THE UNIVERSITY *of* EDINBURGH

Results from the ZEPLIN–III Dark Matter Search Experiment

Paul Scovell

School of Physics and Astronomy

The University of Edinburgh

A thesis submitted for the degree of

Doctor of Philosophy

March 2011

For
My Family

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–III detector utilises liquid xenon (Xe) as a target medium in the search for 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–III experiment, operating in two-phase (liquid/gas) mode, measures both the scintillation and ionisation signatures produced during an interaction.

The first science run (FSR) of ZEPLIN–III was performed during three months in 2008. The run culminated in a published result which excluded a WIMP-nucleon interaction cross-section above 8.1×10^{-8} pb for a 60 GeVc^{-2} WIMP at the 90% confidence level.

ZEPLIN–III then entered an upgrade period where the photomultiplier tube (PMT) array, previously the dominant source of background, was replaced with new, ultra-low background, PMTs. The radio-contamination of components used to make these PMTs has been thoroughly studied and their impact on the background rates in ZEPLIN–III characterised. Additionally, a new 1.5 tonne plastic scintillator veto detector was constructed, increasing the ability to reject WIMP-like signals caused by neutron induced nuclear recoil events and improving the γ-ray discrimination capability of ZEPLIN–III.

The second science run (SSR) of ZEPLIN–III began in June 2010 and continued for 6 months, with a projected upper limit for the interaction cross-section of 1.52×10^{-8} pb for a 55 GeVc^{-2} WIMP at the 90% confidence level.

Acknowledgements

I have found the past three-and-a-half years to be a deeply exciting and rewarding time and I have been proud to be part of a fantastic team working on an experiment that is pushing back the boundaries of human knowledge. The breadth of work that I have been able to perform simply would not be available in other experiments and I am grateful to have been given the opportunity.

The people I should thank first and foremost are my two supervisors, Alex Murphy and Chamkaur Ghag. It is through their encouragement, support and friendship that I have been able to work to the best of my ability. I would also like to thank the current and past members of the Edinburgh Dark Matter group: Emma, Ant and Lea for making the work more fun and the non-physics related antics. In addition to this, I would also like to thank the members of the ZEPLIN-III collaboration who have supported me throughout the course of my studies and have always been willing to help me (or give me more work). I don't believe I will ever be in a collaboration that will have a meeting quite like the one we had in Moscow!

I would also like to extend my thanks to the Nuclear Physics group at the University of Edinburgh and, in particular, Gav. I feel we have a unique group ethos here which seems to centre around nights in the magnificent Garry-B's and five-a-side football on a Thursday (which is always entertaining!). I have enjoyed almost every minute of the last three and a half years and it's thanks to you all. Big thanks also has to go to my flatmate Henry just for being Henry.

Finally, I'd like to thank my girlfriend Kirsten for putting up with me for the past few years. I know that my frequent trips away are a pain but your support has meant a lot to me.

Declaration

This thesis has been composed by myself and no portion has been submitted for any other degree or professional qualification. The work presented in this thesis was carried out within the Nuclear Physics Group at the University of Edinburgh, as part of the ZEPLIN–III collaboration. I personally made a substantial contribution to the work described from Chapter 3 onwards, although the ZEPLIN–III project is a collaborative effort and much work must be attributed to several workers at a time. My work includes the analysis of data from the first science run of ZEPLIN–III, the radio-assaying and simulation of components used in the ZEPLIN–III detector, the manufacture of and the characterisation and analysis of data from the ZEPLIN–III veto detector. Finally, I also worked on the data analysis of the ZEPLIN–III second science run including the verification of simulated backgrounds and the development of data quality cuts and the definition of their efficiencies.

Paul Scovell

Contents

List of Figures

List of Tables

Chapter 1

Introduction

For almost 75 years, cosmology has been evolving to the point where we have what would appear to be an elegant and successful description of our Universe. Many astrophysical observables can now be described by one model containing just 6 parameters. ΛCDM (Cold Dark Matter, where Λ denotes the inclusion of a cosmological constant) explains the accelerating expansion of the Universe, the cosmic microwave background and the observed distribution of matter on the largest scales. ΛCDM is also capable of describing the velocities of galaxies within clusters and the velocity of matter in spiral galaxies.

The major problem facing the ΛCDM model is that the combination of all astrophysical observables produces a result that lends only 4-5% of the mass of the Universe to baryonic matter. This means that a large proportion of the mass of the Universe must be of non-baryonic origin. The ΛCDM model points towards 23% of this being Cold Dark Matter.

Contemporaneously, in particle physics, a proposed solution to the hierarchy problem was developed. Supersymmetry postulates a new family of supersymmetric partners to each of the standard model particles. These particles differ in spin to their standard model partners by $\frac{1}{2}\hbar$. Supersymmetry is described as a broken symmetry as no evidence of supersymmetric particles has been seen in the mass-energy range of the known standard model particles ($<$172.7 GeVc^{-2}, the mass of the top quark) meaning the supersymmetric partners of standard model particles must have masses exceeding those of the standard model. Most Minimal Supersymmetric Models (MSSM) predict that the Lightest Supersymmetric

Particle (LSP) would have a mass in the 100-1000 GeVc^{-2} range and also predict that this particle could be stable. If this is true, then it is reasonable to assume that this particle constitutes the vast majority (if not all) of the non-baryonic mass. Some particle physics models also predict that this LSP will interact with baryonic matter via the weak nuclear interaction.

These independent predictions, when taken together, complement each-other elegantly. A problem discovered in cosmology "solved" by a potential solution to a problem in particle physics and vice-versa. A direct detection of dark matter particles by observable interaction with baryonic matter would solve two of the biggest problems in modern science at the same time.

The ZEPLIN–III collaboration is a multi-national collaboration involving the University of Edinburgh, Imperial College London, the STFC Rutherford Appleton Laboratory, LIP-Coimbra and ITEP Moscow. The collaboration is one of many searching for the elusive signal that would confirm that dark matter not only exists, but interacts with baryonic matter via a known mechanism. The ZEPLIN–III detector is based at the Boulby Potash mine in Redcar and Cleveland, UK. The detector uses liquid Xe as a detector medium and detects interactions using the principles of scintillation and ionisation. The detector is designed to have a sensitivity to a WIMP-nucleon weak interaction cross-section of order $\sim 10^{-9}$ pb (10^{-45} cm^2). Chapter 4 of this thesis will concentrate on the design and commissioning of the detector culminating in a three month engineering science run. Chapter 5 concentrates on the radio-assaying measurements and background simulation of the components used in the upgrade of the ZEPLIN-III system. Chapter 6 concentrates on the ZEPLIN–III veto detector, in particular, detector design and manufacture leading to the detectors calibration and analysis of science data. Chapter 7 concentrates on the re-commissioning of ZEPLIN–III in preparation for a second science run, the detector calibration, daily monitoring, data analysis and finally the calculation of an upper limit on the WIMP-nucleon weak interaction cross-section assuming a null detection.

The ZEPLIN–III collaboration provides a unique situation, not present in other particle physics experiments. The collaboration is sufficiently large to produce cutting edge scientific results but it is also small enough that each and every member is able (and, indeed, is required) to take leading roles in many aspects of the experimental process. Within ZEPLIN–III, I have taken leading roles in

the calibration and parameterisation of the detector, including the characterisation of radiological impurities of the detector components. I have also taken a leading role in defining and parameterising data cuts to remove spurious events, which would be dominant above the detection of any (rare) potential WIMP interactions.

During the upgrade period of ZEPLIN–III, I took a leading role in the radio-assaying of components for the upgraded photo-multipier tube (PMT) array and veto detector. This included measuring the radio-activity of individual components and assessing the resultant impact on the backgrounds in ZEPLIN-III.

During the second science run (SSR) I maintained a leading role in the calibration of and analysis of data from ZEPLIN–III itself. I also took a leading role in building, testing and installing the ZEPLIN–III veto detector. I also took a leading role in the measurement of the response of the plastic scintillator, synchronisation of veto detector data with ZEPLIN–III data and its analysis.

Chapter 2

The Dark Matter Problem

For centuries, humankind has striven to understand how the Universe works on the largest scale (beginning with early astronomical observations) and the smallest scale. Until less than 100 years ago, it was believed that luminous matter constituted all of the mass in the Universe. Modern estimates now place this figure at around 1% of the mass-energy content. This, plus another 4% of mass not found in stars, (*e.g.* interstellar gas) means that only 5% of the entire Universe is composed of baryonic (standard model) matter. Evidence also shows that the mass-energy content of the Universe is comprised of a 28% mass component and a 72% energy component. This suggests that 23% of the Universe is composed of non-luminous, non-baryonic, matter.

Research into the nature of missing (or "dark") matter has been the subject of experimental and theoretical endeavour for almost three quarters of a century and remains one of the greatest unanswered questions in science [1]. The following chapter will outline how cosmology and particle physics have converged to lead us to the assumed Universal mass-energy components, beginning with an overview of the cosmological and particulate theories underlying the ideas and leading on to outlining the actual observations that support the theoretical postulations.

2.1 Cosmological Origins

The Big Bang theory is one of two key principles underpinning the current ideas behind modern cosmology. The other is the so-called cosmological principle which

states that the Universe is both isotropic and homogeneous. In effect, we do not hold a special place in the Universe and, wherever we look, the laws of physics remain the same.

The idea that the Universe is expanding was first potulated by Edwin Hubble in his studies of distant galaxies. Hubble found that these galaxies were receding at a velocity $\vec{v}$ proportional to their distance away $\vec{r}$:

$$\vec{v} = H_0 \vec{r}. \tag{2.1}$$

The constant of proportionality in this equation, H_0, is known as Hubble's parameter at the present time. The most recent measurements using Cepheid variables have set a value of $H_0 = 72 \pm 8$ (km/s)/Mpc [2].

This law suggests that, at earlier times, the galaxies in our Universe were closer together. If this theory is extrapolated back to the beginning of time, it leads to the conclusion that the Universe has evolved from a single point source (size at most $\sim$ the Planck distance) of infinite temperature and density, known as a singularity. The transition from this singularity to a tenable state of matter and energy is known as the Big Bang.

Another source of evidence comes from the observation that the average number of galaxies per unit volume in the Universe is almost constant across the sky. This holds for the present epoch but has also been directly observed in the cosmic microwave background, which was discovered by Penzias and Wilson in 1965 [3]. The Cosmic Background Explorer (COBE) and the Wilkinson Microwave Anisotropy Probe (WMAP) have found no anisotropies above the 10^{-5} and 10^{-6} scale respectively and it is expected that the Planck satellite will further constrain these values.

The theory of General Relativity, attributed to Albert Einstein [4], is a framework by which the gravitation and curvature of space can be equated. At the centre of this theory are the so-called Einstein field equations. For an isotropic and homogeneous Universe, the expansion (or collapse) of space with mass density ρ and pressure p is given by the Friedmann equations. These equations are derived from the Friedmann-Lemaître-Robertson-Walker (FLRW) [2] solution of Einstein's field equations of general relativity.

The relative expansion rate of the Universe can be calculated by solving the

Einstein field equations, as long as the matter and energy content of the Universe is known. Assuming the cosmological principle is valid, energy and matter can be described as a perfect fluid:

$$T_{\mu\nu} = (\rho + \frac{p}{c^2})U_\mu U_\nu + p g_{\mu\nu}, \tag{2.2}$$

where $T_{\mu\nu}$ is a tensor relating energy and momentum, $g_{\mu\nu}$ is a metric tensor, p is the pressure of the fluid and ρ the density. Pressure and density can be related by the equation of state:

$$p = (\gamma - 1)\rho c^2. \tag{2.3}$$

This equation of state can be neatly solved using just three values for γ. In the past, the Universe was dominated by radiation. In this radiation dominated era, $\gamma_r = \frac{4}{3}$ which gives $p_r = \frac{1}{3}\rho_r c^2$. This is a good approximation for the early Universe which was hot and dense and where matter and radiation were in near thermal equilibrium. The present time is described by γ_m=1 which leads to p_m = 0. This is a good description of the current matter dominated Universe where the pressure-free matter dominates over the energy density of the electromagnetic radiation in the Universe. In the future, it is expected that vacuum energy will dominate. In this case $\gamma_v = 0$ giving p_v = -$\rho_v c^2$. It is thought that there was a brief period in the history of the Universe where a similar equation of state existed. During this period, known as inflation and occurring 10^{-36} s after the big bang, vacuum energy pressure caused a rapid expansion.

Solutions of the Einstein field equations lead to the so called "Friedmann-equations" one of which is described below:

$$\underbrace{H^2}_{\text{Expansion}} + \underbrace{\frac{\kappa c^2}{a^2}}_{\text{Curvature}} = \underbrace{\frac{8\pi G}{3}\rho}_{\text{Density}}, \tag{2.4}$$

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

In a flat universe with no curvature ($\kappa = 0$), equation 2.4 reduces to an equation that describes a so-called critical density:

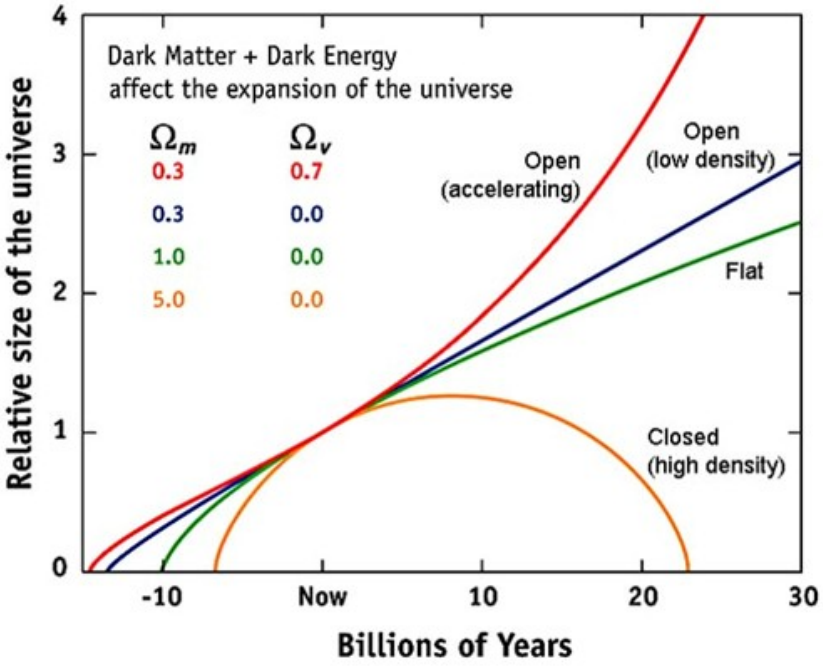

Figure 2.1: A graphical representation of various scenarios for Universal evolution dependent on the mass-energy composition. When comparing to equation 2.6 $\Omega \equiv \Omega_m + \Omega_v$.

$$\rho_c = \frac{3H^2}{8\pi G}. \tag{2.5}$$

If the density is greater than ρ_c, a closed Universe is described which will eventually contract to a "Big Crunch". If the density is less than ρ_c, an open Universe is described which will expand indefinitely. A graphical representation of this can be seen in Figure 2.1. The density parameter Ω is the ratio of the actual Universal density compared to the critical density:

$$\Omega \equiv \frac{\rho}{\rho_c} = \frac{8\pi G}{3H^2}\rho. \tag{2.6}$$

With $\Omega = 1$ representing a flat universe, the individual components due to each constituent of the Universe can be calculated. Ω_{tot} is, therefore composed of a constituent part for radiation (Ω_r), matter (Ω_m) in both baryonic (Ω_b) and non-baryonic (Ω_d) forms, and dark energy (Ω_Λ). The latest calculated values for each

of these components are given in Section 2.3.1. ΛCDM predicts a Universe which is open and accelerating and, even without the inclusion of dark energy, there is still a requirement for a significant dark matter component in the Universe.

2.2 Observational Evidence

2.2.1 Redshift Measurements of Galactic Clusters

In 1933, Fritz Zwicky was the first to propose the idea of missing mass or dark matter [5]. Zwicky looked at eight galaxies in the Coma cluster. He assumed that the cluster was spherically symmetric so that the mutual gravitational forces of the galaxies attract each one towards the centre and that, on average, the outermost galaxies would obey Newtonian mechanics under the influence of a mass equal to that of the cluster. Zwicky determined the average velocity of the galaxies by measuring the Doppler shifting of emitted light and also determined the radius of the mass of the cluster.

Zwicky was then able to use the virial theorem to calculate the total mass needed to produce the galactic motions observed. At the same time, assuming that the amount of light emitted is proportional to the mass of the object, he calculated an independent value for the mass of the cluster. When comparing the two answers, Zwicky realised that the virial mass calculation gave a mass about 400 times greater than the luminous mass. This suggested that a large proportion of the mass was non-luminous. Zwicky did not try to postulate the nature of this dark matter but suggested that gravitational lensing would be a viable technique to measure the mass of galaxies.

2.2.2 Gravitational Lensing

The first gravitationally lensed objects were seen in the 1960s. Even so, it is only relatively recently that Zwicky's predictions that gravitational lensing could be used to determine galactic masses has been proven correct. Gravitational lensing exploits the fact that, in accordance with the laws of general relativity, light will follow the curved space caused by a massive object. This being the case, the object (when viewed from the earth) becomes distorted to an extent directly

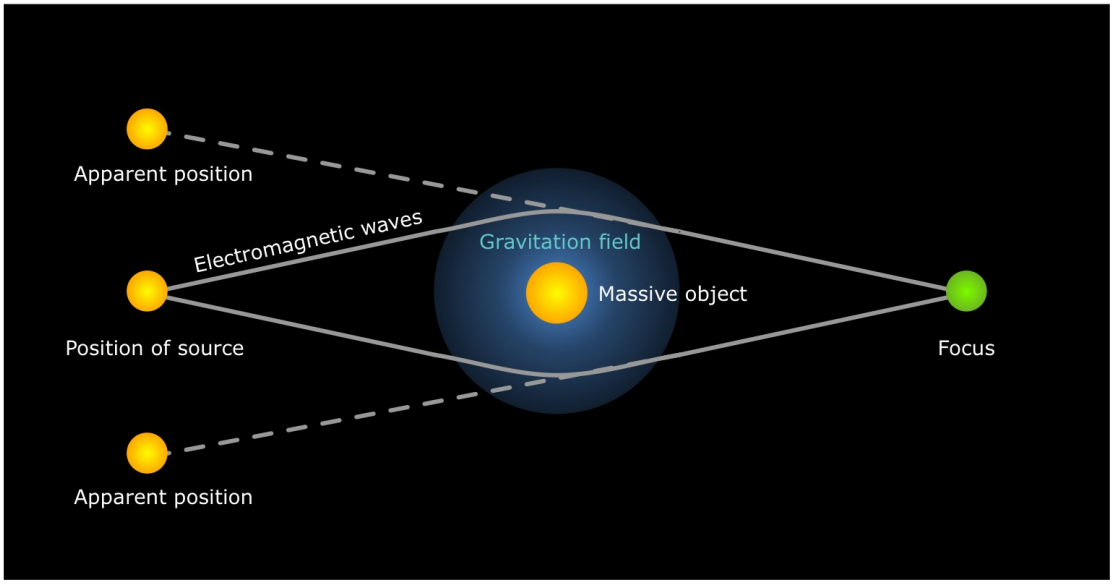

Figure 2.2: Gravitational lensing is used to determine the distribution of mass through the effect of the curvature of space-time due to the mass of the intermediate body. The schematic shows this effect and how a distant object would be observed.

proportional to the mass of the intermediate object (see Figure 2.2).

Strong lensing occurs when the intermediate object is very large and the lensed object is close by. The light from the object is able to take multiple paths, usually causing a series of distorted arcs around the object (see Figure 2.3). Very rarely, if the intermediate object it directly in the line-of-sight between the observer and the distant object, a complete halo may be formed around the object, known as an "Einstein Ring".

The properties of the arced light (curvature, geometry, number of arcs, *etc.*) are used to determine the mass distribution of the lensing object. The cosmological constant can be constrained by determining the volume of space between the observer and the source. Surveys of radio lensing galaxies give a result for the energy density of the Universe (due to the cosmological constant) of $\Omega_\Lambda \sim 0.73$. If the properties of the source object vary with time, the images observed will also vary. Due to the curvature of space-time around the lensing object, there are time delays in the changes in the images that can be used to calculate the Hubble constant, H_0, the value of which is consistent with independent measurements

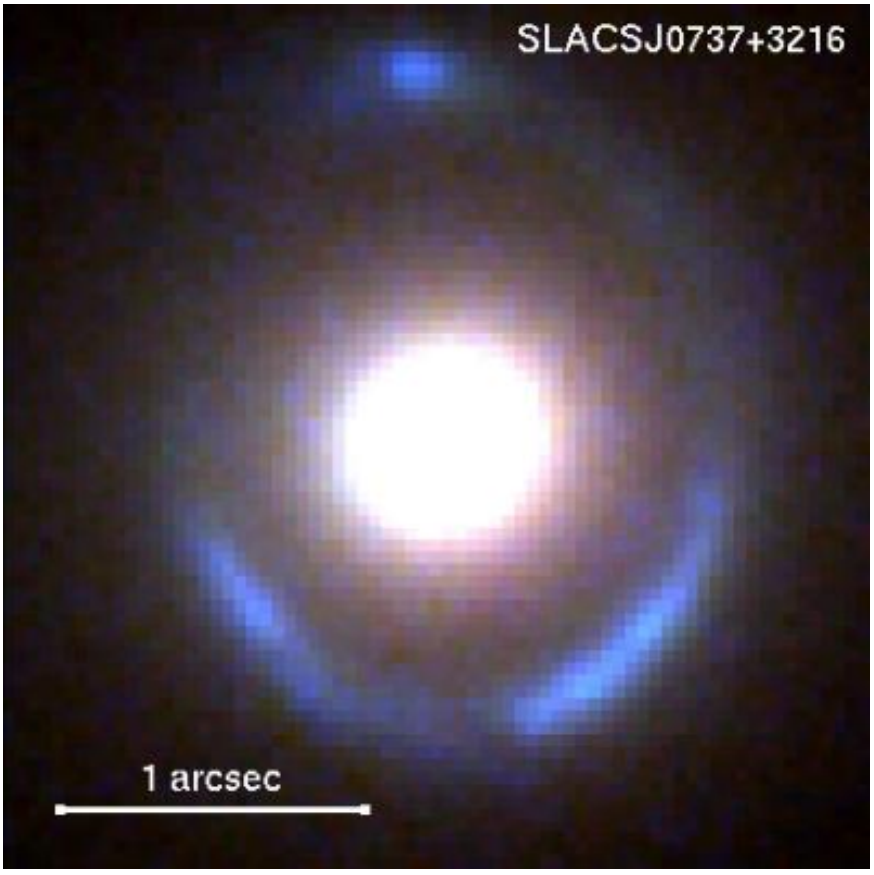

Figure 2.3: Colour composite image of the SLACSJ0737+3216 gravitational lens system, made from Hubble (blue and green) and Keck (red) data. The blue ring is the tiny background galaxy, stretched by the gravitational pull of the foreground lens galaxy at the centre of the image [6].

[7]. The Cosmic Lens All Sky Survey (CLASS) has observed 22 lensing systems in detail and, from determining their parameters, found results consistent with a flat Universe and a cosmological constant with $\Omega_m = 0.31^{+0.27}_{-0.41}$ [8].

2.2.3 Galactic Rotation Curves

Spiral galaxies, including our own galaxy, the Milky Way, and the nearby Andromeda galaxy, M31, consist of a central bulge and a thin disk. The disk is stabilised against radial collapse by angular momentum conservation. By measuring the orbital velocities of the disk using the red-shift of spectral lines, one can determine the orbital velocity as a function of radius. The luminosity of the galaxy falls exponentially with radius so it would be expected that most of the galactic mass would be concentrated within a few scale-lengths of the central bulge. If this were the case, the orbital velocity outside the central bulge would follow Kepler's third law:

$$v_{\rm rot}(r) = \sqrt{\frac{G_N M(r)}{r}}, \tag{2.7}$$

where G_N is the universal gravitational constant and M is the central mass of the galaxy contained within a radius, r. The velocity would then behave in analogy to the Keplarian law $v_{rot} \sim r^{-1/2}$. The galactic curves of around 1000 galaxies have been collected through optical studies [9, 10, 11].

Purely optical studies do not provide the most convincing evidence for unexpected galactic rotation curves, (this arises from the study of the 21 cm radio emission from galactic hydrogen), since optical observations only allow information to be gathered out to about 2 - 4 scale lengths of the central bulge. The emission of hydrogen lines can be observed to much greater radii. It is also possible to observe the ratios of the two CO transitions (wavelengths 2.6 and 1.3 mm) out to large radii to trace the abundance of galactic hydrogen. Figure 2.4 shows that, in almost all cases, the orbital velocity remains constant after a sharp rise near the galactic centre.

The difference between observed and expected rotation curves can be attributed to the gravitational effects of dark matter. A constant orbital velocity can be obtained by placing the galaxy in a spherical 'halo' component so that the total mass $M(r) \propto r$ and the density $\rho(r) \propto r^{-2}$ (see Figure 2.5).

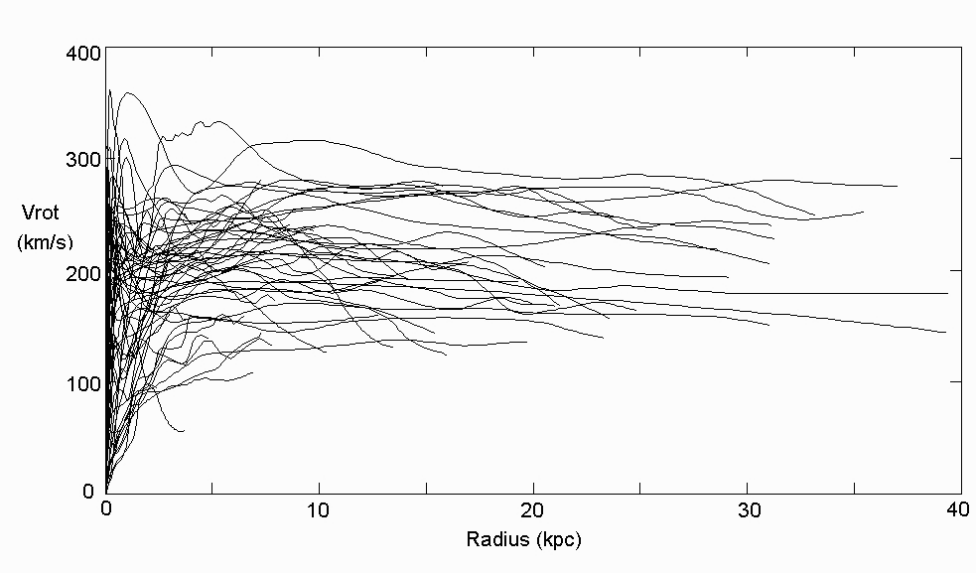

Figure 2.4: Rotation curves for many spiral galaxies obtained by several different methods (optical, H, I, Hα and CO line spectroscopy). The rotation curves generally show a steep rise in velocity close to the centre followed by a constant velocity out to large radii [12].

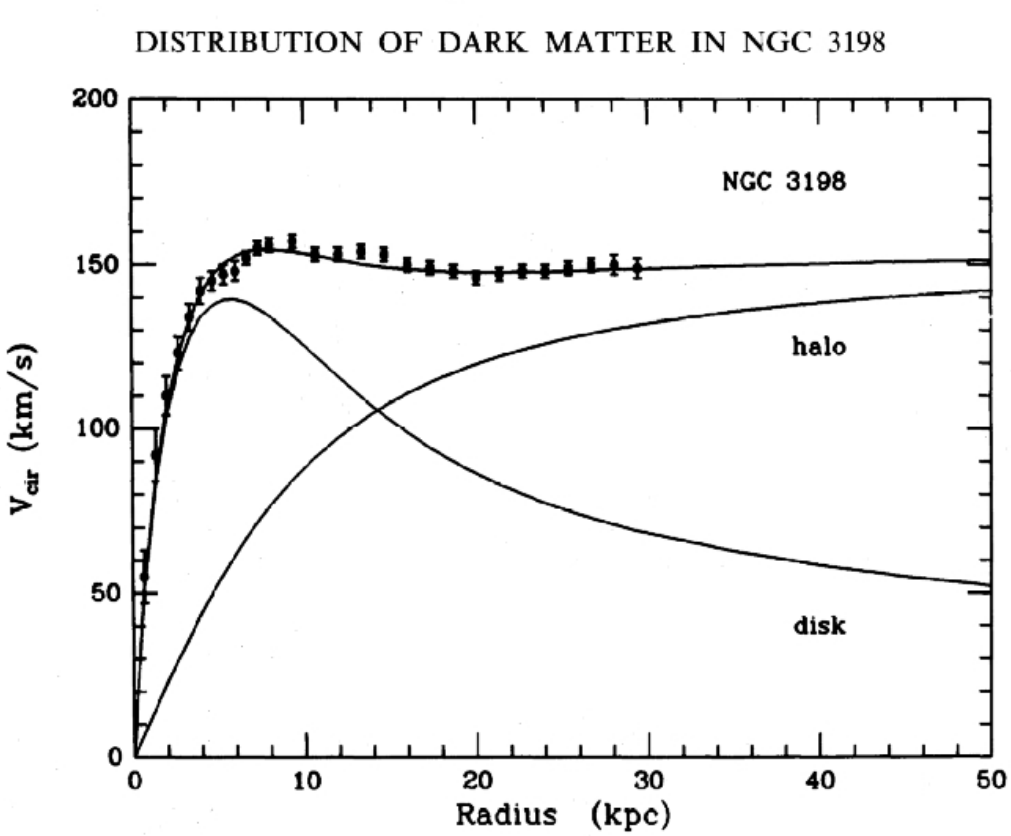

Figure 2.5: Rotation curve for the galaxy NGC 3198. The graph shows the disk and dark matter halo contributions that would lead to the observed galactic rotation curve [12].

2.2.4 Modified Newtonian Dynamics ("MOND")

Modified Newtonian Dynamics (MOND) [13] is an empirical theory proposed to explain observed galactic rotation curves without the need for an unknown dark matter component. It suggests that, at very small accelerations ($a \lesssim cH_0 \approx 7{\times}10^{-10}$ m s^{-2}), there may be deviations from Newtonian gravity. The Newtonian acceleration, α_N can be modified as $a_N \rightarrow a_N\mu(a/a_0)$, where a_0 is a natural constant of rate 1.2×10^{-10} m s^{-2}. The function $\mu \sim 1$ for $a \gg a_0$, and $\mu \sim a/a_0$ for $a \ll a_0$. At large radii, the accelerations experienced by objects become comparable to a_0, and in a simple circular orbit, $v = \sqrt[4]{GMa_0}$. In effect, at large radii, radial dependence vanishes and a flat velocity profile is seen.

When applied to the Coma cluster, however, MOND predicted a mass $\times 4$ greater than the actual mass. Two solutions were proposed where the first was that a_0 is about $4\times$ greater for clusters than for galaxies; the second was that there is some kind of clumping of dark matter on the scale of a cluster but not on the scale of a galaxy. Dark matter clumping would, effectively, rule out cold dark matter but would be compatible with a ~ 2 eVc^{-2} neutrino [14]. The density of these neutrinos would be $\Omega_\nu \approx 0.1$. Having neutrinos as a dark matter candidate would require a large component of additional baryonic dark matter but a 2 eVc^{-2} neutrino mass dark matter component has been ruled out to a $> 95\%$ confidence limit [15] using studies of large-scale structure.

An alternative version of MOND was hypothesised to apply to the evolution of large-scale structure in the Universe [16] which was found to be compatible with observations of the cosmic microwave background and type Ia supernovae. Again, the mass deficit in the Universe was overcome by postulating a 2 eVc^{-2} neutrino with no need for dark matter. However, studies of strong lensing galaxies show an incompatibility with this MOND framework without the inclusion of a significant dark matter component.

A recent addition to the plethora of theorised dark matter alternatives is the idea of Conformal Cyclical Cosmology (CCC) [17]. This theory proposes the existence of a time before the big bang where events happen that directly effect events within our universe. The theory proposes that these events would be observed as concentric rings of anomalously low temperature variance. The paper claims that such concentric formations have been observed with a confidence level

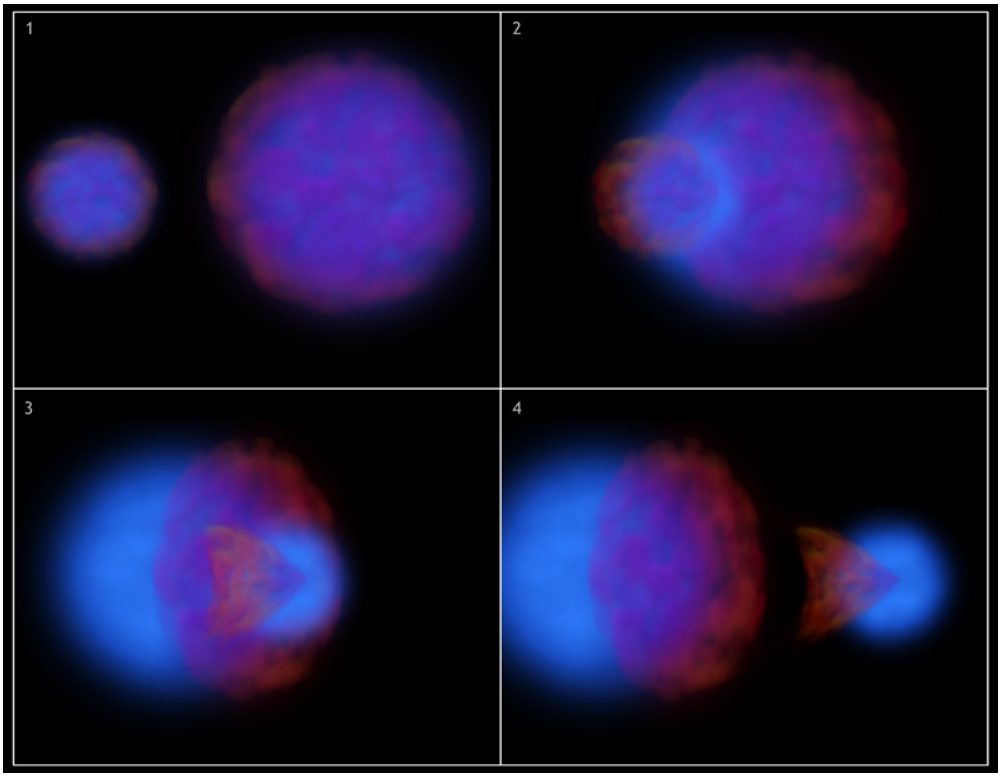

Figure 2.6: An artist's impression of the interaction in the Bullet cluster collision. The interstellar gas collides, causing a shock front in the X-ray spectrum, whereas the dark matter components and, indeed, galaxies pass by each other relatively undisturbed.

of 6σ in both the WMAP 7-year data and the BOOMERanG98 data. These claims have been refuted by several sources [18, 19, 20].

An interesting feature in the Bullet cluster (see Figure 2.6) was discovered in 2006 [21]. The Bullet cluster is a pair of galaxy clusters which have collided. Observed in the visible spectrum, it is possible to see that the galaxies contained in the two clusters pass by each other without being greatly altered, apart from some gravitational slowing. The hot interstellar gas (which makes up the majority of the ordinary, baryonic matter in the cluster) interacts electromagnetically, causing the gases of both clusters to slow much more than the galaxies. In a MOND framework, by definition, dark matter does not exist so it would be reasonable to expect that the lensing mass of the cluster would follow the hot

interstellar gas. However, weak lensing shows that the majority of the mass from each of the two clusters passes through the mass of other and the majority of the mass is located in the same region as the visible galaxies. This provides support that the majority of mass in the cluster is in the form of non-interacting (or, at least, weakly interacting) matter.

Recent studies of the collision dynamics of the bullet cluster have cast into doubt the strength of this argument for dark matter. It has been argued that the shock velocity observed is 23% higher than should be reasonably expected in a cold dark matter scenario [22]. MOND is an idea that will evolve with time and, while it is possible that supersymmetry and weakly interacting dark matter may be discovered in the near future, it is important to remember that alternative theories should be considered.

2.3 ΛCDM

ΛCDM is a model of the Universe built upon the Friedmann equations (see equation 2.4 for one example) which gives a description of the history of the Universe. Directly after the big bang, is a period known as the very early universe. Any ideas concerning this period are purely speculative as accelerator experiments have not been able to probe the energies needed to provide any significant insight. It is proposed (after this very early period) that the Universe then moves through three stages of force unification. The first of these is the Planck epoch (up to 10^{-43} s after the big bang) where, if supersymmetry is correct, the four fundamental forces of electromagnetism, gravitation and the weak and strong nuclear forces would all have the same strength. The Grand Unified Theory (GUT) epoch follows (10^{-43} s $\rightarrow$ 10^{-36} s) in which gravitation begins to separate from the fundamental gauge interactions: electromagnetism and the strong and weak nuclear forces. After this (10^{-36} s $\rightarrow$ 10^{-12} s) comes the electroweak epoch in which electromagnetism and the weak nuclear force are unified. The transition from the GUT to electroweak epochs triggers a period of exponential expansion in the Universe known as the inflationary epoch.

According to the ΛCDM model, dark energy is present as a property of space itself, beginning immediately following the inflationary epoch (10^{-36} s $\rightarrow$ 10^{-32} s), as described by the equation of state. ΛCDM says nothing about the funda-

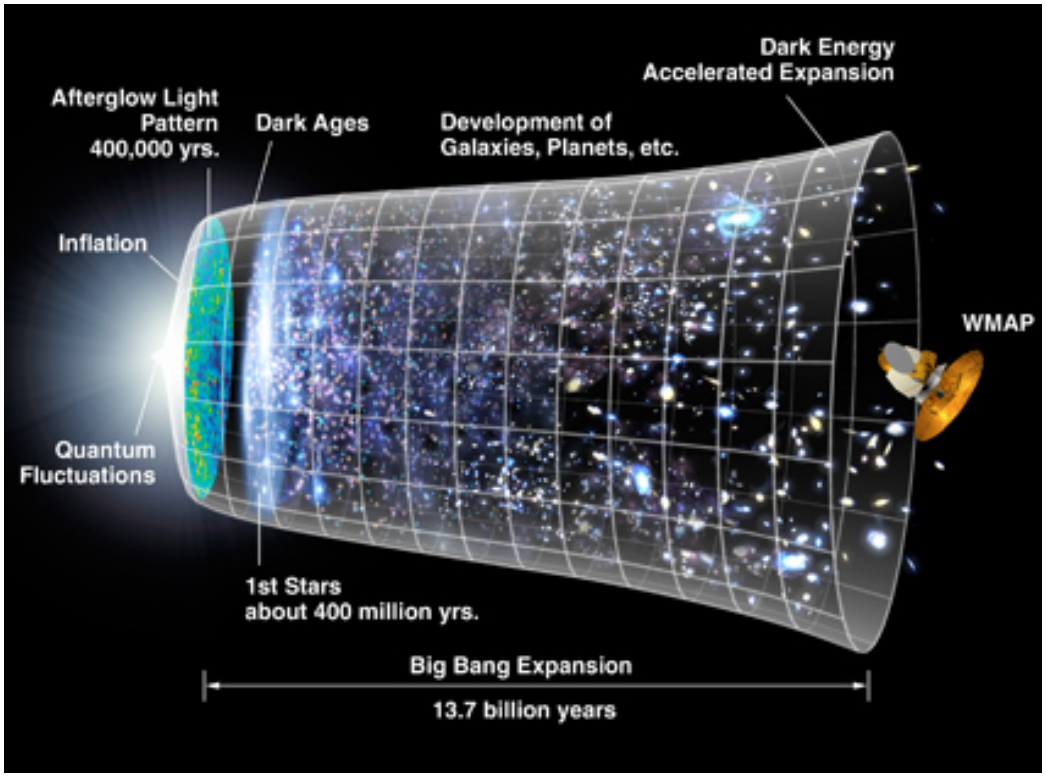

Figure 2.7: A timeline of the evolution of the Universe, starting with the Big Bang on the left and ending with the present day on the right. The transition from an opaque to a transparent universe (the origin of the CMB) is represented by the "Afterglow Light Pattern".

mental physical origin of dark energy but it represents the energy density of a flat universe. Observations indicate that it has existed for at least 9 billion years [23].

The ΛCDM model also explains observations of large-scale structure (galaxies, clusters, superclusters, *etc.*), and the primordial fluctuations in the early universe ($t \lesssim$ 380,000 yr) which were the seeds for the larger structure. As the temperature of the Universe reached $\sim$ 1 eV (3000 K), neutral H atoms were able to form, photons decoupled from baryons and the Universe became transparent to radiation for the first time. Today we see the remnants of this decoupling as the Cosmic Microwave Background (CMB). Figure 2.7 shows a graphic of the evolution of the Universe from the Big Bang to the current epoch.

2.3.1 The CMB

If the Universe evolved from an early, extremely high density stage, as proposed in the Big Bang model, one would expect in this early phase (before galaxy or star formation) a hot, dense, plasma where matter and radiation were in equilibrium. The temperature and, hence, energy of the Universe would have been high enough that electrons were dissociated from atomic nuclei. Photons would undergo continuous Thomson scattering from the free electrons, resulting in an opaque Universe. After the plasma cooled down (due to expansion) to $\sim 10^5$ K, protons and electrons would form neutral hydrogen and the atoms would decouple from the photons. At this point, the Universe would have become transparent. The last scattering photons became a blackbody relic of the Big Bang, permeating the Universe with radiation. As the Universe expanded and cooled, the photons red-shifted to become the Cosmic Microwave Background radiation we observe today.

It was in 1948 that George Gamow, Ralph Alpher and Robert Herman [24, 25] hypothesised that, as a remnant of the primordial plasma, there should be a thermal relic with a black-body spectrum at a temperature of the order of $T \sim 3 - 10$ K. It was only in 1964 that A.A. Penzias and R.W. Wilson discovered an excess radio background at $\lambda = 7.35$ cm corresponding to a temperature of 2.5 - 4.5 K [3]. The most complete (if not the most recent) study of the CMB comes from the Wilkinson Microwave Anisotropy Probe (WMAP) [26, 27] and an all-sky map of temperature anisotropies is shown in Figure 2.8. In this figure, the temperature (and hence, density) fluctuations are obvious.

The temperature fluctuations are only ~ 100 μK. The local galactic plane runs along the major axis of the plot and the Coma cluster is near to the top of the minor axis. The average temperature is in excellent agreement with a single temperature black-body spectrum of 2.725K. The results are the most precise blackbody spectrum ever measured with a temperature of 2.725 ± 0.002 K and illustrate the predictions of the Big Bang theory to an extraordinary degree.

The detection of the CMB and its blackbody spectrum are the best evidence for the Big Bang theory after the discovery of the expansion of the Universe by Hubble. Measurements have shown the CMB to be isotropic to <100 μK, supporting the idea of an homogeneous early universe. If the current epoch is

considered, however, it is obvious that there are massive temperature and density differences between galaxies, galaxy clusters and the space in between. In the ΛCDM model, quantum fluctuations in the early universe can give rise to density fluctuations (following inflation) manifested as over-densities that do not collapse under the force of self gravity until they reach their own particle horizon and every point within the region of over-density is in causal contact with every other point. After this, the perturbation will collapse until it reaches the Jean's length. The radiation pressure will counteract gravity and acoustic oscillations will occur. These are visible in the CMB power spectrum (see Figure 2.9) as a series of peaks with a higher multipole moment than the mean value.

The position of the first peak at $l \approx 100$ gives the scale of the baryon acoustic oscillations ($r_s = 147 \pm 2$ Mpc [28]). The ratio of the heights of the first and second acoustic peaks in Figure 2.9 is a sensitive probe of the baryon density and gives $\Omega_m = 0.0457 \pm 0.0018$ for $h = 0.7$ [28]. Through the Saches-Wolfe effect (gravitational redshift of CMB photons) WMAP is also sensitive to the total matter density of the Universe, Ω_m. The result for one year of data ($\Omega_m = 0.27 \pm 0.06$ [28]) was, more recently, confirmed and improved with the analysis of a five year dataset. This, in combination with large-scale structure and supernova data give a result of $\Omega_m = 0.279 \pm 0.008$ [29]. Ω_m combined with the current best estimate of the age of the Universe, combine to give a value for the Hubble constant of $H_0 = 71^{+4}_{-3}$ km s^{-1} Mpc^{-1}. This result is consistent with the value determined using Type Ia supernovae observations (discussed in Section 2.3.4).

In summary, WMAP (and, shortly, Planck) gives the best estimate for Ω_b, Ω_m and H_0 (see Figure 2.8). These values can be constrained (and indeed others determined) by combining CMB data with observations of Type Ia supernovae and large-scale structure. The current best values for the composition of the Universe according to the ΛCDM model are summarised in Table 2.1.

2.3.2 Large-scale Structure

Rigorous constraints can be placed on the cosmological model as well as the dark matter distribution by observation and analysis of large-scale structure in the Universe. The structure of galaxies along with clusters of galaxies and, indeed, the space between them, can be described using linear perturbation theory [31].

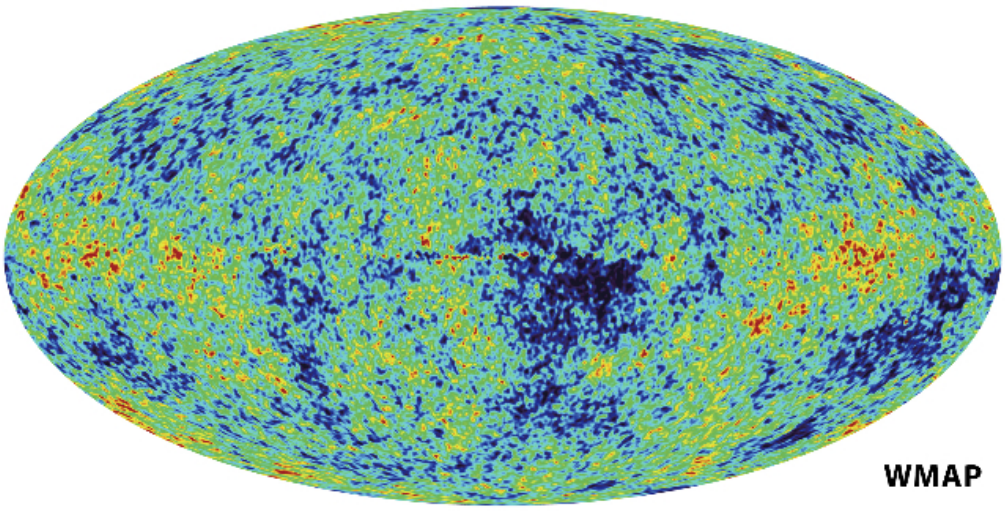

Figure 2.8: WMAP 5-year measurement of temperature anisotropies in the CMB [29]. The temperature difference between blue and red is $\sim \Delta T = 100\mu K$ and corresponds to matter density fluctuations. These perturbations are thought to have been caused by quantum fluctuations in the very early universe which were amplified by inflation and became gravitational potential differences which led to large-scale structure formation.

Parameter	Value
Ω_{tot}	1.005 ± 0.006
Ω_b	0.046 ± 0.002
Ω_m	0.279 ± 0.008
Ω_d	0.233 ± 0.007
Ω_Λ	0.721 ± 0.015
Ω_ν	< 0.013
h	0.71 ± 0.013

Table 2.1: The cosmological parameters estimated by the most recent astrophysical observations.

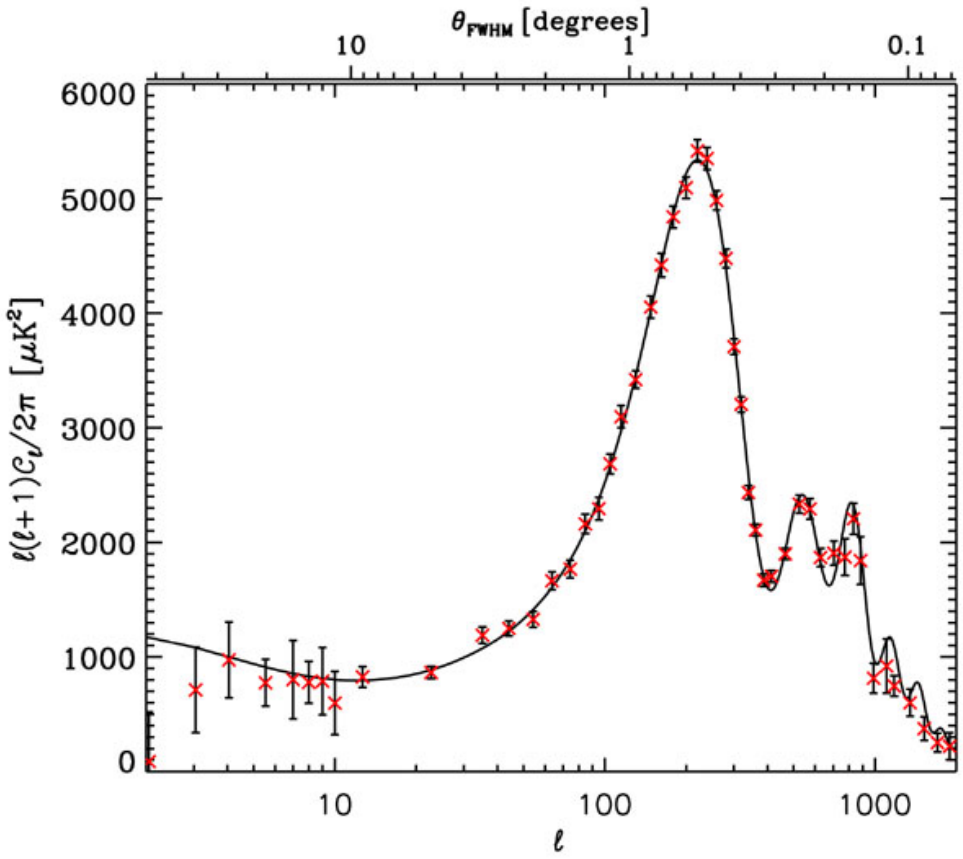

Figure 2.9: CMB angular power spectrum from WMAP 5-year data [30]. The black line is the best fit ΛCDM model to the WMAP data.

In effect, the small anisotropies in the early Universe can be described using large-scale structures seen in the present epoch. If this idea is correct, evidence of baryon acoustic oscillations should be detectable in large-scale structure similar to those seen in the power spectrum of the CMB (Figure 2.9). The Sloan Digital Sky Survey (SDSS) has mapped the positions and absolute luminosities of around 100 million celestial objects [32]. The survey measured the geometry of the Universe through the distance-redshift relationship and through this, it was possible to extract the oscillatory features in the power spectrum.

Approximately 47,000 luminous red galaxies were used for this analysis [33], with a typical redshift of $z = 0.37$. Redshift is a unit of astrophysical distance related to the recession velocity of distant objects. Redshift is determined by observations electromagnetic radiation emitted from distant objects. Due to the Doppler effect, the wavelengths of expected spectral features are shifted in proportion to the recession velocity of the object relative to the observer. Redshift (z) is calculated as follows:

$$z = \frac{\lambda - \lambda_{\text{obs}}}{\lambda}, \tag{2.8}$$

where λ is the expected wavelength of the spectral feature studied and λ_{obs} is the observed wavelength.

Some evidence of baryon acoustic oscillation was seen (at a level of 3.4σ) at a scale of $\sim$ 100 Mpc which is in agreement with the WMAP value quoted in Section 2.3.1. The amplitude of the detected oscillation is small enough that, combined with using a ΛCDM model, it is not compatible with a large baryonic fraction of the total mass of the Universe Ω_m. This means that a large dark matter component is required in the early universe ($z \approx 1100$) to account for the observed large-scale structure in the current epoch.

2.3.3 Big Bang Nucleosynthesis

As described in Section 2.2.1, it is clear that luminous matter only accounts for a small fraction of the total amount of matter in the Universe. Big Bang Nucleosynthesis (BBN) studies can put constraints on the abundances of elements in the primordial universe and, in turn, put constraints on the overall baryonic component of the matter mass. It is also possible to determine if the non-luminous

component is hidden in the form of massive compact halo objects (MACHOs, see Section 2.4.1).

At $t \leq 0.1$ s after the Big Bang, the neutron:proton ratio was maintained in thermal equilibrium. This was due to the Universe being sufficiently hot and dense. Between 0.1 s and 1 s, the ratio followed the form:

$$\frac{n}{p} = e^{-\frac{m_n - m_p}{T}}, \tag{2.9}$$

where n and p are the number densities of neutrons an protons, respectively. The masses of the neutron and proton are given by m_n and m_p respectively and T is the average temperature of the Universe at that epoch.

At ~ 1 s, the temperature fell below the mass difference of the neutron and the proton and the equilibrium reaction rate became lower than the rate of the expansion of the Universe. This meant that the neutron:proton ratio froze out at about 1:6, a ratio which dropped slowly due to neutron decay. The half-life of the neutron is ≈ 615 s and, without further reactions, the Universe today would be made entirely of hydrogen.

At $t \approx 200$ s, BBN began. The reaction that preserves the neutron is deuteron formation ($\mathrm{p} + \mathrm{n} \rightarrow \mathrm{d} + \gamma$). At 200s, the temperature of the Universe was $T \sim 80$ keV. At temperatures higher than this, the fusion reaction is in thermal equilibrium ($\mathrm{p} + \mathrm{n} \rightleftharpoons \mathrm{d} + \gamma$) meaning deuterons photo-dissociate to a proton and a neutron as quickly as they are produced. Once the temperature had dropped to < 80 keV, the production of ^{3}He, ^{4}He and ^{7}Li could follow.

It is possible to estimate the abundances of elements in the Universe today by using calculations of nucleosynthesis at these early times [34]. The abundances are constrained by comparing predicted values with measured values (see Figure 2.10). Studies of deuterium absorption lines in quasar spectra from hydrogen clouds between the Earth and the quasar give a result of $\Omega_b = 0.038 \pm 0.002$. The primordial abundances of ^{4}He and ^{7}Li relative to H imply $\Omega_b = 0.023 \pm 0.007$. These two values, combined with the WMAP value given in Section 2.3.1 provide compelling evidence that, not only is there a lack of baryonic matter needed to account for the total mass content of the Universe, but there is also not enough luminous matter to account for the total baryonic mass content of the Universe.

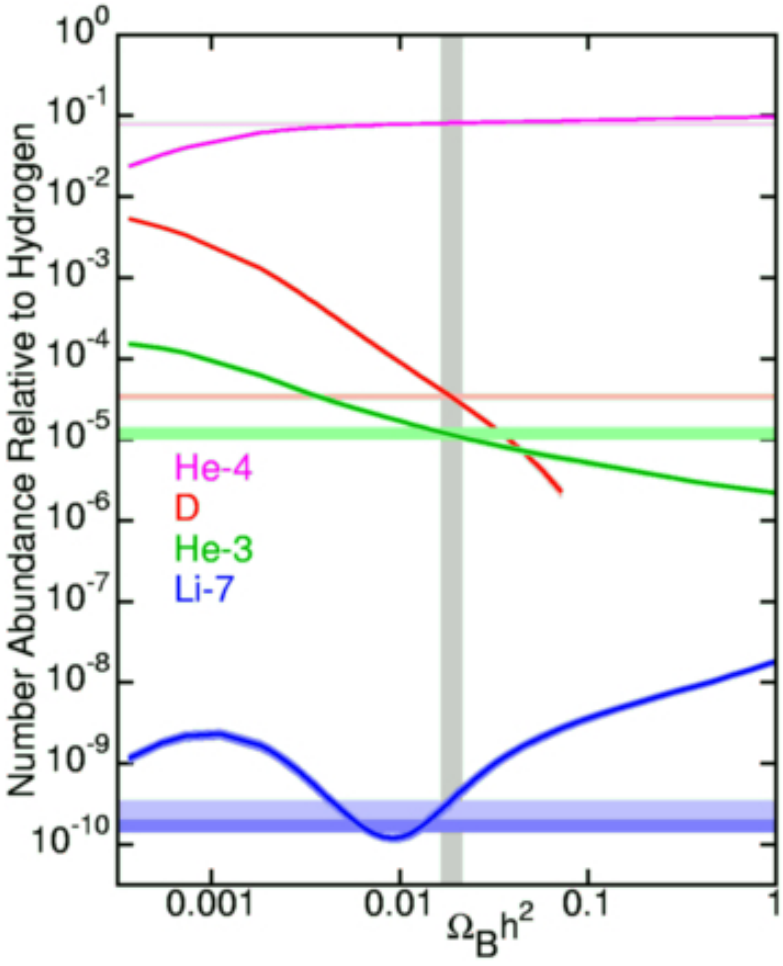

Figure 2.10: The curves show several elemental abundances relative to hydrogen as a function of the baryon fraction of the Universe [35]

2.3.4 Type Ia Supernovae

Type Ia supernovae occur in binary systems with a white dwarf and a companion red giant or main sequence star. They occur when the white dwarf has a mass $<$ 1.4 $M_\odot$ and accretes material from the companion. When the white dwarf mass exceeds the Chandrasekhar limit of 1.4 $M_\odot$, its electron degeneracy pressure is no longer sufficient to prevent gravitational collapse and the white dwarf undergoes a thermo-nuclear explosion. The fact that type Ia supernovae all occur with the same mass means that they can be considered as "standard candles", meaning that each supernova has broadly the same intensity and luminosity. By studying type Ia supernovae to high redshift values, their apparent magnitudes can be used as a measure of distance and expansion rate over cosmological distances [36]. A deviation in the Hubble law will be dependent on a deceleration parameter q_0, which is related to Ω_m and Ω_Λ by:

$$q_0 = -\frac{a(t_0)\ddot{a}(t_0)}{\dot{a}^2(t_0)} = \frac{\Omega_m}{2} - \Omega_\Lambda(t_0). \tag{2.10}$$

For these large redshifts, the luminosity depends on the matter and energy density composition of the Universe. It is therefore possible to measure the apparent luminosities of supernovae as a function of redshift and plot contours for varying matter/energy compositions. The High-Z Supernova Search (HZT) [37] and the Supernova Cosmology Project (SCP) [38] survey type Ia supernovae between redshift $z = 0.01 \rightarrow 1.7$. Results from these projects, combined with values of H_0 from CMB studies, imply that the expansion of the Universe is increasing with time due to a non-zero Λ energy contribution. In effect, distant supernovae are further away than expected [39, 40]. The best fit to the SCP data shows a mass contribution of $\Omega_m \sim 0.25$ and an energy contribution of $\Omega_\Lambda \sim 0.75$ (see Figure 2.11). The HZT finds (assuming $\Omega_{\text{tot}} = 1$) that the mass contribution $\Omega_m = 0.29 \pm 0.05$ and energy contribution $\Omega_\Lambda = 0.71 \pm 0.05$.

Combining the WMAP results with these results gives an overall estimate for the dark energy content of the Universe of $\Omega_\Lambda = 0.72 \pm 0.04$.

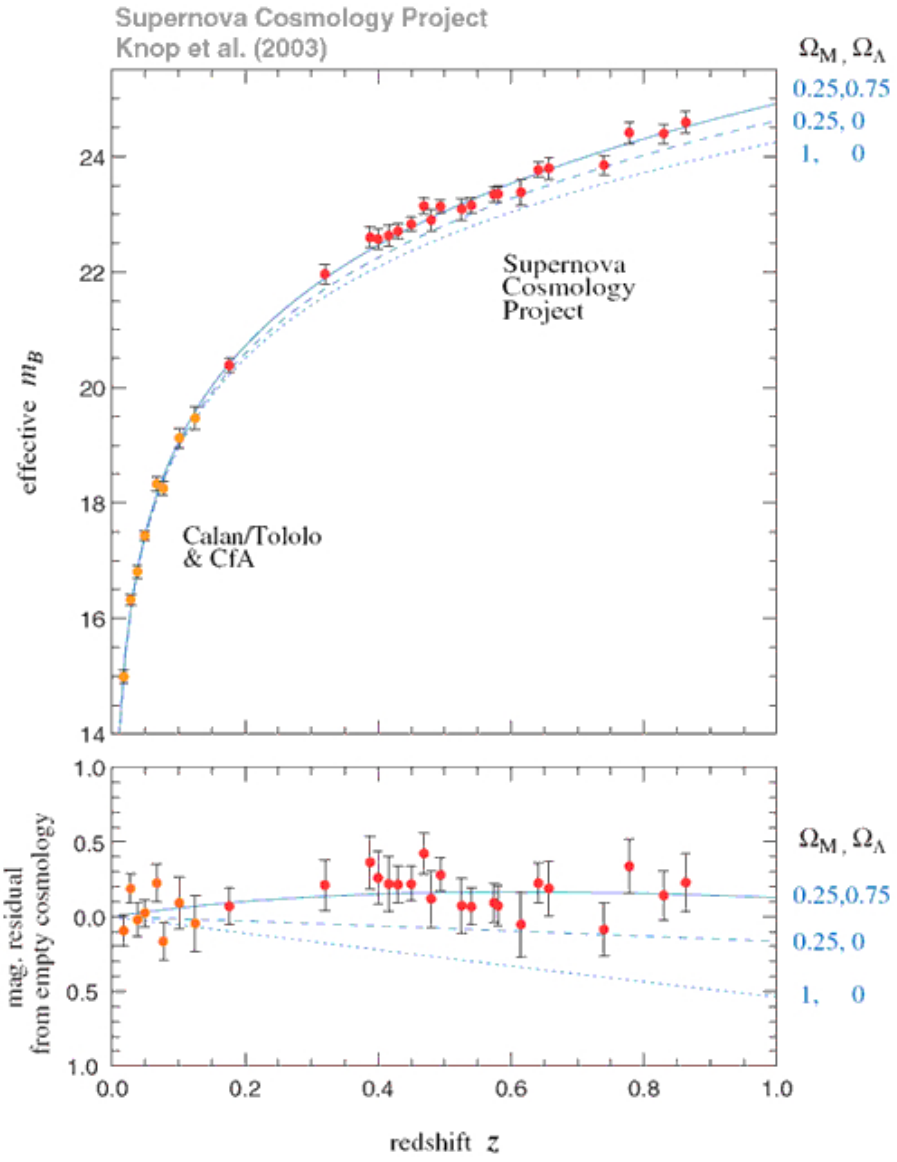

Figure 2.11: The SCP results of measurements of the effective luminosity (m_B) of type Ia supernovae as a function of redshift (z). The best fits to the results suggest a matter:energy ratio of $\sim$ 1:3.

2.3.5 Concordance Cosmology

It is possible to combine the ideas proposed in sections 2.3.1 - 2.3.4 to derive an overall idea of the constraints that can be placed on the mass/energy ratio in the Universe. Figure 2.12 shows the 1, 2 and 3σ allowed regions in $\Omega_m - \Omega_\Lambda$ parameter space with constraints placed by the ideas discussed in Sections 2.3.1 - 2.3.4. Looking deeper at how the matter component is composed, there is a wealth of evidence that a large proportion of the mass of the Universe is non-luminous and also non-baryonic. The next section of this thesis will discuss the possible origin of this matter.

2.4 Dark Matter Candidates

2.4.1 MACHOs

Massive Compact Halo Objects (MACHOs) are any astrophysical objects that could be used to explain the apparent presence of dark matter in galactic halos. Examples of MACHOs are white or red dwarfs, neutron stars and stellar black holes or even unassociated planets (outside a planetary system). MACHOs can be detected when they pass in front of a luminous object. The gravitational potential of the object will bend the light causing the object to appear brighter. This phenomenon is called microlensing and is exactly the same process as gravitational lensing described in Section 2.2.2.

The possibility of the missing mass problem being solved by the presence of large non-luminous but baryonic objects has been ruled out most conclusively in two different ways. Big Bang nucleosynthesis constraints on the abundance of baryonic matter in the early Universe means there cannot be enough baryonic matter in the present epoch to explain the missing mass [42]. The other method by which MACHOs have been ruled out as candidate dark matter is through gravitational lensing as mentioned above. The abundance of MACHOs found by this method is too sparse to account for the missing mass [43].

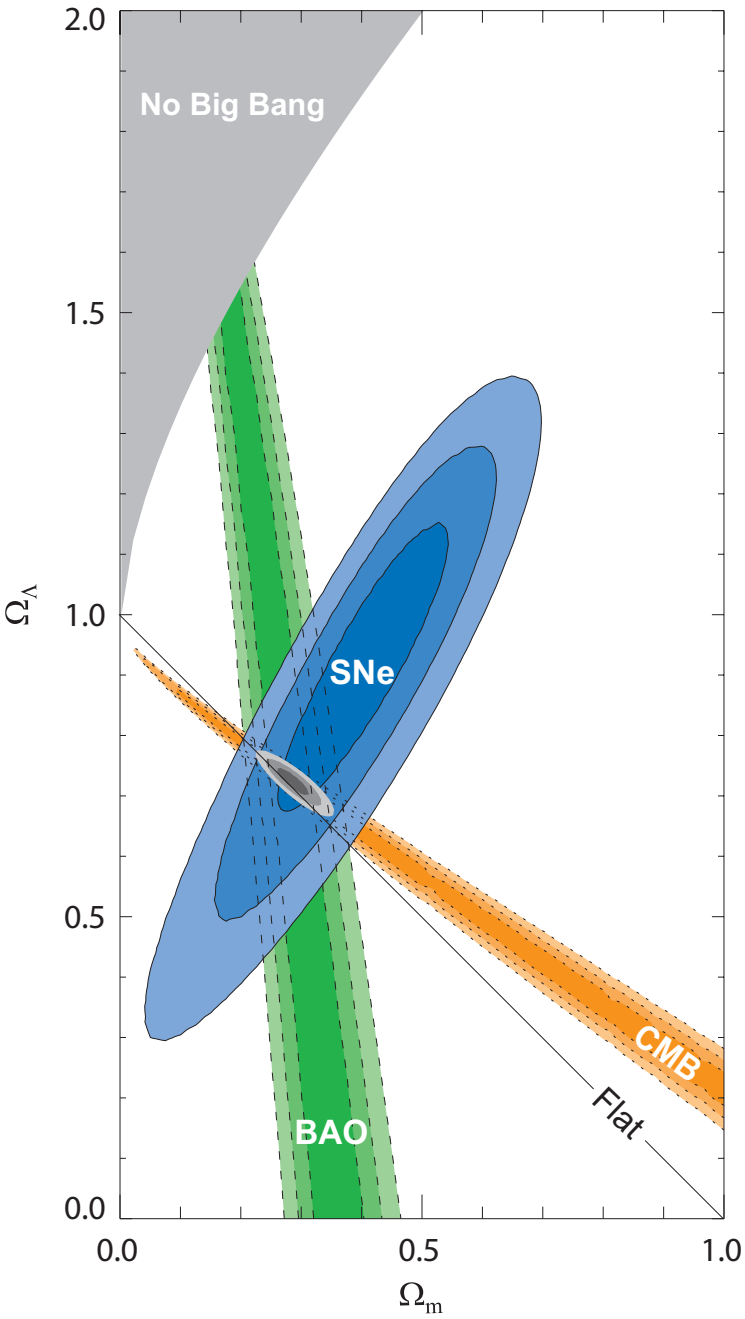

Figure 2.12: The 1, 2 and 3σ allowed regions in $\Omega_m - \Omega_\Lambda$ parameter space. The allowed regions are based on type Ia supernovae constraints (blue), large-scale structure (BAO) constraints (green) and CMB constraints (orange). The combined allowed region is shaded grey [41].

2.4.2 Neutrinos

If dark matter was able to interact via the strong or electromagnetic forces, it is reasonable to expect that dark matter would have bound to neutrons and protons when the temperature of the Universe fell below the energy needed for quark confinement. This would manifest itself in anomalous isotopes of normal matter, which have not been observed [44]. This being the case, the weak nuclear force is the only force, other than gravitation, through which dark matter can interact. Within the the standard model framework, the only particle that will fit this limitation is the neutrino.

It is possible to calculate the contribution of neutrinos to the total mass of the Universe Ω_ν as a function of their combined masses Σm_{ν_i}:

$$\Omega_\nu = \frac{8\pi G}{3H_0^2} n_\nu \Sigma m_{\nu_i}, \tag{2.11}$$

where n_ν is the number density of neutrinos in the current epoch. When neutrinos decoupled from other particles, at time = 1 second after the Big Bang, the Universe still had a temperature $\sim$ 1 MeV. Light neutrinos with mass $<$ 1 MeV are, therefore, relativistic. These relativistic neutrinos are referred to as "hot dark matter". The total density contribution for a Dirac neutrino (half for a Majorana neutrino) can be written [45]:

$$\Omega_\nu = \frac{\Sigma m_{\nu_i}}{93 \text{ eV}} h^{-2}. \tag{2.12}$$

This assumes a neutrino number density of $n_\nu = 113$ cm^{-3} [46, 31].

The experimental limits on the (electron) neutrino mass are dominated by experiments that determine the kinematics of triton decay ${}^3_1H \rightarrow {}^3_2 H + e^- + \bar{\nu}_e$ which give a current upper limit of $m_{\nu_e} < 2.3$ eV [47]. Next generation experiments are expected to improve on this upper limit [48].

Astrophysical observations can be used to place limits on the mass of all neutrinos. The current best limit is set as $\Sigma m_{\nu_i} \leq 0.65$ eV [49]. With masses this small, neutrinos produced at the big bang would be relativistic and would, therefore, damp out the growth of density perturbations needed for large-scale structure evolution. Dark matter must, therefore, be produced non-thermally and be non-relativistic at the point of decoupling [50]. Simulations have been

run that use various starting points for dark matter scenarios and it is a cold dark matter framework that leads to the formation of large-scale structure as seen today (see Figure 2.13).

2.4.3 Axions

Axions were proposed to explain the lack of CP violation in the strong interaction [52, 53]. It was originally proposed that axions would interact strongly with matter, but no evidence of this has been seen. New models for axions have been proposed to circumvent the limits set by dark matter detectors that have not yet been able to detect them [54, 55, 56, 57].

Axions are produced by two distinct mechanisms. At the point when free quarks are bound into hadrons (the QCD phase transition), a Bose-Einstein condensate of axions is produced. These particles would behave as cold dark matter. Axions can also be created through the decay of strings formed at the so-called "Peccei-Quinn" phase transition [58, 59] and this is thought to be the dominant mechanism for axion production as long as inflation occurs after this point.

