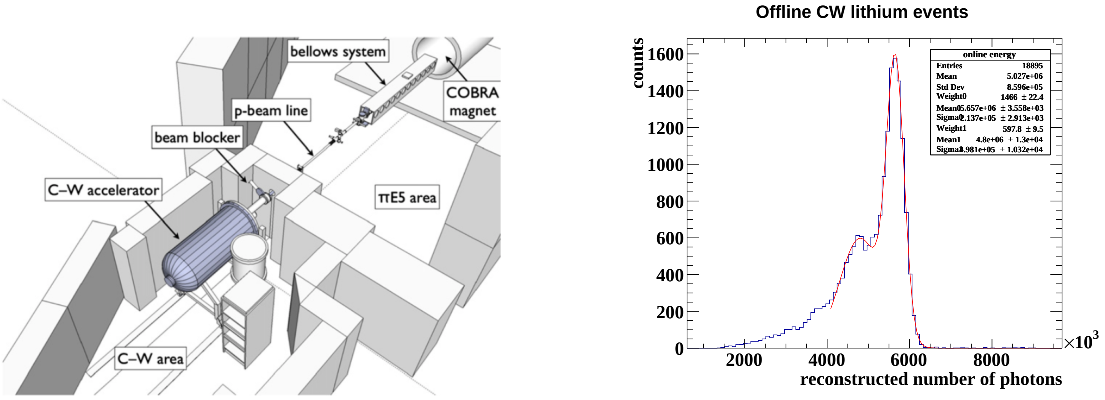

(a) Drawing of Cockroft-Walton proton beam-line.

(b) Lithium-generated gamma lines observed with the Liquid Xenon detector.

Figure 2.21: Cockroft-Walton Liquid Xenon calibration.

Nickel γ generator also during beam when the AmBe is masked by the beam-induced background.

2.4.3 Cockroft-Walton Accelerator for Lithium/Boron Calibration lines

A proton beam[79] is used to produce γ lines from nuclear reaction in a lithium thetraborate target ($Li_2B_4O_7$); a 17.6 MeV monochromatic line from $^7\mathrm{Li}(p,\gamma_{17.6})^8\mathrm{Be}$ and a pair of monochromatic lines $^{11}\mathrm{B}(p,\gamma_{4.4}\gamma_{11.6})^{12}\mathrm{C}$ with 4.4 and 11.6 MeV are produced.

The source of protons is a dedicated Cockroft-Walton accelerator (sketched in figure 2.21a) that can provide a proton beam at energies up to 1100 keV with currents up to 110 μA; the beam propagates in the direction opposite to muon beam (from downstream to upstream) inside a vacuum beam pipe, until it reaches its target. The standard target can be removed by a pneumatic system and a movable bellow can be used to have the $Li_2B_4O_7$ target positioned at the centre of the CoBRa magnet at. The insertion-extraction procedure requires around ten minutes and, during physics data-taking, the Cockroft-Walton will provide a daily measurement of the energy resolution.

While the 17.6 MeV line, shown in figure 2.21b, is used to monitor the detector response for its higher energy and production cross section (5 mb), the boron lines, produced by two successive de-excitement of the carbon atom, provide a clear event signature to tune the relative timing of the Liquid Xenon and pixelated Timing Counter detectors. A dedicated trigger on the coincidence of the two detectors is required to select this category of events otherwise the higher cross section of the lithium line overwhelms the boron ones, having only a cross section of 0.2 mb.

2.4.4 BGO Detector and the Charge Exchange reaction

Extrapolating the energy scale from 17.6 MeV to 52.8MeV can not provide the required precision, therefore a calibration line at energy near the signal is needed. The $\pi^0 \to \gamma\gamma$ decay can be used for this purpose. In the π^0 rest frame, the two photons are emitted with an energy of:

$$E_\gamma^{CM} = \frac{m_{\pi^0}}{2} \simeq 67.5 \text{ MeV} \tag{2.10}$$

π^0 can be produced through a Charge Exchange reaction of negative pions π^- on a liquid hydrogen target, producing a π^0 and a neutron. In this case the neutral pion is produced with a momentum of $\sim$ 28 MeV/c, boosting the decay products depending on their emission angle with respect to the flying direction of the pion. The minimum and maximum endpoints of the resulting γ spectrum in the laboratory frame are:

$$E_\gamma^{min} = \frac{m_{\pi^0}}{2}\sqrt{\frac{1-\beta}{1+\beta}} \simeq 54.9 \text{ MeV} \tag{2.11}$$

$$E_\gamma^{max} = \frac{m_{\pi^0}}{2}\sqrt{\frac{1+\beta}{1-\beta}} \simeq 82.9 \text{ MeV} \tag{2.12}$$

corresponding to events with both photons emitted along the direction of flight of the pion, in this case the opening angle in the laboratory is 180°.

During data-taking with the Charge Exchange reaction the target system is exchanged with a thin container were liquid hydrogen is liquified and an auxiliary BGO calorimeter is placed looking at the inner face of the Liquid Xenon detector from the opposite side of the target. Selecting events with a wide opening angle, it is possible to obtain a clean sample of photons at 54.9 MeV.

The BGO calorimeter consists in an array of 4×4 BGO crystals with a pair of plastic scintillators in front of it as a preshower. For photons converting in the preshower a very good time resolution is measured at the expense of a worse energy resolution, by contrast if there is no hit in the preshower the BGO provides a very good energy measurement but the timing information is worse.

The first MEG II charge exchange data-taking is foreseen for the end of 2019 to provide a first measurement of the capabilities of the updated Liquid Xenon detector.

2.4.5 Mott positron scattering

An alternative tuning of the beamline elements, including the separator, can select a beam of positrons which come from the Dalitz decay $\pi^0 \to \mathrm{e}^+\mathrm{e}^-\gamma$ of neutral pions generated in the target E together with charged ones.

The Dalitz spectrum of the positron covers all the energies up to $\approx m_\pi^0$. The $\pi E5$ beamline elements can be tuned to select a slice of the spectrum centred at approximately the signal energy $\sim$ 53 MeV with a momentum spread of $\simeq$ 250 keV and fluencies of $\sim 10^7\mathrm{e}^+/\mathrm{s}$.

The π^0 decay is an electromagnetic process with a time scale much faster than the charged pion decays, the original 40 MHz time structure of the proton beam is preserved and can be used to remove beam contamination.

When the positrons impinge on the MEG target a scattering process occurs with a distribution given by the Mott scattering[80]:

$$\frac{d\sigma}{d\Omega} = \frac{\alpha^2}{4p_{\mathrm{e}}^2 \sin^4\frac{\theta}{2}}(\cos^2\frac{\theta}{2} - \frac{q^2}{2m^2}\sin^2\frac{\theta}{2}) \tag{2.13}$$

where α is the fine structure constant, p_e is the momentum of the incoming positron, m is the target mass (tipically carbon/hydrogen in MEG II plastic target) q is the transferred four-momentum ($2m\Delta E_{\mathrm{e}}$) and θ is the scattering angle with respect to the beam direction.

Equation 2.13 predicts an average small average scattering angle. However, having a high intensity positron beam, there will be roughly $\approx$ 1 kHz rate in the spectrometer angular acceptance.

This process is completely uncorrelated from the physics being studied and can provide an alternative measurement of the spectrometer resolutions, acceptances and energy scale[81] that otherwise can only be extracted by a fit of the kinematic endpoint of Michel decays.

2.5 Improvements over the MEG experiment

All the major MEG I detector[82] have either been rebuilt from scratch or heavily modified to provide a big improvement on all the observables. Table 2.3 shows the comparison of MEG II versus the former experimental setup, with all observable resolution improving by roughly a factor 2.

	MEG I	**MEGII**	**Improvement**
E_{e^+} (keV)	380	130	2.9
θ_{e^+} (mrad)	9.4	5.3	1.8
ϕ_{e^+} (mrad)	8.7	3.7	2.7
E_γ $w_\gamma > 2;\mathrm{cm}(\%)$	1.7	1.1	1.5
E_γ $w_\gamma < 2;\mathrm{cm}(\%)$	2.4	1.0	2.4
u_γ (mm)	5	2.6	1.9
v_γ (mm)	5	2.2	2.3
$t_{\mathrm{e}^+\gamma}$ (ps)	122	84	1.45

Table 2.3: MEG I resolutions in comparisons of MEG II design values where u_γ and v_γ are photon conversion point and w_γ is the conversion depth. The improvement is calculate as the ratio of the two values.

Details of the improvements for each detector will be discussed in chapter 2.5.1 and 2.5.2. MEG I probability density functions can be modified with the resolution of table

2.3 to generate a set of toy Montecarlo background events. An upper limit on the BR of $\mu^+ \to e^+ + \gamma$ is determined on each toy by a maximum likelihood fit and the expected sensitivity is extracted as the median of the upper limits. Assuming a muon stopping rate of $7 \cdot 10^7$ μ^+/s a foreseen sensitivity of $6 \cdot 10^{-14}$ can be reached in 3 years (figure 2.22), considering a total of 140 data-taking days every year, to account for the accelerator shutdown periods and detector maintenance activities.

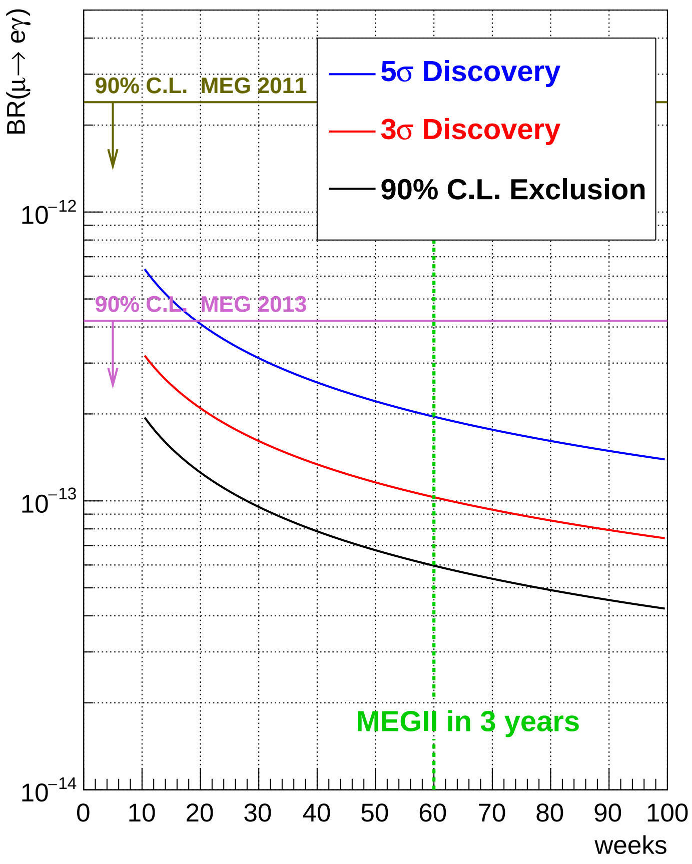

Figure 2.22: Evolution of expected MEG II sensitivity as a function of DAQ time.

2.5.1 Liquid Xenon

As visible from figure 2.23, the major feature of the upgraded Liquid Xenon detector is undoubtedly the MPPC-based inner face, however the lateral PMT holders were re-designed to extend the inner face by 10% beyond the nominal acceptance and to contain more PMTs being recycled from the inner face. In fact, the MEG experiment used the same PMTs on all the faces of the detector, including the inner face that contained 216 PMTs.

The extension of the inner face increases the lateral shower containment of γ at the edge of the acceptance and, together with the new disposition of PMTs, reduces the shower reconstruction uncertainties for that category of γ events.

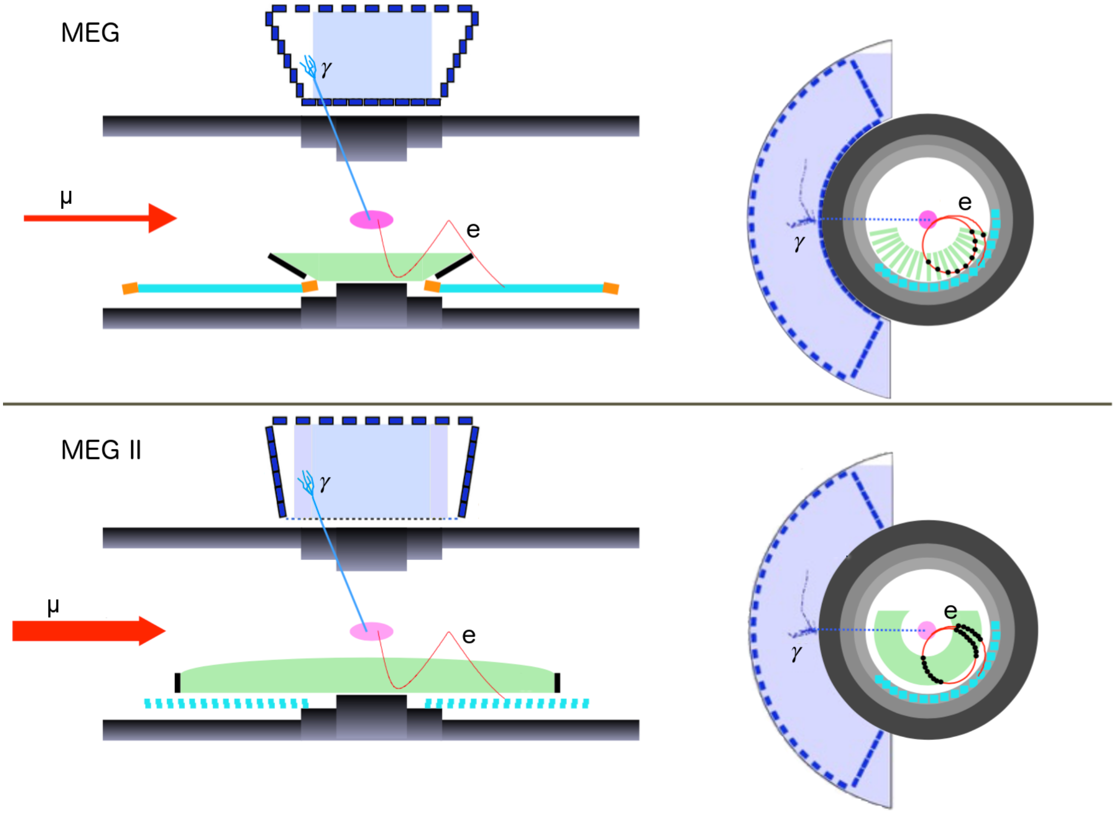

Figure 2.23: MEG and MEG II experiments side by side.

The higher detector segmentation on the inner face provides directly an increase in spatial resolution but also a higher fraction of sensitive area and a reduction of dead material, not only important to improve the resolutions but also to make them less sensitive to the conversion depth. In MEG the light collection of shallow events (conversion point depth < 2 cm) depended heavily on the impact position of the photon: if the shower originated in between the PMTs a big part of the light was lost resulting in a worse energy resolution (see E_γ table 2.3) with respect of of photons having conversion point depth > 2 cm. Such conversion point and depth dependance of the energy resolution is very much reduced, according to Montecarlo simulations, in MEG II.

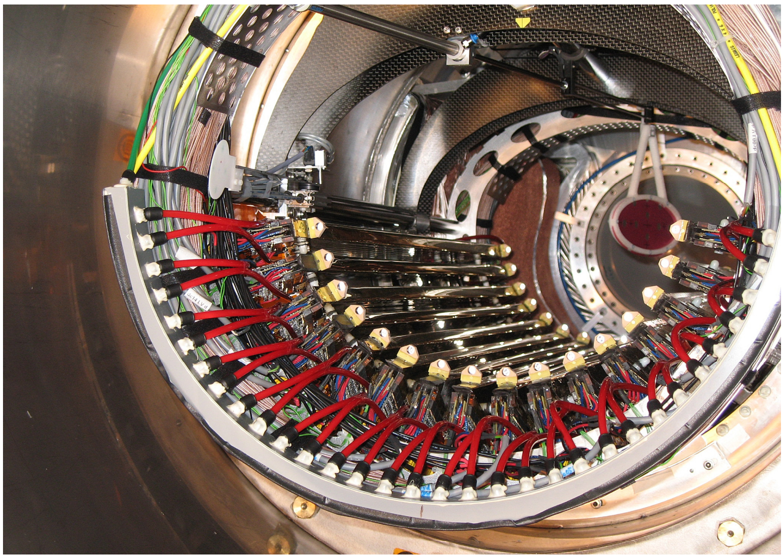

Figure 2.24: Drift chamber system of the former MEG experiment.

2.5.2 Positron Spectrometer

The former MEG spectrometer (depicted in figure 2.24) was made of 16 independent trapezoidal drift chamber modules aligned along Z and disposed in half a circle around the target, paired with a Timing Counter made of two sectors each with 15 bars of scintillator for timing.

The single drift chamber module had a trapezoidal shape with the long side, facing the target, measuring 104 cm, and the other, attached to CoBRa, 40 cm long. The sensitive region was composed of layers of alternating anode and cathode wires each enclosed between two metallised cathode sheets, segmented with a specific pattern (Vernier), whose charge provided additional information to improve the Z estimation.

Several aspects of the spectrometer proved to be inadequate: the drift chambers extended along the beam line for a very minimal amount with the support structure lying in between the tracking area and the Timing Counter and not far from the track of the positrons. A clean track association to Timing Counter was not easy because of the long track extrapolation needed to propagate between the tracking region as visible from figure 2.23 and the Timing Counter and resulted in a small overall tracking efficiency of 30%, to be compared with the 70% expected for MEG II.

Several aspects of the spectrometer proved to be inadequate, as visible from figure 2.23:

- the drift chamber system longitudinal extension was limited measuring only the first part of positron track;
- a long track extrapolation was needed from the drift chamber to the TC;
- the drift chamber support material was not distant enough from the positron track;

- the drift chamber modules suffered for Malter effect and were not sufficiently stable;
- the modularity of the drift chamber system required a precise alignment of the various modules.

All of those aspects have been addressed with the new spectrometer, in particular by the new monolithic CDCH.

The number of hits per track was really small, with an average of 10 hits with a spatial resolution of $\sigma_r^{\mathrm{MEG}} = 210\ \mu$m; one can compare these numbers with those of the new CDCH: 77 hits and a predicted resolution $\sigma_r^{\mathrm{MEGII}} = 110\ \mu$m. Furthermore the stereo geometry of the new tracker provides a better Z coordinate estimation during the track fit resulting in better angular resolutions and the small overall material budget improves the multiple scattering.

The former Timing counter bars were made of tick $4 \times 4 \times 80$ cm^3 fast plastic scintillators being readout at both end by PMTs. Even if the single bar, on the paper, should provide a resolution of ~ 40 ps, the impact of the CoBRa magnetic field and the high variation of optical path due to its large size reduced its performances to ~ 70 ps. By comparison MEG II Tiles have an average intrinsic resolution of only 72 ps for the 40 mm variant and 81 ps for the 50 mm ones, but the insensitivity to the magnetic field and the possibility to average multiple hits pushed the resolution on signal tracks down to ~ 30 ps at the expense of an increased number of channels.

2.5.3 DAQ requirements of the MEG II experimental setup

The MEG II detector introduced an overall segmentation increase with respect to the former experiment to improve the spatial resolution and cope with the increased muon stop rate. Correspondingly the total number of readout channels (detailed list available in table 2.4), not including the auxiliary detectors, jumped from 2694 of MEG by a factor 3 to 8188 in the new system, or 9340 if the full CDCH is readout and not only the wires in acceptance for signal positrons.

Over 50% of MEG II channels will be connected to SiPMs, either alone or in series, each requiring a dedicated source of bias voltage paired with a high gain amplifier compensating for the the overall smaller signal amplitude due to a smaller photostatistics and smaller gain ($8 \cdot 10^5$ for Liquid Xenon MPPCs) if compared to former PMTs (with nominal gain of $1.8 \cdot 10^6$).

It s clear that the simple scaling of the former DAQ system was not trivial at all, especially considering that the electronic boards being developed more than 10 years ago relied on components and technologies no longer in production or up to date with modern standards. The MEG online system was composed of two independent subsystems: the readout was accomplished with a custom designed digitising mezzanine[83] plugged on commercial carrier boards in a VME bus shown in figure 2.25a. Simultaneously a copy of the same signal was sent to the trigger subsystem[84], also composed of VME boards (figure 2.25b), where it was also sampled by 100 Million samples per second (MSPS) analog to digital converters and then processed by Field Programmable Gate Arrays

Detector	Channels	notes
MEG I		
Liquid Xenon detector	846	
Drift Chambers	1728	
Timing Counter	60 × 2	Both analog and a discriminated signal were readout
total	2694	
MEG II		
Liquid Xenon detector	4760	4092 MPPCs and 668 PMTs
Cylindrical Drift Chamber	2304	3456 for full readout
Pixelated Timing Counter	1024	
RDC	100	12 × 2 for the plastics and 78 for LYSO
total	8188	9340 for full CDCH readout

Table 2.4: Comparison of requested amount of readout channels between the MEG and the MEG II experiment.

(FPGA) connected on a tree-like structure to aggregate the information and generate the trigger signal for the whole system.

The double digitisation process employed in the MEG online system simplified the interplay between trigger and readout subsystems being almost independent from each other, on the other hand it resulted in a not optimal usage of the available space nearby the detector, including the required active splitting system.

By a careful tuning of the readout firmware a fixed DAQ dead time of 625 μs was achieved, which was needed to complete the readout of the DRS4 digitiser chip (also employed for the MEG experiment: section 3.1.1), but the maximum throughput of the full DAQ was limited to $\sim$ 40 events per second by the throughput of the VME-based readout system.

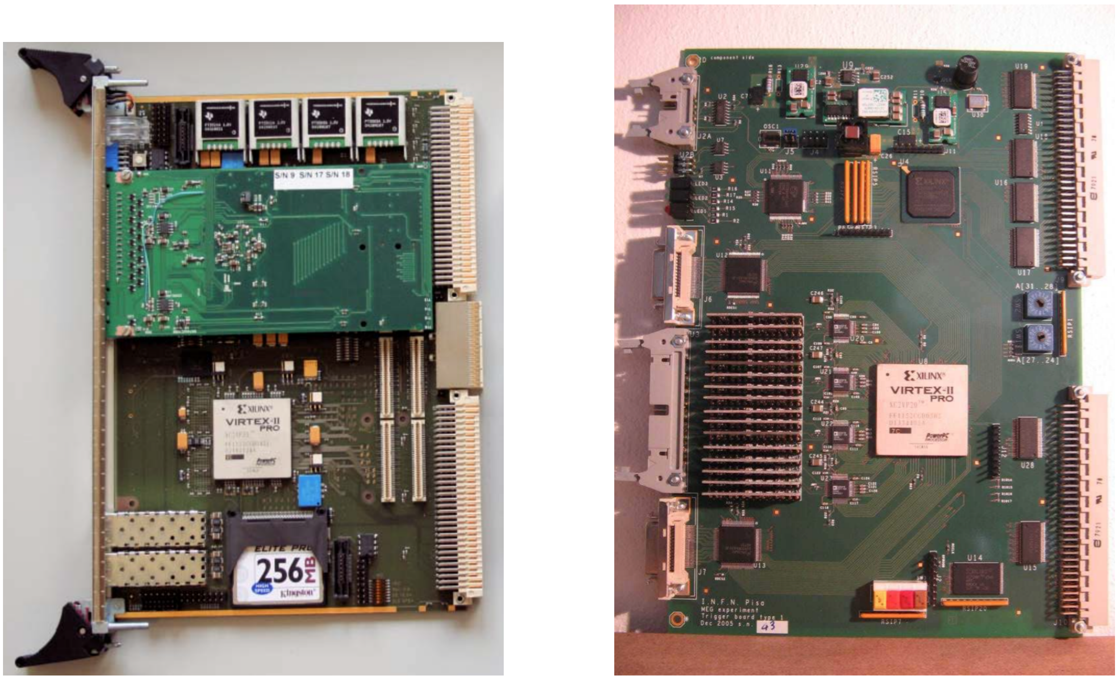

(a) DRS readout electronics.

(b) Trigger digitization card.

Figure 2.25: Former MEG digitisation components.

3 The WaveDAQ Trigger and Data Acquisition System

The responsibility for developing the Trigger and Data acquisition system is shared between PSI, which developed the MEG I data acquisition, and INFN Pisa, who designed and built the former Trigger system. The result is the merging of two previously distinct items into a highly integrated system capable of overcoming the readout speed limitations and the lack of scalability of the former VME-based system.

3.1 WaveDREAM readout Board

The building block of the system is the so-called WaveDREAM (Waveform Domino Readout Module) shown in figure 3.1a, a custom 3U-board designed to feature a complete 16 channel TDAQ system, each channel including front-end stage and voltage supply modules in less than 10×16 cm^2.

The channel front-end is designed to operate an AC-coupled, double stage amplifier (0.5-100 gain) and a programmable Pole-Zero compensation circuit to sharpen each pulse and help suppressing signal pile-up, although at the cost of a reduced charge sensitivity. A sketch of the frontend is shown in figure 3.1b in the dashed area. Different hardware versions of the front-end, depending on the detector, are available to operate on single-ended signals (which is the case of either SiPMs or PMTs), or differential ones (as those coming out of CDCH pre-amplifiers)[1].

WaveDREAMs connected to SiPMs can also source a programmable bias voltage up to 210 V through a piggy- back Cockcroft-Walton multiplier stage (the light green box in figure 3.1b), providing a common HV level for all the channels. Individual voltages can be then tuned channel by channel through a set of digital to analog converters starting from the common HV. If HV is not needed on one or more channels, a solid state switch can disconnect it leaving the adjacent channels powered. Applied voltages as well as driven currents are periodically recorded for slow-control monitoring.

Frontend outputs undergo a single to differential conversion before being digitised. In parallel, signals are discriminated by fast analog comparators with a tunable channel-by-channel threshold. Discriminated values provide an input for trigger algorithms as described later. Signals are sampled by a custom sample-and-hold integrated circuit named Domino Ring Sampler 4 (DRS 4); conversion is accomplished by a commercial 80 MSPS analog to digital converter (ADC).

---

[1] In the latter case, the first amplifier stage is replaced by a differential-to-single-ended conversion so that the same highly compact single-end frontend can be used for the following stages even though the maximum gain is reduced to 20.

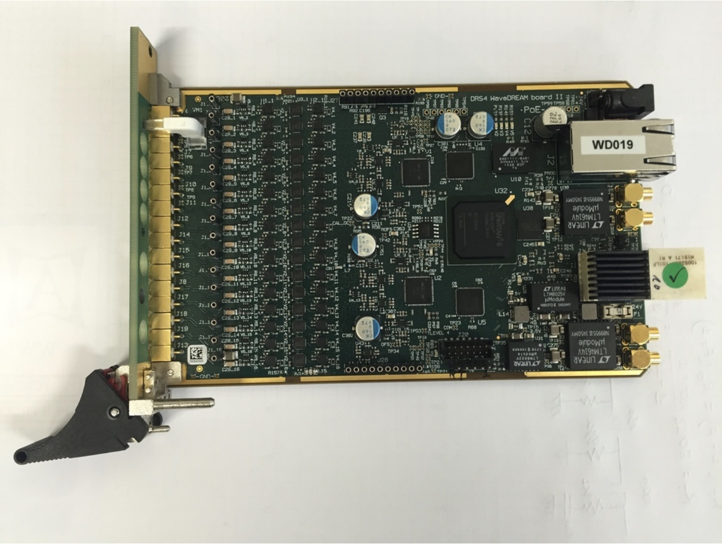

(a) Picture of board; input connectors are visible on the left side followed by frontend stages. On the right side, the board hosts the connector to the crate backplane.

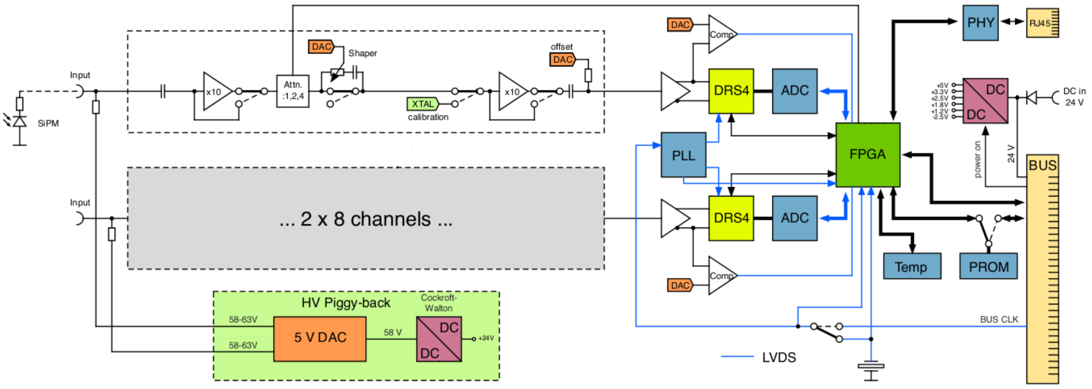

(b) Block schematic; a single channel frontend is shown in the upper left part.

Figure 3.1: The WaveDREAM card.

The full process is mastered by a Xilinx Spartan 6 FPGA[85], which gathers the full data set of ADC samples and operates a finite state machine for data transmission to DAQ software. A Microblaze CPU, instantiated in the FPGA logic, provides access to the various configuration bits from both the external hardware as well as from the internal Finite State Machines.

Being provided with a standalone trigger capability, the WaveDREAM can be also conceived as a 16-channel GSPS oscilloscope coupled to a dedicated software via an onboard ethernet interface. This feature is very useful for fast signal checks and quick bench-top tests, which do not need to involve the complete TDAQ chain.

3.1.1 DRS4 Ring Sampler Chip

MEG II goal sensitivity can be exploited if accurate time and amplitude reconstruction, as well as efficient pile-up rejection, is achieved on individual channel signals. This demands waveform digitisation at GSPS sampling speed on $\sim$ 8000 channels, which would not be feasible with commercial flash-ADCs, neither in terms of costs nor thermal and power requirements.

A viable solution is a dedicated PSI-developed analog chip, named Domino Ring Sampler 4 (DRS 4), shown in figure 3.2. It consists of a 9-channel array of 1024 switched capacitors being conceived as analog cells to store signal waveforms long enough for the trigger to be generated. To obtain such a high sampling speed, the fast write switches of those cells are operated through the inverter chain shown on top in the schematic of figure 3.2b. A logic wave (the so-called Domino wave) propagates through the chain to enable only one cell at a time to be charged at the current input voltage. The stability of the Domino wave is guaranteed to be better than $\sim$ 100 ps thanks to an external phase locked loop (PLL) chip.

The propagation delay between cells, hence the sampling speed, can be regulated by means of variable resistors controlled by means of an external voltage (U_{speed} in figure) between inverters. After a trigger, the Domino wave is stopped: the stored charge in each cell can be fed through a read switch to the input of the external ADC.

Optimal signal-to-noise ratio (in our case 1 mV RMS over 1 V range) demands off-line cell-by-cell amplitude calibration to take into account:

- first-order effects, mostly due to cell-to-cell capacitance spread;
- second-order corrections to compensate for time-dependent charge leakage of a cell (due to the finite resistance of Domino switches).

To compensate for the first effect, a cell-dependent pedestal is subtracted from the raw ADC value; then a cell-dependent weight is applied to compensate the capacitance spread. Charge leakage is corrected by adding a value to each cell according to the readout order[2].

---

[2] If the readout order correction were not applied, the baseline would exhibit a clear negative slope due to part of the charge in the capacitor being lost

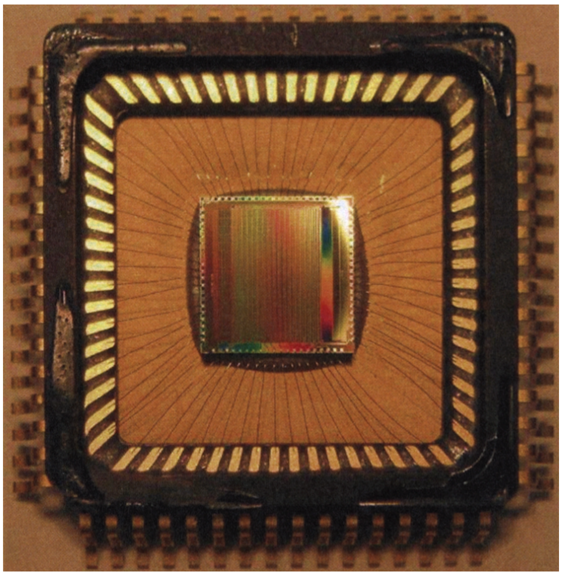

(a) picture

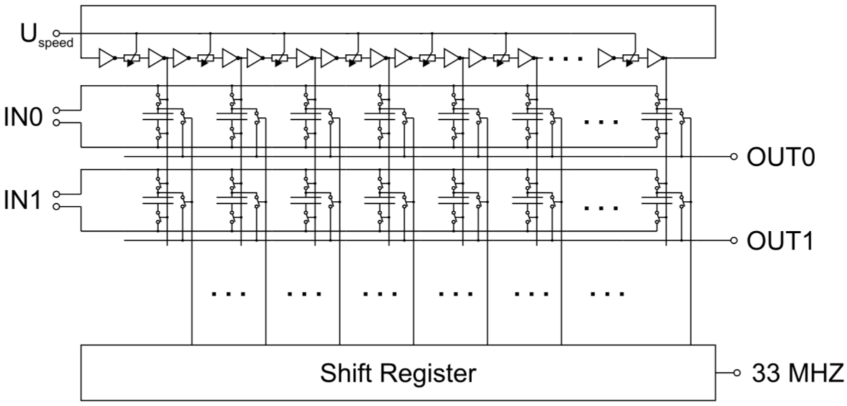

(b) internal diagram

Figure 3.2: The DRS4 asic sampler.

