

THE PHYSICS OF BACKGROUND DISCRIMINATION IN
LIQUID XENON, AND FIRST RESULTS FROM XENON10 IN
THE HUNT FOR WIMP DARK MATTER.

CARL ERIC DAHL

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: TOM SHUTT

SEPTEMBER 2009

© Copyright by Carl Eric Dahl, 2009.

All Rights Reserved

Abstract

The WIMP limit set by the Xenon10 experiment in 2007 signals a new era in direct detection of dark matter, with several large-scale liquid target detectors now under construction. A major challenge in these detectors will be to understand backgrounds at the level necessary to claim a positive WIMP signal. In liquid xenon, these backgrounds are dominated by electron recoils, which may be distinguished from the WIMP signal (nuclear recoils) by their higher charge-to-light ratio. During the construction and operation of Xenon10, the prototype detector Xed probed the physics of this discrimination.

Particle interactions in liquid xenon both ionize and excite xenon atoms, giving charge and scintillation signals, respectively. Some fraction of ions recombine, reducing the charge signal and creating additional scintillation. The charge-to-light ratio, determined by the initial exciton-ion ratio and the ion recombination fraction, provides the basis for discrimination between electron and nuclear recoils. Intrinsic fluctuations in the recombination fraction limit discrimination.

Changes in recombination induce an exact anti-correlation between charge and light, and when calibrated this anti-correlation distinguishes recombination fluctuations from uncorrelated fluctuations in the measured signals. We determine the mean recombination and recombination fluctuations as a function of energy and applied field for electron and nuclear recoils, finding that recombination fluctuations are already the limiting factor for discrimination above ~ 12 keVr (nuclear recoil energy). Below 12 keVr statistical fluctuations in the number of scintillation photons counted dominate, and we project a x6 improvement in background rejection with a x2 increase in light collection efficiency.

We also build a simple recombination model that successfully reproduces the mean recombination in electron and nuclear recoils, including the surprising reversal of the expected trend for recombination with ionization density in low energy electron recoils. The model also reproduces the measured recombination fluctuations to within a factor of two at high energies. Surprisingly, the model suggests that recombination at low energies is independent of ionization density, and our observed discrimination is due not to the different stopping powers of electrons and nuclei as was thought, but rather to a different initial exciton-ion ratio. We suggest two possible physical models for this new result.

Acknowledgements

I'd like to thank first my advisor, Tom Shutt, without whom I might well have become a theorist. I am eternally grateful for your intervention, and for the many, many hours you spent with us in lab.

And to my labmate John Kwong — you say I helped keep you sane, I think you helped me go a little crazy. I hope our paths cross again. Watch out for five-legged bears up in Minnesota.

Thanks to all on the Xenon10 collaboration. To the PI's, Elena Aprile, Rick Gaitskell, Laura Baudis, Dan McKinsey, and Uwe Oberlack, it wasn't always pretty, but it was a grand endeavor. To Aaron Manalaysay, we very nearly cut our theses short up on the Corno. Let's climb it again, but not when there's lightning this time. And the rest of the Hall d'Montagio crew, Simon, Peter, Luiz, and Rick, who was ever with us in the trenches — there were some long nights and longer phone meetings, but you all made it worthwhile. And to all the others I worked with at Gran Sasso — Angel and Maria-Elena, I couldn't ask for better company at the pizzeria. Masaki and family, Alfredo, Kaixuan, Guillaume, Louis, and Roman, you all made it a great time. And for those of you still there, I'm glad you all survived the quake, and I hope you recover soon, especially Alfredo, although even with a broken leg you'd beat us on the football field.

To all my friends at Princeton, the times with you are my fondest memories of grad school. To Janice and Amin, I can't describe how glad I am to have known you — I hope Isabel knows how lucky she is. To Chris and Annika, Rob, Justin, Amol, Glen, Louis, Judy, and Jed, I love you all. And to Asad, Aaron, Mike, Phil and Erika, Matt, Ryan, Beth, Fiona, and Joe, if you take your cleats to conferences, I'll bring cones and a disc. Seriously.

Thanks to everyone I met at Case. To Alex Bolozdynya, who has forgotten more xenon lore than I'll ever know, it's been an honor to work with you. To John L., Zak, and Cathy, thanks for pulling us from our exile atop A.W. Smith. To Pavel, Adam, Jeremy, Sean, Mike, Nate, Sander, and Sashank, you all made our lab a better place, and not just because of the work you did. To Ken and Patrick, it was brief, but all the sweeter for its briefness. And to Dan Akerib and Mike Dragowsky, thanks for lending us Ken and Patrick. Best of luck to all the LUXers.

Thanks to Pat Rothfuss, for a few timely words from a master wordsmith, and to Cristian Galbiati, for a timely reading of the words of a lesser wordsmith.

And finally, I need to thank my parents, who have always offered love, support and encouragement, whichever path I chose. Whether that speaks to my choices or your tolerance or my selective memory, I'm not sure, but I think it's turned out all right. And to Christy, my beloved, how did I ever survive without you? You've stuck with me through it all, and I love you.

To Christy

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	x
List of Figures	xi
1 The Case for WIMP Dark Matter	1
1.1 Evidence for Dark Matter	2
1.1.1 Cosmological Constraints on Ω_m	2
1.1.2 Big Bang Nucleosynthesis	4
1.1.3 Galaxies and Galaxy Clusters	4
1.1.4 Structure Formation	6
1.2 The WIMP Miracle	7
2 WIMP Detection	10
2.1 WIMP recoil spectrum	10
2.2 Direct Detection Methods	13
2.2.1 Sources of Background	13
2.2.2 Annual Modulation	15
2.2.3 Active Shielding	16
2.2.4 Nuclear Recoil Discrimination	18
3 Xenon Time Projection Chambers	23
3.1 Basic Xenon TPC Operation	23
3.2 XenonTPCs as WIMP Detectors	25
3.2.1 WIMP sensitivity	25
3.2.2 Scalability	25

3.2.3	Chemical Purity	26
3.2.4	Radiopurity	26
3.2.5	Active Shielding	26
3.2.6	Discrimination	28
3.3	The Xed Prototype TPC	29
3.3.1	Cryostat	29
3.3.2	Internal Structure	30
3.3.3	Cleaning and Xenon Purification	32
3.3.4	Liquid Level Monitoring	34
3.3.5	Source Placement	34
3.3.6	Data Acquisition	35
3.3.7	Data Processing	35
3.3.8	Single-Phase Operation	36
3.3.9	Other Xed Configurations	38
4	Recombination and Calibration	39
4.1	Recombination Model for Signal Production in Liquid Xenon	39
4.1.1	Signal Production in Electron Recoils	40
4.1.2	Nuclear Recoils	42
4.1.3	Energy Scales	46
4.2	Calibrating Recombination Measurements	48
4.2.1	Method for Calibrating Using Charge-Light Anti-correlation	48
4.2.2	Single-phase, Separate Charge and Light	49
4.2.3	Single-phase, Simultaneous Charge and Light	52
4.2.4	Dual-phase, Simultaneous Charge and Light	55
4.2.5	Recombination versus Field	57
4.2.6	Absolute Light Calibration	60
4.3	Measuring Recombination Fluctuations	62
4.4	Conclusions on Recombination and Calibration	66
5	Recombination and Discrimination	67
5.1	Calibration of Gamma and Neutron Data	68
5.1.1	Determining Calibration Constants	68
5.1.2	Determining Bulk Instrumental Fluctuations	69

5.2	Characterizing Gamma and Neutron Bands	76
5.2.1	Centroid Estimates	76
5.2.2	Band Statistics	79
5.3	Analytical Band Width Decomposition	82
5.4	Monte-Carlo of Fluctuations	84
5.4.1	Building the Monte Carlo	85
5.4.2	Tuning the Monte Carlo	86
5.4.3	Band Decomposition Revisited	89
5.5	Electron Recoil Discrimination	89
5.5.1	Reproducing Xenon10 Discrimination	92
5.5.2	True Discrimination Calculations	93
5.5.3	Discrimination Projections	95
5.6	Conclusions on Recombination and Discrimination	98
6	Recombination Models	101
6.1	The Need for a Recombination Model	101
6.2	Existing Recombination Models	103
6.3	Construction of a Recombination Model	106
6.3.1	Electron recoil track Monte Carlo	107
6.3.2	Nuclear recoil track Monte Carlo	108
6.3.3	Modified Thomas-Imel Model	109
6.4	Comparing the Model to the Data	111
6.4.1	Electron Recoils	111
6.4.2	Nuclear Recoils	113
6.4.3	Combined Fits	113
6.4.4	Compound Recoils	114
6.4.5	Fluctuations	115
6.5	Interpretation of Fit Parameters	117
6.6	Conclusions on Recombination Physics	122
7	Xenon10 and Beyond — Searching for WIMPS with Large Scale Xenon TPCs	124
7.1	Fluid Handling	125
7.1.1	Liquid Level Control	125
7.1.2	Recirculation Rates	126

7.2	Data Handling	127
7.2.1	Data Compression	127
7.2.2	Data Processing	129
7.2.3	Position Reconstruction	132
7.3	Detector Calibration	133
7.3.1	Position Dependent Corrections	134
7.3.2	Absolute S1 and S2 Calibrations	134
7.3.3	Nuclear Recoil Measurements	137
7.3.4	Discrimination Measurements	139
7.4	WIMP Analysis	141
7.4.1	WIMP candidate selection	141
7.4.2	Non-Gaussian Leakage	142
7.4.3	Calculating a Dark Matter Limit	144
7.5	Conclusions on WIMP Hunting	147
A	Electrostatics of Wire Grids	149
B	Maximum Likelihood Histogram Fitting	153
C	RIVAL — A Nuclear Recoil Track Monte Carlo	156

List of Tables

5.1	Calibration constants for Xed dual-phase data	71
5.2	Instrumental fluctuations in the 40 keV inelastic peak in Xed at each field	76
6.1	Modified Box Model fit parameters for the electron and nuclear recoil bands at each field	112
6.2	Modified Box Model fit parameters for the electron and nuclear recoil bands at each field, taking a field independent $\frac{N_{ex}}{N_i}$ for nuclear recoils	115

List of Figures

1.1	Constraints on matter density and the cosmological constant from supernovae, the cosmic microwave background, and galaxy clusters	3
1.2	Primordial abundances of nuclei as a function of Ω_b from big bang nucleosynthesis	5
1.3	Image of the Bullet cluster (1E 0657-56)	6
1.4	A map of large scale structure from the Sloan Digital Sky Survey data	7
1.5	Densities of thermal relics by annihilation cross section	8
2.1	Sensitivity of xenon, germanium, and argon to spin-independent WIMP scattering	14
2.2	Annual modulation in the spin-independent WIMP recoil spectrum in NaI	15
2.3	The annual modulation result from DAMA	16
2.4	Schematic of the 800 kg XMASS detector	18
2.5	Electron recoil discrimination in CDMS	19
2.6	Electron recoil discrimination in DEAP	20
2.7	Electron recoil discrimination in WARP	21
2.8	Muon, neutron, and alpha/WIMP events in COUPP	22
3.1	Electron drift velocity in gas and liquid xenon	24
3.2	Krypton removal using gas chromatography	27
3.3	Gamma and neutron interaction lengths in liquid xenon	28
3.4	Single scatter background rate as a function of position in a liquid xenon detector	29
3.5	The Xed cryostat	30
3.6	The Xed detector	31
3.7	Picture of the Xed active region	32
3.8	Picture of the Xed cryostat and gas handling system	33
3.9	Liquid level monitoring in Xed, showing the periodic growth of bubbles	34
3.10	Example nuclear recoil event in Xed	36

4.1	Published data on nuclear recoil scintillation yields in liquid xenon	45
4.2	Lindhard factor inferred from \mathcal{L}_{eff} and Xed data	46
4.3	Nuclear recoil scintillation yield (\mathcal{L}_{eff}) inferred from Lindhard theory and Xed data	47
4.4	Light yield versus charge yield for 122 keV ^{57}Co gamma events, scanning over drift field	50
4.5	Spec-amp spectrum for single-phase charge-only ^{57}Co data	51
4.6	Mean pre-amp waveform for single-phase charge-only ^{57}Co data	52
4.7	Single phase, simultaneous charge and light measurement of the ^{57}Co source	54
4.8	Charge and light background profiles for single-phase dual-channel data	55
4.9	Decomposition of fluctuations in dual-phase measurements of the 122 keV ^{57}Co peak	57
4.10	Normalized charge yield versus field for 122 keV ^{57}Co gammas, 5.3 MeV ^{210}Po alphas, and 56.5 keV Xe recoils	58
4.11	Electronic stopping power versus energy for electrons, alphas, and Xe nuclei in liquid xenon	59
4.12	Single photo-electron distribution — height versus integral	60
4.13	Single photo-electron spectrum and fit	61
4.14	S1, S2, and recombination fluctuations in the 122 keV ^{57}Co peak versus normalized charge yield	65
5.1	Gamma, neutron, and calibration data in Xed at 876 V/cm, showing electron and nuclear recoil bands and the 40 keV inelastic and 122 keV photo-absorption peaks .	70
5.2	Electron recoil band profiles at low and high energy for each source, as used for S2 calibration adjustments	71
5.3	Gamma data in Xed at 4060, 1951, 522, and 60 V/cm, showing the electron recoil band	72
5.4	Neutron data in Xed at 4060, 1951, 522, and 60 V/cm, showing the nuclear recoil band and 40 keV inelastic peak	73
5.5	Diagram of inelastic neutron scatter and photo-absorption events	74
5.6	Energy spectra of electron and nuclear recoils in Xed from the ^{133}Ba and ^{252}Cf sources	75
5.7	Centroid-subtracted electron and nuclear recoil bands at 876 V/cm	78
5.8	Electron and nuclear recoil band widths, means, and separation in Xed	80
5.9	Estimated inflation in the measured variance in Xed due to centroid uncertainty . .	81
5.10	Illustration of S1 statistical fluctuations in the electron and nuclear recoil bands in Xed at 876 V/cm	83

5.11 Analytical decomposition of electron and nuclear recoil bands into constituent fluctuations at 876 V/cm	84
5.12 Monte-Carlo tuning results — mean recombination versus energy for electron and nuclear recoils at each field	87
5.13 Monte-Carlo tuning results — recombination fluctuations versus energy for electron and nuclear recoils at each field	88
5.14 Monte-Carlo decomposition of electron recoil band widths at each field	90
5.15 Monte-Carlo decomposition of nuclear recoil band widths at each field	91
5.16 Simulated discrimination in the Xenon10 detector, using the Monte Carlo tuned by Xed, compared to the observed discrimination	93
5.17 Simulated WIMP and background data in Xed and in an ideal detector with twice the Xenon10 light collection	94
5.18 Projected Leakage Fraction in Xed, Xenon10, ideal detectors with x2 and x3 the Xenon10 light collection, and an ideal detector with 100% light collection, at 876 V/cm	96
5.19 Projected Relative Leakage Rate in Xed, Xenon10, ideal detectors with x2 and x3 the Xenon10 light collection, and an ideal detector with 100% light collection, at 876 V/cm	97
5.20 Projected Leakage Rate relative to Xenon10 in Xed, ideal detectors with x2 and x3 the Xenon10 light collection, and an ideal detector with 100% light collection, at 876 V/cm	98
5.21 Projected Relative Leakage Rate in an ideal detector with x2 the Xenon10 light collection, at 4060, 1951, 876, 522, and 60 V/cm	99
6.1 Electron recoil tracks simulated using Penelope	108
6.2 Mean and variation in recoil track size as a function of energy for electron and nuclear recoils	109
6.3 Nuclear recoil tracks simulated using RIVAL	110
6.4 Modified Box Model fits to the electron and nuclear recoil band centroids	112
6.5 Modified Box Model fits to the electron and nuclear recoil band centroids, taking a field independent $\frac{N_{\text{ex}}}{N_i}$	114
6.6 40 keV inelastic and 122 keV photo-absorption peak reconstruction at 876 V/cm using the modified box model	116
6.7 40 keV inelastic peak reconstruction at 4060, 1951, 522, and 60 V/cm using the modified box model	117
6.8 Modified Box Model predictions for recombination fluctuations at all fields	118

6.9	Modeled recombination distribution at 80 keVee	119
6.10	Scaling of box size with drift field	120
6.11	Electron temperature versus drift field, compared with the Langevin interpretation of our recombination model	121
7.1	Bell and weir methods for liquid level control	126
7.2	Example heat exchanger for recirculation	128
7.3	Coherent PMT noise and baseline flattening in Xenon10	131
7.4	Fast S2 filter example	132
7.5	Simulated S2 hit pattern in Xenon10	133
7.6	Position dependence of S1 light collection in Xenon10	135
7.7	Position dependence of S2 light production and collection in Xenon10	136
7.8	Single electron S2's in Xenon10	137
7.9	40, 164, and 236 keV nuclear excitation peaks in Xenon10	138
7.10	Nuclear and electron recoil bands in Xenon10	138
7.11	Multiple neutron scatters in Xenon10	139
7.12	Electron recoil band profile in Xenon10	140
7.13	Reverse field region Gamma-X events in Xenon10	143
7.14	Resistor chain Gamma-X events in Xenon10	143
7.15	Xenon10 WIMP search results — energies, charge yields, and positions of WIMP candidate events	144
7.16	Xenon10 limit at 90% C.L. on the spin independent WIMP-nucleon cross section . .	148
A.1	Diagram of wire grid between two plates (electrostatics problem)	150

Chapter 1

The Case for WIMP Dark Matter

The nature of dark matter is one of the longest standing fundamental problems in physics. The first evidence for dark matter came in 1933 when Zwicky measured the velocity dispersion of objects at the edge of the Coma cluster [1], and in the years since, measurements of galactic rotation curves [2], gravitational lensing [3], the cosmic microwave background [4], and even collisions between clusters of galaxies [5] all support the existence of a large invisible matter component to the universe. Our understanding of nucleosynthesis in the big bang indicates that the majority of this missing matter is non-baryonic [6], while our understanding of structure formation rules out relativistic particles such as neutrinos [7]. In short, no particle in the Standard Model can account for the extra mass demanded by these observations.

When we look beyond the standard model for dark matter candidates, one compelling class of particles are WIMPs, or weakly interacting massive particles. Predicted by many supersymmetric extensions of the standard model, WIMPs interact via the weak force like neutrinos, but can have masses from 40 GeV–3.2 TeV [8]. Remarkably, the particle physics describing WIMPs naturally gives a relic WIMP density matching that of dark matter. If WIMPs are the dark matter, then they should be present in our own galaxy, with local densities of a few particles per liter and velocities of several hundred km/s [9].

In this chapter we review the evidence for non-relativistic, non-baryonic dark matter, and describe how WIMPs naturally emerge with the required dark matter density. In Chapter 2 we describe how WIMPs may be detected on earth as they scatter off target nuclei, and we summarize the current experimental efforts to detect or rule out a WIMP signal. Chapter 3 discusses dual-phase xenon time projection chambers, as used in the ZEPLIN, XENON, and LUX experiments, and describes

the instrumentation of our prototype detector Xed. Chapters 4–6 cover our measurements in Xed on the physics of signal production in liquid xenon, with particular attention to the background discrimination capability of these detectors. The final chapter discusses WIMP searches in large xenon time projection chambers, ending with the Xenon10 upper limit on the spin-independent WIMP-nucleon cross section.

1.1 Evidence for Dark Matter

As stated above, there is evidence for dark matter at all scales, from the cosmological to single galaxies. While alternatives to dark matter such as MOND (modified newtonian dynamics) may satisfy one or even many of these observations [10], the overall evidence is overwhelmingly in favor of a universe full of dark matter.

1.1.1 Cosmological Constraints on Ω_m

The mass density in the universe may be determined from a combination of measurements of cosmic expansion and the cosmic microwave background, as shown in Fig. 1.1. In a universe composed of non-relativistic matter and a cosmological constant (known as the Λ CDM model), the Friedmann equations may be written

$$\frac{H^2}{H_0^2} \equiv \frac{1}{H_0^2} \left(\frac{\dot{a}}{a} \right)^2 = \Omega_m \left(\frac{a_0}{a} \right)^3 + \Omega_k \left(\frac{a_0}{a} \right)^2 + \Omega_\Lambda, \quad (1.1)$$

and

$$\frac{1}{H_0^2} \frac{\ddot{a}}{a} = -\frac{1}{2} \Omega_m \left(\frac{a_0}{a} \right)^3 + \Omega_\Lambda, \quad (1.2)$$

where Ω_m and Ω_Λ are the energy densities today of matter and the cosmological constant relative to the critical density $\frac{3H_0^2}{8\pi G}$, H_0 is the Hubble constant today (70.8 (km/s)/Mpc), H and a are the Hubble constant and scale factor at time t , and Ω_k is a term for the spatial curvature of the universe. For open, closed, and flat universes, Ω_k is greater than, less than, and equal to zero, respectively. The relative densities are constructed such that $\Omega_m + \Omega_k + \Omega_\Lambda = 1$, as may be seen by evaluating Eq. (1.1) at $t = t_0$ (today).

Measuring the acceleration of the universe lets us relate Ω_m and Ω_Λ using Eq. (1.2). These measurements are made by measuring redshift as a function of distance, where distances are determined by use of a standard candle (such as supernova IA's [11]) or standard ruler (such as from baryon

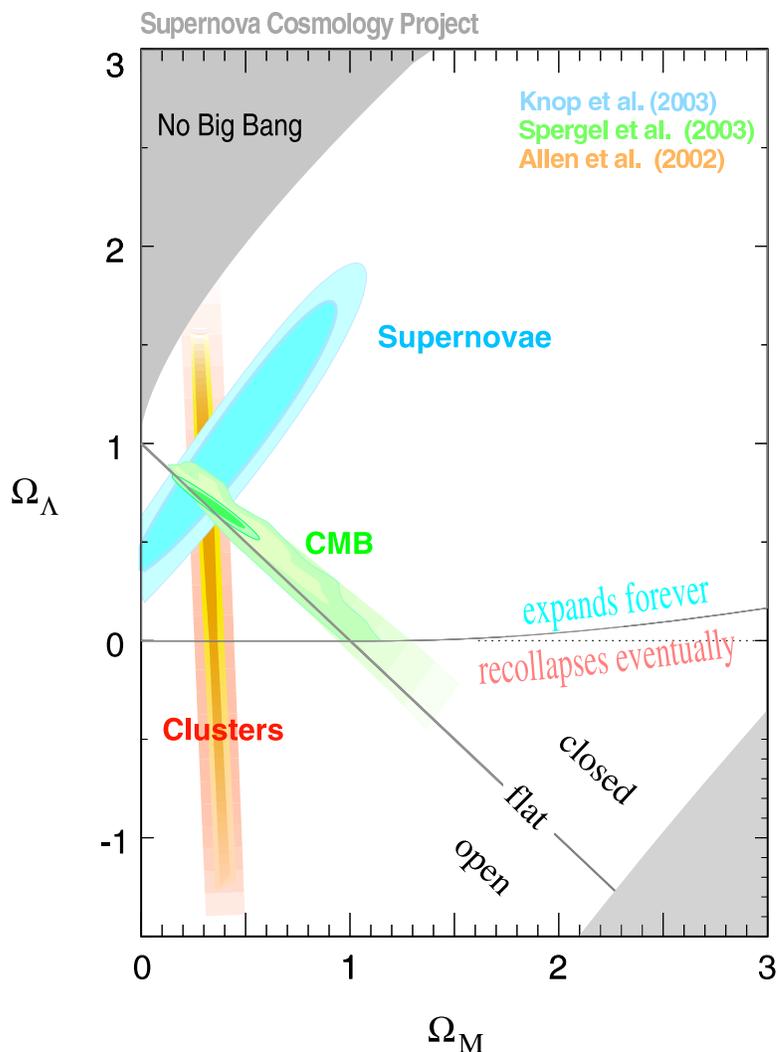


Figure 1.1: Constraints on matter density and the cosmological constant from supernovae, the cosmic microwave background, and galaxy clusters (with big bang nucleosynthesis). Data from [11, 12, 13], figure taken from [14].

acoustic oscillations [15]), with the resulting constraint on Ω_m and Ω_Λ shown in Fig 1.1.

A second constraint comes from anisotropies in the cosmic microwave background (CMB). The first peak in the CMB power spectrum gives the curvature of the universe, which is measured to be flat ($\Omega_k \approx 0$), so we have $\Omega_m + \Omega_\Lambda \approx 1$. Combined with expansion data, this picks out $\Omega_m = 0.28$ and $\Omega_\Lambda = 0.72$ (see Fig. 1.1). Further analysis of the anisotropy power spectrum gives an independent measurement of Ω_m that is consistent with these results [4].

1.1.2 Big Bang Nucleosynthesis

We need to compare our value of Ω_m with the density of ‘normal’ baryonic matter, Ω_b . This can also be extracted from the CMB power spectrum [4], but an alternate method comes from big bang nucleosynthesis [6]. In the very early universe ($t < 3$ min), neutrons and protons are free and in thermal equilibrium. At $t = 3$ min the universe has cooled to the point where nuclei begin to form. The nuclear composition of the universe evolves according to the cross-sections of various nuclear interactions and the thermodynamics of the expanding universe, until the Hubble expansion rate exceeds the nuclear interaction rate. At this point the existing nuclei “freeze-out”, their abundances fixed until nuclear processes restart with stellar formation. The primordial abundances are sensitive to the temperature (time) at which freeze-out occurs, and thus to the baryon density in the early universe. The overall matter density is not important — the universe is radiation dominated at this stage, so non-baryonic matter affects neither the expansion rate nor the nuclear interaction rate. The primordial abundances of ^4He , ^3He , D (^2H), and ^7Li are shown in Fig. 1.2 as a function of Ω_b , along with the measured primordial abundances based on observations of present day abundances [16, 17]. The value of Ω_b given by the measured primordial abundances is consistent with that determined from the microwave background, giving $\Omega_b = 0.046$ and leaving $\Omega_\chi = 0.23$ for non-baryonic dark matter.

1.1.3 Galaxies and Galaxy Clusters

The best measurements of matter densities in clusters come from a combination of x-ray and weak lensing observations, which measure the baryonic and total mass of the cluster, respectively. In weak lensing, one looks at the deformation of galaxies behind the cluster in question. In contrast to strong lensing, where a specific object is identified as being lensed, producing multiple images or gross distortion, weak lensing measures distortions that are smaller than the normal variation in the unlensed object’s appearance — thus it is impossible to say how lensed a single object is. By looking at patterns of distortion in large numbers of lensed objects, however, one can extract the mass of the lens. An overview of weak lensing can be found in [3].

The baryonic mass of a large cluster is dominated by x-ray emitting intracluster gas, and so the total baryonic mass can be estimated from the x-ray intensity [13]. Combined with the total mass given by weak lensing, this gives the ratio of baryonic to dark matter in the cluster. Taking the baryon density from big bang nucleosynthesis and assuming that the matter in the cluster is representative of the matter make-up of the universe, one finds the total mass density. The Ω_m picked

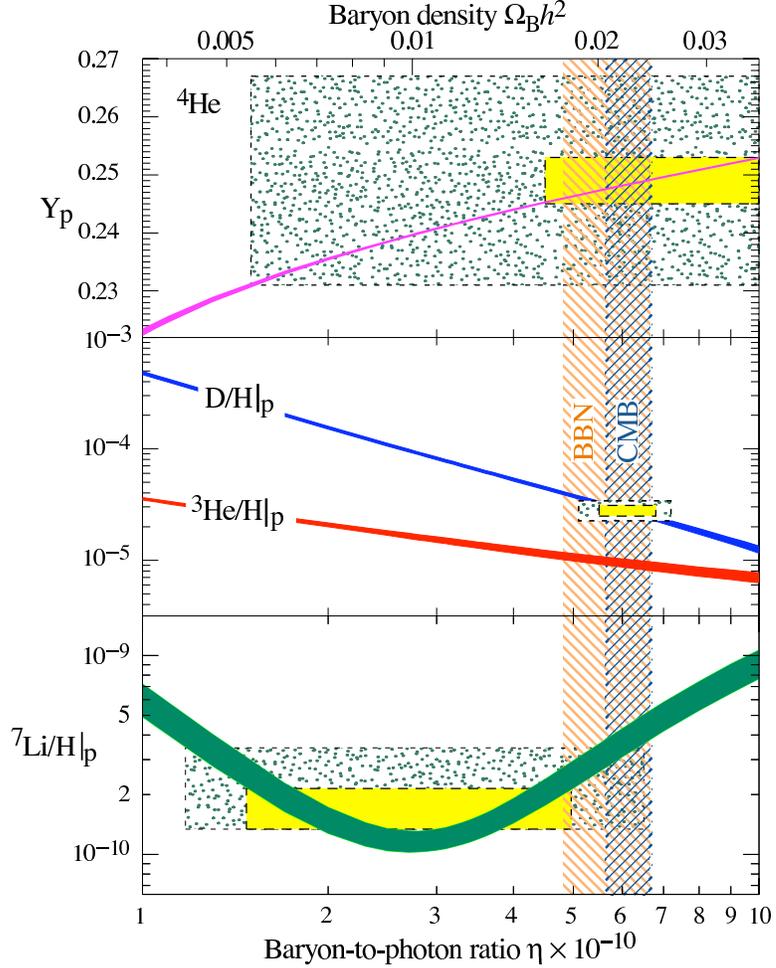


Figure 1.2: Primordial abundances of nuclei as a function of Ω_b according to big bang nucleosynthesis, with observed abundances shown as boxes. The Ω_b given by the cosmic microwave background is also shown. Figure taken from [6].

out by this technique agrees with that determined from expansion and microwave background data, as shown in Fig. 1.1.

There are many methods of measuring dark matter densities based on the dynamics of gravitationally bound objects, including measurements galactic rotation curves [2] and the velocity dispersion of galaxies in clusters [19]. These and all the previous measurements, however, could potentially be explained by modifications to gravity at large length scales. This is not the case for the most recent evidence for dark matter — the Bullet cluster (1E 0657-56) is a recent collision between two galaxy clusters, shown in Fig. 1.3. As before, the baryonic mass and total mass of the colliding clusters are measured using x-ray observations and gravitational lensing, but now the key information is the location of the lens relative to the gas. During the collision, the x-ray gas experiences drag

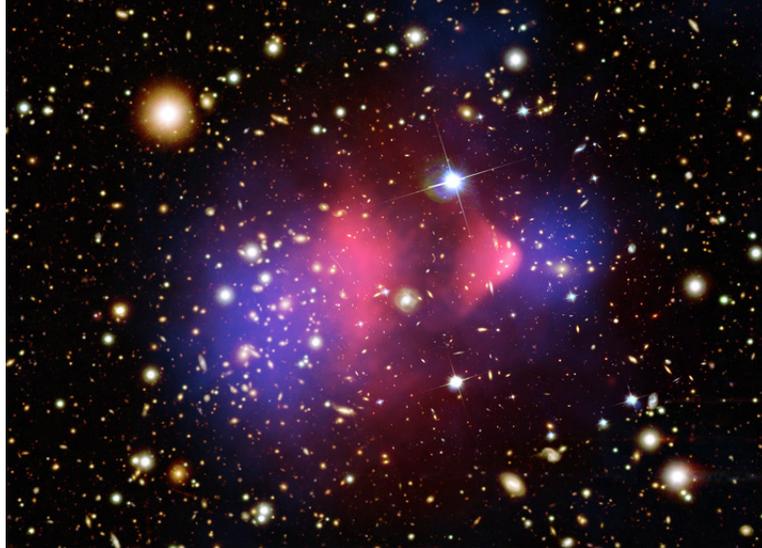


Figure 1.3: The Bullet cluster (1E 0657-56). The blue shows the location of mass, as determined by gravitational lensing, while the red shows the location of the x-ray gas [18]. The x-ray gas experienced drag and formed a shock wave during the collision, but the dark matter passed through unhindered.

forces, forming a visible shock-wave and lagging behind the unimpeded dark matter. The center of gravitational mass and the center of visible mass are now significantly displaced from each other, a feat beyond the reach of modified gravity schemes [5].

1.1.4 Structure Formation

The final piece of evidence for dark matter, as well as a clue towards the nature of dark matter, comes from considering structure formation. Structures grow as overdense regions draw in surrounding matter, creating a positive feedback for primordial density fluctuations. This structure is believed to have grown from the bottom up, with small structures (galaxies) forming first, followed later by the growth of galaxy groups, clusters, and super clusters. Experiments such as the Sloan Digital Sky Survey have revealed the extent of large scale structure in the universe to extend even beyond the super cluster size (see Fig. 1.4). Such large scale structures today require an early start to structure formation. Structure formation due to baryons is suppressed until electron-proton combination occurs, since until this point the photon pressure works against the growth of density perturbations [20]. This does not affect dark matter particles, which may begin forming structures long before combination occurs. Not only does structure formation require dark matter, it specifically requires non-relativistic dark matter. Relativistic particles tend to smooth density perturbations, both delaying the growth of the first galaxies and increasing the physical size of these structures when they

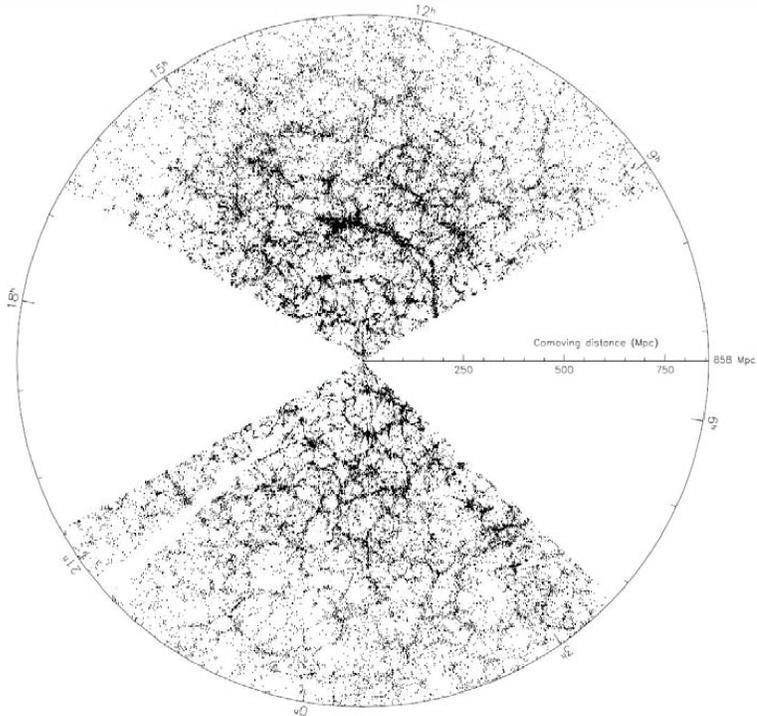


Figure 1.4: A map of large scale structure from the Sloan Digital Sky Survey data. Figure taken from [21].

eventually form [7].

1.2 The WIMP Miracle

We now have a basic recipe for the universe: 4.6% baryons, 23% non-relativistic, non-baryonic dark matter, and 72% vacuum energy (or dark energy). Thus 95% of the universe is unknown, but there are several candidates for the dark matter piece. The two primary candidates are axions and weakly interacting massive particles, or WIMPs (for a discussion of axions, see [22]). The WIMP hypothesis follows from the hypothesis that dark matter is a thermal relic of the big bang, as described in [23]. Let us consider a particle χ with mass M_χ and annihilation cross section σ_a . In the very early universe, when $T \gg M_\chi$, these particles will have the same abundance as any light particle in thermal equilibrium (modulo the number of internal degrees of freedom). As the universe cools to $T < M_\chi$, the equilibrium number of particles n_χ^{eq} falls. The number density then follows

$$\frac{dn_\chi}{dt} = -\langle\sigma_a v\rangle \left[(n_\chi)^2 - (n_\chi^{eq})^2 \right] - 3Hn_\chi, \quad (1.3)$$

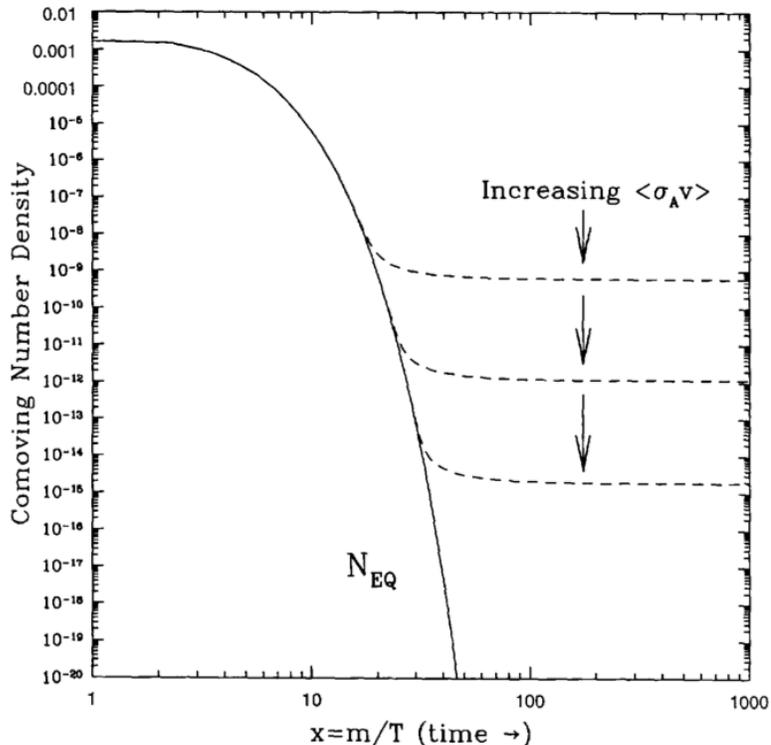


Figure 1.5: Densities of particles in thermal equilibrium and after freezout. The solid line shows the equilibrium density falling as the universe cools, with dotted lines showing the freezout densities for various annihilation cross-sections. Figure taken from [23], originally appearing in [20].

where $\langle\sigma_a v\rangle$ is the thermally averaged velocity times annihilation cross section. The first term on the right hand side of Eq. (1.3) shows that n_χ approaches n_χ^{eq} at a rate determined by the annihilation cross-section times the particle flux, while the second term gives the drop in density due to expansion. The expansion H is determined independently since the early universe is radiation dominated. When annihilation begins the first term is dominant, but as the universe expands and the particle density falls, the expansion term begins to dominate. At this point annihilation stops, and the remaining particles freeze out with a constant co-moving density. Following through the thermodynamics and expansion physics (see [23] for a complete explanation), one finds that the mass density at freeze-out is independent of M_χ and inversely proportional to $\langle\sigma_a v\rangle$, as shown in Fig. 1.5. If the particle is stable, then this relic density will still exist today.

Now the WIMP miracle occurs. If we take an annihilation cross section corresponding to the weak interaction in particle physics [23],

$$\langle\sigma_a v\rangle = \alpha^2 (100\text{GeV})^{-2} \sim 10^{-25}\text{cm}^3\text{s}^{-1}, \quad (1.4)$$

we come to a relic density of $\Omega_\chi \sim 0.1$, within an order of magnitude of what we need for dark matter. In other words, if a stable WIMP exists, it will be present as dark matter at or near the density we require. As a bonus, many supersymmetric extensions of the standard model predict a stable WIMP. In particular, models that conserve R-parity (originally invoked to prevent proton decay) have a stable LSP (lightest supersymmetric particle). This particle is often a partner to the neutral gauge bosons, known as the neutralino [23]. This, or any other WIMP, would at long last solve the dark matter problem.

Chapter 2

WIMP Detection

There are two paths to testing the WIMP hypothesis: direct and indirect detection. In indirect detection, one looks for signatures of WIMPs annihilating in the galactic center, the sun, or in nearby dark-matter dominated galaxies. These signatures include mono-energetic gammas [24], and an excess of neutrinos [25] and cosmic anti-matter [26]. We will focus on direct detection, where WIMPs scattering off of a target nucleus produce measurable nuclear recoils. In the first part of this chapter we calculate the WIMP recoil spectrum a function of the target nucleus and WIMP mass. We then review existing direct detection experiments according to their techniques for distinguishing the WIMP signal from background radiation.

2.1 WIMP recoil spectrum

In this section we follow the discussion of WIMP recoil spectra by Lewin and Smith [9]. The shape of the WIMP recoil spectrum is determined by the WIMP mass M_χ , the target nucleus mass M_T , a nuclear form factor F , and the distribution of WIMP velocities about the target. We write the last as

$$dn = \frac{n_0}{k} f(\mathbf{v}, \mathbf{v}_E) d^3\mathbf{v}, \quad (2.1)$$

where n_0 is the WIMP number density $\frac{\rho_0}{M_\chi}$ (with ρ_0 the local dark matter density), k is a normalization factor such that $\int dn = n_0$, \mathbf{v} is the WIMP velocity with respect to the target, and \mathbf{v}_E is the velocity of the target (earth) in the rest frame of the galaxy. For $\mathbf{v}_E = 0$, we take \mathbf{v} to follow a Maxwell-Boltzmann distribution,

$$f(\mathbf{v}, 0) = e^{-\frac{v^2}{v_0^2}}, \quad (2.2)$$

truncated at the galactic escape velocity v_{esc} . A non-zero \mathbf{v}_E translates this distribution so that $\langle \mathbf{v} \rangle = -\mathbf{v}_E$. Independent of \mathbf{v}_E , we have

$$k = (\pi v_0^2)^{\frac{3}{2}} \left[\operatorname{erf} \left(\frac{v_{esc}}{v_0} \right) - \frac{2}{\sqrt{\pi}} \frac{v_{esc}}{v_0} e^{-\frac{v_{esc}^2}{v_0^2}} \right]. \quad (2.3)$$

The interaction rate per target mass is given by the number of target nuclei per unit mass M_T^{-1} , the scattering cross section per nucleus σ_T , and the WIMP flux vdn ,

$$dR = \frac{1}{M_T} \sigma_T v d n. \quad (2.4)$$

Integrating Eq. (2.4) over WIMP velocities gives the total WIMP scattering rate. To find the differential spectrum, we begin by assuming isotropic scattering in the center-of-mass frame. In terms of the center-of-mass scattering angle θ_{cm} , the recoil energy is given by

$$E_r = M_\chi v^2 (1 - \cos \theta_{cm}) \frac{M_T M_\chi}{(M_T + M_\chi)^2}. \quad (2.5)$$

Isotropic scattering gives a uniform distribution in $\cos \theta$, so for a given WIMP velocity we see a flat recoil spectrum from 0 to $E_{r,max} = \frac{1}{2} M_\chi v^2 \frac{4M_T M_\chi}{(M_T + M_\chi)^2}$. Putting together the pieces thus far, we have a differential recoil spectrum of

$$\frac{dR}{dE_r} = \frac{(M_T + M_\chi)^2}{2M_T^2 M_\chi^2} \frac{\rho_0}{M_\chi} \frac{\sigma_T}{k} \int d^3 \mathbf{v} \frac{\Theta(v_{esc} - |\mathbf{v} + \mathbf{v}_E|) \Theta(E_{r,max} - E_r)}{|\mathbf{v}|} e^{-\frac{|\mathbf{v} + \mathbf{v}_E|^2}{v_0^2}}, \quad (2.6)$$

where the two theta functions limit the integral to velocities present in the WIMP halo and capable of producing a recoil of energy E_r . If we take $v_E = 0$ and $v_{esc} = \infty$, this reduces to

$$\frac{dR(v_E=0, v_{esc}=\infty)}{dE_r} = \frac{1}{M_T} \frac{\rho_0}{M_\chi} \frac{2v_0 \sigma_T}{\sqrt{\pi}} \frac{1}{E_0 r} e^{-\frac{E_r}{E_0 r}} \quad (2.7)$$

where $E_0 = \frac{1}{2} M_\chi v_0^2$ is the mean WIMP energy and $r = \frac{4M_T M_\chi}{(M_T + M_\chi)^2}$. The basic recoil spectrum (before considering v_{esc} , v_E , or the upcoming nuclear form factor) is simply a falling exponential with mean energy $E_0 r$, which is maximized when $M_\chi = M_T$. Lewin and Smith take $\rho_0=0.4 \text{ GeV}/c^2 \text{ cm}^3$, $v_0=230 \text{ km/s}$, $v_{esc}=600 \text{ km/s}$, and

$$v_E = (244 + 15 \sin(2\pi y)) \text{ km/s}, \quad (2.8)$$

where y is (approximately) the time in years since March 2nd [9].

We still need an expression for σ_T , which includes both a coherence term and a nuclear form factor. When the momentum-transfer $q = \sqrt{2M_T E_r}$ is small, so that $\frac{\hbar}{q}$ is much larger than the size of the target nucleus, the scattering is coherent over the entire nucleus. For spin-independent interactions (scalar couplings), where all nucleons contribute to the scattering amplitude with the same sign, this gives a factor of A^2 in the cross section, where A is the number of nucleons in the nucleus. (This assumes identical interactions for neutrons and protons — scalings closer to $(A - Z)^2$ may be seen for spin-independent interactions in a given model. For comparing results between different target materials, these scalings are very similar and we adopt the A^2 scaling.) The cross section also scales like the square of the reduced mass, giving

$$\sigma_T(q=0) = A^4 \frac{(M_n + M_\chi)^2}{(AM_n + M_\chi)^2} \sigma_n, \quad (2.9)$$

where M_n is the nucleon mass, σ_n is the scattering cross section on a single nucleon, and we've taken $M_T = AM_n$. The $q=0$ indicates that this is in the coherent limit.

For spin-dependent interactions (axial couplings), nucleons with opposite spins contribute with opposite signs, giving zero amplitude for paired nucleons. Thus, only nuclei with odd numbers of neutrons and/or protons are sensitive to the spin-dependent cross-section, and there is no A^2 bonus. Comparing spin-dependent interactions between targets is more complicated, depending on the spin structure of the nucleus. See [9] for a review of spin-dependent cross sections.

For higher energy recoils, the distribution of scatterers in the nucleus becomes important, effectively negating the A^2 coherent bonus in spin-independent scattering. We write the reduced cross-section in terms of a form factor,

$$\sigma_T(q) = \sigma_T(q=0) F^2(q), \quad (2.10)$$

where, for spin-independent scattering, $F(q)$ is the Fourier transform of the scatterer (mass) density of the nucleus. Helm gives a convenient analytic form factor for a nucleus with constant density inside radius R_0 convoluted with a Gaussian of width s [27]

$$F(q) = \frac{3j_1(qR_0)}{qR_0} e^{-\frac{(qs)^2}{2}}, \quad (2.11)$$

where j_1 is the spherical Bessel function. Lewin and Smith suggest $s = 0.9$ fm and

$$R_0 = \sqrt{(1.23A^{1/3} - 0.60)^2 + 2.18} \text{ fm.} \quad (2.12)$$

The spin-dependent scattering form factor is again more complicated — see [9, 28] for a discussion of spin-dependent form factors.

Putting together Eq. (2.3,2.6,2.9,2.10,2.11), we find the WIMP recoil spectrum from spin-independent scattering in three common target materials: germanium ($A=73$), xenon ($A=131$), and argon ($A=40$) (see Fig. 2.1). All give featureless exponential spectra, with the highest low-energy scattering rate in xenon. The xenon rate falls more quickly with recoil energy than the others due to loss of coherence over the large nucleus.

2.2 Direct Detection Methods

The primary challenge for any direct detection experiment is to distinguish the WIMP signal from radioactive backgrounds. There are three basic methods for accomplishing this. First, the annual modulation in v_E results in an annual modulation in the WIMP spectrum, which may distinguish the WIMP signal from background. Second, WIMPs will scatter only once in the target volume, while other penetrating particles (gammas and neutrons) can scatter multiple times. Finally, many backgrounds deposit energy in target electrons, while WIMPs deposit energy in target nuclei. Detectors may respond differently to these two recoil types, giving what can be very strong background discrimination.

2.2.1 Sources of Background

Background events in WIMP detectors can come either from penetrating particles originating outside the target, or short-range particles from decays in the bulk or on the surfaces of the target. The former include high energy gammas and neutrons producing Compton scatters and nuclear recoils in the detector, respectively. These backgrounds can be reduced by shielding and selection of low-radioactivity detector components. Neutrinos will also become an irreducible background in this category as detectors improve in sensitivity.

Short range bulk and surface backgrounds include beta decays, x-rays following electron capture decays, internal conversion electrons, and alphas with the associated nuclear recoil. The most problematic are betas, since these have a broad energy spectrum covering the WIMP region, while

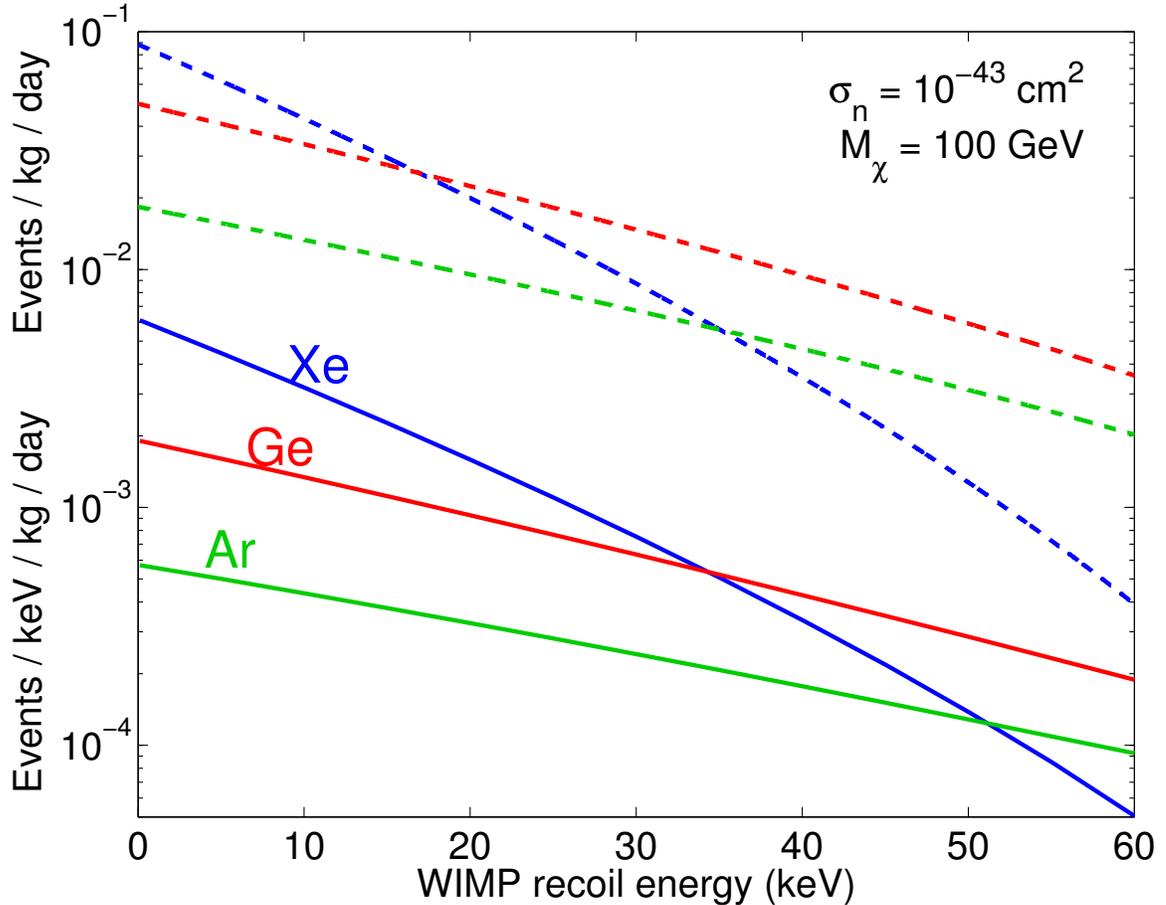


Figure 2.1: The sensitivity of xenon, germanium, and argon to spin-independent WIMP scattering, taking a 100 GeV WIMP and a 10^{-43}cm^2 scattering cross section for a single nucleon. The solid lines show the differential rate (events / keV / kg / day) and the dashed show the integrated rate over the detection threshold threshold (events / kg / day). For detection thresholds below 17 keV, xenon is the most sensitive of the three target materials.

the others have characteristic energies for a given decay that make them identifiable and potentially put them outside the range of interest. Betas with no associated gammas, i.e. betas to the ground state of a stable nucleus, are known as ‘naked’ betas and are particularly troublesome. Alphas are a background only in threshold detectors, since alpha energies (~ 5 MeV) are orders of magnitude higher than WIMP recoil energies. The recoiling nuclei in an alpha decay have energies of ~ 100 keV, at the high-end of the WIMP spectrum. If the alpha emitter is on the surface of the target and the alpha is emitted outward, only the nuclear recoil is measured, giving a potential background. Bulk and surface backgrounds must be removed via purification (sometimes requiring isotopic separation) of the target medium.

After reducing the backgrounds that can be reduced, it is necessary to distinguish the residual

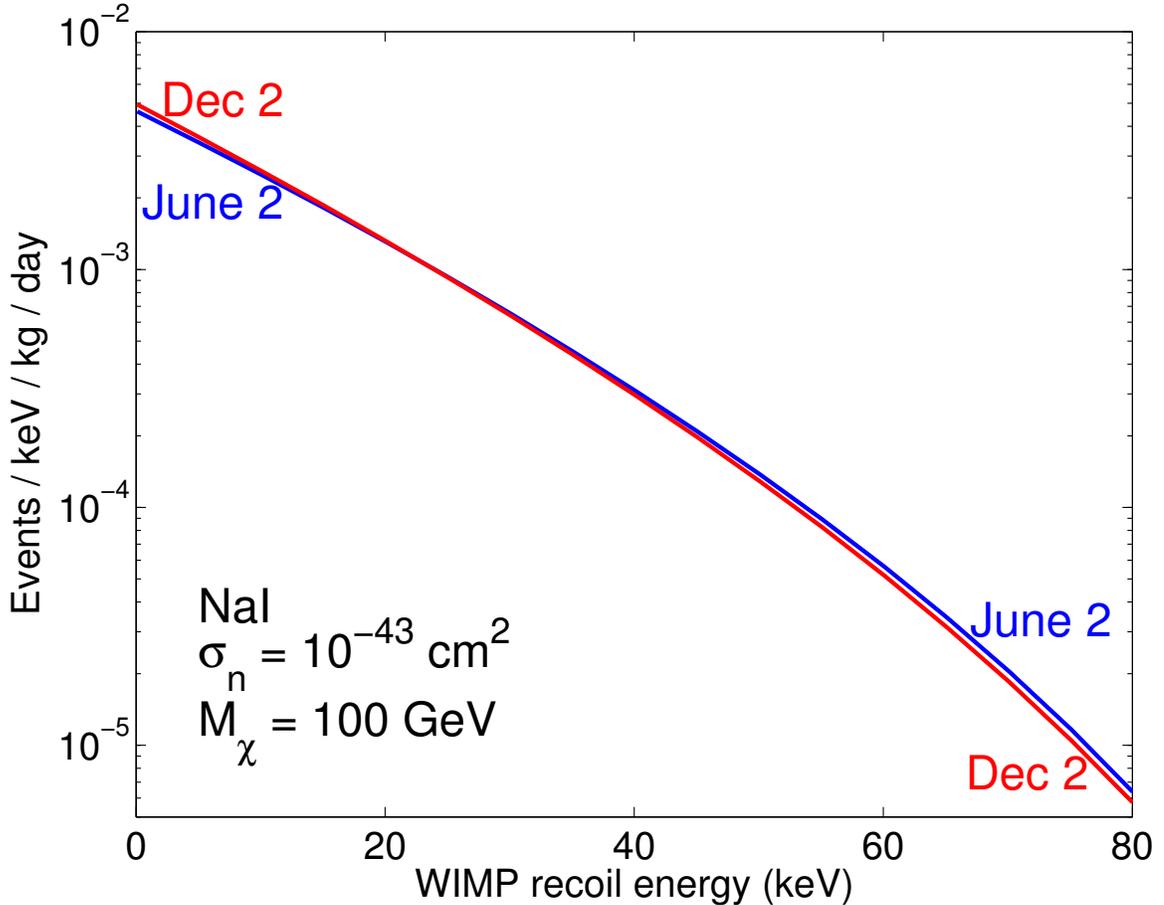


Figure 2.2: The expected annual modulation in the spin-independent WIMP recoil spectrum in NaI. At the lowest energies, there is a $\sim 6\%$ higher scattering rate on Dec 2 (lowest velocity relative to WIMP halo) than on June 2 (highest velocity). The spectra cross at 23 keV, and by 80 keV there is a 10% higher rate in June than in December.

background from a WIMP signal. The following three sections discuss the three basic methods for achieving this.

2.2.2 Annual Modulation

The expected annual modulation in the WIMP spectrum based on the analysis in Section 2.1 is shown in Fig. 2.2. The DAMA/LIBRA experiment uses 250 kg of highly radiopure NaI(Tl) scintillator to look for this modulation signal [29]. DAMA has reported a modulation signal (see Fig. 2.3), but the modulation is incompatible with the limits set by other dark matter experiments in the context of standard WIMP models [30, 31, 32]. It may be that DAMA sees an annual modulation in their background — in particular, the modulation is observed at the same energy as the argon Auger electron that follows a ^{40}K electron capture. The question of the origin of the DAMA

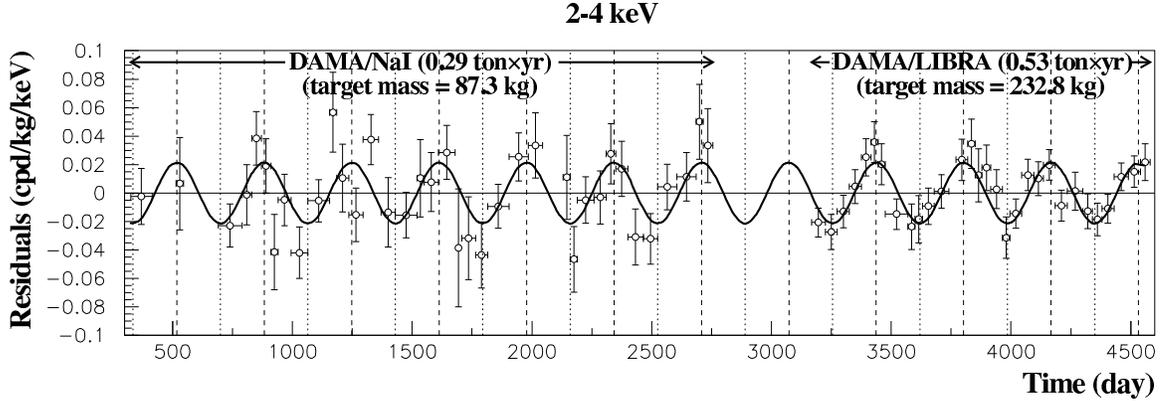


Figure 2.3: The annual modulation result from DAMA. The superimposed sine waves have maxima at Dec 2 (minima at June 2). Figure taken from [29].

signal is unresolved, reflecting the difficulty understanding backgrounds and systematic errors in this technique.

2.2.3 Active Shielding

It is preferable to reject backgrounds on an event-by-event basis. One technique to do so is to use 3-D position reconstruction to reject neutrons and gammas that scatter multiple times in the detector volume. This is known as ‘active shielding’, as the outer layer of the detector serves to tag events originating from outside the detector, while the inner or ‘fiducial’ volume is used for WIMP detection. That is, any event with a scatter in the fiducial volume and no scatters in the outer volume is a WIMP candidate.

Another way to describe this technique is to examine the rate of single-scatter events from penetrating particles as a function of position in the detector. The probability distribution for the distance between scatters of a penetrating particle is given by

$$P(x) = \frac{1}{x_0} e^{-\frac{x}{x_0}}, \quad (2.13)$$

where x_0 is the interaction length of the particle and is a function of particle type and energy. We are interested only in scatters at WIMP recoil energies, which are orders of magnitude below the energies of penetrating neutrons and gammas, so we can consider the incident particle to continue traveling in the same direction with the same energy after each scatter. The probability of a particle

scattering n times along a path of length d through the target is then

$$P_n(d, x_0) = \frac{\left(\frac{d}{x_0}\right)^n}{n!} e^{-\frac{d}{x_0}}. \quad (2.14)$$

This can be proved by induction, taking the base case of no scatters, $P_0 = e^{-\frac{d}{x_0}}$, and the recursion relation

$$P_n(d, x_0) = \int_0^d dx \frac{1}{x_0} e^{-\frac{x}{x_0}} P_{n-1}(d-x, x_0), \quad (2.15)$$

where $dx \frac{1}{x_0} e^{-\frac{x}{x_0}}$ is the probability of the first scatter occurring at x , and $P_{n-1}(d-x, x_0)$ is the probability of having $n-1$ scatters along the remainder of the track. The probability of a single low-energy scatter at any given point along the path is $\frac{P_1(d, x_0)}{d} P_{E < E_0} = \frac{1}{x_0} e^{-\frac{d}{x_0}} P_{E < E_0}$, where the additional factor $P_{E < E_0}$ is the probability that the scatter energy is in the WIMP region of interest.

To determine the total single scatter rate at a given point in the detector, we integrate over all straight paths passing through that point, weighted by the expected activity along that path:

$$R = \int d\phi d\theta \sin\theta \frac{P_{E < E_0}}{x_0} e^{-\frac{d(\theta, \phi)}{x_0}} \int dr A(r, \theta, \phi), \quad (2.16)$$

where θ, ϕ define a line through the point in the detector, $A(r, \theta, \phi)$ is the activity in decays/volume a distance r from the point along that line, and $d(\theta, \phi)$ is the length of the line inside the target. The total rate will be dominated by the path that minimizes $d(\theta, \phi)$. At the edges of the detector, this length can be very short, and backgrounds are correspondingly high. When the minimum d is many times the interaction length, backgrounds are greatly reduced (see the example in Chapter 3).

This style of background rejection is not effective against non-penetrating internal backgrounds such as beta decays, so this technique requires an extremely radiopure target. A prime candidate for this is liquid xenon, which is easily purified, has no long-lived radioactive isotopes, and has interaction lengths of 10 and 15 cm for gammas and neutrons, respectively. The XMASS experiment uses pure liquid chambers surrounded by PMTs for event reconstruction [33]. A 100 kg XMASS detector has been completed, and an 800 kg module is now being constructed (see Fig. 2.4). The next generation of dual-phase xenon time projection chambers will also take advantage of active self shielding, as we will discuss in Chapter 3.

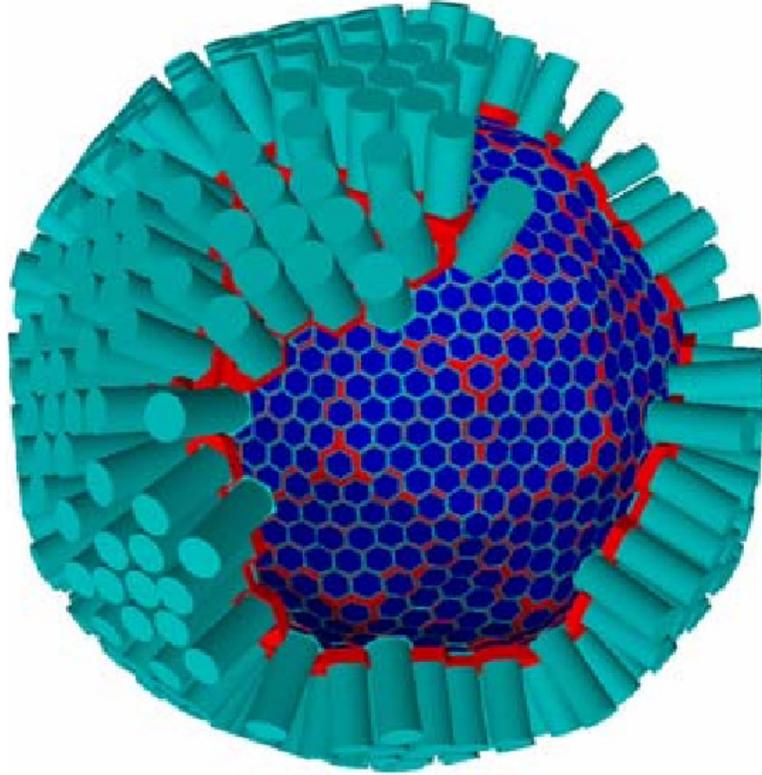


Figure 2.4: Schematic of the 800 kg XMASS detector. Figure taken from [33].

2.2.4 Nuclear Recoil Discrimination

The most successful dark matter detectors to date have been those that discriminate between electron and nuclear recoils. The dominant nuclear recoil background at WIMP energies comes from elastic scattering of neutrons produced in cosmic muon interactions and alpha-n nuclear reactions. These backgrounds are reduced by going deep underground, shielding with hydrogenated material, and practicing the same low background techniques (targeting the U/Th chains) used to reduce internal gamma backgrounds. Neutron backgrounds are several orders of magnitude lower than the Compton scatter background for a given U/Th activity, so a method for rejecting electron recoils can greatly reduce the overall background. Electron recoil discrimination is also the only way to reject naked betas in the target bulk. With one exception, discriminating experiments work by detecting events in multiple channels, using the relative strengths of the different signals to determine the recoil type.

Cryogenic Detectors

Cryogenic detectors operate at temperatures of <100 mK and detect both electronic excitation (either as scintillation, as in CRESST [35], or as ionization, as in EDELWEISS [36] and CDMS

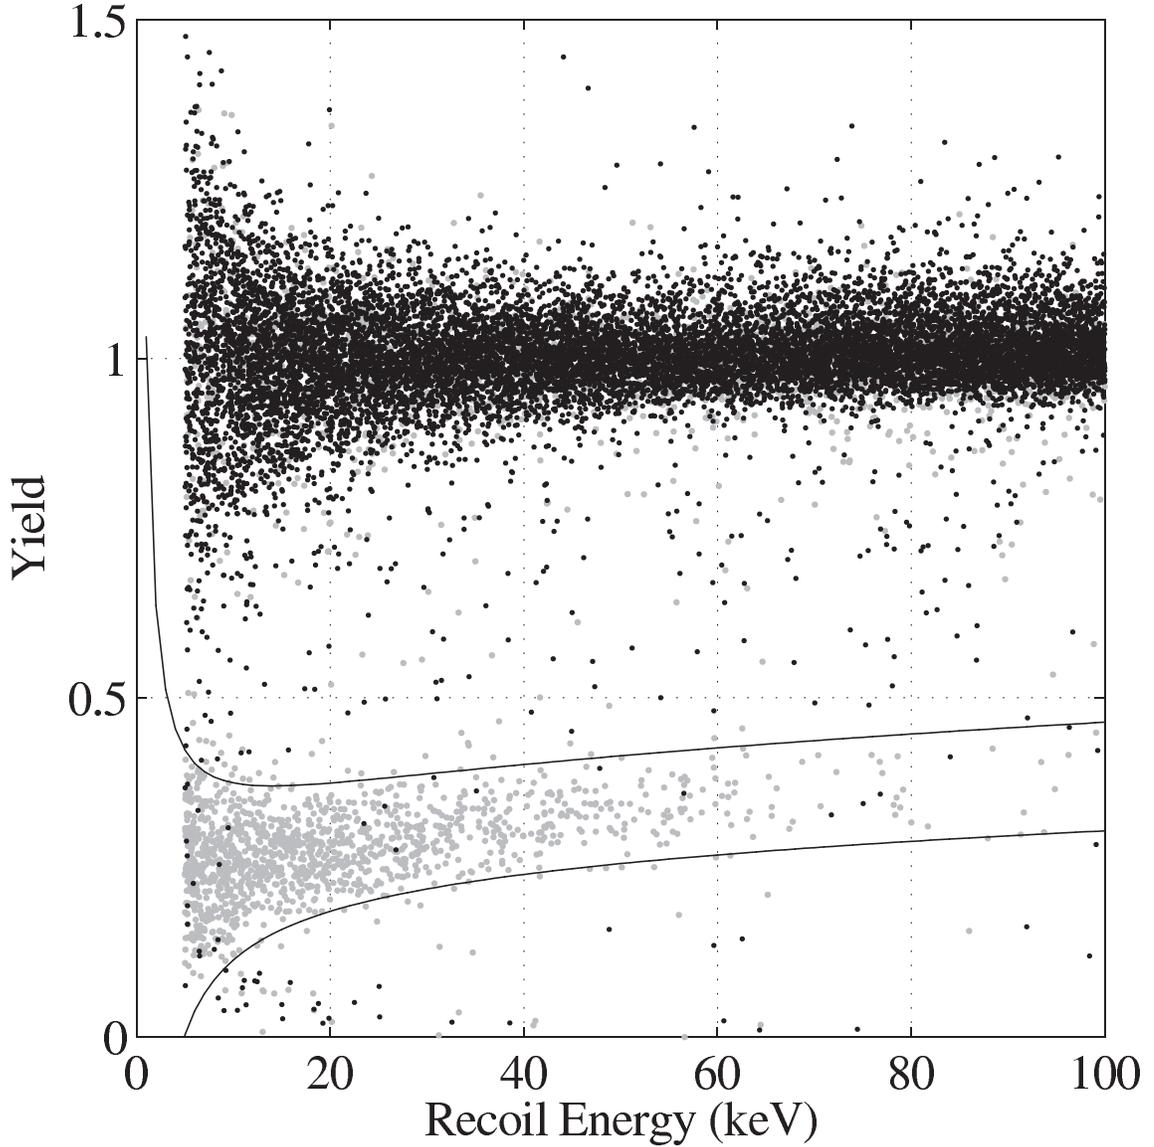


Figure 2.5: Electron recoil discrimination in CDMS. The dark (top) population of events are electron recoils from a gamma source, and the gray (bottom) are nuclear recoils from a neutron source, with the y-axis showing the ratio of the ionization signal to the phonon signal, normalized to electron recoils. The discrimination shown here is before timing cuts, which remove problematic events at the edge of the detector. Figure taken from [34].

