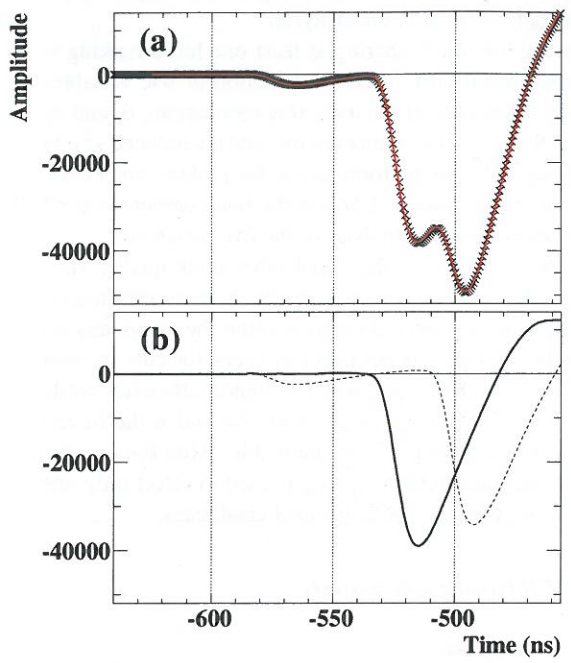

Figure 10 (a) Example of a LXe detector waveform for an event with three γ-rays (2.5, 40.1, and 36.1 MeV). The cross markers show the waveform (with the digital high-pass filter) summed over all PMTs with the coefficients defined in the text and the red line shows the fitted superposition of three template waveforms. (b) The unfolded main pulse (solid line) and the pile-up pulses (dashed).

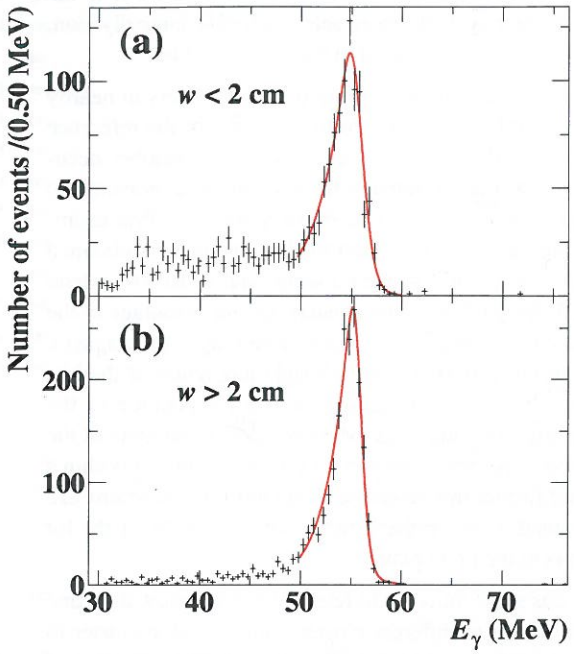

Figure 11 Energy response of the LXe detector to 54.9-MeV γ-rays in a restricted range of (u, v) for two groups of events with different w: (a) $0 < w < 2$ cm (event fraction 42%) and (b) $w > 2$ cm (58%).

3.2 Positron reconstruction

Editor's comments:
Section coordinator: Francesco R., Gordon L., Luca G., Matteo D.
Text: 3.0
Figure: 8.

3.2.1 DCH reconstruction

Editor's comments:
Section coordinator: Francesco

The reconstruction of tracks in the DCH is performed in four steps: hit reconstruction in each single cell, clustering of hits within the same chamber, track finding and track fit.

First of all, raw waveforms from anodes and cathodes are filtered in order to remove known noise contributions of fixed frequencies. Then, a hit is defined as a negative signal appearing in both waveforms collected at the two ends of the anode wire, with an amplitude of at least −5 mV below the baseline, which is estimated, along with its standard deviation σ, from the waveform itself in a region of about 625 ns before the trigger time. The hit time is taken from the anode waveform having the largest signal, as the time of the first sample below -3σ from the baseline.

The samples with amplitude below -2σ from the baseline and in a range of $[-24, +56]$ ns around the peak, are used for charge integration. The range is optimised to minimise the uncertainty produced by the electronic noise. A first estimate of the z coordinate, with a resolution of about 1 cm, is obtained from charge division in the anode wire, and it allows to determine the Vernier cycle the hit is within. If one or more of the cathode pad channels is known to be broken, the z coordinate from charge division is used and is assigned a 1 cm uncertainty. Otherwise, charge integration is performed on the cathode pad waveforms and the resulting charges are combined to refine the z measurement, exploiting the Vernier pattern. Given the charge asymmetries between the upper and lower sections of the inner and outer cathodes:

$$A_{in,out} = \frac{Q^{UP}_{in,out} - Q^{DOWN}_{in,out}}{Q^{UP}_{in,out} + Q^{DOWN}_{in,out}},$$

the position within the $\lambda = 5$ cm Vernier cycle is given by:

$$\delta_z = \arctan(A_{in}/A_{out}) \cdot \frac{\lambda}{2\pi}.$$

At this stage, a first estimate of the position of the hit in the (x, y) plane is given by the wire position.

Once they have been reconstructed, hits from nearby cells and with similar z are grouped into clusters, taking into ac-

count that the z-measurement can be shifted by λ if the wrong Vernier cycle has been selected from charge division. These clusters are then used to build track seeds.

A seed is defined as a group of three clusters in nearby chambers, at large distance from the center of the reference frame ($R = \sqrt{x^2 + y^2} > 24$ cm, where the chamber occupancy is lower and limited to tracks with large momentum) and satisfying proximity criteria in R and z. A first estimate of the track curvature and total momentum is obtained from the coordinates of the hit wires, and is used to extend the track and look for other clusters, taking advantage of the adiabatic invariant P_T^2/B_z for slowly varying axial magnetic fields. In this procedure, the left/right ambiguity of the hits can be solved in most cases, based on the position of the nearby hits along the track. Moreover, a first estimate of the track time (and hence the precise position of the hit within a cell) and further improvement of the left/right solutions can be obtained by minimising the χ^2 of a circle fit of the hit positions in the (x, y) plane.

