

The MEG detector for $\mu^+ \rightarrow e^+ \gamma$ decay search

J. Adam^{1,2}, X. Bai³, A.M. Baldini⁴, E. Baracchini^{3,15,6}, C. Bemporad^{4,5}, G. Boca^{7,8}, P.W. Cattaneo^{7,a}, G. Cavoto⁹, F. Cei^{4,5}, C. Cerri⁴, M. Corbo^{4,5}, N. Curalli^{4,5}, A. De Bari^{7,8}, M. De Gerone^{1,9,10,11,12}, L. Del Frate⁴, S. Doke¹³, S. Dussoni^{4,5,11,12}, J. Egger¹, K. Fratini^{11,12}, Y. Fujii³, L. Galli^{1,4,5}, S. Galeotti⁴, G. Gallucci^{4,5}, F. Gatti^{11,12}, B. Golden⁶, M. Grassi⁴, A. Graziosi⁹, D.N. Grigoriev^{14,19}, T. Haruyama¹⁵, M. Hildebrandt¹, Y. Hisamatsu³, F. Ignatov¹⁴, T. Iwamoto³, D. Kaneko³, K. Kasami¹⁵, P.-R. Kettle¹, B.I. Khazin¹⁴, O. Kiselev¹, A. Korenchenko¹⁶, N. Kravchuk¹⁶, G. Lim⁶, A. Maki¹⁵, S. Mihara¹⁵, W. Molzon⁶, T. Mori³, F. Morsani⁴, D. Mzavia^{16,†}, R. Nardò^{7,8}, H. Natori³, D. Nicolò^{4,5}, H. Nishiguchi¹⁵, Y. Nishimura³, W. Ootani³, K. Ozone³, M. Panareo^{17,18}, A. Papa^{1,4,5}, R. Pazzi^{4,5,†}, G. Piredda⁹, A. Popov¹⁴, F. Raffaelli⁴, F. Renga^{1,9,10}, E. Ripiccini^{9,10}, S. Ritt¹, M. Rossella⁷, R. Sawada³, M. Schneebeli^{1,2}, F. Sergiampietri⁴, G. Signorelli⁴, S. Suzuki¹³, F. Tenchini^{4,5}, C. Topchyan⁶, Y. Uchiyama^{1,3}, R. Valle^{11,12}, C. Voena⁹, F. Xiao⁶, S. Yamada¹⁵, S. Yamamoto¹⁵, S. Yamashita³, Yu.V. Yudin¹⁴, D. Zanello⁹

¹Paul Scherrer Institut PSI, 5232 Villigen, Switzerland
²Swiss Federal Institute of Technology ETH, 8093 Zurich, Switzerland
³ICEPP, University of Tokyo 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
⁴INFN Sezione di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
⁵Dipartimento di Fisica, Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
⁶University of California, Irvine, CA 92697, USA
⁷INFN Sezione di Pavia, Via Bassi 6, 27100 Pavia, Italy
⁸Dipartimento di Fisica, Università di Pavia, Via Bassi 6, 27100 Pavia, Italy
⁹INFN Sezione di Roma, Piazzale A. Moro, 00185 Rome, Italy
¹⁰Dipartimento di Fisica, Università di Roma “Sapienza”, Piazzale A. Moro, 00185 Rome, Italy
¹¹INFN Sezione di Genova, Via Dodecaneso 33, 16146 Genoa, Italy
¹²Dipartimento di Fisica, Università degli Studi di Genova, Via Dodecaneso 33, 16146 Genoa, Italy
¹³Research Institute for Science and Engineering, Waseda University, 3-4-1 Okubo, Shinjuku-ku, Tokyo 169-8555, Japan
¹⁴Budker Institute of Nuclear Physics of Siberian Branch of Russian Academy of Sciences, 630090 Novosibirsk, Russia
¹⁵KEK, High Energy Accelerator Research Organization, 1-1 Oho, Tsukuba, Ibaraki 305-0801, Japan
¹⁶Joint Institute for Nuclear Research, 141980 Dubna, Russia
¹⁷INFN Sezione di Lecce, Via per Arnesano, 73100 Lecce, Italy
¹⁸Dipartimento di Matematica e Fisica, Università del Salento, Via per Arnesano, 73100 Lecce, Italy
¹⁹Novosibirsk State Technical University, 630092 Novosibirsk, Russia

Received: 30 November 2012 / Revised: 8 February 2013 / Published online: 3 April 2013
© The Author(s) 2013. This article is published with open access at Springerlink.com

Abstract The MEG (Mu to Electron Gamma) experiment has been running at the Paul Scherrer Institut (PSI), Switzerland since 2008 to search for the decay $\mu^+ \rightarrow e^+ \gamma$ by using one of the most intense continuous μ^+ beams in the world. This paper presents the MEG components: the positron spectrometer, including a thin target, a superconducting magnet, a set of drift chambers for measuring the muon decay vertex and the positron momentum, a timing counter for measuring the positron time, and a liquid xenon detector for measuring the photon energy, position and time. The trigger system, the read-out electronics and the data acquisition system are also presented in detail. The paper is completed with a description of the equipment and techniques developed for the calibration in time and energy and the simulation of the whole apparatus.

Contents

1	Introduction	2
2	Beam line	3
2.1	Experimental requirements	3
2.2	π E5 beam line and MEG beam transport system	4
2.3	The Beam Transport Solenoid (BTS)	5
2.4	Spectrometer coupling system	5
2.5	Commissioning	6
2.6	Beam line performance	7
2.7	Target	7

^a e-mail: paolo.cattaneo@pv.infn.it

[†] Deceased.

3	COBRA magnet	10	11.4 Bartender	55
3.1	Concept	10	11.5 Database	56
3.2	Design	10	11.6 Comparison with data	56
3.3	Performance	12	12 Summary	56
4	Drift chamber system	12	Acknowledgements	58
4.1	Introduction	12	References	58
4.2	Design of DCH module	13		
4.3	Charge division and vernier pads	14		
4.4	Counting gas	14		
4.5	Pressure regulation system	14		
4.6	HV system	14		
4.7	Readout electronics	14		
4.8	Calibrations	16		
4.9	Tracking	17		
4.10	Performance	17		
5	Timing Counter	20		
5.1	Concept and design of the Timing Counter	20		
5.2	The longitudinal detector	21		
5.3	The transverse detector	22		
5.4	Commissioning and in-situ performance	23		
5.5	Trigger settings	24		
5.6	Timing resolution of TC	25		
5.7	Positron timing resolution	25		
5.8	Position resolution	25		
5.9	Efficiency	26		
6	Liquid Xenon detector	26		
6.1	Physical requirements vs. liquid xenon properties	26		
6.2	Detector design	27		
6.3	Reconstruction methods	29		
6.4	Performance	30		
7	Calibrations	34		
7.1	LXe detector calibration	34		
7.2	Monochromatic positron beam	37		
7.3	Intercalibration of detectors	38		
8	Trigger	41		
8.1	Requirements	41		
8.2	The trigger scheme	41		
8.3	Online algorithms	42		
8.4	Hardware implementation	42		
8.5	Operation and performance	43		
9	Front-end electronics and Data Acquisition system	45		
9.1	Requirements	45		
9.2	Front-end electronics	45		
9.3	DAQ system	47		
9.4	Performance	49		
10	Slow Control	50		
10.1	MSCB system	50		
10.2	DAQ integration	54		
11	Simulation	54		
11.1	The simulation architecture	55		
11.2	Physics event simulation	55		
11.3	Detector simulation	55		

1 Introduction

A search for the Charged Lepton Flavour Violating (CLFV) decay $\mu^+ \rightarrow e^+ \gamma$, the MEG experiment (see [1] and references therein for a detailed report of the experiment motivation, design criteria and goals) is in progress at the Paul Scherrer Institut (PSI) in Switzerland. Preliminary results have already been published [2, 3]. The goal is to push the sensitivity to this decay down to $\sim 5 \times 10^{-13}$ improving the previous limit set by the MEGA experiment, 1.2×10^{-11} [4], by a factor 20.

CLFV processes are practically forbidden in the Standard Model (SM), which, even in presence of neutrino masses and mixing, predicts tiny branching ratios ($BR \ll 10^{-50}$) for CLFV decays. Detecting such decays would be a clear indication of new physics beyond the SM, as predicted by many extensions such as supersymmetry [5]. Hence, CLFV searches with improved sensitivity either reveal new physics or constrain the allowed parameter space of SM extensions.

In MEG positive muons stop and decay in a thin target located at the centre of a magnetic spectrometer. The signal has the simple kinematics of a two-body decay from a particle at rest: one monochromatic positron and one monochromatic photon moving in opposite directions each with an energy of 52.83 MeV (half of the muon mass) and being coincident in time.

This signature needs to be extracted from a background induced by Michel ($\mu^+ \rightarrow e^+ \nu \bar{\nu}$) and radiative ($\mu^+ \rightarrow e^+ \gamma \nu \bar{\nu}$) muon decays. The background is dominated by accidental coincidence events where a positron and a photon from different muon decays with energies close to the kinematic limit overlap within the direction and time resolution of the detector. Because the rate of accidental coincidence events is proportional to the square of the μ^+ decay rate, while the rate of signal events is proportional only to the μ^+ decay rate, direct-current beams allow a better signal to background ratio to be achieved than for pulsed beams. Hence we use the PSI continuous surface muon beam with intensity $\sim 3 \times 10^7 \mu^+/s$ (see Sect. 2).

A schematic of the MEG apparatus is shown in Fig. 1. A magnet, COBRA (COnstant Bending RAdius), generates a gradient magnetic field, for the first time among particle physics experiments, with the field strength gradually decreasing at increasing distance along the magnet axis from the centre.

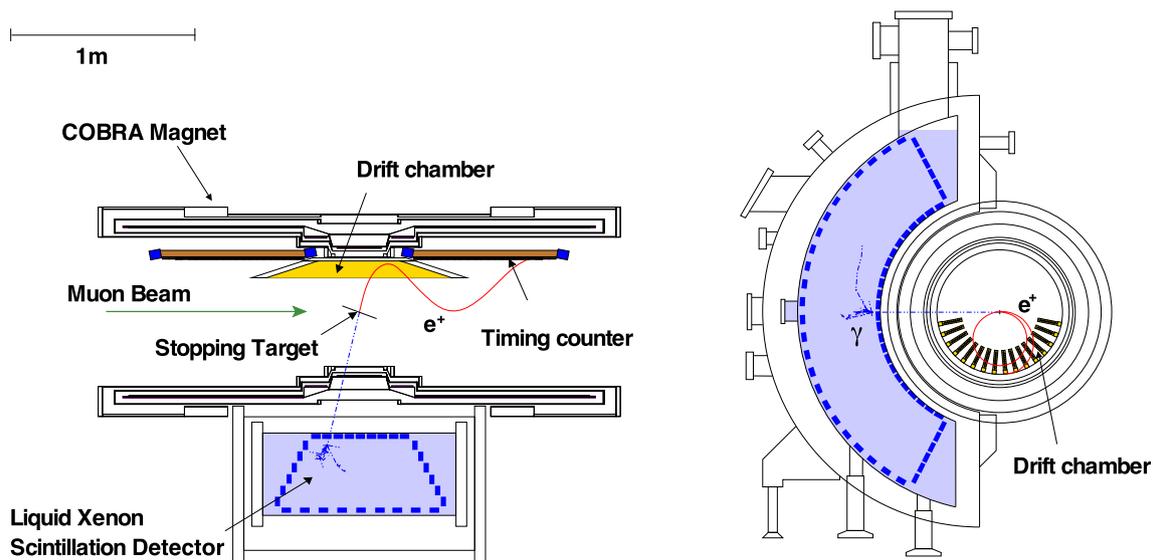


Fig. 1 Schematic view of the MEG detector showing one simulated signal event emitted from the target

This configuration is optimised to sweep low-momentum positrons from Michel decays rapidly out of the magnet, and to keep the bending radius of the positron trajectories only weakly dependent on their emission angle within the acceptance region (see Sect. 3).

The positron track parameters are measured by a set of very low mass Drift Chambers (DCH) designed to minimise the multiple scattering (see Sect. 4). The positron time is measured by a Timing Counter (TC) consisting of scintillator bars read out by PhotoMultiplier Tubes (PMT) (see Sect. 5).

For γ -ray detection, we have developed an innovative detector using Liquid Xenon (LXe) as a scintillation material viewed by PMTs submersed in the liquid. This detector provides accurate measurements of the γ -ray energy and of the time and position of the interaction point (see Sect. 6).

We monitor the performance of the experiment continuously by introducing a variety of calibration methods (see Sect. 7).

We employ a flexible and sophisticated electronics system for triggering the $\mu^+ \rightarrow e^+ \gamma$ candidate events along with calibration data (see Sect. 8).

Special considerations influenced the development of the data acquisition system. We record signals from all detectors with a custom-designed waveform digitiser, the Domino Ring Sampler (DRS) (see Sects. 9 and 10).

We evaluate detector response functions using data with help of detector simulations (see Sect. 11).

In this paper, we mostly refer to a cylindrical coordinate system (r, ϕ, z) with origin in the centre of COBRA. The z -axis is parallel to the COBRA axis and directed along the incoming muon beam. The axis defining $\phi = 0$ (the x -axis of

the corresponding Cartesian coordinate system) is directed opposite to the centre of the LXe detector (as a consequence the y -axis is directed upwards). Positrons move along trajectories of decreasing ϕ coordinate. When required, the polar angle θ with respect to the z -axis is also used. The region with $z < 0$ is referred as upstream, that with $z > 0$ as downstream.

The angular acceptance of the experiment is defined by the size of the LXe fiducial volume, corresponding approximately to $\phi \in (\frac{2}{3}\pi, \frac{4}{3}\pi)$ and $|\cos \theta| < 0.35$, for a total acceptance of $\sim 11\%$. The spectrometer is designed to fully accept positrons from $\mu^+ \rightarrow e^+ \gamma$ decays when the γ -ray falls in the acceptance region.

2 Beam line

2.1 Experimental requirements

In order to meet the stringent background requirements necessary for a high-sensitivity, high-rate coincidence experiment such as MEG, one of the world's highest intensity continuous muon beams, the $\pi E5$ channel at PSI, is used in conjunction with the high-performance MEG beam transport system which couples the $\pi E5$ channel to the COBRA spectrometer (see Sect. 3). This allows more than $10^8 \mu^+/s$, with an optimal suppression of beam-correlated background, to be transported to the ultra-thin stopping-target at the centre of COBRA. The main requirements of such a low-energy surface muon beam are: high-intensity with high-transmission optics at 28 MeV/c, close to the kinematic edge of stopped muon decay and to the maximum intensity of the

production target. A separation quality between muons and beam positrons of $8.1\sigma_\mu$ is so achieved, corresponding to a 12 cm physical separation at the collimator system as shown in Fig. 3. This allows an almost pure muon beam⁴ to propagate to the Beam Transport Solenoid (BTS) (see Sect. 2.3), which couples the injection to the superconducting COBRA spectrometer in an iron-free way. By means of a 300 μm thick Mylar[®] degrader system placed at the intermediate focus, the BTS also minimises the multiple scattering contribution to the beam and adjusts the muon range for a maximum stop-density at the centre of the muon target, placed inside a helium atmosphere, at the centre of the COBRA spectrometer.

2.3 The Beam Transport Solenoid (BTS)

The BTS, a 2.8 m long iron-free superconducting solenoid, with a 38 cm warm-bore aperture, coupling directly to the beam line vacuum, is shown in Fig. 4. With a maximal current of 300 A a field of 0.54 T can be reached, higher than the 0.36 T required for an intermediate focus at the central degrader system for muons. The double-layered 2.63 m long coil is made up of a 45 % Ni/Ti, 55 % Cu matrix conductor with an inductance of 1.0 H so yielding a maximum stored energy of 45 kJ.

On the cryogenic side, a dedicated liquid helium (LHe) port to the central cryoplant is used to transport the supercritical helium at ~ 10 bar to two Joule–Thompson valves

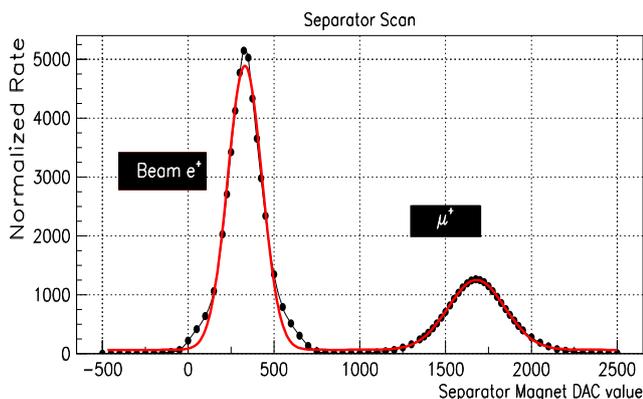


Fig. 3 Separator scan plot, measured post collimator system. The *black dots* represent on-axis, low discriminator threshold intensity measurements at a given separator magnetic field value during the scan, for a fixed electric field value of -195 kV. The *red curve* represents a double Gaussian fit to the data points, with a constant background. A separation of 8.1 muon beam σ_μ is found, corresponding to more than 12 cm separation of beam-spot centres at the collimator. The raw muon peak measured at low discriminator threshold also contains a contribution from Michel positrons due to decays in the detector. These can easily be distinguished by also measuring at high threshold, where only muons are visible on account of their higher energy loss and hence higher pulse-height (Color figure online)

⁴Positron suppression factor ~ 850 under the muon peak.

(in the upper tower in Fig. 4), which are used to liquefy and automatically regulate the level of the 150 ℓ of LHe in the cryostat. A separate liquid nitrogen line cools an outer heat-shield used to minimise the LHe losses.

Furthermore, three iron-free correction dipole magnets (one $\cos\theta$ -type placed upstream of the BTS) one horizontal and one vertical conventional-dipole type, both mounted on the downstream side of the BTS, are employed to ensure an axial injection into both solenoids. The three dipole magnets compensate for a radial asymmetry in the fringe fields induced by the interaction of the stray field from the large-aperture, iron-free COBRA magnet with an iron component of the hall-floor foundations. An example of the optics in Fig. 5 shows the 2σ beam envelopes both vertically and horizontally along the 22 m long beam line, together with the apertures of the various elements and the dispersion trajectory of particles having a 1 % higher momentum (dotted line).

2.4 Spectrometer coupling system

The COBRA spectrometer is sealed at both ends by end-caps which constitute the final part of the beam line, separating the vacuum section from the helium environment inside. These specially designed 1160 mm diameter light-weight (35–40 kg), only 4 mm-thick aluminium constructs

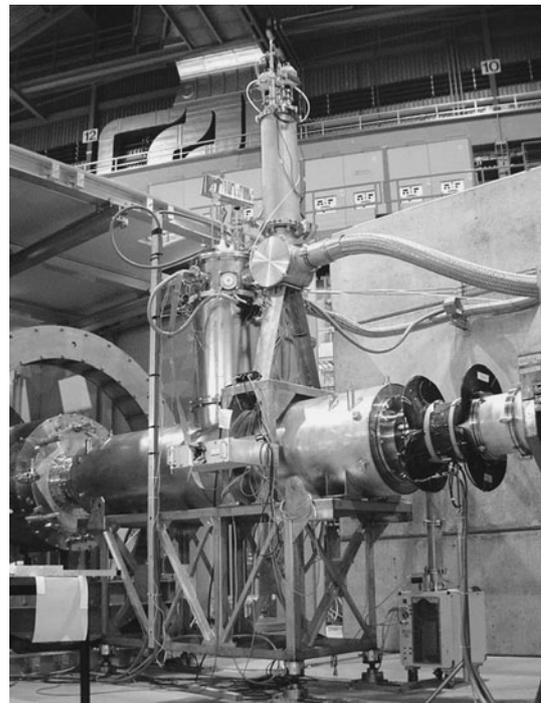
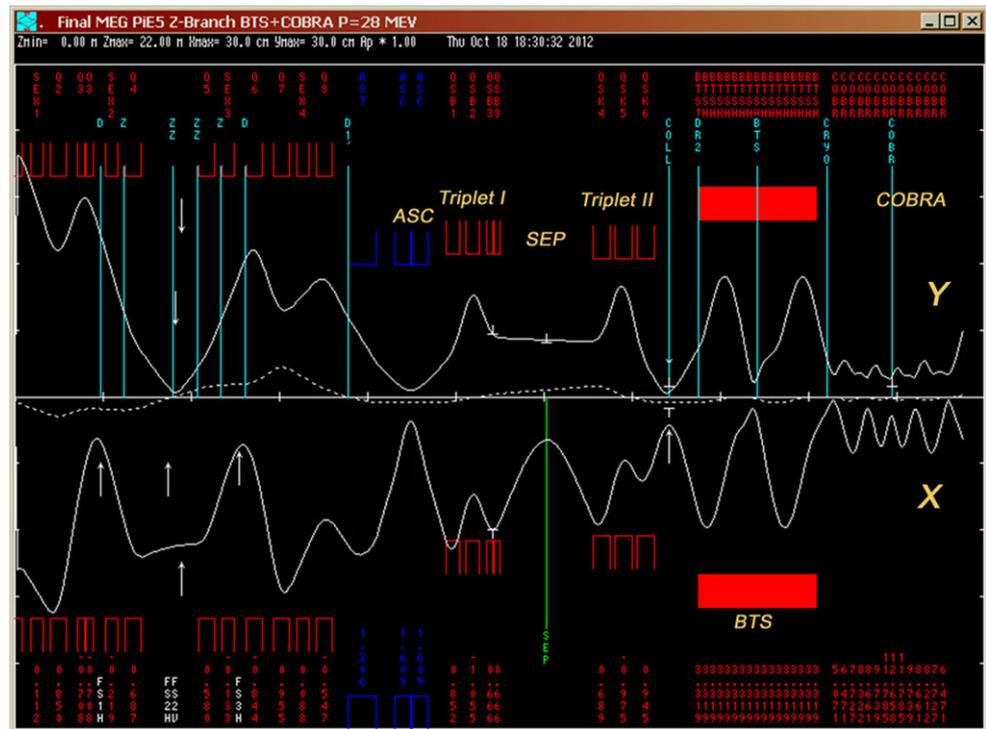


Fig. 4 The Beam Transport Solenoid (BTS), shown coupled to the beam line between Triplet II and the COBRA spectrometer. The main cryogenic tower with the current leads and the Joule–Thompson valve tower with the LHe transfer line from the cryoplant coupled to it as well as the smaller LN₂ transfer line in the background can be seen. The two correction dipole magnets on the downstream side are not shown here

Fig. 5 2nd-order TRANSPORT calculation showing the vertical (*top-half*) and horizontal (*bottom-half*) beam envelopes plotted along the beam-axis from close to target E up to the centre of COBRA, a distance of 22 m. The vertical and horizontal scales differ 33:1. The *coloured elements* show the vacuum chamber apertures of the various beam elements, while the *dotted line* represents the central dispersion trajectory for particles with a 1 % higher momentum (Color figure online)



minimise the material available for possible photon background production, crucial to a background-sensitive experiment. They also integrate the beam vacuum window on the upstream side and an automated thin-windowed (20 μm) bellows insertion system, capable of an extension-length of about 1.7 m. The bellows insertion system is used to introduce various types of targets such as a liquid hydrogen target used in conjunction with the pion beam, or a lithium tetraborate photon production target used together with the C–W accelerator, on the downstream side.

2.5 Commissioning

The commissioning of the MEG beam line was completed with the first “engineering run” in late 2007. During commissioning, several beam optical solutions were studied, including the investigation of both $\pi E5$ branches (“U”-branch and “Z”-branch) as well as different beam-correlated background separation methods: a combined degrader energy-loss/magnetic-field separation method or a crossed ($E \times B$)-field separator (Wien filter). The optimum solution favoured the “Z”-channel with a Wien filter, both from the experimental layout and the rate and background-suppression point of view.

In order to maximise the beam transmission rate and the stopping density in the target, an understanding of the momentum spectrum and channel momentum acceptance is necessary. This was carefully studied by measuring the muon momentum spectrum in the vicinity of the kinematic

edge of stopped pion decay (29.79 MeV/c). As with all phase-space measurements performed at high rate, a scanning technique was used: either one using a small cylindrical “pill” scintillator of 2 mm diameter and 2 mm thickness coupled to a miniature PMT [8], or one using a thick depletion layer Avalanche PhotoDiode (APD), without a scintillator, in the case of measurements in a magnetic field. Both these solutions gave an optimum pulse-height separation between positrons and muons. This allows integrated beam-spot intensity measurements up to 2 GHz to be made.

For the momentum spectrum, taken with the full 7 % FWHM momentum-bite and shown in Fig. 6, the whole beam line was tuned to successive momentum values and the integral muon rate was then determined by scanning the beam-spot using the beam line as a spectrometer. This allows to determine both the central momentum and the momentum-bite by fitting a theoretical distribution, folded with a Gaussian resolution function corresponding to the momentum-bite, plus a constant cloud-muon background over this limited momentum range. The red curve corresponding to the fitted function shows the expected $P^{3.5}$ empirical behaviour below the kinematic edge. The cloud-muon content below the surface-muon peak, which has its origin in pion decay-in-flight in the vicinity of the production target, was determined to be 3.3 % of the surface-muon rate. This was determined from negative muon measurements at 28 MeV/c and using the known pion production cross-section differences [9], since negative surface-muons do not exist due to the formation of pionic atoms

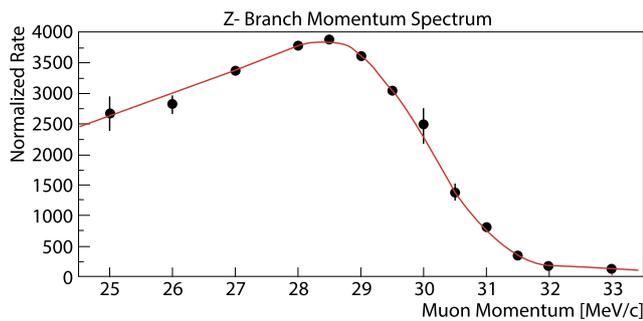


Fig. 6 Measured MEG muon momentum spectrum with fully open slits. Each point is obtained by optimising the whole beam line for the corresponding central momentum and measuring the full beam-spot intensity. The *red line* is a fit to the data, based on a theoretical $P^{3.5}$ behaviour, folded with a Gaussian resolution function corresponding to the momentum-bite plus a constant cloud-muon background (Color figure online)

and hence the small cloud-muon contribution is much easier to measure in the absence of a dominating surface-muon signal.

Another important feature investigated was the muon polarisation on decaying: since the muons are born in a 100 % polarised state, with their spin vector opposing their momentum vector, the resultant angular distribution on decay is an important issue in understanding the angular distribution of the Michel positrons and the influence of detector acceptance. The depolarisation along the beam line, influenced by divergence, materials, magnetic fields and beam contaminants was studied and incorporated into the MEG Monte Carlo, which was used to compare results with the measured angular distribution asymmetry (upstream/downstream events) of Michel positrons. The predicted muon polarisation of 91 % was confirmed by these measurements, which yielded a value of 89 ± 4 % [10, 11]. The main single depolarising contribution is due to the mean polarization of the 3.3 % cloud-muon contamination, which opposes that of the surface muons. However, because the depolarisation in the polyethylene target is totally quenched by the strength of the longitudinal magnetic field in COBRA, a high polarization can be maintained [10].

In order to achieve the smallest beam-spot size at the target, the position of the final degrader as well as the polarity combination of COBRA and the BTS were studied. The smallest beam divergences and hence the smallest beam-spot size at the target were achieved by having a reverse polarity magnetic field of the BTS with respect to COBRA. This is due to the fact that the radial field is enhanced at the intermediate focus between the two magnets, thus effectively increasing the radial focussing power and minimising the growth of the beam divergences through the COBRA volume. The final round beam-spot at the target, with $\sigma_{x,y} \sim 10$ mm is shown in Fig. 7.

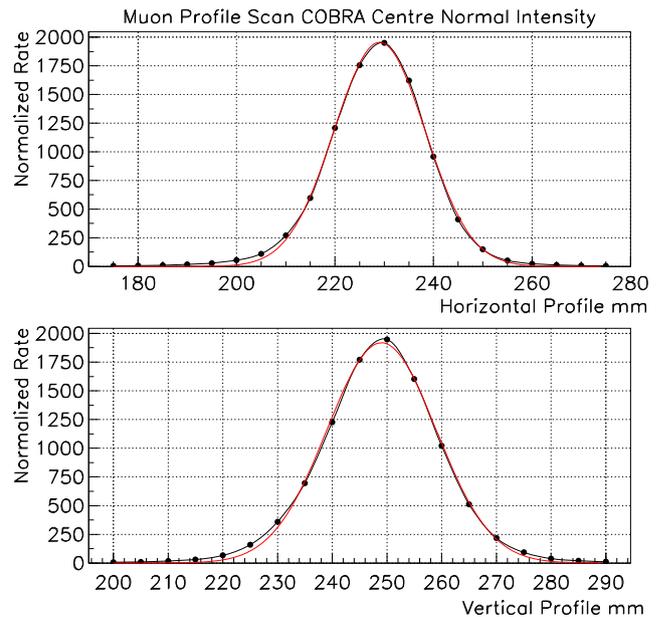


Fig. 7 Scanned horizontal and vertical muon beam-spot profiles measured at the centre of COBRA. The *red lines* show the Gaussian fits to the distributions that yield $\sigma_{x,y} \sim 10$ mm and a total stopping intensity of $\sim 3 \times 10^7 \mu^+/s$ at a proton beam intensity of 2.2 mA and momentum-bite of 5.0 % FWHM (Color figure online)

2.6 Beam line performance

Although the beam line was tuned for maximal rate with minimal positron beam contamination, the final beam stopping intensity was determined by a full analysis of the detector performance and an evaluation of its sensitivity to the $\mu^+ \rightarrow e^+ \gamma$ decay. This yielded an optimal muon stopping rate in the target of $3 \times 10^7 \mu^+/s$.

Furthermore, for both the pion beam and Mott positron beam calibration tunes, mentioned in Sect. 2.1, which involve prompt particle production at the proton target, the central momentum was chosen to give a maximal time-of-flight separation between the wanted particles and the unwanted background particles at the centre of COBRA, based on the use of the accelerator modulo 19.75 ns radio-frequency signal as a clock, which has a fixed time relationship to the production at the proton target. In the case of the Mott positron beam the final momentum should also be close to the signal energy of 52.83 MeV/c. A summary of the final beam parameters is given in Table 1.

2.7 Target

2.7.1 Requirements

Positive muons are stopped in a thin target at the centre of COBRA. The target is designed to satisfy the following criteria:

Table 1 MEG final beam line characteristics

Parameter	Value
Solid Angle $\Delta\Omega$	150 msr
Beam Line length L	22 m
Muon Central Momentum P_μ	28 MeV/c
Pion Central Momentum P_π	-70.5 MeV/c
Positron Central Momentum P_e	53 MeV/c
Muon Momentum-bite $\Delta P/P$	5.0 % (full-width)
Beam-spot size $\sigma_{x,y}$	9–10 mm
beam e^+/μ^+ ratio (Sep. off)	8
μ^+/e^+ separation (Sep. on)	$8.1\sigma_\mu, 12$ cm
Optimal μ^+ stopping intensity	$3 \times 10^7 \mu^+/s$

- 1 Stop >80 % of the muons in a limited axial region at the magnet centre.
- 2 Minimise conversions of γ -rays from radiative muon decay in the target.
- 3 Minimise positron annihilation-in-flight with γ -rays in the acceptance region.
- 4 Minimise multiple scattering and Bremsstrahlung of positrons from muon decay.
- 5 Allow the determination of the decay vertex position and the positron direction at the vertex by projecting the reconstructed positron trajectory to the target plane.
- 6 Be dimensionally stable and be remotely removable to allow frequent insertion of ancillary calibration targets.

Criterion 1 requires a path-length along the incident muon axial trajectory of $\sim 0.054 \text{ g/cm}^2$ for low-Z materials and for the average beam momentum of 19.8 MeV/c (corresponding to a kinetic energy of 1.85 MeV) at the centre of COBRA. To reduce scattering and conversion of outgoing particles, the target is constructed in the form of a thin sheet in a vertical plane, with the direction perpendicular to the target at $\theta \sim 70^\circ$, oriented so that the muon beam is incident on the target side facing the LXe detector (see Fig. 1). The requirement that we be able to deduce the positron angle at the vertex sets the required precision on the knowledge of the position of the target in the direction perpendicular to its plane. This derives from the effect of the track curvature on the inferred angle; an error in the projected track distance results in an error in the positron angle at the target due to the curvature over the distance between the assumed and actual target plane position. This angle error is larger when the positrons are incident at a large angle with respect to the normal vector; for an incidence angle of 45° , an error in the target plane position of 500 μm results in an angle error of ~ 5 mrad. This sets the requirements that the target thickness be below 500 μm (eliminating low-density foam as a target material) and that its position with respect to the DCH spectrometer be measured to $\sim 100 \mu\text{m}$.

2.7.2 Realisation

To meet the above criteria, the target was implemented as a 205 μm thick sheet of a low-density (0.895 g/cm^3) layered-structure of a polyethylene and polyester with an elliptical shape and semi-major and semi-minor axes of 10 cm and 4 cm. The target was captured between two rings of Rohacell[®] foam of density 0.053 g/cm^3 and hung by a Delrin[®] and Rohacell stem from a remotely actuated positioning device attached to the DCH support frame. To allow a software verification of the position of the target plane, as described in Sect. 2.7.5, eight holes of radius 0.5 cm were punched in the foil.

2.7.3 Automated insertion

The positioning device consists of a pneumatically operated, bi-directional cylinder that inserts and retracts the target. Physical constraints on the position of the cylinder required the extraction direction to not be parallel to the target plane; hence the position of the plane when inserted was sensitive to the distance by which the piston travelled. This distance was controlled by both a mechanical stop on the piston travel and a sensor that indicated when the target was fully inserted.

2.7.4 Optical alignment

The position of the target plane is measured every run period with an optical survey that images crosses inscribed on the target plane; this survey is done simultaneously with that of the DCH. The relative accuracy of the target and DCH positions was estimated to be $\pm(0.5, 0.5, 0.8)$ mm in the (x, y, z) directions.

2.7.5 Software alignment

The optical survey was checked by imaging the holes in the target. We define z' as the position parallel to the target plane and perpendicular to the y coordinate and x' as the position perpendicular to the target plane. The y and z' positions are easily checked by projecting particle trajectories to the nominal target plane and determining the z' and y hole coordinates by fitting the illuminations to holes superimposed on a beam profile of approximately Gaussian shape. We note that these target coordinates are not important to the analysis, since a translation of the target in the target plane has no effect on the inferred vertex position. The x' coordinate was checked by imaging the y positions of the holes as a function of the track direction at the target ϕ_e . Figure 8 illustrates the technique. The x' target position is determined by requiring that the target hole images be insensitive to ϕ_e . Figure 9 shows the illumination plot for the 2011 data sample. Figure 10 shows the normalised y distribution of events

with $20^\circ < \phi_e < 30^\circ$ in a slice at fixed z' centred on a target hole (in red) and in adjacent slices (in blue). The difference (in black) is fitted with a Gaussian. Figure 11 shows the ap-

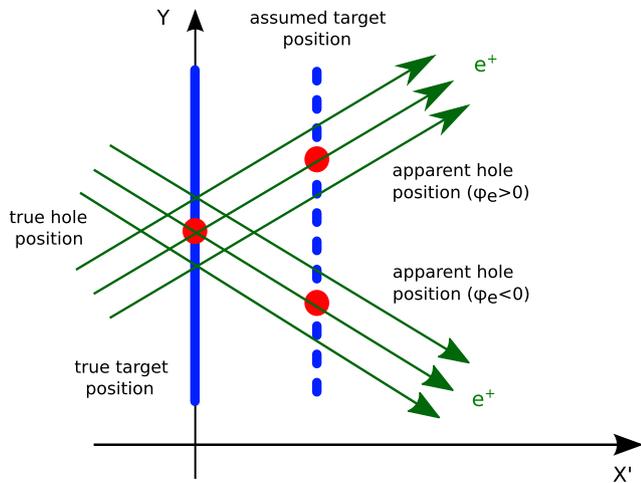


Fig. 8 An illustration of the target hole imaging technique. If there is an offset in the x' direction between the assumed target position and the true target position, the apparent y position of a target hole will depend on the angle of emission at the target ϕ_e . The shift in y is approximately proportional to the shift in the x' direction by $\tan(\phi_e)$

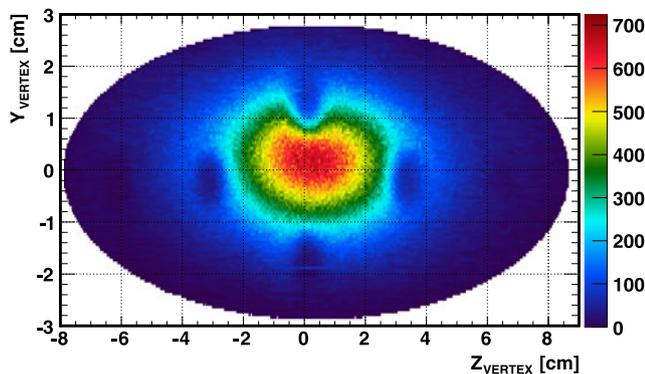


Fig. 9 The target illumination plot for the 2011 data sample

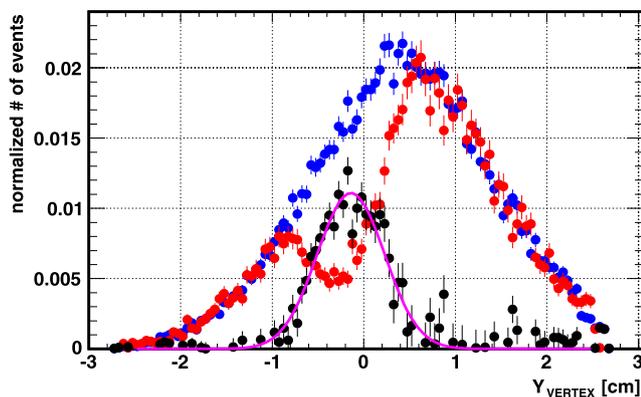


Fig. 10 The normalised y distribution of events with $20^\circ < \phi_e < 30^\circ$ in a slice at fixed z' centred on a target hole (in red) and in adjacent slices (in blue). The difference (in black) is fitted with a Gaussian (Color figure online)

parent y position of one of the target holes versus ϕ_e , fitted with a $\tan(\phi_e)$ function, from which we infer that the x' offset with respect to the optical survey is 0.2 ± 0.1 mm. By imaging all the holes (i.e. finding $x'(y, z')$) both the planarity and the orientation of the target plane was determined.