[30]) and thermal/phonon energy loss. Nuclear recoils lose the majority of their energy through low energy collisions with other nuclei, producing typically one fourth the electronic excitation of an electron recoil with the same energy. This effect was discovered by Lindhard [37], and is discussed in Chapter 4. Electron recoil misidentification rates of less than 1 in 10^{-6} have been reported by CDMS (see Fig. 2.5). The discrimination in these detectors is limited by processes that cause an

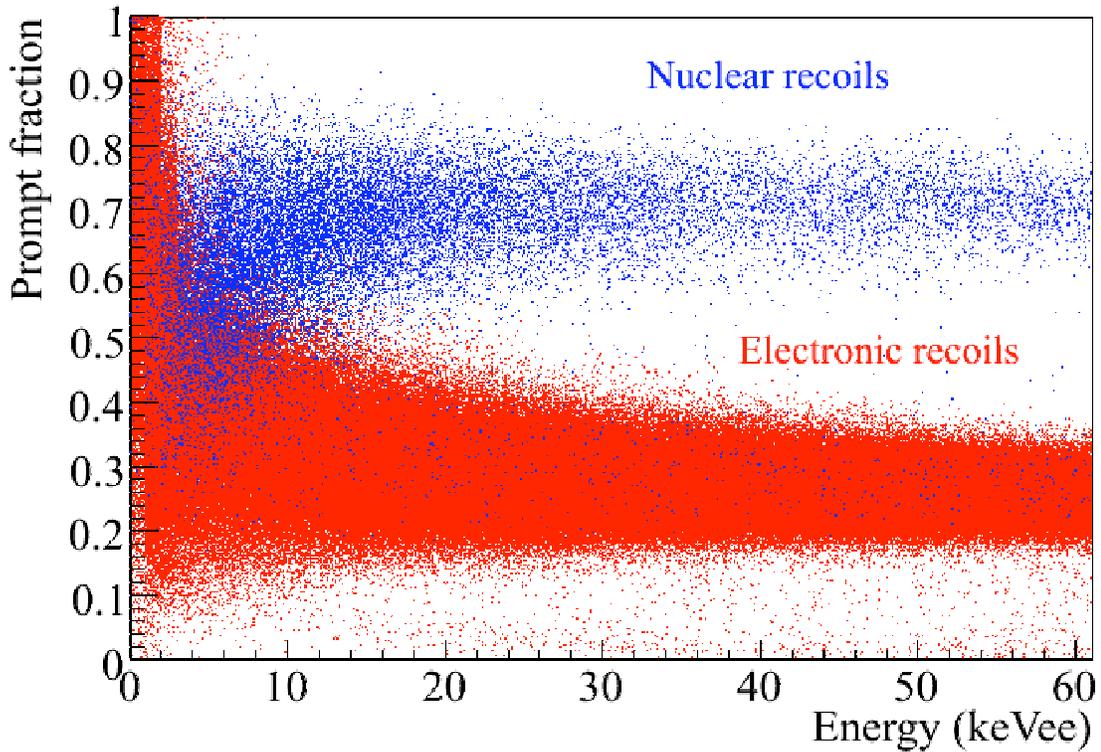


Figure 2.6: Electron recoil discrimination in DEAP (liquid argon) based on pulse shape discrimination. The blue (top) population of events are nuclear recoils from a neutron source, and the red (bottom) are electron recoils from a gamma source. Nuclear recoils produce a greater fraction of excitons in the short-lived singlet state, giving a higher fraction of ‘prompt’ scintillation. Figure taken from [38].

event to lose part of the electronic signal, particularly at the edges of the detector.

Pulse-Shape Discriminating Detectors

Discrimination is also possible when looking at the electronic signal alone. The DEAP/CLEAN collaboration uses the pulse-shape in the scintillation signals in liquid argon and neon to distinguish electron and nuclear recoils [38]. Noble liquids produce scintillation following recombination of ions (see Chapter 4). The resulting excitons can be in a long-lived triplet state or short-lived singlet state, with a preference in nuclear recoils for the singlet state. Nuclear recoils thus have a greater fraction of scintillation light at the start of the pulse, as in Fig. 2.6. DEAP has demonstrated electron recoil discrimination of 7.6×10^{-7} in liquid argon at WIMP energies with 50% acceptance for nuclear recoils. This large discrimination is necessary for an argon based experiment due to the background from the 269 year half-life naked beta emitter ^{39}Ar .

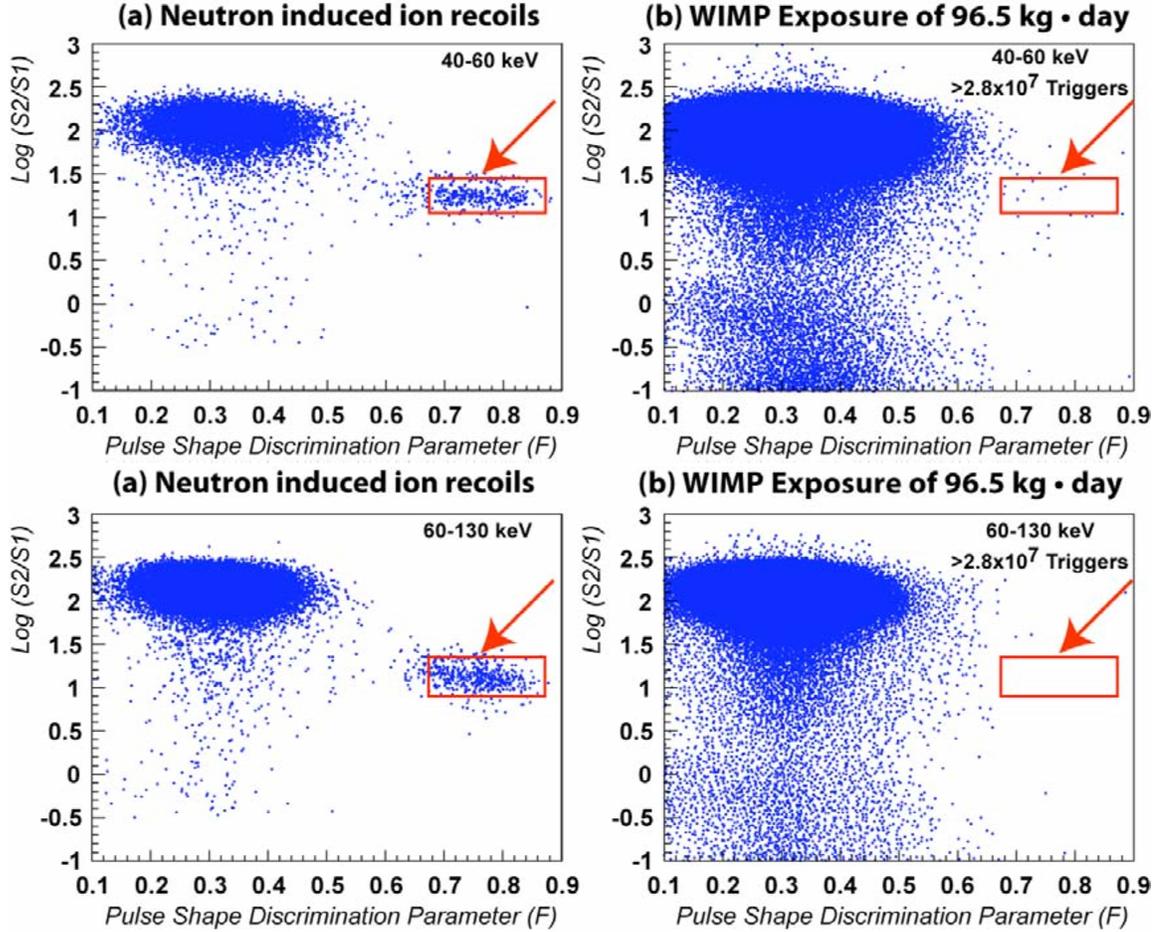


Figure 2.7: Electron recoil discrimination in WARP. The left plots show nuclear recoils from a neutron source, and the right show WIMP data (dominated by ^{39}Ar betas). The top plots show a WIMP recoil energy range from 40–60 keV, and the bottom 60–130 keV. The x-axes show the pulse-shape discrimination parameter and the y-axes the charge-light discrimination parameter. Discrimination at 50% nuclear recoil acceptance is 3×10^{-7} , primarily from the scintillation pulse shape discrimination. Figure taken from [39].

Charge-Light Discriminating Detectors

Another way to split the electronic excitation signal in noble liquids is to apply a drift field and extract electrons from the event before they recombine. The ratio of the charge and light signals can then be used to distinguish electron and nuclear recoils, giving discrimination up to 10^{-3} at 50% nuclear recoil acceptance — this is discussed extensively in Chapters 5 and 6. This technique also has the advantage that the charge signal may be used to give high resolution position information for an event, as described in Chapter 3. This allows XMASS style rejection of gamma and neutron events while rejecting betas (which XMASS cannot reject) using the electron recoil discrimination. The ZEPLIN [40], XENON, and LUX collaborations use this technique in liquid xenon. The WARP [39]

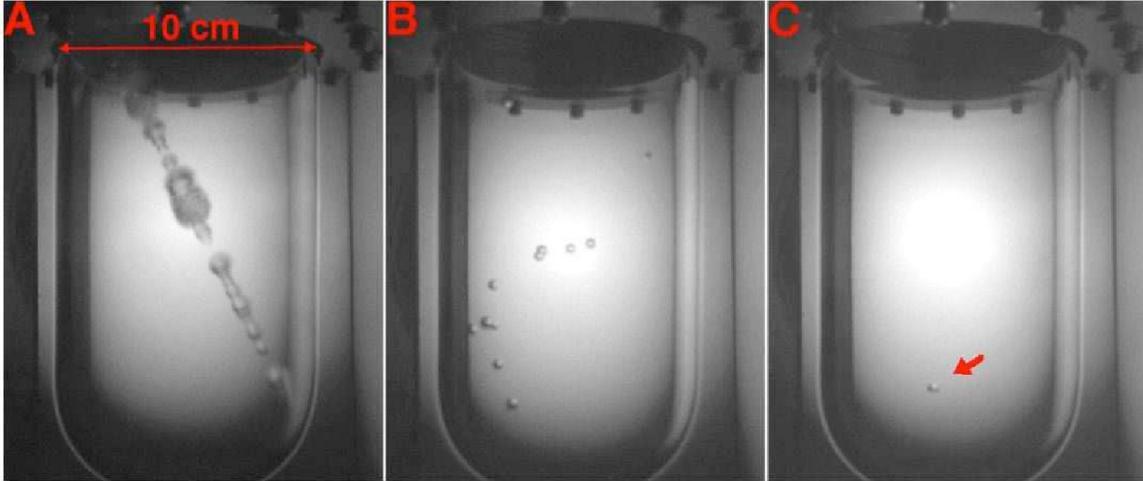


Figure 2.8: From left to right, muon, neutron, and alpha (or WIMP) events in COUPP. Muon tracks and other minimum ionizing particles do not nucleate bubbles in the detector under normal (WIMP search) operating conditions. Neutrons may be identified by multiple scatters, while alphas and WIMPs appear as a single bubble. Pictures taken from [32].

and ArDM [41] experiments use an argon target, taking advantage of both pulse-shape discrimination and charge-light discrimination, but losing the active shielding advantage because of the internal ^{39}Ar background. The combined pulse-shape and charge-light discrimination in WARP is shown in Fig. 2.7.

Superheated Liquid Detectors

The final WIMP detector type uses superheated liquids, either as droplets in a gel (PICASSO [42]), or as a large bubble chamber (COUPP [32]). Energy deposited by a nuclear recoil forms a bubble in the metastable liquid, which then grows until it is observable either as an acoustic signal, as in PICASSO, or visually, as in COUPP. The temperature and pressure of these chambers is set so that electrons and other minimum ionizing particles cannot deposit enough energy in a small enough volume to nucleate a stable bubble. Discrimination in these detectors is extremely powerful — COUPP reports 1 in 10^{10} electron recoils producing a bubble at a nuclear recoil threshold of 10 keV (threshold is also a function of pressure and temperature). The weak point in these detectors is the inability to discriminate between bulk alpha decays and WIMP recoils, but recent measurements by PICASSO suggest that these events may be distinguished by their ultrasonic signature [43]. See Fig. 2.8 for pictures of muon, neutron, and alpha (or WIMP) events in COUPP.

Chapter 3

Xenon Time Projection Chambers

This chapter introduces the dual-phase xenon time projection chamber (TPC), also known as an electron emission chamber. We lay out the basic technique and the advantages of these detectors for direct detection of WIMPs, and describe in detail the prototype chamber ‘Xed’, which we use for the calibration and discrimination measurements in Chapters 4–6.

3.1 Basic Xenon TPC Operation

Interactions in liquid xenon produce both free electrons and 175 nm scintillation photons (the physics behind the production of these two signals is discussed in Chapter 4). The scintillation signal is prompt, produced in the ~ 100 ns following the interaction. The charge signal can be collected by drifting the free electrons under an applied field, with drift velocities of ~ 2 mm/ μ s at typical fields (see Fig. 3.1). The delay between the prompt scintillation and the arrival of electrons at the collector gives the distance to the interaction, hence the name ‘time projection chamber’. In the context of direct detection experiments (as opposed to high energy tracking chambers), these chambers measure point interactions, where a penetrating particle (gamma, neutron, neutrino, or WIMP) or internal decay (alpha or beta) deposits energy in a single recoil extending no more than a few 10’s of μ m (see Chapter 6 for details of energy loss by electron and nuclear recoils), so the charge signal arrives as a short pulse. The prompt scintillation and charge signal are known as S1 and S2, respectively.

There are many methods for measuring the S1 and S2 signals, with the typical setup is shown in Fig. 3.1. Here the S1 is detected by photo-multiplier tubes (PMTs) above and below the target volume. Efficiencies for detecting S1 photons in existing detectors are 5–10%, with $\sim 50\%$ light collection efficiency and $\sim 17\%$ PMT efficiency — see Chapters 4 and 5, as well as [44] for more

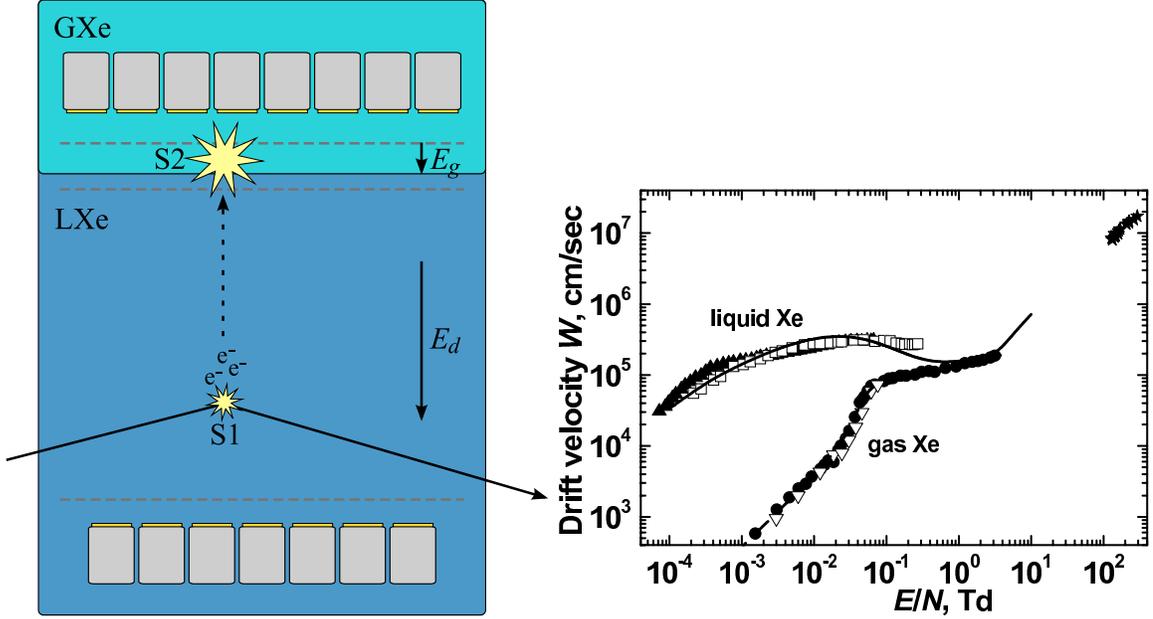


Figure 3.1: Left: Schematic of a dual phase xenon time projection chamber. An incoming particle (gamma, neutron, or wimp) scatters in the liquid, producing a flash of prompt scintillation and free ionization. The free electrons drift to the liquid surface under the applied field E_d and are extracted into the gas, where they produce electro luminescence until being collected on the anode. The two flashes of scintillation light (prompt and electro luminescence) are known as S1 and S2, respectively. The delay between flashes gives the drift time of electrons in the liquid, and thus the depth of the event. Arrays of photo-multiplier tubes above and below the active volume detect the S1 and S2 light. The top array is also used to determine the x-y position of the event, based on the location of the S2 flash. This figure is supplied by John Kwong, appearing in [44]. Right: Drift velocities of electrons in liquid and gas xenon. The x-axis is in drift field over number density of atoms, where $1 \text{ Td} = 10^{-17} \text{ V cm}^2$. At 180 K (a typical operating temperature), $N = 1.3 \cdot 10^{22} \text{ cm}^{-3}$ in liquid and $9.4 \cdot 10^{19} \text{ cm}^{-3}$ in gas [45]. This figure is taken from Atrazhev et al. [46]. The solid lines show the calculations by Atrazhev et al., and the points show data from experiments [47, 48, 49, 50].

discussion on S1 efficiency. To detect the charge signal with the same efficiency, some physical gain is needed. This is easily accomplished in xenon gas, where electrons drifting at high field produce electro luminescence [51]. This light is also 175 nm, and is sometimes called ‘proportional scintillation’ in analogy with a proportional counter, where electrons drifting at high field produce additional ionization.

There is a potential barrier at the phase boundary that must be overcome to extract electrons from liquid into gas [51], and gas fields of 10 kV/cm are necessary for efficient (>98%) extraction [52]. At typical operating pressures ($\sim 2 \text{ atm}$), electrons drifting in this field produce $\mathcal{O}(100)$ scintillation photons per millimeter — this light may be measured with the same PMTs used to detect S1. Detection of single electrons extracted from the liquid has been demonstrated using this technique (see Chapter 7). An array of PMTs above the anode gives the x-y location of the S2 light pulse.

Combined with the depth determined from the electron drift time, this gives full 3-D position reconstruction. Position resolutions of 3 mm in x-y and 1 mm in z have been demonstrated in these detectors [53], and because of the relatively large S2 signals this resolution is essentially independent of event energy, all the way to the S1 detection threshold.

A number of support systems are necessary for dual-phase TPC operation. The most basic of these include cryogenics (typical operating temperatures are 170–190 K) and xenon purification (electronegative impurities can trap drifting electrons). Details on the instrumentation used in the Xenon10 and LUX experiments may be found in [54, 55, 44, 56], and in Section 3.3 we describe in detail our prototype xenon TPC, Xed.

3.2 XenonTPCs as WIMP Detectors

Xenon TPCs excel as direct detection experiments. They are easily scalable, and have the low threshold needed for WIMP recoil sensitivity. Xenon itself is easily purified and strongly self-shielding, with no long-lived radioactive isotopes, so extremely low background rates are possible. The backgrounds that remain are primarily electron recoils and may be discriminated against based on the ratio of the S2 and S1 signals.

3.2.1 WIMP sensitivity

The WIMP sensitivity of xenon compared to germanium and argon for spin-independent interactions was shown in the previous chapter (Fig. 2.1). Xenon benefits from the A^2 term in the coherent spin-independent cross section (Eq. (2.9)). Because of its large nucleus, scattering on xenon loses coherence at high energies, but for sufficiently low energy thresholds (e.g. 17 keV for 100 GeV WIMPs) xenon is the most sensitive of these targets. The energy threshold in liquid xenon is determined by the efficiency for detecting S1 photons. Low energy nuclear recoils in xenon give ~ 10 scintillation photons per keV, and thresholds of 5 keV have been demonstrated (see Chapters 4, 7). Xenon is also sensitive to spin-dependent interactions coupled to neutrons — ^{129}Xe and ^{131}Xe are neutron-odd and make up 26.4% and 21.2% of naturally occurring xenon, respectively.

3.2.2 Scalability

WIMP sensitivity scales with the mass of the detector, and current WIMP limits already show the need for targets of several hundred kilograms, or even several tons. The target material in a xenon

TPC is a liquid and the instrumentation scales with the surface area (top and bottom arrays of PMTs) while the sensitivity scales with volume, making these detectors both easy and economical to scale to very large sizes. The only limitations to size are set by high voltage constraints and the ability to drift charge and see scintillation light over long distances. Of these, the limiting factor appears to be charge drift. A drift length of 4 meters was measured in the Xenon10 experiment [57, 55] — we discuss the problem of charge drift in larger detectors in Chapter 7.

3.2.3 Chemical Purity

Xenon is a noble gas, and so chemical purification is straightforward. Commercial xenon is produced by distillation, and purification in the lab is performed by heated getters, which remove non-noble contaminants at <1 ppb [58]. Purification from other noble elements may be achieved by gas chromatography (see next section).

3.2.4 Radiopurity

There are no long-lived radioactive xenon isotopes. The chief internal background comes from krypton, which is present at 5–10 ppb in the best commercially available xenon. ^{85}Kr is a naked beta emitter (99.6% of decays have no associated gamma) with half-life of 10.77 years, decaying to the stable ^{85}Rb . At WIMP recoil energies, the beta spectrum gives $2.2 \cdot 10^{-3}$ betas per keV per decay (nuclear data from [59]). With a conservative (high) estimate for the relative abundance of ^{85}Kr , one ppb of krypton contamination corresponds to ~ 0.05 events/keV/kg/day [60].

Krypton may be removed from xenon via either distillation or gas chromatography. In the latter, a pulse of xenon contaminated with krypton is forced through a charcoal column by a carrier gas (helium). The greater adsorption of xenon onto the charcoal separates the gasses in the column — the krypton emerges first and is diverted to a trap, followed by the xenon (see Fig. 3.2). A chromatography system built at Case Western has achieved krypton levels of < 3 ppt (parts-per-trillion) [60], sufficient for a 10 ton experiment..

3.2.5 Active Shielding

A major strength for these detectors is their 3-D position reconstruction. As described in Section 2.2.3, neutron and gamma events can be efficiently rejected in a large 3-D position reconstructing detector by identifying multiple-scatters. Interaction lengths for gammas and neutrons in liquid xenon are shown in Fig. 3.3, roughly 15 cm for neutrons and 10 cm for the most penetrating gammas.

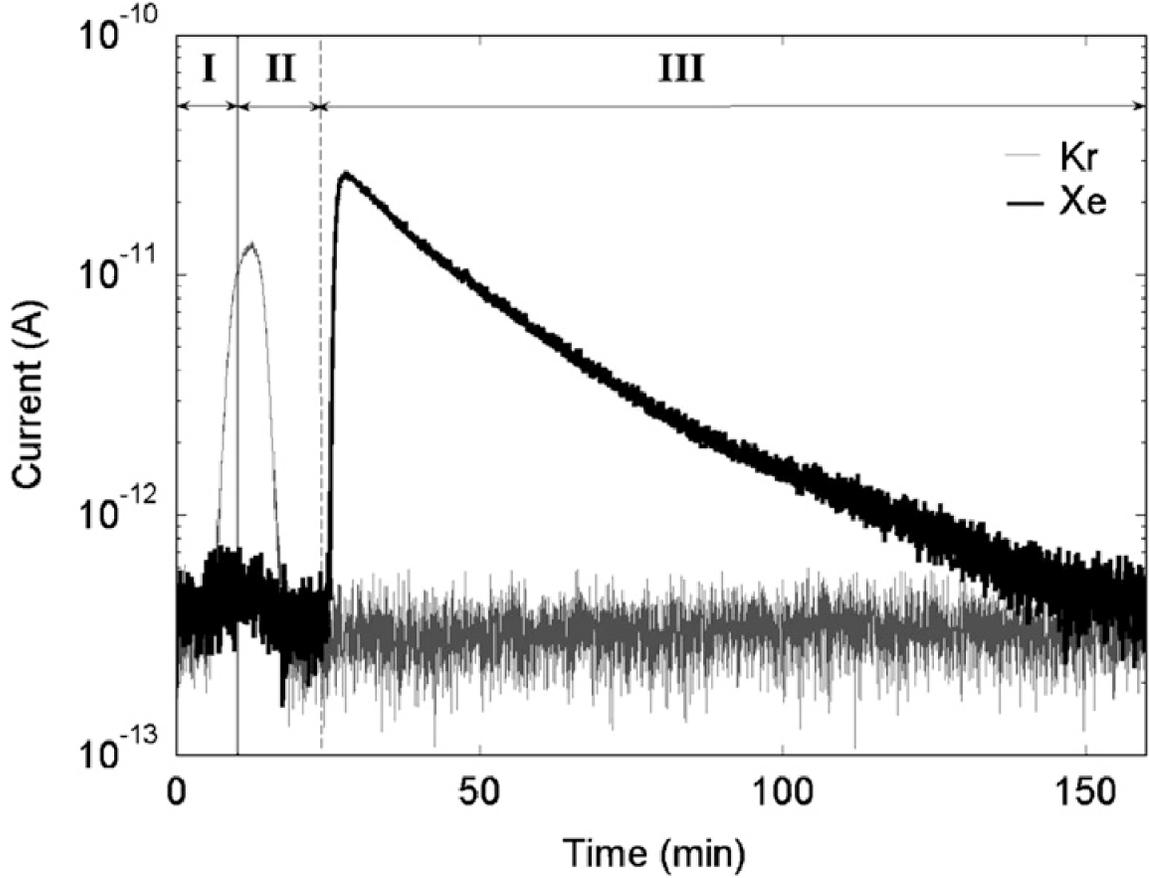


Figure 3.2: Krypton removal from xenon using gas chromatography. This plot shows the krypton and xenon partial pressures (plotted as current in a mass spectrometer) at the outlet of the charcoal column in the Case Western krypton purification system. In phase I, a pulse of xenon+krypton is fed into the charcoal column, in phase II the krypton leaves the column and is diverted to a cold trap, and in phase III the xenon is leaves the column and is recovered. This plot is taken from [60].

From Eq. (2.16), the single scatter rate at a point in the detector due to gammas or neutrons from a given direction is proportional to $e^{-\frac{d}{x_0}}$, where x_0 is the interaction length of the particle and d is the path length in the active region that the particle must travel. For example, in the corner of the detector there are several paths that cross only a small length of active xenon, and penetrating particles along these paths have a significant single scatter probability. Any path through the center of the detector must cross the entire active region, so single scatters there are very unlikely. Integrating over all paths through a given point, weighting by the expected activity along each path as in Eq. (2.16), we get very similar results to a full Monte Carlo background simulation. The expected Compton scatter background as a function of position in a LUX size detector is shown in Fig. 3.4.

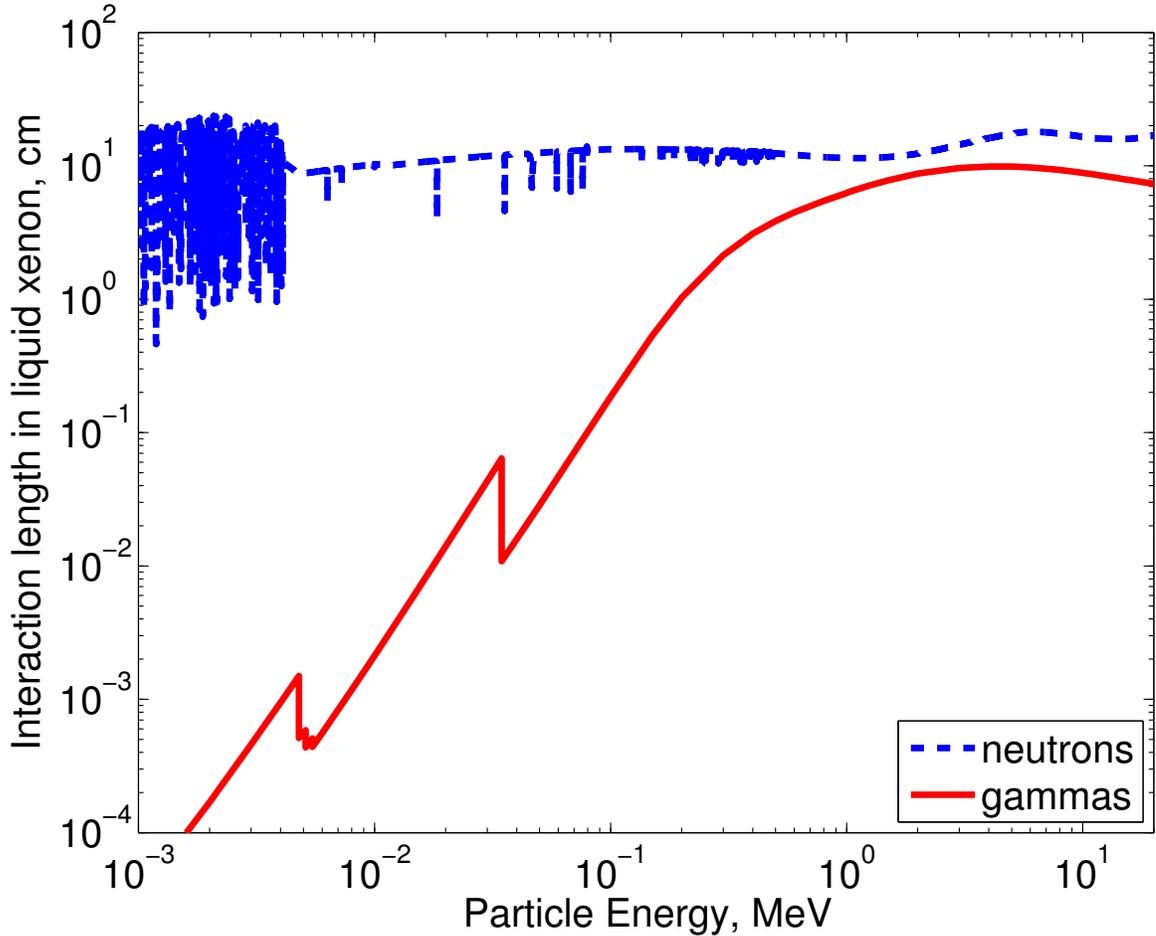


Figure 3.3: Gamma and neutron interaction lengths in liquid xenon. Gamma data comes from [61] and neutron data from [62]. Neutron data is averaged over the stable xenon isotopes, weighted by their abundance. We take a liquid xenon density of 2.84 g/cm^3 ($T=180 \text{ K}$).

3.2.6 Discrimination

Active shielding is the primary method for rejecting for gamma events, but one can also discriminate between electron and nuclear recoils using the ratio of the S1 and S2 signals. This also allows rejection of backgrounds from internal betas (dominated by ^{85}Kr) and at very large scales from charged-current neutrino interactions. Xenon10 has demonstrated 99.4% rejection efficiency with 50% nuclear recoil acceptance, and we project 99.9% rejection at low energies in upcoming experiments — see Chapter 5 for an in depth discussion on electron recoil discrimination.

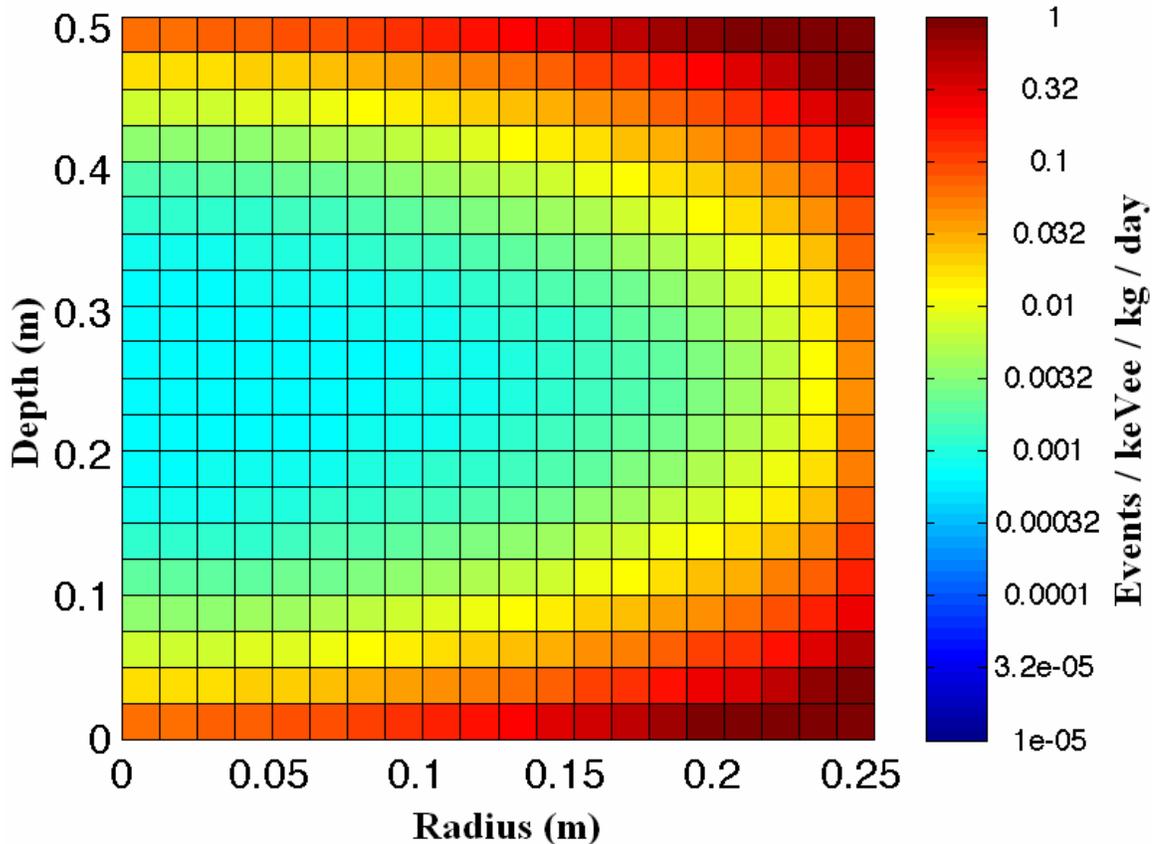


Figure 3.4: An analytic calculation of single scatter background due to gammas in a cylindrical 300 kg (LUX size) liquid xenon chamber. The x-axis gives the distance from the central axis of the detector, while the y-axis gives the height. The rate varies by three orders of magnitude from the hottest point (top and bottom edge) to coldest (center). This plot produced by Tom Shutt.

3.3 The Xed Prototype TPC

The data presented in Chapters 4, 5, and 6 was taken in a prototype xenon TPC constructed at Princeton University and operated at Case Western Reserve University from 2005 on. This detector was dubbed Xed. This section describes the instrumentation, data acquisition, and data processing techniques used with Xed.

3.3.1 Cryostat

The Xed cryostat is shown in Fig. 3.5. This is a vacuum cryostat with an aluminum cold-finger to a liquid nitrogen bath for cooling. The cryostat has two sealed spaces: the inner can, containing the detector, and the outer can, kept at vacuum. The detector hangs from the top flange of the inner can, which in turn hangs from the top flange of the outer can, supported by Vespel[®] rods. Thin

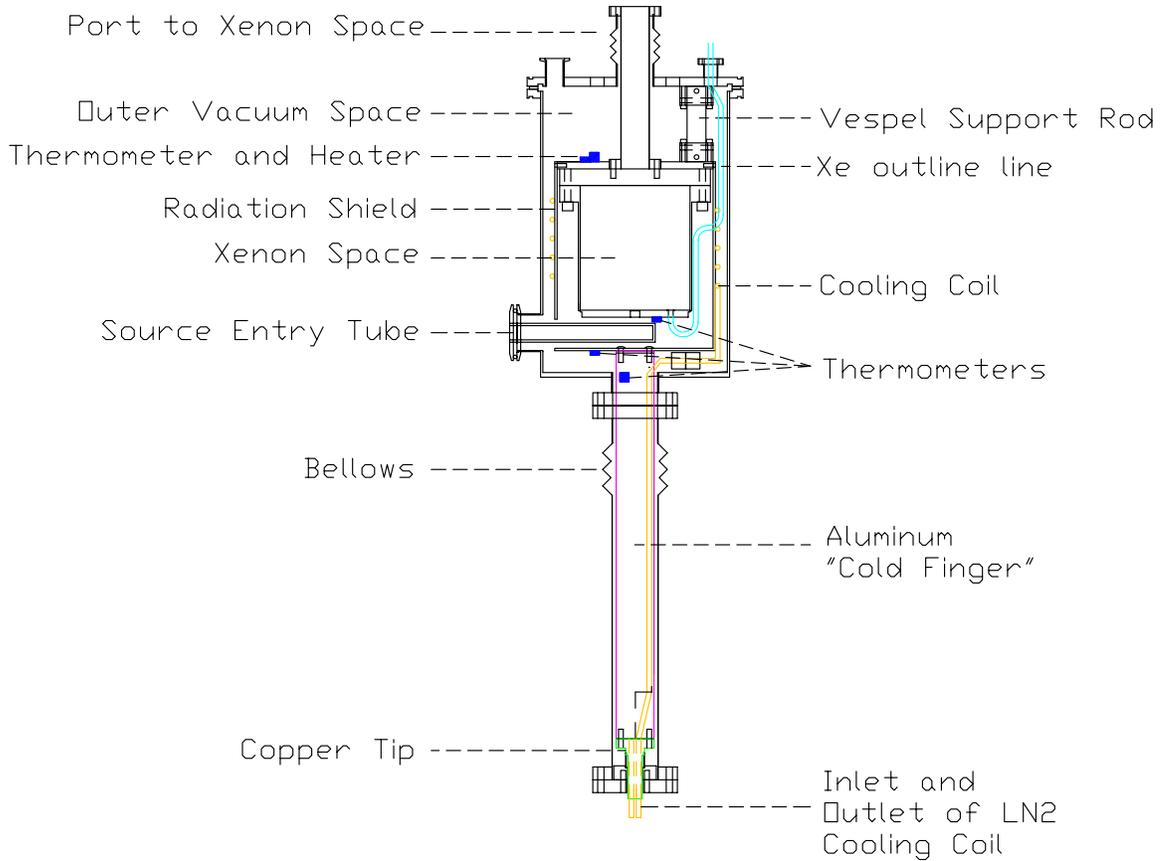


Figure 3.5: The Xed vacuum cold-finger cryostat, designed by Tom Shutt. Figure compiled by John Kwong, as appearing in [44].

walled stainless steel nipples extend from the inner can up to air-side feedthroughs for electrical and plumbing connections. An aluminum radiation shield is mounted to the top flange of the inner can, and the cold-finger extends from the bottom of the radiation shield. Platinum resistor thermometers are mounted at the top of the cold finger, the bottom of the radiation shield, and the top and bottom of the inner can. A PID controlled 50 W heater on the top of the inner can regulates the temperature.

3.3.2 Internal Structure

The internal structure of the Xed detector is shown in Figs. 3.6 and 3.7. Two photo-multiplier tubes (PMTs) form the top and bottom faces of a cylinder with 3.729 cm diameter in the active region, stepping out to 4.534 cm diameter at the top. Five wire grids (in order, ‘bottom’, ‘cathode’, ‘gate’, ‘anode’, and ‘top’) set the electric fields in the chamber. Each grid is a single plane of parallel wires, tensioned and soldered to copper pads on a cirlex disk at 2 mm pitch. The bottom three grids are 40 μm BeCu wires at 40 gm tension (room temperature), and the top two grids are 120 μm

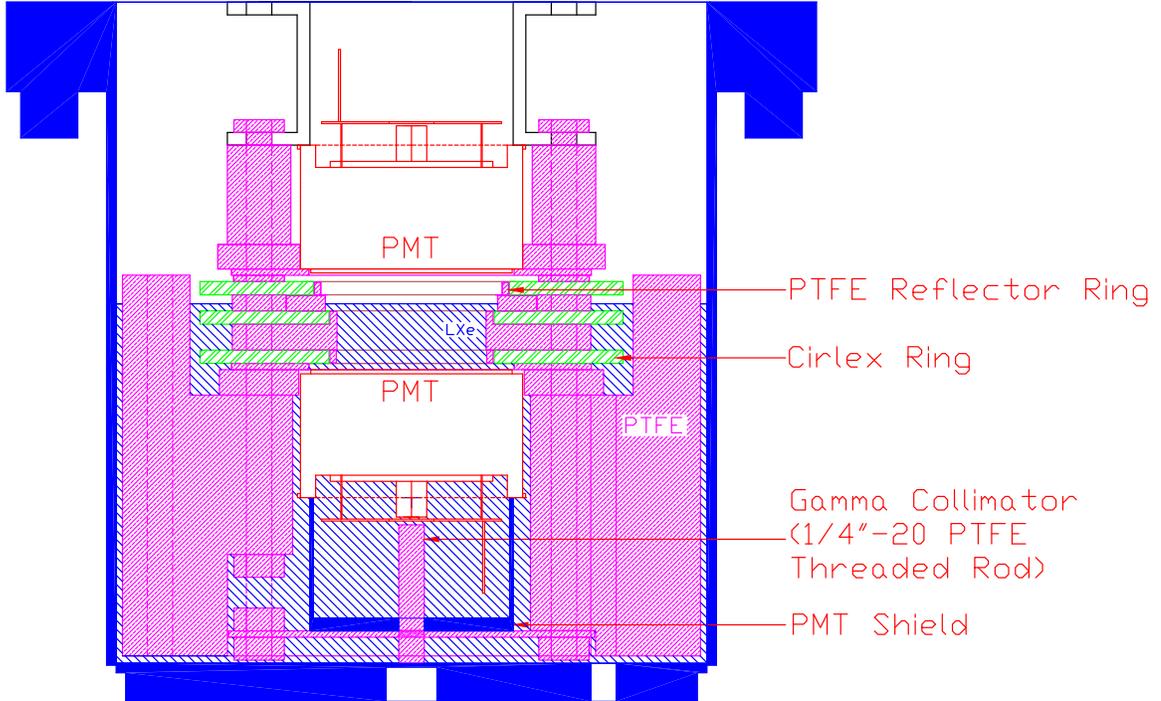


Figure 3.6: The Xed detector. This shows the two PMT configuration — in the single PMT configuration, the bottom PMT and wire grid are removed, and the cathode (next-to-bottom) grid is replaced by a silver-plated copper plate with a ^{210}Po alpha source deposited in its center. The active region in both configurations is 0.964 cm tall and 3.729 cm in diameter. The following materials are color coded: stainless-steel (solid blue), liquid xenon (slashed blue), PTFE (Teflon[®]) (slashed pink), Cirlex[®] (slashed green), and aluminum (black outline). Photo-multiplier tubes are shown in red. This figure produced by John Kwong, appearing in [44].

Au plated Al wires at 100 gm tension. The active volume, between the cathode and gate grids, is 0.964 cm tall, for a total active mass of 30 gm. The cathode, gate, and anode grids set the drift and extraction fields, while the top and bottom grids shield the PMTs from the high field regions. The extraction field is kept at 10 kV/cm in the gas (5.1 kV/cm liquid) with the liquid surface midway between the gate and anode grids. Active region drift fields of up to 4.4 kV/cm are achieved in this setup. Voltages are set using the electrostatic analysis in Appendix A.

Typical events produce similar numbers of primary scintillation photons and electrons. Each electron in turn produces ~ 250 secondary scintillation photons, creating a large dynamic range between small S1s and large S2s. Due to reflection at the liquid surface, the bottom PMT is ~ 4 times more efficient than the top for detecting S1 light. Here we use a Hamamatsu R9288 with a quoted 23.24% quantum efficiency at 175 nm and 70% photo-electron collection efficiency for a total efficiency of 16%, biased at 875 V for a clear single photo-electron signal. The top PMT is a Hamamatsu R6041 with 7% total efficiency, biased at 820 V. S2s of over ~ 300 electrons begin to

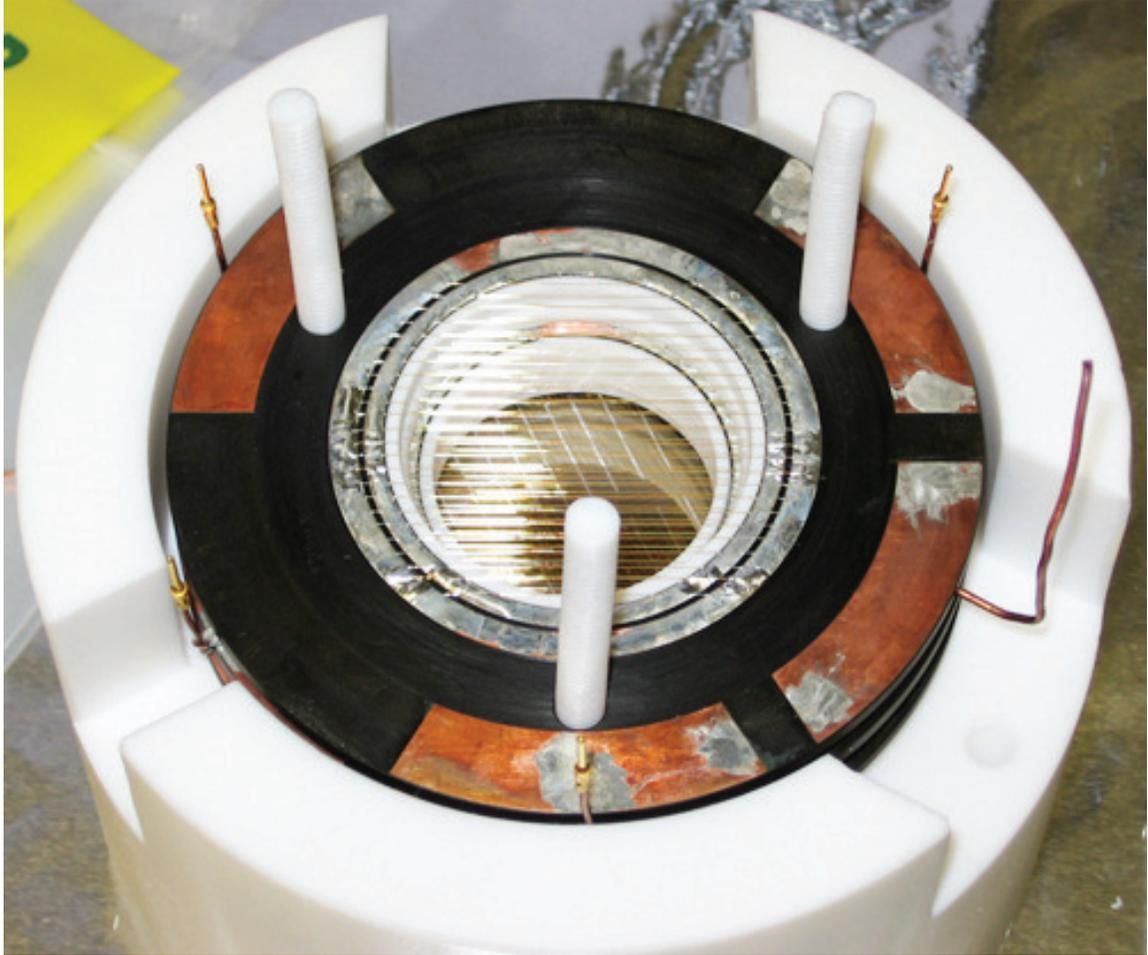


Figure 3.7: Picture showing the Xed active region, with the top PMT removed. Copper clad Cirlex[®] rings hold the wire grids used to set the fields in the chamber. The copper pads at the outside edge of the ring are repeated on the grid and anode planes and form the capacitors used to measure liquid level. The pads at 120° on the grid and anode planes measure the liquid level, and the fourth pads on the anode and top planes form a reference capacitor. This picture originally published in [44].

give a non-linear response in the bottom PMT and are measured with the top PMT. The maximum S1 signal giving a linear response in the bottom PMT corresponds to $\sim 10,000$ scintillation photons, or ~ 600 bottom PMT photon-electrons.

3.3.3 Cleaning and Xenon Purification

Xenon purity begins with cleanliness inside the detector. The detector is assembled in a clean room or at a clean bench (hooded bench with HEPA blower), and all internal detector components except for PMTs are ultrasonically cleaned in acetone followed by ethyl alcohol, with deionized water rinses after both cleanings. After sealing, the detector is heated to 50 C° (limited by the PMT's) and

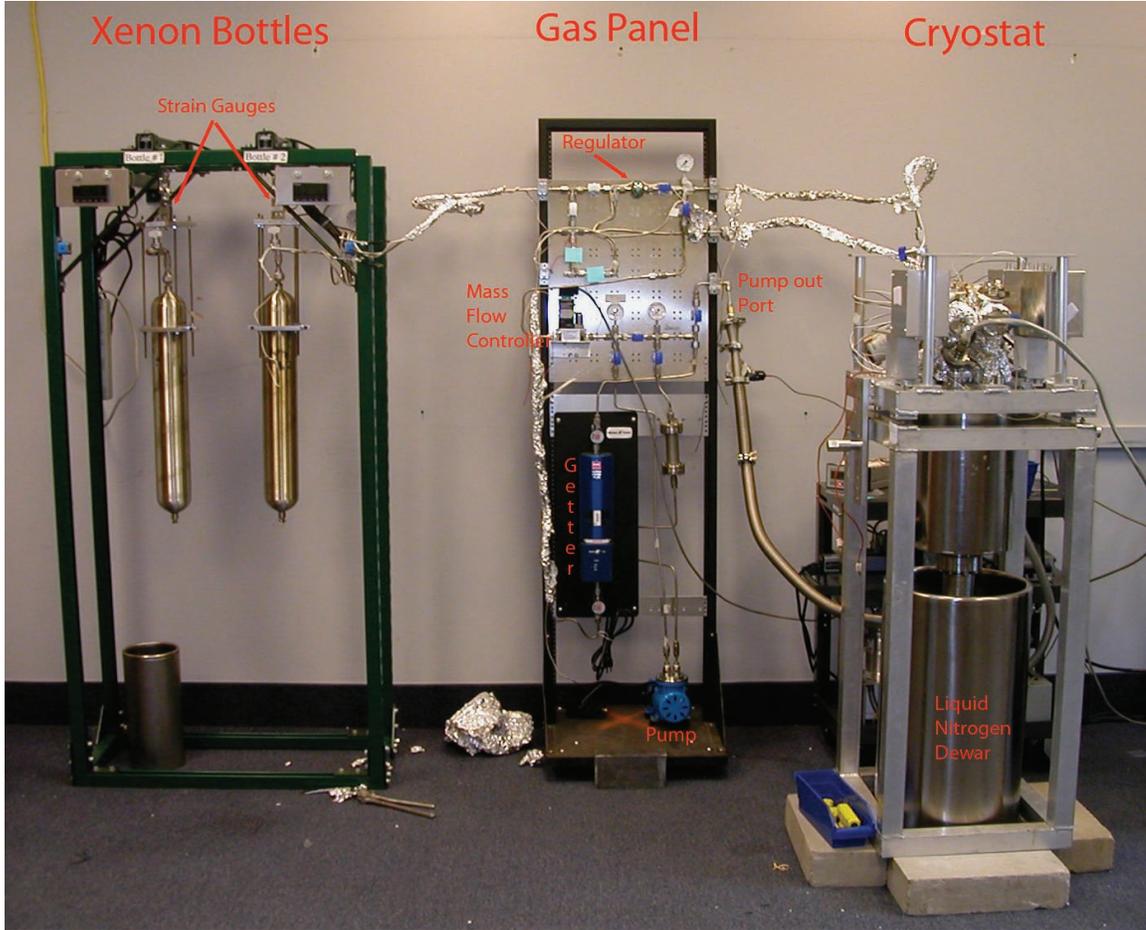


Figure 3.8: Picture showing (from right to left) the Xed cryostat, gas panel, and storage cylinders. Picture annotations by John Kwong, as appearing in [44].

pumped for several days to drive out impurities absorbed in detector materials (in particular H_2O , although the exact impurity content is unknown), typically reaching vacuum pressures of 10^{-5} torr.

Xenon is purified using a SAES[®] MonoTorr[®] heated getter, pictured with the gas handling system in Fig. 3.8. Xenon passes through the getter when moving between the stainless-steel storage cylinders and the main chamber and when recirculating the fluid in the main chamber. At the start of each run, the xenon was either recirculated or repeatedly moved between the storage cylinders and the chamber (batch mode purification) until drift lengths of 30 cm were achieved, determined from the depth dependence of S2 from 122 keV photo-absorption events. Recirculation was restarted periodically during a run as necessary to maintain a 30 cm drift length, but data was not taken while recirculating due to large fluctuations in the liquid level.

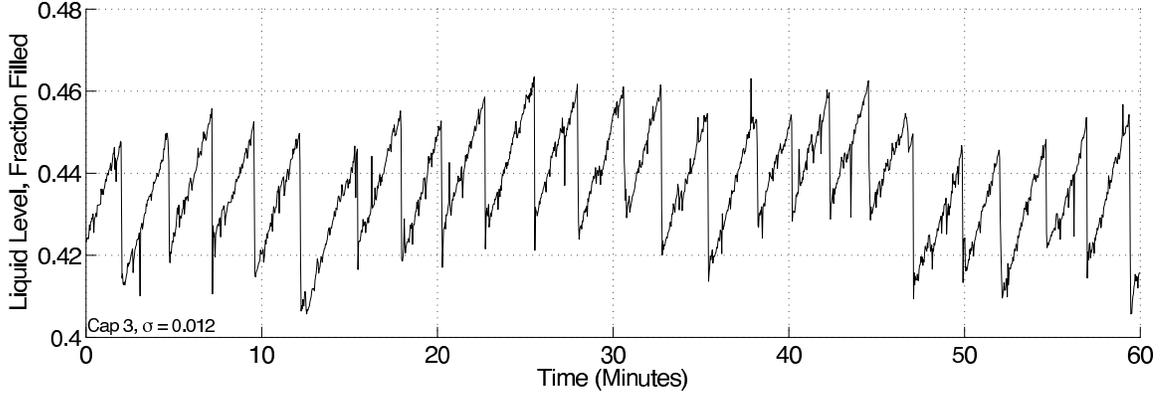


Figure 3.9: Liquid level monitoring in Xed. This shows one of the three parallel plate liquid level monitors displaying the sawtooth pattern that we associate with bubbles trapped under surfaces in the liquid. The same pattern is seen with the same phase and amplitude in all three meters. As the bubble grows, liquid is displaced and the liquid level rises, until the bubble is large enough to escape and the liquid level quickly drops back to its starting point. The y-axis gives the liquid level as a fraction of the distance from the gate to anode (4.67 mm) — the amplitude of the oscillations is $\sim 100 \mu\text{m}$. The bubbles can be stopped by slowly warming the detector from the top. This raises the gas pressure above the vapor pressure of the liquid, halting bubble formation. Plot courtesy of John Kwong, appearing in [44].

3.3.4 Liquid Level Monitoring

For normal operation, the liquid level is kept in the 0.467 cm region between the gate and anode grids. The height and tilt of the liquid surface are measured using three parallel plate capacitors, each with one plate level with the gate and one level with the anode. The commercially available UTI chip reads the capacitances [63], achieving $< 40 \mu\text{m}$ resolution in the liquid height and a sample rate of $\sim 2 \text{ Hz}$. The level meters also monitor the state of the liquid surface, detecting ripples in the liquid surface (due primarily to bubble formation at the bottom PMT base), shifts in the liquid level when fields are applied, etc. Figure 3.9 shows behavior associated with bubbling in the xenon.

3.3.5 Source Placement

The dual-phase data in Chapters 4 and 5 uses three sources: ^{57}Co for photo-absorption events from 122 and 136 keV gammas, ^{133}Ba for Compton scatters from 356 keV gammas, and ^{252}Cf for elastic and inelastic nuclear recoils from neutrons (average energy 2.35 MeV). The ^{57}Co source is inserted into a re-entrant tube in the cryostat, positioning it directly under the center of the active volume. The bottom of the stainless-steel inner can is thinned from 0.375" to 0.039" in a 0.5" diameter disc over the source to allow the 122 keV gammas into the xenon volume. A PTFE rod from the inside bottom of the inner can to the bottom PMT further collimates the gammas, which have an

interaction length of only 3.0 mm in the surrounding xenon. The ^{133}Ba source is placed at the side of the cryostat above the liquid xenon level — the 356 keV gammas have a 2.5 cm interaction length in liquid xenon, illuminating the entire active volume. The ^{252}Cf is housed in a lead castle $\sim 9''$ from the center of the detector, with $4''$ of lead between the source and the active volume and $6''$ of lead behind. The lead serves to shield the detector from the high rate of gammas coming from the source (the heavy lead nuclei scatter but do not moderate the neutrons).

3.3.6 Data Acquisition

Waveforms from the PMTs are amplified by an SRS 350 MHz, x5 amplifier and digitized using 8 bit 500 MHz Acqiris digitizers. We use two digitizer channels with different full-scales for each PMT to capture the full dynamic range. The digitizers are triggered by a discriminator on the bottom PMT with a threshold at 5 electrons in the S2 signal — this is well below the analysis threshold set by the S1 light collection efficiency. High energy events trigger on S1 before triggering on S2 — the digitized waveforms extend a full drift length on either side of the trigger to catch both pulses regardless of which causes the trigger.

3.3.7 Data Processing

The saved traces pass through a software pulse-finder, which doubles as a compression algorithm. Of the 10–15 μs digitizer trace, typically less than 2 μs contains signal (see Fig. 3.10). The pulse finder identifies pulses and saves a compressed version of the file, recording the waveform segments where there are pulses and the sample at which they occur, while saving only the mean and rms noise of the baseline. The pulse finding algorithm first determines the baseline and noise in the trace using the first microsecond of the waveform, being careful not to include any pulses in the baseline sample. The baseline is subtracted from the trace, and the result is put through a time-domain filter. The filter shape is the derivative of a Gaussian with $\sigma=30$ ns. A threshold (determined by the baseline noise) is set on the absolute value of the filtered pulse, and any piece of the waveform over threshold is identified as part of a pulse. Typically this catches the entire S1, but will occasionally miss the flat peak of an S2, catching only the rising and falling edges. A 40 ns buffer is added to either side of each found pulse piece, catching tails and filling in any gap in the S2. Figure 3.10 shows a sample event with the pulse finding results.

After pulse finding, individual pulses are classified as S1 or S2 based on pulse shape. The S2 pulse is 1 μs wide, set by the time it takes electrons to drift from the liquid surface to the anode

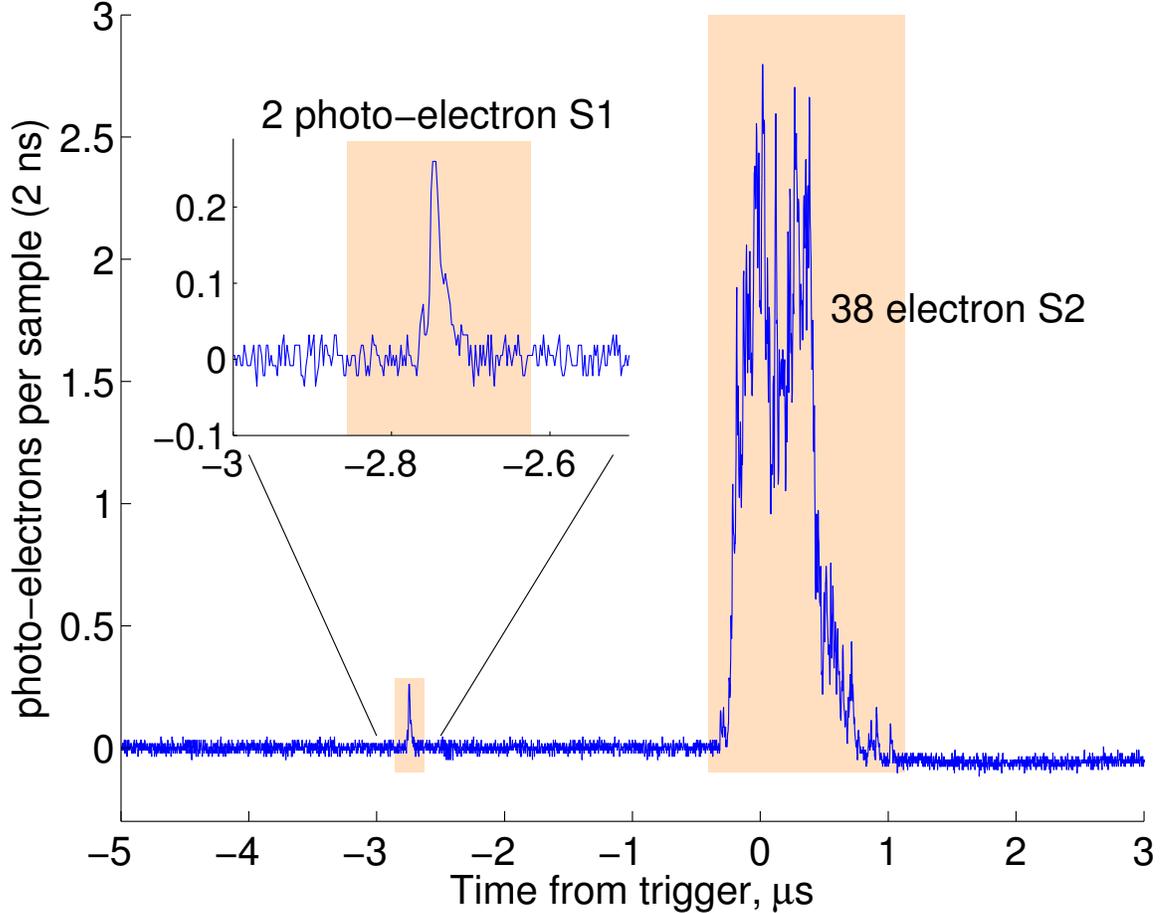


Figure 3.10: Sample waveform from the bottom PMT for a nuclear recoil event in Xed. The shaded regions show where pulses were found, in this case a single S1 and single S2.

grid, and fairly symmetric. The S1 pulse has a ~ 10 ns risetime, set by the PMT risetime, and a ~ 100 ns falltime, set by the recombination timescale — see [44] for details on S1 pulse shapes in liquid xenon.

The integrals of the S1 and S2 pulses gives the light and charge signals, and basic cuts reduce the data to single-scatter events (one S1 and one S2). Depth corrections are applied to the S1 signal to correct for position dependence in the light collection efficiency (with a single top PMT, we can only reconstruct the z-coordinate of the event). After recirculation, no depth dependence is seen in S2. Cuts and corrections were developed by John Kwong and are discussed in [44].

3.3.8 Single-Phase Operation

The detector is occasionally run as an ionization chamber rather than a time projection chamber, with the liquid level raised above the top grid. In this mode, the ‘top’ grid acts as the anode, with

electrons from the active region drifting past both the ‘gate’ and ‘anode’ grids before being collected. A charge sensitive preamp connected to the top grid measures the charge signal directly. This has much lower signal to noise than proportional scintillation but removes dependence on pressure and gas field and allows an absolute calibration of the charge signal.

Single phase data is taken both in the configuration in Fig. 3.6 and in a single PMT configuration with the bottom PMT and bottom grids removed and the cathode grid replaced by a silver-plated copper plate. In the former configuration, the bottom PMT is used to detect S1 while charge data is taken with the pre-amp. In latter, S1 light can still be measured by the top PMT, but charge and light data cannot be taken simultaneously due to cross talk between the PMT and the top wire grid. The advantage of the single PMT setup is the simpler electrode geometry, making it easier to reliably set drift fields. The plate also removes the xenon space below the cathode, where scintillation light is produced but no charge is collected. The cathode plate has a 1 nCi ^{210}Po alpha source deposited in its center, which is used for preamp calibration.

For both configurations, we use an Ortec 142AH preamp to integrate the charge signal. The feedback capacitance (gain) is calibrated through a combination of fall time measurements, gain measurements with pulses through known input capacitances, and measurements of the ionization signal from the alpha source in gas and liquid compared with published results [64]. These measurements are consistent, and give a feedback capacitance of $C_f = 1.15 \pm 0.01$ pF. The preamp output is sent through an Ortec 572 spectroscopy amplifier with 3 μs shaping time and then digitized. For charge-only data in the single PMT configuration, the unshaped preamp output is also digitized.

In single-phase operation we attempt no position reconstruction for events, and so restrict our sources to the ^{57}Co centered below the detector and the ^{210}Po on the cathode in the single-PMT configuration. The 122 and 136 keV gammas from ^{57}Co have interaction lengths in liquid xenon of 3.0 and 3.9 mm, respectively, so that $\sim 96\%$ of these events occur below the ‘gate’ grid (alphas from the ^{210}Po travel only ~ 45 μm , so are also contained in the active region). The ‘gate’ and ‘anode’ grids shield the ‘top’ from slow moving positive ions so the charge signal is independent of the depth of the event (see Appendix A). The fields in the gate-anode and anode-top regions are kept constant, fixing the velocities of electrons in the regions visible to the preamp. The maximum active region drift field is 2 kV/cm, with fields of 4 and 8 kV/cm in the regions above to ensure transparency of the grids to drifting electrons. For light-only data, neither drift speed nor transparency is an issue, so all fields are kept proportional to the field in the active region.

3.3.9 Other Xed Configurations

The Xed detector ran in many other configurations besides the two described. It was used to investigate the use of CsI photocathodes inside the chamber for detecting S1 light, proportional gas gain at the anode for measuring the S2 signal, and the use of waveshifters to increase S1 light collection efficiency. These projects are described by John Kwong in [44]. We will focus on the calibration, recombination, and discrimination measurements made in the single- and dual-phase modes described above.

Chapter 4

Recombination and Calibration

The charge and light signals produced in liquid xenon are best understood in terms of ion recombination at the event site. Recombining ions contribute to, and in fact create most of the primary scintillation light, so that the total ionization signal is split between the charge and light signals. Variations in the recombination fraction induce an exact anti-correlation between the two signals. This anti-correlation may be exploited for calibration and used to distinguish fluctuations in the recombination fraction from the uncorrelated fluctuations associated with measurement of the charge and light signals.

This chapter presents the basic recombination model for signal production and demonstrates the power of the model analysing 122 keV gamma events from a ^{57}Co calibration source in the Case prototype, Xed. We demonstrate the signature anti-correlation by varying the applied drift field and thus the mean recombination fraction, and from this calculate the W-value (average energy expended per ion or exciton) and S1 light collection efficiency. We then decompose the fluctuations seen in the 122 keV peak into uncorrelated S1 and S2 fluctuations and anti-correlated recombination fluctuations, discussing the physics behind each type of fluctuation.

4.1 Recombination Model for Signal Production in Liquid Xenon

The recombination model for signal production in liquid xenon is not new, but past treatments of the model have missed its broad implications. Recombination is often mentioned in the context of energy resolution [65, 66], while the applications in calibration and fluctuation decomposition in

dual-channel detectors have been missed. This is at least partly due to the lack of a precise statement of the model in the literature. We therefore begin this chapter with a careful discussion of signal production in liquid xenon, building the framework within which the measurements in this chapter and the next should be understood.

4.1.1 Signal Production in Electron Recoils

Interactions in liquid xenon create tracks of excited and ionized xenon atoms, with ten to twenty times as many ions as excitons (excited atoms). Excitons form an excited dimer, Xe_2^* , which emits a 175 nm (VUV) photon as it decays [67]. If an ionized electron recombines with a xenon ion, the resulting exciton will also go through this process, so that the total number of scintillation photons will be

$$N_{ph} = a \cdot N_{ex} + b \cdot r \cdot N_i, \quad (4.1)$$

where N_{ph} is the number of scintillation photons, N_{ex} is the number of direct excitons in the track, N_i is the number of ions in the track, r is the fraction of ions that recombine, and a and b are the efficiency with which direct excitons and recombined ions produce scintillation photons, respectively. We expect $a \approx b \approx 1$, but we include them as coefficients here so it is clear later which results are sensitive to these efficiencies. For electron recoils, we define the W-values

$$W_q = E/N_i \quad (4.2)$$

and

$$W_{ph} = E/(a \cdot N_{ex} + b \cdot N_i), \quad (4.3)$$

where E is the energy of the recoil, so that E/W_q gives the total charge yield with zero recombination, and E/W_{ph} gives the total light yield with full recombination. Both W_q and W_{ph} are expected to be energy independent [68]. Due to photons from direct excitation,

$$W_{ph} = W_q / \left(a \cdot \frac{N_{ex}}{N_i} + b \right) \approx W_q / 1.06, \quad (4.4)$$

where the value $\frac{N_{ex}}{N_i} = 0.06$ is taken from theoretical calculation based on absorption measurements and oscillator strengths in solid phase [69, 70]. In the literature, a third W-value for light yield at zero field is often cited. This value is generally 10–50% higher than W_{ph} due to incomplete recombination at zero field as electrons diffuse away from the track [68]. Since it depends on the

recombination fraction, the zero-field W-value varies with recoil energy and track type.

When an electric field is applied, electrons that do not recombine may be extracted from the track and measured separately. The amount of charge extracted is

$$N_q = (1 - r)N_i, \quad (4.5)$$

and, combining Eq. (4.1,4.3,4.5), we find the recombination independent sum

$$E = (N_q + (N_{ph}/b)) \cdot (W_{ph} \cdot b), \quad (4.6)$$

where we have moved all dependence on b into terms with N_{ph} and W_{ph} . This grouping of terms is convenient since absolutely calibrating N_q is much easier than N_{ph} (the former requires only a calibrated preamp, while the latter requires both a photo-sensor calibrated at liquid xenon temperatures and an accurate calculation of light collection efficiency in the chamber, neither of which has yet been achieved). We rename the pieces of Eq. (4.6): $n_e = N_q$, $n_\gamma = N_{ph}/b$ and $W = W_{ph} \cdot b$, giving

$$n_e = (1 - r) \cdot N_i, \quad (4.7)$$

$$n_\gamma = \left(\frac{a}{b} \frac{N_{ex}}{N_i} + r \right) \cdot N_i, \text{ and} \quad (4.8)$$

$$E = (n_e + n_\gamma) \cdot W. \quad (4.9)$$

In words, one electron recombining will reduce n_e by one and increase n_γ by one. (Strictly speaking, a recombining electron has probability b of increasing n_γ by b^{-1} and probability $1 - b$ of leaving it unchanged, giving an average increase in n_γ of one.) We write the units for n_e and n_γ as $[n_e]$ and $[n_\gamma]$, read as numbers of extracted and recombined electrons, respectively. For example, a calibration constant for the scintillation signal may have units $\text{phe}/[n_\gamma]$, meaning photo-electrons per recombined electron. Note that n_γ itself is not the actual number of recombining electrons, as it also includes light from direct excitation, but rather the number of recombining electrons required to produce the observed scintillation signal in the absence of direct excitation. Note also that a and b do not explicitly appear in Eq. (4.9), and no knowledge of a and b is necessary to construct the recombination independent energy scale. We will construct this scale in section 4.2, finding $W = 13.7 \pm 0.2$ eV.

