

xy POSITION RECONSTRUCTION IN DARKSIDE-50

JASON PHILIP BRODSKY

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS
ADVISER: PETER D. MEYERS

SEPTEMBER 2015

© Copyright by Jason Philip Brodsky, 2015.

All rights reserved.

Abstract

The DarkSide-50 experiment seeks to directly detect dark matter in a liquid argon time projection chamber. In this dissertation, I present an algorithm of my design that determines the position of particle interactions with the liquid argon. This position reconstruction algorithm will be used by DarkSide-50 to reject backgrounds, particularly backgrounds from radioactive elements on the detector surface.

The position reconstruction algorithm functions by constructing light response functions (LRFs) that map locations in the detector to the expected distribution of signal in DarkSide-50's 38 photomultiplier tubes. Accurate LRFs cannot be produced by simulations of DarkSide-50's optics because such simulations are known to be flawed. Instead, this algorithm constructs LRFs using an iterative process driven by data. Initial, flawed LRFs are produced using simulated events but then used to produce new LRFs from data events. Multiple generations of LRFs are created from data with each generation driven to better satisfy a known feature of the detector: the dominant argon-39 background is uniformly distributed.

I also discuss a method of discriminating against surface background as an alternative to the common approach of fiducialization. This method considers the difference in goodness-of-fit between the best-fit reconstructed position and the best-fit position at the detector's surface.

I conclude by presenting results on the performance and validity of this algorithm, including some discussion of reconstruction errors.

Acknowledgements

Many, many thanks to Prof. Peter Meyers, my adviser. Peter's attention and aid to my education has been extraordinary. Throughout my time at Princeton, I could count on being able to walk into Peter's office and receive patient counsel, enlightening insights, and enthusiasm for my progress. Peter's technical contributions to my work are too many to count, because he always took the time to thoroughly understand the methods and challenges of each project. His contributions to my development as a student are even greater. Peter pointed me at the best opportunities to learn by doing. He emphasized that my education was important, not just my work, and in doing so gave me the drive to succeed. Peter has been an excellent manager, clear in his standards and considerate of my personal needs. Finally, Peter has been a true role model. I am very fortunate to have studied under someone who sets an excellent example in his engagement with the experiment and his rigorous pursuit of the truth.

One exceptional feature of the Princeton DarkSide group is our three professors, each of whom brings great strengths to the team. Many thanks go out to Profs. Frank Calaprice and Cristiano Galbiati, whose advice and leadership has made DarkSide possible and my work on it productive. Many thanks also to the students and postdocs I have had the pleasure of working with and learning from. They have been an incredible group to work with, both productive and supportive. The DarkSide collaboration is tight-knit, and I have been lucky enough to meet almost everyone working on the experiment. Many thanks to everyone who made DarkSide-50 possible.

Particular thanks to Masayuki Wada and Andrew Watson, who have worked on and contributed greatly to the problem of position reconstruction in DarkSide-50. My own work owes much to having them working in parallel, developing techniques that are very important to my final methods. Thanks also to Chris Stanford, who provided the coincident decay study my results depend on.

Thanks to Vladimir Solovov, one of the authors of the work that inspired mine. My personal communications with him were very helpful in understanding his work and beginning mine.

None of this work would have been possible without the loving support of my family, which began long before I came to Princeton and carried me through all the challenges of graduate school. I am a scientist because my family encouraged that dream. In my wife, Weiyi, I have found someone who shares and supports my ambitions with love, patience, respect, and hard work. She has helped make me the person I am.

DarkSide is supported by National Science Foundation (US, Grants PHY-0919363, PHY-1004072, PHY-1211308, and associated collaborative Grants), DOE (US, Contract Nos. DE-FG02-91ER40671 and DE-AC02-07CH11359), the Istituto Nazionale di Fisica Nucleare (Italy), and the NCN (Poland, Grant UMO-2012/05/E/ST2/02333). This work was supported in part by the Kavli Institute for Cosmological Physics at the University of Chicago through grant National Science Foundation PHY-1125897 and an endowment from the Kavli Foundation and its founder Fred Kavli.

To my mother, Dr. Valerie Cooke, the first scientist in our family,
and to my wife, Weiyi Guo.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction To WIMP Dark Matter	1
1.1 Dark Matter	1
1.2 Detection of WIMPs	3
1.3 Nuclear Recoils Produced by WIMP Interactions	5
1.4 State of the Field of Direct Detection	7
2 The DarkSide-50 Detector	10
2.1 Detector Overview	10
2.2 The Liquid Argon Time Projection Chamber	10
2.2.1 Liquid Argon Scintillation	11
2.2.2 Liquid Argon Ionization and the Time Projection Chamber Design	14
2.2.3 Containing the Argon	15
2.2.4 Photomultiplier Tubes in the Liquid Argon TPC	17
2.2.5 Wavelength Shifter	19
2.2.6 Support Infrastructure of the Liquid Argon TPC	21
2.2.7 Argon-39 and the Underground Argon Program	21
2.3 The Veto Detectors	22
2.4 DarkSide-50 at LNGS	23

2.5	Data Acquisition	24
2.5.1	TPC Data Acquisition	24
2.5.2	Veto Data Acquisition	24
2.6	Fundamental Analysis	25
2.6.1	Baseline Finding	25
2.6.2	Single Photoelectron Calibration	26
2.6.3	Pulse Finding	29
2.7	Events in DarkSide-50	32
2.7.1	WIMP Events	32
2.7.2	Argon-39	33
2.7.3	Surface Backgrounds	33
2.8	Information Available for xy Reconstruction	35
2.9	Simulation of DarkSide-50	36
2.10	Current Status of the DarkSide-50 Detectors	38
3	xy Reconstruction Methodology	41
3.1	Weighted Least Squares Event Reconstruction	41
3.1.1	The χ^2 Distribution	42
3.2	Light Response Functions	44
3.2.1	The Combined Bottom LRF	46
3.2.2	Far Smoothing	47
3.2.3	Determining LRFs from Events with Known Location	48
3.2.4	The Chicken-and-Egg Problem and Its Solution	49
3.3	Prior Work	50
3.3.1	Differences between past work and this work	51
3.3.2	Other Prior Work	52
3.4	Iteration, Convergence and Uniformity	52
3.4.1	Iteration Runaway	52

3.4.2	Nonuniformity	54
3.4.3	Convergence Pressure	56
3.4.4	Remnant Iteration Runaway	57
3.4.5	The Argon-39 Uniformity Assumption	58
3.5	Monte Carlo Adjustment	59
3.6	Implementation	63
3.6.1	LRF Implementation and Scale	63
3.6.2	Nominal Detector Radius	64
3.6.3	Location Evaluation	65
3.6.4	LRF Construction from Monte Carlo	66
3.6.5	LRF Construction from Data	67
3.6.6	Uniformity Pressure	68
3.6.7	API	69
3.6.8	Portability	70
3.7	The Surface Background Cut	71
3.7.1	Surface Backgrounds	71
3.7.2	Fiducial Cuts	71
3.7.3	An Alternative to Fiducialization	72
3.7.4	Surface Cut Implementation	72
4	<i>xy</i> Reconstruction Results	75
4.1	Event Sets Used	75
4.1.1	Simulated High-S2 Events	75
4.1.2	Simulated Low-S2 Events	75
4.1.3	Data Events	76
4.2	Characteristics of Data and Simulated Events	78
4.2.1	Effects of Binning on Uncertainty Estimation	78
4.2.2	Indications of Detector Asymmetry in Data	81

4.3	<i>xy</i> Reconstruction Performance Results	81
4.3.1	Tests on Simulation	82
4.3.2	Uniformity	85
4.4	<i>xy</i> Reconstruction Validation Results	92
4.4.1	Coincident Chain Decays	92
4.4.2	Recovery from Flawed Simulation	97
4.5	Surface Background Cut Results	105
4.5.1	Acceptance vs Rejection in Simulated Events	105
4.5.2	Acceptance vs Rejection in Data Events	109
5	Conclusion	113

List of Figures

1.1	WIMP Interactions.	4
1.2	WIMP nuclear recoil interaction spectrum.	6
1.3	Exclusion limits set by direct detection experiments.	8
2.1	The liquid argon TPC in cross section.	12
2.2	The different pulse shapes of electron and nuclear recoils.	13
2.3	The field rings maintain a voltage gradient, creating an electric field across the liquid argon and a stronger field across the gas-liquid interface.	16
2.4	The PMT packing.	18
2.5	The diving bell and bottom window with TPB applied.	20
2.6	Laser spectrum in one channel.	27
2.7	Illustration of pulse finding algorithm.	31
2.8	Results of the pulse finding algorithm.	31
2.9	Sketch of drift electrons producing S2 light collected by PMTs.	37
2.10	The DarkSide-50 exclusion limit along with limits set by other experiments.	39
3.1	The observed distribution of χ^2	43
3.2	The light response function for PMT 30.	45
3.3	The light response function for PMTs 3 and 17, with the PMT locations illustrated.	46

3.4	Increase in reconstruction error over multiple iterations, without uniformity pressure, Monte Carlo.	54
3.5	Event density by position in iterations 1 (left) and 60 (right) Monte Carlo. . .	55
3.6	Event density by position in iteration 1 in data.	55
3.7	Increase in reconstruction error over multiple iterations, with and without uniformity pressure.	58
3.8	Reconstructed data positions vs. the drift time.	60
3.9	Discrepancy between Monte Carlo and data, as measured by fraction of light in channel 32.	61
3.10	Discrepancy between data and Monte Carlo after adjustment, as measured by fraction of light in channel 32.	62
3.11	The search boxes for events with the most light in PMTs 29 and 35, Monte Carlo.	66
4.1	The standard deviation of measured light in each bin and the and the estimated statistical uncertainty of that measurement, plotted along the x -axis . . .	79
4.2	Asymmetric data.	82
4.3	Average reconstruction error in each bin, high-S2 Monte Carlo, reconstructed from true LRFs.	83
4.4	Average reconstruction error by radius, high-S2 Monte Carlo, reconstructed from true LRFs.	84
4.5	Average reconstruction error in each bin, low-S2 Monte Carlo, reconstructed from true LRFs.	84
4.6	Average reconstruction error by radius, low-S2 Monte Carlo, reconstructed from true LRFs.	85
4.7	Event density by position, 0th iteration LRFs, with uniformity pressure. . . .	87
4.8	Event density by position, 4th iteration LRFs, with uniformity pressure. . . .	88
4.9	Event density by position, 29th iteration LRFs, with uniformity pressure. . .	89

4.10	Event density by position, 30th iteration LRFs, without uniformity pressure.	90
4.11	Fraction of Events by Radius, 30th iteration LRFs, without uniformity pressure.	91
4.12	Number of Bins at Each Event Density, 30th iteration LRFs, without uniformity pressure.	91
4.13	Event density by position, 100th iteration LRFs, no uniformity pressure, Monte Carlo.	93
4.14	A BiPo waveform.	95
4.15	Reconstructed location of ^{214}Bi and ^{214}Po decays.	98
4.16	Reconstructed distance between ^{214}Bi and ^{214}Po decays.	99
4.17	Regions where the “flawed” simulation has the visible light emission probability of TPB reduced.	100
4.18	Reconstruction error as a function of position, using 0th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.	102
4.19	Reconstruction error as a function of position, using 29th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.	103
4.20	Reconstruction error in the north flaw in each iteration using flawed Monte Carlo events.	104
4.21	Reconstruction error in the south flaw in each iteration using flawed Monte Carlo events	105
4.22	Event density by position, using 29th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.	106
4.23	Acceptance at cut thresholds, high-statistics Monte Carlo.	107
4.24	Background rejection at cut thresholds, high-statistics Monte Carlo.	108
4.25	Acceptance at cut thresholds, low-statistics Monte Carlo.	109
4.26	Background rejection at cut thresholds, low-statistics Monte Carlo.	110
4.27	Acceptance at cut thresholds, data.	111
4.28	Estimation of background rejection using pseudoSCM on coincident decays. .	112

Chapter 1

Introduction To WIMP Dark Matter

1.1 Dark Matter

Most of the matter in our universe is not visible. This conclusion is supported by a growing body of evidence developed over the last century. From our present-day galaxy to the beginnings of the universe, we observe the gravitational effects of mass far in excess of the amount visible via electromagnetic processes. Measurements in different contexts confirm that 84.5% of all matter in the universe must be this dark matter, interacting gravitationally and probably weakly, but not via the other fundamental forces.

One of the striking features of the evidence for dark matter is the variety of contexts in which that evidence is observed. The “classical” evidence for dark matter is the rotation of galaxies in clusters and stars within galaxies, both of which require greater gravitational forces than can be generated by luminous matter [1][2]. Since those discoveries, many other techniques have found evidence for dark matter at the galactic cluster scale and in the early universe [3]. The combined results of these studies requires the mass-energy content of the universe to be 73% dark energy, 23% dark matter, and 4.6% baryonic matter.

The collected evidence for dark matter puts some constraints on what it might be. For the dark matter to exist in both the early and present universe it must be stable¹. Early universe models that match the cosmic microwave background require dark matter to not be “hot” that is, relativistic after decoupling from ordinary matter in the early universe [4]. The same models constrain the baryonic matter density to be smaller than the dark matter density, so dark matter cannot be baryonic. Combined with the fact that dark matter must be gravitationally interacting, these constraints rule out all standard model particles as candidates for dark matter.

The constraints on dark matter still allow a number of theoretical candidates of undiscovered particles that could be dark matter [3]. Many of these models allow the dark matter candidate some coupling with standard model particles, allowing for the possibility of observing dark matter interactions besides gravity. Of these candidates, two lend themselves to direct detection by earth-based experiments: axions and WIMPs. Experiments to detect both candidates are underway.

Axions arise from the solution by [5] to the strong CP problem². To meet the constraints on dark matter, axions must have low mass, $< \sim 0.01$ eV, and have very weak coupling to ordinary matter [6]. A number of ongoing experiments are currently searching for axions, primarily through their coupling to photons, such as ADMX [7].

WIMP stands for weakly-interacting massive particle. Many extensions of the standard model posit such particles, including highly-favored possibilities such as minimal supersymmetry [8]. A stable WIMP³ in the GeV-TeV mass range would satisfy the constraints on dark matter. An additional motivation for the WIMP comes from the so-called “WIMP miracle”. WIMPs in the early universe would exist at equilibrium with ordinary matter. As the uni-

¹Some models allow for the early universe dark matter to later decay into the present universe dark matter, but in either case particles that decay faster than the lifetime of the universe cannot be the dark matter around today.

²In brief, QCD allows for substantial violation of CP in the strong sector, but measurements of the neutrino electric dipole moment require the CP violation to be small. The axion theory is one explanation for why strong interactions preserve CP.

³For example, the lightest supersymmetric particle.

verse’s temperature dropped, however, WIMP-WIMP annihilation would become too rare and the equilibrium with baryonic matter would be broken, leaving a fixed “relic” density of WIMPs. Thus, the inferred density of WIMPs in the present day implies a scale for the WIMP’s coupling, which “miraculously” is calculated to be the at the weak scale. The WIMP miracle is no longer as compelling as it once was. Astrophysical searches for WIMP-WIMP annihilation have ruled out the annihilation cross section implied by the WIMP miracle calculations for WIMP masses below 30 GeV [9]. However, this does not exclude WIMP-baryonic matter interaction cross sections detectable by earth-based experiments and does not yet restrict WIMPs above 30 GeV.

My work is on the DarkSide-50 experiment, a detector for WIMP dark matter. The rest of this dissertation therefore focuses on WIMPs and their detection.

1.2 Detection of WIMPs

The weak coupling of WIMPs to ordinary matter is illustrated in Figure 1.1. This diagram can be read with the time arrow going in three directions, each of which indicates a specific type of interaction and also a method of WIMP detection. WIMPs (X) can interact with ordinary matter (SM) by:

- **Annihilation:** $X + X \rightarrow SM + SM$. The energetic ordinary matter produced by the self-annihilation of WIMPs is one avenue for detection. Searches for this process typically look to massive astrophysical objects where particularly high concentrations of WIMPs are expected due to gravity. Such objects can include our sun and earth itself.
- **Production:** $SM + SM \rightarrow X + X$. Particle colliders can produce dark matter from collisions of energetic ordinary matter. Dark matter will not be detected by the collider instrumentation, so production of dark matter can be identified by missing energy and momentum. To distinguish dark matter production from other processes with missing

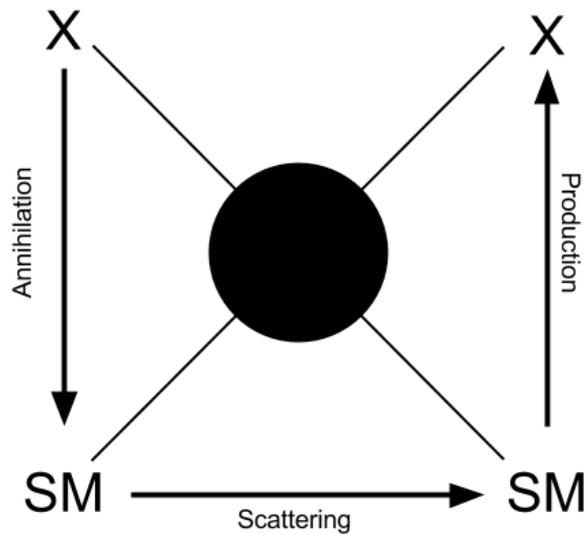


Figure 1.1: WIMP Interactions.

momentum, analysis searches for specific processes, like the “Higgs portal,” for which dark matter production would create a substantial excess signal over backgrounds [10].

- Scattering: $X + SM \rightarrow X + SM$. Dark matter can scatter off ordinary matter, transferring energy and momentum to the ordinary matter which can be detected. This type of detection is called “direct,” as it is an observation of a WIMP interaction directly with the detector.

Direct detection is appealing for a number of reasons. Earth-based experiments offer the opportunity to control backgrounds directly. Direct detection probabilities scale with the mass of the detector’s target medium, so proven technologies can be scaled up to probe lower dark matter cross sections. Direct detection experiments are complementary with annihilation and production experiments both by being sensitive to different WIMP masses and because the cross section for the different processes may differ.

DarkSide-50 is designed to detect WIMPs directly via scattering off the argon nucleus.

1.3 Nuclear Recoils Produced by WIMP Interactions

WIMPs by definition interact with ordinary matter via the weak force. In popular models, the primary coupling is to atomic nucleons, with the WIMP interacting coherently with the entire atomic nucleus⁴. These interactions transfer energy and momentum to the nucleus, causing it to recoil. The recoiling nucleus deposits energy in the surrounding matter that can be detected as a signal of the WIMP interaction.

The differential rate for WIMP interactions with nuclei is:

$$\frac{dR}{dE_R} = N_N \frac{\rho_0}{m_W} \int_{v_{\min}}^{v_{\max}} d\mathbf{v} f(\mathbf{v}) v \frac{d\sigma}{dE_R}$$

where N_N is the number of the nuclei in the target, ρ_0 is the WIMP halo density (at Earth), m_W is the mass of the WIMP, \mathbf{v} is the WIMP velocity (in Earth's reference frame), $f(\mathbf{v})$ is the WIMP velocity distribution, and $\frac{d\sigma}{dE_R}$ is the differential cross section for WIMP-nucleus scattering. v_{\min} is the minimum velocity that could produce a WIMP-nucleus scattering of energy E_R :

$$v_{\min} = \sqrt{\frac{m_N E_R}{2\mu^2}}$$

where m_N is the nuclear mass and μ is the WIMP-nucleus reduced mass, $\mu = (m_N m_W) / (m_N + m_W)$. v_{\max} is the escape velocity of WIMPs in the galaxy. Theory provides models for the WIMP density and velocity distribution, and WIMP detectors can measure the energy spectrum⁵. Such a measurement would fix the cross section for any possible WIMP mass.

WIMPs could have either or both spin-independent and spin-dependent couplings to nucleons. The larger the nucleus the higher the spin-independent (SI) WIMP-nucleus cross section, since there are more nucleons to interact with. Due to the coherent interaction

⁴For the scattering energies likely to occur, the de Broglie wavelength can be large compared to the size of the nucleus, meaning that the WIMP interacts with the nucleons all at once

⁵Assuming they can detect the WIMP at all

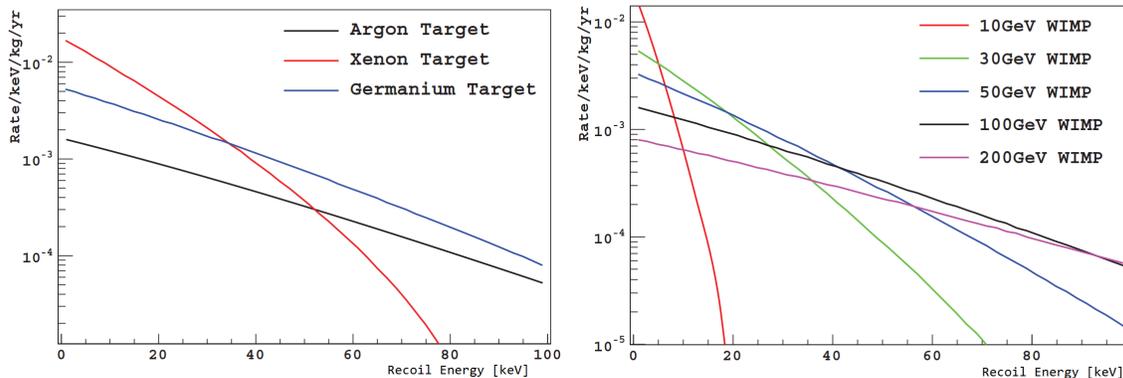


Figure 1.2: Left: the WIMP nuclear recoil interaction spectrum from 100 GeV WIMPs via spin-independent interactions with three different targets: (in decreasing order of atomic mass) xenon, germanium and argon. Right: the WIMP nuclear recoil spectrum on argon from WIMPs of different masses. In both plots, the WIMP-nucleon cross section is assumed to be 10^{-45} cm^2 . Thanks to J. Xu [11].

between the WIMP and the nucleus, the WIMP-*nucleon* scattering amplitudes essentially add⁶. Since the cross section goes as the square of the amplitude, the SI cross section is proportional to A^2 , where A is the atomic mass of the target. Due to the dependency of v_{min} on m_N , heavy nuclei have a penalty at higher interaction energies compared to lighter nuclei. This effect eventually outweighs the A^2 cross section effect, particularly for lighter WIMP mass possibilities. Figure 1.2 demonstrates how the WIMP mass and target mass influence the WIMP nuclear recoil spectrum, assuming purely SI interactions.

The spectrum increases roughly exponentially as scattering energy decreases. As a result, detectors benefit from having a low energy threshold above which they are sensitive to WIMP signals. This incentive competes with low-energy backgrounds that exist in many detector designs.

In the case of spin-dependent (SD) coupling, the coherent interaction with the nucleus causes the WIMP-nucleus cross section to be approximately proportional to $J(J+1)$, where J is the nuclear spin [12]. This makes targets with high nuclear spin much more sensitive to SD dark matter, and argon, which has no nuclear spin, not at all sensitive.

⁶Barring more exotic models. Additionally, at higher momentum transfer, as the de Broglie wavelength shrinks, there can be interference from the different nuclei.

1.4 State of the Field of Direct Detection

Figure 1.3 shows the current state of dark matter searches for spin-independent dark matter interactions.

Lines on this figure indicate the exclusion limits set by experiments that saw no evidence of dark matter. Having seen no dark matter, they can rule out the possibility that dark matter is easy to see—that is, they exclude cross sections above the line. These exclusion limits trend towards higher cross section values for higher masses, as a heavier WIMP implies a smaller number density of WIMPs in the galaxy, given the observed mass density. These exclusion limits also curve up sharply at low masses, as low-mass WIMPs produce few interactions above the energy thresholds of typical detectors.

At masses above 6 GeV, the best limit is set by the LUX experiment [13]. LUX searches for dark matter interactions with 118 kg of liquid xenon, and this result was set using 85.3 live days of data. LUX used a profile likelihood technique to analyze whether the 160 events they observed in their WIMP search energy region were consistent with a background-only hypothesis, finding that to be the case with a p value of 0.35. Figure 1.3 also shows results from the XENON100 experiment, which operates on similar principles to LUX [14]. Below 6 GeV WIMP mass, a series of experiments with lower thresholds and lighter targets set the best limit. These experiments have smaller active masses than LUX and so cannot exclude cross sections as low, but have better sensitivity to low-mass WIMPs.

Three experiments have made claims of dark matter observations, as illustrated by the shaded regions in Figure 1.3. Two of these experiments, DAMA/LIBRA and CoGeNT, observe a signal that modulates over the course of a year in a fashion explainable by the Earth’s modulating velocity relative to the WIMP halo [15][16]. One, CDMSII-Si, saw three events considered to be dark matter [17]. These results exist in tension with the experiments that exclude the masses and cross sections claimed. The claims may be ruled out by the exclusion limits, or modifications to the dark matter astrophysical or coupling model might explain the discrepancy [18][19].

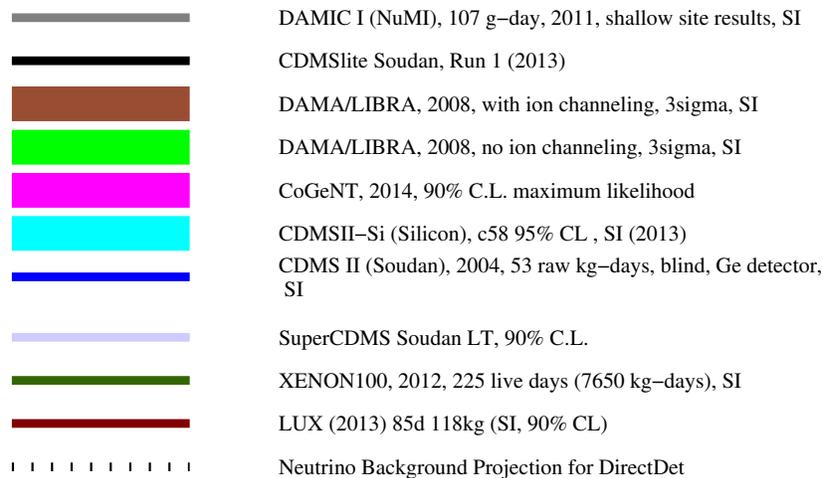
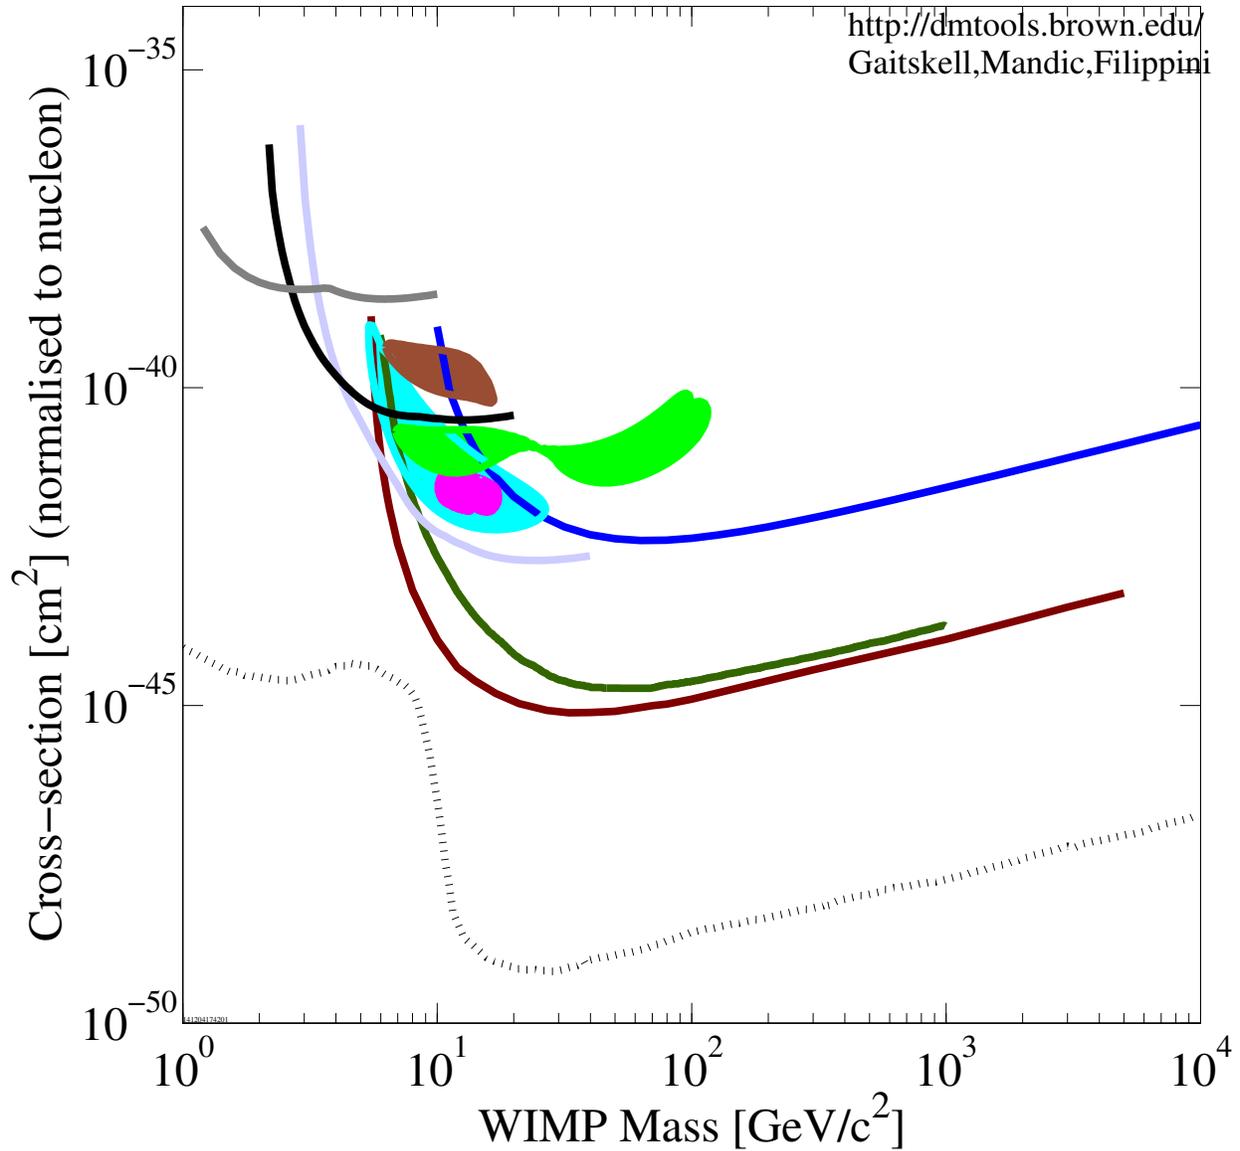


Figure 1.3: Exclusion limits set by direct detection experiments. Cross sections below the lines are excluded by those experiments. Shaded regions indicate experiments that found evidence for dark matter with mass and cross section in those regions.

Future experimental results are expected to either push the exclusion limits considerably further down or to discover WIMPs. Several experiments are currently in operation or construction to accomplish this goal—for example, LUX is expected to release further data from the same detector. In Figure 1.3 I included a then-state-of-the-art result from the CDMS experiment in 2004 to demonstrate the progress of the field in the past 10 years [20]. As detectors are built larger and larger, sensitivity will increase to the point where they exclude all cross sections above the dotted line in Figure 1.3. Below that line, backgrounds from neutrinos pose a substantial problem for many detectors. Neutrino backgrounds are difficult to reduce directly, but it may still be possible to search below that limit.

Chapter 2

The DarkSide-50 Detector

2.1 Detector Overview

DarkSide-50 is a detector for WIMP dark matter that looks for WIMP interactions in a liquid argon time projection chamber (TPC). Surrounding the liquid argon TPC are two additional detectors that detect and veto background events. All detectors are instrumented with photomultiplier tubes (PMTs) to observe light-producing interactions. Signals from the PMTs trigger DarkSide-50's data acquisition electronics, which digitize and save the PMT output over time, producing digital waveforms. The analysis of DarkSide-50 consists of identifying pulses in the waveforms corresponding to the expected behavior of particle interactions, and reconstructing the particle behavior that could create those pulses.

In the following sections, I will describe the various components of the detector and its analysis, with particular attention given to the liquid argon TPC and its analysis, as those are the components engaged in the following chapters on my contribution to that analysis.

2.2 The Liquid Argon Time Projection Chamber

The central component of DarkSide-50 is the liquid argon time projection chamber. This is a vertical cylinder of liquid argon, instrumented to detect scintillation and ionization produced

by particle interactions with the argon. Detection of the ionization component of interactions is enabled by the time projection chamber design, which applies a strong electric field to the argon and a layer of gaseous argon at the top of the cylinder. WIMPs may interact and produce signals in any of DarkSide-50's detectors, but only WIMP interactions in the liquid argon TPC will be acknowledged in our WIMP search, as only in the liquid argon TPC can DarkSide-50 reliably distinguish WIMP signals from backgrounds.

Figure 2.1 shows the liquid argon time projection chamber in cross section.

2.2.1 Liquid Argon Scintillation

Liquid argon scintillates when energy is deposited into it by a particle interaction. This signal has several features that make it useful for the detection of WIMPs and discrimination against background.

In optimal conditions¹, liquid argon scintillation produces 40 photons per keV, of which DarkSide-50 is able to observe an average of 7.9 per keV. When the electric field of the TPC is applied, the light yield is reduced to 7.0 photoelectrons per keV, as the field suppresses scintillation-producing recombination. This high number of photons help reduce the impact of statistical errors on analysis of photon signals.