2.7.6 Multiple scattering contribution from the target

The target itself constitutes the most relevant irreducible contribution to the positron multiple scattering and therefore to the angular resolution of the spectrometer.

The simulated distribution of the μ^+ decay vertex along x' is shown in Fig. 12. The average is slightly beyond the centre of the target. Averaging over the decay vertex position and over the positron direction in the MEG acceptance region, the difference between the positron direction at the μ^+ decay vertex and that at the target boundary follows a Gaussian distribution with $\sigma = 5.2$ mrad.

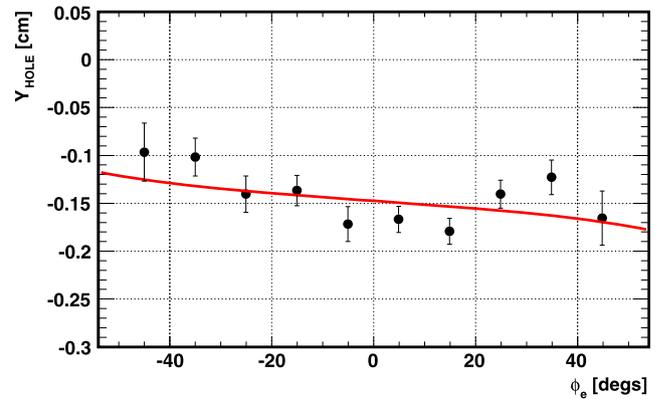


Fig. 11 The apparent y position of a target hole versus ϕ_e fitted with a $\tan(\phi_e)$ function, from which we infer that the x' offset with respect to the optical survey is 0.2 ± 0.1 mm

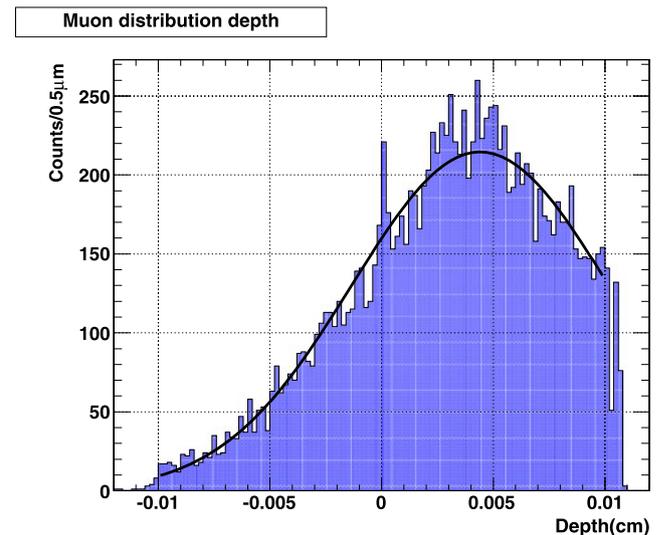


Fig. 12 Distribution of the μ^+ decay vertex along x'

3 COBRA magnet

3.1 Concept

The COBRA (COnstant Bending RAdius) magnet is a thin-wall superconducting magnet generating a gradient magnetic field which allows stable operation of the positron spectrometer in a high-rate environment. The gradient magnetic field is specially designed so that the positron emitted from the target follows a trajectory with an almost constant projected bending radius weakly dependent on the emission polar angle θ (Fig. 13(a)). Only high-momentum positrons can therefore reach the DCH placed at the outer radius of the inner bore of COBRA. Another good feature of the gradient field is that the positrons emitted at $\cos\theta \sim 0$ are quickly swept away (Fig. 13(b)). This can be contrasted to a uniform solenoidal field where positrons emitted transversely would turn repeatedly in the spectrometer. The hit rates of Michel positrons expected for the gradient and uniform fields are compared in Fig. 14 and indicate a significant reduction for the gradient field.

The central part of the superconducting magnet is as thin as $0.197X_0$ so that only a fraction of the γ -rays from the target interacts before reaching the LXe detector placed outside COBRA. The COBRA magnet (see Fig. 15) is equipped with a pair of compensation coils to reduce the stray field around the LXe detector for the operation of the PMTs.

The parameters of COBRA are listed in Table 2.

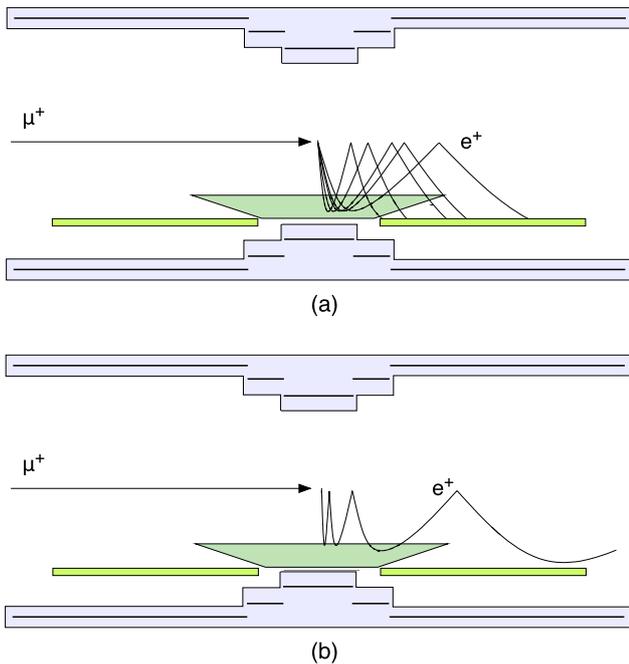


Fig. 13 Concept of the gradient magnetic field of COBRA. The positrons follow trajectories at a constant bending radius weakly dependent on the emission angle θ (a) and those transversely emitted from the target ($\cos\theta \sim 0$) are quickly swept away from the DCH (b)

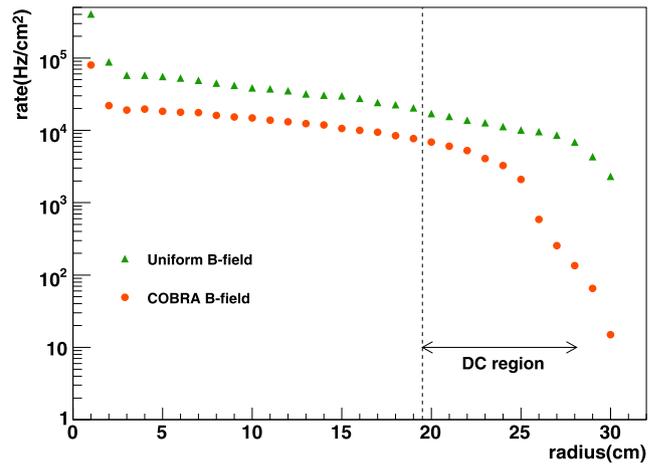


Fig. 14 Hit rate of the Michel positrons as a function of the radial distance from the target in both the gradient and uniform field cases

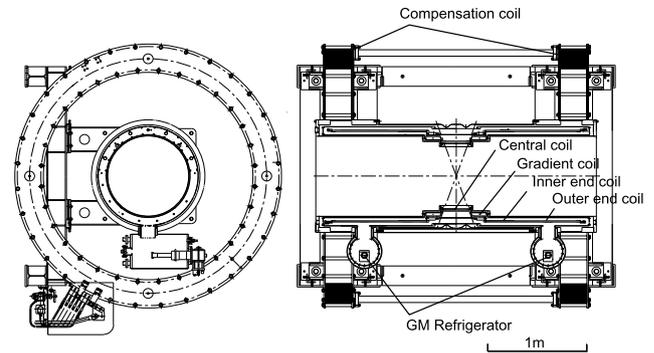


Fig. 15 Cross-sectional view of the COBRA magnet

3.2 Design

3.2.1 Superconducting magnet

The gradient magnetic field ranging from 1.27 T at the centre to 0.49 T at either end of the magnet cryostat is generated by a step structure of five coils with three different radii (one central coil, two gradient coils and two end coils) (Fig. 16). The field gradient is optimised by adjusting the radii of the coils and the winding densities of the conductor. The coil structure is conductively cooled by two mechanical refrigerators attached to the end coils; each is a two-stage Gifford–McMahon (GM) refrigerator [12] with a cooling power of 1 W at 4.2 K. The thin support structure of the coil is carefully designed in terms of the mechanical strength and the thermal conductivity. The basic idea is to make the coil structure as thin as possible by using a high-strength conductor on a thin aluminium mechanical support cylinder. Figure 17 shows the layer structure of the central coil, which has the highest current density. A high-strength conductor is wound in four layers inside the 2 mm-thick aluminium support cylinder. Pure aluminium strips with a thickness of

Table 2 Parameters of the COBRA magnet

Coil	Central	Gradient	Inner end	Outer end	Compensation
Conductivity	Super	Super	Super	Super	Resistive
Inner dia. (mm)	699.1	809.1	919.1	919.1	2210
Outer dia. (mm)	711.6	820.6	929.5	929.5	2590
Length (mm)	240.3	110.4	189.9	749.2	265
z-coordinate of coil centre (mm)	0.8	±235	±405.4	±874.95	±1190
Layers	4	4	3	3	14
Turns (total)	1068	399	240	1548	280
Winding density (Turns/m)	4444.4	3614.1	1263.8	2066.2	1056.6
Inductance (H)	1.64	0.62	0.35	2.29	0.54
Current (A)	359.1	359.1	359.1	359.1	319.2
Energy <i>E</i> (kJ)	106	40	23	148	35
Weight <i>M</i> (kg)	9	4	7	28	1620
<i>E</i> / <i>M</i> (kJ/kg)	11.8	10.0	3.3	5.3	0.02

100 μm are attached on the inner surface of the coil structure in order to increase the thermal conductivity. Several quench protection heaters are attached to all coils in order to avoid a local energy dump in case of a quench. The high thermal conductivity of the coil structure and the uniform quench induced by the protection heaters are important to protect the magnet.

A high-strength aluminium stabilised conductor is used for the superconducting coils to minimise the thickness of the support cylinder. The cross-sectional view of the conductor is shown in Fig. 17. A copper matrix NbTi multifilamentary core wire is clad with aluminium stabiliser which is mechanically reinforced by means of “micro-alloying” and “cold work hardening” [13, 14]. The overall yield strength of the conductor is measured to be higher than 220 MPa at 4.4 K. The measured critical current of the con-

ductor as a function of applied magnetic field is shown in Fig. 18 where the operating condition of COBRA is also shown (the highest magnetic field of 1.7 T is reached in the coils with the operating current of 359.1 A). This indicates that COBRA operates at 4.2 K with a safety margin of 40 % for the conductor.

3.2.2 Compensation coils

The COBRA magnet is equipped with a pair of resistive compensation coils in order to reduce the stray field from the superconducting magnet around the LXe detector (Fig. 15). The stray field should be reduced below 5×10^{-3} T for the operation of the PMTs of the LXe detector. The stray field from the superconducting magnet is efficiently cancelled around the LXe detector since the distribution of the magnetic field generated by the compensation coils is similar to that of the superconducting magnet around that region

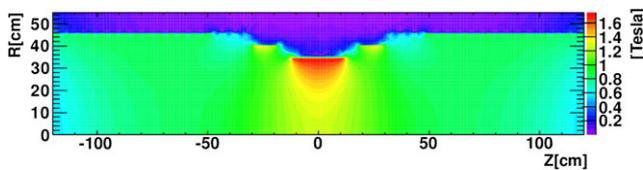


Fig. 16 Gradient magnetic field generated by the COBRA magnet

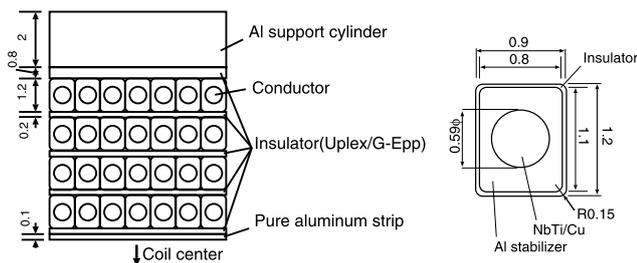


Fig. 17 Layer structure of the central coil (left) and the cross-sectional view of the conductor (right). Units in mm

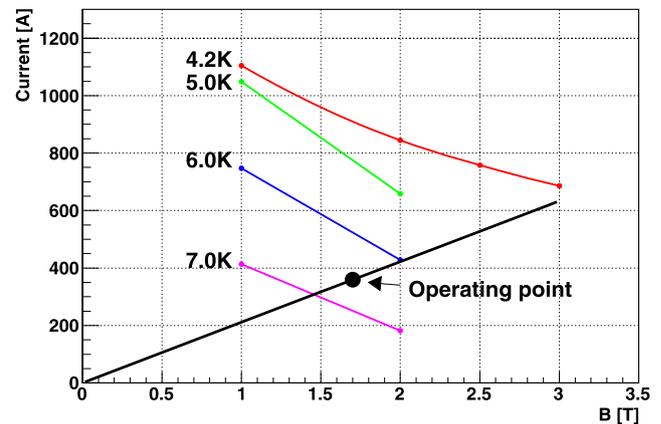


Fig. 18 Measured critical current of the conductor as a function of the applied magnetic field. The operating condition of COBRA is also shown

but the sign of the field is opposite. Figure 19 shows that the expected distribution of the residual magnetic field in the vicinity of the LXe detector is always below 5×10^{-3} T.

3.3 Performance

3.3.1 Excitation tests

The superconducting magnet was successfully tested up to 379 A coil current, which is 5.6 % higher than the operating current. The temperature and stress distribution are monitored by using resistive temperature sensors and strain gauges attached to the coils, respectively. The temperatures of all the coils are found to stay lower than 4.5 K even at the highest coil current.

The strains of the coils are measured by the strain gauges up to 379 A coil current at the central coils and the support cylinder where the strongest axial electromagnetic force is expected. We observe a linear relation between the strains and the square of the coil current, which is proportional to the electromagnetic force acting on the coils. This indicates that the mechanical strength of the coils and the support strength are sufficient up to 379 A.

The effect of the compensation coils is also measured. The gain of the PMTs of the LXe detector, measured using LEDs (see Sect. 7.1.1), is suppressed by a factor of 50 on average in the stray magnetic field from the superconducting magnet, whereas it recovers up to 93 % of the zero field value with the compensation coils.

3.3.2 Mapping of magnetic field

The magnetic field of COBRA was measured prior to the experiment start with a commercial three-axis Hall probe

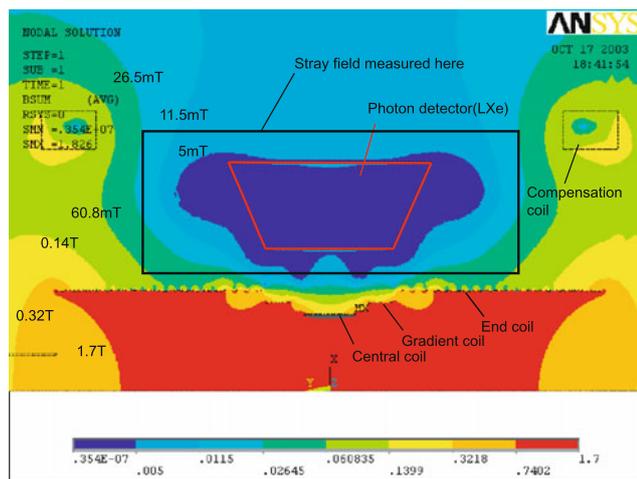


Fig. 19 Distribution of the residual magnetic field around the LXe detector. The PMTs of the LXe detector are placed along the *trapezoidal box* shown in this figure

mounted on a wagon moving three-dimensionally along z , r and ϕ . The ranges of the measurement are $|z| < 110$ cm with 111 steps, $0^\circ < \phi < 360^\circ$ with 12 steps and $0 \text{ cm} < r < 29$ cm with 17 steps, which mostly cover the positron tracking volume. The probe contains three Hall sensors orthogonally aligned and is aligned on the moving wagon such that the three Hall sensors can measure B_z , B_r and B_ϕ individually. The field measuring machine, the moving wagon and the Hall probe are aligned at a precision of a few mrad using a laser tracker. However, even small misalignments could cause a relatively large effect on the secondary components, B_r and B_ϕ due to the large main component B_z , while the measurement of B_z is less sensitive to misalignments. Only the measured B_z is, therefore, used in the analysis to minimise the influence of misalignments and the secondary components B_r and B_ϕ are computed from the measured B_z using the Maxwell equations as

$$B_\phi(z, r, \phi) = B_\phi(z_0, r, \phi) + \frac{1}{r} \int_{z_0}^z \frac{\partial B_z(z', r, \phi)}{\partial \phi} dz', \quad (1)$$

$$B_r(z, r, \phi) = B_r(z_0, r, \phi) + \int_{z_0}^z \frac{\partial B_z(z', r, \phi)}{\partial r} dz'. \quad (2)$$

The computations require the measured values for B_r and B_ϕ only at the plane defined by $z = z_0$. This plane is chosen at $z_0 = 1$ mm near the symmetry plane at the magnet centre where B_r is measured to be small ($|B_r| < 2 \times 10^{-3}$ T) as expected. The effect of the misalignment of the B_ϕ -measuring sensor on $B_\phi(z_0, r, \phi)$ is estimated by requiring the reconstructed B_r and B_ϕ be consistent with all the other Maxwell equations.

Since $r = 0$ is a singularity of Eqs. (1) and (2) and the effect of the misalignment of the B_z -measuring sensor is not negligible at large r due to large B_r , the secondary components are computed only for $2 \text{ cm} < r < 26 \text{ cm}$. The magnetic field calculated with a detailed coil model is used outside this region and also at $z < -106 \text{ cm}$ and $z > 109 \text{ cm}$ where the magnetic field was not measured.

The magnetic field is defined at the measuring grid points as described above and a continuous magnetic field map to be used in the analysis is obtained by interpolating the grid points by a B-spline fit.

4 Drift chamber system

4.1 Introduction

The Drift Chamber (DCH) system of the MEG experiment [15] is designed to ensure precision measurement of positrons from $\mu^+ \rightarrow e^+ \gamma$ decays. It must fulfil several stringent requirements: cope with a huge number of Michel positrons due to a very high muon stopping rate up to

$3 \times 10^7 \mu^+/s$, be a low-mass tracker as the momentum resolution is limited by multiple Coulomb scattering and in order to minimise the accidental γ -ray background by positron annihilation-in-flight, and finally provide excellent resolution in the measurement of the radial coordinate as well as in the z coordinate.

The DCH system consists of 16 independent modules, placed inside the bore of COBRA (see Sect. 3) aligned in a half circle with 10.5° intervals covering the azimuthal region $\phi \in (191.25^\circ, 348.75^\circ)$ and the radial region between 19.3 cm and 27.9 cm (see Fig. 20).

4.2 Design of DCH module

All 16 DCH modules have the same trapezoidal shape (see Fig. 21) with base lengths of 40 cm and 104 cm, without any supporting structure on the long side (open frame geometry) to reduce the amount of material. The modules are mounted with the long side in the inner part of the spectrometer (small radius) and the short one positioned on the central coil of the magnet (large radius).

Each module consists of two detector planes which are operated independently. The two planes are separated by a middle cathode made of two foils with a gap of 3 mm. Each plane consists of two cathode foils with a gas gap of 7 mm. In between the foils there is a wire array containing alternating anode and potential wires. They are stretched in the axial direction and are mounted with a pitch of 4.5 mm. The longest wire has a length of 82.8 cm, the shortest of 37.6 cm. The anode-cathode distance is 3.5 mm. The two wire arrays in the same module are staggered in radial direction by half a drift cell to resolve left-right ambiguities (see Fig. 22). The isochronous and drift lines in a cell calculated with GARFIELD [16] are shown in Fig. 23 for $B = 1.60$ T.

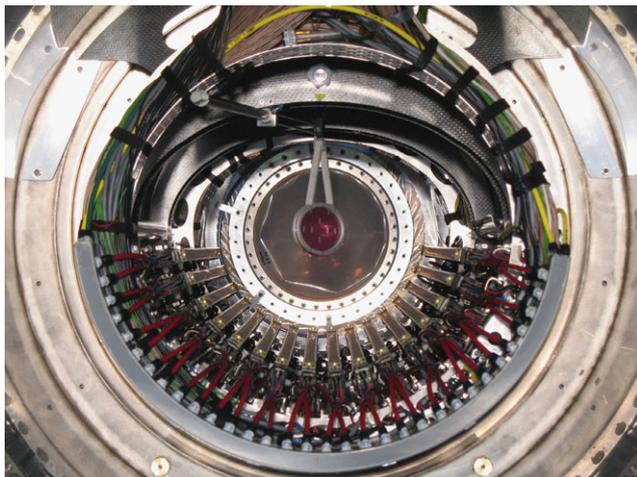


Fig. 20 View of the DCH system from the downstream side of the MEG detector. The muon stopping target is placed in the centre, the 16 DCH chamber modules are mounted in a half circle

The two detector planes are enclosed by the so-called hood cathode. The middle as well as the hood cathodes are made of a $12.5 \mu\text{m}$ -thick polyamide foil with an aluminium deposition of 2500 \AA .

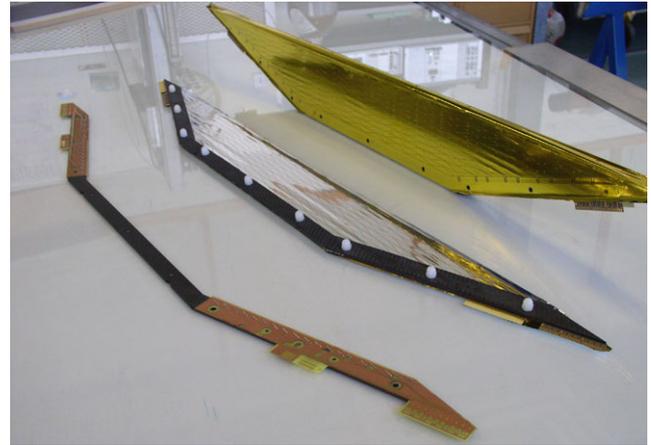


Fig. 21 Main components of a DCH module: anode frame with wires (front), middle cathode and hood cathode (back)

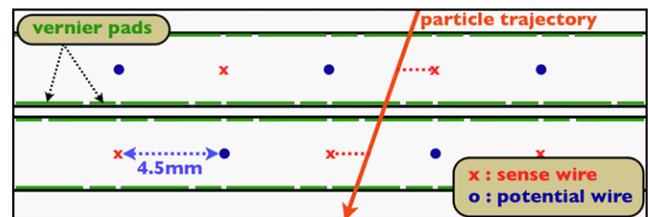


Fig. 22 Schematic view of the cell structure of a DCH plane

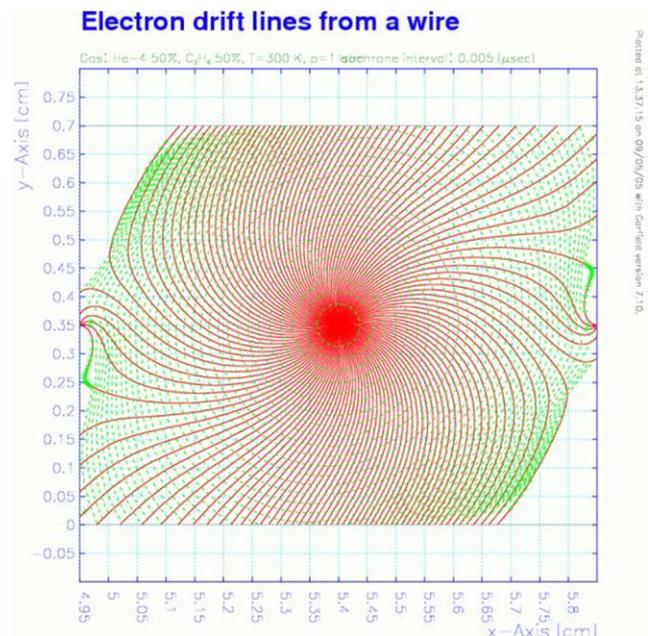


Fig. 23 Isochrons (dashed green) and drift (solid red) lines of a cell with $B = 1.60$ T (Color figure online)

Thanks to such a low-mass construction and the use of a helium-based gas mixture, the average amount of material in one DCH module sums up to only $2.6 \times 10^{-4} X_0$, which totals $2.0 \times 10^{-3} X_0$ along the positron track.

All frames of the DCH modules are made of carbon fibre. Despite being a very light material, it offers the advantage of being rather rigid. Consequently, there is only a negligible deformation of the frames due to the applied pretension compensating the mechanical tension of the wires (1.7 kg) or the cathode foils (12.0 kg).

4.3 Charge division and vernier pads

A preliminary determination of the hit z coordinate, z_{anode} , is based on the principle of charge division. To achieve this goal, the anodes are resistive Ni–Cr wires with a resistance per unit length of 2.2 k Ω /m. Firstly, the z coordinate, is measured from the ratio of charges at both ends of the anode wire to an accuracy better than 2 % of the wire length.

Secondly the information from the cathodes is used to improve the measurement, by using a double wedge or vernier pad structure [17–19] with pitch $\lambda = 5$ cm, etched on the cathode planes on both sides of the anode wires (see Fig. 24) with a resistance per unit length of 50 Ω /m.

Due to the double wedge structure, the fraction of charge induced on each pad depends periodically on the z coordinate. The period λ is solved by means of z_{anode} . In total there are four cathode pad signals for each wire. To increase the capability of this method the vernier pads of one cathode plane is staggered by $\lambda/4$ in the axial direction with respect to its partner plane.

4.4 Counting gas

In order to reduce the amount of material along the positron trajectory, the helium-based gas mixture He:C₂H₆ = 50:50 is adopted as a counting gas [20, 21]. For this mixture the disadvantages of small primary ionisation and large diffusion of helium is balanced by the high primary ionisation, good quenching properties and high voltage (HV) stability



Fig. 24 Double wedge or vernier pads etched on the aluminium deposition of the cathode foil (Color figure online)

of ethane. A reliable HV stability is very important as the DCHs are operated in a high-rate environment and with a very high gas gain of several times 10^5 . The main advantages of this mixture are a rather fast and saturated drift velocity, ~ 4 cm/ μ s, a Lorentz angle smaller than 8° and a long radiation length ($X_0/\rho = 640$ m).

4.5 Pressure regulation system

A specially designed pressure control system manages the gas flows and gas pressures inside the DCH modules and inside the bore of COBRA, which is filled with almost pure He to also reduce multiple scattering along the positron path outside the DCH volume. The pressure regulation value, defined as the pressure difference between inside and outside the DCH, is on average 1.2×10^{-5} bar regulated to a precision better than 0.2×10^{-5} bar. This sensitivity is required as larger pressure differences would induce large distortions in the thin cathode foils and consequently lead to changes in the anode-cathode distances resulting in huge gain inhomogeneities along the anode wire as well as from wire to wire.

4.6 HV system

The DCHs are operated with a voltage of ~ 1800 V applied to the anode wires. In the first years of operation a HV system based on the MSCB system (see Sect. 10.1.4) was used. In this case one power supply [22] was used as the primary HV device which fed 16 independent regulator units with 2 channels each. The primary HV was daisy-chained along the 16 regulator modules.

In 2011 the MSCB-based HV system was replaced by a new commercial one. Two plug-in units [23] with 16 independent channels each, operated in a Wiener Mpod mini crate, were chosen to reduce the high-frequency electronic noise contribution (mainly 14 MHz).

4.7 Readout electronics

Cathode and anode signals are conducted via printed circuit board tracks to the pre-amplifiers, which are mounted at the frames of the DCH modules. The cathode signal is directly connected to the input of the pre-amplifiers, whereas the anode signal is decoupled by a 2.7 nF capacitor as the anode wire is on positive HV potential.

The pre-amplifier was custom-designed to match the geometrical boundary conditions and electrical requirements of the DCH system. The circuit is based on two operational amplifiers with feedback loops and protection diodes. The total gain of the pre-amplifier is ~ 50 and the bandwidth is 140 MHz for the cathode signals and 190 MHz for the anode signals. In addition, the anode output is inverted to match the

required positive input of the DRS read-out chip. The noise contribution of the pre-amplifier is 0.74 mV.

All output signals from the pre-amplifiers are individually transferred by non-magnetic, coaxial cables (Radial MIL-C-17/93-RG178) to the back-end electronics. At the end-cap of the COBRA spectrometer there is a feedthrough patch panel which has resistive dividers splitting the anode signals into two outputs in the proportion 1:9. The larger signal goes to the DRS input, whereas the smaller one is amplified by a feedback amplifier and summed with several other anode outputs to form a DCH self-trigger signal.

4.7.1 Alignment tools and optical survey

As the DCH is a position-sensitive detector, the position of the wires and the pads must be known very precisely. The alignment procedure consists of two parts, a geometrical alignment and a software alignment using cosmic rays and Michel positrons.

During the construction of each single frame the positions of the anode wires and of the cathode pads are measured with respect to an alignment pin located at the bottom left edge of each frame. Each cathode hood is equipped with two target marks (cross hairs) and two corner cube reflectors placed on the most upstream and most downstream upper edge of the cathode hood (see Fig. 25). After assembly of a DCH module the position of these identification marks was measured with respect to the alignment pin, which allows the alignment of the different frames within the assembled DCH module and acts as a reference for the wire as well as the pad positions. Even though the chamber geometry is an open frame construction a geometrical precision of $\sim 50 \mu\text{m}$ was achieved over the full length of the chamber for the position of the wires and the pads.

The support structure between two adjacent DCH modules is also equipped with target marks and corner cube reflectors on both the downstream and the upstream sides.

After the installation of the DCH system inside COBRA, an optical survey was performed. The positions of the target marks and of the corner cube reflectors on the cathode hoods



Fig. 25 Corner cube reflectors and target marks on the downstream side of the DCH modules

and on the support structure were measured with respect to the beam axis and the downstream flange of the magnet using a theodolite (target marks) and a laser tracker system (corner cube reflectors) with a precision of 0.2 mm.

4.7.2 Track-based alignment: the Millipede

The optical survey measurements described above allow the determination of the wire positions in the absolute coordinate system of the experiment. To achieve a better accuracy, we developed a track-based alignment procedure consisting of three steps: (i) an internal alignment of DCH modules; (ii) the placement of the so-obtained DCH frame in the spectrometer, with a fine adjustment of the relative DCH-target-COBRA position; (iii) the alignment of the spectrometer with respect to the LXe detector (see Sect. 4.7.4).

The first step is accomplished by recording cosmic ray tracks during the COBRA shutdown periods, to make the internal DCH alignment independent of the magnetic field map. Events are triggered by dedicated scintillation Cosmic Ray Counters (CRC), located around the magnet cryostat, which also provide a reference for the drift-time measurement. The alignment procedure utilises the reconstructed position of hit DCH modules to minimise the residuals with respect to straight muon tracks, according to the Millipede algorithm [24]. Global parameters, associated with displacements of each DCH module from its initial position, are obtained with an accuracy better than $150 \mu\text{m}$ for each coordinate.

The second step relies on consistency checks based on a sample of double-turn Michel positron tracks in the magnet volume. The relative DCH-COBRA position is fine-tuned until the track parameters related to each single turn (which are reconstructed as independent tracks by the fitting algorithm) are consistent within the fit uncertainties.

4.7.3 Michel alignment

The chamber positions from the optical survey are also cross-checked by using Michel positrons from μ^+ decays. Both the radial chamber alignment and the alignment in z are analysed by looking at the difference between the measured coordinate and that predicted by a fitted trajectory for each chamber plane. Because this diagnostic is insensitive to the absolute chamber positions, the overall location of the support structure in space is constrained by fixing the average radial and z positions over all planes to those of the optical survey. Each chamber plane is then shifted to bring the mean of its pull distribution closer to the average over all planes. The shifts are found to be consistent with the results in Sect. 4.7.2. Additional effects are investigated with Michel positrons such as tilts in the $(r-z)$ plane, and displacement in the $(x-y)$ plane. However no such effects have been found.

4.7.4 Alignment with the LXe detector

Cosmic rays penetrating both the DCHs and the LXe detector are used to measure the position of the DCH frame relative to the LXe detector. From the distribution of the difference between the position of incidence on the LXe detector reconstructed by the PMTs and that extrapolated from the DCHs, the two detectors can be aligned relative to each other with an accuracy of 1 mm in the horizontal and vertical directions.

4.8 Calibrations

4.8.1 Time calibration

Systematic time offsets between signals are calibrated by looking at the distribution of the drift times. It is a broad distribution, whose width is the maximum drift time in the cell, but it has a steep edge for particles passing close to the sense wire. The position of this edge shows systematic offsets for different channels (wires and pads), even in the same cell, due to different delays in the electronics. The determination of the edge position allows us to realign these offsets.

An essential ingredient of this procedure is the determination of the track time, to be subtracted from the time of each hit for determining the drift time. If positron tracks from muon decays are used, the time can be provided by the track reconstruction itself (see Sect. 4.10). Yet, a better performance is obtained if cosmic rays are used, with the track time provided with high precision by the CRC and corrected for the time of flight from the reference counter to the cell. This approach disentangles the calibration from the positron tracking procedure, and guarantees a relative time alignment at the level of 500 ps, comparable with the counter time resolution. The track reconstruction uncertainty due to these residual offset errors is negligible with respect to other contributions.

The same procedure can be used for both wire and pad signals. Alternatively, to avoid problems due to the smaller size of pad signals, wire offsets can be aligned first, and then the average time difference of wire and pad signal peaks in the same cell can be used to align the pads with respect to the wires.

4.8.2 z coordinate calibration

A variety of electronic hardware properties influence the accuracy of the anode wire and cathode pad measurements of the z coordinate. As discussed in Sect. 4.3, the hit z coordinate is first estimated by charge division using the anode wire signals and then refined using the charge distributions induced on the cathode pads. The anode charge asymmetry is defined as

$$A^{\text{anode}} \equiv \frac{Q_d^{\text{anode}} - Q_u^{\text{anode}}}{Q_d^{\text{anode}} + Q_u^{\text{anode}}}, \tag{3}$$

where $Q_{u(d)}^{\text{anode}}$ is the measured anode charge at the upstream (downstream) end. The cathode pad charge asymmetry is defined similarly with $Q_{u(d)}^{\text{cathode}}$ replacing $Q_{u(d)}^{\text{anode}}$, where *cathode* is either the *middle* or the *hood* cathode of the cell. Calculations show that the cathode charge asymmetry depends sinusoidally on z to high precision. One group of calibrations involves adjusting the parameters entering into the anode z calculation; this improve the probability of obtaining the correct pad period from the initial determination of z .

The anode z coordinate is given by:

$$z_{\text{anode}} = \left(\frac{L}{2} + \frac{R}{\rho} \right) A^{\text{anode}}, \tag{4}$$

where L is the wire length, R the pre-amplifier input impedance, and ρ the wire resistivity.

Figure 26 demonstrates the fundamental calibration tool: the measured cathode charge asymmetry versus the measured z_{anode} coordinate for each cell. The sinusoid period is forced to the pad pitch ($\lambda = 5$ cm) by adjusting the factor $(\frac{L}{2} + \frac{R}{\rho})$ and its phase is forced to be 0 for $z_{\text{anode}} = 0$ by calibrating the relative gain between anode charges Q_d^{anode} and Q_u^{anode} .

