

The ZEPLIN Dark Matter Search: Two Phase Xenon as a WIMP Target

Blair Nicholas Victor Edwards

Department of Physics
Imperial College London

A thesis submitted for the degree of Doctor of Philosophy
in the Department of Physics at Imperial College London

July 2009

Abstract

The existence of a significant non-baryonic component to the Universe is widely accepted, with worldwide efforts underway trying to detect this so-called dark matter. The ZEPLIN detectors utilise liquid xenon as a target medium in the search of the expected rare interactions of Weakly Interacting Massive Particles, or WIMPs, with ordinary baryonic matter. The neutralino, arising in supersymmetric extensions to the standard model of particle physics, provides a particularly well-motivated candidate. The ZEPLIN-II and ZEPLIN-III experiments, operate in two-phase mode (liquid/gas), measuring both the scintillation and ionisation signatures produced during an interaction. These instruments form the basis of this thesis.

The ZEPLIN-II experiment was operated underground at the Boulby Underground Laboratory, culminating in a WIMP search run lasting 57 days. Some key operational aspects are discussed, and a full description of the data analysis is given, which yielded a competitive upper limit on the spin-independent WIMP-nucleon scattering cross-section with a minimum of 6.6×10^{-7} pb for a $65 \text{ GeV}/c^2$ WIMP with 90% confidence. Subsequently, a smaller collaboration proceeded with the commissioning and operation of ZEPLIN-III at Boulby. The detector was operated stably for 12 months, culminating in the first science run, which excluded a cross-section above 7.7×10^{-8} pb for a $55 \text{ GeV}/c^2$ WIMP. This placed ZEPLIN-III as one of the world's leading WIMP search experiments.

Along with the WIMP search results, the data collected from these instruments have been exploited to extract information about the underlying xenon physics processes, which will play an important role in design of future systems. This includes the first quantitative measurements of single electron emission in a two-phase noble gas detector, studies of the field dependence of their response and of the anti-correlation between the scintillation and ionisation channels.

Acknowledgements

I consider myself privileged to have spent the past three (and a bit) years working on such interesting and exciting projects. The search for dark matter is at the forefront of science and I thank those who have given me the opportunity to work in the field. The broad range of work has been enlightening and I have been lucky enough to split my time between two outstanding institutions, Imperial College and the Rutherford Appleton Laboratory.

I would particularly like to thank both Nigel Smith and Henrique Araujo for their support, guidance and friendship throughout my PhD. The optimist/pessimist combination you provided was the perfect blend, and I hope I have found the middle ground, being an optimistic realist. One of the most enjoyable features of my time spent at RAL was the team spirit within the dark matter group, this was fostered through regular trips to the R1 coffee lounge. To this end I would like to thank all the other members of the RAL group during my time: Roland, Roy, Tim, Pawel, Mandy, Vito, Peter and George. Likewise I would like to thank the ZEPLIN team at Imperial who I have worked closely with: Tim, Vadim, Richard, Igor, Markus, Claire and Alastair. I dedicate this thesis to the memory of Vadim Lebedenko, father of ZEPLIN-III. Vadim was a lovely character and a brilliant scientist. I feel honoured to have worked with him (however briefly)!

This team spirit also extended throughout the relatively small collaboration and I have enjoyed working with such a bunch of clever and interesting people. Thanks to the Edinburgh, Coimbra and Sheffield groups particularly for their collaboration and the fun times we had in Whitby and at group meetings. Working at Boulby has been a particular highlight and whilst at the time, it is not always fun getting up for the early rides, it is an experience which will stay with me forever!

The funding for my studies as well as the experiments as a whole was provided by STFC. I would also like to thank Cleveland Potash Ltd who operate Boulby mine for supporting the lab and allowing us to undertake our scientific endeavors.

Finally, I would like to thank the people who have supported me through my studies. My parents who have always encouraged me and made me believe in myself. Thanks to Kulvier for her hospitality. But most importantly of all I would like to thank my long-suffering partner Parvinder, who has put up with me “not having a proper job” and my many trips away. Thank you for your belief, love and patience, I could not have made it through without you!

Contents

Abstract	1
Acknowledgements	3
List of Figures	14
List of Tables	15
1 Introduction	16
2 The Dark Matter Problem	19
2.1 Current state of cosmology	19
2.2 Evidence for Dark Matter	24
2.3 Dark Baryons?	32
2.4 Dark Matter Candidates	38
2.5 Indirect WIMP Searches	44
2.6 Direct Detection Experiments	46
2.7 Summary	58
3 The ZEPLIN programme	60
3.1 Principles of WIMP detection using liquid xenon	61
3.2 Boulby Underground Laboratory	71
3.3 ZEPLIN-I	72
3.4 ZEPLIN-II	75

3.5	ZEPLIN-III	86
4	ZEPLIN-II analysis: From raw data to a WIMP search dataset	93
4.1	Data acquisition and pipeline	94
4.2	Data reduction	97
4.3	Event selection	100
4.4	Corrections	105
4.5	Fiducialisation	113
4.6	Energy calibration	120
4.7	Discrimination parameter calibrations	123
4.8	Science data	134
5	ZEPLIN-III: First Science Run	140
5.1	Data acquisition and pipeline	141
5.2	Pulse identification process	144
5.3	Event selection philosophy	145
5.4	Position reconstruction	152
5.5	Detector monitoring from data	156
5.6	^{57}Co energy calibrations	167
5.7	Dead-region events	168
5.8	Fiducial volume cuts	174
5.9	Discrimination parameter calibrations	176
5.10	Energy conversion and energy-dependent efficiencies	183
5.11	WIMP search data	186
6	Calculation of limits on the WIMP-nucleon scattering cross-section	191
6.1	90% confidence limits on observed event rate	192
6.2	WIMP-nucleon scattering cross-sections	213
6.3	Final results from the ZEPLIN detectors	217

7	Xenon Physics with the ZEPLIN Detectors	221
7.1	Single Electron Emission	222
7.2	Photoionisation as a production mechanism	230
7.3	Nuclear Recoil Ionisation Yield	234
7.4	Scintillation light yield as a function of purity	237
7.5	Field-dependent detector response	239
7.6	Scintillation/Ionisation anti-correlation	243
8	Conclusion	251
8.1	Lessons learned from the ZEPLIN detectors	253
8.2	Future work	254
A	Appendix A: Calibration sources	256
	Bibliography	271

List of Figures

2.1	Constraints on the Ω_Λ and Ω_m density parameters.	23
2.2	Galactic rotation curves for 21 different galaxies.	27
2.3	Rotation curve of NGC 3198 with visible and halo contributions.	28
2.4	Observations of the Bullet cluster merger in visible and x-rays.	29
2.5	Schematic of gravitational lensing.	31
2.6	Hubble deep field image of the Abell 2218 cluster showing image distortion.	31
2.7	Map of dark matter from HST COSMOS gravitational lensing survey. . .	32
2.8	Example of microlensing light curve (in red and blue passbands).	34
2.9	All-sky WMAP 5-year measurement of CMB temperature variations. . .	35
2.10	CMB angular power spectrum from WMAP 5-year data (and other data).	36
2.11	Relative abundances of different nuclei produced during BBN as a function of the baryon density.	38
2.12	Magnitude of gauge coupling constants in the standard model and MSSM.	40
2.13	Positron fraction measured by the PAMELA experiment.	45
2.14	Expected differential energy spectrum for different WIMP-masses. . . .	47
2.15	Depth and muon flux for some major underground labs.	49
2.16	Diagram of energy deposition channels utilised in dark matter searches.	50
2.17	Discrimination scatter plot from the CDMS-II 2008 data.	54
2.18	Calibration of EDELWEISS-I to electron and nuclear recoils.	55
2.19	Annual modulation in residual event rate from DAMA experiments. . .	57

2.20	Current world status of constraints placed on spin-independent WIMP-nucleon cross-section.	59
3.1	Expected event rate for different target species.	62
3.2	Schematic of the processes occurring during an interaction in liquid xenon.	64
3.3	Measurements of the relative scintillation efficiency in xenon.	67
3.4	Measurements of electric-field induced variations of the scintillation and ionisation yields for different interacting particles.	68
3.5	Measurements of electron drift velocity in liquid and solid xenon.	69
3.6	Cross-phase emission probability for electrons in LXe.	70
3.7	Electroluminescence yield measured in both cold xenon vapour and in room temperature gas.	71
3.8	Schematic diagram of the Boulby Underground Laboratory.	73
3.9	Schematic diagram of the ZEPLIN-I instrument.	74
3.10	ZEPLIN-I neutron calibration and the effect of a second population on the electron recoil distribution.	74
3.11	Schematic diagram of the ZEPLIN-II detector.	77
3.12	Schematic diagram of the ZEPLIN-II detector within its shielding castle.	79
3.13	Timeline of electron lifetime in ZEPLIN-II during active recirculation.	81
3.14	High-energy α -decay peak in ZEPLIN-II background data.	82
3.15	Timeline of the α -particle event rate measured in the ZEPLIN-II detector.	83
3.16	Decay of α event rate after the end of recirculation.	84
3.17	Tracking of α event rate, with modeling of production and decay.	85
3.18	Schematic diagram of the ZEPLIN-III instrument.	87
3.19	Schematic diagram of the ZEPLIN-III target region.	88
3.20	Simulation of the electric field in the ZEPLIN-III target.	89
3.21	Identification of coincident events and reconstructed positions.	90
4.1	Typical single-scatter electron recoil event in ZEPLIN-II.	95

4.2	Typical single-scatter neutron recoil event in ZEPLIN-II.	95
4.3	Analysis flow diagram for ZEPLIN-II.	96
4.4	UNZAP2 - The ZEPLIN-II data reduction and visualisation tool.	98
4.5	Clustering algorithm applied to ZEPLIN-II waveforms.	99
4.6	Waveform highlighting the physically allowed regions for S1 and S2 pulses.	102
4.7	ZEPLIN-II S1 trigger efficiency.	103
4.8	The diffusion effect on s2tau and empirical correction in ZEPLIN-II. . .	104
4.9	Positions of S1 and S2 for golden events and trigger cut in ZEPLIN-II. .	106
4.10	Fraction of events removed by a coincidence cut.	106
4.11	Typical single photoelectron area spectra in ZEPLIN-II.	108
4.12	Light collection uniformity in Z for ZEPLIN-II.	109
4.13	Pressure dependence of S2 response as a function of target pressure. . .	110
4.14	Electron lifetime measurement during commissioning phase of ZEPLIN-II.	112
4.15	Electron lifetime monitored during the ZEPLIN-II science data-taking period.	113
4.16	S2/S1 spectrum for events from the bottom of the detector showing cath- ode recoil events.	114
4.17	S2/S1 response of cathode recoil events monitored through the science data-taking period.	114
4.18	Drift time distributions for background and ^{57}Co data from ZEPLIN-II.	116
4.19	Calibration of $x-y$ position reconstruction with ^{57}Co data.	118
4.20	Fiducial cuts shown with the reconstructed positions of science data. . .	120
4.21	Energy calibration spectra of S1 and S2 from ^{57}Co	121
4.22	Scintillation response monitored through the science data-taking period.	122
4.23	Drift time distribution showing coincidence events from ^{60}Co calibrations.	124
4.24	Discrimination parameter scatter plot from Z-II electron recoil calibration dataset.	125
4.25	Discrimination parameter scatter plot from Z-II nuclear recoil calibration dataset.	126

4.26	Fitting of 2 keV _{ee} slices in ZEPLIN-II AmBe neutron calibration data. .	128
4.27	Efficiency plot for the hardware upper-level discriminator (‘saturation cut’).	130
4.28	Total Z-II nuclear recoil detection efficiency, both calculated and measured.	132
4.29	Fits to the electron recoil population from the Z-II ⁶⁰ Co calibration. . .	133
4.30	Discrimination parameter scatter plot of the final Z-II science dataset. .	136
4.31	Radial distributions of small-S2 events in the science dataset.	137
4.32	Radial position resolution as a function of S2 signal size in ZEPLIN-II. .	138
5.1	Display of all 31 high sensitivity waveforms from ZEPLIN-III.	142
5.2	Schematic of the trigger and acquisition circuitry for ZEPLIN-III. . . .	142
5.3	Saturation cut efficiency calculated with ¹³⁷ Cs data.	143
5.4	Event visualisation using ZE3RA software.	145
5.5	Example of amplifier overshoot in ZEPLIN-III waveform.	148
5.6	Example of noisy S2 tail in ZEPLIN-III waveform.	148
5.7	Diffusion effect on s2tau and correction applied.	151
5.8	Drift time distributions for ⁵⁷ Co calibrations and background data. . .	153
5.9	Orientation of ZEPLIN-III position reconstruction coordinates.	154
5.10	Bell-shaped response functions of PMT1 and PMT2 in ZEPLIN-III. . .	155
5.11	Typical electron lifetime measurement from the Z-III first science run. .	158
5.12	Electron lifetime monitored through the Z-III first science run.	158
5.13	Electron lifetime measured as a function of energy.	160
5.14	Levelling plot from the original monitoring routine in ZEPLIN-III. . . .	161
5.15	Historical trend of the phase, amplitude and mid-point of detector tilt. .	161
5.16	Levelling plots from the enhanced monitoring routine, measuring as a function of radius.	162
5.17	Plots for monitoring consistency of channel amplification in Z-III. . . .	163
5.18	Historical trend of S1 and S2 energy calibration values during the first science run.	165
5.19	Improved automated fitting method for ⁵⁷ Co energy calibrations. . . .	165

5.20	Typical differential energy spectrum from a daily sample of science data.	166
5.21	Comparison of ZEPLIN-III ^{57}Co calibration with simulations.	167
5.22	Diagram of the Z-III target geometry highlighting detector dead regions.	169
5.23	Radial position of events from beyond the nominal maximum drift time.	169
5.24	Extrapolation of single scatter drift time distribution.	171
5.25	Comparison of s1tau and s1rms from S1-only and single scatter events. .	172
5.26	MSSI removal cuts and associated efficiencies.	173
5.27	Position of the cathode shown by variation of S2/S1 response.	175
5.28	Radial and drift time fiducial cuts in Z-III science data.	175
5.29	Reconstructed spatial distributions from AmBe calibration data.	177
5.30	Fitting in Log_{10} S2/S1 of AmBe calibration data in 1 keV_{ee} slices. . . .	178
5.31	Discrimination parameter scatter plot of Z-III AmBe calibrations.	179
5.32	Parameterisation of nuclear recoil mean and sigma as a function of energy.	179
5.33	Reconstructed spatial distributions from ^{137}Cs calibration data.	180
5.34	Discrimination parameter scatter plot showing both electron and nuclear recoil calibration data.	180
5.35	Discrimination parameter scatter plot from 10% of science data.	182
5.36	Comparison of AmBe calibration data with spectrum from simulations.	184
5.37	Energy conversion factor calculated from comparison method and the associated efficiency spectrum.	186
5.38	Examples of fitted S2/S1 distributions from the WIMP search dataset. .	187
5.39	Comparison of mean and sigma of electron recoil populations from ^{137}Cs calibration and WIMP search data.	188
5.40	Reconstructed spatial distributions of full science dataset.	188
5.41	Final discrimination scatter plot of the full ZEPLIN-III science dataset.	190
6.1	Neymann construction for 90% confidence limits of a Poisson process with known background, using the Feldmann-Cousins ordering principle. . . .	194
6.2	Neymann diagram for the Z-II observation of 29 events with 28.6 expected.	196

6.3	Illustration of the energy scale conversion for the Yellin interval methods.	198
6.4	Final discrimination scatter plot from the XENON10 experiment.	199
6.5	Final discrimination parameter scatter plot of the ZEPLIN-III science data in traditional coordinate space and NR % space.	202
6.6	Example probability density functions for background and signal with the observed ZEPLIN-III events.	204
6.7	Combined μ distributions for different signal and background levels. . .	205
6.8	Cumulative $\ln \mathcal{L}$ distributions with increasing signal levels.	206
6.9	$\ln \mathcal{L}$ distribution from MC simulation for $S = 0$	206
6.10	Probability surfaces for background and signal.	208
6.11	Colour map of $\ln \mathcal{L}$ distribution, varying β and S	212
6.12	Differential WIMP recoil spectrum before and after application of detector characteristics	216
6.13	World status of constraints on $\sigma_{\text{w-n}}$ at Z-II publication (Jan 2007). . . .	218
6.14	Current world status of constraints on $\sigma_{\text{w-n}}$	219
7.1	Sum waveform showing single electrons after an S2 signal.	222
7.2	Sum waveform containing a candidate single electron pulse.	224
7.3	Single electron area spectrum at 1.5 bar gas pressure.	225
7.4	Dependence of single electron response on detector gas pressure.	227
7.5	Radial and depth distributions for single electrons.	229
7.6	Predicted radial distributions for small S2 signals.	229
7.7	Fraction of S1-triggered events where a single electron is observed as a function of S1 area.	230
7.8	Fraction of S1-triggered events where a single electron is observed as a function of electron lifetime.	231
7.9	Measurements of the photoionisation cross-section for O_2^- ions.	233
7.10	ZEPLIN-III waveform containing a single electron candidate.	234
7.11	Typical single electron emission spectrum from ZEPLIN-III.	235

7.12	Nuclear recoil ionisation yield measurements including data from Z-II and Z-III.	237
7.13	Relationship between light yield and electron lifetime in ZEPLIN-II. . .	238
7.14	Onset of field suppression measured with S1 response.	240
7.15	Dependence of S1 signal with electric field in ZEPLIN-III.	241
7.16	Parameterisations of electroluminescence yield, emission efficiency and relative charge extraction.	242
7.17	Dependence of S2 response with electric field in ZEPLIN-III.	244
7.18	Energy spectrum for ^{57}Co in S1 and E^*	246
7.19	Colour density plots of S2/S1 against energy, for ^{57}Co and ^{60}Co with energy described by S1 and E^*	246
7.20	Spectra of the 40 keV $_{ee}$ inelastic scattering feature in S1, E^* and from simulations.	247
7.21	Relative energy resolution from ^{57}Co in E^* for varying k	249
7.22	Flow diagram of channelling of deposited energy.	249
7.23	Relative energy resolution from ^{57}Co in E^* , varying η_{S1} at two different electric fields.	250
8.1	Constraints on WIMP-nucleon scattering before this research compared with now.	252
A.1	Interaction length as a function of energy for γ -rays.	256
A.2	^{57}Co decay scheme.	257
A.3	^{60}Co decay scheme.	258
A.4	^{137}Cs decay scheme.	258
A.5	AmBe neutron energy spectrum.	259

List of Tables

2.1	Properties of noble gases of interest for dark matter targets.	51
3.1	Radon decay chain	83
4.1	ZEPLIN-II centroid reconstruction parameters.	117
4.2	Nuclear recoil detection efficiencies in ZEPLIN-II science data.	131
4.3	Summary of ZEPLIN-II exposure for science background data-set.	134
4.4	Summary of background expectation in the ZEPLIN-II WIMP search box.	139
5.1	Parameters recorded in <i>hbook</i> ntuples by ZE3RA.	146
5.2	Parameters from the <i>Mercury</i> reconstruction algorithm	155
5.3	Summary of calculated efficiencies	185
6.1	Expectation values from the ZEPLIN-II science data.	196
6.2	Skew-Gaussian fit parameters from slicing electron background data.	211
7.1	Experimentally measured electroluminescence yield constants, a and b	226
A.1	γ -rays emitted in the decay of ^{57}Co . [1]	257
A.2	γ -rays emitted in the decay of ^{60}Co . [1]	258
A.3	γ -rays emitted in the decay of ^{137}Cs . [1]	259

Chapter 1

Introduction

Over the past 70 years evidence has been mounting for the existence of a significant amount of mass in the Universe which is invisible in the electromagnetic spectrum, but nevertheless makes its presence felt by its gravitational interactions. Studies of the potential candidates for this missing mass have ruled out a significant unseen baryonic component, with constraints placed on the baryon fraction from measurements of the cosmic microwave background (CMB), the theory of primordial nucleosynthesis and direct statistics in our galaxy. This means that a significant non-baryonic component is required, perhaps in the form of a new elementary particle, with studies of large scale structure and the evolution of the Universe pointing to a Cold Dark Matter (CDM) scenario within the standard Λ CDM cosmological model. The currently favoured class of candidates for CDM is the Weakly Interacting Massive Particle (WIMP), with the neutralino, the lightest particle predicted by many supersymmetric theories, providing a well-motivated WIMP candidate.

Most evidence for dark matter amassed so far consists of observations of its gravitational effects on baryonic objects. The next step required is to detect directly the rare WIMP-baryon interactions and to investigate the properties of galactic WIMPs. To this end, many different techniques are being implemented across the world. One of the leading experimental techniques is the use of liquid noble gases, with detectors utilising two-phase xenon providing some of the world's leading constraints. The evidence for dark matter, potential candidates and detection techniques are reviewed in Chapter 2.

The focus of this thesis is the ZEPLIN dark matter search, which uses two-phase

xenon as a target for WIMP interactions. Chapter 3 provides a description of how radiation and particles interact in liquid xenon, as well as describing both the ZEPLIN-II and ZEPLIN-III instruments. This chapter also includes studies carried out personally investigating emanation of ^{222}Rn and improvement in xenon purity in ZEPLIN-II. The data analysis of the two experiments is presented in Chapters 4 and 5, before calculations of the constraints placed on the WIMP-nucleon scattering cross-section are detailed in Chapter 6. Finally, Chapter 7 presents additional studies, using ZEPLIN data, of the underlying xenon physics at the core of the two-phase technique, including the first quantitative analysis of single electron emission in a two-phase noble gas detector. Such studies are vital for the design and operation of future instruments.

The ZEPLIN programme is a collaborative endeavour with many people and several institutions involved in the development, operation and science exploitation of the experiments. Whilst the work presented here would not have been possible without the hard work of all involved, the project was nevertheless sufficiently small to allow me to play an important role in many different aspects of the programme. These aspects are highlighted throughout this thesis where relevant. My main contributions focused on data analysis for the two experiments, including novel xenon physics studies and development of limit setting procedures.

Within the ZEPLIN-II analysis I personally worked on event selection, calculation of corrections, implementation of the centroid position reconstruction algorithm and analysis of the calibrations and final science data. Progressing to ZEPLIN-III I utilised the knowledge I had gained from the previous analysis to help debug the reduction code, define event selection cuts, implement basic position reconstruction (before others developed more advanced techniques), produce an automated data monitoring suite, investigate detector dead-regions and analyses the final data-sets. Calculation of the ZEPLIN-II limit was done to cross-check the result calculated by others, before development of potential new techniques for ZEPLIN-III and the calculation of the resulting dark matter limit. Finally, I was able, personally, to exploit the data from the experiments to investigate features of two-phase xenon as a detection technique.

The ZEPLIN-II experiment was operated under the umbrella of the UK Dark Matter Collaboration (UKDMC), which consisted of Imperial College London, University of Sheffield, University of Edinburgh and the STFC Rutherford Appleton Laboratory,

along with Portuguese LIP-Coimbra and US collaborators at UCLA, Texas A&M and University of Rochester. The development and operations of ZEPLIN-III were furthered by a smaller collaboration incorporating Imperial College London, STFC Rutherford Appleton Laboratory, University of Edinburgh, LIP-Coimbra and ITEP Moscow.

Chapter 2

The Dark Matter Problem

The quest to understand the world around us stretches back to ancient times, a challenge which continues today with particle physics and cosmology exploring the Universe on the smallest and largest scales. Since the 1930s evidence has been building for the presence of “missing” mass in the Universe, with many different approaches converging on the same conclusions.

A century ago it was believed that luminous matter in the form of stars contained nearly the entire mass of the Universe. Today’s astronomical observations determine that luminous matter contributes only $\sim 1\%$ of the total mass-energy content of the Universe. Further, measurements of the different components of energy-mass density imply that baryonic matter only accounts for about 4% of the total. However, evidence suggests a total matter component of $\sim 30\%$, implying that $\sim 26\%$ is comprised of a non-luminous, weakly interacting, non-baryonic form, termed “dark” matter. In this chapter I will briefly review our current understanding of cosmology, lay out the evidence for the presence of dark matter, discuss possible candidates and potential direct detection methods currently being applied to the problem.

2.1 Current state of cosmology

There are two dominant principles underlying the current understanding of cosmology: the cosmological principle and the hot Big Bang theory [2]. The cosmological principle states that the Universe is both isotropic and homogeneous. This may appear to be

inaccurate when considering terrestrial or galactic scales, but on scales above 100 Mpc the distribution of clusters of galaxies indeed appears to be relatively uniform. In this way, it also fits the Copernican principle, that we (as observers) do not possess a privileged position within the cosmos. The hot Big Bang theory states that the Universe grew from an initial singularity, expanding to create the Universe we observe today.

In the early part of the last century Edwin Hubble first discovered galaxies outside the Milky Way. He proceeded to measure the distance to many galaxies along with their velocities, finding that they were moving away from us with a velocity proportional to their distance, defining this expansion by:

$$\vec{v} = H_0 \vec{r} , \quad (2.1)$$

where \vec{v} is the recessional velocity, \vec{r} is the distance to the galaxy and H_0 is the Hubble parameter, a numerical measure of the expansion rate, often expressed as:

$$H_0 = 100h_0 \text{ km/s/Mpc} , \quad (2.2)$$

where current observations determine the reduced Hubble parameter, $h_0 = 0.72 \pm 0.08$ [3].

This expansion implies that in the past galaxies were much closer together. Taken to its ultimate conclusion, this fact yields the hot Big Bang scenario, where the Universe expands and cools from an initial state of extreme density and temperature. Hubble's observations also fit the cosmological principle as, in a uniformly expanding Universe, an observer will view all objects receding.

Considering the expanding Universe, the rate of expansion can be described by the scale factor, $a(t)$,

$$\vec{r} = a(t) \vec{x} , \quad (2.3)$$

relating the real distance between two points, \vec{r} , to their separation in a co-moving coordinate system (expanding with the Universe). The scale factor is assigned to unity for the present day, $a_0 = 1$. This allows the Hubble parameter to express the expansion

rate thusly:

$$H(t) = \frac{\dot{a}(t)}{a(t)}, \quad (2.4)$$

Under the cosmological assumptions of isotropy and homogeneity, the expansion of the Universe can be described by the Friedmann-Robertson-Walker (FRW) metric, an exact solution of the Einstein Field Equations (EFEs):

$$ds^2 = c^2 dt^2 - a^2(t) \left[\frac{dr^2}{1 - \kappa r^2} + r^2(d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (2.5)$$

where s is the proper distance between the points in space-time, r , θ and ϕ are spherical polar co-moving coordinates, the time dependence is provided by the scale factor, $a(t)$, and κ is a constant describing the curvature of space-time. The curvature has three possible values, $\kappa = -1, 0, +1$, for negative curvature, a flat Universe or positive curvature.

The theory of General Relativity developed by Einstein describes the equivalence of gravitation and the curvature of space-time. General relativity and its importance within cosmology are reviewed in [4, 5]. The Einstein Field Equations describe how the curvature of space-time (expressed by the Einstein tensor, $G_{\mu\nu}$) is governed by the mass-energy content within it (expressed by the energy-momentum tensor, $T_{\mu\nu}$). The field equations can be written as:

$$G_{\mu\nu} - \Lambda g_{\mu\nu} = -\frac{8\pi G}{c^4} T_{\mu\nu} \quad (2.6)$$

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R - \Lambda g_{\mu\nu} = -\frac{8\pi G}{c^4} T_{\mu\nu}, \quad (2.7)$$

where the c is the speed of light and G is the gravitational constant. The Ricci curvature tensor, $R_{\mu\nu}$, the scalar constant, R , the metric tensor, $g_{\mu\nu}$, and the cosmological constant, Λ , determine the geometry of space-time, with the right hand side of Eq. 2.7 describing the energy-mass content. The cosmological constant, Λ , acts as a negative pressure component boosting the expansion of the Universe. This component is now associated with vacuum energy or dark energy, and is believed to constitute the majority of the mass-energy content ($\sim 70\%$) [6].

The Friedmann equations are derived from the EFE, the FRW metric and by assuming that on large scales we can describe the matter in the Universe as a fluid, with density ρ and pressure p , related by the equation of state:

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{\kappa c^2}{a^2} - \frac{\Lambda}{3} = \frac{8\pi G}{3}\rho, \quad (2.8)$$

where both the radiation and matter components are incorporated (such that $\rho = \rho_m + \rho_\gamma$). The equations can be simplified by incorporating the cosmological constant into the density:

$$\rho \rightarrow \rho + \frac{\Lambda c^2}{8\pi G}, \quad (2.9)$$

which thus yields:

$$\underbrace{H^2}_{\text{Expansion}} + \underbrace{\frac{\kappa c^2}{a^2}}_{\text{Curvature}} = \underbrace{\frac{8\pi G}{3}\rho}_{\text{Density}}, \quad (2.10)$$

where the terms describe the expansion, curvature and energy-mass density of the Universe, respectively.

The flat Euclidean Universe ($\kappa = 0$) will occur for a certain density, known as the critical density, ρ_c . From Eq. 2.10, with $\kappa = 0$, the critical density is:

$$\rho_c = \frac{3H^2}{8\pi G}. \quad (2.11)$$

A density greater than ρ_c results in a closed Universe, which will eventually recollapse, and conversely a density below ρ_c gives an open Universe, which will expand forever. A density parameter Ω can be defined as the ratio of the actual density to the critical density:

$$\Omega \equiv \frac{\rho}{\rho_c} = \frac{8\pi G}{3H^2}\rho, \quad (2.12)$$

where H and ρ are both a function of time resulting in an epoch-dependent density parameter, with $\Omega = 1$ representing a flat Universe. Ω can thus be divided into its constituent parts for radiation (Ω_r), matter (Ω_m) in both baryonic (Ω_b) and non-baryonic forms (Ω_d) and dark energy (Ω_Λ). These different density components can be deter-

mined with various experimental and theoretical techniques, allowing for a cosmological picture of the Universe to be developed.




Figure 2.1: Constraints on the Ω_Λ and Ω_m ($= \Omega_m + \Omega_\Lambda$) density parameters [7], displaying the convergence on the values favoured by the Λ CDM model. Constraints from CMB, Type Ia supernovae and galaxy clusters are shown, along with a line representing $\Omega = 1$, defining an open, flat or closed Universe and regions which exclude the hot big bang scenario.

Figure 2.1 shows the constraints placed on Ω_m and Ω_Λ by different observations. Observations of Type Ia supernovae have demonstrated that the expansion of the Universe is accelerating, with the negative pressure from dark energy overpowering the gravitational forces attempting to slow the expansion. Measurements of the CMB and of galaxy clusters also help to place constraints on the matter (both baryonic and non-baryonic) and dark energy components, leading to the currently favoured model, usually referred

to as the Λ Cold Dark Matter (Λ CDM).

The Λ CDM model is often known as the concordance model as it incorporates observations from the CMB, Type Ia supernovae and large-scale structure formation. The Λ CDM model is relatively simple, yet is able to explain and predict all the observations to date. The model consists of the cosmological constant, describing dark energy and matter (in both baryonic and non-baryonic forms), plus a small radiation component. Observations (many of which are described later) allow constraints to be placed on the amounts of each, with the current measured values [8]:

Radiation:	$\Omega_r = 8.4 \times 10^{-5}$
Baryonic matter:	$\Omega_b = 0.0456 \pm 0.0015$
Dark (non-baryonic) matter:	$\Omega_d = 0.228 \pm 0.013$
Total matter:	$\Omega_m = \Omega_b + \Omega_d = 0.2736 \pm 0.013$
Dark energy:	$\Omega_\Lambda = 0.726 \pm 0.015$

2.2 Evidence for Dark Matter

Predictions for the existence of non-luminous matter in the Universe were made as far back as the 1840's. In 1844, Friedrich Bessel predicted the existence of Sirius B by observing the motion of Sirius. Around the same time the presence of Neptune was inferred through observations of anomalous perturbations in the orbit of Uranus. In a similar manner, observations of the motion of visible matter have led to the prediction that a significant amount of non-luminous matter, now known as dark matter, exists in galaxies and clusters.

Evidence for the existence of dark matter has been building since the 1930's [9, 10], with it later becoming clear that baryonic matter alone accounts for a small fraction of the Universe. The baryonic matter component (Ω_b) accounts for all the visible mass of the Universe along with gas, dust and non-luminous astrophysical bodies (such as planets or brown dwarfs). This has led, with mounting evidence, to a picture of the cosmos which includes a substantial amount of dark matter (Ω_d), which has, so far, only manifest itself through gravitational interactions.

2.2.1 Dark matter in galaxies and clusters

The first observations suggesting non-luminous matter on a large scale were carried out by Jan Oort [9], who observed the motions of stars in the Milky Way disk. By measuring the velocity of stars moving vertically (as well as rotationally) through the galactic disk, he was able to calculate the mass required to keep the disk stable. He calculated a mid-plane density of $\rho_0 \sim 0.09 \text{ M}_\odot \text{pc}^{-3}$ (where M_\odot is equal to one solar mass), hence known as the Oort limit, which compares with a luminous value of $\rho_0 \sim 0.03 \text{ M}_\odot \text{pc}^{-3}$, suggesting 3 times more matter than observed. Much more recent calculations by Bahcall (1984) reduced the disparity between the calculated and observed (luminous) values, before further measurements [11, 12] accounted for the required mass in baryonic forms. The balance of evidence seems to suggest no significant dark matter component within the galactic disk, although this is still a point of discussion.

In 1933, Fritz Zwicky argued that traditional methods of measuring galactic masses (photometry and rotation) were inaccurate and highly biased. He asserted that the virial theorem provided a much more accurate measurement of the mass of galaxy clusters. The virial theorem gives a relationship between the time-averaged total kinetic (T) and potential (V) energies in a closed, stable system: $2T + V = 0$, allowing calculation of the mass of a cluster through measurement of its velocity dispersion. Zwicky applied the virial theorem to the Coma cluster, finding that the virial mass was about 400 times greater than that predicted by photometry, suggesting a large fraction of non-luminous matter [10]. This factor has been reduced by recent measurements of the cluster ($M_{\text{cluster}} = 1.6 \times 10^{15} \text{ M}_\odot$ [13] with $M_{\text{galaxies}} \sim 7 \times 10^{13} \text{ M}_\odot$ and $M_{\text{gas}} \sim 9.6 \times 10^{13} \text{ M}_\odot$ [14]), although a significant fraction of dark matter is still implied. Zwicky later coined the term “dark matter” and suggested the potential of gravitational lensing as a viable technique to measure the mass of galaxies [15].

Another significant piece of evidence is based upon the timing argument for calculation of the mass of the Milky Way. According to the Hubble expansion of the Universe, the Milky Way and M31 should be moving apart. However, measurements showed that they are approaching one another with a velocity of $\sim 125 \text{ km/s}$ (at a distance of $\sim 740 \text{ kpc}$). No massive objects were found in their past paths which could have deflected their trajectory, leaving gravitational attraction as the only explanation.

Kahn and Woltjer [16] used this method to estimate the mass of the two-galaxy system finding it to be about 6 times larger than expected from visible matter alone. Similar work by Zaritsky *et al* [17] using the timing argument for the Milky Way with M31 and Leo I arrived at the same conclusion of a significant dark matter component.

Simulations of structure evolution provide a useful tool for determining the composition of galaxies and the Universe from the structure observed today. The first example of this was a compelling study by Ostriker and Peebles (1973) [18], in which they simulated a disk of stars and followed the evolution of the disk under rotation. They observed that above a certain rotation rate all galaxies would form bar instabilities. Such galaxies do exist in the Universe, although they are relatively rare and the majority of galaxies with a sufficient rotation rate show no bar formation. Ostriker and Peebles discovered that introducing a spherical distribution of matter, with significant mass (above 3 times the mass of the disk), around the disk stabilised the evolution and greatly reduced the proportion developing a bar, a finding that spurred on the search for dark matter in the form of spherical halos.

An explosion of work on dark matter ensued, with a multitude of discoveries in the late 1970s and early 1980s. The most well known and convincing piece of observational evidence for dark matter is the measurement of galactic rotation curves. This began in 1939 with work by Horace Babcock, who attempted to map the mass distribution in the galactic disk of M31. Instead of the expected decrease in mass at large radii, matching the light profile, the mass appeared to increase [19]. His observations were largely overlooked by the astronomical community until higher precision measurements were available.

Further work by Rubin & Ford [20] and Roberts & Whitehurst [21] confirmed the flat rotation curve with higher precision. Subsequent studies of many other galaxies yielded flat rotation curves (as shown in Figure 2.2), providing strong evidence for an additional component to the mass profile of the galaxy, with the bulge and disk components unable to account for the observed curves. The universality of the shapes observed suggests the effect to be fundamental in the composition and dynamics of spiral galaxies. Additionally, the high precision definition of a rotation curve allows for the composition of the galactic mass profile to be assessed.




Figure 2.2: Galaxy rotational velocity as a function of distance from the galactic centre for 21 different galaxies [22].

Spiral galaxies usually have two visible components, the disk and the bulge. The disk usually extends to an optical radius of about 10 kpc (for an average galaxy), whereas the bulge usually has a radius smaller than 1 kpc. As a result, at large radii the bulge provides little contribution to the gravitational potential of the galaxy. By assuming a circular orbit around the galactic centre, the rotation velocity of a star can be calculated by equating the gravitational and centrifugal forces acting upon it:

$$\frac{GmM_r}{r^2} = \frac{mv^2}{r}, \quad (2.13)$$

where M_r is the mass contained within the radius r . From this the rotational velocity follows:

$$v(r) = \sqrt{\frac{GM_r}{r}}. \quad (2.14)$$

For the inner-most region of the galaxy, near the bulge, $M_r = \rho \times \frac{4}{3}\pi r^3$, assuming a spherically-symmetric bulge of uniform density, ρ . This results in a rotation velocity increasing linearly with radius, $v(r) \propto r$. However, at large radii, beyond the majority of the visible disk, the complete visible mass lies within the radius, so one would expect $v(r) \propto 1/\sqrt{r}$. However, the measured rotation curves show no evidence of a decrease in rotation velocity at large radii, well beyond the visible disk. The flat curves observed ($v(r)$ constant) imply a linear increase in mass with radius, $M_r \propto r$. This result provides strong evidence for an extensive halo of non-visible matter beyond the visible disk of




Figure 2.3: Rotation curve of NGC 3198 with contributions from the visible disk and dark matter components shown [23].

the galaxy, in good agreement with the postulation from Ostriker and Peebles.

Similarly, evidence is found for a significant dark matter component in elliptical galaxies. Under the assumptions of a spherically symmetric galaxy and hydrostatic balance, measurements of the temperature and density profiles allow for determination of the galactic mass. Such studies of M87 [24] suggest that 99% of its mass is in the form of dark matter. Further evidence using these techniques is described in [25]. Additionally, it is known that elliptical galaxies contain a significant amount of hot x-ray emitting gas [26], which implies temperatures of about 10^7 or 10^8 K, suggesting a velocity for the gas particles significantly greater than the escape velocity of the galaxy derived from visible matter alone. For the gas to be gravitationally bound within the galaxy, a large dark matter component is required to provide the extra mass.

Observations of the Bullet cluster by the Hubble Space Telescope (HST) in the optical, and Chandra in x-rays [27], provide unique empirical evidence for dark matter. The Bullet cluster, a merger of two galaxy clusters, provides a unique opportunity to observe the interactions of the different cluster components during the collision. Figure 2.4 shows the two merging clusters after they have passed through one another, with the visible galaxies remaining as two distinct clusters. This is expected as the

individual galaxies will not interact due to their sparse nature. However, the gas within the clusters (seen in Figure 2.4 from its x-ray emissions) does interact during the merger and therefore lags behind the visible galaxies. In the absence of a dark matter component the gravitational potential should follow the dominant mass component, which is by far the x-ray emitting gas. Weak-lensing measurements allow the gravitational potential to be measured in the system (shown by the contours in Figure 2.4). These clearly follow the visible galactic component of the clusters, rather than the gas. This segregation of the dominant baryonic component and the gravitational potential provides extremely strong evidence for a significant non-baryonic dark matter component. Additionally, the strong spatial separation of the observation means that explanations of this effect based solely on modified gravity theories are simply not viable.




Figure 2.4: Observations of the Bullet cluster merger in visible (left) and x-rays (right), with the gravitational potential, mapped through weak-lensing, shown by the green contours [27].

2.2.2 Sunyaev-Zeldovich effect

Another method for measuring the presence of the x-ray emitting intra-cluster gas, in addition to its thermal bremsstrahlung, utilises its effect on radiation from background sources. CMB photons passing through the hot ionised gas will be inverse Compton scattered, distorting the expected shape of the CMB blackbody spectrum, with some photons being shifted to higher energies; this is known as the Sunyaev-Zeldovich (SZ) effect. This is an excellent complement to the measurements of mass from thermal bremsstrahlung x-ray emission, as the magnitude is dependent upon n (the density of absorbers) compared with n^2 for x-ray emission.

The majority of the mass within galaxy clusters is accounted for by intra-cluster gas (as in the Bullet cluster), making the gas mass fraction a good measure of the baryonic matter component. The gas mass fractions of 38 massive galaxy clusters have been measured using the SZ effect and their x-ray emission [28], finding an excellent agreement: $f_g(\text{x-ray}) = 0.119 \pm 0.003^{+0.007}_{-0.014}$ and $f_g(\text{SZ}) = 0.121 \pm 0.005^{+0.009}_{-0.016}$. This further demonstrates the need for a non-baryonic dark matter component.

2.2.3 Gravitational Lensing

As light travels through the Universe, general relativity dictates that it follows geodesics in curved space, meaning that its path will be curved when passing near a massive object. Gravitational lensing is the exploitation of this effect, whereby the light from a distant astronomical source is distorted by an intermediate massive object, as shown in Figure 2.5. There are three distinct forms of gravitational lensing:

1. Strong lensing - significant distortions, such as arcs, multiple images or Einstein rings, are visible.
2. Weak lensing - distortions are less well pronounced, but information can be gained by analysing the distortion of a number of objects.
3. Microlensing - no distortion of the shape, but focussing of light can brighten the image from a background source, indicating the presence of an intermediate object.

It was Fritz Zwicky who first proposed the application of gravitational lensing to measure the mass of galaxy clusters [15]. Now the technique is widely used and plays an important role in modern astrophysics and cosmology. Strong lensing was the first method utilised with images such as that in Figure 2.6 exhibiting obvious distortion, with multiple images and arcs. One of the many collaborations carrying out gravitational lensing surveys to measure the mass of clusters, the Cosmic All-Sky Survey (CLASS), have utilised the method to measure the total matter density of the Universe: $\Omega_m = 0.31^{+0.27}_{-0.14}$ [29]. Although the uncertainties are large, the mean Ω_m values published by this and other surveys indicate significant dark matter fractions.

Weak lensing relies upon the analysis of the small distortions observed in a number of background sources to calculate the gravitational potential of a foreground object.




Figure 2.5: Schematic of gravitational lensing showing light from a source bent around a massive intermediate object before being observed. *Image credit: Bell Labs, Lucent Technologies*




Figure 2.6: Hubble deep field image of the Abell 2218 cluster, with light from background galaxies distorted by the foreground cluster, forming multiple images and arcs. *Image credit: NASA*

This technique has been used on multiple occasions to demonstrate the presence of dark matter, most notably in the case of the Bullet cluster [27]. Weak lensing has recently been utilised to map the dark matter density in the Universe, providing a new measurement of large scale structure. By studying the HST Cosmic Evolution Survey (COSMOS), utilising the dependence of the lensing profile on the distance from the lensing object, the gravitational large scale structure is resolved in both angle and time [30], creating a three-dimensional map of the dark matter (shown in Figure 2.7).




Figure 2.7: Map of dark matter in the Universe with gravitational lensing observed in the HST COSMOS survey [31]. *Image credit: NASA*

2.3 Dark Baryons?

The astrophysical observations described previously demonstrate the need for a significant fraction of matter in non-luminous forms. For a time, the possibility remained that this missing mass could be accounted for by astrophysical objects, formed of baryons, but in a compact form which cannot be readily observed. Possible candidates were white dwarves, neutron stars, brown dwarves and black holes, commonly known as MAssive Compact Halo Objects (MACHOs).

For a while, MACHOs were believed to contribute a significant proportion (if not the dominant component) of the missing mass. However, there are several major pieces of evidence against dark baryons contributing significantly to the overall mass content. These include: micro-lensing MACHO searches, measurement of the baryonic compo-

ment from the CMB and Big Bang nucleosynthesis (BBN). These methods all point independently to the total baryonic component of the Universe being only $\sim 5\%$.

2.3.1 Microlensing

MACHOs are too small and too dim to be observed directly, but their presence can be detected through their gravitational potential using the technique of microlensing [32, 33]. Small gravitational lenses produce a focussing effect, equivalent to multiple unresolvable images, resulting in an enhancement of the light observed from a background source, rather than significant image distortion. The characteristic symmetric, achromatic light curve signature, shown in Figure 2.8, as a MACHO crosses the line of sight to the background source can be used to distinguish it from other astrophysical sources and determine the mass of the lensing object.

Searches for microlensing objects help to place limits on both the number of MACHOs and their properties. The results of such searches by the MACHO and EROS collaborations monitoring millions of stars in the Large and Small Magellanic clouds have indicated that MACHOs do not contribute a significant fraction of the halo mass, observing far fewer MACHO candidates than would have been expected for the measured halo mass [35].

2.3.2 Cosmic Microwave Background

According to the Big Bang theory, the early matter-dominated Universe consisted of a hot plasma, emitting equilibrium blackbody radiation. Whilst the temperature was above ~ 3000 K ($t \sim 300,000$ years), electrons were dissociated from atomic nuclei, leaving an ionised Universe containing photons, free electrons and baryons. During this hot, dense epoch, photons underwent continuous Thomson scattering from the free electrons, resulting in an opaque Universe. However, as the Universe cooled below the ionisation energy of hydrogen, recombination occurred, with electrons becoming bound into atoms. At this moment, photons became decoupled from matter and the Universe became transparent. The last scattering photons then became a blackbody relic of the Big Bang, permeating the Universe with background radiation. These photons red-shifted as the Universe expanded and cooled, being observed today as the cosmic




Figure 2.8: Light curve (in red and blue passbands) observed in a microlensing event, with the theoretical microlensing model fitted to both colours. It is seen that the ratio of the two colours remains constant [34].

microwave background (CMB) at a blackbody temperature of 2.7 K.

The existence of such a background was first predicted by Gamow, Alpher and Herman in the 1950s [36], but was not discovered until the 1960s. An excess noise was observed by Penzias and Wilson during the operation of a new telescope at Bell Laboratories, with a temperature of 3.5 ± 1 K [37]. Without realising, they had discovered the relic blackbody radiation produced by the early Universe [38]. Early measurements of the CMB showed it to be isotropic with a constant temperature, although it was realised that some perturbations, at the level of 1 in 100,000, must be present to evolve into the large scale structure observed in the Universe today.

Data from the COBE (Cosmic Background Explorer) satellite proved the blackbody nature of the CMB, measuring a temperature of $T = 2.725 \pm 0.002$ K [40], with the




Figure 2.9: All-sky WMAP 5-year measurement of variations in the temperature of the CMB [39].

data from balloon-borne experiments (BOOMERANG [41] and MAXIMA [42]) and from WMAP (the Wilkinson Microwave Anisotropy Probe) confirming this. The small temperature perturbations expected were indeed discovered and they now provide an important tool for cosmology. Figure 2.9 shows an all-sky map of the tiny temperature variations measured from the 5 year WMAP data [39]. Measurements of these anisotropies on different scales can provide information about the early Universe and the conditions at the time of recombination. A power spectrum can be calculated characterising the CMB anisotropies, shown in Figure 2.10, with the peaks in the spectrum determined by the size, composition and state of the early Universe.

Quantum fluctuations in the hot plasma grow into gravitationally-driven oscillations between regions with over-density and under-density, known as baryon acoustic oscillations (BAO). The magnitude of these oscillations depends on the total matter content, allowing constraints to be placed on Ω_m . During such oscillations, baryons are additionally acted upon by radiation pressure, producing a damping effect, allowing constraints to be placed on Ω_b . The results of the WMAP 5 year data find no deviation from the




Figure 2.10: CMB angular power spectrum from WMAP 5-year data [43], along with measurements from the ACBAR [44], Boomerang [41] and CBI [45] experiments. The red line is the best-fit Λ CDM model to the WMAP data.

Λ CDM model to 99% confidence [8], measuring the following parameters [46]:

$$\begin{aligned}\Omega_b h^2 &= 0.02273 \pm 0.00062 \\ \Omega_d h^2 &= 0.1099 \pm 0.0062 \\ \Omega_\Lambda &= 0.742 \pm 0.030 \\ \text{with } h &= 0.719^{+0.026}_{-0.027}\end{aligned}$$

which equate to: $\Omega_b = 0.04397^{+0.00199}_{-0.00204}$, $\Omega_d = 0.2126^{+0.0143}_{-0.0144}$.


2.3.3 Primordial nucleosynthesis

One second after the Big Bang the Universe was extremely hot, keeping all the protons and neutrons free and unbound. As the Universe expanded and cooled, the energies of the free nucleons decreased and eventually at $t \sim 1$ minute ($T \sim 3 \times 10^9$ K), they began to bind forming nuclei. When protons and neutrons left thermal equilibrium, the n/p ratio was determined by the difference in their masses, implying a ratio $\simeq 1/6$. However, in the minute between freeze-out and binding, a fraction of the neutrons decayed ($t_{1/2} \sim 10$ minutes) resulting in a n/p ratio $\sim 1/7$.

Production of the light elements, known as Big Bang nucleosynthesis (BBN) began with the production of deuterium:



Deuterium is relatively easy to destroy and subsequently, deuterium has a strong tendency to form ${}^4\text{He}$ through the following processes:



Whilst the favoured end state for these processes is the stable ${}^4\text{He}$ nucleus, the reactions also produce ${}^3\text{He}$ and tritium (${}^3\text{H}$). The production of heavier elements is inhibited by the lack of stable nuclei with mass 5 or 8, with only very small amounts of ${}^7\text{Li}$ being produced. The continued expansion of the Universe also inhibited the process by decreasing the nucleon number density and precluding the possibility of 3α reactions (which do occur in stars). From the understanding of the n/p ratio and the reactions involved the relative abundances of the different nuclei can be calculated, with these varying with the baryon-photon ratio, as shown in Figure 2.11.

Measurements of the elemental abundances can be made through observations of the light absorption by intergalactic primordial gas clouds. The observed abundances allow the baryon-photon ratio to be constrained (as shown in Figure 2.11) which, in combination with measurements of the photon density from the CMB, allow calculation of the baryon density:

$$n_b = 1.13 \times 10^{-5} \Omega_B h^2 \text{ cm}^{-3}, \quad (2.16)$$

which equates to a constraint on the baryonic matter fraction of:

$$0.021 \leq \Omega_B \leq 0.072 \quad (2.17)$$

The ability to explain the abundances of the light elements by primordial nucleosyn-




Figure 2.11: Relative abundances of different nuclei produced during big bang nucleosynthesis as a function of the baryon density ($\Omega_B h^2$). The green regions show the measured limits on the relative abundances for each species. The current measured bound on the baryon density from WMAP is shown. [47].

thesis is one of the greatest successes of the Big Bang theory, and this places strong constraints on the total number of baryons in the Universe, providing further compelling evidence for a large non-baryonic dark matter component. Further details can be found in [48, 49]

2.4 Dark Matter Candidates

As a result of constraints placed on the baryonic content of the Universe, it is clear that a new particle is required. Such non-baryonic particles must satisfy several conditions to explain the observational evidence. They must be stable on cosmological time-scales, leaving a significant number in the present epoch, they must have the correct relic

density and they must interact rarely with baryonic matter or electromagnetic radiation in a weak or sub-weak interaction.

The menagerie of different theorised exotic dark matter candidates is extensive. As a result, we concentrate on the currently favoured candidate, Weakly Interacting Massive Particles (WIMPs), and overview a number of others only briefly.

2.4.1 Weakly Interacting Massive Particles (WIMPs)

WIMP is a generic term used to describe an electrically neutral, massive (GeV to TeV) weakly interacting particle, assumed to carry a conserved quantum number. In the early Universe these particles would have been created thermally in great numbers. As the Universe expanded and cooled, particle concentration would then have frozen out when the expansion rate overtook the annihilation rate, leaving a relic density of WIMPs. Simple calculations demonstrate that the strength of the weak interaction is of the correct order to produce a relic density of WIMPs contributing substantially to Ω_m .

The standard model of particle physics (SM) [50, 51] is a framework which has proved very successful in describing the fundamental structure of matter, describing the fundamental particles and forces with the exception of gravity. The standard model has successfully predicted the existence and properties of the W and Z bosons, the gluon, top and charm quarks, with particular precision. Despite its dominant position in particle physics over the past decades, the theory is incomplete. Besides not accounting for gravity or explaining the neutrino mass, it requires fine-tuning of the fermion and boson couplings to 1 part in 10^{14} to solve the Gauge Hierarchy Problem [52, 53, 54] and fails to unify the strong, weak and electromagnetic forces at high energy. The Gauge Hierarchy Problem is centred around why the weak force is 10^{32} times stronger than gravity and why the Higgs mass is so much smaller than the Planck mass. Quantum corrections should yield a Higgs mass at the Planck scale and an unnatural level of fine-tuning, with precise cancellation of the basic mass terms with quantum fluctuations, is required.

The theory of supersymmetry (SUSY) [56, 57, 58, 59] provides an elegant solution to some of these issues. The situation is analogous to the discovery of antiparticles,




Figure 2.12: The magnitude of the gauge coupling constants as a function of energy in the standard model (left) and in the minimal supersymmetric model (MSSM) (right) [55]. The gauge couplings $U(1)$, $SU(2)$ and $SU(3)$ represent the symmetries underlying the electromagnetic, weak and strong forces respectively.

with the introduction of an additional symmetry. In SUSY all standard model particles have a supersymmetric partner, or superpartner, with identical quantum numbers but $\Delta\text{spin} = 1/2$. This means that all fermions have bosonic superpartners, identified by the prefix *s*, e.g. the squark and the slepton [60]. Similarly, all bosons have a fermionic superpartners denoted by the suffix *ino*, e.g. the Wino (superpartner of the W boson) and the Higgsino.

SUSY models solve the Hierarchy problem by removing the sensitivity of the Higgs boson mass to the energy scale at which new physics is set through cancellations of the basic mass terms with those of superpartner particles, reducing the level of tuning required. The unification of forces at high energies is expected and vital for a grand unified theory (GUT), but the projected gauge couplings from the SM fail to converge (as shown in Figure 2.12). SUSY models reconcile this problem, with the couplings converging at $\sim 10^{16}$ GeV [61]. This brings theory into much better agreement with experimental limits placed by proton decay [50].