Argon's scintillation comes from the creation and relaxation of two excited states, a singlet and a triplet state. The triplet state's decay to the ground state is forbidden, leading to a substantially longer lifetime for this state. Thus, the pulse of scintillation light released from a particle interaction has two components, a fast component with lifetime of 7 ns, and a slow component with a lifetime of $\sim 1.5 \mu\text{s}$. These two states are generated in different proportions depending on the ionization density of the particle interaction. Nuclear recoils, having a higher ionization density, excite a larger ratio of the short-lived singlet state than the electron recoils produced by β/γ backgrounds. Thus, signals with a smaller fast component and larger slow component as a fraction of the total scintillation are likely β/γ backgrounds,

¹Electron recoils above 1 MeV

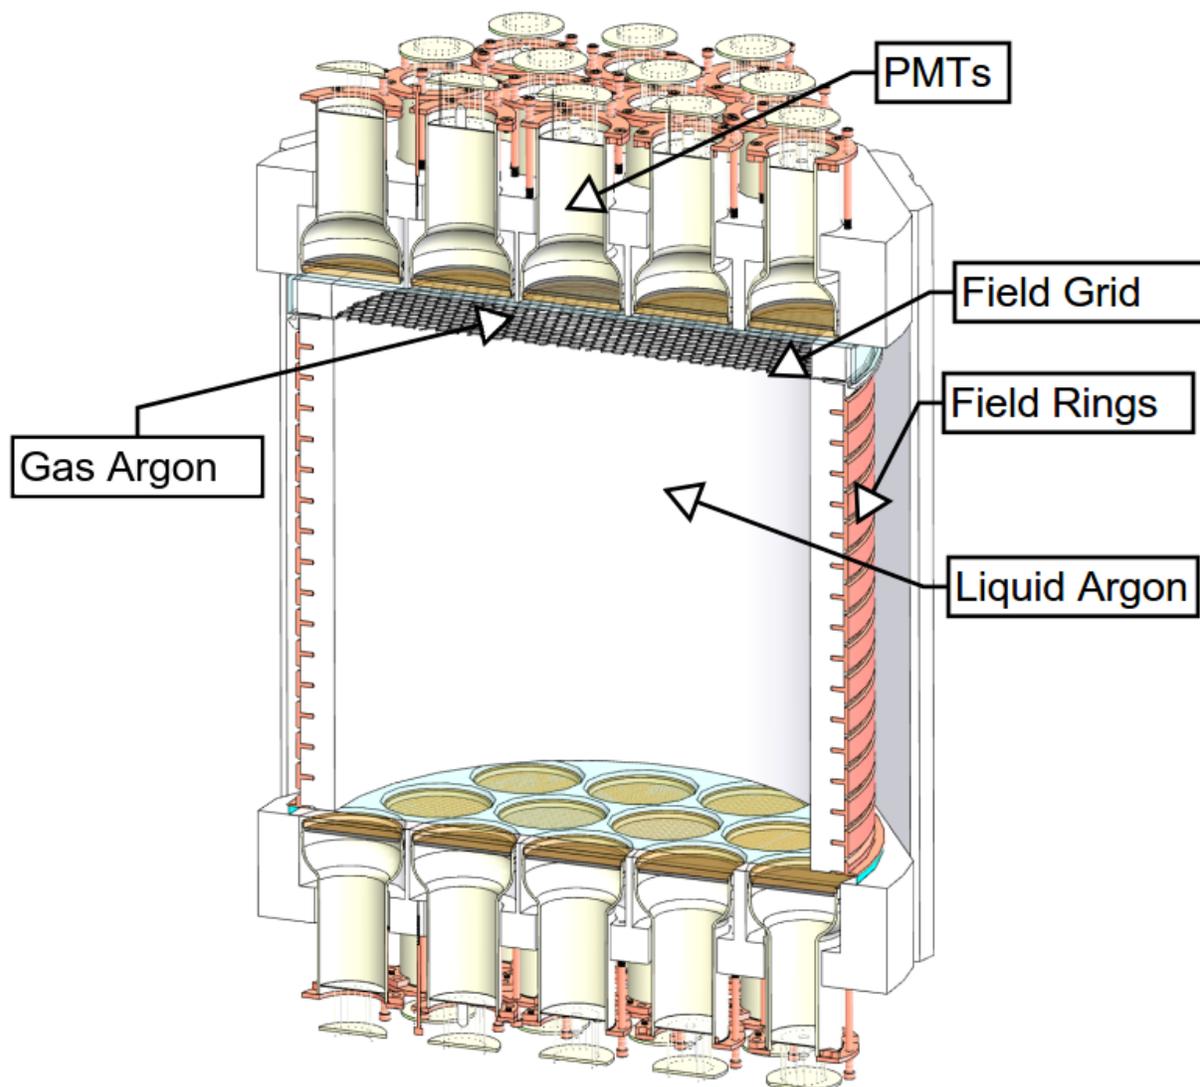


Figure 2.1: The liquid argon TPC in cross section. PMTs observe a volume of liquid argon contained in a cylinder. Copper rings around the cylinder maintain an electric field, assisted by metal grid near the top of the argon volume. Above the grid, a layer of gaseous argon in maintained to observe drifting electrons.

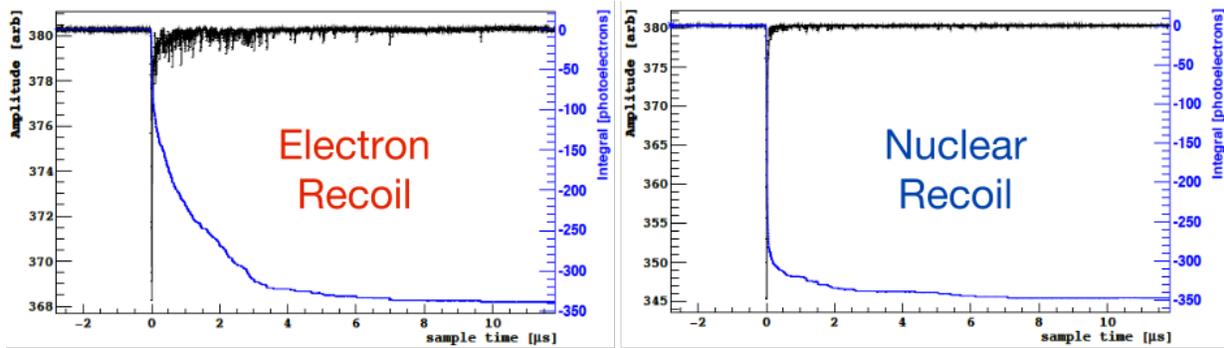


Figure 2.2: The different pulse shapes of electron and nuclear recoils. These examples were chosen to have the pulse area.

while signals with a larger fast component are likely nuclear recoils and so might be WIMP interactions. The ability to distinguish between electron and nuclear recoils by examining the time profile of the scintillation signal was a major motivator for the choice of liquid argon as the active medium of DarkSide-50. Figure 2.2 illustrates this effect. DarkSide-50 relies on the different pulse shapes of the two recoil types to perform pulse shape discrimination (PSD) to reject backgrounds at the $10^6 - 10^8$ level. The power of pulse shape discrimination depends highly on the energy of the event. At low energies, there are fewer scintillation photons emitted, and so more chance that a statistical fluctuation changes the ratio of the fast and slow components.

The difference in ionization density between nuclear and electron recoils has a second, less desirable, effect. The higher ionization density of nuclear recoils leads to quenching of the scintillation signal, in which the deposited energy is transformed into heat rather than light. As a result, nuclear recoils produce less scintillation light than electron recoils, by a factor of 0.28 (or smaller at low interaction energies).

The scintillation light from liquid argon is produced in the vacuum ultraviolet range, peaked at 128 nm. Argon itself is transparent to this wavelength, but the windows of commercially available PMTs are not. To detect scintillation light, it must first be shifted from the ultraviolet to the visible range, process described in Section 2.2.5.

This scintillation signal is the first signal observed when the detector is operated as a time projection chamber, and so it is referred to as S1. S1 signal strength is measured as the number of photoelectrons (p.e.) detected from the event; this is measured via the integral of the signal pulse from the PMTs.

2.2.2 Liquid Argon Ionization and the Time Projection Chamber Design

When energy is deposited into the argon, some of it goes into ionizing argon atoms. Some of the free electrons recombine with the ions, contributing to the scintillation light, but if an electric field is applied some of the freed electrons can be detected as a separate signal. DarkSide-50 uses a two-phase time projection chamber design to collect this ionization signal. We apply a electric field (typically 200 V/cm) vertically across the liquid argon, which causes free electrons in the argon to drift to the top of the argon volume at a speed of approximately 1 mm/ μ s. A small amount of the argon is boiled to produce a 1 cm layer of argon gas immediately above the liquid argon. A stronger electric field is applied across the gas layer and the liquid/gas interface, pulling free electrons from the top of the liquid into and through the gas². Light³ is produced from the gas when the electrons are pulled through it, creating a secondary light signal, S2. The size of S2 indicates the amount of ionization caused by the original particle interaction, and the time between the S1 scintillation signal and S2 indicates the distance the electrons had to drift from the site of the original particle interaction to the top of the argon volume.

The amount of S2 light generated is typically 10 – 50 \times the amount of S1 for electron recoils, with a lower multiplier for nuclear recoils⁴. For reasons not currently completely

²This field is typically 2.8 kV/cm in the liquid.

³Ultraviolet light with the same spectrum as argon scintillation

⁴The multiplier is not yet measured in DarkSide-50, but might be expected to be roughly 1 – 10 \times based on results from the DarkSide-10 prototype detector.

understood, the ratio $S2/S1$ appears to be position dependent: the ratio is higher near the center of the detector and lower near the edges of the detector.

Although electron and nuclear recoils produce different values for $S2/S1$, DarkSide-50 does not use this to discriminate between recoil types, as it is not believed doing so would reliably and substantially increase the discrimination power at low energies over pulse shape discrimination by itself.

A key feature of S2 production worth repeating is that it occurs in the layer of gaseous argon at the top of the detector. This layer is very close to the top PMTs, and so regardless of where on the vertical axis the original event occurred, the S2 signal will be produced close to some of the top PMTs.

The electric field is maintained using a chain of copper rings connected by resistors such that the successive rings form a voltage gradient. The top and bottom of the TPC have transparent electrodes that also contribute to maintaining the field. To produce a sharp step from the drift field in the liquid argon to the extraction field that brings electrons into the gas, a steel grid lies just under the surface of the liquid. Figure 2.3 illustrates this.

2.2.3 Containing the Argon

The active argon is contained within a vertical cylinder designed to maximize light collection.

The wall of the cylinder is made of highly crystalline polytetrafluoroethylene (PTFE), a highly reflective material. PTFE contains no hydrogen, which would interfere with the passage of any neutrons interacting with the argon to the surrounding neutron veto volume (Section 2.3).

Immediately above and below the argon are windows made of fused silica. These clear windows allow transmission of visible (but not far ultraviolet) light from the argon to the PMTs. The windows are coated in indium tin oxide (ITO), a transparent conducting substance. The layers of ITO are connected to a high voltage power supply to be electrodes

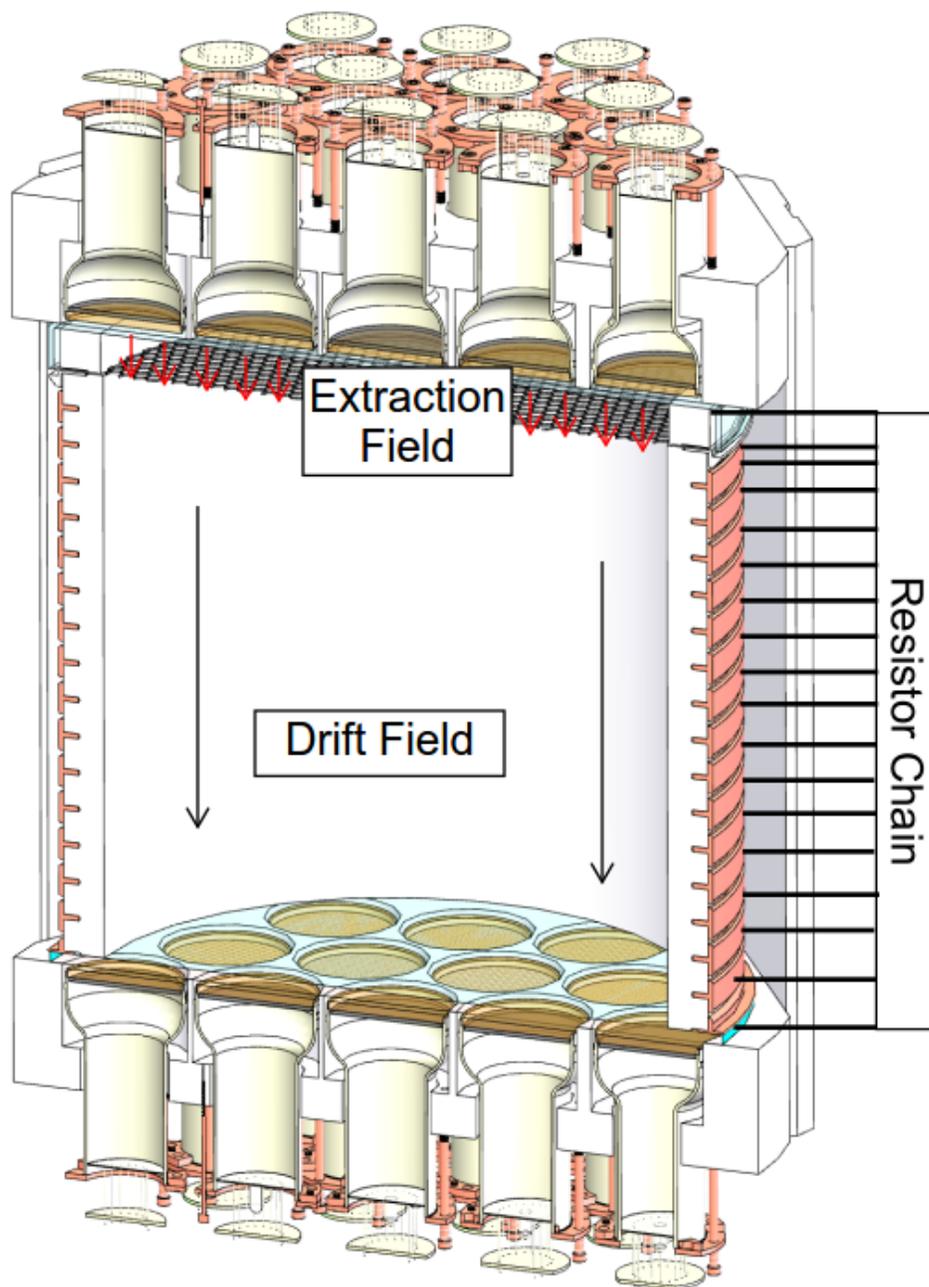


Figure 2.3: The field rings maintain a voltage gradient, creating an electric field across the liquid argon and a stronger field across the gas-liquid interface.

that maintain the electric field in the TPC. The top window has a rim, giving it a “diving bell” shape, which, along with a boiling resistor, maintains the gas layer of argon.

The PTFE wall and the top and bottom windows have their inner (argon-facing) surfaces coated with the wavelength shifting chemical TPB, discussed further in Section 2.2.5.

Above and below the windows are PMTs facing the argon. The PMTs are arranged in a hexagonal pattern, and held in that pattern by a highly crystalline PTFE structure resembling a honeycomb. The PMT holders are thus also the top and bottom reflectors for the TPC, reflecting any light that does not strike a PMT back into the detector for possible collection. The top reflector, in particular, plays a role in the results of the position reconstruction algorithm discussed in later chapters.

The cylinder, the argon it contains, and the PMTs are all contained within a cryostat and surrounded by a buffer volume of liquid argon. The cryostat is designed to facilitate the delivery of liquid argon to the inner detector and the transport of the high voltage needed to operate the TPC.

2.2.4 Photomultiplier Tubes in the Liquid Argon TPC

The TPC is instrumented with 38 Hamamatsu R11065 PMTs, 19 each on the top and bottom of the TPC. These PMTs are three inches in diameter and arranged in a hexagonal packing such that the PMT centers are 3.25 in apart. In this packing, six PMTs, the “corners” of the hexagon, are at the furthest radius from the detector center, and these six corner PMTs are partly cut off by the circumference of the wall of the TPC. Figure 2.4 illustrates this arrangement with a top-down view of the PMTs in their PTFE holder. The cylindrical side wall is superimposed on this sketch to demonstrate how the PMTs fit into the active volume.

The PMTs have high quantum efficiencies, in the 30-36% range. They also are demonstrated to work at liquid argon temperatures, a challenging requirement for PMTs. The R11065 PMTs are also designed and measured to be low in radioactivity to minimize detector backgrounds.

PMT Top Array, Viewed From Top

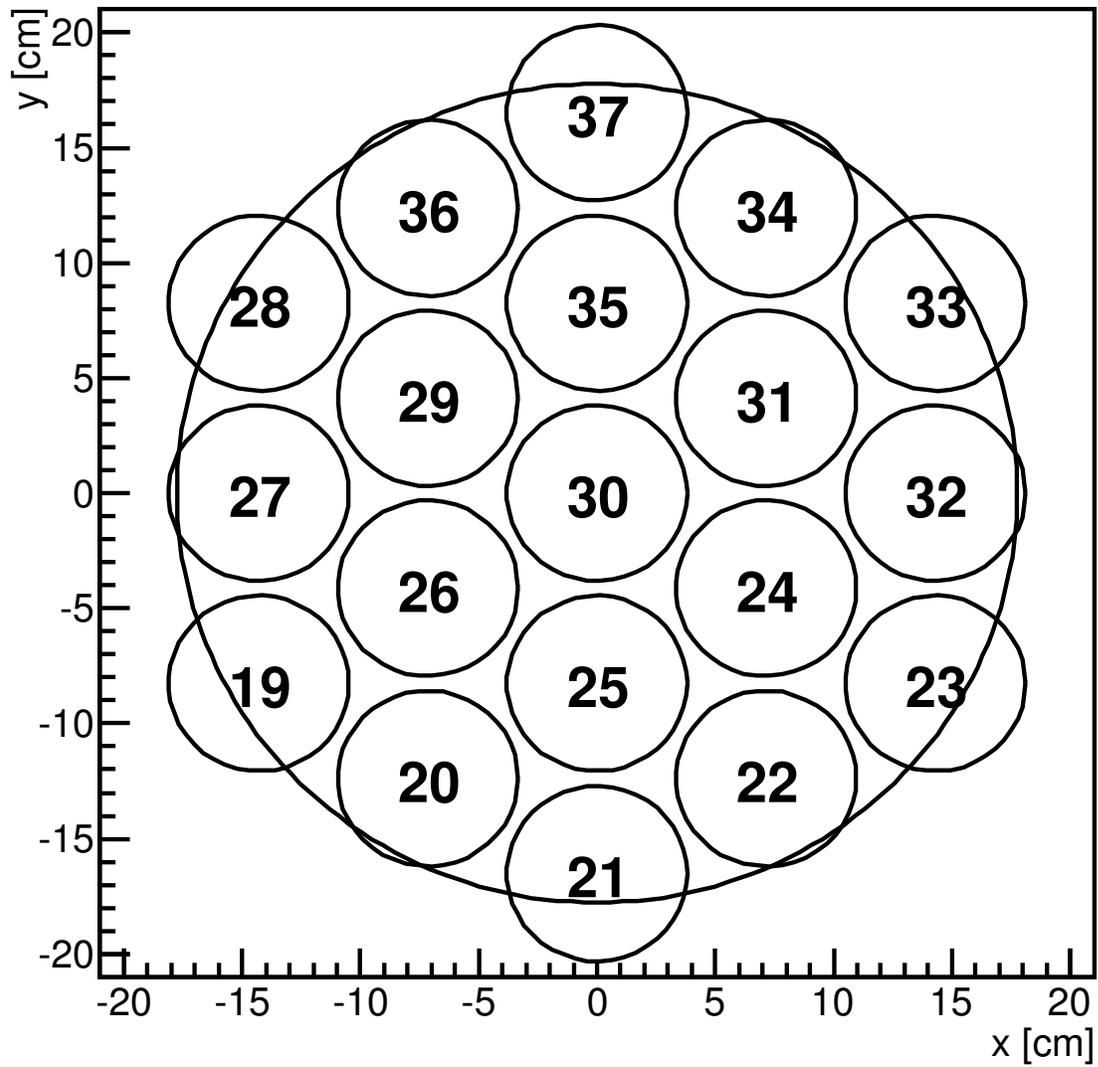


Figure 2.4: The PMT packing.

2.2.5 Wavelength Shifter

The ultraviolet light produced by argon in S1 and S2 cannot pass through the front windows of commercially available PMTs. To make S1 and S2 detectable, the ultraviolet light emitted by the argon must be turned into light in the visible spectrum which can penetrate the PMT windows. This is accomplished by coating the inner surfaces of the TPC with tetraphenylbutadiene (TPB), a wavelength shifting chemical.

TPB absorbs ultraviolet light and reemits visible light peaked around 420 nm. The fluorescence efficiency is 1.22 when absorbing 125 nm UV light [21].

TPB was applied to the DarkSide-50 components by vacuum evaporation. Each detector component was placed into a large vacuum chamber that also contained a crucible full of TPB in powder form. The chamber was evacuated, and the crucible heated, causing the TPB in the crucible to evaporate and then condense onto surfaces in the chamber, including the detector component. Thickness monitors were used to measure the progress of the evaporation, and a calibration campaign was conducted to learn what thickness measurement corresponded to the desired mass of the TPB coating.

As the solid angle per area changes over the span of each detector component, the TPB is not applied completely uniformly to the components. Based on calibration measurements, the top and bottom surfaces of the TPC have $(230 \pm 10) \mu\text{g}/\text{cm}^2$ TPB at the center and $(190 \pm 15) \mu\text{g}/\text{cm}^2$ TPB at the maximum radius. The side wall was coated twice to produce a symmetric coat. It has $(224 \pm 27) \mu\text{g}/\text{cm}^2$ TPB on the top and bottom, and $(165 \pm 20) \mu\text{g}/\text{cm}^2$ TPB at the half-height point. All components were continuously rotated by a motor during the evaporation to promote azimuthal symmetry in the TPB coat.

I personally supervised the calibration campaigns and the coating of some of the detector components. My qualitative impression is that the evaporation process produces some unevenness in the TPB coat just barely visible to the eye and often difficult to make out in photos. Observations of TPB-coated components from DarkSide prototypes after disassem-



Figure 2.5: The diving bell (left) and bottom window (right) with TPB applied.

ably suggest that TPB can also become detached from surfaces during detector operation⁵. It is not known whether unevenness in the TPB thickness leads to significant variation in the optical behavior of the detector—there exists some evidence that the overall light collection efficiency is not very sensitive to the thickness of the TPB [22]. Figure shows the diving bell (top window) and bottom window after TPB application.

In addition to wavelength shifting ultraviolet light, TPB is a diffuse reflector in the visible spectrum. At the thicknesses used in DarkSide-50, the layer of TPB is partially transparent, allowing visible light to reach the PMTs, but some light will be reflected back into the detector by the TPB. As a result, simulations of light collection in the TPC depend heavily on the model used for the TPB.

Scintillation in argon outside the TPC cannot produce signals in the PMTs as outside the TPC there is no wavelength shifter to convert the scintillation light to visible wavelengths that can reach the PMT photocathodes.

⁵DarkSide-50 contains design features intended to reduce the impact of liquid argon filling on the TPB, so it may perform better in this regard than other detectors.

2.2.6 Support Infrastructure of the Liquid Argon TPC

Operation of the TPC relies heavily on extensive support infrastructure. Of particular note are:

- The argon handling system, which purifies, liquifies, and delivers argon to the cryostat containing the TPC.
- The screening program, which measured radioactivity in proposed detector components so the lowest-radioactivity components could be chosen for use.
- The recirculation and boiling system. Argon is boiled to produce the gas layer in the TPC, as well as to return argon to the argon handling system for purification.
- The electrical system, which provides the high voltage needed to operate the PMTs and the even higher voltage required to maintain the electric field in the TPC.
- The PMT electronics. The PMTs are augmented with preamplifiers attached to their bases which allow the PMTs themselves to be operated at lower voltages, avoiding some “flashing” issues experienced with higher voltage operation. These preamplifiers were designed to operate at the cold temperature of liquid argon, and so are referred to as “cold” preamplifiers.

2.2.7 Argon-39 and the Underground Argon Program

When the studies described in this dissertation were performed, the TPC was filled with atmospheric argon. Argon, being about 0.9% of the atmosphere, is cheap and plentiful when distilled from the atmosphere. However, atmospheric argon contains the radioactive isotope argon-39 at 1 Bq/kg in atmospheric argon. Argon-39 undergoes β decay, producing backgrounds in the detector that must be removed via pulse shape discrimination. Although pulse shape discrimination is powerful, 50 Bq of backgrounds in our 50 kg active volume

over the lifetime of the detector puts a strain on the ability of PSD to remove backgrounds without also removing most of the WIMP acceptance.

Argon-39's half life of 269 years is too long to wait for the argon-39 to decay away. The energy endpoint of the β decay is 565 keV, higher than our WIMP search region, but that still leaves much of the β spectrum within the search region. The β emissions do not penetrate out to the veto detectors (Section 2.3), and so cannot be removed that way. To reduce these backgrounds beyond what PSD alone can accomplish, we must reduce the concentration of argon-39 in the liquid argon.

To achieve this goal, DarkSide-50 has a program to distill argon from underground sources. Argon-39 is produced primarily by cosmic ray neutrons, and so argon that has been underground for a long time has substantially reduced abundance of argon-39. Measurements to date suggest that the abundance of argon-39 in underground argon is reduced at least $150\times$ compared to atmospheric argon.

Distilling the large amount of argon needed from the underground source was a substantial undertaking. The production of the underground argon and its insertion into the cryostat was completed during the writing of this dissertation. Studies on the underground argon in DarkSide-50 are taking place at time of writing.

2.3 The Veto Detectors

Surrounding the cryostat containing DarkSide-50's TPC is an organic liquid scintillator detector. Surrounding that is a water-Cherenkov detector. These two detectors operate as vetoes, detecting background that might appear nearly simultaneously in a veto detector and the TPC and ruling out those events as WIMP candidates.

Given current limits on the WIMP-nucleon cross section, a WIMP interacting even once in a detector is a very rare event. Any event consisting of multiple interactions is therefore a background. When highly penetrative backgrounds, especially neutrons and cosmic muons,

interact with the TPC, it is likely that they either will continue outside the TPC and interact with a veto detector or have already interacted with a veto detector. Thus, by excluding any TPC signals in coincidence with events detected in a veto detector, DarkSide-50 can reduce the backgrounds in the WIMP search.

The liquid scintillator veto (LSV) is a sphere of organic liquid scintillator instrumented with 110 PMTs. The liquid scintillator is a mixture designed for the efficient detection of neutrons consisting of trimethyl borate (TMB) and pseudocumene, with 2,5-diphenyloxazole added as a wavelength shifter. The sphere is 4 m in diameter.

The LSV reduces backgrounds when neutrons that interact with the liquid argon do not stop in the argon, the TPC support structure, or the cryostat, but continue out into the LSV and then produce a clear signal there. The design of the TPC and cryostat minimize the chance that neutrons are captured before reaching the LSV. For example, the TPC structure includes little hydrogen to discourage thermalization of neutrons. The boron in the TMB has a large thermal neutron capture cross section, which contributes to the veto efficiency by greatly decreasing the capture time for neutrons, creating capture signals in the veto closer in time to the interaction in the TPC.

The LSV is estimated to veto about 99.5% of all neutrons that interact with the TPC. It also contributes to the rejection of γ backgrounds.

The water-Cherenkov detector is a 10 m tall cylindrical water tank 11 m in diameter, instrumented with 80 PMTs. This detector's primary purpose is to veto cosmogenic muons and other relativistic particles.

2.4 DarkSide-50 at LNGS

DarkSide-50 is located underground at the Laboratori Nazionali del Gran Sasso (LNGS), an underground laboratory operated by the Istituto Nazionale di Fisica Nucleare, Italy's national nuclear physics institute. The underground facility at LNGS is shielded from cosmic

rays by the rock overhead, the equivalent of 3800 m of water [23]. This substantially reduces cosmogenic backgrounds.

2.5 Data Acquisition

2.5.1 TPC Data Acquisition

Signals from the TPC PMTs are amplified by the cold preamplifiers and then conducted out to the control room through wires that pass through the cryostat and both veto detectors. Each PMT's signal is then amplified again and sent both to a discriminator and to a 12-bit, 250 MHz digitizer. An FPGA trigger logic board analyzes the discriminator to determine when to trigger the digitizer and save the digitized PMT signals (the “waveforms”).

In standard WIMP search operation, the discriminator threshold is set to about 0.6 the average amplitude of a single photoelectron, so the discriminator should fire most of the time on even the smallest signals. The trigger fires when three separate PMT signals exceed the discriminator threshold in a 100 ns time window. This trigger is $> 99\%$ efficient for events with $S1 > 60$ photoelectrons.

In WIMP search operation, the digitizers save 440 μ s of the PMT signal for each trigger. This span is considerably longer than either S1 or S2 pulses separately, but is long enough to include both S1 and S2 in a single waveform, even when the event occurred at the bottom of the TPC and so S2 was delayed by the maximum time.

2.5.2 Veto Data Acquisition

The PMTs in both veto detectors are, after amplification, digitized at 10 bits and 1.25 GHz. To reduce the burden of storing the veto waveforms, the veto output is zero-suppressed at the time it is acquired, so only sections of the waveform with peaks exceeding a threshold of 0.25 photoelectrons are stored.

In WIMP search operation, the veto does not trigger itself but is instead triggered by the TPC. This is because only veto signals in coincidence with signals in the TPC are of interest in the WIMP search. When the vetoes are triggered, up to $70\ \mu\text{s}$ of data is taken from the vetoes.

2.6 Fundamental Analysis

In this section, I describe the fundamental analysis needed to identify the size of the S1 and S2 signals in an event.

2.6.1 Baseline Finding

The PMTs no-signal baseline lies close to the high end of the range of the digitizers, and light signals appear as negative pulses below that baseline. To measure the size of pulses, the baseline value must be identified and subtracted.

In DarkSide-50, the baseline value can change slightly not just event by event but over the course of an event. When the baseline is misidentified and the value is set too high or too low, that mistake causes a small error in the measurement of pulse amplitudes, but a large error in the measurement of pulse integrals. So, the baseline finding analysis algorithm tracks the baseline over the course of its movement.

The baseline value is set at each time sample in the digitizer output with a moving average of a 80 ns window. Whenever this window would include any sharp excursion away from the current baseline, the averaging stops, so as not to include signals averaged into the baseline value. Instead, in the samples that are part of or nearby an excursion, the baseline value is interpolated from the nearest points where the baseline could be found. By design, the beginning of a digitized waveform precedes the trigger time by a small amount to ensure⁶ the start of the waveform contains baseline.

⁶except in cases where the tail of a previous pulse caused the trigger

Once the baseline value is set at all points in the waveform, either by averaging or interpolation, that value is subtracted from the raw waveform to produce the baseline-subtracted waveform for each channel.

2.6.2 Single Photoelectron Calibration

The correspondence between the number of photoelectrons collected by a PMT and the pulse size produced varies from PMT to PMT. We calibrate this relationship by producing signals containing small numbers of photoelectrons and identifying the average pulse size for a single photoelectron.

Pulses are measured by integrating the waveform over the pulse length. The maximum amplitude of the pulse from a single photoelectron varies considerably more than the pulse integral, mostly due to aliasing effect. So, integrals are a more precise way of measuring the photoelectron count in a signal.

To measure the average single photoelectron pulse size, we illuminate the TPC with low-intensity light from a pulsed laser, fed into the TPC by an optical fiber, and trigger the detector when the laser fires. The light intensity is low enough that the average number of photoelectrons collected by any PMT in each laser pulse is less than one. As a result, the spectrum of pulse areas observed in each PMT has distinguishable components from laser triggers with zero, one, and sometimes two photoelectrons. Figure 2.6 shows the laser spectrum in one channel. By assuming the distribution of the number of photoelectrons is Poissonian, the mean and variance of the single photoelectron signal distribution is related to the mean and variance of the total distribution and the zero-photoelectron distribution by the average number of photoelectrons per trigger. These numbers can be measured directly allowing for a determination of the single photoelectron mean and variance without many assumptions on the shape of the single photoelectron distribution.

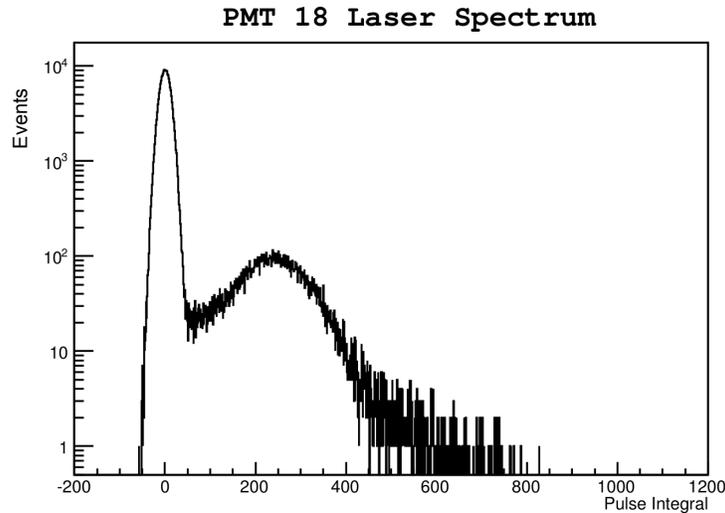


Figure 2.6: Laser spectrum in one channel. Due to the low laser illumination, most triggers see no photoelectrons. The one photoelectron signal is clear, with two photoelectrons adding a small component at the tail.

The PMT signal, q , is determined by two components: the baseline noise present in every trigger, $P_B(q)$, and the signal that depends on the number of photoelectrons captured in that event, $P_S(q)$. The distribution $P_T(q)$ is the convolution of these two components:

$$P_T(q) = P_B(q) * P_S(q),$$

which means that the statistical moments are related:

$$\text{mean}[P_T(q)] = \text{mean}[P_B(q)] + \text{mean}[P_S(q)] \quad (2.1)$$

$$\text{var}[P_T(q)] = \text{var}[P_B(q)] + \text{var}[P_S(q)]. \quad (2.2)$$

$P_S(q)$ is the distribution for all light-dependent signals, and can be broken into terms for each number of photoelectrons:

$$P_S(q) = \sum_{p=0}^{\infty} P_S(q|p) P_L(p),$$

where p is the number of photoelectrons and $P_L(p)$ is the distribution of the number of photoelectrons.

$P_S(q|0) = \delta(q)$, as no light should produce no light-dependent signal. $P_S(q|1)$ is the single photoelectron distribution we are interested in. Higher values of p involve convolving the $p = 1$ distribution with itself. $P_S(q|2) = P_S(q|1) * P_S(q|1)$, and $P_S(q|p) = [P_S(q|1)]^p$ (where exponentiation in this case refers to convolution). Thus, $mean [P_S(q|p)] = p \cdot mean [P_S(q|1)]$, $var [P_S(q|p)] = p \cdot var [P_S(q|1)]$, and so:

$$mean [P_S(q)] = \sum_{p=0}^{\infty} mean [P_S(q|p)] P_L(p) \quad (2.3)$$

$$= mean [P_S(q|1)] \sum_{p=0}^{\infty} p \cdot P_L(p) \quad (2.4)$$

$$= mean [P_S(q|1)] \cdot mean [P_L(p)], \quad (2.5)$$

and similarly

$$var [P_S(q)] = (var [P_S(q|1)] + mean^2 [P_S(q|1)]) \cdot mean [P_L(p)]. \quad (2.6)$$

Plugging equations 2.5 and 2.6 into equations 2.1 and 2.2, we can rearrange and get:

$$mean [P_S(q|1)] = \frac{mean [P_T(q)] - mean [P_B(q)]}{mean [P_L(p)]},$$

and

$$var [P_S(q|1)] = \frac{var [P_T(q)] - var [P_B(q)]}{mean [P_L(p)]} - mean^2 [P_S(q|1)].$$

All of the parameters needed to determine the mean and variance of $P_S(q|1)$ can be measured from the low-illumination data:

$mean [P_T(q)]$ and $var [P_T(q)]$ are straightforward statistical measurements of the spectrum.

$mean [P_B(q)]$ and $var [P_B(q)]$ can be measured by examining the baseline noise at a region of the waveform before the laser pulse appears.