A calibration of the relative gain between cathode pads is performed to achieve optimal z resolution. This procedure uses plots similar to Fig. 26, except that z_{anode} is replaced with the more precisely determined coordinate from fitted tracks z_{track} . A vertical offset indicates a relative difference between upstream and downstream cathode pad gains, which is zeroed by adjusting the cathode pad gains. The systematic uncertainty in the cathode pad gains from this procedure is estimated by comparing the results using z_{track} and z_{anode} as the abscissa, to be $\sigma_z^{\text{DCH}} = 70 \mu\text{m}$.

The chambers, being operated at slight overpressure (see Sect. 4.5), experience bowing of the outer cathode hood foils. The bulge of the hood is largest near $z = 0$ and at middle cell numbers because the foil is only attached at the

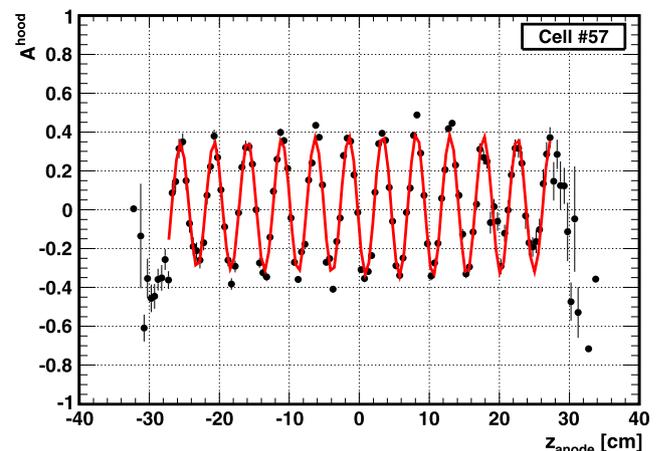


Fig. 26 A fit to the cathode pad asymmetry versus z_{anode} for the cathode hood of a particular drift cell

edges of the chamber. This means that the two cathode foils are not equidistant from the anode wires, but instead the distance varies as a function of z even within the same cell. The induced charge, as well as the magnitude of the charge asymmetry, also differ between the two planes depending on z .

This effect is calibrated by correcting for the mismatch between asymmetry amplitudes. For a given cell,

$$\frac{Q_{\text{middle}}}{Q_{\text{hood}}} = \frac{Q_u^{\text{middle}} + Q_d^{\text{middle}}}{Q_u^{\text{hood}} + Q_d^{\text{hood}}}, \tag{5}$$

which is related to the ratio of asymmetry magnitudes for geometric reasons, exhibits a clear and expected dependence on z , as shown in Fig. 27. The dependence for each wire and a preliminary evaluation of the pad z coordinate are used to rescale the asymmetries on an event-by-event basis for a final evaluation of the pad z coordinate. The charge ratio dependence of the charge asymmetry magnitude ratio is obtained by a scatter plot of those two variables summed over all wires and values of z .

4.9 Tracking

A chain of software algorithms first measures the arrival time and charge on each wire and pad associated with the passage of a charged particle through a cell. The arrival times and charges are converted into spatial coordinates. The information from each cell, referred to as a hit, is cross-checked with other hits in the same chamber, and groups of hits consistent with coming from the same particle are then collected into clusters. Patterns of clusters consistent with coming from the same particle are then collected into tracks.

Finally, a list of clusters associated with each track is passed to a Kalman filter to fit a trajectory. The fit incorpo-

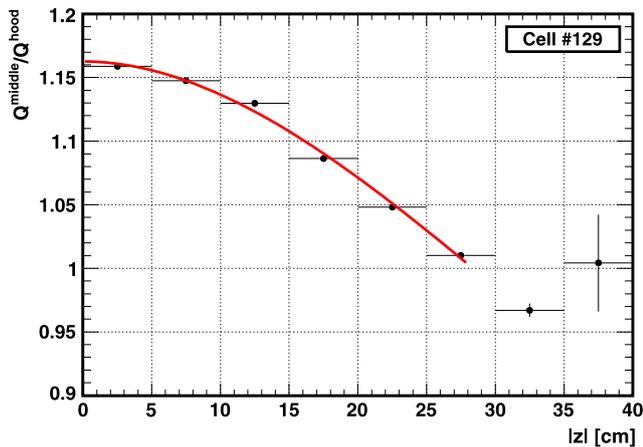


Fig. 27 A plot of the charge ratio versus $|z|$ of a particular drift cell. This is fit to a quadratic polynomial to correct for the relative cathode-to-hood gain

rates average energy loss, multiple scattering, and the measured magnetic field. It also attempts to link multiple turns consistent with belonging to a single trajectory. The Kalman filter provides the track momentum, decay vertex, positron emission angles, the predicted impact location at the TC and the total path length from the target to the TC. An event-by-event indicator of the resolution in each of these quantities is also given.

4.10 Performance

This section reports on the DCH performance regarding spatial, angular, momentum and vertex resolutions.

4.10.1 Single hit resolution

We refer to the uncertainty in a single hit position measurement as the intrinsic resolution, i.e. not affected by multiple scattering. A technique for measuring the intrinsic radial coordinate resolution is illustrated in Fig. 28. The method considers clusters with exactly one hit in each chamber plane, on adjacent cells. A local track circle in the (x, y) plane is calculated from the position of this cluster and the two neighbouring clusters. Each hit is then propagated to the central chamber plane using this trajectory and the difference in the two radial coordinates at the central plane is interpreted as a measure of the radial intrinsic resolution. Figure 29 shows the distribution of $R_{\text{plane } 0} - R_{\text{plane } 1}$. This distribution is fit to the convolution of two double Gaussian functions.

The reason for using a double Gaussian function is that the single hit resolution for drift chambers usually has large tails, which are due to the non-Gaussian fluctuations of the primary ionisation, to the large diffusion and the irregular shape of the time-to-distance relations for hits produced in the peripheral regions of drift cells, to wrongly left/right assignments, frequent for hits at small drift distances, to wrong drift time measurements in some unusually noisy signals, etc. Along with the non-Gaussian material effects (multiple scattering and dE/dx), it produces tails in the track parameter resolutions (angles, momentum and vertex).

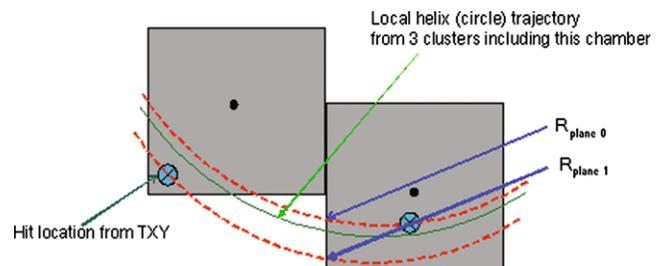


Fig. 28 A diagram of the technique for measuring the intrinsic radial coordinate resolution

The result is a radial coordinate resolution of $\sigma_r^{\text{DCH}} = 210 \mu\text{m}$ in the core (87 %) and $\sigma_r^{\text{DCH}} = 780 \mu\text{m}$ in the tail. The design resolution was parametrised by a single Gaussian with $\sigma_r^{\text{DCH}} = 200 \mu\text{m}$.

A technique for measuring the intrinsic z coordinate resolution is illustrated in Fig. 30. Two-hit clusters belonging to a track segment are selected in the same way described for the radial coordinate resolution measurement to compare the two measured z coordinates. The calculated track circle in the (x, y) plane is used to define a local coordinate system, with the centre of the circle defining the origin. Within this coordinate system, the polar angle ϕ is calculated for each overall cluster position and for each hit within the cluster, as shown in Fig. 30(a). From the three cluster coordinates, a quadratic trajectory for z as a function of the local ϕ is computed. This trajectory is used to project the measured z coordinates of the two hits in question to a common ϕ and the resulting distribution of the difference of the projected z coordinates, $z_{\text{late}} - z_{\text{early}}$, is used to infer the intrinsic z resolution, as pictured in Fig. 30(b). This distribution, shown in Fig. 31, is fit with the convolution of two double Gaussian functions, resulting in a z coordinate resolution of $\sigma_z^{\text{DCH}} = 800 \mu\text{m}$ in the core (91 %) and $\sigma_z^{\text{DCH}} = 2.1 \text{ mm}$ in the tail. The largest known contribution to the z coordinate resolution comes from the stochastic fluctuations of

the baseline in the presence of noise; this is estimated to be $\sigma_{z,\text{noise}}^{\text{DCH}} = 550 \mu\text{m}$ on average. The design z coordinate resolution was $\sigma_z^{\text{DCH}} = 300 \mu\text{m}$.

4.10.2 Angular resolution

The resolutions in the measurements of the positron angles at the target are measured from data by exploiting events where the positron makes two turns in the DCHs. Each turn is treated as an independent track, fitted, and propagated to the beam line where the track angles are compared. The distributions of the difference of the two measured angles in double-turn events are shown in Fig. 32 for θ_e and in Fig. 33 for ϕ_e . The resolution in each turn is assumed to be the same and these distributions are fit to the convolution of a single (θ_e) or double (ϕ_e) Gaussian functions. These functions represent the resolution function of the positron angles.

According to Monte Carlo studies, this method provides a significant overestimate of the true resolution. After cor-

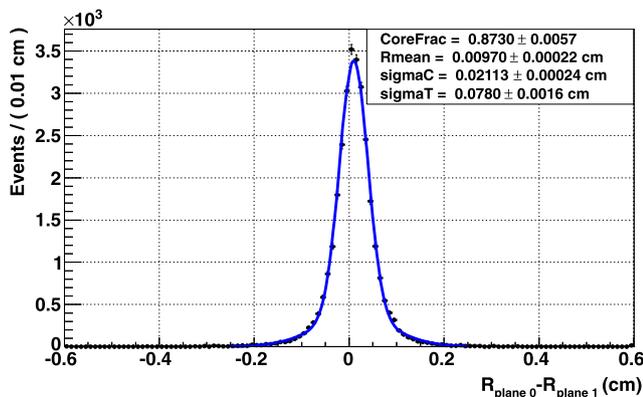
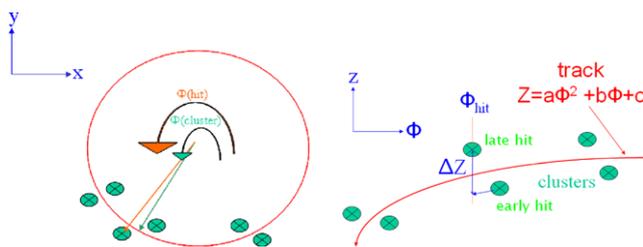


Fig. 29 A fit to the distribution of the difference of projected r coordinates for two-hit clusters



(a) A view of a track segment in the (x, y) plane. (b) A view of a track segment in the z - ϕ plane. See text for details

Fig. 30 A diagram of the technique for measuring the intrinsic z coordinate resolution

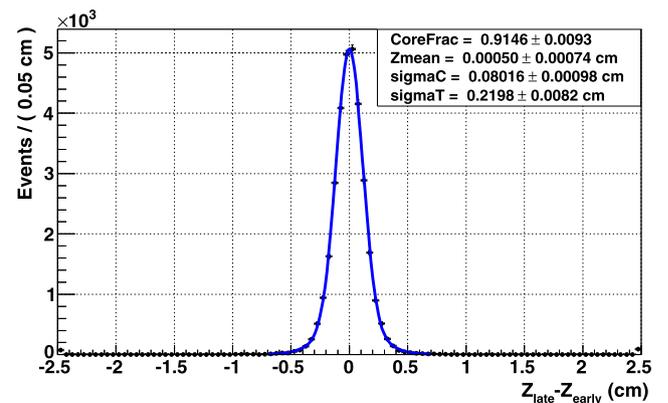


Fig. 31 A fit to the distribution of the difference of projected z coordinates for two-hit clusters

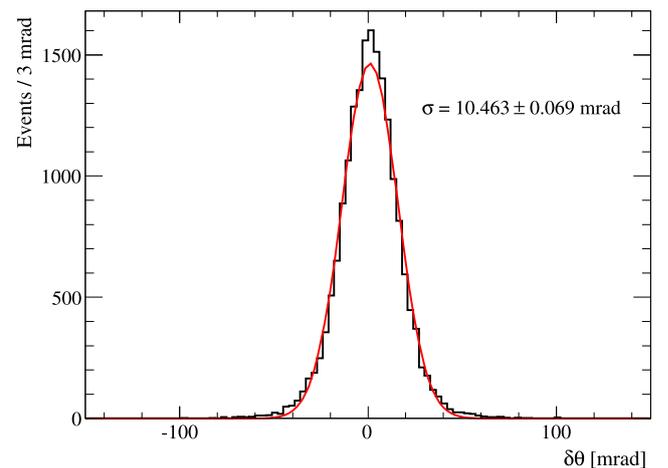


Fig. 32 A fit to the distribution of $\delta\theta_e \equiv \theta_e^{\text{1st turn}} - \theta_e^{\text{2nd turn}}$ on double-turn events. The distribution is fitted with a single Gaussian convolved with itself, and the corresponding width is shown

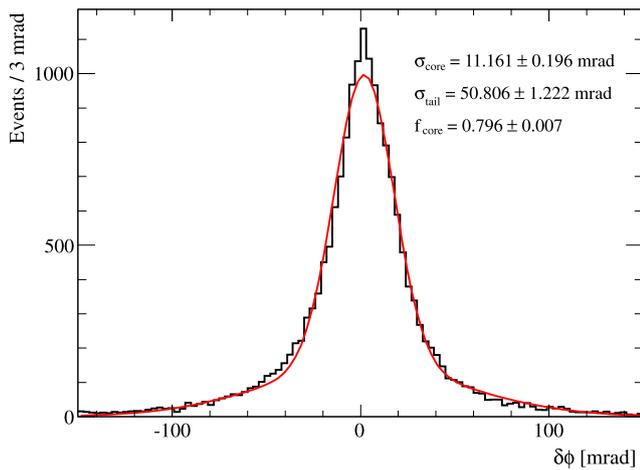


Fig. 33 A fit to the distribution of $\delta\phi_e \equiv \phi_e^{\text{1st turn}} - \phi_e^{\text{2nd turn}}$ on double-turn events. The distribution is fitted with a double Gaussian function convolved with itself, and the corresponding core and tail widths are shown, along with the fraction of events in the core component

recting for this, we obtain a single Gaussian resolution $\sigma_{\theta_e} = 9.4 \pm 0.5$ mrad and a double Gaussian ϕ_e resolution of $\sigma_{\phi_e} = 8.4 \pm 1.4$ mrad in the core (80 %) and $\sigma_{\phi_e} = 38 \pm 6$ mrad in the tail, where the errors are dominated by the systematic uncertainty of the correction. The Monte Carlo resolutions are $\sigma_{\theta_e} \sim 9$ mrad $\sigma_{\phi_e} \sim 8$ mrad, while the design resolutions were $\sigma_{\theta_e, \phi_e} \sim 5$ mrad.

It is also important to stress that these resolutions are affected by correlations among the other positron observables, which can be treated on an event-by-event basis, so that the effective resolutions determining the experimental sensitivity are $\sigma_{\theta_e} = 8.5 \pm 0.5$ mrad and $\sigma_{\phi_e} = 7.7 \pm 1.4$ mrad in the core.

The multiple scattering contribution to these resolution is $\sigma_{\theta_e, \phi_e} \sim 6.0$ mrad, the rest is due to the single hit resolution.

4.10.3 Vertex resolution

The resolution in the position of the decay vertex on the target is dominated by the positron angular resolution. For a proper evaluation of the angular resolution function, a precise knowledge of the correlations between positron angle error and vertex position error is required. The average vertex position resolutions, however, can be measured directly by comparing the projected point of interception at the target plane on double-turn events. The difference in vertex z coordinates of the two turns is fit to the convolution of a single Gaussian function with itself, while that of the vertex y coordinates is fit to the convolution of a double Gaussian function with itself. After the Monte Carlo corrections are applied, the resolutions are $\sigma_{y_e} = 1.1 \pm 0.1$ mm in the core (86.7 %), $\sigma_{y_e} = 5.3 \pm 3.0$ mm in the tail and $\sigma_{z_e} = 2.5 \pm 1.0$ mm. The resolutions for Monte Carlo events are $\sigma_{y_e}^{\text{MC}} = 1.0 \pm 0.1$ mm in the core and $\sigma_{z_e}^{\text{MC}} = 2.9 \pm 0.3$ mm.

The values of σ_{y_e} are corrected for the correlation with the positron energy assumed to be the signal energy.

The design resolution was $\sigma_{y_e, z_e} \sim 1.0$ mm without correcting for correlation.

4.10.4 Energy resolution

The positron energy resolution is measured with a fit of the energy distribution to the unpolarised Michel spectrum multiplied by an acceptance function and convolved with a resolution function:

$$\text{Probability density } (E_e^{\text{measured}}) = (\text{Michel} * \text{Acceptance})(E_e^{\text{true}}) \otimes \text{Resolution.} \quad (6)$$

Functional forms for both the acceptance and the resolution functions are based on the guidance provided by Monte Carlo simulation. The acceptance function is assumed to be:

$$\text{Acceptance}(E_e^{\text{true}}) = \frac{1 + \text{erf}\left(\frac{E_e^{\text{true}} - \mu_{\text{acc}}}{\sqrt{2}\sigma_{\text{acc}}}\right)}{2}, \quad (7)$$

and the resolution function is taken to be a double Gaussian. The acceptance and the resolution parameters are extracted from the fit, as shown in Fig. 34. This gives an average resolution of $\sigma_{E_e} = 330 \pm 16$ keV in the core (82 %) and

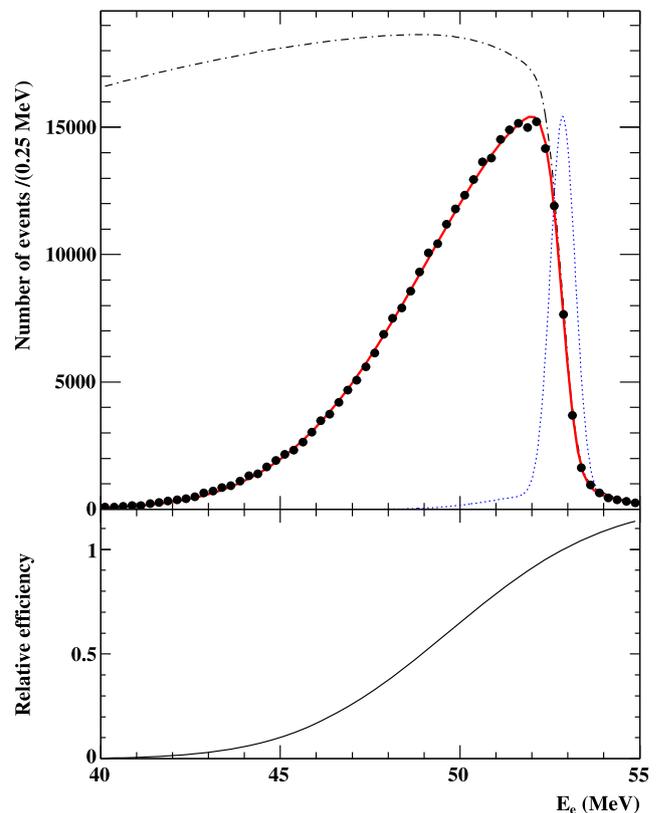


Fig. 34 A fit to the Michel positron energy spectrum. The theoretical spectrum (dashed black), the resolution function (dashed blue) and the acceptance curve (in the bottom plot) are also shown

$\sigma_{E_e} = 1.13 \pm 0.12$ MeV in the tail. There is also a 60 keV systematic underestimation of the energy, to which we associate a conservative 25 keV systematic uncertainty from Monte Carlo studies. This is to be compared with the resolution goal of $\sigma_{E_e} = 180$ keV (0.8 % FWHM).

A complementary approach to determining the positron energy resolution is possible by using two-turn events as for the angular resolution. Figure 35 shows the distribution of the energy difference between the two turns. This is fit to the convolution of a double Gaussian function with itself, the same shape assumed in the fit of the edge of the Michel spectrum. A disadvantage of this technique is its inability to detect a global shift in the positron energy scale. This technique gives an average resolution of $\sigma_{E_e} = 330$ keV in the core (79 %) and $\sigma_{E_e} = 1.56$ MeV in the tail, in reasonable agreement with the results obtained from the fit of the Michel spectrum. A systematic offset of 108 keV between the energies of the two turns also appears; the energy of the first turn is systematically larger than the energy of the second turn. A related effect is the dependence of the measured Michel edge on θ_e . These effects point to errors in the magnetic field mapping.

4.10.5 Chamber detection efficiency

The relative efficiency of each chamber plane is measured as the probability to have a reconstructed hit when its neighbouring plane in the same chamber has at least one hit associated to a track. This probability is called the “hardware” efficiency, while the probability to have a hit associated to the track is called the “software” efficiency. Low-efficiency planes can be traced to operation at voltages below the nominal voltage. For the other planes the average hardware efficiency is ~ 95 % and the average software efficiency is ~ 90 %.

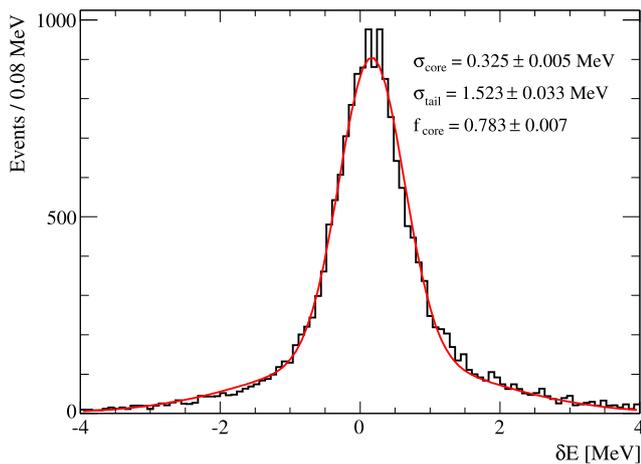


Fig. 35 A fit to the distribution of $\delta E_e \equiv E_e^{\text{1st turn}} - E_e^{\text{2nd turn}}$ on double-turn events. The distribution is fitted with a double Gaussian function convolved with itself, and the corresponding core and tail widths are shown, along with the fraction of events in the core component

4.10.6 DCH efficiency

The absolute efficiency of the DCH system for signal and Michel positrons is difficult to measure. Yet it is not used in the physics analysis, because the $\mu^+ \rightarrow e^+ \gamma$ Branching Ratio is normalised with respect to the number of reconstructed Michel or radiative muon decays, so that the absolute positron reconstruction efficiency cancels. Nevertheless in Fig. 34, the fit to the Michel positron spectrum relative to tracks associated to a TC hit returns an estimate of the relative efficiency versus the positron momentum. In particular the relative efficiency for Michel ($48 \text{ MeV} < E_e$) versus signal positrons is ~ 0.70 . This relative efficiency can also be estimated by Monte Carlo simulation to be ~ 0.75 in good agreement.

The absolute efficiency of DCH only for Michel positron ($48 \text{ MeV} < E_e$) as estimated from Monte Carlo is $\epsilon_{e,\text{DCH}}^{\text{MC,Michel}} \sim 90$ %, while for signal $\epsilon_{e,\text{DCH}}^{\text{MC}} \sim 83$ %. The largest source of inefficiency for the DCH alone are the tracks emitted at $\theta_e \sim 90^\circ$, where multiple turns generate inefficiency in the track-finding algorithm.

5 Timing Counter

The Timing Counter (TC) is dedicated to precisely measuring the impact time of positrons to infer their emission time at the decay vertex in the target by correcting for the track length obtained from the DCH information.

The main requirements of the TC are:

- fast response to be used in the online trigger algorithms and to avoid rate effects;
- online fast and approximate (~ 5 cm) impact point position reconstruction for trigger purposes;
- capability of determining the positron time with an accuracy ~ 50 ps;
- good (~ 1 cm) impact point position reconstruction in the offline event analysis;
- reliable operation in a harsh environment: high and non-uniform magnetic field, helium atmosphere, possibility of ageing effects;
- cover the full acceptance region for signal while matching the tight mechanical constraints dictated by the DCH and COBRA.

These goals were achieved through extensive laboratory and beam tests [25–28].

5.1 Concept and design of the Timing Counter

As mentioned in Sect. 1 and visible in Fig. 1, the TC matches the signal kinematics and is compatible with the mechanical constraints by the placement of one module (sector) upstream and the other downstream.

Each sector is barrel-shaped and consists of two sub-detectors, mounted on top of each other, with full angular coverage for positrons from $\mu^+ \rightarrow e^+ \gamma$ decays when the γ -ray points to the LXe detector (see Fig. 36 for a sketch of a sector).

At the outer radius, the longitudinal detector, consisting of scintillating bars read out at each end by PMTs, is dedicated mainly to time measurements and, thanks to its segmentation along ϕ , also provides a measurement of the positron impact ϕ coordinate. It also provides a measurement of the positron impact z coordinate by exploiting the separate time measurements of both PMTs. An approximate estimate of these variables is obtained online with fast algorithms for trigger purposes. This detector is therefore sufficient to reconstruct all positron variables required to match a DCH track and recover the muon decay time by extrapolating the measured time back to the target.

At the inner radius, the transverse detector, consisting of scintillating fibres, is devoted to determining with high precision the impact z coordinate to improve the matching between the DCH track and TC point.

Each sector is surrounded by a bag made of a 0.5 mm-thick EVAL[®] foil that separates the helium-filled volume of the magnet bore from the PMTs, which are sensitive to helium leakage through the borosilicate window [29, 30]. A buildup of helium inside the PMTs would lead to internal discharge making them unusable. The bag volume is continuously flushed with nitrogen to guarantee a stable operating condition.

5.2 The longitudinal detector

Each sector of the longitudinal detector consists of an array of 15 scintillating bars (Bicron BC404 [25, 31]) with approximate square section and size $4.0 \times 4.0 \times 79.6 \text{ cm}^3$; each bar is read out by a couple of fine-mesh 2" PMTs for high

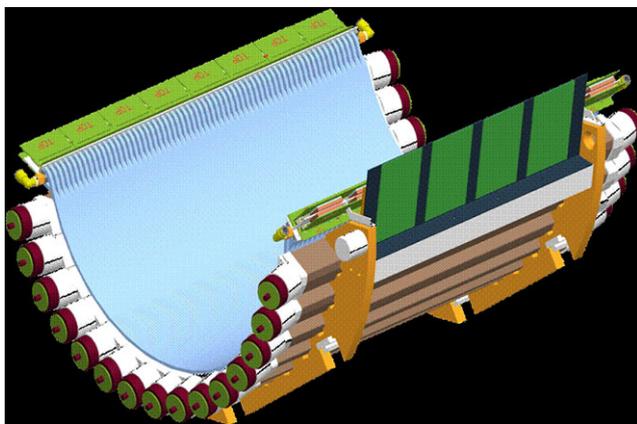


Fig. 36 Schematic picture of a TC sector. Scintillator bars are read out by PMTs. Scintillation fibres are placed on top of the scintillator bars and read out by APDs

magnetic fields [32] glued at the ends. The bars are arranged in a barrel-like shape to fit the COBRA inner profile with a 10.5° gap between adjacent bars matching the DCH pitch. The scintillator type was chosen from among fast scintillator candidates through a series of beam tests by comparing their timing performances. The PMTs were chosen to match the bar size taking into account the magnetic field strength and orientation [33] as well as the mechanical constraints.

The number of bars is defined by requirements on the size of the trigger regions defined in the (ϕ, θ) plane mapped onto the TC; the ϕ width corresponds to the gap between adjacent bars.

The bar thickness was chosen to acquire sufficient energy deposition from positrons to obtain the design time resolution.

The bar width is dictated by the requirement of providing full efficiency for positrons crossing the TC, leaving no gap between bars. Full acceptance of positrons emitted in the experiment's angular range defines the bar length.

Furthermore, the details of the bar geometry are also dictated by the mechanical constraints coming from the TC transverse detector and the DCH system, resulting in the bar geometry in Fig. 37.

5.2.1 Readout electronics

The signals from the PMTs are interfaced with the DRS chip fast digitiser (see Sect. 9.2.1) with the read-out scheme depicted in Fig. 38. Each PMT signal is passively split into three channels with 80 %, 10 % and 10 % amplitude fractions, respectively. The highest fraction is sent to a Double Threshold Discriminator (DTD) specifically designed for time analysis purposes.

DTD is a high-bandwidth, low-noise discriminator with two different tunable thresholds: the lower one allows pre-

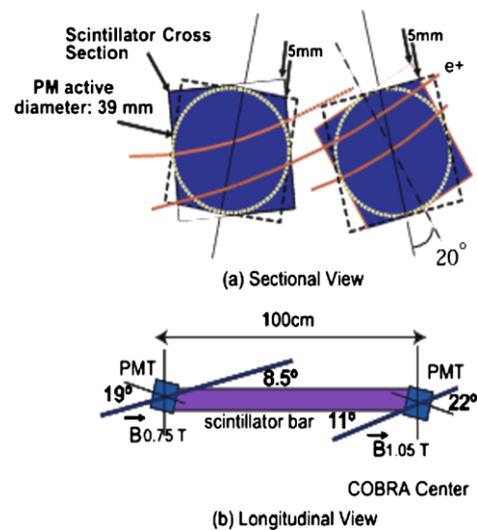


Fig. 37 The layout of a TC bar-PMT assembly

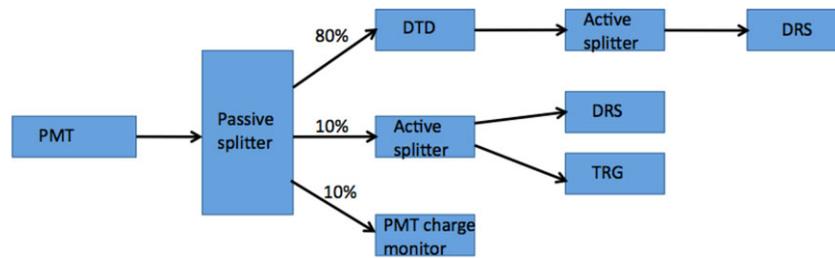


Fig. 38 Schematic picture of the read-out electronics of the TC longitudinal detector. Signals from each PMT are passively split and fed into dedicated electronics channels. Different splitting ratios are used

cise pulse-time measurements (see Sect. 5.4.2) while the higher one acts as a veto against low energy hits. The value of the high threshold determines the detector efficiency and is also related to the trigger threshold as discussed in Sect. 5.4. When the DTD is fired, a standard NIM waveform is generated and then digitised by the DRS boards together with a copy of the PMT waveform. The NIM and PMT waveforms are analysed offline to extract the hit time. Several algorithms can be applied, the best one (used as default) fits the NIM waveform. This method reduces the contribution from the DRS jitter and dynamic range to the overall time resolution.

The first 10 % copy of each signal is duplicated and fed to the trigger system and the DRS: in this stage an active splitter described in Sect. 9.2.2 is used for the level translation needed to match the input ranges of the trigger (see Sect. 8) and DRS boards. The second 10 % signal copy is sent to a charge integrator to monitor the PMT ageing. Since 2011 we implemented a charge monitor based on pulse analysis. Consequently we feed 20 % of the PMT signals to the DRS/trigger inputs.

5.3 The transverse detector

The transverse detector consists of 256 scintillating multi-clad fibres [34] coupled to Avalanche PhotoDiodes (APD) [35] covering the inner surface of a TC sector.

The small sectional area of the fibres ($5 \times 5 \text{ mm}^2$ plus $2 \times 0.5 \text{ mm}$ wrapping thickness totalling 6.0 mm in size along z) satisfies the required space resolution and matches the $5 \times 5 \text{ mm}^2$ APD sensitive area. In order to comply with the mechanical constraints, the fibre ends at the APD side are bent in two different ways: one set has almost straight termination while the other is “S” shaped with small curvature radii ($\sim 2 \text{ cm}$) and shows a significant light leakage, recovered with a high-reflectance wrapping [36].

The APD advantages can be summarised as: small size, insensitivity to magnetic field, and fast response. The main disadvantage is the low gain, at most $\sim 10^3$ if biased near the breakdown voltage, and the high capacitance that make

to optimise time and pulse reconstruction. The PMT charge monitor was deactivated in 2011 and its fraction is added to DRS/trigger input

them sensitive to noise. For this reason a single APD current pulse is read out as voltage across a load resistor and amplified. Bundles of 8 APDs are mounted on distinct boards carrying amplifying stages, power supplies, HV bias and ancillary control signals. Due to mechanical constraints, the fibre bundle mounted on a single front-end board is interleaved with the fibre bundle connected to the adjacent board: the resulting assembly, 16 consecutive fibres read out by two adjacent boards, covers $9.6 \times 40.0 \text{ cm}^2$ and is viewed in the trigger system as a single block.

Each APD has an independent discriminated output, acquired via a FPGA based board and stored. Therefore the transverse detector delivers the list of fired fibres for each event.

The reconstructed coordinate of the hit fibre z_{fibre} can be matched to the z_{bar} measured by the longitudinal detector. The distribution of $z_{\text{bar}} - z_{\text{fibre}}$ for the downstream sector obtained with cosmic rays is shown in Fig. 39. Some tails are related to a $\mathcal{O}(10 \%)$ inefficiency of the APD detector due to its geometry: incomplete coverage and 1 mm gap between fibres.

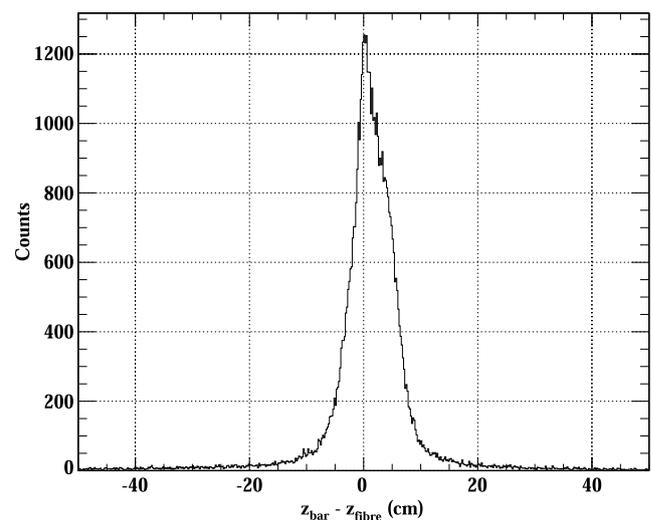


Fig. 39 Distribution of $z_{\text{bar}} - z_{\text{fibre}}$ for the downstream sector

5.4 Commissioning and in-situ performance

The integration of the TC in MEG requires a set of calibrations to determine the detector parameters optimising its performance. Hereinafter only those involving the TC alone are described. Additional calibrations, involving other sub-detectors (e.g. time alignments with the DCH and LXe detectors) are described in Sect. 7.3.1 and in [37].

5.4.1 Gain equalisation

Due to the non-solenoidal configuration of the magnetic field, the inner and outer PMTs are subject to different field intensities and orientations and therefore operate in different working conditions resulting in different gains. Gains need to be equalised for their use in triggering and analysis.

The PMT gain equalisation procedure exploits cosmic rays hitting the TC: due to the uniformity of cosmic ray hit distribution along the bars the charge and amplitude spectra of the inner and outer PMTs are only marginally affected by geometrical effects. It is therefore possible to tune the PMT gains acting on the bias high voltages and equalise the peaks of the Landau distributions (Fig. 40) within 15 % or better. In the trigger algorithm a further software equalisation is performed online to keep a uniform threshold on the whole detector. By means of dedicated weights for each PMT a percent level uniformity is obtained.

5.4.2 Time Walk effect

The Time Walk effect is the dependence on the pulse amplitude of the threshold-crossing time delay as shown in Fig. 41. The analog waveforms sampled at 1.6 GHz are averaged and interpolated separately for each PMT over many events to obtain a template waveform, used to evaluate the correction required in time reconstruction. In the leading

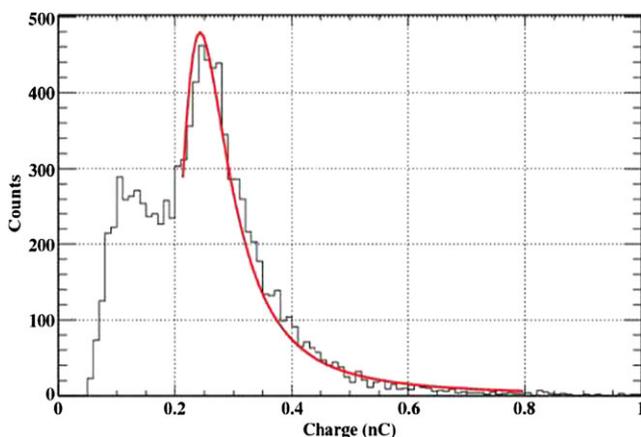


Fig. 40 Example of charge spectrum acquired with a TC bar, superimposed with a Landau function fit

edge region of each template waveform, the dependence of the delay time versus the pulse amplitude is fitted with the empirical function:

$$TW(x) = A + B\sqrt{x} + C \log x, \tag{8}$$

where x is the ratio of the DTD Low Level Threshold (LLT) value to the PMT pulse amplitude and A, B, C are the parameters fit separately for each PMT. These parameters are recalculated every year to account for possible change in the PMT pulse shape.

