Direct searches for WIMP dark matter with ZEPLIN-III

Alastair Edward Currie
Department of Physics
Imperial College London

A thesis submitted for the degree of Doctor of Philosophy
July 2012

Abstract

Work contributing to experimental limits on WIMP dark matter scattering in a liquid xenon target is presented. The ZEPLIN-III detector at Boulby completed 83 days' continuous operation in 2008 and, following a hardware upgrade, a further 319 days in 2010–11. Inelastic dark matter hypotheses to explain the DAMA/LIBRA modulation as scattering from iodine are excluded with $\geq 87\%$ confidence using 2008 data. Upper limits on the elastic WIMP-nucleon cross section are set from the combined exposure, with a minimum of 3.9×10^{-8} pb for 50 GeVc^{-2} WIMP mass at 90% CL, in addition to competitive limits on the WIMP-neutron spin-dependent cross section.

The detector responses to scintillation and ionisation — throughout the fiducial volume and over the duration of the run — are characterised via calibration with ^{57}Co and ^{137}Cs sources, as well as sidebands in search data. Signal spectra are modeled using these responses and an *in situ* measurement of the energy-dependent light and charge yields of nuclear recoils. Analysis software, event selection and background estimation are all described. Confidence intervals based on sparse regions of parameter space (Yellin limits) and the likelihood ratios are implemented for ZEPLIN-III and discussed in the context of rare-event searches with significant background uncertainty.

Declaration

I certify that all material herein which is not my own work has been properly attributed.

—Alastair Currie

Acknowledgements

The work in this thesis was made possible by the efforts of many people. I thank Henrique Araújo for his judgement and dedication as a supervisor. The camaraderie of the ZEPLIN-III team will always stay with me and I feel lucky to have worked among this talented group of people. I am grateful to the other Imperial group members for being such insightful and helpful colleagues: Blair Edwards, Markus Horn, Karl Lyons, John Quenby, Tim Sumner, Claire Thorne and Richard Walker.

Boulby days (and occasional nights) were the defining experience of my time on the project. It was a privilege to work side-by-side with many fine physicists, and to play a part in realising the vision of one in particular, Vadim Lebedenko, who sadly was not there to see it happen. Thanks go to CPL for hosting a unique low-background facility and to the Palmer Lab team for their vital help in keeping the show on the road. Glen Cowan, Roberto Trotta and Neal Weiner were generous with their time and expertise, which were very much appreciated.

Finally, I would like to thank my family (current and soon-to-be) and my fiancée for their invaluable encouragement and support through the last four years and before.

Contents

List of Figures

List of Tables

Chapter 1

Dark matter

> *These considerations show that the great dispersion of velocities in the Coma system harbours a problem that is not yet understood.*
>
> —Fritz Zwicky, 1933 [1]

The idea that familiar luminous matter is outweighed in the universe by matter which neither emits nor absorbs electromagnetic radiation has a long history, but the nature of the dark matter remains mysterious. In the 1930s Zwicky [1] inferred a gravitational mass for the Coma galaxy cluster, by application of the virial theorem, which was many times the mass estimated from surface brightness. Subsequent observations of velocity dispersions within galaxies and clusters, the rotational speeds of spiral galaxies, and the motion of our local group repeated that qualitative conclusion ([2] and references therein): luminous matter could not account for the apparent gravitational potentials, at least without modifying the theories of Newton and Einstein. The most compelling recent pieces of evidence for dark matter — which also favour it over a revised gravitational theory as an explanation of the dynamical results — come from cosmology and gravitational lensing.

This chapter begins by discussing some of these very large scale observations, but returns to galactic dynamics when considering the distribution of dark matter in the vicinity of Earth. Candidate particles are predicted by many physics models, often independently motivated, and a selection of these are outlined, as are some experimental strategies to detect dark matter via non-gravitational effects: production at colliders, searches for annihilation products (indirect detection), and scattering in terrestrial targets (direct detection). The principles and main results of direct detectors in particular are described, setting the scene for the ZEPLIN-III experiment.

1.1 Evidence for dark matter in the universe

The pre-eminent theory of cosmology at present is known as the Λ-CDM model for its two dominant components: cold dark matter (CDM, with cold implying non-relativistic particles even in the early universe) is electromagnetically inert but clumps gravitationally and

is mostly responsible for the distribution of matter within the expanding universe; dark energy (which may be described by a universal 'cosmological constant' Λ in general relativity) is a negative-pressure vacuum energy which determines the rate of expansion. We now describe, in broad terms, the development of the universe and some of the cosmological observations which point to the existence of cold dark matter.

1.1.1 The cosmic microwave background

The history of the universe after the Big Bang (see, e.g. Ref. [3]) is defined by phase transitions as the initial hot, dense state expands and cools. A presumed quark-gluon plasma gave way to protons and neutrons at $T \sim 10^{12}$ K; the light nuclei were made from $\sim 10^{10}$ K and existed, tightly coupled to electrons and to electromagnetic radiation, in a photon-baryon fluid; at 3000 K, corresponding to a 380 000-year-old universe, electrons and nuclei combine into atoms and radiation effectively decouples from matter — the universe becomes transparent when the mean free path of a photon exceeds the size of the observable universe. The spherical shell from which photons produced at the time of decoupling are just now arriving at Earth is known as the surface of last scattering; these red-shifted photons follow a black-body spectrum with temperature of 2.725 K and are known as the cosmic microwave background (CMB).

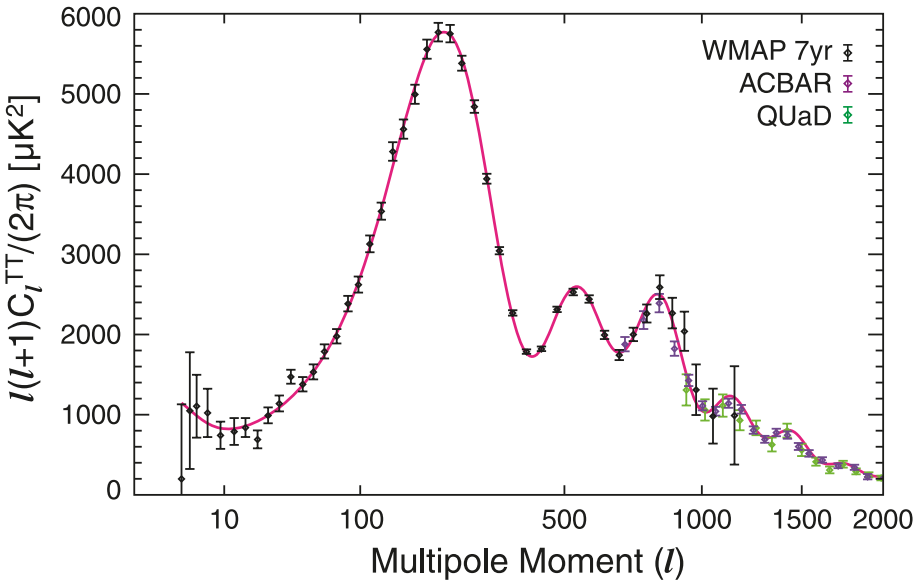

Figure 1.1: The angular power spectrum of the temperature of the CMB from 7-year WMAP data, in addition to data from the ground-based ACBAR and QUaD experiments where they are more precise than WMAP. C_l is the average, over the orders m, of the squared amplitudes of the spherical harmonic expansion coefficients with degree l. The red line is a fit to the 6 parameters of Λ-CDM in a geometrically flat universe (From Ref. [4]).

The main point of interest, in relation to dark matter, is that temperature fluctuations in the CMB encode density fluctuations in the photon-baryon fluid; the statistical properties of these fluctuations in turn depend on the composition of the universe at the time of decoupling. The temperature fluctuations of the surface of last scattering, $\delta T(\theta, \phi)$,

are quite small ($\delta T/T \sim 10^{-5}$), and may be analysed by decomposition into spherical harmonics, Y_l^m. The angular power spectrum of the temperature fluctuations is shown in Fig. 1.1. The spectrum is well described by the Λ-CDM model, which has 6 free parameters including the baryon density, Ω_b, and the cold dark matter density, Ω_c, both measured in units of the critical density for a geometrically flat universe. The 7-year dataset of the Wilkinson Microwave Anisotropy Probe (WMAP) [4] implies

$$\Omega_b h^2 = 0.0226 \pm 0.0005$$

and

$$\Omega_c h^2 = 0.113 \pm 0.004\,,$$

where

$$h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}) = 0.702 \pm 0.014$$

is the dimensionless Hubble parameter.

Radiation exerts pressure against the accumulation of baryonic matter in primordial gravitational potential wells. Qualitatively, the peaks in the angular power spectrum correspond to different acoustic modes of the photon-baryon fluid, excited by density fluctuations after the end of inflation. The fundamental mode (corresponding to the first peak, around $l = 220$) has a wavelength of twice the sound horizon in the fluid at the time of decoupling, and points to the realised critical density. The ratio of peak heights in the even and odd modes is sensitive to the baryon density: inertia in the fluid due to baryons enhances the odd (compression) peaks relative to the even (rarefaction) ones. The amplitude of all the peaks is sensitive to the ratio of total matter to radiation and hence, by subtraction, the non-baryonic matter density. The relative amounts of radiation and matter determine the amount of 'radiation driving', whereby the time-dependence of a radiation-dominated potential increases the amplitude of the acoustic oscillations [5]. In practice, many other physical effects must be accounted for to compare acoustic oscillations with the observed CMB today. These include the gravitational potential that the photons originate in (Sachs-Wolfe effect) and time-varying potentials they pass through before reaching us (integrated Sachs-Wolfe effect). However, the conclusion is unambiguous: dark matter dominates over baryonic matter in the universe as a whole, by a factor of about 5.

1.1.2 Large-scale structure

Although the universe appears homogeneous over very large distances, ~100 Mpc, significant structure is observed in the distribution of galaxies below this scale. Galaxies are gathered into groups and clusters which themselves form superclusters and connecting filaments, separated by voids which can be tens of megaparsecs across. The resulting 'cosmic web' is shown in Fig. 1.2 which represents data from the Sloan Digital Sky Survey (SDSS), a galactic red-shift survey.

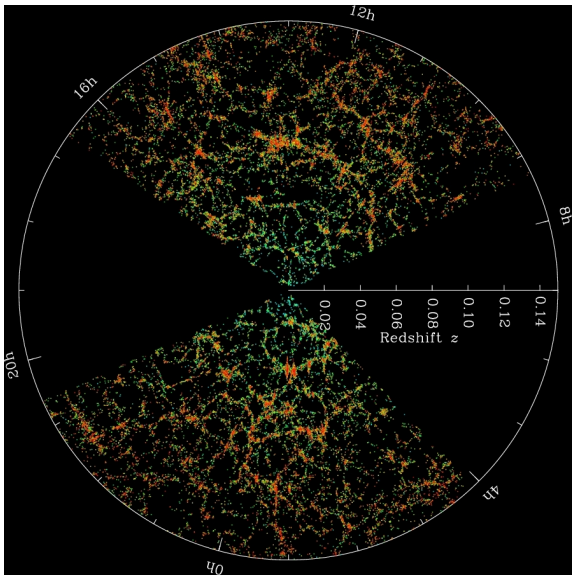

Figure 1.2: The distribution of galaxies in a 2.5-degree slice of declination, out to red-shift 0.15 (600 Mpc) as recorded by the Sloan Digital Sky Survey. Red (green) colour corresponds to the presence of old (young) stars. The large-scale structure of filaments and voids is clearly visible. (From Ref. [6]).

Large-scale structure is principally driven by gravitational instability: initial overdensities in the dark matter distribution establish potential wells and accumulate matter to seed the successive levels of hierarchical clustering. The evolution of early density fluctuations into the large scale structure of the universe today is an important constraint on dark matter properties. A simple consequence is that any particle candidate must be long-lived relative to the 13.7 billion-year age of the universe. The dark matter must also be predominantly cold (i.e. non-relativistic at the time of decoupling) in order to form galaxies and reproduce the cosmic web in numerical simulations of structure formation (reviewed in Ref. [7]) .

Acoustic oscillations of the photon-baryon fluid (as discussed in the previous section in relation to the CMB) are also imprinted in the the large-scale structure of the universe. Initial (post-inflation) overdense regions of dark matter coincide spatially with overdense regions of baryons. The dark matter starts the process of structure formation, by attracting surrounding matter, but feels no radiation pressure; the overdense baryons, however, do and this results in an acoustic wave which propagates until decoupling. That distance of one sound horizon at decoupling shows up as a peak in the 2-point correlation function of galactic matter (after accounting for the intervening expansion). This signature of baryon acoustic oscillation (BAO) in the distribution of galaxies is shown in Fig.1.3, and the overall matter power spectrum is well fit by the predictions of cold dark matter and baryons with the same densities obtained from the CMB [8].

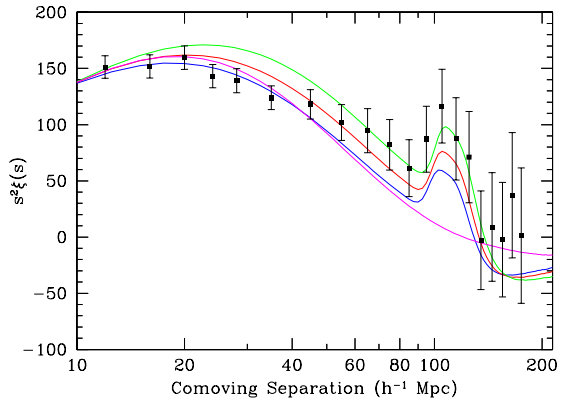

Figure 1.3: The two-point matter correlation function of large-scale structure measured by SDSS, showing the baryon acoustic peak at separation $s = 100\,h^{-1}$Mpc. The green, red and blue lines model $\Omega_m h^2$ equal to 0.12, 0.13 and 0.14 respectively (where $\Omega_m = \Omega_c + \Omega_b$, and with $\Omega_b h^2 = 0.024$ in all cases), showing the sensitivity to total matter density. The magenta line models pure CDM without baryons, hence lacking the BAO peak. (From Ref. [8].)

1.1.3 Gravitational lensing

Further evidence that dark matter is the dominant component of clusters and superclusters comes from recent studies of gravitational lensing (reviewed in Ref. [9]). As predicted by general relativity, the path of light from source to observer can be bent by an intervening gravitational field. There are two broad classes of lensed images, described as strong and weak. In strong lensing, the deflection results in multiple images or distorted arcs from single sources: complete 'Einstein rings' may be observed if a distant source galaxy, a strong gravitational lens and an observer are well aligned. Relatively few lines of sight pass near enough to massive enough objects to cause strong lensing, and weak lensing theory describes the linear transformation of images by small amounts of gravitational magnification, shear and rotation. Statistical averaging of the observed shapes of source galaxies (whose intrinsic physical shapes are uncorrelated) can then reveal the distribution of gravitational matter in a cluster.

Combining such gravitational maps with X-ray observations of clusters allows observers to distinguish the dominant baryonic component (X-ray-emitting plasma) from the dark matter. Clusters which are in the process of merging together can have a centres of baryonic mass displaced from the centre of gravitational mass, because the hot gas clouds of the two colliding clusters experience ram forces; the dark matter, in contrast, may well be collisionless, and is at most very weakly self-interacting. The prime example of such a cluster merger is the Bullet Cluster (1E 0657–558, shown in Fig. 1.4) which was presented as 'A direct empirical proof of the existence of dark matter' [11]. A similar effect is seen in MACSJ0025.4-1222 [12]. Such detailed observations of the origin of baryon-dark matter separation in galaxy clusters — with collisional motion roughly in the observable plane and yet to dynamically relax — are understandably rare, but a recent study of 38 clusters

Figure 1.4: The Bullet Cluster (1E 0657–558) merger shown in visible light. An X-ray image of the hot intergalactic plasma (which out-weighs the visible galaxies very substantially) is shown in pink; a map of total gravitational matter from lensing data is superimposed in blue. Two clusters have passed through each other, approximately in the image plane, at 4 700 km/s. The height of the image corresponds to ~1 Mpc at the distance of the cluster. (NASA image: X-ray analysis from [10], gravitational analysis from [11]).

having both X-ray and strong-lensing data found that around half display a statistically significant offset [13].

1.2 Dark matter candidates

The foregoing paradigm of non-baryonic, cold, gravitationally-interacting matter in a universe with dark energy gives a very successful explanation of the decoupling era, subsequent structure formation and the dynamics of galaxies, clusters and super-clusters as well as numerous other observations. It is then natural to speculate about the fundamental properties of the dark matter, and how it might relate to another conspicuously successful physical theory, the Standard Model of particle physics. The phenomonological constraints on possible candidates, including effects on stellar evolution and big bang nucleosynthesis are collated in Ref. [14]. Here we review two especially promising ideas which have survived all observational tests so far.

1.2.1 Axions

The Standard Model itself contains no satisfactory CDM candidate: neutrinos are non-baryonic, dark and are not massless, but nor are they massive or slow enough to be bound by galaxy-sized objects and thus explain large scale structure. However, the hypothetical axion is a cold dark matter candidate which arises from an independently-motivated extension of the standard model. The Lagrangian of quantum chromodynamics (QCD) includes a non-perturbative term which results in an amount of CP violation governed by a free parameter in the theory, $\bar{\theta}$. This parameter could, *a priori*, take any value from 0 to 2π. However, no CP violation has ever been observed in strong interactions and the most stringent experimental limits, from the electric dipole of the neutron [15], indicate

that $\bar{\theta} < 10^{-9}$. One solution to this fine-tuning problem, due to Peccei and Quinn [16], is to invoke a global U(1) symmetry which is spontaneously broken, at some temperature T_{PQ}, such that the Lagrangian includes an effective potential $V\left(\bar{\theta}\right)$ with a minimum at $\bar{\theta} = 0$. The particle associated with the symmetry breaking (a pseudo-Goldstone boson) is the axion, a pseudo-scalar particle with mass denoted by m_a.

CP conservation in strong interactions would be explained by any value of the axion mass over many orders of magnitude, but the present experimentally-allowed region extends only up to 10^{-2} eV [17]. Masses above around 50 keV are excluded by accelerator searches. Lighter axions would be an efficient means of energy loss in stars and the observed longevity of red giants excludes masses above 0.5 eV. The mechanism of supernovae cooling would also alter, and the duration of the neutrino burst from Supernova 1987a extends the excluded mass range down to 10^{-2} eV.

Although of low mass, axions can fill the role of cold dark matter because they would be produced by processes other than 'freezing out' from chemical equilibrium. The theoretical production mechanisms for non-relativistic axions ('vacuum realignment', whereby a Bose condensate of axions is formed, and the decay of cosmic strings) are outlined in e.g. Ref. [18]. The cosmological density of cold axions from these processes is proportional to $m_a^{-7/6}$ and WMAP sets a lower limit, with significant uncertainty from the details of the QCD phase transition, of about 10^{-5} eV.

Laboratory searches exploit the predicted axion-to-photon coupling and look for the resonant conversion of dark matter axions to photons in a magnetic field [19] or laser light being converted into axions and then reconverted on the other side of an opaque wall [20]. Limited mass ranges have been excluded, but axions remain a feasible dark matter candidate.

1.2.2 WIMPs

A fairly broad class of cold dark matter candidate is the weakly interacting massive particle, or WIMP: a particle which participates in the weak interaction as well as gravity, and is massive enough to freeze out as cold dark matter. WIMPs can arise in Kaluza-Klein extra-dimensional models or supersymmetric extensions to the standard model, among other well-motivated theories [21] — in general, some kind of new physics in the weak sector is suggested by the extreme fine-tuning required of loop corrections to the Higgs mass (the hierarchy problem) [22]. WIMPs are especially interesting as dark matter candidates because annihilation cross sections typical of the weak force would give a thermal relic density of the same order of magnitude as the observed Ω_m.

This extraordinary coincidence of densities, the calculation of which includes such disparate energies as the Planck mass and the CMB temperature, is known as the 'WIMP miracle' and can be understood as follows. In the very early Universe — before even Big Bang nucleosynthesis — there is chemical equilibrium between the WIMPs and their annihilation products. Once the temperature (in natural units) is comparable to the

WIMP mass, the relative density of WIMPs starts to drop exponentially — annihilation continues without enough thermal energy for production from lighter particles. However, at a density determined by the thermally-averaged annihilation cross section $\langle \sigma_A v \rangle$, the mean free path for WIMP-annihilating collisions is of the order of the Hubble radius and chemical equilibrium ends: in other words, the expansion of the universe beats the annihilation rate, and the number of WIMPs is fixed. A numerical evaluation of the co-moving density (i.e. number density with the expansion of the universe divided out) in these conditions is shown in Fig. 1.5, for example values of $\langle \sigma_A v \rangle$. An approximate expression for the density today of a WIMP with mass m_χ, in units of critical density, is [23]:

$$\Omega_\chi h^2 \simeq \left(\frac{3 \times 10^{-27}\,\mathrm{cm}^3\,\mathrm{s}^{-1}}{\langle \sigma_A v \rangle} \right). \tag{1.1}$$

The density for a given $\langle \sigma_A v \rangle$ is independent of the WIMP mass, up to logarithmic corrections, but the annihilation cross section is typically mass dependent. On dimensional grounds, we can write a weak-mediated annihilation cross section as [24]:

$$\sigma_A v = k \, \frac{g_{\mathrm{weak}}^4}{16\pi^2 m_\chi^2} \,, \tag{1.2}$$

where k is a constant of order 1 and $g_{\mathrm{weak}} \simeq 0.65$ is the weak gauge coupling. There is an additional overall v^2 factor for P- rather than S-wave annihilation. Accounting for this, and accepting k=0.5–2, it is found that a single WIMP with mass in the range 10^2–10^3 GeVc^{-2} would reproduce the correct relic abundance by itself. A spectrum of lower-mass WIMPs could also fill the role of CDM.

If dark matter particles interact via the weak nuclear force, this opens the possibility to detect them directly by searching for scattering interactions. The sensitivity of this method will be determined by the density of dark matter in a laboratory, rather than in the entire Universe, and so we now proceed to consider dark matter in the Milky Way.

1.3 Galactic dark matter in the solar neighbourhood

The key measurement which constrains the distribution of dark matter in spiral galaxies such as ours is the rotation curve $v\,(r)$: the orbital velocity of stars and gas as a function of galactocentric radius. Rotation curves are measured via the Doppler shifts of emission spectra: these include optical-wavelength atomic transitions such as the hydrogen lines, those from the hyperfine structure (e.g. the 21-cm line of atomic hydrogen) and rotational transitions in CO (in the millimetre-wave range).

The first rotation curve to suggest non-luminous matter was recorded for the Andromeda galaxy M31 in 1939 [26]. In the 1970s, precise determination of rotation curves for a large sample of spiral galaxies by Rubin et al. [27] was instrumental in bolstering the idea of galactic dark matter. Fig. 1.6 is a recent collation of data for nearby spiral

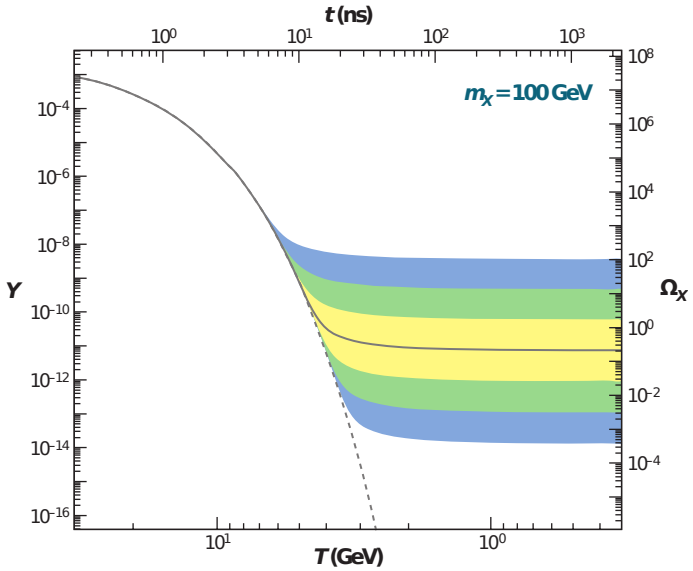

Figure 1.5: The evolution of co-moving WIMP density, Y, in the early Universe, showing WIMP freeze-out. The solid line corresponds to the $\langle \sigma_A v \rangle$ which gives the observed matter density. The shaded bands extend to annihilation cross sections 10, 10^1 and 10^2 times lower or higher. Higher $\langle \sigma_A v \rangle$ reduces the relic density by keeping the WIMP in chemical equilibrium to later times. The dashed line shows the density that would be reached if were equilibrium were maintained. (From Ref. [24].)

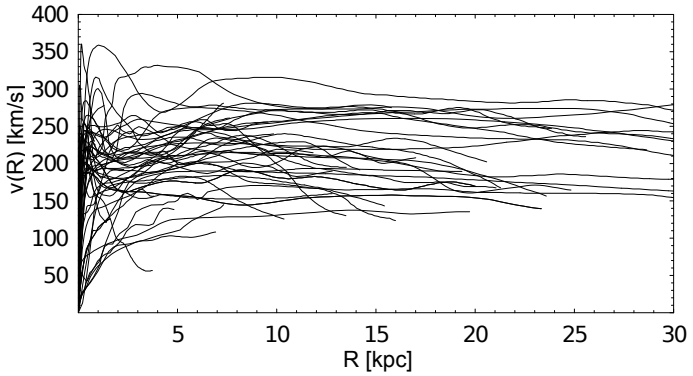

Figure 1.6: Measured rotation curves for nearby spiral galaxies, combining CO, optical and 21-cm data at progressively larger radii. (From Ref. [25].)

galaxies, showing the essentially universal flattening of rotation curves at large galactic radii.

The density of the visible disc of the Milky Way falls exponentially with radial distance, having a decay length of 3.5 kpc, and the central bulge extends only to about 1 kpc [28]; in the absence of dark matter the rotation curve would fall off in a roughly Keplerian fashion ($v \propto r^{-1/2}$) beyond a few kpc. Instead, the rotation curve of the Galaxy stays flat to at least 20 kpc and falls only gently by 60 kpc [29]. A flat rotation curve at large distance, where luminous matter is very sparse, is most simply explained by an isothermal spherically-symmetric halo of dark matter, which corresponds to a density profile $\rho(r) \propto r^{-2}$. Solving the collisionless Boltzmann equation in this case gives a Maxwellian distribution for the velocity of the dark matter particles, with RMS speed equal to $\sqrt{3/2}$ times the circular velocity [30]. To calculate direct detection rates, this distribution must be boosted into the lab frame; the speed of the Earth in the galactic frame modulates annually in the range 215–240 km/s.

The spherical isothermal model (with the velocity distribution truncated to account for the escape velocity of the galaxy) is used as a standard reference halo for direct detection results. Values of 220 km/s for circular velocity, 544 km/s for galactic escape velocity and $0.3\,\mathrm{GeV\,c^{-2}\,cm^{-3}}$ for the local density of dark matter are usually adopted. Many possible halo profiles, either inspired by N-body simulations or the detail of rotation curve measurements, give acceptable fits to the available astronomical data, and they predict slightly different local densities of dark matter. The theoretical rotation curves from a selection of halo profiles are shown in Fig. 1.7. An uncertainty of at least 30% on the dark matter density in the lab, and hence on event rates, should be borne in mind when comparing direct search results to other constraints [31].

1.4 Direct detection

Scattering only via the weak force, WIMPs would cause recoils of atomic nuclei in a target medium, as opposed to the electron recoils produced by most backgrounds. The evidence from the local Galaxy suggests that dark matter particles are present in the lab with a density of order $0.3\,\mathrm{GeV\,c^{-2}\,cm^{-3}}$ and an RMS speed of order $0.001c$. The mean kinetic energy of, e.g., a WIMP of mass $100\,\mathrm{GeV\,c^{-2}}$ (consistent with the thermal relic density) would then be 50 keV. Measuring recoil energies of this order of magnitude within large detector targets, typically in conjunction with extremely powerful background rejection, is the goal of WIMP direct detection.

1.4.1 WIMP-nucleus scattering rates

The signal model in the ZEPLIN-III experiment will be developed in detail in Chapter 4. Here, in order to introduce direct detection experiments in general, we pass over some details and present approximate rates as set out in Ref. [32]. The rate of WIMP-nucleus

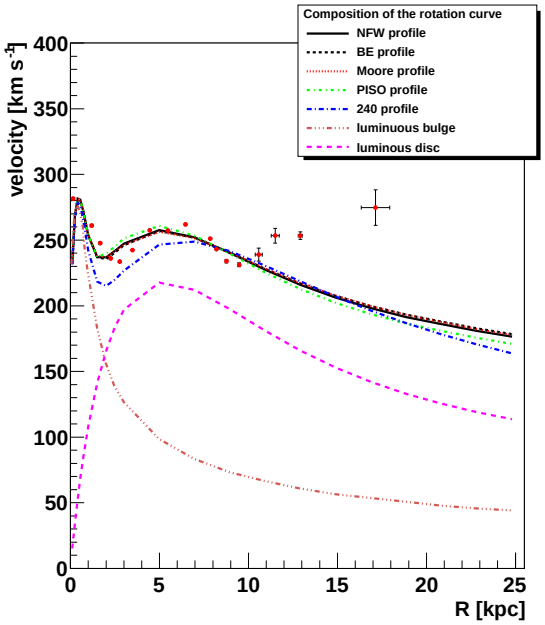

Figure 1.7: The rotation curve of the Milky Way (averaged from many observations) with fits to dark matter profiles. The NFW, BE and Moore profiles are, to various degrees, 'cuspy' in having a density profile which are divergent in the centre. Cusps are found in many N-body simulations of dark-matter structure formation, but may in part be an artefact due to the omission of baryons. PISO and 240 are 'cored' distributions, tending to constant ρ at small radius, in agreement with observations of dark-matter dominated dwarf spiral galaxies. (From Ref. [31], which contains the density profile definitions.)

scattering is determined by the WIMP-nucleus cross section σ_0, the local WIMP density ρ_0 and velocity distribution, and the properties of the target nucleus including the nuclear form-factor. The coherent interaction with the whole nucleus breaks down for sufficiently large momentum transfer q; the form factor $F^2(q)$ accounts for this.

The idealised total scattering rate R_0 is defined by considering a detector at rest in a dark matter halo with an escape velocity $v_{\rm esc} = \infty$. In a spherical isothermal halo with velocity dispersion $\sqrt{(3/2)} \cdot v_0$, the rate per unit detector mass is:

$$R_0 = \frac{2}{\sqrt{\pi}} \frac{\rho_0}{m_{\rm N} m_{\rm W}} \sigma_0 v_0, \tag{1.3}$$

where $m_{\rm N}$ is the mass of the target nucleus and $m_{\rm W}$ is the WIMP mass. Putting in units, and for a target with atomic mass A:

$$\frac{R_0}{1\,{\rm kg}^{-1}\cdot{\rm day}^{-1}} = \left(\frac{390}{A}\right)\left(\frac{1\,{\rm GeV\,c}^{-2}}{m_{\rm W}}\right)\left(\frac{\sigma_0}{1\,{\rm pb}}\right)\left(\frac{\rho_0}{0.3\,{\rm GeV\,c}^{-2}\,{\rm cm}^{-3}}\right)\left(\frac{v_0}{220\,{\rm km\,s}^{-1}}\right). \tag{1.4}$$

Although heavy elements give fewer target nuclei per unit mass, this is more than compensated for by the A^2 dependence of σ_0 in terms of the underlying spin-independent WIMP-nulceon cross section $\sigma_{\rm n}$:

$$\sigma_0 = A^2 \left(\frac{\mu_{\rm N}}{\mu_{\rm n}}\right)^2 \sigma_{\rm n}\,, \tag{1.5}$$

where $\mu_{\rm n}$ and $\mu_{\rm N}$ are the WIMP-nucleon and WIMP-nucleus reduced masses. For WIMPs significantly heavier than the target nucleus, $\mu_{\rm N}^2 \propto A^2$ and hence $\sigma_0 \propto A^4$; heavier elements result in higher scattering rates, even after the inclusion of the form factor (which drops more quickly with $E = q^2/2m_{\rm N}$ for large $m_{\rm N}$).

Accounting for the Earth's velocity, but still neglecting the effect of $v_{\rm esc}$, the differential spectrum in recoil energy E is well approximated by an exponential:

$$\frac{dR}{dE} \simeq c_1 \frac{R_0}{E_0 r} \exp\left(\frac{-c_2 E}{E_0 r}\right); \tag{1.6}$$

here $r \equiv 4 \cdot m_{\rm W} m_{\rm N}/(m_{\rm W}+m_{\rm N})^2$ is a kinematic factor and $E_0 \equiv m_{\rm W} v_0^2/2$ is the characteristic kinetic energy of the halo WIMPs. The numerical factors c_1 and c_2, tabulated in Ref. [32], depend on the Earth's velocity; they are on average 0.75 and 0.56, respectively, and vary through the year within 5% of the mean. The total scattering rate is the integral of this featureless spectrum above some nuclear recoil detection threshold; sensitivity to a light ($\lesssim$10 GeV c^{-2}) WIMP is achieved by only a few technologies — including ionisation-only searches with xenon [33] — due to the steeply falling spectra in this scenario.

1.4.2 Experimental techniques for direct detection

The present global status of limits on the spin-independent WIMP-nucleon scattering cross section — including the results of the ZEPLIN-III project which this thesis describes —

is shown in Fig. 5.7 on page 121. A number of common technical challenges are faced by the leading experiments as they probe cross sections below 10^{-7} pb. Event rates due to α, β, γ and neutron radiation must be controlled by deep underground operation, materials selection, shielding and possibly veto instruments. Energy thresholds of 10 keV and below for nuclear recoils are desirable in general, and essential for sensitivity to light WIMPs. Finally, electron-recoil background events typically need to be distinguished from signal-like nuclear recoils with high discrimination power.

Scattering events in many modern direct detectors are measured using two energy channels out of heat (phonons), scintillation light and ionised charge. For each technology, the relative yields in the available channels depend on the recoiling species, and the most successful background discrimination strategies exploit this. In general the scintillation and ionisation yields for nuclear recoils are lower than for electron recoils, and are thus described as quenched. The phonon signal is found not to be significantly quenched [34]. An event's 'visible' energy in the ionisation or scintillation channel is reported as an electron-equivalent energy in units of keVee; the corresponding nuclear recoil energy, related to electron-equivalent energy by the quenching factor, is sometimes reported in units of keVr to avoid ambiguity. We now survey the leading experimental searches for the spin-independent interaction, which fall into two categories based on the target medium: cryogenic crystal detectors (especially germanium) and liquid noble element detectors (notably xenon and argon). A more complete review of the dozens of current and recent direct detection experiments may be found in Ref. [35]. It should be stressed that, in the event of discovery, data from multiple target materials would be extremely complementary: degeneracies in particle and halo models may be broken by comparing several different elements (see, e.g. Refs [36, 37, 38]).

Cryogenic crystals

Energies on the scale predicted by WIMP-nuclear scattering may be measured as a temperature increase, of order μK, in crystals cooled to tens of mK. Cryogenic bolometers measure either thermal or athermal phonons as a primary energy scale; electron-recoil discrimination is achieved via a second channel. The phonon signal, made up of ~10,000 quanta per keV, can achieve excellent energy resolution, e.g. σ =0.6 keVee at 122 keVee in CRESST-II [39].

The CDMS-II [40] experiment at the Soudan Mine, USA, employed a combination of silicon and germanium crystals as ionisation bolometers. A weak electric field drifts ionisation charge from the interaction vertex to electrodes on the top and bottom of the ~10-mm thick wafers. The main WIMP-nucleon limits are derived from the nineteen 230-g germanium modules; comparison of rates in the two target elements helps to guard against mis-identifying neutron background as a spin-independent WIMP-nucleus signal, since the A-dependence of the cross section would be much stronger for the signal. Silicon modules have also been analysed for the purposes of dedicated low-WIMP-mass limits [41].

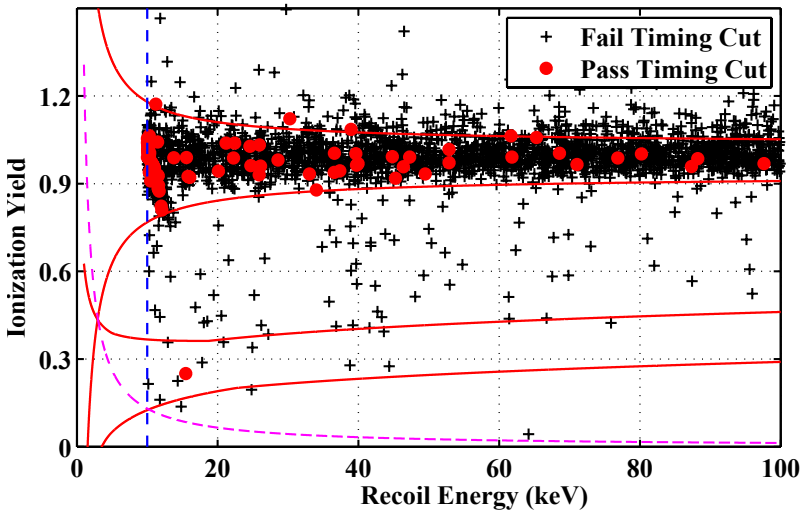

Figure 1.8: Ionisation yield versus recoil energy for search events in one module of the CDMS-II experiment. The solid red lines show ionisation yields of $\pm 2\sigma$ about the mean for electron calibration (centred on 1) and neutron calibration (below). The latter defines the signal acceptance region. The vertical dashed line is the recoil energy threshold, and the sloping dashed line is the ionisation threshold. (From Ref. [40].)

A potentially damaging background in ionisation bolometers is surface events: electron recoils close to the electrodes. Charge collection can be suppressed in this case by trapping in crystal defects and diffusion of electrons or holes to the wrong electrode. A very effective cut on phonon timing is employed: the response of the superconducting transition edge sensors is sufficiently fast to distinguish bulk and surface events. Figure 1.8 shows an example of the discrimination scatter plot for one module. In 2010, CDMS-II published a 90% upper limit on σ_{n} with a minimum of 3.8×10^{-8} pb at 70 GeVc^{-2} WIMP mass, which is still the leading cryogenic result.

EDELWEISS-II , which ran at the Laboratoire Souterrain de Modane, also employs ionisation-sensitive bolometers. It consists of ten 400-g germanium detector modules. An arrangement of interdigitated electrodes was developed to mitigate surface events: the field in the crystal bulk drifts all charge to a fiducial subset of electrodes, whereas events at the edge will also produce signal in other, veto electrodes which alternate with the signal electrodes. The 2011 result [42] achieved electron-recoil discrimination of 3×10^{-5} in the energy range 20–200 keV. On the basis of 5 candidate events and a background estimate of 3.0 in 384 kg·day net exposure, a 90% spin-independent limit of 4.4×10^{-8} pb was set at 85 GeVc^{-2} WIMP mass.

A different approach to cryogenic direct detection is exemplified by the scintillating bolometer experiment CRESST-II. Each detector module consists of two crystals in a single highly reflective housing: a 300-g $CaWO_4$ target and a separate, much smaller 2-g silicon or sapphire-on-silicon crystal to record scintillation light emitted by the target. Signals from both crystals are detected using transition edge sensors to give precise event energies and a light-yield discrimination parameter. The 2011 CRESST-II results [43] report an excess of low-energy events which is claimed to constitute 4.7-σ evidence against the background-

only hypothesis. 67 candidate events were observed and the best fit model predicted 42 background events from a combination of electron recoil, neutrons, α-particles and recoiling Pb nuclei. Two allowed regions in the standard spin-independent elastic WIMP scenario are presented therein, centred on (25.3 GeVc^{-2}, 1.6×10^{-6} pb) and (11.6 GeVc^{-2}, 3.7×10^{-5} pb), but these are very surely excluded by other experiments. It remains to be determined whether deficiencies in the background model or a non-minimal dark-matter signal (e.g. from isospin-violating dark matter [44]) are responsible for the apparent excess.

Cryogenic detectors have distinct advantages in their energy resolution, threshold and discrimination but do not scale up as naturally as some competing technologies. The method of choice to build tonne-scale cryogenic targets is an array of many small modules: there are diminishing returns from crystals above ~1 kg mass, due to the increased heat capacity and cost [45]; also, each additional crystal adds background to the experiment through support structures and readout interfaces. Next-generation cryogenic experiments, with target masses increasing slowly up to 1 tonne, are currently being prepared by the SuperCDMS [46] and EURECA [47] collaborations.

Liquid noble elements

The world-leading class of target at present is the liquid noble elements (hereafter, 'noble liquids') and in particular liquid xenon, as used in the ZEPLIN-III and XENON100 experiments. Single volumes of liquid target scale extremely well, due to self-shielding and the falling ratio of instrumentation to target mass. Noble liquid experiments have produced results using either or both of scintillation and ionisation as energy estimators; energy resolution is not as fine as for the heat signal in cryogenic detectors but, given the featureless signal spectra, little information is sacrificed for WIMP masses above ~10 GeVc^{-2}.

Aside from their common advantages, liquid argon is distinguished by more effective pulse-shape discrimination of the scintillation signal and its relative abundance in the atmosphere (although underground sources must be used to achieve low inherent background by depleting ^{39}Ar, which raises the cost significantly); liquid xenon has better self-shielding properties, higher yields of light and charge and a scintillation spectrum that, while in the VUV range, can be detected without wavelength shifting through the use of quartz-windowed photodetectors.

Contemporary noble liquid dark matter experiments which use both energy channels are of the two-phase emission type, operating as time projection chambers (TPCs). They measure ionisation by applying a strong ($\gtrsim$3kV/cm) electric field to emit electrons from the liquid into a thin layer of vapour above; electrons are detected via electroluminescence. The liquid is the WIMP target and scintillator medium; the gas is a means of amplifying the ionisation signal and also allows 3-D position reconstruction from the drift time of electrons to reach the surface. Both signals are read out by an array of PMTs in the liquid and, optionally, another in the gas. The main discrimination parameter to reject electron-recoil events is the relative yield of ionisation to scintillation. A calibration scatter plot

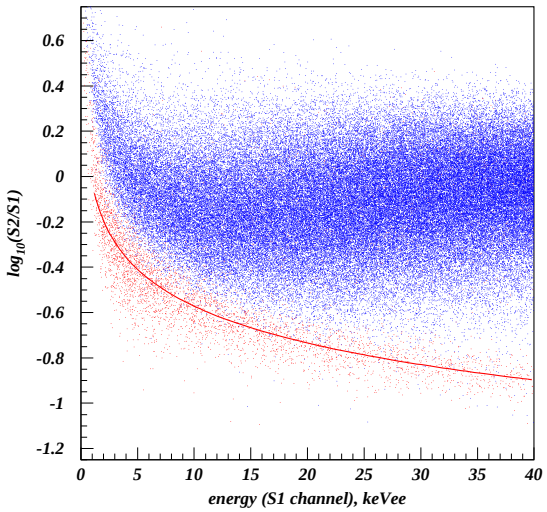

Figure 1.9: The discrimination parameter ($\log_{10}(^{S2}/_{S1})$) versus scintillation electron-equivalent energy (S1) for neutron (red points) and γ-ray (blue points) calibration data in the ZEPLIN-III two-phase xenon experiment. A region of low $\log_{10}(^{S2}/_{S1})$ is used to define signal acceptance for WIMP-induced nuclear recoils. (From Ref. [48].)

from the first run of ZEPLIN-III in 2008 is shown in Fig. 1.9 to illustrate the separation of electron and nuclear recoil populations.

ZEPLIN-II published the first two-phase noble liquid dark matter result in 2007 [49] and has been followed by the XENON10 [50], ZEPLIN-III and XENON100 [51] experiments. The last has a fiducial mass of 48 kg and a background count rate below 5×10^{-3} events per keVee·kg·day, driven down by self-shielding. As of March 2012, the leading 90% confidence limits on σ_{n} from liquid noble elements are the 100-live-day result of XENON100 at 7.0×10^{-9} pb and the ZEPLIN-III second science run at 4.8×10^{-8} pb, in both cases for minima around 50 GeVc^{-2} WIMP mass.

Rapid progress to 100-kg fiducial masses of xenon is demonstrated by the scintillation-only XMASS and the two-phase LUX experiments, both already built. The former relies on the self-shielding power of ~700 kg of xenon surrounding the fiducial volume, while the latter builds on the design of previous XENON and ZEPLIN detectors. The XENON-1T and LUX-ZEPLIN collaborations plan two-phase liquid xenon detectors at the tonne scale and above in the next 5 years, to reach sensitivities of at least 10^{-11} pb and probe much of the favoured parameter space in supersymmetric extensions of the standard model [52]; if a higher cross section is realised in nature, these large detectors will acquire sufficient counts to make detailed studies of the signal properties.

1.4.3 Directionality and annual modulation

The relative motion of the Earth and dark matter halo provides two particular characteristics to distinguish WIMP scattering from many sources of background in a model-independent way: directionality and annual modulation. The boosted WIMP velocity

distribution causes strong asymmetry in the angular distribution of the recoils. If ψ is the angle between the recoiling nucleus and the velocity of the halo in the lab frame, the differential rate is approximately [53]:

$$\frac{\mathrm{d}^2 R}{\mathrm{d}E \mathrm{d}\cos\psi} = \frac{R_0}{2E_0 r} \exp\left(-\frac{(v_\mathrm{E}\cos\psi - v_\mathrm{min})^2}{v_0^2}\right) \tag{1.7}$$

for Earth velocity v_E, where $v_\mathrm{min} = \sqrt{2Em_\mathrm{W}/r}$ is the minimum WIMP velocity for nuclear recoil with energy E. Such an angular dependence would be conclusive evidence of the dark-matter origin of a signal, but reconstruction of the direction and sense of tracks has only been achieved for relatively large recoil energies ($\gtrsim$ 50 keV) and in gaseous targets, [54, 55], severely limiting sensitivity for reasonably sized detectors. The UK's DRIFT Collaboration developed the first gaseous TPCs for directional detection, which are still operating at the Boulby Mine [56].

There is an annual modulation in the component of the Earth's orbital velocity in the direction of the (significantly larger) velocity of the solar system in the galactic halo. The relative speed of the Earth and halo peaks around June 2nd, with the minimum six months later, and the dark matter velocity distribution is boosted accordingly by between 215 and 240 km/s. The amplitude of modulation in the scattering rate varies with v_min, and is of the order of 1–10% [57].

The DAMA/LIBRA experiment, which measures scintillation events in 250 kg of high-purity sodium-iodide crystals, sees annual modulation with high (>8σ) statistical significance over many cycles [58]. Explanation in terms of a single elastic WIMP-nucleon cross section is ruled out by the null results of other direct detection experiments [59, 60]. Chapter 6 will present exclusion limits for an alternative WIMP model — 'inelastic dark matter' — which might have accounted for the inconsistency; almost all of its parameter space was disfavoured by the ZEPLIN-III result and it was later conclusively ruled out by XENON100.

An annual modulation in the rate of low-energy events is also reported in the CoGeNT direct detection experiment [61]. CoGeNT is located in the Soudan Underground Laboratory and measures ionisation signals in a germanium target maintained at 77 K. The modulation spectrum is consistent with a WIMP mass of 7–9 GeVc^{-2} and $\sigma_\mathrm{n} \sim 10^{-4}$ pb; however, its statistical significance, at 2.8σ after 442 live days, is much lower than that of the DAMA result. A recent temporal analysis of events in CDMS cryogenic germanium crystals found no evidence for modulation and disfavoured a WIMP-nucleus scattering explanation of the CoGeNT result at >98% confidence [62].

1.5 Other search strategies

Direct searches are by no means the only route to determine the particle properties of dark matter: candidate particles beyond the Standard Model are a major goal at colliders

and extensive astro-particle searches are made for dark matter annihilation products.