$mean [P_L(p)]$, the average number of photoelectrons, is more complicated to measure. Assuming that $P_L(p)$ is a Poissonian distribution,

$$mean [P_L(p)] = -\ln(N_0/N_T),$$

where N_0 and N_T are the number of zero-photoelectron events and the total number of triggers, respectively. N_T can be counted directly, but N_0 must be estimated by finding a variable that efficiently separates events with photoelectrons from events without. For DarkSide-50, we used the maximum amplitude for this separation variable, since fluctuations in the baseline can produce a single-photoelectron-like integral more easily than a large amplitude⁷. By examining the behavior of the amplitude in both the pre-trigger region before the laser light appears and in the window in which laser light does appear, we can estimate the number of triggers that contain no photoelectrons in the laser window. [24] discusses this estimation process in more detail, including the statistical errors associated with the process.

At the end of this analysis, the mean pulse integral of single photoelectron signals is known for each PMT. The baseline-subtracted waveforms are scaled by this amount, so that an average single photoelectron appearing in one PMT's waveform should now have an integral of 1.

2.6.3 Pulse Finding

When S1 (scintillation) or S2 (electron drift) light is produced, the PMTs receive a pulse of light and produce a corresponding pulse of signal. The pulse finding algorithm identifies these pulses in the digitized waveform.

⁷This is a result of the method of integrating single photoelectron signals, which uses a relatively large fixed integration window to avoid trying to search for the start time of a small pulse

Although the ideal waveform contains exactly two pulses, S1 and S2 (in that order), there are a number of other possible scenarios. S2 is occasionally followed by S3, a tertiary signal believed to be caused by S2 light striking the TPC cathode. The trigger could fire on S2 instead of S1 if S1 occurred during the inhibition window after a previous trigger. Multiple events could occur in coincidence within the same digitization window, producing multiple S1s and S2s. Or, a single particle could scatter multiple times, producing a single S1 but multiple S2s. In some cases, two pulses may overlap, for example when a second S2 occurs during the tail of the previous S2. The pulse finding algorithm searches for any and all pulses rather than assuming a particular scenario.

The pulse finding algorithm works on a summed, zero-suppressed waveform. After the individual PMT waveforms are baseline-subtracted and scaled to the single photoelectron calibration, they are zero-suppressed, setting any sample that is less than 0.1 photoelectrons to be exactly zero, which suppresses any coherent noise that is negligible in any one channel but would be significant when the channels are summed. The waveforms are then summed sample-by-sample.

To speed up processing time, the initial search for pulses is coarse-grained. The summed waveform is integrated, then every 250th point of the integral waveform (i.e. every $1\ \mu\text{s}$) is taken to assemble the “downsampled” integral. This downsampled integral waveform is then differentiated twice to produce the “curvature” waveform⁸. A search through the curvature waveform identifies the coarse-grained time region in which a pulse causes the curvature to fall below⁹ a threshold. Then, a fine-grained search is performed to identify more precisely where the pulse begins by checking where the waveform goes below an amplitude threshold. An amplitude search by itself would fall prey much more easily to noise that exceeded the amplitude threshold but was not otherwise associated with pulse-like behavior. Figure 2.7 illustrates this approach. Figure 2.8

⁸NB: This is the curvature of the integral waveform, and thus related to the slope of the original waveform.

⁹Recall that pulses go in the negative direction relative to the baseline.

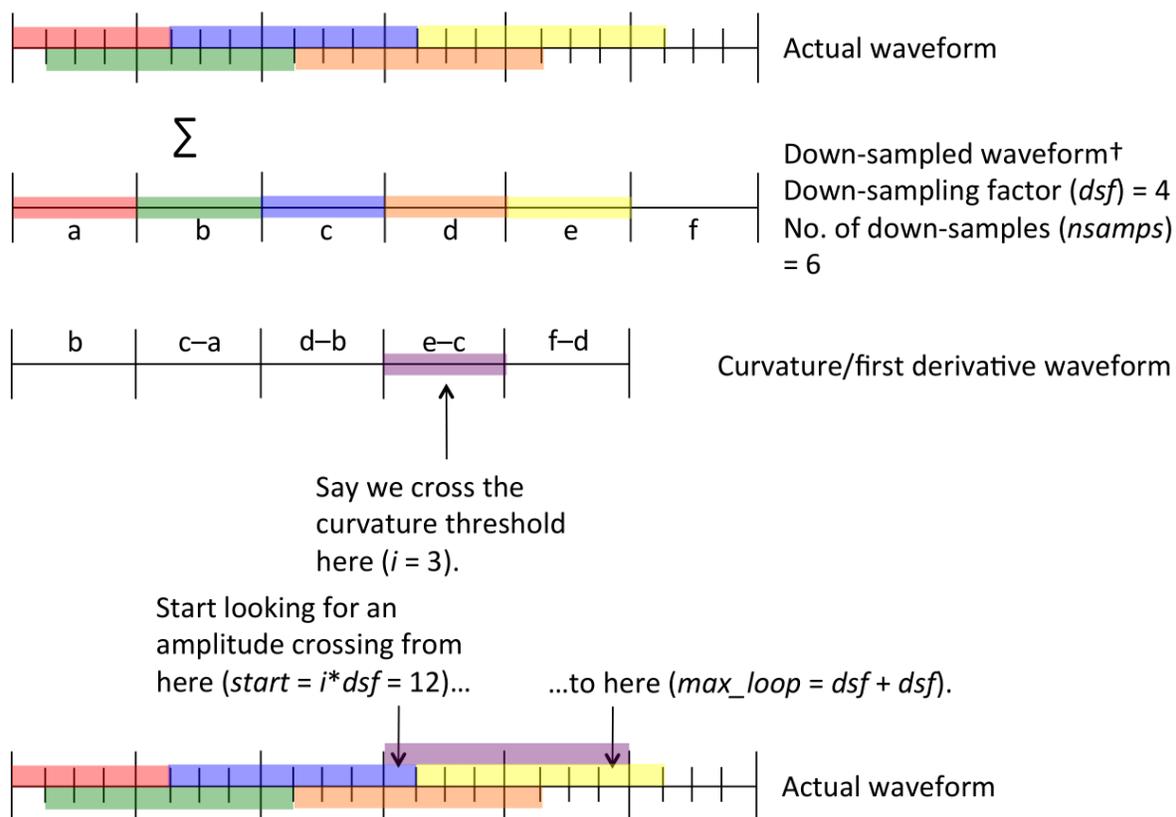


Figure 2.7: Illustration of pulse finding algorithm. Thanks to G. Koh.

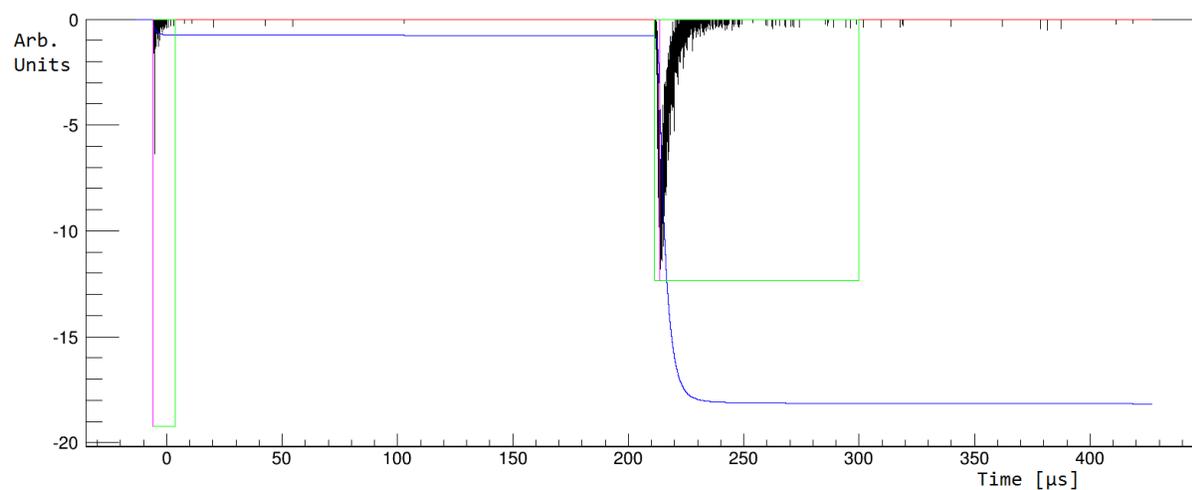


Figure 2.8: Results of the pulse finding algorithm. The green boxes show the identified pulses.

Once a pulse is found, it is checked for a pileup pulse overlapping the first pulse's tail. If the curvature waveform becomes positive, indicating a peak in the first pulse, and then becomes negative again, indicating another pulse, the algorithm concludes that there was a pileup pulse. This search uses the downsampled curvature waveform, and so is not sensitive to pileup pulses separated by less than $1 \mu\text{s}$. Once a pileup pulse has been identified, a fine-grained search is performed to find the point where the second pulse begins.

Assuming a pileup pulse does not interrupt, the pulse end time is determined by when the integral waveform's change falls below a threshold. This is equivalent to ending the pulse when the original waveform stops containing photoelectron-sized signals. Finding the end time of pulses is quite prone to error, and so typically to make measurements of the S1 and S2 size, those pulses are integrated from the start time out to a fixed length set on the typical pulse lengths for S1 and S2: $7 \mu\text{s}$ for S1 and $20 \mu\text{s}$ for S2. This prevents errors in the pulse end identification affecting the pulse size measurement.

2.7 Events in DarkSide-50

There are several important categories of events observed in DarkSide-50. I discuss three to give a sense of scale of what is observed in DarkSide-50.

2.7.1 WIMP Events

WIMP events are of course the most important category of events in a WIMP search. These events have not yet been definitely observed in DarkSide-50 or in any other detector, so this section describes how WIMP events are expected to appear in DarkSide-50.

The fundamental behavior of WIMP-matter interaction is described in Section 1.3. As mentioned there, the WIMP spectrum favors lower energy interactions and so the lower the minimum energy threshold of a detector, the higher the observed WIMP rate. As a lighter target atom than some other WIMP detectors, argon sees more high-energy WIMP

interactions. However, argon has substantial low energy backgrounds¹⁰, discussed in the following parts of Section 2.7, which makes achieving a low energy threshold difficult.

Rather than using a single threshold, DarkSide-50's analysis to date was able to eliminate backgrounds using cuts that gradually reduced acceptance of WIMP signals [25]. WIMP interactions above 54 keV have the highest acceptance, at 73.8%. Between 54 and 38 keV, the acceptance drops off as the cuts needed to eliminate backgrounds start including more of the nuclear recoil region. In the future, DarkSide-50's underground argon program (Section 2.2.7) will reduce argon-39 backgrounds and allow WIMP acceptance at lower energies.

Because WIMP interactions cause nuclear recoils, the scintillation light yield is quenched. WIMP interactions above 57.2 keV have a light yield of 2.23 ± 0.16 p.e./keV, with somewhat lower light yields below that point (about 20% lower at the end of the WIMP search region at 38 keV). Thus, WIMP signals at 54 keV, where the full acceptance begins, produce on average 120 photoelectrons in S1, and from 1 – 10× that in S2. The DarkSide-50 WIMP search so far has set an upper bound of 206.3 keV, corresponding to 460 photoelectrons [25].

2.7.2 Argon-39

Argon-39 β decays deposit up to 565 keV in the TPC. As electron recoils, they have a higher light yield than WIMP signals: 7.0 ± 0.3 p.e./keV. Thus, the argon-39 spectrum goes up to 3955 photoelectrons, although events at that high energy are generally unsuitable for analysis as they saturate the digitizers during the S2 pulse.

2.7.3 Surface Backgrounds

Alpha decays originating from most locations in DarkSide-50 are not backgrounds of concern [6]. Alpha emitters suspended in the active liquid argon will deposit the full decay energy into the active volume. That energy is well above DarkSide-50's WIMP search region, even

¹⁰Both because there is more emission of low-energy backgrounds than higher energies, and because pulse shape discrimination is less powerful at low energies.

taking into account quenching. Alpha emitters inside other detector materials will deposit no energy into the active volume, due to the short range of alphas in solids and liquids.

At one, literally narrow, class of locations, alpha decays are of serious concern. Alpha emitters attached to the surfaces of the active volume can produce events in which much of the event energy is deposited into the inactive material surrounding the liquid argon, but some energy does make it into the argon. This brings the event energy down into the WIMP search region. These events can be hard to differentiate from WIMP signals via pulse-shape discrimination, as they produce pulse shapes that are similar to nuclear recoils (or nearly identical, if most of the deposited energy of the event comes from the recoil of the nucleus that emitted the alpha). The rate of these events should be low, due to strenuous efforts made before the assembly of DarkSide-50 to keep the detector surfaces clean, but there are still some alpha emitters on the surfaces. Some alpha emitters attached before the assembly, despite the precautions, and remain there. Others have their way to the surfaces since the detector began operation. Uranium and thorium in the various components of the detector eventually decays into radon, which can migrate out of those components and into the argon. Although DarkSide-50 has a radon trap that should keep much of this radon out of the inner detector, some may nonetheless reach the TPC and decay there, producing radon daughter isotopes that can stick to the surfaces and emit alphas.

Thankfully, energy deposits from alpha decays on the surface are limited to a narrow region around the edge of the detector. A 5 MeV alpha particle can travel an average of only $46\ \mu\text{m}$ in liquid argon [26]. The emitting nucleus had a factor of roughly 100 more stopping power and carries much less energy, so will go even less far [27]. A cut that identifies energy deposits near the surfaces of the detector will remove these potential backgrounds.

The top and bottom surfaces of the TPC can be cut very efficiently by the z -coordinate reconstruction enabled by the TPC design of DarkSide-50. The drift time of the ionization signal precisely indicates the vertical position of the original energy deposit, and so surface backgrounds from the top and bottom can be removed by cutting very short and very long

drift times. To cut backgrounds from the cylindrical side wall of the detector, another method must be used that can locate an event's position on the x and y coordinates. The DarkSide-50 detector provides useful information about the horizontal position of events as described in Section 2.8, and I present a method for using that information in Chapter 3.

The rest of this dissertation assumes that surface backgrounds from the side walls are a serious concern, based on projections in the design stage of DarkSide-50. It bears mentioning that surface backgrounds in the actual detector may be far below projections. These backgrounds have yet to be observed in DarkSide-50, contrary to some predictions that they would occur at a high enough rate that extant DarkSide-50 data should contain several. So, it is possible some currently unaccounted for feature is already removing these backgrounds. Possibilities include:

- that the efforts to keep the surfaces clean were extremely successful
- that the drift field moves alpha emitters to the cathode or anode instead of the side walls
- that the wavelength shifting coating on the detector surfaces produces scintillation light during an alpha event in a way which alters the event's signal so it does not pass other event selection cuts
- that drift field irregularities extinguish the drift electron signal for surface events

Even if these backgrounds have not yet been observed, DarkSide-50 will take many times more data than it has so far, so we must be prepared for the possibility surface backgrounds will show up in the future.

2.8 Information Available for xy Reconstruction

The position of an event in DarkSide-50 on the x and y axes should be highly correlated with the fraction of S2 light collected in each PMT. The ionization electrons produced by

an energy deposit should drift straight up from the location of the deposit. When those electrons reach the gas phase of the detector and produce S2 light, they will still have the same x and y coordinates of the original event. As the S2 light is produced in the gas phase, it is produced between 0.6 and 1.6 cm from the top PMTs¹¹. Thus, the solid angle of nearby PMTs observed from the origin of the S2 light is much larger than in PMTs further away. Events at a given xy coordinate put much of their S2 light into the PMT directly above, and less and less light into PMTs further and further away on the horizontal plane¹². Put differently, the xy coordinate of an event can be determined by observing how its S2 light was split among the various PMTs.

Figure 2.9 shows the DarkSide-50 TPC in cross section, with the z axis vertical on the page. An energetic particle deposits energy at the yellow starburst labeled “S1.” Electrons are freed in that energy deposit and then drift straight upwards to the gas gap, producing secondary signal light in the gas gap at the yellow starburst marked “S2.” The top PMTs in this sketch are shown with colored bars indicating how much S2 light that PMT collected. The PMT directly above the original event collects the most light, followed by other PMTs in order of how close they are to the original event’s position on the horizontal plane. In a real event, a position reconstruction algorithm should reconstruct the original position by observing the light collection information represented by the colored bars.

2.9 Simulation of DarkSide-50

g4ds is a *Geant4*-based simulation that attempts to simulate many of the physical processes influencing signals in DarkSide-50:

- energetic particle tracking

¹¹S2 light is produced all along the vertical path the electrons travel in the gas. As S1 is produced at the site of the original event interaction in the liquid argon, S1 is produced further from any PMTs (except for events at the bottom of the detector)

¹²See the end of Section 3.2 for a discussion of the differences in light collected in the top and bottom PMTs.

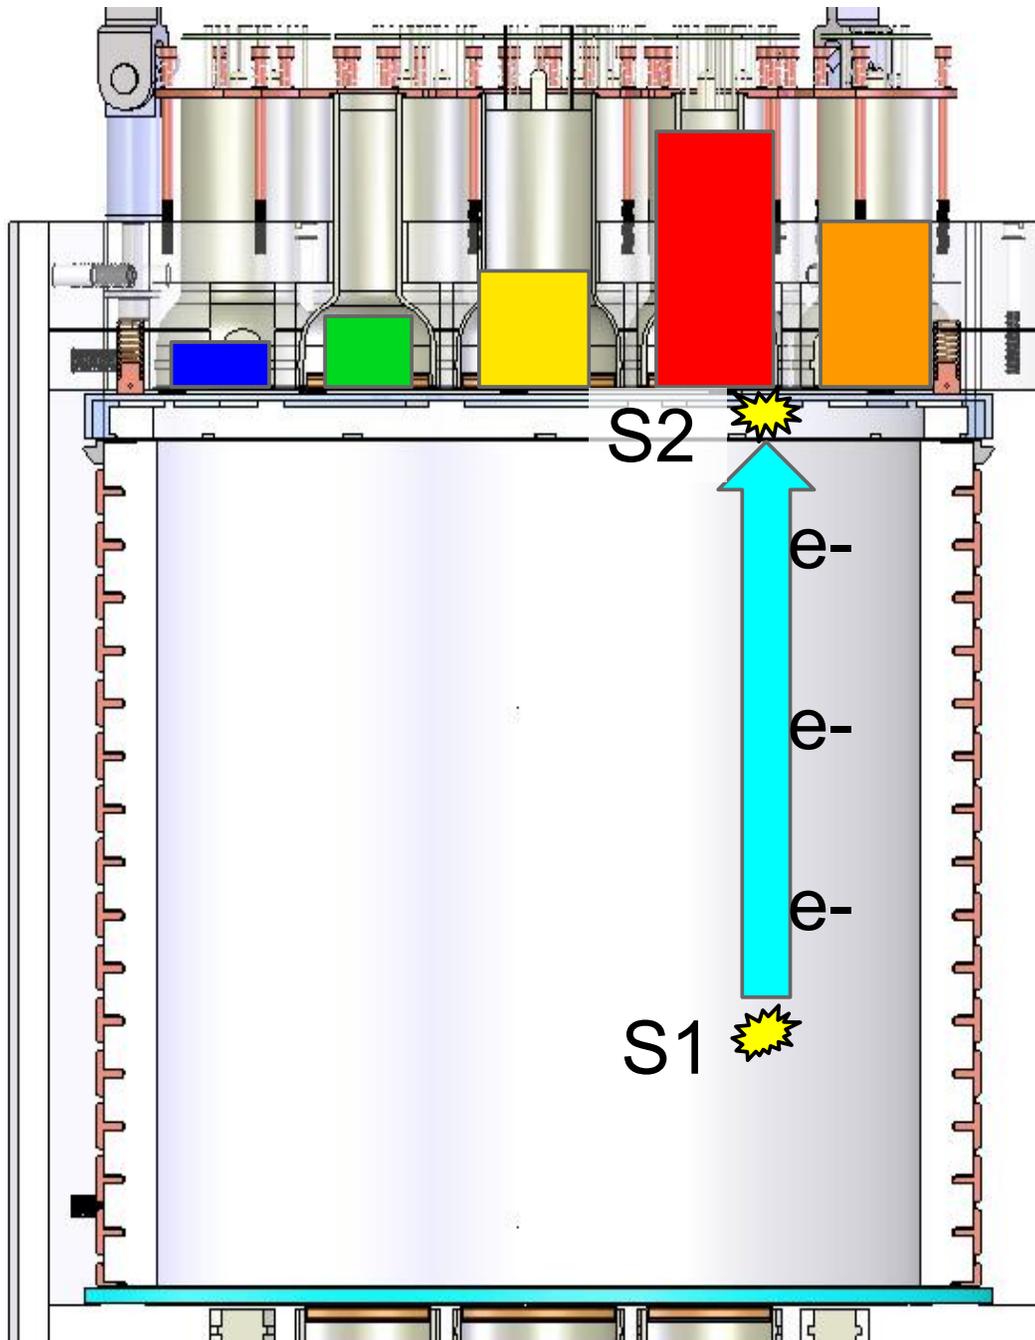


Figure 2.9: Sketch of drift electrons producing S2 light collected by PMTs.

- energy deposition in the liquid argon
- scintillation and drift electron production
- electron drift and diffusion to the gas layer
- S2 light production in the gas layer
- optical photon propagation to the PMTs, including:
 - wavelength shifting by TPB and reflection
 - absorption and refraction by the detector materials.

This is in contrast to some other simulations available to the DarkSide-50 project, including simulations designed to only simulate S2 light propagation.

The developers of *g4ds* have put some substantial effort into comparing *g4ds* S2 behavior with S2 data from the real detector and tuning *g4ds*'s parameters to bring the simulation more in line with the real behavior. That project is, however, incomplete at the of time of this writing. While I expect additional progress to be made in bringing *g4ds* closer to data and expect that progress to improve the position reconstruction, my efforts on this position reconstruction method have always assumed that DarkSide-50's position reconstruction method must not rely on a completely accurate Monte Carlo simulation. My method is designed to work around inaccuracies in the simulation through the iterative process, as discussed in the next chapter.

2.10 Current Status of the DarkSide-50 Detectors

DarkSide-50 finished construction and commissioning in October 2013. At the point, the TPC was filled with atmospheric argon. Atmospheric argon, with its substantial rate of electron recoil backgrounds, provided ample data for testing the response of DarkSide-50 to electron recoil backgrounds.

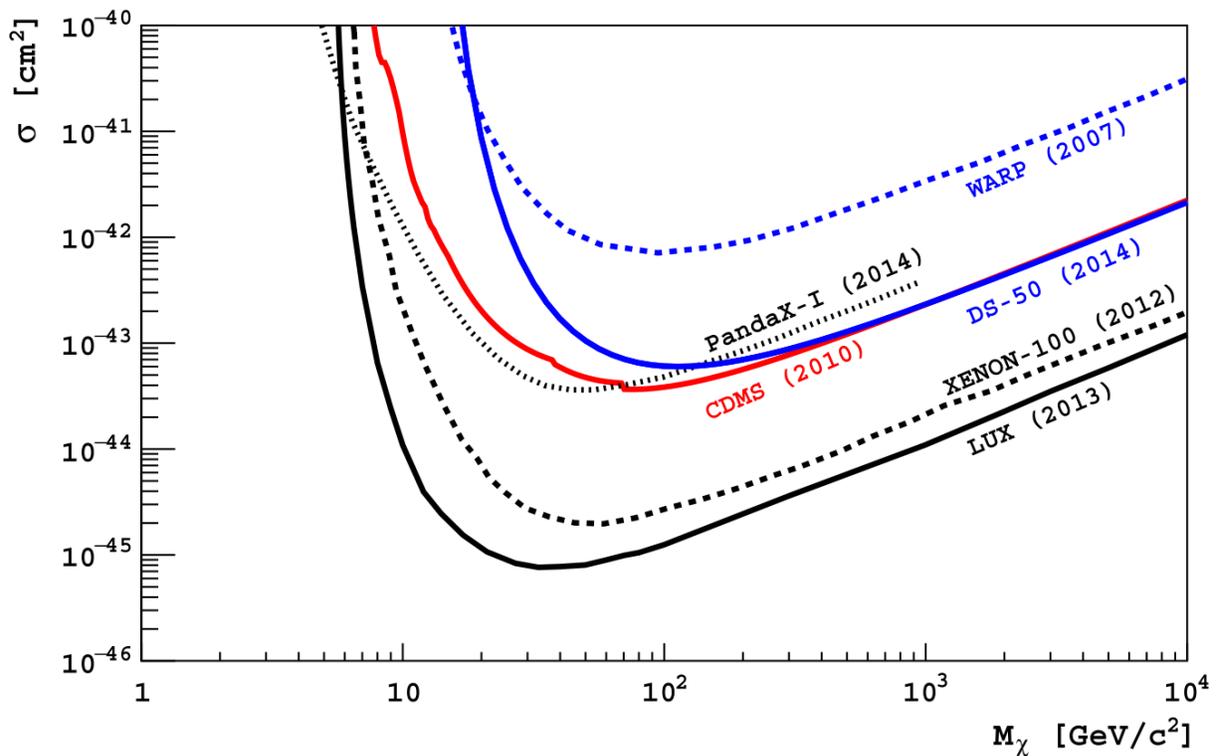


Figure 2.10: The DarkSide-50 exclusion limit along with limits set by other experiments.

From October 2013 to June 2014, the detector was operated in WIMP search mode. This provided a “shakedown” of all systems necessary for a WIMP search, including both vetoes and the pulse shape discrimination. Over this year, 47.1 days of usable data were collected. A full analysis chain was developed, sufficient to analyze and exclude all backgrounds in those 47.1 days. This analysis chain did not include position reconstruction on the x - and y -axes, as the algorithm was not yet ready.

After cuts, the exposure in this sample was (1422 ± 67) kg days. No signals remained in the WIMP search region, and so DarkSide-50 placed an exclusion limit on the dark matter cross section, shown in Figure 2.10. As DarkSide-50 was loaded with atmospheric argon at this time, it had to operate with a reduced WIMP search region to avoid argon-39 backgrounds. This reduced the sensitivity of the detector, so our results exclude less of the cross section space. These results were reported in [25].

Atmospheric argon has a rate of argon-39 backgrounds at least $150\times$ higher than underground argon, so the 47.1 days of atmospheric argon data provided as many argon-39 backgrounds as will be seen in over 19 years of underground argon operation. As a result, the ability of DarkSide-50 to eliminate argon-39 backgrounds in the WIMP search region is no longer only a projection, but is demonstrated.

Since the release of the atmospheric argon data, DarkSide-50 has been loaded with underground argon. We have also performed a calibration campaign with external sources, including neutron sources so the nuclear recoil response could be examined. Changes to the chemical composition of the liquid scintillator veto were performed to improve the neutron veto efficiency. These effort should allow a WIMP search with the full capabilities of DarkSide-50 to begin in the near future.

Chapter 3

xy Reconstruction Methodology

3.1 Weighted Least Squares Event Reconstruction

The weighted least squares (WLS) method evaluates the goodness of fit between a measured event and the expectation of that measurement given some parameter. In the case of position reconstruction, the WLS method evaluates the measured S2 light in each PMT against the expected light in that PMT for events at a given position.

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

where M_i is the measured S2 signal in PMT i for a given event, M_{tot} is the total measured S2 in all PMTs, L_i is the expected light collection (as a fraction of total light) in that PMT for an event at position (x, y) , discussed further in Section 3.2, and $\sigma_{M_i}^2$ is the squared uncertainty of the measurement. The collection and conversion of S2 photons into detected photoelectrons is approximately Poissonian, so we estimate $\sigma_{M_i} = \sqrt{M_i}$.

Minimizing χ^2 determines the position that best matches an event's measured S2 signal with the expected signal. Finding the absolute best estimate of the location parameters (x, y) is not the only goal; it is also desirable to understand how plausible other possible positions are. For example, an event that is plausibly at the surface of the detector is of

concern as a background, even if the minimum χ^2 appears at a location far from the surface. The WLS method allows any position to be compared to other possibilities by comparing χ^2 .

3.1.1 The χ^2 Distribution

χ^2 as defined here is almost, but not quite, the sum of independent normally distributed random variables, the standard definition of a chi-squared measurement. Each channel's light measurement is nearly independent. If the measurement in one PMT fluctuates upwards, it may be because more photons striking that PMT converted to photoelectrons, an independent process. It may also be because more photons took a path to that PMT instead of some other PMT, a correlated process. $\sigma_{M_i}^2$ is estimated, and the error in that estimate throws off the normal distribution of the terms of χ^2 . These differences mean that χ^2 will not end up distributed exactly according to the standard χ^2 distribution. However, that distribution is nonetheless a helpful estimator of the probability of observing a particular χ^2 as the best fit to a data event. This probability function is used to weight analysis of different positions by the probability of each of those positions being the best fit for an event. Sections 3.6.5 and 3.6.6 give some examples of how this probability is used.

The number of degrees of freedom relevant for determining the distribution of χ^2 is 18. There are 19 PMTs on the top of the detector, plus a “super-PMT” made by combining the bottom PMTs (see Section 3.2.1), for a total of 20 terms in χ^2 . The number of degrees of freedom is reduced by two for being minimized in two parameters, resulting in 18 degrees of freedom.

I found the minimum χ^2 for a set of events and fit to that a standard χ^2 distribution, with the number of degrees of freedom as the fit parameter. Figure 3.1 plots the minimized χ^2 for the data as a cumulative distribution, overlaid with the best fitted standard χ^2 distribution with $NDF = 22$, shown in blue. In red is shown the χ^2 distribution with 18 degrees of freedom, for reference.. The fit is only approximate, demonstrating the differences between

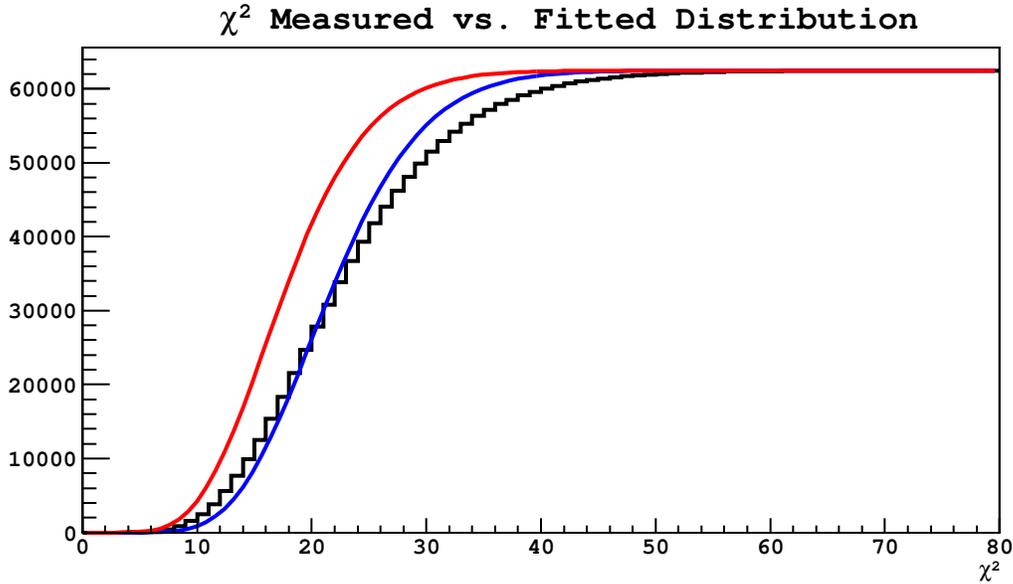


Figure 3.1: The observed distribution of χ^2 , the best fit of the standard χ^2 distribution, achieved with NDF=22 (blue), and the standard χ^2 distribution with NDF=18, for reference (red).

the measured χ^2 variable and the ideal behavior of a chi-squared variable. In this case, the best fit of the chi-squared distribution was at 22 degrees of freedom, not 18. Since the cause of the discrepancy is unknown, I continue to evaluate my χ^2 with 18 degrees of freedom.

The standard χ^2 distribution is used by this algorithm to assign probabilities that different coordinates are the correct location for the event. Section 3.6.5 describes how this is used to adapt the algorithm based on the χ^2 from multiple positions, but that process is not very sensitive to the difference between 18 and 22 degrees of freedom. Section 3.7.4 describes how the standard χ^2 distribution is used in cutting surface backgrounds, but the effect of 18 versus 22 degrees of freedom can be compensated for by changing the cut threshold. Thus, my choice of 18 degrees of freedom, if wrong, does not introduce serious errors.

3.2 Light Response Functions

A light response function (LRF) describes the expected S2 light collection fraction of a given PMT for events at a given point:

$$L_{PMT}(x, y) = \frac{\text{S2 in this PMT}}{\text{Total Observed S2}}$$

for events at position (x, y) . Determining the LRFs is the core task of this xy reconstruction method, since once the LRFs are known, the WLS method can evaluate the position of any event. LRFs are expressed as a fraction of the total S2 signal because they express an optical property independent of the total S2.

The LRFs of the PMTs describe fixed optical and signal properties of the detector. They do not change event by event, unless the optical properties of the detector change, for example by reducing the size of the gas gap or by degradation of the wavelength shifter. Instead, they vary by position. For the top PMTs, it is generally true that for (x, y) closer to the horizontal coordinates of the PMT center, the value of L is larger.

For example, PMT 30 is the center top PMT, located at horizontal position $(0, 0)$. At the center of PMT 30, $L_{30}(0, 0) = 0.34$. That is, on average 34% of all collected S2 light from an event located at $(0, 0)$ ends up collected by PMT 30. As we look further away from the PMT, the LRF becomes smaller. Three centimeters away, $L_{30}(0, 3) = 0.13$, and ten centimeters away, $L_{30}(10, 0) = 0.02$ (Figure 3.2).

