UNIVERSITY OF CALIFORNIA
Los Angeles

Results from the DarkSide-50 Dark Matter Experiment

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Physics

by

Alden Fan

2016

© Copyright by
Alden Fan
2016

ABSTRACT OF THE DISSERTATION

Results from the DarkSide-50 Dark Matter Experiment

by

Alden Fan

Doctor of Philosophy in Physics

University of California, Los Angeles, 2016

Professor David Saltzberg, Chair

While there is tremendous astrophysical and cosmological evidence for dark matter, its precise nature is one of the most significant open questions in modern physics. Weakly interacting massive particles (WIMPs) are a particularly compelling class of dark matter candidates with masses of the order 100 GeV and couplings to ordinary matter at the weak scale. Direct detection experiments are aiming to observe the low energy (<100 keV) scattering of dark matter off normal matter. With the liquid noble technology leading the way in WIMP sensitivity, no conclusive signals have been observed yet. The DarkSide experiment is looking for WIMP dark matter using a liquid argon target in a dual-phase time projection chamber located deep underground at Gran Sasso National Laboratory (LNGS) in Italy. Currently filled with argon obtained from underground sources, which is greatly reduced in radioactive ^{39}Ar , DarkSide-50 recently made the most sensitive measurement of the ^{39}Ar activity in underground argon and used it to set the strongest WIMP dark matter limit using liquid argon to date.

This work describes the full chain of analysis used to produce the recent dark matter limit, from reconstruction of raw data to evaluation of the final exclusion curve. The DarkSide-50 apparatus is described in detail, followed by discussion of the low level reconstruction algorithms. The algorithms are then used to arrive at three broad analysis results: The electroluminescence signals in DarkSide-50 are used to perform a precision measurement of

longitudinal electron diffusion in liquid argon. A search is performed on the underground argon data to identify the delayed coincidence signature of ^{85}Kr decays to the ^{85m}Rb state, a crucial ingredient in the measurement of the ^{39}Ar activity in the underground argon. Finally, a full description of the WIMP search is given, including development of cuts, efficiencies, energy scale, and exclusion curve in the WIMP mass vs. spin-independent WIMP-nucleon scattering cross section plane.

This work was supervised by Hanguo Wang and was completed in collaboration with members of the DarkSide collaboration.

The dissertation of Alden Fan is approved.

Jay Hauser

Huan Z. Huang

David A Paige

David Saltzberg, Committee Chair

University of California, Los Angeles

2016

To my father

TABLE OF CONTENTS

1	Introduction	1
1.1	Evidence for dark matter	1
1.1.1	Motion of stars and galaxies	1
1.1.2	Gravitational lensing	3
1.1.3	Bullet cluster	4
1.1.4	Cosmological parameters	5
1.1.5	Big Bang Nucleosynthesis	8
1.2	Dark matter candidates	12
1.2.1	Baryonic dark matter	12
1.2.2	Neutrinos	13
1.2.3	Axions	13
1.2.4	WIMPs	14
1.3	Detection of dark matter	18
1.3.1	WIMP event rate	19
1.3.2	Direct detection techniques	21
1.3.3	Current status of direction detection	24
2	DarkSide	27
2.1	Backgrounds	28
2.2	Liquid argon	30
2.3	Underground Argon	34
2.4	Dual-phase LAr TPC	35
2.5	DarkSide-50	39

2.5.1	TPC	40
2.5.2	Outer Detectors	47
3	Reconstruction	54
3.1	DarkArt - TPC	54
3.1.1	Technical overview	55
3.1.2	Modules	56
3.2	DarkArt - OD	64
3.3	Single photoelectron mean	65
3.4	Position reconstruction	66
3.5	SLAD	69
3.5.1	Pulse identification	71
3.5.2	Corrections	73
3.6	Commonly used variables	78
3.7	Electronics MC	79
3.8	Reconstruction efficiencies and biases	81
4	S2	86
4.1	Introduction	86
4.2	S2 pulse shape	86
4.2.1	Basic shape	87
4.2.2	Gaussian smearing	87
4.2.3	Degeneracy of Parameters	89
4.3	Fitting S2 pulse shape	93
4.3.1	Maximum likelihood fitting	93
4.3.2	Goodness-of-fit	96

4.4	Gas pocket drift time and zero-diffusion events	99
4.4.1	Data selection	100
4.4.2	Fitting procedure	106
4.4.3	Results	107
4.5	Electron diffusion	113
4.5.1	Data selection	115
4.5.2	Fitting procedure	116
4.5.3	Results	119
4.5.4	Systematics	122
4.5.5	Comparison to literature	127
4.6	Derivations of analytic S2 pulse shape	129
4.6.1	Idealized S2 pulse shape	129
4.6.2	S2 pulse shape with Gaussian smearing	132
5	UAr Depletion	134
5.1	AAr vs. UAr S1 spectra	135
5.2	The DarkSide-50 Monte Carlo	138
5.3	Spectral fitting	139
5.4	Delayed coincidence measurement	140
5.4.1	Event selection	143
5.4.2	Results	152
5.5	Summary	155
6	WIMP search with UAr	157
6.1	Run selection and stability	157
6.1.1	PMT gain	158

6.1.2	Light yield	158
6.2	Cuts	165
6.2.1	TPC cuts	167
6.2.2	Veto cuts	180
6.3	f_{90}	183
6.3.1	Ratio of Gaussians	183
6.3.2	Fitting ER data	187
6.3.3	Fitting NR data	195
6.4	WIMP search region	197
6.5	Energy scale	200
6.6	Exposure and NR acceptance	201
6.6.1	Total live time	201
6.6.2	Total NR acceptance	203
6.7	Exclusion curve	203
6.7.1	Combining curves	208
6.7.2	Projected sensitivity	208
7	Conclusion and outlook	211
7.1	Summary	211
7.2	Future	212
A	The ARIS experiment	215
A.1	Proposed Measurement	217
A.2	Experimental Setup	218
	Bibliography	222

LIST OF FIGURES

1.1	Rotation curve for galaxy NGC 3198.	2
1.2	Gravitational lensing around the galaxy cluster Abell 2218.	3
1.3	The Bullet Cluster	4
1.4	Measurements of the CMB temperature from the Planck experiment	7
1.5	Power spectrum of the CMB temperature anisotropy for different values of Ω_b	9
1.6	Light elemental abundances from Big Bang Nucleosynthesis.	11
1.7	Dark matter relic density as function of time from the early universe	17
1.8	Expected WIMP-induced nuclear recoil spectrum.	22
1.9	Current and projected limits of WIMP direct detection experiments from 2013.	25
2.1	Cartoon of a dual-phase LAr TPC.	36
2.2	Example raw waveforms for typical ER and NR events	38
2.3	Conceptual drawing of the DarkSide-50 system.	39
2.4	Conceptual drawing the DarkSide-50 TPC.	41
2.5	Electric field simulation of the DarkSide-50 TPC	44
2.6	View of the inside of the LSV	49
2.7	View of the inside of the WCD	52
3.1	Example raw waveform for a single channel in DarkSide-50.	55
3.2	Example waveforms showing the baseline overlaid on raw waveform data.	59
3.3	Sum of baseline-subtracted waveforms without zero-suppression, showing coherent noise across channels.	60
3.4	Example sum channel waveform composed of the sum of zero-suppressed baseline-subtracted waveforms.	61
3.5	(x, y) distribution of ^{39}Ar events using barycenter position reconstruction	68

3.6	(x, y) distribution of ^{39}Ar events using WLS position reconstruction	70
3.7	Light yield relative to TPC center, measured from ^{83m}Kr calibration data and ^{39}Ar data	75
3.8	Mean S2/S1 response vs. transverse position of ^{83m}Kr events.	76
3.9	Distribution of S2/S1 vs. t_{drift} for evaluation of electron drift lifetime.	77
3.10	Template single PE response generated by averaging many single PEs	80
3.11	Example interpolation portion of the baseline-finder algorithm.	82
3.12	Bias in reconstructed size of S1 and S2 and f_{90}	84
3.13	Efficiency of DarkArt pulse finder to identify the start of S1 and S2 pulses	85
4.1	Example S2 pulse shape.	88
4.2	Sequence for building a family of different S2 pulses with nearly identical pulse shape.	90
4.3	Degeneracy of parameters in S2 pulse shape.	92
4.4	Cartoon illustrating the fixed length integration window starting at fixed offset from the S2 pulse start.	96
4.5	Example spectra built from MC and real data.	97
4.6	Re-binning of the S2 pulse shape to a flat distribution.	98
4.7	Example event with zero diffusion.	101
4.8	f_{90} vs. total PE plane for candidate zero-diffusion events.	104
4.9	Cross checks that f_{90} -based selection (Eqn. 4.15) successfully isolates zero-diffusion events.	105
4.10	Onset of ADC saturation in merged S1+S2 zero-diffusion events.	106
4.11	Reduced χ^2 of fitted S2 pulse shape to zero-diffusion events.	108
4.12	Example zero-diffusion events with good and poor S2 fits.	109
4.13	2D profile of average T vs. transverse position.	111

4.14	Distribution of slow component lifetime τ_2 , extracted from fits of S2 in zero-diffusion events.	112
4.15	2D histogram of the fast component fraction p vs. r^2	113
4.16	Examples of S2 pulse shape fits for the electron diffusion measurement. . . .	120
4.17	Reduced χ^2 of S2 pulse shape fits to events in diffusion analysis.	121
4.18	σ_L^2 vs. t_{drift} for measurement of electron diffusion.	122
4.19	Diffusion measurement using various data slices.	124
4.20	Diffusion measurement at different drift fields.	126
4.21	T vs. r^2 and p vs. r^2 for standard 2.8 kV/cm and 2.3 kV/cm extraction field data.	128
4.22	Diffusion measurement at different extraction fields.	129
4.23	Effective electron energy ϵ_L vs. drift field.	131
5.1	AAr and UAr S1 spectra at null field and 200 V/cm drift field.	136
5.2	AAr and UAr S1 spectra at 200 V/cm drift field with veto cuts.	137
5.3	Fit to the null field and 200 V/cm drift field S1 spectra of UAr data.	141
5.4	Decay scheme of ^{85}Kr	142
5.5	Simulated waveforms for $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$	143
5.6	f_{90} vs. S1 distribution for simulated $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays and data. .	145
5.7	Example waveform of an event from the uppermost portion of the LAr, giving an S1 and S2 very near each other.	147
5.8	Distribution of f_{5000} vs. peak time for simulated delayed coincidences and data.149	
5.9	Distributions of top-bottom asymmetry parameter a_{S1} and S1 max. frac. for candidate delayed coincidence events.	150
5.10	Candidate $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decay.	153
5.11	Spectrum of $\text{S1}_\beta + \text{S1}_\gamma$ for all candidate events.	154

5.12	β - γ separation in $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays.	154
6.1	Trend of SPE mean over time for each TPC PMT in the (a) bottom array and (b) top array. Plot from X. Xiang.	159
6.2	Low energy peak in UAr data due to ^{37}Ar activity.	161
6.3	^{37}Ar spectrum of integrated 7 μs ROI in expected S1 window	162
6.4	Decay of ^{37}Ar peak over time, fitted with an exponential plus constant . . .	163
6.5	Light yield trend over time as measured from the ^{37}Ar peak in UAr data . .	164
6.6	LSV light yield trend using the ^{60}Co peak	166
6.7	Estimated efficiency of single scatter cut to accept NRs as a function of S1. .	173
6.8	S1 maximum fraction distributions for a few representative (S1, t_{drift}) bins. .	176
6.9	Distribution of asymmetry parameter a_{max} in AAr and UAr data	178
6.10	t_{drift} distribution for single scatter events in the UAr data set.	180
6.11	Toy MC of f_{90} for ERs, fitted with the ratio-of-Gaussians model	186
6.12	Fitted f_{90} distributions for various S1 slices of AAr and UAr data	191
6.13	Evidence for non-single-scatter ER events in UAr data	194
6.14	NR acceptance contours in the f_{90} vs. S1 plane.	198
6.15	ER leakage contours in the f_{90} vs. S1 plane.	199
6.16	WIMP search results in the f_{90} vs. S1	199
6.17	NR acceptance as function of energy	204
6.18	Graphical illustration of steps in the toy WIMP MC.	206
6.19	90 % C.L. exclusion limit for WIMP-nucleon scattering in DarkSide-50	207
6.20	Projected 3 yr sensitivity for DarkSide-50.	210
7.1	Sketch of DarkSide-20k.	213
7.2	Projection sensitivity of DarkSide-20k and ARGO.	214

A.1	Previous measurements of relative scintillation efficiency ($\mathcal{L}_{\text{eff}}$) in other experiments. Red squares (“This work”) refer to Ref. [119].	217
A.2	Schematic of the LAr TPC in the neutron beam with neutron detectors (green) at various angles.	218
A.3	(a) Fully assembled LAr TPC at UCLA. (b) TPC installed within cryostat and integrated into full cryogenics, high voltage, and data acquisition systems.	220

LIST OF TABLES

4.1	Initial values and limits of fit parameters for zero-diffusion events.	108
4.2	Initial values and limits of fit parameters.	118
4.3	Drift velocities in liquid for different drift fields in DarkSide-50.	125
4.4	Diffusion constant D_L measured from different data with different drift and extraction fields.	130
6.1	Measurements of f_{90} medians from SCENE and DarkSide-50.	196
6.2	Measurements of $\mathcal{L}_{\text{eff}}$ from the SCENE experiment.	200
6.3	Tabulation of livetime after successive application of quality cuts	203

ACKNOWLEDGMENTS

Working on DarkSide has been a tremendous collaborative effort, and I am thankful for the opportunity to work with so many people around the world. First and foremost, I would like to thank my adviser, Hanguo Wang. It has been a privilege and a pleasure to observe, learn from, and collaborate with him while designing and building a variety of complex instruments and detectors. I have learned many invaluable lessons from him both in and out of the lab over the last four years. I would like to thank Emilija Pantic, from whom I have learned the most about dark matter searches and on whom I can always rely to answer my many questions. I am indebted to Peter Meyers, who first introduced me to DarkSide seven years ago while serving as my adviser for my first junior paper and senior thesis, as well as Cristiano Galbiati for his nearly superhuman leadership of the collaboration.

There are a great many others within DarkSide that I would like to acknowledge. Stephen Pordes's sharp wit and sharper intellect have been an inspiration. Jeff Martoff's experience and knowledge of nuclear physics seems to know no bounds. Masayuki Wada has shown me many tips and tricks for doing analysis and it has been a joy to work with him. Yann Guardincerri and Ken Herner provided much needed assistance with many data processing tasks. The veto gurus, Stefano Davini and Shawn Westerdale, are responsible for nearly everything related to the vetoes in this thesis. With great help from Xin Xiang and Luca Pagani (the Just Do It crew), the WIMP search analysis was finished in a (somewhat) timely manner. Alissa Monte and Xinran Li provided crucial inputs to the S2 pulse shape and electron diffusion analyses. Luca Grandi and Richard Saldanha were great mentors while working on the reconstruction in the early days of DarkSide-50. I have benefitted from collaboration with many more, including the rest of the Princeton group (Guangyong Koh, Hao Qian, Chris Stanford, and Chengliang Zhu); the Hawaii group (Erin Edkins, Brianne Hackett, and Bernd Reinhold); and the Paris group (Davide Franco and Paolo Agnes).

It has also been a pleasure to work with the UCLA group. I would like to thank Yuri Suvorov for his gracious assistance with anything and everything any time I visited LNGS or he came back to UCLA, Yixiong Meng for many many useful discussions about analysis and

hardware and life, Yi Wang for his unwavering willingness to lend a helping hand, and Xiang Xiao for a helpful new perspective in recent months. Special thanks go to Quentin Riffard, Ben Schlitzer, and Andrew Watson for their assistance with the small TPC at UCLA while I was writing this thesis.

Finally, I would like to thank the great support from my friends and family. Blake has been an exceptional roommate and friend, and I blame him for getting me into climbing. In turn, it has been a joy to climb with Mike and Emily, both indoors and outdoors. The weekly pesto night, with Alli, Meghan, James and Fuson, has provided a regular respite from the stress of graduate school and will surely be missed. I can also always count on my brothers when I need to switch gears or take a break. Lastly, I would like to thank my parents. Surely, I would not be where I am today without their love and guidance.

VITA

2011	A.B. (Physics), Princeton University.
2011	Kusaka Memorial Prize in Physics, Princeton University.
2012	Roy and Dorothy John Fellowship, UCLA.
2012	Teaching Assistant, Physics and Astronomy Department, UCLA.
2013	M.S. (Physics), UCLA, Los Angeles, California.
2011–present	Graduate Student Researcher, Physics and Astronomy Department, UCLA.

PUBLICATIONS

Agnes et al. *The veto system of the DarkSide-50 experiment*. JINST 11 (2016) P03016.

Agnes et al. *Results from the first use of low radioactivity argon in a dark matter search*. Phys. Rev. D 93 (2016) 081101(R).

Agnes et al. *First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso*. Phys. Lett. B. 743 (2015) 456.

Alexander et al. *Light yield in DarkSide-10: A prototype two-phase argon TPC for dark matter searches*. Astropart. Phys. 49 (2013) 44.

CHAPTER 1

Introduction

The nature of dark matter is one of the most important open questions of modern physics. A multitude of evidence has accumulated over the last century for the existence of a non-luminous (dark) matter, and it is now well established that dark matter makes up 27% of the energy density of the universe, and 85% of the matter content. However, all of the evidence relies on its gravitational effects at astrophysical scales. The precise properties of dark matter are still to be unveiled. The leading dark matter candidates are particle in nature, and a host of experiments are underway to detect these elusive particles.

In Sec. 1.1, we highlight some of the historical evidence for the existence of dark matter, and in Sec. 1.2, we describe a few of the most popular dark matter candidates. Finally, in Sec. 1.3 we give an overview of the current status of dark matter searches using direct detection.

1.1 Evidence for dark matter

1.1.1 Motion of stars and galaxies

The first evidence for a missing matter component in the universe comes from Fritz Zwicky, who in the 1930s, calculated the velocity dispersion of galaxies in the Coma cluster by measuring their red shift and applied the virial theorem to estimate the average mass of each galaxy [1, 2]. Combining his results with average luminosity measurements gave a mass-to-luminosity ratio M/L (normalized to solar mass and solar luminosity) several hundred times larger than the expectation of order 1. Zwicky proposed that a non-luminous form of matter could account for the discrepancy and coined the term “dark matter” [1].

Figure 1.1: Rotation curve for galaxy NGC 3198. From Ref. [4].

In the 1970s, Vera Rubin firmly established the need for a dark matter component when she measured the rotation curves of a host of galaxies [3]. Rubin measured rotation curves by selecting edge-on galaxies and comparing the redshift of the 21 cm line from stars on opposite sides (approaching and receding). From Newtonian dynamics, one expects that the velocity of stars and gas in circular orbit should follow $v(r) \propto \sqrt{M(r)/r}$, where $M(r)$ is the mass internal to radius r from the galactic center. Beyond the optical disk, where the bulk of the luminous matter of the galaxy resides, $v(r)$ should have a $1/\sqrt{r}$ dependence. Instead Rubin found that the rotation curves were constant well beyond the optical disk for many galaxies, as in Fig. 1.1, suggesting a spherical dark matter halo with mass density $\rho(r) \propto 1/r^2$. The density will have to fall off faster at some point in order to keep the total mass of the galaxy finite, and the extent of the dark matter halo is still not well known, but the constant profile of rotation curves remains among the strongest evidence for the existence of dark matter.

Figure 1.2: Gravitational lensing around the galaxy cluster Abell 2218. The lensed objects are seen as elongated arcs. From Ref. [5].

1.1.2 Gravitational lensing

Gravitational lensing provides another powerful tool for establishing the existence of dark matter. One of the consequences of General Relativity is that massive objects, such as galaxies or galaxy clusters, distort their surrounding space-time. In gravitational lensing, light from a distant source object behind the massive object is bent, or lensed, on its way to Earth. The image of the source object can be multiplied, amplified, and distorted into arcs or rings, as shown, for example, in Fig. 1.2. The Hubble Space Telescope and other survey telescopes have observed lensed objects around numerous galaxies and galaxy clusters and the inferred masses have consistently indicated that the M/L values are large, 10 to 20 for galaxies and 100 to 300 for galaxy clusters [6] and require the existence of large amounts of dark matter. Gravitational lensing techniques can further determine the mass density

Figure 1.3: The Bullet Cluster is formed from two galaxy clusters that have recently collided. The colored orange regions indicate the location of hot gases, measured from X-ray observations, and the contours indicate the mass distribution of both clusters, measured from gravitational lensing observations. The white bar indicates 200 kpc. From Ref [8].

profiles within galaxies and galaxy clusters, though understanding the precise details of the dark matter profiles is still an area of active research [7].

1.1.3 Bullet cluster

The Bullet Cluster (1E0657-558) [8] is composed of two galaxy clusters that have recently collided (recent on cosmological timescales) and provides another striking piece of evidence for the existence of dark matter. The clusters are observed separately through luminosity measurements, from the Chandra X-ray telescope, and gravitational lensing measurements, from the Hubble Space Telescope, as shown in Fig. 1.3. The luminosity measurements indicate the location of hot gases, which make up most of the clusters' baryonic mass.

Lensing measurements indicate the overall mass distributions of the two clusters. The hot gases are observed to have interacted with each other electromagnetically, causing them to decelerate as the two clusters collided. Meanwhile, the mass is distributed on either side of the luminous matter, suggesting that most of the mass content of the two clusters passed through each other unimpeded. This is strong evidence for a dominant but non-interacting (or at least very weakly self-interacting) dark matter component in both clusters.

1.1.4 Cosmological parameters

The strongest evidence to date for the existence of dark matter comes from measurements of the cosmic microwave background (CMB). These measurements give the most precise accounting for the matter and energy content of the universe, which is conventionally written, for species ‘i’:

$$\Omega_i = \frac{\rho_i}{\rho_{\text{cr}}} \quad (1.1)$$

where ρ_i is the physical density, and ρ_{cr} is the critical density necessary for a flat universe, given by

$$\rho_{\text{cr}} = \frac{3H_0^2}{8\pi G} \quad (1.2)$$

where H_0 is the Hubble parameter and G is the gravitational constant. For a flat universe, $\Omega_{\text{tot}} = 1$; for a closed universe (spherical space-time geometry), $\Omega_{\text{tot}} > 1$; and for an open universe (hyperbolic geometry), $\Omega_{\text{tot}} < 1$. As will be discussed below, we now have strong evidence that the universe is flat, so Ω_x can be interpreted as the fractional density of species ‘i’ in the universe.

Before describing how observations of the CMB give rise to measurements of the dark matter density Ω_{DM} , we give a brief discussion of the early universe. In the Big Bang model of cosmology, the earliest stage of the universe was a period of extremely rapid expansion by a factor 10^{27} over 10^{-32} s (inflation) [9], followed by a slower expansion for another 380,000 years. During this time, the universe was composed of a hot dense plasma of ionized particles. Photons in the sea of charged particles have short scattering lengths, making the baryons and photons tightly coupled and the plasma opaque. The plasma can then be

approximated as a single photon-baryon fluid. The fluid compresses and expands (acoustic oscillations) due to two competing effects: Compression occurs as the fluid falls into gravitational wells left over from inflation. As the fluid compresses, the pressure of the fluid increases due to the high electromagnetic interaction rate until it forces expansion outward. The pressure decreases and gravitational attraction takes over again. As the universe expanded and cooled, the temperature eventually fell below the atomic ionization energy threshold (a few eV), and the ionized plasma formed neutral atoms, primarily hydrogen. This is called recombination, which occurred about 380,000 years after the Big Bang. Due to the absence of charged particles, the photons could free stream until we see them today as the CMB [10]. Thus we measure the CMB from the “surface of last scattering” which has a temperature of ~ 3000 K [9]. From recombination to today, the universe expanded by a factor of ~ 1000 and the temperature dropped by a similar factor. We observe the temperature of the CMB to be 2.73 K [11].

Measurements of the CMB have led to a multitude of discoveries about the universe. One of the early results was the discovery of temperature anisotropy. Due to the acoustic oscillations of the photon-baryon fluid, the temperature at the surface of last scattering has minute variations. The temperature map of the CMB therefore imprints the density fluctuations of the universe at recombination. However, in the 1990s, the COBE experiment discovered that the temperature fluctuations were extremely tiny [12], too small to account for the structure of galaxies and galaxy clusters. Figure 1.4a shows a recent temperature map observed by the Planck experiment [13]. The baryonic density fluctuations implied by the observed CMB fluctuations are too small to account for the gravitational wells necessary to form the structure that we see today. Structure formation must have started before recombination, while the universe was a hot plasma, which can only happen with a neutral form of matter—dark matter—that clumps into gravitational wells before recombination.

The power spectrum of the temperature anisotropies encodes a tremendous amount of information about the universe. Figure 1.4b shows the power spectrum observed by the Planck experiment [13]. For example, one of the crowning achievements is the measurement of the geometry of the universe. The sound horizon, the maximum distance a sound wave

(a)

(b)

Figure 1.4: Measurements of the CMB temperature from the Planck experiment. (a) Full sky temperature map. (b) Power spectrum of the temperature anisotropy with the best-fit prediction from Λ CDM cosmology. From Ref. [13].

in the primordial plasma could have traveled in the time before recombination, is used as a sort of measuring stick. The size of the largest structures (globs) in the temperature map corresponds to the sound horizon and governs the location of the first peak in the power spectrum. Combined with the age of the universe, researchers determined that if the universe is flat, the location of the peak should be $\sim 1^\circ$, corresponding to a multipole moment of $l \sim 200$; if it is open, the peak should be at larger angular scale; and if it is closed, the peak should be at smaller angular scale. The Mobile Anisotropy Telescope measured the first peak to be at $\sim 1^\circ$ [14]. The WMAP and Planck experiments confirmed the result to exquisite precision [10, 13]. It is now well established that the universe is flat. Another result is that the relative amplitudes of the peaks in the power spectrum tell us about the relative amounts of baryonic matter and dark matter in the universe. Figure 1.5 illustrates the point by showing the expected power spectra for different densities of baryons and dark matter (different Ω_b and Ω_{DM}) and comparing to data. In this way, the values of Ω_b and Ω_{DM} can be tightly constrained.

The above arguments are simplifications, but they capture the broad picture. A host of experiments in the last few decades, including COBE, WMAP, and Planck, have measured the CMB with ever increasing precision. Today, the data are fit with a sophisticated 6 parameter model of Λ CDM that includes all the cosmological parameters. The results give the highest precision measurements of the fractional densities of the baryons, dark matter, and dark energy of the universe: $\Omega_b = 0.049$, $\Omega_{DM} = 0.268$, and $\Omega_\Lambda = 0.683$ [16].

1.1.5 Big Bang Nucleosynthesis

While not a measure of the dark matter density, Big Bang Nucleosynthesis (BBN) provides further incontrovertible evidence that baryonic matter constitutes only $\sim 4\%$ of the universe. Immediately after the Big Bang, the universe was a primordial soup of subatomic particles—quarks, leptons, gauge bosons, Higgs bosons, and dark matter—all being created and annihilated in equilibrium. After 10^{-5} s, the quark-hadron transition occurs: as the universe expands and the temperature drops, quarks combine into hadrons (protons and

Figure 1.5: Power spectrum of the CMB temperature anisotropy for different values of Ω_b and Ω_{DM} , normalized to $\Omega_{\text{tot}} = 1$. From Ref. [15].

neutrons) [17]. After 3 minutes, BBN occurs as the temperature drops enough for the formation of stable deuterium, which then allows for the formation of other light elements: ^3He , ^4He , and ^7Li [18].

The abundances of these isotopes at the end of BBN can be predicted using nuclear physics and known reaction rates. A rough estimate of the abundance of ^4He , the most stable of the light elements, can be determined using the neutron-proton ratio n/p , where n is the number of neutrons and p the number of protons, at the time of BBN. Immediately after the Big Bang, there were equal numbers of protons and neutrons, but as neutrons decay into protons, n/p drops over time. Because neutron decay is a weak process, the change is relatively slow, and at the time of BBN, $n/p \approx 1/7$ [18]. Because there are more protons than neutrons, the maximum amount of ^4He that could be produced is $2n$, so the maximum ^4He fraction is $Y = 2n/(n + p) = 1/4$. The prediction of the abundances of the other elements is more complex and is done with the help of numerical simulations of the element production and destruction processes.

Measurements of the elemental abundances after BBN were performed by observing very distant (hence very old) portions of the universe that have small fractions of heavy elements, indicating low levels of chemical enrichment. The elemental abundances in these areas are closest to their primordial values. The observations agree with the predictions extremely well [18], as shown in Fig. 1.6. This result is one of the pillars of the Hot Big Bang model.

Finally, the results from BBN confirm that baryonic matter constitutes only a small fraction of the universe. From General Relativity, we know that the expansion rate of the universe is influenced by the total amount of matter—more matter, faster expansion. Were there more baryonic matter, the expansion rate would have been faster, giving neutrons less time to decay into protons, increasing n/p , increasing Y . Likewise, the deuterium abundance would have been different. Observations of the primordial elemental abundances are consistent only with $\Omega_b \approx 0.04$, as can be seen in Fig. 1.6, with the deuterium abundance placing the strictest constraint. These results reinforce the need for a non-baryonic dark matter component of the universe.

Figure 1.6: Light elemental abundances from Big Bang Nucleosynthesis as a function of baryon density. The colored bands represent predictions. The vertical band is the baryon density from CMB measurements. The black boxes and arrows are from observations. From Ref. [19].

1.2 Dark matter candidates

The existence of a dark matter component of the universe is now well established, but its nature remains unknown. In order to fit the picture painted by the astrophysical evidence described in the previous section, dark matter candidates should satisfy several criteria. They should be electrically neutral, interacting with ordinary matter at most at the weak scale; they should be stable on cosmological timescales, or they would have decayed by now; and they should be cold (non-relativistic), in order to satisfy constraints set by structure formation [20]. Many candidates have been proposed. Particle dark matter candidates are usually developed to solve some other problem in particle physics, and the fact that they can potentially solve the dark matter problem makes them particularly alluring. We highlight some of the most popular dark matter candidates here.

1.2.1 Baryonic dark matter

Early candidates for dark matter were baryonic in form. Non-luminous stellar remnants such as white dwarfs and neutron stars, often referred to as MAssive Compact Halo Objects (MACHOs), were proposed to populate galactic halos [21]. These objects can be looked for via microlensing observations, in which the object passes in front of a distant luminous background object, inducing a transient change in brightness due to gravitational lensing. The microlensing results indicate that MACHOs can account for no more than 20 % of what would be needed to account for all dark matter in the galactic halo [22].

Various forms of black holes have also been proposed to constitute at least some of the dark matter [21]. While supermassive black holes at the centers of galaxies and black holes formed by stellar collapse have largely been ruled out by CMB and BBN constraints, primordial black holes, formed when the universe was still a hot plasma, are still a viable candidate [23]. Primordial black holes are, by nature, dark and for masses greater than 10^{15} g, long-lived since they can survive Hawking radiation until the present [24]. However, recent limits have constrained the possibility for primordial black holes as a significant fraction of the dark matter to very small mass windows [25]. Interestingly, $20 M_{\odot}$ to $100 M_{\odot}$ primordial

black holes remain open as possible explanations for dark matter, and the gravitational waves observed by the LIGO experiment [26] come from $30 M_{\odot}$ black holes [27].

1.2.2 Neutrinos

In the original Standard Model (SM), the neutrino is left-handed and massless and the three neutrino flavors e, μ, τ are separately conserved. But near the turn of the millennium, the discovery of neutrino flavor oscillations proved that they do indeed have mass [28, 29]. Being electrically neutral and weakly interacting, a massive neutrino might appear to be a promising dark matter candidate. For this to be the case, the neutrino mass would need to be ~ 10 eV [30]. However, measurements of the CMB have constrained the sum of neutrino masses to be much smaller (< 0.44 eV [10]). None of the three neutrino species can solve the dark matter problem.

That neutrinos have mass requires new physics beyond the Standard Model. One method for resolving the issue is to add right-handed neutrinos which have zero electric, weak, and strong charges—sterile neutrinos [31]. These neutrinos could mix with the three active neutrinos and would be massive and interact gravitationally. The lightest sterile neutrino makes a plausible dark matter candidate, but it is unstable due to its mixing with the active neutrinos, and it would need to have a lifetime of the scale of the age of the universe. This sets the lightest sterile neutrino mass at a few keV. Sterile neutrinos are also attractive because they can explain the matter-antimatter asymmetry in the universe. Because the sterile neutrino is a decaying dark matter candidate, it is not completely dark, producing a narrow decay line in the X-ray band, which can be looked for [31]. Sterile neutrinos have not yet been completely ruled out.

1.2.3 Axions

The axion is a light pseudoscalar first proposed by Peccei and Quinn to solve the strong CP problem [32]. The Standard Model Lagrangian includes a QCD term

$$\mathcal{L}_{\text{SM}} = \dots + \frac{\theta g^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu} \quad (1.3)$$

where $G_{\mu\nu}^a$ are the QCD field strengths, g is the QCD coupling constant, and θ is a parameter. QCD depends on θ only through the combination of parameters $\bar{\theta} = \theta - \arg \det m_q$, where m_q is the quark mass matrix [33]. Because physics depends on $\bar{\theta}$, the value of $\bar{\theta}$ is determined by experiment. If $\bar{\theta} = 0$, QCD violates P and CP. The term in Eqn. 1.3 by itself is CP violating. Yet the experimental bounds, set by the neutron electric dipole moment [34], yield $\bar{\theta} < 10^{-9}$ [33]. It is a puzzle, referred to as the strong CP problem, why the value of $\bar{\theta}$ is so small. Peccei and Quinn proposed that a spontaneously broken U(1) symmetry, resulting in a Goldstone boson, the axion, could account for the unexpected conservation of CP symmetry in the strong interaction. The value of $\bar{\theta}$ would be allowed to relax to zero, solving the strong CP problem [35]. The allowed parameters for axions imply that they are extremely light and weakly interacting [36]. If they are abundant enough, the axion then makes another dark matter candidate. The main mode for axion detection is its decay to two photons. There are several experimental efforts underway to look for axions in the mass range of 10^{-6} eV to 10^{-3} eV [37, 38].

1.2.4 WIMPs

The most popular class of dark matter candidate is the Weakly Interacting Massive Particle (WIMP). They are just as the name describes: massive particles, typically in the GeV to TeV range, that interact with ordinary matter at the weak scale or below. WIMPs are particularly interesting because they arise naturally in a variety of beyond-Standard-Model theories and they naturally have the correct relic density to account for all the dark matter in the universe. One of the most popular such theories is supersymmetry (SUSY), which aims to solve the hierarchy problem in the Standard Model.¹

¹The hierarchy problem is the question of why the electroweak scale is so small: the Higgs mass, and therefore the scale of electroweak processes, is 125 GeV, while the Planck mass, the scale at which quantum effects of gravity become strong, is 1.2×10^{19} GeV.

1.2.4.1 Supersymmetry

In supersymmetry, every SM particle has a supersymmetric partner, which has the same quantum numbers and gauge interactions as the SM particle but differs in spin by $1/2$. Each boson has a superpartner fermion and vice versa. The superpartners of the fermions have the prefix ‘s’ (e.g. slepton, squark, selectron, stop) and the superpartners of the bosons have the suffix ‘ino’ (e.g. gravitino, wino, Higgsino). Mass eigenstates of the superpartners include the neutralino, a mixture of the photino, zino, and neutral Higgsino, and the chargino, a mixture of winos and charged Higgsinos. Masses of the supersymmetric particles are typically of the order TeV in order to solve the hierarchy problem [39].

By itself, supersymmetry allows for baryon and lepton number violation [39], which is not the case in ordinary SM physics. To resolve this, it is common to add R -parity to SUSY, $R = (-1)^{3B+L+2s}$, where B is baryon number, L is lepton number, and s is spin. For SM particles, $R = 1$, while for their superpartners, $R = -1$. Since every interaction involving SUSY requires an even number of SUSY particles, R -parity then implies that the lightest supersymmetric particle (LSP) is stable. Furthermore, in much of SUSY parameter space, the neutralino is the LSP [40]. Since the neutralino is weakly interacting, electrically neutral, has zero color, and should have mass of order TeV, it is a...weakly interacting massive particle.

That the LSP is stable has significant cosmological implications. All SUSY particles formed in the early universe will have decayed, and we should see a cosmic relic density of LSPs. The LSP could therefore be a dark matter candidate. Will its relic density be high enough to account for all the dark matter observed by, say, CMB measurements? Hint: yes.

1.2.4.2 The WIMP relic density

In the early universe, the temperature is high enough that the massive dark matter is in thermal equilibrium with ordinary matter. As the universe expands and the temperature drops, lighter particles no longer have enough kinetic energy to produce heavier dark matter particles and the gas of dark matter particles becomes so dilute that they cannot find each

other to annihilate. The (co-moving) number density approaches a constant (“freeze out”), which we observe as the dark matter relic density.

The time evolution of the dark matter density is described quantitatively by the Boltzmann equation [41]

$$\frac{dn}{dt} = -3Hn - \langle \sigma_A v \rangle (n^2 - n_{\text{eq}}^2) \quad (1.4)$$

where n is the number density of dark matter particle χ , H is the Hubble constant, $\langle \sigma_A v \rangle$ is the thermally averaged annihilation cross section, and n_{eq} is the dark matter number density in thermal equilibrium. On the right-hand side, the first term accounts for dilution from expansion. The n^2 term arises from the process $\chi\chi \rightarrow \text{SM SM}$ that destroy χ particles, and the n_{eq}^2 term arises from the reverse process, which creates χ particles.

It is convenient to recast Eqn. 1.4 in terms of $Y = n/s$ and $x = m_\chi/T$, where s is the time dependent entropy density of the universe, m_χ is the particle mass and T is the temperature. Using the Friedmann equation to relate the Hubble constant to the energy density and assuming a radiation dominated universe, one obtains

$$\frac{dY}{dx} = -\sqrt{\frac{\pi}{45G}} \frac{g_*^{1/2} m_\chi}{x^2} \langle \sigma_A v \rangle (Y^2 - Y_{\text{eq}}^2) \quad (1.5)$$

where $g_*^{1/2}$ is a parameter which depends on the effective degrees of freedom. See Ref. [40] for a lucid derivation of Eqn. 1.5, which must be solved numerically. Figure 1.7 shows a solution for a 100 GeV particle. In particular, at late times, as the temperature drops to $T \ll m_\chi$, then $x \gg 1$ and $dY/dt \rightarrow 0$, so Y becomes constant. Notice that the equilibrium density Y_{eq} drops off exponentially with time.

We can approximate the final relic density analytically (following Refs. [42] and [43]). First, define the dimensionless quantity

$$\lambda = \sqrt{\frac{\pi}{45G}} g_*^{1/2} m_\chi \langle \sigma_A v \rangle. \quad (1.6)$$

At late times, Y_{eq} drops off exponentially, so Eqn. 1.5 becomes

$$\frac{dY}{dx} \simeq -\frac{\lambda Y^2}{x^2} \quad (1.7)$$

Figure 1.7: Dark matter relic density as a function of temperature T and time t for a 100 GeV particle. The solid gray contour is for an annihilation cross section that yields the correct relic density, and the shaded regions are for cross sections that differ by 10, 10^2 , and 10^3 from this value. Larger cross sections result in lower relic density. The dashed gray contour is the number density of a particle that remains in thermal equilibrium. From Ref. [41].

which can be integrated analytically from freeze-out to today, where $x \rightarrow \infty$, to yield

$$\frac{1}{Y_0} - \frac{1}{Y_f} \simeq \frac{\lambda}{x_f} \quad (1.8)$$

where the subscript 0 denotes present values and the subscript f denotes values at freeze-out. Numerical solutions show that at freeze-out, x_f varies only slightly with particle mass and can be considered a constant $x_f \sim 20$ [41]. The abundance at freeze-out is generally much larger than the present abundance [42], so Eqn. 1.8 becomes

$$Y_0 \simeq \frac{x_f}{\lambda} \quad (1.9)$$

The relic dark matter density is then

$$\Omega_{\text{DM}} = \frac{m_\chi s_0 Y_0}{\rho_{\text{cr}}} \simeq \frac{m_\chi s_0 x_f}{\rho_{\text{cr}} \lambda} = \sqrt{\frac{45G}{\pi}} \frac{s_0 x_f}{\rho_{\text{cr}} g_*^{1/2}} \frac{1}{\langle \sigma_A v \rangle} \quad (1.10)$$

where s_0 is the present day entropy density. To lowest order, the relic density is insensitive to the dark matter mass m_χ and inversely proportional to the annihilation cross section $\langle \sigma_A v \rangle$. Plugging in numerical values yields [42]

$$\Omega_{\text{DM}} \simeq \frac{10^{-39} \text{ cm}^2}{\langle \sigma_A v \rangle}. \quad (1.11)$$

In order to obtain the dark matter relic density of today ($\Omega_{\text{DM}} = 0.268$), one needs annihilation cross sections of order 10^{-39} cm^2 , approximately the weak scale. This is the WIMP miracle. For this reason, WIMPs are one of the most favored dark matter candidates and are the focus of the remainder of this work.

1.3 Detection of dark matter

It is now well established that dark matter exists on astrophysical scales. We observe it through its gravitational effects. The leading candidate for dark matter is a WIMP particle. How can we search for WIMPs other than through their gravitational interactions? What are their particular properties? There are generally three different approaches: collider searches, indirect detection, and direct detection.

In collider searches, one tries to produce dark matter by colliding Standard Model particles together at very high energy in such accelerators as the Large Hadron Collider (LHC). Dark matter will be produced in particle-antiparticle pairs which will escape detection. The signature for dark matter in collider searches is missing energy.

In indirect detection, one looks for Standard Model particles as products of WIMP annihilation. The exact annihilation process depends on the specifics of the dark matter model, but in many cases, WIMPs are Majorana particles (they are their own antiparticle) and will therefore annihilate with each other. Frequently, experiments look for annihilation products such as gamma rays, neutrinos, and antimatter. The signature for dark matter in indirect detection is an excess of Standard Model particles.

In direct detection, a WIMP scatters off a Standard Model particle, and one looks for the recoiling particle. Most often, one looks for WIMP scattering off atomic nuclei. We usually model the WIMP-nucleus interaction as an elastic scatter, so the energy of the recoiling nucleus from a WIMP of mass m_χ is given by

$$E_R = \frac{\mu_\chi^2 v^2 (1 - \cos \theta)}{m_N} \quad (1.12)$$

where $\mu_\chi = m_\chi m_N / (m_\chi + m_N)$ is the reduced mass of the WIMP-nucleus system (assuming a target nucleus of mass m_N), v is the speed of the WIMP relative to the nucleus, and θ is the scattering angle in the center of mass frame. For WIMPs in the tens to thousands of GeV range, and WIMP velocity of 200 km/s, recoil energy is in the 1-100 keV range. These are relatively low energy deposits that WIMP direct detection experiments must be able to observe.