At this stage, in order to retain high efficiency, the same hit can belong to different clusters and the same cluster to different track candidates, and it can result in duplicated tracks. Only after the track fit, when the best information about the track is available, multiple assignments are handled, as described later, to define independent tracks.

As soon as a track candidate is available, a precise estimate of the (x, y) position of the hits associated to that candidate can be extracted. It is obtained by defining the drift time as the difference between the hit and track times, and using tables to extract the position as a function of the drift time and the track angle with respect to the chamber plane. These tables are derived from simulations based on the GARFIELD software [21], and the track angle allows to take into account the deviation of the isochrone curves from circles, due to the square shape of the cells. The reconstructed (x, y) position is continuously updated during the tracking process, as the track information improves.

A track fit is finally performed with the Kalman filter technique [22,23]. We make use of the GEANE software [24] for the propagation of the track and the estimate of the error matrix within the filter. Exploiting the results of a first track fit, hits not included in the track candidate can be added while hits which are inconsistent with the fitted track can be removed. The track is also propagated to the TC in order to be matched with the hits in the bars (see Sect. 3.2.7 for details). The time of the matched TC hit is used to give a more accurate estimate of the drift times, and then the fit is performed again with this refined information. Finally, the track is propagated back to the target, and the intersection with the target foil is taken as the decay vertex (x_e, y_e, z_e), while the extrapolated direction (ϕ_e, θ_e) is taken as the positron direction at the decay vertex. The error matrix of the track parameters at the decay vertex on the target is also computed, to be used in the analysis.

Among the tracks sharing at least one hit, a ranking is performed based on a linear combination of five variables denoting the quality of the track (the momentum, θ_e and ϕ_e errors at the target, the number of hits and the reduced χ^2). In order to optimise the performance of the ranking procedure, the linear combination is taken as the first component from a principal component analysis of the five variables.

A series of cuts on these and other track quality variables are then applied to remove badly reconstructed tracks. Since the analysis will take into account the errors associated to the track parameters event by event, the cuts are kept as loose as possible to preserve the highest efficiency while removing badly reconstructed tracks, for which the fit and the associated errors could be unreliable. After the cuts are applied, the track quality ranking is used to select only one track among the survived duplicated candidates.

3.2.2 DCH missing turn recovery

Editor's comments:
Section coordinator: Gordon

A positron can traverse the DCH multiple times before it exits the spectrometer. An individual crossing of the DCH is referred to as a positron 'turn'. An intermediate merging step in the Kalman fit procedure described previously attempts to identify multi-turn positrons by combining and refitting individually reconstructed turns into a multi-turn track. However, it is possible that not all turns of multi-turn positron are correctly reconstructed or merged into a multi-turn track. If this involves the first turn, i.e. the turn closest to the muon stopping target, this will lead to an incorrect determination of the muon decay point and time as well as an incorrect determination of the positron momentum and direction at the muon decay point, and therefore a loss of signal efficiency.

After the track reconstruction has finished, a Missing First Turn (MFT) recovery algorithm, developed and incorporated in the DCH reconstruction software expressly for this analysis, is used to identify and refit positron tracks with an MFT. First, for each track in an event, the algorithm identifies all hits that could potentially be part of an MFT, based on their z-coordinates and DCH module identifiers with respect to the positron track. The vertex state vector of the track is propagated backward to the point of closest approach with each potential MFT hit, and the hit selection is refined based on the r and z residuals between the potential MFT hits and their propagated state vector positions. Potential MFT candidates are subsequently selected based on the distribution properties and number of remaining selected MFT hits. A new MFT track is reconstructed using the Kalman filter technique based on the selected MFT hits and corresponding propagated state vectors. Finally, the original

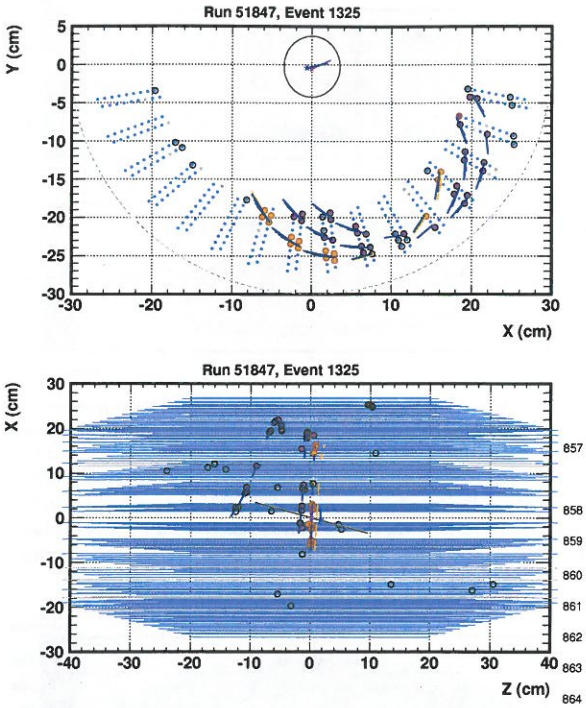

Figure 12 Example of a triple-turn positron in a 2009 event. This positron was originally reconstructed as a double-turn track (magenta hits), but was found to have an MFT (brown hits) by the MFT recovery algorithm. The original double-turn positron vector at the target and the refitted triple-turn positron vector at the target are shown as magenta and blue arrows, respectively.

positron and MFT tracks are combined and refitted using the Kalman filter technique, followed by a recalculation of the track quality ranking and the positron variables and their uncertainties at the target. An example of a multi-turn positron with a recovered MFT is shown in Fig. 12.

The improvement of the overall track reconstruction efficiency due to the use of the MFT recovery algorithm, defined by the ratio of the number of reconstructed Michel positrons with a recovered MFT and the total number of reconstructed Michel positrons, can be measured using data and is shown as a function of E_e and θ_e in Fig. 13. As can be seen from the left figure, the improvement of the track reconstruction efficiency at the signal energy due to the use of the MFT recovery algorithm, averaged over all angles, is ~ 4%. The efficiency improvement decreases with energy because the nominal track reconstruction is more efficient at higher energy. The right figure shows that the efficiency improvement is maximal for positrons emitted perpendicular to the beam direction, as expected since these positrons are more likely to have multiple turns.