The properties of axions are set by the mass m_a. The smaller the mass, the more weakly an axion couples to protons and electrons. Many astrophysical arguments come together to suggest that the mass of the axion is $\leq 10^{-2}$ eV [60]. If it is required that the Universe should not be "over-closed", then $\Omega_a h^2 < 1$. In this case, $m_a > 1$ μeV and if strings play an important role then the limit, set by cosmology, is closer to 1 meV. This means that there is only a small window within which axions can exist of 10^{-3} eV $< m_a < 10^{-2}$ eV [60].

Through the addition of extra dimensions, so-called "Kaluza-Klein" axions can have masses of up to 12 keV and this gives an excellent discovery potential [61]. Axions provide a potential and interesting candidate for cold dark matter. With such small masses, however, axions alone cannot form the dominant matter content of the Universe [62]. Axions are, potentially, discoverable through their weak coupling to electromagnetism [63]. It is possible that the axions could resonantly decay to two photons in the presence of a strong magnetic field and efforts have been made to look for this signature decay [64] but no positive discovery has yet been made.

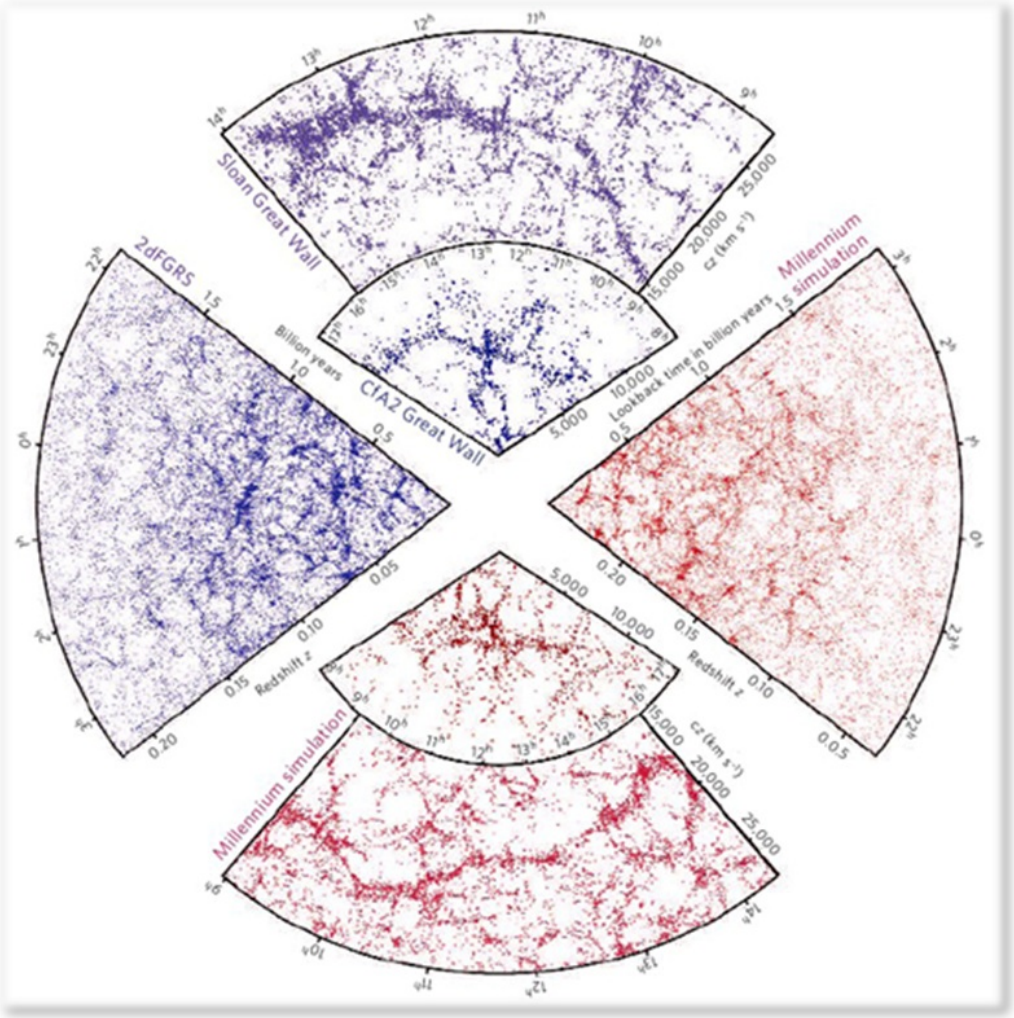

Figure 2.13: Results of the millennium simulation, run at the Durham University Centre for Computational Cosmology [51]. The remarkable similarity between the 2dFGRS (blue) and simulation (red) results can be seen.

Particle	Symbol	Spin	Superparticle	Symbol	Spin
Quark	q	1/2	Squark	$\tilde{q}$	0
Electron	e	1/2	Selectron	$\tilde{e}$	0
Muon	μ	1/2	Smuon	$\tilde{\mu}$	0
Tau	τ	1/2	Stau	$\tilde{\tau}$	0
W	W	1	Wino	$\tilde{W}$	1/2
Z	Z	1	Zino	$\tilde{Z}$	1/2
Photon	γ	1	Photino	$\tilde{\gamma}$	1/2
Gluon	g	1	Gluino	$\tilde{g}$	1/2
Higgs	H	0	Higgsino	$\tilde{H}$	1/2

Table 2.2: The particles of the standard model and their superpartners.

2.4.4 Supersymmetry and WIMPs as DM candidates

2.4.4.1 Supersymmetry

In particle physics, supersymmetry (SUSY) is a symmetry that relates standard model particles of one spin to other particles that differ by half a unit of spin and are known as superpartners. The theory of SUSY, if it is an unbroken symmetry, states that for every type of boson, there exists a corresponding type of fermion with the same mass and quantum numbers. Since no evidence of supersymmetric particles has been seen in particle accelerators and a plethora of standard model particles have, it is expected that SUSY is, in-fact, a broken symmetry. This allows superpartners of standard model particles to have much higher masses.

Standard model bosons (B) are related to their fermion (F) superpartners using the following transformations:

$$Q_\alpha |F\rangle = |B\rangle \text{ and } Q_\alpha |B\rangle = |F\rangle , \tag{2.13}$$

where Q_α are spin-1/2 Hermitian generators. For fermion $\rightarrow$ boson translations, spin goes from $1/2 \rightarrow 0$. For boson $\rightarrow$ fermion translations spin goes from 1 (except the Higgs boson, which has spin 0) $\rightarrow 1/2$ as shown in Table 2.2.

SUSY is a useful tool for solving the so-called "hierarchy problem" in particle physics. In-fact, it was an attempt to solve the hierarchy problem that led to the development of SUSY in the first place. In basic terms, the hierarchy problem asks the question why the weak force is 10^{32} times stronger than gravity. Both

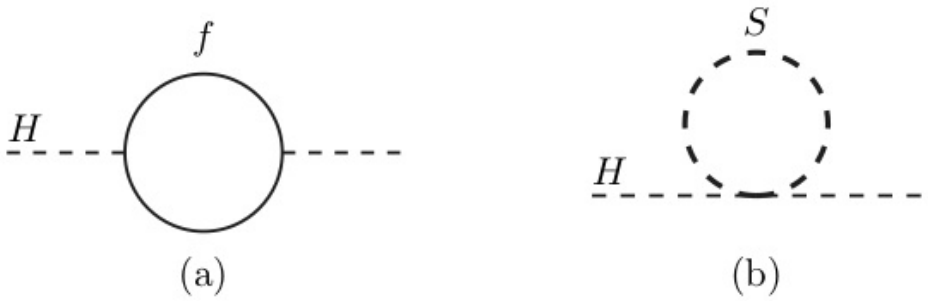

Figure 2.14: One-loop quantum corrections to the Higgs mass parameter Δm_H^2 due to (a) a Dirac fermion f, and (b) a scalar S.

the weak force and gravity are determined by constants of nature, the Fermi coupling constant (G_F) and Newton's constant (G) respectively. The question then arises as to why the Higgs mass is so much lighter than the Planck mass (grand unification energy). Given the large quantum contributions to the Higgs boson mass, one would expect its mass to be close to the Planck mass. This disparity is only possible through some incredibly detailed fine tuning of corrections.

SUSY is the most popular theory for solving the hierarchy problem. It allows the Higgs mass to avoid large quantum corrections by removing the power-law divergences of the radiative corrections to the Higgs mass. Figure 2.14 shows that quantum corrections due to scalar particles cancel with those due to fermions.

Additional motivation comes from the fact that the unification of forces at high energies is expected and vital for a grand unified theory (GUT). A GUT is a theory in which (at high enough energies) the three gauge interactions, electromagnetic, strong and weak, merge into one single interaction with a single coupling constant. In the standard model, the electromagnetic and weak coupling constants converge at $\sim$ 100 GeV, but the strong force does not. Figure 2.15 shows the running of the three coupling constants for the electromagnetic, weak and strong forces with and without a SUSY component. When a SUSY component is assumed, the three forces unify at the 10^{16} GeV scale [45].

The Minimal Supersymmetric Standard Model (MSSM), is an extension to the Standard Model that attempts to add as few new parameters as possible. In

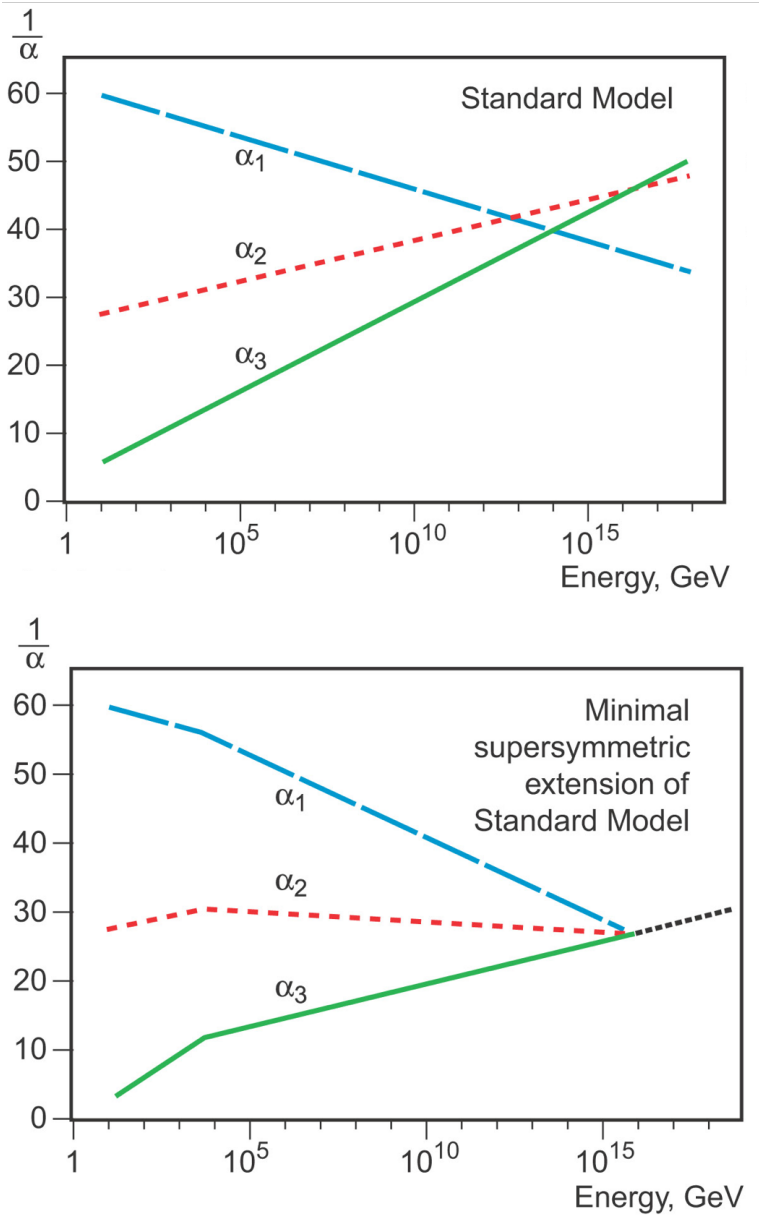

Figure 2.15: Running of the three coupling constants - α_1, α_2 & α_3 - within supersymmetric and non-supersymmetric frameworks. For the supersymmetric model, superpartners have masses around the TeV scale.

the MSSM, the higgsino and gauginos are weak eigenstates, meaning that they are able to mix to form mass eigenstates that are superpartners of the neutral Higgs bosons with gauge bosons. Four combinations are formed and these are known as 'neutralinos'.

One of the problems with MSSM is that the new couplings in the theory can cause lepton and baryon number violating interactions Lepton and baryon number conservation has been tested very precisely by experimentation and the non-conservation would be visible on all energy scales. To solve this problem, a new concept is introduced called R-parity. R-parity is described by the following equation [45]:

$$R = (-1)^{2j+3B+L}, \tag{2.14}$$

where spin is denoted j, baryon number B and lepton number L. Standard model particles have $R = 1$ and supersymmetric particles $R = -1$. With R-parity conserved, the lightest supersymmetric particle (LSP) cannot decay. Annihilation and creation of supersymmetric particles can only occur in pairs and supersymmetric particles will decay to lighter supersymmetric particles until they reach the LSP (see Figure 2.16). The LSP is stable, colourless and has zero charge, making it an ideal candidate for dark matter.

2.4.4.2 WIMPs as Dark Matter

Weakly Interacting Massive Particle (WIMP) is a phrase coined to cover a variety of possible Dark Matter candidates including (but not limited to) LSP candidates such as neutralinos [66] and sneutrinos [62]. For most of the MSSM parameter space, however, the lightest of the four neutralino eigenstates, χ_1^0, is the LSP and this will, henceforth, be referred to as the WIMP. Limits on the mass of the WIMP have been set by accelerators and CMB studies. Accelerators give a lower mass limit of 37 GeVc^{-2} [67] (limited by the maximum centre of mass collision energies obtained and for coupling to W/Z bosons, not photons) and WMAP gives an upper limit of 500 GeVc^{-2} [68] (determined by measurements of Ω_m).

In many models, the WIMP is a Majorana particle. This means that it is a spin $frac12$ particle that is identical to it's anti-particle and is able to self-annihilate ($\chi \equiv \tilde{\chi}$). During the Big Bang, WIMPs were produced along with

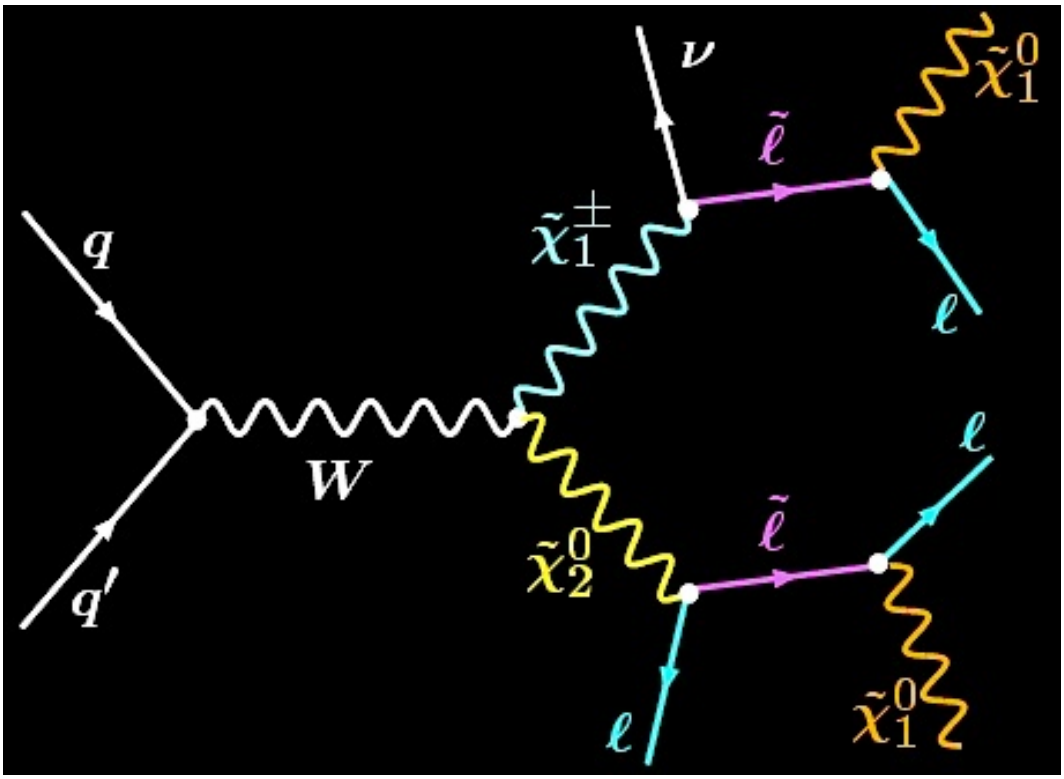

Figure 2.16: potential decay mechanism for the discovery of SUSY at the LHC. the final decay branches see two sleptons each decaying to a lepton-neutralino pair [65].

all other SUSY particles. Whilst the temperature of the Universe was $> m_\chi$, WIMPs were in thermal equilibrium with other particles. The production of WIMPs occurred through, for example, lepton annihilation, $l\tilde{l} \rightarrow \chi\tilde{\chi}$ (where l signifies a lepton), and annihilation of WIMPs occurred through, *e.g.* $\chi\tilde{\chi} \rightarrow l\tilde{l}$.

As the Universe expanded, it reached a scale where the density of WIMPs fell to a point where the probability of annihilation became close to zero. At this point, the cosmological abundance of WIMPs "froze out" to the abundance seen in the present epoch, as shown in Figure 2.17. The present mass-density parameter for the WIMP can be calculated quite easily, given the critical density of the Universe as $\rho_c \sim 10^{-5}h^2$ GeV cm^{-3} [45]:

$$\Omega_\chi h^2 = \frac{m_\chi n_\chi}{\rho_c} \simeq \frac{3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1}}{\langle \sigma_A v \rangle}. \tag{2.15}$$

The WIMP annihilation reation rate can be estimated as $\langle \sigma_A v \rangle \sim 10^{-25}$ cm^3 s^{-1} which leads to a value of Ω_χ of between 0.06 and 0.35. This is in agreement with the missing mass value of $\Omega_d = 0.233 \pm 0.007$ given in Table 2.1.

The fact that the value for Ω_χ derived through particle physics and the value for Ω_d derived through cosmology are in such good agreement is remarkable. This agreement provides the strongest evidence that the missing mass in the Universe can be described most elegantly using a cold dark matter scenario.

2.5 Summary

In summary, the evidence for the presence of a significant dark matter component in the mass of the Universe is almost, though not entirely, conclusive. Complementary studies in both cosmology and particle physics propose a scenario whereby luminous mass only contributes a small fraction of the mass-energy content of the Universe. On a cosmological scale, these studies suggest that the mass of the Universe is dominated by non-luminous, non-baryonic matter which is thought to be cold. For particle physics, the presence of this dark matter may solve the mass-hierarchy problem and help to unify the fundamental forces of nature at the GUT scale.

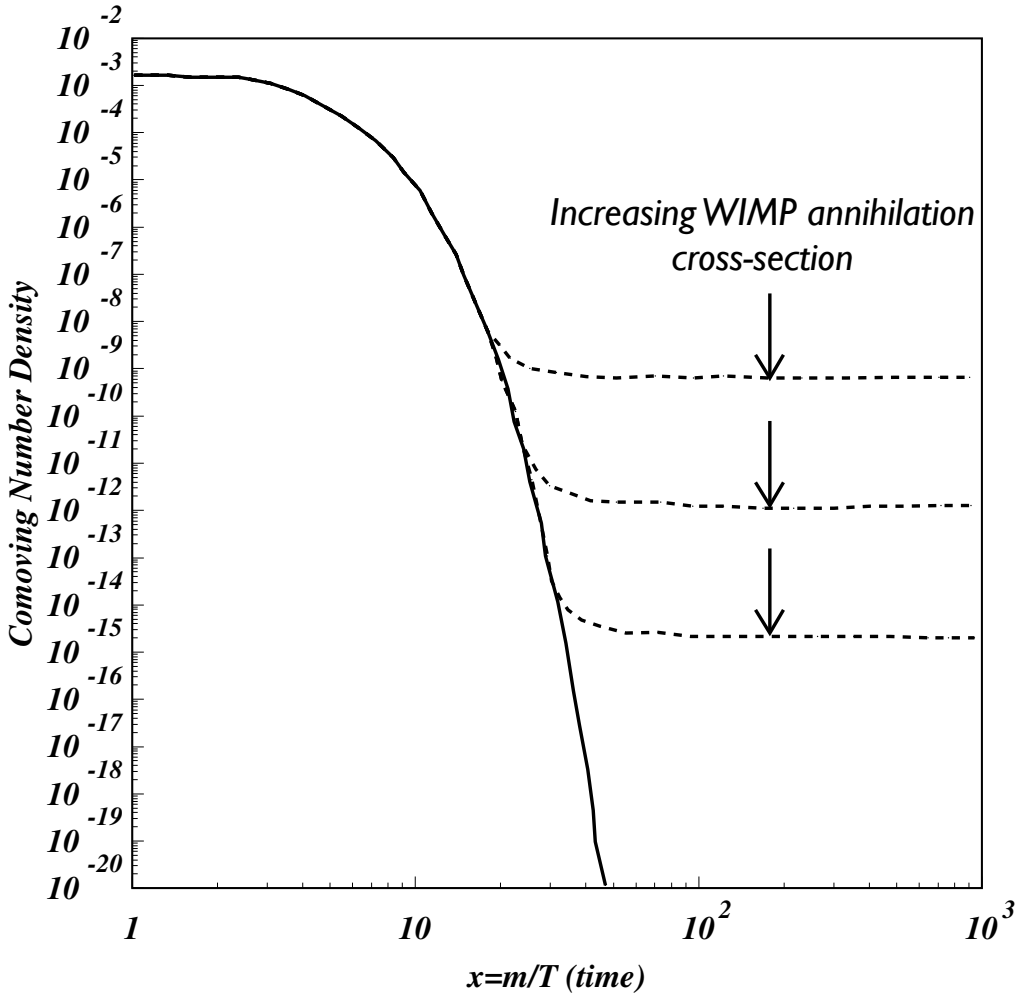

Figure 2.17: The WIMP number density as a function of time. When the size of the Universe reduces the probability of WIMP annihilation to zero, the number density flattens off to a relic density. As the cross-section increases, the number density falls.

Chapter 3

Detection of Dark Matter and the ZEPLIN–III Instrument

The preceding chapter discussed the evidence supporting the postulate that a non-luminous, non-baryonic component comprises ∼23 % of the mass-energy budget of the Universe. The chapter also presented evidence that the dominant form of this matter is a weakly interacting particle with a mass between 10 and 1000 GeVc^{-2}. If this is the case, there are two major pathways to the discovery of this particle: direct and indirect detection. Supersymmetry may be discovered at the LHC, but these data alone could not identify a supersymmetric particle as the dominant galactic dark matter component.

3.1 Indirect Detection Techniques

Indirect dark matter detectors are designed to detect the products of WIMP annihilation. This is dependent on the Majorana mechanism of WIMP self-annihilation to produce γ-rays and particle-antiparticle pairs. The rate of annihilation seen is dependent on the square of the density of dark matter. This being the case, indirect searches concentrate on areas of high mass such as the Galactic centre.

If WIMPs are Majorana particles, it is possible for them to annihilate to form standard model particles. Amongst these particles are anti-protons and positrons which will have an energy spectrum that extends to a maximum equivalent to the

mass of the annihilating particles. Many experiments exploit the fact that the annihilation positrons will, themselves, annihilate with electrons, emitting a telltale 511 keV γ-ray signature (equal to the rest energy of the electron/positron). There are, currently, several experiments looking for this signature, including EGRET [69], FERMI [70], INTEGRAL [71] and MAGIC [72].

Anti-particles can also be produced through the well known interactions between cosmic rays and the interstellar medium. These interactions produce a cosmic-ray background spectrum with an energy-dependent positron fraction. This spectrum is fairly well known and deviations from this could provide evidence for dark matter annihilation. However, any excess could also be due to positron emission from pulsars, so there would be difficulty claiming a detection.

The PAMELA [73] and ATIC [74] experiments have recently observed a positron excess above the expected model for the background (see Figure 3.1 for PAMELA results). The two collaborations have stated that an excess of positrons has been seen that can be attributed to either an astrophysical source or dark matter annihilation products. Some theoreticians have seized on this result, postulating a number of exotic theories that tie this result into a positive detection of dark matter annihilation. However, alternative models have been proposed in which the excess flux is consistent with well known pulsars [75].

3.2 Direct Detection Techniques

As discussed in Section 2.4.4.2, the favoured dark matter model gives a neutralino as the lightest supersymmetric particle that is able to interact weakly with matter. The opportunity to both detect the particle that provides the biggest contribution to the missing mass problem, and to give the first experimental evidence for particle physics beyond the standard model, means that many experiments have been devised to perform a direct detection of WIMPs.

A WIMP will either couple to the spin of a target nucleus or to its mass. These two cases are generally referred to as spin-dependent and spin-independent scattering respectively. Spin-independent scattering benefits from the coherent summation of the scattering amplitudes from each nucleon in the nucleus. Spin-dependent scattering, since paired nucleon spins sum to zero, is only possible on target nuclei with an unpaired nucleon. The particle shell model assumes that

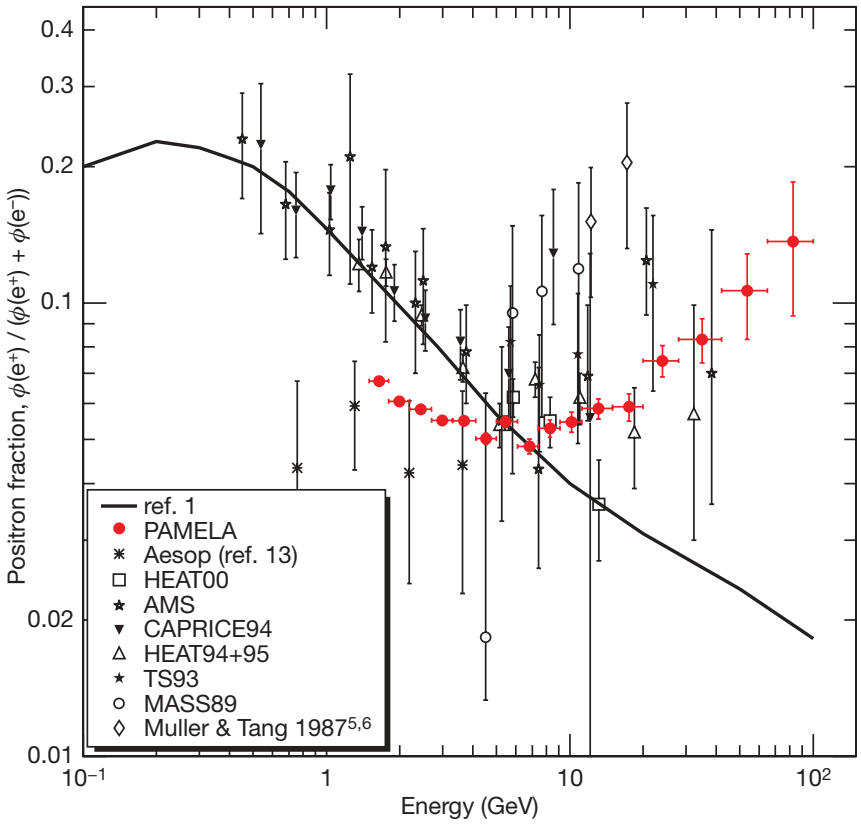

Figure 3.1: The positron fraction measured by the PAMELA experiment compared with other recent experimental data. The solid line shows a calculation for pure secondary production of positrons during the propagation of cosmic rays in the Galaxy. Error bars show 1 s.d.; if not visible, they lie inside the data points [73].

the net nuclear spin is due to the spin of the single unpaired protons or neutrons and, thus, vanishes for even nuclei. The WIMP may also couple to a combination of both the mass and the spin of the target nucleus.

Due to having a relatively large mass it is expected that WIMPs will interact via a WIMP-nuclear recoil. In the absence of excited states of the WIMP or the target nuclei at the low energies accessible due to the kinetic energy of the WIMP, this interaction is expected to be elastic and, as such, the energy transfer can be calculated via simple mechanics. The energy transfer spectrum is dependent on three main factors: the mass of the WIMP, the mass of the target material and the velocity of the WIMP. The dependence of rate on WIMP mass is shown in Figures 3.2.

The expected event rate can be calculated to first approximation quite easily using the following equation:

$$R \propto N_T \frac{\rho_0}{m_\chi} \sigma \langle v \rangle , \tag{3.1}$$

where R is the expected rate, N_T the number of target nuclei, ρ_0 the local dark matter number density, m_χ the WIMP mass, σ the WIMP-nuclear interaction cross section and $\langle v \rangle$ the average WIMP velocity.

This, however, is a rather simplified view as there are several astrophysical and nuclear physics factors that need to be taken into account. The energy dependent differential rate is given as:

$$\frac{dR}{dE_R} \propto N_T \frac{\rho_0}{m_\chi} \sigma \frac{d}{dE_R} \langle v \rangle , \tag{3.2}$$

which is usually described in terms of differential rate unit (dru) which is defined as counts/kg/keV/day. The average WIMP velocity is described by:

$$\langle v \rangle = \int_{v_{\min}}^{v_{\text{escape}}} v f(v) dv, \tag{3.3}$$

and, thus:

$$\frac{dR}{dE_R} \propto N_T \frac{\rho_0}{m_\chi} \sigma \int_{v_{\min}}^{v_{\text{escape}}} v f(v) dv. \tag{3.4}$$

The upper limit on the WIMP velocity is, formally, infinite but, in practice,

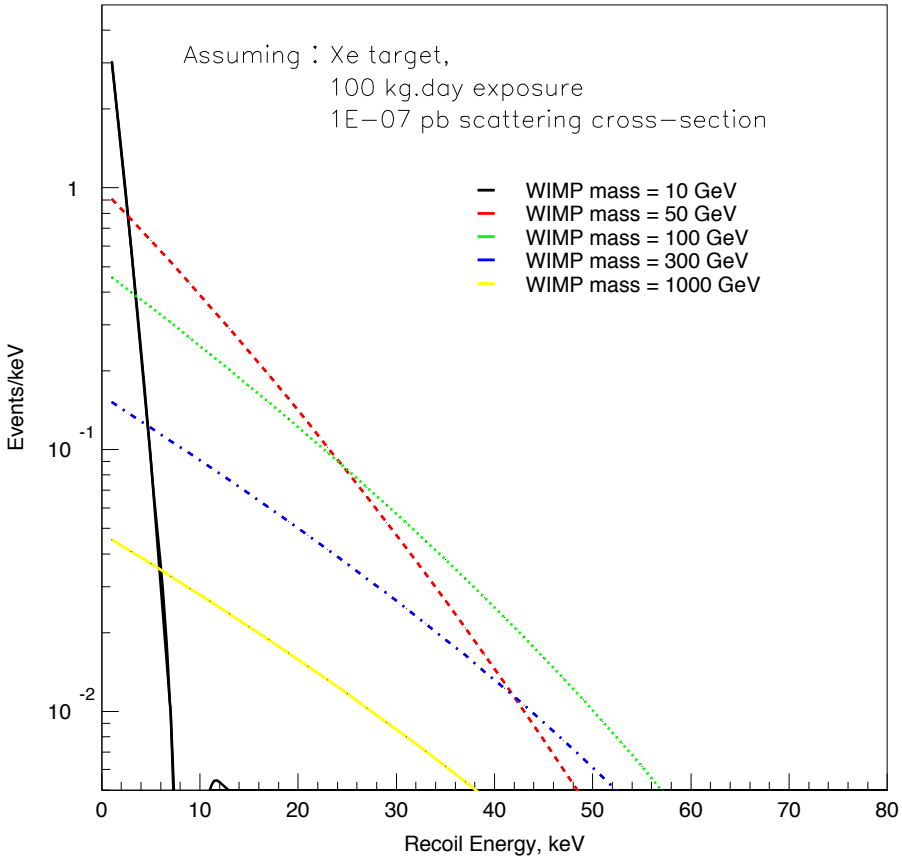

Figure 3.2: The expected differential energy spectra for WIMPs of different masses on a Xe target. The rates assume a spin-independent cross-section of 1×10^{-7} pb and 100 kg.day exposure.

the escape velocity from the Galactic halo is used where 498 km s^{-1} $< v_{\rm escape} <$ 608 km s^{-1} [76]. The parameter $v_{\rm min}$ defines the smallest velocity needed to produce a detectable energy (where E_R^{max} is the maximum recoil energy possible for the minimum velocity v_{min}) and is given by:

$$v_{\rm min} = \sqrt{\frac{E_R^{\rm max} m_N}{2\mu^2}}, \tag{3.5}$$

where m_N is the mass of the target nucleus and μ is the reduced mass of the WIMP-target nucleus system. For a head on collision, with $m_N = m_\chi = 100$ GeVc^{-2} and demanding a minimum nuclear recoil energy of 50 keV, a minimum WIMP velocity of $\sim$ 220 km s^{-1} is necessary.

Assuming a Maxwellian dark matter velocity distribution, the integral term in equation 3.4 can be rewritten:

$$\int_{v_{\rm min}}^{v_{\rm escape}} v f(v) dv \propto \exp\left(\frac{-v^2}{v_0^2}\right). \tag{3.6}$$

Since the WIMP velocity is related to the recoil energy E_R (equation 3.5), we find that the differential event rate for WIMPs falls exponentially with E_R:

$$\frac{dR}{dE_R} = \frac{R}{kE_0} \exp\left(\frac{-E_R}{kE_0}\right), \tag{3.7}$$

where R is the total rate, $E_0 = m_\chi v_0^2/2$ and $k = 4\mu/(m_\chi + m_N)$.

While equation 3.9 provides a good basis for the differential dark matter rate, there are two important corrections which must be considered. The first is a correction for the Earth's velocity relative to the galactic frame and the second corrects for the structure of the target nucleus.

To make a correction to the differential rate due to the velocity of the Earth, one must consider three velocities. $\vec{v}_{\chi,E}$ is the WIMP velocity in the Earth (or detector) frame, $\vec{v}_{\chi,G}$ is the WIMP velocity in the Galactic frame and $\vec{v}_{E,G}$ is the Earth's velocity in the Galactic frame of reference. Unsurprisingly, simple vector subtraction gives us a value for $\vec{v}_{\chi,E} = \vec{v}_{\chi,G} - \vec{v}_{E,G} = 220$ km s^{-1} [77, 78].

The second correction is due to the form-factor of the target nucleus. The form-factor describes the spatial extent of a nucleus. The scattering of electrons from a point-like nucleus is described via Rutherford scattering but this is only

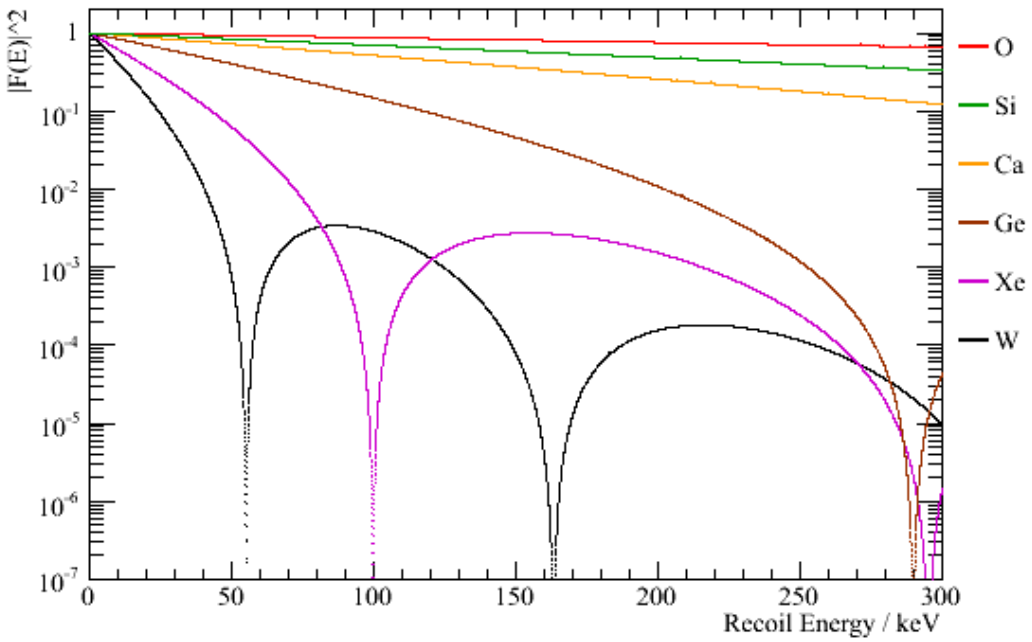

Figure 3.3: Form-factor measurements for several potential detector materials. For Xe, the characteristic diffraction minima can be seen at nuclear recoil energies of 100 keV and ∼300 keV.

correct for small values of momentum transfer (q). For large values of q, the cross-sections become:

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{exp.}} = \left(\frac{d\sigma}{d\Omega}\right)_{\text{Ruth.}} \cdot \left|F(q^2)\right|^2, \tag{3.8}$$

where $F(q^2)$ is the Fourier transform of the (charge or mass) nuclear density, σ is the scattering cross-section and Ω is the scattering angle.

As seen in Figure 3.3, the form-factor drops more slowly for lighter target nuclei, but this fact must also be combined with the results seen in Figure 3.4, where it can be seen that lighter target nuclei expect a lower initial rate [79]. Figure 3.3 also reveals another important factor in dark matter detection. It is essential for detectors to have very low thresholds to take advantage of $F(q)^2 \to 1$ at low recoil energies.

Having taken the Earth velocity and form-factor into account, the differential

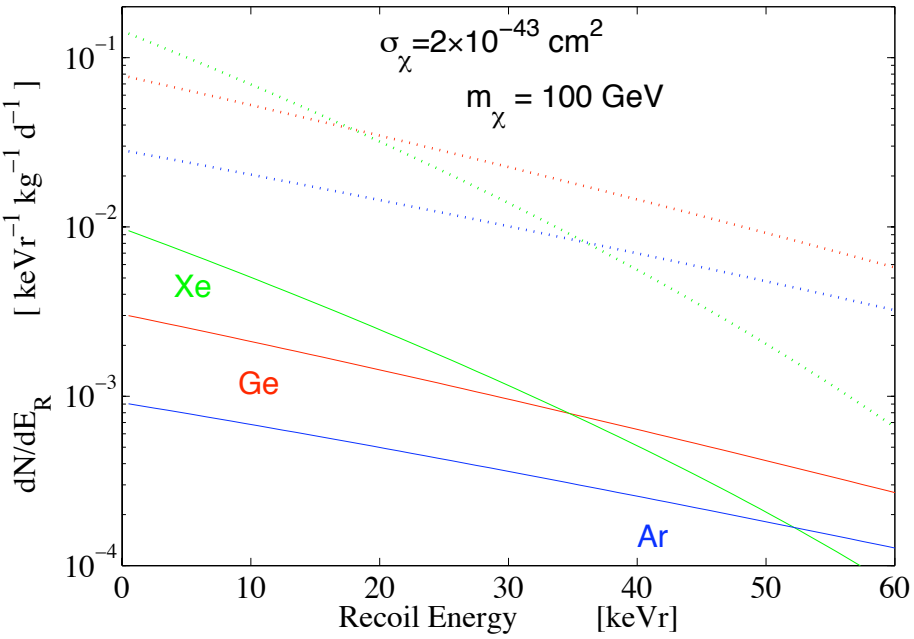

Figure 3.4: Expected event rate for several species of detector material assuming a WIMP mass of 100 GeVc^{-2} [85]. The solid lines represent the differential rate expected per unit energy and the dashed lines represent the integrated rate over all energies.

rate becomes:

$$\frac{dR}{dE_R} = \frac{R}{kE_0} F^2(E_R) \exp\left(\frac{-E_R}{kE_0}\right), \tag{3.9}$$

which shows the multiplicative nature of the form-factor correction. The differential rate can, therefore, be expected to drop as momentum transfer increases with minima at 100 keV and ∼300 keV for a Xe target.

There are numerous experiments world-wide which use various different detection techniques and detector media [80]. The current world leading results have been performed using Xe and Ge detection systems [81, 82, 83, 84].

Figure 3.4 shows the expected recoil spectra for various detector material given a WIMP with mass 100 GeVc^{-2} and a WIMP-nuclear cross-section of $\sigma_{\chi N} = 2 \times 10^{-7}$ pb. If Ge is taken as an example in this plot, it can be seen that the integrated rate above a 10 keV threshold is $\sim 5 \times 10^{-2}$ kg^{-1} day^{-1}. This is equivalent to a rate of 1 interaction per kilogram every 20 days.

Due to the extremely low event rates expected in dark matter detection, it is vital that all sources of background are removed. One of the largest potential sources of background is cosmic-rays. Any surface-based detector would be swamped with cosmic-rays and cosmic-ray induced backgrounds. This being the case, rare event physics (not just dark matter detection) is consigned to some of the deepest laboratories on earth. There are two possible methods for achieving the required depth. One is to base a laboratory in a deep mine and the other is to base a laboratory within a mountain. Each technique can provide thousands of metres of rock protection, often described using the units m.w.e. or metres of water equivalent. Reviews of current underground facilities are given in [86] and [87] and a plot of some of the larger facilities and their depth (m.w.e.) is given in Figure 3.5.

It is also essential that the detector itself, along with its shielding, must be constructed using low background materials (this is discussed in depth in Chapter 5). Finally, even the rock surrounding the facility can introduce unwanted levels of background into the detector, whether it be from rock-borne radioactivity or air-borne from radon. These additional backgrounds are minimised through the installation of shielding and the use of other techniques.

In direct dark matter detection, there are three different energy deposition channels that are used for detection: scintillation, ionisation and lattice vibrations. It is common (though not universal) practice to use two of these channels as an aid to discrimination between the nuclear recoils associated with a WIMP-nuclear event and electron recoils. Figure 3.6 shows the three energy deposition channels and the experiments that use one or two of them as methods for WIMP detection.

3.3 The ZEPLIN–III Detector

ZEPLIN–III (ZonEd Proportional scintillation in LIquid Noble gases) is a third generation detector designed, built and run in the UK. The ZEPLIN project was initially devised under the umbrella of the UKDMC (UK Dark Matter Collaboration) which included Imperial College London, the Rutherford Appleton Laboratory, University of Sheffield, Royal Holloway, Birkbeck College, Queen Mary London, Nottingham University and the University of Edinburgh. Initially, the

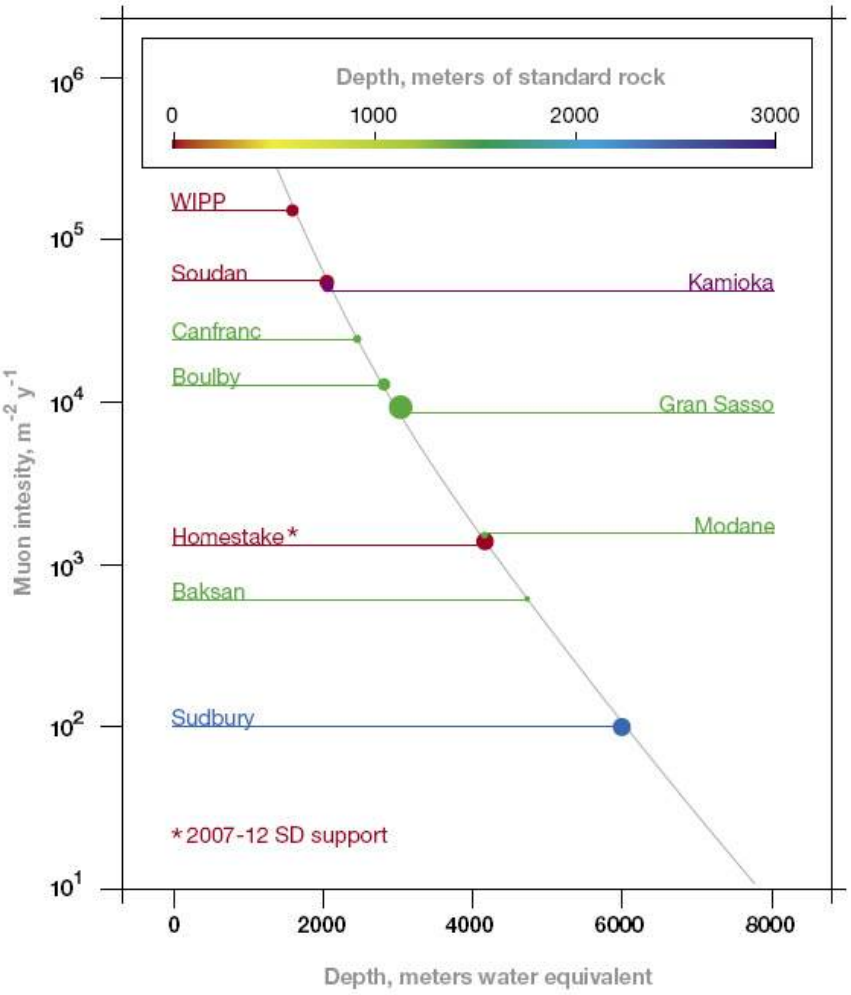

Figure 3.5: The muon intensity (number of muons per square metre per year) as a function of depth in some of the largest underground laboratories. The colouring of the data points gives an idea of the depth of the facilities in terms of rock overburden.

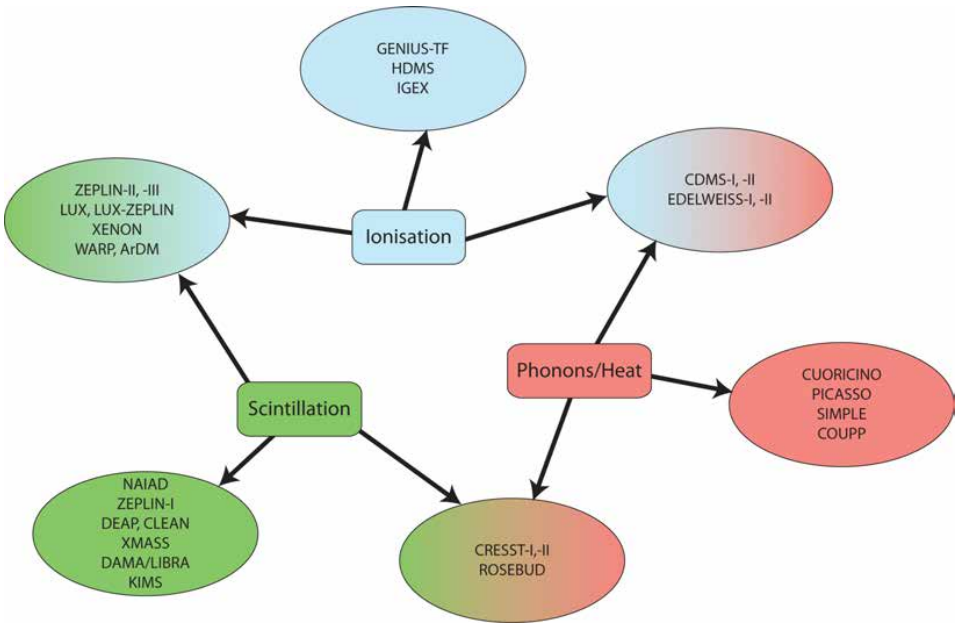

Figure 3.6: The three energy deposition channels for dark matter nuclear recoils: scintillation, ionisation and heat. A number of experiments that take advantage of the different channels are shown for comparison.

collaboration concentrated on using NaI as a detector material [88], leading to the NAIAD experiment (an array of NaI detectors) that gave published results in 2003 [89] and 2005 [90].

The inclusion of international collaborators (UCLA, Texas A&M, University of Rochester, LIP–Coimbra, ITEP Moscow) brought the knowledge, equipment and finance needed to allow the development of more complex detection systems with greater WIMP sensitivity. This marked the beginning of the use of liquid Xe as a detector material [91] along with (including collaboration with Occidental College, Los Angeles) the investigation of directional dark matter signals using the low pressure DRIFT detectors [92, 93].

ZEPLIN–I was the first of the three ZEPLIN detectors and was a single phase (scintillation only) detector which used pulse shape analysis to discriminate between nuclear recoils and electron recoils. It had a liquid Xe mass of 5 kg and operated from mid-2001 until late 2002. ZEPLIN–I set a world leading limit on the WIMP-nuclear cross section of 1.1×10^{-6} pb for an 80 GeVc^{-2} WIMP mass [94].

ZEPLIN–II was a dual phase (scintillation and ionisation) detector with a liquid Xe mass of 31 kg. ZEPLIN–II built upon R&D performed at CERN [95, 96, 97], Imperial College London [98] and ITEP [99]. The whole system was meticulously designed and simulated [100, 101, 102] resulting in an expected background of < 40 nuclear recoil events/year. However, when the data was analysed, it was found that there was much more background than expected. This was found to be due to radon progeny introduced at one of the initial purification stages. Although this had an effect on the final result, a strong limit of 6.6×10^{-7} pb [103] was published in 2007.

At the end of ZEPLIN–II, the UKDMC ceased to exist and the ZEPLIN programme was taken forward by the ZEPLIN–III collaboration which includes groups from Imperial College London, the STFC Rutherford Appleton Laboratory, the University of Edinburgh, LIP–Coimbra and ITEP Moscow. The author joined the collaboration at this point, so this thesis will concentrate on the research and work performed for, and results from, the ZEPLIN–III project, including design, simulation, calibration and data analysis.

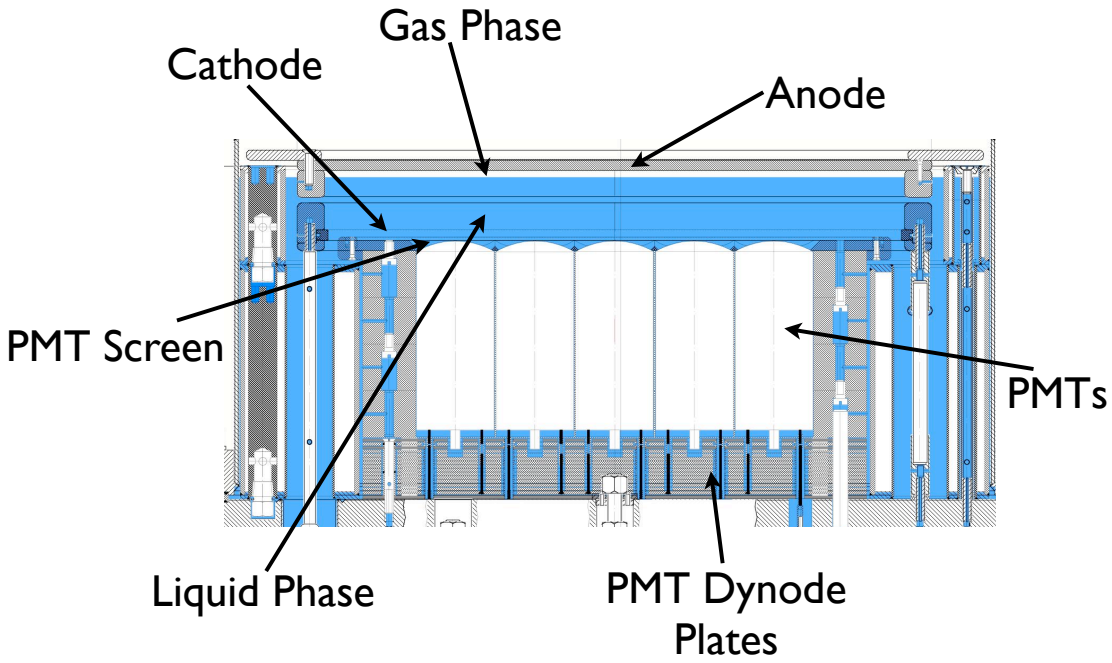

Figure 3.7: Close up view of the ZEPLIN–III target volume. The PMTs, anode, cathode and PMT grid are labelled. In addition to this, the gas and liquid phases of the target vessel are shown.

3.4 The ZEPLIN–III Detector

3.4.1 The Target Vessel

ZEPLIN–III was designed in parallel to ZEPLIN–II. It represented a step away from the ZEPLIN–II design which used a reasonably deep liquid phase (total liquid/gas thickness of 14 cm) and was viewed using PMTs above the volume. ZEPLIN–III, on the other hand, uses a shallow geometry in the liquid phase (total liquid/gas thickness of 40 mm) and the PMTs are submerged within the liquid, viewing the active volume from below. The PMTs were submerged in the liquid to allow a boost in light yield of the instrument which acts to lower the energy threshold for detection. This high light yield is extremely important as the expected energy spectrum for WIMP-nuclear scattering falls exponentially at low energies. A schematic view of the ZEPLIN–III target vessel is shown in

Figure 3.7

The electric field in the liquid and gas phase of ZEPLIN–III can be calculated fairly easily. All that is needed is the liquid level (or gas thickness) and the relative permittivity of the gas and liquid phases. For the liquid phase, the electric field is calculated as follows:

$$E_l = \frac{V}{\left(\frac{\epsilon_l}{\epsilon_g} \times d_g\right) + d_l}, \tag{3.10}$$

where ϵ_l and ϵ_g are the relative permittivities in the liquid and gas phases, respectively, and d_l and d_g are the thicknesses of the liquid and gas phases, respectively. Here V is the potential difference between the cathode and anode used to define the electric field across the Xe volume. The field in the gas is, roughly, double that in the liquid:

$$E_g = \frac{\epsilon_l}{\epsilon_g} \times E_l. \tag{3.11}$$

Unlike other two-phase detection systems, the drift field in ZEPLIN–III is attained by applying a voltage to a wire grid cathode and an anode mirror (the mirror allows light reflection to boost the ZEPLIN–III light yield further) without the need for a dedicated charge extraction grid or field shaping rings [104]. During the second science run (see Chapter 7), ZEPLIN–III operated using a liquid thickness of ∼36 mm, a gas thickness of ∼4 mm and an anode-cathode potential difference of 15 kV. Equations 3.10 and 3.11 yield an electric field in the liquid phase of 3.42 kV/cm and an electric field in the gas phase of 6.69 kV/cm. In the first science run, ZEPLIN–III operated with an electric field of 3.88 kV/cm in the liquid phase and 7.60 kV/cm in the gas phase. The geometry of ZEPLIN–III allows a very uniform electric field across the active region. In addition to the cathode and anode, a second grid, the PMT grid, is found below the cathode grid. The PMT grid is held at the same voltage as the PMTs and is designed to apply a reverse electric field (relative to the active volume) as an aid to rejection of low energy background radiation from the PMTs. Background from the PMTs would produce scintillation signals but electrons would not be able to drift to the liquid surface. This would give rise to signals that contained a primary but no secondary signal. Figure 3.8 shows a simulation of the electric field in ZEPLIN–III

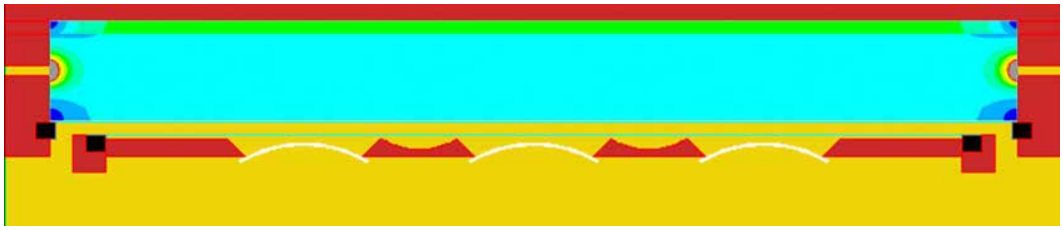

Figure 3.8: A simulation of the electric field found in the active volume in ZEPLIN–III demonstrating the excellent uniformity in the central region. The liquid phase (blue) is at 3.88 kV/cm and the gas phase (green) is at 7.60 kV/cm. The yellow region, directly below the liquid phase, is the reverse field region. Edge effects due to detector surfaces can be observed, highlighting the need for fiducialisation in analysis [105].

and highlights the uniformity.

In order to reduce the levels of background in the detector, the vacuum and target vessels are constructed using oxygen free copper. To form the dome, the copper has been electron beam welded. This is a process whereby a beam of high velocity electrons is used to melt separate pieces of the copper material together. This allows a very clean weld with none of the (relatively) radioactive deposits found in standard welding. Other components used in the construction of the detector have been radio-assayed in order to ascertain and minimise their individual contributions to the radiation budget for ZEPLIN–III (see Chapter 5 for a full description of how the radiation budget is calculated). During the first science run of ZEPLIN–III, 31 2-inch low background PMTs (ETL D730/9829Q) were used. The PMT array provided excellent 3D position reconstruction which helps to aid the rejection of background from surfaces and multiply scattered events (remembering that WIMPs are not expected to multiply scatter due to their low interaction cross-section). The performance of each PMT was characterised with low temperature tests carried out at Imperial College [106], before the array was installed into the detector. The PMTs are surrounded by a tight fitting PMT screen which aids the reduction of PMT crosstalk and reduces the possibility of scintillation occurring in volumes between the PMTs being detected. The PMT array is shown in Figure **??**. Also shown in Figure 3.9 are the numerical values for each pmt.

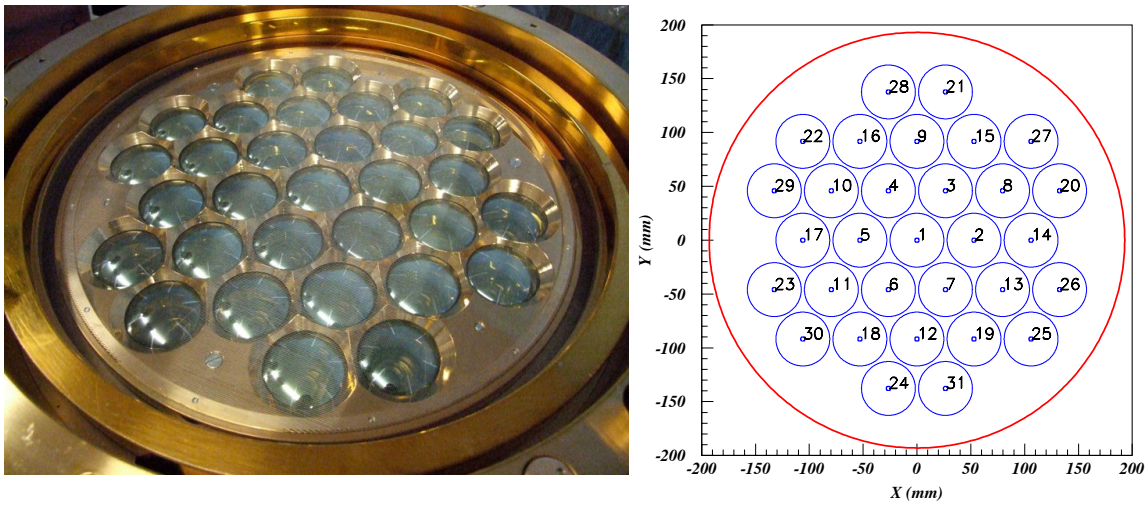

Figure 3.9: The ZEPLIN–III PMT array shown photographically (left) and graphically (right). The copper PMT screen which acts to isolate each PMT from its neighbours is seen in the photograph.

In the first science run, the PMT dominated background in the detector. This being the case, an additional set of ultra-low background PMTs was developed in partnership with ETEL (Electron Tube Enterprises Ltd) designed to reduce the overall background in ZEPLIN–III by a factor of at least 10. Individual components for the PMTs were radio-assayed in order to find the cleanest possible components (see Chapter 5).

3.4.2 Detector Cooling

The Xe is maintained in liquid phase using liquid N_2. There is an internal liquid N_2 reservoir located below the target vessel but within the vacuum jacket (to maintain thermal isolation) capable of holding 36 litres, enough to cool the detector for 1 day without the need for refilling. This is advantageous as it has been observed that the refill process induces spurious signals in the data. There are two thermal links between this reservoir and the target vessel. The first link is a conduction path provided by flexible bundles of thick copper wires thermally anchored to a hollow copper cooling flange attached to the underside of the target vessel. The copper wires are intentionally flexible to decouple acoustic and

mechanical noise from the liquid N_2 reservoir in the target vessel. The other end of the bundle dips into the liquid nitrogen and the thermal impedance depends on the depth of the liquid.

A second thermal path is provided by a direct connection between the nitrogen reservoir and a hollow cooling flange. This allows cold boil-off gas to be used as additional coolant and provides the means for active thermal control with minimum cryogen usage, which is important during stand-alone operation underground.

The thermal stability of the system is of paramount importance to the quality of the data produced. Excellent thermal stability is reflected in the pressure stability of the system and, hence, the stability of the liquid level and gas phase thickness and density. The temperature control unit monitors the temperature of the liquid N_2 input braid. If this temperature exceeds the temperature set-point defined by the unit, the control valves are oriented such that liquid N_2 gas boil-off reaches the cooling flange and the temperature begins to decrease. As the temperature falls below the set-point, the valves reverse and cooling is no longer provided. Naturally, the input braid begins to heat up. The cooling system allows the braid temperature to be maintained to a stability of ΔT<1 K with an oscillation period of ∼30 minutes (seen in Figure 3.10). The time constant for a reaction to temperature change for the pressure in ZEPLIN–III is much longer than the period of the temperature oscillation so pressure remains constant to within 0.01 bar. Figure 3.11 shows the stability of the system. Any pressure variation in the corrected pressure measurement (labelled TARC in Figure 3.11) is corrected in data analysis.

3.5 Liquid Xe as a Detection Medium

Liquid Xe has a number of properties that make it a good choice for dark matter detection. Liquid Xe has a very high scintillation yield and is transparent to its own scintillation light which means, as far as scintillation light is concerned, large volumes can be used without sensitivity loss. Liquid Xe also scintillates in the middle of the UV spectrum (λ = 175 nm) [107] which means the light can be detected by photo-multiplier tubes (PMTs) with reasonable quantum efficiency without the need for wavelength shifting materials [108].

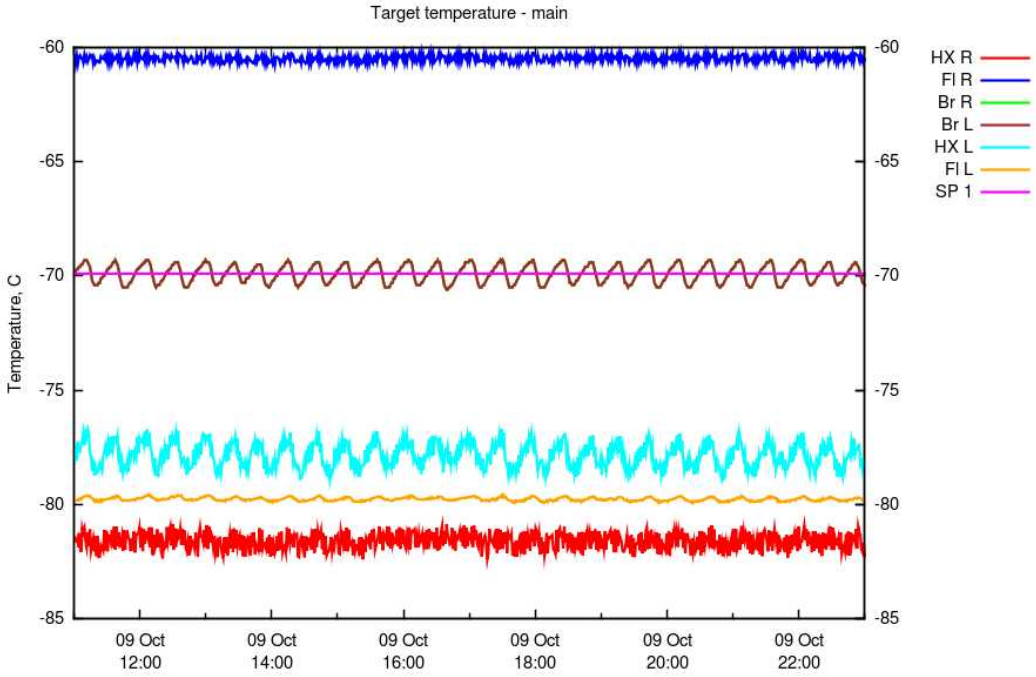

Figure 3.10: Plots of the temperature within ZEPLIN–III over a 12 hour period. The temperature set-point is denoted by the pink "SP 1" line and the braid temperature by the brown "Br L" line. The 30 minute periodicity can easily be seen.

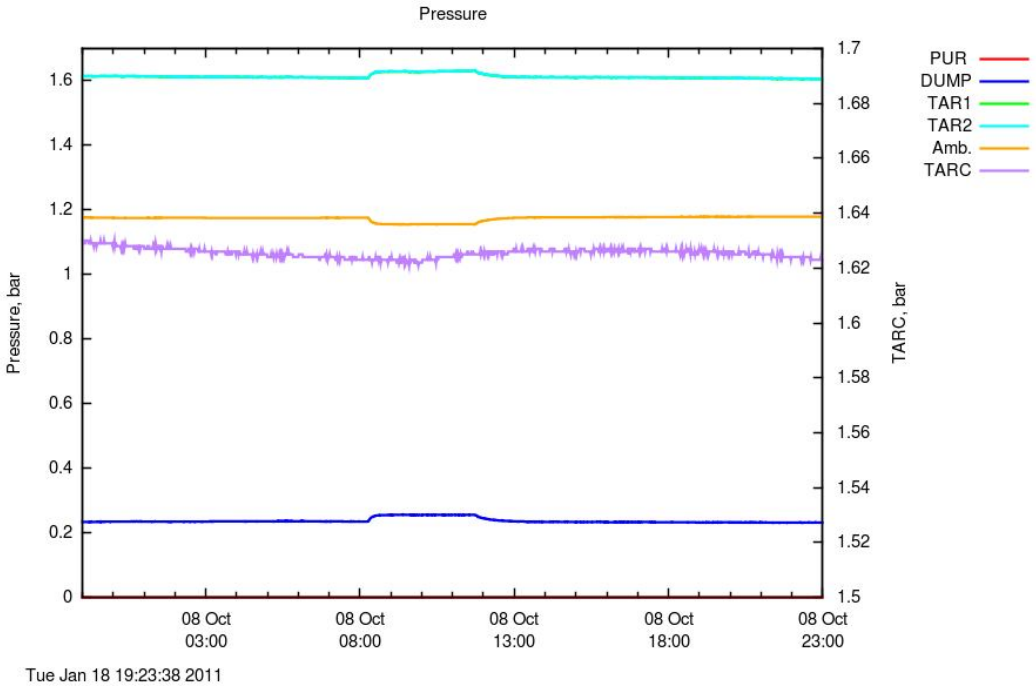

Figure 3.11: Plot of the pressure within ZEPLIN–III over a 12 hour period. The pressure within the target vessel (light blue) varies very slowly. The plot shows an apparent deviation in the pressure of the vessel but this is mirrored in the ambient pressure reading and the purple line (TARC, measured on the right-hand axis) shows the target pressure corrected for changes in ambient pressure which remains approximately flat to within 0.01 bar.

Very high sensitivity is also achieved in the ionisation channel. It is possible to drift free electrons through large volumes of liquid with the only limiting factor being electronegative impurities in the liquid. It is also important to note that Xe is itself relatively radio-pure, intrinsically containing no long lived isotopes. The longest lived isotope in naturally occurring Xe is ^{127}Xe with a half-life $t_{1/2} = 36$ days. This isotope will β^- decay with a typical endpoint energy of $E_{\beta^-} > 287$ keV. The decay is accompanied by at least 1 γ-ray with typically $E_\gamma > 145$ keV. The isotope ^{127}Xe may be produced by the neutron activation of ^{126}Xe, though the isotopic abundance in natural Xe is only $\sim$0.01%. Another form of internal radiation may be from the double-β^- decay of ^{136}Xe which is expected to have a Q value of 2458 keV [109]. The currently measured upper limit on the half-life of this reaction (either 2ν or 0ν) is $t_{1/2} > 4.4 \times 10^{23}$ yr [110].

Another source of background to be aware of is the ^{85}Kr contamination found in all commercial Xe. ^{85}Kr is a β^- emitting isotope with a half-life of $t_{1/2} = 10.8$ yr. The Xe used in ZEPLIN–III is Xe sourced from ITEP Moscow and has an age of >40 years, meaning it has been through $\sim$4 half-lives, greatly reducing the ^{85}Kr contamination. ^{85}Kr is an anthropogenic isotope, related to the nuclear power and weapons industry. This being the case, the purest Xe would be that extracted from the atmosphere prior to the nuclear age.

Liquid Xe is a dense material ($\sim$ 3 g/cm) meaning that a large mass can occupy a relatively small volume. This reduces the size of vessel needed to hold a specific mass when compared to other media such as Ar. This relative reduction in material helps to reduce unwanted backgrounds in a detection system. Large masses of Xe also provide an excellent level of self-shielding. Low energy and surface nuclear recoil backgrounds are absorbed by the outer regions of Xe, leaving an ultra-low background volume of Xe in the centre. This Xe self shielding has been exploited by the XMASS collaboration [111] to produce a single phase liquid Xe system with high sensitivity.

3.5.1 Available Signal and Primary Scintillation

Energy deposition via particle interactions in Xe leave atoms in both excited and ionised states. Excited Xe atoms combine with un-excited Xe atoms to form an exciton in a matter of a few picoseconds:

$$Xe^* + Xe \rightarrow Xe_2^*, \tag{3.12}$$

which decays to:

$$Xe_2^* \rightarrow 2Xe + \gamma, \tag{3.13}$$

where γ is an ultraviolet photon of wavelength 178 nm [107]. This process produces scintillation light with two different time constants, 3 ns for the singlet and 27 ns for the triplet [91] known as 'exciton' luminescence.

At the same time, ionised Xe atoms combine with non-ionised Xe atoms to form an Xe_2^+ molecule [112]:

$$Xe^+ + Xe \rightarrow Xe_2^+. \tag{3.14}$$

In the absence of an electric field, this recombines to form an excited Xe atom and heat:

$$Xe_2^+ + e^- \rightarrow Xe^* + \text{heat}, \tag{3.15}$$

which will then release an additional scintillation photon following equations 3.12 and 3.13. This process is known as 'recombination' luminescence. The two routes for luminescence become important when trying to discriminate between electron and nuclear recoils at zero electric field. It has been observed that the ratio between the two states depends on the type of particle interaction. 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 ionisation track topology, recombination luminescence becomes important, yielding an apparent time constant in the region of 24-46 ns [113, 114] (see Figure 3.12). In this figure, 0-field background data is presented and the time constant of $\sim$30 ns is in excellent agreement with the published range.