The bottom PMTs are further from the emission of the S2 photons, and so their LRFs are less correlated with the distance of the event from the PMT. Figure 3.3 shows the LRF for PMT 3, which is the leftmost bottom PMT on the x-axis, at $(-14.3, 0)$ cm and PMT 17, which is the rightmost bottom PMT on the x-axis, at $(+14.3, 0)$ cm. There is little difference in light collection between the two, despite the very different positions of the PMTs. The most visible pattern is a honeycomb shape, with the “walls” of the honeycomb seeing more light than the “chambers.” This pattern is the result of the design of the DarkSide-50 detector,

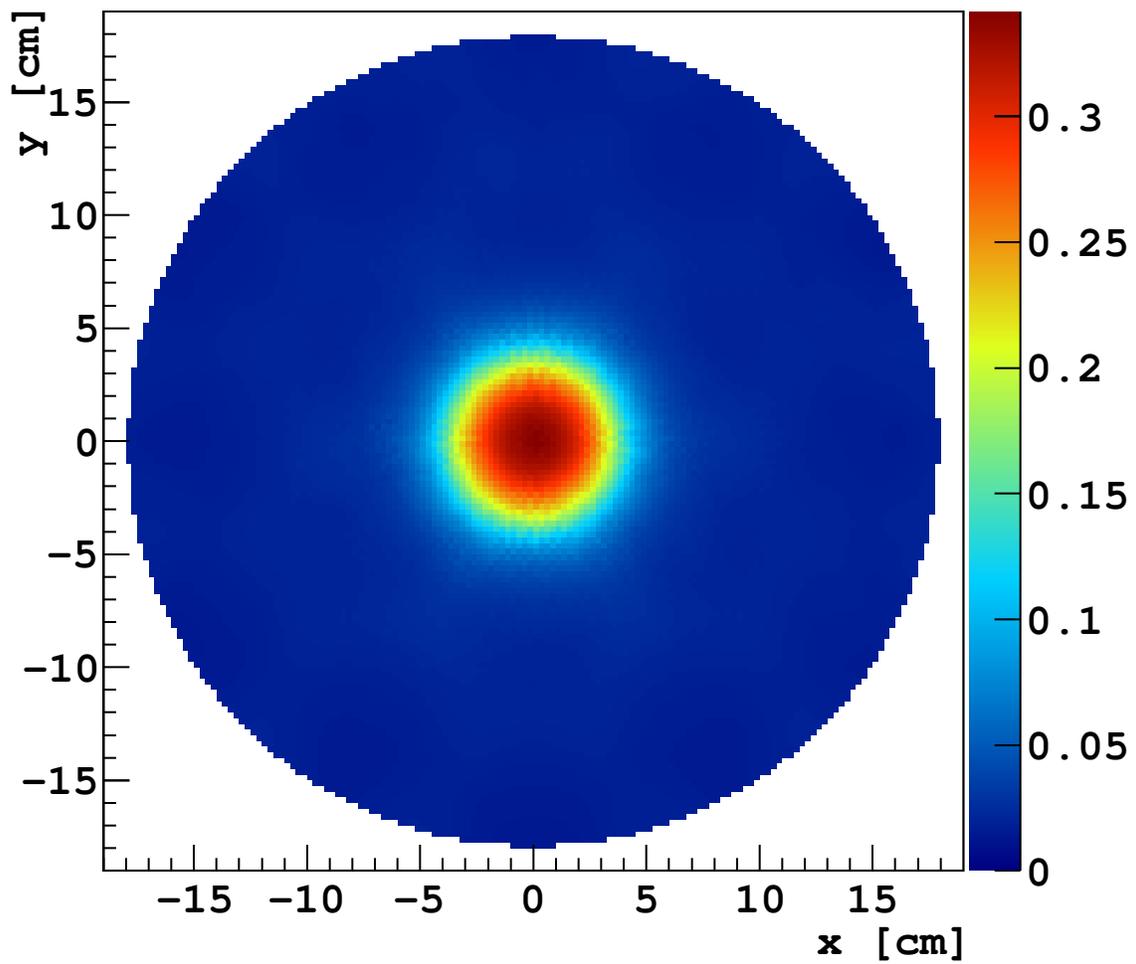
LRF 30

Figure 3.2: The light response function for PMT 30.

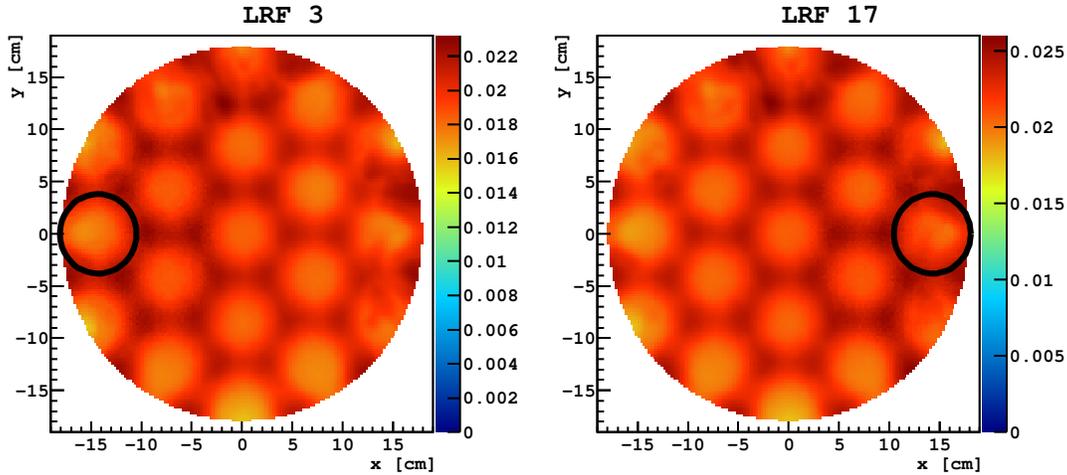


Figure 3.3: The light response function for PMTs 3 and 17, with the PMT locations illustrated.

which has reflective PTFE in the gaps between the PMTs. S2 light generated under this reflector is reflected down in greater amounts than S2 light generated under a top PMT.

3.2.1 The Combined Bottom LRF

Because the bottom PMT LRFs show mostly identical behavior, each individual bottom PMT LRF past the first does not contribute that much new information to the WLS minimization. Each additional LRF adds the risk that a particular location will be favored by the WLS minimization because of a statistical fluctuation in the LRF. For top PMTs, this drawback is far outweighed by the position information provided by the LRF, but not so for bottom PMTs. The reliability of the WLS minimization can be improved by combining all the bottom PMTs into one “super PMT.”

All bottom PMTs are stored individually, but when the WLS method is evaluated, the sum of S2 light in all bottom PMTs is compared against the sum of all bottom LRFs at the location in question. Equivalently, for the bottom PMTs I square the sum of differences instead of summing the squares of differences.

$$\chi^2 = \frac{1}{\sigma_{M_{\text{bot}}}^2} \left(\sum_i^{\text{bot PMTs}} (M_i - L_i(x, y) M_{\text{tot}}) \right)^2 + \sum_i^{\text{top PMTs}} \frac{1}{\sigma_{M_i}^2} (M_i - L_i(x, y) M_{\text{tot}})^2$$

This approach was chosen over treating all 38 PMTs individually because it performed better in tests on simulated events. In these tests, combining the bottom PMTs improved the average reconstruction error by a small amount.

3.2.2 Far Smoothing

I apply a smoothing algorithm to each LRF, limited to the region of the LRF far from the center of the corresponding PMT. Smoothing reduces the impact of statistical uncertainties in each bin by using additional nearby events to inform the value of the bin. However, smoothing also reduces the information content of the LRF by “smoothing out” real optical changes. This downside is especially pronounced wherever the LRF is changing significantly over a short distance, such as close to the PMT and across the boundaries of the top reflector PTFE.

As in Section 3.2.1, I want to keep all useful information in the LRFs while minimizing the impact of LRF uncertainties. Therefore, I apply the smoothing algorithm only to bins at least 12 cm from the center of the PMT. At a given location, the bin in that location will be smoothed in some LRFs and not in others, according to the distance between that location and the LRF’s corresponding PMT.

The smoothing algorithm, when applied, sets the contents of the bin to be a weighted average of nearby bins. The bin’s own value gets the most weight, with less weight given to

from events with the highest photoelectron statistics possible, which limits the uncertainties on each event.

Naturally, no finite set of events can cover the infinite number of coordinates in the detector. There are two possible solutions to this. One is to divide the detector into “bins,” discrete locations where the LRFs are tracked and where positions can be evaluated with the WLS method. This is the approach I used, and is described in more detail in Section 3.6.1. The alternative is to define some interpolation function that can evaluate the LRFs between known events. Fitting an analytic function to the events is one example of this kind of interpolation. This approach was not chosen because not enough was known about the shapes of the LRFs to choose an appropriate analytic form that could fit the data well.

Another important consideration is that only simulated events ever have an absolutely known location. The location for data events is never known *a priori*, but only estimated using the position reconstruction algorithm. LRFs can be built from these reconstructed positions for data events, with the understanding that if the reconstruction was incorrect the LRFs will be inaccurate to some degree.

3.2.4 The Chicken-and-Egg Problem and Its Solution

When the LRFs are known, the WLS method can evaluate the position of an event. When event positions are known, the LRFs can be determined. This is the “chicken-and-egg problem,” in that it seems the best way to reach either of our goals is to have already satisfied the other.

Like the real chicken-and-egg problem, the solution is evolution through iterative generations. The egg of a dinosaur can hatch a proto-chicken, and the egg of the proto-chicken can hatch a more modern chicken. Similarly, a prototype LRF can enable a flawed reconstruction which can be used to build a better LRF.

At the largest scale, the process is this:

1. Produce initial LRFs from simulated events

2. Reconstruct data using initial LRFs
3. Build entirely new LRFs from the results of step 2. Impose some expectations to make the new LRFs better than the previous ones.
4. Repeat steps 2 & 3 as long as each repetition improves the LRFs

The second half of step three is particularly important. Without some intervention by the algorithm, simply the new generation of LRFs will not be any better than the previous generation. However, we cannot simply insist the new generation be “better,” that is, to more accurately describe the light collection in the detector. If we knew exactly what an accurate LRF looked like, we would not be trying to build one through iteration. Instead, the algorithm imposes one feature expected in accurate LRFs.

This approach was first demonstrated in the ZEPLIN-III experiment, discussed in Section 3.3. That experiment imposed a symmetry assumption to create improved LRFs. In the DarkSide-50 experiment, substantial new work was performed to replace the symmetry assumption with a uniformity assumption. That work is discussed in Section 3.4.

3.3 Prior Work

Solovov et al., writing as part of the ZEPLIN-III collaboration, gave the first demonstration of the solution to the chicken-and-egg problem discussed in Section 3.2.4 [28]. They report a scheme in which an initial reconstruction of events is performed with a centroid method,² known to have substantial biases near the edges of the detector. LRFs are built from this initial reconstruction, with the understanding that these LRFs will incorporate the bias of the centroid method. These LRFs are used to perform a new reconstruction of the data, and the results of that new reconstruction create a new generation of LRFs. They repeat these steps until “some convergence criterion is reached” and give some examples of possible

²i.e. barycenter, or center of mass method. An event’s position is estimated by an average of the PMT positions weighted by the S2 light fraction in each PMT.

criteria for convergence. In their case, the LRFs stopped changing after five iterations of the LRF and reconstruction steps.

Solovov et al. demonstrated a resolution between 1.6 and 3 mm FWHM in ZEPLIN-III. They measured this by adding a specially designed insert into the detector that would produce a distinctive pattern with sharp edges in the spacial arrangement of energy deposits when the detector was exposed to a calibration γ -ray source. The resolution was determined by comparing the sharpness of the observed pattern to a step function convolved with a Gaussian.

3.3.1 Differences between past work and this work

Solovov et al. assumed that the light response functions in the ZEPLIN-III detector were functions purely of the distance of the event from the center of the PMT: $L_i(r_i) = L_i(|\vec{x}_{\text{PMT}} - \vec{x}_{\text{event}}|)$. That is, they assumed azimuthal symmetry in light collection³. This does not apply in DarkSide-50. See Section 4.2.2. From this one change in the core assumption, a great number of additional changes proved necessary.

One of the most important consequences of discarding the symmetry assumption is that DarkSide-50's LRFs become two dimensional functions, in contrast with Solovov et al.'s one dimensional LRFs. Solovov et al. fit analytic, smooth cubic spline functions to their data to determine their one-dimensional LRFs. There is no analogous process for fitting a smooth and analytic function to data in two dimensions without making more restrictive assumptions not known to be true in DarkSide-50. On the smooth LRFs of Solovov et al., χ^2 minimization can be done using a minimization algorithm like *MINUIT*, whereas with the binned LRFs describe in this work *MINUIT* performs poorly because a small step in x or y does not always change the value of the function.

The single most important change is that the symmetry assumption must be replaced. As described in Section 3.4.1, if there is no assumption requiring specific behavior of the

³Solovov et al. refer to this as “axial symmetry”

LRFs, repeated iterations will make the LRFs worse rather than better. Therefore, as the Solovov et al. method relies on one known feature of the detector optics, so too must the method for DarkSide-50: we know argon-39 data is uniformly distributed, and so build the LRFs from events whose reconstructed positions were pressured to be uniform. This has significant effects on the convergence behavior, as described in Section 3.4.

Solovov et al. include the initial number of photons generated by an event as a minimization parameter in the WLS method. This approach has potential for building in the light collection information gained in producing LRFs to create a position-dependent correction on event energy, which might improve the energy resolution. However, this approach has not yet been tested in DarkSide-50, and is not presented in this dissertation.

3.3.2 Other Prior Work

Morozov et al., a group that includes the first author of Solovov et al., published a paper in May 2013 which describes the ANTS software package that includes a similar iterative approach for position reconstruction [29]. As in Solovov et al., this approach relies on “axial symmetry,” and so cannot be directly applied to the DarkSide-50 detector.

3.4 Iteration, Convergence and Uniformity

This section describes in abstract my approach to creating improved LRFs through iteration. The details of implementing this approach are described in Section 3.6.

3.4.1 Iteration Runaway

Based on the success of Solovov et al., my first approach to position reconstruction was to rely heavily on iteration alone to improve the LRFs from their flawed beginnings to convergence on a final state. I initially believed the role of the symmetry assumption in Solovov et al. was just to produce LRFs that were easier to work with, and that in dispensing with the symmetry

assumption in DarkSide-50, I only needed to manage the computational challenges of working with two-dimensional LRFs. Instead, I discovered the symmetry assumption played a more fundamental role and without a replacement for it, iteration would not improve the LRFs in DarkSide-50.

Figure 3.4 illustrates the behavior of reconstruction of simulated events in an “iteration only” method in which the method applies no direct pressure to improve the LRFs. I constructed initial LRFs using the true positions of the simulated events, and so the initial LRFs are completely accurate with respect to the optical behavior of the simulated detector—there are no flaws in the LRFs stemming from being built from an inaccurate reconstruction⁴. I then reconstructed simulated events over several iterations, each time producing the next generation of LRFs from the previous round’s reconstruction, but without the measure to actively promote LRF improvement discussed in the following sections. When using simulated events, for which true positions are available, the reconstruction error can be measured for each event by finding the distance between the true and reconstructed positions:

$$\Delta\vec{r} = \sqrt{(x_t - x_r)^2 + (y_t - y_r)^2}.$$

My expectation, given my belief that iteration alone would improve LRFs to reach an accurate final state, was that by beginning with completely accurate LRFs built from truth data, the first iteration would have only small statistical errors and subsequent iterations would get no worse. In other words, if iteration makes imperfect LRFs converge to better ones, it should make perfect LRFs mostly stay the same. Instead, as seen in Figure 3.4, the reconstruction error steadily gets worse. Based on tests such as this, I concluded that iteration alone would not improve the LRFs to create a more accurate reconstruction, and something more was required.

⁴Figure 3.4 omits the point at zero iterations with zero error, to keep the other points on scale.

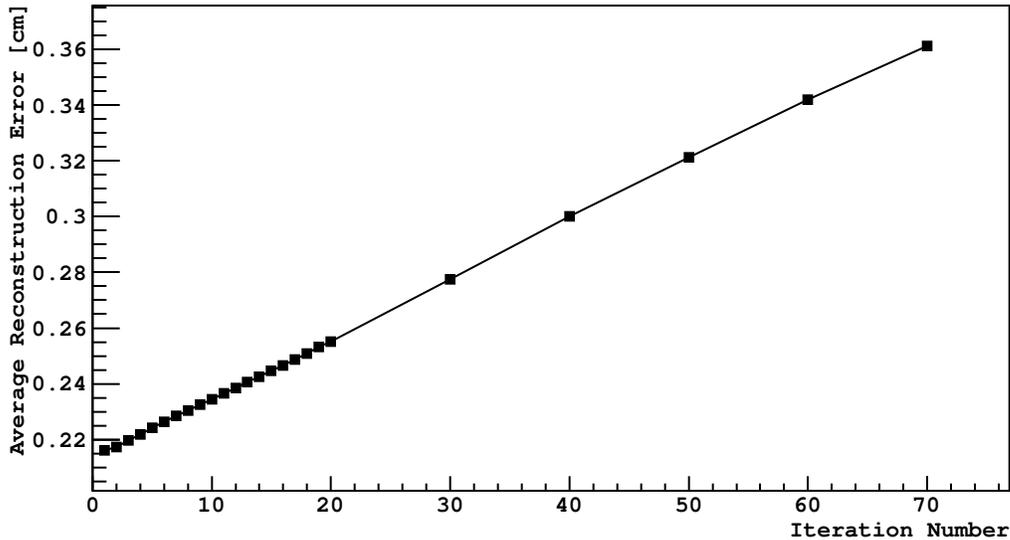


Figure 3.4: Increase in reconstruction error over multiple iterations, without uniformity pressure, Monte Carlo.

3.4.2 Nonuniformity

One consequence of reconstruction error is that the argon-39 background, which in reality is distributed uniformly throughout the detector, tends to be reconstructed nonuniformly⁵. The reconstruction error from iteration runaway can be observed in reconstruction nonuniformity. Figure 3.5 illustrates this, showing the spatial distribution of reconstructed events in iterations 1 and 60 using Monte Carlo data. As seen in Figure 3.4, the error is smallest in the first iteration, which can be seen in the relatively high degree on uniformity. By iteration 60, the nonuniformity is quite severe.

When reconstructing data, the first iteration is expected to have high reconstruction errors, due to discrepancies between the Monte Carlo and data. This is reflected in an observed nonuniformity, as shown in Figure 3.6.

All of the above plots show the nonuniformity created without the benefit of the convergence pressure discussed in the following sections. Using that technique, the iterative

⁵It is possible to have reconstruction error while preserving the uniformity of the original events, for example by swapping the location of pairs of events, but in practice reconstruction error does break uniformity.

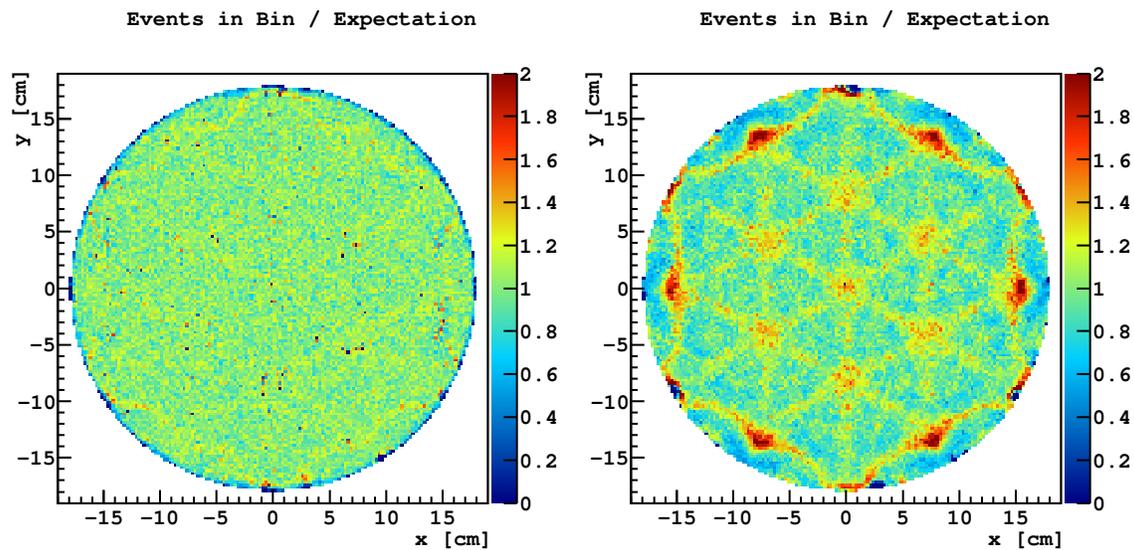


Figure 3.5: Event density by position in iterations 1 (left) and 60 (right) Monte Carlo.

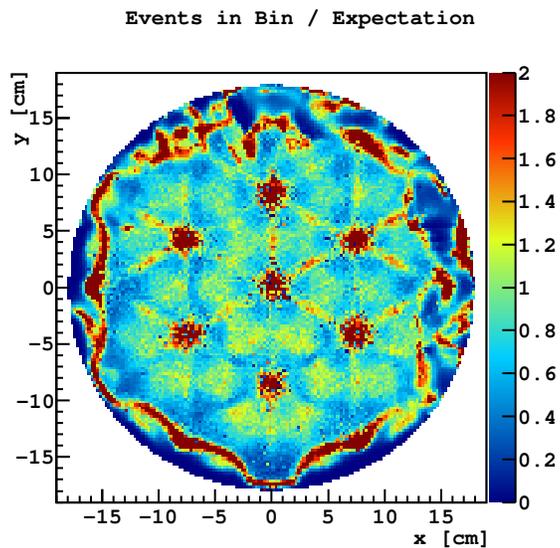


Figure 3.6: Event density by position in iteration 1 in data.

runaway can be slowed, and the uniformity substantially improved. The results of that improvement are discussed in Chapter 4.

3.4.3 Convergence Pressure

If iteration alone does not improve the LRFs, there must be some active part of the construction method that pressures a new LRF to be better than the previous one. However, as discussed earlier, we cannot simply insist the new LRF be more accurate, as we do not, *a priori*, know what an accurate LRF looks like. What is possible, though, is to ensure that LRFs improve on some metric that is a good proxy for accuracy. In particular, if we know some behavior of the detector, we can pressure each new LRF to better match that behavior.

In DarkSide-50, we know one behavior of the detector with near certainty: there is a source of background, argon-39, which is uniformly distributed throughout the detector (see Section 3.4.5). A reconstruction of argon-39 background events should find their positions uniformly distributed, and a good set of LRFs should produce such a reconstruction. Argon-39 is the dominant background in atmospheric argon by a factor of over 100 in the energy range of interest. So, a selection of data events from atmospheric argon will be nearly all argon-39 backgrounds and therefore uniformly distributed.

This remains only a proxy for more accurate LRFs; it is possible an inaccurate set of LRFs will nonetheless reconstruct events uniformly. That said, by pushing for improving uniformity in LRFs over the course of iteration, a set of LRFs with one obvious flaw, nonuniform reconstruction, can be improved to one with no obvious flaws and the potential for other forms of validation.

If the input to the LRF construction is made more uniform, the output of the algorithm will be more uniform—the input being the reconstruction results used to construct the LRFs. The initial LRFs, being flawed, will not reconstruct uniformly. At locations where the flawed LRFs do not match the true LRFs of the detector, few events will be reconstructed, since the real events will rarely match unrealistic LRFs. This creates an overconcentration

of events in other places. If new LRFs were built using this nonuniform reconstruction as input, they would perpetuate the problem. Instead, the nonuniform reconstruction is corrected before being input into the next generation of LRF. Events that would reconstruct at locations of overpopulation are instead input into the next LRF favoring locations without overpopulation.

3.4.4 Remnant Iteration Runaway

Although the LRFs can be induced to converge to a uniform reconstruction over the course of several iterations, repeated iterations still produce some negative effects. The iterative runaway in Section 3.4.1 does not entirely stop. Instead, the increase in reconstruction error is greatly slowed, although it begins at a somewhat higher level in the first iteration. Figure 3.7 demonstrates this remaining behavior.

My hypothesis is that this slow runaway to inaccuracy is the result of multiple possible LRFs satisfying the convergence criterion of uniform reconstruction. Over the course of many iterations, the LRF values random-walk among the possible values consistent with the data and the uniformity condition.

When reconstructing simulated events, the first iteration is based on the true LRFs and so is the most accurate. Successive iterations on simulated events only get worse due to runaway. For data events, the first iteration is believed to be significantly flawed because of discrepancies between the simulated events, which generate the initial LRFs, and data. Iteration improves these flaws by converging towards uniformity, but at the same time the runaway process is slowly increasing some errors. Thus, there is competition between the positive effect of iterations with uniformity pressure and the negative effect of runaway. As the runaway is slow so long as the uniformity pressure is used, the resulting errors can be limited by using only a small number of iterations. As seen in Figure 3.7, 30 iterations can be performed before the runaway error becomes larger than 0.25 cm. As this scale is no larger

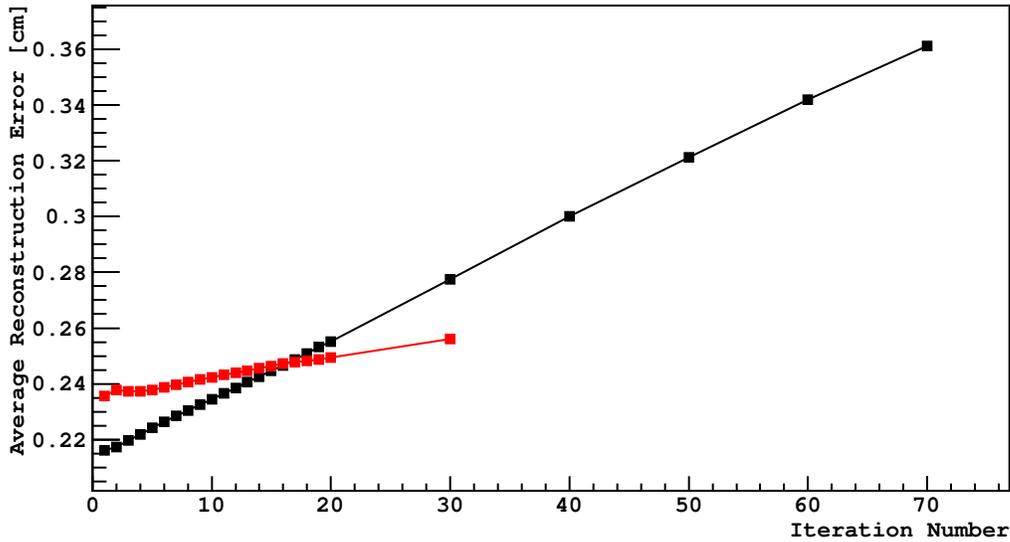


Figure 3.7: Increase in reconstruction error over multiple iterations. In black, without uniformity pressure; in red, using uniformity pressure. Both use Monte Carlo events.

than the binning scale in the LRFs (Section 3.6.1), I consider this an acceptable tradeoff in exchange for the benefits of correcting nonuniform reconstruction.

3.4.5 The Argon-39 Uniformity Assumption

Argon-39 is necessarily uniformly distributed throughout the active argon in the detector. This assumption is not the same as assuming that argon-39-induced S2 signals are uniformly distributed in x and y . While the design of DarkSide-50 includes uniform vertical electric fields that should not change the x and y distribution of S2 signals from the originally uniform energy depositions, it is reasonable to consider the possibility that the physical detector differs from this design goal. Even in this case, we expect only a nonuniform distribution of S2 signals from argon-39 events, which is not the same as nonuniformity of the underlying events. If, for example, the drift field moved all drift electrons from one side of the detector 10% closer to the center of the detector (in x and y), that would have a major impact on the spatial distribution of S2 signals. However, the underlying events would still be uniform. Instead of changing the uniformity assumption, field distortion would change the LRFs. By

changing the relationship between event location and S2 location, field distortion changes the relationship between event location and light collected by the various PMTs. Since this algorithm attempts to find the correct LRFs using the uniformity assumption, it should be able to do so even in the case where field distortion has affected the LRFs.

Despite the robustness of the uniformity assumption, it is possible for distorted drift fields to impede accurate reconstruction. The LRFs are constructed assuming a single set of 38 LRFs is sufficient to describe the optical behavior of all events. If distorted drift fields produce different S2 locations for original events at the same xy , but different z , position, a single set of LRFs will misreconstruct events from z positions where the fields were distorted. I checked for this possibility by examining whether the reconstruction of argon-39 events showed different behavior over the Z-axis. Figure 3.8 shows the results of this study. Each plot shows a set of argon-39 background events on two axes. The vertical position is identified with the drift time, while the position on the x or y axis comes from the reconstruction algorithm. The data is displayed as a two-dimensional histogram (the colored bins), showing qualitatively that the reconstruction does not change substantially at different z positions. The same data is also displayed as a profile histogram (the black bars), which quantitatively shows the mean of the absolute value of the reconstructed position does not change with z .

3.5 Monte Carlo Adjustment

The *g4ds* simulation package (Section 2.9) does not perfectly predict features of light collection observed in the data. One demonstration of this discrepancy can be seen in Figure 3.9, a histogram showing the spectrum of the light fraction ($S2/S2_{total}$) in channel 32, both from Monte Carlo events and data events. Similar discrepancies exist in other channels.

These discrepancies are significant enough that attempting to use Monte Carlo-built LRFs to reconstruct data fails in substantial ways that are hard to recover from in subsequent

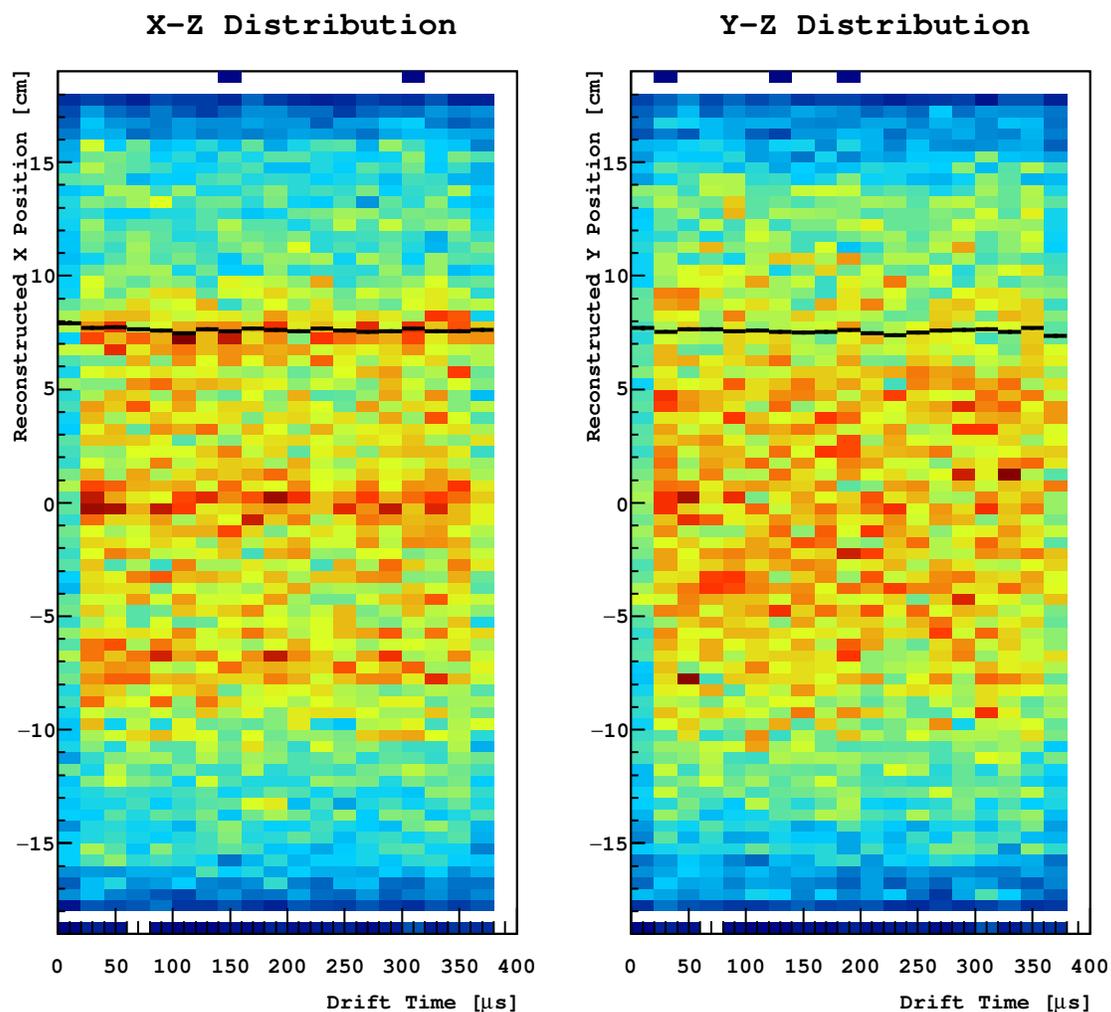


Figure 3.8: Reconstructed data positions vs. the drift time (z -axis position), demonstrating no breakdown in uniformity at different z positions. The black bars show the mean absolute value of the x or y coordinate in each drift time bin.

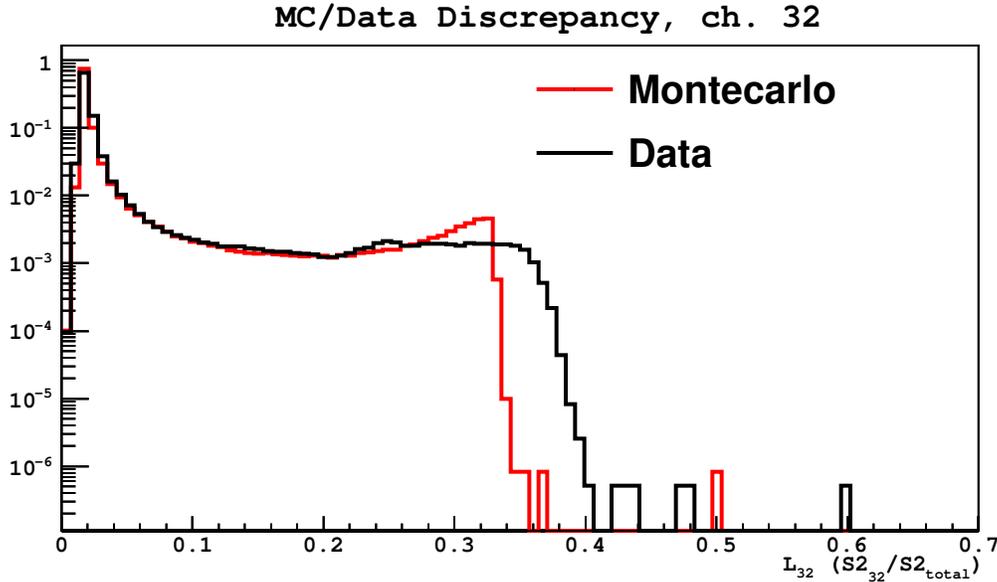


Figure 3.9: Discrepancy between Monte Carlo and data, as measured by fraction of light in channel 32.

iterations. In particular, in many channels the spectrum of data events extends to higher light fraction than is ever seen in the simulation. When reconstructing data with simulated LRFs, these high-light fraction data events do not fit well anywhere. Their best fit is at the point where the simulated events had the highest light fraction, at the center of the PMT, and so all data events with light fraction greater than the simulated maximum pile up at the PMT center. This substantially disrupts uniformity, to the point where many iterations would be required to restore uniformity. Due to the runaway behavior after many iterations it is desirable to produce uniformity in as few iterations as possible.

