Available online at www.sciencedirect.com
ScienceDirect
Nuclear Instruments and Methods in Physics Research A 587 (2008) 101–109
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

The ZEPLIN II dark matter detector: Data acquisition system and data reduction

G.J. Alner[a], H.M. Araújo[b,a], A. Bewick[b], C. Bungau[a,b], B. Camanzi[a], M.J. Carson[c,*], H. Chagani[c], V. Chepel[d], D. Cline[e], D. Davidge[b], J.C. Davies[c], E. Daw[c], J. Dawson[b], T. Durkin[a], B. Edwards[b,a], T. Gamble[c], J. Gao[e], C. Ghag[f], W.G. Jones[b], M. Joshi[b], E.V. Korolkova[f], V.A. Kudryavtsev[c], T. Lawson[c], V.N. Lebedenko[b], J.D. Lewin[a], P.K. Lightfoot[c], A. Lindote[d], I. Liubarsky[b], M.I. Lopes[d], R. Lüscher[a], P. Majewski[c], K. Mavrokoridis[c], J.E. McMillan[c], B. Morgan[c], D. Muna[c], A.S. Murphy[f], F. Neves[d], G.G. Nicklin[c], W. Ooi[e], S.M. Paling[c], J. Pinto da Cunha[d], S.J.S. Plank[f], R.M. Preece[a], J. Quenby[b], M. Robinson[c], C. Silva[d], V.N. Solovov[d], N.J.T. Smith[a], P.F. Smith[a], N.J.C. Spooner[c], T.J. Sumner[b], C. Thorne[b], D. Tovey[c], E. Tziaferi[c], R.J. Walker[b], H. Wang[e], J. White[g], F.L.H. Wolfs[h]

[a]*Particle Physics Department, Rutherford Appleton Laboratory, Chilton, UK*
[b]*Blackett Laboratory, Imperial College London, UK*
[c]*Department of Physics & Astronomy, University of Sheffield, UK*
[d]*LIP-Coimbra & Department of Physics, University of Coimbra, Portugal*
[e]*Department of Physics & Astronomy, University of California, USA*
[f]*School of Physics, University of Edinburgh, UK*
[g]*Department of Physics, Texas A&M University, USA*
[h]*Department of Physics & Astronomy, University of Rochester, USA*

Received 16 May 2007; received in revised form 31 October 2007; accepted 2 November 2007
Available online 12 November 2007

Abstract

ZEPLIN II is a two-phase (liquid/gas) xenon dark matter detector searching for WIMP-nucleon interactions. In this paper we describe the data acquisition system used to record the data from ZEPLIN II and the reduction procedures which parameterise the data for subsequent analysis.

PACS: 95.35.+d; 07.05.Hd; 07.05.Kf

Keywords: Dark matter experiments; Data acquisition; Data analysis

1. Introduction

The ZEPLIN II dark matter detector has been operational at the Boulby Mine underground laboratory since 2005. Its principle aim is to detect and measure the faint nuclear recoil signal from galactic Weakly Interacting Massive Particles (WIMPs). ZEPLIN II is a two-phase (liquid/gas) xenon detector which has an increased sensitivity over previous UK Dark Matter Collaboration experiments NaIAD [1] and ZEPLIN I [2]. By measuring both the primary and secondary scintillation signals produced by particles interacting in the target volume [3–5] this technique has the potential to improve discrimination

*Corresponding author.
E-mail address: m.j.carson@sheffield.ac.uk (M.J. Carson).

0168-9002/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2007.11.005

between nuclear recoils expected from WIMP interactions (also produced in neutron collisions with nuclei) and electron recoils (caused by γ-rays and electron interactions, e.g. β-particles). The differential energy spectrum of these nuclear recoil events is expected to be featureless and smoothly decreasing with detected recoil energies which are less than 100 keV_{ee} (where keV_{ee} is electron recoil equivalent energy) for WIMP masses in the range 10–1000 GeV/c^2 [6]. Dark matter experiments need to be capable of detecting recoil energies of a few keV_{ee} in order to place the most stringent limits on the rate of WIMP interactions. In scintillator experiments this requires sensitivities to single photoelectrons. Due to the extended waveform digitisation and fine resolution, a large volume of data is collected (~8 TBytes/yr, excluding calibration data). The data requires efficient processing and must be stored for subsequent analysis. Reduction procedures capable of parameterising the data have been developed to process the waveforms and output a set of parameters representative of the original waveform.

In this paper we present a brief description of the detector followed by a more detailed discussion of the data acquisition system and data reduction procedures. In our companion papers [7,8] we presented initial results from the first underground run of ZEPLIN II.

2. The detector

ZEPLIN II consists of 31 kg of liquid xenon contained in a 50 cm diameter copper vessel (Fig. 1). The target volume is viewed from above by seven ETL D742QKFLB 130 mm photomultiplier tubes (PMTs) in a close-packed hexagonal pattern. Two steel grids and a mesh act as electrodes which define the electric field in the detector. Two of the grids are positioned above and below the liquid surface; these are the top and bottom grids which define the electroluminescence and charge extraction fields. The wire mesh is located at the bottom of the target defining the drift field in the active volume of the target. PTFE lining the inside of the copper vessel acts both as a support structure for the high voltage grids and a reflector. The target is maintained at liquid xenon temperature by an IGC PFC330 Polycold system [9]. Circulation of the xenon through SAES getters (model PS11-MC500) [10] ensures impurities do not reduce the performance of the target. The copper vessel is surrounded by a stainless steel jacket which provides an insulating vacuum. The detector sits in a liquid scintillator veto system viewed from above by 10 ETL 9354KA 200 mm PMTs. The upper half of the detector/veto system is surrounded by hydrocarbon slabs (separated by Gd-loaded resin sheets) shielding the target from external background neutrons (the liquid scintillator veto acts as its own neutron shield). The entire apparatus is housed in a lead 'castle' which shields against external background γ-rays.

In normal (two-phase) operating mode ZEPLIN II is designed to detect the vacuum ultra-violet scintillation and ionisation charge signal of an interacting particle.