Equations (4.7)–(4.9) describe the partition of electronic excitation into charge and scintillation

signals, given by the initial excitation-ion ratio and recombination fraction. The initial partition into direct excitation and ionization is expected to be independent of energy [71], so the term $\frac{a}{b} \frac{N_{ex}}{N_i}$ should be constant (although it may vary with recoil type, as discussed below). The recombination term depends on the electric field and ionization density in the track, with greater recombination at low fields and in dense tracks. Locally, the ionization density is determined by the electronic stopping power for the recoiling particle, while the local electric field is the applied drift field minus any screening from the ionization in the track. Globally, both of these are influenced by the track structure: since stopping power is energy dependent, the division of energy among daughter recoils will affect the ionization density distribution, and track geometry may impact screening of the electric field. The fraction of charge recombining for electron and nuclear recoils versus energy and drift field is given in Chapter 5, and a model for recombination is presented in Chapter 6.

4.1.2 Nuclear Recoils

In electron recoils the entire recoil energy is lost through electronic excitation, but nuclear recoils also lose energy through elastic collisions with other nuclei. While some of these nuclei create further ionization, most have energies below the ionization threshold. This energy loss is not detectable in liquid xenon, so we see an overall suppression of electronic excitation for nuclear recoils. Lindhard derives the integro-differential equation describing the total electronic energy loss in a nuclear recoil cascade in terms of the differential elastic nuclear scattering cross-section and electronic stopping power for recoiling nuclei [37]. For low-velocity ions ($v < c/137$) Lindhard takes the electronic stopping power to be

$$S = 4\pi\hbar ca_0 Z^{7/6} \sqrt{\frac{E}{Mc^2}}, \quad (4.10)$$

where Z is the atomic number of both the recoiling nucleus and the atoms in the medium, M and E are the mass and kinetic energy of the recoiling nucleus, and a_0 is the Bohr radius. The differential scattering cross section is given by

$$d\sigma = \pi a^2 \frac{d\eta}{\eta^2} f(\eta), \quad (4.11)$$

where

$$\eta = \frac{4\pi\epsilon_0 a}{2e^2 Z^2} \sqrt{EE_r} \quad (4.12)$$

for incident energy E , energy transfer E_r , and screening radius a given by

$$a = 0.626a_0 Z^{-1/3}. \quad (4.13)$$

Lindhard uses the Thomas-Fermi model to compute $f(\eta)$, which is approximated by

$$f(\eta) = \frac{\lambda\eta^{1-2m}}{(1 + (2\lambda\eta^{2(1-m)})^q)^{1/q}} \quad (4.14)$$

with $m = 1/3$, $q = 2/3$, and $\lambda = 1.309$ [72]. Other approximations for stopping powers, differential cross sections, and screening radii are listed in [73].

Lindhard's numerical solution is approximated in [9] by

$$\mathcal{L} = \frac{kg(\epsilon)}{1 + kg(\epsilon)}, \quad (4.15)$$

where the energy dependent factor \mathcal{L} is the fraction of the initial recoil energy eventually lost to electronic excitation, with

$$\epsilon = 11.5 (E_{nr}/keV) Z^{(-7/3)}, \quad (4.16)$$

$$k = 0.133 Z^{2/3} A^{-1/2}, \text{ and} \quad (4.17)$$

$$g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon, \quad (4.18)$$

for a recoil of energy E_{nr} . Making the commonly-used assumption that the portion of energy lost electronically creates excitation and ionization with the same efficiency (W-values) as electron recoils, we write

$$E_{nr} = \mathcal{L}^{-1} \cdot (n_e + n_\gamma) \cdot W. \quad (4.19)$$

Equation (4.15) gives the generic prediction for \mathcal{L} in any material, and it accurately describes ionization from nuclear recoils in semiconductors [74]. The prediction in liquid xenon is more complicated. There is some evidence from our recombination modeling results (see Chapter 6) that $\frac{N_{ex}}{N_i}$ is higher for nuclear recoils than electron recoils. This would lower the W-value for electronic energy loss in nuclear recoils (the W in Eq. (4.19)), since less energy is required to produce an exciton than an ion. Experimentally this has the same signature as an increase in the Lindhard parameter (i.e., an increase in the total electronic energy loss) and can be folded into \mathcal{L} in Eq. (4.19). Except where explicitly stated otherwise, we assume the same W-value for electron recoils and electronic excitation in nuclear recoils when discussing \mathcal{L} .

Instead of the Lindhard factor \mathcal{L} , many experiments measure the ‘effective Lindhard factor’ \mathcal{L}_{eff} , also called ‘relative scintillation efficiency’ [75, 76, 77, 78, 79, 80] (see Fig. 4.1). This factor is defined

as the ratio of the zero-field light yields of electron and nuclear recoils,

$$\mathcal{L}_{eff} = \frac{n_{\gamma nr}}{n_{\gamma er}} \frac{E_{er}}{E_{nr}} \text{ at zero field,} \quad (4.20)$$

where the *nr* and *er* subscripts indicate the nuclear recoil and a reference electron recoil. A standard reference is necessary, since the zero-field $\frac{E_{er}}{n_{\gamma er}}$ is energy dependent — to date this reference has been a 122 keV photo-absorbed gamma from a ^{57}Co source. E_{nr} is usually determined kinematically by measuring the scattering angle of a neutron with known initial energy, with the exception of [79], where \mathcal{L}_{eff} is found from the n_{γ} spectrum at 730 V/cm from an AmBe neutron source compared with the Monte Carlo generated recoil spectrum. This measurement is adjusted to zero-field values using the field dependence of the 56.5 keV nuclear recoil light yield reported in [77].

If recombination were complete at zero field, one would see $\mathcal{L}_{eff} = \mathcal{L}$, but experimentally neither nuclear nor electron recoils undergo complete recombination at zero field. At this time no experiments have published results measuring \mathcal{L} directly, e.g., by measuring both n_{γ} and n_e for kinematically known recoils (this data should appear in [81]), but measurements giving both n_{γ} and n_e for nuclear recoils of unknown energy can be used to relate \mathcal{L}_{eff} to \mathcal{L} . For a given \mathcal{L}_{eff} , the recoil energy of an event can be determined from n_{γ} , taking the measured field dependence for the 56.5 keVr nuclear recoil light yield in [77] to shift from zero-field to at-field light yields. The equivalent electron energy of the event is given by Eq. (4.9). By definition, the ratio of these energies is \mathcal{L} . This is shown in Fig. 4.2, where we use the calibrated neutron measurements in Chapter 5 to compute \mathcal{L} from Sorensen’s \mathcal{L}_{eff} [79]. We use Sorensen’s values because they are based on data taken at field (730 V/cm) — since our data is taken at a similar field, this avoids most of the systematic uncertainty in the field dependence of the light yield for nuclear recoils (i.e., when shifting Sorensen’s values from the zero-field \mathcal{L}_{eff} to the light yield at our field, we are merely undoing the adjustment Sorensen made to produce his zero-field values). The resulting \mathcal{L} is more consistent with a constant Lindhard factor of 0.25 than the \mathcal{L} falling with energy described in Eq. (4.15). Turning the process around, we compute \mathcal{L}_{eff} from the neutron data using both $\mathcal{L} = 0.25$ and the \mathcal{L} given by Eq. (4.15) (Fig. 4.3). The constant Lindhard factor is consistent with Sorensen’s spline, but the theoretical factor is not. This implies that the drop in \mathcal{L}_{eff} from 0.24 around 50 keVr to 0.15 below 10 keVr is due not to a drop in the actual Lindhard value but to an increased fraction of the signal lost to escaping electrons at zero field.

Constant Lindhard factors have also been seen in some scintillators, including NaI [82] and CaF₂ [83], but it is not clear why this should be the case in liquid xenon. One possible explanation

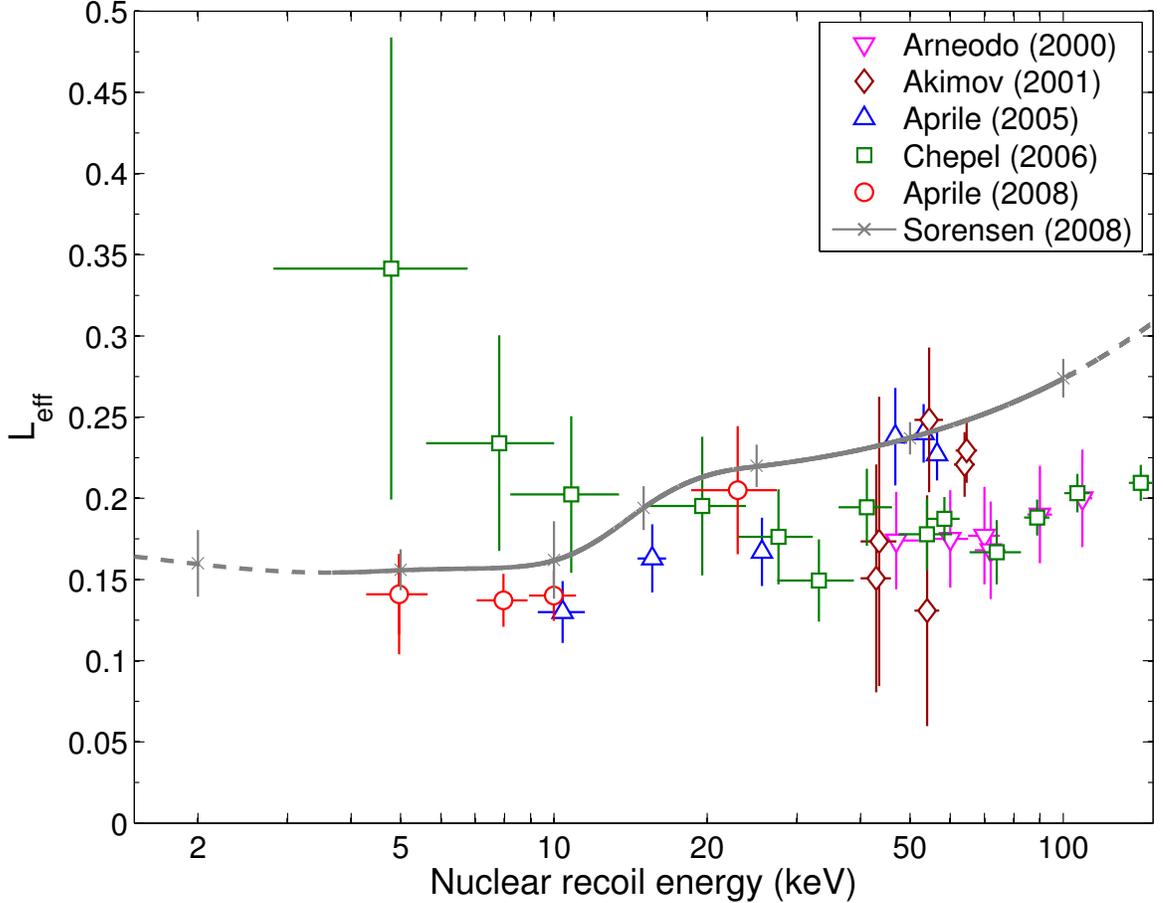


Figure 4.1: Relative scintillation yield (\mathcal{L}_{eff}) of nuclear recoils in liquid xenon, referenced to 122 keV gamma rays from ^{57}Co . The data points are from neutron scattering experiments at zero field, with nuclear recoil energies determined kinematically from the scattering angle [75, 76, 77, 78, 80]. The gray curve is the spline fit for \mathcal{L}_{eff} in [79], which is based on a fit to a recoil spectrum from an AmBe neutron source. The errorbars along the curve indicate the knots in the spline and their uncertainties, and the dashed parts of the curve indicate regions not used in the fit. The Sorensen curve is based on data taken at 730 V/cm and adjusted to zero-field using the field dependence of the 56.5 keV recoil light yield reported in [77]. This plot was compiled by Aaron Manalaysay.

is bi-excitonic quenching — in this process, the usual $\text{Xe} + \text{Xe}^* \rightarrow \text{Xe}_2^* \rightarrow 2\text{Xe} + \gamma$ is replaced with $\text{Xe}^* + \text{Xe}^* \rightarrow \text{Xe}_2^{**} \rightarrow \text{Xe} + \text{Xe}^+ + e^-$. This electron may then recombine to produce a single scintillation photon, giving one photon from two excitons. This process applies to both direct excitons and excitons following recombination, as described in [84], and has been observed in tracks with extremely dense excitation and ionization — alpha tracks experience a 23% reduction in scintillation yield attributed to this effect [68]. If present in nuclear recoils, as is suggested in [85], bi-excitonic quenching would be stronger for higher energy recoils because of their higher electronic stopping power (see Section 4.2.5), tending to flatten the measured Lindhard value — but it would

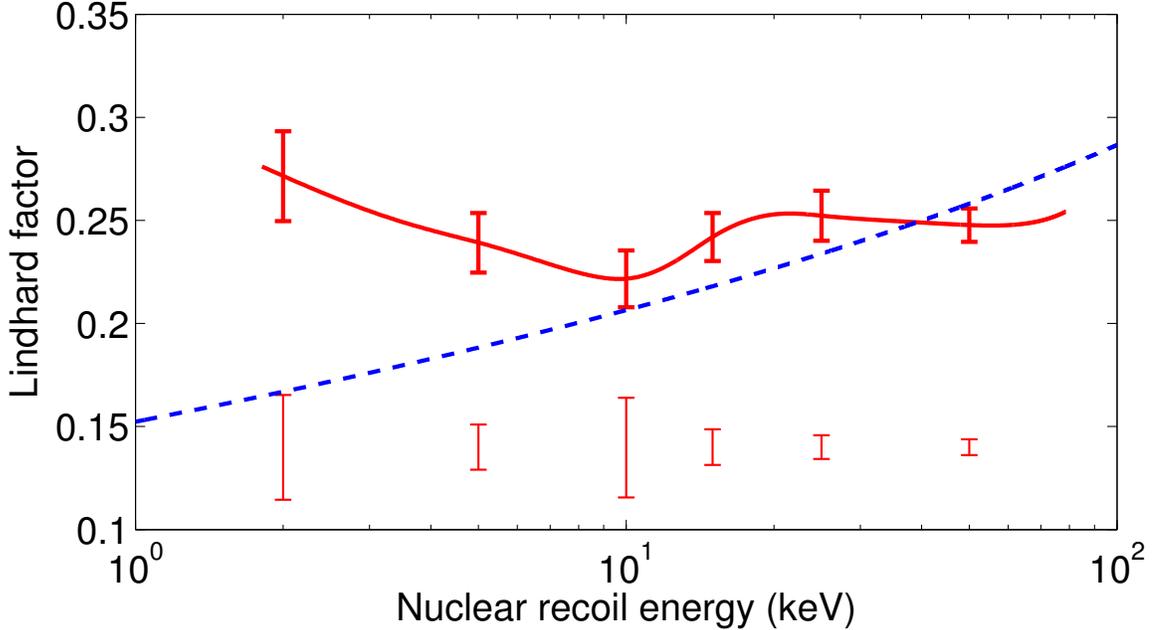


Figure 4.2: The Lindhard factor, or fraction of energy lost to electrons in nuclear recoils, versus nuclear recoil energy. The solid red line takes the nuclear recoil energy determined from n_γ using the \mathcal{L}_{eff} given by Sorensen [79], while energy lost to electrons is given by $W(n_\gamma + n_e)$. Statistical errorbars on the curve come from the uncertainty in \mathcal{L}_{eff} . Systematic errors (shown below) include the systematics quoted by Sorensen and the uncertainty in field-dependence of the light yield. The blue dashed line is the theoretical Lindhard factor as given in Lewin and Smith [9].

also result in a measured value lower than the theoretical prediction. We see a Lindhard value higher than the theoretical prediction, but this could be due to a higher direct-exciton fraction and therefore lower W-value for nuclear recoils. Some combination of these effects may describe the nuclear recoil signal. Dual channel (charge and light) measurements with known energy nuclear recoils are needed to resolve this issue. Not only will this measure \mathcal{L} directly, but if bi-excitonic quenching is a factor, then the efficiency b for creating scintillation via recombination will be lower for nuclear recoils than electron recoils. This would be apparent in anti-correlation measurements, such as those in Section 4.2.

4.1.3 Energy Scales

There is historical confusion on the construction of energy scales in liquid xenon detectors. The custom has been to use n_γ as the energy scale, a practice inherited from detectors where only the light channel was measured. From Eq. (4.8), we see that a linear scale in n_γ will only work for electron recoils if the recombination fraction r is energy independent. This is not the case, even at zero field, especially when extrapolating from a calibration source (typically > 100 keV) down to

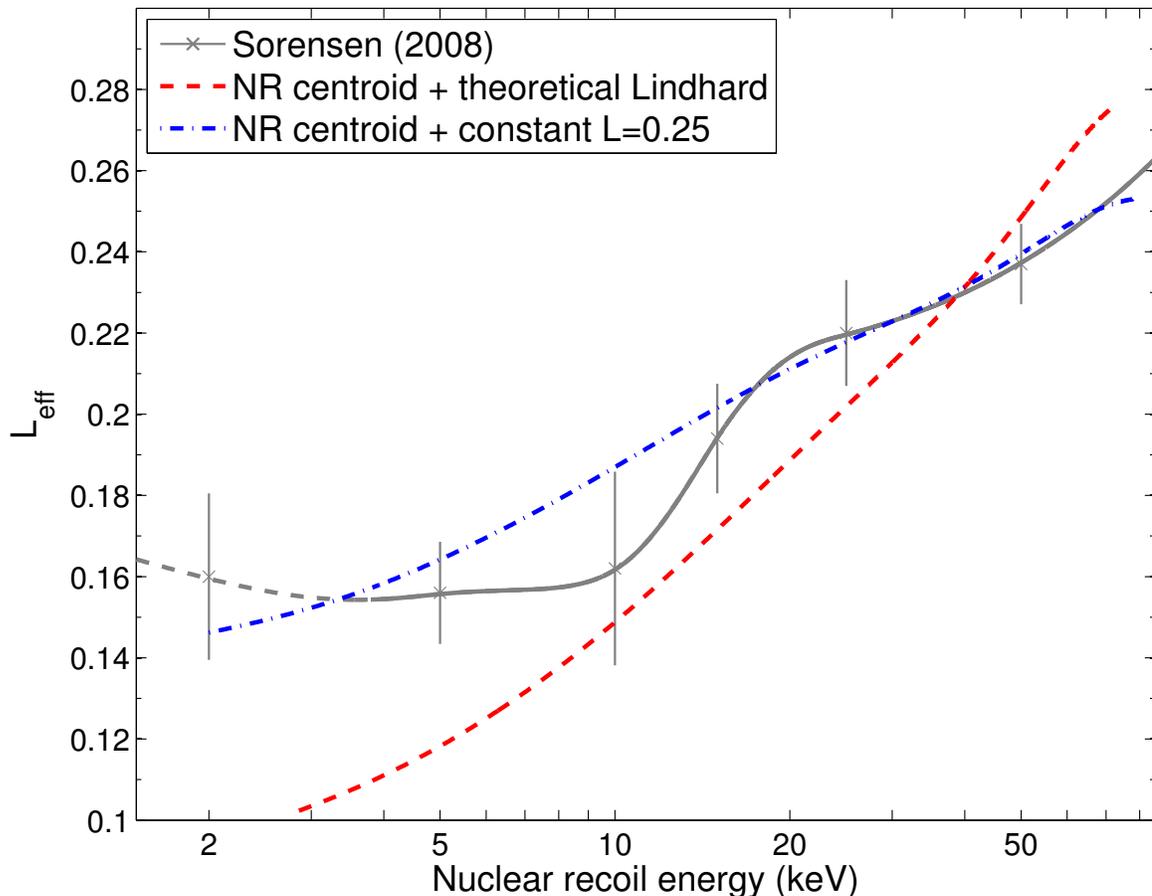


Figure 4.3: Relative scintillation yield (\mathcal{L}_{eff}) based on the measured recombination in nuclear recoils in Xed at 520 V/cm and the theoretical value for \mathcal{L} (red $--$) and $\mathcal{L} = 0.25$ (blue $- \cdot -$). Also shown is the spline fit to \mathcal{L}_{eff} from [79], which is based on nuclear recoil data from Xenon10 at 730 V/cm.

WIMP recoil energies. In fact, even a non-linear energy scale based on n_γ is inadequate, because r is not single-valued versus energy. In particular, as will see in Chapters 5 and 6, a Compton scatter and photo-absorption event with the same total energy will have different recombination fractions, and thus different n_γ 's. The combined energy defined in Eq. (4.9) is recombination independent, and therefore both single valued and linear for electron recoils. In this thesis ‘keVee’ refers to a measured electron equivalent energy using Eq. (4.9).

For nuclear recoils, there is call for an energy scale based on n_γ because the existing calibrations for nuclear recoils are in terms of n_γ . This requires choosing a \mathcal{L}_{eff} and a reference electron recoil (to date a 122 keV photo-absorption event from ^{57}Co). On the other hand, it is also desirable to have a recombination independent energy scale for nuclear recoils, as in Eq. (4.19). This requires choosing an \mathcal{L} . (There is also an implicit assumption of a common W-value for electronic energy loss in electron and nuclear recoils, which may be considered part of the definition of \mathcal{L} .) In this

thesis, ‘keVr’ refers to a measured nuclear recoil energy using Eq. (4.19) with $\mathcal{L} = 0.25$.

All n_γ based energies in this thesis will be clearly labeled as such, with the \mathcal{L}_{eff} or calibration point used to construct the scale explicitly stated. We will also use ‘keV’ when referring to physical (as opposed to reconstructed) energies.

4.2 Calibrating Recombination Measurements

From Eq. (4.9) we see that the sum $n_\gamma + n_e$ is independent of the recombination fraction due to the exact anti-correlation between the charge and light signals. We can turn this around and determine the relative calibration of our measured charge and light signals by varying recombination and finding the recombination independent sum of the two signals. For many measurements this relative calibration is all that is necessary, but we can also use the relative calibration to translate an absolute charge calibration into an absolute light calibration and find the W-value in Eq. (4.9). This section describes the method for calibrating from the charge-light anti-correlation and its application in three configurations of Xed.

4.2.1 Method for Calibrating Using Charge-Light Anti-correlation

We define the normalized charge yield, y , as

$$y = \frac{n_e}{n_\gamma + n_e}. \quad (4.21)$$

We can calibrate y using the anti-correlation between charge and light with varying recombination, making it easier to work with than the actual recombination fraction r . The two are related by

$$r = 1 - \left(\frac{a}{b} \frac{N_{ex}}{N_i} + 1 \right) y, \quad (4.22)$$

containing the difficult to measure $\frac{a}{b} \frac{N_{ex}}{N_i}$. When a recombination fraction is necessary, we take $\frac{a}{b} \frac{N_{ex}}{N_i} = 0.06$, corresponding to the theoretical value for $\frac{N_{ex}}{N_i}$ and $a = b = 1$.

We define the overall measurement gains g_1 and g_2 such that

$$\begin{aligned} S1 &= g_1 \cdot n_\gamma \text{ and} \\ S2 &= g_2 \cdot n_e, \end{aligned} \quad (4.23)$$

where $S1$ and $S2$ are the uncalibrated scintillation and charge signals, respectively. To determine y

from the measured $S1$ and $S2$, we need only the ratio of gains $\frac{g_1}{g_2}$, since

$$y = \frac{S2}{S2 + \frac{g_2}{g_1} S1}. \quad (4.24)$$

We also have, for electron recoils,

$$E \propto \left(S2 + \frac{g_2}{g_1} S1 \right), \quad (4.25)$$

from Eq. (4.9,4.23). This ratio of gains may be determined by measuring $S1$ and $S2$ for two or more peaks with known energies and different values of y . For two peaks,

$$\frac{g_2}{g_1} = \frac{E_a S2_b - E_b S2_a}{E_b S1_a - E_a S1_b}, \quad (4.26)$$

where the subscripts indicate the two peaks being considered. This gives an uncertainty inversely proportional to the separation in y ,

$$\delta [\ln (g_2/g_1)] \propto \frac{1}{|y_a - y_b|}. \quad (4.27)$$

For multiple peaks a least squares approach is needed, but again the measurement depends on finding peaks with good separation in y .

The simplest way to apply Eq. (4.26) is to look at a single peak while varying the applied drift field. Energy then falls out of the equation, and $\frac{g_2}{g_1}$ is simply the (negative) slope of a plot of $S2$ vs $S1$. This is shown in Fig. 4.4 for three sets of measurements of 122 keV ^{57}Co gammas, with drift fields ranging from a few V/cm to 4.4 kV/cm and y from 0.11 to 0.67. The three datasets correspond to the three configurations for Xed described in Chapter 3: single phase separate charge and light measurements, single phase simultaneous charge and light measurements, and dual phase simultaneous measurements.

4.2.2 Single-phase, Separate Charge and Light

The single-phase separate charge and light measurement data is taken in the single PMT configuration of Xed (see Chapter 3). For charge data, this has the smallest field range of the three ^{57}Co datasets, limited at the high end to 2 kV/cm by the high voltage requirements (electrons are collected on the ‘top’ grid rather than the ‘anode’ for direct charge readout, and the extra drift distance limits the fields available) and at the low end to 400 V/cm, below which the shrinking charge signal is unmeasurable in the face of electronics noise. The advantage of this mode is the absolute

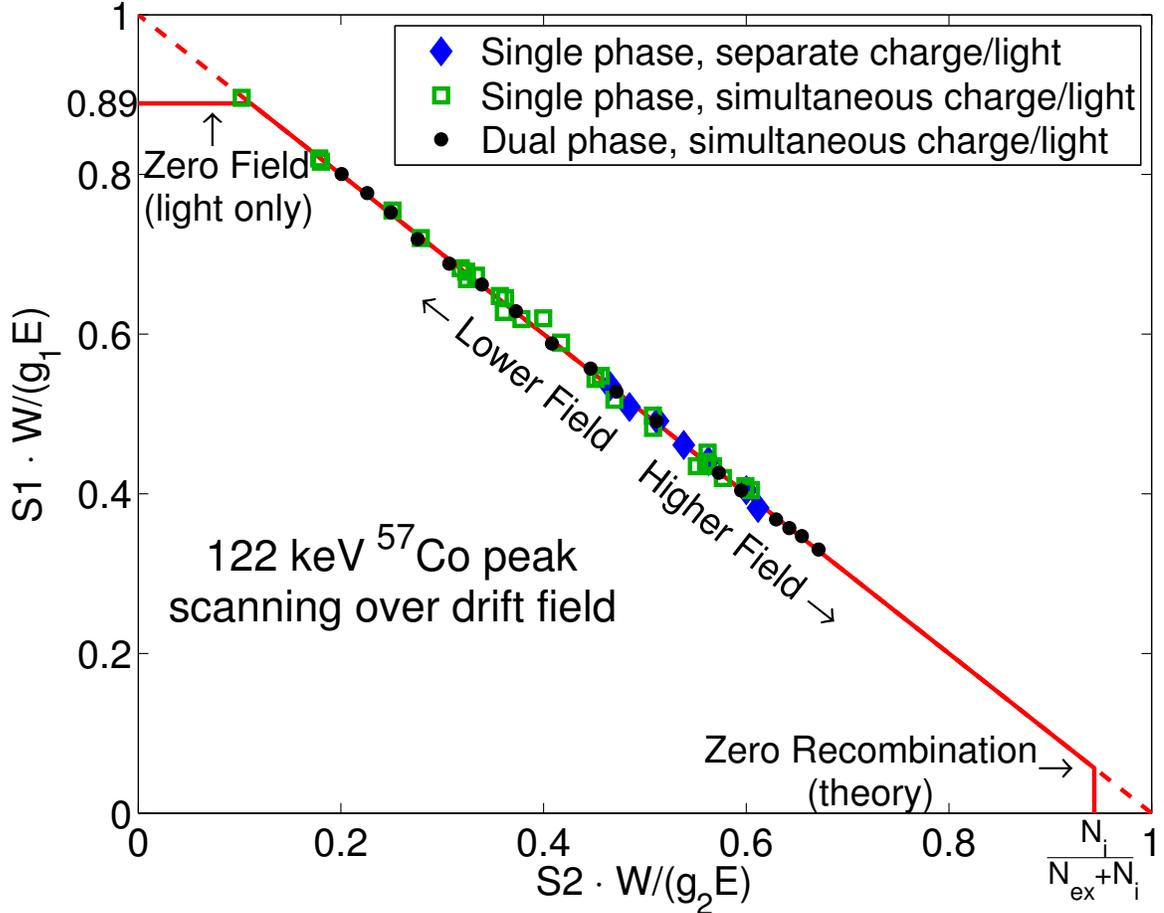


Figure 4.4: Light output vs charge output for 122 keV gamma events at drift fields from 0–4.4 kV/cm. For each of the three datasets, the axes are scaled so that the intercepts of the least-squares linear fit are at 1. E and g_2 are known, allowing one to determine W (13.7 ± 0.2 eV) and g_1 (0.060 ± 0.001 phe/ $[n_\gamma]$ in dual phase) from the unscaled intercepts.

calibration of the charge measurement. To determine charge yields we first perform a Gaussian fit to the background-subtracted specamp spectrum, as shown in Fig. 4.6. We use the Gaussian fit to determine the likelihood for each event to have come from the ^{57}Co source, and average the unshaped preamp traces weighted by this likelihood. The charge signal is read off the mean preamp trace, making g_2^{-1} the preamp feedback capacitance, $C_f = 1.15 \pm 0.01$ pF, measured in Chapter 3. There is an additional factor of 0.874 in g_2 when the spec-amp is connected due to its low input impedance, but this is easily calibrated.

Light data was taken separately, since cross-talk from the single PMT above the top grid swamped the charge signal when the PMT was biased. Light-only data may also be taken over a wider field range, which we use to our advantage in Section 4.2.5. Since there is no need to drift electrons, the drift field only needs to be maintained over the active region, and fields of 4.5 kV/cm are

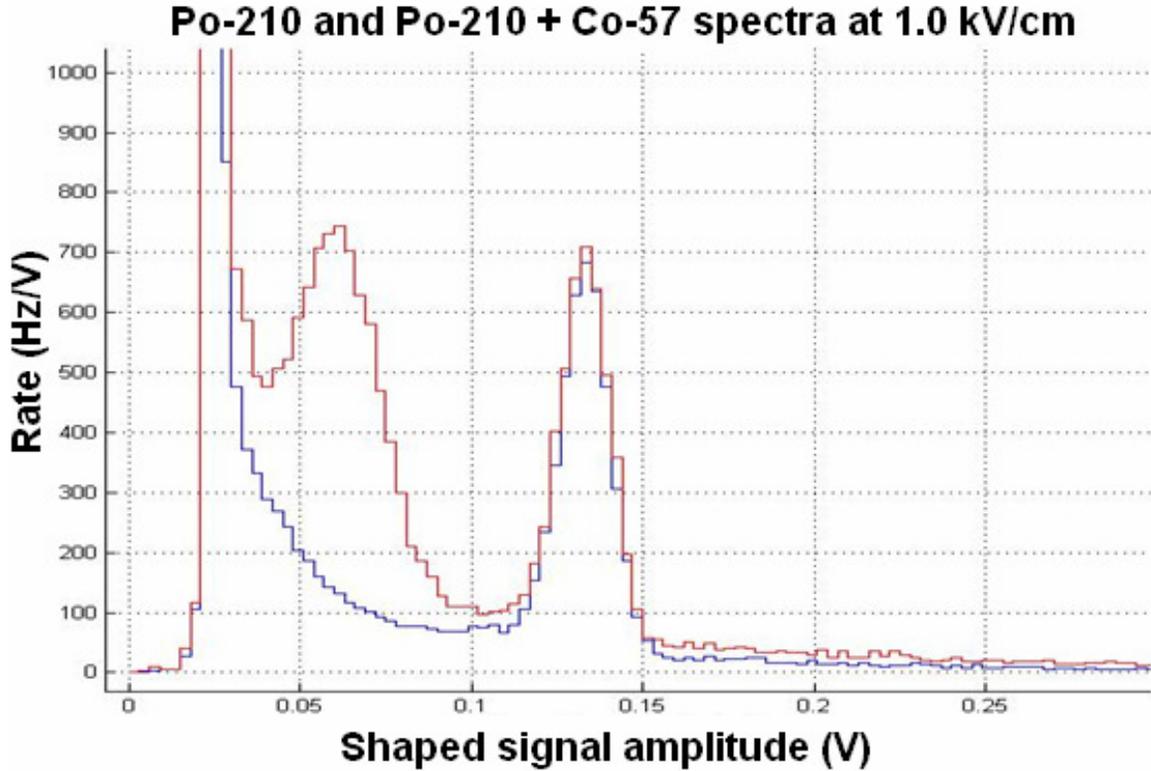


Figure 4.5:]
 Spec-amp spectrum at 1 kV/cm with and without the ^{57}Co source. The higher peak is from the ^{210}Po alpha source inside the detector.

achieved. There is no lower bound to the drift field — the light signal increases as field decreases. The systematic errors in the drift field associated with large field changes across electrodes are eliminated by keeping all fields proportional to the drift field (again, this is possible because there are no requirements on the fields outside the active region). There is also no need for position reconstruction in this setup, since the ^{57}Co gammas are collimated in x-y, cannot interact to produce visible S1 light until they are above the cathode, and have attenuation lengths in liquid of 3.0 and 3.9 mm for the 122 keV and 136 keV lines, respectively, so that events are concentrated at the bottom center of the active region. This is in fact the only setup in which we can reliably measure light yield at zero drift field. In the 2 PMT setup there is no way to reject events below the cathode at zero field, and these events have significantly higher light collection efficiency in the bottom PMT, skewing the light yield measurement.

Plotting the light versus charge at fields where both are available, the data falls on a line as expected from equation (4.9). The independent measurement of g_2 lets us extract W from the (unscaled) x-intercept of the linear fit in Fig. 4.4. This is complicated by the fact that the ^{57}Co

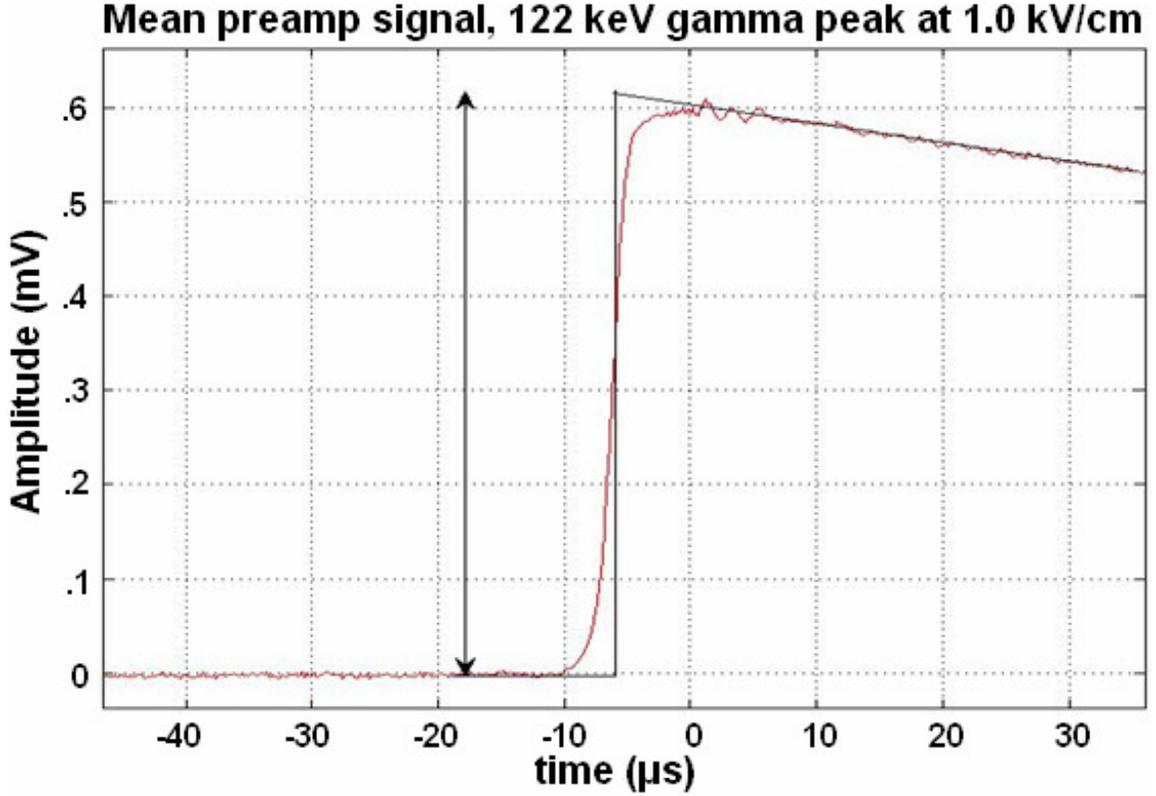


Figure 4.6: The average preamp trace for ^{57}Co events at 1 kV/cm. The height of the step is equal to $\frac{xne}{C_f}$, where $C_f = 1.15 \pm 0.01\text{pF}$ is the feedback capacitance of the preamp, and $x = 0.874$ is the measured drawdown when connecting the spec-amp (which does not have a high-impedance input).

source emits gammas at both 122 and 136 keV, which cannot be distinguished with the relatively poor resolution we see in this mode. At the source, 11% of the gammas are 136 keV, giving a calculated 13% 136 keV piece in the active volume due to the greater penetration of the higher energy gamma through the aluminum, stainless steel, liquid xenon, PTFE, and copper between the source and the active volume. Using a mean event energy of 123.92 ± 0.03 keV, we find $W = 13.7 \pm 0.3$ eV, where the error in W comes primarily from the statistical error in the linear fit. This agrees well with the 13.8 ± 0.9 eV reported by T. Doke et al. for $W_{ph,max}$, the maximum light yield seen from relativistic heavy ions in liquid xenon at zero field [68]. Doke interprets this maximum as corresponding to full recombination and zero bi-excitonic quenching, corresponding to our $W_{ph} = W/b$.

4.2.3 Single-phase, Simultaneous Charge and Light

The single-phase simultaneous charge/light data can push to lower field by using S1 values to select ^{57}Co events, thereby improving the signal-to-noise in the charge channel. The actual fields achieved

at low field settings are unknown because geometric tolerances lead to uncertainties in how much of the high field above the ‘gate’ grid bleeds through to the active region. This configuration also seems to suffer from charge build up on insulating surfaces, evidenced by a hysteresis in the drift field. To reach low fields, the nominal drift field is varied until the charge signal disappears, indicating that the active region field is reversed.

Due to DAQ limitations at the time of taking this data, the unshaped preamp output was not digitized in this mode. Instead, the charge signal is determined for each event from the amplitude of the shaped signal from the spectroscopy amplifier. We use a template fit to give an unbiased estimate Q of the signal amplitude,

$$Q = \frac{\sum (S_i T_i)}{\sum (T_i^2)} \cdot \max(T_i), \quad (4.28)$$

where the waveform S_i is the digitized amplifier output, and T_i is the template waveform constructed from the average ^{57}Co signal. The template is the same for each drift field, but it is translated in time to match the mean ^{57}Co peak time. The digitizers for both the light signal and charge signal are triggered from the PMT (S1), so the timing of the charge signal depends on the electron drift velocity in the active region. The fields above the active region are kept fixed, so that the electron velocity in regions visible to the readout (top) grid are independent of the applied drift field.

The challenge at low fields is to distinguish ^{57}Co events (with small charge signals) from events occurring in the liquid below the cathode (which give zero charge signal, since the electrons drift the wrong direction). To accomplish this, we take advantage of the dual-channel data by finding the ^{57}Co peak in 2-D, as in Fig. 4.7. We fit to a 2-D Gaussian in Q and $S1$ (taken as the integral of the PMT signal) plus a fixed background of events with zero charge signal and a variable flat background in the $S1$ - Q plane. The zero-charge background from events below the cathode is the product of the noise peak in Q and the spectrum in $S1$ of these events. We characterize the Q noise peak by looking at 122 keV events below the cathode, which are guaranteed to have zero charge signal. These comprise over half of the total triggers and are always well separated from the active region ^{57}Co peak due to the high (reverse) field below the cathode — at low fields, active region events have a higher light yield than events below the cathode, while at high fields they are separated by the charge signal. The noise peak was observed to follow a logistic distribution,

$$P(x) = \frac{d}{dx} \frac{1}{1 + e^{-(x-\mu)/s}}, \quad (4.29)$$

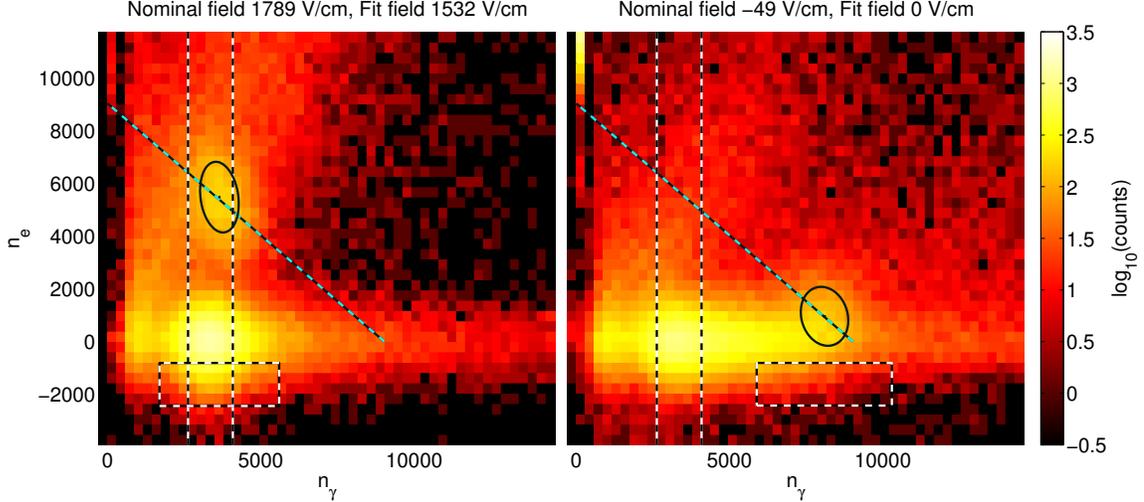


Figure 4.7: Single phase, simultaneous charge and light measurements of the 122 keV (and 136 keV) ^{57}Co peak. The left plot shows data at the highest field reached in this mode, while the right shows the lowest field with non-zero charge signal. The fit drift field is determined by matching the light yield (scaled by the intercept of the linear fit in Fig. 4.4) to the yield-vs-field found in Section 4.2.5. The charge and light yields in the right plot are consistent with what is observed at zero-field. Lowering the nominal field further results in zero charge signal and lower light yield, indicating a rising reverse drift field in the active region. The ellipses show the 2-D Gaussian fits to the peak. The ellipse is defined by $(\mathbf{x} - \boldsymbol{\mu})^\top \cdot \boldsymbol{\Sigma}^{-1} \cdot (\mathbf{x} - \boldsymbol{\mu}) = 1$, where $\boldsymbol{\mu}$ is the mean and $\boldsymbol{\Sigma}$ is the covariance matrix. (This is equivalent to drawing points at $x = \mu \pm \sigma$ in a 1-D distribution.) The vertical black-and-white dashed lines show the slice used to determine the charge noise profile, and the black-and-white box indicates the region used to find the zero-charge $S1$ spectrum. Noise profiles and $S1$ spectra are shown in Fig. 4.8. The diagonal dashed line is the linear fit to peak positions at different fields, as in Fig. 4.4. Histogram bins are 292 [n_γ] by 391 [n_e]. Calibration to n_e and n_γ is done from the intercepts of the linear fit, using the W found in the previous section and an estimated average energy of 124 keV.

as shown in Fig. 4.8. The $S1$ spectrum of zero-charge events is found by taking a slice in Q from 1 to 3σ below the noise peak (where $\sigma = \frac{\pi}{\sqrt{3}}s$ is the square root of the variance of the logistic distribution). We fit the spectrum to a polynomial, and take the product with the noise peak to fix the 2-D zero-charge background. An additional variable flat background is included in the 2-D Gaussian fit to allow for background events in the active region. Both 1-D and 2-D fits are done using maximum-likelihood histogram fits, described in Appendix B. The 2-D fit extends to the estimated half-max of the Gaussian distribution, which in 2-D includes half of the events in the peak.

This data lacks an independent calibration of g_2 , and given the large systematic uncertainty in the drift field, cannot be anchored to the previous data set except by matching the charge intercept of the linear fit in Fig. 4.4. We therefore gain no information on the W -value or field dependence of the charge and light yield. This data does confirm that the exact anti-correlation between the charge and light signals holds all the way to zero field, and it gives a measurement of the near-zero-field

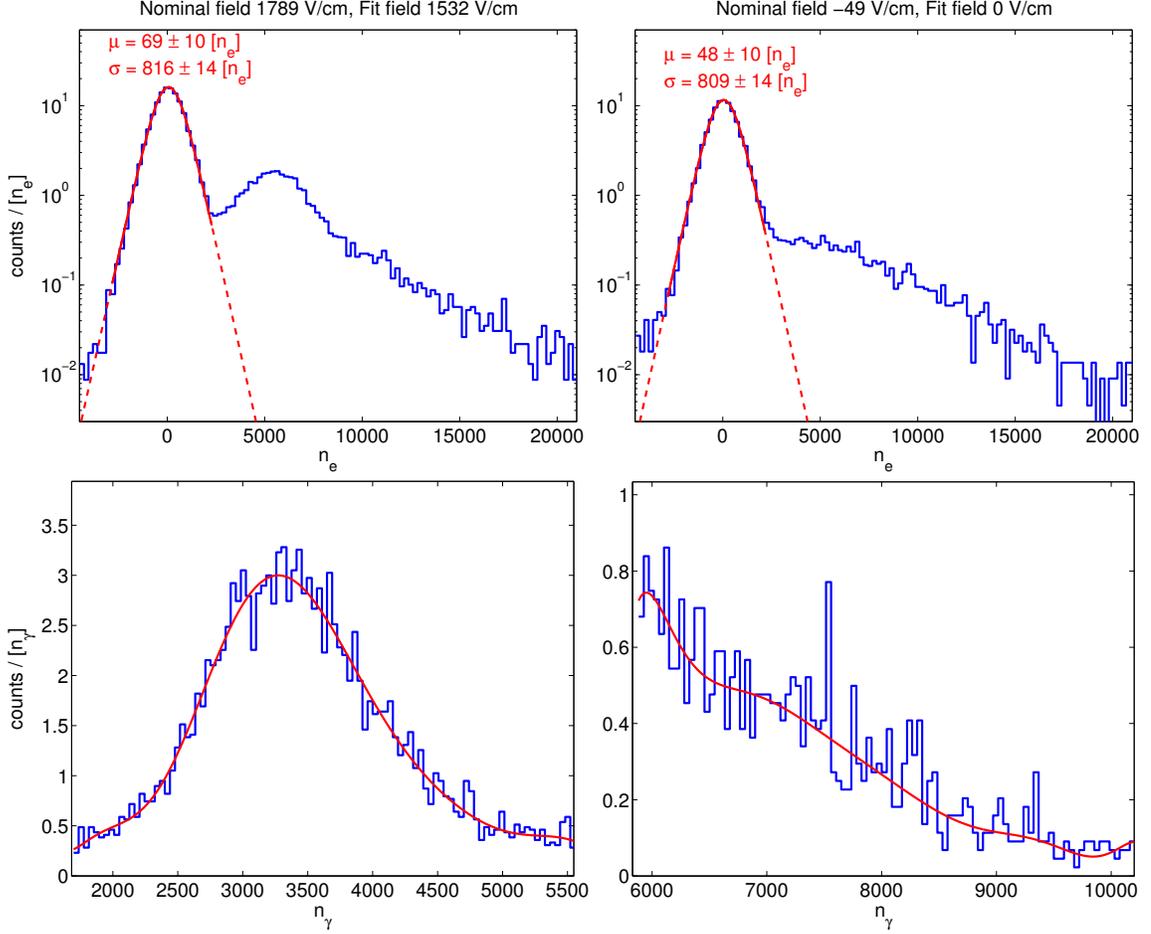


Figure 4.8: Charge (above) and light (below) background profiles for single-phase dual-channel data. The left two plots correspond to the high-field data in Fig. 4.7, and the right two to the low-field data in the same figure. The upper plots show the charge noise peak (blue histogram) and the fit logistic distribution (red curve). The solid part of the curve indicates the region used in the fit. The σ 's reported for both curves are the square root of the variance of the distribution. The lower plots show a polynomial fit to the $S1$ spectrum associated with the zero-charge background. The product of these two backgrounds gives the fixed 2-D background of zero-charge events used for the 2-D Gaussian fit in Fig. 4.7.

charge yield.

4.2.4 Dual-phase, Simultaneous Charge and Light

Dual-phase data also benefits from simultaneous charge/light measurements, as well as higher resolution and lower noise for the charge signal (the noise in the proportional light signal is < 1 extracted electron, compared to ~ 580 and ~ 810 electrons in our single- and dual-channel single-phase setups — resolution is discussed in Section 4.3). There is potential in dual phase for systematic variation in g_2 with drift field since the production of proportional scintillation is sensitive to the electric field

in the gas [51], but if present this would result in a non-linear light-vs-charge plot. Dual-phase data in Xed can reach drift fields up to 4.4 kV/cm since electrons are collected on the ‘anode’ rather than the ‘top’, reducing high voltage requirements.

In dual-phase, it is possible to resolve the 122 and 136 keV ^{57}Co peaks, as seen in Fig. 4.9. To fit the peaks, we first estimate g_1 and g_2 by fitting Gaussians to the S1 and S2 distributions. It is not possible to distinguish the two peaks in either signal alone, so we take a mean energy of 124.25 keV and fix g_1 by matching the light-yield versus field at high fields (>2 kV/cm) to that seen in the single-phase light-only measurements. We then find g_2 at each field by requiring $W(g_1S1 + g_2S2) = 124.25$ keV. From these initial estimates, we construct the recombination independent energy scale, in which the two peaks are distinguishable. Making a cut around the 122 keV peak in the combined energy, we repeat the process (now with a peak of known energy). These estimates are then used to create the axes in Fig. 4.9.

With these calibrated axes, we fit the 122 and 136 keV peaks to two 2-D Gaussians in the recombination independent energy E and $\log(\frac{n_e}{n_\gamma})$. We require the two peaks to have the same covariance matrix, and also fix the ratio of the mean energies of the distributions to 136/122. This leaves eight free parameters in the fit: amplitude, mean, and covariance matrix for the 122 keV peak plus the amplitude and mean $\log(\frac{n_e}{n_\gamma})$ of the 136 keV peak. The background of photo-absorption and Compton scatter events is difficult to determine (especially complicated because of the higher charge yield for Compton scatter events) and is not included in the fit. To avoid background, the fit covers only the center 40% (1σ) of the 122 keV peak and 15% of the 136 keV peak, chosen so that the signal is well above background. The dual-phase points in Fig. 4.4 are found by taking the means of the two peaks in raw S1 and S2 and weighting them by the 87%:13% 122:136 keV ratio seen in the single-phase separate charge-light data. Based on the above calibration, or equivalently, on the scatter about the fit line in Fig. 4.4, we calculate a $\sim 1\%$ rms variation in g_2 that is uncorrelated with the drift field.

The dual-phase data has much smaller error bars and a wider field range than either single phase dataset, and we can take advantage of this to reduce the statistical error on the W-value measurement. The W-value is given by the x-intercept of the linear fits in Fig. 4.4 and by g_2 . We only have an absolute calibration of g_2 in the single-phase charge-only data, but we can use that to fix the mean g_2 in the dual-phase data by requiring that the charge yield versus field agree at high fields (above 2 kV/cm, where the field dependence of the charge yield and uncertainty in the applied field are small). We can then use the much tighter linear fit from the dual-phase data to determine W . This gives a value 0.5% higher than the previous W-value, and significantly reduces

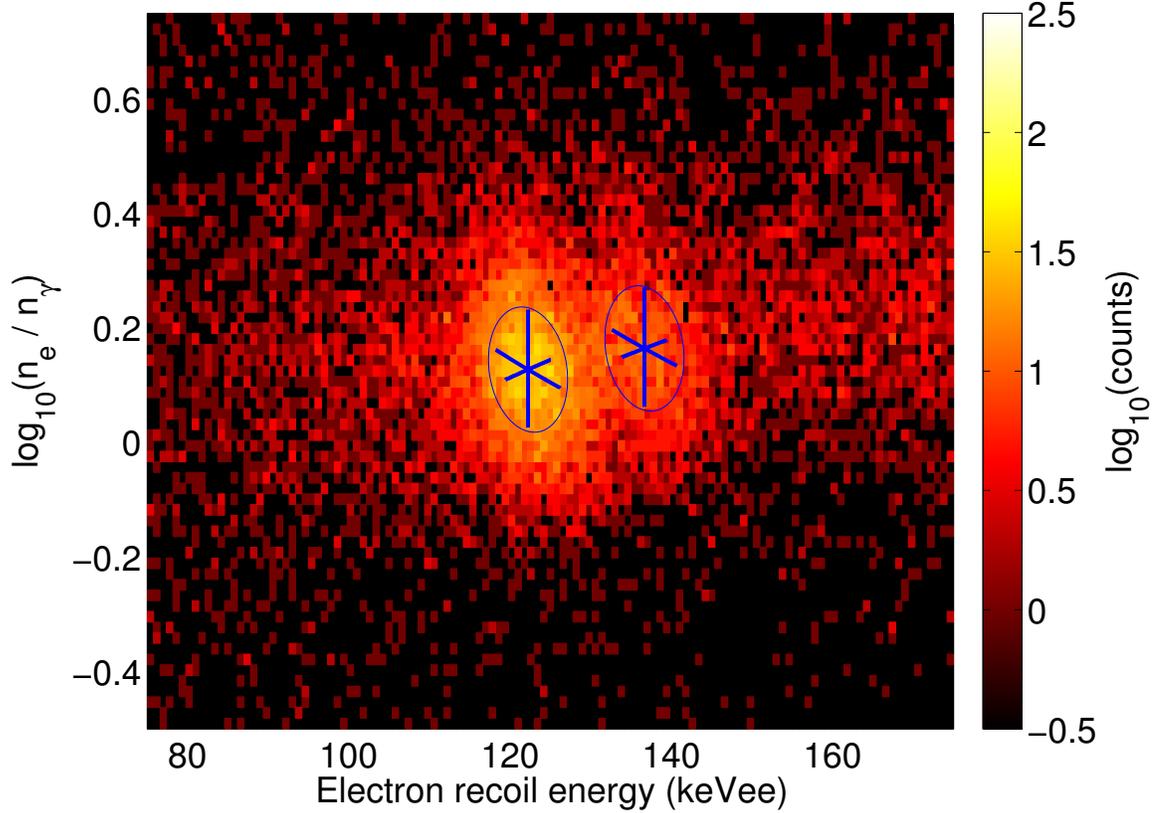


Figure 4.9: Dual-phase ^{57}Co data at 1.23 kV/cm, showing the 122 and 136 keV peaks. The ellipses around the peaks are the 2-D Gaussian fit to the peaks, where the ellipses are defined by $(\mathbf{x} - \boldsymbol{\mu})^\top \cdot \boldsymbol{\Sigma}^{-1} \cdot (\mathbf{x} - \boldsymbol{\mu}) = 1$. The lines inside the ellipse show the decomposition into S1 (upper-left to lower-right), S2 (lower-left to upper-right), and recombination (vertical) fluctuations, with each line extending to $\pm 1\sigma$ for its respective fluctuation. Histogram bin sizes are 0.8 keVee by 0.0187 $\left[\log_{10} \left(\frac{n_e}{n_\gamma} \right) \right]$.

the statistical uncertainty. Our final W-value is $W = 13.7 \pm 0.2$ eV, where the uncertainty is now dominated by the uncertainty in the preamp calibration.

4.2.5 Recombination versus Field

To calibrate other detectors, it is necessary only to have a peak with known energy and y . To this end we find y versus drift field for the ^{57}Co source, so that to calibrate another detector we need only know the drift field and see the ^{57}Co peaks. We use the single phase light-yield-only data for this, inferring y from S1 using the linear fit in Fig. 4.4. (As we did for the W-value, we use the dual-phase data for the linear fit, anchoring this to the single-phase light-only data by matching light yields at fields above 2 kV/cm.) This allows us to measure y at low fields both by removing the uncertainty in drift field associated with the bleed-through of the high fields above the gate and by

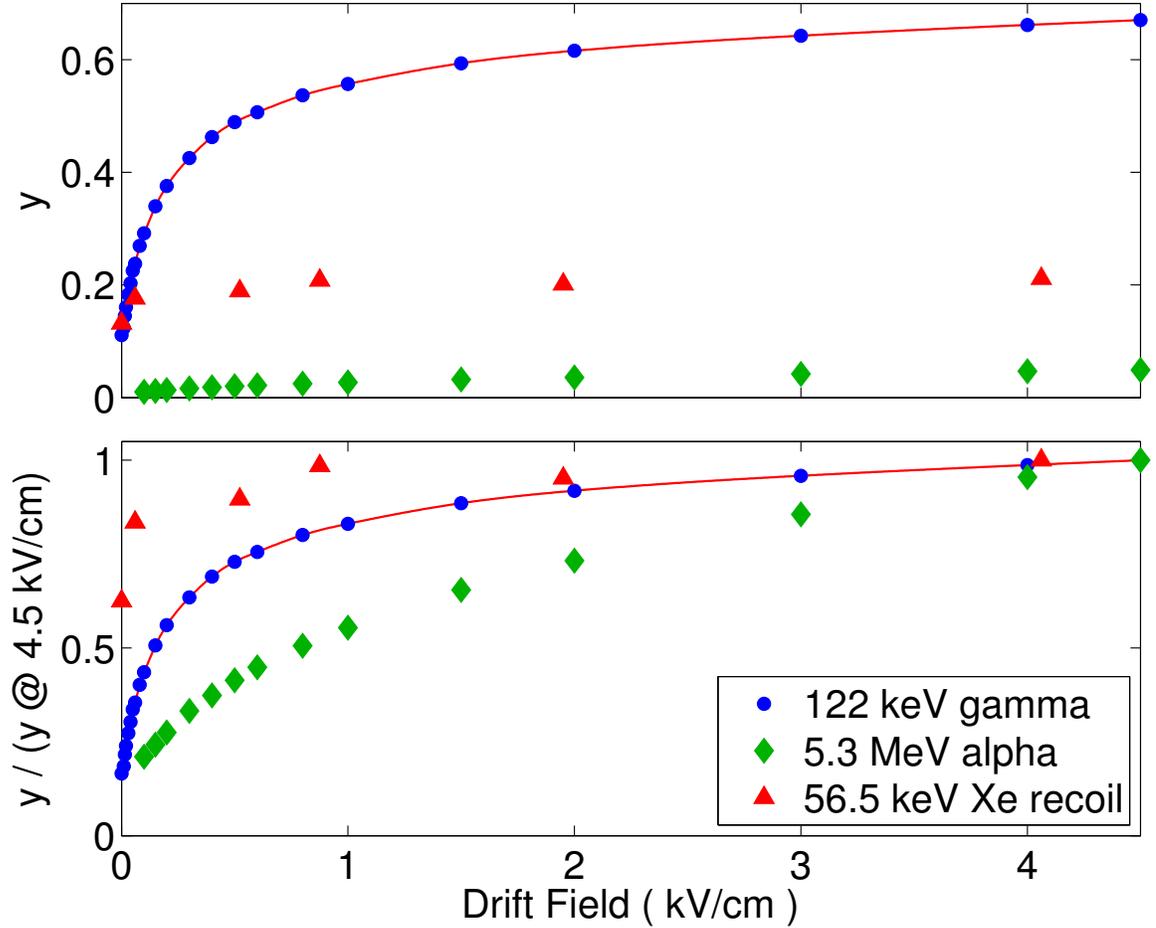


Figure 4.10: Above: Normalized charge yield for ^{57}Co gammas (87.0% 122.1 keV, 13.0% 136.5 keV). ^{210}Po alphas (5.3 MeV) and 56.5 keV nuclear recoils are also shown. The line through the ^{57}Co points is a spline interpolation used for calibration at other fields. Below: The same, but scaled relative to charge yield at 4.5 kV/cm for each recoil, showing the contrasting field dependence of the three recoil types. Errorbars for gamma and alpha data are too small to plot; errorbars for nuclear recoils are the size of plot points in the lower figure.

using the channel (S1) with higher signal to noise. It also allows us to measure the zero field charge yield, i.e., the charge that fails to recombine in the absence of an applied field. The resulting yields are shown in Fig. 4.10, along with yields for the ^{120}Po alpha source and 56.5 keV nuclear recoils. The nuclear recoil yields include data from the nuclear recoil bands in Chapter 5 taking the \mathcal{L}_{eff} given by Sorensen [79], and from Aprile [77] for the zero-field point.

Figure 4.10 demonstrates the two basic trends expected for recombination in liquid xenon. First, recombination decreases as the applied field is strengthened, and second, tracks with higher ionization density experience greater recombination. Along any piece of a track, the ionization density is determined by the electronic stopping power, or the amount of energy lost to atomic electrons

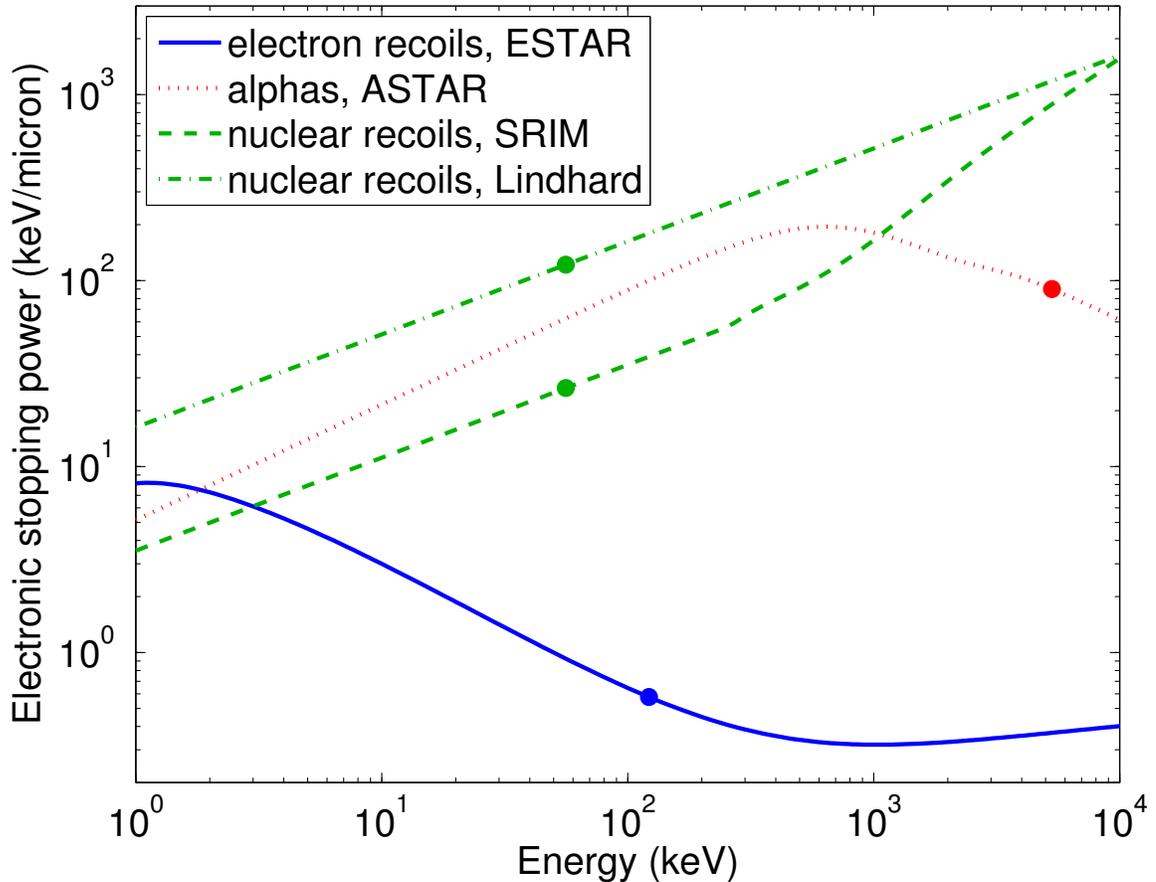


Figure 4.11: Electronic stopping power for electrons [86], alphas [87], and nuclear recoils [88, 37] in liquid xenon. The dots correspond to the starting energies of the recoils shown in Fig. 4.10. Higher electronic stopping power corresponds to higher ionization density, and, in most situations, greater recombination. (Exceptions to this trend are discussed in Chapters 5 and 6.)

per distance traveled. Figure 4.11 shows the stopping power versus energy for electrons, alpha particles, and xenon nuclei in liquid xenon, with the dots on each curve indicating the recoil energies shown in Fig. 4.10. When taking the SRIM data for nuclear recoils, the highest stopping power corresponds to the lowest charge yield (alphas), and the lowest stopping power corresponds to the highest charge yield (electrons). It is worth noting, however, that while the charge yield for alphas and electrons shows a strong dependence on drift field, nuclear recoils, which have an absolute charge yield between that of alphas and electrons, show almost no field dependence. This is the first indication that recombination in small tracks is qualitatively different than in large tracks, a theme that will reappear in Chapters 5 and 6.

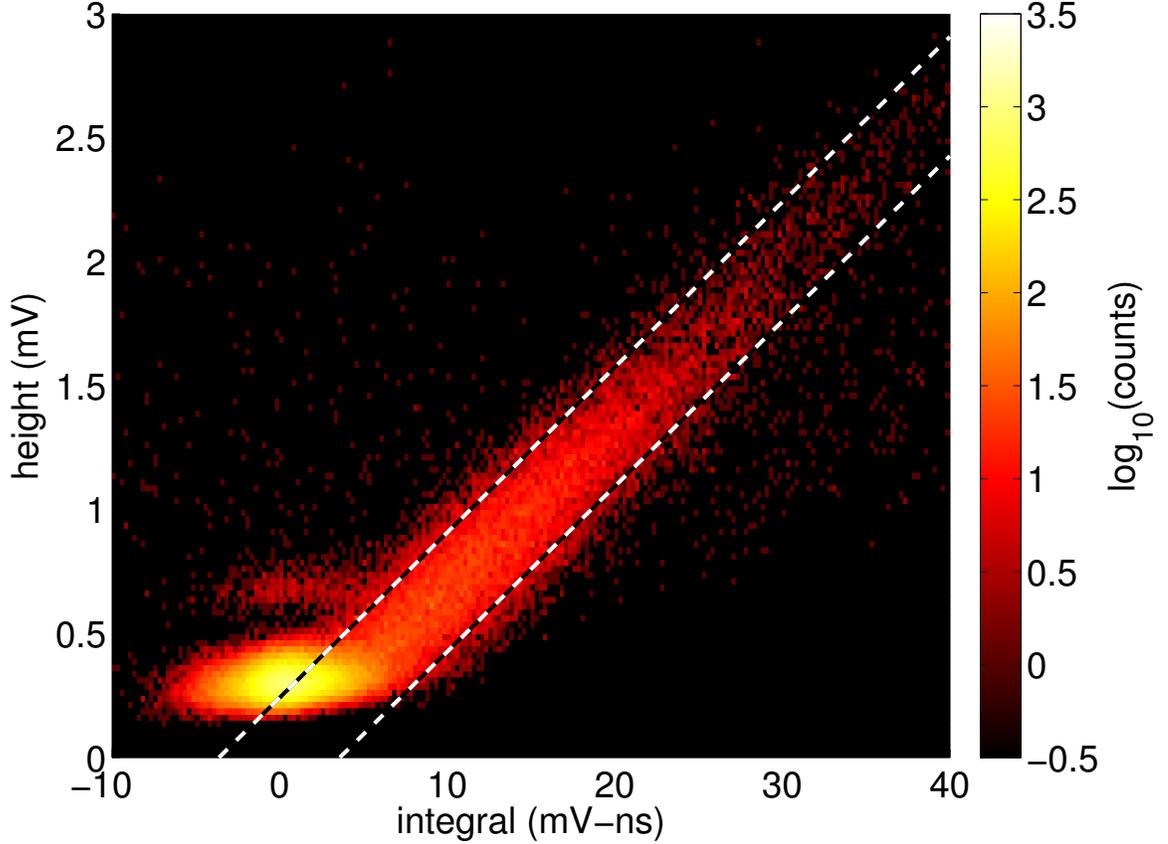


Figure 4.12: Single photo-electrons and noise triggers in the bottom PMT of Xed, during dual-phase operation. PMT bias is 875 V. The dashed lines show the cut around the single photo-electron peak used for the fit in Fig. 4.13.

4.2.6 Absolute Light Calibration

Once W is known, we can use the intercepts of the linear fit to determine g_1 for dual-phase data. Light collection efficiency for S1 is commonly quoted in photo-electrons (phe) per keV at 122 keV at zero field (3.9 phe per keV in dual-phase in Xed), but in finding g_1 we calculate the efficiency in absolute units, phe per $[n_\gamma]$. The single photo-electron signal is typically found using a pulsed LED, tuned so that $\sim 90\%$ of the pulses give no signal. (By Poisson statistics, this ensures that 95% of the pulses that do give signal give only single photo-electrons.) There is no LED in Xed, so to see single photo-electron signals we reduce the trigger threshold until most triggers are due to electronics noise. We do this while at liquid xenon temperatures, where thermal emission of electrons from the photo-cathode and dynodes should be rare, so the light pulses we see are most likely S1s from background radiation interacting with xenon outside the active region.