To avoid this problem, the g4ds output is adjusted before being used as input to construct LRFs. This adjustment is not a “correction” to the Monte Carlo, in that it does not attempt to fix any underlying optical behavior or otherwise guarantee the simulation produces events just like data events. Instead, the adjustment directly changes the light collection numbers output by the simulation, forcing the simulated events to match gross features of the data. Because the final output is adjusted, not the input parameters of the simulation, adjusted events no longer correspond to a simulated series of photon behaviors. Although these

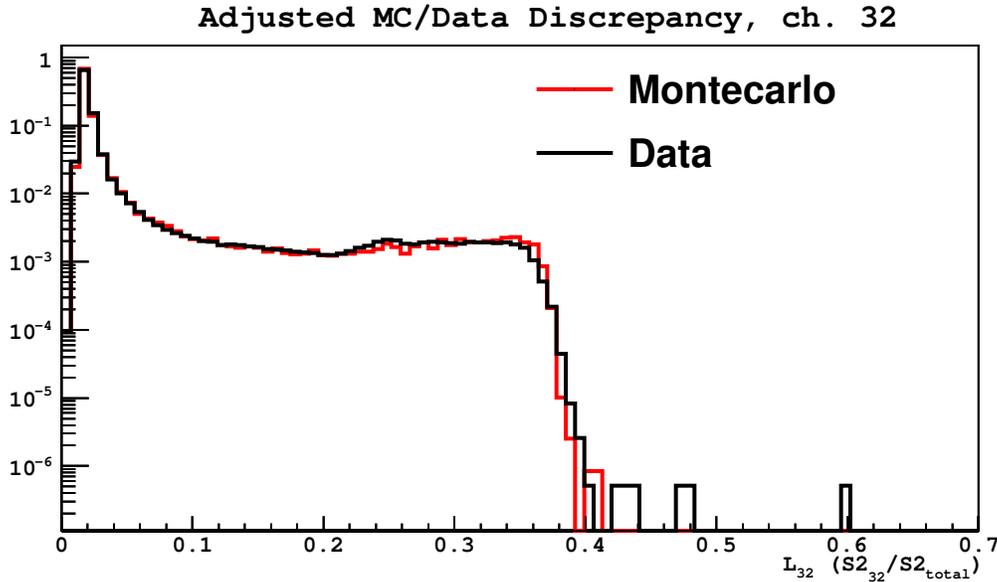


Figure 3.10: Discrepancy between data and Monte Carlo after adjustment, as measured by fraction of light in channel 32.

adjusted events are in this sense less physical than either data or unadjusted Monte Carlo events, their purpose is only to build an initial LRF whose general character matches the data well enough that in no place does the data fail to reconstruct.

The adjustment begins by examining the distribution of M_i/M_{total} (the measured light fraction, see Section 3.1) for each PMT in both simulation and data events. These distributions are divided into 1000 quantiles. For a Monte Carlo event whose light collection fraction in PMT i falls in the n th quantile of the Monte Carlo distribution, the adjustment changes the measured number of photoelectrons in PMT i by the ratio between the n th quantile of the data and the n th quantile of the Monte Carlo.

$$M_{i,\text{adjusted}} = M_{i,\text{original}} \frac{q_{n,\text{data}}}{q_{n,\text{MC}}}$$

The results of the adjustment are demonstrated in Figure 3.10, which, like Figure 3.9, plots the spectrum of light fraction in channel 32. Unlike Figure 3.9, almost no discrepancy remains.

3.6 Implementation

3.6.1 LRF Implementation and Scale

I implement the LRFs as binned functions, as previously mentioned in Section 3.2.3. The bin size is $0.25 \text{ cm} \times 0.25 \text{ cm}$. That is, there is a single value of $L_i(x, y)$ for $0 \leq x, y < 0.25 \text{ cm}$ and similarly a single value of the LRF in each other similar square on the grid.

This sets a fundamental scale for all other processes in the implementation. With the LRF values defined on a 2.5 mm square grid, we also reconstruct events on this grid, evaluate uniformity on this scale, etc.

The scale was set balancing considerations of precision and computation practicality. A smaller scale for the LRF implementation enables higher precision and decreases uncertainties stemming from optical changes over the span of the bin. However, a smaller scale increases statistical uncertainties as the value of each bin is the average of a smaller number of events in that bin. A smaller scale also increases processing time, as there are more possible reconstruction positions to evaluate.

LRFs are implemented using a modified histogram type available in ROOT called a `TProfile2D`. This extends the ordinary 2D histogram behavior to store averages instead of accumulating sums. Events are added into the `TProfile2D` specifying the light fraction and xy coordinates of the event. The `TProfile2D` locates the bin containing those x and y coordinates and adds to an accumulating sum of the light fractions in the bin. It also increments a count of the number of events filled into the bin. When the `TProfile2D` is asked for the value of a bin, it returns the sum of light values divided by the count of events, producing the mean light value of the events in that bin.

It is also possible to apply a weight when filling an event into a bin, allowing the `TProfile2D` to track a weighted mean of the events. Adding an event with some weight increments the appropriate bin's denominator by that weight and increments the light frac-

tion sum by that event's light value times that weight. This capability is used frequently in this method.

A LRF describes the light collection of every location in the detector once at least one event has been filled into each bin in the detector. If a bin in the detector exists with no events in it, its value is set to be an average of nearby bins using the smoothing algorithm discussed in Section 3.2.2 (without the restriction that the bin be far from a PMT). Of course, any bin whose value is informed by only one event is prone to substantial statistical errors compared to a bin containing many events, and so in operation of this method LRFs are always built from sufficient events that every bin has many more than one event in it.

3.6.2 Nominal Detector Radius

The radius of the PTFE reflective cylinder that bounds the active volume of DarkSide-50 is 17.77 cm, while warm. Every part of the DarkSide-50 TPC undergoes thermal contraction when immersed in liquid argon. This thermal contraction is not accounted for in any part of the position reconstruction algorithm. The outer boundary of the detector is defined to be 17.77 cm, essentially redefining the centimeter wherever it appears in the position reconstruction algorithm, both in input and output. This is irrelevant for purposes of surface background rejection, and a small correction to other types of results derived from the position reconstruction. When it is necessary to find results in true centimeters, the nominal centimeters used in this position reconstruction should be multiplied by 0.98.

Not every bin that exists corresponds to a location inside the detector. Some bins intersect the edge of the detector, and so not all of the 0.0625 cm^2 area of the bin lies inside the detector. In total, 16,164 bins lie wholly or partly in the detector.

3.6.3 Location Evaluation

The input of the location evaluation engine specifies the number of S2 photoelectrons⁶ in each PMT in the detector.⁷ This input format can be easily assembled at any stage of the analysis where the S2 values in each channel are known, although it is not a “native” format of other DarkSide-50 analysis tools and so some basic reformatting of information is required.

The first step of the location evaluation is to identify a region of positions in the detector to search for a good match between the input and the LRFs. Evaluating every location in the detector is computationally expensive, and so narrowing the search space saves substantial processing time. The search region is defined to be a square $9.75\text{ cm} \times 9.75\text{ cm}$ centered on the PMT that collected the most S2 light in the event. Simulations indicate that this search region always includes the true position of the event. Figure 3.11 illustrates simulated events for which PMT 29 or 35 collected the most light, plotted by where in the detector they occurred. It also shows boxes indicating the search regions for both PMTs, and a circle showing the boundary of the detector⁸. The events form a hexagon because of the hexagonal packing of the PMTs.

For each position in the search region, the χ^2 metric described in Section 3.1 is evaluated. If $\chi^2 < 90$, that location is stored alongside its χ^2 . This threshold was chosen to reduce the processing time associated with storing (and, later, searching) the positions with very high χ^2 which are poor matches between the input and the LRFs and need not be considered. If no position has $\chi^2 < 60$, the search is done again with a threshold of $\chi^2 < 900$ ⁹.

⁶The number of photoelectrons, not the fraction, although in a real data event the number of photoelectrons is often measured to be a non-integer number as this is measured from pulse integrals, not by actually counting individual photoelectrons.

⁷In code, the format is a `std::map<int, double>`, where the `int` is the PMT ID number and the `double` is the number of photoelectrons measured.

⁸The search region is very tight around the region where events appear. The plotted events all truly fell within this region; they were not forced there by the tightness of the search box.

⁹In a previous version of the algorithm, the second search with the higher threshold was only performed if no position had $\chi^2 < 90$. Some results in this dissertation are based on LRFs constructed with the older version.

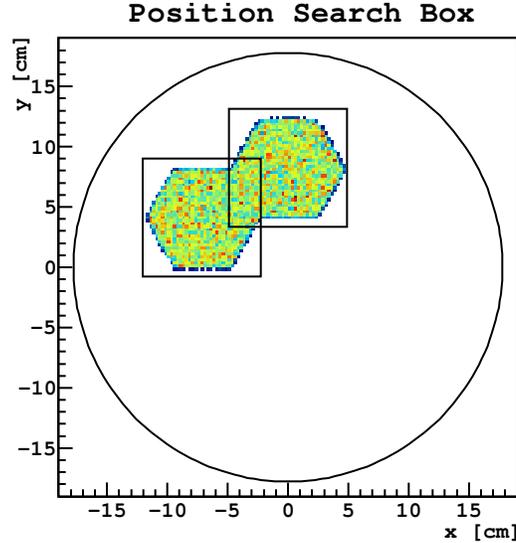


Figure 3.11: The search boxes for events with the most light in PMTs 29 and 35, Monte Carlo.

When the purpose of location evaluation is only to reconstruct the best fit position for an event, the location evaluation engine returns as output that position and the associated best χ^2 . The full list of evaluated positions with χ^2 below the threshold can be used for more purposes, such as LRF building (Section 3.6.5), and the surface background cut (Section 3.7).

3.6.4 LRF Construction from Monte Carlo

This method requires initial LRFs, built using Monte Carlo simulated events. I used a set of 1,215,350 events to construct the initial LRFs. Each event is input into the LRF builder as a list of PMT ID numbers and the number of S2 photoelectrons collected in that PMT. This input is converted to fractions of light collection by dividing each PMT's number of photoelectrons by the total number of photoelectrons. For each of the LRFs (that is, for each PMT), the code then calls the method that adds the event to the LRF, specifying the x and y position and the light fraction in that PMT. ROOT's code then takes over, finding the bin containing that (x, y) coordinate and adjusting the mean value in that bin. Since

this event-addition command is run for each LRF, it is run 38 times in DarkSide-50. Each time, the x and y values are the same, since the event’s true location hasn’t changed, but the S2 light fraction is different because it is the measurement from a different PMT.

When this procedure is performed on an entire set of Monte Carlo events, the value of each LRF in each bin is set at the average light collection fraction of the simulated events that fall in that bin.

3.6.5 LRF Construction from Data

All LRFs beyond the initial set are constructed from data. Once LRFs exist, either the initial Monte Carlo-derived LRFs or later ones built from data, the location evaluation engine can examine each event and evaluate a χ^2 match between that event and each possible location, as described in Section 3.6.3. New LRFs are built from scratch using these events and their reconstructed positions.

The location evaluation engine returns a list of positions and their χ^2 values. Each of these locations is a possible match for the event in question, some more plausible than others. Rather than consider only the best-matched position for the event, the LRFs are updated using each position, weighted by how plausible the match was. Each location in the engine’s output is assigned a weight, equal to the probability of the best fit χ^2 being that location’s χ^2 or greater, given 18 degrees of freedom. The weights are then normalized so the sum of all weights of all locations in the output is one. This ensures that each event has the same total weight, regardless of how many locations had χ^2 below the threshold.

As a result, each event is “spread out” over several locations. The event will influence each LRF the most at the location where the event fits best and will have little to no influence at locations where the event is a poor fit. In other words, from a full set of data events, every bin on the LRFs will be filled with many events. The bin’s value will mostly be determined by the events that best fit there, with some input from events with mediocre fits there.

3.6.6 Uniformity Pressure

As described in Section 3.4, when building LRFs it is important to influence the LRFs to reconstruct the spatially uniform argon-39 background more uniformly than the previous LRFs. This is implemented by changing the weight calculated for each location to favor locations that would otherwise contain fewer reconstructed events.

I perform several “dummy” reconstructions to determine where events would end up if no uniformity pressure was applied. The “dummy” reconstructions are so called because at each of the following steps in this description, the values in the LRFs are forbidden from changing. These dummy reconstructions are used to “look into the future” and check what would happen so the appropriate correction can be made.

After each dummy reconstruction, I look to see at what positions the new LRFs are over- or under-populated—not which LRF bins have the wrong mean value, but which are filled with too many or too few events. From this information, I develop a map of the population in each bin. When the next reconstruction is performed, and new LRFs are made, the weights of each event at each position, originally determined by the χ^2 probability of that position, are adjusted to reduce the weight in positions that were overpopulated and increase it in positions that were underpopulated. The adjustment is the ratio of that bins population over the population if the reconstruction were completely uniform.

Once the weights have been adjusted, an event that had the lowest χ^2 in a bin that was seen to be overpopulated in the dummy reconstruction will still be partially added to that bin, but at a lower weight than would otherwise have happened. Since the weights of each event are normalized, lowering the weight in an overpopulated bin causes the event to be filled more into bins that were not overpopulated. This means that the next generation of LRFs will not be filled with more events in some bins than in others, or at least the population imbalance will be reduced.

After one dummy reconstruction, the map of weight adjustments does not perfectly correct for nonuniformities. A second dummy reconstruction is performed, incorporating

the corrections determined by the first dummy reconstruction. The purpose of this second dummy reconstruction is to analyze what nonuniformities remain after the first set of corrections, and the correction map is updated to reflect these results. A total of seven dummy reconstructions are performed, each one looking for any remaining nonuniformities and adjusting the correction map to reduce those. After those seven dummy reconstructions, a final reconstruction is performed, this time actually building a new generation of LRFs with changed values, incorporating the previous seven rounds of corrections.

Even after seven rounds of refining the corrections, the final real reconstruction may not be completely uniform. No matter how the weights are changed, an event will not be added to LRF bins where it did not have a remotely plausible χ^2 . Creating a more uniform reconstruction eventually requires actually changing the LRF values to change where events find good χ^2 matches. This is why I perform only seven dummy reconstructions before performing a real LRF-changing reconstruction. The next LRFs, made of events placed more uniformly than they were reconstructed by the previous generation, will then reflect that uniformity when used for reconstruction. Thus, the uncorrected reconstruction of the next generation will be more uniform than the previous generation and allow further progress to be made by developing corrections.

As mentioned in Section 3.4.4, the improvements from uniformity pressure compete with runaway behavior that increases reconstruction error. I performed 30 iterations on data with uniformity pressure because 30 iterations on simulated events results in less than 2.5 cm reconstruction error¹⁰.

3.6.7 API

DarkSide-50 collaborators wishing to analyze an event to learn its position will not directly interact with the method for building LRFs. Instead, they will interact with three functions

¹⁰Runaway on data could produce different amounts of error, but that cannot be directly measured.

designed for the purposes of collaborators not intimately involved with *xy* reconstruction development.

1. A function that takes as input an event, and returns the best-fit position and the χ^2 at that position. The input format for the event is as described in Section 3.6.3. The user is responsible for formatting the event S2 measurements to match the design of this function, an uncomplicated process requiring only knowledge of C++'s `std::map` object.
2. A function very similar to #1, which takes an event as input and returns a list of possible reconstruction locations and the χ^2 associated with each location—in essence the same output as function #1 but describing more locations than just the best fit. This function is designed for users who want more information about the event's reconstruction.
3. A function that takes as input the output of function #2, and returns as output the surface cut metric described in Section 3.7.4.

3.6.8 Portability

The current LRFs are stored using ROOT's ability to save an object in memory to disk, and can be retrieved to allow a user to pick up the LRFs and perform position reconstruction from the exact state the *xy* developer left off. The LRF storage method comes bundled with the functions needed to use those LRFs.

The necessary files needed to perform position reconstruction of data are the saved LRF container file `xylocations.root` and the C++ code files defining the behavior of the LRF storage class and the other types stored by that class:

`xylocator.hh`, `xylocator.C`

`xy_pmtdata.hh`, `xy_pmtdata.C`

`xy_pmtdata_map.hh`, `xy_pmtdata_map.C`

The only required software libraries are the Standard Library and ROOT (tested with v5.34/12 and v5.34/23).

Building LRFs requires two more files:

`xy_iterator.hh`, `xy_iterator.C`

3.7 The Surface Background Cut

3.7.1 Surface Backgrounds

As described in Section 2.7.3, DarkSide-50 must use xy position reconstruction to remove backgrounds caused by alpha emitters on the surface. These surface backgrounds deposit energy very close to the surface, within a few dozen micrometers, so the goal of an algorithmic surface background cut is to efficiently identify and remove events from that narrow region near the surface while removing as few as possible events that were not truly that close to the surface. Given the goals of DarkSide-50, efficiently removing surface backgrounds is a higher priority than preserving non-surface events: we want no surface backgrounds to survive the cut and be mistaken for WIMPs, and will accept some substantial loss of acceptance to achieve that goal, if necessary.

3.7.2 Fiducial Cuts

The typical approach to surface background in low-background detectors is to “fiducialize” the active volume, designating some fiducial, that is, trusted, region some distance from the surfaces. Although surface backgrounds only truly occur very close to the surfaces, the fiducial cut may have to exclude all volume within a centimeter or more of the active volume, as the position reconstruction cannot reliably distinguish surface events from events some distance from the surface. In trying to label every region as either completely trusted

or completely cut, this sort of cut can end up with a very unfavorable tradeoff between background rejection and acceptance of the desired events. Tests with simulated data show a fiducial cut of the outer 1.5 cm of radius has a cut efficiency of only 99% while sacrificing 16% of the acceptance. The more sophisticated approach described in the following sections can achieve better than 99.95% background rejection with the same acceptance loss (see Section 4.5).

3.7.3 An Alternative to Fiducialization

I approached designing an alternative to a fiducial cut by considering these principles:

1. Only events that are truly on the surface are a concern. An event that is certainly even a few millimeters from the surface does not need to be cut.
2. Events that are certainly on the surface should be cut
3. Many events are not certainly on the surface, but are plausibly on the surface. They may have decent χ^2 fits at several locations, including surface locations. Depending on how aggressive the surface background cut must be to meet the goals of the detector, the cut can reject events that are very plausibly on the surface or only slightly plausibly on the surface.

3.7.4 Surface Cut Implementation

As discussed in Section 3.6.3, the location evaluation engine provides for each event a list of evaluated positions and the associated χ^2 . From each χ^2 , I can calculate a probability according to the χ^2 distribution for 18 degrees of freedom, as described in Section 3.1.1.

The algorithm finds the largest probability in any position and the largest probability among the bins on the surface. A bin is considered to be on the surface if it or any adjacent bin does not lie entirely inside the 17.77 cm nominal radius of the detector. Equivalently, the “edge bins” include all bins that actually intersect the boundary of the detector, and so

are not whole $0.25\text{ cm} \times 0.25\text{ cm}$ squares, as well as the outermost bins that are wholly in the detector. This ensures that we look for surface events in the smallest possible area near the surface. To reduce this area would require either smaller bins, or relying entirely on the partially cut-off bins to reconstruct accurately.

I define the “surface cut metric”:

$$SCM = P(\chi_{\text{diff}}^2),$$

where P is a S-shaped curve for which $P(0) = 1$, $P(18) = 0.46$, and $P(\gg 18) \approx 0^{11}$.

$$\chi_{\text{diff}}^2 = \chi_{\text{edge}}^2 - \chi_{\text{best}}^2,$$

where χ_{best}^2 is the lowest χ^2 found at any location and χ_{edge}^2 is the lowest χ^2 found at any surface location. In cases where the best fit is at the edge, $SCM = 1$. In cases where no edge positions were considered, for example an event with the most light in the center tube, for which the position search box does not include the edge (Section 3.6.3), $SCM = 0$. Likewise, if no edge position had $\chi^2 < 90$ (or < 900 and no other position had $\chi^2 < 90$), $SCM = 0$ because no edge position was included in the vector of save positions. In all other cases, $0 < SCM < 1$, with cases where edge locations had a better fit having higher SCM .

The choice to base the SCM on the difference of the two χ^2 was made to account for a common “failure mode” I observed causing some simulated events to have a very low SCM despite being close to the edge. For these events, χ_{best}^2 was substantially larger than 18, the number of degrees of freedom. It appeared χ^2 was being offset by a certain amount, likely due to misestimation of the errors (Section 4.2.1). χ_{edge}^2 was often larger than χ_{best}^2 by a relatively small amount. I assumed the offset in χ_{best}^2 should be subtracted to properly estimate how “bad” χ_{edge}^2 was, and so how plausible it was that the event came from the surface. This is

¹¹ P is in fact the probability function for a standard χ^2 variable with 18 degrees of freedom, but χ_{diff}^2 is not a standard χ^2 variable and SCM is not a probability. Here, P is just a convenient function to map values of χ_{diff}^2 to the range $[0, 1]$.

an “off warranty” usage of the χ^2 distribution, in the sense that the mathematically proven properties of the χ^2 distribution do not include guarantees about the probability distribution of the difference of two χ^2 variables with an assumed offset from misestimated errors. Thus, the *SCM* cannot be treated as the actual probability an event was a surface background; instead it is only a metric that summarizes the χ^2 behavior of the event on a scale of 0 to 1.

Surface backgrounds can be cut by rejecting events with *SCM* above some threshold. The higher the threshold, the fewer events will be rejected, improving overall acceptance but reducing the efficiency of rejecting surface backgrounds. I make no claim at this time of the best cut threshold, which must be decided based on the overall mission goals of DarkSide-50, but the results in Section 4.5 illustrate the performance of different cut thresholds. Lower photon statistics reduce the performance of the cut, so it may be desirable to set the cut threshold differently for different event energies.

Chapter 4

xy Reconstruction Results

4.1 Event Sets Used

I used a few sets of events to produce all the results reported in this chapter.

4.1.1 Simulated High-S2 Events

I used the DarkSide-50 Monte Carlo simulations package *g4ds* (Section 2.9) to produce 1,215,350 events with an average of 316,000 photoelectrons per event, corresponding to roughly 890 keV electron recoils in the center of the detector. The events were uniformly distributed throughout the detector. They were generated by starting monoenergetic¹ electrons in the liquid argon.

4.1.2 Simulated Low-S2 Events

Using *g4ds*, I produced a second set of simulated events with fewer photoelectrons per event. This set contains 10^6 events, with an average of 17,250 photoelectrons detected per event, corresponding to about 50 keV electron recoils in the center of the detector. These events were also distributed uniformly throughout the detector. Although referred to as “low-S2”

¹200 keV. Producing 300,000 photons from a 200 keV electron is not representative of the S2 yield of the real detector, but in simulation the process determining the S2 yield had not yet been tuned.

here, these events still have substantially more S2 than even the highest-S2 WIMP signal included in the search, at 4,600 photoelectrons.

4.1.3 Data Events

I used DarkSide-50’s analysis tools to select argon-39 background events from data, assembling a set of 1,945,880 events with an average of 20,924 photoelectrons per event.

The events were selected from DarkSide-50’s atmospheric argon campaign. Much of that campaign was performed with trigger logic known as the “G2” trigger that excluded high-energy argon-39 backgrounds; the data selected here was from runs without the G2 trigger. This campaign was conducted with atmospheric argon in DarkSide-50’s TPC, resulting in a large number of argon-39 decay signals. This high rate of argon-39 events dominates the data collected.

I selected events passing quality and energy cuts to find argon-39 backgrounds with relatively high energy but without digitizer saturation. To apply the position reconstruction algorithm, the S2 signal of the event must be identified correctly. Mistakenly identifying a scintillation signal, a noise signal, or only a piece of a signal as the S2 signal will make position reconstruction impossible. The cuts applied were:

- The event’s baseline was found. Events failing this cut had such significant noise that no pre-trigger baseline could be identified.
- The event contains exactly two pulses, indicating an S1 and S2 signal, or three pulses if the third pulse’s timing indicate an S3 signal². Events failing this cut may contain more than one particle interaction³, or may be failures of the pulse finding algorithm.
- The first pulse arrival time fell between -0.25 and $-0.15 \mu\text{s}$, the expected window for events in the examined runs. This cut helps ensure that our identification of the

²S3 is a third light signal, much smaller than S2, caused by electrons related by the TPC’s cathode when struck by S2 light.

³Either two particles interacting in the same trigger window, or one particle interacting twice

first pulse as the scintillation signal and the second pulse as the drift electron signal is correct. Events fail this cut typically because the first pulse is not the S1 of an event that just occurred but, for example, the tail of the previous event.

- The first pulse has at least 10% of the pulse area in the first 90 ns of the pulse. Events that fail this pulse are not peaked in a manner at all consistent with scintillation signals. This further helps ensure the first pulse is S1, which in turn helps ensure the second pulse is S2.
- Less than 40% of the total S1 is seen by a single PMT. Events failing this cut are likely Cherenkov background events in which light is created in the TPC windows or the PMT windows.
- The event did not saturate the digitizer at any point. Events whose signal amplitude exceeds the dynamic range of the digitizer can have a greatly distorted distribution of signal among the PMTs, as saturation occurs on a PMT-by-PMT basis. That is, the signal in the PMT that saw the most light could saturate, limiting the recorded signal to be less than the true light reaching that PMT, while other PMTs that saw less light would not saturate and record the full light they saw. Since the distribution of light among the PMTs is the source of information for position reconstruction, saturation absolutely ruins the possibility of accurate position reconstruction. The energy cut discussed below nearly entirely eliminates saturating events on its own; only one event (out of 1.9 million) is removed by the saturation cut that would not otherwise be removed by the energy cut.
- An “energy” cut is made by accepting only events with S1 pulse area⁴ between 550 and 750 photoelectrons. S1 pulse area is the basic measurement of event energy in DarkSide-

⁴Corrected for z -axis effects using a standard technique for DarkSide-50

50⁵. The range of 550 to 750 corresponds with an energy range of 79–107 keV electron recoils. The upper end of this range was chosen by choosing the maximum S1 that would never saturate in either S1 or S2. The lower end of the range was chosen to include a large number of data events.

The above cuts stem from the same motivations as the official DarkSide-50 WIMP search cuts described in Ref. [25], but differ in implementation. As one example, [25] does not explicitly cut events that saturate during the S2 pulse, but does cut events based on the pulse shape of the second pulse. Differences between the event selection cuts presented here and in [25] reflect changes in the state-of-the-art of DarkSide-50 analysis cuts over time.

4.2 Characteristics of Data and Simulated Events

4.2.1 Effects of Binning on Uncertainty Estimation

As described in Section 3.1, the uncertainty factor σ_{M_i} is estimated by assuming the distribution of measured S2 photoelectrons (M_i) from events at a single point is Poissonian, thus $\sigma_{M_i} = \sqrt{M_i}$. The Poissonian assumption holds well for statistical uncertainty, but the binning of the LRF creates systematic uncertainties in the difference $M_i - L_i M_{\text{tot}}$. In any bin, the underlying optics change somewhat over the width of the bin, and so some events in the bin experience slightly different light collection than the average, even without statistical fluctuations. This effect is most pronounced when the LRF is changing over a short distance, which happens 3.8 cm from the center of each PMT, i.e., at the edge of the PMT. When the edge of the PMT is close to the edge of the detector, which is true for PMTs 20, 22, 27, 32, 34 and 36, the effect can be quite sharp.

Figure 4.1 illustrates this effect. Using simulated events, I can select for each bin the events truly located in that bin, and find for that bin the actual standard deviation of the

⁵A more precise measurement of energy can be made by considering the anticorrelation of S1 and S2 measurements, but the necessary information to use this technique was not available in time to apply to the selection of these data events.

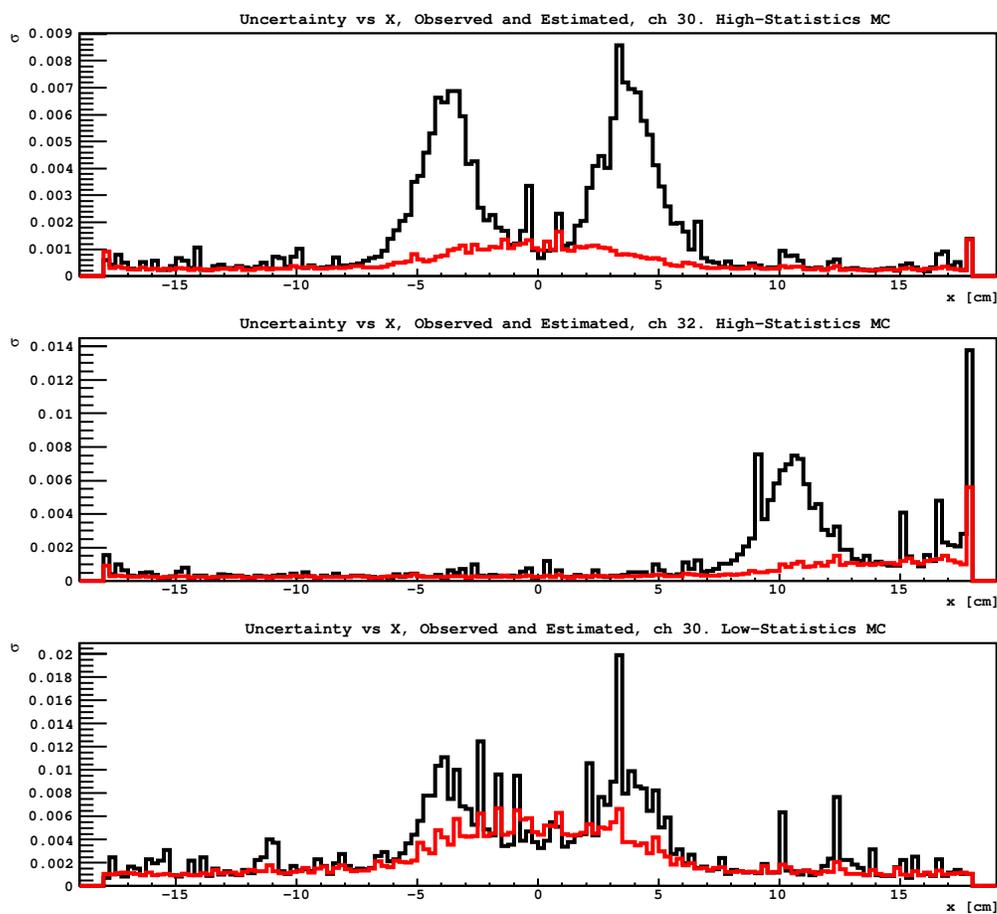


Figure 4.1: The standard deviation of measured light in each bin (black) and the estimated statistical uncertainty of that measurement (red), plotted along the x -axis for PMTs 30 and 32. Monte Carlo, showing both high- and low-S2.

S2 light in a given PMT, which is plotted in black. I can also find the average uncertainty for estimated by $\sigma_{M_i} = \sqrt{M_i}$ for events in that bin, plotted in red. At most positions, the estimates track the measured standard deviation quite well, including at $x = 0$ cm for PMT 30, which is directly under the center of that PMT. In a band around the edge of each PMT, however, the observed standard deviation is substantially higher than the estimated uncertainty. As seen in the third plot, made using the low-S2 Monte Carlo events, the excess standard deviation is relatively smaller when the statistical errors are increased by the reduced S2 statistics. That makes the estimate more accurate on average, although with the reduced S2 fluctuations in the observed standard deviation are more common.

I tested several approaches to a position-dependent uncertainty term that would account for this issue. Unfortunately, all these efforts fell to the pitfall that when the uncertainty term is dependent on position, the best χ^2 might be wherever the uncertainty denominator is highest rather than where the difference term in the numerator is lowest.

The issue therefore remains that events in some locations overestimate the weight of some terms of χ^2 by underestimating the uncertainty. In many cases, the effect of this is limited to moving the best χ^2 over by a single bin. Because an event near the edge of the bin may differ from the LRF's value in that bin even before statistical fluctuations, a very small fluctuation can push it "over the edge" into the next bin. However, an event cannot be pushed more than one bin away by this effect, as the systematic effect of being close to the edge of one bin makes the event closely resemble the next bin over, not bins further away.

A hidden negative effect, however, is that for events that do mis-reconstruct, the χ^2 in the true bin can be unreasonably high. Typically, we expect when an event has statistical fluctuations that move it to another bin, that event will have a χ^2 of order 1 less in the wrong bin than in the true bin, because a fluctuation of one standard deviation in one PMT should change χ^2 by 1. If an event moves due to systematic uncertainty, however, the χ^2 in its true bin can be very large (>50) while the χ^2 in the reconstructed bin has a more typical

value for a good fit ($\sim 18 - 22$). This can give a false sense of confidence that the event is not from the true bin, since the χ^2 is so large.

4.2.2 Indications of Detector Asymmetry in Data

In considering whether to apply symmetry assumptions to the xy reconstruction algorithm, I studied whether the data indicated the detector was symmetric. I performed several studies of the sort illustrated in Figure 4.2. I plotted the S2 fraction in one PMT versus two of its neighbors—in the example illustrated, I plotted PMT 37 against the PMTs to the left (36) and right (34). These plots are made without any use of the position reconstruction algorithm and only show the recorded features of the events.

In the symmetric design of DarkSide-50, the left and right neighbors of PMT 37 should behave identically. Even accounting for the different quantum efficiencies of the two PMTs, the existence of a “notch” in the left plot of Figure 4.2 that does not exist in the right plot is starkly contradictory with the symmetric design. There must exist some optical effect to the left of PMT 37 that does not exist to the right. Since this study was done without any xy reconstruction, the effect cannot be due to errors in the reconstruction algorithm. It is a real feature of the detector optics. Similar effects can be seen with the symmetry neighbors 32/34 and 23/33. There may be several asymmetric optical features in the detector or just one whose effects are seen in multiple PMTs, but the evidence in data is clear that the detector is not as optically symmetric as designed.

4.3 xy Reconstruction Performance Results

In this section, I discuss results demonstrating the ability of this xy reconstruction method to perform as designed. Evidence that this design works—that is, that it enables accurate position reconstruction in the real DarkSide-50 detector—is discussed in Section 4.4.

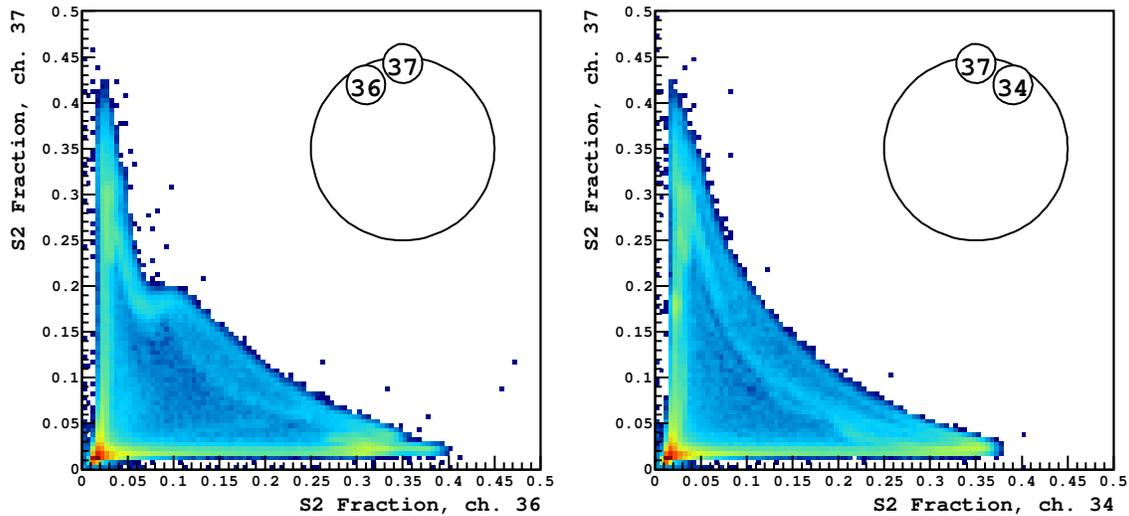


Figure 4.2: Asymmetric data. These plots show the S2 fraction in PMT 37 versus the fraction in one of PMT 37’s neighbors, either PMT 36 or PMT 34. PMTs 36 and 34 are symmetrically located on either side of PMT 37.