For SUSY to maintain the conservation of baryon and lepton number, an additional symmetry is required, known as R-parity, yielding a conserved quantum number [58]:

$$R = (-1)^{3(B-L)+2S}, \quad (2.18)$$

where B is the baryon number, L is the lepton number and S is the spin, with $R = +1$ for ordinary particles and $R = -1$ for their SUSY partners. Conservation of R-parity

has three consequences: production of SUSY particles can only occur in reactions of particle pairs (e.g. $e^+e^- \rightarrow \tilde{e}^+\tilde{e}^-$), heavy superpartners may decay into lighter SUSY particles (e.g. $\tilde{e} \rightarrow e\tilde{\gamma}$) and the lightest supersymmetric particle must be stable, with no kinematically-allowed negative R-parity state for it to decay into. In many flavours the lightest supersymmetric particle (LSP) is the neutralino, a combination of the four mass eigenstates, made from the zino, photino and 2 higgsinos, all of which have the same quantum numbers [62, 63].

The minimal supersymmetric standard model (MSSM) [64] is the SUSY extension to the standard model with the minimal particle content. Despite generally having over 120 free parameters, most MSSM models reduce this number by making well motivated assumptions, reducing the number of free parameters to just five. Whilst in some SUSY models the LSP could be the gravitino [65] or sneutrinos [2], in MSSM the lightest neutralino, $\tilde{\chi}_1^0$, constitutes the LSP.

As a consequence of R-parity conservation, after freeze out the heavier superpartners will decay, leaving the LSP as a relic. The LSP WIMPs will freeze out at $T \simeq m_\chi/20$ (largely independent of the particle properties), meaning that they are non-relativistic at the time of thermal decoupling, yielding a cold dark matter scenario (CDM). As mentioned previously, estimates of the relic density suggest that WIMPs, such as the LSP, should form a significant component of (if not all of) the dark matter.

The most promising method of exploring supersymmetry is with particle accelerators, where missing transverse momentum observed in collisions would provide evidence for SUSY particles. To date no SUSY particles have been found, although accelerator experiments have allowed constraints to be placed on their masses. Current accelerator searches have placed a lower limit on the mass of $\tilde{\chi}_1^0$ at 37 GeVc^{-2} [66]. Calculations using the latest WMAP data provide an upper limit on the mass of 500 GeVc^{-2} [65].

Due to the underlying motivations for supersymmetry and the excellent candidate particle resulting from it, the neutralino is the currently favoured candidate for dark matter, with a huge range of direct and indirect detection techniques being explored. Whilst accelerator searches should provide evidence for SUSY particles within the next few years, detection by non-accelerator experiments is required to show that the LSP is indeed the galactic dark matter. Searches for the rare, weak interactions between WIMPs and ordinary matter is the focus of this thesis.

2.4.2 Other candidates

Whilst WIMPs are the clearly favoured candidate for dark matter, they are not the only possibilities. Here we briefly discuss two of the other potential well-motivated candidates: neutrinos and axions.

Neutrinos

Neutrinos are part of the lepton family in the standard model of particle physics. The three neutrino flavours (electron, muon and tau), all interact through the weak force alone, and the absence of strong or electromagnetic couplings leave them weakly coupled with baryonic matter. In the standard model neutrinos are massless; however, measurements of neutrino oscillations [67, 68] prove them to possess a non-zero mass.

In standard Big Bang theory, a large number of neutrinos would have been produced in the early Universe which, combined with their weakly interacting nature and non-zero mass, meant that they were a promising dark matter candidate [69]. By applying an upper limit to the mass of the neutrino (from β -decay experiments) of $m_{\nu_e} < 2.05$ eV [70], an upper limit can be placed on the total neutrino relic density, $\Omega_\nu h^2 \lesssim 0.07$, which is clearly incompatible with neutrinos dominating the dark matter sector.

Additionally, due to their small mass, neutrinos must have remained relativistic at freeze out, retaining an energy ~ 1 MeV. Due to their weakly interacting nature they cannot be slowed significantly, yielding a Hot Dark Matter (HDM) scenario. Hot dark matter has been ruled out as the dominant form of dark matter from studies of the formation and evolution of large scale structure [71], as the fast moving particles cannot coalesce to form significant gravitationally bound structures. A more stringent constraint can be placed on their relic density by combining observations of large scale structure and CMB anisotropies, implying $\Omega_\nu h^2 < 0.0067$ (95% C.L.). As a result, neutrinos can no longer be considered as the dominant dark matter component, although they may play a role in a mixed dark matter model [72].

Axions

Axions were first proposed as a solution to the strong CP problem in theories of particle physics [73, 74, 75]. The strong interaction in the standard model, governed by QCD, should produce CP violation, although this has not been observed. In fact, constraints placed on the strong CP violation term, through high-precision measurements of the neutron electric dipole moment, require it to be extremely small. The question of why this parameter should be so close to zero, provides a naturalness issue, known as the strong CP problem.

One of the most elegant and compelling solutions to the strong CP problem, the Peccei-Quinn mechanism, was developed in the late 1970's. The axion particle is a result of this mechanism, additionally occurring naturally from many superstring theories. Despite having a very small mass, axions are an interesting candidate for CDM due to their non-thermal production method [76, 77].

Axions have not been detected to date, although experimental efforts are ongoing. One of the most promising approaches for axion searches utilises the axion-to-photon vertex, where a conversion between axions and photons can occur within a strong electric or magnetic field. The CERN Axion Solar Telescope (CAST) uses a strong magnet to search for the predicted flux of solar axions, but has so far only succeeded in placing constraints on the mass and photon coupling [78, 79]. Another approach is to look for changes in the polarisation of light as it passes through a vacuum in a strong magnetic field. Using this technique the PVLAS experiment originally claimed an effect which could have been interpreted as an axion mass of 1–1.5 meV [80], although this was at odds with other experimental and astrophysical constraints, and was later retracted after further studies showed instrumental effects were the cause [81]. Similarly, constraints can be placed on the mass and coupling of axions due to the predicted effects in astrophysical situations, such as energy loss from stars [82, 83, 84].

Despite their extremely small mass and the absence of an experimental discovery, axions remain a potential candidate contributing to the non-baryonic dark matter in the Universe. As with the neutralino, the motivation of the axion hypothesis is deep-rooted in particle physics theory, with their potential as a dark matter candidate being incidental.

2.5 Indirect WIMP Searches

Indirect searches for non-baryonic dark matter look for the products of dark matter annihilations. The amount of radiation produced will be proportional to the annihilation rate, which is in turn proportional to the square of the dark matter density. Therefore, searches for this radiation generally concentrate on regions (often termed amplifiers) such as the galactic centre, or even in compact objects such as stars and planets. This would occur as the dark matter interacts (however weakly) with the dense matter, loses energy and gets captured into a lower energy orbit until the central density increases. Annihilation photons cannot escape from such compact objects, however neutrinos would provide a potential signal.

The flux of γ -rays from neutralino annihilations at the galactic centre can be predicted, although a significant uncertainty arises from the unknown slope of the dark matter profile in the region. The annihilations mainly produce a continuum γ -ray spectrum, but may also emit mono-energetic γ -rays from:

$$\chi\chi \rightarrow \gamma\gamma \quad \chi\chi \rightarrow \gamma Z. \quad (2.19)$$

These mono-energetic lines offer a potential “smoking-gun” signature for the existence of dark matter, although the processes leading to lines imply small fluxes. Different experiments have searched or are currently searching for such γ -ray signatures include EGRET [85], FERMI (formerly GLAST) [86], INTEGRAL [87] and MAGIC [88]. Similarly, the ANTARES and ICECUBE experiments are searching for the neutrinos which may also be a remnant from such dark matter annihilations. Other sources of indirect signals are also under investigation, such as synchrotron radiation from secondary electrons (from neutralino annihilations) at the galactic centre. A review of potential signatures and the mechanisms behind them is given in [2] and references therein.

Anti-particles constitute a small fraction of the total flux of cosmic rays as they are primarily produced through secondary processes in interactions of cosmic-ray nuclei with the interstellar medium. Such secondary processes predict a cosmic-ray background spectrum with a certain energy-dependent positron fraction. However, such anti-particles may also be the products of dark matter annihilations or of astrophysical

objects such as pulsars, these are known as primary sources. Therefore the observation of an excess in the positron fraction could provide an indirect dark matter signature.




Figure 2.13: Positron fraction measured by the PAMELA experiment compared with other previous data, showing an excess above the predicted background spectra (continuous line) above 10 GeV and a deficit below 5 GeV due to variations in solar activity [89] (and references therein).

Recent reports of an excess in the positron fraction from cosmic rays have been interpreted by some as a dark matter signature. The space-based PAMELA experiment, measuring the positron flux in data recorded between 2006 and 2008, observed an excess in the positron/electron ratio above the expected model for the background [89] (shown in Figure 2.13). They concluded that a primary source was required, either astrophysical or from dark matter annihilation. These results follow the observation of another discrepancy by the ATIC balloon experiment, which measured the total flux of electrons plus positrons. Their data showed an excess in the total flux between 300 and 800 GeV, again implying the existence of a primary source [90]. The first results from the FERMI-LAT satellite also show an excess in the total flux in the 300 to 800 GeV region, although to a lesser extent than ATIC. A lower flux is also seen at low energies although this is attributed to variations in the solar activity.

These observations were seized upon by some theoreticians, devising a number of exotic dark matter models to account for the observed excesses. However, as discussed in [91], the excess fluxes observed are clearly consistent with well-known pulsars. Similarly, interpretation of the PAMELA and FERMI-LAT data in terms of dark matter annihilations requires additional boost and enhancement factors. Thus far there has been no convincing indirect evidence for a dark matter annihilation signature, although ongoing searches will continue to probe ever more interesting regions of the parameter space.

2.6 Direct Detection Experiments

WIMPs, in the form of neutralinos (the lightest SUSY particle), are the currently favoured candidate, with their motivation rooted in new particle physics theory beyond the standard model. As a result, there is a concerted effort by many groups across the world to develop potential techniques for detecting galactic WIMPs directly. Due to the weak or sub-weak nature of the interactions of such particles, they will interact with baryonic matter very rarely, producing low energy nuclear recoils when they do scatter.

Dark matter searches currently look primarily for the WIMP-nucleon interactions in the form of nuclear recoils, which are distinct from the majority of radioactive backgrounds (γ and β) which produce electron recoils in the target through interactions with the electronic subsystem. The remaining potentially problematic background for the current generation of instruments are neutrons which produce nuclear recoils, mimicking the signal expected from WIMPs.

WIMP-nucleus scattering is expected to be elastic, and as a result the energy transfer can be calculated through the kinematics of the interaction. This produces an exponentially falling differential energy deposition spectrum which is dependent upon the WIMP mass, the mass of the target species, and the velocity of the WIMPs with respect to the target. Figure 2.14 shows the recoil spectra for different WIMP masses scattering off a xenon target, with the spectra hardening with increasing WIMP mass. The importance of the mass of the target species is shown later in Figure 3.1, where it is clear that an increase in the mass of the target species enhances the expected rate

at low energies. A discussion of the calculation of recoil rates is provided in Chapter 6, following the methods described in [92].

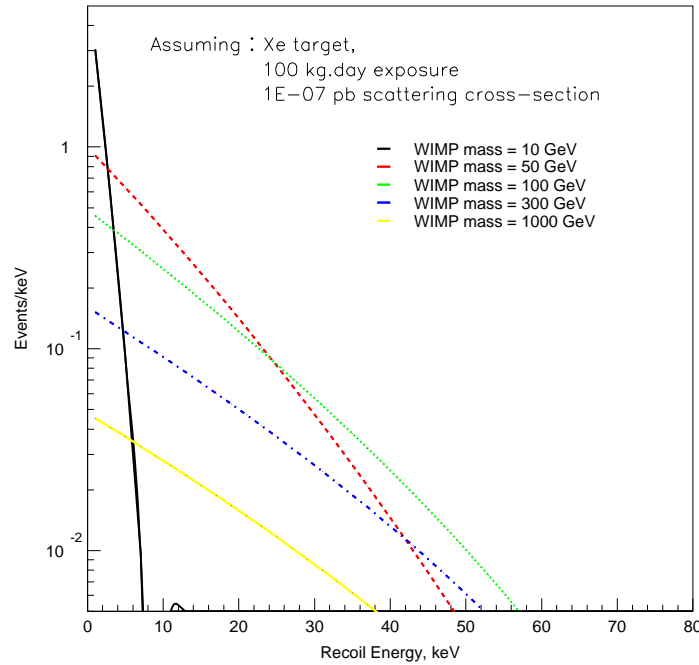


Figure 2.14: The differential energy spectra expected for different WIMP masses scattering on a xenon target. The rates assume a spin-independent scattering cross-section of 1×10^{-07} pb and a 100 kg·day exposure.

WIMP-nucleus scattering results from an expected WIMP-quark scattering channel. For neutralinos, in particular, this can involve squark or Z-boson exchange, for example. A scalar (spin-independent) WIMP-quark interaction would result in a WIMP-nucleus cross-section which is determined by:

$$\sigma = \frac{4m_r^2}{\pi} (Zf_p + (A - Z)f_n)^2 \quad (2.20)$$

where m_r is the reduced mass of the WIMP-nucleon system, $f_{p,n}$ is the coupling term for WIMP-proton or WIMP-neutron. As there is no difference between the coupling terms for protons and neutrons, the cross-section has a A^2 dependence. This clearly favours the use of heavier target species, where the interaction rates will be increased by this enhancement factor. This means that searches are currently setting the most stringent limits on spin-independent cross-section. However, for an axial-vector (spin-dependent) interaction between WIMPs and nucleons, the cross-section is no longer dependent on

the A of the target species, but becomes proportional to $J(J+1)$, where J is the nuclear spin. This removes the advantage of using heavier elements and promotes target species with a favourable spin factor (such as ^{19}F or ^{127}I).

Due to the extremely rare and low-energy nature of interactions between WIMPs and target species, it is crucial that instruments searching for them are operated in extremely low-background configurations. Consequently, careful selection of detector components and external shielding is required. However, on the Earth's surface, any detector will be swamped by cosmic rays and associated backgrounds, rendering the instruments useless for rare-event studies. Consequently, the search for WIMP dark matter has been forced underground into laboratories where the cosmic rays are shielded by hundreds (or thousands) of metres of rock. These underground laboratories are often located in mines or road tunnels where basic infrastructure is already in place. As the different rock compositions at the different sites will vary, giving different effective shielding, it is common to denote the water-equivalent (w.e.) shielding for each site. Figure 2.15 shows the relationship between the muon flux and the water-equivalent depth with different underground laboratories across the world noted. Some of the different underground laboratories have been reviewed in [93] and [94].

A multitude of techniques have been applied to the problem of detecting WIMP-nucleon scattering. To achieve the excellent sensitivity required, high atomic mass A , large target mass and low energy threshold help to boost event rates, with a low background and good discrimination aiding the signal-to-background ratio. Many different target species have been applied to the challenge of detecting WIMPs, each having advantages and disadvantages. A variety of different target species will be important when it comes to characterising the dark matter particles after a detection has been made, as the different species will provide different scattering properties and reduce any associated systematics.

Three different energy deposition channels exist: scintillation (as a result of excitation), ionisation and heat. These signals can be measured either individually or in combination, for single events. Alternatively, the expected directional dependence of a signal, as the Earth (and hence the target), moves through the dark matter halo, can help infer the presence of a signal. In its simplest form this would be observed as an annual modulation in the event rates over and above a constant background, with




Figure 2.15: Relationship between depth and muon flux for some of the major underground laboratories [95].

the eventual aim being to measure the directional information for individual events, facilitating the long-term goal of “WIMP astronomy”, measuring the local halo properties. Figure 2.16 details the different channels for detection and experiments which utilise them. In this section we will briefly review the two current leading WIMP search techniques, liquid noble gases and cryogenic detectors, outlining the detection principles and resulting limits on WIMP-nucleon scattering. We will also discuss other potential techniques being developed.

2.6.1 Noble gas detectors

One of the most promising techniques being applied currently to search for WIMPs is the use of liquid noble gases as a target medium. Particle interactions in the condensed phase produce scintillation and ionisation. Elastic interactions of neutrons or WIMPs cause a xenon nucleus to recoil, whereas γ -rays and β particles interact through electron recoils. The ratio of ionisation/scintillation outputs depends on the type of recoil through differences in track topology caused, ultimately, by the different rates of elec-

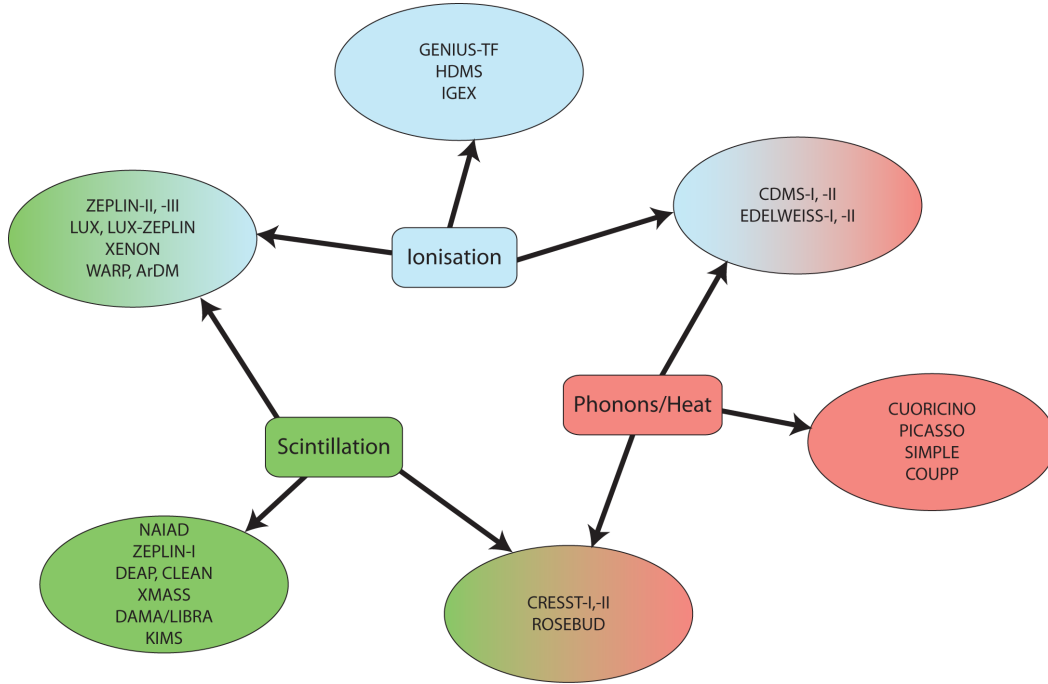


Figure 2.16: Diagram of energy deposition channels (and combinations of channels) utilised in dark matter searches.

tronic and nuclear energy loss and the varying degree of charge recombination in each case. A more detailed description of the excitation processes is given in Chapter 3.

In the case where scintillation alone is measured (single-phase detectors using liquid alone), some discrimination power is available from the scintillation pulse shape, characterised by variations in the apparent scintillation time constant [96]. However, two-phase detectors (with both liquid and gas layers) can measure both the scintillation and ionisation produced. Application of an electric field across the liquid causes the free electrons liberated by a particle interaction to be drifted to a region or surface where they can be detected. The currently-favoured method of measuring the ionisation is to emit the electrons into a thin gas region, where a higher field produces proportional electroluminescence as the electrons drift through the gas. This is a large gain process, with a large number of photons produced per ionisation electron. The two signals can then be recorded by the same array/s of photomultiplier tubes, with no need for different readout technologies for each response channel.

Several noble gases (with properties summarised in Table 2.1) have been considered for WIMP search experiments: xenon, argon, krypton and neon. They each have their

own advantages and are an attractive prospect for WIMP searches for a number of reasons, which are described in Section 3.1 for the case of xenon. For this noble element, both detection methods, single-phase and two-phase, have been demonstrated and used to place constraints on the WIMP-nucleon scattering cross-section. Single-phase xenon was demonstrated by the UKDMC with the ZEPLIN-I detector operating at Boulby [97]. ZEPLIN-I is discussed in more detail in Chapter 3, producing a limit with a minimum of $\sigma_{w-n} = 1.1 \times 10^{-6}$ pb. The Japanese XMASS experiment [98] will consist of an 800 kg liquid xenon target, entirely surrounded by photomultipliers. It will utilise the excellent self-shielding afforded by xenon, creating an extremely low-background central target region. It aims to begin underground operation soon, with a projected sensitivity of about $\sigma_{w-n} = 1 \times 10^{-9}$ pb, assuming the detector operates as designed and remains background free.

Table 2.1: Properties of noble gases of interest for dark matter targets.

	Z (A)	Boiling Point at 1 atm [K]	ρ at BP [g/cm ⁻³]	Scintillation λ [nm]
Ne	10 (20)	27.1	1.21	85
Ar	18 (40)	87.3	1.40	125
Kr	36 (84)	119.8	2.41	150
Xe	54 (131)	165.0	3.06	175

The DEAP/CLEAN collaboration are planning to make use of the same technique, but instead using argon [99, 100]. The previous DEAP and CLEAN collaborations have demonstrated the technologies applied to smaller target volumes. The 7 kg liquid argon DEAP-1 target demonstrated the excellent pulse-shape discrimination between electron and nuclear recoils offered by argon, with γ -ray rejection at the level of 6×10^{-8} . They intend to construct a 3600 kg argon target at SNOLab (Canada) using a similar detector geometry to XMASS, aiming to begin data-taking in 2010.

The use of two-phase noble gas systems has been extensively demonstrated with four WIMP-exclusion limits published. Two-phase xenon has been the leading technology, with the ZEPLIN programme producing limits from two different two-phase detectors. The ZEPLIN-II [101] and ZEPLIN-III [102] detectors are the subject of this thesis, with description of the experimental setups and data analysis chains. ZEPLIN-II produced

the first WIMP search result from a two-phase noble gas instrument, setting a limit of $\sigma_{w-n} = 6.6 \times 10^{-7}$ pb [101]. The first science run of ZEPLIN-III yielded a limit on the cross-section of $\sigma_{w-n} = 8.1 \times 10^{-8}$ pb, placing it as one of the world's leading WIMP search experiments [103].

XENON10 [104], a detector operated at Gran Sasso, used a similar detector geometry to ZEPLIN-II, but with an additional array of PMTs located in the liquid, greatly increasing the light yield and thus enabling a lower energy threshold. They used a slightly lower drift field to that in ZEPLIN-II, demonstrating a similar separation of electron and nuclear recoils. From their 58 day background data set, they were able to place a 90% upper confidence limit, using the Yellin optimum interval method, of $\sigma_{w-n} = 4.5 \times 10^{-8}$ pb for a 30 GeV/c² WIMP [104]. The two-phase xenon experiments operated thus far have been successful in placing meaningful constraints on WIMP-nucleon scattering, as well as demonstrating the power of the two-phase technique for WIMP searches.

The WARP collaboration have operated a detector with a similar concept to that of the two-phase xenon detectors, but utilising argon as the target medium [105]. This has the advantage of supplying an additional discriminant parameter, from pulse-shape analysis given that the difference in scintillation time constants for nuclear and electron recoils is very large in liquid argon. However, due to the lower A and larger threshold compared with xenon, the expected scattering rate will be much lower.

The demonstration of the technology and initial results have been carried out with detectors with an active volume up to ~ 35 kg. The next step in the search for WIMPs is to scale these detectors up using essentially the same technology and operating principles. The development of even lower background photomultiplier tubes was vital for large instruments with copious numbers of PMTs, a task which has been addressed by several groups. The next phase will employ target masses of order 100–300 kg, which should push the constraints on the WIMP properties into the currently favoured areas of the parameter space, namely through the US-led XENON100 and LUX300 instruments. Tonne-scale systems are already being designed, with the UK expected to play a major role with the newly formed LUX-ZEPLIN collaboration.

2.6.2 Cryogenic detectors

The other leading technique is that of cryogenic detectors, utilising crystals cooled to mK temperatures. When a particle interacts in a crystal, the majority of the energy is transferred to lattice vibrations (phonons). This energy is unquenched for nuclear recoils (unlike ionisation and scintillation), allowing (potentially) for very low energy thresholds.

Due to the large unquenched fraction of deposited energy, phonons are an excellent measure of the energy deposited, providing excellent energy resolution. When measurement of the heat is combined with that of another, quenched, excitation channel (scintillation or ionisation), very good discrimination can be achieved. Here the phonon channel provides the energy scale and the ratio of the second channel to the phonon channel provides the discriminant.

The Cryogenic Dark Matter Search (CDMS) is a US-based collaboration currently operating at the Soudan underground laboratory [106, 107]. CDMS have been developing the technique of combined phonon and ionisation readout from Ge and Si crystals and their current experimental setup, CDMS-II, employs 19 Ge (250 g each) and 11 Si (100 g each) Z-dependent Ionisation and Phonon (ZIP) detectors, cooled to ~ 40 mK. Each detector forms a 1 cm thick disk, 7.6 cm in diameter, with ionisation readout from one flat face and phonon readout from the other. Recent improvements in readout techniques, measuring the timing of the signal, allows rejection of surface events where suppression of the ionisation signal from electron recoils mimic nuclear recoils.

The most recent results published by CDMS-II, combined a re-analysis of their 2002 and 2003 data with newer data from 2008. The 2008 data alone showed no events in their WIMP search region (as shown in Figure 2.17), yielding a Poisson 90% confidence upper limit with a minimum at 6.6×10^{-8} pb for a WIMP mass of 60 GeV/c². The analysis of the combined dataset produced a world leading limit with a minimum at 4.6×10^{-8} pb for a 60 GeV/c² WIMP. The experiment has continued acquiring data, but has yet to announce further results.

The EDELWEISS experiment operates at the Laboratoire Souterrain de Modane in the Frejus Tunnel, France. Working on the same detection principle, measuring both phonons and ionisation, they published results from a 320 g Ge detector. As with

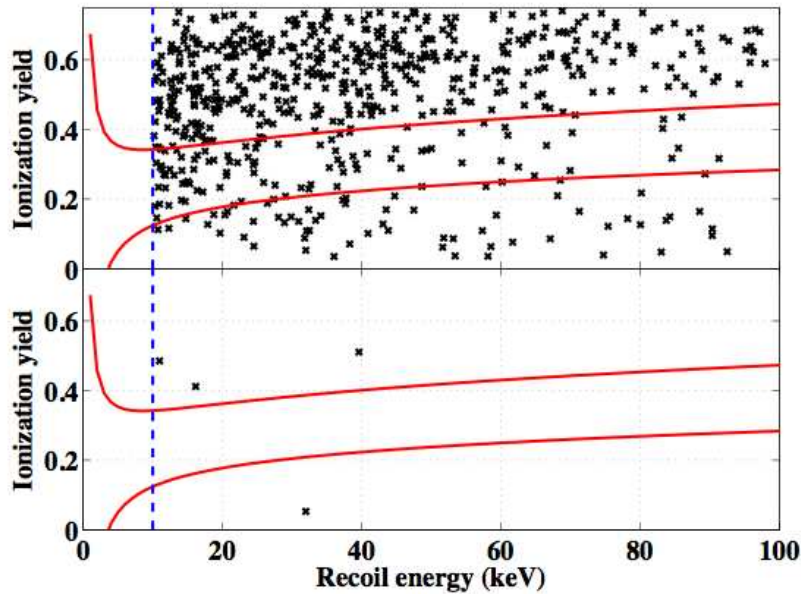


Figure 2.17: Discrimination scatter plot from the CDMS-II 2008 data [107], before (top) and after (bottom) applying an event timing cut. The blue dotted line defines their experimental threshold and the solid red lines define the signal region.

CDMS, they proved the technique to have excellent discrimination between electron and nuclear recoils, shown by the large separation of the populations in Figure 2.18. The 2001 data from EDELWEISS-I [108] produced an upper limit on spin-independent WIMP-nucleon scattering cross-section of 7.2×10^{-6} pb for a $52 \text{ GeV}/c^2$ WIMP. The CRESST experiment [109] has developed detectors measuring phonons in combination with scintillation. They have so far published results from a small initial setup with their limit having a minimum of 6×10^{-7} pb from an exposure of 67 kg-days [109].

The cryogenic experiments have demonstrated excellent discrimination, energy resolution and low energy thresholds. These all suggest this to be a leading technology for WIMP searches, but they are hampered by some serious drawbacks which will hinder their progress in sensitivity. The issue of surface contamination producing suppressed signals in their search regions has been a major limiting factor, although new detector techniques currently being demonstrated show this problem to be greatly mitigated. Increasing the event rate in WIMP search experiments depends of two factors, the A of the target material (a weakness when compared with other target species such as Xe) and the mass of the detector. High-purity crystals are not easy or cheap to produce, with each detector limited to a few hundred grams. Hence, to achieve a large mass

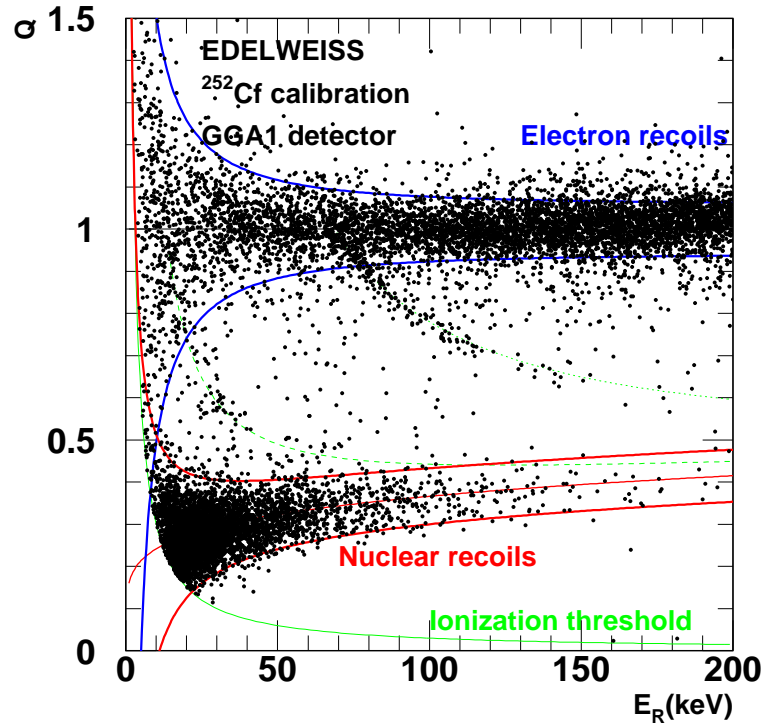


Figure 2.18: Calibration data from the EDELWEISS-I detector showing the electron and nuclear recoils populations, defining the signal (red lines) and background (blue lines) regions and the experimental threshold (green line) [110].

target many detector modules are required and this increases not only the cost, but also the complexity and the background within the cryostat. It is also clear that cryostats themselves are a limiting factor, with long cooling/warming times adding to the length of commissioning periods. However, cryogenic detectors are a key technology which may play an important role in detecting and characterising WIMP dark matter.

2.6.3 Other techniques

Alongside the impressive progress made using liquid noble gas and cryogenic detectors, other techniques are being developed. We summarise briefly two of these: bubble chambers/superheated liquids and directional detection, through both annual modulation searches and event-by-event directionality.

Bubble chambers and superheated liquids

A relatively recent technique in the field of WIMP searches is the use of bubble chambers and superheated liquids. The WIMP target for the PICASSO experiment [111] is superheated C_4F_{10} droplets dispersed in a viscous medium. The superheated (i.e. heated well above boiling point) material is extremely unstable, with any slight perturbation triggering a sudden phase transition from liquid to gas. A WIMP scattering from one of the atoms in a target droplet will result in such a transformation, creating an acoustic pulse through the medium, which can be detected by piezoelectric sensors. Such detectors can be made virtually blind to the electron recoil background, as the much greater rate of energy loss from nuclear recoils triggers bubble formation at much lower energies. PICASSO quote an electron recoil background rejection of 10^{-7} for a nuclear recoil threshold of 5 keV.

The SIMPLE collaboration [112] utilises a similar technique, with C_2ClF_5 as the working medium. Such species provide good targets for spin-dependent interactions (due to the favourable nuclear spin of fluorine), but yield poor sensitivity for spin-independent interactions. In general, superheated droplet detectors offer an excellent prospect of having a target virtually free from electromagnetic background. Despite being some way off producing competitive constraints on spin-independent scattering, they are currently reporting significant spin-dependent results [113], helping to constrain primarily σ_{w-p} .

Directional detection

The most basic form of directional detection is the search for an annual modulation in signal rates. The DAMA experiment uses this approach, forgoing pulse-shape discrimination, and instead searching for an annual modulation in observed event rates. As the Earth moves around the Sun, its velocity with respect to the halo dark matter varies, yielding a modulation in the scattering rate. DAMA initially operated a 100 kg high-purity NaI target for 7 years (DAMA/NaI), before upgrading to 250 kg target (DAMA/LIBRA), for a further 4 year run [114]. Both experiments observed an annual modulation in the signal with 8σ confidence [115], as shown in Figure 2.19.

The effect is clear, although the explanation is not. In the most recent analysis of




Figure 2.19: Residual event rate of single-hit scintillation events in the DAMA/NaI and DAMA/LIBRA experiments in three different energy ranges. The curve shows the best-fit modulation signal to the data [115].

the complete data-set, potential environmental and experimental sources of modulation have been considered and determined to be insignificant. However, the WIMP properties implied by the modulation results are almost totally excluded by other direct detection experiments. For the DAMA results to be compatible with other experimental data, a different interaction mechanism must dominate. This could be manifest through interactions with electrons rather than nuclei, which are not well studied with the other detection techniques. Consequently, theoreticians have been hard at work developing models which can reconcile all the observations, such as inelastic dark matter [116].

The modulation observed by the DAMA experiment is no longer in doubt, and it is indeed the case that many of its characteristics agree with those required by a dark matter signature (correct phase, amplitude, single scatters only). Although, an unknown background or experimental effects could still provide a potential explanation, as could

variability in the potassium contamination which provides a low energy calibration feature at 3.2 keV [114], in the low-energy range of interest. If, however, the modulation is a result of dark matter interactions, new interaction models will probably be required.

The eventual goal of direct WIMP searches would be WIMP astronomy, providing definition of the local dark matter halo. The directional signature could also provide strong support for the detection of dark matter. Annual and diurnal variations will help indicate a dark matter signature, measuring directionality by population. However, the ultimate goal would be directional information on an event-by-event basis. To this end, there are currently several experiments in their early stages, working towards the goal of directional measurements of nuclear recoils. These include the DRIFT programme [117, 118] operating at Boulby mine and R&D with the DM-TPC [119], NEWAGE [120] and MIMAC [121] programmes.

2.7 Summary

The current period in the development of instruments searching for WIMP dark matter is an exciting one, with sections of parameter space favoured by MSSM being probed for the first time, with different and equally promising techniques being pursued. Figure 2.20 shows the current world status (mid-2009) of constraints which have been placed on spin-independent WIMP-nucleon scattering. The currently favoured regions of the parameter space (shaded in green in Figure 2.20) are determined through Monte-Carlo simulations of different MSSM models incorporating current observations from a wide range of experiments [122, 123]. Over time these regions change as they further constraints help inform the models. Three experiments: ZEPLIN-III, XENON10 and CDMS-II, currently provide the most stringent limits at the $\sim 5 \times 10^{-8}$ pb level, beginning to test the currently favoured parameter space [122].

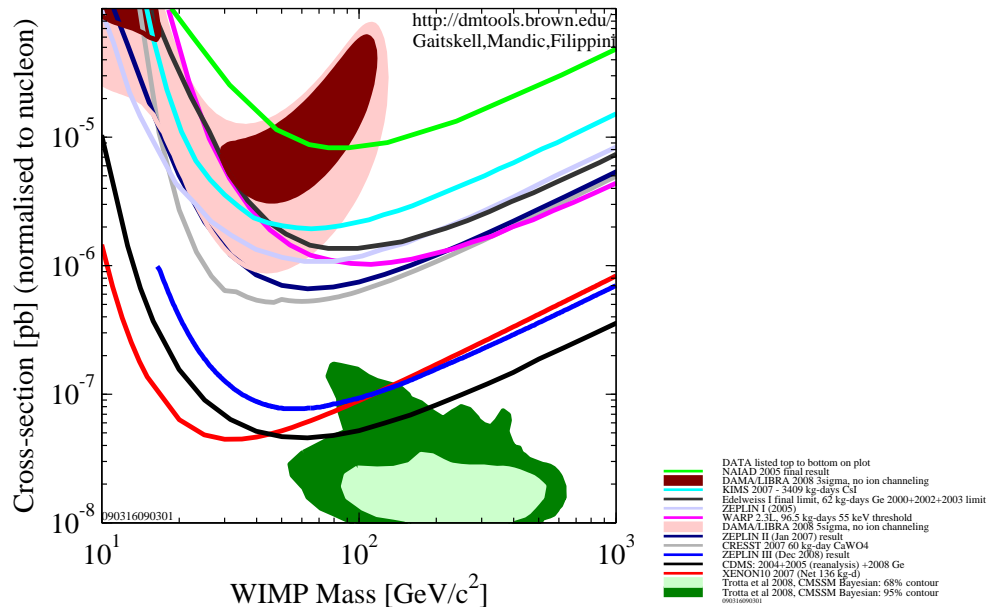


Figure 2.20: Current world status of experimental constraints placed on the spin-independent WIMP-nucleon scattering cross-section, as a function of WIMP mass. The red shaded region shows the claimed detection by the DAMA experiment, with the green regions showing the currently favoured MSSM parameter space [124].

Chapter 3

The ZEPLIN programme

Searches for dark matter in the form of WIMPs began in the UK in the late 1980s with investigation of different potential WIMP search techniques, within the framework of the UK Dark Matter Collaboration (UKDMC). The need for a low-background environment for such studies led to the establishment of the Boulby Underground Laboratory. Initially, the UKDMC, consisting of Imperial College London, the Rutherford Appleton Laboratory, University of Sheffield, Royal Holloway, Birkbeck College, Queen Mary London, Nottingham University and the University of Edinburgh, focused its efforts on utilising NaI scintillators for WIMP searches [125]. The culmination of this work was NAIAD, an array of encapsulated NaI detectors, which acquired data used to set upper limits on the spin-independent WIMP-nucleon scattering cross-section in 2003 [126] and 2005 [127].

In partnership with other international collaborating groups (UCLA, Texas A & M, University of Rochester, LIP-Coimbra, ITEP Moscow), the UKDMC began work on more complex detectors with enhanced sensitivity to WIMP interactions. These included the use of liquid xenon as a target medium [96], with the ZEPLIN detectors, and low-pressure time projection chambers aiming to probe the directionality of WIMP scattering, with the DRIFT detectors [117, 118].

The ZEPLIN programme (ZonEd Proportional scintillation in LIquid Noble gases) utilises liquid xenon as a target for WIMP scattering, with three detectors (ZEPLINs I, II and III) each having their own specific designs and characteristics. ZEPLIN-I used a single-phase (liquid) target measuring scintillation alone and using pulse-shape

discrimination to distinguish between nuclear and electron recoils. ZEPLIN-II and III utilise a two-phase setup measuring both scintillation and ionisation, with the ratio of the two providing enhanced discrimination power. The two-phase systems exploited very different designs and engineering solutions which would enable the best features of each instrument to be selected, effectively prototyping technologies for a future tonne-scale xenon experiment, ZEPLIN-MAX.

In 2007 the UKDMC ceased to exist following the completion of the ZEPLIN-II experiment, with separate collaborations based around the ZEPLIN-III and DRIFT-II experiments resulting. The ZEPLIN-III collaboration consists of Imperial College London, STFC Rutherford Appleton Laboratory, University of Edinburgh, LIP-Coimbra and ITEP Moscow, totalling some 30 active members.

In this chapter I will review the physics processes at work in liquid xenon as a radiation detection medium, provide a brief description of the ZEPLIN-I experiment and describe the ZEPLIN-II and III instruments with some discussion of operational considerations, which I personally played a part in investigating.

3.1 Principles of WIMP detection using liquid xenon

In the previous chapter I described a plethora of different techniques applied to the search for WIMP dark matter. However, the reasons for using liquid xenon as a target species are compelling, some relating to the nature of WIMP-xenon interactions and others relating to its macroscopic properties. Listed below are some of the major advantages of xenon as a WIMP target.

1. Liquid xenon has a high scintillation yield (comparable with NaI), and is transparent to its own scintillation light, meaning that large volumes can be utilised with little loss of sensitivity. Xenon scintillates in the vacuum ultraviolet (VUV), centred at $\lambda = 175 \text{ nm}$ ($\sim 7.1 \text{ eV}$) with a spectral width of about 10 nm [128].
2. Very high sensitivity is similarly achieved in the ionisation channel, where it is possible to drift free ionisation electrons through a large depth of liquid, with the limiting factor being electronegative impurities (which can be removed through purification techniques). The sensitivity of the two-phase technique to single elec-

trons released in a large volume of xenon was demonstrated in ZEPLIN-II (as described in detail in Chapter 7 and [129]).

3. The size of the xenon nucleus ($A \sim 131$) provides a good kinematic match to the favoured WIMP masses, enhancing interaction rates at low energies. Figure 3.1 shows the event rates expected from a perfect detector (i.e. one with 100% detection efficiency and perfect energy resolution) made from some of the different target species used in WIMP searches. Xenon clearly has a favourable event rate at lower energies, for the same target mass.

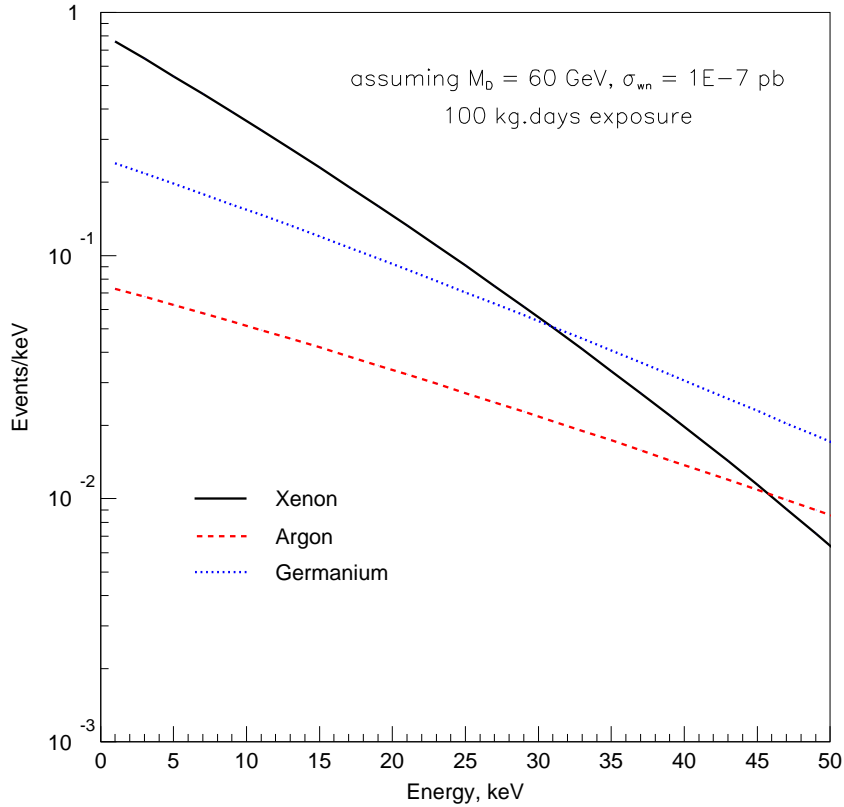


Figure 3.1: Expected event rate from a perfect detector for different target species (for a WIMP mass of 60 GeV).

4. Xenon itself is intrinsically very radio-pure with no long-lived radioactive isotopes. Commercial Xe usually has small amounts of Kr, and a small fraction of this is ^{85}Kr , which is a β -emitter ($T_{1/2} = 10.8$ yr) of anthropogenic origin. Kr can be removed by novel purification techniques; alternatively, pre-nuclear testing sources

of xenon can be sought with very low levels of ^{85}Kr , as is the case of the ZEPLIN - III xenon from ITEP.

5. Liquid xenon is dense, $\sim 3 \text{ g/cm}^3$ at 195 K, which means a large mass detector occupies a relatively small volume. This is extremely advantageous for the construction of large scale detectors.
6. Liquid noble gas targets have the advantage of scalability to larger volumes using the same technology, with a larger vessel required. This scaling also provides additional active self-shielding, where a region of xenon around the edge of the target absorbs the majority of low-energy and surface backgrounds and can be sacrificed to remove them, leaving a central volume with extremely low-background.

The scintillation produced by the interactions of ionising particles in noble gases has been utilised as a mechanism for radiation detection since the 1950's. An ionising particle interacting with an atom of the target medium induces the emission of VUV scintillation photons, with the number proportional to the energy deposited. Although the exact photon yield remains somewhat uncertain, measurements have demonstrated that, on average, between 12 and 40 eV are required to produce a scintillation photon [130, 131]. These photons can be observed by a photomultiplier tube (or an array of tubes) with quartz windows. The digitised pulses can then be characterised by integrated pulse area, representing the energy deposited during the interaction, and pulse-shape. Such characteristics can then be utilised to identify the type of interacting particle.

Early work investigating two-phase emission of ionisation electrons was carried out in the 1940's [132], but the mechanism was not fully exploited as a method for radiation detection until the 1970's with the development of detectors using solid argon [133]. In two-phase xenon, a particle interacting in the liquid xenon (LXe) target produces both prompt scintillation photons and releases ionisation electrons. By applying an external electric field to the liquid, some electrons can be extracted from the interaction site to be detected independently. The currently favoured method of charge detection from a liquid target relies on using electroluminescence to convert the ionisation signal into a proportional photon signal in the gas phase [134]. Upon reaching the liquid surface, electrons can be emitted into the gas phase with near unity efficiency at 5 kV/cm [135].

Once in the gas, they are accelerated by the stronger field there, collisionally exciting atoms to produce many VUV photons through secondary scintillation. This allows both primary and secondary signals to be detected with the same array of photon detectors.





Figure 3.2: Schematic of the processes occurring during an interaction in liquid xenon.

3.1.1 Liquid xenon scintillation

The interactions in the xenon proceed by two separate processes (described in Figure 3.2), both resulting in the production of VUV scintillation photons [136, 137]. The

processes involve either the direct or indirect creation of the excited dimer, Xe_2^* :



With zero electric field - and, to a lesser degree, at any finite field - the Xe^+ and e^- can recombine, as in Eq. (3.2), yielding the same Xe_2^* state as in Eq. (3.1). The radiative decay of this excited dimer produces a VUV photon. Two different time constants, 3 ns and 27 ns [96], are observed for the production of scintillation from this mechanism ('exciton' luminescence), resulting from the lifetimes of the singlet and triplet states of the excited dimer. It has been observed that the ratio between the two states is dependent upon the type of particle interacting. In the case of nuclear recoils, recombination is fast and, as a result, scintillation emission is dominated by exciton luminescence alone. However, in the case of electron recoils, with a longer, sparser ionisation track topology, recombination luminescence becomes important, yielding an apparent time constant in the region of 34-45 ns [138, 137], which is now determined by the mean time associated with recombination in Eq. (3.2). This difference in the observed time constants for nuclear and electron recoils allows some moderate discrimination to be achieved with pulse-shape analysis in xenon, especially at zero field [96].

In addition, the different track topologies result in a variation of the scintillation yield between electron and nuclear recoils. In the case of nuclear recoils from elastic scattering, a portion of the energy of the incident particle (neutron or WIMP) will be transferred to atomic motion (atom cascades), with only a fraction transferring to the electronic subsystem. Consequently, there is a quenching of the observed energy deposit when compared with deposition through electron recoils from γ -rays. This is extremely important for WIMP search experiments as the aim is to detect and measure nuclear recoils (from WIMP interactions). However, energy calibrations of WIMP search

detectors usually rely on γ -rays to provide a mono-energetic beam into the detector. As a result, knowledge of the scintillation yields from the different particle interaction types is required to convert this electron recoil energy (electron-equivalent), E_{ee} , into nuclear recoil energy, E_{nr} , accounting for the lower scintillation yield for the same energy deposition in nuclear recoil.

Measurement of the absolute scintillation yields from different particles is difficult and, as a result, the relative scintillation efficiency (or L_{eff}) is measured as the ratio of the two yields (nuclear recoil and γ -ray). Measurements of L_{eff} are usually carried out using a mono-energetic neutron beam, calculating the recoil energy kinematically from the scattering angle, and comparing the response to that of a particular γ -ray energy (typically 122 keV from ^{57}Co). Due to the precision required at very low scattering angles, low-energy measurements (below $\sim 40 \text{ keV}_{nr}$) of L_{eff} are difficult. Until recently, measurements of L_{eff} had only been made down to $\sim 10 \text{ keV}_{nr}$ with good precision, showing a L_{eff} that was broadly consistent with a constant value of 0.2 [139]. However, recent measurements [140] suggest a falling L_{eff} at low energies. All these measurements and energy calibrations rely on linearity in electron recoil response below 122 keV, which has not been decisively proven. Linearity is observed down to about 30-40 keV_{ee} ; below this, confirmation becomes more challenging with low energy γ -rays failing to penetrate detector walls.

3.1.2 Two-phase detectors

The application of a significant electric field progressively inhibits the recombination process, with a fraction of the ionisation electrons extracted from the interaction site. This reduces the discrimination power from scintillation pulse-shape, through progressive removal of the slower recombination component for electron recoils. It does, however, open a new possibility: discrimination through simultaneous measurement of both scintillation and ionisation.

The fraction of ionisation electrons drifted away from the interaction site is dependent upon the electric field strength. This fraction will vary for different interacting species (nuclear recoils, γ -rays and α -particles) as a result of their differing recoil track configurations. This precipitates a suppression of the scintillation yield, with different field dependencies measured for different particle types. Figure 3.4 shows these




Figure 3.3: Plot showing measurements of the relative scintillation efficiency, L_{eff} , as a function of energy [140]. Symbols correspond to (\circ) - Aprile *et al*, 2009 [140]; (\square) - Chepel *et al*, 2006 [139]; (\triangle) - Aprile *et al*, 2005 [141]; (\diamond) - Akimov *et al*, 2002 [142]; (\times) - Bernabei *et al*, 2001 [143]; (∇) - Arneodo *et al*, 2000 [144]. The solid grey curve is the result from a best-fit analysis of the XENON10 AmBe source between data and MC [145]. Also shown is the theoretical prediction of Hitachi (dashed line) [146].

measurements along with the complementary variation in the ionisation channel. The scintillation yield for electron recoils decreases significantly as the field increases, even beyond the 5 kV/cm strength covered in the plot. Conversely, nuclear recoils demonstrate weaker field suppression due to their differing track topology. At 1 kV/cm, suppression of scintillation from nuclear recoils is only about 10%, compared to about 50% for electron recoils at the same field.

Until relatively recently, nuclear recoils were not known to produce any ionisation and as a result their behaviour is not yet precisely understood. A gap remains in the understanding of their scintillation suppression as only a single energy has so far been measured, 56 keV_{nr}. However, these suppression factors are very important as they determine the nuclear recoil energy calibration of WIMP search data. Further studies of field dependence at different energies are clearly required, perhaps in combination with a study of L_{eff} .

The energy scale in two-phase xenon instruments is determined by the scintillation channel. External γ -ray sources are used to calibrate the response of the scintillation




Figure 3.4: Measurements of the electric-field induced variations of the scintillation and ionisation yields from electron recoils (ER), nuclear recoils (NR) and α -particles [147]. $S(E)/S_0$ and $Q(E)/Q_0$ are the scintillation and ionisation yields compared to their zero field values.

channel to electron recoils, E_{ee} . However, for WIMP search experiments the nuclear recoil energy, E_{nr} , is required. For single-phase experiments this conversion is determined simply by the L_{eff} . However, the same conversion in two-phase instruments (discussed further in Sections 4.6 and 5.10) additionally requires knowledge of the electron and nuclear recoil scintillation suppression factors (s_e and s_n):

$$E_{nr} = \frac{E_{ee}}{L_{eff}} \times \frac{s_e}{s_n}. \quad (3.3)$$

Once electrons are liberated from the interaction site, they drift through the liquid xenon following the electric field lines, with a drift velocity dependent upon the magnitude of the applied field. Figure 3.5 shows measurements of the electron drift velocity as a function of field, showing v_e increasing, eventually reaching a saturation value of $\sim 3 \text{ mm}/\mu\text{s}$ at a few kV/cm [148].

During their drift through the liquid they may be trapped by electronegative contaminants in the xenon. These contaminants can include O_2 , N_2O and SF_6 , each with different attachment cross-sections. The cross-sections for some of these species have




Figure 3.5: Measurements of the drift velocity of electrons in liquid and solid xenon [148].

been measured [149], and found to exhibit characteristic variations with the applied electric field. For O_2 and SF_6 the attachment probability decreases with increasing field, but N_2O shows the opposite behaviour, with attachment more probable at higher fields. The presence of such species in the liquid xenon results in a finite electron lifetime, τ_{e-} , for electrons drifting through the liquid, with $\tau_{e-} \propto [X]^{-1}$, where $[X]$ is the concentration of the contaminant species.

Upon reaching the surface of the liquid, the electric field can lead to cross-phase emission, with electrons passing from the liquid to the gas. The phase boundary between liquid and gas acts as a potential barrier for electrons to cross, with the probability of electrons crossing the boundary dependent upon the energy afforded them by the electric field in the liquid. There is a minimum threshold energy which translates directly to a field threshold, $E_{thres} \simeq 1.5$ kV/cm (i.e. for emission to occur $E_l > E_{thres}$). Above this threshold the fraction of electrons being emitted increases until a field is reached, ~ 5 kV/cm, which affords enough energy to the electrons for $\sim 100\%$ emission.