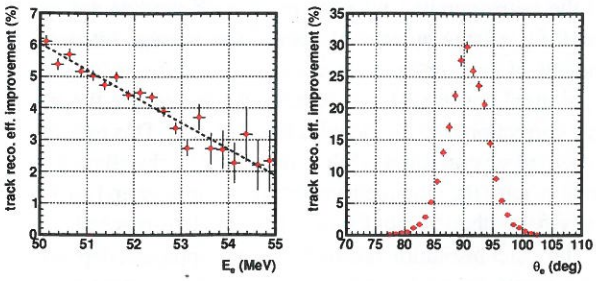

Figure 13 The improvement of the overall track reconstruction efficiency due to the use of the MFT recovery algorithm as a function of E_e (left) and θ_e (right).

3.2.3 DCH alignment

Editor's comments:
Section coordinator: Gordon, Luca

Accurate positron track reconstruction requires precise knowledge of the location and orientation of the anode wires and cathode pads in the DCH system. This is achieved by an alignment procedure that consists of two parts: an optical survey alignment based on reference markers, and a software alignment based on reconstructed tracks.

Each DCH module is equipped with cross hair marks on the upstream and downstream sides of the module. Each module is fastened between carbon blocks on the upstream and downstream sides of the support structure, where each block contains an alignment pin with an optically recognisable centre. Before the start of each data-taking period an optical survey of the cross hairs and pins was performed using a theodolite. The optical survey technique was improved in 2011 by adding corner cube reflectors next to the cross hairs, which were used in conjunction with a laser tracker system. The resolution of the laser method is ~ 0.2 mm for each coordinate.

Two independent software alignment methods are used to cross-check and further improve the alignment precision of the DCH system. The first method is based on the Millepede algorithm [25] using reconstructed cosmic rays. During COBRA shutdown periods, cosmic rays were triggered using dedicated scintillation counters located around the magnet cryostat. The alignment procedure utilises the reconstructed hit positions on the DCH modules to minimise the residuals with respect to straight tracks, according to the Millepede algorithm. The global alignment parameters, three positional and three rotational degrees of freedom per module, are determined with an accuracy better than 150 μm for each coordinate.

The second method is based on reconstructed Michel and Mott positrons and aims to improve the relative radial and longitudinal alignment of the DCH modules. The radial and longitudinal differences between the track position

and the corresponding hit position at each module are recorded for a large number of tracks. The average hit-track residuals of each module are used to correct the radial and longitudinal position of the modules, while keeping the average correction over all modules equal to zero. This process is repeated several times while refitting the tracks after each iteration, until the alignment corrections converge and an accuracy better than 50 μm for each coordinate is reached.

The exact resolution reached by either approach depends on the resolution of the optical survey used as starting position. For a low resolution survey, the Millepede method obtains a better resolution, while the Michel-Mott method obtains a better resolution for a high resolution survey. Based on these points, the Millepede method was adopted for the years 2009-2011 and the Michel-Mott method was used for the years 2012-2013 for which the novel optical survey data are available; in 2011, the first year with the novel optical survey data, the resulting resolution of both approaches is comparable.

3.2.4 Target alignment

Editor's comments:
Section coordinator: Gordon

Precise knowledge of the location and orientation of the target relative to the DCH system is crucial for an accurate determination of the muon decay vertex and the positron direction at the target. The positions of the cross marks on the target foil (see Fig. 2) are surveyed every year using a theodolite, with an estimated accuracy of $\pm(0.5, 0.5, 1.5)$ mm in the (x, y, z) directions. The cross mark measurements and preliminary target hole reconstruction results (see next paragraph) indicate that the target foil developed a gradual non-planar deformation over time. This is confirmed by measurements of the target foil shape performed with a high-precision FARO 3D laser scanner [26] at the end of 2013, as shown in the top panel of Fig. 14. In the track reconstruction, in the propagation of the track to the target, a plane fit to the cross mark measurements is used as a first approximation of the target location. This is improved in a second step in which the non-planarity of the target foil is approximated by fitting the cross mark measurements with a paraboloid $z_t - z_0 = c_x(x_t - x_0)^2 + c_y(y_t - y_0)^2$, where (x_t, y_t, z_t) is the local target reference frame in which x_t (y_t) is the coordinate along the semi-major (semi-minor) axis of the target, and z_t is the coordinate perpendicular to the target plane. The fit parameters in this function are (x_0, y_0, z_0) for the position of the paraboloid maximum, and c_x and c_y for the paraboloid curvatures in the x_t and y_t directions. The paraboloidal fit results of 2013, as shown in the bottom panel of Fig. 14, exhibit the largest non-planarity of all paraboloidal fits. In this fit $c_y = -0.03$ cm^{-1}, which corresponds to a focal length of ~8 cm for the semi-minor axis of the target in 2013.

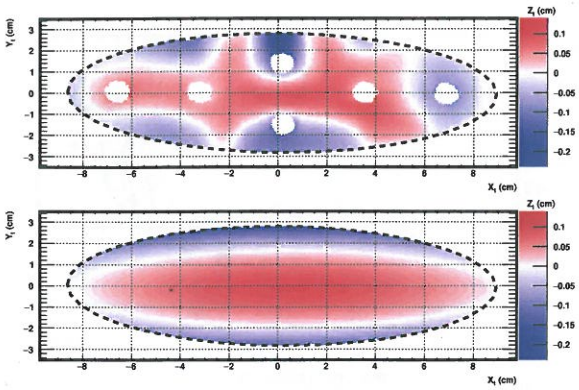

Figure 14 *Top*: FARO scan measurements of the target foil shape in the local target reference frame, in which x_t (y_t) is the coordinate along the semi-major (semi-minor) axis of the target, and z_t as indicated by the colour axis is the coordinate perpendicular to the target plane. *Bottom*: result of the paraboloidal fit of the 2013 cross mark measurements. For comparison, the range of the z_t-axis is the same in both plots and corresponds to the z_t-range of the FARO scan, while the z_t-range of the paraboloidal fit inside the target ellipse is $[-0.17, 0.10]$ cm.