4.3.1 Tests on Simulation

Simulated events provide a useful test that this method behaves as desired in at least one “universe.” While simulated events do not precisely mirror real events, they describe a plausible, if fictional, universe in which the detector behaves

as laid out in the simulation configuration. I also use simulated events with very high photoelectron statistics, minimizing statistical fluctuations that could cause an event to look like it belongs somewhere else. Results from reconstruction of simulated events demonstrate the achievable resolution of the WLS method.

When using simulated events, for which true positions are available, the reconstruction error can be measured for each event by finding the distance between the true and reconstructed positions: $\Delta\vec{r} = \sqrt{(x_t - x_r)^2 + (y_t - y_r)^2}$. I calculated the reconstruction error for high-photoelectron-statistics simulated events, reconstructed using LRFs built using the true positions of those events, without iteration.

Figure 4.3 shows the average reconstruction error for events whose true position lies in each bin. It shows a distinctive pattern of vertical stripes, a result of the location evaluation

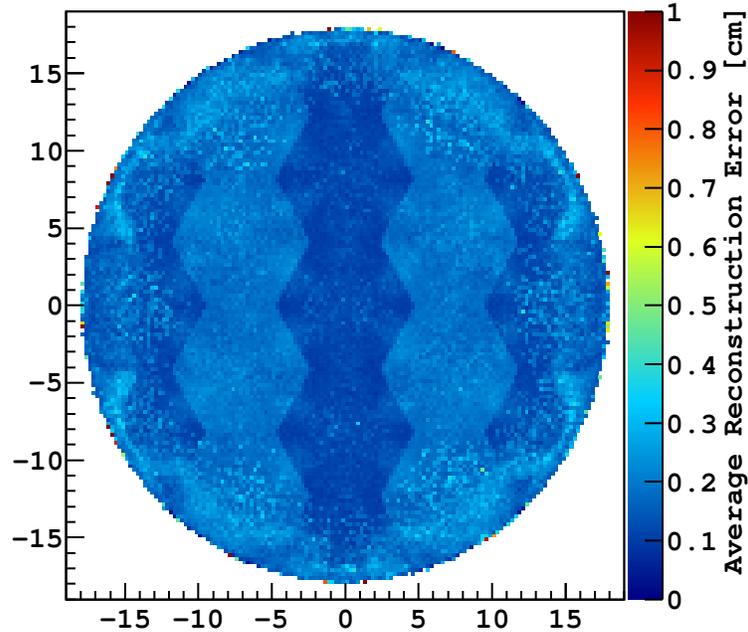


Figure 4.3: Average reconstruction error in each bin, high-S2 Monte Carlo, reconstructed from true LRFs.

engine. The location evaluation engine searches on a grid centered on the PMT, while events are binned on a grid centered on the detector. The PMTs in the lighter blue stripes have a somewhat larger offset between these two grids, which increases the reconstruction error slightly. This effect is at the 0.5 mm level, and so negligible. Figure 4.4 shows the average reconstruction error for events at different radii in the detector.

The average reconstruction error is very small, less than the bin spacing, indicating the majority of events are reconstructed in the same bin they truly belong in. This matches expectations that when statistical errors are limited by high photoelectron statistics and the LRFs are absolutely correct, reconstruction accuracy should be high.

When lower statistics simulated events are used, the reconstruction error is larger. Figures 4.5 and 4.6 show the reconstruction error by bin and by radius for the low-statistics simulated events, again reconstructed with LRFs built from true positions.

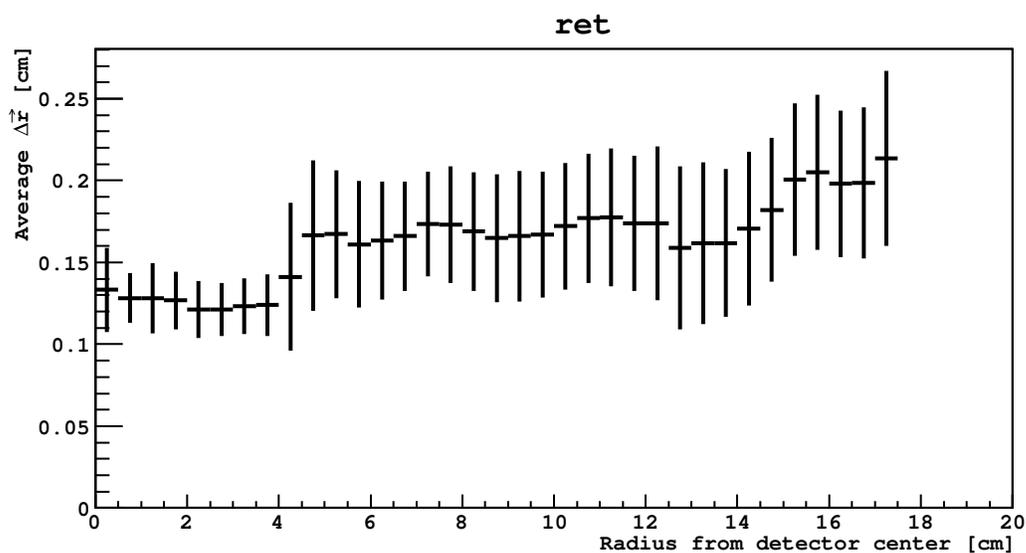


Figure 4.4: Average reconstruction error by radius, high-S2 Monte Carlo, reconstructed from true LRFs. For each bin, the horizontal bar shows the mean and the the vertical bar indicates the standard deviation of reconstruction error over the bins at that radius.

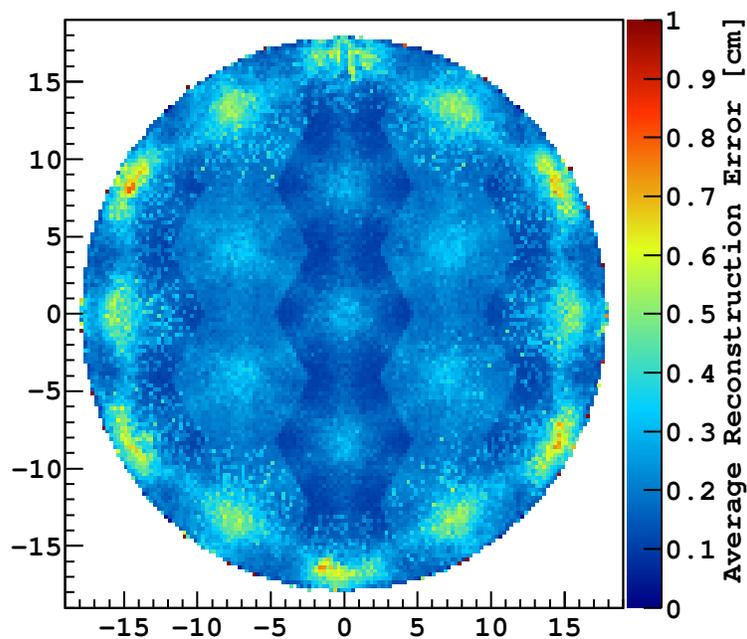


Figure 4.5: Average reconstruction error in each bin, low-S2 Monte Carlo, reconstructed from true LRFs.

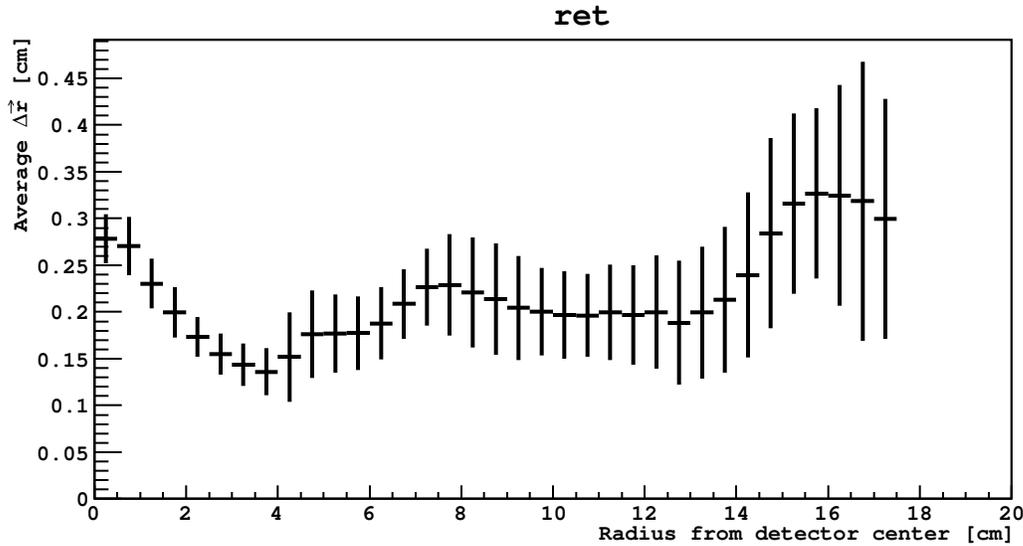


Figure 4.6: Average reconstruction error by radius, low-S2 Monte Carlo, reconstructed from true LRFs. For each bin, the horizontal bar shows the mean and the vertical bar indicates the standard deviation of reconstruction error over the bins at that radius.

4.3.2 Uniformity

As explained in Section 3.4, this method is designed to ensure that after several iterations the light response functions produce an increasingly uniform reconstruction of the data. I built LRFs using the argon-39 background data events, and iterated 30 times. I then analyzed the spacial uniformity of the reconstruction. I present in this section results both with uniformity pressure, in which the reconstructed locations were influenced to be more uniform, and after uniformity pressure, in which the reconstructed locations were not directly influenced to be more uniform, but were found using LRFs created using uniformity pressure. That second case, reconstruction after uniformity pressure, is the procedure that will be used to reconstruct individual data events in the WIMP search.

Figure 4.7 shows each bin of the detector and the density of events reconstructed in that bin using the initial (0th iteration) LRFs and using uniformity pressure. The density is given in terms of the number of events expected in each bin in a completely uniform reconstruction. Near the edge, there is significant nonuniform reconstruction. This is an indication that the

real light collection from events in these areas differs substantially from the simulated light collection that determined the initial Monte Carlo-built LRFs.

By the fifth iterative reconstruction, using the 4th iteration LRFs (Figure 4.8), the uniformity of the reconstruction is substantially improved. Notable “hot spots” remain near $(12, 5)$ and $(-3, 12)$. These spots have been a consistent problem in many versions of the algorithm, and also correspond to known asymmetries in the data (Section 4.2.2). I believe there are physical features in the detector at these locations that cause the optics for events near there to be substantially different than elsewhere in the detector.

Figure 4.9 shows 30th reconstruction with uniformity pressure, using the 29th-iteration LRFs. After this many iterations, the uniformity is improved to the point where the hotspots are just barely visible and only the one at $(-3, 12)$ contains a bigger excess. There is still a 20% excess of events in the center of most PMTs and a 20% deficit in the gaps between PMTs. The nature of this structure remains a mystery. Although I expect some nonuniformity to remain in any finite number of iterations, I have not been able to discover why events specifically pile up in the PMT centers.

Figure 4.10 shows the reconstruction uniformity without direct uniformity pressure using the 30th-iteration LRFs (which were constructed using uniformity pressure). The region around the edge shows some uniformity at the bin-by-bin scale, but no hot spots on the scale seen in Figure 4.7. Figure 4.11 shows the fraction of total events reconstructed at each radius. As the area encompassed by a radius bin scales linearly with radius, we expect a linear increase in events at each radius if they are reconstructed uniformly, and the red line on figure 4.11 shows this expectation.

Figure 4.12 projects Figure 4.9’s color axis, showing the distribution of bins with various excesses/deficits of events. The mean of this distribution is by definition 1, and the standard deviation is 0.218.

Iterated reconstruction of simulated events also shows nonuniformity with similar structure to what is seen in figure 4.10. In the case of simulated events, the nonuniformity gets

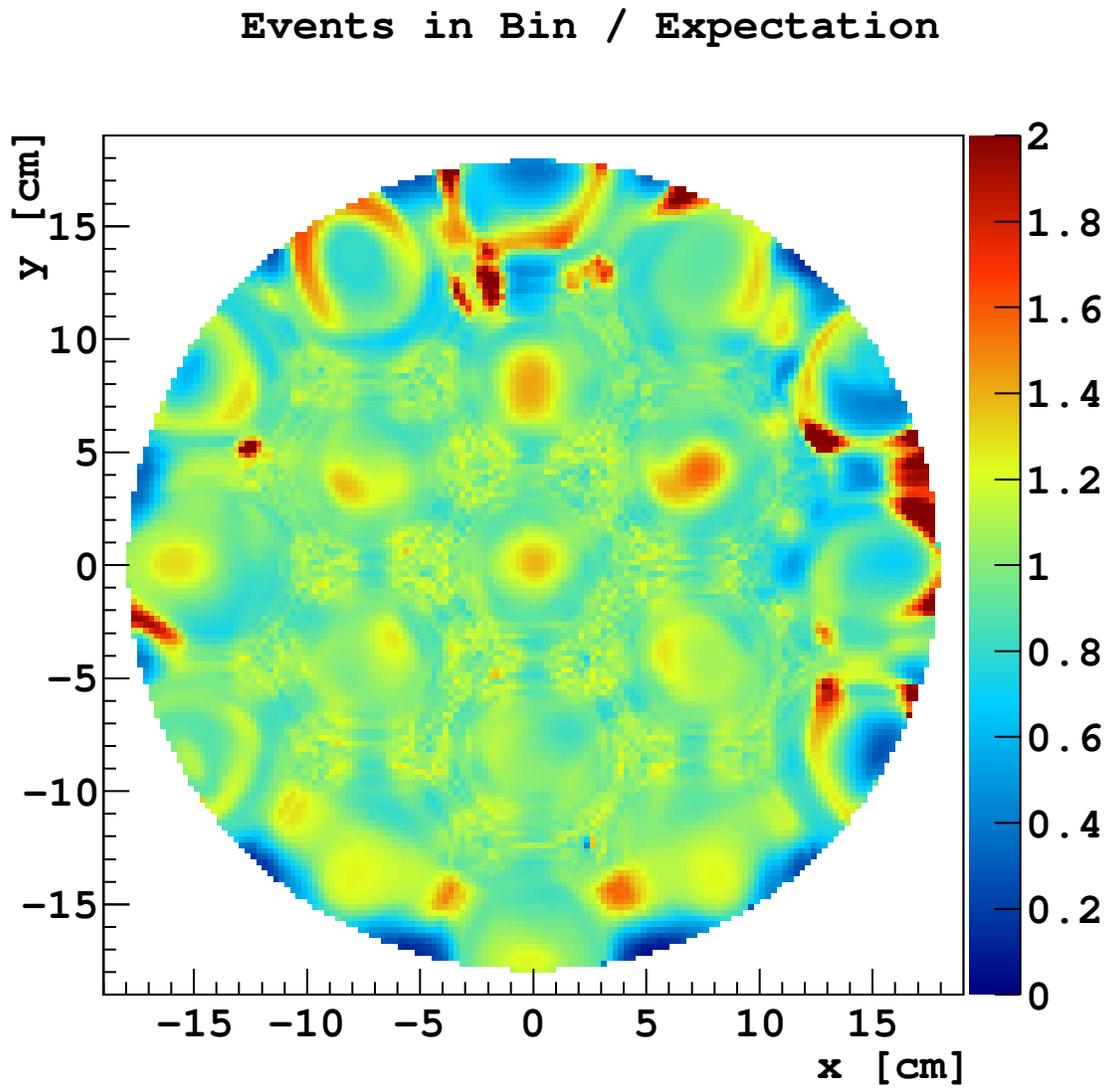


Figure 4.7: Event density by position, 0th iteration LRFs, with uniformity pressure.

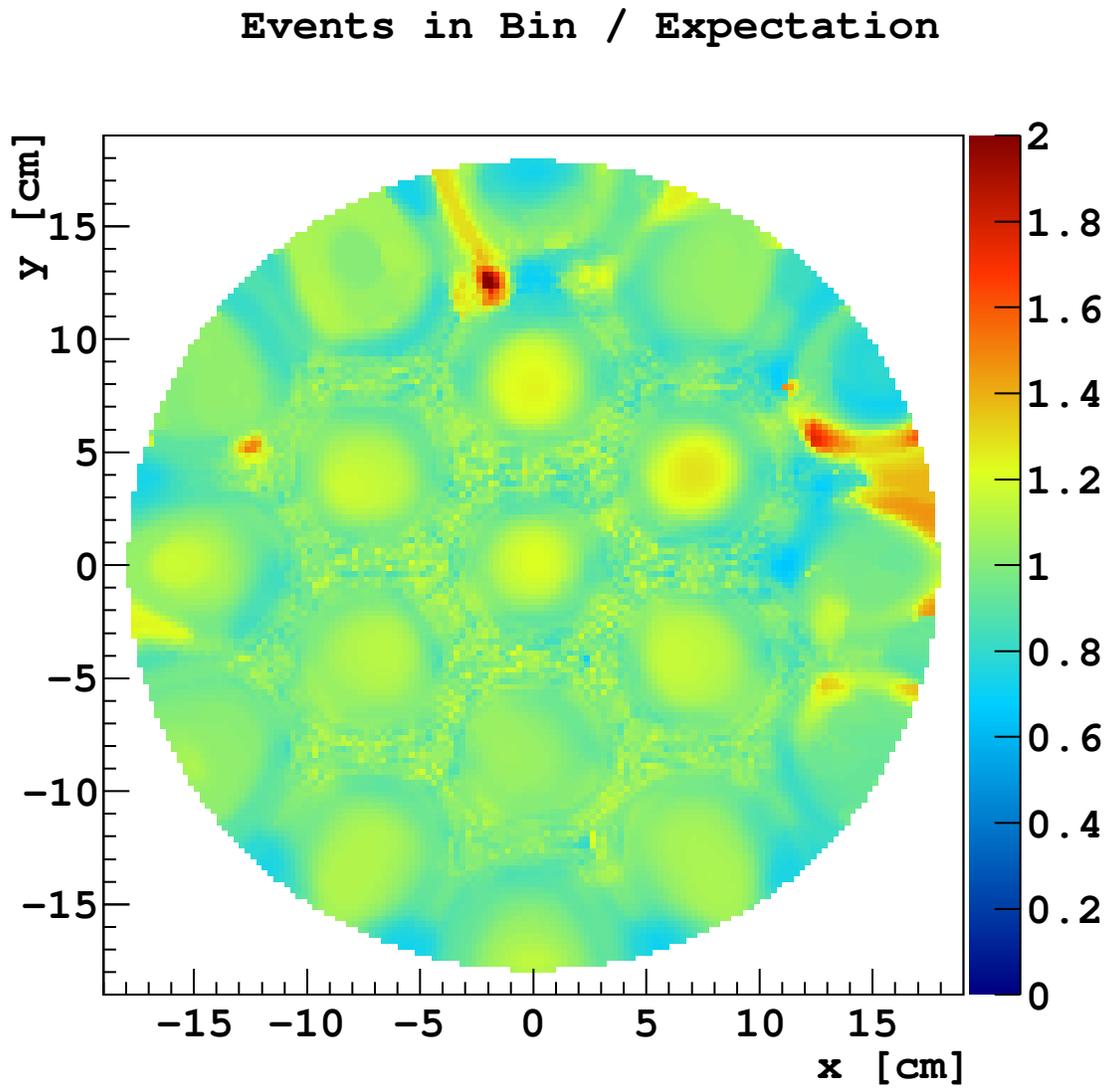


Figure 4.8: Event density by position, 4th iteration LRFs, with uniformity pressure.

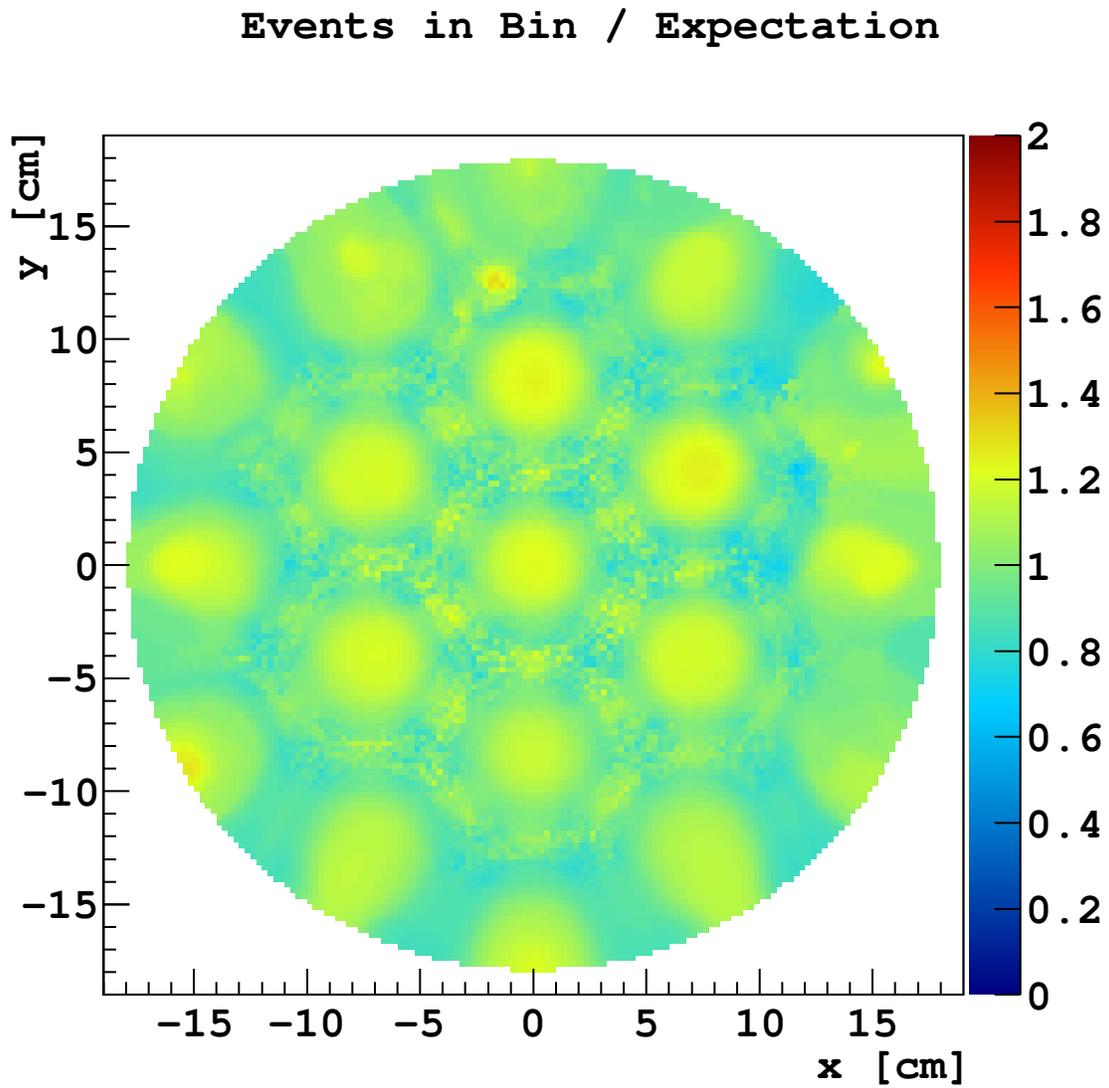


Figure 4.9: Event density by position, 29th iteration LRFs, with uniformity pressure.

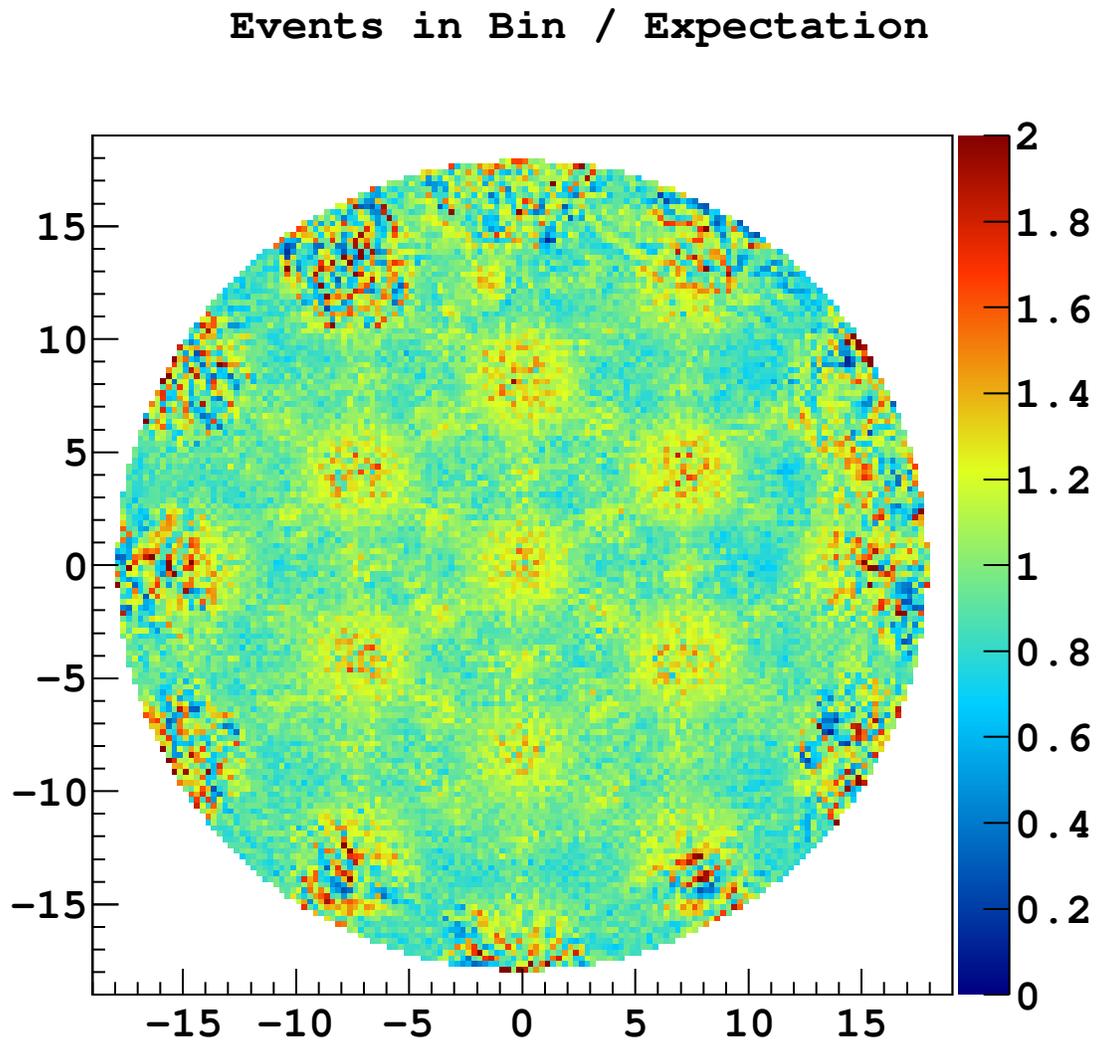


Figure 4.10: Event density by position, 30th iteration LRFs, without uniformity pressure.

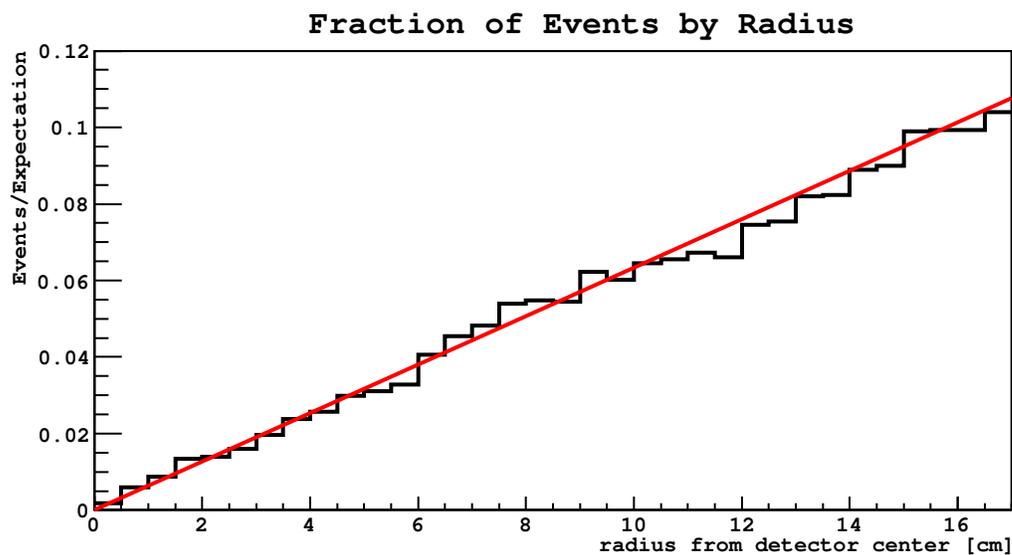


Figure 4.11: Fraction of Events by Radius, 30th iteration LRFs, without uniformity pressure.

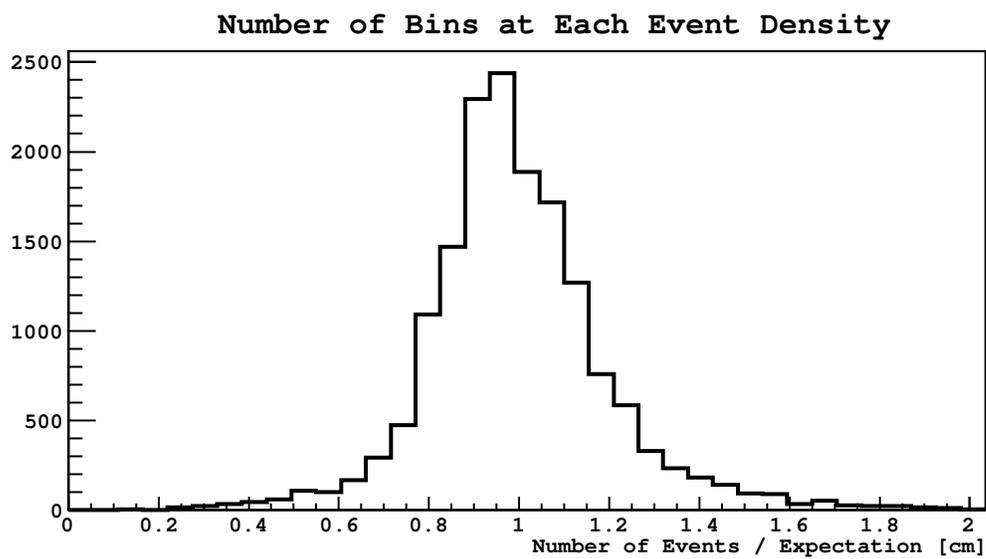


Figure 4.12: Number of Bins at Each Event Density, 30th iteration LRFs, without uniformity pressure. mean= 1, $\sigma = 0.218$.

worse over repeated iterations, due to iterative runaway. I compared reconstructions of simulated events to reconstructions of data by comparing the standard deviation of the analog of Figure 4.12. The reconstruction of data with the 30th-iteration LRFs had a standard deviation of 0.218, and I found the reconstruction of low-S2 simulated events had a similar standard deviation of 0.219. Figure 4.13 shows the nonuniformity of this reconstruction of simulated events. This iteration has an average reconstruction error of 0.44 cm, which implies that the data reconstruction from the 30th iteration LRFs may have similar errors stemming purely from the nonuniformity. There are other indications that the data has somewhat larger reconstruction errors than this, discussed in the following sections. However, the nonuniformity observed in data does not by itself imply reconstruction errors larger than 0.44 cm.

4.4 *xy* Reconstruction Validation Results

4.4.1 Coincident Chain Decays

The “BiPo” coincident chain decays offers a helpful method for validating the precision of the xylocator method. The “BiPo” decays are the radioactive decays of bismuth-214 and polonium-214, part of the upper half of the radon-226 decay chain. Bismuth-214 beta decays into polonium 214, which alpha decays a short time later. The half life of polonium-214 is only $164.3 \mu\text{s}$, which means that most polonium-214 decays occur before the polonium could have moved any significant distance from the decay site of the bismuth-214. Since neither the emitted beta nor the alpha travels far before depositing its energy in the argon, the energy deposits from the two decays in the BiPo occur very nearly at the same position. If the method is precise, it should reconstruct both decays at the same location. This method does not, however, indicate anything about the accuracy of the method, as there is no information available about the true position of either decay.

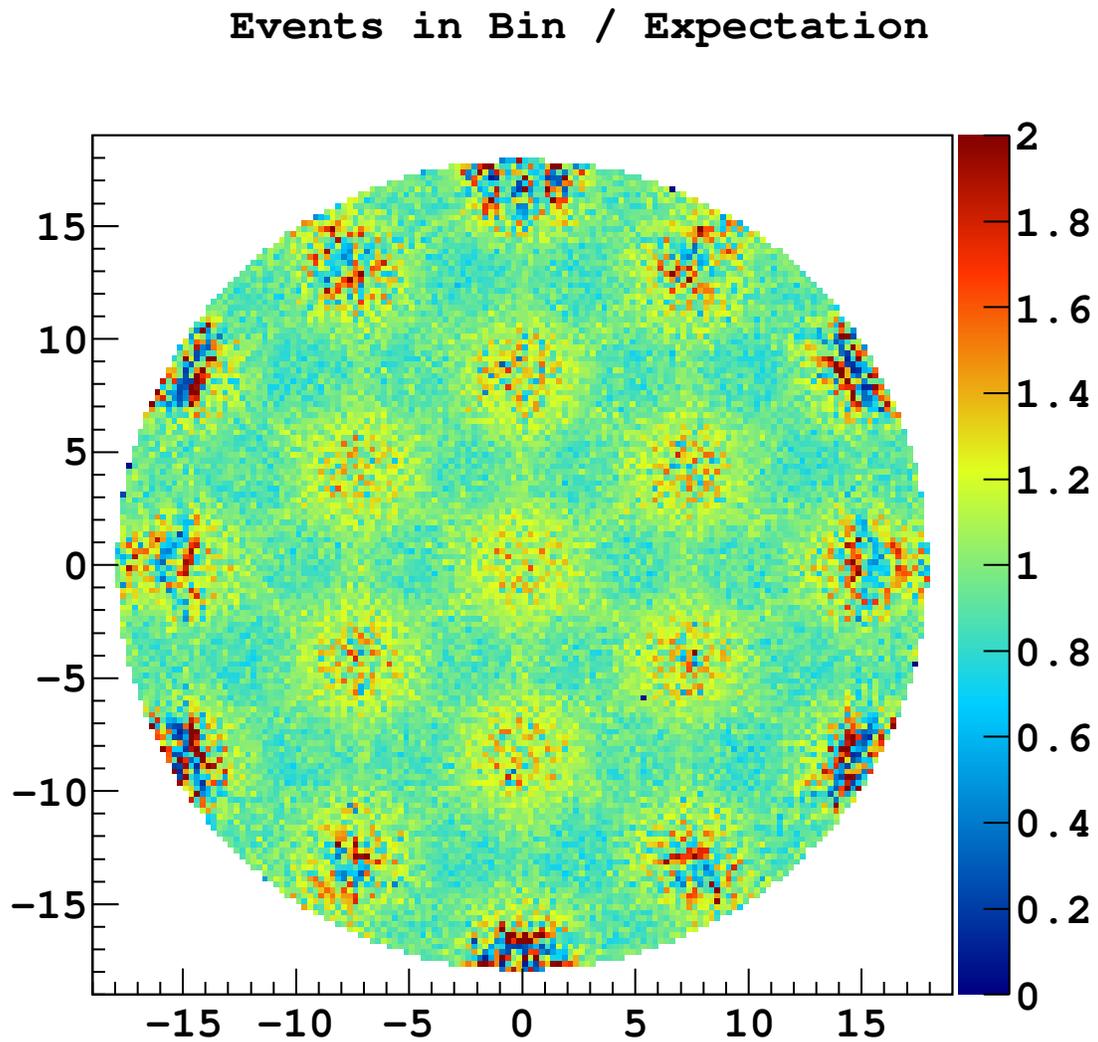


Figure 4.13: Event density by position, 100th iteration LRFs, no uniformity pressure, Monte Carlo. Average reconstruction error at this iteration: 0.44 cm.

