

Abstract of “A Position-Sensitive Liquid Xenon Time-Projection Chamber for Direct Detection of Dark Matter: The XENON10 Experiment” by Peter F Sorensen, Ph.D., Brown University, May 2008.

Recent astrophysical observations indicate that about 23% of the matter in the universe exists as non-luminous, non-baryonic dark matter. General thermodynamic arguments suggest that if an as-yet undiscovered weakly interacting particle were in thermal equilibrium in the early universe, it could have a cosmologically interesting relic abundance today. Dark matter particles are predicted to have collapsed into isothermal halos on a galactic scale. Minimal super-symmetric extensions of the standard model predict the existence of a stable particle with a mass in the range $10 - 1000 \text{ GeV } c^{-2}$, and a an interaction cross section with ordinary matter $\sigma_\chi < 10^{-36} \text{ cm}^2$. For the past two decades, numerous experiments have been deployed with the aim of direct or indirect detection of dark matter. XENON10 is a direct-detection liquid Xe experiment, with event-by-event 3D position reconstruction. Since dark matter does not interact electromagnetically, it's signature in a Xe detector is expected to be a low-energy nuclear recoil. Particle interactions in Xe create scintillation light ($S1$) and ionization. The ionization is drifted across the active Xe target and converted into a proportional scintillation signal ($S2$) by an external electric field. The ratio $y = \log_{10}(S2/S1)$ is larger for β and γ background events than for nuclear recoils, as determined by neutron calibration data. This allows event-by-event discrimination of background from the expected signal. In 2007 XENON10 eclipsed CDMS II as the most sensitive direct-detection experiment, excluding the existence of particle dark matter with a cross-section (normalized to a single nucleon) $\sigma_\chi > 10^{-43} \text{ cm}^2$ at a particle mass of $100 \text{ GeV } c^{-2}$ (90% C.L.). Liquid Xe technology is extremely promising since it is easily scaled to larger target mass, which will allow greater sensitivity to particle dark matter interactions. I discuss the XENON10 detector, deployment, operation, analysis and dark matter exclusion results. I also present a new method to determine the light yield of Xe for nuclear recoils, which is essential for calibrating the nuclear recoil energy scale and understanding the detector threshold.

A POSITION-SENSITIVE LIQUID XENON TIME-PROJECTION
CHAMBER FOR DIRECT DETECTION OF DARK MATTER: THE
XENON10 EXPERIMENT

by
PETER F SORENSEN

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN THE
DEPARTMENT OF PHYSICS AT BROWN UNIVERSITY

PROVIDENCE, RHODE ISLAND
MAY 2008

© Copyright 2008 by Peter F Sorensen

This dissertation by Peter F Sorensen is accepted in its present form by
the Department of Physics as satisfying the dissertation requirement
for the degree of Doctor of Philosophy.

Date _____

Professor Richard J. Gaitskell, Director

Recommended to the Graduate Council

Date _____

Professor Robert E. Lanou, Jr, Reader

Date _____

Professor Ian P. Dell'Antonio, Reader

Approved by the Graduate Council

Date _____

Sheila Bond
Dean of the Graduate School

*“Some people work very hard, but still they never get it right. Well I’m beginning to see the light...
wine in the morning, and some breakfast at night”*

– The Velvet Underground

for Alexandra

Acknowledgements

The XENON10 experiment owes its success to the dedication and perseverance of a relatively small number of scientists, and over the past 6 years I have benefited from their collective wisdom and experience.

I thank my advisor Rick Gaitskell for his tireless pursuit of knowledge: for showing me that every aspect of a problem, regardless of its size or perceived frivolity, should be pursued with the same degree of rigor... even at the expense of sleep and repast. The level of understanding that can be achieved as a result is unparalleled – and exciting. I also thank him for his patience with me, for leading by example, for always finding a moment when time was already in debt, and for giving me the chance to excel.

I thank Masaki Yamashita for initiating me into the cult of liquid Xe detectors, and for answering my incessant questions – usually in 10 seconds or less. I hope I have gained a fraction of the breadth of his experience in this field. Elena Aprile graciously hosted me at Nevis laboratory during the early days of XENON10 R&D. Under her watch I first recorded a flash of Xe scintillation from a small detector I had built. During that time I also received guidance on the intricacies of detector construction from Karl Giboni.

I look forward to the day that I tell Tom Shutt something about dark matter detectors that he does not already know. During operations and analysis I gained insight from discussions with Laura Baudis and Dan McKinsey, and in the pre-deployment testing of PMTs and amplifiers from Adam Bernstein and Celeste Winant. A special thanks to my readers, Robert Lanou and Ian Dell’Antonio for their attention to the manuscript.

The year in Italy was a trip. It is always a pleasant surprise when one’s co-workers become one’s friends, and I am thankful as much for the physics discussions as the camaraderie. In particular: Aaron Manalaysay’s alarming foosball skills, Eric Dahl’s abrupt and playful mastery of bash scripting, Luiz deViveiros’ inability to forget a towel, Simon Fiorucci’s decision to make his own prosciutto, Kaixuan Ni’s pico-tolerance of Moretti, Angel Manzur’s foray into Mexican pizza, Louis Kastens’ ability to share office space with one-and-only $\mathcal{J}$...

I am forever indebted to Abruzzo for honing my appreciation of prosciutto, pecorino, espresso and that aqua vitae Montepulciano d’Abruzzo; to the LNGS engineering staff for working past five on a friday in support of our accelerated time-table; to the Italian shipping company that taught a particularly shameless lesson in weaseling money out of foreigners; to the dogs who kept vigil outside my window nearly every night; and to the guards at the main gate of LNGS who famously referred to my advisor as “il biondo robusto” after several futile attempts to pronounce his name. To those whom I forgot to mention specifically by name, please re-read page iv and show me a little mercy.

I am especially grateful to my friends, my sisters and my parents – who maintained contact with me when I was under a mountain (either literally or figuratively) and unable to be contacted. And to the princess who patiently shared me with an endeavor that didn’t always share alike, I have already dedicated this work.

Contents

1	THE CASE FOR DARK MATTER	1
1.1	A matter of light and dark	2
1.2	Cosmological preliminaries	2
1.3	Astrophysical observations: evidence for the existence of dark matter	4
1.3.1	Cluster redshift measurements	4
1.3.2	Gravitational lensing	4
1.3.3	Rotational velocity of spiral galaxies	5
1.3.4	Elliptical galaxies	8
1.3.5	Modified Newtonian dynamics (MOND)	8
1.3.6	The extent of known dark matter	10
1.4	In support of Λ CDM	12
1.4.1	Observation and interpretation of the CMB	13
1.4.2	Large-scale structure	15
1.4.3	Big Bang nucleosynthesis	16
1.4.4	Type Ia supernovae	16
1.5	MSSM and a natural dark matter candidate	18
1.5.1	Other dark matter candidates	20
1.6	Dark matter detection	21
1.6.1	Indirect detection	21
1.6.2	Direct detection	22
2	THE XENON10 DETECTOR	26
2.1	The XENON10 Detector	27
2.2	Properties of Xe for Particle Detection	28
2.2.1	Available Signal	28
2.2.2	Primary Scintillation ($S1$)	29
2.2.3	Attenuation and Rayleigh Scattering	30
2.2.4	The Ionization Signal ($S2$)	31
2.3	A Typical Event	35
2.4	Xe Purity	37
2.5	Light Collection Efficiency	39

2.5.1	PTFE Reflector	39
2.5.2	Electrode Meshes	39
2.5.3	PMT Efficiencies	40
2.5.4	Geometrical Light Collection	40
2.6	Energy Scales: keVee and keVr	41
2.7	Background Discrimination	43
2.7.1	Pulse-Shape Discrimination	43
2.7.2	Ionization-Scintillation Yield Discrimination	45
2.8	Correlated Recombination Fluctuations	46
2.9	Scintillation yield for nuclear recoils	51
2.9.1	Measurement with a mono-energetic neutron source	51
2.9.2	Measurement with a broad-spectrum neutron source	52
3	INSTALLATION OF XENON10 AT LNGS	53
3.1	XENON10 installation at LNGS	54
3.1.1	Summary of operation at LNGS	57
3.2	Shielding and radioactive backgrounds	58
3.2.1	External sources of electromagnetic background	59
3.2.2	External sources of neutron background	60
3.2.3	Cosmic ray background	61
3.2.4	Neutron production from cosmic rays	62
3.2.5	Shield performance	63
3.2.6	Shield design	64
3.2.7	Xe self-shielding	71
3.2.8	Internal sources of electromagnetic background	71
3.2.9	^{85}Kr contamination	76
3.2.10	Internal sources of neutron background	78
3.3	Detector sub-systems	80
3.3.1	Cryogenics	80
3.3.2	Heat load	83
3.3.3	Safety and slow control	84
3.3.4	Xe gas handling and purification	85
3.3.5	PMTs	87
3.3.6	High voltage requirements	90
4	DATA ACQUISITION AND ANALYSIS	92
4.1	Data acquisition	93
4.1.1	Digitization	93
4.1.2	Time-offset jitter	94
4.1.3	Filter simulations	94
4.2	Data management	94
4.2.1	Zero th -order compression	95

4.2.2	Optimal compression	96
4.3	Trigger	98
4.3.1	$S1$ -based trigger	98
4.3.2	$S2$ -based trigger	99
4.3.3	Thresholds trigger rate and DAQ dead time	99
4.3.4	Single-electron $S2$	101
4.3.5	$S1$ peak-finding efficiency	102
4.4	Analysis algorithms	103
4.4.1	First pass: REEFrunner	103
4.4.2	Second pass: decisions	103
4.4.3	Position reconstruction	104
4.5	Position-dependent corrections	108
4.5.1	$S1$ z Correction	108
4.5.2	$S2$ (x, y) correction	108
4.5.3	PMT relative sensitivity correction	109
4.5.4	$S2$ electron lifetime correction	111
4.6	Data quality cuts	112
4.6.1	Basic quality cuts: QCO	113
4.6.2	Fiducial volume cuts: QC1	118
4.6.3	The need for advanced cuts (QC2)	122
4.6.4	The origin of non-Gaussian tail events	123
4.6.5	Description and performance of QC2 cuts	123
4.7	Acceptance of QC2 cuts for nuclear recoils	140
5	DARK MATTER EXCLUSION LIMITS	142
5.1	XENON10 dark matter exclusion limits	143
5.1.1	The WIMP acceptance box and $\Delta \log_{10}(S2/S1)$ coordinates	143
5.1.2	A comment on expected Gaussian tail events in the box	148
5.1.3	The importance of blind calibration data	148
5.1.4	Calculation of WIMP-nucleon exclusion limits	149
5.2	Discussion of the WIMP candidate events	149
5.2.1	Candidate events from the ROOT analysis	149
5.2.2	Resistor chain events	153
5.2.3	Candidate events from the Matlab analysis	156
5.2.4	Recommendations for future liquid Xe experiments	164
6	NEUTRON CALIBRATION AND $\mathcal{L}_{eff}$	166
6.1	Neutron calibration with an Americium-Beryllium source	167
6.2	Determination of the relative scintillation yield $\mathcal{L}_{eff}$ for liquid Xe	168
6.2.1	Event acceptance	170
6.2.2	Signal identification and $S1$ acceptance	173
6.2.3	Other systematics	174

6.2.4	Behavior of $\mathcal{L}_{eff}$ below 10 keVr	174
6.3	Determination of the ionization yield Q_y for liquid Xe	175
6.3.1	An alternate method	176
6.3.2	The utility of the alternate method	177
6.4	XENON10 energy threshold	178
7	MUON TRACKS IN XENON10	179
7.1	Ionizing Tracks	180
7.2	Saturation	182
7.2.1	Quick refresher on PMT saturation	182
7.2.2	Limits of decoupling capacitors	182
7.3	Energy calibration for muon tracks	184
7.4	Muon energy loss in Xe	184
7.5	Comparison with predicted specific energy loss	186
8	CLOSING REMARKS	188
	Bibliography	190

Chapter 1

THE CASE FOR DARK MATTER

“Poets say science takes away from the beauty of the stars - mere globs of gas atoms. I, too, can see the stars on a desert night, and feel them. But do I see less or more?”

– Richard Feynman

1.1 A matter of light and dark

On the one hand, cosmology today is in a rather contented state. A wide range of astrophysical observations are explained by a single, 6-parameter model known as Λ CDM. It explains the accelerating expansion of the universe, the cosmic microwave background radiation and the observed distribution of matter on the largest scales (Gpc). It further (indirectly) explains the observed velocities of galaxies within galaxy clusters, and the velocity of matter in spiral galaxies. Based on redshift measurements of galaxies on the Coma cluster, it has been known for 3/4 of a century that the luminous matter in that cluster appears to account for only $\sim 0.5\%$ of the cluster mass [174]. This first indication of the existence of some type of non-luminous matter has been repeatedly upheld by observations.

On the other hand, cosmology today is in a rather awkward state. A best fit to the Λ CDM model from analyses combining various astrophysical observations [81, 151] describes a universe in which only 4 – 5% of the energy density of the universe is in the form of atoms (baryons). A further 23% of the energy density appears to be a non-luminous, non-baryonic material known collectively (and to some critics, dull and unimaginatively [158]) as “dark matter”. The remaining 72% appears not to be matter at all, but rather an unknown energy density whose negative pressure causes the observed accelerating expansion of the universe.

On a drastically smaller scale, a new symmetry of space-time (“supersymmetry”) was initially proposed to explain the hierarchy problem [52] of particle physics. Among the predictions of the minimal supersymmetric standard model (MSSM) are a superpartner particle corresponding to each of the known standard model particles, differing in spin by $\frac{1}{2}$. No evidence for such particles has been observed in the mass range of the known particles, which implies that SUSY is a broken symmetry. In most MSSM models, the lightest supersymmetric particle is expected to be stable, with a mass of roughly 100 – 1000 GeV c^{-2} . If supersymmetry is in fact a symmetry of nature, there is compelling evidence that the lightest supersymmetric particle (if it is stable) constitutes some or all of the unknown 23% of non-luminous energy density in the universe today.

This coincidence of particle physics and cosmology, coupled with the alarming realization that we apparently know next to nothing about the nature of the majority of the matter in the universe, has led to numerous collaborative efforts which seek to directly detect particle dark matter. The XENON10 experiment is one of many recent experiments which use a liquified noble gas as the detector target medium. At the time of this writing, it is also the most sensitive direct detection experiment in the world, having recently eclipsed [10] the CDMS experiment [7]. The remainder of Chapter 1 will discuss the compelling astrophysical evidence for dark matter, possible dark matter candidates, and why one should expect to be able to directly detect it. The rest of this thesis will describe how one might go about detecting dark matter, via the technology of the XENON10 detector.

1.2 Cosmological preliminaries

The expansion (or collapse) of space with mass density ρ and pressure p in an isotropic, homogenous universe is given by the Friedmann equations, which were derived from the Friedmann-Lemaître-Robertson-Walker solution (“metric”) of Einstein’s field equations of general relativity. The 1st Friedmann equation can be written in terms of energy density Ω as

$$\left(\frac{H}{H_0}\right)^2 = \Omega_m \left(\frac{a_0}{a}\right)^3 + \Omega_k \left(\frac{a_0}{a}\right)^2 + \Omega_\Lambda \quad (1.1)$$

where the Ω_i ($i = m, k, \Lambda$) define matter, curvature and vacuum energy density relative to the critical energy density:

$$\Omega_m = \frac{\rho_0}{\rho_c} \quad (1.2)$$

$$\Omega_k = -\frac{1}{H_0^2 a_0^2 R^2} \quad (1.3)$$

$$\Omega_\Lambda = \frac{\Lambda}{3H_0^2} \quad (1.4)$$

The critical density for a flat (Euclidean) geometry of the universe is given by $\rho_c = 3H_0^2/8\pi G$, and is obtained by assuming the curvature R and the vacuum energy density Λ are identically zero. In addition to Newton's gravitational constant G , the Friedmann equation includes the scale factor a which relates proper coordinates to co-moving coordinates – and therefore is a measure of the relative expansion (or contraction) of the universe. It is worth noting that the universe is expected to appear isotropic to a co-moving observer.

The present-day scale factor is often defined to be $a_0 = 1$. The Hubble parameter H is defined as the rate of change of the scale factor with time:

$$H \equiv \frac{\dot{a}}{a} \quad (1.5)$$

and so is implicitly a function of time. For nearby objects, the present-day value H_0 can be obtained from Hubble's law $v = H_0 d$, relating the recessional velocity v of a cosmic object to its distance d from the observer. The criteria for “nearby” is imposed by the cosmological redshift z of the distant object. Cosmological redshift occurs because cosmic expansion dilates the wavelength of light as it travels. It is defined in terms of the rate of change of the scale factor of the universe: $1 + z = a_0/a$. Hubble's law is actually the first term in a Taylor expansion in z (or in d), and so is valid for $z \ll 1$. In such a case, the recessional velocity of an object is given by $v = cz$. For larger z , the expansion of the universe must be accounted for. The most robust measurements of H_0 are due to observations of Type Ia supernovae, which are discussed further in Sec. 1.4.4. Even tighter bounds on H_0 arise from combined analyses such as [81], which finds $H_0 = 70.1 \pm 1.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The present-day value of the Hubble constant often appears in dimensionless form h with units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. A detailed discussion of the concepts outlined here can be found (for example) in [125].

It can be seen from Eqs. 1.1 and 1.5 that at $a = a_0$, the relation $\Omega_m + \Omega_\Lambda + \Omega_k = 1$ holds. Recent analyses [81, 151] indicate that $\Omega_k \approx 0$, implying that we (appear to) live in a spatially flat universe. The unknown energy density referred to in Sec. 1.1 is identified with the vacuum energy density $\Omega_\Lambda = 0.721 \pm 0.015$ [81], which is more generally called the cosmological constant. Einstein originally proposed the cosmological constant to allow his theory of general relativity to result in a static universe – and then famously referred to this as his “greatest blunder”. However, its recent necessity in cosmological models and its stubborn defiance of our attempts to understand its nature make it look rather more like a legacy than a blunder. Ω_Λ is also known euphemistically as “dark energy.” Understanding the nature and presence of dark energy in the universe is a fundamental, outstanding problem in cosmology, but it will not be discussed further here.

The total matter energy density of the universe can be divided into two parts: $\Omega_m = \Omega_b + \Omega_c$, where $\Omega_b = 0.0462 \pm 0.0015$ is the baryon energy density and $\Omega_c = 0.233 \pm 0.013$ is the cold dark matter energy density. The values of both parameters are taken from [81]. Independent constraints on the baryon energy density in the universe are obtained from the theory of Big Bang nucleosynthesis (discussed in Sec. 1.4.3), from gravitational lensing (Sec. 1.3.2), and from calculations based on observed luminous matter in galaxies and clusters (Sec. 1.3.1 and 1.3.3). The descriptor “cold” implies matter that was already non-relativistic

when matter and radiation decoupled, about 380,000 y after the Big Bang. After this point, matter began to coalesce to form structure. This is discussed further in Sec. 1.4.

1.3 Astrophysical observations: evidence for the existence of dark matter

1.3.1 Cluster redshift measurements

Measurements of the velocity dispersion of galaxies in the Coma cluster by Zwicky lead to the first published [174] speculation that there is dark matter on the scale of galactic clusters. The Coma cluster had already been studied and catalogued by others, including Hubble, who found that its distribution of luminosity was in agreement with theoretical predictions for an “isothermal gravitational gas sphere” [175]. Under the assumptions that the cluster is stationary (neither expanding nor collapsing) and that Newton’s inverse-square law of gravitation holds, a measurement of the velocity dispersion of individual galaxies gives a measure of the total mass of the system. This is expressed by the virial theorem of classical mechanics as

$$-\overline{E_p} = 2\overline{K_T} = \sum M_i \overline{v_i^2} \quad (1.6)$$

where E_p is the total gravitational potential energy of the cluster, where v is the velocity of the mass M ; the bar denotes a time-average.

Zwicky derived a lower limit of $4.5 \times 10^{10} M_\odot$ on the average galactic mass. He made a detailed assessment of possible uncertainties, but his result that the mass-to-light ratio Υ (in solar units, $M_\odot/L_\odot$) in the Coma cluster of about 500 remained robust in order of magnitude. Observation of the local Kapteyn stellar systems showed a mass-to-light ratio of about 3, from which one concludes that stars account for only about 0.6% of the luminous matter in the Coma cluster. Zwicky did not postulate the nature of the non-luminous matter. However, he did state that while his result based on the virial theorem should be viewed with “considerable confidence”, it would be good to “throw some light on the problem of the density of internebular matter in clusters.” To this end, he suggested that gravitational lensing should be the “simplest and most accurate determination of nebular masses.”

More recent studies of the Coma cluster [172] found $\Upsilon = 182$. Of course, the indirect observation of dark matter is not limited to the Coma cluster; a survey of the virial mass of 89 clusters indicates an average mass-to-light ratio of $250 h$ [71]. The mass-to-light ratio for closure has been calculated in terms of the (B-band) luminosity density as $\simeq 1350 h$ [173]. One can thus conclude that $\Omega_m \simeq 0.19$. This must be taken as a lower bound, since in nearly all observable system the matter density appears to continue well beyond the extent of luminous material.

1.3.2 Gravitational lensing

While it was not until 1960s that the first gravitationally lensed objects (quasars) were observed, it is only recently that Zwicky’s prediction of an accurate determination of nebular (galactic) masses has been made. His qualification “simplest” would likely spark some debate; a review of gravitational lensing in astrophysics is given (for example) by [157]. An image of galaxy cluster CL0024 from the Hubble space telescope is shown in Fig. 1.1 (left). The cluster (at a redshift $z = 0.39$) acts as a strong gravitational lens for a high surface brightness galaxy located directly behind it. Due to this alignment, it was possible to make a detailed map

of the mass distribution within the cluster [155] (Fig. 1.1, right). The scale of distortions of a strongly lensed object allow tight constraints to be placed on the total mass of the lens.




Figure 1.1: **(left)** An image of galaxy cluster CL0024+1654 taken by the Hubble Space Telescope. **(right)** A strong gravitational lensing mass map of cluster CL0024+1654, showing sharp spikes corresponding to the observed luminous matter. A smooth background mass density is observed even in regions where there is no luminous matter. The luminous matter is found to account for about 0.5% of the mass in the cluster. Figure from [49]

The map was made by comparing the strong lensing image with models of the cluster mass, the source luminosity and the Hubble image. The spikes in mass density correspond to the luminous galaxies; what is striking about this map is the smooth, elliptical background mass density – extending out to several Mpc – which has no optical analog. The total mass-to-light ratio for this cluster was measured to be $\Upsilon = 276 \pm 40 h$ [155], referred to the visual band luminosity. Corrected for luminosity evolution, this is equivalent to a zero-redshift value of about $\Upsilon = 390 h$.

A recent combined (weak + strong) lensing analysis of 22 galaxies found an average mass-to-light ratio of $246^{+101}_{-87} h$ [68]. With the additional resolution achieved by weak lensing, their results show that the total mass density profile appears isothermal ($\rho \propto r^{-2}$) out to a radius $r \sim 300 h^{-1}$ kpc. The best fit to the density profile components are shown in Fig. 1.2. A de Vaucouleurs [156] profile is assumed for the stellar density and a Navarro-Frenk-White (NFW) [120] profile for the dark matter halo. The transition to dark matter dominated occurs at about $5 \text{ kpc } h^{-1}$ (indicated by the arrow), and inside this radius the dark matter fraction is calculated to be about 27%. It has been known for some time that the observed mass-to-light ratio tends to increase with the scale of the objects under consideration. Values range from $\Upsilon \approx 5$ for stellar observations (as measured by [68], for example) out to several hundred on the scale of clusters (~ 10 Mpc). From observations of Υ for large-scale structure, the matter energy density has been estimated as $\Omega_m \approx 0.14 - 0.18$ [153]. This is in agreement with the estimates obtained in Sec. 1.3.1, and is also a lower bound to Ω_m for the reasons already discussed.

1.3.3 Rotational velocity of spiral galaxies

It was observed as early as 1939 that the outer regions of the Andromeda galaxy (M31) were rotating faster than predicted by Newtonian dynamics [23]. Forty years later, rotation curves – a mapping of the velocity




Figure 1.2: The density profile $\rho(R)$ for 22 galaxies surveyed in [68]. The total profile (with line width indicating 1σ uncertainty) is roughly isothermal. The individual contributions of stellar and dark matter were modeled separately by de Vaucouleurs and NFW density profiles. The transition to dark matter domination occurs at about $5 \text{ kpc } h^{-1}$ (indicated by the arrow at $R \sim 0.005 \text{ Mpc } h^{-1}$), and inside this radius the dark matter fraction is calculated to be about 27%.

profile with radius from the galactic center – of M33 and NGC300 were made [64], based on observations of neutral hydrogen (HI). These and subsequent measurements (for example, a survey [39] of 25 spiral galaxies in 1978) showed that rotational velocities were decreasing extremely slowly (if at all) as a function of radius from the galactic center.

This is in sharp contrast to the Keplerian $v \propto r^{-1/2}$ that is obtained from equating Newton’s law of gravitation with his second law of motion:

$$v_{rot}(r) = \sqrt{\frac{GM(r)}{r}} \quad (1.7)$$

where $v_{rot}(r)$ is the observed rotational velocity at galactic radius r , and $M(r)$ is the enclosed total mass at the radius. Two recent examples of galactic rotation curves which asymptote to $v \sim \text{constant}$ are shown in Fig. 1.3.

Meanwhile, the measured (visible) surface-mass density of the galaxies shown in Fig. 1.3 falls exponentially by an order of magnitude by $r \approx 1 \text{ kpc}$, and by a further order of magnitude by $r \approx 10 \text{ kpc}$. This is essentially the situation already shown in Fig. 1.2 for the curve labeled “stars de Vauc.”. If luminous matter were the end of the story, then $M(r)$ would also fall exponentially. The data clearly require $M(r) \propto r$, which would be obtained if there is an additional (non-luminous) mass component which asymptotes to $\rho(r) \propto r^{-2}$ at large radii. This is the origin of the assumption that the galaxies are surrounded by an isothermal halo of dark matter.

A collection of gravitationally bound particles of mass M in thermal equilibrium at a temperature T constitute such an isothermal halo. Their phase-space distribution is given by the usual Boltzmann form, with the exponential term modified to include the gravitational potential energy. The density distribution $\rho(r)$ is not trivial to obtain, even with the simplifying assumptions that the dark halo is isotropic and




Figure 1.3: **(top)** The observed rotational velocity v_{rot} for NGC 4303 and NGC 4192 as a function of radius r from the galactic center. They were obtained from high-resolution observations of CO lines, and are reproduced from [142]. Considering the typical galactic density profile the luminous mass (shown in Fig. 1.2), the rotation curve would be expected to follow $v \propto r^{-1}$ at least by $r \approx 10$ kpc. **(bottom)** Images of NGC 4303 (left) and NGC 4192 (right) from SDSS [141]. The arrows indicate a span of about 20 kpc in each case.

spherical. However, it is usually approximated as

$$\rho(r) = \frac{\rho_0 r_c^2}{r^2 + r_c^2} \quad (1.8)$$

where r is the halo radius, r_c is the core radius and ρ_c is the core density. The core radius is defined at the point where the surface brightness has dropped to half its central value, and is about 200 pc for the galaxies shown in Fig. 1.3. This is shown (for example) in [142, 133], and leads to excellent agreement with observations. The apparent requirement of a galactic dark matter halo is not limited to a handful of galaxies; more than 1000 rotation curves are collected in [133], and shown to possess a nearly universal form that depends only on the galactic luminosity.

Given that our sun lies about 8 kpc from the galactic center of the Milky Way, it is of interest to know the local density of the non-luminous matter. A careful modeling of our galaxy, in which eq. 1.8 is further modified to account for anisotropy along the rotational axis, finds a central value of $\rho_{D,nb} = 0.45 \pm 0.18 \text{ GeV c}^{-2} \text{ cm}^{-3}$ (1σ errors extrapolated) [69]. This work considered a sample of about 10^6 “viable” galactic models, and

allowed for up to 30% of the non-luminous halo to be in the form of non-luminous stars. The subscript “ $_{nb}$ ” refers to non-baryonic dark matter. Non-luminous stars or planets (baryonic dark matter) may seem like a prime candidate for non-luminous dark matter; as we will discuss in Sec. 1.3.6, their contribution to the total is sub-dominant. It is customary to conservatively take the lower limit ($\rho_D \approx 0.3 \text{ GeV c}^{-2} \text{ cm}^{-3}$) as the typical value for the local non-baryonic dark matter density.

1.3.4 Elliptical galaxies

The observational technique just described is not generally applicable to elliptical galaxies, since their constituent stars do not move coherently. Understanding the mass distribution of such galaxies has tended to rely on less direct methods, such as relating the observed X-ray temperature to the mass under the assumption of hydrostatic equilibrium. Prominent X-ray emission can originate from the central region of the galaxy, due to the presence of ionized gas. These data are also best explained by a non-luminous halo [103]. It has been known for some time that certain elliptical galaxies contain a large, faint disk of neutral hydrogen HI. Recent observations of NGC 2974 with the Very Large Array radio telescope found the disk to actually be a ring, beginning at about $r = 5 \text{ kpc}$ and extending out to about $r = 12 \text{ kpc}$. This allowed the inferred velocity distribution from the ionized gas ($r \lesssim 2.5 \text{ kpc}$) and the velocity distribution of the neutral hydrogen ring ($r \gtrsim 5 \text{ kpc}$) to be compared with halo models [159]. This is shown in Fig. 1.4.




Figure 1.4: Measurements of the velocity field of ionized gas ($r \lesssim 2.5 \text{ kpc}$) and neutral hydrogen ($r \gtrsim 5 \text{ kpc}$) in NGC 2974 (circles). An NFW halo model is fitted to the data (solid line). Individual contributions from the stars, halo and neutral hydrogen gas are also shown (thin lines). 120 arcsec $\approx 12 \text{ kpc}$. A very similar quality of fit is obtained with either a pseudo-isothermal halo model, or under a modified Newtonian dynamics (MOND) scenario. Figure reproduced from [159].

1.3.5 Modified Newtonian dynamics (MOND)

MOND is an empirical theory proposed [114] to explain the observed rotation curves of spiral galaxies, without the need for an unknown dark matter component. It postulates that at very small accelerations $a \lesssim cH_0 \approx 7 \times 10^{-10} \text{ m s}^{-2}$, there may be deviations from Newtonian gravity. The Newtonian acceleration is thus modified as $a_N \rightarrow a_N \mu(a/a_0)$, where the function $\mu \sim 1$ for $a \gg a_0$, and $\mu \sim a/a_0$ for $a \ll a_0$. At large galactic radii the accelerations experienced by objects become comparable to a_0 , and in the simple case

of a circular orbit, one obtains $v = \sqrt[4]{GMa_0}$, in other words a flat velocity profile with galactic radius. Not surprisingly, MOND does an excellent job of explaining galactic rotation curves; in doing so, the constant a_0 can be determined from the data. Milgrom’s original estimate was $a_0 = 1.2 \times 10^{10} \text{ m s}^{-2}$.

The application of MOND to the Coma cluster resulted in a $\times 4$ over-prediction in the total mass of the cluster [136]. Two resolutions were proposed: either a_0 is about $\times 4$ larger for clusters than for galaxies, or there is some dark matter clumping on the scale of clusters (but not on the scale of galaxies). The latter proposal would effectively rule out cold dark matter, but it is compatible with a $\sim 2 \text{ eV}$ neutrino (see also [137]). The energy density of neutrinos in such a case would be roughly $\Omega_\nu \approx 0.10$, and a significant additional baryonic dark matter component would be required. Note that a 2 eV neutrino is excluded at $> 95\%$ CL by observations of large-scale structure [151], which is discussed further in Sec. 1.4.2. A discussion of neutrinos and baryonic dark matter can be found in Sec. 1.3.6.

A covariant implementation of MOND was applied to the evolution of large-scale structure in the universe [140], and it was found to be compatible with both observations of the cosmic microwave background radiation (see Sec. 1.4.1) and Type Ia supernovae (see Sec. 1.4.4). In this framework, each neutrino has a mass of 2 eV , and there is no need for any unknown dark matter. But it was recently shown that gravitational lensing is incompatible with the MOND explanation of galactic rotation curves, and a 2 eV neutrino as dark matter [62]. A calculation of strong gravitational lensing within the MOND framework found very similar results to the standard strong lensing analysis. In either framework, a sample of six strongly-lensed galaxies are shown to indicate a significant galactic dark matter component (which by definition does not exist in the MOND framework).




Figure 1.5: The Bullet cluster – a collision of two clusters – shows a shock front in the X-ray spectrum (pink smear), indicating the distribution of interstellar gas. The bulk of the mass of the cluster is indicated by the blue smear (obtained from weak lensing), which is strong support for the existence of weakly-interacting cold dark matter. Image courtesy of NASA.

Convincing arguments which indicate the “fin du MOND” are not new. Combined observations of the Bullet cluster in 2006 were announced as “direct proof of dark matter” [119], which would mean an indirect

contradiction of MOND. An understanding of the cluster was obtained in the visible (by HST), the X-ray (by Chandra) and the “unseen” (by weak lensing). This produced a clear picture of a collision between two clusters, which is shown in Fig. 1.5. The stars (which are difficult to pick out from the background, but which follow the blue smear) passed through essentially unscathed. This would be expected based on the measurements of Υ discussed above. The interstellar gas as imaged by Chandra shows a dramatic shock front caused by friction (pink smear). The weak lensing signal (the nail) indicates the total mass distribution of the system, and is indicated by the blue smear. The fact that most of the mass followed the stellar matter is strong support for weakly-interacting cold dark matter.

However, recent studies of the collisional dynamics of the Bullet cluster argued that the observed shock velocity is 23% higher than could reasonably be expected within a cold dark matter framework [11]. In contrast, MOND naturally accommodates the dynamics. The authors interpret this as strong support for MOND. An array of case-study arguments like this can be found in the literature. Clearly the simplicity of the MOND supposition appeals to many people (including myself). The problem in many cases is that arguments in favor of MOND (such as [11]) do not address the larger scope of the evidence for an as-yet unknown form of dark matter. In the extreme (such as [136]), they already contradict it. And in between, they propose entirely new paradigms (such as the $\Omega_\Lambda = 0.95$, $\Omega_\nu = 0.17$ of [140]). I discuss this evidence in detail in Sec. 1.4.

1.3.6 The extent of known dark matter

MACHOs

Nature offers plenty of known dark matter candidates, including brown dwarf stars ($M \lesssim 0.08 M_\odot$) and large gas giants such as Jupiter ($M \approx 10^{-3} M_\odot$). It was stated in Sec. 1.1 that most of the dark matter in the universe is non-baryonic, and this argument stands on firm ground from many angles (see Sec. 1.4). Still, it is worth pausing to examine what evidence is available for baryonic dark matter in our neighborhood. The EROS and MACHO experiments looked for gravitational microlensing in the direction of the Large Magellanic Cloud. Microlensing is distinguished from weak lensing (see [157], for a review) in that the lensed object is too small even for a large number of lensed objects to permit a statistical measure of their concerted distortions (“shear”). The remaining observable is therefore the change in a distant (lensed) object’s luminosity with time.

The EROS and MACHO collaborations reported a combined upper limit of 25% on baryonic dark matter in the galactic halo [3], in the mass range $10^{-7} M_\odot \lesssim m \lesssim 10^{-3} M_\odot$. This limit is in accord with the value of ρ_D quoted above. A subsequent paper by the MACHO collaboration reporting an increased exposure of 5.7 yrs [4] noted “there is intriguing evidence for a population of white dwarfs” that could contribute up to half the dark matter in the Milky Way halo. The evidence referred to [85]. However, subsequent white dwarf detections [86] and re-analysis of the MACHO and EROS data [154] placed an upper limit of 10% (95% CL) on baryonic dark matter in the galactic halo, in the mass range $0.08 M_\odot \lesssim m \lesssim 1 M_\odot$. For the un-initiated, MACHO is an acronym: a Massive Compact Halo Object is a generic term for a clump of baryonic matter in the mass ranges discussed above. One wonders if the critic at-large [158] would approve.

Neutrinos

Since the first convincing evidence for neutrino flavor oscillation [65], standard-model neutrinos are known to have a finite (but presumably small) mass. The results for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations from atmospheric neutrino observations found $\Delta m^2 \approx 3 \times 10^{-3} \text{ eV}^2$ at 90% CL. Recent results from the MINOS collaboration for ν_μ disappearance find a similar result [116]. The KamLAND experiment obtained a value of $\Delta m^2 \approx 8 \times 10^{-5} \text{ eV}^2$ for reactor $\bar{\nu}_e$ disappearance [16]. This sets a lower limit for the scale of the most massive (standard-model) neutrino. The neutrino contribution to the cosmological density (for a Dirac neutrino) is expected to be $\Omega_\nu h^2 = 2m_\nu/91 \text{ eV}$ [125]. Considering the present day value of the Hubble constant ($h = 0.7$) gives $\Omega_\nu \gtrsim 0.002$. This argument considers only the smallest mass of the heaviest of the three (known) neutrino flavors and so is a lower limit.




Figure 1.6: Allowed regions for the neutrino fraction f_ν of the total dark matter density $\omega_d \equiv \Omega_d h^2$. It is assumed that $\Omega_d = \Omega_\nu + \Omega_c$. The red, orange and yellow shaded regions are ruled out by WMAP and SDSS data, leaving the small allowed region. The center of the allowed region corresponds to $\Omega_\nu \approx 0.05 \cdot 0.11 h^{-2} = 0.011$. Constraints on the composition of the universe from WMAP and SDSS are discussed in Sec. 1.4. Figure reproduced from [151]

There is a crucial distinction to be made: neutrinos are relativistic (recent measurements by MINOS find $(v_\nu - c)/c \approx 5 \times 10^{-5}$ [116]) and so are referred to as “hot”. The distinction between “hot” and “cold” in this case is based on a particle’s speed at the time in the early universe when atoms began to coalesce to form structures. This is discussed in Sec. 1.5, in the context of cold dark matter. The distinction is particularly relevant because the formation of small-scale structures ($\sim 1 \text{ pc}$) in the universe would be severely hampered by an excess of hot dark matter. Recent analysis of a subset of the Sloan Digital Sky Survey (SDSS) combined with WMAP 3-year data place an upper limit on the neutrino contribution to the cosmological density as $\Omega_\nu \lesssim 0.024$ [151]. This is shown in Fig. 1.6. The SDSS and WMAP mission are discussed in Sec. 1.4.

1.4 In support of Λ CDM

The Λ CDM model of cosmology builds on the Friedmann equations (one of which is shown in Sec. 1.2), describing the universe in terms of the Friedmann-Lemaître-Robertson-Walker metric as the evolution of a (perfect) fluid with pressure p and density ρ . The initial point of the evolution is the Big Bang, which is then followed by cosmic inflation. The fluid is composed of the Standard Model particles, and some as-yet unknown cold dark matter particles. The model assumes spatial flatness ($\Omega_k = 1$) and a cosmological constant Ω_Λ , or vacuum energy density. The effect of a positive vacuum energy density is a negative pressure: as (empty) space expands, Ω_Λ causes the rate of expansion to increase. This prediction is discussed further in Sec. 1.4.4. A pictorial timeline of this situation is shown in Fig. 1.7.

The Λ CDM model also explains observations of large-scale structure (galaxies, clusters, superclusters, Great Walls), and the primordial fluctuations in the early universe ($t_{dec} \lesssim 380,000$ yr, $z_{rec} = 1089$ [31]) which seeded such growth. The fluctuations are generally assumed to be scale-invariant, adiabatic and Gaussian [147]. The early universe was a plasma of particles and photons; as it expanded, it cooled (adiabatically). As it approached a temperature $T \approx 1$ eV (3000 K) at time t_{dec} , neutral H atoms could form. At this point, photons decoupled from baryons, the universe was no longer opaque to radiation, and today we can observe the snapshot of this ancient epoch in the cosmic microwave background radiation (CMB). The remainder of this section describes the interpretation of recent cosmic survey data in terms of the Λ CDM model. It is worth noting that the model has only six basic parameters, from which others are then derived. Of the six, we discuss H_0 , Ω_m , and Ω_b . In particular, we are interested in $\Omega_c = \Omega_m + \Omega_b$.




Figure 1.7: A pictorial representation of the history of the universe based on the concordance model, with the Big Bang on the left and today on the right. The “afterglow light pattern” refers to the imprint left by the cosmic microwave background radiation, which is discussed in Sec. 1.4.1. Figure reproduced from [168].

1.4.1 Observation and interpretation of the CMB

The Wilkinson Microwave Anisotropy Probe (WMAP) [168], like the COBE satellite before it, measures temperature fluctuations in the microwave (23 – 94 GHz) sky. The fluctuations are an expected effect of photon-baryon decoupling: leading up to t_{dec} , photons still had enough energy to ionize atomic H. At the same time, the first structures began to coalesce under their own gravity. This leads to density fluctuations, and the expectation of small oscillations due to competition between radiation pressure and gravitational collapse. The characteristic scale of these oscillations prior to the decoupling of photons from baryons is thus imprinted in the temperature of the CMB photons. A sky map of the CMB temperature from [81] is shown in Fig. 1.8 (top). The local galactic plane runs along the major axis, and the Coma cluster is near the top of the minor axis. The average temperature is 2.725 K.

The temperature anisotropy at a direction $\mathbf{n}$ in the sky is expanded in terms of spherical harmonics (following [31]),

$$T(\mathbf{n}) = \sum a_{lm} Y_{lm}(\mathbf{n}) \quad (1.9)$$

where the expansion coefficient a_{lm} as observed in the sky is written as $C_l^{sky} = \frac{1}{2l+1} \sum |a_{lm}|^2$. The temperature anisotropy is then expressed as

$$\Delta T = \sqrt{\frac{C_l^{sky} l(l+1)}{2\pi}} \quad (1.10)$$

which can be reduced to the angular power spectrum. The temperature power spectrum is also shown in Fig. 1.8 (bottom), in terms of the multipole moment (or equivalently, the angular scale). The position of the first peak at $l \approx 200$ is interpreted as the scale of acoustic baryon-photon density oscillations ($r_s = 147 \pm 2$ Mpc [31]) at the time of decoupling. From r_s and the redshift at decoupling $z_{dec} = 1089 \pm 1$, the 13.7 Gy age of the universe is inferred.

The ratio of the heights of the first and second acoustic peaks in Fig. 1.8 is a sensitive probe of the baryon density, giving $\Omega_b = 0.0457 \pm 0.0018$ for $h = 0.7$ [31]. WMAP is also sensitive to the total matter density of the universe Ω_m via the early integrated Sachs-Wolfe effect. The Sachs-Wolfe effect is the gravitational redshift of CMB photons. The qualifier ‘early integrated’ refers to redshift that occurred shortly after CMB photons decoupled, and is the dominant source of temperature fluctuations for angular scales greater than about 10° (the left-hand edge of Fig. 1.8). The result $\Omega_m = 0.27 \pm 0.06$ from the first year of WMAP data only [31] is in good agreement with the most recent determination $\Omega_m = 0.279 \pm 0.008$ from 5 years of WMAP data, combined with large-scale structure and supernova data [81]. Combining the age and Ω_m determinations yields a value of the Hubble constant $H_0 = 71_{-3}^{+4} \text{ km s}^{-1} \text{ Mpc}^{-1}$. As I will discuss in Sec. 1.4.4, this is consistent with the (more accurate) determination of H_0 from observations of Type Ia supernova.

In summary, WMAP alone provides a determination of the fundamental Λ CDM parameters Ω_b , Ω_m and H_0 (among others) from a best fit to CMB temperature fluctuations. However, more stringent limits can be placed when combining WMAP data with measurements of Type Ia supernovae and large-scale structure. The most recent best fit is shown in Fig. 1.8 (bottom, red line). The composition of the universe determined from the Λ CDM model in conjunction with recent astrophysical observations [81, 95, 59] are summarized in Table 1.1.




Figure 1.8: A sky map of the CMB temperature, showing fluctuations (anisotropy) $\Delta T \approx 100 \mu\text{K}$ (top). The sky map data are reduced to an angular power spectrum (bottom), which can be compared with cosmological models (discussed in the text). Figure reproduced from [168].

Table 1.1: The composition of the universe according to Λ CDM and the most recent astrophysical observations. Data from [81].

Ω_{tot}	1.005 ± 0.006
Ω_b	0.046 ± 0.002
Ω_m	0.279 ± 0.008
Ω_c	0.233 ± 0.007
Ω_Λ	0.721 ± 0.015
Ω_ν	< 0.013
h	0.701 ± 0.013

1.4.2 Large-scale structure

After the CMB decoupled from matter, peaks in the density fluctuations (described at the beginning of Sec. 1.4) seeded the growth of structure in the early universe. Structures such as galaxies, clusters of galaxies (and the empty regions between them) can be modeled by linear perturbation theory as clustering due to gravitational instability. The model spans from before t_{rec} to the present. The detailed process is discussed (for example) in [125]. In the present discussion, it is sufficient to focus on two things: **(1)** if the theory is correct, observations of large-scale structure should show a clear signature of the acoustic oscillations already referred to in Sec. 1.4.1, and **(2)** the size of the oscillations is sensitive to the baryon content of the total matter density Ω_m . Numerous surveys have been made, including the 2dF Galaxy Redshift Survey and the larger SDSS (images from SDSS are shown in Fig. 1.3).




Figure 1.9: $\Omega_m - \Omega_\Lambda$ parameter space ruled out by data interpreted in terms of the Λ CDM model. Only SDSS and WMAP data are considered. The remaining allowed region is indicated by the small white ellipse. Figure reproduced from [151].

A subset of $\sim 47,000$ luminous red galaxies (LRG) from the SDSS was found to provide the best data sample for this analysis [59]. The typical redshift of the objects surveyed was $z = 0.39$. The authors found a small but clear (3.4σ) signature of baryon acoustic oscillation (BAO) at a scale of $100h^{-1}$ Mpc, which is consistent with (and independent of) the value of r_s obtained by WMAP. They point out that with our present understanding of Λ CDM cosmology, the small amplitude of the oscillation feature is not compatible with a large baryon fraction of Ω_m . Stated another way, there appears to be a requirement for a significant cold dark matter component of the early universe (at $z_{rec} \approx 1100$) in order to account for the observed structure today (at $z = 0.39$). The LRG data by themselves do not place an upper limit to the baryon fraction of Ω_m , or an independent measure of Ω_m . When combined with $\Omega_b h^2 = 0.024$ as obtained from the theory of Big Bang nucleosynthesis, the LRG data point to $\Omega_m \approx 0.27$ [59].

The value of this independent measurement of the BAO (in terms of its constraint on the Λ CDM model) is evident in a combined analysis such as [151]. Using only data from the LRG sample in conjunction with WMAP 3-year data reduces the allowed $\Omega_m - \Omega_\Lambda$ parameter space, as shown in Fig. 1.9.

1.4.3 Big Bang nucleosynthesis

Big Bang nucleosynthesis (BBN) describes a process of nuclear fusion which created the abundance of light elements that we observe today. The discussion here follows a recent review [148]. At times $t \lesssim 0.1$ s after the Big Bang, the universe was sufficiently dense and hot ($T \sim 3$ MeV) that essentially all particles were in thermal equilibrium. The neutron to proton ratio followed the expected form $(n/p) = \exp(-(m_n - m_p)/T)$ until about $t \sim 1$ s ($T \lesssim 1$ MeV). It then remained a roughly constant $1/6$ due to weak interactions with $e^\pm$ and $\nu_e, \bar{\nu}_e$, dropping slowly due to neutron decay.