A 2-D plot of pulse height versus integral shows a clear band for real PMT signals (see Fig. 4.12).

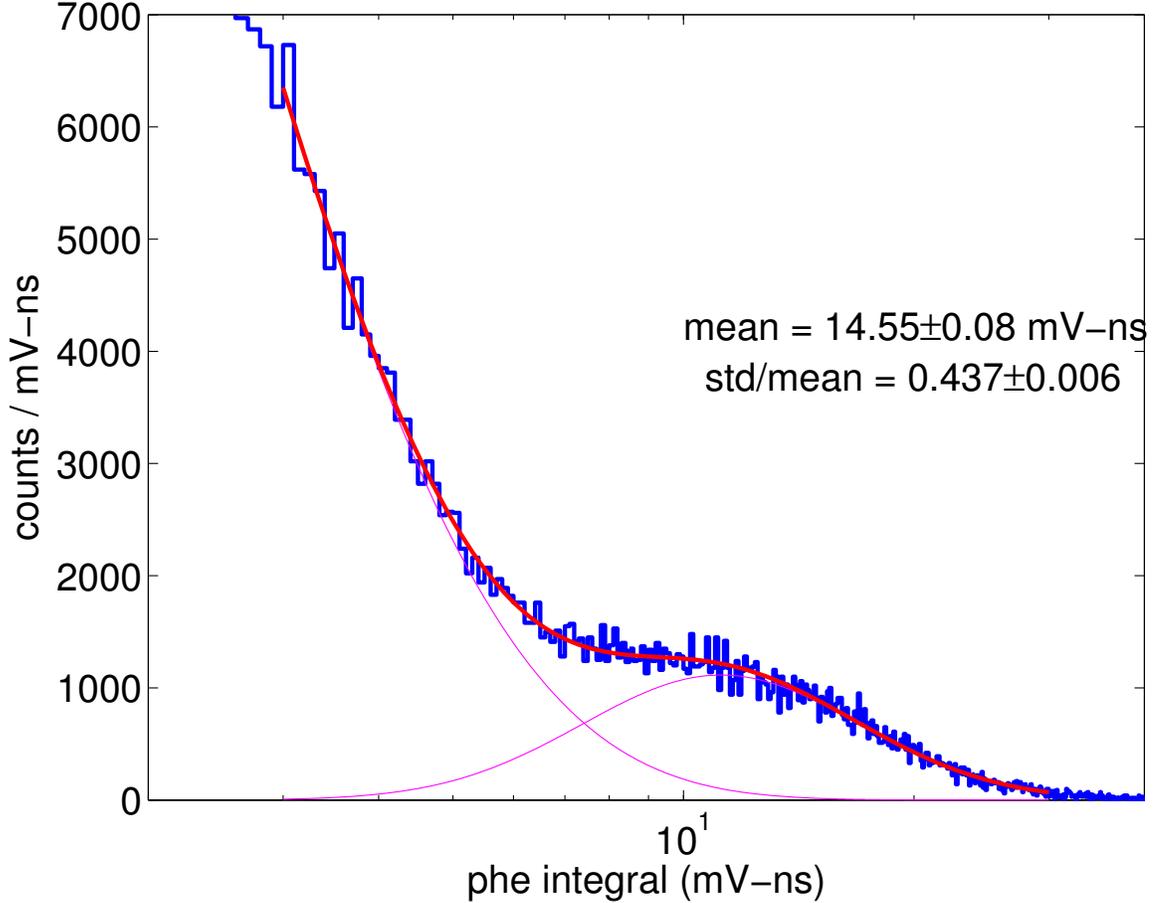


Figure 4.13: The single photo-electron spectrum in the bottom PMT of Xed, during dual-phase operation. PMT bias is 875 V, with the signal going into a 25Ω impedance (50Ω terminations at the PMT and amplifier), followed by a $\times 5$ voltage amplifier. The spectrum is fit to a log-normal distribution over an exponential tail of noise, and the mean corresponds to a PMT gain of 7.27×10^5 . The spectral shape is not well understood — the noise peak in the charge signal in single phase also has exponential tails. A log-normal photo-electron spectrum would be expected for uncorrelated fluctuations in the average gain at each dynode, but these are usually sub-dominant to the Poisson statistics for electrons coming off the first dynode. The mean and variance of the fit log-normal distribution match the statistical fluctuations in S1 measured in Section 4.3, and the mean is also consistent with extrapolations from photo-electron spectra at higher bias, where the peak is clearly separated from noise.

Cutting around the band and plotting the integral spectrum, we see an exponential tail of noise triggers plus a log-normal distribution for the single photo-electron peak (Fig. 4.13). The log-normal distribution is not expected, and the tail at high signal may be due to multiple photo-electron pulses. Taking the mean and variance of the log-normal fit, however, gives good agreement with the statistical S1 fluctuations seen in section 4.3, and the mean is also consistent with extrapolations from higher bias, where the single photo-electron peak is well separated from noise. From the fit mean, we find $g_1 = 0.060 \pm 0.001$ phe per n_γ . Dividing out the reported 23% quantum efficiency

of the bottom PMT and the estimated 70% collection efficiency for the first dynode, we see 37% of recombined ions result in a scintillation photon striking the bottom PMT photo-cathode. Given the large uncertainty in the reflectivity of the PTFE walls of the chamber, this is consistent with $b = 1$.

Finally, we note that if bi-excitonic quenching is present and affects excitons created via recombination, then one would see $b < 1$. This would be evident in plot like Fig. 4.4 as a lower slope for the quenched recoil. This may be a method to test the bi-excitonic quenching hypothesis in nuclear recoils, but it faces two problems. First, as seen in Fig. 4.10, recombination in nuclear recoils is nearly field independent, so it will be difficult to get the spread in y necessary for a good linear fit. Second, the excitons that you gain and lose when shifting field are likely from the edges of the track, while bi-excitonic quenching is expected at the core. Thus, while a slope change between electron and nuclear recoils would be a strong indication of bi-excitonic quenching, it is difficult to rule out bi-excitonic quenching by the absence of the effect.

4.3 Measuring Recombination Fluctuations

As has been noted by several authors [65, 89, 66], recombination fluctuations are the dominant factor in resolution in single channel liquid xenon detectors. These fluctuations produce the same charge-to-light anti-correlation as drift field variations, and the calibrated sum of the light and charge signals in Eq. (4.9) is insensitive to recombination fluctuations. Because of this, the sum energy scale gives much higher energy resolution than can be obtained in a single-channel detector.

Under the assumption that all other fluctuations are uncorrelated in the charge and light signals, it is possible to extract the strength of recombination fluctuations at a peak from this improvement in resolution, or equivalently, from the 2-D charge/light distribution. This assumption holds except for fluctuations in initial ion production, which are suppressed by a theoretical Fano factor of 0.05 [71] and thus subdominant, and for position-dependent effects, which are easily corrected in a 3-D position reconstructing detector. Xed has no x-y position reconstruction, but the ^{57}Co gammas are collimated so that they populate only the center axis of the detector.

The simplest analysis of fluctuations is to find the Gaussian widths of the ^{57}Co 122 keV peak in the measured values of n_e , n_γ , and E , which are given by

$$\sigma_{n_\gamma}^2 = \sigma_R^2 + \sigma_{S1}^2, \quad (4.30)$$

$$\sigma_{n_e}^2 = \sigma_R^2 + \sigma_{S2}^2, \text{ and} \quad (4.31)$$

$$\sigma_E^2 = W^2 (\sigma_{S1}^2 + \sigma_{S2}^2), \quad (4.32)$$

where σ_{S1}^2 , σ_{S2}^2 , and σ_R^2 are the variances in units of quanta measured ($[n_\gamma]^2$ and $[n_e]^2$) due to uncorrelated S1 and S2 fluctuations and recombination fluctuations, respectively. These relations may be found by the same technique used below for Σ , but they should be fairly obvious — n_γ is sensitive to recombination and S1 fluctuations, which add in quadrature since they are independent, and similarly for n_e . The energy E is unaffected by recombination, but responds to both S1 and S2 fluctuations. It follows that

$$\sigma_R^2 = \frac{1}{2} \left(\sigma_{n_\gamma}^2 + \sigma_{n_e}^2 - \frac{\sigma_E^2}{W^2} \right). \quad (4.33)$$

This method is complicated by the nearby 136 keV peak, so we shift to 2-D, where the two peaks are distinct. We consider a two-dimensional distribution of points resulting from three different types of fluctuations, each with a known direction but unknown magnitude: S1 fluctuations move points along lines of constant S2 and vice versa, while recombination fluctuations move points along a line of constant energy (see Fig. 4.9). A 2-D Gaussian has a width described by three parameters, which we can relate to the magnitudes of these three fluctuations. The general multi-variable Gaussian distribution is given by

$$P(\mathbf{x}) = \sqrt{\frac{1}{|2\pi\Sigma|}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\dagger \cdot \Sigma^{-1} \cdot (\mathbf{x}-\boldsymbol{\mu})}, \quad (4.34)$$

where $\boldsymbol{\mu}$ is the mean of the distribution and Σ is the covariance matrix,

$$\Sigma_{ij} = \langle (x_i - \mu_i)(x_j - \mu_j) \rangle. \quad (4.35)$$

We approximate the peak as a 2-D Gaussian in the recombination-independent energy E and $\log(\frac{n_e}{n_\gamma})$, taking the fits from Section 4.2.4. The latter variable is the discrimination parameter used by Xenon10 (up to a constant), and has been seen to follow a Gaussian distribution for electron recoils [31] (see Chapter 7). Considering small fluctuations δ_{S1} , δ_{S2} , and δ_R , the corresponding shifts in E and $\log(\frac{n_e}{n_\gamma})$ are

$$\delta_E = W(\delta_{S1} + \delta_{S2}), \quad (4.36)$$

and

$$\delta_{\ln(\frac{n_e}{n_\gamma})} = \frac{\delta_{S2} - \delta_R}{n_e} - \frac{\delta_{S1} + \delta_R}{n_\gamma}, \quad (4.37)$$

where we have used $\delta_{n_\gamma} = \delta_{S1} + \delta_R$ and $\delta_{n_e} = \delta_{S2} - \delta_R$. The covariance matrix for the distribution is then given by

$$\Sigma_{E,E} = \langle \delta_E^2 \rangle = W^2(\sigma_{S2}^2 + \sigma_{S1}^2), \quad (4.38)$$

$$\Sigma_{E,\ln(\frac{n_e}{n_\gamma})} = \left\langle \delta_E \delta_{\ln(\frac{n_e}{n_\gamma})} \right\rangle = W \left(\frac{\sigma_{S2}^2}{n_e} - \frac{\sigma_{S1}^2}{n_\gamma} \right), \quad \text{and} \quad (4.39)$$

$$\Sigma_{\ln(\frac{n_e}{n_\gamma}),\ln(\frac{n_e}{n_\gamma})} = \left\langle \delta_{\ln(\frac{n_e}{n_\gamma})}^2 \right\rangle = \frac{\sigma_{S2}^2}{n_e^2} + \frac{\sigma_{S1}^2}{n_\gamma^2} + \left(\frac{1}{n_e} + \frac{1}{n_\gamma} \right)^2 \sigma_R^2, \quad (4.40)$$

where we have used $\langle \delta_X^2 \rangle = \sigma_X^2$ and $\langle \delta_X \delta_Y \rangle = 0$ for $X, Y = R, S1, S2$ ($X \neq Y$). Figure 4.9 shows a sample fit of the 122 and 136 keV peaks and their decomposition into constituent fluctuations.

Figure 4.14 shows the breakdown of fluctuations for the dual phase ^{57}Co data, scanning over drift field as in Fig. 4.4. The recombination fluctuations are peaked around 50% recombination ($y \approx 0.47$) as would be expected for a binomial process, but the variance is 100–150 times larger than from binomial statistics. The peak being near 50% indicates that these fluctuations are due not to variations in field (which we expect would be peaked at low y , where field dependence is strong), but rather to random fluctuations in the recombination process. These fluctuations are associated with variations in track structure, as explored in Chapter 6. The uncorrelated S1 and S2 fluctuations may be divided into statistical and ‘instrumental’ fluctuations based on their scaling with signal size. Statistical fluctuations are proportional to the square root of the signal size and include binomial fluctuations for S1 light collection ($\sigma^2 = \frac{(1-g)gn}{g^2}$), Poisson fluctuations for S2 light production ($\sigma^2 = \frac{gn}{g^2}$), and the variation in photo-electron response for both S1 and S2 ($\sigma^2 = \frac{gn\sigma_{\text{phe}}^2}{g^2}$), where for each of these n is the number of quanta creating the signal, gn is average number of photo-electrons created, and σ^2 is measured variance in units of $[n]^2$. Combining these gives

$$\sigma_{S1,\text{stat}}^2 = \frac{1 - g_1 + \sigma_{\text{phe}}^2}{g_1} n_\gamma, \quad \text{and} \quad (4.41)$$

$$\sigma_{S2,\text{stat}}^2 = \frac{1 + \sigma_{\text{phe}}^2}{g_2} n_e, \quad (4.42)$$

where g_1 and g_2 are in phe per quanta and σ_{phe}^2 is the variance of the single photo-electron distribution in phe². As noted in Section 4.2.6, we find $g_1 = 0.060$ phe per $[n_\gamma]$, and the single photo-electron distribution is log-normal with a variance $\sigma_{\text{phe}}^2 = 0.19$. Typical values for g_2 are 13–15 phe per $[n_e]$. The remaining ‘instrumental’ fluctuations are event by event variations in g_1 and g_2 , and thus proportional to signal size.

Figure 4.14 also shows fits for instrumental plus statistical fluctuations in S1 and instrumental fluctuations in S2 (statistical fluctuations in S2 are negligible). The fit statistical piece for S1 fluctuations is $4.48 \pm 0.15\sqrt{n_\gamma}$, consistent with the expected $4.34\sqrt{n_\gamma}$ from Eq. (4.41). The fit instrumental fluctuations in S1 and S2 are $0.0315n_\gamma$ and $0.0395n_e$, respectively. The dominant

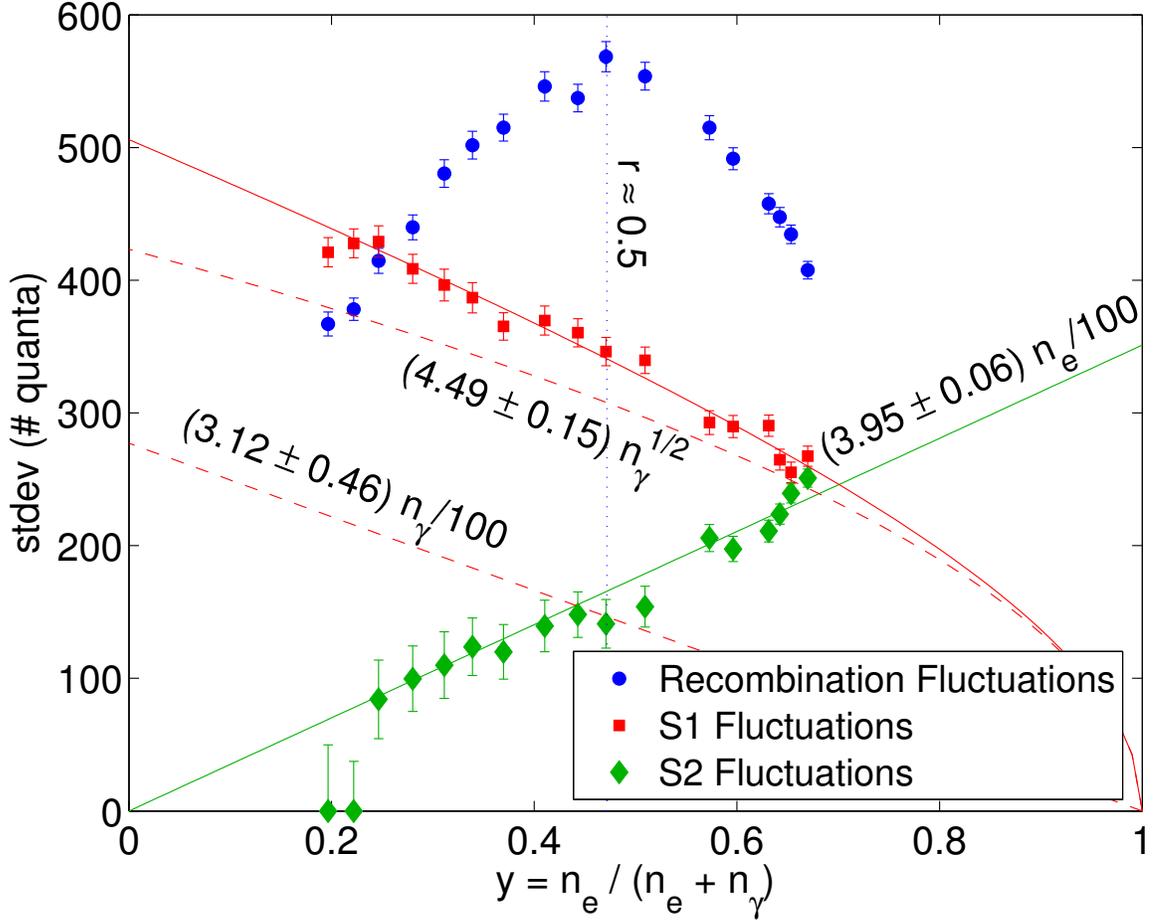


Figure 4.14: Uncorrelated S1 and S2 fluctuations and recombination fluctuations in the ^{57}Co 122 keV gamma peak, scanning over drift field. The y-axis is in # of quanta (n_γ for S1, n_e for S2, and $n_e \leftrightarrow n_\gamma$ for recombination). The total 122 keV signal is 8891 quanta ($n_\gamma + n_e$). Also shown are fits to pure instrumental fluctuations for S2 and instrumental plus statistical fluctuations for S1. The dashed lines show the individual instrumental and statistical contributions to the S1 fit. The two lowest field points are not included in the fits. The dotted vertical line at $r \approx 0.5$ ($y = 0.47$) assumes $a = b = 1$ and takes the theoretical value $\frac{N_{ep}}{N_i} = 0.06$.

instrumental fluctuations in S2 come from ripples on the liquid surface, which can affect both the amount of S2 light produced per electron and the efficiency with which electrons are extracted into the gas phase. Improved S2 resolution has been achieved by slowly warming the detector from the top, which raises the gas pressure above the vapor pressure of liquid and halts bubble formation, visibly calming the liquid surface. Instrumental fluctuations in S1 may also be related to the liquid surface, or may be residual position dependence in light collection efficiency. Instrumental fluctuations in both channels increase significantly for a non-collimated source, due to uncorrected x-y position dependence.

4.4 Conclusions on Recombination and Calibration

We have demonstrated that the 122 keV ^{57}Co peak follows an exact anti-correlation between the charge and light signals, as expected from the recombination model for signal production. This anti-correlation has been verified in both single- and dual-phase detectors, from zero drift field to 4.4 kV/cm. The resulting linear fit, along with the absolute calibration of the charge signal in single-phase, gives the W-value $W = 13.7 \pm 0.2$ eV, where the uncertainty is due to the uncertainty in the preamp calibration. The anti-correlation also lets us translate the absolute charge calibration to an absolute light calibration, and we find that $6.01 \pm .10\%$ of recombined ions result in a S1 photo-electron in the bottom PMT of Xed. Dividing out the PMT efficiency and assuming 100% efficiency for scintillation production from recombining ions, this corresponds to a 37% light collection efficiency.

The same anti-correlation seen when varying the drift field is present in recombination fluctuations. We have decomposed the measured fluctuations in the ^{57}Co peak into recombination fluctuations and uncorrelated charge and light fluctuations based on the known orientation of each fluctuation type. The resulting recombination fluctuations are peaked at 50% recombination, as for a binomial distribution, but the variance is 100–150 times that from binomial statistics. We have divided the uncorrelated charge and light fluctuations into instrumental and statistical fluctuations, where the former scale with signal size and the latter with the square root of signal size. We find that fluctuations in the light signal have an instrumental piece of $3.12 \pm 0.46\%$ and a statistical piece consistent with that predicted by the above light yield and the measured photo-electron spectrum. Statistical fluctuations in the charge signal are negligible, also as expected, while instrumental charge fluctuations are found to be $3.95 \pm 0.06\%$.

We now have a reference for absolute calibration of any detector with known drift field, as well as a consistent picture for fluctuations in the detector. We are ready to extend this picture to WIMP recoil energies.

Chapter 5

Recombination and Discrimination

We now extend the recombination picture to the WIMP region of interest. Our goal is to measure the discrimination power between electron and nuclear recoils at WIMP energies based on the different recombination fractions of the two event types, and to understand this discrimination in terms of fundamental fluctuations as we did the resolution of the ^{57}Co peak in the last chapter. Ultimately, we want to determine the recombination distribution as a function of energy and drift field for both recoil types.

With Xed in dual-phase mode, we study low energy Compton scatters from 356 keV ^{133}Ba gammas and elastic scatters from ^{252}Cf neutrons for electron and nuclear recoils, respectively, looking at drift fields of 60, 522, 876, 1951, and 4060 V/cm. After determining the statistical and instrumental fluctuations in the detector, we extract the intrinsic recombination fluctuations from the residual widths of the electron and nuclear recoil bands. We then predict the discrimination power of other real and envisioned xenon TPCs by adding back the appropriate statistical and instrumental fluctuations.

We note that this work began alongside the construction of the Xenon10 experiment — in fact, when the XENON collaboration was formed, there was neither a physical model nor data supporting any discrimination between electron and nuclear recoils at WIMP energies, aside from the known electronic stopping powers of the two recoil types and the observed trend that, as electronic stopping power increases, so does the recombination fraction. These measurements were the first to verify discrimination between electron and nuclear recoils at WIMP energies, and indeed it turns out that recombination at low energies behaves very differently than expected. Although our original motivation to look for low energy discrimination was incorrect (as we will see in Chapter 6), we

demonstrate that 99.9% discrimination can be achieved in these detectors below 12 keVr.

5.1 Calibration of Gamma and Neutron Data

To study the fundamental processes determining discrimination, there are two calibration tasks to perform at each drift field setting. We need to find gains in the light and charge signals (g_1 and g_2 in the notation from the previous chapter), and we need to determine the instrumental fluctuations in the light and charge signals for events uniformly distributed in the active region (which may be significantly larger than those found for the collimated ^{57}Co events). The first calibration we accomplish using the ^{57}Co peaks measured in Chapter 4, and we use the 40 keV inelastic nuclear recoil peak present in the neutron data to determine instrumental fluctuations (see Fig. 5.1).

5.1.1 Determining Calibration Constants

In the last chapter we found the charge and light yields for the 122 and 136 keV ^{57}Co photo-absorption peaks versus drift field, so that this source could serve as a calibration reference in any detector with a known applied field. Because of the relatively short active region in Xed and the large field changes across wire grids, there are large uncertainties on the actual drift fields achieved. We thus have three unknowns at each field setting: the actual drift field, g_1 , and g_2 .

We have several handles on the calibration problem. We know the nominal drift field at each setting, and at high fields, where field dependence is weak and error in uncertainty in the field setting is low, this gives an acceptably good estimate of the charge and light yields. We also know the g_1 and g_2 found for dual-phase data in Chapter 4, and we have an expectation that g_1 , which depends only on PMT gain, PMT efficiency, and S1 light collection efficiency, be stable and field independent. In addition, we have the field independent relation

$$\frac{S1}{g_1} + \frac{S2}{g_2} = \frac{E}{W}, \quad (5.1)$$

with $W = 13.7 \pm 0.2$ eV found in the previous chapter. This applies to the 122 and 136 keV peaks, and also to the 40 keV inelastic scatter (although the exact energy of this peak is uncertain, due to the unknown nuclear recoil part of the interaction).

We find the most consistent picture for both 40 and 122 keV peaks when we keep g_1 fixed across datasets, setting its value by matching the light yield expected at 122 keV at the nominal drift field for the two highest field datasets. This gives $g_1 = 6.02 \pm 0.03 \cdot 10^{-2} \text{phe}/[n_\gamma]$, which agrees with the

g_1 found for dual-phase data in Chapter 4. Fixing g_1 , we use the 122 keV peak to determine g_2 and the drift field from Eq. (5.1) and the data in Chapter 4, with results shown in Table 5.1.

We found that it was also necessary to allow g_2 to vary between datasets at the same applied field. Each dataset (^{57}Co , ^{133}Ba , and ^{252}Cf) contains low energy Compton scatters and high energy electron recoils (a mix of Compton scatters and photo-absorptions), and we shift g_2 for the ^{133}Ba and ^{252}Cf data such that the charge yield versus energy for these events is the same as in the ^{57}Co data (see Fig. 5.2). The low-energy comparison is limited to 14–30 keVee — below 14 keVee the elastic nuclear recoil band overlaps with the Compton band in the ^{252}Cf data, and above 30 keV the Compton band runs into the inelastic nuclear recoil peak. We also compare above the ^{57}Co peaks, from 160–200 keVee. We have much higher stats at these energies, but we also have an unknown mix of Compton scatter and photo-absorption events, which may be different for the different sources. Since photo-absorption events have a lower average charge yield than Compton scatters (discussed below), there is potential for a systematic offset between sources. We find, however, that matching the bands at low energies gives near agreement at high energies as well.

At 522, 876, and 1951 V/cm a $\sim 5\%$ decrease in g_2 is observed between the ^{57}Co dataset and the ^{133}Ba and ^{252}Cf datasets. No significant change in g_2 is seen between the ^{133}Ba and ^{252}Cf datasets at any field except at 60 V/cm, which is the only field where the electron and nuclear recoil datasets were not taken consecutively. At 4090 V/cm the ^{57}Co and ^{252}Cf sources were used simultaneously, so there is no shift in g_2 between any datasets at that field. The shift in g_2 at the three middle fields is probably related to the total event rate — our trigger rate is ~ 100 Hz for the ^{133}Ba and ^{252}Cf sources, and ~ 25 Hz for the ^{57}Co source. The g_2 values for each dataset and the nominal and calibrated drift fields are shown in Table 5.1, along with limits on the shift in $\log_{10}(\frac{g_2}{g_1})$ (after corrections) between the ^{133}Ba and ^{252}Cf datasets based on Compton scatter events.

The calibrated neutron and gamma data at each field is shown in Figs. 5.1, 5.3, and 5.4.

5.1.2 Determining Bulk Instrumental Fluctuations

In Chapter 4 we found the instrumental fluctuations in the 122 keV peak, but these events were concentrated in the bottom-center of the active region. To measure the instrumental fluctuations present in the electron and nuclear recoil bands we need a peak that is uniform in the active volume. Photo-absorption lines like those from ^{57}Co gammas are unsuitable for this. Above 300 keV the Compton scattering cross section is greater than that for photo-absorption, and the photo-peak is washed out by forward scattering of incoming gammas and high energy Compton scattering inside

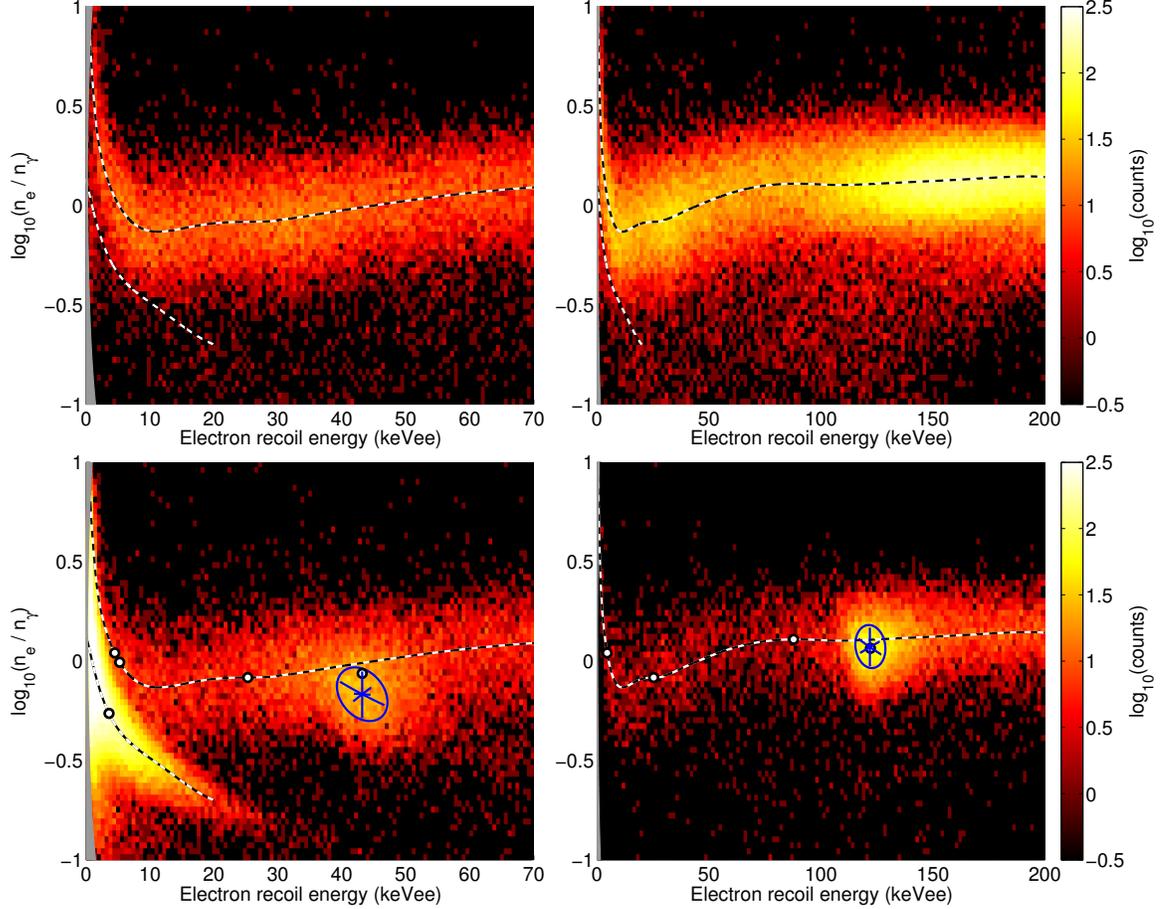


Figure 5.1: Calibrated Xed data at 876 V/cm with all three sources. The top two plots show ^{133}Ba gamma data at low and high energy, the lower-left shows ^{252}Cf elastic and inelastic neutron scatters (plus gammas), and the lower-right shows the ^{122}Co calibration source with the 122 and 136 keV photo-absorption peaks. The electron and nuclear recoil band centroids are also shown. The gray-masked region at the left edge of each plot shows our threshold. At high charge yield this is an S1 threshold corresponding to half a photo-electron, set by the PMT noise peak. At low charge yield this is an S2 threshold of 10 extracted electrons, set by our hardware trigger. An upper S1 threshold can be seen in the high-energy gamma plot, cutting out the lower-right corner of the plot. Bin sizes are 0.6 keVee by $0.03 \left[\log_{10} \left(\frac{n_e}{n_\gamma} \right) \right]$ on the left, with 1.6 keVee bins on the right. Multiply by the x-axis values by 4 for nuclear recoil energies (keVr). Fits to the 40 keV inelastic and 122 keV photo-absorption peaks are shown as ellipses at 1σ , with crosses indicating the decomposition of the fit width into uncorrelated fluctuations as in Chapter 4. The additional horizontal piece of the cross in the 40 keV peak is the smearing due to the variable nuclear recoil. The magnitude of this smearing is unknown — we show it as 1.3 keVee (5.2 keVr), half of the maximum allowed by the observed width and known statistical fluctuations. Both the 40 and 122 keV peaks lie below the electron recoil centroid. The individual recoils making up these peaks are shown as dots on their corresponding centroids, and the dot at the full energy is the sum of the constituent recoils. For the 40 keV inelastic peak, this reconstructed charge yield is significantly higher than the observed yield because many of the constituent recoils arise from the same atom, giving them a lower charge yield than the lone Compton scatters. For the 122 keV peak, the charge yield for the low-energy recoils is similarly overestimated, while the charge yield of the 87 keV photo-electron is underestimated due to photo-absorption events skewing the electron recoil centroid, conspiring to give an accurate reconstruction. Both reconstructions are handled correctly in Chapter 6.

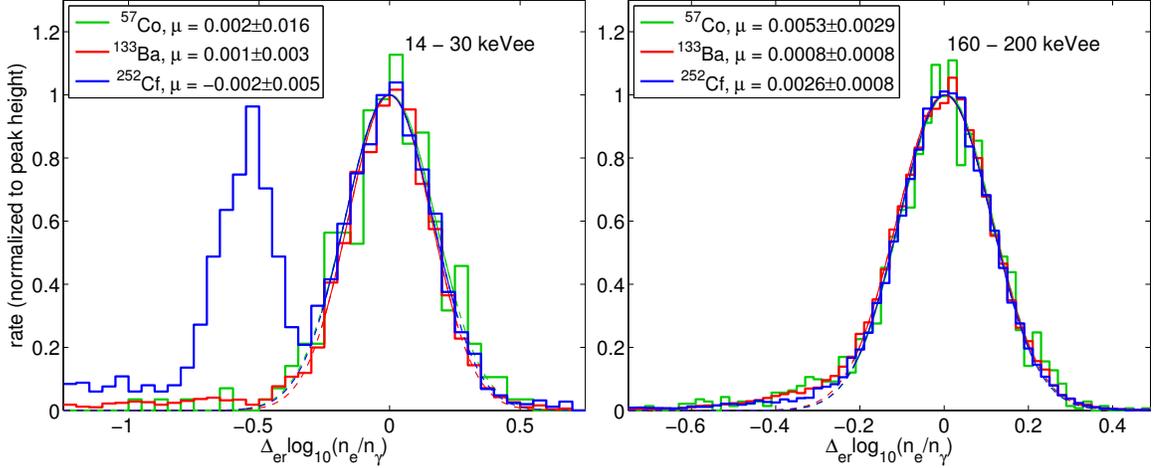


Figure 5.2: Calibrated Xed data at 876 V/cm with each source. The histograms show the charge yield (with the electron recoil centroid subtracted) for each source in energy windows of 14–30 keVee (left) and 160–200 keVee (right), with a Gaussian fit about each peak. We have calibrated g_1 and g_2 for the ^{57}Co dataset using the 122 keV photo-absorption peak, and shifted g_2 for the other sources so that the electron recoil bands align at low energy. Calibration values are given in Table 5.1.

Table 5.1: Calibration values for dual-phase data. Errors in the table are statistical only, based on the ^{57}Co fit and fits to low energy Compton scatters. Errors do not include the uncertainty in the W-value or single photo-electron signal, although these are included later in the systematic error analysis. The final column is the 1σ uncertainty on the shift in $\log_{10}(\frac{g_2}{g_1})$ between the ^{133}Ba and ^{252}Cf datasets, based on low energy Compton scatters and high energy electron recoils (Compton scatters and photo-absorptions).

Drift Field (V/cm)		$g_1(\text{phe}/[n_\gamma])$	$g_2(\text{phe}/[n_e])$			$\delta \log_{10}(\frac{g_2}{g_1})$
Nominal	Calibrated	$^{57}\text{Co}, ^{133}\text{Ba}, ^{252}\text{Cf}$	^{57}Co	^{133}Ba	^{252}Cf	$^{133}\text{Ba} \leftrightarrow ^{252}\text{Cf}$
4000	4060 ± 190		...	13.51(3)	...	0.003
2000	1951 ± 86	\vdots	13.70(4)	13.43(14)		0.003
1000	876 ± 36	$6.02(3) 10^{-2}$	14.57(7)	14.00(15)		0.002
500	522 ± 23	\vdots	14.20(6)	13.50(15)		0.004
100	60 ± 5		17.26(25)	15.13(27)	15.78(26)	0.004

the detector. Lower energy gammas have interaction lengths at or below the dimensions of even our small active region, giving a non-uniform spatial distribution. The neutron interaction length is ~ 15 cm, much larger than our active volume, so we can use the 40 keV inelastic recoil peak (see Fig. 5.4) as our uniform source.

The 40 keV inelastic peak includes 39.6 keV in electron recoils from the excitation of the target nucleus. Typically this energy goes into a 5.3 keV K-shell internal conversion electron, a 29.8 keV

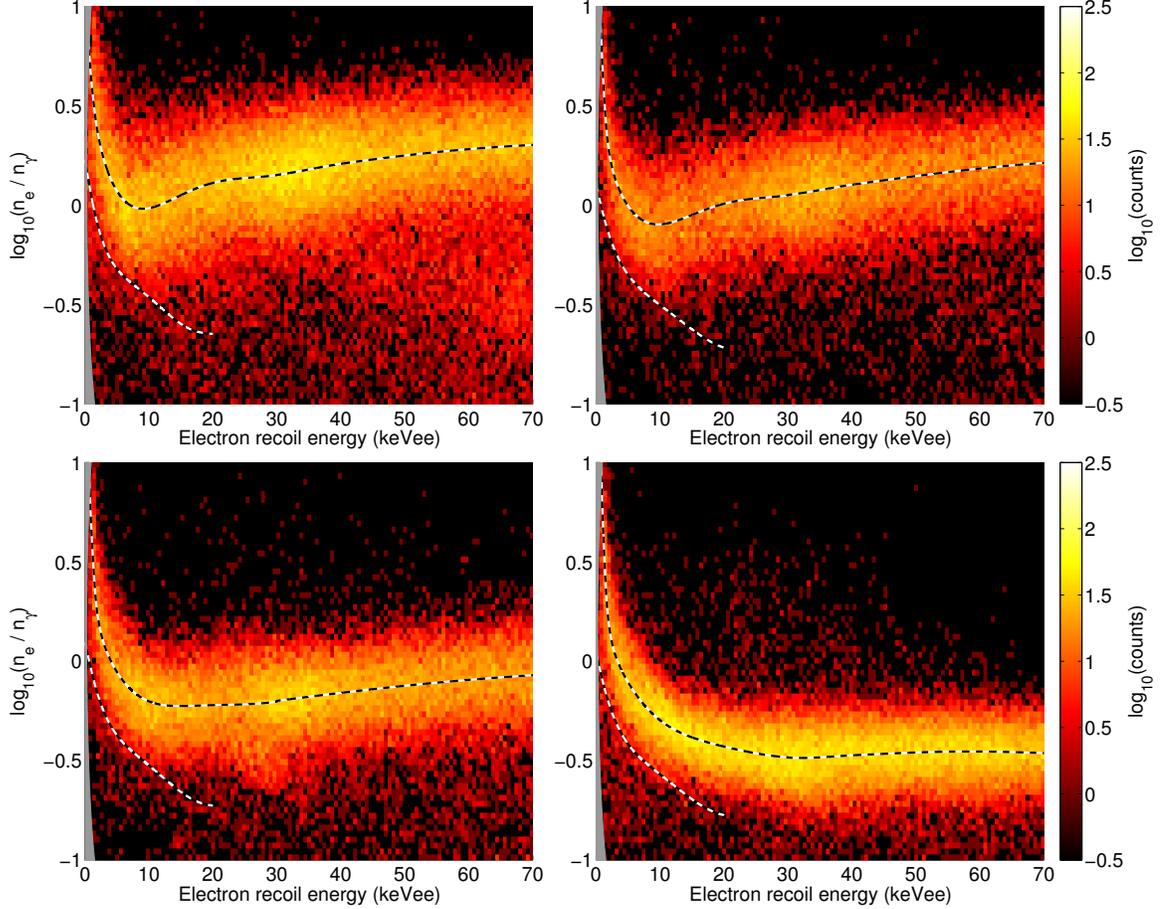


Figure 5.3: Calibrated Xed data at 4060 (top-left), 1951 (top-right), 522 (bottom-left), and 60 V/cm (bottom-right), showing the ^{133}Ba source (Compton scatters). The electron and nuclear centroids are drawn in, the latter taken from Fig. 5.4. The tail of low charge-yield events underneath the electron recoil band is due to events at the edge of the detector, where charge may be lost on the PTFE walls. This tail is removed by a fiducial cut in detectors with x-y position reconstruction.

$K_{\alpha 1}$ x-ray, and a 4.5 keV L_3 -XX Auger electron, as diagrammed in Fig. 5.5. The x-ray travels ~ 0.4 mm before being photo-absorbed by an L-shell electron, giving a 25 keV photo-electron and another 5 keV Auger electron. There is also a variable nuclear recoil component to the peak. The energy spectrum of this recoil is unknown, but the measured elastic recoil spectrum from neutrons is roughly a falling exponential with mean of ~ 6.8 keVr (~ 1.7 keVee) (see Fig. 5.6). This is smaller than the observed width in E of the 40 keV peak, so an exponential shape to the peak may be washed out by other fluctuations. Given this and the unknown shape of the underlying nuclear recoil distribution, we simply fit the peak to a 2-D Gaussian.

We fit the 40 keV peak to a 2-D Gaussian in E and $\log\left(\frac{n_e}{n_\gamma}\right)$ over a fixed 2-D background of Compton scatters and determine the instrumental fluctuations from the fit covariance matrix. The

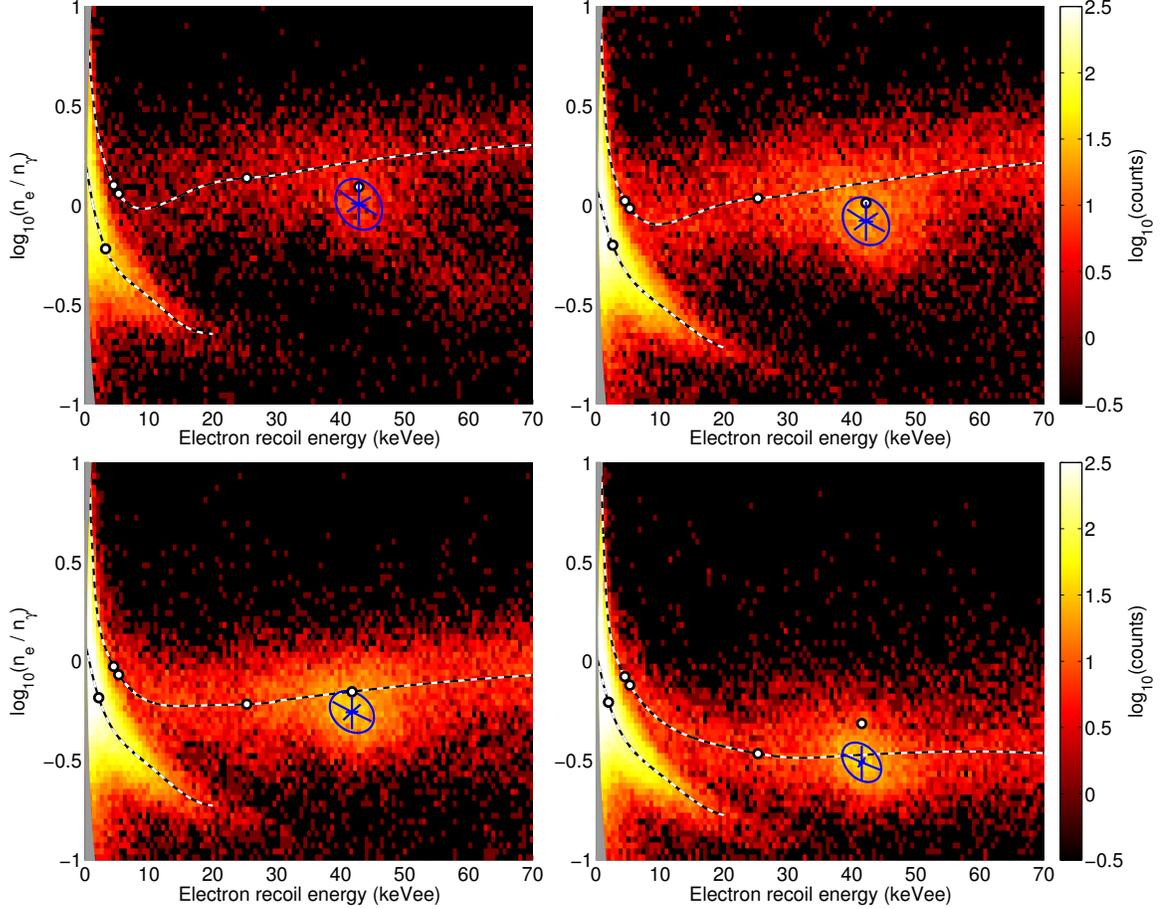


Figure 5.4: Calibrated Xed data at 4060 (top-left), 1951 (top-right), 522 (bottom-left), and 60 V/cm (bottom-right), showing the ^{252}Cf source (Compton scatters and elastic and inelastic neutron scatters). The fits to the 40 keV inelastic peaks are also shown, giving the instrumental fluctuations listed in Table 5.2. The electron and nuclear centroids are drawn in, the former taken from Fig. 5.3. At 4060 V/cm the ^{57}Co source is also present, and we have a lower total ^{252}Cf exposure than at the other fields.

profile of the fixed background in $\log\left(\frac{n_e}{n_\gamma}\right)$ versus E is taken from the gamma data at the same field, and the background energy spectrum is a linear interpolation of the gamma rate above and below the 40 keV peak in the neutron data. The fluctuation decomposition is similar to that for the 122 keV peak from Chapter 4, where the strengths of three linear fluctuations with known orientations in $\left(E, \log\left(\frac{n_e}{n_\gamma}\right)\right)$ were determined from the three elements of the covariance matrix. For the inelastic scattering peak, the variable nuclear recoil adds a fourth fluctuation. The direction of this fluctuation is estimated from the elastic nuclear recoil band, taking

$$\delta \left[\ln \frac{n_e}{n_\gamma} \right] = \left(\frac{\tilde{n}_e}{\bar{n}_e} - \frac{\tilde{n}_\gamma}{\bar{n}_\gamma} \right) \frac{\delta [E]}{E_{nr}}, \quad (5.2)$$

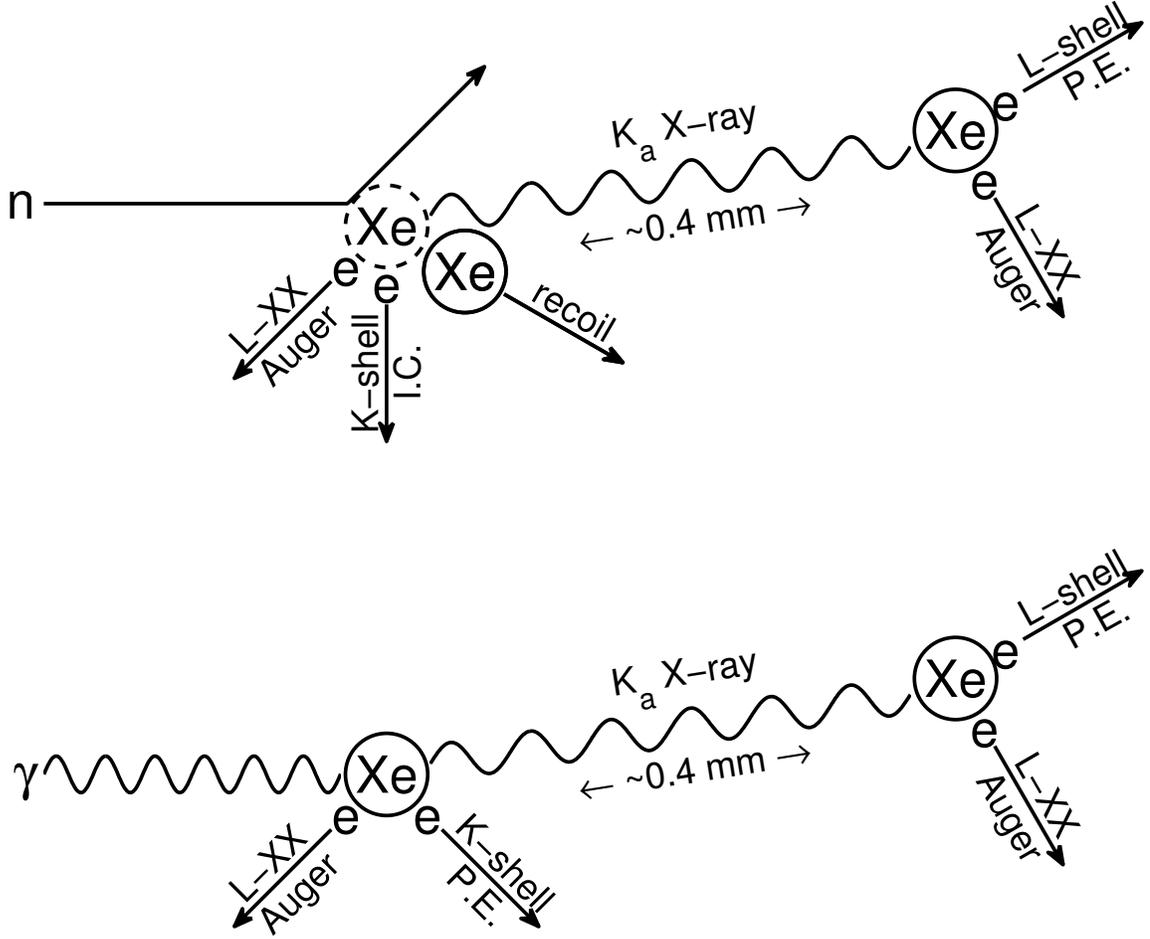


Figure 5.5: Schematic of inelastic neutron scattering (above) and photo-absorption (below) events. Both event types typically knock a K-shell electron from the target atom (a photo-electron in the photo-absorption case, and an internal conversion electron from nuclear de-excitation for the inelastic recoil). This results in a K_a x-ray (29–30 keV) as the hole is filled by an electron from the L-shell, followed by an Auger electron (4–5 keV) as the L-shell vacancy is filled. The x-ray travels a mean distance of 0.4 mm before being photo-absorbed by an L-shell electron, which escapes with ~ 25 keV and results in another Auger electron.

where \bar{n}_γ and \bar{n}_e are the mean of the 40 keV peak, $E_{nr} = W(\bar{n}_\gamma + \bar{n}_e)$ is the mean nuclear recoil energy in keVee (given by the peak energy minus 39.6 keVee), and \tilde{n}_γ and \tilde{n}_e are mean nuclear recoil contribution to \bar{n}_γ and \bar{n}_e , estimated from the elastic nuclear recoil band at E_{nr} .

For a given amplitude of nuclear-recoil fluctuations $\delta[E]$ we can subtract the nuclear-recoil fluctuations from the fit covariance matrix and determine the magnitudes of the uncorrelated S1, S2, and recombination fluctuations as before. Instrumental fluctuations are then found by subtracting the known statistical fluctuations given by Eq. (4.41,4.42) from the total S1 and S2 fluctuations. We vary $\delta[E]$ from zero to the maximum consistent with non-zero instrumental fluctuations for S1 and

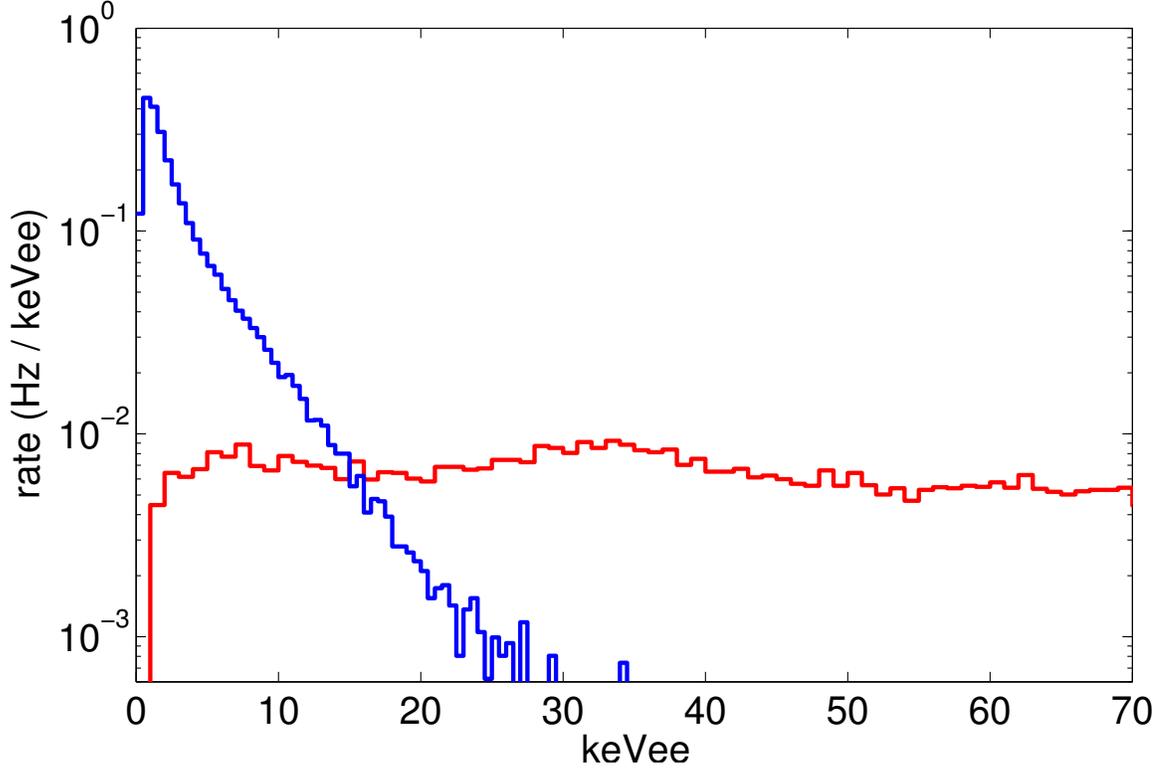


Figure 5.6: Spectra of ^{133}Ba electron recoils (red) and ^{252}Cf nuclear recoils (blue) in the 876 V/cm data, after cuts (drift-time and single-scatter), and selecting events within 0.26 in $\log_{10}\left(\frac{n_e}{n_\gamma}\right)$ of the appropriate centroid (so inelastic scatters and electron recoils are not included in the nuclear recoil spectrum). In simulations, we take the electron recoil spectrum to be constant, and the nuclear recoil spectrum to be a falling exponential.

S2. The resulting ranges for instrumental fluctuations in S1 and S2 are shown in Table 5.2.

Table 5.2 also shows the fit energies and charge yields for the 40 keV peak. The fit energies are consistently higher than the 41.3 keVee expected based on the elastic nuclear recoil spectrum, especially at high field, where the fluctuations in the measured E (dominated by statistical fluctuations in $S1$) are largest. This may indicate a more energetic nuclear recoil piece to the spectrum than was guessed, which would simply shift the mean when energy resolution is poor (higher field), but would skew the fit, which tends to pick out the mode, to lower energy when energy resolution is good (lower field), assuming the spectrum is still exponential. Without nuclear data to support such a spectrum, however, we simply take the instrumental fluctuations as listed in Table 5.2. Because of the way instrumental fluctuations scale with signal size, the large uncertainties become unimportant in the WIMP region.

Table 5.2: Fits to the 40 keV inelastic neutron scatter peak. The errors on E and $\log_{10}(\frac{n_e}{n_\gamma})$ are statistical only. The ranges in S1 and S2 instrumental fluctuations are found by varying the smearing due to the variable nuclear recoil piece of the event. It is unclear why the fit energy varies with field. One possibility is that as statistical S1 fluctuations decrease (due to higher S1 at lower field), the shape of the underlying spectrum (probably exponential) begins to skew the fit.

Field (V/cm)	E (keVee)	$\log_{10}(\frac{n_e}{n_\gamma})$	S1 inst	S2 inst
4060	42.8(3)	0.046(10)	1.7–9.2%	5.8–9.4%
1951	42.2(2)	−0.080(5)	2.4–8.1	6.9–10.3
876	43.2(2)	−0.167(5)	6.5–10.0	0.6–9.1
522	41.7(1)	−0.255(4)	1.9–6.4	4.5–10.0
60	41.5(1)	−0.509(4)	2.0–2.9	1.5–6.4

5.2 Characterizing Gamma and Neutron Bands

The next task is to characterize the electron and nuclear recoil bands in the ^{133}Ba and ^{252}Cf data. Ideally, we should measure the profile of each band in $\log\left(\frac{n_e}{n_\gamma}\right)$ as a function of recoil energy E . In fact, neither Xenon10 nor Xed alone can do this, Xenon10 because of the lack of statistics in the electron recoil band, and Xed because in the absence of x-y information we cannot reject events at the edge of the active region that lose charge and create a long tail at low $\log\left(\frac{n_e}{n_\gamma}\right)$. Fortunately, Xenon10 has shown that this tail is removed with a fiducial cut in a 3-D imaging detector [53], and that the remaining distribution in $\log\left(\frac{n_e}{n_\gamma}\right)$ is Gaussian (see Chapter 7). We therefore study the bands in Xed by truncating the charge-loss tail and fitting Gaussians about the peaks of the distributions in $\log\left(\frac{n_e}{n_\gamma}\right)$.

5.2.1 Centroid Estimates

The actual centroid and width of a recoil band are smooth functions of recoil energy, but faced with a finite amount of data, we must slice the data into discrete energy bins and find the mean and width of the band bin by bin. This gives a mean that is the average centroid value in the bin, weighted by the energy spectrum across the bin, but it systematically inflates the width if the centroid is not constant across the bin. Since our bin size is determined by the available statistics, we cannot escape this problem simply by moving to small bins. Instead, we construct a smooth centroid estimate and

shift to centroid-subtracted units

$$\Delta \log \left(\frac{n_e}{n_\gamma} \right) = \log \left(\frac{n_e}{n_\gamma} \right) - \left\langle \log \left(\frac{n_e}{n_\gamma} \right) \right\rangle_{n_\gamma + n_e}, \quad (5.3)$$

where the brackets indicate the centroid estimate as a function of energy (given by $n_\gamma + n_e$). With an accurate centroid estimate, this completely removes bin-size effects from the calculation. For an imperfect centroid, the systematic offset in the measured variance in a bin is given by

$$\delta [\sigma^2] = \langle \delta [\mu]^2 \rangle - \langle \delta [\mu] \rangle^2 = \frac{\int_{E_1}^{E_2} dE (\delta [\mu])^2}{E_2 - E_1} - \left(\frac{\int_{E_1}^{E_2} dE \delta [\mu]}{E_2 - E_1} \right)^2, \quad (5.4)$$

where $\delta [\mu]$ is the difference between the centroid estimate and the true centroid, and E_1 and E_2 are the bin edges.

Our requirement on the centroid estimate is that $\delta [\sigma^2]$ be much smaller than the statistical uncertainty on the variance for the bin sizes we choose. (This is trivial if we choose very small bins, but we want to choose bins such that our statistical uncertainties are similar to the systematic uncertainties in the calibration.) There is no unique way to fulfill this requirement. We construct our centroid estimates by slicing the data into energy bins with $\sim 2,000$ events per bin, fitting a Gaussian about the peak ($\mu \pm \sim 1.5\sigma$) in each bin, fitting a high order polynomial to the fit means to smooth outliers, and taking a cubic spline between the polynomial fit points. This process is repeated on the centroid-subtracted data with shifted energy bins, until a stable centroid is found, typically after two or three iterations. At low energies, the polynomial and spline use $\log(E)$ rather than E as the x-axis. The final centroid estimates are shown over the data in Figs. 5.1, 5.3, and 5.4, and the centroid-subtracted data is shown in Fig. 5.7. In the next section we will show that the $\delta [\sigma^2]$ from this approach is indeed smaller than our statistical uncertainty.

Although the centroid estimates are primarily coordinate transforms to avoid the bin-width issue, smooth centroids are useful in many areas, and these estimates may be used for qualitative statements or in cases where other errors dominate. For example, the nuclear recoil centroid estimate at 522 V/cm was used in Chapter 4 to convert between \mathcal{L} and \mathcal{L}_{eff} in Figs. 4.2 and 4.3. The centroids also reflect the mean recombination versus energy, since

$$\log \left(\frac{n_e}{n_\gamma} \right) = \log \left(\frac{y}{1-y} \right) = \log \left(\frac{1-r}{\frac{a}{b} \frac{N_{ex}}{N_i} + r} \right) \approx \log \left(\frac{1-r}{0.06+r} \right). \quad (5.5)$$

For electron recoils at energies above ~ 10 keVee, we see that recombination decreases as the field in-

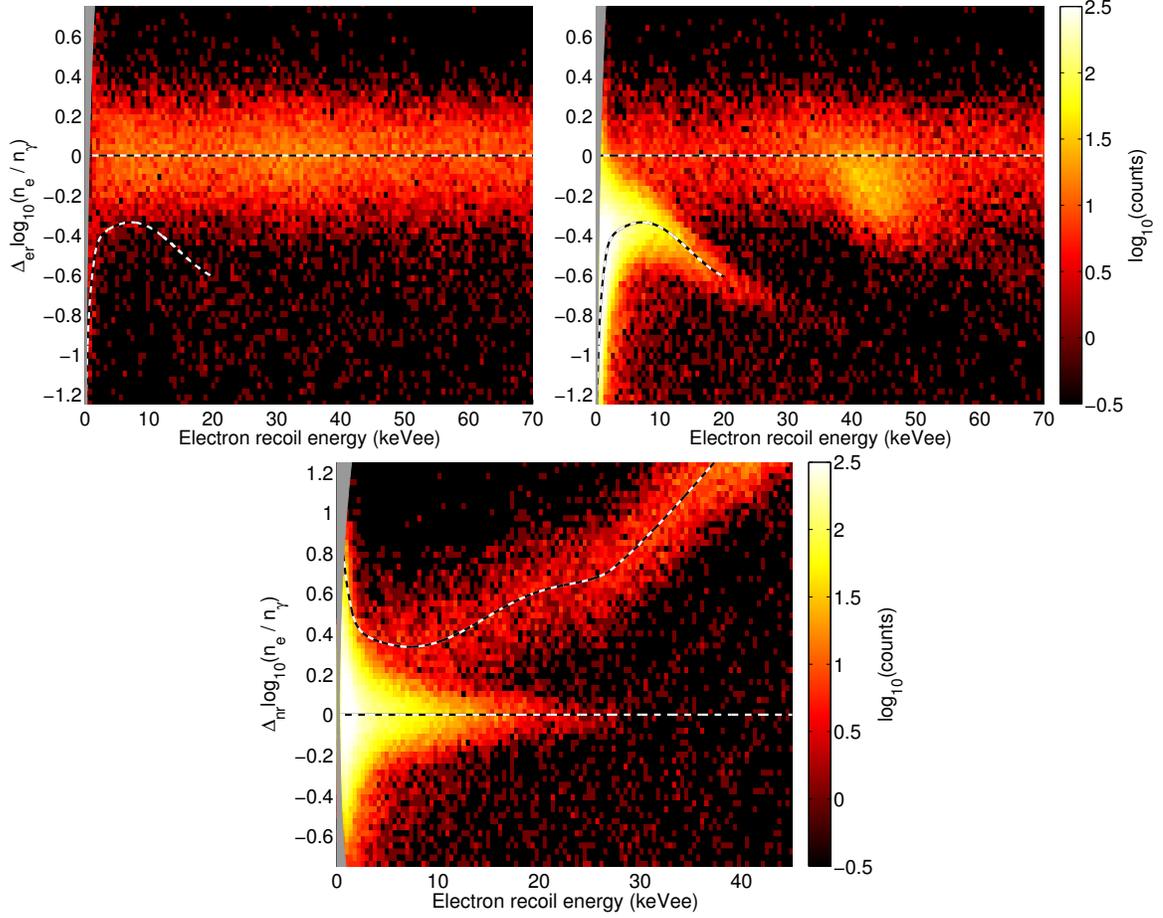


Figure 5.7: Centroid-subtracted data at 876 V/cm. The top two plots have the electron recoil centroid subtracted, while the bottom has the nuclear recoil centroid subtracted. Band widths are measured after subtracting the band’s own centroid, while band separation is measured with the electron recoil centroid subtracted from both bands.

creases and also as energy increases, as we expect based on the electronic stopping power in Fig. 4.11. That recombination reaches a maximum (charge yield reaches a minimum) around ~ 10 keVee is not expected. The electronic stopping power continues to increase as energy decreases, all the way to 1 keVee, so this change in the slope of the centroid breaks the correlation between recombination and stopping power. We also see that field dependence largely disappears for electron recoils at energies below the minimum charge yield, in fact showing behavior very similar to nuclear recoils. These features are explored in much greater detail in Chapter 6.

Continuing to look at the electron recoil centroid in Figs. 5.1 and 5.3, we see that the 40 keV inelastic peak and the 122 keV and 136 keV photo-absorption peaks lie below the Compton scatter band. Unlike Compton scatters, these events produce multiple electron recoils. The inelastic recoil was described earlier — an internal conversion electron and Auger electron are produced at the

primary site, and a photo-electron and Auger electron come from the x-ray absorption site ~ 0.4 mm away. Similarly, the 122 keV gamma is photo-absorbed by a K-shell electron, giving an 87 keV photo-electron and 5 keV Auger electron plus the 30 keV K_α x-ray, which produces another 25 keV photo-electron and 5 keV Auger electron as for the inelastic scatter. These events thus sample the Compton band at energies below the total energy of the event. Since charge yield decreases with energy, the sum charge yield is lower than for Compton events with the same total energy.

If one simply sums the charge yield from the constituent Compton recoils, however, the charge yield is still over-estimated. This is because the upturn in charge yield below 10 keV is associated with overall track size, as will be shown in Chapter 6. The 5 keV Auger electrons, however, are always accompanied by other recoils, and as part of a larger track they experience greater recombination than a lone 5 keV Compton scatter. This effect is borne out in the recombination model presented in Chapter 6.

5.2.2 Band Statistics

Our desired band statistics can now be found by shifting to centroid-subtracted units as in Eq. (5.3), choosing energy bins, and performing the Gaussian fits. We choose our energy bins such that our statistical uncertainties are similar to the systematic errors from the calibration uncertainty. In particular, the last column in Table 5.1 gives the systematic uncertainty on the separation between the electron and nuclear recoil bands, and at low energies we choose bin sizes such that the uncertainty in the mean in each bin is similar to this systematic error. At higher energies we use logarithmic binning to avoid a proliferation of data points (this costs little information, since the band statistics do not change quickly at higher energies). We use the same energy bins at each field.

Within each bin, our Gaussian fits extend to $+2.0\sigma$ and -1.6σ for electron recoils, and $+1.8\sigma$ and -1.6σ for nuclear recoils. The lower bound is set by the tail of low charge-yield events from the edge of the detector, and upper bound for nuclear recoils is to avoid the electron recoil band (present at a lower rate in the nuclear recoil data). We fit each band in its own delta-space to determine its width, and also fit the nuclear recoil band with the electron recoil band centroid subtracted to determine band separation. Band separation and widths are shown in Fig. 5.8.

The band statistics show quantitatively much of what we have already seen from the band centroid estimates, including weak field dependence in band separation at WIMP recoil energies. We can also use these band statistics to verify that our centroid estimates were sufficiently accurate to avoid inflating the measured band widths. Equation (5.4) tells us the width inflation due to

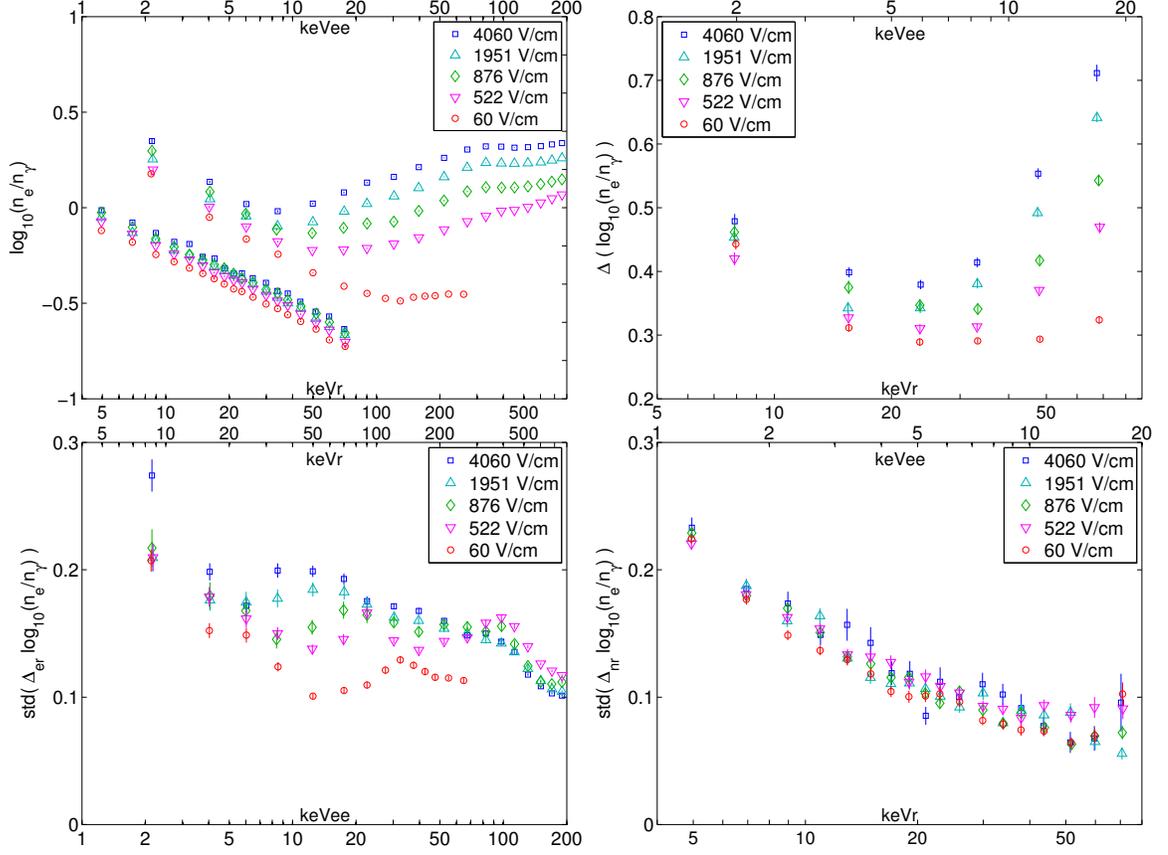


Figure 5.8: Gamma band width (bottom-left), neutron band width (bottom-right), band means (top-left), and band separation (top-right) versus energy. Errorbars show statistical uncertainties. On the means plot, the lower set of points corresponds to nuclear recoils. The systematic uncertainty in band separation, given in the last column of Table 5.1, is similar to the statistical uncertainties shown in the means at low energy.

the error in the centroid estimate — we cannot truly calculate this since we don't know the true centroid, but we estimate $\delta[\mu]$ by taking the difference between two different centroid estimates. For our second centroid estimate, we simply take the spline in $\log\left(\frac{n_e}{n_\gamma}\right)$ versus $\log(E)$ through the means in Fig. 5.8. The resulting $\delta[\sigma^2]$'s are much smaller than the statistical error-bars on our width measurements except for the lowest energy electron recoil bin, where the statistical errors and systematic inflation are comparable (see Fig. 5.9). This is sufficient for the band width decomposition in Section 5.3, and does not effect at all the Monte-Carlo analysis in Section 5.4.

