

12-18-2015

Design, Construction, and Optimization of the Energy Resolution of a Xenon TPC Detector

Nicholas E. Destefano

University of Connecticut, nicholas.destefano@uconn.edu

Follow this and additional works at: <https://opencommons.uconn.edu/dissertations>

Recommended Citation

Destefano, Nicholas E., "Design, Construction, and Optimization of the Energy Resolution of a Xenon TPC Detector" (2015). *Doctoral Dissertations*. 986.
<https://opencommons.uconn.edu/dissertations/986>

Design, Construction, and Optimization of the Energy Resolution of a Xenon TPC Detector

Nicholas Evan Destefano, PhD

University of Connecticut, 2015

PIXeY (Particle Identification in Xenon at Yale) is a two-phase (liquid/gas) xenon prototype detector with 3-kg active mass. The two-phase xenon technology has many applications that include gamma-ray imaging, neutrinoless double beta decay searches, dark matter searches, and medical imaging PET scanners. PIXeY was built to optimize energy resolution, with a number of technological improvements over previous work. Parallel-wire grids, which control the drift and proportional-scintillation fields, are optimized both for light-collection efficiency and field uniformity with optical transparencies between 92% and 97%. High voltage epoxy feedthroughs have been installed with stable operating voltages better than 5 kV on the anode and -10 kV on the cathode. High quantum efficiency Hamamatsu R8778 PMTs (photomultiplier tubes), high-reflectivity Teflon walls, and charge-light anti-correlation techniques are also incorporated.

A number of fiducial cuts and correction factors were used to optimize energy resolution. The light and charge signals were corrected by the spatial location of the event within the detector, taking into account effects such as the electron lifetime, and geometric light collection. The energy spectrum of various sources was studied by varying the cathode, anode, and PMT voltages. Optimal configurations for the drift and scintillation fields were found for energies ranging from 41.5 keV (^{83m}Kr) to 2.61 MeV (^{228}Th), resolving the light signal and keeping the charge signal unsaturated. After quality cuts and optimization of parameters, energy resolutions of $5.4 \pm 0.4 \% \sigma/E$ at 42 keV, $1.81 \pm 0.13 \% \sigma/E$ at 350 keV, $1.71 \pm 0.12 \% \sigma/E$ at 511 keV, $1.16 \pm 0.08 \% \sigma/E$ at 662 keV, $1.01 \pm 0.07 \% \sigma/E$ at 1275 keV, and $0.85 \pm 0.06 \% \sigma/E$ at 2614 keV for 1,000 V/cm drift field were obtained. In addition, the recombination fluctuations at the source energies were studied, with the recombination fluctuations modeled as a binomial variance with an additional prefactor that is a function of the number of ion quanta, which appears to plateau at $E_r = 527 \pm 137$ keV.

Design, Construction, and Optimization of the Energy Resolution of a Xenon TPC Detector

Nicholas Evan Destefano

B.S., University of Connecticut, **2007**

M.S., University of Connecticut, **2010**

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

at the

University of Connecticut

2015

Copyright by
Nicholas Evan Destefano

2015

ii

APPROVAL PAGE

Doctor of Philosophy Dissertation

Design, Construction, and Optimization of the Energy Resolution of a Xenon TPC Detector

Presented by

Nicholas Evan Destefano, B.S., M.S.

Major Advisor _____
Moshe Gai

Associate Advisor _____
Daniel McKinsey

Associate Advisor _____
Peter Schweitzer

University of Connecticut
2015

Acknowledgements

Whenever I've read anything with an acknowledgements section, I always think "who reads this?", although I have been known to skim them for curiosity's sake. Now I realize that this is for them, not necessarily for the outside audience.

At any rate, I'd like to first and foremost thank my advisors Moshe Gai, Dan McKinsey, and Peter Schweitzer as well as my committee members Barry Wells and Alex Kovner for guiding me through this process. I'd also like to thank Ed Eyler for taking the risk on me all those years ago when I was first starting out.

Thanks to all of my co-workers and colleagues who helped me grow through this period, particularly all of my quasi-advisors Ethan Bernard, Chris Wahl, and Blair Edwards. I wish everyone good luck on all of your future projects.

I'd like to thank Vinny and Russ Bernardo and everyone else at the Yale Physics Machine Shop. They made my designed parts a reality, and if it wasn't for them I would not have learned how to design machine parts that can actually be fabricated.

I'd like to thank my friends, in particular Jesse, Jared, Jonathan, Mike, Dan, Colin, Brian, Dave, Nicole, and Isabella for help, emotional support, and enabling my board and video game habits. I've known some of you for a long time and your friendship means a lot to me.

I'd like to thank my wife Shree for support and allowing me the time to finish everything once and for all. I'd also like to thank her parents for all of their help in taking care of my children so that I can finish my studies and start my career.

Last but not least I'd like to thank my parents as well both for the recent support with my children but also for raising me and such.

Work supported by ARI grant CBET-1039206, DHS grant 2011-DN-007-ARI056-02, DOE grant DE-SC0004599 and DE-FG02-94ER40870.



When will you be Dr. Daddy? – Leela 2015

Contents

1	Introduction.....	1
1.1	Comparison with Other Materials.....	1
1.2	Applications Overview	3
1.2.1	Dark Matter.....	3
1.2.2	Neutrinoless double beta decay.....	6
1.2.3	Compton Imaging	12
1.2.4	Portal Monitoring.....	17
1.2.5	PET Scanners	20
1.3	Thesis Outline	22
2	The PIXeY Detector	24
2.1	Principle of Detector Operation.....	24
2.2	Hardware Design of PIXeY	26
2.3	Electrostatic Grids.....	29
2.3.1	Grid Introduction.....	29
2.3.2	Grid Simulation.....	29
2.3.3	Mesh Grid Construction.....	33
2.3.4	Single Wire Grid Construction	35
2.4	High Voltage Design Components	40
2.4.1	Introduction.....	40
2.4.2	The Anode Hardware	40
2.4.3	The Cathode Hardware	44
2.5	Experimental Setup.....	46
2.5.1	Signal Chain.....	46
2.5.2	Data Acquisition and Triggering.....	46
2.5.3	Procedure	47
3	Analysis and Acquisition Code – PIXeYCalc	49
3.1	Introduction.....	49
3.2	Baseline Determination.....	51
3.2.1	First n-Points.....	51
3.2.2	All Non-Pulse Samples	53
3.2.3	Other Considerations and Future Improvements	53

3.3	Pulse Finding	56
3.3.1	Sum Waveform	56
3.3.2	Waveform by Waveform	57
3.3.3	Double Box Method (Sum).....	58
3.3.4	Threshold Testing	58
3.4	Identifying S1 and S2.....	61
3.5	Pulse Declustering	63
4	Detector Diagnostics	66
4.1	Introduction.....	66
4.2	PMT Calibrations.....	66
4.3	LED Saturation Tests	73
4.3.1	Introduction.....	73
4.3.2	Pulse Linearity Studies.....	74
4.3.3	Frequency-Based Saturation	77
4.3.4	Large Pulse Test.....	80
4.4	Measurement of the liquid level.....	86
4.5	Lifetime Measurements.....	88
4.6	S1 vs. Depth Dependence	91
4.7	Quality Factors.....	92
4.7.1	Q-Factor	92
4.7.2	R-Factor	93
4.7.3	TB-Factor	95
4.8	Saturation Plots	97
4.8.1	Introduction.....	97
4.8.2	Saturation per PMT	97
4.8.3	Determining Optimal PMT gains.....	99
4.9	Tuning Gas Field Region	101
4.9.1	Introduction.....	101
4.9.2	Anode Studies with ^{83m}Kr	101
4.9.3	Anode studies with ^{137}Cs	104
4.10	Position-Based Quality Cuts	105
4.10.1	Introduction.....	105
4.10.2	Planar Quality Cuts	105

4.10.3	Depth Quality Cut	108
5	Spectrum Optimization – PIXeYOptimizer	111
5.1	Introduction.....	111
5.2	Exploratory Studies.....	112
5.2.1	Introduction.....	112
5.2.2	Light Collection Profile Investigation.....	112
5.2.3	Light Collection Profile Simulations	118
5.3	Program Design	122
5.3.1	PIXeYIterator.....	122
5.3.2	Optimization Algorithm.....	124
5.3.3	Overview of the Optimization Process	125
5.3.4	Auto-binning and Peak Fitting Studies	127
5.4	Results.....	129
5.4.1	Visual Optimization Examples	129
5.4.2	Concluding Remarks.....	131
6	Energy Resolution Measurements	133
6.1	Introduction.....	133
6.2	Position Dependence.....	134
6.2.1	Introduction.....	134
6.2.2	No Quality Cuts	134
6.2.3	Reduced Datasets	137
6.3	Electric Field Dependence	140
6.3.1	Reduced Datasets	140
6.4	Comparison with Other work.....	141
6.5	Light and Charge Yields	143
6.5.1	Calculation of g_1 and g_2	143
6.5.2	N_γ , N_e , and N_q	146
6.5.3	Mean Recombination Probability	148
6.6	Detector and Recombination Fluctuations	149
7	PIXeYCI – A Liquid Xenon Compton Imager	156
7.1	Introduction.....	156
7.2	Wire Readout	158
7.2.1	Introduction.....	158

7.2.2	Amplifier Design.....	161
7.2.3	Test Boards	162
7.2.4	Noise Analysis	163
7.2.5	Simulated Gain Calculation	165
7.3	Projected Performance	167
Appendix.....		169
A.1.	Theoretical Photon Yield Calculation.....	169
A.1.1.	One dimensional derivation	169
A.1.2.	Generalization to Three dimensions.....	171
A.2	Determination of the Electric Field in the Gas Gap	174
A.3.	Estimating the Electric Field on the Cathode Wire.....	176
A.4.	Uniqueness of Energy Resolution Minimization	178
A.5.	Lemma for Energy Resolution Uniqueness	179
Bibliography		182

1 Introduction

With the threat of nuclear proliferation, the advancement of science and technology, and improvements in the medical sciences there is a need for radiation detection, and more specifically, gamma-ray detection in today's world. Previous technology relied on raw counting, with secondary scans to confirm the presence of radiation. Today's technology is much more advanced, with detectors having energy resolutions better than 1% [1]. In the process, two pathways have developed. One pathway has large, bulk detectors that are inexpensive but have poor properties and the other has smaller detectors with excellent performance but are typically expensive.