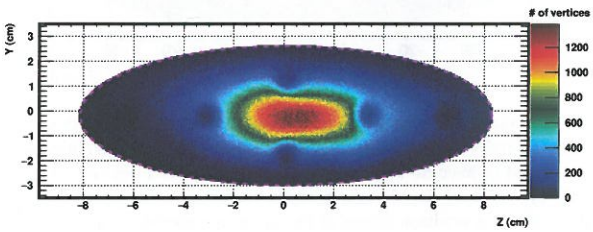

Figure 15 Vertex distribution in 2012 projected on the $z-y$ plane. The four central target holes can be seen as dips around the beam spot in the centre of the plot.

The fitted target position is cross-checked by imaging the holes in the target foil (see Fig. 2) using reconstructed Michel positrons. The target holes appear as dips in projections of the vertex distribution, as shown in Fig. 15. The alignment of the paraboloid in the z_t direction, the coordinate that is relevant in the analysis, is checked by determining the reconstructed y_t position of the target holes as a function of the positron angle ϕ_e. Ideally the target hole positions should be independent of the track direction, while a misalignment in z_t will induce an offset in y_t which to first order is proportional to $\tan(\phi_e)$. For example, Fig. 16 shows the reconstructed y_t position of the left-central target hole in 2011; the fit indicates a z_t misalignment of 1 mm towards the LXe detector. By imaging all four central holes, the systematic uncertainty of the z_t position of the target is estimated at ~0.3 mm.

3.2.5 DCH performance

Editor's comments:
Section coordinator: Francesco Renga

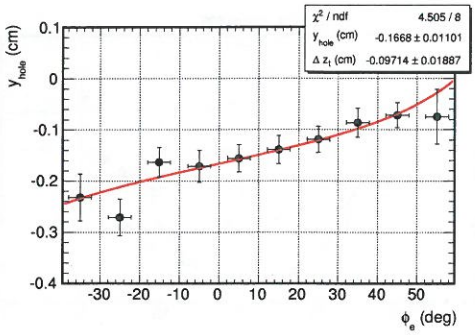

Figure 16 The reconstructed y_t position of the left-central target hole in 2011. The fit indicates that the true hole position is shifted by 1 mm in the negative z_t direction (i.e. towards the LXe detector) with respect to its position according to the fitted optical survey.

We developed a series of methods to extract from data an estimate of the resolution functions, defined for a generic observable p as the distribution of the residuals, $p - p_{true}$.

A complete overview of the performances of the spectrometer can be found in [7], where the methods used to evaluate them are also described in detail. Two methods are used to extract the resolution functions for the positron parameters. The energy resolution function, including the absolute energy scale, is extracted with good accuracy from a fit to the energy spectrum of positrons from Michel decays. A core resolution of about 330 keV is found, with a ~ 18% tail component of about 1.1 MeV width. The resolution functions for the positron angles and production vertex are extracted exploiting tracks that make two turns inside the spectrometer. The two turns are treated as independent tracks and extrapolated to a prolongation of the target plane in between the two. The resulting position and direction of the two turns are compared to extract the vertex and angle resolutions. The same method allows to study the correlations among the variables and to cross-check the energy resolution. However, since the two-turn tracks are a biased sample with respect to the whole dataset, substantial MC-based corrections are needed. Moreover, no information can be extracted about a possible reconstruction bias, which needs to be estimated from the detector alignment procedures described later in this paper. After applying the corrections, resolutions are found of about: 9.4 mrad on θ_e; 8.4 mrad on ϕ_e; 1.1 mm and 2.5 mm on y_e and z_e at the target, respectively.

On the other hand, in order to maximise the sensitivity of the likelihood analysis for the search for $\mu^+ \to e^+\gamma$ (the likelihood analysis in detail will be discussed in Sec. 4.4) instead of using these average resolutions we use the per-event estimate of the uncertainties, as provided by the track fit. The reliability of these uncertainties is verified by estimating on data the distribution of the *pulls*:

$$\text{pull} = \frac{p - p_{\text{true}}}{\sigma_p}$$

where σ_p is the uncertainty provided by the track fit. In case of Gaussian errors, and if the uncertainty is correctly estimated and there is no bias in the reconstruction of p, the pulls are expected to be distributed according to a Normal distribution. Anyway, non-Gaussian behaviours of the interaction of positrons with thin layers of material (e.g. Coulomb multiple scattering at large angles) lead to expected deviations from the Normal distribution, which need to be taken into account in the likelihood analysis. Practically, the pull distribution is used in the likelihood analysis instead of the residual distribution. Additionally, correlations among variables are also described in the likelihood as correlations among pulls. In this case, anyway, we preferred to use analytic formulas instead of the information coming from the fit matrix, thanks to our solid understanding of per-event correlations based on geometrical models.

3.2.6 TC reconstruction

Editor's comments:
Section coordinator: Matteo

The signal from each Timing Counter PMT is processed with a Double Threshold Discriminator (DTD) to extract the arrival time of the scintillating photons minimising the Time Walk effect. A TC hit is formed when both PMTs have signals above the higher DTD threshold. The times t_e^{in} and t_e^{out}, measured by the inner and outer PMT respectively, are extracted by a template fit to a NIM waveform fired at the lower DTD threshold and digitised by a DRS.

Considering a positron impinging on a TC bar at time t_e^{TC}, the relations with the measured PMT times are:

$$t_e^{in} = t_e^{TC} + b_{in} + TW_{in} + \frac{\frac{L}{2} + z_{TC}}{v_{\text{eff}}}$$
$$t_e^{out} = t_e^{TC} + b_{out} + TW_{out} + \frac{\frac{L}{2} - z_{TC}}{v_{\text{eff}}} \qquad (1)$$

where $b_{in,out}$ are offsets, $TW_{in,out}$ 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 taken in the middle of the bar. Adding the two parts of Eq. 1 the result is:

$$t_e^{TC} = \frac{t_e^{in} + t_e^{out}}{2} - \frac{b_{in} + b_{out}}{2} - \frac{TW_{in} + TW_{out}}{2} - \frac{L}{2v_{\text{eff}}}.$$

The time resolution is $\sigma_{t_e^{TC}} \approx 60$ ps.

Subtracting the two parts of Eq. 1 the coordinate of the impact point along the bar is given by:

$$z_{TC} = \frac{v_{\text{eff}}}{2}\left((t_e^{in} - t_e^{out}) - (b_{in} - b_{out}) - (TW_{in} - TW_{out})\right).$$

The position resolution is $\sigma_{z_{TC}} \approx 1.0$ cm.