1.3.1 WIMP event rate

The WIMP direct detection rate in terrestrial detectors depends on several factors: the local dark matter halo density and velocity distribution, the WIMP mass, and the cross section on the target nuclei. The local dark matter density is conventionally placed at $\rho_\chi = 0.3 \text{ GeV/cm}^3$, though the uncertainties are large [44]. The differential recoil rate per unit detector mass for a WIMP of mass m_χ , typically given in units of counts/kg/day/keV,

can be written as²

$$\frac{dR}{dE_R} = \frac{\sigma(q)\rho_\chi}{2\mu_\chi^2 m_\chi} \int_{v_{\min}}^{v_{\text{esc}}} \frac{f(v)}{v} dv \quad (1.13)$$

where $q = \sqrt{2m_N E_R}$ is the nucleus recoil momentum, $\sigma(q)$ is the WIMP-nucleus interaction cross section, the integral gives the mean inverse speed and accounts for the velocity distribution $f(v)$ of WIMPs in the halo, $v_{\min} = \sqrt{m_N E_R / 2\mu_\chi^2}$ is the minimum WIMP velocity able to generate a recoil of energy E_R , and v_{esc} is the maximum WIMP velocity set by the escape velocity of the halo model. The velocity distribution is commonly assumed to be Maxwellian. The cross section can be factored as

$$\sigma(q) = \sigma_0 F^2(q) \quad (1.14)$$

where σ_0 is the cross section at zero momentum transfer and $F(q)$ is a nuclear form factor, which accounts for the finite size of the target nucleus, depends principally on the nuclear radius and recoil energy, and may differ for spin-dependent and spin-independent interactions. We will consider only spin-independent interactions here. The cross section depends on the number of protons and neutrons in the nucleus [46]:

$$\sigma_0 = \sigma_n \left(Z \frac{f_p}{f_n} + (A - Z) \right)^2 \frac{\mu_\chi^2}{\mu_n^2} \quad (1.15)$$

where σ_n is the WIMP-*nucleon* cross section, Z is the number of protons, A is the atomic mass number, f_p and f_n are the WIMP couplings to the proton and neutron, and μ_n is the WIMP-nucleon reduced mass (treating the neutron and proton masses as equal).

Putting it all together, we have³

$$\frac{dR}{dE_R} = \frac{\sigma_n \rho_\chi}{2\mu_n^2 M_\chi} \left(Z \frac{f_p}{f_n} + (A - Z) \right)^2 F^2(q) \int_{v_{\min}}^{v_{\text{esc}}} \frac{f(v)}{v} dv. \quad (1.16)$$

It is typical to assume WIMP couplings to proton and neutrons are the same, so we take $f_p = f_n$. Thus, Eqn 1.16 has an A^2 dependence, which means heavier elements can expect

²See Refs. [39, 45] for derivations of Eqn. 1.13

³A note on units: Writing the c^2 term explicitly, the cross section is in units of cm^2 ; ρ_χ is in units of $\text{GeV}/c^2/\text{cm}^3$; masses are in units of GeV/c^2 ; $F^2(q)$ is unitless; and the integral of the velocity distribution yields units of inverse velocity $(\text{km/s})^{-1}$ (see, for example, Ref. [46]). Putting it all together, dR/dE_R is in units of $\text{cm}^{-1}(\text{GeV}/c^2)^{-2}(\text{km/s})^{-1}$. Notice that the conversion of one of the GeV/c^2 terms to kg allows us to write the rate in units of $\text{kg}^{-1}\text{day}^{-1}\text{keV}^{-1}$.

to see higher event rates for a given σ_n and target mass, at least in the low-recoil-energy regime where scattering is coherent. The velocity integral is worked out in Refs. [45, 46] and results in a roughly exponential spectrum. Figure 1.8a shows the result for standard parameters and different target nuclei. With $f_p = f_n$, the only free parameters in Eqn 1.16 are the WIMP-nucleon cross section and the WIMP mass. WIMP search results are then typically given in the cross section vs. mass plane.

As can be seen from Fig. 1.8a, the expected rates are very low, of order 10^{-4} /kg/day/keV. Direct dark matter experiments today will not observe enough WIMP-induced recoils to measure the full recoil spectrum, but instead attempt to identify and count recoils above some threshold (set by detector efficiency and backgrounds). Therefore, it is common to plot the integrated recoil rate as a function of threshold, as shown in Fig. 1.8b. For even the lowest thresholds, the integrated rate is of the order 10^{-2} counts/kg/day, or a few per kg per year.

1.3.2 Direct detection techniques

We can now start outlining the kind of apparatus necessary for detecting WIMP-induced recoils on a terrestrial scale. The low energy nuclear recoil expected from a WIMP scatter demands a low threshold detector, one that can observe recoil energies of tens of keV or lower. The energy readout must be very efficient to achieve a low energy threshold. The low expected event rate (we call these experiments rare event searches) demands extraordinarily low backgrounds. Frequently, the dominant backgrounds are due to electron recoils from γ -rays and beta particles. WIMP detectors should be able to discriminate between the two types of recoils. Neutrons are also a dangerous background because they can produce the same nuclear recoils expected from WIMPs. WIMP detectors should be composed of ultra low radioactivity components to reduce neutron backgrounds. Detectors can also exploit the multiple scattering of neutrons (either within the detector itself or in a surrounding instrumented volume) due to their finite interaction length to reduce the effective neutron background. WIMPs will interact only once, while the neutrons can interact multiple times.

Figure 1.8: (a) Expected WIMP-induced nuclear recoil spectrum for Xe (black), Ar (red), Ge (green), and Si (blue) targets, assuming a WIMP mass of 100 GeV and WIMP-nucleon scattering cross section of 10^{-43} cm^2 and standard halo parameters. (b) Integrated rates as a function of threshold energy.

Finally, the sensitivity to WIMPs scales with detector volume. The technology of choice should be able to scale to large (tonne-scale) volumes to enable the most sensitive searches.

There are typically three channels into which energy is deposited from WIMP-induced recoils: scintillation, ionization, and heat. Different experiments make use of different channels, and many use a combinations of channels.

Scintillation In a scintillator, some of the deposited energy is transferred to electrons in the material, putting them in an excited state. The electrons de-excite with a characteristic decay time and release optical photons. Generally, the number of photons produced is proportional to the energy deposited, allowing for energy reconstruction. Good scintillator materials are transparent to their own scintillation light. Furthermore, the amount of energy deposited in the scintillation channel and the time profile of the scintillation light may depend on the ionization density of the recoiling particle from the incident radiation. Nuclear recoils have a high ionization density, while electron recoils have low ionization density. The scintillation channel can then be used as a powerful tool to discriminate the nuclear recoil signals from the dominant electron recoil backgrounds. The light is read out using a variety of techniques ranging from traditional photomultiplier tubes (PMTs) to phonon-based photosensors used in cryogenic semiconductor detectors. Examples of scintillators used in dark matter searches include sodium iodide (NaI), calcium tungstate crystals (CaWO_4), and liquid nobles (Ar, Xe, Ne).

Ionization If enough energy is transferred to an atomic electron, the atom can become ionized. The free ionization electrons can then be drifted using a strong external electric field and observed by a variety of techniques. For example, in cryogenic detectors using Si or Ge, the electrons are drifted to the crystal surface and read out using charge amplifiers. In dual-phase liquid noble detectors, the electrons are drifted to a gaseous region, where they are again drifted in a higher field to produce electroluminescence light, which can then be read out using some of the light detection techniques described above. As mentioned above, the ionization density varies for nuclear and electron recoils. Therefore, the ionization channel

is also used as a discriminator between signal and background.

Phonons In a crystal lattice, a recoiling nucleus deposits its energy in a series of collisions with other ions and electrons in the lattice. The ions can be collectively excited to form a phonon excitation. The phonons induce a small temperature change in the material, which can be detected using bolometric techniques. The temperature changes are on the order of $1\ \mu\text{K}$, which can only be observed at milli-Kelvin temperatures.

Another technique for direct detection uses superheated fluids, such as CF_3I or C_4F_{10} , in which a nuclear recoil leads to a local nucleation of a bubble. The bubble can be observed with imaging sensors (video cameras) or acoustic sensors. The temperature and pressure of the fluid can be tuned such that only interactions with high dE/dx , such as nuclear recoils, will produce bubbles, which eliminates the abundant electromagnetic backgrounds that plague other WIMP searches. However, there is little energy information encoded in the bubbles, and so these detectors can only count integrated rates above threshold.

1.3.3 Current status of direction detection

There is a wide variety of experiments that employ the techniques just described, as illustrated by the number of experiments shown in Fig. 1.9. At high WIMP masses ($>100\ \text{GeV}/c^2$) no WIMPs have been observed, and experiments have only been able to set exclusion limits. The limits at high mass have been set by experiments using liquid nobles as the detector medium. The XENON [48] and LUX [49] experiments use dual-phase (liquid-gas) xenon detectors to read out the scintillation and ionization channels. The DarkSide experiment uses a similar technique with liquid argon, and is the focus of this work. Other experiments such as DEAP [50] and XMASS [51] use liquid nobles in single phase, relying on the scintillation channel only; these detectors have the advantage that they can scale to large volume very quickly. Indeed the DEAP-3600 detector contains 3.6 t of liquid argon and is preparing to take data in 2016. The current best limit now comes from the LUX experiment [49], whose 90 % confidence level exclusion limit on the WIMP-nucleon spin-independent cross section

Figure 1.9: A compilation of WIMP-nucleon spin-independent cross section limits (solid curves), hints for WIMP signals (shaded closed contours) and projections (dot and dot-dashed curves) for direct detection experiments that are expected to operate over the next decade. Also shown is an approximate band where coherent scattering of ^8B solar neutrinos, atmospheric neutrinos and diffuse supernova neutrinos with nuclei will begin to limit the sensitivity of direct detection experiments to WIMPs. Finally, a suite of theoretical model predictions is indicated by the shaded regions, with model references included. From Ref. [47].

is at $4 \times 10^{-46} \text{ cm}^2$ for a WIMP mass of $33 \text{ GeV}/c^2$.

In the low mass region, there is tension between hints of WIMP signals that have been observed in some experiments while resounding exclusion of those hints by others. The DAMA experiment [52] uses solid NaI crystals to measure the event rate of low energy recoils over time; they do not have any discrimination techniques to select nuclear recoils. They observe an annual modulation that could be a signature of WIMP dark matter (not discussed here) with a mass of $\sim 7 \text{ GeV}$ and cross section $\sim 10^{-40} \text{ cm}^2$ [53]. However, these results have been excluded by several other experiments, including CDMS, XENON100, and LUX [48, 49, 54].

In the next decade or so, direct detection experiments are expected to reach the so-called neutrino floor. Neutrinos are expected to coherently scatter off atomic nuclei [55], but the cross section and recoil energy of such a process is so small that it has not yet been observed. Coherent neutrino scattering will produce nuclear recoils identical to those expected from WIMPs and thus represents an irreducible background to WIMP direct detection experiments. It will become significantly more challenging to look for WIMPs in the face of the coherent neutrino background. The most abundant neutrinos will come from the sun and will pose the biggest background for experiments probing WIMP masses below $10 \text{ GeV}/c^2$. Therefore, the next generation of direct detection experiments are designed to probe the remaining parameter space outside of the neutrino floor: some experiments, such as SuperCDMS [56], are pushing to ever lower thresholds to probe lower mass WIMPs, while others, such as XENON [57], LZ [58], and DarkSide, are building larger mass detectors to probe smaller cross sections at higher mass.

CHAPTER 2

DarkSide

The DarkSide program consists of a staged series of WIMP direct detection experiments. The main detector technology is the dual-phase (liquid-gas) Time Projection Chamber (TPC) with a liquid argon (LAr) target. The premise of the program is two-fold: ultra low backgrounds and extremely powerful rejection of the remaining backgrounds. DarkSide achieves ultra low backgrounds by placing the detector deep underground at Laboratori Nazionali del Gran Sasso (LNGS) in Gran Sasso, Italy to reduce the cosmic ray flux; surrounding the detector with water and liquid scintillator shields to reduce the flux of neutrons and γ -rays that reach the LAr target; and selecting detector components to be radiopure to reduce the radioactive neutron and electromagnetic backgrounds, including the LAr target itself. DarkSide achieves extremely powerful background rejection through a variety of techniques. Foremost is the use of pulse shape discrimination (PSD) provided by the intrinsic scintillation properties of LAr to reject backgrounds from minimum-ionizing particles, which are the most abundant in DarkSide. The ratio of ionization to scintillation signals provides additional discrimination of background electron recoils against the expected WIMP-induced nuclear recoil signal. The LAr TPC technology also enables 3D position reconstruction, enabling the rejection of surface backgrounds. Finally, the water and scintillator shields surrounding the TPC are instrumented (active) to allow efficient rejection of neutron backgrounds.

In this chapter, we give a general discussion of the main backgrounds in WIMP searches using LAr as the target material, discuss in detail the principles of LAr TPCs, and give a technical description of the DarkSide detectors.

2.1 Backgrounds

The expected signal from a WIMP interaction is a recoiling nucleus—or nuclear recoil (NR)—with energy of tens to hundreds of keV. The expected WIMP interaction rate is on the order of a few (or lower) interactions per ton of target material per year. Any low energy particle interaction from Standard Model physics represents background to a dark matter search. The challenge is to identify the rare WIMP events, if any, in the face of the deluge of backgrounds. The main backgrounds in DarkSide-50, and in liquid noble WIMP detection experiments in general, are:

- “internal” backgrounds: radioactivity from the target material itself that result in β particles or X-rays
- “external” γ backgrounds: γ -rays from radioactive components of the detector materials
- surface α -particle radiation
- radiogenic neutrons: neutrons from radioactive decays of the detector materials
- cosmogenics: cosmic-ray muons and induced neutrons

The internal backgrounds are decays of radioactive isotopes present in trace amounts in the target material. For liquid argon, the main concern is β -decay of ^{39}Ar , which is discussed in further detail in the next section (Sec. 2.2). For liquid xenon, the main internal backgrounds are β -decays from ^{85}Kr and X-rays from ^{127}Xe [59]. In both liquid argon and liquid xenon the radioactive isotopes are thoroughly mixed and uniformly distributed in the target material, and the decay products are fully absorbed within mm of the decay site. The signature in all of these cases is an electron recoil (ER), either directly from β -decay or indirectly from electron capture producing an X-ray that is then absorbed. The internal backgrounds are reduced first by removing as much of the radioactive isotopes as possible from the target material, which can be done using a variety of techniques such as cryogenic distillation [60], chromatographic separation [61], or identification of sources of the

target material naturally reduced in the radioisotopes. However, the radioisotopes cannot be completely removed from the target materials, and the remaining internal ER backgrounds must be rejected by some form of active discrimination, usually accomplished by exploiting the fact that NRs are highly ionizing, while ERs are minimally ionizing. Therefore any observable that depends on the ionization density can be used as a handle to reject ERs.

The external γ backgrounds are usually due to radioactive contaminants such as uranium, thorium, cobalt, and potassium in the detector materials. Gamma rays can interact with the target material via Compton scattering and photoelectric absorption in the energy regime of interest for WIMP interactions, so like the internal backgrounds, the external γ backgrounds produce ERs. Dark matter experiments place considerable effort in constructing the detectors of very radiopure materials, but the decay rate of even the cleanest materials is well above the expected dark matter interaction rate. The external γ backgrounds are rejected via the same ER discrimination techniques mentioned for the internal backgrounds. In addition, γ rays are attenuated by the target material, so the external backgrounds can be reduced by fiducializing the target volume—considering only events from a central portion of the detector. Fiducialization requires event-by-event 3D position reconstruction, which will be discussed in the context of LAr TPCs in general in Sec. 2.4 and in DarkSide in particular in Sec. 3.4.

Surface α radiation comes from α -emitting isotopes, such as polonium in the decay chain of radon, adsorbed on the surfaces defining the detector volume. The α from a decay can be directed into the surface, causing the heavier daughter nucleus to recoil into the target material, mimicking the signal of a WIMP-induced recoil. Surface α backgrounds can be reduced by assembling the detector in a radon-free clean room environment. The remaining background can be rejected by performing event-by-event position reconstruction and excluding those events from near the surface, i.e. fiducialization.

Radiogenic neutrons generally come from spontaneous fission of heavy elements and (α, n) reactions of light elements. Therefore, the neutron emission rate is largely proportional to the contamination of uranium and thorium isotopes in the detector materials. Neutrons are dangerous because they produce NRs, exactly like a WIMP-induced recoil. The neutron

background is primarily reduced by constructing the detector from radiopure materials. The remaining neutron backgrounds can be reduced using a variety of techniques, which are discussed in Sec. 2.5.2.

The main cosmogenic backgrounds (backgrounds associated with cosmic rays) are due to cosmic ray muons and cosmogenically induced neutrons via spallation. These are mainly reduced by placing the detector deep underground, using the Earth overburden for shielding. At 3800 m water equivalent (m.w.e.) depth of LNGS, the cosmic ray muon flux is $(3.41 \pm 0.01) \times 10^{-4} \text{ m}^{-2} \text{ s}^{-1}$ compared to $180 \text{ m}^{-2} \text{ s}^{-1}$ at the Earth's surface [62, 63]. Furthermore, the overburden is extremely effective at reducing the low energy cosmic ray flux, and the muons that manage to penetrate through to the detector are high energy, producing far larger signals than those expected from WIMPs and are easily rejected. Cosmogenic neutrons are produced by spallation in the cosmic ray muon showers in the surrounding rock near the detector. Such neutrons can be greatly attenuated by shielding the detector with materials that have a large neutron interaction cross section [64].

2.2 Liquid argon

Liquid argon has a variety of attractive features as a WIMP dark matter direct detection target, particularly in terms of rejection of the electromagnetic backgrounds described above. For both NRs from neutron and WIMP scatters and ERs from β -decay and photon scatters and absorption, energy is deposited in the form of excitation and ionization. For ERs, these two channels account for nearly all the energy deposition. But for NRs, a significant fraction of the energy is deposited to other channels, such as atomic motion (heat) [65]. Excitation of Ar atoms leads to the production of scintillation photons (this process will be described in more detail below), and the total scintillation yield for NRs is observed to be quenched to 0.25 that of ERs [66, 67]. Due to the quenching, we often use separate units for the energy deposited via NRs and ERs. For ERs, we use keV_{ee} for electron equivalent energy, while for NRs, we use keV_{nr} for nuclear recoil equivalent energy ($1 \text{ keV}_{\text{nr}} \approx 0.25 \text{ keV}_{\text{ee}}$). Argon-based WIMP detection experiments focus on maximizing the efficiency to observe the excitation

and ionization channels, thus maximizing the ability to understand, control, and reject the dominant ER backgrounds.

The excitation and ionization processes for ERs and NRs are mostly understood for liquid argon (and noble liquids in general). The basic picture for a nuclear recoil is that a neutral particle (neutron or WIMP) elastically scatters off an Ar nucleus, and the recoiling nucleus deposits its energy in a track of excited and ionized atoms. In an electron recoil, on the other hand, an electron (typically either from a β -decay or from Compton scatter or photoelectric absorption of a γ -ray) scatters many times off many Ar atoms, also producing a track of excited and ionized atoms. Since the stopping power (dE/dx) for recoiling nuclei is higher than that for electrons, the density of a NR track is higher than the density of the ER track.

The ratio of excited atom production to Ar ion plus electron pair production has been calculated to be 0.21 [68]. The excited atoms form weakly bound dimers (excitons) [69, 70], which radiatively de-excite, producing VUV scintillation photons:

Some of the ionized atoms recombine with an electron and go through a non-radiative de-excitation followed by a radiative de-excitation, also producing VUV scintillation light:

Notice that the last two steps of the ion recombination process are identical to the scintillation process of the initially excited atoms. The scintillation light ($h\nu$) is narrowly peaked at 128 nm [71, 72]. Liquid argon produces about 40,000 VUV photons per MeV of energy deposited for ERs [68], making it a bright scintillator. Fortunately, LAr is transparent to its own scintillation light. Tens of thousands of ionization electrons may also be produced per

MeV of energy deposition, though the observed number of electrons varies significantly with field, energy, and radiation type. The scintillation and ionization channels can be efficiently observed with such detectors as the dual-phase TPC, described in Sec. 2.4.

The excited Ar dimer Ar_2^* can be in either the singlet state or triplet state, each with different decay times of 7 ns and 1.6 μs , respectively [73]. Therefore, we often call the scintillation light associated with the decay of the singlet (triplet) state, the fast (slow) component. The decay times do not depend on the ionization density, but the ratio of the singlet to triplet states does. For NRs, the singlet/triplet ratio (after recombination) is ~ 3 , while for ERs the ratio is ~ 0.3 , allowing for discrimination between NRs and ERs. The origin of the difference in singlet and triplet intensities between nuclear and electron recoils is not well established. Hitachi et al. [73] suggest that superelastic collisions between singlet state excitons and thermal electrons can produce triplet state excitons. Furthermore, the lower dE/dx of ERs can result in recombination taking place over more time compared to NRs, allowing for more of the singlet states to convert into triplets, which reduces the singlet/triplet ratio for ERs. Another possibility, described in [74], is that the ratio of singlets to triplets produced by direct excitation may be different than the ratio produced via recombination. If the singlet/triplet ratio is higher for excitons produced through recombination, then for NRs, with their higher dE/dx and greater recombination [75], the ratio will also be higher.

Since we must analyze the shape of the scintillation time profile to discriminate between NRs and ERs, we call this pulse shape discrimination (PSD). Discrimination power is defined as the number of events that can be rejected as ERs for every event that is incorrectly classified as a NR, given a particular nuclear recoil acceptance level. DarkSide-50 has demonstrated a discrimination power of 10^7 using PSD alone [76]. The observed ratio of ionization to scintillation (i.e. after recombination) is also different between NRs and ERs, leading to a second method of discrimination. The ionization to scintillation ratio enters little in the analyses presented in this work, so we do not discuss it further here.

The combination of discrimination from both scintillation pulse shape and the ratio of ionization to scintillation for LAr is to be compared with discrimination from only the ratio of ionization to scintillation in liquid xenon. Qualitatively, the scintillation and ionization

processes for LXe follow all the same properties as LAr, but they differ quantitatively. The singlet and, especially, triplet decay times are much shorter in LXe—4.3 ns and 22 ns, respectively [73]. The small difference in decay times renders pulse shape discrimination in LXe impractical. Liquid xenon based WIMP dark matter detectors rely primarily on the ionization to scintillation ratio for discrimination of NRs and ERs, and achieve a discrimination power of $\sim 10^3$ [48, 77, 78]. The extreme discrimination power of PSD in LAr is the main motivation for using LAr over LXe as a WIMP search target.

Liquid argon is attractive as a WIMP search target not only because of its spectacular discrimination power. The second main draw of the LAr technology (in fact, the liquid noble technology, in general) is that it is scalable to large target masses. Photosensors that can detect the scintillation light constitute a large part of the cost of noble liquid direct detection experiments, and photosensor coverage scales with surface area, while WIMP sensitivity scales with volume, so there is an economy of scale. Furthermore, the high voltage demands of large volume detectors are not extraordinary—the DarkSide-50 TPC uses not more than 13 kV on its cathode, which does not pose a technology challenge. Finally, the ease of purification (removal of electronegative impurities) allows for long drift times (>1 ms), which are necessary for ton-scale detectors.

Natural argon is also abundantly available, making it cost-effective to build large detectors. However, this argument is invalidated by the presence of ^{39}Ar in natural argon, which is typically obtained from the atmosphere. The majority of argon in the atmosphere is stable ^{40}Ar , produced by electron capture of ^{40}K [79]. There is very little potassium in the atmosphere, and the bulk of argon production occurs underground, from which some of the argon gas then diffuses out. However, in the atmosphere, cosmic ray spallation of ^{40}Ar produces ^{39}Ar , which undergoes β -decay with a Q-value of 565 keV and has a half-life of 269 years. Natural atmospheric argon contains a specific activity of ^{39}Ar of (1.01 ± 0.08) Bq/kg [80], which presents an overwhelming background for rare event search experiments using atmospheric argon. It is then necessary to obtain argon that is reduced in ^{39}Ar .

2.3 Underground Argon

There are several avenues one might pursue to obtain argon reduced in ^{39}Ar . One could try to remove ^{39}Ar from AAr using, for example, centrifuge techniques or thermal diffusion. However, it would be prohibitively expensive and time-consuming to use these techniques to produce the large quantities of argon (tens to hundreds to thousands of kilograms) demanded for a WIMP search.

As already mentioned, ^{40}Ar production occurs primarily underground. While some of the gas diffuses out to the atmosphere, much of the gas remains trapped underground. One might expect that underground argon (UAr) would be reduced in ^{39}Ar because it is shielded from cosmic ray activation. However, ^{39}Ar can be produced underground by neutron capture on ^{39}K , via $^{39}\text{K}(n,p)^{39}\text{Ar}$. The concentration of ^{39}Ar in underground argon then follows the local free neutron flux. Free neutrons can be produced underground both cosmogenically, due to interactions from high energy cosmic ray muons, and radiogenically from spontaneous fission and (α,n) reactions primarily from the decay chains of U and Th. The concentration of U and Th, and therefore also the concentration of ^{39}Ar in underground argon, varies from place to place within the Earth's crust, with measured ^{39}Ar levels ranging from 20 times lower to 16 times higher than the ^{39}Ar level in atmospheric argon [81, 82]. While ^{39}K leads to the production of ^{39}Ar , ^{40}K is needed for the production of ^{40}Ar . Therefore, underground argon with low concentrations of ^{39}Ar can only occur in locations with a low free neutron flux. One might then look for underground argon in locations of the Earth's crust with low U and Th concentrations. However, the gas can diffuse away from their points of origin. Therefore, one must look for sources of underground argon with reduced ^{39}Ar content only by directly measuring the gases.

In 2007, a source of underground argon with reduced ^{39}Ar content was identified in the National Helium Reserve in Amarillo, TX [83]. Subsequently, gases from the Reliant Dry Ice Plant in Bueyeros, NM and from Kinder Morgan CO₂ in Cortez, Colorado were also found to contain low radioactivity argon [84]. In 2010, a plant was installed at Kinder Morgan CO₂

to extract argon from the facility’s gas stream¹, which contained argon at the ~ 500 ppm level [84]. The product gas contained argon at the few percent level and was shipped to Fermilab, where it was further purified using cryogenic distillation to detector-grade argon with ppb levels of impurities [60]. The ^{39}Ar activity of the underground argon was measured in 2011 using a small detector at the KURF underground laboratory, which set an upper limit of 6.6 mBq/kg, or factor 150 reduction compared to ^{39}Ar activity in atmospheric argon. Precise measurement of the ^{39}Ar activity in underground argon was not established until its use in DarkSide-50. The measurement is described in Chapter 5.

2.4 Dual-phase LAr TPC

We now introduce the concept of the dual-phase (liquid-gas) time projection chamber (TPC), which we use to detect the LAr scintillation and ionization signals. In this section, we describe the general principles of a dual-phase LAr TPC, and in the next section we describe the DarkSide-50 TPC in detail. A cartoon of a LAr TPC is shown in Fig. 2.1. The active volume (the volume in which we can detect ionizing radiation) is a monolithic cylindrical body of liquid argon. Above the liquid volume is a thin gas layer. A uniform electric field is applied across the active volume such that ionization electrons that escape recombination are drifted up to the gas layer, where a stronger electric field extracts the electrons out of the liquid and drifts them across the gas layer. The stronger electric field in the gas region causes the electrons to excite (but not ionize) Ar atoms in the gas, which then de-excite and emit VUV scintillation light. This secondary light production is called electroluminescence.

The active volume is instrumented with photomultiplier tubes (PMTs) that observe both the primary scintillation signal, commonly referred to as S1, and the electroluminescence signal, commonly referred to as S2. The PMT faces are coated with a photocathode material, which absorbs visible photons and emits electrons by the photoelectric effect. Such an emitted electron is commonly referred to as a photoelectron (PE), which is cascaded along

¹Kinder Morgan CO₂ extracts gases from underground wells in Doe Canyon, Colorado to obtain CO₂ for use in the petroleum industry.

Figure 2.1: Cartoon of a dual-phase LAr TPC.

the PMT dynode chain. The amount of scintillation light produced is proportional to the energy deposited, so the size of S1, often given in units of PE, gives energy information. Likewise, the amount of electroluminescence light produced, and therefore the size of S2, also given in units of PE, are proportional to the amount of ionization produced. Therefore, S2/S1 allows for discrimination between NRs and ERs.

The cylindrical wall of the active volume is usually made from a highly reflective material to maximize the light collection efficiency of both S1 and S2. The top and bottom of the active volume are defined by two electrodes that form the anode and cathode surfaces that form the drift and extraction electric fields. A grid just below the liquid surface allows the drift and extraction fields to be set independently. Field shaping rings outside the TPC help to make the fields uniform. The inner surfaces of the active volume (the cylindrical wall and the top and bottom surfaces) are coated in wavelength shifter to shift the 128 nm Ar scintillation light to the visible (420 nm). All scintillation photons hit the wavelength shifter first before being detected by the PMTs.

The drift time of ionization electrons through the TPC is generally long (tens to hundreds of microseconds) compared to the time distribution of the S1 and S2 signals (up to a few microseconds). The timing between S1 and S2 gives z position information (the z -axis is commonly taken to point along the drift direction). The top PMTs are placed in a tightly packed array just above the gas region. The S2 signal is highly non-uniform on the top array and the hit pattern of S2 over the PMTs gives (x, y) position information about the primary interaction. In this way, we have 3D position information for each event.

PMTs have a fast response to photons. The FWHM of a single photon is ~ 10 ns for the PMTs in DarkSide. Therefore, the PMT response depicts the photon arrival time spectrum for single scintillation signals, which have characteristic times of several microseconds. We record the anode response vs. time—i.e. we record waveform data—for each PMT for every event. An event is defined as a collection of waveforms from all the PMTs, acquired simultaneously in a window spanning several hundred microseconds, long enough to include the S1 and S2 signals. Recall that the characteristic times for the fast and slow component of S1 are 7 ns and 1.6 μ s. Though we may not have the resolution to reconstruct the fast

Figure 2.2: Example raw waveforms for typical ER and NR events of about the same energy. In both events, the S1 is at $t = -6 \mu\text{s}$ and the S2 is at $t = 55 \mu\text{s}$. (a) Waveform with S1 and S2 for NR. (b) Waveform with S1 and S2 for ER. (c) Zoom to S1 region of NR event. (d) Zoom to S1 region of ER event.

component lifetime, we need only estimate the fast to slow ratio to perform pulse shape discrimination. Figure 2.2 shows example waveforms for a prototypical NR event and a prototypical ER event. The two events were chosen to have approximately the same S1 size, and therefore illustrates, qualitatively, the two discrimination methods, S2/S1 and PSD: the size of the NR S2 is significantly smaller than that of the ER S2, and the tail of the NR S1 is significantly reduced compared to that of the ER S1.

Figure 2.3: Conceptual drawing of the DarkSide-50 system.

2.5 DarkSide-50

The DarkSide-50 experiment is a phased series of dual-phase argon TPCs. The first detector was DarkSide-10, built at Princeton University and moved to LNGS. Its purpose was to develop the dual-phase LAr TPC technology in preparation for large-scale detectors. It successfully demonstrated that we could hold ~ 20 kV of HV at the cathode, and produced an exceptionally high light yield of 9 PE/keV_{ee} at null field [85].

DarkSide-50 is the first physics-capable detector in the DarkSide program. The DarkSide-50 apparatus consists of three nested detectors, as sketched in Fig. 2.3. From the center outward, the three detectors are: the dual-phase argon TPC, which is the dark matter detector; the Liquid Scintillator Veto (LSV), serving as shielding and as anti-coincidence for radiogenic and cosmogenic neutrons, gamma rays, and cosmic muons; and the Water Cherenkov Detector (WCD), serving as shielding and as anti-coincidence for cosmic muons. The detector system is located underground in LNGS at a depth of 3800 m.w.e., where the

cosmogenic neutron flux has been measured by Borexino to be $\sim 1 \text{ m}^{-2} \text{ s}^{-1}$ [62]. DarkSide-50 is located in Hall C, next to Borexino.

The DarkSide-50 TPC was first assembled in April 2013. The goals of the first deployment included evaluation of new R11065-20 PMTs from Hamamatsu, which had significantly lower radioactivity and commissioning of the TPC assembly procedure. However, the new PMTs were found to be non-functional at low temperature and were replaced with functional R11065-10 PMTs. Re-assembly of the TPC was completed in August 2013 and was filled with atmospheric argon in September 2013. The trigger rate was dominated by the large ^{39}Ar activity in the AAr and provided a high statistics sample on which to tune the present analyses. A first dark matter search was performed with atmospheric argon with data taken between November 2013 and May 2014 [76]. The TPC was emptied and refilled with underground argon in March 2015.

2.5.1 TPC

Figure 2.4 shows a sketch of the DarkSide-50 LAr TPC. The active volume is a 36 cm by 36 cm cylinder. The cylinder wall is a monolithic 1" thick PTFE (Teflon) machined piece. The top and bottom of the active volume are defined by fused silica windows. The active volume shrinks by $\sim 1.5\%$ from room temperature to LAr temperature (89 K) due to shrinkage of PTFE at low temperatures. The active volume is viewed by 38 PMTs, arranged in two hexagonal arrays, 19 on the top and 19 on the bottom, and held in place by two PTFE structures at the top and bottom of the TPC. The PTFE structures and the PMT photocathodes directly face the fused silica windows.

The TPC is placed inside a stainless steel cryostat, described below. We now describe various components of the DarkSide-50 TPC and supporting subsystems.

PMTs The PMTs are Hamamatsu model R11065 with 3" diameter photocathode. They are made from low radioactivity materials and have high quantum efficiency (QE) with a special bialkali photocathode designed for operation at liquid argon cryogenic temperature.

Figure 2.4: Conceptual drawing the DarkSide-50 TPC.

The average QE of the PMTs at room temperature is 34% at 420 nm. The PMTs are maintained at negative bias, meaning the PMT photocathode is at -HV and the PMT anode is at ground. Due to undesired light emission effects at low temperature when operating the R11065 PMTs at full bias (-1500 V), we operate the PMTs at reduced bias (about -1200 V) and thus reduced gain at the PMT anode. This required the use of a cryogenic amplifier on each PMT, described later in this section. In order to maintain maximum collection efficiency of the photoelectrons from the PMT photocathode to the first dynode, we use a modified voltage divider configuration on the dynodes.

Diving bell The top window has a cylindrical rim that extends downward by 1 cm to form a diving bell that holds the gas layer of the TPC. After liquid argon is filled past the top PMT array, the gas layer is produced by boiling argon within the cryostat (outside of the TPC active volume) and delivering the gas to the diving bell. The gas then exits the diving bell via a bubbler that maintains the liquid level at the desired height.

TPB All surfaces of the active volume (PTFE wall and fused silica windows) are coated with tetraphenyl-butadiene (TPB), a wavelength shifter. The TPB absorbs the 128 nm scintillation light from LAr and re-emits visible photons, peaked at 420 nm wavelength. The TPB is evaporated onto the PTFE and fused silica surfaces at LNGS prior to assembly. The thickness of the TPB layer on the fused silica windows varies from (230 ± 10) $\mu\text{g}/\text{cm}^2$ at the center to (190 ± 15) $\mu\text{g}/\text{cm}^2$ at the edge of the active volume. The thickness of the TPB on the PTFE wall is (165 ± 20) $\mu\text{g}/\text{cm}^2$ at the center and (224 ± 27) $\mu\text{g}/\text{cm}^2$ at the top and bottom.

Field cage Each side of the fused silica windows is coated with a 15 nm thick layer of indium tin oxide (ITO), a transparent conductor with broad industrial applications such as LCD monitors and thin film photovoltaics. The outer surfaces of the windows (which face the PMTs) are set at the average PMT photocathode potential, approximately -1200 V. The inner surfaces of the windows form the anode and cathode surfaces. The bottom ITO

layer on the top window serves as the grounded anode surface and the top ITO layer on the bottom window serves as the -HV cathode. A 50 μm thick stainless steel grid with hexagonal mesh with 2 mm pitch sits ~ 5 mm below the liquid surface. The mesh has a 95 % optical transparency at normal incidence. Negative high voltage is applied between the grid and cathode to produce a vertical electric field to drift the ionization electrons upward. Voltage is also applied between the anode and grid to extract the drift electrons out of the liquid and produce the electroluminescence signal in the gas layer. The voltages on the grid and cathode surfaces can be tuned to set the drift and extraction electric fields independently. The nominal operating voltages for DarkSide-50 are -12.7 kV on the cathode and -5.6 kV on the grid, giving a drift field of 200 V/cm, extraction field of 2.8 kV/cm, and electroluminescence field of 4.2 kV/cm. (Liquid argon has a dielectric constant of 1.5, and gas argon has a dielectric constant of approximately 1.) Outside the cylindrical PTFE wall, copper rings at graded potentials keep the drift field uniform throughout the active volume. The relative voltages on the individual rings are determined by resistors connecting each adjacent ring, forming a large voltage divider chain. Figure 2.5 shows a simulation of the electric field in DarkSide-50. The graded potentials of the field cage rings were chosen to produce a uniform field with 1000 V/cm drift field and 2.8 kV/cm extraction field. With the 200 V/cm drift field, the field cage is no longer optimized and the stream lines are “pushed” inward at the top of the TPC due to leakage of the extraction field through the grid. However, due to charge buildup on the PTFE walls, the drift field is likely uniform throughout.

The grid and cathode voltages are transferred from Bertan 225-20R and 225-50R DC high voltage power supplies in the cleanroom above the WCD to inside the detector via custom made feedthroughs. The basic design of each HV feedthrough is a stainless steel conductor press fitted into a UHMW polyethylene tube, which is surrounded by a grounded stainless steel shield tube. The feedthrough must be able to hold HV while holding ultra high vacuum. The leak tight seal between the polyethylene and outer ground shield is achieved by a “cryofit” procedure: the outer diameter of the UHMW polyethylene tube is machined to be 1.5 % larger than the inner diameter of the stainless steel ground shield tube. The

Figure 2.5: Electric field simulation of the DarkSide-50 TPC. Uses an axisymmetric model of the field cage. The colored lines represent equipotentials, and the black lines through the TPC represent stream lines along which the ionization electrons drift.

polyethylene is then cooled in a liquid nitrogen bath which causes the material to shrink by 2% and can then be fit into the stainless steel tube while cold. As the polyethylene returns to room temperature and expands to its machined diameter, it forms a vacuum tight seal against the inner surface of the stainless steel. The grounded tube extends from the cryostat flange through the gaseous argon to below the liquid surface. The ground tube keeps the electric fields due to the HV conductor well contained and prevents charge buildup on any surfaces in the gaseous argon, which has a low breakdown point. In the liquid, which has a high breakdown point (hundreds of kV/cm [86]), we can tolerate much higher electric fields and do not need the ground shield to provide such stringent protection for the HV conductor.

Cryostat The TPC is housed inside a double-walled stainless steel cryostat with inner cryostat inner diameter of 50 cm and outer cryostat outer diameter of 64 cm. The volume between the cryostat walls contains a multi-layer mylar insulation blanket and is constantly vacuum pumped to $<10^{-6}$ mbar. The vacuum plus mylar insulation holds the temperature at 89 K inside the cryostat while the outside of the cryostat is at room temperature. The entire cryostat hangs from rods that extend through the LSV to the top of the WCD. The rods are threaded, allowing for precise leveling adjustments of the TPC.

Cryogenics Cooling of the cryostat is done using an external circulation loop. Argon gas is drawn from the cryostat at 30 slpm through the same line dedicated to nearly all the electronics cables. These large surface areas are expected to be the dominant source of impurities in the Ar due to outgassing. The gas is therefore passed out of the detector system, directly to the cryogenic and purification system, located in a radon-suppressed clean room above the WCD. In this way, we reduce the contamination of the “dirty” gas into the “clean” gas and liquid in the cryostat. The gas passes through a SAES Monotorr PS4-MT50-R-2 getter, which removes contaminants such as O₂ and N₂ to sub-ppb levels. The gas is then pre-cooled in a heat exchanger before passing through a cold-charcoal radon trap that is operated in the range 185 K to 190 K. The argon is then liquefied by a liquid-nitrogen-cooled heat exchanger. The liquid argon is then delivered directly to the active

volume of the TPC via a vacuum jacketed stainless steel transfer line. The loop cooling power is controlled to maintain a stable pressure in the cryostat. The pressure oscillates within a band of ± 0.1 mbar around the set point of 1080 mbar.

Electronics and data acquisition While the PMTs are located outside the sensitive volume of the TPC, they are all immersed in LAr and are operated at 89 K. Based on previous experience in DarkSide-10 and extensive laboratory tests, it was found that charge accumulation on internal components of the PMTs induced light emission, rendering the PMTs useless. This is mitigated by decreasing the bias voltage, reducing the PMT gain to $\sim 4 \times 10^5$. The reduced gain necessitates a local pre-amplifier operating at LAr temperature to drive signals through the ~ 10 m of cable to the amplifiers and digitizers in the clean room above the WCD. The cryogenic pre-amplifier is mounted directly on the PMT base. The pre-amplifier has 150 MHz bandwidth and provides a factor 24 voltage gain. The pre-amplifier component materials are selected for low-radioactivity.

At room temperature, inside the clean room, the signals are split into several channels. One copy is amplified $\times 10$ and split again, one to a high speed discriminator, set to 0.6 PE and used to form the TPC trigger and another to a 12-bit 250 MHz digitizer channel (CAEN V1720). Another un-amplified copy is sent directly to a CAEN V1724 14-bit 100 MHz digitizer. The use of two digitizer types extends the dynamic range, providing a linear response between 1 PE to 10 000 PE. The multiple ADCs are synchronized with a common 50 MHz clock and external trigger.

The TPC is triggered via a majority trigger, requiring a preset number of channel discriminators to fire within a 100 ns window. In this way, DarkSide-50 efficiently triggers on S1. During the AAr campaign, a majority threshold of 3 was used, while in the UAr campaign, the threshold was changed to a majority 2. In the AAr campaign, the high rate of ^{39}Ar decays presented a demand for large data acquisition and throughput rates. The trigger rate was ~ 50 Hz. In order to ease the data throughput requirements, parts of the data were pre-scaled using the so-called *G2 trigger*. Events with S1 above ~ 700 PE were suppressed in the data acquisition stage by a factor 33. In the UAr campaign, the trigger

rate was ~ 1.5 Hz, and the G2 trigger was not necessary.