Another important factor to consider is the variation in scintillation yield given by electron recoils and nuclear recoils. The energy loss by nuclear collisions and the higher excitation density of nuclear recoils compared to electron recoils are the factors that decide this difference, known as relative scintillation efficiency (L_{eff}).

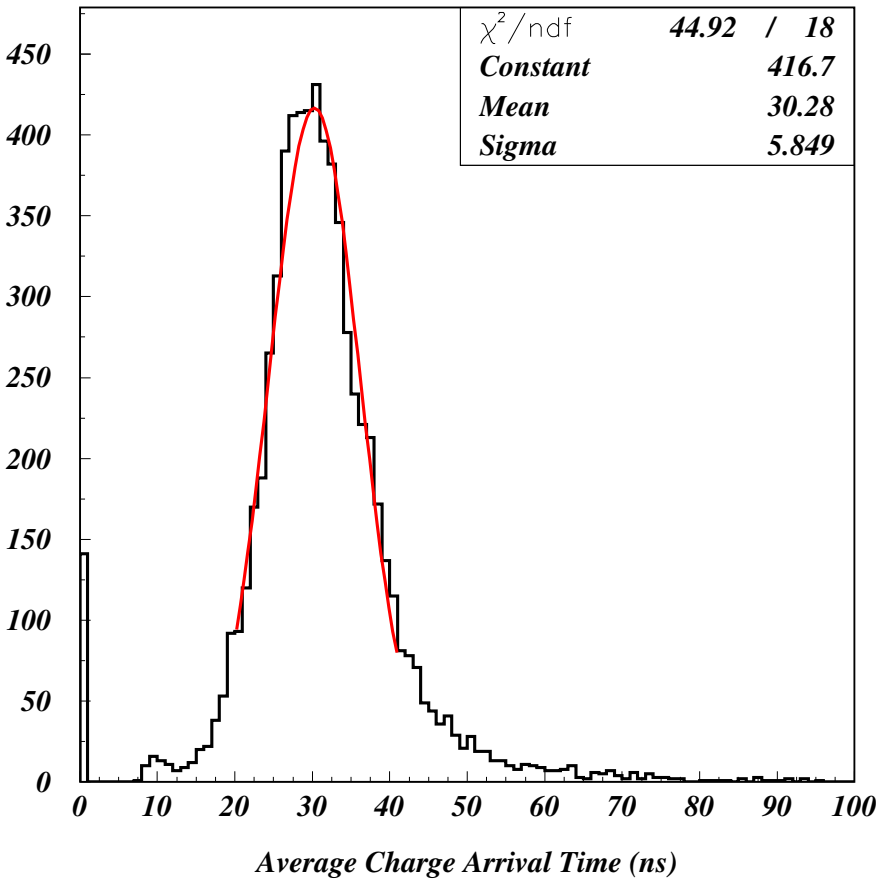

Figure 3.12: The average charge arrival distribution for ZEPLIN–III pulses that are unsaturated. The low background nature of the detector means that the distribution is completely dominated by electron recoil recombination luminescence with an average time of ∼30 ns.

Although dark matter detectors search for nuclear recoil events, mono-energetic calibration is performed using γ-ray sources. Due to this, it is important to understand the conversion between electron recoil energy (E_{ee}) and nuclear recoil energy (E_{nr}). The difference between the energy deposited in a nuclear recoil and an electron recoil is known as either the 'quenching' or the Lindhard effect [115].

The energy scales E_{ee} and E_{nr} are related to each-other as follows:

$$E_{\mathrm{nr}} = \frac{E_{\mathrm{ee}}}{L_{\mathrm{eff}}} \times \frac{S_e}{S_n}, \tag{3.16}$$

where S_e and S_n are field dependent suppression factors for electron and nuclear recoils, respectively.

The measurement of L_{eff} is not simple. Until relatively recently, measurements of L_{eff} had only been taken down to $\sim$ 10 $\mathrm{keV_{nr}}$ [116, 117, 118, 119, 120]. New measurements, however, see signals down to 5 $\mathrm{keV_{nr}}$ [121] (see Figure 3.13 for details of measurements). Measurements down to 10 $\mathrm{keV_{nr}}$ suggested that the value for L_{eff} was, roughly, a constant 0.2. The more recent measurements, however, suggest that L_{eff} may begin to fall at even lower energies. These measurements rely on a linear response to electron recoil energy down to low energies which has not yet been conclusively proven.

A potential reason for the non-linearity seen in L_{eff} measurements has been proposed by Manzur *et al* [124]. They propose that L_{eff} can be split into three discrete effects:

$$L_{\mathrm{eff}} = q_{\mathrm{ncl}} \times q_{\mathrm{esc}} \times q_{\mathrm{el}}. \tag{3.17}$$

The first term of the equation, q_{ncl}, describes the Lindhard effect. The second term, q_{esc}, describes the reduction of light yield due to the escape of electrons produced by ionisation that thermalise outside the Onsager radius and therefore become free from recombination, even in the absence of an electric field [125]. This effect had been observed in electron recoils but it is only relatively recently that it has been suggested the same effect may be a factor in the scintillation yield reduction for nuclear recoils [124]. The final term, q_{el}, describes the scintillation light quenched by biexcitonic collision which describes why the measured L_{eff} are much smaller than those calculated using q_{ncl} only [123]. Biexcitonic collisions involve two excitons combining and producing only a single photon instead of

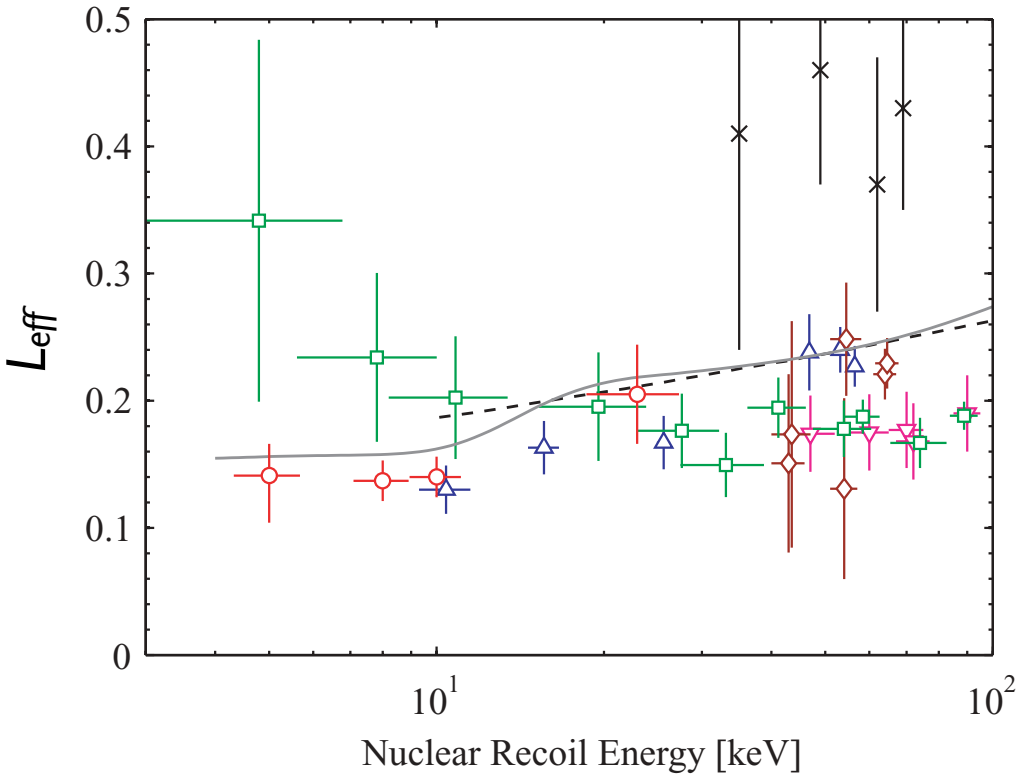

Figure 3.13: The relative scintillation yield in LXe (L_{eff}) measured by several experiments. The different values are as follows: (○) - Aprile *et al*, 2009 [121]; (△) - Aprile *et al*, 2005 [119]; (□) - Chepel *et al*, 2006 [120]; (◇) - Akimov *et al*, 2002 [118]; (▽) - Arneodo *et al*, 2000 [116]; (×) - Bernabei *et al*, 2001 [117]. The solid line is the result from a best fit of the XENON10 AmBe experimental and Monte Carlo data [122]. Also shown is the theoretical prediction of Hitachi (dashed line) [123]

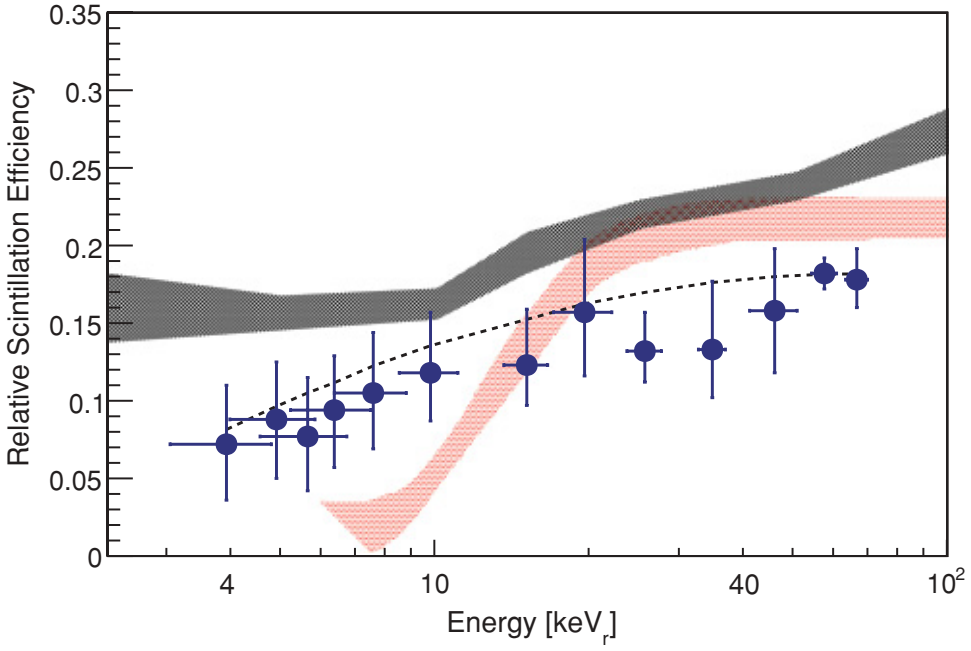

Figure 3.14: L_{eff} values measured by ZEPLIN–III (red shade), XENON10 (grey shade) and Manzur *et al*. The dashed line shows the theoretical value calculated using equation 4.2.

two. These three effects combine to form a good description of a value for L_{eff} which reduces at low energies. Figure 3.14 shows recent data from XENON10, ZEPLIN–III and Manzur *et al* [124]. The ZEPLIN–III result is discussed in more detail in Section 4.10.

3.5.2 Secondary Ionisation Signal

Under an electric field, the ion recombination process is suppressed. This means that the scintillation yield is quenched when compared to its zero field value. By applying an electric field, it is possible to extract electrons from the interaction site and these can then be detected independently. In detectors such as ZEPLIN–III, this is achieved through detection of electroluminescence in which the ionisation signal is converted into proportional photon signal in a gaseous Xe phase [126].

Given standard ZEPLIN–III operating parameters of ∼3.5 kV/cm, electrons

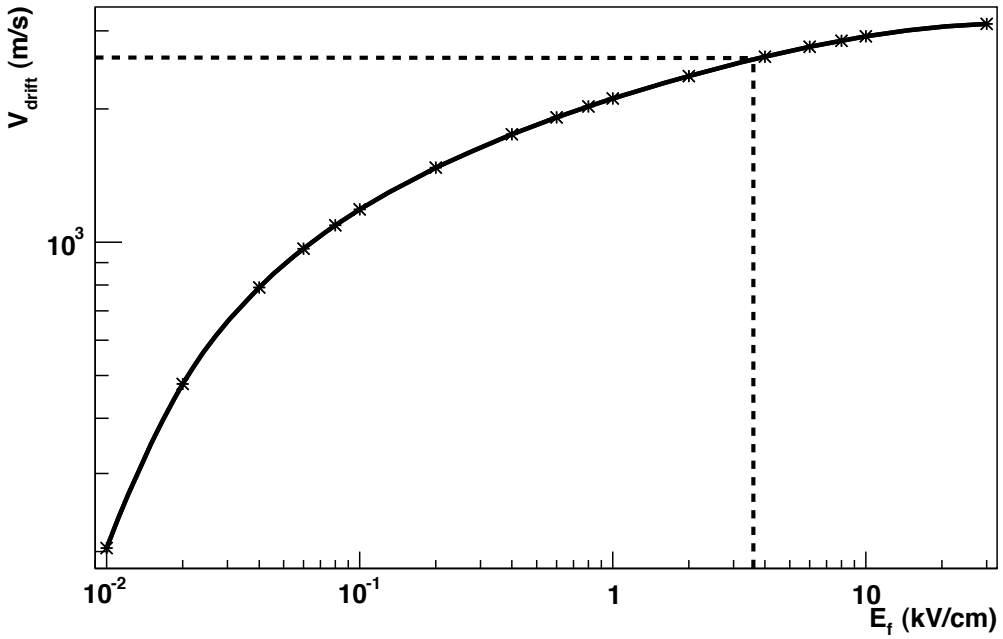

Figure 3.15: The measured electron drift velocity versus electric field in liquid Xe. The equivalent drift velocity for a 3.5 kV/cm drift field is marked.

are drifted through the liquid Xe phase at a velocity of ∼2300 m s^{-1} (see Figure 3.15). The field dependence of suppression varies for different species (see Figure 3.16).When looking at this figure, it is easy to see that the light from electron recoils is heavily suppressed due to the dominance of recombination luminescence. The free electrons that would recombine with the nucleus to form scintillation are instead drifted to the gas phase and ionised. It is no surprise that the ionisation yield for electron recoils increase is inversely proportional to the fall in scintillation yield. It is also apparent that the dependence of scintillation and ionisation yield on electric field strength is much weaker for nuclear recoils, since the scintillation from nuclear recoils is dominated by exciton luminescence. The scintillation from the other heavy species, α-particles, is also dominated by exciton luminescence and therefore displays the same weak field dependence as nuclear recoils.

As electrons drift through the liquid phase towards the gas phase, they may be trapped by electronegative contaminants such as O_2 and N_2O, amongst others. The capture cross-section for these contaminants varies with electric field. The cross-section for capture by O_2 *decreases* with electric field whilst the cross-section for N_2O *increases* with field. The various measurements of the effect of contaminants in liquid noble gases (including Xe) are discussed in [128].

Once the electrons reach the liquid surface, the electric field can lead to cross-phase emission. The phase change acts as a potential barrier to the electrons and they will only be able to cross if they have been given enough energy by the electric field. This being the case, the fraction of electrons extracted to the gas phase depends on the electric field strength. With a field of 5 kV/cm, electrons reaching the liquid surface can be extracted with nearly 100% efficiency [129]. This can be seen clearly in Figure 3.17.

Electrons that overcome the liquid/gas barrier are accelerated through the high electric field in the gas and, occasionally, will collide with a Xe atom. This excites the Xe atoms, generating electroluminescence (proportional scintillation):

$$\begin{aligned} e + Xe &\rightarrow e + Xe^*, \\ Xe^* + 2Xe &\rightarrow Xe_2^* + Xe, \\ Xe_2^* &\rightarrow 2Xe + h\nu, \end{aligned} \tag{3.18}$$

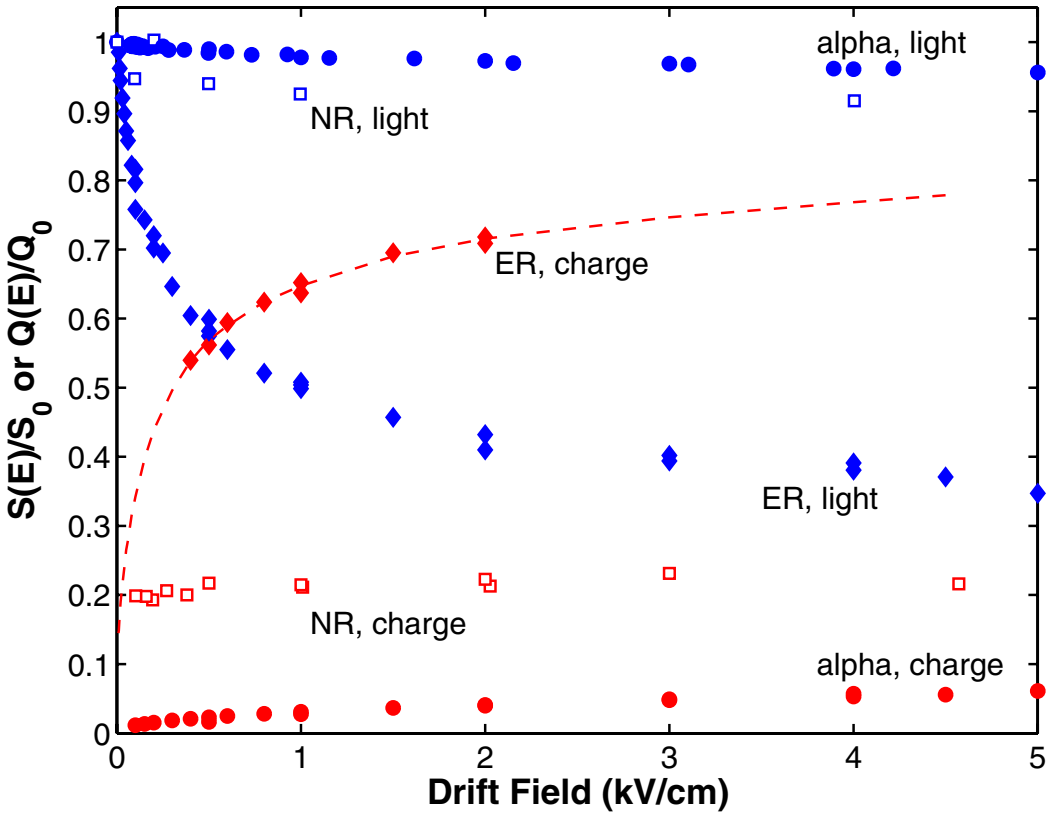

Figure 3.16: Field dependence of scintillation and ionisation yield in LXe for 122 $\mathrm{keV_{ee}}$ electron recoils (ER), 56.5 $\mathrm{keV_{nr}}$ nuclear recoils (NR) and alphas [127].

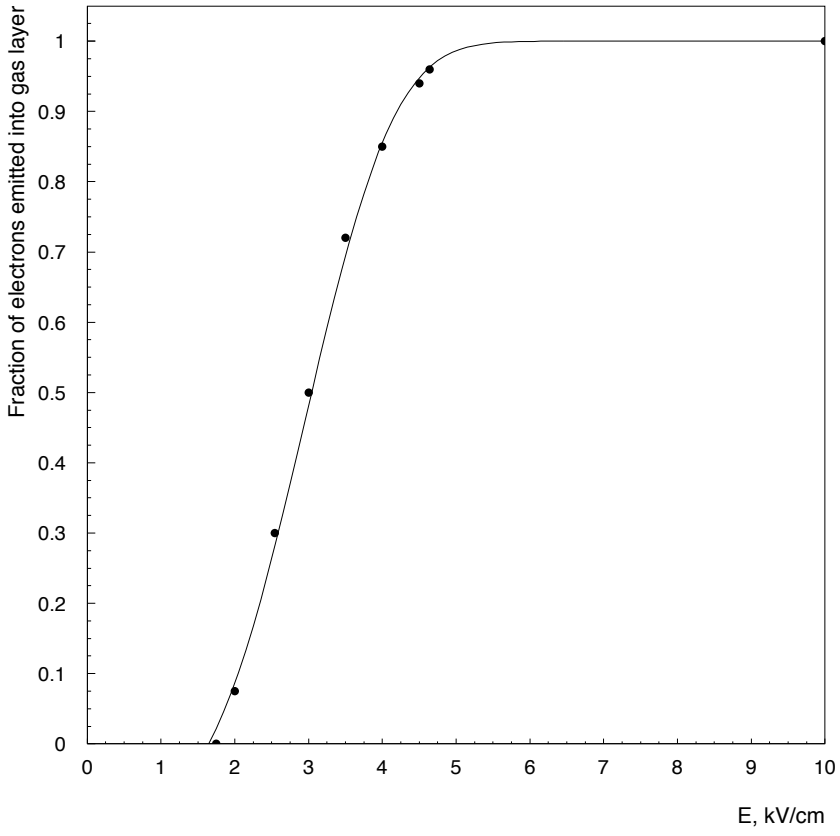

Figure 3.17: Measurement of the cross-phase emission probability for electrons in liquid Xe [129]. The line shows the parameterisation utilised in ZEPLIN–III detector simulations [130].

where $h\nu$ is a delayed electroluminescence signal. The energy deposited in the gas phase is much higher than the energy deposited in the liquid phase, giving rise to secondary signals which are large when compared to primary signals.

Contamination can occur in the gas phase of the detector which will act to suppress the ionisation yield when the detector is operational with its nominal electric field. A contamination of $<1\%$ of N_2 can cause the drift velocity in the gas phase to increase significantly, the increase being most prominent at low electric field values. This means that electrons drift through the gaseous phase of the detector much more quickly, reducing the amplitude of the ionisation signal. At values of electric field where extraction should begin to occur (around 2 kV/cm, see Figure 3.17) in contaminated Xe, the electroluminescence signal size will be much less than expected in pure Xe. This is because the time allowed for energy deposition will decrease and, hence, the signal amplitude will decrease. Figure 3.18 shows the drift velocity for various mixtures of gases, clearly showing that small concentrations of nitrogen can have a big effect on drift velocity. Such an effect was seen in the commissioning period of the second science run.

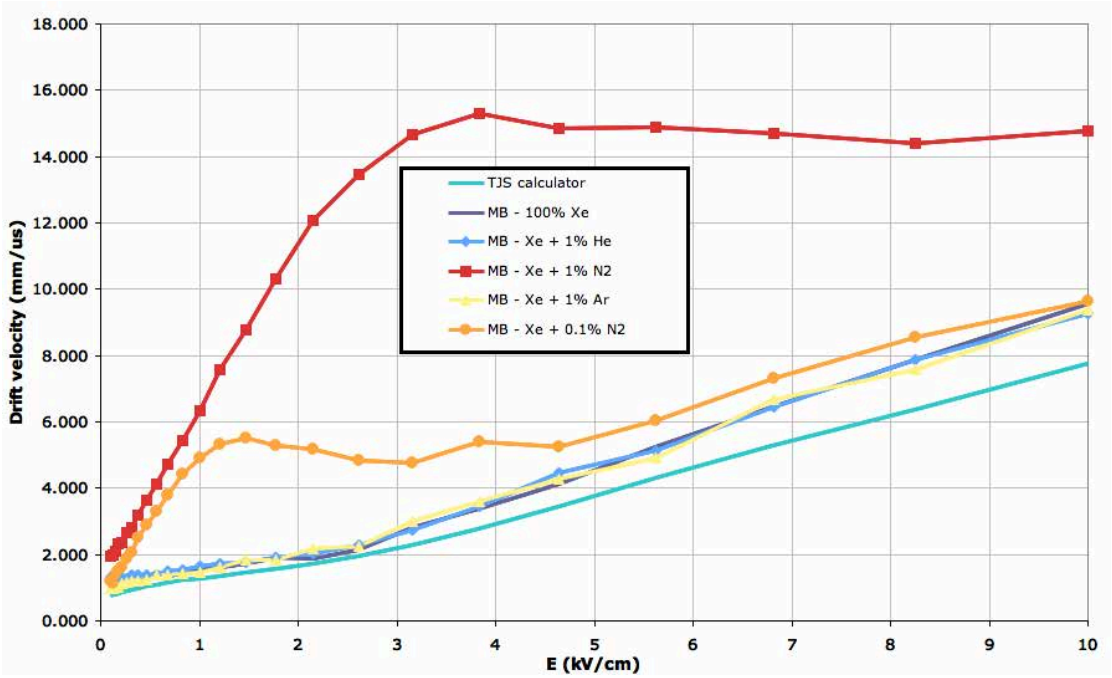

Figure 3.18: Values for electron drift velocity in Xe gas given various levels of contamination as calculated by the Magboltz software (which calculates the transportation of electrons in gas mixtures) [131] in comparison with values calculated for pure Xe (TJS calculator). N_2 gas gives the most significant increase in drift velocity [132].

Chapter 4

ZEPLIN–III First Science Run

The First Science Run (FSR) of the ZEPLIN-III dark matter search experiment started during February 2008 and continued through to June 2008. It consisted of 84 days of high quality, shielded, science data and daily calibrations. Section 3.4.1 gives an overview of the ZEPLIN-III construction though a more detailed description of the hardware is available in [133].

The collaborative nature of the ZEPLIN-III project means that some of the work within this chapter has been led by other people. However, I personally played a major role in most aspects of the FSR including the debugging and optimisation of the data reduction code (ZE3RA), optimisation of the event selecting and parameterising golden code, the ongoing daily analysis of ^{57}Co calibration data, the studying of mis-parameterised events and, most significantly, analysis of the calibration and science data. In order to give a more complete overview of the science run, I have included information about all aspects of the run.

4.1 Data Acquisition

The PMT array of ZEPLIN-III consists of 31 elements; each PMT signal is split between amplified and non-amplified channels, equating to 62 channels of data which must be recorded. The data acquisition is performed using 8-bit ACQIRIS digitisers sampling at a rate of 500 MHz. Waveforms are centred on the trigger point (the system triggers using the sum of the amplified channels) and extend 18 μs before and after. This means that the full drift time of ZEPLIN-III is

included in any waveform. An event that triggers from a secondary gas ionisation (S2) signal with a drift time of 15 μs will have a primary liquid scintillation (S1) signal at $\sim$-15 μs and an event that triggers from an S1 signal with the same drift time would have a S2 signal beginning at $\sim$15 μs that would not extend beyond the maximum 18 μs timeline.

The signals from the 31 PMTs are normalised by passing through attenuators and are then split. One channel is recorded un-amplified as a low sensitivity (LS) channel and the other is amplified by a fast 10x amplifier and is recorded as a high sensitivity (HS) channel. This dual channel recording allows for the recording of large secondary pulses whilst being able to maintain a high level of sensitivity for small primary signals. In software, it is possible to select whether the HS channel or LS channel is used in event parameterisation based on the size of the pulse being studied and if amplifier saturation has occurred.

The amplified HS signals are passed into a summing amplifier before integration (500 ns) and differentiation (500 ns) by a timing filter amplifier. This signal passes to a fast discriminator, set to a -80 mV threshold, which outputs to an external gate generator triggering the digitisers.

This unit is inhibited for 1 ms after each event. Studies of the trigger have shown that the trigger threshold is equivalent to $\sim$11 ionisation electrons emitted from the liquid surface, or $\sim$0.2 keV for nuclear recoils. The threshold was set at this level to avoid excess triggering from single electron emission. The waveforms are recorded by custom LINUX software that writes to file using two processor cores to reduce detector dead time.

Data is then compressed and transported on data tapes to RAL where it is uploaded to a central data repository. The data is reduced using an Apple XGRID system with 28 processor cores combining to give an overall processing speed of 72 GHz. During the FSR, 15.5 TB of compressed data were acquired.

4.2 Pulse Identification

Pulse identification in ZEPLIN-III is performed using custom designed LINUX software. This package is called ZE3RA (**ZE**PLIN **3** **R**eduction & **A**nalysis) and was developed by collaborators at LIP–Coimbra. ZE3RA is designed to ascribe pulse parameters to raw waveforms without making any sort of physical

assumptions. This method of pulse analysis removes any bias in the data analysis process which could jeopardise the blindness of the result. The ZE3RA package includes an event display allowing instant visualisation of recorded events.

The first stage of the reduction process involves an accurate determination of the baseline of the waveform. A 2 μs sample is extracted from the start of the timeline and is used to determine the mean of the baseline as well as the noise level (RMS). The baseline may be misidentified if a pulse with a large width is located at the beginning of the timeline. However, this is not an issue as any waveform displaying these characteristics will not pass the event selection criteria described in Section 4.3. The baseline of each channel is determined individually and a software trigger threshold of 3 $\times$ RMS set. This means pulses must have size $> 3\times$ RMS to be considered for parameterisation.

Once a baseline has been identified and a threshold characterised, pulse identification can proceed. Pulse identification is implemented on both the HS and LS channels giving two sets of pulse parameterisations per ZEPLIN–III event. This is particularly important for small signals where pulses in individual channels could be washed out by PMT noise in the LS channel but would be identifiable in the HS channel and also for large signals that would saturate the HS amplifiers but would still give a clear and complete signal in the LS channel. In addition to individual channels, a HS and LS sum channel is computed for quick identification of events.

Once parameterisation is complete, the pulses are ordered by decreasing area and the 10 largest are retained and written to an *hbook* n-tuple. Within the n-tuple, pulses are identified by two indices (pulse = 1-10, channel = 1-64) and waveforms by a single index (channel = 1-64). The sum channels are identified as indices 32 and 64 for HS and LS, respectively. The parameters identified and recorded by ZE3RA are shown in Table 4.1.

4.3 Event Selection

The event selection process is completed using a so-called "golden code". The code is designed to select events containing a single S1 and a single S2 event. The golden code assumes that WIMPs will only scatter once in the detector (the probability of multiple scatters is vanishingly small, given the elastic scatter

Parameter Block	Parameters
Event	Event number and time
Detector Slow Control	Trigger rate, detector temperature and pressures, liquid level
High Voltage	Electrode and PMT voltages and current
External Environment	Lab temperature and pressure
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

Table 4.1: Parameters recorded in the ZE3RA *hbook* n-tuple.

cross-section) giving only one S1 and one S2.

The event selection process can be split into 4 distinct sections:

- **First Pass Cuts - Golden rule:** to select one S1 and one S2.
- **First Pass Cuts - Voltage cuts:** to remove any data with spurious anode, cathode or PMT voltages or currents.
- **First Pass Cuts - Waveform cuts:** to remove any events with wild waveforms, HS/LS inconsistencies and DAq mistriggers.
- **Second Pass Cuts** - Refine event selection using newly calculated parameters.

Any event that passes all these cuts is then written to a "golden" n-tuple for further analysis. In order to aid in the analysis of data, the golden code also records a number of counters which record the number of events that fail each stage of the golden code, the total livetime of the data reduced and, finally, the number of events which pass the golden code. This is useful for the definition of detector efficiency.

4.3.1 First Pass Cuts

The initial stage of the golden code is concerned with the selection of primary and secondary signals. The first stage of the pulse selection is to look for an S2 signal. This is achieved by stipulating certain parameters by which a pulse may be considered of the right size and form to be an S2. The selection criteria at this point, depend on the mean charge arrival time of the pulse (defined as *s2tau*, the time required to reach 50% of the rise time of the pulse) which must be >150 ns as determined by the width of the gas gap and the strength of the electric field. Also required is that the pulse be greater than a specific integrated area (*s2area* of 0.5 V.ns) This cut specifically prevents the identification of single electron emission from the liquid surface as an S2 pulse (as such emission would satisfy the *s2tau* requirement for an S2 signal).

There is also a process of optical feedback that may occur in the detector whereby electrons are emitted from the cathode grid. As the cathode grid is found at the bottom of the drift volume, the drift time of the electrons will be the full potential drift time of ZEPLIN–III. For this reason, signals near to the maximum drift time (>15 μs) are ignored in pulse selection.

Finally, cuts are also applied to remove amplifier overshoot and the production of small S2 like signals in the tail of a true S2 signal. Amplifier overshoot is what happens when saturation occurs in one of the amplifiers and manifests itself as a low amplitude, wide signal directly after the S2 pulse. This is removed by looking in the low sensitivity (non-amplified) channel to see if the same pulse is visible. If it isn't, it is not counted as an S2 signal. An example of amplifier overshoot can be seen in Figure 4.1.

Large S2 signals are often associated with multiple feedback and single electron signals which may be clustered together by ZE3RA and given parameters sufficient for a naive analysis to identify them as true S2 signals. The golden code will ignore these candidate events if their area is less than 5% of that of the main S2 signal in the waveform. An example of such an event is shown in Figure 4.2.

Once an S2 signal has been identified, the golden code begins the process of identifying S1 pulses. This is a more difficult process for several reasons. Firstly, the noisy tails in S2 signals (which happen quite frequently) can give rise to multiple S1-like signals which occurs only if the signals are not clustered.

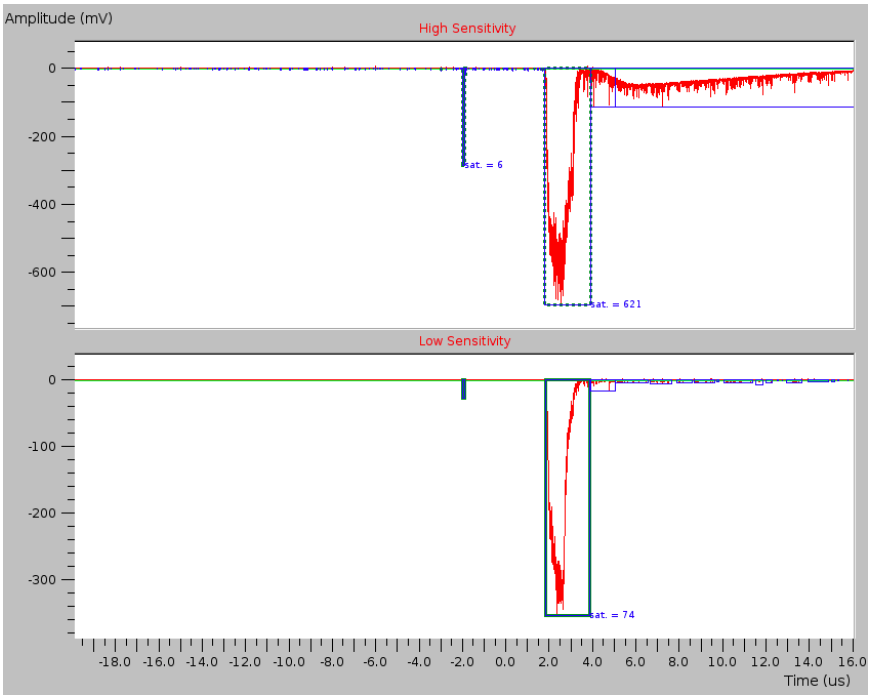

Figure 4.1: A typical example of amplifier overshoot in ZEPLIN-III. A large signal is seen following the ionisation pulse in the HS channel but not in the LS channel.

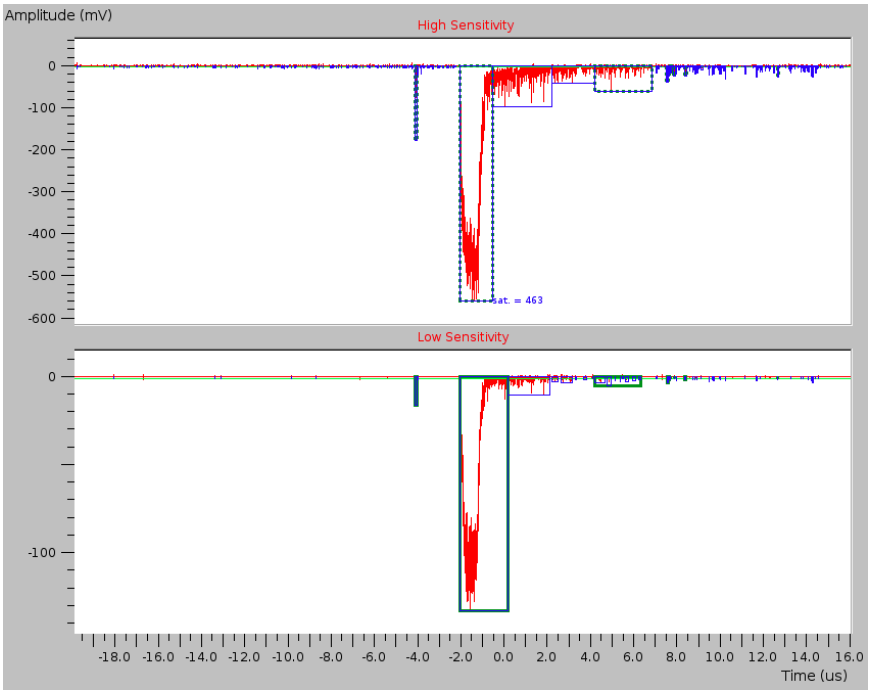

Figure 4.2: A typical example of where ZE3RA has clustered noisy pulse tail events together to give an S2 like signal.

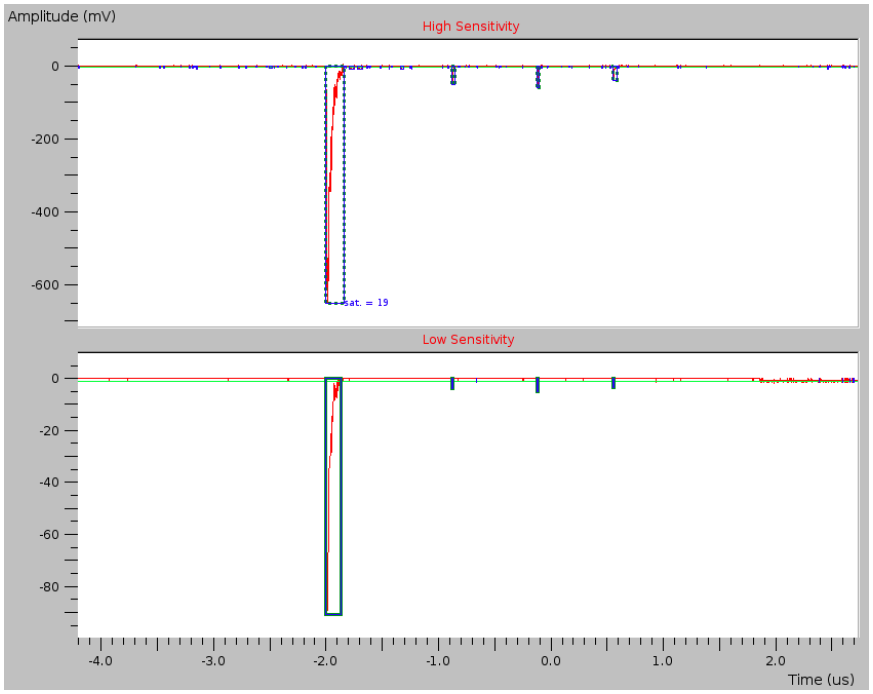

Figure 4.3: A typical example of PMT afterpulsing in ZEPLIN–III. The afterpulsing can be seen as small pulses subsequent to an S1 signal.

To remove the risk of these being identified as S1 signals, a rule is introduced whereby candidate S1 signals must precede the identified S2 pulse.

In addition to the pulses seen after S2 pulses, signals are also seen after S1 signals. These signals are associated with PMT afterpulsing where electrons are emitted from somewhere along the dynode chain. The electrons are accelerated along the remainder of the dynode chain giving rise to an above 3 RMS threshold, S1 like, signal. In order to identify PMT afterpulsing, a survey is taken of all pulses subsequent to an identified S1 signal. If the pulses occur within a specific time-window, and if each signal is associated with a single PMT, then the pulse is not counted as an S1 signal. An example of PMT afterpulsing can be seen in Figure 4.3. Afterpulsing windows are individual to PMTs and, as such, these pulses are easily identified. In addition, the afterpulse candidate must be smaller than the main S1 candidate pulse.

These pulse characterisation criteria are applied to all waveforms in a dataset and the number of S1 and S2 signals identified per waveform. In order to proceed

beyond this point, an event must have only a single S1 and a single S2.

4.3.2 Waveform and Voltage Cuts

Fluctuations in the baseline of the waveforms can cause incorrect parameterisation of pulses in data. To avoid this occurring, the baseline and noise on each channel must fall between fixed values. In addition to this, consistency is required between HS and LS channels in both timing and pulse area.

Voltage cuts are also applied to the data. Any fluctuations in PMT, anode or cathode voltage or current will cause the expected pulse parameters to change. In order to avoid any problems due to this, cuts were placed on all voltage and current values.

4.3.3 Golden Parameters

Following the application of the basic cuts, other pulse parameters are calculated. All parameters are calculated using the HS channel unless saturation is observed, in which case the LS channel is used. The golden variables are listed below.

- File and event information: *fil, dat, day, daqtime, lastime* - File number, date, day of run (since 27/02/08), time since first golden event and time since last ZE3RA event. The *lastime* measurement is used to calculate trigger rates and detector dead-time, even when events have been discarded by the golden code.

- Pulse ID: *s1id, s2id* - Pulse index (1-10) for S1 and S2 pulse.

- Pulse Saturation: *s1sat, s2sat* - Flag for saturation of S1 and S2 signals in the LS channel

- Pulse timing information: *s1time, s2time, dtime* - The time (within the individual timeline) of S1 and S2 pulses and the difference between the two, equivalent to the event drift time.

- Pulse Width: *s1tau, s2tau, s1wid50, s2wid50* - The mean charge arrival time for S1 and S2 pulses and the pulse width at 50% of the amplitude.

- Pulse Size: *s1amp, s2amp, s1area, s2area, s1areape, s2areape* - Amplitudes of S2 and S1 and their corrected areas in both V.ns and number of photoelectrons.

- Pulse Energy: *energy, estar* - Energy measured by S1 area and by a combination of S1+S2 energy.

- Peak PMTs: *s1peak, s2peak* - PMT which sees the most light for S1 and S2 signals.

- S1 Fold: *s1fold* - The n-fold coincidence level of the S1 signal.

- Corrections: *corlife, corpress, corfield, cordiff, corlevel, cortilt* - Correction factors due to electron lifetime, pressure, electric field, charge diffusion, liquid level and detector tilt.

The pulse areas are corrected for light-collection non-uniformity and variations in gain between PMTs through application of a flat-fielding correction. For *s2area*, corrections are applied for the mean electron lifetime (purity), detector tilt, field variations and liquid level variations. These corrections are only applied to S2 pulses as only the electroluminescence signal is dependent on these parameters. The *cordiff* correction is applied to *s2tau* to flatten the relationship between mean charge collection time and drift time. The relationship is not uniform due to the diffusion of the electron cloud through the liquid. This effect is more prominent for events with longer drift times. Figure 4.4 shows this correction which allows for easier data cuts at a later stage.

4.3.4 Second Pass Cuts

Once values for *s1tau, s2tau* and *dtime* have been calculated, it is possible to apply tighter cuts to ensure that all pulses are correctly assigned as an S1 or an S2. A new cut is applied on the mean charge arrival time of S1 (5< *s1tau* <40 ns) and S2 (400< *s2tau* <1100 ns). These cuts are defined by analysis of the mean charge arrival distribution in both cases and ensure that only pulses close to the mean of the distribution are included.

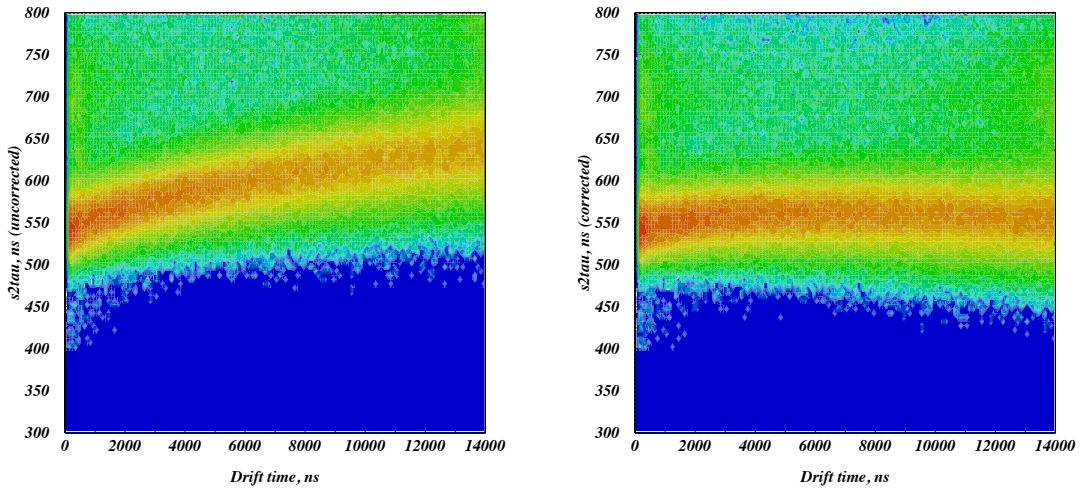

Figure 4.4: The uncorrected (left) and corrected (right) *s2tau* parameter. Flattening this distribution allows for easier data cuts at a later analysis stage

In addition to this, a cut is placed on the electron drift time to remove events with unphysical values. The cut is set (100< *dtime* <18000 ns) to remove events at the liquid surface or below the level of the cathode grid.

4.3.5 Position Reconstruction

The importance of a precise position reconstruction in ZEPLIN–III data cannot be overstated. The position reconstruction is vital for the calculation of light collection non-uniformities in the PMT array and the removal of events that have one S2 signal associated with two (or more) coincident S1 signals (see Section 4.5) which are more likely to appear with long drift times or larger radii.

In ZEPLIN–III, the vertical (z) position of an event is determined by the drift time. As mentioned previously, the full drift time of ZEPLIN-III is 14 μs (the drift time position of the cathode grid) and this equates to a maximum drift distance of 36 mm. The drift time distribution for the FSR data, displayed in Figure 4.5, is relatively uniform with peaks at both low and high values of drift time. This is expected as the PMT array and liquid surface scatters both dominate background in ZEPLIN-III for the FSR. The former provides background radiation and the latter has no Xe self-shielding for external radiation.

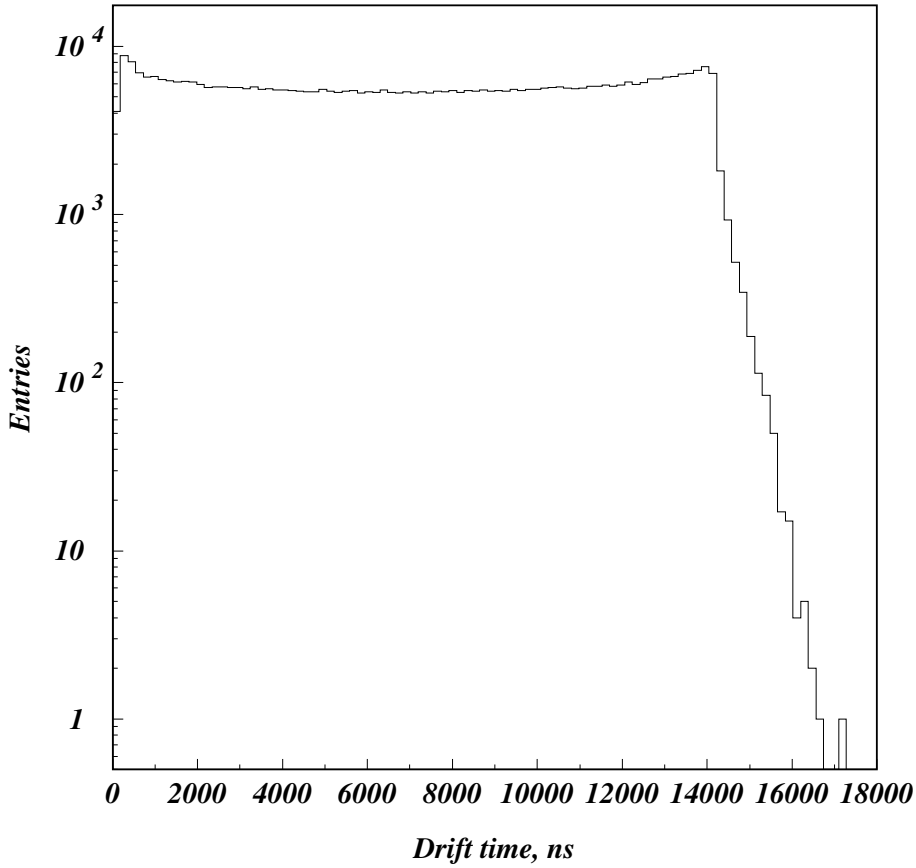

Figure 4.5: The drift time distribution for ZEPLIN-III FSR data. The increased rates at low and high values of drift time are from increased rates due to liquid surface scatters and PMT γ-rays respectively.

Several complementary methods are used to determine the position of an event in the horizontal (xy) plane. Firstly, a simple (yet sufficiently accurate for detector commissioning tests) centroid algorithm was developed to give an xy position calculated by using the PMT location weighted by the signal observed in the PMT (corrected for light collection variability). To improve upon this, a second method was developed that included a correction based upon simulations of expected signal output versus position in the xy plane. This improved confidence in event position out to greater radii than could be expected from a centroid position

calculation alone.

In order to utilise the full radial capability of ZEPLIN–III, a more advanced position reconstruction algorithm (known as *Mercury* reconstruction) was developed. This method utilises the spatial response profile of the PMTs for position reconstruction. A 2D Gaussian response function describes the signal expected in each PMT dependent on the *xy* position of the event. A least squares minimisation is performed in order to match the PMT hit pattern of an event to the response curve of all PMTs simultaneously in order to extract a location and energy for each event. This minimisation is performed on S1 and S2 pulses yielding the energy and *xy* position of each event. The precision of the reconstruction has problems with very small S1 events where few photoelectrons are detected.

The method relies on the fact that each PMT has the same position-dependent response as all others. An example factor that would prevent this from being true would be if the outer PMTs were close to the walls of the detector. Here, reflections from the walls would distort the response function leading to poor position reconstruction. The Gaussian response of PMT1 is shown in Figure 4.6. This, along with the responses of all 31 PMTs, is used to predict the true position and energy of any given pulse in a ZEPLIN–III timeline. The Mercury reconstruction provides a new set of parameters which are used for post reduction analysis of the ZEPLIN–III data. These parameters are described in Table 4.2. Figure 4.7 describes the coordinates used to describe the geometry of the target volume.

4.4 Detector Stability

Many parameters are monitored in order to inform adjustment of the experimental setup or correction of the data. Most of the parameters are monitored by the detector slow control systems (temperature, pressure and liquid level) but there are other parameters that may only be monitored using the data itself.

An automatic routine was developed to analyse daily ^{57}Co data (used to monitor the stability of the Xe purity, detector tilt and energy calibration) and the daily, unblinded 10% of the science run data. 90% of the data remained blinded in order that the final analysis would not be biased by knowledge of any event distributions in the WIMP search region. This was developed to allow a repro-

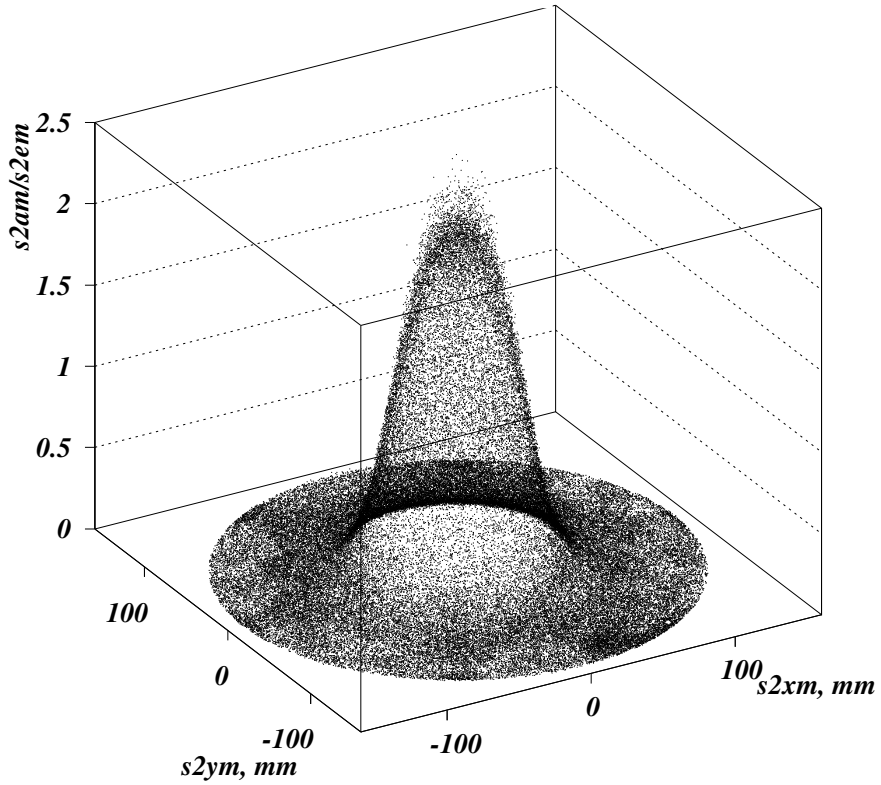

Figure 4.6: The bell-shaped response function of PMT1. The z-axis shows the signal intensity while the x and y axes define event position in the xy plane. The parameters $s2ym$ and $s2xm$ describe the pulse position in the xy-plane while $s2am$ and $s2em$ describe the integrated area of the pulse and the reconstructed energy respectively. These parameters are described in Table 4.2.

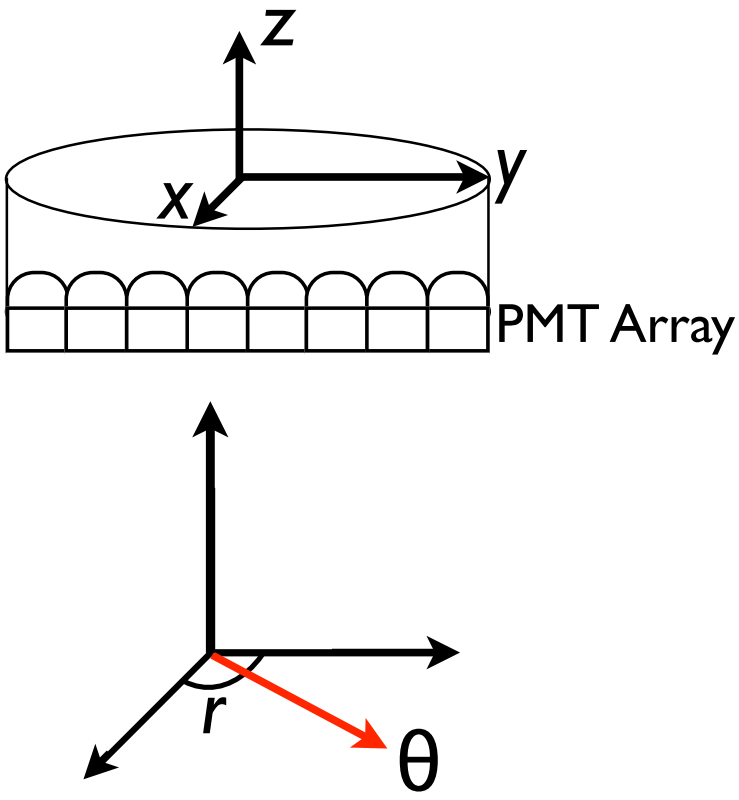

Figure 4.7: The coordinates used in ZEPLIN–III analysis. The upper graphic shows the xy and z (defined using the *dtime* parameter) coordinates and the lower diagram shows the r and θ parameters which describe the target volume geometry in terms of cylindrical polar coordinates. These parameters are described in Table 4.2.

Parameters	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 or S2
s1rmsm, s2rmsm	The RMS of the spatial hit pattern for S1 and S2
s1em, s2em	Event energy calculated from S1 and S2
s1e1m	S1 energy calculated using S2 position
s1chim, s2chim	Measures the χ^2 goodness-of-fit from the least squares minimisation
s12dism	The distance between the event locations measured by S1 and S2

Table 4.2: Additional parameters calculated by the mercury algorithm. The co-ordinate system used is defined in Figure 4.7.

ducible analysis that did not require any user input, thus avoiding any user error or variability in input parameters. The routine used in the FSR is described below, with details, in the following subsections:

1) Process daily ^{57}Co data with no corrections.

- Xenon purity routine - outputs electron lifetime measurements.
- Energy calibration - outputs S1 and S2 energy calibrations in keV/nVs

2) Reprocess daily ^{57}Co data with parameters calculated from 1).

- Measure tilt of the detector - outputs magnitude and phase of detector tilt.
- Channel amplification consistency check.
- Light yield calibration.

3) Process 10% of daily background data - Using S2 and S1 energy calibration and electron lifetime measurements calculated in 1) & 2).

- DRU routine - produces a daily differential energy spectrum.

4.4.1 Xenon Purity Measurements

The purity of the Xe volume is important in the maximisation of available fiducial volume. Trapping of electrons by electronegative contaminants as they drift through the liquid phase of the detector is the major cause of variability in S2 signal size as a function of depth. In order to monitor and correct for this, a measurement of S2/S1 as a function of drift time can be made. Since the S1 signal is (largely) unaffected, the ratio of S1 to S2 will vary with depth, giving a measure of the electron lifetime. An exponential fit is applied to the resultant measurement.

The event rate as a function of depth falls exponentially due to the low energy nature of the γ-rays associated with the β decay of ^{57}Co. Due to this, it is necessary to boost slice statistics at longer drift times by increasing the size of the energy interval used for slicing and fitting of the data. A typical electron lifetime measurement is shown in Figure 4.8 and the historical trend throughout the FSR is shown in Figure 4.9. ZEPLIN–III is designed with an all metal target chamber in order to maintain the purity of the Xe (organic compounds can be outgassed from materials such as PTFE).

Before the FSR, an electron lifetime of $\sim$20 μs was achieved and this improved to around 40 μs at the end of the run. This was due to the application of a constant electric field in the detector whilst the Xe was not recirculated, leaving the system isolated. With no new electronegative contaminants added, the negatively charged ions produced through electronegative attachment drift away from the fiducial volume.

4.4.2 Energy Calibration

The stability of the energy calibration was monitored using the daily ^{57}Co. A Gaussian fit was applied to the 122/136 keV peaks (associated with the β decay of ^{57}Co to ^{57}Ni) which fitted the top of the photo-peak for both S1 (top 50%) and S2 (top 30%) signals.

This peak distribution was compared to the calibration performed on the first day of the FSR and a shift was calculated. A typical calibration histogram is show in Figure 4.10 and the behaviour with time of both S2 and S1 energy calibrations is shown in Figure 4.11. These are stable to within a few percent.

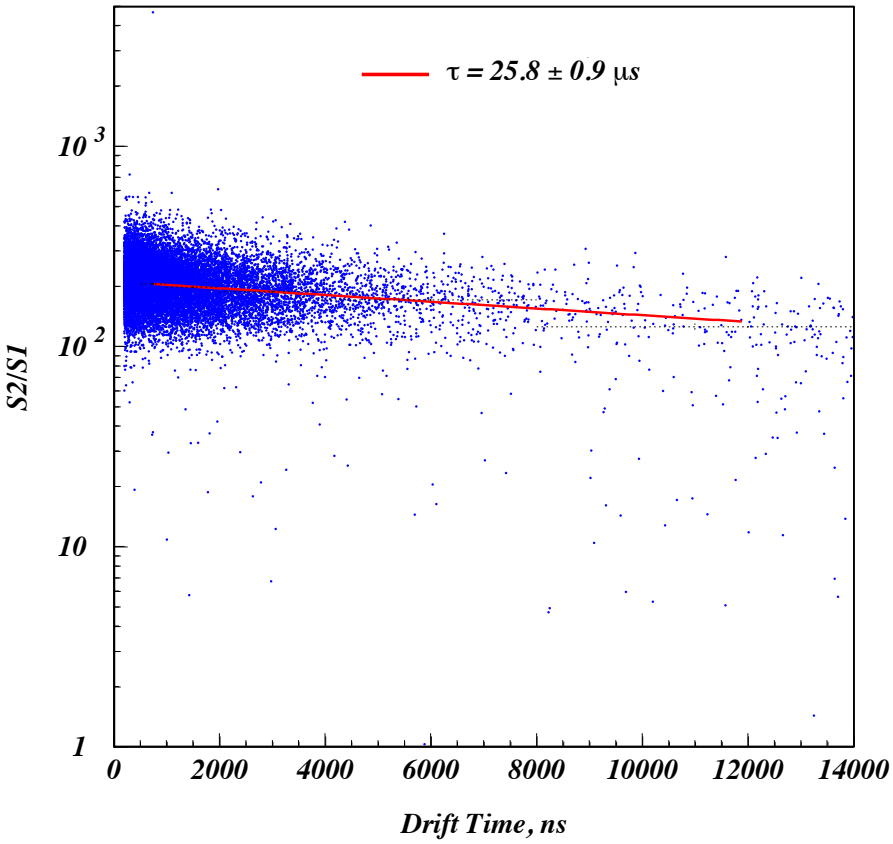

Figure 4.8: A typical daily measurement of the S2/S1 ratio which yields an electron lifetime fit. These data are acquired using a ^{57}Co run.

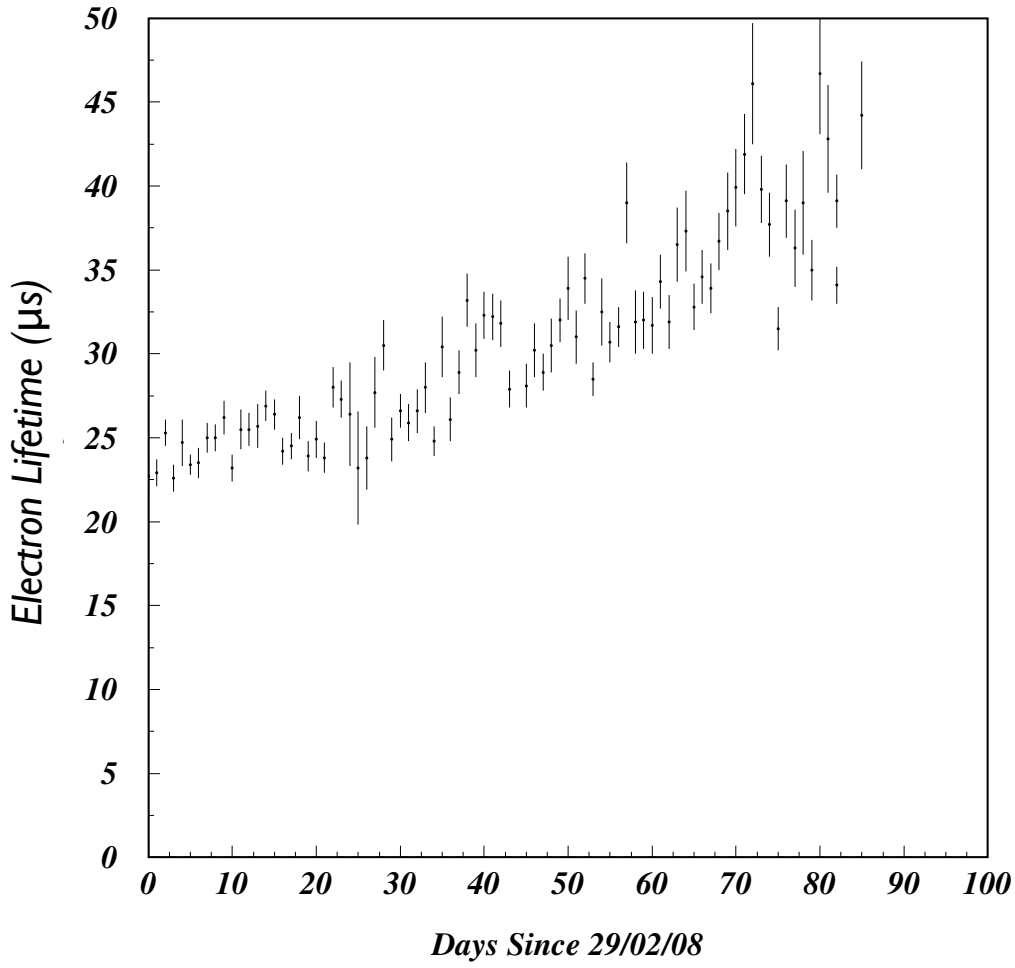

Figure 4.9: The historical trend of electron lifetime throughout the duration of the FSR. An electron lifetime of ∼25 μs equates to a correction of ∼1.82 between the top and bottom of the detector.

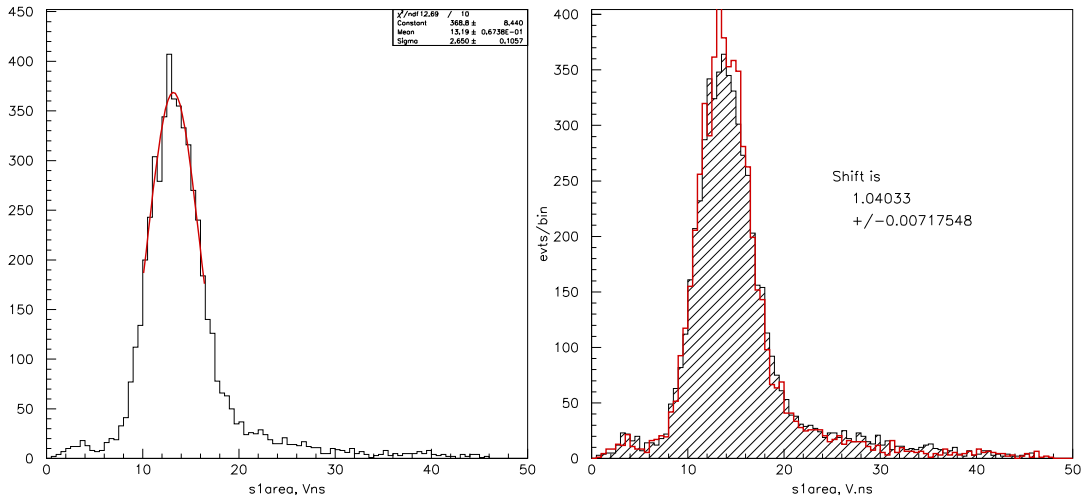

Figure 4.10: Fitting of the S1 spectrum from a ^{57}Co energy calibration. The left hand plot shows the fit to the top 50% of the data and the right hand plot shows the histogram in comparison to the reference data.

The measurement of the 122/136 keV peak allows the energy scale of electron recoils to be set, converting from the DAq scale of nV.s (voltage*ns) to keV. It should be noted that this only sets the scale for electron recoil energies. Nuclear recoils are quenched, so a second energy scale is defined. The process of setting this energy scale is described in Section 4.10.

In addition to measuring the stability of the S1 and S2 energy calibrations, the ^{57}Co data can be used to calculate the energy resolution of ZEPLIN–III. In order to calculate the energy resolution, a Monte-Carlo simulation is performed for a ^{57}Co calibration run. A smearing function is applied to the simulated data until it matches the experimental data. The width of this smearing function is used to define the resolution. Figure 4.12 shows the energy resolution measured in the FSR. The energy resolution is described by the following equation:

$$\frac{\sigma}{E} = \frac{0.6}{\sqrt{E}} \tag{4.1}$$

where σ/E is the energy resolution. For the 122.1 keV peak in ^{57}Co, this is calculated as 5.4%.

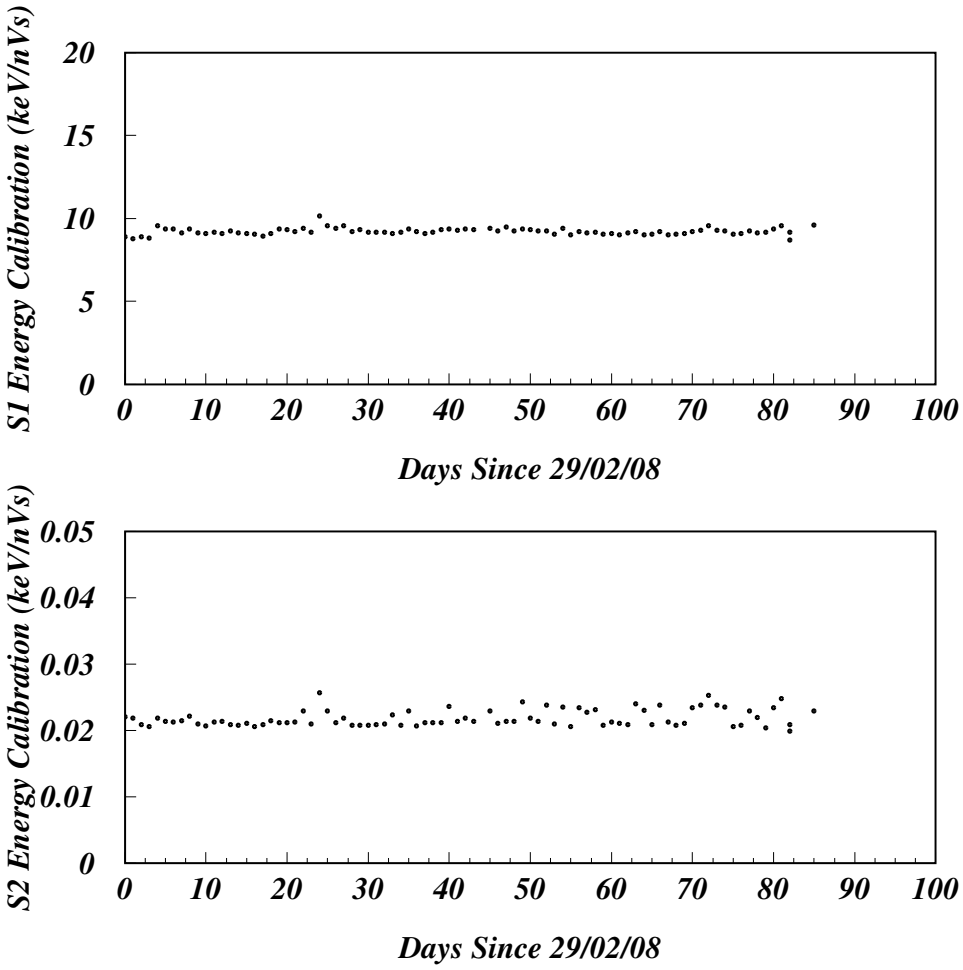

Figure 4.11: The historical trend of the S1 and S2 energy calibrations throughout the FSR.