The time (positions) TC resolution is determined using tracks hitting multiple bars from the distribution of the time (position) difference between hits on neighbouring bars corrected for the path length.

The radial and azimuthal coordinates are available from the bar position.

The TC, therefore, provides the required information to reconstruct all positron variables required to match a DCH track (see Sect, 3.2.7) and recover the muon decay time by extrapolating the t_e^{TC} back to the target.

3.2.7 DCH-TC matching

Editor's comments:
Section coordinator: Francesco

The matching of DCH tracks with hits in the TC is performed as an intermediate step in the track fit procedure, in order to exploit the information from the TC in the track reconstruction.

After being reconstructed within the DCH system, a track is propagated to the entrance of the TC. From here, the track is propagated to the first bar volume it encounters (*reference bar*). If no bar volume is crossed, the procedure is repeated with an extended bar volume. Then, for each TC hit within ± 5 bars from the reference one, the track is propagated to the corresponding bar volume and the hit is matched with the track according to the following matching categories:

1. the TC hit comes from the reference bar, and the difference in z between the track and the hit (Δz_{TC}) is less than 12 cm (the track position being defined as the entrance point of the track in the bar volume).
2. the TC hit comes from another bar whose extended volume is also crossed by the track, and $\Delta z_{TC} < 12$ cm (the track position being defined as the entrance point of the track in the extended bar volume).
3. the TC hit comes from a bar whose extended volume is not crossed by the track, but the distance of closest approach of the track to the bar axis is less than 5 cm and $\Delta z_{TC} < 12$ cm (the track position being defined as the point of closest approach of the track to the bar axis).

Among all successful matching candidates, those with the lowest category are chosen. Among them, the one with the lowest Δz_{TC} is finally used.

The time of the matched TC hit is assigned to the track, which is back-propagated to the chambers in order to correct accordingly the drift time of the hits, taking the track length contribution into account. During the back propagation the Kalman filter procedure is also applied to get the best estimate of the track parameters at the target, including the time t_e.

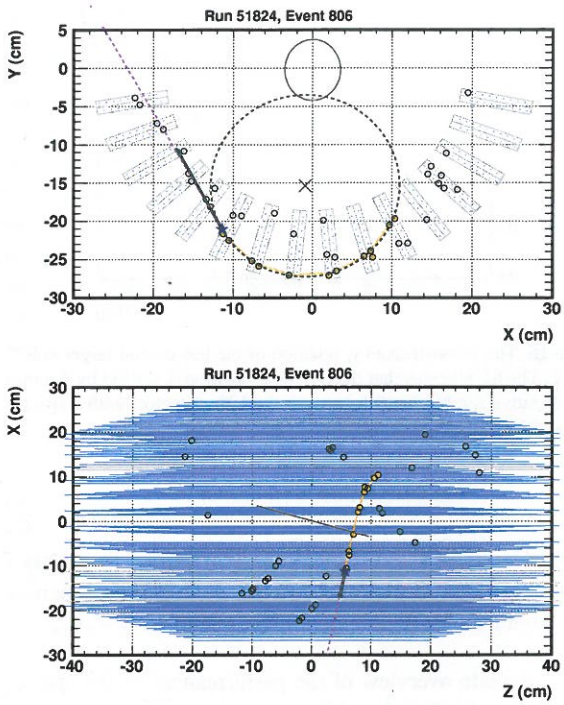

Figure 17 Example of a reconstructed positron AIF candidate in a 2009 event due to downstream muon decay. The reconstructed AIF vertex is indicated as a blue star, the AIF direction is indicated as a green arrow, and the vector connecting the AIF vertex and the γ-ray conversion vertex in the LXe is indicated as a magenta dashed line. Note that the green arrow and the magenta line nearly overlap, as expected for a true AIF event.

3.2.8 Positron AIF reconstruction

Editor's comments:
Section coordinator: Gordon

The γ-ray background around the signal energy is dominated by positron Annihilation-In-Flight (AIF) in the detector[4]. If the positron crosses part of the DCH before it annihilates, it can leave a trace of hits which are correlated with the subsequent γ-ray signal. A pattern recognition algorithm has been developed that can identify these types of positron AIF events. Since positron AIF contributes to the accidental background, this algorithm can help to distinguish accidental background events from signal and radiative muon decay events. The algorithm is summarised in the following.

The procedure starts by building positron AIF seeds from all reconstructed clusters. An AIF seed is defined as a set of clusters on adjacent DCH modules which satisfy a number of minimum proximity criteria. A positron AIF candidate (e^+_{AIF}) is reconstructed from each seed by performing a circle fit based on the xy-coordinates of all clusters in the seed. The

[4] See Sect. 4.3.1 for more details.

circle fit is improved by considering the individual hits in all clusters. The xy-coordinates of hits in multi-hit clusters are refined and left/right solutions based on the initial circle fit are determined by taking into account the time information of the individual hits, which also results in an estimate of the AIF time. The xy-coordinates of the AIF vertex are determined by the intersection point of the circle fit with the first DCH cathode plane after the last cluster hit. If the circle fit does not cross the next DCH cathode plane, the intersection point of the circle fit with the support structure of the next DCH module or the inner wall of COBRA is used. The z-coordinate of the AIF vertex is calculated by extrapolating a quadratic polynomial fit of the xz-positions of the last three clusters in the AIF candidate to the x-coordinate of the AIF vertex. The AIF candidate direction is taken as the direction tangent to both the circle fit and the quadratic polynomial fit at the AIF vertex. Figure 17 shows an example of a reconstructed AIF candidate.

3.3 Combined reconstruction

This section deals with variables requiring signals both in the spectrometer and in the γ-ray detector.

3.3.1 Relative γ-ray–positron angle

Editor's comments:
Section coordinator: Ryu, Yusuke
Text: 0.5
Figure: 1.

As the LXe detector is not capable of reconstructing the direction of the incoming γ-rays, this direction is determined connecting the reconstructed first conversion vertex of the γ-ray in the LXe detector and the reconstructed decay vertex on target: it is defined through its azimuthal and polar angles $(\phi_\gamma, \theta_\gamma)$.