At about $t = 200$ s ($T \sim 80$ keV), (n/p) rapidly approached zero; this was the start of BBN. It lasted less than thirty minutes. The reaction which started first is the fusion of Deuterium, $n + p \rightarrow D + \gamma$. Note that for $T \gtrsim 80$ keV, the arrow points in both directions (D is annihilated as quickly as it is produced). The production of He could then begin, followed by Li. The primordial abundances of these light elements are calculated based on a single input parameter, the baryon abundance $\eta_{10} \equiv 10^{10}(n_B/n_\gamma)$. The results are shown in Fig. 1.10 as abundance relative to H (by number); ^4He is shown as an abundance by mass, Y_P . The width of the colored bands indicate the uncertainty. Y_P is highly insensitive to η_{10} , since the production of ^4He is limited by the neutron abundance.

The primordial abundances are observed in a variety of astronomical objects by a variety of techniques and are not discussed here. A number of recent measurements are shown in [148], with the best estimates from this data indicated as a black dashed line in Fig. 1.10. The WMAP best fit [31] found $\eta_{10} = 6.1_{0.2}^{0.3}$, which is shown as a vertical (green) line in Fig. 1.10. Aside from ^7Li , the agreement is striking. As explained in [148], the disagreement in the case of ^7Li may tell us more about stellar processes than about the success of Λ CDM.

1.4.4 Type Ia supernovae

Supernovae are stellar explosions that can occur in various stages as a stellar core burns through its supply of hydrogen. Type Ia supernovae, however, are thought to occur when a white dwarf star with mass $M < 1.4M_\odot$ accretes material from a binary companion star. If the white dwarf mass exceeds the Chandrasekhar limit of $1.4M_\odot$, its electron degeneracy pressure is no longer sufficient to prevent gravitational collapse. In a period of a few seconds, the core material (mostly carbon) undergoes nuclear fusion, with the release of a




Figure 1.10: The calculated primordial abundances of light elements from Big Bang Nucleosynthesis as a function of the baryon to photon ratio η_{10} (colored bands. D, ^3He and ^7Li are shown relative to H by number, ^4He is shown relative to H by mass (labeled Y_P). The best estimates for the primordial abundances are indicated by the horizontal dashed lines (data from [148]). The WMAP best fit [31] found $\eta_{10} = 6.1^{0.3}_{0.2}$, indicated by the vertical line. Figure reproduced from [148].

tremendous amount of energy ($\approx 10^{63}$ eV). The “standard” conditions that create this explosion lead to a common peak magnitude, and thus make it an excellent choice of standard candle for measuring distance.

For redshift $z \lesssim 0.2$, the apparent brightness scales linearly with redshift (apparent brightness as a function of redshift is referred to as the luminosity distance, D_L). This allows the determination $H_0 = 70$ [126] from Hubble’s law (as described in Sec. 1.2), to within $\sim 10\%$. Supernova measurements are generally reported in terms of the distance modulus $\mu = 5(\log_{10}(D_L) - 1)$. A recent compilation of available Type Ia supernova data [96] is shown in Fig. 1.11. The data used in the combined analysis [81] reported in Table 1.1 made use of Astier *et al.* [22] (Fig. 1.11, red)

For larger redshift, the luminosity distance depends (to first order) on the matter and energy density of the universe. Therefore, the assumption of a particular model universe (say, $\Omega_m = 0.3, \Omega_\Lambda = 0.7$) leads to a specific prediction for the luminosity distance as a function of redshift. This places constraints on the allowed $\Omega_m - \Omega_\Lambda$ parameter space, as shown in Fig. 1.12. The figure also shows the 1, 2 and 3σ allowed regions from the complementary large-scale structure (BAO) and CMB analyses. The combined allowed region from this analysis [96] is the small grey ellipse.




Figure 1.11: A recent compilation of available type Ia supernovae distance modulus μ as a function of redshift. Figure reproduced from [96]. Tick marks represent $z = 0.2$.

1.5 MSSM and a natural dark matter candidate

By now the reader should be convinced that an alarming array of evidence indicates a significant dark matter component in the universe. As discussed in Sec. 1.3.6, a small fraction of that dark matter appears to be baryonic, and an additional small fraction appears to be non-baryonic but hot (neutrinos). The recent observations of the CMB anisotropy and of baryon acoustic oscillations in low-redshift large scale structure looks like “a stark confirmation of the cold dark matter theory for the growth of cosmological structure.” [59]. And yet, the evidence by itself offers very little suggestion of exactly *what* the cold dark matter is – other than that it appears not to be accounted for by the standard model of particle physics.

As discussed in Sec. 1.1, supersymmetric extensions of the standard model predict the existence of new superpartner particles. A detailed discussion of supersymmetry is beyond the present scope, and is reviewed extensively in (for example) [92]. There are two principal reasons that supersymmetry offers a well-motivated candidate for cold dark matter. In most minimal supersymmetric (MSSM) theories, the lightest supersymmetric particle is stable as long as R-parity is conserved. This is not a major caveat, since R-parity violation would allow (among other things) the decay of the proton. The mass of the lightest supersymmetric particle in most models lies in the range $100 - 1000 \text{ GeV } c^{-2}$. The lightest particle is generally also its own anti-particle (it is a Majorana particle). What makes such a particle a compelling candidate is that if its interaction cross-section (with baryonic matter) is at the scale of the weak interaction, it is also predicted to have a cosmologically interesting relic abundance today. The calculation of the relic abundance is discussed more generally in [125], and presented in [92] based on earlier work. It is briefly outlined here.

As discussed in Sec. 1.4.3, the universe just after the Big Bang was sufficiently hot and dense that all particles were effectively in thermal equilibrium. Imagine there exists a new, undetected stable particle χ . At sufficiently high temperatures, its number density is $n_\chi \propto \exp(-m_\chi/T)$. In equilibrium, χ undergoes continuous annihilation and creation reactions with lighter particles such as quarks or leptons ($\chi\bar{\chi} \leftrightarrow q\bar{q}$).




Figure 1.12: The 1, 2 and 3σ allowed regions in $\Omega_m - \Omega_\Lambda$ parameter space based on Type Ia supernovae (blue shaded). Also shown are similar regions for CMB constraints (orange shaded) and large-scale structure (BAO) constraints (green shaded). The combined allowed region is the small grey ellipse. Figure reproduced from [96] (truncated at $\Omega_\Lambda = 1.6$).

The annihilation rate per particle is $\Gamma_A = n_\chi \langle \sigma_A v \rangle$, where σ_A is the reaction cross section and v is the relative velocity of two particles. Eventually the rate of expansion of the universe will exceed the annihilation rate $\Gamma_A < H$, and the particle drops (“freezes”) out of thermal equilibrium. Its number density is then “frozen in.”

If this new particle were to have a weak-scale cross section, it is found that the freeze out would occur at about $T_f \approx m_\chi/20$, almost independent of the specific model assumptions. This shows that it would be non-relativistic (“cold”) by that time. Its cosmological relic abundance in units of the critical density can then be shown to be [92]

$$\Omega_\chi h^2 \approx \frac{0.1 \text{ pb } c}{\langle \sigma_A v \rangle} \quad (1.11)$$

which gives a cosmologically interesting relic density for a cross section characteristic of weak interactions. This result is perhaps surprising, placing a potential solution to a cosmological problem in the lap of a particle physics theory that was developed and motivated by entirely different concerns [52]. A few points should be emphasized. A larger cross section σ_A would lead to a smaller relic abundance today, and vice-versa.

The natural mass scale for particle dark matter from the preceding arguments is $\sim \text{GeV } c^{-2}$. This is not set by the assumption of weak-scale interactions, but rather by the requirement of an interesting relic abundance. If m_χ were larger or smaller by more than about 10^3 , the exponential dependence of the number density n_χ would lead to a relic abundance that either over-closed the universe, or were vanishingly small. An additional lower limit to the particle mass (in many models) can be inferred from the Large Electron-Positron collider (LEP). LEP data exclude a fourth-generation heavy neutrino with mass less than about $45 \text{ GeV } c^{-2}$, based on its absence in the observed width of the Z_0 decay [56].

The lightest supersymmetric particle is occasionally referred to as the LSP, and in the context of cold dark matter as a WIMP. A WIMP, or Weakly Interacting Massive Particle, refers to the general class of unknown particles postulated by supersymmetry. By definition, they interact only via the weak interaction and gravity. The whimsical contrast to the baryonic dark matter candidate discussed in Sec. 1.3.6 is not an accident. However, both in terms of cosmology and the lexicon, it is safe to say that the WIMP would have to come first.

1.5.1 Other dark matter candidates

The WIMP is far from the only cold dark matter candidate, and several other candidates are reasonably well-motivated. A notable case is the axion. It was postulated to resolve the question of why CP-violating phases arise naturally in quantum chromodynamics, but have not been observed. A classic example of this is that a non-zero CP-violating phase would be expected to introduce an electric dipole moment in the neutron, roughly 10^{12} larger than present experimental limits. Peccei-Quinn theory proposes that the CP-violating phase is actually a dynamic field (particle), the axion. As discussed in [40], typical axion masses are expected to be about 10^{-5} eV . A cosmologically interesting relic density can also be obtained in this case, and several experiments are seeking to detect axions [171, 60], so far with negative results.

A summary of the theoretical mass and cross section ranges for several dark matter candidates is shown in Fig. 1.13. An excellent (theoretical) review of the particle dark matter landscape is given in [131]. As an aside, the axion was first postulated by Wilczek [160] to “clean up” the strong CP problem. Axion brand detergent was eventually discontinued because its excessive phosphate content was found to pollute the environment. While this is unrelated (excepting the superstitious) to the present case for the existence of axion dark matter, it does point to the need for better outreach from the physics community to the general




Figure 1.13: A schematic summary of several well-motivated dark matter particles with cosmologically interesting ($\Omega \sim 1$) relic abundance. The interaction cross section with ordinary matter is σ_{int} . The neutrino is hot dark matter and is only shown for reference. Figure reproduced from [131].

public, particularly in regard to [158].

1.6 Dark matter detection

The preceeding discussions bring us to the subject of this thesis: given that 23% of the matter in the universe may be in the form of an as-yet unknown, undetected particle (here-after, a WIMP), how can the experimentalist best hope to detect it?

1.6.1 Indirect detection

One possibility that was indirectly raised in Sec. 1.5 is that WIMP annihilation ($\chi\bar{\chi} \rightarrow q\bar{q}$) might produce a detectable flux of high energy neutrinos [58]. A detectable flux would likely originate from the sun, since the local WIMP density there would be largest due to gravitational trapping. The neutrino energy is expected to be substantially higher than the energies of solar neutrinos, and the directionality of the signal would prevent contamination from atmospheric neutrinos. The IceCube experiment at the South Pole is looking for such a signature [79]. A further possibility is a high-energy γ flux from the cumulative effect of extra-galactic WIMP annihilation. Feynman diagrams for the annihilation into photons involve at least one loop, so the process is suppressed relative to the neutrino channel [92]. This process offers the particularly attractive possibility of a mono-energetic line at $E_\gamma \simeq m_\chi$. The GLAST space telescope (and others) are searching for this signal [170].

1.6.2 Direct detection

It is also predicted that a WIMP could generate a small energy deposition (recoil energy) by elastic scattering off an atomic nucleus. The size of the energy deposition depends on kinematics (including m_χ), and the rate depends on the coupling strength (σ_χ) of a WIMP to ordinary matter (such as quarks). As discussed, there are great uncertainties in both of these properties. But that does not affect the basic approach. A complete treatment of the couplings is given in [92], and of the generation of recoil energy in an earth-bound detector in [101]; I summarize some important results.

A WIMP will either couple to the spin of a target nucleus, or to its mass. These two cases are generally referred to as spin-dependent (SD) and spin-independent (SI) scattering. SI WIMP-nucleus scattering benefits from a coherent summation of the A scattering amplitudes in the nucleus. This is not the case for SD scattering, since paired nucleon spins sum to zero. SD scattering is therefore relegated to target nuclei with an unpaired nucleon. At the same time, the SD coupling per nucleon is larger than for the SI case. The result is that SI scattering amplitudes are expected to dominate the event rate for target nuclei with $A \gtrsim 30$. Importantly, the SI event rate will be higher for targets with larger A (such as Xe).

In the simplest case of a detector at rest in the galactic frame, the expected differential event rate (for isotropic scattering) as a function of recoil energy E_R is given by

$$\frac{dN}{dE_R} = \frac{R_0}{rE_0} e^{-E_R/rE_0} \quad (1.12)$$

where R_0 is the total event rate and $E_0 = m_\chi v_0^2/2$ is most probable kinetic energy of a WIMP. The most probable velocity $v_0 \simeq 230 \text{ km s}^{-1}$ is determined (assuming a Maxwell-Boltzmann velocity distribution) by the orbital speed of our solar system through the galaxy. This is expected to be valid for a galaxy with a flat rotation curve [57] – in other words, a galaxy with an isothermal dark matter halo (see Sec. 1.3.3). For a 100 GeV c^{-2} WIMP, $E_0 \approx 30 \text{ keV}$. The kinematic factor $r = 4m_\chi m_T / (m_\chi + m_T)^2 \simeq 1$ for a 100 GeV c^{-2} WIMP incident on a Xe target. Note that the total event rate R_0 is exactly calculable, but depends on σ_χ (which is not exactly known).

The situation is modified slightly in accounting for the Earth's motion around the sun ($v_E \simeq 244 + 15 \sin(2\pi y) \text{ km s}^{-1}$), and the finite galactic escape velocity $v_{esc} \approx 600 \text{ km s}^{-1}$. The expected event rate is then

$$\frac{dN(v_E, v_{esc})}{dE_R} = \frac{k_0}{k} \left[\frac{dN(v_E, \infty)}{dE_R} - \frac{R_0}{rE_0} e^{-v_{esc}^2/v_0^2} \right] \quad (1.13)$$

where $k_0 = (\pi v_0^2)^{3/2}$ and $k_0/k_1 = 0.9965$ is identically 1 for $v_{esc} \rightarrow \infty$. The Earth's annual motion is accounted for by

$$\frac{dN(v_E, \infty)}{dE_R} = c_1 \frac{R_0}{rE_0} e^{-c_2 E_R/rE_0} \quad (1.14)$$

where central values are $c_1 = 0.751$ and $c_2 = 0.561$ (the coefficients vary by about 6%, and are tabulated in [101]). If sufficient detection statistics were accumulated, a $\sim 3\%$ annual rate modulation would be expected.

So far the calculation has assumed the momentum transfer of the scattering to be $q \rightarrow 0$. In the case $q > 0$, further modifications to the calculated recoil energy spectrum are necessary. The modifications are accounted for by a nuclear form factor $F(q)$, which modifies the differential event rate by $F(q)^2$. The form factor becomes relevant when the wavelength $h/q \lesssim a_r$, where $q = \sqrt{2m_T E_R}$ and a_r is the characteristic nuclear radius. For example, the nuclear radius of Xe is $a_r \approx 6 \text{ fm}$, and a 5 keV recoil energy in Xe gives $h/q \approx 6 \text{ fm}$. As shown in Sec. 5.1, this implies that for Xe the form factor is always relevant. The physical interpretation of $F(q)$ is a loss of coherence in the scattering amplitudes (per nucleon) for $q > 0$, such

that they do not add in phase. The exponential decrease of event rate with E_R should already be enough to motivate the experimentalist to push for the lowest possible detector threshold; in terms of achieving maximal sensitivity, the behavior of $F(q)$ only increases the motivation.

Calculation of the form factor in the case of SD scattering requires knowledge of the quark content of the nucleus, and the expectation value of nucleon spin. The latter is a function of q , and is best obtained by detailed nuclear shell model calculations [129]. This leads to typical uncertainties of 20% – 50% in the event rate. In contrast, the SI form factor is just the Fourier transform of the nucleon density. Various analytical approximations of the nucleon density in Xe [101] lead to small differences in $F(q)$ for recoil energies $E_R \lesssim 30$ keV. For SI scattering in Xe, $F(q_{30}) \approx 0.3F(0)$, where q_{30} denotes the value of q corresponding to a recoil energy of 30 keV. Figure 1.14 shows the predicted differential event rate for SI scattering as calculated from Eq. 1.13 and Eq. 1.14, combined with the coherence factor enhancement and the form factor suppression.




Figure 1.14: The predicted recoil energy spectra (solid lines) of a 100 GeV c^{-2} WIMP in Xe, Ge and Ar based on the considerations discussed in the text. The assumed WIMP-nucleon cross section $\sigma_\chi = 2 \times 10^{-43} \text{ cm}^2$ was the smallest excluded cross section for a 100 GeV c^{-2} WIMP when XENON10 was commissioned in late 2006. The event rate is shown in standard units of counts $\text{keVr}^{-1} \text{ kg}^{-1} \text{ day}^{-1}$. The dashed lines indicate the total integrated event rate above a given energy threshold. For example, a 10 keVr threshold in Ge gives an integrated rate of $5 \times 10^{-2} \text{ kg}^{-1} \text{ day}^{-1}$. In other words: in 1 kg of Ge, 1 WIMP interaction would be expected every 20 days.

There are numerous direct detection experiments world-wide, employing roughly a dozen different target media. A recent review of the field is given in [67]. The isotopes Ge and Xe presently represent the most sensitive direct detection experiments [7, 10], and Ar technology looks promising [33]. The recoil energy spectra for these isotopes were calculated assuming a WIMP-nucleon cross section $\sigma_\chi = 2 \times 10^{-43} \text{ cm}^2$ for a 100 GeV c^{-2} WIMP. This was the smallest excluded cross section for a 100 GeV c^{-2} WIMP when XENON10

was commissioned in late 2006 [7]. The dashed lines in Fig. 1.14 give the corresponding integrated event rate above a particular detector energy threshold. For example, a 10 keVr threshold in Ge gives an integrated rate of $5 \times 10^{-2} \text{ kg}^{-1} \text{ day}^{-1}$. In other words, given $\sigma_\chi = 2 \times 10^{-43} \text{ cm}^2$ and a 1 kg Ge detector, 1 WIMP interaction would be expected every 20 days. Direct detection of dark matter is truly a rare-event search.

A non-detection result is also interesting. Once an experiment has placed a limit on the observed nuclear recoil event rate for a given detector exposure, they can place limits on the WIMP-nucleon cross section (for any given WIMP mass). An example of such a “limit plot” is shown in Fig. 1.15, for the WARP [33] (Ar, pink curve), Zeplin II [5] (Xe, green curve), CDMS [7] (Ge, blue curve) and XENON10 [10] (Xe, red curve) experiments. Also shown are new (un-published) limits from the CDMS experiment (Ge, blue dashed curve), and the projected sensitivity of the superCDMS (Ge, blue dotted) and LUX (Xe, green dotted) experiments. SuperCDMS and LUX expect to report results in 2013 and 2009, respectively. The grey [132] and yellow [61] shaded regions indicate two specific theoretical predictions for possible WIMP parameter space within the constrained MSSM. The theoretical predictions of WIMP parameter space extend down to (and below) $\sigma_\chi = 2 \times 10^{-46} \text{ cm}^2$. This corresponds to an integrated event rate of about 1 event $100 \text{ kg}^{-1} \text{ year}^{-1}$. Such a small event rate is not within the reach of current experiments, which have $\sim \text{kg}$ -scale target masses. The larger target masses of experiments like LUX (0.3 tonne) will probe a significantly larger region of allowed WIMP parameter space.

This thesis will discuss aspects of the design, deployment and analysis of the XENON10 experiment. Considering the sensitivity XENON10 achieved (Fig. 1.15, red curve), it is clearly a very promising new technology for direct detection of WIMP dark matter. What makes it particularly relevant for next-generation direct detection experiments is the relative ease of scaling the XENON10 detector concept to a significantly larger target mass. This can be seen in the projections of the LUX experiment in Fig. 1.15; in contrast, scaling the CDMS technology may not be as easy. This thesis will not discuss plans for a detector scale-up. However many of the results obtained from XENON10 will be directly applicable.




Figure 1.15: Exclusion limits on the WIMP-nucleon cross section σ_χ as a function of WIMP mass, from several recent experiments (solid lines). Two different calculations of allowed WIMP parameter space within the constrained MSSM are indicated (shaded). Current experiments have just begun to probe this parameter space. The projected sensitivity of two next-generation experiments – SuperCDMS and LUX – are also shown. The LUX experiment is also a liquid Xe detector, similar in principle to XENON10 but with more than a factor $\times 100$ increase in target mass.

Chapter 2

THE XENON10 DETECTOR

“Detectors are really the way to express yourself. To say somehow what you have in your guts. In the case of painters, its painting. In the case of sculptors, its sculpture. In the case of experimental physicists, its detectors.”

— Carlo Rubbia

2.1 The XENON10 Detector

XENON10 is a dual-phase (liquid-gas) Xe Time-Projection Chamber (TPC) for particle detection. The active target of XENON10 is 13.5 kg of liquid Xe, enclosed in a Teflon can (15 cm height, 10 cm radius), maintained at $T = 177$ K and $p = 2.1$ atm. (absolute). Particle interactions in Xe cause prompt primary scintillation ($S1$), which are recorded by 89 Hamamatsu R8520 PMTs (48 in the top PMT array, 41 in the bottom PMT array). Particle interactions in the Xe also create ionization electrons. An electric field $E_d = 0.73$ kV cm⁻¹ was applied across the active target by applying -12 kV to the cathode mesh, which was located 1.3 cm above the bottom PMT array. Two electro-formed stainless steel meshes at the top of the active target allowed for the application of a 2nd electric field $E_e \sim 13$ kV cm⁻¹ across the liquid-gas interface. Ionization can then be drifted across the active target and extracted into the gas by E_e . In the gas Xe, the ionization is rapidly accelerated, and the subsequent collisions with Xe atoms cause a large burst of secondary scintillation ($S2$). The detector is shown schematically in Fig. 2.1. Sub-systems are described in detail in Sec. 3.1.




Figure 2.1: Schematic Drawing of XENON10 (K. Giboni). A close-up view of the active target region is shown in Fig. 3.29.

2.2 Properties of Xe for Particle Detection

The crucial features of Xe which make it attractive for dark matter detection experiments are (1) availability of both scintillation and ionization signal, with event-by-event discrimination based on the amount of signal in each channel, (2) high density $\sim 3 \text{ g cm}^{-3}$ and relatively easy cryogenics ($T \sim 177 \text{ K}$) and (3) the absence of any naturally-occurring long-lived radio-isotopes. The longest half-life is $\tau = 36$ days in ^{127}Xe ($Q = 662 \text{ keV}$), which β decays with a typical endpoint energy $E_\beta > 287 \text{ keV}$. The decay is always accompanied by at least 1 γ , with $E_\gamma > 145 \text{ keV}$ typical. The isotope can be produced by neutron activation of ^{126}Xe (abundance $\sim 0.01\%$). ^{136}Xe is expected to undergo double β decay, with a Q value of 2458 keV [128]. The measured half-life limit for this process (either 2ν or 0ν) is $\tau > 4.4 \times 10^{23} \text{ yr}$ [48]. Relevant properties of Xe are summarized in Table 2.1, in which argon (also of interest for dark matter detection) is included for reference. Typical operating conditions of XENON10 during WIMP Search are summarized in Table 2.2.

Table 2.1: Properties of Xe and Ar relevant to particle detection.

Property	Xe	Ar	unit	Ref.
Boiling point	177	87.2	K	[123]
Density (liquid) ρ_L	2.86 ^a	1.40 ^a	g cm^{-3}	[123]
Refractive Index n	1.56, 1.69 ± 0.02 ^b	1.24 ^c		[28, 143]
Dielectric constant (liquid) $\epsilon(\omega)$	2.85 ^b	1.90 ^b		[139]
DC relative permittivity (liquid) ϵ_r	1.88 ^d , 1.96	1.33 ^d		[46, 138]
Dielectric strength	~ 400	$\gtrsim 400$	kV cm^{-1}	[90]
Scintillation wavelength λ_s	174 – 178	128	nm	[54, 26]
Spectral width (FWHM)	14		nm	[91]
Energy / scint. photon W_{ph}	21.6	24.4	eV	[55]
Energy / electron-ion pair W	15.6 ± 0.3	23.6 ± 0.3	eV	[55]
Xe_2^* lifetime (singlet) τ_s	2.2	5.0	ns	[54]
Xe_2^* lifetime (triplet) τ_s	27	1590	ns	[54]
Scint. absorption length λ_{abs}	> 100	> 90	cm	[26]
Rayleigh scattering length λ_R	29 ± 2 ^b	90.0 ^b	cm	[87, 139, 73]
Natural abundance	^{132}Xe (0.269) ^{129}Xe (0.264) ^{131}Xe (0.212) ^{134}Xe (0.104) ^{136}Xe (0.088)	^{40}Ar (0.996) ^{36}Ar (0.003)		[167]
Natural radio-isotope	^{136}Xe	^{39}Ar		[167]
Approx. cost	1000+	2	$\$ \text{ kg}^{-1}$	

^a at boiling point

^b at λ_s

^c for visible light

^d at triple point

2.2.1 Available Signal

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

Table 2.2: Standard operating conditions and properties of XENON10 during WIMP Search. The absolute uncertainty (*syst.*) in the pressure measurement arises because the gauge was set to read $p = 1$ at STP (on-site). LNGS is about 1000 m above sea level. The (*stat.*) uncertainty is the relevant number in terms of the temperature and pressure stability.

Property	Value	Unit
Temperature	177 ± 0.01 (<i>stat.</i>) ± 0.5 (<i>syst.</i>)	K
Pressure	2.11 ± 0.005 (<i>stat.</i>) ± 0.05 (<i>syst.</i>)	atm
Electric field (E_d)	0.73 ± 0.005	kV cm ⁻¹
Max. electron drift (dt)	79.3 ± 0.2	μ s
DAQ clock speed	105	MHz
Total Xe mass	~ 22	kg
Active target (r, z)	(10, 15) ^a	cm
Active target mass	13.5 ± 0.1	kg
Fiducial target (r, z) ^b	(9.45, 8.0 ± 0.1)	cm
Fiducial target mass	5.43 ± 0.03 ^c	kg

^a Uncertainty assumed to be sub-dominant.

^b Concentric with the active target and spaced symmetrically in z .

^c The value 5.4 kg was used.

[110]:



which subsequently decays



with the emission of a vacuum ultraviolet (VUV) photon ($\lambda = 178$ nm), with $FWHM \sim 14$ nm [91]. Ionized Xe atoms combine with un-ionized Xe atoms, also on a time-scale of a few picoseconds [110]. The process is understood as:



and in the absence of an external electric field, the exciton recombines to form an excited Xe atom



which then releases an additional scintillation photon according to Eq. 2.1 and 2.2. In XENON10, the ion recombination process is suppressed by the application of electric field (E_d), which leads to a quenching of the primary scintillation yield with respect to its zero-field value; this is discussed in Sec. 2.9. The number of scintillation photons (N_γ) and number of ionization electrons (N_e) constitute the available signal in Xe-based particle detectors, and their ratio of N_e/N_γ forms the basis for event-by-event background rejection (described in detail in Sec. 2.7).

2.2.2 Primary Scintillation ($S1$)

At zero applied electric field ($E_d = 0$), the time structure of Xe scintillation light is dominated by the electron-ion recombination time of about 45 ns [54]; otherwise, a substructure can be observed, corresponding to the

dissociation of an excited dimer (see Eq. 2.1 and Eq. 2.3). The dimer initially forms in either a singlet or a triplet spin configuration, with the latter measured to have a characteristic decay time $\tau = 27$ ns and the former having $\tau = 2.2$ ns [97]. The decay constant for the triplet state as measured from XENON10 data is consistent with previous measurements, as shown in Fig. 2.2. Because XENON10 digitizes events at 105 MHz, it is difficult to distinguish the singlet decay time (as shown in Fig. 2.2); however it is easily resolvable with a fast (> 1 GHz) digitizer, for pulses with sufficient signal size that Poisson statistics are not a limiting factor.




Figure 2.2: Data from XENON10 showing scintillation decay time corresponding to dissociation of triplet exciton states. The time scale was set to $t = 0$ at the peak pulse height. Photo-electron pulses on individual PMTs (from a single 80 phe S1 event) are shown in thin colored lines. The scaled sum of 319 S1 events (of 80 phe each) is shown (dashed) to indicate the shape of the summed S1 signal. A best-fit decay time to this sum is shown in solid blue. The pulse rise time is due to the electronics bandwidth (discussed in Sec. 4.1).

2.2.3 Attenuation and Rayleigh Scattering

As with other Noble liquids, Xe is essentially transparent to its own scintillation. This is because the energy of a scintillation photon released from the relaxation of the exciton (about 7 eV) is not sufficient to cause atomic excitation (the 1st ionization potential is $E = 12.1$ eV). The attenuation length of photons in Xe can be described by the usual expression $I(x) = I_0 e^{-x/\lambda_{att}}$, where $1/\lambda_{att} = 1/\lambda_{abs} + 1/\lambda_R$. A lower limit on the absorption length was measured by Baldini *et al.*, who found $\lambda_{abs} > 100$ cm [26] at 90% C.L.; the measurement was limited by the maximum dimension of the detector (50 cm), and was also consistent with no absorption ($\lambda_{abs} \rightarrow \infty$). The linear attenuation length is therefore governed by the Rayleigh scattering length, which has been calculated to be $\lambda_R = 30$ cm [139] and measured to be 29 cm [87]. However, since Rayleigh scattering is elastic, the light-collection capability of a Xe-based particle detector will be dominated by λ_{abs} , assuming a suitable VUV reflector is placed on non-instrumented surfaces; Teflon was used in XENON10, as discussed in Sec. 2.5.1. This is of critical importance for future detectors which will

attempt to increase sensitivity to rare-event interactions by scaling-up the detector mass [104] (and hence linear dimension).

2.2.4 The Ionization Signal (S_2)

When an external electric field (E_d , the “drift field”) is applied across the detector volume, the ionization created by an incident particle (see Eqn. 3 above) is separated and the electrons can be drifted through the liquid Xe at a velocity $v_d = 1.89 \pm 0.01 \text{ mm } \mu\text{s}^{-1}$ (at $E_d = 0.73 \text{ kV cm}^{-1}$ and $T=177 \text{ K}$). This value of v_d was measured in XENON10 from the 10% – 90% width of the maximum drift time across the 15 cm active target, which was $79.3 \pm 0.2 \text{ } \mu\text{s}$. The measured value agrees well with previous results [115], as shown in Fig. 2.3. The electron drift velocity is also a function of the liquid temperature: a 10 K decrease (increase) in temperature leads to about 5% increase (decrease) in v_d [32].




Figure 2.3: The measured electron drift velocity versus electric field in liquid Xe, from [115].

Two electrode meshes (See Fig.2.1 and Sec. 2.5.2) at the top of the liquid create a second electric field (E_e , the “extraction field”) across the liquid-gas interface. The drifted electrons are then extracted into the rarified gas phase, where they are accelerated and create electro-luminescence (secondary scintillation). The liquid level was maintained about 2 mm above the lower electrode, and the total gap between the two electrodes was 5 mm. The lower electrode was maintained at -1.15 kV and the upper electrode (the anode) at $+3.15 \text{ kV}$. Considering the liquid Xe relative permittivity $\epsilon_r \approx 2$ quoted in Table 2.1, and the gaseous Xe relative permittivity $\epsilon_{r,G} = 1.0$), the electric field as a function of the liquid level can be calculated; this is shown in Fig. 2.4.

The number of scintillation photons produced by a single electron accelerated across a Xe gas gap of $x \text{ cm}$ has been empirically characterized as a function of the extraction field E_e as [37]

$$n_{ph} = 70 \cdot (E_e/p - 1.0) \cdot x \cdot p, \quad (2.5)$$

where p is the gas pressure (atm.). For a gas gap of 2.5 mm, pressure $p = 2.1 \text{ atm}$ and extraction field $E_e = 12 \text{ kV cm}^{-1}$, this predicts 173 scintillation photons per extracted ionization electron. The width of the S_2 is determined by the width of the gas gap, in which electrons move with a velocity $v_{gas} \sim v_d$ (limited by




Figure 2.4: Electric field strength across liquid-gas interface, under standard operating conditions (Table 2.2). From Eq. 2.5, the number of scintillation photons produced per extracted electron is $n_\gamma = 164 \pm 20$. A schematic of the electrode meshes and the liquid level is also shown (blue shading).

collisions). The summed $S2$ pulse is very nearly Gaussian (see Fig. 2.5) for $S2$ signals smaller than about 122 keVee (keVee is the unit of electron equivalent energy; this is discussed in Sec. 2.6). A FWHM of about $\frac{2}{3} \mu\text{s}$ is typical, from which the typical liquid level is expected to be about 2.5 mm above the lower grid.

Considering the total predicted efficiency $\eta \approx 10\%$ for the conversion of a scintillation photon to a PMT photo-electron (Table. 2.4), a single electron extracted from the liquid would be expected to produce a signal of ~ 17 photo-electrons. XENON10 is sensitive to the electro-luminescence (hence-forth referred to as $S2$) from a single ionization electron, which is found to have $S2 = 24 \pm 7$ (1σ) photo-electrons. This is discussed in Sec. 4.3.4. Considering that Eq. 2.5 was determined from measurements in a different liquid Xe detector, with its own uncertainties in light collection efficiency and PMT quantum efficiencies, the agreement should be considered quite good. Light collection efficiency is discussed in Sec. 2.5.

The Xe level was measured by a triad of capacitive level meters, described in [108]. In practice it is the overall flatness (and stability), rather than the absolute level that is critical. The absolute liquid level relative to the anode grid was tuned so that the $S2$ pulses had a reasonable width, and so that the $S2$ from a 122 keV gamma (as from ^{57}Co) did not saturate the PMTs, Amplifier or ADC.

The width of an $S2$ pulse depends primarily on the size of the gas gap (x in Eq. 2.5) in which the electro-luminescence is generated. It also depends on the details of the particle interaction in the Xe target (the structure and orientation of the energy-loss track, from which the ionization is extracted). It further depends (indirectly) on the depth (z) of the interaction. As a cloud of ionization electrons drifts through the liquid Xe, they diffuse. The degree of diffusion depends on several factors including the elapsed time. The increase in $S2$ pulse width due to diffusion is shown in Fig. 2.6 (left). The pulse widths were obtained by Gaussian fit to about 2000 single scatter elastic nuclear recoil events in the center ($R < 50$ mm) of the detector. The events had energies $5 < E_R < 50$ keVr, and $500 < S2 < 5000$ phe. There is no noticeable dependence of the $S2$ width on the pulse size or type (electron recoil or nuclear recoil). Figure 2.6 (right)




Figure 2.5: **(left)** The electro-luminescence (S_2) signal from a ~ 10 keV γ interaction. The total signal is about 8×10^3 photo-electrons spread across the 89 PMTs (shown as thin colored lines); the summed waveform from all 89 PMTs is shown dashed, scaled by $\times 0.25$. The summed S_2 is well-fit ($\chi^2/d.o.f. = 1.05$) by a Gaussian (red) with $\sigma = 0.22 \mu\text{s}$ (FWHM = $0.52 \mu\text{s}$). **(right)** The corresponding hit pattern on the top PMT array shows a localization around PMT number 14. The position reconstruction is described in Sec. 4.4.3. PMT number 48 failed in August 2006 and is not shown in the figure. The bottom PMT array receives a diffuse wash of signal from S_2 (not shown). The color scale indicates $\log_{10}(\text{phe}/\text{PMT})$.

shows the distribution of pulse times for events with a fixed $\sigma = 0.25 \mu\text{s}$. There is clear tendency to select events with drift time $dt \sim 40 \mu\text{s}$. However, other processes such as the generation of the electro-luminescence photons and the non-uniformity of E_e as it approaches the mesh electrode contribute significantly to the $S2$ width. This currently precludes an effective determination of the z coordinate of interaction based on diffusion broadening. This parameter could be used to greater effect in future analyses.




Figure 2.6: **(left)** Broadening of the $S2$ pulse as it drifts is due to diffusion of the electron cloud. The pulse widths were obtained by Gaussian fit to about 2000 single scatter elastic nuclear recoil events in the center ($R < 50 \text{ mm}$) of the detector. The events had energies $5 < E_R < 50 \text{ keVr}$, and $500 < S2 < 5000 \text{ phe}$. There is no dependence of the $S2$ width on the pulse size. **(right)** The distribution of the event z coordinate for a fixed $S2$ pulse width ($\sigma = 0.25 \mu\text{s}$). While there is a clear tendency of the diffusion broadening to indicate the z coordinate of interaction, other processes contribute significantly to the pulse width. **(bottom)** The FWHM (σ shown along top) of the $S2$ signal as determined by Gaussian fits to ~ 2000 single scatter elastic nuclear recoils, with energies $5 < E_R < 50 \text{ keVr}$. The events were taken from the central region $R < 50 \text{ mm}$. The distribution of $S2$ widths has $\mu = 0.63 \mu\text{s}$ and $\sigma = 0.10 \mu\text{s}$.

2.3 A Typical Event

Figure 2.7 shows the response of XENON10 to a low-energy electron recoil Compton-scatter of about 4 keV electron-recoil equivalent energy (generally abbreviated as keVee; in contrast, nuclear recoil events are measured in keV nuclear-recoil equivalent energy, or keVr. The need for this distinction is discussed further in Sec. 2.7.2 and Sec. 2.9). The primary scintillation signal ($S1$) occurs at $20\ \mu\text{s}$ in the event record. The ionization created by the interaction is then drifted across the active Xe target by E_d , extracted into the gas phase by E_e , creating the $S2$ (Sec. 2.2.4). The elapsed time of $60\ \mu\text{s}$ between $S1$ and $S2$ gives the z coordinate of the interaction as 12.2 cm below the liquid surface (based on $v_d = 1.9\ \text{mm}\ \mu\text{s}^{-1}$). The (x, y) coordinates are obtained from the $S2$ hit-pattern on the top PMT array, which shows a locus near PMT number 18 (Fig. 2.8). Position reconstruction is described in detail in Sec. 4.4.3.




Figure 2.7: A typical low-energy single scatter γ of ~ 4 keVee, showing the sum of the 89 PMT channels. The $S1 = 13$ phe, $S2 = 1939$ phe, and the event drift time is $dt \sim 60\ \mu\text{s}$. Additional detail is shown in Fig. 2.8.

Single photo-electron after-pulsing for a period of $\sim 10\ \mu\text{s}$ following the $S2$ is generally observed. This event shows an additional feature at $100\ \mu\text{s}$, which upon closer inspection is found to be about 18 ($S2$) photo-electrons corresponding to a single ionization electron. Single electron $S2$ pulses are discussed further in Sec. 4.3.4. The fact that this single-electron $S2$ pulse occurs at almost exactly the maximum drift time ($80\ \mu\text{s}$) caused some speculation that a generating mechanism could be photo-electric emission from the stainless steel cathode mesh. The work function of stainless steel is 4.4 eV, compared with the 7 eV energy of a 178 nm UV photon. However, single-electron $S2$ pulses are found (in the WIMP Search data) to occur with equal probability at any point in the event record. Another example event showing an uncorrelated single-electron $S2$ can be found in Fig. 4.6.

A typical low-energy nuclear recoil event looks very similar, however the ratio $S2/S1$ (which in this case is about 190) tends to be smaller by as much as $\times 2$. A detailed discussion (and further example events) can be found in Sec. 2.6.




Figure 2.8: A typical event from a low-energy γ Compton scatter of ~ 4 keVee, showing (Top) the event waveform, each of the 89 channels plotted in a separate color; $S1 = 13$ phe, $S2 = 1939$ phe, and drift time $\sim 60 \mu\text{s}$. The DAQ samples every $(10^3/105)$ ns, so this event window corresponds to about 10^4 samples. A single-electron $S2$ appears at $100 \mu\text{s}$, almost exactly the maximum drift time (Bot. Left) a zoom on the $S1$, showing individual photo-electron pulses; (Bot. Right) the $S2$ hit-pattern on the top PMT array, from which position information is derived (colorbar indicates $\log_{10}(\text{phe}/\text{PMT})$).

2.4 Xe Purity

The lifetime of ionization electrons in liquid Xe depends critically on the concentration of impurities, particularly electro-negative impurities such as O_2 . Other common impurities are summarized in Table 2.3. The rate of attachment of O_2 to ionization electrons drops by about $\times 5$ as the applied electric field is increased from 1 kV cm^{-1} to 10 kV cm^{-1} [24], however a much larger effect on the lifetime ($> \times 100$) can be achieved by reducing the impurity concentration as much as possible (see below).

Xenon can be extracted from liquified air (or by fractional distillation) and is available commercially. For XENON10, research-grade Xe was obtained from Spectra Gasses Inc. with certified purity of 99.9998%. The composition (in *ppb*) is summarized in Table 2.3. The electron lifetime at this level of impurity was measured to be $20 \mu\text{s}$, which corresponds to a characteristic length $\lambda_{e-} = 4 \text{ cm}$. It is undesirable to have the available quantity of ionization electrons reduced by several e-foldings simply from drifting across the full 15 cm active target of the detector. A closed-circuit circulation system was installed to continuously pass the Xe through a purifier (SAES MonoTorr Phase II (model PS4-MT3-R) high-temperature getter). The purification capabilities as stated in the user manual are summarized in Table 2.3. The actual impurity concentrations in XENON10 were not measured, and could have been higher than the stated values due to continual out-gassing of detector components, or diffusion into the gas-handling lines down-stream from the getter. The purification sub-system is described further in Sec. 3.3.4.

Table 2.3: Impurity concentrations in the Xe obtained from Spectra Gases. (as measured by Spectra Gases), and the stated purification capability of the SAES PS-4 for Xe as determined by SAES. The efficacy of the PS-4 for reactive nitric oxide (NO) impurities was not stated by the vendor, but is expected to be similar to the level obtainable for O_2 [135].

Impurity	Spectra Gases Xe [<i>ppb</i>]	After SAES Getter [<i>ppb</i>]
O_2	100	< 1
N_2	< 100	< 1
CO_2	< 100	< 1
H_2O	11	< 1
THC	< 100	< 1
Kr	9	—

After several weeks of continuous purification, the electron lifetime achieved a value $\tau_{e-} > 2 \text{ ms}$, at which point the characteristic length is $\lambda_{e-} > 4 \text{ m}$. This corresponds to a charge loss of about 2% for ionization electrons which are drifted across the full 15 cm, and is easily corrected for. Electron lifetime data are shown in Fig. 2.9. The analysis is presented in Sec. 4.5.4. Error bars reflect the uncertainty from the fit. Noble gasses such as Kr are not removed by the getter; radioactive ^{85}Kr is a major concern for XENON10 and is discussed in Sec. 3.2.9.




Figure 2.9: Electron lifetime improvement due to continuous circulation through a closed-circuit purification system. Day 0 corresponds to July 31, 2006. The purification system is described in Sec. 3.3.4. Measurements from the background data (BG) and the calibration data.

2.5 Light Collection Efficiency

The light collection efficiency for primary scintillation ($S1$) is extremely important because it is the limiting factor in the energy threshold for XENON10. The detector components which affect the total light collection efficiency are discussed here. Note that factors affecting the secondary proportional scintillation ($S2$) yield are the same, but not of as much concern: the amplification of the proportional scintillation is sufficiently large that the $S2$ signal is not limited by the light collection efficiency.

A light collection Monte Carlo simulation (ray-tracing) of the XENON10 geometry was able to reproduce the observed ratio $\mathcal{R}(z) = S1_{top}/S1_{bottom}$ for primary scintillation light, as well as an estimate of the (geometrical) light collection efficiency [106]. This work found that the Teflon reflectivity is degenerate with the absorption length, but it does place lower limits. The absorption length is consistent with $\lambda_{abs} \rightarrow \infty$. Factors affecting this efficiency are discussed below, and summarized in Table 2.4.




Figure 2.10: Photos taken during construction of the XENON10 detector. **(left)** A top-down view of the electrode meshes, the Teflon reflector which defines the active target, and the bottom PMT array of 41 PMTs. **(right)** A different angle shows the 3 top electrode meshes; proportional scintillation ($S2$) is generated between the lower of the two.

2.5.1 PTFE Reflector

The active target region of the detector is surrounded by a cylindrical polytetrafluoroethylene (PTFE, purchased as Teflon[®]) shell, as shown in Fig. 2.10. Measurements of the diffuse reflectivity of Teflon for Xe scintillation light range from $R = 82\%$ [117] to as high as $R = 95\%$ [163]. The latter value was obtained under the assumption $\lambda_{abs} = 100$ cm; as $\lambda_{abs} \rightarrow \infty$, a value $R = 88\%$ was obtained. As demonstrated in Sec. 2.5.4, the use of Teflon increases the primary ($S1$) light yield by $> \times 5$. The light collection Monte Carlo simulation [106] assumed the reflectivity to be 95%.

2.5.2 Electrode Meshes

The stainless steel grids were electro-formed and electro-polished to ensure that all edges are smooth. The thickness of the grid is 0.203 mm, the width of the bars 0.182 mm, and the grid spacing is 2.0×2.0 mm. From these considerations the optical transparency is about 90% if the surfaces are considered to be

perfectly absorbing. Data for stainless steel reflectivity in the VUV are sparse; a measurement at 121 nm shows specular reflection of 10% – 24% (for 15° – 65° incident angle) [89]. In the visible range of the spectrum, measurements are closer to 62% for the same angular range, increasing to > 80% at higher angles [94]. The light collection Monte Carlo simulation assumed the reflectivity of the electrode meshes to be 50% [106].

2.5.3 PMT Efficiencies

XENON10 opted to use the compact, rectangular Hamamatsu R8520 PMT, with metal channel dynodes and a 2.5×2.5 cm window. The average photo-cathode quantum efficiency (QE) for the 89 PMTs installed in XENON10 is 0.237 ± 0.019 at 178 nm [74]. A more detailed discussion of the PMTs can be found in Sec. 3.3.5). Because of its small size, the difficulty of focusing of photo-electrons from the photo-cathode to the first dynode results in a 0.75 ± 0.05 efficiency for collecting these photo-electrons. This is referred to as “first-strike” probability, or simply collection efficiency (CE). Larger PMTs, for example the Hamamatsu R8778, are reported to have $CE = 0.88$. It is important to realize that the QE measurements are made at room temperature. The ZEPLIN III experiment has reported an 18% increase in QE between 293 K and 177 K [17], using a sample of 35 ETL 9829Q PMTs, also with bi-alkali photo-cathode. A similar (12%) improvement was also seen for a sample of 4 Hamamatsu R8778 PMTs [21]. Assuming a 12% increase in the photo-cathode QE, the effective total R*50 PMT QE (including the CE) is expected to be 0.194 ± 0.021 phe/scintillation photon.

2.5.4 Geometrical Light Collection

The greatest uncertainty comes from the effective geometric coverage of the 89 PMTs. As an example, it is instructive to consider an event in the center of the detector and neglect the reflectivity of the Teflon and the SS electrodes. The direct line of sight from this point to the edge of the bottom PMT array subtends about 30° (6.7% of the total solid angle); for the top PMT array, the situation is similar but the refraction (and total internal reflection) at the liquid-gas interface reduces the effective solid angle by more than half. The total effective solid angle of the sum of the top and bottom PMT arrays is then about 10%. The active area of the PMT photo-cathode within each PMT array is only about 65% (see Fig. 3.27, left). This would predict a total geometric light collection efficiency of 6.5%. However, the photon tracking Monte Carlo simulation [106] using the efficiencies summarized in Table 2.4 indicate an overall geometric light collection efficiency $\sim 40\%$.