Electrons emitted into the gas phase are accelerated by the strong electric field there, collisionally exciting xenon atoms to form Xe_2^* , which in turn radiatively decays producing secondary VUV scintillation photons. The production of photons by electroluminescence increases linearly with electric field. At lower fields, electrons do not gain




Figure 3.6: Measurement of the cross-phase emission probability for electrons in liquid xenon [135]. The line shows the parameterisation utilised in ZEPLIN-III detector simulations [150].

enough energy between collisions to excite the xenon atoms, resulting in a threshold field required to produce electroluminescence, ~ 2 kV/cm. Measurements of the secondary scintillation yield have been made in room temperature gas and, more recently, in cold saturated vapour (as found in two-phase systems). In the cold saturated vapour the yield is found to be higher than at room temperature, corrected for the same atom number density [151]. A summary of measurements is given in [152] (and references therein) with Figure 3.7 showing the values both for the cold saturated vapour and room temperature measurements. The electroluminescence yield per unit track length, Y , is found to be a function of the electric field in the gas, E_g , and the gas density, expressed through equivalent pressure for the same gas density at 0°C , P_{eq} :

$$Y = A \times E_g - B \times P_{eq}, \quad (3.4)$$

where A and B are experimentally determined constants, summarised later in Table 7.1. A single electron transversing the gas layer can easily generate several hundred VUV electroluminescence photons at a moderate field. This large gain provides excellent




Figure 3.7: Electroluminescence yield measured in both cold xenon vapour and in room temperature gas [152]. The references in the key are those referred to in [152], with “this work” being the referenced paper.

sensitivity in the ionisation channel, with the signal being clearly distinct from primary scintillation pulses.

Studies of the physical processes in liquid xenon are areas of current research, most notably measurements of L_{eff} . There is still uncertainty about the low-energy behaviour of this parameter in xenon, but also the energy-dependent field suppression. For these purposes a study of the relative scintillation efficiency down to very low energies measured for different electric fields is required, helping to understand some of the issues discussed in Section 5.10. Such studies are important for the analysis of future WIMP search experiments using two-phase xenon.

The excellent properties of liquid xenon have resulted in both the single and two-phase techniques being applied to the search for WIMP dark matter. ZEPLIN-I utilised the moderate discrimination afforded by the differing scintillation pulse-shapes of electron and nuclear recoils, with the more powerful discrimination provided by the two-phase method being exploited by the ZEPLIN-II and III instruments.

3.2 Boulby Underground Laboratory

The extremely rare nature of WIMP-scattering events requires an extremely low-background environment. Although local sources of ambient neutrons and γ -rays can be mitigated in

surface laboratories, this is not the case for cosmic ray-induced backgrounds. In the UK, the best location for this purpose is the Boulby Underground Laboratory at the Boulby mine. Located on the North-Eastern coast of England, between Loftus and Whitby, the Boulby Mine has been mining potash and rock-salt since 1968, and is the deepest active mine in Britain. 1100 m below the surface the salt seam provides an excellent location for low-background experiments, since the halite rock in the salt seams is low in natural radioactivity. Measurements of rock-salt samples show an average contamination of 30 ppb U and 150 ppb Th [153]. The rock overburden (2800 m water-equivalent) provides an excellent shield to cosmic rays, reducing the muon flux by a factor of $\sim 10^6$ (to $(4.09 \pm 0.15) \times 10^{-8}$ muons/cm²/s [154]).

Earlier work by the UKDMC was carried out in the warm and dusty atmosphere of the mine in less-than-ideal laboratory space, but in 2000 new facilities were opened including over 750 m² of air-conditioned clean lab space, including support facilities, heavy-duty crange, power and network infrastructure. The dark matter experiments (ZEPLIN-I, ZEPLIN-II, ZEPLIN-III, DRIFT-I and DRIFT-II) are joined in the underground facility by a low-background radioactivity measurement set-up and the SKY experiment, exploring a link between cosmic rays and cloud formation [155], with future projects currently under consideration.

3.3 ZEPLIN-I

The ZEPLIN-I detector was a single-phase instrument with three photomultipliers detecting the scintillation from interactions in 5 kg of liquid xenon. It was installed underground in mid-2001, operating until late 2002. Calibrations of the detector were carried out during surface testing, but it was not possible to repeat them underground, due to issues with the detector cooling.

The target vessel was made of copper (Cu-101 oxygen free) and lined with a PTFE reflector to enhance the light yield. The instrument was cooled using a commercial Polycold unit circulating coolant around the target, liquifying the xenon. The liquid volume was viewed by three 8 cm PMTs (ETL 9265Q), operating near room temperature, viewing the liquid through quartz windows. The PMTs were located at the end of individual 4 cm turrets, helping to reject low-energy background from the PMTs them-




Figure 3.8: Schematic diagram of the Boulby Underground Laboratory. *Credit: UKDMC*

selves. To reduce background from neutrons and higher energy γ -rays depositing small amounts of energy through Compton scattering in the target, a 1.3 tonne liquid scintillator veto surrounded the instrument. This acted as a neutron absorber/moderator and also provided γ -ray shielding, with the instrument additionally surrounded by 25 cm of Pb to attenuate γ -rays from the cavern.

Data was acquired with 1 ns sampling allowing for accurate measurement of the fast scintillation pulses. To discriminate between electron and nuclear recoils, the difference in scintillation pulse shape is utilised. Figure 3.10 shows an example of the calibration of the instrument with an AmBe neutron source. The plots shows the distribution of the effective time constant estimator, τ_m , in 20-30 keV_{ee} energy range. Scintillation processes in xenon have three distinct time constants (from the singlet and triplet states and charge recombination) which combine to form an effective time constant, dependent




Figure 3.9: Schematic diagram of the ZEPLIN-I instrument located inside the liquid scintillator veto and Pb shielding [97].




Figure 3.10: Top: τ_m distributions from a neutron calibration of the ZEPLIN-I instrument [97], showing events in the observed energy range 20-30 keV from a tagged AmBe source. Bottom: Illustration showing the effect of adding nuclear recoil events to an electron recoil background population, producing a deviation in the low τ_m tail. [97]

on the event topology, measured by ZEPLIN-I as a mean arrival time, τ_m . The peak at lower τ_m results from nuclear recoils (where exciton luminescence dominates), with the electron recoil population having a higher τ_m (with recombination luminescence dominating). The clear difference shown allowed analysis of the shielded background data from the underground science runs of the detector.

The science data were analysed by assessing the lower edge of the electron recoil background population for a change of slope (as shown in Figure 3.10). If there were no nuclear recoils (and hence WIMPs) present, then the slope should fall smoothly, with no inflection. By statistical analysis of this tail for a deviation, an upper limit was placed on the number of nuclear recoils present.

The ZEPLIN-I detector collected background data during three separate science runs, which were combined accounting for variations in energy calibration and trigger threshold (both measured with ^{57}Co calibrations). The combined run time was 91.5 days, providing a final exposure, from the 3.2 kg fiducial mass, of 293 kg.days. From analysis of this data-set a 90% confidence limit on the spin-independent WIMP-nucleon scattering cross-section was set with a minimum at 1.1×10^{-6} pb at a WIMP mass of about 80 GeV, the best in the world upon its release [97].

3.4 ZEPLIN-II

Following the single-phase xenon search of ZEPLIN-I, the focus of the ZEPLIN collaboration shifted to two-phase systems. Occasionally a WIMP particle passing through the detector will scatter off a xenon atom producing a nuclear recoil. This recoil will create prompt primary scintillation (S1), through a combination of direct and indirect excitation, i.e. exciton and recombination luminescence. Some of the ionisation produced by the interaction can be drifted away from the interaction site by an applied electric field. Following emission into the gas-phase, electroluminescence is produced giving the secondary signal (S2). The time between S1 and S2 is that taken by the electron cloud to the drift to the surface, being a measure of the depth of the interaction in the liquid (z coordinate). The location of the interaction in x and y can be determined from the hit pattern of the S2 photons on an array of photomultipliers. Discrimination between electron and nuclear recoils is afforded by the ratio of ionisation and scintillation, due

to the structure of the tracks created by different interacting particles. ZEPLIN-II was the first stage of the two-phase programme, building on prototype R&D work at CERN [156, 157, 158], Imperial College London [159] and ITEP [160].

3.4.1 The ZEPLIN-II instrument

Figure 3.11 shows the geometry of the detector, with the 31 kg liquid xenon target held in a truncated PTFE cone (inner radii of 16.2 cm at the top and 14.2 cm at the bottom) and viewed from above by seven photomultiplier tubes, arranged in a hexagonal array, located in the gas. The cold liquid (-100°C) and gas layer are in thermal equilibrium with any change in liquid temperature producing a variation in gas pressure. The PTFE cone helps to define the active xenon volume and provides good reflection of the VUV photons helping to improve the light yield. As with most two-phase systems both the scintillation and ionisation channels are detected by an array of photomultiplier tubes which, for ZEPLIN-II, were ETL low-background quartz-windowed tubes (D742QKFLB) [161] with a room temperature quantum efficiency of about 17% at 175 nm.

The electric field was created by applying voltage to three grids defining two distinct field regions, the drift and extraction regions. The cathode grid was located at the bottom of the xenon volume with the two extraction grids at the top, either side of the liquid surface, defining the extraction region. The majority of the xenon lies in the drift region, forming the target for the WIMP search, where the field, responsible for drifting ionisation away from the interaction site, was produced by applying voltage between the cathode and lower extraction grids. A 1 kV/cm drift field was applied during science operations. Field shaping rings supported in the PTFE help define a uniform, vertical field in the drift volume, aiding charge collection from near the side walls.

A stronger field is applied between the two extraction grids in order to facilitate cross-phase emission with high efficiency and induce the production of electroluminescence in the gas. The field in this region of liquid is known as the extraction field; with a standard operating value was 4.2 kV/cm, the highest stable field achievable. The extraction grids also define the field in the gas, the electroluminescence field, which accelerates the electrons to produce secondary scintillation by collisional excitation.




Figure 3.11: Schematic diagram of the ZEPLIN-II detector.

Calculation of the extraction and electroluminescence fields, either side of the liquid surface, requires information about the liquid level (or, alternatively, gas thickness) and the dielectric constants of the gas and liquid phases. The field in the liquid, E_l , is given by:

$$E_l = \frac{V}{\left(\frac{\epsilon_l}{\epsilon_g} \times d_g\right) + d_l}, \quad (3.5)$$

where $\epsilon_l = 1.96$ and $\epsilon_g = 1.00126$ are the dielectric constants of the liquid and gas, d_g and d_l are the thicknesses of the gas and liquid layers (where $d_g + d_l = 1$ cm in ZEPLIN-II) and V is the voltage across the extraction grids. The field in the gas is

roughly double that in the liquid:

$$E_g = \frac{\epsilon_l}{\epsilon_g} \times E_l. \quad (3.6)$$

The xenon vessel was constructed from high-purity copper, to minimise the radioactive background, and was surrounded by a vacuum vessel of a cast stainless steel alloy. Rigorous consideration and testing of all materials located around the target was required to minimise internal backgrounds. To reduce the external background, the detector is surrounded by 25 cm of Pb γ -ray shielding and 30 cm of hydrocarbon neutron shielding (shown in Figure 3.12). The neutron shielding consists of a liquid scintillator veto surrounding the target plus a solid hydrocarbon roof, the same shielding setup used for ZEPLIN-I. The most important background for WIMP searches is that of neutrons, which can mimic a WIMP signal in the detector. Neutrons can result from cosmic-ray muon interactions, from spontaneous fission or (α, n) processes from uranium and thorium contamination in the cavern rock, as well as detector and shielding components, with the dominant source for ZEPLIN-II being the rock neutrons, followed by detector components. The expected single scattering neutron event rates for ZEPLIN-II were estimated by extensive simulations [162, 163, 164] and summarised in [101] resulting in a total estimated neutron background of < 40 events/year. This predicted event rate suggested that the detector should have been able to reach sensitivities better than those achieved at the time.

Additional rejection of background was possible using the active liquid scintillator veto. Most neutrons scattering in the xenon will also pass through the veto, producing a signal through scattering or radiative capture by the hydrogen. Neutrons which pass through the solid hydrocarbon roof (instead of the liquid scintillator) can be thermalised by the hydrocarbon and absorbed by Gd loaded into the blocks (0.2% Gd). Neutron captures in the gadolinium produce an average of 4 γ -rays with a total energy of 8 MeV, and these can be detected by the veto. By monitoring signals in the veto, looking for coincidences with events in the detector, it was estimated [101] that up to 60% of the expected neutron events could be vetoed. The veto also aids γ -ray rejection with 14% of the Compton-scattering events vetoed in the science dataset.

The target was cooled to -100°C using a Polycold IGC PFC320 connected to a liq-




Figure 3.12: Schematic diagram of the ZEPLIN-II detector within its shielding castle, where: A denotes the liquid xenon target, B is the liquid scintillator veto, C is the solid hydrocarbon roof and D is the 25 cm of Pb shielding.

uefaction head at the top of the detector. This cooled the target structure convectively, liquifying the xenon and additionally causing xenon ‘rain’. During operations the Polycold proved relatively unstable (due to a small leak of coolant in an inaccessible joint), resulting in continuous variations in temperature causing a variable xenon vapour pressure in the target, leading to additional operational and analysis difficulties.

The xenon gas was stored externally and delivered by a high-purity gas handling and purification system. The xenon was stored in two 27 litre stainless steel cylinders which could be cooled with liquid nitrogen to collect the xenon. This provides a mechanism to force xenon around the system, for example in ‘bottle-to-bottle’ purification and for emptying the detector. In the event of cooling failure of the detector, the liquid would vapourise causing the pressure to rise. If this were to go unchecked, the target vessel, connections or pipework could rupture resulting in damage to the instrument and loss of the xenon. As a result, burst disks were fitted between the detector and two large chambers, from where the xenon gas could easily be recovered in the event of a failure.

This also has the advantage that the impact on the cleanliness of the instrument and the xenon would have been reduced. During the operation of ZEPLIN-II the dumps were never required as a safety mechanism.

3.4.2 Xenon purification

In order to successfully measure the ionisation signal from throughout the full active volume, the concentration of electronegative impurities in the xenon must be extremely small (ppb level or below). These impurities will trap electrons as they drift through the liquid, reducing or removing the ionisation signal. The required level of purity is far higher than that available commercially, so additional purification is required.

In ZEPLIN-II, this involved passing the xenon gas through SAES getters (PS11-MC500 [165]), which remove contaminants such as O_2 , H_2O , CO and CO_2 . Purification of the xenon during initial commissioning utilised a ‘single-shot’ method, passing xenon gas through the getters during filling and emptying of the detector and additionally from bottle-to-bottle. During this period the electron lifetime (measured from data using the method described in Section 4.4.4) remained around $4\ \mu s$, which was far from adequate considering the full drift time in the active volume was $\sim 73\ \mu s$. Before filling the detector, the gas purity was measured independently with an Electron Lifetime Monitor (ELM), a purpose-built test chamber, and found to be of a significantly higher level ($\gtrsim 100\ \mu s$). This roughly constant, poor lifetime measured in the target was thus attributed to contamination within the target. In particular, the large amounts of PTFE (11 kg) created a high out-gas pressure, which was always likely to cause a problem.

The solution to this purity problem was to implement an active recirculation system, involving additional heaters (to boil off a small amount of liquid) and an external recirculation pump. As a result, the xenon could be circulated from liquid in the target, through a getter in the external gas system, and condensed back into the target. With this recirculation running at its maximum stable flow rate (17 lpm), a full charge of xenon ($\sim 40\ kg$) was purified every day. The main advantage of this recirculation path through the target was that a satisfactory electron lifetime could be achieved in spite of the continuous release of internal contamination. During initial operations with the recirculation system the pump could only be run when supervised, so during un-

manned periods a passive mode was engaged, utilising only the heaters to provide a lower flow-rate of ~ 2 lpm.

The affect of the forced recirculation on the xenon purity was dramatic, with a fast improvement observed after only a few days of recirculation. Figure 3.13 shows a timeline of the electron lifetime upon the application of the recirculation, with periods of active and passive operation noted. From the initial value of $\sim 4 \mu\text{s}$, the electron lifetime increased to about 1 ms within a week of operation. The improvement is approximately exponential, suggesting removal of contaminants at a constant rate. Once a high level of purity was achieved, a lower flow rate (~ 7 lpm) was established to keep the liquid surface stable for data-taking whilst maintaining acceptable purity level. The electron lifetime was maintained at about $100 \mu\text{s}$ for the entire 57 day science run.

3.4.3 Radon emanation

In background data acquired with a large energy range, high energy ($\sim \text{MeV}$) events were observed as soon as the detector was first switched on. These were initially identified




Figure 3.13: Timeline of the improvement in electron lifetime in ZEPLIN-II after implementation of the active recirculation system. The red regions denote periods of active recirculation (with pump and heaters) whilst the blue regions highlight passive recirculation (with heaters alone).

as α -decay events from their low ionisation to scintillation ratio. The proposed cause of these α -decays was radon in the liquid xenon or on the PTFE surfaces. This hypothesis was supported by investigation of the time delay between events, hinting at stages in the ^{222}Rn decay chain. However, confirmation of the source of these α -events had to wait until recirculation was halted following the science run.

The rate of α -decays in the detector was measured simply by counting high-energy events (from the size of the scintillation pulses, Figure 3.14). If the source of α -decays was ^{222}Rn introduced by the recirculation system, an effect would be expected when recirculation was halted. Following completion of the science run and calibrations the recirculation was stopped, with data-sets acquired monitoring the response of the purity. These data-sets were additionally utilised to monitor the rate of α -decays, with Figure 3.15 showing the decay rate over a period of 3 months, incorporating the end of the science run. The rate is initially constant at about $2.2 \alpha/s$, followed by a spike (linked with a change of getter) before it again returned to a constant rate of about $2.5 \alpha/s$. Upon termination of the recirculation the rate was seen to decay away.

^{222}Rn is part of the decay chain of ^{238}U with a constant production rate due to




Figure 3.14: High-energy peak (shaded) in ZEPLIN-II background data identified as α -decays, with the main electron recoil background at low-energies.




Figure 3.15: Timeline showing the rate of α -particle events measured in the ZEPLIN-II detector between 1st June and 19th Sept 2006, the period covering the end of the science run (\sim day 7).

the long half-life of ^{226}Ra decay ($T_{1/2} = 1600$ years). Since radon is a gas, it emanates slowly from materials and can therefore decay far from its production site. ^{222}Rn decays producing an α -particle ($E_{\alpha} = 5.59$ MeV), with $T_{1/2} \sim 3.82$ days, leaving ^{218}Po , which is also unstable. The dominant decays in the chain are detailed in Table 3.1, leading to ^{210}Pb which is the end point for the timescales of interest, having $T_{1/2} \sim 22$ years.

Table 3.1: Radon decay chain

Parent	Daughter	Decay	E_{α}	$T_{1/2}$	Branching ratio
^{222}Rn	^{218}Po	α	5.590 MeV	3.82 days	100 %
^{218}Po	^{214}Pb	α	6.115 MeV	3.1 minutes	99.98 %
^{214}Pb	^{214}Bi	β		26.8 minutes	100 %
^{214}Bi	^{214}Po	β		19.9 minutes	99.98 %
^{214}Po	^{210}Pb	α	5.617 MeV	164.3 μs	100 %

As the table shows the intermediate steps in the chain all have relatively short half-lives, meaning that ^{222}Rn will limit the speed of the process. From ^{222}Rn to ^{210}Pb , a total of 3 α -decays occur, such that the overall rate of α -decays is three times the ^{222}Rn activity.

Focussing on the period immediately post-recirculation, if the source of ^{222}Rn were

removed, at a time t_0 , the rate should be described by:

$$R = A_0 e^{-\ln 2 \cdot t / T_{1/2}} + C, \quad (3.7)$$

where A_0 is the initial rate, t is the time since t_0 and C is the rate of any residual α activity. Figure 3.16 shows the post-recirculation data fitted by this equation, resulting in $T_{1/2} = 91.95 \pm 2.33$ hours, which is excellent agreement with the value expected for ^{222}Rn (91.764 hours). The initial activity was found to be $A_0 = 2.889 \pm 0.029$ α/s , with a residual activity $C = 0.0344 \pm 0.0269$ α/s .




Figure 3.16: Rate of α events decaying after the end of recirculation. The data are fitted with an exponential plus constant showing a decay half-life consistent with that of ^{222}Rn .

Examining data from the commissioning phase (pre-recirculation), evidence was found for α -decays with the ^{222}Rn half-life soon after the detector had been filled through a getter. Additionally, the rate of α -decays was tracked in data following the science run when the purity was being recovered. Figure 3.17 shows the radon decay rates (equal to 1/3 of the α -decay rate) tracked over a period of 15 days. At the beginning of the period the target was isolated from the recirculator for a few days. After 24 hours the

recirculation was started ($t = t_0$) causing the rate to gradually increase. This could be described by:

$$R = A_0 e^{-\ln 2 \cdot t / T_{1/2}} + P \left(1 - e^{-\ln 2 \cdot t / T_{1/2}} \right), \quad (3.8)$$

where P is the production rate of the parent species (here ^{222}Rn) and A_0 is the initial rate. Using this equation and some reasonable values for A_0 and P along with the $T_{1/2}$ for ^{222}Rn , a model for the growth is derived (blue solid line in Figure 3.17). After about 120 hours the recirculation path was altered to pass through the alternate getter. This boosted the α -decay rate as ^{222}Rn would have been emanated and stored in the closed getter. When the xenon first passes through the getter, the stored radon is flushed into the detector. After the initial boost the radon decays faster than it is replenished, so the observed rate decreases towards an equilibrium level (the parent production rate), again as described by Eq. 3.8.




Figure 3.17: Timeline following the rate of α events from the decay of ^{222}Rn . The engaging of the recirculation is marked as is the time when the recirculation path was altered to pass through another getter. Models for the expected behaviour are shown by the solid lines.

From monitoring of the α -decay rate in the detector we found extremely strong evidence that the high energy events observed in the ZEPLIN-II data were produced by the radioactive decay chain between ^{222}Rn and ^{210}Pb . It was also determined that the

source of the radon was emanation from the SAES getters used for xenon purification. This was eventually backed up by testing a getter (of the same model) with an external radon detection setup, showing an emanation rate of 0.62 ± 0.04 Rn/s, in good agreement with observed radon activity ($A_0 \simeq 0.96$) in the detector (allowing for some variability between individual getters).

The presence of α -decays from the radon chain was recognised early on during the commissioning phase, although their implications were not fully realised until analysis of the final dataset was completed (see Section 4.8). Plating of α -emitters on the PTFE walls resulted in a low-energy background population, limiting the performance of the detector. As a result, emanation of radon should be considered and mitigated in future systems. This could be achieved either by utilising radon-free getters or by implementing a radon trap between getter and target.

A description of the ZEPLIN-II instrument has been provided highlighting design features important to the science operations of the instrument. Two operational considerations, electron lifetime and radon emanation, which I personally played a role in analysing, have also been discussed. The analysis of the dark matter search data is presented in Chapter 4 with additional xenon physics results discussed in Chapter 7.

3.5 ZEPLIN-III

The third phase of the ZEPLIN programme utilises the same two-phase technique as ZEPLIN-II, but with a different detector design. Construction of ZEPLIN-III took place at Imperial College, where surface tests were also carried out. Following installation and commissioning in the Palmer Underground Laboratory, Boulby, science data were acquired during the first half of 2008. A brief overview of the detector is provided here, concentrating on features of ZEPLIN-III designed to enhance its sensitivity.

Whilst utilising the same detection principles as its predecessor, ZEPLIN-III incorporates several features designed to enhance some of the most vital detector characteristics in dark matter searches. Figures 3.18 and 3.19 illustrate schematically the detector target geometry, with the active xenon region forming a shallow disc viewed from below by PMTs located in the liquid. Total internal reflection of scintillation light at the liquid surface and the flat geometry of the target boost the light yield of the instrument. This

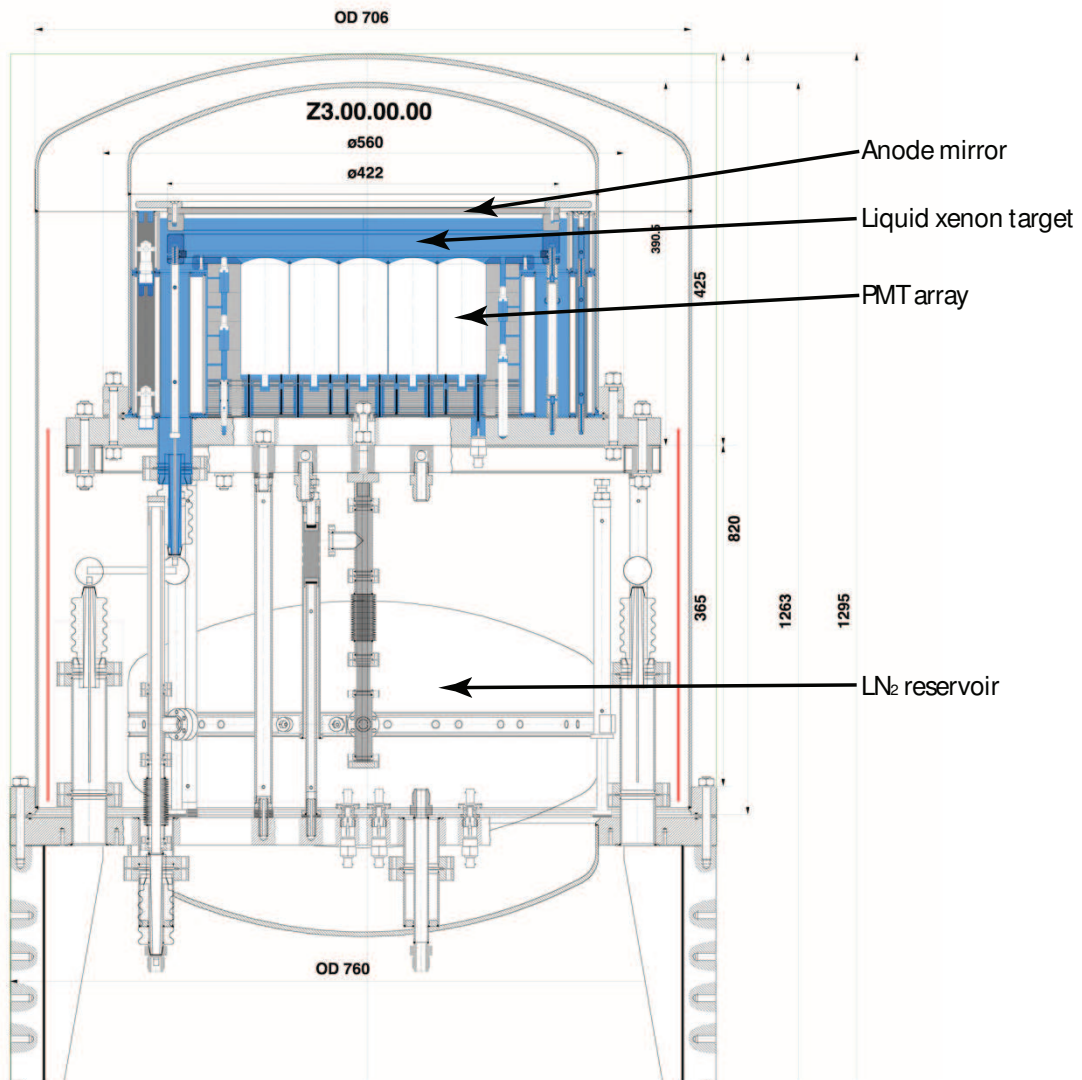


Figure 3.18: Schematic diagram of the ZEPLIN-III instrument.




Figure 3.19: Schematic diagram of the ZEPLIN-III target region, showing the PMT array and anode and cathode defining the target region. An example WIMP interaction is shown, with the production of S1 and S2.

higher light yield lowers the energy threshold of the detector, important when searching for WIMP-nucleon scattering with an expected exponentially falling energy spectrum. The light yield achieved at zero-field was roughly three times that in ZEPLIN-II, providing a significant improvement in threshold.

Another feature in the design of ZEPLIN-III to enhance the signal-to-noise for nuclear recoil detection is the expected improved discrimination between the background and signal populations from the application of a stronger electric field. ZEPLIN-III employs both of these methods, combining to improve the sensitivity of the instrument.

ZEPLIN-III operates with a higher drift field, ~ 4 kV/cm in the liquid, with the aim of enhancing the intrinsic discrimination between electron and nuclear recoils. Figure 3.20 shows simulations of the electrostatic field in the xenon, exhibiting excellent uniformity within the central region above the PMTs. The field in the active volume is produced by applying voltages to the cathode grid and anode mirror (with no need for dedicated extraction grids), where the fields in the liquid and gas can be calculated using Eq. 3.5 and Eq. 3.6, with a gas gap of approximately 4 mm and a drift region of ~ 36 mm. The PMT grid, set to the same voltage as the PMTs, creates a strong reverse field region between the PMT and cathode grids. This reverse field region was designed to aid rejection of low-energy background from the PMTs and to protect the

input optics of the phototubes.

The active xenon region is viewed by an array of 31 2-inch low-background PMTs (ETL D730/9829Q). The array provides excellent 3D position reconstruction (see Section 5.4) aiding in the rejection of background from surfaces and multiply scattered events. Figure 3.21 shows the PMT array with an example of a double-scatter event, where the finer resolution of the array allows the resolution of the two individual scatters. Additionally, the enhanced position resolution facilitates better definition of a low-background central region. The individual performance of the PMTs was characterised with low-temperature testing before installation [166], measuring the individual PMT quantum efficiencies (at xenon scintillation wavelengths) and gains. In the detector they are surrounded by a tight-fitting copper screen and topped with a highly polished “PMT mirror”, to prevent optical cross-talk.

To achieve the low background required, two aspects of the detector are considered. Within the lab the detector target is located inside a shielding castle to reduce background generated in the cavern walls. ZEPLIN-III had a purpose-built shielding castle consisting of 20 cm thickness of high-purity Pb and 30 cm of polypropylene shielding, attenuating the background of both γ -rays and neutrons from the cavern walls by a factor of 10^5 .

The remaining problematic background for the detector is the internal activity of detector materials and components. Consequently, careful consideration was given to the selection of detector materials and clean-room procedures were employed whilst working on the instrument. Oxygen-free copper (type C103) was used to construct the target and vacuum vessels, along with any internal structures where this was possible; this was




Figure 3.20: Simulation of the electric field in the ZEPLIN-III target, demonstrating the excellent uniformity of field in the central region. The central target region (light blue) is at 3.9 kV/cm, with the field in the gas (green) being double.




Figure 3.21: Two coincident events shown in the individual PMT waveforms (left) and the positions of the two events on the PMT array (right), obtained during early surface commissioning tests in 2006.

extremely challenging and expensive due to the difficulty in machining copper with high precision and the cost of the raw material. For other components, the materials and batches involved were individually screened, with many undergoing high-precision radioactivity measurements with a low-background HPGe detector setup underground at Boulby. These measurements allowed for the component contribution to the background to be calculated from simulations. A record of all detector components is maintained, accounting for their contributions to the γ -ray and neutron backgrounds.

As discussed for ZEPLIN-II, electronegative contaminants in the xenon will trap ionisation electrons as they drift through the liquid, introducing a depth dependence on the S2 signal. As with most variations, a correction can be applied to the data *a posteriori*, although it adds yet another factor of uncertainty. During the commissioning phase and first science run of ZEPLIN-III no recirculation system had been implemented, so purification was carried out with a single-shot method. Repeating this procedure, transferring the xenon gas between two stainless steel cylinders through the bakeable all-metal gas system and passing through a SAES getter [165], allowed an acceptable electron lifetime to be reached for the first science run. It was believed that ZEPLIN-III would not require any recirculation during live periods as the xenon-friendly materials and extremely good cleanliness should prevent the purity from degrading. This was shown

to be correct during the first science run, where the electron lifetime even increased steadily during the run! The radon problem encountered by ZEPLIN-II was avoided in ZEPLIN-III, as the xenon was purified bottle-to-bottle where the radon could decay outside the target.

The major cause for the varying response of ZEPLIN-II was the unstable pressure in the target, a consequence of problems with the cooling. ZEPLIN-III uses liquid nitrogen for cooling, with a liquid nitrogen reservoir sitting below the target (see Figure 3.18) connected through two thermal couplings: a braid made from bundles of thick copper wire which dips into the liquid nitrogen and a flow of cold nitrogen boil-off gas into a hollow cooling flange (attached to the bottom flange of the target). The flow of the boil-off gas can be controlled, with a large flow utilised for initial cooling of the instrument, before fine control of the flow allows the temperature to be kept stable. This method has proved very successful with the temperature controlled to within 0.2°C during the surface tests, and a variation $< 0.5^{\circ}\text{C}$ over the whole science run, keeping the xenon pressure stable to within 2%.

Learning from the operation of ZEPLIN-II, most of the major pitfalls were avoided, with the detector proving exceptionally stable. ZEPLIN-III was the first high-field dark matter search instrument in the world and, following optimisation, operated for several months with the fields on and stable. During this period, calibration and WIMP search data were acquired, the analysis of which is described in Chapter 5.

3.5.1 Future plans

Following the first science run of ZEPLIN-III, upgrades to the instrument are underway to further improve its sensitivity before an extended second run. The major component upgrade is the replacement of the PMTs with pin-compatible ultra-low background tubes developed for this purpose by ET Enterprises Ltd [161]. The new PMTs should reduce the background from the PMT array by a factor of ~ 40 , meaning that they will no longer be the dominant source of background in the detector. An active veto will also be installed, used to tag coincident events, aiding rejection of the neutron and γ -ray backgrounds.

Additional upgrades to aid detector stability, automatic operation and calibration are

also being implemented. The reduced background, enhanced operation and extended data-taking period, along with the excellent high-field discrimination observed in the first science run (see Chapter 5), should allow ZEPLIN-III to probe a significant portion of the parameter space favoured by SUSY models. It has also allowed us to validate a number of design solutions which will be very useful for the next generation detectors now being proposed.

Chapter 4

ZEPLIN-II analysis: From raw data to a WIMP search dataset

ZEPLIN-II was the first large-volume, low-background two-phase xenon instrument in the ZEPLIN programme and indeed in the world. Consequently, the analysis of the resulting data provided new challenges in understanding both the scintillation and ionisation signals. The WIMP target consisted of ~ 31 kg of liquid xenon held in a truncated PTFE cone, viewed from above by an array of 7 photomultiplier tubes located in the gas phase. The scintillation produced by interactions in the liquid was detected by the PMTs, forming the primary pulse (S1), with some of the ionisation charge recombining to add to the prompt scintillation signal. The remainder of the ionisation electrons are drifted away from the interaction site by the applied electric field, towards the liquid surface, into a region where a stronger field induces cross-phase emission of the drifting electrons. Once in the gas phase, the electrons produce electroluminescence as they are accelerated through the gas, which the PMTs detect as a delayed signal, forming the secondary pulse (S2).

Discrimination between the electron recoil background and the expected WIMP-nucleon scattering (nuclear recoil) signature is achieved through assessment of the ionisation-to-scintillation ratio of the events, $S2/S1$. As described in the previous chapter, nuclear recoils produce a reduced ionisation signal (in comparison to electron recoils) as a result of their different track topology. Figures 4.1 and 4.2 show example waveforms of typical electron and nuclear recoil events with the same observed energy, i.e.

the same S1 pulse area.

The detector was commissioned underground in late 2005, with data taken to assist optimisation of detector operation and characterisation of detector response. Subsequently, ZEPLIN-II acquired WIMP search data and associated characterisation and calibration data during 2006, with the analysis completed and the final spin-independent [101] and spin-dependent [167] results published in 2007.

Described here is the complete analysis procedure from the raw data to the final WIMP search dataset. I was directly involved in many key aspects of the analysis, to which I will give particular focus, although all areas of the analysis will be described in the requisite depth to provide a complete description of the process.

Figure 4.3 illustrates the complete ZEPLIN-II analysis chain from the relevant raw data-sets to the final limit on the WIMP-nucleon spin-dependent cross-section. The raw data requires processing to identify and parameterise pulses in the recorded waveforms. S1 and S2 pulses must be identified, as a robust description of all recorded pulses is essential for further analysis. Energy and position calibrations are required along with the calculation of corrections to the data, normalising the response of the instrument in both event location and time. A calibrated and equalised data-set facilitates characterisation of the detector response for the observed electron recoil background and the expected WIMP nuclear recoil signal, allowing the science background data to be analysed for the presence of a WIMP signature. The description of the statistical methods utilised for this final step and the process of calculating the WIMP-nucleon spin-independent cross-section are both given in Chapter 6.

4.1 Data acquisition and pipeline

Both scintillation and ionisation processes produce VUV photons ($\lambda \sim 175$ nm) which are detected by the array of seven photomultiplier tubes. As a result, the S1 and S2 signals from an interaction can be recorded in the same waveform with the delay between the two being the time taken for the electron cloud to drift through the liquid. The depth of the liquid (~ 14 cm) and drift velocity at the operating field of 1 kV/cm, $v_d \sim 2$ mm/ μ s, determine a maximum drift time of ~ 70 μ s in the main bulk. To allow for assessment of waveforms before S1 and after S2, 100 μ s was recorded before and




Figure 4.1: Typical single-scatter electron recoil event in ZEPLIN-II, shown in the smoothed sum channel (top). Pulse p1 (left) is a ~ 16 keV_{ee} prompt scintillation pulse (S1) and p2 (right) is the delayed secondary scintillation pulse (S2).




Figure 4.2: Typical single-scatter neutron recoil event in ZEPLIN-II, shown in the smoothed sum channel (top). Pulse p4 (left) is a ~ 16 keV_{ee} prompt scintillation pulse (S1) and p5 (right) is the delayed secondary scintillation pulse (S2).

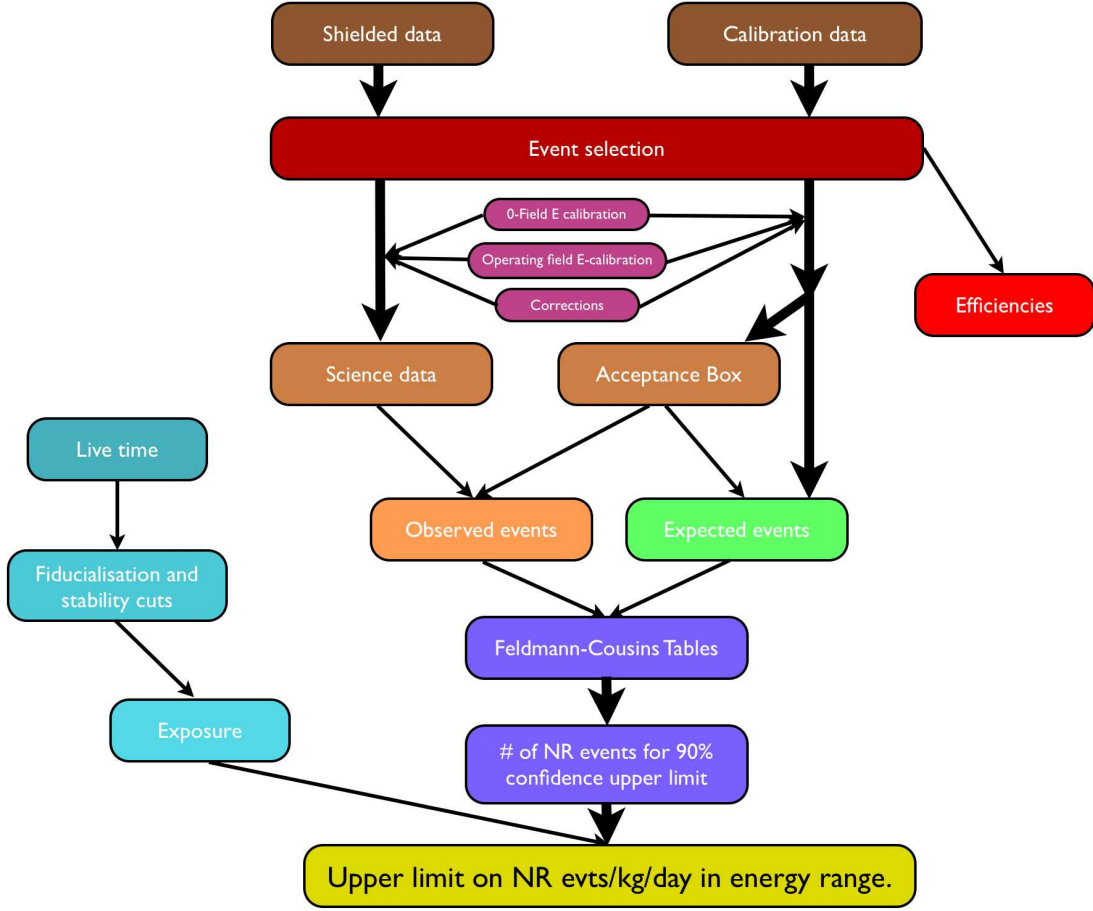


Figure 4.3: Flow diagram showing the connections between the different parts of the analysis and how they combine to produce the science result.

after the trigger point, giving a $200 \mu\text{s}$ waveform allowing for both S1 and S2 triggered events.

In hardware, the seven PMTs have their gains equalised to $\sim 4.5 \text{ mV}$ per photoelectron, improving uniformity of response across the array. The signals from the PMTs are passively split (using a 50Ω splitter), with one line feeding the data acquisition and the other feeding the trigger. The trigger requires 5 of the 7 PMTs to register a signal (after $\times 10$ amplification) above a 17 mV threshold (approximately $2/5$ of a single photoelectron signal). A high-level inhibit is enforced in hardware, based on the signal in the central PMT, to decrease dead-time by preventing triggering of the acquisition for 1 ms after a large pulse. This requirement results from optical feedback effects, producing signals large enough to trigger the system, observed in the detector when a

large amount of light floods the chamber. Large pulses are also vetoed (in software) during acquisition of low-energy data, as they will be above the energy range of interest for WIMP searches. The summed channel of 10 signals from the active veto system is also recorded, tagging events with coincident scatters in the liquid scintillator.

The philosophy of the ZEPLIN-II electronics is to trigger low-energy events from the S2 channel, resolving the small associated S1 signals by looking back through the waveform. For higher energy events the S1 pulse will be large enough to satisfy the 5-fold requirement and thus will trigger the acquisition. As the instrument may be triggered by S1 or S2, 100 μs must be recorded either side of the trigger point, resulting in a 200 μs waveform ($-200 \mu\text{s}$ to $0 \mu\text{s}$) with a central trigger point ($-100 \mu\text{s}$).

Events triggering the system have their $8 \times 200 \mu\text{s}$ waveforms (7 PMTs and the veto) digitised by an 8-bit ACQIRIS system (DC265 M2M digitisers housed in a CC103 crate [168]), sampling at 2 ns intervals. The system was controlled and data written to disk by LINUX-based software, storing 2000 events per raw data file, each consisting of 8 waveforms of 100,000 points. After compression the ~ 250 MB raw data files are transferred to two duplicate 100 GB magnetic digital tapes using a ADIC Scalar 100 tape robot system. One tape was transported to the University of Sheffield for data processing and analysis, with the other archived at Boulby for backup purposes. The data was additionally transferred to an analysis computer underground for rapid assessment of the daily data. During science data-taking, 25 GB of compressed data were acquired daily, with a considerably greater volume acquired during calibration periods.

4.2 Data reduction

The recorded waveforms require processing before they can be easily utilised for characterisation of the instrument. The waveforms must be assessed for the presence of pulses, which are subsequently parameterised, providing a set of numeric values representing each pulse. This reduction of data is completed by a LINUX-based application (UNZAP2) which allows for visual scanning of events (as shown in Figure 4.4) along with the reduction algorithms used to process the waveforms. These algorithms require a set of input parameters to optimise pulse-finding and correct for detector effects. The algorithms for data reduction underwent thorough evaluation and development during

the data analysis phase. The efficiency and sensitivity of this code is vital for the whole analysis, with rare-process searches potentially affected by flaws in the code affecting one event in a thousand or less. Assessment of the algorithms was carried out through visually scanning of large numbers of events.




Figure 4.4: UNZAP2 - The ZEPLIN-II data reduction and visualisation tool. The waveforms from the 7 PMTs are shown along with the sum channel and the signal from the veto.

Differences in the lengths of the signal cables and the characteristics of the individual PMTs result in a delay between the pulse arrival time in different channels. These delays must be accounted for when combining the seven waveforms to form a sum channel, otherwise a distortion of the pulses in the sum channel is introduced leading to incorrect pulse parameterisation, an effect particularly significant for small and fast S1 signals.

To aid pulse-finding on the waveforms smoothing is carried out, reducing the significance of small-amplitude noise fluctuations in the baseline. The time-scale of this smoothing must be optimised, with different values between 5 ns and 50 ns evaluated. For excessive smoothing, low-energy events are ‘washed-out’, removing those of greatest interest for a WIMP search. After evaluation, a value of 12 ns was chosen, allowing for successful reproduction of the single photoelectron (spe) spectra. This smoothed sum channel is used only for the identification of pulses, with the raw waveforms interrogated

for subsequent parameterisation.

Pulse finding proceeds by identifying excursions above a user-defined software threshold, which depends on the full-scale acquired, with a 2 mV ($\sim 1/2$ spe) threshold applied for a 200 mV full scale. For small pulses the waveform can briefly dip below threshold before the end of the pulse, this causes the signal to be split into several individual pulses. As a consequence ‘clustering’ is required: grouping pulses in close proximity into a single pulse. Figure 4.5 shows the effect of clustering, with a clustering time-scale of 50 ns leaving an obvious single pulse split into 4 separate sections (p3, p4, p5 and p6), which are correctly grouped together and identified as a single pulse by the optimised time-scale of 400 ns.

The identified pulses are subsequently parameterised, with up to a maximum of




Figure 4.5: Example of the clustering algorithm applied to the data during pulse finding. A time-scale of 50 ns (top) leaves the pulse as separate components (p3, p4, p5 and p6), whereas the optimised 400 ns time-scale clusters them together to form a single pulse.

10 sets of pulse parameters recorded to HBOOK ntuples [169]. The summed channel is utilised to define the start and end time of the identified pulses, allocating a pulse time window for each. The seven individual waveforms are interrogated within this time window to define the contribution of each to the total signal and parameterise the pulses, with the start time, amplitude, area (pulse integral within the time window), width (start to finish) and full-width at half maximum (FWHM) recorded for every pulse. Additionally, the mean charge arrival time, τ , is calculated for each pulse, as a simple estimator for the scintillation time constant for S1 signals. The number of pulses found and waveform parameters (baseline level, *rms* noise and threshold) are recorded per channel, with the event number, time and operational parameters (xenon pressure, target and ambient temperatures) also recorded for each event.

The resulting “reduced” ntuples contain the parameters for the pulses identified in each of the seven channels along with parameters for the sum channel. They also include a veto flag indicating the presence of a time-coincident signal from the liquid scintillator veto.

4.3 Event selection

The data reduction assigns no physical definition to the pulses, with all pulses parameterised in a consistent manner. They must subsequently be evaluated as candidate S1 or S2 signals (or neither), utilising the distinct physical properties of the two signal types. The prompt scintillation (S1) consists of a fast exponentially decaying pulse ($\tau \sim 30$ ns), located either at the trigger point (S1-triggered event) or in the preceding ~ 73 μ s. The electroluminescence process produces a wider pulse (S2) as the electrons induce secondary scintillation drifting through the gas layer, with a characteristic width of ~ 1 μ s (dependent on the thickness of the gas layer and the electric field there).

WIMP-nucleon interactions are expected to be extremely rare, consequently a WIMP traversing the xenon target will interact only once (if at all), unlike neutrons or high-energy γ -rays, which can scatter multiply in the liquid. As a result, we can define the “golden” rule: an event must have only a single S1 and a single S2. A double scatter event consists of two time-coincident S1 pulses, indistinguishable in the recorded waveforms, and two distinct S2 pulses (assuming different depths for the two interactions).

Even if a waveform has two unrelated interactions (i.e. two S1 and two S2 pulses) it is nearly impossible to reliably determine which S1 and S2 are related, meaning that the event must be discarded - this case is very rare anyway. Despite cutting a significant fraction of events in the data, the golden rule has no efficiency penalty for the WIMP search.

An initial identification of S1 and S2 is carried out using a set of loose, inclusive cuts to retain high-efficiency, before application of the golden rule. Tighter cuts are subsequently implemented to constrain the quality of the events, removing outliers in several parameters and anomalous events. The identification process proceeds by finding single S2 pulses, which is relatively straightforward due to their large size and width, before the time window preceding the identified S2 can be examined for the presence of a single S1.

An S2 pulse must be wide, so a lower cut of 150 ns is placed on the τ parameter (charge mean arrival time). For the purposes of defining the width of a pulse, τ was found to be the most robust parameter, being largely unaffected by the often noisy trailing edge of the pulse. By allowing such a large range of τ at this early stage of event selection, no significant loss of efficiency is incurred. Small electroluminescence signals are observed in the data from the emission of the single electrons. However, these are not ‘WIMP-like’ S2 signals (i.e. not associated with ionisation from a nuclear recoil interaction) and should be ignored in the S2 pulse count. Consequently, a minimum area cut of 1 nVs is enforced, ignoring pulses with fewer than ~ 27 photoelectrons (a single electron produces a mean of 8 photoelectrons). The efficiency of this cut is energy-dependent, quantified by fitting the S2 distribution for events with an S1 in different energy ranges (5-10 keV, 10-15 keV, 15-20 keV and 20-25 keV). Integration of the distributions above 1 nVs yields the efficiency at different energies; 90.3% for 5-10 keV, 99.7% for 10-15 keV and $\simeq 100\%$ for higher energies. S2 signals must also occur in a physically allowed region of the timeline. S2-triggered events will have S2 located at the trigger point ($-97 \mu\text{s}$), with larger S1-triggered events placing S2 between the trigger point (for surface events) and the maximum physical drift time of $\sim 73 \mu\text{s}$. For events at the bottom of the liquid this results in an allowed time window for S2 pulses of $-97 \mu\text{s}$ to $-24 \mu\text{s}$, as shown in Figure 4.6. Requiring the pulse to lie in a physically allowed region of the waveform incurs no efficiency penalty.




Figure 4.6: Waveform highlighting the physically allowed regions for S1 and S2 pulses.

After identification of an S2 signal, pulses were evaluated as candidate S1 signals. These fast scintillation pulses will be narrow, selected with a cut on the mean charge arrival time: $3 < \tau < 150$ ns, complementing the τ cut for S2 signals. This complementarity is important to ensure that pulse identification has no loopholes, allowing pulses to escape being counted. To additionally exclude noise fluctuations and PMT-related artefact signals, observed in a single channel, a requirement that at least 3 PMTs contain pulses above a software threshold of 1.7 mV (~ 0.4 phe) is applied. As with the minimum area requirement on S2, the 3-fold requirement is energy-dependent, cutting low-energy events more preferentially. The energy-dependent efficiency, shown in Figure 4.7, was calculated through dedicated simulations accounting for PMT single photoelectron response (SER), light collection uniformity and energy resolution. Finally, a candidate S1 pulse must lie within a physically allowed region of the timeline, preceding S2 within the maximum drift depth, $0 < \text{dtime} < 73 \mu\text{s}$, as S1-like pulses before this time window must be unrelated to the identified S2.

These simple selection cuts provide pulse identification of S1 and S2 signals, allowing application of the golden rule. These golden events, with one S1 and one S2, are subsequently subjected to several further quality cuts.




Figure 4.7: S1 trigger efficiency calculated from simulations as a function of energy.

An additional drift time cut is implemented to remove events originating near the liquid surface, in the higher-field extraction region or around the lower extraction grid ($dtime > 4 \mu s$). The efficiency of the additional drift time cut is incorporated later during calculation of the fiducial target mass.

The pulse-width cuts are tightened, removing events from the tails of the distributions. For S1 a cut of $3 < \tau_{S1} < 50$ ns effectively selects high-quality scintillation pulses. The efficiency of this tightening was assessed with neutron calibration and background data, showing $\simeq 100\%$ efficiency in the energy range of interest. A similar approach is followed for S2, but the standard τ estimator is first corrected for the effects of diffusion of the electron cloud during its drift to the liquid surface. Diffusion of the electron cloud along the drift direction causes a distortion of the ideal square pulse shape producing a smoother, Gaussian-like one. This broadening effect is demonstrated in Figure 4.8, showing an increase of τ_{S2} with depth, widening the overall τ_{S2} distribution. Applying a depth-dependent correction to τ_{S2} , as shown in the figure, allows for a tighter cut on secondary width, $220 < \tau_{S2}^{cor} < 360$ ns. Application of this cut to the AmBe calibration

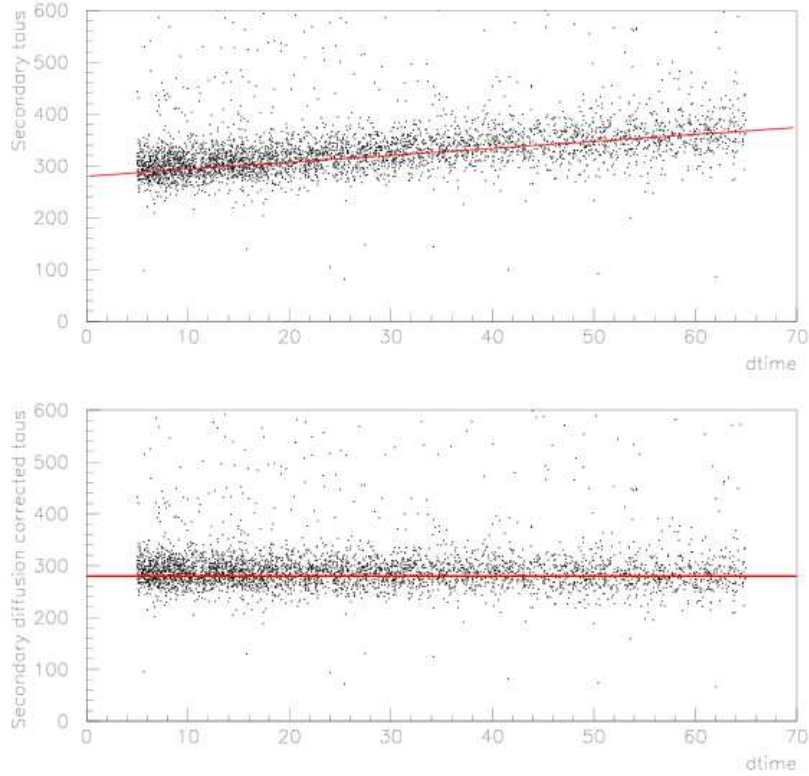


Figure 4.8: The diffusion effect and empirical correction demonstrated through τ_{S2} . τ_{S2} uncorrected (top) and corrected (bottom) for diffusion.

data yielded an efficiency of 90.2% for nuclear recoils in the central region of the detector below 30 keV_{ee}.