It is tempting to calculate discrimination directly from these band widths and separations, as was in fact done when reporting the Xenon10 result [31]. There are implicit assumptions when doing this, the most obvious of which is that the electron recoil band remains Gaussian in the low charge-yield tail, where Xenon10 lacks statistics and Xed is contaminated by edge events. With current

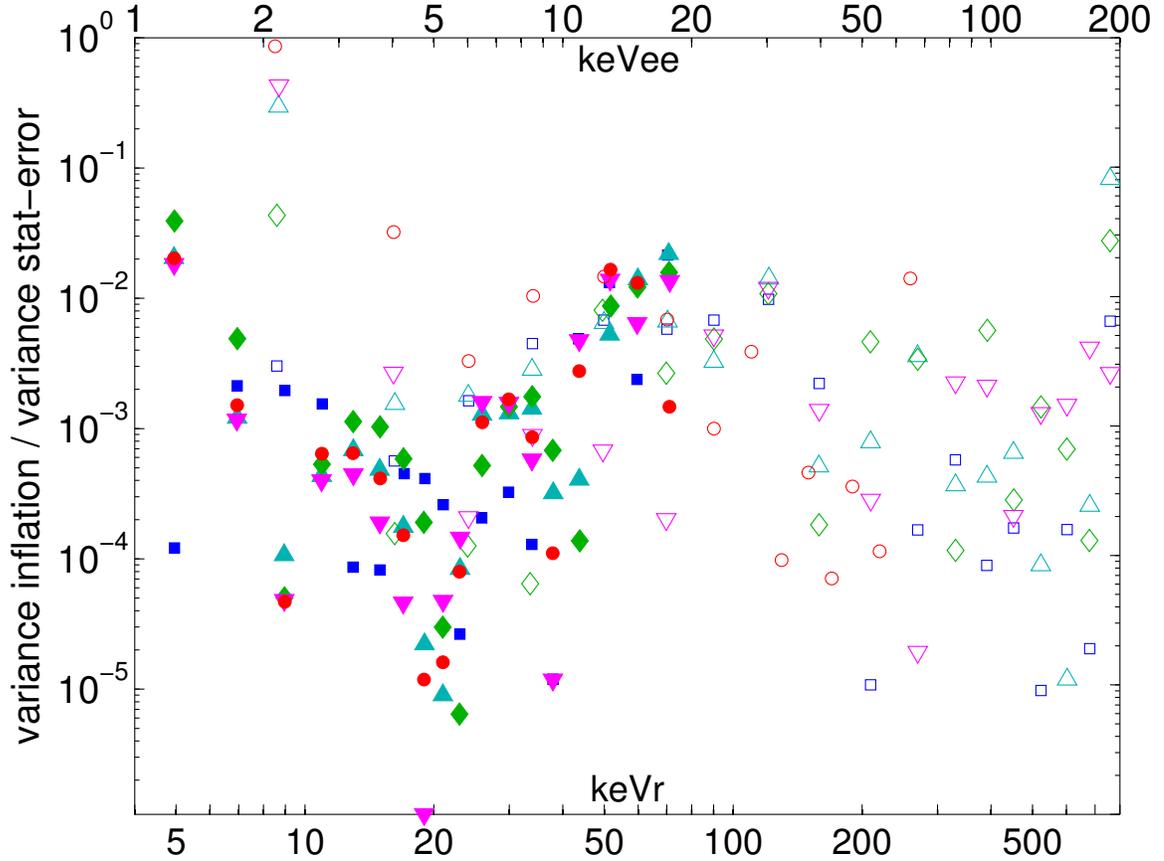


Figure 5.9: Estimated inflation in the measured variance in Xed due to centroid uncertainty, relative to the statistical uncertainties shown in Fig. 5.8. The solid points indicate the nuclear recoil band, empty points the electron recoil band, with data at 4060 (blue \square), 1951 (cyan \triangle), 876 (green \diamond), 522 (magenta ∇), and 60 V/cm (red \circ). The inflation in the variance is at least an order of magnitude below the statistical uncertainty of the measurement except for the lowest energy bin for electron recoils.

data, such an assumption is unavoidable, and we will make a similar one below. The less obvious assumption is that the fluctuations determining discrimination do not move events in energy. This is true for recombination fluctuations (in the Xenon10 case it is true for S2 fluctuations, since there we use an S1 based energy scale), but for generic fluctuations, this definition of discrimination is inadequate.

For example, a useful measure of discrimination is the fraction of electron recoils of a given energy that fluctuate into the WIMP acceptance region. Compare this with the Xenon10 definition, which is the ratio of the number of electron recoil events in the WIMP acceptance region to the total number of electron recoil events observed at the same S1. As we will see, the two dominant fluctuations (recombination and S1) both give a negative correlation between $\log(\frac{n_{\text{e}}}{n_{\gamma}})$ and n_{γ} , so that leakage events are always events that have fluctuated to higher S1. By comparing leakage events to events

at the same S1, the Xenon10 definition of discrimination inflates the number of events considered as candidates to leak into the WIMP region (i.e., the denominator in the discrimination fraction), giving a discrimination power two to five times higher than the more intuitive definition.

The picture becomes more complicated when one is interested in the apparent energy spectrum of leakage events, as is the case when creating a limit using Yellin’s maximum gap method [90]. Since different physical fluctuations move events in different directions, some fluctuations are more likely to give low energy leakage events than others. This is pertinent, since both S1 and recombination fluctuations are dominant in different regions of the WIMP range (as we will see). The band width is therefore insufficient to calculate meaningful discrimination numbers — we also need to decompose the band width into its constituent fluctuations. A Monte-Carlo including the various fluctuations can then give an accurate background and discrimination analysis.

5.3 Analytical Band Width Decomposition

Before building the fluctuation Monte-Carlo, it is illuminating to do an approximate analytical decomposition of the band width. Except for recombination fluctuations, we know how each fluctuation scales with signal size, and we can estimate the effect of each fluctuation of the band width. Adding the known fluctuations in quadrature, we can determine the magnitude of recombination fluctuations from the excess in the measured band width.

We make two approximations, first that S1 and S2 fluctuations are small compared to the curvature of the band centroid, so that the shift in centroid-subtracted value $\Delta \ln \left(\frac{n_e}{n_\gamma} \right)$ associated with the shifts $\delta [S1]$ and $\delta [S2]$ is

$$\delta \left[\Delta \ln \frac{n_e}{n_\gamma} \right] = \left(\frac{\delta [S2]}{S2} - \frac{\delta [S1]}{S1} \right) - W \left(\frac{\delta [S2]}{g2} + \frac{\delta [S1]}{g1} \right) \cdot \frac{d}{dE} \left\langle \ln \frac{n_e}{n_\gamma} \right\rangle_E. \quad (5.6)$$

The first term in the equation is the shift in $\ln \left(\frac{n_e}{n_\gamma} \right)$, while the second is the shift in the centroid value with the changing reconstructed energy of the event. Thus, when $\frac{d}{dE} \left\langle \ln \left(\frac{n_e}{n_\gamma} \right) \right\rangle_E$ is negative, as is the case for both nuclear and electron recoils in the WIMP recoil energy range, S1 fluctuations have a reduced impact on $\Delta \ln \left(\frac{n_e}{n_\gamma} \right)$. That is, S1 fluctuations tend to shift an event along the band without moving away from the centroid. This allows discrimination at much lower energies than would otherwise be possible (see Fig. 5.10).

The second approximation we make is that the band width at E depends on the size of fluctuations

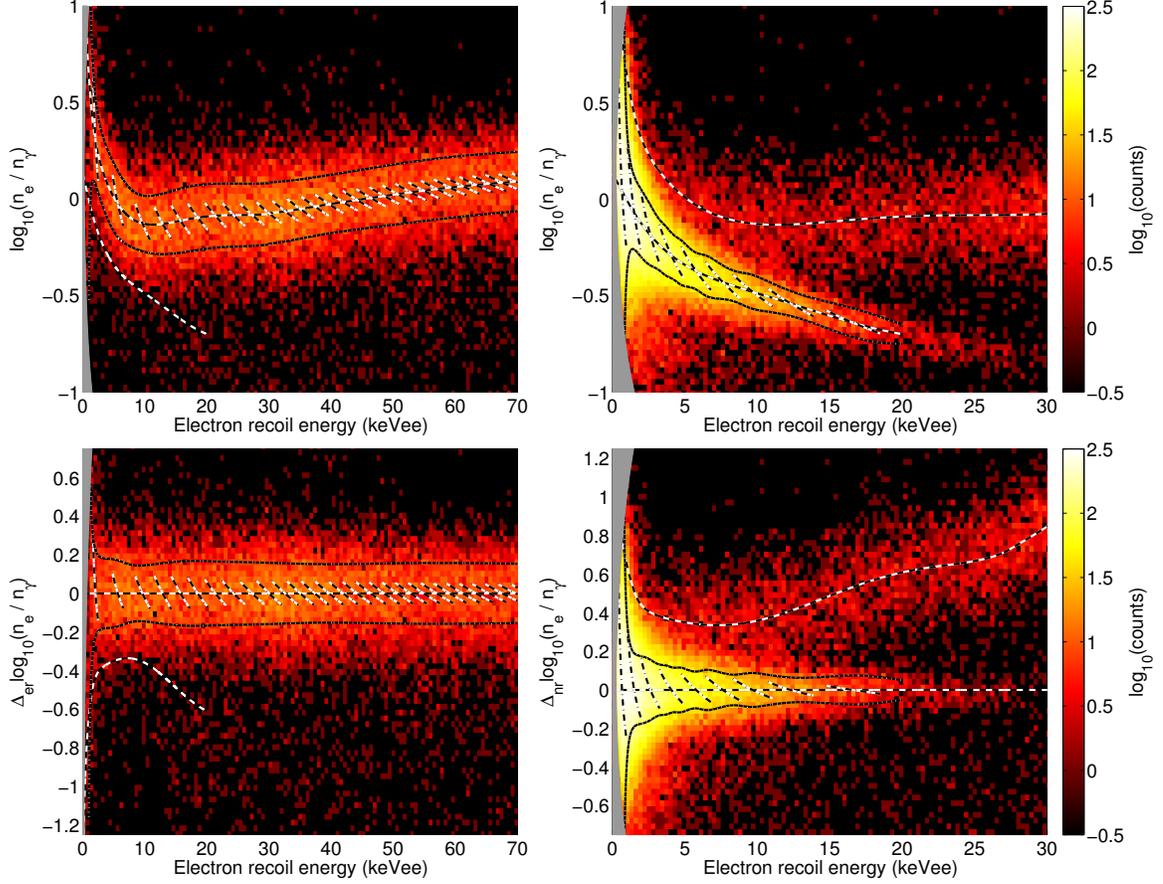


Figure 5.10: Electron (left) and nuclear (right) recoil bands, with centroids subtracted in the lower plots. The dotted lines above and below the centroid indicate the measured 1σ band width. The dashed lines crossing the centroid are lines of constant S2. Fluctuations in S1 move events along these lines, and each crossing line extends to $\pm 1\sigma$ for statistical fluctuations in S1. For electron recoils these fluctuations fall short of the measured band width except at the lowest energies. For nuclear recoils the statistical S1 fluctuations account for almost all of the observed band width. In the WIMP range (~ 10 keVee and below) S1 fluctuations are nearly parallel to the electron and nuclear recoil centroids before centroid subtraction and have a reduced contribution to the band width. (The ‘nearly parallel’ statement does not apply to the centroid subtracted bands since centroid subtraction does not preserve angles, but the contribution of S1 fluctuations to the band width is the same.)

about the centroid at E . This requires that the strengths of S1 and S2 fluctuations, which shift an event in energy, change slowly with energy. Both approximations fail at some point, the former at the minimum electron recoil charge yield where band curvature is high, and the latter near threshold where statistical fluctuations in S1 are increasing rapidly.

Using Eq. (4.41,4.42,5.6) and the instrumental fluctuation values in Table 5.2, we calculate the contributions to band width of statistical and instrumental fluctuations in S1 and S2. These fluctuations are shown in Fig. 5.11 for the 876 V/cm data, along with the measured band widths for both electron and nuclear recoils. The difference between the measured width and the sum of the

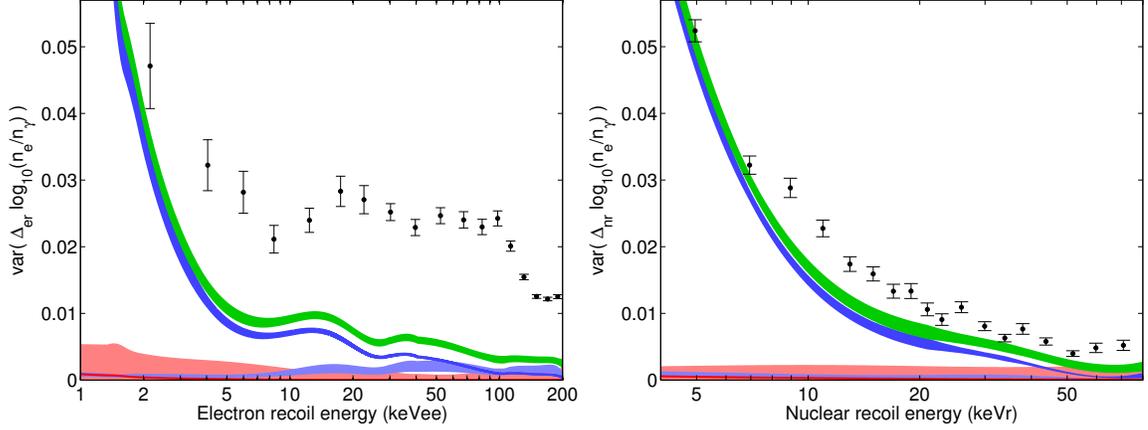


Figure 5.11: Contributions of S1 and S2 statistical and instrumental fluctuations to total band width for electron recoils (left) and nuclear recoils (right) at 876 V/cm. From top to bottom at the left edge, the solid bands are analytical estimates of the effects of total S1/S2 fluctuations (green), S1 statistical fluctuations (blue), S2 instrumental fluctuations (red), S2 statistical fluctuations (red), and S1 instrumental fluctuations (blue). The widths of the bands indicate the systematic error in determining the strength of each fluctuation. The black data points show the measured variance with statistical uncertainties given by the error-bars. The difference between the observed width and the sum of all other contributions (green band) gives the magnitude of recombination fluctuations.

known S1 and S2 fluctuations is inferred to be the recombination fluctuation component. We see immediately that recombination fluctuations are the dominant factor in the electron recoil band width down to 4 keVee, or 16 keVr. The nuclear recoil band, on the other hand, is dominated by statistical fluctuations in S1.

This approach has one additional source of systematic error, which is apparent in the projected contribution of S1 statistical fluctuations to the electron recoil band width. We see large oscillations in the strength of these fluctuations between 10 and 40 keVee, stemming from the second term in Eq. (5.6), where we take the derivative of the centroid estimate. An accurate derivative is a much more stringent requirement on the centroid estimate than we had for band width calculations. The fluctuations in the calculated band width indicate that either our centroid estimate is not up to the task (i.e., the fluctuations in the derivative leading to the band width wiggles are noise), or the oscillations are real, in which case our first approximation fails and this approach is invalid. These difficulties are easily handled in the context of a fluctuation Monte Carlo.

5.4 Monte-Carlo of Fluctuations

In this section, we build and tune a Monte Carlo that simulates the fluctuations in each step of signal generation, from the initial recoil to the signal output by the PMTs. The only significant unknown in

this process is the recombination distribution at a given energy. We make the assumption that this distribution is Gaussian in $\log\left(\frac{1-r}{r}\right)$, where r is the recombination fraction, in order to match the observed Gaussian shape in $\log\left(\frac{n_e}{n_\gamma}\right) \approx \log\left(\frac{1-r}{r}\right)$ in the electron recoil band in Xenon10 at energies where, we know from the previous section, recombination fluctuations dominate. The unknowns in the Monte Carlo are the mean and variance in $\log\left(\frac{1-r}{r}\right)$ as a function of drift field, recoil type, and recoil energy. We tune these parameters so that the Monte Carlo output matches the band statistics from Section 5.2.

5.4.1 Building the Monte Carlo

The Monte Carlo follows events step-by-step through the creation of the measured signal. First the energy of the event in electron equivalent units is randomly sampled from the input spectrum. When tuning to Xed data, we input a flat (constant) spectrum for electron recoils and a piecewise-exponential spectrum for nuclear recoils, fit to the observed spectra. We then calculate the number of ions N_i and direct excitons N_{ex} in the track. We take $N_i + N_{ex}$ from a Gaussian distribution, with a mean determined by the W-value $W = 13.7$ eV, and a variance of $F \cdot (N_i + N_{ex})$, where F is the Fano factor. This has not been directly measured in liquid xenon, so we take the theoretical value $F = 0.05$ [71]. We then separate $N_i + N_{ex}$ into N_i and N_{ex} following a binomial distribution, with a mean $\frac{N_{ex}}{N_i} = 0.06$, also based on theory.

Next we compute the number of ions that recombine. We supply the Monte Carlo with smooth functions giving the mean and variance of $\log\left(\frac{1-r}{r}\right)$ versus electron equivalent energy. Rather than taking the actual recoil energy, the Monte Carlo passes the energy $W_q N_i$ to these functions, taking $W_q = 1.06W$ (see Eq. (4.4)). We sample r from the resulting Gaussian distribution in $\log\left(\frac{1-r}{r}\right)$ and translate this to a number of recombining ions, rounding to an integer number of ions. At this point we have the true number of S1 scintillation photons and S2 electrons for the event, $n_\gamma = N_{ex} + rN_i$ and $n_e = (1-r)N_i$.

Creating the photo-electrons in the PMTs is a two step process. First we determine the gains g_1 and g_2 in $\text{phe}/[n_\gamma]$ and $\text{phe}/[n_e]$, including S1 and S2 instrumental fluctuations as event-by-event fluctuations in the gains. We take g_1 and g_2 to be Gaussian distributions, with means given by the calibration parameters in Table 5.1 and widths given by the midpoints of the instrumental fluctuation ranges in Table 5.2. Once we have the gains, we step through the statistical process of creating photo-electrons. For S1, each scintillation photon has a probability g_1 of creating a photo-electron, so the total number of S1 photo-electrons follows a binomial distribution with n_γ

trials and mean $g_1 n_\gamma$. For S2, the number of electrons successfully extracted into the gas follows a binomial distribution, electroluminescence photons are created in a Poisson process as the extracted electrons travel through the gas, and these photons in turn follow binomial statistics when creating photo-electrons. The first binomial process has a success rate near one (assuming negligible charge loss to impurities and a strong extraction field, as in Xed) and may be ignored. A Poisson process followed by a binomial process is merely a Poisson process with a lower mean than the original, so the number of S2 photo-electrons follows a Poisson distribution with mean $g_2 n_e$. For speed in Matlab the distributions for S1 and S2 photo-electrons are modeled as Gaussians with variances of $g_1(1 - g_1)n_\gamma$ and $g_2 n_e$, respectively, whenever the expected number of photoelectrons is ≥ 10 . This is always true for S2 light, since $g_2 \approx 14 \text{ phe}/[n_e]$. For events with $g_1 n_\gamma < 10$ an actual binomial distribution is sampled to determine the number of S1 photo-electrons.

The final step is to sample the single photo-electron spectrum for each photo-electron to give the measured S1 and S2. For Xed, we take this to be the log-normal distribution found in Chapter 4, with $\frac{\sigma}{\mu} = 0.437$, where μ is the mean and σ is the square root of the variance. The measured S1 and S2 are then scaled to units of $[n_\gamma]$ and $[n_e]$ using the mean g_1 and g_2 values.

5.4.2 Tuning the Monte Carlo

The only inputs to the Monte Carlo not known from other measurements are the mean and width of the recombination distribution as a function of energy, recoil type, and drift field. These parameters must be tuned so that the Monte Carlo reproduces the band statistics found in Section 5.2. More precisely, we want to tune the recombination parameters for each band (both recoil types at each field) so that, when simulating Xed data and performing Gaussian fits identical to those in Section 5.2 (i.e., using the same centroid subtraction, bin edges, and fit regions), we get means and widths consistent with those found in the real data.

For a given band, the mean and width of the recombination distribution should be smooth functions in energy. We constrain both the mean and the log of the variance in $\log\left(\frac{1-r}{r}\right)$ to be cubic splines of $\log(E)$. Knots (discontinuities in the 3rd derivative) in both splines are fixed at the centers of the bins used for the Gaussian fits in Section 5.2, and we impose vanishing 2nd derivatives at the first and last knots. Our initial guess for the mean recombination is the smoothest spline that fits the means from Section 5.2 with a reduced χ^2 of 1, where we have converted the means to $\log\left(\frac{1-r}{r}\right)$ using $\frac{N_{\text{ex}}}{N_i} = 0.06$. By ‘smoothest’ we mean the spline with the minimum integral of the square of the second derivative.

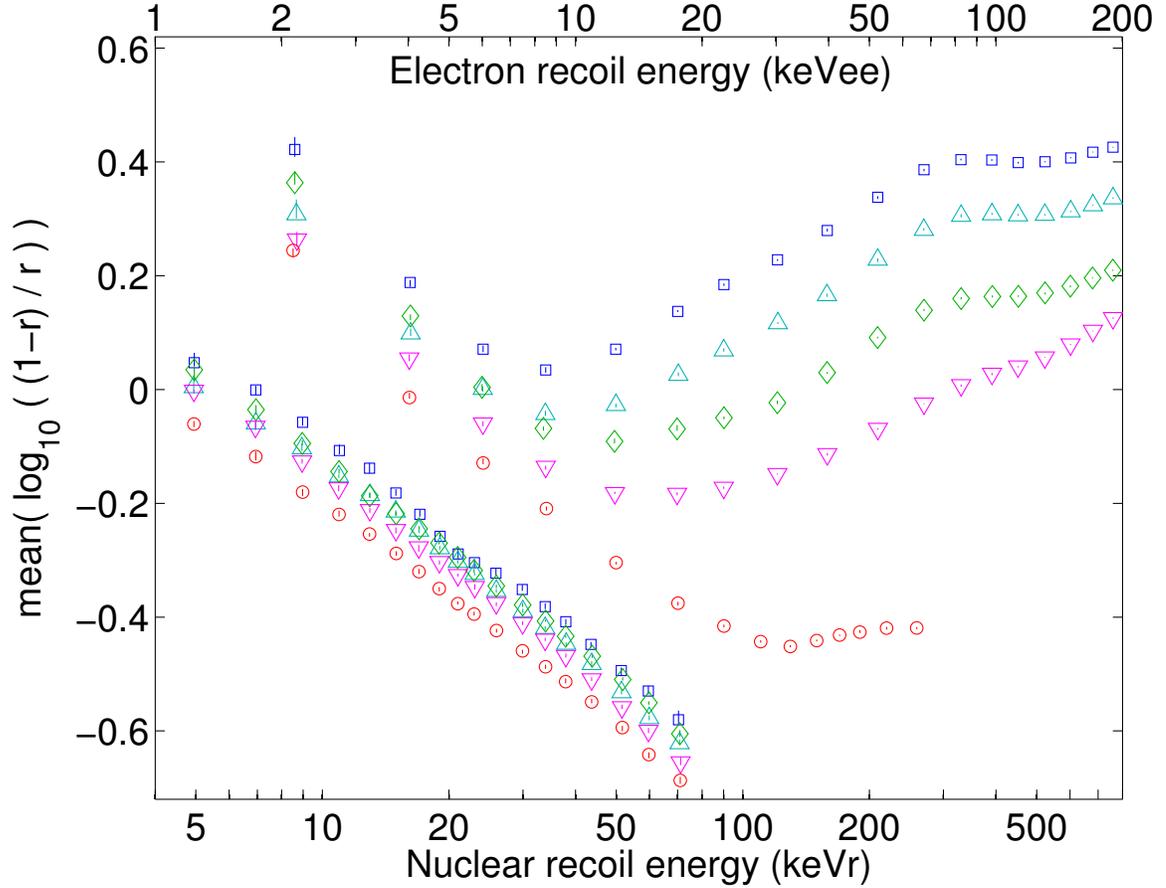


Figure 5.12: Data points show the knots of the spline input to the Monte-Carlo for the mean $\log_{10} \left(\frac{1-r}{r} \right)$, with errorbars showing statistical uncertainties. Fields shown are 4060 (blue \square), 1951 (cyan \triangle), 876 (green \diamond), 522 (magenta ∇), and 60 V/cm (red \circ). The upper sets of points are for electron recoils, and the lower are for nuclear recoils. The recombination fraction r is calculated from the normalized yield y using $\frac{a}{b} \frac{N_{ex}}{N_i} = 0.06$.

We run the Monte Carlo with this mean and zero recombination fluctuations, then create a guess for the recombination fluctuations from the difference between the simulated band variance in each energy bin and the measured variance. Again, the guess is the smoothest spline consistent with the data. We repeat the process, running with the new recombination variance to improve the centroid spline, then again with the new centroid spline for a final shot at the variance. We find that this level of iteration is sufficient to achieve agreement between the Monte Carlo simulations and the Xed data.

The tuning result depends on both the other inputs to the Monte Carlo (g_1 , g_2 , instrumental fluctuations, and the photo-electron spectrum width), and on the band statistics being matched. Both of these have uncertainties that we wish to propagate through the tuning and upcoming discrimination calculations. To find the uncertainties in the tuned parameters, we repeat the tuning

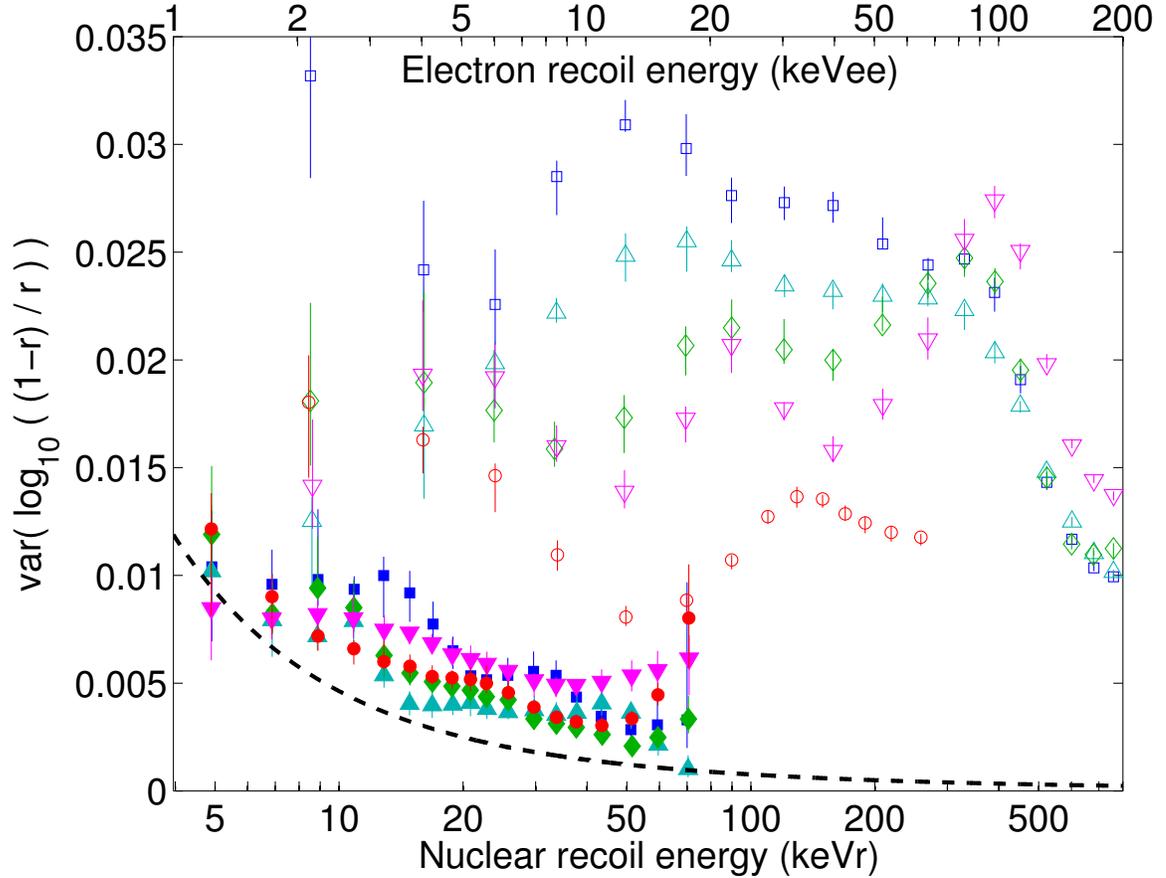


Figure 5.13: The same as Fig. 5.12, but showing the variance in $\log_{10} \left(\frac{1-r}{r} \right)$ rather than the mean. Solid points are for nuclear recoils, empty points are for electron recoils. The dashed black line gives the variance expected from statistical fluctuations for a binomial process with r given by the nuclear recoil centroid at 876 V/cm — the binomial fluctuations expected for electron recoils are similar. The recombination fluctuations observed in nuclear recoils are nearly consistent with the binomial model, but those observed in electron recoils are much larger, especially at high energies. The drop in recombination fluctuations in electron recoils around 100 keVee may indicate photo-absorption events entering the data at that energy (see discussion in Chapter 6).

using sets of inputs and band statistics varied randomly about their measured values according to their uncertainties. We find systematic uncertainties by varying the inputs, statistical uncertainties by varying the band statistics being matched, and total uncertainties by varying both. The tuned parameters are uncertainties are shown in Figs. 5.12 and 5.13. When using the Monte Carlo for discrimination calculations, we run the simulation with every set of tuned parameters found in this error analysis and use the resulting spread in discrimination values to determine the uncertainty in discrimination.

We will look in detail at the physics behind the recombination parameters we have just found in Chapter 6, but it is worth comparing now the measured recombination fluctuations with those

we would see if recombination were a binomial process. This gives the variance $\sigma_{rN_i}^2 = r(1-r)N_i$, which translates into a variance in $\log\left(\frac{1-r}{r}\right)$ of

$$\sigma_{\ln\left(\frac{1-r}{r}\right)}^2 = \frac{1}{r(1-r)N_i}. \quad (5.7)$$

We see in Fig. 5.13 that this is nearly consistent with the nuclear recoil band at low energies, but is much lower than the observed recombination fluctuations in electron recoils, especially at high energies.

5.4.3 Band Decomposition Revisited

We can repeat the analysis from Section 5.3 using the Monte Carlo, turning on and off various fluctuations in the simulation and looking at the resulting band widths. Figs. 5.14 and 5.15 show the simulated band widths from individual and combined statistical and instrumental fluctuations over the original analytic decomposition at each field. Recombination fluctuations still dominate the electron recoil band width down to 4 keVee, and statistical fluctuations in S1 dominate at lower energies and in nuclear recoils at all energies, but the irregularities due to the fluctuations in the derivative of the centroid estimate are mostly gone. Again, the effect of S1 fluctuations at WIMP recoil energies is suppressed by the slopes of the electron and nuclear recoil bands. The simulated band width with all fluctuations included on is also shown in these plots, and it correctly follows the data (as it was tuned to do) except in the lowest energy bin for electron recoils. At all but the highest drift field, the simulated fluctuations in this bin are $1-2\sigma$ higher than observed. This is due to the smoothness requirement we imposed on the recombination fluctuations versus energy, and may cause us to over-estimate leakage by ~ 1 error-bar in the lowest bin in the following discrimination analysis.

5.5 Electron Recoil Discrimination

With the tuned Monte Carlo, it is simple to calculate any discrimination parameter we choose. We can also simulate discrimination in other detectors by inputting the appropriate instrumental fluctuations, S1 and S2 gains, and photo-electron spectrum.

It is possible for a large scale detector to eliminate at WIMP recoil energies all fluctuations except for those from recombination and S1 light collection statistics. Beyond that, if we wish to improve discrimination, we must face these two fluctuations. We can decrease S1 fluctuations by

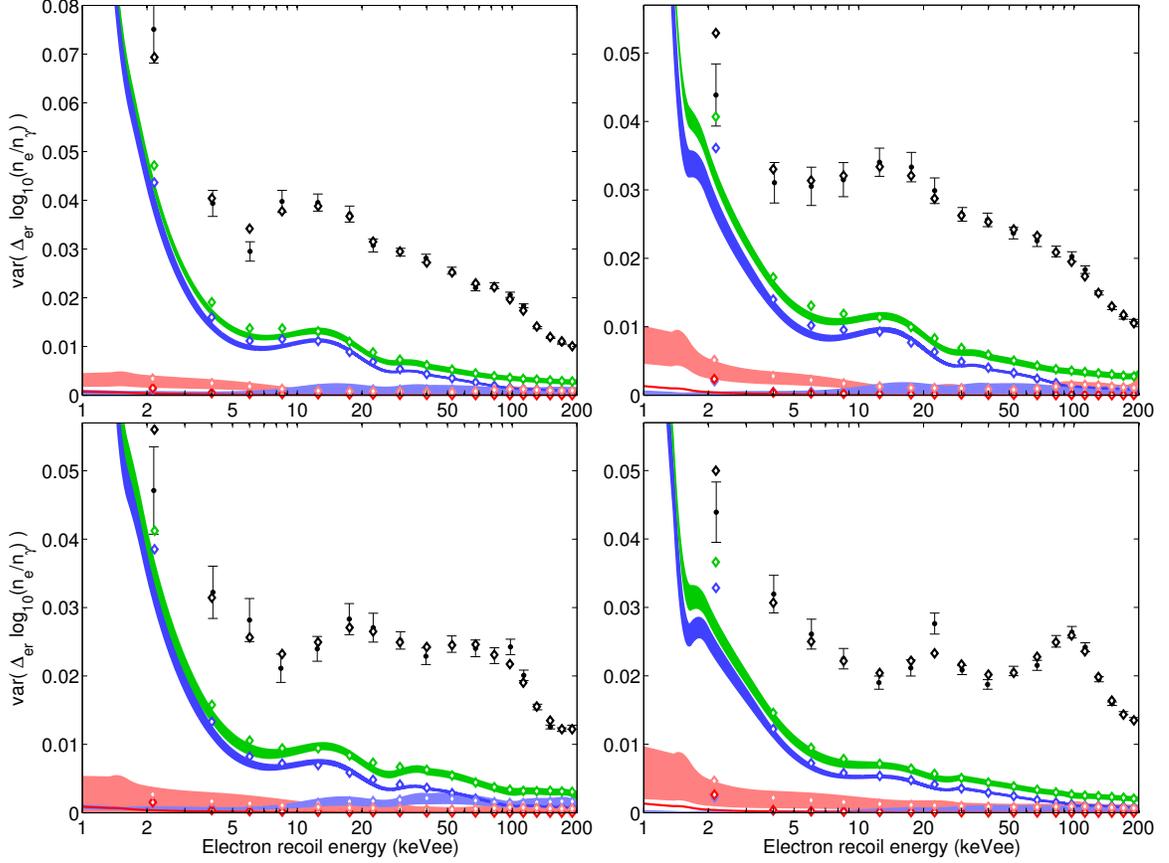


Figure 5.14: Decomposition of the electron recoil band width at 4060 (top-left), 1951 (top-right), 876 (bottom-left), 522 (bottom-right), and 60 V/cm (bottom-left of Fig. 5.15). The solid bands show the approximate analytic decomposition of the band width as in Fig. 5.11, and the points with error-bars show the measured band width. The empty diamonds show the simulated band widths with all fluctuations (black), all fluctuations except recombination (green), and each individual known fluctuation (colors matching the corresponding bands). To the extent that our analytic decomposition was accurate, the colored diamonds lie atop the solid band of the same color. The Monte-Carlo verifies that the oscillations seen in the contributions of known fluctuations were largely artifacts of the analytical calculation, and also corrects the contribution of S1 statistical fluctuations at low energies where our small fluctuation approximation fails. The black diamonds demonstrate the Monte-Carlo tuning, and should lie atop the measured band widths, smoothing over outliers. Note that at all except the highest field, the band width in the lowest energy bin is over-estimated by $1-2\sigma$. This is due to our smoothness requirement for recombination fluctuations, and may cause us to over-estimate leakage in the lowest energy bin in the following discrimination calculations by ~ 1 error-bar.

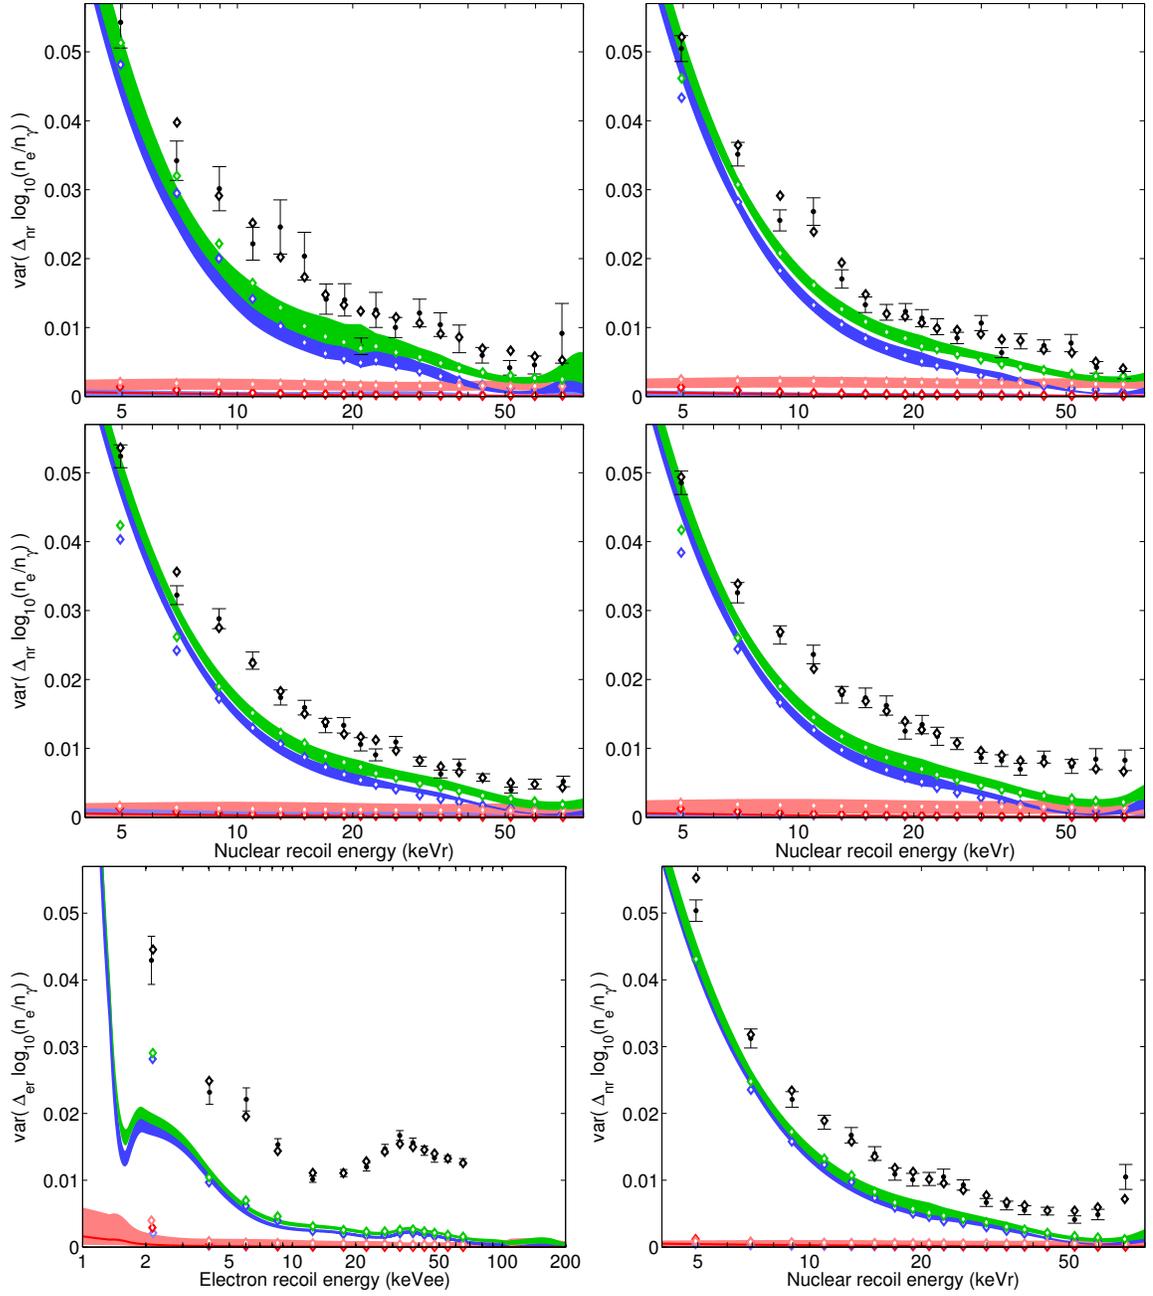


Figure 5.15: Decomposition of the nuclear recoil band width at 4060 (top-left), 1951 (top-right), 876 (mid-left), 522 (mid-right), and 60 V/cm (bottom-right). (Bottom-left shows 60 V/cm electron recoils, continuing from Fig. 5.14.) The bands, data points, and diamonds are as in Fig. 5.14.

increasing the total S1 efficiency, which will significantly improve discrimination at low energies where these fluctuations dominate (and where the WIMP signal is strongest). We have less control over recombination, either to increase band separation or reduce recombination fluctuations. There are only two physical knobs available that might affect recombination: the applied drift field and the liquid density. The former we know is important at energies above the WIMP range, but there is no guarantee that increasing the drift field helps discrimination at low energies. The latter we have not investigated — the liquid density can only vary by $\sim 5\%$ in the operational vapor pressure range of most detectors, and in the model in Chapter 6 this would have a much smaller effect on recombination in the WIMP region than the large shifts available to us in drift field.

5.5.1 Reproducing Xenon10 Discrimination

To test the predictive ability of our Monte Carlo, we first simulate the Xenon10 detector and compare with its published results. We take $g_1 = 0.087 \text{ phe}/[n_\gamma]$ (5.6 phe per keVee at 122 keV at zero field), $g_2 = 23.7 \text{ phe}/[n_e]$, $\sigma_{phe} = 0.5 \text{ phe}$, and zero instrumental fluctuations, the last having been essentially removed by the position dependent calibrations described in Chapter 7. We use the recombination parameters tuned to our 876 and 522 V/cm data — Xenon10 ran at 730 V/cm.

To make the comparison, we precisely follow the Xenon10 prescription for calculating discrimination. We simulate electron recoils assuming a flat energy spectrum, and find the electron recoil centroid in $\log\left(\frac{n_e}{n_\gamma}\right)$ vs n_γ in the simulated data. We then shift to centroid subtracted units as before, but now subtracting the centroid along lines of constant n_γ instead of constant energy. We then slice the data in n_γ , using the same bins as in [31], and fit a Gaussian about the electron recoil peak in each bin. To determine the WIMP acceptance window, we simulate nuclear recoils, taking an exponential energy spectrum with a mean of 8 keVr to match the simulated spectrum from the AmBe source in [79]. We shift to electron-centroid subtracted units and slice in n_γ bins as before, and in each bin choose a threshold to match the nuclear recoil acceptance quoted in [31]. We then calculate, based on the Gaussian fit, the fraction of electron recoil events in each bin falling in the WIMP acceptance region.

Figure 5.16 shows the simulated and reported Xenon10 discrimination. While these are generally consistent, the Xenon10 discrimination shows a more dramatic change with energy than the projections. Given the large error-bars this may be acceptable, but it may also indicate that we are slightly overestimating electron recoil recombination fluctuations in our lowest energy bin (5–12 keVr).

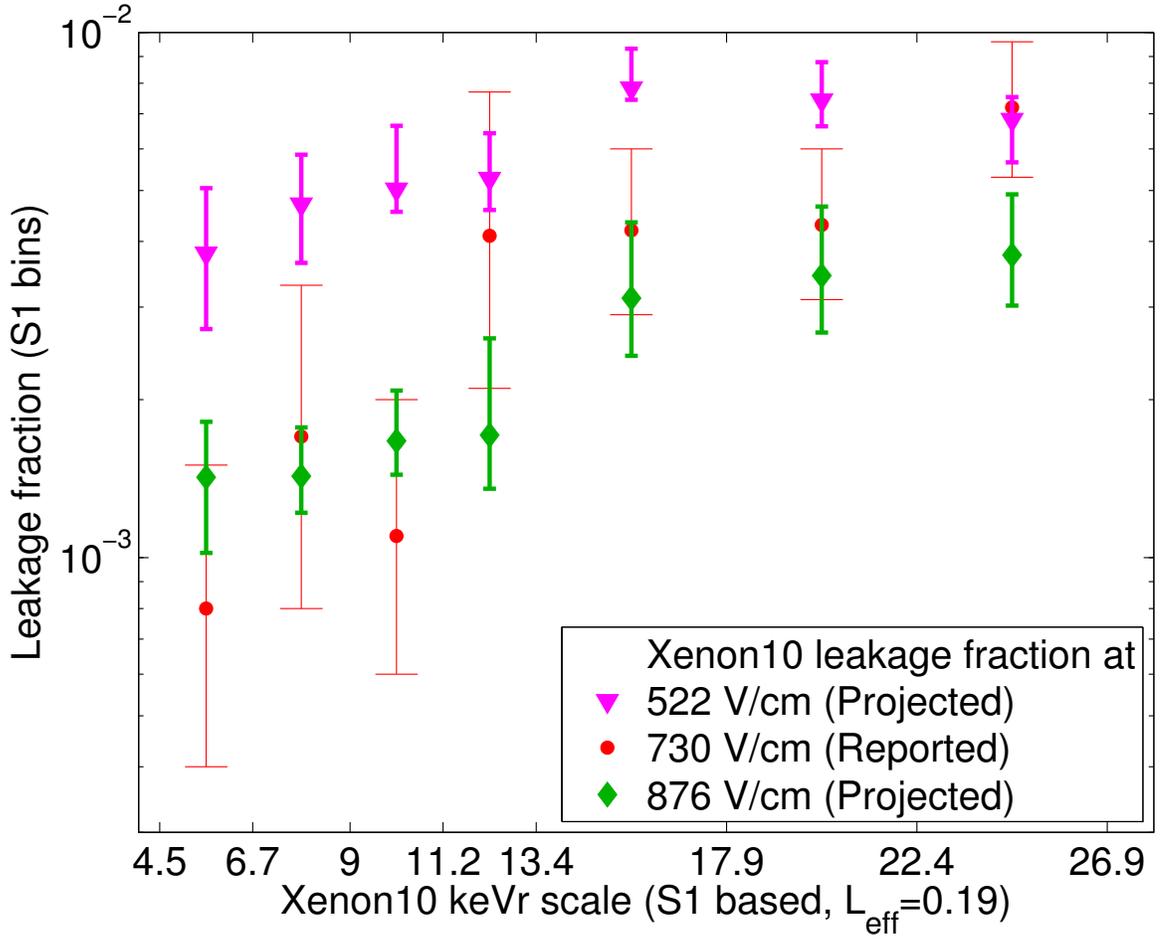


Figure 5.16: Projected discrimination for Xenon10, calculated along lines of constant S1 at drift fields of 876 V/cm and 522 V/cm, compared to the reported Xenon10 discrimination at 730 V/cm [31]. The projected discrimination is calculated for 44%–49% acceptance, varying by S1 bin, as done for the Xenon10 result. The tick marks on the x-axis indicate the binedges used for calculating discrimination. The Xenon10 discrimination generally lies between our projections at higher and lower field, but shows a more dramatic improvement at low energy. Given the large error-bars this may be acceptable, but it may also indicate that we are slightly overestimating electron recoil recombination fluctuations in our lowest energy bin (5–12 keVr) — see Fig. 5.14.

5.5.2 True Discrimination Calculations

With the Monte Carlo, there is no need to restrict ourselves to discrimination calculations that can be pulled from the ‘measured’ data. That is, we need not compare leakage ratios in bins based on measured parameters. We can instead, for example, look at what fraction of electron recoils at a given real energy are measured to be in the WIMP acceptance region. We will refer to this as the ‘Leakage Fraction’. As noted above, we also care about the measured energy spectrum of leakage events. For this we define the ‘Relative Leakage Rate’ as the differential rate of electron recoils in the WIMP acceptance region for a flat electron recoil background, divided by the background rate

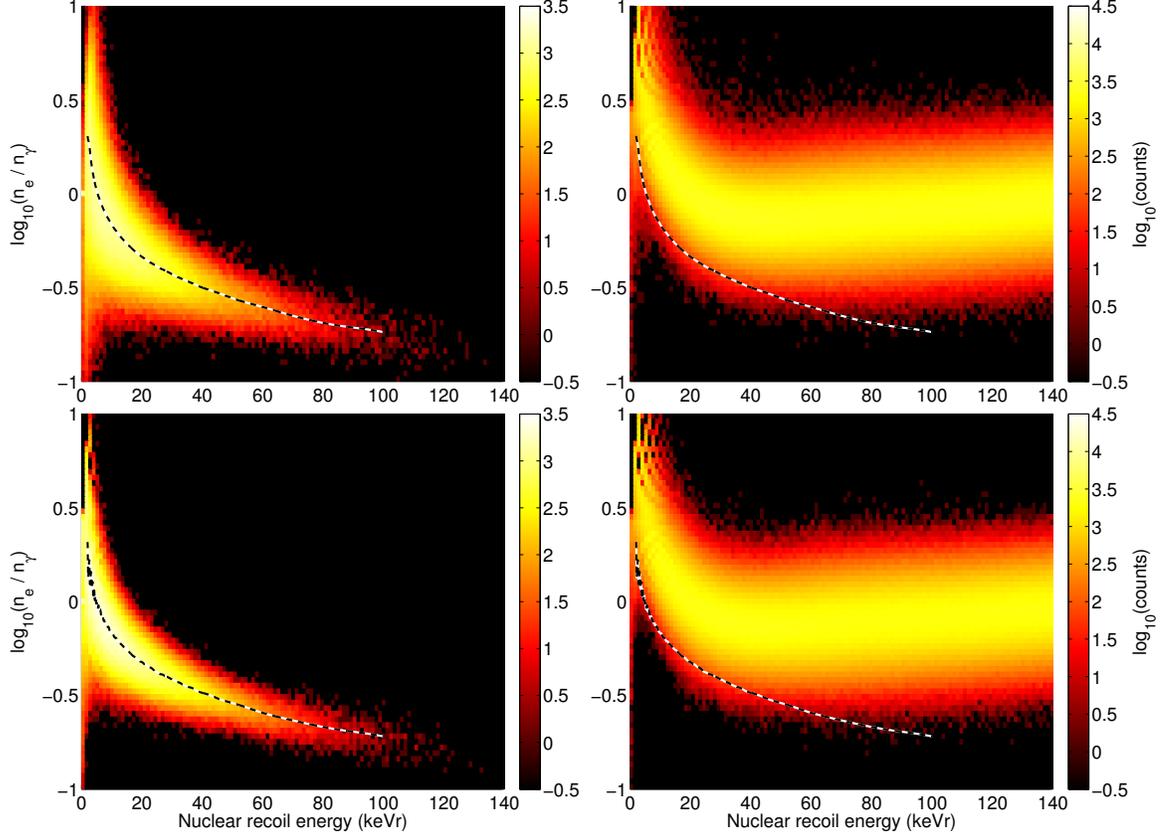


Figure 5.17: Simulated data for Xed (above) and an ideal detector (below) with $g_1 = 0.174 \text{ phe}/[n_\gamma]$ (twice the Xenon10 light collection). The plots on the left show a simulated WIMP signal, nuclear recoils with an exponential spectrum with a mean 12 keVr recoil. The plots on the right show the simulated background, a flat spectrum of electron recoils. The dashed line is the centroid of the simulated WIMP band — the area underneath the centroid is the WIMP acceptance region in our discrimination calculations. The ideal detector simulation takes $\sigma_{phe} = 0$, and in the upper-left of the simulated background the stripes correspond to integer numbers of photo-electrons (in the rest of the plot the separation between stripes is smaller than the histogram bin width).

(in the same units — i.e., when dividing rates, both the leakage rate and background rate should be in events/keVee).

Both the Leakage Fraction and Relative Leakage Rate depend on the definition of a WIMP acceptance region. We choose this region such that, at any measured energy, there is 50% acceptance for WIMP events. To calculate this, we simulate nuclear recoils with a falling exponential spectrum of mean 3 keVee (12 keVr, corresponding to a ~ 100 GeV WIMP), slice the result into small bins in the measured energy, find the median in each bin, and take a smoothing spline in $\log\left(\frac{n_e}{n_\gamma}\right)$ versus $\log(E)$ through the medians. The WIMP acceptance region is the area below the spline (see Fig. 5.17).

Finally, we should reiterate that we are assuming a Gaussian distribution in $\log\left(\frac{1-r}{r}\right)$ when

running the simulations that produce the following discrimination numbers. Xenon10 data suggests that this is valid out to at least 3σ (99.87% discrimination) based on events from 2–20 keVee (see Chapter 7), but with current data we cannot rule out non-Gaussian tails in recombination that develop at low energies or below 3σ .

5.5.3 Discrimination Projections

We now find the Leakage Fraction and Relative Leakage Rate at each our drift fields in five simulated detectors. The first two of these are Xed and Xenon10, with input parameters as described above. The last three are ‘ideal’ detectors, in that they have zero instrumental fluctuations in S1 and S2, and zero photo-electron spectrum width. Negligible instrumental fluctuations are easily achievable for the small signals at WIMP recoil energies in detectors with 3-D position reconstruction — by correcting for the position dependence of the S1 and S2 calibrations, one sees instrumental fluctuations at the level of the collimated 122 keV peak in Xed. The zero photo-electron width is less obvious, but it is also readily achieved for small S1s in large detectors. When a small S1 is divided over many PMTs, the majority of the PMTs see either zero or one photo-electrons, and those that see two are likely to see them separated in time (this is especially true for the next generation detectors, where the photon travel time can be several hundred nano-seconds with multiple reflections, see [44]). As long as one photo-electron is easily distinguishable from zero, the detector becomes digital, giving discrete numbers of photo-electrons. This does not hold for the S2 signal, but even with non-zero photo-electron width statistical fluctuations in S2 are negligible. The ‘ideal’ detector is thus described by a single parameter — the total S1 efficiency, g_1 . We simulate ideal detectors with g_1 of 2x and 3x that of Xenon10, and one with $g_1 = 1$, or 100% light collection efficiency. Twice the Xenon10 light collection is a reasonable goal for next generation detectors, which will already benefit from PMTs with higher quantum efficiencies and photo-electron collection efficiencies. There is also research into materials that are reflective at 175 nm, which could optimistically boost light collection to three times the Xenon10 value. The $g_1 = 1$ detector is not realistic, but illustrates the effect of the inescapable recombination fluctuations.

Figures 5.18 and 5.19 show the Leakage Fraction and Relative Leakage Rate at 876 V/cm for the five detectors. The former looks at the actual recoil energy of leakage events, while the latter looks at the measured energy. Where recombination fluctuations dominate, these numbers are the same, since recombination fluctuations don’t move an event in energy. Figure 5.20 shows the Relative Leakage Rate normalized to the Xenon10 projection, with error-bars indicating the uncertainty in

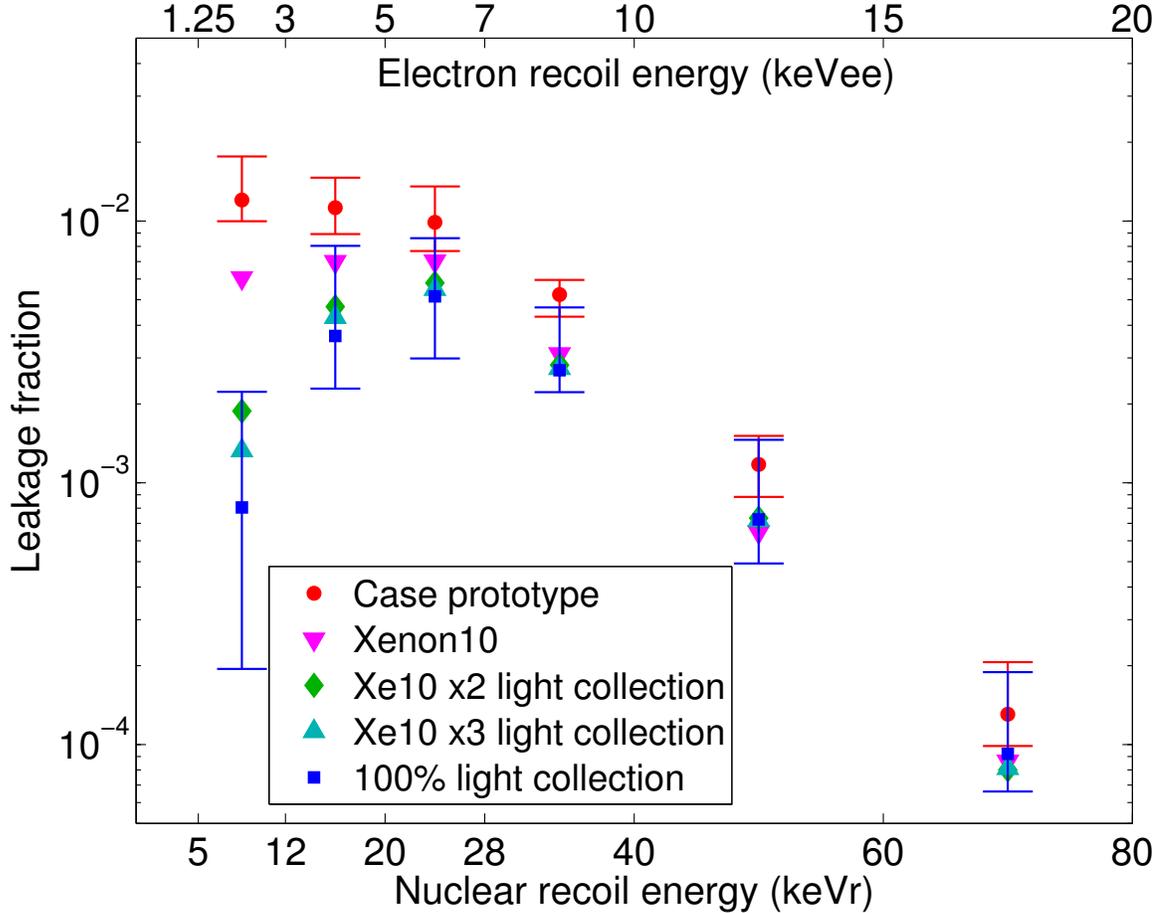


Figure 5.18: The leakage fraction as a function of recoil energy, at a drift field of 876 V/cm in various detectors. Tick marks on the x-axis indicate energy bins used for the calculation. The data points give the leakage in the Case prototype ($g_1 = 0.060$), in Xenon10 ($g_1 = 0.087$), and in ideal detectors with twice the Xenon10, three times Xenon10, and perfect light collection. Twice the Xenon10 light collection is a reasonable goal for next generation detectors. Note the difference between the relative leakage rate and the Xenon10 definition of discrimination — the magenta ∇ 's in this figure and the green \diamond 's in Fig. 5.16 are from the same simulations, and differ by a factor of five at low energies. Error-bars on the Case and $g_1 = 1$ points indicate total uncertainties (dominated by statistical uncertainties). Uncertainties in the different projections are correlated. We may be overestimating leakage in the lowest energy bin by ~ 1 error-bar due to our smoothness requirement for recombination fluctuations versus energy.

the improvement with increased light yield. At most energies there is only room for a x2 reduction in leakage from the Xenon10 value before hitting the barrier of recombination fluctuations. In the lowest energy bin (5–12 keVr) a x6 reduction in leakage is expected when doubling the Xenon10 light collection (error-bars from x3.3–x11), but diminishing returns for further improvements in light collection efficiency. In terms of absolute discrimination, the x2 detector gives $\sim 99.6\%$ rejection in most of the WIMP region, improving to better than 99.8% in the lowest energy bin. Note that our lowest energy bin estimates may be pessimistic — if recombination fluctuations drop as sharply

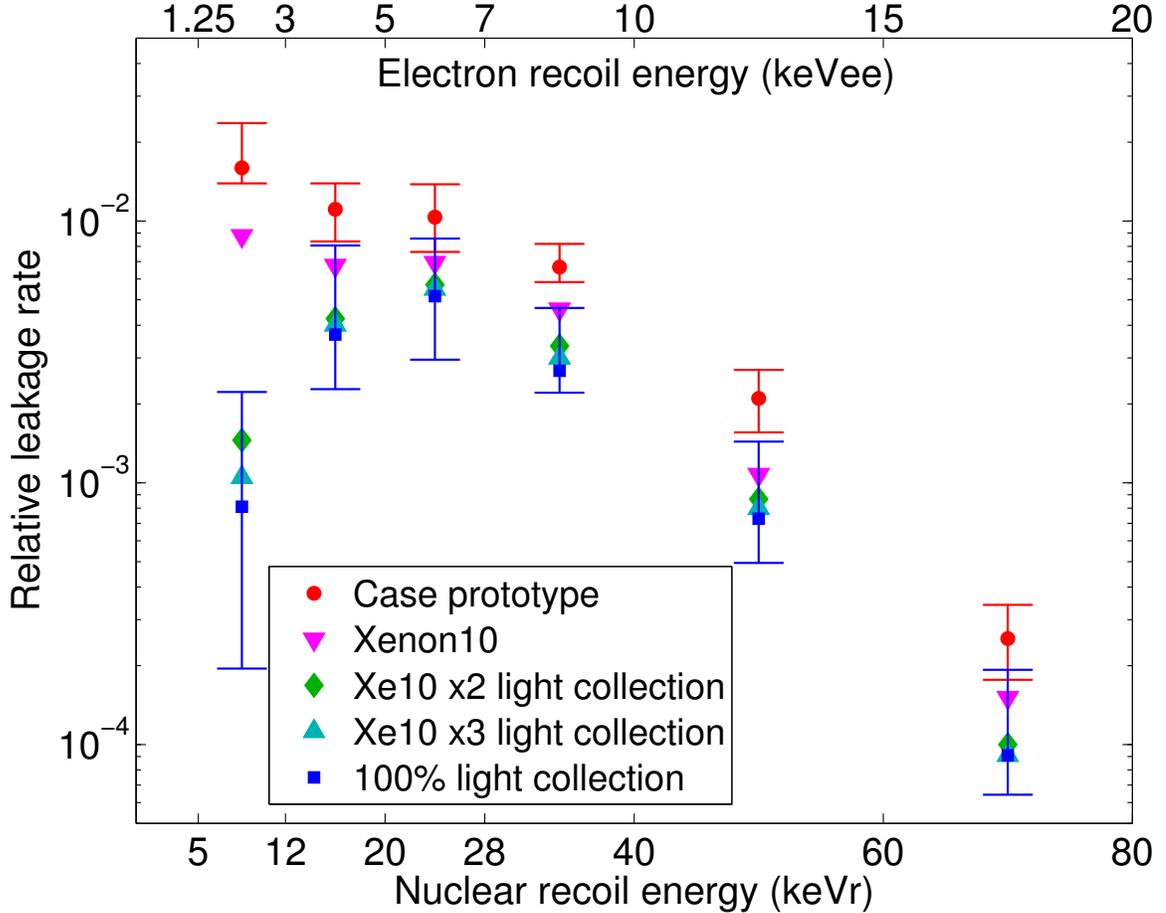


Figure 5.19: The energy spectrum of leakage events resulting from a flat electron recoil background, at a drift field of 876 V/cm in various detectors. These numbers are very similar to those of Fig. 5.18, but here the x-axis indicates the measured energy of the leakage event rather than the true energy. Error-bars on the Case and $g_1 = 1$ data indicate total uncertainties. Uncertainties in the different projections are correlated — see Fig. 5.20 for uncertainties on the relative improvement between detectors. Again, we may be overestimating leakage in the lowest energy bin by ~ 1 error-bar due to our smoothness requirement for recombination fluctuations versus energy.

at low energy as our data may suggest (and was disallowed by our smoothness requirement in the Monte Carlo tuning), the lowest energy discrimination projections may be systematically too high by roughly one error-bar.

To escape recombination fluctuations, our only viable option is to increase the drift field. Figure 5.21 shows the Relative Leakage Rate in the ideal detector with x2 Xenon10 light collection at each of our drift fields. We find relatively little field dependence for discrimination in the WIMP range. There may be a factor of two reduction in leakage in the lowest energy bin when increasing the field from 876 to 1951 V/cm, but this is within errors. Going to 4060 V/cm actually appears to makes discrimination worse, partly because of the smaller S1 signals. On the other hand, our lowest

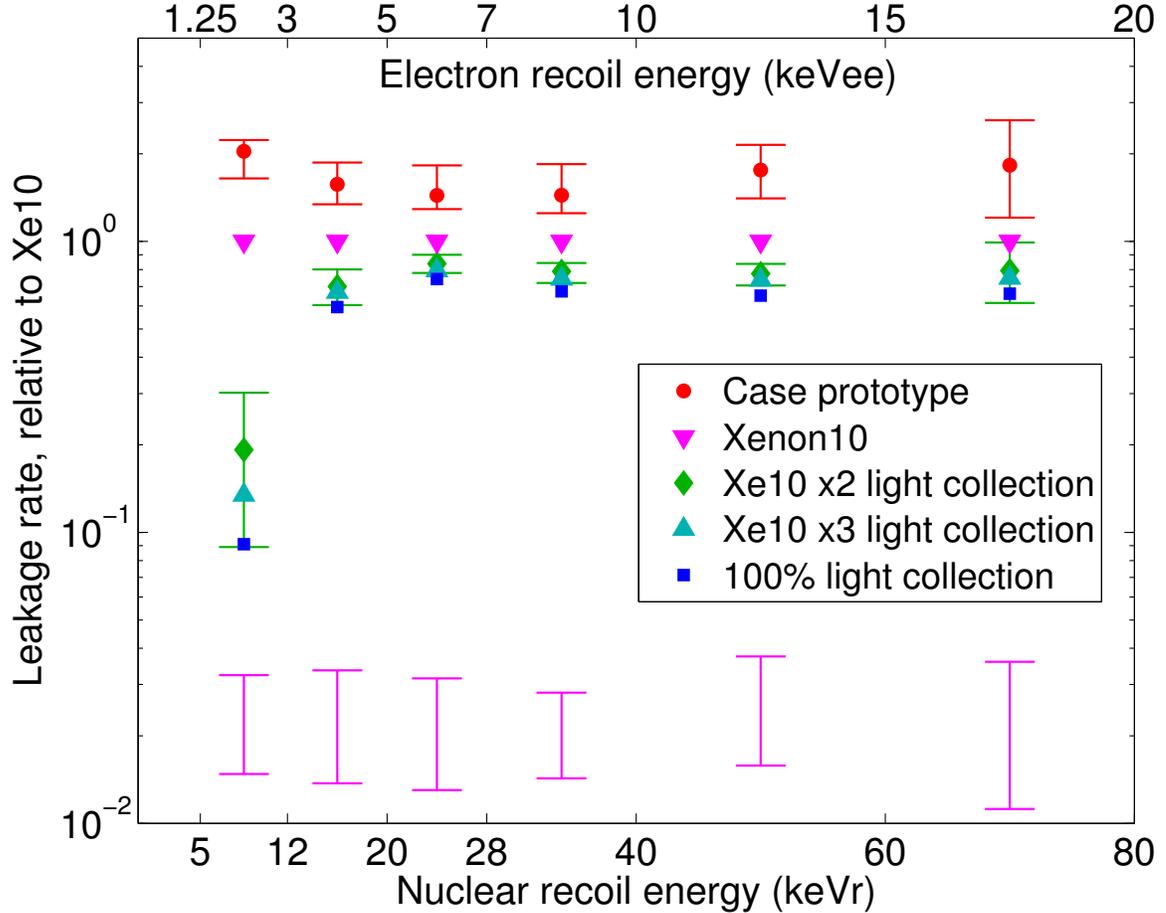


Figure 5.20: The Relative Leakage Rate at 876 V/cm in various detectors, normalized to the Xenon10 leakage rate. The datapoints correspond to those of Fig. 5.19, but the errorbars show only the error in the ratio between leakage rates. The errorbars along the bottom of the graph show the common mode error in each bin. The tick marks on the x-axis indicate bin edges. In the lowest energy bin, a x3.3–x11.0 reduction in leakage is expected between Xenon10 and a detector with twice the Xenon10 light collection.

energy electron recoil data point at 4060 V/cm appears to be an outlier at high band-width, and unpublished results from ZEPLIN III suggest the 4060 V/cm point should follow the trend of the 876 and 1951 V/cm points, giving strong discrimination. The discrimination at low fields is surprisingly good. Since high voltage requirements are one of the hurdles for very-large scale detectors, the ability to discriminate at low fields may be important in the future.