Upon receipt of a trigger, the DAQ records a $440\ \mu\text{s}$ gate of waveform data for each of the 38 channels. The trigger is also inhibited from re-firing for $810\ \mu\text{s}$ (the inhibit window) to prevent re-triggering on the residuals of the signal. The raw data is then passed to the reconstruction software, described in Ch. 3.

Calibration system Calibration of the TPC can be performed using a variety of sources. The sources are placed on an articulated arm that extends from the clean room above the WCD, through an “organ pipe” that extends to the LSV stainless steel sphere, through the liquid scintillator, and parks next to the TPC cryostat. Calibration can be performed with neutron sources (AmBe and AmC) and gamma sources (^{57}Co , ^{133}Ba , ^{137}Cs , and ^{22}Na).

DarkSide-50 can also perform $^{83\text{m}}\text{Kr}$ calibration. A sample of ^{83}Rb is placed on a charcoal substrate which is placed within the gas recirculation system and can be isolated by UHV valves. Rubidium-83 decays to $^{83\text{m}}\text{Kr}$ with 86 d half-life (with branching ratio 75 %). Metastable $^{83\text{m}}\text{Kr}$ de-excites, emitting 32.1 keV and 9.4 keV gamma rays in quick succession, with a half-life of 1.83 hrs. When the source gas is exposed to the gas recirculation system, $^{83\text{m}}\text{Kr}$ mixes with the argon and enters the TPC. In this way, one achieves a uniform calibration source throughout the TPC. The $^{83\text{m}}\text{Kr}$ calibration was used to measure the 8 PE/keV_{ee} light yield at null field and 7 PE/keV_{ee} light yield at 200 V/cm drift field in DarkSide-50.

2.5.2 Outer Detectors

2.5.2.1 The need for active shielding

Perhaps the most troublesome backgrounds for DarkSide-50 are nuclear recoils from single neutron scatters because they are indistinguishable from the expected WIMP interactions in the TPC. PSD and S2/S1 do not help, and the moderate size of DarkSide-50 does not allow a fiducial volume that is sufficiently shielded from neutron-induced backgrounds. Therefore, great care must be taken to reduce the rate of neutron-induced backgrounds in the TPC.

A simple solution is to use passive shielding around the detector, but this leaves several challenges. With only passive shielding, it is difficult to conclusively measure the background levels, making the interpretation of a few observed recoil events as a WIMP signal problematic. Furthermore, passive shielding does not protect against radiation from the detector components themselves, and does not protect against high energy (often ~ 100 MeV) cosmogenic neutrons, which may penetrate deeply through most shielding.

A better method of neutron suppression is the use of an instrumented—or “active”—neutron detector (the neutron veto). If the WIMP detector is surrounded by a volume of material with a high neutron interaction cross section, neutron-induced NRs in the TPC will, with high probability, be coincident with a neutron-induced signal in the veto. In contrast, WIMP-induced NRs in the TPC will not be accompanied by any signals in the veto, i.e. will be in anti-coincidence. In addition to removing neutrons, an active veto also provides *in situ* measurements of the neutron background in the experiment. In DarkSide-50 we actively veto radiogenic and cosmogenic neutron backgrounds with a pair of vetoes: a boron-loaded liquid scintillator veto (LSV) and a water-Cherenkov detector (WCD). The veto systems are described in detail in Ref. [87]. We summarize them here.

2.5.2.2 LSV

The LSV is a 4.0 m diameter stainless steel sphere filled with 30 tonnes of boron-loaded liquid scintillator. The sphere is lined with Lumirror, a reflecting foil used to enhance the light collection efficiency. An array of 110 Hamamatsu R5912 LRI 8” PMTs is mounted on the inside surface of the sphere to detect scintillation photons. Figure 2.6 shows a view of the inside of the LSV. The boron-loaded liquid scintillator has three primary components: pseudocumene (PC), trimethyl borate (TMB), and 2,5-diphenyloxazole (PPO). The liquid scintillator cocktail is composed of 95 % PC by mass, 5 % TMB, and 1.4 g/L PPO.² With

²The 5 % TMB concentration is the present configuration. The initial configuration, used in the first WIMP search campaign with atmospheric argon [76], used a different liquid scintillator cocktail with 50/50 PC and TMB by mass and 2.5 g/L PPO. However, it was found that the LSV exhibited a high rate of ^{14}C (~ 150 kBq). The endpoint of the ^{14}C β spectrum is 156 keV, making α and ^7Li decay products impossible to distinguish over the background and limiting the rejection power of the veto. It was discovered that some of the TMB feedstock was derived from modern carbon, which has a much higher ^{14}C content than

Figure 2.6: View of the inside of the LSV. The metal can in the center is the TPC cryostat. The PMTs and Lumirror reflectors line the wall.

this TMB concentration, the neutron capture time is $\sim 22 \mu\text{s}$.

In order to have a high neutron vetoing efficiency it is important to choose a scintillator cocktail that

- has a large neutron capture cross section
- has a high scintillation yield
- is transparent to its own scintillation light
- the decay time of the scintillator is short compared to light collection time
- the emission spectrum of the scintillator is peaked around where the PMTs are most sensitive and the reflector most reflective

PC is the primary scintillator used in the veto and makes up the bulk of the cocktail. TMB is an organic molecule containing one boron atom. Boron-10, with a natural abundance of 20 %, has a very high thermal neutron capture cross section of 3840 b. When a neutron captures on ^{10}B , two reactions are possible:

The decay to the excited state produces a γ -ray that is easily seen as long as it does not escape into the cryostat before depositing energy into the scintillator. Energy deposits due to the α and ^7Li nucleus, on the other hand, are always contained in the scintillator, due to their high stopping power and consequently short track length. This gives boron a comparative advantage over other loading options such as gadolinium, which only produces high energy γ -rays that may escape the veto without leaving a detectable signal. However, the light output of α and ^7Li nuclei is highly suppressed due to ionization quenching, causing them to

petroleum-derived material. A new batch of petroleum-derived TMB was identified and installed in the LSV in Jan. 2015. As a result the ^{14}C activity was reduced to $\sim 0.3 \text{ kBq}$. The use of the lower TMB concentration was for cost reasons.

scintillate at 50 keV_{ee} to 60 keV_{ee}. Detecting these decay products therefore requires a high light collection efficiency. If the detector can reliably observe these nuclear decay products, it can efficiently detect neutrons that capture in the veto, regardless of their initial kinetic energy, since the energy of these captures products does not depend on the neutron’s initial energy. The measured LSV light yield is (0.54 ± 0.04) PE/keV, making the quenched energy readily detectable.

2.5.2.3 WCD

The WCD is an 11 m diameter, 10 m height cylindrical tank filled with high purity water. The tank was originally part of the Borexino Counting Test Facility [88]. The inside surface of the tank is covered with a laminated Tyvek-polyethylene-Tyvek reflector [89]. An array of 80 ETL 9351 8” PMTs, with 27 % average quantum efficiency at 420 nm, is mounted on the side and bottom of the water tank to detect Cherenkov photons produced by muons or other relativistic particles traversing the water. Figure 2.7 shows an inside view of the WCD.

The WCD is a powerful shield against external background (gamma-rays and neutrons from the surrounding rock), and is also used as a muon detector. The $\sim 1.1 \text{ m}^{-2}\text{h}^{-1}$ muon flux in the LNGS experimental hall corresponds to 2000 muons per day crossing the WCD, 380 muons per day crossing the LSV, and 4 muons per day crossing the TPC. Cosmogenic muons can produce high energy neutrons, which can penetrate several meters of shielding. In order to avoid backgrounds from these high energy neutrons, the WCD acts as a veto to detect the muons that may produce them and therefore leave a detectable coincident signal.

2.5.2.4 Electronics and data acquisition

The anode signals from the 110 PMTs in the LSV and the 80 PMTs in the WCD are amplified and split by means of a custom front-end board. A $\times 10$ amplified signal is sent to 190 channels of National Instruments PXIe-5162 digitizers which sample at 1.25 GHz with a 10 bit resolution. Zero-suppression is performed on the fly and only sections of the waveform around identified peaks above threshold are stored. The zero-suppression threshold is set to

Figure 2.7: View of the inside of the WCD. The stainless steel sphere of the LSV is shrouded in Tyvek reflector at the center. PMTs are mounted on the wall and floor of the WCD.

a level about 0.25 times the amplitude of a single photoelectron pulse for routine data taking. For WIMP search data taking, the trigger for the veto system is given by the TPC trigger, i.e. we use the TPC as a global trigger. To keep the TPC and veto readouts aligned, a pulse per second (PPS) generated by a GPS receiver is sent to the two systems, where it is acquired and interpolated with a resolution of 20 ns to allow offline confirmation of event matching. Due to the potentially long capture time of neutrons in the LSV, the veto acquisition window is set to as long as 200 μ s.

CHAPTER 3

Reconstruction

The DarkSide data consists of basic units called *events*. Each trigger of the TPC corresponds to one event. At the raw data level, each event consists of a collection of raw waveforms, one for each PMT, or channel. For the TPC, there are 38 channels, each digitized with 4 ns sampling on the time axis and 12-bit resolution on the amplitude axis. Figure 3.1 shows an example raw waveform with a typical S1 and S2 signal. For the LSV and WCD, there are 110 and 80 channels, respectively, each digitized with 0.8 ns sampling and 10-bit resolution. The first task for analysis of DarkSide data is to translate raw waveform data for each event into analysis variables, which we use for statistical analysis of the events.

There are two main data reduction steps to go from raw waveform data to physical variables: DarkArt and SLAD (SLim Analysis Data). DarkArt is the first pass of analysis over the raw data and is often referred to as the *reconstruction*. All analyses that require access to waveform data are performed in DarkArt. SLAD performs a further reduction on the DarkArt output files and provides user-friendly plain ROOT files. This chapter describes the algorithms used in the DarkArt and SLAD software and the evaluation of their performance, including efficiencies and biases.

3.1 DarkArt - TPC

The main goal of the DarkArt reconstruction software is to identify scintillation signals in raw data and evaluate a variety of parameters for each. Because every interaction (primary and secondary) typically illuminates every channel to some degree (though not necessarily equally), it is not so straightforward to localize signals in space. The goal of the recon-

Figure 3.1: Example raw waveform for a single channel in DarkSide-50 with a typical S1 and S2. The pulse at $t = 0$ is S1, and the pulse at $t = 250 \mu\text{s}$ is S2.

struction is to localize signals in time. Therefore, we typically consider information from all channels together or from individual channels, but we make no attempt to cluster subsets of channels together.

3.1.1 Technical overview

Before describing the reconstruction algorithms, we give a brief overview of the technical aspects of DarkArt. The DarkArt software is implemented in modular form within the *art* framework [90], which is a Fermilab-developed event-processing framework intended for use by small- to medium- scale particle physics experiments. For each event, there exists a single C++ class called an `art::Event`, which is a generic container to hold *data products*, which are globs (C++ classes) of information that are passed from one *module* to another. Modules are blocks of code (also C++ classes) which house the experiment-specific algorithms that operate on the data products. Since DarkSide-50 raw data are written by DS50DAQ, which is also implemented within the *art* framework, reading of raw data into DarkArt is nearly transparent (Just Works). The output of DarkArt is a custom-made super-product (in

practice, a large C++ class) for each event. For very simple scanning of the variables, DarkArt outputs can be accessed with standalone ROOT. For more complicated data access and in compiled code, the DarkArt headers need to be linked; this is where SLAD enters. The DarkArt source code is publicly available online [91].

3.1.2 Modules

In broad steps, the components of the DarkArt reconstruction algorithms are: identify and remove the baseline offset from each channel, build a sum channel, identify pulses within the sum channel, and compute parameters for each pulse. Generally speaking, each component of the chain is housed in one module. In this section, we describe the main algorithms used in the reconstruction.

3.1.2.1 Baseline finder

The first step of the reconstruction chain is to identify and remove the baseline offset for each channel. PMT signals are negative deviation from baseline noise, so a DC offset is applied such that the baseline noise is near the top of the dynamic range of each channel. The goal of the baseline-finder is to identify the offset, or the *baseline*. The baseline often has time-dependent fluctuations due to electronics noise at a variety of frequencies, so we cannot use a single fixed baseline value event-to-event or even within an event. Even if the baseline had no low frequency components (apart from a DC offset), use of a fixed baseline would introduce large biases to cumulative operations on waveforms (i.e. integrals), which is how all energy information is evaluated. Therefore, we use a boxcar algorithm, or moving average, to track the low frequency variations of the baseline. (“Low”, in this case, is relative to the 250 MHz sampling rate.) We often refer to this as a drifting baseline algorithm, or moving baseline algorithm.

We cannot indiscriminately evaluate a boxcar average everywhere: in the presence of signal in the waveform, the average will be heavily biased downward. The challenge for the drifting baseline algorithm is that *a priori*, we do not know whether any given part of a

waveform is baseline or signal. The baseline finder introduces the first assessment of signal vs. baseline. In the presence of signal, we do not know the exact behavior, so we use a linear interpolation to approximate the baseline. The parameters of the algorithm are tuned such that it is sensitive enough to interpolate under single PE.

Algorithm Step through consecutive samples of the raw waveform, and at each sample determine whether it is in baseline or in signal. If the sample is baseline, the baseline value is defined as the average of the surrounding samples of the raw waveform. In DarkSide-50, we use the surrounding 21 samples (80 ns). If the sample is signal, the baseline value is defined as the linear interpolation between the two adjacent non-interpolated baseline values on either side of the sample.

How do we determine whether a sample is baseline or signal? There are three scenarios by which a sample can be deemed baseline (if a sample is not baseline, it is signal):

At the very start of the waveform. At the very start, we use an algorithm grandfathered in from DarkSide-10. We look for the highest point of the raw waveform in the pre-trigger region and then check that all samples of the first boxcar window are close to that maximum value; this step should be robust against any PE signal in the pre-trigger window since PE signals are only negative fluctuation from baseline.

When the previous sample is baseline. When the previous sample is baseline, the current sample is deemed in baseline when the raw waveform is within $\pm \text{max_amplitude}$ of the previous baseline value. The `max\_amplitude` parameter can be tuned individually for each channel. In DarkSide-50, the average value of the parameter is 8.5 counts and varies by ± 1 .

When the previous sample is signal. When the previous sample is signal, the current sample is deemed in baseline if the raw waveform is within $\pm \text{end_max_amplitude}$ of the previous non-interpolated baseline value. Note that the previous non-interpolated baseline value is, in this case, associated with the sample at least 2 prior, since the previous sample is signal. The value of `end\_max\_amplitude` is 3 counts, smaller than

`max_amplitude` to reduce the bias in the baseline reconstruction, which corresponds to a reduction in the bias of the reconstructed signal size.

The `max_amplitude` threshold is tuned such that samples associated with single PE are classified as signal, and the baseline is interpolated under the single PE spike, as shown in Fig. 3.2a. For larger signals, like S2, the algorithm interpolates the baseline for the entire signal region, as shown in Fig. 3.2b. The algorithm is also configurable to force interpolations over user-defined windows. This feature is useful for single photoelectron calibration.

Once the baseline is defined for all samples in the waveform, we build a baseline-subtracted waveform by subtracting the assumed baseline value from the raw waveform value at each point. All downstream reconstruction steps operate on the baseline-subtracted waveform.

Limitations There are two main limitations to the baseline finding algorithm. First, the algorithm makes only a primitive estimate of the baseline under signal regions. We know that the baseline can fluctuate or drift on the timescale of microseconds, and if a signal lasts a long time, there is no compensation for such behavior. Second, if there is any positive fluctuation of the raw waveform in the pre-trigger region, the baseline finder algorithm is very likely to abort, even if the baseline is otherwise easily identifiable. In the UAr data, a non-negligible fraction of triggers are due to bi-polar noise spikes. Fortunately, this noise is uncorrelated with scintillation signals, and the accidental rate of scintillation with bi-polar noise in the pre-trigger region is negligible.

Furthermore, the parameters require manual tuning. This could introduce a large bias in the reconstructed of baseline. Fortunately, we have estimated the bias to be small, as described in Sec. 3.8. Finally, the algorithm parameters have been assumed to be independent of run. That is, we assume the electronics noise and gain have been relatively stable over time. For the data used in this work, the noise and gain fluctuations have been sufficiently stable that the same baseline-finder parameters can be used throughout.

(a)

(b)

Figure 3.2: Example waveforms showing the baseline (red) overlaid on raw waveform data (black). The green points indicate the endpoints of interpolated regions of the baseline. (a) The baseline is interpolated over single PE signals. (b) The baseline is interpolated over sustained signals such as S2.

Figure 3.3: Sum of baseline-subtracted waveforms without zero-suppression, showing coherent noise across channels.

3.1.2.2 Zero suppression and sum channel

This module creates a virtual sum channel, which stores the normalized sum of zero-suppressed baseline-subtracted waveforms from the real channels. The reason for zero-suppressing the (baseline-subtracted) waveforms of the real channels is that there is baseline noise that is coherent across all the channels, as can be seen in Fig. 3.3, which shows the sum of the baseline-subtracted waveforms without zero-suppression. Were zero-suppressed waveforms not used, single photoelectron pulses can be easily buried in the coherent baseline noise, making the found pulse start and end times less precise and less reliable. Note that this sum channel is used only for pulse-finding and not for computing any reconstruction variables besides pulse start and end times.

Algorithm We first scale the baseline-subtracted waveform of each channel by the measured single photoelectron (SPE) mean response for that channel. Sec. 3.3 describes the determination of the SPE mean response. We then set to zero all portions of each scaled waveform that are above a threshold value of 0.1 PE. Recall that signals in these wave-

Figure 3.4: Example sum channel waveform composed of the sum of zero-suppressed baseline-subtracted waveforms. The non-zero-suppressed portions are due to photoelectrons from an S1 signal.

forms are negative, i.e. below baseline. Then we add together the baseline-subtracted zero-suppressed waveforms. Fig. 3.4 shows an example.

3.1.2.3 Pulse finder

Given the summed waveform, the next step in the reconstruction is to identify clusters of PE, which we refer to as *pulses*, that constitute a scintillation signal. The pulse-finder is in many ways similar to the baseline-finder in that both algorithms attempt to distinguish baseline from signal. The main difference is that the baseline-finder is sensitive to individual PE, whereas the pulse-finding algorithm identifies macroscopic clusters of PE.

The timing of the pulses within any event is not always predictable. Typically, we trigger on the S1, in which case we know where to expect the S1. However, because events can occur throughout the TPC, S2 can appear any time within the acquisition window of an event. We also sometimes trigger on other types of signals. Given the variety of pulses and topologies for them to appear, the pulse-finding algorithm does not make any assumptions

about the timing or time-ordering of pulses. Furthermore, the algorithm does not make any assumptions about the relative amplitude, duration, or shape of the pulses.

The goal of the pulse-finder algorithm is to localize clusters of PE in time. We look for macroscopic changes in behavior, macroscopic in this case meaning tens to hundreds to thousands of PE clustered to within tens of microseconds. We use the sum channel to maximize the information from all channels. Though the pulse-finder is generic, it is tuned to be efficient at finding both S1 and S2. The precise efficiency is a function of the shape and amplitude of the pulse. We discuss the evaluation of the efficiency in Sec. 3.8. The pulse-finder is tuned conservatively, so that while we efficiently find S1s and S2s, we also find many other types of pulses. This conservative approach proves to be useful when investigating non-standard event topologies (anything other than plain S1 plus S2).

The most important aspect of the pulse-finder is to identify every pulse start precisely and accurately. The pulse start times tell us the electron drift time in the TPC, which in turn gives us z position information. We do not care particularly much about the precision of the pulse end because S1 and S2 signals have exponential fall-off. In any case, we use fixed length integrals for most energy estimators, so we ignore pulse end information in most downstream analysis.

Algorithm The main principle of the pulse-finder algorithm is to identify clusters of PE arriving within some short time window. The specific values of the minimum number of PE to be deemed a cluster and the size of the time window are chosen to optimize the efficiency to find S1 and S2 signals. The algorithm uses a two-step iterative approach.

1. We first do a coarse-grained search over the waveform to identify the presence of some sort of pulse. If, over the course of $2\ \mu\text{s}$ (500 samples), the waveform changes by 5 PE, we know that there is some sort of pulse present.
2. We then do a fine-grained search within the $2\ \mu\text{s}$ window to find precisely where the pulse starts. The pulse start threshold is 0.3 PE, small enough that the pulse start coincides with the first PE that appears in the $2\ \mu\text{s}$ window.

3. Once we have established the presence of a pulse, we then determine the pulse end. We first roughly determine where the pulse maximum is by looking for a change in slope of the coarse-grained waveform. Once past the maximum, we check for one of two end conditions.

Presence of pileup. The condition for the presence of pileup is similar to the initial search for the presence of a pulse: we look for a coarse-grained change in the waveform—this time with a larger threshold value of 35 PE—and then perform a fine-grained search for the start time of the pile-up pulse.

Waveform returns to baseline. We look for an absence of signal in a long integration window of 15 μs .

Limitations For some specific topologies of signals that occur very near each other within the waveform, the pulse-finder has difficulty accurately identifying the start time of the second pulse. For example, a multiply scattering gamma can produce two or more S2s that appear very near each other in time. The start time of the first S2 to arrive will be well reconstructed, but the start time of the second S2 will be obscured and therefore carry a greater uncertainty. For a WIMP search, this does not matter because we throw out all events in which two pulses are found so close together. For other analyses, special care must be taken to handle such “pile-up” cases.

Another limitation is that it is not often clear where to define the end of a pulse. The exponential fall-off of the S1 and S2 pulse shapes means the signal falls gently back into baseline, and there is no clearly defined endpoint of the pulse. Uncertainty in the pulse end time is mitigated by the use of fixed length integration windows for pulse area estimators.

3.1.2.4 Parameter evaluation

Once the pulses are identified for each event, we evaluate a variety of parameters for each pulse. We evaluate parameters at both the channel level and the event level. The event-level parameters are typically the sum of the channel-level parameters. For example, the S1 of

an event is given by the sum of the S1 integral over all the channels. We might have used the parameters evaluated using the sum channel to obtain the event-level parameters, but in order to minimize effects due to coherent noise across channels, we use sums of channel-level integrals. We highlight a few of the most important variables here. These variables are evaluated for every pulse.

- `fixed_int1`: 7 μ s integral from pulse start
- `fixed_int2`: 30 μ s integral from pulse start
- `npe`: integral from pulse start to pulse end
- `f90`: ratio of integral of first 90 ns of pulse to `npe`
- `f90_fixed`: ratio of integral of first 90 ns of pulse to `fixed_int1`

3.2 DarkArt - OD

Like the reconstruction of raw data from the TPC, the reconstruction of raw PMT waveform data from the LSV and WCD detectors is done in DarkArt. The analog signals from the LSV and WCD PMTs are treated identically to each other and independently of the TPC signals.

For the analyses presented in this work, the data are taken in slave mode: the trigger of the veto is initiated by the TPC. When the trigger is received, the veto digital electronics record data in an acquisition window with a width of several neutron capture times in the LSV. The veto data reconstruction is detailed in Ref. [87] and briefly summarized here.

Recall that, in contrast to the TPC data, the LSV and WCD raw data are zero-suppressed; we record waveform data only for regions above threshold. Each region is called a *pulse*, distinct from the TPC pulses described in Sec. 3.1.2.3. For a given event, there may be many pulses for each channel spread throughout the acquisition window. The first step is to combine pulses together to reconstruct a complete waveform for the event, accounting for baseline drift, ADC saturation, and gain calibration. Waveforms of individual pulses

from each channel are added together to form a sum waveform. We produce separate sum waveforms for the LSV and WCD. To account for baseline drift, a pre-signal region of each pulse is recorded, and the average of the first 12 ns (15 samples) is subtracted from each sample of the pulse waveform. To account for ADC saturation, we employ a simple triangle correction: Linear approximations of small regions before and after the saturation region are evaluated. The two lines are then extended to their intersection to approximate the pulse waveform beyond the ADC dynamic range. To account for PMT gain, the waveforms are scaled by the respective channel's SPE mean, where the SPE mean is evaluated in a manner similar to that of the TPC PMTs (discussed in Sec. 3.3).

Given the reconstructed waveform for each event, we evaluate two types of charge integral estimators for the WIMP search analysis. First, the sum waveform is integrated in fixed regions of interest (ROIs). The main ROI covers a 300 ns window around the trigger time, as given by the TPC trigger. The purpose of this ROI is to efficiently identify signals in prompt coincidence with primary scintillation in the TPC. Typically, these are gammas that scatter both in the TPC and LSV or correlated gammas emitted at the same time as radiation detected by the TPC (e.g. neutrons). The second charge integral estimator is called the slider charge, which looks for regions of maximal charge in the sum waveform. The algorithm slides a fixed length integration window along a pre-defined range of the sum waveform and records the largest signal. The purpose of the slider integral is to identify neutron capture signals in the LSV, which are in delayed coincidence with the neutron scatter in the TPC. The slider integration window is 500 ns and the algorithm scans from the end of the prompt ROI window to the end of the sum waveform. We also evaluate the total integrated signal in the LSV and WCD for each event. Crossing muons will typically produce very large signals and are easily identified by placing a threshold cut on the total integrated signal.

3.3 Single photoelectron mean

Signals from the LAr TPC are typically measured in units of photoelectrons (PE). A photon hits the PMT photocathode and, with some efficiency, ejects an electron inside the PMT via

the photoelectric effect. The electron (which we call photoelectron) is accelerated to a series of dynodes, where an avalanche of electrons is created. The avalanche of electrons produces an observable change in the voltage at the anode. S1 and S2 can produce tens to hundreds or thousands of PE. We must calibrate the PMT response by evaluating the average size of a single PE. This is accomplished with laser calibration: we inject low intensity pulsed laser light into the TPC. The intensity is tuned such that each PMT sees a photoelectron in only a small fraction of triggers ($\sim 5\%$; the fraction is called the occupancy). When photoelectrons do appear, they always appear in a predictable location within the waveform. The size of a single photoelectron (SPE) signal has some variability.

The general idea for the laser calibration is that in the vast majority of laser triggers, no signal appears. In a small fraction, the laser signal is observed. There is then a very small probability that two photoelectrons are present. We integrate a fixed region of interest in every trigger, where we expect the laser signal to appear, regardless of whether it is visible, and build a spectrum. Most of the time, we integrate noise, which forms the pedestal. Single photoelectrons appear as a shoulder or bump in the spectrum. The mean value of the bump indicates the average SPE size.

3.4 Position reconstruction

We now describe the algorithms for reconstructing the longitudinal (z) and transverse (x and y) positions of any particular scatter. Broadly speaking, the z position is given by the time separation between S1 and S2 and the (x, y) position is given by the hit pattern of S2 on the top PMT array. However, the DarkArt reconstruction does not attempt to identify pulses as S1 or S2; it only generically finds pulses within an event. The position reconstruction algorithms therefore rely on a second layer of analysis on top of DarkArt that performs the identification of S1 and S2. We defer that discussion to the next section (Sec. 3.5), and assume here that S1 and S2 have been identified in each event.

z position

The z position reconstruction is straightforward: the time separation between S1 and S2, t_{drift} , gives the z position via simply:

$$z = v_{\text{drift}} \cdot t_{\text{drift}} \quad (3.1)$$

In the LAr TPC, the drift field is uniform, so the drift velocity is constant throughout the active volume. The drift velocity is determined *in situ* by

$$v_{\text{drift}} = \frac{L_{\text{TPC}}}{t_{\text{drift}}^{\text{max}}} \quad (3.2)$$

where $t_{\text{drift}}^{\text{max}}$ is given by the endpoint of the t_{drift} distribution and is 376 μs for 200 V/cm drift electric field and L_{TPC} is the length of the active volume. This gives a v_{drift} value of (0.93 ± 0.01) mm/ μs .

xy position

Transverse (x, y) position reconstruction is less straightforward. S2 light is produced directly below the top array of PMTs, so illumination of the PMTs is highly non-uniform. We use the hit pattern of S2 over the PMTs, particularly the top array, to estimate (x, y) coordinates.

A naive first approach would be to use barycenter—the weighted average of the transverse PMT locations, where the weights are given by the fraction of S2 light on each channel—but that gives a highly nonuniform (x, y) distribution for AAr data (Fig. 3.5). We expect AAr data to be uniformly distributed since they are dominated by ^{39}Ar events. Instead, we use a weighted least squares (WLS) approach. A simple WLS method is to minimize

$$\chi^2 = \sum_i \frac{1}{\sigma_{M_i}^2} (M_i - L_i(x, y) M_{\text{tot}})^2 \quad (3.3)$$

where, for a given event, M_i is the measured S2 signal in PMT i , M_{tot} is the total measured S2 in all PMTs, L_i is the expected light collection as a fraction of total light in PMT i for an event at position (x, y) , and σ_{M_i} is the uncertainty of the measurement. The sum is over all PMTs. The M_i and M_{tot} are data inputs. The challenge is to determine $L_i(x, y)$, which

Figure 3.5: (x, y) distribution of ^{39}Ar events using barycenter position reconstruction. The x- and y- axes are in arbitrary length units, and the z-axis represents counts. Argon-39 events should be uniformly distributed, which barycenter reconstruction dramatically fails to reproduce.

we call the light response function (LRF). We could use Monte Carlo simulation, but the complex optics of the TPC are difficult to model. This presents a chicken-and-egg problem: If we know the true (x, y) positions, we could generate the LRFs; if we know the LRFs, we can evaluate (x, y) for each event using the WLS method; we know neither. Keep in mind that, for any given real data event, we do not know, *a priori*, the event’s (x, y) position. We determine the LRFs using an iterative approach, introduced by Solovov et al. [92].

Despite the large uncertainties in a MC approach, the initial LRFs are derived from a simulation of the TPC optics for S2. Using these initial LRFs, the (x, y) positions are estimated for a large sample of AAr events using the WLS approach. Using the newly reconstructed (x, y) positions, a new set of LRFs can be generated, which can be used to re-evaluate the (x, y) positions of AAr events, and so on. Eventually, the algorithm is expected to equilibrate, and the LRFs and (x, y) positions will be stable from iteration to iteration. However, in DarkSide-50, this condition is not always reached, so to help the process, we use another input. Because the generation of the LRFs is tuned on high statistics ^{39}Ar -dominated AAr data, we expect the transverse positions to be uniformly distributed. At each step of the iteration, we then impose a uniformity “pressure” which modifies the LRFs to encourage more uniform distribution of the events. Notice that, with such a tuning, the LRFs cannot be generated using UAr data, where the uniformity assumption does not hold due to the dominance of external gamma backgrounds. Finally, once the LRFs are determined, the (x, y) positions can be determined event-by-event. The resulting (x, y) distribution for AAr data is shown in Fig. 3.6.

This and other (x, y) reconstruction methods are under active development within DarkSide-50. One method uses a similar WLS approach, but develops the LRFs in a different manner. Another reconstruction uses principle component analysis and a multidimensional fit.

3.5 SLAD

Though the DarkArt reconstruction is comprehensive, its outputs are not immediately readily usable for many analysis tasks. The outputs are in terms of the reconstruction variables,

Figure 3.6: (x, y) distribution of ^{39}Ar events using WLS position reconstruction. Events are nearly uniformly distributed, as expected for ^{39}Ar decays. We use 1 cm^2 binning as that is the estimated resolution of the transverse position reconstruction.

which often make sense only in the context of the reconstruction algorithms. Furthermore, DarkArt does not attempt to identify S1 and S2 in each pulse, leaving the determination to a later step. As a result, DarkArt calculates an extremely broad set of variables, designed to allow users to build any sort of analysis variables they wish but does not provide a set of variables ready for immediate consumption by general users. While this reduces the need for large DarkArt reprocessing campaigns, the variety of parameters available in the DarkArt outputs verges on excessive and bloats the output files. DarkArt output files are often up to tens of GB per run, making them cumbersome to work with. Finally, reading of the DarkArt output files requires the DarkArt library files for anything beyond the most primitive analysis tasks.

The SLim Analysis Data (SLAD) program aims to solve these issues. On the analysis front, SLAD identifies S1 and S2 pulses for each event and provides analysis-friendly variables, including estimators for the S1 and S2 amplitude, corrections to those estimators, drift time, (x, y) , f_{90} , etc. On the technical front, SLAD provides lightweight data files that can be used easily and quickly. We try to find a balance between providing a broad enough set of variables that SLAD is useful for the vast majority of analysis tasks, but slim enough that the SLAD output data files are portable. (SLAD files can be accessed by plain ROOT.)

3.5.1 Pulse identification

The first main analysis task of SLAD is to identify the S1 and S2 pulses, if any, for each event. There is a great variety of pulse topologies that are commonly observed in the TPC and reconstructed by DarkArt. We list the major classes of events here:

- S1 + S2: By far the most common topology, corresponding to single scatter events in the TPC.
- S1 + multiple S2s: Usually due to multiply scattering gammas. The multiple scatters occur within nanoseconds of each other, so the S1s of all the scatters are observed at the same time. If the scatters occur at different z positions, then the S2 pulses will be separated in time and are distinguishable by the reconstruction.

- S1 + S2 + S2-echo: S2 light can induce photo-ionization of the cathode surface. The ionization electrons then drift the full length of the TPC, are extracted into the gas region, and produce electroluminescence signals, which we call S2-echo (also referred to as S3). S2-echo pulses are easily tagged because they appear at fixed separation from S2, equal to the maximum drift time of the TPC.
- S1 + S2 + S1-echo: Similarly to S2, S1 light can induce photo-ionization of the cathode, producing electroluminescence signals (S1-echo) at fixed separation from S1.
- S1 + S2 + S1-echo + S2-echo
- Multiple S1s and multiple S2s: These events are typically accidental coincidences of events in the TPC. In the AAr data, these occur at the $\sim 2\%$ level, mostly due to ^{39}Ar activity, while in UAr data, the rate of accidental coincidences is lower.
- S1 only: There are a variety of classes of events that can produce an S1 only. One is that there are holes in the TPC walls for liquid argon inlet and outlet. Decays of ^{39}Ar can occur in these cavities, and we observe their S1, but their charge is trapped by the PTFE wall and no S2 is produced. Another class of S1-only events is associated with α 's. Their f_{90} is above the ER band, but the α 's are produced near (or in) the PTFE wall and their ionization signal is trapped.
- Cherenkov: Gamma rays can Compton scatter in the fused silica windows or in the PMT photocathodes, and the recoiling electron then produces Cherenkov light. Events with Cherenkov often appear without any scintillation signal, and the light is usually concentrated in a single PMT.
- S2 only: Events in which the TPC trigger missed the S1 and triggered on the S2. Most often, an interaction occurs in the TPC during an event's inhibit window, and the S2 then appears after the inhibit has released and triggers the DAQ. The signature for such events is typically that their event timestamp is very close to the previous trigger.
- S2-echo only: Events in which the TPC triggered on an S2-echo. This class of events is very similar to the S2 only events. The TPC trigger misses the S1 and the S2 because

they are in an event’s inhibit window, and the S2-echo appears after the inhibit has been released.

- Merged S1 and S2: Typically events near the very top of the TPC, where the S1 and S2 are clearly visible but are so close together that the DarkArt reconstruction merges them into a single pulse. These are discussed in great detail in Sec. 4.4.

The goal of the pulse identification algorithm is to identify, when possible, the S1 and S2 pulses within any given event. The algorithm identifies S1 and S2 only when a small set of conditions are satisfied: in particular, we identify the first pulse as S1 and the second pulse as S2 when either

- DarkArt identifies exactly 2 pulses in the event, or
- DarkArt identifies exactly 3 pulses in the event and the 3rd pulse is consistent with being S2-echo.

We identify a pulse as being an S2-echo when the time difference between the 2nd and 3rd pulse is within $372\ \mu\text{s}$ to $400\ \mu\text{s}$. (The maximum drift time of the TPC is $376\ \mu\text{s}$.) The wide window allows for variability of the pulse finder algorithm to identify the start of the S2-echo pulse. Notice that this scheme does not make any use of pulse shape information or pulse amplitude information. This is to avoid introducing different biases in identifying ER and NR scatters.

In the context of a WIMP search, there is no need to identify S1s and S2s in events with other topologies besides S1+S2 and S1+S2+S2-echo. The DarkArt pulse-finder is tuned well enough that we can rely on it to accurately reconstruct the number of physical scintillation signals. See Sec. 3.8 and 6.2 for discussions of the efficiencies of the pulse-finder and pulse-identification scheme.

3.5.2 Corrections

Once the S1 and S2 pulses are identified within an event, we can apply corrections to them, which account for geometrical variations of the light collection efficiency across the TPC.

The corrections are developed independently. We generate correction functions for S1 and S2 which generally depend on position. Here we assume the 3D position of any single-sited scatter can be reconstructed with reasonable accuracy ($\sim 1\%$ for z and $\sim 5\%$ for (x, y)). The correction functions are hard-coded into the main SLAD program.

S1 corrections

For S1, we apply only a z correction. Due to total internal reflection at the liquid-surface interface and the (slightly) less than unity optical transparency of the extraction grid, more S1 light is collected in the bottom array of PMTs than in the top array, on average. And for events near the bottom of the TPC, more light is collected overall. The light collection efficiency curve as a function of z , relative to the center of the TPC, is shown in Fig. 3.7 and varies by 14% from the top to the bottom. This curve is generated using a combination of data from ^{83m}Kr injected into the TPC and from the ^{39}Ar endpoint. We found that the ^{83m}Kr does not populate the top of the TPC, likely due to the circulation pattern of the liquid, necessitating the use of the ^{39}Ar endpoint to estimate the light yield in that region. Using the same ^{83m}Kr data, we find that there is $<3\%$ variation in the light collection efficiency vs. transverse position at all z positions, so we do not apply an (x, y) correction to S1.

S2 corrections

The S2 detector response in DarkSide-50 is found to have a strong (x, y) dependence. The S2 signal is found to be about three times stronger at the center of the TPC than at the edges, as shown in Fig. 3.8. The cause of the radial dependence has not been firmly established. Possible explanations include a sagging anode window or electromechanical deflection of the grid. Both of these models tend to make a stronger electroluminescence field near the center of the TPC. This effect is supported by measurements of the electron drift time across the gas pocket, discussed in Sec. 4.4. Though the radial dependence is not fully understood, we can correct for it. We normalize the S2 signal to the center of the TPC. The normalization is done event-by-event using a 2D correction map generated using ^{83m}Kr data.

(a)

(b)

Figure 3.7: (a) Light yield relative to TPC center, measured from ^{83m}Kr calibration data (black) and ^{39}Ar data (red). (b) Fifth order polynomial (black) fitted to the combined light yield measurements (red).

(a)

(b)

Figure 3.8: (a) 2D profile of the $S2/S1$ response vs. (x, y) of ^{83m}Kr events. The z-axis is the mean $S2/S1$ response at the given (x, y) position. (b) Distribution of ^{83m}Kr events in the $S2/S1$ vs. r plane (color) and profile of the $S2/S1$ response vs. r (black).

Figure 3.9: Profile of the distribution of $S2/S1$ vs. t_{drift} for evaluation of electron drift lifetime. $S1$ is z -corrected is restricted to 350 PE to 450 PE, and $S2$ is xy -corrected. While we expect exponential decay of $S2/S1$ with respect to t_{drift} due to attachment to electronegative impurities, we observe a non-exponential behavior.

The $S2$ detector response has z variations due to impurities in the LAr and possibly other effects. Impurities in the LAr can capture the drifting electrons, reducing the observed $S2$ signal. The survival probability for electrons to drift all the way to the gas phase follows an exponential distribution, whose mean is referred to as the electron drift lifetime. The electron drift lifetime is measured by fitting an exponential to the profile of the $S2/S1$ vs. t_{drift} distribution. The $S1$ must be z -corrected as described above. Because the $S2$ signal has a strong radial dependence, the $S2$ must be xy corrected. (Alternatively, one could carry out the analysis using small r slices.) The electron drift lifetime in DarkSide-50 is estimated to be >5 ms, and with a maximum drift time of $376 \mu\text{s}$ in the TPC, the total z variation of $S2$ is $<7\%$. The $S2/S1$ vs. t_{drift} distribution is found not to be exactly exponentially distributed, as shown in Fig. 3.9. This indicates there are other effects besides impurities affecting the z dependence of $S2$. Therefore, for most analyses, we do not apply any z correction for $S2$ and accept the $<7\%$ variation as a systematic uncertainty.

3.6 Commonly used variables

We list the most commonly used variables available in SLAD outputs. Integrals here are the sum of the integrals of the individual channels.

- S1: integral of first 7 μ s after start of S1 pulse.
- S2: integral of first 30 μ s after start of S2 pulse
- f_{90} : ratio of integral of first 90 ns to integral of first 7 μ s of S1 pulse¹
- S1 max frac: fraction of S1 light in the dominant S1 channel
- N_{PE} : integral of a pulse from start to end, as found by the pulse-finder
- S1 top-bottom asymmetry: $(S1_{top} - S1_{bot})/(S1_{top} + S1_{bot})$ where $S1_{top}$ ($S1_{bot}$) is the sum of S1 integrals over the top (bottom) PMTs.
- S1 start time: pulse start time of S1, as found by pulse-finder
- S2 start time: pulse start time of S2, as found by pulse-finder
- t_{drift} : difference between S1 start time and S2 start time
- S2 f_{90} : ratio of integral of first 90 ns to integral of first 30 μ s of S2 pulse
- $S1_{corr}$: z -corrected S1
- $S2_{xycorr}$: xy -corrected S2
- x, y : transverse coordinates as determined from (x, y) position reconstruction
- inhibit time: time of inhibit window associated with previous trigger
- live time: time from end of previous event's inhibit to present event's trigger. The sum of the inhibit time and live time gives the time to the previous trigger.

¹Due to 4 ns sampling, 90 ns integration is over 22 samples, which more precisely corresponds to a 88 ns integration window.

3.7 Electronics MC

We now describe the electronics Monte Carlo, which produces simulated raw data and is implemented within the DarkArt software. The ability to produce fake raw waveforms is a powerful tool for studying a variety of effects, including systematic biases of the reconstruction, SPE calibration procedure, f_{90} model, etc.

There are two main ingredients for producing raw data: signal and baseline. The baseline encapsulates all the electronics noise of the amplifiers, front end boards, cables, digitizers, etc. Rather than simulate the electronics noise, we calibrate directly on the DarkSide-50 electronics. We use real baseline data obtained by randomly triggering the DarkSide-50 DAQ. The low background rate of DarkSide-50 ensures that the majority of randomly triggered events are void of scintillation signals. For the ~ 50 Bq of activity in AAr data, about $\sim 2\%$ of random triggers exhibit accidental coincidence with a scintillation pulse.