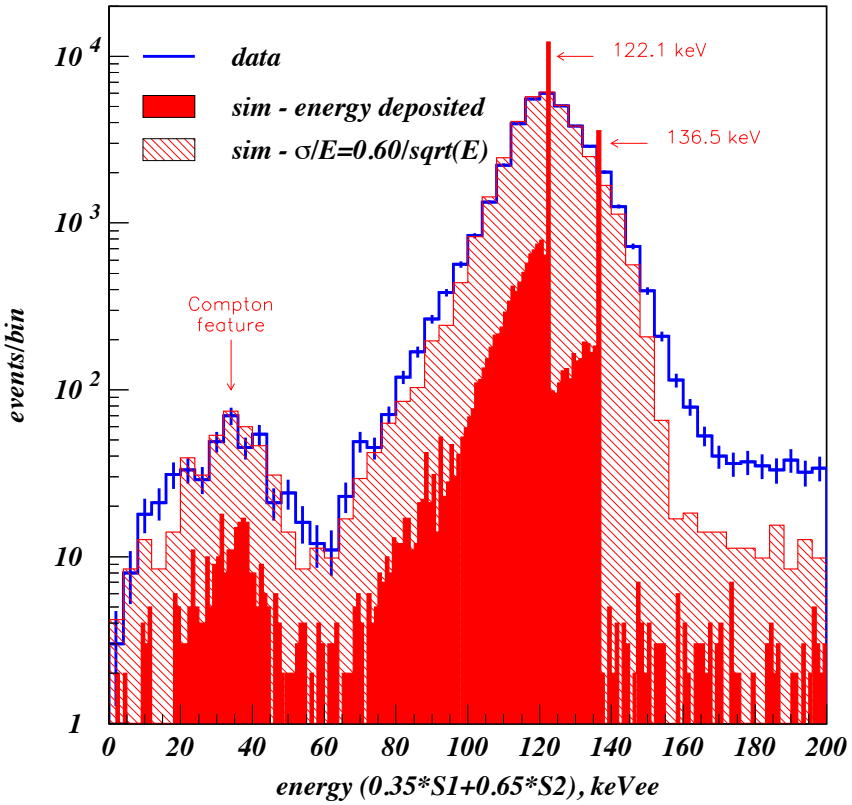

Figure 4.12: Detector response to an external ^{57}Co source. The blue line shows the data taken from a 1-day run, the dark red region is the simulated data with characteristic peaks and Compton feature showing and the lighter red region is the simulation with a smearing function added to match the data. The resolution is shown on the plot [83].

4.4.3 Detector Tilt and Liquid Level

The Boulby facility is located in a seam of rock salt with potash seams above and below. Both rock salt and potash are soft rocks so the laboratory is subject to geological movement. This being the case, it is important to monitor the tilt of ZEPLIN-III in order to correct for the non-uniformity of the gas gap thickness.

The width of S2 pulses is determined by the width of the gas phase in a 2-phase detector (with amplitude determined by the number of electrons accelerated through the gas phase). Any variation in the liquid level of the system will lead to incomparable data if it is not corrected. The liquid level of the detector is monitored using three separate copper rods, each of which has its capacitance measured. As the liquid level increases or decreases, the capacitance also rises and falls linearly.

Also of interest is any variation in gas gap across the surface of the detector which would be associated with a non-uniform S2 width response across the surface of the liquid. The variation in tilt is measured by plotting the width of the S2 pulses at half maximum (*s2wid50*) as a function of the angular position (*s2theta*). To this, a sinusoidal function is fitted and the amplitude and phase recorded.

At the start of the run, the detector tilt measured ∼1 mrad and at the end ∼2 mrad. The detector level was not manually corrected through the FSR as the tilting rate was relatively low and the run was fairly short. Figure 4.13 shows a typical plot of *s2wid50* as a function of *s2theta*. Figure 4.14 shows the historical trend of the detector tilt over the duration of the FSR.

4.4.4 Channel Amplifier Stability Check

The outputs from the 31 PMTs are recorded and analysed assuming that there is a constant 10x difference between low and high sensitivity channels. For average sized pulses, the relationship should be linear. Sensitivity to the small signals in the LS channel is reduced and large signals cause amplifier saturation in the HS channel. These factors cause deviation from the linear relationship but it is possible to check the amplifier stability using only the central region of the plot.

During the FSR, the amplifier stability was checked daily and no deviation from a linear relationship between signal size in the HS and scaled LS channels

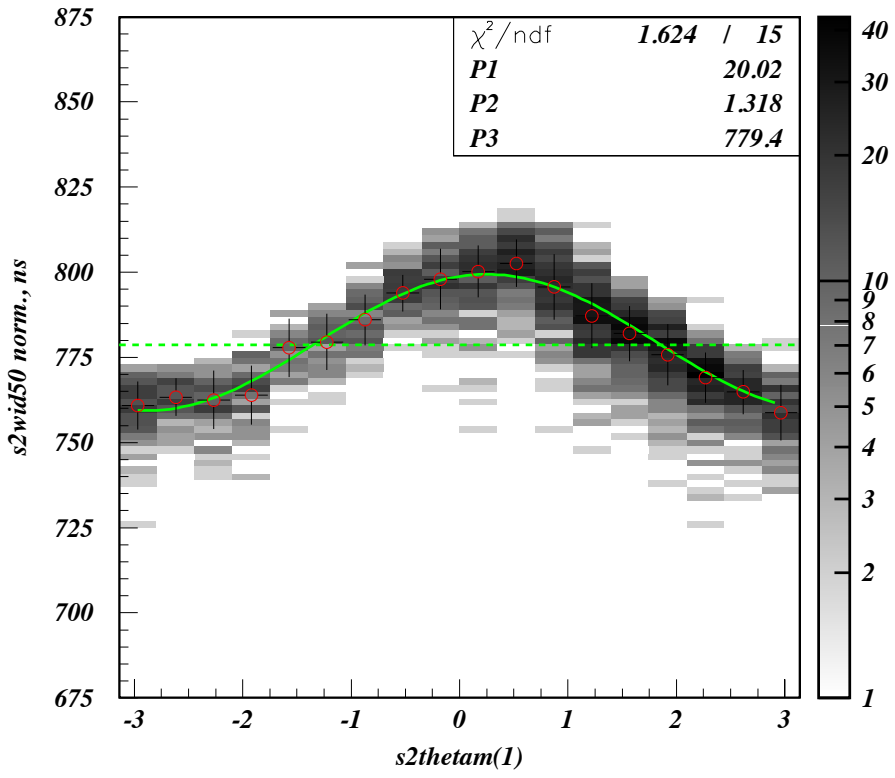

Figure 4.13: A typical levelling plot for ZEPLIN-III during the FSR. A sinusoidal relationship can be observed between *s2wid50* and *s2theta*.

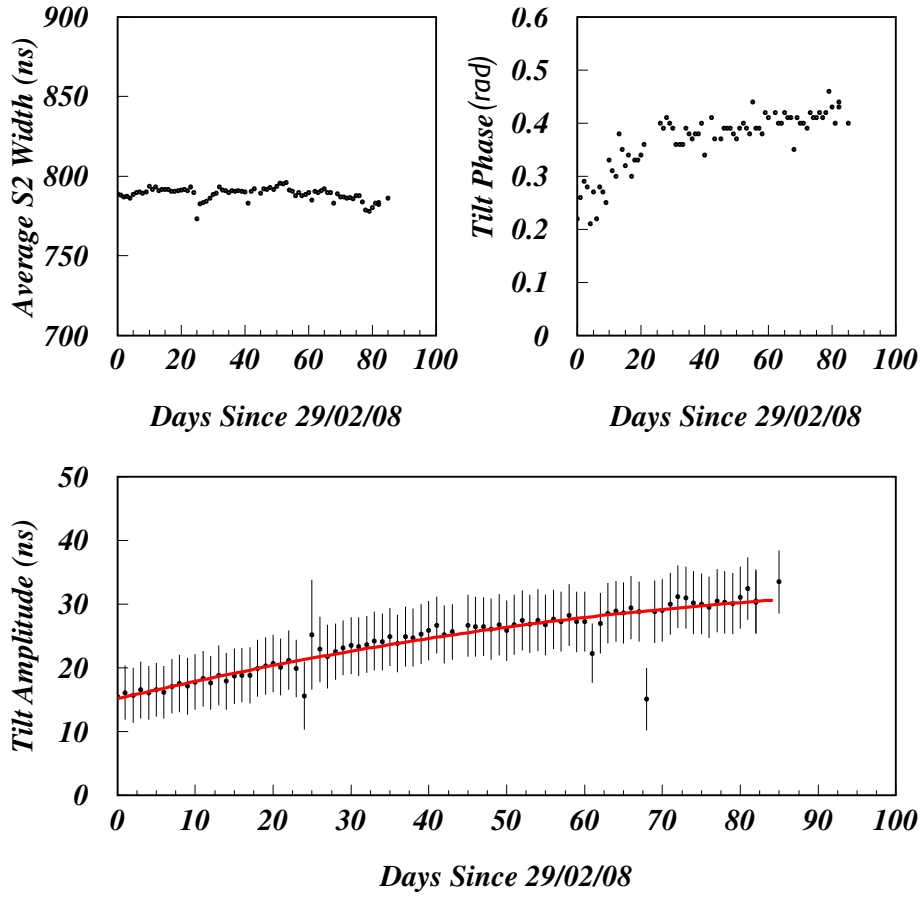

Figure 4.14: The tilt evolution during the FSR.

was observed.

4.5 Multiple Scintillation Single Ionisation Events

Perhaps the most challenging aspect of ZEPLIN–III data analysis is the presence of events in which two (or more) scatters occur in so-called dead regions of xenon. These are regions where ionisation cannot be extracted to the liquid surface, such as regions where no electric field is applied, or a reverse field is applied (such as between the cathode grid and PMT grid). Regions in which Multiple Scintillation, Single Ionisation (MSSI) events are prevalent are highlighted in Figure 4.15. Both scintillation signals are detected but only one ionisation signals occurs.

These MSSI events appear, in data, to be consistent with the parameters required for a golden event. This is because the S1 signals are coincident in time, appearing to be a single, large S1 signal. The additional S1 signal helps to reduce the S2/S1 ratio, artificially pushing the event down in S2/S1 parameter space. In extreme cases, this leakage can contaminate the area of parameter space containing the nuclear recoil band, potentially giving rise to false signals in the WIMP search region.

The severity of the shift in S2/S1 compared to the true value of the individual S2:S1 pair is determined by the size of the second S1 signal in comparison to the initial S1 signal. If the second S1 signal is large, the event is more likely to have similar parameters to those expected of a nuclear recoil event (lower value of S2/S1).

During the FSR, the PMTs form the dominant source of background. This being the case, most of the MSSI events will first scatter in the region labelled A in figure 4.15. In this region, no electric field is applied (the PMTs and PMT grid are held at the same voltage), so electrons are not drifted in recoils that occur in this region. Both region B and region C contribute fewer MSSI events. Events from region B are due to the reverse field region between the cathode and the PMT grid and the events from region C are due to non-vertical field lines found towards the edge of the liquid region that do not allow ionisation electrons to be drifted to the liquid surface. The reverse field region, however, remains important to reduce the background from β-decay and photo-absorption signals from low energy γ-rays emitted by the PMTs.

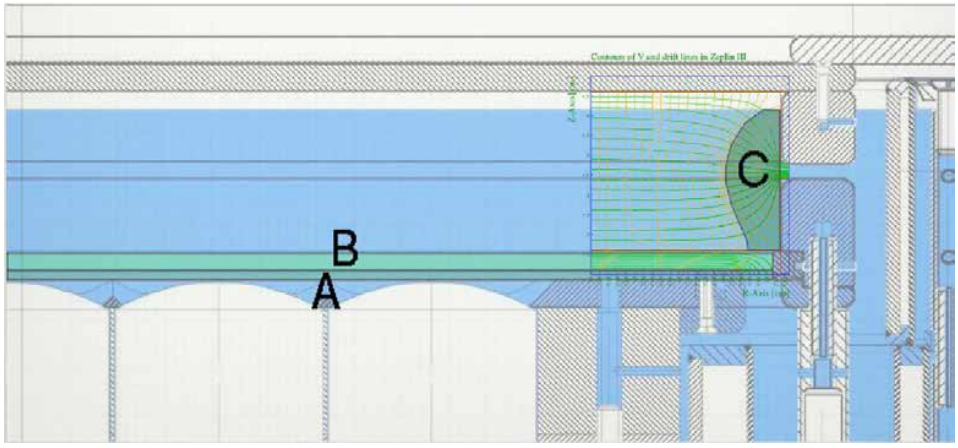

Figure 4.15: Diagram of the ZEPLIN–III target geometry highlighting the three regions in which ionisation signals may be lost. Region A is the region in between PMT faces where there is no electric field. Region B is a region of high reverse field between the cathode and PMT grid. Region C is where non-vertical field lines prevent ionisation signals from reaching the liquid surface.

4.5.1 Removing MSSI Events

Through study of S1 only events (coming, mainly, from the reverse field region), it was possible to define analysis parameters by which the MSSI events deviated from the expected event distributions. The two parameters used for MSSI removal where *s1rmsm* and *s2chim*.

The first parameter, *s1rmsm* gives the spatial RMS value for each reconstructed event. Figure 4.16 shows that a significant fraction of MSSI events from below the nuclear recoil median line have low RMS values in comparison to the main population of events. A cut was developed to remove these events without removing a significant number of events from the AmBe data run used to determine WIMP detection efficiency (see Section 4.7).

The second parameter used to remove MSSI events was the χ^2 goodness-of-fit parameter from the position reconstruction algorithm. Events with multiple vertices due to scattering, should have scintillation hit patterns across the PMTs which do not agree with those expected for single scatter events. Due to this, MSSI events can be cut for events with a χ^2 value above those seen in the main population. This was a very powerful cut which removed most of the MSSI events

with only a small efficiency penalty in the WIMP search region.

4.6 Fiducialisation

The full active mass of liquid Xe in ZEPLIN–III is ∼12 kg. By design, it is not possible to utilise the whole volume due to the dead regions described in section 4.5. These dead regions exist to remove low energy background radiation emitted by the PMT array. To remove events from these regions, cuts are applied to define a fiducial volume in ZEPLIN-III. The parameters used to define these cuts are, in the xy-plane, $s1rhom$ and $s2rhom$ and, in the z direction, $dtime$.

The maximum allowable drift time was found by producing a plot of S2/S1 as a function of drift time. Close to the cathode grid, the electric field increases and becomes non-uniform. This leads to an anomalous S2/S1 response in this region and defines the maximum drift time. Figure 4.17 shows S2/S1 as a function of drift time and the position of the cathode inferred from this.

In the xy-plane, a cut is applied such that only events within a radius of 150mm are accepted. This allows self-shielding from low radiation emitted by the copper displacement rings found at the edge of the detector.

4.7 AmBe Calibration

The detector must be calibrated to expected WIMP signatures, namely, elastic nuclear recoils. To this end, exposure to an americium-beryllium (AmBe) source is performed. Neutron calibrations are essential for both the definition of the WIMP search region with the calculation of efficiencies, and the calculation of the neutron/γ-ray discrimination power. The AmBe calibration used to define the WIMP search region in the ZEPLIN–III FSR was carried out on 24^{th} and 25^{th} May 2008 and gave a combined run time of approximately 5 hours. The potential calibration time for an AmBe calibration is limited due to the neutron activation of metalwork in the ZEPLIN-III target volume. This activation causes the production of excited states of ^{131}Xe and ^{129}Xe which increase background levels in the detector (activation is discussed again in section 7.1.1). The AmBe run length of 5 hours is chosen so that the additional background has returned to the standard background level within a few days.

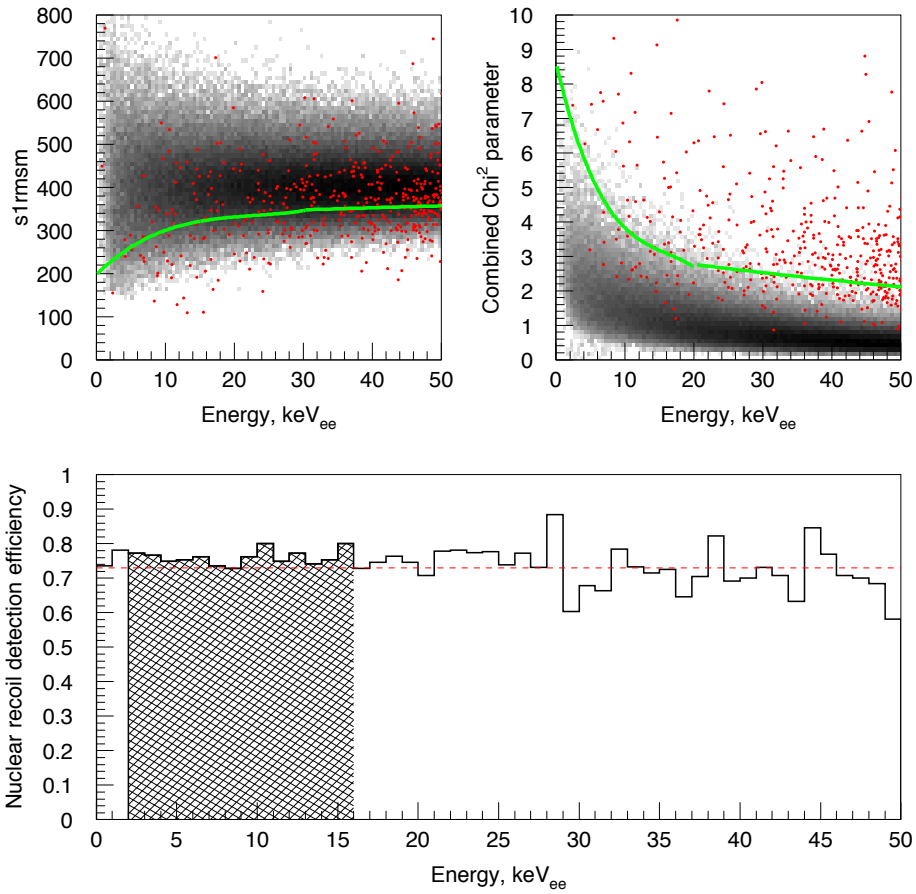

Figure 4.16: Cuts used to remove MSSI events in the FSR. Top left shows the $s1rmsm$ where events with values below the main population are cut. Top right shows the $s2chim$ cut where events with χ^2 values higher than expected are cut. The red points show events where the S2/S1 ratio is below the 50% median of the elastic nuclear recoil band, discussed in Section 4.7. The lower plot shows the energy dependent efficiency of these cuts with the region of interest (2-16 keV$_{ee}$) highlighted.

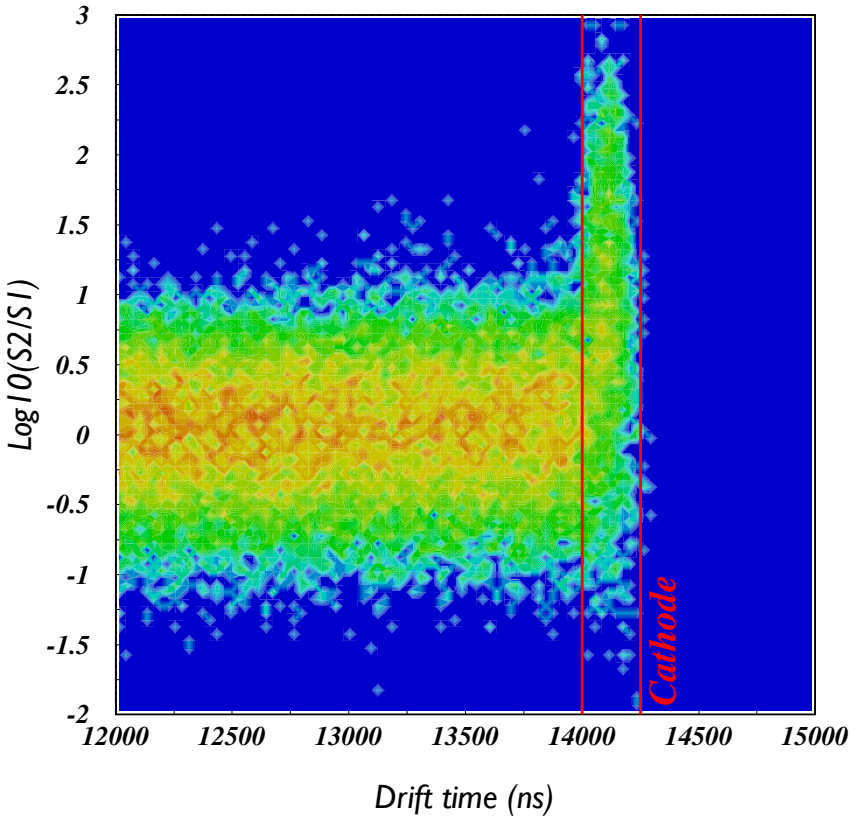

Figure 4.17: The definition of maximum allowable drift time in ZEPLIN–III. The position of the cathode is highlighted.

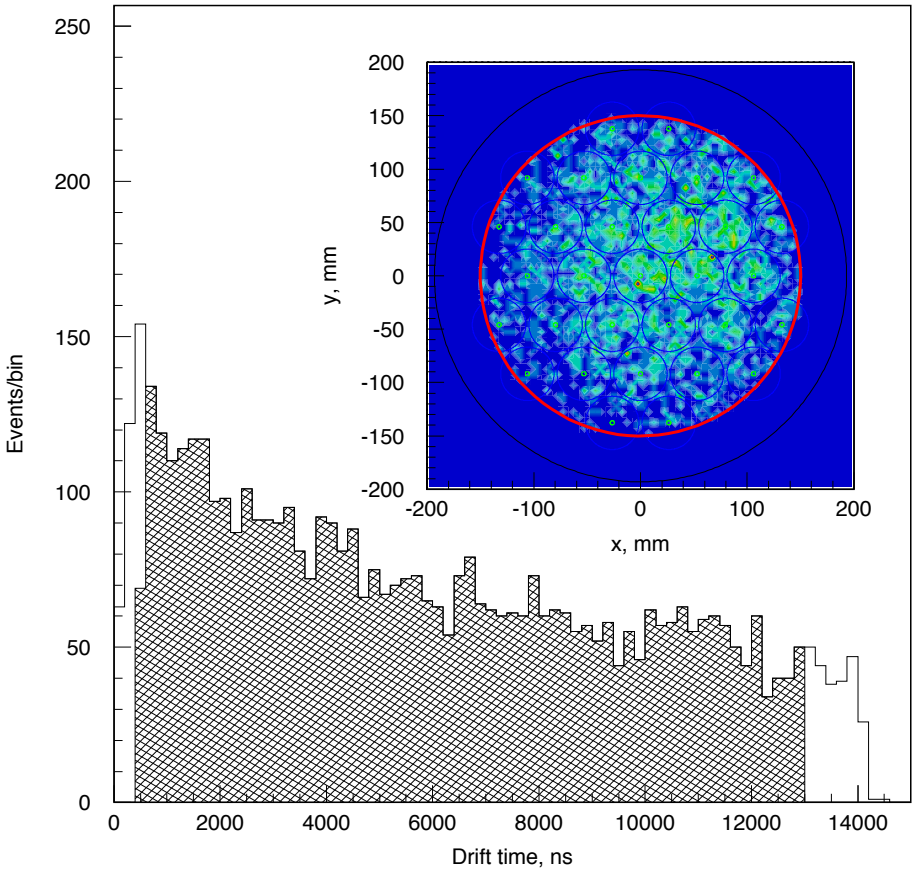

Figure 4.18: The drift time distribution for events in the AmBe calibration run (main plot) and the xy distribution (insert).

Figure 4.18 shows the spatial distribution of AmBe events for this run, the bias to the top right of the array occurs as the source was retracted through the polypropylene shielding to optimise trigger rate (giving acceptable dead-time levels) in the run. The drift time distribution is also consistent with events occurring towards the top of the array as is expected for a run taken with the AmBe source above the fiducial volume. The emission of neutrons occurs in coincidence with γ-rays that preferentially scatter in the upper region of the Xe volume (due to the short scattering length of γ-rays in liquid Xe) whereas the neutrons are distributed uniformly across all depths.

The nuclear recoil signal region is defined by selecting the elastic recoil pop-

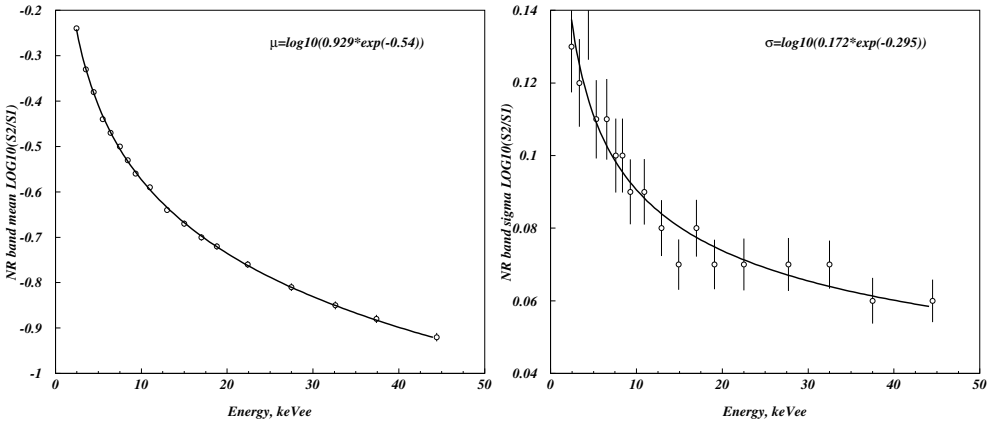

Figure 4.19: Functional fits to the means of each 1 kev_{ee} slice of the AmBe data (left) and to the standard deviation (right).

ulation of the AmBe data and slicing it into 1 keV_{ee} bins. A log-normal fit is applied to these slices and the means and widths of each recorded. For the FSR, the median of the nuclear recoil distribution is described by a power law of the form $\mu = \log_{10}\left((0.929 \pm 0.006)\, e^{-(0.540 \pm 0.003)}\right)$. The width of the distribution is described by $\sigma = (0.179 \pm 0.013)\, e^{-(0.295 \pm 0.028)}$.

Figure 4.20 shows the fitted parameters for μ and σ and Figure 4.20 shows the S2/S1 vs energy parameter space plot for the AmBe calibration run. The median of the nuclear recoil band (also referred to as the 50% line) is highlighted in red and $\pm 1\sigma$ bands are shown in blue.

4.8 ^{137}Cs Calibration

In order to improve the level of statistics and give a comparable representation of the expected electron recoil background, a ^{137}Cs run was taken over a period of 9 days. The source was placed within both the lead and polypropylene shielding and retracted so as to attain an optimal event rate.

The spatial distribution of the ^{137}Cs calibration is shown in Figure 4.21. The data shows a much more uniform distribution in the z direction which is consistent with the expected distribution due to the, relatively, higher energy γ-rays being

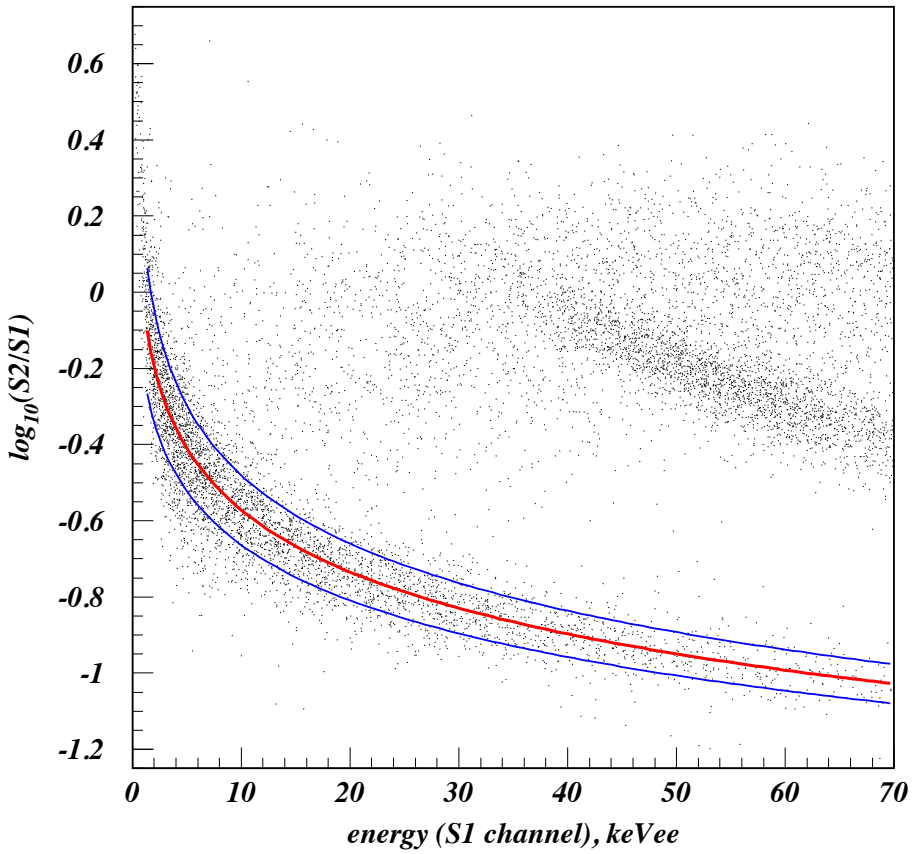

Figure 4.20: The S2/S1 vs energy plot for the FSR AmBe calibration data. The red line highlights the median of the nuclear recoil band and the blue lines highlight the $\pm 1\sigma$ bands. The second event population seen is due to inelastic neutron recoils where the Xe nucleus is excited. The nucleus de-excites with the emission of a characteristic 40 keV γ-ray.

able to probe greater depths of liquid (as compared to ^{57}Co). A comparison of the AmBe nuclear recoil data and ^{137}Cs data is shown in Figure 4.22. 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. The improvement in discrimination at higher fields is one of the most important results in the ZEPLIN–III experiment. The calculation of a numerical discrimination power is discussed in Section 4.11.

In addition to the calculation of discrimination power, ^{137}Cs is used to generate low energy electron recoils in the Xe volume. The higher energy of the γ-ray associated with the decay of ^{137}Cs (662 keV in comparison to 122/137 keV for ^{57}Co) means that low energy recoils can be generated through Compton scattering. This is not possible with a ^{57}Co source as the lower energy γ-rays are at an energy where photo-absorption dominates meaning very few low energy scatters will occur.

4.9 The 10% Science Data Sample and Defining a WIMP Search Region

In order to calculate event parameters, tune data cuts, and monitor stability in ZEPLIN–III, it was decided that 10% (every 10$^{\text{th}}$ file of the dataset) of the science data would be unblinded. In order to avoid biasing the analysis (potentially leading to removal or false identification of a WIMP signal) all science data is blinded. Opening (unblinding) 10% of this data to perform analysis is equivalent to a 10% efficiency loss in the final result. A comparison between the mean and the standard deviation of the ^{137}Cs calibration data and the 10% background data showed excellent agreement. However, the background data showed many fewer outliers at lower values of S2/S1 than the ^{137}Cs data (in effect, the background had proportionally fewer MSSI events). The developed MSSI cuts were more effective at rejecting events in the background data so it was decided that the background data would be used to characterise the electron recoil population. This difference occurs because ^{137}Cs is, effectively, a point source of radiation whereas the true background radiation is emitted from 4π around the detector.

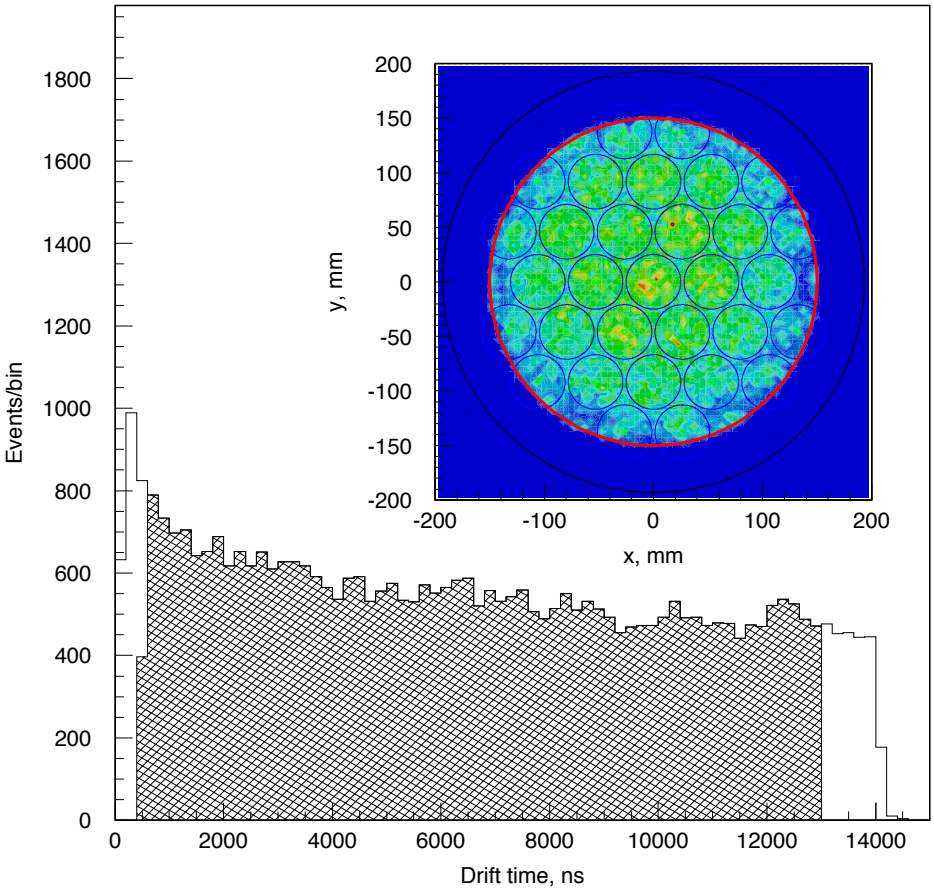

Figure 4.21: The drift time distribution for events in the ^{137}Cs calibration run (main plot) and the xy distribution (insert).

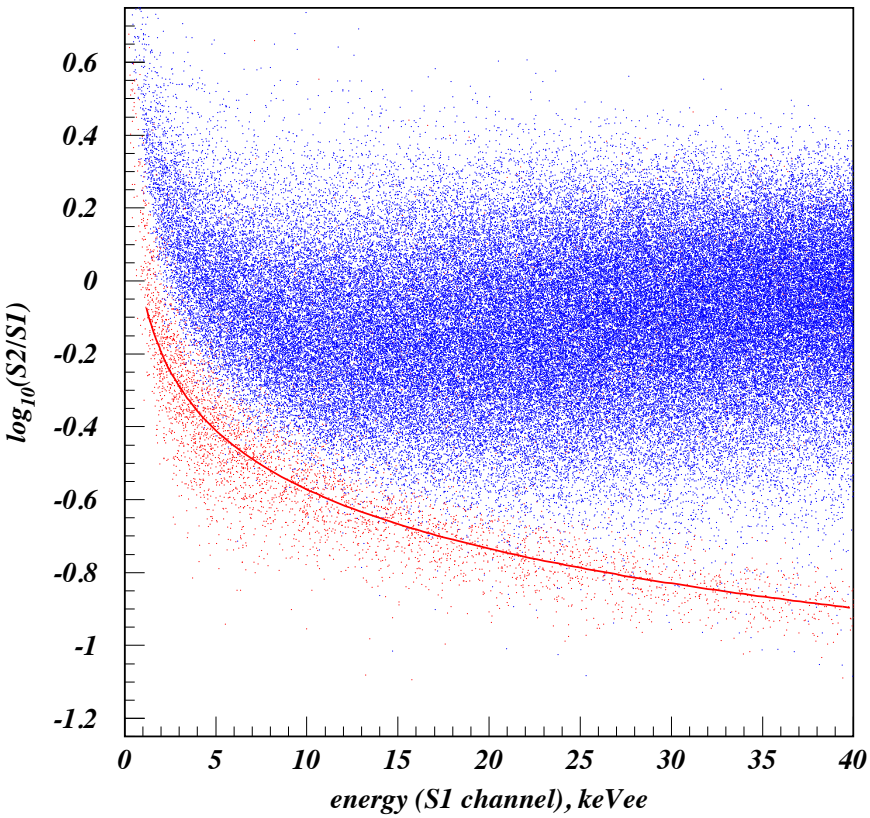

Figure 4.22: Scatter plot showing the nuclear recoil population of the AmBe calibration (red) and the electron recoil population from the ^{137}Cs calibration (blue). The 50% nuclear recoil line is shown in red for reference.

A second reason for the poor representation of the background is the fact that the ^{137}Cs source triggers ZEPLIN-III at a higher rate than background radiation. The difference is due to the increase in resistivity of the bialkali PMT photocathode at low temperatures [134, 135, 136]. An increase in resistivity enhances localised charging of the photocathode and this, in turn, attenuates and distorts electric fields of the input optics. This effect is countered somewhat by the installation of a set of metal tracks below the photocathode. These tracks decrease the resistivity of the photocathode but also increase the non-uniformity by creating regions where the ability of the PMT to neutralise the charge produced in the emission of electrons from the photocathode is changed. The increase in cathode non-uniformity coupled with the localised charging cause an overall decrease in the mean response of the PMT. This reduction in response is dependent on the rate, distribution and intensity of the radiation source. The disagreement between background and ^{137}Cs data is consistent with the increased rate of energy deposition in the liquid Xe volume. Figure 4.23 shows the excellent agreement between mean and standard deviation values between WIMP search data and ^{137}Cs. The differences described only becomes apparent when looking at the tails of the distribution.

By comparison of the 10% dataset and the AmBe nuclear recoil calibration, a WIMP search region was defined. It is necessary to define a region (the "box") that is sensitive to nuclear recoil events but that is free from background contamination. The box was defined with an upper limit in log(S2/S1) phase space set by the median of the nuclear recoil population, $\mu_{nr}(E)$, with a lower limit set at $\mu_{nr}(E) - 2\sigma_{nr}(E)$.

The lower and upper bounds in terms of energy, are set using the detector efficiencies and backgrounds. The lower limit is set at 2 keV$_{ee}$ which is the lower threshold of ZEPLIN–III. Below this, low threshold data shows a precipitous fall in detection efficiency. The upper limit in energy is set by the probability of MSSI leakage below the log(S2/S1) acceptance line which increases with increasing energy while the signal probability is falling exponentially. As a consequence of this, the upper bound is set at 16 keV$_{ee}$. Section 3.2 showed that the expected WIMP-recoil spectrum peaks at low energies and falls at an exponential rate. Due to this, above 16 keV$_{ee}$ there is little to gain in terms of additional signal while the electron recoil band gets wider. This box was defined prior to the unblinding

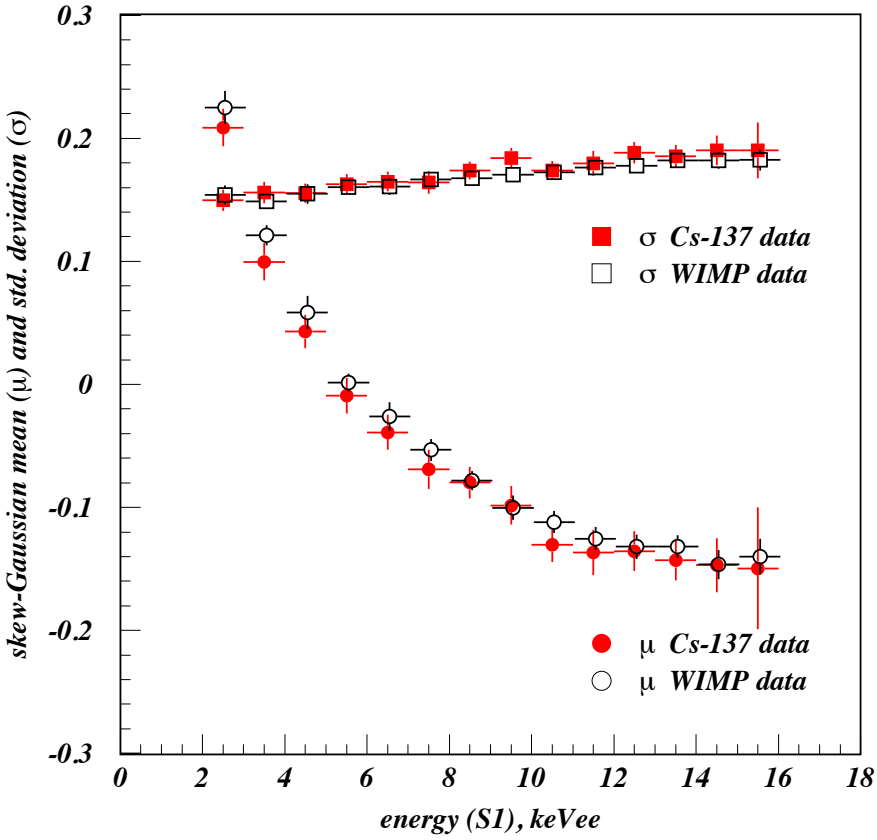

Figure 4.23: Comparison of mean and sigma of electron recoil populations from ^{137}Cs calibration (red) and WIMP search data (black).

of the full dataset and remained unchanged. The box includes 47.7% of nuclear recoil acceptance from 2-16 $\mathrm{keV_{ee}}$ before any cuts have been applied to the data.

4.10 WIMP Search Data

As discussed in Section 3.5.1, the differing energy deposition process of nuclear recoils and electron recoils results in a different scintillation response, requiring a conversion between energy scales. In previous experiments (ZEPLIN-II and XENON10) a flat energy conversion factor between $\mathrm{keV_{nr}}$ and $\mathrm{keV_{ee}}$ was applied.

The conversion in energy is calculated using the following formula:

$$E_{nr} = \frac{S_e}{L_{eff} S_n} E_{ee}, \tag{4.2}$$

where L_{eff} is the zero-field relative scintillation efficiency and S_e and S_n are the field suppression factors for electrons and neutrons. The energy-dependent nuclear recoil detection efficiency is calculated through comparison between experimental and and simulated AmBe spectra. At the operating fields of ZEPLIN-III, the field suppression factors are $S_e = 0.36$ and $S_n = 0.9$ [127].

Using the energy independent field suppression terms, a conversion of $E_{nr} = 2.09 E_{ee}$ can be calculated. Figure 4.25 shows the comparison between simulated (solid red curve) and experimental (blue shaded region) values. The simulated values are scaled using a conversion factor of 2.09. Figure 4.25 shows a clear mis-match between the simulated and experimental data. This mismatch could be described by one of two causes. The first would require a loss in detection efficiency at low energies that had not been accounted for. Each potential source of efficiency loss was checked and verified and it was found that this could not explain the mismatch. The efficiencies for each stage of the ZEPLIN–III FSR are listed in Table 4.3.

The second explanation for this is because the L_{eff} term has an energy dependence. The reason for the energy dependence of this term has already been discussed in Section 3.5.1. It is important to note that the theory behind the energy dependence of the relative scintillation efficiency was constructed after the effect had been seen in both ZEPLIN-III and XENON10. This is why the detection efficiency was investigated before an energy dependent L_{eff} was used.

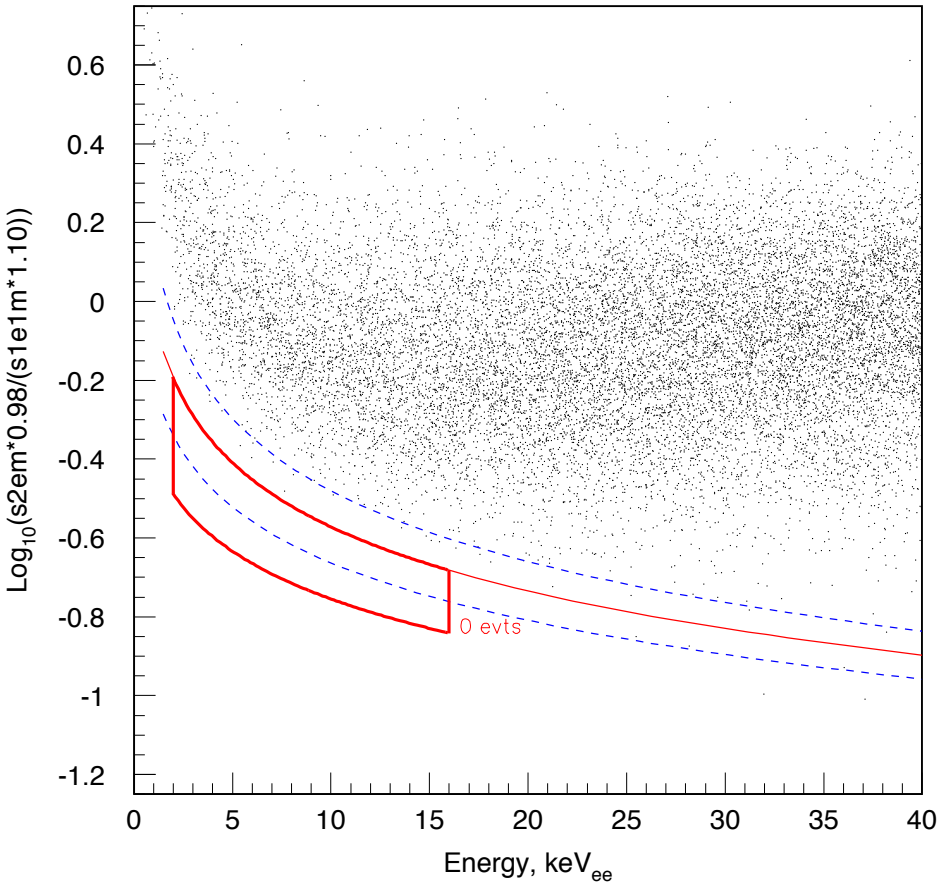

Figure 4.24: Discrimination scatter plot showing the 10% sample of science data from the ZEPLIN–III FSR. The WIMP search region is shown (red) along with $\pm 1\sigma$ (blue dashed) lines. With 10% of data, no events are seen in the WIMP search region. The extended red curve represents the median of the elastic nuclear recoil band of the AmBe data.

Energy Independent Efficiencies		
Effect	**Efficiency**	**Method**
Deadtime	91.7%	Measured
Hardware upper threshold	100%	On-off comparison
ZE3RA pulse finding (1)	96.0%	Visual inspection Hand calculation
Event reconstruction (2,3)	91.9%	Visual inspection
Selection Cuts (5)	73.0%	On-off compare
WIMP box acceptance	47.7%	Calculation
Energy Dependent Efficiencies		
Effect	**Threshold**	**Method**
Hardware (S2) trigger	S1 = 0.5 keV_{ee}	Two data-sets Visual inspection Modelling Pulser tests
Software S1 area	S1 < 1 keV_{ee}	Two data-sets Calculation Scatter plots
Software S1 3-fold	S1 = 1.7 keV_{ee}	Calculation 2-fold data-set analyses

Table 4.3: Calculated efficiencies for the ZEPLIN–III FSR.

Figure 4.26 shows Figure 4.25 replotted as a relative efficiency between simulation and data. It shows a distinct fall below 10 keV_{ee}. The functional fit can be converted into an L_{eff} which is used to give an energy dependent version of Equation 4.2. The energy dependent L_{eff} is shown in Figure 4.27.

The science run lasted for 83 days between 27^{th} February and 20^{th} May 2008 with an average 84% duty-cycle. This gave a raw exposure of 847 kg.days from a 12 kg active target which lead to a post fiducialisation exposure of 126.7 kg.days. Over this time, the detector remained stable, with all operational parameters monitored and corrections applied where needed.

A proportion of 90% of the data had been left blinded with the intention of performing a blind analysis but this was not possible for two reasons. The first of these was that a weakness was found in the Mercury reconstruction code after the data had already been processed once. Secondly, as previously discussed, the ^{137}Cs calibration did not accurately represent the tail of the electron recoil

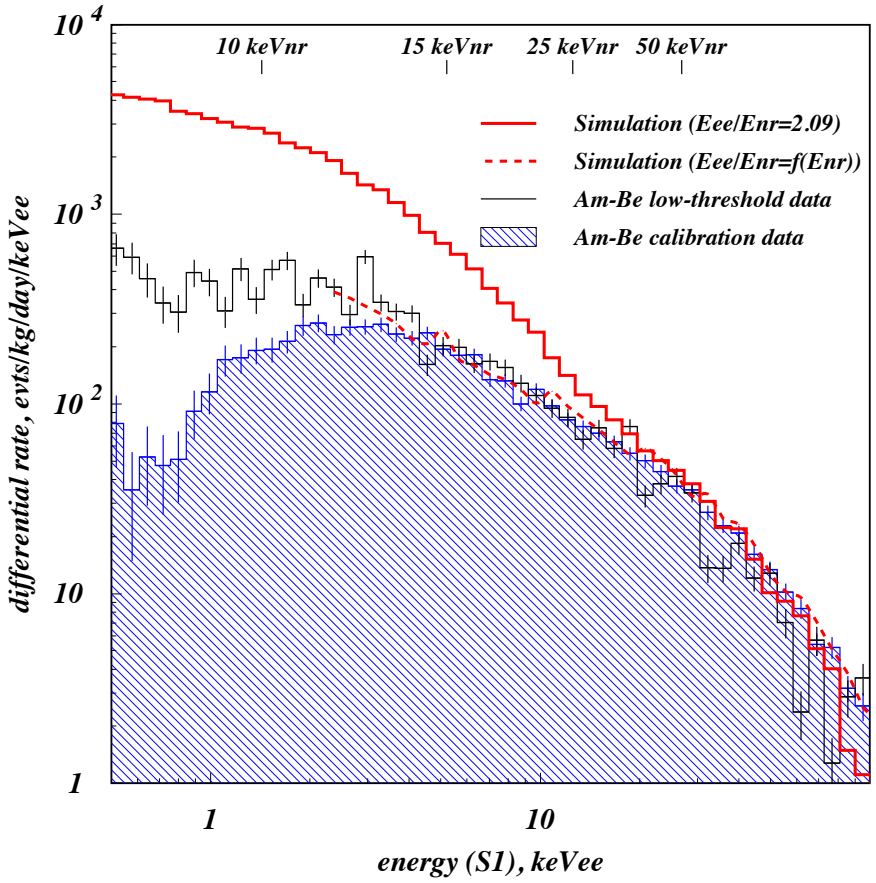

Figure 4.25: Comparison of simulated and experimental data from AmBe calibrations (both standard and low threshold). A large discrepancy is observed below $\sim$10 keV$_{ee}$ which is not explained by instrumental efficiencies. Taken from [83].

distribution, making it impossible to characterise the electron recoil leakage into the WIMP search region. Consequently, a non-blind analysis was performed on the full 100% dataset.

4.10.1 Expected Events in the WIMP Box

Due to the misrepresentation of the data using the ^{137}Cs, the decision was taken to use the electron recoil population from the science data to determine the expected number of events in the WIMP search box. The electron recoil population was

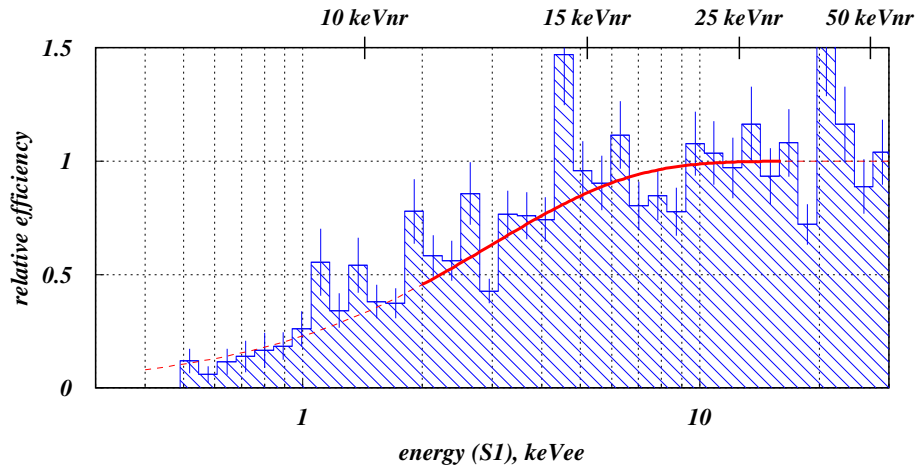

Figure 4.26: Relative efficiency for detection below ∼10 keV_{ee}. The functional fit to the data is used to define L_{eff}. Taken from [83].

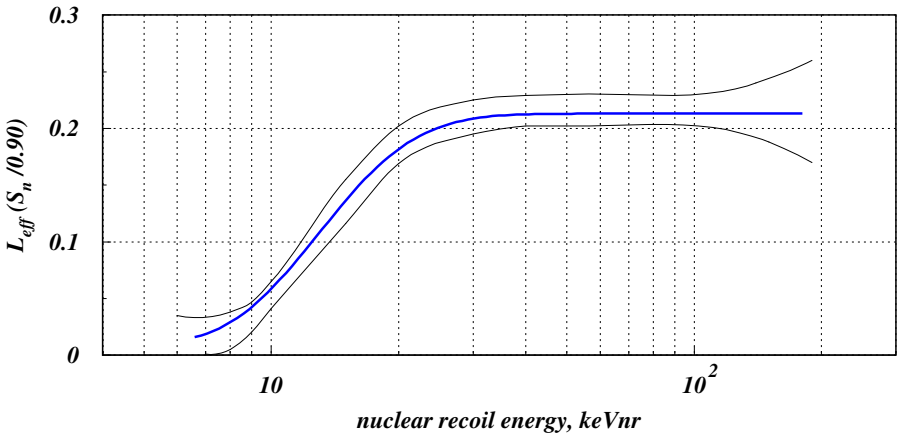

Figure 4.27: The energy conversion factor for the FSR in ZEPLIN-III. The value is a combination of L_{eff} and S_n and $\pm 1\sigma$ bands are shown. Taken from [83].

selected and sliced into 1 $\mathrm{keV_{ee}}$ and the resulting $\log_{10}$(S2/S1) fitted with a Skew-Gauss function of the form:

$$G(x \mid A, x_c, w, \alpha) = \frac{A\sqrt{2}}{w\pi} e^{-\frac{(x-x_c)^2}{2w}} \int_{-\infty}^{\alpha \frac{x-x_c}{w}} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt, \tag{4.3}$$

which reduces to:

$$G(x \mid A, x_c, w, \alpha) = \frac{1}{w\sqrt{2\pi}} e^{\frac{(x-x_c)^2}{2w}} \left[1 + erf\left(\frac{\alpha}{\sqrt{2}} \frac{x-x_c}{w}\right)\right], \tag{4.4}$$

where erf stands for the Gaussian error function.

Three separate methods where utilised to fit each slice of the electron recoil distribution. Firstly, a standard χ^2 minimisation was used. In addition to this, a maximum likelihood (ML) method with a Poisson distribution was used. The motivation for using the ML method was to account correctly for low statistics in the tail of the electron recoil distribution - the region of most interest. The third method is an extension to the ML analysis which consists of finding the best ML fit to the data which simultaneously maximises the probability of getting the observed number of events n_{obs} in each energy slice.

For each energy bin, the fitted $G\left(x; x_c, w, \alpha\right)$ were used to estimate the number of events (μ_{est}) in the box for the limits bounded by the energy slice ($[l_{low}, l_{high}]$):

$$\mu_{est} = \int_{l_{low}}^{l_{high}} G\left(x \mid A, x_c, w, \alpha\right) dx. \tag{4.5}$$

The S2/S1 bounds are defined using the functional form of the AmBe nuclear recoil distribution median. This leads to a stepped box (unlike the curve described in Figure 4.24). The stepped box allows much more simple mathematics without radically altering the final result. Examples of the skew-Gaussian fit are shown in Figure 4.28. The observed and expected values for events in the box are summarised in Table 4.4 and the parameters included in the fits for each energy slice are shown in Figure 4.29. The weighted ML method gives an expectation of (11.64±3.0) events.

The final discrimination plot is shown in Figure 4.30 showing 7 events that fell along the upper edge of the WIMP search region. The fact that they all fall along the upper edge of the search region suggests that they are likely to be due to

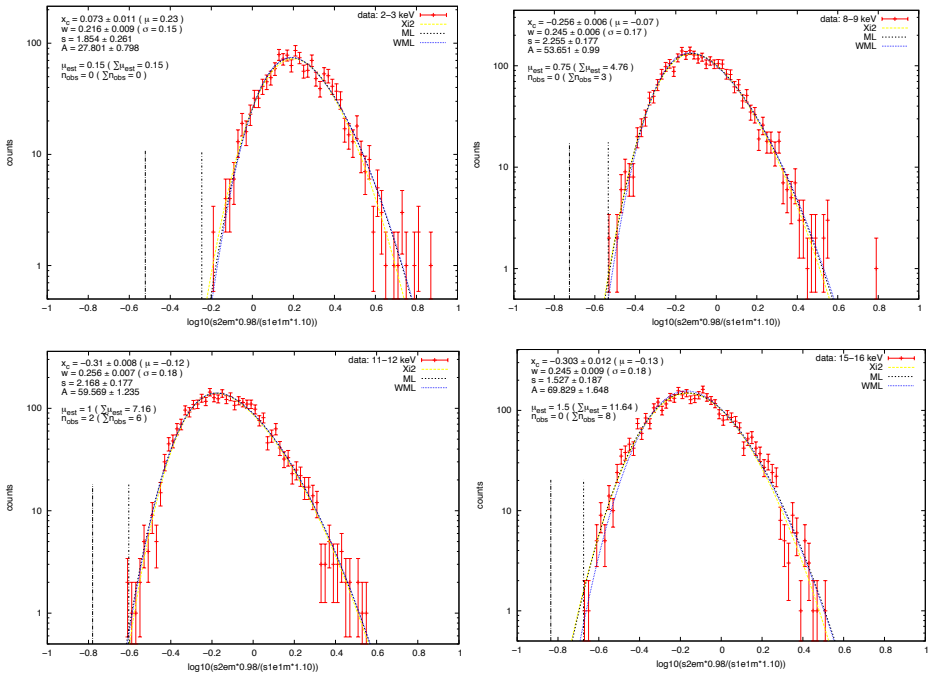

Figure 4.28: Example skew Gaussian fits to the FSR electron recoil distribution.

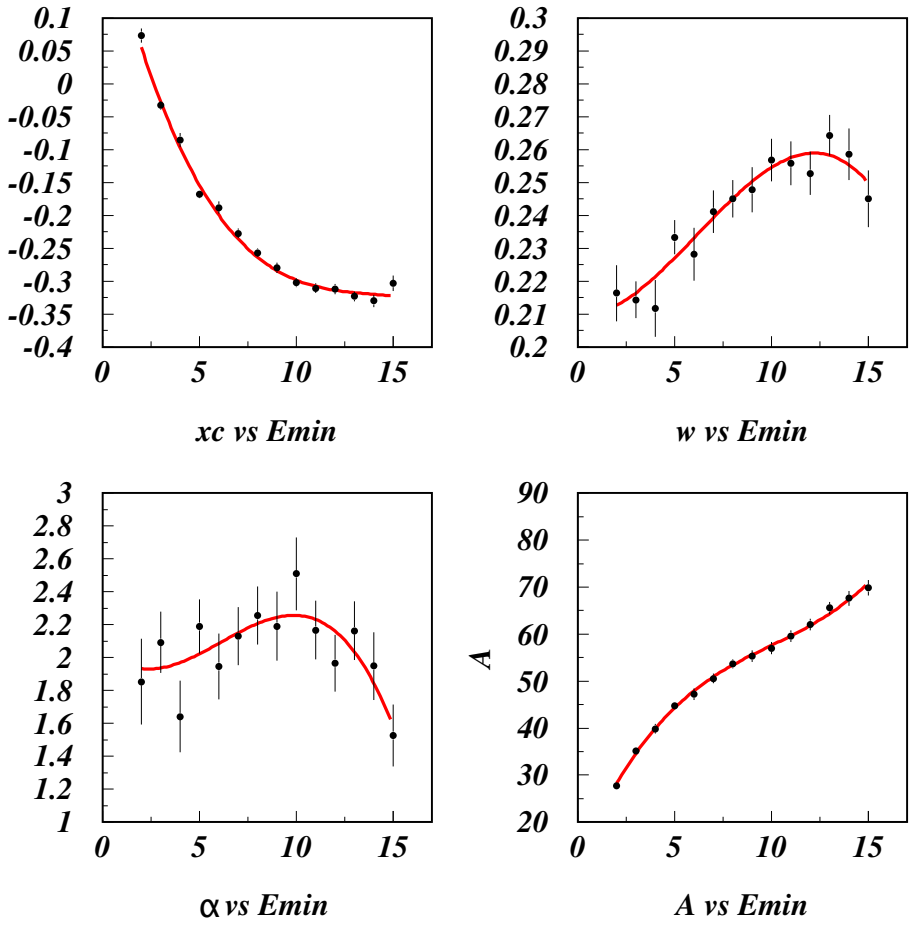

Figure 4.29: Weighted ML parameters. The plots show values for the parameters given in Equation 4.3.

$[E_{low}, E_{high}[$, keV$_{ee}$	I_{low}	I_{high}	n_{obs}	μ_{est}
$[2, 3[$	-0.52	-0.25	0	0.15
$[3, 4[$	-0.57	-0.33	0	0.14
$[4, 5[$	-0.62	-0.38	1	0.96
$[5, 6[$	-0.65	-0.43	0	0.8
$[6, 7[$	-0.68	-0.47	1	0.98
$[7, 8[$	-0.70	-0.50	1	0.99
$[8, 9[$	-0.73	-0.53	0	0.75
$[9, 10[$	-0.75	-0.56	1	1
$[10, 11[$	-0.76	-0.58	0	0.4
$[11, 12[$	-0.78	-0.60	2	1
$[12, 13[$	-0.80	-0.62	1	1.2
$[13, 14[$	-0.81	-0.64	0	0.7
$[14, 15[$	-0.82	-0.66	1	1.08
$[15, 16[$	-0.84	-0.67	0	1.5
		Total	**8**	**11.64**

Table 4.4: The observed and estimated event values in the sliced energy bins.

leakage from the electron recoil population. The spatial distribution of the FSR data with the 7 events highlighted is shown in Figure 4.31. The data can be used to calculate the discrimination power of ZEPLIN-III. The discrimination power describes the ratio of events that are described as nuclear recoil events (by falling within the WIMP search region) compared to the number that fall outside this region. For 2-16 keV$_{ee}$, it can be seen that 7 events fall in the box and ∼36500 events fall outside the box. This equates to a discrimination power of ∼1/5200.

In addition to the electron leakage expectation, simulations are used to calculate the expected rate of nuclear recoil events in the WIMP search box. As has been previously mentioned, the PMT array provides the dominant source of γ-ray background in the FSR. In addition to this, the PMT array provides the dominant source of neutron background through α,n reactions and spontaneous fission from $^{238/235}$U and ^{232}Th. Simulations show an expected rate of about 25 events/year above 10 keV$_{nr}$. This rate reduces to $\lesssim$1 event/year for single scatter events [130]. Although this is a low rate, this limits the length of time that ZEPLIN-III can run before it becomes background limited. This means that, after ∼1 year, it is expected that the data will include an event from a single

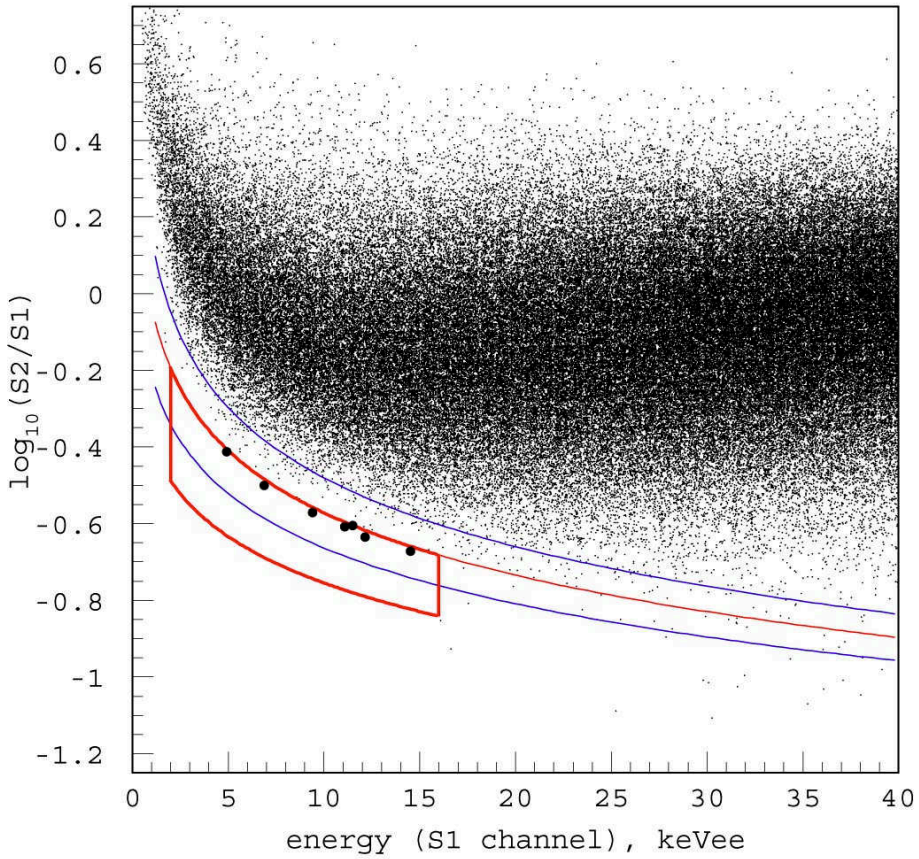

Figure 4.30: Final discrimination plot for the ZEPLIN-III FSR. The 7 events in the box are highlighted.

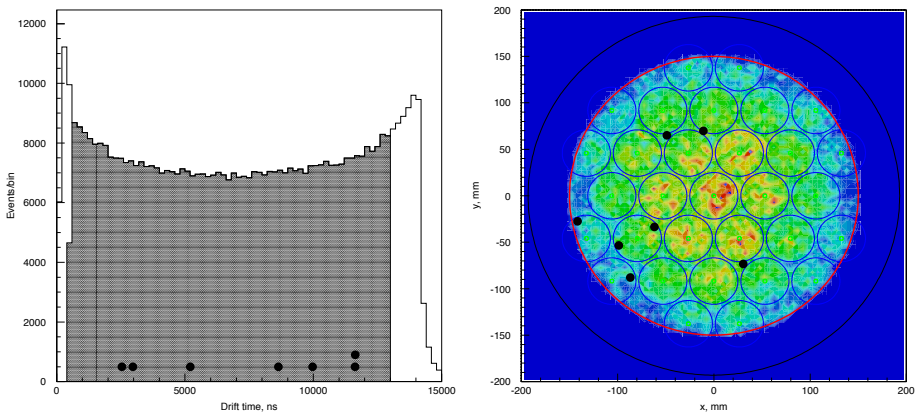

Figure 4.31: Drift time (left) and xy distribution (right) for data in the FSR. The events in the region of interest are highlighted showing a fairly flat positional dependence.

scatter elastic nuclear recoil. In the FSR, this is indistinguishable from a WIMP recoil and, therefore there would be a large uncertainty in the identification of a WIMP-recoil signal. As the FSR lasted only 83 days, it was far from being background limited.

4.11 FSR Limit Calculation

The limit calculation for the FSR data-set is based on three pieces of information about the data plotted in Figure 4.30. The first of these is the reasonable assumption that any expected electron-recoil background falling in the search box should, preferentially, fall in the top part. Based on this, the WIMP search region is divided into two regions which have significantly different probabilities of containing electron-recoil events based on the expectation analysis using the electron-recoil population in the data. Figure 4.32 (taken from [83]) shows the WIMP box divided into nuclear-recoil acceptance percentiles as derived from AmBe data. This is motivated by the fact that a standard Feldman-Cousins [137] analysis given an observed rate lower than the expected rate (11.6±3.0 expected, 7 observed) gives an artificially low signal limit. It is, however, possible to follow a standard Feldman-Cousins analysis by truncating the expected rate to be equal

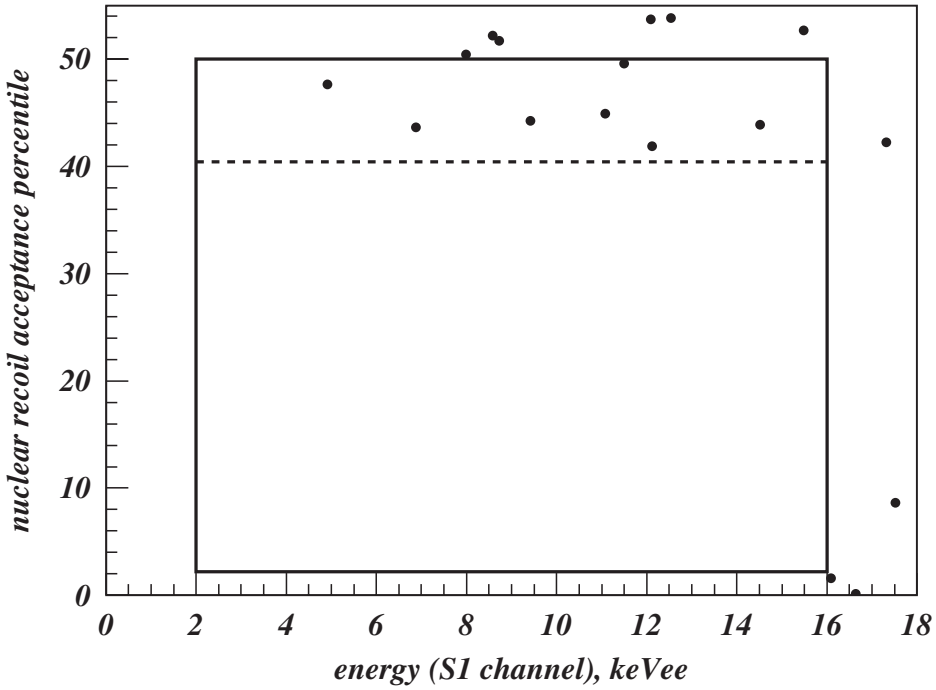

Figure 4.32: The WIMP-search box with the vertical axis remapped onto nuclear-recoil percentiles. The positions of the 7 events falling within the box are shown as well as other events just outside the box. The horizontal dashed line separates the box into two regions with an area ratio of 1:4 [83].

to the observed rate. This method is acceptable as it results in a conservative signal rate upper limit.

Given the nuclear-recoil acceptance percentile representation, any positive WIMP signal should populate this parameter space uniformly. The dashed line is set so as to contain all leakage events above the line and have an empty box below the line. The ratio of box area above and below the dashed line is 1:4.

The second piece of information necessary for the limit calculation is the fact that no events fall in the lower part of the box. The final piece of information is that there may be up to 7 signal events in the upper box.