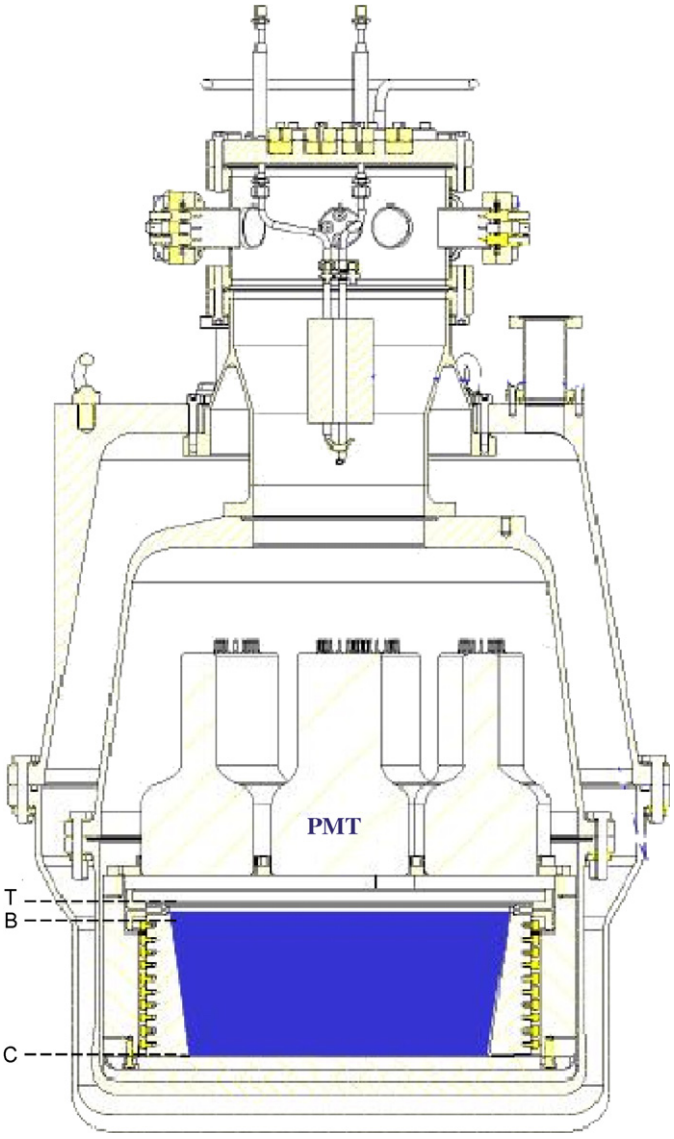

Fig. 1. Schematic of ZEPLIN II showing the relative positions of the electric field grids and the liquid/gas interface where T is top grid, B is bottom grid and C is the cathode.

The scintillation light comes from the initial particle interaction (prompt de-excitation of Xe_2^* dimer to dissociative ground state). Under an applied electric field, a fraction of the ionisation electrons are drifted to the liquid surface where they are removed by the extraction field into the gas phase producing a secondary scintillation pulse through electroluminescence. We adopt the convention of naming the primary and secondary signals S1 and S2, respectively.

The region between the bottom grid and the cathode defines the active volume of the detector and contains a drift field of ~1 kV/cm. The electron drift velocity in this region is about 2.0 mm/μs. Field shaping rings embedded in the PTFE support structure keep the drift field lines parallel. The top grid defines both the extraction and electroluminescence fields in the region between the top and bottom grids. The time delay between the S1 and S2

signals permits the depth of the particle interaction to be reconstructed.

3. Data acquisition system

Fig. 2 shows the signal path from PMT to the data acquisition system (DAQ). The seven PMT signals from the target are first passively split by a 50 Ω splitter: one line is fed to a ×10 amplifier, the other to the signal input of the DAQ (channels ACQ1 to ACQ7 in Fig. 2). PMT gains are equalised (by adjusting the high voltage) to give about 45 mV per photoelectron (pe) output after splitter and amplification. The PMT gains have shown stability over time (see Ref. [11] for details of PMT performance and stability). From the amplifier the signals are fed to discriminator D1 which outputs 50 mV/channel for input signals above 17 mV (approximately $\frac{2}{5}$ of the spe signal). The FWHM of the PMT signals is typically 0.04 V ns [11]. The logic sum of the signals is then fed to discriminator D2 which outputs a NIM pulse to dual timer T2 when five out of seven PMTs detect a signal above the 17 mV threshold level. At the same time, T2 sends a 100 μs square-wave pulse to the veto input of D1 to prevent further triggers until the whole waveform is read by the DAQ. The amplified signal from the central PMT (PMT 1) is also attenuated before going into discriminator D3. Due to the relative placing of the PMTs and the PTFE support/reflector structure, the central PMT sees a larger signal (on average) than the outer PMTs. Large amplitude signals can cause optical feedback in the target giving rise to many noise pulses of long duration. This is due to scintillation light being reflected back into the target giving rise to electrons (from photoelectric emission) which are then drifted to the surface producing electroluminescence signals. These signals can occur for tens of μs after the main S2 signal. Signals exceeding a pre-set threshold of 200 mV (see below for explanation) are vetoed by the dual timer T2 which sends a 1 ms inhibit signal to D1 to prevent optical feedback signals from triggering the system. This also has the effect of reducing the trigger rate by 60%, reducing data processing and data storage requirements. The 10 PMT signals from the veto feed 10 channels of a discriminator/buffer NIM module. The sum of the discriminator outputs is passed into a second discriminator whose threshold is set to output a logic pulse if at least three of the veto PMTs fire in coincidence. This pulse is delayed by 100 ns in a delay line and the output is added to the analog sum of the veto PMT outputs. The veto pulse is recorded in order to ensure that a 3-fold coincidence has been detected which rules out background signals from single PMTs.

The waveform hardware consists of DC265 M2M ACQIRIS digitizers embedded in CC103 ACQIRIS [12] crates. These are based on CompactPCI technology interfaced through the PCI bus. The digital conversion of signals has an 8 bit resolution, a conversion rate of up to 500 MSamples/s, a bandwidth of 150 MHz and a memory of 2 MPoints/channel. Each 200 μs waveform is sampled at 2 ns intervals. LINUX-based software reads out the digitised waveforms which are then written to disc. Monitoring of all target parameters such as temperature and pressure is done with a 64 channel Datascan [13] module via the serial port on the DAQ computer.