The signal can be produced in two different ways: (1) Simulate events in *g4ds*, DarkSide's Monte Carlo code (described in Sec. 5.2, which includes energy deposition, a model of the microphysics of LAr, and the optics. For each event, *g4ds* produces a list of PEs, each specifying which channel it appeared on and at what time relative to the event's $t = 0$. (2) Produce a list of PEs directly within the electronics MC. The hit pattern and time distribution of the PEs for each event can be finely controlled. Though the second method is not very useful for reproducing accurate hit patterns of PEs over the PMTs (which requires simulating the optics of the TPC), it is extremely useful for rapid analysis and is sufficient for a variety of studies. The hit pattern of the PEs over the PMTs and the PE time distribution are independent. Typically, we produce PE times according to one of three time profiles: S1, S2, or laser. The S1 time profile is a two-component exponential decay. The S2 time profile is a two-component exponential convolved with a boxcar function and a Gaussian; see Sec. 4.2 for details of the S2 time profile. The laser time profile is a delta function, implemented as a narrow Gaussian; i.e. laser PE always appear at the same time relative to $t = 0$ for each simulated trigger. The hit pattern of the PEs over the PMTs can be chosen arbitrarily. Examples of some distributions that have been implemented include fixed number of PE per

Figure 3.10: Template single PE response generated by averaging many single PEs. The red curve is a mathematical function chosen to match the data.

channel, Poisson distributed number of PE per channel, and uniform distribution of PE over the channels. The total number of PE per event can also be configured to be fixed, or follow a Poisson distribution, or follow a uniform distribution in some range.

Once the PE times and hit pattern are defined, either from G4DS or by simulating directly in the electronics MC, the next step is to simulate the PMT response for each PE. That is, we assign a waveform contribution to each PE. For simplicity, we assume that the waveform response of each PE is identical. The template waveform is shown in Fig. 3.10 and is an average response of many single PEs. The overall amplitude of each PE must be simulated. The amplitude distribution is given by the SPE spectrum, described in Sec. 3.3. For simplicity, we assume the amplitude distribution is identical for all PMTs and follows the Gaussian plus exponential model described in Ref. [93]. The parameters of the model are chosen to match the average SPE distribution of the DarkSide-50 PMTs.

Once each PE is assigned a waveform response, we build the aggregate waveform response for each channel by stacking the individual PE waveforms. We then overlay the

simulated signal waveform onto the real baseline waveform. Finally, we simulate the ADCs by digitizing the combined simulated signal + real baseline waveform. (Though the y-axis of the simulated PE waveform is in units of ADC counts, the template PE waveform is an average, and so the y-axis is essentially continuously distributed. Real raw data is digitized to the 12-bit resolution of the V1720 ADCs.) Naively, one would perform the digitization by rounding—either conventional rounding or rounding down (floor operation) or rounding up (ceiling operation)—each sample of the waveform to the nearest integer. But this is inaccurate: because the baseline data is taken directly from the DarkSide-50 DAQ, the baseline waveforms are already digitized. Adding the continuously distributed signal to the baseline and then digitizing introduces a large bias. Suppose we use a floor operation to digitize. Then even a tiny PE signal, which could have a response smaller than a single ADC count, is amplified to a full count (recall that signal is a negative deviation from baseline for raw PMT data), and when the simulated waveform is integrated, the total pulse area will be overestimated. Similarly, if we used a ceiling operation, then small PE are truncated, and the pulse area will be underestimated. These inaccuracies can be as large as 3%.

The solution is to add vertical jitter to the simulated waveforms: after overlaying the simulated signal on baseline, add a random offset and then do the floor operation. The random offset is uniformly distributed between 0 and 1 ADC count. (If one wanted to use a ceil digitization, one would subtract the random offset.) This method produces realistic digitized waveforms, while keeping the bias of the digitization procedure to a minimum. The integrated simulated signal then matches the true simulated signal very well.

3.8 Reconstruction efficiencies and biases

We now apply the electronics MC to study reconstruction efficiencies and biases of the DarkArt algorithms. Reconstruction efficiency represents the ability to accurately reconstruct the number of PE in a pulse. Biases here represent systematic biases in energy scale.

The baseline finder is the major source of bias in reconstructing the number of PE per pulse. As described in Sec. [3.1.2.1](#), we interpolate the baseline underneath signal regions of

Figure 3.11: Example interpolation portion of the baseline-finder algorithm. Black is the raw waveform, red is the baseline, and the green points are the endpoints of the interpolated region. The linear interpolation tends to have negative slope, which biases estimators that use integrated windows.

each waveform. Yet we know that the linear interpolation is not a realistic representation of the true baseline—we see long timescale fluctuations in the baseline when there is no signal present, and we have no reason to expect that there should not be fluctuations in the baseline when there is signal present. Furthermore, the baseline finding algorithm tends to eat away part of the signal region: as the algorithm traverses along a signal region of the waveform, it must decide at some point that the waveform has returned to baseline. The criterion for the waveform to return to baseline is that all the samples of the boxcar window be within some range of the last baseline sample. But since signal is negative deviation from baseline (in raw waveforms), the first baseline sample coming out of a signal region tends to be *lower* than the last non-interpolated baseline point. See Fig. 3.11 for an illustration of this effect. The interpolated baseline then tends to have a negative slope, which “eats away” at the integral estimate of the signal amplitude, and we thus expect a slight negative bias in the reconstructed signal amplitude.

We can precisely quantify the bias using the electronics MC. We build fake waveforms with precisely known number of simulated PE, pass the waveforms through the DarkArt reconstruction, and compare the reconstructed number of PE to the simulated number of PE. The bias associated with the baseline finding algorithm depends on the number of interpolated regions of the baseline and on the slopes of the interpolations, which both depend on the time profile and size of the recorded signals. Therefore, the baseline reconstruction bias is different for S1 and S2, and we perform this analysis separately for each. The S1 reconstruction bias is $\sim 0.5\%$ and the S2 reconstruction bias is $< 0.5\%$. The biases as a function of the size of S1 and S2 are shown in Fig. 3.12. The estimation of f_{90} is subject to a similar bias. The number and slopes of the interpolated regions of the baseline in the prompt window (in fact, there is usually only one interpolation for the prompt region in each channel), and therefore also the bias in the baseline reconstruction in the prompt window, are different from those of the total S1 window. The ratio of the prompt to the total is then biased. Using the same simulated S1 waveforms, we estimate the f_{90} bias to be $\sim 1\%$, as shown in Fig. 3.12.

Another source of inefficiency in the reconstruction software is our inability to identify very small pulses. The TPC trigger is set very low, requiring only 2 or 3 (depending on campaign) discriminators to fire within a 100 ns window. In contrast, the pulse finder algorithm of the DarkArt reconstruction software is designed to find macroscopic pulses of PE. Roughly speaking, the pulse finder has a minimum threshold of 5 PE in a $2\ \mu\text{s}$ window. We can use the electronics MC to precisely quantify the pulse finder's ability to identify (the start of) pulses. Again, since the time profiles of S1 and S2 are different, the pulse finder efficiency may be different for each. We analyze each type of signal separately.

We generate fake S1 and S2 waveforms. The sizes of the S1s and S2s are drawn from uniform distributions from 0 PE to 100 PE. The simulated waveforms are then passed through the standard DarkArt reconstruction. We find that the pulse finder is 50 % efficient at identifying S1s of 7 PE and reaches 100 % efficiency around 13 PE. Similarly, the pulse finder reaches 50 % efficiency to detect S2 at 17 PE and 100 % around 38 PE. The efficiency curves for S1 and S2 are shown in Fig. 3.13.

(a)

(b)

Figure 3.12: (a) Bias in reconstructed size of S1 (black) and f_{90} (red). (b) Bias in reconstructed size of S2. The down turn is due to discrepancy introduced by simulated ADC saturation.

Figure 3.13: Efficiency of DarkArt pulse finder to identify the start of S1 (blue) and S2 (red) pulses.

CHAPTER 4

S2

4.1 Introduction

In a dual-phase liquid argon Time Projection Chamber (TPC) such as DarkSide-50 [76], the so-called S2 signal is the electroluminescence signal produced by ionization electrons that have been drifted through the liquid and across the gas pocket above the liquid. The S2 signal encodes a tremendous amount of information about each event and about the TPC itself. In this chapter, we focus on the pulse shape of S2 and apply the results to measure the diffusion of electrons in liquid argon.

In Sec. 4.2 we derive an analytic form of the S2 pulse shape. In Sec. 4.3 we discuss how to fit the analytic shape to data. In Sec. 4.4 we use the S2 pulse shape to characterize the gas pocket in DarkSide-50. In Sec. 4.5 we apply the fitting procedure to various data sets to perform a measurement of electron diffusion in liquid argon.

4.2 S2 pulse shape

The analytic expression for the S2 pulse shape is derived from the following model for the production of light in the gas pocket of the TPC. We assume electrons drift with constant velocity across the gas pocket, producing Ar excimers uniformly along their drift path. The excimers de-excite and produce light according to a two-component exponential [94], similar to the light production in the liquid. If all electrons are extracted from the liquid at precisely the same time, then these two effects define the S2 pulse shape, as described mathematically in Sec. 4.2.1. In reality, electrons of a given ionization cloud are extracted from the liquid

with a distribution of times, which we model by introducing a Gaussian smearing term, as described in Sec. 4.2.2.

4.2.1 Basic shape

The basic, or idealized, form of the S2 pulse shape is a time profile $y(t)$ given by the convolution of a uniform distribution with a two-component exponential:

$$y_{\text{ideal}}(t; \tau_1, \tau_2, p, T) = \text{uniform}(0, T) * (\text{two component exponential}) \quad (4.1)$$

$$= p \cdot y'_{\text{ideal}}(t; \tau_1, T) + (1 - p) \cdot y'_{\text{ideal}}(t; \tau_2, T) \quad (4.2)$$

where

$$y'_{\text{ideal}}(t; \tau, T) = \frac{1}{T} \begin{cases} 0, & \text{if } t < 0 \\ 1 - e^{-t/\tau}, & \text{if } 0 \leq t \leq T \\ e^{-(t-T)/\tau} - e^{-t/\tau}, & \text{if } t > T \end{cases} \quad (4.3)$$

Here, τ_1 and τ_2 are the fast and slow component lifetimes, respectively; p is the fast component fraction; and T is the drift time of the electrons across the gas pocket. We assume that all electrons are extracted out of the liquid at the same time, where we set $t = 0$. The derivation of Eqn. 4.2 and 4.3 is deferred to Sec. 4.6. The two decay constants are expected to differ from those of the liquid. According to Ref. [94], the fast component is about 11 ns and the slow component is about 3.2 μs . An example pulse shape is shown in Fig. 4.1. Notice that T governs the time to the peak of the pulse. The “kinks” in the rising and falling edges are due to the drastically different decay times τ_1 and τ_2 , and the vertical positions of the kinks (on this time-scale) are set by p , while the horizontal positions are set by the total drift time in the gas.

4.2.2 Gaussian smearing

The above model for the S2 pulse shape assumes all electrons are extracted out of the liquid at the same time. In fact, there are many reasons this is not true. Primarily the cloud of electrons is diffuse, with diffusion arising from the cloud drifting through the liquid and from passing through the grid.

Figure 4.1: Example S2 pulse shape with $\tau_1 = 0.05 \mu\text{s}$, $\tau_2 = 3.2 \mu\text{s}$, $p = 0.1$, and $T = 1.5 \mu\text{s}$. Black: idealized form (no smearing). Green: includes Gaussian smearing at $\sigma = 0.3 \mu\text{s}$.

To model the diffusion, we incorporate a Gaussian smearing term into the S2 pulse shape by convolving Eqn. 4.2 with a Gaussian centered at 0 with width σ :

$$y(t; \tau_1, \tau_2, p, T, \sigma) = y_{\text{ideal}} * \text{gaus}(0, \sigma) \quad (4.4)$$

$$= p \cdot y'(t; \tau_1, T, \sigma) + (1 - p) \cdot y'(t; \tau_2, T, \sigma) \quad (4.5)$$

where

$$y'(t; \tau, T, \sigma) = \frac{1}{2T} (y''(t; \tau, \sigma) - y''(t - T; \tau, \sigma)) \quad (4.6)$$

$$y''(t; \tau, \sigma) = \text{erf}\left(\frac{t}{\sqrt{2}\sigma}\right) - e^{-t/\tau} e^{\sigma^2/2\tau^2} \text{erfc}\left(\frac{\sigma^2 - t\tau}{\sqrt{2}\sigma\tau}\right) \quad (4.7)$$

We again defer the derivation to Sec. 4.6. This form is analytically complicated, but it has a simple intuitive interpretation: It is the ideal shape of Eqn. 4.2 with all the sharp features smoothed out, as shown in Fig. 4.1. For a preview of the level of smearing present in S2s in DarkSide-50, see Fig. 4.16, which shows example events from near the top and bottom of the TPC.

4.2.3 Degeneracy of Parameters

To describe an arbitrary S2 pulse, in addition to the pulse shape and smearing parameters, we include three additional parameters in the fit function: a time offset t_0 , a vertical offset y_0 , and an overall scale factor A . So the final function we use for the fit is of the form:

$$y_{\text{fit}}(t; \tau_1, \tau_2, p, T, \sigma, A, t_0, y_0) = y_0 + A \cdot y(t - t_0; \tau_1, \tau_2, p, T, \sigma) \quad (4.8)$$

The form of the S2 pulse shape given in Eqn. 4.8 has an approximate degeneracy: the same shape can be produced using different combinations of T , t_0 , and σ . The degeneracy can be seen visually in Fig. 4.2, where five nearly identical pulse shapes are shown using different parameter values. This degeneracy can result in incorrect parameter estimation if the parameters are all left free in the fit. In this section, we describe how the degeneracy arises. For the analyses in this work, we break the degeneracy by fixing one or more of the parameters in the fits. The choice of parameters to fix depends on the goals of the analysis, so we defer such a discussion to later.

The degeneracy of these parameters was first revealed empirically by playing around with the analytic form of the S2 pulse shape. The following relation seemed to hold:

$$t_0 + T/2 = -1. \quad (4.9)$$

To develop a feel for the degeneracy of T , t_0 , and σ , we build a family of (nearly) identical pulse shapes as follows. We take a reference shape with rise time T_0 , and consider other shapes by varying T about T_0 . Then we construct values for t_0 and σ that “undo” the deformations induced by varying T . In the tail, the S2 pulse shape is a decaying exponential, and a change in T corresponds to a change in amplitude; the decay constant does not change. This can be seen from Eqn. 4.3 and by noting that the smearing has little effect on the tail ($t > T$), as evident in Fig. 4.1. Since t_0 is a time offset, and a time offset in an exponential is equivalent to a change in amplitude, we can always find a value for t_0 that effects the opposite change in amplitude in the tail. Furthermore, since the smearing has no effect on the tail, once the tails are made to match by varying T and finding an appropriate t_0 , σ can be tuned such that the rising edge and peak of the pulse shapes match for different T .

(a)

(b)

(c)

Figure 4.2: Sequence for building a family of different S2 pulses with nearly identical pulse shape. (a) Start with pulse shapes of various T . (b) Align the exponential tails by adjusting t_0 . (c) Align the rising edges and the peaks by adjusting σ .

We now make quantitative statements for the above argument. We first find a function $t_0(T)$ that prescribe the change in t_0 for a given variation in T such that the S2 tail remains unchanged. Since the smearing has no effect on the tail, we can arbitrarily take $\sigma = 0$ for $t \gg T$. Furthermore, because the fast component lifetime is so small compared to the slow component lifetime, the fast component has negligible contribution to the tail. So from Eqn. 4.3, we can write

$$y_{\text{tail}}(t; \tau_2, p, T, t_0) = y(t \gg T; \tau_2, p, T, t_0) \approx \frac{1-p}{T} \cdot e^{-(t-t_0)/\tau_2} \cdot (e^{T/\tau_2} - 1) \quad (4.10)$$

To enforce that varying T about T_0 has no change in the tail, we construct t_0 to satisfy

$$\frac{1-p}{T_0} \cdot e^{-t/\tau_2} \cdot (e^{T_0/\tau_2} - 1) = \frac{1-p}{T} \cdot e^{-(t-t_0)/\tau_2} \cdot (e^{T/\tau_2} - 1). \quad (4.11)$$

Solving for t_0 , we have

$$t_0 = \tau_2 \cdot \ln \left(\frac{T}{T_0} \frac{e^{T_0/\tau_2} - 1}{e^{T/\tau_2} - 1} \right). \quad (4.12)$$

Now we can vary T and t_0 such that, by construction, the tail of the S2 pulse shape remains unchanged. The function is shown in Fig. 4.3 where nominal values for T_0 and τ_2 are used. Furthermore, we can recover the linear relation of Eqn. 4.9 by calculating the lowest order expansion of $t_0(T)$ about T_0 (and assuming $\tau_2 > T_0$ which is usually true in our case):

$$t_0(T) \approx \frac{T_0 - T}{2}. \quad (4.13)$$

(The constant offset term is different because the empirical relation (Eqn. 4.9) was found with $t_0 = 0$ defined differently than here.)

Finally, we vary σ until all the shapes are nearly aligned across all t . Increasing σ lowers the peak and pushes the rising edge out to the left, and decreasing σ raises the peak and pushes the rising edge in to the right. We have not found a quantitative method for finding the values of σ for a given set of t_0 and T that give degenerate S2 pulse shapes. The above procedure is illustrated graphically in Fig. 4.2.

In summary, it is evidently possible to build a family of waveforms that have nearly identical pulse shape. Any change in T can be compensated for in the tail by t_0 and in the

Figure 4.3: Degeneracy of parameters in S2 pulse shape. Blue: Eqn. 4.12 with $T_0 = 2 \mu s$ and $\tau_2 = 3.2 \mu s$. Red: Empirical linear relation between t_0 and T : $t_0 = (2 - T)/2$.

rise and peak by σ . This works for some values of σ , but not all. For example, if $\sigma = 0.1$, there is essentially no smearing and the varying shapes in different T cannot be accounted for by σ . Furthermore, the degeneracy is not exact: all the parameters are necessary to uniquely define the rising edge and the peak of the pulse shape. But the approximate degeneracy is enough to invalidate any fit results that do not account for it. To break the degeneracy of T , t_0 , and σ , one must either operate in a regime in which the degeneracy does not exist (namely, small σ) or fix one or more of the parameters in the fit. The choice of approach depends on the purpose for doing the fits. For example, in Sec. 4.4, we fix σ , while for the electron diffusion study in Sec. 4.5, we fix T .

4.3 Fitting S2 pulse shape

In this section we describe the general procedure for fitting the S2 pulse shape. We motivate the use of maximum likelihoods for fitting waveforms and describe a goodness-of-fit calculation. Though the section is written in the context of fitting S2, the results are applicable to fitting any waveform data in DarkSide-50.

4.3.1 Maximum likelihood fitting

One of the long-standing problems in DarkSide is how to fit analytic functions to waveform data. The difficulty has usually been how to treat uncertainties of each sample. This is important for any topics related to pulse shape fitting, including S1 and S2. Over the course of this analysis, we have come up with a coherent method for how to handle uncertainties for each sample in a given waveform. The main idea is to treat each waveform as a histogram of photoelectron times. This interpretation is how we often think about waveforms intuitively. Here we show how to adopt the interpretation more formally and discuss its range of validity. This allows us to use standard histogram fitting techniques, particularly those prescribed by Baker and Cousins [95].

We start with perhaps a blatantly obvious statement: raw waveforms cannot be interpreted as histograms of PE times. We must have at least a baseline-subtracted waveform,

negated so that signal appears as a positive deviation above baseline. For each event, the baseline-subtracted waveforms of all the channels may be scaled by their respective SPE means and added together to construct a sum waveform in the usual way, as described in Sec. 3.1.2.2. In standard DarkArt reconstruction, we usually zero-suppress the individual channels before building the sum waveform. For this analysis, we will generally use events with large PE statistics, so the use or non-use of zero-suppression is expected to have little impact on the final results. We now build the histogram interpretation of waveforms on the sum waveform, though the arguments are also valid on individual channels.

In most cases, histograms have bin contents that represent counts (or relative frequencies), which are typically integer valued. The events that make up the histogram are typically independent, so the contents are uncorrelated. To interpret the sum waveform as a histogram of PE arrival times, we can consider the x-value of each sample as a bin, and the y-value as the number of PE arriving in that sample. However, at the 250 MHz sampling rate of the DarkSide-50 DAQ, the waveform resolution is higher than the width of a single photoelectron, which is 10 ns, so at 4 ns sampling, the bin contents of adjacent bins are strongly correlated. To reduce the bin-to-bin correlations, we re-bin the sum waveform, where each bin is the sum of 8 consecutive samples (32 ns). Some photoelectrons will span two bins, inducing some bin-to-bin correlations, but at 32 ns binning, the correlations should be greatly reduced.

Based on the assumption that histogram entries are independent, each bin's contents should then follow Poisson statistics (assuming the total number of events in the histogram is not constrained). The contents of a single bin is a Poisson random variable. So despite the fact that our (re-binned) waveforms do NOT have integer bin contents, we can make progress if we can demonstrate the bin contents approximately follow Poisson statistics. There are several reasons why the Poisson bin statistics may not hold: for example, the bin-to-bin correlations are too high or the integrated noise is too high. In the next section, we will demonstrate in what regimes the waveform bin contents follow Poisson statistics, which will set the range of validity for interpreting sum waveforms as histograms of PE times.

Once we have re-binned the sum waveform, we can fit the S2 pulse shape using the maximum likelihood technique described in Ref. [95]. We use ROOT's `TH1::Fit()` method

with option “L”, which implements the method. We evaluate goodness-of-fit using a separate χ^2 calculation. See Sec. 4.3.2.

4.3.1.1 Poisson bin statistics

To validate our interpretation of waveforms as histograms, we should check that the bin contents follow Poisson statistics. The general idea is to look at the integral spectra for various size integration windows, as small as 4 ns, from a large sample of similarly sized S2s. We use both simulated waveforms and real data.

The electronics simulation is described in Sec. 3.7. Here we provide only a brief summary of the configuration of the electronics simulation. We build fake S2 waveforms by selecting PE times from the S2 time profile given in Eqn. 4.4. The total number of PE per event is fixed to a nominal value of 1×10^4 PE, which are then uniformly distributed over the PMTs. Notice that this is not an accurate simulation of the PE distribution over the channels for real S2, but since we are concerned with the sum channel only, this inaccuracy should have little effect on the results. Each PE is given a time profile using the average shape and scaled by a size chosen from a nominal SPE distribution. We use the same SPE distribution for all channels. Each of the SPE time profiles is accumulated onto a baseline waveform. Finally, we have a set of simulated waveforms with S2 signal that can be written and saved to file and processed identically to real data.

For real data, we use the high statistics of ^{39}Ar decays in AAr data, selecting single scatter events with S2 in the range of 9.9×10^3 PE to 10.1×10^3 PE and coming from a t_{drift} range of 100 μs to 130 μs . Since we expect the shape to vary with z due to diffusion, we use a narrow t_{drift} range to ensure that the S2 pulse shapes are similar.

For both real and simulated data, we integrate fixed length windows of the sum waveform starting at fixed offset with respect to the start of S2 pulses, as shown graphically by the cartoon in Fig. 4.4. We then histogram the integrals and fit Poisson distributions to each spectrum. Examples from both simulated and real data are shown in Fig. 4.5, where we use two different integration windows: 5.000 μs to 5.004 μs and 5.000 μs to 5.032 μs . We see that

Figure 4.4: Cartoon illustrating the fixed length integration window (red) starting at fixed offset from the S2 pulse start (blue). Inset: zoomed view of the waveform with the two integration windows highlighted in different shades of red.

for the 4 ns integration window, the spectra are not consistent with Poisson distributions, both in data and MC, while for the 32 ns integration windows, the spectra are consistent. The 32 ns integration window corresponds to re-binning of the sum waveform, combining every 8 samples together. The procedure of building the integral spectra can be repeated with different size integration lengths and different offsets from pulse start. The Poisson nature of the PE statistics holds in all cases with integration windows greater than or equal to 32 ns.

4.3.2 Goodness-of-fit

To evaluate goodness-of-fit of the S2 pulse shape on the waveforms, we evaluate a χ^2 statistic. However, many of the bins have low (fewer than 5) counts, even after 8 sample re-binning, invalidating a direct χ^2 evaluation. To resolve this issue, we re-bin the waveform again, this time using unequal bin widths. We choose the bin edges so that, for the (nominal) S2 pulse shown in Fig. 4.6a, each bin has equal counts (Fig. 4.6b). The binning is configured so that

(a) 4 ns integration window from MC

(b) 32 ns integration window from MC

(c) 4 ns integration window from real data

(d) 32 ns integration window from real data

Figure 4.5: Example spectra built from MC and real data. Black are data. Red are Poisson distributions fitted to the data.

Figure 4.6: (a) Sample S2 pulse shapes with nominal smearing (green) and no smearing (black). (b) Re-binned S2 pulse shapes using unequal bin widths, chosen such that the S2 pulse with smearing has a flat distribution. The black and blue curves have the same binning. (c) Sample S2 from electronics MC (black) with fitted pulse shape (magenta). (d) Re-binned versions of waveform and fit.

the minimum bin width is 32 ns, and the bin edges are truncated to land on 4 ns intervals. For simplicity, we use the same re-binning for evaluating the χ^2 of all events. As the pulse shape varies, the re-binned waveforms will NOT populate the bins with equal counts, as shown in Fig. 4.6b. But their shapes will still be similar enough that the bins will not go below 5 counts. Example re-binned waveforms are shown in Fig. 4.6d.

The χ^2 statistic that we use is the one prescribed by Baker and Cousins [95], reproduced here:

$$\chi^2 = 2 \sum_i y_i - n_i + n_i \ln \left(\frac{n_i}{y_i} \right) \quad (4.14)$$

where the sum is over the bins of the re-binned S2 waveform, n_i is the content of the i th bin, and y_i is the number of PE predicted by the model to be in the i th bin. The form of Eqn. 4.14 is derived from Poisson statistics and reduces to the more familiar form for Gaussian variances in the limit of very large statistics.

4.4 Gas pocket drift time and zero-diffusion events

One of the many pieces of information that the S2 signal encodes is the drift time of the electrons across the gas pocket. In the absence of diffusion, the sharp turn-on and turn-off of the production of Ar excimers in the gas combined with the short decay time of the fast component of the electroluminescence process gives rise to a very steep rising edge of the S2 pulse shape and a steep fall off at $t = T$. If we are able to find events with very little diffusion in DarkSide-50, we can obtain a precise estimate of T by fitting these two sharp features of the S2 pulse shape.

However, in DarkSide-50, we know that there is a strong radial dependence of the S2 signal. (See Sec. 3.5.2.) The strength of the electroluminescence field drastically increases towards the center of the TPC, which is evident both from the magnitude of S2 and from the gas pocket drift time T . We present here the method for determining $T(r)$, the radially dependent gas pocket drift time. We accomplish this by finding a set of events with little to no diffusion of the ionization electron cloud and precisely measuring T by fitting the S2 pulse shape. As mentioned in Sec. 4.2.3, there is a degeneracy with t_0 and σ , but for

events with little to no diffusion, there is little smearing, so the degeneracy is essentially nonexistent. We use events from the very top portion of liquid argon in the TPC, including in the liquid above the grid, requiring the liquid drift time t_{drift} to be $<5 \mu\text{s}$. We will call these “zero-diffusion” events, and an example is shown in Fig. 4.7. We fit the S2 pulse shape to zero-diffusion events without worrying about the degeneracy of the parameters to obtain the radial dependence of T . Once $T(r)$ is known, we can use it as an input to the diffusion analysis, where the smearing is significant and the degeneracy among T , t_0 and σ must be broken.

4.4.1 Data selection

We search for zero-diffusion events in the high statistics AAr data. Because we will fit the S2 pulse shape to individual events, we generally require higher energy events than those used in the WIMP search analysis. The G2 trigger used throughout much of the AAr WIMP search campaign suppresses a lot of the higher energy events, so we restrict this analysis to data without the G2 trigger. In addition, as will be shown in the next section, triggers on the residuals of previous triggers are a nuisance to this analysis. In the earliest days of DarkSide-50, the trigger inhibit was the same length as the acquisition, allowing for many re-triggers. We therefore restrict ourselves to runs with an extended trigger inhibit.

The general guideline for finding zero-diffusion events is to look for good single scatter events from the top of the TPC. More precisely, we look for events with a physical S1, a physical S2, and t_{drift} very small, typically $<5 \mu\text{s}$. The first challenge is that DarkArt is limited in its ability to correctly reconstruct such events. The drift times are so short that the reconstruction software often cannot distinguish the S1 from the S2, instead lumping them together as a single pulse. The “deficiency” occurs for drift times up to $\sim 3 \mu\text{s}$. Rather than re-tune DarkArt to identify the events with extremely short liquid drift time, we extract them from existing reconstructed data: we select events where the S1 and S2 were reconstructed as a single pulse.

This is not to say that all zero-diffusion events will have their S1 and S2 reconstructed as

Figure 4.7: Example event with zero diffusion. (a) Portion of the waveform showing S1 and S2. The S1 is at $t = 0$, and the S2 start is at $t = 2 \mu\text{s}$. (b) Fit to S2 portion after re-binning and re-aligning the x-axis to be at the S2 start.

a single pulse. Events with extremely short liquid drift times are sometimes reconstructed as separate pulses. But because the tail of S1 can extend for many microseconds, often reaching the S2, the S2 start time is often inaccurate. Since we can achieve sufficient statistics with just the merged S1+S2 events, we ignore the events with separately reconstructed S1 and S2.

We now give the strategy for selecting events where the S1 and S2 were reconstructed as a single pulse.

- First, we must consider that the S2 can photo-ionize the cathode and produce an echo electroluminescence signal that the reconstruction can identify. The photo-ionization electrons must travel the full length of the TPC, so the echo signal always appears at a fixed time after the S2, corresponding to the maximum drift time. Therefore, we select events with exactly 1 reconstructed pulse or exactly 2 reconstructed pulses but with the 2nd pulse in the expected time window for echo signals. For 2-pulse events, we require $370 \mu\text{s} < \Delta t_{01} < 400 \mu\text{s}$, where Δt_{01} is the time difference between the first and second reconstructed pulse.
- Since S2 is significantly larger than S1, the total pulse area of the first pulse (or only pulse, as the case may be), which includes S1 and S2, should be much larger than typical S1s.
- The first pulse has an extremely small f_{90} because the ratio is taken with respect to the combined S1 and S2 area.
- The first pulse should exhibit an S1 at the front, and since we trigger on S1, the pulse start time should be at the expected trigger time within the waveform.
- Finally, the time to the previous trigger should be exponentially distributed.

Now that we have a basic strategy for finding zero-diffusion events with the S1 and S2 reconstructed as the same pulse, we discover a large background to this search: S2-triggered events. When an S1 occurs in the inhibit of a trigger, the DAQ will trigger on the subsequent

S2. These S2 triggers have many of the same properties of the merged S1+S2 events: in both classes, the first reconstructed pulse is dominated by S2 (in zero-diffusion events, the S1 is at most a 1% addition on the size of the pulse, especially at the high energies that are of interest for the diffusion analysis); both have artificially tiny f_{90} ; and both sometimes induce an echo pulse if the S2 is big enough. To distinguish these two classes, we use f_{90} as a discriminating parameter. For S2-triggers, the prompt window (first 88 ns) of the first pulse is essentially empty, especially for S2 with significant smearing. For merged S1+S2 pulses, the prompt window contains the prompt portion of the S1. In both classes, the total area is dominated by the S2. Therefore, we expect merged S1+S2 pulses to have a slightly higher f_{90} value compared to pure S2 pulses. In the f_{90} vs. total PE plane, shown in Fig. 4.8 there is a clear separation of events into two populations. We use an empirically derived function to separate the two populations: Accept events with

$$f_{90} > 0.0004 + 0.004 \exp(-\text{NPE}/6000) \quad (4.15)$$

To verify that we have successfully isolated zero-diffusion events from S2-triggered events, we examine the distributions of Δt to the previous trigger and the first pulse start time. We expect that for zero-diffusion events, the live-time (which is essentially equivalent to Δt) should be exponentially distributed, and the first pulse start time should be in the expected trigger window for typical bulk events. Meanwhile, S2-triggered events will have very short live-time because their S1 occurs in the inhibit of the previous event. In addition, the first pulse start time of S2-triggered events will generally fall outside the expected trigger window. Roughly speaking, the DAQ triggers when there are >3 PE in a 100 ns window, while the DarkArt pulse finder “triggers” when there are >5 PE in a $2 \mu\text{s}$ window, and since S2 has a relatively slow rise time (compared to S1), the DAQ, with its very short coincidence window, triggers deeper in the S2 pulse, whereas the pulse finder, with its wide “coincidence window”, triggers earlier in the S2 pulse. The end result is that for S2-triggered events, the pulse-finder will typically find the start of the first pulse to be *before* the time of the trigger. As can be seen in Fig. 4.9, events passing the f_{90} -based selection criterion are consistent with being

Figure 4.8: f_{90} vs. total PE plane for candidate zero-diffusion events. The line shows the function used to separate events with merged S1+S2 pulses (above the line) vs. S2-triggered events (below the line).

Figure 4.9: Cross checks that f_{90} -based selection (Eqn. 4.15) successfully isolates zero-diffusion events. Green: events passing the f_{90} cut. Red: events failing the f_{90} cut.

zero-diffusion events, while events failing are consistent with being S2-triggers. Though the cuts are all correlated, we apply all three of the f_{90} , live-time, and trigger time cuts.

The final set of cuts used for selecting zero-diffusion events is listed here:

- 38 channels
- baselines found
- event $\Delta t > 1.35$ ms
- live-time < 1 s
- good inhibit time¹
- saturation cut: reject events with any ADC saturation anywhere
- trigger time cut: accept events with first pulse start within the expected trigger window
- # pulses cut: accept events with exactly 1 pulse or events with 2 pulses and the 2nd pulse start time is within the expected window for an echo pulse (with respect to the first pulse start time)

¹During these runs (which were *before* the use of the G2 trigger) the inhibit time for some events is much longer than expected. We don't know why this happens, but the fraction of such events is small. To be cautious, we remove long inhibit time events from this analysis.

Figure 4.10: Onset of ADC saturation in merged S1+S2 zero-diffusion events.

- f_{90} cut: keep events with first pulse f_{90} satisfying Eqn. 4.15
- PE range cut: ask for first pulse size to be within 8×10^3 PE to 50×10^3 PE.

We select candidate zero-diffusion events with first pulse size greater than 8×10^3 PE to ensure good PE statistics to fit each event. We restrict the first pulse size to be less than 50×10^3 PE to avoid ADC saturation, as shown in Fig. 4.10.

4.4.2 Fitting procedure

Now that we can select zero-diffusion events, we are ready to fit the S2 pulse shape to each one. We fit Eqn. 4.8 to each event, using the maximum likelihood method described in Sec. 4.3. We describe here further details specific to fitting the S2 pulse shape to zero-diffusion events.

Because of the merged S1+S2 search strategy described above, we know that each candidate zero-diffusion event begins with an S1 signal. The first step is to better estimate the S2 start time, which we do with a simple threshold analysis on the sum waveform. We begin

the threshold search 200 ns after the pulse start in order to avoid the S1, and we estimate the S2 start time to be when the waveform reaches 6 PE. We then shift the x-axis, putting $t = 0$ at the new S2 start time. The fits are then performed in the region $-0.1 \mu\text{s}$ to $20 \mu\text{s}$ with respect to the new x-axis. We use a very small window prior to the S2 start because we know the tail of S1 is present in that region.

Of the 8 parameters in Eqn. 4.8, we can constrain a few of them before giving the fitter free reign over the waveforms. First, we set $\sigma = 0.01 \mu\text{s}$. Ideally, we would set $\sigma = 0$, but because the fast component lifetime τ_1 is very small, we run the risk of overflow of double floating points (see Sec. 4.5.2.1). With $\sigma = 0.01 \mu\text{s}$, the S2 pulse shape still has very little smearing and the effect of the smearing will not be discernible considering the 10 ns spread of the SPE time profile and the 32 ns binning, described in Sec. 4.3. The 3-parameter degeneracy described in Sec. 4.2.3 is not relevant in zero-diffusion events, but even if there were a degeneracy, it would be broken by fixing σ . Second, the slow component term τ_2 can be “pre-fit” using the tail of each waveform, where the fast component is negligible. We fit a simple exponential decay in the range $9 \mu\text{s}$ to $20 \mu\text{s}$ of each event, then in the full fit, we initialize τ_2 to the fitted value and restrict any change in the parameter to at most $\pm 5\%$ of that value. Third, the overall amplitude A is initialized to the total area of the waveform and any change in the parameters restricted to at most $\pm 10\%$. We have now turned an 8 parameter fit into effectively a 5 parameter fit. The remaining parameters are given sensible initial values and wide limits, as shown in Tab. 4.1. We perform the final fit of the S2 pulse shape over the range $-0.1 \mu\text{s}$ to $20 \mu\text{s}$ for each zero-diffusion event.

4.4.3 Results

We fit the S2 pulse shape to 34,712 zero-diffusion events. The goodness-of-fit is evaluated for each event using the procedure described in Sec. 4.3.2. The reduced χ^2 distribution is shown in Fig. 4.11, zoomed to $\chi_{\text{red}}^2 < 6$. About 15% of events have very poor fits with $\chi_{\text{red}}^2 > 1.5$. Examples of S2 fits are shown in Fig. 4.12. Scanning waveforms by eye, we find that candidate zero-diffusion events with poor fits often exhibit a concave rising edge, rather

Parameter	Initial value	Limits
τ_1	0.01 μs	0.001 μs to 0.1 μs
τ_2	pre-fit in tail	$\pm 5\%$ of pre-fit value
p	0.1	0 to 1
T	1.6 μs	0.5 μs to 10 μs
σ	0.01	fixed
A	area of pulse	$\pm 10\%$ of area
t_0	0	$-10 \mu s$ to $10 \mu s$
y_0	0	-100 PE to 100 PE

Table 4.1: Initial values and limits of fit parameters for zero-diffusion events.

Figure 4.11: Reduced χ^2 of fitted S2 pulse shape to zero-diffusion events.

(a)

(b)

(c)

(d)

Figure 4.12: (a) and (b): Zero-diffusion events with good S2 fits ($\chi_{\text{red}}^2 < 2$). (c) and (d): Events with poor S2 fits ($\chi_{\text{red}}^2 \geq 2$).

than the convex shape expected for the model of S2 described in Sec. 4.2. These may be events occurring very close to the grid surface or near the edge of the TPC, where the electric fields can be highly non-uniform. The longitudinal profile of the ionization electron cloud can then become extremely distorted, giving the unexpected S2 pulse shape. We remove such events from the analysis before proceeding with finding the relationship between T and r .

Furthermore, there are some events where the S1 and S2 signals were so close to each other that they are essentially indistinguishable. In these cases, the re-evaluated S2 start time (described in Sec. 4.4.2) is inaccurately placed at exactly 200 ns after the pulse start. The fit relies on having a small portion of the baseline in front of the rising edge of the S2, so such events have unreliable fits. We require that events have a drift time t_d (difference between the S1 start time and re-evaluated S2 start time) of $>0.21 \mu\text{s}$. Similarly, when t_0 is found to be $< -0.1 \mu\text{s}$ (recall that the fit window for the zero-diffusion events is $-0.1 \mu\text{s}$ to $20 \mu\text{s}$), it indicates there is no baseline in front of the S2. Therefore, we require $t_0 > -0.1 \mu\text{s}$.

Taking those events with $\chi_{\text{red}}^2 < 1.5$, $t_d > 0.21 \mu\text{s}$, and $t_0 > -0.1 \mu\text{s}$, we plot T from each fit as a function of transverse position, as shown in Fig. 4.13. The xy positions are evaluated using the merged S1+S2 pulses. The algorithm should be as equally valid on these pulses as on well-reconstructed single scatters since the S1 is small compared to the S2 and the S1 is in very nearly the same location as the S2 (in fact the same xy location and only a small z difference). The relationship between T and r is consistent with a non-uniform electroluminescence field that is strongest at the center of the TPC and gradually becomes weaker towards the edge. There are several possible explanations, including a sagging anode window or a deflecting grid, but it is non-trivial to use these results to disentangle the possibilities. Evidently, the T vs. r^2 profile is well fit by a piece-wise linear function. A priori, we have no reason to expect that it should or should not be of this form. We take the function $T(r)$ to be of the form:

$$T(r) = \begin{cases} C_T + \frac{B_T - C_T}{A_T} r^2 & , \text{ if } r^2 < A_T \\ B_T & , \text{ if } r^2 \geq A_T \end{cases} \quad (4.16)$$

(a)

(b)

Figure 4.13: (a) 2D profile of average T from fits vs x and y . (b) 2D histogram of T vs. r^2 . The histogram is profiled and fit with a piece-wise linear function. The profile is in black and barely visible under the fit (red). Events with low T at large r may arise due to distortions of the electroluminescence field near the TPC edge.

Figure 4.14: Distribution of slow component lifetime τ_2 , extracted from fits of S2 in zero-diffusion events.

Fitting Eqn. 4.16 to the profile of the T vs. r^2 distribution, we find $A_T = 240 \text{ cm}^2$, $B_T = 1.49 \text{ } \mu\text{s}$, and $C_T = 0.91 \text{ } \mu\text{s}$. We could use these results to estimate the gas pocket height as a function of radius, but we will not cover such an analysis in this document. The flat step at large r^2 ($>250 \text{ cm}^2$) may be due to the xy reconstruction, which is known to have difficulty reconstructing edge events.

The fits to the zero-diffusion events can also tell us about the slow component lifetime τ_2 and the fast component fraction p in the gas. Because of the 32 ns binning of the waveforms, we do not have the resolution to estimate the fast component lifetime τ_1 . The distribution of τ_2 is shown in Fig. 4.14. Evidently, the average slow component lifetime is $\tau_2 = 3.43 \text{ } \mu\text{s}$. The distribution of p vs. transverse position is shown in Fig. 4.15. The fast component fraction depends on the electroluminescence field in the gas. Since the electroluminescence field varies significantly in DarkSide-50, so does p . As the electroluminescence field becomes weaker at large radius, the fast fraction increases, as expected. Evidently the relationship

Figure 4.15: 2D histogram of the fast component fraction p vs. r^2 . The histogram is profiled and fit with a piece-wise function. The profile is in black and the fit to the profile is in red.

between p and r is well fit by a piece-wise function of the form

$$p(r) = \begin{cases} C_p + \frac{B_p - C_p}{A_p^2} r^4 & , \text{ if } r^2 < A_p \\ B_p & , \text{ if } r^2 \geq A_p \end{cases} \quad (4.17)$$

Fitting Eqn. 4.17 to the profile of the p vs. r^2 distribution we find $A_p = 215 \text{ cm}^2$, $B_p = 0.089$, and $C_p = 0.075$.