The correction in Eq. (8) is applied on an event-by-event basis. As shown in Fig. 42, this function reproduces well the experimental data.

The optimal LLT value balancing Time Walk immunity and noise rejection was found by evaluating the timing resolution of each bar with dedicated tests (see Sect. 5.6).

Figure 43 shows the comparison between the time resolutions obtained with two different LLT values (10 mV and 25 mV) on the double bar sample (see Sect. 5.6). The difference between the time resolutions for the two thresholds

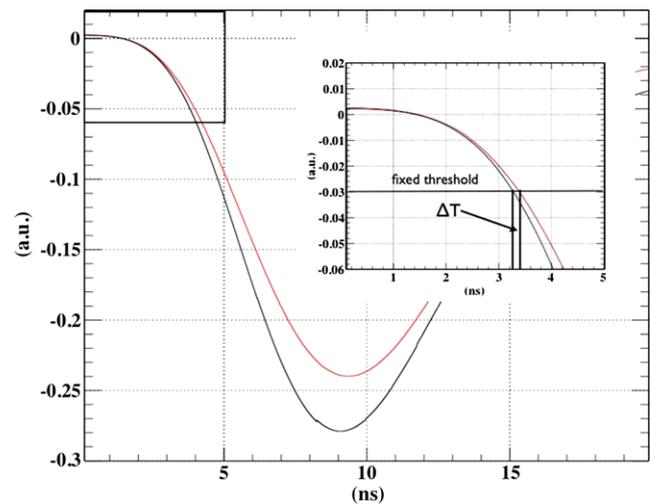


Fig. 41 Graphical representation of the Time Walk effect

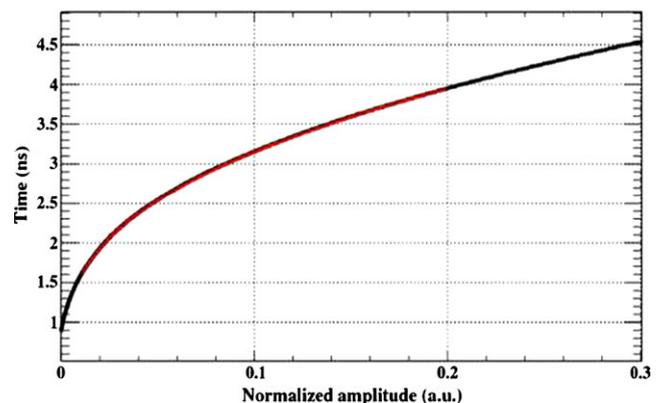


Fig. 42 Time delay versus amplitude relation with fit superimposed

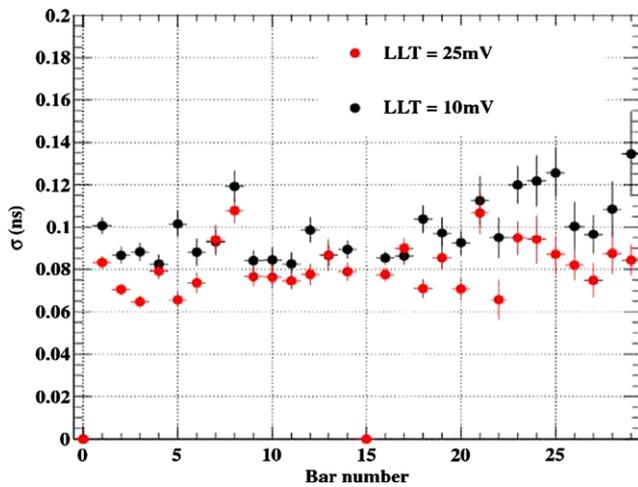


Fig. 43 Timing resolutions for different values of low-level threshold for double bar events

systematically favours 25 mV. From systematic studies of the double bar time resolution for LLT values ranging from 5 mV to 35 mV in 5 mV step, the optimal value of LLT was found to be 25 mV.

5.4.3 Time offsets

Inter-bar time offsets due to the electronic chains are evaluated using cosmic rays. All PMTs are equipped with 10 m of low-loss signal cable to preserve the leading edge of the pulses [38]; the equal length for all PMTs minimises the possibility of time offsets drifting during the data-taking period (several months) and, together with a continuous monitoring of the relevant variables, ensures stable operation. For each bar, the distribution of the time difference between the inner and outer PMTs is acquired. Due to the uniformity and isotropy of cosmic rays this distribution is expected to be flat and centred at zero; the mean values of these distributions are direct measurements of the relative offsets between PMTs on each scintillating bar. The time offsets between different bars and between the TC and the LXe detector are measured using positrons from Michel decays as well as different calibration approaches [39].

Figure 44 shows the effect of accounting for time offsets on a two-dimensional (z versus ϕ) cosmic ray hit map. Using the measured time offsets (lower panel), the z reconstruction algorithm returns the correct position.

5.5 Trigger settings

The TC plays a crucial role in trigger (see Sect. 8.3): it detects the positron and provides a preliminary estimation of its timing and direction.

The trigger position resolution is a key parameter for the direction match with the γ -ray entering the LXe detector

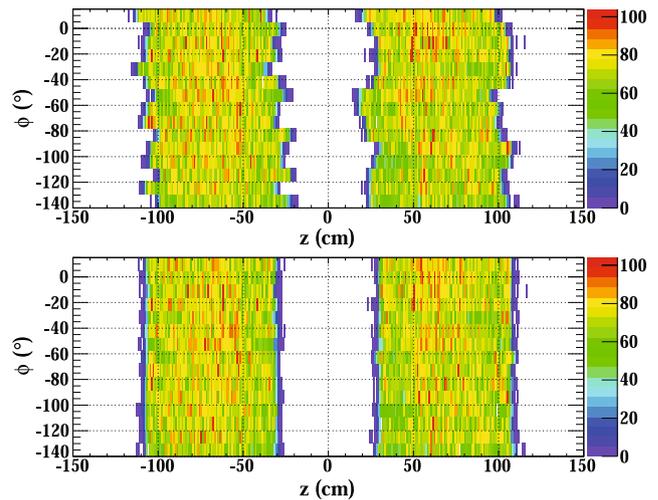


Fig. 44 Cosmic ray hit maps before and after bar time-offset subtractions

and is evaluated by comparing the trigger-level reconstruction of the impact point from Eq. (17) with the one obtained from Eq. (13) in Sect. 5.8. It was $\sigma_z \sim 7.3$ cm in 2009 and has improved to $\sigma_z \sim 5.0$ cm since 2010 after refining the algorithms.

From the point of view of efficiency at the analysis level, the TC trigger and DTD conditional efficiency studies are motivated by the need for all triggered events to fire the DTD to allow event reconstruction. The relevant parameter is the threshold efficiency ϵ_{DTD} defined, for each bar, as the ratio of the number of events with NIM signals to the total number of events triggered on the same bar. The high-level threshold (HLT) is set slightly lower than the trigger one, in order to acquire only events with NIM pulse information. Figure 45 presents the efficiency ϵ_{DTD} versus the HLT. Averaging over all bars yields $\epsilon_{\text{DTD}} \geq 99.9\%$ for $\text{HLT} \leq 400$ mV.

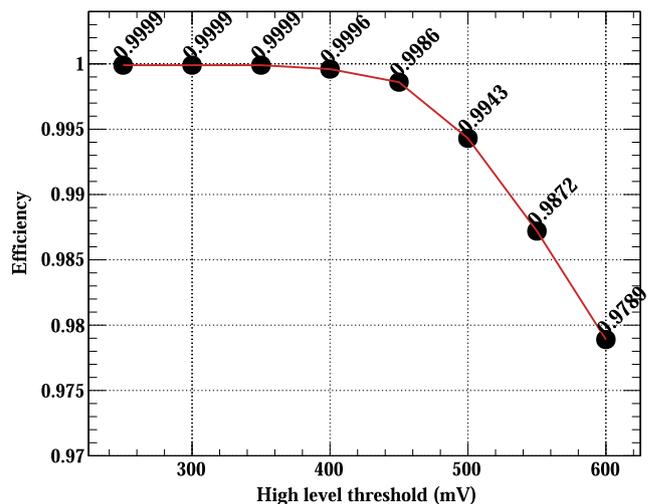


Fig. 45 DTD efficiency vs. high threshold value. The efficiency plateau is reached for $\text{HT} \leq 400$ mV

5.6 Timing resolution of TC

If a positron impinges on a TC bar at time T_{TC} , the times measured by the inner (*in*) and outer (*out*) PMTs read as:

$$t_{in} = T_{TC} + b_{in} + TW_{in} + \frac{\frac{L}{2} + z}{v_{eff}},$$

$$t_{out} = T_{TC} + b_{out} + TW_{out} + \frac{\frac{L}{2} - z}{v_{eff}},$$
(9)

where $b_{in,out}$ are time offsets depending on the read-out chain, $TW_{in,out}$ are contributions from Time Walk effect, v_{eff} is the effective velocity of light in the bar and L is the bar length; the z axis points along the main axis of the bar and its origin is in the middle of the bar. From Eq. (9) the impact time is:

$$T_{TC} = \frac{t_{in} + t_{out}}{2} - \frac{b_{in} + b_{out}}{2} - \frac{TW_{in} + TW_{out}}{2} - \frac{L}{2v_{eff}}.$$
(10)

A more detailed discussion on the determination of the factors b is in Sect. 5.4.3 and of the factors TW in Sect. 5.4.2.

The timing resolution is determined from hits on multiple bars. On a sample of two consecutive bars hit by a positron, the time difference $\Delta T = T_2 - T_1$ is studied. For each bar the impact time is defined as in Eq. (10). The time resolution can be written as

$$\sigma_t = \sqrt{(\sigma_1^2 + \sigma_2^2) + \sigma_{track}^2},$$
(11)

where $\sigma_{1,2}$ are the single bar time resolutions and σ_{track} the contribution from the track length spread. Under the hypothesis that the two bars have similar resolutions $\sigma_1 \simeq \sigma_2$ and neglecting the term due to track propagation, the single bar resolution is estimated from $\sigma_t/\sqrt{2}$. This method overestimates the timing resolution since the track length spread term is not corrected for. Using triple bar events and evaluating the quantity:

$$\Delta T = T_2 - \frac{T_1 + T_3}{2}$$
(12)

the effect of the different path length between bars can be removed at first order, resulting in a better estimate of timing resolution. This is shown in Fig. 46 where the red and black markers represent respectively the resolutions obtained in triple and double bar samples. However, the triple sample has significantly smaller statistics so the double bar sample is used as the reference tool for checking the detector performance.

A certain degradation of the timing performance between the Beam Test ($\sigma_t \sim 40\text{--}50$ ps) and the Physics Run configuration ($\sigma_t \sim 65$ ps) has been observed. A few factors contribute to this effect: the need for a lower PMT gain in order to withstand the high rate and match the dynamic range of the DAQ/electronics chain; a slightly higher value for the

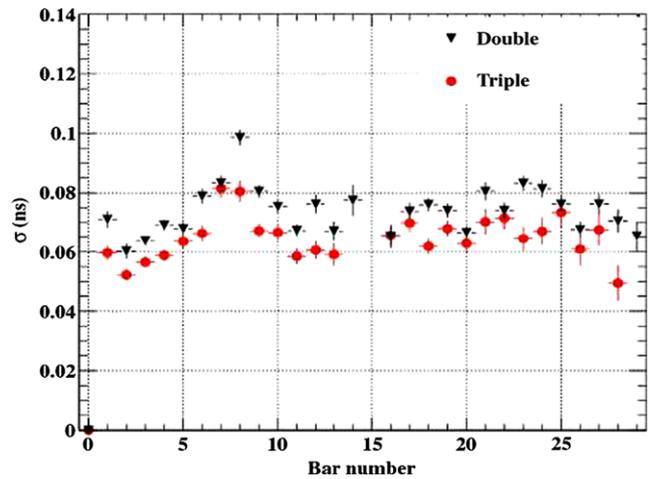


Fig. 46 Timing resolution on double (black markers) and triple (red markers) bar events, with LLT = 25 mV. Due to the different rates for two-bar and three-bar events, the sample used here is different than the one in Fig. 43 (Color figure online)

low threshold due to additional noise from the surrounding environment; intrinsic uncertainties in both double and triple bar estimates; and contributions from DRS calibration.

5.7 Positron timing resolution

In the proposal the timing resolution of the positron was assumed to be dominated by the timing resolution of the TC. That was because of the optimistic assumption on tracking performances, as well as the neglect of the contributions of the calibration of the DRS boards and of TC counter offsets. Assuming a small contribution from the uncertainty on the track length, the design positron timing resolution was to be $\sigma_{te} = 50$ ps.

In Monte Carlo, which incorporates a fairly precise description of DCH and of the positron reconstruction algorithm, an additional contribution due to track length fluctuations $\sigma_{te}^{L,MC} = 50$ ps is present. Added to the TC intrinsic resolution $\sigma_{te}^{TC,MC} = 40$ ps, it totals a positron time resolution $\sigma_{te}^{MC} = 64$ ps.

In the data, those contributions are evaluated to be $\sigma_{te}^{TC} = 60$ ps and $\sigma_{te}^L = 75$ ps. Additionally, an estimated contribution from DRS calibration is $\sigma_{te}^{DRS} = 25$ ps and from TC offset calibration $\sigma_{te}^{cal} = 40$ ps. The total positron resolution is $\sigma_{te} = 107$ ps.

5.8 Position resolution

The positron impact point calculated from the time difference between the two pulses is obtained from Eq. (9) as

$$z = \frac{v_{eff}}{2} \times (t_{in} - t_{out} - (b_{in} - b_{out}) - (TW_{in} - TW_{out})).$$
(13)

The impact point can also be evaluated using the ratio of the charges delivered at the inner and outer PMTs [40]:

$$Q_{in} = EG_{in}e^{-\frac{L}{\Lambda_{eff}} + z}, \tag{14}$$

$$Q_{out} = EG_{out}e^{-\frac{L}{\Lambda_{eff}} - z}, \tag{15}$$

where E is the energy released inside the bar, $G_{in,out}$ takes into account several contributions (i.e. the scintillator yield, PMT quantum efficiency and gain), Λ_{eff} is the effective attenuation length of the bar. Taking the ratio we obtain:

$$\frac{Q_{in}}{Q_{out}} = \frac{G_{in}}{G_{out}}e^{-\frac{2z}{\Lambda_{eff}}}, \tag{16}$$

which leads to:

$$z = \frac{\Lambda_{eff}}{2} \left(\ln \frac{Q_{out}}{Q_{in}} - \ln \frac{G_{out}}{G_{in}} \right). \tag{17}$$

Moreover, from Eq. (14) the energy release in the bar can be estimated without dependence on z :

$$\sqrt{Q_{in}Q_{out}} = E\sqrt{G_{out}G_{in}}e^{-\frac{L}{\Lambda_{eff}}}. \tag{18}$$

Note that the combination of Eq. (17) and Eq. (13) provides a way to evaluate the ratio Λ_{eff}/v_{eff} for each TC bar. Assuming $v_{eff} = 14.0$ cm/ns, the value of Λ_{eff} is extracted from a linear fit as shown in Fig. 47. The Λ_{eff} values extracted in 2011 for all bars fall in the range 40–90 cm. This apparently large dispersion of values is caused by variation in surface reflectivity due to surface machining, hand-made polishing and wrapping, while the bulk attenuation is expected to be constant between the bars.

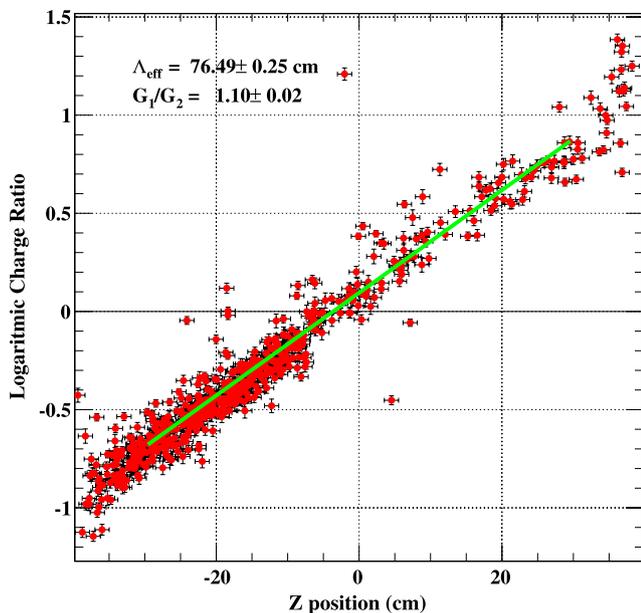


Fig. 47 Logarithm of the PMT charge ratio vs. the z coordinate measured using Eq. (13) for bar number 1. The linear fit returns Λ_{eff} and the ratio of PMT gains

Two methods for impact point reconstruction are used in different stages of the data acquisition chain. The on-line algorithm for z reconstruction in the trigger, requiring fast response and moderate precision, relies on the charge ratio method, while the offline analysis, aiming at the best possible resolution, exploits the PMT time difference. The need for offline calibrations to guarantee the ultimate performance (both for z and time resolutions) is satisfied by using several tools, presented in the following.

5.9 Efficiency

The TC efficiency for Michel and signal positrons emitted from the target can be estimated by Monte Carlo. In the acceptance region the efficiency is $\epsilon_{e,TC}^{MC,Michel} \sim 46\%$ for Michel ($48 \text{ MeV} < E_e$) positrons and $\epsilon_{e,TC}^{MC,Signal} \sim 59\%$ for signal positrons. Such a low efficiency compared to a design value $>90\%$, is due to the DCH support frame that intercepts a large fraction of the positrons exiting the DCH volume; these interacting positrons may lose a large fraction of their energy or change their direction by a large angle due to multiple scattering. Therefore they do not reach the TC and follow a different trajectory from that predicted by the DCH and hence fail the DCH-TC matching cut.

Using the DCH-only trigger (see Sect. 8.3), the TC efficiency versus the reconstructed positron energy was also estimated from data, with results compatible with the Monte Carlo simulation.

The total positron efficiency requires a track in the DCH, a hit in the TC and that the track matches the TC hit. This requirements amount to an efficiency $\epsilon_e^{MC,Michel} \sim 37\%$ for Michel positrons and $\epsilon_e^{MC,Signal} \sim 48\%$ for signal positrons, as estimated by Monte Carlo. A direct measure from data is not possible.

Absolute detection efficiencies are anyway not used in the analysis, as explained in Sect. 4.10.6, therefore the values reported are just indicative.

6 Liquid Xenon detector

6.1 Physical requirements vs. liquid xenon properties

Liquid Xenon (LXe), with its high density (and short radiation length) is an efficient detector medium for γ -rays. Both ionisation charges and scintillation photons are produced by radiation. In MEG, only the scintillation photons are used to simplify the detector construction and to utilise a prompt response of the detector. The general properties of LXe are summarised in Table 3.

The MEG γ -ray detector requires excellent position, time and energy resolutions to minimise the number of accidental coincidences, which are the dominant background process.

Table 3 Liquid xenon properties: W_{ph} is the average energy for scintillation photon, scintillation wavelength is the wavelength where the scintillation spectrum has its maximum

Properties	Value & Unit
Atomic/Mass number	54/131.293
Density at 161.4 K	2.978 g/cm ³
Boiling/Melting point	165.1/161.4 K
Radiation length	2.77 cm
Moliere radius	4.20 cm
Scintillation wavelength	178 nm
W_{ph} for electron/ α	21.6 eV/17.9 eV
Decay time (recombination)	45 ns
Decay time (Fast/Slow component)	4.2 ns/22 ns
Absorption length	> 100 cm

These requirements need high light yield. Fast decay time helps to reduce pile-up events.

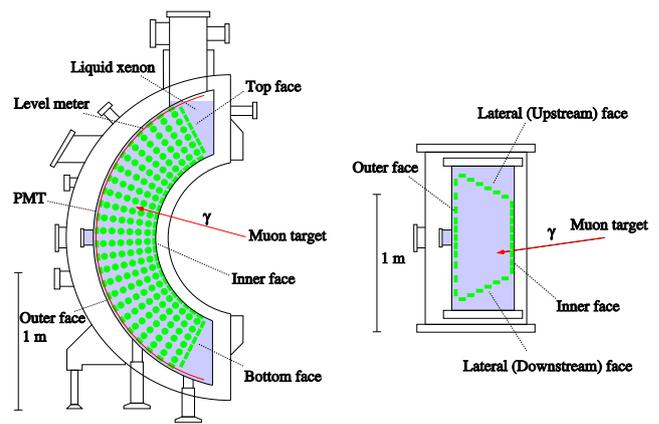
On the other hand, the peak of the LXe scintillation emission spectrum is in Vacuum UltraViolet (VUV), $\lambda \sim 178$ nm, which requires quartz windows in the PMTs, and the operational temperature is ~ 165 K, which requires cryogenic equipment. Furthermore, the scintillation light can be easily absorbed by impurities like H₂O, O₂, and N₂ etc., which must therefore be removed efficiently and continuously.

6.2 Detector design

We developed detector prototypes to perform various feasibility tests. The first used 2.3 l of LXe with 32 PMTs, and the second used 68.6 l of LXe with 228 PMTs [41]. The purpose of the first prototype was a proof-of-principle detection of γ -rays at ~ 2 MeV from radioactive sources in LXe with adequate resolutions. The second prototype was designed to confirm the good resolutions for position, time and energy for γ -rays at ~ 50 MeV and to study LXe detector operation including the purification system toward a full-size detector.

Based on the prototype measurements, the final detector design employing 900 l of LXe was completed in 2004. The construction began in 2005 and was completed in late 2007. A schematic view of the LXe detector is shown in Fig. 48.

The C-shaped structure fits the outer radius of COBRA. The active volume is read out by 846 PMTs submerged directly in LXe to detect the scintillation light, placed on all six faces of the detector (inner, outer, upstream, downstream, top, and bottom), with different PMT densities. The active volume of the detector is ~ 800 l and covers 11 % of the solid angle viewed from the centre of the stopping target. Its depth is 38.5 cm, corresponding to $\sim 14X_0$ and fully containing a shower induced by a 52.83 MeV γ -ray.

**Fig. 48** Schematic view of the LXe detector: from the side (*left*), from the top (*right*)

6.2.1 Cryostat

The LXe detector consists of the inner and the outer vessels. The inner vessel holds 900 l of LXe and the PMT support structure. The outer vacuum vessel is used as a thermal insulation layer. To reduce the material traversed by incident γ -rays, the window of the outer vessel consists of a thin stainless steel plate (0.4 mm thickness), while that of the inner vessel is made of aluminium honeycomb panels covered with carbon-fibre plates to withstand pressure up to 3 bar. The total thickness of the γ -ray entrance window is $0.075X_0$. A picture of the LXe detector is shown in Fig. 49.

In order to monitor the detector condition, 27 temperature, two pressure, and two liquid-level sensors are installed inside the detector. To maintain the LXe temperature and to enable the recondensation of xenon, a custom-designed 200 W pulse-tube cryocooler [42] is installed on the top of the chimney. It does not produce any electrical noise or mechanical vibrations, thanks to the cooling principle. In case more cooling power is needed, e.g. during liquefaction or purification, liquid nitrogen lines are installed to cool down the gaseous volume and the wall of the inner vessel. From the other two chimneys, the cables of PMT high voltages, signals, and sensors etc. are extracted via feedthroughs. A turbomolecular pump is connected to each vessel to evacuate the detector, and a cryopump and a getter pump are installed in the inner vessel to efficiently remove water vapour, which is the most dangerous impurity.

6.2.2 PMTs

In total, 846 2" PMTs [43] are internally mounted on all surfaces and submerged in LXe; the support structure is shown in Fig. 50. In order to allow the operation in LXe, the PMTs are designed to withstand up to 3 bar. The coverage of the inner surface with active photo-cathodes is ~ 35 %. The PMT



Fig. 49 LXe detector under construction. A flange of the outer vacuum vessel was not closed yet when this picture was taken



Fig. 50 View of the LXe detector inside. A total of 846 PMTs are mounted on all surfaces. The face on the *left* side is the incident one for γ -rays from the target

uses a metal-channel dynode structure to achieve a reasonable gain inside a short package (the height is 3.2 cm) and a high tolerance for magnetic fields as large as $\sim 10^{-2}$ T [44], directed either along the transverse or the longitudinal direction. The PMTs are operated at LXe temperature ~ 165 K, and are equipped with a quartz window that transmits VUV photons and a bi-alkali (K–Cs–Sb) photo-cathode sensitive to VUV photons. Aluminium strips are added to the surface of the cathode to reduce the sheet resistance,

which increases at low temperature. Heat dissipation from the base circuit is minimised by optimising the resistor chain ($16\text{ M}\Omega$) with Zener diode protection in the last two dynodes, which keep the voltage constant under a high counting rate. Typical PMTs show Quantum Efficiency (QE) $\sim 15\%$ and gain $\sim 1.8 \times 10^6$ at 850 V. Inside the cryostat, 3.0 m plus 1.6 m RG-196A/U coaxial cables are used for the PMTs signal, and 3.0 m RG-188A/U coaxial and 1.6 m wires for HV. These are connected to the feedthroughs. Outside the cryostat, signal cables are connected to the splitter boards (see Sect. 9.2.2), and divided into trigger (see Sect. 8) and DRS4 boards (see Sect. 9.2.1) which record all the PMT waveforms.

6.2.3 Purification system

As discussed in Sect. 6.1, scintillation light may be absorbed by impurities in LXe, such as H_2O and O_2 at ppm level. We have developed two purification methods, gas-phase [41] and liquid-phase [45], to remove those impurities. Gas-phase purification removes impurities in Gaseous Xenon (GXe) by means of a metal-heated getter. Although the gas-phase purification successfully reduces such impurities, its circulation speed is limited (~ 0.6 l/h). It turned out during the gas-phase purification study that H_2O was the dominant component for the absorption.

Liquid purification was developed to improve the circulation speed and to remove mainly H_2O from LXe by using a cryogenic centrifugal fluid pump [46] and molecular sieves (MS13A). At the normal working point, the flow rate is ~ 70 l/h. Molecular sieves can absorb more than 24 g of water, and the cartridge contains heaters which enable a regeneration of the cartridge.

6.2.4 Storage system

In addition to the detector, two storage systems for LXe and GXe have been developed so that 1000 l of xenon can be stored safely when the detector is not operated [47]. One is a 1000 l Dewar with a pulse-tube cryocooler as well as liquid nitrogen cooling lines and a thermal insulation layer. The heat income is estimated to be less than 20 W, and the cryogenic tank is designed to tolerate a pressure up to 6 bar. As a result, LXe can be stored safely without supplying any cooling power for 100 h. The Dewar is connected to the detector with flexible tubes thermally insulated to allow a rapid LXe transfer.

The other storage system consists of eight high-pressure gas tanks to store GXe, each of which can contain up to 360 kg of xenon corresponding to 120 l of LXe. This can be used in long shutdown periods. In Fig. 51, a schematic view of the LXe detector system including purification and storage systems is shown.

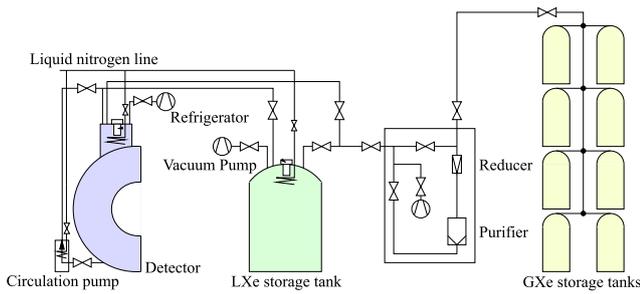


Fig. 51 Schematic view of LXe storage and purification systems as well as of the detector

6.2.5 Detector operation

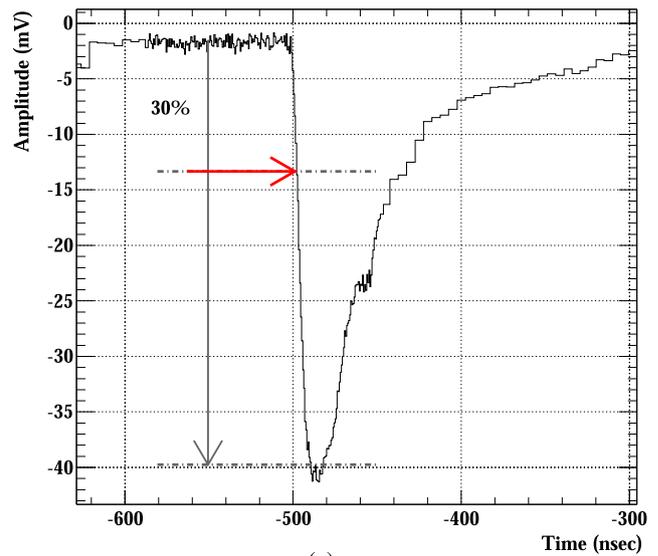
Before starting detector operation, the cryostat is evacuated by means of turbomolecular and getter pumps to minimise impurities inside the detector. Meanwhile, GXe is liquefied into the 1000 l storage tank. LXe can be transferred from the tank to the detector by means of the pressure difference between the pressurised tank and the detector. After the liquid transfer is complete, the pressure of the detector is controlled to be 1.2–1.4 bar by means of the pulse-tube cryocooler. If the cooling power is too much, heaters surrounding the cryocooler are used to reduce the cooling power, and if the cooling power is not sufficient, liquid nitrogen is used to sustain the pressure. In 2010, the pulse tube cryocooler was sufficient to keep 900 l of LXe in the detector, and ~30 W heaters are constantly used to produce additional heat. In case of a problem, or during a short maintenance period, we can transfer the LXe from the detector into the 1000 l tank by means of the centrifugal pump.

6.3 Reconstruction methods

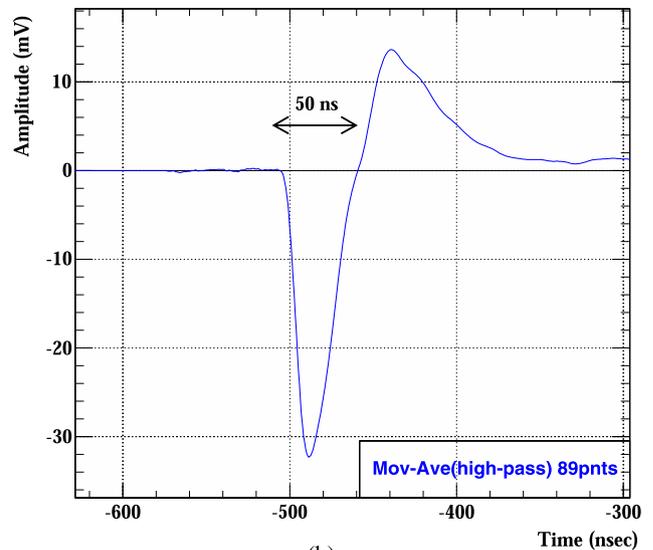
The reconstruction methods and algorithms are described. In the analysis of the LXe detector, a special coordinate system (u, v, w) is used: u coincides with z in the MEG coordinate system; v is directed along the negative ϕ -direction at the radius of the inner face ($r_{in} = 67.85$ cm), which is the direction along the inner face from bottom to top; $w = r - r_{in}$, measures the depth from the inner face.

6.3.1 Waveform analysis

We extract the pulse time and charge from each waveform. A typical waveform is shown in Fig. 52(a). The baseline is estimated by averaging the points in the region before the pulse for each channel on an event-by-event basis. The digital constant fraction method is used to determine the pulse time, defined as the time when the signal reaches a predefined fraction (30 %) of the maximum pulse height. To maximise the signal-to-noise ratio for the determination of the charge, a digital high-pass filter, based on a moving-average



(a)



(b)

Fig. 52 (a) An example of a PMT raw waveform from the waveform digitiser. (b) An example of a waveform with high-pass filter

method, is applied as shown in Fig. 52(b). The charge is estimated by integrating the filtered PMT pulse and is later converted into the number of photoelectrons. The number of average points is 90–100, corresponding to about a 10 MHz cutoff frequency.

6.3.2 Position

The γ -ray conversion point is reconstructed by using the distribution of the light detected by the PMTs near the incident position. The three-dimensional position is determined by fitting the expected light distribution on the PMTs, calculated from the solid angles, to the observed distribution. To minimise the effect of shower fluctuations, the fit is done twice: the first with ~45 PMTs and the second with

~ 15 PMTs around the first fit result. The remaining bias on the reconstructed position along u and w coordinates is corrected using results from the Monte Carlo simulation. The γ -ray direction is defined by the line segment between its reconstructed conversion point in the LXe detector and the reconstructed positron vertex on the target.

6.3.3 Energy

The reconstruction of the γ -ray energy E_γ is based on the sum of the number of photons detected by the PMTs corrected for their geometrical coverage, which depends on the PMT locations: the coverage on the outer face is 2.6 times less than that on the inner face. This correction works well for almost all events except for very shallow ones ($w_\gamma < 3$ cm), which require special treatment because they are very sensitive to the relative positions of each PMT and the conversion point. In this case, the solid angle subtended by each photo-cathode at the conversion point can be a better variable than the coverage for estimating the collection efficiency.

In Fig. 53, the linearity of the relation between E_γ and the number of detected scintillation photons is shown. The conversion factor from the number of photons to energy in order to define an absolute energy scale is obtained from the 55 MeV γ -ray peak taken by the Charge EXchange (CEX) calibration (see Sect. 7.3.2).

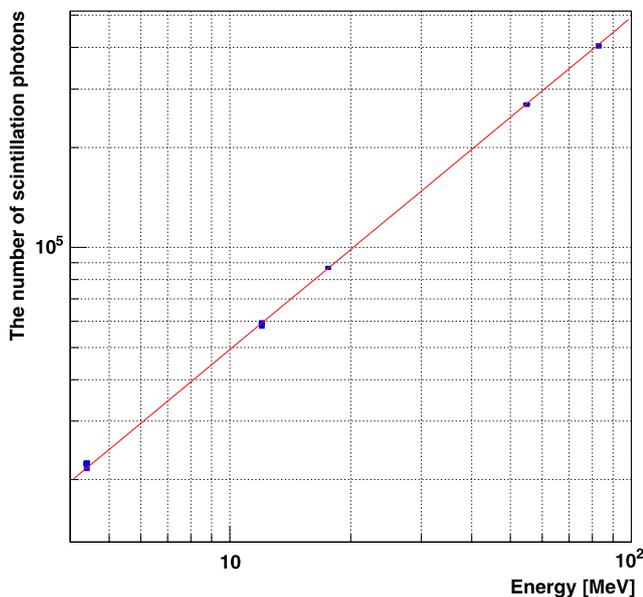


Fig. 53 The number of detected scintillation photons versus the incident γ -ray energy: 4.4, 11.6 and 17.6 MeV γ -rays are measured by C–W calibration (see Sect. 7.3.1), while 55 and 83 MeV γ -rays by the CEX calibration (see Sect. 7.3.2). Error bars assigned to data points are too small to be visible

6.3.4 Timing

Each PMT determines the time of arrival of the scintillation light from the waveform analysis described in Sect. 6.3.1. To relate this time to the γ -ray conversion time, the propagation time of the scintillation light must be subtracted as well as any hardware-induced time offset (e.g. cable length). The scintillation light propagation time is subtracted by exploiting the distance between the reconstructed interaction point and each PMT assuming an effective light velocity ~ 8 cm/ns. After the subtraction of these effects, the γ -ray conversion time is obtained by combining the timings of those PMTs which contain more than 50 photoelectrons and calculating the minimum value of χ^2 defined as a sum of the squared difference between each calculated and reconstructed time. Typically ~ 150 PMTs are used to reconstruct 50 MeV γ -rays. The PMTs with large contributions to χ^2 are rejected during this fitting procedure to remove pile-up effects.

Finally, the reconstructed time is corrected by the time-of-flight between the positron reconstructed vertex on the target and the reconstructed conversion point in the LXe detector.

6.4 Performance

The LXe detector performance is described in the following and summarised in Table 4.