1.5.1 Collider experiments

Many theories beyond the Standard Model predict dark matter (DM) candidates that would be produced at experiments such as the LHC proton-proton collider at CERN. Cross sections for the production of DM candidates from primary interactions at a hadron collider are extremely small — whatever dark matter is, it does not couple strongly to quarks and gluons — but independently-motivated theories typically also predict new particles which do interact via QCD. If sufficiently light, these coloured particles will be produced with reasonable cross section and are likely to decay into final states which include the DM candidate.

The best-studied, and arguably best-motivated, theory of physics beyond the Standard model is that of supersymmetry (hereafter SUSY, and treated pedagogically in e.g. Ref. [63]). A symmetry relating bosons and fermions is interesting on theoretical grounds related to the allowable symmetries in quantum field theory, the possibility of a quantum description of gravity and the unification of the three Standard Model gauge couplings at high energy scales (grand unification).

In supersymmetric extensions of the standard model, the known elementary particles have superpartners with the same couplings but spin different by one-half: squarks for quarks, gluinos for gluons, etc. The superpartners cancel the divergent loop corrections to the Higgs mass, and the hierarchy problem is solved; the amount of residual fine-tuning depends on how much more massive than their partners the SUSY particles are, and this is one of the reasons to expect that they might appear at the LHC [64].

The very simplest SUSY theories make falsified predictions about the lifetime of the proton and flavor-changing decays. These problems can be solved simply via a multiplicative quantum number, known as R-parity, which is equal to -1 for SUSY particles and $+1$ for their standard model partners. R-parity conservation implies that SUSY particles are produced only in pairs, and that the lightest supersymmetric particle (LSP) would be stable, as required of a DM candidate. A number of possible LSPs could in principle fill the role of cosmological dark matter: the sneutrino; the gravitino, partner of the hypothetical graviton; the axino, partner of the hypothetical axion; or the neutralino, an admixture of the neutral Higgsinos, bino and wino which partner the Standard Model Higgs and the B and W^0 gauge bosons. The neutralino is an example of a WIMP, and is accessible to direct detection. The correct neutralino relic density may be obtained by adjusting the contribution of the bino (with its relatively low annihilation cross section) versus the wino and higgsino [65].

Neutralinos will not interact in detectors at the LHC but would show up as an imbalance in the vector sum of detected particles' momenta perpendicular to the beam. The most powerful constraints on neutralino dark matter in the general purpose detectors (ATLAS and CMS) come from searches for missing transverse energy in conjunction with

hadronic jets. Whether neutralino production is observable at the LHC depends principally on the mass of the squarks and gluinos which would decay into it. In the absence of a significant departure from the standard model, a global fit [66] to a minimal version of supersymmetry (the constrained minimal supersymmetric standard model, cMSSM) in the direct detection parameter space is shown in Fig. 1.10. It has been noted that in the so-called 'nightmare scenario' of no physics beyond the Standard Model being discovered at the LHC, all of the remaining cMSSM parameter space would be accessible to tonne-scale direct detectors [67].

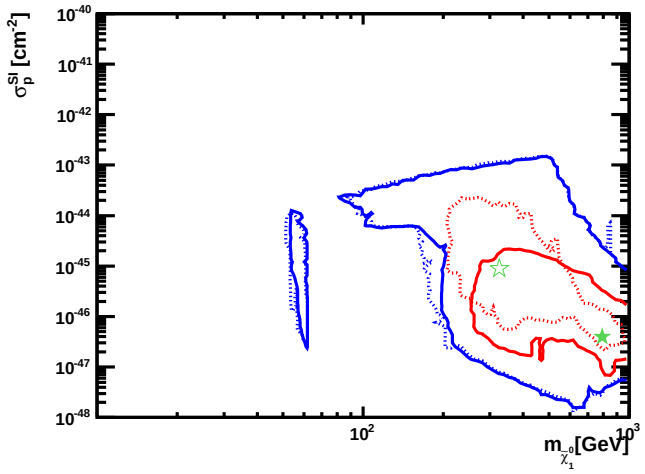

Figure 1.10: Allowed regions of neutralino mass and spin-independent neutralino-proton cross section in a frequentist analysis of the cMSSM. ATLAS and CMS SUSY results from 1/fb are used, together with cosmological, precision electro-weak and XENON100 [51] results. The solid (dotted) regions are calculated with (without) assuming a Higgs mass measurement of (125±1) GeVc^{-2}. Red (blue) correspond to 68% (95%) confidence levels. The flat shape at the top of the 95% regions is due to the XENON100 result. (From Ref. [66].)

1.5.2 Indirect detection

Dark matter which is its own antiparticle, for example a neutralino, will annihilate at a rate proportional to the square of its density in any given region. It may thus be detected by searching for annihilation products in ground-, balloon- or space-based experiments: the method of indirect detection (reviewed in Ref. [68]). Such results have been derived from neutrino, photon and anti-matter (positron, anti-proton and anti-deuteron) measurements. Annihilation neutrinos and photons will travel essentially unimpeded and should point back to areas of high DM density, within the Galaxy or beyond. Charged particles, by contrast, have their trajectories randomised by magnetic fields in the interstellar medium (ISM) and so only spectral information is available for these.

Neutralinos from the galactic halo are predicted to accumulate within the Sun by losing energy in elastic scatters until gravitationally bound. The capture and annihilation

rates balance at an equilibrium density determined by the DM halo parameters and the WIMP-nucleon scattering cross section. Of the (Standard Model) annihilation products only neutrinos will escape from the Sun; the strength of this indirect signal may be used to set cross section limits.

The currently leading measurement of this type was made using the IceCube neutrino telescope: an array of 5 160 photomultiplier tubes lying at 1450–2450 m depth in the Antarctic ice, over an area of about 1 km^2. Among other particle species, it detects muons above ~100 GeV energy via their Cherenkov light, and has an angular resolution of around 1 degree. To avoid the background of atmospheric muons, only neutrinos arriving from the direction of the Sun and passing first through the interior of the Earth are selected. The limits [69] set by IceCube in conjunction with the predecessor AMANDA experiment are shown in Fig. 1.11. These results are somewhat more model-dependent than traditional direct detection limits, due to the unknown annihilation branching fractions, but the approach is uniquely effective at high neutralino mass. An augmentation of the IceCube array with closer photomultiplier spacing, known as DeepCore, has recently been installed and should lower the threshold by a factor of ten in future acquisitions.

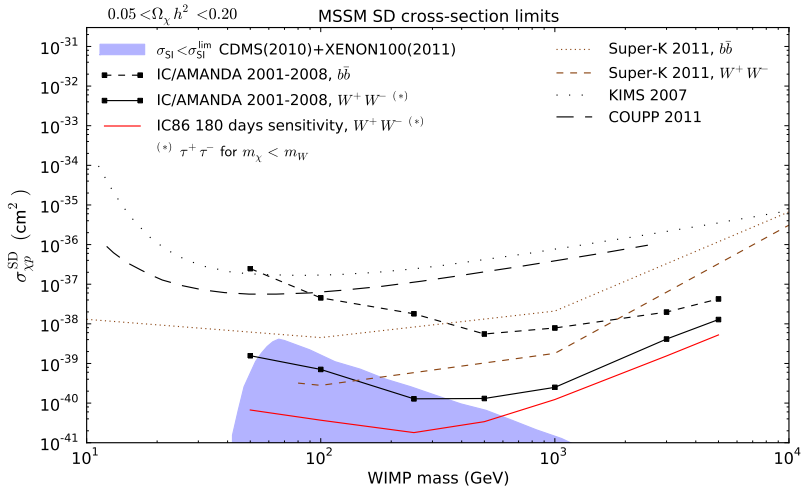

Figure 1.11: Upper limits (90% CL) on the WIMP-proton spin-dependent cross section from the IceCube and AMANDA experiments. Also shown are the results of solar DM annihilation searches from Super-K and direct search results from the KIMS CsI(Tl) detector and COUPP CF_3I bubble chamber. The shaded area is allowed MSSM parameter space and the red curve is an expected sensitivity for the completed IceCube detector. (From Ref. [69].)

An important indirect detection result was obtained by the Fermi Large Area Telescope (Fermi-LAT) in a recent study of 10 dwarf spheroidal satellite galaxies around the Milky Way [70]. Fermi-LAT is a pair-conversion space telescope that uses a tracker-calorimeter arrangement to measure γ-ray spectra from 20 MeV to over 300 GeV. Dwarf spheroidals are dark-matter-dominated systems with very little interstellar gas or recent star formation; as such they provide excellent sources for DM annihilation searches. Given some assumptions about halo shape, the strength of annihilation signal for each galaxy was predicted from velocity dispersion measurements, which then imply DM densities. No significant signals

were observed, and the limit curves, shown in Fig. 1.12 exclude the simple thermal relic annihilation cross section for light WIMPs ($\lesssim 20$ GeVc^{-2}) which decay dominantly to b, τ or W pairs.

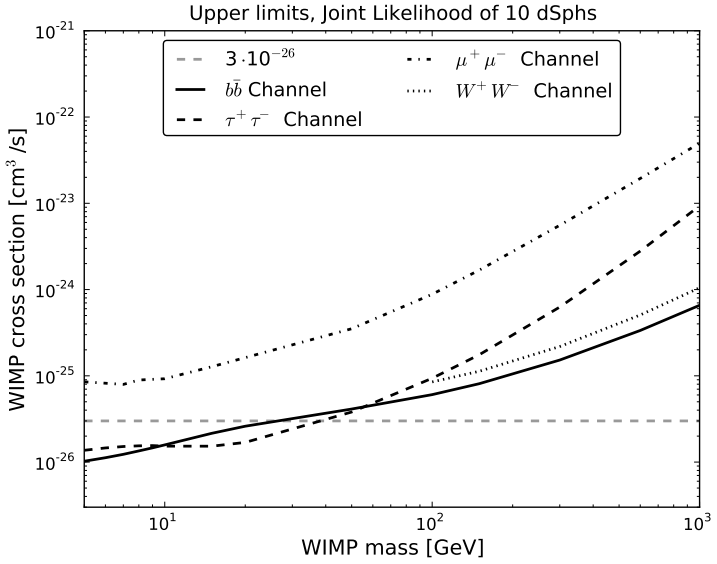

Figure 1.12: Derived 95% CL upper limits on WIMP annihilation cross section for a range of assumed annihilation products, set by the Fermi-LAT search for γ emission from satellite galaxies. (From Ref. [70].)

An analysis of public Fermi-LAT data within a few degrees of the Galactic Centre [71] found strong evidence for an extended source of 0.3–10 GeV γ-rays in addition to known point sources and contributions from the galactic disk. The spectrum is consistent with the annihilation products of WIMPs with a mass in the range 7–12 GeVc^{-2}if annihilation is predominantly to leptons, or 25–45 GeVc^{-2} for hadronic products. Astrophysical explanations have also been proposed [72]: the black hole at the Galactic Centre is expected to generate energetic cosmic ray protons which would go on to produce pions through collisions with the inter-stellar medium [73]; however, the fit to the observed γ-ray spectrum is relatively poor. A population of unresolved pulsars could also be responsible [74], although its density would have to fall much more rapidly with radius than the stellar population in general. This excess and other possible hints of signal in indirect search experiments are reviewed in Ref. [75].

In summary, the Λ-CDM paradigm succeeds in explaining physical phenomena over vastly different distance and time scales. A well-motivated class of candidate for the dark matter, the WIMP, reproduces the required density via a remarkable coincidence of magnitudes. An understanding of dark matter is being sought in underground labs, the Antarctic icecap, high-energy colliders, satellite experiments and elsewhere. Rapid progress is being made in the diverse experimental techniques relevant to the dark matter problem.

The remainder of this thesis is organised as follows. Chapter 2 introduces the ZEPLIN-

III experiment, including the principles of two-phase emission detectors in general. Chapter 3 describes the calibration of the instrument's response and includes measurements of yield functions which determine the energy scale for nuclear recoils. Chapter 4 describes the process of event selection, including the signal and background models which inform it. The WIMP-nucleon elastic cross section results from the second run of ZEPLIN-III are presented in Chapter 5, together with work on statistical techniques for rare-event searches of this type. A search for the signature of inelastic dark matter scattering in the first science run — using many techniques common to the elastic results, but aiming to test a particular explanation of the DAMA annual modulation — is presented in Chapter 6. Finally, the conclusions of the work are summarised.

My personal contributions are highlighted in the preamble to each chapter. Of publications by the ZEPLIN-III collaboration, I contributed text and figures to Refs [48, 76, 77, 78] and was corresponding author of Ref. [79]. In addition to work on analysis I spent many weeks in the lab at Boulby commissioning and operating the experiment. ZEPLIN-III was a highly collaborative endeavour by a team of 30 scientists, engineers and students at Imperial, Edinburgh University, Rutherford Appleton Laboratory, LIP-Coimbra and ITEP, Moscow. Work on my particular areas of responsibility was always informed by extensive discussions with colleagues and I, in turn, contributed to other parts of the programme which are not presented in detail here.

Chapter 2

ZEPLIN-III — hardware, software and operations

Two-phase xenon is one of the leading technologies which can achieve the threshold and background needed to observe WIMP-nucleon scattering definitively in the lab. At present, it is best placed to do so because it scales up well to large target masses. This chapter describes the physics underlying the approach, its specific implementation in the ZEPLIN-III apparatus, and the operation of that system at the Boulby mine.

2.1 Two-phase xenon time-projection chambers

The noble liquids produce detectable light and charge signals from energy deposited by nuclear radiation. In two-phase detectors, a large liquid volume is kept in thermal equilibrium with a thin layer of vapour above it. The condensed phase is a scintillator; ionised charge emitted into the gaseous phase is transduced into a second optical signal. In ZEPLIN-III, the depths of the liquid and vapour regions are about 36 and 4 mm, respectively.

Crucially, the light-to-charge ratio is far lower for recoiling nuclei, which would be caused by WIMP scattering, than for recoiling electrons, which are caused by the vast majority of background interactions. This is the key to two-phase xenon WIMP searches, giving a detector such as ZEPLIN-III, with a background count rate of order 0.1 Hz, sensitivity to a DM signal of perhaps ten nuclear recoils per year. Liquid noble elements in general offer this method of discrimination, and the possibility for large, homogeneous detector volumes. Liquid xenon holds several particular advantages as a target medium:

- A high density of $3\,\mathrm{g\,cm^{-3}}$, resulting in compact targets for a given mass and efficient self-shielding from liquid outside the fiducial volume.

- Essentially no radioactivity from xenon isotopes in their natural abundances. The only decay is that of ^{136}Xe (natural abundance 9%) via two-neutrino double-beta emission, recently observed for the first time by EXO-200 [80] and confirmed by KamLAND-Zen [81]; the latter measured a half-life of $(2.38 \pm 0.14) \times 10^{21}$ years.

- Extremely low overall activity achievable after depletion or distillation of ^{85}Kr.
- Good sensitivity to the spin-independent cross section ($\propto A^2$) due to the high atomic mass; and to spin-dependent WIMP-neutron cross section, as 48% abundance comes from neutron-odd isotopes.
- Moderate cryogenic requirements: having a triple point at 161 K, xenon can be maintained in a two-phase system with liquid nitrogen cooling.
- A high yield of scintillation light, at 72 ± 4 photons per keV electron recoil energy [82], with fast response time.
- The highest ionisation yield of the liquid noble elements: W — the average energy required to produce one electron-ion pair, in the high-field limit where recombination is negligible — is measured as (15.6 ± 0.3) eV using ^{207}Bi internal conversion electrons [83].

Recoiling electrons in liquid xenon lose energy predominantly to ionisation and atomic excitation, the latter leaving a trail of excited molecular states (excimers) produced along the particle track. For recoiling nuclei, processes of excitation-free energy loss — leading to atomic motion rather than any detectable signal — are also significant. The fraction of total nuclear recoil energy that is lost to electrons, either as excitation or ionisation, is known as the quenching factor $f_{\mathrm{n}}(E)$, and is adequately described by the semi-empirical model of Lindhard [84, 85].

In two-phase xenon systems, the scintillation and ionisation produced by energetic particles interacting in the liquid are both measured via VUV light. This concept was pioneered in the 1970s at the Moscow Engineering Physics Institute [86]. The wavelength of peak intensity is 178 nm [87], within the sensitive range of photomultipliers with quartz windows and bi-alkali photocathodes. Luminescence in the liquid (called the primary signal) originates from both direct excitation and charge recombination, described below. Ionisation electrons which are drifted away from the interaction site without recombining may be emitted into the gas phase, accelerated, and thus detected as a pulse of electroluminescence (the secondary signal). The depth in the liquid of the scattering vertex is measured by the drift time, i.e. the time elapsed between primary and secondary signals. We now examine the steps of primary and secondary signal generation in more detail, leading to definitions of the ZEPLIN-III energy variables.

2.1.1 Scintillation signal

VUV scintillation light is produced in xenon by the radiative dissociation of the two lowest states of the excimer (excited dimer) Xe_2^*: these singlet and triplet states have near-degenerate energies and decay times of 4 and 21 ns, respectively. A model of particle interactions following initial excitation and ionisation along recoil tracks in the liquid, illustrated in Fig. 2.1, is found in Ref. [88]. Following direct excitation, the initial exciton

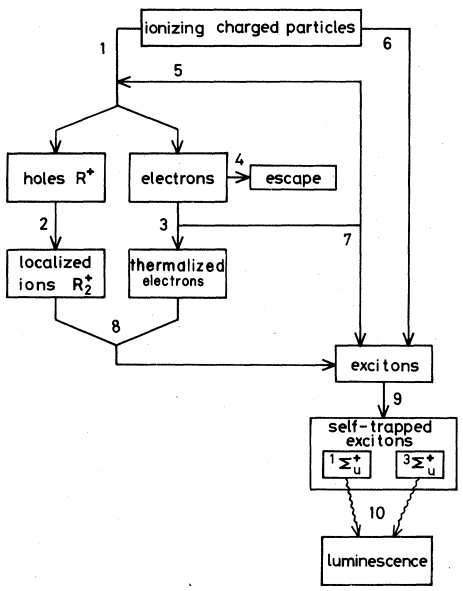

Figure 2.1: Processes leading to scintillation and ionisation in liquid argon, krypton and xenon (generically, element R). 1. Production of electron-hole pairs. 2. Formation of molecular ions. 3. Thermalisation of hot electrons. 4. Escape of electrons from the Coulomb attraction of ions. 5. Production of electron-hole pairs by secondary electrons. 6. Production of excitons by the ionising particle. 7. Production of excitons by secondary electrons. 8. Recombination of free electrons and molecular ions forming excitons. 9. Self-trapping of excitons forming excited molecules. 10. Radiative dissociation of excited molecules emitting VUV photons. (From Ref. [88]).

is rapidly (within a picosecond) trapped to form the excimer state. Following ionisation, the hole of an electron-hole pair is immediately localised as a Xe_2^+ molecular ion; the free electron loses energy to further excitation, ionisation, and phonons (thermalising on a timescale of $\sim$ 1 ns) and eventually either recombines with a molecular ion to form the excimer, or escapes the Coulomb attraction of the track altogether. The escaping charge goes on to be detected as the secondary signal in two-phase detectors.

The electric field dependence of the extracted charge and scintillation yields for different incident species is shown in Fig 2.2. The suppression of recombination along electron recoil tracks increases with electric field, up to at least 5 $\mathrm{kV\,cm^{-1}}$. For nuclear recoils, however, the yields reach their asymptotic high-field values below 1 $\mathrm{kV\,cm^{-1}}$. The difference is attributed to the structure of the tracks: an initial recoiling nucleus generates a branched sequence of secondary recoils. The much lower energy of the secondary recoils causes lower ionisation density, and hence less shielding of the applied field, than in the more cohesive tracks of electrons and α particles [89].

The scintillation variable measured by ZEPLIN-III — denoted S1 and typically chosen to set the energy scale of signal models — is the number of scintillation photons produced in the liquid (including recombination) relative to that for 122 keV electron recoils. Because of this normalisation, it is known as the 'electron-equivalent energy', and is measured in units of 'keVee'. The (non-linear) conversion between nuclear recoil energy and S1, i.e. the scintillation yield for nuclear recoils, will be discussed in section 3.5. For reference,

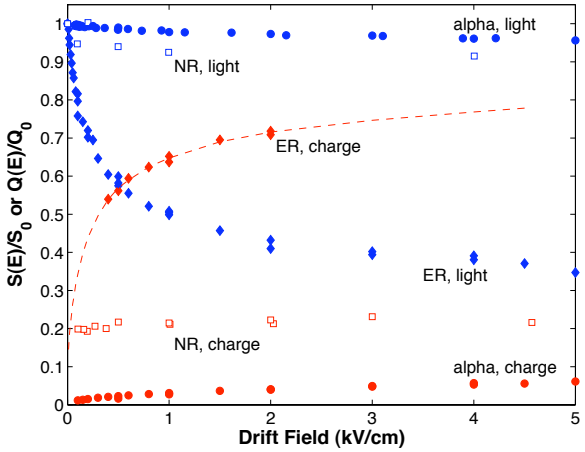

Figure 2.2: Dependence of scintillation and ionisation yields $S(E_\mathrm{l})$ and $Q(E_\mathrm{l})$ on the electric field in liquid xenon, for 122 keV electron recoils (ER), 56.5 keV nuclear recoils (NR) and 5.5 MeV α-particles. S_0 is the zero-field value for each species. For α-particles and electron recoils, Q_0 is the recoil energy divided by the W-value for electrons; for nuclear recoils this is reduced by a Lindhard quenching factor. (From Ref. [90]).

a WIMP search region of 2–12 keVee in scintillation corresponds to 7–29 keV in nuclear recoil energy.

2.1.2 Ionisation signal

The four stages in the detection of the ionisation signal — extraction, drift, emission and electroluminescence — all depend on electric field. In a two-electrode chamber with the cathode immersed in the liquid and the anode in the gas phase, an applied voltage, V, across a depth d_l of liquid and d_g of gas results in field stengths:

$$E_\mathrm{l} = \frac{V}{\frac{\epsilon_\mathrm{l}}{\epsilon_\mathrm{g}} \cdot d_\mathrm{g} + d_\mathrm{l}} \quad \text{and} \quad E_\mathrm{g} = \frac{\epsilon_\mathrm{l}}{\epsilon_\mathrm{g}} E_\mathrm{l} \tag{2.1}$$

in the liquid and gas, where $\epsilon_\mathrm{l} = 1.96$ and $\epsilon_\mathrm{g} = 1.00$ are the respective dielectric constants. During the second science run of ZEPLIN-III, E_l was 3.5 kV cm^{-1} and E_g was 6.8 kV cm^{-1}.

Some fraction of the ionised charge that is extracted from a track will, in practice, attach to electronegative impurities such as O_2 before reaching the liquid surface. This purity-dependent attenuation of the secondary signal is described by a free electron lifetime, τ_e. For a single contaminant species, it is a function of the attachment rate constant k in 1 mol^{-1} s^{-1}, which varies with species and electric field, and of the concentration S in mol l^{-1}: $\tau_e = (kS)^{-1}$ [91]. During the second science run of ZEPLIN-III, the electron lifetime ranged from 14 to more than 40 µs, compared to a maximum drift time in the fiducial volume of 13 µs. Before filling, the lifetime was measured as >100 µs in a dedicated portable lifetime monitor with lower field, suggesting that the dominant impurity has an attachment rate constant which increases with drift speed.

Upon reaching the surface, electrons have to overcome the electron affinity of the liquid (−0.67 eV). The efficiency for ionisation electrons which reach the liquid surface to be emitted into the gas phase therefore depends on the field in the liquid, rising from 50% for 2 kV cm^{-1} to approximately unity at 5 kV cm^{-1} [92, 90].

Electrons emitted into the gas phase generate the secondary signal via the process of proportional scintillation, or electroluminescence. The reduced secondary light yield, $\mathrm{d}L_s/\mathrm{d}x$, is defined as the number of photons generated by a single electron per unit path length, x, divided by the number density of Xe atoms, n [93]. Thus N_e electrons drifted across a gas gap of width d_g result in a number of secondary photons:

$$N_\gamma = n \cdot N_\mathrm{e} \cdot d_\mathrm{g} \frac{\mathrm{d}L_s}{\mathrm{d}x} \,. \tag{2.2}$$

The yield $\mathrm{d}L_s/\mathrm{d}x$ is found to have a linear dependence on the reduced electric field E_g/n above a threshold, as shown in Fig. 2.3. For the operating conditions of the ZEPLIN-III second science run (SSR), roughly 200 photons per emitted electron are expected, and the single electron signal is expected to be readily identifiable in PMT waveforms.

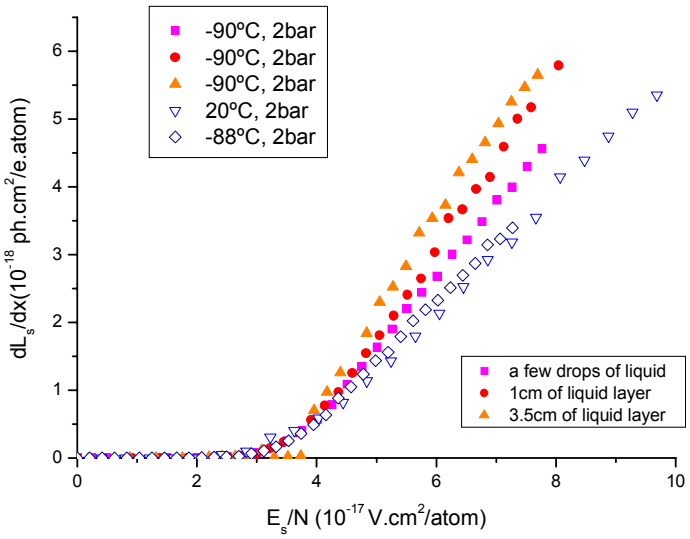

Figure 2.3: Measurements of the reduced secondary light yield in saturated xenon vapour, from Ref. [93]. In the SSR configuration of ZEPLIN-III, the number density of atoms in the 3.5 mm gas gap and the reduced field are 7.4×10^{19} cm^{-3} and 9.2×10^{-17} V cm^2, respectively .

The ionisation variable measured by ZEPLIN-III — denoted by S2 and used primarily to distinguish between electron and nuclear recoils — is the number of ionisation electrons leaving the track relative to that for 122 keV electron recoils. S2 is reconstructed from the electroluminescence signal recorded in the PMTs by accounting for position-dependent light collection, attenuation due to electron attachment and other effects on an event-by-event basis in analysis. The field dependence of the S2 yield for electron recoils is complementary to that of S1 and very weak for nuclear recoils, as discussed above: higher fields remove more charge from the track and so less is available for recombination

luminescence. The yield of charge from nuclear recoils is also energy dependent; the yield function will be described and measured in section 3.5.

In summary, nuclear recoils of tens of keV in a typical two-phase xenon emission detector are recorded as a prompt primary signal of up to tens of photoelectrons, with a characteristic time of ~20 ns, followed by a delay of microseconds due to the drift time in the liquid, then a secondary signal of hundreds of photoelectrons over ~1 µs of drift time in the gas.

2.2 ZEPLIN-III detector hardware

ZEPLIN-III is a two-phase xenon TPC with a disc-shaped active volume containing ~12 kg of liquid. The overall design of the system was driven by three central requirements:

- a threshold low enough to detect a significant fraction of the $\lesssim$100-keV energy deposits of WIMP-nuclear scattering. This relies on good light collection for primary scintillation signals, and sensitivity down to a few electrons in ionisation.
- a low rate of background events in the detector, achieved primarily with radiopure construction, shielding and deep underground operation.
- an extremely low rate of signal-like background, i.e. 'candidate' events which resemble nuclear recoils. This is achieved primarily by using the ionisation-to-scintillation ratio to discriminate against electron recoils, but it has consequences for essentially all aspects of detector construction and data acquisition.

Limits on WIMP scattering are set by comparing signal models with the observed event rate within a search box in the $\log_{10}(\mathrm{S2}/\mathrm{S1})$-S1 parameter space. The first science run of the instrument in 2008 produced a suite of competitive dark matter results [48, 94, 79]. In preparation for a second run the array of PMTs, which had been the dominant source of radiological background, was replaced using a custom designed model (ETEL D766Q) with activity of only 35–50 mBq per tube, 3% that of its predecessor. A new veto detector was also installed around the target, further reducing the effective background. Extensive details of the ZEPLIN-III hardware appear in Refs [95, 96, 78]. An overview of the second science run configuration is presented here, emphasizing upgraded features and those most relevant to the calibration and analysis described in later chapters.

2.2.1 Shielding and veto

ZEPLIN-III operated in the Palmer Laboratory at the Boulby Mine, 1070 m below the surface with an overburden of (2850 ± 20) m water-equivalent [97]. The instrument is shown in Fig. 2.4 surrounded by shielding. The flux of atmospheric muons in the lab is 3.8×10^{-8} cm^{-2} s^{-1}: muon-induced neutrons result in an expectation of under 0.01 events in the search box of the SSR and are in any case vetoed with high efficiency. Radiation originating in the external environment of lab and cavern is mainly due to ^{40}K as well as

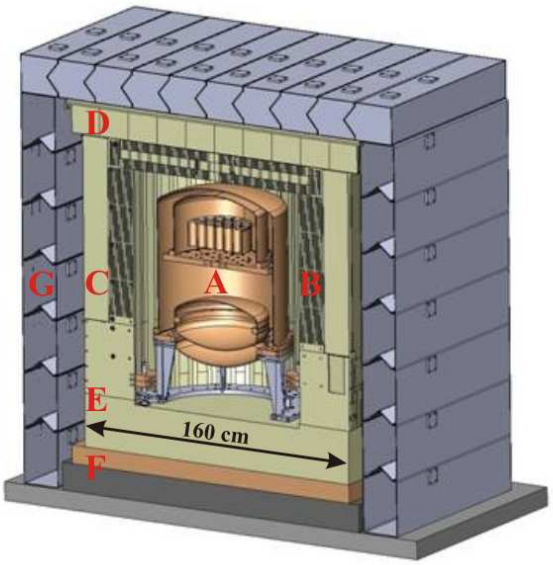

Figure 2.4: Cross section showing the configuration of veto and shielding, from Ref. [96]. A: the ZEPLIN-III instrument. B (hatched): gadolinium-loaded polypropylene sections. C: plastic scintillator wall, consisting of 32 blocks. D: plastic scintillator roof, consisting of 20 blocks. E: the lower polypropylene structure, without gadolinium. F: copper and lead base plates. G: lead castle.

trace amounts of U and Th decay chains in the rock; neutrons in particular originate from (α,n) reactions and ^{238}U fission. This external radioactivity is shielded by 20 cm of lead and 30 cm of plastic (including veto components), attenuating both neutrons and γ-rays by a factor of 10^5. As a result (see Table 2.1) external radioactivity adds negligibly to the rate of electron recoil events caused by decays within the apparatus, although it does contribute somewhat to the neutron background.

The active part of the veto instrument consists of 52 optically-isolated blocks of plastic scintillator, giving >3π solid angle coverage around and above the target. A PMT is bonded to the end of each scintillator block, at the far end from the ZEPLIN-III instrument in order to minimise radioactivity background. Immediately inside the active veto are blocks of gadolinium-loaded polypropylene; γ-rays emitted via radiative capture in this material are used to identify neutrons which may have scattered inside ZEPLIN-III, as discussed in section 4.3.1. The veto operates as an independent instrument and is described in Refs [96, 76]. The barrels of scintillator and Gd-loaded material rest on a pure polypropylene base structure, itself on top of a copper base plate. The outermost shielding layer consists of a lead castle and lead base plate with total mass greater than 60 tonnes.

2.2.2 The WIMP target

The xenon target chamber and a liquid nitrogen coolant reservoir are both enclosed in the same high-vacuum cryostat vessel, as shown in Fig. 2.5. The main thermal contact between the two is made by braids of copper wire, immersed in the liquid nitrogen and thermally anchored to the baseplate of the xenon chamber. Additional (and efficient) cooling power

and fine control are achieved by venting nitrogen boil-off vapour via a heat exchanger attached to the xenon vessel baseplate. A set-point temperature, measured on one of the conducting braids, is adjusted to maintain a target of around −100℃ and a pressure in the range 1.55–1.65 bar. The liquid level is measured during filling by capacitive level sensors located at the side of the active region.

At the bottom of the target chamber is an array of 31 2-inch photomultiplier tubes, submerged in liquid xenon and facing upwards to view the active region and the gas above. The fields E_l and E_g are both applied by an anode plate (at +9.5 kV) in the gas near the top of the chamber and a cathode wire grid (at −5.5 kV) 40 mm below it, in the liquid. The liquid surface lies 36.5 mm above the cathode, giving a maximum drift time in the liquid of 15.4 µs and a 3.5 mm gas gap. Another wire grid just above the PMT array and 5 mm below the cathode is held at −900 V, equal to the photocathode potential, to establish a reverse-field region from which ionisation electrons are not drifted to the surface. An inactive region such as this between the array and the active volume reduces the rate of background events from radioactive impurities in the PMTs; furthermore, this configuration protects the input optics of the photomultipliers from external fields.

The instrument is supported on three legs; two are threaded so that the detector can be relevelled to compensate for geological movement. This is done from outside the shielding via cables. Variation of the duration of the electoluminescence signal in the x-y plane reveals how parallel the liquid and anode surfaces are: if the detector is leaning south then the gas gap will be larger above the north side of the array and so the S2 pulses there will be longer and of larger area. Tilt measurements from daily calibration were used to decide on physical adjustments underground, and to correct for residual tilt during event analysis.

2.2.3 Experiment backgrounds

Table 2.1 summarizes the contributions to radioactive background from components of the system and its environment from a comprehensive study published in Ref. [98]: a long campaign of radioassay and simulation gave results in very good agreement with observed rates. The xenon has an extremely low ^{85}Kr/Xe atom fraction of 4×10^{-21} as it was extracted from underground sources in the 1970s, when the ^{85}Kr abundance was lower than today, and has undergone cryogenic distillation. Structural components of ZEPLIN-III, including the target and liquid nitrogen vessels and the vacuum dome, are made of high-purity oxygen-free (C103) copper totalling some 400 kg. The leading sources of electron-recoil background are the PMTs and the plastic shielding; the low level of neutron background comes mainly from ceramic feedthroughs and the PMTs.

2.2.4 Signal readout

The PMTs are supplied from a common voltage divider, with 10 dynode stages between photocathode and anode. The dynodes are supplied via a stack of equipotential copper

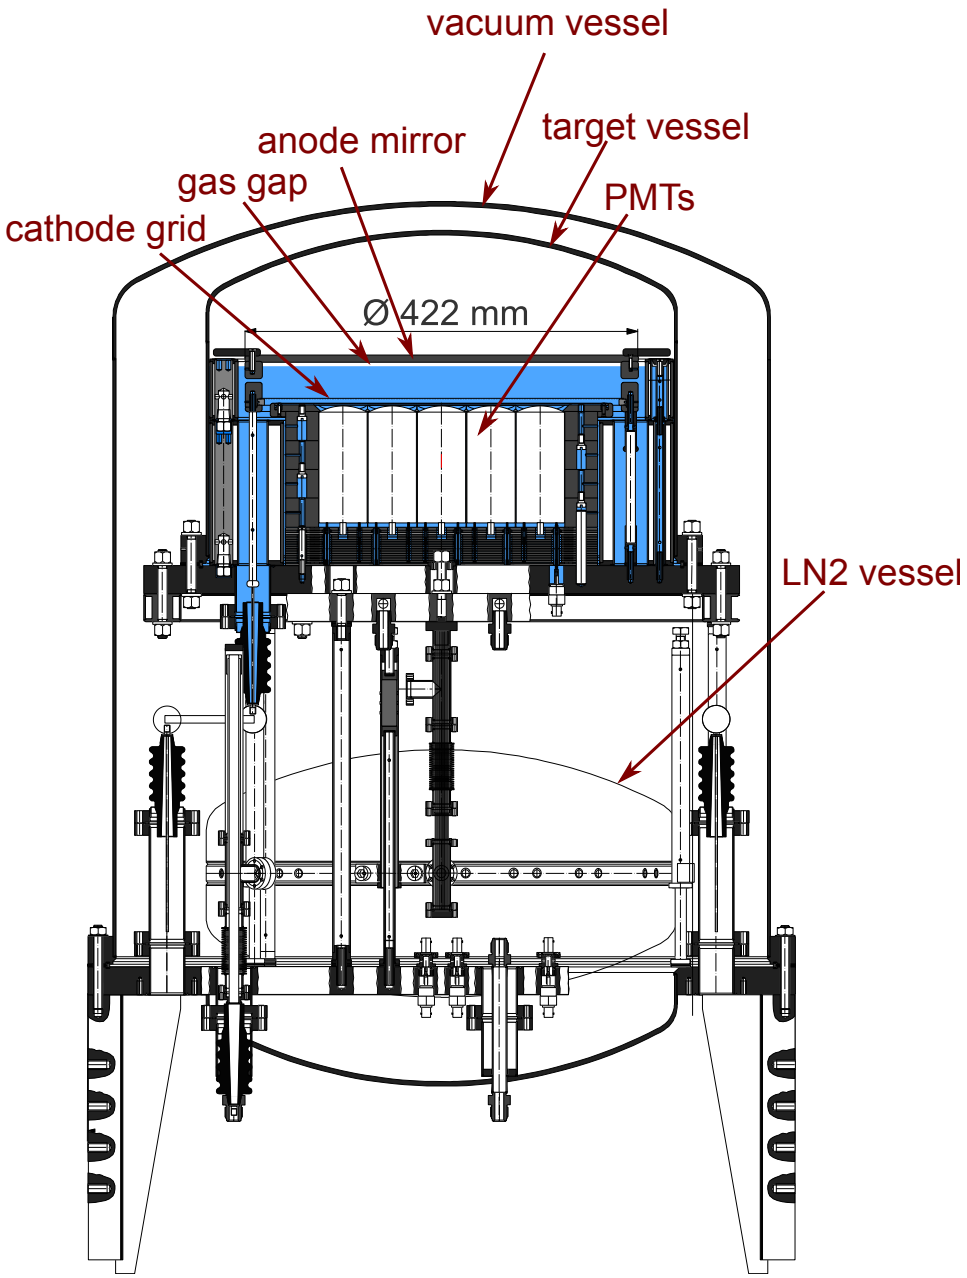

Figure 2.5: Schematic drawing of the ZEPLIN-III instrument, highlighting selected components. Liquid xenon is shown as the blue shaded region.

Table 2.1: Component-level analysis of electron-recoil and single-scatter neutron elastic scattering backgrounds in the SSR, adapted from [98]. Simulated expectations are shown except for the bottom row which contains the observed rate in search data. The efficiencies with which electron-recoil and neutron background events are identified as such by the veto are denoted by 'ptag' and 'dtag', respectively. The search region of the SSR has an expectation of only 0.06±0.01 un-vetoed neutron events (much lower than the SSR total in the table, which counts events over a wider energy range, before efficiencies from analysis cuts and veto tagging).

Material	mass, kg	e-recoil, dru[†]	ptag	n/year[‡]	dtag
^{85}Kr in liquid	–	0.007	~0	–	–
Ceramic feedthroughs	0.9	0.08	0.30	1.35	0.58
Photomultipliers	4.2	0.40	0.26	0.74	0.58
Rock (halite)	–	~0	~0	0.53	0.58
Polypropylene shield	1,266	0.25	0.04	0.10	0.58
Scintillator modules	1,057	0.09	~1	0.03	~1
Copperware	~400	<0.1	0.10	<0.15	0.58
Lead castle	~60,000	0.01	0.54	~0	0.58
^{222}Rn inside castle (1 m^3)	–	0.03	0.19	~0	–
Muon-induced	–	–		~0.3	~1
SSR total [*]		0.86±0.05	0.28	3.05±0.5	0.58
SSR data		0.75±0.05	0.28	n/a	–

[†] events/kg/day/keVee at 10 keVee electron recoil energy
[‡] single scatters in 2 370 kg·days over 5–50 keV nuclear recoil energy
[*] neglecting material with only upper-limits on activity

plates (visible below the PMT array in Fig. 2.5): this has the advantage of avoiding electronics in the xenon chamber (preserving purity) and reducing the number of feedthroughs (minimising background) but does not allow equalisation of the responses at the PMT anodes. The waveform from each PMT is read out via the electronics chain illustrated in Fig. 2.6, which may be adjusted to equalise the responses of the different channels. This process will be described in section 3.1 of the next chapter. High- and low-sensitivity (HS and LS) readouts of each PMT, with ×10 difference in gain, are recorded in 8-bit DAQ channels with 2 ns sampling, using Acqiris (DC265) hardware and custom acquisition software. Search data were acquired with 100 mV full scale in HS and 200 mV in LS; events within the signal region, having relatively small primaries and secondaries, are reconstructed solely from the HS channels. Event acquisition is triggered by the sum of the 31 HS channels: a Canberra 2111 timing filter amplifier, with integration time 500 ns, shapes the summed pulses, which are fed to a discriminator.

Acquisition will trigger on the primary signal if S1 $\gtrsim$ 30 keVee, or otherwise on the secondary — for events with enough S1 photoelectrons to be reconstructed in software, there is essentially unity efficiency from the much larger electroluminescence signal, with a trigger threshold of ~5 emitted electrons. The recorded waveforms span 36 µs, beginning 20 µs before the trigger point, to fully contain the longest drift times whether triggering on S1 or S2. There is a dead-time of ~18 ms, dominated by the DAQ crate hardware.

PMT waveforms, together with slow-control data on pressures, temperatures and the high-voltage system, are written to local disks by the DAQ PC, and constitute about 1 MB per event before compression. These binary files are transfered to analysis machines only after the end of the acquisition, to avoid additional dead-time; for the main WIMP search run, this happened during the daily nitrogen fill.

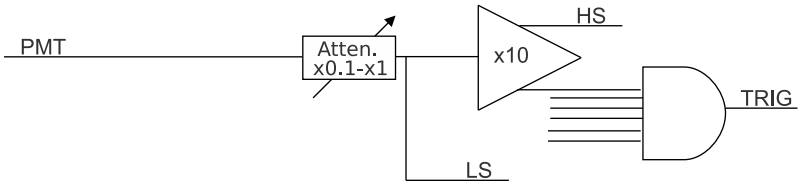

Figure 2.6: Electronics chain for a single PMT, up to the HS and LS DAQ channels and the trigger summing amplifier. For the 12 lowest-gain tubes, an additional ×10 wideband amplifier (Philips Scientific 770) is included before the variable attenuator (Philips Scientific 804).

2.3 Slow control

The complete apparatus is operated via a system of slow control which acts on and logs dozens of parameters, automates many daily operations and implements safety procedures in order to allow remote running of the system and very good stability. The following processes are carried out by the slow control system:

- Readout and control of temperature and pressure in the target, including an over-

pressure protection system capable of leaking xenon to external emergency capture tanks.

- Control of the high-voltage systems, including slow ramping and safety cut-outs to protect equipment in the event of discharging.
- Monitoring of the vacuum in the cryostat and the power consumption of the pump which maintains it.
- Measurement of liquid xenon level via capacitive sensors.
- Monitoring of the lab environment, including temperatures of electronics systems.
- Daily calibration, with automated delivery of a ^{57}Co source to within the shielding.
- Daily filling of the liquid nitrogen reservoir.
- Automatic data acquisition and logging of acquired files to a data repository database.

Operational data was accessible via a web interface which allowed custom plotting of the various parameters. One year of continual two-phase operation is testament to the success of the slow control.

2.4 Event reconstruction

The ZEPLIN-III DAQ records 2×31 PMT waveforms, plus the sum and trigger channels which are used to confirm event matching with the veto DAQ. Event reconstruction is the process of identifying, from the waveforms, the S1 and S2 pulses of single-scatter events in the active volume and then estimating the interaction positions and energies. This is achieved in three stages of software analysis, which are described fully in the following sections:

- **ZE3RA** — the ZEPLIN-III data Reduction and Analysis package, written in C++— finds and clusters together pulses. It characterises them and the full waveforms via numerous parameters but it makes no physical interpretation of the underlying event. The input is raw PMT traces; the output is reduced n-tuples describing up to ten clusters of pulses per event, for all triggered events.
- The **Golden Code**, a FORTRAN program, identifies 'golden' events: those consisting of one S1 followed by one S2 signal which pass a set of mild quality cuts on pulse shape and baseline noise. It also makes a provisional reconstruction of position and electron-equivalent energies. Corrections for position- and time-dependent detector response are applied at this stage. The input is ZE3RA n-tuples and the output is golden n-tuples characterising the S1 and S2 signals and the state of the detector for those events with S1 < 70 keVee. This dataset is much smaller than the ZE3RA dataset of all triggered events in search data, and so more manageable for interactive analysis.

- The **Mercury** program, written in C++, makes maximum-likelihood estimates of the position and electron-equivalent energy of both S1 and S2, as identified by the Golden Code, by fitting PMT response functions to the pulse areas measured by ZE3RA. Multiple fitting algorithms are supported, and parameters describing the quality of fit are returned along with the point estimates.

Data from the veto detector is reduced (by the **raVen** software) and synchronised with golden events (by the **vesync software**), and appended to the outputs of the three ZEPLIN-III codes. These final n-tuples contain over 150 named variables, several of which are one- or two-dimensional arrays with an entry for each PMT or pulse; they occupy roughly 10 kB per event, compressed. An entire search dataset in this reduced form can be analysed interactively in ROOT or PAW. When required, complete waveforms for any triggered event can be easily retrieved from the (multi-terabyte) binary data repository and viewed using an interactive GUI mode of ZE3RA. An example of a low-energy nuclear recoil interaction is shown in Fig. 2.7 on page 51.

2.4.1 Waveform parametrisation

The primary scintillation from a low-energy WIMP-nucleus scatter may generate just a few photoelectrons spread over tens of nanoseconds. The secondary electroluminescence signal may generate tens of photoelectrons over hundreds of nanoseconds. Clustering (over time) and matching (across channels) these signals in the presence of noise and thermal photoelectrons is a significant challenge. It was addressed in the dedicated software package ZE3RA, developed by the Coimbra group and described in detail in Ref. [99]. The task is complicated by the presence of electrical cross-talk and PMT after-pulsing — in a WIMP scattering search, quite unlikely configurations of noise can conspire to destroy sensitivity.

The analysis of each waveform $V(t)$ begins with characterising its baseline. A time window for this purpose is defined by progressively including up to forty 50-ns regions, beginning at the start of the waveform, until there is evidence for something other than baseline noise; this is decided by an F-test on the consistency of the variance in successive regions. The baseline variances and means (σ_i^2 and μ_i in the ith channel) are used in the pulse-finding algorithms and are also written to file for the purpose of detector monitoring and event quality cuts.

A pulse — the lowest-level component in signal identification — is a section of waveform lying below some threshold voltage, in practice $V_{\mathrm{thresh}} = \mu_i - 3\sigma_i$, shown as a solid green line in Fig. 2.7. A cluster is a set of one or more pulses in a single waveform attributed to one physical process in the xenon (e.g. multiple photoelectrons generated by a single scintillation signal). The clustering algorithm implemented in ZEPLIN-III is based on the time interval between the end of one pulse and the beginning of the next. Successive pulses separated by intervals smaller than a certain t_{gap} are clustered together, until no next pulse is found within t_{gap}, or a maximum total extent for the cluster is reached. This

maximum extent is a fixed multiple of the mean signal arrival time (or centre of area, known as τ) in the cluster so far; this constraint mitigates the effect of after-pulsing and amplifier ringing.

ZE3RA uses low-pass filtering as part of the clustering process. Pulses are searched for first in a filtered waveform, $\bar{V}(t)$, and these will automatically cluster multiple pulses in the raw waveform $V(t)$. Several filters are supported in software. In practice a moving average was used, with the sample width adjusted dynamically in the range 70–200 ns — based on the variance of $V(t)$ in the previous sample — to better resolve pulses of various widths.