Over the years, liquid xenon has been developing as a useful detector medium for a variety of purposes. With its high Z, high stopping power, and good energy resolution [2], it has many properties that make it useful as a detector medium. These purposes include dark matter searches [3], neutrinoless double beta decay searches [4], gamma ray detection [5], medical imaging [6], and homeland security applications [7]. These various technologies and applications will be summarized. It will then be proposed how liquid xenon could fit into these applications. With the emphasis of good energy resolution for efficient operation, this will then lead into PIXeY, a liquid xenon R&D detector.

1.1 Comparison with Other Materials

Detector materials are typically placed in two broad categories according to their cost: those that are relatively inexpensive (tens to hundreds of dollars per kg) and those that aren't (thousands of dollars per kg). Typically the cheap detectors are ones that are bulky and can be made very large, like plastic and/or liquid scintillators. The purpose of their operation is for bulk radiation detection. As shown in Table 1.1 however, their energy and position resolution is poor. Most of these detectors are used for raw

radiation counting and sometimes utilize energy windows, such as 8 windows in a 0-3 MeV range [8].

Isotope matching or position reconstruction to determine more information about the radioactive material typically isn't feasible.

	Energy resolution (σ @ 662 keV)	Position resolution (cm)	Density (g/cc)	Cost (\$/kg)	Detector element thickness (cm)	n/gamma discrimination
Plastic scintillator	Poor	~10	1	~30	~20	No
Liquid scintillator	17%	~10	0.9	~30	~20	Yes
Sodium Iodide	5%	~ few cm	3.7	~800	~10	No
CZT	0.33%	~0.1	5.8	~50,000	~1	No
Germanium	0.17%	~0.3	5.3	~20,000	~5	Yes
Liquid xenon	2.5-5%	~0.1	3	~1,000	~20	Yes

Table 1.1: Different detector materials comparing their performance characteristics

Large detectors are most often used today for portal monitoring and other bulk radiation applications [9]. This is used for basic radiation assessment; secondary scans must be frequently employed for radioactive sources that turn out to be benign. This will be gone into further detail in Section 1.2.4.

On the other end of the spectrum are the small, expensive detectors like CZT (Cadmium Zinc Telluride) and Germanium. These detectors are solid state crystalline detectors and have excellent performance in terms of energy and position resolution [1]. However, because of their cost and difficulty in manufacturing [10], the scalability of such a detector is poor. Typically these types of detectors are used in handheld applications, a small detector that's used to assess an area by having a backpack or small device.

Liquid Xenon is a type of detector material that has advantages from both application spaces. It is a relatively inexpensive detector material that combines excellent position reconstruction with the energy resolution that's better than a NaI detector [11]. Position reconstruction is particularly important for Compton imaging detectors, which will be discussed in Section 1.2.3. In addition, because of the process of how the scintillation light is created the material is largely transparent to its own light [12], which means it has good scalability.

Liquid xenon fills an application space where a detector must be large but also where energy and position resolution are important metrics. Throughout this section I will go over various technology areas in the basic sciences and in applied spaces and argue why such a gamma-ray radiation detector is needed.

1.2 Applications Overview

1.2.1 Dark Matter

Astrophysical anomalies and observations hint of the presence of dark matter. This was first observed by comparing galactic rotation curves to theoretical predictions. It was discovered in the 1970s that the galaxies were spinning too fast for the amount of material, which means that all galaxies should be unstable. Early on it was hypothesized that there was some missing matter holding galaxies together. This pioneering work was done in 1932 by Oort [13] and more precisely in 1937 by Zwicky [14]. Since then there has been a gathering of more evidence hinting at dark matter, including big bang nucleosynthesis (BBN), anisotropies in the cosmic microwave background (CMB), baryonic acoustic oscillations (BAO), formation of large structures, mass discrepancies in galactic rotation curves, and distortions in gravitational lensing.

WIMP dark matter is the leading candidate to explain the observations due to the nature of the observations: it does not couple to electromagnetic radiation. Theory predicts that such a particle could

interact via the weak interaction [15]. This missing mass cannot be due to neutrinos, because even though they have mass they are too light to account for the observed small-scale structure [16].

A better survey of the various experimental evidences for Dark Matter has been described elsewhere [17]. Instead what will be done here is a brief introduction to a LXe TPC detector currently searching for WIMP dark matter: LUX.

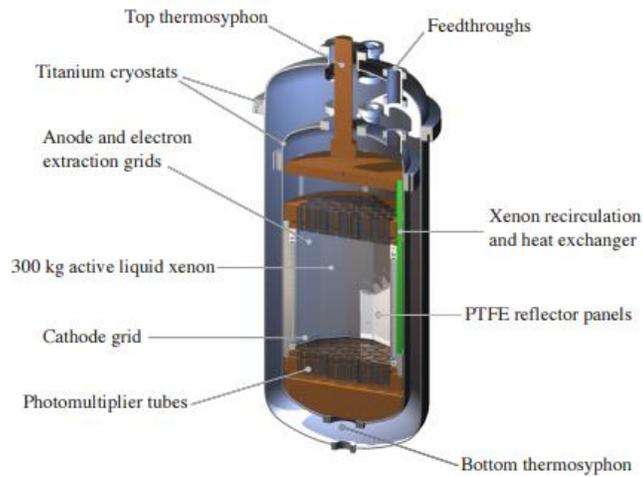


Figure 1.1: Annotated cross-sectional view of the LUX detector.

LUX is a large liquid xenon time projection chamber that was built, commissioned, and has been in operation since 2015 and searching for WIMP dark matter. It has a 370 kg of liquid xenon in its target region and is installed in the Sanford underground mine in Lead, South Dakota. Its main experiment is a rare event search, which means that even though the detector is very large only a few events per run are expected. Experiments typically run for 1-2 years, the DM particle is expected to interact only rarely with the liquid xenon medium due to it interacting only through the weak interaction.

Since a DM event interaction is so rare, great effort is expended to shield the detector from background events and fabricate the detector with low radioactive materials, which would otherwise limit the efficiency of operation. Large trial runs are performed simply to better understand the background

rates in the detector, and there are sophisticated pulse identification techniques to try to remove radioactive background from search runs, such as being able to discriminate between gamma-like events and dark matter neutron-like events [3]. These detectors have largely been successful due to the liquid xenon purification combined with the ability to see very small signals in the proportional scintillation region.

If the detector fails to detect a DM signal the experiment still produces useful results. If the background rates of the detector are well known then the parameter space of the hypothetical DM particles can be constrained by the failure to see any DM signals. This is known as an exclusion curve, which excludes cross-sections greater than the values found since that “should” have been found while the regions below have yet to be probed.

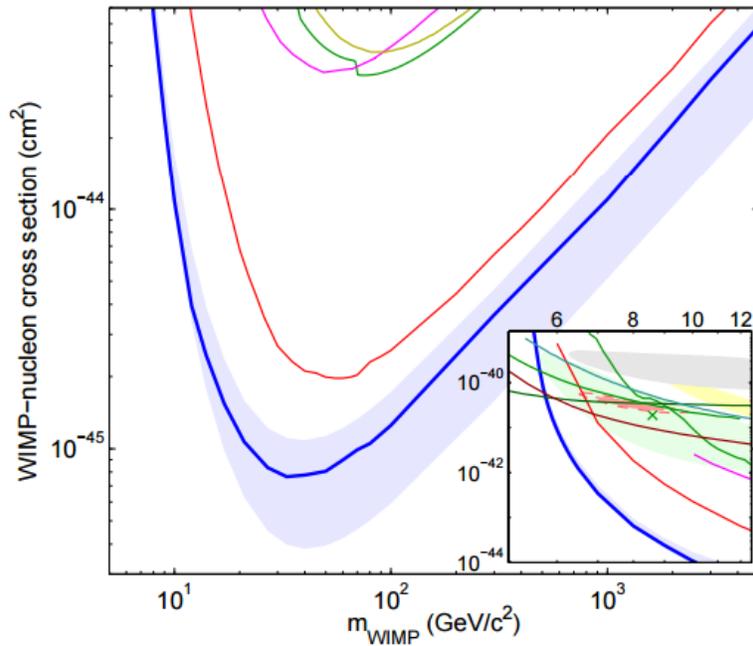


Figure 1.2: Exclusion curve set by the LUX Detector. This exclusion curve is for spin-independent WIMPs, with LUX in the blue band. The inset shows the various limits for low mass WIMSP (<12 GeV). The grey area is the DAMA result [18] and shows the exclusion of this result by various other detectors. Data from the 2014 science run [3].

Liquid xenon has put the best cross-sectional limits on spin independent dark matter over a large mass range. Energy resolution in these detectors is important mainly for establishing the precise location for the boundaries for the exclusion plots.

The success of liquid xenon in the WIMP dark matter sector is what spurred this project into researching other viable options for this technology. It is from here that the idea of using liquid xenon for other application spaces was first considered.

1.2.2 Neutrinoless double beta decay

Liquid xenon detectors are also currently being used in neutrinoless double beta decay searches (hereafter called $0\nu\beta\beta$). Like dark matter, $0\nu\beta\beta$ searches are rare event searches, except on a higher energy scale. For dark matter the energy range of interest is typically ~ 100 eV electron equivalent (eVee) to 10 keVee, while for $0\nu\beta\beta$ the energy range of interest is 2-3 MeVee. Successful observations of $0\nu\beta\beta$ would be the discovery of a new type of radioactive decay, separate from beta decay and even double beta decay.

In normal beta decay, an unstable atomic nucleus spontaneously emits an electron and electron anti-neutrino from the nucleus, converting a neutron into a proton. This process is well-known and has been observed by a number of atomic isotopes. Building off of that model there is a second type of decay wherein this happens twice simultaneously in the same nucleus, so two neutrons are converted into two protons and two electrons and two electron anti-neutrinos are emitted. This process, while theoretically proposed in 1935 by Maria Goeppert-Mayer [19], it has only been recently observed for Ca-48, Ge-76, Se-82, Zr-96, Mo-100, Cd-116, 128-Te, 130-Te, 136-Xe, 150-Nd, and 238-U [20].