A classical one sided limit (90% confidence level) calculation determines the value under which 10% of repeated experiments would return zero events in the lower box and up to 7 events in the upper box. This can be represented in terms

of poisson probabilities as:

$$P\left(n_l = 0, n_u \leq 7|\mu\right) = P\left(n_l = 0|f\mu\right) * \sum_{i=0}^{7} P\left(n_u = i|\left(1-f\right)\mu\right) = 0.1, \quad (4.6)$$

where n_l and n_u are the number of events in the lower and upper box respectively, μ is the WIMP expectation value in the whole box and f is the fraction of the box where the dashed line in Figure 4.32 represents the upper bound.

Over the range $f = 0.74 - 0.84$, the calculated result is $2.30/f$. A calculated fraction for the maximum allowed area under the dashed line that excludes all events is $f = 0.84$. The upper limit is driven almost entirely by the fact that the region under the box is empty and, as such, the second term in the Poisson probability function remains close to unity regardless of the value of f.

The value $\mu = 2.30$ is regarded as the classical 90% upper limit on zero. A two sided 90% confidence level upper limit on zero is given, by a Feldman-Cousins statistical approach, as $\mu = 2.44$. For the published limit, a conservative value of $f = 0.8$ was used as below this, the probability of the box containing background is close to zero.

Given a value for f of 0.8, the 90% confidence upper limit becomes $\mu = 3.05$. With this value of μ, there is a 54% probability that the upper fraction of the box contains no WIMP events, a 33% chance of there being one WIMP event and a 13% chance of there being $\geq$2 WIMP events.

4.12 Calculating Limits

Once an upper limit has been placed on the number of events seen in the box, it is possible to begin the process of setting an upper limit on the cross section for both spin-independent and spin-dependent WIMP-nucleon scattering cross-sections. This section will go step by step through the process of setting the upper limit including the numerical factors, mathematics and corrections needed. This section will only discuss the formulae needed for this specific limit analysis but a full review of this process is provided in [138].

Section 3.2 describes the calculation of an expected differential spectrum for

a WIMP-nucleon scatter. Equation 3.9 gives the expected differential equation for this spectrum. In order to calculate the spin-independent and spin-dependent WIMP-nucleon cross-sections, the form factor term $F^2(qr_n)$ must be modified.

4.12.1 The Spin-Independent Form Factors

In the spin-independent calculation, the form factor is calculated as:

$$F^2(qr_n) = \left[3\left[\sin(qr_n) - qr_n\cos(qr_n)/(qr_n)^3\right]^2\right] e^{-(qs)^2}. \quad (4.7)$$

The exponential term in this equation acts to dampen the form factor. This is due to the fact that the nuclei have, in essence, a 'soft edge'. This means that the charge density falls to zero over the finite skin thickness of the nuclei, resulting in a damping of the form factor. In the case of xenon, $s = 0.9$ fm and $r_n = 5.6$ fm.

4.12.2 The Spin-Dependent Form Factors

In the spin-dependent calculation, the two isotopes ^{129}Xe and ^{131}Xe must be considered. This is due to the fact that these isotopes contain an odd-neutron and, as such, make Xe targets sensitive to the WIMP-neutron interaction.

Having calculated the momentum transfer for each energy recoil, the form factor may be calculated. In the case of the spin-dependent interaction, this is slightly more complicated than in the spin-independent case. The momentum transfer is parameterised using the Bonn-CD structure functions [139] and fitted with a 6^{th} order polynomial which gives a reasonable functional description of the structure functions. There are three spin structure functions to consider S_{00}, S_{01} and S_{11} and these can be trivially combined to give an overall spin contribution for each momentum transfer:

$$S(q) = a_0^2 S_{00}(q) + a_1^2 S_{11}(q) + a_0 a_1 S_{01}(q), \quad (4.8)$$

where, for a neutron, $a_0 = 1.00$ and $a_1 = -1.00$ and, for a proton, $a_0 = 1.00$ and $a_1 = 1.00$ giving a positive S_{01} contribution for the proton and a negative contribution for the neutron.

The form factors for ^{129}Xe and ^{131}Xe can then be calculated using:

$$F^2(q) = \frac{S(q)}{S(0)}, \tag{4.9}$$

where $S(0)$ is the combined Bonn-CD structure function for zero-momentum transfer.

The calculations diverge again at this point as the form factor is folded into the differential rate calculation:

$$\frac{dR}{dE_R} = \frac{c_1 R_0}{E_0 r} e^{-c_2 E_R / E_0 r} F^2(q r_n), \tag{4.10}$$

where E_R is the energy of the recoiling nucleus. This gives the expected spin-independent and spin-dependent differential spectra for a Xe detector but does not allow for the intrinsic characteristics of the detector.

The first detector characteristic considered is the energy resolution of the detector. As described previously, this was calculated as 5.4% at 122 keV_{ee} when contributions from the S1 and S2 channels are combined, this is highlighted in Figure 4.12. The contribution from each energy bin is smeared using a normal distribution. The mean of this distribution is given by the centre of the energy bin (*i.e.* for 1-2 keV_{ee}, the mean would be 1.5 keV_{ee}) and the variance (width) is given by the detector resolution (σ_{ee}) for the specific energy range. This smearing is applied to the entire spectrum and an expected rate contribution in each energy bin due to all other energy bins is calculated.

Finally, the characteristic efficiencies of the detector are folded in. These efficiencies can be found in Table 4.3 and also in [83]. The efficiencies include a calculated efficiency of 73% for the selection cuts including the fiducial cuts and cuts applied to remove MSSI events. The event selection efficiency is calculated by applying selection cuts to the AmBe calibration data and calculating the loss of nuclear recoil events in the WIMP search box.

Once all efficiencies have been included, the expected differential spectrum for WIMP nuclear recoils for both spin-independent and spin-dependent interactions can be calculated. An example spectrum is shown in Figure 4.33.

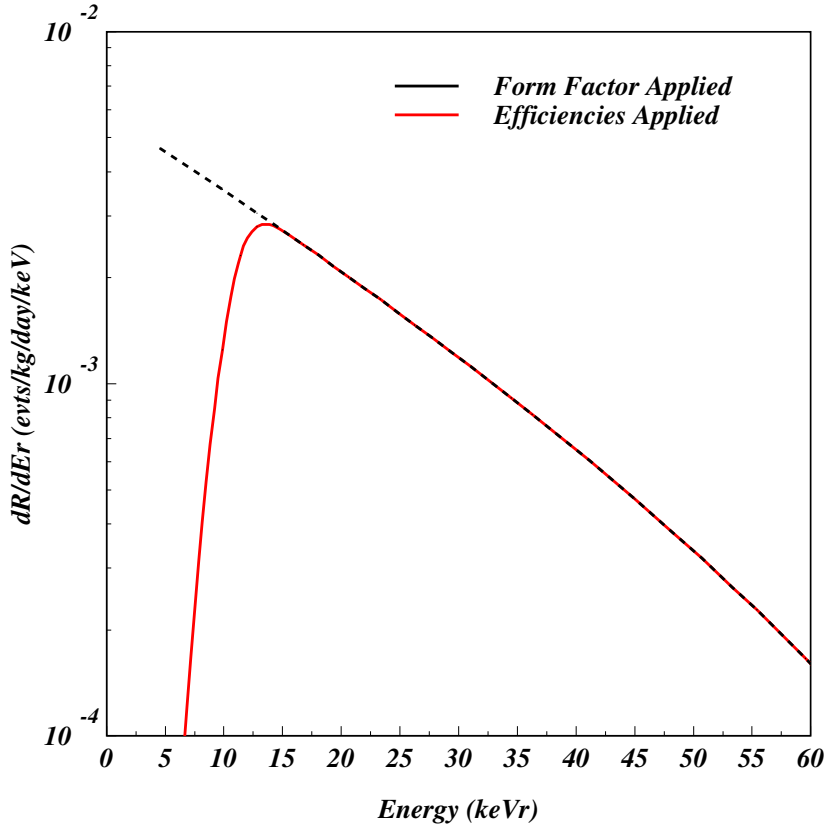

Figure 4.33: An expected differential spectrum for spin-independent WIMP-nuclear recoils. The black line shows the spectrum with just the form-factor correction and the red shows the spectrum after the detector efficiencies have been folded in.

4.13 Comparing Expectation with Data

As with the form factor measurements, the calculations for spin-dependent and spin-independent cross section limits diverge at this point. The spin-independent calculation is relatively simple, relying on only two equations. The first describes the relationship between the measured and expected rates and the WIMP-nuclear interaction. For an assumed local dark matter density of 0.3 GeVc^{-2}cm^{-3} and an average (as described by a Maxwellian distribution) WIMP velocity of 220 km s^{-1}:

$$\frac{R_0}{r} \equiv 94.3 \left(\frac{\sigma_0}{1 \text{ pb}}\right) \left(\frac{1 \text{ GeV}c^{-2}}{\mu}\right)^2 \left(\frac{\rho_D}{0.3 \text{ GeV}c^{-2}\text{cm}^{-3}}\right) \left(\frac{v_0}{220\text{km s}^{-1}}\right) \text{ tru}, \tag{4.11}$$

where R_0 is the total rate from the expected WIMP-nuclear recoil spectrum, r is the observed rate and tru is in units of events/kg/day. This equation is reduced to:

$$\frac{R_0}{r} = \frac{D\sigma_A(\text{pb})}{\mu_A^2}, \tag{4.12}$$

where μ_A is the reduced mass for the WIMP-target system and σ_A is the WIMP-nuclear cross-section.

In dark matter searches, it is customary to give a cross-section in terms of the WIMP-nucleon interaction, $\sigma_{\chi-n}$. This is done to remove the dependence on the target species to give a value that is comparable to, for example, argon or germanium detectors.

In the case of a spin-independent limit calculation, it is simply a case of normalising the calculation to a specific nucleus of atomic mass A:

$$\sigma_{\chi-n} \text{ (pb)} = \left(\frac{\mu_1}{\mu_A}\right)^2 \left(\frac{1}{A^2}\right) \sigma_A \text{ (pb)}. \tag{4.13}$$

where μ_1 is the reduced mass for the WIMP-target system for a target mass $A = 1$.

Equations 4.12 and 4.14 are combined to form a calculation for the WIMP-nucleon cross section that requires no knowledge of the WIMP-nuclear cross-section:

$$\sigma_{\chi-n}\ (\mathrm{pb}) = \left(\frac{\mu_1^2}{A^2 D}\right)\frac{R_0}{r}\ (\mathrm{pb}). \tag{4.14}$$

All that remains in order to calculate a cross-section is to solve the equation for an array of potential WIMP masses. This is usually done in the range of 10-10,000 GeVc^{-2} in order to satisfy a full range of theoretical values. The calculated data are then plotted and a minimum calculated. The cross-section is then quoted as the minimum value but with the specific WIMP mass quoted.

The calculation for the spin-dependent cross-sections is somewhat more complicated. Firstly, the calculation must be performed for both the ^{129}Xe and ^{131}Xe components. The WIMP-nuclear cross-section is calculated using equation 4.12 normalised to the abundance of the specific isotope in the Xe used in ZEPLIN-III. Naturally occurring xenon contains 26.4% and 21.2% of ^{129}Xe and ^{131}Xe (by number), respectively. However, in the ZEPLIN-III xenon, the levels of ^{136}Xe and ^{134}Xe are depleted. The depletion boosts the abundance of ^{129}Xe and ^{131}Xe to 29.5% and 23.7%, respectively.

The isotope specific WIMP-proton and WIMP-neutron interaction cross-sections can then be calculated using the following equation [138]:

$$\frac{\sigma_{p,n}}{\sigma_A} = \frac{3}{4}\frac{\mu_{p,n}^2}{\mu_a^2}\frac{1}{\langle S_{p,n}\rangle^2}\frac{J}{J+1}, \tag{4.15}$$

where p, n denotes proton and neutron respectively, $\langle S_{p,n}\rangle$ is the spin factor and J is the angular momentum of the odd neutron in both isotopes.

The values for WIMP-proton and WIMP-neutron cross-sections for the two Xe isotopes are then combined to give an overall spin-dependent cross-sections for the ZEPLIN-III xenon.

4.14 Results from the ZEPLIN-III FSR

The analysis of the 847 kg.days (∼450 kg.days fiducial) of data from the FSR of ZEPLIN-III gave a signal lower limit consistent with zero and a spin-independent WIMP-nucleon elastic interaction cross-section of 8.1×10^{-8} pb for a WIMP mass of 60 GeVc^{-2} at 90% confidence level [83]. This makes ZEPLIN–III the most sensitive liquid noble gas detector for WIMP masses above 80 GeVc^{-2} and

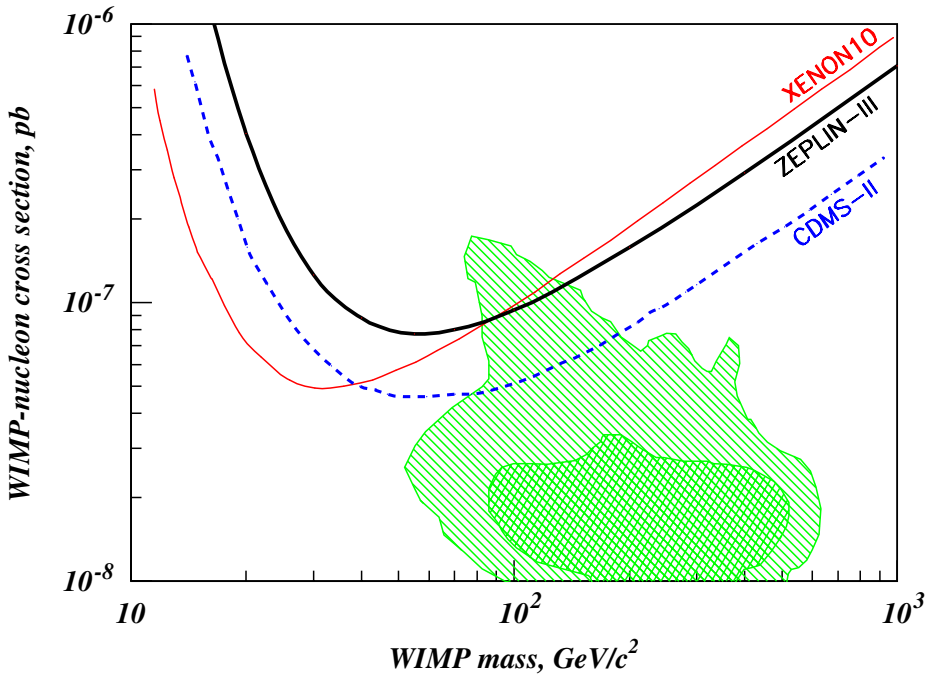

Figure 4.34: 90% confidence upper limit to the WIMP-nucleon elastic scattering cross-section as derived from the FSR of ZEPLIN-III for a spin-independent interaction. For comparison, the experimental results from XENON10 [141] and CDMS-II [142].

also the most sensitive European dark matter device.

In the spin-dependent channel, the combine cross-sections for WIMP-neutron and WIMP-proton interaction cross-sections are $\sigma_n = 1.9 \times 10^{-2}$ pb and $\sigma_p = 4.1 \times 10^{+2}$ pb, respectively [140].

A second analysis was performed using the Bonn-A parameterisation of spin structures. This was done for comparison with other dark matter search experiments and gives more favourable (*i.e.* more sensitive) results. However, the Bonn-CD parameterisation is a more recent and accurate parameterisation of the spin structure functions when compared to Bonn-A.

The results of the first science run of ZEPLIN–III were a remarkable success, delivering a sensitivity almost exactly as expected. The detector also showed the best discrimination power for a 2-phase liquid noble gas detection system. The

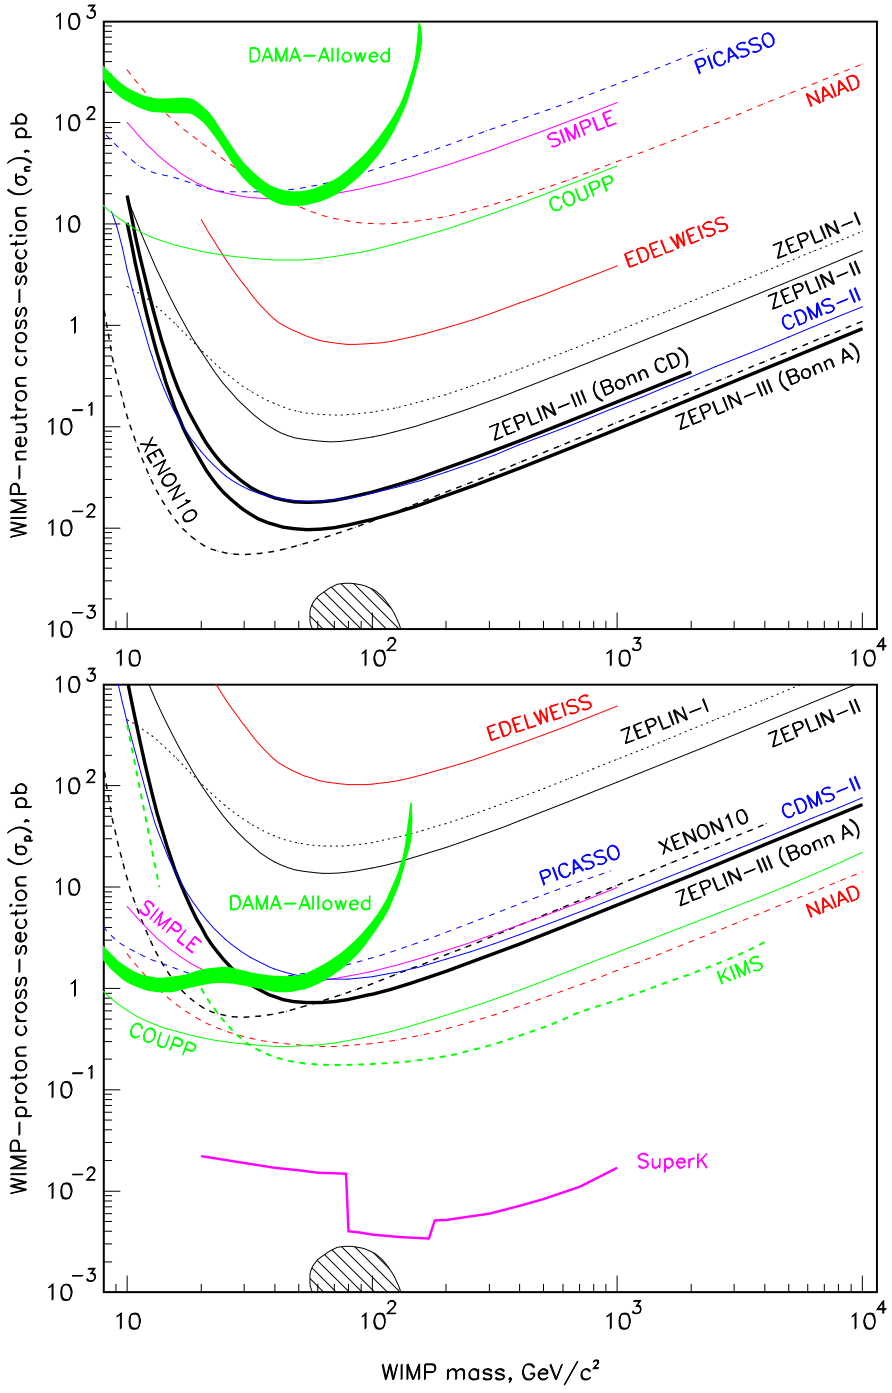

Figure 4.35: Upper limits for the WIMP-neutron and WIMP-proton spin-dependent cross-sections. In addition to ZEPLIN-III, other Xe experiments are shown in black: ZEPLIN-I [143], ZEPLIN-II [143] and XENON10 (Bonn-A) [144]. Also shown are curves from CDMS-II [142], COUPP[145], EDELWEISS [146], KIMS [147], NAIAD [90], PICASSO [148] and SIMPLE [149]. The hatched are is the tip of the 95% probability region for neutralinos in CMSSM [150]

first science run was intended as an engineering run to test all the systems. In order to improve the result from this run, the detector entered a period of upgrade. To improve the sensitivity of ZEPLIN-III a new PMT array was developed with a much lower intrinsic background than that of the array used in the first science run. The next chapter gives detail on the development of the PMT array.

Chapter 5

Component Radiation Measurements

In rare event searches, it is essential to know the radiation budget of the system to a level of accuracy not needed in many other experiments. Not only is it important to know the budget but it is important to know the effect each source of radiation will have on the result as a whole so that expected rates and spectra for both γ-ray induced electron-recoil background and neutron induced nuclear-recoil background can be accurately calculated. Only then can a signal in ZEPLIN-III be assessed as being above the expected level of background or not. For this reason, components that are not shielded from the detection volume must be radio-assayed in order to minimise and understand their impacts. The Boulby underground facility plays host to a high-purity germanium (HPGe) detector. This detector has been widely used for both environmental measurements and component measurements in the past [151]. I personally lead the work on the radio-assaying of the various components for use in the ZEPLIN–III experiment. The radio-assaying tests were performed in order to test the components to be used in the upgrade of ZEPLIN–III in order to reduce and characterise the exact backgrounds in the system. For the second science run, a PMT background reduction of 10x over the background from the first science run PMT array was targeted.

In any naturally sourced material, the dominant source of radioactivity comes from uranium, thorium and potassium contamination. It is also common for

Decay Isotope	Energy (keV)	Branching Ratio	Source Reaction
^{238}U	295	0.436	^{214}Pb→^{214}Th
	352	0.304	^{214}Pb→^{214}Th
	609	0.266	^{214}Bi→^{214}Po
	1765	0.356	^{214}Bi→^{214}Po
^{232}Th	239	0.177	^{212}Pb→^{212}Bi
	583	0.342	^{208}Tl→^{208}Pb
	911	0.428	^{228}Ac→^{228}Th
	2615	0.147	^{208}Tl→^{208}Pb
^{40}K	1460	0.11	^{40}K→^{40}Ca

Table 5.1: The peaks of interest for ^{238}U, ^{232}Th and ^{40}K observed in the HPGe detector.

cosmogenically produced ^{60}Co to provide contamination. The elements ^{238}U and ^{232}Th follow the decay chains shown in figures 5.1 and 5.2.

Following the trend of mass number along each decay chain, it is clear to see that each decay chain proceeds via either α or β^- decay. The process of β^- decay causes emission of an electron, often leaving the daughter nucleus in an excited state. Through a process of de-excitation, γ-rays of varying energies are emitted which are characteristic of the decay from parent to daughter nucleus. A specific parent nucleus may β^- decay and leave an excited daughter nucleus in more than one way, leading to a different de-excitation pattern.

In the decay from $^{214}_{83}$Bi→$^{214}_{84}$Po, the β^--decay will leave an excited nucleus with an energy of 1765 keV, 16.9% of the time. This excited state can decay via emission of a 1765 keV γ-ray directly to the ground state. The decay also leaves a $^{214}_{84}$Po daughter nucleus with an excitation energy of 1729 keV 17.5% of the time. The preferred de-excitation route of this energy level involves 2 characteristic γ-rays. The route taken is to emit a γ-ray with an energy of 1120 keV followed by the emission of a γ-ray with an energy of 609 keV (to the ground-state) giving a total energy of 1729 keV, the initial excitation energy.

Also of importance in the decay of both ^{238}U and ^{232}Th is α decay. A typical α particle from a radioactive decay has energy of 5-6 MeV due to the Geiger-Nuttal law [152]. If this α particle collides with an atom containing a loosely bound neutron (in the case of ^{9}Be, used in an AmBe neutron source, the binding

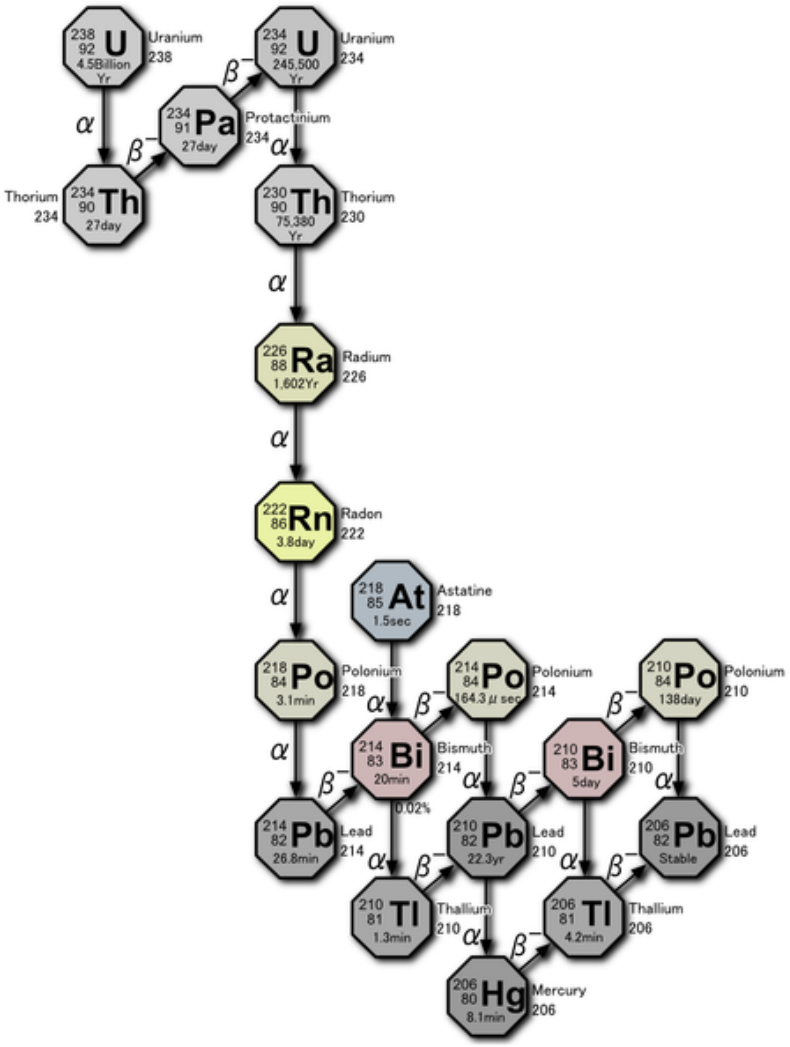

Figure 5.1: The decay scheme for ^{238}U [152].

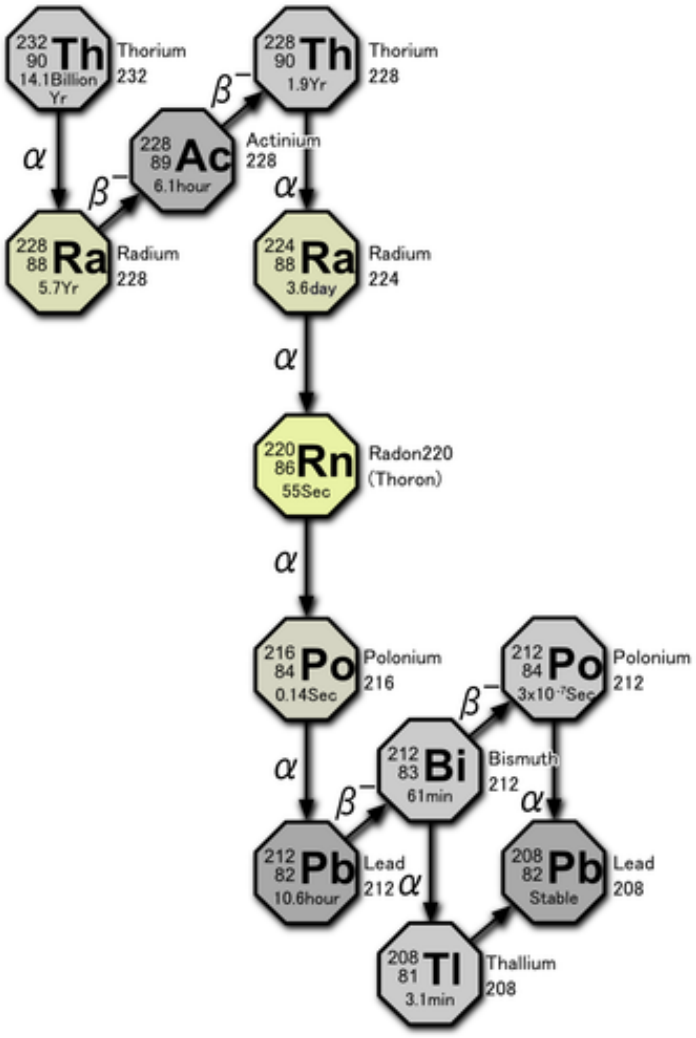

Figure 5.2: The decay scheme for ^{232}Th [152].

energy of the unpaired neutron is ∼1.7 MeV), it is possible that a neutron can be released through an (α,n) reaction [153, 154] for example:

$$\alpha + {}^{9}\mathrm{Be} \rightarrow {}^{12}\mathrm{C} + \mathrm{n}. \tag{5.1}$$

The radioactive Q value for this reaction is 5.7 MeV and the decay chain of ^{238}U provides α particles with energies of up to 7.8 MeV making this reaction energetically possible. Of course, there are other naturally occurring elemental isotopes with loosely bound neutrons so it is possible for neutrons with a broad range of kinetic energies to be produced.

The HPGe detector in the Boulby facility is shielded from rock-borne radioactivity by a combination of 20 cm of lead and 20 cm of copper shielding. Simulation shows this to be effective at excluding >99% of external γ-ray radiation from the system from non-gaseous sources. This being the case, the background seen by the detector is dominated by radiation within the shielding itself and from radiation introduced into the system from ^{222}Rn. As radioactive contaminant energy peaks are visible in both background and sample data sets, it is important to analyse live time normalised, background-subtracted spectra. When an acquired background spectrum is viewed, characteristic energy peaks can easily be seen. Figure 5.3 shows a typical energy spectrum and Table 5.1 summarises the energy peaks seen, their intensities and the original β^--decay source parent-daughter pair.

5.1 Simulation

An HPGe detector is not 100% efficient across all energies. The efficiency of the detector is also affected by radiation self absorption in the sample being measured. The level of self absorption is sample specific and depends on factors including the sample mass, shape, density and position within the detector setup. For this reason, it is important to ascertain an energy dependent efficiency of the detector.

A GEANT4 [155] Monte-Carlo simulation has been created for the HPGe detector located in the low background lab (see Figure 5.4). In this simulation, both the lead (Pb) and copper (Cu) shielding are represented. Also included are the HPGe crystal and the steel enclosure.

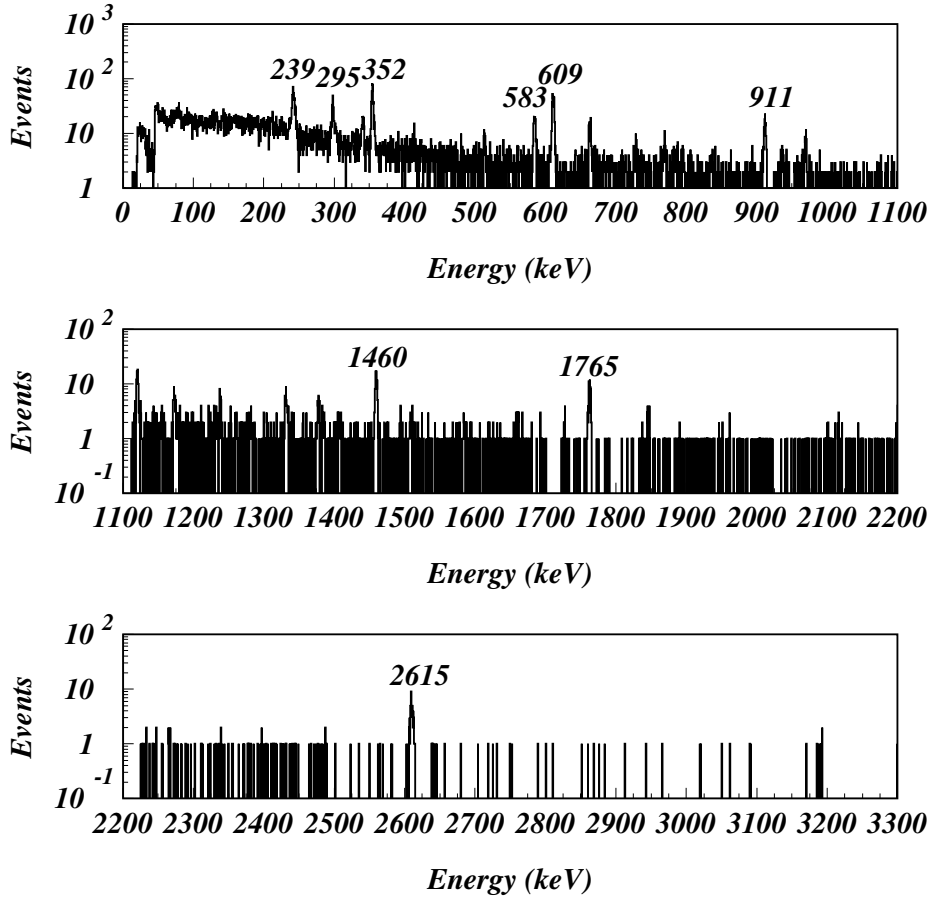

Figure 5.3: Typical background spectrum for the HPGe detector based at the Boulby facility. Key energy peaks are labelled.

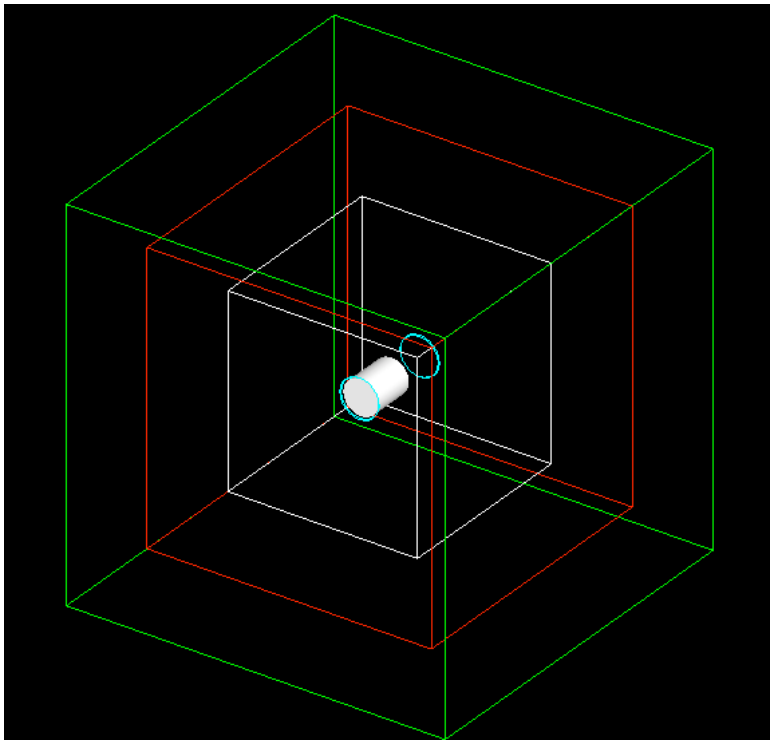

Figure 5.4: The GEANT4 simulation of the HPGe detector. The green box represents the Pb shielding and the red box the Cu shielding. The white cylinder represents the Ge crystal whilst the blue cylinder represents the steel enclosure around the crystal.

Once a simulation of the detector setup has been created, it is possible to add sample materials into the simulation (this is shown in Figure 5.5 with the simulated PMT test shown next to the actual experimental setup). If isotropic γ-rays (with a flat energy spectrum) are then fired from the sample, it is possible to build a picture of the sample specific detection energy as a function of γ-ray energy. This then accounts for self absorption from the sample, affecting the detection efficiency of emitted γ-rays. The germanium crystal is treated as a "sensitive" volume in the simulation which means that any energy deposited within the volume by a γ-ray is recorded. The simulation runs event by event so for each γ-ray emitted, its total deposited energy is recorded.

The simulation is run such that several million isotropic events are fired from a simulated sample (GEANT4 allows chemical structure and density, amongst

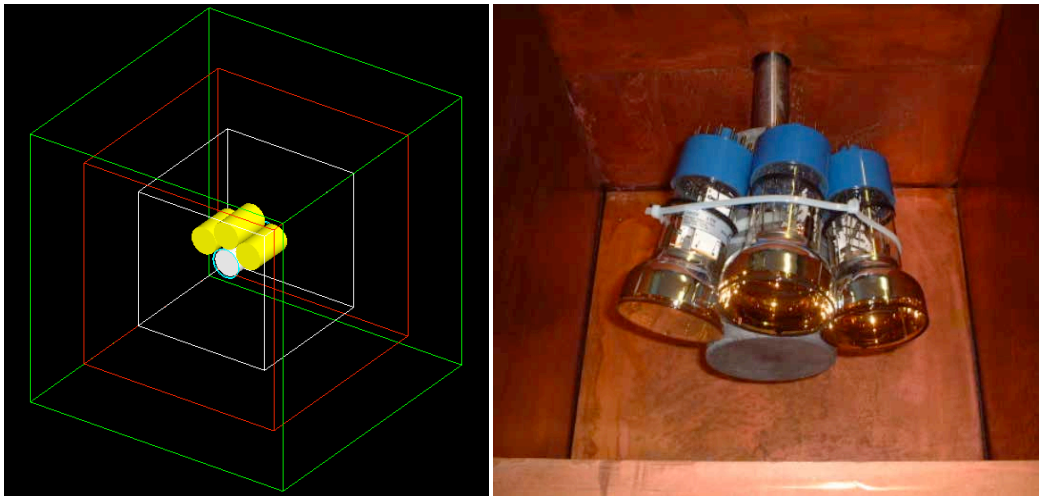

Figure 5.5: Simulation and actual placement of the ZEPLIN–III veto PMTs in the HPGe detector.

other things, to be factored in) and the energy deposited in the sensitive volume can be collected with minimal statistical uncertainty. Figure 5.6 shows the sample specific efficiency for the simulation of PMTs used in the veto detector.

5.2 Calculating a Final Concentration

A Gaussian fit is applied to each of the peaks of interest observed (detailed in Table 5.1) in the recorded HPGe spectrum and the resultant distribution can be integrated to give a number of counts contributing to the peak. The values for the fits obtained using a background spectrum are scaled for exposure time and then subtracted from the sample spectrum to give an overall count rate for each peak of interest. Once the data is acquired and fitted, and the GEANT4 efficiency for that energy applied, a concentration for each of the contaminants can be calculated.

Initially, the number of atoms of ^{232}Th needed (N_i) to give one disintegration per day can be calculated:

$$N_i = \left(\frac{1}{\ln 2}\right) t_{1/2}, \tag{5.2}$$

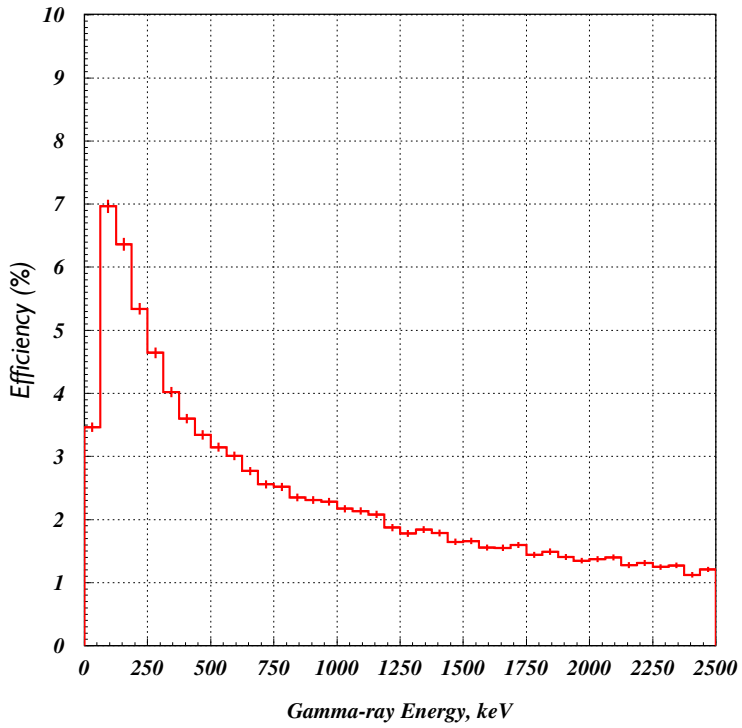

Figure 5.6: The energy dependent efficiency for a sample of 3 veto PMTs placed on the HPGe detector head.

where the half-life, $t_{1/2}$ is given in days. This equation can be modified in order to calculate the mass concentration needed to give one disintegration per day per kilogram:

$$C = \left(\frac{1}{\ln 2}\right) \frac{A_i t_{1/2}}{N_A}. \tag{5.3}$$

where N_A is Avagadro's constant.

At this point, the values should be scaled using the isotope mass fraction. In the case of ^{232}Th this is approximately equal to 1. For a chain of n decays (for Th and U), the above equations must be multiplied by $1/n$. In the case of ^{232}Th, $n = 10$; for U, there is the added complication that 95.6% of primary decays are ^{232}U ($n = 14$), and 4.4% are ^{235}U ($n = 11$), giving $n = 13.87$ overall.

Additional multiplicative factors are included to account for the probability of producing β and γ-ray events (n_γ and n_β) giving an overall equation for mass concentration to give one disintegration per day of:

$$C = \left(\frac{1}{\ln 2}\right) \frac{A_i t_{1/2}}{N_A} \left(\frac{n_\gamma n_{\beta^-}}{n}\right). \tag{5.4}$$

This equation can then be scaled to give a relationship between Bq/kg and parts per million or billion (ppm or ppb). Using the necessary values for potassium, thorium and uranium, the following conversion factors can be calculated:

$$\begin{aligned} 1 \text{ ppm K} &= 0.0309 \text{ Bq/kg}, \\ 1 \text{ ppb Th} &= 0.00406 \text{ Bq/kg}, \\ 1 \text{ ppb U} &= 0.0129 \text{ Bq/kg}. \end{aligned} \tag{5.5}$$

A value for Bq/kg can then be calculated using the detector efficiency, *eff*, branching ratio of each decay, *br*, the mass of the sample, m, and the background subtracted rate of each peak, $rate_{\text{bs}}$:

$$\frac{\text{Bq}}{\text{kg}} = 1.16 \times 10^{-5} \left(\frac{rate_{\text{bs}}}{(eff \times br \times m)}\right). \tag{5.6}$$

5.3 Calculating Errors

The statistical nature of counting events in the HPGe detector where concentrations of contaminants are, by design, low means that a standard statistical error is pertinent. A brief summary of error analysis is included since, for many samples, only upper limits on contaminations can be calculated. Errors on the sample and background counts are independent and can therefore be combined in quadrature as follows:

$$\text{error} = \sqrt{\left(\frac{\sqrt{\text{sample counts}}}{\text{sample livetime}}\right)^2 + \left(\frac{\sqrt{\text{background counts}}}{\text{background livetime}}\right)^2}. \tag{5.7}$$

The final value for the ppm/ppb contamination is calculated using a weighted mean taken from the contamination and the error on the contamination (derived using the error on the counts):

$$\text{ppb/ppm} = \frac{\left(\frac{x_1}{\Delta x_1^2} + \frac{x_2}{\Delta x_2^2} + \cdots + \frac{x_n}{\Delta x_n^2}\right)}{\left(\frac{1}{\Delta x_1^2} + \frac{1}{\Delta x_2^2} + \cdots + \frac{1}{\Delta x_n^2}\right)}, \tag{5.8}$$

where x is the individual contamination calculated for each peak and Δx is the error on each peak. The overall energy can also be calculated, where n is the number of counts:

$$\Delta(\text{ppm/ppb}) = \frac{\sqrt{n}}{\left(\frac{1}{\Delta x_1^2} + \frac{1}{\Delta x_2^2} + \cdots + \frac{1}{\Delta x_n^2}\right)}. \tag{5.9}$$

A summary of the contaminations of various ZEPLIN–III and veto components can be found in Table 5.2, which shows the wealth of information obtained using this technique. The samples measured were components required for the second science run upgrades to the PMT array and components required in the veto detector (discussed in the following chapter). In addition to this, the reduced background expected from the PMTs in the second science run required a deeper knowledge of background from other components such as the electric feedthroughs (made from ceramic material) and the polypropylene shielding.

Table 5.2: Radiological content with statistical uncertainties where appropriate of veto components as assayed either by direct observation of γ-ray emission (HPGe) or through mass spectroscopy techniques (ICP-MS/OES).

Component	Mass, kg	Radiological content		
		U (ppb)	Th (ppb)	K (ppm)
HPGe measurements				
Plastic scintillator	1057.0	0.2±0.3	0.1±0.7	0.2±0.6
PTFE inner wrap	8.9	1.3±0.2	0.2±0.5	1.2±0.4
Silicone	0.1	2.9±0.4	0.5±0.8	5.7±1.1
PTFE tape	3.1	3.2±1.3	6.1±1.1	3.9±1.0
Veto PMTs	6.2	38.0±0.8	21.1±1.2	65.5±2.4
PMT preamplifiers	0.7	8.4±1.7	13.2±2.2	10.1±1.7
PMT base	5.5	12.7±1.4	14.8±2.4	20.2±2.4
Epoxy	70.0	2.5±0.6	0.9±0.3	0.6±0.1
Gd oxide	8.0	0.9±0.1	1.2±0.3	1.7±1.1
ICP-MS/OES				
Copper tape	26.0	1.9±0.2	2.9±0.4	14.0±2.0
PTFE inner wrap	8.8	2.0±1.0	5.0±1.0	<4
Veto PMTs	6.2	30.2±2.2	30.0±3.7	60±2.2
PMT preamplifiers	0.7	10.3±0.5	29.7±3.2	24±3.7
PMT base	5.5	13±3.4	19±2.0	21±3.0
Polypropylene	510	<1	<1	<5
PMT mounting	15.8	30±7.8	<10	<10
Cabling	30.2	110±5.4	20±3.2	29±7.3
Connectors	2.1	<10	<10	<4
Optical gel	0.3	<1	<1	<1
Gd oxide	8.0	2.5±0.5	3.4±0.7	<4

5.4 The Impact of Contaminants in ZEPLIN–III

Once the contamination due to individual radio-isotopes has been found, it is necessary to model the γ-ray and neutron emission from each component within a simulation that includes an accurate ZEPLIN–III geometry in order to determine whether a sample contamination is acceptable for components used in the upgrade and, if so, what the impact in the detector is in terms of both nuclear recoils and electron recoils.

The simulations for the ZEPLIN–III detector [130] were performed primarily using the GEANT4 package, as well as custom built libraries and the ROOT analysis framework [156]. GEANT4 is an object-oriented toolkit, written in the C++ programming language, and is used for the simulation of the passage of particles through matter. Thus, it provides an ideal environment for establishing theoretical rates of particle interaction in the ZEPLIN–III detector, with end-to-end simulation.

The simulation package required the ability to handle very low energy nuclear recoils $\sim$ few keV, as well as the simultaneous presence of neutrons with energies up to $\sim$ 10 MeV. Furthermore, accurate simulation of recoil events in the detector requires each neutron to be handled individually with all interactions recorded event by event. GEANT4 is not the only simulation package available and, as such, other packages such as MCNP [157] or FLUKA [158] were considered. These other packages, however, did not achieve the detail across the whole energy spectrum achieved with GEANT4. In FLUKA, for example, interactions with neutrons with energy below 20 MeV are not treated on an individual basis. Instead they are averaged and treated statistically.

The structure of the simulations can broadly be broken into several sections that deal with aspects of physics processes included in the simulation, geometry of the matter present, the step-by-step running of the simulation, and the output of statistics & results generated in running the simulation.

The object-oriented nature of GEANT4 allows for the simple addition of specific models into the simulation as required by the user. The ZEPLIN–III simulation itself includes a number of libraries and classes to fully represent almost all possible interaction processes. The most important model used in the simulation

is that which allows particle transport through material. The simulation also uses low energy models for γ-rays including Rayleigh and Compton scattering, the photo-electric effect, and γ-ray conversion. The simulation also includes ionisation and Bremsstrahlung for electrons and positrons and pair production and capture for muons. Ionisation models for hadrons and multiple scattering models for all charged particles were also included. Scintillation, electroluminescence and radioactive decay models were registered to simulate the light emission process in the Xe volume.

ZEPLIN–III is designed for the identification of nuclear recoils and, within the liquid volume, these can be produced by four distinct processes with varying significance:

Elastic Scattering - Elastic scattering, leaving the nucleus with no excited states, is the interaction most significant in dark matter searches. Elastic scattering is the favoured process by which WIMPs are expected to interact with atomic nuclei. GEANT4 samples the differential neutron cross-section of the nuclei to give a final state through the use of two supported representations. Where data is available, a tabulation of the differential cross-section as a function of scattering angle is used. Where data is not available, statistical predictions for the cross-sections are used.

Inelastic Scattering - Postulated as another method by which WIMPs can interact with matter [159], the inelastic interaction (n,n$'\gamma$) is the process by which a neutron scatters off a nucleus leaving the nucleus in an excited state. The final state is given by the neutron, recoiling nucleus and daughter products from the decay of the excited nucleus. As well as γ-ray emission, the de-excitation may also result in the emission of an α-particle. The GEANT4 model includes a number of possible final states. The energy and angular distribution of the interaction products are determined by tabulated values or as a series of Legendre polynomials as described for elastic scattering.

Fission - This process is unimportant for component contaminations in ZEPLIN–III. However, fission processes become important for muon-induced neutron studies. In fission, the neutron is captured by the nucleus which then decays by splitting into at least two fragments. For neutron induced fission in GEANT4, the

four most common fission channels are used. Neutron yield is tabulated as a function of both incoming and outgoing neutron energy. The angular distributions are, again, either taken from a tabulated source or inferred using Legendre polynomials. There are no known fissile isotopes of Xe, although, ^{135}Xe (produced as a uranium fission product in nuclear reactors) has the largest known neutron capture cross-section.

Radiative Capture - This is a process which becomes important when considering the capture of thermal neutrons by both hydrogen and gadolinium in the ZEPLIN–III shielding. Radiative capture occurs when a slow moving neutron is captured by the target nucleus with the daughter nucleus de-exciting via the emission of a photon (n, γ). The final state is described using photon multiplicities or photon production cross-sections, as well as the angular distributions of the emitted photons.

The models are chosen and registered in the user defined "Physics List" along with specific particles to be included. At this point, it is also possible to apply energy cuts to be placed on particles. These cuts define an energy, below which a particle is not permitted to generate a secondary particle through its own interactions. This is useful in reducing the processing time per event by neglecting any interactions that have no significant impact on the result of the simulation.

5.4.1 Electron Recoil Background Impact

Using the conversion factors developed in equation 5.5, it is trivial to calculate a γ-ray flux in terms of γ-rays/kg/day. For each of the isotopes, this is as follows:

$$\begin{aligned} \text{1 ppm of K emits 2310 } \gamma\text{-rays/kg/day,} \\ \text{1 ppb of Th emits 958 } \gamma\text{-rays/kg/day,} \\ \text{1 ppb of U emits 278 } \gamma\text{-rays/kg/day.} \end{aligned} \tag{5.10}$$

Given the mass and contamination of the sample, and the branching ratio of each of the isotopes of interest, an energy spectrum for γ-ray emission from the sample can be calculated. The dominant sources of electron recoil background in the second science run of ZEPLIN–III are the polypropylene shielding, the

PMT array, the ceramic high voltage feedthroughs and the veto detector (plastic scintillator, PMTs, PTFE wrapping, *etc.*).

For each of the contributing components, a γ-ray spectrum was fired and the energy deposition in ZEPLIN–III recorded. Figure 5.7 shows the resultant expected single scatter spectrum. The spectrum shows an expected rate of 0.82 ± 0.05 events/kg/day/keV$_{ee}$ for energies below 100 keV$_{ee}$.

5.4.2 Neutron Impact

As ZEPLIN–III is designed to detect the nuclear recoil signature of a WIMP-nucleon interaction, it is of paramount importance to understand the neutron background of the detector. The neutron emission spectrum for components in ZEPLIN–III is, as in all cases, sample dependent. The SOURCES-4A code [160] is a code which calculates the neutron production rates and spectra from (α,n) reactions, spontaneous fission and delayed neutron emission due to radionuclide decay from U & Th content. This code can be modified according to the source material, density and mass to give a spectrum for neutron emission. This value can be normalised in order to give a value for neutrons emitted, per second, per centimetre cubed, per x ppb. The value of x is the value measured in the HPGe detector (see Section 5.2) Typically, 10 ppb is used to generate the neutron spectrum which is then normalised to the actual U and Th contaminations. Figure 5.8 shows the result of the SOURCES-4A run in the case of several ZEPLIN–III components.

An integration is performed across the spectrum to find the total number of neutrons/s/cm^3/10 ppb. This number is then the base number and each energy bin in the spectrum is the relevant proportion of this. In the example of the PMT array, the total neutron flux is (in n/s/cm^3/10 pp) 3.28×10^{-9} for Th and 1.16×10^{8} for U.

The numbers quoted above can be converted into a number of neutrons emitted per year for the ZEPLIN–III PMTs. In the case of U, 21.5 neutrons are emitted per year and for Th, 8.7 neutrons are emitted. Of course, this number is purely statistical (0.5 of a neutron is not possible) but these numbers are scaled to millions of neutrons generated in the simulation. If one million neutrons were generated, this would equate to $\sim$33,000 years of livetime. Given all the processes

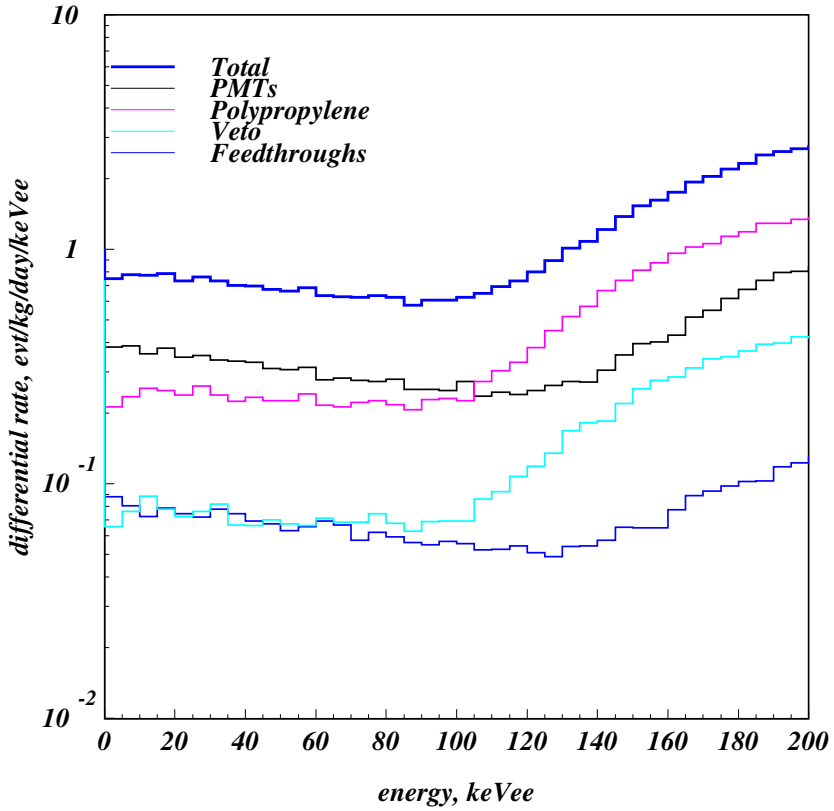

Figure 5.7: Simulated γ-ray background from dominant ZEPLIN–III components. Shown are, the veto (light blue), the ceramic feedthroughs (dark blue), the PMT array (black), the polypropylene shielding (pink) and the total background (bold, dark blue).

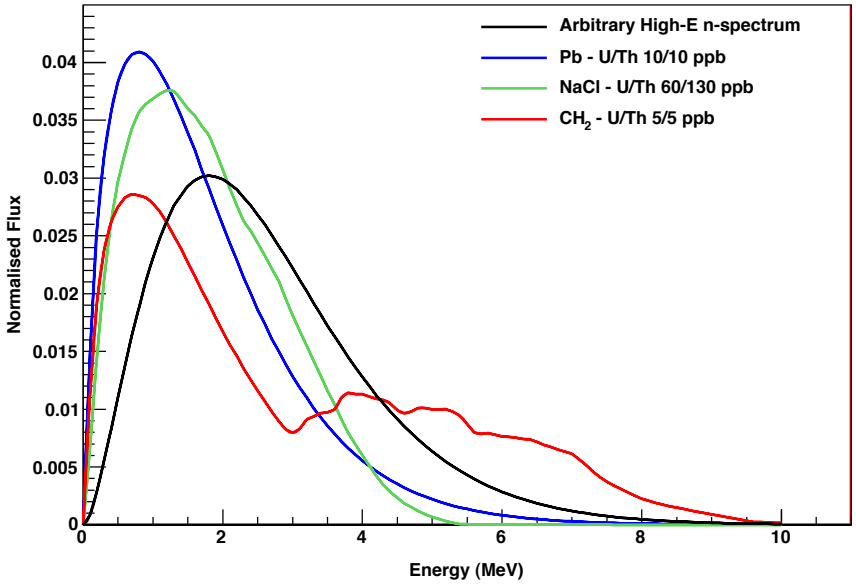

Figure 5.8: Simulated neutron spectra for several sources of external and internal neutrons in ZEPLIN-III. The rate is normalised to the total number of neutrons emitted. Shown are Pb (blue) which makes up the γ-ray shielding in ZEPLIN–III, NaCl (green) which is the main component of the rock surrounding the lab and CH_2 (red) which represents the neutron shielding in ZEPLIN–III.

associated with a single neutron event, it is useful to use a large computing grid to reduce the time taken to reach an answer. Each neutron event takes >5x longer to compute than a single γ-ray event as a neutron is likely to interact many times.

Integrating over energies above 5 keV_{nr}, an expected neutron single scatter rate due to the PMT array of 0.27 events per year is calculated in an 8kg fiducial mass. The energy dependent single scatter rate is shown in Figure 5.9.

Figure 5.10 shows the single scatter rate expected from all the components used in the electron recoil background simulation. In addition to these components, the nuclear recoil rate from the cavern rock is included as this contributes a nuclear recoil background level of the same order of magnitude as the ZEPLIN-III PMTs. The γ-rays associated with the rock surrounding the laboratory are not included in the γ-ray background simulation as they are effectively moderated by the lead shielding and, therefore, only provide a negligible electron recoil background in ZEPLIN–III.

Similar simulations have been performed for all ZEPLIN–III components following exactly the same method. For the second science run, a total single scatter neutron rate of $\sim$0.6 neutrons per year is expected for energies >5 keV_{nr} (with 100% nuclear recoil acceptance). For an assumed veto neutron tagging efficiency of $\sim$60% (see Section 6.3.4 for full calculation), this reduces to $\sim$ 0.2 un-vetoed neutrons per year for energies >5 keV_{nr}, which includes the WIMP search region (defined in Section 4.9).

With the radio-assaying process completed, the PMT array was manufactured ready for use in the second science run. Having reduced the background level of internal components, the sensitivity of ZEPLIN–III can be further improved with the installation of a veto detector, designed to reject signals that are consistent with background. The veto detection system is discussed in the following chapter.

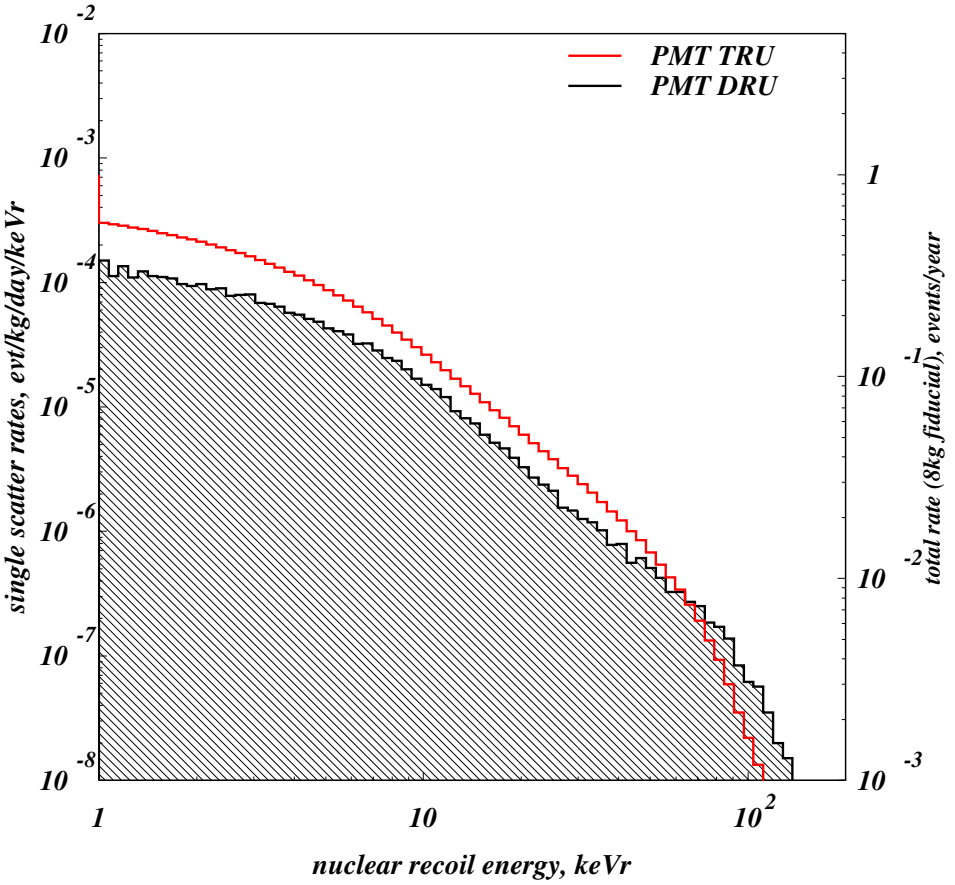

Figure 5.9: The neutron simulation for the ZEPLIN–III PMTs. The black histogram represents the differential rate with its usual units and the red histogram gives the total rate, in tru, integrated across all energies.

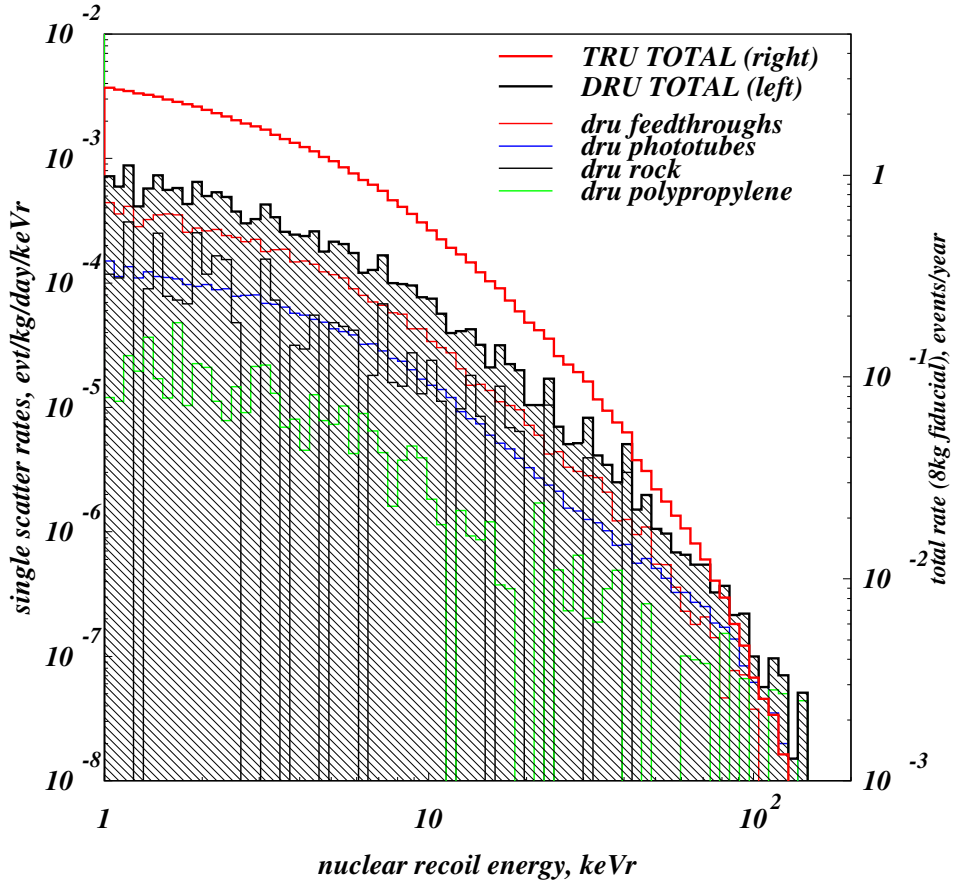

Figure 5.10: The neutron simulation for all ZEPLIN–III components. The individual components are the polypropylene shielding (green), the rock neutrons (black), the PMT neutrons (blue) and the ceramic feedthrough neutrons (red). The shaded black histogram represents the total differential rate and the red histogram gives the total rate.

Chapter 6

The ZEPLIN–III Veto Detector

Following the first science run of ZEPLIN–III, the detector underwent a period of upgrade in order to improve the WIMP-nucleon interaction cross-section sensitivity. Included in this upgrade period was the commissioning and installation of a veto detector designed to tag and reject neutron and γ-ray signals in coincidence with ZEPLIN–III. The ZEPLIN–III veto detector (hereby referred to as, simply, "the veto") was designed, constructed and commissioned by the University of Edinburgh with help from ITEP, Moscow. As such, I led much of the work described here.

The veto is built in a modular structure with slabs of plastic scintillator optically coupled with low background PMTs. These slabs, in turn, rest upon modules of gadolinium loaded polypropylene. The gadolinium is included to aid in the capture of neutrons coming from both the outside and inside of the detector. As well as the tagging and rejection of events in ZEPLIN–III, the veto also acts as a diagnostic tool, identifying and diagnosing spurious events. The background rejection capabilities of the veto are of paramount importance in allowing ZEPLIN–III to reach its second science run (SSR) sensitivity goal of 10^{-8} pb for the spin independent WIMP-nucleon cross-section [161].

6.1 Veto Design and Construction

The veto may be considered in two sections, the barrel section which surrounds ZEPLIN–III vertically and the roof section which sits above ZEPLIN–III. The

barrel section is constructed from 32 modules of gadolinium loaded polypropylene and 32 modules of plastic scintillator. The 32 scintillator modules are a uniform length of 100 cm and provide coverage around ZEPLIN–III. The scintillator module is then seated in its reSPEctive polypropylene base section and clamped using copper retaining plates. The roof section of the veto is constructed from 25 modules that combine to form a circular roof section which rest upon a cylindrical Gd-loaded polypropylene plug. The 1057 kg of scintillator material is polystyrene based UPS- 923A (p-terphenyl 2%, POPOP 0.02%) produced by Amcrys-H, Kharkov, Ukraine. The plastic emits scintillation light at a wavelength of 420nm with a rise time of 0.9 ns and a decay time of 3.3 ns. The complete veto gives $> 3\pi$ steradian coverage of the ZEPLIN–III instrument.

The 32 polypropylene sections and the roof plug of the veto are loaded with gadolinium to a level of 0.5% (by weight) in the form of gadolinium oxide powder. This value was optimised using a Monte-Carlo simulation in which the percentage of gadolinium loading was varied from 0-10%. Figure 6.2 shows that the idealised tagging efficiency (no threshold) rises to a maximum of $\sim$80% for a gadolinium loading of 0.5%.

Natural gadolinium has a neutron cross-section of 49,000 barns (1 barn = 10^{-24} cm^2, which is the highest neutron capture cross-section of any naturally occurring element. Gadolinium itself is a heavy element and ^{152}Gd is a radioisotope which decays via the emission of an α-particle. To reach the stable isotope ^{140}Ce, two further α-particles are emitted. This fact was a cause for concern as there was a potential for these particles to interact with a loosely bound neutron, causing neutron emission. The SOURCES-4A code is used to give an expected neutron emission spectrum for CH_2 (an accurate representation of polypropylene) with a concentration of 10 ppb of ^{152}Gd. An integrated neutron production rate of 1.74×10^{-17} n/s/cm^3 is calculated. Given that the CH_2 is actually loaded with a concentration of 0.5% by weight, the shielding has 0.5×10^7 ppb of naturally occurring gadolinium (actually Gd_2O_3 but the mass fraction of oxygen is included). The isotope ^{152}Gd has a 0.2% abundance in natural gadolinium which means that the concentration of ^{152}Gd in the CH_2 shielding is actually 10^4 ppb. Normalising the 10 ppb neutron production rate to this value gives a neutron production rate of 1.74×10^{-13} n/s/cm^3. This rate is approximately 1000 times less than the rate of neutrons produced by uranium and thorium (1.2×10^{-10} n/s/cm^3) contam-

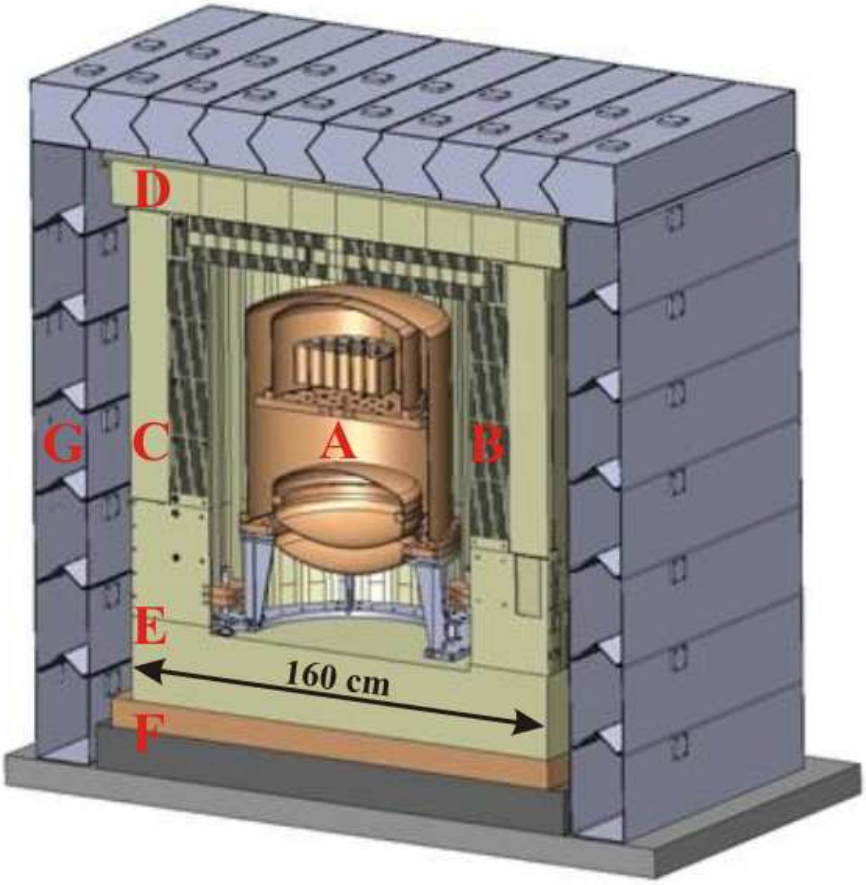

Figure 6.1: A CAD drawing of the veto in situ with ZEPLIN–III and surrounded by the lead castle. In the centre is the ZEPLINIII detector (labelled A) showing the copper vacuum vessel enclosing the Xe target and the LN2 vessel below it. Forming a barrel around ZEPLINIII are the 32 Gd-loaded polypropylene pieces and roof plug (labelled B and hatched). Surrounding these are the active scintillator modules (C) with PMTs housed in cups and recessed into the lower polypropylene structure. The 20 roof modules of scintillator (D) rest on the roof plug. The lower polypropylene structure (E) contains no Gd and rests on a copper and a lead base (F). Finally, the lead castle (G) envelopes the assembly [162].

inants in the CH_2 shielding. This means the effect of the α-particle emitting ^{152}Gd isotope is negligible compared to the effect of the heavy metal contamination of the CH_2 shielding itself.