In the final stage of signal identification within ZE3RA, clusters in the waveforms of different PMTs are matched up as belonging to the same signal. A summed waveform is calculated in software (accounting for relative delays between channels, and distinct from the shaped sum produced in hardware by a summing amplifier) and the clusters in it are matched with the clusters found in individual channels. The matching algorithm, using a figure of merit derived from the integral of the product of waveforms, is detailed in Ref. [99].

The ten largest-area clusters in the sum are selected, and the parameters of all the individual PMT clusters which contribute to these are written to an n-tuple file. (In golden events, clusters outside the top two are typically the result of PMT afterpulsing or dark noise, or single ionisation electrons from photoionisation or optical feedback.) The output thus includes arrays such as `area(i,j)`, containing the area of the cluster in the `i`th PMT waveform that contributes to the `j`th largest cluster in the sum. Other variables include timing parameters, amplitude and measures of symmetry and saturation.

Beyond its data-reduction functions, ZE3RA provides a customisable and information-rich waveform viewer, illustrated in Fig. 2.7, and a GUI interface to runtime reduction options. A blindness manager is also implemented in ZE3RA, so that only open events — known not to be destined for the WIMP search region — are shown when examining waveforms or reduced. The process of unblinding search events is discussed in section 4.4.2.

2.4.2 Event classification

Only a small fraction of triggered events is in fact relevant to the WIMP search, and one role of the Golden Code is to filter out uninteresting data that account for the majority of events in the reduced n-tuples. For a reasonable range of S1 < 70 keVee, about 250 000 events survive the Golden Code, out of more than 10^7 triggered events in the SSR. An initial selection of plausible events (the first pass) is made by requiring that they:

- do not coincide with periods of momentary high count rates or HV discharging, detected through electrode or PMT currents;
- do not begin at large baseline excursions or have particularly large rms noise;
- contain exactly one short cluster followed by one long cluster (τ below and above

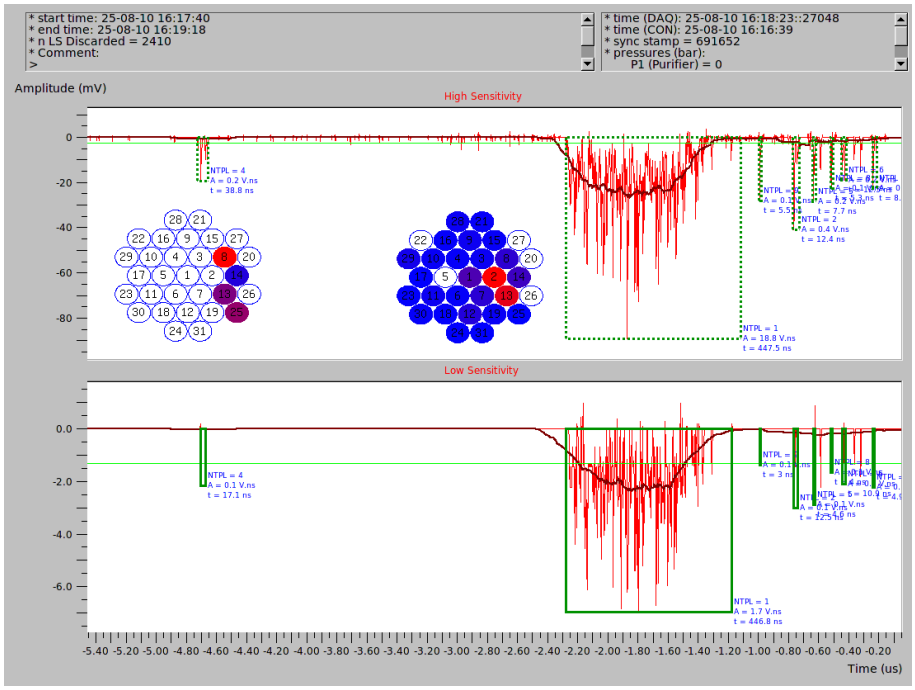

Figure 2.7: A segment of the sum waveform containing a nuclear recoil event, shown in the ZE3RA GUI. This neutron calibration event was reconstructed within the WIMP search box at S1 = 2.9 keVee, S2 = 1.0 keVee. The smoothed waveform is shown in maroon, and the 10 clusters written to n-tuple are boxed in green. The array maps show which tubes contributed to the S1 and S2 signals, red (blue) corresponding to large (small) pulse area. Some afterpulsing can be seen following S2, and it has been correctly identified as separate from the electroluminescence cluster.

300 ns), to be consistent with the structure of single scatters. Exceptions are made via time windows for afterpulsing in the PMTs and optical feedback, whereby light from the large electroluminescence signal causes electrons to be emitted from the cathode grid via the photoelectic effect.

To ensure that only genuine scintillation signals are considered, the candidate S1 cluster must have a pulse area above 1/3 of a photoelectron in at least three PMTs. This three-fold coincidence requirement determines the threshold of the detector in two-phase operation. The full set of first pass criteria are met by 38% of triggered events.; most rejected events at this stage are multiple scatters with more than one S2 pulse or S1-only events caused by interactions from which no ionisation charge can be collected, e.g. below the cathode grid.

The survivors of the first pass have clusters identified as possible S1 and S2 signals; a second pass analysis in the Golden Code makes cuts based on the detailed properties of these two clusters, and estimates the amount of initial scintillation and ionisation. Additional parameters are calculated, such as the centroid of the primary and secondary clusters in x-y, and corrections to detector response due to the time-varying pressure, tilt and purity of the target. PMTs' responses are multiplied by flat-fielding coefficients that compensate for variations in gain and quantum efficiency; the drift time is calculated to give the z-coordinate of the vertex and the electron lifetime correction; and preliminary light-collection factors are applied, using the identity of the peak PMT response as a proxy for x-y position. At this stage the S1 and S2 estimates known as golden energies are obtained by summing weighted areas. The second pass cuts exclude probable mis-identifications of S1 and S2 through various tests of pulse shape, and a requirement that any PMT recording more than 35% of the S1 light must record at least 1.5% of the S2. At this stage, 34% of triggered events remain. Finally, events above a maximum golden energy in S1 (70 keVee) are discarded to leave 2.4% of original triggers. Surviving events, now described by several dozen named variables, are written to golden n-tuples and passed to Mercury for more precise reconstruction.

2.4.3 Reconstructing scintillation and ionisation

The vast majority of events included in the golden n-tuples are good single scatters, with position and energy fairly well estimated in the Golden Code. However better precision, and hence better discrimination against background, can be achieved, especially at low energy, by making maximum-likelihood fits to a detailed detector response model. This is done by the Mercury program, detailed in Ref. [100]. In addition, mis-reconstructed events which have passed the Golden Code can be removed at this stage based on goodness-of-fit. Cuts of this type are powerful for flat detector geometries such as ZEPLIN-III, where multiple scatters in the liquid are often identifiable by the pattern of light across the array.

The reconstruction algorithms rely on light response functions, LRFs: they describe the mean pulse area in the ith PMT per unit electron-equivalent energy, as a function

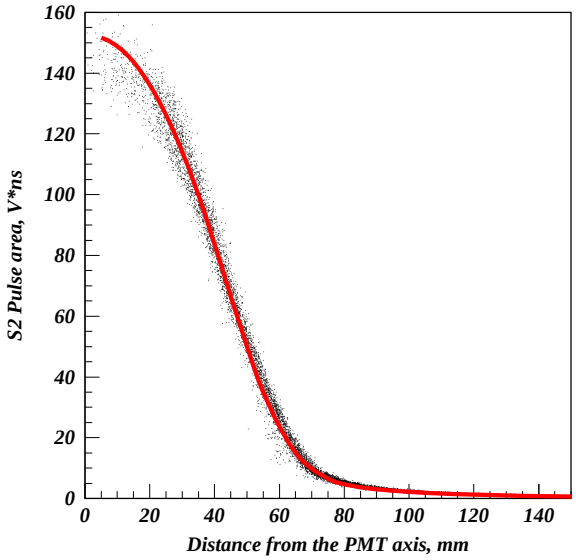

Figure 2.8: The light response function (red line) of a typical PMT, as found by the Mercury program after five iterations of fitting to ^{57}Co photoabsorption events. The data are shown as points at their reconstructed radii, demonstrating good convergence of the spline LRF model. For reference, the radii of the photocathode and PMT envelope are 23 and 26 mm, respectively. (From Ref. [100].)

of distance from the PMT axis, and are denoted by $\eta_i(r_i)$. Either S1 or S2 can be reconstructed from the overall scaling of the response (using the relevant LRFs for primary or secondary light) and E is used here to stand for either. The response functions for S1 light vary smoothly with the z-coordinate of the vertex in the liquid; this is measured independently and to high precision — tens of µm FWHM — by the drift time between the S1 and S2 signals, evaluated in the Golden Code. Thus, z is not a fit parameter within Mercury but it is used as an input to choose the correct response functions. The LRFs fold in the effects of initial scintillation yield, geometrical light collection and quantum efficiency, and are parametrised as splines. An example is shown in Fig 2.8. We find that all channels conform to the same shape of LRF and show good radial symmetry; this is because the side walls of the chamber are relatively far from the array and have low reflectivity. E is varied along with x and y in a fit to the recorded pulse areas.

To calculate a joint likelihood for the observed pulse areas, A_i in the ith PMT, the response model must include distributions of the single photoelectron area, q_i. A gamma distribution is shown in Ref. [101] to be a good description of the single photoelectron distribution in our PMTs, and is adopted in Mercury with parameters derived from an LED light calibration. The discrete Poisson probability to observe n photoelectrons is converted to a gamma distribution in pulse area, with mean $n{\cdot}q_i$; the sum of these functions for all n gives the expected distribution of pulse area in the tube, denoted $p(A_i; E, x, y)$. The overall likelihood function is the product of the 31 p functions, and (E, x, y) are

estimated by their maximum-likelihood (ML) values.

This process of likelihood fitting is essentially how Mercury reconstructs primaries and secondaries, with a few technical refinements. A simplification is possible for the comparatively large S2 signals: the responses of tubes close to the vertex are approximately Gaussian in this case, and so a simple weighted least squares fit can be used. Furthermore, the x and y coordinates are measured with much greater precision for S2 than for S1, but it is known that the true positions must coincide for good events; therefore, the fiducial S1 energy variable is a fit to the S1 responses assuming the position found from S2. Events with a large discrepancy between the two independently-reconstructed positions are cut as likely multiple scatters. Many other parameters describing quality of fit are returned, including overall likelihoods, the expected pulse areas in all tubes for the ML reconstruction, and the number of iterations taken for the minimisation routine to converge.

One of the novel features of Mercury is its ability to obtain the light response functions from calibration events, the positions of which need not be known a priori. In other contexts, e.g. gamma cameras used for medical imaging, a source may be systematically scanned over the volume to establish the LRFs but this is impracticable in ZEPLIN-III. Instead, Mercury is trained on samples of ^{57}Co photoabsorption events of known energy. In the first reconstruction, positions are estimated by a centroid method and LRFs are fitted to the profile of area versus estimated distance for each tube. These preliminary LRFs are then used to re-estimate event positions in the next reconstruction, after which new, improved LRFs are fitted to the area-distance profiles; this process is repeated until convergence is reached. The quality of LRFs obtained in this way was verified in the second science run by reconstructing the shadow of a new copper grid positioned on top of the anode, between the calibration source and the liquid. As shown in Fig. 2.9, this was achieved out to ~150 mm from the array centre. The sharpness of the grid image was used to put an upper limit on the position resolution of S2 at 122 keVee: 1.6 mm FWHM was achieved within the inner 100 mm radius, rising to 3.0 mm at 150 mm. The energy resolution that can be achieved is demonstrated in Fig. 2.10, which shows first science run calibration data reanalysed with a later version of Mercury developed for the second run (v2.3). PMT performance (including uniformity of the photocathodes, electrical pickup in channels requiring large additional gain and the quality of the photoelectron area distributions) was significantly poorer in the SSR and degraded the equivalent resolution by a factor of ~2.

2.5 Underground operations

The Palmer lab was opened in the year 2000 at the Boulby mine, a working salt and potash mine operated by Cleveland Potash Ltd. It is a short (<1 km), if dusty, walk away from the mineshaft and provides an air-conditioned, HEPA-filtered environment with a bridge crane and network connectivity throughout. The lab has hosted a number of other direct DM detection experiments: DRIFT-I [102], DRIFT-II [56] and ZEPLIN-II [49]. It is

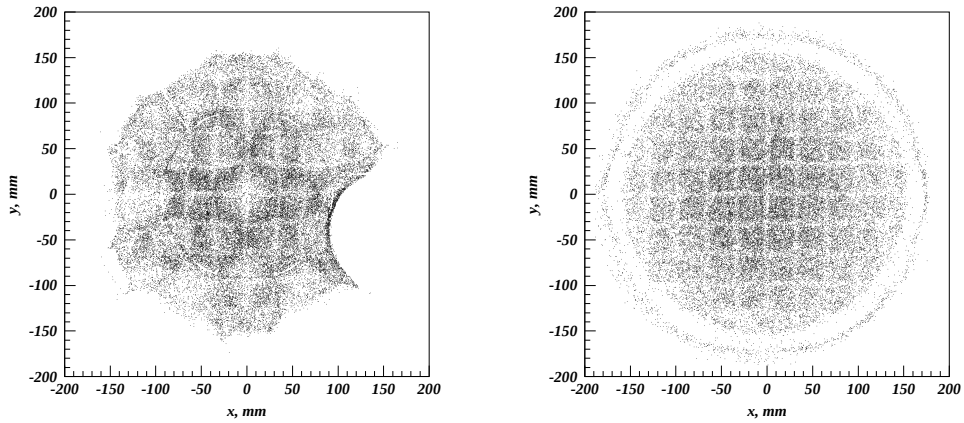

Figure 2.9: Reconstructed positions of ^{57}Co events using a centroid estimate (left) and in Mercury (right), during the second science run. The low-density strips are the shadow of a 5.1-mm thick copper grid, with bars 5 mm wide, located on top of the 7-mm thick copper anode plate; the grid is well-reconstructed even above a failed PMT centred at x = 133 mm, y = −36 mm. (From Ref. [100].)

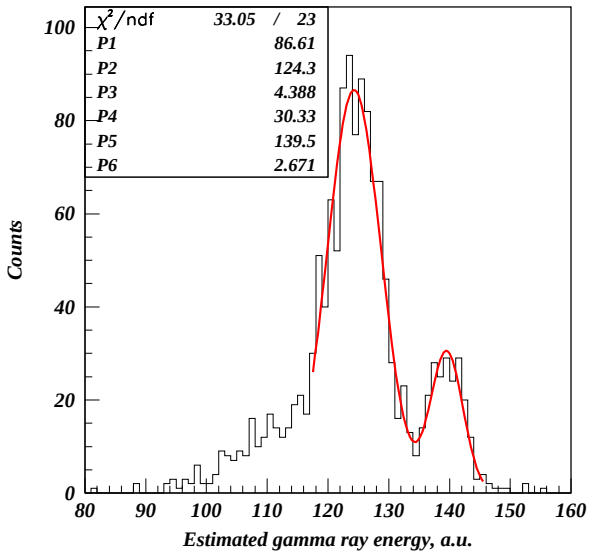

Figure 2.10: Reconstructed spectrum of ^{57}Co events with $r < 50$ mm in a combined Mercury energy variable (0.25·S1 + 0.75·S2) for FSR data. The energy resolution is 3.5% at 122 keVee. The 122 and 136 keV lines are clearly resolved, which is not achieved in the golden energy variables. (From Ref. [100].)

also home to a low-background counting facility and the SKY experiments on ion-induced aerosol nucleation.

2.5.1 Commissioning

The ZEPLIN-III instrument was upgraded and reassembled for the second science run in an underground clean room, located in a side stub of the main Palmer lab. Multiple vacuum failures in the new batch of PMTs — some of which occurred when ZEPLIN-III had already been installed, cooled and filled at its final location — demanded repeated reassembly over a period of months. The final rebuild was completed in January 2010. Commissioning of the full system then continued for several months, during which time the shielding, veto, xenon gas system, cooling system, source delivery mechanism and cabling were all integrated with the target; DAQ channels were calibrated, including relative delays, and PMT responses were equalised. In parallel, the system was vacuum pumped for 2 weeks in preparation for cooling.

An important prerequisite for running in search mode is sufficient xenon purity to reconstruct S2 for WIMP-nucleon events throughout the active volume. An electron lifetime of 14 µs was achieved during commissioning by recirculating the xenon gas through a hot getter. Minimal outgassing from the carefully-chosen construction materials and the 'sweeping' effect of a continuously applied field helped to maintain and improve this purity during the run without external purification.

When the target chamber was first filled for the SSR, it was found that N_2 contamination in the gas was significantly suppressing the electroluminescence signal. This was solved by a process of rapid dilution: xenon from the gas phase was vented to external storage tanks in short bursts, taking gaseous nitrogen with it. Four iterations of gas expulsion and re-equilibration increased the S2 pulse area by a factor of 40 to around 300 photoelectrons per keVee at the surface, sufficient to trigger the system and reconstruct S2 for all nuclear recoil events above threshold in S1. This N_2 contamination is believed to have been introduced with the additional xenon required for the SSR, the upgraded PMTs having slightly smaller volume than the FSR units.

2.5.2 Operations during the second science run

Including commissioning and calibration, the second run of ZEPLIN-III lasted for over one year of two-phase operation. Running the detector for so long, and mainly from hundreds of miles away, relied on extensive systems of slow control, as described in section 2.3. A subset of slow control parameters were included in the DAQ data files and used to make corrections for varying detector response and to identify periods of discharging or other instability. Significant excursions in slow-control parameters were removed by the Golden Code; in addition, analysis of search events with outlying values of $\log_{10}(^{S2}/_{S1})$ sought out more subtle correlations which might have degraded background discrimination, although no significant effects were found.

While the slow control monitored the internal and external environment of the target, the responses to scintillation and ionisation were measured directly by means of γ-ray calibration. Each day, a ^{57}Co source was moved into position above the vacuum vessel by a computer-controlled delivery mechanism, and 35200 events were acquired over 20–30 minutes. The dominant 122 and 136 keV photons emitted by the source interact mainly in the first few millimetres of liquid to generate electron recoils of known energy distribution. The data were immediately analysed with ZE3RA, Golden and Mercury codes on an underground PC, and plots of electron lifetime, tilt and the S1 and S2 responses were automatically e-mailed to the team members on watch.

After the commissioning phase, manual operations underground were carried out on a single day per week by a two-person team. In most weeks, science data were acquired with a duty cycle above 90%: acquisition was interrupted for about one hour per day of ^{57}Co calibration, data transfer and filling of the internal LN_2 reservoir, and for an additional few hours per week while the underground team filled the external LN_2 storage dewar and carried out LED calibration of both sets of PMTs (target and veto). Over the 319 calendar days of the science run, ZEPLIN-III operated in search mode for 271 days of live acquisition time. The following three chapters describe the process of searching for a WIMP signal in the resulting ten million triggered events, beginning with calibration of the scintillation and ionisation signals.

Chapter 3

Detector response

With the design and underlying principles of the ZEPLIN-III instrument established, this chapter describes the work of calibrating the detector, i.e. measurements, hardware adjustments and analysis aiming to reduce bias in the reconstruction of energy deposits. Calibrating nuclear recoil energy from the observed S1 and S2 pulses is divided into two stages: firstly, characterising detector response (i.e. the recorded PMT pulse areas) to a unit of scintillation light or drifted ionisation charge anywhere in the active volume and at any time in the SSR; and secondly, measuring the energy-dependent yields of scintillation and ionisation due to a recoiling nucleus. Within the collaboration I made leading contributions to early measurement and equalisation of the single photoelectron responses, the investigation of energy reconstruction via radioactivity lines, and the correction for drifting gains in the DAQ that were uncovered during the SSR. I also contributed substantially to the measurement of nuclear recoil yields [77], which was led by the Imperial group, for instance in the modeling of efficiencies and evaluating statistical uncertainty.

3.1 Characterisation of photomultiplier responses

Although individual PMT gains are accounted for in the Golden and Mercury reconstruction codes, it is nevertheless helpful to equalise the recorded single photoelectron (SPE) responses across the different tubes. The shaped sum of PMT signals is used to trigger event acquisition, and so very unequal PMT responses would cause the hardware trigger threshold to depend on the event position in x-y. Variation in quantum efficiency will also contribute to trigger variability, but the effect is relatively small: the mean and standard deviation QE for the array are 0.26 and 0.06, respectively.

Approximate equalisation was done in signal readout electronics, based on the spectra of background and thermal photoelectrons (dark counts) during detector commissioning. Later, mean SPE areas were measured more accurately, and monitored over the run, by other methods: they were extracted from search data via Poisson statistics — counting observed zeros to find the mean number of photoelectrons corresponding to an expected pulse area from Mercury — and also from weekly calibration with pulsed LED light.

3.1.1 SPE areas from single-channel triggering

The PMTs were operated at 900 V anode-cathode voltage; higher voltages were found to cause current instability and power supply cut-out. Since the electrode voltages in all tubes are supplied from a common voltage divider, responses were adjusted via the electronics chain following the PMT anodes. Signals pass through this chain before reaching the trigger and low-sensitivity DAQ channels; for each PMT it consists of a wideband rotary attenuator (having 10 linear steps from 0.1 to 1.0 gain) and, optionally, a 10× wideband amplifier or a barrel attenuator (of 6, 10 or 20 dB power attenuation). The high-sensitivity DAQ channels are each fed via a final 10× amplifier.

During detector commissioning SPE areas were measured for each tube in single-channel trigger mode with the ZEPLIN-III DAQ. Data were acquired with no radioactive or light source present, but before the detector shielding had been installed. A trigger threshold of −2 mV was used for most channels, of smaller magnitude than SPE-like pulses observed away from trigger point. Due to baseline noise, channels 2 and 29 required thresholds of −10 mV and −5 mV respectively, which is likely to have introduced some bias to the estimated peaks.

Examples of pulse area spectra, as parametrised by ZE3RA, are shown in Fig. 3.1. The upper panels show distributions in area versus mean charge arrival time (i.e. the time integral of $tV\,(t)$ divided by that of integral of $V\,(t)$); the characteristic sizes and timings of single-photoelectron signals result in peaks in these parameters. The bin of pulse area 0 nVs for PMT2 corresponds to events where ZE3RA finds no pulse in the waveform, i.e. probable triggering on baseline noise. Some spectra shapes were far from ideal SPE curves – having long tails towards high area, or low peak-to-valley ratios – but fitted Gaussian peak positions were obtained for all channels and used for equalisation. Systematic bias in the method and the discreteness of attenuation choice led to some residual variation, but adjustment of the electronics chains achieved a sufficiently good result: a trigger level could be set such that no region of the detector caused either (i) significant loss of efficiency by missing signal-like S2s (see section 4.2.2) or (ii) impractically high trigger rates due to, for example, large responses to single extracted electrons. The set of 30 functioning PMTs installed for the SSR had an extremely wide variation in internal gain, as shown in Fig. 3.2, significantly at odds with the agreed specification. The high electronics gain required for certain channels did increase pickup and amplifier noise compared to the FSR array, which came to pose significant problems during event reconstruction.

3.1.2 SPE areas from Poisson statistics

The SPE responses may be measured with great precision from two-phase datasets by exploiting the event reconstruction achieved in Mercury using iterative fits to PMT light-response functions (LRFs being the expected pulse areas per unit scintillation as a function of vertex position, see section 2.4.3). Both the detector environment and the mechanism of generation (VUV photons as opposed to longer-wavelength light or thermal emission)

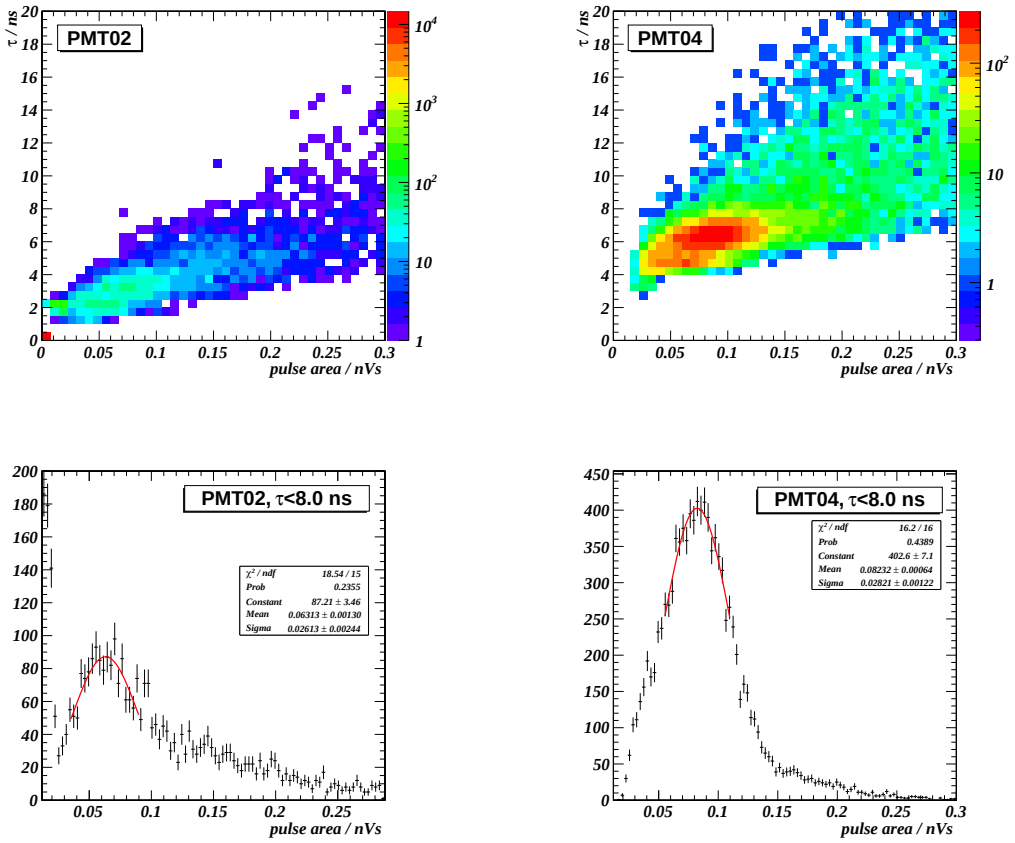

Figure 3.1: Examples of SPE responses in single-channel trigger mode, used for hardware equalisation before the SSR. PMT 2 (left) has the smallest and least distinct response of all tubes; in PMT 4 (right) it is well separated from baseline noise. Top row: pulse distributions in mean charge arrival time, τ, and area, showing the SPE feature. Bottom: pulse-area histograms with Gaussian peak fits.

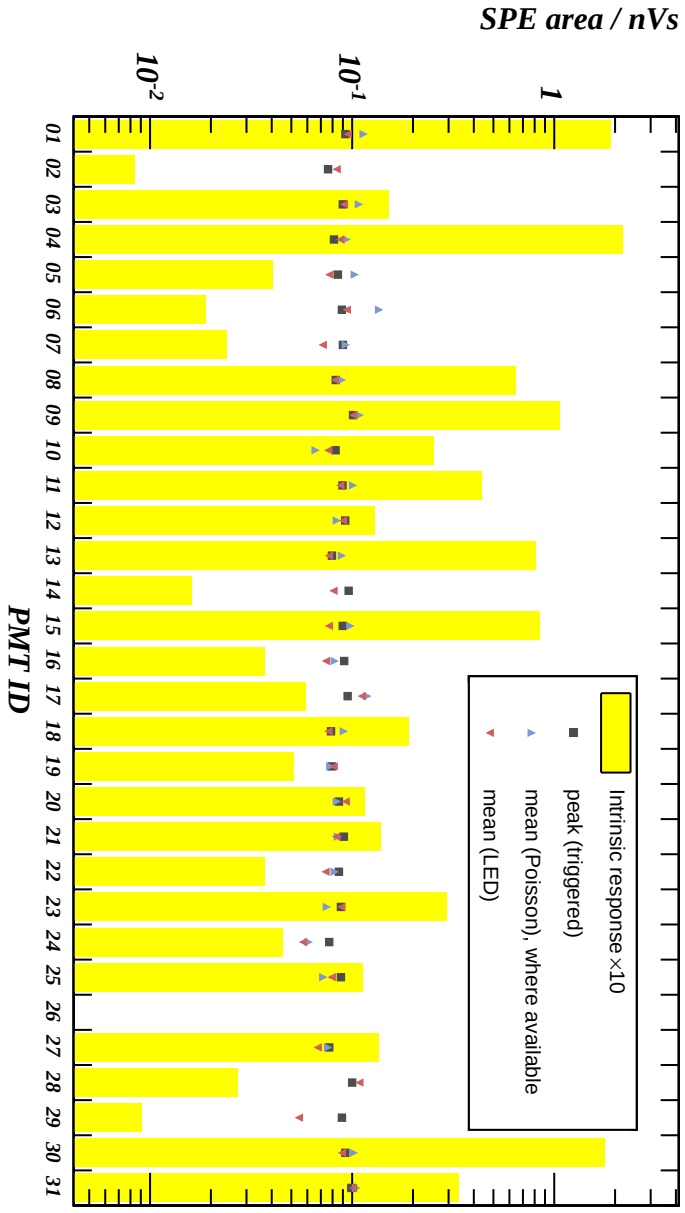

Figure 3.2: The single-photoelecton responses in the SSR array, as estimated by single-channel triggering, Poisson and LED methods (described in the text). Bars show the inferred SPE areas at the anode, found by dividing out the nominal gain of the signal chain.

are systematically the same as for signal events. In addition, linearity of the pulse-area response to small numbers of photoelectrons can be confirmed. The general approach is to treat the reconstructed energy of an event, dominated by PMTs close to the vertex, as a measurement of the expected pulse area in a distant PMT. A bin of expected area in the distant PMT corresponds to a sample of events with the same (small, unknown) expected number of photoelectrons; the expected photoelectron count for events in a bin of expected pulse area (and, hence, the mean area of one photoelectron) can then be determined via the fraction of events in that bin which produced no signal in the distant PMT.

Measurement of PMT i begins with binning reconstructed S1s of medium size (γ-ray calibration events of 30–70 keVee in Fig. 3.3) by the expected pulse area in the ith channel, $\langle A_i \rangle$, based on the Mercury fit. If the ith channel records only a small fraction of the signal, the reconstructed $\langle A_i \rangle$ is approximately fixed by the other channel responses. The S1 pulses in a bin of $\langle A_i \rangle$ are therefore treated as Poisson trials with a constant mean number of SPEs, μ, and hence a constant probability of zero being observed. Under the simplest model — assuming that some threshold amplitude V_{thresh} is smaller than all actual SPEs and larger than all noise — good results are obtained for PMTs with large inherent gain, as shown in Fig. 3.3. In this case, for a bin containing N events of which N_0 have no coincident pulse above V_{thresh}, the estimate of μ is $\log\left({}^{N}/_{N_0}\right)$. Ref. [103] reports the complete analysis in which the RMS noise of the baseline is evaluated and used to derive V_{thresh} for each event. This technique was successfully applied to 26 of the SSR tubes, with the results plotted in Fig. 3.2. In that work it was also demonstrated that the widths of pulse area distributions for a given energy could be obtained by considering the Poisson and SPE contributions to the variance for a given μ; this informs models of energy resolution in the few-photoelectron region.

3.1.3 LED calibration

A third and final measurement of SPE size was allowed by one of the SSR hardware upgrades: a set of light-emitting diodes coupled to the chamber via a quartz optical fibre. During weekly calibration, pulses of 1 µs duration at 100 Hz were shone into the target and the DAQ acquired, in synchrony, timelines of 3.5 µs for all channels. The mean response in any given PMT is small, at most 4×10^{-3} photoelectrons per light pulse. In data analysis, two 0.5-µs windows were selected from each timeline: a dark window before the LED fires and a light window during the LED pulse. Figure 3.4 shows spectra of ZE3RA-parametrised pulses in these two windows; the difference is interpreted as the spectrum of light-stimulated SPE responses.

The three methods described have different practical advantages and systematic uncertainties. Peaks in triggered spectra are quick to acquire and analyse while hardware adjustments are being made, but they are not high-accuracy measurements: systematic error arises from the selection effects of the trigger and neglecting or approximating the spectrum of noise. Poisson-based measurements with search data can characterise the

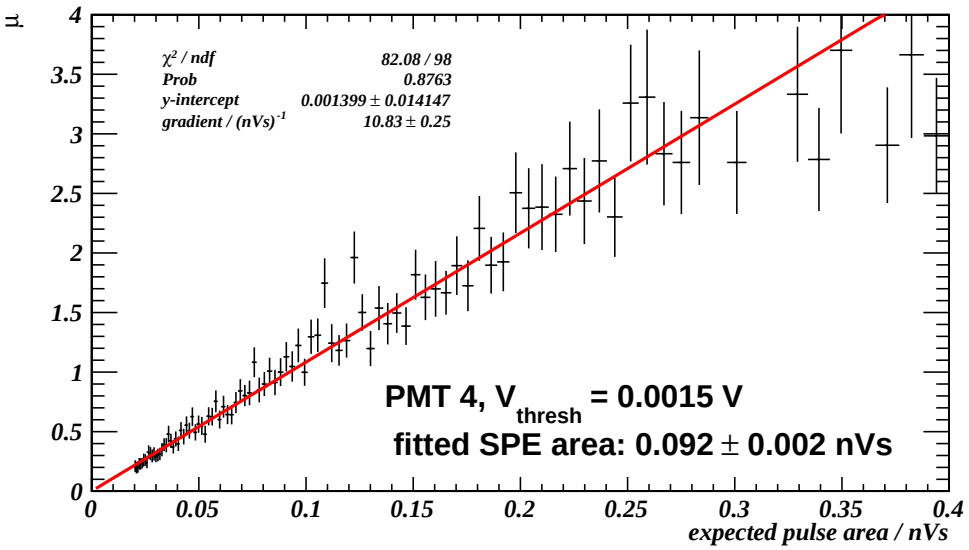

Figure 3.3: A precision measurement of mean SPE area for PMT4 based on Mercury reconstruction of S1 in ^{137}Cs calibration data. The mean number of photoelectrons, μ, for events in a bin of $\langle A \rangle$ is obtained from the fraction of empty waveforms (i.e. having no coincident pulse above V_{thresh} in PMT 4). The error shown is statistical; a larger systematic uncertainty of around 0.005 nVs arises due to the range of V_{thresh} which gives a good fit.

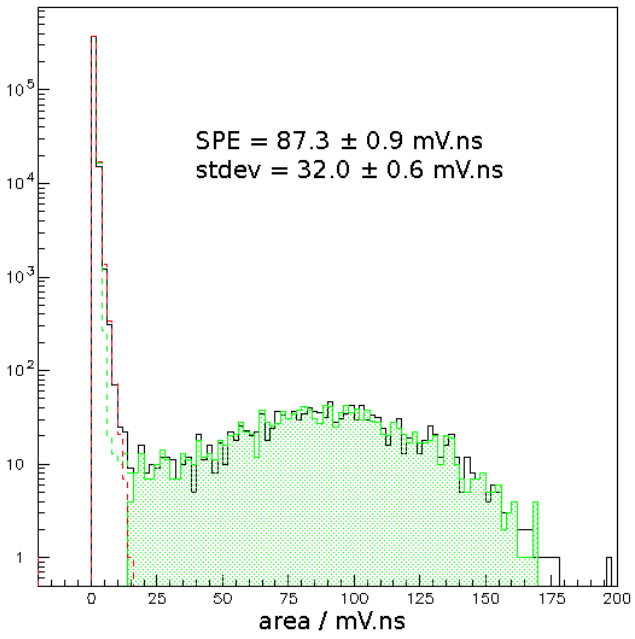

Figure 3.4: Example of photoelectron response during calibration with blue LED light. The red histogram shows the area of the highest-amplitude pulse during the dark window and the green histogram shows the same for the light-coincident window. The black histogram shows the summed area of *all* pulses in the light window, confirming that the probability of two photoelectrons during one light pulse is negligible.

PMTs in exactly the operating conditions of interest; the method can also be used to look for variation in response during high-rate calibration runs. Uniquely, it gives additional information about the linearity of response to multiple photoelectrons. However, it relies on the performance of event reconstruction which may not be good in the early stages of commissioning. Where suitable hardware is available, LED calibration achieves an accurate single-photoelectron measurement that is independent of the rest of the array.

3.2 Channel response drift

The ADC channels of the DAQ system, which record the waveforms from a PMT after they have passed through the amplifier chain, may be internally calibrated against an inbuilt precision voltage source. This is recommended to avoid variation in the gain of internal components, for instance due to temperature, changing the pulse size recorded from an input waveform. It was noted part way through the SSR that the ratio of recorded areas in the HS and LS channels (for a single physical pulse) was drifting in some tubes. This was traced to a change in DAQ software that had removed initial auto-calibration from the routine of data acquisition. Reprogramming resolved this issue for the rest of the run, but the first four months of data already taken required correction in software.

The ADC is not the only stage in the path from energy deposit to recorded pulses that could vary during a run of many months: the post-anode electronics chain, the quantum efficiency and gain of the PMTs or even the light yield of the target, for instance via surface reflectivity, could conceivably drift over time. The S1 signals due to 122 keV photoabsorption are a calibration point that probes all of these effects, and they are within the usable range of both the high- and low-sensitivity DAQ channels in the PMT immediately below the vertex (see Fig. 3.5). These events also define the unit of primary scintillation in which the signal spectrum is calculated. A channel-by-channel gain correction was applied to obtain a constant response to these signals over time. Variations affecting only the S2 response are treated by separate corrections, summarized in section 3.3.2.

Gain correction

On every day of the SSR, an automated slow-control script prompted the source-delivery system to insert a ^{57}Co source to the same position above the vacuum dome, offset by 5 cm of the centre of the array (see Appendix A.1 for the source spectrum). For each channel and in each day's ^{57}Co data, the peak in pulse area is estimated for the S1s of events positioned within 25 mm of the PMT central axis in the x-y plane. The maximum-likelihood S2 position from Mercury is used, giving resolution of < 3 mm FWHM. An example fit to one day's data is shown in Fig. 3.6, and the variation over the run, for two PMTs, is shown in Fig. 3.7.

This peak position measures the response of a channel, in pulse area, to a known

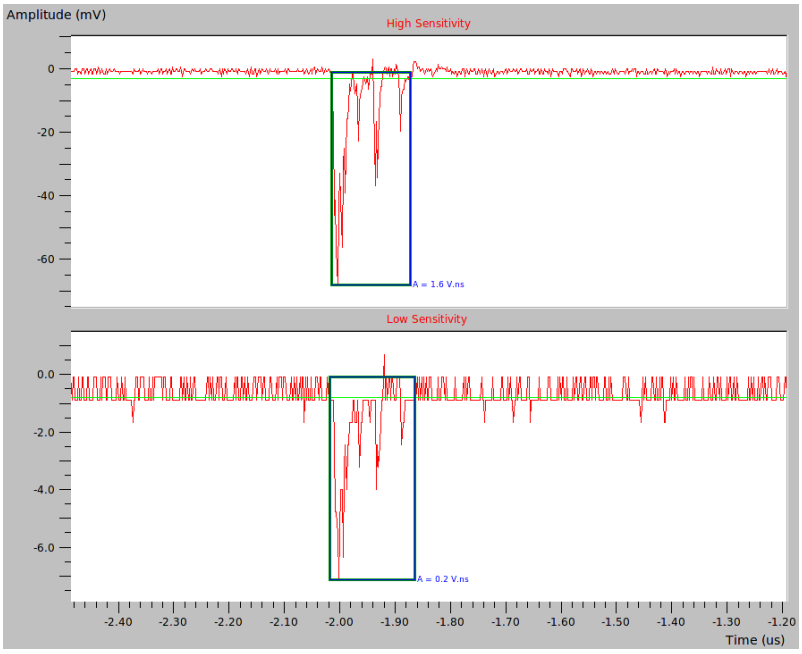

Figure 3.5: Example of an S1 waveform recorded in the high- and low-sensitivity channels of PMT 3, due to 122 keV photoabsorption above that tube. Each channel is digitised with an 8-bit, 200 mV scale during ^{57}Co calibration.

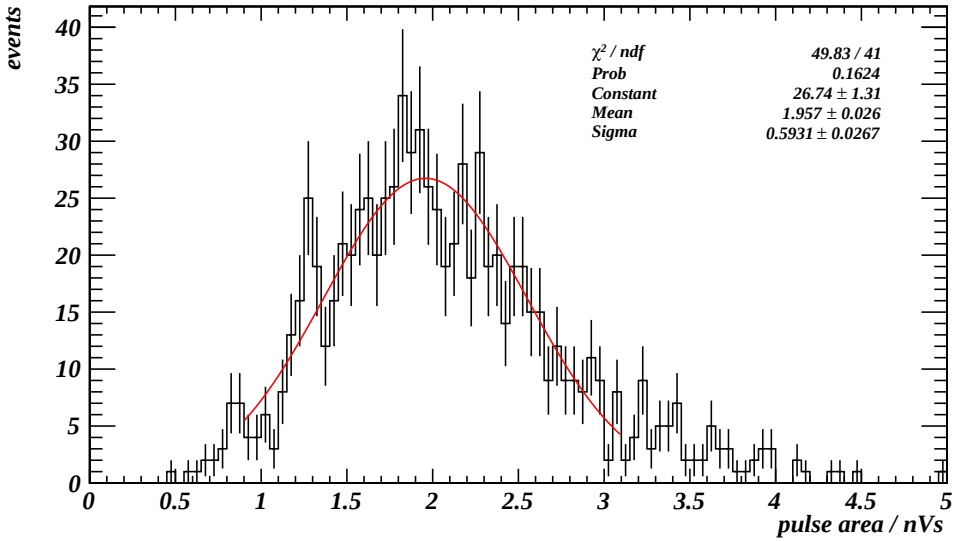

Figure 3.6: S1 pulse areas in the high-sensitivity channel of PMT 3 due to 122 keV photoabsorption above that tube. Data are from a standard daily calibration (1st May 2011). A maximum-likelihood Gaussian fit to the peak, used to calculate the gain correction is shown.

electron recoil energy. The response is built up from the scintillation yield in the liquid, the fraction of scintillation light reaching the photocathode, the quantum efficiency of the tube, the charge arriving at the anode per photoelectron, the gain of the electronics chain, and the internal gain of the DAQ. A shift in a channel's response can be caused by variation in any of these, but correlation across channels can narrow down the physical origin. For example, trends that appear in the HS but not the LS channel of one PMT, and in the period when the DAQ was auto-calibrating, may be ascribed to the HS amplifier; this effect, at the level of $\lesssim 5\%$, was seen for PMTs number 1 and 19. Similarly a trend seen in all tubes could indicate a uniform scaling in incident yield due to a change in electric field or xenon transparency; no consistent global trend was seen in data. The clearest variations in response were due to the drifting internal calibration in the DAQ in the first months of the run. Several channels showed a step change of 5–10% on the day when DAQ autocalibration was first reinstated. In addition several PMTs, including number 3 as shown in Fig. 3.7, exhibited a steady degradation of a few percent over the run, unrelated to the DAQ drift.

In order to implement a channel-gain correction, a reference response for each PMT was defined as the area peak in the high-sensitivity channel on 30th October 2010. On this day, a longer dataset of 88 000 events was acquired. The time variation of the fitted peak position on other days was smoothed using the robust locally-weighted regression routine LOWESS [104, 105]. A set of 2×30 correction factors was calculated for each day by comparing the smoothed peak positions in HS (and 10 times those in LS) to the reference value for the given tube. All pulse areas calculated by ZE3RA from raw waveforms were scaled by these factors before being analysed in the Golden and Mercury codes, giving an RMS correction over all pulses in search data of 4.7%.

3.3 Calibrating scintillation and ionisation variables

From calibrating individual channel responses, we move to the energy variables calculated from the array as a whole, S1 and S2. These estimates of scintillation and ionisation are given relative to the ^{57}Co calibration signal: 122 keVee in S1 (S2) is defined as the mean amount of scintillation light (ionisation charge) that leaves the interaction site after photoabsorption of a 122 keV γ-ray. The full spectrum of γ-rays emitted during ^{57}Co calibration includes a 136 keV line with 0.125 times the intensity of the 122 keV line, and some of those absorbed in the liquid will have Compton scattered on the way. A simulated distribution of energy deposits with depth in the liquid is shown in Fig. 3.8.

3.3.1 Energy variables

The resolutions of the S1 and S2 measurements merge the observed ^{57}Co events into a single peak centred on a mean energy around 124 keV. The calibration of these energy scales is shown in Fig. 3.9, where the simulated and observed peaks are matched, via a

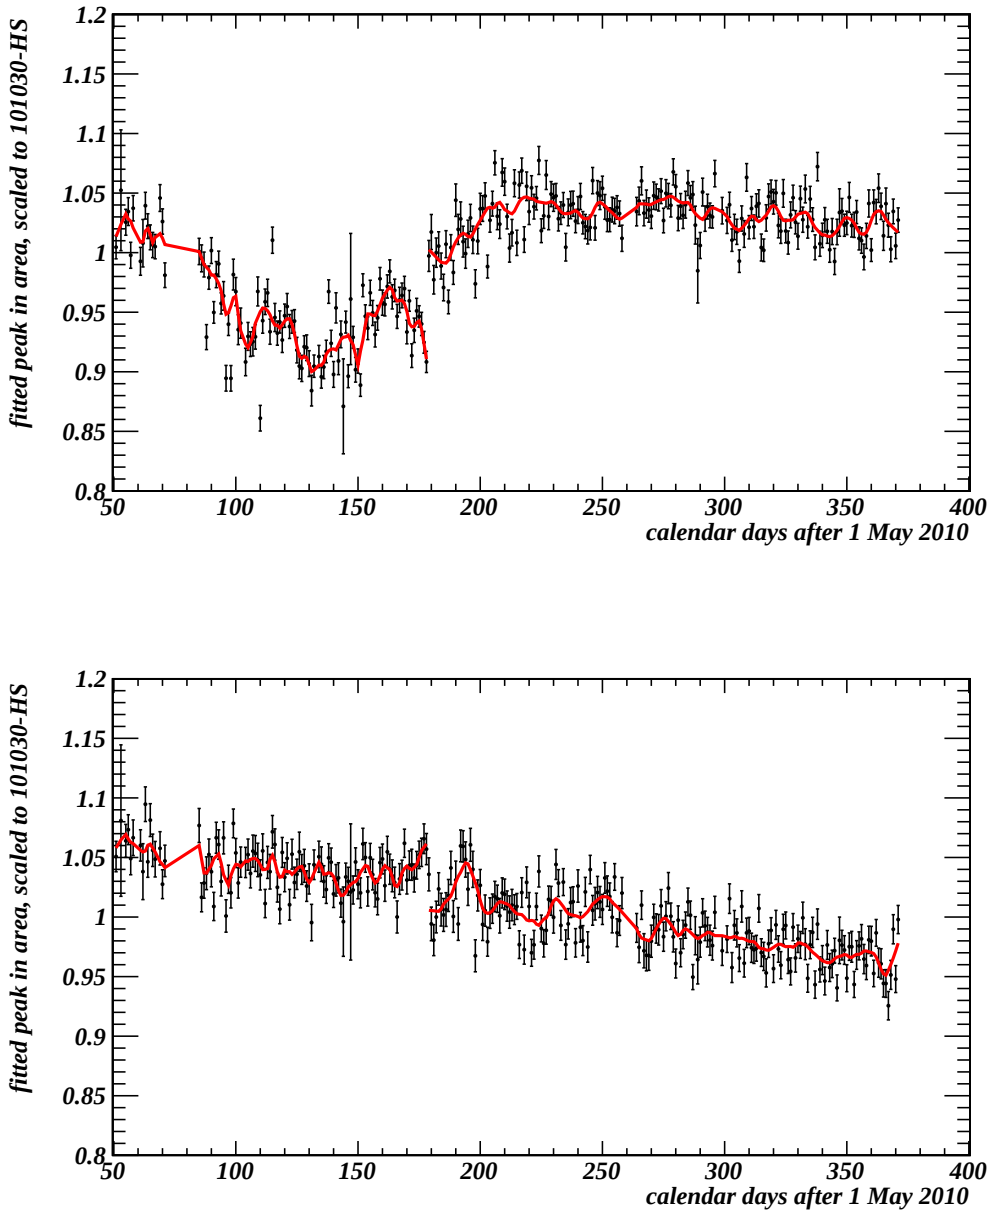

Figure 3.7: Normalised position of the fitted area peak (the reciprocal of the gain correction) for the high-sensitivity channels of PMT 1 (top) and 3 (bottom). The discontinuity in the smoothed line coincides with the resumption of daily DAQ autocalibration on 27th October 2010 (day 179 on the x-axis). The high-sensitivity channel of PMT 1 showed the largest DAQ drift of all channels. The slowly declining response of PMT 3 was also seen in the low-sensitivity channel.