A calibration source is available on the WaveDREAM and can automatically be connected to any channel in order to characterise the chip response. Voltage calibration factors are computed once and stored so that they can be applied to data before being recorded to disk.

A timing calibration is needed to accurately measure the inter-cell delays, at level of $\sim$ 10 ps (one order of magnitude improvement). As a result, time bins in calibrated waveforms are no longer equally spaced and offline pulse reconstruction must properly account for this[3]. The measurement proceeds through an automatic procedure which distributes a sinusoidal signal to all inputs and iteratively corrects the time bins around its zero crossings by approximating the sine with a straight line.

While eight DRS4 inputs are connected to as many input channels, the ninth is used to digitise a copy of a reference clock (the same being used to lock the PLL) and distributed to all the WaveDREAM boards via very low jitter fan-outs ($\sim$ a few ps). Event-dependent time offsets can be computed and applied for channel-to-channel synchronisation.

All of the mentioned calibrations depend on the internal switching performances of the DRS chips and, for this reason, show a dependance on the silicon temperature. A temperature sensor is therefore placed near the DRS to monitor temperature and to warn if the measured value is too different from the one registered at the moment of calibration.

The size of DRS4 analog cell array poses the most stringent limitation to trigger latency. If f_{DRS} is the chosen sampling speed, the Domino wave must be stopped at latest $1024/f_{DRS}$ after the occurrence of an event of interest. Otherwise the charge stored in part of the cells may be overwritten and a few samples of the waveform (in particular those recording the leading edge of detector signals) may be lost. So in the case of $f_{DRS} = 2$ GSPS, the trigger latency must be shorter than 500 ns. Moreover, it would be advisable to preserve a few tenth sampled values before the leading edge for unbiased event-by-event pedestal computation.

In a WaveDAQ setup which involves fewer channels and requires a simpler trigger (maybe based only on comparator outputs) can operate at a faster DRS sampling speed because the trigger computation can be performed in a small time.

3.1.2 Trigger information extraction with ADCs and TDCs

While the DRS is sampling, the same ADC used for digitisation, which would otherwise stay idle waiting for a trigger to arrive, is used to simultaneously sample the signal, with a smaller frequency, to be used as an input for amplitude-based trigger selections.

Simpler threshold-based triggers use the fast analog comparators in the frontend with much shorter latency than ADC conversion. Accurate online timing can be achieved by oversampling the state of the comparator output with the fastest clock available in the current FPGA firmware: a 640 MHz clock is used resulting in $\sim$ 1.56 ps bin timing, much better than that achievable using only the 80MHz ADC samples.

---

[3] For example, before using any Fourier Transformation on the waveform a resampling with a consistent time bin is needed not to produce artefacts.

When used in the trigger, ADC values need to be corrected for possible slow pedestal fluctuations: raw samples are arranged in a shift register to obtain a pedestal value by averaging few previous samples before the current sampling. To avoid pulses to enter the shift register, a raw ADC value will be used in the pedestal computation only if the difference from the expected pedestal value is smaller than a programmable threshold.

3.1.2.1 ADC-DRS channel correlation

The density of the WaveDAQ system is achieved by sharing frontend electronics and ADCs between the trigger computation and the DRS readout in a single board. WaveDREAMs can save both DRS and ADC waveforms for the same event for offline quality checks. In figure 3.3 a comparison of the ADC recorded pulse height with the DRS equivalent is shown for a group of channels (those connected to the SIPMs of MEG II Liquid Xenon detector). A clear correlation is observed for most of the events; a small fraction which looks out of correlation is due to events where a pileup signal occurs.

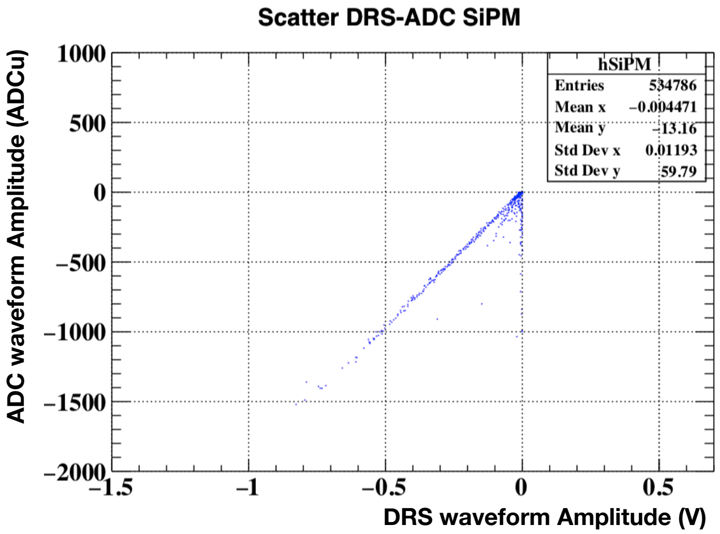

Figure 3.3: ADC-DRS scatter plot.

Online ADC resolution is worsened, in the current prototype WaveDREAM version, by the pickup of noise from the DRS chips. Further details will be given in chapter 4.

3.1.2.2 Single Channel TDC Timing

MEG I trigger[86] could only rely on 100 MHz ADC samples. A coarse (10 ns) time estimate was based on the first sample over a programmable threshold; a fine adjustment (2.5 ns accurate) was performed depending on the slope of the leading edge, as obtained from the first two samples above threshold.

In MEG II the dedicated hardware comparators provide a direct measurement of the crossing time: their output is oversampled with a 8 times faster clock; for each tick of this clock (1.5625 ns period), one bit is recorded containing the comparator state. These bits are arranged in a 8-bit register with the comparator output at equally spaced time bins along the clock period. This provides an estimate of the hit time at the level of the fast clock period without no need to interpolate the signal. The procedure can therefore be applied also to very narrow pulses, like those of Timing Counter (~ 10 ns wide pulses due to Pole-Zero compensation), which could not be reconstructed by the an ADC waveform interpolation.

In figure 3.4, the result of a laboratory test is shown; a 20 ns wide, 200 mV high pulse from a function generator is fed through a passive splitter to two adjacent channels in a WaveDREAM. The difference of the two recorded time is expected to have an RMS of $1.5625\ \text{ns} \cdot \sqrt{2}/\sqrt{12} = 638\ \text{ps}^4$. The RMS of the resulting distribution is 484 ns, slightly better than expected because the same clock is used to sample both channels: which results in a positive correlation of the two estimates.

A more realistic estimate of timing resolution, including also intrinsic detector resolution effects, will be shown in section 4.5.

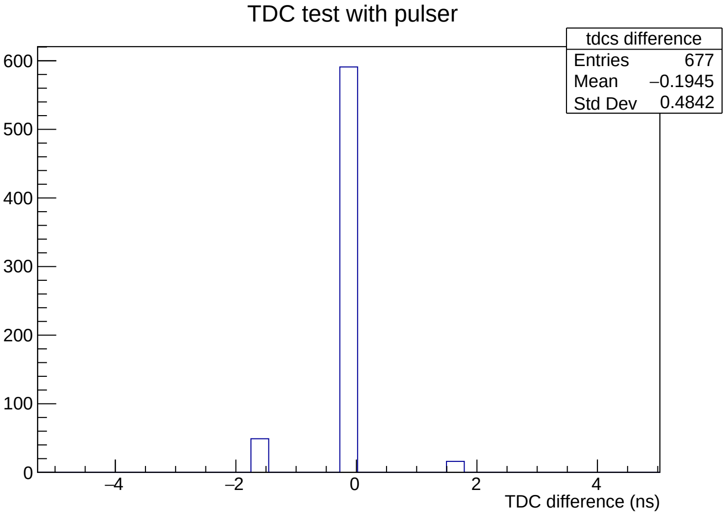

Figure 3.4: Reconstructed TDC time difference of two pulses in a lab setup.

3.2 WaveDAQ Crate and the Crate Management Board

WaveDREAM boards can operate in single mode coupled to up to 16 channels. In larger systems, with many boards needed, all the pieces of information obtained from ADC and

---

[4] $\sqrt{12}$ comes from the RMS of a flat distribution; $\sqrt{2}$ is due to the time difference of two uncorrelated measurements

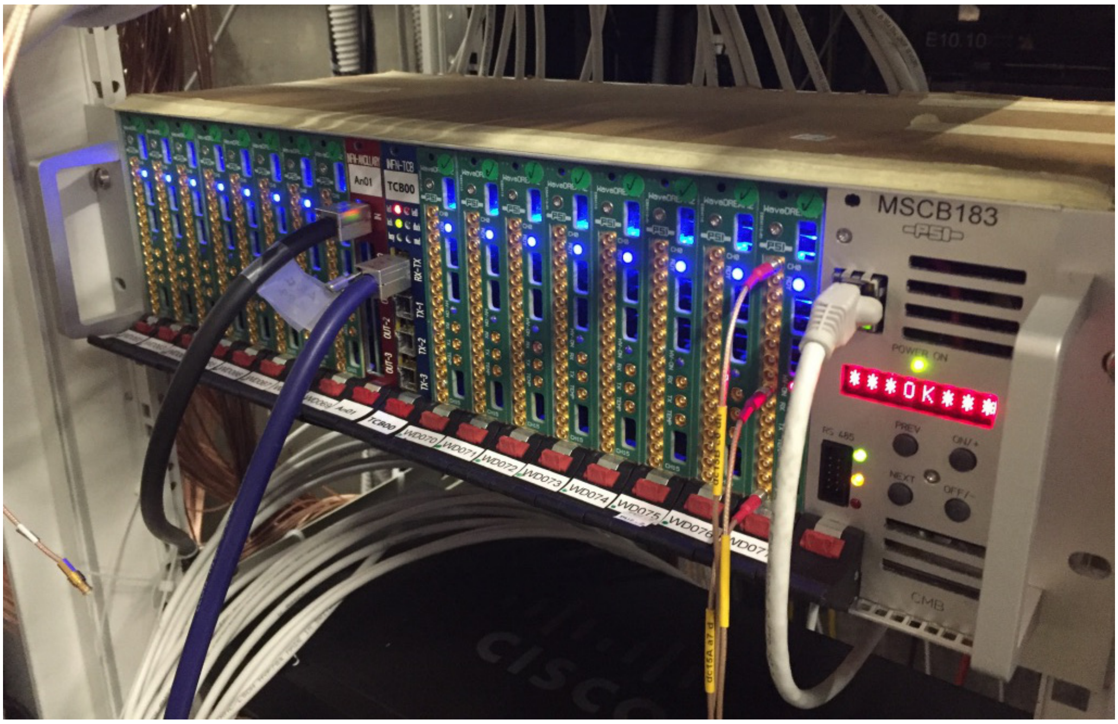

Figure 3.5: WaveDAQ Crate with, on the left, a Crate Management Board.

TDC algorithms, need to be gathered to a global system which is in charge of generating a common trigger signal.

The basic unit of a multiple board WaveDAQ system is a crate like the one shown in figure 3.5. Up to sixteen WaveDREAM cards can be hosted in a single 3U crate with 18 slots. The two central slots are reserved to interface boards, one to the data acquisition and the other to the trigger subsystem (both to be described later).

A Crate Management Board (CMB), the rightmost in the picture, supplies power and provides management for boards within the crate:

- along the backplane, a 24 V shared power rail provides up to 400W with a monitoring of the actual voltage and current.
- The temperature of the crate is recorded in a few points and used to regulate the speed of three fans on the back of the crate. This ensures a good airflow which is mandatory to keep the DRS temperature in the operational range.
- Dedicated lines allow us to remotely power on and off any slot.
- A “Plug In” signal is also available from each slot to mark if a board was physically inserted.
- A soft reset function is handled by the “Init” signal which drives a pin of the FPGA causing it to reboot without power cycle.

The CMB can access the boards in the crate using an adapted version of the widespread Serial Peripheral Interface (SPI) specification[87]. The connection scheme follows the standard SPI interface: a clock signal is used to synchronise data shift out of the transmission master (the CMB) to the slave (any board in a given slot) on the MOSI line while, at the same time, signals are shifted out of the slave to the master on the MISO line. All those signals use Low Voltage Differential Signalling (LVDS) to provide immunity to environmental pickup and fast (~ 1 MHz toggling rate) transfer speeds even on 1m long backplane line. Considering that each slot in the crate is a slave (and the left central slot might also be a master, as described in section 3.7.1) data and clock lines must be shared on the backplane by all of them. Each board, including the CMB, have therefore the possibility to disable their output drivers if they are not involved in the transmission, so as to avoid any bus conflict[5].

To select the slave to be addressed a dedicated "Slot Select" line is available in each slot. Most of the boards in the WaveDAQ contain an FPGA with its own configuration memory; a shared "Flash Select" signal is used to select between a configuration memory operation and an FPGA transaction within the board in the slot with the "Slot Select" signal.

3.3 Trigger Generation

Trigger requirements become challenging with an increasing number of signals coming out of different detectors. For instance, event selections based on charge reconstruction is more relevant in the case of Liquid Xenon detector; while timing reconstruction, based on comparator output, is essential to trigger on pixelated Timing Counter hits.

The use of FPGAs is widespread in particle physics for quick online event selection, as demanded in current experiments. This kind of programmable chips are best suited for demanding computations which can be split into a bigger number of smaller steps, each being registered on the edge of a clock. These registration stages are useful to interrupt a combinatorial path which would be too long to be executed at the nominal clock frequency. On the other side a large number of these registers has the counter effect of increasing the latency which must be kept as small as possible according to the DRS sampling speed.

Also the choice of trigger algorithms must obey a compromise between performances and execution speed; it happens that the best algorithm has to be discarded in favour of a less refined, but faster one. An example is given by calorimetric triggers: the calculation of a signal charge would guarantee the best resolution; however, if the signals have a long tail, the number of samples to add up increases rapidly and the latency is too long

---

[5] Disabling an LVDS driver is not foreseen in this standard as it was developed for a point-to-point transmission. A variant of the LVDS, the "Bus LVDS", is used instead: every board is provided with an output disable and every driver is required to be capable to drive higher currents. Two termination resistors are indeed needed in the Bus LVDS to properly operate the backplane transmission line without signal reflections from the two edges. The Bus LVDS driver must therefore be able to drive twice the current of a standard LVDS driver.

to be suited. A faster, peak finding selection turns out to be the best compromise in that case.

3.3.1 Trigger Concentration Board

The trigger system which collects the pieces of information generated by the WaveDREAMs is arranged in a tree-like structure whose basic element is a custom Trigger Concentration Board (TCB).

Data transmission and processing is operated by an onboard Xilinx Kintex 7[88] FPGA; the huge amount of LVDS transceivers available in that chip[6] allows us to use the same printed circuit board in two different board configurations, “Master” or “Slave”, just by changing front and backplane connectors as shown in figure 3.6. The “Slave” TCB board has four front panel connectors and just a single backplane; the “Master” is tailored to receive backplane signals and therefore needs five connectors and a single one on the front panel. Thanks to this flexibility, it is possible for any TCB to play three different roles depending on their location in the system:

- TCBs in the Master configuration are hosted in the reserved right central slot in a crate, as shown in the two bottom crates in figure 3.7. In this slot, 16 Trigger Serial links, whose implementation is described in section 3.5, gather the trigger information from the 16 WaveDREAM boards hosted in the same crate.
- TCBs in the Slave configuration are located in a dedicated Trigger crate, in slots occupied by WaveDREAMs in the other crates. Each board is connected to up to four of the previous TCBs to collect data from four different crates, as shown in cyan in figure 3.6. Data transmission uses the same Trigger Serial links even though through an LVDS cable.
- The last role is accomplished by a single Master TCB that is placed in the dedicated slot in the Trigger crate. That board is the master trigger board of the full system, that is to say the one to operate final event reconstruction and to generate the trigger.

The master trigger board firmware supports up to 64 independent trigger conditions (listed in table 3.1 and identified by an ID code), each with a dedicated rate counter. Each of these trigger conditions can be prescaled: a down-counter for each of them starts from a programmable value and is decreased any time the condition is met when it reaches zero and is automatically reloaded to the initial value. The prescaling is crucial to mix different triggers according to the desired fraction. Prescaling factors must be tuned as to take the overall DAQ rate under control.

Different latency due to different conditions is compensated by delay generators such that the triggered pulse occurs roughly at the same position in the DRS waveform.

If more than one condition is met at the same time, the priority is given to the one with the smallest rank in the condition list; so condition ID code 0 is assigned to MEG signal

---

[6] The used package provides 900 connections for power and signal input/output. More than 170 LVDS lines are connected for an aggregated throughput of more than 100 Gbps.

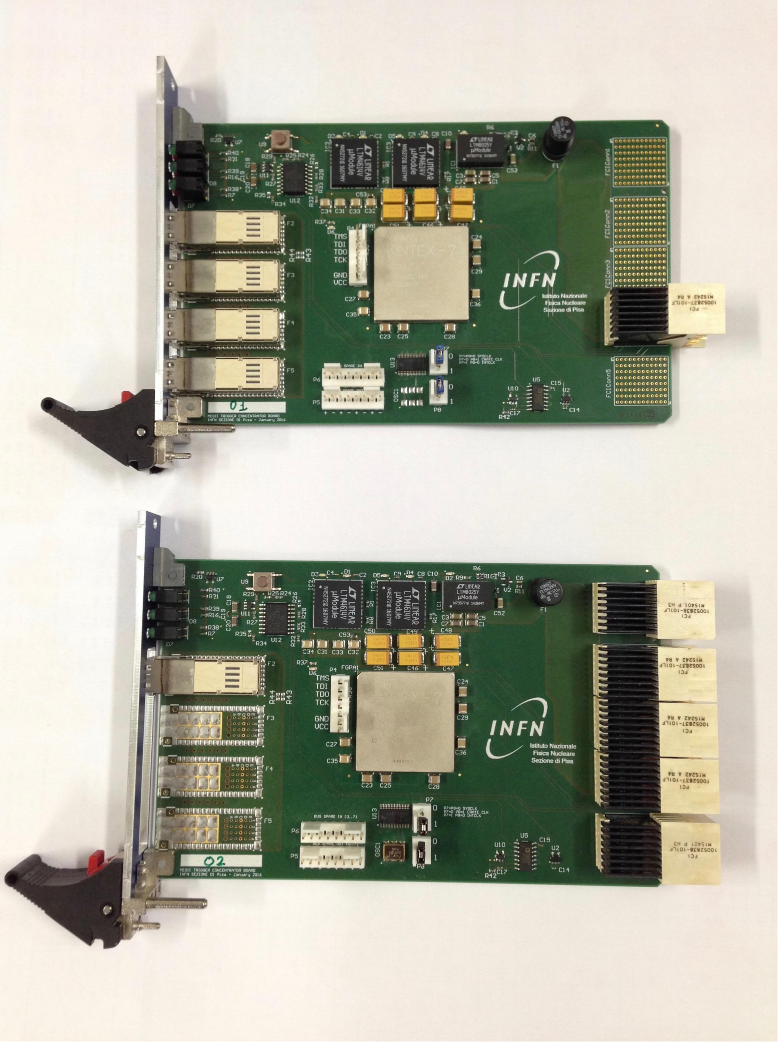

Figure 3.6: Trigger Concentration Board, in the two configurations: “Slave” (top) and “Master” (bottom).

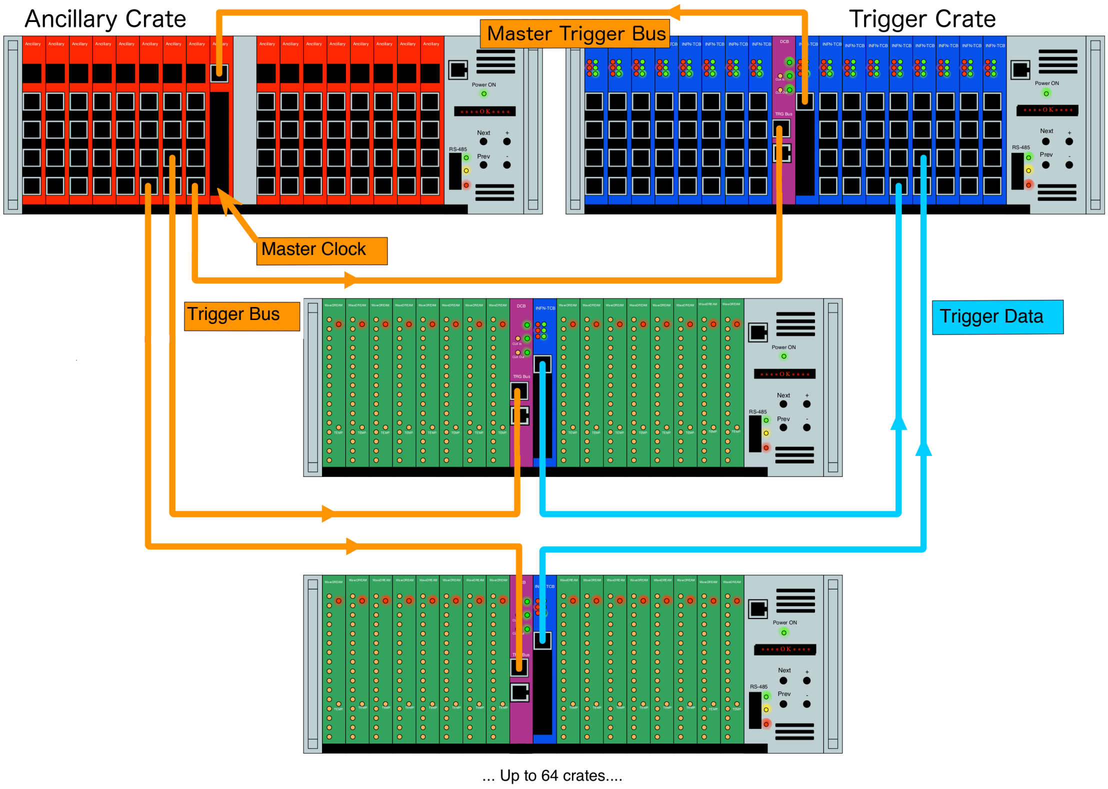

Figure 3.7: Data flow and trigger signals flow in the multi-crate WaveDAQ system: WaveDREAMs in green, TCBs in blue, DCBs in violet and Ancillary board in red.

events, while ID code 63 is reserved to random pedestal trigger. In between increasing ID code is assigned to looser trigger conditions used to acquire detector calibration events.

The condition ID code identify the “Trigger Type” which is used to trace back the condition which fired the trigger. Together with the Trigger Type an progressive “Event Number” is assigned to each event by the master TCB for unambiguous identification.

3.4 Trigger Signals Distribution

When a trigger is generated it is distributed back to all the boards through the connections marked in orange in figure 3.7: a fanout system made by custom designed Ancillary boards (figure 3.8), placed in a dedicated crate, is used. Apart from the trigger, there are others to be distributed:

- a copy of the main low jitter 80 MHz reference clock to be used for DRS sampling and trigger data transmission;

ID	**Name**	
0	MEG	MEG physics trigger
1	MEG LowQ	MEG physics trigger with reduced E_γ cut
2	MEG WideAngle	MEG physics trigger with reduced Direction Match
3	MEG WideTime	MEG physics trigger with reduced $T_{\mathrm{e}\gamma}$
4	Radiative NarrowTime	Radiative trigger: E_γ and $T_{\mathrm{e}\gamma}$
5	Radiative WideTime	Radiative trigger with reduced $T_{\mathrm{e}\gamma}$
6	Michel	Positron combined trigger to be defined
10	Xenon HighQ	E_γ High cut including veto
11	Xenon LowQ	E_γ Low cut including veto
12	Xenon Alfa	Pulse-shape alfa trigger
13	Xenon LED 0	External led pulser signal
14	Xenon LED 1	External led pulser signal
15	Xenon Cosmic	E_γ above veto threshold
20	TC Track	Any Timing Counter track
21	TC Single	Or of all Timing Counter tiles
22	TC Cosmic	TC trigger to be defined
23	TC Laser SYNC	External laser SYNC signal
24	TC Laser Photodiode	Photodiode signal in the laser system
30	RDC-XEC	Coincidence of RDC trigger and Xenon HighQ
31	RDC trigger	Combination of RDC conditions
32	RDC scintillators	Or of plastic scintillators
33	RDC LYSO Sum	Discrimination on sum on LYSO crystals
34	RDC LYSO Single	Or of LYSO crystals
40	CDCH Track	CDCH tracking trigger to be defined
41	CDCH-CRC	Coincidence of CDCH cosmic trigger and CRC bars
42	CDCH Cosmic	Internal CDCH cosmic trigger to be defined
43	CRC Pair	Coincidence of any pair of top and bottom CRC bars
44	CRC Single	Or of all CRC bars
45	CDCH Single	Or of all CDCH channels
50	Pi0 trigger	Charge Exchange trigger for T_γ
51	Pi0 no preshower	Charge Exchange trigger for E_γ
52	BGO no preshower	Combination of BGO conditions without preshower
53	BGO Sum	Discrimination on BGO crystal sum
54	BGO cosmic	Coincidence of a top and a bottom crystal
55	BGO preshower	Combination of BGO conditions with preshower
56	CW Boron	Xenon sum threshold and Timing Counter hit coarse coincidence
58	Neutron Generator	Xenon sum threshold in a window given by Neutron generator
63	Pedestal	Random generated trigger

Table 3.1: Trigger list, see section 3.6 for more details. Trigger IDs not reported in this list are reserved for future use.

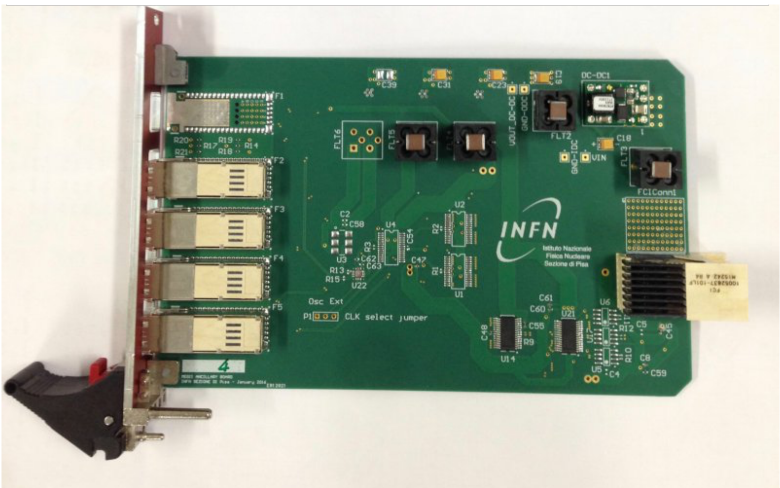

Figure 3.8: Ancillary board.

- a synchronisation signal "SYNC" which provides a common start signal for all the DRS chips as well as a globally synchronous reset;
- a serialised trigger information for the triggered event; which encodes Trigger Type and the Event Number so that they are available on every board to be attached to all data fragment for later event building;
- a Busy signal going back to the trigger board.

All these signals are arranged in a bus called "Trigger Bus" and transmitted to each crate. The other end of the Trigger Bus cable is connected to a board, named Data Concentrator Board (DCB, more details of this board in section 3.7.1), placed in the left reserved slot in the crate. The DCB drives the signals from cable to shared lines in the backplane.

When a trigger is received waveform sampling is stopped; each WaveDREAM starts converting the DRS waveform and, in the meanwhile, asserts an open-drain BUSY line on the backplane. That shared line acts as a wired-or signalling if any board in that crate is busy. The crate BUSY state is back-propagated by the DCB to the Ancillary system through the cable. A mix of few logic chips (inside each Ancillary card) and the backplane wired-or combines the information of the different crates and provides a "System Busy" signal to the main trigger board. Trigger generation is inhibited until all WaveDREAMs release their own BUSY signalling that they are ready to record a new event.

3.5 Trigger Serial Links

The procedure to transmit data in the trigger system must also be tailored according to requirement of minimal latency, maintaining at the same time a bandwidth large enough to contain all the information.

The most recent generations of FPGAs are provided with multi-gigabit transceivers, designed to implement widely used protocols like 10-Gb Ethernet or PCI-express communication. However all those communications exhibit a latency of few hundred nanosecond beyond what is sustainable in the WaveDAQ.

In multi-gigabit transceivers the latency a clock recovery on the receiving end is mandatory: transmitting the clock together with the data is simply not practical in addition to require a dedicated line between transmitter and receiver which could carry data instead. Unfortunately clock recovery demands for additional operation to be performed on data[89]; special sequences of bits, named Commas, are inserted in the stream so that there are enough transitions for the receiver to recognise them and generate a local copy of the transmission clock. This process clearly introduce latency because of the because of the additional operations, furthermore the deserialisation clock will be different for each incoming connection and also different from the local clock used to run the algorithms and, eventually, to forward data to the next stage. Additional latency is then wasted to properly buffer data so that they can be safely moved to the algorithm clock.

Fortunately in the WaveDAQ system we already distribute a very stable shared clock from the Ancillary system to synchronise all of the DRS4 samplings. The transmitter board can therefore serialise data using a local copy of the shared clock (suitably multiplied to drive the transmission) and the receiver deserialises them using its own copy. This scheme[90] solves the need of both clock recovery, without wasting links to forward the transmitter clock, and clock domain crossings, because all the serialisation happens on the shared clock (which is therefore used also for the algorithms).

The choice of not using multi-gigabit transceivers has the drawback of a reduced transmission speed but with the advantage of using the slower “SerDes” serialisation and deserialisation primitives (having transmission frequencies up to 1080 Mbps in Spartan 6) which are available at each FPGA I/O pin[91]. The abundance of data lines in each FPGA partially compensates the reduced throughput of a single line by aggregating more of them in a single link sharing the load.

In the final design, each Trigger Serial link uses 8 LVDS pairs, each sending 8 bit every 80 MHz clock cycle (being limited by the Spartan 6 FPGA available on the WaveDREAM board), for a 5.1 Gbps overall throughput for each link.

Serialised data can reach the receiving board, either from the backplane, if communication is between a WaveDREAM of a TCB Slave and a Master TCB of the same crate, or from the front panel if a cable is used, in case of the communications between TCBs in WaveDREAM crates and TCB Slaves. In both cases it needs to be properly sampled on a local copy of the shared clock. Such procedure is not straightforward: in fact data from different links, depending on the physical link length, may arrive at the receiver at any phase with respect to local clock, because data paths from different sources are not compensated in hardware. To solve this Xilinx Kintex 7 FPGA used in TCBs has a programmable input delay on each I/O line; the delay can be tuned so that bits are sampled when they are stable.

Even if the data are sampled correctly, the unknown delay of the data path can

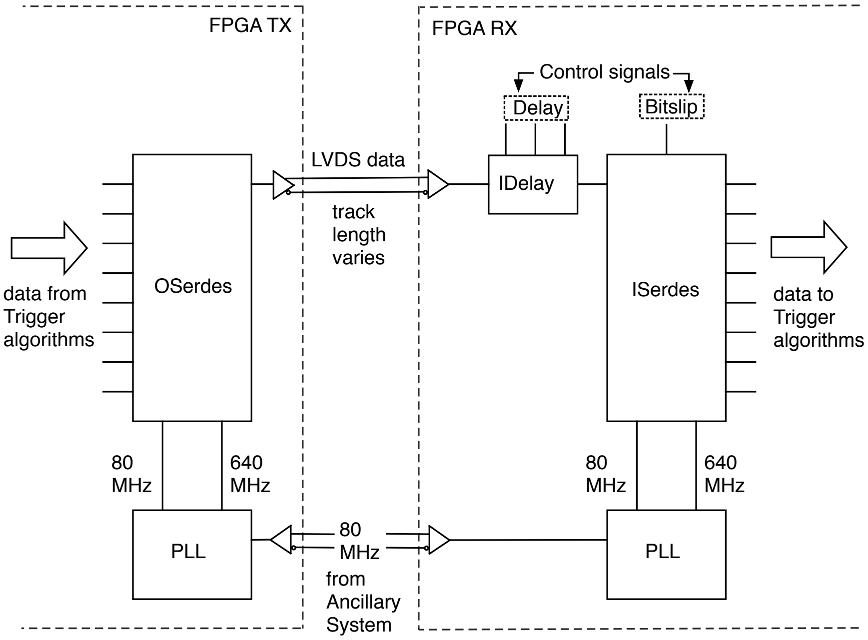

Figure 3.9: Schematic of a single LVDS line in a Trigger Serial link showing the various FPGA primitives and control signals used; the complete serial link contains 8 instances this LVDS line.

cause the deserialisation to start from a bit which was not the first one being sent. To compensate for this the receiving logic cannot be a simple shift register like it is for the transmitter: an additional feature in Xilinx serial receiver, named “Bitslip”, is used to select which bit to start deserialising first.

Figure 3.9 shows the schematic of a single LVDS line being used in Trigger Serial links: the shared 80 MHz clock from Ancillary system is multiplied on each board by a local PLLs to obtain the required bit-rate (640 MHz) clock. Both clocks are fed to OSerdes and ISerdes primitives respectively in the transmitter and receiving FPGAs; these blocks handle the parallel to serial conversion of data and vice versa. An IDelay is also shown before the receiving stage in figure 3.9; it is used to apply the programmable delays to compensate track length offsets (together with Bitslips).