It is also possible to infer the light collection efficiency directly from XENON10 data, based on the W_{ph} -value. The energy required to create a scintillation photon in liquid Xe has been measured (see Table 2.1) for 976 keV relativistic electrons (minimum ionizing particles, $\beta = 0.94$) to be $W_{ph} = 21.6 \pm 2.4$ eV per photon [55]. From XENON10 γ calibration data, the best available comparison are relativistic ($\beta = 0.90$) 662 keV electrons (from ^{137}Cs calibration). These arise from photo-absorption, or single-scatter events from the full energy peak. The scintillation yield in liquid Xe depends on linear energy transfer (LET), and the difference in LET between a 662 keV and 976 keV electron is about 5% [41]. The variation in liquid Xe scintillation yield has been calculated by Doke *et al.* (see [55], Fig. 4), and measured by Barabanov *et al.* [27] to also increase by 5% for this change in LET.

In XENON10 2.20 phe/keVee were measured for 662 keV single-scatter γ events in the center of the detector. From the LET scaling it can be inferred that the yield would be 2.31 phe/keVee for 976 keV

Table 2.4: Factors affecting the $S1$ light collection efficiency (phe/scintillation photon) in XENON10. Geometric light collection was inferred from the data, and independently determined by a Monte Carlo simulation. The inputs to the light collection Monte Carlo simulation used to obtain the geometric light collection are explained in Sec. 2.5. The predicted total $S1$ light collection efficiency is the product of the PMT efficiencies and the geometric light collection.

Component	Efficiency η
R8520 PMT QE	0.237 ± 0.019 ^a
R8520 PMT CE	0.75 ± 0.1
Expected PMT QE increase	1.12 ^b
Geometric light collection	0.4 ± 0.15 ^c
Total predicted	0.08 ± 0.03
Total observed	0.11 ± 0.03

^a at 300 K [74]

^b at 177 K, based on R8778 measurements [21]

^c uncertainty is due Teflon and Steel reflectivity, and λ_R .

electrons. The effect of electric field quenching on the scintillation yield must also be considered (discussed in Sec. 2.8), since the measurement of W_{ph} is valid for zero applied electric field. For the 0.73 kV cm^{-1} electric field used in XENON10, the scintillation yield (for a 662 keV electron) is quenched to 46% of its zero-field value [15]. The measured total $S1$ signal collection efficiency (photo-electrons per scintillation photon) can therefore be deduced as $2.31 \cdot 10^{-3} / 0.46 \cdot 21.6 = 0.108 \pm 0.03$. This result is lower than (but consistent with) with the total predicted light collection (Table 2.4). This could indicate that the reflectivity of Teflon and Steel used as input to the light collection Monte Carlo simulation may be too low. It could also indicate that the PMT collection efficiency (CE) is higher than 0.75, or that the QE increased by more than 12% at 177 K.

2.6 Energy Scales: keVee and keVr

The interaction of a γ or β in the Xe target leads to energy loss via interactions with atomic electrons. This is referred to as an “electronic recoil”, and is often abbreviated as “ER”. An uncharged particle such as a neutron is about 500 times more likely to scatter off a proton than an electron, based on kinematics. Similar considerations apply to a hypothetical 100 GeV WIMP in terms of its probability to scatter from an atomic nucleus as opposed to an atomic electron. This is referred to as a “nuclear recoil”, and is often abbreviated as “NR”. A nuclear recoil energy deposition is much less effective at generating electronic excitation – primary scintillation and ionization – than an electron recoil event of similar energy. This is generally referred to as “quenching”, and leads to the distinct energy scales keVee and keVr (already mentioned in Sec. 2.3), depending on the type of particle interaction. Quenching is the subject of on-going theoretical and experimental investigation, and is discussed further in Sec. 2.9 and Sec. 6.2.

For historical reasons, energy calibrations have tended to be based solely on the primary scintillation ($S1$) signal. Many Xe detectors (XMASS is a well-known example) do not measure the ionization signal at all. Even in dual-phase detectors, it is often desirable to understand the zero-field ($E_d = 0$) primary scintillation yield for calibration purposes. However, an $S1$ -based energy scale gives an (incorrectly) pessimistic




Figure 2.11: The XENON10 primary scintillation (*S1*) light yield in terms of PMT photo-electrons per keVee (Nphe/keVee) was measured at 30 keV (Xe k-shell x-ray), 40 keV (^{129m}Xe), 122 keV (^{57}Co), 164 keV (^{131m}Xe), 236 keV (^{129m}Xe) and 662 keV (^{137}Cs). All measurements are volume-averaged and referred to the center of the active volume. The metastable Xe states were produced during the AmBe neutron calibration, and from a dedicated activation of a small sample of Xe. These are described in Sec. 6 and Sec. 2.8 respectively. The errors are statistical only. The dashed line is a fit to the data, excepting the 236 keV point. The measured Nphe/keVee at 236 keV appears higher because the (^{129m}Xe is actually a 2-step de-excitation, with the emission of a 196 keV γ followed immediately (1 ns) by a 40 keV γ . The energy-averaged light yield that would be expected in this case based on the fit to the other points (the dashed line) is indicated by the red **x**.

assessment of the energy resolution [47]. Further, it is distinctly non-linear. If this is not properly accounted for, the background event rate will also be mistaken. Background event rates are typically measured in differential rate units (dru), with $1 \text{ dru} = 1 \text{ count keVee}^{-1} \text{ kg}^{-1} \text{ day}^{-1}$ (see Sec. 3.2.1). It is also typical to quote rates in $\text{mdru} \equiv 10^{-3} \text{ dru}$.

As a practical matter, the *S1* keVee energy scale is determined by the γ calibration data. The observed non-linearity is similar in shape to the measured scintillation yields from CsI and NaI scintillation [93], as shown in Fig. 2.12. The energy calibration is most directly expressed in terms of the total number of PMT photo-electrons in the full-energy peak, or Nphe/keVee. The keVee energy scale is typically used only as a reference, and often in such a case one picks a single energy point (for example, 122 keV from ^{57}Co) and scales all the data by the phe/keVee light yield for this point (which is incorrect, but linear and often useful). The measured light yield for several γ line energies is shown in Fig. 2.11 (compare Fig. 2.12).

Determination of the keVr energy scale is slightly more complicated. The measurement of the relative scintillation yield of nuclear recoils ($\mathcal{L}_{eff}$) has been made by several groups, and is discussed further in Sec. 2.9. This measurement is relative to the scintillation yield of a known γ line, and 122 keV from ^{57}Co is the typical reference energy in the literature. As shown in Fig. 2.20, there is significant uncertainty in $\mathcal{L}_{eff}$, particularly in the energy range of interest for WIMP Search with Xe detectors (roughly 5-25 keVr).




Figure 2.12: The scintillation yield in Thallium-activated CsI and NaI shows a relative yield (versus energy) similar to that observed in XENON10. The figure is taken from [93] Ch. 8, Fig. 8.8.

In Sec. 6.2 an alternative method is presented, which allows $\mathcal{L}_{eff}$ to be determined from the Americium-Beryllium neutron calibration data. However, in this analysis all results are presented in terms of keVr using a constant $\mathcal{L}_{eff} = 0.19$, with 122 keVee as the reference energy. This choice is inherently linear, is justified in Sec. 2.9 and corresponds to 0.98 $S1$ photo-electrons per keVr. The scaling is discussed in detail in Sec. 6.2. Typical ER and NR events which generate the same amount of primary scintillation ($S1$) are shown in Fig. 2.13.

2.7 Background Discrimination

2.7.1 Pulse-Shape Discrimination

Primary scintillation light in liquid Xe has two distinct decay time constants of about 2.2 ns and 27 ns (see Table 2.1) which are explained as originating from self-trapped singlet and triplet exciton states [97]. The tendency of incident particles to populate these states is found to be a function of linear energy transfer (LET) [82]. Ionizing β and γ particles, which are the dominant source of background, preferentially excite the triplet state. A singlet/triplet intensity ratio $I_s/I_t = 0.05$ was observed by [97]. The tendency to excite the singlet state increases with LET, so that for nuclear recoils caused by neutrons (and as predicted to be caused by WIMPs), $I_s/I_t = 1.6 \pm 0.2$ [82].

This property was used with some success by the Zeplin I collaboration [6]. There are two main problems. The decay time separation is only about $\times 10$, unlike in Ar (about $\times 300$). Additionally, measurement of either the singlet or the triplet decay time is obscured by the electron-ion recombination time $\tau_{rec} = 45$ ns [54]. This effect can be mitigated by applying an external electric field (similar to E_d in XENON10), at which point the available scintillation signal becomes quenched. The quenching is energy-dependent and




Figure 2.13: Typical 5 photo-electron (5 keVr) events in XENON10. The summed signal from all the PMTs is shown; a zoom on the $S1$ signal is inset. **(top)** A low-angle γ scatter from the ^{137}Cs calibration data has $\log_{10}(S2/S1) = 2.71$. **(bottom)** An elastic nuclear recoil from the AmBe neutron calibration data has $\log_{10}(S2/S1) = 2.47$.

LET-dependent, which introduces additional complications. Quenching is discussed further in Sec. 2.9.

The ability to discriminate by pulse shape is thus inherently limited, and is further limited by Poisson statistics at lower energies: where the expected interaction rate for WIMPs is highest, only a handful of photo-electrons are available in the primary scintillation signal (see Fig. 2.13 for an example). The time characteristics also require fast digitization (~ 1 GHz) to resolve the difference in intensity ratio, which increases the cost of a DAQ system (see Sec. 4.1.1).




Figure 2.14: The predicted electronic stopping power (LET) for different particle type in liquid Xe. The higher dE/dx for nuclear recoils in Xe leads to a denser track structure, which in turn leads to increased recombination. The result is a smaller ionization signal $S2$ compared with an electron recoil of the same energy.

2.7.2 Ionization-Scintillation Yield Discrimination

Empirically one observes that the ratio $S2/S1$ is lower (by about $\times 2$) for nuclear recoils than for electron recoils, as shown in Fig. 2.18. This is due to the nature of the energy loss in Xe (at the atomic scale) for the different types of particle. The predicted LET is shown in Fig. 2.14. In XENON10 analysis, the parameter $y = \log_{10}(S2/S1)$ is used to characterize the type of event, instead of $(S2/S1)$. This is due to the observation that the $\log_{10}(S2/S1)$ is well-fit by the hypothesis of a Gaussian distribution, which should be expected.

The rejection power at 50% acceptance for nuclear recoils was calculated from the fit parameters shown in Fig. 2.16 and Fig. 2.17. The rejection approaches 99.9% at the detector energy threshold of about 5 keVr, with an average better than 99.5% across the energy range of interest (roughly 5 – 25 keVr). The error bars indicate the statistical error on the nuclear recoil centroid and the electron recoil centroid and sigma. A separate analysis [130] of the XENON10 data showed a very similar rejection capability (Fig. 2.15, light red).

The energy-dependence of the discrimination parameter as shown in Fig. 2.15 can be understood from Fig. 2.18 (top). The electron recoil -2.58σ (99.5% discrimination) contour (green line) is found to lie above the nuclear recoil centroid (red line) below about 15 keVr. This is fortuitous from the point of view of dark matter detection, since as shown in Sec. 1.6.2 the WIMP recoil spectrum is predicted to be a falling




Figure 2.15: The measured single-scatter electron recoil rejection efficiency at 50% acceptance for single-scatter nuclear recoils, obtained from the parameters of a Gaussian fit to the ^{137}Cs electron recoil distribution (as shown in Fig. 2.16 and Fig. 2.17), in discrete energy bins. Results from a separate analysis [130] were similar and are shown in red (data was not available above 25 keVr). The rejection approaches 99.9% at the detector energy threshold of about 5 keVr, with an average better than 99.5% across the energy range of interest (roughly 5 – 25 keVr).

exponential.

The minimum in the measured ER rejection efficiency at about 25 keVr could be due in part to the non-linear ER scintillation yield. A peak in the ER scintillation yield at 30 keVr (~ 8 keVee) would give systematically lower values of $\log_{10}(S_2/S_1)$ over the energy range of the peak and would lead to the observed effect in Fig. 2.15. However, there are no measurements from γ line sources to support the hypothesis that the scintillation yield in Xe begins to drop again below 8 keVee, as it does for CsI and NaI (Fig. 2.12). The decrease in ER rejection efficiency also appears to be due to the behavior of the NR band centroid (Fig. 2.18, top). The centroid falls smoothly until about 20 keVr, at which point it (almost) plateaus until about 30 keVr before resuming a downward trend. The nuclear recoil calibration data are discussed further in Sec. 6.

2.8 Correlated Recombination Fluctuations

As can be seen from Eq. 2.4, the amount of signal available as scintillation light and ionization is a function of the applied electric field (E_d); for this reason, the trade-off between scintillation and ionization yields has been referred to as “macroscopic” anti-correlation [45]. This is quantified by the ratio $S(E_d)/S_{max}$ for scintillation, and $Q(E_d)/Q_{max}$ for charge collection. Note that S_{max} is obtained at $E_d=0$, while Q_{max} is obtained as $E_d \rightarrow \infty$. For the $E_d = 0.73 \text{ kV cm}^{-1}$ drift field in XENON10, $Q(E_d)/Q_{max} = 0.61$ and $S(E_d)/S_{max} = 0.54$ at 122 keV [14]. The yields are energy-dependent, and do not sum to 1 because each ratio is relative to its respective maximum possible signal. The charge collection $Q(E_d)$ is found to asymptote above $E_d = 10 \text{ kV cm}^{-1}$ [55], at which point the scintillation yield is about 0.30.

Even with E_d fixed, the partitioning of energy into ionization and scintillation also fluctuates on an event-by-event basis. This is shown in Fig. 2.19, and has been referred to as “microscopic” anti-correlation




Figure 2.16: The discrimination parameter $y = \log_{10}(S2/S1)$ versus $S1$, for single-scatter calibration data from the ^{137}Cs source (0.98 $S1$ phe per keVr, for $\mathcal{L}_{eff}=0.19$). Data are from a 5.4 kg (central) fiducial volume. The bin width is $dy = 0.03$. The χ^2 is calculated for the region where the Gaussian fit has $y > 1$.




Figure 2.17: The discrimination parameter $y = \log_{10}(S2/S1)$ versus $S1$, for single-scatter calibration data from the AmBe source (0.98 $S1$ phe per keVr, for $\mathcal{L}_{eff}=0.19$). Data are from a 5.4 kg (central) fiducial volume. The bin width is $dy = 0.03$. The χ^2 is calculated for the region where the Gaussian fit has $y > 1$.




Figure 2.18: **(top)** The discrimination parameter $y = \log_{10}(S2/S1)$ for single-scatter calibration data from ^{137}Cs source (electron recoils, blue) and AmBe source (nuclear recoils, red) versus energy. The $S1$ -based energy scale using a constant $\mathcal{L}_{eff} = 0.19$ corresponds to 0.98 $S1$ phe/keVr. Also shown are the centroids (same color as data points), the -2.58σ (99.5%, green) and -3.09σ (99.9%, blue dash) contours. **(bottom)** The 1σ width of the calibration bands from the fits in Fig. 2.16 and Fig. 2.17. The data were fit with a 5th-order polynomial (solid lines), which was then used to generate the 99.5% and 99.9% contours shown in the top plot.




Figure 2.19: **(top)** Anti-correlation of $S2$ (ionization) and $S1$ (scintillation photons) for 164 keV (^{131m}Xe , $\tau_{1/2} = 1.8$ days) and 236 keV (^{129m}Xe , $\tau_{1/2} = 8.9$ days) γ lines from the neutron-activated Xe calibration data. The small cloud below the 164 keV line is expected (from ^{133}Xe) as discussed in [121]. The $S1$ data were scaled assuming 10.0% total conversion efficiency from scintillation $\rightarrow$ phe (see Table 5.2.3). The $S2$ data were scaled assuming 24.0 phe per ionization electron (see Sec. 4.3.4). Only the bottom PMT array was used for the $S2$, to avoid saturated signals. Account was taken of the observation that 45% of the $S2$ signal appears on the bottom PMT array. **(right)** The same data shown in the more typical form $S1$ versus $\log_{10}(S2/S1)$. This provides a clear example of the contribution of recombination fluctuations to the electron recoil band width. It also shows that the degree of fluctuation is a function of energy: compare the angle of the major axis of the 164 keV and 236 keV ellipses. The color scale for both plots indicates $\log_{10}(\text{bin counts})$. The event rate at 164 keV is about 1000 dru.

[45]. It is perhaps more illuminating to refer to microscopic anti-correlation as “recombination fluctuations”. Recombination fluctuations are the limiting factor in the energy resolution of liquid Xe detectors for γ and β particles. It is more relevant to the present work that recombination fluctuations also lead to unavoidable widening of the $S2/S1$ bands for electron recoils, and are thus a limiting factor in the discrimination power of liquid Xe particle detectors.

In contrast to electron recoils, recombination fluctuations appear to be sub-dominant for nuclear recoils in liquid Xe. In this case, Poisson fluctuations dominate, with a smaller contribution from the single photo-electron resolution of the PMTs (PMTs are discussed in Sec. 3.3.5). The analysis which leads to these statements is not trivial, and a study of the subject is undertaken by E. Dahl in his thesis [47]. Nevertheless, the result can be seen in Fig. 2.18: the nuclear recoil band width decreases with energy, while the electron recoil band width remains roughly constant with energy.

The manifestation of recombination fluctuations (for a fixed level of macroscopic anti-correlation at $E_d = 0.73 \text{ kV cm}^{-1}$) is shown in Fig. 2.19, for γ lines at 164 keV and 236 keV. The data is from neutron-activated Xe, which was introduced into XENON10 after the WIMP Search data-taking was complete. The procedure to activate a sample of Xe for calibration use (via ^{252}Cf for example) has been described elsewhere [121]. The primary activity arises from meta-stable states: ^{131m}Xe (164 keV) decays with $\tau_{1/2} = 11.84$ days, and ^{129m}Xe (236 keV) decays with $\tau_{1/2} = 8.88$ days. Based on reaction cross-section and natural abundance, the 164 keV line is expected to be about $\times 2$ stronger than the 236 keV line immediately following exposure. This data was taken several weeks after the initial exposure.

2.9 Scintillation yield for nuclear recoils

2.9.1 Measurement with a mono-energetic neutron source

As discussed in Sec. 2.6, the scintillation yield from a nuclear recoil is about 20% of that for a 122 keV electron recoil (at $E_d = 0$). This observation is referred to as quenching, and often as “the quenching factor” in the literature. The term $\mathcal{L}_{eff}$ is preferred because “quenching factor” implies a measure of the relative light yield of a nuclear recoil, compared with an electron recoil of the same energy. In fact what is measured is the light yield of nuclear recoils compared to an electron recoil of a single, fixed energy. The standard electron recoil energy reference is 122 keV.

At least 4 measurements of $\mathcal{L}_{eff}$ using tagged neutron scattering have been published in the past 10 years, including a measurement made during XENON10 R&D using a neutron beam at the Columbia Nevis Laboratory [13]. These results are shown in Fig. 2.20. The principle of the measurement is relatively simple: neutrons of mean energy 2.4 MeV were produced by the $^3\text{H}(p,n)^3\text{He}$ reaction at the Columbia RARAF facility, and scatters in the Xe test cell were tagged at 6 different angles using a BC501A scintillator cell. If an incident neutron follows an unobstructed trajectory from its source to the BC501A scintillator, with only a single scatter in the Xe target along the way, then the recoil energy in the Xe is given by

$$E_R = T_n \frac{4M_n M_{Xe}}{(M_n + M_{Xe})^2} \frac{(1 - \cos \theta)}{2} \quad (2.6)$$

where E_R is the nuclear recoil energy, T_n is the incident neutron energy and θ is measured in the center of mass frame. It then remains only to compare the measured scintillation yield at each energy ($S1_{nr}$ in Eq. 2.7) with the scintillation yield measured at 122 keV (L_y in Eq. 2.7). It is typical to quote $\mathcal{L}_{eff}$ in terms of the zero-field ($E_d = 0$) scintillation yield. As discussed in Sec. 2.8, $S_e = 0.54$ is the scintillation yield for a




Figure 2.20: The measured scintillation efficiency for nuclear recoils in liquid Xe at zero applied electric field ($E_d = 0$) from XENON10 R&D (blue stars) [13], along with measurements from [8] (green circles), [20] (red triangles) and [43] (pink crosses). Results from [36] indicate $\mathcal{L}_{eff} > 0.35$ in the energy range 35 – 70 keVr and are not shown. The dashed line at $\mathcal{L}_{eff} = 0.19$ indicates the most reasonable zero-slope fit to the available data. The shaded blue indicates the region of highest sensitivity to scalar WIMP elastic scattering.

122 keV γ at $E_d = 0.73 \text{ kV cm}^{-1}$, relative to $E_d = 0$. The scintillation yield of nuclear recoils is found to be only mildly quenched ($S_n = 0.93$). Thus $\mathcal{L}_{eff}$ is defined as:

$$\mathcal{L}_{eff} = \frac{S_{1nr}}{L_y \cdot E_R} \cdot \frac{S_e}{S_n} \quad (2.7)$$

In practice such a measurement is complicated by neutrons which scatter in other materials or travel directly to the BC501A cell. Depending on the total light collection efficiency of the Xe detector, Poisson fluctuations and other threshold effects may further complicate the measurement. Given the available measurements of $\mathcal{L}_{eff}$ (Fig. 2.20) the most reasonable zero-slope fit is $\mathcal{L}_{eff} = 0.19$. This is shown as a dashed line in Fig. 2.20.

2.9.2 Measurement with a broad-spectrum neutron source

Importantly, only two of the measurements shown in Fig. 2.20 extend to the WIMP Search region of interest, i.e. $E_R < 30 \text{ keVr}$. In Sec. 6.2, it is shown that the scintillation yield for nuclear recoils $\mathcal{L}_{eff}$ can also be determined from the XENON10 neutron calibration data. This analysis is presented in Sec. 6.2 (to be published [145]).

Chapter 3

INSTALLATION OF XENON10 AT LNGS

“I am lying because I know myself as surely as two times two makes four, that it is not underground that is better, but something different, quite different, for which I long but which I cannot find! Damn underground!”

— Fyodor Dostoevsky

3.1 XENON10 installation at LNGS

The XENON10 dark matter direct detection experiment was installed in the (underground) Laboratori Nazionali del Gran Sasso (LNGS) in the Abruzzo region of Italy, about 120 km north east of Rome. The LNGS underground installation provides an average 3100 meters water equivalent (m.w.e) of rock overburden [112] for reduction of cosmic ray muon flux. The XENON10 detector is shown in Fig. 3.4, installed in the shield in the underground cavern at LNGS. The Gran Sasso massif, part of the Apennine Mountains, is shown in Fig. 3.1. An access-controlled turn-off from route S.S. 17 (state road 17), about half way through the tunnel under the massif, allows easy access for scientists and visitors via car or shuttle bus, and for deliveries via truck. A schematic of the underground cavern is shown in Fig 3.2, with the S.S. 17 in the lower right.



Figure 3.1: The Gran Sasso Massif, providing 3100 meters water-equivalent shielding from cosmic rays. Assergi (an ancient Roman village) is in the foreground, and the LNGS above-ground offices in the background (at right). The tunnel under the massif begins in the lower right corner of the photo; a schematic of the underground labs is shown in Fig. 3.2.




Figure 3.2: Schematic of the LNGS underground laboratory. Entrance is via auto tunnel from the Teramo side. XENON10 is installed in an alcove of the interferometer tunnel, which presented some engineering challenges during installation due to the low ceiling height. The remote location later proved fortuitous because of the continuous construction in the main halls during the XENON10 calibration and science data-taking.



Figure 3.3: Initial assembly and testing of XENON10 in March 2006 occurred in a temporary location 50 m from the final installation site. XENON10 operated in this location, unshielded except for the rock overburden, until July 2006 when the shield installation was completed. From left: R. Gomez (Rice University), A. Ferella (LNGS) and the author.





Figure 3.4: XENON10 Installed in Shield in LNGS underground Laboratory. **(left)** XENON10 securely mounted to the sliding shield door, with the door partially closed. The 20 cm HDPE and 20 cm Pb are evident. **(right)** The Xe purification system is mounted to the outside of the shield door (which is fully closed). The DAQ box is visible to the right, and the Xe lines (covered in Al foil) along the top left.

3.1.1 Summary of operation at LNGS

This section provides a general overview of about 1 year of operation. XENON10 was installed in a temporary underground room in late March 2006, where it was fully assembled for testing of the detector and subsystems. Purification of the Xenon, as well as calibration with external gamma sources, and with LED (for single photo-electron response) proceeded in parallel with construction of the shield from April → July 2006, when the shield was completed. The detector was then installed into the shield, and the experiment was commissioned at the end of July 2006.




Figure 3.5: Summary of the XENON10 data accumulation beginning in August 2006. Calibration data (shown in green and blue) was scaled to WIMP Search equivalent live-days based on the trigger rate. Figure courtesy of L. deViveiros.

With the detector operating in the shield, the (expected) limiting radioactive background was the ~ 25 ppm Kr content of the Xe sample; the spectrum was recorded, and on July 29th the Xe was exchanged for a second Xe sample obtained from Spectra Gases and certified to have < 10 ppb Kr. This factor of ~ 1000 reduction in the ^{85}Kr rate marked the beginning of low background operation, with the event rate in the center of the active Xe target at < 2 counts $\text{kg}^{-1} \text{keV}^{-1} \text{day}^{-1}$ (dru).

For the next 3 weeks, the DAQ trigger system was tested and optimized for the lowest possible threshold. Using an $S1$ -based trigger with $n \geq 3$ coincidence on 3 bottom-array PMTs, a higher than expected rate of $S1$ -only events was observed. A stable configuration was achieved by August 23rd with the $S1$ trigger at $n \geq 6$ coincidence, an $S2$ -based “catch-all” trigger, and a high-energy veto. With a global trigger rate of about 4.5 Hz, the first science data taking began. This data run, referred to as WIMP Search 2 (WS2), lasted until the 27th of September. A timeline summary of the data-taking is shown in Fig. 3.5.

The primary motivation for employing an $S1$ -based trigger was that it would simplify the event reconstruction: an $S2$ is large and easy to identify, and occurs $< 80 \mu\text{s}$ after an $S1$ by design. However, even with $n \geq 6$ coincidence, the rate of $S1$ -only events was too high (> 10 Hz depending on the trigger threshold). The $S1$ trigger was removed in early October, marking the beginning of WS3 and “blind” data-taking. In mid-November 2006, high statistics γ calibration data were taken (equivalent to about 90 live days of WIMP Search data), using a ^{137}Cs source with a high-energy veto to limit the accumulated data to the WIMP Search region of interest (< 20 keVee). This was followed by a neutron (AmBe source) calibration data set




Figure 3.6: The ~ 600 dru event rate immediately following installation in the shield (July 2006) was dominated by β decay from ^{85}Kr ($1 \text{ dru} = 1 \text{ count kg}^{-1} \text{ keVee}^{-1} \text{ day}^{-1}$). A Monte Carlo simulation of 10 ppm Kr contamination in the Xe is shown (red dash). The trigger scheme was still under study in July 2006, and the threshold for this data was set rather high. This led to the roll-off in event rate below 100 keVee. The energy scaling of 2.1 phe/keVee was obtained from ^{137}Cs calibration, and before the start of WIMP Search science data increased to 2.2 phe/keVee due to a modification of the PMT gain configuration. As described in Sec. 2.6, it is not correct to use this energy scaling across such a large range. The error introduced by this approximation could have over-predicted the Kr contamination by about $\times 2$.

on December 1st. In both cases the sources were placed inside the shield, near the center of the cryostat. The sources could be inserted through a small, specially-designed sealed tube (6 mm $\varnothing$) which penetrated straight through the shield. In this way, the air-tight seal for the Rn purge was not compromised.

WIMP Search 4 began immediately after the neutron calibration, and ended on February 14th, at which point a small amount of neutron-activated Xe was introduced into the detector. This provided two homogenous internal gamma lines (see Sec. 2.8) for calibration of the position-dependence of the signals. Periodic calibration data were performed with external ^{57}Co and ^{137}Cs sources during WS2, WS3, and WS4 to monitor the Xe purity, study the position dependence of the detector and assess the stability of the detector. Prior to the final analysis, about 20% of the accumulated “blind” data from WS3 and WS4 were “un-blinded” to as a final check of the performance of the analysis algorithms.

3.2 Shielding and radioactive backgrounds

The XENON10 detector is housed inside a cubic “castle” of 20 cm high-density polyethylene (HDPE) inside of 20 cm Pb, which shield against the ambient electromagnetic and neutron backgrounds. The internal cavity dimensions are $90 \times 90 \times 107.5$ cm (x, y, z), and access is obtained by a sliding wall (on rails), on which the detector cryostat is securely mounted (see Fig. 3.7). The total mass of Pb (HDPE) is 33.42 (1.56) tonne.

The specific considerations which lead to the choice of thickness and orientation of each component will be described in the following sections. The Pb and HDPE are supported by a steel frame, with additional steel support struts supporting the detector on the movable wall. This is shown schematically in Fig. 3.7.




Figure 3.7: Engineering drawing of the XENON10 shield, showing the steel frame (grey), Pb bricks with “standard” ²¹⁰Pb activity of 560 Bq/kg (yellow), Pb bricks with low ²¹⁰Pb activity of 17 Bq/kg (purple), and HDPE (cyan). The door is shown in its open position, with the detector cryostat (red) mounted on steel bars (dark grey). Drawing by L. Tatananni (LNGS Engineering staff).

3.2.1 External sources of electromagnetic background

At 3100 m.w.e, the muon rate is about $1 \text{ m}^{-2} \text{ hr}^{-1}$, a factor of nearly 10^6 reduced from the sea-level rate. Nevertheless, an unshielded detector experiences a raw background γ rate of about $10^4 \text{ counts kg}^{-1} \text{ keV}^{-1} \text{ day}^{-1}$ [80, ?] in the energy region of interest for dark matter detection ($< 10 \text{ keV}$). As mentioned in Sec. 2.6, a standard unit of background event rate is the differential rate unit (dru), with $1 \text{ dru} = 1 \text{ count/kg/keV/day}$. This is also sometimes referred to as dru electron-equivalent, or druee; in this work it dru always implies electron-equivalent and is simply referred to as “dru”. The energy spectrum of background events in this range is essentially flat, and due primarily to low-angle Compton scatters from high-energy photons. Major sources include the naturally occurring isotopes ²³⁸U ($\tau_{1/2} = 4 \times 10^9 \text{ y}$), ²³²Th ($\tau_{1/2} = 1 \times 10^{10} \text{ y}$) and ⁴⁰K ($\tau_{1/2} = 1 \times 10^9 \text{ y}$). The decay of these isotopes leads to the emission of penetrating, high-energy γ rays ($> 1 \text{ MeV}$), high-energy β particles (which create additional, lower-energy γ ’s via bremsstrahlung) and α particles. This leads to an unacceptable situation when compared with a signal whose predicted (and experimentally constrained [7]) interaction rate is $< 10^{-2} \text{ counts kg}^{-1} \text{ keV}^{-1} \text{ day}^{-1}$.

In principle any radio-pure material can be used for shielding external γ and β flux, with high-Z materials offering superior performance for equal shield thickness. Commonly used materials include Cu and Pb. In

particular, Cu can offer extremely high radio-purity and is often used as an inner liner. The 20 cm of Pb used in XENON10 affords > 4 orders of magnitude reduction for the highest energy 2.6 MeV photons from ^{232}Th decay (the mean free path for a 2.6 MeV γ in Pb is $\lambda = 2.2$ cm). A detailed Monte Carlo simulation of activity from U and Th in the rock showed a reduction of about 3×10^{-6} , and the raw rate inside the shield from external radioactivity was predicted to be < 0.05 dru (50 mdru). The actual background event rate in XENON10 during WIMP Search was more than $\times 10$ higher, and was attributable to activity from the detector itself (inside the shield). Residual U and Th impurities in the Pb were also sub-dominant. The major consideration for Pb is ^{210}Pb , which is typically found in Pb ore at between 10–1000 Bq/kg; it decays eventually to ^{206}Pb with the emission (among other things) of a 1.16 MeV β . This is discussed further in Sec. 3.2.6.

3.2.2 External sources of neutron background

The flux of neutrons (0 – 10 MeV) in the rock and concrete in the LNGS underground laboratory have been measured many times and the results vary. This is discussed in [161], who performed detailed simulations of the raw neutron production rates and subsequent neutron flux. The production rates were calculated based on measurements of the U/Th concentrations near the three main halls of LNGS. The prediction is a roughly equal rate for (α, n) reactions and spontaneous fission. This is summarized in Table 3.1:

Table 3.1: Calculated raw neutron production rates (neutrons $\text{yr}^{-1} \text{ g}^{-1}$ of rock) at LNGS from spontaneous fission and (α, n) reactions. Measurements of the concentrations of U/Th [ppm] as used in the calculation are also shown. Hall B is very similar to Hall C. Data from [161].

	<i>measurement</i>		<i>simulation</i>	
	^{238}U [ppm]	^{232}Th [ppm]	Fission [n/y/g]	(α, n) [n/y/g]
Rock near Hall A	6.80 ± 0.67	2.17 ± 0.074	3.54	4.38
Rock near Hall C	0.66 ± 0.14	0.066 ± 0.025	0.34	0.38
LNGS concrete	1.05 ± 0.12	0.656 ± 0.028	0.55	0.51

The simulated neutron flux in the experimental halls (using the raw production rates shown in Table 3.1) was found to be dominated by (α, n) reactions. The contribution from spontaneous fission of ^{238}U , ^{235}U and ^{232}Th was about $\times 5$ lower for 1 MeV neutrons (decreasing to more than $\times 10$ lower for neutron energies above 4 MeV). Interestingly, the flux was found to be independent of location (i.e. Hall A versus Hall C) to better than a factor of $\times 2$. The reason is the presence of 0.3 – 1.0 m of concrete lining the walls of the LNGS caverns, which moderates the neutron flux. The typical $(1/e)$ depth at which 63% of the total neutron flux originated was found to be ~ 7 cm for neutron energies $T_n > 1$ MeV. The spectral shape for neutrons from the cavern was found to be essentially a falling exponential, with some residual (weak) features due to the discrete energies of the alpha particle energies (see Fig. 3.8); more than 90% of the neutron flux as it enters the cavern has incident energy $T_n < 5$ MeV.

The cavern neutron flux has been measured in the range 1 – 10 MeV to be $(0.42 \pm 0.12) \times 10^{-6} \text{cm}^{-2} \text{s}^{-1}$ in Hall C [19], and $(0.70 \pm 0.14) \times 10^{-6} \text{cm}^{-2} \text{s}^{-1}$ in Hall A [30]. Considering neutron energies $T_n < 1$ MeV as well, the total flux in the range 0 – 10 MeV was found to be $(3.81 \pm 0.11) \times 10^{-6} \text{cm}^{-2} \text{s}^{-1}$ [30]. This agrees to within 2% with the calculations and Monte Carlo simulations of [161], assuming 8% water content in the concrete (open diamonds in Fig. 3.8).




Figure 3.8: The spectrum of neutron energies entering the experimental halls at LNGS, from the simulations of [161]. The flux is highest in Hall A, assuming 8% H_2O content of the concrete (filled black circles); the corresponding integral flux 1 – 10 MeV is $(0.58 \pm 0.13) \times 10^{-6} \text{cm}^{-2} \text{s}^{-1}$. The same flux is shown, considering fission reactions only (open diamonds), and assuming 16% H_2O content of the concrete (x). The flux in Hall C, assuming 8% H_2O content of the concrete, is found to be $< \times 2$ lower than the flux in Hall A (open circles).

The primary function of neutron shielding material is to moderate the incident neutron energy via collisions. For example, once an incident neutron energy is moderated to $T_n = 200 \text{ keV}$, the typical recoil energy deposited from $\text{Xe}(n,n)\text{Xe}$ elastic scattering is $E_R = 3 \text{ keVr}$. Once moderated to a thermal energy, neutrons are either captured (which typically releases $> 1 \text{ MeV}$ of γ flux) or they decay. From Eq. 2.6, it is clear that neutron moderation is most effective if the shielding material has a mass similar to the neutron. For example, while the average kinetic energy loss for a neutron scattering from Xe is about 1.5%, for a scatter on H an incident neutron loses half its initial energy. Common neutron shielding materials therefore include H_2O and hydrocarbons such as HDPE. Concrete is not typically used, but its high content of Ca and Si make it a nice bonus as a lining material for an underground lab – provided that its internal radio-purity exceeds that of the surrounding rock. The cross section for a 1 (5) MeV neutron in HDPE is $\sigma = 6$ (4) barn [63], with a corresponding mean free path of $\lambda = 4.2$ (6.3) cm. Thus, for 20 cm of HDPE one expects about $\times 100$ reduction of incident neutron flux (for 1 MeV neutrons), and about $\times 60$ reduction for 5 MeV neutrons. In fact the 20 cm of HDPE as implemented in XENON10 provides an additional factor of $\times 10$; this is discussed further in Sec. 3.2.5.

3.2.3 Cosmic ray background

High-statistics measurements from MACRO [29] find a muon flux of $1.04 \pm 0.02 \text{ m}^{-2} \text{ hr}^{-1}$ in Hall B of LNGS. The average muon flux at LNGS has also been assessed at $0.92 \text{ m}^{-2} \text{ h}^{-1}$ [112]. The latter was

obtained by considering the angular dependence of the muon energy (since the rock overburden is not flat), and comparing experimental data from LVD with FLUKA simulations; the equivalent flat overburden was then estimated to be 3100 ± 200 meters water equivalent (m.w.e.). The average muon energy is a function of depth, and has been measured to be $E_\mu = 270 \pm 18$ GeV [9] at LNGS.




Figure 3.9: An example of a muon track in XENON10, from a single, unsaturated PMT. Typically many PMTs saturate so the summed waveform is not a good reflection of the track structure. This is discussed in Sec. 5.2.3. The ionizing track (*S2*) begins immediately after the *S1* (shown inset, zoom in μs only), indicating entry through the top of the detector. The track length indicates that the muon traversed the entire detector, so its zenith angle was $\theta < 53^\circ$.

A total of 69 tracks with sufficient length and signal size to be consistent with muons were found in 58.6 live-days (1406 hr) of WIMP Search data. The (x,y,z) coordinates of the muon entry in the XENON10 fiducial volume were determined from the tracks. Of the 69 tracks, 47 were consistent with muons which entered the active 20 cm ϕ target through the top of the detector (at any angle). These events are clearly distinguishable because the ionizing track begins immediately following the primary scintillation. The muon flux at the XENON10 installation site was calculated to be

$$R_\mu = \frac{47}{\pi(0.10 \text{ m})^2(1406 \text{ hr})} = 1.06 \pm 0.03 \text{ m}^{-2} \text{ hr}^{-1}, \quad (3.1)$$

in good agreement with the previously reported values. The uncertainty was estimated assuming a 1% uncertainty in the total live time, and a 1% uncertainty in the effective area subtended by the detector. The effective area is a reflection of the small uncertainty that arises in determining the delay dt between the muon *S1* signal and the start of the track. Analysis of muon tracks is discussed further in Sec. 7.1.

3.2.4 Neutron production from cosmic rays

Fast neutrons are produced in underground laboratories by a number of processes as muons traverse the surrounding rock. At LNGS depths secondary processes such as π^- spallation (rather than direct muon spallation) dominate [164]. The muon-induced neutron flux from the cavern is about $5 \times 10^{-10} \text{ cm}^{-2} \text{ s}^{-1}$ in the range 0–50 MeV; this falls by 2 orders of magnitude in the energy range 0–500 MeV, and then by about a decade per GeV up to several GeV [112]. Measurements of the muon-induced neutron energy spectrum

by the KARMEN collaboration show that the differential spectrum falls even faster for $T_n < 15$ MeV, and FLUKA simulations are in reasonable agreement [164]. The steep slope of the spectrum is fortuitous, because the XENON10 shielding has little effect on the neutron flux for neutrons with energy above about 30 MeV (see Table 3.2); at the same time, neutrons up about 300 GeV have appreciable probability to deposit WIMP-like recoil energy ($E_R < 30$ keVr).

Table 3.2: Interaction length (λ_I) and raw attenuation of neutrons for 20 cm of HDPE moderator. Data from [63]. The actual attenuation in XENON10 is higher due to neutron elastic scattering on the Pb; see Sec. 3.2.5

E_n [MeV]	λ_I [cm]	Attenuation
1	4.17	$\times 121$
5	6.25	$\times 25$
10	10	$\times 7.4$
30	19.2	$\times 2.8$
100	59.5	$\times 1.4$

Fast neutrons are also produced by muons traversing the shielding material. The production rate at LNGS is estimated to be $5.5 \times 10^{-10} \text{ cm}^{-3} \text{ s}^{-1}$ in Pb, and $4.8 \times 10^{-11} \text{ cm}^{-3} \text{ s}^{-1}$ in HDPE [112]. Considering an average path length in the XENON10 shield (for vertically directed muons) of 101 cm (96 cm) in Pb (HDPE), the raw rate of muon-induced neutron production (integrated over all neutron energies) from the Pb (HDPE) is $5.5 \times 10^{-8} \text{ cm}^{-2} \text{ s}^{-1}$ ($4.6 \times 10^{-9} \text{ cm}^{-2} \text{ s}^{-1}$).

Because of the severe drop in moderator efficacy for neutrons with energies above about 30 MeV, such neutrons are often referred to as “punch-through” neutrons. The most effective counter measures are (a) choosing a deeper site, and (b) deploying a muon veto. For example, the Sudbury mine in Canada (6 km.w.e.) has a muon-induced flux that is $\times 100$ lower than the flux at LNGS. The experimental reach of XENON10 is not limited by the muon-induced neutrons (see Sec. 3.2.5), however, next-generation liquid Xe experiments with fiducial masses > 100 kg will want to consider both (a) and (b).

3.2.5 Shield performance

Electromagnetic shielding performance

As discussed in Sec. 3.2.1, the external gamma flux was attenuated to < 50 mdr (inside the XENON10 shield) by the 20 cm of Pb. The internal contamination of the shielding materials are summarized in Table 3.3. This discussion focuses on the contribution from ^{210}Pb , which is summarized in Table 3.4. Less than 1% of the count rate due to ^{210}Pb bremsstrahlung from the 1.16 MeV beta decay originate outside the inner 2 cm of Pb [80]. With the 5 cm of low-activity Pb liner (from Fonderies de Gentilly; see Table 3.4) used in XENON10, Monte Carlo studies show that (a) the count rate in the detector from ^{210}Pb activity the outer 15 cm (from Plombum) is entirely sub-dominant, and (b) the count rate in the outer 1 cm of the active Xe due to ^{210}Pb bremsstrahlung originating in the 5 cm of liner Pb is as high as 0.9 dru [50].

Due to Xe self-shielding, the ^{210}Pb contribution to the background event rate in the WIMP Search 5.4 kg fiducial target is sub-dominant to other internal detector contributions. This is shown in Fig. 3.17. Other internal contributions are discussed in Sec. 3.2.8. The self-shielding is also discussed further in Sec. 3.2.7.

The unexpected presence of ^{207}Bi and ^{194}Hg in the outer Pb do not significantly impact the event rate in the detector: ^{194}Hg emits a 40 keV β that is effectively attenuated. The 1064 keV γ from ^{207}Bi has a non-trivial penetration length, but Monte Carlo studies show that the rate from the measured radio-impurity (240 mBq/kg) lead to an event rate less than $\times 0.1$ compared with the rate from the ^{210}Pb bremsstrahlung.

Neutron shielding performance

As a consequence of the incident neutron energy spectrum (shown in Fig. 3.8), the neutron recoil spectrum in the XENON10 detector is also predicted to be an essentially featureless exponential – similar to the expected WIMP spectrum (Sec. 1.6.2) – but with a steeper (by about $\times 3$) log-slope of 20 keVr / decade. From Eq. 2.6 it should be clear that it is only necessary to moderate an incident neutron’s energy down to about 100 keV, at which point the neutron would be unable to generate a detectable signal in the XENON10. The placement of the Pb outside the HDPE is found to attenuate the raw flux by about an order of magnitude more than that predicted for the 20 cm of HDPE alone.

This increase in shielding efficiency is due to neutron inelastic scattering on the Pb. The scattering reduces the incident neutron energy, thus increasing the efficacy of the HDPE moderator. This fact was also observed by [42], and is discussed further in [50]. Taking the measured neutron flux of $2 \times 10^{-6} \text{ cm}^{-2} \text{ s}^{-1}$ and an energy spectrum characteristic of equal parts (α, n) and fission neutron sources, Monte Carlo simulations of the neutron rate in XENON10 [50] predict < 0.1 neutron/2 months/10 kg with recoil energy $5 < E_R < 25$ keVr. The simulations predict a similar rate for muon-induced neutron production in the Pb shield. As discussed in Sec. 3.2.4, neutrons from muons in the surrounding rock are expected at $\times 10$ below this.

3.2.6 Shield design

Engineering details

The XENON10 shield was designed to be as simple as possible, without compromising the effectiveness or accessibility. All surfaces in the construction were fit snugly to ensure that trapped air (which can contain Rn gas) was eliminated from the inside of the structure. The Pb was furnished in $20 \times 10 \times 5$ cm bricks with < 1 mm tolerance, which were simply stacked in place (machined as needed for a precise fit), and sandwiched between the steel frame and the HDPE. The HDPE was furnished in sheets with thickness 2.5 cm or 3.75 cm and therefore could be used in such a structural capacity. Thicker HDPE may have been preferable from a construction point of view. However, the cost (\$/kg) was found to increase rapidly with sheet thickness. A limited number of steel bolts (about 0.1 kg) were passed from the outside of the steel frame to the inside of the cavity, in order to ensure the rigidity of the HDPE (and thus Pb) placement. The HDPE itself required a total of 15 screws (< 0.1 kg), in addition to bi-adhesive tape to keep the sheets tightly pressed together. All interior (to the Pb) surfaces were cleaned with ethanol prior to installation.

To ensure the structural integrity of the ceiling, the inner 5 cm of Pb was cast into 4 plates ($65 \times 65 \times 5$ cm each), which were suspended from a 25 mm steel plate. This is shown in Fig. 3.10. The steel plate in turn was supported directly by the steel frame with 16 bolts (see also Fig. 3.7), and the remaining 15 cm of common Pb were stacked on it. Samples of the internal steel components were screened concurrently with the construction. The ~ 200 kg detector itself was supported on steel bars which were an integral part of the shield door construction. Most of the steel showed better than expected radio-purity, particularly in ^{60}Co , and so would not contribute significantly to the background rate. This is summarized in Table 3.3.




Figure 3.10: **(top)** A view from the top of the shield during construction, showing the steel plate supported from the central I-beam. The view angle of the top photo is rotated by 90° with respect to that of the lower photo. **(bottom)** The author ponders solitary confinement as the shield nears completion. Two (of the four) cast Pb plates are visible edge-on, bolted to the 25 mm steel plate above the structure. A view of the one of the bolts which hold the cast Pb plates to the steel plate is shown inset. A small gap in the ceiling layer of HDPE is visible; this occurred at the 4 edges by the door, and was rectified with a small number of steel screws.



Figure 3.11: The entire shield was constructed on a 15 cm platform of high molecular weight polyethylene (“PE500”), which was structurally supported (sandwiched) by steel plates. The surface was leveled with steel shims. This was a precautionary step to ensure that neutrons could not be a limiting source of background. In contrast, additional shielding on the other 5 sides could have been added (but was not) as an after thought. The contractors were not always pleased by our exacting design requirements. This photo was taken shortly after discussing the requirements for cleaning all internal surfaces prior to installation. One of the more polite responses was “non ci credo!”