6.4.1 Position resolution

The performance of the position reconstruction is evaluated by a Monte Carlo simulation and is validated in dedicated CEX runs by placing lead collimators at several positions in front of the LXe detector. Figure 54 shows the result of position reconstruction with a 1.8 cm thick lead collimator installed in front of the LXe detector during the CEX run. The lead collimator has two 1 cm wide slits along the u coordinate at $v = 0$ cm and $v = 6$ cm, and this figure shows the projection onto the v coordinate. Events with the energy between $50 \text{ MeV} < E_\gamma < 80 \text{ MeV}$, and with the γ conversion depth $w_\gamma > 2$ cm are selected. Since the thickness of

Table 4 LXe detector performance summary

Energy resolution [$w_\gamma > 3$ cm]	1.6 %
Energy resolution [$0.8 < w_\gamma < 3$ cm]	2.0 %
Energy resolution [$0 < w_\gamma < 0.8$ cm]	2.7 %
Timing resolution	67 ps
Position resolution	$5(u_\gamma, v_\gamma), 6(w_\gamma)$ mm
Detection efficiency	63 %

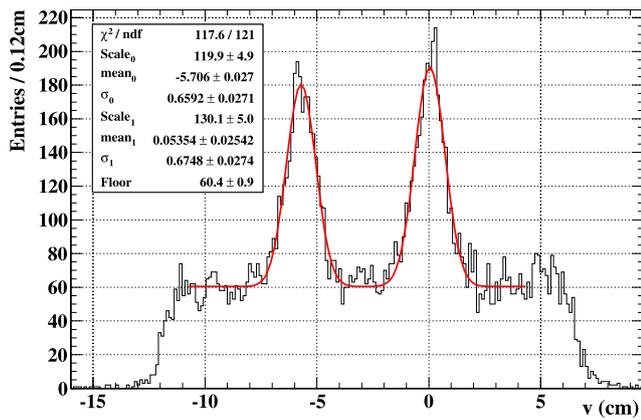


Fig. 54 Reconstructed position distribution with a lead collimator in CEX runs. There are two slits 1 cm wide in the 1.8 cm thick collimator

the lead collimator is not sufficient to stop 55 MeV γ -rays, the floor events which are penetrating the lead collimator are also observed. A double Gaussian function plus a constant term is fitted to extract the position resolution of 54.9 MeV γ -rays, and the results are 6.6 mm, 6.7 mm in this example. The average resolution at different positions is 6.9 mm. This result contains the effect of the slit width itself and of the spread of π^0 decay points, and is to be compared with the average resolution (6.5 mm) from a Monte Carlo simulation of the same configuration. The quadratic difference (1.8 mm, expected to come from the PMT QEs calibration uncertainty) between the data and the Monte Carlo simulation is added into the position resolution map built from simulation.

Taking into account the difference between Monte Carlo simulation and the data, the average position resolutions are estimated to be $\sigma_{(u_\gamma, v_\gamma)} \sim 5$ mm and $\sigma_{w_\gamma} \sim 6$ mm, respectively, comparable with the design position resolutions. They are also close to $\sigma_{(u_\gamma, v_\gamma)}^{\text{MC}} \sim 4$ mm and $\sigma_{w_\gamma}^{\text{MC}} \sim 6$ mm estimated by Monte Carlo simulation.

6.4.2 Energy resolution

The energy response of γ -rays at the signal energy is extracted from the CEX calibration. A small correction is applied to take into account the different background conditions between the muon and the pion beams, and the opening angle between the two γ -rays.

The response function of the detector for monochromatic γ -rays is asymmetric with a low-energy tail due to mainly two reasons. One is the interaction of γ -rays in the material in front of the LXe active volume, and the other is the shower leakage from the front face. Figure 55 shows the LXe detector response to 54.9 MeV γ -rays. The distribution is fitted with an asymmetric function $F(x)$ convolved

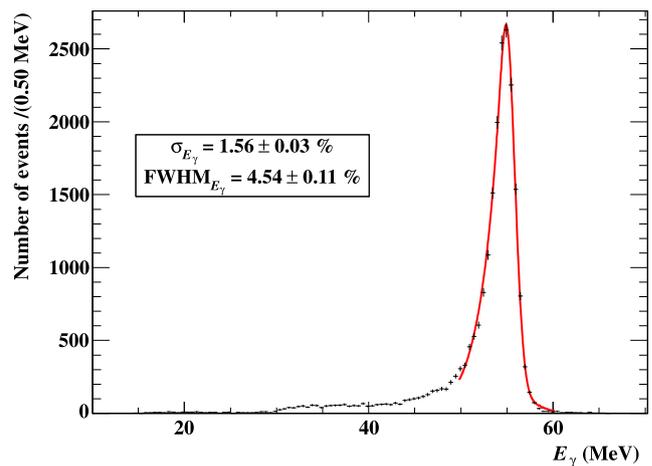


Fig. 55 Energy response of the LXe detector to 54.9 MeV γ -rays for $w_\gamma > 2$ cm in a restricted range of (u_γ, v_γ) . The fitting function is described in the text. The resolution is $\sigma_{E_\gamma} = 1.56$ % and $\text{FWHM}_{E_\gamma} = 4.54$ %

with the pedestal distribution $h(x)$ in the CEX run. $F(x)$ is given by

$$F(x) = \begin{cases} A \exp\left(-\frac{(x-x_0)^2}{2\sigma_{E_\gamma}^2}\right) & (x > x_0 + \tau), \\ A \exp\left(\frac{\tau}{\sigma_{E_\gamma}^2}(\tau/2 - (x - x_0))\right) & (x \leq x_0 + \tau), \end{cases}$$

where A is a scale parameter, x_0 is a peak position parameter, τ is a transition parameter and σ_{E_γ} is a resolution parameter indicating the distribution width on the high-energy side. Since $F(x)$ shows the intrinsic resolution of the detector without pedestal contribution, it can be used for any realistic environment with a different pedestal distribution.

A 3-dimensional mapping of these parameters is incorporated into the likelihood function for the final analysis since they depend on the position of the γ -ray conversion, mainly on its w_γ coordinate. As an example, the average energy resolution is measured to be $\sigma_{E_\gamma} = 1.6$ % ($3 \text{ cm} < w_\gamma$), 2.0 % ($0.8 \text{ cm} < w_\gamma < 3 \text{ cm}$) and 2.7 % ($0 \text{ cm} < w_\gamma < 0.8 \text{ cm}$) in 2011. Except for the acceptance edge along v coordinate ($\sigma_{E_\gamma} \sim 2.5$ % with $|v| > 68.2$ cm and $w_\gamma > 3$ cm), the energy resolution depends weakly on the u_γ and v_γ coordinates. This number is to be compared with the energy resolution of $\sigma_{E_\gamma}^{\text{MC}} = 1.2$ % for ($2 \text{ cm} < w_\gamma$) evaluated by Monte Carlo simulation. The reason of this slightly worse resolution is not fully understood. The behaviour of PMTs such as gain stability, angular dependence etc., or optical properties of LXe such as convection might be possible sources.

The design resolution was $\sigma_{E_\gamma} = 1.7$ % over all w_γ .

6.4.3 Timing resolution

To investigate the intrinsic time resolution of the LXe detector due to photoelectron statistics, two PMT groups (even PMT IDs or odd PMT IDs) are defined, and the times of the

same event are reconstructed by these two groups independently. Then the intrinsic time resolution is estimated by the time difference between these two results, being dominated by photoelectron statistics, while the effects of electronics, the position reconstruction and the event-by-event shower spread cancel out resulting in $\sigma_{t_\gamma}^{\text{phe}} = 36$ ps at 55 MeV.

The absolute timing resolution is evaluated from the time difference between two γ -rays, emitted back-to-back from the π^0 decay during CEX calibration (see Sect. 7.3.2), reconstructed by the LXe detector and by a reference preshower counter. Figure 56 shows the measured time difference distribution having a resolution of $\sigma_{t_{\gamma\gamma}}^{\pi^0} = 119$ ps, which includes not only the LXe detector timing resolution but also the contributions due to the uncertainty of the π^0 decay position and to the reference counter. The former is evaluated to be $\sigma_{t_\gamma}^{\pi^0} = 58$ ps by the π^- beam spread ($\sim 8 \times 8$ mm² beam spot size). The reference counter has two plastic scintillator plates and both sides are read out by 2'' fine mesh PMTs. The timing resolution of each plate is estimated by the time difference between the two PMTs, and the timing resolution of the counter is estimated by these resolutions by taking into account the correlation of the plates. Finally, the resolution of the reference counter is evaluated to be $\sigma_{t_\gamma}^{\text{ref}} = 81$ ps. The absolute timing resolution of the LXe detector is estimated by subtracting these contributions, resulting in $\sigma_{t_\gamma} = 65$ ps at 55 MeV. Figure 57 shows the energy-dependent timing resolution of the LXe detector. Black squares show the measured timing resolution $\sigma_{t_{\gamma\gamma}}^{\pi^0}$, while the red circles show the LXe detector timing resolution σ_{t_γ} . The black and the red smooth curves are the fit results, and their functions are shown in the figure. A vertical dotted line shows the γ -ray signal energy (52.83 MeV). The timing resolution

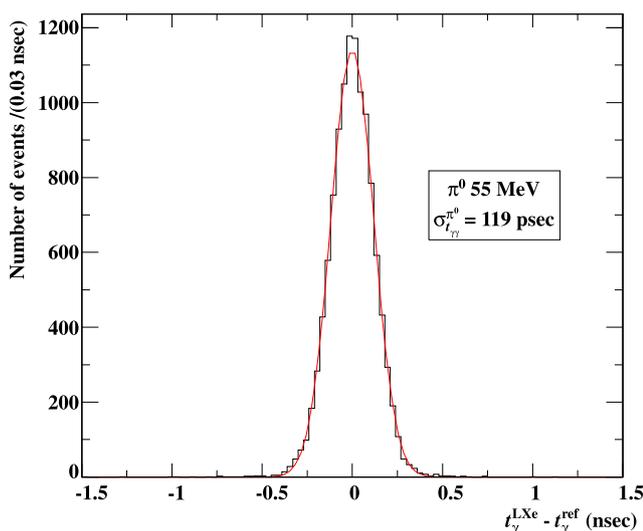


Fig. 56 Time difference reconstructed by the LXe detector and a reference preshower counter for 54.9 MeV γ -rays

improves at higher energy, which indicates that the photoelectron statistics still contributes significantly. This small energy dependence is taken into account in extracting the timing resolution for signal γ -rays to obtain $\sigma_{t_\gamma}^{\text{LXe}} = 67$ ps. This number is in good agreement with the timing resolution of $\sigma_{t_\gamma}^{\text{LXe,MC}} = 69$ ps evaluated by Monte Carlo simulation.

The breakdown of the time resolution is as follows:

$$\sigma_{t_\gamma}^{\text{LXe}} = \sigma_{t_\gamma}^{\text{phe}} \oplus \sigma_{t_\gamma}^{\text{ele}} \oplus \sigma_{t_\gamma}^{\text{TOF}} \oplus \sigma_{t_\gamma}^{\text{sho}}, \quad (19)$$

where $\sigma_{t_\gamma}^{\text{phe}}$ is defined above, $\sigma_{t_\gamma}^{\text{ele}} = 24$ ps from electronics contribution, $\sigma_{t_\gamma}^{\text{TOF}} = 20$ ps from the γ -ray time of flight uncertainty (which corresponds to depth reconstruction uncertainty), and $\sigma_{t_\gamma}^{\text{sho}} = 46$ ps from position reconstruction uncertainty and the shower fluctuation.

The final resolution on the γ -ray timing σ_{t_γ} is obtained combining $\sigma_{t_\gamma}^{\text{LXe}}$ with the additional $\sigma_{t_\gamma}^{\text{tar}}$ spread due to the uncertainty in the muon decay vertex on the target as measured by extrapolating the positron at the target plane (see Sect. 4.10.3). This spread is no more than $\sigma_{t_\gamma}^{\text{tar}} \sim 5$ ps giving a negligible contribution. The results are $\sigma_{t_\gamma} = 67$ ps and $\sigma_{t_\gamma}^{\text{MC}} = 69$ ps.

These results are to be compared with the design resolution $\sigma_{t_\gamma}^{\text{LXe}} \sim 43$ ps, that was calculated taking into account approximately only the contribution from position reconstruction uncertainty. Taking into account also the uncertainty in the muon decay vertex the design resolution is again $\sigma_{t_\gamma} \sim 43$ ps.

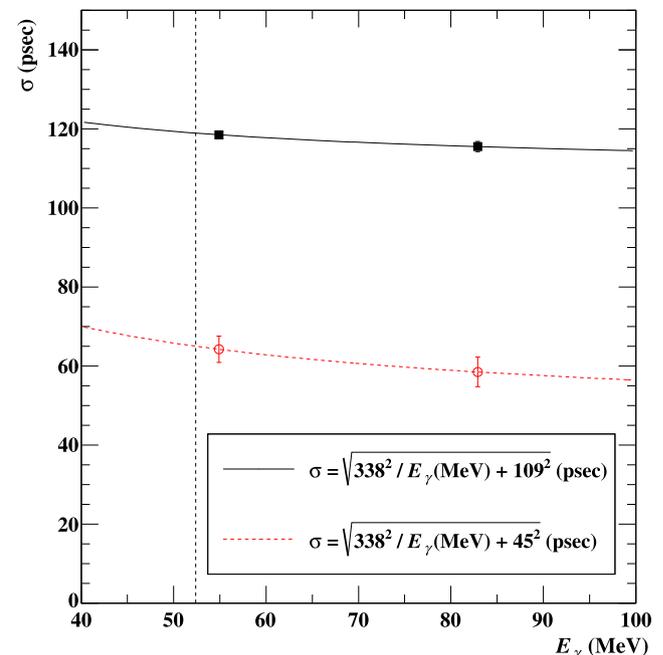


Fig. 57 Energy-dependent timing resolution of the LXe detector (see text for details)

6.4.4 Detector uniformity and energy scale

A spatial non-uniformity of the detector, due mainly to its geometry, is measured by means of the C–W accelerator (see Sect. 7.3.1) and CEX (see Sect. 7.3.2) and corrected to equalise the detector response.

The E_γ scale is constantly monitored by looking at the reconstructed 17.6 MeV energy peak from photons produced by C–W protons on Li (see Sect. 7.3.1). The E_γ scale and its resolution are cross-checked by fitting the γ -ray background spectrum to a theoretical spectrum obtained with a Monte Carlo simulation including radiative muon decay, positron annihilation-in-flight and γ -ray pile-up, folded with the detector resolution (see Sect. 6.4.7). The energy scale and the energy resolution are free parameters in the fit, and we can compare the fit results with the calibrated energy scale and the measured energy resolution in CEX. The systematic uncertainty on the energy scale on the basis of these measurements is $\sim 0.3\%$, in which the dominant source is from the detector uniformity correction ($\sim 0.2\%$).

6.4.5 Rejection of pile-up events

It is important to recognise and unfold pile-up events in a high muon intensity environment in order to avoid a loss of signal efficiency. At $3 \times 10^7 \mu^+/s$ beam rate, around 15 % of triggered events suffer from pile-up. The γ -ray pile-up signature is identified by three main methods, the shape of the waveform sum of all PMTs, the charge distribution of the inner and outer face PMTs, and the χ^2/NDF distribution in the time reconstruction.

The first method identifies multiple γ -ray events with different timing by searching for peaks in the summed waveform of all PMTs. If a pile-up γ -ray is detected by this method, its contribution is subtracted by using a template waveform, and the waveform of the main γ -ray is used for energy estimation.

The second method identifies multiple γ -ray events in different positions by searching the PMT charge distribution on the inner and outer faces for spatially separated peaks. If the event has two peaks, the main peak is retained after removing the secondary γ -ray. A pile-up removal method was developed, to accomplish this efficiently. A position-dependent table containing the average charge of each PMT in presence of 17.6 MeV γ -rays is prepared beforehand. Once a pile-up event is identified, the energy of the event is estimated by fitting the PMT charges to the table without using PMTs around the secondary γ -ray. Then, the PMT charges around the secondary γ -ray are replaced with the charges estimated by the fit. Finally, the energy is reconstructed with the replaced PMT charges as described in Sect. 6.3.3.

In the third method multiple γ -ray events with different timings are identified using the algorithm of the first method.

If an event has multiple γ -rays, the χ^2/NDF in the time reconstruction becomes large. Although the PMTs with large contributions to the χ^2 are discarded when the γ -ray timing is reconstructed (see Sect. 6.3.4), all the PMTs observing more than 50 photoelectrons are used to identify pile-up events effectively. If the pile-up is not identified in the summed waveform or in the light distribution, and is only identified in the timing distribution, the event is discarded. Similarly, if the pile-up is only found in the light distribution but the subtracted energy is negative or larger than 10 % of the total energy, the event is discarded as well. This cut has a 98 % efficiency, estimated by the time sideband data around the signal region.

6.4.6 Rejection of cosmic ray events

Cosmic ray events are rejected using topological cuts because these events mostly enter the detector from the outer face. Therefore the ratio of the charges collected on the inner and the outer faces of the detector is smaller in cosmic ray events than in muon-beam events, which enter the detector from the inner face. The reconstructed depths of cosmic ray events are significantly larger than those of signal γ -rays for the same reason. Figure 58 shows the two-dimensional scatter plot of the charge ratio vs. depth.

In this two-dimensional parameter space, the selection criteria are defined to maximise cosmic ray rejection efficiency while keeping a signal efficiency of 99 %. The cut discards 56 % of cosmic ray events.

The combined analysis efficiency of this and the pile-up cuts is 97 %.

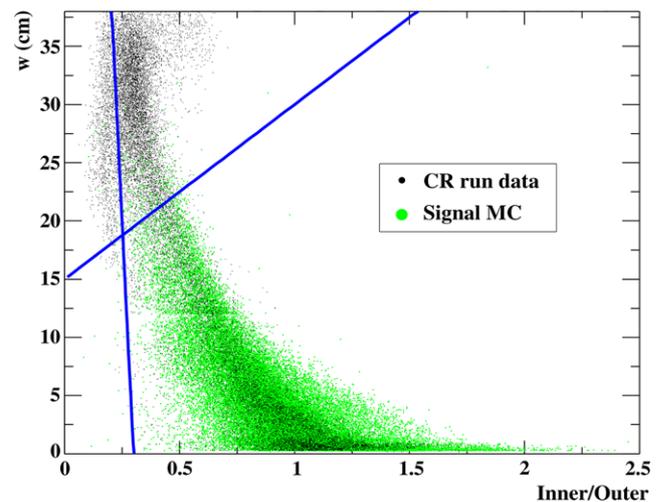


Fig. 58 Two-dimensional plot of the charge ratio collected on the inner/outer faces vs. the γ -ray interaction depth. *Black points* show the data from dedicated cosmic ray runs, and *green* ones the signal Monte Carlo events. The two *lines* (blue in the online version) show the selection criteria for rejection of cosmic ray events (Color figure online)

6.4.7 γ -ray background spectra

In Fig. 59 the γ -ray background data are shown, while the green line shows the fitted spectrum from the Radiative Muon Decay (RMD) plus positron Annihilation-In-Flight (AIF) spectrum, the blue line the γ -ray RMD + AIF spectrum taking into account the effect of pile-up, the black line the cosmic ray spectrum, and the red line the fit result taking into account all the background spectra. In the 51–55 MeV region with $w_\gamma > 2$ cm, the RMD + AIF spectrum accounts for 94–96 % of the total background spectrum, cosmic ray spectrum for less than 1 %, and the remaining components from pile-up, energy response with higher energy tail, etc. for 3–5 %.

6.4.8 Detection efficiency

The fiducial volume of the LXe detector is defined as $|u| < 25$ cm, $|v| < 71$ cm and 0 cm $< w_\gamma < 38.5$ cm. The γ -ray detection efficiency is estimated using the Monte Carlo simulation, and is confirmed by the measurement in CEX runs and by the γ -ray background rate in physics runs. The low-energy tail in Fig. 55 comes from interactions with the material in front of the active volume and shower escape from the inner face. On the path of a γ -ray entering the LXe active volume, there are the drift chamber frame, its support structure, the COBRA magnet, an entrance window of the LXe cryostat, the PMTs and their support structure, etc. The largest source of inefficiency is due to the COBRA magnet, 14.2 %. Within the analysis window 48.0 MeV $< E_\gamma < 58.0$ MeV, the average efficiency in Monte Carlo simulation is found to be $\epsilon_\gamma = 65$ %, taking into account the positron event distribution. This efficiency was measured using NaI single trigger data in the CEX run. By tagging an 83 MeV γ -ray from π^0 decay with the NaI detector, the detection efficiency of the detector for 55 MeV γ -rays is found to be $\epsilon_\gamma = 64$ – 67 %, consistent with the Monte Carlo estimation.

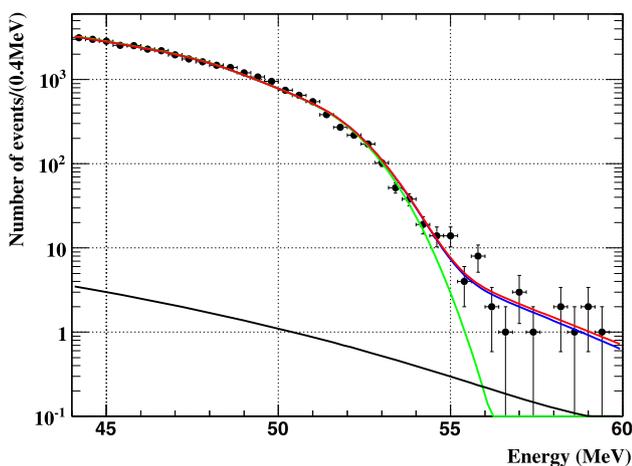


Fig. 59 γ -ray background spectrum during physics run in 2010 (see text for details)

Taking into account the analysis efficiency of 97 % as discussed in Sect. 6.4.5), the combined γ -ray analysis and detection efficiency is $\epsilon_\gamma = 0.65 \times 0.97 = (63 \pm 3)$ %. This value is to be compared with an expected value $\epsilon_\gamma \sim 60$ %.

7 Calibrations

For each kinematic variable it is necessary to know the resolution, supplemented by the absolute scale (for the energies) or the position of the zero (for the relative time and direction) and the degree of stability of these parameters over the experiment lifetime. We refer to the knowledge of a single quantity at a given time as “calibration”, and call the ensemble of these periodically repeated calibrations “monitoring”.

Continuous monitoring of the apparatus is essential for two reasons: firstly for early detection of misbehaviour of the apparatus and secondly the knowledge of the parameters has a direct influence on the systematic uncertainties on physical variables, acceptances and thresholds.

The issues of calibration and monitoring were incorporated into the design of MEG at an early stage, so that several auxiliary devices were designed, assembled and integrated for this purpose [49].

The measurement of positron and γ -ray kinematic variables can be obtained from one detector (the γ -ray energy and positron momentum) or from the combination of the information from several detectors (e.g. the relative time and direction). Hence we are prepared to deal with calibration of each detector separately, or of multiple detectors simultaneously.

Firstly the devices involved in the calibration of the LXe detector are presented: LEDs and point-like α -sources inside the LXe detector, a movable AmBe source and a 9.0 MeV γ -ray generator made of a nickel/polyethylene sandwich coupled to a pulsed neutron generator.

Then the calibration of the spectrometer by means of Mott scattering of positrons is described.

Finally the devices that allow the simultaneous calibration of multiple detectors: the 1 MeV Cockcroft–Walton (C–W) accelerator and the liquid hydrogen and NaI/BGO setup for the Charge EXchange (CEX) reaction are discussed.

A complete calibration cycle is composed of LED, α -source, lithium, boron, neutron generator runs. It takes about three hours to complete, and it is usually performed two/three times per week.

7.1 LXe detector calibration

Unlike the spectrometer, where the magnitude of the positron momentum is determined uniquely by the knowledge of the geometry of the DCH and of the magnetic field, the value

of the γ -ray energy must be extracted from the number of photons detected by the PMTs surrounding the LXe volume, once the proper proportionality factor is known. This factor contains the LXe light yield, the photo-cathodic coverage, the PMT gains and the Quantum Efficiencies (QE). All these quantities may depend on time: some more than others (e.g. the purity of LXe may change, some PMTs may be switched off for some runs).

For this reason a number of calibration lines are available to check the energy scale over the full energy range. Low-energy calibrations by means of radioactive sources are easier and performed more frequently. Although they are of limited use to set the absolute energy scale in the signal region, they are helpful in finding gross variations of LXe purity. At the opposite end, 54.9 MeV γ -rays from π^0 decays make possible to directly measure the detector response, uniformity and resolution close to the signal energy. A drawback of this calibration method is the need to change the beam polarity and momentum (from μ^+ to π^-) and the usage of a liquid hydrogen target.

Table 5 presents a list of these lines, which span a broad energy range:

1. In the low-energy region 4.4 MeV γ -rays from an AmBe source and 5.5 MeV α -particles from ^{241}Am sources deposited on thin wires are used to monitor the PMT QEs and the LXe optical properties on a daily basis. In addition, 9.0 MeV γ -rays from capture by nickel of thermalized neutrons produced by a neutron generator are also available (see Sects. 7.1.2–7.1.5).
2. In the intermediate-energy region a C–W accelerator is used, two/three times per week, to accelerate protons, in the energy range 400–900 keV, onto a $\text{Li}_2\text{B}_4\text{O}_7$ target. γ -rays of 17.6 MeV energy from $^7\text{Li}(p, \gamma_{17.6})^8\text{Be}$ monitor the LXe detector energy scale and resolution, while time-coincident 4.4 MeV and 11.6 MeV γ -rays from $^{11}\text{B}(p, \gamma_{4.4}\gamma_{11.6})^{12}\text{C}$ are used to intercalibrate the relative timing of the LXe detector with the TC (see Sect. 7.3.1).
3. In the high-energy region measurements of γ -rays from π^0 decays produced by the π^- CEX reaction in a liq-

uid hydrogen target are performed once/twice a year (see Sect. 7.3.2).

7.1.1 LEDs and gain evaluation

PMT gains are estimated by using blue LEDs immersed in the LXe at different positions. In total, 44 LEDs are installed in the detector. To minimise the position dependence of the detected photons, 11 LEDs are flashed simultaneously. In dedicated gain measurement runs, performed every second day on average, LEDs are flashed at several different intensities, and the PMT gains are evaluated from the photoelectron statistics following the method described in [41]. There are two kinds of known PMT gain instabilities, a long-term gain decrease, and a rate-dependent gain shift. The former is typically 0.1 %/day during physics runs and 0.4 %/day during CEX calibration runs (see Sect. 7.3.2).

The latter is observed when starting to use the μ beam, the typical shift being $\sim +2\%$. It was found that the LED light intensity was stable enough to check the long-term stability of each PMT gain. To monitor and correct long-term instabilities, constantly flashing LED data have been taken during physics runs since 2009 at a rate of ~ 0.2 Hz: the LED peak position for each PMT, and its variation with the time, is used to bridge the PMTs gain variation between two consecutive dedicated gain measurement runs.

7.1.2 Point-like α -sources

A calibration technique based on a lattice of ^{241}Am point-like α -sources was developed and applied, for the first time, in a prototype LXe detector [48]. The sources are prepared by fixing small portions of a ^{241}Am foil to a gold-plated tungsten wire by a thermo-compression method. The resulting wire diameter, after the source mounting, is $< 150\ \mu\text{m}$.

In the LXe detector we mounted five wires, each one hosting five point-like sources of ~ 1 kBq/source, for a total activity of ≈ 25 kBq. The wires are positioned in a staggered fashion to optimise the range of angles and distances from which they are viewed from the PMTs (see Fig. 60).

Table 5 Typical calibrations performed to determine the LXe detector performance (energy scale, linearity, etc.) with their energy range and frequency

Process		Energy (MeV)	Frequency
Charge exchange	$\pi^- p \rightarrow \pi^0 n$	54.9, 82.9	yearly
	$\pi^0 \rightarrow \gamma\gamma$		
Charge exchange	$\pi^- p \rightarrow n\gamma$	129.0	yearly
Radiative μ^+ decay	$\mu^+ \rightarrow e^+\gamma\nu\nu$	52.83 endpoint	weekly
Proton accelerator	$^7\text{Li}(p, \gamma_{17.6(14.8)})^8\text{Be}$	14.8, 17.6	weekly
	$^{11}\text{B}(p, \gamma_{4.4}\gamma_{11.6})^{12}\text{C}$	4.4, 11.6	weekly
Nuclear reaction	$^{58}\text{Ni}(n, \gamma_{9.0})^{59}\text{Ni}$	9.0	daily
AmBe source	$^9\text{Be}(\alpha_{241\text{Am}}, n)^{12}\text{C}_*$	4.4	daily
	$^{12}\text{C}_* \rightarrow ^{12}\text{C}\gamma_{4.4}$		

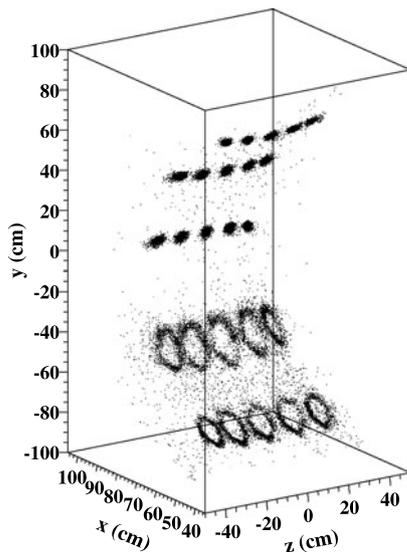


Fig. 60 Reconstructed position of the 25 α -sources inside the LXe detector. In this particular run the detector was partially filled with liquid xenon. The different reconstructed shape of α -sources in gas and in liquid xenon is visible (see text for details)

In GXe the sources are reconstructed as 3-dimensional spots of ~ 1 cm diameter (the range of α -sources in GXe is ~ 8 mm). In LXe, due to the much smaller range (~ 40 μm , comparable to the wire diameter) part of the scintillation light impinges on the wire itself, which projects a shadow on the PMTs opposite to the α -source emission direction. This produces reconstructed positions of α -sources in the form of rings surrounding each wire. It is still possible to identify all 25 sources and furthermore one can select the preferential α -source emission direction (e.g. towards the front face) by applying a topological cut on the rings.

Dedicated α -source calibration runs are collected on a daily basis. Thanks to their different pulse shapes in LXe, the trigger is able to classify each event as either low-energy γ -ray or α -source also during normal data taking (beam-ON).

The constant monitoring of the combined peak from all α -sources is used to control the light yield and transparency of LXe. A comparison of the observed light distribution on each PMT with the one obtained by a detailed Monte Carlo simulation results in a measurement of the LXe attenuation length [41]. This has turned out to be much larger than the dimension of our detector during all physics runs ($\lambda_{\text{Abs}} > 300$ cm).

7.1.3 Quantum efficiency evaluation

The relative QE of each PMT is also measured by using the α -sources. LEDs are not suitable for this task since the LED wavelength is different from that of LXe scintillation light, and the PMT response depends on the wavelength.

QE⁵ is first evaluated from the observed α -peak charge compared with that from simulation [48] and then corrected by using the γ -rays from Li peak events (see Sect. 7.3.1) to minimise systematic uncertainties related to the α -source positions and angles between sources and PMTs. The measured time dependence of the QE of each PMT has also been corrected for.

7.1.4 Americium/beryllium source

Since LXe has a different response for α -particles and γ -rays, A low-energy calibration point (see Fig. 61) is provided by 4.4 MeV γ -rays from an $^{241}\text{Am}/\text{Be}$ source. A Gaussian fit (blue spectrum) is used to extract the peak position and the energy resolution. A 50 kBq AmBe source [50] stored in a lead repository can be moved, by means of a compressed air circuit, to the front of the LXe detector. This allows a rapid check of the LXe light yield and of the correct time constants of the scintillation waveforms of α -particles and γ -rays in the same energy range. This calibration is possible only during beam-OFF periods due to the high rate of low-energy γ -rays present during normal data taking. In the latter condition the method described in the Sect. 7.1.5 is used.

7.1.5 Neutron Generator

A 9.0 MeV γ -ray line from capture of thermal neutrons in nickel $^{58}\text{Ni}(n, \gamma_{9,0})^{59}\text{Ni}$ by using a Neutron Generator (NG), allows a rapid and frequent calibration and monitoring of the LXe detector in normal (beam-ON) conditions.

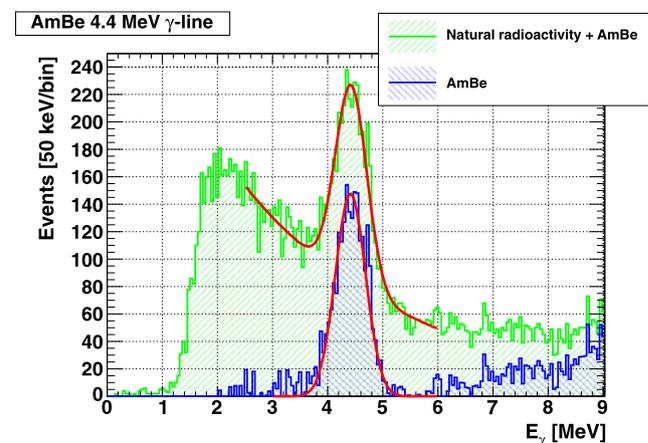


Fig. 61 The AmBe γ -ray spectrum measured by the LXe detector. A Gaussian and exponential fit is superimposed on the green spectrum. A Gaussian fit is used after subtracting the exponential natural radioactivity background

⁵The measured QE is actually the product of the photocathode quantum efficiency times the collection efficiency of the first dynode. The distribution of the measured QE of all PMTs in the (10–20) % range.

The LXe detector response can be studied when electromagnetic radiation associated with the beam (radiative muon decay and annihilation in flight) illuminates the detector. It is therefore appropriate for checking the stability of the LXe detector behaviour in case of variable μ -beam intensity. Neutrons of 2.5 MeV kinetic energy are produced by a pulsed D–D neutron generator, based on the $d(d, {}^3\text{He})n$ nuclear reaction, with a tunable frequency (1–100 Hz) and intensity (up to $2.5 \times 10^4 n/\text{pulse}$) [51]. The neutron generator is placed at the centre of a polyethylene block placed in front of the LXe detector. The calorimeter-facing side consists of alternating slabs of nickel and polyethylene. A large fraction of the neutrons are thermalized in polyethylene and then captured by nickel nuclei. The sandwich structure is optimised to allow a large fraction of the 9.0 MeV γ -rays from the capture of thermal neutrons to reach the LXe detector (see Fig. 62). This structure and the materials are able to accommodate several requirements: to minimise the unthermalized neutron flux, to avoid background from fast neutron interactions in the detector, and to minimise the thermal neutron flux out of the moderator-nickel box in order to avoid material activations.

Nickel was selected instead of other materials because of its relatively large probability ($\sim 34\%$ [52]) of emitting a single capture γ -ray at 9.0 MeV.

The pulsed mode of the NG is crucial to reach a good signal-to-noise ratio when the muon beam is operated. Thermal neutron capture events follow the NG emission by a typical average delay of 50–100 μs . Therefore a special “neutron generator” trigger is used, in which a 100 μs gate is opened 15 μs after the NG pulse. In this way the uncorrelated γ -rays from the beam are reduced, and prompt fast neutron interactions are suppressed. The neutron generator intensity is tuned in such a way as to have only one photon per pulse on average in the LXe detector.

A typical γ -ray spectrum is shown in Fig. 63. A double Gaussian fit is used to extract the peak position and energy resolution. Besides the 9.0 MeV line, the γ -ray continuum from other Ni level is visible above threshold. Data are acquired with the muon beam ON and OFF. No energy offset is observed between these data sets, if the gain calibration with

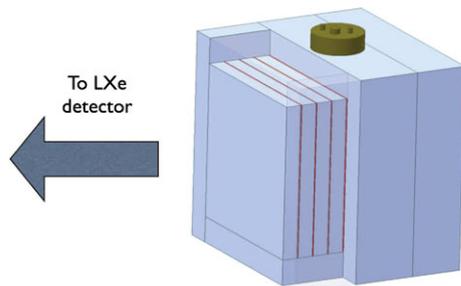


Fig. 62 The nickel slabs (red), moderator (blue) and neutron generator (green) (Color figure online)

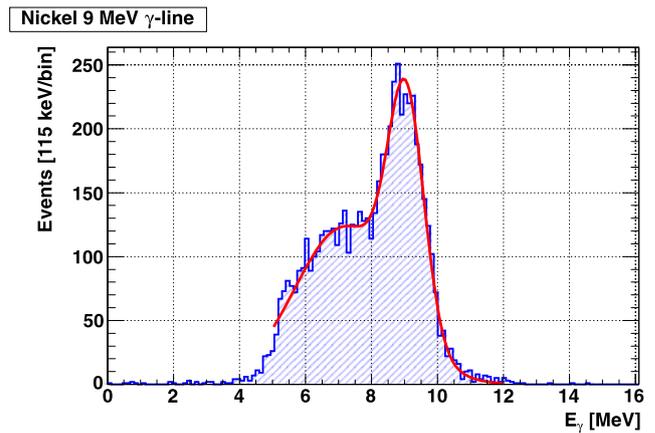


Fig. 63 The 9.0 MeV γ -ray line from neutron capture in ${}^{58}\text{Ni}$ as reconstructed in LXe detector

beam on/off is applied, proving that the LXe detector is well calibrated. Figure 64 shows the consistency of this method with the other ones.

7.2 Monochromatic positron beam

The abundant positron component of the MEG beam (eight times more intense than the μ^+ component, but normally separated and rejected) is used for selecting a monochromatic positron beam. The positrons hit the MEG target and are scattered into the COBRA spectrometer by coherent Mott scattering on carbon.