Signals exceeding an amplitude of 200 mV are vetoed by the DAQ electronics. This corresponds to a high energy cut-off of approximately 100 keV_{ee}. In addition, an upper cut of 180 mV is implemented in software. Signals above this threshold are above the energy range of interest for dark matter searches but the effect, in terms of efficiency, must be estimated. Fig. 3 shows the efficiency as a function of energy calculated from data taken with and without the software threshold—but keeping the additional 200 mV saturation cut on the central PMT in both cases. The binned spectrum from data taken with the software threshold is divided by the binned spectrum from data taken without the software threshold to produce the efficiency which is 100% up to 30 keV_{ee}.

In order to investigate the fraction of events lost due to DAQ dead time dedicated pulser measurements were performed. The pulser was set to output a pulse of amplitude 0.81 V with 0.2 ms duration. For a waveform of 100,000 samples at 2 ns per sample the maximum recordable rate was 22 Hz, corresponding to a dead-time of ~50 ms. In a non-paralyzable model, where events occurring during dead periods do not extend the dead-time, the actual event rate n is given by $n = m/(1 - m\tau)$ [14], where m is the measured event rate and τ is the dead-time. The typical measured background rate in the target is 2 Hz which corresopnds to a loss of ~10% of events (since $m/n = 1/(1 + n\tau)$). For data taken with an AmBe neutron source located approximately 1 m above the target the actual event rate was 34.4 Hz which corresponds to a loss of 67%. For a ^{60}Co source located in the same position the loss was 78% with an actual event rate of 70 Hz.

4. Data reduction

A raw data file with 2000 events is approximately 250 MB compressed (each waveform has 100,000 points and there are seven PMT channels plus one veto channel for each event). Approximately 25 GB of data (excluding calibration runs) is recorded each day. The data are then written to magnetic digital tape on an ADIC Scalar 100 tape robot system [15]. Each tape can store 100 GB of data. The data are subsequently transferred to an Apple XGrid and Linux cluster capable of reducing ~0.6 TB of data per day.

Raw data are reduced with a LINUX-based application which reads in the binary data files and outputs a set of numeric parameters representing each pulse found on each waveform. The software consists of an event viewer, allowing the examination of each trace in each PMT and reduction algorithms which process the waveforms from each channel. All peaks in each waveform must be