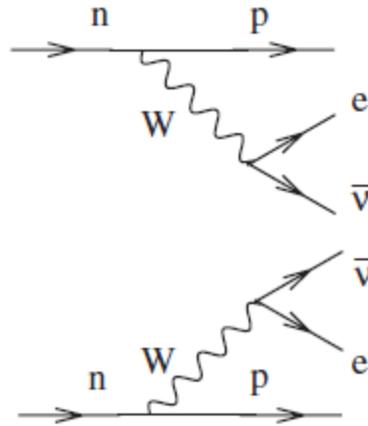


Figure 1.3: Diagram of 2-neutrino double beta decay. The emitted W bosons decay into an electron and electron anti-neutrino pair, much like in ordinary beta decay.

Interestingly, the decay model for this process is unchanged if the neutrino is its own anti-particle, a Majorana particle, named after the first physicist that proposed it [21]. But if that's true, then it's possible that the two neutrinos in two neutrino decay can annihilate one another via a virtual exchange, which is a process called neutrinoless double beta decay. As the name implies, this type of decay would be similar to the 2-neutrino double beta decay, except no neutrinos are emitted, which means that the two electrons carry the full energy of the Q-value of the decay.

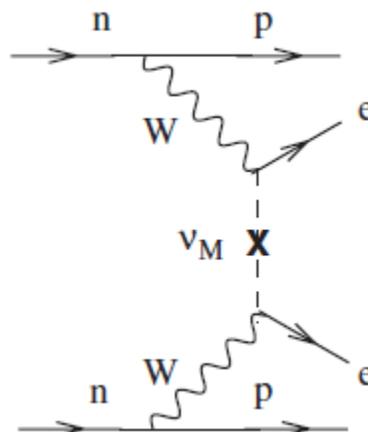


Figure 1.4: Diagram of neutrinoless double beta decay. The neutrinos annihilate via virtual exchange, which can only happen if the neutrino is its own anti-particle

If such a process was experimentally observed, not only would it answer a basic science question about the nature of neutrinos, but it also would provide a way to constrain the absolute mass of the neutrinos. It was discovered that neutrinos have mass [22], but so far projects that seek to determine the masses of the neutrinos can only measure the mass differences of neutrinos [23]. Since only the differences are known, this leads to two paradigms or “hierarchies” as a model of how the masses are ordered.

It is currently unknown which hierarchy is correct for neutrinos. However if neutrinos are Majorana particles then the value of electron neutrino mass can be probed. This is ascertained from the half-life of the neutrinoless double beta decay, which allows the effective double beta decay mass to be calculated. This double beta decay mass places limits on the absolute value of the neutrino masses, not the mass differences.

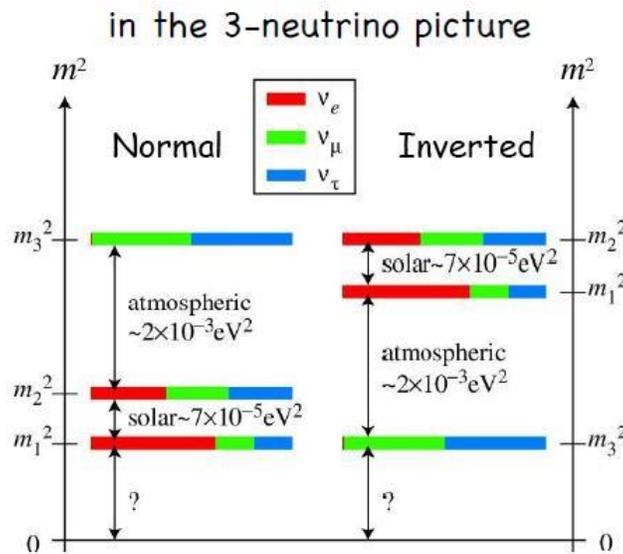


Figure 1.5: Different hierarchical models of neutrinos. Observations of solar and atmospheric neutrinos measure the m_1 to m_2 and the m_2 to m_3 mass splittings respectively [24].

The equation for the half-life of neutrinoless double beta decay is [25]:

$$\left(T_{\frac{1}{2}}^{0\nu}\right)^{-1} = C_{mm} \frac{\langle m \rangle^2}{m_e^2} \quad (1.1)$$

with m_e being the mass of the electron, $\langle m \rangle$ is the Majorana mass and C_{mm} is the nuclear matrix element, which depends on the model. The Majorana mass can then be used to probe the individual neutrino masses:

$$\langle m \rangle = |m_{ee}^1| + e^{i\phi_2}|m_{ee}^2| + e^{i\phi_3}|m_{ee}^3| \quad (1.2)$$

where m_{ee}^1 , m_{ee}^2 , and m_{ee}^3 are the mass eigenstates of the neutrino, and the phi terms are the relative Majorana phases connected with CP violation. Thus with these two equations one is able to probe the types of masses that the neutrinos can have.

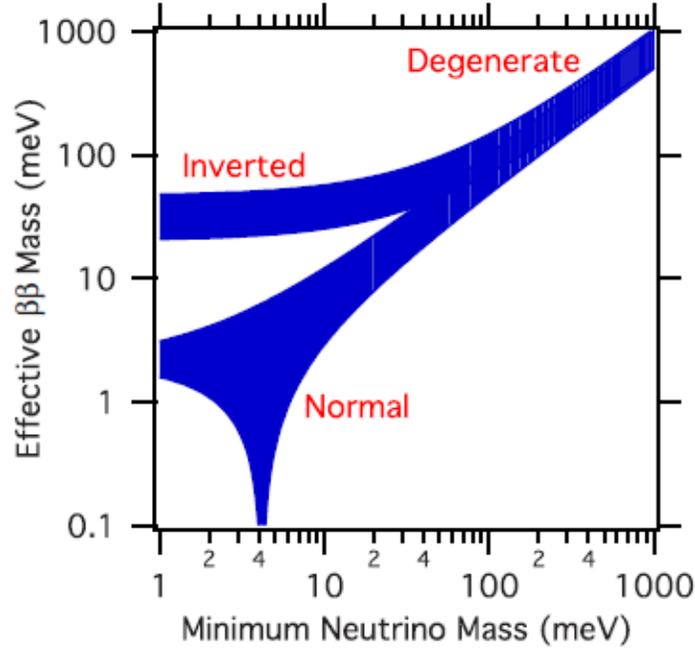


Figure 1.6: Allowed neutrino masses for the different mass hierarchies.

There has been some recent success in finding two-neutrino double beta decay using liquid xenon [4] [26]. It is lucky that xenon has a two neutrino double beta decay mode, since xenon is a good detector medium for radiation spectroscopy which allows this two neutrino double beta decay mode to be characterized. Since the neutrino can carry some energy with it away from the decay, these sorts of experiments look for broad distributions that are characterized by two-neutrino double beta decay. However, for neutrinoless double beta decay there are no emitted neutrinos so the electrons carry all the energy associated with the decay. This means that the ideal signal would be a delta function at the Q-value of decay, which for xenon is 2.46 MeV.

However, finite detector resolution blurs the peak at the Q-value, and there are also many background signals in that region including gamma-rays from ^{232}Th , ^{40}K , ^{60}Co , ^{65}Zn , and ^{238}U which are close to the Q-value of decay. These two factors limit the performance of the detector, which in turns limits the quality of measurement being performed. Since $0\nu\beta\beta$ is a rare event search, this makes for a less efficient detector that would take more time to reach the same half-life sensitivity.

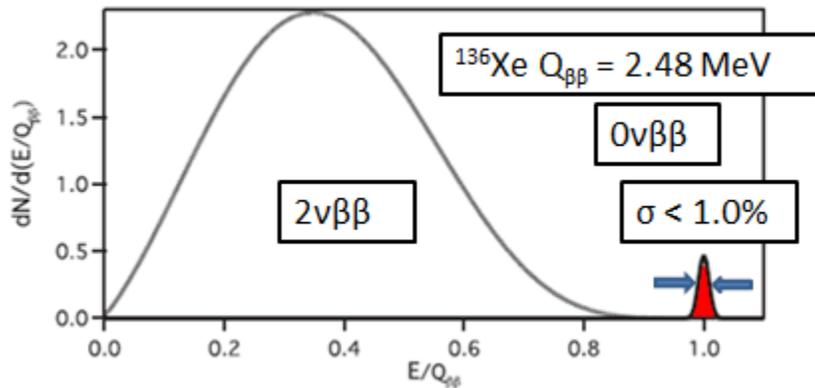


Figure 1.7: Schematic spectrum of a $0\nu\beta\beta$ signal. The two neutrino signal is an unavoidable signal and the neutrinoless signal would be a spike at the Q-value of decay.

While today's detectors have sufficient energy resolution to distinguish the two-neutrino signal from the hypothetical neutrinoless signal, background and background/signal discrimination remains a

challenge. This means that energy resolution enhancement are important to the study of neutrinoless double beta decay, since the energy resolution at the Q-value would improve half-life sensitivity by $\sqrt{\sigma}$, where sigma is the energy resolution, for a statistical process. This would aid in a leading to discovery or at the least constrain the theoretical lifetime and hence the range of neutrino masses possible, since better discrimination would be possible between background sources and the signal at the Q-value.