There are therefore a 5 bit delay and a 3 bit “Bitslip” values to be adjusted for each LVDS line. In total we have slightly less of 10000 values to be tuned in the full system and even only loading them from a database would be prohibitive at each startup of the system. So a dedicated Finite State Machine (FSM) was developed to perform the link calibration in parallel at any system reboot; as the checks are established inside the receiving FPGA without need of any external intervention, the procedure is extremely quick (less then 1 s). The serial links calibration starts at the initial configuration of the system and is based on a known pattern being generated by each transmitter: the 8-bit word consists of a first 4-bit fixed part followed by a 4-bit counter value which is

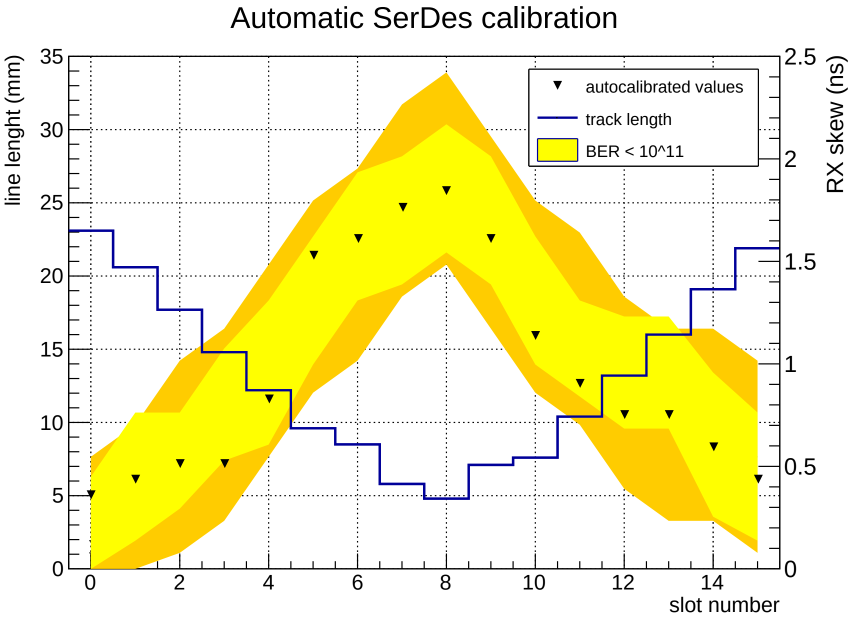

Figure 3.10: Trigger link delays selected for all slots by the auto-calibration FSM, track length is shown for comparison. Yellow band was measured to have Bit Error rate less than 10^{-11} in WaveDREAM-TCB transmission, Dark band is the same but between TCBs.

simultaneously reset by the SYNC pulse. The FSM looks for the fixed part and sets the delay value to receive it correctly and, after that, finds the Bitslip to have the correct bit ordering.

Another problem arises as the FSM may select values that makes different input links have different latency. This is very problematic on the trigger system which relies on the synchronicity of data propagation in order to combine reliably informations. A second FSM is then used to identify and correct such behaviour: once all the links are established the 4-bit counter of each link is used to measure the relative latency and to equalise all the transmission to the slowest one.

Figure 3.10 shows the values selected by the auto-calibration FSMs for the 16 backplane Trigger Serial links of a TCB Master. Each slot delay value clearly anticorrelates with the track length and all values lie within the yellow bands, which represents the values with bit error rates less than 10^{-11} in the WaveDREAM-TCB (light yellow) and TCB-TCB (dark yellow) transmissions. The TCB-TCB band extends beyond WaveDREAM-TCB one because of the increased switching performances of the Kintex 7 compared with the Spartan 6 in the WaveDREAM.

3.6 MEG Triggers

The main MEG Trigger (ID code = 0) relies on very different algorithms to be combined together: calorimetric reconstruction and time-stamping of the photon shower is performed on all Liquid Xenon signals; at the same time, the impact position and timing is extracted from pixelated Timing Counter pulses generated by ~ 100 kHz positron rate.

Each trigger selection criterium can be considered as an event cut, therefore the efficiency of all the trigger cuts must be known and properly monitored. For this reason, for each selection, two trigger conditions are defined, one being at the nominal value and one looser; the latter, registered with a proper prescaling, can be used to acquire a “pure” data sample to check the performances and the efficiency of the more stringent cut.

3.6.1 E_γ

The calorimetric reconstruction in the Liquid Xenon detector is the most effective tool in selecting event candidates. The background rate depends indeed to the square of the energy window as shown in section 1.4.2.2. A completely homogeneous calorimeter, with a huge amount of photosensors collecting a fraction of the light, makes the reconstruction not trivial if the best possible energy resolution has to be pursued.

Each of the 4760 channels are multiplied in the FPGA (using dedicated DSP blocks[92]) by a scaling factor accounting for channel-by-channel spread in both gain and detection efficiency[7].

Another factor has to be applied to take into account the local ratio between dead material and sensitive area. PMTs in the outer face are more spread and then is more likely for a scintillation photon to hit the support structure instead of the instrumented area. The collected light fraction is indeed different between outer face and lateral face PMTs which are closely packed and, for this reason, they should be weighted differently.

As already described the best estimate of the number of detected photons would be the charge integral. However the time needed to integrate the long tail of MPPCs would increase the trigger latency out of the system constraints. Because of this the online estimate relies on the summed pulse height instead of the charge. Unfortunately the pulse shape of the MPPC signals is different from PMT ones because of the bigger capacitance of the silicon detectors and the presence of additional charge sources (noise and after-pulses). MPPC weights must therefore include a compensation for the different charge to pulse height ratio. Moreover, PMT waveforms must be delayed by few samples to account for the different signal risetime.

The weighted sum so obtained is discriminated with three threshold: a “High” and a “Low” threshold are associated with the nominal and the loose trigger condition respectively. The sum is also compared with a third, higher, “Veto” threshold to reject events with large energy deposit due to cosmic muons. For the veto to work properly, the veto discrimination needs to be active few samples earlier than the others, during the rising

---

[7] Detection efficiency is called Photon Detection Efficiency (PDE) in case of MPPCs and Quantum Efficiency (QE) for the PMTs

edge, and stay on a little longer so as to avoid to trigger on the pulse tail. The High and Low discriminator outputs are delayed a few clock cycles, enough to wait the Veto condition to be met.

High and Low threshold conditions (with the corresponding Veto) are assigned trigger ID code 10 and 11 in the condition list (table 3.1). A Cosmic trigger is also available (at ID code 15) to monitor the unselected spectrum and for calibration purposes, which request a pulse height above the Veto threshold.

3.6.2 Direction Match

The Direction Match is the request of back-to-back photon and positron to be detected in Liquid Xenon detector and in pixelated Timing Counter respectively; Montecarlo simulations of signal events shows indeed that the γ-interaction vertex in Liquid Xenon detector is correlated with a certain region of the pixelated Timing Counter hit by a back-to-back 52.8 MeV positron (this correlation is spread out, at some extent due to matter effects on positrons). So, for each reconstructed position for the γ in Liquid Xenon detector, it is possible to define a set of TC Tiles, whose width depends on the desired efficiency.

The trigger is asserted only if the reconstructed $e^+ - \gamma$ pair is among the expected ones for a signal-like event. The list of possible pairs forms a “Direction Match table” which is stored inside FPGA BRAM memory blocks[93]. The table can be then indexed by means of the reconstructed values and returns a single bit output if such pair satisfy the Direction Match condition.

Also in this case, two different tables are loaded into the firmware at the compilation phase to fulfil either nominal or loose selection criteria.

3.6.2.1 Photon reconstruction

The photon interaction vertex, projected to the inner face, is reconstructed by using only analog comparators output. In the former experiment an estimator of that position was that of the PMT having the highest pulse height in the inner face. Such association is suited since the average transverse fluctuations of the shower are of the order of the Mòlière radius, which is 5.2 cm in liquid xenon, and is well matched with the old 2-inch diameter of PMT tube.

In the MEG II detector, the new 15×15 mm^2 MPPCs have a much higher granularity and the single channel is more sensitive to fluctuations of energy deposit on the inner face. Also to reduce the size of the Direction Match map, MPPCs were connected to the same WaveDREAM in 4-by-4 groups, later referred as Patches. The area covered by a Patch on the inner face is 6×6 cm^2 MPPCs, which matches better the Mòlière radius.

For each Patch a multiplicity is computed as the number of MPPCs over threshold. The Patch having the highest multiplicity is selected as estimate of the interaction point.

As described later, the Patch with the highest multiplicity is also the best one to use for timing.

3.6.2.2 Positron reconstruction

Positron reconstruction in the Timing Counter gets complicated by the design itself of the detector, which maximise hit multiplicity along the same particle track to improve time resolution. However, the hits due to a track need to be somehow disentangled from the possible pile-ups. As confirmed in 2018 Run, TC hits (as reconstructed offline to belong to the same track) shows on average 9 hits/track. However there is a large spread of the number of hits and, to keep the trigger efficiency as high as possible, no selection based on the number of hits can be operated.

In the online selection a peculiar feature of the indexing of pixelated Timing Counter is used: for 52.8 MeV positrons within the acceptance, cabling is such that the tile with the smallest ID is crossed first. So the smallest Tile ID over threshold is an estimator of the first impact point. Then, the time of the selected tile is used to drop all the hits within a programmable window, as they are assumed to belong to the same track.

According to experiment Montecarlo simulations, the biggest time span of hits belonging to the same track does not exceed 2 ns; it corresponds to positrons turning several times in the CoBRa volume, each time crossing the CDCH and hitting the TC at increasing $|z|$ positions. To be fully efficient on merging also those kind of events, two track candidates needs be selected each clock cycle; this way the $\sim$ 100 kHz positron rate can be fully analysed.

3.6.3 $T_{\mathrm{e}\gamma}$

An online estimator of the relative time is obtained by exploiting the timing of Liquid Xenon and Timing Counter channels which were selected during the reconstruction of the impact positions. As an effect of the either photostatistics or tracking geometry, those channels are the ones to provide the most precise timing information.

Positron time T_{e} is just the average time of both ends of the selected pTC Tile; including more tiles in the computation is made difficult by the various hit topologies of tracks observed inside the Timing Counter. Some events shows in fact non-contiguous hits due to positron passing between the Tile plastic scintillator and the aluminium support structure, which complicates hit clustering. This design choice fixes the online resolution to be the average of two 1.5625 ns bin TDCs, for a theoretical RMS of 319 ps assuming the quadrature sum of the two hits.

The photon interaction is localised within a 4x4 MPPC Patch so the averaging of the signals within the same patch results in a factor $\sqrt{16}/\sqrt{2}$ better resolution if all MPPCs are over threshold. As shown in data analysis of 2018 data (section 4), this condition is currently achieved only for very shallow energy deposits; the algorithm has to take care of a variable number of channel over threshold in the patch not to loose efficiency or resolution. A Patch multiplicity is calculated to cope with it and sent to the master TCB together with the sum of all the detected hit time in the Patch.

Photon time T_{γ} is the average patch time, which is the ratio of the sum of all time for each MPPC over threshold divided by the multiplicity; similarly TC time also need to be divided by 2 (the two channels reading out the Tile).

With some algebra the relative cut in photon-positron time $T_{e\gamma}$ can be transformed to an expression only in terms of sums and products:

$$
\begin{aligned}
&|T_{\text{Xenon}} - T_{\text{TC}}| < \text{Threshold}\\
&|\frac{\sum_{\text{chn. over Thr.}} T_{\text{chn}}}{\#_{\text{chn. over Thr.}}} - \frac{\sum_{\text{tile chn.}} T_{\text{chn}}}{2}| < \text{Threshold}\\
&\frac{2\cdot\sum_{\text{chn. over Thr.}} T_{\text{chn}} - \#_{\text{chn. over Thr.}}\cdot\sum_{\text{tile chn.}} T_{\text{chn}}}{2\cdot\#_{\text{chn. over Thr.}}} \pm \text{Threshold} \gtrless 0\\
&2\cdot\sum_{\text{chn. over Thr.}} T_{\text{chn}} - \#_{\text{chn. over Thr.}}\cdot\sum_{\text{tile chn.}} T_{\text{chn}} \pm 2\cdot\#_{\text{chn. over Thr.}}\text{Threshold} \gtrless 0\\
&2\cdot\sum_{\text{chn. over Thr.}} T_{\text{chn}} - \#_{\text{chn. over Thr.}}\cdot(\sum_{\text{tile chn.}} T_{\text{chn}} \pm 2\cdot\text{Threshold}) \gtrless 0
\end{aligned}
$$

This transformation is very important to map the logic inside the FPGA: modern chips contains hundreds of dedicated hardware multiplier and adders[94][8].

Since multiplication by powers of 2 is equivalent to bitwise shifts, the final selection requires only two adder/subtracter and a single multiplier for each time comparison.

3.7 Data Readout Scheme

The main drawback of having a full waveform digitisation is the huge event size: each WaveDREAM can generate up to 89 kB/evt as reported in table 3.2; having planned 528 WaveDREAMs in the MEG II system, the maximum data size is in the order of 46 MB/evt, if everything is readout, with a minimal DRS-only event being as much as 14 MB/evt.

Type of data		**size per event**
DRS Waveform	1024 samples × 12 bits × 18 channels	27.6 kB
ADC Waveform	2048 samples × 12 bits × 16 channels	49.2 kB
TDC Waveform	512 samples × 8 bits × 16 channels	8.2 kB
Trigger Data	512 samples × 64 bits	4 kB
Total		89 kB

Table 3.2: Amount of data being generated by a single WaveDREAM board.

The Microblaze processor in the WaveDREAM is not suited to handle this amount of data directly: samples from the readout FSM are stored in a local buffer capable of containing up to three complete events, to provide de-randomisation of the data stream and to compensate for short-term bursts of events. Another FSM packs the data with headers to be streamed to the ethernet interface (if the board is used in a standalone setup) or to the backplane connector, when used inside a crate.

---

[8] Performing division instead is a very demanding operation that is usually accomplished in multiple clock cycles

A similar schema is also employed in TCBs, so that pieces of information on the trigger being generated can be saved. The main difference is that each TCB is different from an other, each performing dedicated calculations. For this reason the data to be readout are different from TCB to TCB.

The selection of data to be saved in a TCB is managed by a programmable FSM (being almost evolved into a custom-designed CPU core). It is activated by a trigger and performs a list of operations (among those of table 3.3) moving the data to a set of four event buffers for the subsequent readout by the DAQ software.

Operation	**Operands**	**Description**
STOP	none	stops execution
COPY	two 32-bit addresses	copies a 32 bit value from the two points in the board address space
BLOCK_COPY	two 32-bit addresses and a 16-bit amount	copies a set of consecutive 32-bit words between two addresses
DIRECT_WRITE	a 32-bit value and a 32-bit destination address	write a given constant value to a specific address
JUMP	10-bit command address	jumps to a given command to be executed
JUMP_IF	32-bit address, 32-bit value, 10-bit command address	jumps to the given command address only if data at the address matches the bits requested

Table 3.3: Operations accessible to TCB readout FSM.

An header has to be attached by any board to data fragments before transmission; by design, the header consists in two parts with the first one being fixed and containing basic informations:

board Type discriminates WaveDREAMs from TCBs and among board revisions;

serial Number identifies uniquely a board in the system;

board position containing a Crate identifier and the position inside that crate;

packet Number is a progressive packet number identifying the packet;

packet Type to easily disentangle multiple data format;

flags marking the beginning and the end of both data types and event stream;

payload Length and Offset to mark the amount of data being transferred and their offset (when multiple packets needs to be combined together to form a larger block of data which could not fit in a single packet).

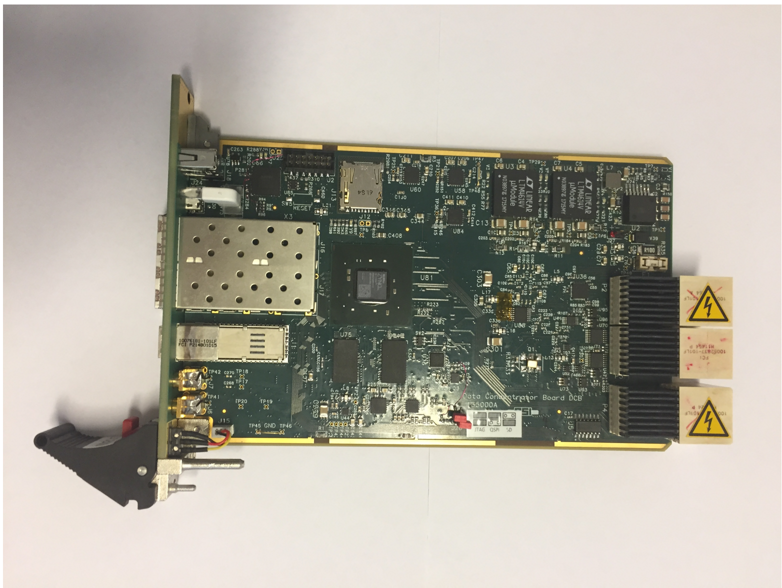

Figure 3.11: Data concentrator board.

The second part of the header contains board-specific information, such as the DRS trigger cell, sampling frequency and temperature, so that data are always paired with the board configuration being currently running thus assuring consistency. This part of the header also contains serialised trigger information (Event Number and Trigger Type) transmitted with the Trigger system, so that a unique key is available for each event.

Data coming from any board can be checked for completeness by comparing the number of packets received between the Start-Of-Event and the End-Of-Event Flags with the Packet Number in the header. Dropping of a packet can also be identified and tagged so that only complete events are saved. This feature helps in addressing the unreliability of ethernet packet transfer employing UDP packets; the protocol was selected to be used in data transport to the DAQ backend machine because of its simplicity and small overhead requested.

If data is sent on the backplane, serial links similar to the ones used in the Trigger computation are used; the Data Concentrator Board collects the packets from all the slots and forwards them to the main DAQ.

3.7.1 Data Concentration Boards

As described in section 3.4, the Data Concentrator Board (DCB), shown in figure 3.11, distributes the trigger signals and the clock to the boards in the crate. An onboard Zynq 7000 System-On-Chip[88] also receives data from all other boards in the crate and provides a connection to the DAQ.

The Zynq chip is a combination of a standard FPGA fabric with a two core ARM Cortex A9 CPU capable of running a Linux operating system. This unique feature allows the FPGA to perform operations such as handling and combining of the data streams, while the ARM CPU provides access and management for all the boards in the system. The DCB can indeed take the control of the SPI on the backplane to access all the other boards and configure them as well as to download new firmwares with a much higher speed than that possible with the CMB[9].

The double role of crate control and data concentrator is visible from the front panel connectivity: two ethernet connections are available, one to CPU connection, the second for a direct FPGA stream; DAQ ethernet flux can be physically separated from the configuration and the slow control access, thus improving reliability. The ethernet connection uses Small Factor Pluggable (SFP) cages so that the interface module can be easily swapped to be compatible with either copper or fibre connections, depending on the deployment needs (currently both this cages use the 1Gb Ethernet protocol). However the circuit board was designed with the option to install a Zynq with faster transceivers supporting 10Gb Ethernet connections on SFP+ modules. Alternatively the current boards could be updated, with only firmware modifications, to support 5Gb Ethernet.

3.7.2 Online computing and network infrastructure

Currently the DCB firmware is under development and the part handling data stream is not yet deployed for production. Each WaveDREAM is therefore directly connected to the PSI network with a direct ethernet connection. Few Cisco SG110 gigabit ethernet switches are used for the first level of aggregation ($\sim$ 1 crate), before sending data to a Cisco Catalyst 3200 from PSI network. While the SG110 can nominally switch up to 35.7 million packets per second (corresponding to a throughput of 48Gbps), which theoretically should be sufficient, packet drops were observed well below the saturation of the link. This was found to be due to packets being all generated with a fixed timing relative to the trigger on all the WaveDREAMs of the system: all the packets entered the network at the same time (down to 100 ns level) and this overwhelmed the instantaneous switching capacity of the first layer switches. Two programmable delays were then added to reduce the instantaneous throughput, the first delay shifts the very first packet of the event and the second one sets the inter-packet delay. By fine tuning these values, packets can be arranged in a round robin and the data-flow scaled up to 1Gbps, being limited by the current 1 Gbps ethernet interface on the Megon01 backend online machine.

Figure 3.12 shows the achieved performances of the current WaveDAQ system composed of 90 WaveDREAMs and 10 TCBs: data rate to disk scaled up to 1.1 Gbps with an event rate of 42 Hz, limited by the available network speed. These performances are possible also thanks to a newly developed multithreaded readout architecture which will be described in the next section.

---

[9] Currently the download of a Kintex 7 firmware image (11 MB in size) takes 15 minutes, being limited by the CMB speed driving SPI bus; the DCBs can instead be programmed to download locally the image, not requiring anymore the online machine to drive the writing operations, and will move data on the SPI bus $3-5$ times faster.

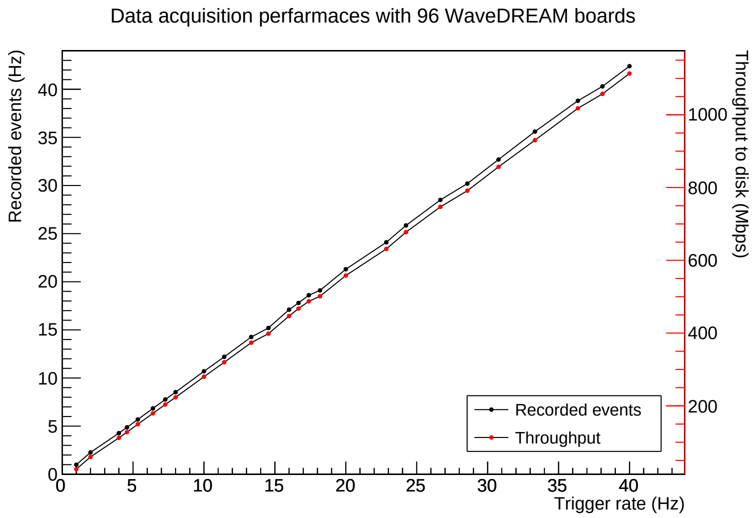

Figure 3.12: Observed Data Acquisition throughput during 2019 run for different trigger rates. Maximum throughput to disk was observed ad 1.1 Gbps at 42 events per seconds.

3.7.3 Multithreaded readout software

The WaveDAQ software library, handling WaveDAQ system configuration and data readout, was developed following an Object-Oriented Programming paradigm: the full system is represented by a "WDSystem" object containing a list of "WDCrate" objects, each abstracting a WaveDAQ crate. Finally up to 18 "WDBoard" object are stored in each WDCrate providing a base shared class for a generic electronic board in the system. Derived class of the WDBoard (linking additional external libraries when needed) are specially developed for WaveDREAM, TCB or DCB configuration and management.

The description of the system, in terms of board naming and positioning in a crate, is stored in an XML (eXtensible Markup Language) file[95], which also contains pairs of key-value to be used to describe parameters to be applied to boards. Having to handle several types of data, the value is stored in memory as a string of characters and is parsed to decimal and hexadecimal integers, floating point numbers or booleans depending on the needs.

In order to fully take advantage of modern multithread CPUs, the readout library for the WaveDAQ was developed by fully exploiting the standard thread library available since revision 2011 of the C++ standard[96]. The DAQ library provides a set of specialised threads and buffers and is linked by the WaveDAQ library to be operated simultaneously to the hardware. Hardware boards can be in a RUNNING or in a BUSY state, representing if the data-taking is in progress or not; similarly, each thread can be in either a run or stop state[10].

---

[10] An additional feature of having code executed at the transition from RUNNING to BUSY (data-taking run starting) or on the vice versa (data-taking stopping) is available.

The threads used for readout are designed so that the task being executed is the smallest possible, to exploit at its maximum the hardware parallelism. The threads are named after the function they perform:

Collector links to an UDP socket which receives packets from the WaveDREAMs and, in the future, also from TCBs and DCBs. It also decodes the packet header and, according to its contents, reads the packet content. Parsed packet data is pushed to an external buffer allowing, if needed, the possibility to have more Collector threads, listening on different UDP ports, to increase the throughput.

TCB reader is a temporary thread designed to poll the buffer state of a TCB and read it out simulating a TCB packet being received. This will be accomplished in the hardware by the DCBs in the future and, at that point, the TCB reader will be replaced by the Collector thread.

Event Builder collects the packets belonging to the same event according to the Event Number. The code is designed to cope with the mixing of packets coming from various events by having a local buffer containing events being built. When all the packets from every boards are received a WDEvent object is created containing all the received information. There are two conditions which may cause an event to be discarded: the oldest event is dropped if the local buffer exceeds its maximum capacity (by default 20 events) or any events older than a successfully built one (by default by 10 events).

Worker applies the voltage waveform calibration to provide the system with the best waveform for further processing, such as zero suppression algorithms.

Writer is an optional thread designed to provide a full DAQ system for small application and system tests: it writes data into a binary file.

Supervisor is the main thread which spawns the others and controls the RUNNING/BUSY state of the system.It also handles the configuration of the boards.

This architecture was shown to be capable of sustaining a rate of $\sim$ 40 Hz with 90 WaveDREAMs all transmitting the DRS Waveforms(figure 3.12), with a measured throughput of 995 Mbps on network and more than 140 MB/s to disk (400 GB/h). To put this last number in perspective, the full storage currently available for MEG II is 1.2 PB of filesystem paired with 400 TB of slower storage but will last roughly 4 month if this amount of data is continuously written without any zero suppression.

The clear bottleneck of the current library is the 1 Gbps ethernet connection; by scaling that limit to the final system of 518 WaveDREAMs, it would limit the maximum DAQ speed to $\sim$ 7 Hz. This may already suffice for the physics case of MEG II but the 10 Gb connection available on the readout computer will provide additional bandwidth to account for other trigger lines dedicated to other physics channels.

3.7.4 MIDAS data acquisition Software

The backend of the data acquisition is handled by the "Maximum Integrated Data Acquisition System" (MIDAS) package[97], developed by PSI and TRIUMF laboratories, with several application in medium-sized particle physics experiments (including the former MEG). Even if the first development dates back to early '90, the choice of building around several independent programs being combined by interprocess communications or network sockets proved to be well suited for multi-core CPUs and modern ethernet-based DAQs.

The main component is the "Online DataBase" (ODB) which allows all programs to share experiment configuration parameters and monitoring values, like thresholds, applied HVs and corresponding currents. "Frontend programs" are responsible to receive data from hardware, to pack it into Banks and to upload hardware configuration parameters described in the ODB. A Frontend can operate either by polling the hardware for data or by periodic trigger; it can also be used to interface to slow control systems that are not part of the main DAQ (HV modules, beamline parameters, CDCH gas system...). The "Logger" program writes the data to disk in a binary format which is then zipped to reduce the data size. The "LazyLogger" program finally moves the zipped files from the local 9TB disk to the offline storage.

The WaveDAQ library is linked into the main Frontend program, configures ODB to the boards and synchronises the thread and hardware readout state to the global MIDAS run. Data are packed into a structured Bank-based data stream, with each data identified by a 4-character string while being handled by the MIDAS infrastructure: all DRS waveforms are packed into a single bank (DRSV), separated by appropriate headers also carrying channel identification and the corresponding trigger cell. With a similar structure ADC, TDC and additional trigger pieces of information from WaveDREAMs are packed in ADCW, TDCW and TRGW Banks.

Having to process different algorithms, the data coming from TCBs are board-dependent and pre-packaged by the programmable FSM in the firmware, together with the four characters to be used as Bank name. Currently only the final trigger board produces Banks; the most important one is the TRGI which contains basic event information to be transmitted at every event, such as the list of fired triggers, the Event Number and the total and live time. The TRGC Bank contains trigger scalers and the proton current scaler used to monitor the number of stopped muons and the overall detector and trigger stability. Finally a dedicated TXEC and TSPX Banks contain the estimators for the Liquid Xenon detector and Timing Counter observables respectively.

3.7.5 Offline software and cluster structure

The offline event processing is build on top of the ROME framework (Root based Object oriented Midas Extension) that has beed developed to interface MIDAS with CERN ROOT libraries during the previous phase of MEG experiment[98].

ROME generates a full set of C++ classes to provide event data processing in a scheme defined in an XML file. That file contains the definition of all MIDAS Banks, so that it generates all the required formatted readout from the data file; it also contains a specification of the output ROOT files and data TTrees together with a set of chained Tasks representing the various analysis steps. Each Task can be enabled or disabled at runtime without need to rebuild the code.

The ROME-generated analysis software can use other sources from MIDAS output file: it can also connect directly to the experiment MIDAS data stream, to perform online analysis for shifters, or it can read data from existing ROOT files, so that the same analysis can also be run on simulated data.

All the offline event processing is accomplished by a dedicated cluster of 8 computing nodes (320 cores) paired with 1.2 PB of IBM GPFS (General Purpose File System) data storage with an additional 400 TB of slower NFS (Network File System) storage.

3.7.5.1 Offline software structure and waveform simulation

Being based on the ROME framework (which eases the overall code management and data input/output), the offline software for the MEG II experiment consists of three programs[99]: “Gem4” event simulator, “MEGBartender” mixing software and “MEGAnalyzer” for the main analysis.

Gem4 is a simulation of the full MEG II experiment based on Geant 4 toolkit[100]. It can simulate signal, RMD and standard muon decays as well as the other calibration events. The energy deposit calculated by Geant4 is converted into a number of optical photons (accordingly to the light yield) to be tracked in Liquid Xenon and in Timing Counter Tiles. If the energy is deposited in the CDCH volume, helium ions are generated to drift inside the CDCH cells according to models from dedicated simulations based on the Garfield libraries[101].

Photon and Ions hit time in each MEG II sensor or wire are then passed to the MEGBartender, which folds them with a detector-specific response to finally generate a waveform for each readout channel. The software is named from the additional feature to mix events from different Gem4 simulations; the MEGBartender is then able to simulate the detector response for various pileup levels without the need of computing a full Gem4 simulation for each pileup condition.

Finally data are passed to the MEGAnalyzer which performs the event reconstruction either from MEGBartender generated waveforms or from real MIDAS data. The same program, when directly attached to the MIDAS data taking system, provides the online event monitoring for shifters.

3.7.5.2 Trigger offline processing

Trigger information can either be generated by MEGBartender software using a rebinning of the DRS waveforms or fed into the MEGAnalyzer from the MIDAS file.

The “ReadTRG” and “ReadData” tasks read and unpack MIDAS Banks into ROME data structures when analysing MIDAS data. Two tasks can then access waveform data: “TRGWaveformAnalysis” searches and selects features of the ADC, TDC and TCB waveforms to be saved for subsequent macro-based analysis; “TRGSimulation” calculates waveforms of interest in the trigger chain starting from ADC and TDC data that are the input of the trigger calculation also in the hardware.

The C++ calculations made in TRGSimulation is not designed to have a 100% correspondence to the hardware one: simulating all the calculations of the trigger at the register level would have a very high calculation cost. Only a "logical" simulation is performed paying attention to causality while processing waveforms: data at later samples must not affect samples being calculated earlier since this is impossible in the actual hardware. The option of having a full register-level simulation is currently under discussion.

A "TRGOfflineReconstruction" Task is foreseen for the future to carry out offline-level analysis from TRGWaveformAnalysis reconstructed hits, being helpful as crosscheck for the main DRS-based analysis.

3.7.5.3 Trigger monitoring tabs

A set of ROME monitoring graphic interfaces were developed to be used to check trigger performance both offline or during shifts: the "TRGTriggerRates" of figure 3.13a shows the rate of all 64 triggers together with the total and live time of the run and the Event Number of the event being processed. All of the scalers can be configured to calculate the differential rate in order to easily identify drops of beam intensity which can sporadically occur.

Detector-specific tasks for trigger monitoring were also developed, like the "TRGXEC-Monitor" of figure 3.13b which can display various waveforms both from the hardware or the simulation software; FFT functions are available to monitor noise level.