5.6 Conclusions on Recombination and Discrimination

We have now described discrimination as well as resolution in terms of the fundamental fluctuations present in liquid xenon detectors. From the gamma and neutron data taken in Xed, we have extracted

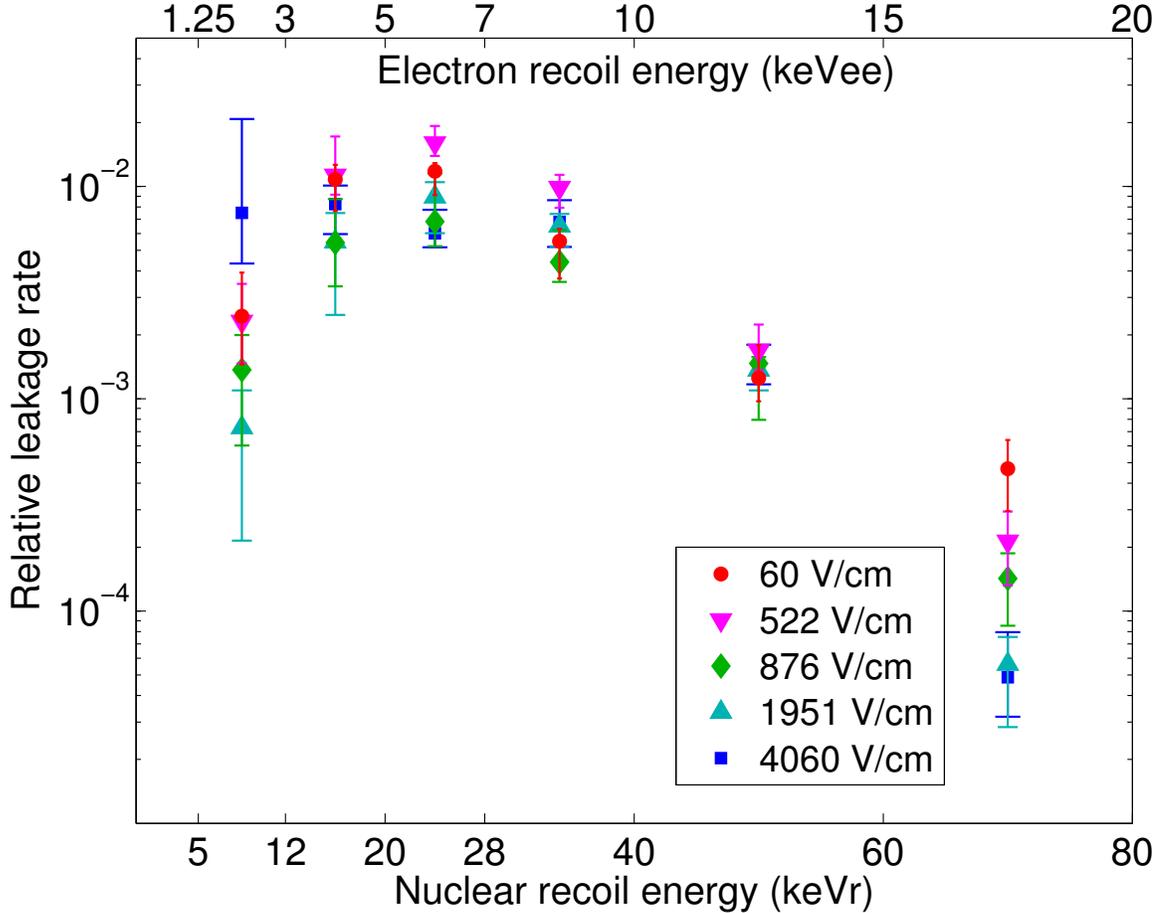


Figure 5.21: The energy spectrum of leakage events resulting from a flat electron recoil background in an ideal detector with twice the Xenon10 light collection ($g_1 = 0.174$) at various fields. Errorbars indicate statistical uncertainties, and uncertainties in projections at different fields are *not* correlated. Systematic errors are smaller than the statistical errors shown. As before, we may be overestimating leakage by ~ 1 error-bar in the lowest energy bin at intermediate fields.

the unknown recombination fluctuations by tuning a Monte Carlo so that it reproduces the observed electron and nuclear recoil band widths. We have seen that the electron recoil band, and thus discrimination, is dominated by the intrinsic fluctuations in recombination at most energies, thanks in large part to the suppression of S1 fluctuations by the slopes of the electron and nuclear recoil bands. Below 12 keVr S1 fluctuations dominate, and improved discrimination may be achieved by increasing the S1 light collection efficiency. A factor of two increase in light collection over Xenon10 will result in a x3.3–x11 decrease in the Relative Leakage Rate below 12 keVr.

We have also shown that increases in drift field appear not to significantly improve discrimination in the WIMP energy range. Both electron and nuclear recoils show very little field dependence in this range — increasing the drift field increases the charge yield in both bands, and while it slightly

increases band separation, it also increases recombination fluctuations. The increased charge yield also results in smaller S1 signals and therefore stronger S1 fluctuations.

Thus far we have considered recombination and recombination fluctuations empirically. Our vague notions that recombination increases with increasing stopping power and decreases with increasing field held up at high energies, but fall apart in the WIMP region, where field dependence for both recoil types is weak, and electron recoils actually have the wrong slope, with recombination decreasing as energy falls and stopping power increases. Recombination fluctuations in electron recoils at all energies are orders of magnitude greater than a binomial process would suggest. We turn now towards the physics behind recombination, in an attempt to explain these unexpected features.

Chapter 6

Recombination Models

In the last chapter we measured the mean and variance of the recombination distribution versus energy and drift field for electron and nuclear recoils. Of the underlying physics, we have discussed only the basic trend that recombination should increase in tracks with greater electronic stopping power (denser ionization) and decrease when a stronger electric field is applied, and we vaguely attribute recombination fluctuations to variations in track structure. In this chapter we present a simple model based on Monte Carlo track simulations that quantitatively matches the mean recombination observed in electron recoils and comes within a factor of two of the measured fluctuations (or factor of ~ 3 in the variance). Most importantly, this model reproduces the features in the data which are counter to the basic trends mentioned above. The success of the model illuminates the underlying recombination physics, paving the way for future models that may more accurately reproduce recombination fluctuations.

6.1 The Need for a Recombination Model

The physics of recombination in liquid xenon is not well understood. Existing recombination models can reproduce the field dependence of recombination in alpha and high energy electron recoils, but thus far no model has quantitatively explained the energy dependence of recombination or the magnitude of recombination fluctuations, especially at WIMP recoil energies. As noted earlier, when the XENON collaboration was formed, there was neither a physical model nor data supporting any discrimination between electron and nuclear recoils at WIMP energies, aside from the known electronic stopping powers of the two recoil types and the observed trend that, as electronic stopping power increases, so does the recombination fraction.

We saw in the Chapter 5 that, while discrimination is observed at low energies, the expected recombination trend does not hold for electron recoils in the WIMP region — the recombination fraction for electron recoils decreases as energy falls below 10 keVee, even though the electronic stopping power continues to rise. This actually saves our discrimination in the WIMP region. Early data with thresholds above 10 keVee painted a very dismal picture for background rejection, and the fortuitous turn-up in the electron recoil band was a great surprise.

Although this unknown recombination physics rescues us from the electron recoil background, it is not a perfect rescue. In Chapter 5 we determined that fluctuations in the recombination fraction are the dominant process causing leakage of electron recoils into the WIMP acceptance region. To match Xenon10 data we model these fluctuations as Gaussian in $\log\left(\frac{1-r}{r}\right)$, where r is the recombination fraction, but our statistics at the low charge-yield (high recombination) tail of this distribution are very poor, and so our true discrimination capability is unknown. For this reason Xenon10 did not use background subtraction when reporting a WIMP limit [31].

The poor understanding of recombination in liquid xenon is not important when setting a WIMP limit without background subtraction, but it becomes a major stumbling block when faced with a potential discovery. It is necessary to calibrate with an electron recoil source, gathering an order of magnitude more calibration events than WIMP background events, but this is non-trivial in a large-scale detector (as we will discuss in Chapter 7). Our understanding of different sources of fluctuations in Chapter 5 makes it feasible to perform the discrimination measurement in a smaller detector and translate the results to the detector used for WIMP searches, but such a technique will be more convincing with a physics understanding of the recombination fluctuations being measured.

Finally, there are still open design questions for very large scale liquid xenon TPC's. For example, we saw in Chapter 5 that it may not be necessary to go to high drift fields for background discrimination. If we can understand the physics behind this statement, we may confidently relax the high voltage requirements in future detectors. Similarly, understanding the tails of the recombination distribution will improve our ability to project discrimination in large detectors, which in turn sets the requirements for radio-purity and ^{85}Kr contamination. For Xenon10 these issues were not considered since only modest voltage and backgrounds were required, but for larger detectors with discovery potential, these questions become important.

We will therefore attempt to build a physical model for recombination and recombination fluctuations in liquid xenon. The ultimate goal is to quantitatively match the recombination fluctuations already observed, and then to use the model to give insight into the tails of the recombination distribution. Before that, we must reproduce and understand the weak field dependence and anomalous

energy dependence of recombination in electron recoils at WIMP energies.

6.2 Existing Recombination Models

The physics determining the recombination rate in a population of electrons and holes (or negative and positive ions) in a low-mobility medium was put forth by Langevin in 1903 [91]. Langevin considers a system where recombination times are dominated by diffusion, so that once a pair is close enough that the electrostatic potential between them overcomes diffusion, recombination happens quickly. This critical radius is given by

$$\frac{e^2}{4\pi\epsilon r_c} = kT, \quad (6.1)$$

where kT gives the thermal energy of the electron-hole pair. Outside r_c (called the Onsager or Coulomb radius), diffusion is the dominant mode of transport between the pair, but inside r_c drift due to the electric field from the charges dominates. The total recombination rate is the rate at which electrons and holes diffuse into each others' Onsager radii. This rate is proportional to the densities of the two species,

$$R = \alpha N_+ N_-, \quad (6.2)$$

where R is the recombination rate per unit volume, N_+ and N_- are the number densities of positive and negative charge carriers, and α is the recombination coefficient. This coefficient can be estimated by considering the flux of charge carriers across the surface of the sphere of radius r_c about a charge carrier of the opposite sign. This flux is determined by diffusion, but by construction transport by diffusion and drift under the electric field are balanced at r_c , so we can instead consider the drift velocity

$$v = \mu \frac{e}{4\pi\epsilon r_c^2}, \quad (6.3)$$

where μ is the combined mobility of the two charge carriers. Multiplying by the surface area and carrier densities, we get

$$R = \frac{\mu e}{\epsilon} N_+ N_-, \quad (6.4)$$

giving $\alpha = \frac{\mu e}{\epsilon} = \frac{D e^2}{kT \epsilon}$, where D is sum of the electron and hole diffusion constants, and the final equality is given by the Nernst-Einstein relation.

Jaffé first considered recombination in ion tracks in 1913 [92], writing the differential equations

$$\frac{\partial N_+}{\partial t} = D_+ \nabla^2 N_+ - \mu_+ \mathbf{E} \cdot \nabla N_+ - \alpha N_+ N_- \quad (6.5)$$

$$\frac{\partial N_-}{\partial t} = D_- \nabla^2 N_- + \mu_- \mathbf{E} \cdot \nabla N_- - \alpha N_+ N_-, \quad (6.6)$$

where D_{\pm} are the diffusion constants for the two charge carriers, μ_{\pm} are the mobilities, and \mathbf{E} is the electric field (both applied and from the other charge carriers). For each species there is a diffusion term, a drift term, and Langevin's recombination term. Since writing these equations, Jaffé and others have found solutions with various approximations and initial conditions. Of particular interest to us is the work by Kramers (1952) [93] and Thomas and Imel (1987) [94]. Thomas and Imel specifically consider liquid argon and xenon, simplifying the above equations to

$$\frac{\partial N_+}{\partial t} = -\alpha N_+ N_- \quad (6.7)$$

$$\frac{\partial N_-}{\partial t} = -v \frac{\partial N_-}{\partial z} - \alpha N_+ N_-, \quad (6.8)$$

where we have taken a constant electric field in the $-z$ direction so that $v = \mu_- \mathbf{E}$ (i.e., we ignore the fields from the charges themselves). The charge carriers in liquid xenon are positive ions and electrons, with ions having much lower mobility and diffusion constants than electrons, so the positive charge distribution is treated as stationary. From here on, we will drop the subscripts in μ and D , always meaning the electron mobility and diffusion constant. Thomas and Imel also ignore electron diffusion, which we will discuss in Section 6.5, leaving a negative charge distribution drifting with constant velocity past a fixed positive charge distribution.

Equations (6.7,6.8) are exactly solvable for the initial condition $N_-(t=0) = N_+(t=0) = N_0$, as done by Kramers [93]. We substitute Eq. (6.7) into the right hand side of Eq. (6.8), giving

$$\frac{\partial N_-}{\partial t} = \frac{\partial}{\partial t} \left[\frac{v}{\alpha} \frac{\partial}{\partial z} [\ln N_+] + N_+ \right]. \quad (6.9)$$

Integrating from $t = 0$ and substituting for N_- on the left hand side gives a first order equation for N_+ ,

$$\frac{\partial}{\partial t} [\ln N_+] = -v \frac{\partial}{\partial z} \left[\ln \frac{N_+}{N_0} \right] - \alpha N_+. \quad (6.10)$$

To solve this, we define $Y = \frac{N_0}{N_+}$ and change variables to $u = t + \frac{z}{v}$ and $w = t - \frac{z}{v}$, simplifying

Eq. (6.10) to

$$\frac{\partial Y}{\partial u} = \frac{\alpha}{2} N_0, \quad (6.11)$$

with the initial condition $Y(u+w=0) = 1$. Integrating from $u = -w$ gives

$$Y(u, w, x, y) = 1 + \frac{\alpha}{2} \int_{-w}^u du' N_0 \left(x, y, \frac{v}{2} (u' - w) \right), \quad (6.12)$$

$$Y(t, x, y, z) = 1 + \frac{\alpha}{v} \int_{z-vt}^z dz' N_0(x, y, z'), \quad (6.13)$$

and finally,

$$N_+(t, x, y, z) = \frac{N_0(x, y, z)}{1 + \frac{\alpha}{v} \int_{z-vt}^z dz' N_0(x, y, z')}. \quad (6.14)$$

The positive charge density at a point decreases over time as negative charge drifts up from below and recombines. The total charge escaping recombination is given by taking $t \rightarrow \infty$ and integrating over all space,

$$n_e = \int d^3x \frac{N_0(x, y, z)}{1 + \frac{\alpha}{v} \int_{-\infty}^z dz' N_0(x, y, z')}. \quad (6.15)$$

Kramers and Thomas-Imel differ when choosing the initial condition N_0 . Kramers follows Jaffé and considers an infinite column with Gaussian cross section, as in the track of a relativistic heavy charged particle. Thomas and Imel are interested in smaller tracks, taking

$$N_0 = \begin{cases} \frac{N_i}{8a^3} & |x|, |y|, |z| < a \\ 0 & \text{otherwise} \end{cases}, \quad (6.16)$$

or N_i ions spread uniformly in a 3-D box with sides of length $2a$. Taking Thomas and Imel's N_0 and performing the integrals in Eq. (6.15) gives

$$\frac{n_e}{N_i} = \frac{1}{\xi} \ln(1 + \xi), \quad \text{where} \quad (6.17)$$

$$\xi = \frac{N_i \alpha}{4a^2 v}. \quad (6.18)$$

Thomas and Imel show that this functional form fits the observed field dependence of alphas and 368 keV ^{113}Sn internal conversion electrons in liquid argon and xenon when they take $\xi \propto \mathbf{E}^{-1}$. They interpret this as a drift velocity effect ($v \propto \mathbf{E}$), with a and α field independent. This is curious, since at the fields considered (up to 2 kV/cm in xenon and 30 kV/cm in argon) electrons in liquid argon/xenon do not have a constant mobility. Indeed, electron drift velocity changes very slowly with the applied field [46]. We will return to this point in Section 6.5 and give a more consistent

interpretation of the Thomas-Imel result.

The recombination models following Jaffé’s differential equations consider ‘columnar’ or ‘cluster’ recombination, i.e., recombination between a population of positive and negative charges. Other models such as that by Onsager [95] consider recombination between an electron and the ion from which it originated. We don’t expect this ‘initial recombination’ to be dominant in liquid xenon, since if it were, we would see no correlation between ionization density and recombination fraction, and we would also see only binomial recombination fluctuations (since the recombination of every ion pair would be an independent event) — both contrary to observations. In fact, an ‘initial recombination’ fraction would be indistinguishable from a shift in $\frac{N_{ee}}{N_i}$, except that the former may have some field dependence. Initial recombination also depends strongly on the initial energy given to ionized electrons, as it requires the electrons to thermalize before leaving the Onsager radius. We will see that there is a possibility for initial recombination in nuclear recoil tracks, but defer to Onsager’s paper for a complete discussion.

For completeness, there are also models that describe recombination as an interaction between a cylindrical core of positive ions and classically orbiting electrons, with a small probability for recombination each time an electron passes through the track core. These models do not apply to the small, scattered tracks we see at WIMP recoil energies, but interested readers may find them in [84, 96].

6.3 Construction of a Recombination Model

Where the existing recombination models focus on the field dependence of recombination, we wish to focus on the energy dependence and on recombination fluctuations — that is, we are interested primarily in changes in recombination at a fixed drift field. Recombination takes place on time scales of 10s to 100s of nanoseconds [44], many orders of magnitude longer than it takes the initial track of ions to form, so all information that can influence recombination is encoded in the positions of the ions and electrons after the recoil energy has been spent. We therefore divide the recombination modeling problem into two pieces. First we develop Monte Carlo methods to generate recoil tracks in liquid xenon, creating spatial maps of the resulting electron and ion distributions. We then construct a method for determining the expected recombination fraction for a given ionization map.

6.3.1 Electron recoil track Monte Carlo

We create electron recoil tracks using Penelope [97], a Monte-Carlo program for electro-magnetic interactions based on a model of generalized oscillator strengths. Penelope was designed for applications in medical physics and specializes in low energy interactions, where it is more accurate than Geant4 and EGS [98]. Penelope follows cascades of electrons down to 100 eV, reporting all energy loss along the cascade. For each simulated recoil, we translate the Penelope output into a spatial map of ions by placing an ion at every location where a new cascade electron is created. We also place one or more ions at points where energy loss below the 100 eV tracking threshold is reported, with the number of ions placed equal to the energy lost divided by W_q , rounding down, but always placing at least one ion (even if the energy loss is less than W_q). This last requirement is reasonable because Penelope records energy losses only for inelastic scatters, with a minimum energy loss around the ionization energy of the atom. Elastic scatters are simulated in Penelope as ‘random hinges’ with zero energy loss.

Figure 6.1 shows several sample tracks from Penelope. We see that there are occasional hard scatters, with the incoming electron giving several keV or more to another electron. It is this division of energy among daughter recoils that we expect to give rise to recombination fluctuations. Since stopping power increases with every energy loss, these branches result in two dense tracks where we would otherwise see one sparse track. For example, a 10 keV electron that loses its energy without any hard scatters samples the stopping power from 0–10 keV, while the same recoil that immediately loses half its energy to another electron samples the stopping power from 0–5 keV twice, resulting in a higher average ionization density. Even daughter recoils with a few hundred eV result in a dense spot of ionization along the track. This is reflected in the ion positions we take from Penelope, which when fed through our recombination model should result in higher recombination in tracks with many hard scatters.

We also see that electron recoil tracks are not at all straight lines. When interpreting the parameters of our recombination model below, it will be useful to know the spatial extent of these tracks as a function of energy. As a measure of track size we take the rms radius of the track, $R_{\text{rms}} = \sqrt{\langle |\mathbf{x} - \langle \mathbf{x} \rangle|^2 \rangle}$, where the brackets indicate the mean over the ions in the track. Figure 6.2 shows the mean and variation in R_{rms} as a function of electron recoil energy.

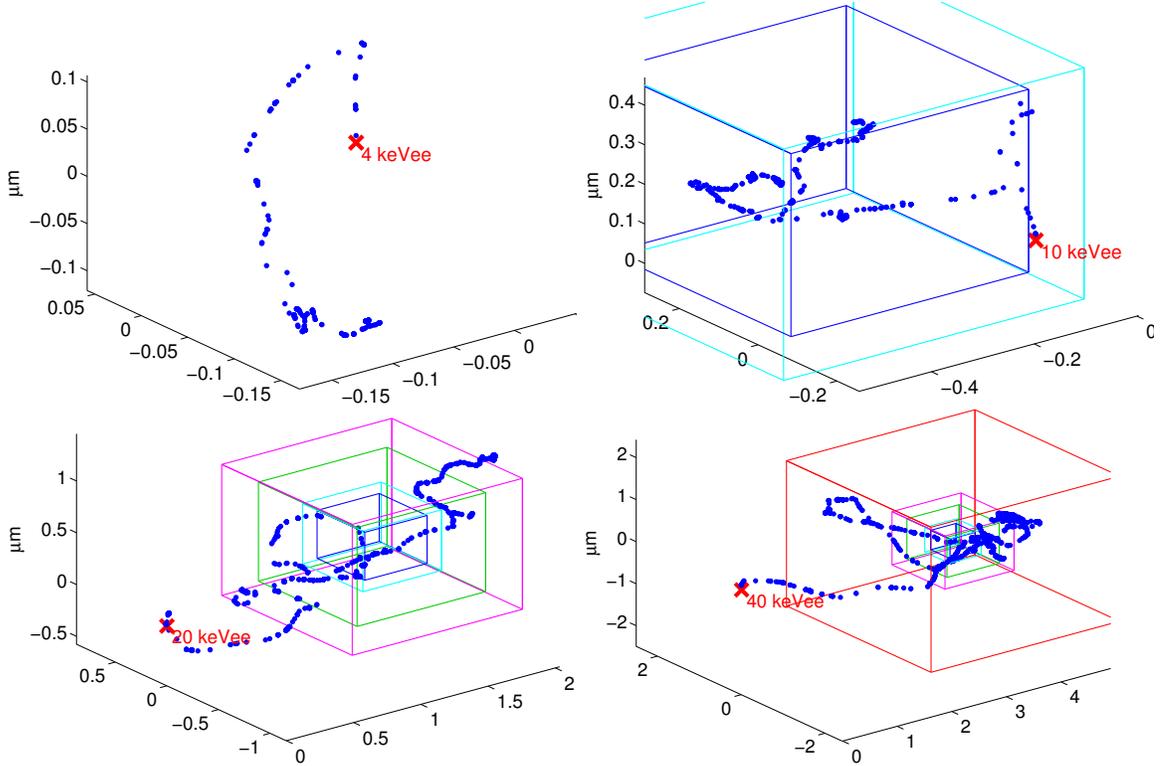


Figure 6.1: Example electron recoil tracks at 4, 10, 20, and 40 keVee, simulated using Penelope. Each blue dot is one xenon ion, and the red X marks the location of the initial recoil. The boxes correspond to the box sizes (a) in our recombination model, as listed in Table 6.2.

6.3.2 Nuclear recoil track Monte Carlo

For nuclear recoils, the Monte Carlo program SRIM/TRIM [88] creates cascades of recoils using the universal nuclear scattering cross sections given by Ziegler [99]. Although this program does keep track of energy lost via electronic channels, it is difficult to extract this information from the output available to the user. We therefore create a Monte Carlo to suit our purposes, dubbed RIVAL (Recoiling Ions in Various Atomic Liquids). This program also generates nuclear recoil tracks, following the primary and daughter recoils down to a set energy threshold and keeping track of energy lost via low-energy nuclear collisions and electronic stopping along the way. It turns out that the details of the nuclear recoil simulation are unimportant in our recombination model, but see Appendix C for a complete description of the Monte Carlo.

Sample nuclear recoil tracks are shown in Fig. 6.3. As with electron recoils, hard scatters give the track a tree-like structure. Since stopping power in nuclear recoils decreases as energy is lost, such branches have the opposite effect on ionization density as in electron recoils — every branch ends in a sparse scattering of ions. Figure 6.2 shows the rms radii of nuclear recoil tracks as a function

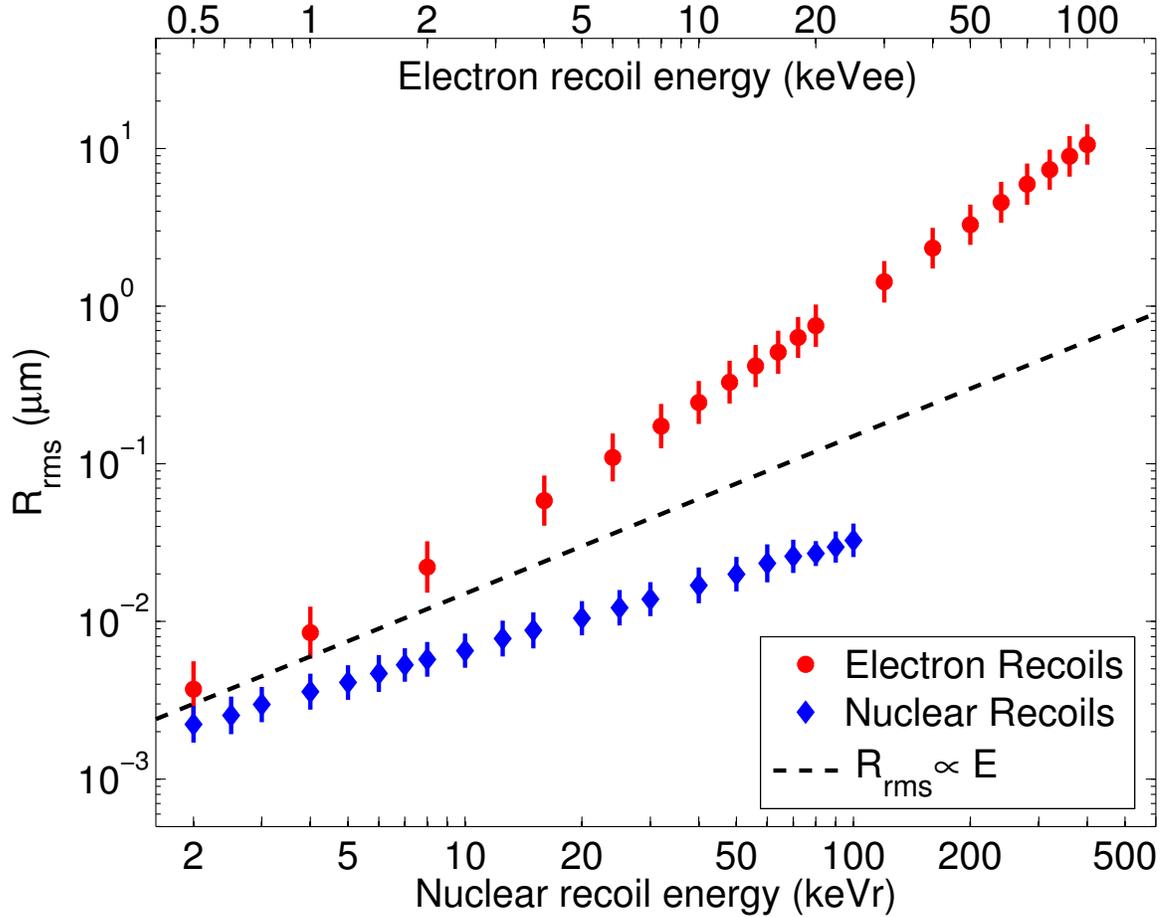


Figure 6.2: Track size as a function of energy for electron and nuclear recoils. We use the rms radius of the track as a measure of track size. The points show energies at which tracks were simulated in Penelope (electron recoils, red \circ) or RIVAL (nuclear recoils, blue \diamond). Error-bars indicate the standard deviation in track size at each energy. The black dashed line has a slope of 1 (track size proportional to energy) — for electron recoils, stopping power falls as energy increases, so track size grows faster than energy. For nuclear recoils, stopping power increases with energy, so track size rises more slowly.

of energy. Note that our biggest nuclear recoils (100 keVr) are only as big as a ~ 3 keVee electron recoil. This is well within in the range where we see weak field dependence and anomalous energy dependence for recombination in electron recoils.

6.3.3 Modified Thomas-Imel Model

We now need a prescription to go from a spatial map of ions to a recombination fraction. We begin with the same differential equations as Thomas-Imel model, with the solution given by Eq. (6.15).

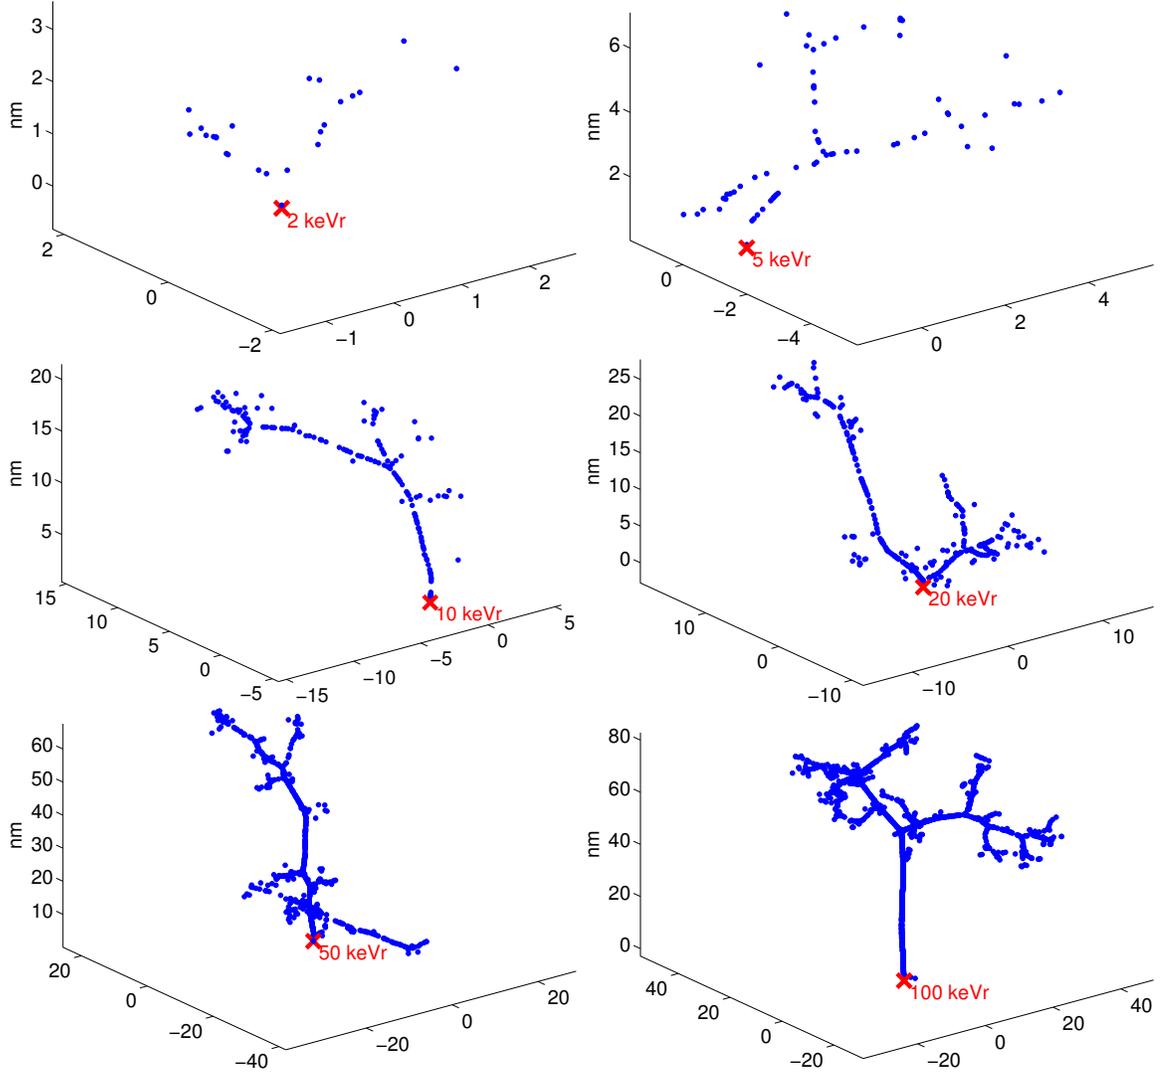


Figure 6.3: Example nuclear recoil tracks at 2, 5, 10, 20, 50, and 100 keVr, simulated using RIVAL. Each blue dot is one xenon ion, and the red X marks the location of the initial recoil. All of these tracks are much smaller than the box sizes found in Section 6.4.

We then take the initial condition

$$N_0 = \sum_{k=1}^{N_i} \begin{cases} \frac{1}{8a^3} & |x - x_k|, |y - y_k|, |z - z_k| < a \\ 0 & \text{otherwise} \end{cases}, \quad (6.19)$$

where $\{x_k, y_k, z_k\}$ are the locations of the ions in the track as given by the recoil simulations. Note that in the original Thomas-Imel condition a determines the total size of the track, but for us the track geometry is given in the spatial distribution of ions and a is the smearing applied at each ion. This construction is the simplest way to include the Monte Carlo tracks in the recombination

calculation, and we assign no a priori physical interpretation to the box size a . If the size of the track is much smaller than the box size, our initial condition becomes the same as the Thomas-Imel condition, and the solution is given by Eq. (6.17). This model is entirely deterministic, so that recombination fluctuations must come from the variations in structure output by the track Monte Carlos.

6.4 Comparing the Model to the Data

The integral in equation (6.14) may be exactly calculated for the initial condition given in Eq. (6.19). In practice we snap the ions to a grid with spacing of $0.1a$ for faster computation (this does not significantly affect the result). The drift velocity and recombination coefficient enter the solution together, so the two free parameters in the model are $\frac{\alpha}{v}$ and the box size a . We simulate electron recoils at energies from 0.5–100 keVee and nuclear recoils from 0.5–100 keVr and find the resulting recombination over a range of $\frac{\alpha}{v}$ and a . We find that, at each field, we can match the observed mean electron recoil recombination versus energy with fixed values of these two parameters. With one adjustment we extend the model to nuclear recoils as well. We also look at the recombination fluctuations produced by the model, finding that while they have roughly a third the measured variance, they do reproduce some of the features in fluctuation strength seen in the data.

6.4.1 Electron Recoils

The first hurdle for our model is to reproduce the energy dependence of recombination in electron recoils. The model naturally produces a maximum recombination fraction when overall track size R_{rms} is close to the box size a . Any track structure at length scales smaller than a is washed out by the box-smearing, and for tracks much smaller than a the model reverts to the original Thomas-Imel model with the recombination fraction determined entirely by the ratio $\frac{N_i \alpha}{a^2 v}$, completely independent of the stopping power. Thus, below a threshold set by the box-size, recombination falls and charge yield rises as energy and N_i decrease.

We fit the electron recoil centroid at each field with a fixed a and $\frac{\alpha}{v}$, minimizing χ^2 with respect to the measured recombination up to 60 keVee. Figure 6.4 shows the resulting fits, with fit parameters listed in Table 6.1. The fit box-size is determined by the location of the charge-yield minimum, while $\frac{\alpha}{v}$ is varied to match the measured recombination fraction. That we also fit the slope of the electron recoil band above and below the minimum charge yield is a clear success for the model. The deviation between model and data above ~ 80 keVee is likely due to photo-absorption events in

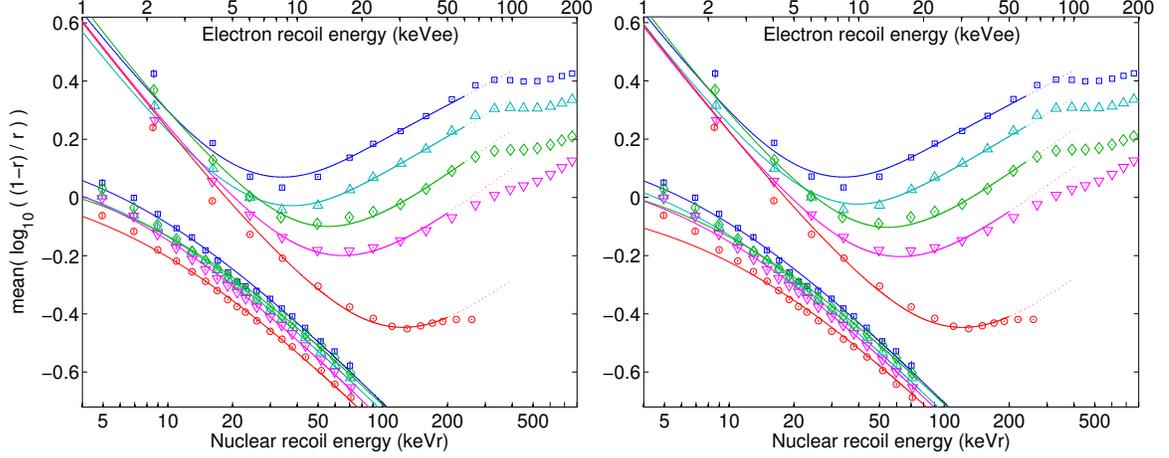


Figure 6.4: Fits of the modified box model to the electron and nuclear recoil band centroids. Data points show the recombination fractions measured in Chapter 5, and lines show the fits to our recombination model. In the left plot electron and nuclear recoil bands are fit independently, while on the right they share the same $\frac{\alpha}{a^2v}$ at each field. The dashed section of the electron recoil line indicates the region not used in the fit. The discrepancy between model and data in this region may be due to the presence of photo-absorption events in the data. Fields shown are 4060 (blue \square), 1951 (cyan \triangle), 876 (green \diamond), 522 (magenta ∇), and 60 V/cm (red \circ), with electron recoils above and nuclear recoils below. The y-axis takes r calculated using $\frac{N_{ex}}{N_i} = 0.06$, the expected value for electron recoils. We allow $\frac{N_{ex}}{N_i}$ to vary with field for nuclear recoils. Fit values are given in Table 6.1.

Table 6.1: Box model parameters giving the best fit to the electron recoil centroid, nuclear recoil centroid, and both centroids at each field. Fits are shown in Fig. 6.4. Each band at each field has two free parameters — for electron recoils, these are the two box model parameters, while for nuclear recoils they are a single box model parameter and the ratio of direct-excitons to ions (before recombination). We also perform a simultaneous fit to both bands, using the same box model parameters for electron and nuclear recoils. The exciton-ion ratio for electron recoils is always held at $\frac{N_{ex}}{N_i} = 0.06$. Errors shown are statistical. Systematic errors are larger, and may be estimated from the spread between the different fitting schemes (see also Table 6.2).

Drift Field E (V/cm)	Electron Recoils		Nuclear Recoils		Both Recoils		
	a (nm)	$100\frac{\alpha}{a^2v}$	$100\frac{\alpha}{a^2v}$	$\frac{N_{ex}}{N_i}$	a (nm)	$100\frac{\alpha}{a^2v}$	$\frac{N_{ex}}{N_i}$
4060 ± 190	201(3)	3.17(5)	3.19(5)	0.711(17)	200(3)	3.18(4)	0.716(13)
1951 ± 86	244(4)	3.71(5)	3.37(3)	0.817(12)	263(3)	3.48(3)	0.781(10)
876 ± 36	484(6)	2.96(3)	3.08(3)	0.865(11)	473(4)	3.03(2)	0.883(9)
522 ± 23	648(8)	3.35(3)	3.76(3)	0.786(10)	604(5)	3.54(2)	0.846(8)
60 ± 5	2111(18)	3.39(2)	4.02(5)	0.945(15)	2032(17)	3.48(2)	1.111(8)

the data, discussed in detail in Section 6.4.5. Finally, the weak field dependence at low energy we now see as weak field dependence in the ratio $\frac{\alpha}{a^2v}$. Field dependence will be the key to the physics interpretation of the model in Section 6.5.

6.4.2 Nuclear Recoils

We see an immediate problem when applying the model to nuclear recoils. The key to reproducing the electron recoil band is the smearing that washes out structures smaller than the box size, but all of our nuclear recoil tracks are also smaller than the box-size. Thus, at energies below the minimum electron charge yield, electron and nuclear recoils are indistinguishable in our model. (This is also why the details of our nuclear recoil track Monte Carlo were unimportant.) Even allowing our fit parameters to vary between recoil types, we cannot match the nuclear band data. For small tracks our only free parameter is $\frac{\alpha}{a^2v}$, which adjusts the intercept of a line in $\log\left(\frac{1-r}{r}\right)$ vs $\log(E)$ (as in low-energy electron recoils), but the slope is fixed by the model. This slope does not match that of the nuclear recoil band.

We can fit the nuclear recoil band by introducing the initial exciton-ion ratio $\frac{N_{ex}}{N_i}$ as a free parameter. The physical motivation for such a parameter may be either initial recombination of electrons that never leave the Onsager radius, as in [95], or an actual shift in the direct exciton to ion ratio in the recoil track. Certainly, the mechanisms for electronic energy loss in nuclear and electron recoils are very different. Ionized electrons in nuclear recoil tracks may have very low energies, and therefore be subject to initial recombination. In [100], ionization in nuclear recoils is attributed to electron promotion when energy levels shift as the recoiling nucleus (ion) crosses the target atom, followed by Auger electron emission. This is effectively an adiabatic process, with atom velocities much lower than velocities of the orbiting electrons. Since only a subset of atomic states shift in a manner that will cause ionization, this may reduce the number of ionization channels available to nuclear recoils, giving a larger fraction of direct excitons. Whatever the underlying mechanism, this parameter allows us to fit the nuclear recoil bands, as shown in Fig. 6.4, with fit parameters in Table 6.1.

6.4.3 Combined Fits

The electron and nuclear recoil fits have one parameter in common, $\frac{\alpha}{a^2v}$, the fit values of which are very similar for the two bands. We therefore perform simultaneous fits to both bands using three parameters (a , $\frac{\alpha}{a^2v}$, and $\frac{N_{ex}}{N_i}$), with results very similar to those found when fitting the bands separately (see Fig. 6.4 and Table 6.1). Going one step further, we fit all bands at all fields, allowing a and $\frac{\alpha}{a^2v}$ to vary with field but taking a single value for $\frac{N_{ex}}{N_i}$ for nuclear recoils (see Fig. 6.5 and Table 6.2). This shifts the fit box model parameters significantly at low and high fields and by eye gives cleaner trends with drift field (the fits at different fields no longer fall on top of each other

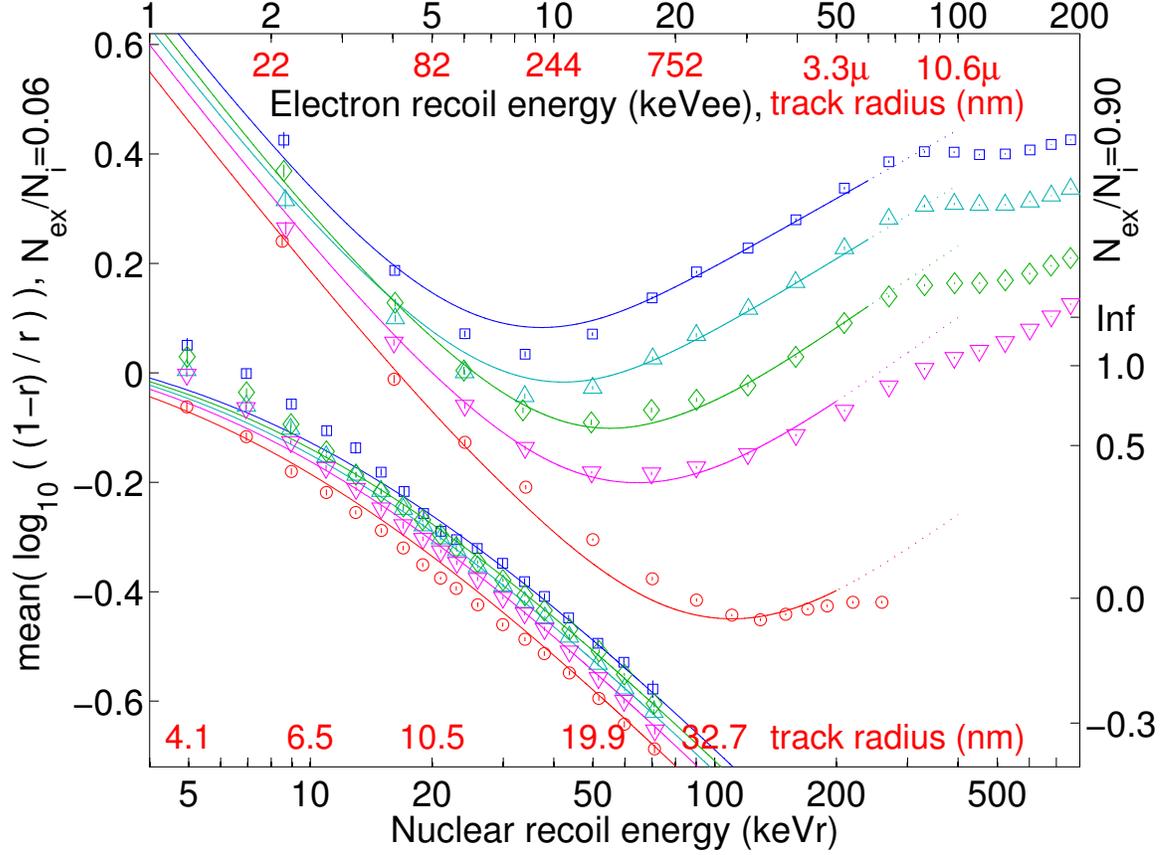


Figure 6.5: Fits of the modified box model to the electron and nuclear recoil band centroids, as in Fig. 6.4, but taking a field independent value for $\frac{N_{ex}}{N_i}$ for nuclear recoils. These fits show smoother field dependence and better match the electron recoil bands at low energies, but do not fit nuclear recoils below 10 keVr (2.5 keVee) as well as the previous fits. The right y-axis shows $\log_{10}(\frac{1-r}{r})$ calculated using the fit $\frac{N_{ex}}{N_i}$ for nuclear recoils, while the left uses $\frac{N_{ex}}{N_i} = 0.06$. Fit parameters are given in Table 6.2.

at low energies), but does not match the nuclear recoil data at low energies as well as the previous fits. For all fits, we keep $\frac{N_{ex}}{N_i} = 0.06$ for electron recoils — allowing this to vary as well does not significantly change the fit.

6.4.4 Compound Recoils

We now apply our model with the fit parameters from Table 6.2 (combined electron and nuclear recoil fits with field-independent $\frac{N_{ex}}{N_i}$) to the inelastic scatter and photo-absorption peaks. As described earlier, these peaks have two interaction sites (primary scatter and x-ray absorption), each with multiple recoils. We find the recombination at each site by concatenating the ion positions of all recoils originating at that site and sending the total ion list to the recombination model. Summing

Table 6.2: Box model parameters giving the best fit to both centroids at all fields, allowing a and $\frac{\alpha}{a^2v}$ to vary with field but keeping $\frac{N_{ex}}{N_i}$ for nuclear recoils fixed. For electron recoils we take $\frac{N_{ex}}{N_i}=0.06$. Errors shown are statistical. Systematic errors are larger, and may be estimated from the spread between these values and those in Table 6.1. The fits with these parameters are shown in Fig. 6.5.

Drift Field E (V/cm)	All Recoils		
	a (nm)	$100\frac{\alpha}{a^2v}$	$\frac{N_{ex}}{N_i}$
4060 ± 190	231(2)	2.80(2)	
1951 ± 86	293(2)	3.19(2)	\vdots
876 ± 36	478(3)	2.99(1)	0.899(4)
522 ± 23	632(5)	3.42(2)	\vdots
60 ± 5	1726(13)	3.86(2)	

the signals from the two interaction sites, we find the locations of the peaks, shown in Figs. 6.6 and 6.7. When reconstructing the 40 keV inelastic scatter, we include a nuclear recoil piece from 0–20 keVr. Note that, although the 122 keV reconstruction includes a recoil at 87 keVee where the model does not fit the band, the peak reconstruction is still successful. This supports the hypothesis that the electron recoil band is contaminated by photo-absorption events at those energies.

6.4.5 Fluctuations

Figures 6.6 and 6.7 also show the fluctuations in recombination given by the model for the 40 keV and 122 keV peaks. Our modeled fluctuations consistently have about one third the measured variance, extracted from the 2-D Gaussian fits to the peaks. The fluctuations predicted for the electron recoil band are shown in Fig. 6.8. Immediately we see a problem at low energies — since the model is insensitive to structure at scales smaller than the box size, small tracks show zero recombination fluctuations. This is perhaps accurate for nuclear recoils, where the observed fluctuations are near those expected from binomial statistics (which are not included in the model), but it is not consistent with electron recoils. At higher energies we do produce fluctuations, but again they have roughly one third the measured variance.

The sharp drop in recombination fluctuations at all fields around 100 keV is our strongest indication that the electron recoil band is contaminated by photo-absorptions at high energies. Because photo-absorption events have significant structure imposed on them initially (the full energy being divided into two Auger electrons, an x-ray photo-electron, and the initial photo-electron), there is

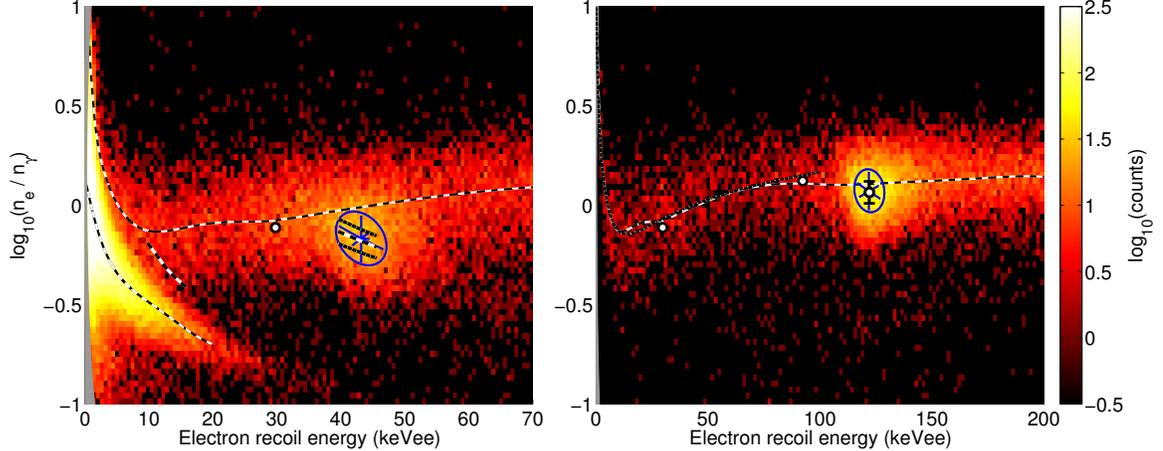


Figure 6.6: Reconstruction of the 40 keV and 122 keV peaks at 876 V/cm using our recombination model with the fit parameters in Table 6.2. Left: The 40 keV peak is shown with the 2-D Gaussian fit from Chapter 5 and decomposition into constituent fluctuations drawn in blue. The vertical component of the blue cross indicates the recombination fluctuation component. Electron and nuclear recoil centroids are drawn as dashed lines, with the dashed line between the centroids from 10–15 keVee showing the predicted recombination at the primary scatter (0–22 keVr nuclear recoil, 5 keVee internal conversion electron, and 5 keVee Auger electron). The dot at 30 keVee shows the predicted recombination at the x-ray absorption site (25 keVee photo-electron and 5 keVee Auger electron). The dashed line from 40–45 keVee is the sum of these two components. The dotted lines above and below the sum show the 1σ variation in recombination predicted by the model, which in this case is roughly half of the observed fluctuations (one fourth the variance). Predictions at other fields (see Fig. 6.7) also give from one half to one fourth of the observed variance. Right: The 122 keV photo-absorption peak, with the Gaussian fit and decomposition in blue. The dashed line is the electron recoil centroid, and the dotted line the recombination predicted for Compton scatters. The discrepancy between the two above ~ 80 keVee may be due to photo-absorption events in the data. The dots at 92 and 30 keVee show the predicted recombination at the primary absorption and x-ray absorption sites, respectively, with the dot at 122 keVee showing the sum. The black error-bars on the 122 keVee point indicate the predicted recombination fluctuations, which have one third of the observed variance.

much less phase space for fluctuations than for a Compton recoil with the same energy. This is reflected both in the model and the data when looking at the 122 keV peak — both show roughly half the variance for the peak that they do for the Compton band below 100 keV. That the model and the data show the same relative decrease in fluctuations for an effect that we know is due to track structure indicates that our problem is not a failure to create structure variations, but rather a failure to propagate them fully through the recombination model.

Although we do not quantitatively reproduce the observed recombination fluctuations, it is worth looking at the shape of the modeled recombination distribution. We cannot do this in the WIMP region since we produce no recombination fluctuations there, but Fig. 6.9 shows the recombination distribution at 80 keVee. We have only simulated 1,000 events at each energy and so do not have large statistics, but we see a definite tail at low charge-yield. We cannot draw any conclusions about

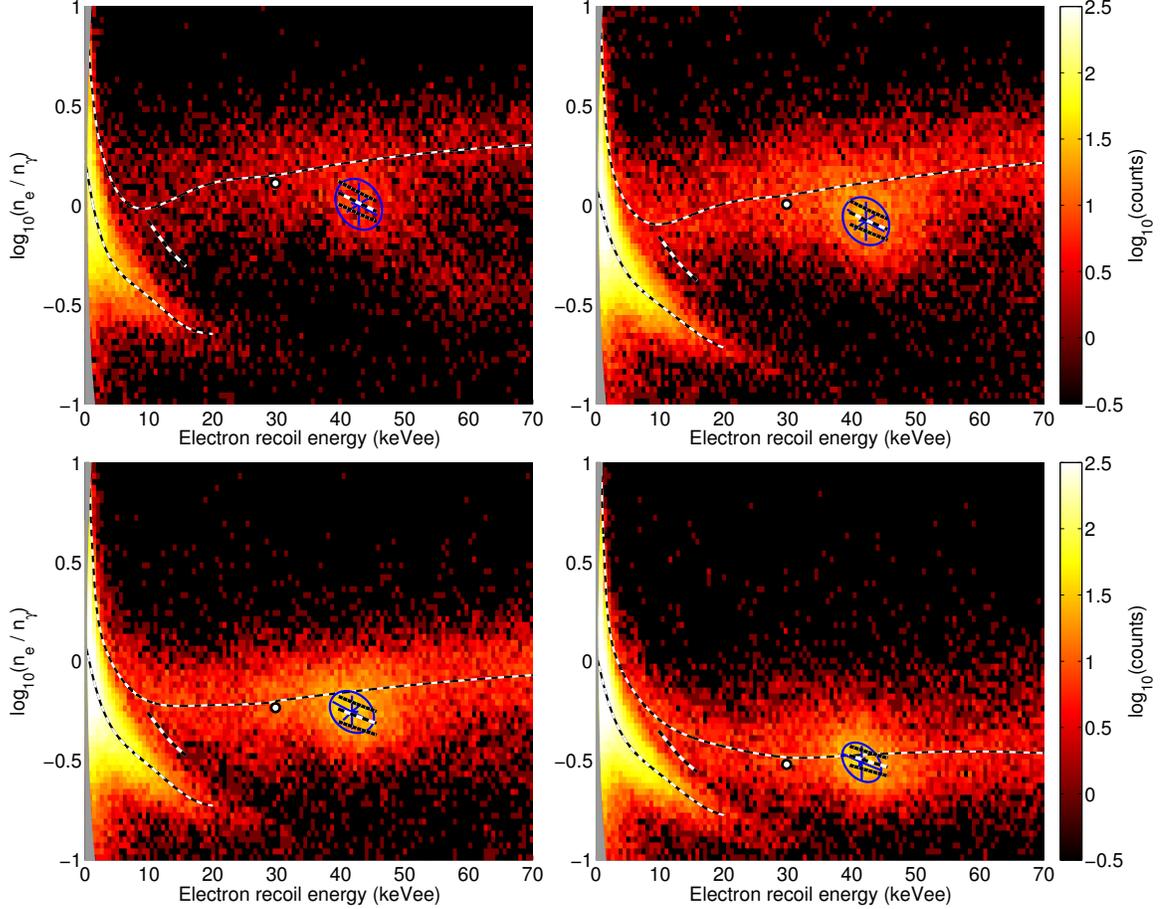


Figure 6.7: Reconstruction of the 40 keV inelastic recoil at 4060 (top-left), 1951 (top-right), 522 (bottom-left), and 60 V/cm (bottom-right). See the caption of Fig. 6.6 for an explanation of the plots.

the shape of the recombination distribution at WIMP energies, but this emphasizes the need for a quantitative understanding of these fluctuations.

6.5 Interpretation of Fit Parameters

We have seen that this model reproduces the shape of the electron recoil centroid, but we do not yet have a physical interpretation of the box size. The key to this lies in the field dependence of the model parameters. Figure 6.10 shows the box size following a power law relation $a \propto E^{-\gamma}$, where E is the magnitude of the applied field, with $\gamma=0.48-0.56$ depending on the set of fits used. This scales similarly to the radius R_F within which the field from the ion is stronger than the applied

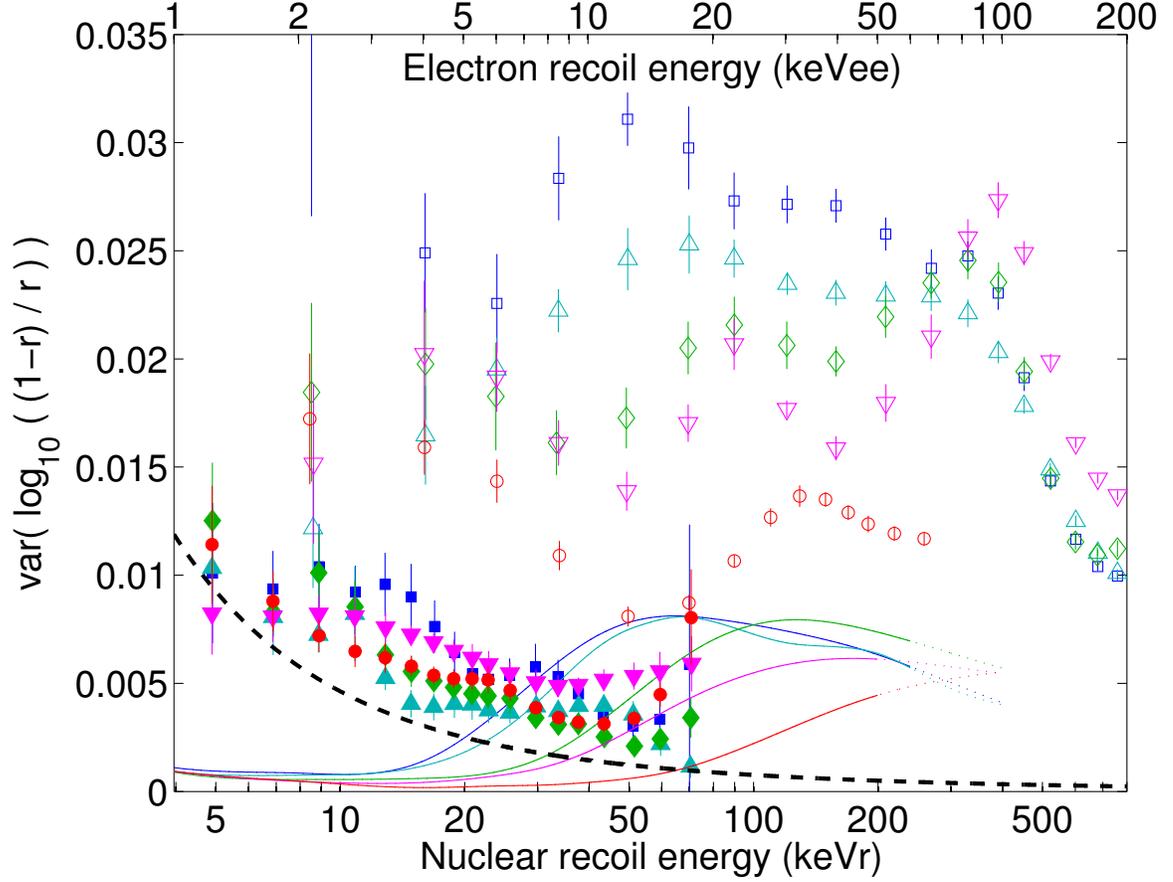


Figure 6.8: Model predictions for recombination fluctuations at all fields. Data points show the recombination fluctuations measured in Chapter 5, and lines show the fluctuations in electron recoils given by our recombination model, using the electron recoil fit parameters in Table 6.1. The model produces no recombination fluctuations in nuclear recoils. Fields shown are 4060 (blue \square), 1951 (cyan \triangle), 876 (green \diamond), 522 (magenta ∇), and 60 V/cm (red \circ), with empty points for electron recoils and solid for nuclear recoils. The dashed black line shows the fluctuations expected for a binomial process (slightly dependent on the mean recombination — here we have taken the nuclear recoil centroid at 876 V/cm). The drop in measured recombination fluctuations at 100 keVee is likely due to photo-absorption events in the data. The 122 keV photo-peak has a measured recombination variance of ~ 0.011 in these units.

drift field,

$$R_F = \sqrt{\frac{e}{4\pi\epsilon\mathbf{E}}} = 86\text{nm}\sqrt{\frac{\text{kV/cm}}{\mathbf{E}}}, \quad (6.20)$$

where we have taken $\epsilon=1.96\epsilon_0$. The fit values for a are roughly 5 times the values for R_F . Similar interpretations, such as the impact parameter at which field lines converge on the ion ($2R_F$), may come closer. Regardless of the exact interpretation, order of magnitude and scaling with drift field are consistent with a being an electrostatic length scale.

This forces a reinterpretation of v . In the model, the role of v is to determine the amount of time

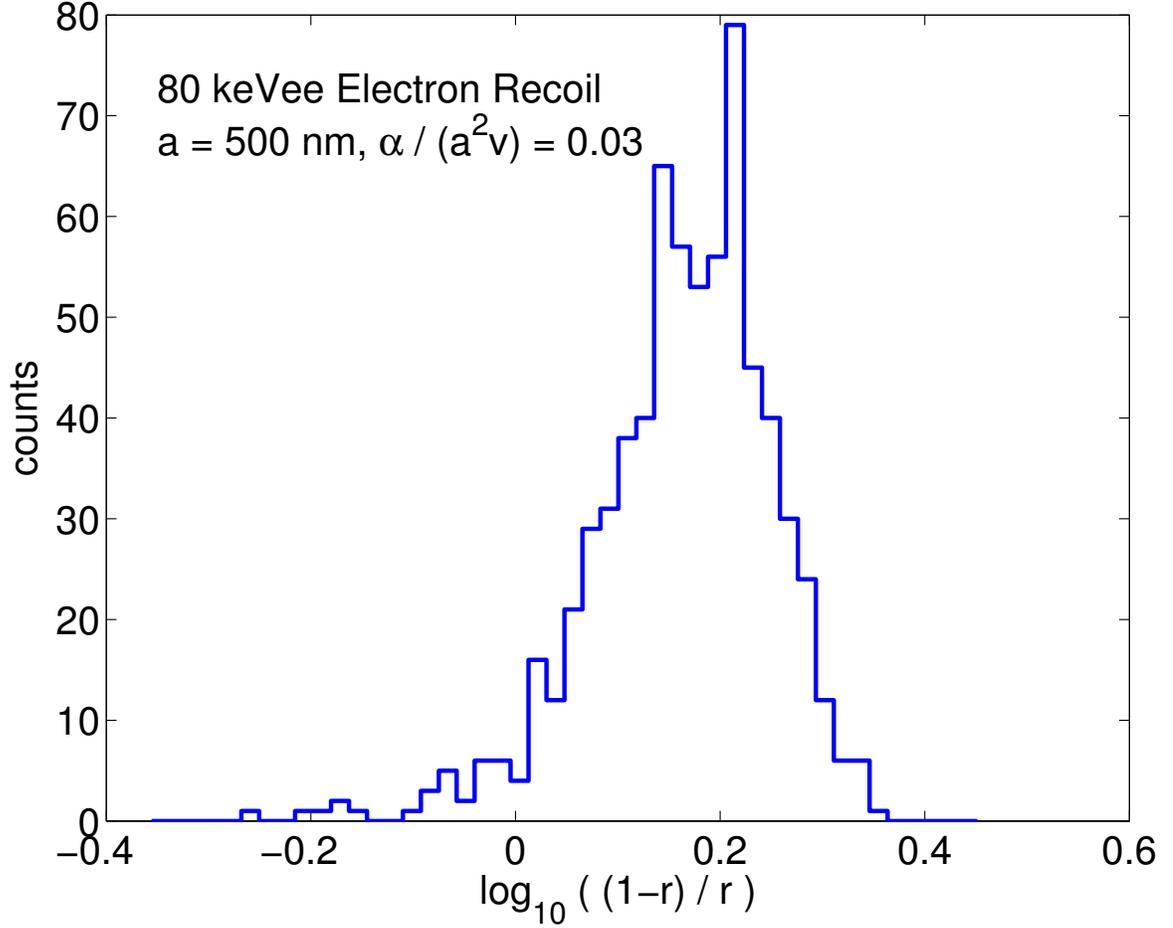


Figure 6.9: The modeled recombination distribution at 80 keVee, taking $a = 500$ nm, corresponding to a drift field of ~ 800 V/cm. There is a clear tail of low charge-yield events.

a passing electron spends inside a given ion's box ($t = \frac{2a}{v}$) during which it may recombine. (The drift velocity also moves electrons to the box, of course, but that time scale does not enter in the recombination fraction.) If the field inside the box is dominated by the field from the ion, then v cannot be the drift velocity due to the applied field. The actual mechanism for electron transport inside the box depends on the box size. Diffusion is the primary mode of transport at length scales greater than the Onsager radius $\frac{e^2}{4\pi\epsilon kT}$, within which drift under the ionic field dominates, and smaller than $\frac{kT}{eE}$, beyond which drift under the applied field dominates. The box scale is between these two distance (R_F is the geometric mean of the two) so transport in and out of the box is primarily by diffusion. The time spent in the box is then of order $\frac{(2a)^2}{2D}$, where $D \approx 80$ cm²/s is the diffusion constant for electrons in liquid xenon [101]. The v entering the model should then be roughly $\frac{D}{a}$.

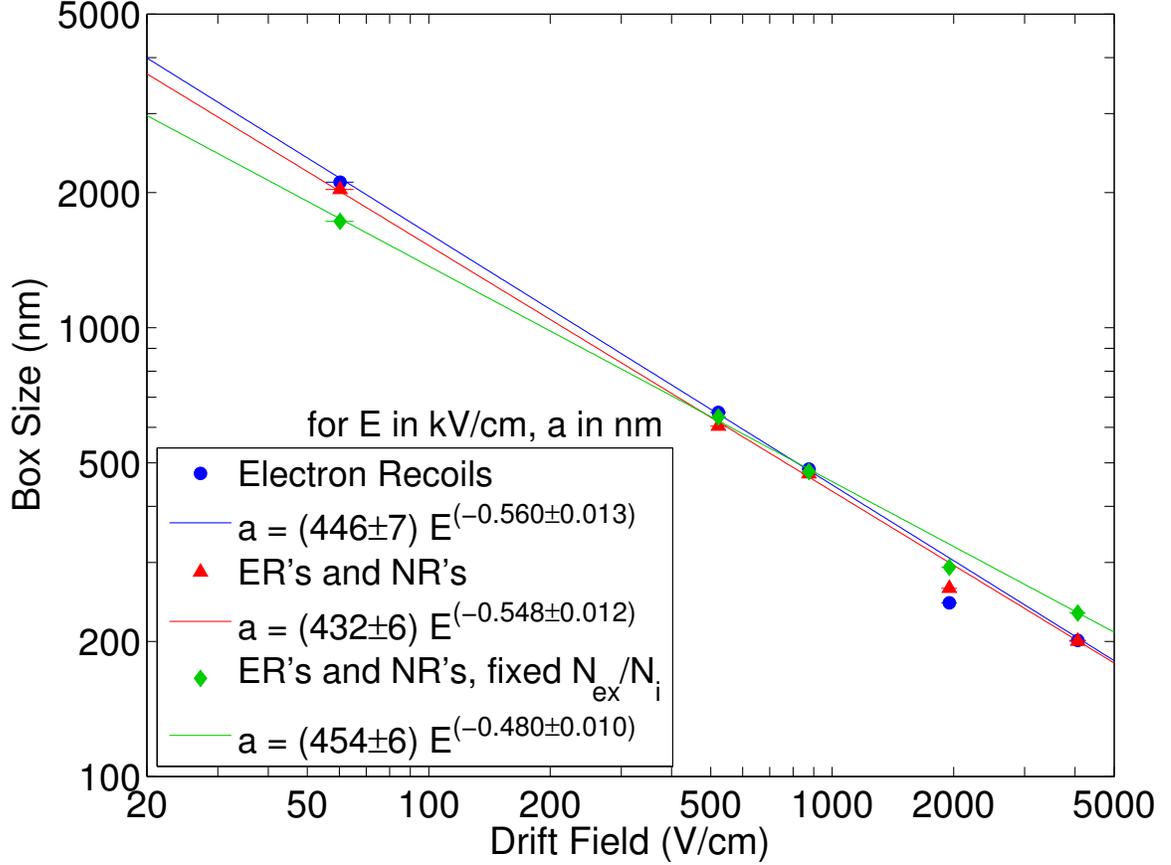


Figure 6.10: Box size versus drift field for our three fit schemes. The results at 1951 V/cm are not included in the linear fits. The fit with fixed $\frac{N_{ex}}{N_i}$ for nuclear recoils is consistent with the $\frac{1}{\sqrt{E}}$ scaling we expect for electrostatic length scales. Errors are determined from the uncertainties in the calibrated drift fields.