BiPo coincidences are easy to search for in DarkSide-50 data because a data acquisition waveform triggered by the bismuth-214 decay will usually also include the polonium-214 decay. Thus we can search each waveform for the signature of these coincident decays⁶:

1. The event begins with a β -like S1. The pulse shape discrimination parameter $f90$ must satisfy $0.1 \leq f90 \leq 0.4$.
2. There must be at least four pulses, S1 and S2 for the β and α . Additional S2 pulses can arise from coincident γ emission from the decays, but if there are many more than four pulses conclusive identification of the BiPo is hard. Thus, we require $4 \leq \text{pulses} \leq 8$.
3. Pulses are identified as S1 if they have $f90 \geq 0.1$, and as S2 if $f90 < 0.1$. Pulses with less than 100 photoelectrons are not counted as either S1 or S2.
4. There must be exactly two S1 pulses, and at least two S2 pulses.
5. The second S1 pulse is from the α decay, which should have a prompt pulse shape and a large energy. It must have $f90 \geq 0.35$ and at least 5000 photoelectrons.
6. The drift time of the two decays should be the same, and thus the time between the S1 pulses should be the same as the time between the S2 pulses. We require the drift times match to $\pm 5 \mu\text{s}$ ⁷. If more than two S2 pulses exist, the BiPo search checks every pair of S2 pulses for a drift time match with the S1 pulses. We require exactly one pair of S2s with matching drift time and identify them as the β and α S2s.
7. If either the β or the α S2 saturated the digitizer in any channel, the event is thrown out. This saturation would alter the distribution of S2 light among the channels, distorting the position reconstruction. However, saturation in S2 is very rare, as even large S2 pulses do not have a particularly high maximum amplitude.

Figure 4.14 shows a very clean BiPo waveform.

⁶Thanks to C. Stanford for this analysis.

⁷The true drift times should be much closer than this, but coincident γ emission with the decays can throw off the identification of drift time.

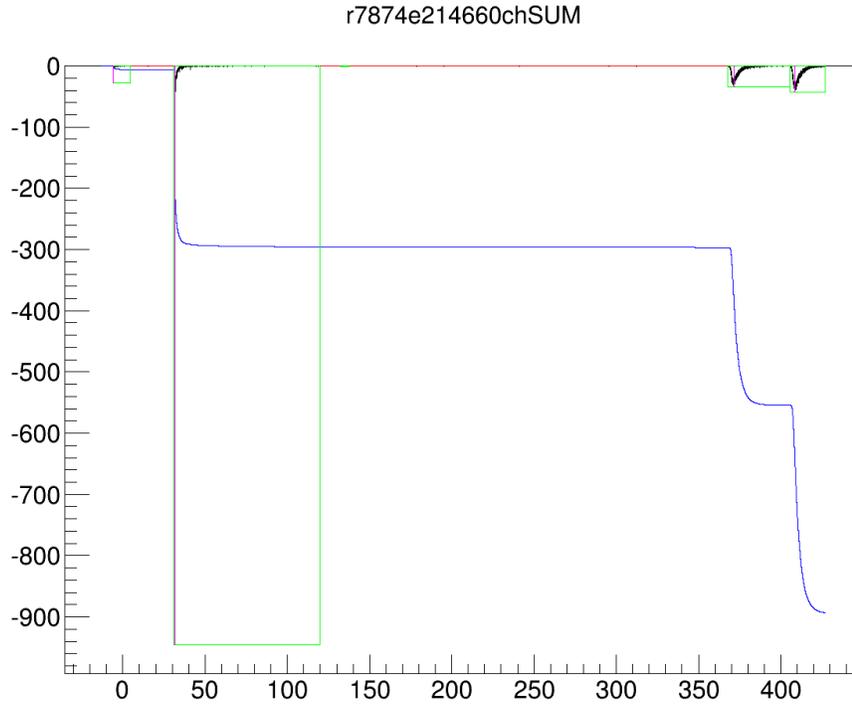


Figure 4.14: A BiPo waveform.

BiPo coincidences are not common in DarkSide-50 due to an efficient radon trap that limits the amount of radon entering the detector. Since bismuth-214 is part of the upper section of the radon-226 decay chain, radon-226 in contact with the detector before the operation of the radon trap has since decayed past the bismuth-214 stage to lead-210 or its daughters. New radon-226 is produced in the detector via decays of uranium-238 in the detector materials, but much of the radon-226 produced cannot diffuse into the argon because such diffusion is very slow at liquid argon temperatures. Most warm surfaces from which radon-226 could diffuse into the argon are part of the gas recirculation system, which passes the radon trap before returning to the detector.

A second kind of coincident decay can also be identified in DarkSide-50. Radium-224, part of the thorium-232 decay chain, undergoes alpha decay into radon-220. Radon-220 then alpha decays into polonium-216 with a half life of 55s, which in turn alpha decays with a half-life of 0.14s. These decays are separated enough in time that they do not occur in the

same digitized waveform, but the signature of three consecutive alphas is distinct enough to select them and use them for the same purpose as BiPos. These “RnPo” decays are selected by:

1. Identify all high-energy alphas in the data by their prompt pulse shape and large $S1$. $f90 > 0.35$ and $S1 > 5000$ photoelectrons.
2. For each set of three alphas consecutive in the list of alphas, examine the time spacing. If the second alpha follows the first by $< 7 \times (55 \text{ s})$, that pair might be a radium-224 \rightarrow radon-220 sequence. If the third alpha follows the second by $< 7 \times (0.14 \text{ s})$, that might be a radon-220 \rightarrow polonium-216 sequence. Trios that do not satisfy these spacings are cut from the selection.
3. The second and third alphas in the trio must have a proper $S2$ pulse. They must have a second pulse with $f90 < 0.1$ and more than 100 photoelectrons. These $S2$ pulses must also not saturate.

More strict cuts on the RnPos are possible, but unnecessary. The total rate of alpha signals in the detector is low enough that three alphas in tight sequence are always from a RnPo, and so also have the same drift time.

Figure 4.15 shows the 27 BiPos and RnPos discovered in the DarkSide-50 data so far plotted by reconstructed position. Each pair of coincident decays is shown as two points with the same color⁸. Figure 4.16 is a histogram of the reconstructed distance between the two decays in each pair. With perfect reconstruction, every pair should reconstruct to the same location, but this algorithm shows some finite resolution as most pairs do not reconstruct to the same location. The red line is a fit to a Gaussian with a mean of zero, and the width of this Gaussian indicates the resolution of the algorithm implied by these coincident decay results. The best fit resolution was $0.61 \pm 0.09 \text{ cm}$. The data does not

⁸Although when the two decays reconstruct at the same location, the two points on Figure 4.15 look like one point.

closely match the fitted Gaussian, suggesting the possibility that reconstruction errors are not normally distributed, though more statistics are required to conclude that with certainty. One result of the non-Gaussian behavior is the mean of the histogram is somewhat smaller than the expected mean of a half-Gaussian with $\sigma = 0.61$ cm. The observed mean is 0.39 cm, while 0.48 cm is expected.

4.4.2 Recovery from Flawed Simulation

When reconstructing data, this algorithm begins with LRFs derived from simulated events, and then real data events are used to produce new iterations of the LRFs. As the data events are known to differ in some ways from the simulated events, it is believed the iteration method can produce LRFs that are accurate despite the initial LRFs being inaccurate. This process can be tested by substituting simulated “flawed” events for data events.

I simulated a set of 10^6 events with an average of 17,190 photoelectrons per event, similar to the standard low-statistics Monte Carlo set. This set is distinguished by the presence of “flaws” in the simulation, deliberate alterations to the simulated TPB. These flaws were designed to simulate the possibility that in the physical detector the TPB is degraded in some locations, either by falling off the surface or by oxidization. TPB degradation is one way in which the physical detector could differ from the detector design implemented in *g4ds*, and so these events show examples of the kind of behavior we believe could exist in data and might cause a problem by being unaccounted for in the initial LRFs.

The “flaws” were implemented by adjusting the TPB wavelength shifting code to reduce the probability that a visible photon is emitted after the absorption of a ultraviolet photon. This reduction applies only when the ultraviolet photon was absorbed in one of three regions (units in cm):

- The “north” flaw. In the region $y > 15$, $-2 < x < 2$, the emission probability is reduced along a horizontal gradient. At $x = -2$ or 2 , the emission probability is normal, at

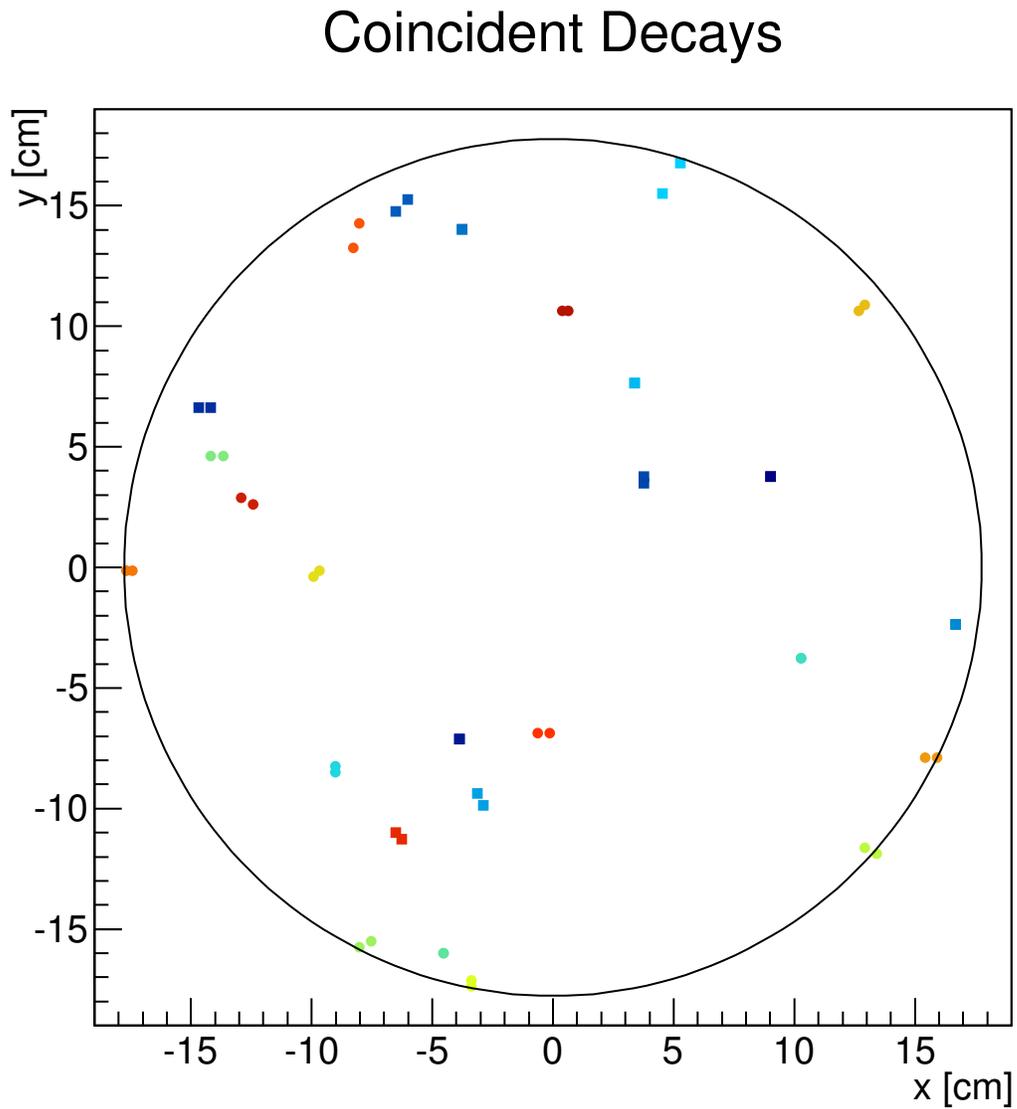


Figure 4.15: Reconstructed location of ^{214}Bi and ^{214}Po decays. Same-color pairs of points indicate the two decays in a BiPo event. Squares are BiPos, circles are RnPos.

Coincident Decays Reconstruction Position

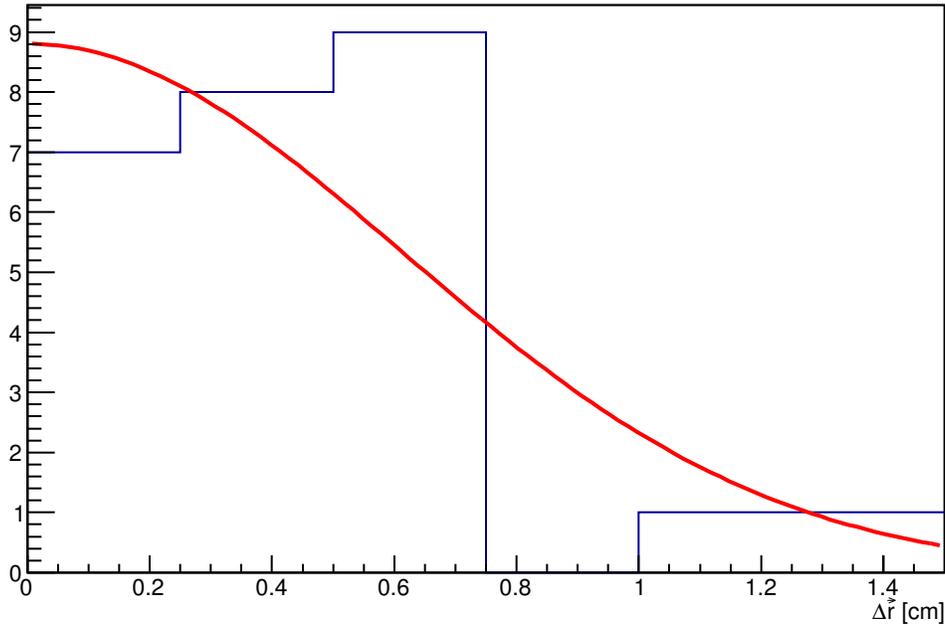


Figure 4.16: Reconstructed distance between ^{214}Bi and ^{214}Po decays, and fit to a resolution $= 0.61 \pm 0.09$ cm. The mean of the histogram is 0.39 cm.

$x = -1$ or 1 it is half normal and at $x = 0$ the emission probability is zero. This flaw covers much of the area under PMT 37.

- The “west” flaw. In the region $x < -15$, $3 < y < 5$, the emission probability is reduced along a vertical gradient. At $y = 3$ or 5 , the emission probability is normal, at $y = 3.5$ or 4.5 it is half normal and at $y = 4$ the emission probability is zero. This flaw is largely not under any PMT, instead being mostly under the reflector between PMTs 27 and 28.
- The “east” flaw. In the region $r > 16$, $y > -15$, $-2 < x < 2$, the emission probability is reduced linearly as the radius increases. At $r = 16$ the emission probability is normal, and by $r = 17.77$, the emission probability is zero. This flaw is under PMT 21, which covers this section of the detector’s edge.

The TPB behavior remains normal in all other regions, and in the altered regions the TPB still reflects, absorbs, and transmits light as usual.

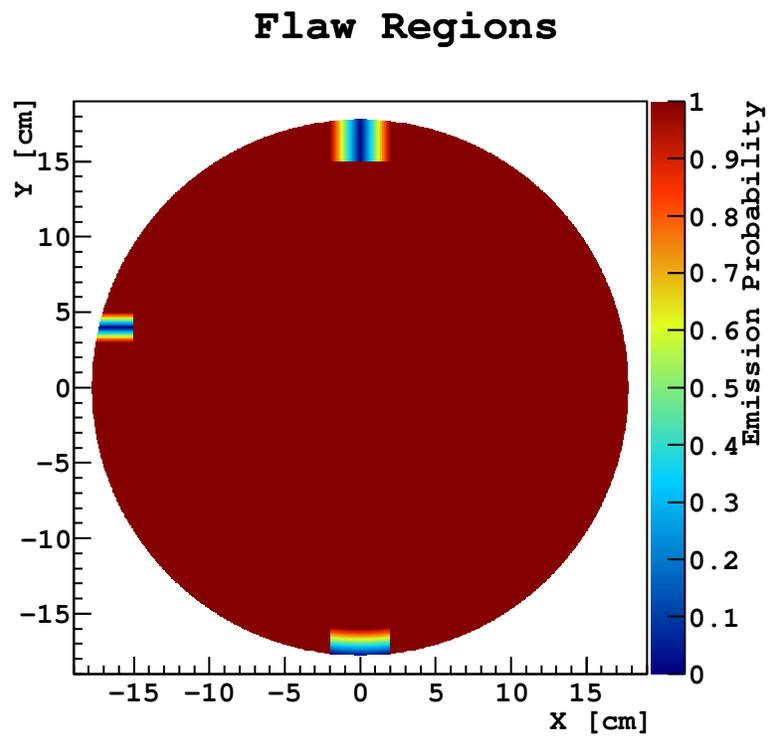


Figure 4.17: Regions where the “flawed” simulation has the visible light emission probability of TPB reduced.

The flawed simulation also incorporates a number of tuning adjustments made to $g4ds$ after the high-statistics events were created. This creates discrepancies between the light collection spectra of the flawed and high-statistics simulated event sets similar to the discrepancies seen between the high-statistics simulated events and data (Section 3.5). These discrepancies appear in all PMTs, not just the ones near the flaw. Thus, tests of the reconstruction of the flawed data can also test whether the adjustment scheme for Monte Carlo discrepancies works.

These flawed events are reconstructed via the following scheme, identical to the scheme for reconstructing data events except for the substitution of flawed simulated events:

1. Initial LRFs are built from high-statistics simulated events without flaws.
2. Second-generation LRFs are built from flawed simulated events.
3. Further generations of LRFs are built from flawed simulated events, iteratively.

As the reconstructed events originated from simulation, the reconstructed positions can be compared to the true positions to measure the error of the reconstruction. Figures 4.18 and 4.19 show the reconstruction error plotted across the detector during the first and 30th reconstruction, respectively (using the 0th- and 29th-iteration LRFs).

In Figure 4.18, the north and south flaws show up as regions with exceptionally high reconstruction errors of several centimeters. This shows that events that belong under the flaw are instead being reconstructed several centimeters away. In the case of the south flaw, the events under the flaw are reconstructed around $y = -15$, which causes the events that belong around $y = -15$ to move away from that region due to uniformity pressure. That creates the “two blob” effect in the reconstruction error in the south end of the detector. The west flaw appears as a yellow dot around $(-15, 4)$. Either because of the smaller size or the placement under the reflector, the west flaw induces smaller reconstruction error than the other two flaws. Elsewhere in the detector the reconstruction errors are minimal,

Reconstruction Error with Flaws, Iteration 1

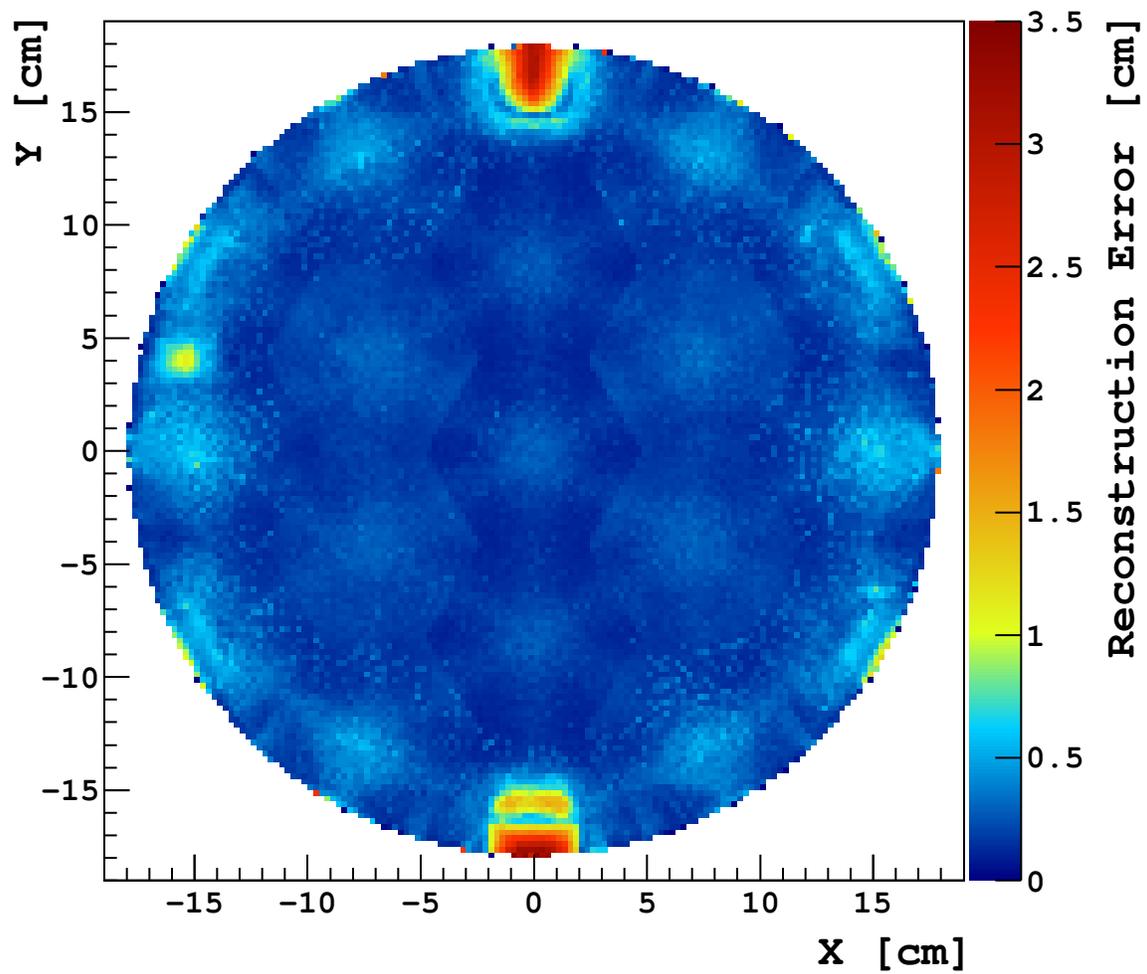


Figure 4.18: Reconstruction error as a function of position, using 0th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.

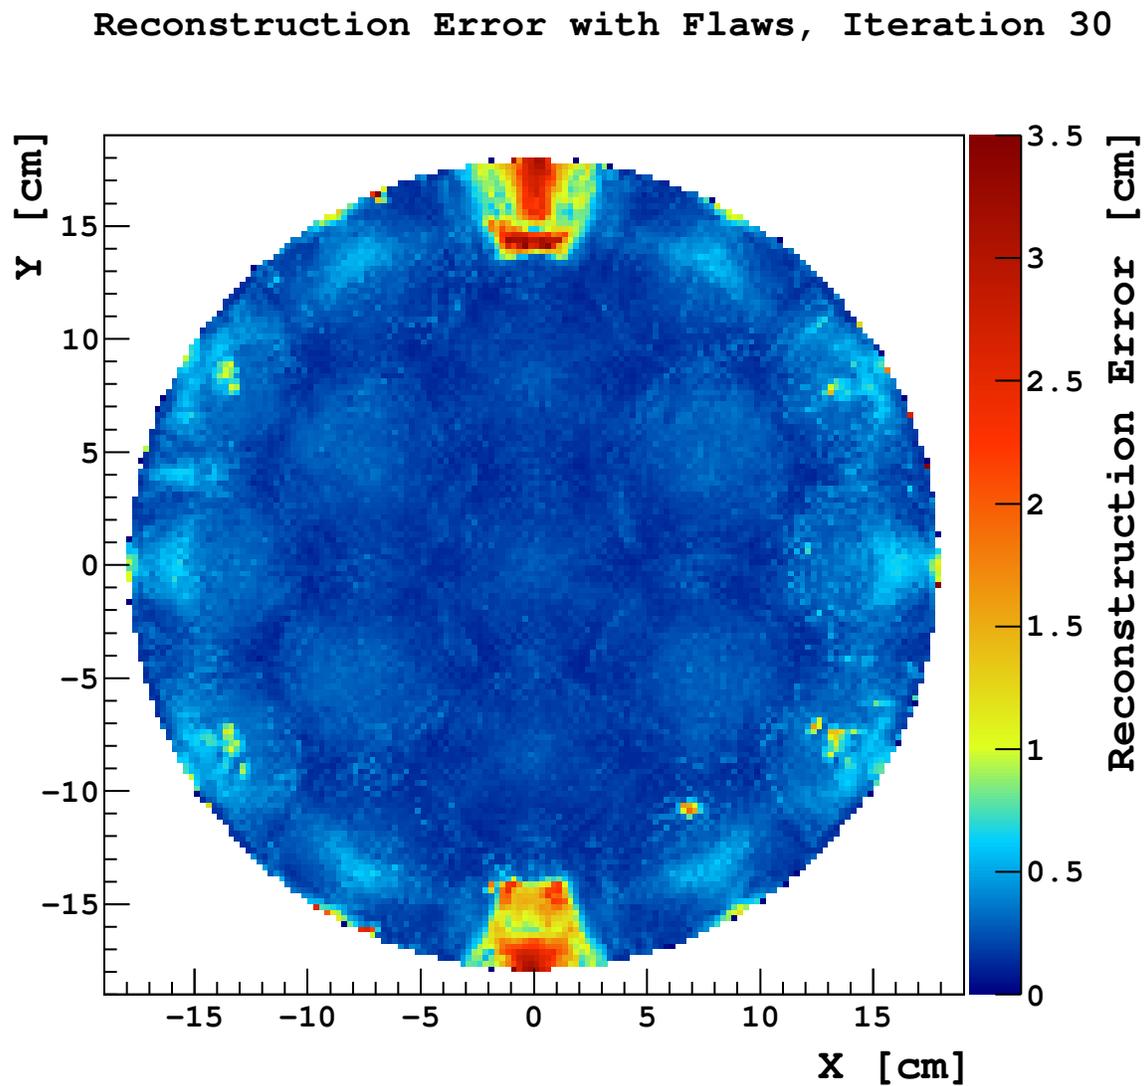


Figure 4.19: Reconstruction error as a function of position, using 29th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.

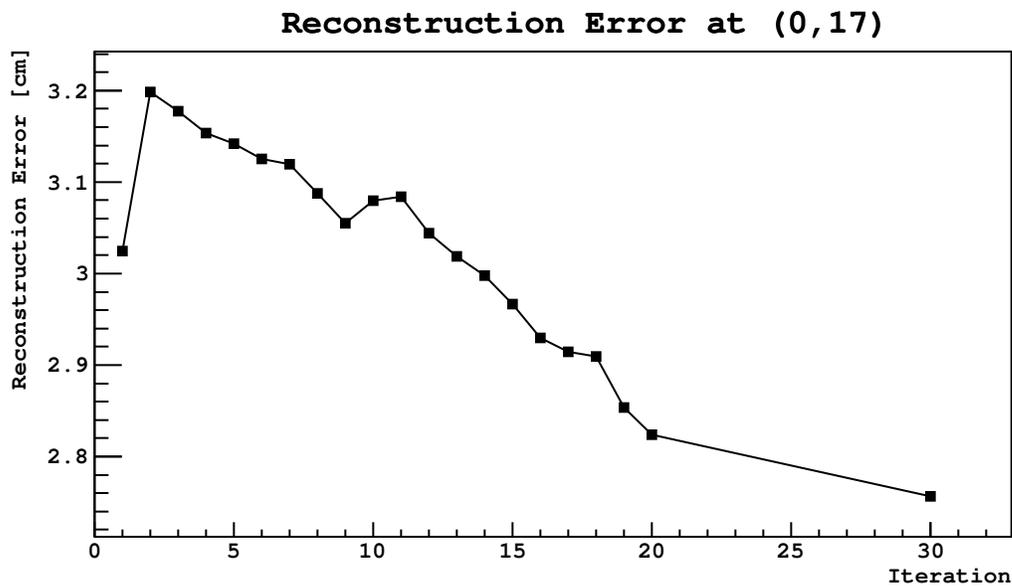


Figure 4.20: Reconstruction error in the north flaw in each iteration using flawed Monte Carlo events.

indicating, happily, that the spectra discrepancy adjustments performed well across all non-flawed regions of the detector.

30 iterations later, in Figure 4.19, the west flaw now has a reconstruction error of only about 0.6 cm, a marked improvement from the 1.1 cm reconstruction error seen in the first iteration. The north and south flaws have also improved, but by a smaller amount and at the cost of extending the region affected by reconstruction error, as events from a broad area shuffle around trying to resolve the nonuniformity introduced by the reconstruction error. Figures 4.20 and 4.21 show this improvement in the reconstruction error in the flaws.

My conclusion from these results is that the iteration approach can fix some flaws, but others, if they existed in data, would remain uncorrected in a reasonable number of iterations. Thankfully, the irreparable north and south flaws leave equally irreparable marks on the uniformity results. Figure 4.22 shows the uniformity after 30 iterations of the flawed events. This is important because it implies if the real data had similar irreparable flaws, there would be a sign of that in the uniformity results for data—and no such severe nonuniformity exists in the reconstruction of data.

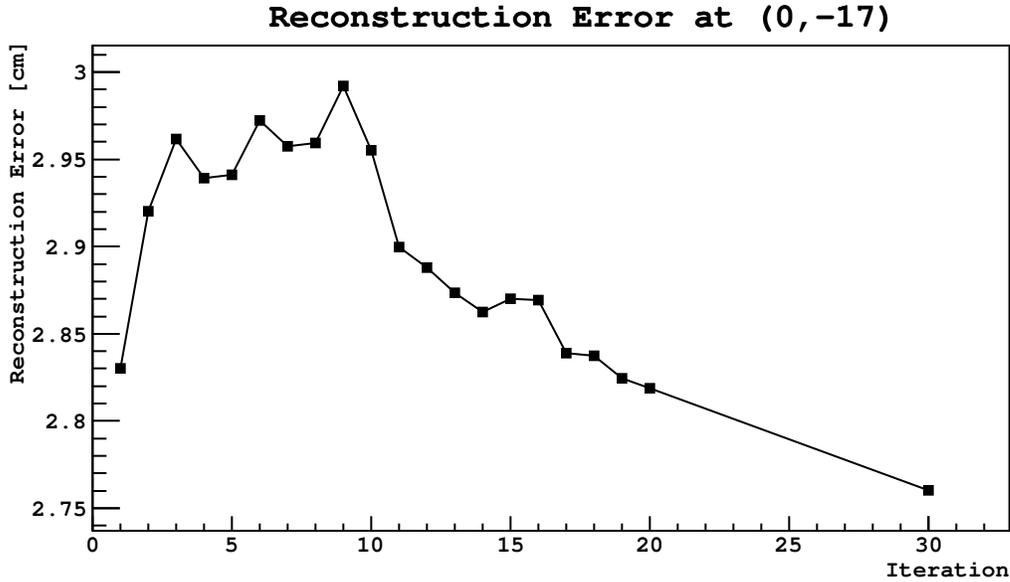


Figure 4.21: Reconstruction error in the south flaw in each iteration using flawed Monte Carlo events

4.5 Surface Background Cut Results

I applied the position reconstruction and surface background analysis to the several different event sets. This produces a value of the surface cut metric (*SCM*) for each event. As discussed in Section 3.7.4, the *SCM* ranges from 0 to 1, and is higher for events that more plausibly reconstruct at the edge of the detector. Using these events, I examine the consequences of cutting all events with *SCM* greater than some threshold. For all sets, I analyze the fraction of the total events cut by different thresholds. The fewer events cut, the better the cut's acceptance and the better the detector's sensitivity to WIMPs. For simulated events, I can analyze the efficiency of removing events truly at the surface, as this requires known true positions for the events.

4.5.1 Acceptance vs Rejection in Simulated Events

I considered three different sets of simulated events.