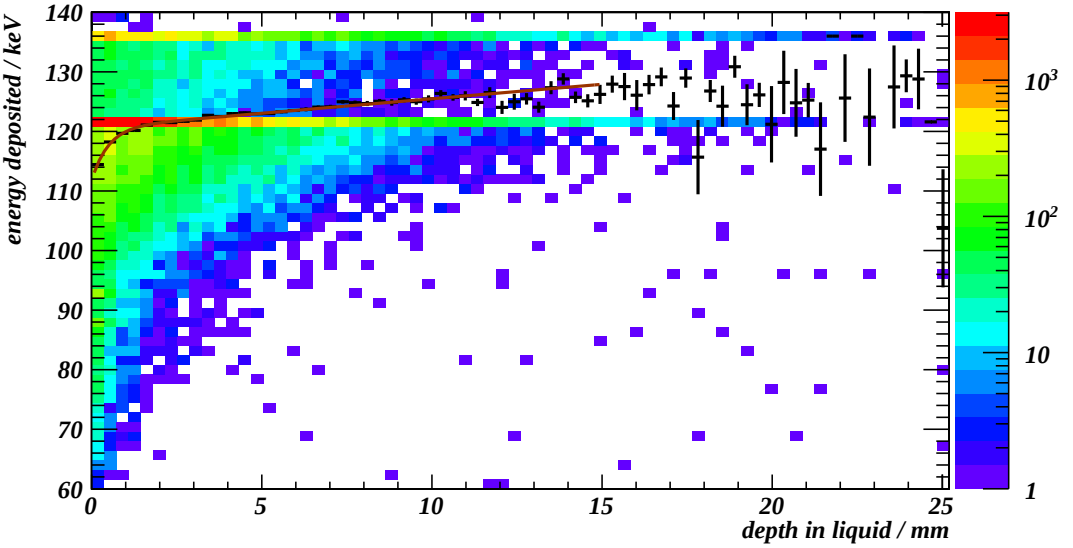

Figure 3.8: The simulated distribution of events in energy and depth during ^{57}Co calibration is shown in colour. Superimposed are the mean energies for each depth bin, and an empirical (straight line plus exponential) fit to their variation. The fiducial volume of the detector extends from 2.4 to 31.2 mm depth in the liquid, but daily calibration events are concentrated towards the surface.

simple resolution model. In terms of mean pulse area in the sum channel, 1 keVee in S1 corresponds to 0.13 nVs and 1 keVee in S2 (at the surface) corresponds to 27.1 nVs.

Also shown in Fig. 3.9 is a combined energy estimator, known as E^*, which is a linear combination of S1 and S2. It exploits the anti-correlation introduced by the recombination of ionisation charge which can contribute to either S1 or S2 but not both (see section 2.1). The theoretical variance of $E^* = k \cdot \mathrm{S1} + (1-k)\mathrm{S2}$, at fixed true energy, is:

$$\sigma_*^2 = k^2\sigma_1^2 + (1-k)^2\,\sigma_2^2 + 2k\,(1-k)\,\sigma_1\sigma_2\rho\,, \tag{3.1}$$

where σ_1^2 and σ_2^2 are the variances of S1 and S2, and ρ is the correlation coefficient:

$$\rho \equiv \frac{\mathrm{cov}\,[\mathrm{S1},\mathrm{S2}]}{\sigma_1\sigma_2}. \tag{3.2}$$

σ_*^2 has a minimum of

$$\frac{\sigma_1^2\sigma_2^2\left(1-\rho^2\right)}{\sigma_1^2+\sigma_2^2-2\sigma_1\sigma_2\rho} \tag{3.3}$$

when

$$k = \frac{\sigma_2^2-\sigma_1\sigma_2\rho}{\sigma_1^2+\sigma_2^2-2\sigma_1\sigma_2\rho}. \tag{3.4}$$

The improvement in resolution at the ^{57}Co photoabsorption peak is better than would be expected for statistically independent energy measurements, and the optimum k is found

to be 0.25. An energy resolution of $\sigma_*/E^* = 8.1\%$ at 122 keVee was obtained during the SSR (as illustrated in Fig. 3.9); the equivalent value for FSR data re-analysed with the same software is 4.5%, confirming the relatively poor performance of the SSR array.

The ZEPLIN-III WIMP search box is a region in S1-S2 parameter space, and so adopting E^* as an energy scale instead of S1 would not change the number of counts within the box. However, in a discovery scenario the parent energy distribution of signal events, and hence WIMP and halo properties, would be more precisely reconstructed by using a combined energy variable with minimal variance for nuclear recoils, as advocated in Ref. [85].

3.3.2 Corrections derived from daily calibration

The PMT response functions and overall energy scales within Mercury are tuned to reconstruct S1 and S2 for 122 keV electron recoils in the reference ^{57}Co dataset. Variation in the response to S1 during the run is corrected, on a channel-by-channel basis, via the gain correction described in section 3.2. There are, in addition, sources of day-to-day variation in S2 response, i.e. the amount of electroluminescence light generated per unit ionisation drifted away from a recoil track. The most significant effects, all corrected for via daily calibration, are:

- tilting of the detector — with the liquid surface not exactly parallel to the anode, the path length through the gas varies in x-y plane;
- variation in vapour pressure, caused by temperature changes — this affects the electroluminescence yield;
- electrons attaching to impurities in the liquid — this results in a finite 'electron lifetime' and an exponential decrease, with drift time, of the fraction of ionisation that reaches the surface.

The lifetime of drifting electrons in the liquid is measured via fits to the ^{57}Co data. The known depth-dependence of the energy peak, caused by the energy-dependent attenuation length, is obtained from simulation as in Fig. 3.8; it is used to normalise the peak in observed electroluminescence to the mean electron recoil energy for different slices of depth. The variation with depth of the peak in electroluminescence per unit energy is fitted to find an exponential decay constant, which is the reciprocal of the electron lifetime in the liquid, as shown in Fig. 3.10. There is potential for systematic error due, for example, to any PMT saturation being less severe for deep events, but the statistical precision of the decay constant measurement is $\lesssim 7\%$ each day.

A novel method allowing a complementary lifetime estimate — physically independent, although with larger statistical error — was developed within the collaboration based on the survival probability of individual electrons on their way to the surface. These electrons are produced by the photoionisation of impurities by the VUV photons. The electroluminescence signal of a single extracted electron can be clustered and counted

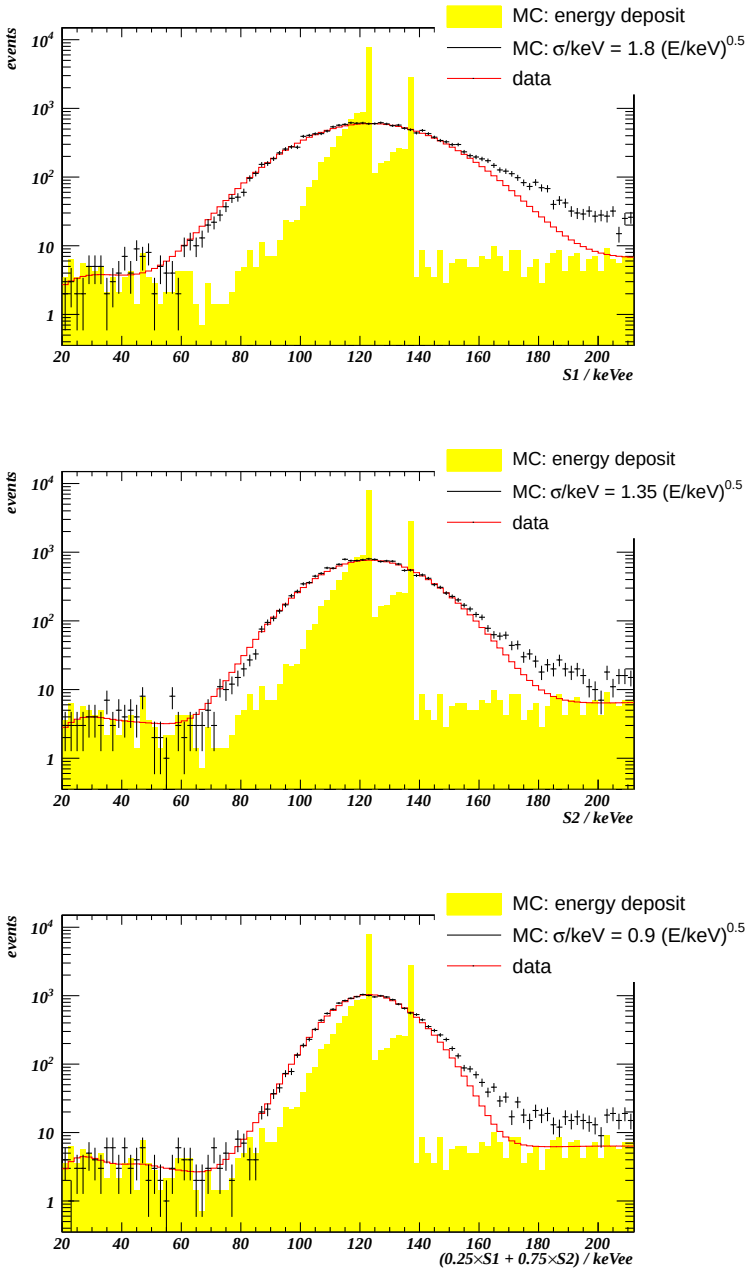

Figure 3.9: Simulated and measured spectra in S1, S2 and E^* for the reference ^{57}Co dataset. A Gaussian resolution model gives good agreement around the peak. There is an excess above simulation on the high-energy side, particularly in S1, corresponding to non-Gaussian tails in detector response and a small contribution from un-subtracted background: the background rate rises from below 1 dru at 20 keVee to 3 dru (corresponding to 1.5 events per bin in the above histograms) at 200 keVee.

with good efficiency in ZEPLIN-III. When they are generated by the light of primary scintillation, the time delay to the single electron pulse corresponds to a drift time, and hence the depth of the photoionisation vertex in the liquid. In ZEPLIN-III this production mechanism happens uniformly through the active volume, due in part to the long mean free path of VUV photons in liquid xenon [106]. The observed depth distribution of single electrons reaching the surface thus encodes their survival probability. Note that the uncertainty in the lifetime measured via photoionisation increases greatly for long lifetimes, because the rate of photoionisation itself drops as impurities leave the liquid.

The single electron signal will be overwhelmed by the much larger S2 signal when the latter is present, so S1-only events — with the recoil vertex in a dead region of the detector — are selected for this measurement. The results (Fig. 3.11) are moderately consistent with those obtained from ^{57}Co, suggesting that systematic bias in the daily measurement is not large [107]. As in the FSR, there was general improvement in electron lifetime while field was applied to the target and sharp decreases coinciding with power cuts and HV trips: the negative ions produced by the process of attachment are swept away from the fiducial volume and so the concentration of electronegative impurities decreases over time.

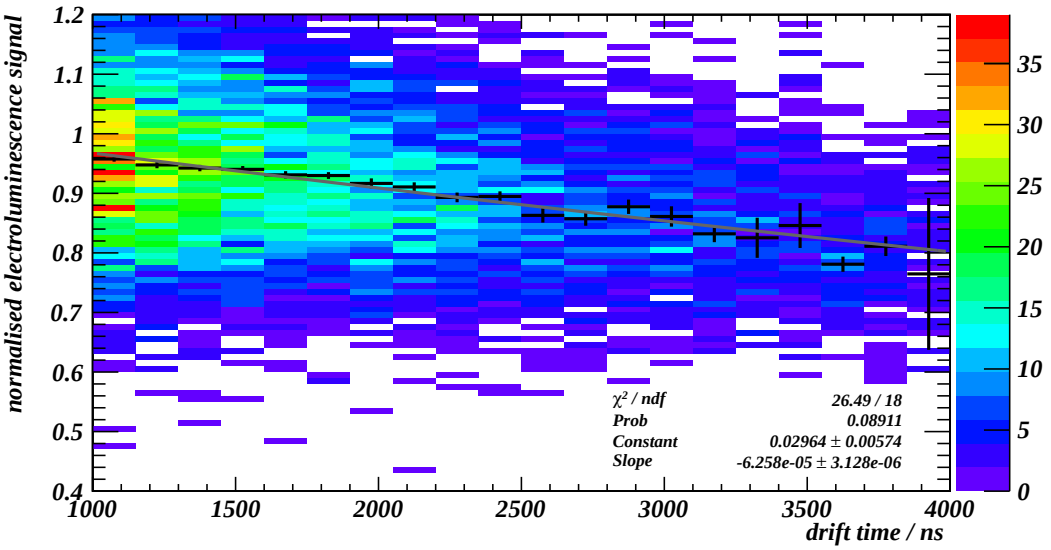

Figure 3.10: Electron lifetime measurement from daily ^{57}Co calibration. The normalised electroluminescence signal size is the uncorrected value of S2, divided by the mean recoil energy at the depth of the event from simulation. The exponential fit shown measures the decay constant as (63 ± 3) ms^{-1}, giving a lifetime of (16.0 ± 0.8) µs.

3.4 Tests of response uniformity

Automated ^{57}Co calibration provides large samples of events under very consistent conditions but they are, in some respects, quite different to the WIMP signal. Photoabsorption

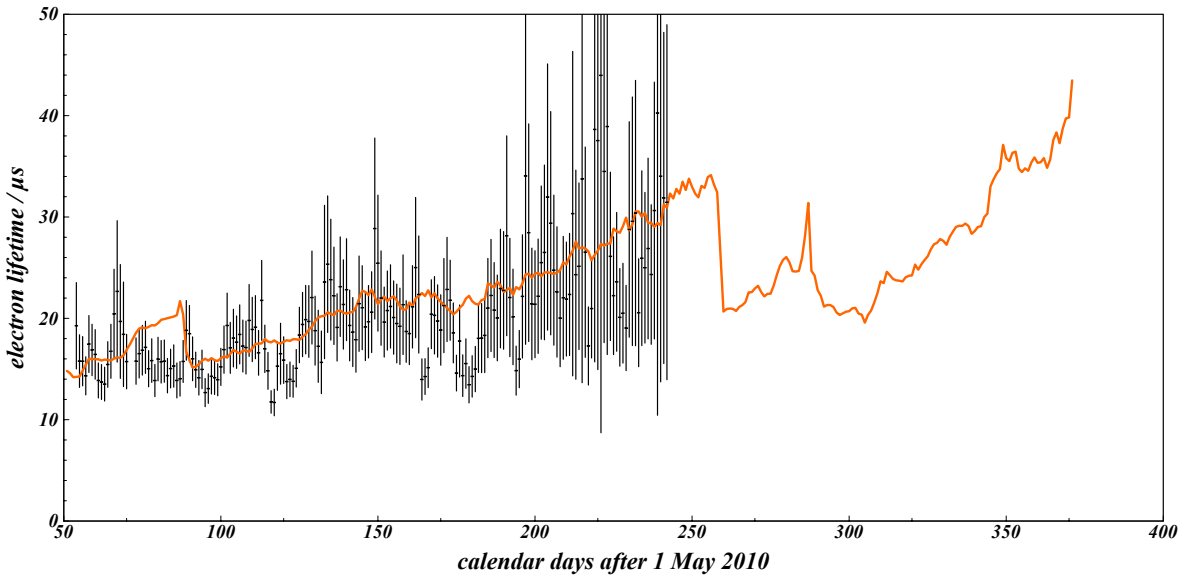

Figure 3.11: Electron lifetime during the SSR, measured via ^{57}Co (line) and single electrons (points with error bars). For reference, the fiducial volume extends to 13 µs and the cathode lies at 15.4 µs. The level of agreement between the two methods was even better in the FSR (not shown), where the statistical uncertainty in the single-electron method was smaller due to the higher background rate of S1-only events [107].

vertices are mostly close to the liquid surface and both interaction energies and mean photocurrent during acquisition are significantly higher than in the signal region of search data. As the reconstruction software, energy scales and day-by-day corrections were developed, their performance was also tested via specific background populations in search data. Events generated by activation of the xenon following neutron calibration, by radioisotopes originating in the cathode grid material and by background γ-rays in the bulk were all investigated to test the response uniformity of the chamber.

3.4.1 Events due to xenon activation

Metastable xenon nuclei generated during neutron calibration disperse evenly throughout the volume and are therefore valuable tests of event reconstruction. The isotopes ^{129m}Xe and ^{131m}Xe are produced by both neutron capture and inelastic scattering of fast neutrons. They decay with half-lives of 8.9 and 11.8 days, emitting γ-rays of 236 and 164 keV, respectively [108]. For the WIMP search dataset, only events with Golden-code S1 below 70 keVee are written to golden ntuples and analysed further. However, for 30 search days in August 2010 events up to a golden-code E^* of 500 keVee were processed as a dedicated high-energy dataset, used to study background radioactivity events away from the search region. These data were acquired either side of a day of neutron and γ-ray calibration on the 25th.

Activation lines were characterised from the days following the August neutron calibration, which consisted of 6.0 hours exposure to a 1321 ± 14 n/s AmBe source. A longer neutron run would have improved precision in nuclear recoil and activation analyses but,

as ever, minimising the background in the search region is more important for experimental sensitivity. Only γ-ray calibration data was acquired in the 16 hours immediately afterwards, and the background below 20 keVee increased only slightly in the remaining days of August: from 49.1 ± 1.5 to 55.1 ± 3.1 golden events per live day, before and after calibration.

The known γ-ray energies are recovered quite well in the combined energy variable E^* after background subtraction, as shown in Fig. 3.12. This supports the conclusions that:

- the mean photocurrent drawn during SSR ^{57}Co calibration does not significantly bias the response of the PMTs compared with 'dark' conditions;
- the analysis software maintains good linearity even when the most-illuminated PMT is ignored due to DAQ scale saturation, as is the case for around one third of S2s in the 236 keV peak.

The discrepancies in reconstructed energy are slightly larger than would be expected from the statistics of the fit, possibly due to the simplified resolution model or small systematic uncertainties arising from time-dependent corrections. Figure 3.13 shows the reconstructed activation lines as a function of position in the fiducial volume, fitting to events in a range of depth or radius. No position-dependent bias is evident, but there is some loss of resolution in the last 10 mm of radius due to worsening light collection.

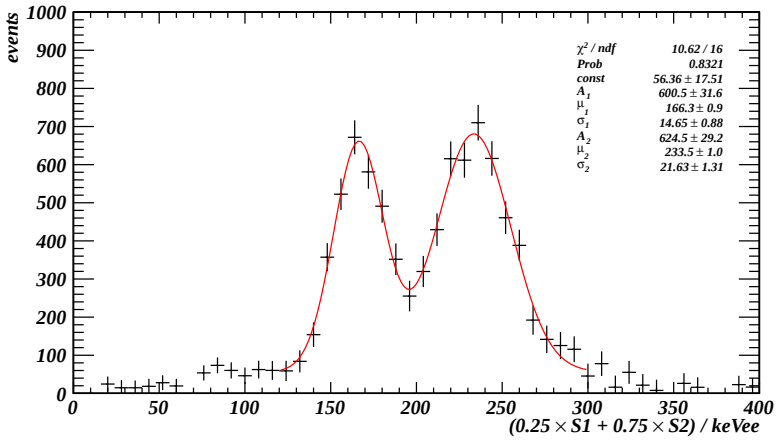

Figure 3.12: Distribution in E^* of activation events in the 6 days following neutron calibration (5.7 live days of acquisition), obtained by subtracting the background measured in the preceding 24 days (22.8 live days). The fitted spectrum shown is the sum of two Gaussians, in fair agreement with the known γ-ray energies of 164 and 236 keV from meta-stable xenon nuclei.

3.4.2 Events due to radon progeny on the cathode

During the period of xenon purification before the FSR, the gas was contaminated with a small amount of ^{222}Rn emanated from the getter (model SAES PS11-MC500) at the

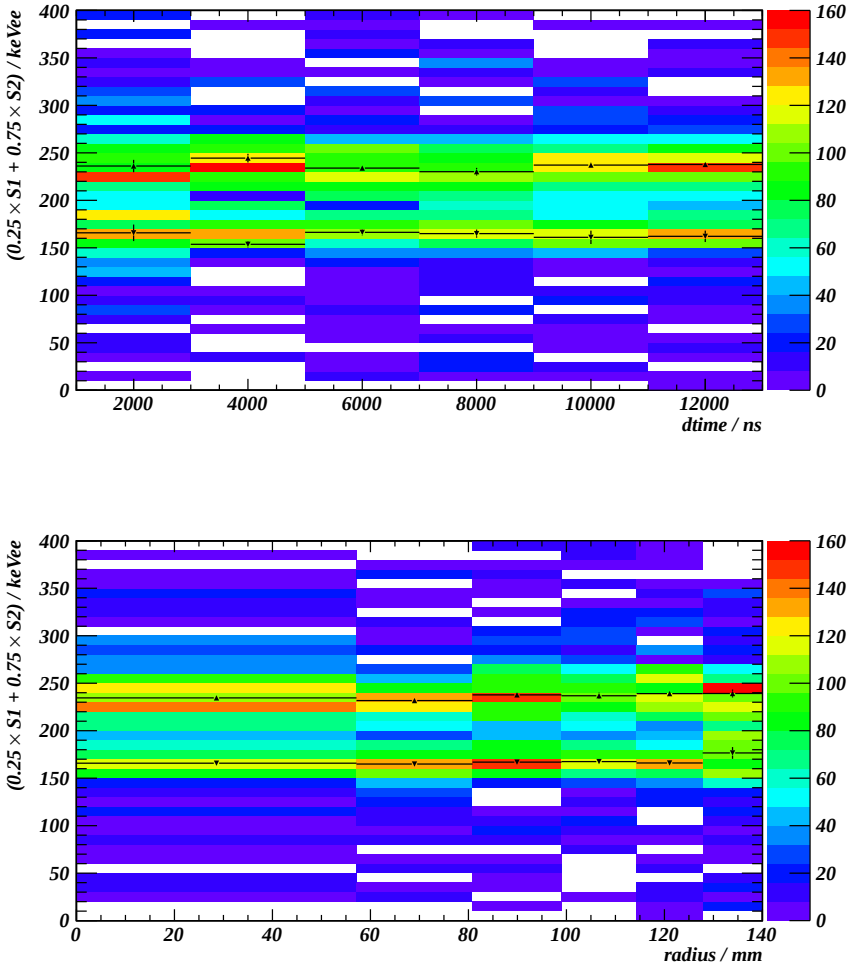

Figure 3.13: Distribution in E^*, at varying depth (top) and radius (bottom), of events in the 6 days following neutron calibration, after subtracting the background from the preceding 24 days. Superimposed are fitted peak positions in the energy spectrum, as in Fig. 3.12, for each slice along the abscissa.

rate of 0.6 atoms per second. The origin of the contamination was found via the time-dependent rate of α events in ZEPLIN-II: a step increase was observed when a getter of the same model was introduced to the recirculation path, corresponding to radon building up in the getter while isolated and then being flushed into the xenon; after removal of the getter, the α event rate dropped with the half-life expected for ^{222}Rn [109]. The decay chain is as follows, with relevant Q values, decay modes and half-lives in brackets:

$$\begin{aligned} ^{222}\mathrm{Rn}\,(\alpha, 3.8\ \mathrm{d}) \rightarrow & ^{218}\mathrm{Po}\,(\alpha, 3.1\ \mathrm{min}) \rightarrow ^{214}\ \mathrm{Pb}\,(\beta^-, 27\ \mathrm{min}) \rightarrow ^{214}\ \mathrm{Bi}\,(\beta^-, 20\ \mathrm{min}) \\ & \rightarrow ^{214}\mathrm{Po}\,(\alpha, 164\ \mu\mathrm{s}) \rightarrow ^{210}\mathrm{Pb}\,(63\ \mathrm{keV}\ \beta^-, 22\ \mathrm{y}) \rightarrow ^{210}\ \mathrm{Bi}\,(1.2\ \mathrm{MeV}\ \beta^-\ , 5\mathrm{d}) \\ & \rightarrow ^{210}\mathrm{Po}\,(5.4\ \mathrm{MeV}\ \alpha, 138\ \mathrm{d}) \rightarrow ^{206}\ \mathrm{Pb}. \end{aligned}$$

The positive ions of ^{214}Bi and ^{214}Po produced by β^--decay drift to the cathode and decay, leaving ^{210}Pb implanted in the wire grid. A more radiopure getter (model SAES PS4-MT3) was used in preparation for the second run, and so there has been little new deposition by this mechanism since May 2008. The decay rates of different species in an initially pure sample of ^{210}Pb are shown in Fig. 3.14; by the time of the SSR any ^{210}Pb deposited during the FSR has decayed sufficiently to reach transient equilibrium, where the ratio of activities from the three decay steps is constant. The background from these decays is thus expected to have an energy distribution which is constant in time, with gradually decreasing rate throughout the SSR. Additional backgrounds from intrinsic uranium and thorium contamination in the wire material also contribute in principle, but affect mainly higher-energy populations than those analysed here.

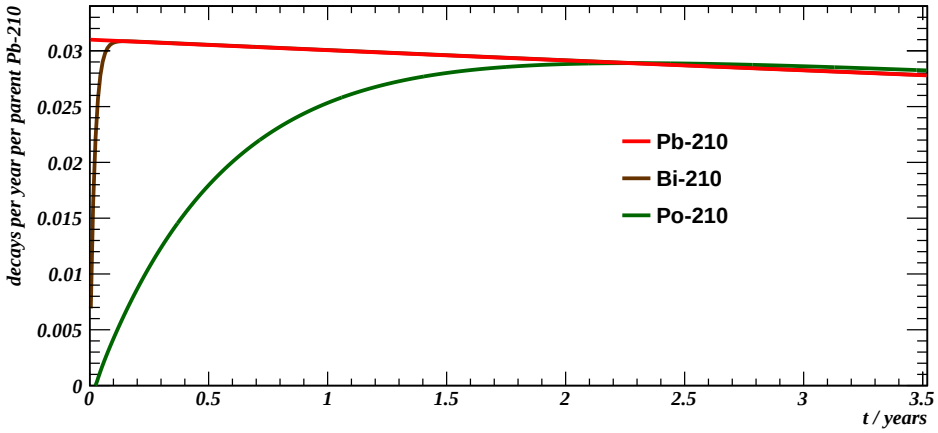

Figure 3.14: Decay rates in the ^{210}Pb sub-chain.

The fiducial volume extends from 5.7 mm above the cathode, beyond the range of the detectable decay products in the ^{210}Pb sub-chain. However, they do deposit energy in the intervening region, from where ionisation charge is still extracted, giving rise to several classes of events in non-fiducial background data. α particles from ^{210}Po decay emitted

into the active volume give S1 and S2 signals well above the range of interest for WIMP searches and these are observed at a rate of around 20 per day; if, instead, the daughter nucleus is emitted into the active volume then a relatively low-energy nuclear-recoil event is observed. Low-energy electron recoils close to the cathode wire are generated mainly by the decay of ^{210}Pb, which has two paths: a β spectrum up to $E_{\max}$ = 63.1 keV (BR=14%); and a β spectrum up to $E_{\max}$ = 16.6 keV accompanied by a 46.5 keV γ-ray with attenuation length of around 200 µm.

The detector response to energy deposits close to the wire grid differs systematically from that in the fiducial volume. Recombination is suppressed by the high electric field within ~100 µm of the wires (diameter 100 µm): more ionisation charge is drifted away from the track and the scintillation yield drops even further than for fiducial events. The response in photoelectrons per unit scintillation light is also reduced, due to obscuration by the grid itself. Reconstruction of S2, interpreted as an amount of drifted charge, should not be significantly biased but it must be noted that the unit of 122 keVee is no longer the expected S2 for 122 keV electron recoils since this is field-dependent. Inefficiency due to an S2 being missed by the hardware trigger — which would imply a purity-dependent spectrum of recorded events — may be neglected: no cathode events which satisfy the 3-PMT coincidence requirement in S1 have a raw, summed S2 pulse below the 5.0 nVs size needed for 100% efficiency. Although the ionisation spectrum of low-energy cathode events is uncertain, depending critically on implantation depth and decay angle, it is expected to be constant over time and so it provides a test of the electron lifetime correction at the depth where it is largest. Because of the varying purity, broadly improving during the SSR, the correction factor applied to cathode events at 15.4 µs drift time ranges from 1.5 to 2.8, with mean of 2.0 and standard deviation of 0.3.

In a discrimination plot of $\log_{10}(\mathrm{S2}/\mathrm{S1})$ versus S1 (Fig. 3.15) the nuclear and electron recoil populations show up clearly but, as discussed, displaced from where the same energy deposits in the fiducial volume would have appeared: the ionisation-to-scintillation ratio is systematically higher where a strong field prevents recombination and light collection is poor. The S1-dependence of the population means was approximately scaled away by using a new discrimination variable, $\log_{10}(\mathrm{S2}/\mathrm{S1}^{0.23})$, as shown in Fig. 3.16. Figure 3.17 shows the stability of the two resulting peaks during the run. The variation of each peak position individually is consistent with fluctuations in the fits. The sample correlation coefficient between the two peak positions, 0.47, could indicate an uncorrected systematic variation, but is also reasonably consistent with underlying statistical independence (p-value of 0.051 for a two-tailed Student test of H_0: $\rho = 0$).

3.4.3 Scintillation-to-ionisation ratio of background events

A final investigation of response uniformity focused on low-energy electron recoil background events in the fiducial volume. A blinded search dataset, detailed in section 4.4.2, was used: near-final cuts were applied, but the small fraction of events (later revealed

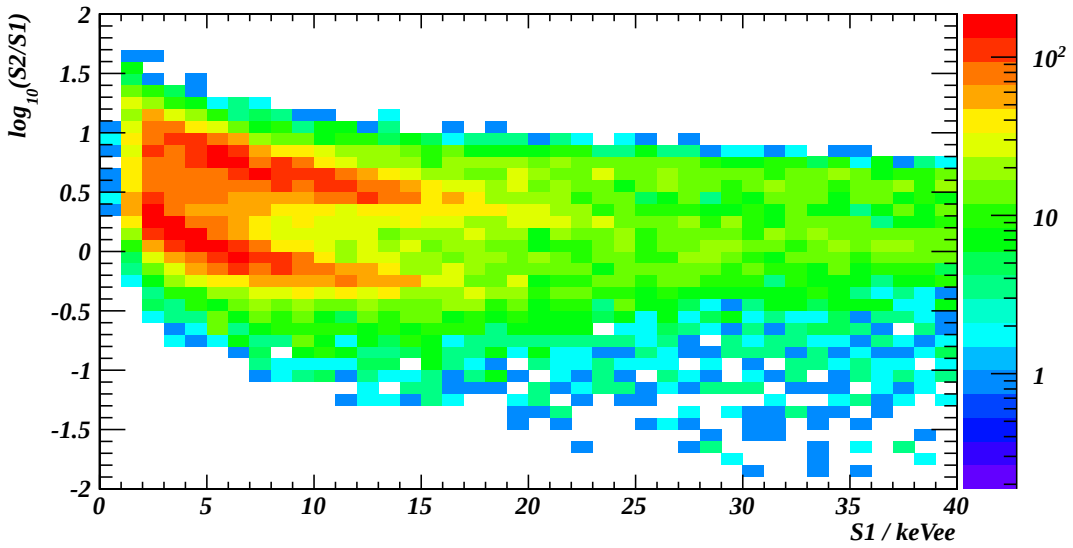

Figure 3.15: Distribution of search events with drift time >15 µs, in S1 and S2. These variables are not accurate estimates of recoil energy close to the cathode grid but they are still, at a fixed position, proportional to the amounts of scintillation and ionisation. The higher- and lower-S2 populations at S1 $\lesssim$ 15 keVee are respectively electron and nuclear recoil events due to decays on the cathode surface, mostly attributed to ^{210}Pb progeny.

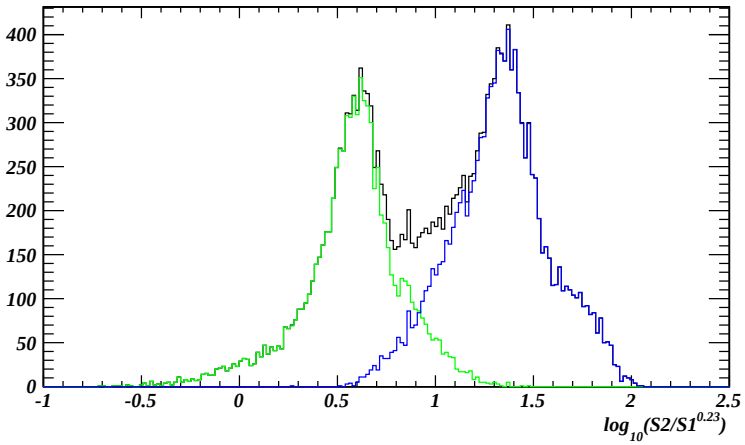

Figure 3.16: Distributions of the events in Fig. 3.15 in an S1-scaled ionisation variable: black, for all events; green and blue for events with the peak PMT for S2 measured in high- and low-sensitivity DAQ channels, respectively.

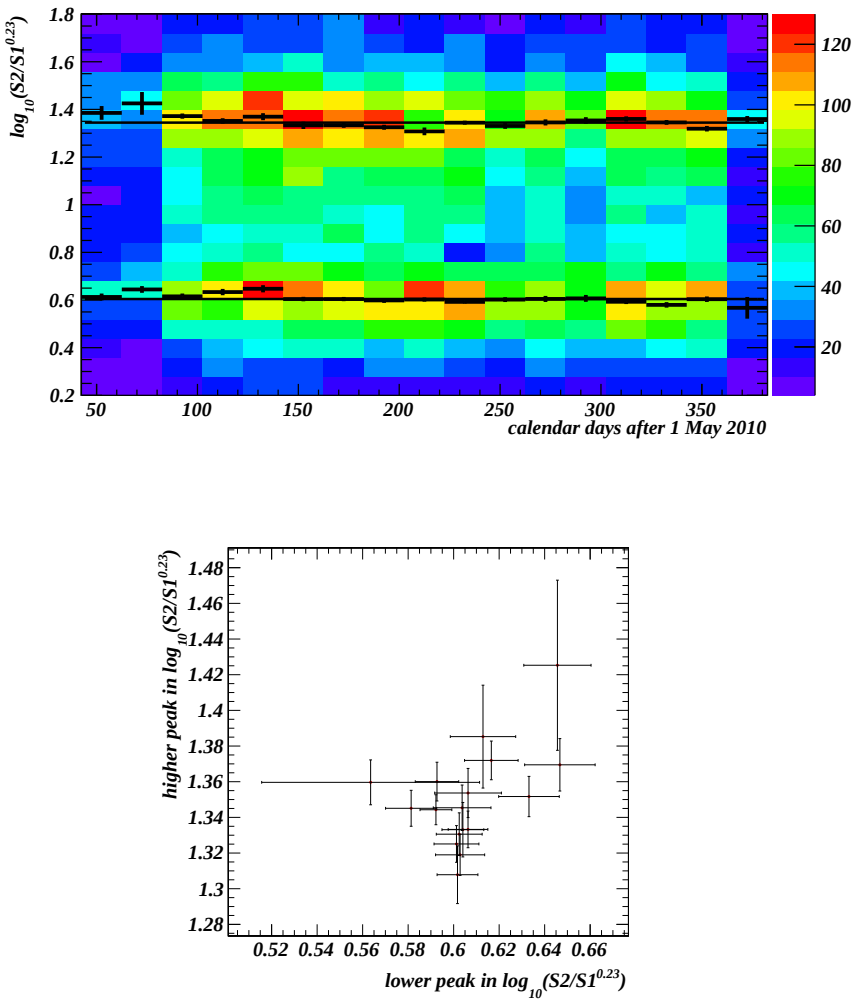

Figure 3.17: Top: time variation of the distribution in Fig. 3.16 during the SSR; fitted peak positions within each slice along the abcissa (points) and over all data (lines) are superimposed. The varying number of counts per bin along the x-axis reflects the distribution of detector live-time — corrected for this, the rate of cathode events versus time is consistent with a straight line. Raw counts are plotted here so that the Poisson uncertainty in each bin is evident. Bottom: the 20-day time intervals as points in the plane of ER-peak-position versus NR-peak-position, showing only a weak correlation between the two ($\rho = 0.47 \pm 0.23$).

to be 0.3% for $2 < \mathrm{S1/keVee} < 20$) which are un-vetoed and lie below the median (S2/S1) for nuclear recoils were omitted. Using an approach similar to the analysis of cathode events, the mean value of $\log_{10}(\mathrm{S2/S1})$ at given S1 (in this case fitted by a polynomial) was subtracted to give a discrimination variable which is flat in S1, as shown in Fig. 3.18. The position of the peak in this $\Delta\log_{10}(\mathrm{S2/S1})$ variable was reconstructed consistently throughout the fiducial range of depth and radius with statistical precision $\lesssim 0.03$, shown in Fig. 3.19.

In conclusion, variation during the run and throughout the fiducial volume was minimised by the campaign of calibration and correction described in the preceding sections. The achieved energy resolution, measured at 122 keV, is worse than that of the FSR detector configuration and this is attributed to the poorer optical performance of the PMTs. Nonetheless we confirmed — using, where possible, low-rate and low-energy data — that the mean response of the chamber to ionisation and scintillation exhibits good uniformity in the range of $(\mathrm{S1}, \mathrm{S2}, x, y, z)$ where the WIMP signal is sought.

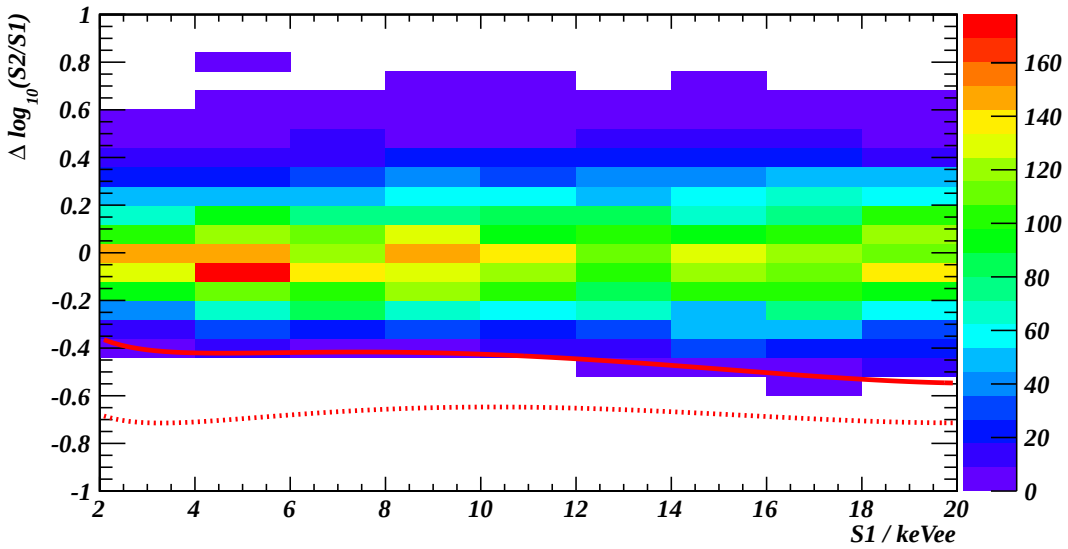

Figure 3.18: Distribution of the ~7000 SSR search events in the range $2 < \mathrm{S1/keVee} < 20$, using a flattened discrimination variable: $\log_{10}(\mathrm{S2/S1})$ minus its population mean at given S1, versus S1. The blinded region, extending 2-σ downwards from the mean of the nuclear recoil population (section 4.4.2), is indicated by the red lines.

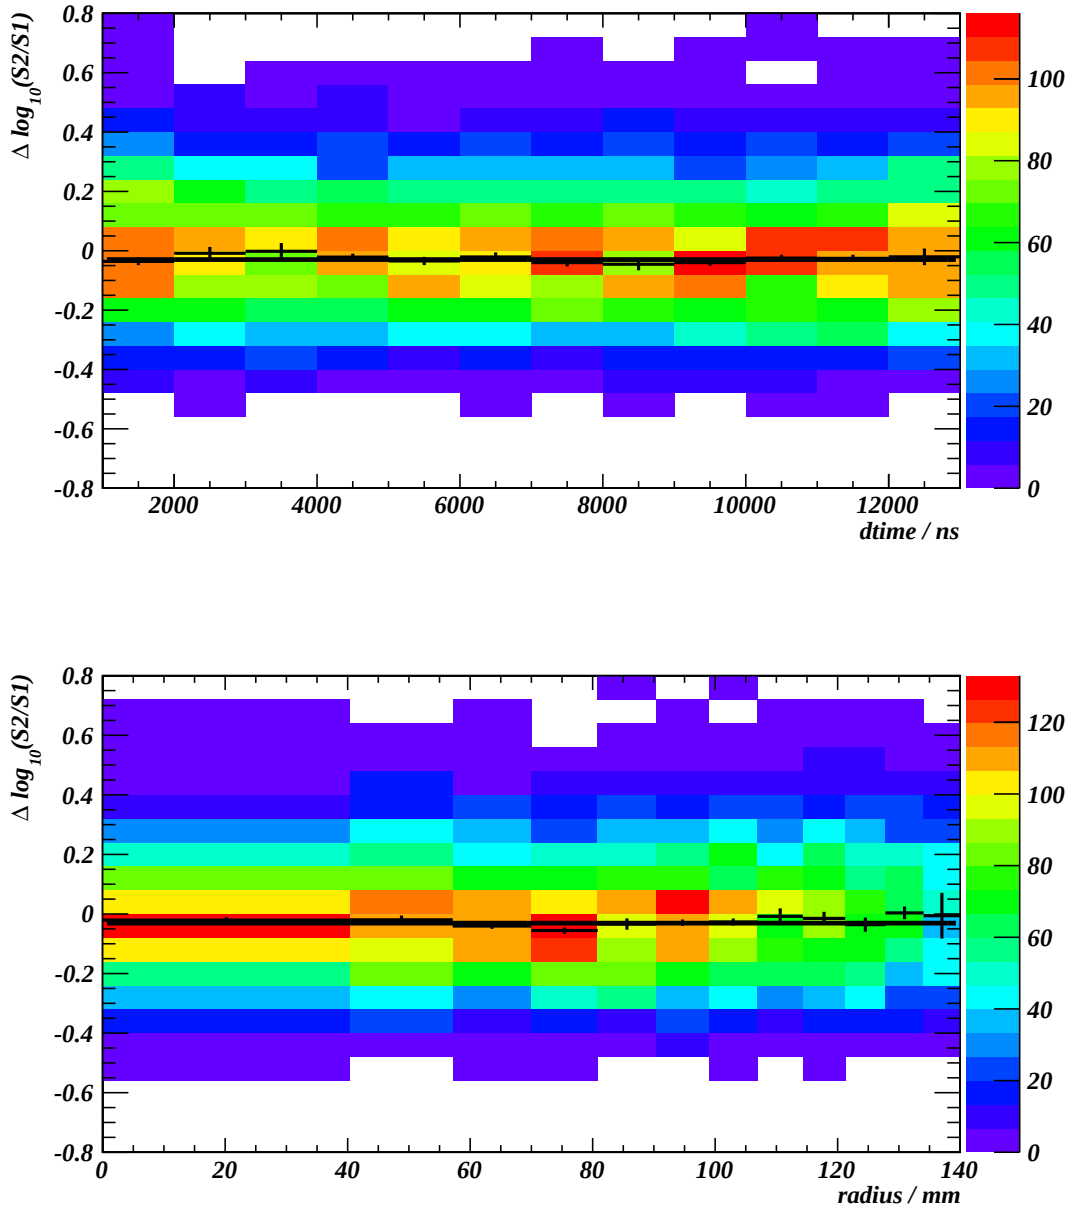

Figure 3.19: Distribution in $\Delta \log_{10}(\mathrm{S2}/\mathrm{S1})$, at varying depth (top) and radius (bottom), of search events in the range 2 < S1/keVee < 20. Superimposed are fitted peak positions for each slice along the abscissa. The drop in density at large radius is partly a result of the radioactivity-induced background decreasing at the edge of the array, away from the PMT materials [98]; in addition, quality cuts on the reconstruction of S1 are failed more frequently in regions of relatively low light yield.

3.5 Scintillation and ionisation yields for nuclear recoil

The last step in calibrating the response of ZEPLIN-III to nuclear recoil events is to relate recoil energy to the quantities measured in the detector: scintillation and ionisation in units of the ^{57}Co calibration point. This scintillation efficiency at recoil energy E, $\mathcal{L}_{\rm eff}(E)$, is defined as the number of scintillation photons produced per unit energy for nuclear recoils divided by that for 122 keV electron recoils, all at zero applied field. It is related to the electron-equivalent energy of the scintillation signal, S1, via:

$$\langle \mathrm{S1} \rangle = \frac{S_{\rm nr}}{S_{\rm ee}} \mathcal{L}_{\rm eff}(E)\, E, \tag{3.5}$$

where $S_{\rm nr}$ and $S_{\rm ee}$ are the field-induced suppression factors (section 2.1.1) for nuclear and electron recoils. The charge yield $\mathcal{Q}_{\rm y}$ is the number of electrons extracted from the track per unit nuclear recoil energy. It is related to the electron-equivalent energy of the ionisation signal, S2, via:

$$\langle \mathrm{S2} \rangle = W_{\rm e} \frac{q(|\mathbf{E}|)}{q_\infty} \mathcal{Q}_{\rm y}(E)\, E, \tag{3.6}$$

where $q(|\mathbf{E}|)/q_\infty$ is the fraction of ionisation charge extracted from nuclear recoil tracks at finite field and $W_{\rm e}$ is the mean electron-recoil energy required to generate an electron-ion pair (which may then recombine or drift).

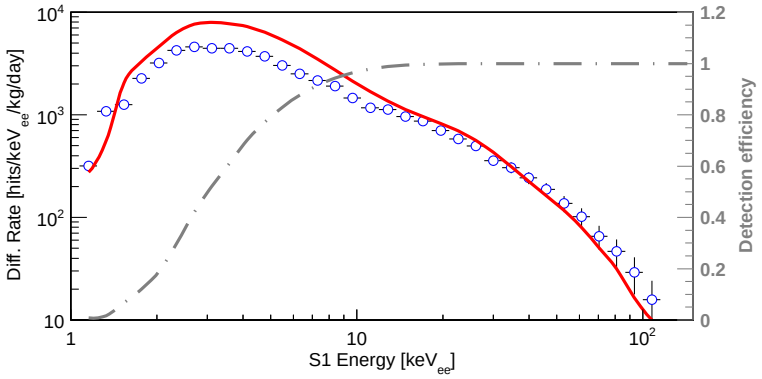

Figure 3.20: Simulated (solid line) and observed (points) nuclear recoil spectra in S1 during AmBe calibration, under the (incorrect) assumption of constant $\mathcal{L}_{\rm eff} = 0.19$. The very poor agreement points to the non-linear relation between nuclear recoil energy and S1.