The angular distance between the positron and γ-ray, defined as the relative angle between the γ-ray momentum vector and the opposite positron momentum vector at the reconstructed vertex decay, is treated separately in the azimuthal and the polar angle:

$$\theta_{e\gamma} = (\pi - \theta_e) - \theta_\gamma,$$
$$\phi_{e\gamma} = (\pi + \phi_e) - \phi_\gamma,$$

taking into account the correlation between them.

As there are no physics channel for measuring the relative angle resolution directly, it is obtained combining 1) the position resolution of the γ-ray detector and 2) the position and angular resolutions of the positron spectrometer. There are correlations among the errors of the positron observables because the observables are reconstructed extrapolating the positron tracks to the target. Due to the correlations, the relative angle resolutions are not the quadratic sum of the γ-ray and positron angular resolutions. The $\theta_{e\gamma}$ resolution is evaluated to be (15.0 – 16.2) mrad depending on the year of data taking by taking into account the correlation between z_e and $\theta_{e\gamma}$. Since the true positron momentum and $\theta_{e\gamma}$ of the $\mu^+ \to e^+\gamma$ signal are known, ϕ_e and y_e can be corrected using the reconstructed energy of the positron and $\theta_{e\gamma}$. The $\phi_{e\gamma}$ resolution after correcting those correlations is (8.9 – 9.0) mrad depending on the year.

There are no direct calibration sources for the relative angles and, thus, the accuracy basically relies on the detector alignment. The systematic uncertainty of the positron emission angle comes from the accuracy of the relative alignment among the magnetic field, the DCH modules, and the target (see Sect. 3.2.3 and 3.2.4 for the alignment methods and the uncertainties). In particular, the relative position of the target and the deformation from the plane directly affects the emission angle measurement and are found to be one of the dominant sources for the systematic uncertainty on the relative angles. For instance, an uncertainty of 0.5 mm in x corresponds to ~4 mrad in $\phi_{e\gamma}$. The target position and shape parameters are treated as nuisance parameters in the likelihood analysis to take the uncertainty into account.

The γ-ray direction relies on the relative alignment between the spectrometer and the LXe detector. While they are aligned separately, the relative alignment is cross-checked by directly measuring two types of physics processes that leave hits in both detectors:

- Positron AIF events,
- Cosmic rays measured without the COBRA magnetic field.

Each of the two measurements independently provides information of the displacement in the longitudinal direction (δz) relative to the survey with a precision better than 1.0 mm, whereas they are subject to large systematic uncertainties on the other degrees of freedom (shift in x and y, rotations, and shrink or expansion) due to the non-uniform distribution of those events and different shower development. The two results in z (δz = 2.1 mm for AIF and δz = 1.8 mm for cosmic rays) agree well, resulting in an average $\delta z = 2.0 \pm 0.4$ mm. This displacement suggests a systematic error of the detector alignment and survey. We take the average between the survey and the average of the measurements as the longitudinal displacement, namely, we apply a δz = 1.0 mm shift of the LXe detector along the z-direction, and, assuming a flat probability distribution of the true value of the displacement in the range [0 mm, 2 mm], assign the uncertainty of $2\,\text{mm}/\sqrt{12} = 0.58$ mm. This uncertainty corresponds to 0.85 mrad in the centre value of $\theta_{e\gamma}$ (converted by the radial position of the LXe detector $r \sim 69$ cm). On the other hand, there is no measurement sensitive to a few mm displacement in any other degree of freedoms and no measure-

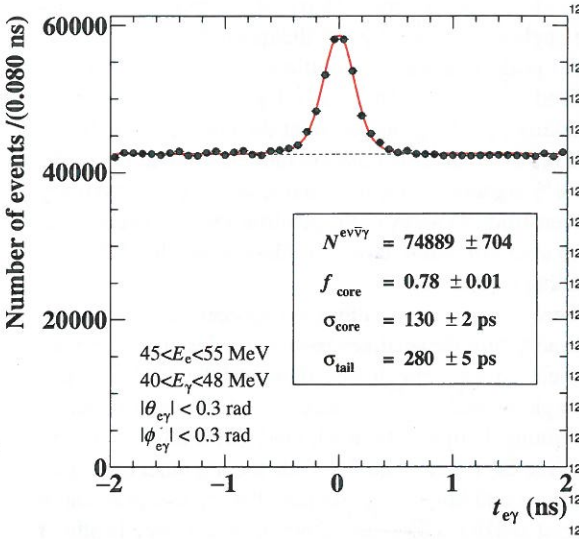

Figure 18 Distribution of $t_{e\gamma}$ for MEG standard trigger. The peak is from radiative muon decays, the flat component is from accidental coincidences.

ment suggests any other significant displacement compared to the associated uncertainties. Therefore, we regard the observed value of δz as an estimate of the systematic uncertainty of the survey while keeping all other survey results for the alignment. Finally, we assign the same angular uncertainty estimated for $\theta_{e\gamma}$ to $\phi_{e\gamma}$.

3.3.2 Relative γ-ray–positron timing

Editor's comments:
Section coordinator: Cecilia Voena
Text: 0.5
Figure: 1.

The relative time $t_{e\gamma} = t_\gamma - t_e$ is defined as the difference of the γ-ray time (see Sect. 3.1.2) and the positron times (see Sect. 3.2.6) calculated at the target. The relative time is calibrated using the radiative muon decay peak observed in the energy side-band[5] shown in Fig. 18. The centre of this distribution is used to correct the time offset between the TC and LXe detectors. The position of the RMD-peak corresponding to $t_{e\gamma} = 0$ was monitored constantly during the physics data-taking period and found to be stable to within 15 ps. In order to obtain the resolution on $t_{e\gamma}$ for signal events, the resolution of Fig. 18 must be corrected for the γ-ray energy dependence as measured in the CEX calibration run and for the positron energy dependence (from MC simulation), resulting in $\sigma_{t_{e\gamma}} = 122 \pm 4$ ps.

The dominant contributions to the $t_{e\gamma}$ resolution are the positron time-of flight resolution (75 ps), the TC intrinsic time resolution (65 ps), and the LXe time resolution (64 ps).