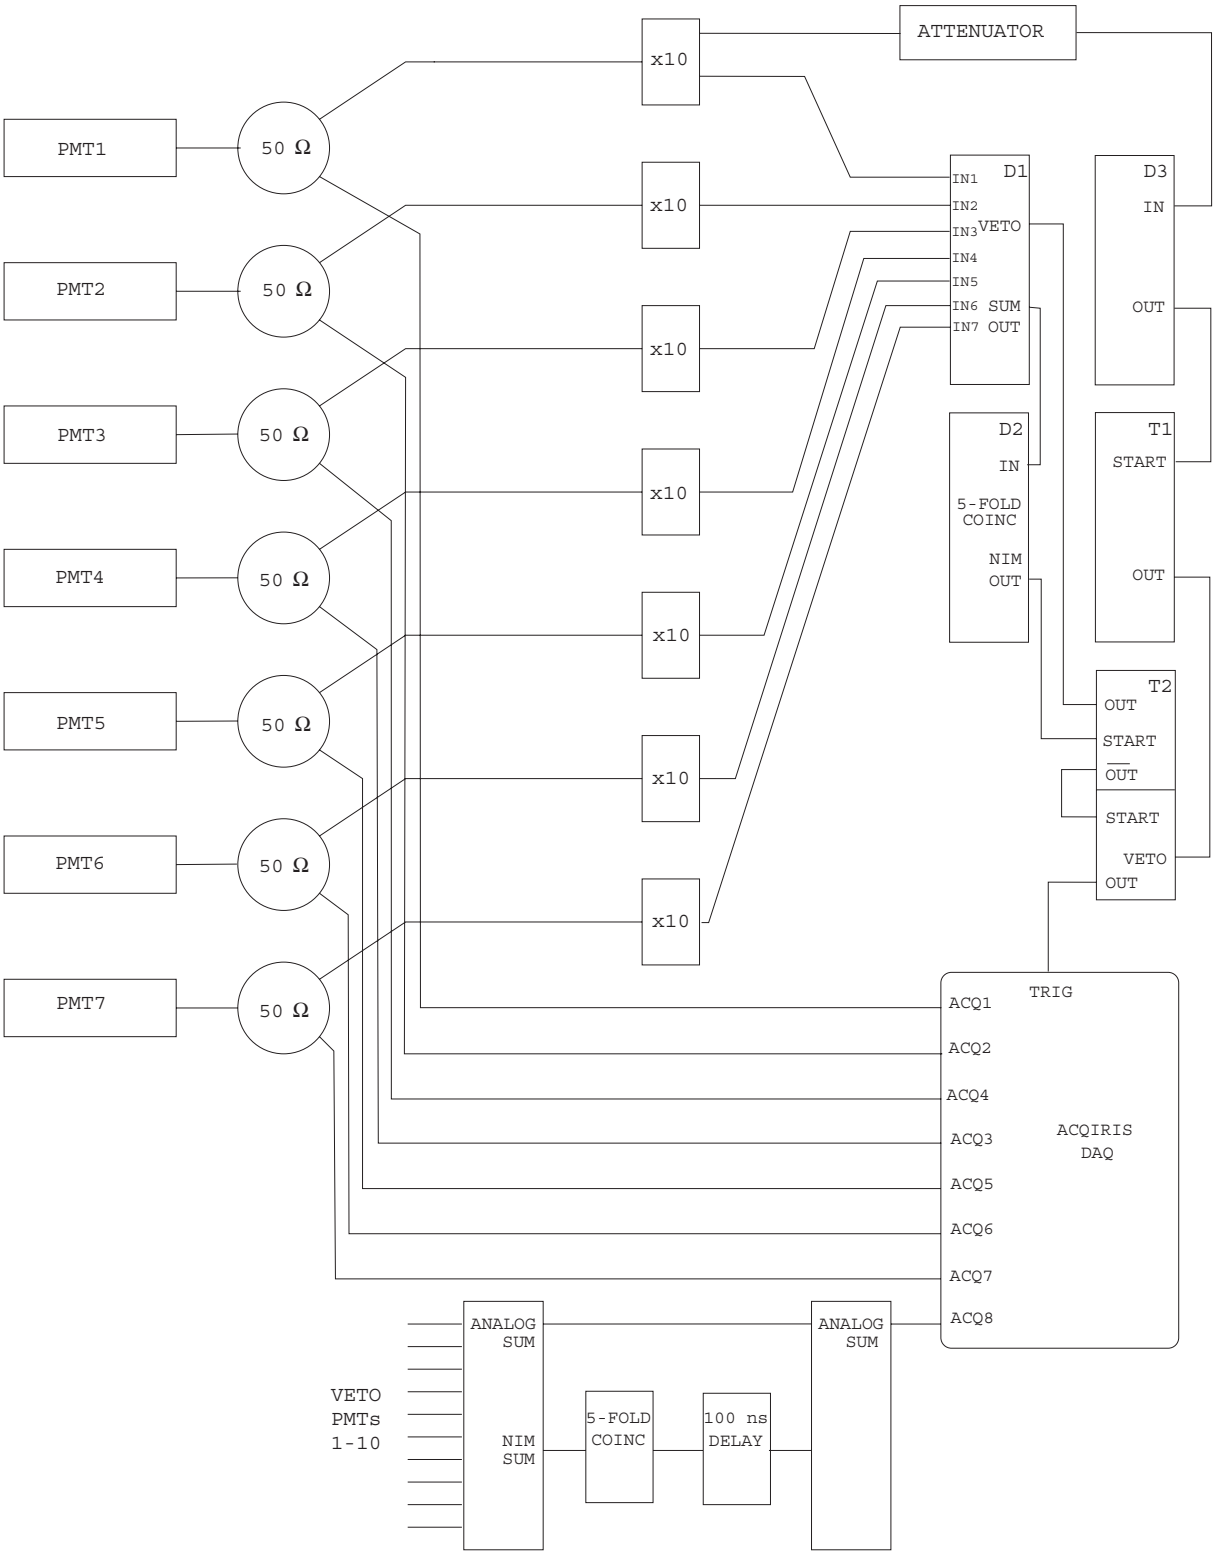

Fig. 2. Signal path from target and veto PMTs to DAQ via trigger electronics.

identified and parameterised according to height, width, area and time-constant. Pulse parameters are then written to reduced data files in HBOOK ntuple format [16] for subsequent analysis. The user must specify a number of input variables for the peak-finding algorithms, these are discussed in turn.

4.1. Input parameters

Different length signal cables from the PMTs to the DAQ and different PMT characteristics can induce delays in the pulse arrival times in each channel. In addition, differences in transit times in the PMTs can induce delays

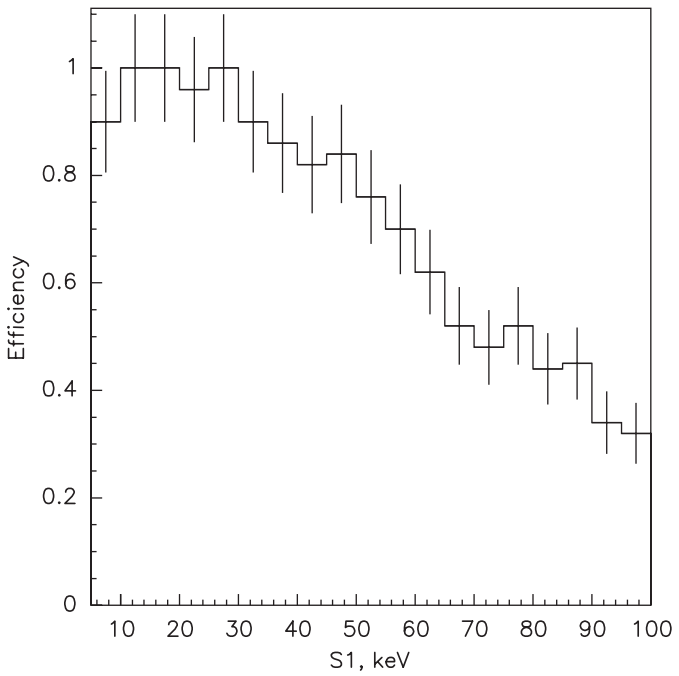

Fig. 3. Software saturation cut efficiency as a function of energy. Efficiency is 100% up to 30 keV_{ee}. (*Note:* horizontal scale is electron recoil equivalent energy.)

up to ~10 ns. Pulses detected coincidently in several channels can appear spread out in the summed waveform if these delays are not corrected. This can lead to peaks in the summed waveform being incorrectly parameterised as separate pulses. Fig. 4 shows the distribution of pulse mean arrival times in each PMT relative to the central PMT (PMT 1) for uncorrected data. Channels are shifted in software by their delay with respect to PMT 1 and the summed waveform calculated.

To facilitate peak-finding a smoothing function is applied to the summed waveform. This ensures that small amplitude fluctuations do not get mis-identified as valid signal pulses. The amplitude h_t at each sample is smoothed as

$$h_s = \frac{\sum_{t=-t_{sm}/2}^{t_{sm}/2} h_t}{N} \quad (1)$$

where t_{sm} is the smoothing timescale and is an input parameter to the reduction and N is the number of points over the interval. Fig. 5 shows a single photoelectron spectrum from the central PMT parameterised with different values of t_{sm} from 5 up to 50 ns. The larger value of $t_{sm} = 50$ ns can be seen to 'wash-out' lower energy pulses resulting in a shift in the peak of the spectrum to higher energies. A value of $t_{sm} = 12$ ns was chosen as it reproduces the measured value of the spe calibrations (0.038 V ns/pe).