Figure 3.12: **(left)** The shield door (inside) during construction. The steel bars which support the detector are just below the conduit for the cryostat vacuum. The PMT signal cables (white) had already been cleaned, wrapped and installed in the shield door. Cu tubing feed-throughs for the PTR are visible next to the PMT cables. The Pb bricks are stacked to about $\frac{3}{4}$ height. The interlock “stairs” of the Pb layers (to prevent line-of-sight γ penetration) is evident. **(right)** The shield door (outside, lower) during construction. The other side of the steel support bars are visible, along with the cryostat vacuum conduit. A single steel bolt through the Pb was required to keep the HDPE sheets pressed tightly together. The gear for opening and closing the shield had a high ratio, so that opening the shield was easily accomplished in about 60 s.

Radio-purity

Table 3.3: Radioactivity of the materials used to construct the XENON10 shield. Activities are quoted as 90% C.L. upper limits [mBq/kg]. Activity of ^{210}Pb is summarized in Table 3.4. Activities of the external steel frame are not tabulated.

Material (<i>required by</i>)	Mass [kg]	U ^a	^{232}Th	^{40}K	^{60}Co	^{207}Bi	^{194}Hg
Common Pb	26440	187	9	20	1	240	120
Low-activity Pb	5780	60	8	28	0.2	—	—
HDPE	1560	145	14	64	—	—	—
Steel plate (<i>ceiling</i>)	330	165	18	42	5	—	—
Steel bars (<i>detector support</i>)	mass	13	14	28	13	—	—
Steel bolts (<i>cast Pb / steel plate</i>)	2	283	62	41	8	—	—
Steel screws (<i>HDPE</i>)	< 0.5	113	14	28	13	—	—
Bi-adhesive tape (<i>HDPE</i>)	< 0.5	272	13	31	—	—	—

^a Contains contributions from both ^{235}U and ^{238}U .

As mentioned in Sec. 3.2.1, care was taken to minimize the natural radioactivity contained in the Pb (and all shield construction materials). The measured 90% C.L. upper limits on isotopic activity (in mBq/kg) of materials used in the XENON10 shield are summarized in Table 3.3 (excluding ^{210}Pb activity). The cost of Pb is non-trivial considering the requisite mass, and is significantly higher for lower ^{210}Pb activity. This is summarized in Table 3.4. Additionally, the cost of Pb rose by $\times 2$ between the December 2005 procurement and the decommissioning of XENON10 in August 2007 (and has continued to rise). The virgin HDPE sheets were purchased from K-Mac Plastics at a cost of \$6/kg (all costs quoted are exclusive of carriage charges). It is worth noting that recycled HDPE is available for as little as \$1/kg, but this is not suitable for internal (to the Pb) use because the radio-purity can be orders of magnitude worse than that of virgin HDPE.

Table 3.4: Cost of Pb (per kg) compared with ^{210}Pb activity (per kg) as available in December 2005. The quoted cost assumes the Pb is cast in brick form, $20 \times 10 \times 5$ cm with < 1 mm tolerance.

Cost [\$ / kg]	^{210}Pb Activity [Bq/kg]	
	Stated by Vendor	Measured at LNGS ^c
\$2.50	700 ^a	560 ± 90
\$6.05	< 30 ^b	17 ± 5
\$19.84	< 3 ^{a,d}	N/A^d

^a Source: Plombum, Poland

^b Source: Fonderies de Gentilly, France

^c Using HPGe counter (M. Laubenstein)

^d *Not used* in XENON10

The U/Th/K activity in Pb has been reported to be < 1 mBq/kg (corresponding to 81 ppt U / 246 ppt Th / < 33 ppb K) for a wide variety of Pb sources [80]. Measurements at LNGS showed U/Th/K activities of < 60/8/28 mBq/kg (Fonderies de Gentilly) and < 187/9/20 mBq/kg (Plombum), which nevertheless are predicted to contribute about 3 orders of magnitude less than the ^{210}Pb to the event rate; this is summarized in Table 3.3. The higher than expected U activity in the Pb from Plombum also coincided with a similar

level of activity from ^{194}Hg and ^{207}Bi ; these isotopes are indicative of exposure to high-energy neutrons [80], or > 10 MeV protons [66] and were not anticipated. As described in Sec. 3.2.5, they did not significantly impact the event rate in XENON10.

Rn exclusion

The concentration of ^{222}Rn in the LNGS cavern was measured with a DurrIDGE RAD-7 to be as high as 140 Bq/m^3 in the location of the XENON10 installation. This is about $\times 2$ higher than measured in other subterranean sites, and is somewhat surprising considering the $40 \times 10^3 \text{ m}^3/\text{hr}$ flow of external air into the LNGS underground laboratory; a partial explanation is that the XENON10 installation is in a far corner, away from the main halls and ventilation ducts. Suitable Rn-exclusion measures were taken. With the shield door closed, a latex gasket made an air-tight seal such that the internal cavity could sustain a slight over-pressure (about 1.2 atm). N_2 gas (boil-off from a liquid N_2 dewar) was vented directly into the top of the shield inner cavity via stainless steel tubing, and forced to exit through the RAD-7 monitoring system at the bottom of the inner cavity. The flow rate was about 0.5 standard liters per minute (SLPM). With this arrangement, the Rn concentration inside the shield was measured to decrease to a stable value of $< 5.5 \pm 1 \text{ Bq/m}^3$ in about 30 hours, as shown in Fig. 3.13. This is actually an upper limit set by the measurement sensitivity of the RAD-7.




Figure 3.13: The Rn concentration (Bq/m^3) inside the shield cavity, with the LN boil-off purge in place. $T = 0$ corresponds to the closing of the shield door; a stable value of $< 5.5 \pm 1 \text{ Bq/m}^3$ is achieved after > 30 hours. Plot courtesy of S. Fiorucci.

Instrumentation feed-through

The PMT signal cables, HV cables, tubing for the Xe filling and circulation, N_2 (purge gas), He (for the PTR) and all other electronic instrumentation cables were mounted (buried) in the movable shield door. With the exception of the 6 mm $\varnothing$ blind tube for inserting the calibration source and the 6 mm $\varnothing$ tubes for the 2 leveling feet of the cryostat, all cables and tubing were fed through a z-bend (dog-leg) in the Pb. The z-bends were designed so that at least 5 cm of Pb shielding would be seen by an incident external γ , as shown in Fig. 3.14 (left). A total of 5 z-bend feed-throughs with diameter 40 mm were required, plus a 50 mm z-bend feed-through for the cryostat turbo pump. The HV cable for the cathode volatge had a




Figure 3.14: **(top left)** Top-down view of the HV cable for the cathode voltage, embedded in the shield door during construction. The j-bend radius of curvature was limited by the specification of the cable (all other feed-throughs used a z-bend). **(top right)** Top-down view of the PMT signal cables embedded in the shield door during construction. The air gaps were filled with Pb shavings from the 17 Bq/kg Pb. **(bottom)** View of the shield door from outside, during construction. The cryostat vacuum conduit is out of the frame (below), and the PMT signal cables had not yet been installed (out of the frame, above).

specified maximum bend radius of 28 cm, so it was passed through a separate 15 mm diameter hole, with a gentle $R = 28$ cm curve. After all cables and tubes were mounted in the shield door, any remaining air gaps were filled with shavings of low-activity Pb (if sufficient space remained) and then sealed with a silicon polymer. Figure 3.14 (left) shows the PMT signal cables mounted in the door, prior to filling the air gaps with Pb shavings. A schematic of the shield door feed-throughs is shown in the Appendix.

3.2.7 Xe self-shielding

The density of liquid Xe under XENON10 operating conditions is 2.86 g cm^{-3} , and as mentioned in Sec. 2.2 there are no long-lived radio-isotopes. The detector mass itself can therefore be relied on for additional self-shielding from external electromagnetic background. Besides the radio-purity of the Xe, this technique requires only the ability to define a fiducial volume (discussed in Sec. 4.4.3). The fiducial cuts implemented in XENON10 are discussed in Sec. 4.6.2.

The Compton cross section (“incoherent”) dominates for photons with $E_\gamma \gtrsim 500$ keV, and below this photo-electric absorption dominates. The mean free path of a 1 MeV γ in Xe is $\lambda_{Xe} = 6.0$ cm. For comparison the mean free path of a 1 MeV γ in steel (Fe) is $\lambda_{Fe} = 2.1$ cm, in liquid Ar $\lambda_{Ar} = 12.6$ cm and in H_2O $\lambda_{\text{H}_2\text{O}} = 2.1$ cm. The electromagnetic cross section and the mean free path λ are shown in Fig. 3.16) as a function of energy. The importance of self-shielding for Xe-based detectors for rare-event-search cannot be over-emphasized. The biggest impact of self-shielding against typical external radioactivity is realized in the first few cm. Despite its modest size, XENON10 saw a $\times 3$ reduction in event rate due to the outer 2 cm of self-shielding Xe. The measured event rate in the XENON10 detector 5 – 25 keVr during WIMP Search data-taking is shown in Fig. 3.15 as a function of cm of self-shielding Xe (in radius, top and bottom).

Earlier work [144] assumed the energy scale obtained from the ^{137}Cs calibration data (2.2 phe/keVee) to be valid at low energies, and as a result the calculated event rate was too low by about $\times 2$. Considering Fig. 2.11 and Fig. 2.12, the energy scaling at the detector threshold was taken to be 3.8 phe/keVee, so that the WIMP Search energy range is about 5 – 25 $S1$ phe $\approx$ 5 – 25 keVr $\approx$ 1.3 – 6.6 keVee. From this, the event rate in the center of the detector (Fig. 3.15) was about 0.8 dru, with only basic (QC0) quality cuts applied. The error bars in Fig. 3.15 reflect uncertainty in the (x, y) position reconstruction (varying from 2.4 mm at the edge to ~ 1 mm in the center [122]), Poisson fluctuations in bin counts and 0.5 keVee uncertainty in the energy range. Uncertainty in the position reconstruction in z is sub-dominant. The event rate in each individual (variable mass) concentric shell of Xe drops sharply across the outer 1 cm, to about 1 dru when there are > 2 cm of self-shielding Xe in all directions. The integrated rates (red circles) are the relevant number for the total background event rate. For understanding the contribution of each component to the observed background rate in the XENON10 detector, the shell rates provide a more sensitive measure. This is because they more closely elucidate the slope of the penetration of the incident radiation (as shown in Fig. 3.17).

3.2.8 Internal sources of electromagnetic background

The internal γ and β background is dominated by the U/Th/K/Co activity from the stainless steel of the detector and cryostat, the PMTs and their associated circuitry (PMT “bases”), the ceramic high-voltage feed-through and the ^{85}Kr in the Xe,. Other smaller contributions arise from the bremsstrahlung of the ^{210}Pb in the shield, the ^{222}Rn in the cavity, and the Teflon. All materials were selected for radio-purity by




Figure 3.15: **(top)** The background event rate in the active target during the WIMP Search. The integrated rate (red circles) is defined for the enclosed volume, so that (for example) with 2 cm of self-shielding Xe in (r, z) , the active target mass with $R < 8$ cm and $2 < z < 13$ cm has a rate of ~ 1.2 dru. The width of the bars indicate the extent of the shell (black lines) in (r, z) , so that (for example) the outermost shell ($0 - 1$ cm self-shielding Xe) has $9 < R < 10$ cm and similarly encloses the outermost 1 cm in z . **(bottom)** The energy spectrum of the WIMP Search data in the 5.4 kg fiducial target, after basic quality cuts (QC0). The event rate in the WIMP Search energy range ($4.5 < E_R < 26.4$ keVr, indicated by the red box) is 1.10 ± 0.03 dru. Scaling to keVee using 3.8 phe/keVee is found to agree very well with the combined $(S1 + S2)$ energy scale, which is discussed in [47].




Figure 3.16: **(top)** Photon cross-sections in Xe [35]. The photoelectric and Compton (“incoherent”) cross sections dominate the attenuation. **(bottom)** The mean free path of photons in Xe, for a liquid density of 2.86 g cm^{-3} . A 1 MeV photon has $\lambda \sim 6 \text{ cm}$.




Figure 3.17: A Monte Carlo simulation of the background event rate in XENON10, by component. The rates are for concentric (variable-mass) shells, with 0 cm indicating the outer edge of the active Xe. The width of the bar indicates the extent of the shell in (r, z) . The event rates correspond to the radio-impurity of each component as shown in Table 3.5. The fall-off of the ^{210}Pb is sharper than the other components, as is expected for a beta source. The internal ^{85}Kr (not shown) would appear as a flat line on this plot. The simulations are described in detail in [50].

checking available literature; in most cases, identical copies (or the actual items used in XENON10) were screened in a HPGe detector to assess their radio-purity.

The specific steel used in construction of the detector vessel and vacuum cryostat was not screened for radio-purity. Several steel samples similar to those used in the detector construction were screened at the SOLO counting facility at Soudan mine [150]. Other samples of steel from various applications of the shield installation were screened at LNGS (see Table 3.3). The U/Th/K contamination of the detector vessel and vacuum cryostat were assumed to similar to the results from an MDC ConFlat flange sample. The Co contamination was conservatively assumed to follow the worst-case measurement. The ^{60}Co contamination of the vacuum cryostat was assumed to be $\times 5$ lower than for the detector vessel, since the former was > 12 years old (^{60}Co has $\tau_{1/2} = 5.27$ y).

The principal radio-impurity concentrations are summarized in Table 3.5, along with the Monte Carlo-predicted event rate in the 5.4 kg fiducial target used for WIMP Search analysis (2 cm self-shielding Xe in r , and 2.8 cm in z – see Sec. 4.6.2). As shown in Fig. 3.18, the predicted event (orange dash-dot) rate from the values in Table 3.5 is lower than the measured event rate (black, with uncertainty) by about $\times 2.5$. There are several possible explanations for this. The stainless steel is the prime suspect, since the actual steel used in XENON10 was not counted. The feed-through, PMT and PMT base contributions are less likely to be higher than tabulated, since the representative samples were counted. The ^{210}Pb is also unlikely to have contributed more than the tabulated event rate, since representative samples were counted. Additionally, the slope of the ^{210}Pb bremsstrahlung spectrum is steeper (Fig. 3.17) than the slopes obtained for the

U/Th/K impurities in the stainless steel, PMTs and feed-throughs. A higher event rate from the Pb is not able to give good agreement with the data.




Figure 3.18: The background event rate ($1.3 < E_R < 6.6$) keVee in the active target during the WIMP Search. The rates are shown for concentric (variable-mass) shells, with 0 cm indicating the outer edge of the active Xe (as in Fig. 3.15 and Fig. 3.17). The width of the bars indicate the extent of the shell in (r, z) . The WIMP Search data (black, with errors) shows a higher event rate than predicted by the radio-impurity concentrations in Table 3.5 (orange dash-dot). Increasing the Monte Carlo-predicted event rate due to U/Th/K by $\times 3$ produced slightly better agreement (red dash). Adding the a homogenous (flat) background event rate of 0.43 dru from ^{85}Kr gave the best agreement (solid blue). The simulations are described in detail in [50].

The second motivation for assuming the steel U/Th/K contamination to be higher than Table 3.5 shows is the results of other steel samples in Table 3.3. Assuming the U/Th/K contamination of the steel to be $\times 3$ higher than shown in Table 3.5 gives slightly better agreement with the data (Fig. 3.18, red dash), but the slope still wrong. This is because the contribution from the homogenous internal ^{85}Kr must be included. Assuming an event rate of 0.43 dru from the ^{85}Kr gives the best agreement between the Monte Carlo (Fig. 3.18, solid blue) and the data (black). It is important to realize that Table 3.5 is meant to give a sketch of the major contributions to the radioactivity budget of XENON10 – rather than an exact accounting – of the event rates due to each component. The argument presented in Fig. 3.18 does not provide conclusive evidence that the stainless steel radio-impurity concentrations are $\times 3$ higher than shown in Table 3.5. What it shows is that the total γ event rate from U/Th/K is higher than assumed in Table 3.5, that the ^{210}Pb bremsstrahlung event rate is consistent, and that the ^{85}Kr event rate is clearly significant. In reality, multiple sources probably contribute to the higher total U/Th/K γ rate. The Kr contamination is discussed in Sec. 3.2.9.

Table 3.5: Radioactivity by isotope of the materials used to construct the XENON10 detector. Activities are quoted as 90% C.L. upper limits [mBq/kg] unless otherwise stated. The PMT and PMT base activities are quoted per unit (rows 5–7), and for all 89 PMTs (row 8). Only the principal sources of expected background are tabulated. The event rate is quoted in the WIMP Search fiducial target volume (5.4 kg Xe), in the energy range $1.3 < E_R < 6.6$ keVee ($5 < E_R < 25$ keVr). The predicted total event rate from these activities is lower than observed, and is discussed in Sec. 3.2.8.

Material	Mass [kg]	U ^a	²³² Th	⁴⁰ K	⁶⁰ Co	Event Rate [dru]
Vacuum cryostat steel ^b	110	29	4	6	25	0.044
Detector vessel steel ^b	33	29	4	6	100	0.176
Teflon	10	0.4	0.2	60	—	0.001
Kyocera A479SS ^c	3	250	40	75	4	0.022
PMT (R8520-06-AL) ^f	1.932 ^d	0.21	0.29	11	0.61	
PMT (R8520-M4F) ^f	0.115 ^e	17.2	3.5	12.7	3.9	
PMT bases ^h		2.71	0.66	0.14	0.03	
89 PMTs (weighted avg. + bases)	2.047	3.87	1.13	11.2	0.83	0.103
²¹⁰ Pb (17 Bq/kg) ^g	6370	—	—	—	—	0.045
Xe (7 ppb ⁸⁵ Kr)	23	—	—	—	—	0.430
²²² Rn (< 5 Bq/m ³)	< 0.6	—	—	—	—	< 0.010
Total (Monte Carlo):						0.83
Total (data):						1.10 ± 0.03

^a Contributions from both ²³⁵U and ²³⁸U.

^b Assumed radioactivity values (exact item not counted); the outer cryostat was > 12 y old so the ⁶⁰Co contribution was assumed to be ×5 lower.

^c or similar (exact items not counted). 5 total in XENON10. Counted at LNGS.

^d 84 PMTs × 23 g/PMT. Counted at LNGS.

^e 5 PMTs × 23 g/PMT. Counted at SOLO.

^f Activity is quoted per PMT.

^g Activity due to the inner 5 cm liner Pb (the outer Pb is sub-dominant).

^h Activity per base (14 resistors and 3 capacitors each). Counted at SOLO. Mass per base ∼ 3 g.

3.2.9 ⁸⁵Kr contamination

The contamination of Kr in the Xe used during WIMP Search was measured by Spectra Gases, Inc. to be < 9 ppb. The actual contamination was found to be ∼ 7 ppb, by comparing the predicted background event rate (versus self-shielding Xe) with the XENON10 data. The comparison is shown in Fig. 3.18. The natural abundance of Kr in the atmosphere is 1.10 ± 0.01 ppm [12]. The radioisotopes ⁸¹Kr ($\tau_{1/2} = 210,000$ y, $Q_\beta = 281$ keV) and ⁸⁵Kr ($\tau_{1/2} = 10.76$ y, $Q_\beta = 687$ keV) are both produced primarily by fission, and also by fast neutron capture (from cosmogenic neutrons). ⁸¹Kr is not much of a concern because its specific activity is 2×10^4 lower than the specific activity of ⁸⁵Kr. It is estimated that a typical commercial nuclear power plant produces 10^{16} Bq of ⁸⁵Kr per year [18], which is released into the atmosphere during fuel re-processing. This leads to a concentration in the atmosphere of about ⁸⁵Kr/Kr ∼ 10^{-12} (corresponding to about 1.5 Bq m⁻³ [111]), which is continually replenished.




Figure 3.19: The simulated event rate and spectral shape in XENON10 due to ^{85}Kr , for various levels of Kr contamination [50]. Spectra Gases quoted an upper limit of 9 ppb Kr for the Xe used for WIMP Search. The actual contamination is estimated to be closer to 7 ppb (orange curve; see also Fig. 3.18 and discussion in text), corresponding to about 430 mdru. The Xe purified by Case Western (< 3 ppt Kr) is predicted to yield an event rate of $\sim 2 \times 10^{-4}$ dru in the energy range $1.3 < E_R < 6.6$ keVee.

Commercial Xe has at best a few ppb Kr, where the separation is performed by distillation or adsorption-based chromatography. The high-temperature getter technology used in XENON10 for Xe purification (described in Sec. 3.3.4) does not remove noble gasses, so the background from the 687 keV beta decay of ^{85}Kr is a potentially serious limitation to the sensitivity of a Xe-based rare-event search. The Monte Carlo predicted β spectra for several concentrations of Kr contamination in the Xe are shown in Fig. 3.19, and compared with XENON10 data in Fig. 3.6. The upper limit of 9 ppb quoted by Spectra Gases (Table 2.3) would correspond to an event rate of at least 550 mdru. As shown in Fig. 3.18, the event rate in the center of the detector asymptotes to about 800 mdru.

A charcoal-column Kr-removal system was developed by Case Western collaborators for XENON10 [38]. After WS4 (from which the initial XENON10 results were drawn), the Xe from Spectra Gases was exchanged for Xe purified by this system, with < 3 ppt Kr contamination. A comparison of background spectra just before and after switching from Xe with < 9 ppb Kr to Xe with < 3 ppt Kr (with all other detector components unchanged) showed a reduction in the event rate 10 – 100 keVee of about $\times 2$. This suggests that the contribution to the background event rate (measured in the center of the detector was ~ 0.4 dru, or about 7 ppb Kr).

3.2.10 Internal sources of neutron background

Neutrons can be generated inside the detector via spontaneous fission of U and Th, or via (α, n) reactions. Spontaneous fission is a lesser concern for events originating inside the shield because up to 9 MeV of prompt γ 's are also produced. This allows for substantial veto efficiency. The principle source of α particles is also U and Th impurities. The rate of neutron production via spontaneous fission is comparable to the rate of (α, n) production in LNGS cavern rock and concrete (as discussed in Sec. 3.2.2). The (α, n) cross-section is a function of α -particle energy and the target element. It tends to be higher for lighter elements (lower Z), because in these cases less energy is required for the α -particle to overcome the Coulomb barrier of the atomic nucleus. The thick-target neutron yield is shown in Fig. 3.20 for several elements of interest (data from [78]).




Figure 3.20: Neutron yield per $10^6 \alpha$ -particles for thick targets, for several principle elements that appear in the construction of the XENON10 detector, in the PMTs and their associated components. Data from [78] except those marked^a, which were scaled from [149] by assuming equal measurements for Al (limited data were available). Data for C (not shown) is roughly $\times 2$ compared with O.

Since the neutron yield data in Fig. 3.20 are for thick targets, the effect of energy loss ($-dE/dx$) as the α -particle travels is already accounted for. The internal neutron production rate can be calculated directly, given (1) the radio-impurity concentrations listed in Table 3.5, (2) a knowledge of the elemental composition of the detector components and (3) the alpha particle energies per U or Th decay. The α particle energies from ^{238}U and ^{232}Th decays were obtained from the GEANT4 database and binned in 0.2 MeV increments to match the yield data from [78]. The energy profile is shown in Fig. 3.21. For ^{238}U (^{232}Th) there are about 8 (6) α -particles per top-chain decay.




Figure 3.21: Summary of α -particles per decay for ^{238}U (left) and ^{232}Th (right), as determined by GEANT4 (blue), and intensity as used for this calculation (red). A bin width of 0.2 MeV was used to match the data shown in Fig. 3.20.

The elemental composition of the Hamamatsu R8520-06-AL PMT are summarized in Table 3.6. The elemental composition of the stainless steel and the electrical feed-throughs are summarized in Table 3.7. The radioactivity of the PMTs was assumed to be equally distributed throughout the PMT. The radioactivity of the PMT bases was assumed to lie entirely in the resistors and capacitors, which have substrates (85% of the mass) made of Alumina (Al_2O_3). The Cirlex base substrate was assumed to be radio-pure. The radioactivity of the feed-through was assumed to be equally distributed between the stainless steel body and the Alumina insulator.

The (α, n) yield from the other principal elements in stainless steel (Co, Ni, Cr) are not tabulated in [78]; however, they are available in [149] for higher energies. There is a systematic offset probably due to detection efficiency, but the tabulation of Al yields in both cases allows a relative comparison using Al as a standard candle. Data from [149] were scaled in this manner, and are shown in Fig. 3.20. Data from ^{58}Ni are a factor of 100 lower than ^{60}Ni , and are not shown. No (α, n) data was found for Cr. It is a reasonable assumption that its yield is lower than Fe. The calculation of the (α, n) yield from stainless steel assumed 100% Fe composition, which seems conservative.

The neutron yield from spontaneous fission is much easier to calculate. The branching fraction to spontaneous fission for ^{238}U is 5.45×10^{-7} [99]. The total internal ^{238}U activity can be calculated from the activities and masses given in Table 3.5. Contributions from ^{232}Th were neglected because the branching fraction is a factor of 10^4 less than for and ^{238}U . Contributions from ^{235}U were also neglected because it was found to be present at a level of about 2% relative to ^{238}U . In the interest of calculating an upper limit to the neutron background rate, it was assumed that 2 neutrons are generated per ^{238}U spontaneous fission.

The total neutron yield from internal (α, n) and spontaneous fission in XENON10 are shown in Table 3.8. A total of about 14 emitted neutrons per month (considering all neutron energies) were predicted from (α, n) , and about 15 neutrons per month from spontaneous fission. Very few (if any) would be expected to single scatter in the Xe fiducial target.

Many factors must be considered in calculating the number of neutrons which single scatter in the XENON10 target per month. These include spectrum of emitted neutron energies, the fact that the neutrons are emitted into a 4π solid angle, the Xe (and other detector material) cross sections for elastic scattering,

Table 3.6: Elemental composition of each Hamamatsu R8520-06-AL PMTs, as deployed with voltage divider network on Cirlex base.

Component	Mass [g]	Composition
Kovar housing & pins	13.0	Fe (55%) Ni (29%) Co (16%)
Stainless steel electrodes	7.0	Fe (70%) Cr (18%) Ni (9%)
Synthetic glass for window	2.0	SiO ₂ (100%)
Borosilicate glass for stem	1.0	SiO ₂ (67%) B ₂ O ₃ (18%) ^a
Cirlex	1.44	C (69%) O (23.7%) ^b
KOA resistors	0.06	Al ₂ O ₃ (85%) ^c
AVX capacitors	0.11	Al ₂ O ₃ (85%)

^a Also: Na₂O (6%) Al₂O₃ (4.3%).

^b Also: N (7.3%) H (2.6%).

^c Also: Ag (2.7%) Ni (2.9%) Sn (2.5%).

Table 3.7: Elemental composition of electrical feed-through and stainless steel detector construction materials.

Component	Mass [kg]	Composition
Stainless steel	143	Fe (70%) Cr (18%) Ni (9%)
Kyocera feed-through	3.0	Al ₂ O ₃ (10%) Fe (63%) ^a

^a Also: Cr (14%) Ni (7%).

and the effectiveness of the ability to veto multiple scatters. Perhaps the best way to understand the raw scattering rate and is via Monte Carlo simulation. The neutron yield quoted in Table 3.8 was used as input to a detail simulation [50]. Preliminary estimates indicate that $\ll 1$ neutron/2 months/10 kg would single scatter with a recoil energy $5 < E_R < 25$ keVr. Considering also the prediction for the external neutron rate (Sec. 3.2.5), the total single-scatter neutron rate ($5 < E_R < 25$ keVr) was expected to be sub-dominant (for a live time of ~ 2 months).

3.3 Detector sub-systems

3.3.1 Cryogenics

XENON10 was maintained at a stable temperature of $T = -93^\circ \text{C}$ (177 K) for the duration of the WIMP Search data-taking and the calibration data-taking. The temperature stability was better than $\pm 0.05^\circ \text{C}$. At this temperature the vapor pressure is $p = 2.11 \pm 0.01$ atm. Stability is crucial because the Xe vapor pressure varies with the temperature, and as mentioned in Sec. 2.2.4 the secondary scintillation (S_2) yield depends on both factors. The exact temperature is not critical, and in fact a range of temperature and pressure combinations (consistent with Fig. 3.22) were explored. The final operating conditions were eventually set such that the efficiency of the cryogenic system was optimized, while maintaining a comfortable (< 2.5 atm) vapor pressure. “Comfortable” means safely below the maximum pressure rating of the PMTs, which is 5 atm, and safely above the sublimation point. Typical operating conditions for XENON10 are indicated by a red circle on the phase diagram in Fig. 3.22.

At these temperatures it is easy to design a cryogenic system using a cold-finger (usually made of Cu),

Table 3.8: Calculated neutron production per month in XENON10 from $^{238}\text{U}/^{232}\text{Th}$ activity in XENON10. The yields assume the material masses and radioactivities from Table 3.5, and the compositions as summarized in Table 3.6 and Table 3.7.

Component	neutron production per month		
	(α, n)		fission
	^{238}U	^{232}Th	^{238}U
89 PMTs ^a	0.25	0.24	0.29
89 Bases ^b	2.95	1.22	0.56
5 HV feed-through	2.61	0.91	2.12
143 kg stainless steel	2.94	1.75	11.72
10 kg Teflon	0.78	0.50	0.01
Total (by type)	9.53	4.62	14.70
Total: 28.85			

with one end immersed in LN and the other making solid mechanical contact with the detector vessel. A stable temperature can be maintained by fitting heating resistors on the cold-finger, and using a temperature controller (such as the LakeShore model 304) with active feed-back to control the current through the resistors. This method, used extensively during R&D, can achieve very stable temperatures (also $\pm 0.05^\circ\text{C}$). However, the liquid N_2 boil-off rate from the dewar leads to a system which requires daily refilling. For this reason, XENON10 opted to use a Pulse-Tube Refrigerator (PTR).

A fixed amount of cooling power is provided by the PTR cold head, which is bolted to the SS top flange of the XENON10 detector (refer to Fig. 2.1). Temperature stability is maintained by a LakeShore model 304 Temperature Controller, with feed-back control of several heating resistors; these are bolted to the cold head. Thermal insulation is provided by maintaining a vacuum (about 1×10^{-3} Torr) between the detector and the SS outer cryostat. The PTR cold head is connected to a 3.5 kW GM-type He compressor via insulated SS flexible lines, and to a buffer volume via a single Cu line. This PTR was developed at KEK and Nihon University in Japan specifically for liquid Xe cryogenics, and was manufactured by Iwatani Corp. It can deliver up to 100 W of cooling power at 165 K [77]. The ratio of cooling power to work input (coefficient of performance) for a PTR is less than for a Carnot cycle and is given by $\xi \equiv Q/W = T_L/T_H$. The primary reason is dissipation in the valves (see Fig. 3.23). The expression for the coefficient of performance predicts a $\sim 10\%$ increase in cooling power when operating at 177 K.

There are many types of PTR, and the following discussion refers specifically to the Gifford-McMahon (GM)-type PTR used in XENON10 (though many of the principles are common to other types). The cooling power is produced by creating smooth pressure variations in a He-filled tube (~ 20 atm.), at a frequency of about 1 Hz. One end of the tube is connected to the compressor, the other to the heat exchanger (labeled “ X_H ” in Fig. 3.23) and then the buffer volume, and in between connection is made to the cold head (labeled “ X_L ” in Fig. 3.23). In each of these regions there is good thermal contact between the He gas and the enclosing material, but in the tube itself the He is maintained in thermal isolation. Helium behaves essentially as an ideal gas under these conditions, and since the pressure variations are slow they are also adiabatic (and reversible). It follows that entropy is constant, and the general relation for the molar entropy S_m

$$TdS_m = C_p dT - T\alpha_V V_m dp, \quad (3.2)$$




Figure 3.22: Phase Diagram for Xe (data from [127]) and liquid Xe density (data from [123]). For dual-phase operation as in XENON10, the comfortable operating window is indicated by the red circle. Higher pressures are excluded out of consideration for the PMTs (which are rated to a maximum of 5 atm, above which they will crush). Lower temperatures dictate a lower vapor pressure, which reduces the efficiency of the *S2* proportional scintillation (Sec. 2.2.4). It is also necessary to stay safely above the 161 K, to exclude the possibility of sublimation.

can be written as

$$dT = \frac{T\alpha_V V_m}{C_p} dp, \quad (3.3)$$

where T is the temperature of the He, α_V is the (volumetric) coefficient of thermal expansion, V_m is the molar volume and C_p is the molar heat capacity. Since $\alpha_V > 0$ for He, this shows that adiabatic expansion leads to cooling.

The compressor is decoupled from the PTR frequency and runs at 50 Hz, which allows for a smaller working volume and greater efficiency. The low and high pressure outputs from the compressor (labeled “comp” in Fig. 3.23) are passed from a rotary valve to the regenerator (“reg”). The frequency of the rotary valve thus determines the PTR frequency, by controlling the motion of the pistons which induce the pressure variations. Helium enters the regenerator at room temperature (“ T_H ”), and due to the high heat capacity of the regenerator, exits at the working temperature of 177 K (“ T_L ”). Heat is then removed from the cold head (“ X_L ”), and passed via a heat exchanger (“ X_H ”) to a water cooling loop in the compressor. The valve (not labeled) leading to the buffer volume (“buffer”) is opened and closed at the same frequency (though phase-shifted) as the PTR, to allow the adiabatic expansion (and hence cooling) to occur.

The PTR cold head (air side), the buffer volume and the rotary valve are shown in Fig. 3.24. The reader may have noticed that in the preceeding discussion there is no mention of anything “pulsing.” In fact, if the pressure variation in the tube did pulse (rather than vary smoothly), the efficiency of the PTR would most likely be very low. However inappropriate, the name has stuck.




Figure 3.23: General schematic representation of GM-type PTR (from F. de Waele [51]). The components are, from left to right: a compressor (“comp”), high and low pressure reservoirs for the compressor (“ p_h ”) and (“ p_l ”), a rotary valve (“valve”), a regenerator (“reg”), a cold head (“ X_L ”), a tube containing the He exchange gas (“tube”), a heat-exchanger (“ X_H ”), a valve (not labeled) and a buffer volume (“buffer”). The function of these components are described in the text.

3.3.2 Heat load

The heat load Q which must be dissipated to keep XENON10 in a steady-state during normal operating conditions consists of the 3 usual suspects: radiative (Q_r), conductive (Q_d) and convective (Q_v), with an additional heat load from the constant Xe circulation system (Q_{circ}). The convective heat load scales as $\sqrt{p}$, where p is the vacuum pressure in the cryostat. For a standard gas (N_2) and vacuum cavity size, the convective heat load can be looked up in cryogenic engineering tables. The outer cryostat vessel was maintained at $T_H = 283$ K, and the inner detector vessel was maintained at $T_L = 177$ K. With the vacuum turbo pump mounted outside the shield, the vacuum was comfortably maintained at a steady $p = 2.7 \times 10^{-3}$ Torr. At this pressure, the predicted convective heat load was about $Q_v = 50$ W. The radiative heat load is calculated from the Stefan-Boltzman relation

$$Q_r = \epsilon \sigma A (T_H^4 - T_L^4), \quad (3.4)$$

where $\epsilon = 0.15 \pm 0.05$ is the emissivity of stainless steel [76], and $\sigma = 5.67 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ is the usual black-body constant. This implies a radiative heat load $Q_r = 23 \pm 8$ W. The conductive heat load is calculated from

$$Q_d = \kappa \frac{A}{L} (T_H - T_L), \quad (3.5)$$

where A is the cross-sectional area through which heat can flow, L is the perpendicular distance between T_H and T_L , and $\kappa \simeq 15$ for stainless steel. The surface area of the inner detector is $A \sim 0.5 \text{ m}^2$ (40 cm height, 15 cm radius). It was estimated that $Q_d < 10$ W. The actual total heat load was lower than predicted because a Super Insulation liner was installed between the cryostat and detector.

The heat load due to the Xe circulation depended on the mass flow rate of Xe, which was maintained at a stable 2.3 standard liters per minute (SLPM). Although the SAES getter (discussed in Sec. 3.3.4) operated at 673 K, the Xe gas was cooled to ambient temperature by the time it reached the detector. This is because the inlet and outlet lines are tightly coupled, specifically as they pass through the shield door. The relevant specific heat is $C_p = 0.269 \text{ J g}^{-1} \text{ K}^{-1}$ and the latent heat of vaporization $L_V = 96.3 \text{ J g}^{-1}$ [146]. Considering the “standard” Xe gas density of 5.4 mg cm^{-3} at 300 K and 1 atm, this corresponds to a




Figure 3.24: **(left)** Air-side view of the PTR cold head, which is made of stainless steel, with a flexible ridged outer connector. The buffer volume is also made of stainless steel, with a smooth finish (to the left of the cold head). A thin black tube can be seen connecting the buffer volume to the cold head. **(right)** The PTR rotary valve is mounted at a distance of about 1 m from the cold head. Helium supply and return (high and low pressure) lines are visible on the left.

circulation of about 8 kg of Xe per day. The heat load from the constant circulation of Xe is then calculated to be $Q_{circ} = m(C_p dT + L_V) = 11.2$ W.

A final (sub-dominant) heat load to consider is due to the PMTs. The power dissipation in each PMT due to Joule heating is $Q_{PMT} = (780 \text{ V})^2 / 12.6 \text{ M}\Omega = 0.05$ W. The maximum heat load from the 89 PMTs in XENON10 was therefore about 4 W. One way to remove this load would be to reduce the current flowing through the PMT voltage divider. For example, if each resistor value in the chain were increased by $\times 10$, the PMT biasing requirements would be maintained but the power dissipation (for 89 PMTs) would be < 0.5 W. Such a scenario would require a careful assessment that the PMT linearity was not affected. In principle it should not be, since the reserve charge provided by the capacitors would not change. This is discussed in Sec. 3.3.5.

A second option that should be considered is to move all the PMT resistive and capacitive elements (the bases) outside the detector vessel. Doing so would have the additional benefit of attenuating the radioactive backgrounds associated with the PMT base. In XENON10, the internal γ background from the bases exceeded that of the (newer) R8520-06-AL PMT, as summarized in Table 3.5. Additionally, the predicted neutron background from the bases exceeded the contribution from all the PMTs, regardless of type. But the reduction of the neutron background simply from moving the bases outside the metal detector vessel would probably be modest. In this regard, it would be best to move them completely outside of the shield.

3.3.3 Safety and slow control

Several safety precautions were implemented in the XENON10 installation, to allow for continuous operation in the absence or presence of workers. The ~ 23 kg of Xe occupies a volume of about 8 L at 177 K and 2.1 atm., but at 300 K and 1 atm. occupies about 4000 L (0.4 m^3). If the cooling system were to fail, one risks destroying about \$100k of PMTs, and possibly losing about \$23k of Xe (purchase price; as of this writing, the replacement cost would be about $\times 2$ higher). The last line of defense against an over-pressure scenario

was provided by a rupture disk, with a burst pressure of 4 atmospheres. While this would save the PMTs and the over-pressure hazard, it would also begin to vent Xe to the atmosphere.

The most likely, extensively verified response to a failure of the PTR cooling system was provided by a Cu coil that was wrapped around the inner detector vessel. One end was connected to a liquid N₂ dewar via a solenoid valve, and the other side vented to the atmosphere. Under normal operation, the solenoid valve remained closed. A digital pressure gauge with its own dedicated UPS monitored the chamber pressure. If the PTR failed (or if a power failure exceeded the reserves of the LNGS underground UPS, which was about 8 hours), the pressure in the detector would begin to rise in a matter of minutes. At 2.5 atm., the solenoid valve would be signaled to open, flooding the cooling coil with N₂ for a few minutes. The liquid N₂ would rapidly cool the detector, lowering the temperature and pressure in a matter of minutes. The valve would be signaled to close again once the pressure dropped below 1.7 atmospheres. The process was verified to cycle stably with the detector in a safe, dynamic steady-state for as long as liquid N₂ remained in the dewar (several hours). This system gave workers plenty of time to arrive underground (even considering the 30 minute drive from Paganica), or to take a deep breath before proceeding, if they were already on-site.

An additional fall-back option existed whenever workers were present. It is a relatively simple task to recuperate the Xe into its storage bottles. An open-mouth dewar filled with liquid N₂ would be hoisted into place around a storage bottle to create a cryo-pump. This was the reason for the second low-pressure liquid N₂ dewar. Both dewars were available at all times just outside the XENON10 underground box, and their level was continuously monitored and topped up. As a final precaution, the slow-control monitoring system was configured to automatically send text messages to people's cell phones if any critical parameter value returned 3 consecutive readings outside a pre-set "safe" range. The slow-control system monitored (at ~ 10 s intervals) the chamber pressure and temperature (in several physical locations outside the vessel), the circulation flow rate, the cryostat vacuum, the bias voltages of the PMTs and grids, and several other relevant parameters. It is described in detail in [108].

3.3.4 Xe gas handling and purification

All inner surfaces of the gas handling system were either 304 stainless steel (SS) or Teflon; most of the SS tubing was electro-polished. Two SS storage cylinders were located outside the XENON10 box, and passed to the purification panel via a two-stage regulator followed by about 7 meters of 1/4" SS tube. The Xe was then passed through the purification panel (shown in Fig. 3.25 below). The ~ 12 Swagelok valves (SS-4BG-V51) allow 3 modes of operation: (1) filling the detector with Xe, (2) recuperating the Xe back into the storage bottles, and (3) circulating Xe out of the detector, through the purifier and back into the detector (closed cycle). In the first 2 modes, the Xe flow was established and maintained by setting the temperature of the source lower than that of the destination. The temperature and pressure requirements for these processes can be seen in Fig. 3.22.

Once the detector was full (which would take about 8–10 hours), continuous circulation was established as follows: liquid Xe was circulated out of the bottom of the detector through a buffer volume (~ 4 L), during which it would vaporize and equilibrate with the temperature of the SS tubing (283 K). A KNF model N143.12E double-membrane diaphragm pump was used for this process. Its active surfaces – aluminum and neoprene – provided the only departure from the SS and Teflon environment. Considering Fig. 2.9, any impact this had on the purity was sub-dominant. A Teledyne-Hastings 302 flow controller was located between the buffer volume and the circulation pump. The Xe gas was then passed through the purifier, a



Figure 3.25: **(left)** The Xe gas handling & purification panel, mounted on the XENON10 shield door. **(right)** Close-up of the SAES PS4 high-temperature getter.

SAES PS4-MT3-R high-temperature getter. The purification capabilities are summarized in Table 2.3. The Xe was finally re-condensed by the PTR and passed back into the detector.

The SAES PS4-MT3-R utilizes a reactive Zr alloy cartridge (heated to 400°C) to bind impurities as they pass through the cartridge. The alloy is heated because this allows the trapped impurities to diffuse into the bulk of the Zr (rather than remaining on the surface), thus increasing the useful life time of the purifying cartridge. The purification capabilities of the getter (see Table 2.3) are valid for Xe gas flows as high as 20 standard liters per minute (SLPM) [134]. However, SAES engineers recommend a flow rate closer to 2 SLPM for optimal purification [135], due to the latent heat of Xe. In earlier work on liquid Xe particle detectors, and even some R&D work at present, Oxisorb purification cartridges were used. These use a proprietary adsorbent which chemically binds O_2 (and as a side-effect, H_2O). The stated purification capabilities of the Oxisorb is a bit less than for the SAES getter, at < 5 ppb O_2 and < 30 ppb H_2O (information for other impurities was not found). Experiments using Oxisorb to purify Xe have reported an electron lifetime > 2.5 ms [100] (the electron lifetime in XENON10 was similar), however the Oxisorb has also been found to introduce Rn contamination into the liquid [5, 100], making it unsuitable for a low-background experiment.

The KNF diaphragm pump could sustain a maximum flow of 30 SLPM with zero gradient from inlet to outlet; in practice there is a pressure differential of about 1 atmosphere. XENON10 maintained a stable flow




Figure 3.26: Schematic of the Xe gas purification system. The items contained on the purification panel (see Fig. 3.25) are highlighted in blue. Figure courtesy of Coimbra Univ.

rate of 2.3 SLPM. The reason for this limitation was to avoid running the PTR too close to its maximum available cooling power. In practice the cooling power of the PTR is fixed, and the applied cooling power is varied by heating elements mounted to the cold head. The heaters were maintained at about 30% of their maximum value. This gave the temperature controller sufficient headroom to maintain a stable equilibrium. A schematic of the purification system is shown in Fig. 3.26.

3.3.5 PMTs

XENON10 uses 89 Hamamatsu R8520-06-AL photo-multiplier tubes (PMTs) to view the Xe scintillation light. The tubes were optimized by Hamamatsu to operate immersed in the liquid Xe (177 K and 2.1 atm) and can withstand up to 5 atm pressure. They can resolve the output pulse corresponding to a single photo-electron ejected from the photo-cathode as a result of incident scintillation. Each tube is compact ($25 \times 25 \times 35$ mm), with a 10-stage metal-channel dynode, a synthetic silica window and a bi-alkali UV-sensitive photo-cathode. As summarized in Sec. 2.5.3, the typical photo-cathode quantum efficiency was measured by Hamamatsu to be in the range 20 – 25% at 178 nm. Due to the compact geometry, the focusing of photo-electrons onto the first dynode (the collection efficiency) was about 75%.

PMT gain was calibrated every 1 – 2 weeks with a pulsed LED [108] throughout the WIMP Search data-taking; stability was found to be better than $\pm 2\%$ over a period of several months. A typical gain of about 2×10^6 is achieved at ~ 780 V bias, with a typical relative standard deviation $\sigma/\mu \equiv \sigma_r = 0.58 \pm 0.08$ from a Gaussian fit to the single photo-electron pulse-area distribution. These values are the average for the 89 PMTs used in XENON10. Gain variation of $\times 2$ was typical across the 89 PMTs (up to $\times 3$ in extremis). A typical single photo-electron spectrum at 2×10^6 gain is shown in Fig. 3.28. The mean of the Gaussian is about 16 mV ns referred to the PMT output (these were the standard units used for XENON10 integrated




Figure 3.27: **(left)** Head-on view of a PMT mounted on the copper plate of the Top Array; the photo-cathode gives the silica window an amber hue, and the focusing grid is visible behind it. **(center)** a Cirlex PMT base, designed and fabricated at LLNL. Surface-mount (voltage-dividing) resistors and pin-cages are visible; capacitors are mounted on the reverse. **(right)** PMTs are firmly suspended by their connection to the pin-cages; serial numbers are inscribed on the side of each tube.

PMT pulses). The charge Q_0 leaving the PMT anode (the anode current) was converted into a voltage pulse across the PMT load resistance R_L :

$$Q_0 = \frac{V \Delta t}{R_L} = \frac{16 \times 10^{-3} \text{ V}}{50 \Omega} \times 10^{-9} \text{ s}, \quad (3.6)$$

where the 50Ω load resistance was determined by the combination of the PMT load resistance ($R_{PMT} = 10 \text{ k}\Omega$) and the cable/electronics input load resistance ($R_{in} = 50 \Omega$) according to $R_L = R_{PMT} \cdot R_{in} / (R_{PMT} + R_{in}) \simeq 50 \Omega$. For a single photo-electron leaving the photo-cathode, this gives a total charge $Q_0 \simeq 2 \times 10^6$ electrons.