The final cut in the initial selection process is based upon the peaking time of the S1 pulse in different channels. After correction for cable delays, the contribution to S1 from different PMTs should coincide with high precision, even allowing for statistical effects with small numbers of photons. Some anomalous events were observed for which unusual structure in the S1 pulse was characterised by a large separation (in time) of the peaks in different PMTs. For the purpose of removing these events, a parameter ‘tpeak(*i*)’ is defined as the peak time in each channel. Events were subsequently rejected if the maximum difference between these values was $\Delta t_{\text{peak}}(i) > 140$ ns. Assessment of the nuclear recoil calibration data yielded an efficiency of $\simeq 100\%$ for the *tpeak* cut.

Beyond the golden selection cuts, two further quality cuts were applied to the data selecting high-quality events. The requirement for both S1 and S2 pulses to be located

within the physical regions of the waveforms does not exclude the possibility of a pulse triggering the system but not being identified as the S1 or S2 in a golden event. To remove such anomalous occurrences, a requirement is placed on the golden events stating either S1 or S2 is located at the trigger point, between $-97 \mu\text{s}$ and $-96 \mu\text{s}$, shown in Figure 4.9. The efficiency of this cut is $\simeq 100\%$, with very few golden events falling outside this window.

A population of nuclear recoils were observed originating from the cathode as well as lower extraction grid wires, often with small S1 signals failing the 3-fold S1 requirement. This is expected since the wires obscure interactions occurring just below them from where very few S1 photons can be collected. These events fall outside the fiducial volume and should be excluded from the final data-set, but accidental coincidences between S1-only events and apparent S2-only events (as S1 is not identified) cause mismatched events to be wrongly placed within the fiducial volume. The S2/S1 ratio of these events is low as the two signals are unrelated, with S1 larger than that of the missed two-fold coincidence. To reject such accidental sub-threshold cathode events, a two-fold search is applied in a drift time window $66 < \text{dtime} < 73 \mu\text{s}$ before the S2 signal, checking for sub-threshold S1 pulses. The efficiency of this requirement was assessed by moving the $7 \mu\text{s}$ window through the target volume (searching for two-fold signals), and calculating the fraction of events removed in each window, as shown in Figure 4.10. The region containing cathode events saw $\sim 3.6\%$ rejected, but in the fiducial volume, the requirement removed a constant 0.3% of events. The efficiency for this cut was defined conservatively as 99.7% , although it is likely that some of those rejected are also coincidences with small sub-threshold events, correctly removed by the code.

4.4 Corrections

For all events recorded over an extended data-taking period to be considered together as a single data set, their response in S1 and S2 must be consistent for equivalent interactions. Measures were implemented during the design and operation of ZEPLIN-II to promote a uniformity of response throughout the target volume and the data-taking period. First-order equalisation of PMT gains is implemented in hardware, although this equalisation is fine-tuned further in analysis. Likewise, the xenon is purified to improve




Figure 4.9: Positions of S1 (left) and S2 (right) for golden events, with the window for the trigger cut shown (red lines).




Figure 4.10: Fraction of events removed by a coincidence cut (i.e. a two-fold S1 search within a 7μ s window) applied at different depths throughout the volume.

S2 uniformity over the full detector depth, but again event-by-event correction of the data is required. Uniformity of response through the data-taking period is dependent upon the stability of the detector: keeping pressure, temperature and applied electric fields constant is of paramount importance. Unfortunately, due to stability problems in the cooling system, the pressure varied greatly over the period of data-taking (from 1.2 to 1.9 bar), requiring significant correction of S2 signals. Variation of the applied electric fields will affect detector response, for both S1 and S2. Although, the fields remained relatively constant throughout the data-taking period, requiring no corrections, some instabilities were observed, resulting in the removal of data from certain days. These factors that can affect the detector response have been evaluated, with the signals from the detector corrected on an event-by-event basis as required.

4.4.1 Photomultiplier gain equalisation

The response of individual photomultiplier tubes can be assessed by measuring the size distribution of their single photoelectron (spe) signals. Single photoelectrons can be spontaneously emitted from the PMT photocathode and multiplied by the dynodes, with their magnitude being a measure of the PMT response to a single quantum of scintillation. Equalisation of PMT gains is achieved in hardware by adjusting the applied voltage slightly for each individual PMT, resulting in a more uniform response from the array. This uniformity can be improved further by measuring the spe spectra for the 7 individual PMTs and using the mean response from each tube to construct a weighted sum for both S1 and S2. Figure 4.11 shows an example of the spe spectra measured in the 7 PMTs, with similar measurements made through the data-taking period in order to monitor stability. The mean spe area was 30–40 pVs for all the PMTs, with a small variation in spe response resulting in a time-dependent weighting being applied to the data.

4.4.2 Light collection uniformity

The design of the ZEPLIN-II target, using reflective PTFE surfaces, was aimed at providing good uniformity of light collection throughout the volume. To confirm this uniformity or allow further correction of any spatial variation found, one requires a

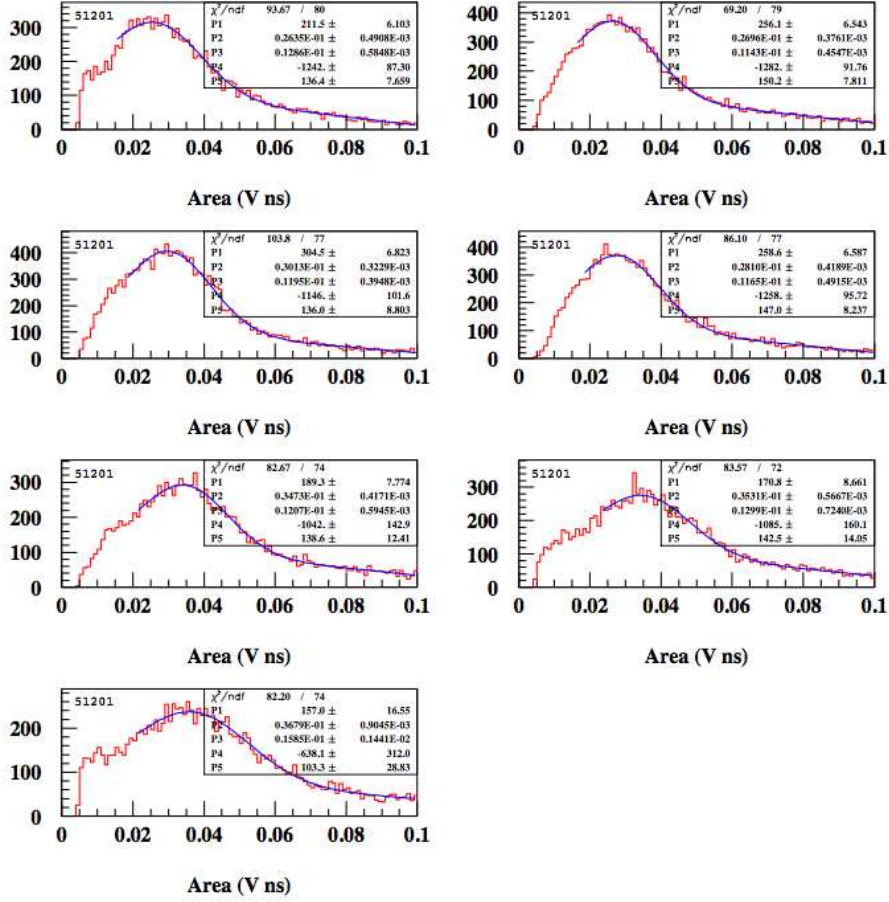


Figure 4.11: Typical single photoelectron area spectra measured for the seven individual PMTs in ZEPLIN-II, used for equalisation of PMT gains.

feature in the energy spectrum. In ZEPLIN-II this is conveniently provided by the α -decay population from the decay of ^{222}Rn .

Fitting the peak in the S1 spectrum from α -events at different depths in the target volume allows the response to be mapped as a function of z . Figure 4.12 shows the mean S1 response for different drift times within the central region of interest. Simulation of the instrument had suggested a slight increase ($\sim 3\%$) in light collection efficiency towards the bottom of the chamber (dependent on the xenon purity, through variation of the photon absorption length). However, measurements with the α -decay population showed a slight increase in the middle of the detector. This variation is small (up to $\sim 3\%$ from the mean) and as a result no correction was applied.

A similar approach to measurement of the light collection as a function of radial




Figure 4.12: Mean S1 response for α -decays at different drift times, highlighting the uniformity of S1 light collection in z .

position was not possible, as the corresponding S2 signals saturate the PMT photocathodes, resulting in a significant bias of the centroid position reconstruction algorithm. However, light collection simulations (which gave reasonable agreement in z) suggest any variation in efficiency to be small in the xy plane. Likewise, the S2 light collection efficiency was calculated from such simulations, and found to decrease away from the centre of the instrument. This decrease is small within the central region of interest (from 27.5% at $r = 0$ mm, to 24% at $r = 100$ mm), and no radial light collection corrections were deemed necessary.

4.4.3 Pressure correction

During operation of ZEPLIN-II, a continual small coolant leak resulted in instability of the cooling power, causing the temperature and pressure in the target to vary. The target pressure was monitored throughout the science data-taking period, varying between ~ 1.2 bar and ~ 1.9 bar, with thermal equilibrium between the liquid and gas

providing a direct link between liquid temperature and gas pressure. Variation of liquid temperature causes a small thermal expansion or contraction of the liquid, changing the thickness of the gas gap, and as a result varying the amount of electroluminescence produced by the emitted electrons. The inverse relation between pressure and S2 response is magnified by the microscopic electroluminescence process as $Y \propto -P$, as discussed in Section 3.1.




Figure 4.13: Pressure dependence of S2 response as a function of target pressure, measured with the ^{57}Co combined 122 keV and 136 keV photopeak during a dedicated calibration run.

The exact relationship between pressure and S2 response was evaluated with a specific ^{57}Co calibration, where the target pressure was deliberately varied between ~ 1.2 bar and ~ 1.9 bar. The photo-peak in the S2 spectrum was located at different pressures (as shown in Figure 4.13), exhibiting the expected relationship. From these measurements an empirical correction was determined.

4.4.4 Xenon purity correction

Electronegative contaminants in the liquid xenon suppress the ionisation signals in two-phase systems through trapping of electrons during their drift to the liquid surface, causing a drift-time-dependent S2 response. Figure 4.14 shows this suppression measured with α -events located throughout the volume. Due to the energy-dependent behaviour of the S2/S1 ratio, it is advantageous to utilise a feature in the energy spectrum (or a tight energy range) to measure the suppression. Using the α population, the S2/S1 distribution is sliced in drift time and fitted by a log-normal distribution. An exponential fit is made to the distribution means, acquiring a function for the depth dependence:

$$\frac{S2}{S1} = A \times e^{-\text{dtime}/\tau_e}, \quad (4.1)$$

where A is the unsuppressed S2/S1 value and τ_e is the mean electron lifetime in the liquid. Section 3.4.2 described the xenon purification which took place in ZEPLIN-II and the dramatic improvement in electron lifetime achieved with active recirculation.

The xenon purity was monitored with specific background runs using a large acquisition full-scale, recording the high-energy α -events specifically. This method was complemented by a second technique for the science background data-set. The population of nuclear recoil events observed on the cathode and lower extraction grids was attributed to the decay of α -emitters deposited on the grids, where only the nuclear recoil component of the decay exits the wire. It is assumed that these two populations will both produce the same S2/S1 ratio, as a result of the same production mechanism, allowing the electron lifetime to be attained using a two-point exponential fit to these two populations.

A database of electron lifetime was populated using the two-point exponential fitting method for 4-5 hour slices of the science data-set. Figure 4.15 shows the gradual decrease in electron lifetime over the science run, from $\sim 150 \mu\text{s}$ at the start of the run, where the maximum correction (for events from the bottom of the detector) was $\times 1.6$, to $\sim 100 \mu\text{s}$ at the end, with a maximum correction of $\times 2$. A larger recirculation flow-rate would have prevented this decrease, although stability of the liquid surface is better ensured by the lower flow-rate, with the slow degradation adequately corrected in the analysis.

The average electron lifetime measured during the data-taking period was $112 \mu\text{s}$, with the lifetime measured separately for the AmBe calibrations, $403 \pm 91 \mu\text{s}$ on the first day and $212 \pm 26 \mu\text{s}$ on the second, and ^{60}Co calibrations, $235 \pm 32 \mu\text{s}$.

Although the nuclear recoil populations used in the two-point fitting method are subjected to higher electric fields (due to their close proximity to the grid wires) the drift of the resulting electrons will be subjected to the same conditions as bulk events. This assertion was substantiated by agreement between the two-point fitting and α population measurements.

4.4.5 Absolute correction

After correction of all measurable effects on the data, a small residual variability in S2 was still observed for events from the cathode grid over the course of the science data-set. The cause of this variability was uncertain, but could be attributed to small changes in




Figure 4.14: Electron lifetime measurement from the exponential decrease of S2/S1 with drift time, using α -events from ZEPLIN-II. This was obtained prior to purification by recirculation, when a very poor electron lifetime of $\sim 4 \mu\text{s}$ was obtained. The lower extraction grid corresponds to a drift time of $3 \mu\text{s}$, clearly seen in the figure.

the electroluminescence field or suppression of cross-phase emission by surface charging or contamination, as well as imperfect corrections.

The population of nuclear and electron recoil events on the cathode is likely to produce a charge spectrum which remains constant in time, thus providing a normalisation feature in the data. The data were examined below the central target region of interest (drift time $> 67 \mu s$) and a double Gaussian distribution fitted, one to the electron recoil background and the other to the nuclear recoil population, as shown in Figure 4.16. A database was populated by extracting the mean S2 response for the cathode nuclear recoil population in time slices through the data-set, shown in Figure 4.17. The S2 pulses were normalised using the database, providing a final correction which allows a further improvement in energy resolution of the S2 channel.

4.5 Fiducialisation

The self-shielding afforded by large-volume liquid noble gas detectors can provide an extremely low-background WIMP-target region, with edge effects (e.g. incomplete charge extraction, varying electric field) and external backgrounds rejected with much greater efficiency. Fiducialisation of the active volume, selecting a central region, requires re-




Figure 4.15: Electron lifetime monitored throughout the science data-taking period of ZEPLIN-II.




Figure 4.16: S2/S1 spectrum for events from the bottom of the detector ($dtime > 67 \mu s$) for background data (black), AmBe (red) and ^{60}Co (blue) calibrations. Double-Gaussian fitting incorporating cathode recoil events (lower S2/S1) and electron recoil background, with the mean and sigma given for the cathode recoil population from each dataset.




Figure 4.17: S2/S1 response of cathode recoil events monitored through the science data-taking period before application of the absolute correction.

construction of the event location within the detector in x , y and z . Event position reconstruction in three dimensions utilises two different techniques, one measuring the depth in the liquid, z , the other calculating the x - y position.

4.5.1 Vertical position reconstruction

The depth of the interaction in the liquid, z , is measured directly by the time taken for the electron cloud to drift to the liquid surface. The drift time is measured as the time difference between the start of the scintillation and ionisation signals, translated into an absolute depth using the drift velocity in LXe and using calibration points in the target (e.g. the grids). For the purposes of fiducialisation the drift time parameter is sufficient, although the depth is required for calculation of the mass within the fiducial volume.

Figure 4.18 shows drift time distributions from background and ^{57}Co calibration data. The drift time distribution observed during ^{57}Co calibrations clearly confirms that the majority of the events originate from the bottom of the xenon volume. Low-energy γ -rays emitted by ^{57}Co are significantly blocked by the full thickness of the copper target vessel. Consequently, to allow the 122 and 136 keV γ -rays to reach the liquid target, small recesses were machined into the base-plate of the vessel. Nevertheless, these low-energy γ -rays are still predominantly absorbed within the bottom ~ 1 cm of xenon in the target, producing the observed depth spectrum. Conversely, the background data exhibits a much more uniform distribution with excess background populations from the cathode grid, lower extraction grid and PMTs. By placing a cut on the drift time, such excess background populations can be rejected.

The calibration of drift time to absolute depth requires consideration of the different fields present in the active volume, causing changes in the drift velocity. The cathode grid appears in the science data at drift time of $\sim 73 \mu\text{s}$, with the bottom extraction grid located at $\sim 3 \mu\text{s}$, dictating a full drift depth of $\sim 70 \mu\text{s}$. The distance between the two grids is 140 mm, yielding a drift velocity of $2.0 \text{ mm}/\mu\text{s}$ (in agreement with previously published measurements at 1 kV/cm [148]), allowing for calculation of the fiducial z region selected. The main limitation on the accuracy of the position reconstruction in the z direction is the diffusion of the electron cloud during its drift to the surface,




Figure 4.18: Distributions of drift time for background data (left) and ⁵⁷Co data (right) from ZEPLIN-II.

implying that the position resolution degrades with increasing depth of the interaction, although this effect is relatively minor.

4.5.2 Position reconstruction in the horizontal plane

To locate an event in x and y , the hit pattern of the signal across the PMT array is utilised. In principle, a PMT will observe a fraction of the total light emitted based upon its proximity to the event location in x and y . This assumption is utilised in the application of a centroid position reconstruction algorithm, providing fast processing speeds and a robust approach.

Calculation of the x and y location combines the position of each PMT in the array, x_i and y_i , weighted by the signal observed in the individual PMT, $S(i)$:

$$x = \frac{\sum_{i=1}^7 w_i x_i S(i)}{\sum_{i=1}^7 S(i)} \quad y = \frac{\sum_{i=1}^7 w_i y_i S(i)}{\sum_{i=1}^7 S(i)} \quad (4.2)$$

where w_i are additional weighting factors applied to account for differences in individual PMT response. These factors correct for distortion of the position response distribution across the array, yielding a reconstructed distribution which represents the

hexagonal shape of the PMT array. The main limitations on position reconstruction in ZEPLIN-II are the low number and large size of the PMTs, limiting the precision of the reconstructed $x-y$ position. Either S1 or S2 can be utilised to reconstruct the x and y position of an event, but far greater power is held in the S2 signal due to the larger number of photons produced and closer proximity of the signal to the PMT array. In addition, the high light collection uniformity in the liquid phase makes reconstruction of S1 position even harder.

Due to the nature of the centroid reconstruction algorithm, where events cannot be allocated positions outside the outer PMTs, the coordinates returned have arbitrary units. Consequently, for simplicity, it was decided to apply arbitrary co-ordinate values for the PMT locations, summarised in Table 4.1, along with the PMT weighting factors used in the final analysis. These arbitrary values can then be calibrated later.

Table 4.1: ZEPLIN-II centroid reconstruction: PMT positions, x_i and y_i , and weighting factors, w_i .

PMT	x_i	y_i	w_i
1	0.0	0.0	3.00
2	2.0	0.0	1.05
3	1.0	-1.7	1.20
4	-1.0	-1.7	0.80
5	-2.0	0.0	0.83
6	-1.0	1.7	0.76
7	1.0	1.7	0.79

From the x and y positions of an event, the polar coordinates can be extracted, where an approximate correction can be applied to correct for an aberration created by the hexagonal PMT layout, through adjustment of radial coordinate, r by:

$$r_{\text{cor}} = (1.0 - 0.1 \cos(6\theta)) \times r \quad (4.3)$$

For calculation of the fiducial mass, a calibration of the radial position parameter, r , is required. As with the calibration of depth, reference points in the detector are required; in this case, these are provided by the recesses in the base-plate. During calibration of the instrument with a ^{57}Co source, placed below the detector, the low-energy γ -rays are able to penetrate into the xenon through the recesses in the base-plate,

forming two concentric rings, at radii of 7.5 cm and 15 cm, providing useful reference points for calibration.

Figure 4.19 shows the reconstructed positions of events during a ^{57}Co calibration run. The positions of the recesses are clear from the data and agree well with the actual locations. The position of the rings allows r to be calibrated, resulting in ~ 200 mm per arbitrary unit. The radial position resolution can be estimated from the width of the peaks in the radial distribution, yielding an accuracy of $\sigma_r \simeq 1$ cm. This is an upper limit on the underlying resolution as the S2 signals drifting from the detector bottom will have undergone the greatest lateral diffusion.




Figure 4.19: Reconstructed x – y position distribution for ^{57}Co data. The two concentric rings represent the positions of two sets of recesses in the base-plate of the detector.

4.5.3 Fiducial cuts

Selection of a fiducial volume with a reduced background requires placement of cuts on the z and r parameters. The process of defining the fiducial volume is a balance between removing background and retaining a large target volume, carried out through

evaluation of the calibration data and a 10% sample of the science data.

The total active target, defined by the reflective PTFE, forms a truncated cone with volume:

$$V = \frac{1}{3}\pi d (r_1^2 + r_1 r_2 + r_2^2) , \quad (4.4)$$

where $r_1 = 162$ mm is the upper radius, $r_2 = 145$ mm is the lower radius and $d = 140$ mm is the depth, resulting in a total volume of $\simeq 10,400$ cm³. Assuming a xenon density of 2.88 g/cm³ at $P = 1.6$ bar, the total mass is ~ 30 kg, with a small additional volume above the lower extraction grid of ~ 600 cm³ (~ 1.6 kg).

Evaluation of the drift time distribution from the 10% science data sample reveals significant background populations from the lower-extraction and cathode grids. These populations must be removed with high efficiency as they contain a large number of nuclear recoils caused by radon progeny events on the wires. The events originating from the lower extraction grid and the photomultipliers were already rejected by the $\text{dtime} > 4 \mu\text{s}$ cut applied during event selection. A similar cut is applied to remove the events originating from the bottom of the detector, $\text{dtime} < 64 \mu\text{s}$. These cuts leave a fiducial drift time of $60 \mu\text{s}$ ($\simeq 120$ mm), containing a total mass of 24 kg.

Similar considerations were applied to the definition of a radial cut, assessing the background populations evident from the walls of the detector, shown in Figure 4.20. The radial distribution exhibits an approximately constant background within a central region of the detector, accounting for the $1/r$ correction for the volume effect of increasing radius. Removal of the increased background from the detector walls is achieved with a radial cut, $r < 0.467$, which equates to a radius of ~ 80 mm.

The central fiducial region defined by cuts on z and r forms a cylinder of xenon as ionisation drifts vertically in this central region away from the detector walls. Again assuming a density of 2.88 g/cm³, this results in a final fiducial mass of 7.2 kg. The final fiducial volume is small compared with the total mass of xenon, illustrating the importance of low-background internal surfaces and components, clearly demonstrating the advantage of going to larger volume detectors with a higher volume-to-surface ratio.




Figure 4.20: Reconstructed position distributions for the science background data in $x - y$ (left) and radius (right). The radial fiducial cut is shown by the black ring (left) and shaded region (right).

4.6 Energy calibration

The low-energy calibration of $>$ kg-scale xenon targets is a major challenge, given the poor penetration of low-energy γ -rays into a dense medium. Calibration of the energy scale in ZEPLIN-II utilises the dominant 122 keV and 136 keV γ -rays from ^{57}Co (see Table A.1), requiring the target vessel to be thinned in places. Despite this, these γ -rays only penetrate a short distance into the xenon ($\lambda \simeq 0.35$ cm), providing direct calibration of only a small fraction of the volume (see Figure 4.18). Furthermore, we note that 122 keV is still well above the energy range of interest for WIMP searches with xenon (< 30 keV $_{ee}$). Current knowledge of scintillation processes in xenon suggest photon-production linearity down to only about 30 keV $_{ee}$. No conclusive data exists at energies below this, so an assumption of linearity in the scintillation channel is required down to very low energies. A brief description of the radioactive sources used for calibration purposes is given in Appendix A.

The only lower energy volume-distributed feature available for calibration purposes is a 40 keV inelastic scattering feature from ^{129}Xe induced during neutron calibrations (discussed in Sections 4.7.2 and 7.6). Comparison of this feature in simulation and data (calibrated from 122 keV) shows good agreement, lending weight to the assumption of a linear energy response down to at least ~ 40 keV $_{ee}$.

The spectra from both S1 and S2, shown in Figure 4.21, exhibit the combined photopeak from the 122 keV and 136 keV γ -rays, with a population stretching to higher energies, up to ~ 700 keV. The higher energy events are a result of unsubtracted background and of the higher energy ^{57}Co γ -rays, whose very small proportion is enhanced by their greater penetrative power.




Figure 4.21: Spectra of S1 (top) and S2 (bottom) from ^{57}Co calibrations. The photopeak results from the combined response to 122 keV and 136 keV γ -rays emitted by the external calibration source.

In ZEPLIN-II the prompt scintillation signal (S1) alone was used to define the interaction energy. The combined photopeak in the S1 spectrum was fitted by a double Gaussian distribution, incorporating the two significant γ -ray lines, allowing for extraction of a numerical value for the scintillation response to a 122 keV γ -ray. This value provides calibration of the energy scale and allows calculation of the photoelectron yield of the instrument (at the operating field of 1 kV/cm). These ^{57}Co calibrations were repeated daily during the science data-taking period, monitoring the stability of the yield (shown in Figure 4.22). The average measured yield at the operating field of 1 kV/cm was 0.55 phe/keV, with reasonably good stability of the S1 calibration peak throughout




Figure 4.22: Scintillation (S1) response of the instrument monitored through the science data-taking period.

the run requiring no time-dependent energy calibration or correction of the S1 area.

Nuclear recoils produce a different scintillation response to electron recoils (as generated by γ -ray interactions). Consequently, energy calibration of nuclear recoils requires a conversion from the electron-equivalent scale measured. As described in Chapter 3, a nuclear recoil produces less scintillation light than an electron recoil of the same energy. At zero electric field this is characterised by the relative scintillation efficiency, L_{eff} . Measurements of L_{eff} from various groups suggest a constant value of approximately 0.19 in the energy range of interest for ZEPLIN-II. At the ZEPLIN-II operating field (1 kV/cm) additional field induced factors apply to both electron and nuclear recoils due to ionisation track topology. The suppression for nuclear recoil signals has been measured previously for a recoil energy of 55 keV [147], yielding a value of $s_n = 0.93$ at 1 kV/cm. The value for electron recoils can be calculated by comparison of the 1 kV/cm calibrations with a zero-field equivalent. A zero-field calibration spectrum yields a photoelectron yield of 1.1 phe/keV, double that at field. Consequently, the suppression factor for electron recoils is $s_e = 0.5$ at 1 kV/cm. For ZEPLIN-II, at 1 kV/cm, the

conversion of electron-equivalent energy, E_{ee} , to nuclear recoil energy, E_{nr} , is then:

$$E_{nr} = \frac{E_{ee}}{L_{eff}} \times \frac{s_e}{s_n} = \frac{E_{ee}}{0.36}. \quad (4.5)$$

The energy resolution of the detector was measured by comparing the width of the ^{57}Co photopeak with simulations, accounting for all γ -ray energies, giving $\sigma_E = (1.80 \pm 0.04) \times \sqrt{E_{ee}}$, with E_{ee} in keV_{ee} . This takes into account statistical fluctuations in the number of photoelectrons created and other stochastic factors. The measured energy resolution is required for comparison with theoretical WIMP spectra, during calculation of the WIMP-nucleon scattering cross-section (see Chapter 6).

4.7 Discrimination parameter calibrations

Once the relevant output from the instrument has been selected, corrected, calibrated and fiducialised, the dataset is ready for a WIMP search. In two-phase xenon, the discriminatory factor exploited is the ratio of ionisation to scintillation produced by the different interaction types. This S2/S1 ratio is usually plotted against S1-derived energy when assessing the different populations present in the dataset. In this parameter space, a WIMP-search region must be defined and the nature of the expected electron recoil background determined.

Favoured WIMP models result in interactions with baryonic matter taking the form of nuclear recoils from an elastic WIMP-nucleon scatter. The detector response to this signature can be calibrated with a neutron source inducing nuclear recoils in the liquid xenon target. Similarly, the electron recoil background population was calibrated using a ^{60}Co source, emitting 1.17 MeV and 1.33 MeV γ -rays, populating the low-energy spectrum through Compton-scattering in the target. From these calibrations a search region was defined and the expected background (in that region) calculated. The search region is not expected to be event-free, due to partial overlap of the electron recoil population, a potential background of residual neutrons and any additional unknown backgrounds. From simulations based on the detailed detector geometry and component radioactivity, < 0.5 neutron events are expected in the WIMP-search region in this particular data-set. As a result, the neutron background is not added to the overall

total expectation in the region, but should be considered in the case of an apparent unknown background.

Calibrations of the instrument were carried out at a high rate, providing sufficient data to accurately define the populations in the S2/S1 space. The high rate utilised in the ZEPLIN-II calibrations, however, caused a large number of overlapping events in the data, where unrelated S1-only (or sub-threshold) and rare S2-only pulses are identified together. This association of unrelated pulses produces a uniform scatter of events across the discrimination plots. The level of accidental coincidences can be assessed by investigating events appearing beyond the maximum physical drift time. Figure 4.23 shows the drift time of events from the ^{60}Co calibration, with a small constant population appearing beyond the maximum physical drift time, $\sim 73 \mu\text{s}$. Assuming this level to be constant, it is clear that the population will not significantly affect the fitting of the main distributions.




Figure 4.23: Drift time distribution from ^{60}Co calibration data, showing a population appearing beyond the maximum physical drift time, resulting from accidental coincidences causing event confusion.

4.7.1 Electron recoil populations: ^{60}Co γ -rays

To attribute any events observed in a science data-set to electron recoils, they must match the response of the instrument to a known electron recoil population. The response of the instrument to the presence of low-energy electron recoils is defined through Compton-scattering of high energy γ -rays from an external calibration source. For this purpose, a $10\ \mu\text{Ci}$ ^{60}Co source was placed within the neutron shielding, emitting 1.17 MeV and 1.33 MeV γ -rays (see Appendix A). Data were acquired over a 19.3 hour run, yielding 7795 golden, fiducial events below $40\ \text{keV}_{ee}$. Ideally, calibration of the background populations for rare-event search should have high statistics, allowing precise definition of the overlapping tails of the distributions. However, these high statistics calibrations prove impractical due to the huge data volume required. Figure 4.24 shows the resulting scatter plot with the electron recoil population having $S2/S1$ between 100 and 500.




Figure 4.24: Discrimination parameter scatter plot for the electron recoil calibration dataset acquired using a ^{60}Co γ -ray source.

4.7.2 Nuclear recoil calibration: AmBe neutrons

To populate the nuclear recoil region of the S2/S1 parameter space, mimicking a WIMP-nucleon scattering population, neutrons are produced by (α, n) interactions in a 0.1 GBq AmBe source (see Appendix A). The source is manually inserted within the neutron shielding, ~ 1 m away from the liquid xenon target. Neutrons interact with xenon atoms through elastic scattering, producing nuclear recoils, or inelastic scattering, producing isotope-specific γ -rays. The pure elastic recoils, mimicking a WIMP signal, have the lowest S2/S1, forming a band stretching down to low energy. A second population is also found just above 40 keV, resulting from inelastic scattering from ^{129}Xe , where the nuclear recoil component is accompanied by a γ -ray of a fixed energy, in this case 39.5 keV. Further populations at higher energy are produced by inelastic scattering from ^{131}Xe (~ 80 keV) and additionally the PTFE cone (^{19}F at ~ 110 keV and ~ 197 keV), along with electron recoils from γ -rays emitted by the AmBe source. A discussion of the 40 keV inelastic feature is given in Chapter 7.




Figure 4.25: Discrimination parameter scatter plot for the nuclear recoil calibration dataset acquired using an AmBe neutron source.

The AmBe science calibration was completed over the course of two days, with 6.3 hours total exposure. The scatter plot in Figure 4.25 shows events in the fiducial volume passing all selection cuts up to 40 keV_{ee}. Characterisation of the detector response to nuclear recoils is achieved by slicing in 2 keV_{ee} bins (shown in Figure 4.26). The nuclear recoil population in the individual slices were fitted by log-normal distributions, describing the mean and width of the S2/S1 response. The means of the distributions were parameterised by a power law, $S2/S1 = 345.3E_{ee}^{-0.4673}$, shown in Figures 4.24 and 4.25 by the solid line. Integration of the log-normal distributions provides a calculation of the fraction, A_{nr} , of nuclear recoils below a given value of $y = S2/S1$, the nuclear recoil acceptance:

$$A_{nr}(y) = \int_{-\infty}^y f(y; \mu, \sigma) dy = \int_{-\infty}^y \frac{1}{y\sigma\sqrt{2\pi}} e^{-\frac{(\ln(y)-\mu)^2}{2\sigma^2}} dy = \frac{1}{2} \left(1 + \operatorname{erf} \left[\frac{\ln(y) - \mu}{\sigma\sqrt{2}} \right] \right), \quad (4.6)$$

where μ and σ are the mean and standard deviation of S2/S1.

Definition of the expected signal region of the parameter space provides a basis on which to define a WIMP search box. Comparison of electron and nuclear recoil populations from the calibrations shows the discriminatory power held in the S2/S1 parameter. A constant S2/S1 lower bound for the search box was defined, $y = 40$, excluding the lower parameter space which contains little recoil acceptance. The upper limit can be defined for the search box based on either constant nuclear recoil acceptance or constant background leakage. In the case of ZEPLIN-II, a constant-acceptance upper bound was adopted. Following assessment of the overlap of the populations, the upper bound was set at 50% nuclear recoil acceptance, $A_{nr} = 0.5$, calculated by integration of the sliced distributions above $y = 40$.

The roughly exponentially falling energy spectrum of the expected WIMP-nucleon signal informs the definition of bounds on the energy range of the search box. The lower energy bound should be reduced as far as practicable to include the most significant part of the recoil spectrum. In practice, however, this lower limit is constrained by the energy threshold and efficiencies of the detector, which remove the signal power at very low energies. This reduced power, combined with the increased probability of observing spurious events at very low energies, led to a lower limit on the energy of

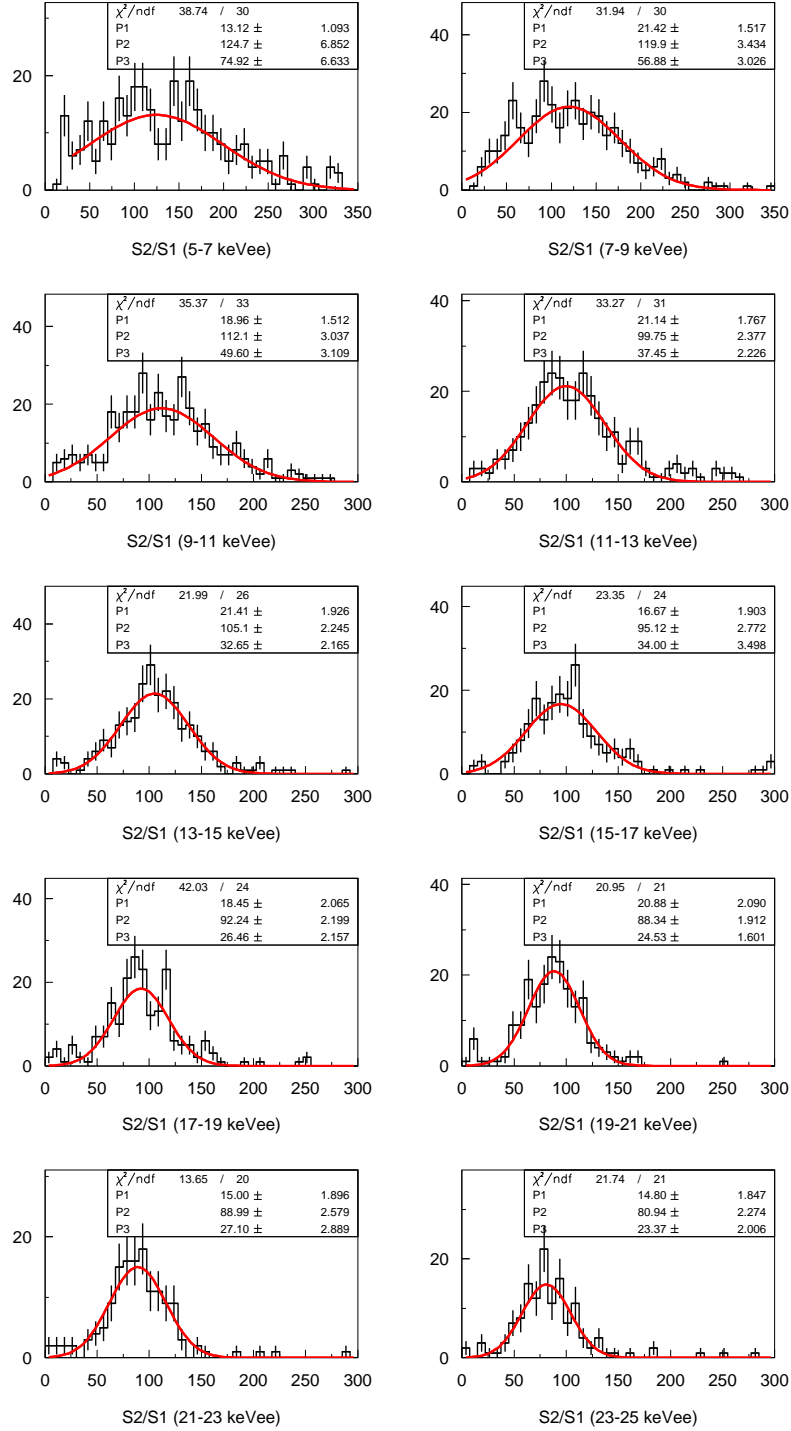


Figure 4.26: 2 keV_{ee} slices of the ZEPLIN-II AmBe neutron calibration, showing fits of the nuclear recoil population with log-normal distributions.

the box of 5 keV_{ee}. The upper limit is likewise determined by the combination of the expected signal spectrum and detector efficiencies, which begin to decrease at energies above 30 keV_{ee} due to a hardware cut. An upper energy limit of 20 keV_{ee} was chosen, excluding only $\lesssim 2\%$ of the signal above 5 keV_{ee}.

The WIMP search box containing 50% of nuclear recoil acceptance above S2/S1 = 40, between 5 keV_{ee} and 20 keV_{ee} is overlaid onto the calibration data (Figures 4.24 and 4.25), showing only a small overlap with the electron recoil population. The box, defined using calibration data, was checked against the 10% background data sample and fixed before analysis of the full data-set.

4.7.3 Efficiencies in the acceptance region

For the calculation of WIMP-nucleon scattering cross-sections we must reconstruct the ‘true’ event rate from the one which is experimentally observed. The observed rate is affected by detector characteristics such as efficiencies, thresholds and energy resolution. All processes which reduce the probability of detecting a nuclear recoil event must be incorporated, yielding an overall nuclear recoil detection efficiency.

The efficiencies of the selection cuts applied to the data have been individually measured (or calculated) and are summarised in Table 4.2. Hardware effects also result in efficiency loss, and must also be included in the overall efficiency. In the latter category, the most significant is a saturation cut applied in hardware during background acquisition; this is energy-dependent, as shown in Figure 4.27, but has unity efficiency below 30 keV_{ee}.

Events in the target are rejected if a coincident signal is observed in the liquid scintillator veto, thus removing $\sim 50\%$ of low-energy neutron scatters in the xenon. A small probability exists that a coincidence between events in the target and veto is accidental, leading to an efficiency of 99.2% during background acquisition (91.5% during AmBe calibrations and 61.8% during ^{60}Co calibrations).

During data acquisition, communication between the ACQIRIS digitisers and the DAQ computer constituted a small bottleneck introducing a dead-time effect. It was calculated that the dead-time fraction for background acquisition was $\sim 10\%$ (an efficiency of 90%), with larger dead-time fractions for calibration data: 37% for the AmBe




Figure 4.27: Efficiency plot for the hardware upper-level discriminator (‘saturation cut’) from comparison with data where no upper-level discriminator was applied.

calibration and 22% for ^{60}Co calibration.

All these efficiencies were combined to yield the final energy-dependent efficiency for nuclear recoil detection in the science data-set of ZEPLIN-II. The solid histogram in Figure 4.28 shows the overall efficiency calculated from the individually measured efficiencies, increasing with energy to reach a maximum efficiency of $\sim 78\%$ around 13 keV_{ee}. The saturation cut reduces the efficiency above $\simeq 30$ keV. The total detection efficiency in the WIMP search box ranges from 34% at 5 keV_{ee} to 78% at 20 keV_{ee}. The detection efficiency can also be assessed through comparison of calibration data with Monte Carlo simulations. Simulation of the AmBe calibrations yielded a spectral shape in good agreement with that of the calibration. By normalising the simulated spectrum to match the calibration spectrum in the energy range of constant efficiency, the shape of the efficiency curve can be calculated. Figure 4.28 also shows the efficiency curves calculated from such a comparison with simulated AmBe data, as well as ^{60}Co data. All three measurements of efficiency show good agreement, which is very reassuring.

Table 4.2: Summary of nuclear recoil detection efficiencies in ZEPLIN-II science data, including software event selection cuts and hardware effects.

First pass event selection cuts	
S2 τ cut	$\simeq 100\%$
S2 position cut	100%
1 nVs minimum S2 area	$f(E)$ 90.3% between 5-10 keV _{ee}
S1 τ cut	$\simeq 100\%$
S1 position cut	100%
3-fold S1 requirement	$f(E)$
Golden rule	100%
Second pass event selection cuts	
Tight S2 τ cut	90.2%
Tight S1 τ cut	$\simeq 100\%$
S1 Δt_{peak} cut	$\simeq 100\%$
Additional selection cuts	
Trigger-point cut	$\simeq 100\%$
Cathode event check	99.7%
Hardware cuts	
Saturation cut	$\simeq 100\%$ below 30 keV _{ee}
Veto accidentals	99.2%
Dead time	90% for background data acquisition




Figure 4.28: Plot of the total nuclear recoil detection efficiency in ZEPLIN-II. The solid histogram represents the combined efficiency calculated from the different cuts (summarised in Table 4.2), with the solid circles and empty squares representing an efficiency calculated from comparison of the AmBe and ^{60}Co calibrations with simulations.

4.7.4 Expectations

The expected number of background events is important when statistically analysing the observation in the WIMP search box for the presence of a signal. Comparison between events observed in the WIMP search box in the science data-set and the electron recoil population defined by the calibrations, requires a calculation of an expectation for the observed background. Integration of the electron recoil calibration within the box yields the fraction of electron recoils leaking into the search region. Due to the energy-dependent S2/S1 behaviour of the electron recoil population (and box upper bound), the energy range is sliced into two bins, 5-10 keV_{ee} and 10-20 keV_{ee} .

The S2/S1 values plotted in Figure 4.29 are offset by the energy-dependent 50% nuclear recoil acceptance value, $k(E)$. The electron recoil S2/S1 distributions within

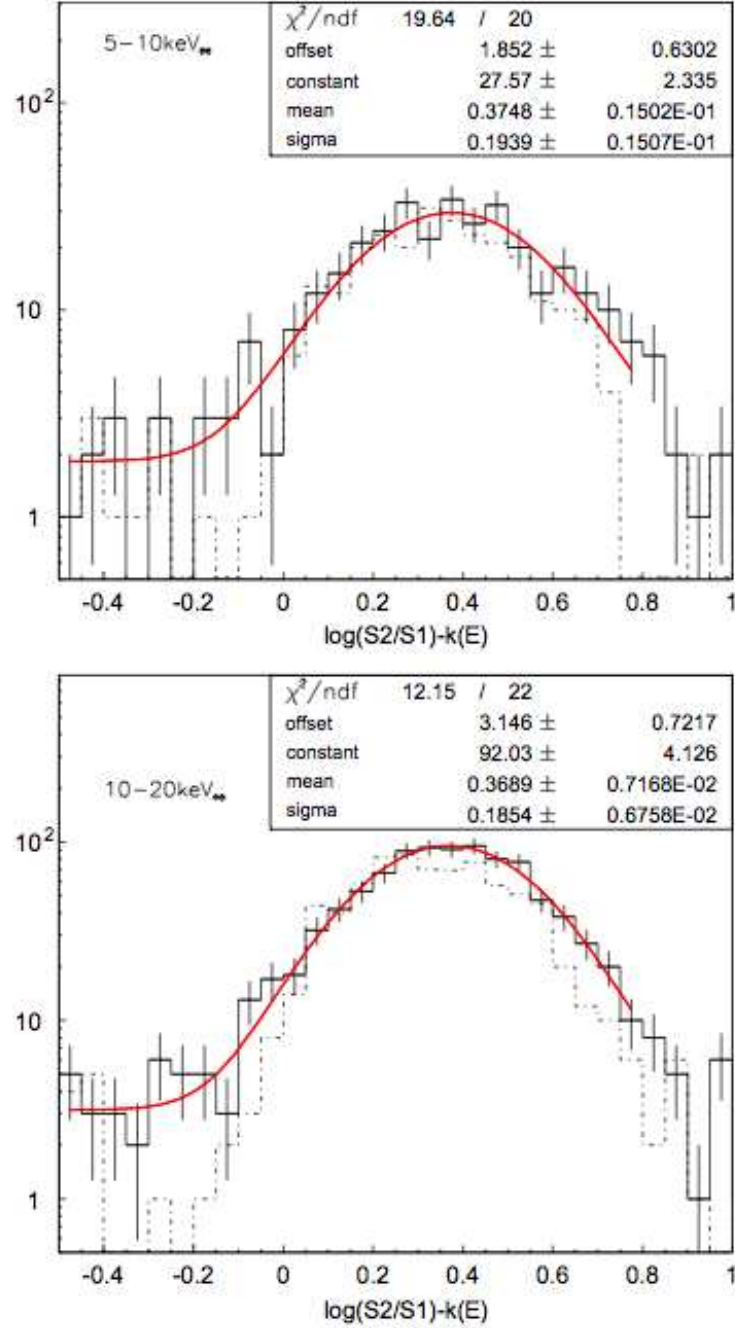


Figure 4.29: Slices of the electron recoil population from the ^{60}Co calibration of ZEPLIN-II. The discrimination parameter plotted as $\log(S2/S1) - k(E)$, where $k(E)$ is the energy-dependent value of the 50% nuclear recoil acceptance, i.e. the upper boundary of the acceptance box.

the slices were fitted by log-normal distributions with a small constant component accounting for the uniform distribution of random coincidence events. Integration of the log-normal distributions below 0 (i.e. inside the search box) yields the number of leakage events in the calibration, a value which requires further normalisation based on the relative exposure time of the calibration and science data-sets. Following normalisation we calculate an expectation of 4.2 ± 2.4 electron recoil events in the box from 5-10 keV_{ee} and 11.9 ± 2.7 events in the 10-20 keV_{ee} range.

4.8 Science data

The science data-set was acquired over a period of 57 days, amassing a total of 1767 kg·days of data from the 31 kg target. A live time of about 77% was achieved, after accounting for daily calibration periods (⁵⁷Co and α -events), brief offline periods of maintenance and removal of the 10% data sample. Further cuts were applied to ensure acceptable uniformity of detector response, removing days containing electric field instabilities, resulting in a final exposure of 31.2 days of science data (967 kg·days for the 31 kg active volume). Fiducial cuts reduce the fiducial mass to 7.2 kg of liquid xenon, giving a final exposure of 225 kg·days. The exposure calculation is detailed in Table 4.3.

To avoid bias when analysing the science dataset for the presence of a signal, a blind analysis procedure was intended, with the event selection, quality and fiducial cuts defined with the commissioning and calibration data, and cross-checked with a

Table 4.3: Summary of ZEPLIN-II exposure for science background data-set.

	Exposure	Cuts applied (either operational or fiducial)
Calender runtime (31 kg target mass)	57 days (1767 kg·days)	Complete length of science run in final background configuration
Science data run (31 kg target mass)	44.2 days (1370 kg·days)	Background exposure, excluding periods of calibrations and maintenance
Stable operation (31 kg target mass)	31.2 days (967 kg·days)	Days experiencing E-field instabilities
Drift time fiducial volume (26 kg target mass)	31.2 days (811 kg·days)	Fiducial cut in drift time, z , applied
Radial fiducial volume (7.2 kg target mass)	31.2 days (225 kg·days)	Fiducial cut on radial position, r , applied

10% sample of the science data. For this purpose, every tenth data file was selected and analysed “unblind”. The WIMP-search box was defined from comparison between electron and nuclear recoil calibration data-sets with no knowledge of the complete science data-set, again cross-checked with the 10% sample. Subsequently, after definition of all cuts, corrections and the WIMP search region, the remaining 90% of background data was processed.

Following processing of the “blind” science data-set (selecting golden events, applying corrections for PMT gain, finite electron lifetime, pressure and absolute variations), a discrimination scatter plot was produced from events within the fiducial volume. Dozens of spurious events were observed and it was clear that a loop-hole in the event selection algorithms was allowing anomalous events to pass the cuts, a result of the code not implementing the pre-defined selection philosophies correctly. This defect could not be identified in the 10% data sample due to lack of statistics. As a result, the code had to be corrected after all the data had been revealed, resulting in an “unblind” analysis of the final science data-set. However, the definition of the acceptance region and cut definitions remained unchanged.

The final discrimination scatter plot for the science background data is shown in Figure 4.30, with and without events vetoed by coincident signals from the liquid scintillator. In total 29 unvetoed events were observed in the acceptance region (14 from 5–10 keV_{ee} and 15 from 10–20 keV_{ee}). The electron recoil background population was observed as expected from ⁶⁰Co calibrations; however, an additional unexpected population of events was observed with low S2/S1, forming a band with an approximately constant S2 signal. This signal was clearly incompatible with nuclear recoil events from the fiducial volume, and it was unlikely to arise from WIMP interactions or a neutron background.

Further study of these events clearly placed their origin at the detector walls, leaking into the radial acceptance region due to the poor radial position resolution for small S2 signals. Figure 4.31 shows the radial distributions of these small S2 signals in the two energy ranges (5–10 keV_{ee} and 10–20 keV_{ee}) showing a wide Gaussian distribution centred at ~ 0.73 (in arbitrary units), the position of the PTFE wall. The width of the reconstructed radial distribution is affected by the size of the S2 signals, as a result of the centroid position reconstruction algorithm. Through fitting of the wall-event

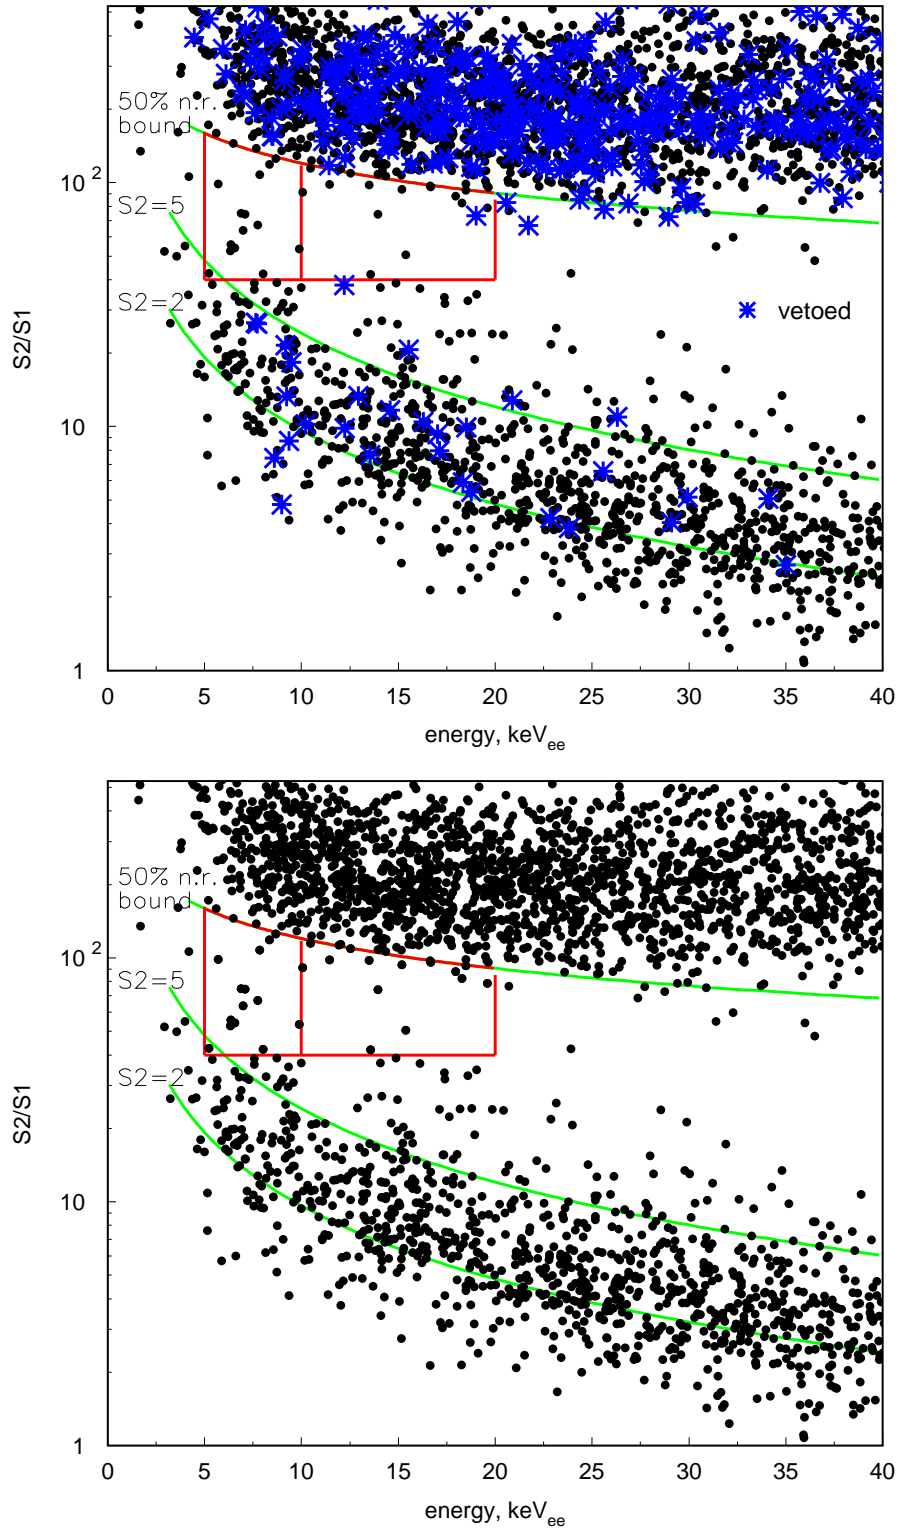


Figure 4.30: Discrimination parameter scatter plot of the final science dataset from ZEPLIN-II, with (top) and without (bottom) vetoed events. In the top plot vetoed events are symbolised by a star. The acceptance region and 50% nuclear recoil acceptance line are shown, together with lines of constant $S2$ ($S2 = 2$ and $S2 = 5$).

population, the radial resolution can be measured as a function of S2 area (shown in Figure 4.32), with the resolution clearly degrading for very small S2 pulses.