4.5 Electron diffusion

We now apply the techniques of fitting the S2 pulse shape to measure electron diffusion. As a cloud of ionization electrons drifts through the liquid, random fluctuations of the (thermalized or nearly thermalized) electrons will cause the cloud to diffuse over time. The diffusion in the longitudinal and transverse directions (relative to the drift direction) need not be the same. In DarkSide-50, we are sensitive to the longitudinal diffusion, which manifests as a smearing of the S2 pulse shape: for more longitudinally diffuse electron clouds, the turn-

on and turn-off of the production of argon excimers in the gas pocket as the cloud is extracted from the liquid surface and drifted into the anode surface is “softened”, corresponding to washing out of the sharp features of the S2 pulse shape. We can fit the S2 pulse shape to individual events and extract the amount of smearing for each. The electron cloud should be more diffuse for events deeper in the TPC, so the smearing should increase as a function of drift distance in the liquid.

If we assume that the initial size of a cloud of ionization electrons is negligibly small compared to the eventual size due to diffusion, then, letting the electrons originate from a point $(x, y, z) = (0, 0, 0)$ at time $t = 0$, their distribution after drifting a time t_d is given by [96]

$$n(\rho, z, t_d) = \frac{n_0}{4\pi D_T t_d \sqrt{4\pi D_L t_d}} \exp\left(-\frac{\rho^2}{4D_T t_d}\right) \exp\left(-\frac{(z - vt_d)^2}{4D_L t_d}\right) \quad (4.18)$$

where n_0 is the number of initial ionization electrons, v_d is the drift velocity in the liquid, D_T is the transverse diffusion coefficient, D_L is the longitudinal diffusion coefficient, $\rho^2 = x^2 + y^2$, and z is defined parallel to the drift direction. We assume n_0 is constant for each event, i.e. no loss of electrons by electronegative attachment to impurities. Even if we account for the electron drift lifetime, the S2 pulse shape should not be affected except for an overall scale factor.

From Eqn. 4.18, we see that the longitudinal profile of the electron cloud is a Gaussian wave which broadens over time:

$$\sigma_L^2 = 2D_L t_d \quad (4.19)$$

where σ_L is the width of the wave. When the width of the wave grows slowly compared to the drift velocity in the liquid, the diffusion of the electron cloud manifests as a simple Gaussian smearing of the S2 pulse shape. The goal of this analysis is to measure D_L , which we do by evaluating the smearing σ as a function of drift time t_d for many events. The smearing is extracted from the S2 fit and the drift time comes directly from the reconstruction. We will motivate a method to break the degeneracy of the S2 pulse shape and take into consideration the variation of the S2 rise time T and fast component fraction p with respect to transverse position.

4.5.1 Data selection

The principle data used for this analysis are the abundant ^{39}Ar decays from AAr data at standard 200 V/cm drift field and 2.8 kV/cm extraction field, the same data set used in Sec. 4.4. As before, we require relatively high energy events to ensure enough PE statistics in S2 to perform the fit, so we avoid G2 triggered data. The use of the same data in these two sections is also a matter of convenience. We use additional sets of data as cross-checks and systematics measurements of the diffusion, including data at different drift fields (100 V/cm and 150 V/cm) and different extraction field (2.3 kV/cm).

To perform the measurement of the diffusion constant, we use well-reconstructed single scatter ^{39}Ar events. The event selection criteria are therefore very similar to the WIMP search criteria, though we can loosen many of the cuts since we are performing a statistical analysis.

The cuts used are as follows:

- 38 channels
- baselines found
- live-time + inhibit time > 1.35 ms
- live-time < 1 s
- good inhibit time (for details, see the same cut in Sec. 4.4.1.)
- trigger time: accept events with first pulse start time within expected trigger time window
- saturation cut
- # of pulses: accept 2-pulse events or 3-pulses events with S3
- r slice: r in $[9, 12]$ cm
- S2 slice: S2 in $[40\text{k}, 50\text{k}]$ PE

To reduce possible systematics due to variations of T with respect to r , we select events in a narrow r slice: 9 cm to 12 cm. Finally, we select events with maximum possible PE statistics: 40×10^3 PE to 50×10^3 PE. We use uncorrected S2 for the PE range cut. The selected events have a mean S1 of 1000 PE with RMS 150 PE, corresponding to (140 ± 20) keV electron recoils. We will repeat the analysis on different r and S2 slices to estimate the systematics.

4.5.2 Fitting procedure

We perform a fit of the S2 pulse shape on every event that passes the event selection. There are 8 parameters in the fit (Eqn. 4.8), as described in Sec. 4.2 and 4.3. We must break the approximate degeneracy between T , σ , and t_0 . We do so by fixing T . Here we describe the choice of initial values and limits for each of the 8 parameters.

- As shown in Sec. 4.4, T varies with transverse position. We fix T on an event-by-event basis, evaluating $T(r)$ as given by Eqn. 4.16.
- For each event, we pre-determine the value of the baseline offset y_0 by fitting a flat line to the pre-signal region of $-5 \mu\text{s}$ to $-1 \mu\text{s}$. The baseline is expected to be consistent with zero by nature of the DarkArt baseline finder. The baseline value in the full fit is fixed to the value determined here.
- The fast component lifetime should be independent of t_d . However, τ_1 cannot be well-constrained due to the resolution of our waveforms. In any case, the fit should be insensitive to the exact value of τ_1 (because the fast component is washed out with any non-negligible amount of smearing), so we fix $\tau_1 = 0.01 \mu\text{s}$.
- The slow component lifetime should also be independent of t_d , but since the quality of the fits is sensitive to the value of τ_2 (it is the principle shape parameter in the long tail of S2), we do not fix it globally. Instead, we determine τ_2 prior to the full S2 fit by fitting an exponential to the tail of the S2 pulse in the region $9 \mu\text{s}$ to $20 \mu\text{s}$. This range is chosen to guarantee that we fit the exponential to the tail of the S2 pulse even in

events with the highest diffusion, where the peak is farthest from the pulse start. The fit function is $y = Ae^{-t/\tau}$. We constrain the value of τ_2 in the full S2 fit to be within $\pm 5\%$ of the fitted value of τ .

- We do not expect the fast component fraction to vary with respect to t_d , but it varies with electroluminescence field. In DarkSide-50, the electroluminescence field varies quite a bit, and, correspondingly so does p with respect to transverse position. Like T , we fix p on an event-by-event basis, evaluating $p(r)$ as given by Eqn. 4.17. In any case, the fast component is washed out in the pulse shape when smearing is non-negligible, so the S2 fits should not be very sensitive to the value of p .
- The initial value of σ is given by the ICARUS value of diffusion, $D_L = 4.8 \text{ cm}^2/\text{s}$ [97]. More specifically, $\sigma_{\text{init}} = \sqrt{2D_L t_d}/v$. The limits for σ are $0.01 \text{ } \mu\text{s}$ to $10 \text{ } \mu\text{s}$. Drift times in DarkSide-50 vary up to hundreds of microseconds. In this analysis, we consider drift times from tens of microseconds to $380 \text{ } \mu\text{s}$ for 200 V/cm drift field data to $700 \text{ } \mu\text{s}$ for 100 V/cm drift field data, corresponding to expected values of σ from $0.1 \text{ } \mu\text{s}$ to $0.7 \text{ } \mu\text{s}$ to $0.9 \text{ } \mu\text{s}$, well within the limits of the parameter in the fit.
- The amplitude parameter A is initialized to the size of the pulse and is allowed to vary within $\pm 10\%$.
- The time offset parameter t_0 is expected to vary with each event: for events with more diffusion, the DarkArt pulse finder will find the pulse start relatively earlier with respect to the pulse peak. From initial attempts at fitting many events, we empirically find that t_0 varies linearly with σ : $t_0 = -0.25 + 3.06\sigma$, which we use to set the initial value of the time offset: $t_{0,\text{init}} = t_0(\sigma_{\text{init}})$. We leave t_0 as a free parameter in the fits and give it wide limits of $-10 \text{ } \mu\text{s}$ to $10 \text{ } \mu\text{s}$.

The initial values and limits of all the fit parameters are summarized in Table 4.2. Of the 8 parameters in the fit, 6 of them are fixed or tightly constrained in the final fit of the S2 pulse shape. The remaining free parameters are σ and t_0 .

Parameter	Initial value	Limits
τ_1	0.01 μs	fixed
τ_2	pre-fit in tail	$\pm 5\%$ of pre-fit value
p	$p(R)$	fixed
T	$T(R)$	fixed
σ	$\sqrt{2D_L t_{\text{drift}}}/v_{\text{drift}}$	0 μs to 10 μs
A	area of S2 pulse	$\pm 10\%$ of S2 area
t_0	$\max(-0.25 + 3.06\sigma_{\text{init}}, 0)$	-10 μs to 10 μs
y_0	pre-fit in pre-signal region	fixed

Table 4.2: Initial values and limits of fit parameters.

For each event, we re-define the x-axis such that $t = 0$ is at the S2 pulse start, as determined by DarkArt. We then truncate the waveform leaving only the $-5 \mu\text{s}$ to $20 \mu\text{s}$ region about the newly defined $t = 0$. We re-bin the truncated waveform to 32 ns binning, convert into a TH1F, and fit the S2 pulse shape with options LMR+.

4.5.2.1 Floating point precision

We must take extra care when numerically fitting the S2 pulse shape due to the very small fast component lifetime. As can be seen in Eqn. 4.7, when τ is very small, as is the case for the 11 ns lifetime of the fast component in gas, the exponential term can blow up, easily surpassing the maximum double floating point precision ($10^{\pm 308}$). To understand where the overflow errors can occur, we examine Eqn. 4.7 more closely. We do not worry about the erf term because it is always finite. The exponential term, which we combine into $\exp(\sigma^2/2\tau^2 - t/\tau)$, can be very large, but since it is modulated by the erfc term, the product is always finite. We verify this numerically, as follows. Let

$$z_1 = \sigma^2/2\tau^2 - t/\tau \tag{4.20}$$

$$z_2 = (\sigma^2 - t\tau)/\sqrt{2}\sigma\tau \tag{4.21}$$

which are the arguments of the exponential term and the erfc term, respectively. We evaluate z_1 and z_2 throughout the parameter space spanned by $\{-10 < t < 30, 0.001 < \tau < 10, 0.01 < \sigma < 2\}$, which encompasses the regimes of interest for this analysis. We find that for $z_1 > 308 \ln(10) = 709$, we always have $z_2 > 10$. For the sake of nice numbers, we round 709 to 700. Now, the complementary error function has the following asymptotic approximation at $z = \infty$ [98]:

$$\text{erfc}(z) \approx \frac{e^{-z^2}}{\sqrt{\pi}z} \quad (4.22)$$

At $z > 10$, erfc is approaching zero and the approximation is $>99\%$ accurate. Therefore, we find that when the exponential term is large, we can use the asymptotic approximation of erfc. Carrying out the product of the exponential and erfc terms with the asymptotic approximation, the exponential term completely cancels, leaving behind a Gaussian:

$$\exp(z_1 > 700) \text{erfc}(z_2) \approx e^{-t^2/2\sigma^2} \frac{1}{\sqrt{\pi}z_2} \quad (4.23)$$

which is finite. We do not worry about the exponential term hitting the other end of the floating point limit ($z_1 < -709$) because in this case, we are multiplying a tiny number by some finite value of the erfc term, so the product will be essentially zero.

In summary, when $z_1 > 700$, we use the asymptotic approximation of erfc; otherwise, we use the full form of Eqn. 4.7. The approximation has been verified to be valid in the parameter space of interest for this analysis. Therefore, we are guaranteed to avoid double floating point overflow errors, while maintaining numerically accurate representations of the S2 pulse shape.

4.5.3 Results

Using the data described in Sec. 4.5.1, there are 89.5×10^3 events that pass the cuts described in Sec. 4.5.1. We fit the S2 pulse shape to each one. Fig. 4.16 shows examples of some of the fits. Most events have good quality fits, with 94.5% having a reduced χ^2 smaller than 1.5, as shown in Fig. 4.17.

To understand the diffusion of the ionization electron cloud, we extract the smearing parameter σ for each event. First, we convert the smearing to a length scale. Assuming no

(a) Event from near the top of the TPC.

(b) Event from near the bottom of the TPC.

Figure 4.16: Examples of S2 pulse shape fits for the electron diffusion measurement. Left: sum channel of raw waveforms. Right: Fits to S2.

Figure 4.17: Reduced χ^2 of S2 pulse shape fits to events in diffusion analysis.

additional smearing of the S2 pulse shape in the electroluminescence region (i.e. there is little diffusion in the gas phase), the physical length σ_L of the electron cloud just below the grid is related to the fit parameter σ via $\sigma_L = v_d \sigma$, where v_d is the drift velocity in the liquid. From Eqn. 4.19 we expect that σ_L^2 should be linearly proportional to t_d , which is verified in Fig. 4.18. The diffusion constant is then easily evaluated by fitting a line to the profile of the σ_L^2 vs. t_d distribution:

$$\sigma_L^2 = \sigma_0^2 + 2D_L t_d \quad (4.24)$$

where the σ_0 term accounts for any systematic smearing independent of drift time, possibly due to a small amount of diffusion in the gas phase or non-negligible initial size of the ionization electron cloud. In DarkSide-50, σ_0 is small relative to σ_L . Fitting to the t_d region of $50 \mu\text{s}$ to $300 \mu\text{s}$, the diffusion constant is found to be $D_L = 4.31 \text{ cm}^2/\text{s}$. The uncertainty from the fit is negligible due to the high statistics. The total uncertainty on D_L is systematics dominated.

Figure 4.18: σ^2 vs. t_{drift} for measurement of electron diffusion. We extract the Gaussian smearing term σ from the S2 pulse shape fits, convert to length scale via $\sigma_L = v_d \sigma$.

4.5.4 Systematics

We estimate the systematic uncertainty on the diffusion coefficient in a few different ways. We vary the fit parameters for extracting the diffusion constant, and we repeat the analysis on various data sets. We use different r and S2 slices from the same set of runs used to produce the results of the previous section, as well as data taken at different drift and extraction fields.

4.5.4.1 Vary fit range

The value of the diffusion constant is sensitive to the range of t_d used in the linear fit. Using various fit windows within the t_d range of 50 μs to 300 μs , we find the diffusion constant varies by $\pm 5\%$. Upon further inspection, we empirically find that the two parameter function $\sigma_L^2 = (\sigma_0 + \sqrt{2D_L t_d})^2$ fits the data extremely well. We do not yet know why the widths prefer to be added linearly and not quadratically.

4.5.4.2 Vary r and S2 slices

Ideally, D_L should be independent of r and $S2$. The analysis chain is applied identically to the same runs (5496 to 7262 of the 50-day campaign) using the same cuts, but selecting events in different r and S2 slices. We choose 8 additional slices:

- r in the ranges $[0,3)$, $[3,6)$, $[6,9)$, $[12,15)$, $[15, 18)$ cm all with S2 in the range $[40e3, 50e3]$ PE.
- S2 in the ranges $[10e3, 20e3)$, $[20e3, 30e3)$, $[30e3, 40e3)$ PE all with r in the range $[9,12)$ cm.

The event-by-event S2 fit procedure is identical to Sec. 4.5.2. The results are shown in Fig. 4.19. The reduced χ^2 of all slices are good (<1.5), and the extracted diffusion constants agree to within 4% for the various r slices and within 6% for the various S2 slices. There is a systematic bias for larger D_L with respect to larger r and larger S2. The source of the bias is unknown. One might expect that towards the edge of the TPC, the electrons spend more time in the gas pocket due to the weaker electroluminescence field (and maybe also a physically larger gap to traverse) and therefore would undergo more diffusion. But that should appear as an overall offset to the σ vs. t_d distribution. Instead, we see both increasing offset and slope at increasing r . We do not yet have any plausible explanation for the increasing diffusion with respect to S2.

4.5.4.3 Vary drift field

At low drift electric fields as in DarkSide-50, the electrons are thermal (i.e. have no extra energy from the field) and diffusion of the electron cloud should, in principle, be a direct function of the time spent in liquid (and not of the speed of the drift). Therefore, we should be able to repeat the analysis on data at different drift fields, and, after normalizing by the drift velocity, the measured diffusion constants should be consistent with each other. We use atmospheric argon background data taken at two different drift fields, 100 V/cm and 150 V/cm, to compare to nominal 200 V/cm drift field data. All are taken with 2.8 kV/cm

(a)

(b)

Figure 4.19: (a) Diffusion measurement using various R slices with identical S2 slice. (b) Diffusion measurement using various S2 slices with identical R slice.

extraction field.

The event selection criteria are nearly identical to those used in the main analysis; we must adjust the npulses cut to reflect that S2-echo occurs at a different location within the waveform. Furthermore, due to reduced statistics, we take a wider slice in the r vs. S2 plane: for both 100 V/cm and 150 V/cm data, we use r in the range 0 cm to 18 cm, and S2 in the range 10×10^3 PE to 50×10^3 PE. The event-by-event fit procedure is identical to that of the standard drift field data. In particular, since the electroluminescence field should not be affected by the drift field, we use the same $T(r)$ and $p(r)$ functions given by Eqn. 4.16 and 4.17, respectively. The results are shown in Fig. 4.20. After normalizing by the drift velocities, which are tabulated for the different drift fields in Tab. 4.3, we find that the diffusion constants, given in Tab. 4.4, agree with each other to within 3%.

Drift [V/cm]	v_d [mm/ μ s]
100	0.53
150	0.75
200	0.93

Table 4.3: Drift velocities in liquid for different drift fields in DarkSide-50.

4.5.4.4 Vary extraction field

We can also obtain a handle on the systematics of the measurement of the diffusion constant by looking at data taken at different extraction fields. The diffusion constant should be independent of the extraction field.

Due to operational constraints, high statistics data were taken at only one other extraction field, 2.3 kV/cm. We repeat the analysis chain applied to standard extraction field data, but we must regenerate the $T(r)$ and $p(r)$ functions: The electron drift time across the pocket and the fast component fraction depend on the electroluminescence field. We repeat the analysis of Sec. 4.4 with no modifications. The T and p distributions change, as shown in Fig. 4.21, but still appear to be consistent with the piece-wise functional forms of

(a)

(b)

Figure 4.20: (a) Results of diffusion measurement for data at different drift fields. (b) After normalizing for drift velocity

Eqn. 4.16 and 4.17. We find $A_T = 215 \text{ cm}^2$, $B_T = 1.64 \text{ } \mu\text{s}$, $C_T = 1.13 \text{ } \mu\text{s}$ and $A_p = 247 \text{ cm}^2$, $B_p = 0.154 \text{ } \mu\text{s}$, $C_p = 0.085 \text{ } \mu\text{s}$ for 2.3 kV/cm extraction field data.

With new $T(r)$ and $p(r)$ functions in hand, we repeat the analysis chain of Sec. 4.5.1 and 4.5.2 and determine the σ_L^2 vs. t_d distribution. Due to the lower statistics relative to standard field data, we again extend the R and S2 slices to include 0 cm to 18 cm and 10×10^3 PE to 50×10^3 PE, respectively. We keep the S2 slice the same as in the measurements for the various drift field data in order to keep the PE statistics consistent with the other data sets. With the reduced electroluminescence field, we are probing a slightly higher energy range. The profile of the resulting σ^2 vs. t_{drift} distribution is shown in Fig. 4.22. We see that there is an overall shift in the distribution, which is expected since, with the lower electroluminescence field, the electrons spend more time in the gas. The slope, and therefore also D_L , is consistent with the results of other data sets.

4.5.4.5 Summary of systematics

The longitudinal diffusion constant extracted from the various data sets are summarized in Tab. 4.4.

We do not give statistical uncertainties on D_L or σ_0 because they are small due to the high statistics in the fits. The average value of the diffusion constant is $D_L = (4.3 \pm 0.2) \text{ cm}^2/\text{s}$, where the uncertainty is systematics dominated.

4.5.5 Comparison to literature

In Ref. [99], the longitudinal diffusion is parametrized by the electron energy

$$\epsilon = \frac{kT}{e}. \quad (4.25)$$

At low electric fields, as in DarkSide-50, the Einstein-Smoluchowski relation holds

$$D = \frac{kT}{e} \mu \quad (4.26)$$

where μ is the electron mobility, defined as $\mu = v_d/E$. At very low fields, the electrons are thermal and ϵ is constant with $kT = 0.0075 \text{ eV}$. As the field is increased, the electrons are

(a)

(b)

Figure 4.21: Profile of (a) T vs. r^2 and (b) p vs. r^2 distributions for standard 2.8 kV/cm extraction field data (blue) and 2.3 kV/cm extraction field data (red).

Figure 4.22: Profile of σ_L^2 vs. t_{drift} for 2.3 kV/cm extraction field data (red) and standard 2.8 kV/cm extraction field data (blue).

no longer thermal, and $kT = eD/\mu$ is taken to define the electron temperature. The effective electron energy associated with the longitudinal diffusion is given by

$$\epsilon_L = \frac{D_L}{\mu}. \quad (4.27)$$

Since μ is field-dependent, we evaluate ϵ_L separately for the different drift field data. The drift velocities at each field are given in Tab. 4.3. The results are shown in Fig. 4.23, along with results from other experiments and models.

4.6 Derivations of analytic S2 pulse shape

4.6.1 Idealized S2 pulse shape

We need to do the following convolution:

$$y_{\text{ideal}}(t; \tau_1, \tau_2, p, T) = f(t; T) * g(t; \tau_1, \tau_2, p) \quad (4.28)$$

Drift [V/cm]	Extr. [kV/cm]	R [cm]	S2 [10^3 PE]	D_L [cm^2/s]	σ_0^2 [cm^2]
200	2.8	[0, 3]	[40, 50]	4.241	2.546e-04
200	2.8	[3, 6]	[40, 50]	4.262	2.556e-04
200	2.8	[6, 9]	[40, 50]	4.274	2.585e-04
200	2.8	[9, 12]	[40, 50]	4.312	2.748e-04
200	2.8	[12, 15]	[40, 50]	4.382	2.839e-04
200	2.8	[15, 18]	[40, 50]	4.414	3.292e-04
200	2.8	[9, 12]	[30, 40]	4.261	2.492e-04
200	2.8	[9, 12]	[20, 30]	4.177	2.252e-04
200	2.8	[9, 12]	[10, 20]	4.074	1.929e-04
150	2.8	[0, 18]	[10, 50]	4.416	2.792e-04
100	2.8	[0, 18]	[10, 50]	4.378	2.801e-04
200	2.3	[0, 18]	[10, 50]	4.345	3.336e-04

Table 4.4: Diffusion constant D_L measured from different data with different drift and extraction fields.

where $f(t; T)$ is a uniform distribution on $[0, T]$:

$$f(t; T) = \frac{1}{T} \begin{cases} 0, & t < 0 \\ 1, & 0 \leq t \leq T \\ 0, & t > T \end{cases} \quad (4.29)$$

and $g(t; \tau_1, \tau_2, p)$ is the standard two-component exponential distribution:

$$g(t; \tau_1, \tau_2, p) = p \cdot h(t; \tau_1) + (1 - p) \cdot h(t; \tau_2) \quad (4.30)$$

with

$$h(t; \tau) = \begin{cases} 0, & t < 0 \\ \frac{1}{\tau} e^{-t/\tau}, & t \geq 0 \end{cases}. \quad (4.31)$$

The fast and slow component terms demand identical computations, so for convenience, we need only consider the case of $p = 1$. Then we let $\tau = \tau_1$. We now proceed with the

Figure 4.23: Effective electron energy ϵ_L vs. drift field. Li data and parametrization are from [99] and ICARUS data are extracted from [97]. The model is that of Atrazhev and Timoshkin [100].

convolution.

$$I(t; \tau, T) = y_{\text{ideal}}(t; \tau, \tau_2, p = 1, T) = \int_{-\infty}^{\infty} f(s)h(t-s) ds = \frac{1}{T} \int_0^T h(t-s) ds \quad (4.32)$$

We split the integral into three parts:

$$I(t) = I_1(t) + I_2(t) + I_3(t) \quad (4.33)$$

where

$$I_1 = I(t < 0) = 0 \quad (4.34)$$

$$I_2 = I(t > T) \quad (4.35)$$

$$I_3 = I(0 \leq t \leq T). \quad (4.36)$$

The first integral vanishes because for $t < 0$ and $s > 0$, $h(t-s < 0)$ is zero. For the second integral, we always have $t > s$, so

$$\begin{aligned} I_2(t > T; \tau, T) &= \frac{1}{T} \int_0^T \frac{1}{\tau} e^{-(t-s)/\tau} ds \\ &= \frac{1}{T} \left(e^{-(t-T)/\tau} - e^{-t/\tau} \right). \end{aligned} \quad (4.37)$$

For the third integral:

$$\begin{aligned}
I_3(0 \leq t \leq T; \tau, T) &= \frac{1}{T} \int_0^T h(t-s; \tau) ds \\
&= \frac{1}{T} \int_0^t h(t-s; \tau) ds + \frac{1}{T} \int_t^T h(t-s; \tau) ds \\
&= \frac{1}{T} \int_0^t \frac{1}{\tau} e^{-(t-s)/\tau} ds \\
&= \frac{1}{T} (1 - e^{-t/\tau})
\end{aligned} \tag{4.38}$$

where we have used that $h(t-s < 0; \tau) = 0$ in the second line. Finally, we have

$$I(t; \tau, T) = \frac{1}{T} \begin{cases} 0, & t < 0 \\ 1 - e^{-t/\tau}, & 0 \leq t \leq T \\ e^{-(t-T)/\tau} - e^{-t/\tau}, & t > T \end{cases} \tag{4.39}$$

which is exactly Eqn. 4.3. Noting that

$$y_{\text{ideal}}(t; \tau_1, \tau_2, p, T) = p \cdot I(t; \tau_1, T) + (1-p) \cdot I(t; \tau_2, T) \tag{4.40}$$

we recover Eqn. 4.2.

4.6.2 S2 pulse shape with Gaussian smearing

We need to carry out the following convolution:

$$y(t; \tau_1, \tau_2, p, T, \sigma) = y_{\text{ideal}}(t; \tau_1, \tau_2, p, T) * k(t; 0, \sigma) \tag{4.41}$$

where $k(t; 0, \sigma)$ is a Gaussian pdf with mean 0 and width σ . We again need only consider the case of $p = 1$ for convenience, and let $\tau = \tau_1$. Carrying out the integrals:

$$J(t; \tau, T, \sigma) = y(t; \tau, \tau_2, p = 1, T, \sigma) \tag{4.42}$$

$$= I(t; \tau, T) * k(t; 0, \sigma) \tag{4.43}$$

$$= \int_0^\infty I(s; \tau, T) k(t-s; 0, \sigma) ds \tag{4.44}$$

$$= \frac{1}{T} \int_0^T (1 - e^{-s/\tau}) k(t-s) ds + \frac{1}{T} \int_T^\infty (e^{-(s-T)/\tau} - e^{-s/\tau}) k(t-s) ds \tag{4.45}$$

$$= \underbrace{\frac{1}{T} \int_0^T k(t-s) ds}_{J_1} - \underbrace{\frac{1}{T} \int_0^T e^{-s/\tau} k(t-s) ds}_{J_2} + \underbrace{\frac{e^{T/\tau} - 1}{T} \int_0^\infty e^{-s/\tau} k(t-s) ds}_{J_3} \tag{4.46}$$

where $I(t; \tau, T)$ is defined in Eqn 4.32. With repeated change of variables, we can identify the integrals as error functions or complementary error functions, and we find ²

$$J_1 = \frac{1}{T} \left[\operatorname{erf} \left(\frac{t}{\sqrt{2}\sigma} \right) - \operatorname{erf} \left(\frac{t-T}{\sqrt{2}\sigma} \right) \right] \quad (4.47)$$

$$J_2 = \frac{1}{2T} \exp \left(\frac{\sigma^2 - 2t\tau}{2\tau^2} \right) \left[\operatorname{erf} \left(\frac{t\tau - \sigma^2}{\sqrt{2}\sigma\tau} \right) + \operatorname{erf} \left(\frac{\sigma^2 - \tau(t-T)}{\sqrt{2}\sigma\tau} \right) \right] \quad (4.48)$$

$$J_3 = \frac{e^{T/\tau} - 1}{2T} \exp \left(\frac{\sigma^2 - 2t\tau}{2\tau^2} \right) \operatorname{erfc} \left(\frac{\sigma^2 - \tau(t-T)}{\sqrt{2}\sigma\tau} \right) \quad (4.49)$$

$J(t)$ can now be simplified to

$$J(t; \tau, T, \sigma) = \frac{1}{2T} \left[\operatorname{erf} \left(\frac{t}{\sqrt{2}\sigma} \right) - \operatorname{erf} \left(\frac{t-T}{\sqrt{2}\sigma} \right) - e^{\sigma^2/2\tau^2} e^{-t/\tau} \left\{ \operatorname{erfc} \left(\frac{\sigma^2 - t\tau}{\sqrt{2}\sigma\tau} \right) + \operatorname{erfc} \left(\frac{\sigma^2 - (t-T)\tau}{\sqrt{2}\sigma\tau} \right) \right\} \right] \quad (4.50)$$

We now have an analytic form of the S2 pulse shape:

$$y(t; \tau_1, \tau_2, p, T, \sigma) = p \cdot J(t; \tau_1, T, \sigma) + (1-p) \cdot J(t; \tau_2, T, \sigma) \quad (4.51)$$

which is readily identified as Eqn. 4.4.

² $\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-s^2} ds$ and $\operatorname{erfc}(z) = 1 - \operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-s^2} ds$

CHAPTER 5

UAr Depletion

The DarkSide-50 experiment is the first to perform a dark matter search using underground argon (UAr). As described in Sec. 2.3, the need for underground argon arises from the large contamination of ^{39}Ar in atmospheric argon (AAr). The isotope β -decays with a half-life of 269 yr and a Q-value of 565 keV. The WIMP-induced recoil energy is expected to be in the range of tens to hundreds of keV, so any ^{39}Ar in the target material poses a potentially dangerous background to a WIMP search. Moreover, argon in the atmosphere contains ^{39}Ar at the level of $(8.0 \pm 0.6) \times 10^{-16}$ g/g [80], which induces an overwhelming background for large scale Ar-based WIMP searches. Indeed, with the DarkSide-50 TPC filled with AAr, the trigger rate is 50 Hz and is dominated by ^{39}Ar decays.

The identification of underground sources of argon makes feasible ton scale dark matter detectors using LAr that can reach the neutrino floor. The ^{40}Ar isotope is produced by (α, n) reactions from the decay chains of U and Th, present in the rock underground. Shielded by the atmosphere and rock overburden, the underground argon contains a significantly reduced relative abundance of ^{39}Ar . In 2007, underground sources of argon were discovered in the National Helium Reserve in Amarillo, TX [83] as well as in Doe Canyon, Colorado [84]. The DarkSide-50 TPC is now filled with UAr from the Doe Canyon wells. Prior to DarkSide-50, the best measurement of the ^{39}Ar activity in UAr was only able to place an upper limit (or a lower limit if speaking about the ^{39}Ar reduction factor) [101]. DarkSide-50 has now performed the most sensitive measurement of the ^{39}Ar activity in UAr, to date.

Once DarkSide-50 was filled with UAr, one of the chief first goals was to estimate the ^{39}Ar reduction factor. In this section we discuss some of the early UAr data in DarkSide-50 as they pertain to measuring the ^{39}Ar activity. The activity is estimated by spectral fitting of

the primary scintillation signal. In UAr, the trigger rate is dominated by the γ background arising from the radioactivity of the detector materials. We cannot efficiently distinguish ^{39}Ar β -decays from the dominant γ events, so we must estimate the ^{39}Ar by fitting the β spectrum underneath the other background components. The γ spectra are generated by a Geant4-based Monte Carlo. We describe here in broad strokes how the spectra are produced and the fit procedure used to estimate the ^{39}Ar activity in UAr.

5.1 AAr vs. UAr S1 spectra

The earliest UAr data were taken at null field. The null field S1 spectrum for both UAr and AAr data are shown in Fig. 5.1a. Only very basic cuts are applied to remove, for example, events in which the DAQ was misbehaving or the electronics noise was unusually high. The data are also z -corrected, à la Sec. 3.5.2, using the top-bottom asymmetry of S1 as a proxy for t_{drift} , which increases the energy resolution. The spectra are normalized by live time and mass to show specific activities. The large feature below ~ 4000 PE in the AAr data is due to the high activity of ^{39}Ar decays. By eye, one can see the ^{39}Ar activity is greatly reduced in UAr data. Events beyond the ^{39}Ar endpoint are mostly due to γ -rays. The gamma peaks match extremely well between AAr and UAr data, indicating that the TPC light yield was unchanged between the two eras of DarkSide-50. The matching of the LY (of 7 PE/keV_{ee} at 200 V/cm drift field) was further confirmed by *in situ* ^{83m}Kr calibration.

Figure 5.1b shows the S1 spectra at 200 V/cm drift field. Only basic cuts are applied with the additional requirement of only single scatter events (they have one S1 and one S2). Again, the data are normalized by live time and mass, and again, the ^{39}Ar activity is seen to be reduced from AAr to UAr data. While the high energy gamma peaks are no longer visible in the field-on data because of the fluctuations in the separation of the light and charge signals, the spectra are consistent beyond the ^{39}Ar endpoint.

That the activity beyond the ^{39}Ar endpoint is due to γ -rays can be seen by applying simple veto cuts. The veto cuts are described in more detail in Sec. 6.2.2. For the purposes of this discussion, we need only know that the principle of the veto cuts is to apply anti-

(a)

(b)

Figure 5.1: AAr and UAr S1 spectra at (a) null field and (b) 200 V/cm drift field. The spectra are normalized by live time and by mass.

Figure 5.2: AAr and UAr S1 spectra at 200 V/cm drift field with veto anti-coincidence cuts.

coincidence conditions between the TPC and LSV: signals in the TPC with coincident signals in the LSV are vetoed. The veto anti-coincidence cuts are designed to remove neutron events, but they predominantly remove gammas which scatter (at least) once in the TPC and (at least) once in the LSV. Gamma-rays predominantly originate from materials outside the active TPC volume (hence, we often refer to them as *external gammas*). That the veto cuts remove mostly external gammas can be seen from Fig. 5.2: below the ^{39}Ar endpoint, where the activity is dominated by ^{39}Ar (in AAr data, at least) which produces TPC signals only, the veto anti-coincidence cuts have little effect, as expected, while above the ^{39}Ar endpoint, the spectra are composed of gamma scatters, and the veto cuts have a significant impact. Not all gammas scatter in both the LSV and TPC, and so the veto cuts do not remove the high energy activity completely. Another effect worth point out is that after the veto cuts, the UAr spectrum exhibits a slight excess above the AAr spectrum at the ^{39}Ar endpoint. Though the effect is not strong, it indicates the presence of some activity in the UAr that was not present in the AAr. This excess will be discussed in more detail below.

Spectral fits are performed on the UAr S1 data to determine the constituent backgrounds that make up the spectra shown in Fig. 5.1. The fits are performed simultaneously on the

null field and 200 V/cm drift field data, as well as on the t_{drift} spectrum of 200 V/cm data. The spectra of the constituent backgrounds (betas and gammas) are generated via Monte Carlo simulation. Before describing the spectral fits, we give an overview of the DarkSide-50 simulation.

5.2 The DarkSide-50 Monte Carlo

The Monte Carlo (MC) code of DarkSide-50 is called *g4ds* and is a comprehensive simulation package. It includes models for the components of the TPC, LSV, and WCD, including geometry and materials. Spectra are generated for each of the most radioactive isotopes (Ar, U, Th, Co, Rn, Mn, K) in the components containing the largest quantities of those isotopes (PMTs, cryostat, fused silica, LAr). The raw energy spectra are determined by the decay chains of each of the isotopes. The γ -rays are propagated throughout the full geometry, recording the energy depositions in the active volume of the TPC. Of course, the ^{39}Ar β spectrum is also included in the energy depositions.

g4ds includes a model for the microphysics of each energy deposition. The model dictates the distribution of energy into scintillation and ionization. A fraction of the ionization electrons are recombined, called the recombination fraction. In *g4ds*, the recombination is assumed to be energy dependent, and the initial ionization vs. excitation fraction is assumed to be energy independent. That is, the recombination fraction absorbs all the energy dependence between recoil energy and observed S1. This can be contrasted with other models found in the literature (see, for example, work from Lenardo et al. [102]) in which the recombination is assumed to be independent of energy and initial ionization to scintillation ratio is energy dependent.

The scintillation photons are then propagated using an optical model for the TPC and LSV, accounting for the attenuation lengths, absorbance, index of refraction, reflectance, transmittance, and specularity of the various detector materials. Finally, the photons are converted to the observable photoelectrons, accounting for the QE of each PMT. The remaining ionization channel is largely ignored for the results presented in this work. The conversion

from ionization electrons to observed S2 is actively being tuned, as of this writing.

g4ds was tuned on the high statistics of ^{39}Ar decays in AAr. The largest uncertainties in the simulation are associated with the model of the microphysics described above and with the optical properties of the detector materials. The simulation parameters are tuned so that such observable parameters as the top-bottom asymmetry and channel occupancy of the PMTs matches between data and MC. The tuning was verified using data from a variety of calibration sources.

Given the simulated observable number of photoelectrons in S1 and S2, one can then pass the resulting photoelectrons through the electronics MC to generate simulated raw waveforms (described in Sec. 3.7), on which we could run the standard DarkArt reconstruction. The electronics MC and reconstruction quickly become very resource-intensive when simulating large numbers of events, so to ease computing needs, a clustering algorithm was introduced on MC data that is tuned to match the “pulse” detection of DarkArt reconstruction.

5.3 Spectral fitting

We now give an overview of the spectral fitting procedure. We wish to know the component backgrounds that make up the UAr spectra shown in Fig. 5.1. Internal betas are simulated uniformly throughout the active volume. At the outset, internal betas are expected to come from ^{39}Ar only, but as we shall see, this is not the case. External gamma backgrounds are simulated from a variety of detector components, as described above. The individual beta and gamma backgrounds are each simulated independently in both null field and field-on conditions. We build simulated spectra of S1 at 200 V/cm drift field, S1 at null field, and t_{drift} at 200 V/cm. The field-on S1 spectrum is fairly featureless, necessitating simultaneous consideration of the other two spectra. Once the spectral shapes are determined, they are added together in a weighted sum with the weights left as free parameters (separate sums for the three different types of spectra). We then perform a multidimensional fit by varying the weights simultaneously in all three spectra.

Fits using only the ^{39}Ar spectrum as the model for the internal betas failed to converge.

In particular, the fitter had trouble reproducing the data near the endpoint of the ^{39}Ar spectrum (cf. the discussion in Sec. 5.1). A variety of possible solutions were considered, including adding other activities such as ^{42}Ar ($\tau_{1/2} = 32.9$ yr, $Q_{\beta^-} = 600$ keV) and ^{85}Kr ($\tau_{1/2} = 10.8$ yr, $Q_{\beta^-} = 687.1$ keV) and distorting the energy dependence of the light yield curve. The addition of ^{85}Kr gave the best agreement between data and MC at the ^{39}Ar endpoint. The fitted spectra for S1 at null field and 200 V/cm drift field are shown in Fig. 5.3.

Using the spectral fits, the total ^{39}Ar activity in the DarkSide-50 UAr is found to be (0.73 ± 0.11) mBq/kg, and the total ^{85}Kr activity is (2.05 ± 0.13) mBq/kg. Due to the high statistics in the data, the uncertainties are systematics dominated. The systematic uncertainties are evaluated by varying the fit procedure. For example, individual components of the fit were removed from the fit, such as ^{235}U from the fused silica and ^{40}K from the cryostat. Also, the fit region of the t_{drift} spectrum was varied. Though the ^{85}Kr component was first hinted at in the spectral fits, its existence was confirmed by its decay to the metastable ^{85m}Rb state, discussed in detail next.

5.4 Delayed coincidence measurement

Evidence for the presence of ^{85}Kr in UAr first arose from spectral fitting of the scintillation signals but was confirmed by a delayed coincidence analysis. The ^{85}Kr isotope predominantly β -decays to stable ^{85}Rb , but it also decays to the metastable state ^{85m}Rb with a branching ratio of 0.43%. Metastable ^{85m}Rb has a half-life of 1.015 μs and de-excites by emitting a 514 keV γ -ray. The decay scheme of ^{85}Kr is shown in Fig. 5.4. The decay to the metastable state provides a delayed coincidence signature: a pair of S1s separated by a mean time of 1.46 μs . The first S1 comes from the β and the second comes from the γ . A search was performed on a 70 d data set, the same data used for the first WIMP search with UAr, as will be described in Ch. 6. In this section, we develop the selection criteria to tag delayed coincidence events, estimate the efficiency of the search (via MC), and compare the results against the expected rate coming from the spectral fits.

(a)

(b)

Figure 5.3: (a) Fit to the null field S1 spectrum of UAr data. The overall MC fit (red) to the UAr data (blue) is shown. The AAr data (black) are shown for comparison. The remaining colors show the component backgrounds. (b) Fit to the 200 V/cm drift field S1 spectrum of UAr data. Only the ^{39}Ar and ^{85}Kr components are shown for the fit.

Figure 5.4: Decay scheme of ^{85}Kr .

The general search strategy is to find a method to tag $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays with very little or no background. The main restriction is that we want to estimate the efficiency of the selection criteria using MC. Therefore, the search should be robust against the weaknesses of the MC. We are not so worried if a cut takes a steep hit in acceptance, as long as its efficiency to tag delayed coincidences can be reliably estimated via MC.

The delayed coincidence search is developed by looking simultaneously at simulated $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays and at real data. The simulation is used to tune the cuts and to estimate the efficiency of the cuts. We simulate 2×10^4 decays uniformly throughout the TPC. Each simulated event consists of a β that promptly produces scintillation light and a γ , produced with an exponentially distributed time delay with respect to the β and given an isotropically distributed direction. The energy of the β is drawn from a continuous distribution with end point 173 keV , and the energy of the γ is fixed at 514 keV . The γ is allowed to Compton scatter, once or multiple times; be absorbed by the photoelectric effect; or escape the TPC completely. The scintillation light from the β and γ scatters are propagated through the standard optical simulation of $g4ds$, which produces a hit pattern of

Figure 5.5: Simulated waveforms for $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$, showing the delayed coincidence of the S1s.

photoelectrons over the channels for each event. Each photoelectron is defined in the simulation by a channel and a time relative to the trigger, usually $t = 0$. We simulate only the primary scintillation signals because the delayed coincidence search is based on only these signals. We then use the electronics MC to produce simulated waveforms for each event. Fig. 5.5 shows an example waveform for a simulated $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decay. We pass the simulated waveforms through the standard DarkArt reconstruction chain.