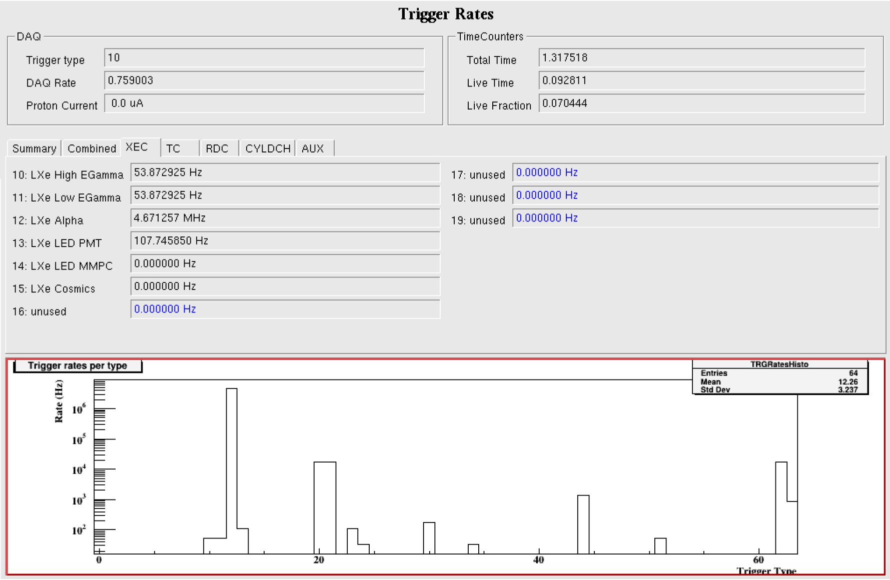

Trigger Rates

DAQ
Trigger type: 10
DAQ Rate: 0.759003
Proton Current: 0.0 uA

TimeCounters
Total Time: 1.317518
Live Time: 0.092811
Live Fraction: 0.070444

10: LXe High EGamma: 53.872925 Hz
11: LXe Low EGamma: 53.872925 Hz
12: LXe Alpha: 4.671257 MHz
13: LXe LED PMT: 107.745850 Hz
14: LXe LED MMPC: 0.000000 Hz
15: LXe Cosmics: 0.000000 Hz
16: unused: 0.000000 Hz
17: unused: 0.000000 Hz
18: unused: 0.000000 Hz
19: unused: 0.000000 Hz

(a) Trigger rates

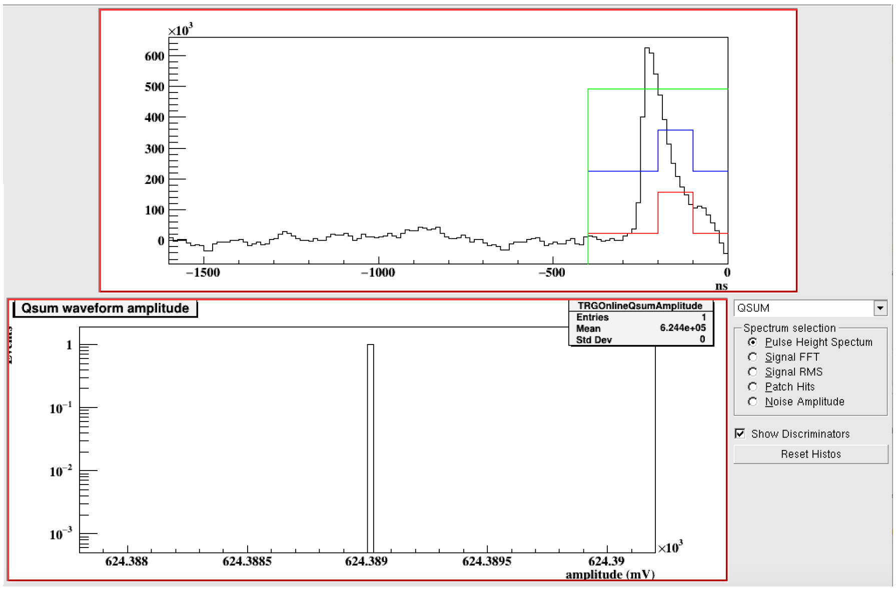

QSUM

Spectrum selection
● Pulse Height Spectum
○ Signal FFT
○ Signal RMS
○ Patch Hits
○ Noise Amplitude

☑ Show Discriminators

(b) Trigger xenon monitoring tab

Figure 3.13: Trigger online monitoring interfaces.

4 Run 2018

In this chapter a preliminary estimate of the online resolutions is given. At the same time a study of the experimental conditions was carried out resulting in improvements to the online event selection, which proved to be fundamental in defining the final trigger algorithms reported in section 3.6.

The 2018 data taking was a preliminary test to gain experience in operating the various systems, since for the first time all the detectors were installed.

The muon beam availability started at the end of November 2018 and lasted three weeks until the end of December. Michel and RMD events were recorded; the lithium Cockroft-Walton calibration line was used for the first time in MEG II; the Liquid Xenon detector response was checked, and Cosmic Ray events were used to measure the CDCH gas gain.

The DAQ electronics available at that time was limited to 96 prototype WaveDREAM board for a total of 1536 readout channels. they were distributed as follows:

- 16 boards for pTC, instead of 64 for a complete readout;
- 64 boards, corresponding to 1024 channels, to Liquid Xenon detector; 640 for MPPCs and 384 for PMTs, less than 1/4 of the total;
- 12 board to digitise 192 CDCH signals, 96 cells out of 1728;
- 1 board for the CDCH cosmic ray counters and 1 board for studies on a new preshower for the BGO detector.

with the remaining channels used for auxiliary signals such as the Liquid Xenon LED and pTC Laser trigger pulse, coming from the two respective generators.

Since the CDCH-specific differential WaveDREAM boards were missing , differential to single end custom converter circuits were employed to readout drift cells from both ends of the wires using the single-end WaveDREAM boards. Unfortunately this additional converter card generated a ground loop adding a 10 mV sinusoid on top of the CDCH signals, which were of the same order of magnitude.

Finally the DCBs (described in section 3.7.1) were still being developed and could not be included in the system: they were replaced by a modified Ancillary board that just distributed the trigger signals to the WaveDREAM boards; each card was directly connected to the ethernet network for readout.

The multithreaded readout software, described in section 3.7.3, was not yet adapted to be used inside a MIDAS frontend and this limited the maximum DAQ rate to 3 Hz for the DRS-only events and below 1 Hz when ADC and TDC information were also readout. The limited rate and beam time resulted in a very reduced dataset to be used

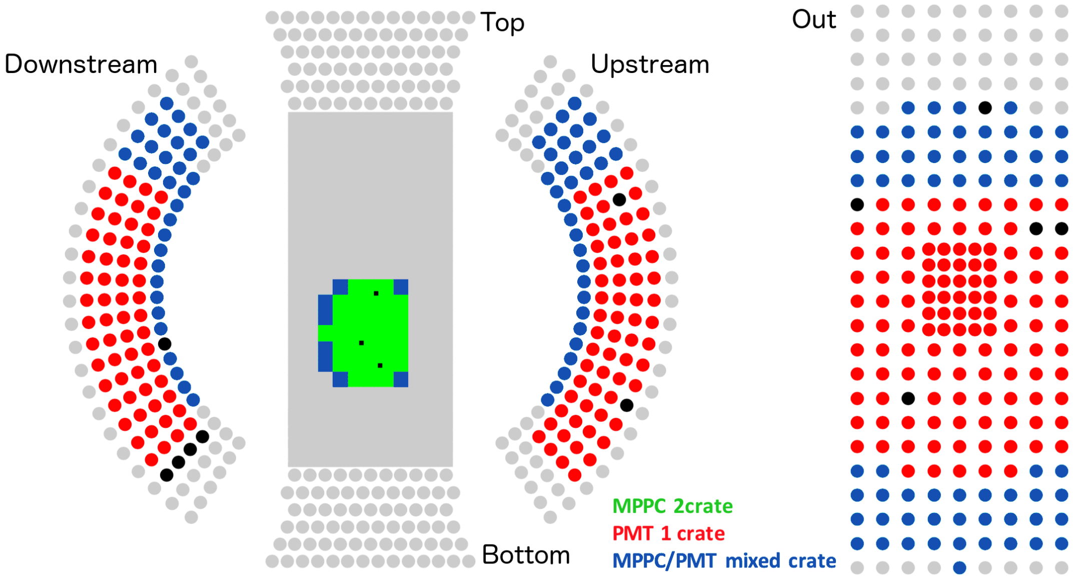

Figure 4.1: Liquid Xenon channels connected in run 2019, a total of four crates were connected, the blue MPPCs were not used online.

for initial performance estimates, a longer and more stable 2019 run promises an enlarged dataset for further studies.

4.1 Detectors and other items involved in 2018 Run

The Liquid Xenon detector channels were connected to four WaveDAQ crates as indicated in figure 4.1. Two crates were entirely dedicated to MPPC and PMT, the third one was a mixed half-MPPC half-PMT crate optimising the few channels for the offline analysis purposes. At the trigger level the first two crates were used as planned while the third one only partially: the PMTs were added to the E_γ computation while the MPPCs could not be exploited

The online photon conversion point estimate, described in section 3.6.2.1, was used to select events hitting the central part of the instrumented inner face, 3×3 Patches wide. This was pretty successful having only few events hitting outside the area of interest due to fluctuations mimicking a peak in the instrumented area. Once the full inner face will be instrumented such category of events could not exist anymore because the correct conversion point will be selected instead.

Considering that only 1/4 of the TC Tiles could be connected to readout boards, the high Z part of the downstream module was selected as shown in figure 4.2. All these channels could be illuminated by the laser system to provide a stable time reference for the offline performance checks[1].

---

[1] the innermost 48 Tiles (one third of the inner half of a module) in fact do not have the laser fibre

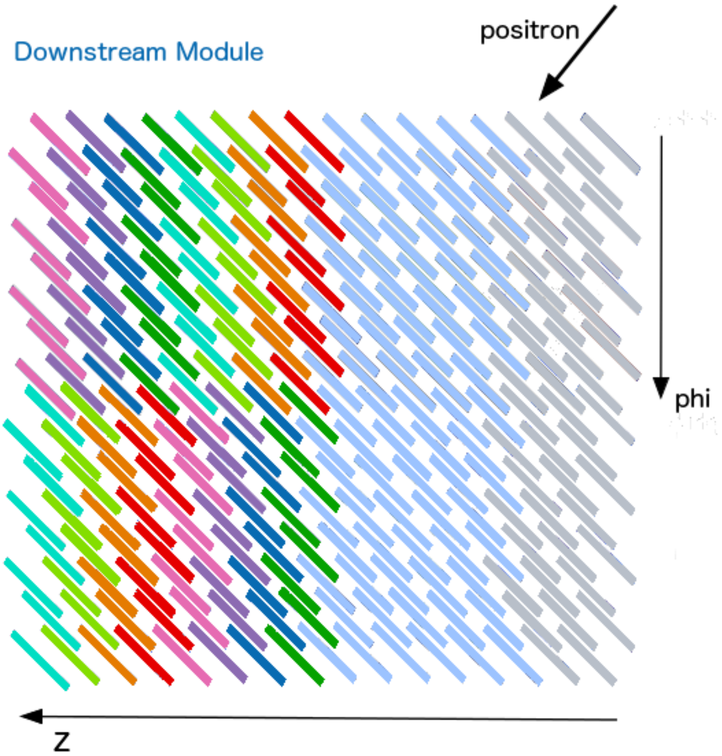

Figure 4.2: TC channels connected in run 2019, only downstream module is shown. Each colour is connected to a different WaveDREAM board. Light blue Tiles are not connected and grey ones identifies counters without the laser fibre.

The CDCH was also operated in the 2018 run, with the outermost four layers with a bias voltage large enough ($\sim$ 1500 V) to have good signal amplification while the inner ones had a lower voltage due to electrostatic instabilities. The high voltage map used in preparation for the run is shown in figure 4.3a with blue cells being operated at a reduced voltage and white ones being disconnected. Unfortunately a cathode wire broke inside the CDCH volume during the run and the high voltage was subsequently reduced to the configuration reported in figure 4.3b. In the same image the channels that were readout are marked with a yellow circle. Nevertheless a characterisation of the detector gain was possible which resulted in a need for a further stretch of the chamber to reach an operational level (gas gain $\sim 10^5$). It was in fact established that such needed high voltage value ranged between 1500 V for the innermost cell layer and 1580 V for the outmost ones. The CDCH was later tensioned so that 100 V safety margin exists, on top of such operational high voltage level, before observing any electrostatic instabilities; the safety margin will account for detector ageing effects and for possible variations at 1% level of gas mixture or external atmospheric pressure.

---

inserted because of the reduced space between the scintillators and the aluminium support structure.

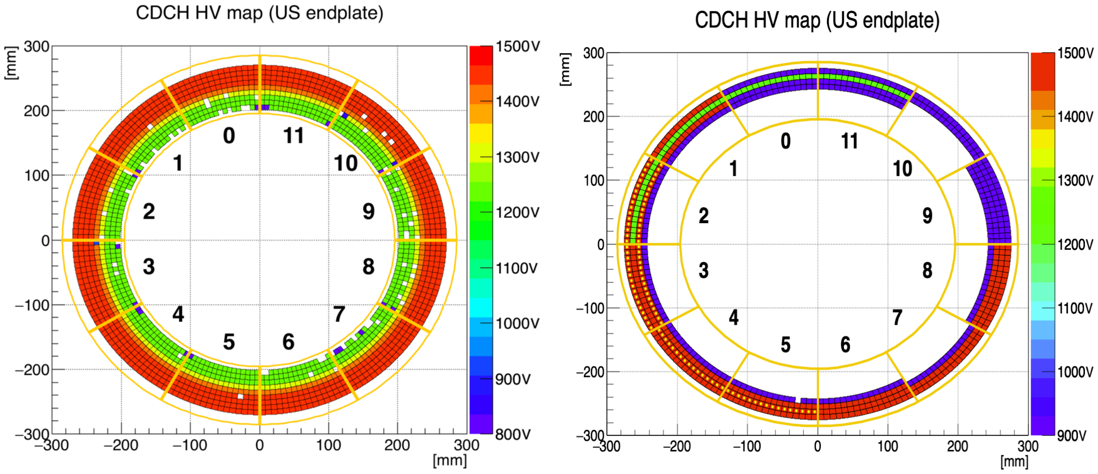

(a) Nominal

(b) After one wire broke during the run

Figure 4.3: CDCH high voltage settings in run 2018 and readout channels (yellow circles).

With signals few 10 mV on top of a 10 mV noise from the converter boards any trigger based on CDCH signals was not possible; therefore external trigger sources were needed. For Michel events from muon decays the Timing Counter was used, while, to select cosmic rays, a set of eight cosmic ray counters (CRC) scintillator bars were used. Four counters were placed above and four below the CoBRa magnet; the CRC trigger required the coincidence of one top counter and any of the bottom bars.

4.2 Calorimeter E_γ trigger

The offline-reconstructed spectra of the Liquid Xenon detector are shown in figure 4.4, with the units being “detected photons” because a calibration of the energy scale is not yet established. Two different spectra are shown because the Trigger boards were configured to generate two different trigger conditions having different thresholds on the weighted sum of the sensor waveforms, which corresponded to roughly 30 and 40 MeV respectively.

The two conditions were operated simultaneously: the Trigger Type and the list of fired triggers were attached to the readout data stream. During the offline processing the events could be correctly identified and filtered according to the necessities.

Considering that the low threshold condition, being less selective, had a higher trigger rate, it was prescaled by a factor 400 during the data-taking. In the offline comparison of figure 4.4, the relative weight of the two categories of events was rescaled accordingly.

By taking the ratio of the two histograms, shown in figure 4.5a, the selection efficiency of the high threshold can be obtained and fitted with an error function[2]: the resulting

---

[2] This procedure actually performs a relative efficiency measurement with respect to the low energy

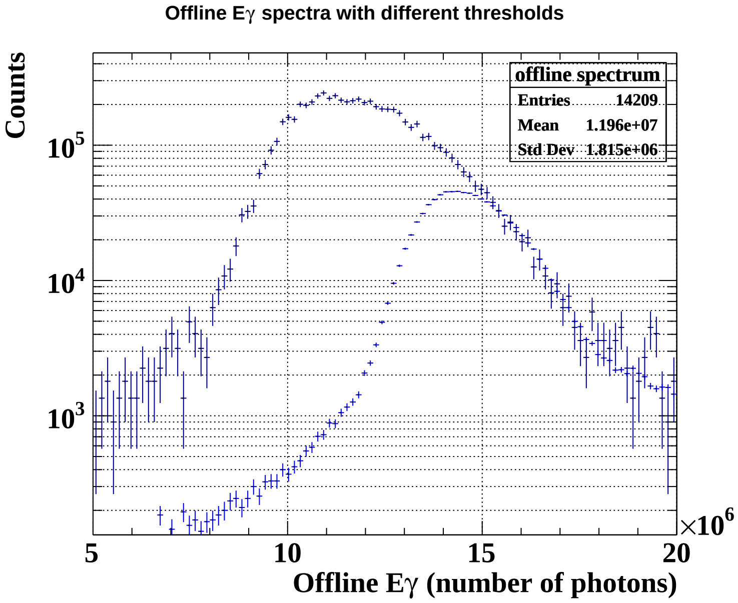

Figure 4.4: Offline-reconstructed Liquid Xenon detector spectra at two different trigger threshold with muon beam on the target.

cut value is $1.4 \cdot 10^7$ detected photons and the resolution is 7%.

The two data samples were recorded simultaneously but by means of two independent online selections, this results in independent bin by bin fluctuations in the two reconstructed spectra. Especially at the endpoint, where the amount of collected data is very sparse, fluctuations can cause the low threshold spectrum to record less events than the high threshold one, resulting in a fictitious efficiency above 1 in the ratio.

Using this same analysis, while taking physics data, the online high threshold can be monitored both in performance and in cut value to assess the data quality and highlight any drift of the apparatus.

There are few open issues that could explain the measured 7% resolution being worse than that expected by Montecarlo for the full calorimeter (a study of the expected MEG II online resolution is given in section 5.1.1):

- The contribution to the energy deposit from channels not being connected to the TDAQ system, only partially mitigated by the selection of photon showers centred in the instrumented part of the inner face.
- The higher than expected level of electronic noise on both DRS and ADC data in

---

threshold which is assumed as minimum bias trigger. The complete efficiency of the trigger chain should include a confirmation of such minimum bias trigger. This is possible by comparing the recorded spectrum of Michel endpoint and various calibration sources with the expected ones from simulations. Those studies were carried out for the MEG I experiment and showed no evidence of trigger bias effects, similar checks will be repeated in MEG II when the complete readout will be available.

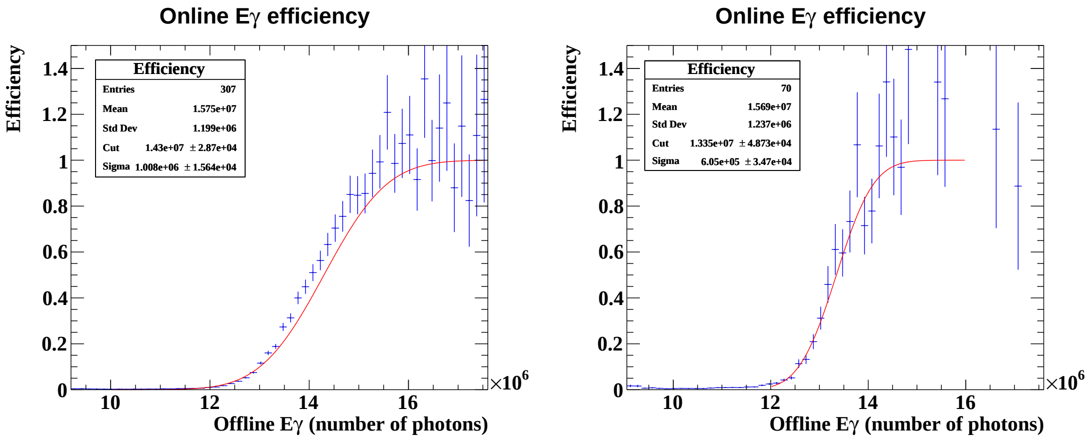

(a) Gain 1

(b) Gain 2.5

Figure 4.5: Trigger response functions with an analog frontend gain of 1 or 2.5.

the prototype WaveDREAM.

- The not yet optimised weights used in the sum to correct for gain and quantum efficiencies.
- The data transmission inside the Trigger system not correctly compensated for transmission latency. When summing data from different sources latency variation can result in a random shift by one clock sample in the ADC summing.
- The smaller than expected collection efficiency of MPPCs on the inner face.
- An unknown part of the observed fluctuation could be due to the the offline energy being also reconstructed with worse than expected precision.

Since the light collection efficiency on the inner face was lower than expected it was partially compensated, for a small part of the run, by an analog gain of 2.5 being selected for the WaveDREAMs connected to MPPCs. The threshold function, obtained with the same procedure discussed above and shown in figure 4.5b, gives a resolution of 4% in this condition. The cut value measured with 2.5 gain slightly moved to $1.3 \cdot 10^7$ detected photons, the difference from gain 1 measurement is due to the rough scaling of trigger sum weight to account for the increased input gain.

A crosscheck analysis can be performed using only the events triggered with the low threshold: a subsample also firing the high threshold trigger can be selected using the list of fired trigger that is extracted from the final trigger board, resulting in the two spectra of figure 4.6. This analysis has the advantage to use the same dataset for both the triggered and untriggered sample, avoiding the bin fluctuations previously quoted. However this procedure is affected by the small amount of low threshold event that could be recorded in 2018.

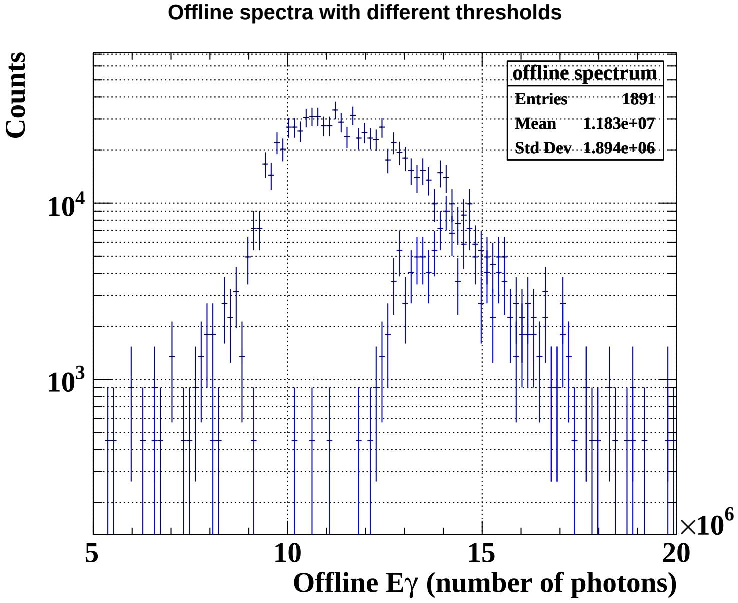

Figure 4.6: Example of spectrum selection using the list of fired triggers.

4.2.1 CW Lithium calibration gamma lines

In the 2018 run, the Cockroft-Walton accelerator was operated for the very first time within the MEG II setup and the generated lithium calibration line were used for online noise studies: two independent datasets of 10000 events with ADC readout enabled were recorded, one with the DRS digitisation running in parallel and the other with the DRS chips switched off.

The online values that were selected for the two datasets are reported in figures 4.7. The one in figure 4.7a has a very poor resolution of 10% due to the noise injected on the ADCs by the DRS sampling. By contrast, that in figure 4.7b shows a vert nice resolution of 3.4% on the 17.5 MeV line and an additional bump is visible at smaller energies due to the broader 14.6 MeV line.

The values of the two peaks are $1.98 \cdot 10^5$ and $1.69 \cdot 10^5$ detected photons respectively, their ratio is compatible with what expected.

By scaling the observed resolution at 17.5 MeV to 52.8 MeV according to the inverse square root of the energy, a resolution of 2% for the signal photon is expected if DRS were not running without any further optimisation being attempted. Physics analysis requires the DRS digitisation to achieve good timing and energy resolutions, the studies of the performance of an ADC-only data-taking showed however that a large amount of noise was induced by DRS sampling.

In figure 4.8 an example of two ADC waveforms, coming from the same MPPC, is shown for the two datasets where the height of Cockroft-Walton events have the same amplitude. A peak-to-peak noise of eight ADC-unit (2 mV) is observed in the case of the DRS being operated that is reduced to below one mV in the ADC-only example.

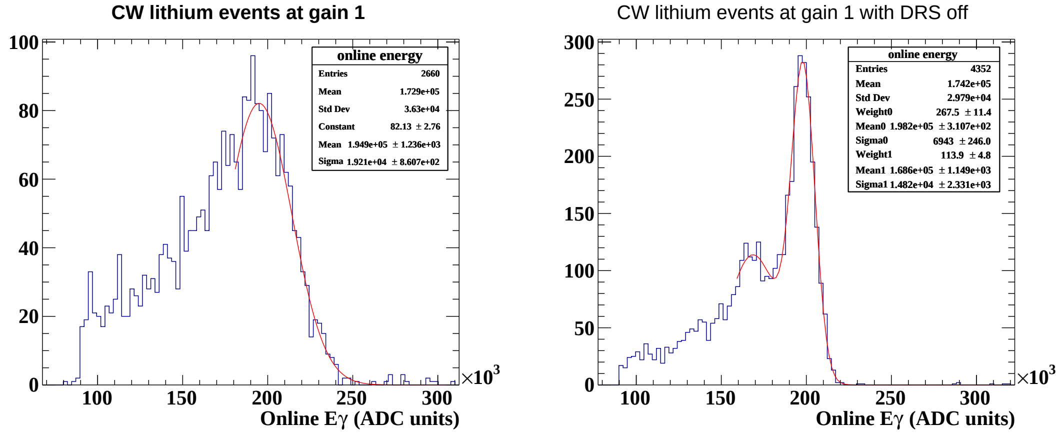

(a) DRS4 readout on.

(b) DRS readout off.

Figure 4.7: Cockroft-Walton lithium online spectra.

The observed noise clearly shows a periodic coherent structure which is present not only within a single WaveDREAM but, with the same phase and amplitude, across the the full system. This fact can be understood since it was later recognised to be caused by the DRS4 digitisation process and all the chips in the system run synchronously by design to achieve the required timing performances.

Following this study, and the corresponding offline related ones, some flaws in the prototype WaveDREAM design were identified and could be addressed before the final production of more than 500 boards. In particular, the two main noise sources will be fixed by a larger inductance providing the DRS output common mode and a low-pass RC filter in between the DRS and the ADC.

Sixteen pre-production boards were delivered in early October 2019 and were employed in 2019 data-taking; a complete comparison of the new board performance is ongoing. A very preliminary test is reported in figure 4.9: a single WaveDREAM was configured to register pedestal events, in the plots the RMS of the ADC waveform is plotted as a function of number of channel being summed up. There can be two contributions to the sum, the uncorrelated noise will sum up with the square root of the number of channels due to its statistical nature, on the other end a completely coherent noise will scale linearly with the number of channels being summed. The two noise sources can thus be measured and compared by fitting the scaling of the sum RMS with the number of channels, in the fit an additional term independent on the number of channel is added for completeness.

The coefficients resulting from the fits are shown in table 4.1. The constant term is always so small to be negligible when few channels are summed up. The non coherent component is at level of few 100 μV per channel and slightly improved with the new board revision, not far from the expected level for a 12-bit ADC with 2V dynamic range[3].

---

[3] The ADC bin size is 488 μV, the expected RMS is then 141 μV not including other noise sources

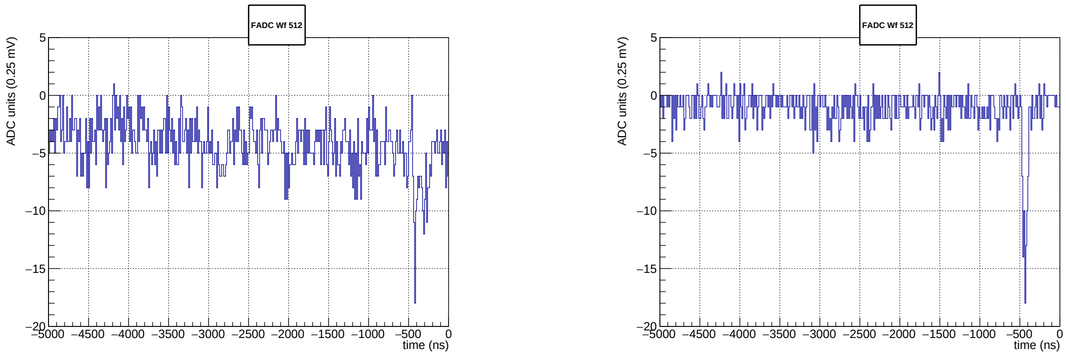

(a) DRS4 readout on.

(b) DRS readout off.

Figure 4.8: Cockroft-Walton lithium ADC waveforms with two different states of DRS readout.

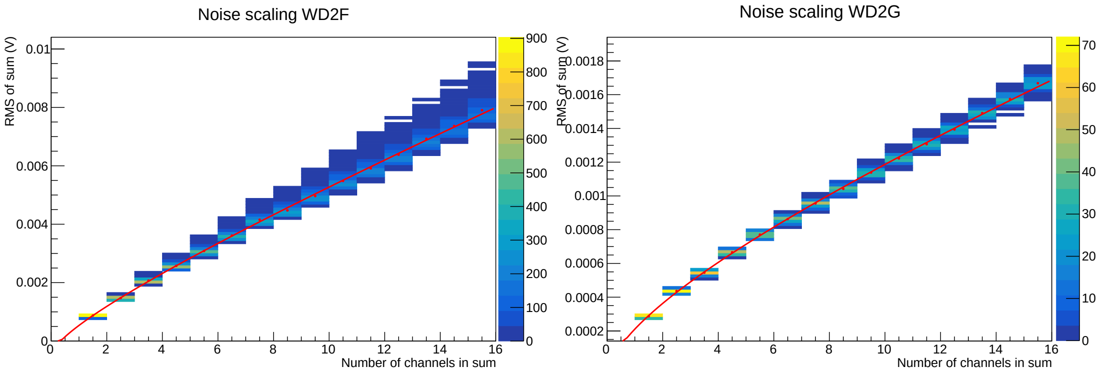

(a) Prototype WaveDREAM

(b) Pre-production WaveDREAM

Figure 4.9: Scaling of RMS noise depending on the number of summed channels.

WaveDREAM revision	**constant**	**non coherent** $\propto \sqrt{N_{\text{ch}}}$	**fully coherent** $\propto N_{\text{ch}}$
prototype	~ 0 V	$5 \cdot 10^{-4}$ V	$3.9 \cdot 10^{-4}$ V
pre-production	$4 \cdot 10^{-5}$ V	$2.0 \cdot 10^{-4}$ V	$5.6 \cdot 10^{-5}$ V
pre-production, DRS off	~ 0 V	$2.8 \cdot 10^{-4}$ V	~ 0 V

Table 4.1: Scaling of the different noise components in the various WaveDREAM revisions. All coefficients are bounded to be positive duirng the fit; if the obtained value is $< 10^{-6}$ V the contribution is negligible and ~ 0 V is reported in the table.

---

beyond the quantisation noise.

The coherent component reduced more than an order of magnitude, such improvement have a huge impact when scaling to the final 4778 channel sum: the final sum, which would have been as high as 1.9 V with the prototype WaveDREAM boards, now is expected to be 280 mV. A caveat should be added on this simple extrapolation to the final channel number: the multiplicative weights are not taken into account which will introduce a non trivial contribution in the sum of the different channels. In addition the pedestal subtraction logic may be tuned according to the noise of the new pre-production WaveDREAMs providing some degree of noise cancellation.

Considering that the single channel noise is now very small a second-level noise subtraction on the final trigger sum, where the noise is higher, will be considered in the future.

4.3 Photon conversion point estimation

During the 2018 run, the photon conversion point estimation was not computed using the reconstruction described in section 3.6.2.1. The idea at that time was to use the same amplitude-based algorithm used in MEG: the MPPCs in the same Patch were summed up and the one with the highest pulse height was selected.

The reason for the new multiplicity-based algorithm is due to how the Liquid Xenon online time trigger was expected to operate. This includes the requirement that the online timing measurement must be fully efficient also for deep events as explained in the next section.

4.4 Liquid Xenon timing

The Liquid Xenon timing was expected to be extracted by the average of all 16 channels in the Patch selected by the photon conversion estimation, assuming that all of them would trigger the hardware comparator producing a corresponding timestamp. With the noise situation observed in the 2018 run, a high threshold of 30 mV was the smallest possible value with all the channels above the noise level[4].

With such high threshold, only a very limited subset of the events had all the 16 channels above threshold (multiplicity distribution are shown in figure 4.10 for gain 1 and gain 2.5) resulting in a huge efficiency drop if such assumption was kept.

In order to solve this issue, the final algorithm of 3.6.2.1 will propagate, together with the sum of timestamps, the multiplicity of channels. In this way we will be able to select even deep events with only few channels above threshold, they will anyhow have a worse timing resolution having less channels with which perform an average.

The 5 bits necessary to contain the multiplicity value could not fit in the 2018 trigger data stream if the amplitude-based position estimation had to be pursued. The solution was found by using the multiplicity itself to estimate the conversion point, no information was needed from ADCs therefore reducing the amount of data to be propagated.

---

[4] An hardware limitation of the prototype WaveDREAM made the discrimination twice as sensitive to voltage offsets of the electronics components being used in the discrimination process.

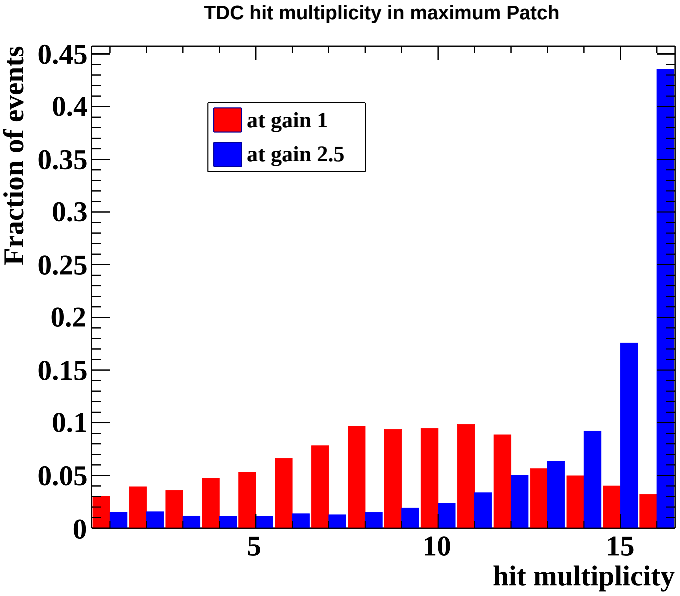

Figure 4.10: Liquid Xenon detector maximum hit multiplicity within a Patch.

The final algorithm have the positive side effect to also select the Patch with the best timing whereas the old one might have selected a suboptimal Patch: the one with the highest pulse height may not actually be the one with the biggest number of hits.

The multiplicity-based algorithm also has a smaller latency to select the conversion point, because the logic does not need to wait for the ADC to process the analog signal and calculate the digital value.

4.4.1 Performances

The performances of the novel online impact position reconstruction (multiplicity-based) needed to be carefully checked with respect of the former amplitude-based one.

During the data-taking 60000 muon-induced photon events were registered with a WaveDREAM input gain of 1 including, in addition to the usual DRS waveform, also the ADC and TDC ones. For such events the ADC and multiplicity based estimators can be computed offline and compared with the official offline reconstruction that calculates the conversion position by using the barycentre of the reconstructed number of photons.