Taking this value for v and the Langevin recombination coefficient $\alpha = \frac{De^2}{\epsilon kT}$ we expect

$$\frac{\alpha}{av} = \frac{e^2}{\epsilon kT}. \quad (6.21)$$

It is important to note that the temperature in Eq. (6.21) is the electron temperature. Above ~ 10 V/cm electron mobility in liquid xenon begins to fall, and higher drift fields increase the electron temperature. Equation (6.21) therefore has a hidden field dependence. It is not obvious which electric field is important in this picture. We care about electrons inside the box, where the ionic field dominates, but electrons may not remain in the box long enough to thermalize with the higher field. From our above analysis, times in the box range from 1 ns at 60 V/cm to 10 ps at 4060 V/cm. Thermalization times for hot electrons in liquid xenon have been measured to be ~ 6.5 ns [102], so at least for high fields (small boxes) the electron temperature is determined by

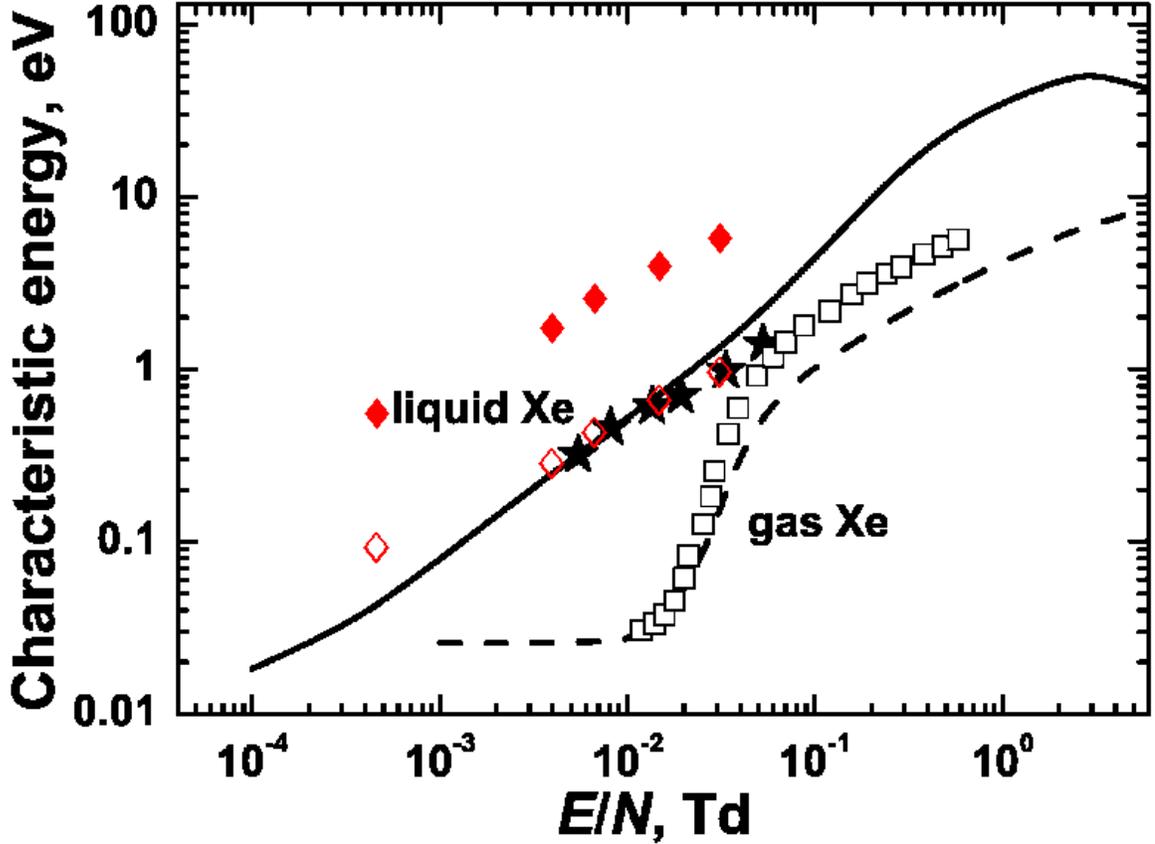


Figure 6.11: Characteristic electron temperature ($\frac{eD_T}{\mu}$, where D_T is the transverse diffusion coefficient) versus drift field. The black-and-white portion of this plot is copied from Atrazhev et al. [46]. The curves indicate calculations by Atrazhev et al., while the liquid and gas data points come from [101] and [103], respectively. The solid red diamonds indicate kT from Eq. (6.21) and Table 6.2, and the empty red diamonds are the same, divided by 6 to show the consistent scaling with data. At our lowest field point, the electron temperature may be influenced by the field of the ion. 1 Td = 10^{-17} V cm², and we have taken $N = 1.3 \times 10^{22}$ cm⁻³ (T=180 K).

the applied drift field. Figure 6.11 shows the characteristic electron energy $\varepsilon_c = \frac{eD_T}{\mu}$ versus drift field (where D_T is the transverse diffusion coefficient), compared to kT from Eq. (6.21) with our fit parameters in Table 6.2. We see $kT \approx 6\varepsilon_c \propto E^{0.6}$, so scaling and order of magnitude support the Langevin interpretation of our recombination model.

We can apply this interpretation to the original Thomas-Imel model and the result $\xi = \frac{N_i \alpha}{4a^2 v} \propto E^{-1}$ in large tracks. In the context of our model, $\frac{\alpha}{a^2 v} \propto E^{-0.1}$ and $a \propto E^{-0.5}$, and so the field dependence in large tracks comes from the varying number of ions in each box as the box-size changes. Since electrons drift through the track, an ion may recombine with electrons arising anywhere in a column with volume proportional to a^2 . Taking $N_i \propto a^2 \propto E^{-1}$ gives us the scaling that Thomas and Imel observed.

We can use the same arguments to predict the dependence of recombination on density. We have $T \propto \rho^{-0.6}$ and, from the Clausius-Mossotti relation, $\epsilon \propto \rho^{0.6}$. This gives $\frac{\alpha}{a^2 v} \propto \rho^{0.3}$, indicating that in small tracks, recombination is three times more sensitive to density than it is to the applied field. We have to be careful looking at larger tracks. Linear dimensions in a track scale like ρ^{-1} , but we can simplify matters by multiplying every unit of length by ρ so that, in our new units, the track geometry is density independent. This gives $a \propto \rho^{0.7}$, and so $N_i \propto \rho^{1.4}$, finally giving $\xi \propto \rho^{1.7}$ for large tracks. In Xed we can vary ρ by $\sim 5\%$ by changing the temperature, which would have the same effect on recombination in large tracks as an 8.5% change in the drift field. This should be observable, and hopefully future work will test this prediction.

We do not have enough information to draw conclusions on the nuclear recoil parameter $\frac{N_{ex}}{N_i}$. It appears to be field dependent, which supports an Onsager recombination interpretation, but we need an estimate of the energies of ionized electrons in nuclear recoil tracks to calculate the expected initial recombination fraction. The direct-exciton interpretation predicts a field independent $\frac{N_{ex}}{N_i}$, which does not fit as well below 10 keVr, but this interpretation does help solve the Lindhard problem. In Chapter 4 we saw that the measured Lindhard factor at low energies is higher than expected. A higher $\frac{N_{ex}}{N_i}$ for nuclear recoils than electron recoils would reduce the average electronic energy loss spent per exciton/ion, giving a larger Lindhard factor at all energies. This shift is not enough to produce agreement between the measured and theoretical Lindhard values, but it significantly closes the gap.

6.6 Conclusions on Recombination Physics

Our modified box model has successfully reproduced the mean recombination seen in electron recoils, and with one extra parameter matches nuclear recoils as well. The minimum in the electron recoil charge yield occurs when track size falls below an electrostatic length scale. For a single ion, this length scale roughly defines the region in which the ion's field is stronger than the applied field. Once an electron enters this region, recombination occurs as described by Langevin: electron transport is dominated by diffusion until it either enters the Onsager radius of the ion and recombines, or until it wanders back into the region where the applied field is dominant and escapes.

In our model, recombination is insensitive to track structure below the critical electrostatic length scale, and so the rising stopping power in low energy electron recoils is irrelevant. By the same token, so is the higher stopping power of nuclear recoils. In fact, it appears that at WIMP energies, our discrimination is not due to stopping power at all, but rather to the different modes of electronic

excitation open to nuclear recoils, resulting in a higher $\frac{N_{ex}}{N_i}$. This may be due to either immediate recombination of very low energy ionized electrons or to a higher fraction of direct excitons.

Field dependence enters our model by changing the critical electrostatic length scale as well as in the electron temperature, which in turn affects the Onsager radius and therefore the Langevin recombination coefficient. These effects conspire to give the weak field dependence we observe in small tracks.

We do not accurately reproduce recombination fluctuations, consistently falling short of the measured variance by a factor of three, but the physics interpretation of our model suggests several improvements which may address this problem. The most glaring problem is our mistreatment of diffusion, which we have replaced with a linear drift. We may also want to separate the initial distributions of ions and electrons. Although computationally expensive, it could be worthwhile to build a recombination Monte Carlo in which individual electrons are followed until they either enter the Onsager radius of an ion or escape the track.

In the end, while a physics understanding of the detector supports a good calibration and may be necessary to claim a discovery, it is possible to set WIMP limits in xenon without a perfect understanding of recombination. We end by describing how we have accomplished this in Xenon10 and discuss the remaining challenges for discovery-oriented large scale xenon TPCs.

Chapter 7

Xenon10 and Beyond — Searching for WIMPS with Large Scale Xenon TPCs

In Chapter 3 we described xenon TPCs in general and the Xed prototype in detail. Constructing and operating a large scale detector in order to produce a WIMP limit involves an entirely new set of challenges, from shielding and radio-purity to cryogenics, fluid handling, data acquisition and processing, calibration, and WIMP search analysis. Many of these issues we successfully faced in Xenon10, as addressed in the theses of fellow collaboration members [54, 55, 44, 81], while others become problematic only in the next generation of large detectors — some of these (e.g., Kr removal and large-scale cryogenics) are solved [60, 56, 104], and some are still being worked out. We make no attempt to give a comprehensive list in this thesis, but rather survey a few interesting problems and describe how they have been or are being faced by the Xenon10 and LUX collaborations.

We look at four categories of problems: fluid handling (now with constant recirculation), data handling (now with multiple PMTs), calibration (in a detector larger than neutron and gamma interaction lengths), and WIMP analysis, ending with the Xenon10 limit on the spin-independent WIMP-nucleon cross section.

7.1 Fluid Handling

In Xed we achieved sufficient purity for charge drift through ultrasonic cleaning, baking at ultra-high vacuum, and batch xenon purification at the start of each run, with periods of recirculation through the purifier as needed. We did not circulate xenon while taking data. In large detectors, where cleaning is more difficult and a greater drift length is needed, constant recirculation is necessary. This creates two issues which we did not face in Xed: dynamic liquid level control and the heat load from the returning xenon.

7.1.1 Liquid Level Control

Shifts in the liquid level height affect the strength of the gas extraction field (see Appendix A) as well as the path length in gas over which extracted electrons generate electroluminescence. A stable S2 signal therefore requires a stable liquid surface. This is non-trivial even without constant recirculation — in Xed we observed bubbles trapping under surfaces, displacing liquid and raising the liquid level. The situation is worse in large detectors, where there are many PMT's in the liquid whose bases are heat sources that can generate bubbles. Temperature changes can also cause significant liquid level shifts in large detectors (liquid xenon has a volume expansion of $2.5 \times 10^{-3} \text{K}^{-1}$ [45]). Recirculation and the resulting heat load from condensing xenon add to these problems, and it becomes impossible to maintain a stable liquid surface without some form of liquid level control.

There are currently two working methods to set a fixed liquid level, both relying on constant circulation of xenon. The first is the diving bell, implemented in Xenon10 and illustrated in Fig. 7.1. In this setup the extraction region and top array of PMTs sit inside a bell, with the incoming purified xenon gas fed into the top of the bell. The gas bubbles out the bottom of a vertical tube to the side of the bell and equal pressures force the liquid surface inside the bell to sit level with the end of outlet tube. By moving the outlet tube up and down, one can change the liquid level inside the bell. Changes to the total xenon volume shift the liquid height outside the bell.

One potential drawback to the diving bell is that it requires an incoming flow of xenon gas, so condensing the xenon before returning it to the main chamber is not an option. The gas pressure inside the bell is also higher than that in the main chamber, which will drive condensation inside the bell. This warms the liquid surface in the bell, resulting in a possibly undesirable thermal gradient. These issues are avoided when the liquid level is controlled using a spill-over dam, or weir, also shown in Fig. 7.1. Here the level is controlled using the outgoing liquid rather than the incoming gas. The weir can be designed to draw liquid from anywhere in the chamber, not necessarily from the surface.

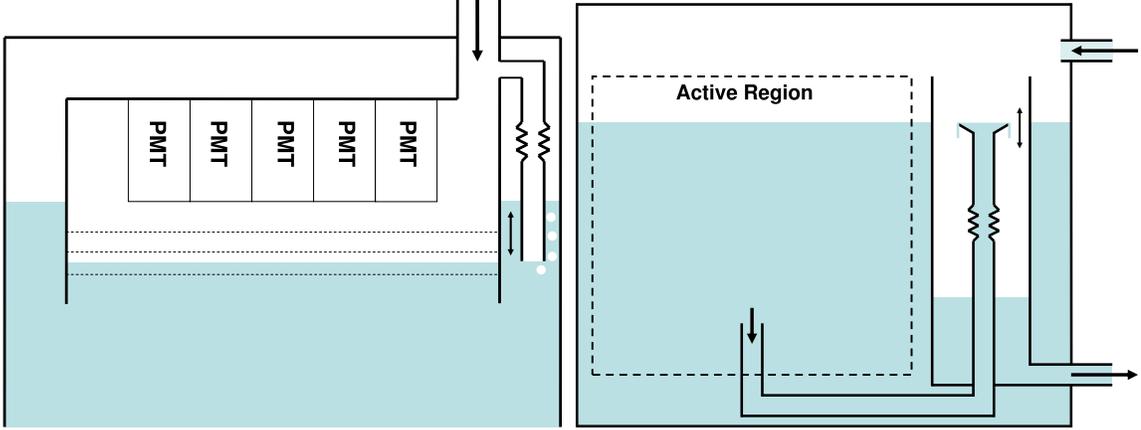


Figure 7.1: Two methods for liquid level control. Left: Diving bell schematic. Xenon gas returning from purification comes into the bell containing the top PMT array and electrodes, and bubbles out the inverted standpipe at the side of the bell. The liquid level inside the bell is set by end of the standpipe, and changes in the total liquid volume shift the liquid level outside the bell. The pressure inside the bell is higher than the outside pressure by the difference in liquid height inside and outside the bell. This will drive condensation in the bell, heating the liquid until the vapor pressure matches the bell gas pressure. If the both the liquid inside and outside the bell are in equilibrium with the gas, the temperature difference will be ~ 0.03 K per cm of column height [45]. Right: Weir schematic. Xenon is drawn from the weir for purification and returned to the main chamber as either liquid or gas. The liquid level height is set by the weir spillover height, and changes in the total xenon volume change the amount of xenon in the weir and outlet tube. The setup illustrated here shows how the weir may draw xenon from any point in the chamber, not necessarily from the liquid surface. Both the weir and bell can be built with movable parts (indicated by double-ended arrows) to adjust the liquid level. Xenon10 used a bell and LUX is using a weir to control liquid level.

Changes in the total xenon volume change the amount of liquid inside the weir outlet. Constant circulation is still necessary for the weir to respond to underfull as well as overfull conditions.

7.1.2 Recirculation Rates

The recirculation rate in a detector determines not only the time it takes to achieve the purity necessary for charge drift, but also the maximum drift length reached. Xenon acts as a solvent for many impurities. If we consider a constant rate of impurities dissolving into the liquid, a well-mixing model gives

$$\frac{dn}{dt} = R - \frac{n}{V}r, \quad (7.1)$$

where n is the number of impurities dissolved in the liquid, V is the liquid volume of the system, r is the liquid volume rate of recirculation, and R is the rate at which impurities enter the liquid.

The solution gives

$$n = n_0 e^{-t/\tau} + R\tau, \quad (7.2)$$

where $\tau = \frac{V}{r}$ is the time it takes to recirculate the entire liquid volume once. The steady-state level of impurities is proportional to this time constant. If recirculation operates in batch mode rather than in a well-mixed mode, then any given volume element spends a time τ between purifications, giving $n \approx \frac{R\tau}{2}$. R itself will decrease with time as the detector is cleaned by the circulating xenon, but the timescale for this in Xenon10 was months, and a plateau was eventually reached. If we assume R scales with the surface area of the detector and we require the impurity density $\frac{n}{V}$ to scale like one over the drift distance required, then τ should be roughly detector independent. The recirculation rate therefore needs to scale with volume of the detector.

In Xenon10, we saw drift lengths up to 4 meters (drift times of 2 ms) with a recirculation rate of 2.5 slpm ($\tau = 16$ hours) [57, 55]. The recirculation rate was limited by the heat load from the returning gas, which is dominated by the heat of condensation, 96 kJ/kg at 170 K. For a target τ of 1 day this corresponds to $\frac{P}{M} = 1.1 \frac{\text{W}}{\text{kg}}$, where P is the heat load from condensing xenon and M is the total xenon mass in the detector. The cryogenic systems used by Xenon10 and LUX expect parasitic heat loads on the main chamber of < 50 W, so condensing xenon is a problem in detectors of more than a few tens of kilograms. With the weir system, it is possible to liquify the purified xenon using a dedicated cooler before returning it to the main chamber, but another solution is to exchange heat between the outgoing evaporating xenon and either the incoming condensing xenon or the condensing gas in the main chamber. Various heat exchange designs are being tested in the LUX engineering runs, with one early version shown schematically in Fig. 7.2.

7.2 Data Handling

Data in a large detector is fundamentally different from what we dealt with in Xed, in that the S1 and S2 signals are split over many PMTs. The long drift times between the S1 and S2 pulses in a large detector also require long digitizer traces to capture events. This section explores the changes in data acquisition and processing that are required to handle the increased data flow and interpret multi-PMT data.

7.2.1 Data Compression

With 88 PMTs and drift lengths of $\sim 80 \mu\text{s}$, the raw digitized PMT waveforms from Xenon10 take more than two orders of magnitude more disk space per event than Xed data (this despite using a 5x slower digitizer). The vast majority of this is flat baseline — digitized waveforms are $\sim 156 \mu\text{s}$ long (to accommodate both S1 and S2 triggers), of which only $\sim 2 \mu\text{s}$ actually contains signal. This clearly

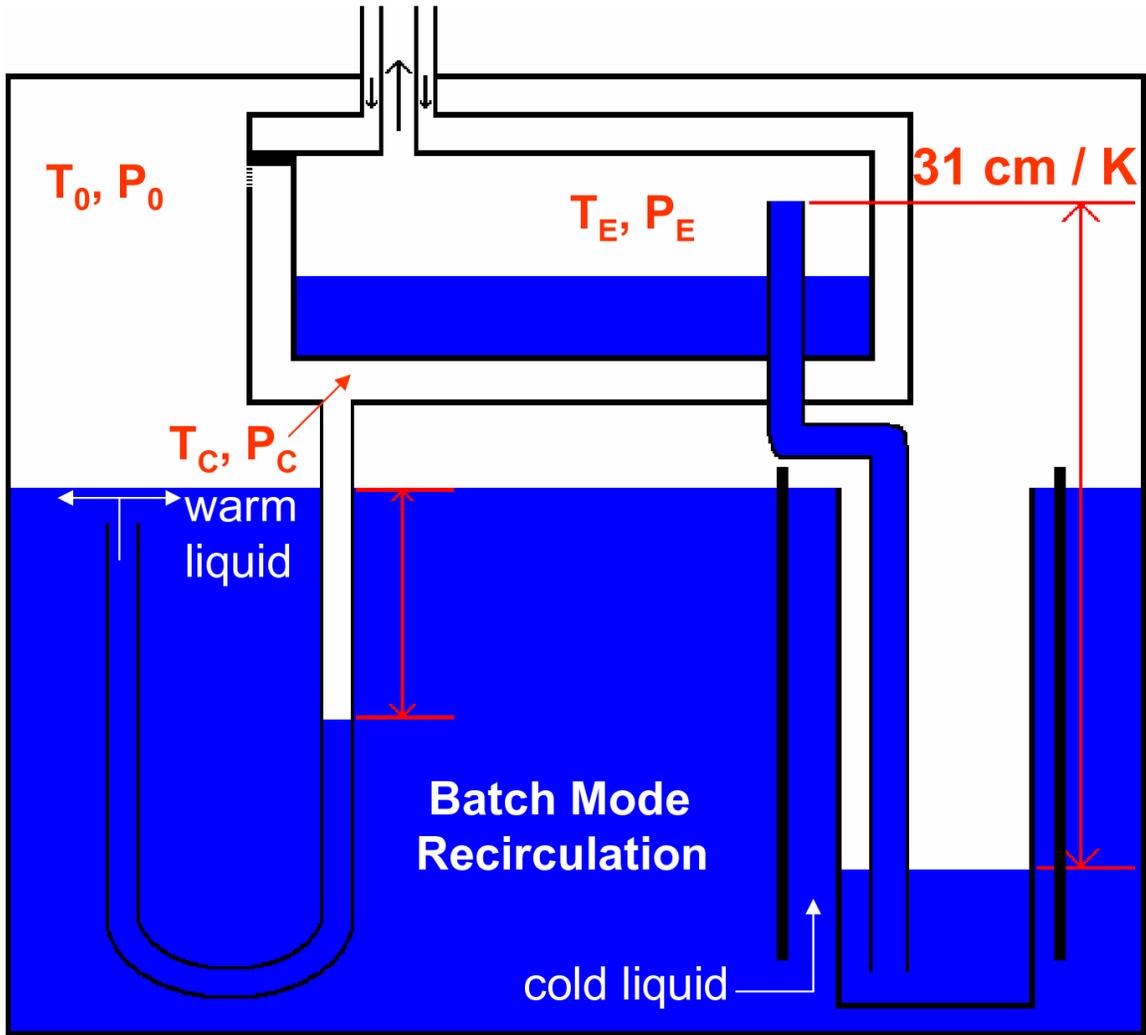


Figure 7.2: Schematic of a prototype heat exchanger used by LUX. The recirculation pump draws gas from the evaporator, which is fed liquid by the weir. In this setup, the weir level is set by the vapor pressure difference between the evaporator and the main chamber $P_0 - P_E$, which at 170 K gives a column height per ΔT of 31 cm/K [45]. Changes in the total xenon volume change the liquid volume inside the evaporator. The purified gas returns to the condenser, drawn here surrounding the evaporator. As drawn, all incoming xenon will condense before returning to the chamber, but this need not be the case. The condenser vapor pressure drives the column height between the liquid surface in the condenser outlet and the main chamber. T_E and T_C are determined by the recirculation rate and the heat flow between the evaporator and condenser. Increasing the recirculation rate increases the evaporative cooling on the evaporator, lowering T_E until the heat input from the condenser balances the evaporative cooling. As T_E drops, the pressure difference $P_0 - P_E$ increases and the weir empties. If the weir empties completely, the system enters a failure mode where it recirculates the chamber gas rather than the liquid. In the first run of this heat exchanger, the maximum flow rate before the weir emptied was roughly one fifth the target recirculation rate. Subsequent designs have increased the heat flow between condenser and evaporator (reducing the necessary $T_C - T_E$) and added safeguards against the empty-weir failure mode. Ultimately, any scheme in which liquid is drawn from the chamber and evaporates while still in thermal contact with the chamber will solve the heat load problem.

calls for data compression, similar to what was used in Xed (see Chapter 3). The compression needs to be applied before the data is written to a remote disk, otherwise disk access and network speeds limit the rate of both data taking and analysis. The compression therefore needs to be part of or directly in line with the data acquisition (DAQ) software.

The Xed compression algorithms are $\sim 5x$ slower (on a single cpu) than the DAQ rate in Xenon10, so faster compression schemes are necessary. The method decided on in Xenon10 was to first remove baseline noise in each waveform by setting any samples within ± 4 ADC (analog to digital converter) counts of the baseline to be exactly equal to the baseline. This threshold removed most of the white noise from the signal but kept electronics glitches (10–20 ADC counts, varying by PMT) and was well below the single photo-electron peak (~ 100 ADC counts). The data was then compressed using gzip (available in the free zlib library [105]). Since the baseline flattening resulted in long strings of bytes with the same values, this generic compression algorithm gave a $x14$ reduction in file size, and was fast enough that it did not limit the data acquisition rate. A $x20$ compression was demonstrated with more aggressive baseline flattening, as used in the first step of data processing in the next section.

With this compression, the limiting factor in the Xenon10 DAQ rate (maximum of 20 Hz) was the data transfer speed between the digitizers and the DAQ computer. While this is more than sufficient for low-background data, it is desirable to have a higher rate during calibrations. The LUX experiment will implement data compression directly in the digitizer firmware so that empty baseline data is not sent to the DAQ computer, increasing the maximum DAQ rate [106].

7.2.2 Data Processing

Since in the data compression in Xenon10 does not identify pulses in waveforms (as was the case for Xed), the first step of data processing is peak finding. (This is also the case for LUX, despite the rudimentary peak finding implemented by the digitizer firmware.) Again, the long waveforms make a processor intensive filter such as the one used to identify pulses in Xed too costly, requiring many processors to keep up with the DAQ.

The peak finding in Xenon10 operated on the sum of the digitized waveforms from all PMT's. Though much faster than scanning each PMT separately, this does create a noise problem. In a single PMT, a single photo-electron is clearly distinguishable from baseline, but summing the noise from 88 PMT's makes this more difficult. This is especially true in the face of coherent noise between PMT's, as shown in Fig. 7.3. The solution is to remove this noise before summing, performing a

more aggressive baseline flattening than was done in the DAQ. On each PMT we impose a threshold determined by that PMT's electronics noise, typically at 10–20 ADC counts (still a fifth or less of the mean photo-electron height). This effectively removes baseline noise from the summed signal.

We search for pulses in the summed flattened waveform, looking independently for S1's and S2's. To find S1 candidates we simply find the maximum of the waveform and step out to either side of the peak until the edges of the pulse are found. We then find the next maximum not contained in the first pulse, and so on until either no points above a set threshold remain or the maximum number of S1 candidates allowed in the analysis have been found. The S2 search is done the same way, but on a filtered waveform. To build the S2 filter, we first define the box filter,

$$b_i^{(a)}(y) = \sum_{j=i-a}^{i+a} y_j, \quad (7.3)$$

and the max value filter,

$$m_i^{(a)}(y) = \text{MAX} \{y_{i-a}, y_{i-a+1}, \dots, y_{i+a}\}, \quad (7.4)$$

where y is the waveform being filtered. The S2 filter is then

$$f^{(a_1, a_2)}(y) = b^{(a_2)}(y) - m^{(a_2 - a_1)}(b^{(a_1)}(y)). \quad (7.5)$$

In words, the filter takes the integral of the pulse in a box of width $2a_2 + 1$ centered on the sample in question, minus the maximum integral in any box of width $2a_1 + 1$ contained within the first box. Thus, the filter gives zero response for any pulse shorter than $2a_1 + 1$. We typically set $a_2 = 100$ samples ($1 \mu\text{s}$, the mean S2 width) and $a_1 = 10$ samples (100 ns, the typical S1 width), so that the filtered signal is completely insensitive to S1-like pulses yet still sensitive to single electrons extracted from the liquid surface (see Fig 7.4). This filter is also extremely fast, requiring only addition and comparison operations.

There are many variations of the above methods — for example, in the Xenon10 analysis, the S1 peek finder looked for S1's only prior to the largest S2 and ranked the S1 candidates strictly by height. An early LUX analysis scheme searched the entire waveform for S1 candidates and ranked them by their height minus their height in the S2 filter (with an appropriate scaling factor). The rules for determining the edge of a pulse in particular must be tuned to avoid either clipping the end of a pulse with a long tail or combining two distinct pulses into one. In general these details

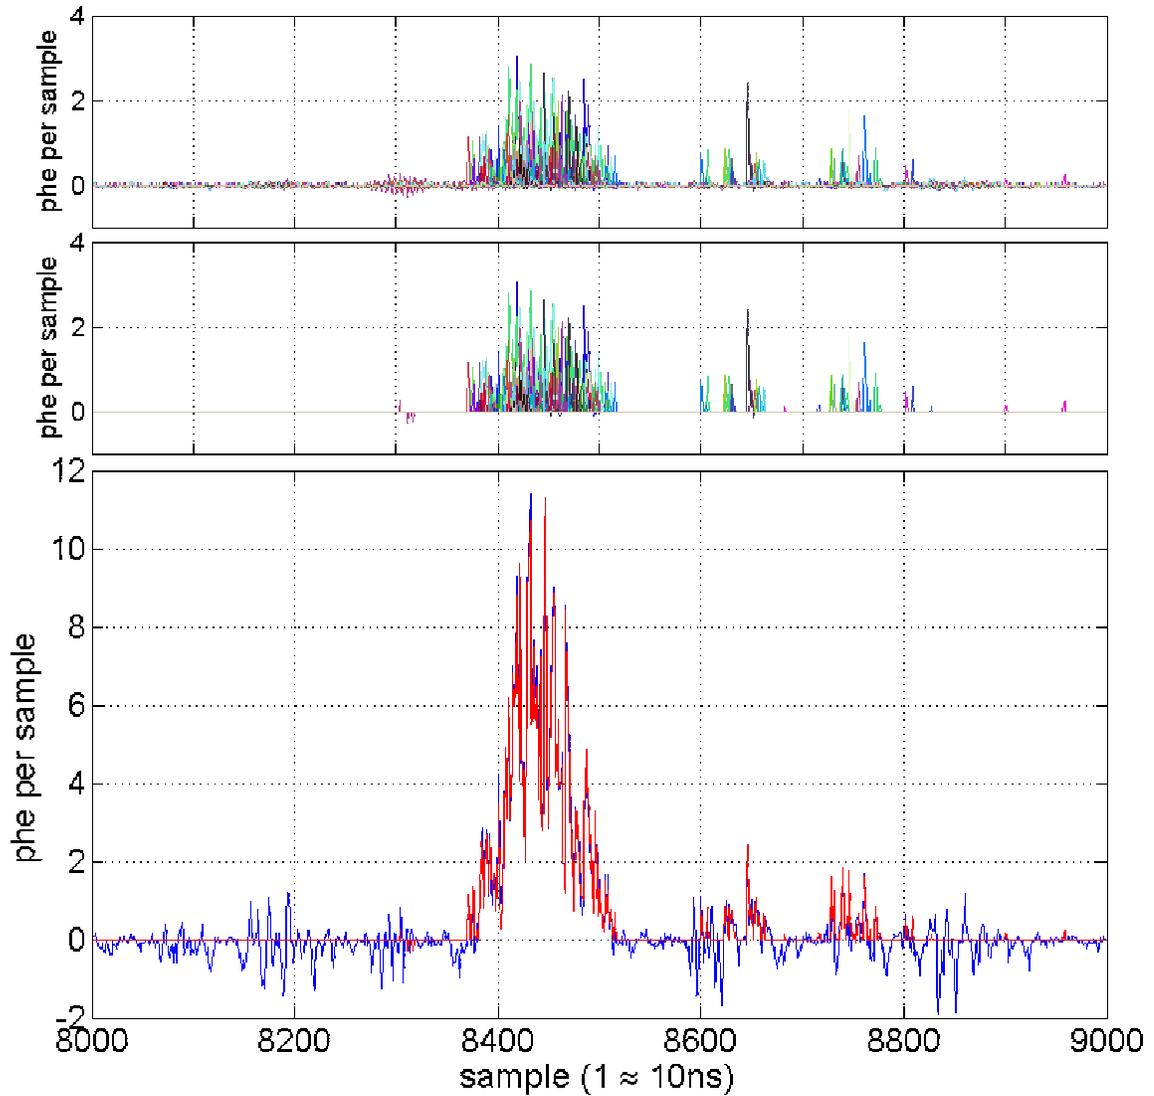


Figure 7.3: A small S2 in Xenon10, showing- Top: An overlay of all 88 PMT waveforms as output by the DAQ (with the basic 4 ADC count flattening). Middle: The same 88 waveforms with the more aggressive baseline flattening (10-20 ADC counts) used by the analysis. Bottom: The sum of PMTs with the DAQ flattening (blue) and the analysis flattening (red). In the top plots, the signal between samples 8600 and 8800 is clearly light in the detector, seen by multiple PMTs (in fact these appear to be late electrons coming off the liquid surface). When summing all PMTs with only the DAQ baseline flattening, these signals are lost in electronics noise, but they are clearly visible in the sum of the more aggressively flattened data.

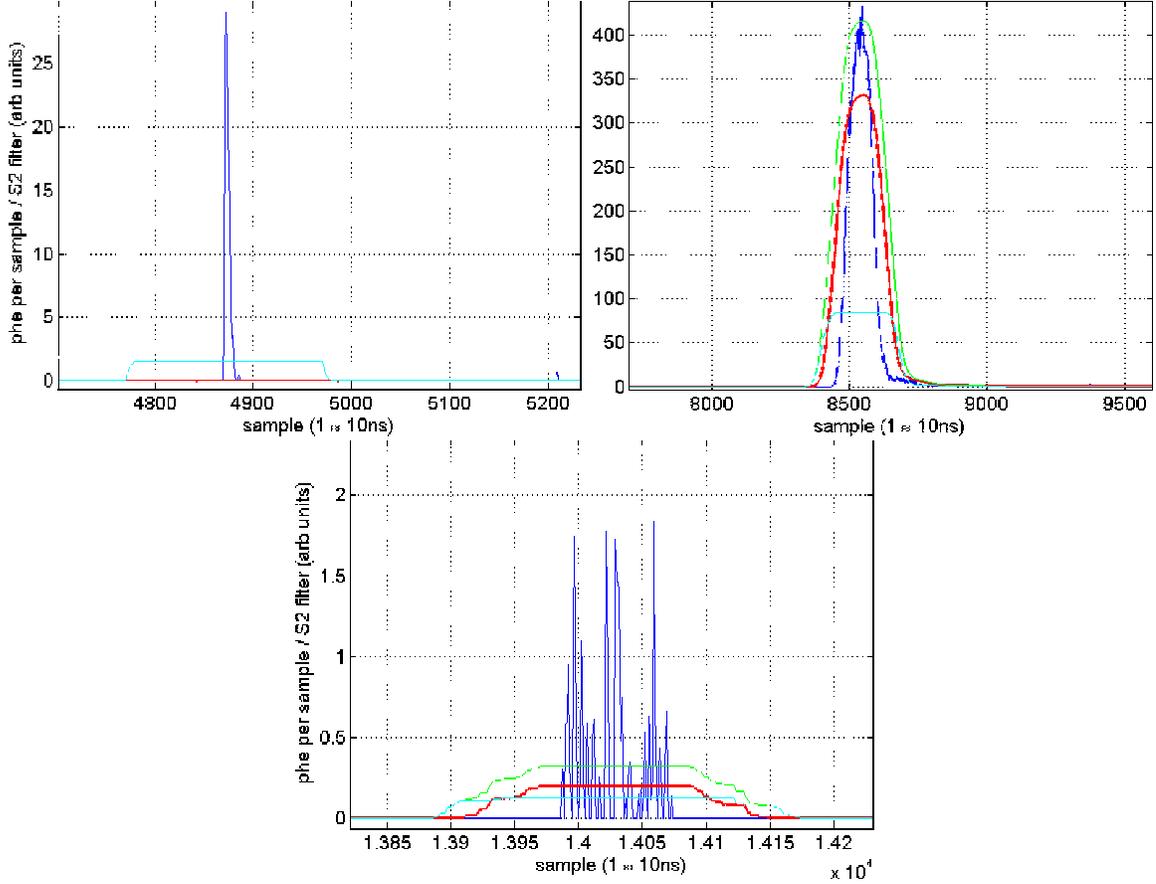


Figure 7.4: Examples of the S2 filter output on Xenon10 pulses, showing- Top-left: a typical S1 pulse, Top-right: a typical S2 pulse, Bottom: a single-electron S2. In each plot, the dark blue line shows the input (flattened and summed) waveform, the green shows $b^{(a_2)}(y)$, the cyan shows $m^{(a_2-a_1)}(b^{(a_1)}(y))$, and the red shows the final filter output (green minus cyan), where we have taken $a_1 = 10$ samples (100 ns) and $a_2 = 100$ samples (1 μ s). In the S1 example, the entire signal is contained within a 200 ns box, so the green and cyan lines are exactly on top of each other, and filter output is identically zero.

will change depending on detector performance.

7.2.3 Position Reconstruction

The great advantage of a multi-PMT detector is the ability to reconstruct event locations in 3-D by measuring the x-y location of the S2 light. This reconstruction is done using the χ^2 or maximum-likelihood methods described in [55], comparing the observed hit pattern in the top PMT array to that expected as a function of the x-y position of the event. The expected hit pattern is given by simulations (including effects such as reflections from the liquid surface and walls of the chamber), as shown in Fig. 7.5. The position reconstruction is verified by comparing the resulting

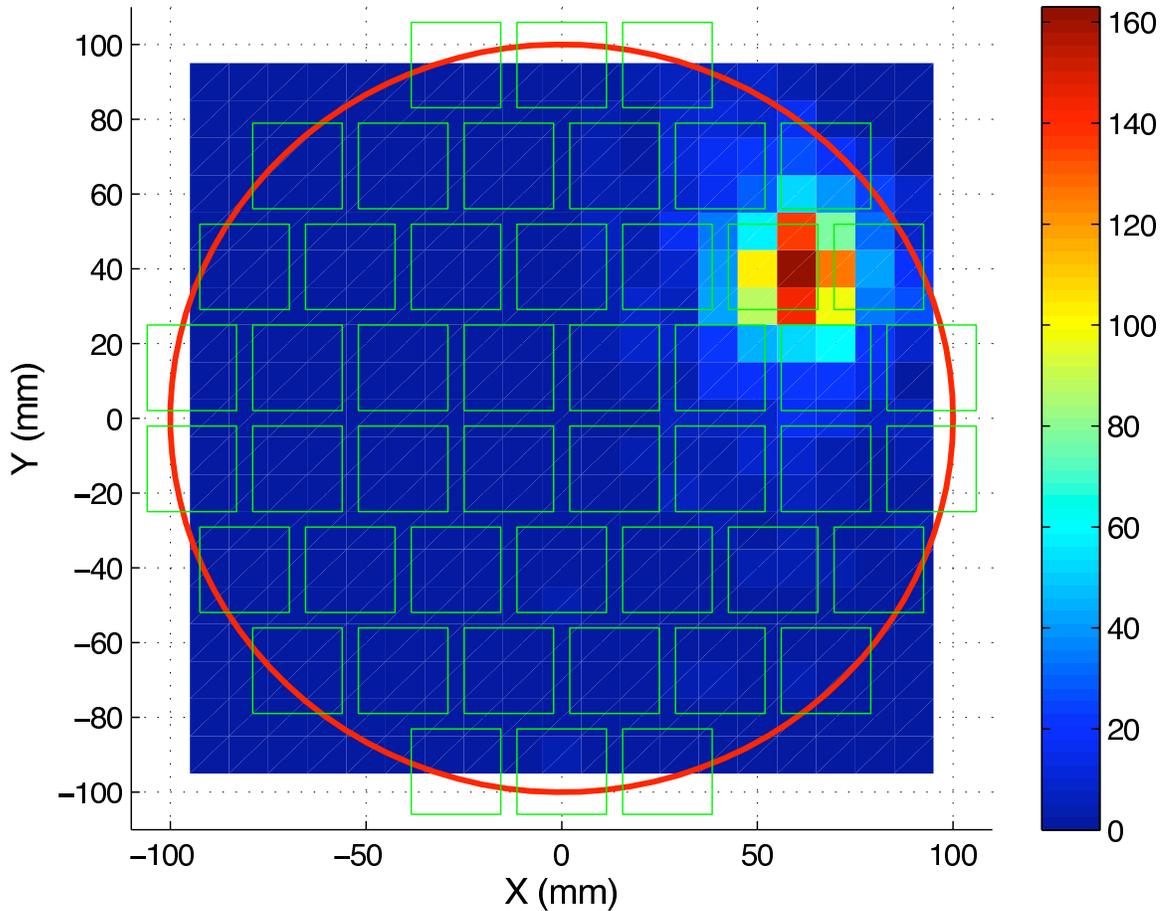


Figure 7.5: Monte Carlo simulation of S2 light in Xenon10. The red circle indicates the active region of the detector, and the green squares are the top-array PMTs. The color histogram shows the number of S2 photons arriving in each bin in the top-PMT plane for an event at $X \approx 60$ mm and $Y \approx 40$ mm. A single S2 is typically concentrated in one or two top PMT's. Plot courtesy of John Kwong, originally appearing in [44].

spatial distribution events with that expected from a Monte Carlo for a given source. The most useful comparison is to a source known to be uniform in the detector, such as the activated xenon calibration discussed in Section 7.3.2.

7.3 Detector Calibration

There are three calibration tasks to be performed in any large detector. The position dependence of the charge and light gains must be measured and corrected for, the absolute charge and light gains measured, and the nuclear and electron recoil bands characterized to determine the WIMP acceptance region and background leakage rate. The challenge for each of these is to produce the necessary events in the middle of the detector. Band characterizations are particularly difficult in a

large detector.

7.3.1 Position Dependent Corrections

Position dependent corrections require a peak (fixed energy event) occurring throughout the active volume in the detector. Photo-absorption peaks are inadequate, since the low energy gammas required for a clean peak do not penetrate to the center of the detector, but there are several nuclear transitions in xenon giving internal conversion electrons that may be used. The ^{129}Xe isotope (26.4% abundance) has two accessible excited states, one at 39.58 keV (0.97 ns half-life) which may be excited by inelastic neutron scattering (see Chapter 5), and one at 236.14 keV (8.88 day half-life) that undergoes a two-step decay (196.56 and 39.58 keV), both typically through internal conversion. There is also a metastable state in ^{131}Xe at 163.93 keV (11.84 day half-life) that gives a single internal conversion electron (nuclear data from [59]). The two meta-stable states are produced by exposing a xenon sample to a neutron source [107], and the activated xenon is then added to the detector. This was done following the WIMP search runs in the Xenon10 experiment. An alternative method using ^{m83}Kr (half-life 1.8 hours) is described in [104].

Figure 7.6 shows the position dependence in r and z of the 164 keV S1 in Xenon10. This dependence comes from the varying light collection efficiency with event position, and does not appear to depend on azimuthal angle. The x-y position dependences of the 40 and 164 keV S2's are shown in Fig. 7.7. The variation in S2 with position is due to the warping of the anode grid, and can be seen in the S2 pulse-width as well as the integral. A shorter S2 pulse (i.e., shorter path length in gas) corresponds to a larger S2, indicating that the increased field with the warping grid has a greater effect than the reduced gas region. Also note the change in the S2 map between the neutron data (40 keV peak) and activated xenon data. This may be due to liquid level or bell pressure changes when the activated xenon was added. These S1 and S2 maps are used to adjust measured S1 and S2 values to the S1 and S2 an event would have had, had it occurred in the center of the detector. See [44, 55] for details on the construction of position dependent correction maps.

7.3.2 Absolute S1 and S2 Calibrations

In Section 5 we used the 122 keV photo-absorption peak as a reference to find the absolute S1 and S2 calibration in Xed. This can be done in a large detector as well, but requires robust position-dependent corrections since the photo-absorption events will occur only at the edge of the detector. The 164 and 236 keV activation lines occurring throughout the detector would be ideal, but to date

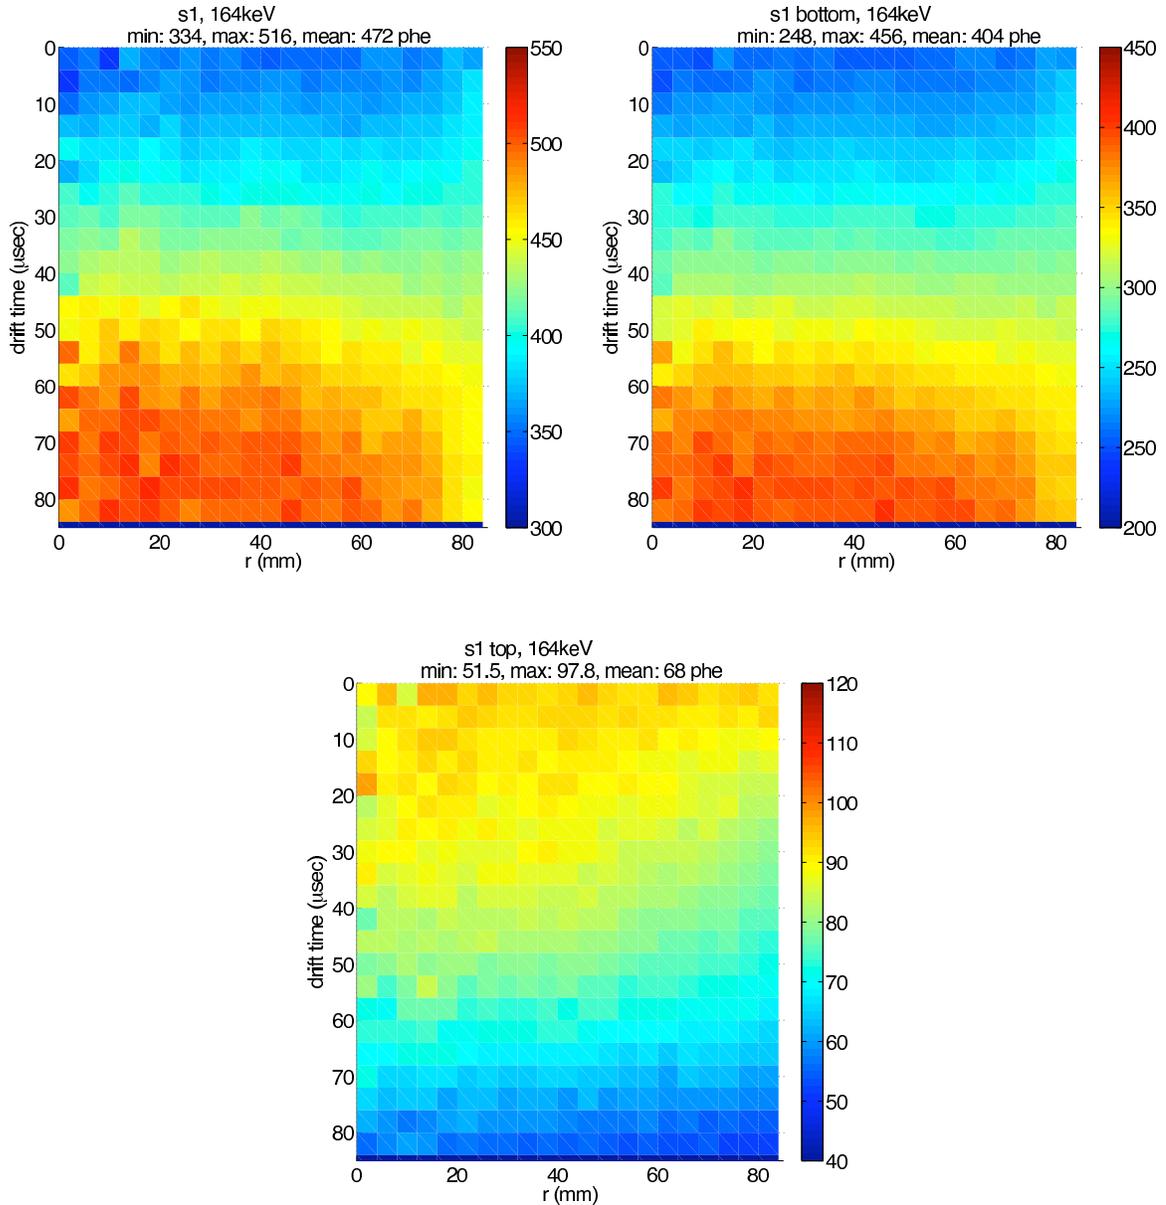


Figure 7.6: Position dependence of the 164 keV ^{131m}Xe decay S1 signal in Xenon10. The top-left shows the S1 in all PMTs, the top-right the S1 in bottom PMTs only, and the bottom the S1 in top PMTs only. The depth is given in terms of drift time ($\sim 2 \text{ mm}/\mu\text{s}$). No significant azimuthal dependence was observed. Plot courtesy of John Kwong, originally appearing in [44].

the field dependence of the charge and light yield for these peaks has not been measured.

An alternate calibration technique is to measure the single-electron S2's from electrons popping off the liquid surface. These are present at some level in all Xenon10 data, and their source is uncertain. Most of these electrons occur at the edge of the detector, indicating that they may be charge that was trapped on the walls and later freed, but many also occur in the fiducial region.

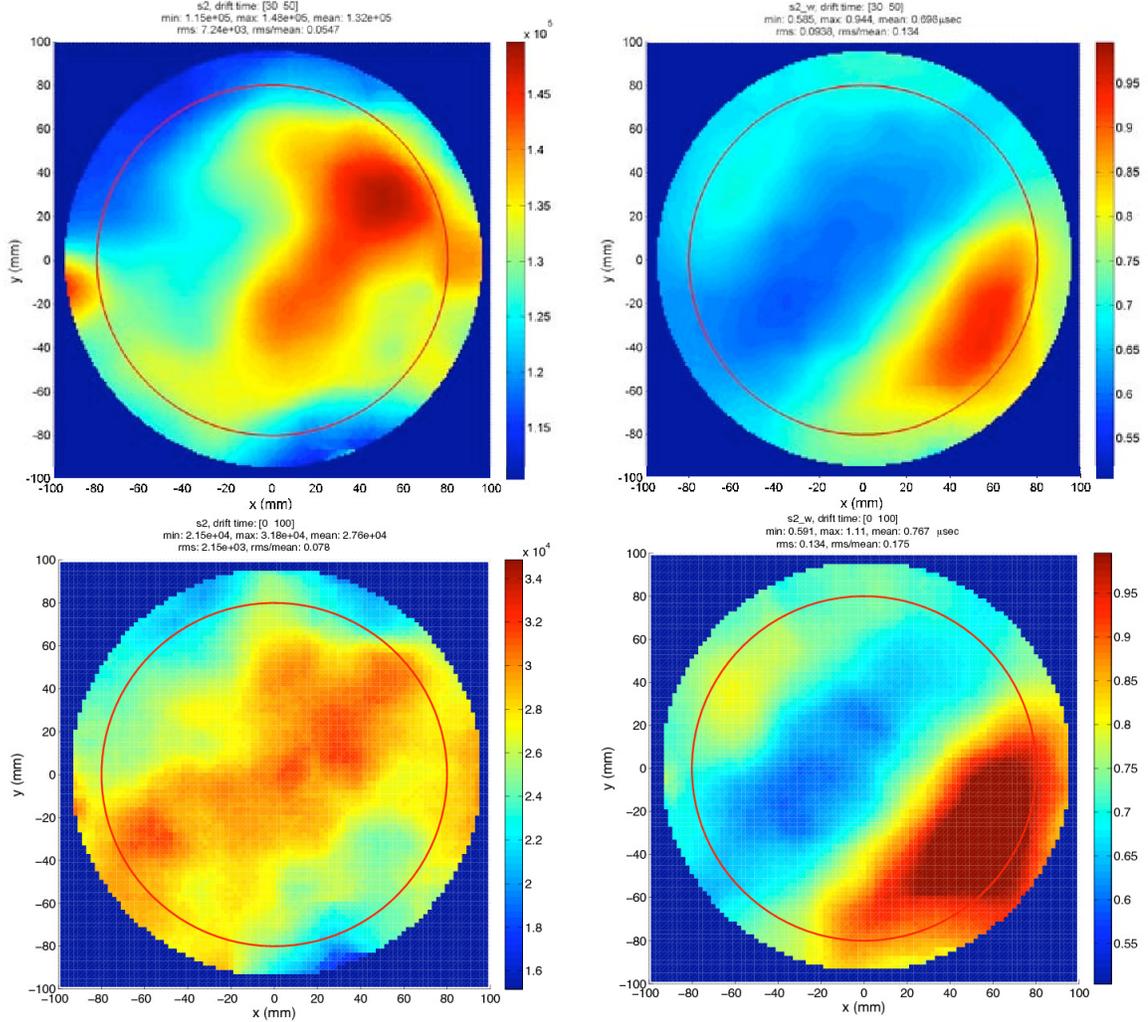


Figure 7.7: Position dependence of the S2 integral (left) and width (right) for 164 keV (above) and 40 keV (below) events in Xenon10. The asymmetric patterns are likely due to warping of the anode mesh. A wider gas gap (larger S2 width) corresponds to a smaller S2 amplitude, indicating that the changing electric fields with grid separation are the dominant factor in the S2 variation. Plot courtesy of John Kwong, originally appearing in [44].

These may be electrons from previous events that failed to extract into gas, although the mode for subsequent extraction is unclear. Higher event rates give a higher rate of these single electrons.

Figure 7.8 shows the single electron signal in the Xenon10 neutron data. These pulses are identifiable as S2's based on timing ($1 \mu\text{s}$ pulse width) and the division of light between top and bottom PMT arrays. Taking this for our absolute S2 calibration (i.e., assuming 100% efficiency for drifting electrons from the event site and extracting them into the gas) we can find the S1 calibration from any peak with known energy, using the W-value found in Chapter 4. With the absolute calibration, we can separate out the uncorrelated S1, S2, and recombination fluctuations

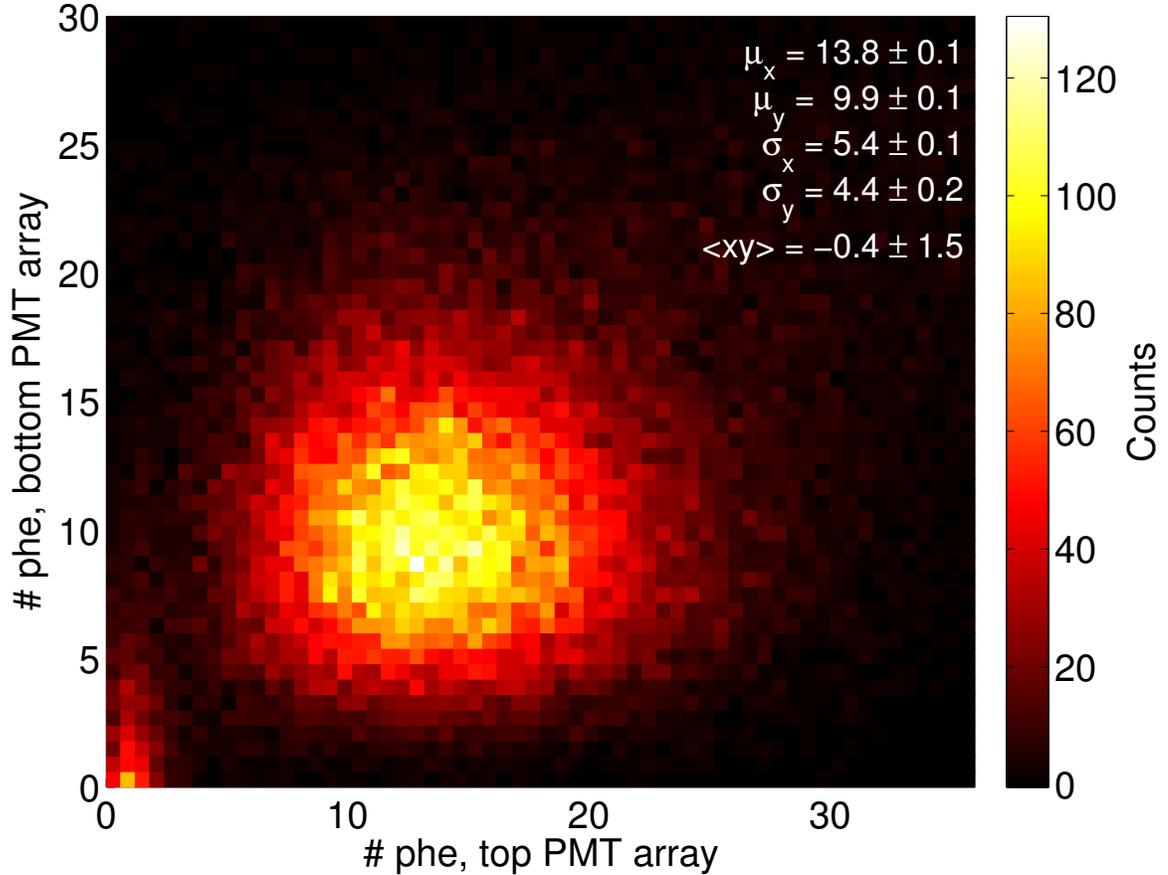


Figure 7.8: Single electron S2's in the Xenon10 neutron calibration data. The plot shows the total signal in the bottom and top PMT arrays for S2-like pulses (width of $\sim 1 \mu\text{s}$) within the fiducial radius (80 mm), after position corrections. The division of signal into top and bottom PMTs verifies that this light originates in the gas gap (S2 region).

as before, shown in Fig. 7.9 for the 40, 164, and 236 keV peaks in Xenon10.

7.3.3 Nuclear Recoil Measurements

Nuclear recoil calibration is necessary to determine the WIMP signal and measure the efficiency of any cuts applied to the data. These calibrations are done exactly as in Xed, using a ^{252}Cf or AmBe neutron source (see Fig. 7.10), but with the complication that many events are now multiple scatters. The neutron interaction length in liquid xenon is $\sim 15 \text{ cm}$ [62], so Xenon10 still sees single scatters throughout the detector volume, but next generation detectors may have to use multiple-scatter or edge events for calibration.

Multiple scatters give a single S1 that is sum of the S1's from each individual scatter, plus one S2 for each scatter (assuming the scatters are separated spatially). If we take the sum of the individual

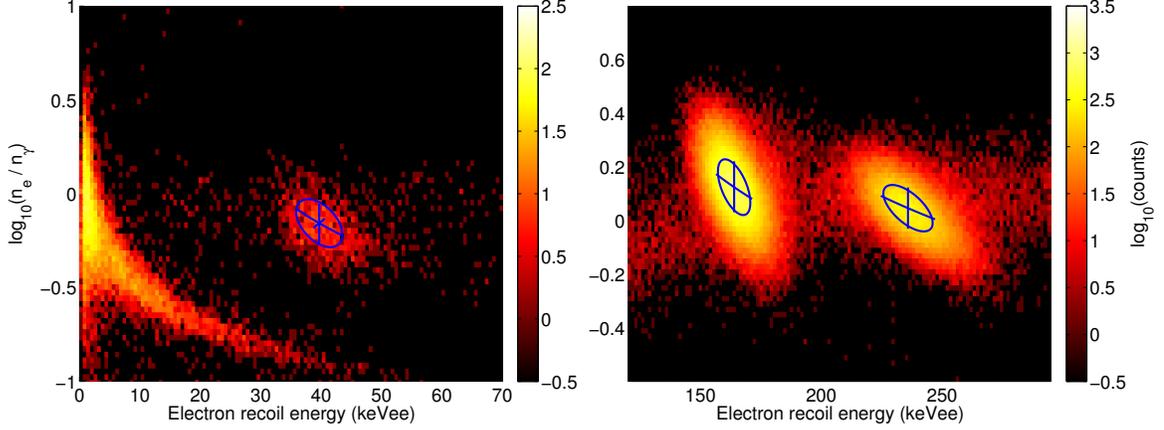


Figure 7.9: Fits to the 40, 164, and 236 keV nuclear excitation peaks in Xenon10, using the single electron peak for S2 calibration. After position dependent corrections, S2 fluctuations are evidently very small, but S1 fluctuations are larger than expected. This is somewhat suspect — the activated xenon data suggests a non-linearity in the S1 response. To fit both 164 and 236 keV peaks with linear calibrations in S1 and S2, the signal from a single electron should be 27.3 photo-electrons (using Eq. (4.26)), while the single electron peak in the activated data appears at 20.1 photo-electrons. This large discrepancy can be resolved by allowing the S1 calibration to fall by $\sim 5\%$ between the 164 and 236 keV peaks. Alternatively, there may be less than 100% extraction efficiency for electrons at the phase boundary. This would produce a difference between the S2's corresponding to a single-electron extracted from the event site and a single-electron extracted from the liquid surface. These issues are unresolved in Xenon10, as they do not significantly affect the WIMP limit.

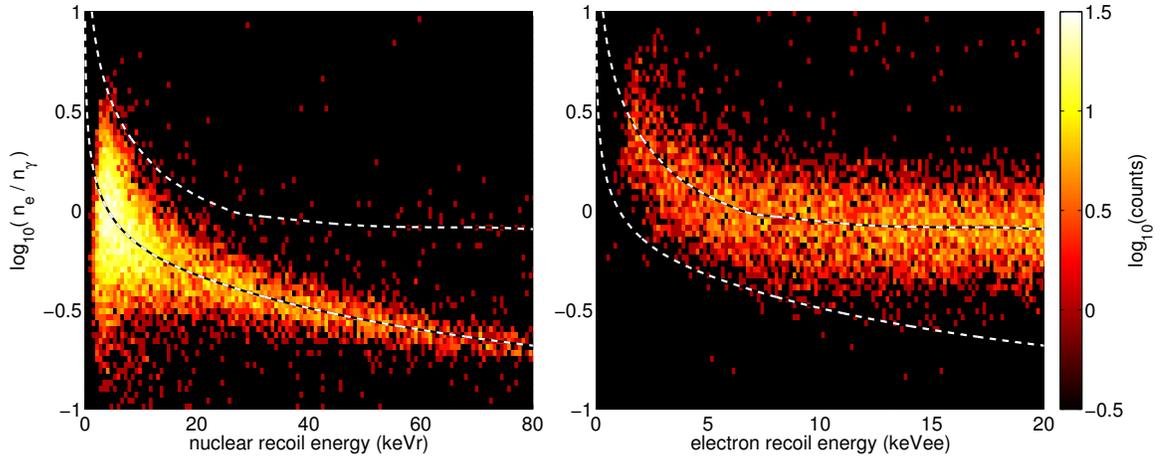


Figure 7.10: Nuclear and electron recoil bands in Xenon10. The nuclear recoils (left) are elastic neutron scatters from an AmBe neutron source, and electron recoils (right) are Compton scatters from ^{137}Cs . The data is shown after all position-dependent corrections, basic cuts, fiducial cuts, and quality cuts (see Section 7.4.1). As in the previous chapters, we plot using the recombination independent energy scale, taking $\mathcal{L} = 0.25$ for nuclear recoils (note, this is not the scheme used in the Xenon10 standard plots). The S1 threshold (2 photo-electrons) and S2 threshold (300 photo-electrons) set in the analysis are clearly visible in the upper-left and lower-left of the plots, respectively.

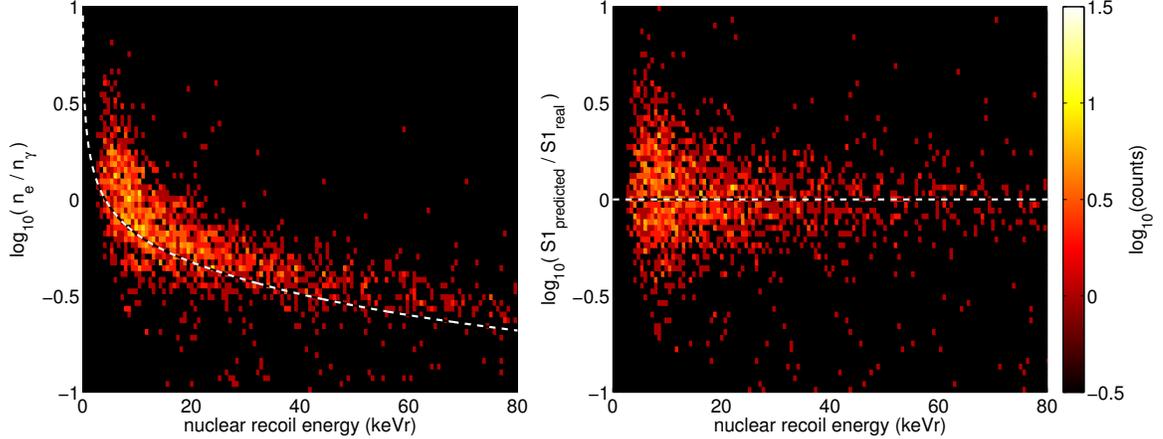


Figure 7.11: Double elastic scatters by AmBe neutrons in Xenon10. The left plot shows the double-recoil band in $\log_{10}\left(\frac{n_e}{n_\gamma}\right)$ versus nuclear recoil energy, taking the sum of the two observed S2's to determine n_e . The dashed line indicates the single-scatter centroid from Fig. 7.10. The right plot compares the observed S1 to that expected based on the two S2's observed. The mean and width of the distribution in $\log\left(\frac{S1_{\text{exp}}}{S1_{\text{obs}}}\right)$ is consistent with the single scatter nuclear recoil band. Next generation detectors may need to rely on multiple scatters when calibrating the nuclear recoil band in the center of the detector.

S2's as the charge signal for the event, the resulting multiple scatter band lies above the single scatter band (see Fig. 7.11). This can be quantified by comparing the observed S1 to the sum of the expected S1's associated with each S2. Thus, the charge yield of the multiple scatter band can be predicted from and used to verify the expected yield of the single scatter band. Determining cut efficiencies on single scatters using the multiple scatter band may be more problematic, especially for cuts based on PMT hit patterns (see Section 7.4.1).

7.3.4 Discrimination Measurements

Electron recoil band measurements face the same problem as nuclear recoil measurements — single scatter events in the center of the detector from outside sources are very rare (the maximum gamma interaction length in xenon is ~ 10 cm, corresponding to 4 MeV gammas [61]). In fact the situation is worse than for nuclear recoils, which typically forward scatter, giving the low energy recoils we are interested in. The Compton scattering spectrum is relatively flat and extends to a much higher energy, so that multiple Compton scatters tend to have at least one energetic scatter that swamps the signal from any WIMP range scatters.

The Xenon10 detector is small enough to still see single scatter events throughout, but collecting statistics at low energies was difficult. A high-energy veto was required to prevent triggering on

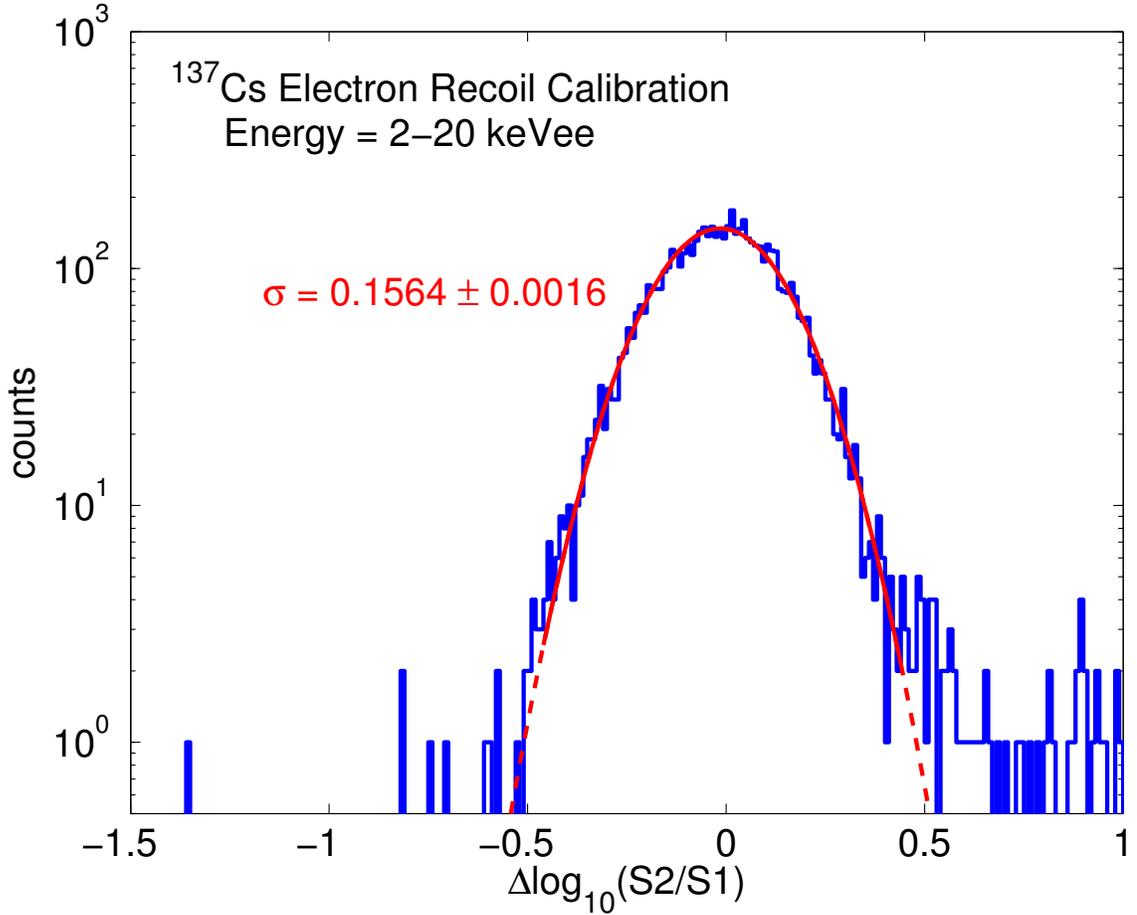


Figure 7.12: Electron recoil band profile in Xenon10, based on the ^{137}Cs calibration. The histogram shows the distribution in $\Delta \log_{10} \frac{n_e}{n_\gamma}$ for electron recoils from 2–20 keVee. The red line shows the Gaussian fit to the peak, with a dashed line in regions not included in the fit. The fit region extends to $\pm 3\sigma$. The tail at low charge yield may include Gamma-X events that were not tagged by the cuts. The tail at high charge yield may also be due to strange pathologies, but since shifts to higher yield do not result in leakage, these events were not studied.

events above the interesting energy range so that DAQ speed did not limit data taking. This was accomplished, but the detector itself began to display poor behavior when the physical event rate was over 100 Hz. This behavior included light flashes in the detector unassociated with clean events (probably spurious S2's from electrons popping off the liquid surface), especially in the millisecond following a large S2. With the physical trigger rate for quality data limited to ~ 50 Hz, the electron recoil calibration statistics gathered were in the end equal to the electron recoil background in WIMP data. Discrimination numbers were produced, as well as a band profile (see Fig. 7.12), but no background subtraction was attempted.