These small S2 events are attributed to ^{222}Rn -progeny interactions, where daughter species produced by decay of ^{222}Rn (emanated from the SAES getters) migrate to the PTFE walls, cathode and extraction grids, under the applied electric field, becoming implanted on the surface. Often, these species decay with only the nuclear recoil component escaping from the surface. Similar nuclear recoil events were observed from the

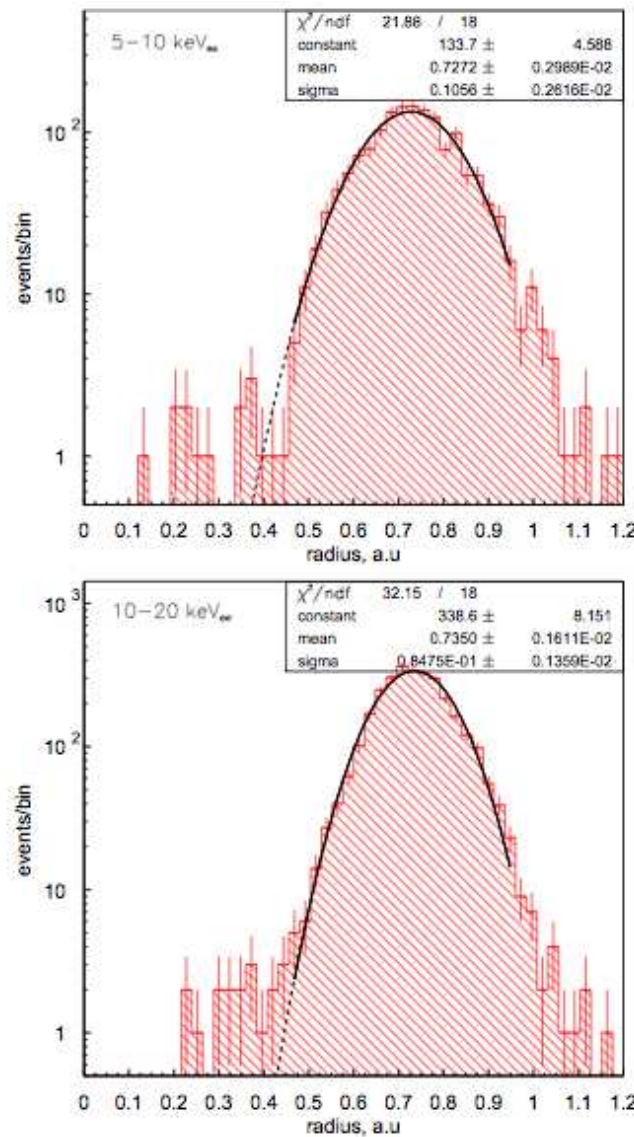


Figure 4.31: Reconstructed radial distributions of small S2 events from the science dataset. Gaussian distributions are fitted outside the fiducial radius (solid line) and extrapolated into the acceptance region (dotted line) to calculate an expectation.




Figure 4.32: Plot showing the worsening of radial position resolution with decreasing S2 signal size. Radon progeny (and other background) interacting near the walls of the detector was used.

cathode and lower-extraction grids, where the proximity to the grid wires provides a higher field to extract the charge. Close to the PTFE walls extraction of charge is suppressed either at the interaction site, whilst drifting through the volume (due to close proximity to the wall) or during electroluminescence (in the peripheral region of the gas layer), producing a smaller S2 pulse.

To account for the additional background population in the search for a WIMP signal, the number of events expected in the search box from this new background must be calculated. This calculation should be done excluding information from within the search box. The well-defined radial distribution of the small S2 events allows for characterisation of the population with a Gaussian distribution fitted beyond the radial fiducial cut (i.e. outside the acceptance region). The distribution is extrapolated and integrated

within the fiducial radius, yielding an expectation on the number of ^{222}Rn -progeny events within the box. This technique was used for the two energy ranges defined previously (shown in Figure 4.31), predicting 10.2 ± 2.2 events in the 5–10 keV $_{ee}$ bin and 2.3 ± 0.5 events in 10–20 keV $_{ee}$.

The total expectation in the WIMP-search box (detailed in Table 4.4) were calculated by summing the γ -ray expectation derived from the ^{60}Co calibrations and the ^{222}Rn -progeny expectation giving a total expectation of 28.6 ± 4.3 background events in the box, compared with the 29 events observed in the data. A statistical analysis on the presence of a signal can be carried out from the observed and expected number of events, with a limit on the signal rate calculated from the experimental exposure (in kg·days). The limit on rate is subsequently compared with the theoretical WIMP-nucleon scattering spectrum calculated for a xenon target with the calculated detector characteristics (energy resolution and energy-dependent efficiencies). By applying a Feldman-Cousins technique (as described in Chapter 6) using the total expectation and observation, gives a 90% upper limit of 10.4 events, which translates to 0.092 evts/kg/day. The process of calculating a limit on the WIMP-nucleon spin-independent scattering cross-section is detailed in Chapter 6.

Table 4.4: Summary of background expectation in the ZEPLIN-II WIMP search box.

Energy range (keV $_{ee}$)	Observed events	γ -ray (^{60}Co data) (\mathbb{A})	γ -ray (bkgd data)	^{222}Rn -progeny (\mathbb{B})	Total ($\mathbb{A} + \mathbb{B}$)
5-10	14	4.2 ± 2.4	5.6 ± 4.6	10.2 ± 2.2	14.4 ± 3.3
10-20	15	11.9 ± 2.7	13.0 ± 6.0	2.3 ± 0.5	14.2 ± 2.7
				Total	28.6 ± 4.3

Chapter 5

ZEPLIN-III: First Science Run

Following the completion of the ZEPLIN-II experiment, ZEPLIN-III was commissioned at the Boulby underground laboratory. ZEPLIN-III utilises the same two-phase detection principle as ZEPLIN-II, sensing both prompt scintillation and drifted ionisation with an array of photomultiplier tubes. However, it operates with a stronger drift field and has the PMTs located in the liquid, viewing the target volume from below (as shown in Figure 3.19). The work described here relates to the analysis of the first science run.

Initial construction was completed at Imperial College at the beginning of 2006, where surface testing was also carried out [170]. The detector was subsequently transported to Boulby (December 2006) and reassembled underground. Commissioning began in mid-2007, when detector characteristics, both operational and technical, were optimised. A first science run of the detector lasted from February until June 2008, including 84 days of shielded science data, daily ^{57}Co calibrations, AmBe neutron calibrations and ^{137}Cs Compton calibrations, along with other datasets of interest. A description of the instrument and related systems is given in Section 3.5, with a detailed review of the hardware provided in [102]. Detailed simulations of the instrument's performance were carried out during the design phase [150]. The results of the first science run have been published [103, 171], and other papers are in preparation describing other significant results.

There were many similarities between the ZEPLIN-II and -III detectors which aided the commissioning and data analysis of ZEPLIN-III. However, there were also important differences which introduced additional challenges, including light collection

non-uniformity, detector dead-regions, the increased number of PMTs and dual-range acquisition. The analysis path for ZEPLIN-III follows along similar lines to that of ZEPLIN-II, but with added complexity in certain areas.

I have personally played a role in many aspects of the analysis: organising the data pipeline, storage and processing, helping to debug and optimise the data reduction software (ZE3RA), optimising the golden code, applying basic position reconstruction, creating automated monitoring routines, studying dead region events, analysis of the calibrations and final science data. Brief descriptions of other pieces of work (carried out by others within the collaboration) have been included to illustrate the complete analysis process leading to the final dark matter result.

5.1 Data acquisition and pipeline

As the complexity of detectors increases with more information being extracted, the data-rate can become a challenge. The increase in the number of channels recorded from ZEPLIN-II (7 PMTs, single-range) to ZEPLIN-III (31 PMTs, dual-range), posed new challenges in acquisition, transfer, handling and storage. Data acquisition of the 62 channels was achieved with 8-bit ACQIRIS digitisers (ACQIRIS DC265), sampling at 500 MS/s (2 ns samples); example waveforms are shown in Figure 5.1. The waveforms recorded are centred on the trigger point, extending 18 μ s before and after, recording events triggered by S1 or S2. The PMTs are powered by a single HV supply, with individual attenuators (Phillips Scientific 804) applying a first-order normalisation of the PMT responses.

The normalised signals from the 31 PMTs are passively split, with one channel recorded unamplified, the low sensitivity channel (LS), and the other passing through a fast 10x amplifier (Phillips Scientific 770) before acquisition, the high sensitivity channel (HS). The amplified HS signals are again split, feeding the acquisition and trigger electronics. The dual range, with a factor of 10 between high and low sensitivity, in addition to different full scale settings available for each range, allows the recording of large S2 signals whilst retaining good sensitivity for the small S1 pulses, providing a larger dynamic range.

The trigger electronics are shown in Figure 5.2. The 31 high sensitivity signals feed




Figure 5.1: Display of all 31 high sensitivity waveforms from ZEPLIN-III.

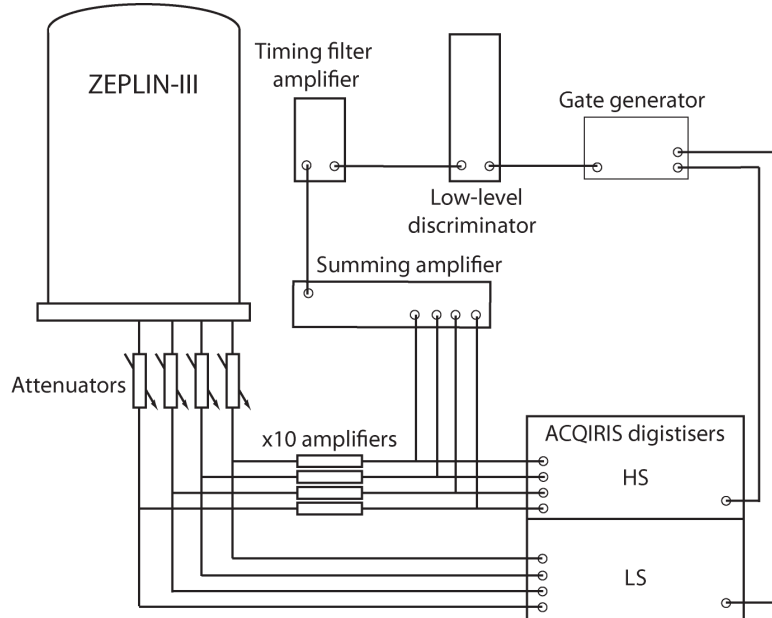


Figure 5.2: A schematic detailing the trigger and acquisition circuitry for ZEPLIN-III.

a summing amplifier before integration (500 ns) and differentiation (500 ns) by a timing filter amplifier (Canberra 2111). This signal passes to a fast discriminator (NE 4684), set to -80 mV, which outputs to an external gate generator triggering the digitisers. This unit is inhibited until the DAQ computer confirms that the crate is ready for trigger and, in addition, for 1 ms after each event. Detailed studies of the trigger have shown that the trigger threshold is equivalent to ~ 11 ionisation electrons emitted from the liquid surface, or ~ 0.2 keV_{ee} for nuclear recoils, and is triggered exclusively by S2 pulses for recoils up to 40 keV_{ee} (and by S1 pulses above that energy). The threshold was set at this level to avoid excess triggering from single electron emission or low energy events with S1 below software threshold.

Waveforms are recorded using specifically-designed LINUX-based software. This software reads out the digitisers into a FIFO-type memory buffer, and writes out events to disk using two separate processors, reducing the overall dead time. To reduce the data-rate recorded to disk, a software saturation cut can be applied in the DAQ, rejecting events which saturate the low sensitivity channel. This removes only high energy events, above the range of interest, leaving an efficiency of 100% below ~ 70 keV_{ee} (shown in Figure 5.3).

After compression, data were transferred to two identical 100 GB LTO1 data tapes. One was transported to our central repository at RAL, whilst the other tape was stored

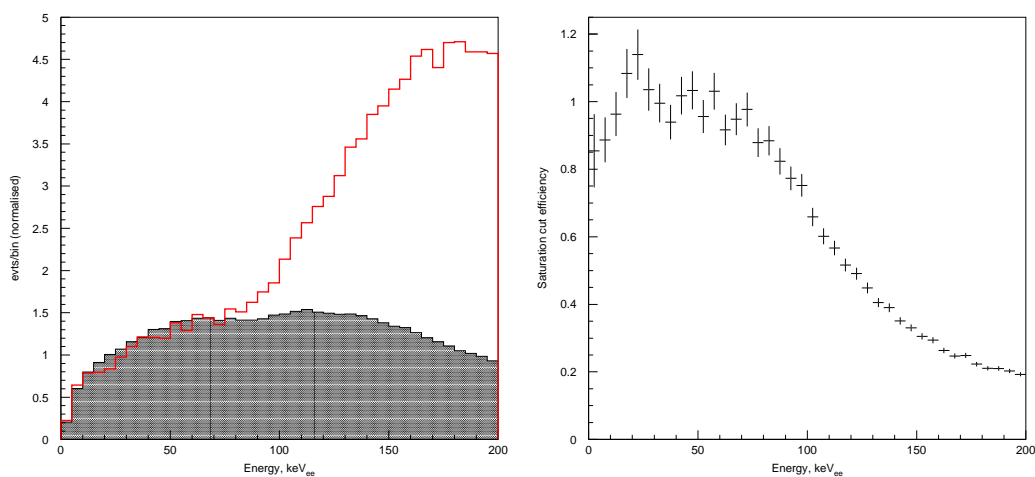


Figure 5.3: ^{137}Cs energy spectrum acquired with no saturation cut (red line) and applying the standard saturation cut (shaded grey) (left) and the derived saturation cut efficiency curve (right).

at Boulby for backup purposes. The first step of the analysis, data reduction (see Section 5.2), is completed utilising an Xgrid system [172], with 28 processors, harnessing a combined processing power of 72 GHz. The data management and processing have been an important part of the project, in fact more of a challenge than originally envisaged. The practicalities of managing large volumes of data must be carefully considered for future large scale systems. The methods of waveform digitisation, data acquisition and data compression can all assist in this, along with adoption of large volume handling facilities afforded by larger projects, such as the LHC Grid.

During the first science run of ZEPLIN-III and associated commissioning runs, about 60 TB of uncompressed raw data were acquired, with a lossless compression factor of ~ 4 achieved using gzip, leaving a total of 15.5 TB of compressed data. Additional compression algorithms are being developed to reduce the volume further.

5.2 Pulse identification process

The acquired raw data comprises 62 individual waveforms, which must be processed to characterise any pulses, in both the high and low sensitivity channels. For ZEPLIN-III, a piece of software, ZE3RA (ZEplin 3 Reduction & Analysis), was specially developed for this purpose. ZE3RA parameterises waveforms without ascribing any physics meaning to pulses, which is important to avoid biasing the data analysis at this early stage. ZE3RA also provides an event display (Figure 5.4), which is vital to enable a progressive understanding of the recorded events. A brief overview is given of the characteristics of ZE3RA which influence the quality of the reduction and the parameters, output into *hbook* ntuples, which are provided for further data analysis.

Before accurate pulse-finding and parameterisation can take place the baseline of each waveform must be characterised. A $2\ \mu\text{s}$ section at the start of the waveform (1000 samples) is analysed for this purpose allowing the baseline level and noise (RMS) to be determined. A threshold of $3 \times \text{RMS}$ is set independently in each channel.

Following characterisation of the baseline and definition of a threshold, pulse finding can proceed, with a search for pulses above threshold in each waveform. Pulses identified in this way are combined to create two intelligent sum channels (HS and LS), providing much clearer sum waveforms for pulse identification and parameterisation. This is

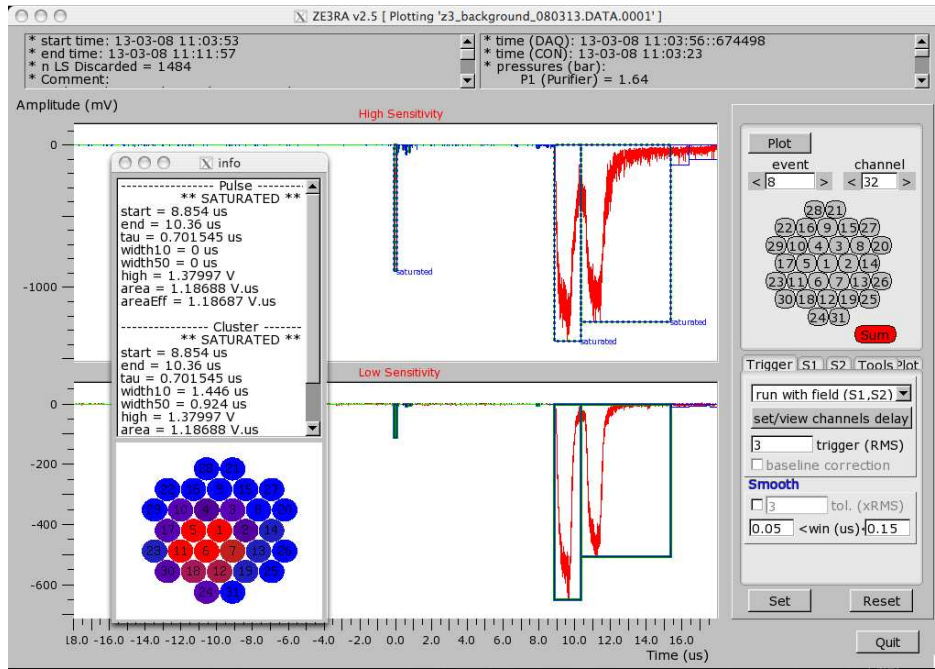


Figure 5.4: Event visualisation using ZE3RA software, showing the sum channel for a waveform in the high and low sensitivity. The inlaid window provides pulse information and basic position information based on pulse size between different PMTs.

especially important for small signals which could be diluted by noise in PMTs seeing no signal. An algorithm is applied to correlate pulses between high and low sensitivity channels, using a statistically-motivated timing/shape coincidence analysis.

For each event the pulses identified are parameterised, ordered by decreasing area (in the HS sum channel), with the 10 largest stored in *hbook* ntuples for further analysis. The parameters recorded are listed below in functional blocks. Pulses are identified by two indices ($id = 1...10$, $channel = 1...64$) and waveforms by a single index ($channel = 1...64$). The number of channels (64) includes the 2×31 signals plus two sum waveforms (HS and LS).

5.3 Event selection philosophy

As with ZEPLIN-II, a golden code was applied to the reduced data, selecting events with one good S1 and one good S2, simplifying further analysis and providing a coherent event selection philosophy for all analyses contributing to the final science result. The event selection philosophy and cuts are described here.

Table 5.1: Parameters recorded in *hbook* ntuples by ZE3RA.

Parameter block	Parameters
Event	Event number and time
Detector slow control	Trigger rate, detector temperatures and pressures, liquid level
High voltage	Electrode and PMT voltages and currents
External environment	Temperatures and pressures of the lab
Waveform	Baseline and RMS noise values, number of pulses stored
Pulse	Pulse saturation flag Start time of pulse Pulse amplitude Pulse width at 10% and 50% of maximum Pulse area integrated above threshold and total integrated pulse area Mean charge arrival time

For a selected dataset, the ntuples produced by ZE3RA are read into the code, with a series of physically motivated event selection cuts applied (listed below), before the surviving events are written to “golden” ntuples.

- First pass cuts: simple S1 and S2 identification cuts \rightarrow golden rule enforced.
- Voltage cuts: check that voltages (grids and PMTs) are within acceptable ranges.
- Waveform cuts: stability check on baseline, noise and other waveform properties.
 - New “golden” variables are calculated for surviving events.
- Second pass cuts: final test of events using newly calculated parameters

Events passing all the cuts are written to “golden” ntuples which can be easily utilised for further analysis. The golden code implements counters to track the number of events which pass or fail each section, aiding assessment of the efficiency of the analysis for nuclear recoil detection. This information, along with the number of files and live time of the data set, are reported at the end of the process.

5.3.1 First pass cuts

Characterisation of the pulses parameterised by ZE3RA as either an S1, S2 or neither, is achieved through the application of a series of cuts based upon the physics processes and detector characteristics. As with ZEPLIN-II, the relatively straightforward S2 search is carried out first, followed by a “look-back” search for the much smaller S1 pulses.

Pulses are first compared with the expected properties of an S2 pulse, of which the most basic characteristic is the large width, requiring application of a cut on the mean charge arrival time, τ ($\tau > 150$ ns). As with ZEPLIN-II, a minimum area ($s2area > 5$ V.ns) cut is applied to ignore single electron emission signals during S2 identification. Likewise, two additional pulse checks are required to account for detector artifacts in the data: amplifier overshoot and optical feedback.

Optical feedback induces emission of electrons from the cathode grid, which drift through the liquid producing a second, small electroluminescence pulse. These signals are not valid S2 pulses, as they are not caused by a direct particle interaction within the fiducial volume, and should therefore not be counted as such. To account for such events, signals near the maximum drift time (> 13 μ s) from the main S2 are ignored during pulse counting.

Amplifier overshoot occurs after saturation due to large signals and causes a low-amplitude, wide signal directly after a real S2, as shown in Figure 5.5. We may discount these signals due to their non-proportionality between high and low sensitivity channels, with no overshoot observed in the low-sensitivity channel.

A final cut is applied to remove very small S2-like pulses produced in the tail of large S2 signals. The tails of large S2 show multiple feedback and single electron signals (see Figure 5.6), which can be clustered together by ZE3RA resulting in wide S2-like pulses. Such events are ignored if their area is less than 5% that of the main S2, as real double scatters of this magnitude will not substantially affect the discrimination parameter.

Following identification of S2 pulses, remaining pulses are examined as potential S1 signals. Prompt scintillation pulses are fast, requiring a cut on the τ parameter, which was chosen to complement the τ cut on S2 ($\tau < 150$ ns). Measurement of small signals requires good discrimination against waveform noise, with a 3-fold requirement enforced: a pulse must exceed a threshold of $\sim 1/3$ of one photoelectron in at least 3 PMTs. This requirement leads to an S1 software threshold of 1.7 keV_{ee}.

Most S2 pulses have a noisy tail which provides multiple S1-like candidates (as seen in Figure 5.6). Consequently, we invoked the requirement that S1 occurs in the time preceding S2. Similarly, S1 signals are often followed by an after-pulsing signature, a signal-induced artifact internal to the PMTs which could potentially be incorrectly




Figure 5.5: Example of amplifier overshoot. Large overshoot visible in the high sensitivity, but not the low sensitivity, following a large S2 pulse.




Figure 5.6: Example of the noisy tail on a large S2 pulse being clustered into multiple S2-like pulses.

identified as an additional S1 pulse. PMT afterpulsing is likely to result from ionisation of residual gas within the photomultiplier, producing additional pulses a few hundred ns after the original scintillation signal. As a result, after an S1 pulse has been identified, all subsequent S1 candidates must be evaluated as potential after-pulsing. This is done by examining the PMTs observing each S1, allowing through fast pulses observed at a particular time delay from an S1 signature in that channel, but taking into account that no afterpulses should exist in channels where S1 was not observed.

These characterisation checks are applied to all pulses parameterised by ZE3RA in a single event, counting the number of S1 and S2 pulses present. The golden rule is then applied, allowing only events with one S1 and one S2 to pass. This embodies the fact that WIMPs are not likely to multiple-scatter, and also rejects events where signal confusion might compromise their interpretation.

5.3.2 Voltage and waveform cuts

Any low-frequency fluctuations of the waveforms will affect the pulse-finding and parameterisation algorithms and, as a result, reduce the robustness of the parameters calculated. Consequently, waveform cuts were applied to ensure that the waveform descriptors remained within acceptable bounds. Additional cuts are applied to check the consistency between the high and low sensitivity channels, in both pulse area and time alignment.

Variations in the voltages applied to the grids or PMTs will affect the response of the instrument by altering the electric fields or the PMT gain. Despite the excellent stability observed during the first science run, cuts were placed on these voltages and currents to ensure a stable response.

5.3.3 Golden variables

Once these basic cuts have been applied to select single S1 and S2 events with high quality waveforms, new parameters can be calculated. These new, “golden” variables are listed below. Parameters such as pulse area, which require summing of the contributions from different PMTs, use the HS channel for all PMTs unless they are saturated, when the LS channel is taken.

- Event time information: *fil*, *dat*, *day*, *daqtime*, *lastime* - File number, date, day of run (since 27/02/2008), time since first golden event, and time since last ZE3RA event.
- Pulse ID: *s1id*, *s2id* - Pulse index for S1 and S2
- Saturation flag: *s1sat*, *s2sat* - Flag for saturation of LS channel
- Timing information: *s1time*, *s2time*, *dtime* - Time of S1 and S2 in the waveform and the drift time of the event.
- Pulse width: *s1tau*, *s2tau*, *s1wid50*, *s2wid50* - Mean charge arrival time for S1 and S2 and width at 50% amplitude.
- Pulse size: *s1area*, *s2area*, *s1amp*, *s2amp* - Corrected areas of S1 and S2 and their amplitudes.
- Energy: *energy*, *estar* - energy measured by S1 area and from combined S1-S2.
- Peak PMTs: *s1peak*, *s2peak* - PMT which sees most signal from S1 and S2.
- Corrections: *corlife*, *corpress*, *corfield*, *cordiff*, *corlevel*, *cortilt* - Multiplicative correction factors applied to the data.
- Centroid Position Reconstruction: Event position information (described in 5.4).
- Corrected Centroid Position Reconstruction: Event position information (described in 5.4).

Pulse areas are the most important parameters as they define the energy scale and the discrimination parameters. The raw pulse areas are corrected for light collection non-uniformity and slight variations in gain between PMTs, through application of a flat-fielding correction. In the case of *s2area*, corrections are also applied for the finite electron lifetime, detector tilt, field variations and changes in liquid level. These corrections are discussed later, although most (except the lifetime correction) were not required in analysis of the first science run. A correction factor, *cordiff*, is applied to *s2tau* correcting for diffusion of the electron cloud. Figure 5.7 shows the uncorrected and corrected *s2tau* against drift time; the clear flattening of *s2tau* allows for a more effective second pass cut.




Figure 5.7: Plots of the uncorrected (left) and corrected (right) $s2\tau$ parameters as a function of drift time, showing the diffusion effect and its correction.

5.3.4 Second pass cuts

The newly defined and corrected parameters allow for tighter quality cuts to be applied. For example, narrower cuts on τ are applied to both S1 ($5 < s1\tau < 40$ ns) and S2 ($400 < s2\tau < 1100$ ns) selecting high quality, well parameterised events. These more stringent selection cuts will remove some “good” events, although this is accounted for by the calculated efficiencies.

A cut on drift time is applied to remove events near the liquid surface and anomalous events with unphysical drift times ($100 < dtime < 18000$ ns). The minimum pulse area for which the LS can be utilised (instead of HS) is constrained to prevent small pulses containing an anomalous spike in a single PMT from passing the cuts. This is important as the parameterisation of such a pulse (from the LS channel) will be poorer, leading to inaccuracy in the event reconstruction. This situation is unlikely, but for a rare event search every eventuality must be accounted for.

5.3.5 Golden code efficiencies and cross-checking

Nuclear recoil detection efficiencies are an important consideration in the analysis process as these define how able we are to detect WIMPs. Calculation of efficiencies for the final analysis steps, not yet described, is relatively simple, but assessment during initial event selection is less straightforward. The aim of the selection code is to identify “good” golden events so efficiencies must be calculated by checking which events

do or do not pass the cuts. This was achieved by visually scanning a large number of calibration events and classifying them “manually” (i.e. a golden event, S1 only, multiple scatter, etc). From this classification and the counters within the golden code, the efficiencies were calculated from the nuclear recoil population in the AmBe calibration data. A breakdown is given below, resulting in an overall acceptance of 88.2%.

- First pass cuts - 96.0%
- Waveform cuts - 98.8 %
- Voltage cuts - 99.7 %
- Second pass cuts - 93.3 %

5.4 Position reconstruction

Position reconstruction in ZEPLIN-III is enhanced over that in ZEPLIN-II by the smaller PMTs, the greater light yield and the geometry of the target. The accurate three-dimensional reconstruction of interaction locations holds increased importance for ZEPLIN-III, as it is required for correction of light collection non-uniformities and removal of spurious events from multiple scatters in detector dead regions.

As with ZEPLIN-II, the z position of an event is defined by its drift time. The cathode grid appears at $14.0 \mu\text{s}$, equating to a drift depth of 36 mm. Figure 5.8 shows a typical drift time distribution from a science run, where the dominant source of background is the PMT array, and a ^{57}Co calibration, with the calibration source placed above the target.

Implementation of centroid position reconstruction algorithms (described for ZEPLIN-II in Section 4.5) in ZEPLIN-III provided a simple and robust method of gaining position information in the xy plane. The algorithms utilise the PMT locations weighted by the signal observed in each, providing adequate results for commissioning of the detector. The centroid algorithms were enhanced by utilising more complex weightings, calculated from simulations, yielding the *corrected centroid*, delivering a more even spatial distribution extending to greater radii. Whilst these methods were implemented in the golden code, assigning basic position parameters to the events, better precision was required for the rare event search. To this end, two different methods of reconstruction,




Figure 5.8: Drift time distributions with ^{57}Co source located above the target (left) and from background data with events more evenly distributed through the volume (right).

the *Template* and *Mercury* algorithms, were developed and will be briefly discussed here.

The orientation of the reconstructed position coordinates with regards to the PMT positions and lab environment are shown in Figure 5.9. The PMT array is arranged with PMT1 at the centre, with increasing PMT number moving outwards. The x direction designated for reconstruction passes through the centre of PMTs 17 to 14, with the y axis perpendicular to that.

The Mercury profile method utilises the spatial response profiles of the PMTs to reconstruct position. A 2-dimensional gaussian response function (shown Figure 5.10) describes the signal expected in each PMT from an event dependent upon its xy location, with the amplitude of the function defined by the event energy. By carrying out a least squares minimisation, matching the event hit pattern with the bell-shaped curves from all PMTs simultaneously, a location and energy can be extracted for each event. This process is carried out with both S1 and S2 pulses, yielding the xy positions of the two signals which should coincide for a single scatter event. However, as with the centroid method, the smaller signals from S1 produce a poorer position resolution. This method of reconstruction relies on the fact that the position-dependent response of each PMT is the same, which is a good assumption in ZEPLIN-III where the outer PMTs are far from the detector walls. However, if this were not the case, the bell-shaped response function would be distorted by reflection from the walls. In addition, this method also




Figure 5.9: Orientation of the position reconstruction coordinate system with respect to the PMT array and laboratory.

allows us to obtain a set of flat-fielding coefficients, resulting from the convergence of the least-squares fit.

The Template method was developed with the basic prescription detailed in [173]. A Monte Carlo simulation generates events at different positions across the volume, producing a PMT-array response template for each position and energy. Subsequently, for each event in the data, the pulses are tested against these templates, finding the one which best fits the event distribution using maximum likelihood. The templates are generated with simulations for S1 and S2 allowing the position to be reconstruction from either. This method outputs position, energy and goodness-of-fit values similar to those produced by the mercury profile method.

For the first science run of ZEPLIN-III the Mercury reconstruction was preferred. This decision was due to the independence of the method from detailed MC simulations,




Figure 5.10: Visualisation of the bell-shaped response function of PMT1 (left) and PMT2 (right) in ZEPLIN-III. The vertical axis shows the signal intensity, with the perpendicular axis being the x and y coordinates of the PMT array.

Table 5.2: Parameters from the *Mercury* reconstruction algorithm

Parameter	Description
$s1xm, s1ym$ $s2xm, s2ym$	x and y positions measured from S1 or S2
$s1rhom, s1thetam$ $s2rhom, s2thetam$	r and θ polar coordinates as measured by S1 and S2
$s1rmsm, s2rmsm$	The RMS of the spatial hit pattern for S1 and S2
$s1em, s2em$ $s1e1m$	Event energy calculated from S1 and S2 S1 energy calculated using S2 position
$s1chim, s2chim$	Measures of the χ^2 goodness-of-fit from the least squares minimisation
$s12dism$	The distance between the event locations measured by S1 and S2

which were still being perfected for the template method. Further assessment of the methods will be carried out before analysis of the second science run. Figure 5.29 shows the xy distribution of events from the AmBe calibration dataset, where the source was located to one side of the detector, compared with the more evenly distributed ^{137}Cs calibration (as shown later in this chapter in Figure 5.33) and science data sets (as shown in Figure 5.40).

Following initial inspection of the final science dataset, a weakness in the Mercury reconstruction method was exposed. Reconstruction of event energy is dependent upon the accuracy of the position reconstruction during the least squares minimisation process. The poor position resolution for small S1 signals introduced an inaccuracy in the energy determined from S1 alone, $s1em$. This was solved by fitting for the S1 energy,

but utilising the higher resolution position from S2, yielding an improved measure of S1 energy, *s1e1m*.

5.5 Detector monitoring from data

As with any experiment the stability of the operating conditions must be monitored, to inform adjustment of the experimental setup or correction of the data. Monitoring of detector parameters such as temperatures and pressures is done using a slow control system. However, there are other operational factors which can only be accurately probed using data from the instrument. In this section we describe the stability tests carried out on a daily basis, their implications for the data and any corrections subsequently derived.

During the first science run, a number of analysis routines were developed to monitor operational parameters such as xenon purity, detector tilt, light yield and DRU (differential rate unit) background. These routines required multiple user inputs. However, it was soon realised that during routine operations, a more robust and automated stability monitoring procedure was required, to remove user-dependent systematics and to relieve the operational team of a time consuming task. An automatic routine was therefore developed to analyse daily ^{57}Co calibrations and a 10% sample of background data. This routine processed the appropriate datasets with the golden code using the correct parameters required for each of the different tests. The different steps in the process are listed below, followed by a detailed description.

1. Process daily ^{57}Co data (no purity correction, default ecal)
 - Xenon purity routine - outputs electron lifetime measurement
2. Reprocess daily ^{57}Co data (calculated purity correction, default ecal)
 - Detector tilt routine - outputs magnitude and phase of detector tilt and average liquid level.
 - Channel amplification consistency check
 - Light yield calibration - Outputs the interaction rate from ^{57}Co and light yield measurements.

3. Process 10% of daily background data with golden code (previously calculated purity correction and energy calibrations for S1 and S2).

- DRU routine - produces differential energy spectrum for 10% science data.

The following sections describe the different routines, including the methods utilised during the first science run and any improvements prepared for the second run. A discussion is also provided of the corrections which may arise from the different features in the data and how they are calculated.

5.5.1 Xenon purity

As explained in the discussion of ZEPLIN-II, trapping of electrons by electronegative contaminants as they drift through the liquid xenon is a major factor varying the size of the S2 signal. For this reason it is crucial that we measure the xenon purity (i.e. the free electron lifetime) with reasonable accuracy, monitor it throughout the science run and calculate the accurate correction it required. This is done with the same basic method as in ZEPLIN-II, by slicing the distribution of S2/S1 as a function of drift time and fitting an exponential to the resulting slice means.

The purity measurements tend to show some variability depending on the slicing that is applied in drift time. In ZEPLIN-II equal bins were used due to the uniformly-distributed α population used for the measurement. However, in ZEPLIN-III the depth distribution of interactions from ^{57}Co is exponentially falling from the surface of the liquid, as shown in Figure 5.8, meaning that constant-width slices would have decreasing numbers of events and progressively larger statistical errors. As a result, slicing with a varying bin size was carried out, with the width of slice increasing with drift time. This method along with fixed start and end points for the slicing reduced systematics from the different inputs of individual users. Figure 5.11 shows an example of a typical electron lifetime measurement from the first science run.

Figure 5.12 shows the historical variation of the electron lifetime measured throughout the first science run. ZEPLIN-III was designed such that the purity of the xenon would not degrade over time due to the clean, metal-dominated construction of the target. This was shown to be true, with the electron lifetime actually improving over the course of the 84 day run. An electron lifetime of about $20\ \mu\text{s}$ had been achieved




Figure 5.11: A typical electron lifetime measurement from the first science run of ZEPLIN-III, extracted by the original automatic monitoring routine.




Figure 5.12: Trend in electron lifetime measured using ^{57}Co through the course of the science run, where a lifetime of $25\ \mu\text{s}$ equates to a correction factor of 1.82 for events from the bottom of the detector.

by external cleaning with cold getters before filling of the detector. Throughout the run the lifetime steadily improved up to $\sim 40 \mu\text{s}$. This was attributed to the extremely low outgas rates achieved in the clean target and continuous application of the strong electric field: if no new electronegative contaminants are introduced into the system, the negative ions produced through attachment cannot trap a further electron and will be drifted away from the fiducial volume by the strong field. The lifetime was studied more precisely after completion of the science run, before calculation of a correction.

Monitoring during the first science run was carried out using the daily ^{57}Co calibrations, providing a consistent dataset for stability monitoring. Assessment of the background data and ^{137}Cs calibrations indicated that an adjustment to the measured electron lifetime was required. Lower values of electron lifetime were measured from the science data taken on the same day, where the majority of events had energies $< 50 \text{ keV}_{ee}$. As a result, the electron lifetime was measured for different energies in the ^{137}Cs calibration data. Figure 5.13 shows the electron lifetime, measured as a function of energy, from the ^{137}Cs science calibration data. A clear dependence is observed down to low energies, introducing the possibility of a purity correction dependent upon both time and energy. This variation is attributed to PMT non-linearity caused by the large S2 signals. The larger signals from interactions at the top of the liquid, affected less by electron attachment, will have their observed S2 signal undergo a greater reduction (from non-linearity) compared to a smaller S2 from an interaction at the bottom of the liquid. This leads to an apparent increase in electron lifetime with increasing energy, as the fitted S2/S1 slope is flattened by the PMT non-linearity.

It can be seen from Figure 5.13 that the electron lifetime measurements are boosted at higher energies by up to about 50%. As a result, the electron lifetimes measured at 122 keV_{ee} with ^{57}Co data were adjusted by a factor of $2/3$, which was consistent with the electron lifetime measured at lower energies with the science data.

5.5.2 Liquid level and detector tilt

The amplitude of S2 signals is dependent upon the number of electrons which are drifted through the gas and the magnitude of the field accelerating them, with the pulse width primarily defined by the gas region thickness. As a result, the total amount




Figure 5.13: Electron lifetime measured as a function of S1-defined energy, measured from ^{137}Cs science calibrations.

of light produced is dependent upon all these factors. Therefore any variation in the thickness of the gas gap will result in a change in S2 response. A long-term change in the overall liquid level introduces a time-dependent S2 response, with an additional position dependence introduced by any tilt of the target. If one side of the target has a larger gas gap the S2 signals produced there will be greater. As a result of these two potential variables, the overall liquid level and detector tilt must be monitored.

After filling of the instrument, the tilt of the detector is assessed using the levelling routines described here and the instrument adjusted accordingly using three external levelling screws. The thickness of the gas gap is best probed by the width of the S2 pulses, measuring the transit time of electrons across it. The tilt is therefore measured by assessing the variation in S2 width across the xy -plane.

The FWHM of S2 pulses is plotted as a function of the angular position in the detector, as shown in Figure 5.14, with a sinusoidal trend about the mid-level indicating the presence of a tilt. During the first science run the tilt magnitude and phase, as well as the mean liquid level, were thus determined on a daily basis. Figure 5.15 shows the historical trend of the tilt parameters, with the increase in the tilt amplitude and phase




Figure 5.14: Example levelling plot from the original monitoring routine, showing a sinusoidal behaviour of $s2wid50$ as a function of $s2theta$.

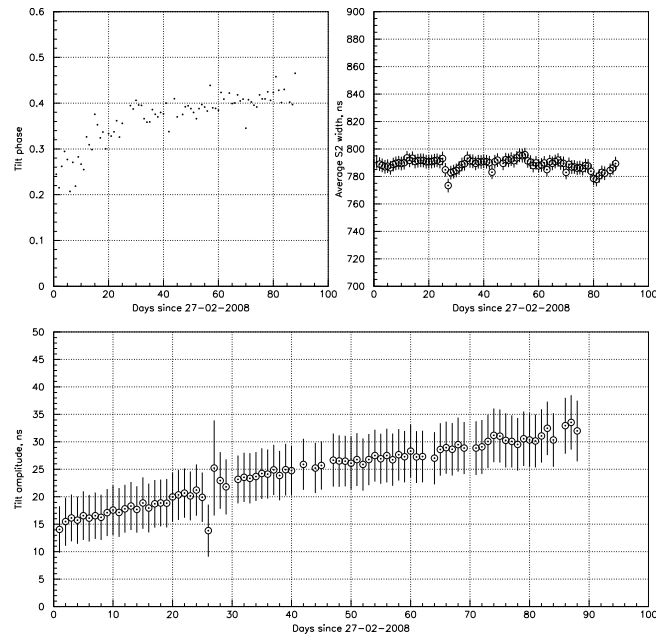


Figure 5.15: Historical trend plot of the phase of detector tilt (top-left), average liquid level (top-right) and detector tilt amplitude (bottom) during the first science run.

attributed to local geological movement. The varying tilt of the detector introduces additional time-dependence to the position-dependent S2 response. The tilt of the detector at the start of the science run was ~ 1 mrad, increasing to ~ 2 mrad by the end of the run. We decided not to correct this tilt during the run, but to correct its effects in the data analysis.

Whilst this simple method proved sufficient for monitoring stability during the first science run, a more precise method was developed for the second science run, taking advantage of the advanced position reconstruction algorithms developed during the first run. The magnitude of any variation in *s2wid50* will clearly depend upon the radial position of the event. Therefore by cosine-fitting for events in radial slices, the magnitude of the tilt and the phase can be extracted more precisely (Figure 5.16). As the S2 width variation increases with increasing radius, slicing into concentric circles yields higher quality cosine fits. Figure 5.16 shows the extracted parameters plotted against radius, where the slope of the amplitude plot describes the detector tilt, with the phase and baseline defined by their constant values.

This method also provides a confirmation of the range of linearity for the position reconstruction algorithm, which is confirmed in Figure 5.16 out to the fiducial radius of 150 mm using the *Mercury* least-squares reconstruction algorithm.




Figure 5.16: Example of levelling plots from the enhanced monitoring routine, with the cosine amplitude, phase and baseline measured at different radii.

5.5.3 Channel amplification stability check

As the 31 PMTs are read together and analysed under the assumption of a $10\times$ difference between low and high sensitivity, it is important to check the amplification for each PMT individually. Figure 5.17 shows the method utilised during the first science run, plotting the signal from the high sensitivity against that from the low sensitivity for each PMT. If the amplification is correct the points should follow a linear relationship in the middle energy region of the plot. The sensitivity to small signals is reduced in the LS channel and large signals begin to saturate in the HS channel, causing the points to depart from the linear relationship at the extremes. However, this is not a problem for the analysis as the code selects the appropriate channel to utilise.

The amplification checks carried out during the first science run were done ‘by eye’, using the plots shown, but a further development would be to apply a statistical test to the data outputting a figure of merit for linearity. Other similar checks could be developed to monitor the baselines, noise and other characteristics of the waveforms.




Figure 5.17: Example plots showing high sensitivity signal against that in the low sensitivity channel for each PMT, used to check consistency of channel amplification.

5.5.4 Light yield stability

Monitoring of the energy calibration was carried out using the daily ^{57}Co calibrations in a similar method to that used in ZEPLIN-II, using the position of the main γ -ray peak in both S1 and S2. Before automation, a user input was required for the bounds of the fitting, producing additional systematics in the result. A method was required with no user input, initially achieved by using a single gaussian fit to the 122/136 keV combined peak with fixed fit bounds. This worked as a very basic check, but fits were often poor, so little could be gained in the way of a correction. As a result, an improved automated method was developed to provide more reliable fitting and enhanced monitoring of detector response. Figure 5.18 shows the stability of the S1 and S2 peaks to within a few percent, measured from the daily ^{57}Co runs.

The improved method utilises a fit to the top of the photo-peak in S1 (top 50% of distribution) and S2 (top 30%), producing a more reliable and precise measurement of the peak positions. The shift from a reference calibration (chosen as the first day of the science run) was calculated, with the histogram subsequently normalised. By comparing the shifted and reference histograms, the shape of the distributions can also be assessed. Figure 5.19 shows the plots produced by the improved method of fitting. A statistical check should eventually be implemented to assess the consistency of calibration runs.

5.5.5 Low-energy electron recoil background rates

The electron recoil background was monitored in the detector by producing a differential energy spectrum. For this purpose the background event rate is measured in differential rate units, DRU (events/kg/day/keV $_{ee}$). A simple routine to plot the differential spectrum using the 10% daily sample of science data was utilised to monitor the background levels throughout the run. The DRU value at 30 keV $_{ee}$ was extracted as a monitoring point to complement the spectrum produced (shown in Figure 5.20). The levels of low-energy background remained within statistical errors throughout the run.

5.5.6 Corrections

As with ZEPLIN-II, corrections were applied for factors affecting the response of the detector. Variations in PMT response and light collection non-uniformities across the




Figure 5.18: Historical trend plot of the energy calibration values for S1 (left) and S2 (right) throughout the first science run.

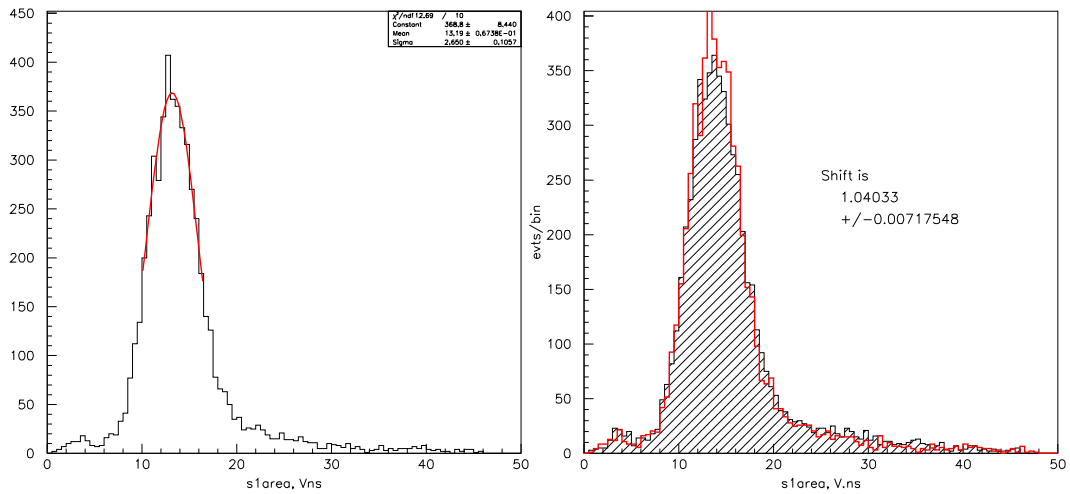


Figure 5.19: Improved method of ^{57}Co energy calibration with the fitting to the S1 spectrum (left), and the adjusted histogram (red) for shape comparison with reference spectrum (shaded) (right).




Figure 5.20: A typical differential energy spectrum produced from a daily 10% sample of the science data, showing the level of background in DRU (events/kg/day/keV_{ee}).

PMT array were corrected by fitting of common cylindrical response functions to the S2 signals in the *Mercury* iterative process, resulting in an array flat-fielding correction. No pressure correction was required, due to the excellent thermal stability achieved during the run. The temperature variation of less than 0.5°C produced only small pressure variations of less than 2% over the first science run. As mentioned previously, the variation in tilt was not corrected for during the primary analysis, although this is under further investigation. The remaining correction, for the finite electron lifetime, was applied in the same manner as ZEPLIN-II (see Section 4.4.4), with the electron lifetime applied being $2/3$ of that measured with the ^{57}Co monitoring data (due to the PMT non-linearity for large S2).

5.6 ^{57}Co energy calibrations

Detailed analysis of ^{57}Co calibrations showed good agreement with simulations, not only for the two main γ -ray lines, but also for the expected Compton feature at around 35 keV, shown in Figure 5.21. The high energy excess above 150 keV is attributed to unsubtracted background. From fitting to the 122 keV peak in S1, a light yield of 1.8 phe/keV_{ee} was measured at the operating field in the liquid of 3.9 kV/cm. This is compared with 5.0 phe/keV_{ee} at zero-field, a larger field suppression (74%) than observed in ZEPLIN-II ($\sim 50\%$) due to the higher operating field.




Figure 5.21: Comparison of ZEPLIN-III ^{57}Co calibration (blue line) with simulations of the instrument (red shaded areas). Experimental data is shown as the blue histogram, with solid red showing the pure simulated spectrum and hatched red the simulated spectrum smeared with an energy resolution.

The energy resolution was also measured from the ^{57}Co data, with individual S1 and S2 peaks giving $\sigma = 16.3\%$ and $\sigma = 8.8\%$ respectively, at 122 keV. The resolution in S2 is better than that in S1, as $\sim 64\%$ of the visible energy is deposited through the ionisation channel. Utilising the combined measure of energy, E^* (discussed in Chapter 7.6), a resolution $\sigma = 5.4\%$ was achieved, which is one of the best reported from a two-phase xenon detector.

5.7 Dead-region events

One of the major challenges in the analysis of ZEPLIN-III data was the removal of double-Compton events where one of the interactions occurs in a dead region of the detector. These multiple-scintillation single-ionisation (MSSI) events appear in the data as a golden event, as the two scintillation pulses are coincident in time, appearing as a single S1, with only one of the interactions producing an S2. The nature of these events, with multiple S1 combining to produce a single, larger S1, yields a reduction in S2/S1 value. This causes MSSI events to leak down from the main electron recoil population, which can, in extreme cases, contaminate the nuclear recoil signal region. The magnitude of the decrease in S2/S1 is dependent upon the dominance of the dead region scatter, i.e. if the scatter in the dead region produces a much larger S1 than the active scatter.

A dead region in the detector is any LXe volume from which ionisation cannot be extracted. This includes regions with no applied field, a reverse field, or field lines which do not reach the liquid surface. In ZEPLIN-III there is a large volume of xenon outside the target volume which has no applied field, although the majority of this has little or no light collection coupling to the PMT array, causing no contamination of events. However, there are three distinct dead regions of importance, highlighted in Figure 5.22.

The layer of xenon between the PMT array and PMT grid (*A* in Figure 5.22) has no field, since both components are kept at the same voltage. In the xenon layer just above, between the PMT and cathode grids (*B*), there is a strong reverse field. The region around the outer edge of the target (*C*), has non-vertical field lines, shown in Figure 5.22, which creates a region where ionisation is drifted, but does not reach the surface. The boundary of this region can be located by assessing events with apparently non-physical drift times (i.e. $dtime > 14 \mu s$). Figure 5.23 shows the radial position of events as a function of their drift time, showing that events beyond $14 \mu s$ have a reconstructed radius of ~ 180 mm. These events are those with ionisation drifted along a curved field line, taking longer to reach the surface. Note also the slight increase of this radius with increasing drift time, as the field lines become increasingly curved.

During the first science run, when the PMT array provided the dominant source of background, the regions directly above the PMT array will contribute the largest

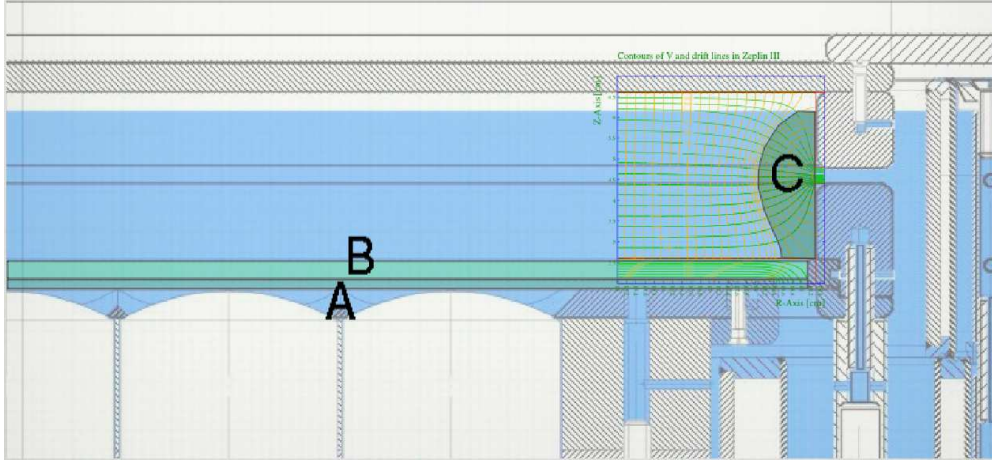


Figure 5.22: Diagram of the ZEPLIN-III target geometry highlighting detector dead regions. Region *A*, between PMT array and PMT grid, has no electric field. Region *B*, between PMT and cathode grids, has a strong reverse field. In region *C*, non-vertical field lines do not reach surface for emission into gas phase. The orange lines show the electric field lines and the green lines show the equipotentials.




Figure 5.23: Radial position of events from beyond the nominal maximum drift time of the central region.

number of MSSI events. MSSI events where the dominant scatter is in one of these regions (A or B) should have uncharacteristically tight scintillation photon hit pattern on the PMT array. On the other hand, MSSI events with a dominant scatter in region C will have an unremarkable hit pattern, but will have an xy S1 location significantly different from that of the fiducial S2. For MSSI events where neither scatter is totally dominant the hit pattern will be inconsistent with a single scatter, as the scintillation light was produced from two different locations in the detector.

MSSI events could potentially be a limiting factor to the sensitivity of the instrument, so it is vital to understand them and devise cuts to remove them from the data. Here we discuss two possible methods utilised to study the properties of these events, and the cuts which were devised to reject them.

5.7.1 S1-only events

In the 10% sample of science data, a large proportion ($\sim 80\%$) of events triggering the data acquisition contained only a prompt scintillation pulse. Although some of these are β s from ^{40}K emitted from the PMTs and interacting in the liquid xenon just outside the PMT envelopes, many are γ -rays interacting in the reverse field (Regions A and B). This large contribution from the PMT γ -ray background is clear from Figure 5.24. By extrapolating the exponentially increasing background, one would expect 60% of single scatters to occur below the cathode. By studying these S1-only events we can understand the potential properties of MSSI events with the dominant scatter below the cathode grid.

The xy reconstruction of S1-only events locates them close to the centre of individual PMTs, suggesting their production close to the PMT array. This is compared with single scatters from the active region, where the reconstructed S1 position are much more evenly distributed. These S1-only events would be expected to have a similarly tight hit pattern on the PMT array as an MSSI event from that region. This effect is well quantified by the RMS of the S1 spatial distribution. Figure 5.25 shows S1 position RMS calculated from the centroid reconstruction method and shows that light from the S1-only events has a less diffuse hit pattern on the PMT array. The S1-only events with small $s1rmse$ values are located closer to an individual PMT, although




Figure 5.24: Drift time distribution of single scatters in the active volume, fitted with an exponential plus a constant to demonstrate large number of potential S1-only events from below the cathode. The blue dashed lines show the positions of the cathode grid (left), PMT grid (centre) and PMT array (right).

there is a limitation of this effect as S1-only events must pass the 3-fold event selection requirement, which is not the case for the MSSI events.