The behaviour of $\mathcal{L}_{\rm eff}$ in the energy range of the dark matter signal has only been revealed in the past few years: a spectrum-matching measurement made in the course of the ZEPLIN-III FSR [48] was among the first observations of a declining scintillation efficiency below 20 keV. The non-linearity of the response is clearly evident in Fig. 3.20, a comparison of (SSR) neutron calibration data with a simulation that assumes constant scintillation efficiency. The nuclear recoil energy threshold of xenon detectors is directly determined by $\mathcal{L}_{\rm eff}$, and is of particular importance when considering putative light-WIMP

signals seen in other targets [43, 110]. The implications for the sensitivity of current and future xenon experiments are great, and were the subject of intense discussion in the community [111, 112]. More data were clearly required. Measurements of the yield functions in liquid xenon fall into two broad categories. For beam measurements, typically made with small dedicated detectors, a monoenergetic neutron source is used and events of known recoil energies are selected kinematically via the angle of the outgoing neutron. For spectrum-matching measurements, such as those described here, the recoil energies produced by a broad-spectrum source are simulated and matched to the observed spectrum via a fitted yield.

The WIMP signal in ZEPLIN-III is modeled as a spectrum in S1 with a distribution in $\log_{10}(\mathrm{S2}/\mathrm{S1})$ from nuclear recoil calibration. This distribution, at given S1, may be obtained from AmBe data without knowing the underlying recoil energies which generate the signals. One could in principle reverse the roles of S1 and S2 — calculating a spectrum in S2 via $\mathcal{Q}_\mathrm{y}$ and taking $\log_{10}(\mathrm{S2}/\mathrm{S1})$ at given S2 from AmBe data — but this is a less natural basis since the dominant inefficiencies are S1-dependent. We therefore focus on the measurement of scintillation efficiency but report the result of a similar measurement of $\mathcal{Q}_\mathrm{y}$. Both are published in Ref. [77].

3.5.1 Nuclear recoil calibration data

Calibration data were acquired with a 20 MBq AmBe source positioned inside the shielding and on top of the vacuum vessel, within 5 cm of the central axis of the PMT array. The rate of neutron emission from the source was measured as (1321±14) s^{-1} (National Physical Laboratory, May 2009); the ZEPLIN-III hardware trigger recorded a count rate of ~ 200 s^{-1}. A total of 10.2 hours of exposure was recorded in three acquisitions — one in each of June, July, and August 2010 — and the data were combined after confirming that the distributions of events in S1 and S2 were statistically consistent. A shallower fiducial volume of mass 3.5 kg was used, smaller than for WIMP search, to guard against potential hardware trigger inefficiency from small ionisation signals with long drift times.

The yield functions are measured by comparing measured and simulated spectra of nuclear recoils from single elastic scatters; however, event acquisition is also triggered by multiple scatters in the active volume, inelastic scattering accompanied by de-excitation γ-rays, and electron recoils due to γ-rays from the AmBe source. These classes of events are removed from the dataset by the Golden code, mild cuts on Mercury reconstruction quality, and selection in the S2-S1 plane, similar to the nuclear recoil region seen in Fig. 4.6 on page 98. The resulting distribution of single-scatter nuclear recoils in S1 is shown in Fig. 3.21.

3.5.2 Simulation, resolution and efficiency

Neutron calibration was simulated for a model of the entire ZEPLIN-III experiment [113] with the Geant4.9.3 toolkit [114]. Updated Xe(n,n) scattering cross sections from

ENDF/B-VII database [115] were used in place of the Geant4.9.3 default (from ENDF/B-VI). The neutron spectrum of the source (shown in Appendix A.3) extends to a maximum of 11 MeV, with a mean of 4.1 MeV. It was simulated from 100 keV up; this cutoff corresponds to a maximum nuclear recoil energy of 3 keV, well below the threshold of the detector.

A spectrum in the observable S1 corresponding to the simulated spectrum of recoils in E was calculated, for given $\mathcal{L}_{\text{eff}}$, by accounting for resolution and efficiency. (A similar process is entailed in modeling the WIMP signal, as will be outlined in section 4.2.) First, the test $\mathcal{L}_{\text{eff}}$ is used to construct a pre-resolution spectrum in S1 via equation 3.5. Resolution in S1 is modeled as a Poisson-like distribution, with the factorial replaced by a continuous Γ function, which approximates variations in the number of photoelectrons and their areas.

The dominant efficiency is the probability for three-fold PMT coincidence. During reconstruction the LRFs, which encode PMT response, are used to estimate an unknown vertex by maximising, over x, y and S1, the likelihood of the observation. This calculation was inverted for the efficiency calculation by drawing from the LRFs as probability distributions, at simulated points $(\text{S1}, x, y, z)$. The fraction of calibration events expected to pass the three-fold requirement was thus found as a function of S1.

The $\mathcal{L}_{\text{eff}}$ curve obtained by spectrum matching relies, above all, on the accuracy of the S1-dependent efficiency applied to simulation. The efficiency obtained from Mercury LRFs was cross-checked by another method which, although more crude, makes more direct use of AmBe data. In the rare event search, 3-PMT coincidence in the primary is required because a small number of systematically mis-reconstructed events could greatly worsen sensitivity. However, the vast majority of S1-like signals with 2-fold coincidence *are* due to genuine scintillation; the ratio of 2- to 3-fold events in bins of S1 can be obtained from AmBe data. Monte Carlo events with light yield tuned to match this observation will reveal the ratio of 1-fold to 3-fold, which determines the (post-resolution) efficiency but is not directly accessible in data. Low energy primaries were simulated and binned by a simplified reconstructed S1 based on photoelectron-counting. The simulated scintillation yield, in photons per keVee, was tuned such that the ratios of 3-fold to 2-fold coincidence matched the data. The 3-fold to 1-fold ratio was then obtained from the simulated events, and agreed with the Mercury method to within $\lesssim 4\%$ in the range of S1 used to fit $\mathcal{L}_{\text{eff}}$.

3.5.3 Spectrum matching and results

The function $\mathcal{L}_{\text{eff}}(E)$ is modeled as a cubic spline with nodes at 0.5, 2.5, 6.0, 15, 50 and 200 keV. The values of $\mathcal{L}_{\text{eff}}$ at the nodes are then varied in a χ^2 fit of the simulated spectrum in S1 to the observation in the range $2.5 < \text{S1/keVee} < 100$. The best-fit values give the spectrum shown in Fig. 3.21. Statistical uncertainty is evaluated via the $\Delta\chi^2$ statistic: for samples sufficiently large to assume Gaussian errors, the difference between the best-fit χ^2 statistic and that given by the true model parameters follows the χ^2 distribution with

degrees of freedom (dof) equal to the number of parameters of interest [116]. The 68%-point in the CDF of the χ^2 distribution for 6 dof is Q_c =7.01. Once the best-fit $\mathcal{L}_{\text{eff}}$ has been found, splines satisfying $\chi^2 - \chi^2_{\text{min}} < Q_c$ define a 1-σ confidence interval in the space of the node values. A rough representation of this statistical uncertainty is shown by the band in Fig. 3.22, which is the area occupied by ~100 such acceptable splines obtained by scanning a 6-dimensional grid.

Neutron calibration data from the first science run were also processed with current software tools. This re-analysis revealed a previously un-noticed inefficiency in FSR event selection, caused by a consistency requirement between the high- and low-sensitivity channels which disproportionately cut small, but nevertheless good, S2 pulses. Resolution in S1 and S2 was improved by the current reconstruction algorithms. Systematic uncertainty in the activity of the source which had been used in 2008 was reduced by calibration against the SSR source in a high-purity germanium detector at Boulby. The revised FSR $\mathcal{L}_{\text{eff}}$ curve, included in Fig. 3.22, is consistent, within statistical errors, with that of the SSR. The fiducial range of the $\mathcal{L}_{\text{eff}}$ measurement begins at 7 and 9 keV, respectively for the FSR and SSR, due to the uncertainty in efficiency for low numbers of S1 photons.

The ITEP group developed a novel analysis which bypassed simulation and efficiency curves by taking others' measurements of $\mathcal{Q}_{\text{y}}$ as an input for measuring $\mathcal{L}_{\text{eff}}$. The mean scintillation signal is found for events in bins of S2, which are mapped to bins of E via equation 3.6 and an assumed $\mathcal{Q}_{\text{y}}$. All events with a single S2 are considered, whether or not S1 is identified by the Golden code; any pulses preceding S2 are then summed to give the mean amount of scintillation light without any threshold effect from S1. A background contribution, from thermal SPEs and single extracted electrons, is estimated from regions preceding an S2 by more than the maximum physical drift time. Without a Mercury reconstruction of S1, the amount of primary scintillation is estimated directly from summed, flat-fielded pulse areas; only events in the inner 120 mm radius are included, to mitigate the effect of varying light collection in this conversion.

The energy threshold of this dataset depends only on the hardware trigger and the raw pulse area of S2; it is therefore dependent on the depth and radius of the scattering vertex, but is everywhere greater than 0.5 keVee in S2. Where S1 is not identified, the drift time is unknown; an average lifetime correction, weighted by the depth distribution of low-energy nuclear recoils which *are* reconstructed as golden events, is applied instead.

The results of this alternative measurement are shown in Fig. 3.22 and agree with the spectrum-matching result for $\mathcal{L}_{\text{eff}}$. The uncertainty is somewhat larger, due to the absence of depth information for events without a clustered S1 and the systematic error on the assumed $\mathcal{Q}_{\text{y}}(E)$, among other effects.

A spectrum-matching analysis of the ionisation yield, $\mathcal{Q}_{\text{y}}$, was also carried out. The main technical difference from matching the $\mathcal{L}_{\text{eff}}$ spectrum arises in the efficiency as a function of S2, which is in fact determined by the expectation for S1: i.e. low-energy events missing from a particular bin of S2 were probably lost due to the S1 signal not being reconstructed. Fits in S2 were therefore made only above 1.0 keVee and 1.2 keVee,

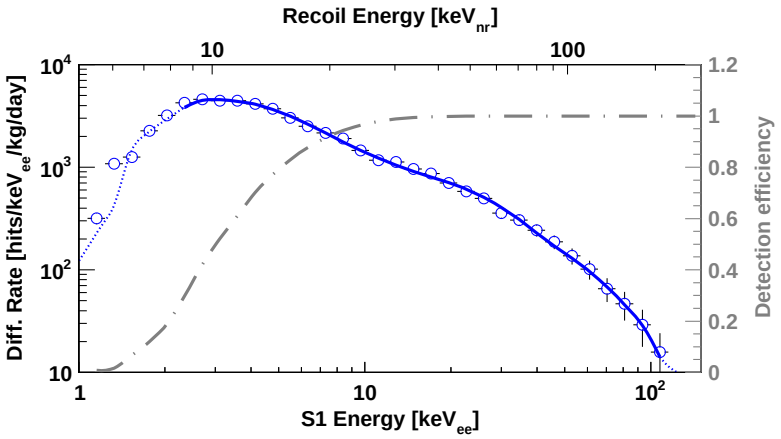

Figure 3.21: Simulated (solid line) and observed (points) nuclear recoil spectrum in S1 during AmBe calibration, together with the S1-dependent detection efficiency (dot-dashed line) from [77]. The recoil energy inferred from the best fit $\mathcal{L}_{\rm eff}$ is shown along the top axis.

respectively for the FSR and SSR, despite the hardware trigger in both cases being fully efficient well below 1 keVee in S2. Spline points for the function $\mathcal{Q}_{\rm y}(E)$ were set at 0.5, 4.0, 10, 30, 75 and 250 keV recoil energy. The best-fit spline and indicative confidence interval, as for $\mathcal{L}_{\rm eff}$, are shown in Fig. 3.23.

In conclusion, the SSR and re-analysed FSR $\mathcal{L}_{\rm eff}$ measurements, made with different sources and different PMTs, are mutually consistent. Furthermore, they are in good agreement with recent beam measurements [118, 117] which extend to 4 keV recoil energy, and report the same steady decrease with decreasing energy. A consensus has been established that two-phase xenon detectors with thresholds ~ 2 keVee in S1 have sensitivity to WIMP masses at least as low as 10 GeV c^{-2}. In the next chapter, we proceed to model exactly how the WIMP signal, low-mass or otherwise, would appear in ZEPLIN-III.

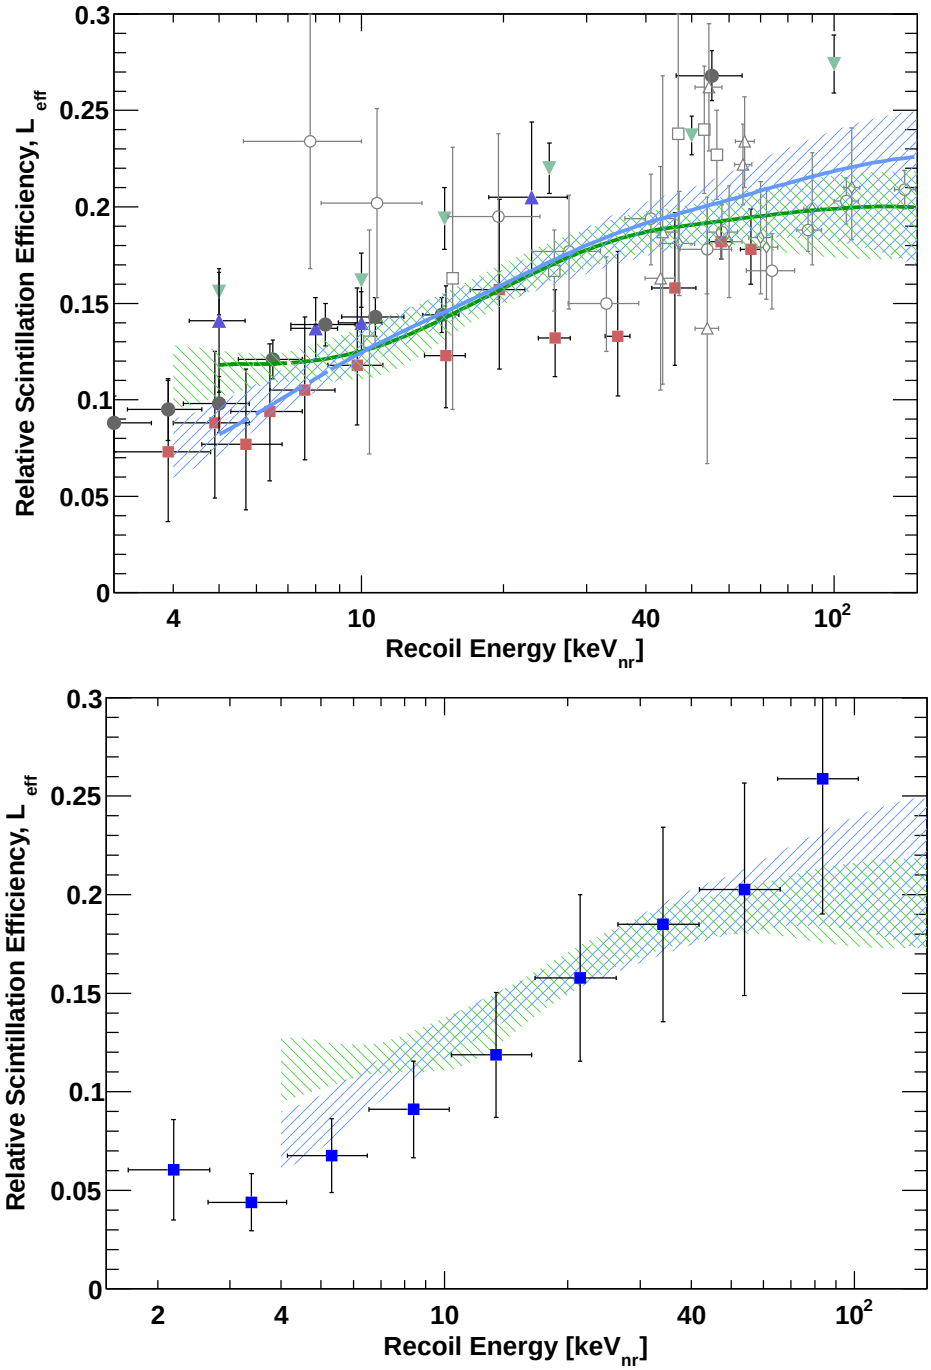

Figure 3.22: Measurements of the scintillation efficiency, $\mathcal{L}_{\text{eff}}$. Green and blue hatched areas in both plots correspond to the 68% indicative confidence belts for FSR and SSR spectrum-matching result. Top: direct measurements are also shown [(●)[117], (■)[118], (▲)[119], (○)[120], (△)[121], (▽)[122] and (◇)[123]] along with the spectrum-matching result from the XENON10 experiment (▼)[124]. Bottom: ZEPLIN-III simulation-independent results are also shown (□). From [77].

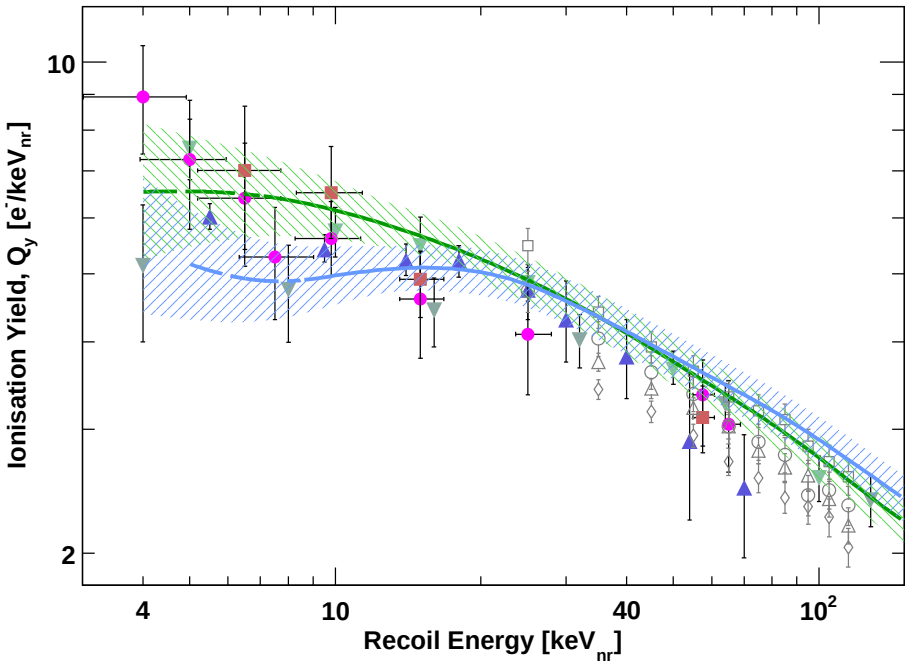

Figure 3.23: Measurements of the ionisation yield for nuclear recoils, $\mathcal{Q}_y$. Green and blue hatched areas are 68% indicative confidence belts for FSR and SSR spectrum-matching result. Also shown are previous measurements at 1.0 kV/cm (●) and 4.0 kV/cm (■) from Ref. [118], at 2 kV/cm (□), (△), 0.3 kV/cm (○) and 0.1 kV/cm (◇) from Ref. [123] along with spectrum matching results at 0.73 kV/cm (▲) [124] and (▼) [125]. From [77].

Chapter 4

Signal, background and event selection

The preceding chapter described (i) calibration of hardware and analysis software to reduce bias and variance when reconstructing scintillation and ionisation signals in the liquid, and (ii) calibration of yields to relate the energy of nuclear recoils to observed signal sizes. This work is now applied to build a signal model for WIMP-nucleus scattering, as it would appear in ZEPLIN-III search data, starting from basic properties of the dark matter particles and the local halo. Characteristics of background events in ZEPLIN-III are also described, including information provided by the veto detector which was a new addition for the second run. Finally, analysis cuts aiming to maximise sensitivity are outlined and the search data are described and revealed. I wrote code to model the WIMP signal in ZEPLIN-III. I contributed to the detailed 'post mortem' analysis of anomalous events in calibration data and search sidebands which informed the development of cuts. The tagging of search events by the veto is reported for completeness, but I was not involved in that analysis. During many weeks over the run I was responsible for the remotely-run daily tasks, or part of the weekly underground team.

4.1 WIMP-nucleus recoil spectrum

A halo of WIMPs, with speed distribution in the target frame $f_{\mathrm{t}}(v)$, will scatter off target atoms to produce nuclear recoil energy E_{R} at a rate determined by the differential cross section $\frac{\mathrm{d}\sigma}{\mathrm{d}E_{\mathrm{R}}}$. The differential rate per unit target mass is:

$$\frac{\mathrm{d}R}{\mathrm{d}E_{\mathrm{R}}} = \frac{1}{m_{\mathrm{N}}} \frac{\rho_0}{m_{\mathrm{W}}} \int_{v>v_{\mathrm{min}}} v f_{\mathrm{t}}(v) \frac{\mathrm{d}\sigma}{\mathrm{d}E_{\mathrm{R}}}(v, E_{\mathrm{R}}) \, \mathrm{d}v \,, \tag{4.1}$$

where m_{W} and m_{N} are the WIMP and nucleus masses, respectively, and ρ_0 is the dark matter density in the laboratory. The lower integration limit comes from the kinematics of non-relativistic elastic scattering:

$$v_{\mathrm{min}} = \sqrt{\frac{m_{\mathrm{N}} E_{\mathrm{R}}}{2\mu_{\mathrm{N}}^2}} \tag{4.2}$$

is the lowest WIMP speed which can cause recoil energy E_{R}, where μ_{N} is the reduced mass of the WIMP-nucleus system.

As the momentum transfer, $q = \sqrt{2m_{\mathrm{N}}E_{\mathrm{R}}}$, increases, the finite size of the nucleus results in a loss of coherence; this is accounted for by form factors, $F^2(q)$ with $F^2(0) = 1$, defined separately for spin-dependent and -independent interactions. We introduce a zero-momentum-transfer or 'standard' cross section σ_0, following [23], to separate the form factor within the cross section:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_{\mathrm{R}}}(v, E_{\mathrm{R}}) = F^2(q)\,\frac{\mathrm{d}\sigma_0}{\mathrm{d}E_{\mathrm{R}}}(v, E_{\mathrm{R}})\,. \tag{4.3}$$

For isotropic scattering the cross section is uniform in the cosine of θ^*, the scattering angle in the centre-of-mass frame. If, in addition, σ_0 — the total cross section after integration over θ^* — is independent of v then $\frac{\mathrm{d}\sigma_0}{\mathrm{d}E_{\mathrm{R}}}$ takes the following simple form:

$$\frac{\mathrm{d}\sigma_0}{\mathrm{d}E_{\mathrm{R}}}(v, E_{\mathrm{R}}) = \frac{\mathrm{d}\sigma_0}{\mathrm{d}\cos\theta^*}\frac{\mathrm{d}\cos\theta^*}{\mathrm{d}E_{\mathrm{R}}} = \frac{\sigma_0}{2}\frac{m_{\mathrm{N}}}{\mu_{\mathrm{N}}^2 v^2}, \tag{4.4}$$

leading to the rate per target mass in terms of σ_0:

$$\frac{\mathrm{d}R}{\mathrm{d}E_{\mathrm{R}}} = F^2(q)\,\frac{\rho_0\sigma_0}{2m_{\mathrm{W}}\mu_{\mathrm{N}}^2}\int\limits_{v>v_{\min}}\frac{f_{\mathrm{t}}(v)}{v}\mathrm{d}v. \tag{4.5}$$

The WIMP-nucleus recoil spectrum thus depends on elementary particle physics via σ_0 and m_{W}; on nuclear physics via $F^2(q)$, and on galactic astrophysics via ρ_0 and the WIMP velocity distribution. The density ρ_0 is determined by fits to observed dynamical properties of the galaxy via an assumed shape of the DM halo. Weber and de Boer [31] report that, depending on the choice of profile, local halo densities in the range 0.2–0.4 GeV c^{-2} cm^{-3} are consistent with current data, also noting that this is a lower limit as the roughly spherical halo may be augmented by a dark disc [126] of co-rotating dark matter from earlier accretion of satellite galaxies.

Zero-momentum-transfer cross sections

The cross section for elastic WIMP-nucleus scattering has, in general, two components from spin-dependent and spin-independent interactions. Scalar and vector interactions result in a spin-independent cross section:

$$\sigma_{0,\mathrm{SI}} = \frac{4\mu_{\mathrm{N}}^2}{\pi}\left(\left[Zf_{\mathrm{p}} + (A-Z)f_{\mathrm{n}}\right]^2 + \frac{B_{\mathrm{N}}^2}{256}\right). \tag{4.6}$$

Here Z and A are the atomic and mass numbers of the target nucleus; f_{p} and f_{n} are the effective scalar couplings to the proton and neutron (calculated for an MSSM neutralino in Ref. [23]). The vector couplings to the valence quarks, $\alpha^{\mathrm{V}}_{q=\mathrm{u,d}}$, determine the vector

WIMP-nucleus interaction:

$$B_{\mathrm{N}} \equiv \alpha_{\mathrm{u}}^{\mathrm{V}}\,(A+Z) + \alpha_{\mathrm{d}}^{\mathrm{V}}\,(2A-Z)\ . \tag{4.7}$$

For comparison with theoretical models and results from other target materials, limits are calculated on the scalar WIMP-*nucleon* cross section, σ_{n}, assuming $f_{\mathrm{p}} = f_{\mathrm{n}}$ and $B_{\mathrm{N}} = 0$. In this case:

$$\sigma_{0,\mathrm{SI}} = A^2 \left(\frac{\mu_{\mathrm{N}}}{\mu_{\mathrm{n}}}\right)^2 \sigma_{\mathrm{n}}\,, \tag{4.8}$$

where μ_{n} is the WIMP-nucleon reduced mass.

Axial-vector interactions result in a spin-dependent cross section [23]:

$$\sigma_{0,\mathrm{SD}} = \frac{32 G_{\mathrm{F}}^2 \mu_{\mathrm{N}}^2}{\pi} \frac{J+1}{J} \left[a_{\mathrm{p}} \langle S_{\mathrm{p}} \rangle + a_{\mathrm{n}} \langle S_{\mathrm{n}} \rangle\right]^2 . \tag{4.9}$$

Here, $\langle S_{\mathrm{p,n}} \rangle$ are the matrix elements for the spin of, respectively, the protons and neutrons in the nucleus; G_{F} is the Fermi coupling constant and J is the total angular momentum of the nucleus. Ref. [127] presents estimates of the spin matrix elements for ^{129}Xe and ^{131}Xe from a shell model which reproduces, within 2%, the measured magnetic moments of these isotopes.

The coupling constants $a_{\mathrm{p,n}}$ are the sum over quark flavor q of the WIMP-quark couplings d_q, in units of $\sqrt{2}G_{\mathrm{F}}$, multiplied by the fraction of the proton or neutron spin which that quark carries, i.e. [128]:

$$a_{\mathrm{p}} = \sum_{q=\mathrm{u,d,s}} \frac{d_q}{\sqrt{2}G_{\mathrm{F}}} \Delta_q^{(\mathrm{p})}, \quad a_{\mathrm{n}} = \sum_{q=u,d,s} \frac{d_q}{\sqrt{2}G_{\mathrm{F}}} \Delta_q^{(\mathrm{n})} \tag{4.10}$$

with

$$\frac{1}{2} \langle \mathrm{p}|\, \overline{q}\boldsymbol{\gamma}\gamma_5 q\, |\mathrm{p}\rangle = \Delta_q^{(\mathrm{p})} \langle \mathrm{p}|\, \mathbf{s}_{\mathrm{p}}\, |\mathrm{p}\rangle\,, \quad \frac{1}{2} \langle \mathrm{n}|\, \overline{q}\boldsymbol{\gamma}\gamma_5 q\, |\mathrm{n}\rangle = \Delta_q^{(\mathrm{n})} \langle \mathrm{n}|\, \mathbf{s}_{\mathrm{n}}\, |\mathrm{n}\rangle\,. \tag{4.11}$$

Limits on the spin-dependent couplings are often reported in the a_{p}–a_{n} plane at fixed WIMP mass [129], with cross-section limits corresponding to allowed regions via 4.9. The relatively low values of $\langle S_{\mathrm{p}} \rangle$ for xenon isotopes give ZEPLIN-III little sensitivity to the WIMP-proton interaction, so only the much more competitive limit on WIMP-neutron cross section for the case $a_{\mathrm{p}} = 0$ is reported.

Form factors

The form factor for spin-independent scattering $F_{\mathrm{SI}}^2(q)$ is determined from electron-nucleus scattering data and may be parametrised as [130]:

$$F_{\mathrm{SI}}^2(q) = \left(\frac{3 j_1(q r_{\mathrm{n}})}{q r_{\mathrm{n}}}\right)^2 \exp\left(-q^2 s^2\right)\ ; \tag{4.12}$$

here j_1 is a spherical Bessel function of the first kind, $s = 1.0$ fm, and $r_n^2 = R^2 - 5s^2$ with $R = 1.2A^{1/3}$ fm. This corresponds to the (squared, normalised) Fourier transform of a nuclear density that is uniform out to radius r_n with a Gaussian shell of thickness s.

There are three independent functions which determine the spin-dependent form factor. In the so-called isospin decomposition these are the isoscalar, isovector and interference-term structure functions: $S_{00}(q)$, $S_{11}(q)$ and $S_{01}(q)$. They are defined fully in [131] and calculated using the Bonn-CD [132] nucleon-nucleon potential in [127]. The structure functions are summed with coefficients determined by the WIMP-nucleon couplings [128]:

$$S(q) = a_0^2 S_{00}(q) + a_1^2 S_{11}(q) + a_0 a_1 S_{01}(q) \tag{4.13}$$

with

$$a_0 = (a_p + a_n), \quad a_1 = (a_p - a_n). \tag{4.14}$$

The spin-dependent form factor is thus obtained:

$$F_{SD}^2(q) = \frac{S(q)}{S(0)}. \tag{4.15}$$

WIMP phase space

The canonical assumed WIMP speed distribution in the rest frame of the galactic DM halo is Maxwellian — as would be the case for an isotropic isothermal DM halo with density proportional to r^{-2} — but truncated at the local galactic escape velocity v_{esc}:

$$f_h(\mathbf{v}) = H(v_{esc} - |\mathbf{v}|) \frac{1}{N_{esc}} \left(\frac{3}{2\pi\sigma_v^2}\right)^{3/2} \exp\left(-\frac{3|\mathbf{v}|^2}{2\sigma_v^2}\right), \tag{4.16}$$

where $H(\cdot)$ is a Heaviside step function and N_{esc} is a normalisation factor. In this standard halo, the RMS speed σ_v is $\sqrt{3/2}$ times the local circular speed, which is measured as $v_c = 220 \pm 20$ km s^{-1} [133]. A study of high-velocity stars in the RAVE survey sets a 90% confidence interval for v_{esc} of 498–608 km s^{-1} [134]. This level of uncertainty does not significantly change elastic scattering rates for moderate WIMP mass ($\gtrsim 50$ GeV c^{-2}), but the effect can be strong for low masses or particle interaction models particularly sensitive to the tail of the velocity distribution [135].

The detector moves in the halo frame with a velocity, $\mathbf{v}_E(t)$, which is the sum of the velocity of the Sun in the halo frame and that of the Earth relative to the Sun. It is parametrised in Ref. [136], which also contains the result of the integral in equation 4.5 when the velocities described by equation 4.16 are boosted into the target frame. Defining $x \equiv v_{min}/v_c$, $y \equiv |\mathbf{v}_E|/v_c$ and $z \equiv v_{esc}/v_c$:

$$\int\limits_{v>v_{\min}} \frac{f_{\rm t}(v)}{v}\mathrm{d}v = \begin{cases} \frac{1}{v_{\rm c}y}\,, & z<y,\ x<|y-z| \\ \frac{1}{2N_{\rm esc}v_{\rm c}y}\left[\mathrm{erf}(x+y)-\mathrm{erf}(x-y)-\frac{4}{\sqrt{\pi}}ye^{-z^2}\right], & z>y,\ x<|y-z| \\ \frac{1}{2N_{\rm esc}v_{\rm c}y}\left[\mathrm{erf}(z)-\mathrm{erf}(x-y)-\frac{2}{\sqrt{\pi}}(y+z-x)e^{-z^2}\right], & |y-z|<x<y+z \\ 0\,, & y+z<x\,. \end{cases} \tag{4.17}$$

The annual modulation of the WIMP signal is caused by the changing $|\mathbf{v}_{\rm E}|$ as the Earth's velocity relative to the Sun rotates with respect to Sun's velocity in the halo frame. There is a WIMP wind in the laboratory, and the wind speed changes through the year. Fig. 4.1 shows the velocity integral in equation 4.5, as function of $v_{\min}$, evaluated at the two annual extrema.

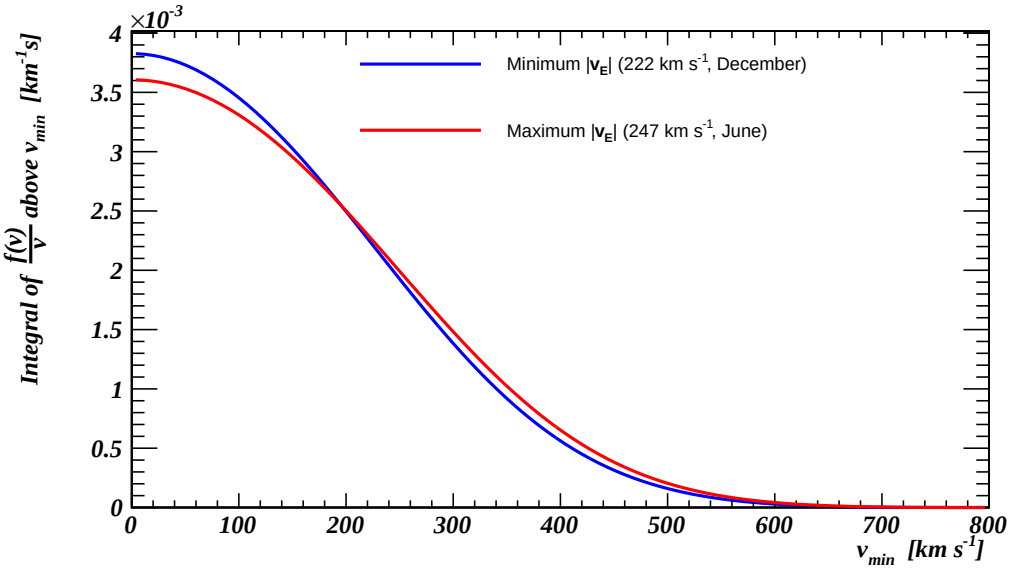

Figure 4.1: The WIMP velocity integral which determines scattering rates via equation 4.1, at the annual extrema of Earth's velocity in the halo. The rate maximum occurs in June (December) for $v_{\min}$ above (below) about 200 km s^{-1}.

4.2 Modeling signal in ZEPLIN-III

WIMP-nucleon cross sections, form-factors and a halo model together predict a differential rate in $E_{\rm R}$ of nuclear scattering events, illustrated in Fig. 4.2. What is measured by ZEPLIN-III, however, is a rate in reconstructed S1 and S2 of those events which are recorded and pass all selection cuts. Evidently, a signal model must account for instrumental resolution and efficiency. Fluctuations in S1 and S2 for nuclear recoils are not significantly anticorrelated at fixed recoil energy [137], in contrast to electron recoils at

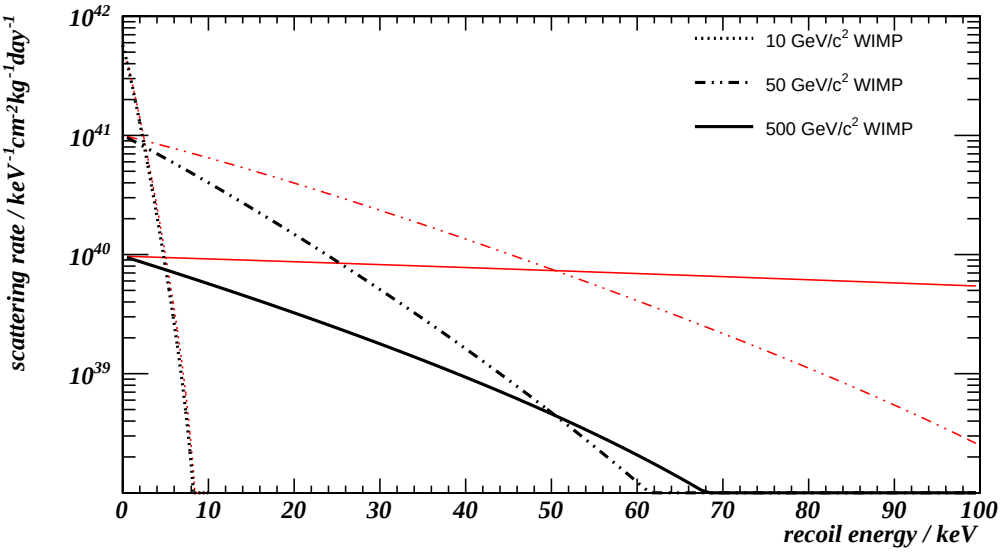

Figure 4.2: Spectra (in black) of spin-independent WIMP-nucleus scattering in a xenon target per unit energy, exposure and WIMP-nucleon cross section, and for the annual mean $|\mathbf{v}_E|$ in the standard halo model. Red lines show the spectra before accounting for the nuclear form factor.

~ 100 keV. The signal model is therefore built up in two stages: a signal spectrum in S1, and a distribution of S2 for nuclear recoils at fixed S1 obtained from neutron calibration.

4.2.1 Signal rate in S1

Calculation of the signal spectrum in S1 begins with the differential rate in equation 4.1; this is what would be recorded in the case of unit efficiency and perfect energy resolution. Next, the scintillation yield, parametrised by $\mathcal{L}_{\text{eff}}$, is used to calculate a differential rate per unit S1, again for idealised detector response. Statistical fluctuations in the various stages of signal generation and detection will affect the amount of scintillation reconstructed by Mercury, and so they must be convolved with the spectrum in S1. In the energy range of interest for WIMP-nucleus recoil, the main sources of variance are the Poisson-distributed number of photons and photoelectrons produced, and the PMT responses to the latter. The photoelectron distributions obtained in section 3.1 along with the light yield from section 3.5 result in an RMS deviation of $\sigma_{\text{S1}}\,[\text{keVee}] \approx 1.0\sqrt{\text{S1}\,[\text{keVee}]}$.

Finally, the efficiency for reconstruction as a golden event in the final dataset is applied. Inefficiency is incurred because of the required three-fold coincidence among the PMTs, as well as the various physics and reconstruction-quality cuts described in section 4.4. The PMT coincidence efficiency for ZEPLIN-III is encoded in the Mercury light-response functions and photoelectron distributions, since these give the distributions of PMT signal size as a function of (E, x, y, z). The fraction of events with at least 3 PMTs recording signal area above 1/3 of a photoelectron can be predicted, as a function of S1, for all points

within the fiducial volume. The average value — weighted uniformly over the volume, as the WIMP signal events would be distributed — produces the efficiency η_1 (S1), shown in Fig. 4.9 on page 105. Examples of signal spectra in S1, before and after detector effects are accounted for, are shown in Fig. 4.3.

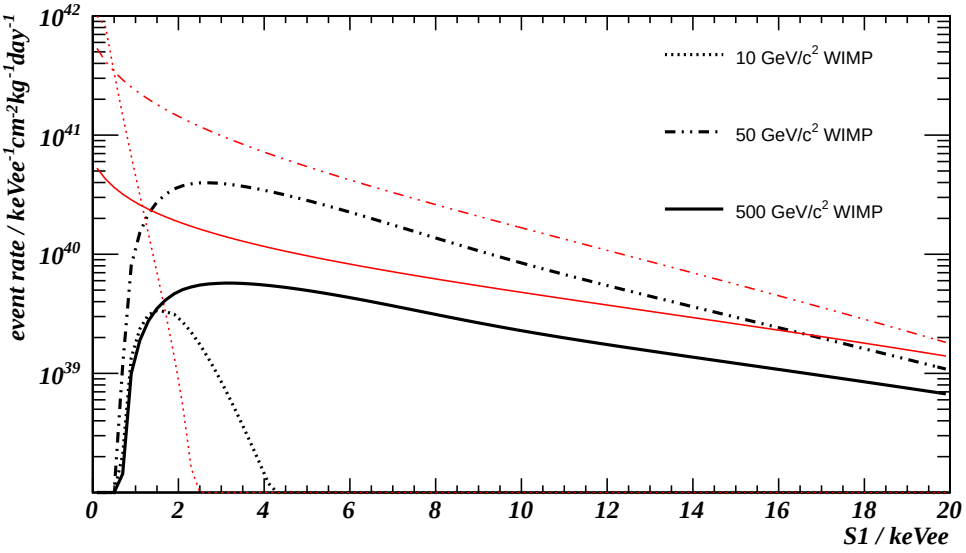

Figure 4.3: Signal spectra (in black) of primary scintillation reconstructed in ZEPLIN-III, S1, due to spin-independent scattering of 3 WIMP masses. Red lines show the spectra in ⟨S1⟩, before resolution and efficiency are accounted for.

4.2.2 Nuclear recoils in S2

The distribution of S2 at fixed S1 for nuclear recoils is found, empirically, to be approximately log-gaussian. Gaussian fits, as shown in Fig. 4.5, were made to the distribution in $\log_{10}(\mathrm{S2}/\mathrm{S1})$ of AmBe data in 2-keVee bins of S1 from 2 to 12 keVee. All three AmBe calibration runs, described in section 3.5, were combined. The variations of the fit parameters with S1 were, in turn, parametrised by polynomials to give a smoothly-varying distribution as shown in Fig. 4.6. The signal model is based on fits to events passing the final set of cuts, to be described in section 4.4.

The fitted log-gaussian distribution of nuclear recoil events may be used to calculate a signal-acceptance coordinate, $a\,(\mathrm{S2/S1;S1})$, which is the cumulative distribution of nuclear recoil events in S2/S1 at given S1. Plotting search events in these coordinates, as in Fig. 5.4 on page 116, makes the presence or absence of signal more immediately evident, and illustrates the discrimination power of the scintillation-to-ionisation ratio.

For events with S1 $\lesssim$ 50 keVee, the shaped sum of the PMT pulses due to primary scintillation is too small to activate the hardware trigger: it is the later, secondary signal that appears at trigger point, with S1 in the pre-trigger part of the waveform. If S2 is

sufficiently small then an event may be missed entirely, and this possible inefficiency was investigated via three questions:

- What amount of detected S2 light, i.e. summed PMT pulse area, will reliably trigger acquisition?
- What is the yield of pulse area per unit ionisation in the liquid, as a function of day, depth and radius?
- What fraction of WIMP-nucleus signal events with S2 within the search region could therefore fail to trigger?

Software emulation of the trigger electronics has shown that the threshold pulse area for fast, S1-like pulses is equal, within a few percent, to that from pulses with $\tau \simeq 500$ ns. In terms of pulse area, the trigger threshold for secondaries is therefore taken to be the same as that observed for primaries. 90% of primaries with raw pulse area of 4.0 nVs are found at the trigger point and above 5.0 nVs all golden events trigger on the primary. The yield of raw pulse area per unit ionisation decreases towards the outer circumference of the fiducial volume. Its minimum, in the outer 2.5-mm-thick annulus, is (16.2 ± 0.2) nVs per keVee in S2 for events near the surface. The yield also decreases with depth, more strongly for days with shorter electron lifetime in the liquid, due to drifting charge being captured by electronegative impurities. These effects together give the overall distribution of the SSR exposure in raw area per unit ionisation, shown in Fig. 4.4.

The low-S2 corner of the search box corresponds to 0.58 keVee ionisation leaving the interaction site, and so the hardware trigger is fully efficient where 0.58 keVee multiplied by the S2 area yield exceeds 5.0 nVs. Overall, less than 3% of the SSR exposure in kg.days suffers from S2 inefficiency anywhere in the search box, and the change to total expected event rates is negligible.

4.3 SSR data

The second science run extended over 319 days from June 23, 2010 to May 7, 2011. In this time over 11 000 binary files, each of 1000 events and ~ 400 MB compressed size, were acquired in search mode. During the daily nitrogen fill, data were copied from the ZEPLIN-III DAQ PC to an underground analysis PC. They were then written to duplicate magnetic tapes, with one retained in the Boulby surface building and the other copied to disk at RAL. Search data were not analysed during daily detector checks, except for noting the DAQ trigger rate. Timelines in search-quality files were hidden from view in ZE3RA until the staged unblinding of events tagged by the veto or placed outside the search box (see sections 4.3.1 and 4.4.2).

Files were classed as search-quality or not (the latter known as ‘bronze’ data) based on the parameters of the laboratory and internal detector environments as well as the proper functioning of detector subsystems and absence of discharging within the chamber.

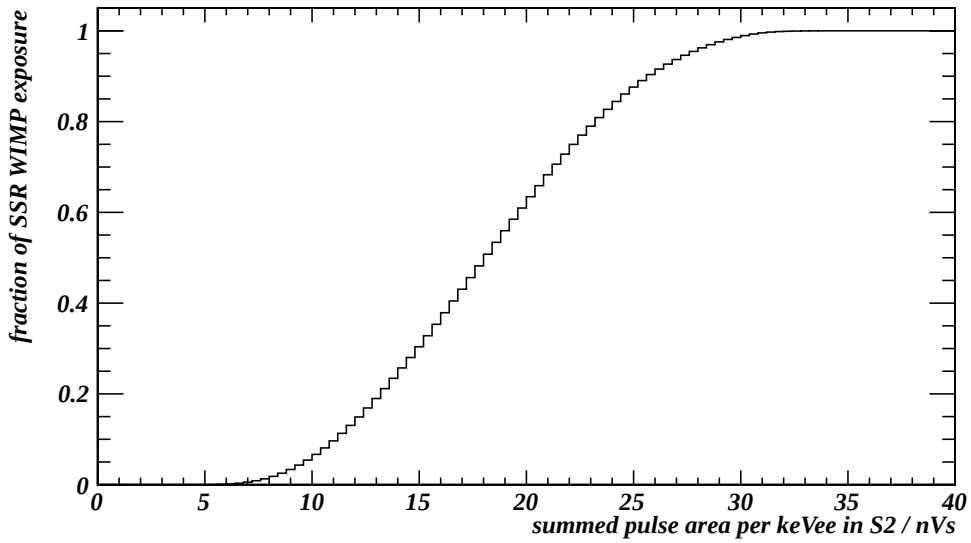

Figure 4.4: The cumulative distribution in ionisation response, as seen by the trigger, of the SSR exposure. The raw pulse area per unit S2 for surface events is profiled in cells of x and y via Mercury reconstruction. The distribution of lifetime corrections for events uniform in depth is also calculated, using the measured electron lifetime for each day. Combining the two effects, a response for each element of WIMP exposure (in x, y, z and t) is obtained, giving the distribution shown. The low-S2 edge of the WIMP search region triggers with 100 (90) % efficiency where the response exceeds 8.6 (6.9) nVs / keVee.