[5] Side-bands are defined in Sect. 4.2.2

3.3.3 γ-ray-AIF analysis

Editor's comments:
Section coordinator: Gordon
Text: 0.5
Figure: 1.

In order to determine if a γ-ray originates from positron AIF, the following three quantities are calculated for each possible e^+_{AIF}–γ pair from all reconstructed e^+_{AIF} candidates and γ-rays in the event: the θ and ϕ angle differences between the AIF candidate direction and the vector connecting the γ-ray and the AIF vertex (θ_{AIF} and ϕ_{AIF}), and the time difference between the γ-ray and the AIF candidate (t_{AIF}). If there are multiple e^+_{AIF} candidates per event, a ranking of e^+_{AIF}–γ pairs is performed by minimising a χ^2-function based on these three observables.

An example plot of ϕ_{AIF} vs. θ_{AIF} for the highest ranked e^+_{AIF}–γ pairs per event in a random sample of 2009 events is shown in Fig. 19. The peak at the centre is caused by γ-rays originating from positron AIF in the DCH. The peak has a tail in the negative ϕ_{AIF} direction since the AIF vertex is reconstructed at the first DCH cathode foil immediately after the last hit in the e^+_{AIF} candidate. However, if the last hit is located in the left plane of a DCH module, it is equally likely (to first order) that the AIF occurred in the first cathode foil of the next DCH module.

The observables θ_{AIF} and ϕ_{AIF} are combined into a 1D "distance" from the peak where the correlation between θ_{AIF} and ϕ_{AIF} and the structure of the two peaks are considered. The smaller the distance, the more likely the event is a true AIF background. We cut events falling within 0.7σ from either of the two peaks. The estimation for the fraction of rejected AIF background events is 1.9%, while loosing 1.1% of signal events. The fraction of rejected RMD events is the same as that of signal. The cut based on the AIF analysis is employed in the physics analysis to remove outlier events which happen to be signal-like with an AIF γ-ray.

4 Analysis

Section coordinator: Fabrizio, Wataru
Text: 8.
Figure: 14.

4.1 Dataset

Editor's comments:
Section coordinator: Toshiyuki
Text: 0.5
Figure: 0. Data were accumulated in the years 2008–2013. The data accumulated in 2008 were presented in [4], but the quality of those data was degraded by problems with the

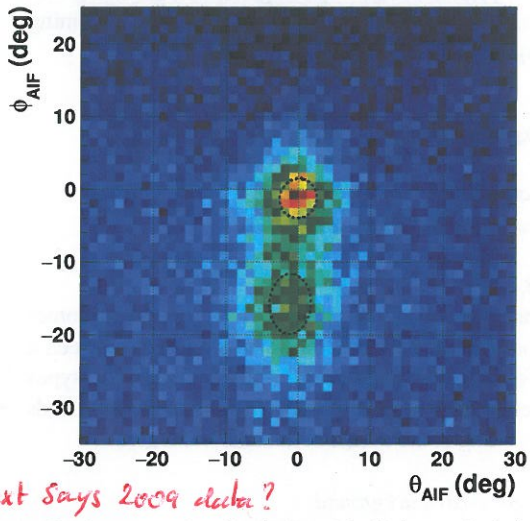

Figure 19 The ϕ_{AIF} vs. θ_{AIF} distribution of the highest ranked $e^+_{AIF} - \gamma$ pairs per event in a sample of year 2011 events. The peak in the centre of the plot is caused by γ-rays originating from positron-AIF in the DCH. The events located inside the dashed line contour are removed by the AIF cut.

tracking system and therefore they are not considered in this analysis.

Figure 20 shows the accumulated stopped muons on target as a function of time. There have been usually planned accelerator shutdown periods of between 4–5 months which are used for accelerator shutdown at PSI in winter for about 4–5 months when the detector maintenance, modification, repair work etc. were performed. During routine physics data taking many kinds of calibration data were taken without losing live-time by introducing different types of triggers with proper pre-scaling factors as well as the main physics trigger.

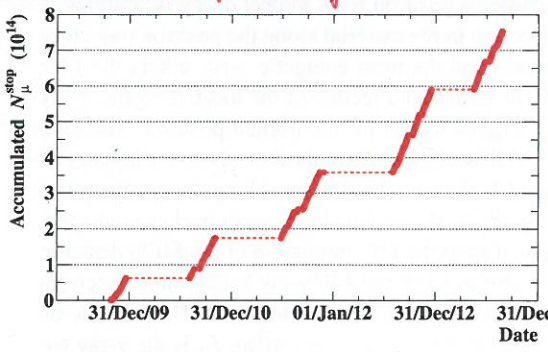

Figure 20 The accumulated number of stopped muons on target as a function of time

In total, 7.5×10^{14} stopped muons on target in the 2009–2013 were recorded. The analysis based on the 3.6×10^{14} stopped muons on target in the 2009–2011 were already published[6]. The data from 2.3×10^{14} stopped muons on target in 2012, and from 1.6×10^{14} in 2013 are newly added in this analysis, thus completing the full data set.

4.2 Event selection

Editor's comments:
Section coordinator: Fabrizio, Wataru
Text: 0.5
Figure: 1.

4.2.1 Preselection

The MEG analysis strategy is based on a combination of blind and likelihood analysis. The blind analysis was chosen to prevent any biases in the evaluation of the expected background in the signal region and the likelihood analysis was preferred to the simpler box analysis in order to avoid boundary effects at the borders of analysis regions and to improve the sensitivity by correctly taking into account, for each event, the probabilities of being due to signal, RMD or accidental background.

The $\mu^+ \to e^+\gamma$ event is characterised by a $e^+\gamma$-pair, simultaneously emitted in opposite directions. Since the positron mass is negligible with respect to the muon mass, the muon energy is equally shared between the positron and the γ-ray, both having an energy $E_e = E_\gamma = m_\mu/2 = 52.83$ MeV. The $\mu^+ \to e^+\gamma$ event signature is therefore very simple and the sensitivity of the experiment is determined by the capability of rejecting positron–γ-ray background pairs, of various origins, which can mimic it. Positron and γ-ray energies, E_e and E_γ, positron–γ-ray relative time $t_{e\gamma}$ and relative angles, $\theta_{e\gamma}$ and $\phi_{e\gamma}$, are the handles to distinguish possible $\mu^+ \to e^+\gamma$ candidates from spurious pairs. In the MEG official analysis $\theta_{e\gamma}$ and $\phi_{e\gamma}$ are treated separately, with independent distributions, since these variables can have different experimental resolutions. However, in an other analysis, used as cross check, the relative stereo angle $\Theta_{e\gamma}$ is looked for instead of the two relative polar angles: we will discuss later the features of the MEG official analysis and of that used as cross-check and compare the results.