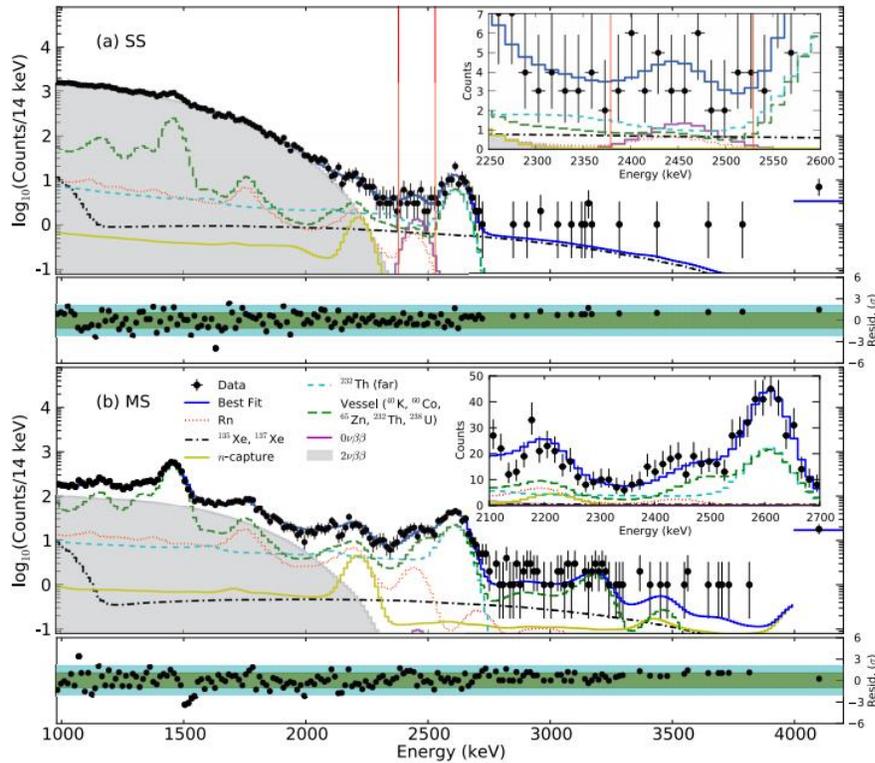


Figure 1.8: The case for energy resolution. The following is data taken from Exo200 during its run for single scatters (top) and multiple scatters (bottom). The energy resolution of the Exo200 detector near the Q-value is about 1.5%; the energy resolution improvement would directly improve sensitivity (red bars) because it would allow for better background rejection.

LXe TPCs have been able to constrain the parameter space of $0\nu\beta\beta$. As shown in Figure 1.9, each improved xenon experiment is pushing the Xe limit to the right on this figure and each improved Ge experiment pushes the limit up. The grey region indicates the controversial Klapdor claim of discovery [27] with the current black line limit set by Gerda [28]. The diagonal lines are the various matrix element

models that would give rise to different values for the germanium and xenon $0\nu\beta\beta$ half-lives. As the performance of these detectors and new detectors are improved, various models and the Klapdor claim can be tested.

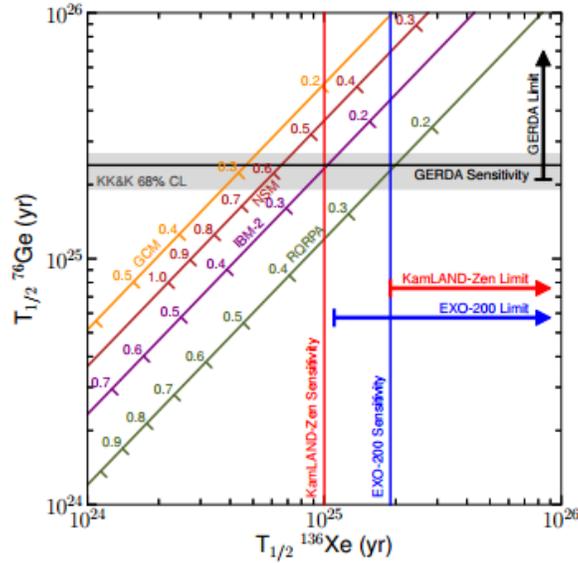


Figure 1.9: Modern exclusion limits set with the Exo200, Kamland-Zen, and GERDA experiments [4].

1.2.3 Compton Imaging

Compton imaging as a technique will be introduced, which can be applied to many other applications. The LXeGrit detector will be introduced, which was mainly used for astrophysical observations, as a starting off point that first used the Compton imaging concept for LXe TPCs. This will then be broadened to other application spaces since Compton imaging enables many of these technologies.

Compton imaging is possible because, as a detector material, gamma-rays may sometimes deposit their energy into multiple locations. Typically gamma-rays in the energy range of few keV to 2-5 MeV either deposit their energy into a single site photoabsorption or through multi-site interactions known as Compton scattering. The physics is well-known for connecting these multi-scatter sites to the original

gamma-ray energy and trajectory, an analysis that is not possible if the gamma-ray deposits its energy in a single location.

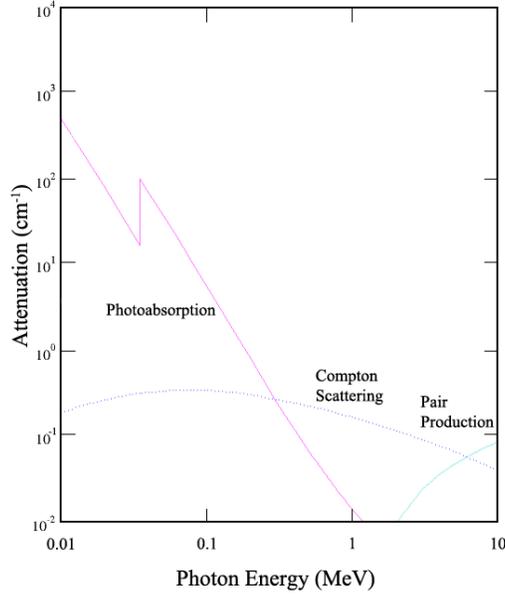


Figure 1.10: Example of xenon attenuations as a function of photon energy. The energy range of interest is 0.01 MeV to ~3 MeV. In these ranges photoabsorption and Compton scattering are two competing processes. Since the energy range of interest does not go beyond 3 MeV pair production is usually not considered, as at these energies pair production cross-sections are small.

The energies are related to the scattering angle via the Compton formula [29]:

$$E'_\gamma = \frac{E_\gamma * m_e c^2}{E_\gamma(1 - \cos \phi) + m_e c^2} \quad (1.3)$$

where E'_γ is the energy of the scattered gamma-ray, E_γ is the initial energy of the gamma-ray, and ϕ is the scattering angle. If there are only two scatters, then one obtains:

$$\phi = \arccos \left[1 - m_e c^2 \left(\frac{1}{E_\gamma - E_1} - \frac{1}{E_\gamma} \right) \right] \quad (1.4)$$

where E_1 is the energy lost in the first Compton scatter.

Thus if one is able to measure the position of each event (x , y , and z for each scatter) and the energy of each deposition then the Compton scattering angle can be calculated. This, combined with the position information, can be used in order to reconstruct the position and energy information of a radiation source. The correct sequence of the two interactions must be known in order to correctly construct the event.

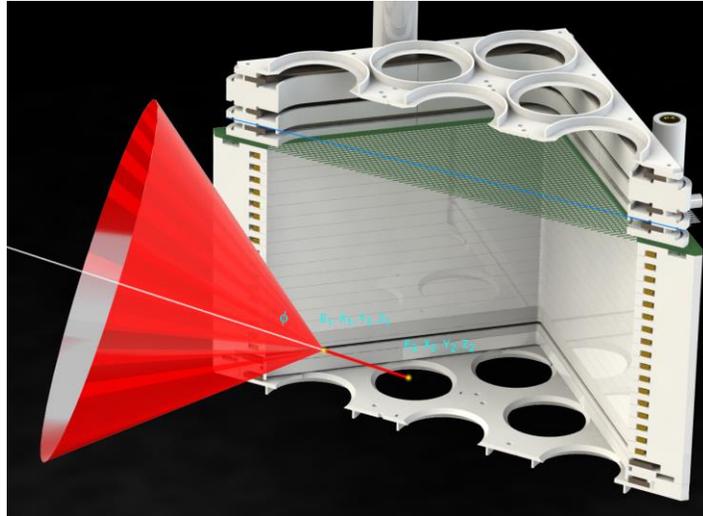


Figure 1.11: Illustration of the Compton imager concept. The red cone represents the Compton back-projected circle onto the image. The four dimensional coordinates (x , y , z , E) are shown. The charge sensitive wires that read out the xy position information are shown in the green grid.

A Compton imaging detector must be able to detect not only the energies of the scatters, but the position of the scatters as well. The energies are extracted via the signals from the PMTs, as usual. The depth of the event (z -coordinate) is inferred from the time difference between the prompt light and the proportional scintillation signal. However the horizontal information (in this case the xy plane) must be known as well. This can roughly approximated using the weighted sum of the PMT signals, however resolutions of this type are usually poor, with values of ~ 1 cm [2]. Instead, an additional readout system using a series two sets of parallel wires is used to measure the xy position by the method of image charges. Using this method a position resolution of \sim mm scale has been demonstrated, with possibly better position resolutions [30]. With the combined information from the PMTs and the wire readout it is

possible to calculate the Compton scattering angle. The positron information combined with the Compton scattering angle back-projects a cone onto the image, since the other scattering angle is unknown.

The process described above is the reconstruction of a single event. During acquisition, multiple events with different energies and scattering angles would be back-projected onto an image. If the radiation is a dispersed background source, then the signal would be many back-projected cones with little to no correlation in their position. However if there is a radioactive source present then the cones will overlap at a particular location. After many reconstructed events an image is formed of the radioactive source if there is one present.



Figure 1.12: Example of multiple back-projections onto an image. The color in this picture represents energy. Although only two discrete energies are shown, a continuum of energies is possible.

The LXeGrit detector was one of the first detectors to put the idea of Compton imaging to practice using liquid xenon [30]. The Liquid Xenon Gamma-Ray Imaging Telescope (LXeGRIT) was a balloon-borne experiment imaged gamma-ray emission from cosmic sources in the 0.15 – 10 MeV energy range. The experiment was tested in three flight campaigns from the Northern Hemisphere. Since it was built as a balloon borne instrument it had to withstand the launch and landing procedures, in effect also showing that LXe TPCs could be made rugged.