A first sanity check of the algorithms is reported in the plot of figure 4.11 where the offline reconstructed position is coloured differently according to the result of the multiplicity based selection. A clear clustering of events is shown proving that the multiplicity estimator is actually sensitive to the photon position and a similar plot can be made also in the case of amplitude-based conversion point estimation. In figure 4.11, the events being reconstructed in the halo (outside -7 cm $< Y <$ 15 cm and

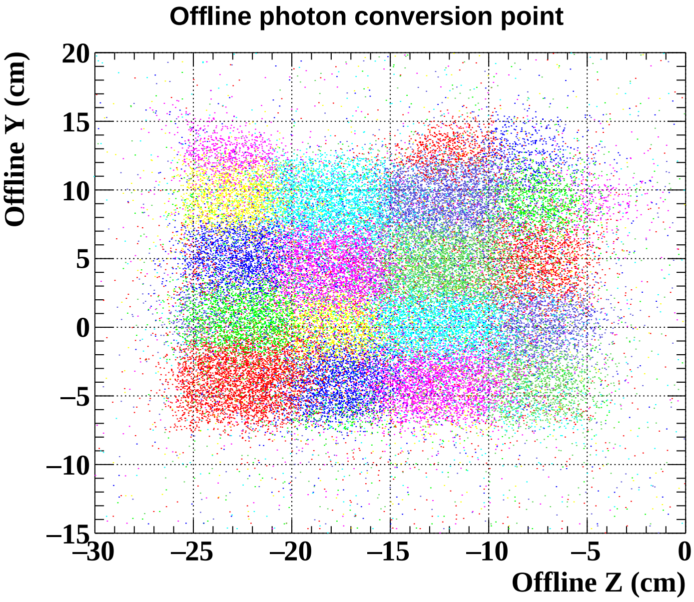

Figure 4.11: Offline reconstructed position for different online estimates, shown with different colours.

-25cm $< Z < -5$ cm) are due to photons impinging far from the instrumented area[5].

The Y and Z projection of a fixed Patch ID of figure 4.11 are reported in figures 4.12 for the multiplicity based reconstruction and the same plot for the amplitude reconstruction is overimposed to perform a direct comparison of the two algorithms. The two amplitude and multiplicity based reconstructions have a very similar behaviour, showing for both algorithms the core of the events being located within ±3 cm from the average value in each direction. This is compatible with a flat resolution within the patch (that is 6×6 cm^2) with a gaussian tails at the edges.

Even considering the largest dataset (gain 1) the number of events for each Patch ranges from 3000 events in the central region to few hundreds in the surrounding Patches where only spill-over events were recorded. This severely limits the possibility to extract the resolutions on a single Patch level. To avoid artefacts from spill-over events the study is then limited to the central 2x2 Patch group; the corresponding distributions are fitted with a double error function extracting the position of the two edges. Using the average of the two edges as estimate of the centre of the Patch, the Z and Y distributions of the 4 central Patches are aligned and summed up obtaining the plots of figure 4.13a and 4.13b for the Y and the Z direction respectively. The obtained distribution is composed by a step-like shape of well reconstructed events on top of a continuous background of badly reconstructed ones. No major difference is observed for the two algorithms and the resolution is evaluated as the sigma of the error functions; such value changes from

---

[5] In the offline analysis events outside the instrumented area on the inner face are correctly reconstructed because also the PMT on the lateral faces are used for the barycentre calculation.

$\sigma_Y^{\mathrm{ADC}} = 0.79 \pm 0.04$ cm, $\sigma_Z^{\mathrm{ADC}} = 0.74 \pm 0.03$ cm using the amplitude-based trigger to $\sigma_Y^{\mathrm{TDC}} = 0.84 \pm 0.04$ cm, $\sigma_Z^{\mathrm{TDC}} = 0.77 \pm 0.04$ cm. It is worth mentioning that the observed sizes of the active area, defined as the distance between the two edges, looks smaller than the expected but does not change with different trigger logic. It is in fact observed to be 4.4×5.1 cm^2 in comparison with the physical size of a Patch being 6×6 cm^2. This effect could be explained as an overall calibration effect or as a few millimetre efficiency drop at the edge of the two patches. In this analysis at least a channel over threshold is indeed requested in order to select the multiplicity-based Patch and, given the gain 1 used, events with 0 hits are removed. A more complete analysis will be performed in the future.

4.5 Track recontruction with Timing Counter

The online positron track selection and time measurement in the pTC are as important as the reconstruction of the photon variables: in 2018 the pTC events were triggered by requiring any of the Tiles being above threshold.

The trigger computed quantities are, also in this case, compared to the offline computed ones. Timing from the DRS waveforms is obtained by means of the standard MEG II analysis involving a constant fraction discriminator at both ends and the subsequent average of the two values giving the Tile timing.

Such offline reconstructed values are compared with the online extracted ones for the Tile selected by the online impact position selection. In figure 4.14b the time difference is reported as a function of the Tile ID showing two distinct groups of Tiles. That effect is caused by an improper distribution of the synchronisation and trigger signals in the 2018 TDAQ setup, mainly due to the lack of DCB boards. They can be clearly recognised because the jump in the reconstructed time is exactly 16 TDC units, corresponding to a system clock cycle.

In addition to the jumps, figure 4.14b also shows a clear decreasing slope in the reconstructed offset, at about ~ 4 TDC units (~ 3 ns) on the full range of Tiles used in 2018. That is due by the signal path in the PCB, holding the Tile in position on the module structure, not having the same track length for all Tiles and thus requiring in the final system some channel-to-channel online offset cancellation. The Tiles having the higher ID are indeed the ones located at larger $|z|$ and, because of this, they are nearer to the end of the pTC where the cables to the TDAQ are connected.

After compensating for this Tile ID dependent time offset the hits align and the resulting distribution is shown in figure 4.15, an overall hit resolution of 554 ps is obtained.

Since the channel resolution is 1.56 ns$/\sqrt{12} = 450$ ps, the theoretical coincidence of both ends should be a factor $\sqrt{2}$ better resulting in 318 ns.

The difference between this value and the obtained resolution can be attributed, in addition to the omnipresent noise, to a suboptimal circuitry being used in the prototype WaveDREAMs to perform the signal discrimination[6]. Nevertheless the obtained

---

[6] Only half of the signal amplitude is used in the prototype WaveDREAMs in the comparison resulting in a worse signal to noise ratio.

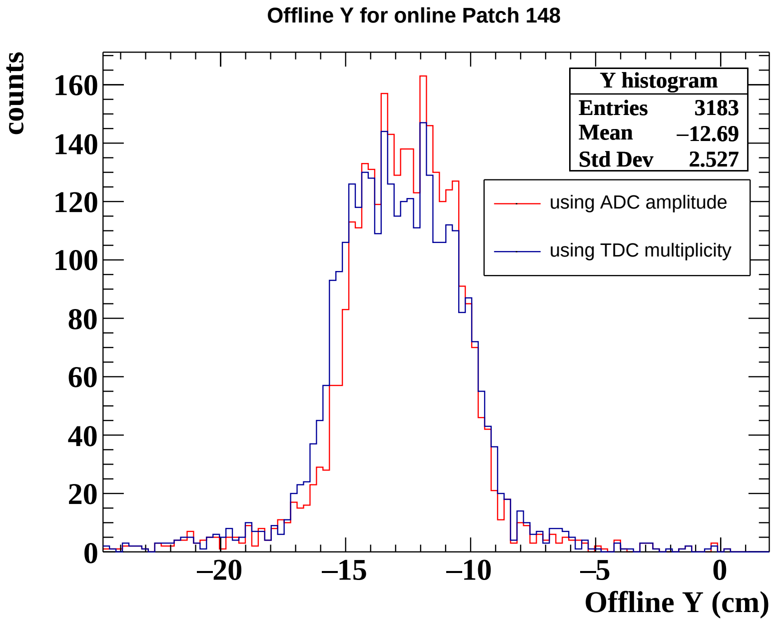

(a) Y

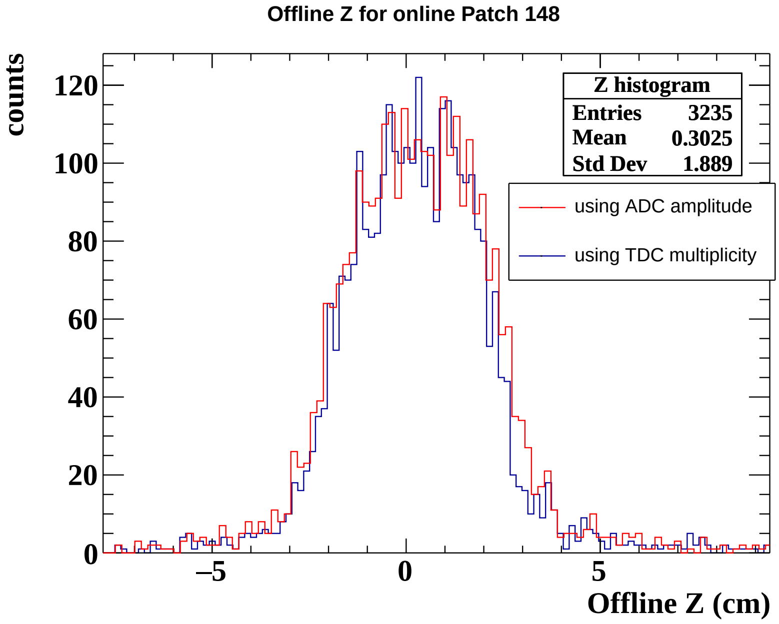

(b) Z

Figure 4.12: Comparison of positron conversion point estimations using TDC and ADC for a given Patch.

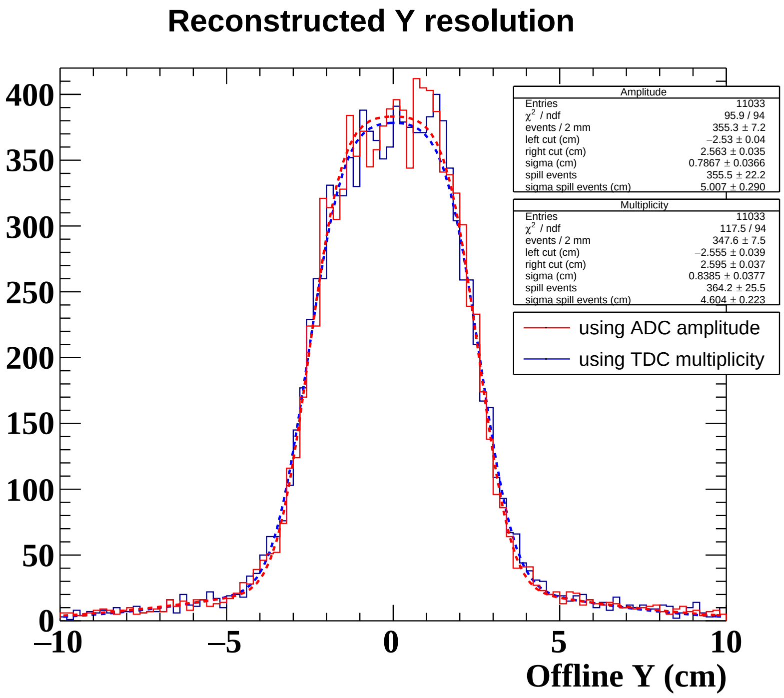

(a) Y

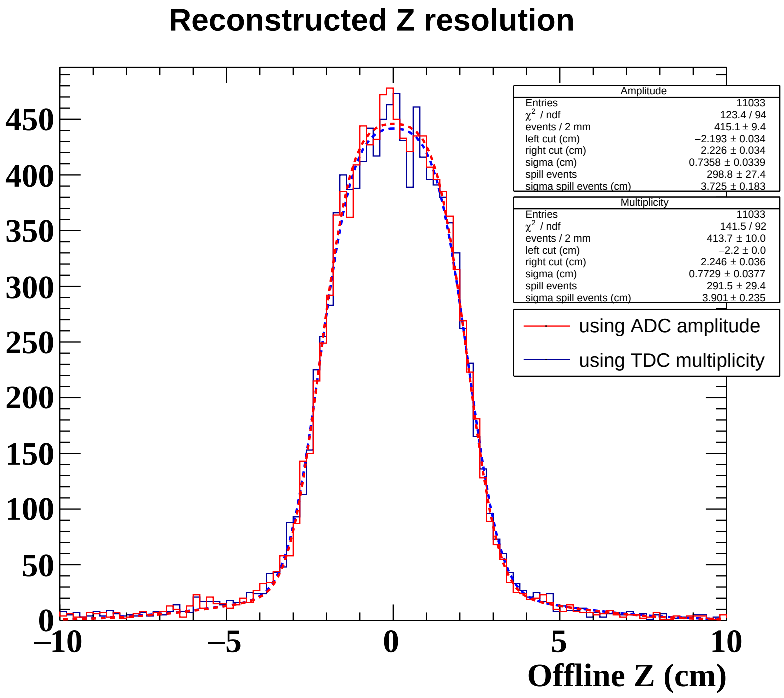

(b) Z

Figure 4.13: Online photon position resolution.

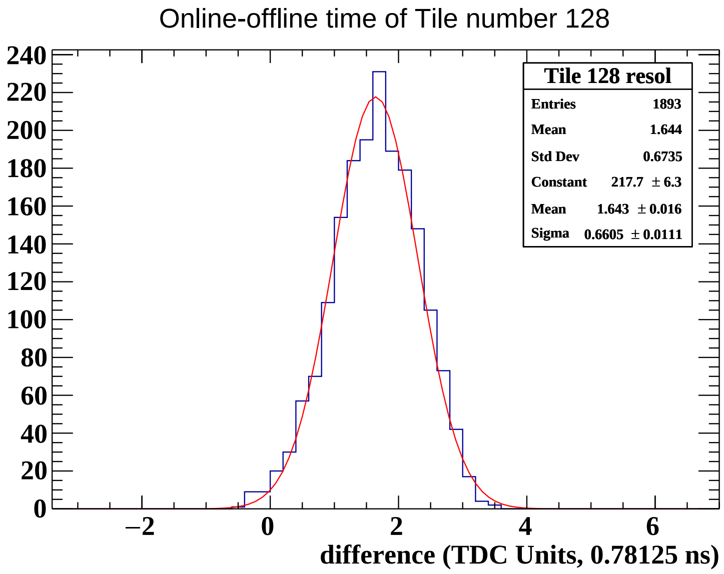

(a) single Tile

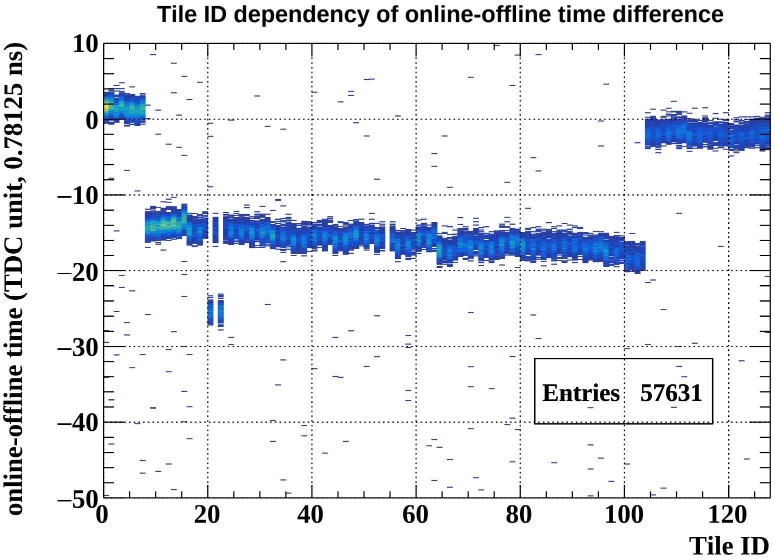

(b) dependency on Tile ID

Figure 4.14: Online positron time resolution.

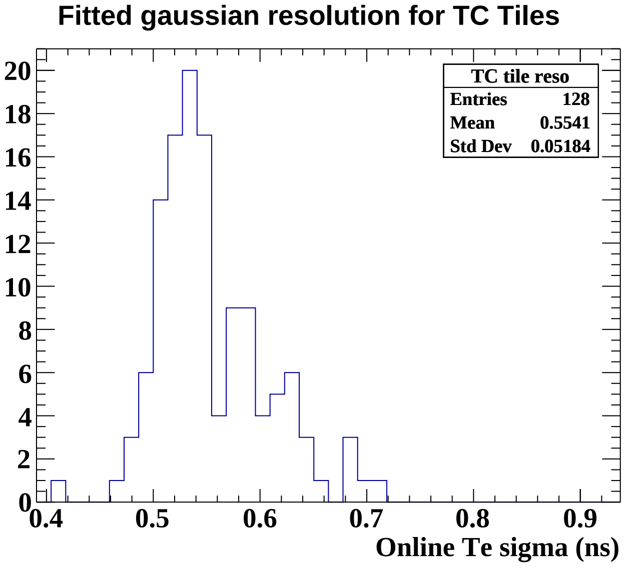

Figure 4.15: Single Tile online resolution obtained for each counter from the gaussian fit like figure 4.14a.

resolution is more than adequate for the MEG II necessities.

4.6 Online pulse shape discrimination

As previously described in section 2.4.2.1, a lattice of Americium α sources is submerged in Liquid Xenon calorimeter. To properly select their signals on top of the gamma background a pulse shape trigger is used: the calibrated sum of all the 32 PMTs, being enabled in that particular run, is shown in figure 4.16 in case of an α-like and a γ-like event. Photomultipliers are actually more sensitive to signal shape variations because of their smaller capacitance when compared to MPPCs. This makes their response much faster and allow to clearly differentiate between α and γ. The small undershoot after the signal is an artefact of the online pedestal subtraction logic; a small part of the signal enters the pedestal calculation shift register and gets subtracted from the subsequent baseline after the shift register length.

A peak finder is implemented in the FPGA searching for 8 consecutive samples where the second one is the largest of all but also above a programmable threshold to cut pedestal fluctuations; in this way a “snapshot” of the waveform is taken with a peak in a very specific position. The peak value and the charge are calculated from the snapshot window and the charge to amplitude ratio is selected with a trick very similar to the one

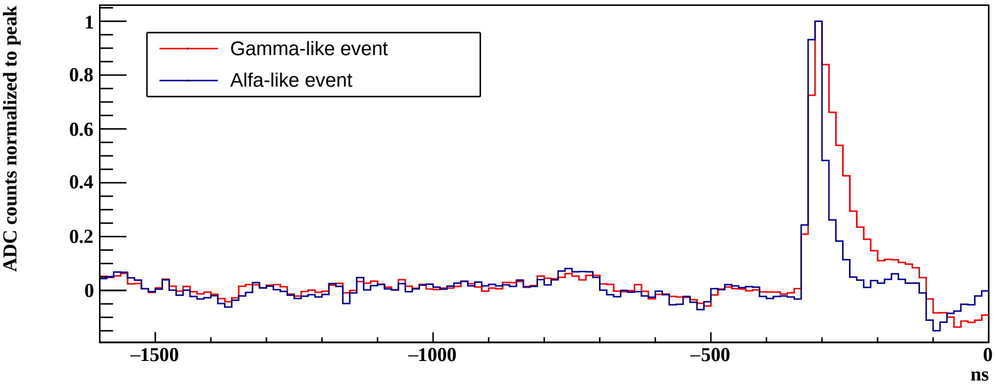

Figure 4.16: PMT sum online waveform for alfa-like and gamma-like events.

used for the $T_{e\gamma}$ trigger, to convert the ratio into a product:

$$\frac{\text{Charge}}{\text{Amplitude}} < \text{THR} \tag{4.1}$$

$$\text{Charge} - \text{THR} \cdot \text{Amplitude} < 0 \tag{4.2}$$

By checking the sign of the value Charge − THR · Amplitude an easier to implement condition is made to perform a cut on the charge to amplitude ratio.

Charge and amplitude are recorded into dedicated readout memories for offline checks of the algorithm: in figure 4.17a the offline computed ratio is shown. The two populations are on average 2 apart (in Q/A ratio units) with a combined sigma of the two gaussian being 0.54.

By producing the same plot without the DRS sampling the two populations, in figure 4.17b, are better separated keeping an average difference of 2 but with a combined sigma of 0.38, much smaller than the former one. The non gaussianity of the left peak in plot 4.17b is due to the source being an alfa but on a distant wire from the used PMTs, having a smaller absolute value and resulting in a broader distribution.

4.7 Trigger latency

The trigger latency can be measured easily by checking the signal position in the DRS window after a trigger is generated. The trigger time and the corresponding DRS stop time actually define the time zero of the event as visible in an example waveform from the Liquid Xenon in figure 4.18a. The observed latency accounts for ADC conversion, trigger information transmission, computation and trigger decision distribution to WaveDREAMs. In 2018 the longest trigger computation, also involving ADC waveform processing, was the Liquid Xenon energy trigger that also performed the mentioned Liquid Xenon position estimate. A way to monitor the overall Liquid Xenon latency and not just for a single channel may be performed by an alternative offline reconstruction

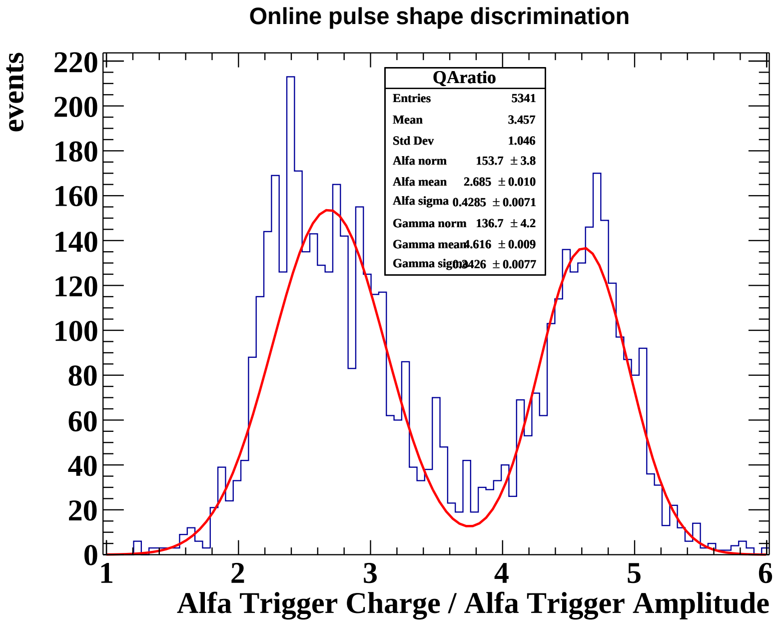

(a) DRS sampling

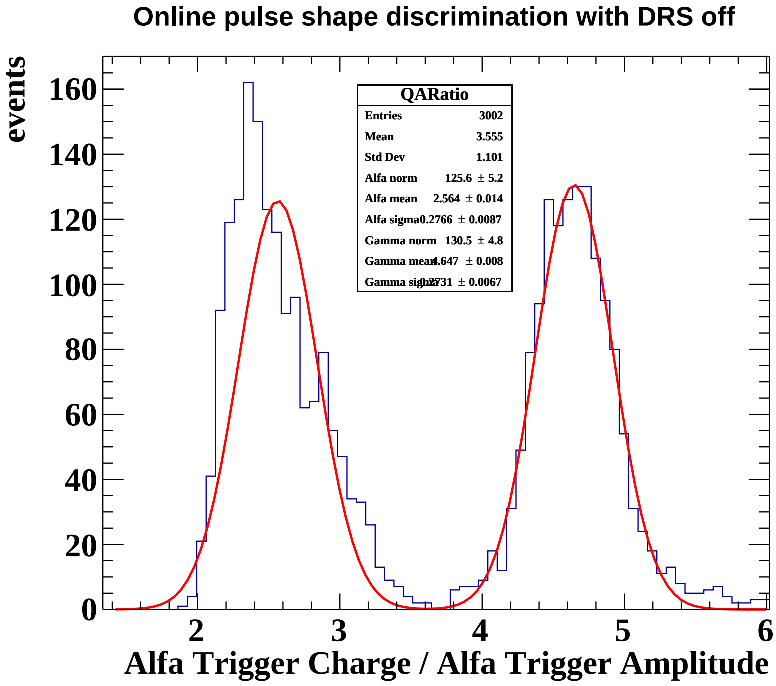

(b) No DRS sampling

Figure 4.17: Charge to amplitude trigger performances for alfa discrimination.

that calculates the sum of all DRS waveform in a fashion similar to that used in the trigger. The distribution of the peak positions of the summed waveform is shown in blue in figure 4.18b, giving an average latency value of 677 ns. For this reason the DRS sampling speed in Run 2018 was set to 1.2 GSPS, corresponding to a DRS digitisation window of 853 ns, also to provide enough samples in front of the signal for the offline analysis to extract and minimise the noise components evident in figure 4.18a.

The various components contributing to the latency are roughly: 200 ns for data transmission in the trigger tree, 250 ns for data processing both inside WaveDREAMs and TCBs and 200 ns needed by the current ADC converter to digitise the data (an Analog Device AD9637). The new WaveDREAM revision will instead use a different ADC (the Linear Technology LTM9009-14) with similar performances but a different internal design that shortens the conversion latency down to 60 ns. The effect is visible in the red plot of figure 4.18b, with the latency given by the same trigger but generated from pre-production WaveDREAM having the new ADCs. The very same algorithm, running on the new boards, results in a latency of 558 ns. Some tweaks in the algorithm is needed, and already identified, to gain few clock cycles in latency and operate the system at the design speed 1.8 GSPS, which corresponds to a DRS window of 569 ns.

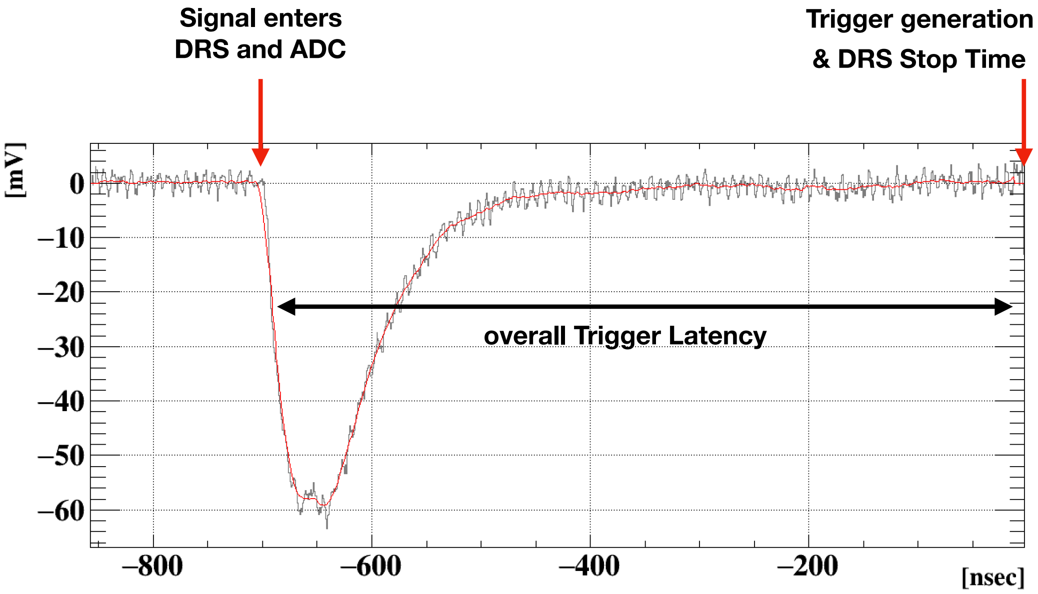

(a) Example of Liquid Xenon waveform

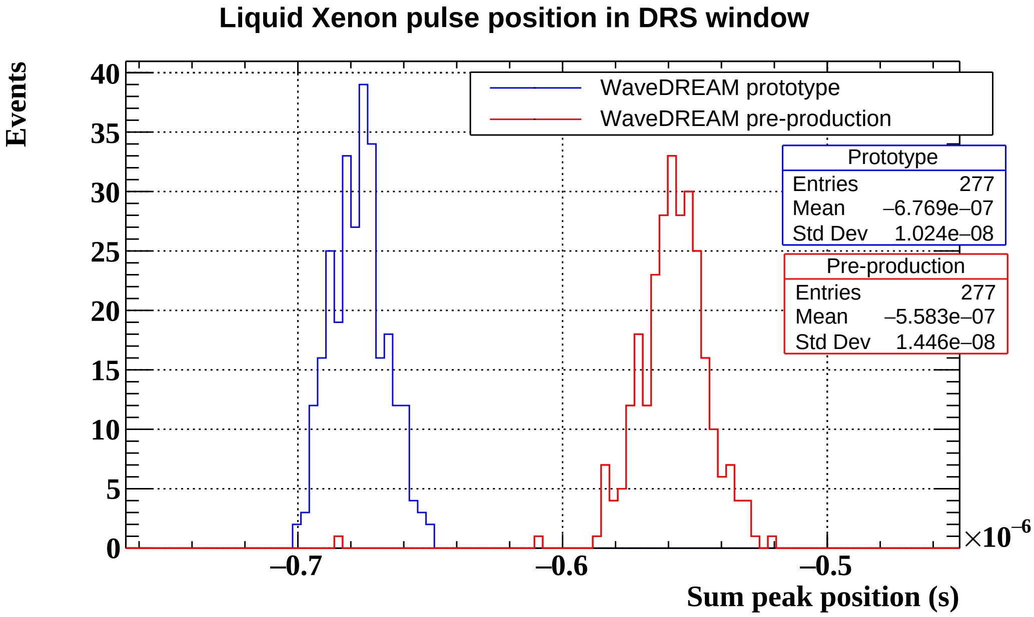

(b) Peak position in Liquid Xenon sum waveform

Figure 4.18: Extraction of peak position of DRS signals to monitor trigger latency.

5 Rate Estimations

Given the huge MEG II event size a careful estimation of online resolution of trigger observables and their effects on the trigger rate is mandatory. However our current simulation of the detector response, especially regarding the noise level and its coherency, is not at the required level to correctly reproduce the observed resolutions both in online and offline and this results in an overestimate of some performances.

5.1 $\mu \to \mathrm{e}\gamma$ DAQ Bandwidth requirements

A set of 10000 Montecarlo signal events were generated according to the standard MEG II design configuration to benchmark the online reconstruction possibilities and provide a baseline for further developments. Trigger ADC and TDC waveforms were generated in MEGBartender software with a respective sampling frequency of 80 MHz and 640 MHz starting from the simulated detector waveforms with a low pass filter at 50MHz to mimic the board ADC analog bandwidth. The 50 MHz value represent the design one used in the prototype WaveDREAM, the new hardware revision will include an additional filtering stage in between the DRS and the ADC with a cut frequency of 30MHz to mitigate the the DRS-induced noise observed in run 2018.

Starting from ADC and TDC waveform the full trigger simulation is carried out resulting in estimates of the various observables, in particular focusing on the E_γ and $T_{\mathrm{e}\gamma}$ extraction.

5.1.1 E_γ performance

In figure 5.1 the photon energy for signal events is plotted against the MPPC ID having the maximum amplitude. The collected light tends to increase when the conversion point moves from the central IDs toward the lateral IDs (the top and the bottom of the Liquid Xenon calorimeter). This effect was also observed with the standard DRS4 reconstruction on Montecarlo signal simulation and its likely due to the shower containment. A calibration factor is extracted by a parabolic fit to linearise the detector response.

The calibration of Liquid Xenon energy response as a function of the interaction point on the inner face is not implemented in the current trigger scheme however the final trigger board, where the Liquid Xenon sensor photodetectors is computed, also estimate the photon conversion point for the Direction Match trigger condition. A dedicated look-up table could provide the correction value for the detected Patch ID and to properly rescale the sum. Such process may add between two and three clock cycles to the trigger computation latency, so the full extent of this effect will be first studied with the complete detector readout and only later applied to the reconstruction.

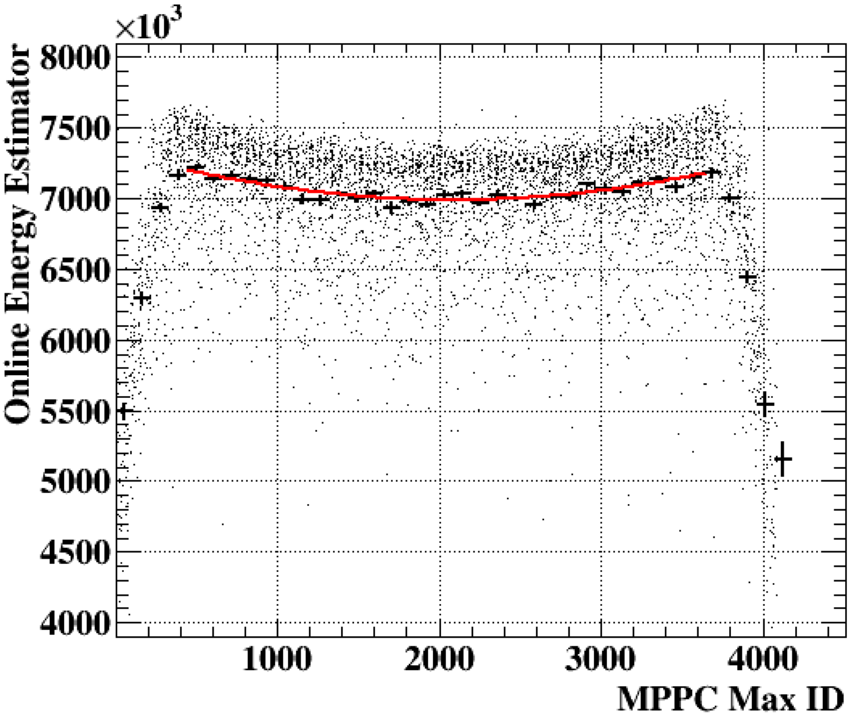

Figure 5.1: Online Montecarlo energy of signal events in ADC units versus maximum amplitude MPPC ID.