When neutrons enter the polypropylene shielding, they are moderated via the process of hydrogen scattering. The energy loss per interaction is maximised due to neutrons and protons having almost identical masses. Once they have reached thermal energies, it is possible for them to undergo radiative capture. The isotope upon which this most often occurs is ^{157}Gd (15.7% natural abundance). This is due to the fact that the ground state of ^{158}Gd is the most stable isotope of gadolinium. The capture of a neutron onto ^{157}Gd to form ^{158}Gd is energetically favourable and requires minimal angular momentum transfer. In effect, the wave-function for ^{157}Gd + n is almost identical to the wave-function for ^{158}Gd.

Although natural gadolinium has a neutron capture cross-section of 49,000 barns, ^{157}Gd alone has a neutron capture cross section of 225,000 barns for thermal neutrons. This is around 66 times as large as the capture cross section of the ^{10}B(n,α)^{7}Li reaction. Boron loading of liquid scintillator has been proposed as an alternative method by which to capture thermal neutrons in dark matter experiments [164].

The neutron capture reaction can be described as follows:

$$^{157}\mathrm{Gd} + \mathrm{n} \rightarrow ^{158}\mathrm{Gd} + \gamma + 7.94\ \mathrm{MeV}. \tag{6.1}$$

The excited state of ^{158}Gd typically decays via the emission of 3-4 γ-rays with a total energy equal to that of the Q-value, 7.94 MeV. The experimental spectrum for this decay is shown in Figure 6.3. The γ-ray emission has a characteristic time of (10.4$\pm$0.3) μs (see Figure 6.4) from the time the neutron enters the loaded material. This characteristic time is dependent on the geometry of ZEPLIN–III. In a purely hydrocarbon system, the mean emission time would be much longer and would lead to the release of a single 2.2 MeV γ-ray following radiative capture by H.

The gadolinium loading is achieved through the machining of grooves into the polypropylene sections and the backfilling of the sections using a mixture of gadolinium oxide and epoxy resin. Simulations showed that for neutron tagging, provided the pitch of the grooves did not exceed $\sim$15 mm, the difference between

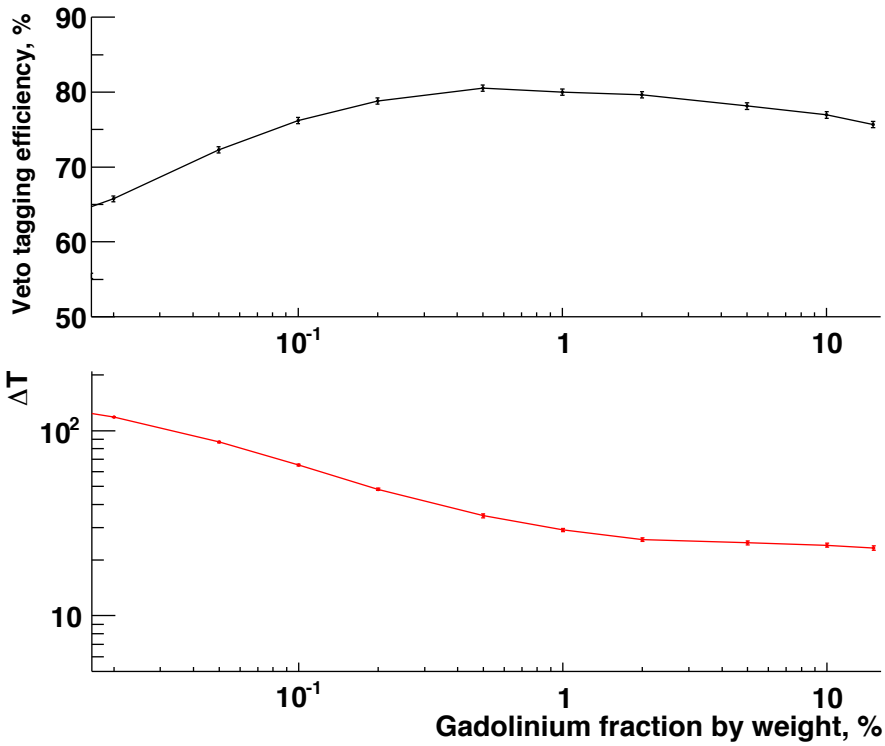

Figure 6.2: The veto tagging efficiency and mean γ-ray emission times for various gadolinium loading fractions. Data is taken from an idealised simulation with no energy threshold [163].

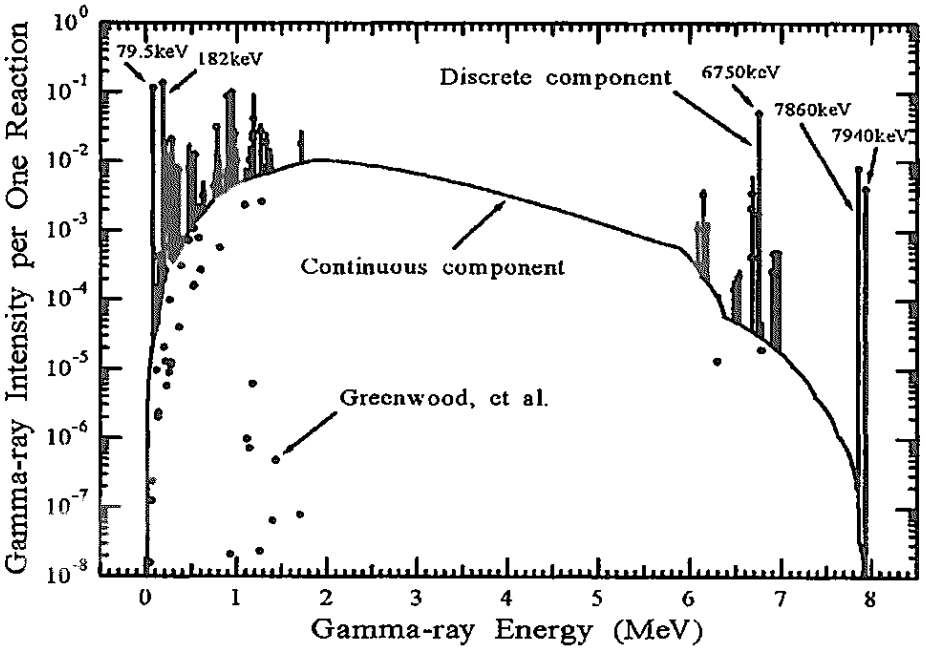

Figure 6.3: The energy spectrum of prompt γ-ray emission from ^{157}Gd(n,γ)^{158}Gd [165].

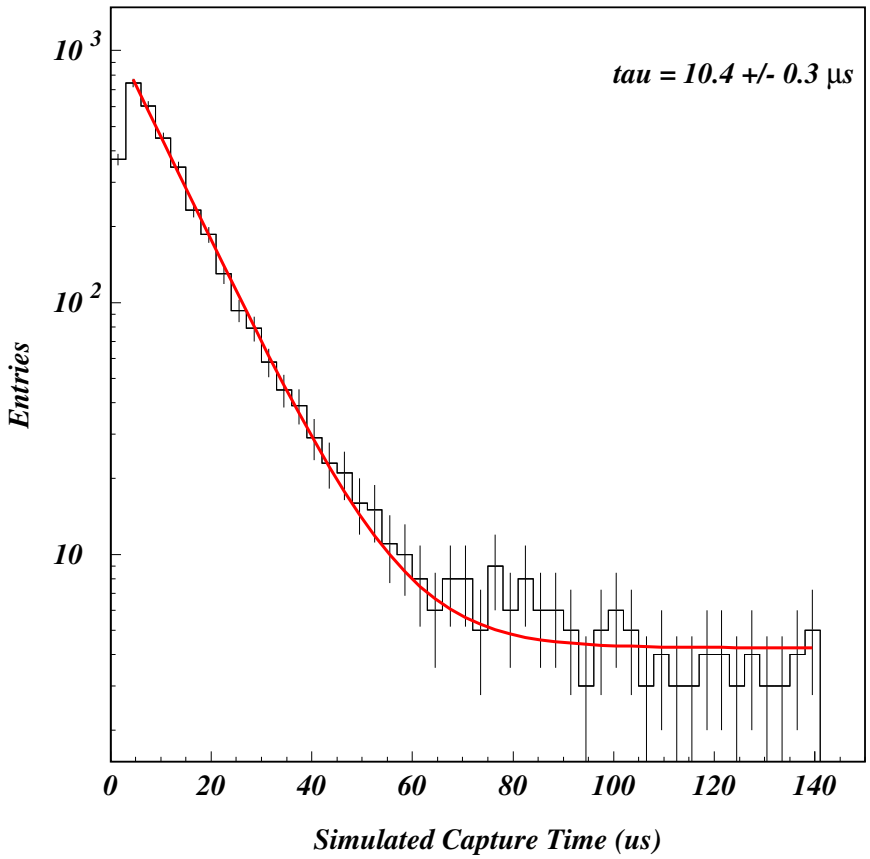

Figure 6.4: The time between a nuclear recoil signal in ZEPLIN-III and delayed γ-ray signal in the veto. A characteristic capture time of 10.2 μs is observed.

grooves and a uniform distribution of gadolinium is negligible. Tests show that the gadolinium oxide powder remains in a uniform suSPEnsion during the period of time taken for the epoxy to dry meaning that neutron capture is equally efficient across the whole volume [163].

The veto scintillator slabs are optically coupled to individual 3-inch ETEL 9302KB PMTs. These low background PMTs were selected in order to reduce the radioactive impact in both ZEPLIN–III and the veto. Given the fact that the active Xe volume of ZEPLIN–III is situated towards the top of the detector, the scintillator slabs were orientated such that the PMTs were as far from the Xe volume as possible.

The PMTs are connected to pre-amplifiers developed at the University of Edinburgh. These pre-amplifiers serve two purposes. The first, and most important of these, is to match the impedance of the PMT output signal to the input impedance of the data acquisition system (DAq) analog-to-digital converter (ADC) cards. As the PMTs were bought "off the shelf", they have an output impedance of 1MΩ whereas the ADCs require an input impedance of 50Ω. The pre-amplifiers also provide a factor 10 gain increase which allows better resolution of low energy events but limits the full energy scale. Details of the veto energy range can be found in Section 6.2.

Also optically coupled to each of the veto scintillator slabs is a single fibre-optic cable. This is installed so that, once a week, a data run can be performed whilst flashing an LED of consistent voltage which serves to monitor the light attenuation length stability of the plastic scintillator. This is discussed in more detail in Section 6.3.

Finally, once the veto is constructed around ZEPLIN–III, the 20 cm lead shielding can be built up. The veto and lead shielding is designed so that ZEPLIN–III and all its ancillary systems (gas system, levelling and electronics) can be accommodated. To serve this purpose, slots are machined into the base of the veto which act as feedthroughs for pipes and wires. Any unused slots are filled with wedges that sit flush within the groove and any used slots have machined wedges inserted that allow the passage of pipes and wires out of the detector but minimise the loss of shielding coverage. The lead roof of the castle is at a height such that the curvature of the fibre-optics feeding into the roof slabs is not so great as to impede transmission of light or, indeed, fracture the fibre-optic cable.

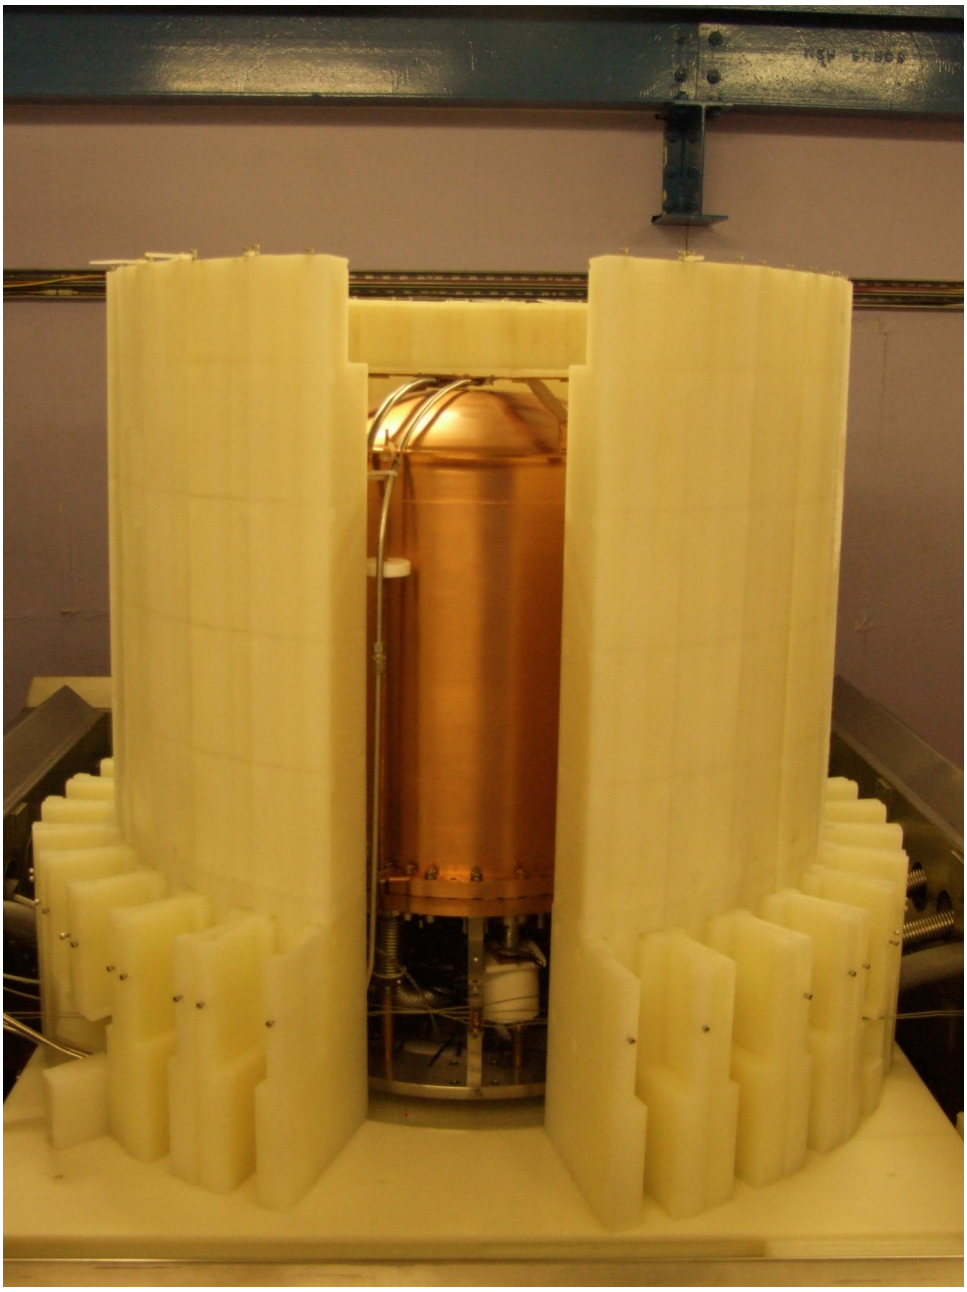

Figure 6.5: The polypropylene veto barrel pictured surrounding ZEPLIN–III. The recesses in which the veto PMTs rest can be seen toward the bottom of the barrel sections.

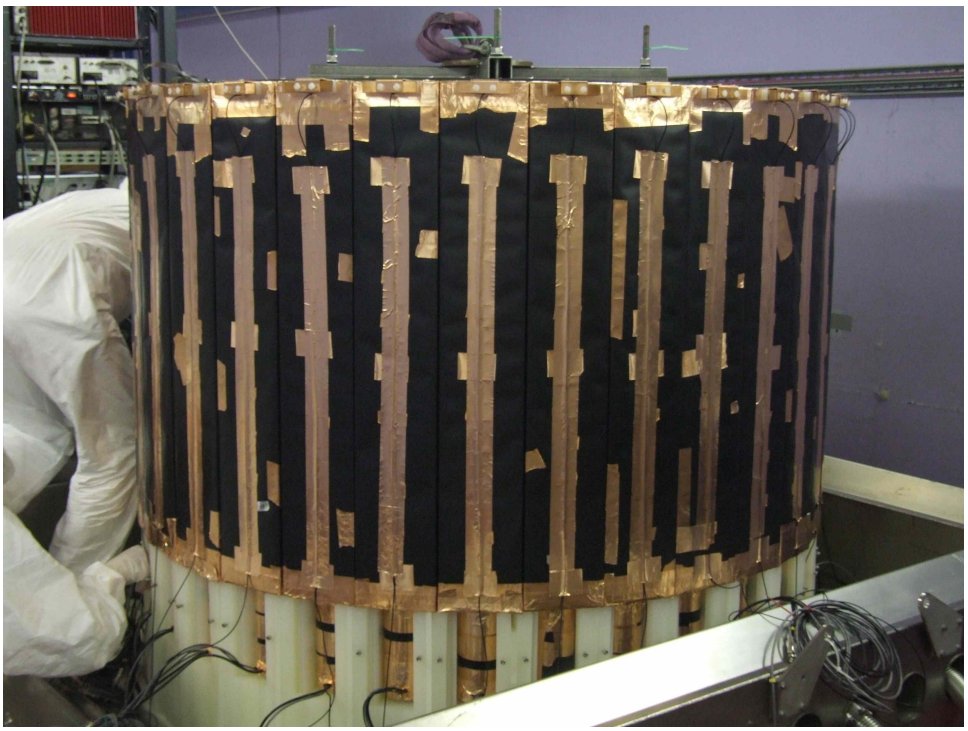

Figure 6.6: The scintillator sections installed into the polypropylene. The four cables: high voltage, signal, pre-amp and LED fibre-optic can be seen.

Figure 6.7: The lead castle built up to the level of the veto roof, the lead roof is built onto this leaving enough room for the LED fibre-optic cables which can be seen attached to the lower part of the roof scintillator slabs.

6.2 Veto Calibration

The calibration of the veto was performed both inside and outside the lead castle. The programme of calibration began with the testing of individual PMTs in order to measure their individual quantum efficiencies and the gain dependence of the single photoelectron signal (SPE). Subsequently, tests were run to determine the full energy range of the slabs at given PMT gain levels and a final operational voltage was set.

6.2.1 PMT Quantum Efficiencies

The quantum efficiency of a PMT is a probability function defined as the ratio of input light photons to output electrons from the PMT photocathode. This quantum efficiency is dependent on an array of factors but the factor on which it depends the most is the wavelength of the incident light. The quantum efficiency of a PMT is always quoted as a percentage efficiency at a SPEcific wavelength. For ETEL type 9302KB PMTs, the quantum efficiency for a typical unit is (as quoted by the manufacturer) 30.1% at 360nm [108]. As mentioned previously, the veto emits light at 420 nm so it is important to know the quantum efficiency at this wavelength.

Three of the PMTs were sent to ETEL for individual quantum efficiency testing. The quoted values are the averages across the entire photocathode surface (typically, the efficiency varies with position on the photocathode). As the scintillator plastic emits light at 420nm, it was requested that ETEL send values for quantum efficiency at this wavelength. The three values given were 26.52%, 25.46% and 26.37%.

In order to calculate the quantum efficiencies of all PMTs, a dedicated LED run was performed for each with a blue LED illuminating the photocathode. The run lasted 100s and the LED and PMTs were held at fixed distances from each other. The gains of the PMTs were equalised across the runs using the position of both the SPE and LED peaks, and it is the integrated area of the LED peak that gives the quantum efficiency. All measurements were normalised to the absolute quantum efficiency of the 3 PMTs measured by ETEL.

This method can be applied to all PMTs and an average value (due to three separate reference points) can be calculated. The result is plotted in Figure 6.11.

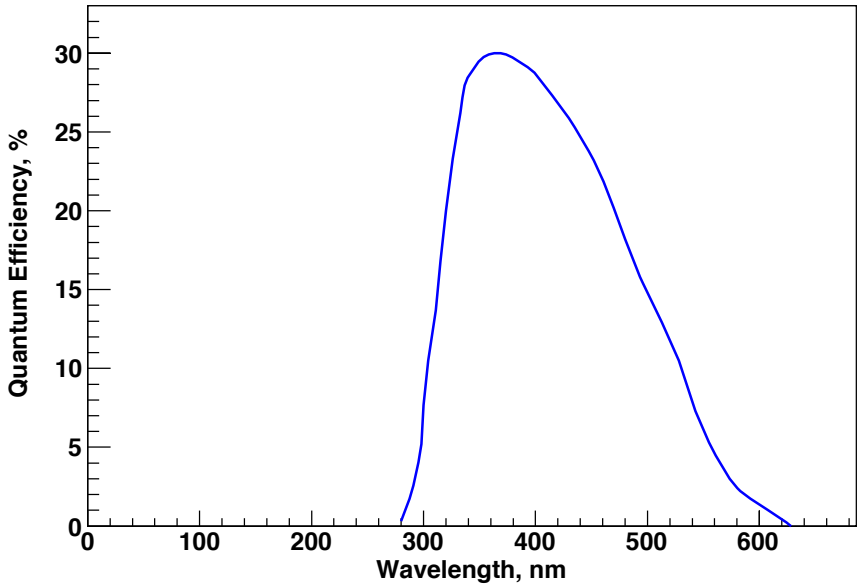

Figure 6.8: The wavelength dependent quantum efficiency for a typical ETEL type 9302KB PMT. The peak quantum efficiency is measured to be 30.1% at 360nm [108].

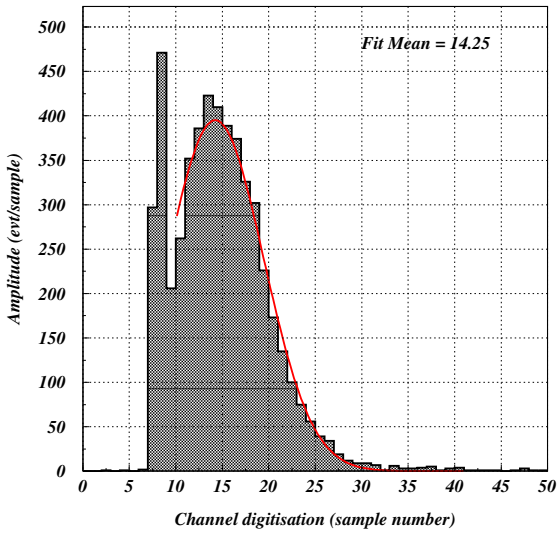

Figure 6.9: The SPE signal from one of the veto PMTs. The SPE feature peaks at a acquisition channel number of 14.

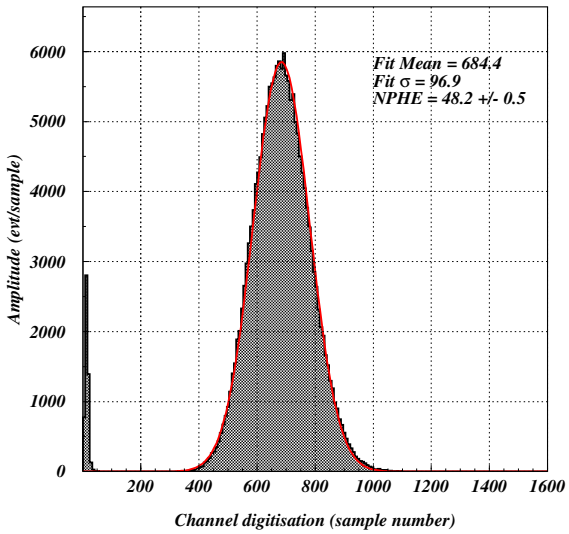

Figure 6.10: The LED signal in the same PMT as used in Figure 6.9. Given the LED feature peaks at acquisition channel 684. Setting a threshold of 10 phe is equivalent to an acquisition channel of 140, which is well removed from the SPE feature and means only LED signal is counted.

The average quantum efficiency for the PMTs used in the veto is 26.2% at 420nm.

6.2.2 Setting an Energy Scale

Once the quantum efficiency of each PMT has been determined, it is matched to one of the veto slabs. The matching of PMT to slab was decided based on the gains and quantum efficiency of each PMT and the bulk attenuation length (BAL) of each slab. The attenuation length is defined as the length of plastic needed for a beam of light to reach $1/e$ of its original intensity. A single PMT was selected for the tests and a ^{22}Na source was also used. ^{22}Na produces characteristic back to back 511 keV γ-rays due to positron emission and subsequent electron-positron annihilation, as well as 1275 keV γ-rays due to the positron emission leaving an excited state.

Energy SPEctra were acquired for varying positions along the length of the slab and the position dependent size of the 511 keV signal was mapped. If this data is fitted with an exponential function of the form $e^{x/l}$, the value of l gives the characteristic length for a photon of light to reach $1/e$ of it's original size (the BAL). Prior to calibrating each slab, the gains of each PMT were adjusted such that the SPE signal from each PMT fell at the same pulse height. This is done to normalise the response of the PMTs to SPE level signals. The energy range in terms of the number of photoelectrons can, in theory, be set trivially. This is accomplished by calculating the position of the SPE peak in the data in terms of a DAq ADC channel. A maximum number of photoelectrons can then be calculated given that the ADC cards have a full scale of 2^{14} channels (14 bit ADC cards).

In practice, it is not so simple. The shape of an acquired veto pulse is such that the pulse height does not scale linearly with number of photoelectrons. Using the acquired SPEctra, it is possible to work out the number of photoelectrons that the mean of the LED distribution represents by using the following relationship:

$$n_{\mathrm{phe}} = \left(\frac{\mu}{\sigma}\right)^2, \tag{6.2}$$

where μ is the mean of the LED distribution and σ its width.

For the range of LED voltages, a plot of pulse height against calculated number of photoelectrons can be produced. A functional fit to this provides the correction

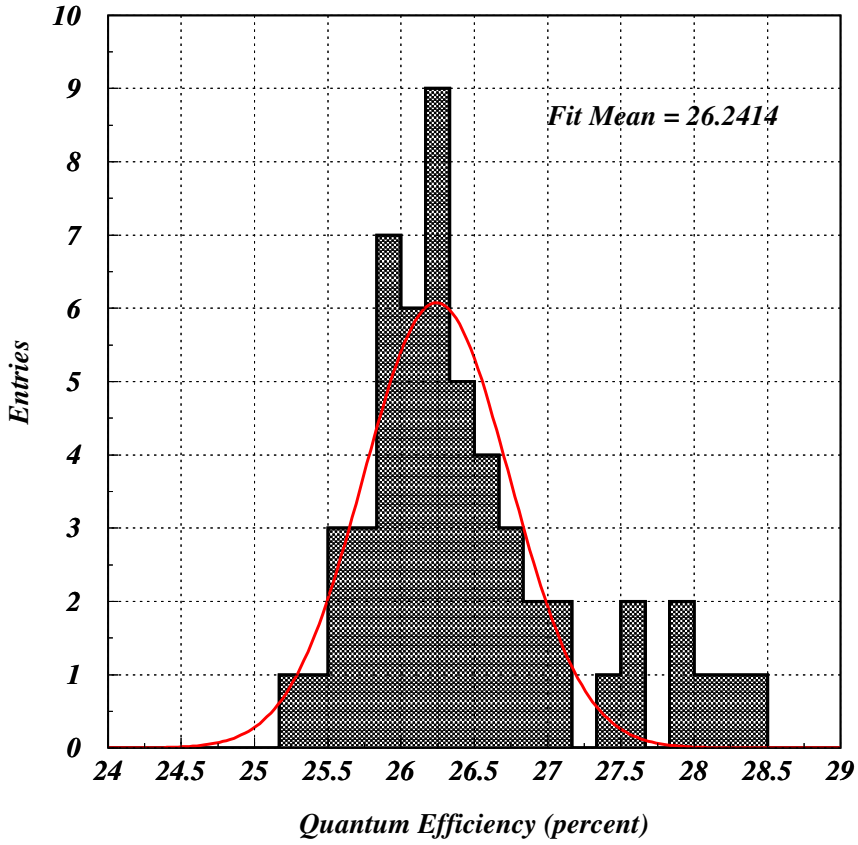

Figure 6.11: The measured quantum efficiency for all veto PMTs deduced using the values measured by ETEL. The fit shows an average value (from a gaussian fit) of 26.24% at 420nm.

to the data needed to remove the pulse height non-linearity. The veto DAq has a voltage full scale of 2.25 V. Two volts of this full scale is dedicated to positive going signals while the remaining 0.25 V is used to record pulse undershoot (large pulses often undershoot the baseline during their decay). This means that, in terms of ADC channels, the maximum pulse height achievable is $2/2.25 \times 2^{14}$ which is 14,564. If this maximum channel number is inserted into the correction function, a maximum photoelectron level of 85.1 is calculated.

The exact full-scale can be verified by the detailed simulation of the experimental setup. Using a simulation of a single veto slab with a ^{137}Cs source placed at the end, it is possible to estimate a spectrum for the expected signal in a veto slab. The simulation takes into account all the physical characteristics of the veto slab, including photon generation and propagation, and also the quantum efficiency of the PMT attached. Once a simulated spectrum has been produced, it can be compared with the (corrected) spectrum from an experimental run. This source data allows an energy range in terms of electron volts to be set for each slab. The ^{137}Cs spectrum peaks at 32 photoelectrons, this peak corresponds to the Compton edge of the ^{137}Cs spectrum at $\sim$400 keV and is associated with the β-decay of ^{137}Cs into a metastable state of ^{137}Ba. Given a full scale of 85 photoelectrons, it is possible to say that our energy full scale for this slab is 1.06 MeV. However, this energy range is only SPEcific to an energy deposition at the far end of the slab. As the source is moved closer to the PMT, the number of photoelectrons seen by the PMT increases for any SPEcific energy so the full energy range decreases. The absolute decrease is proportional to the BAL of the module. Although an 85 photoelectron full range is possible, the veto must be sensitive to both low energy γ-rays and high energy neutrons/muons. In order to achieve a useful sensitivity at low energies, the veto slabs were set such that a full range of 65 photoelectrons was available.

6.3 Veto Science Data Analysis

6.3.1 Data Synchronisation

Data from ZEPLIN–III and the veto are recorded using independent DAq systems. In order to allow the efficient vetoing of unwanted signals, it is essential

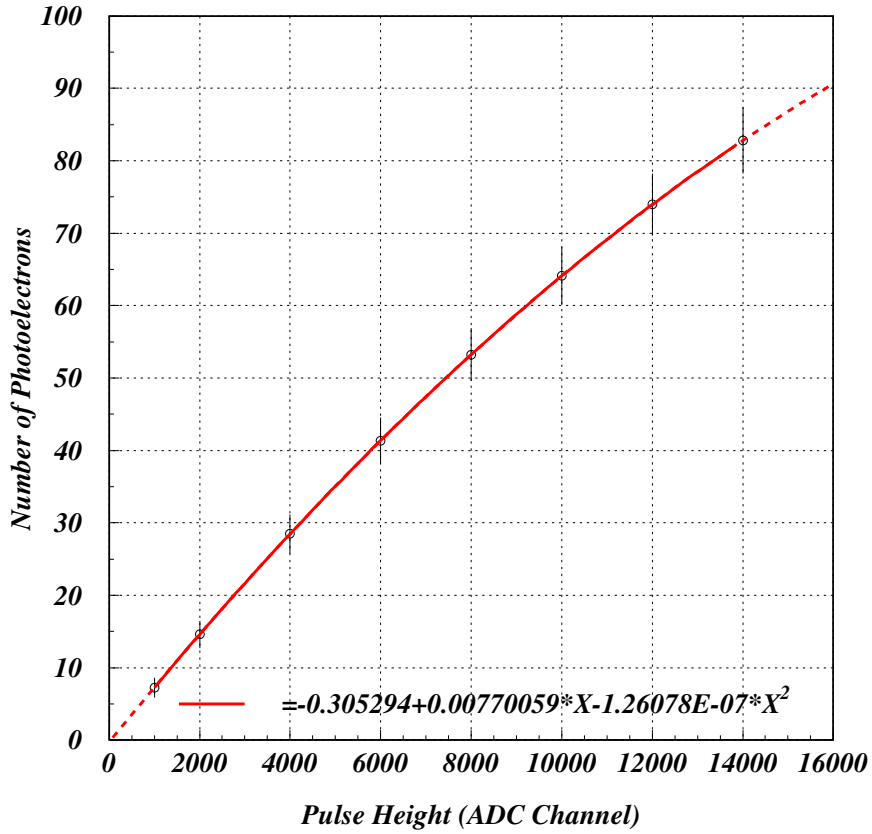

Figure 6.12: The non-linearity in veto pulse height due to the sampling rate of the veto DAq. The fitted polynomial can be applied to data in order to correct the pulse height.

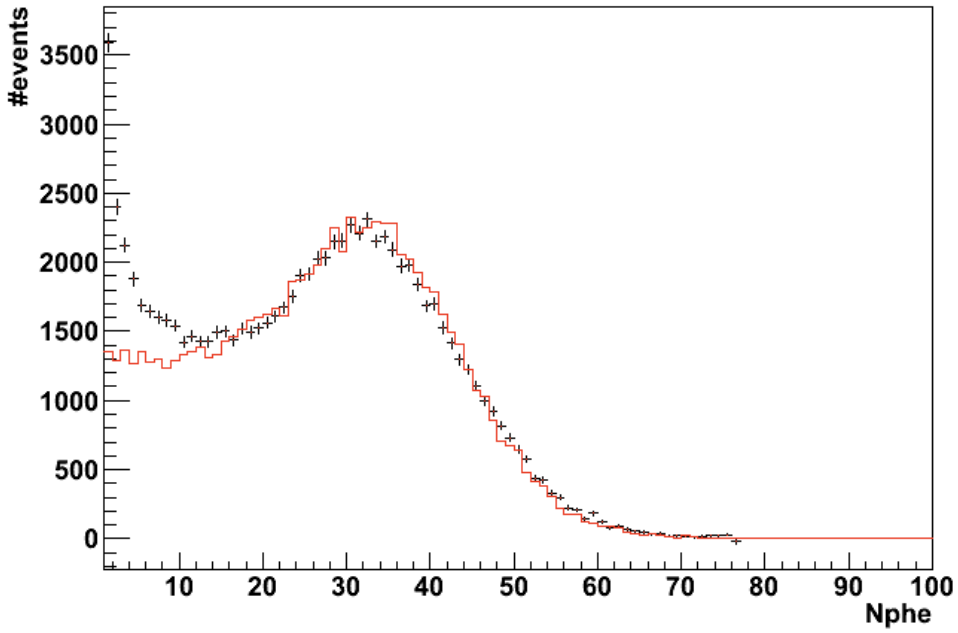

Figure 6.13: A comparison between data and simulation of a ^{137}Cs. The ∼400 keV Compton edge associated with the decay of ^{137}Cs is seen at 32 phe. The position of this peak sets the energy scale. A measurement of ^{60}Co data shows a linear relationship between γ-ray energy and number of photo electrons detected. The increase in rate below $\sim$ 10 photoelectrons can be attributed to a combination of the intrinsic SPE background of the PMT and an increased SPE rate induced by the presence of a source and afterpulsing.

that the two data streams remain synchronised. In the case of ZEPLIN–III this is achieved using a dedicated synchronisation unit developed at the University of Edinburgh.

The unit contains a 1 MHz clock and 32-bit counter that sends an identical digital stamp to both acquisition systems following a trigger from ZEPLIN–III or the veto detector. The 1 MHz generator allows a 1 MHz counting rate. This generator signal feeds into two 32-bit counters, one for each of the DAq systems. With a rate of 1 MHz, $\sim$70 minutes passes before the two counters reset to zero. The 32-bit output of the two counters is read out by two ribbon cables, one of which connects directly to the veto DAq and the other of which is fed into the ZEPLIN–III DAq system. A typical example of an event synchronised using the unit timestamp is shown in Figure 6.14.

Independently of the synchronisation unit, a number of further methods ensuring synchronisation are implemented in the analysis of matching ZEPLIN–III and veto detector events. All the ZEPLIN–III computers, including those used by the data acquisitions systems, are synchronised to the same time server. As such, since the average trigger rate is only 0.4 Hz, the first synchronisation check simply ensures that the absolute time of an event given by the data acquisition clocks agree to within 4 ms of the resolution of the time server.

The time server synchronisation is subject to slow drifting so, rather than rely solely on this, the time difference between the two DAq systems (taken from the computer's clock) is monitored and this may be no more than 1 ms different to the preceding or subsequent event. Figure 6.15 shows this time difference over a period of several days. The synchronisation of the data is performed off-line (after the data has been acquired) so comparison with events that occur later in time is trivial.

Finally, pulse parameters from identical timelines fed into both ZEPLIN–III and the veto detector must agree. This check takes advantage of the summed signal from all ZEPLIN–III PMTs that is shaped and used to trigger the acquisition systems. This timeline is recorded in both the ZEPLIN–III and the veto detector data acquisitions allowing an event-by-event direct comparison of unique parameters such as the number of pulses recorded, their start times, and their amplitudes. An example of a synchronised event is shown in Figure 6.16. The parameters of the two trigger signals are used for data matching.

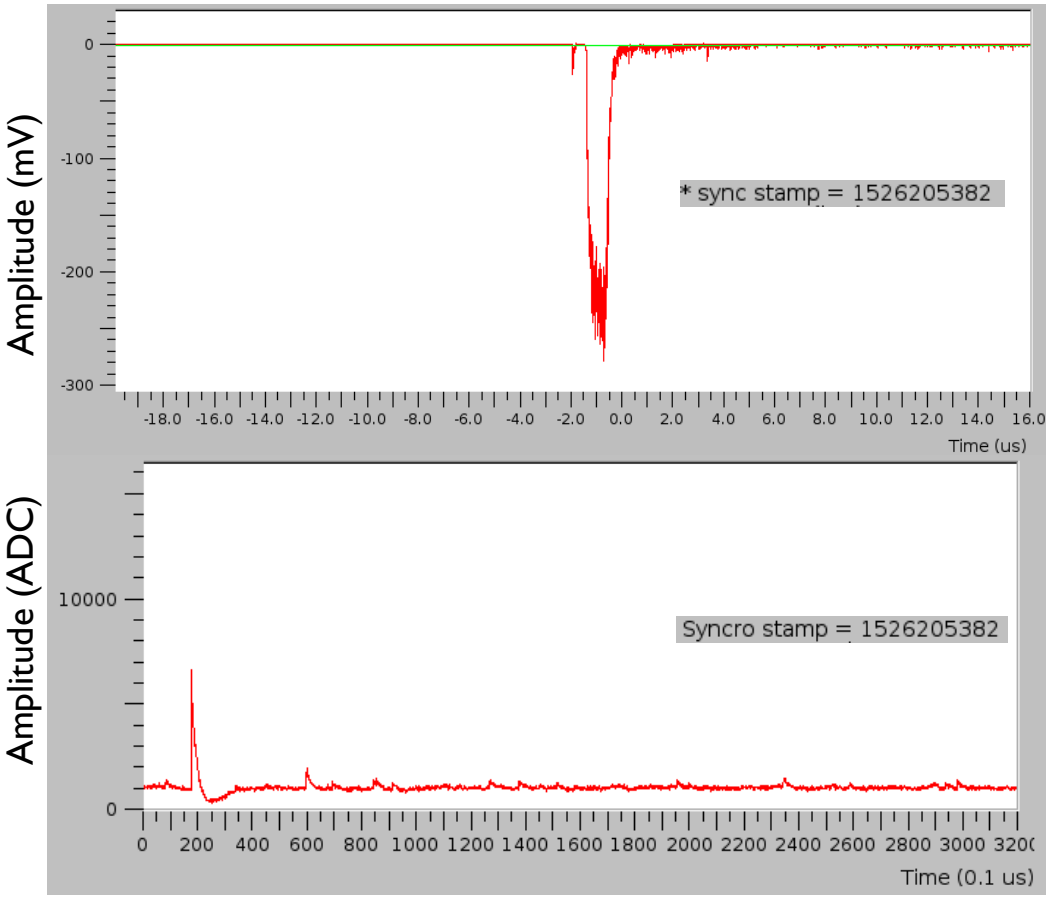

Figure 6.14: A comparison of timelines in ZEPLIN–III (top) and the veto (bottom). This shows a typical electron recoil event with a prompt coincidence signal in the veto (after timeline offset correction). In this event, the two synchronisation values are identical.

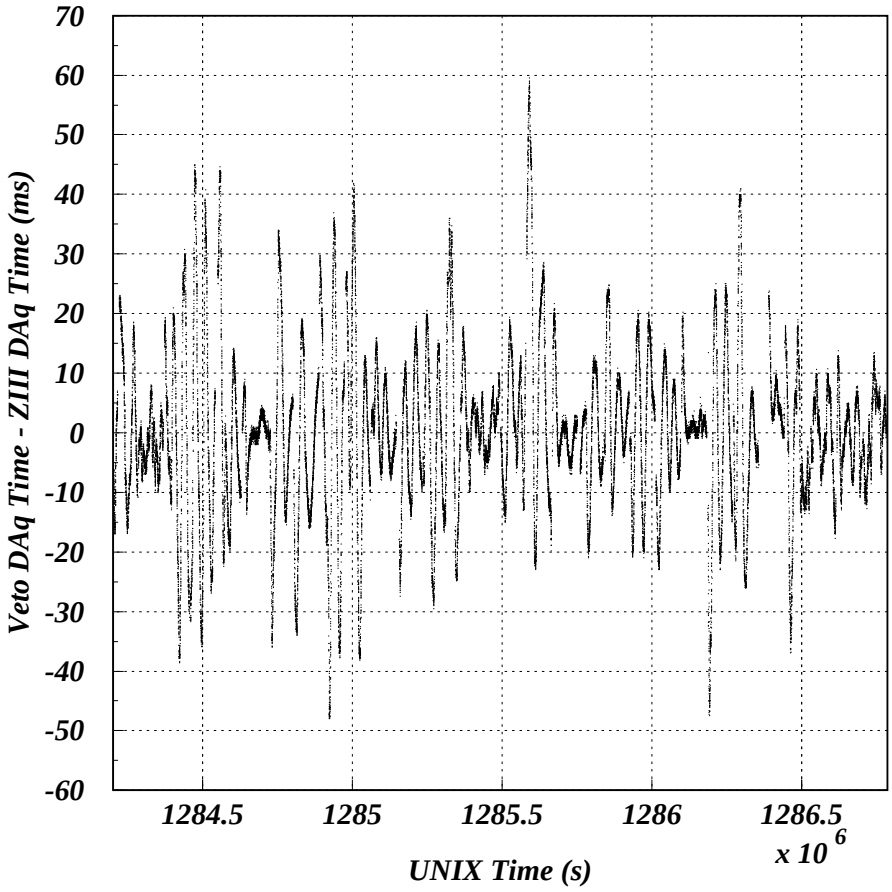

Figure 6.15: The difference between the ZEPLIN-III and veto DAq timestamps for data post 100801. An event whose time difference is consistent with surrounding events is registered as synchronised.

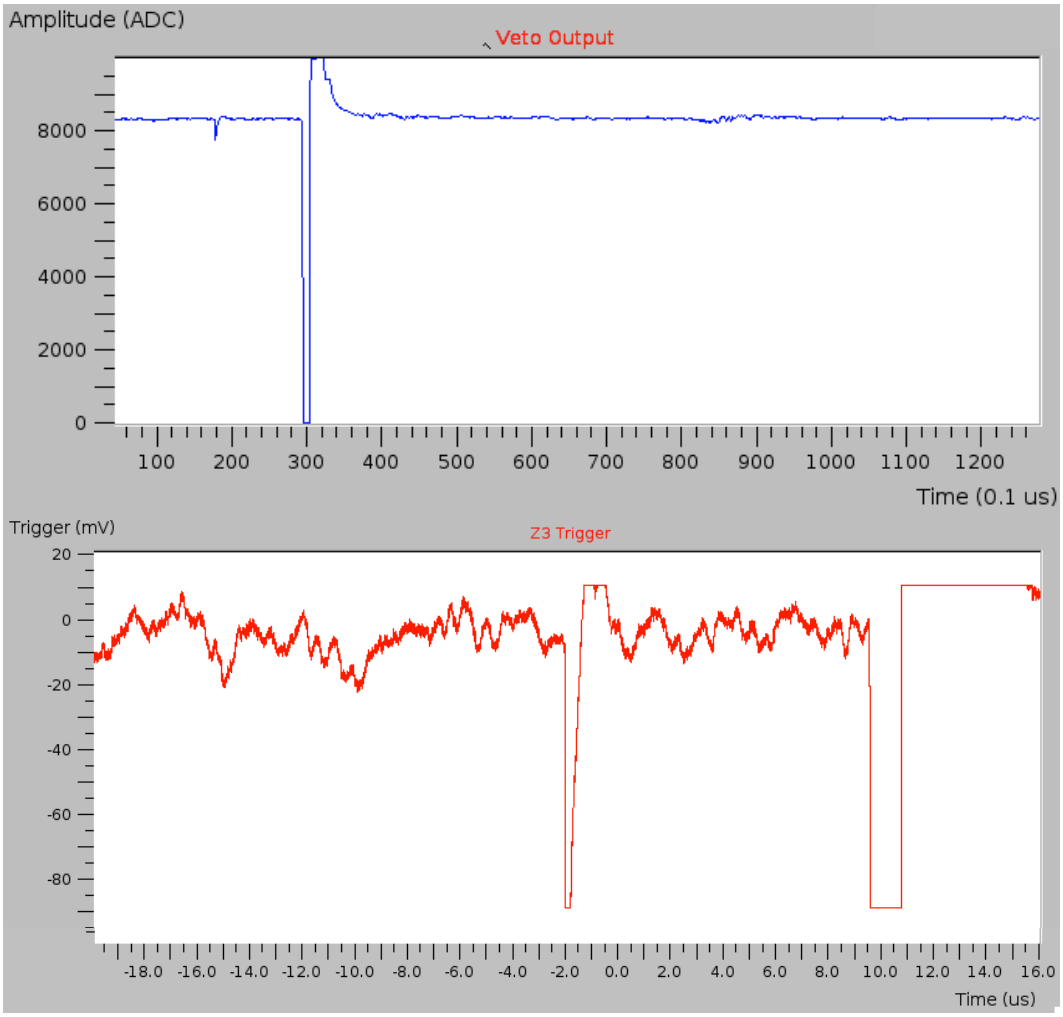

Figure 6.16: Veto (top) and ZEPLIN–III (bottom) events synchronised using ZEPLIN–III sum output. A time difference of 12 μs can be seen between each of the two pulses in the timeline.

More than 97% of all ZEPLIN–III triggers from the analysed data have successfully matched veto system timelines. The less than 3% of ZEPLIN–III events that do not have synchronised veto timelines are of course not removed from further analysis, but simply have no veto capability available to them, which effects the final tagging efficiencies. Figure 6.17 shows the synchronisation of the ZEPLIN-III data over the duration of the Second Science Run (SSR). There are days where the tagging efficiency is reduced but these can be identified as days where extended calibration data were taken or where operating parameters moved away from their acceptable limits (*e.g.* increases in event rate or detector temperature/pressure fluctuations). Figure 6.18 shows the proportion of events which were synchronised as a function of time through the SSR.

The default operation of the veto detector has been in a combination of "slave' and "master' modes. In slave mode, timelines from all modules are recorded when an external trigger sourced from the ZEPLIN–III instrument is received. As such, the veto detector is guaranteed to record data when ZEPLIN–III itself has recorded an event, obviating the need for a separate hardware-enforced trigger and associated efficiency loss at low energies. The additional diagnostic benefit of this is that the veto detector provides an independent and continuous measurement of the event rate in ZEPLIN–III, without probing blinded data. In master mode, when certain conditions are met, the veto system triggers independently of ZEPLIN–III. In the present configuration this happens when three or more veto module PMTs register at least 10 photoelectrons in the same event. The threshold has been set so as to record an event rate of approximately 2 Hz from self triggers, in keeping with data volume constraints (the veto detector generates $\sim$20% the data volume of the ZEPLIN–III instrument). This allows for an independent measure of background from the detector surroundings, eSPEcially ambient neutrons from the laboratory rock. A multiplicity (*i.e.*, the number of veto modules in which a pulse is simultaneously seen) of three modules with a 10 photoelectron threshold is preferable over a multiplicity of two modules with the higher threshold of 32 photoelectrons (maintaining a 2 Hz event rate) since the efficiency for detection of ambient neutrons is almost 10% greater in the former arrangement.

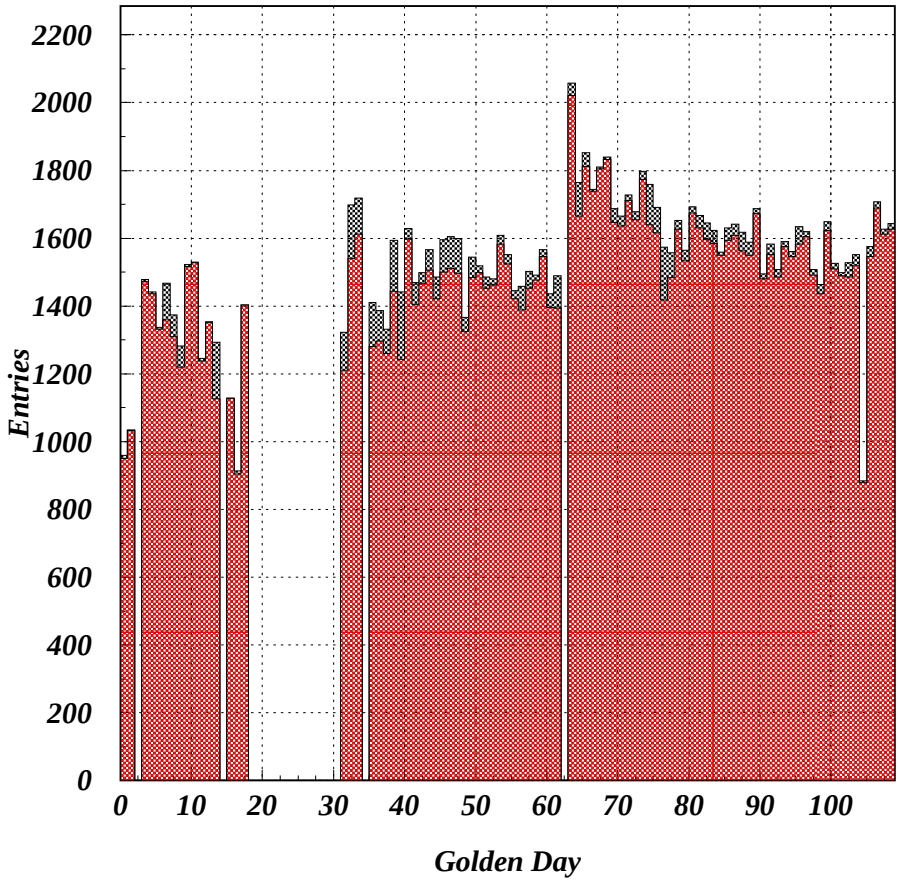

Figure 6.17: A comparison between total events (black) and synched events (red) for the entire golden dataset up to 101010.

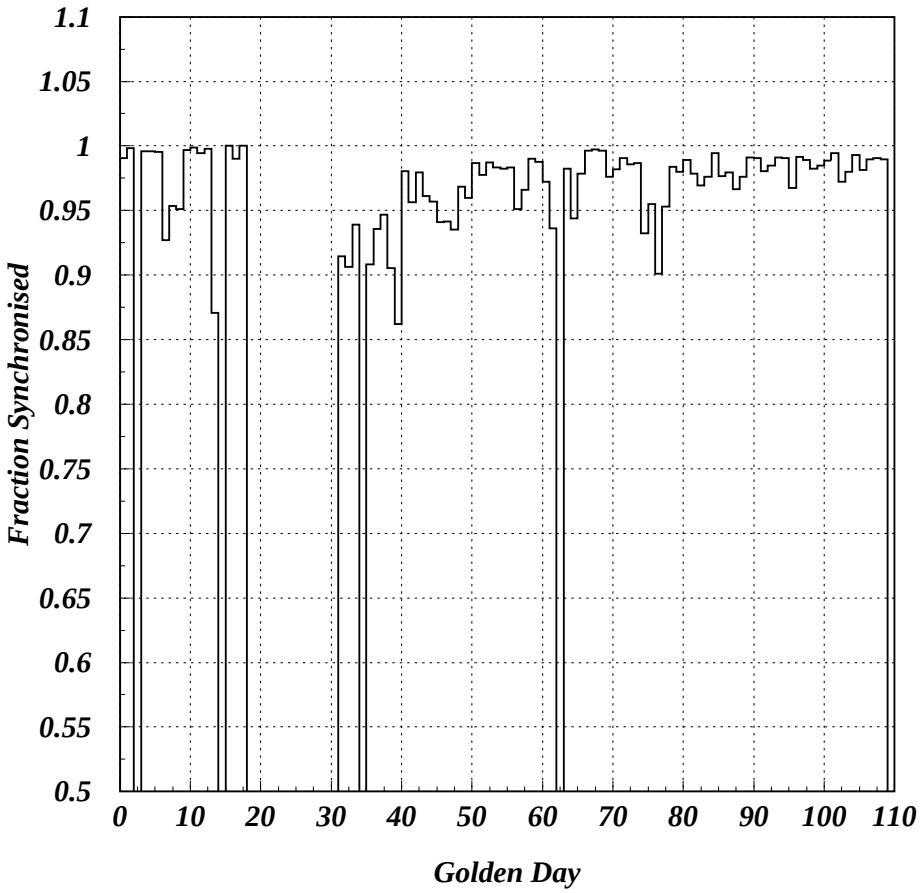

Figure 6.18: The synchronisation expressed as a fraction for the entire golden dataset up to 101010.

6.3.2 Detector Rate

It is possible to determine the intrinsic rate of the veto independently of ZEPLIN–III. Figure 6.19 shows the cumulative background rate spectrum in the veto, highlighting several features of interest. The spectrum extends to approximately 65 photoelectrons (equivalent to several MeV of energy deposited across the array). The fall-off at this point is due to pulse height saturation, determined by the range of the ADCs and the gain settings of the PMTs.

The first feature of interest is at low photoelectron values. This distribution can be attributed to the single photo-electron rate, intrinsic to each PMT. This feature is due to the thermal emission of single electrons from the PMT photocathode which then are accelerated and multiplied across the dynode chain, leaving a characteristic peak at a calibrated value of 1 photoelectron.

The second feature of the spectrum in Figure 6.19 is an exponentially falling background due to radiological contamination within the PMTs. This feature is related to the decay of ^{40}K and dominates for values between 2-15 photoelectrons. The ^{40}K decay has an 89% β-decay branching ratio with an end-point energy of 1.31 MeV. For a refractive index of 1.49 at 400 nm, β radiation emitted from the potassium generators in the PMT behind the photocathode will produce Cherenkov photons in the window when the electron energy exceeds 178 keV. The potassium generators are a byproduct of the photocathode evaporation process. The potassium generators are installed so that the glass envelope of the PMT may be produced as a single piece to reduce the risk of vacuum loss in the PMT. Once the envelope has been produced, a current is applied across the potassium generators which causes the evaporation of the potassium. The rate due to Cherenkov light is supplemented by Bremsstrahlung radiation in the window and, to a lesser extent due to β radiation with energies in the tail of the distribution that are able to escape the window and interact in the veto scintillator. The slope and magnitude of this feature are entirely consistent with expectation [166].

Finally, the third feature of the spectrum is the exponential distribution of pulses at higher energies. This component is due to γ-ray emission from U and Th decay chains and γ-rays associated with the decay of ^{40}K, and are of particular importance for ZEPLIN–III in terms of background. The rate of γ-ray emission includes contributions from all of the veto components such as the plastic scintilla-

tor and the PMTs. The spectrum shows a rate of ∼80 Hz with no photoelectron threshold. This is consistent with the simulated expectation presented in Ref. [162] providing evidence that the radio-contamination values calculated in the radio-assay process (see Section 5) are correct. The veto itself contributes negligible background to ZEPLIN–III at a rate of 0.03 single scatter neutron events per year in the WIMP search region.

The background signal rate is important as it can be used to define the probability of an event in ZEPLIN-III being accidentally vetoed. This accidental tagging rate is determined individually for both prompt γ-ray detection (associated with a γ-ray induced electron recoil in ZEPLIN–III) and delayed γ-ray detection (associated with γ-ray emission following the radiative capture of a neutron). The accidental tagging rate depends on three factors: the energy threshold for acceptance, the pulse timing and slab multiplicity.

6.3.3 γ-ray Tagging

Following a γ-ray Compton scatter in the Xe target, the electron recoil energy deposition may cause both the ZEPLIN–III and veto DAq systems to trigger (see Section 6.3.1). Following the interaction in the Xe, the scattered γ-ray may exit ZEPLIN–III and interact in the veto scintillator, producing a prompt coincident signal. Parameters for the identification of these prompt γ-ray signals must be determined. This includes determining the optimised prompt coincidence window and the setting of an appropriate veto detector threshold. The identification parameters must optimise the tagging efficiency whilst minimising the rate of accidental coincidences. Once identified, a prompt γ-ray tagging efficiency can be determined.

In order to determine a prompt acceptance window, it is necessary to determine the time difference between ZEPLIN-III and veto timelines. Prompt signals will be coincident with the S1 signal in ZEPLIN-III. Figure 6.20 shows the distribution of S1 pulse times, peaking at -2.02 μs with a width, σ =0.01 μs. The same process is used to determine the time of coincident prompt pulses in the veto. Figure 6.21 shows a histogram of veto pulse start times. The peak at 17.6 μs represents pulses coincident with S1 triggered events in ZEPLIN–III. The peak has a FWHM of 0.23 μs which provides a minimum acceptance window for prompt

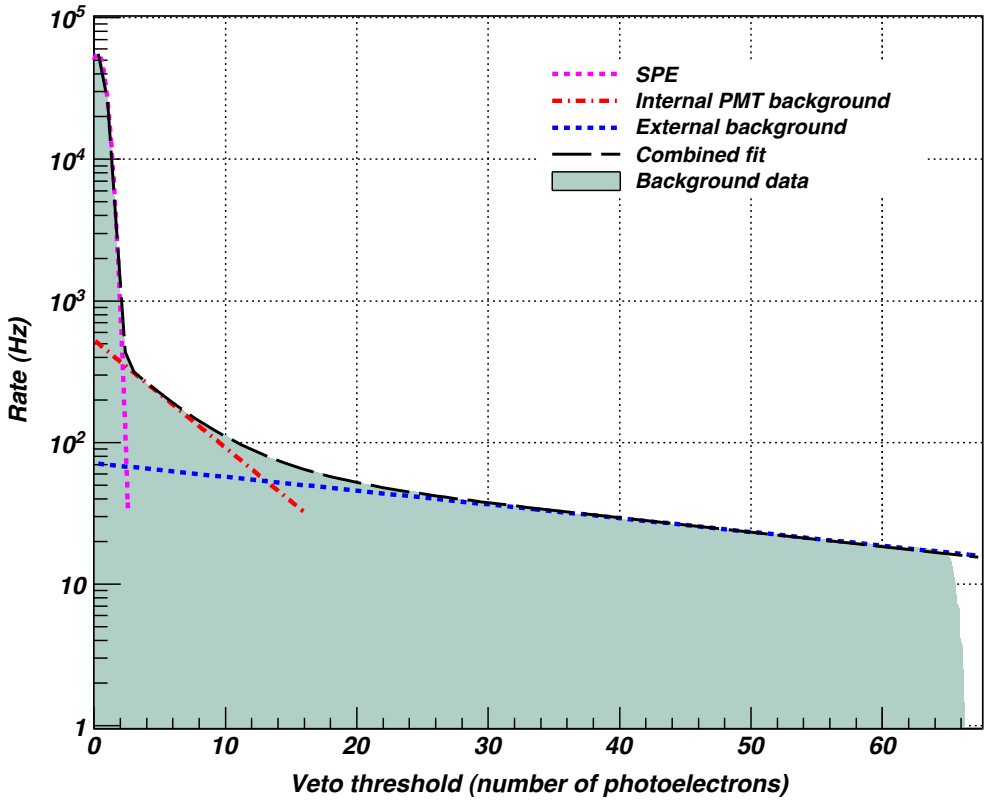

Figure 6.19: The shaded area shows the background signal rate in the full veto detector array as a function of photoelectron threshold. All pulses are considered irreSPEctive of multiplicity, and, as such, the rate shown is cumulative with decreasing applied threshold. Statistical errors are too small to be seen on this scale. The data rate is truncated at 65 photoelectrons where pulse height saturation takes effect. The fit to the data is composed of three well resolved components made up of SPEs, internal background from the veto detector PMTs, and γ-ray background from radiological contamination within surrounding materials and the veto detector itself.

signals of 0.4 μs, largely due to the sampling rate of 100 μs^{-1}. However, other methods may be used to optimise this window and these are discussed later.

Having determined the position of the S1 trigger pulse in ZEPLIN–III and the position of the corresponding prompt coincidence pulse in the veto, the process of prompt coincidence tagging may begin. Before this can be done, it is important to consider low energy recoil events where the ZEPLIN–III DAq is triggered by the S2 signal. In order to identify these events, the acceptance window is not in a fixed position on the veto timeline. Rather, the position of the window associated with the position of the identified S1 pulse time.

The first step to tagging prompt events is to define an acceptance window. In order to optimise the size of this window, a fixed threshold of 2 photoelectrons is set (signals must be $\geq$2 photoelectrons to be considered) and the size of the window, centred on the mean of the prompt veto pulse distribution, is slowly increased.

Figure 6.22 shows this results of increasing this window. The tagging rate increases rapidly as the window increases until it reaches a point where any additional tagging is due to additional accidental tagging alone. The inset in Figure 6.22 shows how it is possible to extract the accidental tagging rate in the prompt window using the accidental rate increase for larger windows. A line is fitted to the data between 0.8 and 1.2 μs which is extrapolated down to an acceptance window of 0 μs. The difference between this 0 value and the value of the line at the chosen window size defines the accidental rate. For an acceptance window of 0.4 μs, the accidental rate is 0.3% for a threshold of $\geq$2 photoelectrons. For a window larger than this, any increase in tagging efficiency is due to an increase in accidental rate.

It is possible to verify the accidental rate calculated using the data in Figure 6.22 by applying the 0.4 μs acceptance window to a region of the ZEPLIN-III timeline which is far from the expected prompt coincident signal region ($\sim$17.7 μs). Figure 6.23 shows the results of setting a fixed acceptance window whilst increasing the energy threshold. For a threshold of $\geq$2 photoelectrons, the tagging efficiency is 0.4%. There is a $\sim$25% discrepancy between this result and that of the method described by the previous method. This is, however to be expected due to the fact that, in Figure 6.22 a tagging efficiency above the accidental rate of 27.85% is seen. This means that roughly 1/4 of all accidentally tagged pulses

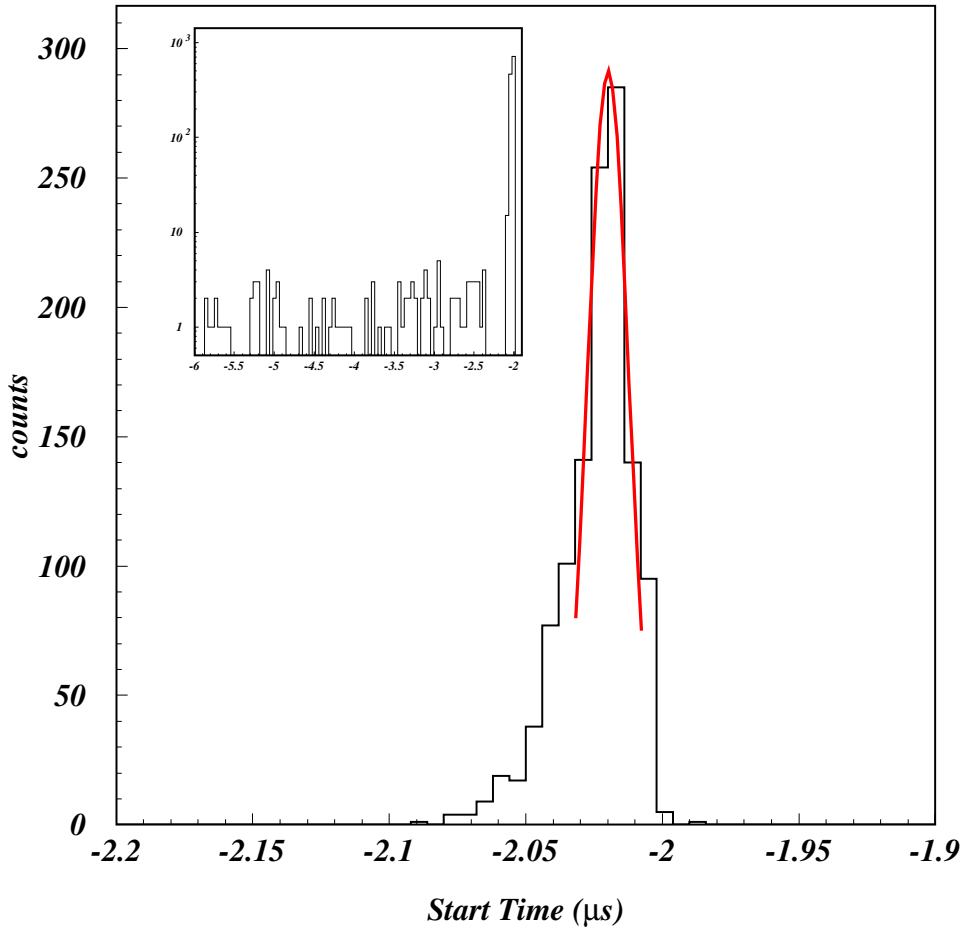

Figure 6.20: The primary signal (S1) start time in ZEPLIN–III. The large plot shows that, for S1 triggers, the S1 start time must be greater than ∼-2.075 μs. The inlayed plot, on a logarithmic scale, shows that this is well separated from S2 triggered events

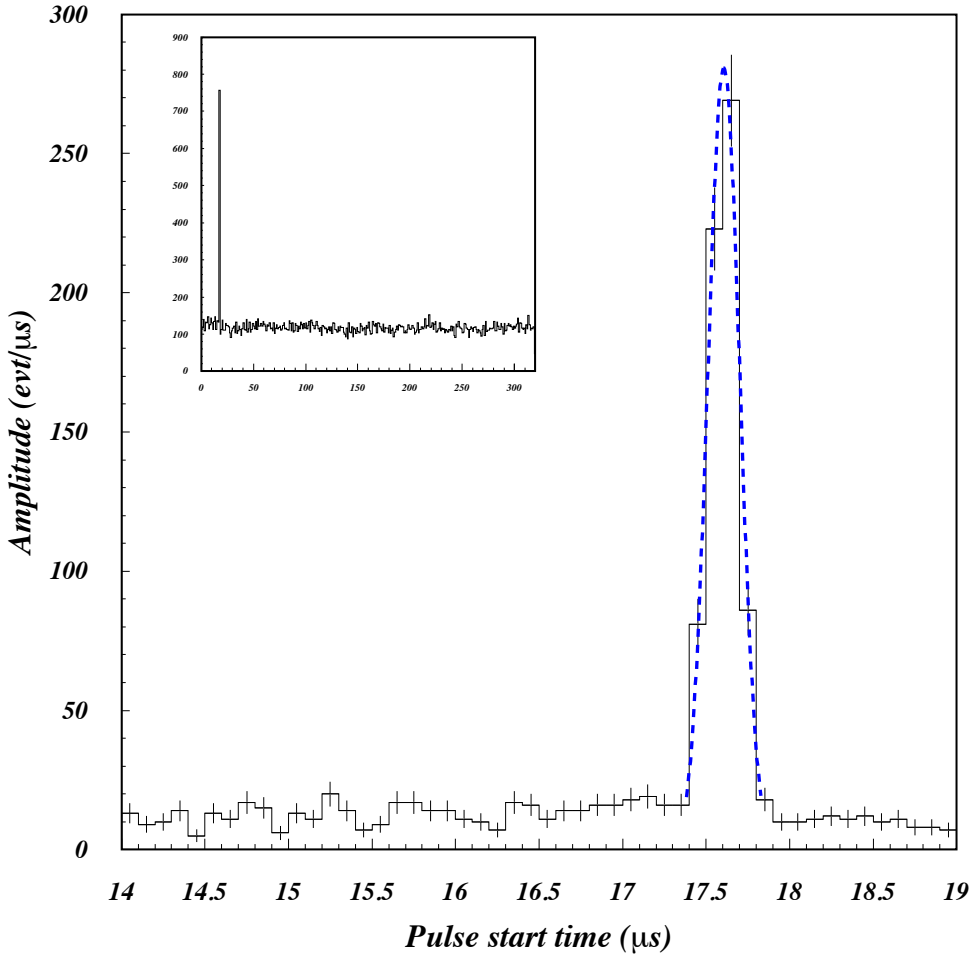

Figure 6.21: A histogram of the pulse start times from 14-19 μs in the veto detector when triggered by ZEPLIN–III. The peak at 17.6 μs is from pulses in the veto detector that are in prompt coincidence with the S1 signals that triggered ZEPLIN–III. Signals in prompt coincidence with the S1 where ZEPLIN–III was triggered by S2, however, are contained within a 16 μs region before the peak. Pulses occurring after this peak cannot be in prompt coincidence with events in ZEPLIN–III, and represent the background rate in the veto detector itself. Insert: the pulse start times in the veto detector for the full 320 μs timeline recorded. The single peak at low times is that seen at 17.6 μs, with the region around it enlarged in the main histogram.