5.4.1 Event selection

We now develop the search algorithm for the delayed coincidence signature of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays. There are several features of the ^{85}Kr decays that help us in the search. The β has up to 173 keV energy, so the DAQ will efficiently trigger on it, and the DarkArt pulse finder will efficiently find it. As already mentioned, the γ is well-separated from the β (in time) of order a microsecond. This is the main feature of the search as there are few other classes of events exhibiting the same waveform topology. Finally, we exploit the fact that the delayed S1 is almost always larger than the S1 on which we trigger (i.e. the γ deposits

more energy than the β).

There are also several challenges to the search. The γ is likely to multiply scatter within the active volume, inducing many S2 pulses, so the search must be robust against the appearance of one or many S2 pulses, and the γ can escape the TPC entirely, without depositing any energy. We cannot distinguish a lone β of a ^{85}Kr decay to ^{85m}Rb from a β of ^{39}Ar decay or ^{85}Kr decay to ^{85}Rb . Escaping gammas constitute a large portion of the search inefficiency. Finally, the DarkArt pulse-finding algorithm has trouble separating two S1s closer together than a few μs , instead lumping them together as a single pulse. We cannot rely on DarkArt's native pulse-finding results to separately identify the S1 of the β and the S1 of the γ . Rather than retune DarkArt, we make use of existing reconstruction variables. Fortunately, DarkArt is efficient at finding the β , so it becomes a matter of creative use of the reconstruction variables to identify the β/γ pairs.

We start with one of the most commonly used parameters in DarkSide-50: f_{90} . The S1 of the delayed γ pulls down the value of f_{90} to below the main ER f_{90} band, as verified in MC data and shown in Fig. 5.6a. There are three populations of events. The population with $\text{S1} < 800$ PE and $f_{90} \approx 0.3$ are escaping gammas. Recall that S1 is the integral of the first 7 μs after the pulse start, regardless of where the pulse finder decided the end of the pulse is. The band with $\text{S1} > 800$ PE and $f_{90} \approx 0.3$ are events with S1 from both the β and γ but the ^{85m}Rb decay was so fast that the two S1s are completely overlapping. And the main population with $f_{90} < 0.2$ have well-separated S1s and constitutes the bulk of the simulated delayed coincidences. Turning to real data, we begin with only the most basic quality cuts: we require 38 valid channels in each event, the baseline to be found in all channels, and the trigger time not be too close to the previous trigger to avoid triggers on the residual of the previous event. For the remaining events, we require the number of reconstructed pulses to be ≥ 2 , which includes the vast majority of events. The f_{90} vs. S1 distribution is shown in Fig. 5.6b, where there are many events with reduced f_{90} . To avoid the very large population of the main ER band, we require $f_{90} < 0.2$ and $\text{S1} > 400$ PE and take the hit in efficiency due to the rejection of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays with an escaping γ and ultra short de-excitation times. The S1 cut is to avoid the large population of low energy events with

(a)

(b)

Figure 5.6: f_{90} vs. $S1$ distribution for (a) simulated $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays and (b) data including all events with ≥ 2 pulses.

f_{90} that spans the full range; these events are typically triggers on S3 or on the tail of S2.

We continue by exploiting the time separation between the β and γ and the expected amplitude difference between their S1s. We look at the peak time of the waveform of the sum channel in the first 5 μs after the pulse start¹. We call this parameter T_{peak} . The use of the fixed time window makes the search robust against DarkArt finding the β and γ separately or together. The T_{peak} parameter is a good estimator for the delay between the two S1s: because the γ is typically larger than the β and Ar scintillation light has a very fast rise time and a short prompt decay time, the value of T_{peak} closely approximates the arrival time of the γ . The peak time of ordinary single scatter S1s (e.g. ^{39}Ar and ^{85}Kr decays) is $<0.1 \mu\text{s}$ from the pulse start, while in ^{85}Kr decays to ^{85m}Rb , T_{peak} is much greater, up to several microseconds. For the delayed coincidence search, we require $T_{\text{peak}} > 0.05 \mu\text{s}$. We also require $T_{\text{peak}} < 4 \mu\text{s}$; this portion of the cut is motivated later.

The T_{peak} cut introduces efficiency losses in several ways. The first efficiency loss is due to escaping gammas, in which case we observe only the β ; in this case, the peak time is governed by the first S1, and the event does not pass the peak time cut. There is little we can do to regain these losses since they are indistinguishable from the huge population of ^{39}Ar and ^{85}Kr decays and single scatter gammas. The second loss is due to the γ arriving within 50 ns of the β , which occurs in 3% of decays. The efficiency losses due to the first two effects subsume the inefficiencies of the f_{90} cut. The third efficiency loss of the peak time cut arises from the S1 of the γ being smaller than the β . This occurs when the γ Compton scatters a small number of times, depositing only a small amount of energy, and then escapes the TPC. The peak time is then governed by the β and will be $<50 \text{ ns}$. This effect contributes 3% to the efficiency loss. The final efficiency loss occurs from rejecting events with $T_{\text{peak}} > 4 \mu\text{s}$ and accounts for 6.5%.

There is a class of non- ^{85}Kr events that satisfy the peak time criterion: events with S1 immediately followed by S2. The drift time of these events is typically $<5 \mu\text{s}$. An example is shown in Fig. 5.7. These are events that occur in the uppermost portion of the

¹Such a variable is not explicitly evaluated in DarkArt, but it can be effectively constructed from existing variables by considering all pulses that start within 5 μs after the trigger time.

Figure 5.7: Example waveform of an event from the uppermost portion of the LAr, giving an S1 and S2 very near each other. This event is a background to the delayed coincidence search of ^{85}Kr decays to ^{85m}Rb .

LAr. There are several parameters that we can use to separate the delayed coincidences of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays from the ultra short drift time events: (a) the pulse shapes of the first $5 \mu\text{s}$ are starkly different, and (b) the distributions of light over the PMTs are different. We begin with the pulse shape. We construct a new parameter, f_{5000} , which is the ratio of the integral of the first $5 \mu\text{s}$ to the first $7 \mu\text{s}$ after the pulse start. With a delay of $1.46 \mu\text{s}$ between the two S1s, we expect f_{5000} to be almost unity for $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays: the S1 of the γ will have mostly decayed away after a few microseconds. If the β and γ deposit all their energy, the expected value of f_{5000} is 0.95 (with average β/γ separation). As the delay between the β and the γ increases, f_{5000} decreases, as shown in Fig. 5.8a. For very short drift time events, we expect f_{5000} to be significantly smaller: the electroluminescence signal is significantly broader and has a longer tail than S1, extending beyond $5 \mu\text{s}$. The FWHM of S2 is typically $\sim 3 \mu\text{s}$. Even for an event with essentially zero drift time, the S2 pulse extends significantly beyond $5 \mu\text{s}$ and the maximum expected value of f_{5000} for very short drift time events is smaller than that of delayed coincidence events, as a function of T_{peak} . In

this way, we can separate delayed coincidence S1s from the short drift time events. Indeed, Fig. 5.8b shows a population of events in data consistent with being delayed coincidences, as predicted by MC. The lower band is associated with short drift time events. There is also a large population of events with extremely small peak time; these are single scatter events. Because f_{5000} and T_{peak} are correlated, we select $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays using a cut in the plane of the two parameters. We empirically place a cut between the two bands, defined by the function

$$f_{5000} > 0.955 - 0.015T_{\text{peak}} + 0.00058T_{\text{peak}}^2 - 0.0032T_{\text{peak}}^3. \quad (5.1)$$

As peak time increases, we see that the delayed coincidence band begins to merge with the very short drift time events. We need additional cuts in additional parameters to further separate the signal from background. We exploit the top heavy distribution of S2 light over the PMTs compared to the relatively uniform distribution of S1 light. We define an asymmetry parameter

$$a_{\text{S1}} = \frac{S1_{\text{top}} - S1_{\text{bot}}}{S1_{\text{top}} + S1_{\text{bot}}} \quad (5.2)$$

where $S1_{\text{top}}$ ($S1_{\text{bot}}$) is the sum of $S1_{\text{fixed},i}$ over the top (bottom) channels. The distribution of a_{S1} for events passing the f_{90} and peak time criteria are shown in Fig. 5.9a. Also shown are the distributions of a_{S1} for typical S2 pulses, taken from well-reconstructed single scatter events, and for simulated $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays. The distributions are scaled by hand to match each other. The bulk distribution to the right is consistent with being S2. The distribution of a_{S1} for simulated decays relies heavily on the model of the optics of the TPC. It has long been a challenge of DarkSide-50 to achieve an accurate optical simulation, so the disagreement between data and MC is not a severe cause for concern. To select $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays, we place a cut on the asymmetry parameter at $a_{\text{S1}} < 0.1$.

We can also use the distribution of the fraction of S1 light in the dominant channel (the so-called S1 maximum fraction parameter) to better isolate our target signal. Because all the light is produced near the very top of the TPC for very short drift time events, all the light should be concentrated in one or a few channels. In contrast, the primary scintillation light of the β and γ of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays will be more uniformly distributed

(a)

(b)

Figure 5.8: Distribution of f_{5000} vs. peak time for (a) simulated delayed coincidence events and (b) data, including all events that pass the f_{90} cut.

(a)

(b)

Figure 5.9: Distributions of (a) top-bottom asymmetry parameter a_{s1} and (b) max. fraction for candidate delayed coincidence events passing the f_{90} and peak time cuts (black), typical S2 signals (red), and simulated delayed coincidence events (blue). The curves are scaled by hand to match each other.

over the PMTs. As shown in Fig. 5.9b, the bulk of the candidate delayed coincidence events passing the f_{90} and peak time cuts are consistent with being S2 signals. Similar to the a_{S1} distribution, it is difficult to draw any conclusions from the disagreement between the S1 max. frac. distributions in data and MC. Events near both the top and bottom of the TPC have extreme S1 max. frac. values. To avoid rejecting good delayed coincidence events from the bottom of the event, we reject events as short drift time events if they have both $a_{S1} > 0.1$ and S1 max. frac > 0.1 .

One must be cautious in making more aggressive cuts in the asymmetry and max frac parameters. These parameters rely heavily on the optical model of the TPC in the MC, which is not at present very finely tuned. If we make more aggressive cuts, the estimate of the efficiency of the selection criteria becomes unreliable. The cuts are at present tuned such that the efficiency losses are estimated to be negligible, while removing enough of the short drift time events that the f_{5000} cut is efficient for β/γ separation of up to $4 \mu s$.

In summary, the cuts used to select $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays are

- 38 channels
- baseline found
- event Δt
- $f_{90} < 0.2$
- $S1 > 400 \text{ PE}$
- peak time $T_{\text{peak}} > 0.05 \mu s$ and $< 5 \mu s$
- $f_{5000} > 0.955 - 0.015T_{\text{peak}} + 0.00058T_{\text{peak}}^2 - 0.0032T_{\text{peak}}^3$
- asymmetry parameter $a_{S1} < 0.1$
- S1 max. frac < 0.1

5.4.2 Results

The delayed coincidence search was conducted over 70.9 live-days of data with UAr and yielded 1717 candidate $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays. An example candidate event is shown in Fig. 5.10.

The S1 spectrum of these candidate events is shown in Fig. 5.11. Recall the S1 is an integral over a $7\ \mu\text{s}$ window, and therefore includes the majority of the S1s of the β and the γ . The full absorption peak is visible, and a Gaussian fit to the peak yields a mean of 3560 PE. The broad spectrum below the full absorption peak is composed of the β and the Compton continuum of the γ . The drop-off at low energy is due to efficiency losses of the peak time cut, and the hard cut-off is due to the $\text{S1} > 400\ \text{PE}$ cut.

The distribution of the β/γ separation is shown in Fig. 5.12. A simple maximum likelihood fit of an exponential decay yields a mean ^{85m}Rb lifetime of $(1.51 \pm 0.06)\ \mu\text{s}$, in agreement with the literature value of $1.46\ \mu\text{s}$, within errors. The quoted uncertainty is directly from the fitter.

5.4.2.1 Efficiency

In order to translate the observed 1717 candidate events into an estimate of the activity of ^{85}Kr decays to the metastable ^{85m}Rb state, we must estimate the global efficiency for the search algorithm to identify the decays. As has been described throughout this section, there are several sources of inefficiency:

- Escaping gammas, leaving a single S1 of the β . This efficiency loss is estimated to be 15 %.
- Time separation between the β and γ is extremely short ($< 50\ \text{ns}$). These events do not pass the peak time cut. The efficiency loss due to the peak time cut subsumes the equivalent loss due to the f_{90} cut. The efficiency loss is estimated to be 3 %.
- The S1 of the γ is smaller than the S1 of the β . These events do not pass the peak time cut. The efficiency loss is 3 %.

(a)

(b)

Figure 5.10: Candidate $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decay. (a) Full sum waveform. The S1s are near $t = 0$, and the rest of the signals are the S2s from the β and the multiple scatters of the γ . (b) Zoom of $-10 \mu\text{s}$ to $10 \mu\text{s}$. The peak at $-6 \mu\text{s}$ is the S1 of the β and the peak at $-5 \mu\text{s}$ is the S1 of γ .

Figure 5.11: Spectrum of $S1_{\beta} + S1_{\gamma}$ for all candidate events.

Figure 5.12: Distribution of time separation between β and γ of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays. The exponential fit (red) gives a decay constant of $(1.51 \pm 0.06) \mu\text{s}$.

- The γ is $>4 \mu\text{s}$ from the β . These events do not pass the peak time cut. The efficiency loss is 6.5 %.
- $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays that do not pass the f_{5000} , asymmetry, or max. frac cuts. These efficiency losses are estimated to be negligible.

The efficiency losses are all estimated from MC. The loss due to escaping gammas is taken directly from $g4ds$, by counting all events in which the γ had zero scatters. The losses due to the other cuts are estimated by passing the simulated raw waveforms through the entire reconstruction chain and applying the identical cuts used for the delayed coincidence search in data.

The global detection efficiency is estimated by taking the ratio of the number of MC events surviving all the event selection criteria to the number of simulated MC events. Of the 20k simulated events, 15055 of them survive, corresponding to a global detection efficiency of 75.9 %, in agreement with the combined efficiencies described above.

5.4.2.2 Rate

In 70.9 d of UAr data, we observe 1717 delayed coincidence events. Considering the global detection efficiency of 75.9 %, we derive the observed rate of 31.9 ± 0.8 ^{85}Kr decays to the metastable ^{85m}Rb state per day, where the uncertainty is the Poisson statistical uncertainty only². From the spectral fit described in Sec. 5.3, we expect 35.3 ± 2.2 decays per day via the 0.43 % branching ratio, where the uncertainty is systematics dominated. The two results agree to within uncertainties. The systematic uncertainties on the observed rate of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays have not been evaluated.

5.5 Summary

We have measured the ^{39}Ar activity in the UAr of DarkSide-50 to be (0.73 ± 0.11) mBq/kg, the highest precision measurement to date. This corresponds to a depletion factor of 1400

²Note that this is an updated value with respect to the one given in Ref. [103].

of ^{39}Ar activity in UAr relative to AAr. In the process, we discovered the presence of ^{85}Kr at a level of (2.05 ± 0.13) mBq/kg in the UAr, which was first hinted at during the fitting of the UAr S1 spectra and was confirmed by the observation of ^{85}Kr decays to the metastable ^{85m}Rb state. One could improve the delayed coincidence search (especially the efficiency of the search algorithm) using a likelihood ratio approach. The present search was designed to be relatively simple and done in a timely manner for publication.

^{85}Kr is also expected to be in atmospheric argon. However, such a measurement is complicated by the high 1 Bq/kg ^{39}Ar activity. The main β -decay channel for ^{85}Kr is buried beneath the overwhelming ^{39}Ar activity, making it difficult to use spectral fitting of AAr data to measure the ^{85}Kr activity. And the accidental coincidence rate of ^{39}Ar decays within a $4\ \mu\text{s}$ window is high enough to obfuscate the delayed coincidence search of $^{85}\text{Kr} \rightarrow ^{85}\text{Rb} + \beta + \gamma$ decays.

The large reduction of ^{39}Ar activity in UAr is a crucial ingredient that enables the LAr TPC technology to reach the neutrino floor, as will be discussed in Ch. 7. Radioactive ^{85}Kr can be produced underground by (α, n) reactions at levels similar to ^{39}Ar [81], and for future detectors, it will be necessary to remove the ^{85}Kr from UAr. For the next generation of DarkSide, DarkSide-20k, the ^{39}Ar will be further depleted in UAr by cryogenic distillation. The ^{85}Kr contamination will be automatically removed as well. There are no plans to remove ^{85}Kr from the UAr in DarkSide-50. In any case, the contamination is not an issue due to the spectacular rejection power of f_{90} , as discussed in Ch. 6.

CHAPTER 6

WIMP search with UAr

The second WIMP search in DarkSide-50 and the first search using underground argon was performed in 2015. The WIMP search data were acquired between April and July of 2015 and the results were published in April 2016 [103]. This chapter describes in detail the analysis, which led to the strongest WIMP exclusion limit with a LAr target, to date. We describe the data selection criteria and a few of the stability checks performed for the TPC. We describe in detail the cuts used for selecting single scatter events and their efficiencies, devoting considerable attention to the f_{90} discrimination parameter. Finally, we describe the procedure for translating the WIMP search results into a dark matter limit.

6.1 Run selection and stability

The nominal configuration of DarkSide-50 in WIMP search mode was as follows: the TPC was filled with UAr; the drift electric field was 200 V/cm, and the extraction field was 2.8 kV/cm; the gas pocket was maintained at nominal 1 cm thickness; the gas pressure above the liquid was 15.6 psi, which set the LAr temperature at 89 K; the gas recirculation loop was continuously flowing at 28.5 slpm; the LSV was filled with a cocktail of 95 % by mass pseudocumene (liquid scintillator), 5 % TMB, and 1.4 g/L PPO (wavelength shifter). The DAQ system was configured as follows: the TPC was triggered with a *majority-2* or *majority-3* mode, requiring the coincidence of any two (or three) discriminated channels to fire within 100 ns¹; upon receipt of a trigger, the digitizers opened a 440 μ s acquisition gate for each channel, and further triggers were inhibited for 810 μ s to prevent re-triggering on

¹The WIMP search data taking began with majority-3 and switched to majority-2 early in the campaign.

the residuals of a primary trigger²; the TPC trigger was also sent to the outer detector DAQ system, where a 200 μs gate was acquired for each channel of both the LSV and WCD. Laser calibration runs for the TPC were taken approximately every 12 hrs. Data taking was interrupted briefly for a $^{83\text{m}}\text{Kr}$ calibration campaign shortly after the start of the WIMP search campaign but was otherwise essentially continuous. The run ID of the first run is 11856 and the last run is 13181.

We examined the stability of a variety of detector components over the course of the WIMP search campaign. A few of the checks are highlighted here.

6.1.1 PMT gain

We check the stability of the TPC PMTs primarily by looking at the trend of the SPE mean. We observe a decrease in the SPE mean, and therefore a decrease in PMT gain, over time, as shown in Fig. 6.1. The PMT “fatigue” is a known effect [104]. Fortunately, the (at least) daily calibration of the PMT gain allows us to track the change, and the fatigue does not affect the light yield of the TPC in any significant way.

6.1.2 Light yield

The light yield in the TPC can be checked using the high energy gamma lines, as discussed in Chapter 5, or by $^{83\text{m}}\text{Kr}$ injected into the circulation loop. These methods were used primarily to verify that the LY is unchanged between the AAr and UAr eras of DarkSide-50. However, they do not provide ideal checks on the LY trend over time. The $^{83\text{m}}\text{Kr}$ calibration was performed at only a single time point, and the high energy gamma lines are statistics limited. Fortunately, we can make use of a low energy peak from ^{37}Ar .

6.1.2.1 ^{37}Ar

Low energy ^{37}Ar decays were discovered unexpectedly at the start of the UAr WIMP search campaign. ^{37}Ar is a cosmogenically activated isotope of argon, activated while the UAr

²All pre-scales of the so-called G2 trigger were set to 1.

(a)

(b)

Figure 6.1: Trend of SPE mean over time for each TPC PMT in the (a) bottom array and (b) top array. Plot from X. Xiang.

inventory was at the Earth’s surface at Fermilab for upwards of a year and while it was shipped across the Atlantic Ocean. The isotope decays with a half life of 35 d via electron capture, emitting a 2.4 keV X-ray. Therefore, ^{37}Ar activity manifests as a peak at about 17 PE, near the low energy threshold of the TPC, as shown in Fig. 6.2. A rough estimate of the cosmogenic activation of the UAr sample was performed [105], which accounted for the time spent at the Earth’s surface at Fermilab and in transit, and gave a ^{37}Ar activity of ~ 120 decays/kg/day. The activity of the peak, determined crudely by dividing the number of events below 30 PE by the live time from a few runs early in the WIMP search campaign, was found to be ~ 30 counts/kg/day. The observed and predicted activities agree to within an order of magnitude. That the peak is so near the low energy threshold of the TPC may account for the lower observed activity. We must be wary that the peak is not arising artificially due to an efficiency threshold, either in the TPC trigger or in the pulse finder of the reconstruction. For this reason, the trigger was changed from majority-3 to majority-2 about 2 weeks after the start of the WIMP search campaign, to try to lower the energy threshold of the TPC. Data taken in the majority-2 mode strictly include events that would have triggered under majority-3, though the improvement in energy threshold is not straightforward to quantify.

To verify that the peak is not due to threshold effects, we compare integral spectra of the pulses with spectra generated by integrating a $7\ \mu\text{s}$ ROI near the trigger time. When the pulse finder identifies the S1 of an ^{37}Ar decay as a pulse, the pulse area and ROI closely agree. However, because the ROI calculation does not rely on the pulse finder, the ROI spectrum can reveal events with S1 signals below the pulse finder threshold. We look for events in which only one pulse is found and require that its start time be after the expected trigger time. That is, we look for events in which the reconstruction finds an S2 but no S1, which should occur when the DAQ triggers on an S1 that is too small to be found by the pulse finder. The S2 is at least 10 times larger than the S1 (at the TPC edge; the factor is even larger at the TPC center, cf. Sec. 3.5.2), so the pulse finder will always find the S2. (There is no evidence of events in which the pulse finder missed both the S1 and the S2 of an ^{37}Ar decay.) The ROI spectrum is shown in Fig. 6.3. While the peak in majority-2 triggered data

(a)

(b)

Figure 6.2: Low energy peak in UAr data due to ^{37}Ar activity. (a) Full energy spectrum. The lowest bin population is easily mistaken for zero-integral events. (b) Zoom to low energy region shows there is, in fact, a peak.

Figure 6.3: ^{37}Ar spectrum of integrated $7\ \mu\text{s}$ ROI in the expected S1 window. We show the live-time-normalized distributions for majority-3 triggered data (blue) and majority-2 triggered data (red) separately. We include both 2-pulse events and 1-pulse events where the found pulse is S2. See text for details.

is reduced due to the finite half-life of ^{37}Ar combined with the fact that majority-2 triggered data were acquired after majority-3 triggered data, a shoulder appears to the left of the peak, suggesting that the peak is real and not due to threshold effects either from the trigger or the pulse finder. Note that it is more challenging to give a quantitative measurement of the trigger threshold.

Perhaps the most convincing argument that the low energy peak is due to ^{37}Ar is that the peak decays over time in rough agreement with the expected mean lifetime of 50.5 d, as shown in Fig. 6.4. The points in the figure correspond to one run each and are produced by dividing the number of events with 2 pulses (S1 and S2) with S1 below 30 PE by the live-time for each run. The error bars consider Poisson statistics only. The data are fitted separately for majority-3 triggered data and majority-2 triggered data since they have different efficiencies for observing ^{37}Ar decays. The fit form is $A + Be^{-\lambda t}$ and the fitted means are (33 ± 5) d and (34 ± 6) d for majority 2 and 3 triggered data, respectively. The fitted decay constants are in

Figure 6.4: Decay of ^{37}Ar peak over time. Because the majority-3 triggered data (blue, fit in black) and majority-2 triggered data (red, fit in green) may have different efficiencies for ^{37}Ar events, we evaluate the decay constant separately for the two types of triggered data. The fit is performed with an exponential plus a constant. The gap is the ^{83m}Kr calibration campaign.

the right ballpark but are lower than the expected value, which may be due to competing low energy effects that are also changing over time. No systematic uncertainties were evaluated.

Finally, ^{37}Ar activity provides a convenient low energy peak to track the behavior of the TPC over the first few months of running with UAr. Due to possible threshold effects, one must be cautious about using the ^{37}Ar peak as an absolute measure of LY, but its relative change over time is valuable. The light yield is found to be constant, within statistical errors, as shown in Fig. 6.5. By the end of the UAr WIMP search data taking, the ^{37}Ar is nearly completely decayed away, and the statistics are too low to reliably fit the low energy peak.

Figure 6.5: Light yield over time as measured from the ^{37}Ar peak in UAr data. Groups of consecutive runs were concatenated to ensure enough statistics in the low energy peak to fit with a Gaussian. Plot from X. Xiang.

6.1.2.2 ^{60}Co

The stability of the LSV performance was also studied. The light yield was measured by fitting the ^{60}Co peak, which is present in the stainless steel cryostat and emits two coincident γ -rays of energies 1.17 MeV and 1.33 MeV, which provide a convenient coincidence tag between the TPC and LSV. Using the ^{60}Co peak to establish the precise light yield of the LSV requires a sophisticated spectral fit, but a simple Gaussian plus exponential form is sufficient for tracking the relative trend over time. Figure 6.6 shows that the LSV light yield is stable over the UAr WIMP search campaign.

A variety of other parameters were also checked, including the consistency of the TPC and LSV electronics noise, the stability of the drift and extraction fields, and the argon purity. None of these parameters indicated anything abnormal about the performance of DarkSide-50. After these data selection criteria, 71.5 live-days of data remained in the WIMP search data set.

6.2 Cuts

We now discuss the event selection criteria, or *cuts*, on the remaining data. The main goal of the cuts is to find candidate WIMP events, which, qualitatively, should satisfy a few conditions. A WIMP event should:

- be a single scatter in the TPC,
- be in anti-coincidence with the LSV and WCD,
- be in the correct energy range, and
- have f_{90} consistent with a NR.

The philosophy of the cuts, then, is to keep as many WIMP-like events as possible, while removing non-WIMP events. Notice that β -decays from, say, ^{39}Ar satisfy all the criteria less the f_{90} one. When developing each cut, it is often useful to keep in mind the concrete

Figure 6.6: (a) ^{60}Co peak visible in the LSV prompt spectrum. Blue is data and red is a Gaussian plus exponential fit. (b) Trend of the Gaussian mean over the UAr WIMP search campaign. The gap corresponds to ^{83m}Kr operations in the TPC. Plots from X. Xiang.

notion of keeping ^{39}Ar events, rather than trying to think about keeping WIMP events that, *a priori*, may or may not be present. Indeed, since the sensitivity of DarkSide-50 is far less than that of existing limits from LUX and XENON, we do not expect to find any WIMP-like events. It can be dangerous to develop cuts with the expectation of finding zero signal events, as one faces the (often subconscious) bias to tune the thresholds until all potential signal events are removed, especially since this is not a blind search. Therefore it is useful to develop cuts (apart from the f_{90} cut) with the mindset of keeping ^{39}Ar decays in order to reduce bias.

Once we have performed the WIMP search, we go through the steps to translate the result into an exclusion curve in the WIMP mass vs. WIMP-nucleon cross section plane. This involves several ingredients, including evaluating the efficiency of each cut on either the exposure or signal acceptance, modeling a canonical WIMP signal using standard astrophysical parameters, and combining into a WIMP exclusion curve.

We evaluate the efficiency of the cuts on a case-by-case basis. The philosophy for evaluating efficiency is to answer the question: how good is this cut at keeping WIMP events? Since we do not have a sample of WIMP events to use (both because we do not expect to have any and because these are the events we are searching for!), we use various proxies. In some cases, the proxy can be neutron events from AmBe or AmC data, other times it is the high statistics ^{39}Ar events from AAr data, while others it is the UAr sample itself. The choice of proxy depends on the parameters used in the cut.

6.2.1 TPC cuts

6.2.1.1 NChannels

Definition Number channels present in the event is 38.

Purpose Occasionally, one or more of the 8 front-end boards in the DAQ hardware malfunction, and fewer than the expected 38 channels are recorded. This cut removes such spurious events from the analysis. The rate of these events is extremely small; >99.9% of

events survive this cut. Events removed by this cut usually occur near end of a run, just before a DAQ crash. When the DAQ crashes, the integrity of the majority of the raw data from the run remain in tact.

6.2.1.2 Baseline found

Definition The baseline is found on the sum channel.

This is a somewhat misleading requirement since we do not actually perform any baseline finding on the sum channel (see Sec. 3.1.2.2). Instead, for the baseline-found variable of the sum channel to indicate success, the baselines must be successfully found on all the individual channels. Therefore, this cut effectively requires the baseline to be valid for all channels.

Purpose All reconstruction steps downstream of the baseline finder make sense only on the baseline-subtracted waveforms. Therefore, we require that the baselines be valid for all channels. Failure modes of the baseline finder include: large fluctuations in electronics noise; trigger on the tail of a large signal (typically a muon or a large S2); bipolar noise. 86.4 % of events surviving the previous cut pass this cut. The vast majority of rejected events do not constitute any loss of acceptance for nuclear recoils.

6.2.1.3 Event Δt

Definition Time to previous recorded trigger is >1.35 ms. The time to the previous trigger is reconstructed by the sum of live time and inhibit time.

Purpose We occasionally re-trigger on the residual signal of a previous trigger, typically an echo pulse (see Sec. 3.5.1). The $810 \mu\text{s}$ inhibit window is supposed to suppress such re-triggers, but events that occur within the inhibit can lead to triggers on echo pulses or tails of S2. The 1.35 ms threshold is chosen as a conservative value that encompasses all the various inhibit and acquisition window configurations used during the AAr WIMP

search campaign³ and was kept at the same value for the UAr search in order to maintain consistency. Of events surviving all previous cuts, 94.7% pass this cut.

6.2.1.4 File I/O

Definition Live time is <1 s.

Purpose During the early days of DarkSide-50 (the AAr campaign), the DAQ was under heavy active development. One of the misbehaviors was that it took a long time to close one raw data file and open the next one. While no triggers were recorded during the file closing and opening procedures, the intervening time was (incorrectly) counted as livetime. To avoid bloating the livetime calculation, we removed those events with exceptionally long livetimes. Such events were almost exclusively associated with the first event of a newly open file. The DAQ was improved to reduce the file closing and opening procedure, so during the UAr campaign, this cut had no effect. The cut was kept for historical purposes.

6.2.1.5 Veto present

Definition There exists a valid veto event with the TPC event, where the GPS timestamps match to within 100 ns.

Purpose Since we use a single global trigger from the TPC to acquire data from all detectors, we expect a one-to-one correspondence between the TPC and LSV/WCD data streams. There are several ways in which the data streams might become misaligned in the final data set. The TPC and veto acquisition systems are independent, and occasionally one of the subsystems may malfunction. This can cause misalignment at the raw data level (the n th event of a TPC file may not have the same event ID as the n th event in the LSV file). In this case, we must use some other means than event sequence for associating TPC and

³During the AAr campaign, the event Δt cut did not completely fulfill its purpose due to the use of the G2 trigger. In the intervening time between two recorded triggers, many triggers may have been rejected by the prescale. The second recorded trigger could then be on the residual of a rejected trigger and have nothing to do with the previous recorded trigger.

veto events. We use the GPS timestamp (not really a GPS timestamp, but a 50 MHz global clock synced to a GPS pps signal). The data streams may also become misaligned when one of the reconstruction algorithms fails on an event, causing a missing event from the DarkArt output. Again, the GPS alignment is robust against such effects. Any time there is a TPC event with missing veto event (or vice versa), we throw out the event.

The GPS timestamp alignment parameter is governed by differences in cable lengths and insertion delays between the two acquisition subsystems. Of events surviving all previous cuts, 99.1 % pass this cut.

6.2.1.6 Single scatter cut

Definition Number of reconstructed pulses in the TPC is 2 or 3 if the 3rd pulse is S3.

Purpose This cut coincides with the pulse identification scheme described in Sec. 3.5.1. WIMP events are expected to be single scatters within the TPC. Single scatters typically have a single S1 and a single S2. We reject events where there is evidence for more than one scatter, usually due to multiply scattering gammas. The scatters will occur no more than a few ns apart, so the S1s of each scatter will overlap each other, but because the scatters will typically be at different z , their S2s will be separable. There is a small probability that multiple scatters will be unresolvable because they occur at the same z ; S2s separated by less than $\sim 2 \mu\text{s}$ are usually reconstructed as a single pulse.

However, not all reconstructed pulses are necessarily S1 or S2. The most common case is S3, which is due to photoionization of the cathode by the S2 light. The ionization electrons are drifted across the full length of the TPC, producing electroluminescence signals à la S2. The S3 light always appears with fixed delay relative to S2 at the maximum drift time of the TPC. We accept all events with S2+S2+S3. S1 light can also photoionize the cathode, and we sometimes see S1-echo at fixed delay relative to S1. But because S1 is smaller than S2 (at least for ER), events with discernible S1-echo are rarer. To keep the analysis simpler, we ignore events with S1-echos.

33.5 % of events surviving all previous cuts pass this cut. The majority of events removed by this cut are multiple scatter gammas and do not count as loss of acceptance for NRs, as discussed below.

Efficiency We estimate the efficiency of the single scatter cut to accept NRs. The two main sources of inefficiency for this cut are single scatters piled up with other signal and single scatters that are without S2. Such single scatter events are rejected by the cut, as defined above. Though the cut also removes multiple scatters and other anomalous classes of events, they are not counted as loss of acceptance because they have some feature that excludes them as single scatters.

We estimate the efficiency by evaluating the ratio

$$\text{efficiency} \simeq \frac{\# \text{ events passing the cut}}{\text{total } \# \text{ single scatters}} \quad (6.1)$$

where the total number of single scatters includes events passing the cut and single scatters removed by the cut. The ratio is evaluated as a function of S1. This estimate is only an approximation: the expected WIMP spectrum is not the same as the background spectrum, so the overall efficiency of the cut will be inaccurate. But, as will be shown, the estimate is conservative, and we will find that the efficiency curve is approximately flat vs. S1, so the overall acceptance will not be wildly inaccurate.

The task now is to estimate the number of single scatters rejected by the cut. We look directly at the UAr data. We cannot use AAr data because the accidental rate is different. We first estimate the number of single scatters that were rejected because they were piled up with some other signal. There are two types of signal that a single scatter can be piled up with: another primary scintillation signal can occur in accidental coincidence anywhere in the acquisition window, or an S1-echo may be present. An event has an S1-echo if there is a pulse found to start within the expected echo window ($372 \mu\text{s}$ to $400 \mu\text{s}$) relative to the S1. We need only consider 3-pulse and 4-pulse events when looking for S1-echo: 3-pulse events will have S1+S2+S1-echo and 4-pulse events will have S1+S2+S1-echo+S2-echo. Events with a greater number of pulses may also contain S1-echo (or S2-echo), but they will either

be handled by the accidental coincidence case or can be ruled out as non-single-scatters (if there are 5 or more pulses in an event, and none other than the first is an S1, then at least one of the remaining pulses is something other than S1, S2, S1-echo, or S2-echo). As a function of S1, nuclear recoils produce S1-echo at the same rate as electronic recoils. The S2-echo production rate will not be the same for NR as ER, as a function of S1, because NR S2 is much smaller than ER S2. But since we are already accepting events with S2-echo, using ER to estimate the acceptance loss due to rejection of S1-echo should be accurate to NR.

We next tag accidental coincidences with another scintillation signal. In AAr, the ^{39}Ar rate was high enough that the accidental rate was non-negligible. But in UAr, the activity is significantly reduced and the accidental rate is very low. We estimate the accidental rate by searching through all events with ≥ 3 pulses (and the 3rd pulse is not an S3) and looking at all pulses in the event beyond the first. We determine if the pulse looks like an S1 using a few criteria, such as f_{90} and FWHM. If any of the pulses (besides the first) looks like an S1, we assume the event has an accidental coincidence of (at least) two single scatters, and the event is counted in the denominator of Eqn. 6.1.

There is also a population of events in which only a single pulse was found, and the pulse looks like a proper S1—the pulse start time is at the expected time for S1, the f_{90} are consistent with S1, and the pulse area spectrum is consistent with ER background. But the S2 is nowhere to be found. By looking at f_{90} , it appears that these single-pulse scintillation events are both ERs and NRs. The NR-like events are likely α events coming from the surfaces of the active volume. The ER-like events could be γ scatters or β decays coming from “dead” regions of the TPC. There are various holes and cutouts in the TPC for LAr inlet and outlet, and events in these regions can produce visible primary scintillation, while the ionization charge will be trapped on a surface. We count all these as loss of acceptance.

Finally, the total number of single scatters is the sum of events passing the single scatter cut + events with S1-echo + events with accidental coincidence + events with S1 but no S2. The ratio of events passing the the cut to the total, as a function of S1, is shown in Fig. 6.7 and spans 94 % to 99 % in the range 0 PE to 500 PE. The decreasing behavior of

Figure 6.7: Estimated efficiency of single scatter cut to accept NRs as a function of S1.

the efficiency is due to the fact that S1-echo pulses are more likely to occur for larger S1, enhancing the denominator of Eqn. 6.1.

Why do we not move the rejected single scatters into accepted single scatters? These events are single scatters that are in some way “dirty” (they have S1-echo or they are in coincidence with another signal). If we accept these events, we risk introducing non-single-scatters to the WIMP search data set; rejecting them and counting them against loss of acceptance is more conservative. The present definition of the cut keeps the overall WIMP search simpler.

6.2.1.7 S1 start time

Definition S1 pulse start time is within $-6.1 \mu\text{s}$ to $-6.0 \mu\text{s}$.

Purpose When the DAQ receives a trigger, it sets the trigger time at $t = 0$. However, due to cable lengths and insertion delays, the physical signal on which the trigger fired may not be at exactly $t = 0$ but should have consistent offset. Since we typically trigger on S1 and the prompt portion of S1 is very fast, the pulse start time as found by the pulse finder accurately

reconstructs the S1 start time. We therefore require the S1 start time to be within a narrow window about the expected time of the S1 start, which is at $t = -6 \mu\text{s}$. Occasionally, the TPC is triggered by something other than S1, such as the tail of S2 or an S2-echo. This cut removes such triggers: because the pulse finder looks for macroscopic clusters of PE, the start time of the first pulse is often outside the expected time for triggers on S1. Of events passing surviving all previous cuts, 99.7 % pass this cut.

Efficiency The efficiency of this cut to accept NRs is evaluated from neutron events in AmBe data and is found to be unity in the relevant energy range for WIMPs (<500 PE).

6.2.1.8 S1 saturation

Definition S1 pulse does not saturate any of the ADCs.

Purpose If the S1 saturates any of the V1720 digitizers of the TPC, the S1 integrals will be biased. We therefore require that none of the digitizers be saturated. We do not place a non-saturation requirement on S2 because we do not use it for any energy estimation in this analysis. Of events surviving all previous cuts, 94.2 % pass this cut. Notice that this cut is applied before the S1 range is defined, and the events removed by this cut are all at extremely high energy.

Efficiency The efficiency of this cut to select NRs is evaluated using neutron events from AmBe calibration data and is found to be unity everywhere in the energy range relevant for this WIMP search (below 500 PE).

6.2.1.9 S1 maximum fraction

Definition S1 maximum fraction (S1MF) is smaller than a pre-defined threshold that depends on t_{drift} and S1. The S1 maximum fraction is defined as the ratio of the S1 light in the dominant S1 channel to the total S1 light. The threshold is chosen for 95 % acceptance.

Purpose One of the major backgrounds for the WIMP search in DarkSide-50 are ER events with coincident Cherenkov light. These are typically gammas that multiply scatter, once in a detector material other than the active volume and once in the active volume. When the scatter in the non-active region is in one of the fused silica windows (anode or cathode) or in one of the PMT photocathodes or in the Teflon, Cherenkov light can be produced. This class of events is dangerous because the Cherenkov light is very fast, and, in coincidence with the ER of the scatter in the active region, will bias the f_{90} up towards the NR band. Furthermore, the scatter in the non-active region will not produce an electroluminescence signal, so the S2/S1 will be biased low, again towards the NR band. The Cherenkov light is typically up to 50 PE so the effect is relevant for the WIMP region of a few tens to hundreds of photoelectrons.

The main tool in tagging events with Cherenkov light is to exploit the fact that, in the case of Cherenkov light produced in the anode or cathode window or in a PMT face, the observed light is abnormally concentrated in a single PMT. The idea of the cut, then, is to remove events where the fraction of light in a single channel is abnormally high. For each event, we need only look at the fraction of light in the dominant channel for the S1 pulse. The parameter is called S1 maximum fraction or S1MF. The nominal distribution of S1MF varies as a function of position and energy: for events closer to the top or bottom PMT array, the S1 light will naturally be more concentrated in a single channel, and for events at lower energy, the S1MF distribution will be broader due to fluctuation statistics. Therefore, we define the cut as a function of t_{drift} and S1. We divide the t_{drift} vs. S1 plane into a grid, and, for each cell in the grid, place a threshold at the 95 % quantile of the S1MF distribution from AAr data. We do not develop the thresholds on the UAr data because the statistics are too low. Figure 6.8 shows the S1 max fraction distribution for a few (S1, t_{drift}) bins.

One of the main complications of this cut is PMT afterpulsing, which occurs when electrons produced along the PMT dynode chain ionize some residual gas in the PMT. The ions drift back to the photocathode and produce a spike of charge. The ion drift time is typically on the order of $1 \mu\text{s}$, and the afterpulse signal is typically 15 PE. Like Cherenkov light, PMT afterpulses are concentrated in a single PMT. Therefore, when we place a threshold on

Figure 6.8: S1 maximum fraction distributions for a few representative $(S1, t_{\text{drift}})$ bins. The red lines indicate the threshold values on S1 maximum fraction for each bin.

S1MF, we are removing events with Cherenkov and events with afterpulsing. It is therefore not straightforward to evaluate the efficiency of the cut to remove events with Cherenkov.