This method has several potentialities which were extensively studied with MC simulations:

- A measurement of the momentum and angular resolutions of the spectrometer with trajectories (double-turn tracks) similar to the ones of the signal.
- A measurement of the muon polarisation based on the comparison between the angular distribution of positrons

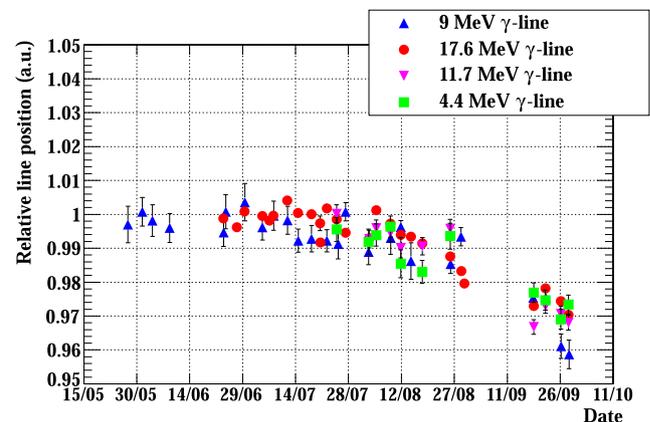


Fig. 64 The relative position of the NG 9.0 MeV γ -ray line compared with other γ -ray line relative positions, as a function of time

from muon decays at rest and that of the Mott scattered positrons.

- A measurement of the spectrometer acceptance.
- An independent check of the spectrometer alignment.

A preliminary experimental study of the positron beam properties and of Mott scattering events was performed. The average momentum of Mott events is centred close to the incident beam positron momentum of 53 MeV/c with a measured one-sigma spread of 450 keV/c, this being largely due to the minimum achievable momentum-slit opening.

The incident positron beam, associated with prompt production in the proton target, is bunched and synchronous with the accelerator 50 MHz RF-signal. This signal can be used to measure the time-of-flight of particles from the production target to the spectrometer and so distinguish beam positrons from other backgrounds, as shown in Fig. 65. Here a sharp peak due to beam positrons is seen at a relative time of -15 ns within the modulo 20 ns timing window (50 MHz). The broad peak on the right is associated with Michel positrons originating from cloud-muons interactions in the MEG target. These muons are produced from pion decay-in-flight in the surrounds of the proton production target. The smaller peak on the left is consistent with the timing expected from pions produced in the production target. Finally the constant background is associated with Michel positrons originating from muon decays after the Wien filter. With appropriate timing cuts imposed, the selected Mott positron line is shown in Fig. 66 with superimposed a fit with two Gaussian, yielding a momentum resolution of 450 keV/c (one sigma) on the high-energy side.

7.3 Intercalibration of detectors

7.3.1 The Cockcroft–Walton accelerator

A 1 MeV proton C–W accelerator [53] is coupled to MEG for calibrating and monitoring the performance of the LXe

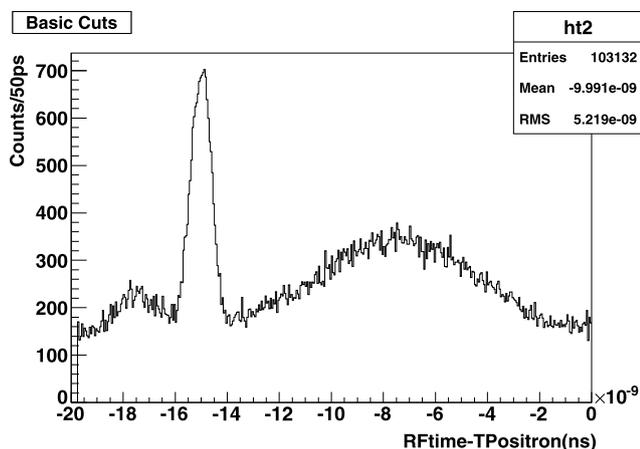


Fig. 65 Positron timing distribution relative to the accelerator 50 MHz RF-signal (see text)

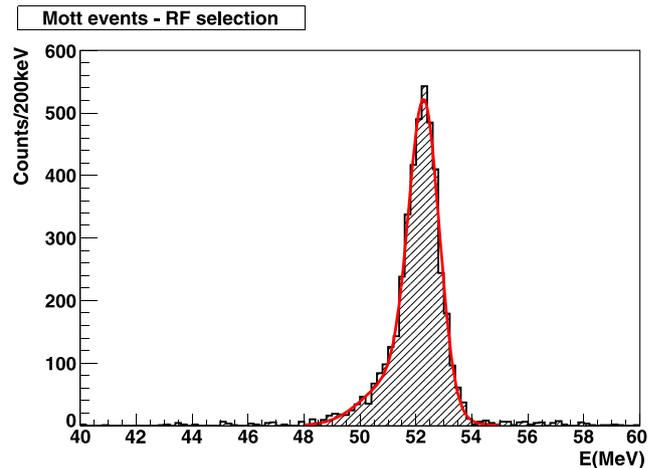


Fig. 66 The momentum distribution of RF-selected Mott events

detector, the relative inter-bar timing of the TC and the relative timing between the TC and LXe detector [39, 54]. Its properties are listed in Table 6.

The C–W accelerator is placed in a separate and independently radiation-surveyed area. in which it can be serviced and tested (see Fig. 67).

Calibration measurements are carried out with γ -rays from two different nuclear reactions: ${}^7\text{Li}(p, \gamma_{17.6}){}^8\text{Be}$ produces monochromatic 17.6 MeV γ -rays and a broad resonance centred at 14.8 MeV; ${}^{11}\text{B}(p, \gamma_{4.4}\gamma_{11.6}){}^{12}\text{C}$ produces a pair of simultaneous γ -rays of 4.4 MeV and 11.6 MeV. The two reactions have a threshold at a proton energy of 440 keV and 163 keV, with a peak cross section of $\sim 6 \times 10^{-3}$ barn and $\sim 2 \times 10^{-4}$ barn, respectively. To increase the number of generated γ -rays, especially for the reaction on boron, protons of energy higher than the threshold are used. These protons are captured during the slow-down process in the target, leading to an increase in the event rate without increasing the p -beam intensity.

The proton beam energy is therefore set to 500 keV for ${}^7\text{Li}(p, \gamma_{17.6}){}^8\text{Be}$ and to 900 keV for ${}^{11}\text{B}(p, \gamma_{4.4}\gamma_{11.6}){}^{12}\text{C}$.

Table 6 Characteristics of the MEG Cockcroft–Walton accelerator

Proton beam properties	MEG C–W
Energy [keV]	300–1000
Energy spread (FWHM) [keV]	<0.5
Angular divergence (FWHM) [mrad × mrad]	<3 × 3
Spot size at 3 meter (FWHM) [cm × cm]	<3 × 3
Energy setting reproducibility [%]	0.1
Energy stability (FWHM) [%]	0.1
Range of the average current [μA]	1–10
Current stability [%]	3
Current reproducibility [%]	10
Duty cycle [%]	100

A lithium tetraborate ($\text{Li}_2\text{B}_4\text{O}_7$) target (p -target to distinguish from the μ -target described in Sect. 2.7) is used to generate γ -rays from both reactions.

During calibration runs, the target, contained in a vacuum pipe connecting the C–W accelerator to the MEG area, is positioned at the centre of COBRA (see Figs. 67 and 68). The p -target is oriented at 45° relative to the proton beam direction, to reduce the amount of material on the path of the γ -rays directed to the LXe detector.

During normal data taking the p -target is positioned downstream outside the COBRA spectrometer. When starting a calibration, the μ -target is removed from the beam line by means of a compressed helium system and the p -target is inserted to the centre of COBRA by means of an extendable bellows system of ~ 2 m stroke. The insertion (or extraction) is computer controlled and takes ten minutes. At the end of the test the inverse operation is performed, and the μ -target reinserted. The reproducibility of its positioning has been visually inspected and surveyed to be better than our spatial resolutions.

Steering magnets and monitors are available along the proton beam line (see Fig. 68) for centring the beam on

the p -target and for measuring the proton beam properties. The data from the lithium reaction are recorded by a low-threshold trigger, while a LXe-TC coincidence trigger is used to record the two boron γ -rays (see Fig. 69).

By means of the C–W calibration lines the energy scale of the experiment is constantly monitored, as are possible drifts in the relative timing between the LXe detector and the TC bars. This allows knowing of the energy scale in the LXe detector at a few-per-mil level, and a time alignment better than 20 ps.

7.3.2 π^- beam and charge exchange set-up

To calibrate the LXe detector at an energy close to that of the signal we use γ -rays from neutral pion decay ($\pi^0 \rightarrow \gamma\gamma$). A neutral pion is produced in the CEX reaction of negative pions on protons at rest $\pi^- p \rightarrow \pi^0 n$. The resulting π^0 has a momentum of ~ 28 MeV/c in the laboratory frame and decays immediately to two γ -rays. The photons are emitted back-to-back in the π^0 rest frame with an energy of

$$E_\gamma^* = \frac{m_{\pi^0}}{2} \simeq 67.5 \text{ MeV}.$$

In the laboratory frame, the photon energies are

$$E_{\gamma 1,2} = \gamma \frac{m_{\pi^0}}{2} (1 \pm \beta \cos \theta^*), \tag{20}$$

where β is the π^0 velocity and θ^* the center-of-mass angle between the photon and the π^0 direction.

Differentiating Eq. (20), the energy spectrum of the two photons in the laboratory frame

$$\frac{dN}{dE_\gamma} = \frac{dN}{d \cos \theta^*} \times \frac{d \cos \theta^*}{dE_\gamma} \tag{21}$$

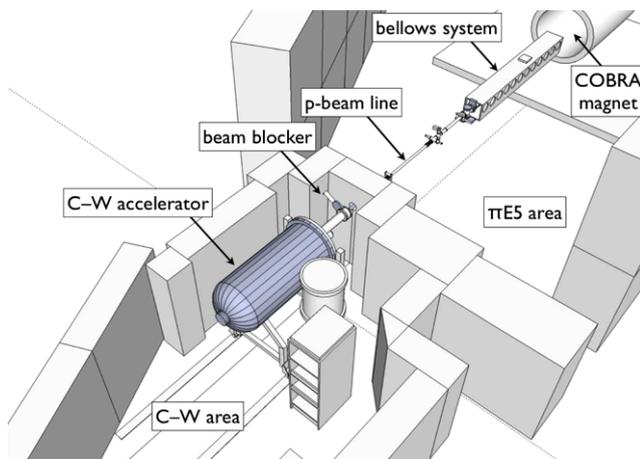


Fig. 67 Schematic layout of the area where the Cockcroft–Walton accelerator is placed, with respect to the π E5 area



Fig. 68 The Cockcroft–Walton accelerator beam line

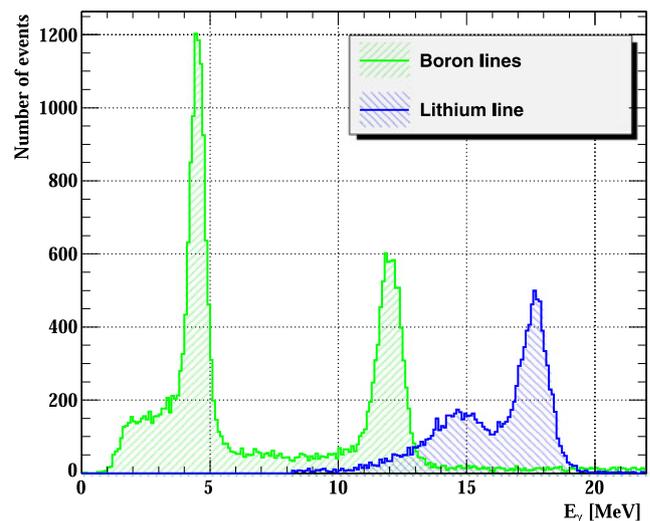


Fig. 69 Measured calibration lines from the reactions $^{11}\text{B}(p, \gamma_{4.4}\gamma_{11.6})^{12}\text{C}$ (green) and $^7\text{Li}(p, \gamma_{17.6})^8\text{Be}$ (blue) (Color figure online)

is flat (because $dN/d\cos\theta^*$ is constant) between the energies

$$E_{\gamma,\text{Min}} = \frac{m_{\pi^0}}{2} \sqrt{\frac{1-\beta}{1+\beta}}, \quad E_{\gamma,\text{Max}} = \frac{m_{\pi^0}}{2} \sqrt{\frac{1+\beta}{1-\beta}}. \quad (22)$$

A strong correlation exists between the laboratory relative angle between photons and energies. In particular the extremal energies are obtained for photons which are emitted at relative angle 180° in the laboratory, that is along the π^0 flying direction.

This strong correlation between energy and opening angle can be used to precisely define the energy of one photon by tagging the direction of the opposite one. By selecting back-to-back γ -rays in the laboratory, by means of an auxiliary detector opposite to the LXe detector (an array of NaI or BGO in our case), the energies are at the edge of the spectrum, and the higher the collinearity, the more monochromatic the two γ -rays. For instance by requiring an opening angle larger than 175° (170°) the energy spread $\sigma_{E_\gamma}/E_\gamma$ is better than 0.2 % (1 %). A few percent of the time one of the two γ -rays converts into a e^+e^- pair, either internally (Dalitz decay) or in the target material.

Additionally the π^- radiative capture on protons produces a 129 MeV γ -ray coincident with a 9.0 MeV neutron.

A negative pion beam of 70.5 MeV/c is chosen, partly due to the maximum current available for the BTS and partly because at this momentum a maximum TOF separation between pions and all other particles is achieved at the centre of COBRA. The measured intensity is 1.7 MHz at a proton current of 2 mA (lower intensities selectable via slit setting) with a round spot-size of 8 mm (sigma), fully illuminating the target cell.

A Liquid-Hydrogen (LH₂) target is used for the CEX reaction. It is a cylindrical cell of 50 mm diameter, 75 mm length placed at the end of a 2 m long support. Both the LH₂ target (cold) window and the vacuum (warm) window are made of 135 μm Mylar in order to minimise π^- multiple scattering and absorption outside the hydrogen-filled volume. The LH₂ target with dedicated transfer lines is shown in Fig. 70. During normal operation the cell is filled with 150 cc LH₂ at 20 K, kept cold by a continuous flow of liquid helium. The usage of a mechanical refrigerator is not possible due to the high magnetic field.

The operation of the target is relatively safe, since all hydrogen is contained in a ~ 100 l reservoir (the “buffer”) at 2.5 bar at room temperature. When liquid helium cools the target cell, hydrogen is liquefied at the centre of COBRA. It takes three days to set up the LH₂ target and a few hours to liquefy the hydrogen. A 250 l liquid-helium Dewar allows operation of the target for a little more than 48 hours.

To select γ -rays with a back-to-back topology a sodium iodide (NaI) detector (2008–2009) or a Bismuth Germanate (BGO) detector (2010–2012) is placed opposite to the LXe detector.

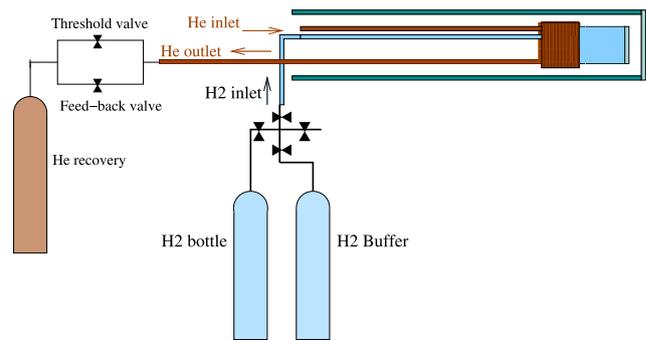


Fig. 70 Scheme of the liquid-hydrogen target with dedicated transfer lines. The 100 litres buffer is filled by means of a standard hydrogen bottle up to 2.5 bar: this quantity is sufficient to have a liquid volume of 150 cc in the cell. When the cell is cooled down the hydrogen flows from the buffer to the target cell

This “opposite side” detector is placed on a movable stage (see Fig. 71) that is remotely controlled to allow the mapping of the entire front face of the LXe detector. When displaced from the $z = 0$ plane, the NaI detector is always tilted to face the centre of the LH₂ target.

The NaI detector consists of nine crystals of NaI(Tl) scintillator and two layers of lead/plastic scintillator for timing measurements (the lead/plastic scintillator can be removed during energy measurements). Each crystal has a size of $62.5 \times 62.5 \times 305$ mm³ ($11.8X_0$) and is read out by a 10×10 mm² APD [55]. The lead/scintillator detector is made of two plastic scintillators measuring $60 \times 60 \times 7$ mm³



Fig. 71 The NaI mover

preceded by a lead converter 5 mm thick, and is placed in front of the central NaI crystal when needed.

The NaI detector was replaced for the 2010 CEX run by a BGO detector composed of 4×4 crystals, because of the better energy and position resolution of the latter. This translated into a better monochromaticity in the back-to-back topology and into a better efficiency. A typical correlation spectrum of the energy seen by the LXe detector vs. the energy seen by the “opposite side” detector in a back-to-back topology is shown in Fig. 72.

7.3.3 Radiative muon decay (RMD)

Observation of the RMD process $\mu^+ \rightarrow e^+ \gamma \nu \nu$ demonstrates that the apparatus can detect coincident $e-\gamma$ events, and measures directly the global time resolution and the relative offset. At the first stages of the experiment we took runs at low beam intensity to reduce the accidental background; later, the time resolution was good enough to also clearly see the RMD peak at the nominal beam intensity of $\sim 3 \times 10^7 \mu^+/\text{s}$. Figure 73 shows the distribution of the $e-\gamma$ relative time $t_{e\gamma}$ for RMD decays using the standard MEG trigger.

In order to obtain the resolution on $t_{e\gamma}$ for signals, the resolution in Fig. 73 must be corrected for the γ -ray energy dependence shown in Fig. 57 and for the positron energy dependence using Monte Carlo resulting in $\sigma_{t_{e\gamma}} = 122 \text{ ns}$.

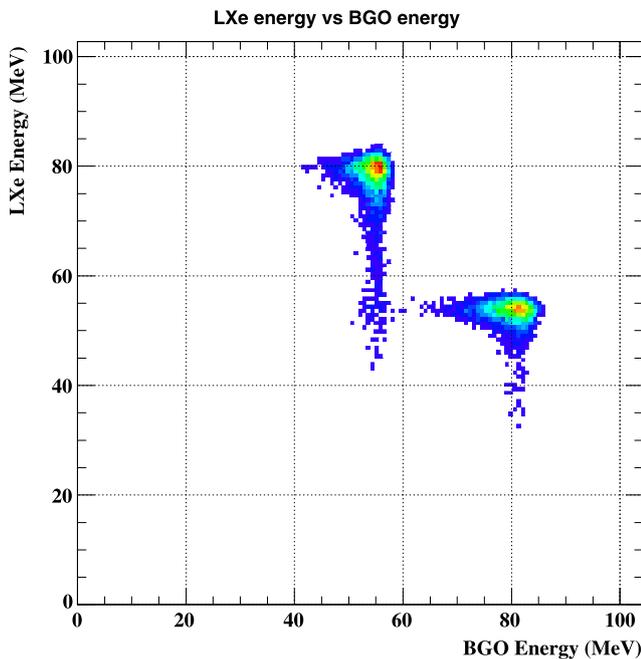


Fig. 72 Energy recorded by the LXe detector vs. the energy recorded by the BGO detector during a CEX calibration run. The extremal energies of the calibration γ -rays at 54.9 MeV and 82.9 MeV anti-correlated in the two detectors are visible

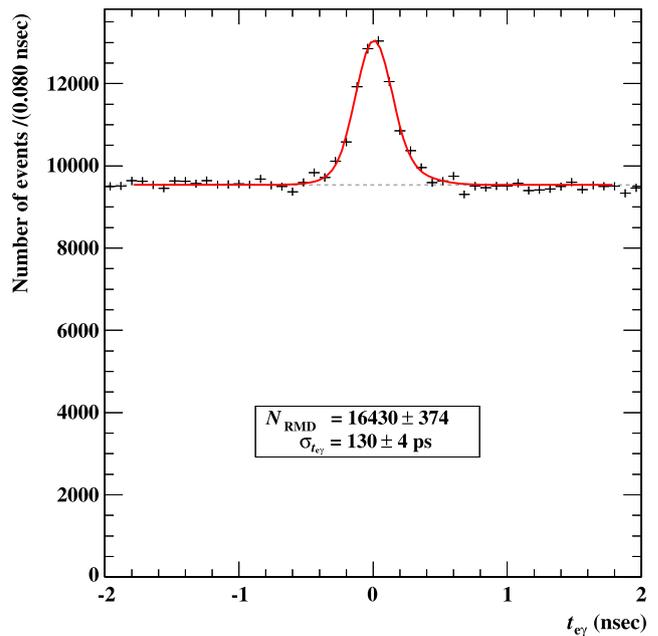


Fig. 73 Distribution of $t_{e\gamma}$ for MEG standard trigger. The peak is from RMDs, the flat component is from accidental coincidences

8 Trigger

8.1 Requirements

An experiment to search for ultra-rare events in a huge beam-related background needs a quick and efficient event selection, which demands the combined use of high-resolution detection techniques and fast front-end, digitising electronics and trigger. The trigger system plays an essential role in processing the detector signals to find the signature of $\mu^+ \rightarrow e^+ \gamma$ events in a high-background environment.

The most stringent limitation to the trigger latency originates from the matching with the DRS (see Sect. 9.2.1), whose digitisation needs to be stopped at latest 450 ns after the event occurrence. If the trigger were later, the charge stored in part of the DRS cells may be overwritten and a few samples of the waveform (in particular those recording the leading edge of PMT signals) may be lost. The speed objective must be achieved while keeping the signal efficiency close to unity. On the other hand, the event selection needs to provide a significant background rejection. The trigger rate, as explained in Sect. 8.5, should be kept below 10 Hz so as not to overload the DAQ system, which would result in a significant increase in the dead time.

8.2 The trigger scheme

In order to accomplish these goals, we developed a system consisting of VME boards arranged in a tree-structure (see Sect. 8.4), based on on-board Field-Programmable Gate Arrays (FPGA) to process the detector signals and reconstruct

the observables needed for the event selection [56]. The signals are digitised by means of 100 MHz flash-ADCs so as to obtain an estimate of the signal amplitude (at the level of a few-per-mil accuracy) and timing (with a few ns resolution). The advantage of choosing an FPGA-based digital approach is manifold:

- operation of reconstruction algorithms to match with the above speed and efficiency requirements;
- availability of RAM memories for storing the digital data stream (for both signal input and algorithm output) into cyclic buffers, which makes the trigger system an independent digitiser system as well, to be used as a back-up for the DRS;
- versatility of the trigger scheme, as it is possible to implement selection criteria, other than for $\mu^+ \rightarrow e^+\gamma$ decays, related to calibration and other event types; the acquisition of pre-scaled triggers associated with those events during normal data taking enables us to monitor the stability of both detector response and trigger efficiency as a function of time.

8.3 Online algorithms

The set of observables to be reconstructed at trigger level includes:

- the γ -ray energy;
- the relative e - γ direction;
- the relative e - γ timing.

The stringent limit due to latency prevents use of any information from the DCH; the electron drift time toward the anode wires may be as long as 200 ns, too long to match the latency requirement. Therefore a reconstruction of the positron momentum cannot be obtained at the trigger level.⁶ Even if there is no explicit requirement on the e^+ in the trigger algorithm, the requirement of a TC hit is equivalent asking for a positron with momentum $\gtrsim 45$ MeV.

The γ -ray energy is the most important observable to be reconstructed. The steep decrease of the spectrum at the endpoint (equal to the signal energy for $\mu^+ \rightarrow e^+\gamma$) emphasises its role in suppressing the background. For this reason the calibration factors for LXe PMT signals (such as PMT gains and QEs) are continuously monitored and periodically updated in a dedicated database. The energy deposited in LXe is estimated by the linear sum of PMT pulse amplitudes, which is accomplished by cascaded adder stages (with the

⁶Although excluded from the $\mu^+ \rightarrow e^+\gamma$ trigger, DCH anode signals are collected indeed by the system to provide stand-alone triggers associated with tracks in the DCH. For each chamber, anode signals are grouped in two sets (as coming out of inner and outer wires) and fanned-in linearly before being fed to Type 1 boards. The track condition is met if four out of five consecutive chambers are hit.

proper bit resolution) up to the master trigger board. The sum of about 250 channels (one for each inner PMT, one to four PMTs for other LXe faces) is sensitive to coherent noise in the electronics. To cope with that, adder stages are preceded by the subtraction of the signal pedestal, updated online by the average of out-of-pulse samples.

The interaction time of the γ -ray in LXe is extracted by a fit of the leading edge of PMT pulses, with a resolution of the order of 2 ns. The same procedure allows us to estimate the time of the positron hit on the TC with similar resolution. The relative time, a useful tool to suppress our accidental background, can thus be obtained from their difference; fluctuations due to the time-of-flight of each particle to their respective detectors are within the resolutions.

The amplitudes of the inner-face PMT pulses are also sent to comparator stages to extract the index of the PMT collecting the most photoelectrons. This provides a robust estimator of the interaction vertex of the γ -ray in LXe, the Molière radius being similar to the PMT size (see Table 3). This vertex and the target centre provides an estimate of the γ -ray direction. On the spectrometer side, lacking any DCH information, one has to rely on the coordinates of the TC online reconstructed hit. Under the assumption of the momentum being that of the signal and the direction opposite to that of the γ -ray, by means of Monte Carlo simulated events, the PMT index is associated with a region of the TC indexed by a bar number and a z -segment. If the online z -coordinate of the TC hit (see Sect. 5.5) falls in this region, the relative e - γ direction is compatible with the back-to-back condition.

8.4 Hardware implementation

The trigger system is organised in a three-layer hierarchical structure, as shown in Fig. 74, which consists of:

- (1) VME 6U boards (so-called “Type 1”), each to receive and sample 16 input analog detector signals at

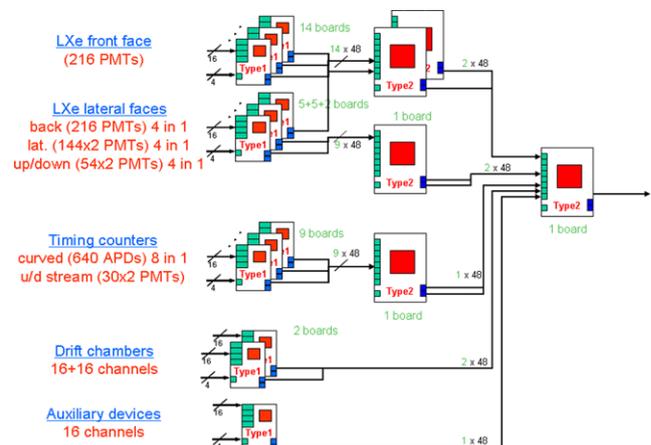


Fig. 74 Hierarchy of the trigger system: (from left to right) Type 1 boards; Type 2 intermediate boards; Type 2 master board

100 MHz by the 10-bit flash-ADCs [57] to be processed by means of digital algorithms implemented on a Virtex XC2VP20 FPGA chip [58];

- (2) VME 9U boards (named “Type 2”), each equipped with a Virtex XC2VP20 FPGA chip [58], to collect digital pieces of information from Type 1 boards associated with different detectors and combine them to obtain estimates of the kinematic observables of interest (namely energy, direction and time of flight) of decay particles.

The input stage of Type 1 boards consists of plug-in Front-End (FE) mini-cards to receive the signals from detectors (in AC coupling) via the active splitters (see Sect. 9.2.2). A low-pass filter with cutoff at 33 MHz (to match Nyquist’s criterion) is obtained by means of an RC integrator through the feed-back branch of an AD8138 differential driver. The common-mode voltage of the output signal can be further adjusted by using a programmable AD5300 DAC to match the dynamic range of the FADCs (which is set to 0–2 V in the case of LXe inner-face PMTs and 0–1 V for all other channels).

The FPGA digital output is arranged in a 48-bit word, encoded with charge, time and other pieces of information (which depend on the detector) relevant to event reconstruction. The output bus is transmitted to the upper-layer boards by means of a DS90CR483 Low Voltage Differential Signal (LVDS) serialiser clocked at 100 MHz, the same frequency as the FADC sampling and FPGA algorithm execution.

The FPGA inputs and outputs are both recorded into cyclic memory buffers, as shown in Fig. 75. The comparison of the output memory content with the result of a simulator running on the same input data enabled debugging of the execution of reconstruction algorithms.

At the intermediate layer, Type 2 boards collect and combine the information transmitted by nine Type 1 boards via the companion DS90CR484 LVDS deserialisers. Again algorithm output arranged in 48-bit buses (up to two per board) is transmitted to the trigger master board via LVDS serialisers and there received by deserialisers. Finally, the task of the master is to assemble all event characteristics

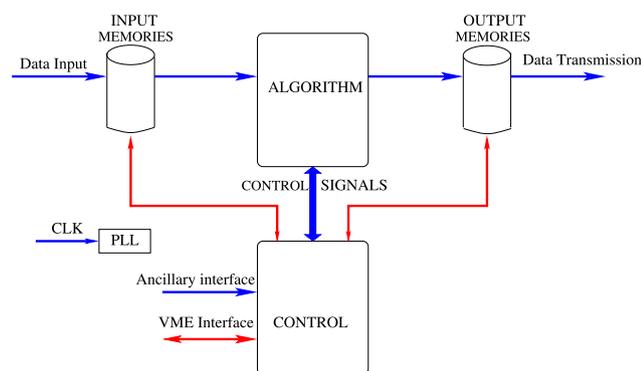


Fig. 75 FPGA firmware structure

and generate a Stop signal whenever a trigger condition is met. Moreover, this board is the one to receive Busy signals raised by each digitiser during data transfer and to issue a Start command as soon as that Busy condition has been released by all digitisers.

The synchronous operation of the system (FADC digitisation, algorithm execution and data flow) is guaranteed by a set of “Ancillary” boards distributing clock signal and control (Start and Stop) signals to all digitising electronics. A master ancillary board hosts a SARONIX SEL3935 oscillator to deliver the 19.44 MHz master clock which, although being not essential for trigger synchronisation, is quite important for ensuring accurate detector timing. On the trigger side, each board utilises a PLL (based on a CY7B994V Roboclock [59]) to multiply the input clock frequency by a factor 5 so as to generate 16 replicas of the 100 MHz clock signal throughout a board. The relative phase between these can be programmed to compensate for the relative timing jitter between chips (and, in the case of the FPGA, even between lines in a bus) and ensure proper synchronisation over the whole trigger tree.

Trigger boards are hosted in three VME crates and accessed by front-end code running on as many online PCs. Each board is equipped with a CoolRunner XC2C384 chip [58] as an interface between all the chips and the VME bus. Input and output data are read out through a standard, double-edge block transfer VME master-slave access to ensure a throughput rate close to 80 MB/s. Single word (both read and write) VME commands are used to access configuration registers to set behavioural parameters (such as thresholds, look-up table entries, calibration factors, clock skews, etc.) as well as to get the status of control signals (namely the Busy) from all the digitisers.

8.5 Operation and performance

The γ -ray response function is obtained with CEX calibration close to the signal energy and continuously monitored (although at lower energies) by using C–W proton-induced γ -ray lines; the resolution is estimated to be 3 % at 52.83 MeV. The equivalent energy threshold for $\mu^+ \rightarrow e^+ \gamma$ events is computed by comparing the energy spectrum with the one obtained at lower thresholds (from 30 MeV to 40 MeV) thanks to a pre-scaled trigger operated in parallel with the main one. The ratio of the two spectra is shown in Fig. 76 and follows a step behaviour smeared with a Gaussian function with an energy resolution in agreement with that estimated by the fit of the 55 MeV line from CEX data. The threshold on the γ -ray energy measured online is set to 45 MeV, which guarantees almost full efficiency ($\gtrsim 99.5\%$) for signal events; on the other side, the single rate drops from a few hundred kHz (the bulk of which is related to γ -rays from radiative muon decays) to 1.5 kHz.

To assess the relative $e\text{-}\gamma$ timing, a 20 ns-wide window is fixed to look for coincidences between the two particles. Relative offsets are calibrated by reconstructing events due to simultaneous γ -rays (one detected in LXe, the other in the TC) from $B(p, \gamma\gamma)C$ (see Sect. 7). The resolution is estimated a posteriori by fitting the edges of the window steps with an integral Gaussian function and amounts to 3 ns, in agreement with the single-detector timing estimate. The coincidence window is wide enough to achieve full efficiency.

The $e\text{-}\gamma$ Direction Match (DM) enables a further reduction of the accidental background trigger rate, which is proportional to the square of the opening angle [60], by one order of magnitude at least, as expected by detector and trigger simulation, as shown in Table 7.

In Table 8 a subset of the available selections both for physics and calibration are listed following their priority order. For calibration triggers refer to Sect. 7 for details.

The final goal is to optimise the overall DAQ efficiency, defined as the product of the trigger efficiency times the DAQ live time. Maximising the former is clearly advisable, but may be wasted if accompanied by a significant increase in the trigger rate and, as a result, in the dead time. The best trade off between the two effects is shown in Fig. 77; due

Table 7 Trigger rates as a function of progressively applied selection criteria (γ -ray energy, relative $e\text{-}\gamma$ timing and direction) at $3 \times 10^7 \mu^+/\text{s}$ stop rate on target. Measured values agree within 20 % with expectations based on detector simulation and reconstruction algorithm emulation

Selection	Measured (s^{-1})	Expected (s^{-1})
$E_\gamma > 2 \text{ MeV}$	3×10^5	2×10^5
$E_\gamma > 45 \text{ MeV}$	1.5×10^3	1.3×10^3
$\Delta t_{e\gamma} < 20 \text{ ns}$	150	120
$\Delta\theta_{e\gamma} \sim \pi$	12	10

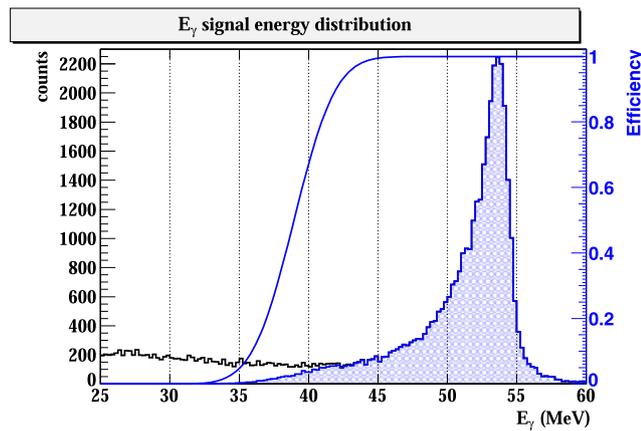


Fig. 76 Spectrum of E_γ measured offline by the LXe detector, as obtained with the threshold at 45 MeV on the online measured γ -ray energy, superimposed to a lower-threshold one. The corresponding efficiency curve, resulting from the ratio of the two spectra, is also shown

to limitations inherent in the DAQ (the digitiser read-out lasting $\sim 24 \text{ ms/event}$), the TC acceptance region is modified until the DAQ efficiency reaches the maximum allowed value. Therefore the only possibility left to overcome such limitations was to reduce the dead time by implementing the multi-buffer read-out scheme, as described in Sect. 9.4.1. After this major improvement, we could loosen the DM-

Table 8 List of the most important triggers available for physics and calibration runs (see text for details)

Name	E_γ	$ \Delta t_{e\gamma} $	DM
MEG	45 MeV	10 ns	narrow
MEG Low E_γ	40 MeV	10 ns	narrow
MEG Wide DM	45 MeV	10 ns	wide
MEG Wide $ \Delta t_{e\gamma} $	45 MeV	20 ns	narrow
Radiative Decay	45 MeV	10 ns	no

Name	Selection criteria
LXe alone	Threshold on E_γ
DCH alone	DCH hit multiplicity
TC alone	TC hit multiplicity
π^0 decay	Coincidence LXe & BGO
α	Pulse shape discrimination
LED	Discrimination on driver
Neutron Generator	Discrimination on driver
Pedestal	Random trigger

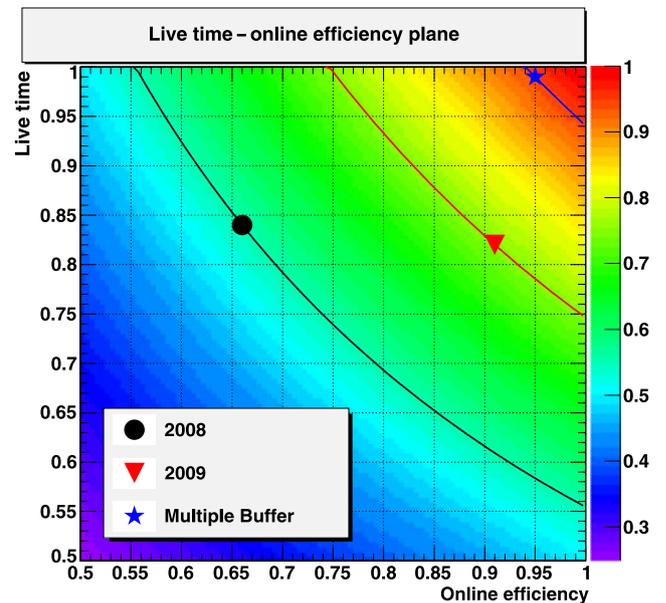


Fig. 77 DAQ efficiency as a function of trigger efficiency and DAQ live time. Superimposed markers correspond to optimal values reached during 2008, 2009/2010 and 2011/2012 runs (marked as “Multiple Buffer”)

match selection criteria and the trigger efficiency increased to 97 %.