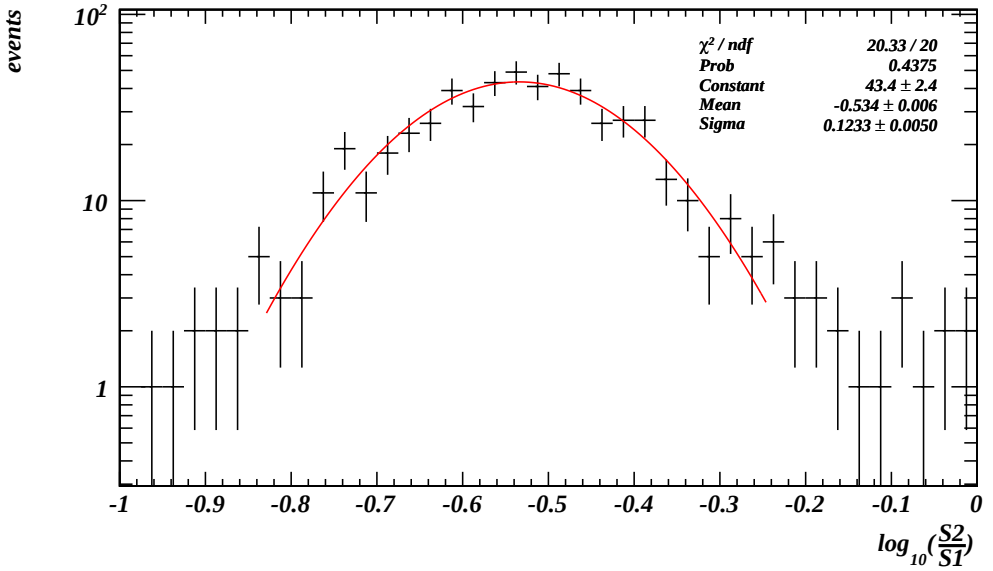

Figure 4.5: An example of the gaussian fits to nuclear recoil events in AmBe calibration data: the 6–7 keVee bin in S1.

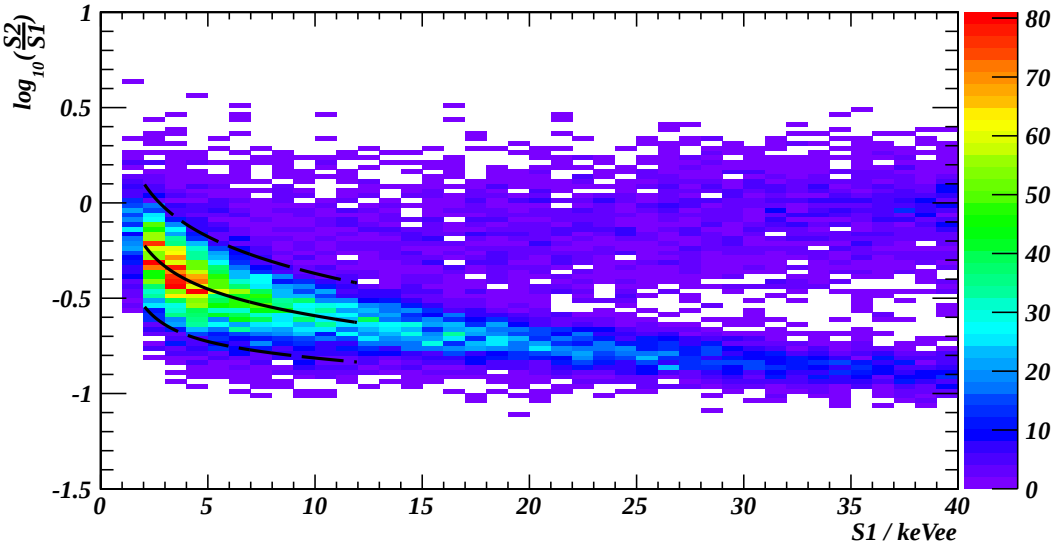

Figure 4.6: Scatter plot for combined AmBe calibration data, with parametrised mean (solid line) and 2-σ bounds (dashed lines) for nuclear recoil events in the search region of S1.

The commissioning phase in June 2010, time following power cuts, and a failure of the connection to PMT1 all resulted in bronze data which were immediately open to analysis. These helped to refine all stages of analysis software in a process that included visually inspecting the PMT traces of thousands of events, especially at low S2/S1. By the end of the run, search-quality data were collected from a time exposure of 265.5 effective days, as shown in Fig. 4.7. The final Golden code returned around 250 000 events with $\mathrm{S1_{Golden}}$ below 70 keVee having exactly one S1 and one S2, before fiducial volume and other cuts. These comprise the WIMP search dataset for the second science run. Events up to 500 keVee in E^{*}_{Golden} were included in a set of around 350 000 events from 30 days in August 2010 that was used to study higher-energy background (as described in the previous chapter).

4.3.1 Veto data

The 52 signal traces from the veto instrument are recorded in a separate data acquisition system from the ZEPLIN-III (target) DAQ. An additional channel in the veto DAQ records the shaped sum of the target PMT signals, in order to confirm event synchronisation. Acquisition of an event is initiated either by the target hardware trigger, by a muon trigger derived from the summed signal in the veto roof modules, or when 10 photoelectrons are recorded in coincidence in each of three veto modules. The veto and target data corresponding to the same event may be synchronised using a high-precision time-stamp fed to both systems from an external synchronisation unit. This unit was not always operational, but equally unambiguous results were achieved by using an internet time

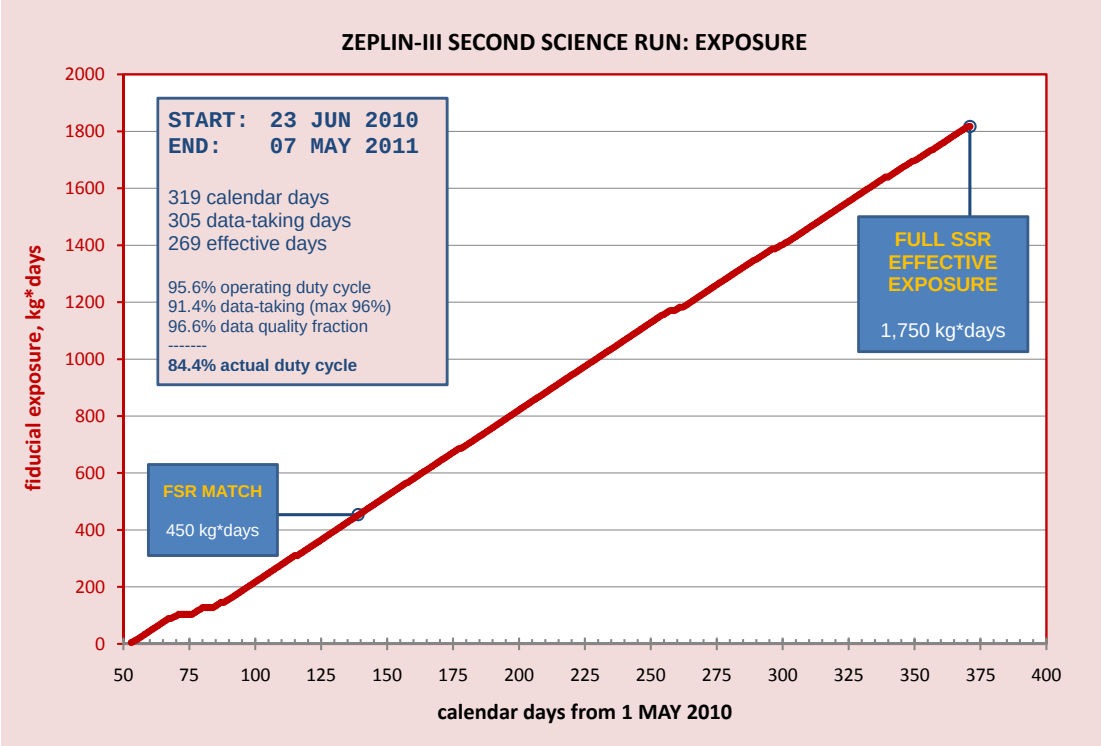

Figure 4.7: Accumulation of search-quality data during the second science run (plotted on the basis of a preliminary fiducial volume of 6.5 kg, 30% larger than the one eventually adopted).

server, accurate to 4 ms, in conjunction with matched time intervals between target triggers and matched pulse timings in the target sum channel; as target triggers occur at a rate of only 0.4 s^{-1}, these methods together leave a negligible fraction of unsynchronised events [76].

In the context of the dark matter search, the role of the veto is to identify a fraction of background events as definitely not being due to WIMP-nucleon scattering. This reduces the experimental background directly, but also provides a sample of events that can be used to characterise backgrounds and refine analysis software. As described in Ref. [76], the signals in the veto PMTs are reduced to a system of event classifications, or 'tags', with known probabilities for different physical event types. ZEPLIN-III events may be un-tagged, prompt-tagged, or delayed-tagged by the veto.

When a γ-ray scatters in the target and in the veto, the two scintillation pulses will be effectively simultaneous, although the ZEPLIN-III trigger itself will follow up to 16 µs later if it is derived from the S2 electroluminescence. A prompt veto tag is set when coincident scintillation, defined as at least 2 photoelectrons ($\gtrsim$40 keV energy deposited) in total across all modules, is found within ±200 ns of the S1 pulse in the target. The probability of a search event in the target being *accidentally* prompt-tagged, due to an unrelated scatter in the veto, is 0.4%. Using a longer time window improves the tagging efficiency for background but only as would be expected from accidental coincidences, consistent with the 100 ns sampling time of the veto DAQ.

Within the gadolinium-loaded polypropylene modules, neutrons are moderated and then undergo radiative capture by ^{157}Gd with a mean capture time of (10.7 ± 0.5) µs [76]. The subsequent de-excitation of ^{158}Gd emits up to 8 MeV of energy split, typically, among 3 or 4 γ-rays which may then be detected in the active modules of the veto. Because of the time take taken to moderate and capture the neutron, scintillation in the veto is not simultaneous with any primary scintillation that the neutron may have caused in the xenon target. A delayed veto tag is therefore used to mitigate neutron-induced background. It is set when an event is not prompt-tagged and a total, across all veto PMTs, of at least 10 photoelectrons occurs in any 400-ns window up to 70 µs after the S1 in the target. Random coincidences will meet the latter criterion in 1.0% of search events, and so 1.0% of WIMPs will be accidentally delayed-tagged; the fraction is smaller for γ-ray background events, as prompt-tagging removes many of these events from consideration for delayed-tagging. The efficiency for the delayed tagging of neutron events is found, from the AmBe calibration data, to be 58.8%. In addition, neutron background events have a 1.7% probability to be prompt-tagged, mainly due to elastic scattering in the veto scintillator. All the veto tagging probabilities are summarised in table 4.1: the veto removes significant fractions of the neutron and γ-ray backgrounds, with little loss in signal efficiency.

After the veto analysis was finalised, tagged ZEPLIN-III events at all values of S1 and S2 were unblinded. Delayed-tagged events were not found below the nuclear recoil median in $\log_{10}(\mathrm{S2}/\mathrm{S1})$, and the total number was consistent with the expected fraction of random coincidences for electron-recoil background. The fraction of search events,

within the fiducial volume and having S1 below 100 keVee, that were prompt tagged is $(28.1 \pm 0.2)\%$; this is compatible with the weighted probability of $(27.4 \pm 0.6)\%$ obtained from Monte Carlo simulation of the dominant γ-ray backgrounds, where both figures include accidental coincidences.

Table 4.1: Simulated and observed veto tagging probabilities

Event type	Prompt tag [%]	Delayed tag [%]
neutron background, simulated	1.5 ± 0.1	59.2 ± 0.1
neutron calibration, simulated	”	”
neutron calibration, observed	1.7 ± 0.1 [a]	58.8 ± 0.5
γ-ray background, simulated [b]	27.4 ± 0.6	0.7
search data, observed [b]	28.1 ± 0.2	0.69 ± 0.02
WIMP signal, calculated	0.4	1.0

[a] extrapolated from distribution of delayed γ-rays
[b] S1 < 100 keVee

4.3.2 Physical origin of search events

Simulations of background events in the ZEPLIN-III SSR, using the GEANT4 toolkit [114] and based on radioactivity measurements of the materials in and around the detector, are detailed in [113]. Neglecting components for which only upper limits on activity are available, they predict an electron recoil background of 0.85 ± 0.05 keV^{-1}kg^{-1}day^{-1} below ~ 100 keV in the fiducial volume. This is in excellent agreement with the 0.75 ± 0.05 keVee^{-1}kg^{-1}day^{-1} in E^* which is observed for search data, demonstrating the thorough understanding of background and detector response necessary for a rare event search. The close agreement between the simulated prompt tagging fraction for γ-ray backgrounds and that observed for search events also confirms that γ-rays, rather than β-particles, dominate the low-energy event rate in ZEPLIN-III.

The same radioactivity measurements which predict the electron-recoil background imply neutron single scatters in the fiducial volume with recoil energy 5–50 keV at a rate of 3.1 ± 0.5 scatters per year, mainly due to uranium and thorium in the ceramic feedthroughs and PMTs. This corresponds to an expectation of only 0.06 ± 0.01 background events in the search box after exposure, trigger efficiency, cuts and veto tagging are taken into account. The neutron background is independently, but more weakly, constrained by the search data, specifically the observation of only one delayed-tagged event over all energies up to the 84th percentile in $\log_{10}(\mathrm{S2/S1})$ for nuclear recoils; this implies an upper limit of 0.75 events in the search box (90% CL).

The events which dominate the background to the WIMP search are thus not neutrons but the low-S2/S1 outliers of the electron recoil background. These outliers may be well-reconstructed events that have fluctuated to much higher S1 or lower S2 than average;

this is more likely for low recoil energies, where the S1 signal is only a few photoelectrons. Systematic effects can also lead to spuriously low S2/S1 values: an event with a different topology may be reconstructed under the assumption that it is a simple single scatter; a vertex may occur above a patch of particularly poor quantum efficiency in a non-uniform photocathode (S1 light typically being more widely distributed across the array than S2 light); or PMT signals due to scintillation light may be clustered together with electronic pickup or optical noise (e.g. afterpulsing or Cherenkov light radiated by a β from within a PMT passing through the quartz window). A potentially damaging example of anomalous topology is multiple-scintillation-single-ionisation (MSSI) events. These occur when a background γ-ray, for example, scatters once in the region of applied field from where ionised charge is extracted, and one or more times in a so-called dead region where only the S1 signal is generated, e.g. below the cathode grid. The effectively simultaneous primary signals will be clustered together, lowering the reconstructed S2/S1. Analysis cuts to minimise electron recoil background in the signal region are described in the next section.

4.4 Selection cuts

The chain of ZE3RA, Golden and Mercury reconstruction programs calculates many dozens of event parameters in addition to the estimators of scintillation, ionisation and position. These characterise, for example, the pulse areas and shapes of S1 and S2 in individual PMTs, and the goodness-of-fit for the reconstructions of the vertex. The SSR event selection cuts fall into two broad categories: those that guard against systematic mis-reconstruction, and those that aim to remove well-reconstructed electron-recoil single scatters to minimise the chance of statistical fluctuation into the search region. Unvetoed search events far from the search region and vetoed events everywhere in S2/S1 were compared with AmBe data to evaluate the efficiency of cuts for background- and signal-like events, typically as a function of S1.

4.4.1 Cuts on quality of reconstruction

When searching for a rare process, a cautious attitude must be taken against systematic mis-reconstruction of even small numbers of events. The state of the detector at the time of the event leads to a first category of quality cuts. Variation in the external environment of the target — including the cryostat, shielding cavity and lab — can alter detector response by increasing electronics noise or moving the two-phase system away from equilibrium. Outliers in lab temperature, for instance, were therefore removed from the dataset. Events are also cut based on the low-level properties of PMT waveforms: for example, large discrepancies between pulse areas clustered by different algorithms, which may be characteristic of pickup. These cuts are accounted for as part of the energy-independent efficiency, η_0.

Mercury reconstructions rely on the assumption of a single scattering vertex and several cuts aim to remove events for which this assumption does not hold. Primary or secondary signals that are unusually tightly or widely distributed among PMTs, or highly asymmetric pulse shape in secondaries can indicate a multiple-vertex event (as well as one affected by noise). The distance between the S1 and S2 position, reconstructed separately, is also used to mitigate MSSI events. The reduced χ^2 of the two signal fits and the individual contribution from the least-likely PMT response quantify the goodness-of-fit of Mercury reconstruction and proved effective in removing outlying electron-recoil events.

4.4.2 Cuts targeting well-reconstructed background events

The fiducial region in z is determined by the spatial distribution of background compared with the uniform density expected of WIMP-nucleon signal events. As shown in Fig. 4.8, the rate of electron recoil events due to decays in the veto and shielding materials drops rapidly in the first few millimetres of liquid depth due to self-shielding. Radon progeny embedded in the cathode grid result in a similar peak at the bottom of the active volume. The fiducial volume was set at 1–13 µs drift time, corresponding to 2.4–30.8 mm below the liquid surface. Selection in x-y is motivated by discrimination rather than rate of background: resolution in S1 and S2 degrades towards the edge of the array, widening the distribution of reconstructed S2/S1. Before the final unblinding it was decided, on the basis of the x-y distribution of events near the search box, to require a reconstructed S2 within 140 mm of the central axis. The fiducial target mass is then 5.1 kg.

Scintillation pulse shape may be used to discriminate signal from electron recoil background events. As discussed in section 2.1, direct scintillation light in liquid xenon is emitted in the radiative dissociation of the excimer , Xe_2^*, which exists in either the singlet or triplet state. The decay times of the two states are measured to be 4 and 21 ns, respectively. Hitachi et al. [138] found an increase in relative intensity of the singlet state with linear energy transfer (LET, i.e. -dE/dx) in a study of energy loss by electrons, α-particles and fission fragments; this is explained therein by super-elastic collisions with thermal electrons, which convert the singlet state to the (lower energy) triplet state: $S + e^- \rightarrow T + e^-$. In the relatively dense nuclear recoil tracks, free electrons are quickly removed by recombination with Xe_2^+ and the intensity ratio of the two states, S/T, remains close to its primitive value of 1/3. However, in electron-recoil tracks charge recombination is less immediate, with mean time ~45 ns, resulting in a significant fraction of the singlet-state excimers scattering with thermal electrons; the triplet state is populated at the expense of the singlet and so the scintillation pulse is extended in time. The delay of the recombination time in itself further lengthens S1 for electron recoils, especially at low or zero field. The combined effect was the basis of discrimination in the single-phase ZEPLIN-I detector [121, 139].

Pulse-shape discrimination was implemented in the SSR via the full-width-at-half-maximum of the S1 signal; this provides a robust parametrisation of scintillation timing,

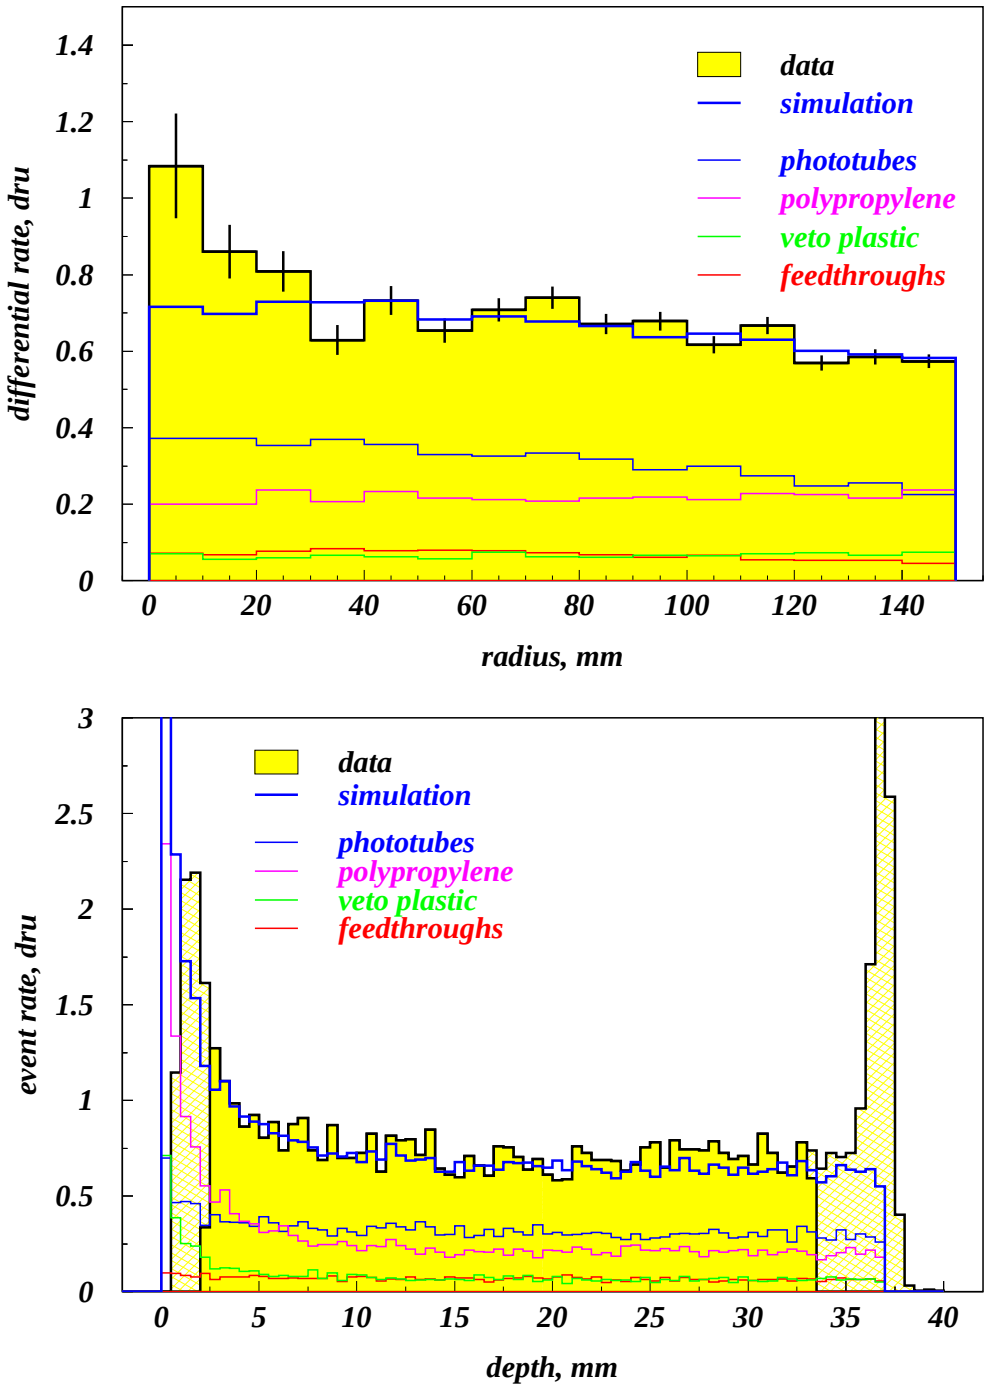

Figure 4.8: Spatial distributions of search events, from Ref. [98]. The discrepancy in the first bin of radius is attributed to higher-than-average radioactivity in the central PMT. Regions outside the fiducial volume in depth are indicated by pale yellow.

less sensitive to pickup and ringing than the integrated mean arrival time (see section 6.4) that was used with the FSR array. The measured efficiencies of the cut, set to require a width below 19 ns, are shown as η_2 in Fig. 4.9. Pulses consisting of few photoelectrons give only poor pulse-shape resolution, and so this cut is most effective for S1 $\gtrsim$10 keVee. The efficiency for signal events is plotted in Fig. 4.9, which also shows the relative contributions of the three sources of energy-dependent inefficiency (S1 reconstruction, timing, and pulse quality cuts).

Event selection based on $\log_{10}(\mathrm{S2/S1})$ — by far the most powerful cut against electron recoil events — was set through a progressive unblinding of the unvetoed search data. Events with electron equivalent energy S1 > 20 keVee or NR acceptance $a(\mathrm{S2/S1;S1}) > 0.84$ were opened at first. These are far from the search box that was initially defined from calibration data, so that adjusting its boundaries as the analysis was refined would not compromise blinding. Once the reconstruction codes were finalised, events at $a > 0.5$ were opened to assess whether a search box extending to the median of signal acceptance, as in the first run, was viable. Fits to the background distribution in the range 2–12 keVee predicted approximately 3 events with $0.45 < a < 0.5$, suggesting limited sensitivity to signal there; the region was discarded from the search, opened and used to inform analysis cuts. The final blinded box was defined by $2 < \mathrm{S1/keVee} < 12$ and $0.023 < a < 0.45$. Fits to the background distribution at higher a gave a prediction of 6 ± 4 events expected in the box.

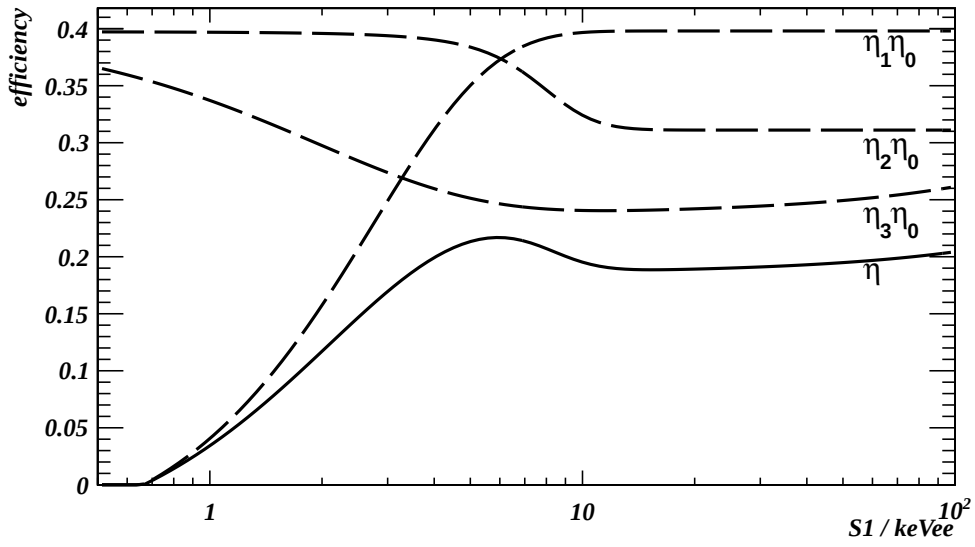

Figure 4.9: Efficiencies for single-scatter nuclear recoil events within the fiducial volume. The 39.8% energy-independent efficiency, dominated by the cut on $\log_{10}(\mathrm{S2/S1})$, is denoted by η_0. The efficiencies of the three-fold PMT coincidence requirement, the S1 width cut and the quality cuts are, respectively, η_1, η_2 and η_3. The overall detection efficiency (solid line) is $\eta = \eta_0\eta_1\eta_2\eta_3$. The WMIP search region is 2–12 keVee.

Table 4.2: Cumulative cut efficiencies for SSR search data. These data are likely to consist entirely of electron-recoil background events, against which the cuts are largely aimed. The efficiency for signal-like nuclear recoil events is treated in Fig. 4.9.

Golden events (single S1, S2) in fid. volume, 2–12 keVee:	6194
after environmental, noise cuts:	6053 (97.8%)
after veto cuts:	4458 (72.0%)
after S1 pulse-shape cut:	3801 (61.4%)
after Mercury goodness-of-fit cuts:	2902 (46.7%)
within search box (0.023–0.45 acceptance in $\log_{10}(S2/S1)$):	8 (0.28%)

4.4.3 Observed search events

Upon opening, the search box was found to contain 12 events. Analysis cuts had already removed all vetoed events from the region. At this stage, goodness-of-fit cuts were reviewed and, in some cases, tightened. An additional temperature, of part of the liquid-nitrogen delivery system, was included in the list of environmental parameter cuts after deviations were found to correlate with events below the median in $\log_{10}(S2/S1)$. Compromising blindness of the search result, albeit only for a handful of cuts, is clearly disappointing; one of the arguments in favour of the SSR upgrade was that, by applying the lessons of the FSR analysis to a very similar instrument, we might better anticipate pathological events and sources of variance. However, the SSR instrument was not as similar as had been planned: the lower attainable field and, especially, far greater variability in PMT response (both across the array and for individual tubes) degraded performance. Measured by the fractional leakage below the nuclear recoil median (assuming no WIMP signal) the discrimination power decreased from 1:7800 to 1:280 in the two runs. The final analysis cuts produce the search dataset shown in Fig. 4.10, with 8 events in the search box. (Table 4.2 quantifies the power of the various classes of analysis cut on background events.) The statistical interpretation of this observation in terms of WIMP-nucleon scattering is treated in the next chapter.

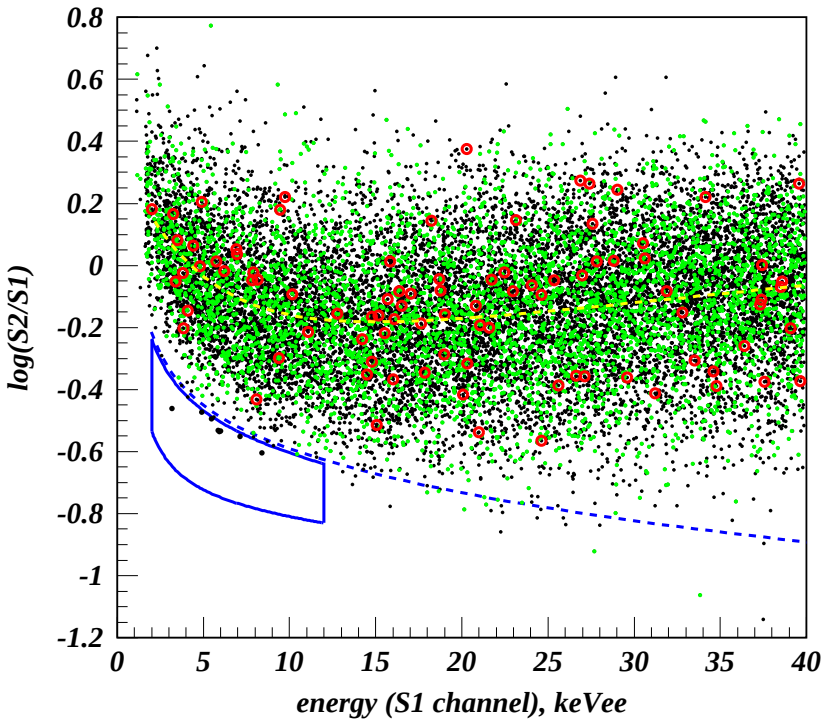

Figure 4.10: The final search dataset of the second science run, from [78]. Green points are prompt-tagged events; red circles are delayed-tagged events. The median $\log_{10}(^{S2}/_{S1})$ for nuclear recoil (background) events is indicated by a dashed blue (yellow) line. The solid blue lines define the search region, which contains 8 events.

Chapter 5

Limits on cross section — statistical techniques and SSR results

Direct searches for dark matter aim to make inferences about the cross section for WIMP-nucleon scattering, σ_{n}; a convenient and intuitive parameter of interest, directly proportional to σ_{n} for a given WIMP mass and assumed halo, is the expected number of signal events, μ, in a search region. In this chapter the properties of frequentist confidence intervals (CIs) for μ are set out, and a number of possible constructions are evaluated with reference to ZEPLIN-III searches. Finally, event rate limits for the FSR and SSR are obtained, and the resultant limits on the WIMP-nucleon cross section are presented. I performed the main statistical analysis of event rate and some of the analysis of the ^{137}Cs calibration. I made one of two independent calculations of cross section within the collaboration, cross-checking the spin-independent limits.

5.1 Frequentist confidence intervals

Scientific hypotheses are subject to falsification; the statistical framework of Neyman and Pearson [140] leads to statements of falsification of which at most some known fraction are mistaken, over repeated trials. If α denotes the fraction of trials which would mistakenly exclude some hypothesis, H, were it to be true, then $(1-\alpha)$ is the coverage of the test under H. A model with one or more unknown parameters is known as a composite hypothesis, or a family of hypotheses; the set of parameter values not excluded by some test is a confidence region with confidence level equal to the minimum (over the model parameters) of coverage.

In the case of direct WIMP detection the parameter of interest is, typically, the WIMP scattering cross section. The signal model can also include uncertain parameters not of primary interest describing, for example, the phase space of the dark matter particles, the detector response and backgrounds. These are known as nuisance parameters. It is not practical to systematically explore coverage for models with many nuisance parameters, but profiling — testing only in conjunction with the nuisance parameter values which maximise likelihood given the parameters of interest — gives CIs which are expected to

have correct coverage over all parameter space [141]. Coverage is the defining characteristic of a frequentist test; it is then desirable to reject hypotheses as often as possible if they are wrong, i.e. to maximise power.

Confidence regions may be constructed by calculating from the observed data (for example events and their measured S1) a test statistic t (for example the largest gap between events in S1) with known distribution under a given hypothesis. A range of values of t which have joint probability $\geq (1-\alpha)$ under H is known as an acceptance region of that hypothesis, and the confidence region is the set of hypotheses for which the observed t is within the acceptance region. The test statistic is simply a means of ordering possible observations for inclusion into the acceptance region.

Ordering principles which accept the zero-signal hypothesis for all values of t will result in single-sided limits: confidence intervals for signal-strength parameters will necessarily extend from 0 to an upper limit. Although discovery potential is sacrificed, CIs of this type can be used in the presence of un-modeled backgrounds; if additional background events can only increase the observed upper limit then (conservative) coverage is obtained. A two-sided CI, capable of excluding the zero-signal hypothesis, necessarily relies on some knowledge of the experimental background.

5.2 Yellin statistics

A family of test statistics due to Yellin [142] is based on the fraction of signal expected in the largest empty region of signal acceptance. The idea was originally applied to the case of a single experimental parameter, recoil energy, and was known as the 'maximum gap' statistic but it has since been generalised to N experimental parameters [143, 144]. In a space in which signal is uniformly distributed, let g be the fraction of signal expected in the largest N-cuboid (orthotope) observed to contain zero events. A signal expectation hypothesis, μ, is rejected as being too high when the observed g is in the upper tail of the distribution that would be expected if μ were true. Therefore an un-modeled background will lead to conservative upper limits, because the largest region with no events of any kind cannot exceed the largest region with no signal events.

To obtain g, one moves from the experimental parameters $s_1, \ldots s_N$ in which the signal PDF is $f_N(\mathbf{s})$ to signal-acceptance coordinates $r_1, \ldots r_N$ in which the signal PDF is unity. Note that there is freedom of choice in this step; before seeing the data experimenters are entitled to choose, from the infinite possible mappings, the one they think will best separate signal and background. In the canonical mapping [143], one defines the projection of $f_N(s_1, \ldots s_N)$ onto the first p signal parameters inductively as f_p for $0 \leq p \leq N-1$:

$$f_{p-1}(s_1, \ldots s_{p-1}) = \int_{-\infty}^{+\infty} f_p(s_1, \ldots s_{p-1}, \xi) \mathrm{d}\xi. \tag{5.1}$$

Then the coordinates $r_{1\ldots N}$ are given by:

$$r_p\,(s_1,\ldots s_p) = \frac{1}{f_{p-1}}\int_{-\infty}^{s_p} f_p(s_1,\ldots s_{p-1},\xi)\mathrm{d}\xi. \tag{5.2}$$

The largest empty region of a given shape in signal-acceptance coordinates is used as a test statistic. For $N = 1$ this is necessarily a line segment, the maximum gap. In higher dimensions, any shape may, in principle, be used; the N-cuboid is adopted for ease of computation. For $N = 2$ it is the largest empty rectangle in r_1–r_2, known as the 'maximum patch'.

In the case of ZEPLIN-III, the principal backgrounds near the search box are the tail of single-scatter electron recoils low in $\log_{10}(\mathrm{S2}/\mathrm{S1})$, and MSSI events which typically have higher reconstructed S1 values than signal events. Thus even without quantitative background models, the maximum patch with $s_1 = \mathrm{S1}$ and $s_2 = \log_{10}(\mathrm{S2}/\mathrm{S1})$ can set limits which are less conservative than simple counting statistics. The more separated the background and signal PDFs are, the lower maximum-gap limits are expected to be.

An algorithm for finding the maximum patch among a set of events in r_1–r_2 is given in Apendix II of [144]; a slightly modified version of this was applied to Monte-Carlo-generated trials to derive $h_j\,(g)$: the CDFs of g when exactly j events are drawn from a uniform distribution in signal-acceptance space. The CDF of g under a signal-only hypothesis μ is given by

$$G\,(g) = \prod_{j=0}^{\infty} \Pr\,(j|\mu)\cdot h_j\,(g)\,. \tag{5.3}$$

The observation of an empty rectangle covering a fraction g of acceptance space thus corresponds to a limit on μ as shown in Fig. 5.1. Also shown for comparison is the single-sided Poisson limit that would be obtained by managing to specify the position of the patch beforehand as the search box (or — as in Refs [56] and [145] — picking a search box after seeing the data but not accounting for this in the quoted coverage).

In the Yellin formalism multiple search regions, for example from the first and second runs of ZEPLIN-III, could naïvely be combined by summing the exposures and finding gaps in the combined set of events. However, if the background rate in one region significantly exceeds that in another, then a straight superposition can easily lead to an upper limit above the lowest of the individually-calculated limits — events from one run can mask large empty patches in another that would otherwise be strong constraints on signal. An alternative test statistic for combining regions, less sensitive to un-modeled backgrounds, can be constructed from the p-values against the signal hypothesis in each region separately. Consider M regions of which the ith contains a fraction e_i of the signal distribution. The signal strength parameter μ is now the sum over all regions; the individual expectations are $\mu_i = \mu\cdot e_i$. We define the test statistic

$$A = \prod_{i=1}^{M} p_i, \tag{5.4}$$

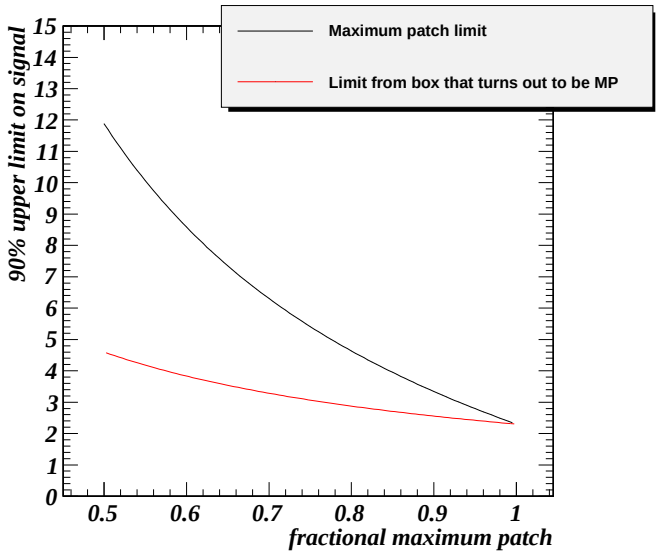

Figure 5.1: Upper limit (in black) derived on expected signal events in 100% acceptance, as a function of the observed maximum patch; shown for comparison (in red) is the limit (again for 100% acceptance) that would be obtained from Poisson statistics had the same empty patch been specified in advance as the search box.

where

$$p_i = 1 - G\left(g_i\right). \tag{5.5}$$

Although numerically the same as Fisherian p-values against the test hypothesis, they are not being used in that sense — rather, their product is used as a statistic to construct (Neyman-Pearson) acceptance regions. In the limit of a large hypothesised μ, each p_i would be distributed uniformly between 0 and 1. However in practice the distribution of each p_i does not extend down to 0, because of the possibility of zero events being observed. An empty observed region gives the (minimum possible) value of $p_i = e^{-\mu_i}$ and happens with probability $e^{-\mu_i}$. Taking this probability mass into account, the distribution of each p_i becomes:

$$p_i \sim z_i \cdot \delta\left(p_i - z_i\right) + \Theta(p_i - z_i), \tag{5.6}$$

where $z_i \equiv e^{-\mu_i}$ denotes the probability of zero events being observed in the ith bin. For the $M = 2$ case, and labeling such that $\mu_1 > \mu_2$, the distribution of A can be calculated analytically:

$$\Pr\left(A < a\right) = \begin{cases} 0 & a < z_1 z_2 \\ z_1 z_2 + z_1\left(\frac{a}{z_1} - z_2\right) + z_2\left(\frac{a}{z_2} - z_1\right) + \int_{z_1}^{a/z_2}\left(\frac{a}{x} - z_2\right)\mathrm{d}x & z_1 z_2 < a < z_1 \\ z_1 + z_2\left(\frac{a}{z_2} - z_1\right) + \left(a - z_1\right)\left(1 - z_2\right) + \int_{a}^{a/z_2}\left(\frac{a}{x} - z_2\right)\mathrm{d}x & z_1 < a < z_2 \\ a - a\log\left(a\right) & z_2 < a \end{cases}$$

$$\Pr(A < a) = \begin{cases} 0 & a < z_1 z_2 \\ a - a\log\left(\frac{z_1 z_2}{a}\right) & z_1 z_2 < a < z_1 \\ a - a\log(z_2) & z_1 < a < z_2 \\ a - a\log(a) & z_2 < a \end{cases} . \tag{5.7}$$

Fig. 5.2 illustrates the contributions to $\Pr(A < a)$ in the p_2–p_1 plane, when the p_is are distributed as in equation 5.6. Four cases of a are shown, all for a signal hypothesis under which region 1 has a 10% chance of being empty and region 2 has a 30% chance of being empty. We note, for example, that $\Pr(A < 0.02) = 0$ (top left case), and that $(A < 0.2)$ is satisfied by all observations in which region 1 is empty (bottom left case).

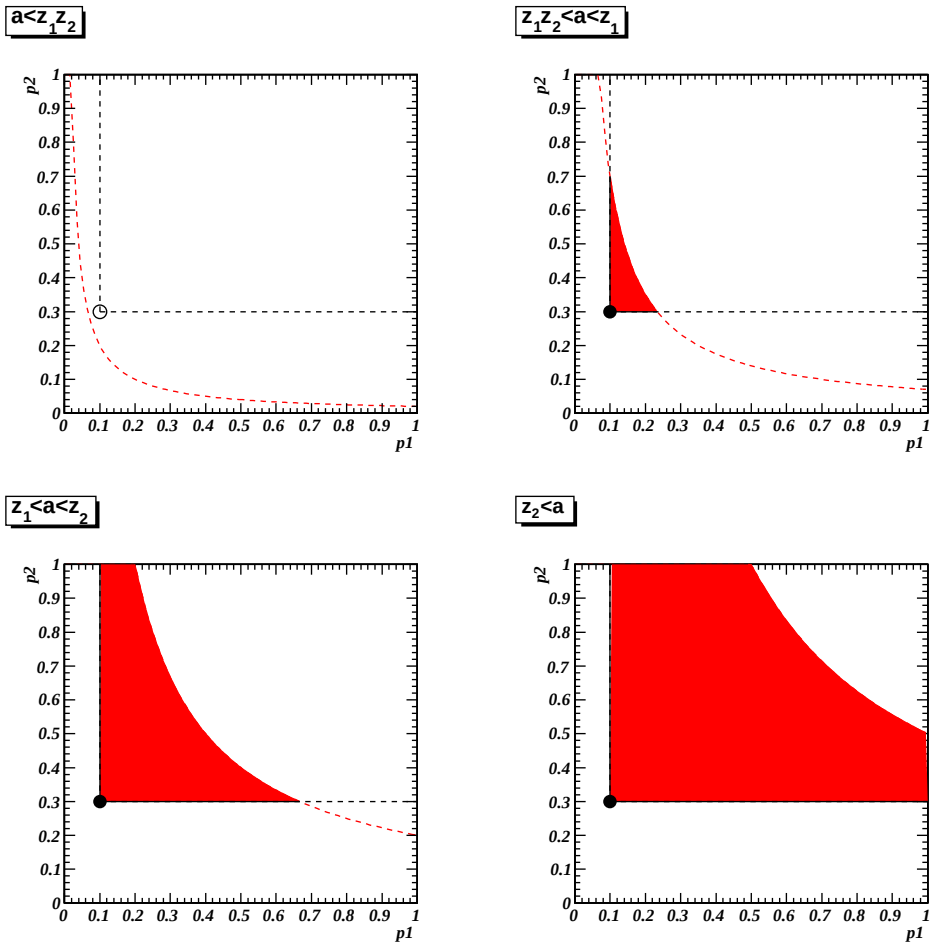

Figure 5.2: Illustrations of $\Pr(p_1 \cdot p_2 < a)$ for the four cases of a, with $z_1 = 0.1$, $z_2 = 0.3$ throughout. Black lines show where one or other experiment has 0 events (solid where included in $\Pr(A < a)$); black circles show where both experiments have 0 events (solid if included in $\Pr(A < a)$).

A rate of signal is rejected as too high at $100(1-\alpha)\%$ confidence level when $\Pr(A < a_{\text{obs}})$ under the signal-only hypothesis is below α. In this way, conservative single-sided limits can be computed from multiple ZEPLIN-III observations. Although empty regions are

not destroyed by the addition of new data, combining regions via A can still give a higher limit than one of the regions alone. One can gain discovery potential, improved power, and more exact coverage by using estimates of background in the search region, to which we now turn.

5.3 Estimating background in the search box

Single-vertex elastic scatters by neutrons which do not also scatter in the veto happen at an extremely low rate but are indistinguishable from signal events. Radiological measurements along with simulations provide an estimate of this irreducible background: 0.06±0.01 events within the SSR search box [98]. The total event rate of ZEPLIN-III in the SSR was dominated by γ-rays emitted in the decay chains of U and Th as well as ^{40}K β-decay, as described in section 4.3.2. The cuts described in chapter 4 result in a search box that excludes the great majority of electron recoils, but still has some finite background expectation. Three approaches to estimating the residual background are described here: the sample of prompt-tagged vetoed events, calibration with 662 keV γ-rays, and extrapolation of the distribution of search events above the box in $\log_{10}(\mathrm{S2}/\mathrm{S1})$.

5.3.1 Vetoed events

Due to the extremely low cross section for interaction, a single WIMP will in practice never scatter in both the veto and the xenon target; in contrast, events due to γ-ray and neutron backgrounds will often be accompanied by near-coincident energy deposit(s) in the veto. As discussed in section 4.3.1, cutting ZEPLIN-III events coincident with veto events is a powerful way to reduce background, with inefficiency for signal only in proportion to the rate of accidental coincidence. If the veto cut is independent of other cuts then the veto-tagged events may also be used to characterise background distributions and rates in an effectively signal-free sample. First, the statistical analysis of vetoed samples in low-background counting experiments is considered. An exact p-value method is applied to the case of equal veto efficiency for all backgrounds, showing how veto efficiency and target exposure may be traded off when designing for some level of discovery significance. In two-phase xenon experiments there are typically two backgrounds in the search region with different veto efficiency, requiring a modified method. Finally, a reason why vetoed events are possibly not a good sideband for background in the SSR search region is discussed.

Consider a low-background search experiment where all background populations are vetoed with the same known efficiency, v. Vetoed events are a signal-free sample which, under the null hypothesis H_0, contains an expected number of events $v/(1-v)$ times that of the unvetoed search data. Let the observed search region contain n_u un-tagged events and n_t tagged events. Following Pearson and Clopper ([146], via [147]), a p-value against

H_0 is given by the binomial probability for n_u being so high given H_0 and $n_{tot} = n_u + n_t$:

$$p_{Bi} = \sum_{j=n_u}^{n_{tot}} \Pr(j; n_{tot}, 1-v) = \frac{B(1-v, n_u, 1+n_t)}{B(n_u, 1+n_t)}, \tag{5.8}$$

where the numerator and denominator are respectively incomplete and complete beta functions. The number of untagged events that would be needed for a 3-σ discovery (p-value of 1.35×10^{-3}) is plotted in Fig. 5.3 for low-background cases of 0, 1 or 2 vetoed events. The limiting value at unity veto efficiency is one event — this effectively corresponds to an experiment with no background. It may not be a veto detector that constrains the

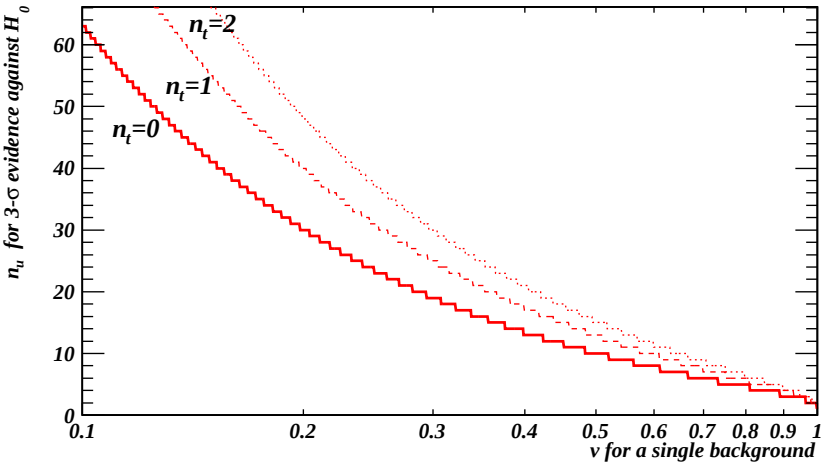

Figure 5.3: The plot shows how many untagged events are required to achieve 3-σ evidence against the zero-signal hypothesis as a function of veto efficiency (assumed here to be the same for all backgrounds).

background rate: any sideband with relative exposure τ can be treated analogously by replacing v with $(\tau/1+\tau)$. When different backgrounds are vetoed with different efficiencies, the zero-signal hypothesis does not imply an exact binomial probability for tagging events; however one can conservatively adopt the lowest of the tagging efficiencies.