A further confirmation of this hypothesis is the $s1\tau$ (analogous to the scintillation time constant) measured for the two populations. The S1-only events have a slightly higher mean tau value, suggesting a contribution of events occurring in a zero-field region, where near-complete charge recombination increases $s1\tau$. However, a large number of the S1-only events will still have similar S1 tau, as they arise in the reverse field region between cathode and PMT grids.

5.7.2 Leaking MSSI events

The data-sets recorded in the low background configuration of ZEPLIN-III exhibit the expected populations from electron and nuclear recoils, but with associated MSSI events. This poses a potential problem for ZEPLIN-III if events leaking from the electron recoil

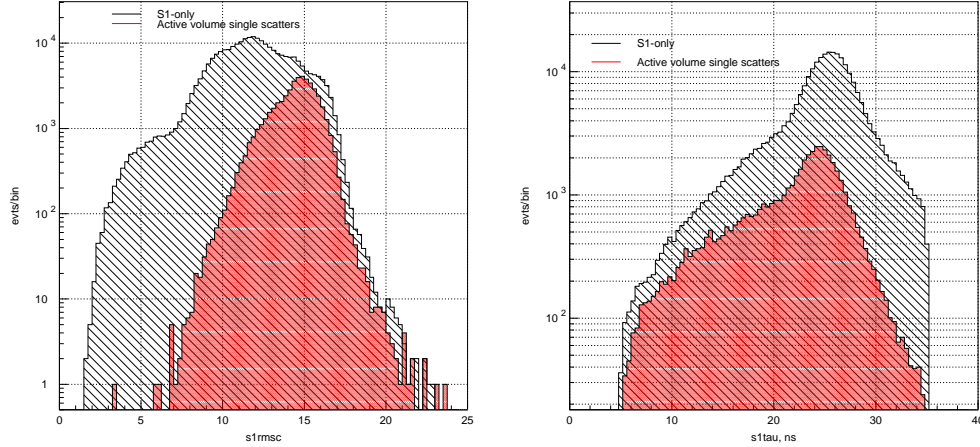


Figure 5.25: S1 centroid position reconstruction RMS plot for S1-only events and active volume single scatters (left). S1 tau distributions for S1-only and active volume single scatters (right).

background reach into the signal region. This population, selected graphically from calibration scatter plots, can be used to help develop cuts for the removal of MSSI events.

It is observed that MSSI events increase in number at higher energies, which is fortunate as the low energy range holds the significant WIMP sensitivity. However, they are not completely absent at low energies, and specific cuts are still required to reject them. The reduction in their frequency at low energies is a result of the low probability of two low-energy scattering events occurring in the detector, with the larger of the two occurring in a dead region.

5.7.3 Removing MSSI events

Utilising what we have learnt from our study of S1-only and the higher energy MSSI events, along with our understanding of detector dead regions, we may construct cuts to remove a large fraction of the critical lower-energy MSSI events. From the S1-only events, we demonstrated that those originating below the cathode grid have a much tighter hit pattern on the PMT array, manifested in smaller spatial RMS values. Plotting this quantity ($s1rmsm$) for events from the main population and for events leaking below the main electron recoil population, as shown in Figure 5.26, it is clear that, as expected, a significant fraction of the leaking MSSI events have a small $s1rmsm$.




Figure 5.26: Top: $s1rmsm$ (left) and the combined χ^2 parameter (right) plotted against energy for events from the main electron recoil population (shaded) and graphically selected MSSI events leaking below the electron recoil band (red). The appropriate MSSI removal cuts are shown (green line). Bottom: Nuclear recoil detection efficiency of the MSSI cuts, calculated with an on-off comparison of nuclear recoils from AmBe calibrations.

As a result, an energy-dependent cut (also shown in the figure) was formulated, which allows for the loss of spatial resolution observed at low energies.

An additional cut was applied utilising the χ^2 goodness-of-fit parameter from the Mercury position reconstruction algorithm. For events with multiple vertices, the scintillation photon hit pattern across the PMT array should have poor agreement with the expected response for a single scatter. As a result, we can reject MSSI events by cutting those with an excessive χ^2 value. It was found that the best measure of χ^2 for discrimi-

nating against leaking MSSI events was a parameter combining the S1 and S2 χ^2 values. As with the RMS cut, an energy-dependent cut was formulated (shown in Figure 5.26) to reject a large fraction of leaking events, whilst retaining a significant efficiency. The two cuts have a combined nuclear recoil detection efficiency of 73% between 2-16 keV_{ee}, as shown in Figure 5.26.

The mismatch of the reconstructed S1 and S2 xy locations could additionally be utilised to remove MSSI events, which would be particularly powerful for rejecting region C vertices. An additional cut on $s12dism$ was trialled, but was found to add no significant rejection power and as a result was disregarded for the analysis of the first science run. When the dominant source of background is no longer the PMT array, an $s12dism$ cut could prove important.

5.8 Fiducial volume cuts

The active target mass of liquid xenon in ZEPLIN-III is $\simeq 12$ kg, although some of this is located in peripheral regions with non-uniform fields and low light collection or in close proximity to sources of background. Consequently, cuts were placed on the location of events in the xy plane and z direction, selecting a central reduced-background region of uniform response.

Ideally, a cut on drift time would remove the majority of the dominant electron recoil background from the PMT array. However, the relatively short depth of the ZEPLIN-III target (~ 14 μ s) precludes this. The fiducial cut on the drift time is thus set to remove background from the cathode grid only.

Investigation was required of the maximum allowable drift depth before the proximity to the cathode grid becomes significant. Very close to the cathode grid the strength of the electric field increases and becomes non-uniform, leading to an anomalous S2/S1 response; similar irregularity is expected from S1 light collection. Figure 5.27 illustrates the position of the cathode grid in drift time, by measuring S2/S1 response, with the stronger electric field close to the cathode grid wires suppressing recombination and hence increasing the S2/S1 ratio of events from that region. This effect is clearly visible between 14.0 μ s and 14.2 μ s. As with ZEPLIN-II, a small population of nuclear recoil-




Figure 5.27: S2/S1 response as a function of drift time highlighting the position of the cathode (shown by the vertical lines).




Figure 5.28: Drift time (left) and radial (right) distributions from science data with the fiducial region shaded.

like events are observed from the cathode grid, these appear with low S2/S1 values which can pose a problem for the WIMP search. As a result, a conservative drift time cut is placed at $13 \mu\text{s}$ to remove all such events.

Selection in the xy plane utilises a radial cut, removing non-uniform regions at the detector edge, also considering the linearity of the position reconstruction out to large radii. Figure 5.16 demonstrates linearity of the radial positions from the mercury reconstruction out to 150 mm, showing that beyond the edge of the PMT array, reconstruction of event location becomes more difficult, resulting in a less accurate energy calculation. Consequently, the radial cut selected was placed at 150 mm ($s2rhom < 150$). These fiducial cuts define a region with a reduced-background and good uniformity of response, leaving a fiducial mass of 6.52 kg (for a liquid xenon density of 2.9 g/cm^3).

5.9 Discrimination parameter calibrations

Calibration of the predicted nuclear recoil signature from WIMPs in the fiducial region just defined is carried out with an AmBe (α, n) source. The source was inserted inside the polypropylene shield, above the detector, but retracted to one side to optimise the event rate. Neutrons interact with the Xe producing elastic nuclear recoils and inelastic scatters. Calibration of the expected electron recoil background used a ^{137}Cs source, emitting mainly 662 keV γ -rays which populate the low energy region of the scatter plot through Compton scattering. To cross-check the background calibration, a 10% sample of the science data was also analysed. In an ideal situation the calibration dataset would contain many more events than the science data-set, allowing for high-statistics characterisation of the tails of the electron recoil population which are key to the sensitivity of the rare event search. However, due to the huge volume of data that would be required, this proved impractical.

5.9.1 Nuclear recoil population: AmBe neutron calibrations

The AmBe dataset used to calibrate the first science run of ZEPLIN-III lasted for about 5 hours (24th and 25th May 2008) producing about 4841 pure elastic nuclear recoils in the fiducial volume (all energies, all cuts applied). Additional AmBe calibrations were carried out prior to the run, at a slightly lower pressure, and after the run with a lower




Figure 5.29: AmBe spatial distributions (for events between 2 and 16 keV_{ee}), drift time (left) and xy plane (right).

threshold. Neither were used to calibrate the science run, but were analysed nonetheless and found to be consistent with the main calibration.

Figure 5.29 shows the xy distribution of events from the calibration, exhibiting a noticeable spatial bias due to the source position. The drift time distribution (also shown) demonstrates the expected trend towards the top of the target, again consistent with the source location. After application of fiducial and MSSI cuts, the scatter plot (Figure 5.31) shows the three familiar interaction populations: elastic nuclear recoils, inelastic nuclear recoils (from ^{129}Xe) and electron recoils. The behaviour of the nuclear recoil population with energy is similar to that observed in the ZEPLIN-II and XENON10 experiments, but with ZEPLIN-III exhibiting tighter distributions and less flaring at low energy, with both of these features attributed to the higher operating field. The inelastic populations are also reminiscent of those described previously for ZEPLIN-II in Chapter 4.

Characterisation of the nuclear recoil “signal” population from the AmBe calibration utilised a similar method to that used in ZEPLIN-II, selecting the elastic recoil population and slicing in 1 keV_{ee} bins. In each slice, S2/S1 was fitted with a log-normal




Figure 5.30: Examples of discrimination parameter distributions in 3 energy slices (low, middle and high energy bins within acceptance region) of the nuclear recoil population from AmBe neutron calibrations.

distribution (Figure 5.30), yielding the energy dependence of the population means and standard deviations (Figure 5.32). The mean S2/S1 of the nuclear recoil band was thus described by a power law of the form $\mu = \log_{10}((0.930 \pm 0.007)E^{-(0.540 \pm 0.003)})$, with its width parameterised as $\sigma = (0.180 \pm 0.010)E^{-(0.295 \pm 0.024)}$, where E is measured in keV_{ee}. These parameterisations map the signal efficiency in the S2/S1 vs energy parameter space, allowing definition of a WIMP search region.

5.9.2 Electron recoil population: ¹³⁷Cs Compton calibrations

The ¹³⁷Cs calibration was carried out over the period of 9 days, attaining 122 hours of data, achieving good statistics for characterisation of the expected electron recoil background. The source was projected inside the shielding (again retracted to optimise event rate), populating the energy spectrum down to about 2 keV_{ee} through Compton scattering in the fiducial volume. The spatial event distributions are shown in Figure 5.33.

Figure 5.34 shows a comparison of the electron recoil population (from ¹³⁷Cs calibration) with the nuclear recoil population (from AmBe calibration). The electron recoil population exhibits similar behaviour to that observed in previous two-phase Xe experiments, but with tighter distributions, a more pronounced upturn at low energies and larger separation from the nuclear recoil population. This improved separation is one of the most important results of the ZEPLIN-III experiment, displaying greatly improved discrimination at higher field.




Figure 5.31: Discrimination parameter scatter plot of AmBe data, showing the nuclear recoil elastic scattering population, with the mean (red line) and 1σ (blue line) bands shown, and the inelastic scattering population at large $S2/S1$.




Figure 5.32: Parameterisation of nuclear recoil population mean (left) and standard deviation (right) measured from AmBe neutron calibrations.




Figure 5.33: ^{137}Cs spatial distributions in drift time (left) and the xy plane (right), for events between 2 and 16 keV_{ee} .




Figure 5.34: Discrimination parameter scatter plot showing the nuclear recoil population from the AmBe neutron calibration (red) and the electron recoil population from the ^{137}Cs calibration (blue). The 50% nuclear recoil acceptance line is shown in red.

Characterisation of the electron recoil population utilises the same energy binning as the AmBe calibration, but this time fitting a skew-gaussian distribution:

$$f(x) = \frac{A}{\omega\sqrt{2\pi}} e^{-\frac{(x-x_c)^2}{2\omega^2}} \left[1 + \operatorname{erf} \left(\frac{\alpha}{\sqrt{2}} \frac{x-x_c}{\omega} \right) \right], \quad (5.1)$$

where $\mu = x_c + \omega\delta\sqrt{2/\pi}$ is the mean and the standard deviation is given by $\sigma^2 = \omega^2(1 - 2\delta^2/\pi)$, when $\delta = \alpha/\sqrt{(1 + \alpha^2)}$. High quality fits were achieved to the complete S2/S1 distributions using the skew-gaussian, as shown further ahead, although the low S2/S1 tail exhibits more outliers (attributed to MSSI events) than the science data. This was attributed to the differing angular distributions for γ -rays between science and calibration data, creating difficulties in mimicking the MSSI population exactly.

5.9.3 The 10% science dataset sample

As with ZEPLIN-II, 90% of the science dataset was initially retained with the intention of carrying out a blind analysis, with the remaining 10% (every 10th file) utilised for monitoring of detector stability and development of the analysis procedures. Comparison of the skew-gaussian fits to the ^{137}Cs calibration and 10% science data sample showed their μ and σ to be in good agreement. The science data sample scatter plot (Figure 5.35) contains proportionally fewer outliers than the ^{137}Cs data, attributed to the calibration not mimicking the γ -ray background accurately. This disagreement is most apparent in the region below 30 keV_{ee}, where the MSSI cuts successfully remove outliers in the science data, with less effective rejection in the ^{137}Cs calibration. As a result of this disagreement in the tails, it was decided to use the complete science dataset itself to characterise the electron recoil population, bearing in mind that the two datasets produced statistically consistent values for μ and σ bin by bin and differed only far below the mean.

5.9.4 Definition of a WIMP search region

When searching for a WIMP signal we examine a region of the parameter space containing a large fraction of signal with little background contamination, achieved by defining a WIMP search box. Using a comparison between the 10% of science data




Figure 5.35: Discrimination parameter scatter plot showing the 10% sample of science data from the ZEPLIN-III first science run, with the WIMP search region shown (red box) along with the nuclear recoil mean (red) and 1σ (blue dashed) lines. Note that no events are observed in the WIMP search region in the 10% data sample.

and the nuclear recoil calibration, the upper bound was chosen as the mean of the NR population, $\mu_{nr}(E)$, with a lower limit set at $\mu_{nr}(E) - 2\sigma_{nr}(E)$. As with ZEPLIN-II, the energy range of the search region is determined after consideration of the detector efficiencies and backgrounds. The lower energy threshold in ZEPLIN-III allows a lower energy bound of 2 keV_{ee}. The upper energy limit was decided taking into consideration the distribution of MSSI events. The probability of an MSSI event increases with increasing energy, whilst the signal probability is exponentially decreasing. Consequently, the upper energy limit was defined as 16 keV_{ee}. This region of the parameter space includes 47.7% of nuclear recoil acceptance from 2 to 16 keV_{ee}, the dominate energy range for a WIMP signal in the detector. This box was defined as part of the intended blind analysis procedure, and remained unchanged after the opening and subsequent adjustments to the analysis.

5.10 Energy conversion and energy-dependent efficiencies

As discussed in Chapter 3.1, the differing energy deposition processes of nuclear and electron recoils result in a different scintillation response, requiring a conversion between energy scales. In ZEPLIN-II, as with the initial analysis of the XENON10 experiment, a flat energy conversion factor between keV_{ee} and keV_{nr} was applied. The conversion in energy is calculated from:

$$E_{nr} = \frac{S_e}{L_{eff} S_n} E_{ee}, \quad (5.2)$$

where the zero-field relative scintillation efficiency for nuclear recoils, L_{eff} , and field suppression factors, S_e and S_n , are utilised (see Section 3.1).

In ZEPLIN-II, an energy-independent $L_{eff} = 0.19$ was used along with appropriate field suppression factors, which at the ZEPLIN-III operating field are $S_e = 0.36$ and $S_n = 0.9$ [147]. As discussed in Section 3.1, previous measurements of the relative scintillation efficiency had only been achieved down to approximately 10 keV_{nr} , with the different measurements showing some small disagreement, with an average L_{eff} of 0.19 providing a reasonable approximation. The field suppression factor, S_n , was measured at one recoil energy (56 keV_{nr}), with no energy dependence measured. This value has also been taken as constant. Using Eq. 5.2 with the energy-independent terms, we find an energy conversion of $E_{nr} = 2.09 E_{ee}$. However, there is no firm reason to believe that the three factors involved are indeed energy-independent.

A reasonable way of measuring the energy-dependent nuclear recoil detection efficiency of the instrument is through comparison between the experimental and simulated AmBe spectra. Figure 5.36 shows the experimental spectrum (shaded blue), scaled using a flat conversion factor of 2.09, compared to the simulated recoil spectrum (solid red curve). A clear mismatch is observed below about 20 keV_{ee} , suggesting a poor detection efficiency. However, this is at odds with other calculations of efficiency, detailed in Table 5.3, both energy-dependent and independent, from both hardware and software, by a very large factor.

The hardware trigger has been studied using a pulser and through examination of events in the data. This included the simulation of the effect of the trigger electronics applied to the high sensitivity sum timeline to derive a realistic trigger waveform.




Figure 5.36: Comparison of simulated data with nuclear recoil spectrum from AmBe calibrations (with standard threshold and low threshold). A large discrepancy is observed below ~ 10 keV_{ee} which is not explained by instrument efficiencies.

These studies clearly demonstrate that the hardware trigger has a threshold below 1 keV_{ee} when triggering on S2 and could not possibly be affecting energies $> 10 \text{ keV}_{ee}$. Analysis of reduced-threshold AmBe calibrations (black line in Figure 5.36) shows the same spectral shape as the main calibration above 4 keV_{ee} , again highlighting that the hardware trigger is not responsible for any mismatch. The software S2 area threshold ($S2 > 5 \text{ V.ns}$) is calculated from the data to equate to less than 1 keV_{ee} , with the software 3-fold requirement on S1 dominating with a threshold of about 1.7 keV_{ee} . These clearly do not account for the mismatch observed with the simulations. The presumed energy-independent cuts have also been assessed for energy dependence with none found.

An alternative approach to understand the mismatch is to investigate the possibility of a non-linearity in the S1 and/or S2 channels. This could be manifested through a varying energy conversion from E_{ee} to E_{nr} , either through the relative scintillation efficiency, field suppression factors, or both. This approach was initially applied to the

Table 5.3: Summary of calculated efficiencies

Energy-independent efficiencies		
Effect	Efficiency	Method of calculation
Deadtime	91.7%	Measured
Hardware saturation cut	100%	On-Off comparison
ZE3RA pulse finding	96.0%	Event scanning
Event reconstruction	91.9%	Event scanning
Living-dead cuts	73.0%	On-Off comparison
WIMP-box acceptance	47.7%	Calculation
Energy-dependent efficiencies		
Effect	Threshold	Method of calculation
Hardware trigger	$< 1 \text{ keV}_{ee}$	Additional dataset Event scanning Pulser measurements
Software S2 criteria	$< 1 \text{ keV}_{ee}$	Calculation Scatter plots
Software S1 3-fold requirement	1.7 keV_{ee}	Calculation Comparison with 2-fold

XENON10 data [174], where a similar, but smaller, discrepancy was found, by carrying out a maximum likelihood fit matching of the data to the calibration (with known efficiencies applied) by varying the shape of the energy conversion factor. The resulting conversion factor exhibits a decrease at low energy relative to the latest measurements of relative scintillation efficiency [140]. The same method of varying L_{eff} was utilised for the case of ZEPLIN-III, with the fit procedure effectively varying the product $s_n \times L_{eff}$ as we do not know how either varies with field at the relevant energies. The shape of the energy conversion factor was allowed to vary in the maximum likelihood fit, with the known efficiencies applied and some constraints (to provide realistic outcomes), yielding the conversion factor shown in Figure 5.37. Application of the variable energy conversion factor to the calculated efficiencies, produces the efficiency curve shown in Figure 5.37. The decrease in the conversion factor at low energies is more dramatic than that observed in XENON10, but note the 6 times stronger operating field in ZEPLIN-III.

It was decided that the varying energy conversion method would be applied as the favoured approach, but with the previously adopted constant conversion also considered, although inconsistent with our data.

5.11 WIMP search data

The science run lasted for 83 days, between 27th February and 20th May 2008, with an 84% livetime. This gave a raw exposure of 847 kg·days from the 12 kg active target, resulting in 126.7 kg·days effective exposure once the WIMP acceptance region and all energy-independent efficiencies and fiducialisation had been applied. Over this period the detector remained stable, with all operational parameters monitored and corrections applied where necessary.

As described previously, 90% of the dataset had been kept unanalysed with the

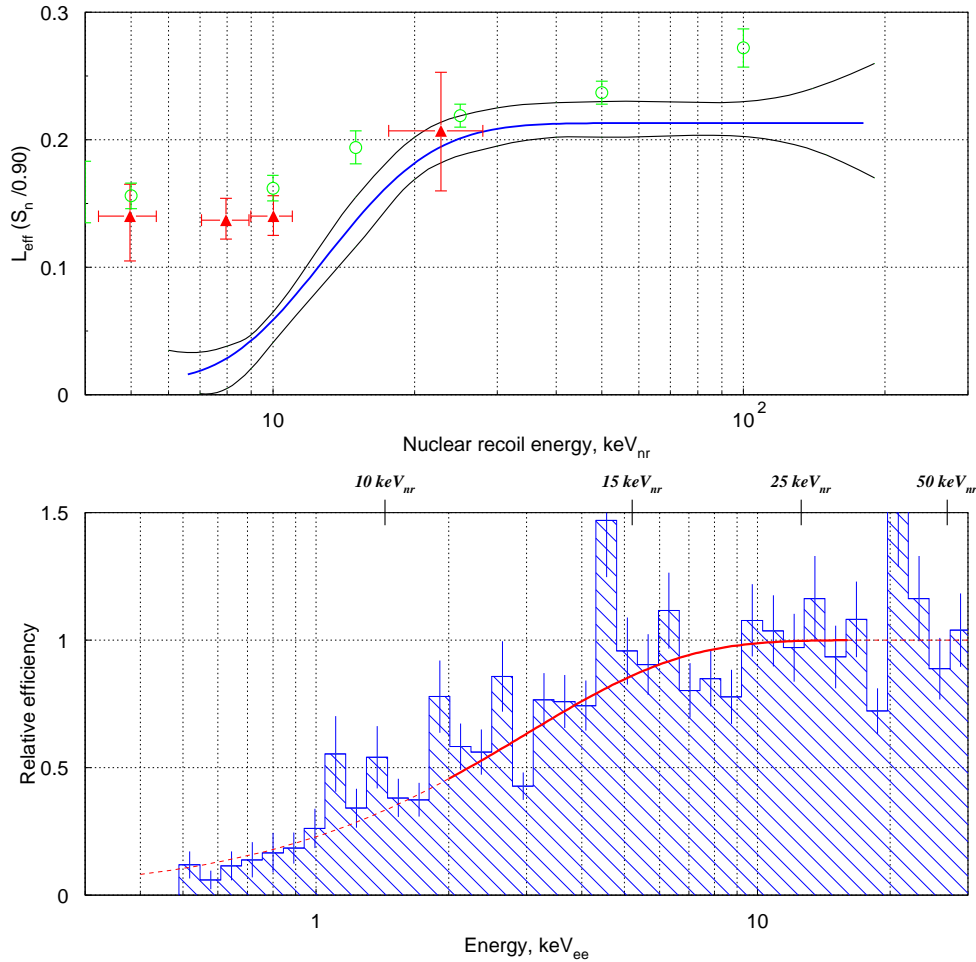


Figure 5.37: Top: Energy conversion factor (combined L_{eff} and S_n) (top) and efficiencies (bottom) derived using varying energy conversion. The red triangles show recent zero-field measurements from [140] and green circles denote the L_{eff} measured from the XENON10 data [145]. Bottom: Efficiency curve constructed from the calculated efficiencies and the variable energy conversion factor shown above.

intention of producing a blind analysis, with the unblinded 10% sample used to optimise the data analysis, event selection cuts and help define the WIMP search region.

The intended blind procedure would have seen the fixing of all parameters and cuts before analysis of the remaining data. This intention was not followed through for two reasons. Firstly, a weakness in the Mercury reconstruction code was discovered, discussed in Section 5.4, and corrected after the data had passed through the analysis once. Secondly, the electron recoil calibration (with ^{137}Cs) did not well represent the tails of the background distributions, making it impossible to characterise the electron recoil population and calculate the expected background in the WIMP search region, with the calibration data containing a higher event rate in the tails of the electron recoil population. Consequently, a non-blind analysis procedure was followed utilising the full 100% of science data.

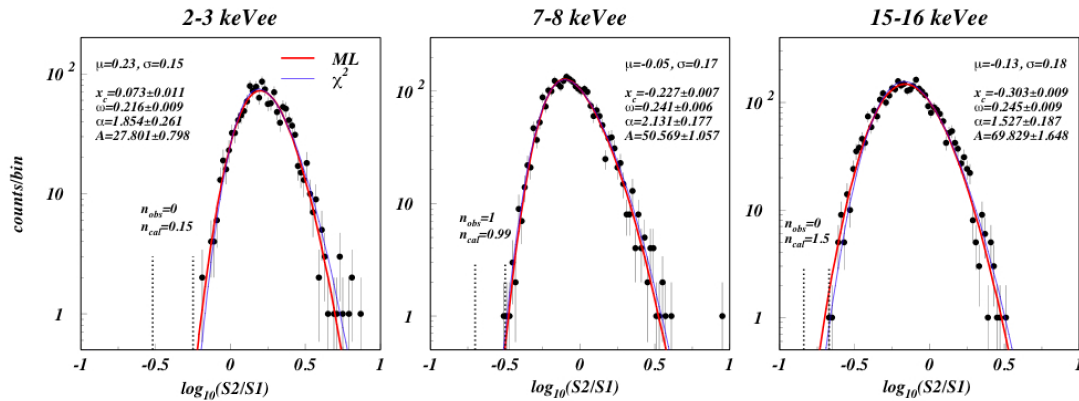


Figure 5.38: Examples of energy-sliced S2/S1 distributions from the WIMP search dataset. The dotted vertical lines define the bounds of the nuclear recoil acceptance region in each bin.

After correction of the reconstructed energy parameter and associated adjustments of the MSSJ rejection cuts, the electron recoil population was fitted with skew-Gaussian distributions, shown in Figure 5.38, using the same method applied to the ^{137}Cs calibration. However, here the fits were made outside the WIMP search box and projected into the search region for calculation of the expected number of background events, yielding a prediction of 11.6 ± 3.0 events in the box. The μ and σ of the background electron-recoils show excellent agreement with those of the ^{137}Cs calibration, despite the excess tail events in the calibration (Figure 5.39).

The final discrimination parameter scatter plot from the first science run is shown in




Figure 5.39: Comparison of mean and sigma of electron recoil populations from ^{137}Cs calibration (red) and WIMP search data (white).

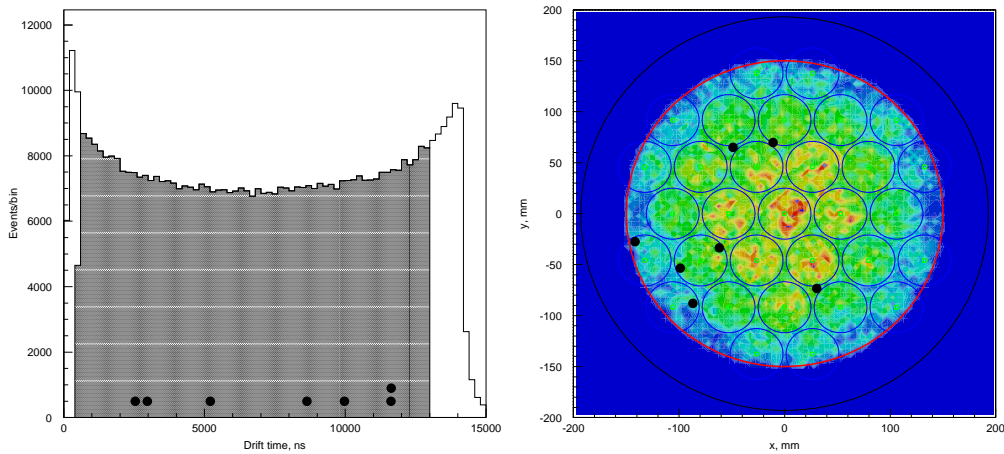


Figure 5.40: Drift time (left) and xy (right) distributions of the full ZEPLIN-III science dataset, with events in the region of interest highlighted.

Figure 5.41, with 7 events falling in the WIMP search region. The location of these seven events along the upper edge of the WIMP search box suggests that they are likely to be the result of leakage from the electron recoil population. The spatial event distributions from the dataset are shown in Figure 5.40, with the locations of the 7 events highlighted. Four of the events are located in the lower left quadrant of the xy space, the section of the plane with the smallest gas gap. This would suggest the need for a tilt correction, which is now underway, although this will not significantly change the final result.

The scatter plot of the full background dataset is the last stage in the data analysis, with the final result of the dark matter search being a limit on the possible WIMP properties. This final stage in the process is described in Chapter 6, where statistical techniques are applied to the observation incorporating the forms of the background and signal populations, detector efficiencies and energy conversion (Figure 5.37).




Figure 5.41: Final discrimination scatter plot of the full ZEPLIN-III science dataset, highlighting the 7 events observed in the WIMP search region.

Chapter 6

Calculation of limits on the WIMP-nucleon scattering cross-section

The commissioning and operation of dark matter detectors and subsequent data analysis provide high-quality datasets to search for a WIMP signal. To arrive at a final dark matter result there is one final stage in the process. The resulting observations must be analysed statistically for the presence of a signal and an associated scattering cross-section calculated.

There are different statistical approaches available for the analysis of the observations, each having its own strengths and weaknesses. In this section will be presented a review of two techniques, the Feldman-Cousins method (F-C) [175] and the Yellin interval methods [176, 177], used previously by LXe WIMP search experiments. In addition, we describe the technique applied to the ZEPLIN-III experiment, the Binned Maximum Likelihood (BML) method, which hopes to improve the quality of the resulting confidence bounds by incorporating as much information as possible about the observed events. This technique will be of great importance when WIMP signatures are finally observed. I have developed the BML from an initial generic implementation into a more detailed method incorporating additional event information.

In the analysis described here, the signal is defined as the number of nuclear recoils

in the data, resulting from both WIMP interactions or neutron background. In the experiments in question, the neutron background has been assessed with simulations and calculated as ~ 1 event per dataset, although significant error (typically ± 0.5) allows this to be consistent with 0 (a mean background of one event is deliberately chosen, determining the desired exposure for each data set). As a result, one should bear in mind that the signal arising from this analysis can be contaminated by a small recoil background.

The second part of the process is to convert the resulting 90% confidence limits on the observed event rate into limits on the spin-independent WIMP-nucleon elastic scattering cross-section. This involves comparison of the observed energy spectrum with the recoil spectrum expected for a given WIMP mass and galactic halo model. This process has been well defined previously, but requires careful application to each given detector and dataset. Whilst spin-independent interactions are expected to be observed first, spin-dependent cross-sections are also calculated from the science data, but these are discussed elsewhere [171].

6.1 90% confidence limits on observed event rate

As detailed in Chapters 4 and 5, one of the main data analysis outputs from the ZEPLIN experiments is a plot of the events observed in the science dataset in 2D parameter space of the main discrimination parameter against observed energy. In order to extract a statistical statement about the probability of having found a WIMP event we define a region of interest in this parameter space. We choose a region of little or no contamination from background (electron-recoil) events, whilst still accounting for a significant portion of the predicted signal probability. For the two-phase noble gas detectors run to date, this region has included the lower 50% of the NR band (from neutron calibrations) in an energy range below 20 keVee. The process of setting and calculating this box was described for ZEPLIN-II and III in the relevant analysis chapters.

Once such a region has been defined, we can statistically analyse the observation within it to determine the 90% confidence boundaries on the magnitude of any observed signal. When doing so we need to take into account the electron recoil background we expect to see in the region and the number of events we actually observe. In the BML

and Yellin methods extra information about events is included when calculating the probability that they are drawn from a signal population: in the case of the Yellin method, the energy of the events is accounted for; with the BML method both the energy and $\log_{10}(S2/S1)$ are utilised to help evaluate the observations.

6.1.1 Feldman-Cousins method

The method laid out by Gary Feldman and Robert Cousins [175] was a development of the classical Neymann process for constructing confidence intervals [178]. They introduced a new system of ordering, known as likelihood-ratio ordering, which solved many of the original criticisms levelled against the classical method. The original process struggled in two particular cases, one of which is the exact situation we face in dark matter searches: a Poisson process in the presence of a known background. A major problem of the classical method occurs when used improperly, in a situation known as “flip-flopping” [179].

If a researcher were to look at his/her data *before* deciding whether to set a two-sided 90% confidence interval *or* a pure 90% upper limit, then, after revealing the data, the classical limit set would suffer from under-coverage. The resulting confidence interval would include less than 90% confidence, a serious flaw in the result. The Feldman-Cousins method solves this problem of flip-flopping by guaranteeing coverage. At small observed event numbers, n (compared with the known background), the limits become pure upper bounds, giving a contiguous set of confidence intervals, removing the danger of flip-flopping. Their ordering method also removes empty sets which are also an issue in certain situations (which we need not discuss here).

To calculate a 90% confidence interval using the F-C method, we must first construct a Neymann diagram (using the likelihood ratio ordering) incorporating our known background. From this we can extract confidence boundaries for any actual observation. Figure 6.1 is an example of a Neymann diagram for a known background of 1 event. A Neymann diagram is constructed horizontally by calculating the 90% probability interval of possible observations for each possible value of the signal mean, μ . We calculate the probability of observing n events, given the true signal, μ , and the known




Figure 6.1: A Neymann construction for 90% confidence limits of a Poisson process with a known background, utilising the Feldmann-Cousins ordering principle.

background, b , with:

$$P(n|\mu) = \frac{(\mu + b)^n e^{-(\mu+b)}}{n!}, \quad (6.1)$$

then we set μ_{best} to the (physically allowed) mean which maximises $P(n|\mu)$:

$$\mu_{best} = \begin{cases} (n - b) & \text{if } n \geq b \\ 0 & \text{if } n < b \end{cases}$$

The ratio of these two probabilities,

$$R = \frac{P(n|\mu)}{P(n|\mu_{best})}, \quad (6.2)$$

becomes the parameter applied to order the possible values in constructing the 90% probability intervals. The values of n are added to the interval in order of decreasing R , until the probability contained in the interval exceeds the required level (in our case 90%). The process is then repeated for each possible signal level.

Having constructed the diagram, by calculating the confidence intervals for many different values of μ , the 90% confidence limits can be read for any number of observed events. The confidence interval on the observation is read vertically, as shown by the solid red line in Figure 6.1. In the example shown, a known background of 1.0 with an observation of 5 would result in a two-sided 90% confidence interval between 1.2 and 9.0.

Coverage was mentioned as a consideration when constructing confidence intervals. Ideally, each interval would contain exactly 90% of the probability, but this is not possible in the case of a Poisson process due to the discrete nature of the possible outcomes. We therefore suffer from over-coverage of our intervals which introduces a level of conservatism to the limits produced, meaning that the resulting confidence levels are slightly higher than 90%.

Another weakness in the use of the F-C method for our search experiments is that it utilises no information about the individual events observed. The only inputs into the calculation of the hypothetical recoil signal are the expected number of background events and the experimentally observed number of events; these are summed over the entire region of interest. In the general case of WIMP search experiments, this is a shortcoming, the impact of which depends on the distribution of the observed data with relation to the background (and signal) prediction. In the case where the observed events fit well with the predicted background, the resulting confidence interval will be overly conservative.

ZEPLIN-II: Feldman-Cousins analysis

The ZEPLIN-II data had two background populations, electron recoils from γ -rays and ^{222}Rn -progeny “wall events”, combining to give a relatively uniform distribution of events throughout the region of interest, designated as the lower 50% of the nuclear recoil band in the energy range 5–20 keV_{ee}. In order to characterise the backgrounds the region was split into two energy slices, 5–10 and 10–20 keV_{ee}, which accounted, to some extent, for variations in the behaviour of the populations with energy. The γ -ray population from the ^{60}Co calibration was fitted with a log-normal distribution (cross-checked by also fitting the electron recoil background population above 50% recoil); these

curves were integrated in the region of interest to provide an expected background from γ -rays. As described in Section 4.8, the radon-progeny events are gaussian in the radial coordinate (due to poor radial reconstruction resolution). Their radial distribution was fitted outside the fiducial radius, and projected into the region of interest to calculate an expectation in the box. Table 6.1 summarises the expectation values for the two backgrounds in each of the two energy slices; the γ -ray expectations from calibration and data agree within errors. These four values (calibration and radon-progeny) are combined to provide an overall expectation for the number of events in the combined region; the errors shown are statistical only, as derived from the fits.

Table 6.1: Expectation values from the ZEPLIN-II science data.

Energy Range	Observed	γ -ray (^{60}Co) (1)	γ -ray (data)	Rn-progeny (2)	Total (1+2)
5-10 keV _{ee}	14	4.2 ± 2.4	5.6 ± 4.6	10.2 ± 2.2	14.4 ± 3.3
10-20 keV _{ee}	15	11.9 ± 2.7	13.0 ± 6.0	2.3 ± 0.5	14.2 ± 2.7
Total					28.6 ± 4.3




Figure 6.2: Neymann diagram using the Feldman-Cousins ordering principle for the ZEPLIN-II observation of 29 events with 28.6 expected.

Once the expectation in the box was calculated the F-C approach can be implemented. ZEPLIN-II had an observation of 29 events, with a background expectation of 28.6 ± 4.3 events. These values are rather large, in fact above those tabulated by Feldman and Cousins in [175], but can be calculated with the F-C procedure, leading to the result shown in Figure 6.2. This results in bounds for a 90% confidence interval of 0.0 and 10.4 events, giving an upper limit on rate of 0.092 evts/kg/day and showing the observation to be consistent with a null result, in the 50% recoil acceptance region over 225 kg-days. The FC values can also be calculated using the *TFeldmanCousins* class within the ROOT analysis framework [180].

6.1.2 Yellin interval methods

The interval methods developed by S. Yellin were designed as procedures for setting upper limits in experiments where there is potential for unknown background populations [176]. They seek to set an upper limit on the hypothetical signal by assessing the observed event distribution for the separation between events in a certain parameter, e.g. energy. These gaps can include either no events (the simple Maximum Gap method) or contain n events (the Optimal Interval method).

The shape of the WIMP recoil spectrum can be calculated for a given WIMP mass, with the absolute rate dependent on the scattering cross-section. Using the shape of the expected signal spectrum, the energy scale can then be transformed giving a uniform expectation (dR/dE_r) in the energy scale (as shown in Figure 6.3). Now by finding the largest gap between two events (or an event and the bounds of the search window), a limit can be placed on the possible signal. This is done by increasing the cross-section until the expected number of events in the interval becomes inconsistent with the observation at the 90% level. Utilising the largest empty interval between events is known as the Maximum Gap method.

However, the most powerful limit will not necessarily be set for an interval containing no events. There is a balance between the number of events included and the size of the interval. Finding the interval which provides the best limit is known as the Optimal Interval method. A full description of both procedures are laid out in [176]. There are certain limitations to consider with the original Yellin methods: a limit of 54 events in the experimental spectrum and restriction to only one dimension (i.e. only energy




Figure 6.3: Example of the transformation of energy scale to provide constant dR/dE_r , with maximum gap and optimal interval shown [181].

spectrum considered). These problems have been addressed and an extension to higher statistics is presented in [177], along with some discussion about extending the method to multiple dimensions.

XENON10: Yellin maximum gap analysis

XENON10 is a competing WIMP search experiment also utilising two-phase xenon [104]. It operated in the Gran Sasso National Laboratory, Italy, with a 15 kg liquid xenon target, viewed from above *and* below by two arrays of PMTs. The detector collected 58.6 live days of science data, with the analysis and results described in [104, 174].

Figure 6.4 shows the final discrimination scatter plot of the XENON10 data. Their WIMP search region was chosen as the lower half the nuclear recoil band between 4.5 and 26.9 keV_{nr}. The region chosen did not account for the full lower 50% of nuclear recoil acceptance, having a lower band (defined as $\mu - 3\sigma$ in each energy slice) which excludes an additional event. In the region ten events are observed, five of which (numbers 3,4,5,7 and 9 in Figure 6.4) are stated to be statistically consistent with leakage from the electron recoil band (seven were predicted). The other events are discounted as nuclear recoils *a posteriori*, attributed to a noise glitch (no. 1) and the weakness of an anomalous hit pattern cut (2,6,8 and 10), used to reject double Compton-scattering (MSSI) events. However, there is no statement about the confidence of the null result, other than stating that the events not consistent with background were “not favoured as evidence for WIMPs.”




Figure 6.4: Final discrimination scatter plot from the XENON10 experiment, plotting $\Delta \log_{10}(S2/S1)$ against nuclear recoil energy (calibrated with field suppression factors and a constant L_{eff} of 0.19). The WIMP search region is defined between the blue lines (lower half of nuclear recoil band) and between 4.5 and 26.9 keV_{nr} [104].

To calculate a limit from the data, the Yellin Maximum Gap method was used on all 10 observed events, with the large and very obvious gap between events 1 and 2 producing an excellent upper limit on the signal. The decision to use a Yellin analysis was made after the null result had been reached. Consequently, flip-flopping occurs resulting in reduced coverage of the calculated limit. The confidence level of the published limit is believed to be in the region 80-85% and as such the true 90% confidence limit would be greater than that published.

This use of the Yellin method in a WIMP search experiment also raises the question of which statistical methods are applicable to a *search* experiment. To avoid flip-flopping, the use of a Maximum Gap method should be stated before the unblinding of the data. If this is the case, however, an admission is implicit that the aim of the experiment is not to find a signal, just to place an upper limit on the presence of one.

The event distribution in the XENON10 search region suggests the Feldman-Cousins method as a more appropriate approach, avoiding flip-flopping and providing a statement of confidence in the null result. From the 7 events expected and 10 observed, the F-C tables give a two-sided 90% confidence interval between 0.00 and 9.50, showing the

result consistent with zero and giving an upper limit on the rate of ~ 0.07 events/kg/day. This results in a limit on the WIMP-nucleon spin-independent scattering cross-section of $\sim 1 \times 10^{-7}$ pb, significantly higher than the 4.5×10^{-8} pb published. On the other hand, if the anomalous events were removed and a F-C limit derived from 5 events observed with 5 expected, the upper limit would be reduced to $\sim 7 \times 10^{-8}$ pb.

6.1.3 Binned Maximum Likelihood method

As WIMP search experiments begin to probe further into the parameter space of interest, giving them real discovery potential, it is important that we utilise as much information as possible about events we observe, to provide the most informative verdict on the possibility of a WIMP signal. To achieve this goal a more powerful technique was clearly required. Whilst the Yellin method does use the energy of the events in its choice of the optimal gap, it does not take all observed events into the calculation of the result and can only place upper limits on a signal (as any actual event is treated as a potential signal). The implementation of a binned maximum likelihood method allows for as much information as possible to be utilised. Naturally, for BML to be used a good knowledge of all possible backgrounds is required, with their distributions well characterised. As a result of these characteristics the BML has discovery potential, whereas Yellin does not.

In this section we begin by discussing the concept of likelihood, before explaining a simple implementation utilising only information about the S2/S1 discrimination parameter of the events (the 1-D implementation); then we expand the method to two dimensions by incorporating the energy of events into the calculations. This method is expandable into any number of dimensions (limited only by processing power), incorporating any other discriminatory factors, e.g. pulse shape or event location.

The likelihood, \mathcal{L} , is calculated as the probability that the observation of a given set of data arose from a theorised, parent distribution. The probability of observing a certain number of events in a single bin from a predicted probability distribution function can be calculated; multiplying together the probabilities for all bins gives the likelihood:

$$\mathcal{L} = \prod_{i=1}^k P(n_i | \mu(x)). \quad (6.3)$$

However, it is often more useful to calculate the log-likelihood value,

$$\ln \mathcal{L} = \sum_{i=1}^k \ln P(n_i | \mu(x)), \quad (6.4)$$

simplifying many calculations with the product becoming a sum. This is particularly useful when probabilities are small and/or many bins are involved in the calculation, which is the case for the analysis described here.

In the case of binned maximum likelihood, the probability function for the number of counts in each bin will usually follow poissonian or gaussian statistics. The number of events expected in each bin is calculated from the parent distribution, describing the populations over the wider parameter space. In our case, bins will have a small number of counts (mostly zeros), so Poisson statistics will prevail within each bin:

$$\ln \mathcal{L} = \sum_{i=1}^k \ln \left(\frac{e^{-\mu_i} \mu_i^{n_i}}{n_i!} \right). \quad (6.5)$$

An event in our acceptance region could have arisen from the scatter of a WIMP (or neutron) or from the electron recoil background, so both of these possibilities will contribute to the probability of observing an event in each bin. Therefore the calculation of the mean for each individual bin becomes:

$$\mu(x) = S \cdot P_s(x) + B \cdot P_b(x), \quad (6.6)$$

where x is the discrimination parameter (e.g. $x = \log_{10}(S2/S1)$), $P_s(x)$ and $P_b(x)$ are the probability density functions (normalised to 1) for the signal and background populations, and S and B are *potential* variables in the maximum likelihood test, being the hypothetical number of signal and background events observed in the region tested.

In the standard parameter space chosen for discrimination scatter plots, $\log_{10}(S2/S1)$ vs energy, the signal population is described by a gaussian in each energy bin, with the background best described by a skew-gaussian (in ZEPLIN-III). In order to simplify the problem, we convert the standard parameter space to provide constant nuclear recoil acceptance percentage on the y -axis, A_{nr} . This involves converting the $\log_{10}(S2/S1)$

values into a parameter calculated from the behaviour of the normally-distributed AmBe neutron calibration of the signal response.




Figure 6.5: Final discrimination parameter scatter plot of the ZEPLIN-III first science run in the traditional coordinate space (left) and the NR % space (right).

Figure 6.5 shows the results of the ZEPLIN-III first science run (and colour contours from the AmBe calibrations) in the familiar scatter plot coordinate space and the same data represented in constant recoil acceptance space. The transformation requires inputs of the mean and sigma of the nuclear recoil population as a function of energy (see Chapter 5); the nuclear recoil acceptance fraction becomes:

$$A_{nr} = 1 - \frac{1}{2} \operatorname{erfc} \left(\frac{1}{\sqrt{2}} \frac{y_s - \mu(E)}{\sigma(E)} \right), \quad (6.7)$$

where erfc is the complementary error function.

The background distributions have been defined by fits to the data and calibrations, with errors associated with the resulting fit parameters. The region in which the background is to be defined is on the tail of the overall electron recoil population, where variation of the fit parameters, within errors, will have a significant effect on the number of background events predicted by extrapolation. When calculating limits on the presence of a signal, an over-estimate of the background will suppress the final result. Looking back at F-C, predicting and observing 1 event results in a 90% confidence interval of 0.0 to 3.36 events, but if the background prediction had been 3 events, the interval

would have been 0.0 to 2.19, over-constraining the signal. If the form of the background in the region is known with high precision (with a small error in the number of events predicted), this suppression is acceptable and the fixed background can be used. But in the case where the number of events predicted is affected by significant systematic or statistical uncertainty, the background can (and should) be allowed to vary as a parameter in the maximum likelihood analysis, as long as it remains within the error bounds on the original fit parameters. The parameter chosen to vary the background can be any one of the fit parameters. For the description of the procedure, the amplitude is allowed to vary as it directly affects the number of background events predicted (B).

As a result, to analyse the data we test the observation against different hypothetical parent populations by varying both S and B in Eq. (6.6). Iterating over many possible populations, we find the combination of S and B which maximises the log-likelihood value, giving the most probable number of signal and background events observed in the region of interest. For a dark matter result, a 90% confidence upper limit (or 90% confidence bounds) needs to be set on the number of signal events. The method for varying the number of background events depends on the situation, it may be more appropriate to vary other parameters describing the background, rather than the amplitude.

Once the most probable combined parent distribution for the observed events has been found, the background is fixed, i.e. we accept the best estimator, B_{max} , from the initial maximum likelihood analysis. We then examine the signal parameter and set a limit on the number of recoil events observed. There are standard procedures for calculating the confidence limits from likelihood analyses (e.g. the likelihood ratio test). In our case, the best estimator of the signal is at (or very near) the boundary of the physical conditions (i.e. we cannot observe negative signal) and we are in a low statistics regime, so the most robust way of calculating confidence intervals is to carry out a set of Monte Carlo experiments and assess the resulting likelihood distributions.

One-Dimensional Implementation

The first step in developing a maximum likelihood analysis of the data was to simplify the problem to one dimension. The region of interest, 2-16 keV_{ee} and below 50% nuclear recoil acceptance, was divided into 49 bins (1-50%) each containing 1% of NR

acceptance, as shown in Figure 6.6. Exploiting some of the discrimination power held in the distribution of the events, it can be seen from Figure 6.6 that all the events in the region of interest occur in the top 8% (between 42% and 50% nuclear recoil acceptance).




Figure 6.6: Observed events from the ZEPLIN-III first science run in the WIMP search region (2-16 keV_{ee}) in the nuclear recoil acceptance parameter and example probability density functions associated with background and signal.

Figure 6.6 includes an example of the probability density functions (normalised to 1) which represent the signal and background populations. The red line shows the signal p.d.f. with constant probability (defined by the choice of coordinate space) and the background p.d.f. (blue) which is a transformation from a skew-gaussian in the traditional coordinate system. From the p.d.f.s and the choice of B and S , μ is calculated for each bin. Figure 6.7 shows possible combined μ distributions for different choices of B and S , examples of the distributions which are used to calculate the likelihood values.

To find the most probable overall parent distribution for the observation, we carry out a maximum likelihood analysis, using Eq. (6.5), varying S and B and finding the maximum log-likelihood value, $\ln \mathcal{L}_{max}$. This yields the best estimators for the number of signal events, S_{best} , and likewise for the background, B_{best} (and hence μ_{best} for each bin). From here we need to extract the relevant confidence intervals, done by simulating a large number of experiments (in our case 10,000), for different values of the signal,

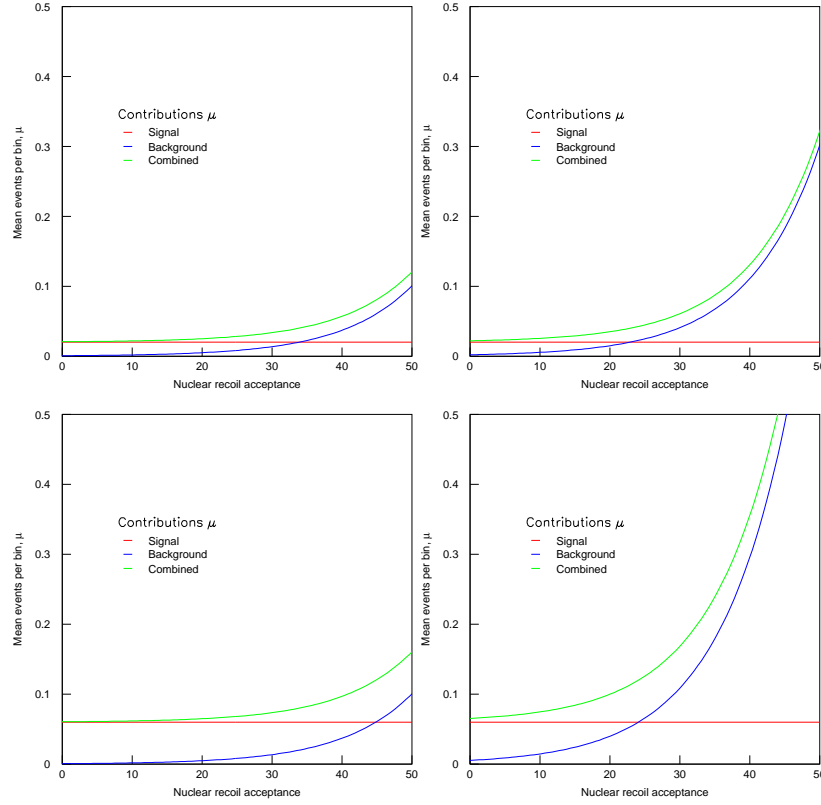


Figure 6.7: Example of combined μ distributions for different levels of signal and background. $S=1$, $B=1$ (top-left), $S=1$, $B=3$ (top-right), $S=3$, $B=1$ (bottom-left), $S=3$, $B=8$ (bottom-right).

S , and holding B constant. To generate an experiment, μ_i can be calculated for each bin and used to throw a random number (from a Poisson distribution with mean μ_i) of counts in that bin. Doing this for all the bins produces a MC experiment, where the mean number of events per experiment is $(S + B)$.

Each experiment thus generated can be tested against the ‘true’ population distributions (using μ_{max}) giving a certain value of $\ln \mathcal{L}$. From the large number of experiments, a $\ln \mathcal{L}$ distribution can be constructed for each tested signal. The 90% confidence upper limit can be set by finding the value of S above which 90% of the experiments are incompatible with the observation, i.e. those with $\ln \mathcal{L} < \ln \mathcal{L}_{max}$. Figure 6.8 shows the cumulative $\ln \mathcal{L}$ distributions for different levels of S . By iterating the signal value, S can be found where the 90% value matches $\ln \mathcal{L}_{max}$.

Limits calculated with such Monte Carlo simulations are truly frequentist in nature, providing full coverage. The method also allows for the setting of two-sided boundaries




Figure 6.8: Cumulative $\ln \mathcal{L}$ distributions from simulated experiments varying S between 2.0 and 4.0 in steps of 0.1. The red line shows the intersection of the cumulative 90% with $\ln \mathcal{L}_{best}$.




Figure 6.9: $\ln \mathcal{L}$ distribution from MC simulation for $S = 0$ (when $S_{best} = 0$), showing the measured $\ln \mathcal{L}$ (red line) and central $1\text{-}\sigma$ bounds (shaded).

in the presence of a non-zero signal. This would simply be done by decreasing the signal (from S_{best}) and simulating experiments. The significance of the null result ($S=0$) can be assessed by simulating experiments with $S=0$ and assessing the position of the observed likelihood in the resulting likelihood distribution. This is also a potential method of deciding whether a one- or two-sided limit should be generated. Figure 6.9 shows the $\ln \mathcal{L}$ distribution produced by MC simulations for $S = 0$ (when $S_{\text{best}} = 0$), showing the measured $\ln \mathcal{L}$ well within the main distribution. If it were located to one extreme it would show that the background may be artificially suppressing the limit. The other extreme would show that a non-zero signal was close and suggest a two-sided limit should be set.