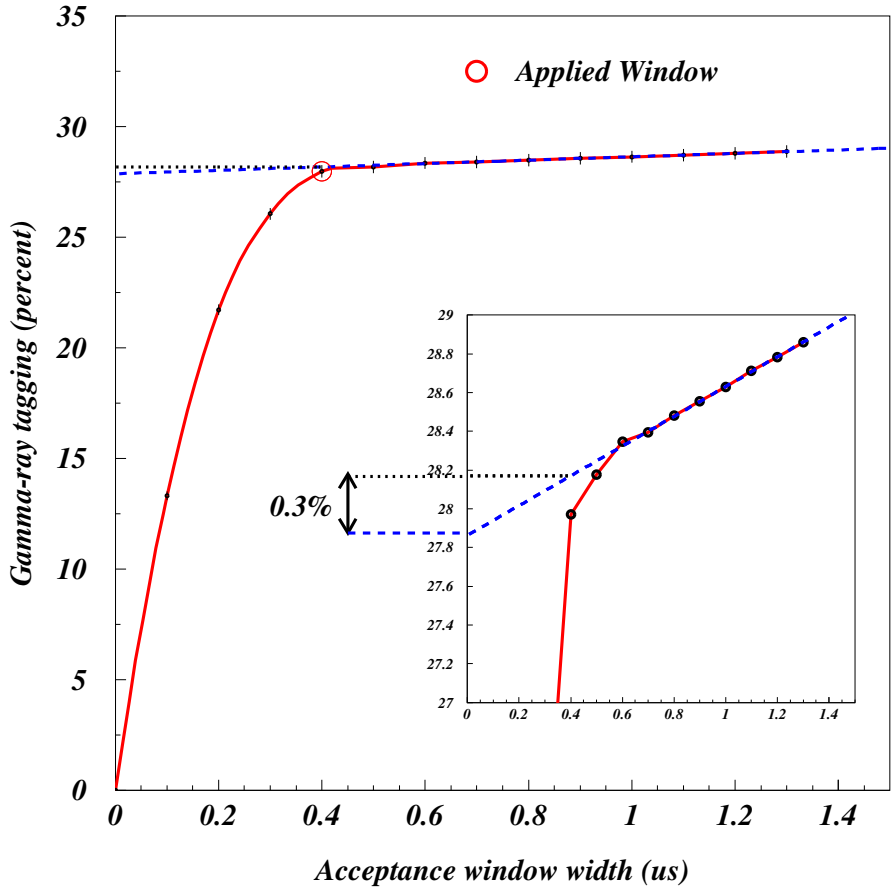

Figure 6.22: The efficiency for tagging prompt coincidence events between ZEPLIN–III and the veto detector, as a function of acceptance window width. The efficiency rises rapidly as the window is opened to 0.4 μs, but beyond this any increases in efficiency merely reflect the linear inclusion of accidental coincidence events. The inlayed plot shows that the accidental rate distribution can be extrapolated to give an accidental rate of 0.3% for a window of 0.4 μs.

will occur in coincidence (within the 0.4 μs defined by the window) with a true prompt tagging pulse.

The tagging efficiency can be increased further by the application of fiducial cuts in ZEPLIN-III. Without any cuts and a threshold of $\geq$2 photoelectrons, the tagging efficiency is $\sim$25.5%. Using the same cuts defined in Section 4.6, the prompt tagging efficiency becomes (28.0$\pm$0.2)%. Figure 6.24 shows the effect of the application of fiducial cuts.

The prompt tagging efficiency represents a significant improvement over the veto system used in the ZEPLIN–I and ZEPLIN–II experiments and exceeds initial design SPEcifications. This is as a result of successful operation at the low threshold of only 2 photoelectrons (equating to $\sim$40 keV of energy deposition in the scintillator) and a narrow coincidence window mitigating the effects of background rate in a large tonne-scale external veto device that also doubles as shielding for the ZEPLIN–III instrument.

In addition to background rejection, the significant γ-ray tagging efficiency further enhances the effectiveness of the veto detector as a diagnostic aid. In particular, it provides an independent estimate of the γ-ray background in ZEPLIN–III, which is eSPEcially important at low energies where particle discrimination is not perfect. If a small population of WIMP candidates is found in ZEPLIN–III, the (lack of) prompt tagging can rule out (to some confidence level) a significant γ-ray component. In addition, a discrepancy between predicted and observed tagging efficiencies in the science exposure could indicate that electron recoil backgrounds have a significant β contribution, such as that expected from ^{85}Kr or surface contamination. The 28% of the science data tagged as prompt γ-ray events can also be unblinded and used to characterise the background of the SSR without jeopardising the blindness of the data in the WIMP search region.

6.3.4 Neutron Tagging

The main purpose of the veto is to provide tagging of nuclear recoil events. Without the ability to tag neutron induced nuclear recoil events, it would be impossible to distinguish between low energy neutron scatters and WIMP interactions. In the case of a positive signal, the significance of the veto is greatly enhanced (see section 6.5) As described in Section 6.1, neutrons are captured by ^{157}Gd leading

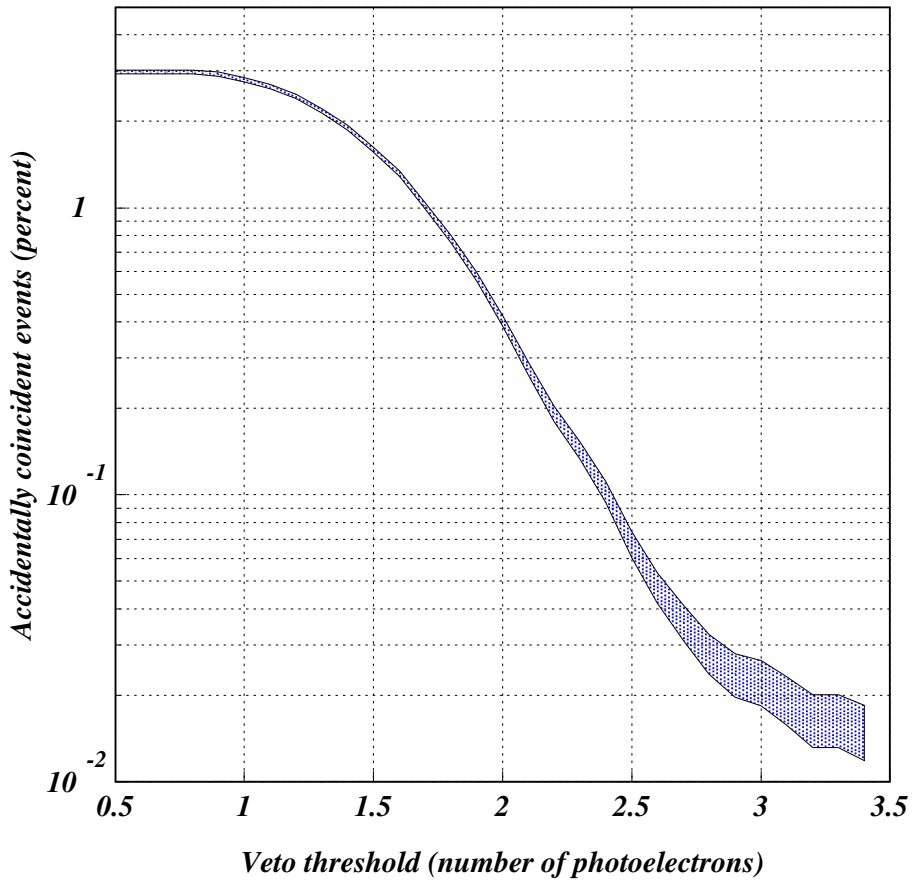

Figure 6.23: A plot of the percentage of events accidentally tagged in a 0.4 μs window as prompt coincidences with signals in ZEPLIN–III, as a function of veto detector threshold. Vertical errors, cumulative with increasing photoelectron value, are indicated by the filled area. The windows selected to measure this rate deliberately exclude the prompt window and the pre-trigger region where genuine prompt coincidences with ZEPLIN-III pulses are found. For the prompt tag threshold of 2 photoelectrons equivalent signal in the veto detector, the probability that a tagged event is not correlated with the signal in ZEPLIN–III is 0.4%.

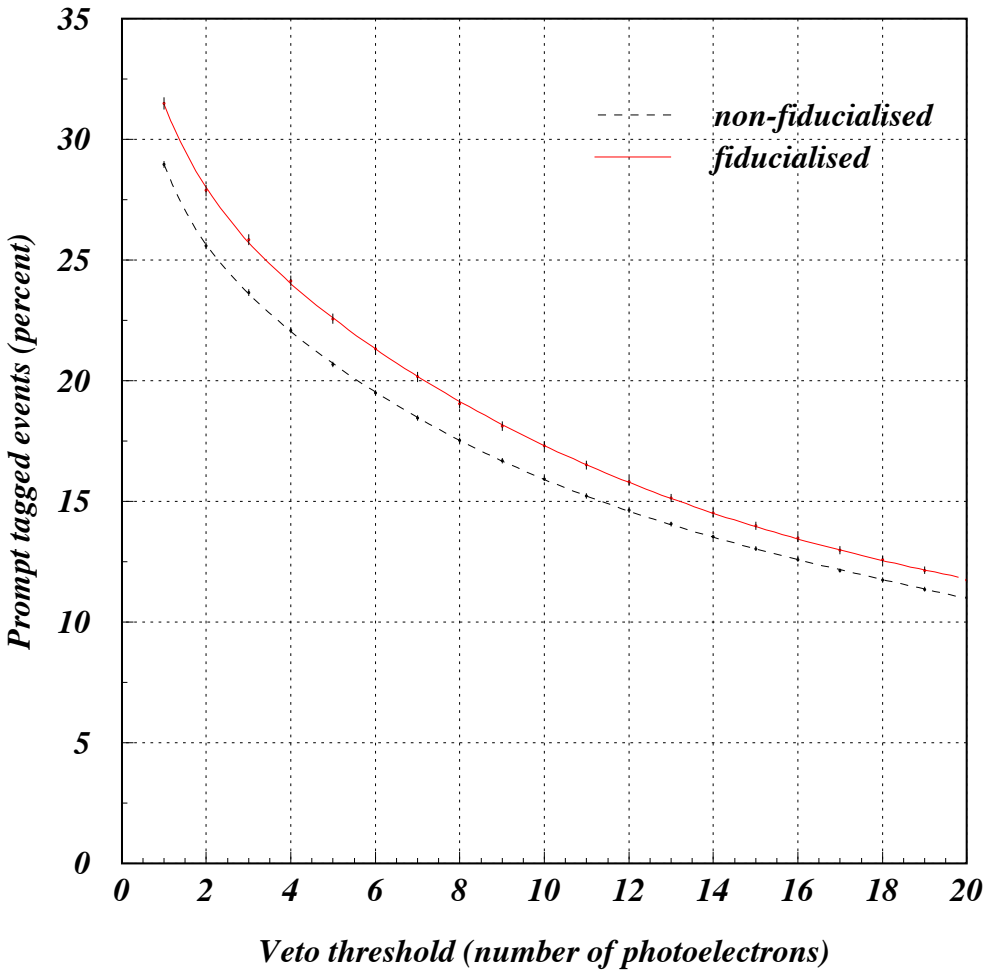

Figure 6.24: The percentage of coincident events satisfying the prompt tag as a function of veto detector threshold, for all synchronised ZEPLIN–III events in the target Xe as well as for the fiducialised target. No additional cuts or restrictions have been applied to the the data.

to the emission of delayed γ-rays. Figure 6.4 shows the expected distribution of neutron capture times.

As discussed in Section 5.4.2, there is an expectation of 0.4 single scatter neutron induced nuclear recoils per year in the WIMP search region. This being the case, the neutron tagging power of the veto must be calculated using AmBe calibration data. Using the same method as described in Section 4.7, the elastic nuclear recoil population is selected. This selection was taken across the full available energy range and between $\pm 2\sigma$ of the mean of the population. Figure 6.25 shows the AmBe calibration and the selected elastic nuclear recoil population.

Having selected elastic nuclear recoil events, it is possible to determine the neutron tagging efficiency. Figure 6.4 suggests that an acceptance window of 70 μs should allow the effective capture of most neutrons. As with the tagging of prompt γ-rays, it is important to determine the level of accidental tagging rate as a function of energy threshold. Figure 6.26 shows that a 70 μs acceptance window with a threshold of $\geq$10 photoelectrons gives an accidental tagging efficiency of 0.8%. This accidental rate is determined in the same way as the prompt tagging accidental rate is in Figure 6.23. The 70 μs window is moved to a region of the veto timeline that is far away from the expected signal region, giving a region populated only by a background signal.

An accidental tagging efficiency of $<$1% was judged to be a reasonable accidental rate. As previously mentioned, only 0.4 single scatter neutron induced nuclear recoils are expected per year in the WIMP search region and, because of this, it is unreasonable to remove large amounts of background data in order to tag a neutron event that may not occur over the duration of the SSR. The 10 photoelectron threshold with no slab multiplicity requirement is found to yield an efficiency and accidental coincidence rate equivalent to a multiplicity of 2 modules recording a total of at least 8 photoelectrons. The rate of accidental coincidences for delayed tags remains low despite a lower photoelectron threshold as a result of the stronger dependence on multiplicity for background γ-rays that make up the accidental rate. Combining the two selection criteria, *i.e.*, higher photoelectron threshold with no multiplicity and lower photoelectron threshold with a multiplicity of 2, does not result in a gain in efficiency since they each tag correlated datasets. For the same reason, however, the accidental coincidence rate is increased. As such, the delay tag selection criteria is only that a combined

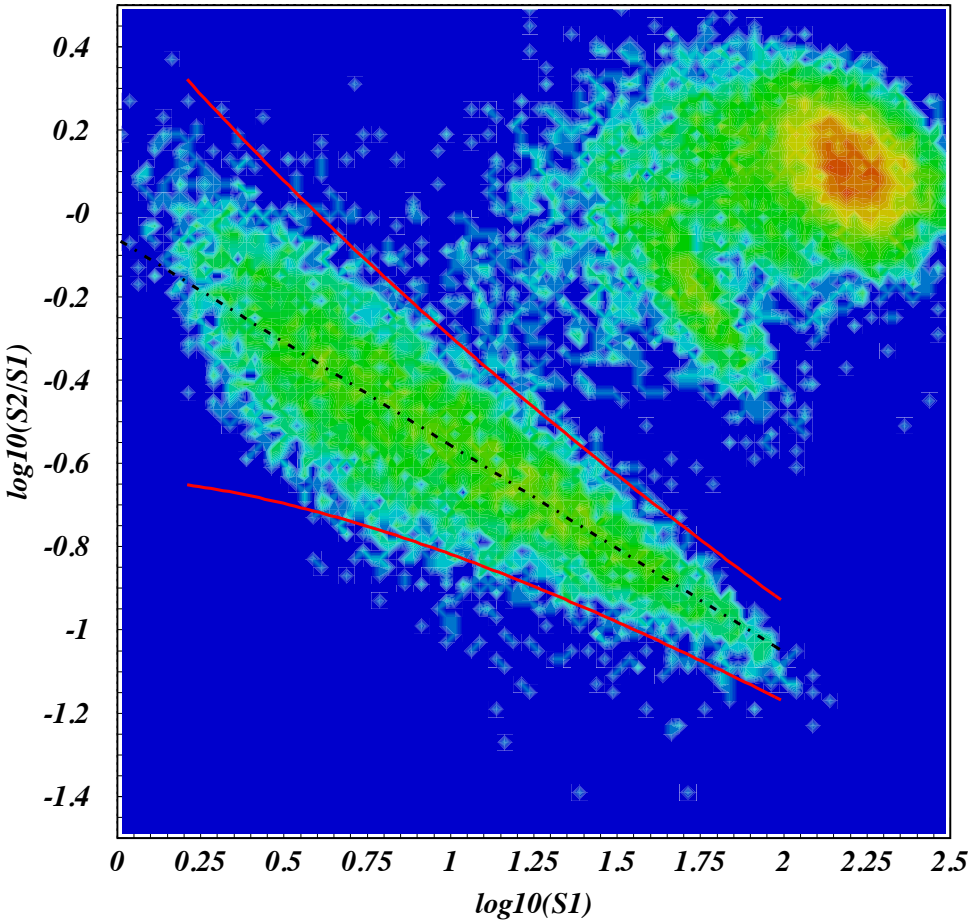

Figure 6.25: A density plot of S2/S1 against S1 for events from an AmBe calibration dataset showing the separation between nuclear and electron recoil bands. Only fiducialised single scatter events are shown with no additional cuts to the data. The mean of the density distribution of the elastic nuclear recoils is indicated by the dashed line. The solid lines to either side mark 2σ from the mean and define the events selected for analysis of the response of the veto detector to neutrons - primarily the tagging, delay time, multiplicity, and energy distributions.

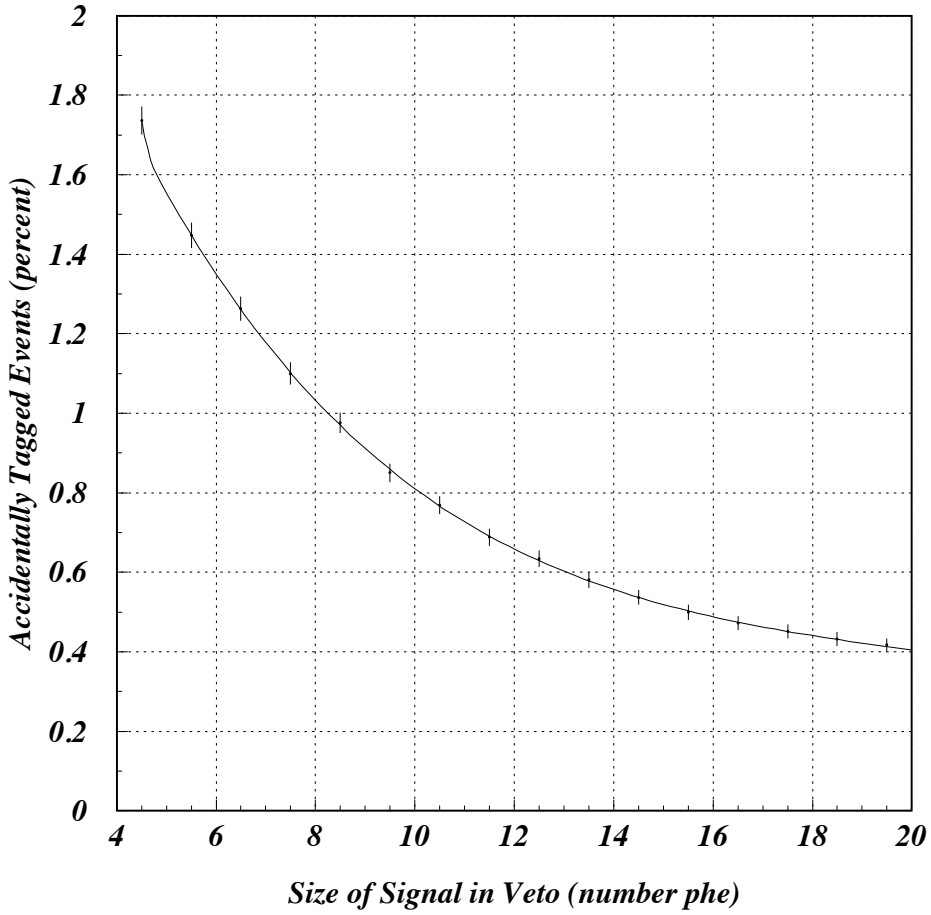

Figure 6.26: Accidental tagging percentage for a 70 μs acceptance window as a function of energy threshold. A threshold of $\geq$10 phe gives an accidental tagging rate of 0.8%. A line is included to guide the eye.

minimum of 10 photoelectrons equivalent signal be detected within the delay window, with no minimum multiplicity condition, although any pulses distributed across multiple modules must be in coincidence with all others.

Having determined the parameters that produce an acceptable level of accidental tagging, the window is applied to the region of the veto timeline immediately following the prompt window. Figure 6.27 shows the time difference between the S1 pulse in the ZEPLIN–III timeline and the pulse in the veto timeline that is identified as the delayed coincident pulse from the neutron capture. Fitting to the data and removing the flat background associated with the accidental tagging rate gives excellent agreement with the simulated capture time. The characteristic times of the fits to both the simulated and experimental data are (10.4±0.3) μs and (10.7±0.5) μs and are in agreement within errors.

The characteristic capture times can be used to determine the Gd concentration in the polypropylene shielding. Figure 6.28 shows the characteristic capture time for thermalised neutrons as a function of Gadolinium concentration. A characteristic capture time of (10.7±0.5) μs is equivalent to a Gd concentration of (0.42±0.03)% (by weight) which is very close to the 0.5% Gd loading fraction required to give the optimised neutron tagging efficiency (as shown in Figure 6.2).

Given the confidence that neutrons were being tagged as expected, it is possible to look at the energy dependent neutron tagging efficiency. The energy of the γ-rays emitted following the capture of thermalised neutrons is independent of the energy of the neutron recoil energy. This is due to the fact that the angular distribution of the interacting neutrons in the xenon is destroyed through proton recoils in the polypropylene and thermalisation prior to capture. Figure 6.29 shows the dru spectrum for nuclear recoil events in the AmBe dataset. The spectrum falls as expected for higher energy recoils, this is because neutron scatters will preferentially scatter with the deposition of a small amount of energy transferred to the recoiling Xe nucleus (because of the large difference in the size of a neutron and a Xe nucleus). The rate of tagged events follows the AmBe spectrum at a ∼constant fraction and this is verified in Figure 6.30. The average tagging fraction for the 2-20 keV_{ee} WIMP search region is (58.8±0.5)% and this tagging fraction remains constant between 0-100 keV_{ee} as expected, although the uncertainties in the tagging efficiency do increase with falling statistics.

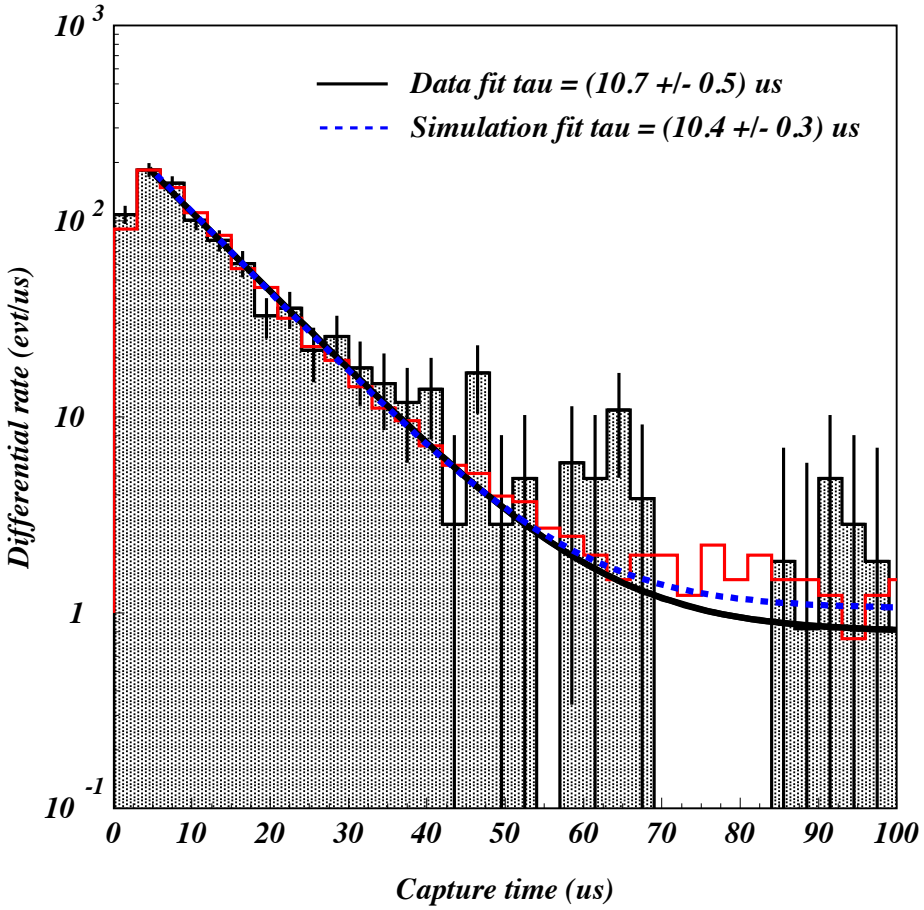

Figure 6.27: The measured distributions of pulse times in the veto array, relative to the S1 signal from nuclear elastic scatters in ZEPLIN–III, for events that satisfy neutron selection cuts is shown as the open histogram. A the fit to the data, indicated by the solid line, yields a characteristic decay time of (10.7±0.5) μs. The shaded histogram overlaid is the predicted distribution from Monte Carlo simulations, showing good agreement, with a characteristic timescale of (10.4±0.3) μs.

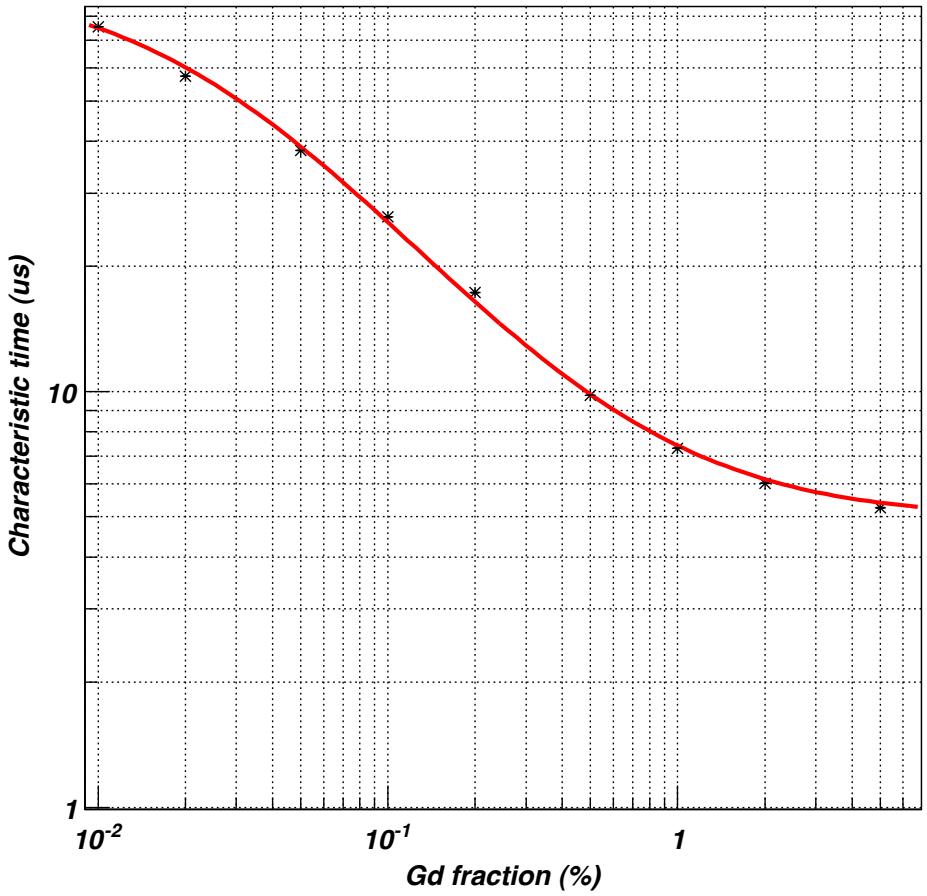

Figure 6.28: A plot of the characteristic time for delayed coincidence signals from neutrons against Gd concentration within the passive shielding, as calculated by Monte Carlo simulations. The solid line is a fit to the data points. The measured characteristic time of (10.7±0.5) μs corresponds to an average Gd concentration of (0.42±0.03)% within the Gd-loaded polypropylene.

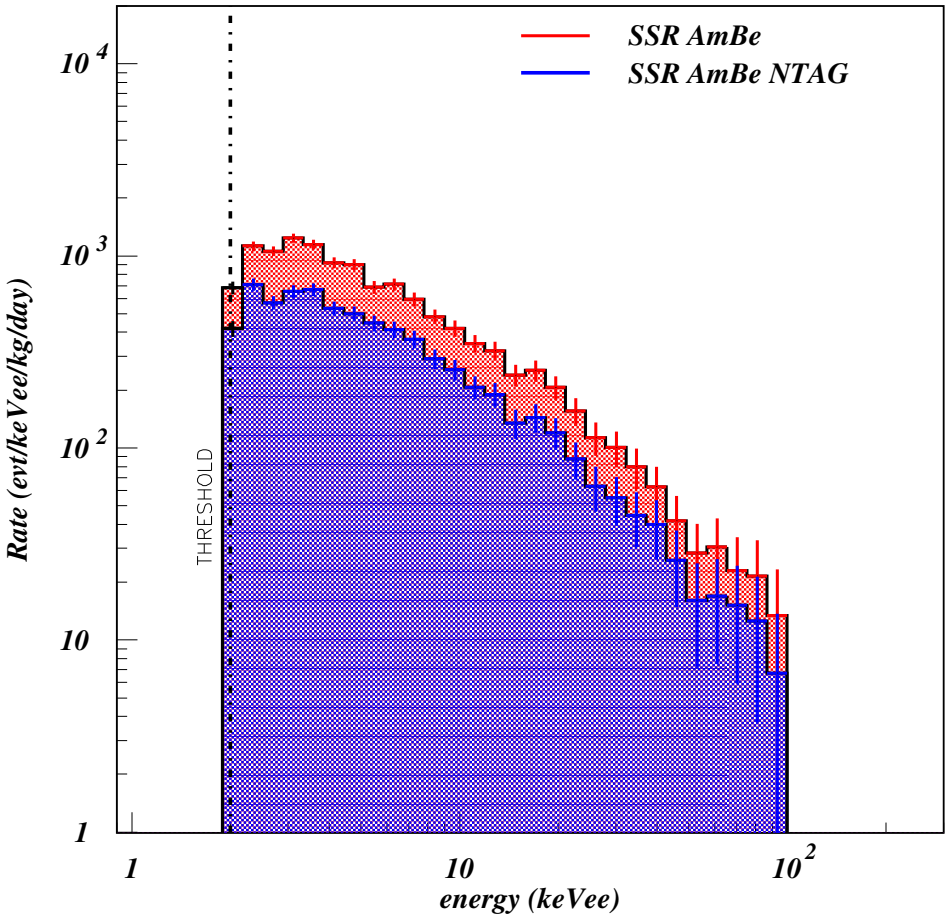

Figure 6.29: A dru plot for elastic nuclear recoil events (red) compared with tagged nuclear recoil events (blue).

For a definition of the full neutron tagging efficiency, captures occurring in the prompt window must also be considered. In data, such neutron captures would be vetoed by the prompt tag but it is not easy to get the rate of these tags using the AmBe calibration data.

The emission of a neutron from the AmBe source occurs through the Be (α,n) reaction. In coincidence with this reaction, high energy γ-rays (mostly 4.44 MeV) are emitted. These γ-rays are identified as prompt tags in the veto and their rate dominates any neutron captures occurring within the prompt window. The high energy γ-rays are indistinguishable from those emitted in the de-excitation of the ^{158}Gd nucleus after radiative capture and, as such, it is impossible to measure the number of neutron tags in the prompt acceptance window directly.

It is possible, through extrapolation of the delayed time distribution for neutron capture, to estimate the fraction of neutron tags in the prompt window. A Gaussian fit to the lower edge of the distribution predicts an additional (1.7±0.1)% of neutron tags falling in the prompt window. Monte Carlo simulation predicts a tagging fraction of (60.7±0.1)% for a 10 photoelecton threshold of which (1.5±0.1)% are within the prompt window [162]. The combined (prompt+delayed) tagging fraction is measured as (60.5±0.5)% which is in excellent agreement with the simulated value.

6.4 Tagging Fractions in Science Data

Having defined the parameters required to tag prompt and delayed events, the tagging requirements can be applied to WIMP search data. This can be done without jeopardising the blindness of the WIMP search data by SPEcifying that the WIMP search region remain blinded except for vetoed events which are, necessarily, not signal. Figure 6.32 shows the energy dependent prompt (PTAG) and delayed (DTAG) tagging rates as a function of recoil energy in ZEPLIN–III. As the WIMP search region (and, in-fact, all data below the nuclear recoil mean) remained blinded, the tagging criteria are applied to electron recoil data only. Both the PTAG and DTAG rates should increase or decrease relative to the background rate. The right hand plot in Figure 6.33 shows the energy dependence of delayed tagging is, indeed flat, with an efficiency of (0.75±0.17)% which is consistent with the 0.8% expected.

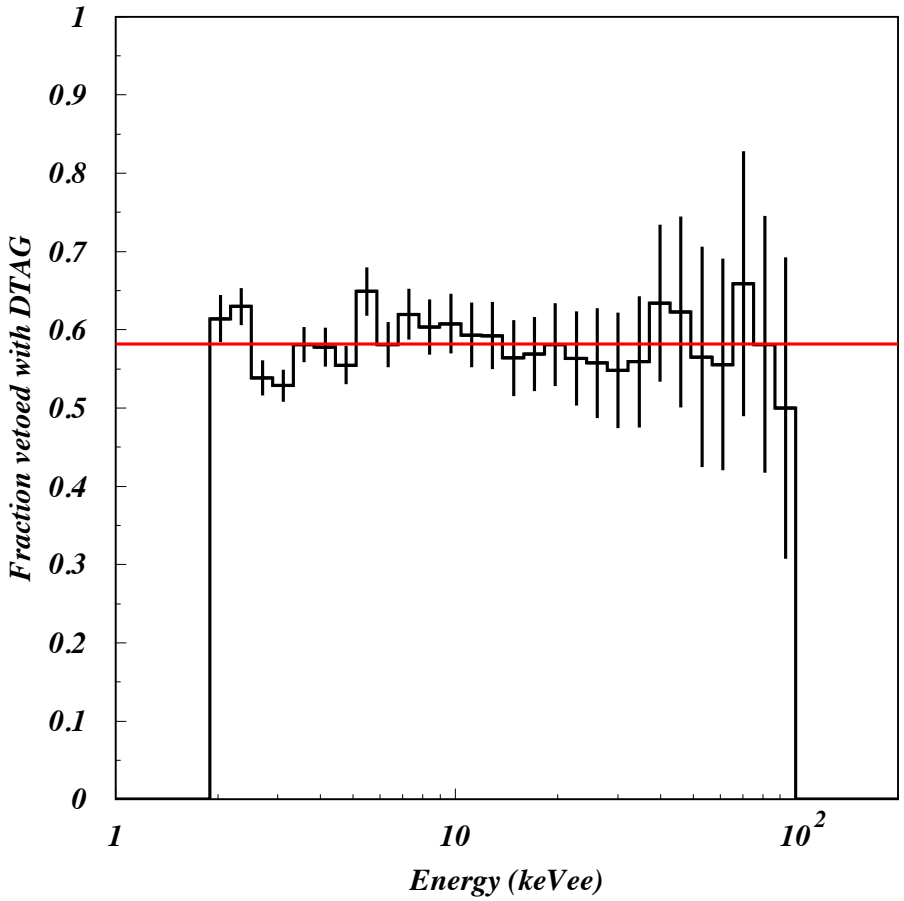

Figure 6.30: The delayed neutron tagging efficiency as a function of nuclear recoil energy deposition in ZEPLIN–III. The average tagging efficiency is (58.8±0.2)% for single scatters in the 2-20 keV_{ee} energy WIMP acceptance window in ZEPLIN–III. This efficiency remains constant since the probability for detection of γ-rays following radiative capture of the neutron and the de-excitation of ^{158}Gd is independent of the neutron energy following the scatter in the Xe.

The left hand plot in Figure 6.33 shows the energy dependent tagging fraction of the prompt tagging. In the WIMP search region of 2-20 keV_{ee}, a fit to the data gives a tagging fraction of (26.8±1.3)%, in agreement with the expected fraction when the 97% synchronisation efficiency is accounted for. However, it is clear to see that there is a deviation from the flat tagging efficiency at >160 keV_{ee}. The reason for this is due to neutron activation of ^{129}Xe and ^{131}Xe. The peak which causes the deviation seen is from ^{131m}Xe with an energy of 164 keV. Given that the absorption length in liquid Xe at an energy of 160 keV is ∼5.9 mm [167], it is unlikely that a γ-ray emitted through the de-excitation of ^{131m}Xe would leave the Xe volume and, if it did, the energy would be below the threshold needed to give a signal in the veto plastic scintillator. This means that dips in tagging efficiency should be expected at any energy where a known γ-ray is emitted. Above ∼295 keV, the path length for Compton scattering becomes shorter than that for photoelectric absorption so the effect of apparent loss of tagging efficiency is mitigated. The attenuation length for γ-ray radiation is shown in Figure 6.31.

6.4.1 Simulated Tagging Efficiency

The dominant sources of background in the SSR of ZEPLIN-III were calculated in Section 5. Figure 5.7 shows their individual contributions in the ZEPLIN–III background as a function of dru. The simulation shows a cumulative rate of 0.731 dru at 10 keV_{ee} and is made up of four components, the PMTs (0.35 dru - of which 92% comes from the glass envelope and 8% comes from the potassium generators), the veto scintillator (0.077 dru), the ceramic feedthroughs (0.074 dru) and the polypropylene shielding (0.23 dru).

The ZEPLIN–III simulation is run in order to ascertain the tagging efficiencies for γ-rays emitted from each of these components and a weighted average taken to give the expected tagging efficiency for all contributions. The polypropylene shielding gives us an expected tagging fraction of only 4% but this is mitigated by the γ-ray emission in the veto scintillator plastic giving a 99% tagging efficiency (the break-up of the nucleus giving large energy depositions in the scintillator that is tagged with near unity efficiency). The simulated tagging efficiencies are shown in Table 6.1 and the weighted average of all tagging efficiencies gives an expected tagging efficiency of (27.0±0.6)% which is in good agreement with the measured

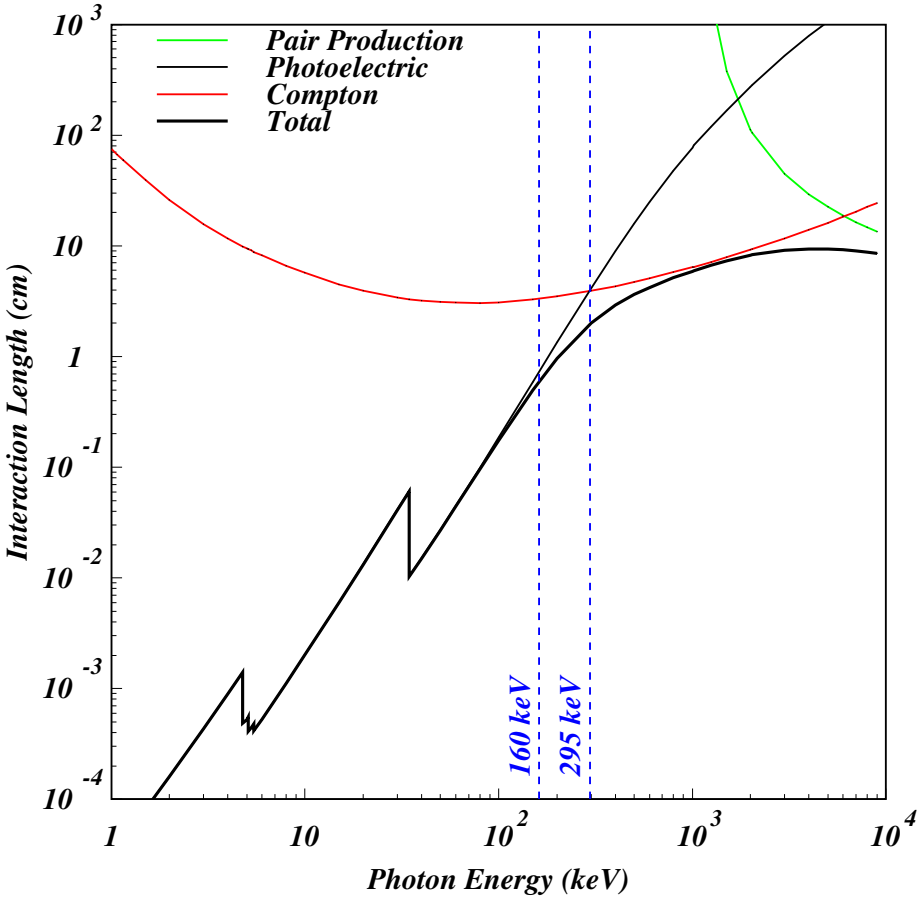

Figure 6.31: The interaction length of photons in liquid Xe. The key defines the three different processes with the thick black line giving the total interaction length. At the ∼160 keV energy of the ^{131m}Xe γ-ray, the interaction is dominated by the photoelectric effect. Also marked is 295 keV, the point at which the interaction lengths for the photoelectric effect and Compton scattering are identical. Above this energy, Compton scattering begins to dominate, leading to γ-rays that can be detected by the veto.

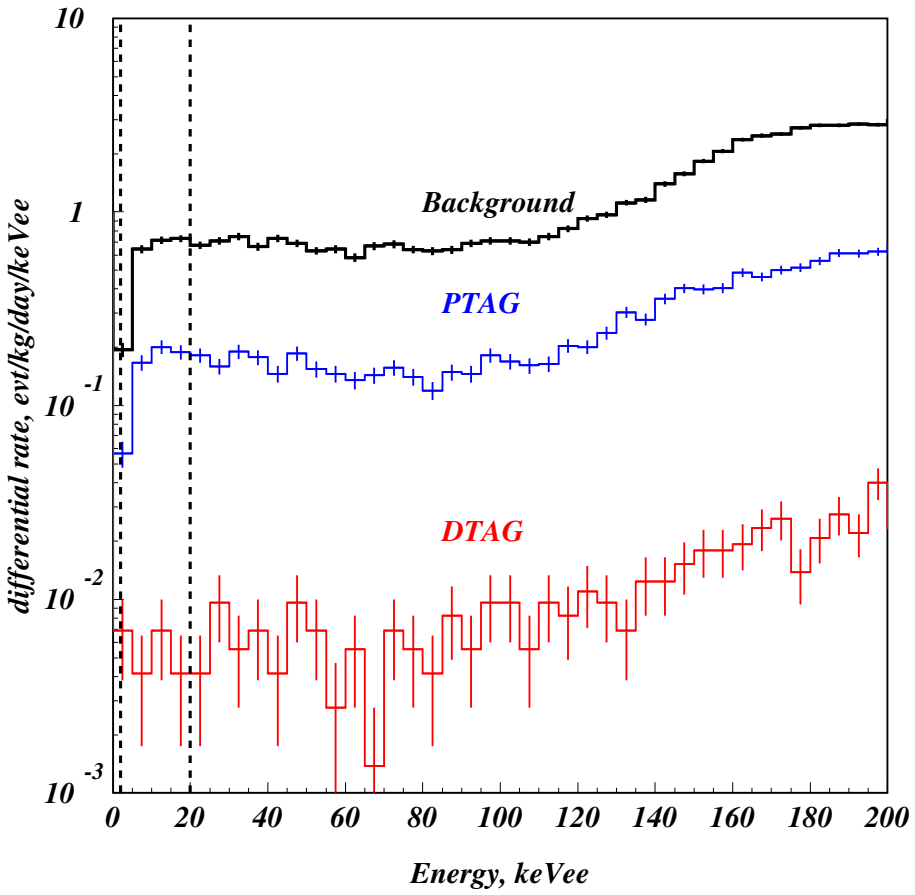

Figure 6.32: A histogram of the differential rate of background events in the fiducial target of ZEPLIN–III in the SSR (solid line - labelled Background). This rate is over an order of magnitude lower than that in the first science run as a result of the upgrades to the instrument. Overlaid are histograms of the differential rates of all events tagged as only prompt coincidences (labelled PTAG), approximately 0.3 of the total rate, and events tagged as only delayed coincidences (labelled DTAG). The latter population is comprised of accidentally coincident delayed tags, and represent 0.8% of the electron recoil background.

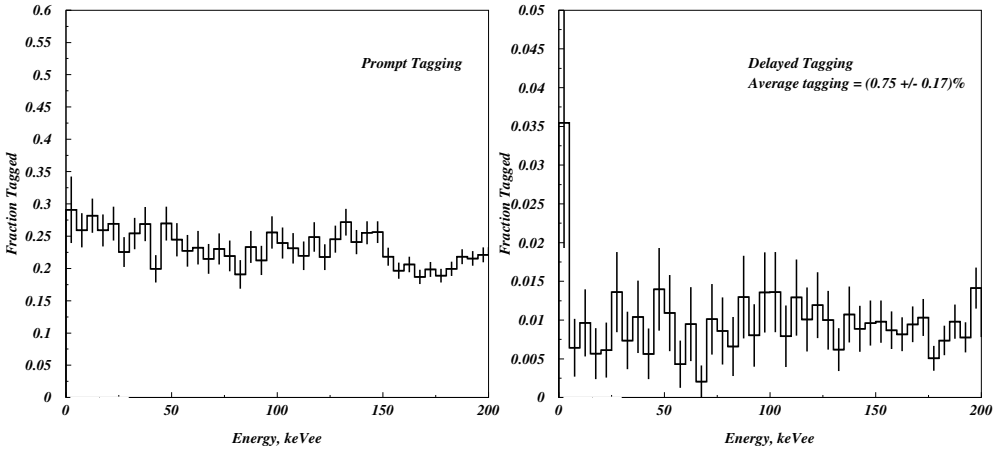

Figure 6.33: The prompt tagging rate as a fraction of total rate (left). The tagging fraction is approximately uniform with dips in the tagging fraction above 140 keV_{ee} due to the photo-absorption of γ-rays emitted from excited states of ^{129}Xe and ^{131}Xe due to the neutron activation of Xe. The delayed tagging fraction is shown on the right. This rate is $\sim$0.8% and is not affected by the photo-absorption peaks seen in the prompt tagging distribution. This is entirely consistent with the accidental tagging rate.

Component	dru	Tagging Fraction
PMT envelope	0.322	0.251±0.007
PMT K generators	0.028	0.316±0.006
Plastic Scintillator	0.077	0.990±0.001
Shielding	0.230	0.042±0.004
Ceramics	0.074	0.297±0.002
Total	**0.731**	**0.270±0.006**

Table 6.1: Simulated tagging fractions for γ-ray background contributing components. The total tagging fraction is an average weighted by the individual dru contributions from each component.

tagging efficiency of (28.0±0.2)%. The simulation does not account for the 0.4% accidental tagging rate which, when included in the simulated calculation, gives excellent agreement within errors.

6.4.2 Position Dependence

In addition to calculating an energy dependent tagging fraction for electron recoil events in background, it is also possible to define a position dependent tagging fraction. The position dependence can be split into two contributions, an xy and a depth dependence. The depth dependence is measured by slicing the Xe volume into drift time sections and applying the prompt tagging parameters. The upper right-hand plot in Figure 6.34 shows the tagging fraction for each of these slices. The tagging fraction increases for events closer to the liquid surface. The dependence can be further split into contributions from the barrel only and the roof only. The dependence is shown in Figure 6.35. There appears to be a very slight depth dependence for prompt γ-ray signals tagged in the roof and a rather stronger depth dependence for events tagged in the barrel.

The radial dependence can be explained by the fact that γ-rays which scatter towards the top of the liquid volume have a wider solid angle available to reach the scintillator in the barrel without passing through any detector materials such as the copper base of the target chamber. This means that, on average, γ-rays which scatter higher in the liquid volume are less likely to be attenuated, giving a signal in the veto.

The xy dependence is measured by taking a radial slices and applying the prompt tagging parameters. The lower right-hand plot in Figure 6.34 shows the tagging fraction for each of these slices. The radial dependence is much less obvious than the depth dependence in the inner most radial slice, (23.5±2.5)%, and the outer most radial slice, (25.5±0.5)%, are within errors of each other. As with the depth dependence, Figure 6.36 shows that any radial dependence is dominated by tagging in the barrel. This figure suggests that tagging is boosted in the barrel for events which occur in the outer radii and is expected as γ-rays scattering in the central Xe volume have, as with the depth dependence, a smaller solid angle for scattering that does not pass through detector materials.

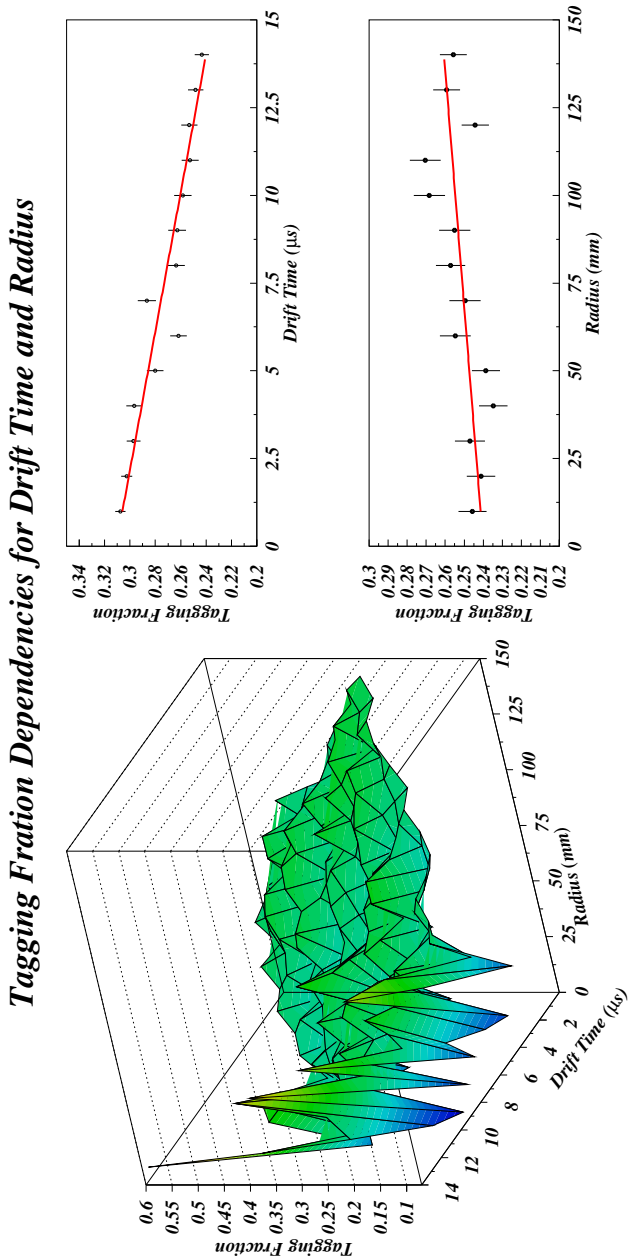

Figure 6.34: The position dependence of the prompt tagging fraction in the ZEPLIN–III SSR. The upper right-hand plot shows the tagging fraction as a function of drift time ranging from (31±0.5)% for surface events down to (25±0.5)% for the deepest events. The lower right-hand plot shows the tagging fraction as a function of radius. This has a slightly smaller dependence, ranging from (23.5±2.5)% for the central Xe volume to (25.5±0.5)% the initial value is, therefore, within errors of the final value. The left-hand plot shows the tagging fraction as a function of both radius and drift time. This remains relatively flat across all values of drift time and radius.

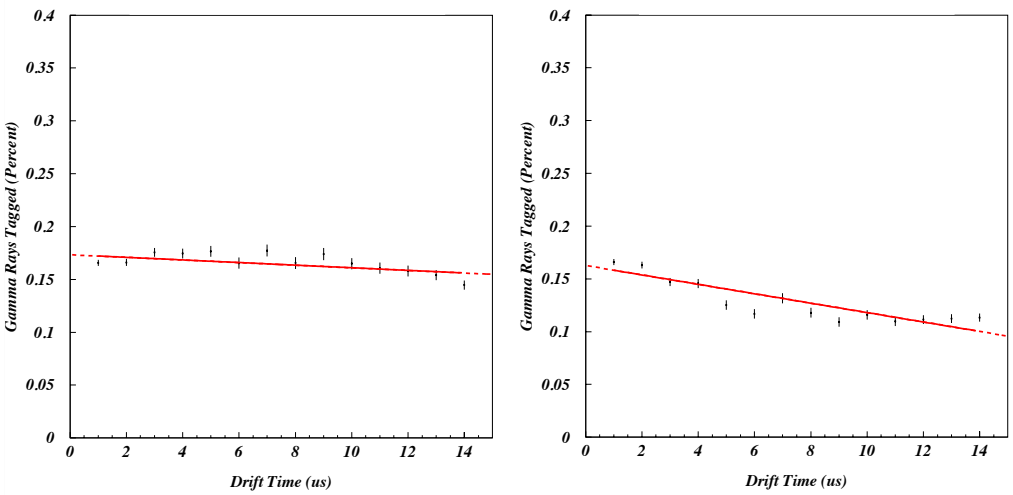

Figure 6.35: The depth dependent tagging fraction for the roof (left) and the barrel (right) of the veto. The roof shows very little depth dependence which suggests that the depth dependence shown in Figure 6.34 is dominated by tagging in the barrel.

6.5 Using the Veto to Confirm a Signal

A generic veto detector is a powerful tool in setting a confidence level for the detection of a WIMP signal. The efficiency, η, for vetoing a background signal gives the relative exposure of vetoed to un-vetoed events in a background sample: $\frac{\eta}{1-\eta}$. Once this ratio has been calculated, it is possible to set a confidence interval for the signal using the profile likelihood ratio (PLR) [168] as implemented in the ROOT class TRolke [169]. Figure 6.37 shows the number of signal events that would constitute a 3σ evidence for a signal as a function of η in the case of an experiment with a single background.

In ZEPLIN–III and similar experiments, the background is described by both nuclear-recoil and electron-recoil populations. These each have different tagging efficiencies which would have to be summed. The discrimination of a 2-phase detector means that the electron-recoil background and, to a lesser extent, the nuclear-recoil background are distributed differently in more parameters than just veto tagging efficiency. For example, as discussed in Section 3.2, S2/S1, pulse-shape, energy and multiplicity are key additional parameters that may be used to

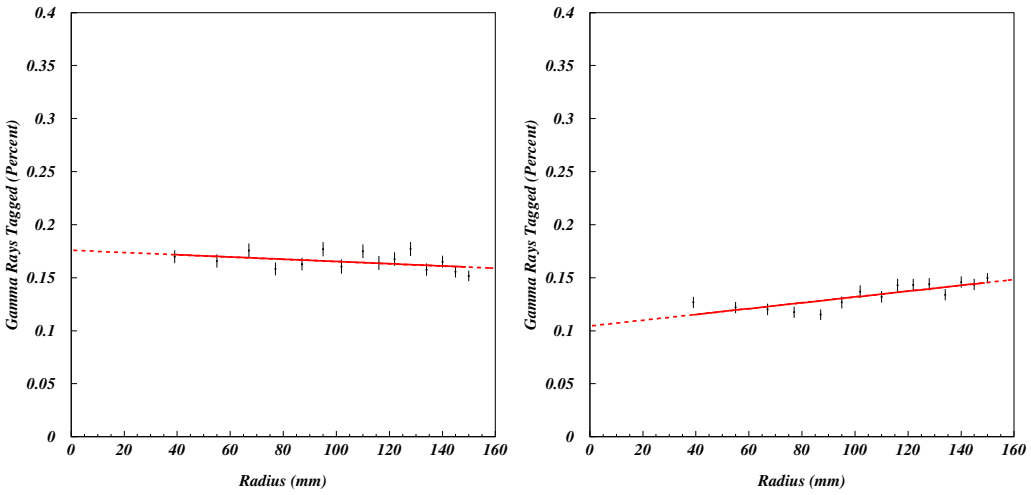

Figure 6.36: The radially dependent tagging fraction for the roof (left) and the barrel (right) of the veto. As with the depth dependence, the roof shows very little radial dependence which suggests that the radial dependence shown in Figure 6.34 is, again, dominated by tagging in the barrel.

reject a background only hypothesis in the WIMP search region. The combination of the counting method used in Figure 6.37 and these additional discrimination parameters would require fewer un-vetoed events to give a 3σ signal result. With an efficiency of 60% for rejection of nuclear recoil background, only 5 un-vetoed events would be required to give 3σ evidence of a WIMP signal. For a dominant background of γ-rays with a tagging efficiency of 26%, $\sim$13 events are required. The limit is for a veto with 100% efficiency where only 1 un-vetoed event would be sufficient.

The veto detector is not designed to simply reject neutron induced nuclear recoil signals. This chapter has shown that it also provides identification and diagnosis of electron-recoil background events. This becomes useful in the un-blinding of ZEPLIN–III science data (discussed in Chapter 7) where some 30% of the data (due to the combination of prompt tagging and accidental delayed tagging) may be viewed without risk of jeopardising the blindness of the WIMP search region and without reducing the overall exposure of the SSR. During the FSR, only 10% of the data was used and this had to be removed from the exposure due to the signal region being undefined and, therefore, unblinded. This allows

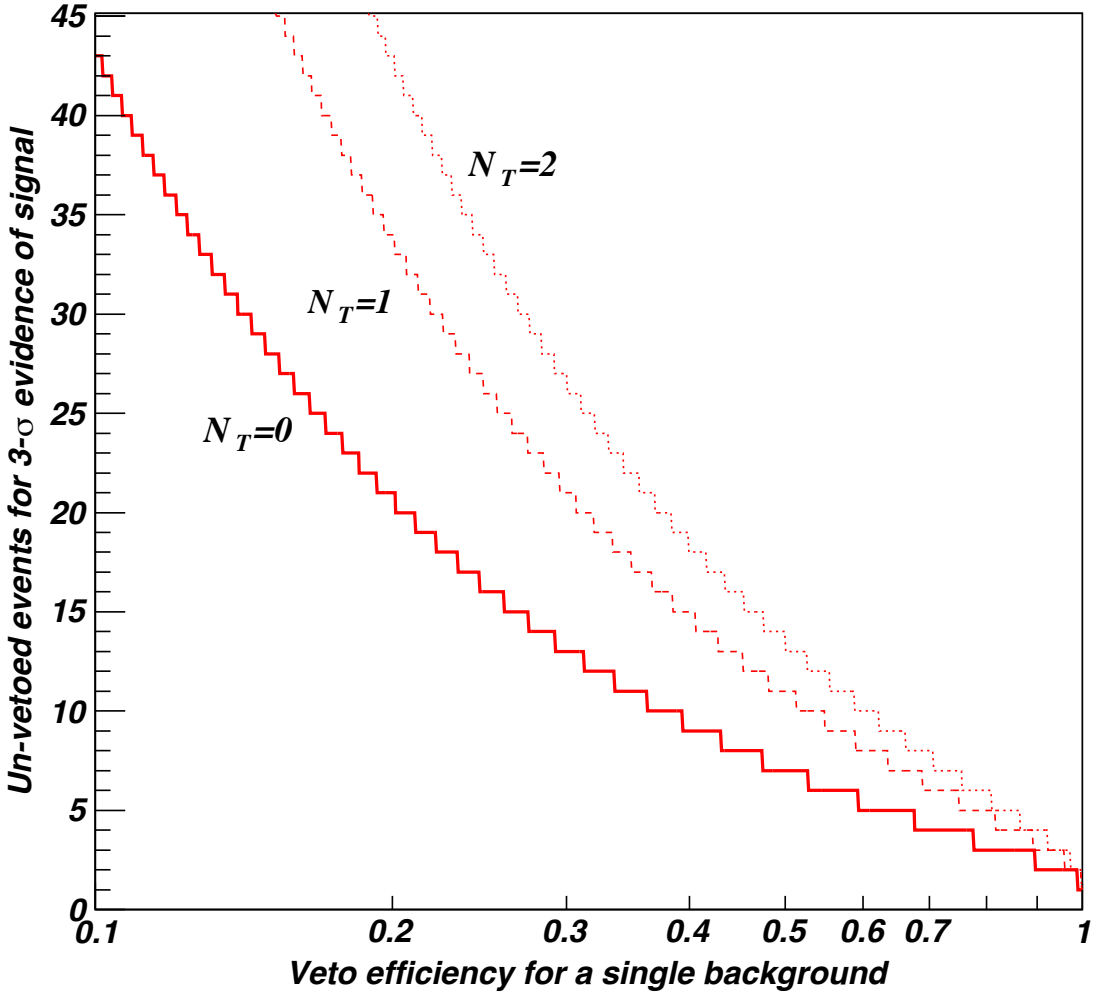

Figure 6.37: The effect of veto efficiency on discovery power for a generic rare event search with a single background and no additional discrimination. N_T is the number of tagged events observed and calculated confidence levels are calculated as in [168].

the study and removal of MSSI events (see Section 4.5) through the definition and application of data quality cuts. With the additional power gained in the case of a positive signal, the veto represents an important and highly effective addition to the ZEPLIN–III detector.

Chapter 7

ZEPLIN–III Second Science Run

7.1 ZEPLIN–III Upgrades

The Second Science Run (SSR) represents the final science goal of ZEPLIN–III. In order to achieve a sensitivity beyond that of the First Science Run (FSR), upgrades were made to many of the systems. The two foremost upgrades came in the form of the installation of the veto detector (discussed in depth in Chapter 6) and the installation of new, ultra low background PMTs (the component assaying and simulated impact of these is discussed in Chapter 5). In addition to this, several improvements were made to the calibration systems and improvements were made to some of the ancillary systems in order to improve both detector characterisation, monitoring and run duty cycle (the length of time, per day, the detector runs in WIMP search mode). The following chapter highlights work done by myself during the final months of my PhD studies together with other members of the ZEPLIN–III collaboration.

7.1.1 PMT Upgrades

Firstly, and most importantly, the PMT array was upgraded. The PMTs used in the first science run were replaced with the tubes developed by ETEL using the radio-assayed components, details of which were discussed in Chapter 5. The PMTs contributed the dominant source of background in the FSR and it was expected that the new array would reduce the background rate to $\sim$10 % that of the FSR. Figure 7.1 shows a comparison between FSR and SSR simulations.

The FSR background rate was dominated by the rate from PMTs whereas, in the SSR, the expected background from the PMTs is comparable to that of the backgrounds from the polypropylene shielding, the veto detector and the ceramic electronic feedthroughs.

The signal blinded data for the period 01/08/10-31/08/10 was used for verification of the predicted backgrounds and electron equivalent energy scale. The dataset used consisted of 680.7 hours live-time, which equates to 181.5 kg.days of data. A neutron calibration was performed during this period on 25/08/10. During a neutron calibration, it is expected that a degree of neutron activation will occur in the Xe volume [170]. The activation occurs through thermal neutron capture onto ^{128}Xe and ^{130}Xe or by fast neutron inelastic scattering onto ^{129}Xe and ^{131}Xe. The resultant ^{129m}Xe and ^{131m}Xe states have half lives of 8.9 and 11.8 days respectively. ^{129m}Xe de-excites with the emission of a 236 keV γ-ray and ^{131m}Xe de-excites with the emission of a 164 keV γ-ray. The half lives of the two de-excitation modes are such that they will return to ∼background level within days of activation and the energies of the γ-rays are far from the WIMP search region between 5-50 keV$_{\rm nr}$. Figure 7.2 shows the two activation peaks expected [170] and Figure 7.3 shows energy spectra both before and after the AmBe calibration. The activation peaks can be seen clearly.

The data taken from before the AmBe calibration run and the simulated background show good agreement (see Figure 7.4). There is a deviation from agreement seen at the level of the excitation peaks (164keV) but this is due to residual activation from a previous AmBe calibration taken during July.

Finally, the background can be compared to the levels seen in the FSR of ZEPLIN–III. Figure 7.5 shows the comparison between the two science runs. The FSR, dominated by the PMT array, shows an average single scatter event rate of (11.33±0.18) events/kg/keV$_{\rm ee}$/day whereas the SSR backgrounds shows an average of (0.58±0.07) events/kg/keV$_{\rm ee}$/day. This represents a 20x reduction in background between the first and second science runs.

Figure 7.5 shows an apparent decrease in the background reduction above 100 keV$_{\rm ee}$, but this is artificial. In the FSR, a saturation cut was applied in the ZEPLIN–III DAq system in order to reduce the amount of output data. Storage limitations were a factor in the application of the cut. Since this cut was applied in hardware, it is not possible to recover the data for the purpose of this

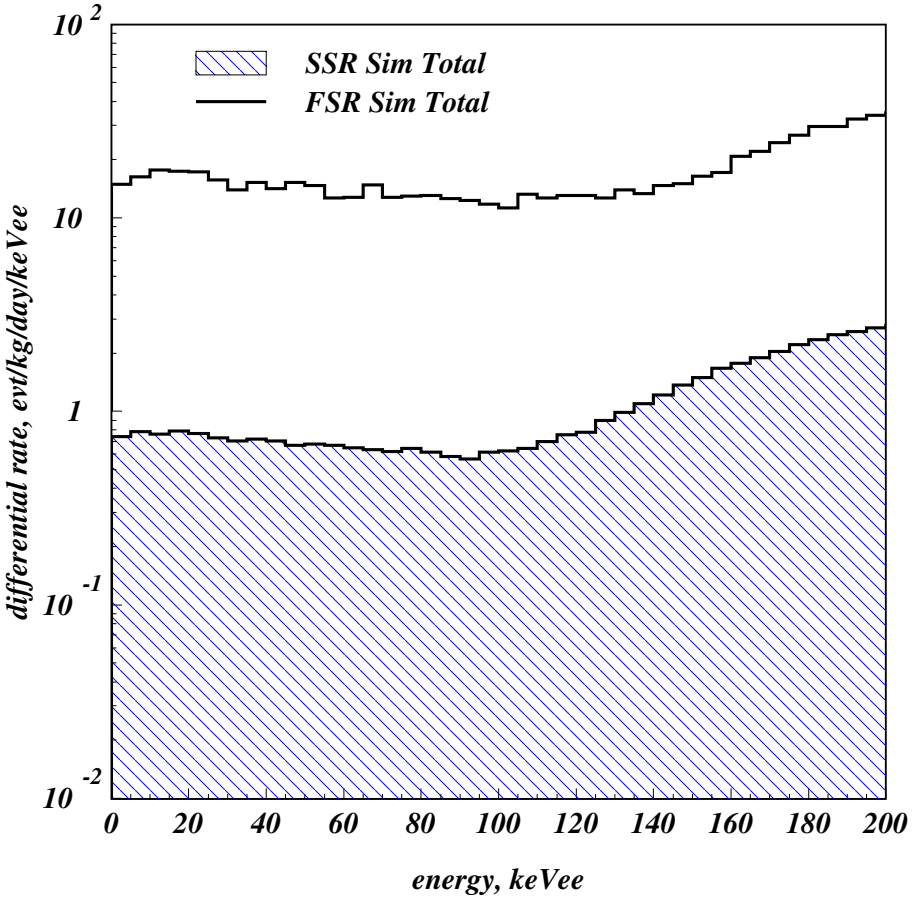

Figure 7.1: Comparison between simulations of expected electron recoil background in the FSR and SSR of ZEPLIN-III. In the FSR, the rate was dominated by the PMTs but in the SSR, the background rate from the PMTs is comparable to that of other sources.

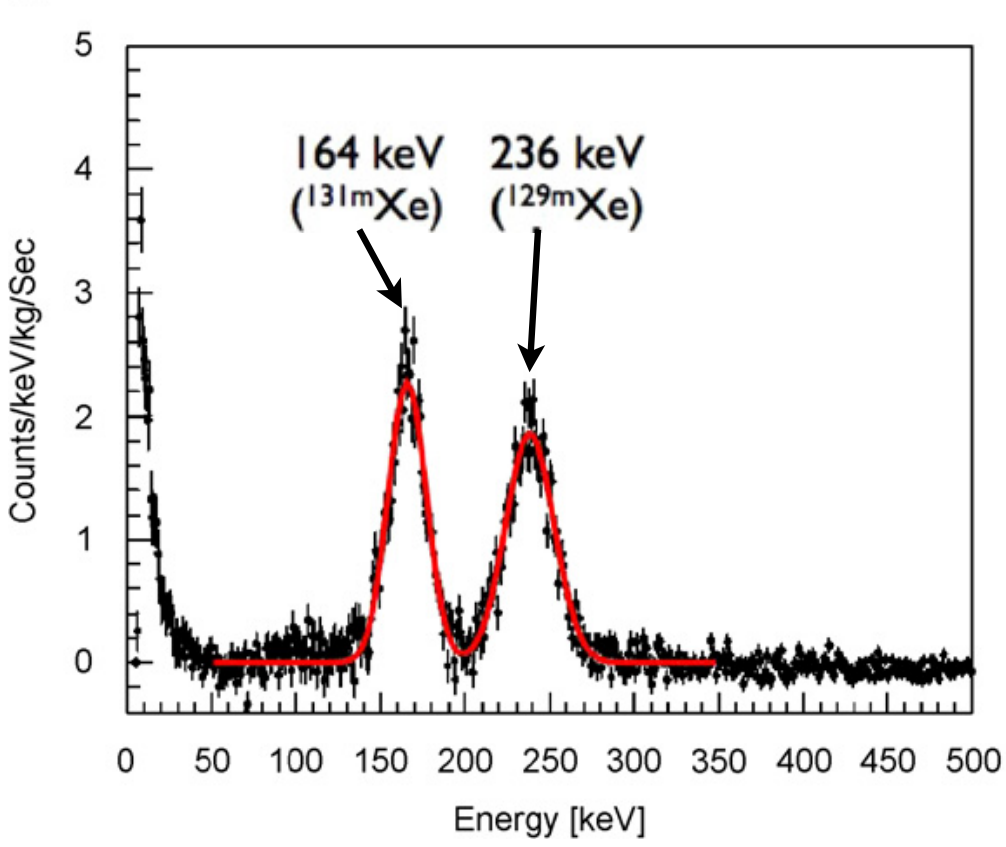

Figure 7.2: An excitation spectrum for a Xenon sample taken after 5 days of activation by a ^{252}Cf source. The 164keV and 236keV levels due to de-excitation γ-rays from ^{131m}Xe and ^{129m}Xe can be seen.