9 Front-end electronics and Data Acquisition system

9.1 Requirements

The Data Acquisition system (DAQ) is dedicated to recording the signals from all detectors of the experiment. The requirements for this high-precision and high-rate experiment are challenging. Energy has to be measured with an accuracy better than a few per mil, even for small signals. The timing accuracy of the electronics has to be better than 40 ps, and the pile-up detection has to distinguish events separated in time down to 10 ns.

The energy from the LXe detector is determined by adding several hundred channels. Even a small level of common noise on these channels, such as 50 Hz line noise, would result in a huge offset of the summed signal. It has therefore been decided to digitise the waveforms of all detectors with high resolution and high speed. This allows an efficient pile-up detection and an event-by-event baseline measurement to eliminate any common low-frequency noise. In order to obtain the required timing resolution, the signals must be digitised with a sampling speed above one Giga-Samples Per Second (GSPS).

9.2 Front-end electronics

While FADCs with sampling speeds beyond one GSPS are commercially available, they are impractical to use because of their costs, space and power requirements considering the 3000 channels of the experiment. An alternative is the use of Switched Capacitor Arrays (SCA). They sample input signals in analog storage cells at high speed. After a trigger, the cells can be read out at much lower speed in the MHz range, and digitised by a commercial ADC with typically 12-bit resolution. Since no suitable SCA was available when the experiment was designed, the design of a new SCA chip fulfilling all requirements was started [61]. The latest version of this chip called Domino Ring Sampler 4 (DRS4) has been used successfully in the experiment since 2009.

The sampling frequency was chosen as 1.6 GHz for the TC and LXe detectors, which perform high-precision timing measurements, and 0.8 GHz for the DCH, which has less stringent timing requirements.

9.2.1 The DRS4 Switched Capacitor Array chip

To obtain the high sampling speed of a SCA chip, the write switches of the analog sampling cells are operated from the inverter chain shown in Fig. 78, and called the domino wave

circuit. This eliminates the need to use a GHz clock, which would be very hard to distribute over thousands of channels.

The variable resistors between the inverters, together with the parasitic input capacitance of the next inverter, a delay element, which is used to vary the sampling speed with a control voltage U_{speed} . The analog inputs and outputs of the DRS4 chip are differential, which reduces the crosstalk between channels. The read-out shift register can be operated up to 33 MHz. However, we use only half this frequency, since the chip read-out dead time is not an issue and the lower frequency improves the analog performance of the chip. Figure 79 shows the functional block diagram of the DRS4 chip.

The analog control voltage to operate the inverter chain is obtained from a Phase-Locked Loop (PLL), which fixes the sampling speed with high precision to an external reference clock over a wide temperature range. The chip contains eight data channels and a ninth timing channel used to sample the external reference clock directly. Since the same domino wave circuit operates all nine channels, any timing jitter in that circuit can be measured in the timing channel and corrected offline to improve the timing accuracy. The analog output of all nine channels is multiplexed into a single ex-

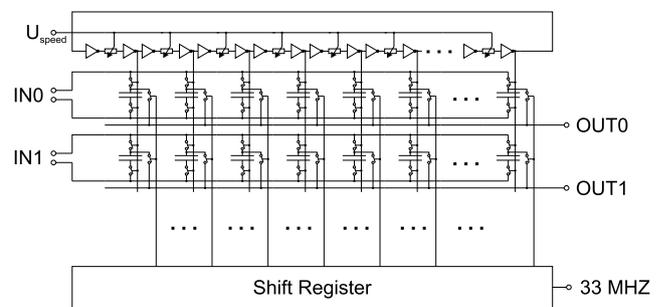


Fig. 78 Simplified schematics of the inverter chain of the DRS4 chip

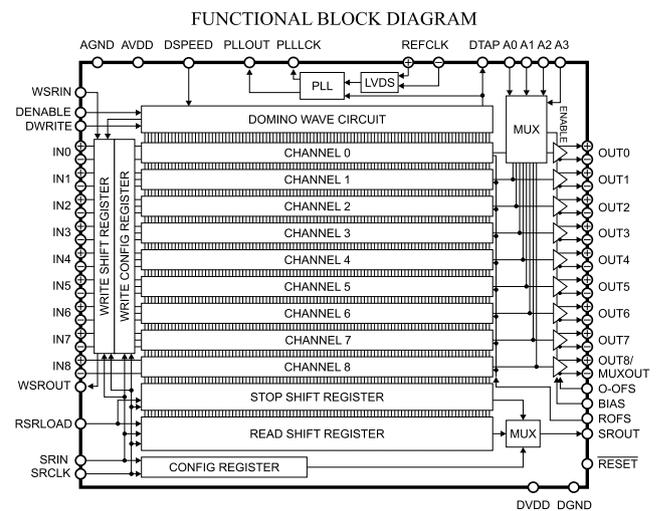


Fig. 79 Functional block diagram of the DRS4 chip

ternal 12-bit ADC, which significantly reduces the cost and space requirements.

9.2.2 Splitters

A system of active splitters receives the signals from the LXe detector and from the TC through the passive splitters. Each splitter board receives 16 signals through coaxial cables and makes two copies for the DRS and trigger. Differential signals are used to improve the noise immunity. Output signals are sent to digitization systems through flat cables (high-density twisted-pairs cable, 0.635 mm pitch, for the DRS and low-density twisted-pairs cable, 1.270 mm pitch, for the trigger). The gain for each channel is one and the bandwidth is about 1 GHz for the DRS output and 100 MHz for the trigger output. Careful choice of the components and their placement ensure a small integral non-linearity. A dedicated test measurement resulted in a combined non-linearity of the splitter system and the DAQ boards of $\pm 0.3\%$. The crosstalk between each channel for the full bandwidth section is $< 0.8\%$. Furthermore, the splitter board provides an analog sum $4 \div 1$ of the LXe detector's lateral face signals, which is used by the trigger system. A block diagram of the splitter section is shown in Fig. 80.

9.2.3 DAQ boards

The DAQ boards are based on the VME PMC Carrier (VPC) (PCI Mezzanine Card (PMC)) board developed at PSI [62]. This general-purpose VME board contains two Virtex-II Pro FPGAs [58], one of which is used to implement the VME64x protocol, while the second contains application-specific code. The board contains two PMC slots, which carry two DRS4 mezzanine boards with a total of 32 channels. The high channel density makes it possible to fit 640 channels into a single VME crate. The experiment uses three VME crates to host 1728 channels for the DCH read-out, and two crates to host 846 channels for the LXe detector read-out and 120 channels for the TC read-out.

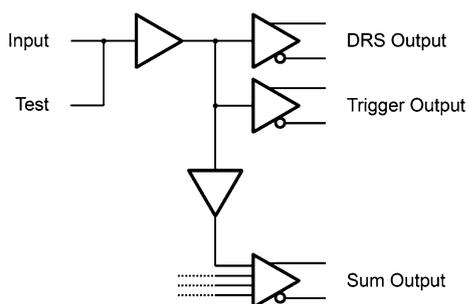


Fig. 80 Block diagram of a splitter channel for the LXe detector

9.2.4 Trigger bus

The synchronisation between the DAQ boards and the trigger is done via the trigger bus, which is a cable running between the master trigger unit and the DAQ boards. The first 12 lines send signals from the trigger to the boards. They contain the main trigger, a synchronisation signal, five bits carrying the trigger type and five bits with the least significant bits (LSB) of the event number. The trigger type is one of the 32 types defined in the main trigger (see Sect. 8 and Table 8). The boards and the front-end program can decide to handle different events flexibly: for example keeping the full waveforms for physics events, while analysing on-line the waveforms and storing only integrals and timings for calibration events.

The last eight lines of the trigger bus send signals from the DAQ boards to the trigger. The main signal is the “busy” line, blocking the trigger until the board read-out is finished. Since all lines of the trigger bus run between FPGAs, their meaning can be changed by modifying the firmware only. This has proved to be very flexible and was successfully used several times to implement new functionality such as a delayed trigger.

9.2.5 DAQ synchronisation

While the trigger signal can be successfully used to start the read-out of all DAQ boards, the timing it carries is rather poor. The bus structure of the signal distribution causes this signal to arrive at slightly different times at each board. In addition the signal is routed through the FPGAs to trigger the DRS4 chips. This causes a jitter of the signal of ~ 100 ps, which exceeds the experiment timing requirements. To overcome this problem, a dedicated high-precision clock distribution was implemented. A low-jitter 19.44 MHz master clock is generated in the trigger and fanned out on a clock tree with 200 Low Voltage Differential Signalling (LVDS) lines, each of them running to an individual mezzanine board. The clock is fed into a clock conditioner chip [63] on each mezzanine board, which removes most of the jitter picked up by the cable, leaving a residual clock jitter below 4 ps. One output of the chip is directly digitised by a dedicated timing channel in each DRS4 chip and used as a timing reference for each event.

The performance of the system was tested with a 50 ns square pulse with rise and fall times of 2.5 ns. It was passively split and the two copies were sent to two arbitrary channels of the digitising electronics. Their relative timing was measured by waveform analysis. In 2009 the timing accuracy of the DRS4 chip was found to be ~ 110 – 170 ps. This poor timing had two causes. The first was the choice of the external components for the loop filter of the PLL (see Sect. 9.2.1). A re-optimisation of these components

improved the timing significantly. The second cause was the continuous digitisation of the DRS4 chip output by the ADC. That generated digital noise which affected the timing jitter of the DRS4 chip. A firmware modification was then applied to stop the ADC when the DRS4 chip is actively sampling data, and starting it only during read-out. Minimising these two effects, the timing accuracy became $\sim 40\text{--}50$ ps, comparable to or smaller than the detector timing resolution so as not to degrade the experiment's performance. All mezzanine boards were modified in June 2010, so all data since then were taken with good timing resolution.

9.3 DAQ system

The data acquisition system (see Fig. 81) is dedicated to reading out the front-end electronics, combining the information into events and storing them on disk. It consists of a cluster of PCs connected to the front-end electronics and an event-building PC connected through a Gigabit Ethernet switch. They run the DAQ software, which handles all aspects of event-based data acquisition and slow control.

9.3.1 Online cluster

The VME boards of the front-end electronics and trigger system are located in nine VME crates. Since they produce a significant amount of data, each crate is read out by a separate PC via the Struck SIS3100/1100 VME controller over a 15 m optical cable [64]. The PCs contain two dual-core Xeon CPUs running at 2.8 GHz under Scientific Linux 5.1. The multi-core functionality is important since the PCs that read out the DRS chips have a high computing load for the online calibration and zero suppression of the raw waveform data. If the front-end program run on a single core, the maximum event rate would be only 12 Hz. Running separate threads for the VME readout, calibration and network transfer of calibrated events increases the rate to 38 Hz. The PCs have redundant dual power supplies and hot-pluggable mirrored hard disks, so that the failure of a power supply or a disk would still allow data taking to continue. The front-end PCs are connected to the central back-end PC via a Gigabit Ethernet switch. This PC uses the same technology, but it has more disks. A Redundant Array of Independent Disks (RAID 5) system with a net capacity of 2 TB can buffer about 2–3 days of data taking. A dedicated process copies the data from these disks constantly to the central PSI computing centre, where a dedicated computer cluster analyses the data (offline cluster). The offline cluster contains 16 nodes, each with four CPU cores, and a disk system with a total of 280 TB, and is used both for data analysis and for simulation.

9.3.2 MIDAS software

The data acquisition software used for the MEG experiment is based on the Maximum Integrated Data Acquisition System (MIDAS) [65]. MIDAS consists of a library written in C and C++, several tools and a general program framework to assist the development of experiment-specific front-end programs as well as analysis programs for online and offline use. The MIDAS library implements the following functionality:

- *System Layer*. All Operating System (OS) specific functionalities are covered by a set of functions using conditional compilation. In this way the upper level of all MIDAS programs can use these general functions and compile on different OSes. Over the last twenty years, many OSes have been supported. Some of them have been phased out such as MS-DOS and VMS, others such as OSX (Darwin) have been added. Since the system layer is relatively small (about 7 % of the MIDAS code), new systems can be added relatively quickly.
- *Online Database (ODB)*. The online database is a hierarchical database which stores all relevant information in a central location. For performance reasons it is implemented completely in shared memory, so that all experiment processes can exchange data very rapidly at an access rate ~ 1 MHz. The MEG ODB is about 5 MB in size and shared by about 20 processes. The ODB contains configuration parameters for the trigger system, DAQ VME boards and event building, as well as the current values for the slow control data and event statistics. A special feature of the ODB are “hot-links”. A process can register to be notified through a call-back routine when a specific value in the ODB changes.
- *Buffer Manager (BM)*. The buffer manager is a set of library functions, over which processes can exchange event-based data. Buffers are segments of shared memory to which processes can attach as data producers or consumers. Data is passed very efficiently in a first-in first-out (FIFO) manner either locally or over the network using a dedicated TCP/IP protocol. On the Gigabit Ethernet network of the MEG experiment a net transfer rate of 98 MB/s has been measured. Each consumer can either register to receive only a subset of all events (such as for monitoring) or request all events (such as the logger writing data to disk). In the latter case the consumer automatically slows down the data producer in case the data cannot be digested as fast as produced (commonly known as “back-pressure”).
- *System Messages (SM)*. Using the BM with a dedicated buffer, each process sends or receives text messages to inform the users about errors or status changes such as run start/stops. All SM are centrally recorded during the duration of the experiment.

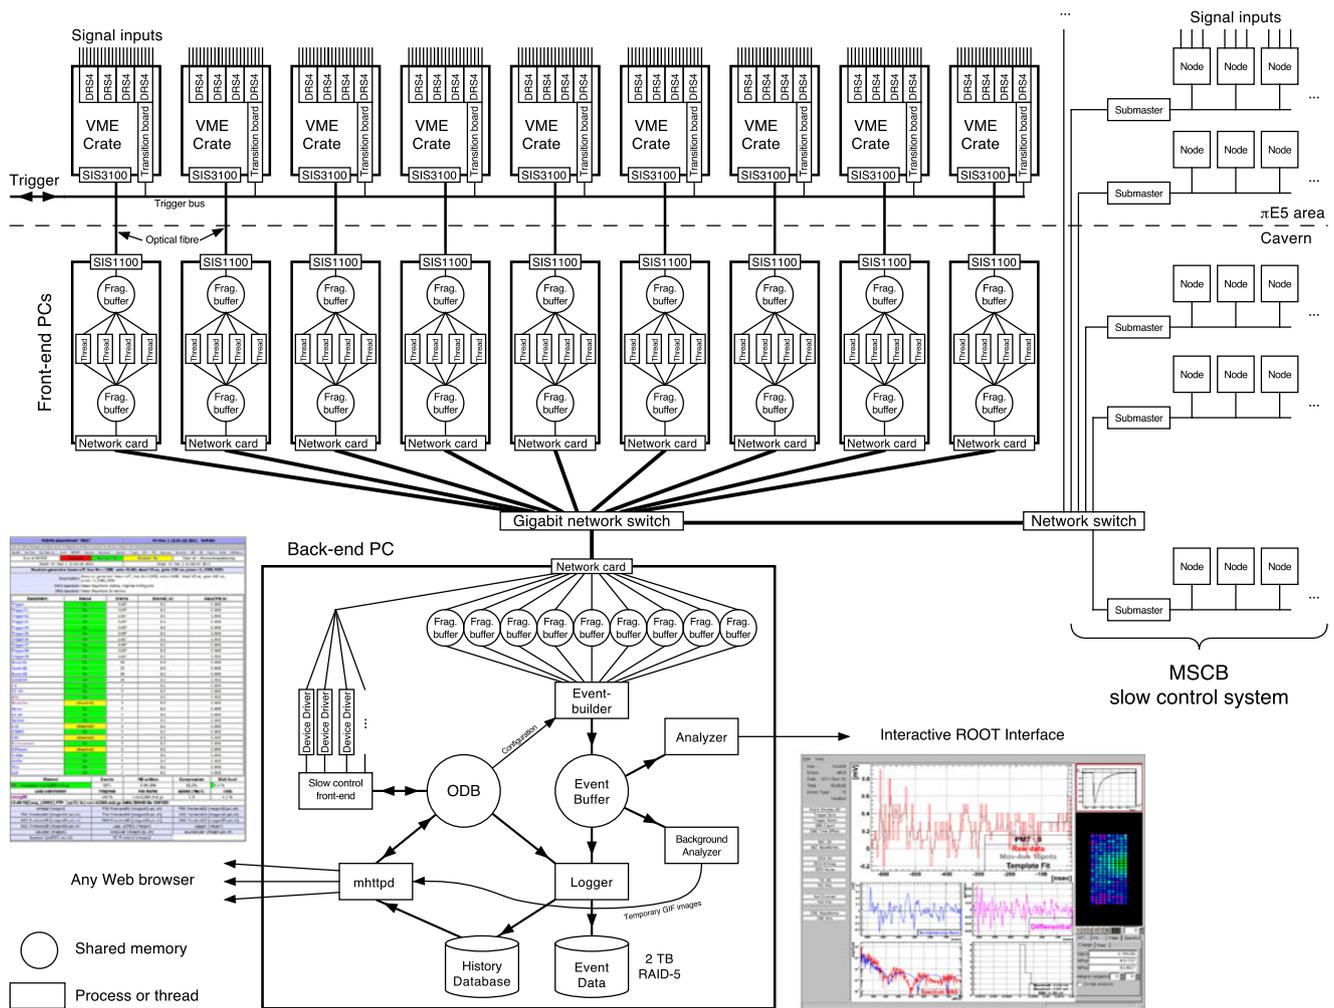


Fig. 81 Simplified schematic of the DAQ system. The individual components are described in the text. The upper part resides next to the detector, the lower one just outside the experimental area. The electronic logbook, the system messages and the alarm system are omitted for clarity

- *Slow Control System (SC)*. MIDAS contains a framework to access “slow control” devices (see Sect. 10). The framework takes care of transferring information between a device and the ODB (read-out) and between the ODB and the device (control) by means of a device driver implementing two functions for reading and writing.
- *History System (HS)*. While the ODB contains only the current value of all slow control values, it is desirable to see changes of the variables over time. The history system implements a database for storing and retrieving all slow control variables over the duration of the experiment. An optimised binary database is capable of writing up to 10^6 variables per second to disk, and retrieving the time variation of any variable over days or months in a couple of seconds.
- *Event Builder (EB)*. The event builder is a dedicated process which combines event fragments from nine different front-end computers into a single event. Software and hardware counters are used to ensure that only fragments

belonging together are combined. In case of a fragment mismatch the event builder stops the current run and issues an error message.

- *Logger*. The logger is an experiment-independent program which stores the events produced by the EB on disk. In addition, it writes the slow control data from the ODB into the HS utilising the above described hot-link scheme. The logger contains basic run control functionality such as stopping a run after a predefined number of events and automatically starting the next run. The logger writes data to a local disk array. A second logging process copies the data through a FTP interface to a central computing cluster, from where they are transferred to a tape archive.
- *Alarm system*. The alarm system is a part of the MIDAS library which runs in each process attached to the experiment. It checks individual values in the ODB or for the presence of a process. If a value exceeds a predefined limit or if a process dies, the alarm system can produce a system message, stop the run and execute a shell script,

- which in turn can send emails or text messages to mobile phones.
- *Online analyzer.* An analysis program based on the ROME framework [66, 67] runs in the online environment sampling and analysing a subset of the data stream. It combines a full physics reconstruction with a single event display. Using various panels, sub-detector hit patterns as well as short-term histograms can be inspected by the shift crew to identify potential problems. A second copy of the program runs on the offline cluster to analyse all events in near-to-real time. Histograms containing all events are available 10–20 minutes after the data has been taken, and allow monitoring of effects which become visible only with adequate statistics. More information about the analysis software can be found in [68, 69].
 - *Web interface.* A dedicated web server (mhttpd) is used to control the whole experiment through web browsers. It contains an interface to the ODB, so the experimenter can view or change any ODB variable. In addition, the interface implements panels for run control (start/stop a run plus run status display), system messages, the alarm system, the history system and a process interface to start and stop individual programs. While the mhttpd program is experiment independent, “custom pages” can be defined to monitor and control a specific part of the experiment (see Sect. 10.2.2). The web interface is also available from off-site locations for remote monitoring.
 - *Run control.* While the standard MIDAS system has basic run control functionality, more sophisticated run sequencing is needed. For example, for the LED calibration of the LXe detector one needs a sequence of several runs, each with a different LED flashing amplitude. To fulfil these needs, a dedicated utility called “runsubmit” was written. This tool reads an Extensible Markup Language (XML) script and executes a sequence of tasks. It can start and stop runs after a certain number of events, and change values in the ODB. This allows the tool to change the trigger settings and prescaling factors, configure the data read-out mode and even change slow control values such as high voltages or beam line settings. Using the sequencing, several runs can be taken automatically with different settings.
 - *Electronic Logbook (ELOG).* Instead of using paper logbooks, the collaboration uses an electronic logbook based on the ELOG program [70]. This program, originally an integral part of MIDAS, is now available as a stand-alone package. The usage of an electronic logbook has proved extremely helpful for communication between the on-site shift crew and off-site collaboration members, as well as for the retrieval of old logbook entries tagged with certain keywords. An additional feature is a shift checklist, filled out once by each shift, requiring the inspection of critical elements in the experimental area not covered by

automatic sensors. This checklist is a web form integrated into the ELOG program and filled out using a tablet computer in the experimental area. This allows a truly paperless operation of the experiment.

The MIDAS software runs very stably; in MEG the DAQ system is only stopped for configuration changes. If left in production mode, it runs for weeks, making most effective use of the beam time.

9.4 Performance

The performance of the DAQ system can be defined as the fraction of recorded events for a given trigger rate. This includes dead time caused by the read-out, lost events during data transport, events not recorded between runs, and events lost during the downtime of the DAQ system due to re-configurations or crashes. Muon beam downtime and calibration runs do not contribute to DAQ inefficiency.

9.4.1 Dead time during event read-out

The DRS is stopped on each trigger and read out before being re-activated. One waveform contains 1024 samples, each stored in a 16-bit value, digitised at 16.5 MHz. Each DRS chip uses four channels for detector waveforms and one for the global reference clock. Two DRS chips are read out through a multiplexer by a single ADC, which requires $61 \text{ ns} \times 1024 \times 5 \times 2 = 625 \text{ } \mu\text{s}$. Each VME board contains two mezzanine boards each having four DRS chips. This amounts to $1024 \text{ (bins)} \times 5 \text{ (channels/chip)} \times 4 \text{ (chips/mezzanine)} \times 2 \text{ (Bytes/sample)} = 40 \text{ kB}$ per event and mezzanine board. A full VME crate contains 20 boards with 40 mezzanine boards, which totals to 1.56 MB per crate and event. The SIS3100/1100 VME interface has a setup time of $125 \text{ } \mu\text{s}$ per DMA transfer and then a line speed of 83 MB/s. The transfer of the 40 data blocks therefore takes $125 \text{ } \mu\text{s} \times 40 + 1.56 \text{ MB}/83 \text{ MB/s} = 23.8 \text{ ms}$. The data read-out rate from the trigger system is smaller and therefore does not pose a bottleneck. Before 2011, the front-end electronics was only re-enabled after all data were read out through the VME bus, thus causing a dead time of $23.8 \text{ ms} + 0.625 \text{ ms} = 24.4 \text{ ms}$ per event. This limited the maximum event rate to 41 Hz and caused a live time of 83.2 % at our average event rate of 6 Hz, given the Poissonian time structure of event triggers. At the end of 2010, a multi-buffer read-out scheme was deployed in the firmware of the DRS and trigger DAQ boards. This scheme uses a FIFO buffer between the front-end read-out and the VME transfer, capable of holding up to three events for the DRS part and up to four events for the trigger part. This allows us to re-enable the front-end electronics after the read-out time of $625 \text{ } \mu\text{s}$ and do the data transfer in the background. If another trigger happens during the VME transfer of the data, it

is written into the FIFO buffer without interfering with the data currently being transferred. The probability that three or more events occur during the 23.8 ms of the VME transfer is 5×10^{-4} at a trigger rate of 6 Hz and can be neglected. The dead time is therefore determined by the probability to have another event during the DRS read-out time of 623 μ s, which causes a dead-time of 0.4 % and leads to a live time of 99.6 %.

9.4.2 Other sources of dead time

The data read out through VME are transferred through Gigabit Ethernet to the back-end PC. Since the TCP/IP protocol is used, no packets are lost. Since a network switch with enough buffers and backplane bandwidth has been chosen, the probability of data collisions and necessary network retransfers is very small. Constant data rates of 98 MB/s over hours at the input port of the back-end computer have been observed, which are more than enough for the typical overall data rate of 20 MB/s. At the back-end PC, the event fragments are combined into full events. The average event size of a typical run is 2.4 MB, consisting of ADC waveforms from the trigger system (12 %), DRS waveforms for the LXe detector (30 %), the TC (8 %) and the DCH (45 %). The remaining 5 % covers special waveform channels for timing synchronisation and the slow control data such as high voltage currents, temperatures and so on. The full event is then compressed using the bzip2 algorithm, which results in a final event size of 0.9 MB (62 % compression ratio). A special parallel version of this algorithm is utilised that makes use of all eight CPU cores, so that the full 10 Hz event rate can be compressed in real time.

A typical MEG run takes about six minutes. This time was chosen to produce raw data files \sim 4.8 GB large. The time it takes to stop and start a new run is six seconds, leading to an additional dead time of 1.7 %. Although this time is small, there are plans to decrease it further by using multi-threaded control programs for starting and stopping runs.

10 Slow Control

Slow Control in this context means the monitoring and control of the physical quantities in the experiment, that vary “slowly” over time, such as temperatures, pressures or high voltages. It is typically sufficient to read these values once every few seconds and record them together with the data from the various detectors. This allows us to find during the offline analysis any correlation between changes in detector performance and variations in ambient quantities such as the temperature. Other quantities need adjustments during the data taking: the beam line needs to be tuned differently for calibrations and for data taking and the high voltages for

DCH and PMTs need to be changed from time to time in order to compensate for drifts in gain and efficiency.

Most experiments use many different systems for the slow control subsystems. This approach has the problem of combining the data from different subsystems into the central DAQ system. While this is viable in large collaborations, it is difficult in smaller experiments such as MEG to maintain many different subsystems. Thus the collaboration decided since the beginning to design a common slow control system for the whole MEG experiment. The challenge of such a design is to cover all requirements of the experiment with a single approach. A field bus was clearly necessary to connect spatially separated sensors, and local intelligence was required for control loops such as automatic high-voltage ramping and trip recovery. The system had to be flexible and easily extendable, very reliable and still operational with one part under maintenance.

10.1 MSCB system

The new system was named Midas Slow Control Bus (MSCB), indicating the close integration into the MIDAS DAQ system. It consists of a communication layer through an RS-485 field bus and through Ethernet, and uses 8-bit microcontrollers as communication end points and for local control loops. A crate-based high-voltage system was designed based on this concept. A general-purpose control box SCS-2000 was built which can be equipped with various daughter cards for up to 64 channels of various analog or digital inputs/outputs.

10.1.1 Communication

Physical quantities in MEG are measured at different locations. To collect them in a centralised system, an electronic connection between the sensors and the DAQ is necessary. Since analog connections are prone to electronic noise pickup, it is desirable to digitise the quantities close to the sensors and transfer the information digitally. For the MSCB system we choose a field bus based on the RS-485 standard. This bus allows a multi-drop topology, where many devices are connected in parallel to the same physical bus, which uses a 5 V TTL differential signal. The used baud rate is 115.2 kBaud/s, which is low enough so that bus lengths below 100 meters do not require termination. Measurement nodes were designed to use a 8-bit microcontroller which implements the MSCB protocol and connects to the serial bus via a RS-485 transceiver. The low-level bus protocol follows exactly the RS-232 standard (one start bit, one stop bit, no parity), so that the internal microcontroller UART (Universal Asynchronous Receiver/Transmitter) is used. On the other side the microcontroller communicates to various ADCs and DACs through different protocols such as SPI

(Serial Peripheral Interface bus) or I2C (Inter-Integrated Circuit), as shown in Fig. 82.

The used RS-485 bus transceivers (MAX1483) have a reduced bus load, so that up to 256 nodes can be connected to a single physical bus. To connect to the data acquisition PC, a so-called “submaster” was built to act as a master on that bus segment. On one side it connects to the MSCB bus, on the other side it encapsulates all MSCB telegrams into network packets and communicates with any PC through standard 10Base-T Ethernet. The Ethernet submaster concept has proved to be extremely versatile. MSCB nodes belonging together logically are connected to one MSCB bus segment and interfaced via a dedicated Ethernet submaster. The MEG experiment uses a total of 15 submasters, which are connected to the various subsystems and to the central experiment network switch. If one segment is down for maintenance, the other segments can still be operational, which was very useful during commissioning phases. The submasters can also be accessed remotely for debugging, development and safety, which can happen in parallel to the normal DAQ operation. In case the central DAQ system is down, a remote backup PC can temporarily take over the slow control monitoring, without the need to change any physical cable or reboot any system.

10.1.2 Communication protocol

The protocol for the MSCB system was designed for minimal overhead. A sequence of just a few bytes is enough to read or write a slow control value. Therefore it is possible to monitor many hundred nodes per second in spite of the relatively slow baud rate of 115.2 kBaund on the RS-485 bus. Each packet starts with a one-byte command, followed by optional arguments and a 8-bit CRC code. Usual commands are reading and writing of values, obtaining information about variables, rebooting a node or updating the node firmware through the RS-485 bus.

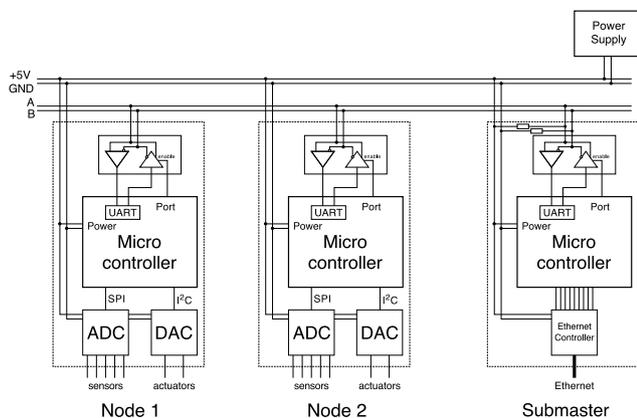


Fig. 82 Simplified communication schematics of the MSCB system

Since many nodes can be connected to the same bus, an addressing scheme has to be used to communicate with individual nodes. Each node contains a unique 16-bit address stored in the EEPROM of the node microcontroller. To address a certain node, an address packet containing the node address has to be sent over the bus. If a node recognises its own address, it switches to the “addressed mode” and interprets all consecutive commands until another node becomes addressed.

Figure 83 shows an example of the communication between a submaster and a node. A first packet addresses the node, while a second packet requests a certain slow control value from that node. The reply to this command contains a command code indicating an acknowledge, followed by the actual value and the CRC code.

The MSCB protocol is encapsulated by the Ethernet submaster into UDP packets. A software library on the PC side communicates with the submaster using these UDP packets. The library is written in C and available for common operating systems such as Linux, Windows and OSX. In addition, virtual instruments (VI) have been written to export the functionality to the LabVIEW system from National Instruments. All MSCB commands require an acknowledge. The commands writing a value get a one-byte acknowledge from the node after applying the new value, while the read command returns the actual value read. This scheme protects not only lost or corrupted packets on the RS-485 bus, but also lost UDP packets. If an acknowledge is not returned in a certain time window, the packet is re-sent several times with an increasing timeout. This protocol has been optimised for local area network connections, where timeouts of the order of a few hundred milliseconds are needed to achieve maximal throughput. Tests have shown that this is more effective than a TCP connection. A simple password authentication scheme ensures that none of the several experiments at PSI using the MSCB system talk to each other accidentally. The connectionless protocol ensures that the system will recover

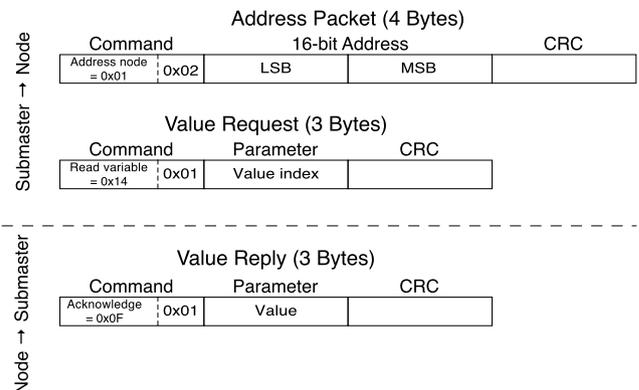


Fig. 83 Typical communication protocol example between a submaster and a node

automatically after a network failure or a reboot of a node or the submaster.

The MSCB system acquires data and controls values by exchanging certain values called network variables. Each MSCB node can implement up to 256 network variables. Each variable has a name, a type (such as byte or 4-byte floating point) and a physical unit. This variable definition resides on each node and can be queried using the MSCB protocol. Since all necessary information can be obtained directly from the MSCB nodes, it is possible to use general-purpose programs to connect to the MSCB system and display all network variables of a certain submaster without using a central configuration database. Adding new nodes or removing nodes is trivial. The MIDAS web interface (mhttpd) described in Sect. 9.3.2 contains a dedicated MSCB page which can be used as an user interface to all nodes of the MEG experiment, as in Fig. 84.

10.1.3 Local intelligence

Each MSCB node and submaster requires a microcontroller to implement the MSCB protocol and to realise local intelligence needed e.g. in control loops. The requirements for those microcontrollers are low power consumption, small size, integrated analog electronics such as ADCs and DACs and user reprogrammability. This led to the C8081F121 microcontroller [71] used in most MSCB nodes, which contains 128 kB of EEPROM, 8 kB of RAM, a 12-bit ADC and two 12-bit DACs and runs at a clock speed of 100 MHz.

The microcontroller is programmed in C using the Keil PK51 development kit [72]. The software is divided into a framework implementing the MSCB protocol and general housekeeping functions and the application-specific part. Both parts communicate via network variables described

above. The available program space of 128 kB is more than enough to accommodate the C runtime library including floating point functions, the MSCB framework consisting of about 3500 lines of C code, and still leaves space for several thousand lines for application-specific functionality such as regulation loops. The DCH pressure control system for example uses elaborate code including sensor linearisation, alarm functions and Proportional-Integral-Differential (PID) loops to keep the gas pressure stable to better than 10^{-5} bar, all implemented in about two thousand lines of C code.

One special feature of the MSCB system is that it can be reprogrammed over the network. The microcontrollers can write to their own code space from the user application. A small boot loader has been realised and is located in the upper memory of the code space, implementing a subset of the MSCB protocol. Using this boot loader, the main program space can be re-programmed. Firmware updates of all MSCB devices are therefore very easy. Since the protocol is encapsulated fully in the Ethernet communication layer, the update can even be sent through the Ethernet submaster from remote locations. A typical update contains 50 kB of code and requires ten seconds. After the update, the node reboots in about 0.5 seconds, reducing downtime to a minimum.

10.1.4 High-voltage system

The MEG experiment requires high-voltage supplies for the PMTs of the LXe and the TC, the APDs of the TC and the DCH. The MEG experiment requires an accuracy and long-term stability better than 0.1 V, which was hard to achieve with commercial systems at the time of the design of the experiment. Therefore a high-voltage supply based on the MSCB system was designed. Since the experiment has many channels which need about the same high voltage, it was easier to make a voltage regulator which is fed from a common external source. All 848 channels for the LXe detector are for example powered from a single high-current commercial high-voltage source, and the voltage is adjusted for each channel separately by a MSCB voltage regulator. The regulator uses a string of optocouplers, each capable of sustaining 240 V. A total of twelve optocouplers allows voltage regulation from zero to 2800 V. The LEDs of the optocouplers are controlled by a regulation loop which compares the output voltage with the demand value obtained from a DAC, as shown in Fig. 85.