Two-dimensional extension

The one-dimensional method just described is a powerful way of analysing event distributions, but it can be further improved by adding additional information to the analysis. In the case of our experiments the obvious information to introduce is the event energy, E . This is beneficial as the background and signal separation (discrimination power) vary with energy, as does the expected signal distribution. We now divide our region of interest into 1% bins in the NR % space (as before) and additionally into 1 keV $_{ee}$ bins in energy from 2 to 16 keV $_{ee}$. This divides our region into a grid of 686 cells, where each individual cell will have different probabilities of seeing signal and background.

The probability density functions for signal and background now become surfaces in the two dimensional analysis, as illustrated in Figure 6.10. The background probability density functions are calculated from the background data outside the region of interest. The data is sliced into 1 keV $_{ee}$ bins and fitted (in the case of ZEPLIN-III, with a skew-gaussian distribution). These fits are extrapolated into the region of interest and used to calculate the probability density function.

The calculation of signal distributions is more elaborate than in one dimension. It must account for the theoretical WIMP recoil spectrum, which varies with WIMP mass. The signal p.d.f. is based on the exponentially-falling ‘ideal’ recoil spectrum convolved with the detector energy resolution and detector efficiencies, to produce a distribution of the form shown in Figure 6.10. The detailed calculation of the WIMP recoil spectrum

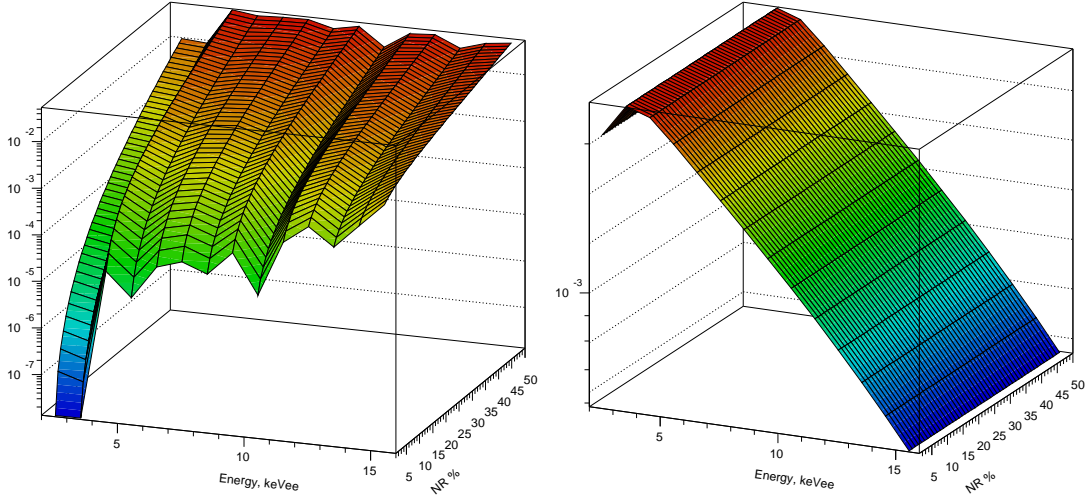


Figure 6.10: Probability surfaces for background (left), derived from the skew-gaussian fitting, and signal (right), derived for a WIMP mass of $50 \text{ GeV}/c^2$ incorporating the detector energy resolution and efficiencies. The jagged nature of the background results from statistical variation between the fits to the sliced data.

is covered in Section 6.2. We assume, with no major error, that a neutron background population would be similarly distributed.

Once the p.d.f.s have been calculated for signal and background, μ for each cell is calculated for the different values of S and B , in the same way as for the one dimensional case. The $\ln \mathcal{L}$ can be calculated as previously with Eq. (6.5), but now we must iterate over all 686 cells. The procedure thus far has varied the amplitude of the background p.d.f. but this does not have to be the case. Any parameter used to describe the background distribution can be allowed to vary, as long as it leaves the resulting background distribution compatible with the overall background population beyond the region of interest. In the case of ZEPLIN-III, varying the amplitude would be incompatible, but varying the width of the population within the error bounds (from the initial whole-distribution skew-gaussian fits) worked well, i.e. the final parameter value was still within errors of the original fit value.

The process of extending the BML method from one dimension to two shows that the only potentially challenging aspect is to characterise the populations in any new parameters which we have chosen to introduce. In the case of the background distribution this is relatively simple, utilising the science data beyond the region of interest and associated electron recoil calibrations. In the case of the signal, the nuclear recoil

calibrations and theoretical WIMP properties need to be considered. As mentioned previously, other potential discrimination parameters could be, in the case of WIMP searches, scintillation pulse shape or event location within the detector.

Another possible extension of the implementation could be to extract directly the scattering cross-sections as a result of the analysis. However, by calculating the 90% confidence limit on rate and then proceeding to calculate the cross-sections separately leaves the process more transparent.

The BML analyses presented here were coded in FORTRAN; carrying out the conversion between the different discrimination ($\log_{10}(S2/S1)$ and y_{nr}) and energy (E_{ee} and E_{nr}) parameters, iterating finely over values of S and B and allowing for repetition of the analysis for different WIMP masses, M_D . The observed number of events per bin/cell and efficiencies are read from external files. The code could easily be extended to include other discrimination variables and the calculation of WIMP-nucleon scattering cross-section could be incorporated.

ZEPLIN-III: Two-dimensional binned maximum likelihood

The excellent discrimination achieved in ZEPLIN-III meant that the distribution of events within the region of interest was not expected to be uniform (for small/no signal); given that there would be statistical power held in the distribution of the events, it was decided to design and implement an analysis of the data using BML.

The ZEPLIN-III background was characterised by slicing and fitting the data (in 1 keV_{ee} bins) as described previously. The most robust fitting method was produced using a skew-gaussian function, which produced high-quality fits to the whole γ -ray distribution in all slices (Figure 5.38). It should be noted that there is as yet no theoretical motivation behind this particular function, but the quality of the fits (for all bins in the science data, as well as the ^{137}Cs) is very high. The fitting to the background is carried out beyond the region of interest, meaning the fits must be extrapolated to predict the background. Ideally enough calibration data could be acquired for the fitting to directly characterise the background within the region. To adequately fit the background within the region would require a calibration dataset with ten times more events than the science data. However, as the maximum possible rate for these calibrations was

only about ten times the background data-rate, the calibration run would need to last as long as the science run itself (83 days), which is impractical.

The skew-gaussian function is defined by:

$$P(x) = \frac{2}{\omega\sqrt{2\pi}} e^{-\frac{(x-x_c)^2}{2\omega^2}} \int_{-\infty}^{\alpha \frac{x-x_c}{\omega}} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt, \quad (6.8)$$

which can be written in the form:

$$P(x) = \frac{1}{\omega\sqrt{2\pi}} e^{-\frac{(x-x_c)^2}{2\omega^2}} \left[1 + \text{erf} \left(\frac{\alpha}{\sqrt{2}} \frac{x-x_c}{\omega} \right) \right], \quad (6.9)$$

where the parameters are as defined in Section 5.9.2.

$P(x)$ is multiplied by an amplitude parameter, A , to fit the energy slices. Fitting to the data (rather than calibration) leaves open the possibility of biasing the fits by the presence of a signal, although in the case where a fit is made to the complete distribution, the influence of a small signal will not be significant. Table ?? shows the parameters obtained from the fits to the 14 slices of interest, with the contribution to each cell calculated using the transformation into the NR % space (Eq. 6.7). Extrapolating the fitted distributions results in a prediction for the expected number of background events from the electron recoil population of 11.6 ± 3.0 in the region of interest, compared with the 7 observed. If our description of the background is to be trusted, our result entails a significant under-fluctuation of $\sim 2\sigma$ from the predicted mean. However, we point out that the functional dependence found to characterise the electron recoil population is not yet theoretically motivated, and therefore we must consider the possibility that there is a systematic uncertainty in our description of its tails far from the mean. This uncertainty must be treated conservatively.

Were we to accept the prediction of 11.6 events, without allowing for variation within errors, this would over-constrain the signal; therefore we must vary one of the parameters describing the background (within the error bounds of the fits). Allowing the amplitude alone to vary would make the fit to the observed events in the region incompatible with the background outside (i.e. the resulting best estimator for the amplitude parameter would be outside the allowed error bounds). But varying the width parameter results

Table 6.2: Skew-Gaussian fit parameters from slicing electron background data.

E (keV _{ee})	x_c	ω	α	A
2-3	0.073 ± 0.011	0.216 ± 0.009	1.31 ± 0.18	27.8 ± 0.8
3-4	-0.033 ± 0.006	0.214 ± 0.006	1.48 ± 0.13	35.1 ± 0.7
4-5	-0.086 ± 0.011	0.212 ± 0.009	1.16 ± 0.16	39.9 ± 1.1
5-6	-0.168 ± 0.006	0.233 ± 0.005	1.55 ± 0.11	44.8 ± 0.8
6-7	-0.188 ± 0.010	0.228 ± 0.008	1.38 ± 0.14	47.2 ± 1.3
7-8	-0.228 ± 0.007	0.241 ± 0.006	1.51 ± 0.13	50.6 ± 1.1
8-9	-0.257 ± 0.006	0.245 ± 0.007	1.59 ± 0.13	53.7 ± 1.0
9-10	-0.280 ± 0.008	0.248 ± 0.007	1.55 ± 0.15	55.3 ± 1.3
10-11	-0.302 ± 0.007	0.257 ± 0.006	1.77 ± 0.16	57.0 ± 1.3
11-12	-0.311 ± 0.008	0.256 ± 0.007	1.54 ± 0.13	59.6 ± 1.2
12-13	-0.312 ± 0.008	0.253 ± 0.007	1.39 ± 0.12	62.1 ± 1.2
13-14	-0.323 ± 0.007	0.264 ± 0.006	1.53 ± 0.13	65.6 ± 1.3
14-15	-0.330 ± 0.010	0.259 ± 0.008	1.38 ± 0.15	67.6 ± 1.6
15-16	-0.303 ± 0.012	0.245 ± 0.009	1.20 ± 0.13	69.8 ± 1.6

in ω remaining within the error bounds of the original skew-gaussian fits, therefore remaining compatible with the population as a whole. As a result, we define a new variable, β , as a parameter multiplying the width, ω , of the distributions. For each value of β , the amplitude of the distributions was adjusted to keep the total number of background events (in the whole distribution) constant.

The signal p.d.f. is calculated from the theoretical WIMP recoil spectrum. This varies with WIMP mass, so the whole process of limit setting must be iterated for different hypothetical masses. The spectrum is smeared with the detector resolution at low energies, $\sigma = 1.0\sqrt{E_{ee} [\text{keV}_{ee}]}$, and the detector efficiency function applied before normalising the sum of the bins to unity. Consequently, the probability in each bin is such that, if a signal event occurs in the region of interest, it will be observed in that bin. Calculating the p.d.f. for signal and background produces the probability surfaces shown in Figure 6.10.

The predicted mean observation in each bin is now calculated using $\mu = B(E, y_{nr}, \beta) + S \cdot P_s(E, A_{nr})$, where $B(E, A_{nr}, \beta)$ is the number of background events predicted in each cell from the skew-gaussian distributions for a given value of β , with the maximum likelihood analysis proceeding by iterating over β and S . The log-likelihood for each possible combination is calculated with Eq. 6.5, with the combination maximising $\ln \mathcal{L}$ providing the most probable parent distribution for the observation. Figure 6.11 is a colour map




Figure 6.11: Colour map of $\ln \mathcal{L}$ distribution, varying β and S . The likelihood increases from blue to red, with the best estimator point (i.e. maximum likelihood) marked with a point.

of the log-likelihood for different combinations of β and S . The best estimator for the number of signal events observed is $S_{best}=0$, indicating no evidence of a WIMP signal in the data, along with the background parameter $\beta_{best}=0.945$ (a reduction of 5.5% in the width of the original fit), producing an estimate of 7.13 background events in the region. $S_{best}=0$ suggests a null result from the data, so an upper confidence limit can be placed on the WIMP-nucleon scattering rate, given the observed background.

Once the best estimators for the ‘true’ parent distribution have been found, the 90% confidence limits can be extracted using the MC method previously described. The β parameter was fixed at 0.945 and the signal increased systematically (starting at 2.0), with 10,000 MC experiments generated at each signal level. The likelihood that the generated experiment arose from the ‘true’ parent distribution was calculated, giving a 90% confidence limit of 2.9 events (0.023 evts/kg/day) for a WIMP mass of 55 GeV. Varying the WIMP mass gives the expected variation in the limit on rate, from 2.45 for a 10 GeV WIMP to 3.0 for a 1000 GeV WIMP. The variation in limit slows as the

WIMP mass increases, resulting from the flatter recoil spectrum at higher masses. For comparison, note that the 90% confidence limit on zero events in the simple Poisson case is 2.3.

It was later understood that the binned maximum likelihood method described here is not as statistically rigorous as possible, and a more correct approach would be to apply a full profile likelihood methodology. Work is currently underway to develop the profile likelihood method for the second science run of ZEPLIN-III and future dark matter search experiments. This involves applying a slightly different test statistic, where the limit setting MC are generated using the best estimators from a constrained maximum likelihood fit, compared with the universal maximum likelihood parameter applied here. Subsequently, the profile likelihood is calculated for the MC experiments in the same way as for the observed data [182]. This change is in the MC limit setting stage and not in the initial signal test, as a result it is not believed that this method should significantly alter the final result. An additional limitation on the use of maximum likelihood methods with such data sets is the requirement for an accurate calibration of the backgrounds which has proved extremely difficult in dark matter experiments to date.

6.2 WIMP-nucleon scattering cross-sections

Thus far this chapter has described the statistical treatment used to calculate the 90% confidence interval on the signal rate observed in the experiment. But WIMP search experiments are interested in the cross-section for WIMP-nucleon scattering and most probable mass of WIMP particle. In the case of a positive observation we would be able to constrain these properties; in the case of a null result, we define upper limits. This section discusses the process of translating the 90% limit on the rate of recoil events into a limit on WIMP properties.

The calculated probability of a WIMP interaction with baryonic matter in a detector will clearly be dependent upon the type of WIMP and the model assumed for the galactic halo. In order to be able to compare direct detection experiments in a consistent manner, a “standard” galactic halo model is considered [92, 101]. This consists of an isothermal, cored dark matter halo, with a Maxwellian velocity distribution of WIMPs, giving a

flat rotation curve (i.e. $\rho \propto r^{-2}$) with a central region of constant density. These cored models result from parameterisations of the observed rotation curve, with the following parameters taken for calculations: characteristic velocity $v_0 = 220$ km/s, a local galactic escape velocity $v_{esc} = 600$ km/s, an Earth velocity relative to the DM distribution $v_E = 232$ km/s, a local DM density $\rho_D = 0.3$ GeV/c²/cm³. Alternative halo models based on N-body simulations, such as those from Navarro, Frenk and White [183] and Moore *et al.* [184], show an increase in density at the centre of the galaxy, known as a cusp. The main differences in these models, however, occur very near the galactic centre and will have little effect on absolute (i.e. non-directional) Earth-bound detection rates. There is still no certainty about the local structure and hence local halo properties. Clumpiness of the dark matter may result in a higher or lower local density, boosting or suppressing interaction rates. Whilst different models for the halo exist, the “standard” isothermal model will be used to compare experiments. A detailed description of the calculations will not be given here. Instead, we overview the most important results following the procedures set out in the classic review by Smith & Lewin [92].

From the “standard” galactic halo model, a theoretical recoil spectrum can be derived by calculating the kinematics of the collisions between a WIMP and a target nucleus. This shows a smoothly decreasing, featureless exponential differential spectrum, described by:

$$\frac{dR}{dE_R} = \frac{R_0}{E_0 r} e^{-E_R/E_0 r}, \quad (6.10)$$

where dR/dE_R is the energy-dependent event rate per keV_{nr}, R_0 is the total event rate, E_0 is the most probable incident kinetic energy of the WIMP:

$$E_0 = \frac{1}{2} M_D v_0^2, \quad (6.11)$$

and r is a kinematic factor,

$$r = \frac{4M_D M_T}{(M_D + M_T)^2}, \quad (6.12)$$

with M_D and M_T being the mass of the WIMP and the target nucleus, respectively.

However, Eq. (6.10) assumes a detector stationary with respect to the galaxy and therefore some corrections are required. A detector on Earth will be subjected to the

rotation of the Earth, the motion of the Earth around the Sun, as well as the motion of the Sun around the galactic centre. Additionally, the finite size of the target nucleus requires the application of a form factor correction, $F^2(E_R)$, which is dependent on the target nuclei used. By incorporating these additional factors, Eq. (6.10) becomes:

$$\frac{dR}{dE_R} = F^2(E_R) \frac{c_1 R_0}{E_0 r} e^{-c_2 E_R / E_0 r}, \quad (6.13)$$

where c_1 and c_2 account for the motion of the Earth, with average values of 0.751 and 0.561, respectively, calculated from the velocities previously stated and varying throughout the year (values tabulated in [92]).

The nuclear form factor is a function of the effective nuclear radius, the recoil energy, E_R , and the atomic weight of the target species, A . The currently adopted description of the form factor, the Helm approximation [185], uses the solid sphere approximation with a term to adjust for soft edge effects, giving:

$$F(qr_n) = \frac{3[\sin(qr_n) - qr_n \cos(qr_n)]}{(qr_n)^3} e^{-(qs)^2/2}. \quad (6.14)$$

where the momentum transfer is $q = (2 M_T E_R)^{1/2}$, r_n is the nuclear radius and s is the skin thickness parameter, with values of $r_n = 5.6$ fm and $s = 0.9$ fm assumed for xenon.

We now have the spectrum one would expect to observe in a “perfect” detector, that is, one with perfect energy resolution and 100% detection efficiency (the blue dashed line in Figure 6.12). Of course this is not the case in reality, so we must incorporate these detector characteristics into our predicted spectrum before comparing it with our observation. We first apply an energy resolution smearing to the prediction as this is intrinsic to the detector set-up (and, in some part, analysis procedures). The spectrum is convolved with a gaussian function representing the energy resolution determined by calibration data. The detector efficiencies are then applied to the resulting spectrum, producing the WIMP spectrum that would be observed by the real detector (the black line in Figure 6.12). When applying these efficiencies it is often easier to separate them into two categories (as discussed previously): energy dependent and energy independent. The latter can be applied separately as a constant factor in the calculation of

accumulated exposure, with only the former being applied directly to the predicted spectrum. The final step in comparing the spectrum with the observation is to convert the energy scale between E_R and E_{ee} using the relative scintillation efficiency and field suppression. The conversion of energy scales has been discussed previously in the context of ZEPLIN-III in Section 5.10.




Figure 6.12: Example of a differential WIMP recoil spectrum before (blue dashed line) and after application (solid black curve) of detector characteristics. The hatched area is that of the region of interest which is integrated for comparison with the observation.

In the analysis method used in the ZEPLIN experiments, the limit on the signal rate is calculated for a specific energy range by finding the limit on the signal and dividing it by the final exposure (with any energy-independent efficiencies folded in), to obtain a rate in events/kg/day. To calculate the limit on the interaction cross-section, we need to compare the limit on rate with the predicted spectrum to derive a limit on the total event rate, R_0 . This is done by integrating the predicted spectrum for $R_0 = 1$ (the hatched region in Figure 6.12) and then scaling it to the observed rate. This is related to the total WIMP-nucleus cross-section, σ_A , through:

$$\frac{R_0}{r} = \frac{D \sigma_A}{\mu_A^2}, \quad (6.15)$$

where $\mu_A = M_D M_T / (M_D + M_T)$ is the reduced mass of the colliding particles, and D is a numerical factor equal to 94.3 for the assumed local halo dark matter density of 0.3 GeV/cm^3 [92].

For comparison between different experiments with different target species, it is standard practice to report the final results in terms of the WIMP-nucleon cross-section, σ_{W-n} ,

$$\sigma_{W-n} = \left(\frac{\mu_1}{\mu_A} \right)^2 \left(\frac{1}{A} \right)^2 \sigma_A \quad (6.16)$$

where μ_1 is the reduced mass for $A=1$. An extra boost for detectors using target species with high A is the coherent enhancement factor, $\sigma_A \propto A^2 \sigma_{W-n}$, resulting from the scattering amplitudes of A nucleons adding in phase (for low momentum transfer). The calculation of the WIMP-nucleon cross-section also facilitates comparison against theoretical WIMP models, e.g. CMSSM [122]. The results on WIMP-nucleon cross-section are reported as a curve in plots of cross-section against WIMP mass, such as those shown in Figures 6.13 and 6.14. The shape of the curve and position of the minimum point are determined by the target nucleus, energy threshold and analysis techniques.

6.3 Final results from the ZEPLIN detectors

The results from ZEPLIN-II and the first science run of ZEPLIN-III were calculated using different statistical methods as a result of the specific conditions and the prevailing knowledge at the time. ZEPLIN-II was analysed using the traditional F-C method and the ZEPLIN-III data using the 2-D BML method. Here we present the final dark matter results from both detectors.

The F-C analysis of the ZEPLIN-II science data provided a two-sided 90% confidence interval on the number of recoil events observed, showing it to be consistent with zero. The upper limit of this interval provided a limit on the rate of $0.092 \text{ events/kg/day}$. To calculate the WIMP-nucleon spin-independent scattering cross-section the method described above was implemented. The standard WIMP halo model was utilised along with the appropriate detector efficiencies and energy resolution (see Chapter 4). The resulting limit curve had a minimum of $6.6 \times 10^{-7} \text{ pb}$ for a WIMP mass of $65 \text{ GeV}/c^2$,

and is shown in Figure 6.13 along with the world status at publication of the ZEPLIN-II result [101].

Although no further exploitation of the ZEPLIN-II instrument was possible, it was projected that if the Rn-progeny events could be removed (through cleaning of detector surfaces) and by eliminating ^{222}Rn emanation from the SAES getters, a sensitivity of 2×10^{-7} pb could be reached with a similar running time, extending down to about 1×10^{-7} pb with a longer science run of about 5 months. Any improvements in detector stability and corrections would also help improve the result.

ZEPLIN-II was the first two-phase noble gas instrument to publish a dark matter result, with only the CDMS-II result excluding more of the parameter space. ZEPLIN-II also contributed to rule out the DAMA evidence region, interpreted in terms of spin-independent interactions causing nuclear recoils. The predictions for ZEPLIN-III (based upon simulations) indicated a minimum of 9×10^{-8} pb for a $70 \text{ GeV}/c^2$, shown in Figure 6.13 by the dashed line. The focus of the collaboration was thus shifted to the new experiment.

The BML analysis of the ZEPLIN-III first science run produced a most likely signal of 0 events and a 90% confidence upper limit on the number of recoils which varied




Figure 6.13: World status (at ZEPLIN-II publication, Jan 2007) of experimental (direct detection) limits on WIMP-nucleon spin-independent scattering cross-section [124]. The solid lines show experimental direct detection limits previously published, the red region shows the reported signal from DAMA and the green region shows the areas of the parameter space from currently favoured MSSM models.

with WIMP mass, between 2.45 for $M_D=10$ GeV and 3.0 for $M_D=1000$ GeV. The rate is calculated by factoring in the acceptance percentage (47.7%) and the effective exposure (267.9 kg/days), resulting in an upper limit between 0.192 events/kg/day and 0.235 events/kg/day. This is higher than ZEPLIN-II due to significant additional signal power probed at low energy by the lower threshold.

The WIMP-nucleon spin-independent scattering cross-section is then calculated, using the same standard galactic halo model, the 90% confidence limit on rate and the detector specific efficiencies, energy resolution and varying quenching factor (as described in Chapter 5). This results in a limit curve which has a minimum at 7.7×10^{-8} pb for a WIMP mass of 55 GeV/c² [103]. The main change in procedure from the ZEPLIN-II limit calculation was the requirement to calculate the 90% confidence limit on rate for each different WIMP mass and then feed these into the calculation separately. The final limit curve is shown in Figure 6.14 with the current world status (at the time of writing).




Figure 6.14: Current world status of experimental (direct detection) limits on WIMP-nucleon spin-independent scattering cross-section [124].

If the constant L_{eff} energy conversion had been utilised instead, with the mismatch between simulated and experimental recoil spectra interpreted in terms of a large efficiency loss of unknown origin, then the limits on rate would have been slightly higher, varying from 2.52 events at 10 GeV/c² to 3.1 at 1000 GeV/c² (3.0 at 55 GeV/c²). The resulting limit on σ_{W-n} is also higher, with a minimum of 1.13×10^{-7} pb for a

WIMP mass of $52 \text{ GeV}/c^2$ ($\simeq 50\%$ higher). Clearly, the choice of energy conversion (and the resulting efficiencies) has an effect on the final result, although this is found to be relatively minor.

At low WIMP mass, the ZEPLIN-III result is the second best limit published from a two-phase noble gas instrument, but surpassing XENON10 in sensitivity above a WIMP mass of $100 \text{ GeV}/c^2$. However, the results from the XENON10 experiment are not directly comparable with other experiments as the confidence limits quoted lack full coverage of 90%. This is due to the choice of adopting the Yellin Optimal Interval method after the data had been analysed, leading to flip-flopping (see Ref [25] in [104]). The coverage is expected to be in the region of 80-85% although a precise calculation is not given. Additionally, the published XENON10 result utilised a flat relative scintillation efficiency, $L_{eff} = 0.19$; later studies by their team [174] analysed the data using a varying L_{eff} resulting in a slight increase in the limit. Both these factors result in a higher limit curve than that shown in Figure 6.14. A way to remove the flip-flopping problem in XENON10 would be to use the F-C method on the published data. This results in a limit curve with a minimum of approximately $1 \times 10^{-7} \text{ pb}$.

The result from ZEPLIN-III achieved the predicted sensitivity with, arguably, the cleanest data-set published from a two-phase xenon detector. The ZEPLIN-III, XENON10 and CDMS-II results are all beginning to probe the regions of parameter space favoured by SUSY neutralino models. The shaded areas show favoured regions of parameter space as derived from a bayesian analysis of CMSSM models [122]. This suggests that the next phase of detector development will bring instruments able to probe large sections of the favoured parameter space, potentially making the first detection of WIMP dark matter.

Chapter 7

Xenon Physics with the ZEPLIN Detectors

The primary goal of the ZEPLIN detectors is to search for a WIMP signature or, in the event of a non-discovery, to set limits on the properties of potential WIMP particles. However, we are able to learn more about xenon physics from the instruments, mainly due to their large active volumes and location in a low-background environment. This chapter details work carried out to investigate aspects governing the physics underlying two-phase Xe detectors, which are of interest to the development of this technology, as well as to enhance the quality of our own data.

A review of studies relevant to the physics of these devices was presented in Chapter 3, including a discussion of the different physical processes at work in two-phase systems. Presented here is a novel quantitative study of single electron emission in noble liquids, helping to demonstrate the sensitivity of the two-phase technique, with consideration of photoionisation as a possible cause (this work was published in [129]). Other physics studies also described here include an analysis of the anti-correlation between scintillation and ionisation signals (with reference to data from both ZEPLIN-II and III), the electric field dependence of the detector response and scintillation light yield variation with xenon purity. My involvement in these studies was a dominant one.

7.1 Single Electron Emission

High-efficiency detection of a single quantum associated with a particular response channel is a clean sign of both the sensitivity and the quality of the measurement technique. Presented here is the first study quantifying the response to the cross-phase emission of a single ionisation electron in a noble gas detector. In two-phase argon, a three-stage Gas Electron Multiplier (GEM) has recently been reported to achieve single-electron sensitivity, although this study relied on electrons photo-produced in the first GEM (in the gas) rather than emitted from the liquid [186]. Sensitivity to single ionisation electrons is important in experiments searching for very small, rare events. On a technical level, it allows for direct measurement of the ionisation yields of different interacting particles, such as nuclear and electron recoils, and may help with the study of photoionisation processes in LXe. Understanding the origin of these electrons may highlight new backgrounds for experiments relying on the detection of even smaller ionisation signatures than those considered of interest in WIMP dark matter searches, which motivated this study. One such example is the proposed detection of sub-keV energy deposits from coherent neutrino scattering [187]. This work was carried out using the ZEPLIN-




Figure 7.1: Seven-PMT sum waveform containing a secondary electroluminescence pulse, with several potential single electrons trailing S2.

II instrument, in which the liquid surface lies between two meshes, 10 mm apart, used to generate a strong electric field to produce electroluminescence photons from emitted electrons. These extraction electrodes consist of a woven stainless steel mesh of $30\text{ }\mu\text{m}$ wire, at a pitch separation of $500\text{ }\mu\text{m}$. Free electrons produced in the target bulk below this region are drifted towards the extraction region by a vertical field of 1 kV/cm ; upon reaching the surface, they are emitted into the gas by the stronger electric field in the extraction region of $\sim 4.8\text{ kV/cm}$ and accelerated across the 2-3 mm gas gap by a field twice as strong. Both the primary (scintillation) and secondary (electroluminescence) signals are independently detected by seven photo-multiplier tubes (PMTs). The time taken for the electrons to drift through the liquid provides a separation in time proportional to the depth in the detector.

This analysis was carried out using the data from the 31-day shielded run of ZEPLIN-II. During this run the average background rate from γ -ray interactions was $\sim 2\text{ evt/s}$ ($> 5\text{ keV}$). An unexpected population of very small secondary-like signals following large secondaries was apparent during early tests, as illustrated in Figure 7.1, which motivated this work. Some of these signals were also observed *between* the primary and secondary pulses associated with normal events and these were selected for further analysis, for the reasons explained below.

7.1.1 Single electron response

The detector response to a single electron emitted into the gas region is designed to be small, yielding fewer than 10 photoelectrons across the PMT array, since excessive gain brings about dynamic range limitations. In addition, the faintest nucleon recoils one wishes to record for the purpose of WIMP searches are determined by the limited number of S1 photons detected and there is little to be gained by recording S2 signals below the software threshold of the S1 response.

To search for such small signals, quiet timelines, free from spurious noise, overlapping events and optical feedback effects are required. For this reason the search for candidate single electron signals was carried out on the low-background dataset used for WIMP searches. From this dataset, only events triggered by the primary scintillation signal were selected in order not to bias the choice of candidate signals. Allowing S2 triggers

would bias against single electron pulses which could have triggered the system between primary and secondary, and allowing self-triggers by small pulses would obviously bias towards larger signals.

The ZEPLIN-II trigger function was a 5-fold coincidence above $1/3$ of a single PMT photoelectron (phe) ($\gtrsim 80\%$ efficiency per channel) within a narrow coincidence window, as already described in Section 4.1, adopted specifically to avoid triggering by single electrons, minimising the impact on data volume. For this reason, the absolute rate of single electron emission in the dataset cannot be calculated, only the rate relative to primary scintillation triggers.




Figure 7.2: Seven-PMT sum waveform containing a candidate single electron signal (p5) visible between the primary scintillation pulse (p4) and the electroluminescence pulse (p6).

Figure 7.2 shows a typical $80 \mu\text{s}$ waveform (sampled at 2 ns) containing a primary scintillation pulse and a secondary electroluminescence pulse, with a single electron candidate observed in the intervening time. The latter is detected as a collection of individual PMT photoelectrons spread over a ~ 500 ns period, the time required for electrons to cross the high-field gas region.

In the analysis described in Chapter 4, the data were corrected for the finite electron lifetime in the LXe as well as operational parameters which affect the gain of the ionisation channel (such as variations of pressure, liquid level and electric fields). In this

analysis the purity correction, compensating the secondary signal for electron trapping by electronegative impurities during their drift to the surface, is not required as a single electron will either reach the surface or be trapped, meaning no *partial* loss of signal. The other operational parameters are considered as variables.

A pulse area histogram of all small signals detected between primary and secondary pulses is shown in Figure 7.3, with the conversion from nVs to photoelectrons (phe) calibrated by the mean measured for each PMT independently. The figure shows a clear population which we attribute to single electrons, along with an exponential noise pedestal. The distribution of single electrons is fitted with a Gaussian. Although photoelectron statistics suggest a Poisson distribution, further broadening occurs due to electronic noise and other fluctuations. The spectrum shows a mean of 8.8 ± 0.4 phe and a width $\sigma = 5.0$ phe, corresponding to a gain of over 200 VUV photons/electron, as discussed later.




Figure 7.3: Single electron spectrum at 1.5 bar gas pressure. The continuous line shows the fit to the entire spectrum; the Gaussian and exponential components are also shown. The inset shows the resulting fit parameters.

7.1.2 Predicting single electron response behaviour

The number of electroluminescence photons created per electron depends on the gas pressure, electric field and gas thickness. The electroluminescence yield of xenon, i.e.

the number of VUV photons produced per electron per cm travelled, has been studied mainly for room-temperature gas (see [152] and references therein), as described in Chapter 3. A dependence is found of the form $Y = aE - bP_{eq}$, where a and b are experimental parameters, E is the field in the gas and P_{eq} is the equivalent pressure for the same gas density at 0 °C. It is known that the photon yield in the cold, saturated vapour is higher than that in the warm gas (for the same density). This effect is clearly shown in [152], where the room temperature measurement is consistent with some published results, namely [188], but the yield in the cold vapour is clearly higher. Table 7.1 summarises experimentally measured values of a and b .

Table 7.1: Experimentally measured electroluminescence yield constants, a and b .

Ref.	Temperature	a	b
Santos <i>et al.</i> [189]	20 °C	139	100
Monteiro <i>et al.</i> [152]	20 °C	140	116
Akimov <i>et al.</i> [188]	20 °C	70	56
Fonseca <i>et al.</i> [151]	20 °C	86	73
	-90 °C	137	125
Aprile <i>et al.</i> [190]	-95 °C	120	154

For the mesh design and the thickness of the gas layer in ZEPLIN-II, a parallel and uniform electric field can be assumed without significant error. This allows calculation of the absolute yield in ZEPLIN-II. If the small signals under scrutiny correspond to the emission of a single electron, this yield must agree with that measured for the cold vapour and show the same pressure dependence.

Gas pressure, P , is a key operational parameter affecting the single electron response, doing so in two complementary ways: it affects electroluminescence at a microscopic level and is also directly linked to the liquid level through thermal expansion of the liquid. The liquid is in thermal equilibrium with the gas due to the slow variation in temperature. As a result, any change in temperature causes a proportional variation of the gas thickness. Besides variation of the path length, this thermal expansion also causes a small variation of the field in the gas, at the level of 10%, which can be ignored for this purpose.

Previous yield measurements can be compared directly with ZEPLIN-II results by factoring in the thickness of the gas region, d , the light collection efficiency, η , and the




Figure 7.4: Dependence of single electron response on detector pressure, compared with predictions from previous measurements of electroluminescence yield in saturated vapour. The shaded bands represent uncertainties assigned to these predictions, including an ad-hoc 10% error assumed for parameters a and b (not given in the literature).

PMT quantum efficiencies, QE . The gas thickness is calculated from the drift time of background interactions occurring very near the lower extraction grid just under the surface; the light collection was simulated by Monte Carlo; there is some uncertainty about the variation of PMT QE down to low temperatures for this particular phototube model, but we believe it to be small. The predicted mean number of photoelectrons per electron is calculated from:

$$\underbrace{SE(P)}_{phe/e^-} = \underbrace{Y(P) \times d_g(P)}_{ph/e^-} \times \underbrace{\eta \times QE}_{phe/ph}, \quad (7.1)$$

where we consider a volume-averaged light collection fraction $\nu = 0.236 \pm 0.018$, approximate low-temperature quantum efficiency $QE = 0.17 \pm 0.03$ and $Y(P)$ is the electroluminescence yield.

Figure 7.4 shows the mean single electron response as a function of pressure, demonstrating a dependence clearly consistent with expectation over the range 1.2–1.9 bar,

the variation observed during this run. Both the absolute number of photons and the pressure behaviour agree well with predictions based on independent measurements of the absolute yield and, combined, provide strong evidence that the candidate population is indeed from single electrons. From our data it was not possible to accurately calculate the experimental parameters a and b for the electroluminescence yield due to the subtle dependence of the field on pressure. For a convincing measurement, one must be fixed and the other allowed to vary systematically.

7.1.3 Production of single electrons

In ZEPLIN-II position reconstruction in the horizontal plane lacked precision for signals as small as those considered here, as discussed in Chapter 4. However, the reconstructed radial distribution for these events, shown in Figure 7.5, is consistent with generation over all radii, and is clearly incompatible with that of small ionisation signals originating from the detector walls (corresponding to a radius of $r = 0.7$ a.u.). The latter is the known background of nuclear recoils arising from plating of the PTFE walls with α -emitters from ^{222}Rn decay [101], which was a major limitation in ZEPLIN-II.

For a source uniformly distributed throughout the entire liquid bulk and a detector with perfect position resolution, a linear increase with r would be expected (Figure 7.6). To predict the form of a uniform volume distribution in the case of very poor resolution, we convolve the x and y coordinates of a population distributed uniformly in x - y with gaussian functions. By also applying the same smearing method to a population from the detector walls, we are able to replicate the shape of both the single electron distribution and Rn-progeny population, suggesting that the source of single electrons is uniformly distributed throughout the bulk. The depth (drift time) distribution (Figure 7.5) also shows relatively uniform production throughout the LXe bulk.

The fact that large secondaries appear to be followed by multiple single electron pulses suggests their production may be related to the number of VUV photons in the chamber. The quiet part of the waveforms found between primary and secondary signals, together with the measured number of primary scintillation photons, allow testing of this hypothesis in a quantitative manner. Figure 7.7 shows the fraction of events where a single electron is observed as a function of energy, where the latter is




Figure 7.5: Radial (top) and depth (bottom) distributions. In the top plot the solid line shows the single electron population and the dashed line the distribution of small ionisation signals (\sim few electrons) from detector walls (^{222}Rn -induced events).




Figure 7.6: Prediction of the radial distribution of a population uniformly distributed across the volume (solid line) and from a population of wall events (dashed line).

measured by the number of scintillation photons generated in the liquid. A clear energy dependence is observed, suggesting that the production of single electrons could be, at least in part, due to photoionisation processes in the liquid, possibly of contaminant species still remaining at \sim ppb concentrations.




Figure 7.7: Fraction of S1-triggered events where a single electron is observed as a function of primary signal area (normalized to γ -ray energy).

This hypothesis gains some additional weight when the rates are examined as a function of the purity of the xenon. Figure 7.8 shows a decrease in the observed rate of single electron emission with increasing electron lifetime (i.e. reduced presence of impurities), despite the small range of electron lifetimes observed during the ZEPLIN-II science run. If the process was due to photoionisation of electronegative impurities, the expected dependence would be a linear relationship with the concentration of such an impurity.

7.2 Photoionisation as a production mechanism

The mean free path (mfp) of the $\simeq 7.1$ eV photons for generating an electron in LXe can be estimated from Figure 7.7. The slope of the trend line suggests that, on average, approximately 800 keV is required to produce a single electron. Combining this with a scintillation yield of 39 photons/keV (measured at 1 kV/cm), this corresponds to an average of 32,500 photons per electron. Scintillation photons can escape from the liquid surface or be absorbed in the surrounding PTFE or in the electrodes. A Monte Carlo simulation places the mean escape length from the LXe at ~ 25 cm for photons generated




Figure 7.8: Fraction of S1-triggered events where a single electron is observed as a function of electron lifetime during the science run.

isotropically and uniformly. This value is not significantly affected by bulk absorption for the xenon purity considered here, or by Rayleigh scattering. This indicates a photon mean free path for photoionisation of $\sim 1.0 \times 10^6$ cm.

Several chemical species, including xenon itself, could be responsible, but impurities in the liquid are the most likely. Unfortunately, the mfp depends on both the microscopic cross-section and the atom number density, so none can be ruled out with certainty since either (or both) quantities may be unknown.

Sub-threshold photoionisation, either of impurities or of Xe atoms, dimers and higher order polymers, cannot be ruled out, even if it is unlikely. Some of the most abundant impurities are electronegative species (e.g. O_2 , H_2O , N_2O , etc) responsible for the finite electron lifetime. Although these have ionisation energies typically above 12 eV in the gas phase [191], this does not rule out completely the possibility of photoionisation at lower energies. Their concentration can be estimated from the rate of electron attachment. From known attachment cross-sections [149] at an average energy of 0.1 eV

for electrons drifting in a 1 kV/cm field [192], we estimate that an O_2 concentration of ~ 60 ppb or a N_2O at just ~ 8 ppb would produce the $\sim 100 \mu\text{s}$ electron lifetime observed during this run. The photoionisation cross-section needed to explain the measured mfp is only ~ 0.001 Mb (compared to 10-100 Mb typically found well above threshold).

In the case of intrinsic LXe photoionisation (9.3 eV threshold [193]) a cross-section as low as $35 \mu\text{b}$ is sufficient. So, despite the sub-threshold process being very unlikely, the interaction probabilities required are also extremely small; in addition, the non-zero width of the scintillation emission will play a favourable role, with the tails of the emission and absorption distributions intersecting at some level. However, note that no dependence on the electron lifetime is expected in this instance.

Alternatively, minute amounts of species with low ionisation thresholds may be responsible, of which there are many candidates, organic and inorganic. In addition to neutral species, it is also conceivable that negative ions previously created by electron attachment can be photoionised during their long drift towards the anode, which is as slow as ~ 0.7 cm/s for a O_2^- ions in LXe [194]. During the science data taking period of ZEPLIN-II the fields were on for 108 “standard” days. A “standard” day includes 23 hours of background exposure, with ~ 2 MeV/s deposited creating some $\sim 30,000$ O_2^- ions per second, plus 1 hour of ^{57}Co calibration, depositing ~ 12.5 MeV/s, creating $\sim 187,500$ O_2^- ions per second. This would produce about 3×10^9 O_2^- ions per day, totaling approximately 3.2×10^{11} O_2^- ions over the whole 108-day run. Accounting for the xenon volume gives a number density of about $3.2 \times 10^7 \text{ cm}^{-3}$, resulting in a required cross-section for photoionisation of about $6 \times 10^{-14} \text{ cm}^2$. Figure 7.9 shows the published values of photoionisation cross-section for O_2^- over the measured energy range up to 3.5 eV (well below the xenon scintillation energy of 7 eV). Even though these measurements are below the energy of interest, generous extrapolation would still leave the cross-section several orders of magnitude away from the required value, suggesting the concentration of O_2^- accumulated would be too low to explain the observed rate.

Finally, we mention photoionisation induced by the well-known ‘n=1’ LXe exciton, which lies below the intrinsic threshold at 8.4 eV [196, 197, 198]. Although excitons do not cause photoionisation directly, they can transfer the excitation to impurities onto




Figure 7.9: Measurements of the photoionisation cross-section for O_2^- ions [195].

which they become trapped and ionise them in a Penning-type interaction.

In summary, a population of small signals in the ZEPLIN-II data was identified with the emission of single electrons into the gas region, probably caused by photoionisation of a yet underdetermined contaminant species in the LXe. The detection of single electrons demonstrates the excellent sensitivity of the ionisation channel in two-phase xenon systems.

7.2.1 Single electron emission in ZEPLIN-III

Following the study with ZEPLIN-II data, a search for single electron emission in the first run of ZEPLIN-III has been carried out. A visual scan of science data events suggested many electron candidate pulses (an example is shown in Figure 7.10). As in ZEPLIN-II, there were also potential single electron pulses following many secondary signals.

A clear optical feedback signal was also identified, consisting of a few electrons originating at the cathode grid. This was expected due to the large number of S2 photons produced, the photoelectric coefficient of stainless steel ($\sim 10^{-5}$) and the geometry of the wire grid.

We can predict the size of a signal from a single electron emitted into the gas phase as previously (Eq. 7.1). An average cold QE of $\sim 30.0\%$ was measured experimentally [166], the light collection for electroluminescence is approximately 20% (from simulations), a




Figure 7.10: A waveform of the sum of all 31 PMTs in the high sensitivity channel of ZEPLIN-III, showing an S1 and S2, with a small single-electron-like pulse in the intervening time.

gas gap of ~ 4 mm was estimated and the data were taken with a field of about 7.8 kV/cm (in the gas) at an average pressure of ~ 1.63 bar. These values would suggest that a single electron emitted into the gas would produce about 300 photons as it traverses the gas gap, resulting in ~ 18 phe detected across the PMT array.

A single electron analysis similar to that carried out with the ZEPLIN-II data was applied to the data from the ZEPLIN-III first science run, producing the spectrum of the detector response to single electron emission shown in Figure 7.11. The typical area for a single electron pulse was roughly 1.2 nVs, equivalent to about 25 phe (assuming a single phe area of 0.047 nVs), which is larger than our original prediction. This disagreement could easily be accounted for by uncertainties in the gas thickness (± 0.5 mm) and S2 light collection efficiency ($\pm 5\%$), along with the variation of the prediction on the choice of coefficients for the electroluminescence yield. This preliminary work on single electron emission in ZEPLIN-III is currently being developed.

Following the first science run, very long timelines were acquired without a signal-derived trigger, to search for single electrons not produced by events triggering the system. This will allow for the processes to be understood in a wider context, searching for other causes of single electron emission (including other potential rare physical processes). These data will be analysed in due course.




Figure 7.11: A typical single electron emission spectrum measured from ZEPLIN-III. The spectrum is fitted with an exponential noise pedestal, a gaussian for the single electron peak, and another gaussian corresponding to the possible emission of two electrons. (supplied by P.Majewski)

7.3 Nuclear Recoil Ionisation Yield

Measurement of the ionisation yield from nuclear recoils is important in the context of two-phase searches for WIMPs, as WIMP-nucleon interactions produce low-energy nuclear recoils in the xenon. By combining measurements of the detector response to a single electron emitted from the liquid surface with nuclear recoil calibration data, the ionisation yield for elastic scattering can be extracted. Measurements in the literature were carried out with two different test chambers [199], utilising a calibrated charge amplifier to measure the ionisation signal, and with data from the XENON10 experiment [145], utilising the same two-phase technique as ZEPLIN-II and III. Given the difficulty in calibrating charge amplifiers and the associated systematic errors, the direct calibration of charge response is of great interest.

The nuclear recoil ionisation yield, Υ , measured in electrons per keV_{nr} , is calculated from the measured mean $S2/S1$ parameter of the nuclear recoil population, $\langle \frac{S2}{S1} \rangle$:

$$\Upsilon = \frac{\langle S2 \rangle}{\eta SE} \times \frac{\mu}{\langle S1 \rangle_{s1cal}} = \left\langle \frac{S2}{S1} \right\rangle \frac{\mu}{\eta SE_{s1cal}} \quad (7.2)$$

where we require the single electron response of the detector, SE [nVs^{-1}], corresponding

to the experimental conditions at the time, the fraction of electrons emitted into the gas phase, η , the energy calibration, $s1cal$ [$\text{keV}_{ee}/\text{nVs}$], and the nuclear recoil energy conversion factor, $\mu = S_e/(L_{eff} S_n)$ [$\text{keV}_{nr}/\text{keV}_{ee}$]. The ZEPLIN-II AmBe calibrations gave a parameterisation for the mean S2/S1 response for nuclear recoils:

$$\left\langle \frac{S2}{S1} \right\rangle = (265 \pm 10) \times E_{ee}^{-0.4 \pm 0.14} [\text{nVs} / \text{nVs}]. \quad (7.3)$$

The cross-phase emission efficiency for electrons is determined by the field in the liquid. Published measurements [135] suggest an emission efficiency of $\eta \sim 95\%$ at the appropriate field in ZEPLIN-II of 4.7 kV/cm. Loss of drifting electrons through trapping by electronegative contaminants is accounted for by correction of S2 response before parameterisation. The conversion of energy scales from electron equivalent to recoil energy was discussed in the previous chapter, with a constant factor of 2.8 $\text{keV}_{nr}/\text{keV}_{ee}$ assumed in this calculation.

The nuclear recoil ionisation yield measured in ZEPLIN-II, shown in Figure 7.12, appears lower than previously published values. Potential causes for this disagreement involve a mechanism inhibiting cross-phase emission of electrons, either through contamination of the liquid surface, or variation of the electric field. Despite correction of the AmBe data for the finite electron lifetime, there may potentially be some effect due to the variation of electron lifetime with event energy (or size of the S2 pulse). The electron lifetime is calibrated in the ZEPLIN detectors with γ -ray or high-energy α interactions, both having significantly larger S2 signals than the low-energy nuclear recoils of interest. However, the cause of the discrepancy with the results in the literature has not yet been resolved.

As described previously, data from ZEPLIN-III was analysed for emission of single electrons. Using those measurements and the nuclear recoil calibration (described in Section 5.9.1), the ionisation yield was calculated for both a constant and a varying quenching factor (as shown in Figure 7.12). The values using a constant quenching factor are in reasonable agreement with the published data, especially considering the stronger electric field. However, utilising the varying quenching factor causes the yield to level off, or even drop, at lower energies. This work is preliminary however, and a more detailed study of the single electron emission spectra is underway. Measurements of the




Figure 7.12: Nuclear recoil ionisation yield measured from the ZEPLIN-II (solid black line) and preliminary results from ZEPLIN-III (black symbols) applying a constant or varying relative scintillation efficiency, compared with the previous published results [199] and data from XENON10 [145].

ionisation yield from the ZEPLIN instruments provide complementary information to the data already published and future measurements should be possible with ZEPLIN-III probing field variation of the yield, previously only measured up to 2 kV/cm.

7.4 Scintillation light yield as a function of purity

During active xenon purification in ZEPLIN-II, described in Chapter 3, the purity improved dramatically, with electron lifetimes increasing from 4 μ s to about 1 ms. ^{57}Co calibrations were taken during this period, allowing the light yield to be monitored. The light yield, measured by fitting to the combined 122/136 keV peak from ^{57}Co , was calibrated using the average measured single photoelectron area, 0.035 nVs. Figure 7.13 shows a clear increase in the measured scintillation output with increasing electron lifetime (measured with α -decay events) up to a maximum value of 0.75 phe/keV $_{ee}$.

The exact source of this dependence is unknown, although it cannot result from a




Figure 7.13: Plot showing the relationship between light yield and electron lifetime in ZEPLIN-II, where light yield values are normalised to that at zero field.

variation in the final detection of photons which is unaffected by the purity of the liquid. This leaves two possibilities; loss of photons during transport from the interaction site to the PMTs or during the production process itself. During transport the VUV scintillation photons can be absorbed by dissolved molecular impurities within the xenon, reducing the photon attenuation length and hence the light yield. This is known to be one of the main limitations on the transparency of liquid noble gases to their own scintillation light [200]. It has been shown that H_2O at the ppm level is the most probable impurity for absorption of VUV photons in LXe [201], a hypothesis which would agree with the increased light yield observed during purification. The production process itself may provide an alternate method for reducing prompt scintillation, if an increased concentration of electronegative impurities inhibits the recombination process, reducing the number of photons produced.

While the observed degradation of scintillation light yield with decreasing xenon purity may imply that the ZEPLIN-III light output would be strongly affected, this is not the case for two reasons. ZEPLIN-III operates with a stronger electric field,

which significantly reduces the electron lifetime due to the field dependence of electron attachment to impurities. Consequently, the observed electron lifetime $\sim 40 \mu\text{s}$ at 4 kV/cm would be significantly larger at 1 kV/cm. Secondly, the light output is degraded by reduction of the photon attenuation length and, as a result, the geometry of the target volume will play an important role in determining the magnitude of the effect.

Further studies were not possible using ZEPLIN-II, as the purity remained high (above 100 μs) for the remainder of the science run. Equally studies have not been possible with ZEPLIN-III as the range of electron lifetimes required was not available, due to the stable purity in the instrument. Additional studies will be carried out if purity is not achieved initially in future operation of ZEPLIN-III.

7.5 Field-dependent detector response

Application of electric field is the cornerstone of two-phase operation in noble gas detectors, with a number of physical processes at work in the detectors. For our understanding of these instruments, allowing us to design new systems, optimise operation and analysis techniques, we require knowledge of the field-dependent detector response.

During commissioning of ZEPLIN-III, data were acquired with a ^{57}Co source, measuring the field-dependent response of the S1 and S2 channels. The radioactive source was placed above the detector and data acquired initially with no field in the liquid. After acquisition of a small zero-field dataset, the acquisition continued whilst voltages were increased at a rate of 180 V/min on the cathode and 120 V/min on the anode (an overall rate of 300 V/min between the two). Upon reaching the nominal operating voltages of the detector (17 kV between anode and cathode), a larger dataset was acquired at full field.

An interaction in the xenon produces scintillation and ionisation, a fraction of which will recombine producing additional scintillation photons. As a result, the response of the detector in the S1 channel is dependent upon the electric field applied across the target volume, as an increased field suppresses recombination. This field suppression has been measured previously [199] and can be used for comparison with the ZEPLIN-III data.

Figure 7.14 shows the onset of field suppression, exhibiting a decrease of the S1

response of the detector with event number in the dataset. No field was applied to the volume for approximately 5000 events, providing a zero-field calibration for this investigation. The zero-field calibration was fitted in Figure 7.14 giving a mean of 29.4 nVs corresponding to the combined 122/136 keV peak (5 phe/keV_{ee}); this is shown in the figure by the solid black line. After approximately 5000 events the electric field was applied, causing suppression of the recombination, resulting in a decrease in the S1 response.




Figure 7.14: Plot showing the onset of field suppression, with S1 response as a function of event number (left) and the zero-field events fitted as a calibration (right).

This study, undertaken during the commissioning phase, required matching of the electric field strength to the event number. In later datasets the applied voltages were recorded event-by-event, but this had not yet been implemented during the commissioning period. To calculate the field, Eq. 3.5 is applied assuming a gas gap of 4 mm (a liquid thickness of 36 mm). Using the number of files acquired and the ramping rates, an approximate calibration from event number to applied field was possible.

Figure 7.15 shows the S1 response as a function of applied electric field, exhibiting a decrease with increasing field. The solid black line represents the predicted response applying the previously published suppression measurements [199] to the zero-field calibration, exhibiting good agreement with the data. The vertical gaps in the data are the results of PMT power supply trips, with missing time periods translating to missing coverage of field strength. The region marked by the red lines is defined by the higher trigger threshold present during the commissioning phase. Below the horizontal line,

S1 pulses are no longer large enough to trigger the system, and the associated S2 pulses only become large enough to trigger above ~ 1.7 kV/cm, resulting in a region with no events. This effect is also evident from the increased fraction of higher energy events below 1.7 kV/cm, where the larger S2 triggered events account for a larger proportion of the triggering events.




Figure 7.15: Dependence of S1 signal with electric field in ZEPLIN-III. The black line represents the prediction from previously published measurements of field suppression. The red mark the S2 and S1 trigger thresholds.