The large value of σ_r arises from the fact that the fluctuation in the PMT gain is dominated by Poisson fluctuations in the yield of secondary photo-electrons from the 1st dynode. Because of the compact size of the R8520 PMT, this involves a trade-off between the collection efficiency and the secondary photo-electron yield from the 1st dynode. The former requires a lower voltage, the latter a higher voltage. It also depends on the material choice for the 1st dynode. The result for the R8520 is a secondary electron yield from the 1st dynode of $\delta = 4.5$ electrons on average. The (Poisson) relative standard deviation from this stage is $\sqrt{\delta}/\delta = 50\%$. The contribution from subsequent dynode stages must also be considered, each weighted by a factor of $1/\delta^{N-1}$, where N is the dynode number [93]. It is easily shown by summing the geometric series that the relative standard deviation is then predicted to be $\sigma_{r,phe} = 1/\sqrt{\delta - 1} = 0.535$. The measured value of $\sigma_r = 0.58$ is slightly higher due to the random time-offset jitter of the digitizer, which introduces an additional relative standard deviation $\sigma_{r,jitter} = 0.20$. This is discussed in detail in Sec. 4.1.3.

The voltage divider network and associated capacitors were contained on a Cirlex printed circuit board, custom-made for XENON10. A voltage divider current of $62 \mu\text{A}$ was typical across the full $12.6 \text{ M}\Omega$ chain. The inter-dynode resistance was $R = 1 \text{ M}\Omega$ except between the photo-cathode and 1st dynode, and between the 1st and 2nd dynodes; each of these had a $R = 2 \text{ M}\Omega$. The final stage between the last dynode and the anode had $R = 0.5 \text{ M}\Omega$, and the final load resistance (to ground) was $0.1 \text{ M}\Omega$. These were the standard values recommended by Hamamatsu [74]. The last 3 stages included a 22 nF stabilizing capacitor in parallel to prevent large signals from depleting the available reserve charge.




Figure 3.28: The single phe spectrum from a PMT used in XENON10, obtained from LED-calibration. The mean gain value and relative standard deviation σ_r were taken from Gaussian fit (blue line). This PMT shows a lower-than-average value of σ_r , and achieved a gain of 2×10^6 at 756 V (also slightly below the average of 780 V). The noise peak also follows a Gaussian distribution (red line). Plot courtesy of A. Manzur [108].

The reserve charge in the final-stage capacitor must exceed about $\times 100$ of the total charge draw Q_0 from a pulse for the output to be linear to within $\pm 3\%$ [75]. The reserve charge is

$$Q_R = 22\text{nF} \cdot 780\text{ V} \cdot 0.5\text{ M}\Omega / 12.6\text{ M}\Omega \simeq 680\text{ nC}. \quad (3.7)$$

As an example, the total $S2$ pulse from a 122 keV (662 keV) γ photo-absorption produces about 10^5 (6×10^5) phe. Considering the typical PMT gain, this is about 32 (192) nC of total charge. The $S2$ is typically concentrated in up to 4 of the top-array PMTs, and the maximum signal received by one PMT is $< 10\%$ (5% typical). The typical maximum charge draw on a PMT is then about 2 (12) nC, so the linearity requirement is well-satisfied at 122 keV. At 662 keV, it is well-satisfied in most of the PMTs, but beginning to push up against the linearity threshold in the 4 top-array PMTs nearest to the event (x, y) locus. For signals larger than this, PMT linearity in the $S2$ signal could not be assumed. A recovery time $\tau \equiv RC = 11\text{ ms}$ is required for the final-stage capacitor to re-charge to $1/e$ of its maximum value. During WIMP Search the raw event rate was 2.5 Hz (typically 400 ms between events), so the re-charge criterion was (typically) easily satisfied.

The signal wires were fed through the cryostat and the detector using 25-pin Burndy connectors. Connections between the PMTs and the PS776 $\times 10$ amplifiers were made with an RG178 coaxial cable (Alpha Wire P/N 9178B) with silver-coated conductor and a braid coverage (shielding) of 96%. The dielectric was made of Tetra Fuoro Ethylene (TFE), and the insulating sheath of Fluorinated Ethylene Propylene (FEP). This wire was chosen because it affords excellent shielding, $50\text{ }\Omega$ impedance, and was known not to have an adverse impact on the liquid Xe purity.

3.3.6 High voltage requirements

XENON10 had two separate high voltage requirements: the PMTs, and the grids which created the electric fields (E_d and E_e) across the active target. Each PMT was biased to a maximum of 900 V using a CAEN A1733. The unit has a maximum output of $200 \mu\text{A}$ at 4 kV, with voltage resolution of 1 V and ripple $< 30 \text{ mV}$ peak-peak. A custom RC filter was placed in-line to block frequencies above $\sim 10 \text{ Hz}$. The PMT HV was brought through the cryostat and detector walls using Kyocera (Al_2O_3) multi-pin feed-throughs with Burndy-style connectors. Air-side connections were made using Redel connectors and 5 multi-strand, insulated HV cables. In vacuum and in the detector, connections were made with single-strand silver-plated Kapton-insulated wire (KAP1, available from MDC). Kapton has a dielectric strength of 80 kV mm^{-1} and extremely low out-gassing.




Figure 3.29: A schematic of the inner detector region, showing the location of the cathode, anode and shielding grids (green horizontal lines), and the location of the feed-through (vertical orange line). The SS walls of the detector are not shown; the area immersed in liquid Xe is shaded purple. A detailed description of each component is provided in the text (Sec. 3.3.6)

The voltages for the anode and grids of the drift region were also provided by the CAEN A1733. At the air-side, the HV for the anode and grids terminated in standard SHV connectors; through the vacuum it was carried by Teflon-insulated wire. Inside the detector, the bare metal from the feed-through made a

direct connection to a second set of metal rods which connected to the anode and grids (see Fig. 3.29). This second set of connectors was embedded in the Teflon can. It was possible to leave the feed-through connectors bare because liquid Xe is an excellent insulator, with a dielectric strength of 40 kV mm^{-1} (Table 2.1).

The cathode grid bias was provided by a Heinzinger HV power supply model PNC, capable of -100 kV with voltage stability and ripple both $< 0.01\%$. The HV was brought into the detector by a custom-made [70] feed-through, consisting of a single steel rod, and a single piece of hollow, cylindrical Teflon insulator (about 4 mm thickness). The dielectric strength of Teflon is 19.7 kV mm^{-1} [46]. The Teflon was a cylindrical shell whose inner diameter was slightly smaller than the steel rod; the rod was pressure-fit into the Teflon so that no gaps remained between the steel and the Teflon. It was found that such gaps lead to HV discharge (dielectric barrier breakdown), and a high rate of spurious scintillation in the Xe. At the detector and cryostat junctions, a Swagelok barrel with a tapered inner diameter was fixed to the Teflon, and the Teflon was swaged by tightening a nut against the barrel. The entire apparatus was mounted on a standard mini-conflat flange. The feed-through was tested to hold $> 20 \text{ kV}$ in air.

The construction of such custom-made feed-throughs are particularly important for future liquid Xe rare-event searches. Commercial feed-throughs that hold $> 20 \text{ kV}$ are generally constructed of large masses ($\sim \text{kg}$) of ceramic (Al_2O_3 , dielectric strength 13.4 kV mm^{-1} [46]), which can contain high levels of radioactive ^{40}K . The effect of ^{40}K on the electron recoil background is not negligible (see Sec. 3.2.8). The simplicity of this particular custom design allows for excellent radio-purity. For even better radio-purity, OFHC Cu could be used in place of the steel.

Chapter 4

DATA ACQUISITION AND ANALYSIS

“All truths are easy to understand once they are discovered; the point is to discover them.”

– Galileo Galilei

4.1 Data acquisition

While charge-sensitive signal amplification and read-out is standard for many particle physics applications (and even some other WIMP Search experiments such as WARP), XENON10 opted for voltage amplification and thus fast digitization (105 MHz, or ~ 10 ns / sample). This was a natural choice for 2 reasons: the single-electron response of the PMTs is already comfortably above the baseline noise, and therefore no gain in sensitivity should be expected from charge amplification. Additionally, the integration time constants ($\sim \mu\text{s}$) typical of charge-sensitive amplification could risk confusing $S1$ pulses with $S2$ pulses. Voltage amplification allows preservation of the discrete arrival time of photo-electrons, which provides additional information about events in the detector. Both the $S1$ and $S2$ signals are read-out through the same electronics chain, which consists of a Hamamatsu R8520-AL PMT (gain $\sim 2 \times 10^6$ at 750 V), a Phillips Scientific PS776 linear $\times 10$ amplifier, and a Struck SIS3301 14-bit digitizer. Each amplifier channel has two outputs, one for direct digitization and the other for triggering decisions.

4.1.1 Digitization

The 14-bit Struck SIS3301 105 MHz analog-to-digital converter (ADC) have a maximum record length of 16384 samples (in multiple-event mode; this mode provided a lower acquisition dead time than single-event mode [50]). Event waveforms of 16350 samples were recorded ($\sim 156 \mu\text{s}$). The maximum drift time across the active Xe target at $E_d = 0.73$ kV/cm was $80 \mu\text{s}$, and in the final $S2$ -based trigger condition (discussed in Sec. 4.3.2), the trigger delay was set so that all physical $S1$ events were safely within the event record. The choice of 105 MHz sampling frequency was dictated by cost, availability and the need for the dynamic range afforded by a 14-bit ADC. The dynamic range requirement is due to the variation between a ~ 3 phe $S1$ at the detector energy threshold, and a ~ 3000 phe $S2$ signal from a typical background γ event at ~ 4 keVee (see for example Fig. 4.6).




Figure 4.1: A simulated typical single photo-electron pulse from the R8520 PMT (black line), and the same pulse passed through an RC filter with time constant $\tau = 5$ ns (blue line). The jitter due to the random sampling time-offset of the 105 MHz digitizer (red circles) with respect to the start of the pulse is shown for two possible random time-offset scenarios (left and right). The pulse area is calculated by simply summing the digitized values. The “true” area of the filtered pulse (blue line) is $16 \text{ mV}\cdot\text{ns}$, while the area of the pulse after 105 MHz digitization (red line) is $17.2 \text{ mV}\cdot\text{ns}$ (left) and $11.6 \text{ mV}\cdot\text{ns}$ (right).

4.1.2 Time-offset jitter

A single photo-electron pulse from the R8520 PMT is essentially Gaussian, with FWHM ~ 4.5 ns. A 105 MHz digitizer samples every 9.5 ns, which at first sight would seem to introduce the possibility to completely miss a single phe pulse. However, it is possible to maintain full sensitivity to single phe pulses with an appropriate choice of front-end analog filter. Figure 4.1 shows a (simulated) typical single phe pulse at a gain of 2×10^6 (black line), and also the same pulse passed through an RC filter with a time constant $\tau = 5$ ns (blue line). This corresponds to a -3 dB decrease at $1/(5 \times 2\pi) = 30$ MHz (the filter pole is 30 MHz). The maximum pulse height after the filter is reduced by about $\times 2$, with a typical measured pulse height of 1.2 mV. After the PS776 $\times 10$ amplification, the typical pulse height is 12 mV. The ADC has a full-scale range of 2 V, so 12 mV corresponds to about 98 ADC bins.

The reason the typical pulse height is 1.2 mV rather than the 1.6 mV maximum shown in Fig. 4.1 is that the digitizer does not always capture the peak maximum, but rather has a random time-offset with respect to the start of each pulse. Two cases of this random time-offset are illustrated in Fig. 4.1. This is a generic effect, un-correlated with the measurement of a single phe. Given a fixed-size pulse input (for example from a pulse-generator) and a fixed ADC clock-speed (105 MHz), the resulting measured pulse area will show a jitter with a standard deviation that depends on the characteristic width of the input pulse. The standard deviation of this distribution can be reduced by selecting an appropriate filter shaping frequency.

4.1.3 Filter simulations

A simulation of the random time-offset effect was used to predict an optimal time constant $\tau = 5$ ns to minimize the jitter. The input to the simulation was a fixed-size single phe pulse, modeled as a Gaussian (as shown in Fig. 4.1). The variation in the measured pulse area $\sigma_r \equiv \sigma/\mu$ is shown as a function of the filter time constant (τ_{RC}) in Fig. 4.2 (left, blue), along with the associated (normalized) reduction in pulse height (right, blue). The effect of a filter with Gaussian shaping is also shown (black dashed). The Gaussian filter was considered for XENON10 because of its superior characteristics, but was not used because of the cost and time required to build a custom 89-channel filter. In contrast, Struck was able to fit a simple RC filter to the input of the ADC boards at no additional cost, and with minimal (< 1 week) delay. As a general note, a filter with a pole (-3 dB point) of $< \times 0.5$ the sample frequency should always be used to avoid signal aliasing.

As can be seen from Fig. 4.2, the $\tau = 5$ ns filter pole results in a relative standard deviation $\sigma_{r,jitter} = 0.20$, with a pulse height reduction of about $\times 2$. This was deemed the optimal compromise in terms of maintaining the single phe pulse height, with a minimal effect on the intrinsic resolution. As discussed in Sec. 3.3.5, the intrinsic resolution is $\sigma_{r,phe} = 0.535$ for the R8520 PMT single photo-electron spectrum. The combined effect is then predicted to be $\sigma_r = \sqrt{\sigma_{r,jitter}^2 + \sigma_{r,phe}^2} = 0.57$. The measured average for the 89 PMTs was $\sigma_r = 0.58$, an increase of only about 8% as a result of the ADC time jitter effect.

4.2 Data management

The event rate for XENON10 has a maximum of 20 Hz, determined by the download speed of the VME to the (Linux-based) acquisition PC. With data stored in binary *unsigned int* format (2 bytes/sample), the event size is $16384 \cdot 88 \cdot 2 = 2.88$ MBytes. During WIMP Search data-taking, the data acquisition rate was




Figure 4.2: **(left)** The variation in measured pulse area ($\sigma_r \equiv \sigma/\mu$) of an identical single photoelectron pulse (as shown in Fig. 4.1) is a function of the RC filter pole. The time constant $\tau = 5$ ns used for XENON10 resulted in $\sigma_r = 0.20$. **(right)** The associated pulse-height reduction as a function of filter pole. For XENON10 the reduction was about $\times 2$. A filter that provides Gaussian (as opposed to the asymmetric RC) shaping is found to be a superior choice in terms of maximizing pulse height and minimizing the variation in measured pulse area.

2.5 Hz. This lead to the alarming data accumulation rate of 620 GB day^{-1} ; the XENON10 RAIDs totaled 9 TByte, implying a potential 14-day limit to the run-time.

Due to the aggressive schedule of XENON10 operations, the installation of the fiber-optic network connection from the underground installation to the LNGS computing center did not occur until November 2006. Until this point, the connection speed was limited to 10 MBit/s. A series of bash shell scripts were written to automatically compress the data after it was recorded to disk, and then *rsync* it from the underground DAQ PC (with 800 GB available locally) to the RAIDs in the computing cluster. The network transfer speed until Novemebr 2006 created a potential bottle-neck which also threatened to limit the run-time.

4.2.1 Zeroth-order compression

An in-house data-compression scheme was developed [47], but it was found to be too CPU-intensive to keep up with the rate of in-coming data. Environment-specific compression (i.e. ROOT tree or Matlab .mat) are known to achieve lossless compression of better than $\times 4$, but these were rejected because XENON10 maintained 2 parallel analysis chains (one in each environment), and using either was found to introduce compatibility issues for the other.

The Unix utilities *bzip2* and *gzip* were tested: Data was stored in files containing 104 events each, for a 300 MByte file size. The *bzip2* gave a typical compression ratio of $\times 4.4$ (68 MByte file size) in a minimum elapsed time of 120 seconds, while the *gzip* achieved $\times 3.4$ compression (88 MByte file size) in a minimum time of 80 seconds. Initially, the *bzip2* was chosen for its superior compression. In this case, the transfer time for a single file over the 10 Mbit/s network was about 60 seconds. With the *gzip*, the minimum transfer time for a single file was closer to 80 seconds.

In either case the situation lead to the underground disks filling up in a matter of days, thus ending the

data run (note that this situation was resolved before the beginning of the blind WIMP Search data taking). Because it was faster, *gzip* allowed for a longer data-taking duration. However, its lesser compression ratio decreased the amount of time before the RAID was full. The usual practice of writing data to tape was not of any practical use in solving these problems, since this operation is significantly slower. The actual run times were shorter than predicted by about $\times 2$, due to network traffic and the fact that the *gzip* operation on the dual-cpu DAQ PC tended to slow down when the PC was asked to handle 3 tasks simultaneously (the data acquisition, the *gzip* compression and the *rsync* to the RAID in the computing center).




Figure 4.3: **(Top)** Event waveforms from a single PMT for 2 different low-energy gamma events. Each event contains one single photo-electron (phe) pulse. Some bi-polar noise (ringing) is evident in the earlier data (blue), but this does not affect the single phe determination. **(bottom)** The distribution of ADC samples for each event. A total of 3000 samples are included in each histogram, while the top plot shows a zoom on only 500 samples. The single phe pulses each have 2 samples between 60 – 90 ADC bins. The ± 4 ADC bin baseline flattening is evident in the data from 24 August 2006 (red); later data (after 7 December 2006) used ± 8 ADC bin baseline flattening (effect illustrated in dashed lines on histogram).

4.2.2 Optimal compression

It was determined that *gzip* achieves much better compression (and in a shorter time period), if more of the bytes in each data file were identical. In principle, all “baseline” (non-pulse) data represents the same value (zero), though in reality this value suffers an RMS fluctuation about zero. This realization lead to a simple real-time baseline-flattening algorithm. The smallest feature that must be preserved in any XENON10 event is a single photo-electron pulse. Typical single phe pulse heights are about 10 – 14 mV (82 – 115 ADC bins)

after $\times 10$ amplification and the 30 MHz RC shaping at the ADC input. For each channel, the first 50 samples of the baseline were used to establish the reference baseline (zero). The baseline flattening threshold was set so that any samples which fluctuated within ± 4 ADC bins of the zero were set identically to zero. A histogram of the counts in each ADC bin from a typical event (in a single channel) is shown in Fig. 4.3. A 3000 sample portion of the full event record is shown in the histogram. Two different events are shown: (blue) without any baseline flattening, and (red) with the ± 4 ADC bin flattening. Data after 7 December 2006 used ± 8 ADC bin flattening (no example is shown in Fig. 4.3).

It should be clear from Fig. 4.3 that the ± 4 ADC bin threshold was conservative. Still, it gave a $\times 4$ boost to compression speed, so that *gzip* processes completed in about 20 seconds. The compression factor increased to about $\times 14$ (21 MByte file size). The *rsync* saw a similar speed boost, and the duration of continuous data-taking was thus limited only by the available space on the RAID. After extensive verification that no event information was being lost (*S1* spectra were verified to behave identically with or without the flattening), the flattening thresholds were boosted to a more aggressive ± 8 ADC bins, yielding a compression of $\times 21$ (file size 14 MBytes).




Figure 4.4: Distribution of the off-line baseline flattening thresholds for the 89 PMTs, as used in the final analysis. The thresholds were implemented symmetrically about zero. The additional flattening thresholds lead to the highest accuracy in determining the *S1* signal. Events were written to disk using more conservative thresholds ($< \pm 8$ ADC bins) as discussed in the text.

After a sufficient statistical sample of low-energy data were accumulated, an additional level of baseline flattening was implemented in the off-line analysis. This technique ensured the most accurate representation of the actual *S1* signals by removing any extra noise from the integration window surrounding *S1* pulses. This is discussed further in Sec. 4.4). An optimal threshold for the off-line baseline flattening was determined on a channel-by-channel basis, by sampling the noise level for each channel. The range used was $\pm 11 - 24$ ADC bins (and in 2 particularly noisy channels, as high as 31 ADC bins). This is still comfortably below the 1σ lower-bound of a typical single photo-electron pulse. For a 12 mV (98 ADC bin) pulse, the 1σ lower bound is $98 \cdot (1 - \sigma_r) = 41$ ADC bins. The typical resolution of a single photo-electron $\sigma_r = 0.58$ is discussed in Sec. 3.3.5.

4.3 Trigger

4.3.1 S1-based trigger

Trigger systems based on both the $S1$ and $S2$ signals were implemented and studied. The $S1$ trigger was set to $n \geq 6$ coincidence in an 80 ns window on any of the bottom 41 PMTs. Since the Xe scintillation decay time is 27 ns (see [54] and also Fig. 2.2), this window leads to about 95% acceptance for the photons from a Xe scintillation pulse. The trigger threshold levels were optimized for each individual PMT and correspond to about 80% acceptance for a single photo-electron pulse in any given PMT. The $S1$ trigger logic (and the DAQ and trigger in general) are discussed in detail in [50].




Figure 4.5: **(left)** The $S1$ energy spectrum from WS2, showing ~ 0.80 counts keVee⁻¹ kg⁻¹ day⁻¹ in the fiducial target volume ($r < 8$ cm, $2 < z < 14$ cm). The data were scaled to keVee using 2.2 phe/keVee. The spectrum roll-off below 4 phe (left, pink) is due to the minimum coincidence requirement of $n \geq 2$ (see Fig. 4.8 for more detail). **(right)** The measured $S1$ trigger efficiency (blue) and fraction of $S2$ triggers (pink); Monte Carlo simulation results for $S1$ coincidence $n \geq 6$ (dashed) and $n \geq 3$ (dash-dot). Simulation input parameters discussed in text.

Due to the detector geometry, the bottom PMT array receives 80% of the available scintillation photons; as discussed in Sec. 3.3.5 the PMTs have an effective QE of 0.194, and a $\sigma_r = 0.58$ for the typical spread in the PMT single photo-electron peak. A Monte Carlo simulation accounting for these factors, and also Poisson fluctuations in the number of photo-electrons generated at the PMT photo-cathode, is compared with the measured $S1$ trigger efficiency over about 20 live days of (September 2006) WIMP Search data in Fig. 4.5.

4.3.2 $S2$ -based trigger

An $S2$ -sensitive trigger was achieved by integrating (with $\tau = 1 \mu\text{s}$, roughly the FWHM of a typical $S2$ pulse) the sum of the central 30 PMTs in the top PMT array. This sum was monitored by a voltage threshold discriminator, and it was found that the threshold could be pushed low enough to be sensitive to the $S2$ pulse from a single ionization electron extracted into the gas Xe. Typical $S2$ pulses from a single electron consist of 24 phe spread across $0.8 \mu\text{s}$, and they appear to be uncorrelated with valid events in the fiducial volume (discussed further in Sec. 4.3.4). They are found to occur with equal probability at any point in an event waveform. A typical background γ event is shown in Fig. 4.6. The pulse at $\sim 103 \mu\text{s}$ is a single-electron $S2$.

During the non-blind WIMP Search 2 (WS2), the $S2$ -sensitive trigger threshold was $> 80\%$ efficient for single electron $S2$ events. This is shown in Fig. 4.7. Because of this, about 20% of digitized events contained only a single electron $S2$ pulse, and no $S1$ pulse. If single-electron $S2$ pulses were to correspond to a particle-induced energy deposition in the Xe target volume, the associated $S1$ would consist of only a handful (literally, < 5) of scintillation photons. This can be inferred from the AmBe calibration data (Sec. 6): the smallest $S2$ pulse from a 3 keVr (~ 3 $S1$ phe) single elastic nuclear recoil is about 250 phe (about 10 ionization electrons). Considering a total $S1$ efficiency (Table 2.4) of 10%, the such an event originally generated about 30 scintillation photons. From this one finds that a single-electron $S2$ would correspond to a nuclear recoil event with about 3 scintillation photons – a factor $\times 10$ below the $S1$ detection threshold in XENON10. The $S2$ trigger offers a significantly lower energy threshold ($\lesssim 2$ keVr) than does the $S1$ trigger (about 20 keVr, assuming 1 phe/keVr). This is shown in Fig. 4.5, and discussed further in Sec. 6. Even if the $S1$ trigger is operated with $n \geq 3$ coincidence, as shown in Fig. 4.5 (right), the threshold was only predicted to reach about 10 keVr. Additionally, operating the $S1$ trigger with a coincidence requirement less than $n \geq 6$ was found to have a high (> 5 Hz) rate of $S1$ -only events, with no associated $S2$ pulse in the event record. These events are due to the additional ~ 12 kg of passive Xe surrounding the active volume, and are rejected in the software analysis if they are digitized.

4.3.3 Thresholds trigger rate and DAQ dead time

After the (non-blind) WIMP Search 2 (WS2 ended 27 September 2006) it was confirmed that a single electron $S2$ pulse was too far below the $S1$ threshold to be useful for WIMP Search, and the threshold of the $S2$ trigger was increased to trigger on about 4 ionization electrons. It was explicitly verified with a square-wave pulser ($1 \mu\text{s}$ width) that the trigger efficiency was 100% for an $S2 > 100$ phe (4 ionization electrons), and about 75% for an $S2$ of about 80 phe (3 ionization electrons). The raw trigger rate during (blind) WIMP Search 3 using the $S2$ -sensitive trigger was 2.5 Hz. No high-energy veto was deployed. A DAQ dead-time of 7.5% was incurred due to the download time required by the VME connection to the PC. An additional 0.5 % dead-time was estimated due the effect of a rolling hold-off circuit [50] which prevented the DAQ from re-triggering until at least $200 \mu\text{s}$ after the last trigger. This rate remained stable for > 9 months of data-taking (Oct 2006 $\rightarrow$ May 2007), with interruptions only for planned calibrations.




Figure 4.6: The summed event record for a typical ~ 4 keVee background γ scatter, showing a 61 μ s drift time between $S1$ and $S2$, and also an isolated 25 phe single-electron $S2$ pulse at 103 μ s. The primary $S2$ pulse (at 80 μ s) triggered the DAQ in this event, however the $S2$ trigger was sensitive to single-electron $S2$ pulses as small as ~ 10 phe (see Fig. ??). The PMT hit-patterns are indicated for each pulse; the lower subplots show a zoom (and PMT hits on individual channels) on each pulse.

4.3.4 Single-electron $S2$

The generating mechanism of the observed single-electron $S2$ events is uncertain. Speculations include impurity-trapping of drifting electrons (with delayed release), and spontaneous boiling of electrons from the stainless steel cathode mesh into the liquid Xe. Photo-electric generation from the cathode mesh was considered, but the single-electron pulses do not appear to show a time-correlation with either $S1$ or $S2$ pulses. Further, it is difficult to see how they could be due to genuine energy depositions (scatters) in the active target because the rate (when measured in dru) is inconsistent with the observed rate 1 – 20 keVee (it is too high by about $\times 1000$).




Figure 4.7: Single electron $S2$ events from 2.52 live days of WS2 background data. Pulses with random event times (black) are shown separately from those which triggered the DAQ (grey). The former are well-fit by a Gaussian distribution (red), with a shoulder clearly attributable to double electron $S2$ events (blue). The population of double electron $S2$ events has $\mu_{2e-} = 2 \mu_{1e-}$, and $\sigma_{2e-} = \sqrt{2} \sigma_{1e-}$. An additional shoulder of pulses corresponding to 3 electrons is evident, but was not fit with a Gaussian due to the low statistics. The measured $\sigma = 6.3$ is 25% higher than that expected due to Poisson fluctuations; the excess is predicted by the width ($\sigma_r = 0.58$) of the PMT single photo-electron distribution. The distribution of pulses which triggered the DAQ (grey) is asymmetrical due to the trigger efficiency, which begins to drop slightly for events with $S2 < 15$ phe. Based on the asymmetry (comparing bin counts), the trigger efficiency during WS2 was still > 0.80 for single-electron $S2$ events with $S2 = 10$ phe. After WS2, a slight modification to the liquid level decreased the single electron $S2$ size by about 11%, to 24 phe per liquid electron.

The raw rate of spontaneous single-electron $S2$ events inside $R < 8$ cm (the z -coordinate is indeterminate without an $S1$) was found to be 0.22 ± 0.01 Hz, based on 2.5 live-days of WIMP Search data (from WS2, early september). This rate includes $S2$ events with 1, 2 and 3 electrons, as shown in Fig. 4.7. The pulses are clearly identifiable by their width, size and distinct hit pattern: about 59% of the proportional scintillation light falls on the top PMT array, and is localized above the (x, y) electron extraction point (see Fig. 4.6).

The $S2$ spectrum at threshold is shown in Fig. 4.7, with a peak corresponding to 1 electron and a shoulder corresponding to 2 electrons; there is an outline of an additional shoulder corresponding to 3 electrons. The events are separated according to those which triggered the DAQ (grey) and those found at random in event

records (black). The distributions for 1 and 2 electrons found at random locations in the event record are well-fit by a Gaussian. The measured value for a single electron is $\sigma_{1e-} = 6.3$ phe, which is predicted to within 5% by Poisson fluctuations ($n = 27$) in combination with the width ($\sigma_r = 58\%$) of the PMT single photo-electron distribution (discussed in Sec. 3.3.5). The expected width is $\sigma_{1e-} = \sqrt{n \cdot (1 + \sigma_r)}$. The typical value for single-electron $S2$ events during WS2 was 27 phe. After WS2, a slight change to the liquid level caused a decrease in the single electron $S2$ pulse size. This value measured from the AmBe calibration data and throughout the (blind) WS3 and WS4 data was 24 phe per liquid electron.

4.3.5 $S1$ peak-finding efficiency

Once an $S2$ triggered the DAQ, it was necessary to look back in the preceeding $80 \mu\text{s}$ to find the associated $S1$ pulse. In order to be confident that a legitimate $S1$ was found, the analysis required $n \geq 2$ coincidence, with a minimum of 0.35 phe/PMT. A Monte Carlo simulation was used to model the $S1$ peak-finding efficiency. As explained in Sec. 4.3.1, it accounted for the width of the PMT single photo-electron spectrum (Sec. 3.3.5), DAQ sampling effects (Sec. 4.1.2), the liquid Xe scintillation decay time of 27 ns and the 300 ns search-window imposed by the analysis software. Each ‘photon’ generated by the simulation was given a flat probability of 0.8/41 to hit a bottom-array PMT, and 0.2/48 to hit a top-array PMT. This simplifying assumption was found to be adequate for modeling $S1$ events that occur in the central 5.4 kg of the active volume. The quantum efficiency and single photo-electron width of each PMT were treated as identical (taking the mean for the 89 PMTs). The simulation result that the $S1$ peak-finding efficiency approaches 100% for $S1 > 4$ phe agrees to better than 1σ with the WIMP Search data, as shown in Fig. 4.8.




Figure 4.8: The $S1$ spectrum from WS4 58.6 live days of data, with 1σ error bars indicating Poisson fluctuation in bin counts, is compared with a Monte Carlo simulation (red line) of the efficiency to find an $S1$ pulse, given an $S2$ pulse trigger. The simulation input parameters are discussed in Sec. 4.3.5, and are identical to the inputs used for the $S1$ trigger efficiency simulation, but with a 300 μs search window instead of the 80 ns window imposed by the $S1$ -trigger discriminator.

4.4 Analysis algorithms

Analysis of XENON10 data was performed independently by two groups, one using Matlab and the other using ROOT; this discussion (and except where noted, this thesis in general) describes the Matlab-based analysis chain only. The analysis was broken down into 2 principle steps.

4.4.1 First pass: REEFrunner

The first step applied a generic, non-discriminatory parameterization known as REEF to each waveform. The acronym is short for “**R**eally **E**qual **F**ooting” (with emphasis on “**E**qual”). It refers to the assumption that XENON10 waveforms consist of a sequence of pulses, some of which may be wide ($S2$), some narrow with a fast rise-time ($S1$), and some which may be un-desireable (such as high-voltage discharge, or pile-up). The parameters extracted for each feature in an event record are the pulse areas and heights, and timing information such as the 50% rise and fall, the 10% rise and fall, and the position of the pulse maximum. The algorithm operated on each waveform after it had been summed across all 89 PMT channels, for maximum sensitivity to pulse shape, size and duration; however, information about pulse areas were stored individually for each PMT.

The speed of the first pass analysis was limited by the time required to load data from disk (about 30 ± 5 seconds per 104-event file, depending on CPU load and RAM usage). The application of the REEF algorithm to each event typically took less than 20 seconds. The REEF approach was originally motivated (in part) by the desire to avoid time-intensive data re-processing, by having a maximal amount of information available after a single pass of the analysis.

It was realized in August 2006 (shortly before WS2) that because very small (especially in the limit of single-electron $S2$) $S2$ pulses consist of only a few tens of photo-electrons spread out over about $1 \mu\text{s}$, such pulses tend to be discrete rather than continuous (see Fig. 4.6 or Fig. 4.9). The original REEF pulse-finding algorithm did not make sufficient provision for such pulses, and it was modified to include a *double – sliding – box – filter*. This filter integrated each candidate pulse with two different time-constants of 200 ns and $2 \mu\text{s}$. The time constants were chosen to correspond to the maximum width of $S1$ and $S2$ pulses. The effect for large (> 1000 phe) $S2$ pulses was negligible, but for very small $S2$ pulses the difference between with the two different time constants allowed a discrete single-electron $S2$ pulse to be distinguished from an $S1$ pulse. The algorithm is discussed in detail in [47].

A typical example event from the AmBe calibration data is shown in Fig. 4.9. The $S1$ (~ 19 phe) occurs at $44 \mu\text{s}$, followed by a single electron $S2$ pulse (~ 26 phe) at $62 \mu\text{s}$, and the main $S2$ (~ 1630 phe) at $87 \mu\text{s}$. With the original pulse-finding algorithm, a portion of the single electron $S2$ was flagged as an $S1$ pulse, due to the discrete nature of the photo-electrons. With the *double – sliding – box – filter* there the ambiguity was eliminated, due to the $\times 5$ difference in time-structure.

4.4.2 Second pass: decisions

With the time-structure and pulse-sizes of a waveform determined, a second-pass algorithm was run to decide which pulse was $S1$, which was $S2$, what was the corresponding drift-time, and a wash of other relevant reduced quantities (RQs). In principle, this step should be easy: a valid WIMP signal in XENON10 should contain only a single $S1$ and $S2$ pulse, with a drift time that is physically allowed by the dimensions of the detector ($0 - 80 \mu\text{s}$). However, the multiple $S1$ and $S2$ pulses – occasionally piled-up – found in some




Figure 4.9: A typical single elastic nuclear recoil from the AmBe calibration data. A zoom on the $0.25 \mu\text{s}$ extent of the $S1$ (which occurs at $t \sim 44 \mu\text{s}$ in the trace) is inset (top left); a zoom on the $1 \mu\text{s}$ extent of the single electron $S2$ (which occurs at $t \sim 62 \mu\text{s}$ in the trace) is also inset. The discrete nature of the photo-electrons in the single-electron $S2$ make it possible to mistake a portion of the pulse for an $S1$, *unless* a filter is employed in the pulse-finding algorithm. See Sec. 4.4 for details on the filter.

events presented a small challenge for approximately 1% of digitized events. But XENON10 is a rare-event search which collected about 10 million events during the blind WIMP Search (WS3 and WS4). The need for unambiguous identification of events cannot be over-stated.

A number of quality cuts were developed based on calibration data. These are discussed further in Sec. 4.6. The second-pass algorithm was tuned over a period of several months on the γ calibration data, and then fine-tuned on the AmBe calibration data. In this way, the analysis was honed for the acceptance of valid nuclear recoil events, the identification of background electron recoil events and the rejection of spurious events. The eventual production version of the decisions code was reached after visual inspection of many thousands of event waveforms (which is still a small fraction of the total data sample).

4.4.3 Position reconstruction

As discussed in Sec. 2.2.4 (and exemplified also in Fig. 4.6), the $S2$ signal is generated in the gas gap just above the (x, y) point where the ionization electrons were extracted from the liquid. This tends to give a diffuse wash on the bottom PMTs, but a highly localized hit-pattern on the top PMTs. The (x, y) position of interaction for each event was determined by comparing the $S2$ hit-pattern on the top PMTs with a simulated map.

The map was generated using GEANT4, by creating localized events ($\sim 10^5$ photons/event) at regular 1 mm intervals in (x, y) for the entire XENON10 gas gap region [122]. For each point on the grid, some

number of photons $\nu_i(x, y)$ were registered on the i^{th} PMT. The simulation accounted for the XENON10 geometry as well as the optical properties of all materials. The most-likely position of interaction was then chosen by a least-squares (often referred to as χ^2) minimization of

$$\chi^2(x, y) = \sum_{i=1}^{47} \frac{[n_i - \nu_i(x, y)]^2}{\sigma_i^2} \quad (4.1)$$

where n_i is the number of photo-electrons detected by the i^{th} PMT, and $\sigma_i^2 = \sigma_{n,i}^2 + \sigma_{g,i}^2$ is the total variance in the measured number of photo-electrons per PMT. The first term $\sigma_{n,i} = \sqrt{n_i}$ accounts for (Poisson) statistical fluctuations of photo-electron emission from the PMT photo-cathode, and the second term $\sigma_{g,i} = \sigma_r \sigma_{n,i}$ accounts for the standard deviation of the gain, as determined by the single-photoelectron spectrum ($\sigma_r = 0.58$ and is discussed in Sec. 3.3.5). This method of position reconstruction is also described in [122].

The χ^2 tends to be well-behaved, as shown in Fig. 4.11. The most-likely position of interaction was chosen by taking the (x, y) positions corresponding to the 5 smallest values of χ^2 (orange triangles), and taking the average. The underlying grid from the simulation is indicated by blue dots in Fig. 4.11. It was found that the position reconstruction results were similar to within 1% if a sparse grid (2 mm spacing) was used. The sparse grid reduced computation time by a factor of $\times 10$.




Figure 4.10: The event rate (dru) vs. enclosed radius for ^{137}Cs gamma calibration data (black circles), compared with Monte Carlo simulation (blue line), for events which deposit < 30 keVee in the active volume. Error bars account for ± 1 keVee uncertainty in energy determination, ± 1 mm (± 2.4 mm) uncertainty in reconstructed position in the center (edge) of the active volume, and $\sqrt{N}$ uncertainty due to fluctuations in bin counts. The discrepancy for each bin is shown on the plot; the slight undershoot in the bin $8 < R < 9$ cm is due to the finite extent of the PMT array. Two methods of position reconstruction are shown: least-squares (χ^2), and maximum likelihood. The performance is very similar for γ calibration data. For the smaller $S2$ signals generated by nuclear recoils, the maximum likelihood method is preferred.




Figure 4.11: An illustration of the position reconstruction. **(left)** The S2 hit-pattern on the top PMTs; the color scale indicates $\log_{10}(N_{phe})$, the blue + indicates the most-hit-PMT, and the green x indicates the most-likely position as determined by Eq. 4.1. **(middle)** A zoom on the region surrounding the reconstructed region, now also showing the underlying grid from the simulation (small blue dots), and the 5 lowest values of χ^2 (orange triangles); the chosen most-likely position of interaction (green x) is the mean of these 5 points. **(right)** The χ^2 values for a single event, showing about 6200 (simulation) grid locations in a roughly 40×40 mm area surrounding the most-hit PMT; outlying points on the grid (not shown) have even higher values of χ^2 . The 5 points on the grid which have the lowest values of χ^2 are highlighted (orange triangles).

Clearly the position reconstruction “works,” as can be seen in Fig. 4.11. But beyond visual inspection, its success can be assessed by comparing the event rate as a function of position for calibration data. High-statistics gamma calibration data was acquired using an external ^{137}Cs source. The global trigger rate without the high-energy veto was about 80 Hz, and with the high-energy veto (rolling-off at about 100 keVee) the trigger rate was 20 Hz. The event rate in dru is compared with a Monte Carlo simulation in Fig. 4.10. The rate is measured in concentric cylinders, using the entire z -range of 15 cm, for single-scatter events which deposit < 30 keVee in the active volume. The error bars allow for ± 1 keVee uncertainty in energy determination, ± 1 mm (± 2.4 mm) uncertainty in reconstructed position in the center (edge) of the active volume, and $\sqrt{N}$ uncertainty due to fluctuations in bin counts.

The agreement is very good, except in the region $8 < R < 9$ cm. This effect was expected as a consequence of the finite extent of the PMT array. In XENON10, we saw a tendency of the algorithm to mis-reconstruct events in the outer 1.5 cm, with a bias towards the outer edge. The effect leads to an undershoot of about 8% in the bin $9 < R < 10$ cm. All analysis used only the inner 8 cm (radially).

Method of maximum likelihood

A maximum likelihood method can be used in place of the χ^2 method described above. The χ^2 method relies on an accurate determination of the variance σ_i^2 , which in turn was estimated from the data. Since the variance depends on the Poisson variable n_i (the number of photo-electrons registered on the i^{th} PMT), it will be rather poorly estimated if $i \lesssim 5$. This would lead to an incorrect weighting of the contribution of the various PMTs when calculating the sum in Eq. 4.1. As an example, only about 6% of events in the WIMP Search data have $S2 < 2000$ phe. In contrast, events with $S2 \sim 1000$ phe are typical of a ~ 5 keVr nuclear recoil (Sec. 6). While the 4 top array PMTs nearest to the (x, y) locus of an event typically receive 5% of the $S2$ each, most PMTs receive an average of $< 1\%$ of the $S2$. For a 1000 phe event, most PMTs still receive ~ 10 phe. But the $S2$ fluctuates on an event-by-event basis, and for a low-energy nuclear recoil, it is not uncommon for several top-array PMTs to receive ~ 1 phe.

The most-likely position of interaction is calculated by maximizing the likelihood function

$$L(\nu_i(\theta)) = \prod f(n_i; \nu_i(\theta)) \quad (4.2)$$

where $f(n_i; \nu_i(\theta))$ is the probability to obtain the expected value ν_i on the i^{th} PMT. The expected value ν is a function of the estimator θ , which in the present case is the (x, y) position in the simulation grid. In general the maximum likelihood method is applied to un-binned data. However, the data is already binned into 48 PMTs. Maximizing the likelihood function in terms of binned data is equivalent to maximizing the likelihood ratio $\lambda(\theta) = f(n_i; \nu_i(\theta))/f(n_i; n_i)$ [166]. It is also equivalent to minimizing [25]

$$-2\log_e(\lambda(\theta)) = 2 \sum \left[\nu_i(\theta) - n_i + n_i \log_e \left(\frac{n_i}{\nu_i(\theta)} \right) \right], \quad (4.3)$$

where the quantity $-2\log_e(\lambda(\theta))$ follows a χ^2 distribution. It is clear from Eq. 4.3 that no explicit estimation of the variance is needed. The maximum likelihood method was not applied to the WIMP Search data prior to un-blinding. Globally, there is very little difference between the maximum likelihood and χ^2 methods; this can be seen in Fig. 4.10. However, on an event-by-event basis there can be differences of up to 5 mm. It is recommended that the maximum likelihood method be used for position reconstruction.

4.5 Position-dependent corrections

4.5.1 S1 z Correction

The collection efficiency for scintillation photons ($S1$) is dependent on the location of an event in the Xe target, for reasons discussed in Sec. 2.5. Generally, the collection efficiency for each PMT array decreases with an event's distance from the array. It is possible to correct for this dependence using a γ line source. Until the beginning of WS4 (01 Dec 2006), the best corrections were obtained using ^{137}Cs calibration data. The AmBe calibration data provided a more homogenous calibration source via the 40 keV γ from $\text{Xe}(n,n,\gamma)\text{Xe}$ inelastic scattering. After WS4 (but before finalizing the blind analysis parameters), a sample of activated Xe (the general procedure is discussed in [121]) that was introduced into XENON10 provided the best $S1$ calibration standard, via the 164 keV γ line from ^{131m}Xe .

The analysis is simple: the full-energy peak from the γ line is fit in slices of z , and the peak positions so obtained are then fit with a 2^{nd} -order polynomial. The data are adjusted by the transformation $S1_{t,b} \rightarrow S1_{t,b}/P_{t,b}(dt) \times P_{t,b}(40)$, where $P_{t,b}$ is the polynomial for either the top or bottom PMT array. The result is normalized to the center of the detector ($dt = 40 \mu\text{s}$). Note that a drift time of $80 \mu\text{s}$ corresponds to $z = 0$ (the bottom of the Xe target). The $S1$ z calibration obtained from the 164 keV data is shown in Fig. 4.12.




Figure 4.12: The z dependence of the $S1$ light collection, as evidenced by the varying photo-electron yield (per keV) with drift time. The data are from the 164 keV gamma line from activated ^{131m}Xe . As expected, the light collection on the bottom (top) PMT array increases (decreases) with drift time.

4.5.2 $S2$ (x, y) correction

As described in Sec. 2.2.4, the $S2$ proportional scintillation is generated by the high-field ($E_e \sim 11 \text{ kV/cm}$) in the gas gap above the liquid Xe. Because the electric field E_e is created by applying a potential difference to the stainless steel mesh grids, it is sensitive to any change in the grid spacing. For proportional scintillation, eq. 2.1 predicts that n_{ph} (the number of proportional scintillation photons produced) changes linearly with E_e . The mesh grids were electro-formed and are susceptible to a small degree of sagging. For example, a

decrease of 2 mm in the gas gap (say from 3.5 mm to 1.5 mm) is predicted to decrease n_{ph} by about 27%, from 210 to 152 proportional scintillation photons per electron. Figure 4.13 shows the $S2$ response from the 40 keV inelastic γ obtained from the AmBe calibration data. The color scale indicates number of PMT photo-electrons $\times 10^4$. A typical 25% change in $S2$ from the center to edge ($R \sim 90$ mm) is observed, and is easily corrected.




Figure 4.13: The $x - y$ map of $S2$ response to the 40 keV inelastic single-scatter γ from AmBe calibration data. The color scale indicates $S2$ phe $\times 10^4$. The correction map was made by [98]

The $S2(x, y)$ correction was not applied to the WIMP Search data by the Matlab analysis chain prior to unblinding, though it was applied by the ROOT analysis chain. In hindsight, this looks like a considerable oversight. At the time, it was deemed to be of secondary importance. The reason was that XENON10 used an $S1$ -based energy scale, so this correction did not influence the determination of event energy or the WIMP Search window. The effect of this correction on the discrimination is significant in the sense that it reduces the Gaussian width of the $\log_{10}(S2/S1)$ distribution in any given energy bin (compare Sec. 2.7 and Fig. 2.18).

However, the effect of this correction on a non-Gaussian background event is very modest: it will still appear to be non-Gaussian after the correction. Leading up to the unblinding of the XENON10 data, there was a much sharper focus on eliminating non-Gaussian background events from the WIMP Search window. Some leakage of the Gaussian background into the WIMP Search window was expected, and at was thought for some time that XENON10 might employ background subtraction based on this expectation. In summary this correction should certainly be applied, and it has been applied to all data presented in this thesis except the result of the Matlab blind analysis in Fig. 5.2. This is discussed further in Sec. 5.1.

4.5.3 PMT relative sensitivity correction

The relative sensitivity of the top and bottom PMT arrays was found by comparing the $S1$ signal in nearest-neighbor PMTs, for events that were determined to lie within $R < 20$ mm of the midpoint between the 2

PMTs; this is shown in Fig. 4.14 and Fig. ???. The relative response of each nearest-neighbor pair of PMTs is fit with a Gaussian (Fig. 4.14, right). Once all nearest-neighbors have been compared (there are as many as 6 nearest-neighbors for a PMT in the center of the array), the matrix of relative values is minimized to find the global relative sensitivity (RS) for the array. The RS for the top PMT array is shown in Fig. 4.15; for each PMT, the RS is indicated in red in the lower left corner, and shown graphically (light red square). The RS tends to vary by $< 5\%$ across the array, with PMT 5 and PMT 25 as notable (less sensitive) exceptions.