All pulses above a user-defined software threshold are tagged, up to a maximum of 10. The software threshold depends on the full-scale of the DAQ and is set to 2 mV (almost $\frac{1}{2}$ pe) for data acquired with a full scale of 200 mV. For an 8-bit digitizer the smallest resolution at this scale is 0.8 mV/bit. A 2 mV threshold ensures that identified pulses are above the noise level. The range was chosen as a compromise between energy resolution for S1 signals and the dynamic range for S2 pulses. The smoothed amplitude h_s is used only to locate peaks on the summed waveform, all other parameterisation uses the unsmoothed original data.

We define a clustering timescale which allows closely spaced peaks on the waveform to be grouped together to form a single pulse. A low energy S2 signal can appear as separate peaks spread out over several μs (the total width of S2 is determined by the distance between the top grid and the liquid surface). Fig. 6 shows the effect of different clustering values on the same pulse. A clustering of 50 ns causes peaks `p3`, `p5` and `p6` to be identified as separate pulses from `p4`. Increasing the clustering timescale to 400 ns correctly groups all peaks as a single pulse.

4.2. Output parameters

The baseline for each waveform is calculated on an event-by-event and channel-by-channel basis. An initial baseline is calculated from the mean of the first 500 data points. However, this is not sufficient because the baseline tends to wander by a few mV during the event [11]. A box-car smoothing algorithm with a width of 5000 ns is used to calculate a wandering baseline to compensate for this effect. As the wandering baseline is calculated, any data point deviating by more than 5× the RMS noise from the current value of the baseline is excluded from the baseline calculation. This prevents the wandering baseline from following the slow S2 signals from the gas phase of the data.

An initial scan is made of the waveform on the sum channel for pulses with amplitudes above the software threshold. The start time of each pulse $t_s(i)$ (the time at which the pulse rises above the threshold level, see Fig. 7) where i is the index of the pulse (up to the maximum of 10) is then used to find the actual start time of the pulses $t_p(i)$ on the baseline. The difference between the end and start time defines the pulse widths $w(i)$. Once all pulses on the summed waveform are identified, each individual channel is scanned in the time windows $w(i)$. Each $t_p(i)$ is used to define a time $t_0(i)$, at which 10% of the total charge of the pulse is detected. This is then used to calculate the charge mean arrival time τ as in Eq. (2)

$$\tau = \frac{\sum_{t=t_0}^{t_p+w} h_t \cdot (t - t_0)}{\sum_{t=t_0}^{t_p+w} h_t} \quad (2)$$

where h_t is the amplitude at time t within the time window. The FWHM of the pulse is calculated starting from the time at which h_t is at maximum and tracing outwards in both directions to the times at which the pulse height falls to half of the maximum h_t. Another measure of the full width at half maximum (LFWHM) is calculated by tracing

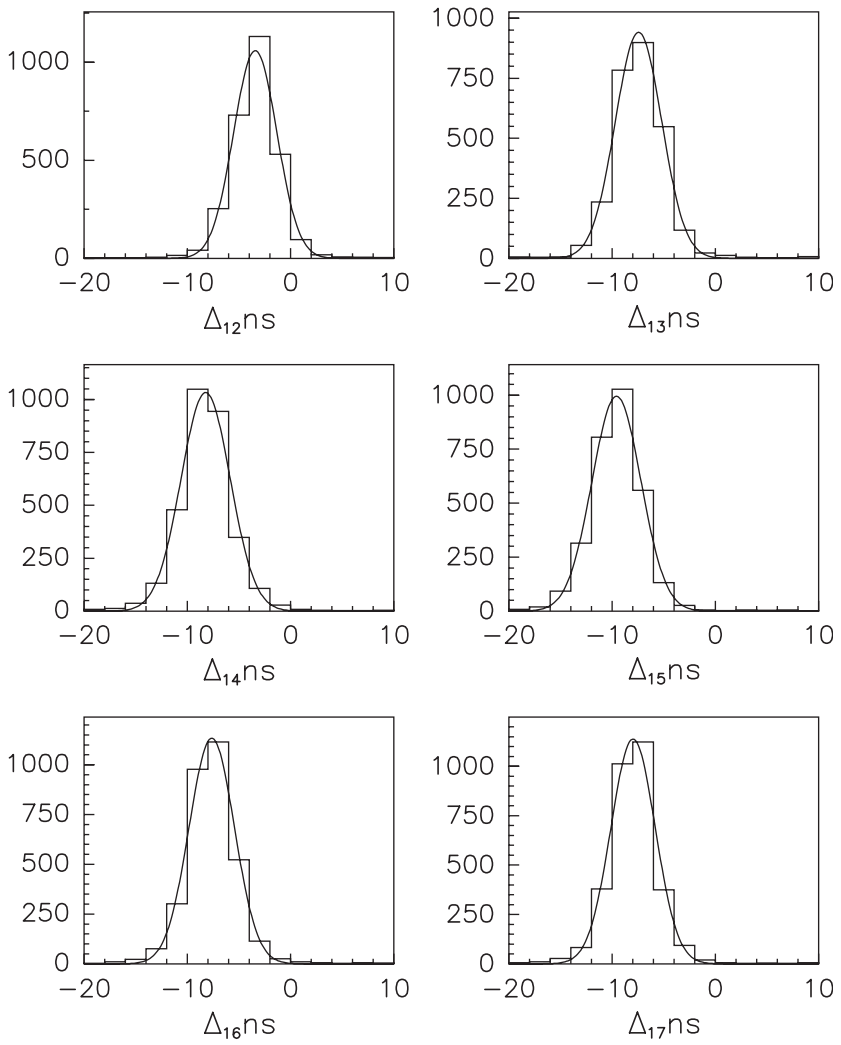

Fig. 4. Mean signal delay Δ_{1i} between PMT 1 and PMT i. The Δ_{1i} are 3.9, 7.9, 8.7, 10.1, 6.3 and 6.9 ns for $i = 2, \ldots, 7$, respectively.

the pulse height to its half maximum value by starting at both the beginning and end of each pulse and moving in towards the maximum pulse height. These parameters are useful for subsequent off-line analysis as, in practice, they can have slightly different values.