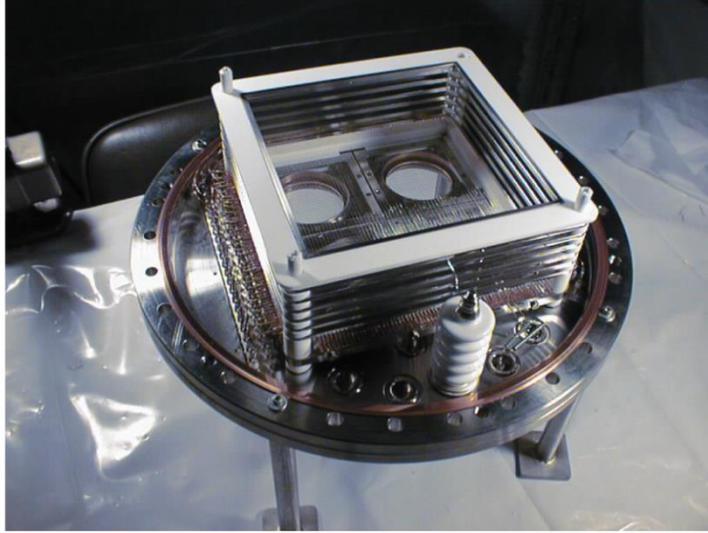


Figure 1.13: View of the inner region of LXeGRIT. The device was built to survive balloon flights.

The LXeGRIT detector had reduced energy resolution performance, which limited the performance of the Compton imaging. This is because the energy fluctuations are correlated to the fluctuations in the Compton scattering angle. When reconstructing an image using the back-projected cones there is going to be some finite width associated with those cones. The width of that circle is proportional to the energy resolution. This will be shown by doing error analysis of the Compton scattering formula.

If we assume that the energy resolution errors are uncorrelated, then we have the following formula for the uncertainty in the scattering angle:

$$\delta\phi^2 = \sqrt{\left(\frac{\partial\phi}{\partial E_\gamma}\right)^2 \delta E_\gamma^2 + \left(\frac{\partial\phi}{\partial E_1}\right)^2 \delta E_1^2} \quad (1.5)$$

with the partial derivatives being functions of the energies E_l and E_γ . It is assumed that the energy resolution is proportional to the energy in the following way [11]:

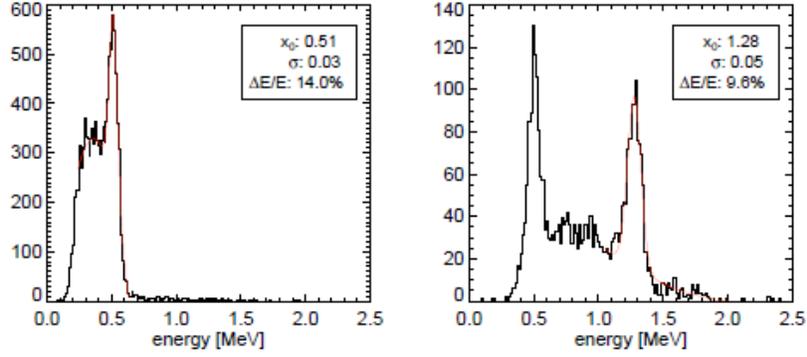


Figure 1.14: Limitations of LXeGRIT performance. Above is the charge signal performance of LXeGRIT at energies of 511 keV and 1.28 MeV. Energy resolution will limit the performance of a system due to relationship between energy and Compton scattering angle.

$$\delta E = \frac{A}{\sqrt{E}} \quad (1.6)$$

for some constant A . This constant should be a function of energy only, not the particular energy E_γ or E_I .

Thus if we plus this back into the above we get:

$$\delta\phi^2 = \sqrt{\left(\frac{\partial\phi}{\partial E_\gamma}\right)^2 \frac{A^2}{E_\gamma} + \left(\frac{\partial\phi}{\partial E_I}\right)^2 \frac{A^2}{E_I}} \quad (1.7)$$

$$\delta\phi^2 = A \sqrt{\left(\frac{\partial\phi}{\partial E_\gamma}\right)^2 \frac{1}{E_\gamma} + \left(\frac{\partial\phi}{\partial E_I}\right)^2 \frac{1}{E_I}}$$

Thus we see that ϕ 's variance is a function of some combination of the energies but it's also linearly proportional to the constant A , which is a measure of the overall energy resolution performance of the system. Thus we see that the energy resolution performance of a detector is critical to the imaging quality of a Compton imager.

1.2.4 Portal Monitoring

Beyond basic science applications there are applied areas in which LXeTPCS could be used. One such area in which liquid xenon could be used as a viable technology is for portal monitoring.

Currently, when there is a person or piece of cargo that crosses a country's border the transportation vehicle or piece of equipment usually goes through a radiation detector. This is done to observe, detect, and catalogue any radiological information, specifically if the cargo in question has any sort of threatening radioactive material (Highly Enriched Uranium, Weapons Grade Plutonium, etc.). Such systems are deployed internationally, but are mainly found in the U.S.

For these portal monitors the vast majority are either plastic or organic scintillators. These types of detectors can be made cheaply and thus are suitable for large bulk applications, however as explained in Section 1.1 their performance, specifically energy resolution, is poor. Most of these detectors are threshold or window-mode detectors meaning that they focus on very broad energy ranges. Since these types of detectors cannot discriminate between benign radioactive sources (such as K-40 found in bananas, uranium found in soil/rocks, etc.) any cargo containing radioactive material must go through a secondary scan before ultimately being cleared. This makes a loss of efficiency, since very nearly all radioactive alarms are from benign sources that must be searched and cleared.

Another way to look at this is looking at the rates of false positives, false negatives, true positives, and true negatives. In this context, a true positive would be the positive confirmation of dangerous radioactive material, a true negative would be the positive confirmation of the absence of it, a false positive would be a confirmation of dangerous material when there is none, and a false negative would be a negative confirmation of material when there really is some. Since one in a thousand false negatives are required for this application the thresholds are set in such a way as to balance towards more false positives. So while ideally, no radioactive material can be smuggled between borders this means that extra time, effort, and money must be spent on clearing all of the false positives, since the actual rate of true positives is very low.



Figure 1.15: A truck going through a portal monitor [31].

There have been many efforts to improve upon this situation, which have included using higher performance materials to replace the plastic and liquid scintillators [32]. One idea was to use high purity germanium detectors to replace the PVTs. While these types of detectors have excellent energy resolution [33], they are also expensive to make. By utilizing germanium radioactive material can be identified by template matching [34] without or reducing the need for secondary screening. Ultimately the cost of large-scale deployment of HPGe detectors across the US was deemed cost prohibitive and it was dropped from the Department of Homeland Security's Advanced Spectroscopic Portal program.

It is possible that liquid xenon could replace the current technologies, have much better energy resolution, and still be scalable like the organic or plastic scintillators. In addition by utilizing Compton imaging it would be able to reduce the rates of false positives and overall increase travel efficiency through portal systems. It could potentially do all of these things while still being a relatively inexpensive system.

1.2.5 PET Scanners

Positron emission tomography (PET) is an advanced scanning technique that takes advantage of the repeatability of nuclear isotopes in order to do imaging of functional processes in the body. The system detects pairs of 511 keV gamma-rays emitted from a β^+ source, which is introduced into the body in a biologically active molecule. Pairs of coincident gamma-rays are gathered using timing information across multiple detectors, which makes a reconstruction of the active area possible. This is used in tandem with an X-ray scan performed at the same time to image the surrounding area.

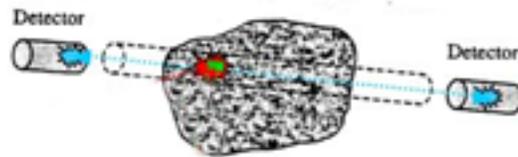


Figure 1.16: Principle of PET scanner operation. Two detectors have signals matched in time that capture the twin 511-keV signals. This creates a line of coincidence, which when added with many signals creates an image that can be reconstructed.

Currently NaI detectors are used for PET scanning operations. As shown in Section 1.1, these are crystalline detectors with an energy resolution of approximately $5\% \sigma/E$ at 662 keV, and are relatively inexpensive. The position reconstruction with each crystal is a few cm so what's typically used is a ring of detectors, and the signals are matched both with an energy response window as well as a time window. This creates a coincidence between the two signals, which creates the LOR (line of response) needed for image reconstruction.

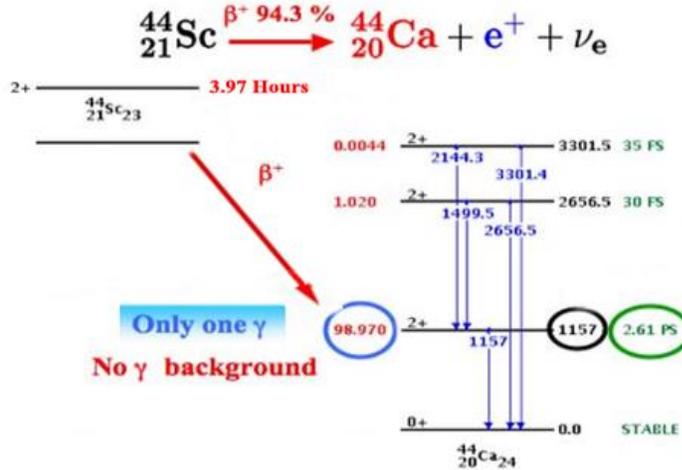


Figure 1.17: Schematic of a ${}^{44}\text{Sc}$ decay. The decay into ${}^{44}\text{Ca}$ produces both a positron and a gamma-ray of 1.2 MeV. Because of the fast time between decays this can be used in a coincidence detector.

The goal of the enhancement for an LXe TPC would not be to replace the NaI detectors, but to rather work in complement to them. As proposed in [35], an LXe TPC would make use of other isotopes that emit gamma-rays in addition to the positron signal. One such source would be ${}^{44}\text{Sc}$, although there are many others, such as ${}^{14}\text{O}$, ${}^{14}\text{C}$, ${}^{42}\text{mSc}$, ${}^{48}\text{V}$, ${}^{50}\text{Mn}$, ${}^{54}\text{Co}$, ${}^{55}\text{Co}$, ${}^{60}\text{Cu}$, ${}^{66}\text{Ga}$, ${}^{67}\text{Ge}$, ${}^{70}\text{As}$, and ${}^{73}\text{Se}$ [36]. For all of these sources a gamma-ray is emitted in addition to the β^+ emission, here with a 2.61 ps lifetime as shown in Figure 1.17. The gamma-ray can be used as a third line of coincidence using Compton back projection, which would reduce each event from a line of coincidence to a small 3D region. This in turn would increase the efficiency and reduce the reconstruction time necessary in order to make an image.