After applying the position dependent correction, the spectrum of photodetector sum amplitude, being the estimator of online photon energy, is fitted in figure 5.2a with a piecewise function:

$$\mathrm{p0}\cdot\begin{cases} e^{-\frac{1}{2}\cdot(\frac{x-\mathrm{p1}}{\mathrm{p2}})^2} & x>\mathrm{p1} \\ e^{-\frac{1}{2}\cdot(\frac{x-\mathrm{p1}}{\mathrm{p2}+\mathrm{p3}\cdot(x-\mathrm{p1})})^2} & x\leq \mathrm{p1}\end{cases} + \mathrm{p4}\cdot(1-\frac{1}{2}\cdot\mathrm{Erf}(\frac{\mathrm{x-p1}}{\mathrm{p2}})-0.5) \tag{5.1}$$

representing a gaussian distribution with an exponential left tail overimposed on a step function. The tail and the step function account for the intrinsic reconstruction inefficiency and losses due to the material between the target and the Liquid Xenon active volume.

The energy resolution at signal energy $\sigma(E_\gamma)/E_\gamma$ can then be extracted by the ratio of p2 (modelling the right gaussian tail) and the peak position p1.

An important parameter to be tuned in order to achieve the best possible online resolution is the relative calibration factor to properly scale the sum of PMT signals with respect to the MPPC one. The plot of figure 5.2b shows the obtained resolution as a function of this intercalibration factor. The best value is achieved for a value of 2.48 where a 1% energy resolution is obtained. This is obviously only a lower limit on the online reconstruction capabilities being extracted from a very optimistic Montecarlo simulation, however it definitely shine the capabilities of the MEG II TDAQ system.

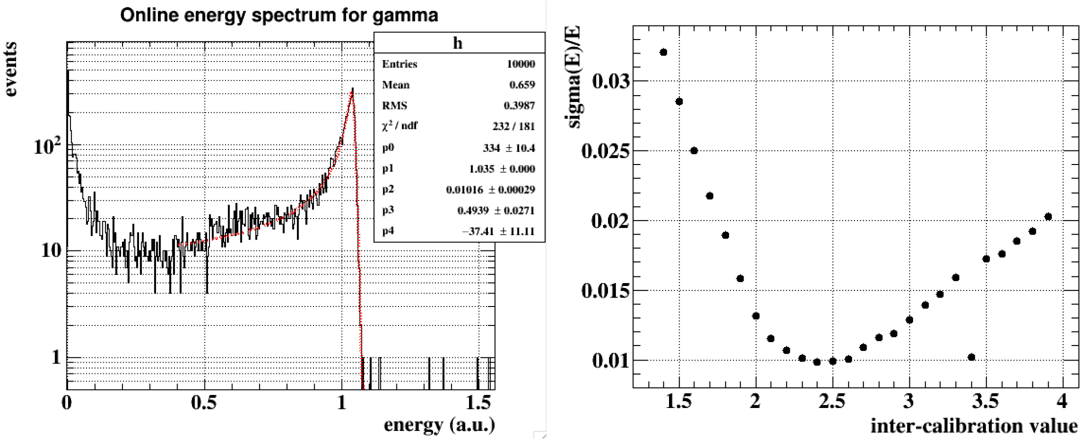

(a) Example of signal Montecarlo fit.

(b) Online energy resolution versus PMT-MPPC intercalibration .

Figure 5.2: Online energy resolution with different MPPC-PMT intercalibrations.

The intercalibration factor can be extracted experimentally both from the Cockroft Walton lithium line measurements or, better, from the calibration line given by the Charge Exchange reaction.

5.1.2 $T_{e\gamma}$ performance

During the Montecarlo waveform simulation by MEGBartender the signal event always is placed in at a fixed absolute time, resulting in pulses inside the digitised waveform with a timing depending only on the response of each detector. The window is then fixed according to the absolute time and, such feature, allows an independent study of T_γ and T_e.

With the standard MEGII Montecarlo configuration the electronic noise was small enough to allow a 5 mV threshold and, in that case, the Patch identifying the conversion point had all its 16 channels above threshold for a big fraction of the events (roughly 80%). The time is then averaged on all sixteen channels resulting, as shown later, in a very impressive time resolution estimate. After the 2018 run, we know that this condition might not be realistic, however in the following study we will focus only on events having all discriminators over threshold, to probe the best-possible online performances.

Even with such a small threshold, time-walk effect is present as shown in figure 5.3a. The time walk is a well known behaviour of leading edge discriminators like the analog comparator we have in the WaveDREAM boards, and it consists in a systematic underestimate of the peak time of the signal depending on its amplitude: higher signals will cross a fixed threshold earlier than smaller signals even thought their timing is the same.

The time-walk effect can be corrected by fitting and subtracting such dependency. In

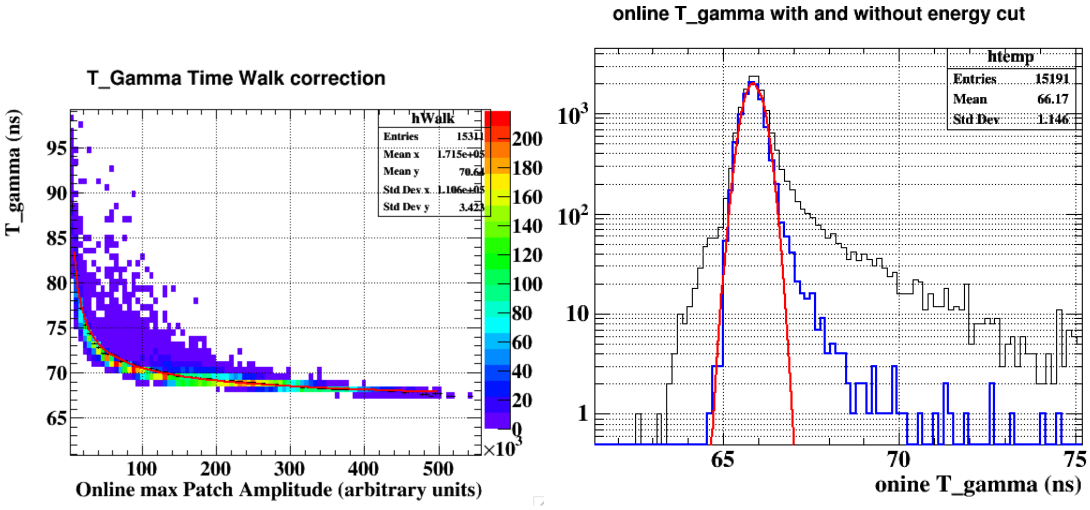

(a) Online reconstructed T_γ versus Maximum Patch energy.

(b) Time Walk correction effect on T_γ

Figure 5.3: Correction of online time walk in the Liquid Xenon detector.

this study an arbitrary phenomenological parametrisation was used:

$$T_\gamma^{\text{observed}} = T_\gamma^{\text{truth}} + \frac{k}{\sqrt{\text{Maximum Patch Amplitude}}} \tag{5.2}$$

The observed resolution on photon timing, assuming 16 channels over threshold, is reported in figure 5.3b before (black histogram) and after the time walk correction (blue histogram). Given the ideal noise assumptions made in the waveform simulation, the observed time resolution of 284 ps should be considered not as an expected value but rather as a lower limit. T_γ resolution in the range of 500-800 ps is a more reasonable estimate considering that in section 4.5 a resolution of 500 ps was obtained with the pixelated Timing Counter, where only two channels were used in the average[1]. Timing Counter signals are, indeed, higher amplitude and faster rising edge due to the SiPMs being used that have a smaller capacitance than Liquid Xenon MPPCs. For this reason, Tile signals do not require time walk correction that, in Liquid Xenon case, may not be possible in the hardware at the level of precision needed to completely compensate the effect, due to unavoidable quantisation errors.

For what concerns the T_e, its Montecarlo distribution receives a contribution from the time-of-flight distance from the target which depend on the impact position on the Timing Counter as shown in figure 5.4a (higher Tile ID implies larger $|z|$).

---

[1] The 800 ps is a conservative estimate which considers a single channel over threshold with the same electronic timing resolution observed in pTC signals (being 700 ps). An additional worsening of 100 ps has been added to cope with time-walk effect, because Liquid Xenon MPPC signals exhibit a slower rising edge and cannot be sharpened by the pole-zero circuit in the WaveDREAM boards, resulting a bit more noisy and less sharp than pTC signals.

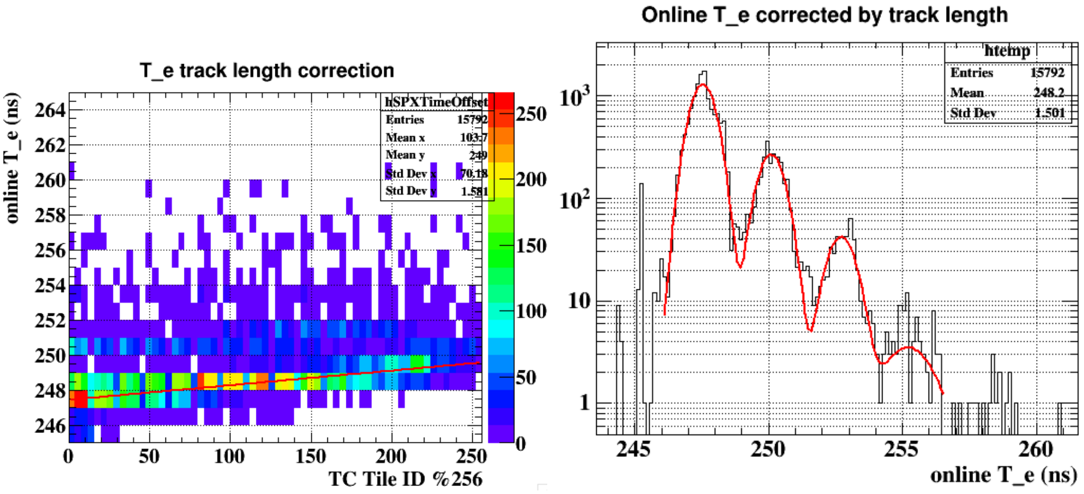

(a) Online reconstructed T_e versus TC hit position. (b) Time-of-flight corrected Montecarlo T_e distribution

Figure 5.4: Dependency of the T_e online time on impact position on TC and corrected T_e distribution.

After a linear correction based on the pixelated Timing Counter Tile ID the obtained T_e distribution is shown in figure 5.4b and presents multiple gaussian peaks equally spaced by ~ 2 ns due to different turns of the positron inside the CDCH volume before producing the first hit in the TC. Each gaussian have a sigma of 460 ns, being the single Tile hit resolution, not far from the 500 ps that were observed in section 4.5 during the 2018 run where the Tile ID.

The Tile ID dependent correction is not yet implemented in the hardware however they apply to both of the DRS and the online analysis and cancel out in the difference.

After the two time reconstructions have been individually tuned, the plot of $T_{e\gamma}$, shown in figure 5.5, is obtained, presenting the same three peaks with populations of 78%, 20% and 2% of the total generated sample and an overall resolution, of each gaussian between 560 and 600 ns.

5.1.3 Expected rate for MEG trigger

Combining the mentioned resolutions into an estimate of rate is difficult because of the uncertainties due to the final noise situation and the yet missing validation of such resolution with a fully instrumented equipment.

The trigger needs to be fully efficient in the region used for the offline analysis defined in the $T_{\mathrm{e}\gamma}$-E_γ plane, not to introduce bias in the events. The size of such a region is highly dependent on the achieved experimental resolution for each observable. Both offline and online distributions can be influenced by the noise situation, affecting differently the trigger requirements: in the first case a wider analysis region must be used and, in the

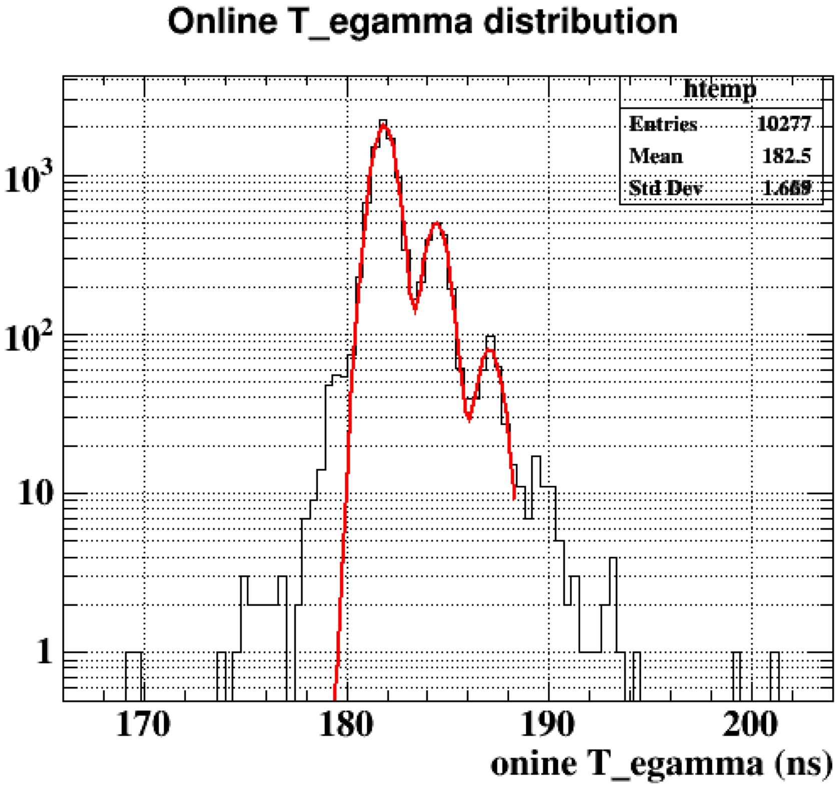

Figure 5.5: Online Montecarlo $T_{\mathrm{e}\gamma}$ distribution for signal events.

latter case, loosen cuts resulting in a reduced purity of the triggered sample to maintain the same level of efficiency.

For the MEG experiment the analysis region was defined in the $T_{\mathrm{e}\gamma}$ and E_γ plane, as shown by the recorded MEG I dataset in figure 5.6, with four different windows: the events with 48 MeV $< E_\gamma <$ 58 MeV and $|T_{\mathrm{e}\gamma}| < 1$ ns constituted the signal region and, because of this, were blinded. The energy region with 46 MeV $< E_\gamma <$ 48 MeV formed the Energy Side-Band where a peak of event around $T_{\mathrm{e}\gamma} = 0$ can be observed, being formed by RMDs, providing the relative normalization of this category of events with respect to the underlying flat Accidental background. A pair of Timing Side-Bands were defined with 1 ns $< |T_{\mathrm{e}\gamma}| <$ 3 ns providing a region where to study the accidental background under the reasonable assumption of being the same as inside the signal region.

The trigger was so required to be fully efficient in the area given by $|T_{\mathrm{e}\gamma}| < 3$ ns and $E_\gamma > 46$ MeV, which corresponds to $\pm 23\sigma_{T_{\mathrm{e}\gamma}^{\mathrm{MEGI}}}$ on the time axis and $\pm 7\sigma_{E_\gamma^{MEGI}}$.

By using the same resolution improvements applied in the sensitivity calculation (from table 2.3) and scaling the MEG I analysis region accordingly, a new region of $|T_{\mathrm{e}\gamma}| < 2$ ns and $E_\gamma > 49$ MeV can be defined as a reasonable estimate of what will be used in the MEG II offline analysis.

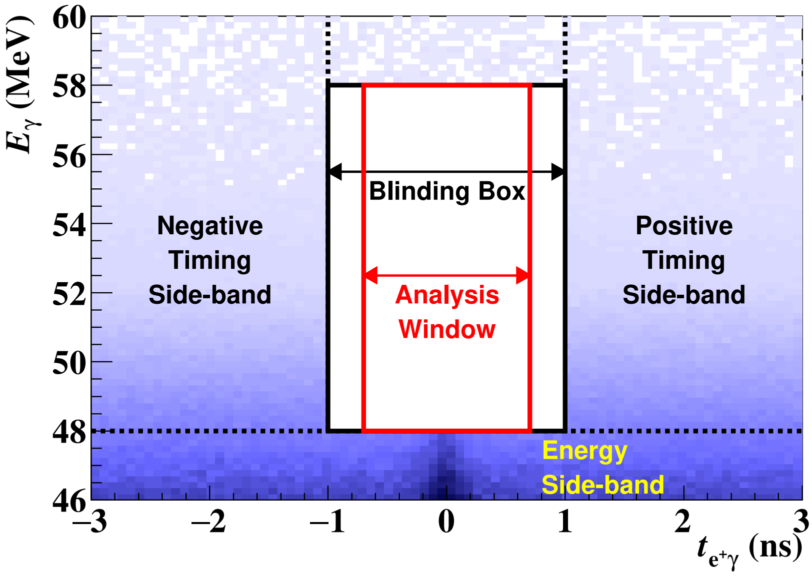

Figure 5.6: MEG I Analysis window with sidebands and blinding box definition.

To estimate the effect on the trigger rate, in figure 5.7, the nominal $T_{e\gamma}$ efficiency for the MEG I Experiment is compared with the MEG II one, under the two assumptions of keeping the MEG I analysis region or the more selective one. In both cases a conservative resolution of 800 ps on the online $T_{\mathrm{e}\gamma}$ is assumed with, in the MEG II case, a shift of the selection region (4 or 6 ns wide) to the right by 2.5 ns in order to account for the second peak in the reconstructed $T_{\mathrm{e}\gamma}$. The third smaller peak is dropped because of the difficulty of the offline analysis in disentangling such a high amount of hits in the high MEG II pileup environment. All the MEG II thresholds have been estimated to have a $> 99\%$ efficiency in the required analysis regions.

A similar plot can be made for E_γ and is reported in figure 5.8a, in that case assuming 1.5% online resolution for MEG II to account somehow for the noise situation. The efficiency curves of figure 5.7 and 5.8a are transformed into rate estimates under the following assumptions:

- The increased MEG II muon stop rate will cause an additional boost to the already dominating accidental background so, for the rate estimate, the RMD event component will be neglected.
- The $T_{\mathrm{e}\gamma}$ efficiency distribution, when applied on the flat accidental distribution, results in an event distribution identical to the efficiency curve. The trigger rate scales with the ratio of the integral of the efficiencies curves.
- The effect of E_γ efficiency instead needs to be weighted on the expected spectrum, that greatly decreases approaching the endpoint; in the current estimate only RMD

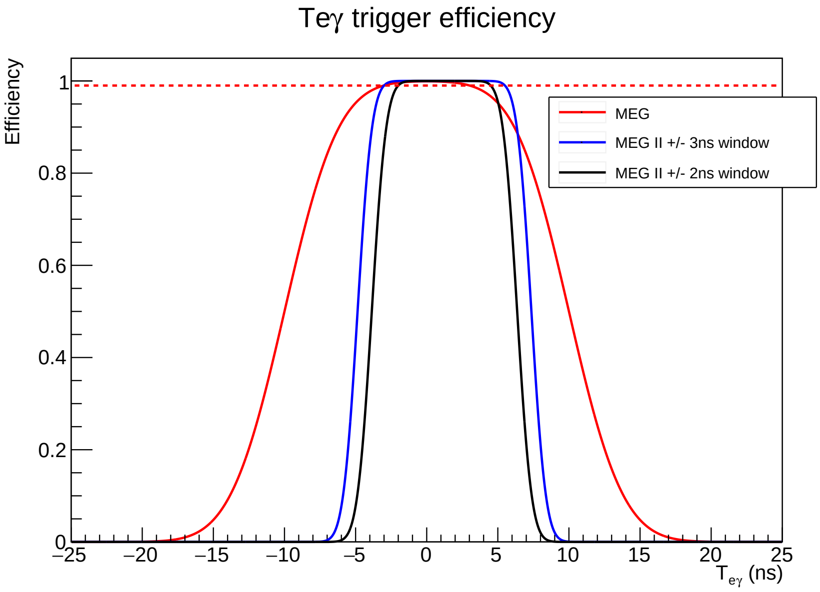

Figure 5.7: Achieved MEG I $T_{e\gamma}$ trigger efficiency and expectd MEG II ones.

photons will be considered because of their increased relative contribution due to the lighter spectrometer.

The expected energy distribution is shown in figure 5.8b.

Considering separately the effects of the two efficiencies and the increase of beam intensity from MEG I rate of $R_\mu^{\mathrm{MEG}} = 3 \cdot 10^7\ \mu^+/\mathrm{s}$ to $R_\mu^{\mathrm{MEGII}} = 7 \cdot 10^7\ \mu^+/\mathrm{s}$ as well as scaling of the accidental background with R_μ^2, the effect on the trigger rate is summarised in table 5.1 and can be expected to range between one half and twice the MEG I rate which was 12 Hz[86].

source	**effect respect MEG I rate** with same analysis region	**effect respect MEG I rate** with updated analysis region
beam	×5.44	×5.44
E_γ	×0.59	×0.23
$T_{e\gamma}$	×0.61	×0.51
Total	×1.95	×0.63

Table 5.1: Effect on trigger rate, relative to MEG I, of increased beam intensity and revised E_γ and $T_{e\gamma}$ triggers.

An additional reduction of the trigger rate comes from a more selective Direction Match table, in particular due to the higher segmentation of the new Timing Counter. That contribution is difficult to estimate due to the lack of a full online tracking of the positron. However this can be estimated by looking at the Direction Match table

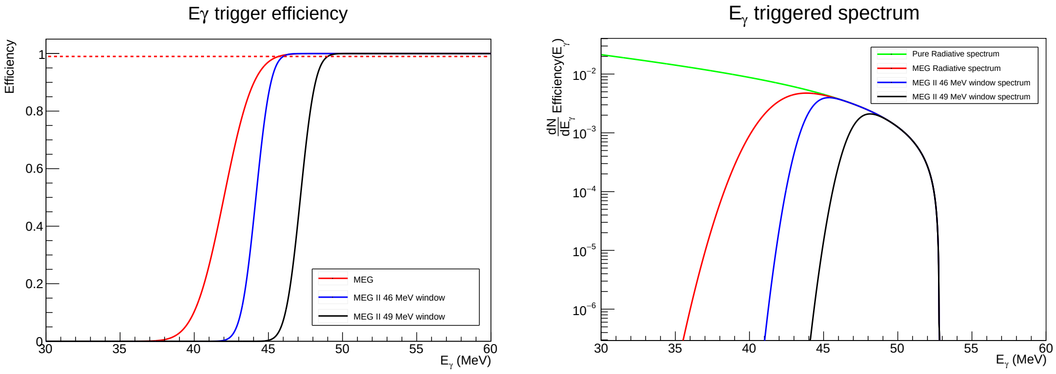

(a) Online MEG I E_γ efficiency and the expected MEG II ones.

(b) MEG I online selected RMD spectrum and MEG II ones

Figure 5.8: Effect of different E_γ trigger selections for MEG II compared with MEG I.

generated for MEG II, so that it covers 97% pTC-Liquid Xenon channel pairs. For a given Liquid Xenon Patch, in average, the 6% of the pTC tiles would be fire a coincidence if, accidentally, a positron impinges on them. By comparison, the MEG Direction Match table, also build to cover 97% of signal events, will have 19% fo the TC surface enabled for a given PMT on the inner face, as shown in figure 5.9. That is expected to result in a further factor 2 reduction of the accidental and RMD acceptance with the actual value needing to be extracted from data when the full Positron Spectrometer and the Liquid Xenon detector will be readout.

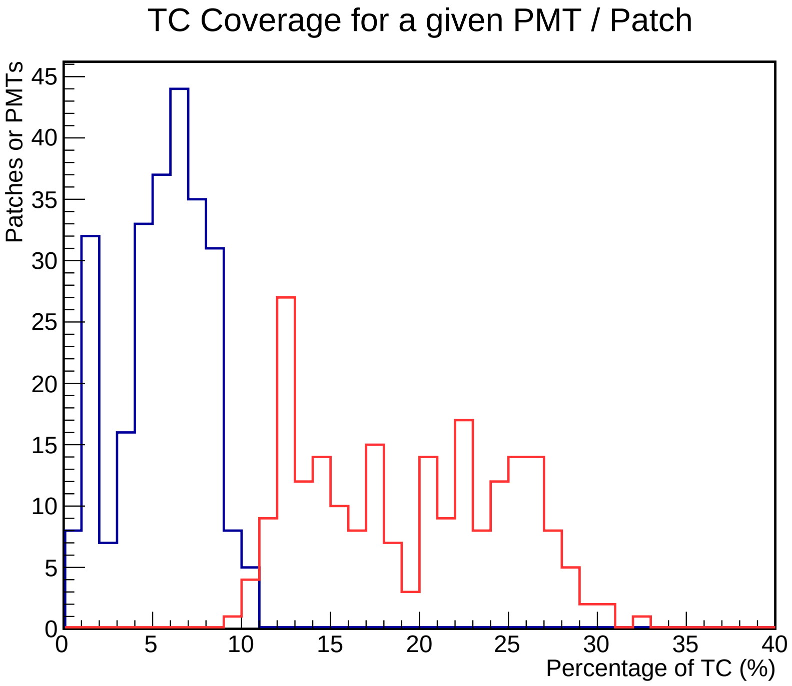

Figure 5.9: Fraction of Timing Counter enabled in Direction Match for each PMT (MEG I) in red or Patch (MEG II) in blue on the inner face.

6 MEG II CLFV search beyond $\mu \to e\gamma$: the $\mu \to eX$ channel

Even though the MEG and MEG II experiments are tailored to the needs of the $\mu \to e\gamma$ search, the excellent performances of the detectors are suitable to probe other CLFV channels. All lepton flavour violating processes have already very stringent bounds; any aim at improving them requires huge datasets to be recorded; possibly with a dedicated online event selection to reject uninteresting muon decays while preserving full efficiency on the channels of interest.

The trigger capabilities of the new MEG II TDAQ system are flexible enough to accommodate additional selections beyond the ones already planned: $\mu \to e\gamma$ event selection, including the ancillary trigger conditions for calibration and monitoring, only occupies $\sim 50\%$ of the available resources in the final trigger board and much less in the other TCBs. There is more than enough space to fit other event selections but such additional data taking should not have impact on the main search of the MEG II experiment.

On the data acquisition side, each event being triggered implies a fixed TDAQ deadtime of 400 μs, required to readout the DRS4 digitisers and record the event. Each WaveDREAM will keep its BUSY signal active during the readout deadtime period and the trigger boards will wait for it to be released before passing in the RUNNING state and starting the search for another event.

The fixed readout dead time sets an upper limit on the DAQ rate to 2.5 kHz that is reached when the trigger system generates a new trigger immediately after the dead time is elapsed. This TDAQ rate cannot be sustained while performing rare physics searches where the trigger has to analyse lots of event in order to select the few interesting for the offline processing; only the time spent in the RUNNING state is in fact employed in analysing physics data. For this reason the RUNNING time is a key parameter to be maximised and a useful observable to monitor its behaviour is the “live fraction”, defined as the ratio of RUNNING time with respect to total “wall clock” time.

The direct effect of a reduction of live fraction is an increased amount of time for the online event selection to process the same amount of physics events. In the end, a small live fraction leads to a longer duration of the experiment to record the same amount of data for the offline analysis.

The expected rate for the $\mu \to e\gamma$ event selection was described in chapter 5. In the worst case of 24 Hz trigger rate, it will produce a live fraction of 99.0% and, according to the sensitivity estimates of section 2.5, three years will be needed to integrate the required statistics. By reducing the live fraction to 90% the experiment will need to run three months more, although the trigger rate could be raised up to 240 Hz.

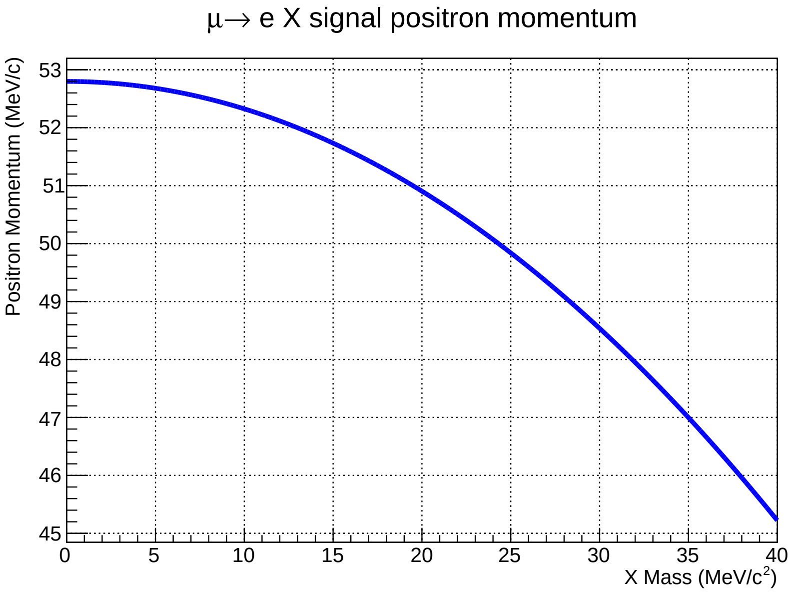

Figure 6.1: Positron momentum as a function of the X mass in the $\mu \to eX$ decay.

Part of the 216 Hz obtained on top of the $\mu \to e\gamma$ trigger will be used to record calibration and monitoring data while the rest, more interestingly, might be allocated to probe additional CLFV channels. For this reason an analysis of the channels within the reach of MEG II performances should be completed before the data-taking begins, so to properly configure the TDAQ system and to maximise the physics explorable with the MEG II experiment.

In this thesis the possibility of a $\mu \to eX$ search will be discussed with MEG II, similar feasibility studies are currently being carried out for the $\mu \to e\gamma\gamma$, $\mu \to eX\gamma$ and other rare/exotic decays. In this chapter a sensitivity study for the $\mu \to eX$ search in MEG I will be presented, with the goal to understand the challenges of this search; an extrapolation to the MEG II experiment will be subsequently given.

6.1 $\mu \to eX$ and the MEG experiment

The expected positron momentum in the $\mu \to eX$ process is obtained from the signal energy of equation 1.14 and is shown in figure 6.1 as a function of the X mass.

The MEG experiment detected positrons emitted almost perpendicularly to the polarisation of decaying muons within a small angular acceptance defined by the photon acceptance of the Liquid Xenon detector (which is opposite to the positron direction). The non uniform magnetic field and the radial extension of the tracking volume introduce a minimum momentum acceptance cut of $\sim$ 48 MeV/c.

From figure 6.1, the MEG I acceptance corresponds to exploring masses in the range $m_X \lesssim 30$ MeV/c^2.

Even if the positrons at kinematic endpoint ~ 52.8 MeV/c are within the acceptance, figure 6.1 shows that for small m_X values a small variation in the momentum corresponds to a large variation in X mass, therefore setting an upper limit for relatively small masses (< 13 MeV/c^2) requires for a very good knowledge of the positron momentum reconstruction.

Since most theoretically interesting models[48, 102] predict small m_X and $A = -1$ the remaining part of this study will focus on a Michel-like angular distribution and a peak search region near the kinematic endpoint. For X masses in the range 13 MeV/c^2 $< m_X <$ 30 MeV/c^2 the shape of Michel endpoint can be used to extract the momentum resolution (as it was done by TWIST) as well as the absolute momentum scale. The sensitivity would improve for high masses but, as reported later, a software cut was already applied on MEG I data used for this analysis and this limited the possibility to study in more detail the high mass region.

In the case of $A \neq -1$, the angular information can be combined to the momentum to increase the experiment sensitivity. However in MEG I the spectrometer is perpendicular to the muon polarisation with a small acceptance therefore limiting this possibility.

6.2 MEG I

The MEG I experiment was not specifically tailored toward a $\mu \to eX$ search, even though the event topology in the spectrometer is exactly the same of the nominal $\mu \to e\gamma$ search. The key difference is that $\mu \to eX$ is based on a single track topology and, because of this, it strongly depends on the correct characterisation of the detector response. $\mu \to e\gamma$ is instead searched as a coincidence of two particles; this coincidence, in addition to providing a further suppression for backgrounds, enabled to define time and energy Side-Bands as described in section 5.1.3. In the standard MEG analysis Side-Bands provide a pure background sample, which was used to measure the background probability density functions and a normalisation factor for the extrapolation of the background in the signal region. Side-Bands methods can not be applied in the $\mu \to eX$ search which relies on a single observable (positron momentum).

The complete MEG I dataset is composed of $7.5 \cdot 10^{14}$ muons that stopped at the target in a five year period. Only a limited subset of these muon decays fired the $\mu \to e\gamma$ trigger and, after pre-selection quality cuts, the remaining dataset consists in $6 \cdot 10^6$ events.

Pre-selection cuts were designed for the $\mu \to e\gamma$ analysis and included:

- reconstructed positron track quality;
- positron momentum 50 MeV/c $< P_e <$ 56 MeV/c;
- existence of a photon in the acceptance of the Liquid Xenon detector;
- photon energy 46 MeV $< E_\gamma <$ 60 MeV;
- relative timing $|\Delta T_{e\gamma}| < 7$ ns.