The response of S2 can be investigated in a similar manner, except now multiple processes must be considered. As with S1, the applied field alters the fraction of electrons drifted from the interaction site. However, it also affects the cross-phase emission efficiency of electrons and the number of secondary photons produced in the electroluminescence process. There may additionally be an effect on the xenon purity through variations in the electron attachment cross-sections of the contaminant species, although this affect will be minimal in the small drift depth of most ^{57}Co events.

Previously published data measuring these processes have been parameterised for simulations of the instrument during the design and construction phase [150]; these parameterisations are shown in Figure 7.16. The fraction of electrons drifted increases

immediately upon the application of an electric field, but the other two processes have a minimum field threshold before they can proceed. The onset of electroluminescence occurs at about 1.4 kV/cm in the liquid (dependent on the parameterisation used) when the electrons acquire enough kinetic energy to collisionally excite the gaseous xenon atoms. The cross-phase emission only occurs at a still higher field, about 1.7 kV/cm (although the parameterisation at the onset is based on only one data point), when the electrons acquire enough energy to overcome the cross-phase potential barrier (LXe has a positive electron affinity relative to the vapour phase). As a result, this is the field at which S2 pulses begin to be produced, although they may only be observed by the detector at a marginally higher field as the pulses must be large enough to trigger the system.

The overall S2 response is thus determined by:

$$S2 = \underbrace{N_e(E) \times \eta(E) \times (Y(E))}_{\text{energy-dependent factors}} \times d_g \times \nu \times QE, \quad (7.4)$$

where N_e is the number of electrons drifted from the interaction site, η is the cross-phase emission efficiency and Y is the electroluminescence yield, all of which are energy-dependent; d_g is the thickness of the gas layer, ν is the light collection and QE is the PMT quantum efficiency. Using the parameterisations from Figure 7.16 and approximations for the energy-independent factors, the energy-dependent secondary response of the instrument can be calculated.

Figure 7.17 shows the S2 response of the detector as a function of the applied field, along with the prediction for the energy-dependence (black line, normalised to data at 2.7 kV/cm). The general behaviour is in reasonable agreement, especially considering that the calibration of electric field strength was only approximate. The reasonable agreement with predictions of the field dependence in both response channels gives us extra confidence in the simulations used to help analyse the science data.

Further data have been acquired with ZEPLIN-III to investigate the field dependence in more detail. These data had the applied voltages (and therefore field) recorded event by event for an accurate calibration, along with better measurements of the absolute factors (e.g. gas thickness). A slower rate was applied to the voltage ramping improving




Figure 7.16: Parameterisations of the three main field-dependent processes in producing S2. (a) Electroluminescence yield, $Y = 137(2.0 \times E) - 315$, where E is field in the liquid (in kV/cm) and a pressure of 1.6 bar is used, (b) Emission efficiency into gas phase, $\eta = 1 - \left(1.1 \times e^{-0.017 \cdot E^{3.45}}\right)$, (c) Relative charge extraction from the interaction site, $\frac{N_e}{N_e(\infty)} = 1 / \left(\frac{0.24056}{E} + \frac{0.4}{E} + 1\right)$, for 122 keV γ -rays.

the statistics of the measurements and the field was ramped down at the same rate to monitor the process for any hysteresis.

7.6 Scintillation/Ionisation anti-correlation

During an interaction in the liquid, the energy deposited by electron recoils is split between two processes, atomic excitation and ionisation. These translate into the two response mechanisms generating S1 and S2 signals. The relative division of the mean number of quanta into each channel depends on the electric field applied at the interaction site, with statistical fluctuations between the two worsening the energy resolution




Figure 7.17: Dependence of S2 signal size with electric field, compared with a prediction (black line) made from the parameterisations shown in Figure 7.16.

in the individual S1 and S2 channels. The fluctuations occur as the ionisation created by the interaction will either drift away to the liquid surface, adding to the S2 signal, or recombine to produce additional prompt scintillation photons, increasing the size of S1. These fluctuations between excitation and ionisation cause an anti-correlation between S1 and S2 for a given total energy deposit.

In an effort to recover a more precise measurement of the energy deposited in an event, a combined energy parameter, E^* , can be created from a mixture of the S1 and S2 parameters. The initial study using ZEPLIN-II data utilised a simple combination of S1 and S2:

$$E^* = S1 + k \cdot S2, \quad (7.5)$$

where the coefficient k normalises the relative gain of the responses, and was therefore taken as the ratio of the mean S1 and S2 responses from the ^{57}Co photopeak. This assumption is a reasonable approximation for the drift field applied in ZEPLIN-II (1 kV/cm), where the field suppression suggests that 50% of the energy is transferred to ionisation.

Figure 7.18 shows the improved resolution of E^* over S1 alone, demonstrated with a ^{57}Co calibration. The E^* peak is noticeably tighter, with an improvement in resolution from 25% to 16.5%, a tightening also observed in the colour plots in Figure 7.19. The ^{57}Co colour plots show the anti-correlation effect clearly, where the region of the photopeak is tilted exhibiting an increased S2 for a lower S1 (and vice versa), when S1 alone is used to determine energy. By employing E^* in place of S1, the population is tightened and the tilt removed. The ^{60}Co Compton calibration can be thought of as a series of tilted regions at all possible deposited energies, with the population tightened by utilising E^* , removing the tilt.

The most interesting application of E^* in ZEPLIN-II was the resolution of the 40 keV_{ee} feature from inelastic neutron scattering off ^{129}Xe . As mentioned earlier, when an AmBe source is placed near the detector, populations of events induced by neutron interactions appear. Pure elastic recoils are the dominant population (being our calibration for an expected WIMP signal), but inelastic scattering populations are also present. These are observed with a gamma of a certain energy accompanied by a recoil component which adds to the total energy deposition (albeit quenched). In xenon, the two main inelastics observed are at 40 keV_{ee} (from ^{129}Xe) and 80 keV_{ee} (from ^{131}Xe).

The presence of these inelastic populations is predicted by detailed simulations of the detector and observed in the data. However, using S1-determined energy it is not possible to resolve the inelastic features (Figure 7.20, top). The feature around 40 keV_{ee} is studied as the higher energy feature is weaker and suffers from more contamination by the electron recoil population. To study the 40 keV feature we must first exclude the pure elastic scattering population as it dominates the part of the energy spectrum we are interested in, this is done graphically. Figure 7.20 shows the energy spectrum observed from S1 alone, exhibiting no peak, and the improvement in resolution gained by implementing E^* . The peak becomes clearly resolved, at least in part due to the improved resolution, and in good agreement with the simulated prediction. The only significant discrepancy from the simulation is an enhanced rate of inelastic scatters. This deficit in the GEANT4 simulations had previously been observed by other experiments suggesting the inelastic cross-sections implemented in GEANT4 for xenon may be too low by a factor of ~ 2 .

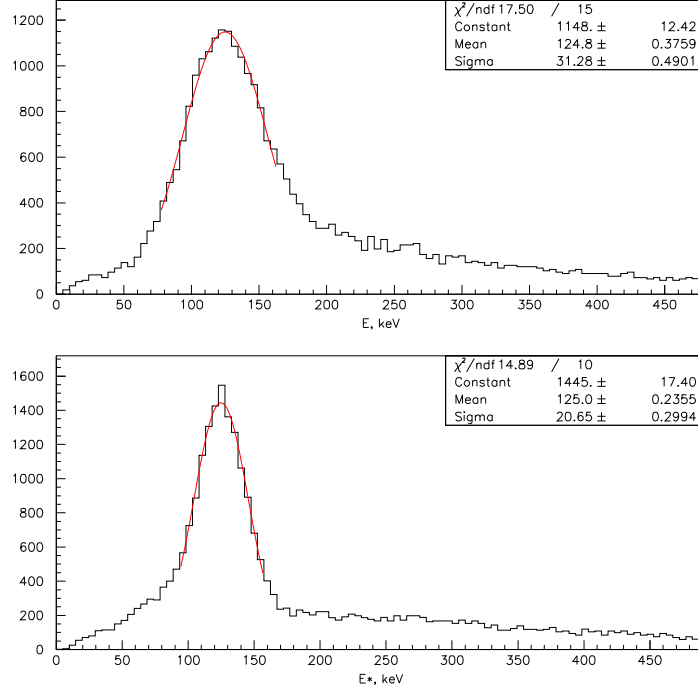


Figure 7.18: Energy spectrum for ^{57}Co in S1 (top) and E^* (bottom) with the combined 122/136 keV photopeak fitted. The energy scale was calibrated by the position of the photopeak.




Figure 7.19: Colour density plots of $S2/S1$ against energy, for ^{57}Co (left) and ^{60}Co (right) with energy described by $S1$ (top) and E^* (bottom).

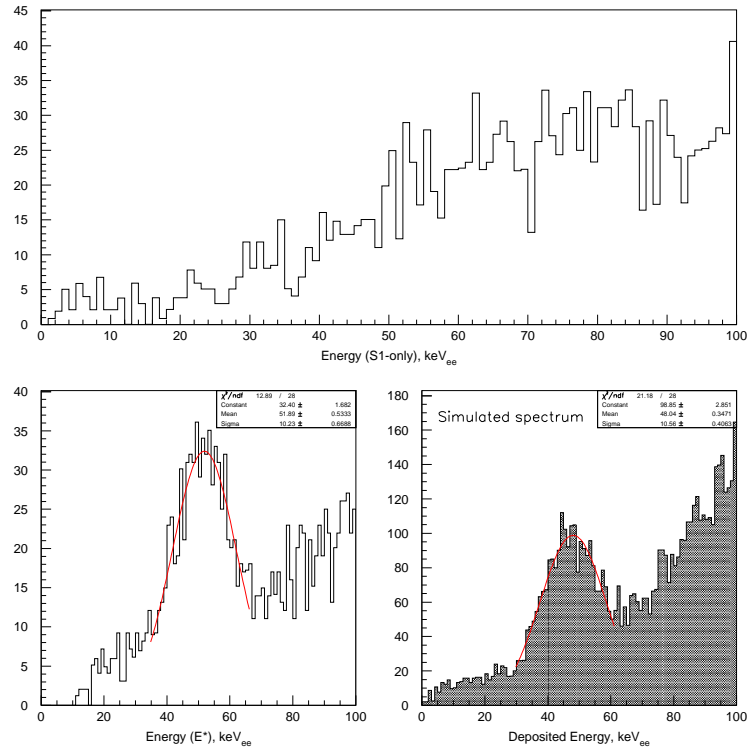


Figure 7.20: Plots showing the 40 keV_{ee} inelastic feature in energy measured in S1 only (top) and E* (left) compared with the simulated spectrum from GEANT4 (right). Note that the mean energy deposit is > 40 keV due to the coincident nuclear recoil contribution.

Despite the clear improvement in energy resolution achieved using E*, it was not applied in dark matter search analysis for ZEPLIN-II. To remove the anti-correlation effect correctly, a different combination coefficient should be applied for nuclear and electron recoils, as the physical processes determine different splitting of deposited energy in each (and there is a non-electronic energy deposit in the case of nuclear recoils). As we cannot say categorically if an event is an electron or nuclear recoil in our background dataset (for each event in isolation), we cannot apply different coefficients. Consequently, applying E* instead of S1 in our final scatter plots adds no new information, meaning no improvement in discrimination. Further studies may show that the improvement in resolution could slightly enhance the final dark matter result (see Chapter 6), although applying a coefficient determined from electron recoils will effect the energy resolution for nuclear recoils to a different extent.

During the commissioning phase of ZEPLIN-III the same method of calculating E*

was tested, providing an improvement in resolution (σ/mean) from $\sim 17\%$ for S1 or S2 individually to $\sim 12\%$ for E^* , at an electric field of ~ 2.5 kV/cm in the liquid.

A new method was devised to optimise the coefficient k to be applied. This optimisation method relies on the resolution of the photopeak being partially caused by the anti-correlation effect. We proceed by varying the combination coefficient and fitting the resulting E^* photopeak with a simple gaussian, yielding a plot (Figure 7.21) of resolution against k . For very small values of k the resolution tends to that of S1 alone (as the S2 contribution is negligible) and the reverse occurs for large k as S2 alone dominates. In the intervening range, the resolution improves as the anti-correlation is corrected, with the best resolution allowing optimisation of k . For calibration data acquired at ~ 2.5 kV/cm in the liquid, a combination coefficient approximately equal to the ratio of the peaks is found to be optimal. This improvement follows the form expected by considering the resolution of the combined parameter:

$$\sigma^2(S1 + k S2) = \sigma^2(S1) + \sigma^2(k S2) + 2 \times \text{Cov}(S1, k S2) \quad (7.6)$$

$$\sigma^2(E^*) = \sigma^2(S1) + k \sigma^2(S2) + 2k (\langle S1 \times S2 \rangle - \langle S1 \rangle \langle S2 \rangle), \quad (7.7)$$

where $\langle S1 \rangle$ and $\langle S2 \rangle$ are the means of the S1 and S2 photopeaks from ^{57}Co , $\sigma(S1)$ and $\sigma(S2)$ are the corresponding energy resolutions and $\langle S1 \times S2 \rangle$ is the mean of the product S1 and S2. The predicted form for the curve is shown in Figure 7.21, showing reasonable agreement with the calibration data.

Before further investigation of the coefficient k , the combination process, accounting for the physical processes at work, should be considered. When calibrating the detectors for energy, S1 alone is usually utilised, but this is not physically correct as the energy deposited during the interaction is split between excitation and ionisation. Strictly speaking, energy calibration should utilise:

$$E^* = (S1_{\text{cal}}^* \times S1) + (S2_{\text{cal}}^* \times S2) \text{ where } S1_{\text{cal}}^* = \frac{\eta_1 \times (125 \text{ keV})}{\langle S1 \rangle} \quad (7.8)$$

with an equivalent for $S2_{\text{cal}}^*$, calibrating the photopeak from ^{57}Co with only a fraction of the 125 keV deposited. Figure 7.22 describes the process leading to the fractions η_1 and η_2 . η_r , the fraction of electrons recombining, is clearly dependent on the electric




Figure 7.21: Plot of the relative energy resolution obtained from the combined 122/136 keV peak in ^{57}Co as a function of the coefficient k assumed. The blue line shows the expected form of the improvement in resolution. The red and green lines show the resolutions of S1 and S2 alone.

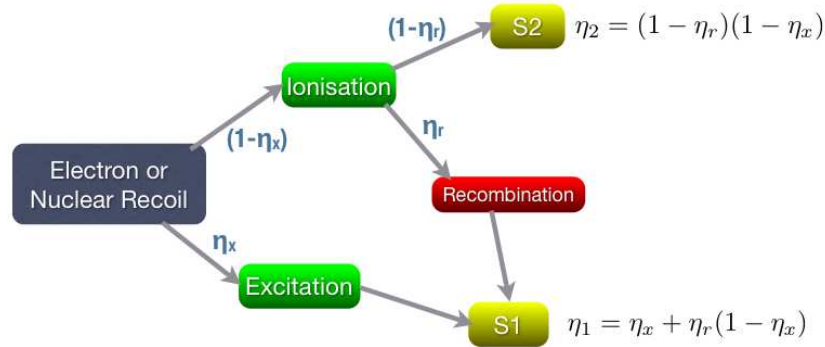


Figure 7.22: Flow diagram showing the processes involved in the separation of deposited energy.

field applied, suggesting that an increase in field would lead to a larger η_2 .

It is difficult to directly compare data-sets with different experimental parameters (e.g. pressure, liquid level or Xe purity) as these will affect the response of S2, varying the resolution. During testing of ZEPLIN-III, two comparable data-sets were acquired with different applied fields. Figure 7.23 shows the results of an optimisation achieved

by varying η_1 (and assuming $\eta_1 + \eta_2 = 1$) for both data-sets. As before the resolution is seen to improve in the mid-range with the pure S1 and S2 resolutions approached for very small S1 or S2 contributions. As expected the data-set with higher applied field is found to have a smaller η_1 , with recombination inhibited by stronger drifting of charge.




Figure 7.23: Plot of energy resolution (measured as σ/mean) of the combined 122/136 keV peak in ^{57}Co as a function of η_1 , the fraction of energy from S1. The red line describes data taken with 11 kV applied (a field of 2.5 kV/cm) and the blue line is 16 kV (3.7 kV/cm).

This effect is very interesting for both understanding the physical processes involved and for potentially enhancing the final dark matter results from such detectors. Dedicated ^{57}Co calibrations will be taken at different fields allowing for the field-dependence of the combination coefficient to be mapped. Further studies will also assess the possibility of applying these coefficients to the nuclear recoil population, possibly calculating a nuclear recoil specific coefficient.

By studying the different physical processes in the data, a clear understanding of the detector response can be achieved, aiding the analysis of the data and helping to determine features such as acquisition and field settings. Further studies will be carried out with data from ZEPLIN-III including specifically tailored data-sets aiming to assess how these effects can be utilised to improve dark matter sensitivity.

Chapter 8

Conclusion

A large body of compelling evidence has built up for the existence of a significant non-baryonic cold dark matter component in the Universe. As a result, many different techniques are being applied in a race to make the first direct detection. One of the most promising approaches is the use of two-phase noble gas detectors measuring both the scintillation and the ionisation response channels. The UK-led ZEPLIN collaboration has applied this two-phase technique in xenon with two different detectors, ZEPLIN-II and ZEPLIN-III.

Both experiments were operated at the Boulby Underground Laboratory, where 1100m of rock overburden and the low-radioactivity salt-seam location provide an excellent environment for low background studies. We have described the successful operation of both instruments and the complete data analysis chains leading to the two dark matter results. The ZEPLIN-II detector operated underground for 20 months, collecting WIMP search data for 57 days, yielding a final fiducial exposure of 225 kg·days. Despite being limited by an unexpected background of ^{222}Rn -progeny events, it posted a competitive upper limit on the spin-independent WIMP-nucleon cross-section with a minimum at 6.6×10^{-7} pb for a 65 GeV/c² WIMP. Figure 8.1 shows the constraints which had been placed on WIMP-nucleon scattering before October 2005 and the current day, showing the huge progress made in the past few years.

Following the operation of ZEPLIN-II, ZEPLIN-III was commissioned underground. Utilising its different detector geometry and high-field technology it operated for 12 months in its initial underground configuration, collecting WIMP search data for 83 days, with

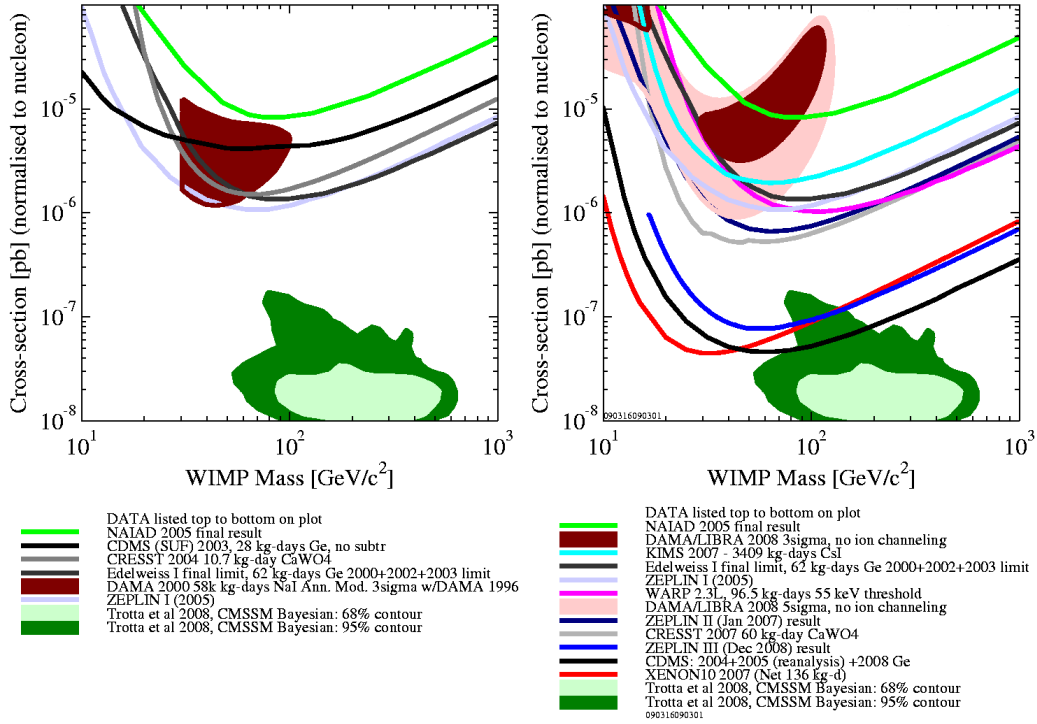


Figure 8.1: The world status of constraints on spin-independent WIMP-nucleon scattering at the start of this PhD research, October 2005 (left), and at the end, July 2009 (right).

a final fiducial exposure of 267.9 kg-days. The data from its first science run placed a limit on the WIMP-nucleon cross-section with a minimum at 7.7×10^{-8} pb for a 55 GeV/c² WIMP. This result placed ZEPLIN-III at the same level as the two other world-leading experiments, XENON10 and CDMS-II.

Both detectors yielded additional interesting measurements of the xenon physics processes underpinning their response to radiation. Most notably, we were able to make the first quantitative measurement of single electron emission in a two-phase noble gas detector, proving the extraordinary sensitivity of this technique in the ionisation channel. We were also able to measure the behaviour of the S1 and S2 responses as a function of electric field, including the level of their anti-correlation.

8.1 Lessons learned from the ZEPLIN detectors

The operation and analysis of the ZEPLIN detectors have not only placed interesting and competitive constraints on WIMP-nucleon scattering cross-sections, they have also yielded a wealth of knowledge about two-phase experimental techniques. This knowledge will be vital for the design of future large-scale detectors.

A particular success for both instruments (but for different reasons) was the xenon purity. After initial difficulties, ZEPLIN-II achieved an excellent electron lifetime through active xenon recirculation. ZEPLIN-III achieved a reasonable electron lifetime through bottle-to-bottle transfers, but the excellent cleanliness, xenon-friendly materials and strong electric field sustained (and in fact improved) the xenon purity through the science run. This highlights that both factors, cleanliness and recirculation, will be important in much larger future systems. In ZEPLIN-II, the radon-progeny events which limited the science result were introduced into the system by the purification getters. Next-generation detectors must avoid this, using either (ideally) radon-free getters or by implementing a radon trap before the target.

The geometry of the two targets was very distinct, each with their own advantages. The flat, disk-like geometry of ZEPLIN-III, with the PMT array located in the liquid, produced a better light yield than that achieved in ZEPLIN-II. However, the deeper ZEPLIN-II geometry, with a reflective PTFE cone defining the target region, produced a more uniform light collection throughout the instrument and allowed for a larger target mass. Two major lessons have been learned from the ZEPLIN two-phase targets: that any detector dead regions will produce MSSl events which can potentially limit sensitivity, and that the self-shielding afforded by a larger volume of xenon greatly aids in background reduction. The greater volume and deeper geometry of the ZEPLIN-II target meant that the background rate in the centre of the detector was in fact lower than that in ZEPLIN-III, despite the more careful consideration of materials in the latter. In fact, as a result of the self-shielding afforded by the xenon, the instrument still managed to produce a respectable dark matter result, despite the unexpected background of surface events.

One of the main features of ZEPLIN-III was its use of a much stronger electric field than in similar experiments. This significantly improved the discrimination between

electron and nuclear recoils. Ideally, a similar electric field should be exploited in future systems, although the difficulty in applying the large voltages required over much greater distances is likely to make this impractical. Even ZEPLIN-III suffered initial difficulties with the nominal design voltages, although a small re-design enabled stable fields to be maintained throughout data-taking. Another consideration is that the applied electric field should be uniform throughout the target region, as variations in field have consequences for the response of both the S1 and S2 channels. Finally, we have demonstrated that two-phase xenon systems can be operated stably over long periods, a doubt that had been expressed by some in the field.

8.2 Future work

ZEPLIN-III is currently undergoing an upgrade which will facilitate a year-long, low-background second run, projecting a sensitivity of $\sigma_{w-n} \leq 10^{-8}$ pb. This upgrade will include the replacement of the original phototubes, which dominated the background in the first science run, with new ultra-low background ones, aiming to reduce the background rate by more than an order of magnitude. An active veto, in the final stages of construction, is to be installed around the instrument, aiming to reject neutron background by detecting scatters in coincidence with events in the target. An active recirculation system has been commissioned to circulate xenon through the system before filling, in order to speed up the cleaning of both the xenon gas and the target surfaces before the second run. Additional systems are also being developed to allow automated operation of the instrument to reduce the operational burden and further improve the duty cycle of an extended run. In addition to the dark matter search, more xenon physics measurements will be carried out to investigate the field dependence of the response channels and the discrimination power, along with further work on single electron emission. Perhaps more critically, we intend to make new measurements of the energy-dependent quenching factor for nuclear recoils, either using ZEPLIN-III itself or dedicated prototypes at a neutron beam.

In the field of direct searches at large, the XENON collaboration are currently commissioning XENON100, extending the basic design of XENON10 to a larger scale system with a fiducial mass of ~ 65 kg (from a total xenon mass of 170 kg), projecting an order of

magnitude improvement over the XENON10 result, $\sigma_{w-n} \sim 2 \times 10^{-9}$ pb. Similarly, the LUX collaboration are constructing two-phase xenon systems to be deployed at Homestake. CDMS-II are due to publish on more data shortly, with Super-CDMS in the early stages of new detector manufacture. Of the other experiments, XMASS is currently under construction, projecting a limit of $\sigma_{w-n} = 10^{-10}$ pb using single-phase xenon, whilst the DEAP-CLEAN collaboration also have funding for a large-scale detector to be located at SNOLab.

Following the second science run of ZEPLIN-III, it is planned that the collaboration will join with LUX to form LUX-ZEPLIN, with the aim to produce tonne-scale detectors. Planning and design work has begun for LZ3, a 3-tonne target to be deployed at the Stanford Underground Science Laboratory (SUSEL) at the Homestake Mine (South Dakota, USA), with design studies being planned for the even larger LZ20, a 20-tonne system to be deployed at the new Deep Underground Science Laboratory (DUSEL), also at the Homestake Mine.

The prospects of a detection within the next 5–10 years look promising, with the scaling up of existing technologies complementing expected results from the LHC. The Large Hadron Collider will provide clues as to the nature of the dark matter by searching for evidence of supersymmetry and constraining the masses of any newly discovered candidate particles. The evidence from the LHC will allow WIMP search experiments to be refined and the analyses optimised to look in the expected part of the mass range. However, large-scale WIMP search experiments will be required to confirm that any newly discovered particle indeed solves the galactic dark matter problem.

Appendix A

Appendix A: Calibration sources

The calibration of dark matter detectors requires their exposure to radioactive sources, mimicking the γ -ray (electron recoil) background or the expected nuclear recoil spectrum from WIMP or neutron elastic scatters or neutron background interactions (nuclear recoil). γ -ray sources are also utilised to calibrate the energy scale in such detectors (as described in Sections 4.6 and 5.10), as well as providing a diagnostic tool for monitoring of detector characteristics. The photon interaction lengths and types for the calibration sources utilised are shown in Figure A.1.




Figure A.1: Interaction length as a function of energy for γ -rays, with the γ -ray calibration sources utilised in this thesis marked.

A.0.1 Cobalt-57 γ -ray source

^{57}Co is utilised as the primary mode of energy calibration. The relatively low, dominant γ -rays at 122 keV and 136 keV are just able to penetrate through the walls of the target vessels and interact in the first few cm of liquid xenon. The decay scheme for ^{57}Co is shown in Figure A.2 with the γ -ray energies given in Table A.1.





Figure A.2: The decay scheme of ^{57}Co . [1]

Table A.1: γ -rays emitted in the decay of ^{57}Co . [1]

γ -ray energy (keV)	Relative intensity (%)
14.41	9.16
122.06	85.60
136.47	10.68
230.4	0.0004
339.69	0.0037
352.33	0.0030
366.8	0.0012
570.09	0.0158
692.41	0.149
706.54	0.0050

A.0.2 Cobalt-60 γ -ray source

Cobalt-60 is utilised for calibration of the expected electron recoil background within the instrument. Cobalt-60 undergoes β -decay, where the relatively low energy β -particles are easily shielded, resulting in excited states of ^{60}Ni . This then leads to the emission of two cascaded γ -rays, 1.17 MeV and 1.33 MeV, which are able to populate the background region through low-angle Compton scattering in the liquid xenon. The decay scheme and table of gamma-ray energies for ^{60}Co are shown in Figure A.3 and Table A.2.

Figure A.3: The decay scheme of ^{60}Co . [1]Table A.2: γ -rays emitted in the decay of ^{60}Co . [1]

γ -ray energy (keV)	Relative intensity (%)
346.93	0.0076
826.28	0.0076
1173.237	99.9736
1332.501	99.9856
2158.77	0.00111
2505	2.0E-6

A.0.3 Caesium-137 γ -ray source

Caesium-137 is also utilised for calibration of the electron recoil backgrounds. It β -decays to a metastable state of barium-137 with a half-life of 30.07 years. The β -particles are easily shielded, but the ^{137m}Ba decays to its ground state with a half-life of 2.55 minutes, emitting γ -rays of 662 keV. These can then be used to populate the electron recoil background region through low-angle Compton scattering in the liquid xenon. The decay scheme and table of gamma-ray energies for ^{137}Cs are shown in Figure A.4 and Table A.3.



Figure A.4: The decay scheme of ^{137}Cs . [1]

Table A.3: γ -rays emitted in the decay of ^{137}Cs . [1]

γ -ray energy (keV)	Relative intensity (%)
283.5	5.8E-4
661.657	85.1

A.0.4 AmBe neutron source

The expected interactions between WIMPs and baryonic matter produce nuclear recoils in the target. As a result, the expected detector response can be calibrated by exposing the instrument to a source of neutrons, also producing nuclear recoils. The neutron source chosen for calibration of the ZEPLIN-II and ZEPLIN-III detectors was an AmBe source. Americium-241 decays with a half-life of 432.2 years, emitting α -particles with an average energy of 5.47 MeV. 59.5 keV γ -rays and 13.9 keV x-rays are also emitted in the alpha decay. By mixing ^{241}Am with beryllium oxides, a neutron source can be constructed, producing neutrons through an (α, n) reaction.



These neutrons have energies ranging from 0 to about 11 MeV, with an average of about 4-5 MeV, with the expected spectral shape from an AmBe source shown in Figure A.5.




Figure A.5: Measured neutron energy spectrum from a 370 GBq AmBe neutron source. [202]

Bibliography

- [1] Korea Atomic Energy Research Institute. <http://atom.kaeri.re.kr/>
- [2] Bertone, G, Hooper, D & Silk, J. *Physics Reports*, 405 279–390 (2005)
- [3] Freedman, W, Madore, B, Gibson, B *et al.* *The Astrophysical Journal*, 553 47–72 (2001)
- [4] Peacock, J. *Cosmological Physics (Cambridge University Press)* (1999)
- [5] Roos, M. *Introduction to Cosmology (Wiley)* (1997)
- [6] Perlmutter, S, Turner, M & White, M. *Physical Review Letters*, 83(4) 670–673 (1999)
- [7] Knop, R, Aldering, G, Amanullah, R *et al.* *The Astrophysical Journal*, 598 102–137 (2003)
- [8] Komatsu, E, Dunkley, J, Nolte, MR *et al.* *The Astrophysical Journal Supplement*, 180 330–376 (2009)
- [9] Oort, J. *Bull. Astron. Inst. Neth*, 6 249–287 (1932)
- [10] Zwicky, F. *Helvetica Physica Acta*, 6 110–127 (1933)
- [11] Kuijken, K & Gilmore, G. *Monthly Notices of the Royal Astronomical Society*, 239(33) 651–664 (1989)
- [12] Bienayme, O, Robin, A & Creze, M. *Astronomy and Astrophysics*, 180 94–110 (1987)
- [13] Geller, M, Diaferio, A & Kurtz, M. *The Astrophysical Journal Letters*, 517 L23–L26 (1999)

- [14] Mohr, J, Mathiesen, B & Evrard, A. *The Astrophysical Journal*, 517 627–649 (1999)
- [15] Zwicky, F. *Physical Review*, 51 290 (1937)
- [16] Kahn, FD & Woltjer, L. *The Astrophysical Journal*, 130(3) 705–717 (1959)
- [17] Zaritsky, D, Olszewski, EW, Schommer, RA *et al.* *The Astrophysical Journal*, 345 759–769 (1989)
- [18] Ostriker, J & Peebles, P. *The Astrophysical Journal*, 186 467–480 (1973)
- [19] Babcock, H. *Lick Observatory Bulletin*, 498 41–51 (1939)
- [20] Rubin, V & Ford, W. *The Astrophysical Journal*, 159 379–403 (1970)
- [21] Roberts, M & Whitehurst, R. *The Astrophysical Journal*, 201 327–346 (1975)
- [22] Rubin, V, Thonnard, N & Ford, W. *The Astrophysical Journal*, 238 471–487 (1980)
- [23] Van Albada, T, Bahcall, J, Begeman, K *et al.* *The Astrophysical Journal Letters*, 295(2) 305–313 (1985)
- [24] Fabricant, D, Lecar, M & Gorenstein, P. *The Astrophysical Journal*, 241 552–560 (1980)
- [25] Saglia, R, Bertin, G, Bertola, F *et al.* *The Astrophysical Journal*, 403(2) 567–572 (1993)
- [26] Forman, W, Jones, C & Tucker, W. *The Astrophysical Journal*, 293 102–119 (1985)
- [27] Clowe, D, Bradač, M, Gonzalez, AH *et al.* *The Astrophysical Journal*, 648 L109–L113 (2006)
- [28] Roque, SL, Bonamente, M, Carlstrom, J *et al.* *The Astrophysical Journal*, 652(2) 917–936 (2006)
- [29] Chae, K, Biggs, A, Blandford, R *et al.* *Physical Review Letters*, 89 151301 (2002)
- [30] Massey, R, Rhodes, J, Ellis, R *et al.* *Nature*, 445 286–290 (2007)
- [31] Rhodes, JD. *American Astronomical Society Meeting*, 207 1309 (2005)

- [32] Paczynski, B. *Annual Review of Astronomy and Astrophysics*, 34 419–59 (1996)
- [33] Gould, A. *Microlensing 2000: A New Era of Microlensing Astrophysics*, 239 3–17 (2001)
- [34] Alcock, C, Akerlof, CW, Allsman, RA *et al.* *Nature*, 365 621 (1993)
- [35] Tisserand, P, Guillou, LL, Afonso, C *et al.* *Astronomy and Astrophysics*, 469(2) 387–404 (2007)
- [36] Alpher, R, Gamow, G & Herman, R. *Proceedings of the National Academy of Sciences* (1967)
- [37] Penzias, AA & Wilson, RW. *The Astrophysical Journal*, 142 419–421 (1965)
- [38] Dicke, RH, Peebles, PJE, Roll, PG *et al.* *The Astrophysical Journal*, 142 414–419 (1965)
- [39] Hinshaw, G, Weiland, JL, Hill, RS *et al.* *The Astrophysical Journal Supplement*, 180 225–245 (2009)
- [40] Mather, J, Fixsen, D, Shafer, R *et al.* *The Astrophysical Journal*, 512 511–520 (1999)
- [41] Jones, W, Ade, P, Bock, J *et al.* *The Astrophysical Journal*, 647 823–832 (2006)
- [42] Jaffe, A, Abroe, M, Borrill, J *et al.* *New Astronomy Reviews*, 47 727–732 (2003)
- [43] Nolta, MR, Dunkley, J, Hill, RS *et al.* *The Astrophysical Journal Supplement*, 180 296–305 (2009)
- [44] Reichardt, C, Ade, P, Bock, J *et al.* *The Astrophysical Journal*, 694 1200 (2009)
- [45] Readhead, A, Mason, B, Contaldi, C *et al.* *The Astrophysical Journal*, 609 498–512 (2004)
- [46] Dunkley, J, Komatsu, E, Nolta, MR *et al.* *The Astrophysical Journal Supplement*, 180 306–329 (2009)
- [47] Coc, A, Vangioni-Flam, E & Descouvemont, P. *The Astrophysical Journal*, 600 544–552 (2004)
- [48] Boerner, G. *The Early Universe (Berlin: Springer)* (1988)

- [49] Klapdor-Kleingrothaus, H & Zuber, K. *Particle Astrophysics (IoP Publishing)* (2000)
- [50] Rolnick, W. *The Fundamental Particles and Their Interactions (Addison-Wesley Publishing)* (1994)
- [51] Halzen, F & Martin, A. *Quarks and Leptons: An Introductory Course in Modern Particle Physics (NY Chichester Wiley)* (1984)
- [52] Gildener, E. *Physical Review D*, 14(6) 1667–1672 (1976)
- [53] Weinberg, S. *Physical Review D*, 13(4) 974–996 (1976)
- [54] Weinberg, S. *Physical Review D*, 19(4) 1277–1280 (1979)
- [55] Murayama, H. *arXiv:hep-ph/0002232* (2000)
- [56] Collins, PDB, Martin, AD & Squires, EJ. *Particle physics and cosmology (Wiley)*, 496 (1989)
- [57] Wess, J & Bagger, J. *Supersymmetry and Supergravity (Princeton University Press)*, 259 (1992)
- [58] Jungman, G, Kamionkowski, M & Griest, K. *Physics Reports*, 267 195–373 (1996)
- [59] Martin, S. *hep-ph/9709356* (1997)
- [60] Wess, J & Zumino, B. *Nuclear Physics B*, 70 39–50 (1974)
- [61] Ellis, J, Kelley, S & Nanopoulos, D. *Physics Letters B*, 260 131 (1991)
- [62] Mohapatra, R. *Physical Review D*, 34(11) 3457–3461 (1986)
- [63] Mohapatra, R & Pal, P. *Massive Neutrinos in Physics and Astrophysics (World Scientific)* (1991)
- [64] Kane, G, Kolda, C, Roszkowski, L *et al.* *Physical Review D*, 49(11) 6173 (1994)
- [65] Ellis, J, Olive, K, Santoso, Y *et al.* *Physics Letters B*, 565 176–182 (2003)
- [66] ALEPH collaboration. *Physics Letters B*, 499 67–84 (2001)
- [67] Fukuda, Y, Hayakawa, T, Ichihara, E *et al.* *Physical Review Letters*, 81(8) 1562–1567 (1998)

- [68] Maltoni, M, Schwetz, T, Tortola, M *et al.* *Physical Review D*, 68 113010 (2003)
- [69] Bergstrom, L. *Reports on Progress in Physics* (2000)
- [70] Lobashev, V. *Nuclear Physics A*, 719 153–160 (2003)
- [71] Bond, J, Efstathiou, G & Silk, J. *Physical Review Letters*, 45(24) 1980–1984 (1980)
- [72] Davis, M, Summers, F & Schlegel, D. *Nature*, 359 393–396 (1992)
- [73] Peccei, RD & Quinn, HR. *Physical Review Letters*, 38 1440 (1977)
- [74] Weinberg, S. *Physical Review Letters*, 40(4) 223–226 (1978)
- [75] Wilczek, F. *Physical Review Letters*, 40(5) 279–282 (1978)
- [76] Rosenberg, L & van Bibber, K. *Physics Reports*, 325 1–39 (2000)
- [77] van Bibber, K & Rosenberg, L. *Physics Today*, 30 (August 2006)
- [78] Andriamonje, S, Aune, S, Autiero, D *et al.* *Journal of Cosmology and Astroparticle Physics*, 04 10 (2007)
- [79] Arik, E, Aune, S, Autiero, D *et al.* *Journal of Cosmology and Astroparticle Physics*, 02 008 (2009)
- [80] Zavattini, E, Zavattini, G, Ruoso, G *et al.* *Physical Review Letters*, 96 110406 (2006)
- [81] Zavattini, E, Zavattini, G, Ruoso, G *et al.* *Physical Review D*, 77 32006 (2008)
- [82] Kim, J. *Physics Reports*, 150(1-2) 1–177 (1987)
- [83] Turner, M. *Physics Reports*, 197 167 (1990)
- [84] Raffelt, G. *Physics Reports*, 198(1-2) 1–113 (1990)
- [85] de Boer, W, Sander, C, Zhukov, V *et al.* *Physics Letters B*, 636(1) 13–19 (2006)
- [86] Gehrels, N & Michelson, P. *Astroparticle Physics*, 11 277–282 (1999)
- [87] Winkler, C. *New Astronomy Reviews*, 50(7-8) 530–533 (2006)
- [88] Albert, J, Aliu, E, Anderhub, H *et al.* *The Astrophysical Journal*, 639 761–765 (2006)

- [89] Adriani, O, Barbarino, GC, Bazilevskaya, GA *et al.* *Nature*, 458(7238) 607–609 (2009)
- [90] Chang, J, Adams, JH, Ahn, HS *et al.* *Nature*, 456(7220) 362–365 (2008)
- [91] Profumo, S. *arXiv:astro-ph/0812.4457* (2008). Submitted to Phys. Rev. D
- [92] Lewin, J & Smith, P. *Astroparticle Physics*, 6(1) 87–112 (1996)
- [93] Bettini, A. *Journal of Physics: Conference Series*, 120(8) 082001 (2008)
- [94] Coccia, E. *Journal of Physics: Conference Series*, 39 497–504 (2006)
- [95] Morales, J & Beltran, B. *The Identification of Dark Matter - 5th International Workshop* (2005)
- [96] Davies, G, Davies, J, Lewin, J *et al.* *Physics Letters B* (1994)
- [97] Alner, G, Araújo, H, Arnison, G *et al.* *Astroparticle Physics*, 23 444–462 (2005)
- [98] Abe, K. *Journal of Physics: Conference Series* (2008)
- [99] Hime, A. *Intersections of Particle and Nuclear Physics: 9th Conference CIPAN2006 (AIP Conference Proceedings)*, 205 (2006)
- [100] Pasuthip, P. *Masters Thesis (Queen's University, Canada)* (2009)
- [101] Alner, G, Araújo, H, Bewick, A *et al.* *Astroparticle Physics*, 28 287–302 (2007)
- [102] Akimov, D, Alner, G, Araújo, H *et al.* *Astroparticle Physics*, 27 46–60 (2007)
- [103] Lebedenko, V, Araújo, H, Barnes, E *et al.* *Physical Review D*, 80(5) (2009)
- [104] Angle, J, Aprile, E, Arneodo, F *et al.* *Physical Review Letters*, 100(2) 21303 (2008)
- [105] WARP Collaboration. *Astroparticle Physics*, 28 495 (2008)
- [106] Abrams, D, Akerib, D, Armel-Funkhouser, M *et al.* *Physical Review D*, 66 122003 (2002)
- [107] Ahmed, Z, Akerib, DS, Arrenberg, S *et al.* *Physical Review Letters*, 102 11301 (2009)
- [108] Benoit, A, Bergé, L, Broniatowski, A *et al.* *Physics Letters B*, 513 15–22 (2001)

- [109] Seidel, W, Angloher, G, Bauer, M *et al.* *Journal of Physics: Conference Series*, 120 (2008)
- [110] Martineau, O, Benoit, A, Bergé, L *et al.* *Nuclear Inst. and Methods in Physics Research, A*, 530(3) 426–439 (2004)
- [111] Aubin, F, Auger, M, Azuelos, G *et al.* *MEDEX 07 - AIP Conference Proceedings* (2007)
- [112] Giuliani, F, Morlat, T, Ramos, A *et al.* *Nuclear Physics B (Proceedings Supplements)*, 173 129–132 (2007)
- [113] Archambault, S, Aubin, F, Auger, M *et al.* *arXiv:0907.0307* (2009)
- [114] Bernabei, R, Belli, P, Bussolotti, A *et al.* *Nuclear Inst. and Methods in Physics Research, A*, 592(3) 297–315 (2008)
- [115] Bernabei, R, Belli, P, Cappella, F *et al.* *The European Physical Journal C-Particles and Fields*, 56(3) 333–355 (2008)
- [116] Tucker-Smith, D & Weiner, N. *Physical Review D*, 64(4) 3502 (2001)
- [117] Alner, G, Araujo, H, Ayad, R *et al.* *Nuclear Inst. and Methods in Physics Research*, 535 644–655 (2004)
- [118] Alner, G, Araújo, H, Bewick, A *et al.* *Nuclear Inst. and Methods in Physics Research*, 555 173–183 (2005)
- [119] Sciolla, G, Ahlen, S, Dujmic, D *et al.* *arXiv:astro-ph/0805.2431* (2008)
- [120] Nishimura, H, Miuchi, K, Hattori, K *et al.* *Astroparticle Physics*, 31 185–191 (2009)
- [121] Mayet, F, Guillaudin, O, Grignon, C *et al.* *arXiv:astro-ph/0905.4654* (2009)
- [122] de Austri, R, Trotta, R & Roszkowski, L. *J HEP*, 05 002 (2006)
- [123] Trotta, R, de Austri, R & Roszkowski, L. *New Astronomy Reviews* (2007)
- [124] R.Gaitskell & V.Mandic. <http://dendera.berkeley.edu/plotter/entryform.html>
- [125] Smith, P, Smith, N, Lewin, J *et al.* *Physics Reports*, 307 275–282 (1998)
- [126] Ahmed, B, Alner, G, Araújo, H *et al.* *Astroparticle Physics*, 19 691–702 (2003)

- [127] Alner, G, Araújo, H, Arnison, G *et al.* *Physics Letters B*, 616(1-2) 17–24 (2005)
- [128] Jortner, J, Meyer, L, Rice, S *et al.* *The Journal of Chemical Physics*, 42(12) 4250–4253 (2004)
- [129] Edwards, B, Araújo, H, Chepel, V *et al.* *Astroparticle Physics*, 30 54–57 (2008)
- [130] Doke, T & Masuda, K. *Nuclear Instruments and Methods in Physics Research A*, 420 62–80 (1999)
- [131] Chepel, V, Lopes, M & Solovov, V. *Radiation Physics and Chemistry*, 74 160–167 (2005)
- [132] Hutchinson, G. *Nature*, 162 610–611 (1948)
- [133] Dolgoshein, B, Lebedenko, V & Rodionov, B. *JETP Letters*, 11(11) 351–353 (1970)
- [134] Bolozdynya, A. *Nuclear Instruments and Methods in Physics Research A*, 422 314–320 (1999)
- [135] Gushchin, E. *Sov. J. Experim. and Theor. Phys.*, 49 856 (1979)
- [136] Kubota, S, Nakamoto, A, Takahashi, T *et al.* *Physical Review B*, 17(6) 2762–2765 (1978)
- [137] Hitachi, A, Takahashi, T, Funayama, N *et al.* *Physical Review B*, 27(9) 5279–5285 (1983)
- [138] Kubota, S, Hishida, M & Raun, J. *J. Phys. C*, 11 2645–2651 (1978)
- [139] Chepel, V, Solovov, V, Neves, F *et al.* *Astroparticle Physics*, 26(1) 58–63 (2006)
- [140] Aprile, E, Baudis, L, Choi, B *et al.* *Physical Review C*, 79(4) 45807 (2009)
- [141] Aprile, E, Giboni, K, Majewski, P *et al.* *Physical Review D*, 72 072006 (2005)
- [142] Akimov, D, Bewick, A, Davidge, D *et al.* *Physics Letters B*, 524 245–251 (2002)
- [143] Gokalp, A, Bernabei, R, Belli, P *et al.* *EPJ direct*, C11 1–8 (2001)
- [144] Arneodo, F, Baiboussinov, B, Badertscher, A *et al.* *Nuclear Inst. and Methods in Physics Research A*, 449 147–157 (2000)

-
- [145] Sorensen, P, Manzur, A, Dahl, C *et al.* *Nuclear Inst. and Methods in Physics Research A*, 601 339–346 (2009)
- [146] Hitachi, A. *Astroparticle Physics*, 24 247–256 (2005)
- [147] Aprile, E, Dahl, C, de Viveiros, L *et al.* *Physical Review Letters*, 97 081302 (2006)
- [148] Miller, L, Howe, S & Spear, W. *Physical Review*, 166(1) 871–878 (1968)
- [149] Bakale, G, Sowada, U & Schmidt, W. *The Journal of Physical Chemistry*, 80(23) 2556 (1976)
- [150] Araújo, H, Akimov, D, Alner, G *et al.* *Astroparticle Physics*, 26 140–153 (2006)
- [151] Fonseca, A, Meleiro, R, Chepel, V *et al.* *Nuclear Science Symposium Conference Record*, 1 572–576 (2004)
- [152] Monteiro, C, Fernandes, L, Lopes, J *et al.* *Journal of Instrumentation*, 2 P05001 (2007)
- [153] Kudryavtsev, VA, Lightfoot, PK, McMillan, JE *et al.* *arXiv:hep-ex/0301038* (2003)
- [154] Robinson, M, Kudryavtsev, V, Lüscher, R *et al.* *Nuclear Inst. and Methods in Physics Research, A*, 511(3) 347–353 (2003)
- [155] Svensmark, H. *Astronomy & Geophysics*, 48 1.18–1.24 (2007)
- [156] Wang, H. *Physics Reports*, 307 263–267 (1998)
- [157] Wang, H. *PhD Thesis, University of California* (1999)
- [158] Cline, D, Curioni, A, Lamarina, A *et al.* *Astroparticle Physics*, 12 373–377 (2000)
- [159] Howard, A, Bewick, A, Davidge, D *et al.* *Proc. of the 3rd International Workshop on the Identification of Dark Matter*, 457 (2001)
- [160] Akimov, D, Araujo, H, Batyaev, V *et al.* *Proceedings of the 4th International Workshop on the Identification of Dark Matter*, 371–376 (2003)
- [161] Electron Tubes Ltd. <http://www.et-enterprises.com/>
- [162] Carson, M, Davies, J, Daw, E *et al.* *Astroparticle Physics*, 21(6) 667–687 (2004)

- [163] Araujo, H, Kudryavtsev, V, Spooner, N *et al.* *Nuclear Inst. and Methods in Physics Research A*, 545 398–411 (2005)
- [164] Bungau, C, Camanzi, B, Champer, J *et al.* *Astroparticle Physics*, 23(1) 97–115 (2005)
- [165] SAES Getters. <http://www.saesgetters.com>
- [166] Araújo, H, Bewick, A, Davidge, D *et al.* *Nuclear Inst. and Methods in Physics Research, A*, 521(2-3) 407–415 (2004)
- [167] Alner, G, Araújo, H, Bewick, A *et al.* *Physics Letters B*, 653(2-4) 161–166 (2007)
- [168] Agilent Technologies. <http://www.acqiris.com>
- [169] HBOOK. http://wwwasdoc.web.cern.ch/wwwasdoc/hbook_html3/hboomain.html
- [170] Sumner, T, Akimov, D, Alner, G *et al.* *Nuclear Physics B (Proceedings Supplements)*, 173 108–112 (2007)
- [171] Lebedenko, VN, Araujo, HM, Barnes, EJ *et al.* *Physical Review Letters*, 103 151302 (2009)
- [172] Apple XGrid. <http://www.apple.com/server/macosx/technology/xgrid.html>
- [173] Lindote, A, Araújo, H, da Cunha, JP *et al.* *Nuclear Inst. and Methods in Physics Research A*, 573 200–203 (2007)
- [174] Sorensen, P, Manzur, A, Dahl, CE *et al.* *arXiv:astro-ph/0807.0459* (2008)
- [175] Feldman, G & Cousins, R. *Physical Review D*, 57(7) 3873–3889 (1998)
- [176] Yellin, S. *Physical Review D*, 66 032005 (2002)
- [177] Yellin, S. *arXiv:0709.2701* (2007)
- [178] Neyman, J. *Philosophical Transactions of the Royal Society of London. Series A*, 236 333–380 (1937)
- [179] Cowan, G. *Statistical data analysis (Oxford Science Publications)*, 197 (1998)
- [180] ROOT Data Analysis Framework. <http://root.cern.ch/drupal/>
- [181] Sanglard, V. *ILIAS N3/Eurograd Dark Matter Training Session*

-
- [182] Currie, A. *ZEPLIN-III collaboration internal note - z3 RepSta 3 v1 090612* (2009)
- [183] Navarro, JF, Frenk, CS & White, SDM. *The Astrophysical Journal*, 462 563–575 (1996)
- [184] Moore, B, Governato, F, Quinn, T *et al.* *The Astrophysical Journal*, 499 L5–L8 (1998)
- [185] Helm, R. *Physical Review*, 104(5) 1466–1475 (1956)
- [186] Bondar, A, Buzulutskov, A, Grebenuk, A *et al.* *Nuclear Inst. and Methods in Physics Research A*, 556 273–280 (2006)
- [187] Hagmann, C & Bernstein, A. *IEEE Transactions on Nuclear Science*, 51(5) 2151–2155 (2004)
- [188] Akimov, D, Burenkov, A, Churakov, D *et al.* *arxiv:hep-ex/9703011* (1997)
- [189] Santos, F, Dias, T, Stauffer, A *et al.* *Journal of Physics D. Applied Physics*, 27 42–48 (1994)
- [190] Aprile, E, Giboni, K, Majewski, P *et al.* *IEEE Transactions on Nuclear Science*, 51(5) 1986–1990 (2004)
- [191] *CRC Handbook of Chemistry and Physics, 88th Edition 2007-2008*
- [192] Gullikson, E & Henke, B. *Physical Review B*, 39(1) 1–9 (1989)
- [193] Reininger, R, Asaf, U, Steinberger, I *et al.* *Physical Review B*, 28 3193–3199 (1983)
- [194] Schmidt, W, Hilt, O, Illenberger, E *et al.* *Radiation Physics and Chemistry*, 74(3-4) 152–159 (2005)
- [195] Lee, L & Smith, G. *The Journal of Chemical Physics*, 70(4) 1727–1735 (1979)
- [196] Beaglehole, D. *Physical Review Letters*, 15(13) 551–553 (1965)
- [197] Reshotko, M, Asaf, U, Ascarelli, G *et al.* *Physical Review B*, 43(17) 14174–14181 (1991)
- [198] Hitachi, A. *The Journal of Chemical Physics*, 80(2) 745–748 (1984)
- [199] Aprile, E, Dahl, C, de Viveiros, L *et al.* *Physical Review Letters*, 97 081302 (2006)

-
- [200] Surko, C, Packard, R, Dick, G *et al.* *Physical Review Letters*, 24(12) 657–659 (1970)
- [201] Baldini, A, Bemporad, C, Cei, F *et al.* *Nuclear Instruments and Methods in Physics Research Section A*, 545 753–764 (2005)
- [202] Marsh, J, Thomas, D & Burke, M. *Nuclear Instruments and Methods in Physics Research A*, 366 340–348 (1995)