In the first stage of MEG analysis, events are pre-selected with loose cuts, requiring the presence of (at least one) positron track and $|t_{e\gamma}| < 4$ ns, a relative timing window more than 25 times larger than the experimental resolution. This procedure reduces our data size to $\sim$ 16 % of the recorded events. No requests are done on γ-ray and positron energies or relative directions. Such loose cuts ensure that even in presence of not yet optimised calibration constants the chance of losing a good $\mu^+ \to e^+\gamma$ event is negligible.

4.2.2 Blinding

Every time the pre-selected events are processed, events falling in the window in the $(t_{e\gamma}, E_\gamma)$ plane defined by $|t_{e\gamma}| < 1$ ns and 48 MeV $< E_\gamma <$ 58 MeV ("Blinding Box") are hidden. The MEG Blinding Box is shown in Fig. 21.

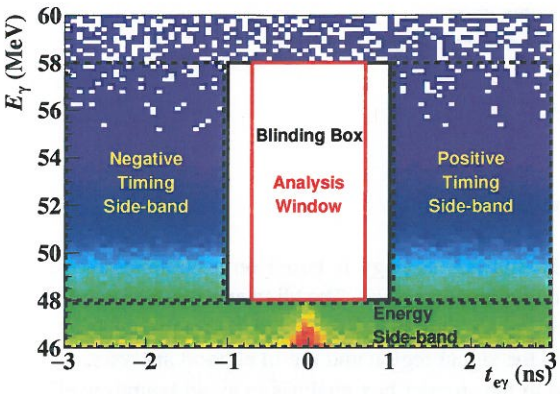

Figure 21 The MEG blinding box and side-bands.

Events with arbitrary γ-ray energy but $|t_{e\gamma}| > 1$ ns fall in "timing side-bands", the left side-band corresponding to $t_{e\gamma} < -1$ ns and the right side-band to $t_{e\gamma} > 1$ ns, while events with arbitrary relative timing but γ-ray energy outside the Blinding Box interval fall in the "energy side-band". RMD events, with zero relative timing, belong to the energy side-band and, as stated in Sect. 3.3.2, are used to accurately calibrate the timing difference between LXe detector and TC. Events in the timing side-bands are very likely accidental events; so, their positron and γ-ray energy spectra and relative angles distributions are uncorrelated. We also define "angle side-bands" the regions corresponding to 50 mrad $< |\theta_{e\gamma}| <$ 150 mrad or 75 mrad $< |\phi_{e\gamma}| <$ 225 mrad, which are used for self-consistency checks of the analysis procedure.

Side-band events are studied in detail to optimise algorithms and analysis quality, to estimate the background in the signal region and to evaluate the experimental sensitivity by using toy MC simulations, as will be explained later. At the end of optimisation procedure, the Blinding Box is open and a maximum likelihood fit is performed to extract the number of Signal ($N_{\rm sig}$), RMD ($N_{\rm RMD}$) and ACCIdental background ($N_{\rm ACC}$) events. The likelihood fit is performed on events falling in the "Analysis Window" defined by 48 MeV $< E_\gamma <$ 58 MeV, 50 MeV $< E_{\rm e} <$ 56 MeV, $|t_{e\gamma}| < 0.7$ ns, $|\theta_{e\gamma}| < 50$ mrad and $|\phi_{e\gamma}| < 75$ mrad. The size of the analysis window was chosen, on the basis of experimental resolutions, to be between five and twenty sigmas on all physical variables in order to prevent any risk of losing good events and to maintain the number of events to be fitted at a reasonable amount. The same fitting procedure is preliminarily applied to equal size regions in the timing and angle side-bands (with appropriate shifts on relative timing or angles) to check the whole calculation.

4.3 Background study

Editor's comments:
Section coordinator: Toshiyuki, Yusuke
Text: 1.5
Figure: 3.

The background in search for the $\mu^+ \to e^+\gamma$ decay comes either from RMD or from an accidental overlap between a Michel positron and a photon from RMD or AIF. All types of the background were thoroughly studied in the side-bands without looking at the events in the analysis region.

4.3.1 Single γ-ray background

Editor's comments:
Section coordinator: Toshiyuki

High energy single γ-ray background events are mainly produced by two processes: RMD and AIF of positrons. The contribution from the external Bremsstrahlung is negligibly small in our analysis region. RMD is the Michel decay with the emission of a γ-ray, also called inner Bremsstrahlung. The integrated fraction of the spectrum of γ-ray from RMD in the signal region is roughly proportional to the square of the energy resolution[27]. AIF γ-ray background events are produced when a positron from Michel decay is annihilated with an electron in the material along the positron trajectory in two γ-rays and the most energetic γ-ray enters the LXe detector. The emission direction of the most energetic γ-ray is closely aligned to that of the original positron. The total amount of AIF background events depends on the layout and the material budget of the detector along the positron trajectory. Figure 22 shows the single γ-ray background spectra calculated from the MC simulation of the MEG detector. The green circles show the AIF γ-ray background spectrum and the red crosses show that due to RMD. The variable on the horizontal axis is $y = 2E_\gamma/m_\mu$ where E_γ is the γ-ray energy and m_μ is the muon mass. The integrated γ-ray yield per decay above y is plotted on the vertical axis (the maximum allowed value for y is slightly smaller than one for RMD and slightly larger than one for AIF, due to the electron mass). The RMD γ-ray fraction is 55%, and the AIF γ-ray fraction is 45% in the $y > 0.9$ region. From Fig. 22, AIF becomes dominant in the $y > 0.92$ region. Since the energy spectra decrease rapidly as a function of y near the signal region, the reduction of the single γ-ray background requires a good energy resolution. In addition to the RMD and AIF components in the signal region, there are contributions from pile-up γ-rays and cosmic ray components, at