The principal disadvantage of using vetoed data to estimate ZEPLIN-III background is that they were open throughout the analysis and were used to develop the cuts described in section 4.4. There is a risk of underestimating the true background rate because of this tuning. The observation of 0 vetoed and 8 un-vetoed events in the search box, given (28.1 ± 0.2)% probability to tag γ-ray events, is suggestive but not conclusive evidence that this did in fact happen.

5.3.2 Calibration with γ-rays

Another estimate of the electron-recoil background in the search box comes from data taken in the presence of a ^{137}Cs γ-ray source (see Appendix A.2 for the decay properties). Two datasets, totaling 17.5 hours, were acquired with a 5 kBq source on top of the vacuum dome within 40 mm of the central axis of the array. Triggering and data acquisition were configured as for search data. The hardware trigger rate was 140 s^{-1}, compared with 0.4 s^{-1} in WIMP search mode; although necessary to acquire sufficient data in a practical amount of time, this increased rate led to systematic differences in detector response.

One known source of rate dependence in PMT response is the charging of the cold photocathodes [148]. Photocurrent, drawn from its outer conductive ring, causes a voltage drop across the photocathode that is proportional to its surface resistivity; the latter increases dramatically for bialkali photocathodes upon cooling. The resulting distortion of field lines lowers the efficiency for collecting photoelectrons on the first dynode, and hence the PMT response. The fractional effect is greater for secondaries than for the much smaller primary scintillation signals, and hence the mean ratio S2/S1 can be reduced when average photocurrent is high. In spite of conductive fingers which had been deposited on the PMT windows to mitigate this effect, the mean S2/S1 was indeed found to be lower in ^{137}Cs data than in search data: by $(6 \pm 2)\%$ in the range $2 < (\mathrm{S1/keVee}) < 12$.

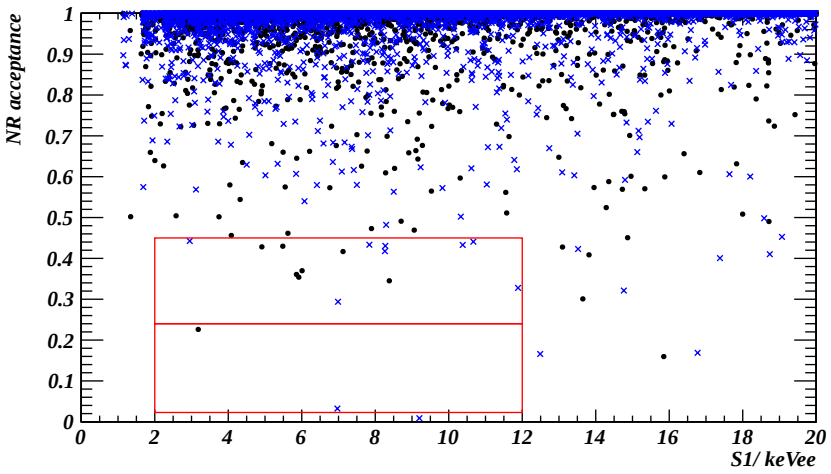

Figure 5.4: ^{137}Cs (blue cross) and search data (black dot) events in signal acceptance $a(\mathrm{S2/S1; S1})$ versus S1. The partition of the search box is discussed in section 5.5.1. A WIMP signal would be uniformly distributed along the vertical axis.

The ^{137}Cs data shown in Fig. 5.4 contain 9 events within the search box, the small cathode charging correction having reduced this number by 1. The relative exposure of the calibration, estimated as the total number of events at all $\log_{10}(\mathrm{S2/S1})$ in ^{137}Cs data divided by that in the search data, is $\tau = 0.96 \pm 0.02$. Therefore, these data predict 9 ± 3 γ-ray events in the search box.

5.3.3 Extrapolation of fitted distribution

A quite different approach to estimating electron-recoil background is to model its density in S1 and S2 via an analytic function with a handful of fit parameters. These parameters replace the background expectations in individual bins as nuisance parameters, which can reduce the computational demands for finely-binned data. The physically-motivated expectation of distributions which are continuous in S1 and S2 is also used rather than ignored. However a totally correct result depends on choosing a family of functions which includes the real distribution of background, and this can be hard to demonstrate.

Empirically, the distribution of $\log_{10}(\mathrm{S2}/\mathrm{S1})$ for electron recoils is consistent with skew-gaussian at fixed S1, as observed in the FSR [48]. This was true in the SSR, within errors, for ^{137}Cs calibration data, for vetoed events and for surviving search events. The skew-gaussian PDF is:

$$f_S\left(x; x_c, \omega, \alpha\right) = \frac{1}{\omega\sqrt{2\pi}} \mathrm{e}^{-\frac{(x-x_c)^2}{2\omega^2}} \left[1 + \mathrm{erf}\left(\frac{\alpha}{\sqrt{2}} \frac{x - x_c}{\omega}\right)\right], \tag{5.9}$$

with $x = \log_{10}(\mathrm{S2}/\mathrm{S1})$ and the parameters x_c, ω and α setting respectively the position, width and skewness. Fits of this form were made to un-vetoed search data above the box in order to estimate the electron recoil background within it: good fits were obtained for each 2-keVee bin of S1 and the sum of expectations in the box, 6.5 ± 2.4 events, is consistent with the Cs-137 calibration. The statistical treatment of this and other background estimates to derive two-sided intervals is developed in the following section.

5.4 Statistics based on likelihood

The test statistic that maximises power against H_0 under an alternative hypothesis H_1 at fixed significance is the ratio of likelihoods $\mathcal{L}\left(\mathrm{H}_0\right)/\mathcal{L}\left(\mathrm{H}_1\right)$ [149]. If hypotheses are entirely specified by a signal strength parameter then the Feldman-Cousins acceptance region [150] for a signal level μ_{test} may be built up by including observations in decreasing order of

$$R_{\mathrm{FC}} = \frac{\mathcal{L}\left(\mu_{\mathrm{test}}\right)}{\mathcal{L}\left(\hat{\mu}\right)}, \tag{5.10}$$

where $\mathcal{L}\left(\hat{\mu}\right)$ is the maximum likelihood (ML) of the observation under any signal level.

Significant background uncertainty requires two modifications be made to the Feldman-Cousins formalism. Firstly, the model now includes nuisance parameters to describe the background, denoted by $\boldsymbol{\beta}$: the elements of $\boldsymbol{\beta}$ may be the expectations in individual bins or the parameters of a density function. Secondly, the observation, and hence the likelihood, includes subsidiary measurements of background. The ratio of the likelihood maximised over $\boldsymbol{\beta}$ *given* μ_{test} to the likelihood maximised over $\boldsymbol{\beta}$ *and* μ, known as the profile likelihood

ratio, is then used to construct acceptance regions [151]:

$$R_{\mathrm{PLR}} = \frac{\mathcal{L}\left(\mu_{\mathrm{test}}, \hat{\hat{\beta}}\right)}{\mathcal{L}\left(\hat{\mu}, \hat{\beta}\right)}, \tag{5.11}$$

where the double-hat denotes the constrained maximum-likelihood value at fixed $\mu = \mu_{\mathrm{test}}$. As discussed in section 5.1, it is sufficient to confirm coverage for signal levels in conjunction with their constrained ML backgrounds, rather than for all possible values of μ and β.

5.5 SSR limit calculation

The final SSR result was derived from a coarsely binned profile-likelihood-ratio (PLR) analysis using extrapolation and ^{137}Cs calibration as auxiliary measurements of electron-recoil background. The neutron background estimate and uncertainty from simulation is included, but does not significantly affect the outcome. This analysis did not rely on detailed knowledge of the background PDF but made some use of its non-uniformity and produced two-sided CIs. In this section the choice of binning, the likelihood ratio function and its distribution are explained in detail, then the SSR results are presented.

5.5.1 Subdivision of the search box

Based on the rest of the search data, the background density within the search box is expected to decrease steeply — in particular, more steeply than the signal density — from high to low $\log_{10}(\mathrm{S2}/\mathrm{S1})$. To exploit this information the box was divided in two along a line of signal acceptance in $\log_{10}(\mathrm{S2}/\mathrm{S1})$, with the division line set by a sensitivity study. Skew-gaussian binned ML fits were made to events above the search region in six 2-keVee slices of S1, with the resulting summed estimates of background shown in Fig. 5.5. For simplicity a two-bin Feldman-Cousins technique was used to derive simulated limits, i.e. the background estimates were treated as exact; while not a correct model, this was deemed adequate for deciding on binning, and does not affect the coverage of the eventual PLR result.

The distribution of F-C limits generated under different binning indicates a shallow minimum in mean sensitivity around a 24% division, as seen in Fig. 5.5; this was therefore adopted. The events within the box were not considered in this decision, so as to maintain coverage; the observed limit happens to be a slight upward fluctuation from the mean signal-free limit.

5.5.2 PLR test statistic

The likelihood function entering into equation 5.11 for the SSR result consist of three parts: approximated Gaussian likelihoods for the electron recoil and neutron background rates in the ith bin, $b_{i,\mathrm{SG}}$ and $b_{i,\mathrm{n}}$, given the uncertainties with which they are estimated; Poisson

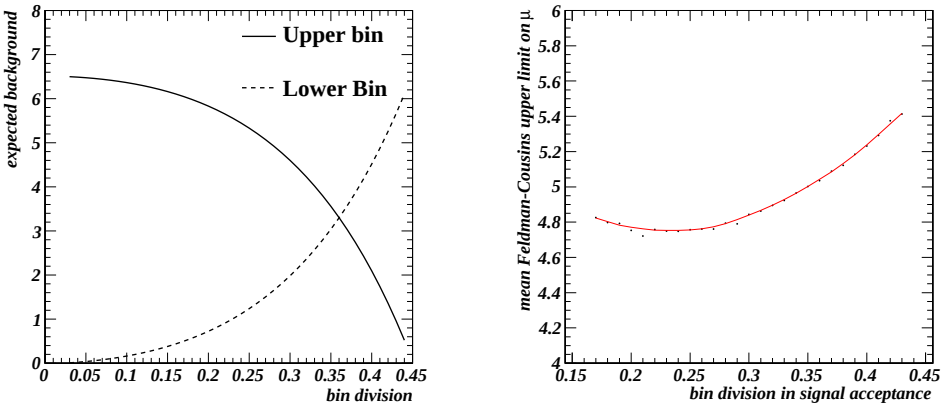

Figure 5.5: Left: background estimates based for a 2-bin search box, based on data outside of the box. These were used to generate expected limits as a function of binning. Right: sensitivity versus search box division from a 2-bin Feldman-Cousins analysis. A division at 24% was adopted for the final, profile-likelihood ratio analysis.

likelihoods for the observed numbers of γ-calibration events, m_i, in a relative exposure of τ given $\beta_{i,\gamma}$; and Poisson likelihoods for the observed numbers of search events, n_i, given the background, $\beta_i = \beta_{i,\gamma} + \beta_{i,\mathrm{n}}$, and signal, μ_i, expected:

$$\begin{aligned}\mathcal{L}(\mu,\boldsymbol{\beta}) &= \mathcal{L}_{\mathrm{box}}(\mu,\boldsymbol{\beta})\cdot\mathcal{L}_{\mathrm{Cs}}(\boldsymbol{\beta})\cdot\mathcal{L}_{\mathrm{extrap.}}(\boldsymbol{\beta}) \\ &= \prod_{i=1}^{2}\Pr(n_i;\mu_i+\beta_i)\cdot\Pr(m_i;\tau\beta_{i,\gamma}) \\ &\quad\times\frac{1}{\sigma_i\sqrt{2\pi}}\exp\left(\frac{\left(\beta_i-\left(b_{i,\mathrm{SG}}+b_{i,\mathrm{n}}\right)\right)^2}{2\sigma_i^2}\right). \end{aligned} \tag{5.12}$$

It is convenient to work in terms of $q \equiv -2\log R_{\mathrm{PLR}}$, which has a χ^2 distribution for asymptotically large samples [152]. In the case of the SSR, the distribution was obtained by Monte Carlo: q was computed for random values of n_i, m_i and $b_{i,\mathrm{SG}}$ generated from the likelihood function in equation 5.12. A signal rate is thus rejected when the observed q is above the 90th percentile of simulated trials assuming μ_{test} and its constrained ML background, $\boldsymbol{\beta}$. Sampling $b_{i,\mathrm{SG}}$ from Gaussian distributions implies a finite probability for un-physical negative values (around 0.3% in this case, as the background expectation is 2.7 standard deviations above zero); these trials were discarded. Also, the maximum-likelihood values of model parameters were constrained to their physically-allowed regions for both observed and Monte Carlo evaluations of 5.11.

5.5.3 WIMP-nucleon cross section results

The zero-signal hypothesis is not excluded and, as shown in Fig. 5.6, the observed value of q exceeds the 90th percentile of MC trials for $\mu > 5.1$. Therefore a 90% CL confidence interval is set at 0–5.1 events in the search box. The signal models derived in section 4.2

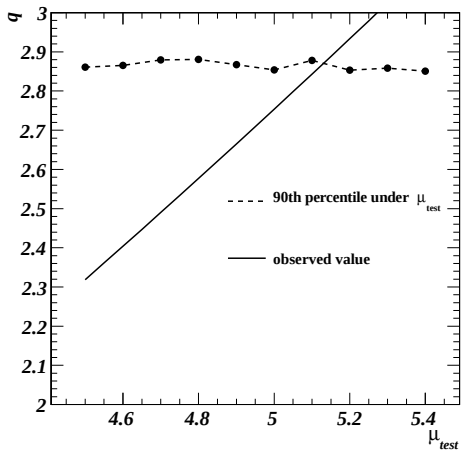

Figure 5.6: Observed and critical values of the PLR statistic q for the SSR search box, versus expected signal.

then imply limits on spin-independent WIMP-nucleon cross section and the spin-dependent WIMP-neutron cross section as shown in figures 5.7 and 5.8.

The first science run was re-analysed with the current analysis codes and updated $\mathcal{L}_{\text{eff}}$ estimate. The re-analysed FSR search box contained 5 events; 7.9 ± 4.8 are expected from background sources, including 1.2 ± 0.6 neutron events. A binned profile likelihood ratio limit was set just as for the SSR, but omitting the ^{137}Cs calibration as it over-estimated leakage into the box by a factor of ~5. The discrepancy may be due to rate-dependent photocathode charging effects [153]. The two search regions were combined via a four-bin PLR calculation to set a 90% confidence interval of 0–6.0 signal events in total.

In conclusion, ZEPLIN-III achieved, in the region $m_{\text{W}} = 15$–$70\,\text{GeV}\,\text{c}^{-2}$, the second most stringent direct detection limit on spin-independent WIMP-nucleon cross section. The limits on WIMP-neutron spin-dependent cross section for $m_{\text{W}} > 10\,\text{GeV}\,\text{c}^{-2}$ are similarly competitive. It is one of a small group of experiments which together disfavour a canonical WIMP interpretation of the DAMA modulation and the excess of events at low energy in CRESST-II, and begin to encroach on allowed regions of constrained supersymmetric extensions of the Standard Model.

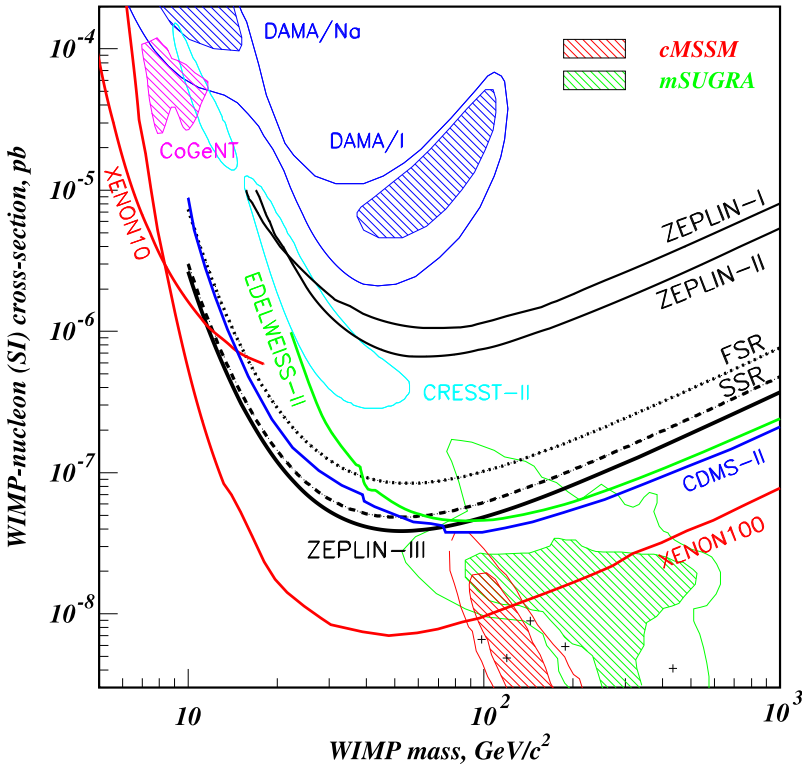

Figure 5.7: Limits on the spin-independent WIMP-nucleon cross section, at 90% CL, from ZEPLIN-I [139], ZEPLIN-II [49] and ZEPLIN-III (FSR, SSR and combined); CDMS-II [40]; EDELWEISS-II [42]; XENON10 (S2-only) [33]; and XENON100 (100.9 live days) [51]. Dark blue regions are 3- and 5-σ explanations of the 2008 DAMA/LIBRA results [154]. The magenta region is a 90%-CL explanation of the CoGeNT result [110]. The cyan regions are 2-σ explanations for CRESST-II result [43]. 68% and 95% favoured regions from a 2008 Bayesian analysis [155] of mSUGRA with flat priors and a 2011 frequentist analysis [52] of the very constrained MSSM [156] in light of early LHC data are also shown, respectively in green and red. From [78].

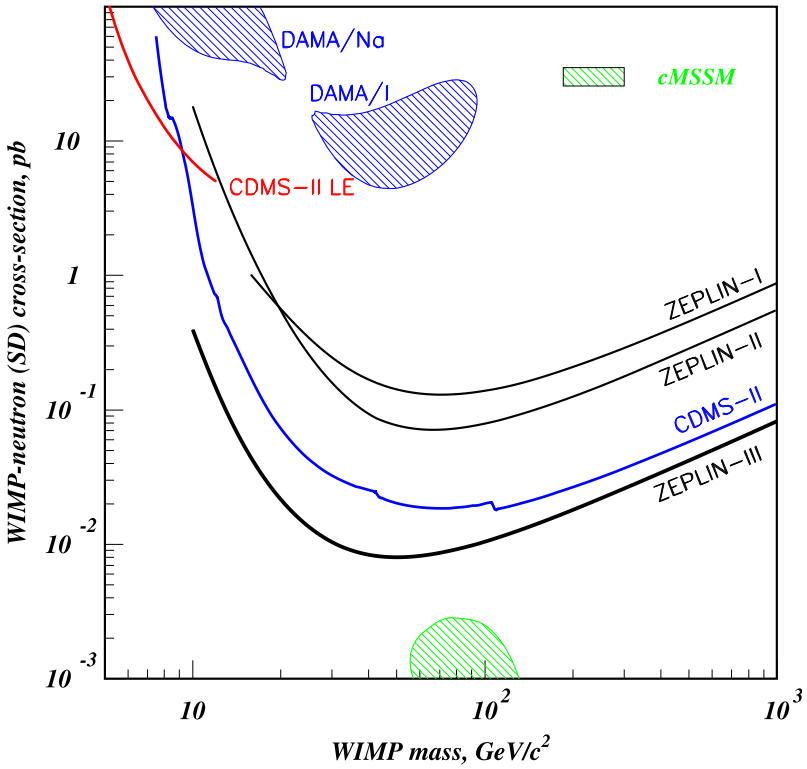

Figure 5.8: Limits on the spin-dependent WIMP-neutron cross section from ZEPLIN-I [139], ZEPLIN-II [49] and ZEPLIN-III (combinined FSR and SSR); CDMS-II [40], including a dedicated low-energy analysis [157]. Dark blue areas are 3-σ explanations of the 2008 DAMA/LIBRA results [154]. The green area is the tip of the 95% credible region in a 2007 Bayesian analysis of the constrained MSSM [158]. The XENON100 collaboration have not yet presented spin-dependent limits; they are expected, based on the 100-live-day data [51] to be a few times lower than the ZEPLIN-III curve. From [78]

.

Chapter 6

Limits on inelastic dark matter from the first science run

The DAMA/NaI and DAMA/LIBRA modulation result — neither confirmed as dark matter signal nor explained as background after 13 annual cycles — is in considerable tension with other direct detection results in the canonical WIMP scattering paradigm of chapter 4. The present chapter presents search results, using FSR data, for one of several alternative models of dark matter which may be able to account simultaneously for the modulation and other experiments' limits on total rate: inelastic dark matter. I led the analysis by the Imperial College group, building on the main FSR elastic search analysis [48], and wrote the paper describing the work, which was published in 2010 [79].

6.1 The 2010 DAMA results

DAMA/NaI and its successor DAMA/LIBRA are direct detection experiments which search for dark matter interactions via the scintillation signal in thallium-doped sodium iodide crystals. There is no discrimination between energy deposited via electron or nuclear recoils, but target masses are large at 87 and 242 kg, respectively, and their NaI(Tl) crystals have the best radiopurity of any used in the field of direct detection. Event energies are reported in electron-equivalent energy, an amount of scintillation light normalised to electron-recoil calibration; nuclear recoil events with recoil energy E will be reconstructed at a quenched energy $E_{\mathrm{ee}} = qE$, with q equal to 0.08 ± 0.02 for iodine recoils and 0.25 ± 0.03 for sodium [159]. Light yields in the individual detector modules have a range of 5.5 to 7.5 photoelectrons per keVee. The background rate in the DAMA/LIBRA detector is shown in Fig. 6.1; away from noise, in the region above 2 keVee, it is of order 1 count per keVee·kg·day with a slight peak due to X-rays and Auger electrons around 3.0 keVee. This corresponds to the K-shell binding energy of ^{40}Ar, which is produced by ^{40}K electron-capture in the crystals. [160].

The DAMA experiments search for the annual modulation signature of dark matter scattering by analysing the time-dependence of residual rates, i.e. event rates minus their long-term means. An annual modulation is observed for single-hit events (see Fig. 6.2)

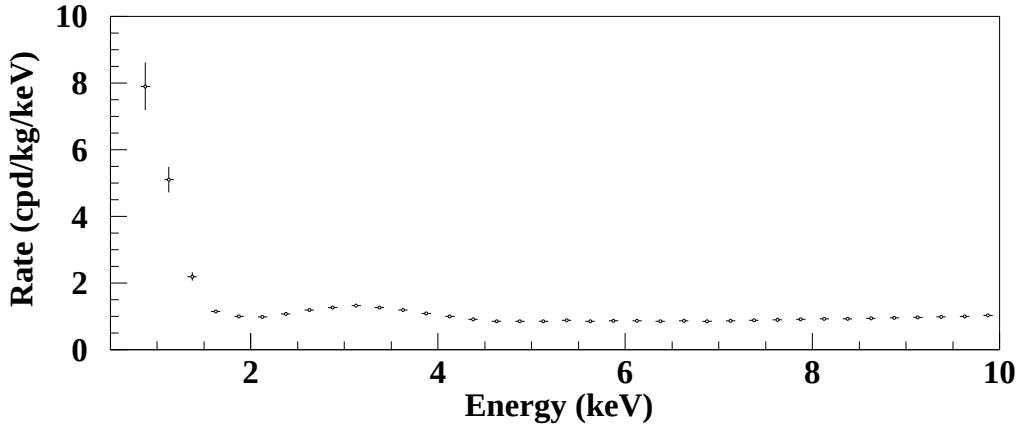

Figure 6.1: The spectrum, in electron-equivalent energy, of single-hit scintillation events within the DAMA/LIBRA modules in search mode (from [160], using 0.53 tonne·yr of data).

and 8.9-σ evidence is reported against the no-modulation hypothesis. In Ref. [58], the modulation in each 0.5-keVee energy bin is parametrised as $S_m \cos\omega\,(t - t_0)$, fixing t_0 = 152.5 d and $\omega = 2\pi$ yr^{-1} as expected for dark matter scattering. A maximum-likelihood fit is made to obtain the modulation amplitude S_m as a function of energy, which is shown in Fig. 6.4.

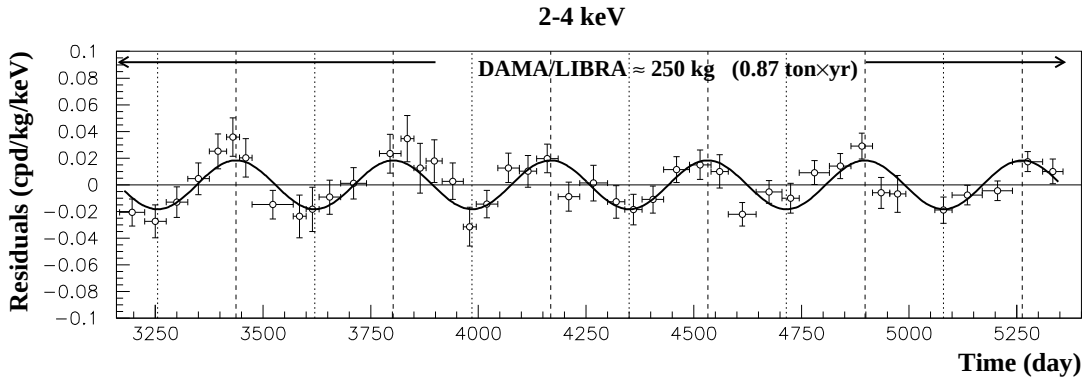

Figure 6.2: Time variation of the residual rate of single-hit scintillation events measured in DAMA/LIBRA (from [58]). The zero of the day axis is January 1st 1995, and the keV unit refers to electron-equivalent energy.

6.2 Inelastic dark matter

Several direct detection experiments have ruled out the interpretation of the DAMA modulation as nuclear recoils due to the elastic scattering of WIMPs in a Maxwellian halo, e.g. [40],[50],[49]. A model which can simultaneously account for the positive and null results is so-called inelastic dark matter [161] (IDM). It consists of WIMPs which scatter pre-

dominantly into a near-degenerate higher-mass state. Theoretical models leading to this phenomenology in the context of warped extra dimensions and extensions of the MSSM are presented in [162].

Scattering with energy transfer E due to a WIMP of ground state mass m_χ and mass splitting δ requires a minimum relative speed (cf. equation 4.2) of

$$v_{\rm min} = \frac{1}{\sqrt{2m_{\rm N}E}}\left(\frac{m_{\rm N}E}{\mu_{\rm N}} + \delta\right), \tag{6.1}$$

where $m_{\rm N}$ is the nucleus mass and $\mu_{\rm N}$ is the reduced mass of the WIMP-nucleus system. A non-zero δ results in a recoil spectrum that goes to zero at low energy and can naturally reproduce the shape of the modulation spectrum in DAMA. The rate is also more sensitive to the upper tail of the WIMP velocity distribution which, as shown in Fig. 4.1, leads to a larger fractional modulation. For fine-tuned parameter values, the kinematic threshold could in fact *only* be reached by DM particles in the summer, giving total modulation. In general, the region of IDM parameter space which can resolve the tensions between experiments is that for which the fractional modulation of signal is large and, especially, where scattering is kinematically accessible to DAMA but not to germanium targets which see no significant signal. This occurs if the DAMA rate modulation is due to scattering from iodine nuclei.

WIMPs slower than $(2\delta/\mu_{\rm N})^{0.5}$ in the lab frame will not scatter inelastically at all and so, for a given local escape velocity, more m_χ-δ parameter space is accessible to heavier target nuclei. Germanium and xenon target nuclei are compared in Fig. 6.3, which shows how DAMA-explaining dark matter can be undetectable to low-A targets. Conversely, systematic uncertainty in the expected relative rates in different targets due to nuclear form factors and WIMP velocity distributions grows with the difference in atomic mass [163] — a factor of two discrepancy between halo models can arise for the relative scattering rates of tungsten and iodine targets. ZEPLIN-III is well suited to test IDM models that would, by predicting a modulated rate of scattering against iodine nuclei, explain the DAMA observation. The scintillation yields from iodine and xenon nuclear recoils mean that a peak around 3 keVee in DAMA (the most modulated bin is 2.5–3 keVee) corresponds to ~40 keV recoil energy or S1 ~18 keVee in ZEPLIN-III. The compatibility of the two observations is now investigated, beginning with full IDM signal models for the two experiments.

6.3 Signal models

The same models of the DM halo, WIMP-nucleon scalar interactions and nuclear form factors as in section 4.2 were used in the inelastic case, with appropriate substitution of $v_{\rm min}$, to predict differential recoil spectra $\frac{{\rm d}R}{{\rm d}E}$ for sodium, iodine and xenon. Detector response models were then applied for comparison with the modulated and absolute rate observed in DAMA and ZEPLIN-III, respectively.

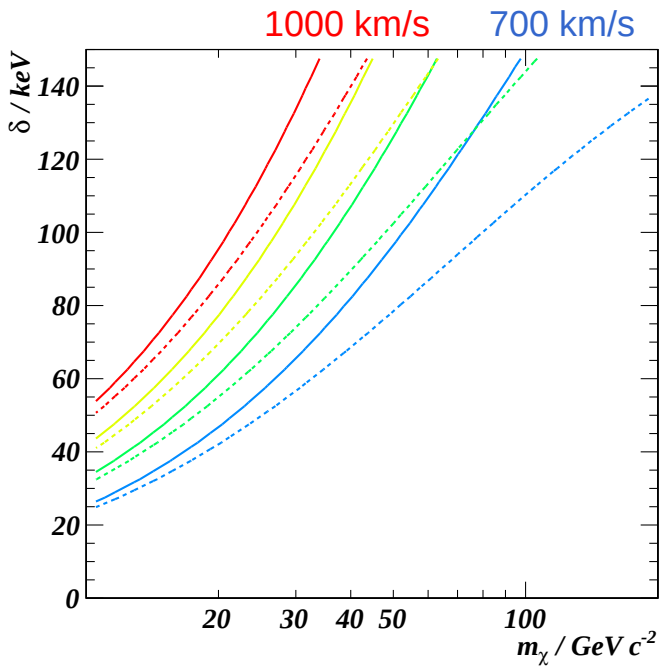

Figure 6.3: The mass splitting δ, as a function of mass, above which inelastic dark matter is kinematically inaccessible for four WIMP-nucleus speeds (700, 800, 900 and 1000 km/s). Dashed and solid lines represent Ge (A = 73) and Xe (A = 131) targets respectively. Values for an iodine target would lie marginally below those for xenon. The maximum relative speed between a terrestrial detector and WIMPs in the halo is approximately 800 km/s, the sum of the Earth's velocity and the local escape velocity.

6.3.1 DAMA/NaI and DAMA/LIBRA

The light yield of nuclear recoil events is accounted for by the quenching factors described above. Gaussian resolution models in electron-equivalent energy for DAMA/NaI and DAMA/LIBRA, as reported in [160] and parametrised in [164], were applied and summed in proportion to the experiments' exposures. The expected value of S_m in each bin of electron-equivalent energy is then compared with observation, as in Fig. 6.4. The $\Delta\chi^2$ statistic for this fit in comparison with the global minimum χ^2 was used to define confidence regions in the parameter space of m_χ, δ, and scalar WIMP-nucleon cross section σ_n. Figure 6.5 shows part of the area in m_χ-δ for which some σ_n is within the 99% confidence region, i.e. $\Delta\chi^2 < 11.3$ for 3 fitted parameters.

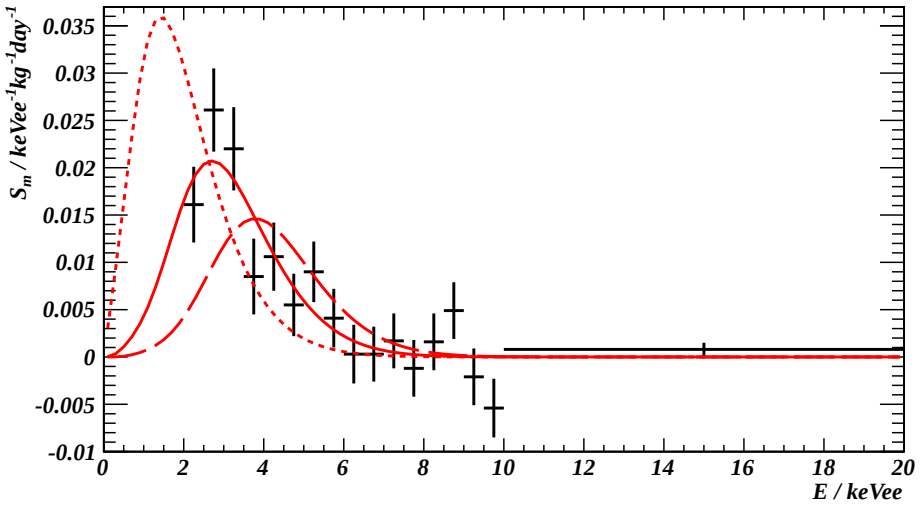

Figure 6.4: The fitted value of the modulation amplitude parameter S_m in bins of visible energy from combined DAMA data [58]. Bins above 10 keVee, in which the fitted value of S_m fluctuate statistically around zero, are amalgamated. The modulation predicted by IDM scattering with the best-fit cross sections at example points in δ-m_χ are also shown: the global best fit (65 GeV c^{-2}, 124 keV) with a solid line; where δ is too low (50 GeV c^{-2}, 80 keV) with a dotted line; and where δ is too high (80 GeV c^{-2}, 150 keV) with a dashed line.

6.3.2 ZEPLIN-III FSR

The response of the detector in its FSR configuration is qualitatively the same as that during the SSR, described in chapter 3. However, in the FSR three times the number of photoelectrons per keVee in S2 were produced; this is due to higher applied field — 17 rather than 15 kV from anode to cathode — a slightly larger gas gap and higher mean quantum efficiency. The number of photoelectrons per keVee in S1 was also higher, by around 40%.

A search region of 20–80 keV recoil energy was chosen for the IDM analysis to include

the majority of events predicted by the quenched, inelastic WIMP-iodine scattering interpretation of the DAMA result. The corresponding range in S1 was calculated using the FSR measurement of $\mathcal{L}_{\rm eff}$ as 8.4–38.3 keVee, well above threshold. Energy resolution (for nuclear recoils) is modeled as Gaussian with $\sigma/\mathrm{keV} = 1.5\sqrt{E/\mathrm{keV}}$, dominated by statistical fluctuations in the number and size of the photoelectrons of the S1 signal.

Most of the DAMA-allowed parameter space predicts event rates in ZEPLIN-III far higher than the observed in the initial FSR search region. However, in the region of high fractional modulation the absolute rate in xenon can drop below 0.1 events per kg.day, as shown in Fig. 6.5. A dedicated analysis of the FSR data, extending to higher energies than the search for elastic WIMP-nucleus scattering, was developed to explore these remaining regions.

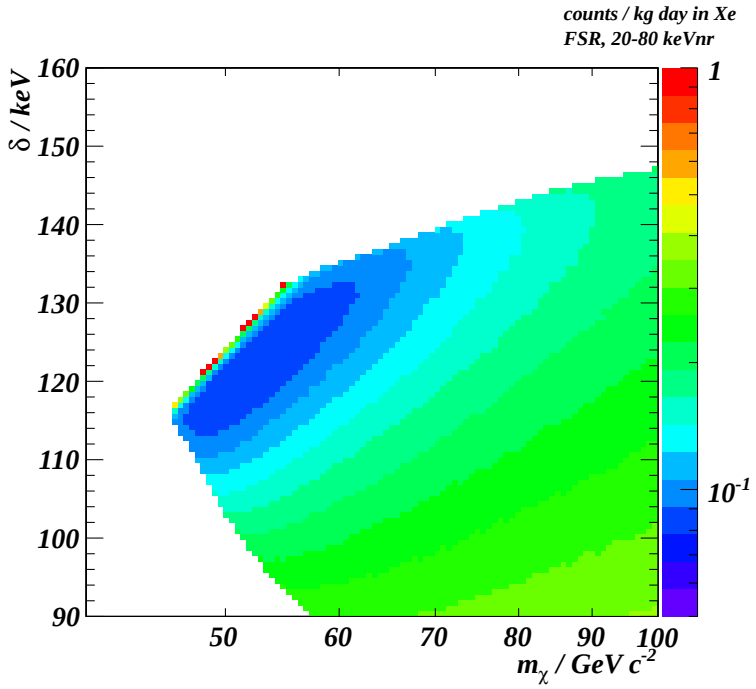

Figure 6.5: The scattering rate expected in a xenon target at the lowest cross section consistent, at 99% confidence level, with explaining the observed DAMA/LIBRA modulation spectrum. White regions indicate that no cross section gives an adequate fit to DAMA: in the top-left corner of the plot this is because scattering is kinematically inaccessible (cf. Fig. 6.3); elsewhere it is because the modulation occurs at the wrong energy (cf. Fig. 6.4). We show only the hardest-to-exclude part of the DAMA-allowed δ-m_χ parameter space, which extends to higher m_χ and lower δ.

6.4 Event selection in ZEPLIN-III

Nuclear recoils and electron recoil background events are distinguished, as in the elastic analysis, through the ratio of scintillation and ionisation signal size. The other cuts applied

in the standard, elastic analysis [48], mainly based on quality of reconstruction and aimed at removing MSSI events, were retained. Additional discrimination was achieved here using scintillation pulse shape, which is faster for nuclear recoils than for electron recoils (section 4.4.2). The timing of a primary pulse is described by τ_1, the mean arrival time of pulse area:

$$\tau_1 \equiv \frac{\int t \cdot V(t)\,\mathrm{d}t}{\int V(t)\,\mathrm{d}t} \quad ,$$

for a pulse beginning at $t = 0$, in a summed PMT trace $V(t)$. The relatively large number of primary photoelectrons in the inelastic signal region, together with the 2-ns sampling of the PMT traces, gives sufficient resolution in τ_1 to significantly improve sensitivity to IDM.

Figure 6.6 shows the separation of the two types of calibration data in an example energy bin. The timing of neutron calibration events within each 5-keVee bin from 5 to 40 keVee is well described by gamma distributions in $1/\tau_1$, as observed in ZEPLIN-I [139]. Fitting a polynomial in S1 to the medians of the gamma distributions produces a cut $\tau_1 < \tau_{1,\mathrm{crit}}(\mathrm{S1})$ with 50% signal acceptance. The power of this cut to reduce electron recoil background increases with energy, as seen in Fig. 6.7, mainly due to a narrowing in the τ_1 distribution of electron recoil events.

AmBe calibration data were then used to obtain the S2/S1 distribution of elastic nuclear recoil events which pass the timing cut, as a function of S1. As in section 4.2.2, the $\log_{10}(\mathrm{S2/S1})$ distribution was fitted by a Gaussian in each energy bin. The energy dependence of the Gaussian parameters was then fitted by a power law to define a search box with 47.7% signal acceptance extending 2 standard deviations down from the mean. After efficiencies from dead time, pulse-finding, event reconstruction [48] and the new cuts on S2/S1 and τ_1, the net exposure for signal events is 63 kg · days, with 5% uncertainty due to neutron calibration statistics.

Events surviving all selection cuts are shown in Fig. 6.8. Six events were found in the search box, out of $1.3{\times}10^5$ at all $\log_{10}(\mathrm{S2/S1})$ and τ_1. This is consistent with the 7 out of $8.5{\times}10^4$ events observed in ^{137}Cs γ -ray calibration data, suggesting the candidate events are well explained as the tail of the electron-recoil background population. As noted in section 5.5.3, there was systematically greater leakage to low S2/S1 in ^{137}Cs data compared with search data at the lower energies of interest for elastic WIMP scattering (attributed to photocathode charging effects in some PMTs); this was not apparent in the inelastic search region, perhaps because of the smaller probability for the recorded signals to be dominated by a single tube due to statistical fluctuations. In any case, the nature of the maximum patch statistic makes the cross section limits absolutely independent of background estimates. The simulated neutron background, dominated by uranium and thorium decay chains in the photomultipliers, is only 0.2 events. Without pulse shape discrimination, 27 events would have appeared in the search box for double the effective exposure.

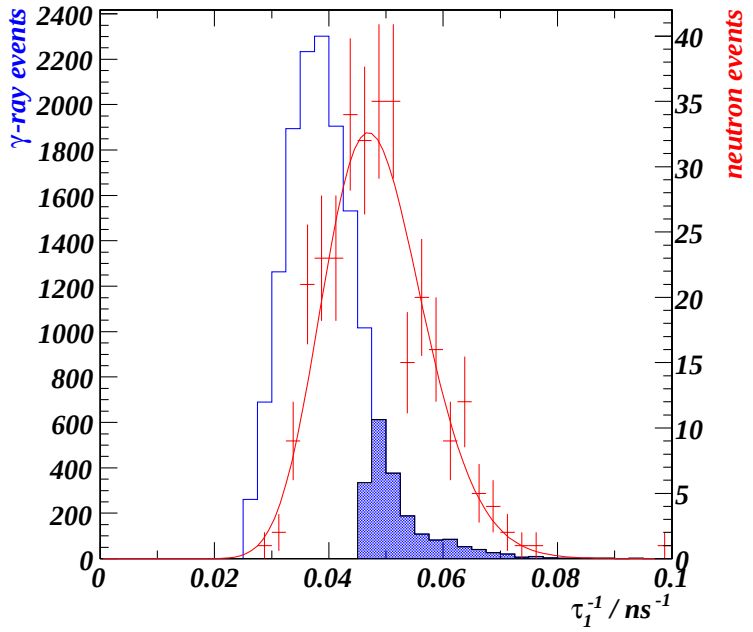

Figure 6.6: Illustration of the τ_1 cut in an an example bin of S1, 20-25 keVee. The points and right-hand scale show single-scatter elastic neutron calibration events, fitted by a gamma distribution. The outlined histogram and left scale show electron-recoil calibration events, with a shaded region corresponding to the 13% of these which pass the cut.

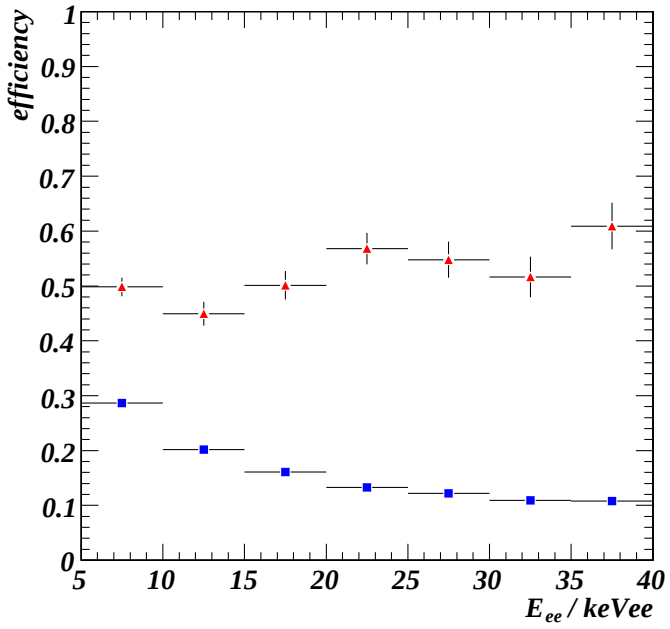

Figure 6.7: The efficiency of the τ_1 cut for electron (blue square) and nuclear recoil (red triangle) calibration events, in bins of S1 from 5 to 40 keVee.

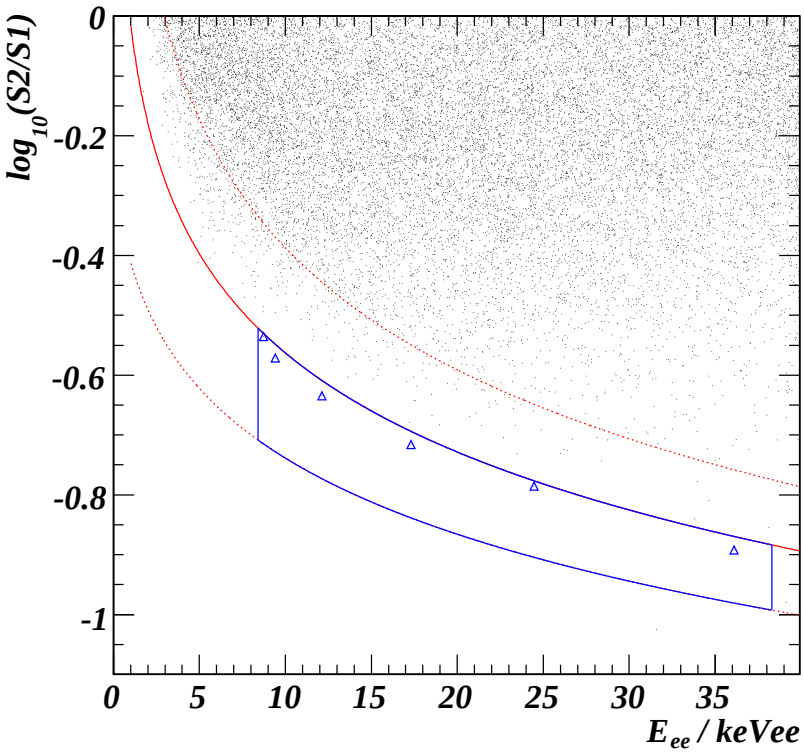

Figure 6.8: Search events passing all cuts except those on S1 and $\log_{10}(\mathrm{S2}/\mathrm{S1})$. The solid curve is the fitted median $\log_{10}(\mathrm{S2}/\mathrm{S1})$ for nuclear recoil events, with dashed lines at 2 standard deviations above and below. The search range in energy is delimited by vertical lines. Six events are found within the search box.

6.5 Results and discussion

The maximum patch statistic (see section 5.2) was used to set single-sided limits on the signal rate within the search box. Events were mapped onto a plane of uniform signal density by integrating the signal spectrum in E and the fitted profile in S2/S1. Under DAMA-explaining signal models, the largest empty rectangle in the re-mapped search box has a fractional acceptance of 0.73–0.75; the precise acceptance mapping depends, via the spectrum shape, on m_χ and δ. This implies a 90% CL limit of 5.8–5.4 expected signal events in the box. The resultant limits on σ_n for $v_{\rm esc} = 550\,{\rm km\,s^{-1}}$ are plotted in Fig. 6.9.