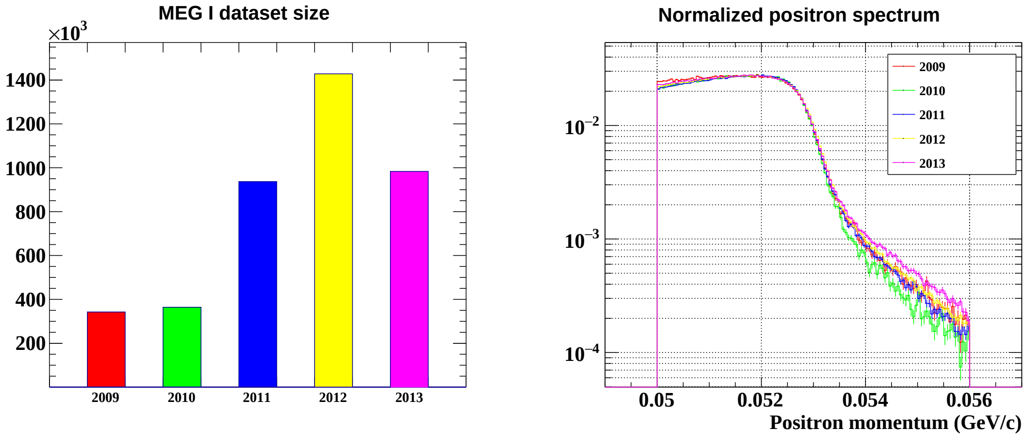

(a) Positron statistic by year

(b) Normalized positron spectrum

Figure 6.2: MEG I Michel statistics and spectra.

For the $\mu \to eX$ search only the first of these cuts is really needed while the others just reduce the statistics but do not alter the spectrum. As described in section 1.4.2.2, the main source of backgrounds for the $\mu \to e\gamma$ analysis is the accidental one: therefore any cut on the photon is not relevant for the uncorrelated positron. The pre-selection cut on positron momentum P_e restricts the explorable mass region to 0 MeV/c^2 $< m_X <$ 25 MeV/c^2.

As described in section 5.1.3, Radiative Michel decays can be identified on the basis of the $T_{e\gamma}$ distribution which, differently from the accidental events, is centered at zero and has a sigma of 130 ps. They can be easily removed by selecting off-time pairs of particles[1], with a time cut $|\Delta T_{e\gamma}| > 500$ ps. The final dataset after this selection contains $4 \cdot 10^6$ events.

6.2.1 Positron spectrum

MEG I physics data-taking went on for five years (2009-2013); during this period the spectrometer was not completely stable: there were frequent changes in the number of active drift chamber modules and, in the last two years of data-taking, wrinkles of up to 3 mm formed on the target surface.

In figure 6.2b the recorded positron spectra for the various years is shown after normalising in the range 51-53 MeV/c. It is not surprising, due to the previous mentioned effects, that the tail of events above the kinematic endpoint shows a variation with time.

The amount of collected data varies among the different years: in figure 6.2a the sizes of the different datasets are reported. The largest data sample is for 2012 and the

---

[1] In the preselected dataset the abundance of radiative decays was measured[22] to be 1 / 100. The core of the distribution is completely removed by the cut, however some leftovers from the tail component gives a spurious contamination of radiative dacays at a level of 3 events every 10^4 Michel events (less than 300 events in the final 2011 sample).

smallest ones for 2009 and 2010. We will limit the following analysis to 2011 data because on one side it was the most stable year from the point of view of the spectrometer and on the other we will show that this dataset is more than sufficient to saturate the MEG I systematics. This selection reduces the dataset by a factor 4 but simplifies the need to follow yearly variations of the experimental apparatus.

For the $\mu \to \mathrm{e}\gamma$ analysis the shape of Michel decays was assumed to have a general polynomial form and it was fitted by using the time side-bands. This procedure provided a data-driven estimate for the positron momentum distribution of the $\mu \to \mathrm{e}\gamma$ accidental background. However polynomial coefficients do not contain any physics information and this procedure requires a pure background dataset to perform the fit. This last requirement cannot hold for the $\mu \to \mathrm{e}X$ search, where there is no way to reject the signal. For this reason a physics-driven parametrisation of the positron spectrum must be employed.

The spectrum is modelled starting from the Michel spectrum $\mathcal{M}(P_e)$ including first order radiative corrections, which were already described in section 1.4.1.1. Radiative corrections are very important for this search since the endpoint shape is largely affected by them[2] as can be seen in figure 1.3.

The positron momentum distribution is integrated in the nominal spectrometer angular acceptance ($-0.35 < \cos(\theta_\mathrm{e}) < 0.35$ and $|\phi_\mathrm{e}| < \pi/3$)[3]. This distribution is multiplied by the acceptance function $A(P_e)$ modelling the spectrometer momentum acceptance. The acceptance is chosen to be an error function because it is sufficient to reproduce the observed data, as shown later:

$$A(P_\mathrm{e}) = \frac{1}{\sqrt{2\pi}A_\mathrm{slope}} \int_{-\infty}^{P_\mathrm{e}} e^{-\frac{(x-A_\mathrm{cut})^2}{2A_\mathrm{slope}^2}} \mathrm{d}x \tag{6.1}$$

with A_cut representing the central cut value and A_slope the steepness of the cut.[4] Finally the spectrum is convoluted with the spectrometer response function $\Sigma(P'_\mathrm{e} - P_\mathrm{e})$ to a monochromatic momentum P'_e. The current parametrisation of this function is composed of three gaussians, with the same central value. A shift in the central value can account for an offset in the reconstructed momentum scale which will prove to be important later on in this analysis.

An example of the theoretical Michel spectrum $\mathcal{M}(P_\mathrm{e})$, in the nominal angular acceptance, is given in figure 6.3a. Figure 6.3b shows the final momentum distribution $(\mathcal{M} \times A) \otimes \Sigma$ obtained by using the values of the parameters fitted to the data.

---

[2] Second order radiative corrections to the Michel decay are currently being computed. Their effect, which is expected to be at the 10^{-3} level, is further reduced by the non complete polarisation of the $\pi E5$ muon beam.

[3] The angular acceptance on reconstructed tracks is requested indirectly by the pre-selection cuts: the photon is requested in the nominal acceptance and the Direction Match trigger performs a subsequent selection of back-to-back tracks.

[4] Considering that angular variables present acceptance cuts, the acceptance should be implemented in the $P_\mathrm{e} - \theta_\mathrm{e}$ plane so to properly account for the momentum-angle correlation given by the Michel decay. As shown later, the momentum-only acceptance description however satisfactorily describes the positron spectra.

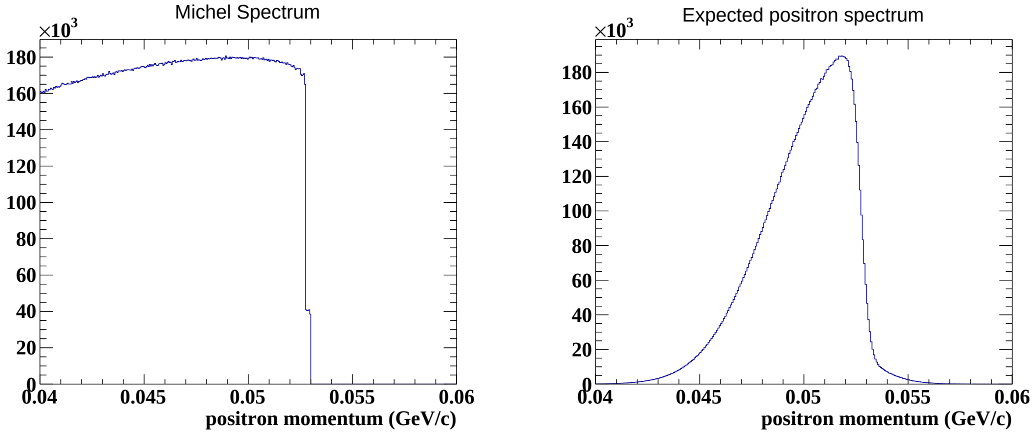

(a) $\mathcal{M}$

(b) $(\mathcal{M} \times A) \otimes \Sigma$

Figure 6.3: Expected positron spectrum (6.3b), compared with a pure Michel spectrum (6.3a).

To properly account for poorly reconstructed events, each positron track is sorted into High and Low Quality categories that will be analysed independently. The selection cuts for the High Quality category depend on year-by-year detector performances and were established for the $\mu \to e\gamma$ positron reconstruction:

- at least 9 hits;
- small fitted uncertainties on the observables, $U(\theta_e) < 0.7$ mrad, $U(\phi_e) < 1.95$ mrad, $U(P_e) < 75$ keV/c;
- good matching between drift chambers track and TC position;
- track $\chi^2 < 6$ computed Kalman filter track fitting[103].

The expected spectrum $(\mathcal{M} \times A) \otimes \Sigma$ can be separately fitted to both categories to extract the values of the spectrum parameters, as shown in figure 6.4.

The results of the fits are reported in table 6.1.[5] The resolutions are indeed better for High quality tracks and the momentum scale changes of 13 ± 9 keV/c between the two categories (compatible with no offset).

Perhaps surprisingly, the acceptance of Low Quality tracks, given by the value A_{cut}, extends to smaller energies than for High Quality tracks and the transition is sharper (A_{slope} is smaller for Low Quality). This behaviour can be understood considering that quality selection was tuned for 52.8 MeV/c positron: a lower momentum track has a smaller bending radius and, in the end, it is more likely to miss few drift chamber modules resulting in a lower quality reconstruction.

---

[5] The momentum acceptance cut for the Low Quality tracks is fixed at the value of a previous estimate given by the collaboration because the data cut at 50 MeV/c prevents a re-fit of this parameter.

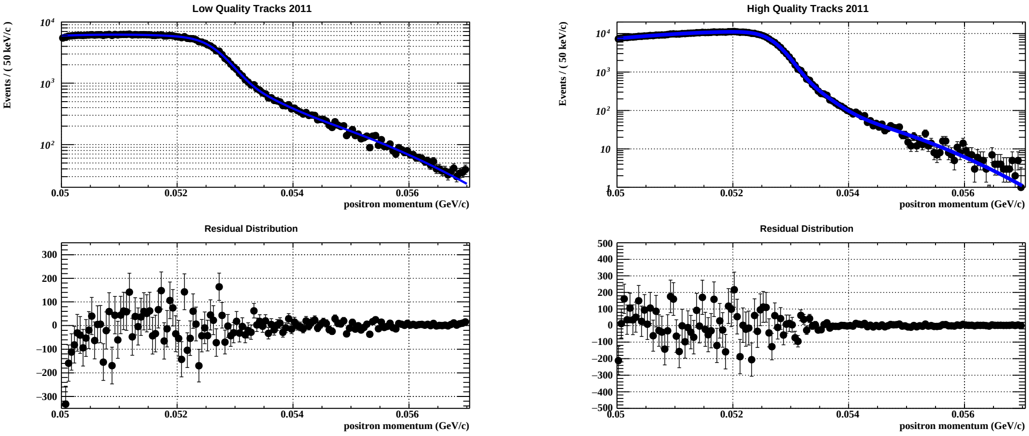

(a) Low Quality

(b) High Quality

Figure 6.4: Fit of positron acceptance and resolutions to MEG I data for Low and High Quality positrons separately.

Parameter	**High Quality**	**Low Quality**
Resolution Core σ_{core}	267 ± 6 keV/c	320 ± 20 keV/c
Resolution Mid σ_{mid}	580 ± 30 keV/c	800 ± 100 keV/c
Resolution Tail σ_{tail}	1.57 ± 0.08 MeV/c	2.1 ± 0.1 MeV/c
Core fraction f_{core}	$88 \pm 3\%$	$82 \pm 5\%$
Tail fraction f_{tail}	$11 \pm 3\%$	$14 \pm 5\%$
Momentum Scale P_{scale}	-103 ± 4 keV/c	-116 ± 5 keV/c
Acceptance cut A_{cut}	49.9 ± 0.1 MeV/c	47.6 MeV/c (fixed)
Acceptance slope A_{slope}	2.2 ± 0.2 MeV/c	1.8 ± 0.1 MeV/c

Table 6.1: Fitted MEG I spectrum parameters on 2011 Michel spectrum.

A different modelling of the Michel spectrum shape has a big impact on the $\mu \to eX$ measurement: in figure 6.5 the same data used for table 6.1 are fitted with a model having only two gaussians in the response function Σ. In the residuals below the plot an excess of few hundreds events at ~ 52.8 MeV/c (followed by a depletion of similar amount) is visible, due to the non perfect modelling of the high energy tail. Such parametrisation is the one used to extract the positron signal probability density function for the $\mu \to e\gamma$ analysis [6].

It is important to emphasise, at this point, the importance of an absolute momentum scale measurement and, with minor relevance, the precise knowledge of the other spectrum parameters: the plot of figure 6.6a shows the effect of an underestimate (of

---

[6] For $\mu \to e\gamma$ the requirements on the positron modelling are much more relaxed with respect to the $\mu \to eX$ analysis, as discussed previously. The level of precision of a simplified modelling is more than adequate for the purposes of the $\mu \to e\gamma$ analysis.

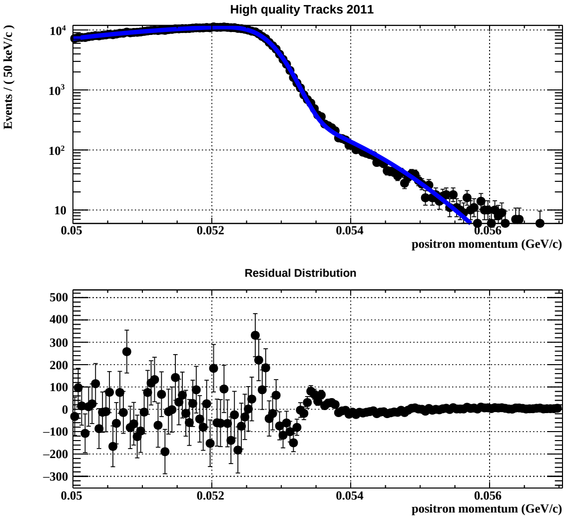

Figure 6.5: Fit of 2011 High Quality tracks with a response function composed of only two gaussians.

50 keV/c) of the momentum scale P_{scale}; the residuals, shown below the spectrum, distribute in a gaussian-like fashion that easily mimics the signal. A smaller effect can also be introduced by other parameters, for example, in the core resolution σ_{core}; a 50 keV/c overestimate of the resolution results in a bipolar-like feature on the residuals, as visible in figure 6.6b (bottom). Such plots prove the very strong correlation existing between the parameters describing the spectrum and the yield of a possible signal at the endpoint. The procedure to fit the spectrum shape directly on data could hide a possible signal as well as produce features that resemble it. The values of table 6.1, for this reason, cannot therefore be used directly as a background estimate.

In the TWIST experiment the momentum bias and its uncertainty were estimated by compiling a complete list of all the possible sources of bias, including, for example, the propagated bias from the magnetic field absolute value uncertainty. The total uncertainty in the momentum bias was computed by means of a Montecarlo simulation which takes into account the correlation among the different sources of bias. The TWIST momentum bias quoted is -4.7 ± 6.1 keV/c.

A different possibility for MEG would be Mott scattering described in section 2.4.5, which was designed to measure the spectrometer acceptance. Unfortunately this method does not have the required momentum precision for measuring the momentum bias, since it has an overall momentum spread of few 100 keV/c as discussed later[7].

Therefore also MEG I is forced to follow the same strategy of TWIST and provide an a priori estimate of the momentum scale uncertainty.

---

[7] Furthermore Mott scattering data were not taken frequently enough in MEG I to provide any useful estimation of the detector stability over time.

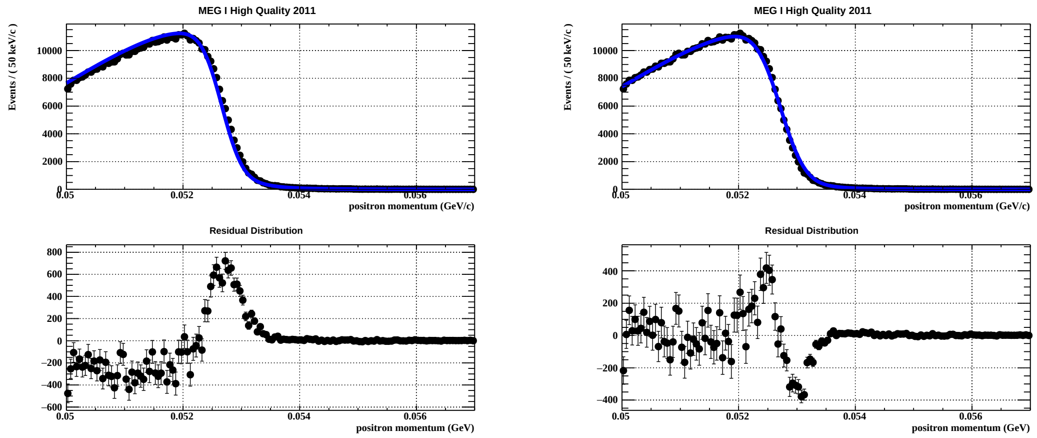

(a) 50 keV/c momentum scale underestimate

(b) 50 keV/c resolution overestimate

Figure 6.6: Effect of a parameter offsets on the reconstructed spectrum.

6.2.2 MEG I systematics

There are several effects which may result in biases of the spectrometer response function or in alterations of A and Σ distributions:

1. **DCH modules instability**: MEG I drift chamber modules experienced high voltage instabilities which caused modules to suddenly shut down altering the spectrometer acceptance and resolutions (because of the missing tracking layer). The effect of the different resolution can clearly be observed in the different high energy tails of the plots in figure 6.2b. The inefficiency following the high voltage instability of a module is however not expected to produce momentum scale biases because the tracking performances of the others modules were unaltered.

2. **DCH modules alignment**: the alignment of the sixteen DCH modules was measured by means of a laser tracker survey of corner cubes placed at the ends of each module, with precisions of the order of 200 μm. Two software alignments were performed, the first one employing positron tracks and the second with cosmic rays ones; the relative positioning uncertainties decreased to $\sim$ 50 μm. As long as there is no coherent shift in the positioning of all modules, the alignment uncertainty should only account for a worsening of the resolution and not a bias in the absolute value. The magnitude of a coherent shift of all modules, which can cause a bias in the momentum, is suppressed by the square root of the number of drift chamber modules that recorded hits (10 on average for the 2011 run), being $\delta R \simeq 50\ \mu\text{m}/\sqrt{10} = 16\ \mu\text{m}$. The uncertainty in the momentum scale bias is proportional to the relative error in the radius $\delta P_e/P_e = \delta R/R = 16\ \mu\text{m}/15\ \text{cm} = 1.1 \cdot 10^{-4}$ where 15 cm is the radius for 52.8 MeV/c positron. The estimated effect on the momentum bias uncertainty is therefore $\sim$ 5 keV/c.

3. **Muon beam position on target**: as already described in section 2.1.1, the beam positioning on target may vary for reasons outside the direct control of the experiment (such as an offset in the impact position of primary proton beam on target E) or for an undetected failure of a magnetic element of the muon beamline (which happened in 2012). Drift of few millimetres were observed in MEG I in the course of the data-taking. Such drifts can indeed affect acceptance because of the limited angular and radial acceptance of the MEG I spectrometer. It may also have a limited impact on the momentum scale because the target is slanted: an offset in the muon decay along Z can bias the tracks toward one side of the spectrometer, inducing upstream-downstream differences in the spectrometer response. Asymmetries of the spectrometer where however checked by means of the positron angular distribution[59] to be smaller than possible global uncertainties. We consider therefore this effect to be much smaller than the global ones (points 2, 5 and 6).

4. **Target position**: the target positioning was found to be unstable in MEG I, with offsets of several millimetres after each insertion cycle. This constituted a major problem for the extrapolation of the track direction to the target. Given the correlation of the track parameters in the fit, emission angle uncertainties can propagate to the momentum one. As already said, an offset of the decay point can result in a bias if upstream-downstream asymmetries in the spectrometer are present. As in the previous point however this effect can be considered to be small relative to the dominating ones (points 2, 5 and 6).

5. **Target deformation**: during the 5 years period of MEG I the target deformed by as much as 3 mm from the plane given by the support frame. In similarity with the two previous points this can alter the decay position of muons; this deformation introduces a systematic error also in the evaluated energy loss inside the target. The length of the positron track inside the 205 μm polyethylene-polyester target ranges from 218 μm, when the muon stop on the side looking at the Liquid Xenon detector, to zero, if it decays on the spectrometer side; the energy loss difference between the two cases is 61 keV. This value has to be averaged on the muon stopping depth distribution inside the target which drastically depends on the amount of material across the muon path and, in the end, to any target alteration. This effect can therefore be estimated to be $\sim$ 10 keV/c.

6. **CoBRa magnetic field knowledge**: CoBRa magnetic field was surveyed twice during the MEG I experiment when the detectors were not installed. The magnetic field interpolation has a precision of 10^{-3} and each point used was measured by Hall sensor probes calibrated with 10^{-4} precision in comparison to a reference NMR probe. A magnet field offset $\delta B/B \sim 3 \cdot 10^{-4}$ directly propagates to a momentum bias producing a momentum scale uncertainty of $\sim$ 15 keV/c for a 52.8 MeV/c positron.

Considering the various contributions the MEG I spectrometer momentum scale turns out to have an uncertainty of $\sim 20 - 30$ keV/c, depending on how the different contributions add (quadrature-linear sum)[8].

6.2.3 MEG I Sensitivity

Few key components are needed in order to proceed to a sensitivity estimate for the $\mu \to eX$ search with the MEG I experiment: the first one being the background probability density function, for which we will use the fitted values of table 6.1. That fit is the only possibility in MEG I to model the background but, as said earlier, using the same data for the limit extraction may artificially hide or enhance the signal. To address such loophole, in particular for what concerns the momentum scale, the uncertainty of 30 keV/c in the momentum bias will later be used.

The other key ingredient is the signal probability density function that can be also be extracted from the spectrum coefficients of table 6.1: the signal positron is produced at a fixed energy $P_e^{sig}(m_X)$ which depends on the X mass as described in the introduction of section 1.6.2. The signal probability density function is then given by $\Sigma(P_e - P_e^{sig}(m_X))$ that is the spectrometer response function Σ evaluated for the monochromatic signal positron. The yield of the signal can also be altered by the spectrometer momentum acceptance which is identical to the acceptance A used to describe the Michel spectrum.

The sensitivity can be evaluated through a cut-and-count analysis despite a likelihood fit would indeed be more sensitive to shape alterations. Developments are in progress to build a likelihood-based analysis, mainly having the advantage to avoid efficiency drop due to finite bin size which will be described later. However it was noticed that the systematics due to momentum scale uncertainty dominates the upper limit extraction just like in the cut-and-count case; the limits obtained by a preliminary version likelihood analysis was found in any case to be consistent with the ones by the cut-and-count calculation.

Cut-and-count analysis defines a signal bin $\pm 1.64\sigma_{\text{core}} \sim 500$ keV/c (see table 6.1) centred at the expected signal positron momentum. The bin therefore has a known efficiency (slightly smaller than 90% because of the long tails in the reconstruction) to collect a signal, the precise bin efficiency is calculated by the integral of the signal pdf in the bin range. The expected number of background events is given by the integral b of the background pdf within the bin multiplied by the amount of collected statistics N. Given that $N \sim 10^6$ in the 2011 run period and the bins are pretty large, the number of counts is big enough to motivate a gaussian approximation of the poissonian bin content. The upper limit on the signal (at the 90% confidence level) is thus given by $1.64 \cdot \sqrt{Nb}/\epsilon$ where ϵ is the signal efficiency, which has two separate components: the absolute signal acceptance A[9], due to the spectrometer, and the bin containment, which accounts for

---

[8] These estimates are an upper limits to the momentum scale uncertainty; a full set of Montecarlo simulations would be needed to evaluate the correlation of the different contributions as it was done by the TWIST collaboration.

[9] Even though the acceptance function A is the same for background and signal events, the effect is different and does not cancel in the ratio: the signal momentum span is rather limited (~ 1 MeV/c) and thus is sensitive to the local efficiency value; the background events instead distribute on all allowed momentum values and its number scales with the efficiency weighted on the spectrum distribution.

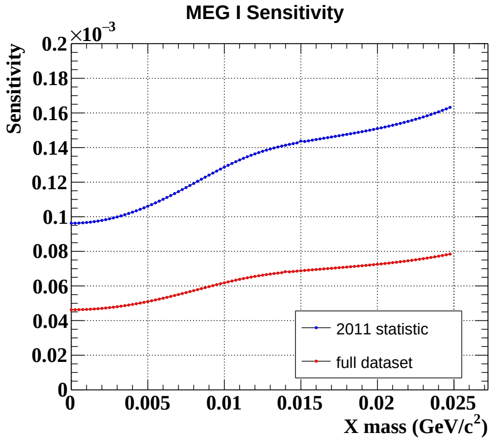

Figure 6.7: Expected $\mu \to eX$ sensitivities for the MEG I experiment.

the limited bin size.

The sensitivity calculation procedure can be repeated for various hypothesis of X mass, by shifting correspondingly the momentum of the signal positrons. No correction for the Look-Elsewhere effect is applied on the sensitivity curve because the mass hypothesis tested in the upper limit calculation is yet to be defined.

In figure 6.7 the sensitivity curves are shown as a function of the X mass in case of 10^6 events above 50 MeV/c (2011 run) or $4 \cdot 10^6$ (full MEG I dataset) [10]. The difference of the two configurations shows an increase in sensitivity of $\sim \sqrt{4}$ due to the reduction in background fluctuation with the increased statistics.

The plots of figure 6.7 account in fact only for statistic fluctuations and not for uncertainties due to systematics. Given that the momentum scale is known to be precise at $\sim$ 30 keV/c level, the high anti-correlation between this parameter and the signal upper limit for massless X can dramatically worsen MEG I sensitivity. The momentum scale can be parametrised as a gaussian with sigma $\sigma(P_{\text{scale}}) = 30$ keV/c; for the 90% C.L. sensitivity calculation, the momentum scale is allowed to vary in the 90% central interval of the gaussian (± 49 keV/c).

The sensitivity calculation described previously is changed as follows: background content in the signal bin is computed for several values of momentum scale within the allowed range. The maximal variation of bin content, with respect to the null offset hypothesis, is then added in quadrature to the contribution coming from the poissonian fluctuation in the calculation of the upper limit. The resulting sensitivity worsens particularly for small X masses, since in this case the signal bin is near the kinematic endpoint of the spectrum and, in that region, a small variation of the momentum scale causes a huge fluctuation of the bin content.

Using a cut-and-count analysis, which is based on a finite signal bin size, an additional worsening of the limit is produced by the actual signal pdf being also shifted by the

---

[10] The "full dataset" curves do not take into account year-by-year variation of the spectrum and, for this reason, should only be considered as best case scenario if a larger dataset is employed

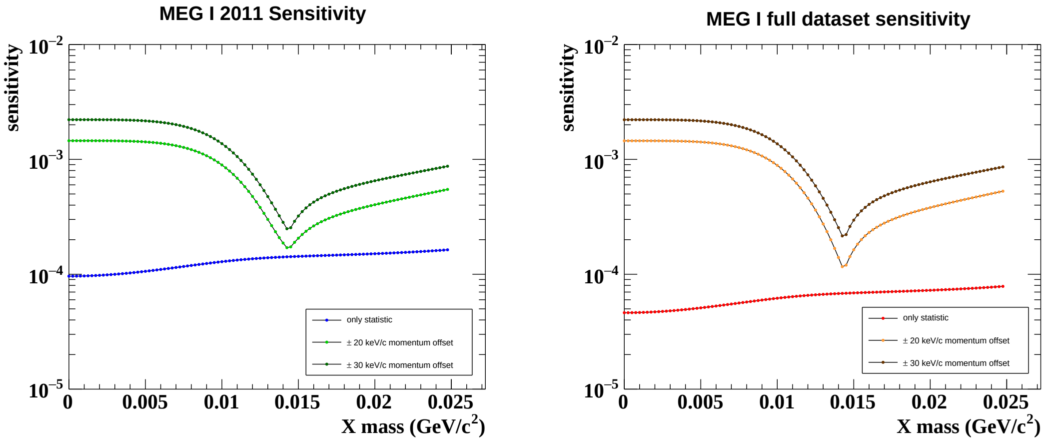

(a) 2011 statistic (10^6 events)

(b) full dataset statistic ($6 \cdot 10^6$ events)

Figure 6.8: Expected $\mu \to eX$ sensitivities for the MEG I experiment including a 20 keV/c or 30 keV/c uncertainty on the momentum scale.

momentum scale in the analysis bin. The bin efficiency drops by as much as 8% and needs to be recalculated for each momentum scale.

Figure 6.8a and 6.8b show the obtained sensitivities for the 2011 statistics and for the full dataset (assuming all data were calibrated as well as 2011 ones). The previous estimates, involving only the statistical term are shown for comparison.

The impact of the systematic has two different contributions, one dominating for small masses and the other for very high masses; the transition between the two effects is at $m_X \approx 15$ MeV/c^2 which corresponds to a signal momentum of $P_{\rm e}^{\rm sig} \approx 51.5$MeV/c. This can be understood in terms of the shape of the reconstructed Michel spectrum which peaks at exactly that energy because of the kinematic endpoint at higher momenta and the spectrometer acceptance at smaller ones. The slope of the Michel spectrum produces a fluctuation of expected background events in the signal region if a non null momentum shift is applied. This effect is more evident for the small m_X region, which has the highest slope due to the kinematic endpoint.

The plots of figure 6.8 show however that, including the momentum scale uncertainty, the 2011 statistics is more than enough to saturate MEG I sensitivity. The limits that can be extracted from MEG data are of the order of 10^{-3}, far worse than the existing limits of $\sim 10^{-5}$ by TWIST [56]. The difference is due to the mentioned systematic effects that were not appropriately kept in full control in MEG I.

For this reason we will not proceed to quote an upper limit on the MEG I data. Given the importance of the control of systematics in this measurement, assessing that the estimates of section 6.2.2 are precise requires a major work in modelling and validating them on a full Montecarlo production keeping also in account correlations among the sources of systematics. Such work is pointless given the the two orders of magnitude in sensitivities which separates MEG I sensitivity from the current TWIST upper limit.

6.3 MEG II

Even though MEG I sensitivity is not competitive with existing limits, the experience gained with MEG I, in particular concerning the importance of a precise knowledge of momentum scale, is really valuable to understand the possibilities of MEG II in the $\mu \to eX$ search.

6.3.1 Key improvements

MEG II will address all the systematics enumerated in the list of section 6.2.2 since they all play a role, even though with different relevance, also in the $\mu \to e\gamma$ analysis.

1. As already described, the upgraded spectrometer will use a single volume tracking drift chamber in place of the array of planar drift chambers used in MEG I and the completely new design will directly address the instabilities, which were observed in the former planar modules.

2. The monolithic nature of the new drift chamber will guarantee that the placement of the wires is known at level of $\sim$ 30 μm in the radial direction. Also the Z position will be better constrained: when the wire load is applied to the carbon fibre structure, a shrinking of $\sim$ 300 μm on a total length of 2 m was measured and can be compensated. The contribution to the momentum scale due to the wire placement is expected to be reduced below 10^{-4} level (a factor 2 improvement with respect to MEG I systematics).

3, 4. Beam and target shape will be monitored several times a day; this will help to track the history of the acceptance of the spectrometer and keep record of any possible upstream-downstream biases.

5. MEG II target is thinner than the former one (140 μm instead of $205 \mu m$) but slightly more dense; the combined effect being a maximum energy loss inside the target of 42 keV, 30% smaller than for MEG I.

6. The magnetic field knowledge is the most important item: the magnet is the same but an improvement of one order of magnitude in the magnetic field knowledge is required not to constitute the limiting factor. This issue can be addressed with a more precise calibration of the sensor (MEG II will use HallCube sensor from ETH/PSI[104], to provide a punctual 3D measurement of the magnetic field). A magnetic field precision down to the 10^{-4} level can be achieved, decreasing the momentum bias uncertainty at the endpoint down to 5 keV/c; at that level the magnet power supply stability starts becoming relevant: the coil currents stability was indeed measured to be better than $2 \cdot 10^{-4}$. A higher precision can be reached by a continuous monitoring of the B-field values so that the tracks can be corrected offline for the variation in magnetic field.

In MEG I, Mott calibration data were established starting from 2012 but data takings were only carried out yearly as a crosscheck of the spectrometer performances. Unfortunately it is not possible to obtain a monochromatic positron beam with an accuracy on the momentum of 5 keV/c due to the amount of material across the beam line (degrader, vacuum window and target); however Mott data can provide a fast crosscheck for long term variations as well as a precise measurement of the spectrometer acceptance and momentum resolution.

An additional calibration source to be explored is the Dalitz decay of neutral pions produced during charge exchange reactions $\pi^0 \to e^+e^-\gamma$, that happens with a branching ratio of 10^{-3}. However the interplay of spectrometer and calorimeter observables needed to reconstruct such a process may be difficult. The Double Dalitz decay $\pi^0 \to e^+e^-e^+e^-$ is the only neutral pion decay that only contains charged particles but its branching ratio is just 10^{-5} and it would be very difficult to collect a sizeable statistics in a short time.

Systematic	**MEG I**	**MEG II**
modules alignment	6 keV/c	negligible
target energy deposit correction	10 keV/c	5 keV/c
magnetic field	15 keV/c	5 keV/c
	20-30 keV/c	6-10 keV/c

Table 6.2: Momentum scale uncertainties in MEG I and MEG II. Only the most relevant points of section 6.2.2 are shown.