Figure 4.14: **(left,top)** The response to $S1$ is compared for PMT 6 and PMT 12 (highlighted in red); events for the comparison are drawn from within the blue circle, $R < 20$ mm from the midpoint of the 2 PMTs. **(center,top)** Scatter plot of the response of each PMT relative to the other; events lying on the green line ($slope = 1$) had equal response in each PMT. Events were required to produce ≥ 3 phe in each PMT to be considered. **(right,top)** The scatter plot is binned and fit with a Gaussian; the mean from this fit (here $\mu = -0.01$) is taken as the relative sensitivity between the 2 PMTs. **(bottom plots)** This is an example in which the relative sensitivity of the 2 PMTs is substantially different from 1 (specifically, PMT 5 is about $\times 2$ weaker than PMT 6).

This same analysis applied to the bottom array with a very similar result in terms of the magnitude of the correction. Only PMT 75 was found to be weaker than the others, by about 15%. The solution which found the global RS for each array could not yield a global normalization. Normalization was applied afterward, such that each array has a mean value of $RS=1$. This presented a systematic uncertainty if one were to apply the RS correction to both PMT arrays. To avoid such a situation, the RS correction was only applied to the top PMT array, where the information bears on position reconstruction in addition to an event $S2$ value. The relative sensitivity correction was found to be less important than the $S2(x,y)$ correction.




Figure 4.15: The relative sensitivity of the top array PMTs, as determined by the method described in the text.

4.5.4 $S2$ electron lifetime correction

As discussed in Sec. 2.4, it is important that the lifetime τ_{e-} of ionization electrons in the liquid Xe target be stable, and long enough that the loss of signal (due to impurity attachment) is negligible. This condition can be stated in terms of the mean free path for ionization electrons $\lambda_{e-} = v_d \cdot \tau_{e-}$, as $\lambda_{e-} \gg h$. The length scale is set by $h = 15$ cm, the full drift length of the Xe target. The condition is reasonably satisfied for $\lambda_{e-} > 1.5$ m, which corresponds to $\tau_{e-} \simeq 800 \mu\text{s}$. The loss due to attachment in this case is 10%. After ~ 6 weeks of continuous purification, τ_{e-} was found to reach a plateau at 2 ms ($\lambda_{e-} = 380$ cm). An electron lifetime $\tau_{e-} > 2$ ms gives a 3.9% maximum reduction of the ionization signal for an event which happens at the bottom of the active volume. This effect is easily adjusted for by the transformation $S1 \rightarrow S1/e^{-dt/\tau_{e-}}$, where dt is just the drift time in μs of the event (equivalent to the z -coordinate of interaction).

The Xe purity is best assessed by finding the mean value of the $S2$ signal for a gamma line source (such as ^{137}Cs , or the 40 keV peak in the AmBe calibration), in slices of z . This is shown in Fig. 4.16 for the 40 keV peak in the AmBe calibration data. The peak mean values from Fig. 4.16 are shown vs. their respective bin centers in Fig. 4.17. The electron lifetime is calculated by fitting an exponential to the data, as shown in Fig. 4.17. The error on tau is calculated from the fit. During WIMP Search data-taking, γ calibration data was taken periodically (about 1-week intervals).

In order to continuously monitor the electron lifetime, an alternative method was tested: instead of fitting the decrease in $S2$ across z , the decrease in $\log(S2/S1)$ was measured, with the slope of a linear fit giving tau. An obvious draw-back of this method is a systematic error due to including the $S1$ value. However, the $S1$ is corrected for z -dependence of light collection, and normalized to the center of the detector (see Sec. 4.5.1). The advantage of the method is that no γ line source is required, and the electron lifetime τ_{e-} can




Figure 4.16: Gaussian fits to the 40 keV $S2$ peak in the AmBe calibration data, in slices of z . A slight decrease is observed from top (shorter drift time) to bottom of the detector; from this, the electron lifetime τ_{e-} can be extracted. $1\text{-}\sigma$ uncertainty on the mean is calculated from the covariance matrix from the fit.

be obtained from WIMP Search data (at an energy above the WIMP Search region of interest). By testing both methods on several Cs (as well as the AmBe) data sets, it was determined that the systematic error is roughly a constant, and gives a τ_{e-} value that is too low by $\times 1.2 \pm 0.2$. The electron lifetime improvement from the beginning of WS2, into the beginning of WS4 is shown in Fig. 2.9. About half the data points in this plot were obtained by the $\log(S2/S1)$ method.




Figure 4.17: Data from the fits (Fig. 4.16) to $S2$ in slices of z (red circles), shown with the underlying $S2$ data (blue dots), and the best fit exponential (red line). The electron lifetime is $\tau_{e-} = 2175$ for this data. $S2$ data from the first and last $5\text{ }\mu\text{s}$ are excluded from the fits. The $S2$ from nuclear recoils (~ 3000 phe) is not used for this analysis.

4.6 Data quality cuts

The event template for a galactic WIMP which scatters from a Xe nucleus in XENON10 is shown in Fig. 4.18. This event is a single scatter elastic nuclear recoil from the AmBe calibration data. A particle interactions with multiple scatter vertices (characterized by a single $S1$ and 2 or more $S2$ pulses) is one of the more

obvious examples of a particle interaction which should *not* be considered as a WIMP candidate. This is justified because WIMPs (whatever they are) interact so weakly that we have yet to observe a single scatter — let alone two in $80 \mu\text{s}$. In contrast, the dominant γ background often results in multiple vertices within the detector. Of course, XENON10 is capable of many other interesting signals that do not fit the template, as well as particular pathologies which must be rejected. These are described below.




Figure 4.18: A classic example of the expected signature of a WIMP interaction with a Xe nucleus in XENON10. The interaction must have only one scatter vertex (one $S2$ pulse), and an $S2/S1$ ratio typical of a Xe nuclear recoil. This is event number 29 from the AmBe neutron calibration data.

The cuts are grouped into three levels: the zeroth-order (QC0, short for Quality Cuts level 0) group of cuts asks little more than that events fit the template for a particle interaction in the Xe target volume. The 1st and 2nd levels (QC1 and QC2) are targeted at improving the discrimination of the detector. Each cut is discussed in turn.

4.6.1 Basic quality cuts: QCO

The cuts outlined in this section represent the basic quality cuts as implemented in the blind analysis of XENON10 data. The cuts were initially developed during the first months of operation, which were almost entirely devoted to γ calibration. The acceptance of the QC0 cuts for genuine events was found to be $\sim 99\%$, independent of event energy. The acceptance of QC0 for nuclear recoils is shown in Fig. 6.5.

Good baselines

So that an accurate assessment of the baseline can be made, none of the first 100 samples (in each channel) may exceed $\times 3$ of the baseline flattening threshold for that channel. As discussed in Sec. 4.2.2, the thresholds are set individually for each channel. The thresholds are summarized in Fig. 4.4.




Figure 4.19: An event which did not pass the “good baselines” requirement. Inset (left) is the first $0.5 \mu\text{s}$ of the trace, showing what is most likely a single photo-electron in a single channel (this exceed $\times 3$ of the flattening threshold). All the other channels had no excursions beyond the 8 ADC bin hardware threshold, and so are recorded as zero. The baseline with the single photo-electron is slightly higher than zero. For comparison, the $S1$ pulse is also shown inset (center). In this case, the error in the baseline determination lead to only a minor error in the pulse $S1$ and $S2$ determination.

Trace edges

In order to make an accurate determination of the pulse areas, it is imperative that no portion of a pulse over-run the edges of the event record. The cut requires the $S1$ start-time ($S1_t$) $> 0.5 \mu\text{s}$, and the $S2$ start-time ($S2_t$) not in the last $2 \mu\text{s}$ of the trace. Pulse start-times are defined at the 10% rise of the leading edge.

$S2$ fraction of trace area

In a clean event such as Fig. 4.18, the $S2$ pulse should account for a majority of the total event area. This would not be the case if (for example) the baseline had wandered significantly, or if there were a large amount of light created by an electrical breakdown. The safe-guard against this is to require that $(S2 \text{ area}) / (\text{total event area}) > 0.5$.

$S2$ top versus bottom

Both data and ray-tracing Monte Carlo simulations both agree that about 58% (42%) of the $S2$ signal should fall in the top (bottom) PMT array. Events in which an $S2$ -like pulse is identified, but with a disproportionate amount of signal in the bottom array almost always implies a mis-characterized event. For example, the $S1$ pulse from an α particle can have sufficient width that it gets tagged by the $S2$ pulse-finder; however, such a pulse has $\sim 80\%$ of its signal in the bottom array. Another very typical example is shown in Fig. 4.21, and is speculated to be the tail end (photon shot noise) of a large pulse (probably $S2$). To remove such pulses, but also allow for reasonable fluctuations in the top/bottom split of the $S2$ signal, the cut requires $\Sigma(\text{Top}) / \Sigma(\text{bottom}) > 0.54$.




Figure 4.20: A ~ 150 keVee event in which the $S2$ pulse runs off the end of the digitizer. The $S1$ and $S2$ are shown inset (and offset from zero) in blue, magnified $\times 10$ in μs . This pathology is possible because the maximum event drift time is $> \times 2$ the maximum record length. Below about 75 keVee, $> 99\%$ of triggers are caused by $S2$ pulses. Therefore, in WIMP Search mode the trigger is delayed so that an $S1$ will not happen before the start of the digitizer. As this event shows, the converse is possible; this pathology is more relevant for γ calibration than for WIMP Search.

$S2$ width

The $S2$ width must be bounded. In addition to the physical timing parameters, the *double – sliding – box – filter* (described in Sec. 4.4) generates width parameters in filter-space. For the blind analysis, the physical width of $S2$ pulses was required to be $0.45 \mu\text{s} < S2_w < 2.25 \mu\text{s}$. The upper bound on the width was particularly relevant in its ability to remove a class of events whose origin is still uncertain, as shown in Fig. 4.21 (bottom). However, this was not a good choice in terms of maximizing acceptance for low-energy nuclear recoils, since as pulses become smaller the discrete arrival time of photo-electrons can lead to an incorrectly low assessment of pulse width. This is exemplified in Fig. 5.13.

Single $S1$

Only a single $S1$ pulse is allowed in the event window corresponding to the physical drift time before the $S2$. This concern arises due to the ~ 12 kg of passive (no electric field) Xe. $S1$ pulses originating in the passive Xe (which cannot have an $S2$) can be time-coincident with a scatter in the active volume. It can then be impossible to determine which $S1$ pulse corresponds to the $S2$ pulse.

Single $S2$

Only a single $S2$ pulse is allowed in the event record. This is an *a priori* assumption for WIMP Search, and is also necessary during calibration, for example in order to understand the position-dependence of the light collection for $S1$ signals. A secondary $S2$ pulse is required to be at least 5% of the primary $S2$ for it to be considered a multiple scatter.




Figure 4.21: **(top)** A classic example of an “event” which triggered the DAQ, but clearly does not fit the desired event template. The rate of these events is low ($< 0.5\%$ of triggers), and additionally they tend to be flagged by more than one of the QC0 cuts. **(bottom)** An event with anomalous width, with top (blue, offset +10) and bottom (red) PMT arrays plotted separately. Less than $\sim 0.5\%$ of digitized events exhibit this pathology. The cause is uncertain, but is speculated to be high voltage discharge.

PMT saturation

Saturation is defined as any PMT having an entry in the maximum digitizer bin. For WIMP Search, the expected signal size makes it reasonable to require no ADC saturation; for γ calibration, this requirement needs to be relaxed. PMTs on the top array just above the (x, y) vertex tend to saturate for scatters above ~ 200 keVee.

Pulse aspect ratio

Typical $S2$ pulses have a bounded aspect ratio (*pulse height* / *pulse area*). Figure 4.22 shows an event in which a particle scattered just below the liquid-gas interface: the $S1$ at $80 \mu\text{s}$ is immediately followed by an $S2$. This is often characterized as a single $S2$ pulse by the analysis algorithm, however a simple check on the aspect ratio confirms that it is anomalous. For reference, the $S2$ pulse shown in Fig. 4.18 is very nearly the same number of phe, but its height is lower by half.




Figure 4.22: An event in which a scatter occurs just below the liquid-gas interface: the $S1$ at $80 \mu\text{s}$ is immediately followed by an $S2$. This is often characterized as a single $S2$ pulse by the analysis algorithm, however a simple check on the aspect ratio confirms the pathology. Such an event is only a problem if a valid (coincidental) $S1$ is also found, such as the pulse at $\sim 45 \mu\text{s}$.

$S1$ pulse shape requirements

- **$S1$ coincidence:** a realistic threshold for our ability to resolve genuine, small $S1$ from random single phe events is established via a minimum coincidence requirement. In the $250 \mu\text{s}$ time-window surrounding an $S1$ candidate pulse, $(S1_i > 0.35 \text{ phe}) \geq 2$ was required. The subscript i runs over the 89 PMTs.
- **$S1$ width:** In order to ensure that the $S1$ pulse found by our algorithm is a legitimate $S1$ (as opposed to a small $S2$), **$S1$ Width** cut requires its 50% width ($S1_w$) $< 0.30 \mu\text{s}$.
- **$S1$ rise time:** 10%-100% rise time ($S1_{rt}$) $< 0.05 \mu\text{s}$.

***S2* minimum size**

For *S2* pulses smaller than ~ 300 phe, the *S1* signal tends to be too small to detect. This is discussed in Sec. 6. Such events tend to fall below the 3σ WIMP acceptance box, but nevertheless it is preferable to exclude them from the analysis. The cut requires (*S2 size*) > 300 phe.

4.6.2 Fiducial volume cuts: QC1

A major advantage of a 3D imaging TPC such as XENON10 is the ability to draw a (smaller) fiducial target volume within the active (instrumented with electric field) volume. There are two primary reasons why this is desirable:

- **Self-shielding** Excluding the outer few cm of Xe in (r, z) allows the detector to exploit the inherent self-shielding afforded by the high density (about 3 g/cc, see Table 2.1) and radio-purity of Xe. As discussed in Sec. 3.2.7, the mean free path of a 1 MeV photon is $\lambda = 5.8$ cm, which leads to a flux reduction of about 30% in 2 cm. For lower-energy photons the flux reduction is greater 5.2.3 [stopping power Fig], and contributes to the high quality of the nuclear recoil data from the XENON10 AmBe calibration (see Fig. 5.2.3).
- **Edge effects** Besides the benefits of self-shielding, avoiding the edges of the detector is crucial to maximize the discrimination of nuclear recoils from the electron recoil background. The primary pathology associated with the edges of the active volume is the loss of ionized electrons (*S2*). This can happen because electric field lines terminate on the Teflon walls or on the mesh electrodes (discussed in [72]). One would expect that such an effect be relegated only a thin (few mm) shell at the very edge of the active volume; in practice the extent of the problem is limited by the accuracy of the position reconstruction, which is at its worst at the edges of the detector (discussed in [122]).

Definition of the (blind) fiducial volume

The fiducial volume as used for the blind analysis (and all analysis presented in this work) of the WIMP Search data defined a 5.4 kg Xe target, and was drawn in advance of the unblinding as:

- radially: $r < 80$ mm
- vertically by drift time as $15 \mu s < dt < 65 \mu s$
- an additional arc was drawn to exclude events near the edge PMT48, which was dead since the beginning of WS2. The arc is defined in mm, for an event with coordinates (x, y) as: $((x - X_{48} + 13)^2 + (y - Y_{48} - 14)^2) < 47^2$. The coordinates of the center of PMT48 are (X_{48}, Y_{48}) .

The fiducial volume used for the XENON10 blind analysis (henceforth referred to as simply “the fiducial volume”) was determined by a study of the ^{137}Cs calibration data (equivalent to about 60 live-days) and the unblind background data from WS2 (xx live-days) and WS3/WS4 (16 live-days). Only the ^{137}Cs calibration data is shown; the WS data behaves very similarly. The primary goal was to minimize the leakage of non-Gaussian background events into the WIMP Search acceptance box. The procedure used to draw the acceptance box is detailed in Sec. 5.1. Here we simply note that it differs from Fig. 2.18 only by a linear change of coordinates, and present the ^{137}Cs data in this modified form (Fig. 4.23) as motivation for the chosen definition of the fiducial volume.




Figure 4.23: Electron-recoil (γ) events from 60 live-day equivalent of ^{137}Cs calibration data remaining inside the blind fiducial volume (5.4 kg Xe) are shown in transformed coordinates, such that the centroid of the (Gaussian) electron recoil centroid lies at zero. The nuclear recoil centroid from the AmBe calibration data is shown in the same coordinates (upper red line), along with its respective -3σ contour (lower red line). The transformation is described in Sec. 5.1.1. The WIMP Search acceptance box, with 50% acceptance for nuclear recoils, is bounded by these red lines. There are 1866 events in the energy range 4.4 – 26.9 keVr (between the black vertical lines). Events ringed in blue/red are γ events “leaking” into the WIMP Search acceptance box that are statistically consistent with the Gaussian electron-recoil population (see Table 4.1). The green curve is the lower 99.5% contour from the electron recoil population (-2.58σ), and the blue curve is the lower 99.9% contour from the electron recoil population (-3.09σ).




Figure 4.24: Top-down view of the (x, y) distribution of single-scatter electron recoil events from 60 live-day equivalent of ^{137}Cs calibration data, in the energy window $4.5 - 26.9$ keVr (assuming $\mathcal{L}_{eff}=0.19$). The bottom PMT array is not drawn. The solid black line indicates the edge of the Teflon, and the solid blue lines indicate the 5.4 kg fiducial target volume chosen for the blind analysis. The arc around PMT 48 represents an excluded region to compensate for the dead PMT. The blue and red circles indicate events remaining in the acceptance box after all cuts, as described in Fig. 4.23. **The choice of fiducial volume** can be understood from the distribution of non-Gaussian events remaining after all cuts (QC0 and QC2) other than fiducial cuts (QC1). Pink circles indicate events with $\log_{10}(S2/S1)$ values that are excluded from the electron recoil population at the 99.9% level (-3.09σ), and that passed all QC0 and QC2 cuts. These would be “leakage” events, in the sense that they are distributed throughout the WIMP Search acceptance box (shown in Fig. 4.23). Reminder: events outside the fiducial volume are not shown in Fig. 4.23.




Figure 4.25: A cross-section $r - z$ view of single-scatter electron recoil events from the ^{137}Cs calibration data in the energy window $4.5 - 26.9$ keVr (assuming $\mathcal{L}_{eff}=0.19$). This is the same data sample and color-coding as shown in Fig. 4.24. See the caption of Fig. 4.24 for details. Shown here, but not in Fig. 4.24 are potential “leakage” events that were cut by QC2 (green circles). This shows the performance of the QC2 cuts outside the fiducial volume. In particular, all non-Gaussian events in the region $R < 80$ and $65 \mu\text{s} < dt < 80 \mu\text{s}$ were removed. The high rejection efficiency in this region is partially attributed to the non-blind analysis of the ^{137}Cs calibration data (which is an expected effect, since this data was used to train the cuts).

4.6.3 The need for advanced cuts (QC2)

In addition to basic quality cuts (QC0) and fiducial volume cuts (QC1), advanced cuts targeted at pathologies which generate non-Gaussian tail events (QC2) were applied to the ^{137}Cs data shown in Fig. 4.23. The motivation for developing QC2 cuts can most easily be understood by re-examining the ^{137}Cs data as presented in Fig. 4.23, but without QC2 applied; this is shown in Fig. 4.26.




Figure 4.26: Electron recoil (γ) events from 60 live-day equivalent of ^{137}Cs data, with cuts QC0 and QC1 only. This is identical to the data as shown in Fig. 4.23, except that here cuts QC2 are not applied. There are 2273 events in the energy range 4.4 – 26.9 keVr, so the acceptance of the QC2 cut for the ^{137}Cs calibration data is 0.821 in this energy window. There are 24 events (compared with 7 in Fig. 4.23) from the Gaussian electron recoil population that are “leaking” into the WIMP Search acceptance box (blue/red circles). A dramatic difference is seen in the number of distinctly non-Gaussian events (red circles); this is summarized in Table 4.1.

The conclusion based on comparing Fig. 4.23 and Fig. 4.26 (as summarized in table 4.1) is that without the QC2 cuts, there is distinct excess of events with low values of $\log_{10}(S2/S1)$ that are not consistent with the Gaussian-distributed electron recoil population. It is also clear that the QC2 cuts do an almost-too-good job at removing these events (compare predicted with actual percentages in Table 4.1). This should not be surprising: the ^{137}Cs calibration data (along with the non-blind WS data, not shown) constituted the “training” set for the cuts.

Table 4.1: Performance of cuts QC2 on ^{137}Cs calibration data. Actual number of events 4.5 – 26.9 after cuts QC0 and QC1 are tabulated, with percentage shown as (%). The 99.5% (-2.58σ) and 99.9% (-3.09σ) curves are single-sided, so that it is predicted to find 1 event in 1000 below the 99.9% curve. The predicted percentages are shown in the final column for comparison.

	QC0 & QC1	+QC2	predicted
Total events 4.5 – 26.9 keVr	2273	1866	–
Events below ER 99.5%	23 (1.01%)	6 (0.32%)	0.50%
Events below ER 99.9%	11 (0.48%)	1 (0.05%)	0.10%

4.6.4 The origin of non-Gaussian tail events

The most obvious mechanism for an electron recoil event with an abnormally low (non-Gaussian) value of $y = \log_{10}(S2/S1)$ is a loss of some portion of the ionization created by a γ scatter. There are two principal scenarios; the first is tended to in large part by QC1 (discussed in Sec. 4.6.2). This refers to events whose vertex is in a region of the detector where the external electric field lines diverge or end, particularly near the walls of the Teflon or near the anode or cathode grids. The second scenario is more subtle, and arises when a γ double-scatters, but an $S2$ pulse is generated for only one of the vertices. The potential for this to happen is (at present) an unavoidable feature of a liquid Xe TPC detector such as XENON10. The need for ~ 1 kV/cm of electric field requires that a large potential be applied to the cathode, directly above the lower PMT array. This is shown in Fig. 3.29. A gap between the cathode and the lower PMT array is thus unavoidable. In XENON10 it was set (perhaps conservatively) to 1.3 cm, and in this region ionization drifts away from the gas gap. Thus about 8% of the Xe between the top and bottom PMT arrays can generate $S1$ but no $S2$. In general, an event in which an multiple scatter vertex generates $S1$ but no $S2$ is referred to as a “Gamma-X” event. A pictorial example of the problem is shown in Fig. 4.28.

4.6.5 Description and performance of QC2 cuts

In order to better understand the performance of the QC2 cuts on the ^{137}Cs calibration data, it was helpful to consider a slightly larger fiducial volume. This allowed for an increase in the data sample size, without contamination of events from the edges of the active volume. Data presented in Fig. 4.27 is identical to that in Fig. 4.23, with three slight modifications: events from a larger, 8.4 kg fiducial volume are considered, events consistent with the Gaussian population are not tagged with blue circles and non-Gaussian events are tagged according to which cut removed them.

It should be mentioned that prior to unblinding, the QC2 cuts were assessed by their performance on all events below the nuclear recoil centroid (top red curve in Fig. 4.27). This was a mistake, and it is not repeated here. The reason is that events which leak into the acceptance box simply because they are on the low side of the Gaussian distribution are (by definition) not targeted by QC2. To be clear: an event that is consistent with the Gaussian population (for example, a blue-circle in Fig. 4.23) *could* be a “Gamma-X” event, but it is not a good idea to set QC2 cut parameters based on events which are (say) 10% likely to be part of the Gaussian population.

The performance of the QC2 cuts is assessed only on events which are excluded from the Gaussian population at the 99.9% level (below the blue curve in Fig. 4.27). It is found in retrospect that three of the

QC2 cuts (*PAF*, *S1W*, *PGX*) acquired an inflated sense of importance, and should not have been used. The cuts are discussed (and listed in the legend of Fig. 4.27) according to their utility, which is defined loosely in terms of the number of non-Gaussian events they cut. The acceptance of the QC2 cuts is discussed in Sec. 4.44.




Figure 4.27: Single-scatter γ calibration data (^{137}Cs) from a larger 8.4 kg fiducial volume, in the energy window 4.5 – 26.9 keVr. Cuts QC0 and QC2 are also applied. Each non-Gaussian event (below the -99.9% contour) that was cut by QC2 is marked according to which cut removed it. Several events are targeted by more than one cut. The cuts are indicated by abbreviated names in the legend, and described further in the text. Six events remaining after all cuts are indicated by red circles (note that five of these are outside the blind 5.4 kg fiducial volume)




Figure 4.28: A schematic summary of the Gamma-X problem: **(left)** a normal double-scatter γ event with $S1 = 13$ phe at $35 \mu s$, and $S2 = 2789$ phe (the sum of the two $S2$ pulses) has $\log_{10}(S2/S1) = 2.34$, which is very typical of γ scatters. **(right)** if the second scatter ($S2 = 1156$ phe at left) had occurred in the reverse-field region (RFR) below the cathode grid, there would be no $S2$ for that vertex – however the $S1$ signal in both cases would be roughly the same. The result would be an event which looks like a single-scatter with $\log_{10}(S2/S1) = 2.10$, which is excluded from the distribution of electron recoils at more than 3σ (see Fig. 2.18).




Figure 4.29: The $r - z$ view of the events shown in Fig. 4.27. The 8.4 kg fiducial volume is indicated by blue dashed lines, and the 5.4 kg (blind analysis) fiducial volume is indicated by solid blue lines. Events outside the 8.4 kg fiducial volume are not shown in Fig. 4.27, and are indicated here in grey.

S1 reverse-field region cut (RFR)

The *S1 RFR* Gamma-X cut is the single most-important QC2 cut, and the origin of its pathology has already been exemplified in Fig. 4.28. In summary the cut attempts to identify γ events which at first-glance look like a single scatter (only one *S2* pulse), but on closer inspection may actually have a second scatter vertex in the reverse-field region (below the cathode grid). The means to identify this pathology relies on the *S1* hit pattern. The *S1* hit pattern for a 23 keVr event (indicated by a $+$ at $\Delta \log_{10}(S2/S1) = -0.54$ in Fig. 4.27) is shown in Fig. 4.30. It is easy to tell that one of the bottom PMTs has received a disproportionate fraction (8 phe, or 35%) of the total *S1* signal. This is the classic signature of an event in which a scatter vertex lies only a few mm above a bottom-array PMT (in the reverse-field region). Non-Gaussian background events arising from this pathology were predicted during XENON10 R&D. A detailed study of *RFR* events was made using a GEANT4 simulation of the XENON10 detector, and is described in [50]. Here the focus of the study is more phenomenological.




Figure 4.30: The *S1* hit pattern for an event tagged by the QC2 *RFR* cut. The anomalous concentration of signal in a single PMT on the bottom array is highly unlikely for a “normal” single-scatter event. The event is indicated by a $+$ at (23 keVr, $\Delta \log_{10}(S2/S1) = -0.54$) in Fig. 4.27). The *S2* for this event located it at $r = 62$ mm and the drift time is $55 \mu\text{s}$ (see Fig. 4.29).

An index $G_i = S1_i / S1_{bot} \cdot \sqrt{S1}$ was defined to reflect the extent to which an individual event might have suffered a scatter in the reverse-field region. The R8520 PMTs have an active area of about 6 cm^2 , and are densely packed in the bottom PMT array (shown in Fig. 2.10 and discussed in Sec. 2.5.4). It is therefore important to consider the possibility that a scatter in the reverse-field region will have a significant fraction of its *S1* signal split across up to 4 PMTs. The subscript $i = 1 \dots 4$ and $S1_i$ is the *S1* signal in each of the 4 bottom array PMT which received the most (2^{nd} -most, etc) signal for that event. The factor $\sqrt{S1}$ is added to remove the otherwise sharp energy dependence G_i . The index G_1 for the ^{137}Cs calibration data is shown in Fig. 4.31.

The same Monte Carlo simulation that was used to model the *S1* peak-finding efficiency (see Sec. 4.3.5) was easily adapted to model G_i ; the result is shown in Fig. 4.32 for $i = 1$. The $n \geq 2$ coincidence requirement was applied to the Monte Carlo. Figure 4.33 compares the ^{137}Cs calibration data with the Monte Carlo

prediction. The cut band for G_1 (red line, Fig. 4.31 and Fig. 4.32) was initially drawn such that the Monte Carlo predicted a roughly constant acceptance of 95%. It was then further tuned according to its ability to cut non-Gaussian events in the ^{137}Cs and un-blinded WIMP Search data.




Figure 4.31: The index $G_i = S1_i/S1_{bot} \cdot \sqrt{S1}$ for $i = 1$ from the ^{137}Cs calibration data. The cut band as used in the blind analysis is indicated by a solid red line; 4 non-Gaussian events (red circles) are removed by this cut alone. Similar plots for G_2, G_3 and G_4 are shown in Fig. 4.35 – Fig. 4.37. Events excluded from the electron recoil population at more than 99.9% are indicated by red circles ($\circ$); the 11 red circle events between 4.5 – 26.9 keVr are also shown in Fig. 4.26 (in the same energy window). Severely non-Gaussian events with $\log_{10}(S2/S1)$ below the nuclear recoil -3σ contour are indicated by a blue $\times$.

The performance of the *RFR* cut for the cases $i = 2 \dots 4$ are shown in Fig. 4.35 – Fig. 4.37. In these three cases, the cut implementation is identical to the $i = 1$ case, but with an additional feature. The feature is a second, more aggressive cut band (systematically lower G_i), with the caveat that this second cut band only applies if the excess $S1$ in the bottom array is localized. Localization is defined simply as $\sigma_{S1}(x) = \Sigma(S1_i \cdot x_i)/S1_{bot}$, and similarly for y . The point of this is to increase the efficacy of the cut, while maximizing the cut acceptance. Hit patterns for event numbers 5 and 8 are shown in Fig. 4.34 to illustrate the localization check. Both events are non-Gaussian; event number 5 shows sufficient localization to fail cut (see Fig. 4.35), while event number 8 does not show sufficient localization (and passes all the G_i cuts). It can be seen from Fig. 4.31 and Fig. 4.35 – Fig. 4.37 that the cuts are not completely orthogonal. For example, event number 11 was not cut by G_1 , but it is cut by G_2, G_3 and G_4 . On the other hand, event number 3 was cut by G_1 and G_3 but not the other two.




Figure 4.32: The index $G_i = S1_i/S1_{bot} \cdot \sqrt{S1}$ for $i = 1$ from a Monte Carlo simulation of $S1$ photo-electron signal generation. The simulation is described in Sec. 4.3.5. The phe scale was converted to keVr by the same scaling applied to the data, and the $n \geq 2$ coincidence requirement was applied to the simulation data. The arc (of Monte Carlo data points) along the upper left of the plot corresponds to events in which all of the bottom array signal is concentrated in a single PMT; such an event can then only satisfy the $n \geq 2$ coincidence requirement if a top-array PMT is also hit. These events are clearly expected from Poisson fluctuations, do not appear to coincide with non-Gaussian events, and are not cut (the cut band is truncated by 5 keVr). The cut band (red line) drawn here is identical to the cut band as drawn in Fig. 4.31, and the acceptance of this cut as predicted by the Monte Carlo is indicated along the bottom of the plot (in 5 keVr increments through the WIMP Search energy window).




Figure 4.33: Comparison of the ^{137}Cs calibration data (of Fig. 4.31, black line, with errors) and the Monte Carlo simulation (of Fig. 4.32, blue line) in the WIMP Search energy window 4.5 – 26.9 keVr. Events with only a single hit in the bottom PMT array are not included. The agreement is very reasonable considering the simplifying input assumptions (see Sec. 4.3.5). The excess of counts in the data for $G_1 > 2$ are indicative of the reverse-field region (*RFR*) pathology.




Figure 4.34: The *S1* hit patterns for event numbers 5 (left) and 8 (right) illustrate the localization check made for the *RFR* cuts G_2 , G_3 and G_4 . The event numbering was first shown in Fig. 4.26. Event number 5 shows high localization of the *S1* in the bottom PMT array and is cut; event number 8 does not show significant localization and passes all the *RFR* cuts.




Figure 4.35: **(left)** The index $G_i = S1_i/S1_{bot} \cdot \sqrt{S1}$ for $i = 2$ from the ^{137}Cs calibration data. The case $i = 1$ is shown in Fig. 4.31. The cut bands as used in the blind analysis are indicated by the red lines; 5 non-Gaussian events (blue $+$) are removed by this cut. Events with G_2 above the dashed red line are always cut; events with G_2 above the solid red line are only cut if they also show a high degree of localization. **(right)** Localization σ_{S1} is defined separately in x as $\sigma_{S1}(x) = \Sigma(S1_i \cdot x_i)/S1_{bot}$ (and similarly in y). Events are considered highly localized if $(\sigma_{S1}x)^2 + (\sigma_{S1}y)^2 < 30^2$. Events are also considered highly localized if σ_{S1} in either x or $y < 14$; this criteria gives heightened sensitive to multiple PMT hits along the edge of the detector.




Figure 4.36: **(left)** The index $G_i = S1_i / S1_{bot} \cdot \sqrt{S1}$ for $i = 3$ from the ^{137}Cs calibration data. **(right)** Localization σ_{S1} is defined separately in x as $\sigma_{S1}(x) = \Sigma(S1_i \cdot x_i) / S1_{bot}$ (and similarly in y). A detailed description is already provided for the case $i = 2$ (Fig. 4.35).




Figure 4.37: **(left)** The index $G_i = S1_i/S1_{bot} \cdot \sqrt{S1}$ for $i = 4$ from the ^{137}Cs calibration data. **(right)** Localization σ_{S1} is defined separately in x as $\sigma_{S1}(x) = \Sigma(S1_i \cdot x_i)/S1_{bot}$ (and similarly in y). A detailed description is already provided for the case $i = 2$ (Fig. 4.35).

S1 edge-to-all ratio cut (*E2A*)

The *S1 E2A* cut is similar in principle to the *S1 RFR* cut. In this case, an index $E_{frac} = S1_{edge}/S1$ was defined to reflect the extent to which an individual event might have scattered at the very edge of the detector, prior to scattering in the fiducial target. As described in Sec. 4.6.2, events at the very edge of the active target are expected to have reduced efficiency for drifting ionization to the liquid-gas interface. An example hit pattern of a non-Gaussian event electron recoil event that was cut by the *E2A* cut is shown in Fig. 4.38. Events that are cut by *E2A* are often (but not always) also cut by *RFR*. The performance of the *E2A* cut is shown in Fig. 4.40 for the ^{137}Cs calibration data.




Figure 4.38: The *S1* hit pattern for an event tagged by the QC2 *E2A* cut. The anomalous concentration of signal near the edge on the bottom array is highly unlikely for a “normal” single-scatter event. The event narrowly missed the threshold for the *RFR* cut.




Figure 4.39: **(top)** The index $E_{frac} = S1_{edge}/S1$ for the ^{137}Cs calibration data. The cut band as used in the blind analysis is indicated by a solid red line; 5 non-Gaussian events (red circles) are removed by this cut alone. Events excluded from the electron recoil population at more than 99.9% are indicated by red circles ($\circ$); the 11 red circle events between 4.5 – 26.9 keVr are also shown in Fig. 4.26 (in the same energy window). Severely non-Gaussian events with $\log_{10}(S2/S1)$ below the nuclear recoil -3σ contour are indicated by a blue x. **(bottom)** The index $E_{frac} = S1_{edge}/S1$ from a Monte Carlo simulation of $S1$ photo-electron signal generation. This is the same simulation already described in Sec. 4.3.5 and shown in Fig. 4.32 in the context of the RFR cut. The phe scale was converted to keVr by the same scaling applied to the data, and the $n \geq 2$ coincidence requirement was applied to the simulation data. Events with $E_{frac} = 1$ in the data were predicted by the Monte Carlo. These events are clearly expected from Poisson fluctuations, do not appear to coincide with non-Gaussian events, and are not cut (the cut band does not extend below 5 keVr). The cut band (red line) drawn here is identical to the cut band as drawn in Fig. 4.40, and the acceptance of this cut as predicted by the Monte Carlo is indicated along the top of the plot (in 5 keVr increments through the WIMP Search energy window).

S1 top fraction cut (*S1TF*)

The *S1TF* cut is also an *S1* hit-pattern cut, like the *RFR* and *E2A* cuts. In this case, an index $T_{frac} = S1_{top}/S1$ was defined to reflect the extent to which an individual event had an anomalous concentration of *S1* on the top PMT array. This could indicate an event which had scattered in the gas just below the top PMT array, prior to scattering in the fiducial target. Scatters in the ~ 1 cm Xe gas region between the top grid and the top PMT array gas are suppressed compared to scatters in the active target by the relative mass of Xe in the two regions. The densities of Xe at 177 K $\rho_{liquid} = 2.86 \text{ g cm}^{-3}$ and $\rho_{gas} = 0.02 \text{ g cm}^{-3}$ result in a factor of $\times 145$, and the gas volume is about $\frac{1}{15}$ compared to the active target. Scatters in the gas just below the top PMT array were therefore expected to be suppressed by about $\times 2000$. Note that in this region, the electric field is $E \approx 0$ because the grid and the PMT photocathodes are both biased at about 1 kV.

The *S1TF* cut was not expected to have a large effect. However, since its acceptance $\epsilon \approx 0.99$ (shown in Fig. 4.44) it was deemed a prudent inclusion. It did not cut any non-Gaussian events in either the ^{137}Cs calibration data or the WIMP Search data. It did cut one non-Gaussian event (that was not already cut by *RFR* or *E2A*) in the ^{137}Cs calibration data when considering the larger 8.4 kg fiducial target. The implementation of the *S1TF* cut is shown in Fig. 4.40 for the ^{137}Cs calibration data.




Figure 4.40: The index $T_{frac} = S1_{top}/S1$ for the ^{137}Cs calibration data. Events with $T_{frac} < 5 \times 10^{-3}$ were set to this value. The cut band as used in the blind analysis is indicated by a solid red line. The cut band is truncated at 5 keVr to allow for expected fluctuations. As discussed in the text, the cut was not expected to remove any non-Gaussian events (red circles). Events excluded from the electron recoil population at more than 99.9% are indicated by red circles ($\circ$); the 11 red circle events between 4.5 – 26.9 keVr are also shown in Fig. 4.26 (in the same energy window). Severely non-Gaussian events with $\log_{10}(S2/S1)$ below the nuclear recoil -3σ contour are indicated by a blue $\times$. The dashed line indicates a cut band that was considered, but not implemented. Event numbers 2 and 5 were already cut by *E2A*.




Figure 4.41: The index $T_{frac} = S1_{top}/S1$ from a Monte Carlo simulation of $S1$ photo-electron signal generation. Events with $T_{frac} < 5 \times 10^{-3}$ were set to this value. This is the same simulation already described in Sec. 4.3.5 and shown in Fig. 4.32 and Fig. 4.41. The phe scale was converted to keVr by the same scaling applied to the data, and the $n \geq 2$ coincidence requirement was applied to the simulation data. A small number of events with $T_{frac} = 1$, and a significant number with $T_{frac} < 5 \times 10^{-3}$ were predicted by the Monte Carlo. This is simply a reflection of the fact that the bottom PMT array receives about 80% of the $S1$ signal. Since these events are expected from Poisson fluctuations, the cut band does not extend below 5 keVr). The cut band (red line) drawn here is identical to the cut band as drawn in Fig. 4.40, and the acceptance of this cut as predicted by the Monte Carlo is indicated along the bottom of the plot (in 5 keVr increments through the WIMP Search energy window).

Pulse area fraction cut (PAF)

The pulse area fraction (PAF) cut was intended to assess the level of excess photon noise in an event. In general, there is very little. Figure 2.8 and Fig. 2.13 are typical in this regard. Another way to state this is that the total area of an event contained by the $S1$ and $S2$ pulses should be very close to one. This is expressed as $P_{frac} = (S1 + S2)/A$, where A is total area of the event record. Typical reasons for $P_{frac} < 1$ are $S2$ after-pulsing and single-electron $S2$ pulses, both of which are typified in Fig. 2.8. The motivation for this cut was that any departure from the expected level of photon noise could indicate a problem with the event. The PAF cut is actually a more aggressive version of the cut “ $S2$ fraction of trace area” already described in Sec. 4.6.1.

The behavior of P_{frac} for the ^{137}Cs calibration data is shown in Fig. 4.42. Almost no events lie below the cut band (red line). Event number 8 contained more than normal photon noise for $\sim 80 \mu\text{s}$ following the $S2$. Its value of P_{frac} is abnormal, but there is no other obvious problem with the event. A closer inspection reveals that 55% of the after-pulsing signal comes from three of the PMTs (numbers 25, 49 and 51). Higher than normal noise in these three PMTs had been noticed during routine data-quality examinations. A very slight, but systematic decrease in the quantum efficiency had also been noticed.

The cut band for the PAF cut was set so that the cut acceptance would be essentially flat with energy.




Figure 4.42: **(top)** The index $P_{frac} = (S1 + S2)/A$ for the ^{137}Cs calibration data. A is total area of the event record. Typical reasons for $P_{frac} < 1$ are $S2$ after-pulsing and single-electron $S2$ pulses, both of which are typified in Fig. 2.8. The cut band as used in the blind analysis is indicated by a solid red line. **(bottom)** The event record for number 8 shows nothing unusual except the $S2$ after-pulsing. 55% of the after-pulsing signal is concentrated in three of the PMTs (numbers 25, 49 and 51).

In retrospect this cut does not belong in the QC2 group, even with its more aggressive threshold. It should be retained, but as a general data quality cut at the level of QC0.

***S1* 10% Width Cut (*S1W*)**

A cut was implemented on the *S1* pulse width, where the width was defined between the points where the pulse rose (and fell) to 10% of its maximum height. The motivation for this cut was the observation of a small (~ 10) population of events with systematically low values of $\log_{10}(S2/S1)$, and anomalously long ($> 200 \mu\text{s}$) full widths at 10% max. An example from the ^{137}Cs calibration data is shown in Fig. 4.43. The event looks normal. The *S2* after-pulsing is slightly elevated, but well within the cut band defined in Fig. 4.42. The *S1* pulse consists of 12 phe. The cut hinges on the late arrival of the last photo-electron (Fig. 4.43, inset, yellow). Since the pulse shape is discreet, the probability to observe this photo-electron roughly 250 ns after the arrival of the first photo-electron is $\exp(-250/27) = 10^{-4}$. Figure 4.18 shows a more typical time profile for low-energy *S1* events.




Figure 4.43: An event in the ^{137}Cs calibration data with anomalous *S1* pulse width. The probability to observe the arrival of the last photo-electron roughly 250 ns after the arrival of the first photo-electron is $\exp(-250/27) = 10^{-4}$. A small (~ 10) sample of events with late-arriving photo-electrons were observed to also have low values of $\log_{10}(S2/S1)$.

The *S1W* cut required the *S1* full width at 10% max to be $< 150 \mu\text{s}$. The probability to observe a photo-electron more than 150 ns after the arrival of the first photo-electron is 0.39%. Ten events in 2000 with the last photo-electron arriving more than 200 μs after the first is about $\times 10$ more than expected. Nevertheless, a correlation with low values of $\log_{10}(S2/S1)$ is difficult to explain. Additionally, the non-Gaussian events removed by the cut were also removed by either *RFR* or *E2A*. This is shown in Fig. 4.27 (and a similar situation was observed in the non-blind WIMP Search data). In retrospect there is no convincing argument to retain this cut.

Partial Gamma-X cut (*PGX*)

The partial Gamma-X cut was conceived as a special case of the *RFR* Gamma-X cut described above and exemplified in Fig. 4.28. Simulations of the electric field near the cathode grid showed that regions from which no ionization would escape could extend up to 1 mm of the grid [72]. In the event that a γ scattered in such a region and then again in the fiducial volume, two possible outcomes have already been discussed: either two *S2* vertices would be observed, or the event would be a Gamma-X event, and possibly vetoed by the *RFR* cut. A third possibility is that a small fraction of the ionization would be extracted, with the rest lost due to the electric field orientation. Since the threshold for declaring a multiple scatter was set at 5% of the size of the primary *S2*, the cut would be sensitive to *S2* vertices of about $0.05 \cdot 4000 \simeq 200$ phe.

The signature of a *PGX* event would be a small secondary *S2* pulse at $dt \approx 80 \mu\text{s}$ after the *S1* (the maximum drift time). One non-Gaussian event in the ^{137}Cs calibration data was tagged by the *PGX* cut, as shown in Fig. 4.27. It had a secondary *S2* of ~ 65 phe ($2 - 3$ ionization electrons). Assuming a partial Gamma-X event could occur with equal probability anywhere within a 1 mm thick region of Xe immediately above the cathode grid, the rate of such events would be about 1/13 compared with Gamma-X events which scattered in the 1.3 cm bulk of the reverse-field region. This is an upper limit, since the shape of the field lines have not been considered. The one event cut by *PGX* in the ^{137}Cs calibration data was also cut by *E2A* (Fig. 4.27). No non-Gaussian events in the blind WIMP Search data exhibited this signature, though 1 non-Gaussian event in the non-blind WIMP Search did.

From a pragmatic point of view, the argument to retain this cut looks weak. This conclusion was reached even before unblinding the WIMP Search data. A primary reason for retaining it was the realization that the effectiveness of the *RFR* cut (discussed above) should diminish with distance from the bottom PMT array, and be least just at the cathode grid.

4.7 Acceptance of QC2 cuts for nuclear recoils

Because XENON10 is searching for elastic scatters of WIMP dark matter with Xe nuclei, it is the acceptance of each cut for nuclear recoils that is relevant. The QC2 cut acceptance as obtained from the AmBe calibration data is shown in Fig. 4.44. The acceptance was measured relative to QC0 and QC1, and only for events with $\Delta \log_{10}(S2/S1)$ values between the -3σ and $+3\sigma$ contours of the nuclear recoil centroid. For reference, the AmBe calibration data is shown in Fig. 4.44 (top) in the same shifted coordinates already described for Fig. 4.23. An assessment of the acceptance of QC0 is shown in Sec. 6.




Figure 4.44: **(top)** The acceptance of the suite of 6 QC2 cuts for nuclear recoils, relative to QC0 and QC1. Only nuclear recoil events with $\Delta \log_{10}(S2/S1)$ between the -3σ and $+3\sigma$ contours were counted. The $S1$ acceptance (black dash) is indicated as a reference (see Sec. 4.3.5). The WIMP Search energy window is indicated by the solid vertical lines. Cuts are indicated by abbreviated names in the legend, and described further in Sec. 4.6.5. The total average acceptance fraction of the QC2 cuts in the WIMP Search window is 0.860. **(bottom)** AmBe calibration data in the 5.4 kg fiducial volume, shown in shifted coordinates such that the γ calibration centroid is at zero. The nuclear recoil centroid and -3σ contours are drawn in piece-wise steps (red lines), while the electron recoil -0.995 (green) and -0.999 (blue) exclusion contours are drawn as continuous splines between bin centers. There are 4206 nuclear recoil events in the range $4.5 - 26.9$ keVr, counting only events with $\Delta \log_{10}(S2/S1)$ between the -3σ and $+3\sigma$ contours. The $+3\sigma$ contour is not drawn (in the range $4.5 - 26.9$ keVr the contour roughly follows the electron recoil centroid).

Chapter 5

DARK MATTER EXCLUSION LIMITS

“I have only one eye, I have a right to be blind sometimes... I really do not see the signal!”