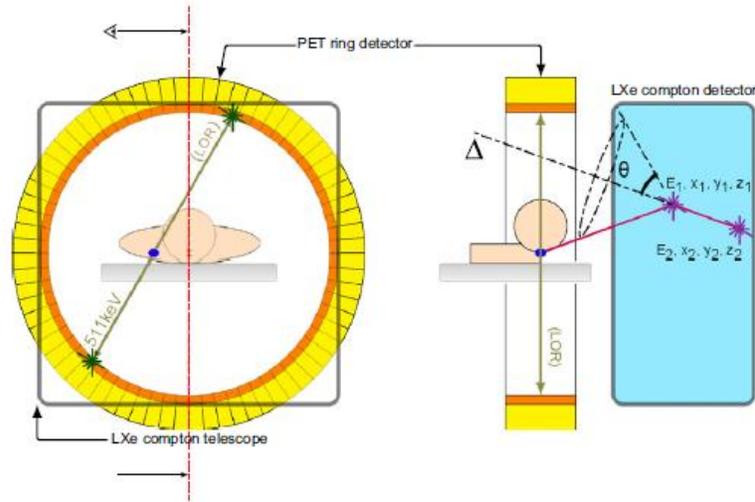


Figure 1.18: Visualization of how a PET system would work, as shown in (reference). The yellow ring represents the standard NaI detectors and the blue box represents the LXe TPC. Rather than each event being a LOR as shown on the left each event is a much smaller 3D region as shown on the right.

1.3 Thesis Outline

All of these applications require precise knowledge of detector functionality and operation, including a detailed knowledge of the underlying physics of the detector. PIXeY, which stands for Particle Identification in Xenon at Yale, was built to study these effects. It is an R&D detector built to test critical parameters in detector design that deal with energy resolution, gamma/neutron discrimination, electron extraction efficiency, and low energy thresholds using ^{37}Ar . Theoretical models predict energy resolutions that are better than what has been achieved thus far, and the mechanism setting the lower bound on energy resolution is currently unknown [37].

In order to determine why the energy resolution values do not match theoretical predictions one must first find the dominant factor that sets the energy resolution of a system for a given energy, which is not known at energies above ~ 50 keV [38]. Since there are many applications where the energy resolution at MeV scales or higher is important [39], the knowledge of the underlying physics is desired.

I propose a thesis on the study, optimization, and modeling of the energy resolution in two-phase liquid-gas xenon time projection chambers (LXe TPCs). The goal will be to determine the parameters and circumstances that optimize energy resolution and derive an empirical model that matches the energy resolution as a function of all detector parameters. This thesis will point out the critical components that went into the hardware design, purpose-built software that was designed from the ground up to produce accurate results, and finally optimization algorithms that enhanced the current measurements by tuning the software detector parameters. This will culminate into a publication reporting on the energy resolutions obtained from this detector and finally modeling of the intrinsic fluctuations within a liquid xenon time projection chamber.

2 The PIXeY Detector

2.1 Principle of Detector Operation

The events that are able to be reconstructed with this detector involve elastic scattering with xenon atoms in the active region. This can be either electron or nuclear scattering and it depends on the incoming particle [40]. When a gamma-ray or other particle interacts with the liquid xenon volume it produces a localized cluster of excited xenon atoms, ionized xenon atoms, and electrons. The excited xenon atoms form excited dimer (exciton) molecules with unexcited atoms. These molecules have a half-life of 27ns or 3ns for the triplet and singlet states respectively [12]. Upon decaying it produces light with a wavelength of 175 nm, which is called the prompt S1 signal. Since this is produced by a dimer and not the atom itself it means that the bulk of the xenon is transparent to this light.

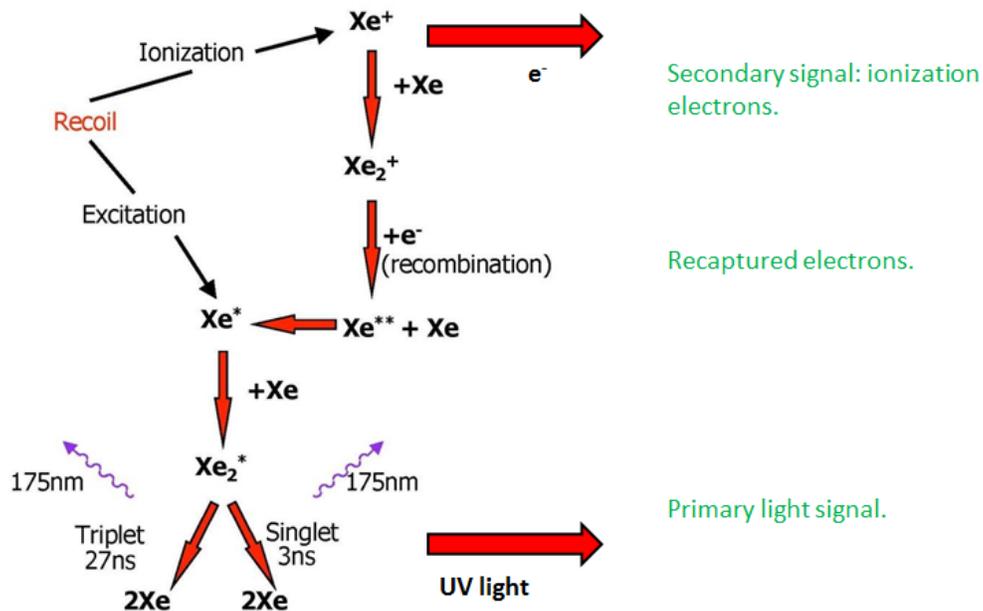


Figure 2.1: Schematic of production of the charge and light signals inside the detector. The primary light signal is produced by the decay of a short-lived Xe_2 molecule. The charge signal is produced by extracting the free electrons.

The remaining charges can be detected by adding an external electric field in the liquid region. These charges drift upward to the liquid-gas interface where they are extracted by the electric field from the anode. The interactions with the drifting electrons and xenon atoms in the gas produce a secondary light signal called proportional scintillation, or the S2 signal. The anode voltage is tuned to yield a reduced field of approximately 3.3 V/cm-torr in the gas medium, which is considerably smaller than 20 V/cm-torr where electron multiplication is known to occur [41]. Since no charge multiplication is produced in the gas the amount of light produced is proportional to the number of electrons drifting through the gas by a high degree of precision, which has less noise than using charge counting techniques [42].

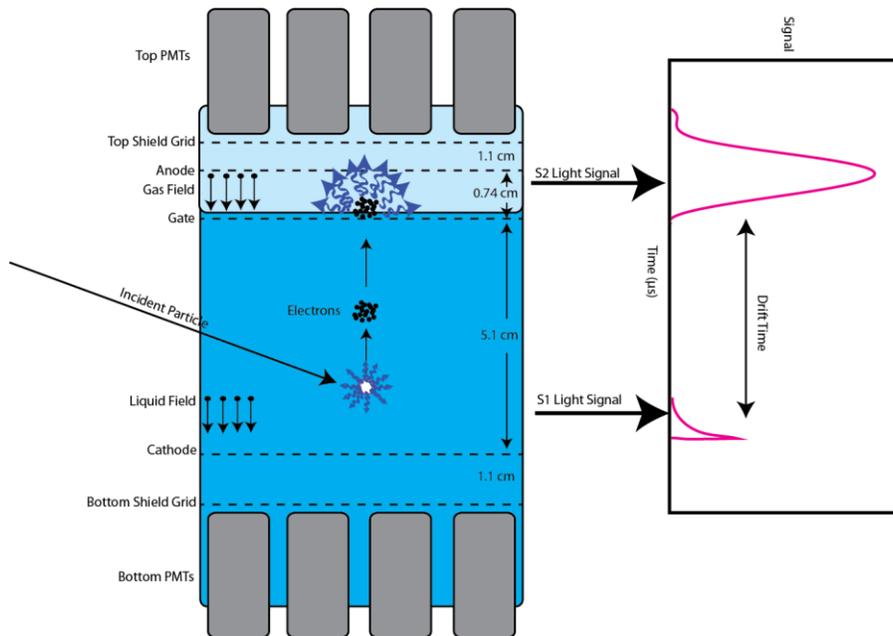


Figure 2.2: Visual schematic of the operation of the detector. The drift field is set up by the cathode grid with negative high voltage, and the gate grid which is at ground potential. The field in the gas is set by the grounded gate grid and the anode grid, which is set to a positive voltage. The S1 and S2 signals are read out by an array of photomultiplier tubes (PMTs) on the top and bottom of the detector. The depth of the event can be determined by calculating the time separation between the S1 and S2 signals.

It was recently discovered that the charge and light signal are anti-correlated [43]; a better energy resolution than either S1 or S2 alone can be obtained by performing a linear combination of both signals.

This will allow for a deeper systematic study of the detector properties of liquid xenon, since this approach ideally cancels out the recombining fluctuations in light and charge.

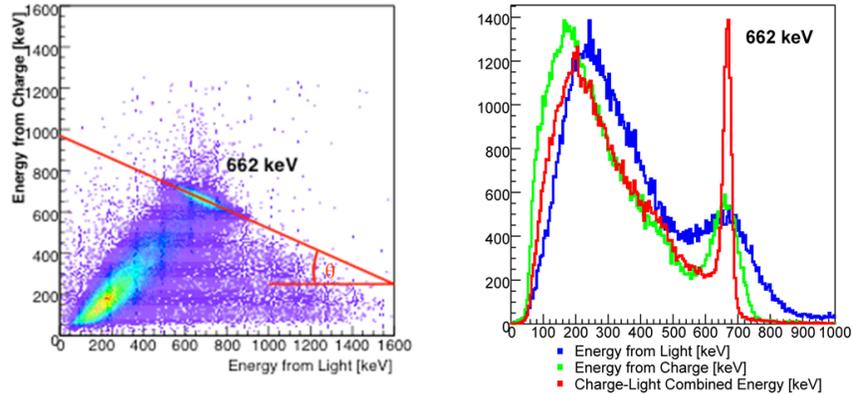


Figure 2.3: Observation of charge/light anti-correlation in the Xenon10 experiment [43]. By adding the S1 and S2 signals according to the anti-correlation angle theta a better energy resolution was obtained.