For larger detectors still, electron recoil calibration with a gamma source becomes unmanageable,

except at the edges of the detector where single scatters are still possible. There is interest in using an internal beta source for calibration, such as ^3H (18.3 keV endpoint, 12.33 year half-life). If xenon recirculation operates in batch mode, the source could simply be injected in the recirculation loop downstream of the purifier, travel through the volume, and be removed on the next pass through the purifier. Such techniques are currently under development for LUX. Without this or a similar solution, discrimination measurements in large detectors will depend on edge events and/or measurements in smaller detectors.

7.4 WIMP Analysis

At last the purpose of a large scale detector is to detect WIMP recoils or set a limit on the WIMP-nucleon cross section, and in this final section we outline the steps of a WIMP analysis. We end with the Xenon10 limit on the spin-independent WIMP-nucleon cross-section.

7.4.1 WIMP candidate selection

The first step in a WIMP analysis is to determine the rules for accepting events as WIMP candidates. These cuts on the data fall into three general categories, which we will term Basic Cuts, Background Cuts, and Quality Cuts. Basic Cuts are cuts with an assumed 100% acceptance for WIMP events, such as rejecting multiple scatters and obvious noise triggers. Background Cuts include the cut in $\log\left(\frac{S_2}{S_1}\right)$ to remove electron recoils and a fiducial cut to remove single scatters at the edge of the detector. Quality cuts target any recurring pathological behaviors in legitimate background events that may cause them to leak into the WIMP acceptance region. The efficiency of Background and Quality Cuts must be measured using the nuclear recoil data and included when calculating WIMP recoil rates. All cuts must be frozen before unblinding WIMP search data, to avoid tailoring cuts to remove specific WIMP candidates from the data.

Background Cuts are chosen to optimize limits based on WIMP exposure (favoring weak cuts) and estimated background (favoring strong cuts). Often this is done so that the expected background in the WIMP search data is ~ 1 event. In Xenon10, we simply set the $\log\left(\frac{S_2}{S_1}\right)$ threshold for 50% acceptance as in the simulations in Chapter 5, while the fiducial volume was set primarily to exclude Gamma-X events (see Section 7.4.2 and Fig. 7.15).

Quality cuts fall into three general categories: PMT hit patterns, pulse timing characteristics, and general detector health. Hit pattern cuts target events that would have been rejected with an accurate position reconstruction, but were misidentified. These include the Gamma-X cuts listed

in Section 7.4.2 and cuts targeting multiple-scatters with the same drift-time. Pulse timing cuts use various measurements of pulse width, rise time, and fall time to reject events with atypical S1's and S2's. General health requirements include low baseline noise and a limit on the total signal not contained in either S1 or S2 (i.e., the spurious photo-electron rate). Many of these cuts are also used to distinguish bona-fide events from non-events arising from electronics glitches or scatters in the gas or outside the active region, and to some extent it is a judgement call at what point the cut becomes a Quality Cut (and therefore included in the WIMP efficiency) rather than a Basic Cut.

The general procedure followed by Xenon10 is to develop Quality Cuts targeting leakage events in electron recoil calibration data and unblinded background data, verify that these cuts do not significantly reduce nuclear recoil acceptance in the neutron data, and then apply the cuts to a second set of electron recoil data to verify their effect on leakage events. Many proposed cuts fail this final test, turning out to have been statistical flukes in the earlier data. (In fact, an additional high statistics electron recoil calibration dataset would have helped tune the cuts in Xenon10, as was apparent when analyzing the WIMP search data. This points out the need for high statistics in-situ electron recoil calibration.) The final sets of cuts used for the Xenon10 analysis are discussed in detail in [55]. Of these, the most important to future detectors are the 'Gamma-X' cuts, discussed next.

7.4.2 Non-Gaussian Leakage

The primary source of leakage in Xenon10, aside from the normal fluctuations of electron recoils into the WIMP region due to light-collection and recombination as described in Chapter 5, was the 'Gamma-X' event. These events consist of a single Compton scatter in the active region of the detector, plus a second scatter in one of the non-active liquid regions where S1 light is produced but no charge is collected. These include the reverse-field region between the cathode and bottom PMT array, the liquid outside the walls of the active region, the liquid <1 mm inside the active region wall, and the liquid in the space left for the resistor chain between the field shaping rings. By collecting the S1 light from two scatters and the S2 from one, the $\frac{S2}{S1}$ ratio of these events falls and they may be classified as nuclear recoils. Since these are multiple scatter events they tend to appear at higher energies, but they are problematic at the high end of the WIMP region.

Gamma-X events are targeted by the fiducial cut (the active region scatter tends to be near the cathode or edge of the detector) and by their S1 hit pattern, looking for either a concentrated hit pattern in the bottom array (for reverse-field region scatters) or an excess of light in edge PMT's

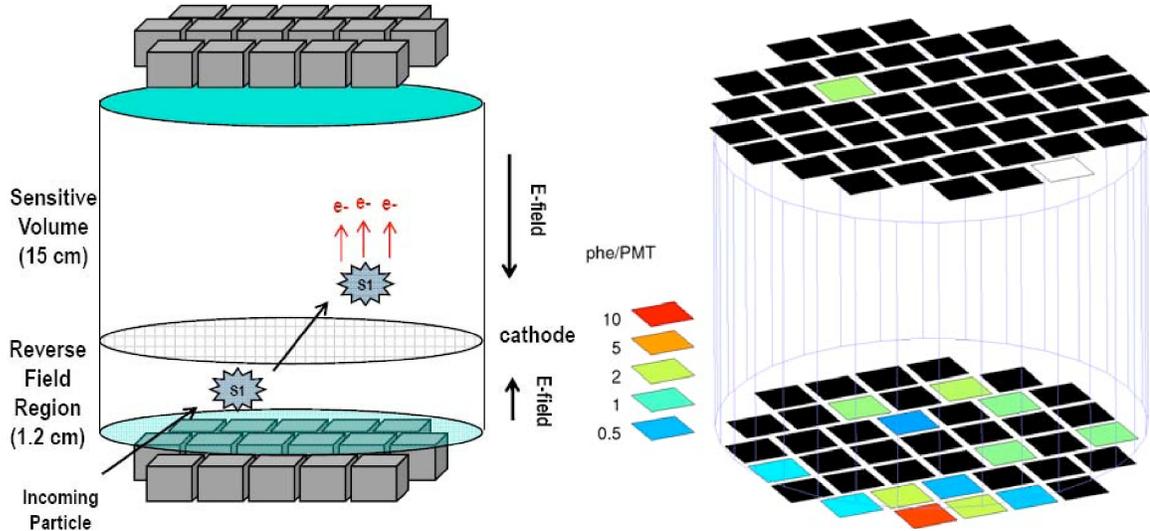


Figure 7.13: Left: An illustration of a Gamma-X event in Xenon10, where the missing-S2 scatter occurs reverse field region between the cathode and bottom PMTs. Right: The S1 hit pattern from an event identified with this behavior. This figure on the left is courtesy John Kwong, appearing in [44], and the figure on the right is courtesy of Peter Sorensen, originally appearing in [55].

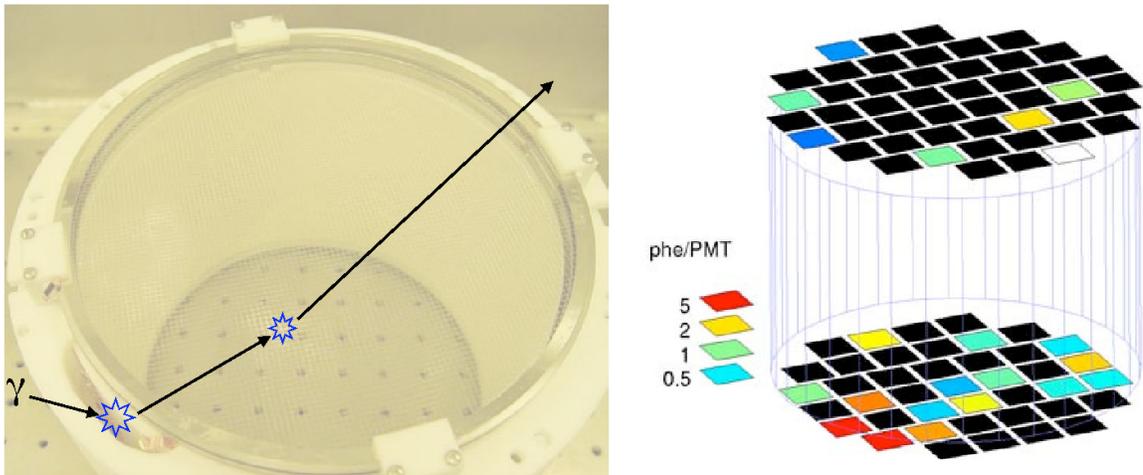


Figure 7.14: Left: An illustration of a Gamma-X event in Xenon10, where the missing-S2 scatter occurs in the hole for the resistor chain setting the voltages on the field shaping rings. Right: The S1 hit pattern from an event identified with this behavior. This figure is courtesy of Peter Sorensen, originally appearing in [55].

(for outer- and resistor-chain region scatters). The primary Xenon10 analysis targeted these events using thresholds on the ratio of S1 light in the top and bottom PMT arrays and on the rms spread in the signal among bottom array PMTs, while a secondary analysis set thresholds on $\frac{S1_i \sqrt{S1}}{S1_{bot}}$ and $\frac{S1_{edge}}{S1}$, where $S1$ is the total S1 signal in photo-electrons, $S1_{bot}$ is the S1 signal in the bottom PMT array, $S1_i$ is the signal in the most-hit PMT, and $S1_{edge}$ is the S1 signal in bottom-edge PMT's. See

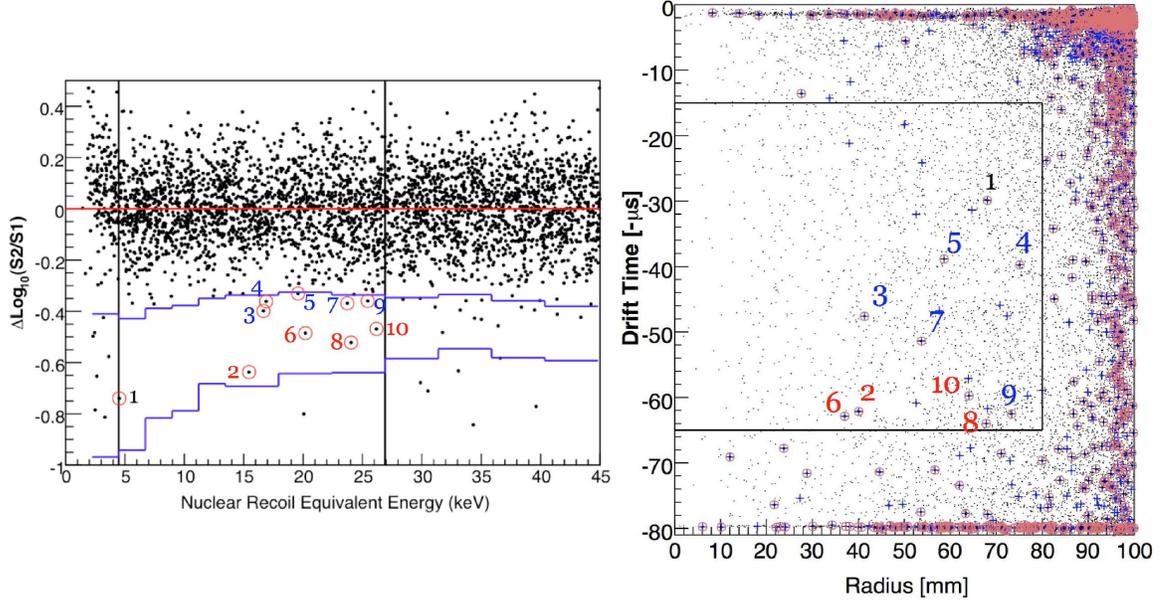


Figure 7.15: The Xenon10 WIMP search results, originally appearing in [31]. On the left, events passing Basic, Quality, and Fiducial cuts are plotted by energy and charge yield, showing ten WIMP candidates in the WIMP acceptance region (these events are discussed in the text). On the right, events passing Basic cuts are shown by position in the detector, with '+'s indicating events in the WIMP acceptance region that were rejected by the Quality cuts, and \oplus 's indicating events in the WIMP acceptance that were not rejected. Most of the WIMP region events (+ 's and \oplus 's) outside the fiducial region towards the wall and bottom of the chamber are attributed to Gamma-X events. The Yellin maximal gap method used by Xenon10 to report a limit picks out the gap between events 1 and 2, where the WIMP signal is expected to be highest, to set the limit.

Figs. 7.13 and 7.14 for example Gamma-X hit patterns from reverse field region and resistor-ladder events.

7.4.3 Calculating a Dark Matter Limit

Once a WIMP acceptance region is defined, a fiducial region chosen, and Basic and Quality Cuts fixed, the WIMP search data is unblinded, giving (possibly) a population of WIMP candidate events (shown in Fig. 7.15 for Xenon10). These events are compared to the expected WIMP recoil spectrum for a given WIMP mass (making assumptions about the WIMP velocity distribution and local density, as in Chapter 2) to produce a limit on the WIMP scattering cross-section.

Limits are generally presented as confidence intervals on the WIMP scattering cross-section as a function of WIMP mass. A confidence interval for a physical parameter σ is defined by a set of intervals $[\sigma_1, \sigma_2]$ corresponding to the possible results of the experiment, where, for any actual value of σ , there is a probability α that the experiment will produce a result giving an interval such

that $\sigma \in [\sigma_1, \sigma_2]$, where α is the confidence level. That is, for any real σ , there is a subset $X^{(\sigma)}$ of experimental results that will give a confidence interval containing σ , and we require

$$\forall \sigma : \sum_{x \in X^{(\sigma)}} P(x|\sigma) = \alpha, \quad (7.6)$$

where x indicates a possible result, and $P(x|\sigma)$ is the likelihood function, or the probability of the result x given the physical parameter σ . Confidence intervals are not unique — there are many choices of $\sigma_1(x)$ and $\sigma_2(x)$ that satisfy Equation 7.6.

If we ignore the energies of the candidate events and the shape of the expected WIMP recoil spectrum, then our ‘result’ is simply the number of events observed, and our confidence intervals are determined by Poisson statistics. Speaking in terms of the expected number of WIMP events ν rather than the cross-section σ , the likelihood for observing n events is

$$P(n|\nu) = \frac{e^{-(\nu+b)} (\nu+b)^n}{n!}, \quad (7.7)$$

where b is the expected number of background events (if known). If one decides a priori to set an upper limit (i.e., $\nu_1 = 0$), then ν_2 is given by: (see [108] for a detailed construction)

$$1 - \alpha = \sum_{k=0}^n \frac{e^{-(\nu_2+b)} (\nu_2+b)^k}{k!}. \quad (7.8)$$

Of course, most experiments also want to claim a discovery if a signal is detected. Feldman and Cousins present a technique for setting confidence intervals without deciding beforehand whether to give a one or two sided interval by choosing intervals based on ratios of likelihoods [108]. For example, in the zero background case with no events detected, Eq. (7.8) gives the 95% confidence limit at 3.00 events, while Feldman and Cousins give 3.09, paying the penalty for the right to choose a two-sided interval had there been a signal. The Feldman-Cousins approach has other advantages as well — Eq. (7.8) can give negative upper limits on ν (resulting in an empty confidence interval) if significantly fewer events are observed than expected in the background. On the one hand, this unphysical limit indicates a poor fit to the background model, but Feldman and Cousins allow one to still produce a physical confidence interval.

If the background is unknown or not well known, as was the case in Xenon10, one must take $b = 0$ in Eq. 7.8 (or in the Feldman-Cousins construction), limiting one’s ability to set a limit in the face of leakage events. If the leakage events have a different energy spectrum than the expected

signal, however, this information may be used to improve the limit, as done by Yellin [90]. The simplest of Yellin’s methods (and the one used by Xenon10) looks at the ‘maximum gap’ between leakage events in the data. Given a set of n WIMP candidate events with energies $\{E_i\}$, there is a corresponding set of ‘gaps’,

$$x_i = \int_{E_i}^{E_{i+1}} dE' N_{M_\chi, \sigma}(E'), \quad (7.9)$$

where i runs from 0 to n , with E_0 the lower bound of the WIMP acceptance region and E_{n+1} the upper bound, and $N_{M_\chi, \sigma}$ is the differential spectrum (in number of events per unit energy) expected for WIMP recoils in the specified acceptance region. In other words, x_i is the number of WIMP events expected in the $(i+1)$ th gap in the WIMP search data. For a given total number of expected WIMP events, $\mu = \sum_{i=0}^{n+1} x_i$, the probability that the largest gap is smaller than x is

$$C_0(x, \mu) = \sum_{k=0}^{\lfloor \frac{\mu}{x} \rfloor} \frac{(kx - \mu)^k e^{-kx}}{k!} \left(1 + \frac{k}{\mu - kx} \right), \quad (7.10)$$

where $\lfloor \frac{\mu}{x} \rfloor$ indicates the largest integer less than or equal to $\frac{\mu}{x}$. Equation (7.10) is derived by Yellin [90]. It has the expected properties, $C_0(0, \mu) = 0$, $C_0(\mu, \mu) = 1 - e^{-\mu}$, and $C_0(x, \mu) = 1$ when $x > \mu$ (the largest possible gap is μ , corresponding to zero events, with likelihood $e^{-\mu}$). It also satisfies the recursion relation (when $x < \mu$),

$$C_0(x, \mu) = \int_0^x dx' e^{-x'} C_0(x, \mu - x'), \quad (7.11)$$

where the terms in the integral are the probability $dx' e^{-x'}$ that the first event occurs after a gap x' and the probability $C_0(x, \mu - x')$ that in the remainder of the acceptance region the largest gap is smaller than x .

Yellin’s maximal gap method is generally used to give an upper limit (single-sided confidence interval). The maximum gap x and total events expected μ are functions of the WIMP mass and cross section (entering in Eq. (7.9)), and the single-ended confidence interval at WIMP mass M_χ for σ is $[0, \sigma_2]$, where $C_0(x(M_\chi, \sigma_2), \mu(M_\chi, \sigma_2)) = \alpha$. For zero measured events ($x = \mu$), this gives $\mu(M_\chi, \sigma_2) = -\ln(1 - \alpha)$, the same as in Eq. (7.8), but when there are leakage events the Yellin method gives the better limit. Yellin also presents the ‘optimal gap’ method in [90], in which gaps containing 1, 2, .. n events are also considered when constructing the limit — the additional spectral information results in a slightly tighter limit. It should also be possible to use Eq. (7.10) or the equivalent in Yellin’s optimal gap method to produce two-sided confidence intervals using

the procedure laid out by Feldman and Cousins for the Poisson case, but this is not widely done. (The strength of the Yellin method is in setting a limit with an unknown background, and if one's background is unknown it is inappropriate to produce a two-sided confidence interval.)

The Xenon10 experiment (and future xenon experiments) benefit greatly from the maximal gap treatment when setting limits. The WIMP signal is strongest at low energies, where both the leakage due to recombination fluctuations (Chapter 5) and the leakage due to Gamma-X events are smallest. In the final Xenon10 data (Fig. 7.15), events 3, 4, 5, 7, and 9 are likely normal electron recoil leakage, while 2, 6, 8, and 10 are likely Gamma-X events. Event 1 appears to be an extremely low energy electron recoil with an electronics glitch that was interpreted by the analysis as an S1. Regardless of these a posteriori explanations, all these events are treated as possible WIMPs when setting a limit. The Yellin method picks the gap between events 1 and 2 to set the cross section limits. There are systematic uncertainties in the size of the gap, due primarily to uncertainties in the nuclear recoil energy scale which enters in Eq. (7.9). The limit shown in Fig. 7.16 for spin-independent interactions takes $\mathcal{L}_{eff} = 0.19$ (with the usual 122 keV photo-absorption reference), and the changes in the limit with subsequent \mathcal{L}_{eff} measurements are given in [79, 80].

7.5 Conclusions on WIMP Hunting

As the Xenon10 experiment has demonstrated, xenon TPCs are a powerful detector technology for direct detection of WIMPs. As these detectors continue to grow in scale, the largest hurdles appear to be purity for charge drift (and the required high recirculation rates) and electron recoil calibration. Xenon10 shows that it is possible to set a strong limit without background subtraction using Yellin's maximum-gap method, but a better understanding of backgrounds will be necessary to claim a discovery. This will require either a high statistics electron recoil calibration throughout the detector volume, or a calibration at the edge of the detector or in a smaller detector coupled with a solid understanding of the fluctuations leading to leakage events, as was put forward in Chapters 4–6. A robust model for and calibration of Gamma-X events will also be necessary.

Future experiments should be performed with potential discovery in mind. If we did not reasonably expect WIMPs to exist, surely 20 years of null results would have discouraged us by now. Theories have adapted over time to the lower limits, but WIMPS remain in the minds of many the primary dark matter candidate, and we are assured by theorists that we are within a factor of 10^2 (or surely 10^3) of either a discovery or a very significant null result, at least for spin-independent interactions. Whether using xenon TPCs or other technologies, dark matter searchers must be prepared

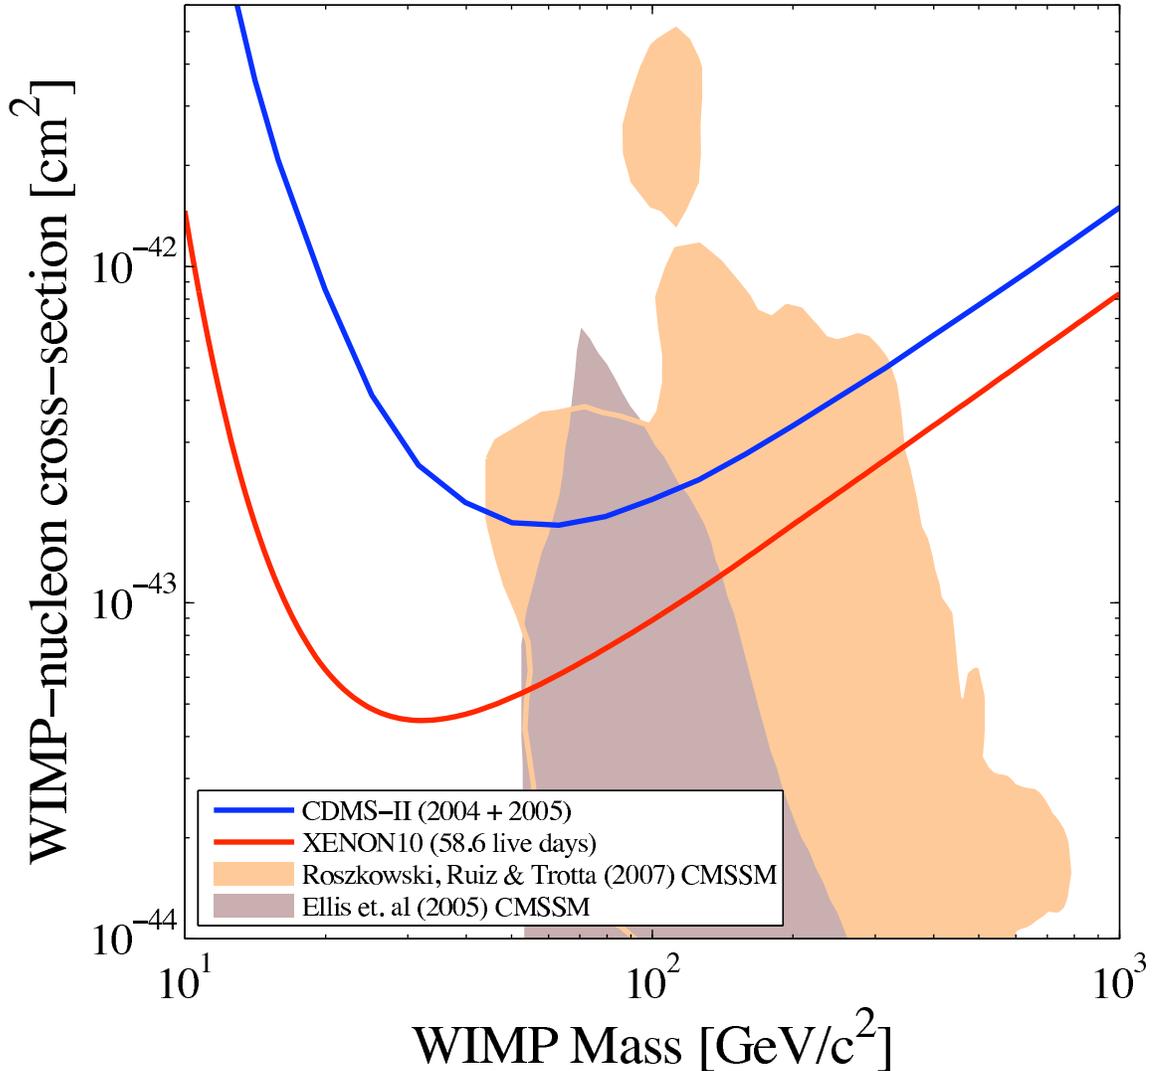


Figure 7.16: The Xenon10 90% C.L. limit on the spin-independent WIMP-nucleon cross section, originally appearing in [31]. The next-strongest limit at the time (CDMS [109]) is also shown. CDMS has since released a new limit, passing the Xenon10 limit for WIMP masses above 40 GeV [30].

at last for a positive signal! Right now it depends who you ask whether, in their heart of hearts, a handful of WIMP events appearing in the detector would be a dream come true or a nightmare. In Xenon10 we were thrilled when we unblinded the data and saw no events in our most sensitive region, and rightly so given the goals of the detector! But now we need to shift modes. When the box is opened and a WIMP signal is seen, we should respond with even greater jubilation — we are not there yet. The aim of this work has been to lay down the physics of background discrimination in liquid xenon, bringing us closer to the understanding we need in order to one day unveil our data and shout “Eureka! We have found it!”

Appendix A

Electrostatics of Wire Grids

Electrostatic simulations of xenon TPCs are generally done using software packages such as Maxwell [110], but for the simple case of parallel wire grids (as in Xed) the analytical calculation is straightforward. We consider first a chamber with two parallel plates separated by a wire grid (see Fig. A.1). The bottom electrode is held at potential V_1 and is a distance b_1 below the grid, and the top is at potential V_2 a distance b_2 above the grid. The grid consists of wires of radius r with spacing a , held at voltage V_0 . By imposing periodic boundary conditions, one can solve for the potential everywhere in the chamber, as done in [111]. At z satisfying

$$\frac{e^{2\pi|z|/a}}{2} \gg 1, \quad (\text{A.1})$$

where $z = 0$ at the wire grid plane, the electric fields are uniform and may be written in terms of the ‘sheet’ potential of the grid, \tilde{V}_0 ,

$$E_1 = \frac{V_1 - \tilde{V}_0}{b_1}, \quad E_2 = \frac{\tilde{V}_0 - V_2}{b_2}. \quad (\text{A.2})$$

The sheet voltage is related to the wire voltage by

$$V_0 = \tilde{V}_0 + \Delta E \frac{a}{2\pi} \ln\left(\frac{a}{2\pi r}\right), \quad (\text{A.3})$$

with $\Delta E = E_2 - E_1$. Given the voltages of the top and bottom plates and the wire grid, one can substitute E_1 and E_2 from Eq. (A.2) into Eq. (A.3) and solve for \tilde{V}_0 .

We generalize to a series of n wire grids with pitch a_i , wire radii r_i , and voltage V_i , spaced

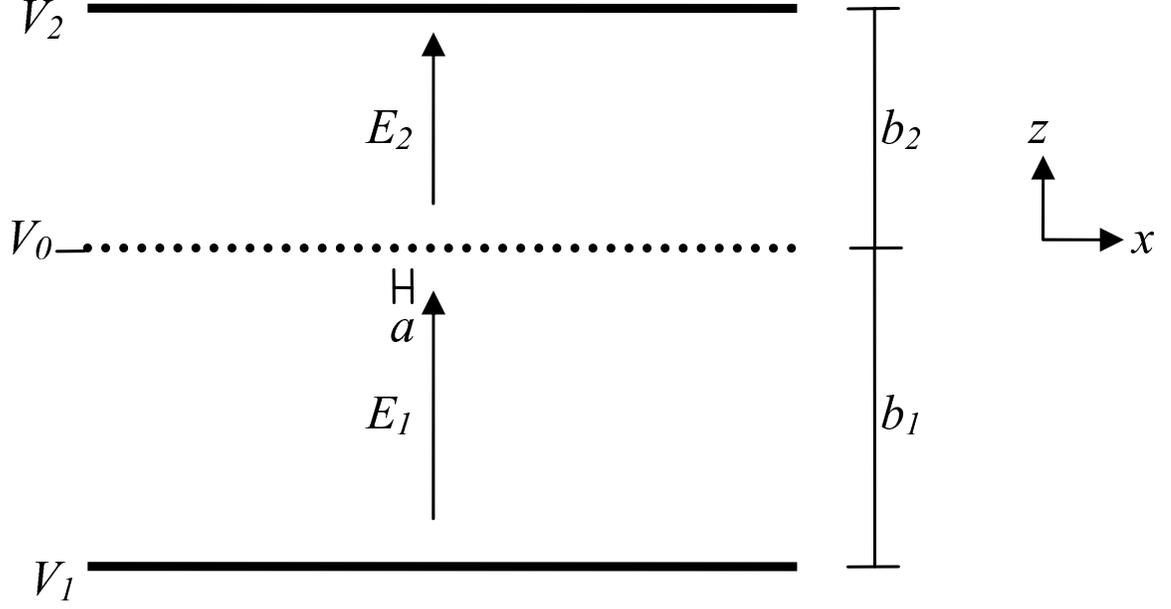


Figure A.1: Wire grid between two parallel plates, with wires perpendicular to the page

between two plates with distances d_i between the i th and $(i + 1)$ th grid (d_0 is the distance between the bottom plate and the bottom grid, and d_n the distance between the top grid and top plate). The problem is to find the fields E_i and sheet voltages \tilde{V}_i given the voltages on each electrode. The above equations now give us

$$E_i = \frac{\tilde{V}_i - \tilde{V}_{i+1}}{d_i}, \quad (\text{A.4})$$

where the E_i 's follow the same numbering as the d_i 's, and

$$V_i = \tilde{V}_i + (E_i - E_{i-1}) X_i, \quad (\text{A.5})$$

where

$$X_i = \frac{a_i}{2\pi} \ln \left(\frac{a_i}{2\pi r_i} \right). \quad (\text{A.6})$$

The bottom and top plates have $X_0 = X_{n+1} = 0$. Substituting Eq. A.4 into Eq. A.5, we have

$$V_i = -\frac{X_i}{d_{i-1}} \tilde{V}_{i-1} + \left(1 + X_i \left(\frac{1}{d_{i-1}} + \frac{1}{d_i} \right) \right) \tilde{V}_i - \frac{X_i}{d_i} \tilde{V}_{i+1}, \quad (\text{A.7})$$

with $V_0 = \tilde{V}_0$ (bottom plate) and $V_{n+1} = \tilde{V}_{n+1}$ (top plate). Thus we can write $V_i = M_{ij} \tilde{V}_j$, summing over repeated indices, where M is a $(n + 1) \times (n + 1)$ matrix containing the coefficients in Eq. (A.7). Inverting the matrix M we get $\tilde{V}_j = M_{ji}^{-1} V_i$, giving us the sheet voltages \tilde{V}_j and thus the electric

fields in terms of the wire grid voltages V_i .

Next we need to add the liquid-gas barrier to the problem. Suppose the liquid level is between the k th and $(k + 1)$ th electrodes, at a fractional distance x . The fields above and below the liquid surface are given by

$$E_{k,liq} = \frac{\tilde{V}_k - V_\ell}{x d_k}, \quad E_{k,gas} = \frac{V_\ell - \tilde{V}_{k+1}}{(1-x) d_k}, \quad (\text{A.8})$$

where V_ℓ is the potential at the liquid surface. Since $E_{k,gas} = \epsilon_r E_{k,liq}$, where $\epsilon_r \approx 1.96$ is the ratio of the dielectric constants in liquid and gas xenon, this potential is given by

$$V_\ell = \frac{\tilde{V}_k \epsilon_r (1-x) + \tilde{V}_{k+1} x}{\epsilon_r (1-x) + x}, \quad (\text{A.9})$$

The liquid surface modifies Eq. (A.5) for V_k and V_{k+1} , the former using $E_{k,liq}$ for E_k , and the latter $E_{k,gas}$. Using Eq. (A.8,A.9) we find the analog of Eq. A.7 for V_k and V_{k+1} , giving new coefficients in the corresponding rows in M . We still simply invert the matrix M to find the sheet voltages in terms of the applied voltages.

There are two other problems of interest, aside from calculating the fields between grids. First, we need the condition to ensure that a grid is transparent to drifting charges. This requires that no field lines end on the grid, or that the surface charge density on the wire have everywhere the same sign as the drifting charge. Using the single-grid setup for ease of notation, the mean charge density around the wire is given by

$$\sigma_0 = \frac{\lambda}{2\pi r} = \frac{\sigma_s a}{2\pi r} = \frac{\Delta E \epsilon a}{2\pi r}, \quad (\text{A.10})$$

where λ is the linear charge density along the wire, σ_s is the average charge density in the wire grid plane, and the last equality is Gauss' law. The background electric field $E_B = \frac{E_1 + E_2}{2}$ adds a sinusoidal charge distribution to the wire,

$$\sigma_B = 2\epsilon E_B \cos \theta, \quad (\text{A.11})$$

where θ is the angle around the wire. For transparency, we require $|\sigma_0| > |\sigma_B|$ so that the surface charge on the wire is all one sign. This gives the transparency condition

$$\left| \frac{\Delta E}{E_B} \right| > \frac{4\pi r}{a}, \quad (\text{A.12})$$

or, for transparency to upward-drifting charge,

$$|E_2| > |E_1| \frac{a + 2\pi r}{a - 2\pi r}. \quad (\text{A.13})$$

The final problem concerns direct charge readout with a charge sensitive preamp attached to one of the grids. The preamp integrates the image current due to the drifting electrons, but a negative image can remain after the electrons are collected due to the slow moving positive ions still in the chamber, reducing the total signal. In Xed we drift electrons past two grids before collecting them to shield the preamp from this positive charge. To calculate the shielding efficiency, we use Green's Reciprocity Theorem,

$$\int dV \rho_1 \phi_2 = \int dV \rho_2 \phi_1, \quad (\text{A.14})$$

where ρ_1 and ρ_2 are two charge distributions resulting in potentials ϕ_1 and ϕ_2 , and the integrals are over all space. To determine the image charge on an electrode from a charge q at \mathbf{x} , we take ϕ_1 and ρ_1 to be the chamber with the i th grid at potential V_i , all other grids grounded, and no free charges. We take ϕ_2 and ρ_2 to be the chamber with all electrodes grounded and a charge q at \mathbf{x} . Since ρ_1 is non-zero only on the grids, and ϕ_2 is zero on the grids, the first integral in Eq. A.14 vanishes. Evaluating the second integral, we find

$$qV_i(\mathbf{x}) + q_i V_i = 0, \quad (\text{A.15})$$

where q_i is the image on the i th grid, and $V_i(\mathbf{x})$ is the potential at \mathbf{x} with the i th grid at V_i . The ratio $-\frac{q_i}{q} = \frac{V_i(\mathbf{x})}{V_i}$ can be calculated from the relevant coefficients of M^{-1} above, since $V_i(\mathbf{x})$ is determined by the sheet voltages of the grids above and below \mathbf{x} . If \mathbf{x} is between the k th and $(k+1)$ th grids, at a height z above the k th grid, then

$$V_i(\mathbf{x}) = \tilde{V}_k \frac{d_k - z}{d_k} + \tilde{V}_{k+1} \frac{z}{d_k}, \quad (\text{A.16})$$

which, since $\tilde{V}_j = M_{ji}^{-1} V_i$, gives

$$-\frac{q_i}{q} = M_{ki}^{-1} \frac{d_k - z}{d_k} + M_{k+1,i}^{-1} \frac{z}{d_k}. \quad (\text{A.17})$$

Appendix B

Maximum Likelihood Histogram Fitting

The fits in Chapters 4 and 5 to peaks and bands in the data are maximum likelihood fits to finely binned histograms, performed using the MATLAB[®] *fminunc* routine. This appendix describes the fit technique.

A maximum likelihood fit is a generalization of the χ^2 fit to data with non-Gaussian error-bars. For a set of physical parameters Θ , the likelihood of a result X is the conditional probability of the result for a given set of parameters,

$$L(X, \Theta) = P(X|\Theta). \quad (\text{B.1})$$

The maximum likelihood estimate for Θ is the set of parameters maximizing $L(X, \Theta)$ for the measured result. In the context of the usual χ^2 fit, the result X is a set of independent measurements x_i , each sampled from a Gaussian distribution with width σ_i and mean μ_i , where the μ_i 's are functions of Θ . The likelihood is then given by

$$L = \prod_i \frac{1}{\sigma_i \sqrt{2\pi}} e^{-\frac{(x_i - \mu_i)^2}{2\sigma_i^2}}, \text{ or} \quad (\text{B.2})$$

$$\ln(L) = - \sum_i \left[\ln(\sigma_i \sqrt{2\pi}) + \frac{(x_i - \mu_i)^2}{2\sigma_i^2} \right] = -\frac{\chi^2}{2} + \text{constant}. \quad (\text{B.3})$$

Thus, finding the maximum likelihood estimate for Θ is same as minimizing χ^2 .

When fitting a histogram, our result is a set of bin counts n_i , each of which follows a Poisson distribution with mean $\nu_i(\Theta)$, giving a likelihood function

$$L = \prod_i \frac{e^{-\nu_i} \nu_i^{n_i}}{n_i!}, \quad (\text{B.4})$$

and the log-likelihood

$$\text{constant} - 2 \ln(L) = \chi_P^2 = \sum_i \begin{cases} 2 \left(\nu_i - n_i + n_i \ln \left(\frac{n_i}{\nu_i} \right) \right) & n_i > 0, \\ 2\nu_i & n_i = 0, \end{cases} \quad (\text{B.5})$$

where the left-hand side of Eq. (B.5) is the equivalent of χ^2 in Eq. (B.3), with the *constant* chosen so that the right-hand side is zero when $\nu_i = n_i$. We call this quantity χ_P^2 , for Poisson chi-square.

When fitting a set of data points to a distribution, it is possible to find the maximum likelihood fit directly without binning the data, but it is often computationally simpler to bin the data and minimize χ_P^2 . The two fits converge as bin size goes to zero. The histogram approach is especially helpful when fitting a sub-region of the data, as when fitting a Gaussian about a peak and ignoring tails — one simply excludes the bins outside the fit region from the sum in Eq. (B.5). Typically we chose bins such that $\langle \nu \rangle \approx 1$. In general, smaller bins preserve more information from the original data and come closer to the unbinned maximum likelihood fits, but rounding errors can be problematic when $\nu_i \ll 1$ (in particular, χ_P^2 blows up when $\nu_i = 0$ and $n_i > 0$) and huge numbers of bins slow down the fit.

Once bins are chosen, we write functions for ν_i in terms of θ_a (where θ_a are the parameters making up Θ). It is also useful to explicitly find $\frac{\partial \nu_i}{\partial \theta_a}$ and $\frac{\partial^2 \nu_i}{\partial \theta_a \partial \theta_b}$. The minimization routine *fminunc* performs better if given the gradient and Hessian of the function being minimized, which in terms of the above derivatives are

$$\frac{\partial \chi_P^2}{\partial \theta_a} = \sum_i \left[2 \left(1 - \frac{n_i}{\nu_i} \right) \frac{\partial \nu_i}{\partial \theta_a} \right], \quad (\text{B.6})$$

and

$$\frac{\partial^2 \chi_P^2}{\partial \theta_a \partial \theta_b} = \sum_i \left[2 \left(1 - \frac{n_i}{\nu_i} \right) \frac{\partial^2 \nu_i}{\partial \theta_a \partial \theta_b} + \frac{2n_i}{\nu_i^2} \frac{\partial \nu_i}{\partial \theta_a} \frac{\partial \nu_i}{\partial \theta_b} \right]. \quad (\text{B.7})$$

Errors on fit parameters may be determined from the Hessian in Eq. (B.7) evaluated at the fit point. As in the normal χ^2 fit, the 1σ contour for fit parameters occurs at $\chi_P^2 = \chi_P^2 + 1$. Approximating χ_P^2 to be a parabola, the error on a one dimensional fit is then

$$\sigma^2[\theta] = 2 / \frac{d^2 \chi_P^2}{d\theta^2}. \quad (\text{B.8})$$

This is equivalent to assuming the likelihood function L is Gaussian in θ and finding the variance. If we have many fit parameters, then Eq. (B.8) holds for the eigenvectors of the Hessian. If these have the normal basis $\tilde{\theta}_b = \sum_a U_{ab} \theta_a$ (where U is unitary), then the error on the original fit parameters is

$$\sigma^2 [\theta_a] = \sum_b \left[(U_{ab})^2 \sigma^2 [\tilde{\theta}_b] \right], \quad (\text{B.9})$$

which is simply the usual error propagation, taking the errors in the $\tilde{\theta}_b$ to be independent.

Appendix C

RIVAL — A Nuclear Recoil Track Monte Carlo

In Chapter 6 we refer to RIVAL, a home-built nuclear recoil track simulator that produces maps of the ionization resulting from nuclear recoils in liquid xenon. The detailed results of this simulation turn out to be unimportant in the context of the chapter, but for reference we document the Monte Carlo here.

RIVAL (Recoiling Ions in Various Atomic Liquids) tracks nuclear recoil cascades, following daughter recoils down to a set tracking threshold, recording energy lost to sub-threshold nuclear recoils and electronic stopping along the way. The inputs to the Monte Carlo include the tracking threshold energy E_{th} , the electronic stopping power as a function of recoil energy $S_e(E)$, and the differential elastic scattering cross section between nuclei, $\sigma(E, E_r)$, where E is the energy of the incoming nucleus and E_r is the energy given to the target nucleus. Nuclear scatters below the threshold energy are not followed, so we define the nuclear stopping power,

$$S_n(E) = \int_0^{E_{th}} dE_r E_r \sigma(E, E_r), \quad (\text{C.1})$$

and the total ‘hard’ nuclear scattering cross-section,

$$\sigma_h(E) = \int_{E_{th}}^E dE_r \sigma(E, E_r). \quad (\text{C.2})$$

Note that we use units of energy \times length² for the stopping power, rather than energy/length. Throughout this discussion, track lengths are multiplied by the number density of nuclei in the

liquid, so that density does not enter into the equations. Include factors of density as appropriate to give normal units.

The Monte-Carlo requires the probability distribution $P(\ell|E)$ of path lengths ℓ to the next hard scatter for a nucleus with energy E , where by ‘hard’ scatter we mean a recoil resulting in another tracked nucleus. To calculate this, we first find the range of the nucleus in the absence of hard collisions,

$$R(E) = \int_{E_{th}}^E dE' \frac{1}{S_{tot}(E)}, \quad (\text{C.3})$$

where $S_{tot}(E) = \sum S_i(E)$ is the total stopping power, or in this case, $S_e(E) + S_n(E)$. The lower limit of the integral in equation (C.3) is our tracking threshold E_{th} — when a recoil reaches this energy, we stop tracking and assign the remaining energy loss to nuclear stopping. The energy of a nucleus that began with energy E and has since traveled distance ℓ without a hard collision is

$$\tilde{E}(E, \ell) = R^{-1}(R(E) - \ell), \quad (\text{C.4})$$

and we have,

$$P(\ell|E) = \sigma_h \left(\tilde{E}(E, \ell) \right) e^{-\int_0^\ell d\ell' \sigma_h(\tilde{E}(E, \ell'))}. \quad (\text{C.5})$$

The exponential is the probability for no hard scatters over the path length ℓ , and the term $\sigma_h \left(\tilde{E}(E, \ell) \right)$ is the probability density for a hard scatter to occur at ℓ . The cumulative distribution function (cdf) $C_l(\ell|E)$, which is the function actually used in the Monte-Carlo, is given by

$$C_l(\ell|E) = 1 - e^{-\int_0^\ell d\ell' \sigma_h(\tilde{E}(E, \ell'))}. \quad (\text{C.6})$$

The cdf will have a maximum value $C_l(R(E)|E) < 1$, indicating that there is a finite probability for the nucleus to fall below the threshold energy (i.e., travel its full range) without undergoing a hard collision. Once the distance ℓ to the next hard scatter has been determined (by taking the inverse-cdf of a random number from 0 to 1), the energy lost along the way due to a given stopping power is

$$\Delta E_i(\ell|E) = \int_0^\ell d\ell' S_i \left(\tilde{E}(E, \ell') \right), \quad (\text{C.7})$$

where the subscript indicates the mode of energy loss. If $\ell = R(E)$ we add E_{th} to ΔE_n . One can verify that $\Delta E_{tot} = E - \tilde{E}(E, \ell)$. Note that we are ignoring any straggling (statistical variations) in the electronic or sub-threshold nuclear stopping powers.

Having taken the recoiling nucleus to the next hard scatter (if there is one), the cdf for recoil

energies is

$$C_r(E_r|E) = \frac{\int_{E_{th}}^{E_r} dE'_r \sigma(E, E'_r)}{\sigma_h(E)}, \quad (\text{C.8})$$

where E is now the energy of the nucleus immediately before the scatter. The recoil energy is generated by taking the inverse-cdf of a random number from 0 to 1. The process now repeats with two recoiling nuclei of energies E_r and $E - E_r$, and so on, generating a tree of recoils. For each branch on the tree we record the ΔE_e and ΔE_n .

The spatial geometry of the tree is also readily calculated. Each segment between hard collisions is taken to be a straight line, and the scattering angle at each collision is determined kinematically. In the center-of-mass frame, the scattering angle between two equal mass objects is

$$\theta_{cm} = 2 \arcsin \left(\sqrt{\frac{E_r}{E}} \right), \quad (\text{C.9})$$

which, with the target ion initially at rest, gives a lab frame scattering angle

$$\theta = \arctan \left(\frac{\sin(\theta_{cm})}{1 + \cos(\theta_{cm})} \right), \quad (\text{C.10})$$

where θ is the angle (from 0 to $\pi/2$) between the initial and final trajectories of the incoming nucleus. The target nucleus recoils at $\pi/2 - \theta$ from the initial trajectory on the opposite side, and there is a random azimuthal angle from 0 to 2π .

Of the above equations, the only input functions used directly by the Monte Carlo are $C_l(\ell|E)$, $C_r(E_r|E)$, $\Delta E_e(\ell|E)$, and $\Delta E_n(\ell|E)$. We create lookup tables for these four functions, using the built-in Matlab numerical integration package. The tables have fixed logarithmic spacing in E , ℓ , and E_r , allowing fast interpolation by the Monte Carlo.

For the results in Chapter 6, we use Eq. (4.10) and (4.11) for the electronic stopping power and differential elastic scattering cross section, and take a tracking threshold of 10 eV. It is worth noting that the procedure laid out above is exactly equivalent to the process described by Lindhard's integro-differential equation for electronic energy loss in a nuclear recoil cascade [37] if we take our threshold energy to zero.

For use in our recombination model, we need a map of ionization sites in the track. We place ions randomly along track pieces, weighted by electronic energy loss. Using the total electronic energy loss to determine the number of ions to place simply reproduces the Lindhard result, which does not match data. Instead, at each simulated recoil energy we assign the measured number of ions at that energy to the average total energy loss in the simulated tracks at that energy. This is fixing the

symptom, not the problem — the problem presumably lies in the stopping power and cross section we have used. We also don't have an actual measurement of the number of ions at a given energy, rather, we have a measurement of the number of ions plus direct excitons. We take $\frac{N_{ex}}{N_i} = 0.06$, as for electron recoils, but as indicated in Chapter 6, this may not be accurate for nuclear recoils. None of these issues affect the conclusions in Chapter 6, where all we require from these simulations is the overall track size, which at all simulated energies much smaller than than the critical box sizes in our recombination model.

Bibliography

- [1] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helvetica Physica Acta*, vol. 6, pp. 110–127, 1933.
- [2] M. Persic, P. Salucci, and F. Stel, “The Universal rotation curve of spiral galaxies: 1. The Dark matter connection,” *Mon. Not. Roy. Astron. Soc.*, vol. 281, p. 27, 1996, astro-ph/9506004.
- [3] A. Refregier, “Weak Gravitational Lensing by Large-Scale Structure,” *Ann. Rev. Astron. Astrophys.*, vol. 41, pp. 645–668, 2003, astro-ph/0307212.
- [4] D. N. Spergel *et al.*, “Wilkinson Microwave Anisotropy Probe (WMAP) three year results: Implications for cosmology,” *Astrophys. J. Suppl.*, vol. 170, p. 377, 2007, astro-ph/0603449.
- [5] D. Clowe *et al.*, “A direct empirical proof of the existence of dark matter,” *Astrophys. J.*, vol. 648, pp. L109–L113, 2006, astro-ph/0608407.
- [6] C. Amsler *et al.*, “Review of particle physics,” *Physics Letters B*, vol. 667, no. 1-5, pp. 1 – 6, 2008. Review of Particle Physics.
- [7] G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees, “Formation of galaxies and large-scale structure with cold dark matter,” *Nature*, vol. 311, pp. 517–525, 1984.
- [8] K. Griest, M. Kamionkowski, and M. S. Turner, “Supersymmetric dark matter above the w mass,” *Phys. Rev. D*, vol. 41, pp. 3565–3582, Jun 1990.
- [9] J. D. Lewin and P. F. Smith, “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil,” *Astropart. Phys.*, vol. 6, pp. 87–112, 1996.
- [10] J. D. Bekenstein, “Relativistic gravitation theory for the modified newtonian dynamics paradigm,” *Phys. Rev. D*, vol. 70, p. 083509, Oct 2004.

- [11] R. A. Knop *et al.*, “New constraints on Ω_M , Ω_Λ , and w from an independent set of eleven high-redshift supernovae observed with hst,” *Astrophys. J.*, vol. 598, p. 102, 2003, astro-ph/0309368.
- [12] D. N. Spergel *et al.*, “First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters,” *Astrophys. J. Suppl.*, vol. 148, pp. 175–194, 2003, astro-ph/0302209.
- [13] S. W. Allen, R. W. Schmidt, and A. C. Fabian, “Cosmological constraints from the x-ray gas mass fraction in relaxed lensing clusters observed with chandra,” *Mon. Not. Roy. Astron. Soc.*, vol. 334, p. L11, 2002, astro-ph/0205007.
- [14] R. A. Knop *et al.* Figure available at <http://www-supernova.lbl.gov/public/papers/knop03/SCP2003SNeCMBClust.pdf>.
- [15] D. J. Eisenstein *et al.*, “Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies,” *Astrophys. J.*, vol. 633, pp. 560–574, 2005, astro-ph/0501171.
- [16] N. H. M. Crighton, J. K. Webb, A. Ortiz-Gill, and A. Fernandez-Soto, “D/H in a new Lyman limit quasar absorption system at $z=3.256$ towards PKS1937-1009,” *Mon. Not. Roy. Astron. Soc.*, vol. 355, p. 1042, 2004, astro-ph/0403512.
- [17] J. M. O’Meara *et al.*, “The Deuterium to Hydrogen Abundance Ratio Towards the QSO SDSS1558-0031,” *Astrophys. J.*, vol. 649, pp. L61–L66, 2006, astro-ph/0608302.
- [18] Image Credit: X-ray: NASA/CXC/CfA/M.Markevitch *et al.*; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe *et al.*; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe *et al.* Image available at <http://chandra.harvard.edu/photo/2006/1e0657/>.
- [19] M. Girardi, G. Giuricin, F. Mardirossian, M. Mezzetti, and W. Boschin, “Optical mass estimates of galaxy clusters,” *Astrophys. J.*, vol. 505, p. 74, 1998.
- [20] E. Kolb and M. Turner, *The Early Universe*. Redwood City: Addison-Wesley, 1989.
- [21] I. Gott, J. Richard *et al.*, “A Map of the Universe,” *Astrophys. J.*, vol. 624, p. 463, 2005, astro-ph/0310571.
- [22] R. Battesti *et al.*, “Axion searches in the past, at present, and in the near future,” *Lect. Notes Phys.*, vol. 741, pp. 199–237, 2008, 0705.0615.

- [23] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter,” *Phys. Rept.*, vol. 267, pp. 195–373, 1996, hep-ph/9506380.
- [24] L. Bergstrom, P. Ullio, and J. H. Buckley, “Observability of gamma rays from dark matter neutralino annihilations in the Milky Way halo,” *Astropart. Phys.*, vol. 9, pp. 137–162, 1998, astro-ph/9712318.
- [25] H. Landsman, “IceCube, the world’s largest dark matter detector,” 2006, astro-ph/0612239.
- [26] M. Boezio *et al.*, “First Results from the PAMELA Space Mission,” 2008, 0810.3508.
- [27] R. H. Helm, “Inelastic and elastic scattering of 187-mev electrons from selected even-even nuclei,” *Phys. Rev.*, vol. 104, pp. 1466–1475, Dec 1956.
- [28] J. Engel, “Nuclear form-factors for the scattering of weakly interacting massive particles,” *Phys. Lett.*, vol. B264, pp. 114–119, 1991.
- [29] R. Bernabei *et al.*, “First results from dama/libra and the combined results with dama/nai,” *Eur. Phys. J.*, vol. C56, pp. 333–355, 2008.
- [30] Z. Ahmed *et al.*, “Search for Weakly Interacting Massive Particles with the First Five-Tower Data from the Cryogenic Dark Matter Search at the Soudan Underground Laboratory,” *Phys. Rev. Lett.*, vol. 102, p. 011301, 2009, 0802.3530.
- [31] J. Angle *et al.*, “First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory,” *Phys. Rev. Lett.*, vol. 100, p. 021303, 2008, 0706.0039.
- [32] E. Behnke *et al.*, “Improved Spin-Dependent WIMP Limits from a Bubble Chamber,” *Science*, vol. 319, pp. 933–936, 2008, 0804.2886.
- [33] K. Abe, “The xmass experiment,” *J. Phys. Conf. Ser.*, vol. 120, p. 042022, 2008.
- [34] D. S. Akerib, M. S. Armel-Funkhouser, M. J. Attisha, C. N. Bailey, L. Baudis, D. A. Bauer, P. L. Brink, R. Bunker, B. Cabrera, D. O. Caldwell, C. L. Chang, M. B. Crisler, P. Cushman, M. Daal, R. Dixon, M. R. Dragowsky, D. D. Driscoll, L. Duong, R. Ferril, J. Filippini, R. J. Gaitskell, R. Hennings-Yeomans, D. Holmgren, M. E. Huber, S. Kamat, A. Lu, and R. Mahapatra, “Exclusion limits on the wimp-nucleon cross section from the first run of the cryogenic dark matter search in the soudan underground laboratory,” *Phys. Rev. D*, vol. 72, p. 052009, Sep 2005.

- [35] G. Angloher *et al.*, “Limits on WIMP dark matter using sapphire cryogenic detectors,” *Astropart. Phys.*, vol. 18, pp. 43–55, 2002.
- [36] V. Sanglard *et al.*, “Final results of the EDELWEISS-I dark matter search with cryogenic heat-and-ionization Ge detectors,” *Phys. Rev.*, vol. D71, p. 122002, 2005, astro-ph/0503265.
- [37] J. Lindhard *et al.*, “Integral equations governing radiation effects,” *Mat. Fys. Medd. Dan. Vid. Selsk.*, vol. 33, no. 10, 1963.
- [38] W. H. Lippincott *et al.*, “Scintillation time dependence and pulse shape discrimination in liquid argon,” *Phys. Rev.*, vol. C78, p. 035801, 2008, 0801.1531.
- [39] P. Benetti *et al.*, “First results from a dark matter search with liquid argon at 87-K in the Gran Sasso underground laboratory,” *Astropart. Phys.*, vol. 28, pp. 495–507, 2008, astro-ph/0701286.
- [40] V. A. Kudryavtsev, “Dark matter experiments at boulby mine,” *Nucl. Phys. Proc. Suppl.*, vol. 118, p. 526, 2003.
- [41] V. Boccone, “Recent updates on the ArDM project: A Liquid Argon TPC for Dark Matter Detection,” 2008, 0810.4490.
- [42] M. Barnabé-Heider, M. D. Marco, P. Doane, M.-H. Genest, R. Gornea, R. Gunette, C. Leroy, L. Lessard, J.-P. Martin, U. Wichoski, V. Zacek, K. Clark, C. Krauss, A. Noble, E. Behnke, W. Feighery, I. Levine, C. Muthusi, S. Kanagalingam, and R. Noulty, “Response of superheated droplet detectors of the picasso dark matter search experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 555, no. 1-2, pp. 184 – 204, 2005.
- [43] F. Aubin *et al.*, “Discrimination of nuclear recoils from alpha particles with superheated liquids,” *New J. Phys.*, vol. 10, p. 103017, 2008, 0807.1536.
- [44] J. Kwong, *The Development of Two-Phase Xenon Dark Matter Detectors*. PhD thesis, Princeton University, 2009.
- [45] “Nist reference fluid thermodynamic and transport properties database,” 2008. Data available at <http://webbook.nist.gov/chemistry/fluid/>.

- [46] V. Atrazhev, A. Berezhnov, D. Dunikov, I. Chernysheva, V. Dmitrenko, and G. Kapralova, “Electron transport coefficients in liquid xenon,” in *Dielectric Liquids, ICDL 2005*, pp. 329–332, IEEE, 2005.
- [47] E. Gushchin, A. Kruglov, and I. Obodovsky *Sov. Phys., JETP*, vol. 55, pp. 650–660, 1980.
- [48] S. S.-S. Huang and G. R. Freeman, “Electron mobilities in gaseous, critical, and liquid xenon: Density, electric field, and temperature effects: Quasilocalization,” *The Journal of Chemical Physics*, vol. 68, no. 4, pp. 1355–1362, 1978.
- [49] E. B. Wagner, F. J. Davis, and G. S. Hurst, “Time-of-flight investigations of electron transport in some atomic and molecular gases,” *The Journal of Chemical Physics*, vol. 47, no. 9, pp. 3138–3147, 1967.
- [50] J. L. Pack, R. E. Voshall, A. V. Phelps, and L. E. Kline, “Longitudinal electron diffusion coefficients in gases: Noble gases,” *Journal of Applied Physics*, vol. 71, no. 11, pp. 5363–5371, 1992.
- [51] A. I. Bolozdynya, “Two-phase emission detectors and their applications,” *Nucl. Instrum. Meth.*, vol. A422, pp. 314–320, 1999.
- [52] E. Gushchin, A. Kruglov, and I. Obodovsky, “Hot electron emission from liquid and solid ar and xe,” *Sov. Phys., JETP*, vol. 82, pp. 1485–1490, 1982.
- [53] J. Angle *et al.*, “3D position sensitive XeTPC for dark matter search,” *Nucl. Phys. Proc. Suppl.*, vol. 173, pp. 117–120, 2007, astro-ph/0609714.
- [54] K. Ni, *Development of a Liquid Xenon Time Projection Chamber for the XENON Dark Matter Search*. PhD thesis, Columbia University, 2006.
- [55] P. Sorensen, *A Position-Sensitive Liquid Xenon Time-Pro jection Chamber for Direct Detec-tion of Dark Matter: The XENON10 Experiment*. PhD thesis, Brown University, 2008.
- [56] A. Bolozdynya, A. Bradley, S. Bryan, C. Dahl, J. Kwong, J. Mock, M. Usowicz, and T. Shutt, “Cryogenics for the lux detector,” *IEEE NSS*, 2009. accepted for publication.
- [57] A. Ferella, *Study of the Electron Lifetime in a dual-phase Xenon Time Projection Chamber for Dark Matter Research*. PhD thesis, Università degli Studi dellAquila, 2006.
- [58] saes[®] MonoTorr[®] PS3-MT3-R/N specifications, available at <http://www.saespuregas.com>.

- [59] L. Ekström and R. Firestone, “Www table of radioactive isotopes.” database version 2/28/99 from URL <http://ie.lbl.gov/toi/index.htm>.
- [60] A. I. Bolozdynya, P. P. Brusov, T. Shutt, C. E. Dahl, and J. Kwong, “A chromatographic system for removal of radioactive Kr-85 from xenon,” *Nucl. Instrum. Meth.*, vol. A579, pp. 50–53, 2007.
- [61] M. Berger, J. Hubbell, S. Seltzer, J. Chang, J. Coursey, R. Sukumar, and D. Zucker, “Xcom: Photon cross sections database nist standard reference database 8 (xgam),” 1998. Data available at <http://physics.nist.gov/PhysRefData/Xcom/Text/XCOM.html>.
- [62] “Evaluated nuclear data file (endf) retrieval and plotting.” Data available at <http://www.nndc.bnl.gov/sigma/>.
- [63] SMARTeC, “Universal transducer interface.” Datasheets available at http://www.smartec.nl/interface_uti.htm.
- [64] E. Aprile, R. Mukherjee, and M. Suzuki, “Ionization of liquid xenon by am-241 and po-210 alpha particles,” *Nucl. Instrum. Meth.*, vol. A307, pp. 119–125, 1991.
- [65] E. Conti *et al.*, “Correlated fluctuations between luminescence and ionization in liquid xenon,” *Phys. Rev.*, vol. B68, pp. 054201–054205, 2003, hep-ex/0303008.
- [66] E. Aprile, K. L. Giboni, P. Majewski, K. Ni, and M. Yamashita, “Observation of Anti-correlation between Scintillation and Ionization for MeV Gamma-Rays in Liquid Xenon,” 2007, 0704.1118.
- [67] T. Doke *Port. Phys.*, vol. 12, p. 9, 1981.
- [68] T. Doke *et al.*, “Absolute scintillation yields in liquid argon and xenon for various particles,” *Jpn. J. Appl. Phys.*, vol. 42, pp. 1528–1545, 2002.
- [69] M. Miyajima *et al.*, “Average energy expended per ion pair in liquid argon,” *Phys. Rev.*, vol. A9, pp. 1438–1443, 1974.
- [70] T. Takahashi *et al.*, “Average energy expended per ion pair in liquid xenon,” *Phys. Rev.*, vol. A12, pp. 1771–1775, 1975.
- [71] T. Doke, A. Hitachi, S. Kubota, A. Nakamoto, and T. Takahashi, “Estimation of fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon,” *Nucl. Instrum. Methods*, vol. 134, p. 353, 1976.

- [72] K. B. Winterbon, P. Simund, and J. B. Sanders, “Spatial distribution of energy deposited by atomic particles in elastic collisions,” *Mat. Fys. Medd. Dan. Vid. Selsk.*, vol. 37, no. 14, pp. 1–73, 1970.
- [73] A. Wambersie *et al.*, “Stopping of ions heavier than helium,” *J. of the ICRU*, vol. 5, no. 1, 2005.
- [74] C. Chasman, K. Jones, H. Kraner, and W. Brandt, “Band-gap effects in the stopping of $ge72^*$ atoms in germanium,” *Phys. Rev. Lett.*, vol. 21, pp. 1430–1433, 1968.
- [75] F. Arneodo *et al.*, “Scintillation efficiency of nuclear recoil in liquid xenon,” *Nucl. Instrum. Meth.*, vol. A449, pp. 147–157, 2000.
- [76] D. Akimov *et al.*, “Measurements of scintillation efficiency and pulse-shape for low energy recoils in liquid xenon,” *Phys. Lett.*, vol. B524, pp. 245–251, 2002, hep-ex/0106042.
- [77] E. Aprile *et al.*, “Scintillation response of liquid xenon to low energy nuclear recoils,” *Phys. Rev.*, vol. D72, p. 072006, 2005, astro-ph/0503621.
- [78] V. Chepel *et al.*, “Scintillation efficiency of liquid xenon for nuclear recoils with the energy down to 5-keV,” *Astropart. Phys.*, vol. 26, pp. 58–63, 2006.
- [79] P. Sorensen *et al.*, “Determination of the scintillation and ionization yield of liquid Xe from the XENON10 experiment,” *Nucl. Inst. Meth.*, 2008, 0807.0459. accepted for publication.
- [80] E. Aprile *et al.*, “New Measurement of the Relative Scintillation Efficiency of Xenon Nuclear Recoils Below 10 keV,” *Phys. Rev.*, vol. C79, p. 045807, 2009, 0810.0274.
- [81] A. Manzur. PhD thesis, Yale University, 2009.
- [82] N. J. C. Spooner *et al.*, “The Scintillation efficiency of sodium and iodine recoils in a NaI(Tl) detector for dark matter searches,” *Phys. Lett.*, vol. B321, pp. 156–160, 1994.
- [83] G. J. Davies *et al.*, “The Scintillation efficiency for calcium and fluorine recoils in CaF-2 and carbon and fluorine recoils in C-6F-6 for dark matter searches,” *Phys. Lett.*, vol. B322, pp. 159–165, 1994.
- [84] A. Hitachi, T. Doke, and A. Mozumder, “Luminescence quenching in liquid argon under charged-particle impact: Relative scintillation yield at different linear energy transfers,” *Phys. Rev.*, vol. B46, pp. 11463–11470, 1992.

- [85] A. Hitachi, “Properties of liquid xenon scintillation for dark matter searches,” *Astropart. Phys.*, vol. 24, pp. 247–256, 2005.
- [86] M. Berger, J. Coursey, M. Zucker, and J. Chang, “Stopping powers and ranges for electrons,” 2005. <http://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>.
- [87] M. Berger, J. Coursey, M. Zucker, and J. Chang, “Stopping powers and range-tables for helium ions,” 2005. <http://physics.nist.gov/PhysRefData/Star/Text/ASTAR.html>.
- [88] J. Ziegler, “The stopping and ranges of ions in matter,” 2008. A collection of tools for calculating the transport of ions in matter, available at <http://www.srim.org/>.
- [89] T. Shutt, C. E. Dahl, J. Kwong, A. Bolozdynya, and P. Brusov, “Performance and fundamental processes at low energy in a two-phase liquid xenon dark matter detector,” *Nucl. Instrum. Meth.*, vol. A579, pp. 451–453, 2007, astro-ph/0608137.
- [90] S. Yellin, “Finding an upper limit in the presence of unknown background,” *Phys. Rev.*, vol. D66, p. 032005, 2002, physics/0203002.
- [91] P. Langevin *Ann. Chim. Phys. VII*, vol. 28, p. 433, 1903.
- [92] G. Jaffé, “Zur theorie der ionisation in kolonnen,” *Ann. Phys. IV*, vol. 42, pp. 303–344, 1913.
- [93] H. Kramers, “On a modification of jaffe’s theory of column-ionization,” *Physica*, vol. 18, pp. 665–675, 1952.
- [94] J. Thomas and D. A. Imel, “Recombination of electron-ion pairs in liquid argon and liquid xenon,” *Phys. Rev.*, vol. A36, pp. 614–616, 1987.
- [95] L. Onsager, “Initial recombination of ions,” *Phys. Rev.*, vol. 54, pp. 554–557, 1938.
- [96] A. Mozumder, “Free-ion yield and electron-ion recombination rate in liquid xenon,” *Chem. Phys. Lett.*, vol. 245, 1995.
- [97] F. Salvat, J. M. Fernandez-Varea, and J. Sempau, *PENELOPE - A Code System for Monte Carlo Simulation of Electron and Photon Transport*. Nuclear Energy Agency, Organisation for Economic Co-operation and Development, 2003.
- [98] S. Hoedel, *Novel Proton Detectors, Ultra Cold Neutron Decay and Electron Backscatter*. PhD thesis, Princeton University, 2003.

- [99] J. F. Ziegler, J. P. Biersack, and U. Littmark, *The stopping and range of ions in solids*. Pergamon Press, 1985.
- [100] U. Fano and W. Lichten, “Interpretation of $ar^+ - ar$ collisions at 50 keV,” *Phys. Rev. Lett.*, vol. 14, pp. 627–629, 1965.
- [101] T. Doke, “Recent developments in liquid xenon detectors,” *Nucl. Inst. Meth.*, vol. 196, pp. 87–96, 1982.
- [102] U. Sowada, J. Warman, and M. de Haas, “Hot-electron thermalization in solid and liquid argon, krypton, and xenon,” *Phys. Rev.*, vol. B25, pp. 3434–3437, 1982.
- [103] T. Koizumi, E. Shirakawa, and I. Ogawa, “Momentum transfer cross sections for low-energy electrons in krypton and xenon from characteristic energies,” *J. Phys. B*, vol. 19, pp. 2331–2342, 1986.
- [104] L. W. Kastens, S. B. Cahn, A. Manzur, and D. N. McKinsey, “Calibration of a Liquid Xenon Detector with Kr-83m,” 2009, 0905.1766.
- [105] M. Adler and J. Gailly, “zlib.” see <http://zlib.net>.
- [106] L. deViveiros. PhD thesis, Brown University, 2009.
- [107] K. Ni *et al.*, “Preparation of Neutron-activated Xenon for Liquid Xenon Detector Calibration,” *Nucl. Instrum. Meth.*, vol. A582, pp. 569–574, 2007, 0708.1976.
- [108] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Phys. Rev. D*, vol. 57, pp. 3873–3889, Apr 1998.
- [109] D. S. Akerib *et al.*, “Limits on spin-independent WIMP nucleon interactions from the two-tower run of the Cryogenic Dark Matter Search,” *Phys. Rev. Lett.*, vol. 96, p. 011302, 2006, astro-ph/0509259.
- [110] Ansoft, “Maxwell,” 2002. see <http://www.ansoft.com/maxwellSV/>.
- [111] K. T. McDonald, “Notes on electrostatic wire grids.” Internal note, available at <http://puhep1.princeton.edu/mcdonald/examples/grids.pdf>, 2003.