– Horatio Nelson

5.1 XENON10 dark matter exclusion limits

After all cuts were frozen based on a study of the ^{137}Cs calibration data and the non-blind WS data, the final XENON10 WIMP-search results were obtained by simply “opening the box”. In other words, the 58.6 live days of data from WS3 and WS4 were looked at for the first time, with all the cuts pre-set according to the calibration data. The ROOT analysis chain was designated by the collaboration as the “primary” analysis for the purpose of reporting the XENON10 WIMP search results. However, both analyses (ROOT and Matlab) were run blind, so WIMP candidate events from both will be discussed. The results from the ROOT (Matlab) analysis are shown in Fig. 5.1 (Fig. 5.2), with 10 (17) events in the WIMP-search box.




Figure 5.1: The XENON10 58.6 live day (blind) WIMP search result from the primary (ROOT) analysis. There are 1805 single-scatter events in the energy range 4.5 – 26.9 keVr. Ten WIMP candidates remained in the acceptance box. Of these, five events are consistent with the Gaussian electron recoil background (indicated by blue numbers); four events are excluded from the Gaussian background population at $> 99.9\%$ (indicated by red numbers). The event just below 5 keVr should not have been accepted, because the $S1\ n \geq 2$ coincidence requirement for this event was met only because of a coherent noise pickup. All the events are labeled with a subscript R to distinguish them from the event numbering in my analysis.

5.1.1 The WIMP acceptance box and $\Delta\log_{10}(S2/S1)$ coordinates

The WIMP search results and the ^{137}Cs and AmBe calibration data are shown in terms of the discrimination parameter $y' = \Delta\log_{10}(S2/S1)$, rather than $y = \log_{10}(S2/S1)$. The transformed coordinates are particularly




Figure 5.2: The XENON10 58.6 live day (blind) WIMP search result from the secondary (Matlab) analysis. There are 1716 single-scatter events in the energy range 4.5 – 26.9 keVr. Seventeen WIMP candidates remained found in the acceptance box. Of these, twelve events are fully consistent with the Gaussian electron recoil background (blue circles). Event numbers 1, 2 and 5 with $\Delta \log_{10}(S2/S1) < -0.6$ (red circles) are clearly inconsistent with the electron recoil background. Event numbers 3 and 4 lie very close to the 99.9% contour (blue line), and could be reasonably expected fluctuations of the electron recoil background population.

convenient because they give an expectation value of $y' = 0$ for background events. The transformation was accomplished by subtracting the electron recoil band centroid, as determined by the ^{137}Cs calibration, from each data set.

The WIMP search acceptance box was drawn by making a Gaussian fit to the transformed discrimination parameter $y' = \Delta \log_{10}(S2/S1)$ for the single-scatter AmBe calibration data, in discrete bins. The bin edges were set at 2, 3, 4, 5, 6, 8, 10 and 12 keVee, using 2.2 $S1$ phe/keVee as obtained from the ^{137}Cs full-energy peak. This encloses the nuclear recoil equivalent energy range 4.5 – 26.9 keVr. The fits to $y = \log_{10}(S2/S1)$ are shown in Fig. 2.18 for the untransformed coordinates. The Gaussian fits for the transformed coordinates are not significantly different aside from a shift in the centroid. The ^{137}Cs calibration data centroid values in each bin are shown in Table 5.1 for both y and y' . The actual values of y' in each bin are slightly different from the expectation of $y' \equiv 0$. This is an expected consequence of the finite bin width, since all data within a given bin were transformed according to the bin centroid. The ER-subtracted centroids of the WIMP Search and AmBe calibration data are also shown. The latter define the top of the 50% acceptance box for the WIMP Search.

The width (σ) of the y and y' bands in each bin are compared in Fig. 5.3 The ^{137}Cs calibration data with

the $S2(x, y)$ correction applied (Table 5.1, column 2) is used as a standard reference (shown in blue). The band width is not significantly changed by the coordinate transformation. For reference the width before the $S2(x, y)$ correction, and the width of the WIMP Search data are also shown.

Table 5.1: The centroid values of the bands $y = \log_{10}(S2/S1)$ and $y' = \Delta \log_{10}(S2/S1)$ for each bin, as determined by a Gaussian fit. The values shown in columns 2 and 3 correspond to the centroids shown in Fig. 2.18 from the fits shown in Fig. 2.16 and Fig. 2.17. The values in column 4 are slightly different from zero, as expected (discussed in the text). The values in column 5 correspond to the top of the 50% acceptance WIMP Search box in Fig. 5.4. The typical error on the mean is ± 0.006 . The energy scaling is given by $S1/0.98 \text{ phe} = 1 \text{ keVr}$.

bin center ($S1 \text{ phe}$)	^{137}Cs y	AmBe y	^{137}Cs y'	AmBe y'	WIMP Search y'
3.0	2.885	2.462	+0.031	-0.390	+0.006
5.0	2.744	2.334	-0.019	-0.423	-0.004
7.0	2.661	2.283	-0.003	-0.397	-0.027
9.0	2.596	2.229	-0.013	-0.385	-0.007
11.5	2.549	2.189	+0.011	-0.362	-0.008
14.5	2.495	2.151	+0.000	-0.342	+0.002
17.5	2.464	2.130	+0.005	-0.334	+0.013
20.5	2.443	2.098	+0.009	-0.334	-0.002
23.5	2.426	2.080	+0.005	-0.338	-0.013
27.5	2.398	2.066	-0.004	-0.335	-0.005
32.5	2.387	2.028	+0.001	-0.364	-0.011
37.5	2.369	1.989	-0.004	-0.389	-0.016
42.5	2.357	1.966	-0.007	-0.395	-0.000
47.5	2.357	1.938	+0.002	-0.411	+0.006

The WIMP search acceptance box was defined between the -3σ and centroid of the AmBe calibration data, giving 50% acceptance for nuclear recoils. Discrete bins were used to avoid any systematic effects or artifacts from the fitting procedure. The energy window for the WIMP-search was also chosen *a priori*. The $S1$ signal was used to set the energy scale. On the low side, 4.4 phe (4.5 keVr assuming $\mathcal{L}_{eff} = 0.19$) was chosen because the efficiency for finding $S1$ pulses is > 0.997 for $S1 > 4.4 \text{ phe}$ ($S1$ acceptance is discussed in Sec. 4.3.5). The high side of the energy window was decided based on the fact that the predicted sensitivity of Xe to elastic scattering of WIMPs has already fallen by $\times 10$ by 30 keVr (as shown in Fig. 1.14). A second motivation was that the ^{137}Cs data shows increased leakage into the WIMP-search acceptance box above about 30 keVr. This effect is discussed in Sec. 2.7; in particular, see Fig. 4.23 and Fig. 4.27.




Figure 5.3: The width (σ) of the calibration bands from Gaussian fits in each $S1$ phe bin. The bin centers and band centroids are given in Table 5.1. The bin centers are offset slightly on each plot for clarity. The ^{137}Cs calibration data with $S2$ (x, y) correction applied (blue points in each plot) were already shown in Fig. 2.18 and Fig. 2.16, and are used as a standard for each comparison. The data were fit with a 5th-order polynomial (solid lines), excluding the first point. **(top)** Comparison with the ^{137}Cs calibration data after the coordinate transformation $y \rightarrow y'$ (violet) shows no significant difference to the width. **(middle)** Comparison with the transformed WIMP Search data (black) also shows no significant difference. **(bottom)** Comparison with the ^{137}Cs calibration data without the $S2$ (x, y) correction shows a systematic difference in the width. This leads to the difference in Gaussian tails in Fig. 5.2 as compared with Fig. 5.4.




Figure 5.4: XENON10 58.6 live day (blind) WIMP search result, with the post-unblinding addition of the $S2(x, y)$ correction. This very slightly modifies the event values of $y' = \Delta \log_{10}(S2/S1)$, and narrows the width of y' in each bin. In comparing Fig. 5.2, it can be seen that only six Gaussian tail events now remain in the WIMP search box (expected). Also, the -3 NR contour is closer to the NR centroid (expected). **Also shown** for each discrete energy bin are the observed (in black) and predicted (in blue, with uncertainties) electron recoil background leakage into the WIMP acceptance box (see also Table 5.2). The predicted leakage is in good agreement with the observed leakage, considering all but event numbers 1 and 2. **(at left)** A histogram of the distribution of y' for events in the range 4.5 – 26.9 keVr before QC2 cuts (light blue) and after all cuts (black). A Gaussian fit to the distribution (red) has $\chi^2 = 43$ and $d.o.f. = 50$. Event numbers 1 – 4 are labeled along the axis. The 5^{th} event (from Fig. 5.2) is excluded from the nuclear recoil population at more than 3σ and is not shown.

5.1.2 A comment on expected Gaussian tail events in the box

In comparing Fig. 5.1 and Fig. 5.2, there is clearly an excess of Gaussian tail events (12 versus 6) in the electron recoil background distribution. These events are labeled with blue numbers in Fig. 5.1, and with blue circles in Fig. 5.2. The reason for the discrepancy is clear from Fig. 5.3, (bottom): the $S2(x, y)$ correction discussed in Sec. 4.5.2 was not applied to the Matlab analysis prior to unblinding.

Table 5.2: The predicted number of leakage events in the WIMP Search data, obtained from the parameters of a Gaussian fit to the electron recoil population (^{137}Cs calibration data) in each bin. The prediction is shown with and without the $S2(x, y)$ correction.

bin center ($S1$ phe)	predicted number of leakage events	
	without $S2(x, y)$ correction	with $S2(x, y)$ correction
5.5	$0.6^{+0.3}_{-0.2}$	$0.9^{+0.4}_{-0.3}$
7.7	$0.5^{+0.3}_{-0.2}$	$0.9^{+0.4}_{-0.3}$
9.9	$0.9^{+0.4}_{-0.3}$	$1.1^{+0.5}_{-0.4}$
12.1	$0.6^{+0.3}_{-0.2}$	$0.9^{+0.4}_{-0.4}$
15.4	$1.0^{+0.4}_{-0.3}$	$2.1^{+0.7}_{-0.6}$
19.8	$1.8^{+0.6}_{-0.5}$	$3.0^{+0.9}_{-0.8}$
24.2	$2.1^{+0.7}_{-0.6}$	$3.4^{+1.0}_{-0.9}$
Total	$7.4^{+1.8}_{-1.5}$	$12.1^{+1.9}_{-1.6}$

In retrospect this seems like a substantial oversight, but prior to unblinding the primary focus was on eliminating non-Gaussian pathologies. Some leakage of events from the Gaussian tail of the electron recoil background distribution into the acceptance box was expected, and is shown in Fig. 5.4. The XENON10 collaboration considered employing a statistical background subtraction in the final analysis. This idea was ultimately rejected by the collaboration because the available statistics in the γ calibration data were only $\times 1.3$ greater than the statistics of the WIMP search data. The argument for statistical subtraction would have been stronger if the number of events in the calibration data had been a factor $\times 10$ larger.

5.1.3 The importance of blind calibration data

Based on the ^{137}Cs calibration data, about six (Fig. 4.23) rather than twelve (Fig. 5.2) events consistent with the Gaussian electron recoil background population were expected in the WIMP search data (blue circles in Fig. 4.23). The reason for this discrepancy is also clear: the ^{137}Cs calibration data was a “training set” for the cuts targeted at non-Gaussian pathologies. In deciding cut parameters and thresholds, the line between a Gaussian and a non-Gaussian event can be blurry near the 99.9% contour (blue curve in Fig. 5.2 and Fig. 5.2). Several of the cuts were clearly tuned to give the best possible performance on the “training set”. This is natural and inevitable. The remedy for such bias is to test the cuts on a second, blind calibration data set. XENON10 did not have time to acquire a second, high-statistics γ calibration data set in advance of the unblinding. As mentioned, a portion (~ 16 live days) of the WIMP Search data was fully examined prior to the (main) unblinding the remaining 58.6 live days. This was sufficient to determine that there were no data quality issues with the blind WIMP search data, but it was not sufficient to elucidate the importance of the $S2(x, y)$ correction for minimizing Gaussian leakage into the WIMP search acceptance box.

The XENON10 WIMP search data is shown in Fig. 5.4 with the $S2(x, y)$ correction applied after unblinding. All other parameters are identical to those used in the blind analysis, except for the Gaussian fits to the electron recoil and nuclear recoil distributions (since these depend on the $S2(x, y)$ correction). This leads to slight changes in the values of $\Delta \log_{10}(S2/S1)$ for each event, as well as in the contours of the WIMP-search acceptance box. The main impact is on the Gaussian population (blue circles), which has halved in number (from twelve to six). The five non-Gaussian events (red circles) from the blind analysis are all still evident on the plot, however event number 5 is no longer inside the acceptance box because the width of the $\log_{10}(S2/S1)$ band has decreased. The predicted number of electron recoil background events which leak into the WIMP acceptance box can be calculated from the parameters of the Gaussian fit in each discrete energy bin. This is compared with the observed number of event in Fig. 5.4. The total prediction of $7.4^{+1.8}_{-1.5}$ events in the energy range $4.5 - 26.9$ keVr agrees well with the observation of 8 ± 1 events which are clearly consistent with the electron recoil population. The events which are outside the 99.9% contour of the electron recoil distribution are discussed in Sec. 5.2. Note that there is no difference in the WIMP-nucleon exclusion limits calculated from the data as shown in Fig. 5.2 versus Fig 5.4, as discussed in Sec. 5.1.4.

5.1.4 Calculation of WIMP-nucleon exclusion limits

The spin-independent (scalar coupling) WIMP-nucleon exclusion limits were obtained from the data shown in Fig. 5.1. All ten events were considered as WIMP candidates, and the 90% C.L. upper limit was calculated using the “maximum gap” method [169]. This method results in a conservative upper limit for the WIMP-nucleon cross section, and makes no assumptions about the expected shape of either the signal or the background distributions. The standard galactic halo assumptions for the distribution of dark matter are discussed in Sec. 1.6.2. The WIMP-nucleon cross section versus WIMP mass exclusion limits are shown in Fig. 5.5. For a WIMP mass of 100 GeV c^{-2} , the excluded cross section upper limit is $\sigma = 8.8 \times 10^{-44} \text{ cm}^2$.

The WIMP-nucleon exclusion limits were calculated assuming the relative scintillation yield for nuclear recoils $\mathcal{L}_{eff} = 0.19$. This simplifying assumption is discussed in Sec. 2.6 and Sec. 6.2. The most likely form of $\mathcal{L}_{eff}$ based on the XENON10 neutron calibration is also discussed in Sec. 6.2, and results in a slightly weaker exclusion limit, passing through $\sigma = 1.1 \times 10^{-43} \text{ cm}^2$ at a WIMP mass of 100 GeV c^{-2} .

5.2 Discussion of the WIMP candidate events

Events consistent with the Gaussian electron recoil background population will not be discussed here. Although these events are technically “WIMP candidates” and were treated as such in both blind analyses (see also Sec. 5.1.4), they represent a natural and expected background population.

The ROOT analysis QC2 cuts were not discussed in Sec. 4.6.5, but are described in [10]. The cuts represent a simple way to statistically characterize anomalous hit patterns, based on the RMS value of the $S1$ signal. They were effective, but they do not specifically address the pathologies which lead to the non-Gaussian population.

5.2.1 Candidate events from the ROOT analysis

Five of the events remaining in the ROOT analysis WIMP Search acceptance box (Fig. 5.1, events $1_R, 2_R, 6_R, 8_R$ and 10_R) are excluded from the electron recoil distribution at $> 99.9\%$. This result generated the exclusion limits reported in [10]. Four of these five events were cut in the Matlab analysis. Each




Figure 5.5: Spin-independent WIMP-nucleon cross section upper limits at 90% C.L. versus WIMP mass. The upper curve (blue) is the previous most sensitive published exclusion limit, from the CDMS II experiment [7]. The lower curve (red) is from the XENON10 experiment [10]. The shaded areas indicate allowed WIMP parameter space calculated from the constrained minimal supersymmetric model (CMSSM). Figure reproduced from [10].

event is considered in turn.

- **Event number** 1_R only passed the $n \geq 2$ $S1$ coincidence requirement in the ROOT analysis because of coherent noise pickup. As discussed in Sec. 4.3.5, this requirement is essential in order to be confident that the $S1$ pulse can be identified with the $S2$ pulse. The coherent noise pickup on this event was ignored in the Matlab analysis because of the baseline-flattening (discussed in Sec. 4.2.2). The event was rejected by the Matlab analysis at the QC0 level.
- **Event number** 2_R was common to both analyses, and is discussed in Sec. 5.2.3.
- **Event number** 6_R was rejected by the Matlab analysis RFR cut. It was found to be indicative of a specific pathology referred to as a “resistor-chain” event. This is discussed in Sec. 5.2.2. The event waveform is shown in Fig. 5.6 with the $S1$ and y' values as determined by the Matlab analysis.
- **Event number** 8_R was also rejected by the Matlab analysis RFR cut, and was also found to be indicative of the “resistor-chain” pathology (see Sec. 5.2.2). The event waveform is shown in Fig. 5.6 with the $S1$ and y' values as determined by the Matlab analysis.

- **Event number** 10_R showed an anomalous $S1$ hit pattern concentrated on the edge PMTs and was tagged by the $E2A$ cut. It was within 2% of the cut threshold for the RFR cut. The event waveform is shown in Fig. 5.7. The event is shown in Fig. 5.6 with the $S1$ and y' values as determined by the Matlab analysis.




Figure 5.6: XENON10 58.6 live day (blind) WIMP search result. The data are identical to the analysis as presented in Fig. 5.4, with the addition of six non-Gaussian events that were tagged by QC2. These events are marked according to the cut ($E2A$ or RFR) which removed them. The events marked (in black) as 6, 8, 10 all remained in the ROOT analysis acceptance box after all cuts. The QC2 cuts are described in Sec. 4.6.3.




Figure 5.7: The waveform for **event number 10_R** (global event ID xev05_20061009T1947_0025238) showed an anomalous *S1* edge hit pattern (inset, top right) and was cut by *E2A*. It was almost cut by *RFR* but was just above the threshold.

5.2.2 Resistor chain events

The need for a reverse-field region between the cathode and the lower PMT array was already discussed in Sec. 4.6.4. This is not the only design requirement of a liquid Xe TPC which can cause anomalous events. In order to keep the drift field uniform across the active target, electric field shaping rings were located at regular intervals (discussed in Sec. 3.3.6). The high voltage resistors for the voltage dividing network (the “resistor chain”) were located in a small hollow region within the Teflon can, as shown in Fig. 5.8 (top). The hollow region is filled with liquid Xe, which provides additional dielectric insulation against electrical breakdown. It also introduces the possibility for a γ to scatter in the hollow region, and then again in the active target. This is shown schematically in Fig. 5.8.




Figure 5.8: **(top)** A photograph of the Teflon can that defines the active Xe target, in a clean room prior to installation in XENON10. The hollow region which houses the resistor voltage divider is in the lower left. A schematic example of a γ double scatter is shown, with one vertex in the hollow region and the other in the active region. **(bottom)** The $S1$ hit pattern from a representative event indicating the signature: the anomalous concentration of $S1$ near the hollow resistor chain region.

Shortly before unblinding the WIMP Search data, a small (~ 10 events in all the ^{137}Cs calibration data and the non-blind WIMP Search data, in the range $E_R < 40$ keVr) subset of RFR events was characterized. The defining characteristic (in addition to being cut by RFR) was an anomalous $S1$ hit pattern adjacent to the hollow region housing the resistor chain. The Teflon can did not define a perfectly light-tight volume, as was evident from the high rate of $S1$ -only events (from passive Xe outside the active volume). Although the hollow region surrounding the resistor-chain was not left open to the active Xe, it was also not perfectly light-tight. Figure 5.9 and Fig. 5.10 show the event waveforms and $S1$ hit patterns (top left) for event numbers 6_R and 8_R . The concentration of $S1$ signal near the hollow region is indicated by the arrow. The event waveforms show clean, single-scatter events (Fig. 5.10); only the anomalous $S1$ hit pattern reveals the nature of the event.




Figure 5.9: The waveform for **event number 6_R** (global event ID xev05_20070119T1654_0498858), showing a zoom on the S1 (inset) and the hit patterns for the S1 (at left) and S2 (at right). The arrival of a single photo-electron $\sim 250 \mu\text{s}$ after the first photo-electron is discussed in Sec. 4.6.5.




Figure 5.10: The waveform for event number 8_R (global event ID xev05_20061221T1648_1099695), showing a zoom on the $S1$ (inset) and the hit patterns for the $S1$ (at left) and $S2$ (at right).

5.2.3 Candidate events from the Matlab analysis

As mentioned in Sec. 5.1. the Matlab analysis was also run fully blind. The five events remaining in the WIMP Search acceptance box (Fig. 5.2) after all cuts are shown in Fig. 5.11 – Fig. 5.15. Each event is considered in turn, including how it was interpreted by the ROOT analysis. The location of all the events remaining in the WIMP Search box are shown in Fig. 5.16. The five events discussed here are indicated by red circles.

- **Event number 1** had an *RFR* index G_1 below the solid line in Fig. 4.31. As described in Sec. 4.6.5, the cut band for G_1 was truncated at about 6 keVr to avoid cutting events in which only 1 PMT on the bottom array was hit. These events correspond to the arc spanning from $(x, y) \approx (1, 1)$ to $(x, y) \approx (10, 3)$ in Fig. 4.31. The arc is slightly more visible in the simulation shown in Fig. 4.32. This event was cut by the ROOT analysis at the QC2 level because of its $S1$ hit pattern in the top PMT array. Similar $S1$ hit patterns can be found in the AmBe nuclear recoil data, but they account for $< 5\%$ of events with $E_R \sim 5$ keVr.
- **Event number 2** was almost cut by *RFR*. Its *RFR* index G_3 was above the solid line in Fig. 4.36 (right), flagging it as a potential Gamma-X event. However, it (just) failed the localization requirement (Fig. 4.36, left). It was found that a relaxation of the localization requirement – with a corresponding decrease of $< 1\%$ in the *RFR* cut acceptance – would have cut this event. The event occurred close to the bottom of the fiducial target, which is where Gamma X events due to the reverse-field region preferentially appear. This event was also found in the ROOT analysis acceptance box (event number 2_R).
- **Event number 3** contains two lessons for improving the efficacy of the basic (QC0) quality cuts. The $S1$ occurred at $t \sim 80 \mu s$, just before the $S2$. It was considered part of the $S2$ by the pulse-finding algorithm. The event was accepted because an anomalously large coherent noise pickup at $t \sim 20 \mu s$ faked an $S1$ pulse. No coherent noise of this size had been observed in the ^{137}Cs calibration data. In retrospect this pathology is easy to spot: the probability to observe a 5 phe $S1$ consisting of ~ 0.3 phe/PMT in 18 PMTs is vanishingly small. The second problem with this event is the mischaracterization of the actual $S1$ pulse. As discussed in Sec. 4.4, two very closely-spaced $S2$ pulses are often flagged as a single pulse. The double scatter can then be identified by its width, with no loss of acceptance for WIMP Search. In the present case, the width appeared normal for a single $S2$. The “pulse aspect ratio” cut described in Sec. 4.6.1 was designed specifically for this type of event. In all similar events observed in the ^{137}Cs calibration data, the pulse aspect ratio was found to be anomalous because the $S1$ pulse height was much larger (compare Fig. 4.22). A better (or perhaps additional) means of tagging this pathology would be to check the rise time of the $S2$ pulse. Because of the discrete arrival times of photo-electrons for low energy recoils, a pulse rise time check should be made after passing an appropriate filter over the pulse. The *double – sliding – box – filter* described in Sec. 4.4 is very apt for this. This event was cut by the ROOT analysis QC0.
- **Event number 4** was also almost cut by *RFR*. Its *RFR* index G_1 was just below the solid line in Fig. 4.31, which is not surprising considering the $S1$ hit pattern shown in Fig. 5.2.3. It was found that a tightening of the G_1 cut band – with a corresponding decrease of $< 1\%$ in the *RFR* cut acceptance – would have cut this event. This event also occurred close to the bottom of the fiducial target. The ROOT analysis QC2 cut this event.

- **Event number 5** passed all cuts in both analyses, however the (neural network) position reconstruction algorithm used by the ROOT analysis found this event outside the fiducial volume, at $R > 85$ mm. The χ^2 position reconstruction used by the Matlab analysis found $R = 78$ mm. A closer inspection of the $S2$ hit pattern on the top PMT array shows 8 PMTs with no signal (shown dark), and a further ~ 20 PMTs with about 1 phe/PMT. This is a clear example of why the maximum likelihood method should be used in place of the χ^2 for position reconstruction (discussed in Sec. 4.4.3). The maximum likelihood method was implemented in the Matlab analysis after unblinding the WIMP Search data. It found the radial position of event number 5 to be $R = 82$ mm.




Figure 5.11: The waveform for event number 1 (global event ID xev05_20070126T1727.0190943), showing a zoom on the $S1$ (inset) and the hit patterns for the $S1$ (at left) and $S2$ (at right).




Figure 5.12: The waveform for event number 2 (global event ID xev05_20070105T1420_0429076), showing a zoom on the S1 (inset) and the hit patterns for the S1 (at left) and S2 (at right).




Figure 5.13: The waveform for event number 3 (global event ID xev05_20070126T1727_0609929), showing a zoom on the S1 (inset) and the hit patterns for the S1 (at left) and S2 (at right).




Figure 5.14: The waveform for event number 4 (global event ID xev05_20070105T1420_0140244), showing a zoom on the $S1$ (inset) and the hit patterns for the $S1$ (at left) and $S2$ (at right).




Figure 5.15: The waveform for event number 5 (global event ID xev05_20061009T1947.0745600), showing a zoom on the $S1$ (inset) and the hit patterns for the $S1$ (at left) and $S2$ (at right).




Figure 5.16: The $r - z$ view of the events shown in Fig. 5.2. The 5.4 kg blind fiducial volume is indicated by solid blue lines. Event numbers 2 and 4 are clustered near the bottom of the fiducial volume, where Gamma-X events with a scatter in the reverse-field region preferentially appear. Event number 3 was not properly reconstructed, so its z coordinate is wrong (it should have been at $dt \sim 1 \mu s$). The radial position of event number 5 was reconstructed at $R > 80$ by the ROOT analysis, and also by the (post-unblinding) maximum likelihood reconstruction (Sec. 4.4.3).

5.2.4 Recommendations for future liquid Xe experiments

The events consistent with the Gaussian electron recoil background population were not discussed in Sec. 5.2.3. These events are technically “WIMP candidates” and were treated as such in the blind analysis (see also Sec. 5.1.4), but they represent a natural and expected background population. Next-generation liquid Xe dark matter detection experiments are expected to have three options for handling this background:

- (1) reduce the overall background rate so that Gaussian leakage into the acceptance box is sub-dominant.
- (2) Amass a sufficient quantity of calibration data (at least a factor $\times 10$ compared with WIMP Search) such that a statistical background subtraction can be employed.
- (3) Define an acceptance box such that the acceptance for nuclear recoils decreases commensurate with the decrease in discrimination power, to maintain zero background.

Taken together, the two analysis chains used in XENON10 offer several clear lessons for future liquid Xe dark matter detection experiments in regard to non-Gaussian events. The specific details are discussed in the previous section, in the context of each event which passed all cuts (but should not have). The lessons can be summarized in three parts, and stand in stark contrast to the protest of Horatio Nelson (at the beginning of this chapter).

- The need for robust basic quality cuts (QC0) – strictly requiring candidate events to conform to the template (Fig. 4.18) – cannot be overstated. Event number 3 (Fig. 5.13) and event number 1_R (Sec. 5.2.1) serve as poignant reminders that non-observance of a signal pathology in $\sim 10^7$ calibration events does not exclude its possible observance in a separate sample of $\sim 10^7$ WIMP Search events (this was the total number of digitized events during the 58.6 live day exposure, at all energies). Fast digitization, voltage amplification and multiple PMTs yield a wealth of information about each event record. With this information, it is possible to fully characterize the expected event template. There is no need for even a single mistaken event out of ~ 1800 , or even 18000. It should also be possible to maintain the acceptance of the event template (QC0) for genuine events at $\sim 99\%$.
- The success of the advanced cuts QC2 is crucial. In particular, the *RFR* and *E2A* cuts based on *S1* hit pattern were found to eliminate a significant fraction of non-Gaussian electron recoil background events. These cuts in their present form could have been pushed harder to increase the discrimination power, as evidenced by event number 2 (Fig. 5.12) and event number 4 (Fig. 5.14). Only a modest decrease in acceptance would have been incurred to remove these 2 events: $\sim 84\%$ instead of 86% . It would be prudent to push the cuts even harder, however. A QC2 cut acceptance of 80% would still lead to an acceptable kg-day exposure, with the effective removal of all non-Gaussian events. In concert with the lessons learned regarding the basic quality cuts QC0, liquid Xe detectors could achieve a total background discrimination of $> 99.9\%$ for non-Gaussian events in the range $5 < E_R < 30$ keVr. The Gaussian population would still follow Fig. 2.15, and points 1 – 3 (above) must be considered.
- The QC2 cuts *RFR* and *E2A* implemented in the Matlab analysis were extremely effective. As mentioned, they could have removed all the non-Gaussian events in the WIMP Search data with only a marginal decrease in acceptance. The other QC2 cuts discussed in Sec. 4.6.5, however well-motivated, were not found to be significant. The QC2 cuts implemented by the ROOT analysis were less effective than the *RFR/E2A* combination, and further, it is not clear that they show as much potential. Future work on characterizing anomalous *S1* hit patterns could build on the concept

described in Sec. 4.6.5 for the *RFR* cut. In particular, a maximum likelihood analysis of the *S1* hit pattern could be performed, considering the information obtained from the size of the *S2* signal and the event vertex (x, y, z) . This could provide the most thorough characterization of any anomalous concentration of signal, while allowing for normal fluctuations.

Chapter 6

NEUTRON CALIBRATION AND $\mathcal{L}_{eff}$

“Everything is complicated; if that were not so, life and poetry and everything else would be a bore.”

— *Wallace Stevens*

6.1 Neutron calibration with an Americium-Beryllium source

The XENON10 detector was exposed for 12 hours, with a live fraction 0.92, to a $3.7 \text{ MBq} \pm 15\%$ Americium-Beryllium (AmBe) source emitting 220 n/s . The neutron rate is based on a yield of $6 \times 10^{-5} \text{ n/Bq}$ [165]. The exposure occurred with the detector installed in the shield, in low-background conditions at Laboratori Nazionali del Gran Sasso. As mentioned in Sec. 3.1, the cosmic muon flux was reduced by about 10^6 compared with surface conditions [112] due to the 3100 meters water equivalent rock overburden. The XENON10 shield (20 cm Pb outside 20 cm high density polyethylene) reduced the cavern γ flux by more than 10^5 , and the cavern neutron flux by about $10^{2.5}$. The source (attached to a steel rod) was inserted through a 7 mm diameter hole in the shield and positioned next to the detector cryostat, sandwiched between an additional 5 cm Pb as shown schematically in Fig. 6.1.




Figure 6.1: Side view schematic of the XENON10 experiment (to scale), showing 20 cm Pb outside 20 cm polyethylene and a cut-away rendering of the geometry from the GEANT4 simulation. The AmBe source was inserted through a 7 mm hole in the shield and positioned next to the cryostat, behind an additional 5 cm of Pb. Ancillary systems and cabling are omitted for clarity. The simulation is described in detail in [108]. Detector cut-away view courtesy of [108].

Initial neutron energies range from below 0.1 MeV to 11 MeV, with a mean at 4.3 MeV [109]. The neutron energy spectrum is known with an accuracy of $\pm 3\%$ (per 0.1 MeV bin) for a source of this strength [152], and the resulting Monte Carlo nuclear recoil spectrum is insensitive to such variations. The active target of XENON10 has a 10 cm radius and 15 cm height, with a Xe mass of 13.7 kg. The analysis presented

here uses only nuclear recoils which occurred in a 5.4 kg fiducial volume, with 8 cm radius and 9.3 cm height. This is the same fiducial volume used for the blind analysis of the WIMP-search data, as reported in [10], surrounded by a minimum of 2 cm of self-shielding Xe. Despite the features in the AmBe source energy spectrum [109], the spectrum of neutron energies as they enter this 5.4 kg fiducial volume is a featureless exponential, falling 1 decade in 3.5 MeV.

The source also emits 148 γ /s at 4.4 MeV from the de-excitation of the ^{12}C final state, and a 0.06 MeV γ (branching ratio of 36%) which is not relevant due to the 5 cm of Pb. The single-scatter γ rate in the fiducial volume < 100 keV electron-recoil energy (keVee) is roughly flat with energy and is reduced $\times 40$ by the 5 cm of Pb. The Monte Carlo predicts an additional $\times 3$ reduction due to the 2 cm of self-shielding Xe. Prior to discrimination, the single scatter γ rate 0 – 100 keVee is < 2 cts/0.25 keVee over the full exposure. In contrast, the single scatter elastic nuclear recoil rate is higher by as much as $\times 400$, as shown in Fig. 6.2. Inelastic scatters on Xe provide additional prominent γ lines at 40 and 80 keVee, which are easily avoided by the $S2/S1$ discrimination; this is shown in Fig. 6.4.

Event data were recorded at a rate of 6.5 Hz during the 12 hour exposure. The $S2$ -sensitive trigger is described in Sec. 4.3.2. A single ionization electron extracted to the gas corresponds to an $S2$ signal of 24 ± 7 photo-electrons (phe) [144]. The trigger was verified with a known voltage input to have full efficiency for an $S2$ corresponding to 4 electrons, and about 75% acceptance for 3 electrons. Analysis was restricted to $S2$ pulses corresponding to at least 8 electrons, which is more than 3σ below the typical $S2$ pulse at $S1 = 3$ phe (shown in Fig. 6.4). The $S1$ signal was found by look-back after an event was digitized, with the signal acceptance determined by the size of single phe pulses from the PMTs, and the requirement of a minimum of 0.35 phe/PMT in 2 or more PMTs ($n \geq 2$ coincidence), in a 0.3 μs time window. The $S1$ peak-finding efficiency is predicted to be better than 0.985 at 1 keVee (3 $S1$ phe). This was described in Sec. 4.3.5, and is indicated by the dotted line in Fig. 6.5. A high energy veto was set on the $S1$ signal to avoid digitizing γ scatters with energies above ~ 150 keVee. At this level, the veto was $\times 1.5$ above the high-energy tail of the elastic nuclear recoil spectrum.

6.2 Determination of the relative scintillation yield L_{eff} for liquid Xe

As described in Sec. 2.9, the relative (to 122 keVee) scintillation yield of nuclear recoils has been measured by many groups. The analysis in this thesis has been presented in terms of a constant (flat) $\mathcal{L}_{eff} = 0.19$. The definition of $\mathcal{L}_{eff}$ was given in Eq. 2.7, and is repeated here in terms of nuclear recoil equivalent energy (in unit of keVr):

$$E_R = \frac{S1}{L_y \cdot \mathcal{L}_{eff}} \cdot \frac{S_e}{S_n} \quad (6.1)$$

where $\mathcal{L}_{eff}$ is the scintillation yield of Xe for nuclear recoils, relative to the zero-field scintillation yield for electron recoils at 122 keVee. The use of this standard reference energy avoids any systematic error from the non-linear scintillation yield of Xe for electron recoils. $S1$ for 122 keVee gamma rays, after position-dependent corrections, gives a volume-averaged light yield $L_y = 3.0 \pm 0.1(\text{syst}) \pm 0.1(\text{stat})$ phe/keVee at the drift field $E_d = 0.73$ kV/cm. Scintillation quenching by the external electric field was measured to be $S_e = 0.54 \pm 0.01$ and $S_n = 0.93 \pm 0.04$ for $E_d = 0.73$ kV/cm [13].




Figure 6.2: The spectrum of single scatter nuclear recoils from 11 hr live exposure to an AmBe neutron source (black line, with 1σ errors), and the spectrum from a detailed Monte Carlo of the experiment, obtained from the best-fit $\mathcal{L}_{eff}$ curve shown in Fig. 6.3 (red line) according to Eq. 6.1. The $S1$ peak-finding efficiency for the bin centered on 1.125 keVee (3.4 phe) is > 0.997 . Bins < 1 keVee were not used to obtain the fit, and are shown in grey. The translation from keVr to keVee is shown for the 7 $\mathcal{L}_{eff}$ spline points in Fig. 6.3 (red circles along axis). Also shown are the Monte Carlo spectra without the effect of the $S1$ peak-finding efficiency (light dash), and for a constant $\mathcal{L}_{eff} = 0.19$ (solid blue line).

The Monte Carlo recoil energy spectrum was converted from keVr to keVee, via Eq. 6.1, and convolved with the measured $S1$ energy resolution. $\mathcal{L}_{eff}$ was then varied to give the best agreement between the data and Monte Carlo spectra, using the binned maximum likelihood method [166]. The form of $\mathcal{L}_{eff}$ was modeled by a piece-wise cubic spline, which provides optimal freedom with continuous 1^{st} and 2^{nd} derivatives. Spline points were fixed at 2, 5, 10, 15, 25, 50 and 100 keVr and unconstrained in $\mathcal{L}_{eff}$. Spline endpoints (also unconstrained in $\mathcal{L}_{eff}$) at 0.8 and 500 keVr are not shown; $d\mathcal{L}_{eff}/d(keVr) = 0$ at the endpoints. Using fewer spline points was found to over-constrain the solution. Conversely, using more spline points merely added fluctuations (within the statistical error bars) to the $\mathcal{L}_{eff}$ curve shown in Fig. 6.3. The best fit $\mathcal{L}_{eff}$ shifts by $< \pm 0.005$ if the specific keVr values of the spline points are translated (in either direction), and this is included in the statistical uncertainty (see Table 6.1).

The single-scatter nuclear recoil spectrum is shown in Fig. 6.2 (black line, with 1σ error bars), along with the best-fit Monte Carlo spectrum (solid red line). The fit has $\chi^2 = 403$ with $d.o.f. = 385$, with a p-value of 0.254. For comparison, the spectrum assuming a constant $\mathcal{L}_{eff} = 0.19$ is also shown (blue line, $\chi^2 = 549$, $p < 10^{-6}$). To avoid additional systematic uncertainty, bins in which the $S1$ peak-finding efficiency (for the bin center) is predicted to be < 0.997 were not included in the maximum likelihood fit, and are indicated in grey (the Monte Carlo spectrum for these bins is shown as a red dashed line).




Figure 6.3: The $\mathcal{L}_{eff}$ curve (at $E_d = 0.73$ kV/cm) corresponding to the best fit shown in Fig. 6.2 (red line), and 6 spline points (red circles) with 1σ error bars. The errors are statistical only, and are summarized along with systematic effects in Table 6.1. Also shown are data from [13] (triangles), [43] (squares) and the theoretical prediction of [83] (dash-dot). The $\mathcal{L}_{eff}$ solution is dashed below ~ 3 keVr (1 keVee), to indicate the energy regime not used to obtain the best fit. A constant $\mathcal{L}_{eff} = 0.19$ is also shown (solid blue line).

The $\mathcal{L}_{eff}$ curve corresponding to the best-fit is shown in Fig. 6.3 (solid red line) as a function of recoil energy, along with the spline points (filled red circles). Also shown are data from [13] (triangles), [43] (squares) and the theoretical prediction of [83] (dash-dot line). The keVee equivalent value of each spline point is indicated in Fig. 6.2 (red circles, along the axis). The 1σ (stat) error bars on the best-fit $\mathcal{L}_{eff}$ spline points were estimated by a Monte Carlo method that is standard for multi-parameter fits: the bin counts of the best-fit recoil spectrum (solid red line in Fig. 6.2) were allowed to Poisson-fluctuate. The resulting spectrum was then treated as input data, and fit with the same method described above. About 10^4 such “experiments” were performed, and the resulting $\mathcal{L}_{eff}$ values at each of the spline points were found to be normally distributed.

The overall event rate normalization between the Monte Carlo and data was treated as a free parameter in the $\mathcal{L}_{eff}$ minimization. Allowing for the DAQ dead time of 8% and the event acceptance described in Sec. 6.2.1, the Monte Carlo predicts an absolute single scatter neutron rate consistent with the uncertainty on the source strength.

6.2.1 Event acceptance

As described in Sec. 4.1.1 and Sec. 4.6.1, events had a record length of $156 \mu\text{s}$ with a 105 MHz ADC, and were required to consist of an $S1$ pulse followed by an $S2$ pulse. The pulses are easily distinguishable by their




Figure 6.4: The discrimination parameter $\log_{10}(S2/S1)$ for nuclear recoils in Xe. The parameter is Gaussian except for a small ($< 2\%$) secondary population below the -3σ contour. The analysis thresholds for light ($S1$) and charge ($S2$) are indicated by dashed lines. The centroid of the distribution for electron recoils (red line) is expected to lie slightly above the center of the 40 and 80 keVee inelastic populations [47]. γ events with $E > 80$ keVee are not shown.

width and shape: $S1$ decays with $\tau = 27$ ns [54], and $S2$ is roughly Gaussian with FWHM $\sim 0.8 \mu\text{s}$. No $S1$ pulse was distinguishable for 7% of events inside the central 8 cm radius (the z -coordinate is indeterminate without the $S1$). A further 4% of events had an $S1$ candidate pulse with all its signal concentrated in a single PMT (fails the $n \geq 2$ coincidence requirement). This is expected due to Poisson fluctuations of the $S1$ signal associated with low-energy nuclear recoils: the number of events considered as having no $S1$ is predicted to better than 15% by the deficit of events between the best-fit Monte Carlo, and the same spectrum without the $S1$ peak-finding efficiency (the red-dashed lines below 1 keVee in Fig. 6.2).

Additional quality cuts placed loose constraints on acceptable $S1$ and $S2$ pulse widths, the distribution of scintillation light between top and bottom PMT arrays, and required the first and last 50 samples of each event to be featureless. These cuts, collectively referred to as QC0, have a combined acceptance $> 99\%$ for genuine scatters in the fiducial volume (shown in Fig. 6.5, discussed in Sec. 4.6.1 and reported in [10]). A $\pm 3\sigma$ cut on the discrimination parameter $y = \log_{10}(S2/S1)$ was applied to select the elastic scatters, as indicated in Fig. 6.4. This excludes the small ($< 2\%$) non-Gaussian population of events below the -3σ band in Fig. 6.4.

Two other quality cuts were applied to the AmBe data. The first cut (“excess Photon Noise”) required that the proportion of signal concentrated in the primary $S2$ pulse be > 0.85 ; pulses failing this cut tend to have an excess of single-electron $S2$ events (or other spurious photon noise), which could complicate the determination of pulse parameters. This cut was not applied to the WIMP Search data; conversely, none




Figure 6.5: Event acceptance for single-scatter nuclear recoils in the fiducial volume. Average acceptance 0.5 – 100 keVee is indicated in the legend. The cuts are described in Sec. 6.2.1. The $S1$ peak-finding efficiency is discussed in Sec. 6.2.2

of the advanced QC2 cuts were applied to this analysis of the AmBe data. The second cut (“single $S1$ ”) required that only 1 possible $S1$ -like pulse be found before the $S2$ pulse; a secondary $S1$ candidate was required to be at least $\times 0.25$ as many phe as the primary. The rate of $S1$ -only pulses in the absence of the AmBe source is > 100 Hz, and the spectrum falls exponentially. This leads to the mild increase in acceptance of the “single $S1$ ” cut with energy, as shown in Fig. 6.5. Below about 3 keVee the acceptance rises again due to the $\times 0.25$ phe threshold for declaring multiple $S1$ candidates.

The “single $S1$ ” cut was already discussed in Sec. 4.6.1 as a QC0 cut. It is presented separately here only because its acceptance for the AmBe data is less than 99%, due to the higher data rate. The acceptance of the “single $S1$ ” cut for the WIMP Search data (included in the definition of QC0) leads to a QC0 cut acceptance $> 99\%$ across the energy range of interest ($5 < E_R < 25$ keVr). Each of these cuts had an acceptance > 0.97 for elastic nuclear recoils in the fiducial volume. The cuts are not completely orthogonal, and the combined acceptance for nuclear recoils (including QC0) is 0.930. The best-fit $\mathcal{L}_{eff}$ solution is not significantly changed if any (or all) of these cuts are not used, however the quality of fit is worse by 10-20%.

Single-electron $S2$ pulses appear in about 50% of elastic nuclear recoil events, and about 4% of events exhibit a small secondary $S2$ pulse of > 2 electrons but below the analysis threshold of 8 electrons (200 phe). A single scatter nuclear recoil of this size would fall below the detector $S1$ threshold (see Fig. 6.4), therefore $S2$ pulses of this size were not treated as multiple scatters. This is accounted for by the $S1$ peak-finding efficiency.

6.2.2 Signal identification and $S1$ acceptance

A 0.75 keVee elastic nuclear recoil event with $S2 = 899$ phe and $S1 = 2.3$ phe (about 1 phe per PMT in two of the bottom-array PMTs) is shown in Fig. 6.6. Baseline fluctuations of up to 0.10 ± 0.03 phe have been set to zero; the threshold is set individually for each PMT based on RMS noise. Coherent noise pick-up is rejected because photo-electron pulses are checked on individual PMTs. The pulse at $100 \mu\text{s}$ is a typical single-electron $S2$ (22 phe spread out over $\sim 0.8 \mu\text{s}$).




Figure 6.6: An elastic nuclear recoil event at 0.75 keVee. The $S1$ signal of 2.3 phe at $30 \mu\text{s}$ is shown (inset) $\times 10$ in y and $\times 100$ in x . The small pulse at $100 \mu\text{s}$ is a 22 phe single-electron $S2$. Baseline fluctuations of up to 0.10 ± 0.03 phe have been set to zero.

The $S1$ peak-finding efficiency (Fig. 6.5, dotted) was obtained by a detailed simulation of the phe counting statistics, including PMT response and DAQ sampling effects. It predicts a flat 100% efficiency above 2 keVee. The low-energy Compton-scatter background observed in the γ calibration data is found to be flat in the energy range 2 – 20 keVee. Under the assumption that this background should remain flat in the range 0.2 – 2 keVee, the simulation prediction and the observed roll-off in the data are consistent at the 1σ level.

The $\mathcal{L}_{eff}$ solution presented in Fig. 6.3 uses the simulation prediction as shown in Fig. 6.5 to model the $S1$ peak-finding efficiency. In order to estimate the possible systematic effect of the $S1$ peak-finding efficiency on the $\mathcal{L}_{eff}$ result, an $S1$ efficiency curve shifted higher (lower) by 20% was also considered. The effect on the $\mathcal{L}_{eff}$ solution at 2 keVr was a systematic decrease (increase) of 0.012. At all other spline points, the shift was at most half this amount. This is summarized in Table 6.1. We also find that the best fit $\mathcal{L}_{eff}$ points shift by < 0.005 if the analysis threshold of 1 keVee is lowered to 0.5 keVee.

Table 6.1: Summary (for each spline point) of $\mathcal{L}_{eff}$ best fit at $E_d = 0.73$ kV/cm, statistical error and systematic error arising from uncertainty in $S1$ acceptance and Xe(n,n)Xe elastic scattering cross-section data (σ_{el}), if using ENDF/B-VII database instead of the default JEFF-3.1 for Xe(n,n)Xe scattering. All errors are quoted in absolute terms.

keVr	$\mathcal{L}_{eff}$	<i>Uncertainty</i>		
		statistical	$S1$ acceptance	σ_{el} data
2	0.165	± 0.014	± 0.012	-0.015
5	0.161	± 0.011	± 0.006	-0.008
10	0.167	± 0.012	± 0.005	-0.007
15	0.202	± 0.011	± 0.006	-0.010
25	0.222	± 0.012	± 0.005	-0.008
50	0.221	± 0.009	± 0.004	-0.006
100	0.234	± 0.010	± 0.004	-0.005

6.2.3 Other systematics

To assess the uncertainty due to the nuclear cross section databases for Xe(n,n)Xe scattering, Monte Carlo recoil spectra were obtained separately using the ENDF/B-VII and the JEFF-3.1 as input to the GEANT4.9.0 simulation [108]. The spectra are consistent to better than 3% per bin for recoil energies above 2 keVr, but the ENDF/B-VII database yields 6% less counts than the default JEFF-3.1 in the energy range 1 – 2 keVr. The effect of 3% fluctuations per bin on the resulting best-fit $\mathcal{L}_{eff}$ is less than 1/2 the statistical errors quoted in Table 6.1. The increase in counts 1 – 2 keVr is found to cause a systematic decrease in the best fit $\mathcal{L}_{eff}$ spline points of at most $\mathcal{L}_{eff} = 0.015$ at 2 keVr and slightly less at the other energies; this is summarized in Table 6.1.