2.2 Hardware Design of PIXeY

PIXeY was designed with a number of hardware improvements made over previous detectors, many of which are now standard for liquid xenon time projection chambers. Teflon walls are used throughout the detector due its high reflectivity in liquid xenon [44]. The Teflon was also used for the surrounding plastic construction, even parts that wouldn't come into contact with the scintillation light. This was done to match the thermal contraction of all materials to prevent thermal tensioning.

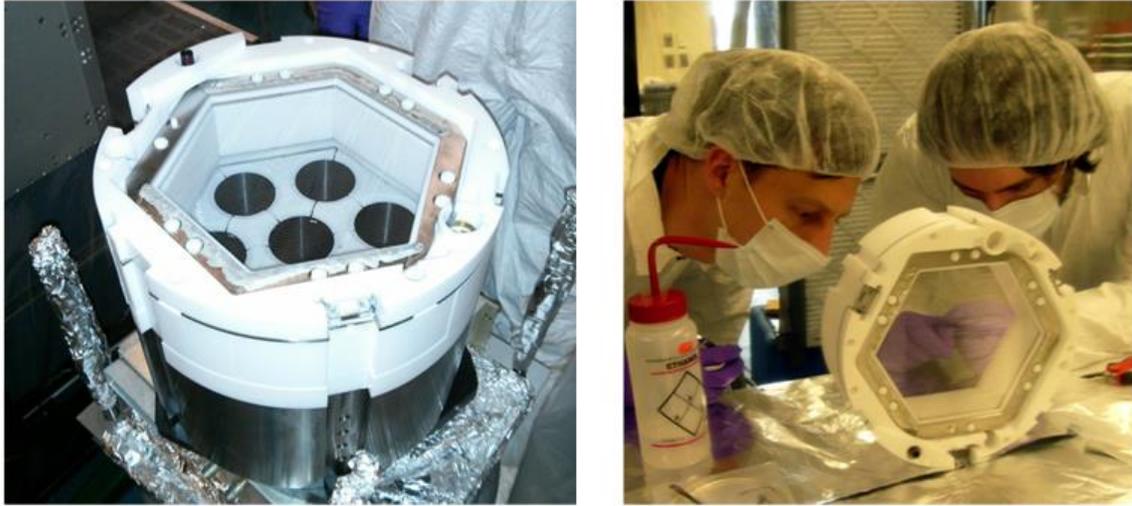


Figure 2.4: Left: inner view of the detector. Teflon walls are used throughout, even for parts that do not come in direct contact with scintillation light. Right: construction of PIXeY was done in clean room environment in order to minimize the amount of contaminates.

The PMTs that were used for PIXeY were Hammamastu R-8778 PMTs. They have a 2 cm viewing view and are approximately 12 cm long and a light detection efficiency of about 35% [45]. There were 14 PMTs used in total, with a patterned array of 7 on the top and bottom. The PMTs are directly immersed in the liquid with no wavelength shifters needed since they can directly detect the scintillation light.



Figure 2.5: “Russian doll” images of PIXeY. Left image: inner detector showing top PMT block, bottom PMT block, and field region. Middle: Stainless steel outer can which forms fits around the inner region, in order to keep the amount of liquid xenon used to a minimum. Right: Outermost can which insulates the inner can. The operational temperature of the detector is 175K, which liquefies the xenon.

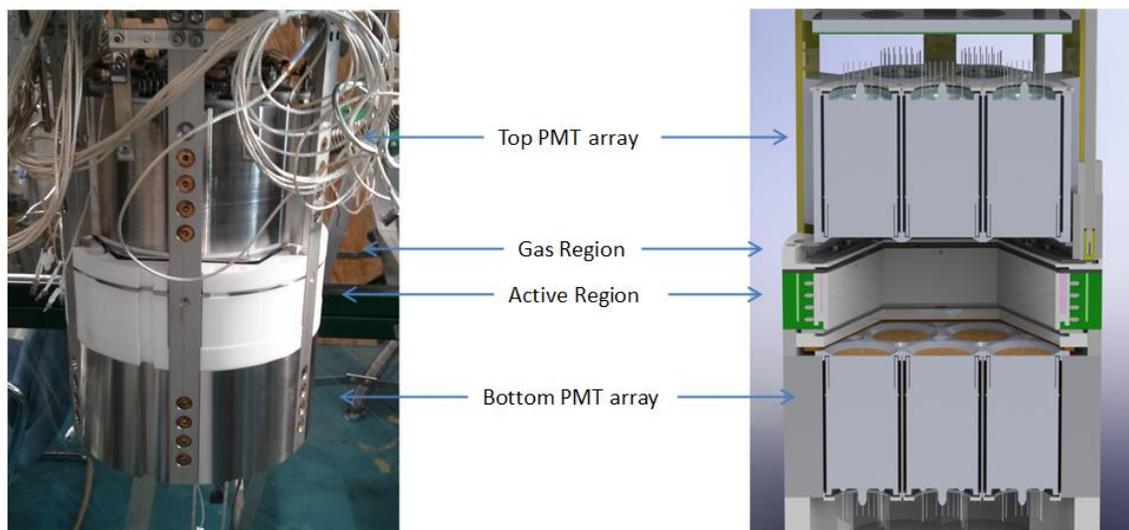


Figure 2.6: Annotated view of PIXeY, with cross-sectional view on the right and the matching external view on the left.

In the following sections I will be going over key components that went into the design and construction of PIXeY. In particular the items that will be detailed are related to the electric fields that are

established in the detector and the methodology of achieving the high voltages and uniform electric fields will be discussed.

2.3 Electrostatic Grids

2.3.1 Grid Introduction

A critical design component in a time projection chamber is the voltage connector systems and planar surfaces that establish the electric fields in the detector. For PIXeY there are two field regions: the drift field, where the electron cloud drifts through the liquid; and the proportional scintillation field, where the electrons drift through the gas. Because a wide variety of fields were tested, the feedthroughs and the grids need to be designed to tolerate a high voltage environment. High voltage tests were performed to simulate detector operation and look for electric field breakdowns. The results of those tests informed the design and construction of each grid frame.

2.3.2 Grid Simulation

The wire grids were optimized by two parameters: electric field uniformity and optical transparency. The grids need to be transparent so that the light created by the S1 and S2 signals can be maximally collected, and they also need to create a uniform electric field in order to reduce field variations in the light collection statistics. If the grids are too transparent, then the wires that establish the field will be too sparse and thus there will be significant field variations. If the wires are too close together then optical transparency will be small. In addition, tension on the wire is important in order to reduce the deflection when voltages are applied. Thus the wire tension, size, and pitch must all be tuned for optimal energy resolution.

Ideally, a detector would have a perfectly uniform electric field; the electric field would have no positional dependence within the detector and would be described by a planar field:

$$E = \frac{V}{d} \quad (2.1)$$

where V is the voltage difference between the surfaces and d is the distance between them. Many calculations assume a perfectly uniform electric field. In the gas and liquid portion of the detector, the speed at which the electron drifts is a function of the drift field in the detector [46] [47], so variations in the electric field would produce variations in drift field speed.

If there is a significant field variation in the liquid portion of the detector, then the drift time between pulses cannot be correlated well to the physical location of the event, which is needed in order to compare different events for quality selection purposes. However in practice the main consideration for the electric field in the liquid is the variation due to the edges of the detector.

Grid spacing becomes much more important when one has to consider the effects on the proportional scintillation, or gas region. Because the gap between the gate and anode grid is ~ 6 mm and the spacing between the wires is mm spacing, any variation in the microscopic electric field at the wire surface could have an impact on overall signal performance. Assuming an ideal setup, the equation for the number of photons created during the proportional scintillation phase is [48]:

$$N_{ph} = \alpha \left(\frac{E}{p} - \beta \right) xp \quad (2.2)$$

where α and β are constants that depend on the type of material, E is the electric field in the gas, x is the gas gap distance, and p is the pressure. A modified version of this equation was studied, which is:

$$N_{ph} = \alpha(V - \beta xp) \quad (2.3)$$

which is explained in Appendix A.1. Any variations in these values will result in a variation in the number of photons, which in turn will impact the energy resolution.

A simulation of the electric fields in the detector was performed using the Maxwell simulation software. It is a commercial off-the-shelf program that can do electrostatic, magnetostatic, and some

electrodynamic simulations [49]. These simulations were done to determine the optimal liquid level placement that separates the liquid and gas regions in the detector, as well as inform on optimal wire pitches. The optimization constraints are the gas-to-anode-gap d_{gas} and field uniformity at the liquid surface. If d_{gas} is too big then electric field variations at the wire surface will negatively affect the S2 signal size, if the gap is too small then the S2 size resolution will be dominated by signal error. Thus the liquid level must be carefully chosen so as to reduce electric field uniformities at the surface of the liquid but also allow a large enough gap for a large proportional scintillation signal.

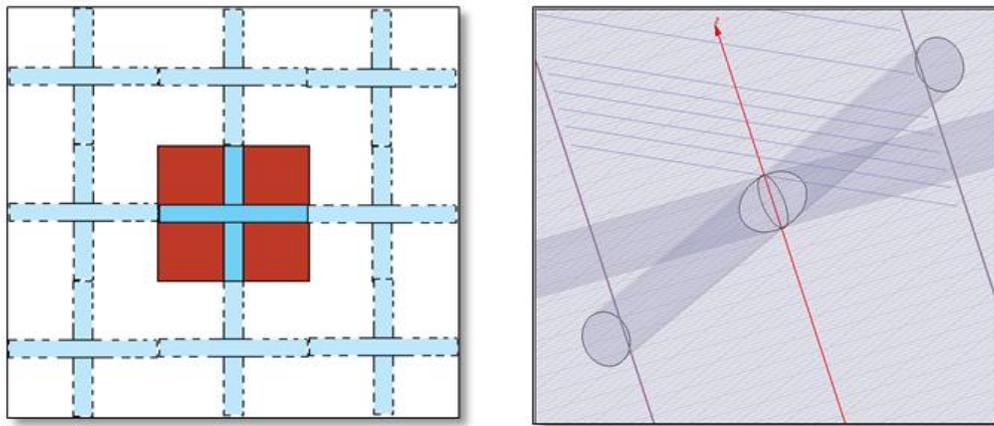


Figure 2.7: Example of modeling a wire mesh grid using the Maxwell simulation software. Left: The red box shows the actual simulation region, with the dotted lines being the images that are created due to the boundary conditions on the box. Right: The diagonal lines across the cross are the voltage probe regions. How the voltage changes across those probe lines will be studied.