The output voltage is measured by means of an ADC and a voltage divider consisting of two resistors with 41 MOhm and 33 kOhm. The current is determined by measuring the voltage drop across a resistor R_{sense} with a differential amplifier. Using 24-bit sigma-delta ADCs (AD7718 from Analog Devices), very good resolutions can be achieved. The

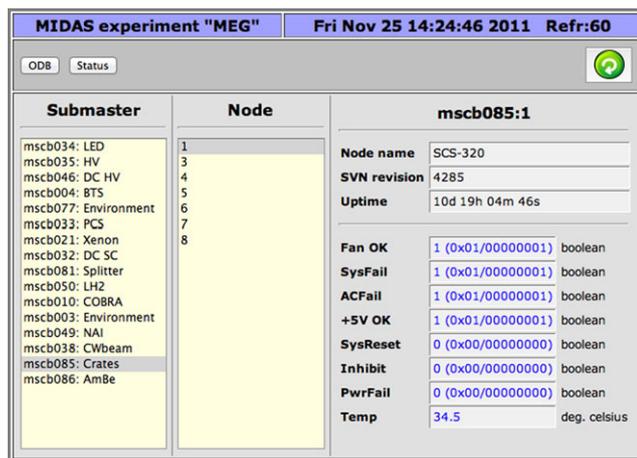


Fig. 84 Web interface page to the MSCB system. The left column shows all available submasters, the second column all nodes connected to the selected submaster, and the right column all network variables of the selected node

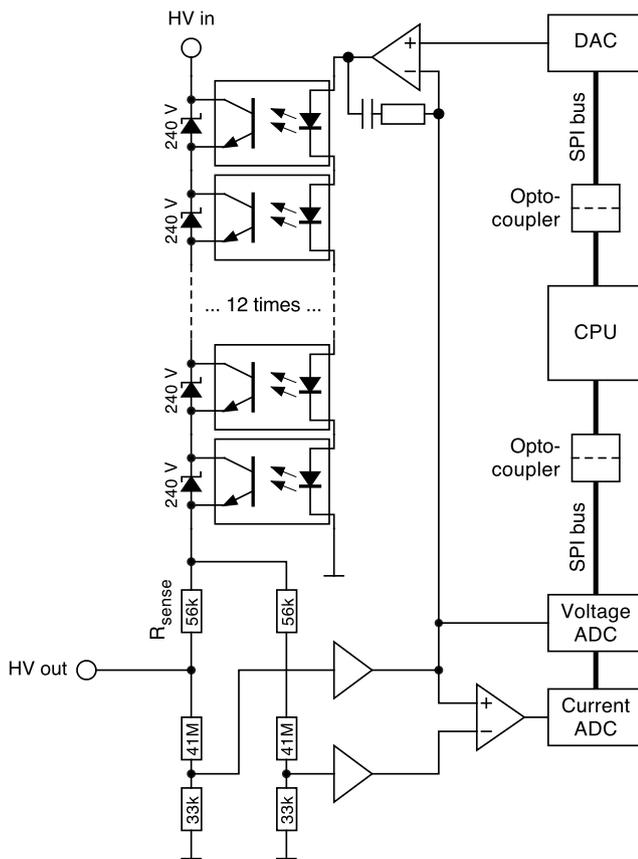


Fig. 85 Simplified schematics of the high-voltage regulator. The 240 V Zener diodes protect the optocouplers, so that no more than 240 V drop across the transistors of the optocouplers

ADCs and the DAC are controlled by the microcontroller through the SPI bus via optocouplers for noise immunity. The accuracy depends only on the resistors of the voltage divider and the current sense resistor, which can be calibrated using an external multimeter. The calibration is then stored in the nonvolatile memory of the microcontroller. Accuracies of 30 mV in voltage and a few nA in current have been achieved at 2000 V output voltage.

10.1.5 RS-232 and GPIB interface nodes

Although the MSCB system handles many sensors and actuators directly, some devices, e.g. COBRA, have their own control systems. These devices typically have a RS-232 or GPIB interface. To integrate these devices into the MEG slow control system, MSCB interface nodes were designed for the RS-232 and the GPIB busses, containing only the microcontroller with the RS485 interface chip (ADM2486 from Analog Devices), and a second interface chip for the RS-232 bus (MAX232 from Maxim) or the GPIB bus (directly connected to the CPU). A 5V DC-DC converter galvanically decouples the bus power from the local power. This allows the nodes to be powered from the bus without

the danger of introducing ground loops between different devices through their earthing system.

The interface nodes have two modes of operation. The so-called transparent mode encapsulates the device communication into MSCB data packages. The PC side contains routines which allow direct writing or reading from the RS-232 or GPIB bus. The advantage of this mode is that new devices can be connected and tested without modification of the MSCB firmware since this system works only as the transport layer. The disadvantage is that each device needs a dedicated driver on the PC side which understand the device-specific ASCII protocol.

The second mode encapsulates the device specific protocol in the firmware of the interface node. If connected to a power supply for example, the voltage and current of the power supply are polled periodically and put into MSCB network variables. These variables can then be read out without the knowledge of the device-specific protocol to query the voltage and current of the power supply. Parameters in attached devices can be changed in a similar way. Writing to the associated network variable triggers a callback routine in the microcontroller firmware which generates a write command to the device through the RS-232 or GPIB bus. Using this mode, all devices attached to the MSCB system can be talked to in the same way (namely network variables) independent of the specific device. This mode has therefore been chosen for all devices in the MEG experiment.

10.1.6 SCS-2000 control box

The MSCB-based SCS-2000 control box was designed for analog and digital inputs and outputs. The box contains an Eurocard-size main board with a microcontroller, a battery-buffered real-time clock and a 1 MB static RAM. User interaction happens through a LCD display and four push buttons. The SCS-2000 unit can either be inserted in a 3 HE crate or be housed in a stand-alone box with an external 24 V power supply. A single box can accommodate eight daughter cards, each with eight input/output channels of various kinds. Analog input cards typically host a 24-bit sigma-delta ADC with eight input channels. Other cards were developed for analog and digital input and output, including a capacitance meter for the liquid xenon level gauge or a 10 V LED pulser for the LXe detector calibration.

A generic firmware scans all plugged-in daughter cards and displays their parameters on the LCD display. Using the four buttons on the front panel, the user can scroll through all values and change values of those cards which support digital or analog output. This is very useful for controlling devices locally with an attached SCS-2000 without the need of using a laptop or other computers. The parameters are visible as network variables and are accessible through the MSCB bus.

10.2 DAQ integration

Most slow control parameters in the MEG experiment are acquired locally by SCS-2000 units or other MSCB nodes, but they need to be integrated into the full DAQ system. Measured values have to be written to disk together with the event data. Furthermore, the temporal changes of parameters have to be accessible by the shift operators in order to ensure the stability of the experiment. This is achieved by a dedicated slow control front-end program, together with the MIDAS online database and history system, which are described in Sect. 9.3.2.

10.2.1 MIDAS slow control front-end

A dedicated process runs on the back-end computer which connects to all slow control submasters through the Ethernet network. All controlled devices are partitioned into separate entities called equipments, such as high-voltage supplies, beam line magnets or ambient temperature sensors. Each equipment is read out periodically with a dedicated thread inside the front-end process for optimal read-out frequency. The readout threads copy all parameters into the ODB, where they become immediately available to other processes or to the shift operator via the mhttpd web interface. Control parameters are placed into the ODB as well, and linked to their devices via hot-links (see Sect. 9.3.2). If the user or a process changes a control parameter in the ODB, a callback routine in the slow control front-end is called, which in turn writes the new value to the associated slow control device. By dumping the ODB into a snapshot file, the current status of the whole experiment can be captured at a specific time, and be reconstructed later from that dump.

In addition, the slow control front-end generates events from the slow control parameters and injects them periodically via the buffer manager into the event stream written to disk for later analysis.

10.2.2 MIDAS custom pages

The whole mhttpd web interface is non-specific, meaning that the same pages are used for different experiments. This is useful for monitoring the general status, starting and stopping runs and changing ODB variables. There is however the need for experiment-specific pages such as the one shown in Fig. 86. Such pages condense information from one or several subsystems into a single display and provide the shift operator with a quick overview of a part of the experiment. To realise these so-called “MIDAS custom pages”, three techniques are implemented into the mhttpd web server:

- The HTML language is extended by an `<ODB [path] [flags]>` tag. When the mhttpd process parses the user-written HTML file for the custom page, the ODB tag is

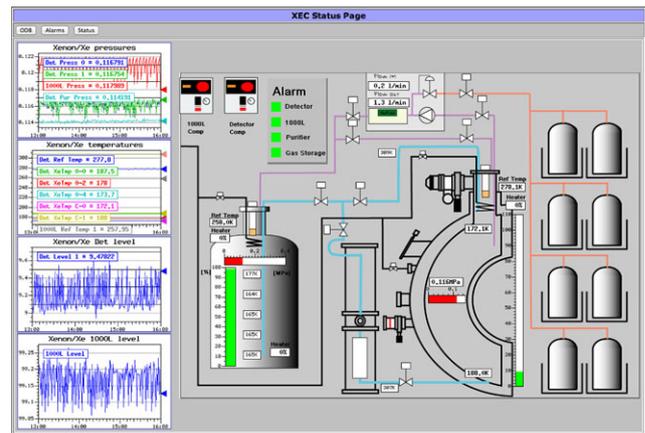


Fig. 86 Custom web page of the LXe detector cryogenic system. The plots on the left are taken from the history system

replaced by the actual value of the parameter stored in the ODB indicated by its path. Optional flags allow various number formats and the possibility to change the ODB value through the web interface.

- Dynamic GIF images can be created by mhttpd. They consist typically of a static background image combined with overlays of labels, bar graphs or coloured areas based on simple script statements. The image is created in memory using the GD library [73] and directly streamed to the browser for maximum efficiency. The right side of Fig. 86 is based on this technique.
- An Asynchronous JavaScript and XML (AJAX) library is available which contains functions for asynchronous communication with mhttpd. These functions can be used to change parts of a web page using the Document Object Model (DOM) in the background without the need to reload the complete page. Functions are implemented for reading and writing ODB values, to start/stop a run and to generate alarm and system messages.

The combination of the three methods has proved to be extremely powerful in designing web pages which almost look and behave like a dedicated program, but have the advantage that only a web browser is required to use them.

The usage of custom pages became a very efficient tool to control and monitor the complete experiment with only a web browser. This was a prerequisite for running remote shifts, which were introduced in 2011 in the MEG experiment.

11 Simulation

This section provides an overview of the packages used in the Monte Carlo (MC) simulation of the MEG experiment.

The MC simulation is based on the physics event simulation (*megeve*), the detector simulation (*gem*), the readout simulation and the event mixing (*megbartender*). These

packages make use directly or indirectly (via *megdb2cards*) of the MySQL [74] based Data Base (DB) (more details are in [68, 69]).

11.1 The simulation architecture

The simulation architecture was designed with a number of constraints from the time schedule, manpower and technical skills available at the start of the project. Therefore the physics and detector simulation packages were written in FORTRAN 77, while *megbartender* was written in C++.

The FORTRAN 77 packages were organised following a modern programming paradigm, i.e. to use an Object-Oriented (OO) approach organised in a custom-designed framework called *rem*. The basic unit manipulated by the framework is a Module, corresponding to an OO class.

Modules can be either sub-detector simulation sections or service tools such as graphics or input/output. The *megeve* and *gem* packages are steered through cards read by the FFREAD package.

11.2 Physics event simulation

In a Module, *megeve*, several event types of relevance for the experiment can be generated. A non-exhaustive list including exclusive and inclusive decays is (particles in brackets are not generated)

Signal	$\mu^+ \rightarrow e^+\gamma$
Signal e^+	$\mu^+ \rightarrow e^+(\gamma)$
Signal γ	$\mu^+ \rightarrow (e^+)\gamma$
Decay $e^+2\gamma$	$\mu^+ \rightarrow e^+\gamma\gamma$
Michel e^+	$\mu^+ \rightarrow e^+(\nu\bar{\nu})$ [60]
Radiative $e^+\gamma$	$\mu^+ \rightarrow e^+\gamma(\nu\bar{\nu})$ [60, 75]
Radiative γ	$\mu^+ \rightarrow (e^+)\gamma(\nu\bar{\nu})$ [60, 76]
AIF γ	$\mu^+ \rightarrow e^+(\nu\bar{\nu}), e^+e^- \rightarrow \gamma(\gamma)$
Cosmic μ^+	μ^+ from cosmic rays
Calibration 2γ	$\pi^-p \rightarrow \pi^0(n), \pi^0 \rightarrow \gamma\gamma$
Calibration Dalitz	$\pi^-p \rightarrow \pi^0(n), \pi^0 \rightarrow \gamma e^+e^-$
Boron Calibration	$^{11}\text{B}(p, 2\gamma)^{12}\text{C}$
Lithium Calibration	$^7\text{Li}(p, \gamma)^8\text{Be}$
Neutron Generator	$^{58}\text{Ni}(n, \gamma)^{59}\text{Ni}$
Mott Scattering	$e^+A \rightarrow e^+A$

When appropriate the generation includes first order radiative correction and μ^+ polarisation effects [60].

The μ^+ decay vertices may be forced to be in the target or each μ^+ is propagated through the beam line starting at the last quadrupole of Triplet II (see Fig. 2) 5647 mm upstream of the centre of COBRA to the decay vertex.

11.3 Detector simulation

The beam propagation, interaction in the target and propagation of the decay products in the detector are simulated with

a FORTRAN 77 MC program (*gem*) based on the GEANT3 package [77], which describes the sub-detector responses up to the read-out.

The program is heavily modularised using the *rem* framework. That allows easy addition of new Modules for implementing the various parts of the simulation. Settings dependent on the run-period are copied from the MySQL database into the FFREAD cards, and used in *gem*. The output is written in ZEBRA format [78] and the code for dealing with the ZEBRA package is automatically generated starting from a bank description in DZDOC format.

Each sub-detectors is described by a Modules: DCH, Longitudinal TC, Transverse TC and LXe detector.

Additional Modules, e.g. COBRA, gas volume, μ target, upstream and downstream elements outside COBRA, complete the description of the material distribution.

The GEANT3 routines simulate the energy release in the detectors, while the conversion into electrons/photons and their propagation up to the read-out devices are managed in *gem* user routines.

The LXe detector Module includes a very detailed simulation of the propagation of each scintillation photon in the detector volume, accounting for reflection, refraction and polarisation effects. The output is the photoelectron current on the PMT photocathodes. In the TC Modules (longitudinal and transverse sections) the photoelectron currents on PMT (APD) photocathodes are calculated with an analytical expression for propagation of the scintillation photons. In the DCH Module the GEANT3 energy hits are split into clusters along the track. Then the electrons are drifted toward the sense wires following the position-time relation obtained with GARFIELD [16], including diffusion effects. The output is wire currents. Figure 87 displays a simulated signal event.

The *gem* package stops short of simulating the read-out, which is performed by Bartender.

11.4 Bartender

Between the detector simulation and the reconstruction program an intermediate program, called Bartender, is required for the processing of MC data.

The Bartender serves different roles:

- Conversion of ZEBRA files into ROOT [79] files
- Readout simulation
- Event mixing

11.4.1 Format conversion

Bartender reads several ZEBRA files written by *gem* and converts them into ROOT format. Simulation-specific data such as kinematics of generated particles, true hit information etc. can be streamed to a “sim” Tree in separate ROOT files for further studies.

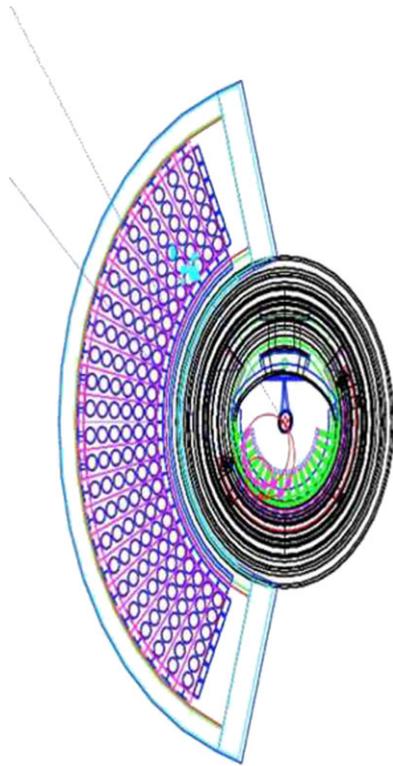


Fig. 87 A simulated event $\mu^+ \rightarrow e^+ \gamma$

11.4.2 Readout simulation

Bartender simulates the detector read-out electronics and makes waveforms. For example, for the LXe detector, waveforms are made by convolution of the single photoelectron response of a PMT with hit-time information of scintillation photons simulated by *gem*. PMT amplification, signal attenuation, saturation of the read-out electronics, noise etc. are taken into account. Simulated waveforms are encoded in the same manner as in the experiment and written to a “raw” Tree in ROOT files.

11.4.3 Event mixing

Bartender makes a mixture of several sub-events; rates of each event type are set with a configuration file. To study the combinatorial background events, sub-events are mixed with adjustable relative timing with respect to each other and with respect to the trigger; random and fixed timing can be selected. This allows simulation of many different pile-up configurations with a limited number of event samples simulated in the detector.

11.5 Database

Run-dependent parameters such as geometry, calibration and analysis constants are stored in a relational database

based on MySQL (DB). The DB is used by the DAQ front-end, analysis and simulation. For simulation, the dedicated program *megdb2cards* reads the DB and writes the FFREAD cards required by *gem*. Therefore the DB is used consistently by all packages.

11.6 Comparison with data

Figure 88 shows the normalised momentum spectrum of the e^+ measured in the spectrometer for data and MC events. Figures 89(a)–89(b) shows the hit distribution on the longitudinal Timing Counter for data and MC events. Figure 90 shows the inclusive γ -ray spectrum measured in the LXe detector for data and MC events. The MC events are separated into different sources.

12 Summary

The MEG detector described in this paper has been operational since 2008. In the period 2009–2011, it has collected $\sim 3.63 \times 10^{14}$ muon decays.

In the Table 9 the design, Monte Carlo and obtained resolutions are summarised. The obtained resolutions are, for some observables, worse than the design values. The differences are, at large extent, understood and the Monte Carlo resolutions match quite well the obtained ones.

The influence of the differences in resolutions on the analysis sensitivity is approximately evaluated firstly calculating the number of expected signal events for a given branching ratio (B)

$$N_{\text{sig}} \propto \epsilon_{\gamma} \times \epsilon_e, \quad (23)$$

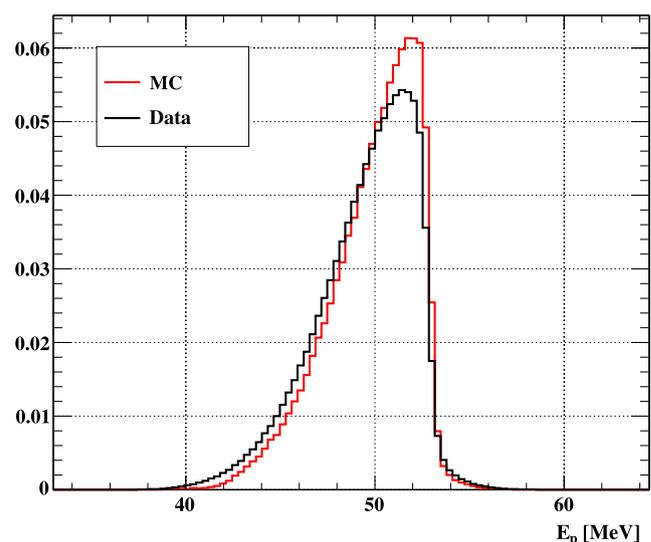
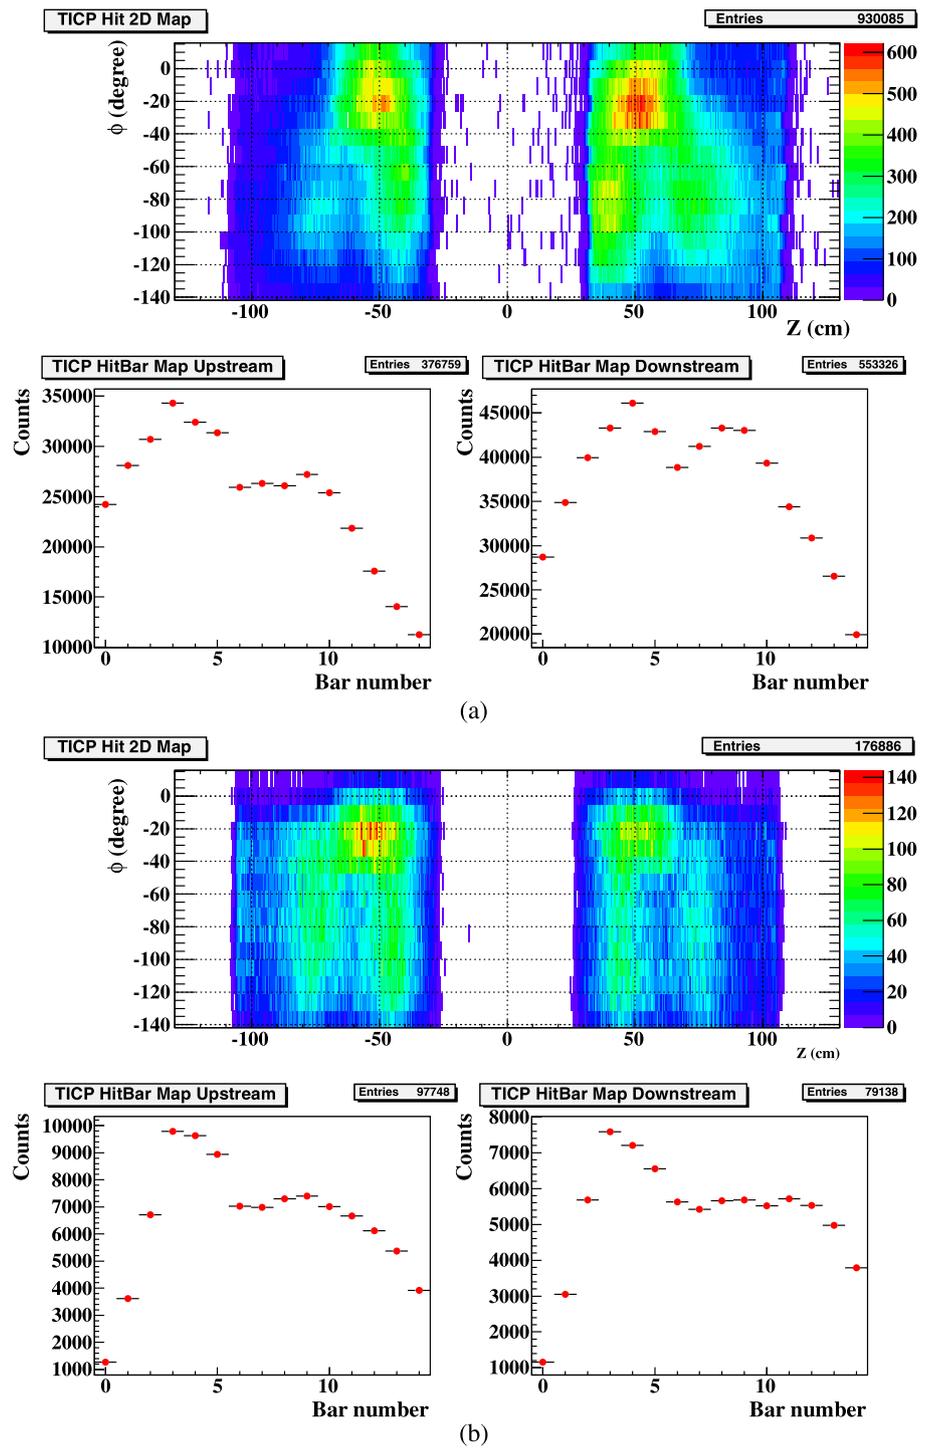


Fig. 88 Normalised momentum spectrum measured in the spectrometer for data and MC simulation

Fig. 89 Hit distribution in longitudinal TC for data (*above*) and MC simulation (*below*)



from Table 9, the decrease in N_{sig} is ~ 2 due to ϵ_e . Secondly, the number of accidental coincidences is

$$N_{acc} \propto \sigma_{E_\gamma}^2 \times \sigma_{E_e} \times \sigma_{t_{e\gamma}} \times \sigma_{\Theta_{e\gamma}}^2. \quad (24)$$

This number increases a factor of ~ 10 compared to the design value. Those effects combined implies that the sensitivity of the present MEG experiment is ~ 5 compared to the design value.

The results from the analysis of data collected up to 2010 are already published and set the most stringent limit to date on the existence of the $\mu^+ \rightarrow e^+ \gamma$ decay. Data collected in 2011 with statistics comparable to the 2009–2010 sample are being analysed.

The experiment run in 2012 collecting a data sample comparable in size to the 2011 sample.

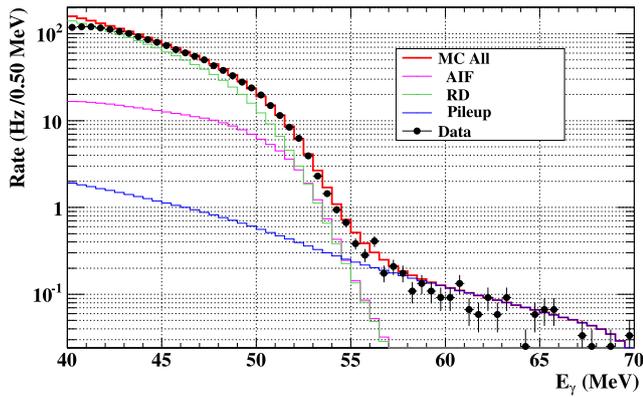


Fig. 90 Inclusive γ -ray spectrum in LXe. Data (dots) and MC (lines). MC spectrum is separated into contributions from AIF, RMD and pile-up from multiple γ -rays

Table 9 Resolutions expressed in Gaussian standard deviations (core or single side, when required) and efficiencies for all observables: comparison between expected and obtained

Variable	Design	Monte Carlo	Obtained
Resolutions			
Positron (e)			
σ_{E_e} (keV)	200	315	306
$\sigma_{\phi_e, \theta_e}$ (mrad)	$5(\phi_e), 5(\theta_e)$	$8(\phi_e), 9(\theta_e)$	$9(\phi_e), 9(\theta_e)$
σ_{z_e, y_e} (mm)	1.0	$2.9(z_e)/1.0(y_e)$	$2.4(z_e)/1.2(y_e)$
σ_{t_e} (ps)	50	65	102
Photon (γ)			
σ_{E_γ} (%)	1.2	1.2	1.7
σ_{t_γ} (ps)	43	69	67
$\sigma_{(u_\gamma, v_\gamma)}$ (mm)	4	5	5
σ_{w_γ} (mm)	5	6	5
Combined (e-γ)			
$\sigma_{t_{e\gamma}}$ (ps)	66	95	122
$\sigma_{\theta_{e\gamma}}$ (mrad)	11	16	17
Efficiencies			
ϵ_e (%)	90	40	40
ϵ_γ (%)	60	63	63
ϵ_{trg} (%)	100	99	99

Acknowledgements We are grateful for the support and co-operation provided by PSI as the host laboratory and to the technical and engineering staff of our institutes. This work is supported by DOE DEFG02-91ER40679 (USA), INFN (Italy) and MEXT KAKENHI 16081205 (Japan).

We thank Prof. David Stoker of the University of California, Irvine, for his careful proofreading of the manuscript and its improvements.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribu-

tion, and reproduction in any medium, provided the original author(s) and the source are credited.

References

1. Y. Hisamatsu, Search for the lepton flavour-violating decay $\mu \rightarrow e\gamma$ the MEG experiment probes GUT scale physics with MeV particles. *Eur. Phys. J. C* **52**(2), 477–485 (2007)
2. J. Adam et al., A limit for the $\mu^+ \rightarrow e^+\gamma$ decay from the MEG experiment. *Nucl. Phys. B* **834**, 1–12 (2010)
3. J. Adam et al., New limit on the lepton-flavour violating decay $\mu^+ \rightarrow e^+\gamma$. *Phys. Rev. Lett.* **107**, 171801 (2011)
4. M.L. Brooks et al., New limit for the lepton-family-number non-conserving decay $\mu^+ \rightarrow e^+\gamma$. *Phys. Rev. Lett.* **83**(8), 1521–1524 (1999)
5. R. Barbieri, L. Hall, A. Strumia, Violations of lepton flavour and CP in supersymmetric unified theories. *Nucl. Phys. B* **445**(2–3), 219–251 (1995)
6. J.F. Crawford et al., Production of low energy pions by 590 MeV protons on nuclei. *Helv. Phys. Acta* **53**, 497–505 (1980)
7. P.-R. Kettle, Optimal production target orientation for surface muon channels. SIN Rep., TM-31-20, October 1980
8. Hamamatsu Photonics, R7400/R9980-110: sales.hamamatsu.com/index.php?id=13189569
9. J.F. Crawford et al., Measurement of cross sections and asymmetry parameters for the production of charged pions from various nuclei by 585-MeV protons. *Phys. Rev. C* **22**, 1184–1196 (1980)
10. J. Adam et al., Measurement of muon polarization in the MEG experiment (2012, in preparation)
11. A. Baldini et al., Muon polarization in the MEG experiment. Predictions and measurements. MEG Technical Note TN068, March 2011
12. SHI 4K Cryocoolers SRDK-408D: www.shicryogenics.com/
13. A. Yamamoto, Y. Makida, Y. Doi, T. Kondo, K. Wada, S. Meguro, Development towards ultra-thin superconducting solenoid magnets for high energy particle detectors. *Nucl. Phys. B* **78**(1), 565–570 (1999)
14. A. Yamamoto et al., A thin superconducting solenoid magnet for particle physics. *IEEE Trans. Appl. Supercond.* **12**, 438–441 (2002)
15. M. Hildebrandt, The drift chamber system of the MEG experiment. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **623**(1), 111–113 (2010)
16. GARFIELD, CERN program Library Long Writeups W5050
17. D.F. Anderson et al., A simple vernier method for improving the accuracy of coordinate read-out in large wire chambers. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **224**(1–2), 315–317 (1984)
18. D. Green et al., Accurate 2 dimensional drift tube readout using time division and vernier pads. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **256**(2), 305–312 (1987)
19. J. Allison et al., The diamond shaped cathode pads of the OPAL muon barrel drift chambers. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **310**(1–2), 527–534 (1991)
20. S. Uno et al., Study of a drift chamber filled with a helium-ethane mixture. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **330**(1–2), 55–63 (1993)
21. H. Hirano et al., A high-resolution cylindrical drift chamber for the KEK B-factory. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **455**(2), 294–304 (2000)
22. ISEG EHQ 103M: www.iseg-hv.com/
23. ISEG EHS 8630p-305F: www.iseg-hv.com/

24. V. Blobel, Software alignment for tracking detectors. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **566**, 5–14 (2006)
25. S. Dussoni, M. De Gerone, F. Gatti, R. Valle, M. Rossella, R. Nardó, P.W. Cattaneo, The Timing Counter of the MEG experiment: design and commissioning. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **617**(1–3), 387–390 (2010)
26. R. Valle, Development of the Timing Counter for the MEG experiment. PhD thesis, University of Genoa (2006)
27. S. Dussoni, In search for lepton-flavour violation: the MEG experiment. PhD thesis, University of Genoa (2006)
28. M. De Gerone et al., Development and commissioning of the Timing Counter for the MEG experiment. IEEE Trans. Nucl. Sci. **59**(2), 379–388 (2011)
29. V.O. Altemose, Helium diffusion through glass. J. Appl. Phys. **32**(7), 1309 (1961)
30. B. Hartman et al., Gamma-ray large area space telescope (GLAST) large area telescope (LAT) ACD phototube helium sensitivity. LAT-TD-00720-D1 (2002)
31. Saint-Gobain: www.detectors.saint-gobain.com/Plastic-Scintillator.aspx
32. Hamamatsu Photonics, R5924: <http://sales.hamamatsu.com/en/products/electron-tube-division/detectors/photomultiplier-tubes/part-r5924-70.php>
33. M. Bonesini et al., Behaviour in high magnetic fields of fine-mesh photodetectors for fast time-of-flight detectors. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **567**, 200–204 (2006)
34. Saint-Gobain: www.detectors.saint-gobain.com/fibers.aspx
35. Hamamatsu Photonics, S8664-55: jp.hamamatsu.com/products/_virtual/sel_ssd_apd_s2/S8664-55index_en.html
36. M. De Gerone et al., Development and characterization of scintillating fiber-APD based detector. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **610**, 218–221 (2009)
37. M. De Gerone et al., The MEG timing counter calibration and performance. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **638**, 41–46 (2011)
38. Huber + Suhner Enviroflex_400: www.hubersuhner.com
39. J. Adam et al., Calibration and monitoring of the MEG experiment by a proton beam from a Cockcroft–Walton accelerator. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **641**, 19–32 (2011)
40. G.F. Knoll, *Radiation Detection and Measurement* (Wiley, New York, 1988), p. 329. Chapter 10.VII
41. A. Baldini et al., Absorption of scintillation light in a 100 l liquid xenon γ -ray detector and expected detector performance. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **545**(3), 753–764 (2005)
42. T. Haruyama, K. Kasami, H. Inoue, S. Mihara, Y. Matsubara, Development of a high-power coaxial pulse tube refrigerator for a liquid xenon calorimeter. AIP Conf. Proc. **710**(1), 1459–1466 (2004)
43. Hamamatsu Photonics, R9869: www.hamamatsu.com/
44. S. Mihara, MEG liquid xenon detector. J. Phys. Conf. Ser. **308**, 012009 (2011)
45. S. Mihara et al., Development of a method for liquid xenon purification using a cryogenic centrifugal pump. Cryogenics **46**(9), 688–693 (2006)
46. BarBer-Nichols Inc.: www.barber-nichols.com
47. T. Iwamoto et al., Development of a large volume zero boil-off liquid xenon storage system for muon rare decay experiment (MEG). Cryogenics **49**, 254 (2009)
48. A. Baldini et al., A radioactive point-source lattice for calibrating and monitoring the liquid xenon calorimeter of the MEG experiment. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **565**(2), 589–598 (2006)
49. A. Papa, Search for the lepton flavour violation in $\mu \rightarrow e\gamma$. The calibration methods for the MEG experiment. Edizioni ETS, University of Pisa, 2010
50. Eckert & Ziegler: www.ezag.com
51. ThermoFisher: www.thermofisher.com/global/en/home.asp
52. www.nds.iaea.org
53. High Voltage Engineering Europe: www.highvolteng.com
54. A. Papa, The CW accelerator and the MEG monitoring and calibration methods. Nuovo Cimento B **122**(6), 627–640 (2007)
55. Hamamatsu Photonics, S8664-1010: jp.hamamatsu.com/products/_virtual/sel_ssd_apd_s2/S8664-1010/index_en.html
56. L. Galli et al., An FPGA-based trigger system for the search of $\mu^+ \rightarrow e^+\gamma$ decay in the MEG experiment. J. Instrum. **8**, P01008 (2013)
57. Analog Devices, AD9218 ADC: www.analog.com
58. Xilinx: www.xilinx.com
59. Cypress Semiconductor Corporation: www.cypress.com/
60. Y. Kuno, Y. Okada, Muon decay and physics beyond the standard model. Rev. Mod. Phys. **73**(1), 151–202 (2001)
61. S. Ritt, R. Dinapoli, U. Hartmann, Application of the DRS chip for fast waveform digitizing. Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. **623**(1), 486–488 (2010)
62. B. Keil et al., The PSI VPC board—first applications of a common digital back-end for electron and proton beam instrumentation at PSI, in *Proceedings of DIPAC05*, Lyon, France (2005), p. 14
63. National Semiconductor, LMK03000, Precision Clock Conditioner: www.national.com/
64. Struck Innovative Systeme GmbH: www.struck.de/pcivme.htm
65. MIDAS: midas.psi.ch
66. M. Schneebeli, R. Sawada, 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)
67. ROME: midas.psi.ch/rome
68. R. Sawada, P.W. Cattaneo, MEG simulation and analysis software, in *IEEE Nuclear Science Symposium Conference Record*, Knoxville, TN, October 2010
69. P.W. Cattaneo, F. Cei, R. Sawada, M. Schneebeli, S. Yamada, The architecture of MEG simulation and analysis software. Eur. Phys. J. Plus **126**(7), 1–12 (2011)
70. ELOG: midas.psi.ch/elog
71. Silicon Labs: www.silabs.com
72. Keil PK51 Development Kit: www.keil.com/c51/pk51kit.asp
73. GD Graphics Library: www.boutell.com/gd/
74. MySQL: <http://www.mysql.com>
75. Y. Kuno, Y. Okada, Proposed $\mu \rightarrow e\gamma$ search with polarized muons. Phys. Rev. Lett. **77**(3), 434–437 (1996)
76. Y. Kuno, A. Maki, Y. Okada, Background suppression for $\mu \rightarrow e\gamma$ with polarized muons. Phys. Rev. D **55**(5), R2517–R2520 (1997)
77. GEANT, CERN program Library Long Writeups W5013
78. ZEBRA, CERN program Library Long Writeups Q100/Q101 (1995)
79. ROOT: <http://root.cern.ch>