The RMS noise of the waveform is the mean deviation of the data from the baseline in the pre-trigger region

$$\text{Noise}_{\text{RMS}} = \sqrt{\frac{\sum h_t^2}{N}} \tag{3}$$

where N is the number of points in the pre-trigger region. Pulse area A (in units of V ns) is defined as the integrated area of the pulse in the time window and is intended as the best measure of the total charge associated with the pulse.

$$A = \sum_{t=t_{\text{p}}}^{t_{\text{p}}+w} A_t \tag{4}$$

where A_t is the discrete area at time t. The summed veto signal is parameterised according to its height, arrival time and total integrated charge following the same procedure for signal pulses from the target. Table 1 lists all pulse parameters calculated.

5. Data analysis

Reduced data files consist of parameters for each pulse found on each of the seven PMT channels defined by those pulses found on the sum channel. All parametersied waveforms must be analysed to extract the required S1 and S2 signals. The secondary ionisation signal will be delayed from the primary scintillation signal by an amount which depends upon the depth of the interaction in the target. The drift velocity of ionisation electrons (which is determined by the drift field) is typically 2.0 mm/μs. The longest drift time is the time taken for ionisation electrons from the cathode to reach the bottom grid and is about 75 μs. Since the DAQ can trigger on either the S1 or S2 signal, waveforms are recorded 100 μs on each side of the trigger position. The waveform time window was chosen to be $[-200, 0]$ μs with the trigger at -100 μs. Events which trigger on the primary will have a secondary signal in the

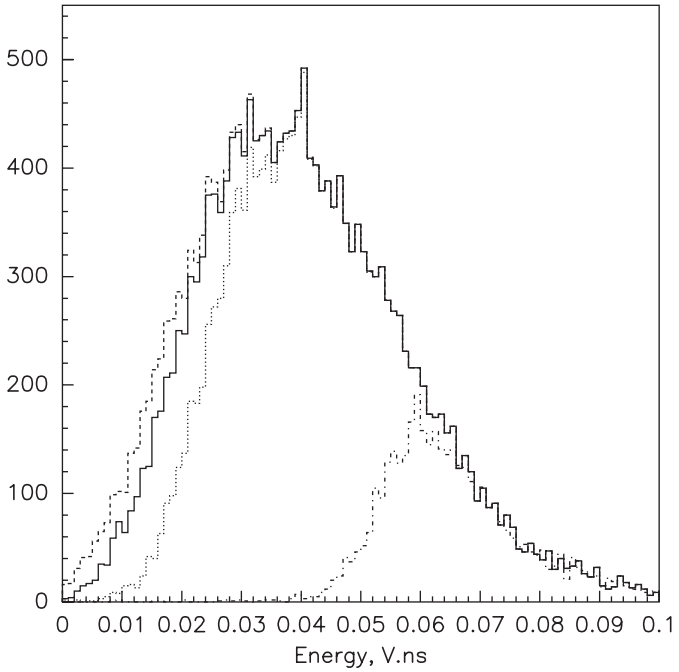

Fig. 5. Single photoelectron spectra for PMT 1 with different values of t_{sm}: 5 ns (dashed line), 12 ns (solid line), 20 ns (dotted line) and 50 ns (dash-dot line). The spectrum peaks at ~0.038 V ns.

range $[-100, 0]$ μs and those events which trigger on the secondary will have a primary scintillation signal in the range $[-200, -100]$ μs. In either case, secondary ionisation signals only occur in the range $[-100, 0]$ μs. The first step in the analysis involves scanning this region for pulses on the sum channel which are greater than 1 V ns in energy (see below). All pulses greater than 1 V ns in this range are tagged as possible S2 signals. Next, the $[-200, -100]$ μs region of the wavefom is scanned for the primary sciltillation pulse. Coincident pulses with an amplitude greater than 1.7 mV in any three out of seven channels are tagged as possible S1 candidates. An additional cut of $\tau > 150$ ns for S2 and $2\,\text{ns} < \tau < 50$ ns for S1 is applied (the efficiencies for these and other selection criteria are discussed in Ref. [7]). We reject those events with more than one S2 signal as WIMPs do not multiple scatter. We also reject events with more than one primary scintillation pulse. All events with *only* one S1 and *only* one S2 are accepted.

The 1 V ns selection criteria ensures that small single electron electroluminescence signals do not get mis-identified as S2. Fig. 8 shows a typical cluster of single electron pulses spread over 600 ns with a total energy of 0.5 V ns. Each peak appears in different channels at different times (not shown) but the clustering parameter groups them together in a single pulse. Low energy nuclear recoils are expected to produce more than one ionisation electron so the probability of rejecting genuine S2 signals is low. To investigate this further, we plot the distribution of S2 for S1 in a restricted energy range in Fig. 9. The distribution peaks at ~10 V ns and is roughly Gaussian