The result may also be presented in terms of the confidence level at which the ZEPLIN-III upper limit excludes IDM models within the 99% CL allowed region from DAMA data. This is illustrated in Fig. 6.10. The minimum confidence level (around 52 GeV c^{-2}, 122 keV) is 87%. This result is insensitive to the uncertainty in $\mathcal{L}_{\rm eff}$ at low energy — there is one event below the nominal range in S1, but it lies high enough in $\log_{10}(\mathrm{S2/S1})$ that it would not alter the fractional maximum patch if included. The 1-σ systematic uncertainty in the net exposure, due to neutron calibration statistics, varies the minimum CL in the range 85–89%. Variation in the local escape velocity within the 90% CI from [134] shifts the point of minimum CL by about 5 GeV c^{-2} and 10 keV in the δ-m_χ plane, but does not change its percentage value. The inelastic explanation of DAMA is not conclusively ruled out, but it is constrained to a very narrow region of parameter space.

At publication, this result was the tightest constraint presented using xenon, a target nucleus whose similarity to iodine minimises uncertainty from the assumed halo model and form factor. The FSR inelastic limit was thus an important corroboration of an analysis of CRESST data [164] which had set 95% exclusion limits on the basis of a tungsten target. In 2011, results from 100 live days of XENON100 [165] definitively ruled out this inelastic dark matter model as an explanation of the DAMA modulation.

One consequence of the minimum scattering velocity of IDM is an additional indirect detection signal from the Sun. Elastically scattering WIMP dark matter with $\sigma_{\rm n} \gtrsim 10^{-47}\ {\rm cm}^2$ which is gravitationally captured by the Sun thermalises and collects in the interior. By contrast, once they are slower than the threshold for scattering from iron, IDM particles will remain in orbit. A significant amount of WIMP co-annihilation is then expected outside of the Sun (the escape velocity at the photosphere surface is approximately 280 km/s), producing electron and positron lines accessible to indirect detection experiments [166]. Recent analysis of FERMI electron and positron data yielded strong limits on IDM but only in the mass range above 70 GeV c^{-2}, due to the effects of geomagetic deflection [167].

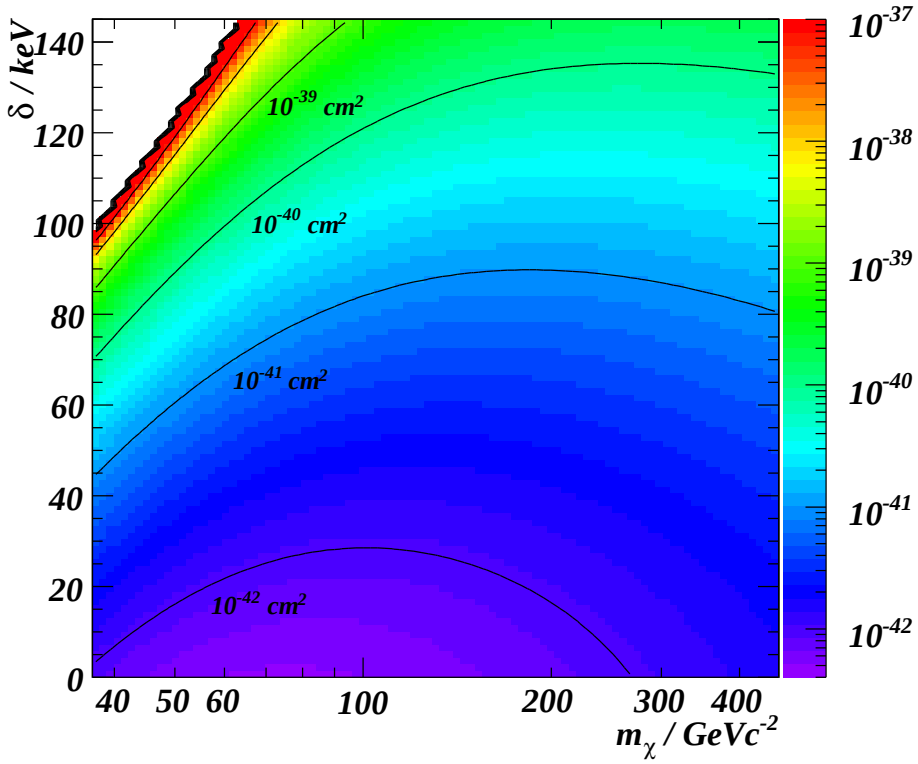

Figure 6.9: 90% limit on σ_n/cm^2 as a function of mass and splitting. The upper left region predicts no inelastic scattering during the run.

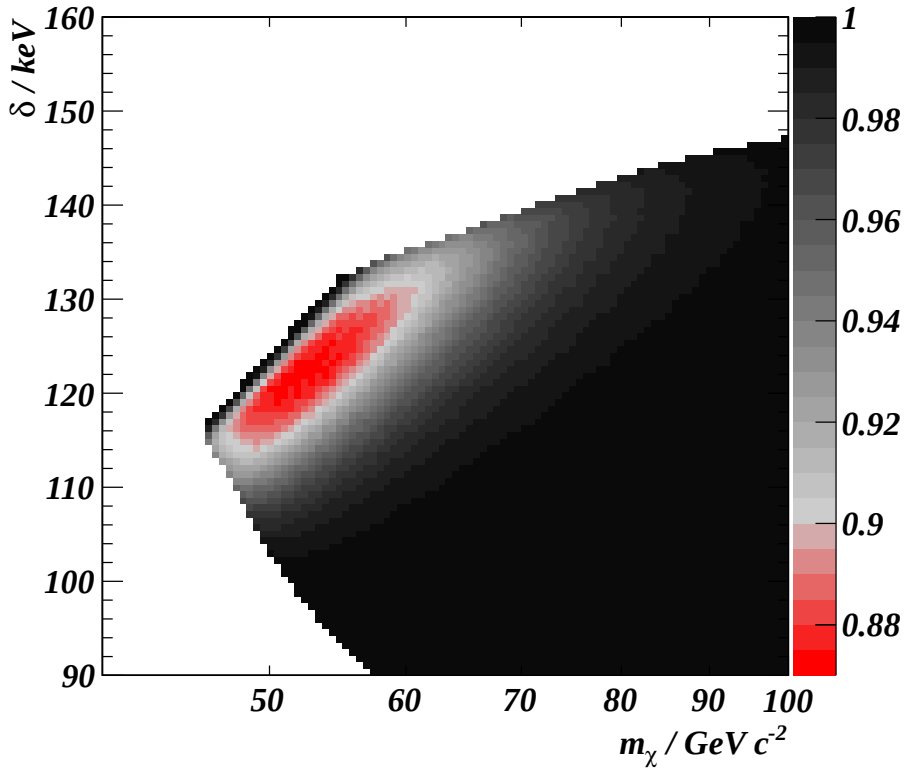

Figure 6.10: In δ-m_χ parameter space, the confidence level at which the FSR observation excludes the minimum inelastic WIMP-nucleon cross section consistent (99% CL) with causing the DAMA modulation by scattering from iodine.

Chapter 7

Conclusions

Understanding the nature of dark matter has emerged as one of the central problems in modern physics: the Λ-CDM model continues to pass ever-stricter experimental tests against the observed cosmic microwave background and large-scale structure; lensing studies map out the distribution of dark matter in galaxy clusters; and dynamics within galaxies also suggest a dominant component of invisible matter. Any non-gravitational interactions, however, have yet to be conclusively observed. The particle properties of dark matter are being sought in collider experiments, in searches for annihilation products and via the technique of direct detection, which is sensitive to a particularly well-motivated class of dark matter candidate: WIMPs. The preceding chapters have detailed the context, methodology and results of one of the world's most sensitive direct detection experiments, ZEPLIN-III — a two-phase emission detector that was the culmination of more than a decade of dark matter searches with liquid xenon at the Boulby mine.

One set of conclusions, contained in chapters 5 and 6, relates to the properties of dark matter. The limits on spin-independent WIMP-nucleon cross section, at 90% CL, from the combined runs of ZEPLIN-III (shown on page 121) reach a minimum of 3.9×10^{-8} pb for a WIMP mass of 50 GeVc^{-2}. These, along with other recent null results [40, 51, 33], are clear evidence that the excesses seen by DAMA and CRESST-II are not caused by canonical spin-independent WIMP-nuclear scattering. The WIMP-neutron spin-dependent limits are the most stringent results published by a direct search experiment and continue progress towards cMSSM-favoured regions in this channel. A modified WIMP phenomenology [161] had been proposed in 2001 and could account for both the DAMA modulation and the upper limits from other direct detection experiments: in this model, scattering leads predominantly to an excited WIMP state. The amplitude of modulation is enhanced and rates in light targets are suppressed due a target-mass-dependent minimum collision speed. Inelastic dark matter scattering from iodine had remained a viable explanation for almost a decade. The result presented in chapter 6 strongly disfavoured all but a very small, fine-tuned region of parameter space in which the minimum velocity for scattering from iodine is very close to the galactic escape velocity; even this was excluded at 87% CL for all cross sections consistent at 99% CL with causing the DAMA result.

A second set of conclusions relates to the experimental techniques and event analysis that underpins the cross section limits. Some conclusions are unequivocally positive. The year-long continual cold operation of ZEPLIN-III for the second run was possible because the collaboration set new standards for the slow-control and systems automation of a two-phase detector. The existence of detectable scintillation in liquid xenon for the lowest-energy nuclear recoils had been questioned, and the yield measurements described in section 3.5 were an important step in clearing this doubt.

Other conclusions are more cautionary. The PMT design for the second science run was, in retrospect, too ambitious a step in low-background construction for the manufacturers. The changed material properties of the tubes meant that previous photosensitisation procedures were no longer reliable. The extreme range of inherent gains, poor single-photoelectron responses and non-uniform photocathodes were the dominant reason for the drop in electron-recoil discrimination from 7800:1 in the FSR to 280:1 in the SSR. Production problems and PMTs which failed in liquid xenon, despite having passed cold testing by the manufacturer, led to the array being delivered 18 months behind schedule, and even then having 30 rather than 31 functioning tubes.

Several developments in calibration and event reconstruction for two-phase systems were made in the course of the ZEPLIN-III SSR. I worked in particular on correcting for channel-by-channel gain drift and on the use of background populations to assess the performance of these and other corrections using search data. Good uniformity in the mean response was achieved, but this on its own is not sufficient for the level of background discrimination needed — the inherent variance in response was significantly degraded from the first run (reflected in the near-doubling of FWHM energy resolution at 122 keVee).

The capability of two-phase xenon emission detectors for dark matter searches has been demonstrated emphatically. Low-background photomultipliers of the quality envisioned (but not realised) for the SSR are now readily available and the knowledge of xenon physics, detector techniques and event analysis achieved during the ZEPLIN-III programme are already being applied to the LUX experiment as it undergoes commissioning. In the coming years this technology is likely to lead field of direct detection, advancing to sensitivities beyond 10^{-11} pb.

Appendix A

Calibration sources

A.1 Cobalt-57

Gamma rays emitted by a ^{57}Co source were used to define the S1 and S2 electron-equivalent energy scales in ZEPLIN-III and to derive time-dependent corrections via daily calibration data (§3.3). ^{57}Co decays by electron capture to the 706 keV and 136 keV excited states of ^{57}Fe. The decay scheme and the γ ray energies and intensities are shown in Fig. A.1 and table A.1.

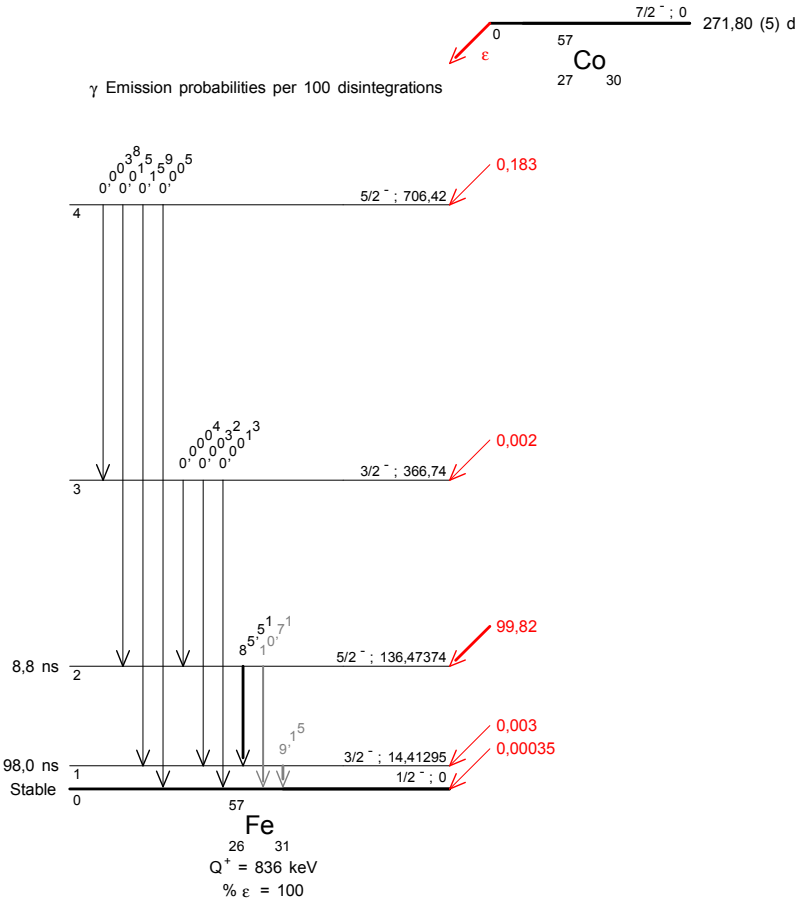

Figure A.1: Decay scheme of ^{57}Co [168].

Table A.1: γ-rays emitted by ^{57}Co

Energy [keV]	Relative intensity [%]
14.41	0.0915
122.06	0.8551
136.47	0.1071
230.27	0.0004
339.67	0.0038
352.34	0.0032
366.74	0.0013
569.94	0.015
692.01	0.159
706.42	0.0050

A.2 Caesium-137

Gamma rays emitted by a ^{137}Cs source were used to characterise the background from electron recoils in the ZEPLIN-III target (§5.3.2). ^{137}Cs decays by β^- emission, mainly to a metastable state of ^{137}Ba which then emits a 662-keV γ ray. The decay scheme and the γ ray energies and intensities are shown in Fig. A.2.

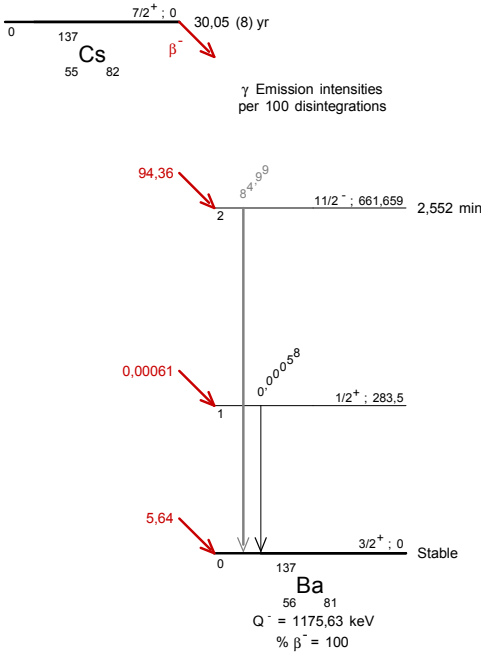

Figure A.2: Decay scheme of ^{137}Cs [169].

A.3 Americium-beryllium

Neutrons emitted by AmBe sources (which are a mixture of ^{241}Am and beryllium oxides) were used to measure the response of ZEPLIN-III to nuclear recoils (§3.5). Americium-241 decays via α emission to ^{237}Np, with a mean α-particle energy of 5.5 MeV. Subsequent (α,n) reactions with beryllium nuclei produce the neutron spectrum shown in Fig. A.3. These reactions also generate a rich spectrum of high-energy γ-rays, at a rate comparable to the neutrons; they are attenuated by lead shielding, and electron-recoil events caused by their interactions in the target are cut from neutron calibration data on the basis of the scintillation-to-ionisation ratio.

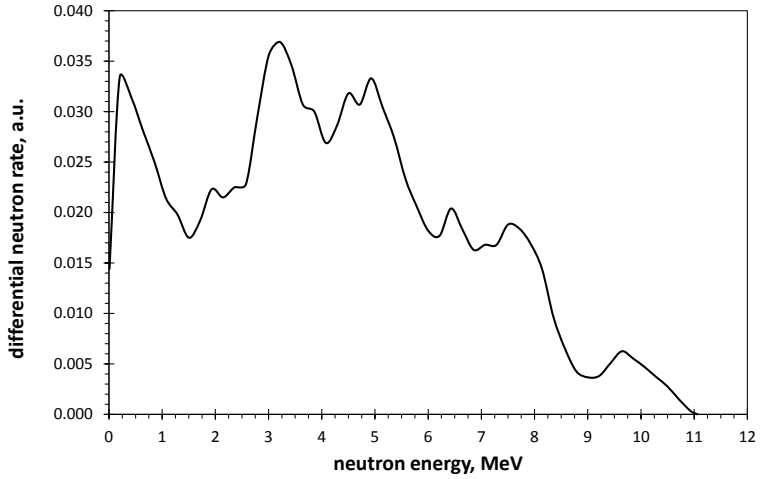

Figure A.3: Spectrum of neutron energies emitted, through (α,n) reactions, by AmBe source [170].

Appendix B

List of acronyms

BAO	Baryon acoustic oscillations
CDM	Cold dark matter (i.e. non-relativistic at freeze-out in the early universe)
CMB	Cosmic microwave background
dru	Differential rate unit ($\mathrm{keVee}^{-1}\mathrm{kg}^{-1}\mathrm{day}^{-1}$)
FSR	First science run of ZEPLIN-III (2008)
keVee	unit of electron-equivalent energy (i.e. amount of scintillation or ionisation relative to 122 keV photoabsorption)
$\mathcal{L}_{\mathbf{eff}}$	Energy-dependent scintillation yield for nuclear recoils at zero field (relative to 122 keV photoabsorption)
LRF	Light response function (i.e. expected fraction of emitted light recorded by a given photomultiplier, as a function of vertex-PMT separation in x,y)
$\mathcal{Q}_{\mathbf{y}}$	Energy-dependent ionisation yield for nuclear recoils at infinite field (electrons per keV recoil energy)
S1	Scintillation signal in the liquid (as a variable: the amount of scintillation in units of keVee)
S2	Signal resulting from ionisation charge drifted from the interaction site (as a variable: the amount of such charge in units of keVee)
SSR	Second science run of ZEPLIN-III (2010–11)
WIMP	Weakly-interacting massive particle

Bibliography

[1] F. Zwicky, *Die rotverschiebung von extragalaktischen nebeln*, *Helv. Phys. Acta* **6** (1933) 110.

[2] J. Einasto, *Dark matter*, in *Encyclopedia of life support systems* (O. Engvold, R. Stabell, B. Czerny, and J. Lattanzio, eds.), Eolss Publishers, 2009. [arXiv:0901.0632].

[3] A. Liddle, *An introduction to modern cosmology*. Wiley, 2003.

[4] E. Komatsu *et al.*, *Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: cosmological interpretation*, *Astrophys. J. Supplement Series* **192** (2011) 18.

[5] W. Hu and S. Dodelson, *Cosmic microwave background anisotropies*, *Ann. Rev. Astr. Astrophys.* **40** (2002) 171–216.

[6] The SDSS Collaboration, *Sloan Digital Sky Survey Gallery*, *http://www.sdss.org/gallery*, 2012.

[7] V. Springel, C. Frenk, and S. White, *The large-scale structure of the universe*, *Nature* **440** (2006) 1137, [astro-ph/0604561v1].

[8] D. J. Eisenstein *et al.*, *Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies*, *Astrophys. J.* **633** (2005) 560–574, [astro-ph/0501171].

[9] R. Massey, T. Kitching, and J. Richard, *The dark matter of gravitational lensing*, *Rep. Progr. Phys.* **73** (2010) 086901.

[10] M. Markevitch *et al.*, *A textbook example of a bow shock in the merging galaxy cluster 1E0657-56*, *Astrophys. J.* **567** (2002) L27, [astro-ph/0110468].

[11] D. Clowe, M. Bradač, and A. Gonzalez, *A direct empirical proof of the existence of dark matter*, *Astrophys. J.* **648** (2006) L109–L113.

[12] M. Bradač, S. Allen, T. Treu, and H. Ebeling, *Revealing the properties of dark matter in the merging cluster MACS J0025.4–1222*, *Astrophys. J.* **687** (2008) 959, [arXiv:0806.2320].

[13] H. Shan, B. Qin, B. Fort, and C. Tao, *Offset between dark matter and ordinary matter: evidence from a sample of 38 lensing clusters of galaxies*, *MNRAS* **1139** (2010) 1134–1139.

[14] M. Taoso, G. Bertone, and A. Masiero, *Dark matter candidates: a ten-point test*, *JCAP* **03** (2008) 022.

[15] C. Baker *et al.*, *An Improved experimental limit on the electric dipole moment of the neutron*, *Phys. Rev. Lett.* **97** (2006) 131801, [hep-ex/0602020].

[16] R. Peccei and H. R. Quinn, *CP conservation in the presence of pseudoparticles*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.

[17] K. Nakamura *et al.*, *Review of particle physics*, *J.Phys.G* **G37** (2010) 075021.

[18] P. Sikivie, *Axions*, in *Particle dark matter* (G. Bertone, ed.), Cambridge University Press, 2010.

[19] S. Asztalos *et al.*, *A SQUID-based microwave cavity search for dark-matter axions*, *Phys. Rev. Lett.* **104** (2010) 041301, [arXiv:0910.5914].

[20] C. Robilliard *et al.*, *No light shining through a wall*, *Phys. Rev. Lett.* **99** (2007) 190403, [arXiv:0707.1296].

[21] G. Servant, *Dark matter at the electroweak scale: non-supersymmetric candidates*, in *Particle dark matter* (G. Bertone, ed.), Cambridge University Press, 2010.

[22] L. Susskind, *Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory*, *Phys. Rev.* **D20** (1979) 2619–2625.

[23] G. Jungman, M. Kamionkowski, and K. Griest, *Supersymmetric dark matter*, *Physics Reports* **267** (1996) 195.

[24] J. L. Feng, *Dark matter candidates from particle physics and methods of detection*, *Ann. Rev. Astr. Astrophys.* **48** (2010) 495–545.

[25] Y. Sofue and V. Rubin, *Rotation curves of spiral galaxies*, *Ann. Rev. Astr. Astrophys.* **39** (2001) 137.

[26] H. Babcock, *The rotation of the Andromeda Nebula*, *Lick Observatory Bulletin* **19** (1939) 41–51.

[27] V. C. Rubin, W. K. Ford, and N. Thonnard, *Extended rotation curves of high-luminosity spiral galaxies IV – Systematic dynamical properties, SA through SC*, *Astrophys. J.* **225** (1978) L107–L111.

[28] D. Maoz, *Astrophysics in a nutshell*. Princeton University Press, 2007.

[29] X. Xue *et al.*, *The Milky Way's circular velocity curve to 60 kpc and an estimate of the dark matter halo mass from the kinematics of 2400 SDSS blue horizontal-branch stars*, *Astrophys. J.* **684** (2008) 1143.

[30] A. M. Green, *Astrophysical uncertainties on direct detection experiments*, [arXiv:1112.0524].

[31] M. Weber and W. de Boer, *Determination of the local dark matter density in our Galaxy*, *Astronomy and Astrophysics* **509** (2010) A25.

[32] J. Lewin and P. Smith, *Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil*, *Astroparticle Physics* **6** (1996) 87.

[33] J. Angle *et al.*, *A search for light dark matter in XENON10 data*, *Phys. Rev. Lett.* **107** (2011) 051301, [arXiv:1104.3088].

[34] A. Benoit *et al.*, *Measurement of the response of heat-and-ionization germanium detectors to nuclear recoils*, *Nucl. Instrum. Methods* **A577** (2007) 558–568.

[35] R. W. Schnee, *Introduction to dark matter experiments*, in *Proceedings of the Theoretical Advanced Study Institute in Elementary Particle Physics* (C. Csaki and S. Dodelson, eds.), p. 775, World Scientific, 2009. [arXiv:1101.5205].

[36] M. Drees and C.-L. Shan, *Model-independent determination of the WIMP mass from direct dark matter detection data*, *JCAP* **06** (2008) 012.

[37] M. Pato *et al.*, *Complementarity of dark matter direct detection targets*, *Phys. Rev.* **D83** (2011) 083505, [arXiv:1012.3458].

[38] M. T. Frandsen, F. Kahlhoefer, C. McCabe, S. Sarkar, and K. Schmidt-Hoberg, *Resolving astrophysical uncertainties in dark matter direct detection*, *JCAP* **01** (2012) 024, [arXiv:1111.0292].

[39] R. Lang *et al.*, *Electron and gamma background in CRESST detectors*, *Astroparticle Physics* **32** (2010) 318–324.

[40] Z. Ahmed *et al.*, *Dark Matter Search Results from the CDMS II Experiment*, *Science* **327** (2010) 1619–1621, [arXiv:0912.3592].

[41] D. Akerib *et al.*, *A low-threshold analysis of CDMS shallow-site data*, *Phys. Rev.* **D82** (2010) 122004, [arXiv:1010.4290].

[42] E. Armengaud *et al.*, *Final results of the EDELWEISS-II WIMP search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes*, *Physics Letters* **B702** (2011) 329–335.

[43] G. Angloher *et al.*, *Results from 730 kg days of the CRESST-II Dark Matter Search*, [arXiv:1109.0702].

[44] J. Kopp, T. Schwetz, and J. Zupan, *Light dark matter in the light of CRESST-II*, *Journal of Cosmology and Astroparticle Physics* **03** (2012) 001.

[45] G. Gerbier and J. Gascon, *Cryogenic detectors*, in *Particle dark matter* (G. Bertone, ed.), Cambridge University Press, 2010.

[46] D. Akerib *et al.*, *The SuperCDMS proposal for dark matter detection*, *Nucl. Instrum. Methods* **A559** (2006) 411–413.

[47] H. Kraus *et al.*, *EURECA — the European future of dark matter searches with cryogenic detectors*, *Nuclear Physics B — Proceedings Supplements* **173** (2007) 168–171.

[48] V. Lebedenko *et al.*, *Result from the First Science Run of the ZEPLIN-III Dark Matter Search Experiment*, *Phys. Rev.* **D80** (2009) 052010, [arXiv:0812.1150].

[49] G. Alner *et al.*, *First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two-phase xenon detector for dark matter detection*, *Astroparticle Physics* **28** (2007) 287–302, [astro-ph/0701858v2].

[50] J. Angle *et al.*, *First results from the XENON10 dark matter experiment at the Gran Sasso National Laboratory*, *Phys. Rev. Lett.* **100** (2008) 021303, [arXiv:0706.0039].

[51] E. Aprile *et al.*, *Dark matter results from 100 live days of XENON100 Data*, *Phys. Rev. Lett.* **107** (2011) 131302, [arXiv:1104.2549].

[52] O. Buchmueller *et al.*, *Implications of initial LHC searches for supersymmetry*, *Eur. Phys. J.* **C71** (2011) 1634, [arXiv:1102.4585].

[53] D. N. Spergel, *The motion of the Earth and the detection of WIMPs*, *Phys. Rev.* **D37** (1988) 1353.

[54] D. Dujmic *et al.*, *Improved measurement of the 'head-tail' effect in nuclear recoils*, *Journal of Physics: Conference Series* **120** (2008) 042030.

[55] S. Burgos *et al.*, *First measurement of the head–tail directional nuclear recoil signature at energies relevant to WIMP dark matter searches*, *Astroparticle Physics* **31** (2009) 261–266.

[56] E. Daw *et al.*, *Spin-dependent limits from the DRIFT-IId directional dark matter detector*, *Astroparticle Physics* **35** (2012) 397–401.

[57] K. Freese, J. A. Frieman, and A. Gould, *Signal modulation in cold dark matter detection*, *Phys. Rev.* **D37** (1988) 3388.

[58] R. Bernabei *et al.*, *New results from DAMA/LIBRA*, *Eur. Phys. J.* **C67** (2010) 39–49.

[59] C. McCabe, *DAMA and CoGeNT without astrophysical uncertainties*, *Phys. Rev.* **D84** (2011) 043525, [arXiv:1107.0741].

[60] P. J. Fox, J. Kopp, M. Lisanti, and N. Weiner, *A CoGeNT modulation analysis*, *Phys. Rev.* **D85** (2012) 036008, [arXiv:1107.0717].

[61] C. Aalseth *et al.*, *Search for an annual modulation in a p-type point contact germanium dark matter detector*, *Phys. Rev. Lett.* **107** (2011) 141301, [arXiv:1106.0650].

[62] Z. Ahmed *et al.*, *Search for annual modulation in low-energy CDMS-II data*, [arXiv:1203.1309].

[63] I. Aitchison, *Supersymmetry in Particle Physics: An Elementary Introduction*. Cambridge University Press, 2007.

[64] J. Ellis and K. Olive, *Supersymmetric dark matter candidates*, in *Particle dark matter* (G. Bertone, ed.), Cambridge University Press, 2010.

[65] N. Arkani-Hamed, A. Delgado, and G. Giudice, *The well-tempered neutralino*, *Nuclear Physics* **B741** (2006) 108–130.

[66] O. Buchmueller, R. Cavanaugh, and A. D. Roeck, *Higgs and supersymmetry*, [arXiv:1112.3564].

[67] G. Bertone, D. Cumberbatch, R. Ruiz de Austri, and R. Trotta, *Dark matter searches: the nightmare scenario*, *JCAP* **01** (2012) 004, [arXiv:1107.5813].

[68] T. A. Porter, R. P. Johnson, and P. W. Graham, *Dark matter searches with astroparticle data*, *Ann. Rev. Astr. Astrophys.* **49** (2011) 155–194.

[69] R. Abbasi *et al.*, *Multi-year search for dark matter annihilations in the Sun with the AMANDA-II and IceCube detectors*, *Phys. Rev.* **D85** (2012) 042002, [arXiv:1112.1840].

[70] M. Ackermann *et al.*, *Constraining dark matter models from a combined analysis of Milky Way satellites with the Fermi Large Area Telescope*, *Phys. Rev. Lett.* **107** (2011) 241302, [arXiv:1108.3546].

[71] D. Hooper and L. Goodenough, *Dark matter annihilation in the Galactic Center as seen by the Fermi Gamma Ray Space Telescope*, *Phys.Lett.* **B697** (2011) 412–428, [arXiv:1010.2752].

[72] D. Hooper, *The empirical case for 10 GeV dark matter*, [arXiv:1201.1303].

[73] M. Chernyakova, D. Malyshev, F. A. Aharonian, R. M. Crocker, and D. I. Jones, *The high-energy, arcminute-scale Galactic Center gamma-ray source*, *The Astrophysical Journal* **726** (2011) 60.

[74] K. N. Abazajian, *The consistency of Fermi-LAT observations of the galactic center with a millisecond pulsar population in the central stellar cluster*, *Journal of Cosmology and Astroparticle Physics* **03** (2011) 010.

[75] L. Bergstrom, *Dark matter evidence, particle physics candidates and detection methods*, [arXiv:1205.4882].

[76] C. Ghag *et al.*, *Performance of the veto detector incorporated into the ZEPLIN-III experiment*, *Astroparticle Physics* **35** (2011) 76–86.

[77] M. Horn *et al.*, *Nuclear recoil scintillation and ionisation yields in liquid xenon from ZEPLIN-III data*, *Physics Letters* **B705** (2011) 471–476.

[78] D. Akimov *et al.*, *WIMP-nucleon cross-section results from the second science run of ZEPLIN-III*, *Physics Letters* **B709** (2012) 14–20.

[79] D. Akimov *et al.*, *Limits on inelastic dark matter from ZEPLIN-III*, *Physics Letters* **B692** (2010) 180–183.

[80] N. Ackerman *et al.*, *Observation of two-neutrino double-beta decay in ^{136}Xe with the EXO-200 detector*, *Phys. Rev. Lett.* **107** (2011) 212501.

[81] A. Gando *et al.*, *Measurement of the double-beta decay half-life of ^{136}Xe with the KamLAND-Zen experiment*, *Phys. Rev. C* **85** (2012) 045504.

[82] T. Doke *et al.*, *Absolute scintillation yields in liquid argon and xenon for various particles*, *Japan. J. Appl. Phys.* **41** (2002) 1538–1545.

[83] T. Takahashi *et al.*, *Average energy expended per ion pair in liquid xenon*, *Phys. Rev.* **A12** (1975) 1771–1775.

[84] J. Lindhard, V. Nielsen, and M. Scharff, *Integral equations governing radiation effects*, *Mat. Fys. Medd. Dan. Vid. Selsk.* **33** (1963) 1–42.

[85] P. Sorensen and C. E. Dahl, *Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter*, *Phys. Rev.* **D83** (2011) 063501, [arXiv:1101.6080].

[86] B. Dolgoshein and V. Lebedenko, *New method of registration of ionizing-particle tracks in condensed matter*, *JETP Letters* **11** (1970) 351.

[87] J. Jortner, L. Meyer, S. A. Rice, and E. G. Wilson, *Localized Excitations in Condensed Ne, Ar, Kr, and Xe*, *J. Chem. Phys.* **42** (1965) 4250.

[88] S. Kubota, M. Hishida, M. Suzuki, and J. Ruan(Gen), *Dynamical behavior of free electrons in the recombination process in liquid argon, krypton, and xenon*, *Phys. Rev.* **B20** (1979) 3486–3496.

[89] E. Aprile and T. Doke, *Liquid xenon detectors for particle physics and astrophysics*, *Rev. Mod. Phys.* **82** (2010) 2053–2097.

[90] E. Aprile, K. Giboni, P. Majewski, K. Ni, and M. Yamashita, *Proportional light in a dual-phase xenon chamber*, *IEEE Trans. Nucl. Sci.* **51** (2004) 1986–1990.

[91] G. Bakale, U. Sowada, and W. F. Schmidt, *Effect of an electric field on electron attachment to sulfur hexafluoride, nitrous oxide, and molecular oxygen in liquid argon and xenon*, *J. Phys. Chem.* **80** (1976) 2556–2559.

[92] E. M. Gushchin *et al.*, *Electron emission from condensed noble gases*, *Sov. Phys.-JETP* **49** (1979) 856–859.

[93] A. Fonseca *et al.*, *Study of secondary scintillation in xenon vapour*, *IEEE Nuclear Science Symposium Conference Record* **7** (2004) 572–576.

[94] V. Lebedenko *et al.*, *Limits on the spin-dependent WIMP-nucleon cross-sections from the first science run of the ZEPLIN-III experiment*, *Phys. Rev. Lett.* **103** (2009) 151302, [arXiv:0901.4348].

[95] D. Akimov *et al.*, *The ZEPLIN-III dark matter detector: Instrument design, manufacture and commissioning*, *Astroparticle Physics* **27** (2007) 46–60.

[96] D. Akimov *et al.*, *The ZEPLIN-III anti-coincidence veto detector*, *Astroparticle Physics* **34** (2010) 151–163.

[97] H. Araújo *et al.*, *Measurements of neutrons produced by high-energy muons at the Boulby Underground Laboratory*, *Astroparticle Physics* **29** (2008) 471–481.

[98] H. Araújo *et al.*, *Radioactivity backgrounds in ZEPLIN–III*, *Astroparticle Physics* **35** (2012) 495, [1104.3538].

[99] F. Neves *et al.*, *ZE3RA: The ZEPLIN-III Reduction and Analysis Package*, *JINST* **6** (2011) P11004, [arXiv:1106.0808].

[100] V. N. Solovov *et al.*, *Position reconstruction in a dual phase xenon scintillation detector*, [arXiv:1112.1481].

[101] F. Neves *et al.*, *Position reconstruction in a liquid xenon scintillation chamber for low-energy nuclear recoils and γ-rays*, *Nucl. Instrum. Methods* **A573** (2007) 48–52.

[102] G. Alner *et al.*, *The DRIFT-I dark matter detector at Boulby: design, installation and operation*, *Nucl. Instrum. Methods* **A535** (2004) 644–655.

[103] F. Neves *et al.*, *Calibration of photomultiplier arrays*, *Astroparticle Physics* **33** (2009) 15, [arXiv:0905.2523].

[104] W. S. Cleveland, *Robust Locally and Smoothing Weighted Regression Scatterplots*, *J. Amer. Statist. Assoc.* **74** (1979) 829– 836.

[105] W. S. Cleveland, *LOWESS : A Program for Smoothing Scatterplots by Robust Locally Weighted Regression*, *Amer. Statist.* **35** (1981) 54.

[106] B. Edwards *et al.*, *Measurement of single electron emission in two-phase xenon*, *Astroparticle Physics* **30** (2008) 54–57.

[107] E. Santos *et al.*, *Single electron emission in two-phase xenon with application to the detection of coherent neutrino-nucleus scattering*, *JHEP* **12** (2011) 115, [arXiv:1110.3056].

[108] K. Ni *et al.*, *Preparation of neutron-activated xenon for liquid xenon detector calibration*, *Nucl. Instrum. Methods* **A582** (2007) 569–574.

[109] B. Edwards, *The ZEPLIN Dark Matter Search: Two Phase Xenon as a WIMP Target*. PhD thesis, Imperial College London, 2009.

[110] C. E. Aalseth *et al.*, *Results from a search for light-mass dark matter with a p-type point contact germanium detector*, *Phys. Rev. Lett.* **106** (2011) 131301.

[111] J. I. Collar and D. N. McKinsey, *Comments on "First dark matter results from the XENON100 experiment"*, [arXiv:1005.0838].

[112] J. I. Collar, *A realistic assessment of the sensitivity of XENON10 and XENON100 to light-mass WIMPs*, [arXiv:1106.0653].

[113] H. Araújo *et al.*, *The ZEPLIN-III dark matter detector: Performance study using an end-to-end simulation tool*, *Astroparticle Physics* **26** (2006) 140–153.

[114] S. Agostinelli *et al.*, *Geant4—a simulation toolkit*, *Nucl. Instrum. Methods* **A506** (2003) 250–303.

[115] M. Chadwick *et al.*, *ENDF/B-VII.0: Next generation evaluated nuclear data library for nuclear science and technology*, *Nuclear Data Sheets* **107** (2006) 2931–3060.

[116] G. Cowan, *Statistical data analysis*. Oxford University Press, 1998.

[117] G. Plante *et al.*, *New Measurement of the Scintillation Efficiency of Low-Energy Nuclear Recoils in Liquid Xenon*, *Phys. Rev.* **C84** (2011) 045805, [arXiv:1104.2587].

[118] A. Manzur *et al.*, *Scintillation efficiency and ionization yield of liquid xenon for mono-energetic nuclear recoils down to 4 keV*, *Phys. Rev.* **C81** (2010) 025808, [arXiv:0909.1063].

[119] E. Aprile *et al.*, *New measurement of the relative scintillation efficiency of xenon nuclear recoils below 10 keV*, *Phys. Rev.* **C79** (2009) 045807, [arXiv:0810.0274].

[120] F. Arneodo *et al.*, *Scintillation efficiency of nuclear recoil in liquid xenon*, *Nucl. Instrum. Methods* **A449** (2000) 147–157.

[121] D. Akimov *et al.*, *Measurements of scintillation efficiency and pulse shape for low-energy recoils in liquid xenon*, *Phys.Lett.* **B524** (2002) 245–251, [hep-ex/0106042].

[122] V. Chepel *et al.*, *Scintillation efficiency of liquid xenon for nuclear recoils with the energy down to 5keV*, *Astroparticle Physics* **26** (2006) 58–63.

[123] E. Aprile *et al.*, *Simultaneous measurement of ionization and scintillation from nuclear recoils in liquid xenon as target for a dark matter experiment*, *Phys. Rev. Lett.* **97** (2006) 081302, [astro-ph/0601552].

[124] P. Sorensen *et al.*, *The scintillation and ionization yield of liquid xenon for nuclear recoils*, *Nucl. Instrum. Methods* **A601** (2009) 339–346.

[125] P. Sorensen, J. Angle, E. Aprile, and F. Arneodo, *Lowering the low-energy threshold of xenon detectors*, in *Proceedings of IDM2010*, p. 017, SISSA, 2010. [arXiv:1011.6439].

[126] J. I. Read, G. Lake, O. Agertz, and V. P. Debattista, *Thin, thick and dark discs in* Λ*CDM*, *MNRAS* **389** (2008) 1041–1057.

[127] P. Toivanen, M. Kortelainen, J. Suhonen, and J. Toivanen, *Large-scale shell-model calculations of elastic and inelastic scattering rates of lightest supersymmetric particles (LSP) on I-127, Xe-129, Xe-131, and Cs-133 nuclei*, *Phys. Rev.* **C79** (2009) 044302.

[128] V. Barger, W.-Y. Keung, and G. Shaughnessy, *Spin Dependence of Dark Matter Scattering*, *Phys. Rev.* **D78** (2008) 056007, [arXiv:0806.1962].

[129] D. Tovey, R. Gaitskell, P. Gondolo, Y. A. Ramachers, and L. Roszkowski, *A new model independent method for extracting spin dependent (cross-section) limits from dark matter searches*, *Physics Letters* **B488** (2000) 17–26, [hep-ph/0005041].

[130] R. H. Helm, *Inelastic and elastic scattering of 187-MeV electrons from selected even-even nuclei*, *Phys. Rev.* **104** (1956) 1466–1475.

[131] J. Engel, *Nuclear form factors for the scattering of weakly interacting massive particles*, *Physics Letters* **B264** (1991) 114–119.

[132] M. Hjorth-Jensen, T. Kuo, and E. Osnes, *Realistic effective interactions for nuclear systems*, *Phys. Rept.* **261** (1995) 125–270.

[133] F. Kerr and D. Lynden-Bell, *Review of galactic constants*, *MNRAS* **221** (1986) 1023–1038.

[134] M. C. Smith *et al.*, *The RAVE survey: constraining the local Galactic escape speed*, *MNRAS* **379** (2007) 755–772.

[135] C. McCabe, *The astrophysical uncertainties of dark matter direct detection experiments*, *Phys. Rev.* **D82** (2010) 023530, [arXiv:1005.0579].

[136] C. Savage, K. Freese, P. Gondolo, and D. Spolyar, *Compatibility of DAMA/LIBRA dark matter detection with other searches in light of new galactic rotation velocity measurements*, *JCAP* **09** (2009) 036.

[137] E. Aprile, K. L. Giboni, P. Majewski, K. Ni, and M. Yamashita, *Observation of anticorrelation between scintillation and ionization for mev gamma rays in liquid xenon*, *Phys. Rev.* **B76** (2007) 014115, [arXiv:0704.1118].

[138] A. Hitachi *et al.*, *Effect of ionization density on the time dependence of luminescence from liquid argon and xenon*, *Phys. Rev.* **B27** (1983) 5279–5285.

[139] G. Alner *et al.*, *First limits on nuclear recoil events from the ZEPLIN-I galactic dark matter detector*, *Astroparticle Physics* **23** (2005) 444–462.

[140] J. Neyman, *Outline of a theory of statistical estimation based on the classical theory of probability*, *Philos. Trans. Roy. Soc. London* **A236** (1937) 333–380.

[141] C. Chuang and T. Lai, *Hybrid resampling methods for confidence intervals*, *Statistica Sinica* **10** (2000) 1–32.

[142] S. Yellin, *Finding an upper limit in the presence of unknown background*, *Phys. Rev.* **D66** (2002) 032005, [physics/0203002].

[143] S. Yellin, *Extending the Optimum Interval Method*, [arXiv:0709.2701].

[144] S. Henderson, J. Monroe, and P. Fisher, *The Maximum Patch Method for Directional Dark Matter Detection*, *Phys. Rev.* **D78** (2008) 015020, [arXiv:0801.1624].

[145] P. Benetti *et al.*, *First results from a dark matter search with liquid argon at 87K in the Gran Sasso underground laboratory*, *Astroparticle Physics* **28** (2008) 495–507.

[146] C. J. Clopper and E. S. Pearson, *The use of confidence intervals or fiducial limits illustrated in the case of the binomial*, *Biometrika* **26** (1934) 404.

[147] R. Cousins, J. Linnemann, and J. Tucker, *Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a Poisson process*, *Nucl. Instrum. Methods* **A595** (2008) 480–501.

[148] J. McMillan, *The low-temperature performance of fast photomultipliers*, *J. Phys.* **E22** (1989) 377.

[149] J. Neyman and E. Pearson, *On the problem of the most efficient tests of statistical hypotheses*, *Philos. Trans. Roy. Soc. London* **A231** (1933) 289–337.

[150] G. J. Feldman and R. D. Cousins, *A unified approach to the classical statistical analysis of small signals*, *Phys. Rev.* **D57** (1998) 3873–3889, [physics/9711021].

[151] A. Stuart and J. K. Ord, *Kendall's advanced theory of statistics, Vol. 2.* Edward Arnold, 5th ed., 1991.

[152] S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, *The Annals of Mathematical Statistics* **9** (1938) 60–62.

[153] H. Araújo *et al.*, *Low-temperature study of 35 photomultiplier tubes for the ZEPLIN III experiment*, *Nucl. Instrum. Methods* **A521** (2004) 407–415.

[154] R. Bernabei *et al.*, *First results from DAMA/LIBRA and the combined results with DAMA/NaI*, *Eur. Phys. J.* **C56** (2008) 333–355.

[155] R. Trotta, F. Feroz, M. Hobson, L. Roszkowski, and R. Austri, *The impact of priors and observables on parameter inferences in the constrained MSSM*, *JHEP* **12** (2008) 024.

[156] J. R. Ellis, K. A. Olive, Y. Santoso, and V. C. Spanos, *Very constrained minimal supersymmetric standard models*, *Phys. Rev.* **D70** (2004) 055005, [hep-ph/0405110].

[157] Z. Ahmed *et al.*, *Results from a low-energy analysis of the CDMS-II germanium data*, *Phys. Rev. Lett.* **106** (2011) 131302, [arXiv:1011.2482].

[158] L. Roszkowski, R. Ruiz de Austri, and R. Trotta, *Implications for the Constrained MSSM from a new prediction for $b \to s\gamma$*, *JHEP* **07** (2007) 075, [arXiv:0705.2012].

[159] G. Gerbier *et al.*, *Pulse shape discrimination and dark matter search with NaI (Tl) scintillator*, *Astroparticle Physics* **11** (1999) 287–302.

[160] R. Bernabei *et al.*, *The DAMA/LIBRA apparatus*, *Nucl. Instrum. Methods A* **592** (2008) 297–315.

[161] D. Smith and N. Weiner, *Inelastic dark matter*, *Physical Review* **D64** (2001) 043502.

[162] Y. Cui, D. E. Morrissey, D. Poland, and L. Randall, *Candidates for inelastic dark matter*, *JHEP* **05** (2009) 076.

[163] M. Kuhlen *et al.*, *Dark matter direct detection with non-Maxwellian velocity structure*, *JCAP* **02** (2010) 030.

[164] K. Schmidt-Hoberg and M. W. Winkler, *Improved constraints on inelastic dark matter*, *JCAP* **09** (2009) 010.

[165] E. Aprile *et al.*, *Implications on inelastic dark matter from 100 live days of XENON100 data*, *Phys. Rev.* **D84** (2011) 061101.

[166] P. Schuster, N. Toro, N. Weiner, and I. Yavin, *High Energy Electron Signals from Dark Matter Annihilation in the Sun*, *Phys. Rev.* **D82** (2010) 115012, [arXiv:0910.1839].

[167] M. Ajello *et al.*, *Constraints on dark matter models from a Fermi LAT search for high-energy cosmic-ray electrons from the Sun*, [arXiv:1107.4272].

[168] V. Chechev and N. Kuzmenko, *Decay Data Evaluation Project – Co-57,* *http://www.nucleide.org/DDEP.htm*, 2004.

[169] R. Helmer and V. Chechev, *Decay Data Evaluation Project – Cs-137,* *http://www.nucleide.org/DDEP.htm*, 2006.

[170] *ISO 8529-1: Reference neutron radiations — Part 1.* International Organization for Standardization, Geneva, Switzerland, 2001.