We define the cut to have uniform 95 % acceptance across the t_{drift} vs. S1 plane. To help determine where the threshold should be placed, we develop a new parameter that is better at distinguishing between events with Cherenkov and events with afterpulsing. First, we define $s_{\text{max},p}$ and $s_{\text{max},l}$, where $s_{\text{max},p}$ is the fraction of prompt light in the dominant channel relative to total prompt light, and $s_{\text{max},l}$ is the fraction of late light in the dominant channel to total late light. More specifically,

$$s_{\text{max},p} = \frac{S1_{\text{prompt}}^{i_{\text{max}}}}{\sum_i S1_{\text{prompt}}^i} \quad (6.2)$$

$$s_{\text{max},l} = \frac{S1_{\text{late}}^{i_{\text{max}}}}{\sum_i S1_{\text{late}}^i} \quad (6.3)$$

where i_{max} is the ID of the channel that saw the most S1 light, and $S1_{\text{prompt}}^i$ ($S1_{\text{late}}^i$) is the total prompt (total late) S1 light of channel i and the sums are over all channels. For bulk

single scatter events, the distribution of prompt light over the PMTs should be very similar to the distribution of late light over the PMTs, so $s_{\max,p}$ should be similar to $s_{\max,l}$. For events with Cherenkov, the prompt light should be more concentrated in a single PMT than the late light, regardless of the location of the scatter in the active volume. Conversely, for events with a PMT afterpulse, the late light should be more concentrated in a single channel than the prompt light, regardless of the location of the scatter in the active volume. We build an asymmetry parameter $a_{\max}$:

$$a_{\max} = \frac{s_{\max,p} - s_{\max,l}}{s_{\max,p} + s_{\max,l}}. \quad (6.4)$$

For normal single scatters, we expect $a_{\max} \approx 0$, but statistical fluctuations may induce a large spread. For events with Cherenkov, we expect $a_{\max} > 0$, and for events with PMT afterpulsing, we expect $a_{\max} < 0$. The production of Cherenkov light and PMT afterpulsing are independent, so the rate of events with both is negligible. Figure 6.9 shows that with an S1MF cut set at 95 % acceptance, the $a_{\max}$ distribution of UAr data matches that of AAr data. We see a peak in the UAr data at $a_{\max} \approx 0.8$ prior to the S1MF cut, which is associated with events with Cherenkov light. The sharp peaks at $a_{\max} = \pm 1$ are due to statistical fluctuations where $s_{\max,p}$ or $s_{\max,l}$ goes to zero. The S1MF thresholds chosen for uniform 95 % acceptance is shown to remove events with Cherenkov at least to the relative level of AAr.

Efficiency The 95 % acceptance is derived from ER data. The behavior of S1MF should be essentially identical for ERs and NRs. Therefore we use the ER events as a proxy for NRs and we quote the NR acceptance as 95 %. The 95 % NR acceptance is also verified using neutrons from AmBe data.

6.2.1.10 S2 cuts

Definition Require xy -corrected pulse area of S2 to be >100 PE and f_{90} of S2 pulse to be <0.2 .

(a)

(b)

Figure 6.9: Distribution of asymmetry parameter $a_{\max}$ in AAr and UAr data (a) before and (b) after the S1 max fraction cut (see text for details). The data are normalized to $-0.2 < a_{\max} < 0.2$. The blue curve is AAr (^{39}Ar -dominated) data, and the red curve is UAr data. The peak in UAr data around $a_{\max} \approx 0.8$, associated with events with coincident Cherenkov light, weakens after application of the S1 max. frac. cut.

Purpose Though we do not use the electroluminescence signal for much of the WIMP search analysis beyond defining drift time, we place some very basic cuts on S2 to remove events in which there was something horribly wrong with the 2nd pulse. The minimum S2 cut is intended to remove events in which the 2nd pulse is abnormally small, inconsistent with being an S2 from either a NR or ER. Similarly, the cut on S2 f_{90} is intended to remove events in which the 2nd pulse has abnormally large f_{90} . The f_{90} of S2 is expected to be minuscule (<0.01) due to the slow rise time of the S2 pulse. This cut removes events where the S2 f_{90} is not minuscule. The effects of both of the S2 cuts are negligible.

Efficiency The efficiency of both S2 cuts to accept NRs is evaluated using AmBe data. Though S2 is reduced for NRs, the cut on S2 size is place far below the expected size of S2 for NRs. The efficiency is $>99\%$.

6.2.1.11 Fiducialization

Definition t_{drift} is in $40\ \mu\text{s}$ to $334.5\ \mu\text{s}$.

Purpose External gammas constitute the majority of the ER background in the UAr WIMP search campaign. The PMTs are the most radioactive component in DarkSide-50. Therefore, gamma scatters are concentrated at the top and bottom of the TPC, as shown in Fig. 6.10. We place a z cut to remove events near the surfaces. The cut was tuned during the AAr campaign, when the cut was used to handle a reconstruction artifact occurring at low t_{drift} . The bug in DarkArt was addressed by the time of the start of the UAr campaign, but we keep the fiducial cut the same as before to reduce the bias that would be introduced by tuning the cut again. We do not fiducialize in the radial direction as there is no evidence of side-wall events contaminating our UAr sample. The lack of radial cut also is intended to match the AAr WIMP search analysis.

We do not evaluate a NR efficiency for this cut, since the effect of this cut is accounted for in the calculation of the exposure. The fiducial mass remaining after the t_{drift} cut is (36.9 ± 0.6) kg, where the dominant uncertainty arises from the uncertainty on the shrinkage

Figure 6.10: t_{drift} distribution for single scatter events in the UAr data set after application of all cuts less the fiducial volume cut. The blue dashed lines represent the fiducial volume cut.

of the Teflon body of the TPC when cooled from room temperature to cryogenic temperature.

6.2.2 Veto cuts

The general philosophy of the veto cuts is to remove events with correlated signals in the TPC and the vetoes. We expect the topology of a WIMP interaction to be a TPC signal in anti-coincidence with the LSV and WCD; that is, there should be TPC signal only (on which we trigger) and no signal in the LSV or WCD. Neutrons, which produce NRs that are identical to the expected WIMP signal in the TPC, are likely to also interact in the LSV, so we rely on the vetoes to remove neutron events.

There are several topologies in which a neutron can scatter in the TPC and leave signals in the LSV, including a neutron scatter in the TPC and capture in the LSV or neutron thermalization in the LSV and scatter in the TPC (and capture on an inactive component). The veto cuts are designed to remove events consistent with these topologies.

The acceptance losses of the veto cuts are generically due to (non-neutron) background activity in the LSV, such as ^{14}C , ^{60}C , ^{208}Tl , etc. Events with accidental coincidence of these decays with a TPC signal are rejected by these cuts, and the accidental rate sets the NR acceptance level of the veto cuts.

The veto cuts are described in detail in Ref. [87] and summarized here.

6.2.2.1 Prompt cut

Definition Remove events with >1 PE in a 300 ns region of interest (ROI) around the time of prompt coincidence with the TPC. See Sec. 3.2 for a description of the evaluation of the ROI integral.

Purpose The prompt veto cut is intended to remove events with neutron thermalization signal in LSV. In practice, the cut predominantly removes coincident gamma scatters in the TPC and LSV. The gammas come from radioactivity in the detector components. That the prompt cut removes gammas is harmless (in fact it helps to reduce the ER backgrounds). The efficiency of the prompt cut to remove neutrons is estimated using AmBe/AmC data to be $>98\%$. The NR acceptance loss of the prompt cut is due to accidental coincidences and is estimated to be $\leq 1\%$.

6.2.2.2 Delayed cut

Definition Remove events with slider charge >6 PE anywhere in the 200 μs following the prompt time. See Sec. 3.2 for the definition of slider charge.

Purpose The purpose of the delayed cut is to remove events with neutron capture signals in the LSV. The slider window is 500 ns and the search window is from the end of the prompt window to the end of the veto acquisition gate, usually 200 μs . The cut is designed to detect neutron capture gamma rays, particularly from the $^7\text{Li}^*$ final state and the $\alpha+^7\text{Li}$ from capture on ^{10}B in the LSV. Thermal neutrons have a capture time of 22 μs , so if a neutron

captures on ^{10}B there is a high probability it is visible in the LSV. The efficiency of the cut to remove neutron capture events is estimated using AmBe plus MC to be $>99.1\%$. The NR acceptance loss is due to accidental delayed coincidences in which random backgrounds create signal above 6 PE threshold in slider search region. The NR acceptance loss is $\sim 16\%$.

6.2.2.3 Pre-prompt cut

Definition Remove events with pre-prompt signal >3 PE. Pre-prompt signals are evaluated with a 500 ns sliding window from the start of the acquisition gate to the start of the prompt window.

Purpose The purpose of this cut is to remove events with signals that precede a neutron scatter in the TPC, for example external neutrons entering the LSV from the WCD. The efficiency of this cut to remove neutron events has not been calculated but neutrons removed by this cut represent a small fraction of the neutron background. The NR acceptance loss is due to accidentals and is small, $\sim 0.1\%$.

6.2.2.4 Muon cut

Definition Remove events within 2 s after a muon. Events are muon-like if they exhibit large amplitude signals in the LSV or WCD.

Purpose Muons can produce heavy nuclei that will decay away with a lifetime of a few hundred milliseconds and produce a high energy neutron in the process. The purpose of this cut is to eliminate such delayed neutrons. The NR acceptance loss due to muons is negligible.

6.3 f_{90}

After application of all TPC and veto cuts described in Sec. 6.2, we are left with a sample of events with a single scatter in the TPC and no scatter in the veto. These are predominantly ERs from β -decays of ^{39}Ar and ^{85}Kr and external gammas that scatter once in the TPC. The goal now is to find the few, if any, NRs remaining in the sample, which could be due to WIMPs. (Considering that other experiments have already set significantly more stringent WIMP limits than DarkSide-50's sensitivity, we are not expecting to find any WIMP-like signals.)

We use two of the remaining handles at our disposal: energy and pulse shape. Here we discuss pulse shape. There is another discriminating parameter, $S2/S1$, but DarkSide-50 does not make use of it in Ref. [103] because the discrimination power of f_{90} is sufficient to set a competitive limit. We expect NR and ER events to have different f_{90} distributions, both with an $S1$ dependence. The goal is to develop an f_{90} cut to remove ERs that maximizes acceptance for NRs while minimizing leakage of ERs past the cut. We therefore need to characterize both the NR and ER f_{90} distributions, both of which depend on $S1$, primarily through fluctuation statistics. We work in the f_{90} vs. $S1$ plane. There are two approaches to characterizing the NR and ER f_{90} distributions: (1) *in situ* calibration with source data, or (2) analytic modeling of the distributions. For both the NR and ER distributions, we choose to use analytic models. The models are tuned to data, either from *in situ* calibration sources or from separate measurements transferred to DarkSide-50. We use the same basic model for both NRs and ERs, though the demands of the model are different for each.

Once the analytic model is in hand, then for any given cut through the f_{90} vs. $S1$ plane, we can evaluate the cut's acceptance for NR events and the leakage of ER events past the cut. We develop the model first, then describe how to set the cut.

6.3.1 Ratio of Gaussians

The development of the f_{90} model here follows that of Ref. [106]. For any given energy deposit, assume the number of photoelectrons in the prompt and late windows, N_p and N_l ,

respectively, are independent and Gaussian distributed with means μ_p and μ_l and variances σ_p^2 and σ_l^2 , respectively. Then the total number of photoelectrons $N_{\text{tot}} = N_p + N_l$ is also Gaussian distributed with mean $\mu_{\text{tot}} = \mu_p + \mu_l$ and variance $\sigma_{\text{tot}}^2 = \sigma_p^2 + \sigma_l^2$. We wish to find the distribution of the ratio $X = N_p/N_{\text{tot}}$.

Hinkley [107] has worked out the probability density function for the ratio of two correlated Gaussian random variables. The correlation ρ between N_p and N_{tot} is

$$\rho = \frac{\sigma_p}{\sigma_{\text{tot}}} \quad (6.5)$$

We use the approximation given by Eq. 9 of Ref. [107], which holds when $0 < \sigma_{\text{tot}} \ll \mu_{\text{tot}}$. Recasting in terms of prompt and late means and variances only, one obtains

$$P(x) = \frac{\sigma_l^2 \mu_p x + \sigma_p^2 \mu_l (1 - x)}{\sqrt{2\pi} [\sigma_l^2 x^2 + \sigma_p^2 (1 - x)^2]^{3/2}} \times \exp \left[-\frac{(\mu_l x - \mu_p (1 - x))^2}{2(\sigma_l^2 x^2 + \sigma_p^2 (1 - x)^2)} \right] \quad (6.6)$$

where we use x for the variable f_{90} for cleaner notation. In practice, it is useful to reparametrize μ_p and μ_l as

$$\mu_p = \hat{f} \mu_{\text{tot}} \quad (6.7)$$

$$\mu_l = (1 - \hat{f}) \mu_{\text{tot}} \quad (6.8)$$

where $\hat{f} = \mu_p/\mu_{\text{tot}}$ and is the mean value of f_{90} . Furthermore, we assume that N_p and N_l are dominated by Poisson counting statistics, but to allow for additional sources of variance, we decompose the prompt and late variances as

$$\sigma_p^2 = \mu_p + \sigma_{p,\text{add}}^2 \quad (6.9)$$

$$\sigma_l^2 = \mu_l + \sigma_{l,\text{add}}^2 \quad (6.10)$$

where $\sigma_{p,\text{add}}^2$ and $\sigma_{l,\text{add}}^2$ are additional terms due to SPE variance, electronics noise, etc. The ratio-of-Gaussians PDF then takes the form

$$P(x) = \frac{\frac{1}{\mu_{\text{tot}}} \hat{f}(1 - \hat{f}) + \frac{1}{\mu_{\text{tot}}^2} (\sigma_{p,\text{add}}^2 (1 - \hat{f})(1 - x) + \sigma_{l,\text{add}}^2 \hat{f}x)}{\sqrt{2\pi} \sigma^3} \times \exp \left[-\frac{(x - \hat{f})^2}{2\sigma^2} \right] \quad (6.11)$$

where

$$\sigma^2 = \frac{1}{\mu_{\text{tot}}} (\hat{f}(1 - \hat{f}) + (x - \hat{f})^2) + \frac{1}{\mu_{\text{tot}}^2} (\sigma_{p,\text{add}}^2 (1 - x)^2 + \sigma_{l,\text{add}}^2 x^2). \quad (6.12)$$

Notice that σ^2 depends on x . This form allows us to develop some intuition for the ratio-of-Gaussians distribution. First, consider the ideal case when the additional noise terms are zero, $\sigma_{p,\text{add}}^2 = \sigma_{l,\text{add}}^2 = 0$:

$$P(x) = \frac{\frac{1}{\mu_{\text{tot}}} \hat{f}(1 - \hat{f})}{\sqrt{2\pi}\sigma^3} \times \exp \left[-\frac{(x - \hat{f})^2}{2\sigma^2} \right] \quad (6.13)$$

One sees an expression resembling that of the large N limit of a binomial distribution, scaled by $1/\mu_{\text{tot}}$, i.e. a Gaussian with mean $\hat{f}$ and variance $\hat{f}(1 - \hat{f})/\mu_{\text{tot}}$. The $(x - \hat{f})^2/\mu_{\text{tot}}$ term in σ^2 represents the departure from the binomial, which arises from the fact that N_{tot} is variable. The variances are dominated by the Poisson counting statistics, and we find that $\sigma_{\text{tot}}^2 \approx \sqrt{\mu_{\text{tot}}}$ is sufficiently small that the approximation is very accurate, even down to $\mu_{\text{tot}} = 10$ PE.

For any given energy deposit, the number of observed photoelectrons N_{tot} is a random variable. We usually do not know the energy of any particular interaction, and events have a broad spectrum of energies. In practice, we fit Eqn. 6.6 to slices in S1. For any particular slice, there are events with many different energies, so the f_{90} distribution for that slice is a mixture of different PDFs. This leads to distortions in the f_{90} distributions. In particular, it causes Eqn. 6.6 to overestimate the tails of the distributions. To see this, we use a simple toy Monte Carlo to simulate photoelectron times for ERs with $\mu_{\text{tot}} = 42.5$ PE, $\hat{f} = 0.338$, and no additional variance terms beyond Poisson counting statistics. The simulated f_{90} vs. S1 distribution is shown in Fig. 6.11a. Taking a slice of 40 PE to 45 PE and fitting Eqn. 6.6, we see that the ratio-of-Gaussians model overestimates the tail of the f_{90} distribution, as shown in Fig. 6.11b. Furthermore, we can model a flat energy spectrum by letting the mean of the N_{tot} distribution be uniformly distributed between 10 PE to 100 PE in the simulation. The events in the 40 PE to 45 PE slice are then from a variety of simulated energies, yet the f_{90} distribution for the slice is still not well modeled by a ratio of Gaussians.

Despite the overestimation of the tails, we proceed using the ratio-of-Gaussians model for defining the f_{90} cut because it has been used in the literature before (Refs. [106, 108, 109]). Though the practice of fitting the f_{90} PDF to data from a narrow S1 slice breaks the assumption of the model, it leads to a conservative estimate of the leakage of ER events past

(a)

(b)

Figure 6.11: Toy Monte Carlo simulation of f_{90} for ER data with $\mu_{\text{tot}} = 40$ PE and $\hat{f} = 0.338$.

(a) f_{90} vs. $S1$ distribution for simulated ERs. (b) f_{90} distribution of the 40 PE to 45 PE slice, fitted with Eqn. 6.6 (red).

the f_{90} cut due to the model overestimating the tails of the f_{90} distribution of each S1 slice. Furthermore, the overestimation is visible only deep in the tails of the f_{90} distributions. The agreement in the bulk of the distribution is fairly good and is sufficient for modeling the NR f_{90} response, which demands agreement at only the 90% quantile.

6.3.2 Fitting ER data

The next step towards establishing an f_{90} cut is to tune the ratio-of-Gaussians model to the DarkSide-50 data. There are four parameters in Eqn. 6.11: μ_{tot} , $\hat{f}$, $\sigma_{p,\text{add}}^2$, and $\sigma_{l,\text{add}}^2$. We treat each S1 slice independently, and we fix μ_{tot} to the central S1 value of each slice. Beyond the variance due to Poisson counting statistics, the additional sources of prompt and late variance include the finite SPE resolution, a TPB emission term, and electronics noise.

$$\sigma_{p,\text{add}}^2 = \mu_p(\sigma_{\text{SPE}}^2 + \sigma_{\text{TPB}}^2) + \sigma_{p,\text{elec}}^2 \quad (6.14)$$

$$\sigma_{l,\text{add}}^2 = \mu_l(\sigma_{\text{SPE}}^2 + \sigma_{\text{TPB}}^2) + \sigma_{l,\text{elec}}^2 \quad (6.15)$$

The form of this decomposition is motivated in the next section.

6.3.2.1 Variance terms in f_{90} model

We begin from the primary scintillation process. Let us model the number of UV scintillation photons as Poisson distributed:

$$N_{\text{UV}} \sim \text{Pois}(\mu_{\text{UV}}) \quad (6.16)$$

where μ_{UV} is the mean number of UV photons. The following arguments hold regardless of whether N_{UV} represents the number of prompt, late, or total UV photons.

We now develop a probabilistic model for the TPB conversion of UV photons to visible photons. From Ref. [110], we know that, on average, each UV photon is converted into 1.2 visible photons. However, the conversion of one UV photon to ≥ 3 visible photons is essentially kinematically forbidden. Therefore, the conversion to one or two visible photons is essentially a binomial process. We model the number of visible photons as

$$N_{\text{vis}} = X_1 + X_2 + \dots + X_{N_{\text{UV}}} \quad (6.17)$$

where

$$X_i \sim 1 + \text{Bin}(1, p) \quad (6.18)$$

and $p = 0.2$. The X_i are independent and identically distributed, and X_i and N_{UV} are independent. We have $E[X] = 1 + p$ and $\text{Var}[X] = p(1 - p)$. If the kinematic cutoff of TPB were higher, and one UV photon could be converted to more than 2 visible photons, we would replace the binomial process with a multinomial.

Notice that N_{vis} is a compound distribution. The law of total expectation says

$$\mu_{\text{vis}} = E[N_{\text{vis}}] = E[E[N_{\text{vis}}|N_{\text{UV}}]] \quad (6.19)$$

$$= E[N_{\text{UV}}E[X]] \quad (6.20)$$

$$= E[N_{\text{UV}}]E[X] \quad (6.21)$$

$$= \mu_{\text{UV}}(1 + p) \quad (6.22)$$

and the law of total variance says

$$\sigma_{\text{vis}}^2 = \text{Var}[N_{\text{vis}}] = E[\text{Var}[N_{\text{vis}}|N_{\text{UV}}]] + \text{Var}[E[N_{\text{vis}}|N_{\text{UV}}]] \quad (6.23)$$

$$= E[N_{\text{UV}}\text{Var}[X]] + \text{Var}[N_{\text{UV}}E[X]] \quad (6.24)$$

$$= E[N_{\text{UV}}]\text{Var}[X] + \text{Var}[N_{\text{UV}}]E[X]^2 \quad (6.25)$$

$$= \mu_{\text{UV}}p(1 - p) + \mu_{\text{UV}}(1 + p)^2 \quad (6.26)$$

$$= \mu_{\text{UV}}(1 + 3p). \quad (6.27)$$

Notice that, to remove the TPB effects, we can set $p = 0$, and we recover Poisson statistics.

We now model the conversion of visible photons to photoelectrons (ignoring collection efficiency of the TPC). We follow a similar procedure to that of TPB treatment. We make a simplifying assumption that single PE are Gaussian⁴ distributed with 40% resolution. The number of PE is modeled as:

$$N_{\text{PE}} = Y_1 + Y_2 + \dots + Y_{N_{\text{vis}}} \quad (6.28)$$

⁴If we want a more sophisticated model of the SPE pdf, it's easily accommodated by replacing the Gaussian. We need only know the SPE resolution.

where

$$Y_i \sim \text{Gaus}(1, \sigma_{\text{SPE}}) \quad (6.29)$$

and $\sigma_{\text{SPE}} = 0.4$. The Y_i are independent and identically distributed, and Y_i and N_{vis} are independent. Notice that N_{PE} is another compound distribution with mean

$$\mu_{\text{PE}} = \text{E}[N_{\text{PE}}] = \text{E}[N_{\text{vis}}]\text{E}[Y] \quad (6.30)$$

$$= \mu_{\text{vis}} \quad (6.31)$$

and variance

$$\sigma_{\text{PE}}^2 = \text{Var}[N_{\text{PE}}] = \text{E}[N_{\text{vis}}]\text{Var}[Y] + \text{Var}[N_{\text{vis}}]\text{E}[Y]^2 \quad (6.32)$$

$$= \mu_{\text{vis}}\sigma_{\text{PE}}^2 + \sigma_{\text{vis}}^2. \quad (6.33)$$

With some simple algebraic manipulation of the previous results, we can rewrite the variance as

$$\sigma_{\text{PE}}^2 = \mu_{\text{vis}}\sigma_{\text{SPE}}^2 + \mu_{\text{vis}}\frac{1+3p}{1+p} \quad (6.34)$$

$$= \mu_{\text{PE}}(1 + \sigma_{\text{TPB}}^2 + \sigma_{\text{SPE}}^2) \quad (6.35)$$

where we have defined $\sigma_{\text{TPB}}^2 = (1+3p)/(1+p) - 1 = 2p/(1+p)$. With the addition of an electronics noise term added quadratically, we recover the form of the noise terms given by Eqns. 6.9 and 6.14 and Eqns. 6.10 and 6.15.

The σ_{SPE}^2 terms reflect the variance in the observed number of PE due to the finite single PE resolution of the PMTs. All the PMTs have roughly equal SPE response and have approximately 40 % resolution, as determined by the laser calibration data. Therefore, we set $\sigma_{\text{SPE}} = 0.4$.

The σ_{TPB}^2 terms reflect the variance in the observed number of PE due to the amplification effect arising from the TPB. Using the average 1.2 amplification factor found in Ref. [110], we expect that $\sigma_{\text{TPB}} = 0.58$. However, we empirically find that $\sigma_{\text{TPB}} = 0.49$ gives a better fit to the DarkSide-50 f_{90} widths. The values are in approximate agreement.

The electronics noise terms are extracted directly from data. Because the prompt window is so short, the prompt electronics noise is not expected to vary with energy. The prompt

noise term is fixed to the small value of 0.2. The noise in the late window is more dependent on energy. In the nearly $7\ \mu\text{s}$ late window, there is an interplay between integrated baseline and signal. In regions of pure baseline, the moving average algorithm of the baseline finder ensures that the integrated noise is zero on every channel, while in regions with photoelectrons, the baseline is linearly interpolated beneath the signal and we cannot guarantee that the baseline noise integrates to zero. Therefore, for higher energy events, where there are more linearly interpolations to estimate the baseline, the contribution of electronics noise to the f_{90} variance is greater. We empirically fit the late noise term with a 3rd order polynomial.

6.3.2.2 Establishing f_{90} means

With the additional variance terms in hand, the only free parameter in Eqn. 6.6 is the mean value of f_{90} . There are a variety of ways one could obtain $\hat{f}$: fit directly to the UAr data, fit to the AAr data, or use values from the literature. Because the AAr data has extremely high statistics compared to the UAr data, we choose to fit the AAr data. Furthermore, as will be discussed in Sec. 6.3.2.3, there is evidence that the UAr data is not populated purely by single scatter ER events. Fitting to the single scatters of essentially purely ^{39}Ar decays of AAr data, we reduce the uncertainty on the f_{90} mean. Once the f_{90} means are established, we need to scale the model to the statistics of the UAr data in order to estimate the leakage of ER events past the f_{90} cut. We use a simple procedure of scaling the model to have area equal to the total number of UAr events in each S1 slice. Figure 6.12 shows examples of the analytic model fitted to AAr data and scaled to UAr data. The agreement between the UAr data and the scaled fit is good. In particular, the effect of overestimation of the tails is weakened because of the lower statistics in UAr data.

We are enabled to do the simple scaling of the analytic model from AAr to UAr data in large part because we know the LY did not change between the two campaigns. In fact, we require many parameters to be the same, including the electronics noise, SPE resolution, and TPB effects. Though some of these parameters varied between the AAr and UAr campaigns, their changes were not large enough to invalidate the scaling procedure, especially considering

Figure 6.12: f_{90} distributions for various S1 slices of AAr data (black) and UAr data (black). The analytic model (red) for f_{90} is fitted to the AAr data and scaled to the statistics of the UAr data.

that their contributions to the variance are subdominant to the Poisson counting statistics.

The scaling of the analytic model to the total number of events in each S1 bin is a conservative procedure. The model is intended only for single scatter events, yet we have evidence for non-single-scatter ER events in the UAr data, as described in the next section. Scaling to the total number of events will tend to further overestimate the tails of the f_{90} distributions, leading to more conservative placement of the f_{90} cut.

6.3.2.3 Non-single-scatter ERs in UAr data

The analytic ratio-of-Gaussians model is intended to describe the f_{90} distribution of single scatter events only. In the AAr data, the abundance of ^{39}Ar decays is so overwhelming that even deep in the tails of the f_{90} distributions, we are confident the statistics are dominated by single scatters and the model applies (with the caveat of overestimation of the tails due to taking narrow S1 slices). However, in UAr data, there is evidence for non-single-scatter ER events contributing to the high side of the f_{90} distributions. Non-single-scatter ER events can have a variety of topologies, the most common of which are unresolved multiple scatter gammas and gammas with coincident Cherenkov light. That there are relatively more gamma events in UAr data is expected: we do not expect the rate of external gammas to change between the two campaigns, while the rate of β decays has dropped by a factor 400 compared to AAr data. Here we describe the basic evidence for a higher fraction of non-single-scatter gammas in the UAr data with f_{90} biased higher relative to the main ER band.

Unresolved multiple scatter gammas can bias f_{90} upwards compared to single scatter ER events. The topology of unresolved gammas is two or more scatters at the same z . Typically, multiple scatter gammas will interact at two distinct z positions. The S1s from each scatter will be piled on top of each other, but the S2s will be separated and the single scatter cut will remove such events. However, when the two scatters occur on the same z plane (to within a few mm), we cannot resolve the S2s, instead reconstructing the two S2s into a single pulse. (We do not make any attempt to use the xy reconstruction to identify events

with multiple scatters at the same z but different transverse position.) Furthermore, It is now well known that the f_{90} means turn upwards at low energy [106, 108], and an event composed of two scatters will necessarily have an S1 with an f_{90} sampled from two lower energy bins. Therefore, the combined S1 has an f_{90} biased upwards.

The other main topology of ER events that could have large f_{90} are gammas with coincident Cherenkov light. As described in Sec. 6.2.1, typically, the gamma will Compton scatter once in either one of the fused silica windows, in a PMT photocathode, or in the Teflon producing a recoil electron that produces Cherenkov light. The gamma may also scatter in the active volume. The Cherenkov light and the primary scintillation signal will be coincident (to within a few ns). The Cherenkov light is very fast and the addition of Cherenkov light to the primary scintillation can bias the observed f_{90} value upwards. Though the S1 maximum fraction cut is designed to remove events with Cherenkov light coming from the fused silica windows or the PMT photocathodes, the efficiency of the cut to remove such events was never well established. Furthermore, the cut is not efficient at removing events with Cherenkov coming from the Teflon.

Both classes of non-single-scatter events are associated with gammas, and in DarkSide-50 essentially all gammas are of external origin (originating from outside the active volume). This gives us several signatures that lead us to believe there is a higher fraction of non-single-scatter gammas in UAr data than in AAr data that leads to enhancement of the high side of the f_{90} distributions. Most of the events tagged by the LSV prompt cut (see Sec. 6.2.2) are thought to be gammas scattering once in the LSV and (at least) once in the TPC. Events that fail the cut should then be correlated with a higher fraction of high f_{90} events. Indeed, this is what we observe, as shown in Fig. 6.13. Furthermore, gammas have a finite scattering length in UAr (rough number), so more of the gamma scatters should be concentrated near the edge of the TPC. Conversely, events from the center of the TPC should be more purely single scatter β decays. Indeed we see matching of the f_{90} tails between AAr and the UAr data from the central core of the TPC ($r < 10$ cm and $135 \mu\text{s} < t_{\text{drift}} < 235 \mu\text{s}$). There is also a 1% upward shift in the ER f_{90} mean from AAr data to UAr data. The shift is likely due to a higher fraction of unresolved multiple scatter gammas in UAr data. We do not

Figure 6.13: Evidence for non-single-scatter ER events in the UAr data. Shown are the f_{90} distribution of events with $60 \text{ PE} < S1 < 120 \text{ PE}$ from AAr data (black) and UAr data that pass (green) and fail (pink) the veto coincidence cuts. The data are normalized by area. Plot from X. Xiang.

account for this shift when we scale the analytic model from AAr to UAr. Such a small shift has negligible effect on the placement of the f_{90} cut.

Evidently there are more non-single-scatter ER events in UAr data than in AAr data which can contribute to the high side of the f_{90} distributions. Though we know they are associated with external gammas, we do not yet have a detailed understanding of their origin or spectrum. We do not incorporate them any further in this analysis other than to note their existence. They do not prevent us from drawing a WIMP search region. The main concern is their contribution to the background, but since the ratio-of-Gaussians model (as we use it) is known to overestimate the f_{90} tails, at least in the low energy region, and since we will place the cut far from the ER distribution, we expect that the background estimate is still small ($\ll 1$).

6.3.3 Fitting NR data

Just as we use an analytic model to describe the f_{90} distribution of ERs, we use the same model to describe the f_{90} distribution of NRs, with the only difference being in the f_{90} means. (We need not worry about the normalization factor because we use quantiles for NR distributions rather than event counts.) We keep all parameters of the ratio-of-Gaussians other than $\hat{f}$ the same. All the additional noise parameters (SPE variance, TPB variance, electronics noise) are carried over from ERs to NRs. The challenge is to determine the f_{90} means for NRs in DarkSide-50. Whereas we have abundant single scatter events in the AAr data to tune the ER f_{90} response, we have no clean sample of NRs in DarkSide-50. *In situ* neutron calibration data from AmBe and AmC sources are contaminated with ER events and inelastic scatters. Fortunately, the only parameter we need is the f_{90} mean, which is an intrinsic property of LAr for a given drift electric field, and the SCENE experiment has measured the f_{90} means at a variety of fields [109].

The SCENE experiment exposed a small LAr TPC (with approximately 3" diameter by 3" height cylindrical active volume) to a pulsed mono-energetic neutron beam from the Tandem Van de Graaff accelerator at the University of Notre Dame. The neutrons were

Source	Experiment	Energy [keV _{nr}]	f_{90} median
${}^7\text{Li}(p, n)$	SCENE	16.9	0.583
${}^7\text{Li}(p, n)$	SCENE	25.4	0.642
${}^7\text{Li}(p, n)$	SCENE	36.1	0.672
${}^7\text{Li}(p, n)$	SCENE	57.2	0.720
AmBe	DarkSide-50	75	0.727
AmBe	DarkSide-50	100	0.750
AmBe	DarkSide-50	125	0.758
AmBe	DarkSide-50	150	0.763
AmBe	DarkSide-50	175	0.772
AmBe	DarkSide-50	200	0.770

Table 6.1: Measurements of f_{90} medians from SCENE [109] and AmBe calibration data in DarkSide-50.

produced by the ${}^7\text{Li}(p, n)$ reaction, with the proton beam from the accelerator impinging on a ${}^7\text{Li}$ target. Scattered neutrons were detected by liquid scintillator counters at select angles and distances with respect to the TPC. The scattering angle defined the recoil energy and the time-of-flight from the TPC to the counter allowed to select a pure sample of neutron scatters. These events were then studied to determine scintillation efficiency and f_{90} means for NRs. The measurements were performed at a variety of recoil energies up to 57.3 keV_{nr}. The f_{90} distributions for NRs were verified to match the ratio-of-Gaussians model to the level of statistics available from SCENE. The SCENE f_{90} means at the various recoil energy points are given in Tab. 6.1. The conversion from recoil energy to observable S1 is described in Sec. 6.5. Between the SCENE data points, we linearly interpolate to obtain the f_{90} means.

The SCENE data is available only up to 57.3 keV_{nr}. For higher energies, we use the *in situ* AmBe calibration data to estimate the f_{90} mean. In this regime, the overlap of the ER and NR bands is minimal. The contamination of inelastic neutron scatters and coincident neutron plus gamma scatters is also estimated to be small enough that the bulk estimate

of the f_{90} mean is not heavily biased. We use the f_{90} medians in place of f_{90} means to further reduce the bias. The f_{90} medians for a selection of NR energies using AmBe data in DarkSide-50 is shown in Tab. 6.1.

In summary, we use a combined set of SCENE measurements and AmBe calibration data to describe the f_{90} means across the full energy region of interest for the WIMP search. The ratio-of-Gaussians model is then used to describe the widths of the f_{90} distributions.

6.4 WIMP search region

Now that we can describe the f_{90} distributions, we can now define the WIMP search region in the f_{90} vs. S1 plane. We define the region so as to maximize NR acceptance, while minimizing the ER background. We require the ER background to be $\ll 1$ in order to maintain the zero-background requirement. In practice, we choose the box such that there are 0.1 total ER leakage events into the box.

Beyond 100 PE, the ER and NR f_{90} bands are very well separated, so in the high energy regime, we can maximize NR acceptance without worrying about ER leakage. A few NR acceptance contours are shown in Fig. 6.14. We use the 90 % acceptance contour as the f_{90} cut in the high energy region. Events with f_{90} greater than this contour are considered NRs and therefore potential WIMP candidates. We truncate the acceptance at 90 % because the SCENE statistics are not enough to verify the agreement of the analytic model to the data any deeper in the tails. In any case, extending the acceptance contour to 99 % provides only a marginal gain in sensitivity. The ER leakage beyond the 90 % NR acceptance contour for $S1 > 100$ PE is negligible.

Below 100 PE, the ER and NR f_{90} distributions overlap. The two main factors are: (1) fluctuation statistics giving broader widths, and (2) the ER and NR f_{90} means turn towards each other at low energies. So we must balance NR acceptance and ER leakage. This is the critical region for sensitivity to lower mass WIMPs. We use the analytic model for f_{90} to derive a “leakage contour” in the f_{90} vs. S1 plane, which is a curve with fixed ER leakage per bin. The leakage is defined as the integral of the f_{90} distribution from the threshold

Figure 6.14: NR acceptance contours in the f_{90} vs. $S1$ plane.

value to 1. A few ER leakage contours are shown in Fig. 6.15. We choose the contour with 0.01 leakage events per 5 PE $S1$ bin. This leakage contour intersects the 90% NR acceptance contour at 95 PE and reaches $f_{90} = 1$ at 20 PE, corresponding to 14 bins and a total estimated ER leakage of 0.14 events.

The lower edge of the $S1$ range is chosen in conjunction with the f_{90} cut. We want to push the WIMP search region to as low energy as possible to increase sensitivity to lower mass, but there is a limit: at low energies, the ER and NR f_{90} bands overlap so much that the zero-background requirement kills the NR acceptance. We set a lower $S1$ limit of 20 PE, approximately where the ER leakage curve hits $f_{90} = 1$. The upper edge of the $S1$ range is 460 PE, chosen to match that of the AAr campaign. The range was chosen in the AAr campaign rather arbitrarily. In any case, the upper range does not matter a lot because of the exponential fall-off of the WIMP spectrum.

The WIMP search region is shown in Fig. 6.16, along with all the UAr events passing the cuts described in Sec. 6.2. There are no events in the WIMP search region.

Figure 6.15: ER leakage contours in the f_{90} vs. S1 plane.

Figure 6.16: WIMP search results in the f_{90} vs. S1 plane with 70 live-days of data obtained with UAr. Red curves indicate the NR acceptance contours. The WIMP search region is outlined in blue.

Source	Experiment	Energy [keV _{nr}]	$\mathcal{L}_{\text{eff}}$
${}^7\text{Li}(p, n)$	SCENE	16.9	0.202
${}^7\text{Li}(p, n)$	SCENE	25.4	0.224
${}^7\text{Li}(p, n)$	SCENE	36.1	0.265
${}^7\text{Li}(p, n)$	SCENE	57.2	0.282

Table 6.2: Measurements of $\mathcal{L}_{\text{eff}}$ from the SCENE experiment [109]. Above 57.3 keV_{nr}, $\mathcal{L}_{\text{eff}}$ is conservatively assumed to be constant.

6.5 Energy scale

Should we see any candidate WIMP events, we would like to know the observed recoil energy. There are a couple ways one can translate the event observables (S1 and/or S2) into recoil energy. We set the energy scale using $\mathcal{L}_{\text{eff}}$, the scintillation efficiency of NRs relative to that of a standard calibration point. The scintillation efficiency is expected to be dependent on recoil energy and on applied drift field. The SCENE experiment has measured the scintillation efficiency of NRs in LAr at 200 V/cm drift field [109], and both SCENE and DarkSide-50 have measured the scintillation efficiency of ERs from ${}^{83m}\text{Kr}$ in LAr at null field. Therefore, we use the null field ${}^{83m}\text{Kr}$ light yield as the cross-calibration point between the two experiments:

$$\mathcal{L}_{\text{eff}, {}^{83m}\text{Kr}}(E_{\text{nr}}, \mathcal{E}_{\text{d}}) = \frac{S1_{\text{nr}}(E_{\text{nr}}, \mathcal{E}_{\text{d}})/E_{\text{nr}}}{S1_{\text{Kr}}(\mathcal{E}_{\text{d}} = 0)/E_{\text{Kr}}} \quad (6.36)$$

where E_{Kr} is 41.5 keV, E_{nr} is the nuclear recoil energy and $\mathcal{E}_{\text{d}}$ is the drift field. If we assume any particular event at 200 V/cm drift field with observed scintillation signal S1, then its energy is easily obtained from Eqn. 6.36: $E_{\text{nr}} = S1/\text{LY}/\mathcal{L}_{\text{eff}}$. The relative scintillation efficiency $\mathcal{L}_{\text{eff}}$ is detector independent.

The SCENE data points exist up to 57 keV_{nr}, as shown in Tab 6.2. Above this point, we conservatively assume $\mathcal{L}_{\text{eff}}$ is constant. Measurements from other LAr TPC experiments [67] indicate this assumption is approximately correct.

6.6 Exposure and NR acceptance

For each of the cuts defined above, we evaluate the NR acceptance. There are two ways to account for NR acceptance: we can consider the cuts as affecting the exposure, or we can consider the effects on the WIMP acceptance. The exposure is defined as fiducial volume mass times live time. The fiducial volume mass is governed by the fiducial cut, which is a z cut only. The fiducial volume mass is (36.9 ± 0.6) kg, where the uncertainties are dominated by the uncertainty in the diameter of the Teflon wall at cryogenic temperature. The live time represents the total amount of time that DarkSide-50 was sensitive to WIMPs and is discussed in more detail below. The total live time is found to be 70.9 d, giving a total exposure of (2616 ± 43) kg d.

The NR acceptance represents the probability that we would observe a WIMP, should one interact in the TPC during the WIMP search. The NR acceptances are discussed with their respective cuts in Sec. 6.2 and the total NR acceptance as a function of energy is discussed below.

6.6.1 Total live time

Evaluation of livetime: recall that for each trigger we record a livetime and an inhibit time. The livetime is the elapsed time from the end of the previous trigger's inhibit to the present event's trigger time. The inhibit time is the length of the inhibit after the previous event's trigger. (The DAQ begins reading out the data immediately after the acquisition window. The inhibit window is longer than the acquisition, and so the present event's inhibit is not known at the time of readout. Therefore, a given event's inhibit variable is actually the inhibit window for the previous event.) With the G2 trigger, the livetime is the sum of all livetimes to the previous recorded trigger, and the inhibit time is the sum of all inhibit times to the previous trigger.

To evaluate the total livetime, we use the following procedure. Start with the total livetime of *all* recorded events. This represents the total time that a WIMP could have triggered the TPC, assuming all detector components were operating perfectly 100 % of the

time. Of course, this is an ideal scenario, and there were some instances that one sub-system or another was malfunctioning. The basic quality cuts are intended to remove such instances, and we evaluate the effects of the cuts as loss of livetime. The specific tabulation for each cut is as follows and summarized in Table 6.3:

NChannels cut: Events fail this cut when there are fewer than 38 channels in the raw data file. This indicates a DAQ problem, which could have been present for some or all the time since the last trigger. To be safe, we consider ourselves dead all the way back to the previous trigger. For events that pass the cut, we do not remove anything from the total livetime.

Baseline found cut: Events fail this cut when the DarkArt baseline finder fails to find a baseline. The primary cause of failure is the presence of bipolar noise. We should NOT penalize ourselves for such events; we were happily sensitive to WIMPs in the time from the previous event up to a bipolar pulse. For events failing and passing this cut, we do not remove anything from the total livetime.

Live+Inhibit time cut: An event fails this cut when it is too close to the previous trigger. That is, we consider ourselves insensitive to WIMPs for 1.35 ms following any trigger. The inhibit time is a fixed 810 μ s for UAr data, so the cut is equivalent to requiring livetime to be $>540 \mu$ s. For events failing the cut, we should remove the livetime from the total (which, by definition, will be $<540 \mu$ s). For events passing the cut, we should remove all 540 μ s in which we consider ourselves dead.

Veto present cut: Events fail this cut when there was a problem with the veto DAQ. Either there is no corresponding veto event in the veto data stream, or the GPS timestamps are too far misaligned. The DAQ problem may have occurred any time since the previous trigger. To be safe, we consider ourselves dead all the way back to the previous trigger. So for events failing this cut, we remove the event's livetime. But we have already removed 540 μ s as a result of the previous step! So for events failing this cut, remove the event's livetime less 540 μ s. For events passing the cut, we do not remove anything from the total livetime.