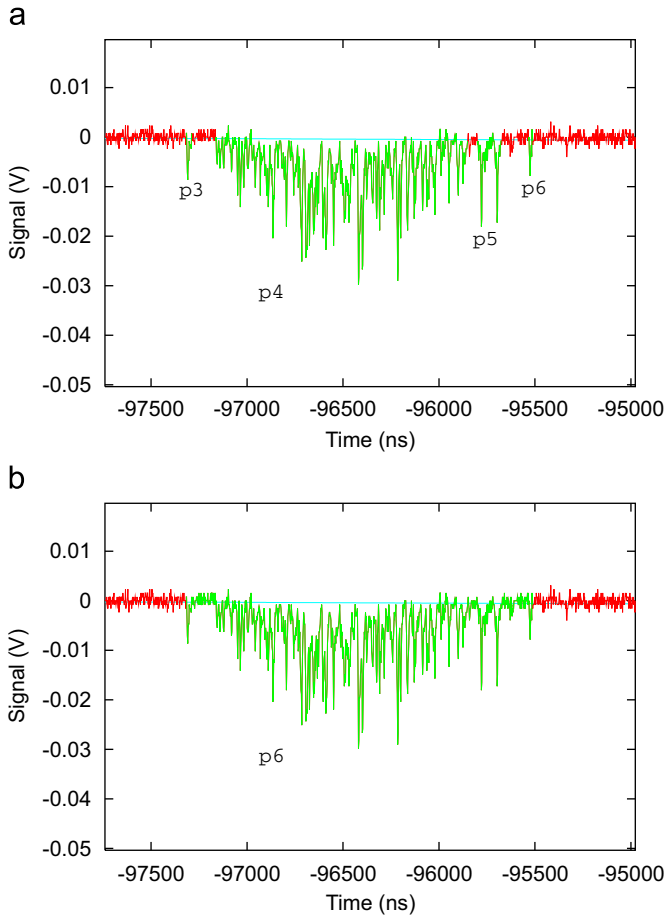

Fig. 6. Ionisation pulse parameterised with 50 ns clustering (a) showing that four different pulses (p3 to p6) have been identified. With 400 ns clustering (b) these have been correctly grouped into a single pulse p6.

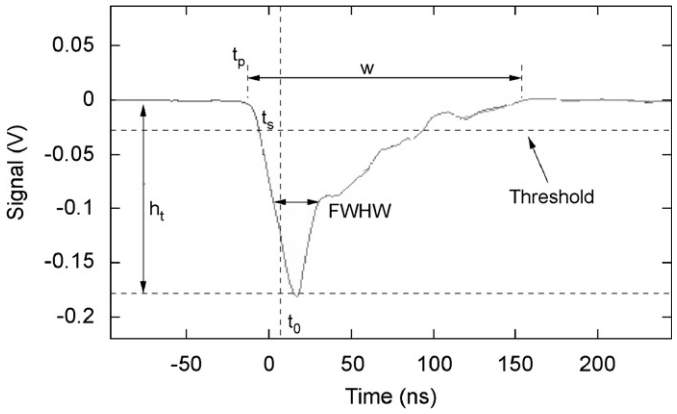

Fig. 7. Basic pulse parameters. h_t is the signal amplitude above the baseline, t_p is the start time of the pulse, t_s is the start time of the pulse at the pulse finding threshold and t_0 is the time at which 10% of the total charge has arrived. w is the pulse width defined by the start and end time of the pulse.

in shape with a noise contribution impinging on the distribution from the left. Fitting an exponential plus Gaussian to the S2 spectrum allows us to calculate the efficiency of the 1 V ns cut by integrating the area of the

Table 1
Waveform parameters calculated by reduction procedure

Parameter	Explanation
B	Baseline level for each channel
$\text{Noise}_{\text{RMS}}$	RMS noise level for each channel
t_p	Start time of all pulses on each waveform
t_0	Time of arrival of 10% of the pulse charge
w	Width of each pulse in each channel
A	Area of each pulse in each channel
h	Height of each pulse in each channel
FWHM	Pulse FWHM measured from maximum pulse height
LFWHM	Pulse FWHM measured from pulse start and pulse end time
τ	Charge mean arrival time
t_v	Start time of veto pulse
h_v	Height of veto pulse
A_v	Integrated area of veto pulse

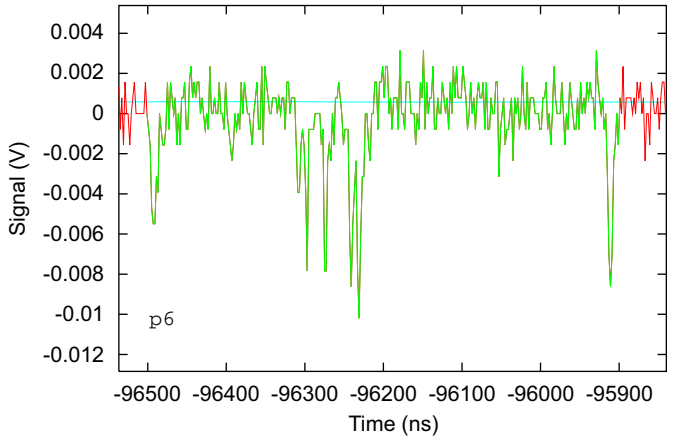

Fig. 8. Typical noise event due to a single electron extracted from the liquid giving rise to an electroluminescence signal in the gas.

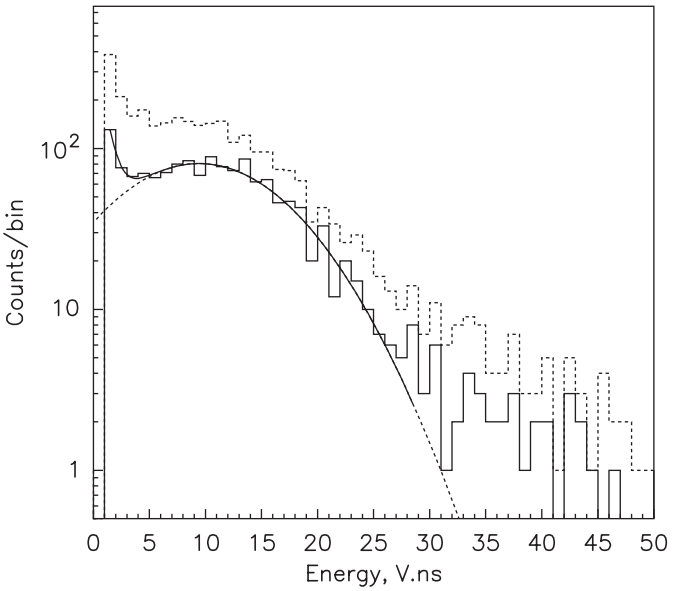

Fig. 9. AmBe S2 distribution for S1 between 5 and 10 keV_{ee} (dashed histogram). The solid histogram is the distribution excluding events near the top and bottom grids. Superimposed is the Gaussian fit (plus exponential noise) to the data.

Gaussian above 1 V ns and dividing by the total area. Requiring a minimum of 1 V ns for S2 results in loss of approximately 8% of events (92% efficiency) between 5 and 10 keV_{ee}. However, the fitted Gaussian extends to below zero, which is unphysical, and so this loss is overestimated. For higher energies the efficiency is close to 100%.