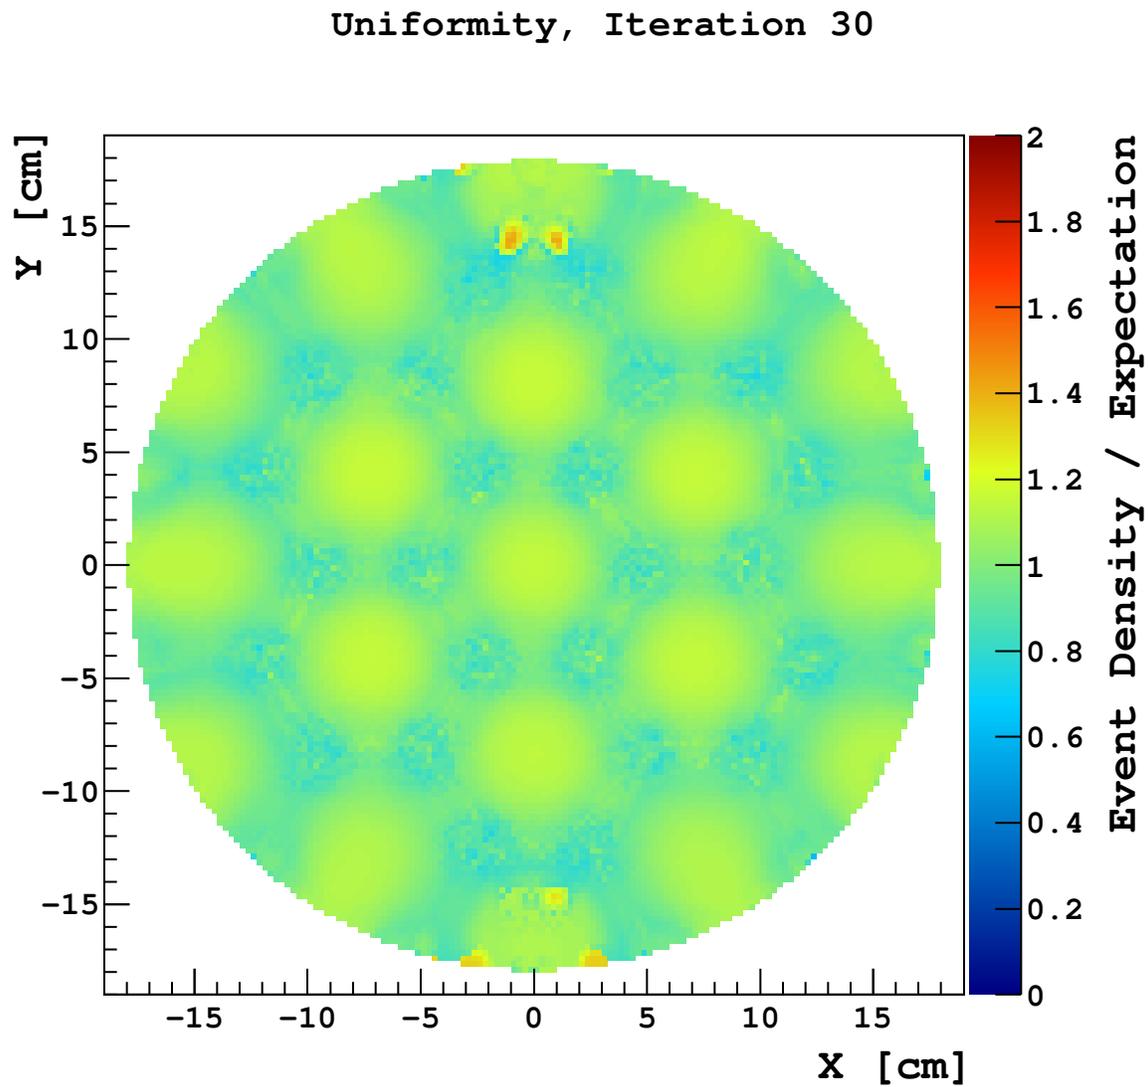


Figure 4.22: Event density by position, using 29th-iteration LRFs, with uniformity pressure, flawed Monte Carlo events.

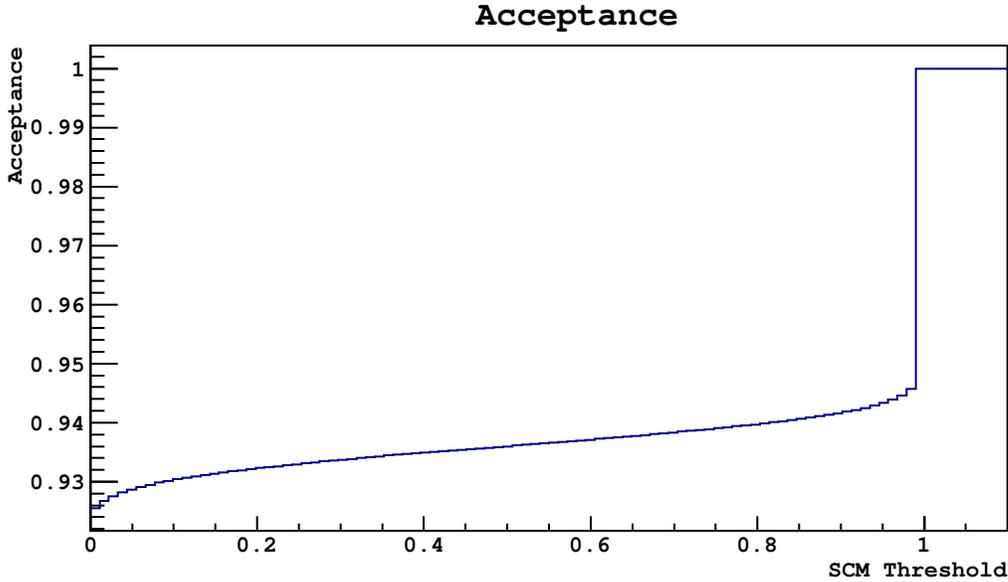


Figure 4.23: Acceptance at cut thresholds, high-statistics Monte Carlo.

High-statistics simulated events have excellent reconstruction accuracy and precision, so I expected they would also have excellent surface background cut performance. However, while cut acceptance was very good, background rejection was hampered by an unexpected effect.

Figure 4.23 shows the dependence of cut acceptance on the SCM threshold. If all events with $SCM > 0$ are cut,⁹ 92.6% of events are still accepted. As the threshold is raised, more and more events pass the cut. Naturally, the acceptance goes to 100% if the threshold is $SCM > 1$, as the SCM cannot be larger than 1.

Background rejection is measured by identifying the simulated events whose true position was within $50 \mu\text{m}$ of the outer radius of the detector and counting how many are caught by the cut at each threshold. Background rejection for the high-statistics events was substantially worse than expected, as seen in Figure 4.24. The highest achievable background rejection only cuts backgrounds to a factor of 10^{-3} , and that is achieved only by setting the SCM threshold almost as low as it can go. On close examination of the individual surface

⁹Recall that $SCM = 0$ either when the event's search region did not include the edge or when no edge position had a χ^2 below the storage threshold of 90.

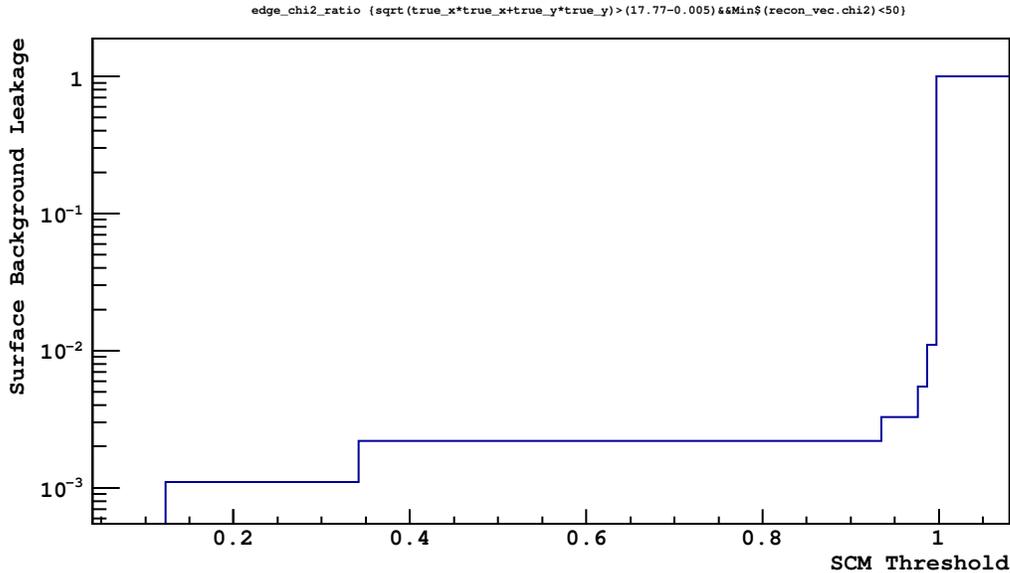


Figure 4.24: Background rejection at cut thresholds, high-statistics Monte Carlo.

background events with $\text{SCM} < 0.2$, I observed they all had their best χ^2 one or two bins from their true position at the edge. By itself is this is not surprising given the ~ 0.2 cm reconstruction error. What is surprising is that the true position had a χ^2 much worse, often several times the best fit χ^2 . This explains the low SCM, which ought to be low when the edge position's χ^2 is several times the best χ^2 . However, the result is puzzling given the expected behavior of χ^2 : an ordinary statistical fluctuation in an event, of the sort needed to cause the event to mis-reconstruct by one bin, should produce only a modest change in χ^2 . As the χ^2 formula divides each term by the variance, a typical fluctuation should increase χ^2 by order 1.

This behavior was explained by examining the events that truly belong in the edge bins. For these events, the standard deviation of the S2 fraction in the closest PMT was observed to be $4-8\times$ the statistically expected value (determined from \sqrt{n} behavior). This additional error has a systematic source: the change in the LRF on a scale shorter than the length of one bin (Section 4.2.1). This systematic error was sub-dominant in past studies using lower-S2 event sets. When it dominates in high-S2 samples, the error estimates in the χ^2 , which only include statistical errors, are too low. This causes χ^2 to be too high in the edge bins if

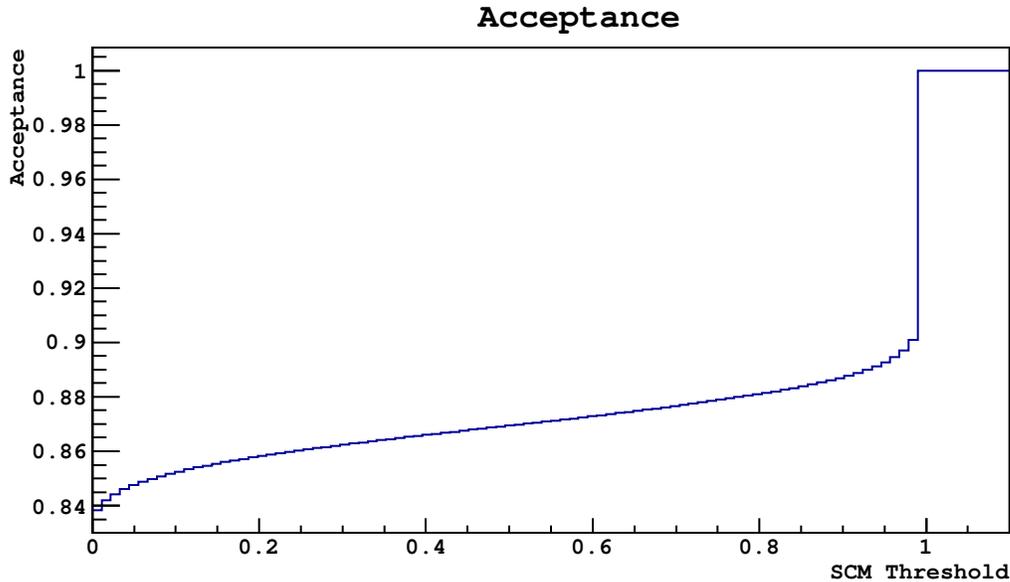


Figure 4.25: Acceptance at cut thresholds, low-statistics Monte Carlo.

the events fluctuate away from the typical behavior for that location. That, in turn, reduces SCM, often to zero. This causes some events to leak through the background cut, and may also contribute to the very high acceptance.

In low-statistics simulated events, the effect of the systematic error in the edge bins is greatly reduced, leading to better background rejection. However, the loss of precision at lower statistics reduces the acceptance.

Figure 4.25 shows the acceptance with the lower statistics events, and Figure 4.26 shows the background rejection. It is now possible to reduce the surface backgrounds to a factor of 4/10,000 of their original value while accepting 85% of events.

4.5.2 Acceptance vs Rejection in Data Events

It is not possible to measure background rejection in data events directly, as there is no independent way of identifying surface backgrounds. However, overall acceptance is straightforward to measure by identifying which data events in a sample would pass a surface background cut with the threshold set at a given level. The data events tested in this section are

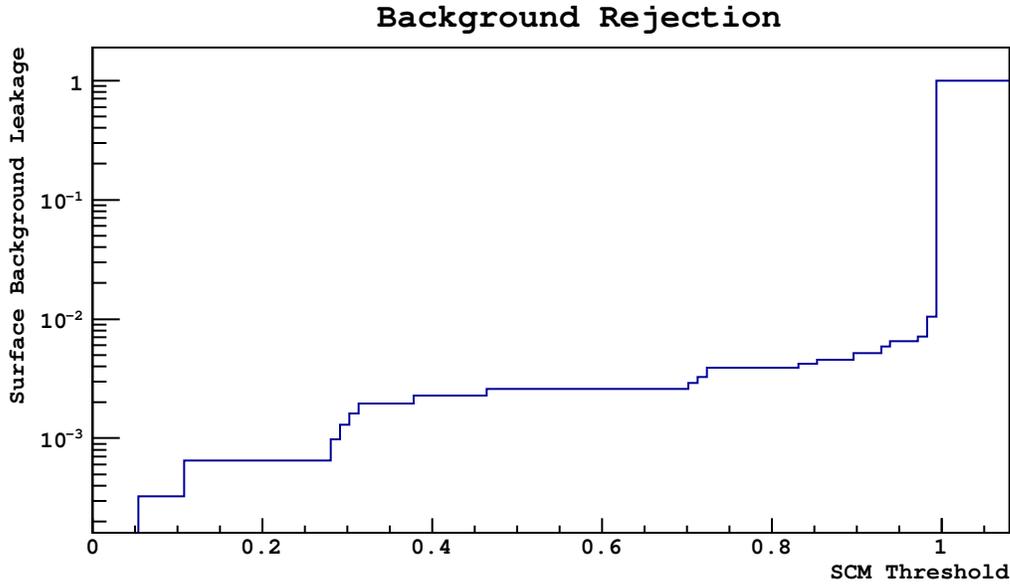


Figure 4.26: Background rejection at cut thresholds, low-statistics Monte Carlo.

the sample described in Section 4.1.3, and so have already undergone event quality cuts and a S1 cut.

Figure 4.27 shows the acceptance as a function of cut threshold. The acceptance is lower than the low-statistics Monte Carlo. This is not necessarily an indication the data LRFs are less accurate than the Monte Carlo LRFs. LRF inaccuracies by themselves could make events fit either better or worse at the edge than they truly should, so could either lower or raise the acceptance. Instead, the lower acceptance in data indicates that the data LRFs indicate smaller optical differences between surface events and non-surface events, so it is easier for a statistical fluctuation of a non-surface event to make it look like a surface event. These smaller optical differences could be real properties of the detector, or results of inaccurate LRFs.

I cannot measure background rejection in data events directly, because there is no set of known surface background events. It is possible to estimate the surface background rejection using the coincident decays discussed in Section 4.4.1. These decays provide pairs of events believed to be at the same location. If one decay in a pair reconstructs close to the edge, the other decay might truly have come from that position, and so we pretend one decay's

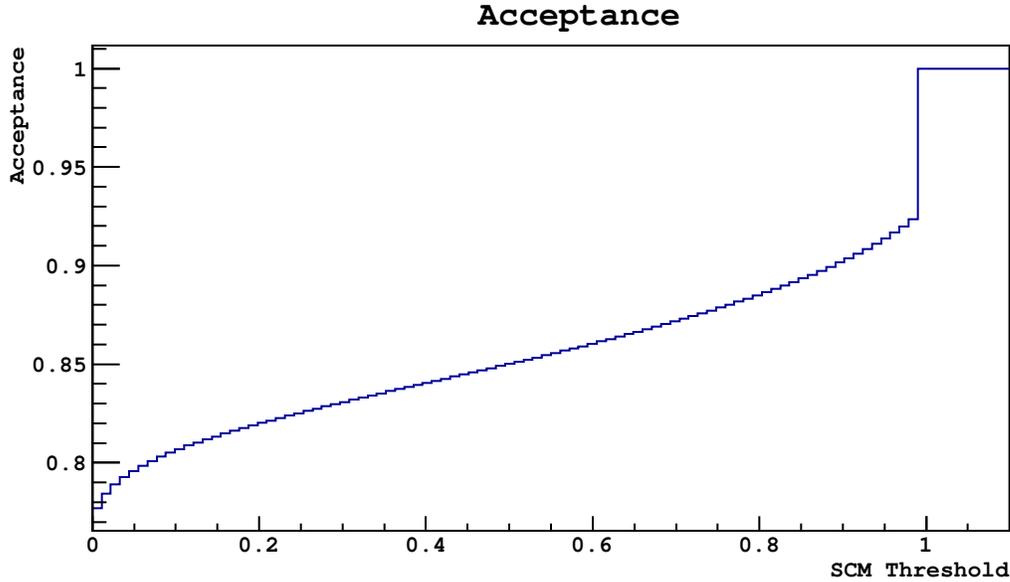


Figure 4.27: Acceptance at cut thresholds, data.

location was the surface location and the other decay was a surface event, and see how well the “surface event” reconstructs at the “surface location.”

More specifically: For a pair of decays, decay A is the one with lower χ^2 , and therefore more plausibly reconstructed at the true position, and decay B is the one with higher χ^2 . If decay A reconstructed within 1 cm of the surface, I calculate a pseudoSCM for decay B :

$$\text{pSCM} = P(\chi_A^2 - \chi_B^2),$$

where χ_A^2 is the χ^2 of decay B at the reconstructed location of decay A , and χ_B^2 is the χ^2 of decay B at decay B ’s reconstructed position—i.e. the best χ^2 for decay B anywhere. From the 11 pairs with decay A near the edge, I create a distribution of pseudoSCMs. I accumulate this histogram to indicate the fraction of events that would pass a cut on the pSCM, shown in Figure 4.28. With only 11 such decays, there are not sufficient statistics to see whether rejection of 10^{-3} or better can be achieved in data. This estimate with coincident decays does suggest that background rejection in data may be worse than what is observed in simulated

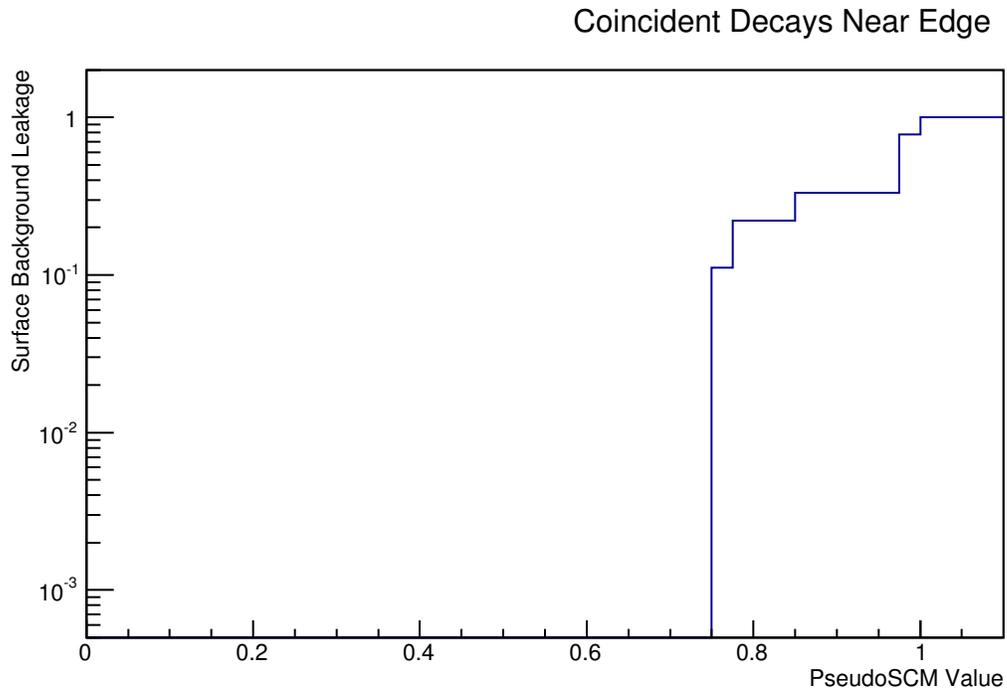


Figure 4.28: Estimation of background rejection using pseudoSCM on coincident decays.

events, although it remains possible that at lower *SCM* values the data background rejection catches up.

Chapter 5

Conclusion

DarkSide-50 is a detector with great promise. Success at the current stage will demonstrate the ability of liquid argon TPC technology to conduct background-free WIMP searches, offering a competitive alternative to other technologies. The position reconstruction algorithm presented here will identify surface backgrounds in DarkSide-50 and thus help keep DarkSide-50 background-free. The algorithm shows several promising measures of success, but further analysis is possible. There also exists great potential to extend the application of this algorithm to other purposes in DarkSide-50.

No one test provides a clear-cut measurement of the accuracy of the algorithm when applied to real data. However, tests on coincident decays demonstrate precision of 0.6 ± 0.09 cm. The coincident decay data also suggests that to the extent that some events are not reconstructed accurately, the SCM cut will still allow surface backgrounds to be cut with only reasonable acceptance losses. The strong performance of the algorithm on simulated events also suggests accurate reconstruction of data is possible.

The “bogeyman” of detector flaws producing an inaccurate reconstruction still exists, but tests on simulated data suggest that some detector flaws will be handled well by the iterative approach, producing no increased reconstruction errors. Other flaws will not be handled well, but if these existed there would be evidence such significant flaws exist in the

uniformity results. Further studies should search for “stealth” flaws that produce significant errors but do not show up in the uniformity results.

One interesting potential approach to handling the possibility of irregularities in the wavelength shifter is to attempt to simulate a nonuniform wavelength shifter, basing the distribution of shifting efficiency on the reconstruction results. Currently, the data reconstruction in early stages shows “hot spots” at locations where the simulation does not resemble the data. It may be possible to identify these regions in a quantitatively rigorous way and assign them modified wavelength shifting efficiency in the simulation. If the resulting simulated events match the data and do not produce hot spots, this technique would allow for a reconstruction of the nonuniformity of the wavelength shifting layer based on the effects in data. That result could inform the design of future detectors by revealing if specific areas of the wavelength shifting layer were damaged in the operation of DarkSide-50.

Another interesting future study involves attempting to utilize position-dependent errors. The uncertainty in some bins of the LRF is greater than others because those bins span a sharp change in the light collection of the detector. Currently, this information is not used in finding the best-fit position for events because allowing position-dependent errors to drive the fit results in events being assigned to positions with large errors instead of to positions where the LRFs match the event. However, position-dependent errors could play a valuable role in the surface cut. Surface positions with higher errors when taking into account position-dependent errors may be considered more plausible when evaluating which events are plausibly from the surface.

The position reconstruction algorithm is expected to play a role in the future in cutting another type of background: multiple-scatter backgrounds. Any particle that interacts twice in the TPC, at two different positions, must not be a WIMP, as WIMPs interact too rarely. Often, a multiple-scatter background can be identified using the TPC’s excellent z -axis reconstruction, as many multiple scatters will occur at different z -positions. However, some of these backgrounds will have both scatters occur at the same z -position, which would produce

only a single S2 pulse. However, this pulse would contain different xy position information than single-scatter events, so some modification to the xy reconstruction algorithm could discriminate against multiple-scatter backgrounds.

This xy reconstruction algorithm finds the position of events based on optical information, despite uncertainties in the optical behavior of DarkSide-50 reflected in discrepancies between simulation and data. In doing so, it faces one of the fundamental challenges of working with low-background detectors: the closed environment can result in uncertainties about what is happening inside the detector. This algorithm demonstrates one way to “look inside” DarkSide-50 and understand what happens there.

Bibliography

- [1] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.
- [2] V. C. Rubin, W. K. J. Ford, and N. . Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/. *The Astrophysical Journal*, 238:471–487, June 1980.
- [3] J. L. Feng. Dark Matter Candidates from Particle Physics and Methods of Detection. *Annual Review of Astronomy and Astrophysics*, 48(1):495–545, 2010.
- [4] E. Komatsu, C. L. Bennett, C. Barnes, R. Bean, C. L. Bennett, O. Doré, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill, G. Hinshaw, N. Jarosik, A. Kogut, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, H. V. Peiris, K. M. Smith, D. N. Spergel, G. S. Tucker, L. Verde, J. L. Weiland, E. Wollack, and E. L. Wright. Results from the Wilkinson Microwave Anisotropy Probe. *Progress of Theoretical and Experimental Physics*, 2014(6):06B102, June 2014.
- [5] R. D. Peccei and H. R. Quinn. CP Conservation in the Presence of Pseudoparticles. *Physical Review Letters*, 38(25):1440–1443, June 1977.
- [6] B. M. Loer. *Towards a depleted argon time projection chamber WIMP search: Darkside prototype analysis and predicted sensitivity*. Ph.D., Princeton University, United States – New Jersey, 2011.
- [7] S. J. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz, L. J. Rosenberg, G. Rybka, J. Hoskins, J. Hwang, P. Sikivie, D. B. Tanner, R. Bradley, and J. Clarke. SQUID-Based Microwave Cavity Search for Dark-Matter Axions. *Physical Review Letters*, 104(4):041301, January 2010.
- [8] L. Bergström. Dark Matter Evidence, Particle Physics Candidates and Detection Methods. *Annalen der Physik*, 524(9-10):479–496, October 2012. arXiv: 1205.4882.
- [9] G. Steigman, B. Dasgupta, and J. F. Beacom. Precise relic WIMP abundance and its impact on searches for dark matter annihilation. *Physical Review D*, 86(2):023506, July 2012.
- [10] A. D. Simone, G. F. Giudice, and A. Strumia. Benchmarks for dark matter searches at the LHC. *Journal of High Energy Physics*, 2014(6):1–27, June 2014.

- [11] J. Xu. *Study of Argon from underground sources for dark matter detection*. Ph.D., Princeton University, United States – New Jersey, 2013.
- [12] G. Bertone, editor. *Particle Dark Matter: Observations, Models and Searches*. Cambridge University Press, 1 edition edition, November 2013.
- [13] LUX Collaboration. First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility. *Physical Review Letters*, 112(9):091303, March 2014.
- [14] E. Aprile, M. Alfonsi, K. Arisaka, F. Arneodo, C. Balan, L. Baudis, B. Bauermeister, A. Behrens, P. Beltrame, K. Bokeloh, E. Brown, G. Bruno, R. Budnik, J. M. R. Cardoso, W.-T. Chen, B. Choi, D. Cline, A. P. Colijn, H. Contreras, J. P. Cussonneau, M. P. Decowski, E. Duchovni, S. Fattori, A. D. Ferella, W. Fulgione, F. Gao, M. Garbini, C. Ghag, K.-L. Giboni, L. W. Goetzke, C. Grignon, E. Gross, W. Hampel, F. Kaether, A. Kish, J. Lamblin, H. Landsman, R. F. Lang, M. Le Calloch, C. Levy, K. E. Lim, Q. Lin, S. Lindemann, M. Lindner, J. A. M. Lopes, K. Lung, T. Marrodán Undagoitia, F. V. Massoli, A. J. Melgarejo Fernandez, Y. Meng, A. Molinario, E. Nativ, K. Ni, U. Oberlack, S. E. A. Orrigo, E. Pantic, R. Persiani, G. Plante, N. Priel, A. Rizzo, S. Rosendahl, J. M. F. dos Santos, G. Sartorelli, J. Schreiner, M. Schumann, L. Scotto Lavina, P. R. Scovell, M. Selvi, P. Shagin, H. Simgen, A. Teymourian, D. Thers, O. Vitells, H. Wang, M. Weber, and C. Weinheimer. Dark Matter Results from 225 Live Days of XENON100 Data. *Physical Review Letters*, 109(18):181301, November 2012.
- [15] R. Bernabei, P. Belli, F. Cappella, R. Cerulli, C. J. Dai, A. d’Angelo, H. L. He, A. Incicchitti, H. H. Kuang, J. M. Ma, F. Montecchia, F. Nozzoli, D. Prospero, X. D. Sheng, and Z. P. Ye. First results from DAMA/LIBRA and the combined results with DAMA/NaI. *The European Physical Journal C*, 56(3):333–355, August 2008.
- [16] C. E. Aalseth, P. S. Barbeau, J. Colaresi, J. I. Collar, J. D. Leon, J. E. Fast, N. E. Fields, T. W. Hossbach, A. Knecht, M. S. Kos, M. G. Marino, H. S. Miley, M. L. Miller, J. L. Orrell, and K. M. Yocum. Search for An Annual Modulation in Three Years of CoGeNT Dark Matter Detector Data. *arXiv:1401.3295 [astro-ph, physics:hep-ex, physics:hep-ph]*, January 2014. arXiv: 1401.3295.
- [17] CDMS Collaboration. Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *arXiv:1304.4279 [astro-ph, physics:hep-ex, physics:physics]*, April 2013. arXiv: 1304.4279.
- [18] K. Freese, M. Lisanti, and C. Savage. Annual Modulation of Dark Matter: A Review. *arXiv:1209.3339 [astro-ph, physics:hep-ph]*, September 2012. arXiv: 1209.3339.
- [19] R. Bernabei, P. Belli, F. Montecchia, F. Nozzoli, F. Cappella, A. Incicchitti, D. Prospero, R. Cerulli, C. J. Dai, H. L. He, H. H. Kuang, J. M. Ma, X. H. Ma, X. D. Sheng, Z. P. Ye, R. G. Wang, and Y. J. Zhang. Investigating electron interacting dark matter. *Physical Review D*, 77(2):023506, January 2008.

- [20] CDMS Collaboration. First Results from the Cryogenic Dark Matter Search in the Soudan Underground Laboratory. *Physical Review Letters*, 93(21):211301, November 2004.
- [21] V. M. Gehman, S. R. Seibert, K. Rielage, A. Hime, Y. Sun, D. M. Mei, J. Maassen, and D. Moore. Fluorescence efficiency and visible re-emission spectrum of tetraphenyl butadiene films at extreme ultraviolet wavelengths. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 654(1):116–121, October 2011.
- [22] The ArDM Collaboration, V. Boccone, P. K. Lightfoot, K. Mavrokoridis, C. Regenfus, C. Amsler, A. Badertscher, A. Bueno, H. Cabrera, M. C. Carmona-Benitez, M. Daniel, E. J. Daw, U. Degunda, A. Dell’Antone, A. Gendotti, L. Epprecht, S. Horikawa, L. Kaufmann, L. Knecht, M. Laffranchi, C. Lazzaro, D. Lussi, J. Lozano, A. Marchionni, A. Melgarejo, P. Mijakowski, G. Natterer, S. Navas-Concha, P. Otyugova, M. d. Prado, P. Przewlocki, F. Resnati, M. Robinson, J. Rochet, L. Romero, E. Rondio, A. Rubbia, N. J. C. Spooner, T. Strauss, J. Ulbricht, and T. Viant. Development of wavelength shifter coated reflectors for the ArDM argon dark matter detector. *Journal of Instrumentation*, 4(06):P06001, June 2009.
- [23] G. Bellini, J. Benziger, D. Bick, G. Bonfini, D. Bravo, M. B. Avanzini, B. Caccianiga, L. Cadonati, F. Calaprice, P. Cavalcante, A. Chavarria, A. Chepurinov, D. D’Angelo, S. Davini, A. Derbin, A. Empl, A. Etenko, K. Fomenko, D. Franco, C. Galbiati, S. Gazdara, C. Ghiano, M. Giammarchi, M. Göger-Neff, A. Goretti, L. Grandi, C. Hagner, E. Hungerford, A. Ianni, A. Ianni, V. Kobychyev, D. Korablev, G. Korga, D. Kryn, M. Laubenstein, T. Lewke, E. Litvinovich, B. Loer, P. Lombardi, F. Lombardi, L. Ludhova, G. Lukyanchenko, I. Machulin, S. Manecki, W. Maneschg, G. Manuzio, Q. Meindl, E. Meroni, L. Miramonti, M. Misiaszek, R. Möllenberg, P. Mosteiro, V. Muratova, L. Oberauer, M. Obolensky, F. Ortica, K. Otis, M. Pallavicini, L. Papp, L. Perasso, S. Perasso, A. Pocar, G. Ranucci, A. Razeto, A. Re, A. Romani, N. Rossi, R. Saldanha, C. Salvo, S. Schönert, H. Simgen, M. Skorokhvatov, O. Smirnov, A. Sotnikov, S. Sukhotin, Y. Suvorov, R. Tartaglia, G. Testera, D. Vignaud, R. B. Vogelaar, F. Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 2013(08):049, August 2013.
- [24] J. Brodsky, C. Cao, A. Chavarria, S. Davini, A. Fan, L. Grandi, Y. Guardincerri, K. R. Herber, G. Koh, A. Kurlej, L. Marini, M. E. Monzani, L. Pagani, K. Randle, B. Reinhold, R. Saldanha, M. Wada, A. Watson, S. Westerdale, and J. H. Yoo. Early Analysis on DS-50 Commissioning Run, March 2014.
- [25] P. Agnes, T. Alexander, A. Alton, K. Arisaka, H. O. Back, B. Baldin, K. Biery, G. Bonfini, M. Bossa, A. Brigatti, J. Brodsky, F. Budano, L. Cadonati, F. Calaprice, N. Canci, A. Candela, H. Cao, M. Cariello, P. Cavalcante, A. Chavarria, A. Chepurinov, A. G. Cocco, L. Crippa, D. D’Angelo, M. D’Incecco, S. Davini, M. De Deo, A. Derbin, A. Devoto, F. Di Eusano, G. Di Pietro, E. Edkins, A. Empl, A. Fan, G. Fiorillo, K. Fomenko, G. Forster, D. Franco, F. Gabriele, C. Galbiati, A. Goretti, L. Grandi, M. Gromov, M. Y.

- Guan, Y. Guardincerri, B. Hackett, K. Herner, E. V. Hungerford, A. Ianni, A. Ianni, C. Jollet, K. Keeter, C. Kendziora, S. Kidner, V. Kobychyev, G. Koh, D. Korablev, G. Korga, A. Kurlej, P. X. Li, B. Loer, P. Lombardi, C. Love, L. Ludhova, S. Luitz, Y. Q. Ma, I. Machulin, A. Mandarano, S. Mari, J. Maricic, L. Marini, C. J. Martoff, A. Meregaglia, E. Meroni, P. D. Meyers, R. Milincic, D. Montanari, A. Monte, M. Montuschi, M. E. Monzani, P. Mosteiro, B. Mount, V. Muratova, P. Musico. First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso. *Physics Letters B*, 743:456–466, April 2015.
- [26] M. Berger, J. Coursey, M. Zucker, and J. Chang. ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.2.3), 2005. National Institute of Standards and Technology, Gaithersburg, MD.
- [27] H. Paul. Stopping Power for Light and Heavier Ions: Graphs, Data, Comments and Programs.
- [28] V. Solovov, V. Belov, D. Akimov, H. Araujo, E. Barnes, A. Burenkov, V. Chepel, A. Currie, L. DeViveiros, B. Edwards, C. Ghag, A. Hollingsworth, M. Horn, G. Kalmus, A. Kobayakin, A. Kovalenko, V. Lebedenko, A. Lindote, M. Lopes, R. Luscher, P. Majewski, A. Murphy, F. Neves, S. Paling, J. Pinto da Cunha, R. Preece, J. Quenby, L. Reichhart, P. Scovell, C. Silva, N. Smith, P. Smith, V. Stekhanov, T. Sumner, C. Thorne, and R. Walker. Position Reconstruction in a Dual Phase Xenon Scintillation Detector. *IEEE Transactions on Nuclear Science*, 59(6):3286–3293, December 2012.
- [29] A. Morozov, I. Defendi, R. Engels, F. a. F. Fraga, M. M. F. R. Fraga, A. Gongadze, B. Guerard, M. Jurkovic, G. Kemmerling, G. Manzin, L. M. S. Margato, H. Niko, L. Pereira, C. Petrillo, A. Peyaud, F. Piscitelli, D. Raspino, N. J. Rhodes, F. Sacchetti, E. M. Schooneveld, V. Solovov, P. V. Esch, and K. Zeitelhack. Adaptive algorithms of position and energy reconstruction in Anger-camera type detectors: experimental data processing in ANTS. *Journal of Instrumentation*, 8(05):P05002, May 2013.