Cut	Amount to remove		Residual livetime [days]
	Accepted events	Rejected events	
No cuts	-	-	71.6
NChannels	nothing	livetime	71.6
Baseline found	nothing	nothing	71.6
Live+inhibit time	540 μ s	livetime	71.5
Veto present	nothing	livetime - 540 μ s	70.9

Table 6.3: Tabulation of livetime after successive application of quality cuts

6.6.2 Total NR acceptance

The nuclear recoil acceptances are described for each cut in Sec. 6.2. The combined NR acceptance is evaluated by simply multiplying each cut’s acceptance. This assumes the cuts are uncorrelated, which is a conservative assumption. The combined NR acceptance as a function of energy is shown in Fig. 6.17, as well as separate curves for the f_{90} cut and all other cuts. The main acceptance loss comes from the veto cuts, due to accidentals.

6.7 Exclusion curve

For any particular WIMP mass m_χ and spin-independent WIMP-nucleon scattering cross section σ , and assuming the standard halo model, one can predict the recoil energy spectrum, as described in Sec. 1.3.1. The question to ask is, given DarkSide-50’s ability to detect recoils of a finite energy range with finite efficiency, is the observation of zero recoil events in a given exposure consistent with the expected rate spectrum for this particular mass and cross section? To answer this, we must consider the number of background events that we expect to observe. Fortunately, as described previously in this chapter, the expected background is small ($\ll 1$).

For each WIMP mass, we vary σ until we find the minimum cross section that DarkSide-50 can exclude at the 90 % confidence level (C.L.). Combining the results for many WIMP

Figure 6.17: NR acceptance as function of energy. Evaluated separately for TPC cuts (less f_{90}), veto cuts, and f_{90} cut.

masses defines the exclusion limit in the m_χ vs. σ plane. We use the standard convention that for zero observed events and zero expected background, we can exclude at the 90 % C.L. the cross section that corresponds to 2.3 expected signal events.⁵ Importantly, σ is an overall scale factor in Eqn. 1.16, so we need not vary σ through a wide variety of values; we can evaluate the recoil spectrum using a large value of σ and scale the spectrum to the 90 % C.L. value.

For each WIMP mass, we run a toy MC to determine the cross section excluded at the 90 % C.L. The MC simulates WIMPs in the f_{90} vs. S1 plane, accounting for the detector effects of DarkSide-50, including exposure, the NR scintillation efficiency, the f_{90} distribution for NRs, a simple detector resolution function, and NR acceptance. For each WIMP mass, we do the following steps:

1. Derive the dR/dE_R vs. E_R spectrum corresponding to WIMP mass m_χ and an artificially large cross section σ . We use a large value of σ so that at the exposure used in

⁵The more modern frequentist approach given by Feldman and Cousins [111] dictates that the 90 % C.L. should correspond to 2.44 signal events. We used the more historical value of 2.3 for Ref. [103].

this analysis, the simulated signal will have very high statistics.

2. Generate WIMP events corresponding to the exposure used for the WIMP search; i.e. generate a sample of E_R .
3. Convert E_R to an observable $S1_{\text{nr}}$ using Eqn. 6.36; i.e. multiply by LY and $\mathcal{L}_{\text{eff}}$.
4. Smear each $S1_{\text{nr}}$ to account for detector resolution. We use a simple Poisson distribution with mean $S1_{\text{nr}}$.
5. Account for the efficiency of the cuts to accept NRs (apart from the f_{90} cut which is handled separately) by keeping only a portion of the events. The fraction is $S1$ -dependent and is given by the total cut efficiency curve shown in Fig. 6.17.
6. For each E_R , generate a value of f_{90} by drawing from the f_{90} distribution corresponding to $S1_{\text{nr}}$. We use the analytic model of the f_{90} distribution described in Sec. 6.3.
7. Draw the f_{90} vs. $S1_{\text{nr}}$ distribution for the simulated WIMP events and determine the number of events N_{acc} that land inside the WIMP box defined in Sec. 6.4. The value of N_{acc} should be large since we used an artificially large value of σ in step 1.
8. Scale the value of σ used in step 1 by $2.3/N_{\text{acc}}$. The scaled cross section is the DarkSide-50 90 % C.L. exclusion limit for a WIMP of mass m_χ .

The main steps of the toy MC are illustrated in Fig. 6.18. The procedure is repeated for all WIMP masses. We probe WIMP masses of 10 GeV to 10 TeV. This procedure sweeps out the 90 % CL exclusion contour for the null result of the 70 d WIMP search with zero background performed with UAr in DarkSide, shown in Fig. 6.19. The exclusion limit reaches a minimum value of $3.1 \times 10^{-44} \text{ cm}^2$ at 100 GeV. The results are compared to the previous results from DarkSide-50 [76] as well as recent results from other direct detection experiments. The present results represent the most stringent WIMP dark matter limit using a liquid argon target, to date.

(a)

(b)

(c)

(d)

Figure 6.18: Graphical illustration of steps in toy WIMP MC. (a) Steps 1 and 2. The expected WIMP event rate for a WIMP mass of 100 GeV and WIMP-nucleon cross section of 10^{-38} cm^2 . (b) Steps 3-5. Event rate translated to observable S1 spectrum. (c) Step 6. Simulated WIMP f_{90} vs. S1 distribution. (d) Step 7. Count number of events in the WIMP search region (red).

Figure 6.19: Spin-independent WIMP-nucleon cross section 90 % C.L. exclusion plots for the DarkSide-50 UAr campaign (dashed red), DarkSide-50 AAr campaign (dotted red) and combination of the UAr and AAr campaigns (solid red). Also shown are results from LUX [49] (solid black), XENON100 [48] (dashed black), PandaX-I [112] (dotted black), CDMS [54] (solid green), PICO [113] (solid cyan), ZEPLIN-III [114] (dash dotted black) and WARP [115] (magenta).

6.7.1 Combining curves

The null results of the 70 d WIMP search with UAr can be combined with the previous null results of the WIMP search with AAr in DarkSide-50. We derive a combined exclusion limit using a similar toy MC approach. We run two simulated “experiments”—one for the UAr exposure and one for the AAr exposure. More specifically, we run steps 2-7 of the above toy MC for each of the two experiments. We use the same light yield, $\mathcal{L}_{\text{eff}}$ and f_{90} model (including noise parameters) for both simulations. The main differences between the two experiments are as follows: First, the UAr exposure is 2616 kg d, as described in Sec. 6.6, while the AAr exposure is 1738 kg d.⁶ Second, the NR acceptance curves are different: the UAr NR acceptance curve is energy dependent, while the AAr NR acceptance (apart from the f_{90} acceptance) is treated as a constant 0.82, independent of energy. Third, in step 7 of the toy MC, we draw different WIMP search boxes for each experiment, governed principally by the different background ER statistics, and we redefine the value of N_{acc} to be the sum of the number of events that land inside the respective WIMP boxes. Finally, in step 8, the simulated cross sections are again scaled by $2.3/N_{\text{acc}}$ with the newly defined N_{acc} . The 90 % C.L. exclusion curve for the combined experiments in the WIMP mass vs. cross section plane is shown in Fig. 6.19 and has a minimum of $2 \times 10^{-44} \text{ cm}^2$ at 100 GeV.

6.7.2 Projected sensitivity

Following the 70 d WIMP search described in this chapter and published in Ref. [103], the DarkSide-50 experiment has continued operating in a stable, low-background configuration. The experiment is designed for a nominal 3 yr exposure. We now describe the procedure for generating the projected sensitivity for such an exposure. The sensitivity curve is derived using a toy MC approach, similar to that described in Sec. 6.7.

The ER background will be larger in a 3 yr exposure compared to that of the 70 d

⁶The AAr exposure given here is defined as mass $\times$ time, where the fiducial mass is the same 36.9 kg and the live time is 47.1 d. This is slightly different from the quoted exposure in Ref. [76] where the exposure is defined as mass $\times$ live time $\times$ NR acceptance, and the NR acceptance is treated as independent of energy and so is folded into exposure, rather than handled separately as we do here.

exposure. We maintain the zero-background constraint on our WIMP search region, which necessitates that we redefine the WIMP acceptance region. In particular, we re-derive the ER leakage contour to account for the larger statistics by simulating the ER background for a 3 yr exposure of DarkSide-50. The number of simulated ER events is determined by simply scaling the number of events in the 70 d UAr exposure to 3 yr. The energy of each event is drawn from the S1 distribution of the 70 d UAr exposure, and the f_{90} of each event is drawn from the f_{90} distributions of the AAr exposure. Despite the evidence for non-single-scatter ER backgrounds in the UAr data, we continue to use the f_{90} model tuned to AAr data and scaled to UAr data, where now the model is scaled to the simulated 3 yr UAr data set. Given the simulated ER events in the f_{90} vs. S1 plane, we re-derive the ER leakage contour by again evaluating 0.01 leakage events per 5 PE S1 bin using the same procedure as described in Sec. 6.3 and 6.4. The leakage contour is moved upwards relative to that in Fig. 6.15, due to the larger ER background. The WIMP search region for the 3 yr exposure is then defined in the same way as described in Sec. 6.4, using the new ER leakage contour and the same NR acceptance contours as described in this chapter.

With the new WIMP search region in hand, we evaluate the projected sensitivity for the extended exposure using the same toy MC procedure as described in Sec. 6.7, the only differences being the 3 yr exposure in step 2 and re-derived WIMP box in step 7. The resulting sensitivity curve is shown in Fig. 6.20.

Figure 6.20: Projected 3 yr sensitivity for DarkSide-50 (blue). See Fig. 6.19 for descriptions of the other curves.

CHAPTER 7

Conclusion and outlook

7.1 Summary

The three detectors of DarkSide-50—the dual-phase LAr TPC, the LSV, and the WCD—have been operating in a stable low-background configuration for over a year. Data taking has been nearly continuous since April 2015, when the TPC was first filled with underground argon. The full analysis chain of DarkSide-50 has been described, starting from the acquisition of raw data and covering DarkArt, the first layer of reconstruction software; SLAD, the second layer of data reduction; and a variety of specific analyses, including a detailed study of S2 pulse shape, the reduction of ^{39}Ar in UAr, and finally, the first WIMP search performed with UAr.

The reconstruction algorithms are found to be robust and to have a large dynamic range, able to identify single photoelectrons (using the baseline finder) as well clusters of photoelectrons spread over many microseconds (using the pulse finder). The reconstruction software can efficiently identify S1 and S2 signals from the TPC and reconstruct their pulse area to within $<1\%$. The SLAD software then provides a broad set of variables for easy consumption by general users.

We then turn our attention to the electroluminescence signals of the TPC, where we develop a model to describe their time distribution. Applying the model to study events from the uppermost layer of the TPC, we find a strong radial dependence of the electron drift time across the gas pocket, consistent with the radial dependence of the S2 amplitude. The radial dependence of the S2 pulse shape is then used as an input to measurement of the longitudinal electron diffusion in liquid argon. The measurement is performed at 200 V/cm, 150 V/cm,

and 100 V/cm drift electric fields and results in the smallest uncertainties compared to other diffusion measurements in literature.

The S1 spectra from AAr and UAr data are then compared, which revealed a variety of interesting results. The stability of DarkSide-50 as a whole between the two epochs is demonstrated by the consistency of the light yields via the alignment of the gamma peaks. A custom Geant4 Monte Carlo tuned on the high statistics of AAr data is then used to fit the background contributions to the UAr S1 spectrum. An unexpected component of ^{85}Kr is found, which illustrates the precision of the MC code. The presence of ^{85}Kr is independently confirmed by measuring its 0.43 % decay branch to metastable ^{85m}Rb which gives a beta and gamma ray in delayed coincidence.

Finally, the first WIMP search using UAr is described. The development of the cuts and their efficiencies to accept nuclear recoils are covered in detail. We then discuss the analytic model used to describe the f_{90} distributions for both nuclear and electron recoils; describe the procedure for defining the WIMP search region in the f_{90} vs. S1 plane; and describe how to translate the null results of the WIMP search to an exclusion limit in the WIMP mass vs. WIMP-nucleon scattering cross section plane. The limit is the strongest to date using a LAr target, at the time of publication.

7.2 Future

The success of DarkSide-50 overall and its use of UAr in particular has spurred the development of DarkSide-20k, a next generation experiment using a dual-phase LAr TPC. If DarkSide-20k is successful, the DarkSide program will continue on with ARGO, consisting of a hundred-tonne-scale TPC. These projects will aim to probe the higher mass region (1 TeV to 10 TeV) of the WIMP mass vs. WIMP-nucleon scattering cross section parameter space, where they have the potential to reach the neutrino floor.

Guided by the successful strategies of DarkSide-50, DarkSide-20k will employ a 30 t (20 t fiducial volume) dual-phase TPC, surrounded by a newly constructed 7 m diameter spherical liquid scintillator veto and 14 m diameter, 12 m height cylindrical water tank. A

Figure 7.1: Sketch of the DarkSide-20k TPC at the center, surrounded by the liquid scintillator veto and the water Cherenkov detector.

sketch of the DarkSide-20k detectors is shown in Fig. 7.1. Importantly, instead of PMTs for light readout in the TPC, DarkSide-20k is planning to use silicon photomultipliers (SiPMs), which will enable extremely low radio-backgrounds (the largest radiogenic backgrounds in DarkSide-50 are due to the PMTs) and significantly higher light yield. Perhaps the most challenging aspect to SiPMs is the high dark count rate that will be induced by the thousands of channels necessary in a multi-tonne detector. However, a significant R&D effort between the DarkSide collaboration and SiPM producer FBK has yielded promising results, achieving dark count rates per unit area at LAr temperature comparable to those of PMTs.

DarkSide-20 aims to achieve a 100 t yr background-free exposure using underground argon from the same Doe Canyon, CO gas wells as for the UAr in DarkSide-50. The argon is to be

Figure 7.2: Projected sensitivity of DarkSide-20k and ARGO, assuming zero background; ^{39}Ar depletion factors of 10^4 and 10^5 , respectively; and exposures of 100 t yr and 1000 t yr, respectively.

further depleted of ^{39}Ar contamination by a factor of ~ 10 via the Aria project, which will use cryogenic distillation to isotopically purify the argon. This process will automatically remove the ^{85}Kr contamination discovered in the UAr in DarkSide-50.

With a 100 t yr background-free exposure and argon depleted in ^{39}Ar by a factor 10^4 relative to AAr, DarkSide-20k can achieve a WIMP-nucleon sensitivity of 10^{-47} cm^2 for a WIMP mass of 1 TeV, as shown in Fig. 7.2. The sensitivity curve is generated using a very similar toy MC approach to that described in Sec. 6.7.2. Figure 7.2 also shows the projected sensitivity curve for ARGO, a proposed 200 t fiducial volume LAr TPC that aims for a 1000 t yr background-free exposure. To compensate for the order of magnitude larger mass, the underground argon will need to be depleted by another order of magnitude, to a factor 10^5 reduction compared to atmospheric argon. The resulting projected sensitivity, assuming zero background, is then a factor of ten improved relative to DarkSide-20k, reaching a WIMP-nucleon cross section of 10^{-48} cm^2 at 1 TeV WIMP mass, touching the neutrino floor.

APPENDIX A

The ARIS experiment

A critical aspect to LAr TPC dark matter search experiments is the limited knowledge of the LAr response for nuclear recoil energies below 100 keV, in both the scintillation and ionization channels, and as function of the drift electric field. It is mandatory to improve the knowledge of the mechanisms at the basis of the S1 and S2 pulses, since the recombination effect is non-linearly energy dependent. To achieve this, we plan to measure the LAr response at different NR energies, especially in the 10 keV to 100 keV range, and with different intensities of the electric field, which also affects the electron-ion recombination probability. Also the scintillation time pulse shape is affected by the recombination effect, since the dimer produced by the ion-electron recombination de-excites in a different way with respect to an argon molecule directly excited by particle interactions.

Furthermore, columnar recombination [116] models suggest that the magnitude of the recombination effect should, in some circumstances, vary with the angle between the field and the track direction. A difference in the electron-ion recombination effect is, in fact, expected when the ionizing track is either parallel or perpendicular to the electric field. In the first case, the electrons cross a “column” of ion-electron pairs produced by the ionizing track, maximizing the recombination probability. In the second case, the crossed “column” volume is strongly reduced, and the recombination probability is minimized. The net effect of the electron-ion recombination is the reduction of the ionization signal (S2), and the consequent enhancing of the primary scintillation pulse (S1). An accurate measurement of the sum ($S1 + S2$) and the ratio ($S1/S2$) of the two signals may provide an indication on the track direction.

The ARIS (Argon Recoil Ionization and Scintillation) experiment proposes to expose a

small-scale dual-phase LAr TPC, capable of 3D position reconstruction, in front of the IPNO neutron beam [117] in Orsay, France to fully characterize the LAr response and to search for directional effects. The neutrons will be observed in coincidence by the LAr TPC and by neutron detectors (NDs), placed at several angles with respect to the beam-TPC axis to kinematically constrain the NR energy. Such measurements are nearly impossible to do in full-scale dark matter detectors such as DarkSide-50, where nuclear recoils are very rare; it is cumbersome and unwieldy to install neutron sources; and the fields and other TPC parameters cannot be readily changed.

Similar measurements were performed in 2013 and 2014 by the SCENE experiment [109, 118] with a TANDEM accelerator at Notre Dame Institute for Structure and Nuclear Astrophysics. The reaction ${}^7\text{Li}(p,n){}^7\text{Be}$ used by SCENE generates an isotropic neutron flux which strongly limited the statistics and increased the background. The exceptional collimation reachable with the LICORNE source, in association with the quasi-monoenergetic regime, makes the IPNO beam the ideal and unique facility to perform such a measurement. To make a comparison, full simulations of the proposed setup in front of the LICORNE source (assuming 13.25 MeV ${}^7\text{Li}$ energy, a neutron rate of 245 kHz, and including the full kinematics) demonstrated that we will reach a signal rate >100 higher than the SCENE one with similar or better signal to background ratios.

The dual-phase TPC that we intend to expose in front of LICORNE has been designed and built at UCLA and is currently in commissioning phase. It is equipped with seven 1-inch photomultiplier tubes (PMTs) on the top and one 3-inch PMT on the bottom, strongly increasing the energy resolution relative to SCENE and making accessible the 3D position reconstruction of each scatter. In comparison, the SCENE TPC was equipped with only one 3-inch PMT on the top, which does not allow reconstruction of the xy-position of the interaction.

Figure A.1: Previous measurements of relative scintillation efficiency ($\mathcal{L}_{\text{eff}}$) in other experiments. Red squares (“This work”) refer to Ref. [119].

A.1 Proposed Measurement

We propose to characterize several aspects of the LAr response. In particular, the physics program is aimed to reach the following goals:

1. **Precision measurement of the relative scintillation efficiency of NRs as function of the energy, in the range 16 keV to 130 keV.** As shown in Fig. A.1, results present in literature are controversial especially below 50 keV [119].
2. **Map of the PSD estimator (f_{90}) as function of the energy and the drift field.** The f_{90} variable is defined as the number of detected primary scintillation photons in the first 90 ns with respect to the total number of detected photons in S1. This parameter is used to discriminate the ER background in WIMP search experiments. In particular, the precise measurement of its distribution is critical to define the acceptance region of the expected WIMP candidates.
3. **Measurement of recombination probability as a function of the energy and drift field.** We will derive its behavior by combining S1 and S2 scintillation signals.

Figure A.2: Schematic of the LAr TPC in the neutron beam with neutron detectors (green) at various angles.

The result will be compared with g_4ds , working for electron recoils but never tested on a NR data set.

4. **Measurement of recoil directionality.** We will measure the scintillation and ionization yields, at fixed NR energies, for perpendicular and (almost) parallel NR directions with respect to the drift field. As already mentioned, we expect an enhancement of the recombination effect for parallel NR tracks with respect to the orthogonal ones, which should result in a difference of the yields. This measurement has the potential to deeply impact the future of direct dark matter experiments, extending their sensitivities to extremely small WIMP elastic scattering cross sections, where also astrophysical neutrinos are expected.

A.2 Experimental Setup

The neutrons produced by the IPNO Tandem ^7Li pulsed beam on the LICORNE source first interact with the liquid argon target in the TPC and then, in coincidence, with a neutron detector, as shown in Fig. A.2. There are seven neutron detectors currently available at the IPNO facility.

The LAr TPC has a design that closely follows the one used in DarkSide-50. The diam-

eter and height of the LAr target are chosen to limit contamination from multiple neutron scattering to an acceptable level. The active volume is contained within a 7.6 cm diameter, 7.6 cm tall, right circular Poly-Tetra-Fluoro-Ethylene (PTFE, or Teflon) cylinder and capped by quartz windows. The LAr is viewed through the windows by one 3-inch Hamamatsu R11065 PMT on the bottom and an array of seven 1-inch Hamamatsu R8520 PMTs on the top. The windows are coated with indium tin oxide (ITO), a transparent conductor, which constitute the anode and cathode surfaces. Copper rings embedded in the PTFE cylinder maintain field uniformity. All internal surfaces of the detector are coated with the wavelength shifter Tetra-Phenyl-Butadiene (TPB), which converts the LAr scintillation light from the vacuum UV range (128 nm) into the blue range (420 nm). A stainless steel mesh grid is positioned 8 mm below the anode and defines the top of the active volume. The gas pocket is designed to be 5 mm thick, so the grid is below the liquid surface. The grid, anode, and cathode surfaces can be tuned to set the drift field up to 1 kV/cm and the extraction field above the grid up to 4 kV/cm. The TPC is mounted on a stainless steel flange which sits on top of a vacuum jacketed stainless steel cryostat. Figure A.3 shows the TPC and its supporting infrastructure at UCLA.

Operation of the TPC requires a variety of supporting subsystems for high voltage, Ar liquefaction, gas recirculation, online Ar purification, data acquisition, and system monitoring. High voltage (HV) is passed into the cryostat via custom made feedthroughs. The maximum HV needed for this experiment is 12 kV, which is passed into the cryostat through a feedthrough that has been tested to 20 kV. A Cryomech PT90 Cryocooler condenses Ar into liquid, which is transferred to the cryostat via a vacuum jacketed transfer line. Argon gas is recirculated using a Q-drive pump in a closed loop from the cryostat through a SAES MonoTorr getter (to remove impurities) and back to the condenser. The cryostat pressure, gas flow rate, and various temperatures are monitored by a LabView slow control. The PMT signals are read by a 12-bit 250 MHz CAEN V1720 digitizer, the same model used in DarkSide-50. The DAQ software is a stripped down version of the one used in DarkSide, and the reconstruction software, like DarkArt, is based in Fermilab's *art* framework and shares much of the same design architecture.

(a)

(b)

Figure A.3: (a) Fully assembled LAr TPC at UCLA. (b) TPC installed within cryostat and integrated into full cryogenics, high voltage, and data acquisition systems.

The TPC has successfully been filled with research grade ultra high purity (atmospheric) argon, and data has been taken with ^{133}Ba and ^{57}Co sources placed near the cryostat. Commissioning of the TPC is ongoing in preparation for an exposure to the neutron beam at the IPNO facility in October 2016.

BIBLIOGRAPHY

- [1] F. Zwicky, “The redshift of extragalactic nebulae”, [Helv. Phys. Acta](#) **6**, 110 (1933).
- [2] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae”, [Astrophys. J.](#) **86**, 217 (1937).
- [3] V. Rubin, “Dark matter in spiral galaxies”, [Sci. Am.](#) **248**, 96 (1983).
- [4] T. S. van Albada, J. N. Bahcall, K. Begeman, and R. Sancisi, “Distribution of dark matter in the spiral galaxy NGC 3198”, [Astrophys. J.](#) **295**, 305 (1985).
- [5] NASA, ESA, and J. Richard, *Abell 2218*, (2008) <https://www.spacetelescope.org/images/heic0814a/>.
- [6] G. Bertone, ed., *Particle Dark Matter: Observations, Models, and Searches* (Cambridge University Press, 2010).
- [7] G. Bertone, D. Hooper, and J. Silk, “Particle dark matter: evidence, candidates and constraints”, [Phys. Rep.](#) **405**, 279 (2005).
- [8] D. Clowe et al., “A direct empirical proof of the existence of dark matter”, [Astrophys. J.](#) **648**, L109 (2006).
- [9] M. S. Turner and J. A. Tyson, “Cosmology at the millennium”, [Rev. Mod. Phys.](#) **71**, S145 (1999).
- [10] C. L. Bennett et al., “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results”, [Astrophys. J. Suppl. Ser.](#) **208**, 20 (2013).
- [11] D. J. Fixsen, “The Temperature of the Cosmic Microwave Background”, [Astrophys. J.](#) **707**, 916 (2009).
- [12] G. F. Smoot et al., “Structure in the COBE Differential Microwave Radiometer First-year Maps”, [Astrophys. J.](#) **396**, L1 (1992).
- [13] P. A. R. Ade et al., “Planck 2013 results I. Overview of products and scientific results”, [Astron. Astrophys.](#) **571**, 48 (2014).

- [14] A. D. Miller et al., “A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from $l = 100$ to 400”, [Astrophys. J. **524**, L1 \(1999\)](#).
- [15] K. Garrett and G. Duda, “Dark Matter: A Primer”, [Adv. Astron. **2011**, 968283 \(2011\)](#).
- [16] P. A. R. Ade et al., “Planck 2013 results. XVI. Cosmological parameters”, [Astron. Astrophys. **571**, A16 \(2014\)](#).
- [17] D. Boyanovsky, H. J. de Vega, and D. J. Schwarz, “Phase Transitions in the Early and Present Universe”, [Annu. Rev. Nucl. Part. Sci. **56**, 441 \(2006\)](#).
- [18] G. Steigman, “Primordial Nucleosynthesis in the Precision Cosmology Era”, [Annu. Rev. Nucl. Part. Sci. **57**, 463 \(2007\)](#).
- [19] S. Burles, K. M. Nollett, and M. S. Turner, “Big-Bang Nucleosynthesis: Linking Inner Space and Outer Space”, [arXiv:astro-ph/9903300 \(1999\)](#).
- [20] M. Taoso, G. Bertone, and A. Masiero, “Dark matter candidates: a ten-point test”, [JCAP **2008**, 022 \(2008\)](#).
- [21] B. Carr, “Baryonic Dark Matter”, [Annu. Rev. Astron. Astrophys. **32**, 531 \(1994\)](#).
- [22] C. Alcock et al., “Experimental Limits on the Dark Matter Halo of the Galaxy from Gravitational Microlensing”, [Phys. Rev. Lett. **74**, 2867 \(1995\)](#).
- [23] P. H. Frampton, M. Kawasaki, F. Takahashi, and T. T. Yanagida, “Primordial black holes as all dark matter”, [JCAP **2010**, 023 \(2010\)](#).
- [24] S. W. Hawking, “Particle creation by black holes”, [Commun. Math. Phys. **43**, 199 \(1975\)](#).
- [25] B. J. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, “New cosmological constraints on primordial black holes”, [Phys. Rev. D **81**, 104019 \(2010\)](#).
- [26] B. P. Abbott et al., “Observation of gravitational waves from a binary black hole merger”, [Phys. Rev. Lett. **116**, 061102 \(2016\)](#).
- [27] S. Bird et al., “Did LIGO detect dark matter?”, [Phys. Rev. Lett. **116**, 201301 \(2016\)](#).
- [28] Y. Fukuda et al., “Evidence for Oscillation of Atmospheric Neutrinos”, [Phys. Rev. Lett. **81**, 1562 \(1998\)](#).

- [29] Q. R. Ahmad et al., “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”, [Phys. Rev. Lett. **89**, 011301 \(2002\)](#).
- [30] Y. B. Zeldovich, J. Einasto, and S. F. Shandarin, “Giant voids in the Universe”, [Nature **300**, 407 \(1982\)](#).
- [31] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, “The Role of Sterile Neutrinos in Cosmology and Astrophysics”, [Annu. Rev. Nucl. Part. Sci. **59**, 191 \(2009\)](#).
- [32] R. D. Peccei and H. R. Quinn, “CP Conservation in the Presence of Pseudoparticles”, [Phys. Rev. Lett. **38**, 1440 \(1977\)](#).
- [33] S. J. Asztalos et al., “Searches for Astrophysical and Cosmological Axions”, [Annu. Rev. Nucl. Part. Sci. **56**, 293 \(2006\)](#).
- [34] C. A. Baker et al., “Improved Experimental Limit on the Electric Dipole Moment of the Neutron”, [Phys. Rev. Lett. **97**, 131801 \(2006\)](#).
- [35] C. Vafa and E. Witten, “Parity Conservation in Quantum Chromodynamics”, [Phys. Rev. Lett. **53**, 535 \(1984\)](#).
- [36] J. Preskill, M. B. Wise, and F. Wilczek, “Cosmology of the invisible axion”, [Phys. Lett. B **120**, 127 \(1983\)](#).
- [37] H. Peng et al., “Cryogenic cavity detector for a large-scale cold dark-matter axion search”, [Nucl. Instrum. Methods A **444**, 569 \(2000\)](#).
- [38] K. Zioutas et al., “A decommissioned LHC model magnet as an axion telescope”, [Nucl. Instrum. Methods A **425**, 480 \(1999\)](#).
- [39] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter”, [Phys. Rep. **267**, 195 \(1996\)](#).
- [40] M. Schelke, “Supersymmetric Dark Matter: aspects of sfermion coannihilations”, PhD thesis (Stockholm University, 2004).
- [41] J. L. Feng, “Dark Matter Candidates from Particle Physics and Methods of Detection”, [Annu. Rev. Astron. Astrophys. **48**, 495 \(2010\)](#).

- [42] S. Dodelson, *Modern Cosmology* (Academic Press, San Francisco, 2003).
- [43] B. M. Loer, “Towards a Depleted Argon Time Projection Chamber WIMP Search: DarkSide Prototype Analysis and Predicted Sensitivity”, PhD thesis (Princeton University, 2011).
- [44] J. Bovy and S. Tremaine, “On the Local Dark Matter Density”, [Astrophys. J. **756**, 89 \(2012\)](#).
- [45] J. D. Lewin and P. F. Smith, “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil”, [Astropart. Phys. **6**, 87 \(1996\)](#).
- [46] C. Savage, K. Freese, and P. Gondolo, “Annual modulation of dark matter in the presence of streams”, [Phys. Rev. D **74**, 043531 \(2006\)](#).
- [47] P. Cushman et al., “Snowmass CF1 Summary: WIMP Dark Matter Direct Detection”, [arXiv:1310.8327 \(2013\)](#).
- [48] E. Aprile et al., “Dark Matter Results from 225 Live Days of XENON100 Data”, [Phys. Rev. Lett. **109**, 181301 \(2012\)](#).
- [49] D. S. Akerib et al., “Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data”, [Phys. Rev. Lett. **116**, 161301 \(2016\)](#).
- [50] M. G. Boulay, “DEAP-3600 Dark Matter Search at SNOLAB”, [J. Phys. Conf. Ser. **375**, 012027 \(2012\)](#).
- [51] K. Abe et al., “XMASS detector”, [Nucl. Instrum. Methods A **716**, 78 \(2013\)](#).
- [52] R. Bernabei et al., “First results from DAMA/LIBRA and the combined results with DAMA/NaI”, [Eur. Phys. J. C **56**, 333 \(2008\)](#).
- [53] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, “Compatibility of DAMA/LIBRA dark matter detection with other searches”, [JCAP **2009**, 036 \(2009\)](#).
- [54] R. Agnese et al., “Improved WIMP-search reach of the CDMS II germanium data”, [Phys. Rev. D **92**, 072003 \(2015\)](#).

- [55] J. Billard, E. Figueroa-Feliciano, and L. Strigari, “Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments”, [Phys. Rev. D **89**, 023524 \(2014\)](#).
- [56] *SuperCDMS (experiment homepage)*, <http://cdms.berkeley.edu/>.
- [57] *XENON1T (experiment homepage)*, <http://www.xenon1t.org/>.
- [58] *LZ (experiment homepage)*, <http://lz.lbl.gov/>.
- [59] D. S. Akerib et al., “Radiogenic and muon-induced backgrounds in the LUX dark matter detector”, [Astropart. Phys. **62**, 33 \(2015\)](#).
- [60] H. O. Back et al., “First Commissioning of a Cryogenic Distillation Column for Low Radioactivity Underground Argon”, [arXiv:1204.6061 \(2012\)](#).
- [61] D. S. Akerib et al., “Chromatographic separation of radioactive noble gases from xenon”, [arXiv:1605.03844 \(2016\)](#).
- [62] G. Bellini et al., “Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth”, [JCAP **05**, 015 \(2012\)](#).
- [63] G. Bellini et al., “Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth This”, [JCAP **08**, 049 \(2013\)](#).
- [64] A. Wright, P. Mosteiro, B. Loer, and F. Calaprice, “A highly efficient neutron veto for dark matter experiments”, [Nucl. Instrum. Methods A **644**, 18 \(2011\)](#).
- [65] J. Lindhard, V. Nielsen, M. Scharff, and P. V. Thomsen, “Integral equations governing radiation effects”, *Mat. Fys. Medd. Dan. Vid. Selsk.* **33**, 1 (1963).
- [66] D. M. Mei, Z. B. Yin, L. C. Stonehill, and A. Hime, “A model of nuclear recoil scintillation efficiency in noble liquids”, [Astropart. Phys. **30**, 12 \(2008\)](#).
- [67] D. Gastler et al., “Measurement of scintillation efficiency for nuclear recoils in liquid argon”, [Phys. Rev. C **85**, 065811 \(2012\)](#).
- [68] T. Doke, K. Masuda, and E. Shibamura, “Estimation of absolute photon yields in liquid argon and xenon for relativistic (1 MeV) electrons”, [Nucl. Instrum. Methods A **291**, 617 \(1990\)](#).

- [69] R. S. Mulliken, “Potential curves of diatomic rare gas molecules and their ions with particular reference to Xe_2 ”, [J. Chem. Phys. **52**, 5170 \(1970\)](#).
- [70] M. Martin, “Exciton Self-Trapping in Rare-Gas Crystals”, [J. Chem. Phys. **54**, 3289 \(1971\)](#).
- [71] O. Cheshnovsky, R. Baruch, and J. Jortner, “Temperature dependence of rare gas molecular emission in the vacuum ultraviolet”, [Chem. Phys. Lett. **15**, 475 \(1972\)](#).
- [72] J. Jortner, L. Meyer, S. A. Rice, and E. G. Wilson, “Localized Excitations in Condensed Ne, Ar, Kr, and Xe”, [J. Chem. Phys. **42**, 4250 \(1965\)](#).
- [73] A. Hitachi et al., “Effect of ionization density on the time dependence of luminescence from liquid argon and xenon”, [Phys. Rev. B **27**, 5279 \(1983\)](#).
- [74] B. J. P. Jones, “Sterile Neutrinos in Cold Climates”, PhD thesis (MIT, 2015).
- [75] S. Amoruso et al., “Study of electron recombination in liquid argon with the ICARUS TPC”, [Nucl. Instrum. Methods A **523**, 275 \(2004\)](#).
- [76] P. Agnes et al., “First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso”, [Phys. Lett. B **743**, 456 \(2015\)](#).
- [77] V. N. Lebedenko et al., “Results from the first science run of the ZEPLIN-III dark matter search experiment”, [Phys. Rev. D **80**, 052010 \(2009\)](#).
- [78] D. S. Akerib et al., “First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility”, [Phys. Rev. Lett. **112**, 91303 \(2014\)](#).
- [79] J. H. J. Poole and C. F. G. Delaney, “Origin of Atmospheric Argon and the Radioactive Decay Constants of Potassium-40”, [Nature **167**, 680 \(1951\)](#).
- [80] P. Benetti et al., “Measurement of the specific activity of ^{39}Ar in natural argon”, [Nucl. Instrum. Methods A **574**, 83 \(2007\)](#).
- [81] B. E. Lehmann, S. N. Davis, and J. T. Fabryka-Martin, “Atmospheric and subsurface sources of stable and radioactive nuclides used for groundwater dating”, [Water Resour. Res. **29**, 2027 \(1993\)](#).
- [82] H. H. Loosli, “A dating method with ^{39}Ar ”, [Earth Planet. Sci. Lett. **63**, 51 \(1983\)](#).

- [83] D. Acosta-Kane et al., “Discovery of underground argon with low level of radioactive ^{39}Ar and possible applications to WIMP dark matter detectors”, [Nucl. Instrum. Methods A **587**, 46 \(2008\)](#).
- [84] H. O. Back et al., “First Large Scale Production of Low Radioactivity Argon From Underground Sources”, [arXiv:1204.6024 \(2012\)](#).
- [85] T. Alexander et al., “Light yield in DarkSide-10: A prototype two-phase argon TPC for dark matter searches”, [Astropart. Phys. **49**, 44 \(2013\)](#).
- [86] R. Acciarri et al., “Liquid Argon Dielectric Breakdown Studies with the MicroBooNE Purification System”, [JINST **9**, P11001 \(2014\)](#).
- [87] P. Agnes et al., “The veto system of the DarkSide-50 experiment”, [JINST **11**, P03016 \(2016\)](#).
- [88] G. Alimonti et al., “The Borexino detector at the Laboratori Nazionali del Gran Sasso”, [Nucl. Instrum. Methods A **600**, 568 \(2009\)](#).
- [89] F. P. An et al., “The muon system of the Daya Bay Reactor antineutrino experiment”, [Nucl. Instrum. Methods A **773**, 8 \(2015\)](#).
- [90] C. Green et al., “The art framework”, [J. Phys. Conf. Ser. **396**, 022020 \(2012\)](#).
- [91] DarkSide Collaboration, *DarkArt source code*, (2012) <https://cdcvns.fnal.gov/redmine/projects/darkart/repository>.
- [92] V. N. Solovov et al., “Position Reconstruction in a Dual Phase Xenon Scintillation Detector”, [IEEE Trans. Nucl. Sci. **59**, 3286 \(2012\)](#).
- [93] R. Dossi, A. Ianni, G. Ranucci, and O. J. Smirnov, “Methods for precise photoelectron counting with photomultipliers”, [Nucl. Instrum. Methods A **451**, 623 \(2000\)](#).
- [94] C. Amsler et al., “Luminescence quenching of the triplet excimer state by air traces in gaseous argon”, [JINST **3**, P02001 \(2008\)](#).
- [95] S. Baker and R. D. Cousins, “Clarification of the use of Chi-square and likelihood functions in fits to histograms”, [Nucl. Instrum. Methods **221**, 437 \(1984\)](#).

- [96] L. G. H. Huxley and R. W. Crompton, *The Diffusion and Drift of Electrons in Gases* (Wiley, New York, 1974).
- [97] P. Cennini et al., “Performance of a three-ton liquid argon time projection chamber”, [Nucl. Instrum. Methods A **345**, 230 \(1994\)](#).
- [98] *Erfc*, <http://functions.wolfram.com/GammaBetaErf/Erfc/06/02/>.
- [99] Y. Li et al., “Measurement of longitudinal electron diffusion in liquid argon”, [Nucl. Instrum. Methods A **816**, 160 \(2016\)](#).
- [100] V. M. Atrazhev and I. V. Timoshkin, “Transport of Electrons in Atomic Liquids in High Electric Fields”, [IEEE Trans. Dielectr. Electr. Insul. **5**, 450 \(1998\)](#).
- [101] J. Xu et al., “A study of the trace ^{39}Ar content in argon from deep underground sources”, [Astropart. Phys. **66**, 53 \(2015\)](#).
- [102] B. Lenardo et al., “A Global Analysis of Light and Charge Yields in Liquid Xenon”, [IEEE Trans. Nucl. Sci. **62**, 3387 \(2015\)](#).
- [103] P. Agnes et al., “Results from the first use of low radioactivity argon in a dark matter search”, [Phys. Rev. D **93**, 081101\(R\) \(2016\)](#).
- [104] Hamamatsu, *Photomultiplier Tubes: Basics and Applications* (Hamamatsu Photonics K.K., 2006).
- [105] C. J. Martoff and P. D. Lewin, “COSMO - a program to estimate spallation radioactivity produced in a pure substance by exposure to cosmic radiation on the earth”, [Comput. Phys. Commun. **72**, 96 \(1992\)](#).
- [106] W. Lippincott et al., “Scintillation time dependence and pulse shape discrimination in liquid argon”, [Phys. Rev. C **78**, 035801 \(2008\)](#).
- [107] D. V. Hinkley, “On the ratio of two correlated normal random variables”, [Biometrika **56**, 635 \(1969\)](#).
- [108] M. G. Boulay et al., “Measurement of the scintillation time spectra and pulse-shape discrimination of low-energy beta and nuclear recoils in liquid argon with DEAP-1”, [arXiv:0904.2930 \(2009\)](#).

- [109] H. Cao et al., “Measurement of scintillation and ionization yield and scintillation pulse shape from nuclear recoils in liquid argon”, [Phys. Rev. D **91**, 092007 \(2015\)](#).
- [110] V. M. Gehman et al., “Fluorescence efficiency and visible re-emission spectrum of tetraphenyl butadiene films at extreme ultraviolet wavelengths”, [Nucl. Instrum. Methods A **654**, 116 \(2011\)](#).
- [111] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals”, [Phys. Rev. D **57**, 3873 \(1998\)](#).
- [112] X. Xiao et al., “Low-mass dark matter search results from full exposure of the PandaX-I experiment”, [Phys. Rev. D **92**, 052004 \(2015\)](#).
- [113] C. Amole et al., “Dark Matter Search Results from the PICO-2L C₃F₈ Bubble Chamber”, [Phys. Rev. Lett. **114**, 231302 \(2015\)](#).
- [114] D. Y. Akimov et al., “WIMP-nucleon cross-section results from the second science run of ZEPLIN-III”, [Phys. Lett. B **709**, 14 \(2012\)](#).
- [115] P. Benetti et al., “First results from a dark matter search with liquid argon at 87 K in the Gran Sasso underground laboratory”, [Astropart. Phys. **28**, 495 \(2008\)](#).
- [116] D. R. Nygren, “Columnar recombination: a tool for nuclear recoil directional sensitivity in a xenon-based direct detection WIMP search”, [J. Phys. Conf. Ser. **460**, 012006 \(2013\)](#).
- [117] J. N. Wilson et al., “The LICORNE neutron source and measurements of prompt γ -rays”, [Phys. Procedia **59**, 31 \(2014\)](#).
- [118] T. Alexander et al., “Observation of the dependence on drift field of scintillation from nuclear recoils in liquid argon”, [Phys. Rev. D **88**, 092006 \(2013\)](#).
- [119] W. Creus et al., “Scintillation efficiency of liquid argon in low energy neutron-argon scattering”, [JINST **10**, P08002 \(2015\)](#).