In conclusion the expected momentum scale uncertainty for MEG II will be between 5 keV/c and 10 keV/c according to the breakout of table 6.2, comparable to the 6.1 keV/c uncertainty of TWIST. The possibility to use Mott scattering to measure detector acceptance and resolution and pion-based reactions for the momentum scale could provide MEG II with a complete data-driven bounded momentum scale with respect to the a priori estimates needed for TWIST and MEG I.

6.3.2 MEG II Sensitivity

Since no data from MEG II detectors are yet available, Montecarlo simulation, also including waveform digitisation, was used to provide an estimate of the detector response.

In figure 6.9, Montecarlo events were fitted in the range of momentum above 45 MeV/c, as for the MEG I Michel fit of section 6.2.1, using in this case a single gaussian resolution for Σ. A momentum resolution of 120 keV/c is obtained and the spectrometer acceptance A extends to lower energies with the A_{cut} being as small as 46 MeV/c, because the new tracker extends to smaller distances from the target. According to the signal positron momentum for the $\mu \to eX$ process which was described in section 6.1, masses of X up to 40 MeV/c^2 are within the acceptance of MEG II.

Two sensitivity calculations, using the fitted resolution and acceptance, are shown in figure 6.10a and 6.10b respectively for a statistic of 10^7 and 10^8 events. The calculations were performed similarly to the ones developed for MEG I; three curves are shown:

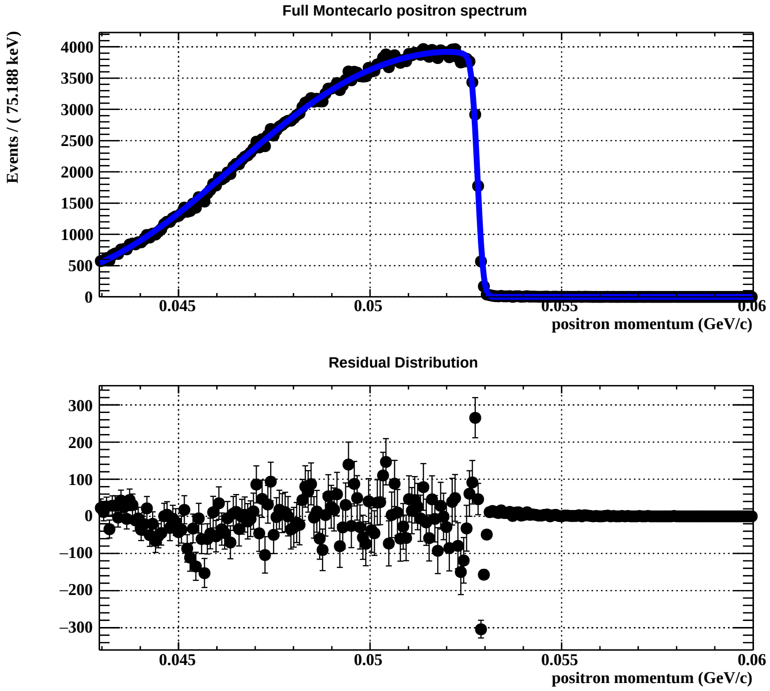

Figure 6.9: MEG II positron background model fitted on Full Montecarlo: resolution 120 ± 2 keV/c, acceptance cut 46.38 ± 0.03 MeV/c, acceptance sigma 3.15 ± 0.05 MeV/c.

the first one has no momentum scale shift (to show the bare statistical upper limit) the others are respectively for a 5 keV/c and 10 keV/c momentum scale uncertainty. While an accuracy of 5 keV/c is required in order to reach a sensitivity comparable with TWIST (shown in dashed red in the plots of figure 6.10), the statistics of 10^8 is more than sufficient, especially at low masses, since systematics dominate.

6.3.3 DAQ plan for a $\mu \to eX$ search in MEG II

Beyond the simple saturation of the systematics the amount of required events for $\mu \to eX$ needs to maintain a manageable size: if we assume a 569 ns DRS window (1.8 GSPS sampling frequency) and a rate of $7 \cdot 10^7$ muons per second, there will be 5-6 positron tracks in the CDCH and 3-4 hit clusters in the pTC for each event. Online zero suppression may discard all channels which do not record any hit, leaving 800 channels of CDCH (70 cells per track with a waveform from both ends) and 23 pTC channels, assuming an average of 9.2 tiles for each track which were observed in the pre-engineering runs [76].

This corresponds to a zero-suppressed event size estimation, for spectrometer-only waveforms, of 1.2 MB, in contrast with 6.7 MB with no suppression considered. Based on this assumption for the data size, table 6.3 shows the expected amount of storage

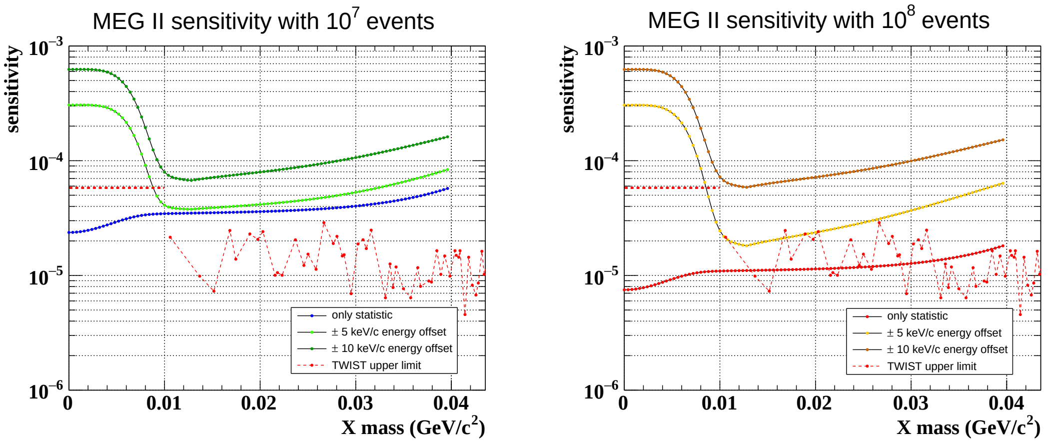

(a) 10^7 events

(b) 10^8 events

Figure 6.10: MEG II sensitivity to $\mu \to eX$ for two dataset sizes. Existing limits by the TWIST experiment are shown in dashed red.

Total number of events	**Dataset size**	**Note**
10^7	12 TB	≈ full MEG I dataset
10^8	120 TB	goal dataset
10^9	1.2 PB	$\mu \to e\gamma$ dataset size

Table 6.3: Dataset sizes for a MEG II $\mu \to eX$ search.

needed for various hypotheses of dataset size. The required 10^8 events would then occupy 120 TB of disk space, a factor 10 smaller than the $\mu \to e\gamma$ dataset.

This estimate does not take into account any event recovery from the $\mu \to e\gamma$ dataset; this is preferable because some alterations of the spectrum may also come from trigger biases: the positron angular distribution can be altered by the direction match trigger if some local inefficiencies (dead channels or reduced gain) on the calorimeter inner face sensors occur. The effect of trigger bias systematics can be removed by using a dedicated minimum bias trigger on the positron itself.

The required dataset size could be recorded in less than one week of full DAQ usage, assuming 200 Hz readout. Additional beam time might be needed to collect, before and/or after the data-taking, Mott data and other spectrometer calibrations.

In conclusion MEG II has the potential to probe the $\mu \to eX$ at a level of 10^{-5} with the same sensitivity of the best current upper limits. This requires an excellent understanding of the performances of the positron spectrometer, in particular track reconstruction inside the CDCH. The knowledge gained pursuing the $\mu \to eX$ search could boost other searches that rely on the spectrometer information.

Conclusions

A revived interest in CLFV searches resumed in the last decade, not only in the $\mu \to e\gamma$ channel but in the entire field of searches with muons. Such interest was pioneered by the same MEG experiment, which constitute the foundation of the much more sensitive MEG II setup. MEG I looked for signal of SuperSymmery theories before a direct searches at accelerators were possible and established current best limit on $\mu \to e\gamma$ branching ratio of $4.2 \cdot 10^{-13}$ (90% C.L.).

There is still consensus on the importance of lepton flavour violation searches, whose result shaped the Standard Model of particle physics being what we know today, and the MEG II experiment will be a leading actor on such stage.

A limit on the $\mu \to e\gamma$ branching ratio of $\mathrm{BR} = 6 \cdot 10^{-14}$ at 90% confidence level is expected to be established by MEG II in 3 years of data taking once all the detectors are assembled.

The performances of the experiment are assured by the bleeding edge technologies employed to achieve the goal resolutions:

- VUV-sensitive large area silicon photo detectors were deployed, for the first time, on a large scale, with the goal of exploiting the marvellous scintillation performances of liquid xenon. Liquid noble gas scintillators, and in particular xenon, have high light yield, typical of inorganic scintillators, paired with the fast time response, typical of organic ones.
- The monolithic drift chamber presents a wire density far beyond what was achieved in similar devices, exhibiting a small amount of material along the positron track.
- A 30 ps timing detector already showed a very nice stability over more than one month continuous data taking, thanks also to a dedicate laser-based calibration.
- A new array of monitoring devices ranges from beam and target positioning tools to the ones already extensively employed in MEG I, despite that, still essential and performant.
- The new integrated Trigger and Data Acquisition system houses, in the same space of the former experiment, 3.4 times the number of channels each with analog conditioning possibilities. It is capable of sustaining a huge data throughput, being limited in the end by the available disk space.

During my Ph.D. I took part in first line in the evolution of MEG II TDAQ from a prototype stage to an (almost) complete system capable of sustaining three engineering runs in which I also took part actively with more than 1024 channels in acquisition. All the trigger algorithms and offline processing code are finalised in their design and are waiting a sufficient number of readout channels to provide a reasonable acceptance for a final integration test. In particular the restyling of the readout code was a much needed improvement, that I developed from scratch, which can sustain line-speed data and will be tested on a more capable network configuration in the next year.

The new Trigger and Data Acquisition system is flexible and performant enough to allow other searches to be conceived in parallel with the experiment golden channel: the MEG II sensitivity on the $\mu^+ \to e^+X$ process was studied as a possible additional channel, requiring a huge control on experimental systematics to be competitive with other searches (to reach a BR sensitivity $\sim 10^{-6}$).

The Majoron/Familon channel is just an example of many possible channels within the reach of MEG II, not only including muon decays but which could also be generated by the Cockroft-Walton proton accelerator. A careful study of the possibilities will be performed before the physics data taking starts in 2021; the TDAQ will be tuned to fulfil such additional requests and exploit at maximum the extraordinary performances of the new experimental apparatus.

Bibliography

[1] A. Baldini, Y. Bao, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, M. Cascella, P. Cattaneo, G. Cavoto, *et al.*, “Search for the lepton flavour violating decay$\mu^+ \to \mathrm{e}^+\gamma$ with the full dataset of the MEG experiment,” *The European Physical Journal C*, vol. 76, no. 8, p. 434, 2016.

[2] A. Baldini, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, P. Cattaneo, G. Cavoto, F. Cei, M. Chiappini, *et al.*, “The design of the MEG II experiment,” *The European Physical Journal C*, vol. 78, no. 5, p. 380, 2018.

[3] L. Galli, A. Baldini, F. Cei, M. Chiappini, M. Francesconi, M. Grassi, U. Hartmann, M. Meucci, F. Morsani, D. Nicolò, *et al.*, “WaveDAQ: An highly integrated trigger and data acquisition system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 936, pp. 399–400, 2019.

[4] S. L. Glashow, “Partial-symmetries of weak interactions,” *Nuclear Physics*, vol. 22, no. 4, pp. 579–588, 1961.

[5] S. Weinberg, “A model of leptons,” *Physical review letters*, vol. 19, no. 21, p. 1264, 1967.

[6] Y. Kuno and Y. Okada, “Muon decay and physics beyond the standard model,” *Reviews of Modern Physics*, vol. 73, no. 1, p. 151, 2001.

[7] R. Barbieri, L. Hall, and A. Strumia, “Violations of lepton flavour and CP in supersymmetric unified theories,” *Nuclear Physics B*, vol. 445, no. 2-3, pp. 219–251, 1995.

[8] S. F. King, “Neutrino mass models,” *Reports on Progress in Physics*, vol. 67, no. 2, p. 107, 2003.

[9] J. N. Bahcall and R. Davis, “Solar neutrinos: a scientific puzzle,” *Science*, vol. 191, no. 4224, pp. 264–267, 1976.

[10] F. Cei and D. Nicolo, “Lepton flavour violation experiments,” *Advances in High Energy Physics*, vol. 2014, 2014.

[11] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Physical Review Letters*, vol. 81, no. 8, p. 1562, 1998.

[12] Q. R. e. a. Ahmad, “Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, Jun 2002.

[13] I. Esteban, M. Gonzalez-Garcia, A. Hernandez-Cabezudo, M. Maltoni, and T. Schwetz, “Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ 23, δ cp, and the mass ordering,” *Journal of High Energy Physics*, vol. 2019, no. 1, p. 106, 2019.

[14] P. Minkowski, “$\mu \to e\gamma$ at a rate of one out of 10^9 muon decays?,” *Physics Letters B*, vol. 67, no. 4, pp. 421–428, 1977.

[15] S. P. Martin, “A supersymmetry primer,” in *Perspectives on supersymmetry*, pp. 1–98, World Scientific, 1998.

[16] A. Canepa, “Searches for supersymmetry at the Large Hadron Collider,” *Reviews in Physics*, p. 100033, 2019.

[17] K. Hirata, T. Kajita, T. Kifune, K. Kihara, M. Nakahata, K. Nakamura, S. Ohara, Y. Oyama, N. Sato, M. Takita, Y. Totsuka, Y. Yaginuma, M. Mori, A. Suzuki, K. Takahashi, T. Tanimori, M. Yamada, M. Koshiba, T. Suda, K. Miyano, H. Miyata, H. Takei, K. Kaneyuki, Y. Nagashima, Y. Suzuki, E. Beier, L. Feldscher, E. Frank, W. Frati, S. Kim, A. Mann, F. Newcomer, R. V. Berg, and W. Zhang, “Experimental limits on nucleon lifetime for lepton+meson decay modes,” *Physics Letters B*, vol. 220, no. 1, pp. 308 – 316, 1989.

[18] E. Dudas, G. Von Gersdorff, J. Parmentier, and S. Pokorski, “Flavour in supersymmetry: horizontal symmetries or wave function renormalisation,” *Journal of High Energy Physics*, vol. 2010, no. 12, p. 15, 2010.

[19] W. Altmannshofer, R. Harnik, and J. Zupan, “Low energy probes of pev scale sfermions,” *Journal of High Energy Physics*, vol. 2013, no. 11, p. 202, 2013.

[20] M. Tanabashi, K. Hagiwara, K. Hikasa, K. Nakamura, Y. Sumino, F. Takahashi, J. Tanaka, K. Agashe, G. Aielli, C. Amsler, *et al.*, “Review of particle physics,” *Physical Review D*, vol. 98, no. 3, p. 030001, 2018.

[21] S. H. Neddermeyer and C. D. Anderson, “Cosmic-ray particles of intermediate mass,” *Phys. Rev.*, vol. 54, pp. 88–89, Jul 1938.

[22] A. Baldini, Y. Bao, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, P. Cattaneo, G. Cavoto, F. Cei, *et al.*, “Measurement of the radiative decay of polarized muons in the meg experiment,” *The European Physical Journal C*, vol. 76, no. 3, p. 108, 2016.

[23] R. R. Crittenden, W. Walker, and J. Ballam, “Radiative decay modes of the muon,” *Physical Review*, vol. 121, no. 6, p. 1823, 1961.

[24] W. Bertl, S. Egli, R. Eichler, R. Engfer, L. Felawka, C. Grab, E. Hermes, N. Kraus, N. Lordong, J. Martino, *et al.*, "Search for the decay $\mu^+ \rightarrow e^+e^+e^-$," *Nuclear Physics B*, vol. 260, no. 1, pp. 1–31, 1985.

[25] A. Crivellin, S. Najjari, and J. Rosiek, "Lepton flavor violation in the standard model with general dimension-six operators," *Journal of High Energy Physics*, vol. 2014, no. 4, p. 167, 2014.

[26] L. Calibbi and G. Signorelli, "Charged lepton flavour violation: an experimental and theoretical introduction," *arXiv preprint arXiv:1709.00294*, 2017.

[27] C. F. Powell, "Mesons," *Reports on Progress in Physics*, vol. 13, no. 1, p. 350, 1950.

[28] C. M. G. Lattes, H. Muirhead, G. P. Occhialini, and C. F. Powell, "Processes involving charged mesons," *Nature*, vol. 159, no. 4047, p. 694, 1947.

[29] E. Hincks and B. Pontecorvo, "On the absence of photons among the decay products of the 2.2 microsecond meson," *Canadian Journal of Research*, vol. 28, no. 1, pp. 29–43, 1950.

[30] S. Lokanathan and J. Steinberger, "Search for improbable meson decays," in *PHYSICAL REVIEW*, vol. 98-1, pp. 240–240, AMERICAN PHYSICAL SOC ONE PHYSICS ELLIPSE, COLLEGE PK, MD 20740-3844 USA, 1955.

[31] S. Frankel, J. Halpern, L. Holloway, W. Wales, M. Yearian, O. Chamberlain, A. Lemonick, and F. Pipkin, "New limit on the e+ γ decay mode of the muon," *Physical Review Letters*, vol. 8, no. 3, p. 123, 1962.

[32] B. Pontecorvo, "Electron and muon neutrinos," in *Neutrino physics*, Joint Institute for Nuclear Research (USSR), 1991.

[33] G. Danby, J. Gaillard, K. Goulianos, L. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos," *Physical Review Letters*, vol. 9, no. 1, p. 36, 1962.

[34] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak interactions with lepton-hadron symmetry," *Physical review D*, vol. 2, no. 7, p. 1285, 1970.

[35] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, "Evidence for the 2 π decay of the k_2^0 meson," *Physical Review Letters*, vol. 13, no. 4, p. 138, 1964.

[36] Y. Kuno and Y. Okada, "Proposed $\mu \rightarrow e\gamma$ search with polarized muons," *Physical review letters*, vol. 77, no. 3, p. 434, 1996.

[37] T. Abe, I. Adachi, K. Adamczyk, S. Ahn, H. Aihara, K. Akai, M. Aloi, L. Andricek, K. Aoki, Y. Arai, *et al.*, "Belle II technical design report," *arXiv preprint arXiv:1011.0352*, 2010.

[38] N. Berger, S. Dittmeier, L. Henkelmann, A. Herkert, F. M. Aeschbacher, Y. Ng, L. Noehte, A. Schöning, and D. Wiedner, "Ultra-low material pixel layers for the Mu3e experiment," *arXiv preprint arXiv:1610.02021*, 2016.

[39] U. Bellgardt, G. Otter, R. Eichler, L. Felawka, C. Niebuhr, H. Walter, W. Bertl, N. Lordong, J. Martino, S. Egli, *et al.*, "Search for the decay $\mu^+ \to e^+e^+e^-$," *Nuclear Physics B*, vol. 299, no. 1, pp. 1–6, 1988.

[40] L. Bartoszek, E. Barnes, J. Miller, J. Mott, A. Palladino, J. Quirk, B. Roberts, J. Crnkovic, V. Polychronakos, V. Tishchenko, *et al.*, "Mu2e technical design report," *arXiv preprint arXiv:1501.05241*, 2015.

[41] Y. Kuno, "A search for muon-to-electron conversion at j-parc: The comet experiment," *Progress of Theoretical and Experimental Physics*, vol. 2013, no. 2, p. 022C01, 2013.

[42] A. Grossheim, R. Bayes, W. Faszer, M. Fujiwara, D. Gill, P. Gumplinger, R. Henderson, A. Hillairet, J. Hu, G. Marshall, *et al.*, "Decay of negative muons bound in {sup 27} al," *Physical Review. D, Particles Fields*, vol. 80, no. 5, 2009.

[43] C. Dohmen, K.-D. Groth, B. Heer, W. Honecker, G. Otter, B. Steinrücken, P. Wintz, V. Djordjadze, J. Hofmann, T. Kozlowski, *et al.*, "Test of lepton-flavour conservation in $\mu \to$ e conversion on titanium," *Physics Letters B*, vol. 317, no. 4, pp. 631–636, 1993.

[44] W. Bertl, R. Engfer, E. Hermes, G. Kurz, T. Kozlowski, J. Kuth, G. Otter, F. Rosenbaum, N. Ryskulov, A. Van Der Schaaf, *et al.*, "A search for μ-e conversion in muonic gold," *The European Physical Journal C-Particles and Fields*, vol. 47, no. 2, pp. 337–346, 2006.

[45] F. Wilczek, "Axions and family symmetry breaking," *Physical Review Letters*, vol. 49, no. 21, p. 1549, 1982.

[46] M. Hirsch, A. Vicente, J. Meyer, and W. Porod, "Majoron emission in muon and tau decays revisited," *Physical Review D*, vol. 79, no. 5, p. 055023, 2009.

[47] Z. G. Berezhiani and A. Rossi, "Majoron decay in matter," *Physics Letters B*, vol. 336, no. 3-4, pp. 439–445, 1994.

[48] F. Björkeroth, E. J. Chun, and S. F. King, "Flavourful axion phenomenology," *Journal of High Energy Physics*, vol. 2018, no. 8, p. 117, 2018.

[49] R. D. Peccei and H. R. Quinn, "CP conservation in the presence of pseudoparticles," *Physical Review Letters*, vol. 38, no. 25, p. 1440, 1977.

[50] S. Adler *et al.*, "Measurement of the $K^+ \to \pi^+\nu\overline{\nu}$ branching ratio," *Phys. Rev. D*, vol. 77, p. 052003, Mar 2008.

[51] J. Collaboration, “A new search for the $K_L \to \pi^0 \nu \nu^-$ and $K_L \to \pi^0 X^0$ decays,” *Progress of Theoretical and Experimental Physics*, vol. 2017, 02 2017. 021C01.

[52] R. Ammar *et al.*, “Search for the Familon via $b^\pm \to \pi^\pm x^0$, $b^\pm \to k^\pm x^0$, and $b^0 \to k_s^0 x^0$ decays,” *Phys. Rev. Lett.*, vol. 87, p. 271801, Dec 2001.

[53] E. C. Gil, E. M. Albarrán, E. Minucci, G. Nüssle, S. Padolski, P. Petrov, N. Szilasi, B. Velghe, G. Georgiev, V. Kozhuharov, *et al.*, “The beam and detector of the NA62 experiment at CERN,” *Journal of Instrumentation*, vol. 12, no. 05, p. P05025, 2017.

[54] J. Heeck and W. Rodejohann, “Lepton flavor violation with displaced vertices,” *Physics Letters B*, vol. 776, pp. 385–390, 2018.

[55] R. Henderson, Y. I. Davydov, W. Faszer, D. Koetke, L. Miasoedov, R. Openshaw, M. Quraan, J. Schaapman, V. Selivanov, G. Sheffer, *et al.*, “Precision planar drift chambers and cradle for the TWIST muon decay spectrometer,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 548, no. 3, pp. 306–335, 2005.

[56] R. Bayes, J. Bueno, Y. I. Davydov, P. Depommier, W. Faszer, M. Fujiwara, C. Gagliardi, A. Gaponenko, D. Gill, A. Grossheim, *et al.*, “Search for two body muon decay signals,” *Physical Review D*, vol. 91, no. 5, p. 052020, 2015.

[57] A. Jodidio, B. Balke, J. Carr, G. Gidal, K. Shinsky, H. Steiner, D. Stoker, M. Strovink, R. Tripp, B. Gobbi, *et al.*, “Search for right-handed currents in muon decay,” *Physical Review D*, vol. 34, no. 7, p. 1967, 1986.

[58] R. Bolton, M. Cooper, J. Frank, A. Hallin, P. Heusi, C. Hoffman, G. Hogan, F. Mariam, H. Matis, R. Mischke, *et al.*, “Search for rare muon decays with the crystal box detector,” *Physical Review D*, vol. 38, no. 7, p. 2077, 1988.

[59] A. Baldini, Y. Bao, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, P. Cattaneo, G. Cavoto, F. Cei, *et al.*, “Muon polarization in the MEG experiment: predictions and measurements,” *The European Physical Journal C*, vol. 76, no. 4, p. 223, 2016.

[60] M. Berger, J. Coursey, M. Zucker, and J. Chang, “the nist estar database,” 2004.

[61] S.-G. Crystals, “Physical properties of common inorganic scintillators.”

[62] V. Chepel and H. Araújo, “Liquid noble gas detectors for low energy particle physics,” *Journal of Instrumentation*, vol. 8, no. 04, p. R04001, 2013.

[63] S. Ogawa and M. I. Collaboration, “Liquid xenon calorimeter for MEG II experiment with VUV-sensitive MPPCs,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 845, pp. 528–532, 2017.

[64] W. Ootani, K. Ieki, T. Iwamoto, D. Kaneko, T. Mori, S. Nakaura, M. Nishimura, S. Ogawa, R. Sawada, N. Shibata, *et al.*, “Development of deep-UV sensitive MPPC for liquid xenon scintillation detector,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 787, pp. 220–223, 2015.

[65] A. Ghassemi, K. Sato, and K. Kobayashi, “Mppc,” *Hamamatsu Photonics*, 2018.

[66] W. Ootani, W. Odashima, S. Kimura, T. Kobayashi, Y. Makida, T. Mitsuhashi, S. Mizumaki, R. Ruber, and A. Yamamoto, “Development of a thin-wall superconducting magnet for the positron spectrometer in the MEG experiment,” *IEEE Trans. Appl. Supercond.*, vol. 14, pp. 568–571, June 2004.

[67] G. Chiarello, “The full stereo drift chamber for the MEG II experiment,” *Journal of Instrumentation*, vol. 12, no. 03, p. C03062, 2017.

[68] G. Chiarello, A. Baldini, G. Cavoto, F. Cei, M. Chiappini, A. Corvaglia, M. Francesconi, L. Galli, F. Grancagnolo, M. Grassi, *et al.*, “The construction technique of the new MEG II tracker,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2018.

[69] G. Chiarello, C. Chiri, A. Corvaglia, F. Grancagnolo, M. Panareo, A. Pepino, C. Pinto, and G. Tassielli, “A high performance front end for MEG II tracker,” in *2015 6th International Workshop on Advances in Sensors and Interfaces (IWASI)*, pp. 55–57, IEEE, 2015.

[70] ISEG-HV, *ISEG EHS 8630p-305F: http://www.iseg-hv.com/.*

[71] M. Cascella, F. Grancagnolo, and G. Tassielli, “Cluster counting/timing techniques for drift chambers,” *Nuclear Physics B-Proceedings Supplements*, vol. 248, pp. 127–130, 2014.

[72] M. De Gerone, F. Gatti, W. Ootani, Y. Uchiyama, M. Nishimura, S. Shirabe, P. Cattaneo, and M. Rossella, “Design and test of an extremely high resolution Timing Counter for the MEG II experiment: preliminary results,” *Journal of Instrumentation*, vol. 9, no. 02, p. C02035, 2014.

[73] G. Boca, P. Cattaneo, M. De Gerone, M. Francesconi, L. Galli, F. Gatti, J. Koga, M. Nakao, M. Nishimura, W. Ootani, *et al.*, “The laser-based time calibration system for the MEG II pixelated Timing Counter,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 947, p. 162672, 2019.

[74] A. Baldini, E. Baracchini, G. Cavoto, M. Cascella, F. Cei, M. Chiappini, G. Chiarello, C. Chiri, S. Dussoni, L. Galli, *et al.*, “Single-hit resolution measurement with MEG II drift chamber prototypes,” *Journal of Instrumentation*, vol. 11, no. 07, p. P07011, 2016.

[75] M. Venturini, A. Baldini, E. Baracchini, F. Cei, S. Dussoni, L. Galli, M. Grassi, D. Nicolò, G. Signorelli, F. Tenchini, *et al.*, "Ageing tests for the MEG II drift chamber," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 824, pp. 592–594, 2016.

[76] M. De Gerone, A. Bevilacqua, M. Biasotti, G. Boca, P. Cattaneo, F. Gatti, M. Nishimura, W. Ootani, G. Pizzigoni, M. Rossella, *et al.*, "A high resolution Timing Counter for the MEG II experiment," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 824, pp. 92–95, 2016.

[77] R. Sawada, "Performance of liquid xenon gamma ray detector for MEG," *Nucl. Instrum. Methods A*, vol. 623, no. 1, pp. 258–260, 2010.

[78] A. Baldini, C. Bemporad, F. Cei, S. Dussoni, F. Gatti, M. Grassi, T. Haruyama, Y. Hisamatsu, T. Iwamoto, S. Mihara, *et al.*, "A radioactive point-source lattice for calibrating and monitoring the liquid xenon calorimeter of the MEG experiment," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 565, no. 2, pp. 589–598, 2006.

[79] J. Adam, X. Bai, A. Baldini, E. Baracchini, C. Bemporad, G. Boca, P. Cattaneo, G. Cavoto, F. Cei, C. Cerri, *et al.*, "Calibration and monitoring of the MEG experiment by a proton beam from a Cockcroft-Walton accelerator," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 641, no. 1, pp. 19–32, 2011.

[80] L. Hulthén, N. Mott, and H. Massey, "The theory of atomic collisions," 1949.

[81] G. Rutar, C. Bemporad, P. W. Cattaneo, F. Cei, L. Galli, P.-R. Kettle, and A. Papa, "A dedicated calibration tool for the MEG and MEG II positron spectrometer," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 824, pp. 575–577, 2016.

[82] J. Adam, X. Bai, A. Baldini, E. Baracchini, C. Bemporad, G. Boca, P. Cattaneo, G. Cavoto, F. Cei, C. Cerri, *et al.*, "The MEG detector for $\mu^+ \to e^+\gamma$ decay search," *The European Physical Journal C*, vol. 73, no. 4, p. 2365, 2013.

[83] S. Ritt, R. Dinapoli, and U. Hartmann, "Application of the DRS chip for fast waveform digitizing," *Nucl. Instrum. Methods A*, vol. 623, no. 1, pp. 486–488, 2010.

[84] L. Galli, F. Cei, S. Galeotti, C. Magazzù, F. Morsani, D. Nicolo, G. Signorelli, and M. Grassi, "An FPGA-based trigger system for the search of $\mu^+ \to e^+ + \gamma$ decay in the MEG experiment," *Journal of Instrumentation*, vol. 8, no. 01, p. P01008, 2013.

[85] Xilinx, *DS160 Spartan-6 Family Overview.*

[86] L. Galli, A. Baldini, P. Cattaneo, F. Cei, M. De Gerone, S. Dussoni, F. Gatti, M. Grassi, F. Morsani, D. Nicolo, *et al.*, "Operation and performance of the trigger system of the MEG experiment," *Journal of Instrumentation*, vol. 9, no. 04, p. P04022, 2014.

[87] P. Barry and P. Crowley, *Modern embedded computing: designing connected, pervasive, media-rich systems.* Elsevier, 2012.

[88] Xilinx, *DS180 7 Series FPGAs Data Sheet: Overview.*

[89] R. Giordano and A. Aloisio, "Fixed-latency, multi-gigabit serial links with Xilinx FPGAs," *IEEE Transactions on Nuclear Science*, vol. 58, no. 1, pp. 194–201, 2011.

[90] M. Francesconi, A. Baldini, F. Cei, M. Chiappini, L. Galli, M. Grassi, U. Hartmann, F. Morsani, D. Nicolò, A. Papa, *et al.*, "Low latency serial communication for MEG II trigger system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 936, pp. 331–332, 2019.

[91] Xilinx, *UG471 7 Series FPGAs SelectIO Resources User Guide.*

[92] Xilinx, *UG389 Spartan-6 FPGA DSP48A1 Slice User Guide.*

[93] Xilinx, *UG473 7 Series FPGAs Memory Resources.*

[94] Xilinx, *UG479 7 Series DSP48E1 Slice User Guide.*

[95] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, "Extensible markup language (xml) 1.0," 2000.

[96] C. S. Committee *et al.*, "iso/iec 14882: 2011, standard for programming language c++," tech. rep., Technical report, 2011. http://www. open-std. org/jtc1/sc22/wg21, 2011.

[97] S. Ritt, P. Amaudruz, and K. Olchanski, "The MIDAS data acquisition system," *Midas online:¡ https://midas. psi*, 1997.

[98] M. Schneebeli, R. Sawada, and S. Ritt, "ROME-a universally applicable analysis framework generator," in *Proceedings of the International Conference on Computing in High Energy and Nuclear Physics (CHEP06), Mumbai, India*, 2006.

[99] P. Cattaneo, F. Cei, R. Sawada, M. Schneebeli, and S. Yamada, "The architecture of MEG simulation and analysis software," *Eur. Phys. J. Plus*, vol. 126, no. 7, pp. 1–12, 2011.

[100] S. Agostinelli, J. Allison, K. a. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. . Barrand, *et al.*, “GEANT4-a simulation toolkit,” *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.

[101] R. Veenhof, “Garfield, a drift chamber simulation program,” in *Conf. Proc. C*, vol. 9306149, p. 66, World Scientific, 1993.

[102] L. Calibbi, F. Goertz, D. Redigolo, R. Ziegler, and J. Zupan, “Minimal axion model from flavor,” *Phys. Rev. D*, vol. 95, p. 095009, May 2017.

[103] J. Rauch and T. Schlüter, “Genfita generic track-fitting toolkit,” in *Journal of Physics: Conference Series*, vol. 608, p. 012042, IOP Publishing, 2015.

[104] C. Wouters, V. Vrankovi, C. Rssler, S. Sidorov, K. Ensslin, W. Wegscheider, and C. Hierold, “Design and fabrication of an innovative three-axis Hall sensor,” *Sensors and Actuators A: Physical*, vol. 237, pp. 62 – 71, 2016.