The $S1$ resolution for nuclear recoils is dominated by binomial fluctuations in the collection of scintillation photons. The intrinsic single photo-electron resolution of the PMTs (average $\sigma/\mu = 0.58 \pm 0.05$) leads to an $S1$ resolution that is about 16% worse (at all energies) than that expected solely from the binomial fluctuations. This result is directly verified from the width of the $\log_{10}(S2/S1)$ band [47]. The light collection efficiency for $S1$ is found to be 8%. Considering the uncertainty in the PMT single phe response, and taking a conservative estimate of $\pm 20\%$ uncertainty in the $S1$ light collection, the uncertainty in the $S1$ resolution is calculated to be $< \pm 2\%$. If the resolution were better (worse) by this amount, the $\mathcal{L}_{eff}$ curve in Fig. 6.3 would translate higher (lower) by less than 1/4 the statistical error (this is not included in Table 6.1). The quality of the fit is not significantly changed.

6.2.4 Behavior of $\mathcal{L}_{eff}$ below 10 keVr

In Fig. 6.7 we show the shape of the simulated single scatter nuclear recoil spectrum if $\mathcal{L}_{eff}$ were to continue to drop below 10 keVr. This was done by taking the best-fit solution shown in Fig. 6.3, and forcing the spline point at 2 keVr to have $\mathcal{L}_{eff} = 0.10$. The prediction of about 22% fewer events per bin for all bins 1 – 3 keVr is more than 5σ inconsistent with the data ($\chi^2 = 769$).

The expectation that $\mathcal{L}_{eff}$ should continue to fall with energy is due to the Lindhard model [102] for the partitioning of recoil energy into atomic motion and electronic excitation (so-called nuclear quenching). In Ge the agreement between experimental ionization quenching and the Lindhard prediction is very good, above 10 keVr [34]. Below 10 keVr measurements give values as much as $\times 2$ above the Lindhard prediction




Figure 6.7: The best fit to the data (see Fig. 6.2) is compared with the quality of the fit if $\mathcal{L}_{eff}$ is forced to continue falling below 10 keVr. (inset) The best-fit $\mathcal{L}_{eff}$ curve (see Fig. 6.3), and the same curve constrained to fall to $\mathcal{L}_{eff} = 0.10$ at 2 keVr (blue line).

[113]. However, the Lindhard model applies to the total electronic excitation. Scintillation light in Xe arises from the portion of electronic excitation which recombines.

To account for this difference, Hitachi proposed a model of biexcitonic electronic quenching [83]. His prediction for the total quenching of Xe scintillation is lower than the Lindhard prediction (when applied directly to scintillation) by about 1/3, and is shown in Fig. 6.3 (dash-dot line). It has also been noted [84, 107] that for nuclear recoils below ~ 10 keVr in Xe, the underlying Thomas-Fermi interaction potential may be a poor approximation, and the Lindhard model for nuclear quenching may not be valid. This is roughly where our $\mathcal{L}_{eff}$ solution ceases to fall, and may indicate the point at which the electronic stopping becomes comparable to the nuclear stopping [84].

6.3 Determination of the ionization yield Q_y for liquid Xe

The ionization yield Q_y is uniquely determined from the $S1$ and $S2$ data, once $\mathcal{L}_{eff}$ is known. This is shown in Fig. 6.8 (blue triangles) for each of the $\mathcal{L}_{eff}$ spline points, using the $\mathcal{L}_{eff}$ solution shown in Fig. 6.3. The ionization yield Q_y obtain from the $\mathcal{L}_{eff}$ solution for $E_d = 0.73$ kV/cm agrees with previous measurements from [14], which are shown in Fig. 6.8 (filled and open circles/squares, errors not shown). The filled and open data points indicate measurements from 2 different detectors, while the circles (squares) represent measurements at $E_d = 0.10$ ($E_d = 2.0$) kV/cm. The modest variation of Q_y with applied electric field is expected since $\mathcal{L}_{eff}$ (and by inference Q_y) has been measured [13] to have only a weak ($< 10\%$) dependence on the applied electric field.




Figure 6.8: The ionization yield Q_y of Xe for nuclear recoils at electric field $E_d = 0.73$ kV/cm, from the $\mathcal{L}_{eff}$ solution shown in Fig. 6.3 (blue triangles), and from the alternate method described in Sec. 6.3.1 (red diamonds). The alternate method, corrected for the resolution (as described in the text), is shown dashed. Also shown are data from [14] (filled and open circles/squares, errors not shown); the filled and open data points indicate measurements from 2 different detectors, while the circles (squares) represent measurements at $E_d = 0.10$ ($E_d = 2.0$) kV/cm.

6.3.1 An alternate method

As discussed in Sec. 2.2.4, the proportional scintillation ($S2$) used to measure the ionization signal from a nuclear recoil is sensitive to single ionization electrons extracted into the gas phase. $S2$ pulses are clearly identifiable over > 5 orders of magnitude. About 41% of the signal appears as a diffuse wash in the bottom PMT array, and 59% is localized in the top PMT array above the (x, y) extraction point. The free electron lifetime for an ionization electron in Xe during the experiment was measured to be $\tau > 2$ ms, which implies a 2.5% probability of attachment for an electron drifting across the fiducial volume ($dt = 50 \mu\text{s}$). However, this effect scales with the distance an electron travels and was easily corrected for (Sec. 4.5.4).

The multiplicity ($m = 1 \dots 3$) of an elastic nuclear recoil event is defined here as the number of scatters (vertices) with recoil energy above a given threshold. Thus, the analysis in Sec. 6.2 was restricted to events with $m = 1$ and an $S2$ threshold of 200 phe (8 ionization electrons). The multiplicity is a sensitive function of the specified recoil energy threshold, so that a specific $S2$ phe threshold ($S2_{thr}$) in the data can be correlated with a keVr equivalent threshold (E_{thr}) from the Monte Carlo. Due to the typical $0.7 \mu\text{s}$ FWHM of $S2$ pulses, scatters with $dz < 3$ mm are not easily resolvable, regardless of their separation in (x, y) . This criteria is duplicated in the Monte Carlo analysis.

The χ^2 distribution for the data versus Monte Carlo comparison of scatter multiplicity is shown as a function of E_{thr} in Fig. 6.9 (right). In all cases, only scatters in the 5.4 kg fiducial volume are counted; in




Figure 6.9: **(left)** The number of single, double, etc. scatters in the fiducial volume with $S2$ (per vertex) > 1600 phe (black circles) is compared with the Monte Carlo prediction, only counting scatters with energy deposition $E > E_{thr}$ (red line). The comparison is a sensitive function of the Monte Carlo energy threshold, below which a scatter is not counted. The best fit corresponds to $E_{thr} = 14.2 \pm 0.7$ keVr. **(right)** The χ^2 distribution (red circles) for the comparison shown at left is a function of the recoil energy threshold imposed in the Monte Carlo. A 2^{nd} -order polynomial fit (black line) is used to extract the minimum and the 1σ uncertainty (indicated by dashed lines).

this particular example, each scatter vertex was required to have $S2 > 1600$ phe ($S2_{thr} = 1600$). The best fit is obtained for $E_{thr} = 14.2 \pm 0.7$ keVr. The overall normalization and the recoil energy threshold are free parameters, so the minimum $\chi^2 = 2$ gives a p value of 0.16 ($d.o.f. = 1$). The ionization yield per keVr for nuclear recoils can then be calculated as $\mathcal{Q}_y = S2_{thr}/S2_{1e-}/E_{thr}$.

The comparison exemplified in Fig. 6.9 was made for a range of $S2_{thr}$ to obtain the ionization yield for nuclear recoils as a function of energy, as shown in Fig. 6.8 (red diamonds). The example (Fig. 6.9) corresponds to the data point at 14.2 keVr in Fig. 6.8, with $\mathcal{Q}_y = 1600/24/14.2 = 4.7$ ionization electrons/keVr. The error bars include the uncertainty from each fit (Fig. 6.9, right) and the uncertainty in the mean of the single-electron $S2$ peak. The χ^2 distributions for all data points shown in Fig. 6.8 exhibit a clear minimum, a $\chi^2/d.o.f. \sim 1$ and are well-fit by a 2^{nd} -order polynomial.

6.3.2 The utility of the alternate method

The major benefit of the alternate method described in Sec. 6.3.1 is that it provides an independent (albeit indirect) check on the $\mathcal{L}_{eff}$ solution shown in Fig. 6.3. The alternate method (Fig. 6.8, red diamonds) predicts a $\mathcal{Q}_y$ that is systematically lower by 12% (at 25 keVr) than is inferred from the best-fit $\mathcal{L}_{eff}$ solution. This is due to the comparison of E_{thr} from the Monte Carlo (which has perfect resolution) with $S2_{thr}$ from the data. The expected variation in $S2$ due to Poisson fluctuations in the number of ionization electrons is 9% at 25 keVr. The alternate method, adjusted to account for the effects of Poisson fluctuations, is shown as a dashed line in Fig. 6.8.

From 2 – 14 keVr, the alternate method (incorrectly) predicts a constant $\mathcal{Q}_y \sim 5$ (not shown in Fig. 6.8). In this regime, the method is limited by the light collection for primary scintillation ($S1$). This somewhat unexpected result can be understood from Fig. 6.4. Lines of constant $S2$ lie parallel to the line corresponding to $S2 = 200$ phe (8 ionization electrons), which is shown dashed in the lower left. The line

for $S2 = 1600$ phe intersects the $+3\sigma$ contour at about 0.5 keVee. This means that for $S2 > 1600$ phe, the multiplicity comparison is not affected by the loss of low-energy events in which $S1$ fails the $n \geq 2$ coincidence requirement (or fluctuates below detection threshold). Conversely, for $S2 < 1600$ phe (14.2 keVr), Poisson fluctuation in $S1$ leads to a loss of events. In contrast, this loss of events is fully accounted for by the $S1$ peak-finding efficiency in the determination of $\mathcal{L}_{eff}$.

6.4 XENON10 energy threshold

The energy threshold of a liquid Xe detector using proportional scintillation for $S2$ is typically determined by its sensitivity to primary scintillation ($S1$) and by $\mathcal{L}_{eff}$. The most likely form of $\mathcal{L}_{eff}$ was determined based on the nuclear recoil calibration of the XENON10 experiment. The energy calibration obtained from the form of $\mathcal{L}_{eff}$ is not substantially different from that obtained using the simplifying assumption of $\mathcal{L}_{eff} = 0.19$. The energy threshold of XENON10 with $\mathcal{L}_{eff} = 0.19$ was found to be about 3 keVr, for an event acceptance > 0.99 . This can be seen from Fig. 4.8, with the scaling of 0.98 $S1$ phe/keVr discussed in Sec. 2.6.

The energy threshold of XENON10 (for an event acceptance > 0.99) was found to be closer to 4 keVr based on the $\mathcal{L}_{eff}$ determined from the AmBe calibration data. This can be seen from Fig. 6.2 and Fig. 6.3. The absolute ionization yield of Xe for nuclear recoils was also determined. For a Xe detector using proportional scintillation to measure the ionization signal, this clearly shows that the ultimate threshold of Xe detectors for nuclear recoils is as low as ~ 1 keVr, and is presently limited by the light collection for primary scintillation ($S1$). The technique used to determine $\mathcal{L}_{eff}$ for primary scintillation ($S1$) is also expected to be applicable to the secondary scintillation ($S2$) or to the energy determined from a combined ($S1 + S2$) scale. Future work should explore this possibility.

Chapter 7

MUON TRACKS IN XENON10

“Who ordered that?”

– Isidor Isaac Rabi

7.1 Ionizing Tracks

XENON10 was not intended to study muon interactions in Xe. The DAQ was tuned for energy deposition $\lesssim 150$ keVee, and for larger energy depositions, saturation of the ADC is observed. For even larger (say, $\gtrsim 1$ MeV) PMT saturation can also occur. As discussed in Sec. 3.2.3, a total of 69 muon tracks were observed in 58.6 live days of WIMP Search data. The energy deposition along an ionizing track can be measured. Saturation effects present a major problem in accurately reconstructing muon energy, but a method to obtain a measure of the specific energy loss of muon tracks in Xe (in spite of the saturation) is described here.




Figure 7.1: An typical track from a muon which entered through the top of the detector and traversed its entire length. The track is shown for the sum of all the PMTs (red), for 43 of the top-array PMTs (black) and for only PMT number 14 (blue). The track for PMT 14 was already shown in Fig. 3.9, along with a close-up of the $S1$ signal. Copious δ -ray production seen along the track. The red arrow indicates a single δ -ray of about 250 keVee (referenced to the track baseline).

An example track from a muon which entered through the top of the detector is shown in Fig. 7.1. The entry through the top is inferred from the minimal delay between the $S1$ and the start of the ionizing track. That the zenith angle of the track was $\theta < 53^\circ$ can be inferred from the fact that the muon traversed the entire detector. The track is shown for sum of all the PMTs (red), for 43 of the top-array PMTs (black) and for only PMT number 14 (blue). Only 43 top-array PMTs are shown because the 4 PMTs directly above the muon entry point saturate heavily. The track for PMT 14 was already shown in Fig. 3.9, along with a close-up of the $S1$ signal. The $S1$ signal saturates the ADC in nearly all bottom-array PMTs, and several of the top-array PMTs. The downward slope seen in the summed tracks is indicative of PMT (as opposed to ADC) saturation. The track runs off the end of the digitizer ($\lesssim 5 \mu\text{s}$ lost), because in WIMP Search operation it was assumed that interesting events would trigger by their $S2$; muons, in contrast, trigger exclusively by their primary scintillation.

Relativistic muons are very nearly minimum ionizing particles, meaning their specific energy loss $\langle -dE/dX \rangle$ is close to the minimum for any ionizing particle. The concept of specific energy loss is very similar to that




Figure 7.2: Two further examples of muon tracks. In these cases, PMT 14 saturated as a result of closer proximity to the track. **(top)** As in Fig. 7.1, this muon entered through the top of the detector. The ionization signal received in PMT 14 tends to increase for $80 \lesssim dt \lesssim 90 \mu\text{s}$, during which it periodically saturates the ADC. After $dt \sim 100 \mu\text{s}$, the PMT response is saturated. **(bottom)** A muon which entered the detector at $z \approx 9$ cm below the top of the active detector volume (as determined by the delayed onset of the ionization track). The ionization signal received in PMT 14 tends to increase during the first $\sim 5 \mu\text{s}$.

of linear energy transfer *LET* (or stopping power) shown in Fig. 2.14. The difference is that specific energy loss includes energy depositions due to bremsstrahlung, which can occur at a distance from the track of the primary particle. LET only includes the loss due to the primary particle track [93]. The specific energy loss is a function of target Z and also the energy of the incident particle. The typical energy of muons traversing the LNGS underground cavern is $E = 270 \pm 18$ GeV [112]. For a 270 GeV muon track in Xe, $\langle -dE/dx \rangle = 1.9 \pm 0.1$ MeV g⁻¹ cm² [165]. Muons are clearly not a problematic background for XENON10; however, as discussed in Sec. 3.2.4, neutrons produced along their tracks could be.

7.2 Saturation

The pulse shape along the length of all the observed muon tracks is distorted by PMT saturation, and truncated by ADC saturation. The ADC saturation is exemplified in Fig. 7.2 (top) by the response of PMT 14, in the range $85 \lesssim \mu s \lesssim 95$. PMT saturation is exemplified by the downward slope evident in the summed waveform. Specific channels at a distance from the track may avoid saturation, as shown for PMT 14 in Fig. 7.1 and Fig. 7.2 (bottom). Completely unsaturated channels give a clear picture of the roughly constant energy loss (excepting δ -rays) along the muon ionizing track. Of course this comes with a loss of light collection efficiency, and is therefore difficult to compare to known energy calibrations.

7.2.1 Quick refresher on PMT saturation

As discussed in Sec. 3.3.5, a typical PMT in XENON10 is operated with 780 V across a 12.6 M Ω voltage dividing network. The bias current was therefore about $I_b = 62$ μ A. As discussed in Sec. 3.3.5, to ensure a linearity of better than $\pm 3\%$ in pulse-mode operation, the total charge draw of the pulse Q_0 (at the anode) must be a factor $\times 100$ less than the reserve charge Q_R in the final-stage capacitor [75].

The example of an S2 pulse from a 122 keV photon was given in Sec. 3.3.5 of. About 1×10^5 phe are generated ($Q_0 \simeq 32$ nC) with a full width at half-maximum of ~ 0.6 μ s, and the 4 PMTs over the vertex tended to receive about 25% of the signal between them. This lead to a typical maximum charge draw of $Q_0 \approx 2$ nC per PMT. The PMT response for typical 122 keV signals was not limited at all because the S2 width was small (pulse-like), and the requisite charge Q_0 was provided by the reserve charge $Q_R \approx 680$ nC in the capacitors. The capacitors (often called “decoupling” capacitors) were in parallel with the resistor at each of the last 3 dynode stages, and the bulk of the benefit came from the last stage. Assuming a sufficiently large capacitance and a count rate (or pulse duration) much less than the recharge time of the capacitor, a PMT would then be limited primarily by space-charge saturation. For the R8520 PMT the space-charge saturation limit is specified to be > 30 mA for $\pm 2\%$ deviation [74].

7.2.2 Limits of decoupling capacitors

For a pulse with extended duration, the specification in terms of the current draw is that $I_0 < 100 I_b$ [75]. To understand the severity of this limitation, imagine the same S2 from a 122 keV γ had an extended width. This would correspond to an average instantaneous current

$$I_0 = \frac{1}{16} \frac{10^5 \text{ phe}}{0.6 \mu s} \cdot G \cdot q \approx 3 \text{ mA} \quad (7.1)$$

where $G = 2 \times 10^6$ is the typical PMT gain for a single phe, and $q = 1.6 \times 10^{-19}$ C. This current value is a factor $\times 50$ larger than the bias current I_b .




Figure 7.3: The response of a single top-array PMT for the muon track shown in Fig. 7.1. The $S1$ for this PMT barely saturated the ADC ($< 5\%$ of pulse lost). **(left)** A zoom on the beginning of the track. The average instantaneous charge draw (indicated by the arrow) from the PMT anode is about 2 mA. **(right)** The entire waveform begins to decay with $\tau \approx 33 \mu\text{s}$ due to the depletion of the reserve charge Q_R . The magnitude of τ varies, depending on the instantaneous current at the anode. Compare the track seen by PMT 14 in Fig. 7.1.

A similar situation is exemplified by a muon track in Fig. 7.3 for a single top-array PMT. This is the same muon track shown in Fig. 7.1. The average instantaneous current draw on the PMT during the first $2 \mu\text{s}$ (signal enclosed by the rectangle) is $I_i \approx 2 \text{ nC } \mu\text{s}^{-1} \approx 2 \text{ mA}$. This current load is sufficient to completely deplete Q_R on the final-stage capacitor in $t_R \sim 334 \mu\text{s}$. The support of the other two decoupling capacitors must be considered, along with the effect of anode space-charge saturation. The result is that the PMT response begins to fall after a time $t \lesssim 5 \mu\text{s}$ with a time constant determined by these factors and the amount of scintillation seen by the PMT. The curve shown in Fig. 7.3 for this PMT has $\tau = 33 \mu\text{s}$ and is meant as a guide to the eye only. Other PMTs show a different τ commensurate with the scintillation signal they saw.

A more severe example is provided by Fig. 7.4, for a typical bottom-array PMT. In this case, the $S1$ is heavily saturated. A simple 2-pole pulse with $\tau_{\text{rise}} = 10 \text{ ns}$ and $\tau_{\text{fall}} = 27 \text{ ns}$ (Fig. 7.4, left, red dotted) has a total area of $\sim 4 \times 10^5 \text{ phe}$. Scaling for a minimum ionizing γ or β at 2.1 phe/keV (see Sec. 2.5.4) gives an equivalent of about 200 MeV for this PMT alone. The corresponding inferred charge draw on the anode is 140 nC, or 20% of Q_R . Given the pulse duration, the instantaneous current draw on the anode is $I_i \approx 1.5 \text{ A}$, a factor of $\times 50$ above the space-charge saturation limit specified by Hamamatsu. Considering space-charge saturation, in which a 20% sag in anode output can be observed for a peak anode current of only 100 mA [75], 200 MeV equivalent signal for this PMT is very much a lower limit. The constant plateau at about 2 phe/sample is seen in nearly all the bottom-array PMTs, after a period of about $50 \mu\text{s}$. The corresponding instantaneous current is $I_i \approx 67 \mu\text{A}$; presumably this is all the PMT can manage once Q_R is depleted.

With an understanding of these limitations, it is possible to extract some useful information from the muon tracks.




Figure 7.4: The response of a single bottom-array PMT for the muon track shown in Fig. 7.1. **(left)** The $S1$ for this PMT is heavily saturated. A simple 2-pole pulse with $\tau_{rise} = 10$ ns and $\tau_{fall} = 27$ ns (red dotted) has a total area of $\sim 4 \times 10^5$ phe. Scaling for a minimum ionizing γ or β at 2.1 phe/keVee (see Sec. 2.5.4) gives an equivalent of about 200 MeV for this PMT alone. The corresponding inferred charge draw on the anode is 140 nC, or 20% of Q_R . **(right)** The full track in this channel shows a more abrupt fall than observed in the top-array PMT (Fig. 7.3, left), and then maintains a minimum output of about 2 phe/sample. Compare the track seen by PMT 14 in Fig. 7.1.

7.3 Energy calibration for muon tracks

Because of the location of the gas gap (where the $S2$ is generated) with respect to the top PMT array, the $S2$ signal for γ , β or neutron recoils tends to be highly concentrated in 4 PMTs. The muon tracks are no exception. Based on low-energy recoils from the AmBe calibration data (in which saturation is not an issue), the remaining 43 top-array PMTs receive $28.5\% \pm 2.4\%$ (1σ error) of the $S2$. The same data shows that the $S2$ light collection ratio for the top and bottom PMT arrays is $\mathcal{R}_{TB} \equiv S2_{top}/S2_{bottom} = 1.44$. The final requisite piece of information is the $S2$ response from the full-energy peak of single-scatter ^{137}Cs calibration data, since the 0.662 MeV electron is also very nearly a minimum ionizing particle. A typical value of $S2_{bottom} = 2.67 \times 10^5$ phe is found, considering only the (unsaturated) bottom PMT array. The implication is $S2_{top} = 3.84 \times 10^5$ phe on the top-array PMTs. The scaling factor for the 43 top-array PMTs which received the least signal is then $\mathcal{E}_{43} = 0.662 \text{ MeV} / (3.84 \times 10^5 \text{ phe})$. This factor should be increased by 5% to account for the difference in LET between a γ of 0.662 MeV and a minimum ionizing γ of 1 MeV (as discussed in Sec. 2.5.4).

7.4 Muon energy loss in Xe

The specific energy loss of muons in Xe was calculated using only 43 of the 47 top-array PMTs (ignoring the non-functional PMT 48). This choice is motivated by the fact that the top-array shows less severe – and in some cases no – saturation. The 4 top-array PMTs directly above the muon entry into the Xe target are ignored since they also saturate heavily. This approximation is somewhat less justified as the muon

track angle becomes significantly different from zero. To limit the deviation, the analysis only considered muons which entered through the top and exited through the bottom of the active target. Such a track looks similar to that of Fig. 7.1. The (x, y) position of the muon was reconstructed along the length of its track, following the methods described in Sec. 4.4.3. This allowed the actual path length of the muon track to be calculated (rather than assuming it to be 15 cm).




Figure 7.5: The measured specific energy loss along the first 10 μs of the muon track. for 26 muon events in XENON10. The value predicted by a Monte Carlo simulation of 270 GeV c^{-2} muons traversing XENON10 is about $1.7 \text{ MeV cm}^2 \text{ g}^{-1}$.

The average muon energy loss along the first 10 μs of the track was calculated as:

$$\text{MeV cm}^{-1} = 1.05 \mathcal{E}_{43} N_{phe} N_s \eta_{43}^{-1} \quad (7.2)$$

where $N_s = 553$ digitizer samples per cm, $\mathcal{E}_{43} = 0.662 \text{ MeV} / (3.84 \times 10^5 \text{ phe})$ and the signal efficiency for the 43 least-hit top-array PMTs is $\eta_{43} = 0.285$. The initial average value of the summed (43 top PMTs) waveform N_{phe} (measured in phe/sample) is exemplified by the red rectangle in Fig. 7.3; the actual value is given for the track in Fig. 7.1 as $N_{phe} \approx 650$.

It was possible to extract the initial value of N_{phe}/sample for 26 tracks. Most of the events which were cut did not meet the requirement of entry/exit through the top/bottom of the active target. An additional $\sim 10\%$ of events were cut because they showed excessive saturation in the first $\sim 10 \mu\text{s}$ of the track. The energy loss rate (MeV cm^{-1}) from Eq. 7.2 was converted to the specific energy loss $\langle -dE/dX \rangle$ ($\text{MeV g}^{-1} \text{ cm}^2$) by accounting for the liquid Xe density ($\rho = 2.86 \text{ g cm}^{-3}$). The results are shown in Fig. 7.5. The expected values are shown in Fig. 7.6. The prediction from a Monte Carlo simulation is $\langle -dE/dX \rangle = 1.7 \pm 0.4 \text{ MeV g}^{-1} \text{ cm}^2$. The XENON10 measurement is about a factor 35% lower, which

should be considered very good agreement considering the uncertainties in the analysis. The bulk of the uncertainty is attributed to the PMT saturation (not ADC saturation), the approximation that the 4 most-saturated top PMTs receive 28.5% of the signal, and the correction for the track angle.

7.5 Comparison with predicted specific energy loss

The specific energy loss as determined by a Monte Carlo simulation of 1000 downward-directed muons ($E_\mu = 270 \text{ GeV c}^{-2}$) traveling along the axis of the XENON10 detector is shown in Fig. 7.6. The simulation is bounded by the predictions of $1.9 \text{ MeV g}^{-1} \text{ cm}^2$ and $3.6 \text{ MeV g}^{-1} \text{ cm}^2$ from [165], as shown in Fig. 7.6. The energy deposition along z from the Monte Carlo was scaled to XENON10 drift time using $v_d = 1.9 \text{ mm } \mu\text{s}^{-1}$. The lower energy deposition in the first $10 \text{ } \mu\text{s}$ is also (often) observed in the XENON10 data, in channels with no PMT or ADC saturation. The effect is most likely a result of the downward-directed momentum of δ -ray production (and other secondary particles) along the muon track.

The Monte Carlo results are bounded by the predicted values of $\langle -dE/dX \rangle$ from the Particle Data Group (PDG) [165]. The lower bound **(a)** ignores radiative energy losses, and the upper bound **(b)** accounts for radiative losses. In the former case, data were given for Fe, Sn and Pb (PDG Fig. 27.3 in “Passage of particles through matter”). The value for Xe was interpolated with minimal uncertainty. In the latter case (PDG Fig. 27.20), data were only given for Fe, and this value was scaled (lower) by 10% based on the fractional difference inferred from PDG Fig. 27.3. This scaling may not be fully justified, and the upper bound is indicated for reference only.

Typical muon tracks in XENON10 show sustained ionizing tracks (in the first $\sim 10 \text{ } \mu\text{s}$) of roughly 5000 phe/sample, considering all PMTs (Fig. 7.1 and subsequent figures). Considering the ADC and PMT saturation, the actual value is probably higher than 1×10^4 phe/sample. This order of magnitude estimate is interesting because at the same time, an ionizing track of as few as ~ 2 phe/sample would be resolvable in XENON10. Such a track would correspond to a specific energy loss $\langle -dE/dX \rangle \approx 400 \text{ eV g}^{-1} \text{ cm}^2$. No known particle has such a low specific energy loss.




Figure 7.6: A Monte Carlo simulation of the energy loss of 270 GeV muons in XENON10 (red +). The blue rectangle indicates the region from which the mean energy loss (as shown in Fig. 7.5) was extracted from the data. The arrow indicates the mean value of $\langle -dE/dX \rangle$ in this region, as obtained from the simulation. The predicted value of $\langle -dE/dX \rangle$ from the Particle Data Group (PDG) [165] is also shown (a) ignoring radiative energy losses, and (b) accounting for radiative losses. In the former case, data were given for Fe, Sn and Pb (PDG Fig. 27.3 in “Passage of particles through matter”). The value for Xe was interpolated with minimal uncertainty. In the latter case (PDG Fig. 27.20), data were only given for Fe, and this value was scaled (lower) by 10% based on the fractional difference observed in PDG Fig. 27.3

Chapter 8

CLOSING REMARKS

“We are at the very beginning of time for the human race. It is not unreasonable that we grapple with problems. But there are tens of thousands of years in the future. Our responsibility is to do what we can, learn what we can, improve the solutions, and pass them on.”

– Richard Feynman

The XENON10 detector was deployed as a proof-of-principle experiment, the first step towards larger, more sensitive detectors for the direct detection of dark matter. The technology was found to work extremely well, and as of this writing XENON10 had published the most sensitive WIMP-nucleon exclusion limit. The XENON10 result, along with motivation for dramatic improvements in detector sensitivity to nuclear recoils, were shown in Fig. 1.15.

A survey of the principles of operation of XENON10 has been given, and the characteristic signals that arise from particle interactions in liquid Xe have been illustrated. Efficient reduction and rejection of background, which is crucial for a successful direct detection experiment, has also been discussed. The dominant electron recoil background population was found to be Gaussian in the discrimination parameter $y = \log_{10}(S_2/S_1)$. The discrimination power between the background and the expected nuclear recoil signal was shown to approach 99.9%, with a 50% acceptance for nuclear recoils.

A detailed study was made of the non-Gaussian tails of the background distribution. It was found that these tail events – which had been predicted – could be fully characterized and rejected. Rejection of non-Gaussian tails is crucial in order to minimize contamination of the expected signal region, and to maximize the sensitivity of an experiment. Suggestions for future dark matter direct detection experiments using liquid Xe were given in Sec. 5.2.4.

Achieving the lowest possible recoil energy threshold is also crucial for maximizing a liquid Xe detector's sensitivity to particle dark matter interactions, since the expected recoil energy spectrum falls exponentially. The energy threshold of XENON10 – as with other detectors using liquid noble gases – was limited by its efficiency for the collection of primary scintillation light (S_1). Future experiments are expected to face a similar situation, although they will benefit from improvements in photodetector technology. Understanding the nuclear recoil energy scale is also essential. In Sec. 6 a new method was presented for determining the scintillation efficiency for nuclear recoils. The results of this analysis indicate that the energy threshold for XENON10 was about 4 keVr.

Experimental physicists have been saying for more than a decade that we are on the verge of discovering dark matter. Over the same period, and with each null result, theorists' predictions for the WIMP-nucleon cross section have been in a gradual but steady retreat. So where does that leave us now, with a factor of $\times 100$ improvement in detection sensitivity on the horizon? For some, the next year may bring the tingle of a discovery in the making; for others, the continued undulation of a discovery that was made more than a decade ago. One thing is certain: people will continue to work very hard. But will they see the light?

Bibliography

- [1] P. Adamson *et al.*, (MINOS Collaboration), Phys. Rev. D **76**, 072005 (2007).
- [2] Q. R. Ahmad *et al.*, Phys. Rev. Lett. **89** 011301 (1998).
- [3] C. Alcock *et al.*, Ap. J. **499** L9 (1998).
- [4] C. Alcock *et al.*, Ap. J. **542** 281 (2000).
- [5] G. J. Alner *et al.*, (ZEPLIN-II Collaboration), arXiv:0701858v2 [astro-ph].
- [6] G.J. Alner *et al.*, Astropart. Phys. **23** 444 (2005).
- [7] D. S. Akerib *et al.*, Phys. Rev. Lett. **96** 011302 (2006).
- [8] D. Akimov *et al.*, Phys. Lett. B **524** 245 (2002).
- [9] M. Ambrosio *et al.*, Astropart. Phys. **19** 313 (2003).
- [10] J. Angle *et al.*, Phys. Rev. Lett. **100** 021303 (2008).
- [11] G. W. Angus and S. S. McGaugh, Mon. Not. R. Astron. Soc. **383** 417 (2008).
- [12] N. Aoki and Y. Makide, Chem. Lett. **34** 1396 (2005).
- [13] E. Aprile *et al.*, Phys. Rev. D **72** 072006 (2005).
- [14] E. Aprile *et al.*, Phys. Rev. Lett. **97** 081302 (2006).
- [15] E. Aprile *et al.*, Phys. Rev. B **76** 014115 (2007).
- [16] T. Araki *et al.*, Phys. Rev. Lett. **94** 081801 (2005).
- [17] H.M. Araujo *et al.*, NIM A **521** 407 (2004).
- [18] *Human Health Fact Sheet*, Argonne National Laboratory, August 2005.
- [19] F. Arneodo *et al.*, Nuovo Cim. **112A** 819 (1999).
- [20] F. Arneodo *et al.*, NIM A **449** 147 (2000).
- [21] Y. Ashie, Masters Thesis (not available in English), University of Tokyo (2004).
- [22] P. Astier *et al.*, A&A **447** 31 (2006).
- [23] H. W. Babcock, Lick Obs. Bull **19** 41 (1939).
- [24] G. Bakale *et al.*, J. Phys. Chem. **80** 2556 (1976).
- [25] S. Baker and R. Cousins, NIM **221**, 437 (1984).
- [26] A. Baldini *et al.*, NIM A **545** 753 (2005).

- [27] I. R. Barabanov *et al.*, NIM A **254** 355 (1987).
- [28] L. M. Barkov *et al.*, NIM A **379** 482 (1996).
- [29] Y. Becherini *et al.*, Astropart. Phys. **23** 341 (2005).
- [30] P. Belli *et al.*, Nuovo Cim. **101A** 959 (1989).
- [31] C. L. Bennett *et. al.*, Ap.J. Supp. **148** 1 (2003).
- [32] P. Benetti *et al.*, NIM A **329** 361 (1993).
- [33] P. Benetti *et al.*, arXiv:0701286v2 [astro-ph].
- [34] A. Benoit *et al.*, NIM A **577** 558 (2007).
- [35] M. J. Berger *et. al.*, (XCOM database). <http://physics.nist.gov/PhysRefData/Xcom/Text/XCOM.html>
- [36] R. Bernabei *et al.*, Eur. Phys. J. **C11** 1 (2002).
- [37] A. Bolozdyna, NIM A **422** 314 (1999).
- [38] A. I. Bolozdynya *et al.*, NIM A **579** 50 (2007).
- [39] A. Bosma AJ **86** 1791 (1981).
- [40] R. Bradley *et al.*, Rev. Mod. Phys. **75** 777 (2003).
- [41] W. R. Bruce *et al.*, Rad. Research **19** 606 (1963).
- [42] M. J. Carson *et al.*, Astropart. Phys. **21** 667 (2004).
- [43] V. Chepel *et al.*, *Astropart. Phys.*, **26** 58 (2006).
- [44] Comasud S.R.L., Teramo, Italia.
- [45] E. Conti *et al.*, Phys. Rev. B **68** 054201 (2003).
- [46] CRC Handbook of Chemistry and Physics, 87th Edition, CRC Press, Boca Raton (2006).
- [47] E. Dahl, Ph.D. Thesis, Princeton University (2008).
- [48] M. Danilov *et al.*, Phys. Lett. B **480** 12 (2000).
- [49] I. Dell’Antonio, *private communication*. see also [155].
- [50] L. de Viveiros, Ph.D. Thesis, Brown University (2009).
- [51] A.T.A.M de Waele, private communication.
- [52] S. Dimopoulos and H. Georgi, Nucl. Phys B **193** 150 (1981).
- [53] T. Doke *et al.*, NIM **134** 353 (1976).
- [54] T.Doke *et al.*, NIM A **420** 62 (1999).
- [55] T. Doke *et al.*, Jpn. J. Appl. Phys. **41** 1543 (2002).
- [56] J. Drees Int. J. Mod. Phys. A **17** 3259 (2002).
- [57] A. K. Drukier, K. Freese and D. N. Spergel, Phys. Rev. D **33** 3495 (1986).
- [58] J. Edsjo, Nucl. Phys. B Proc. Supp. **143** 435 92005).
- [59] D. J. Eisenstein *et. al.*, Ap. J. **633** 560 (2005).
- [60] C. Eleftheriadis *et al.* (CAST collaboration), arXiv:0706.0637 [hep-ex].

- [61] J. Ellis *et al.*, Phys.Rev. D **71** 095007 (2005).
- [62] I. Ferreras, M. Sakellariadou and M. F. Yusaf, Phys. Rev. Lett. **100** 031302 (2008).
- [63] M. Fragopoulou *et al.*, J. of Phys.: Conf. Series **41** 514 (2006).
- [64] K. C. Freeman, Ap. J. **160** 881 (1970).
- [65] Y. Fukuda *et. al.*, Phys. Rev. Lett. **81** 1562 (1998).
- [66] R. Gaitskell, *private communication*.
- [67] R. J. Gaitskell, Annu. Rev. Nucl. Part. Sci. **54** 315 (2004).
- [68] R. Gavazzi *et al.*, Ap. J. **667** 176 (2007).
- [69] E. Gates *et al.*, Ap. J. **449** L123 (1995).
- [70] Made by K. Giboni at Nevis Labs.
- [71] M. Girardi *et al.*, Ap. J. **530** 62 (2000).
- [72] R. Gomez, Masters Thesis, Rice University (2007).
- [73] L. Grandi, Ph.D. Thesis, University of Pavia (2006).
- [74] Hamamatsu Corp. R8520 PMT data sheet (2006).
- [75] Hamamatsu PMT Handbook,
http://www.sales.hamamatsu.com/assets/applications/ETD/pmt_handbook/pmt_handbook_complete.pdf
 (see p 163).
- [76] B.A. Hands, *Cryogenic Engineering*. London: Academic Press, 1986. pp 89-121.
- [77] T. Hauyama *et al.*, Presented at 13th Int'l. Cryocooler Conf., New Orleans (2004).
- [78] R. Heaton *et al.*, NIM A **276** 529 (1989).
- [79] C. de los Heros, arXiv:0802.0147v1 [astro-ph].
- [80] G. Heusser, Ann. Rev. Nucl. Part. Sci. **45** 543 (1995).
- [81] G. Hinshaw *et al.*, arXiv:0803.0732v1 [astro-ph].
- [82] A. Hitachi *et al.*, Phys. Rev. B **27** 5279 (1983).
- [83] A. Hitachi, Astropart. Phys. **24** 247 (2005).
- [84] A. Hitachi, J. Phys.: Conf. Ser. **65** 012013 (2007).
- [85] R. Ibata *et al.*, Ap. J. **524** L95 (1999).
- [86] R. Ibata *et al.*, Ap. J. **532** L41 (2000).
- [87] N. Ishida *et al.*, NIM A **384** 380 (1997).
- [88] JENDL-3.3, <http://www.ndc.jaea.go.jp/jendl/jendl.html>
- [89] M.C. Johnson, Applied Optics **7** 879 (1968).
- [90] H. M. Jones and E. E. Kunhardt, J. Phys. D: Appl. Phys. **28** 178 (1995).
- [91] J. Jortner *et al.*, J. Chem. Phys., **42** 4250 (1965).
- [92] G. Jungman, M. Kamionkowski and K. Griest, Phys. Rep. **267** 195 (1996).
- [93] G. Knoll, *Radiation Detection and Measurement*, Wiley & Sons, New York NY (2000).

- [94] H. Koch *et al.*, http://www.euro-inox.org/pdf/paper/Koch-Otto_EN.pdf
- [95] E. Komatsu *et al.*, arXiv:0803.0547 [astro-ph].
- [96] M. Kowalski *et al.*, arXiv:0804.4142v1 [astro-ph].
- [97] S. Kubota *et al.*, J. Phys. C **11** 2645 (1978).
- [98] J. Kwong, Ph.D. Thesis, Princeton University (2009).
- [99] LBNL Isotopes Project, <http://ie.lbl.gov/toi>
- [100] F. LePort *et al.*, (EXO Collaboration), NIM A (in press).
- [101] J. D. Lewin and P. F. Smith, Astropart. Phys. **6** 87 (1996).
- [102] J. Lindhard, Mat. Fys. Medd. K. Dan. Vidensk. Selsk. **33** 1 (1963).
- [103] M. Loewenstein and R. E. White, Ap. J. **518** 50 (1999).
- [104] LUX Collaboration Proposal (2007).
- [105] LUX Collaboration, <http://luxdarkmatter.org>
- [106] D. Mallin *XENON10 Internal Note* (2008).
- [107] A. Mangiarotti *et al.*, arXiv:physics/0610286
- [108] A. Manzur, Ph.D. Thesis, Yale University (2009).
- [109] J.W. Marsh *et al.*, NIM A **366** 340 (1995).
- [110] M. Martin, J. Chem. Phys. **54** 3289 (1971).
- [111] D. N. McKinsey and C. Orzel, NIM A **545** 524 (2005).
- [112] D.-M. Mei and A. Hime, Phys. Rev. D **73** 053004 (2006).
- [113] Y. Messous *et al.*, Astropart. Phys. **3** 361 (1995).
- [114] M. Milgrom, Ap. J **270** 365 (1983).
- [115] L.S. Miller *et al.*, Phys. Rev. **166** 871 (1968).
- [116] MINOS colabration arXiv:0708.1495v2 [hep-ex].
- [117] M.Miyajima *et al.*, NIM B **63** 297 (1992).
- [118] V. Musilova *et al.*, Cryogenics **45** 529 (2005).
- [119] NASA press release 06-297.
- [120] J.F. Navarro *et al.*, Ap. J. **490** 493 (1997).
- [121] K. Ni *et al.*, arXiv:0708.1976v1 [physics.ins-det].
- [122] K. Ni, Doctoral Thesis, Columbia University (2006).
- [123] NIST Thermophysical Properties of Fluid Systems, <http://webbook.nist.gov/chemistry/fluid/>
- [124] P. J. E. Peebles, *The Large-Scale Structure of the Universe*, Princeton University Press, Princeton NJ (1980).
- [125] P. J. E. Peebles, *Principles of Physical Cosmology*, Princeton University Press, Princeton NJ (1993).
- [126] S. Perlmutter and B. P. Schmidt, arXiv:0303428v1 [astro-ph].

- [127] V. A. Rabinovich *et al.*, “*Thermophysical Properties of Ne, Ar, Kr and Xe*”, Hemisphere Publishing Corp. (1988).
- [128] M. Redshaw *et al.*, Phys. Rev. Lett. **98** 053003 (2007).
- [129] M. T. Ressell and D. J. Dean, Phys. Rev. C **56** 535 (1997).
- [130] M Yamashita, K. Ni, G. Plante and A. Manzur, [XENON10 ROOT-based analysis chain].
- [131] L. Roszkowski, Pramana J. Phys. **62** 1 (2004).
- [132] L. Roszkowski *et al.*, arXiv:0705.2012 [hep-ph].
- [133] P. Salucci and M. Persic, ASP Conf. Ser. **117** 1 (1997).
- [134] SAES Pure Gas, Inc., <http://www.saesgetters.com>
- [135] SAES Pure Gas, Inc., *private communication* (2008).
- [136] R. H. Sanders, Mon. Not. R. Astron. Soc. **342** 901 (2003).
- [137] R. H. Sanders, Mon. Not. R. Astron. Soc. **380** 331 (2007).
- [138] R. Sawada *et al.*, Cryogenics **43** 449 (2003).
- [139] G. M. Seidel *et al.*, NIM A **489** 189 (2002).
- [140] C. Skordis *et al.*, Phys. Rev. Lett. **96** 011301 (2006).
- [141] Sloan Digital Sky Survey, <http://sdss.org>
- [142] Y. Sofue *et al.*, Publ. Astron. Soc. Japan **55** 59 (2003).
- [143] V.N. Solovov *et al.*, NIM A **516** 462 (2004).
- [144] P. Sorensen *et al.*, Proc. 6th Intl. Workshop for Identification of Dark Matter, World Scientific (2007).
- [145] P. Sorensen *et al.*, *manuscript in preparation for* Phys. Rev. D. (2008).
- [146] Spectra Gases, Inc. <http://www.spectragases.com>
- [147] D. N. Spergel *et al.*, Ap. J. Supp. **148** 175 (2003).
- [148] G. Steigman, Annu. Rev. Nucl. Part. Sci. **57** 463 (2007).
- [149] P. H. Stelson and F. K. McGowan, Phys. Rev. **133** B911 (1964).
- [150] SOudan LOw Background Gamma Counting Facility, <http://particleastro.brown.edu/SOLO>
- [151] M. Tegmark *et al.*, Phys. Rev. D **74** 123507 (2006).
- [152] D.J. Thomas, *private communication* (see also [109]).
- [153] J. Tinker *et al.*, Ap. J. **631** 41 (2005).
- [154] S. Torres *et al.*, arXiv:0804.2207v1 [astro-ph].
- [155] J. Tyson, G. Kochanski and I. Dell’Antonio, Ap. J. Lett. **498** 107 (1998).
- [156] G. de Vaucouleurs, Ann. d’Astrophys. **11** 247 (1948).
- [157] J. Wambsganss, Living Rev. Relativity **1** (1998).
- [158] B. Watterson “*The Complete Calvin and Hobbes*”, Andrews McMeel Publishing (2005).
- [159] A-M. Weijmans *et al.*, Mon. Not. R. Astron. Soc. **383** 1343 (2008).

- [160] F. Wilczek, Phys. Rev. Lett. **40** 279 (1978).
- [161] H. Wulandari *et al.*, Astropart. Phys. **22** 313 (2004).
- [162] M. Yamashita, *private communication*.
- [163] M. Yamashita *et al.*, NIM A **535** 692 (2004).
- [164] Y.-F. Yang *et al.*, Phys. Rev. D **64** 013012 (2001).
- [165] W.-M. Yao *et al.*, Journal of Physics, G **33**, 1 (2006). (*see Fig 27.3*).
- [166] W.-M. Yao *et al.*, Journal of Physics, G **33**, 1 (2006). (*see p 297*).
- [167] WebElements, <http://webelements.com/>
- [168] WMAP NASA Explorer mission, <http://map.gsfc.nasa.gov>
- [169] S. Yellin, Phys. Rev. D **66** 032005 (2002).
- [170] G. Zaharijas and D. Hooper, Phys. Rev. D **73** 103501 (2006).
- [171] E. Zavattini *et al.*, (PVLAS collaboration), arXiv:0706.3419 [hep-ex].
- [172] J. Zhao *et al.*, Chin. A&A **15** 260 (1991).
- [173] E. Zucca *et al.*, A&A **326** 477 (1997).
- [174] F. Zwicky, Helv. Phys. Acta **6** 110 (1933).
- [175] F. Zwicky, Ap. J. **86** 217 (1936).