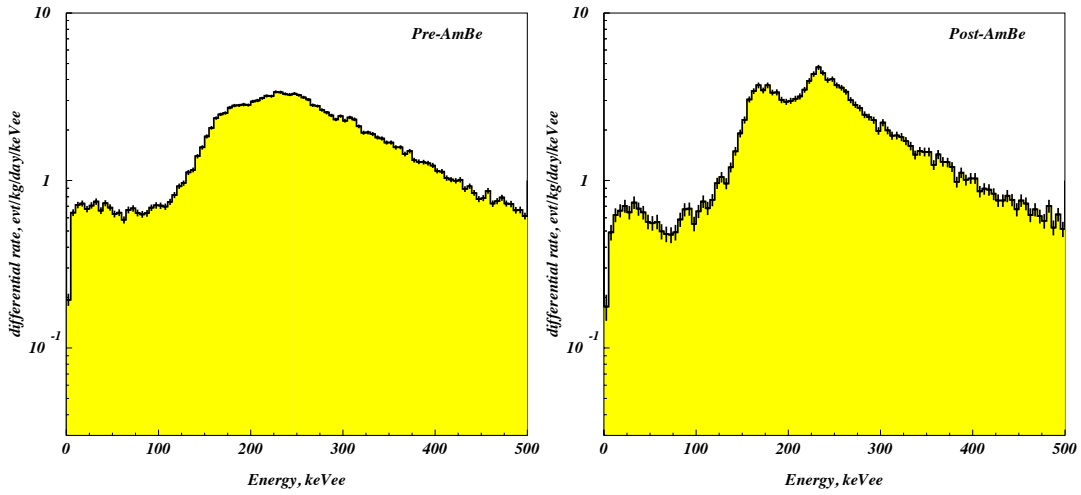

Figure 7.3: A comparison between data taken before (left) and after (right) the AmBe calibration of 100825. A clear increase can be seen at the 164keV and 236keV levels due to de-excitation γ-rays from ^{131m}Xe and ^{129m}Xe respectively.

plot. Furthermore, as described in Section 6.4.1, the veto tagging efficiency is in excellent agreement with simulation if contamination levels are adopted that match the background seen in ZEPLIN–III (shown in Figure 7.4).

7.1.2 Additional Upgrades

In addition to the installation of the new PMT array and the veto detector, several upgrades were made to ancillary systems. In order to ensure the reproducibility and consistency of daily ^{57}Co calibration runs was optimised, a source delivery system was retrofitted to the lead castle externally and to the ZEPLIN–III copper dome internally. This system, combined with an electronic insertion device, ensures that the ^{57}Co source is located in the same position for all daily calibrations.

A so-called "phantom-grid" was installed above the Xe volume in order to further improve the calibration and position reconstruction of ZEPLIN–III data. The phantom grid is shown in Figure 7.6. A dedicated ^{57}Co run was taken in order to fully characterise the phantom grid for use in the position reconstruction

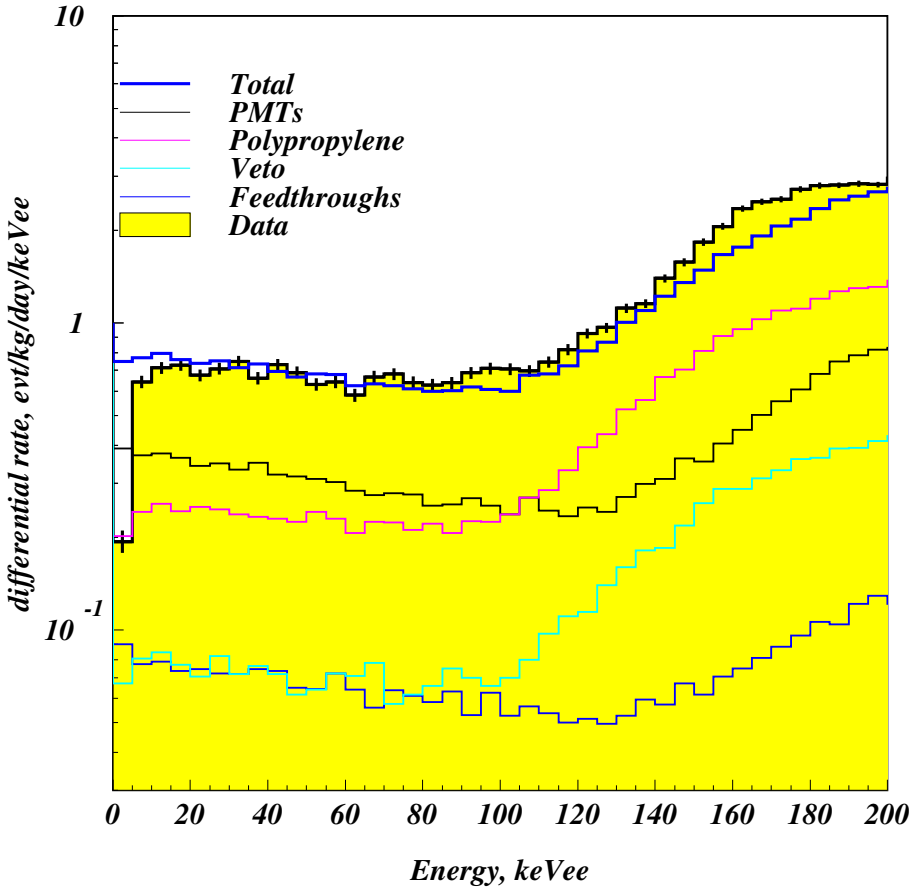

Figure 7.4: A comparison of data and simulated electron recoil background in the ZEPLIN–III SSR. The data is represented by the yellow histogram and the simulation is represented by the thick blue line. An excellent agreement can be seen up to the point where the ^{131m}Xe de-excitation γ-ray causes a rate increase. The increased rate is residual activation from an AmBe calibration performed in July.

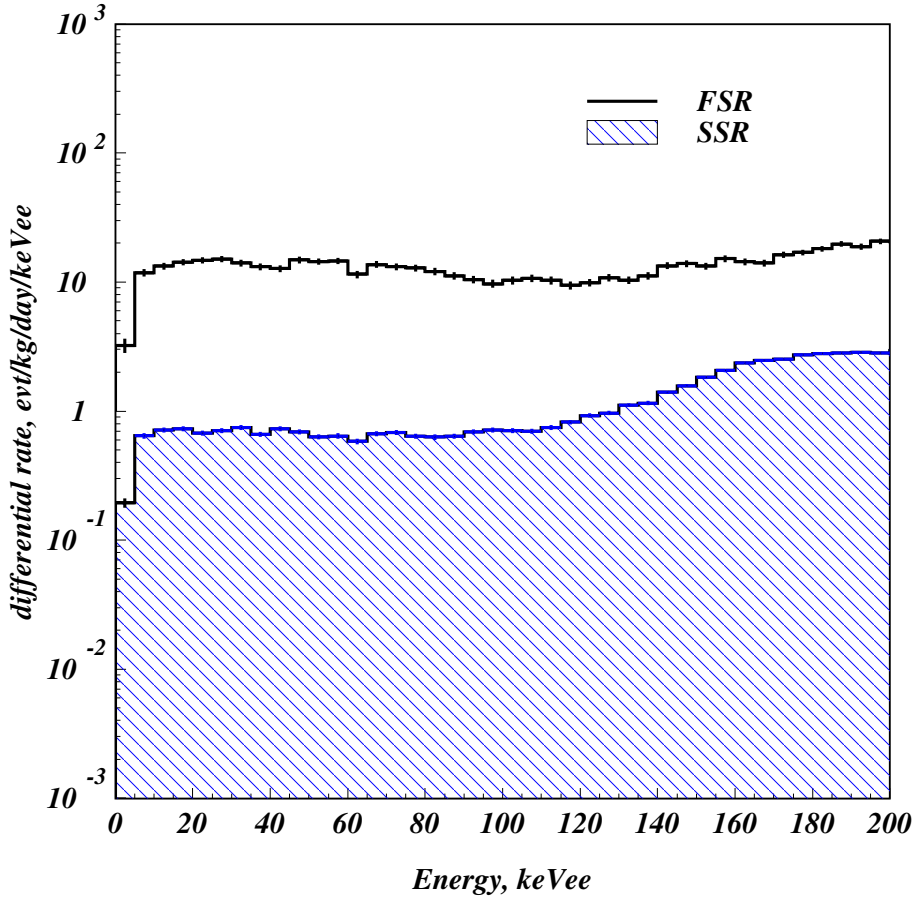

Figure 7.5: Comparison of true backgrounds from the FSR and SSR of ZEPLIN–III. The black line represents the FSR simulation and the shaded blue region the SSR simulation. The backgrounds begin to converge above $\sim$100 $\mathrm{keV_{ee}}$ due to the hardware saturation cut applied to reduce the data rate in the FSR. Due to the lower overall rate, no such saturation cut is applied in the SSR.

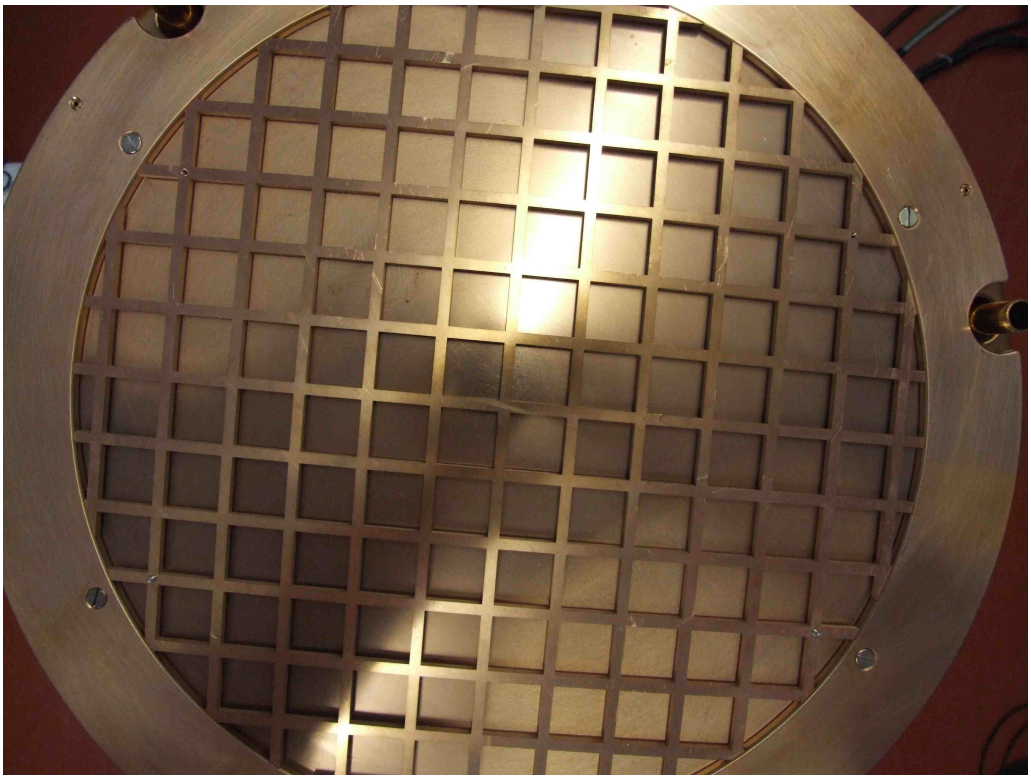

Figure 7.6: The phantom grid installed above the ZEPLIN-III Xe volume. The grid is placed in an exact position to allow comparison of ^{57}Co data with simulation.

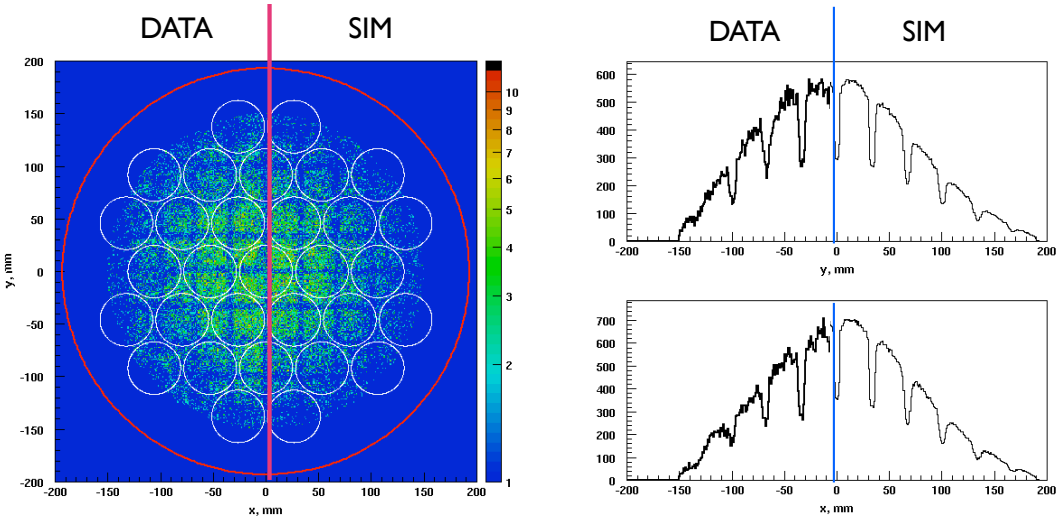

Figure 7.7: A comparison of data and simulation for the phantom grid ^{57}Co data. In the xy-plane (left), the difference between data and simulation is almost indiscernible and the agreement in the individual x and y distributions (right) is excellent.

analysis. A comparison of ^{57}Co data versus simulation is shown in Figure 7.7 which shows the agreement between simulated and observed results. It is the simulated result that is used to define the response of the PMT array as discussed in Section 4.3.5. The phantom alone gives a position sensitivity of ∼2 mm in the xy-plane which is further improved by the position reconstruction algorithm.

Finally, the data acquisition and calibration mechanism was fully automated. This allows for the optimisation of the detector duty cycle and minimises human interference which may alter the operational parameters of the detector. ZEPLIN–III ran with a duty cycle of 96% which is a 12% improvement over the FSR. The lost 4% is taken up by the daily ^{57}Co runs and the weekly PMT LED calibration (described in [171]).

The upgrade to ZEPLIN–III was completed and on 26$^{\text{th}}$ June 2010, the SSR began. This run lasted until 14$^{\text{th}}$ January 2011, giving 190 days (with calibration days removed) of science data. The stability monitoring and calibration of the SSR of ZEPLIN–III followed exactly the same method as the FSR (see Chapter 4) so no discussion of the methods is found in this chapter. The remainder of

this chapter will discuss the cuts used in and efficiencies calculated for the SSR and the calculation of a projected, preliminary WIMP-nucleon scattering cross-section. This result is preliminary because, at the time of writing, analysis of the full dataset is incomplete.

7.2 Golden Analysis

Although the background from the SSR PMT array is 20 times lower than that of the FSR, they are prototype tubes. The array is much less uniform than that found in the FSR, with quantum efficiencies ranging from 16.4% to 35.6% and gain responses varying by factors of >10. The varying response of the PMTs causes erroneous results in position and energy reconstruction and more sophisticated cuts are needed to remove these.

7.2.1 Setting a WIMP Search Region

As in the FSR, the WIMP search region is defined using an AmBe calibration. The elastic recoil population is selected, sliced into energy ranges and fitted. Figure 7.8 shows a selection of these energy slices that are fitted using a Gaussian distribution.

Having calculated the mean and widths of the distribution, the resultant values can be plotted and fitted to give the description of the elastic recoil mean and width. Figure 7.9 shows the mean and sigma measurements and Figure 7.10 shows the resultant WIMP search box plotted from 0-16 $\mathrm{keV_{ee}}$ and extending from the mean (μ) of the distribution to $\mu - 2\sigma$. In addition, a larger box is defined from 0-40 $\mathrm{keV_{ee}}$ and for the S2/S1 range of $\mu^{+\sigma}_{-3\sigma}$. This larger box is defined as part of the unblinding process. Once data quality cuts have been defined, all data outside the large box will be unblinded. If no problems are seen then the large box (minus the small box) is unblinded. Finally, if everything is as expected in the large box, the small box (WIMP search region) is unblinded.

With the addition of the veto, the unblinding of the data followed a different path to that chosen in the FSR. Instead of using 10% of the data for the definition of the cuts, all data tagged by the veto - (28.0±0.2)% from prompt tags and 0.8% accidental delayed tags - was unblinded. This represented about 30% of the data

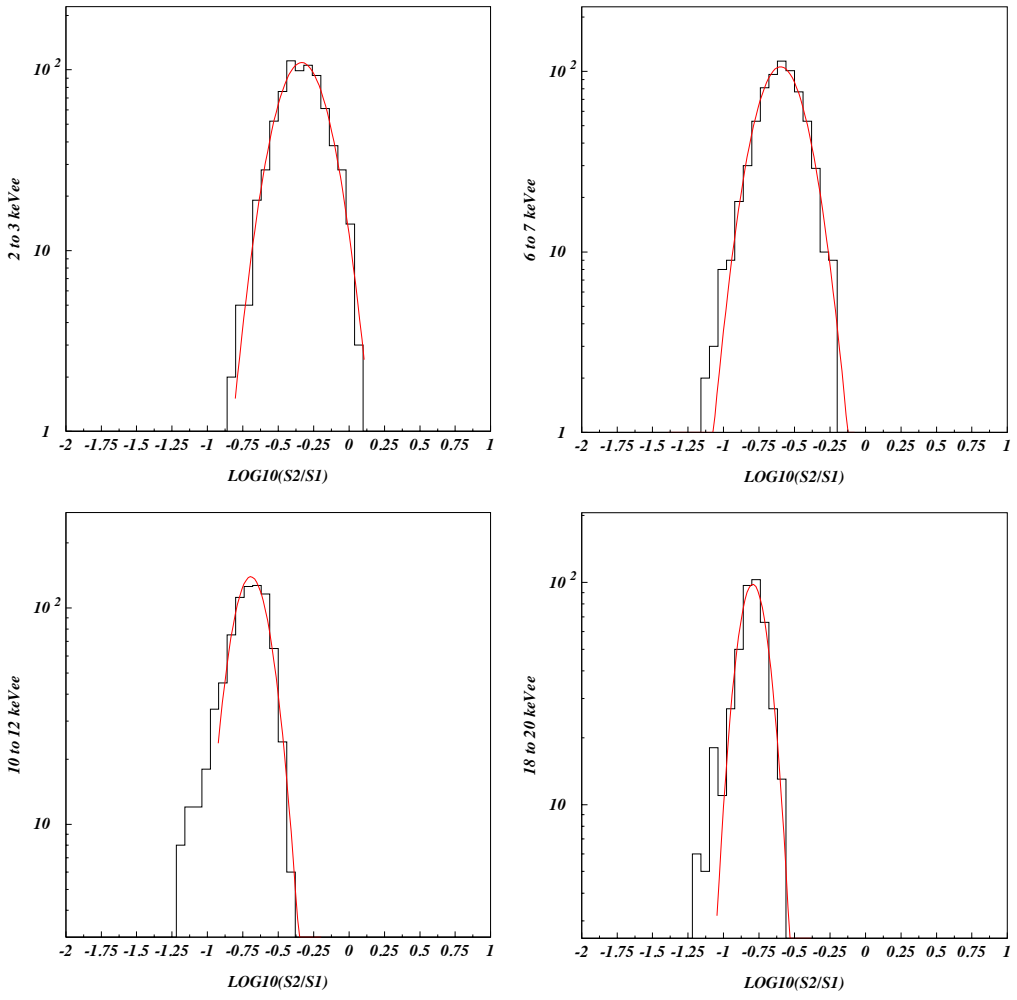

Figure 7.8: Fits to energy ranges in the AmBe elastic recoil population fitted with Gaussian distributions (the fit ranges are selected by hand). The size of energy slices increases at higher energies to account for a lower number of events.

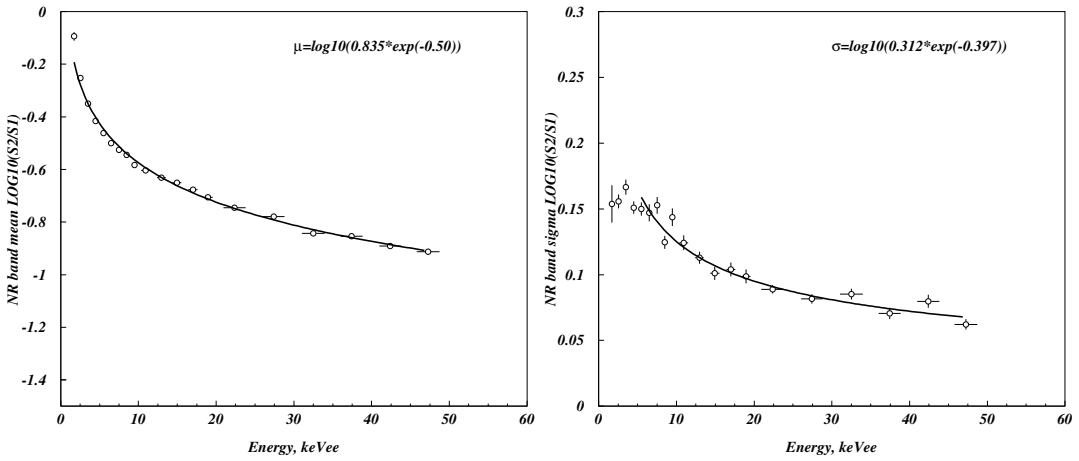

Figure 7.9: The results of the AmBe energy slice fitting are plotted. Fits to the data give the parameters for the mean and standard deviation, σ, of the elastic nuclear recoil population.

set where the events selected would not be consistent with WIMP recoils. It would be possible to unblind all data outside the WIMP search region but the development of cuts may lead to small changes in the distribution of the AmBe elastic recoil population and changes in the energy scales, so this has not yet been done.

All parameters associated with the tagging of ZEPLIN–III events by the veto are fixed and, as such, it was safe to unblind these data without fear of jeopardising the blindness of the WIMP search region since a WIMP recoil will not be tagged by the veto.

7.2.2 Data Quality Cuts

In the SSR, the fiducial cuts on the Xe volume are slightly less stringent than were implemented in the FSR. In the z, or drift-time, axis, a range of $(0.5 < dtime < 14)$ μs was used, which equates to a distance of 32 mm, as well as a radial cut of 160mm. This equates to a fiducial volume of 7.6 kg. The 190 days of data and $\sim$96% duty cycle leads to a fiducial exposure of 1312 kg.days.

The vetoed data was plotted and the box defined in Section 7.2.1 was applied.

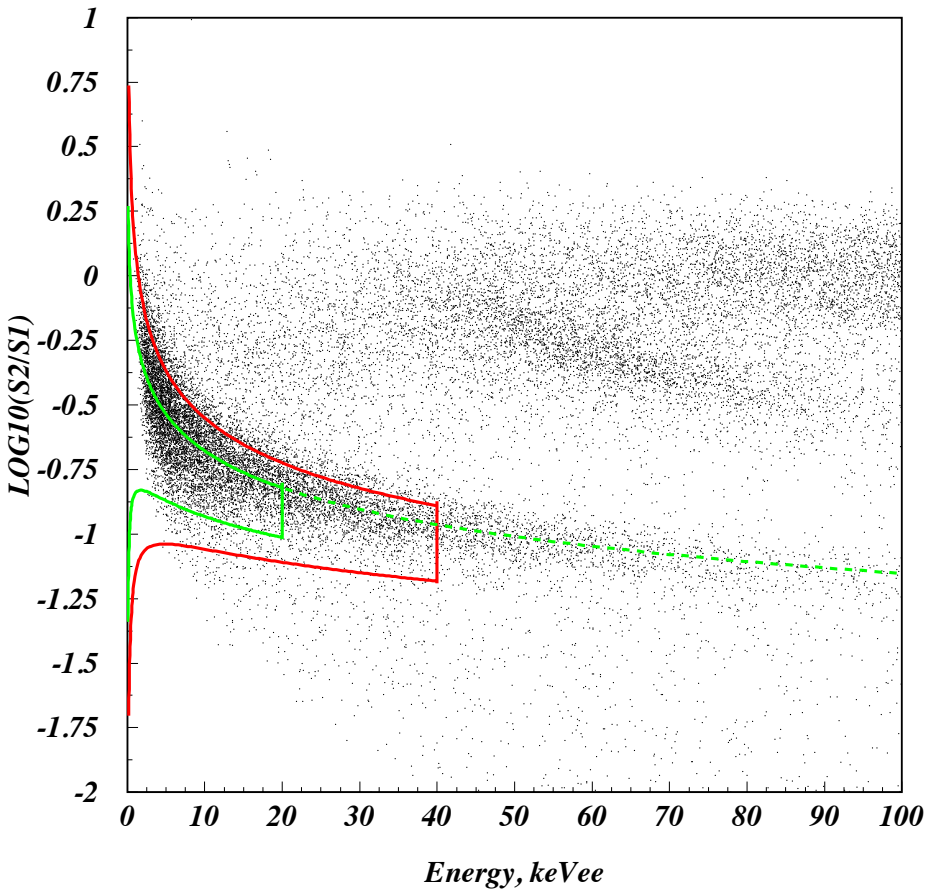

Figure 7.10: The WIMP search region superimposed on the AmBe calibration data. The green line represents the mean while the green box represents the WIMP search region. The red box is defined as part of the unblinding procedure.

Further quality cuts were then defined. Most of the applied cuts only removed a few outlying events but a particularly powerful cut was found using the energy dependence of the $s2rmsm$ parameter. Figure 7.11 shows the distribution of the AmBe data and the cut applied which removes unwanted events from inside the WIMP search region without removing a large proportion of the AmBe elastic nuclear recoil population. The cut removes $\sim$25% of events in the WIMP search region with a loss of only 6% of the AmBe elastic recoil band. The energy dependent efficiency of the cut is shown in Figure 7.12. The efficiency is linear with energy in the region of interest

After the application of additional cuts, the number of events in the WIMP search region was reduced to zero as seen in Figure 7.13. For the remainder of this chapter, a quality cut efficiency similar to that seen in the FSR (70%) is assumed. It is anticipated that the final efficiency will be greater than this when analysis and quality cuts are finalised. The preliminary results thus far, presented in Figure 7.13, highlight the quality of the SSR data.

7.3 Setting a Limit

Given the fact that the vetoed data is clean and represents $\sim$ 30% of the data in the full ZEPLIN–III data-set, it is possible to calculate a projected limit if the unblinding procedure does not reveal any events above expectation in the WIMP search region. As analysis of some of the key parameters in the calculation is still ongoing, some assumptions have to be made.

The first assumption is that the WIMP search region will contain the same number of events as the expectation value calculated from analysis of the electron recoil population which would make the data set consistent with 0 observed events. If the data shows, as in the FSR, that the number of expected events is greater than the number observed, it is reasonable to reduce the number of expected events as this will result in a conservative upper limit. If the number of expected events is greater than observed, Feldmann-Cousins analysis will give an artificially low limit. Using a standard Feldman-Cousins statistical method (see Section 4.11), 0 observed events is equivalent to an upper limit of 2.44 events at a 90% confidence level.

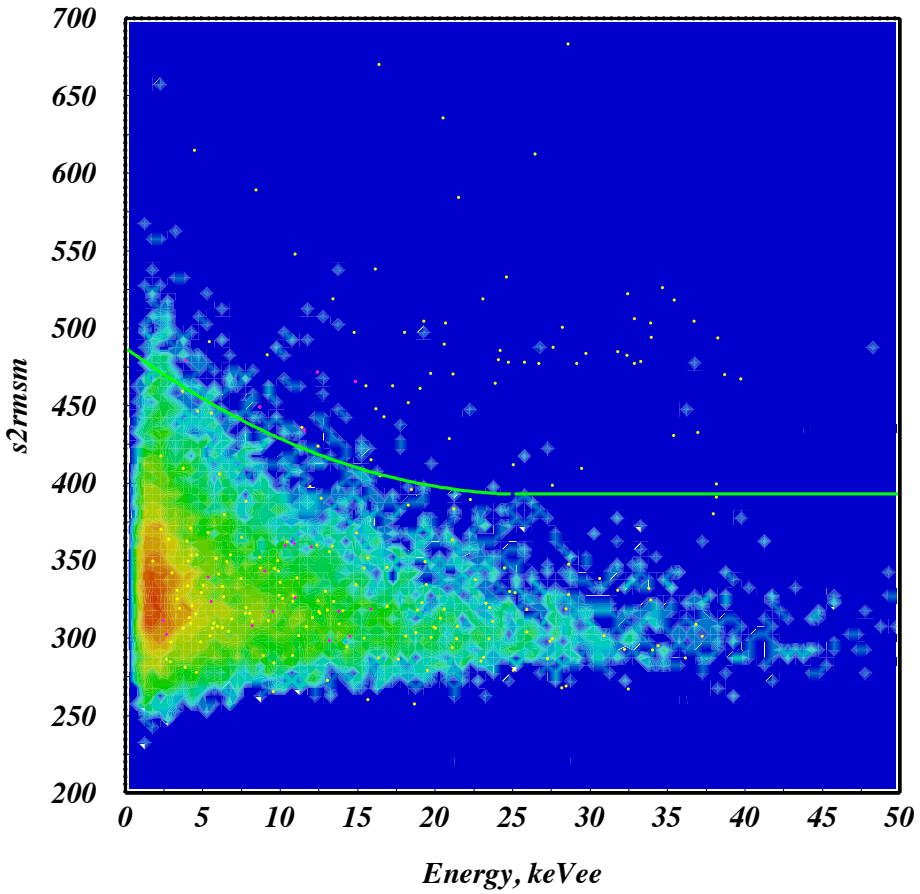

Figure 7.11: The $s2rmsm$ distribution as a function of energy in AmBe data. Superimposed on top are leakage and MSSI events that fall close to or in the large box as described previously. The green line represents the cut applied, above which events are rejected as poorly parameterised. No other cuts are applied.

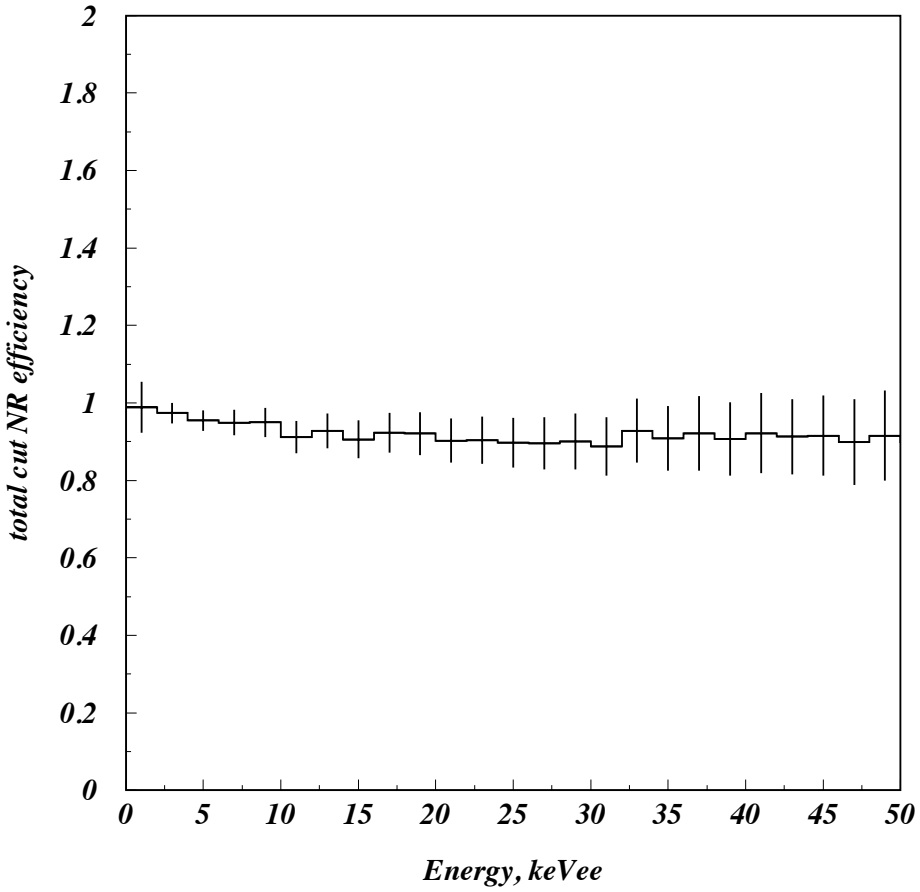

Figure 7.12: The energy dependent efficiency of the $s2rmsm$ parameter cut. The *"total cut NR efficiency"* represents the fraction of events left in the AmBe elastic nuclear recoil population after the cut has been applied.

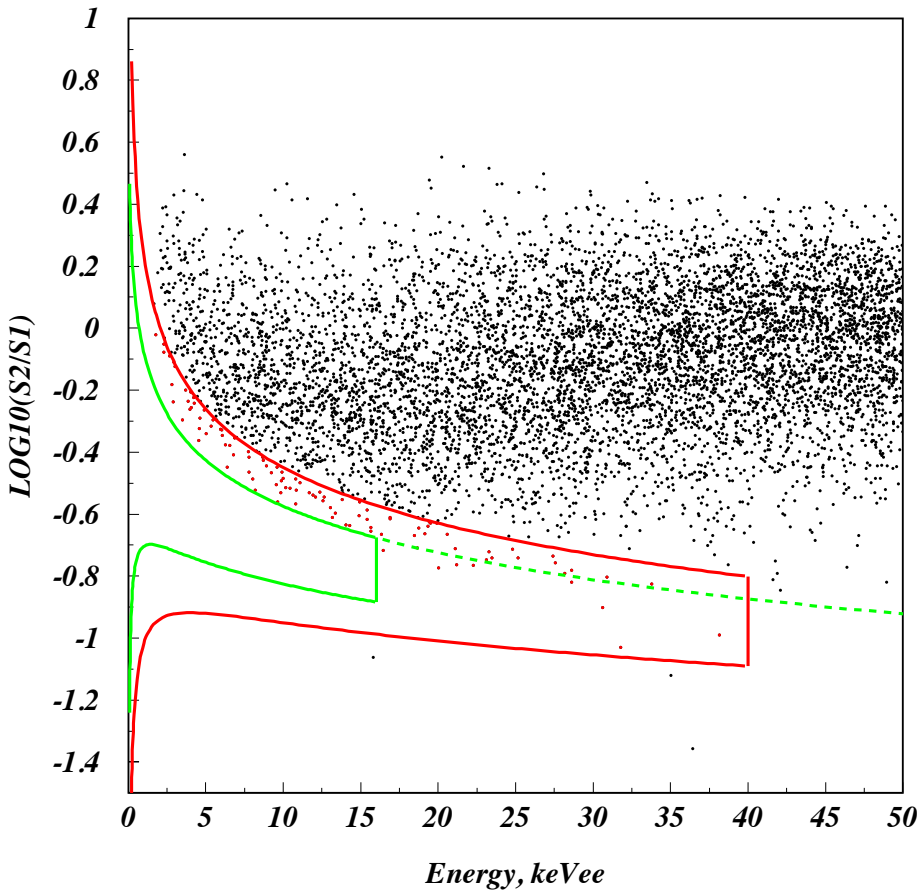

Figure 7.13: A post fiducialisation discrimination plot using preliminary cuts to remove all events that fall within the WIMP search region. This shows that, even without optimisation of cuts, all vetoed events can be removed from the WIMP search region.

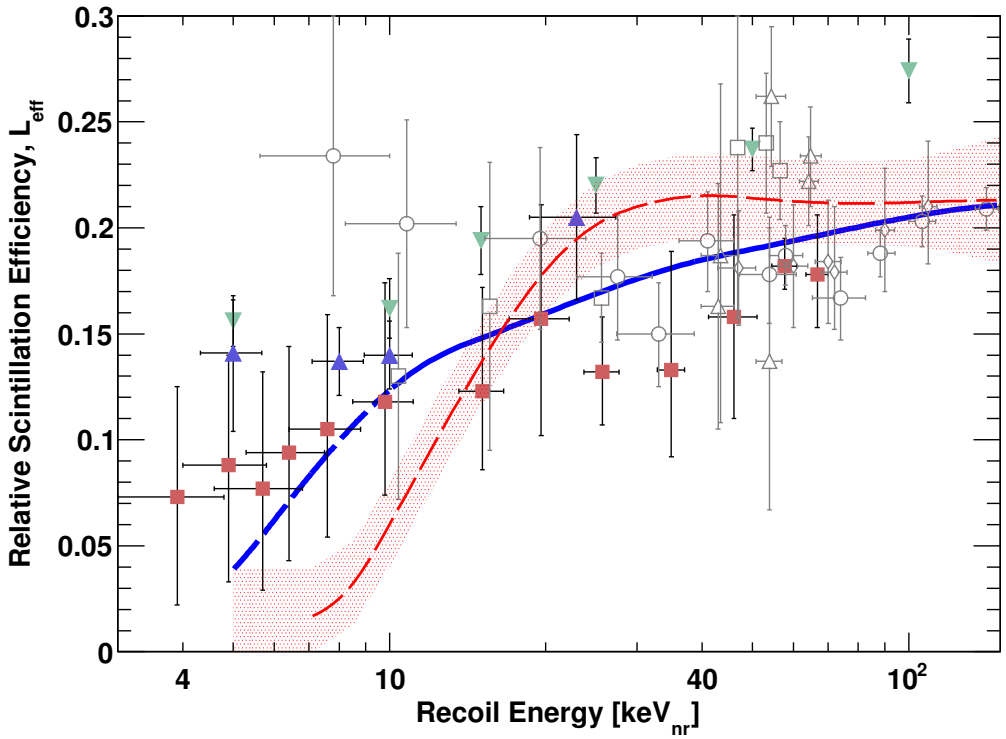

Figure 7.14: The measured L_{eff} curve for the SSR is shown in blue. The curve shows a less extreme fall off than measured for the FSR (red).

7.3.1 Relative Scintillation Efficiency

For the SSR, a new relative scintillation efficiency is calculated. The curve for the FSR showed the same characteristic fall as measurements by other experiments but was much more severe (see Figure 3.13). The analysis of AmBe data in the SSR did not show quite as extreme a deficit between data and simulation at low energies and this was found to be due to a systematic error that was removing low energy events due to mis-identification of noise pulses as true signals.

The efficiency calculated for the data quality cuts can be combined with the fiducialised exposure of 1312 kg.days to give a final corrected exposure of 918.4 kg.days. This assumes that all the other efficiencies listed in table 4.3 increase to 100%. The detector dead time, due to the 10x reduction in rate, does, in-fact increase to ∼100% and improvements in the reconstruction and pulse finding algorithms has pushed the other efficiencies up to close to 100%.

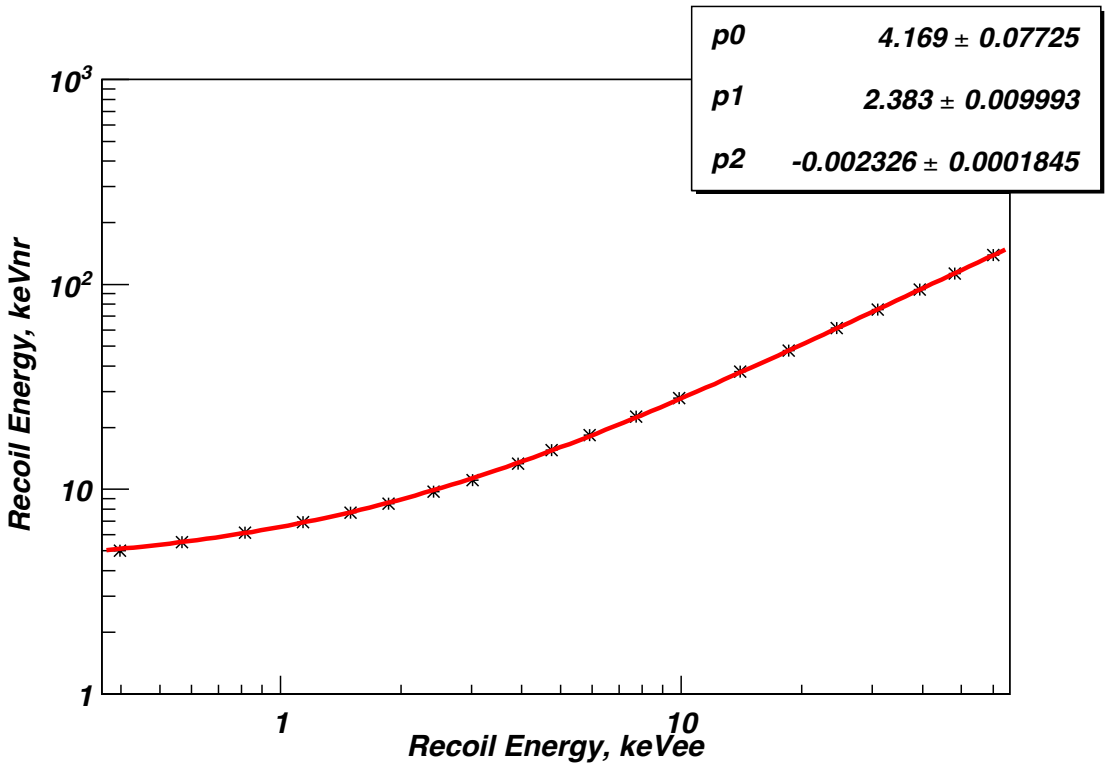

Figure 7.15: The energy conversion between $\mathrm{keV_{ee}}$ and $\mathrm{keV_{nr}}$ used in the SSR. This conversion is calculated using the L_{eff}. A 2^{nd} order polynomial fit ($E = p_0 + p_1 x + p_2 x^2$) describes the data. The fit parameters are shown in the statistics box.

7.4 Calculating Results

Given the efficiencies and the energy conversion factors, it is possible to calculate projected limits for the spin-independent WIMP-nucleon cross section. The spin-independent WIMP-nucleon cross-section is calculated to be 1.52×10^{-8} pb for a 55 GeVc^{-2} WIMP. This would provide the most sensitive measurement of the spin-independent WIMP-nucleon cross-section, surpassing the current most sensitive measurement by the XENON100 experiment with a cross-section upper limit of 3.3×10^{-8} pb for a 55 GeVc^{-2} WIMP. Figure 7.16 shows the comparison between this projected result and the nearest competitive results. The projected result would be the first to probe the 95% allowed region for a WIMP signal in the mSUGRA model.

Given that ZEPLIN–III was originally meant to run for 2 years (delivery of the PMTs by ETEL was severely delayed), it is possible to predict the limit reached if it had been possible for this length of run to be performed. A 190 day run represents ∼25 % of a full 2 year run and, as such, it is reasonable to expect that, given identical operational parameters, an exposure of ∼3700 kg.days (given a 70% data quality cut efficiency) would be recorded. Again, assuming that the number of events seen in the box matches the expected number, a 2 sided 90% confidence level upper limit on the number of WIMP induced nuclear recoil signal events in the box of 2.44 could be defined. These assumptions lead to a projected upper limit on the spin-independent WIMP-nucleon cross-section of 3.79×10^{-9} pb for a 55 GeVc^{-2} WIMP.

This calculated limit does not represent the published limit for the SSR, merely a projection of what may be achieved given a basic analysis of the data and a few key assumptions. It is expected that the full dataset will be unblinded shortly, after which a more rigorous analysis may be performed. In addition to this, improved limits on the spin-dependent cross-sections for the WIMP-neutron and WIMP-proton cross-sections will be calculated and published. Of course, this assumes a limit from a null result. ZEPLIN-III is a discovery device and has a sensitivity that begins to probe the favoured parameter space for a WIMP signal. The vetoed data necessarily excludes any WIMP signal so it is entirely possible that the unblinding of the full SSR dataset may provide a positive signal.

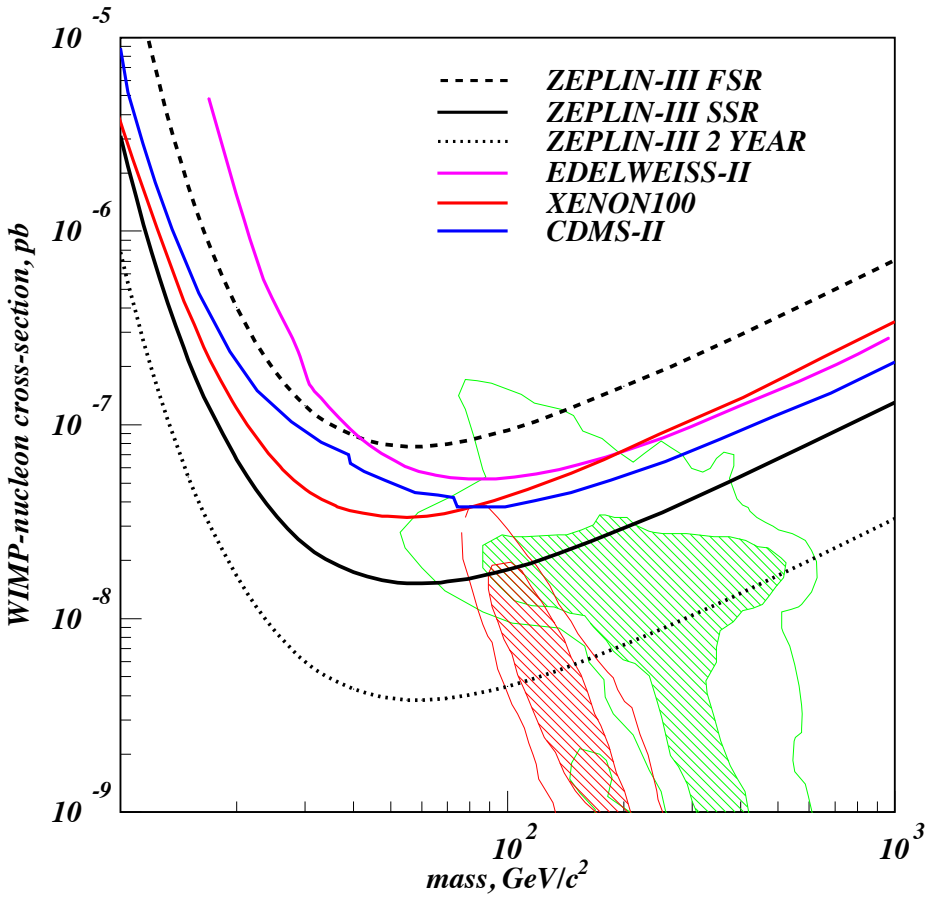

Figure 7.16: The projected result from the SSR of ZEPLIN–III (solid black curve). The result would represent the most sensitive measurement of an upper limit on the spin-independent WIMP-nucleon cross-section so far and would be the first to probe the 95% allowed region for a WIMP signal in the mSUGRA model (green hatched region). The detector is also capable of probing the 95% allowed region under the CMSSM model (red hatch). Also included is a projected limit for a 2 year run of ZEPLIN–III.

Chapter 8

Conclusions

A wealth of evidence is available to support the hypothesis that the WIMP constitutes the dominant dark matter component in the universe. Three of the five most sensitive upper limits on the spin-independent WIMP-nucleon cross-section have been produced by two-phase xenon detectors probing signals through measurement of scintillation and ionisation. The ZEPLIN–III detector has the potential to detect dark matter for the first time, or to produce the world's most stringent limit to date.

The ZEPLIN–III experiment is based at the Boulby mine facility where 2800 m.w.e of rock provides a reduction in muon flux by a factor of $\sim 10^6$. In addition to this, the placement of the lab in a seam of low background salt and a combination of lead and polypropylene shielding reduces the γ-ray and neutron flux by a factor of $\sim 10^5$ giving a near ideal location for a rare event search experiment.

The first science run of ZEPLIN–III began in February 2008 and lasted for 3 months. During the analysis of the data, an important limitation of the ZEPLIN–III design was observed (although it had already been predicted). The main source causing leakage of nuclear recoil background were events with two scattering vertices, one of which occurs in a so-called "dead" region of the detector. This is a region in which no electric field operates or where the electric field is distorted meaning charge cannot be extracted from the liquid surface and accelerated through the gas phase to give a secondary, ionisation, signal. These events, coupled with an initial weakness in the position reconstruction algorithm, meant that the analysis of the first science run was performed un-blinded.

Despite this limitation, ZEPLIN–III produced a competitive spin-independent 90% confidence level upper limit of 8.1×10^{-8} pb for a 60 GeVc^{-2} [83]. In the the spin-dependent channel, ZEPLIN–III produced a limit of 1.9×10^{-2} pb which exceeded the previous most sensitive spin-dependent result for an upper limit on the WIMP-neutron cross-section. In addition to this, the first science run yielded an upper limit on the WIMP-proton cross section of $4.1 \times 10^{+2}$ pb [140].

After the first science run, ZEPLIN–III entered an upgrade period. This included the installation of a new PMT array. The PMT array gave almost a 20x reduction in the electron recoil background whilst reducing the single-scatter nuclear recoil background (in the WIMP search region) to 0.6 neutrons per year. To further reduce this background, a tonne scale plastic scintillator veto detector was installed around the Xe volume. The veto detector is measured to have a neutron tagging efficiency of $\sim$61% which allows the 0.6 neutrons per year emitted from ZEPLIN–III components to reduce to 0.2 un-vetoed events per year.

The veto [162] also provides the ability to tag prompt γ-rays associated with electron-recoil backgrounds in ZEPLIN–III which yields two important results. The first of these is that the background is reduced. The second, and more important, result is that $\sim$30% of the data can be unblinded without the risk of unblinding potential WIMP signal events. This allows the development of data quality cuts to remove spurious MSSI events that reduce discrimination and (in the un-vetoed 60% of the data) can even give events that appear to be consistent with a WIMP recoil. Fitting the electron recoil population of the vetoed events can also be used to predict the number of electron recoil events that will be seen in the WIMP search region, which is another key feature. This is all achieved with no loss of exposure.

In addition to the installation of the veto, the PMT array of ZEPLIN–III was upgraded. The PMT array was developed in collaboration with ETEL during the first science run of ZEPLIN–III with all components being radio-assayed and simulated in order to test the impact each would have on the final backgrounds in the second science run of ZEPLIN–III. As with the installation of the veto, the low intrinsic radioactive background of the upgraded PMT array serves more than one purpose. The first of these is that the nuclear recoil background is reduced. This means that ZEPLIN–III could, in theory, run for two years before any WIMP-recoil signal sensitivity becomes background limited.

The second advantage of the improved PMT array is the reduction of the γ-ray induced electron recoil background. This reduction improves the discrimination of ZEPLIN–III and also reduces the rate of MSSI events. In addition to this, the reduction in background leads to an overall reduction in the intrinsic trigger rate of the ZEPLIN–III detector. The 90% reduction with respect to the first science run reduces the dead time of the system to near 0 and reduces the data storage needs by 90%.

Matching the measured electron recoil background to that expected from simulation yields a fantastic agreement (independently corroborated by the veto). This confirms that the radio-assaying and simulation of the components used in the upgraded PMT array, the veto detector and the ceramic high voltage feedthroughs gave correct results and that the backgrounds in ZEPLIN–III are well understood.

Having run for $\sim$6 months (long before the system has the potential to become background limited), the second science run of ZEPLIN–III ended. The main challenge in the analysis of the second science run data was the fact that the upgraded PMT array, although reducing the backgrounds to a rate 20x less than the first science run, was comprised of prototype tubes. Unfortunately, this meant that the response of the array in regards, for example, gain or light yield was not particularly uniform. This non-uniformity made the development of the position and energy reconstruction algorithm more difficult than it had been in the first science run but this was helped, in part, by the installation of a "phantom grid" above the Xe target volume. This grid is used to cast a shadow on the ^{57}Co data which, through comparison to simulated data, helps to define the expected signal response of the detector.

After the run was complete, the vetoed data was analysed in order to develop data quality cuts to remove MSSI events. A fully empty box was produced with an efficiency greater than 70%, defined using an on-off application of the cuts to AmBe calibration data. Using this conservative efficiency, a second science run fiducial exposure of 1312 kg.days was calculated. At the time of submission, this was the status of the analysis of the second science run data.

This being the case, a projected upper limit on the spin-independent WIMP-nucleon cross-section was calculated. The limit was calculated using the assumption that the full unblinded dataset would also contain 0 events in the WIMP

search region and that this would match the expected number of events from the analysis of the electron recoil background extrapolated into the WIMP search region.

This assumption, along with the application of a new relative scintillation efficiency curve yields a projected result for an upper limit on the spin-independent interaction of 1.52×10^{-8} pb for a 55 GeVc^{-2} WIMP. This result would represent the most sensitive measurement of the WIMP-nucleon scattering cross-section to date and would be the first result to probe the 95% allowed region for a WIMP signal in the mSUGRA model of dark matter. The same assumptions used to calculate the limit for the 190 day run can also be used to calculate a projected limit in the case that ZEPLIN–III was able to run for 2 years. In the absence of any neutron induced nuclear recoils, an upper limit of 3.79×10^{-9} pb can be calculated, again, for a 55 GeVc^{-2} WIMP. These projected results suggest that, although the upgraded PMT array has caused problems, ZEPLIN–III has the potential to reach the sensitivity predicted during its design phase. However, it is possible that unforeseen efficiency losses may increase these upper limits somewhat. It is pertinent to add that the unblinding of the data may reveal a signal above the expected background consistent with WIMP recoils as the projected sensitivity allows for this.

In the event of a null result, the future for the development of 2-phase Xe detectors is bright. The XENON100 experiment is currently running in the Gran-Sasso laboratory in Italy. With a $\sim$40 kg fiducial volume, this promises a sensitivity of below 1×10^{-9} pb. In addition to this, the LUX collaboration are in the process of characterising a detector with a 350 kg total volume of Xe. This is currently undergoing surface tests at the Homestake facility, South Dakota, USA, with a view to being installed underground in the near future. Finally, the XMASS detector is operating in the Kamioka facility in Japan, though it is not entirely clear what stage this single phase detection system is at in terms of yielding a result.

There are also tonne scale detectors in both the design and development stages. XENON1T appears to be the closest to being constructed, but other detectors, from the LUX collaboration and others, are also under development. It seems that 2 phase liquid noble gas, especially Xe, detectors are those favoured to give a possible discovery of dark matter through WIMP-nucleon scattering.

References

[1] C. Seife (2005). *Science*, **309**(5731):78.

[2] W. Freedman, *et al.* (2001). *The Astrophysical Journal*, **553**:47–72.

[3] A. Penzias, *et al.* (1965). *The Astrophysical Journal*, **142**:419–421.

[4] A. Einstein (1916). *Annalen der Physik*.

[5] F. Zwicky (1933). *Helvetica Physica Acta*, **6**:110–127.

[6] P. Marshall, *et al.* http://arxiv.org/abs/0710.0637.

[7] T. Kundic, *et al.* (2004). *The Astrophysical Journal*, **607**:665.

[8] D. M. Chen (2004). *Astronomy & Astrophysics*, **418**:387–392.

[9] M. Persic, *et al.* (1996). *Monthly Notes of the Royal Astronomical Society*, **281**:27.

[10] V. C. Rubin, *et al.* (1980). *The Astrophysical Journal*, **238**:471.

[11] V. C. Rubin, *et al.* (1985). *The Astrophysical Journal*, **289**:81.

[12] Y. Sufue, *et al.* (1999). *The Astrophysical Journal*, **523**:136–146.

[13] M. Milgrom (2002). *Scientific American*, **52**:42–50.

[14] R. H. Sanders (2007). *Monthly Notes of the Royal Astronomical Society*, **380**:331.

[15] M. Tegmark, *et al.* (2006). *The Physical Review D*, **74**(123507).

[16] C. Skordis, *et al.* (2006). *The Physical Review Letters*, **96**(011301).

[17] V. G. Gurzadyan, *et al.* ArXiv:1011.3706v1.

[18] I. K. Wehus, *et al.* ArXiv:1012.1268.

[19] A. Moss, *et al.* ArXiv:1012.1305.

[20] A. Hajian. ArXiv:1012.1656.

[21] D. Clowe, *et al.* (2006). *The Astrophysical Journal*, **648**:L109–113.

[22] G. W. Angus, *et al.* (2008). *Monthly Notes of the Royal Astronomical Society*, **383**:417.

[23] A. G. Riess, *et al.* (2007). *The Astrophysical Journal*, **659**(1).

[24] G. Gamow (1946). *The Physical Review*, **70**:572–573.

[25] G. Gamow (1948). *The Physical Review*, **73**:803–804.

[26] E. Hubble (1929). *Proceedings of the National Astronomical Society*, **15**:168.

[27] NASA WMAP site. URL `http://map.gsfc.nasa.gov/`.

[28] C. L. Bennett, *et al.* (2003). *The Astrophysical Journal Supplement*, **148**:1.

[29] G. Hinshaw, *et al.* (2009). *The Astrophysical Journal Supplement*, **180**(2):225–245.

[30] M. R. Nolta, *et al.* (2009). *The Astrophysical Journal*, **180**:296–305.

[31] P. J. E. Peebles (1993). *Principles of Physical Cosmology*. Princeton University Press.

[32] D. G. York, *et al.* (2000). *The Astronomical Journal*, **120**:1579–1587.

[33] D. J. Eisenstein, *et al.* (2005). *The Astrophysical Journal*, **633**:560.

[34] D. N. Schramm, *et al.* (1998). *Reviews of Modern Physics*, **70**(1):303–318.

[35] P. Bonifacio, *et al.* (2003). *Astronomy & Astrophysics*, **361**:2427–2434.

[36] S. Perlmutter, *et al.* (2003). volume 598. Springer.

[37] J. L. Tonry, *et al.* (2003). *The Astrophysical Journal*, **598**:102.

[38] R. A. Knop, *et al.* (2003). *The Astrophysical Journal*, **598**:102.

[39] A. G. Riess, *et al.* (2001). *Proceedings of the Third International Workshop on the Identification of Dark Matter*. World Scientifice.

[40] A. G. Riess, *et al.* (2004). *The Astrophysical Journal*, **607**:665.

[41] M. Kowalski (2008). *The Astrophysical Journal*, **686**:749.

[42] P. Taylor (1994). *The Hidden Universe*. Praxis Publishing.

[43] N. Palanque-Delabrouille (2001). *New Astronomy Reviews*, **45**:395–399.

[44] P. F. Smith, *et al.* (1979). *Nuclear Physics B*, **149**:525.

[45] G. Jungmann, *et al.* (1996). *Physics Reports*, **267**:195.

[46] M. Davis, *et al.* (1981). *The Astrophysical Journal*, **250**:423–431.

[47] C. Krauss, *et al.* (2005). *European Physics Journal*, **C40**:447–468.

[48] J. Angrik, *et al.* (2004). Katrin Design Report 2004. Technical report, KATRIN Collaboration.

[49] S. Hannestad, *et al.* (2007). *Journal of Cosmology and Astroparticle Physics*, **8**:15.

[50] J. M. Overduin, *et al.* (2003). *Dark Sky, Dark Matter*. Institute of Physics Publishing.

[51] V. Springel, *et al.* (2005). *Nature*, **453**(7042).

[52] S. Weinberg (1978). *The Physical Review Letters*, **40**:223.

[53] F. Wilczek (1978). *The Physical Review Letters*, **40**:279.

[54] J. E. Kim (1979). *The Physical Review Letters*, **43**:103.

[55] M. A. Shifman, *et al.* (1989). *Nuclear Physics B*, **166**:493.

[56] A. R. Zhitnitsky (1980). *Soviet Journal of Nuclear Physics*, **31**:260.

[57] M. Dine, *et al.* (1981). *Physics Letters B*, **104**:1955.

[58] R. L. Davis (1986). *Physics Letters B*, **180**:225.

[59] R. L. Davis (1989). *Nuclear Physics B*, **324**:167.

[60] G. G. Raffelt (1997). *European School of High-Energy Physics*, pages 235–278.

[61] B. Morgan (2004). *Dark Matter Detection with Gas Time Projection Chambers*. Ph.D. thesis, The University of Sheffield.

[62] G. Bertone, *et al.* (2005). *Physics Reports*, **405**:279–390.

[63] P. Sikivie. *The Physical Review Letters*, **51**:1415.

[64] H. Peng, *et al.* (2000). *Nuclear Instruments and Methods in Physics Research A*, **444**:569–583.

[65] John Strologas – Supersymmetry. URL `http://www-cdf.fnal.gov/~strolog/Supersymmetry.html`.

[66] J. Ellis, *et al.* (2004). *Physics Letters B*, **588**:7.

[67] R. Barate, *et al.* (2001). *Physics Letters B*, **499**:67.

[68] J. Ellis, *et al.* (2003). *Physics Letters B*, **565**:176.

[69] W. de Boer, *et al.* (2006). *Physics Letters B*, **636**(1):13–19.

[70] N. Gehrels, *et al.* (1999). *Astroparticle Physics*, **11**:277–282.

[71] C. Winkler (2006). *New Astronomy Reviews*, **50**(7-8):530–533.

[72] J. Albert, *et al.* (2006). *The Astrophysical Journal*, **639**:761–765.

[73] O. Adriani, *et al.* (2009). *Nature*, **458**(7238):607–609.

[74] J. Chang, *et al.* (2008). *Nature*, **456**(7220):362–365.

[75] S. Profumo (2008). *arXiv:astro-ph/0812.4457*.

[76] M. C. Smith, *et al.* (2006). *Galaxy Evolution across the Hubble Time Proceedings*, (235).

[77] F. Donato, *et al.* (1998). *Astroparticle Physics*, **9**:247–260.

[78] F. J. Kerr, *et al.* (1986). *Monthly Notes of the Royal Astronomical Society*, **221**:1023–1028.

[79] X. Roca-Maza, *et al.* (2008). *The Physical Review C*, **78**(044332).

[80] R. J. Gaitskell (2004). *Annual Review of Nuclear and Particle Science*, **54**:315.

[81] E. Armengaud, *et al.* (2010). *Physics Letters B*, **687**:294–298.

[82] J. Angle, *et al.* (2009). *The Physical Review D*, **80**(115005).

[83] V. N. Lebedenko, *et al.* (2009). *The Physical Review D*, **80**(052010).

[84] J. Cooley, *et al.* (2010). *Science*, **327**(5973):1619–1621.

[85] P. F. Sorensen (2008). Ph.D. thesis, Brown University.

[86] E. Coccia (2006). *Journal Of Physics: Conference Series*, **39**:497–504.

[87] J. Morales, *et al.* (2005). In *The Identification of Dark Matter - 5th International Workshop*.

[88] P. Smith, *et al.* (1998). *Physics Reports*, **307**:275–282.

[89] B. Ahmed, *et al.* (2003). *Astroparticle Physics*, **19**:691–702.

[90] G. Alner, *et al.* (2005). *Physics Letters B*, **616**(1-2):17–24.

[91] G. Davies, *et al.* (1994). *Physics Letters B*.

[92] G. Alner, *et al.* (2004). *Nuclear Instruments and Methods in Physics Research A*, **535**:644–655.

[93] G. Alner, *et al.* (2005). *Nuclear Instruments and Methods in Physics Research A*, **555**:173–183.

[94] G. Alner, *et al.* (2005). *Astroparticle Physics*, **23**:444–462.

[95] H. Wang (1998). *Physics Reports*, **307**(263-267).

[96] H. Wang (1999). Ph.D. thesis, University Of California.

[97] D. Cline, *et al.* (2000). *Astroparticle Physics*, **12**:373–377.

[98] A. Howard, *et al.* (2001). In *Proceedings of the Third International Workshop on the Identification of Dark Matter*, page 457.

[99] D. Akimov, *et al.* (2003). Proceedings of the Fourth International Workshop on the Identification of Dark Matter. pages 371–376.

[100] M. Carson, *et al.* (2004). *Astroparticle Physics*, **21**(6):667–687.

[101] H. Araújo, *et al.* (2005). *Nuclear Instruments and Methods in Physics Research A*, **545**:398–411.

[102] C. Bungau, *et al.* (2005). *Astroparticle Physics*, **23**(1):97–115.

[103] G. Alner, *et al.* (2007). *Astroparticle Physics*, **28**:287–302.

[104] E. Aprile, *et al.* (2008). *The Physical Review Letters*, **100**(021303).

[105] B. Edwards (2009). *The ZEPLIN Dark Matter Search: Two Phase Xenon as a WIMP Target*. Ph.D. thesis, Imperial College London.

[106] H. Araújo, *et al.* (2005). *Nuclear Instruments and Methods in Physics Research A*, **545**:398–411.

[107] J. Jortner, *et al.* (2004). *The Journal of Chemical Physics*, **42**(12):4250–4253.

[108] ET Enterprises Ltd. URL `http://www.et-enterprises.com/`.

[109] M. Redshaw, *et al.* (2007). *The Physical Review Letters*, **98**(053003).

[110] M. Danilov, *et al.* (2000). *Physics Letters B*, **480**:12.

[111] K. Abe, *et al.* (2009). *Astroparticle Physics*, **31**:290–296.

[112] M. Martin (1971). *The Journal of Chemical Physics*, **54**:3289.

[113] S. Kubota, *et al.* (1978). *The Physical Review B*, **17**(6):2762.

[114] S. Kubota, *et al.* (1978). *The Journal of Physics C*, **11**:2645–2651.

[115] J. Linhard, *et al.* (1963). *Matematisk-Fysiske Meddelelser*, **33**.

[116] F. Arneodo, *et al.* (2000). *Nuclear Instruments and Methods in Physics Research A*, **449**:147.

[117] R. Bernabei (2001). *EPJ direct*, **3**:1.

[118] D. Akimov, *et al.* (2002). *Physics Letters B*, **524**:245.

[119] E. Aprile, *et al.* (2005). *The Physical Review D*, **72**(072006).

[120] V. Chepel, *et al.* (2006). *Astroparticle Physics*, **26**:58.

[121] E. Aprile, *et al.* (2009). *The Physical Review C*, **79**(045807).

[122] P. Sorensen, *et al.* (2009). *Nuclear Instruments and Methods in Physics Research A*, **601**:339.

[123] A. Hitachi (2005). *Astroparticle Physics*, **24**:247–256.

[124] A. Manzure, *et al.* (2010). *The Physical Review C*, **81**(025808).

[125] T. Doke, *et al.* (2002). *Japanese Journal of Applied Physics*, **41**:1583.

[126] A. Bolozdynya (1999). *Nuclear Instruments and Methods in Physics Research A*, **422**:314–320.

[127] E. Aprile, *et al.* (2006). *The Physical Review Letters*, **97**(081302).

[128] K. Yoshino, *et al.* (1976). *The Physical Review A*, **14**(1):438–444.

[129] E. Gushchin (1979). *Soviet Journal of Experimental and Theoretical Physics*, **49**:856.

[130] H. Araújo, *et al.* (2006). *Astroparticle Physics*, **26**:140–143.

[131] S. F. Biagi, *et al.* (1999). *Nuclear Instruments and Methods in Physics Research A*, **421**:234.

[132] P. Majewski (2010). Private Communication.

[133] D. Y. Akimov, *et al.* (2007). *Astroparticle Physicss*, **27**(46-60).

[134] H. Araújo, *et al.* (1998). *IEEE Transitions on Nuclear Science*, **45**(3):542–528.

[135] R. B. Murray, *et al.* (1960). *IRE Transitions on Nuclear Science*, **NS-7**:80–86.

[136] H. Araújo, *et al.* (2004). *Nuclear Instruments and Methods in Physics Research*, **A521**:407–415.

[137] G. J. Feldman, *et al.* (1998). *The Physical Review D*, **57**(3873).

[138] P. Smith, *et al.* (1996). *Astroparticle Physics*, **6**(1):87–112.

[139] P. Toivanen, *et al.* (2008). *Physics Letters B*, **666**(1).

[140] V. N. Lebedenko, *et al.* (2009). *The Physical Review Letters*, **103**(151302).

[141] J. Angle, *et al.* (2008). *The Physical Review Letters*, **100**(021303).

[142] Z. Ahmed, *et al.* (2009). *The Physical Review Letters*, **102**(011301).

[143] G. Alner, *et al.* (2005). In N. Spooner, editor, *Proceedings of IDM'05*, page 218.

[144] J. Angle, *et al.* (2009). *The Physical Review Letters*, **101**(091301).

[145] E. Behnke, *et al.* (2008). *Science*, **319**(933).

[146] A. Benoit, *et al.* (2005). *Physics Letters B*, **616**(17).

[147] H. S. Lee, *et al.* (2009). *The Physical Review Letters*, **102**(011301).

[148] M. Barnabé-Heider, *et al.* (2005). *Physics Letters B*, **624**(186).

[149] F. Giuliani, *et al.* (2004). *Physics Letters B*, **588**(151).

[150] L. Roskowski, *et al.* (2007). *Journal of High Energy Physics*, **07**(075).

[151] S. J. Plank (2008). *The DRIFT Dark Matter Project: Directionality, Sensitivity, and Environmental Backgrounds*. Ph.D. thesis, The University of Edinburgh.

[152] K. Krane (1987). *Introductory Nuclear Physics*. John Wiley and Sons.

[153] E. A. Lorch (1973). *International Journal of Applied Radiation and Isotopes*, **24**:585–591.

[154] A. D. Vijaya, *et al.* (1973). *Nuclear Instruments and Methods in Physics Research*, **111**:435–440.

[155] S. Agostinelli (2003). *Nuclear Instruments and Methods in Physics Research A*, **506**(3):250–303.

[156] R. Brun, *et al.* (1998). *Linux Journal*, **51**.

[157] J. F. Briesmeister (1997). MCNP - A general Monte-Carlo N-particle code. Los Alamos National Laboratory.

[158] A. Fasso, *et al.* (2001). Proceedings of the Monte Carlo 2000 Conference.

[159] D. Y. Akimov, *et al.* (2010). *Physics Letters B*, **692**(180).

[160] W. B. Wilson, *et al.* SOURCES-4A Technical Report 13639-LA. Technical report, Los Alamos.

[161] T. Sumner (2005). *New Astronomy Reviews*, **49**(2-6):277–281.

[162] D. Y. Akimov, *et al.* (2010). *Astroparticle Physics*, **34**:151–163.

[163] E. J. Barnes (2010). *A veto for the ZEPLIN-III Dark Matter Detector*. Ph.D. thesis, The University of Edinburgh.

[164] A. Wright, *et al.* http://arxiv.org/pdf/1010.3609v1.

[165] Y. Sakurai, *et al.* (2002). *Journal of Nuclear Science and Technology*, **2**:1294–1297.

[166] A. Wright (1983). *Journal of Physics E*, **16**.

[167] V. N. Solovov, *et al.* (2004). *Nuclear Instruments and Methods in Physics Research A*, **516**:462.

[168] W. A. Rolke, *et al.* (2003). *Nuclear Instruments and Methods in Physics Research A*, **551**:493–503.

[169] R. Brun, *et al.* (1997). *Nuclear Instruments and Methods in Physics Research A*, (389):81–86.

[170] K. Ni, *et al.* (2007). *Nuclear Instruments and Methods in Physics Research A*, **582**(2):569–574.

[171] F. Neves, *et al.* (2009). *Astroparticle Physics*, **33**:13–18.