In order to convert the observed pulse energy (in mV) to keV, energy calibrations were performed daily with a ^{57}Co source (two spectral lines at 122 and 136 keV) located beneath the target volume and delivered by a dedicated source delivery mechanism. 10,000 events were recorded with waveforms 200 μs in duration and with a drift field of 1 kV/cm. S1 and S2 signals were extracted from the parameterised data as described. Fig. 10 shows a typical calibration spectrum with a fit to the peak at 2.54 V ns. This corresponds to ~67 pe (1 pe ≈ 0.038 V ns) which gives a light-yield of ~0.55 pe/keV for this data. At 0-field the lightyield was measured to be 1.1 pe/keV [7].

The data analysis is affected by the hardware and software triggers, and by various cuts applied. Trigger and cut efficiencies have been evaluated using data and/or simulations. Trigger efficiency has been simulated taking into account the following factors: (i) measured spectra of the single photoelectron pulses [11]; (ii) hardware and software trigger conditions described above; (iii) measured light collection; (iv) uniformity of the light collection determined by the analysis of high-energy alpha-events; (v) measured energy resolution of the detector. Trigger efficiency is about 35% at 5 keV_{ee} (at the analysis threshold) and reaches 90% at 10 keV_{ee}. Other cuts, discussed in detail in Ref. [7], reduce the overall efficiency by about

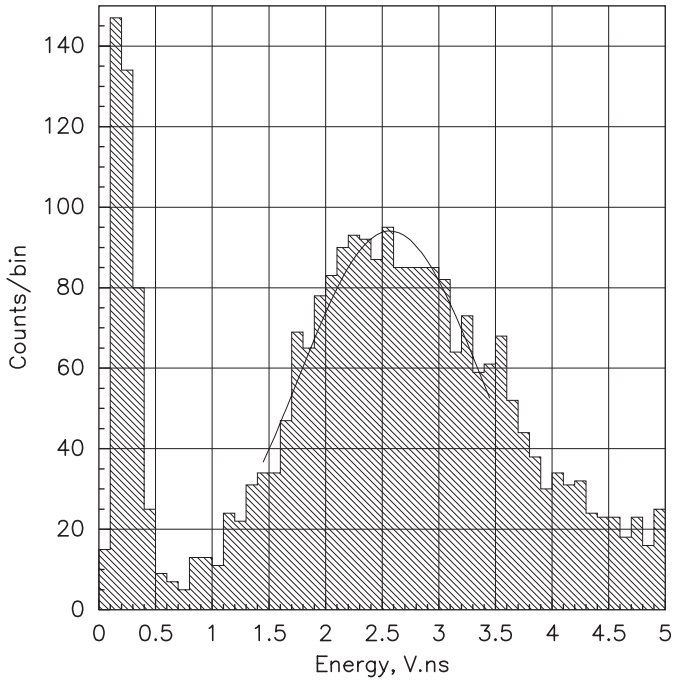

Fig. 10. Energy spectrum from ^{57}Co calibration run with a drift field of ~1 kV/cm in the target. The 122 keV spectral line is a 2.54 V ns. The energy resolution is ~20% (see Ref. [7] for further details).

20%. The calculated efficiencies agree well with the measured shapes of the spectra from gamma (Co-60) and neutron (AmBe) calibrations (see Ref. [7] for detailed discussion).

6. Summary

We have described the data acquisition system for the ZEPLIN II dark matter experiment. Data are recorded with 8-bit ACQIRIS digitizers with a full scale of 200 mV. This corresponds to a cut-off of approximately 100 keV_{ee}, well above the region of interest for dark matter searches. Techniques capable of extracting low energy pulses with good efficiency have been presented and show the 122 keV peak of ^{57}Co at 2.54 V ns. For a primary signal in the energy range 5–10 keV_{ee} the distribution of S2 signals peaks at about 12 V ns. Approximately 8 TB of dark matter data per year are recorded. Data reduction procedures have been developed to parameterise the data; up to 0.6 TB of data per day can be reduced.

Acknowledgements

This work has been funded by the UK Particle Physics and Astronomy Research Council (PPARC), the US Department of Energy (Grant number DE-FG03-91ER40662) and the US National Science Foundation (Grant number PHY-0139065). We acknowledge support from the Central Laboratories for the Research Councils (CCLRC), the Engineering and Physical Sciences Research Council (EPSRC), the ILIAS integrating activity (Contract R113-CT-2004-506222), the INTAS programme (Grant number 04-78-6744) and the Research Corporation (Grant number RA0350). We also acknowledge support from Fun-dação para a Cîncia e Tecnologia (project POCI/FP/FNU/63446/2005), the Marie Curie International Reintegration Grant (Grant number FP6-006651) and a PPARC PIPPS award (Grant PP/D000742/1). We would like to gratefully acknowledge the strong support of Cleveland Potash Ltd., the owners of the Boulby mine, and J. Mulholland and L. Yeoman, the underground facility staff.

References

[1] G.J. Alner, et al., Phys. Lett. B 616 (2005) 17.
[2] G.J. Alner, et al., Astropart. Phys. 23 (2005) 444.
[3] G.J. Davies, et al., Phys. Lett. B 320 (1994) 395.
[4] H. Wang, Ph.D. Thesis, UCLA, 1999.
[5] D.B. Cline, et al., Astropart. Phys. 12 (2000) 373.
[6] J.D. Lewin, P.F. Smith, Astropart. Phys. 6 (1996) 87.
[7] G.J. Alner, et al., Astropart. Phys. 28 (2007) 287.
[8] G.J. Alner, et al., Phys. Lett. B 653 (2007) 161.
[9] ⟨www.polycold.com⟩.
[10] ⟨www.saesgetters.com⟩.
[11] G.J. Alner, et al., 2007, in preparation.
[12] ⟨www.acqiris.com⟩.
[13] ⟨www.msl-datascan.com⟩.
[14] G.F. Knoll, Radiation Detection and Measurement, third ed., Wiley, New York.
[15] ⟨www.adic.com⟩.
[16] ⟨wwwasdoc.web.cern.ch/wwwasdoc/hbook_html3/hboomain.html⟩.