Attempting to do a full simulation of the detector with the details of the wire grid surfaces was too computationally intensive, due to the immense amount of meshing surfaces that would need to be created by the simulation software in order to study the micro-effects of the grid wires themselves. Instead the problem was thought of in two regimes: the macroscopic effects of the detector geometry and the microscopic effects of discrete wire spacing. A much simpler simulation was conceived which was that of an infinite wire plane of discrete wires. In order to model this, a single cross is put into a box with Neumann boundary conditions. This can be used to study the micro-effects on the electric field near the

wire surface while being much less computationally intensive since only a single duplicated region is simulated.

The variation for each equipotential line was calculated as the maximum likelihood standard deviation of that line. The variation in the potential lines was modeled as a log-linear plot as a function of distance away from the grid with the following form:

$$\text{Log}(Variation) = A * d_{wires} + B \quad (2.4)$$

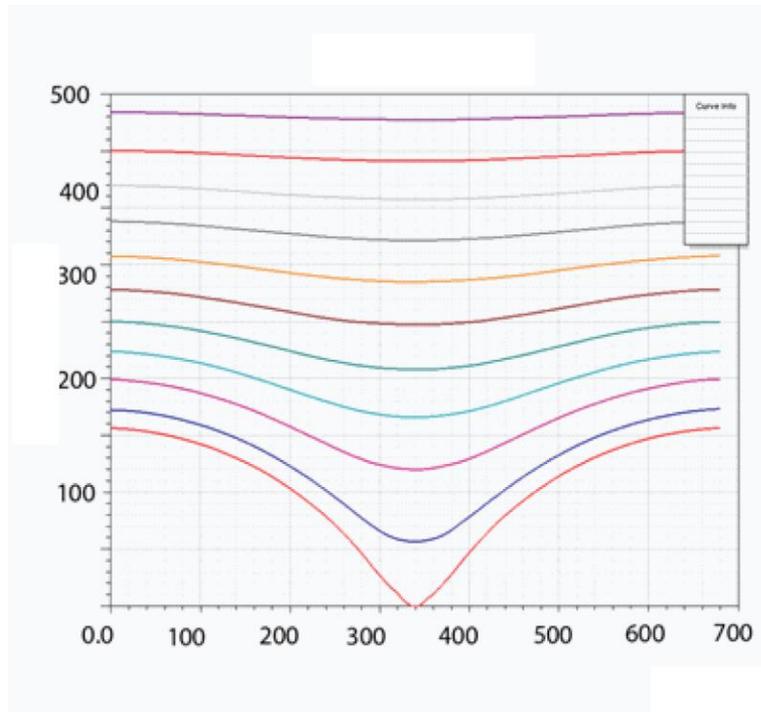


Figure 2.8: Equipotential lines that are the results shown in Figure 2.7. Each line is constant distance away of 50 μm . As the distance is further increased away the wires the equipotential lines more resemble a flat field, which would be an ideal sheet field.

with the values of A and B calculated from the extrapolation in Figure 2.9.

$$\text{Log}(Variation) = (1.92 \pm 0.02)d_{wires} + (-0.00771 \pm 0.00011) \quad (2.5)$$

At the wire surface the field variation is nearly a factor of two of the base value. However at just 100 μm away from the wire surface the variation is only 10%. The variation is 1% after 250 μm away from the wire, which is half the mesh spacing in this simulation. Thus the problem of variations near the wire surface can be circumvented if the liquid level is simply set to be slightly above the wire surface. The liquid level was tuned so that the variations in the electric field are largely shielded in the gas region¹. For our studies the lower bound on the liquid level was chosen to be at least 1mm above the gate grid surface. Eventually this value was raised up to be 3mm based on S2 pulse sizing.

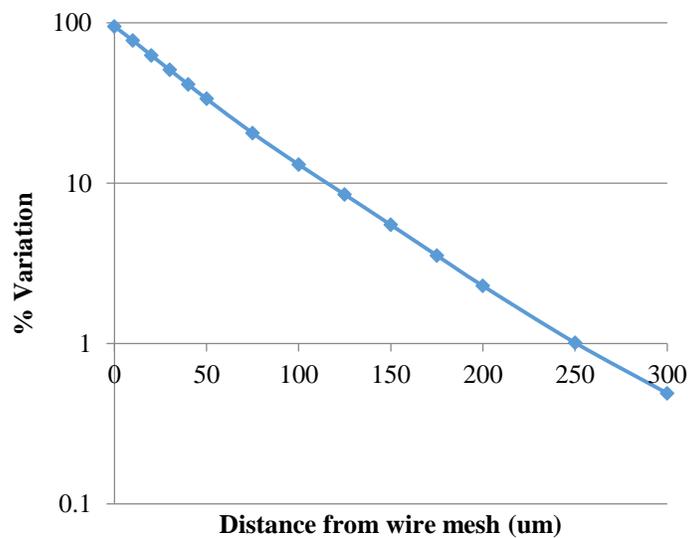


Figure 2.9: Percent variation as a function of the distance from the wires. The relationship follows a log linear relationship and the variation is 1% when 250 μm away from the wire surface.

2.3.3 Mesh Grid Construction

Different types of grid fabrication techniques were researched, mainly dealing with wire mesh or grids that were constructed from spools of wire. One method that was researched was for wire mesh grid fabrication. For a wire mesh grid the mesh must be mounted to a frame in such a way as to retain uniform

¹Since there is a grid on the top and not a charge plate, there were unavoidable fluctuations in charge production from field uniformities at the top grid.

tension throughout the mesh. A stretching rack was used to stretch the mesh in two planar directions. Weights were held on the bottom of the ropes to act as the tension for the rack.

In order to fabricate a mesh grid the following procedure was followed. The wire mesh was first put on a pre-stretching device to hold the mesh in place. Four metal bars were soldered to the mesh, which would later act as the holder for the stretching system. Once that soldering operation was completed, the mesh was cut so that only that piece of wire mesh remained. The metal bars were then attached to the stretcher so that the wire mesh would undergo its final tension. Once the desired amount of tension was in place, a metal grid frame on a hot plate is raised to the tensioned wire mesh. The mesh is soldered to the frame while being heated by the hot plate. Once the frame is cooled the mesh is cut and excess mesh is removed.

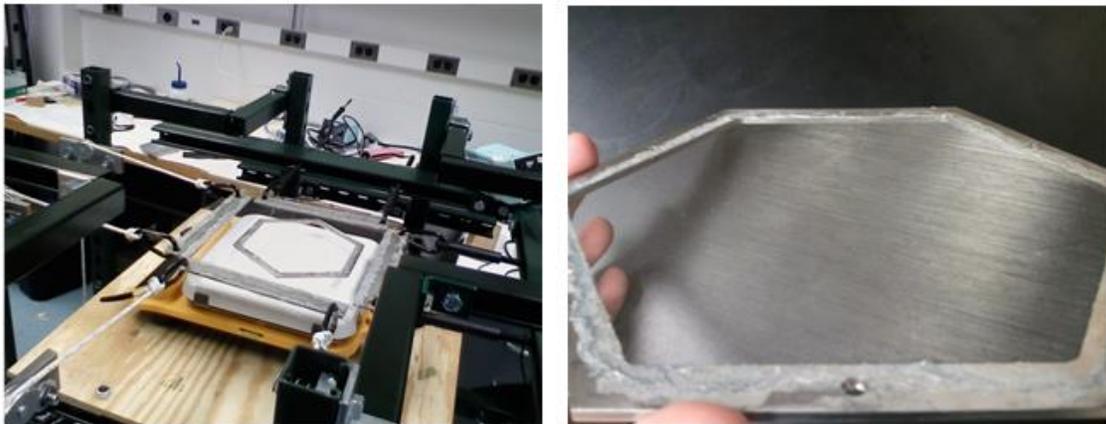


Figure 2.10: Early attempts at creating electric field grids (left) and one completed mesh grid (right). The hot plate, mounted mesh, and stretcher are visible.

The heating operation was the most difficult task. Stainless steel was used for the grid frame and the mesh due to its corrosion resistance, but it is also a difficult material to solder [50]. In addition Tin-Silver solder was used instead of a Lead-Tin solder due to minimization of Lead contaminants in the liquid xenon [51]. In order to melt the solder, a hot plate was used to heat the entire frame before soldering. When the mesh was soldered to the frame and then cooled “bubbles” in the mesh would form

which were locations where the mesh had completely lost its tension. This was most likely due to the fact that the frame and mesh was not cooled uniformly and the pretension process was not uniform. Similar methods have been successful but with epoxy [52], and was successful in mounting the mesh to the frame. Heating appears to be the key factor in failures.

2.3.4 Single Wire Grid Construction

The second technique studied and the eventual technique that was implemented in PIXeY was using a single wire wrapped around a loom and soldered to a frame. The single wire setup has advantages over mesh: a bubble effect is less likely to occur since the wires only need to be soldered in two locations. Also the electric field generated by the wires is conceptually simpler to understand. To make use of these advantages, a reliable method of holding the wires at uniform tension and also keeping that tension during the soldering process was studied.

As has been done elsewhere [53], a device was built that wrapped a wire under a fixed amount of tension on a loom with the grid frame suspended on the inside of the loom. After the wire is wrapped around the appropriate number of times the frame is adjusted so that it is just touching one surface of the wires, as preparation for the soldering operation. In order to melt the solder a blowtorch soldering gun was used to apply large amounts of local contact heat to melt the solder and attach it to the frame, which was used instead of a hot plate in order to reduce frame warping.

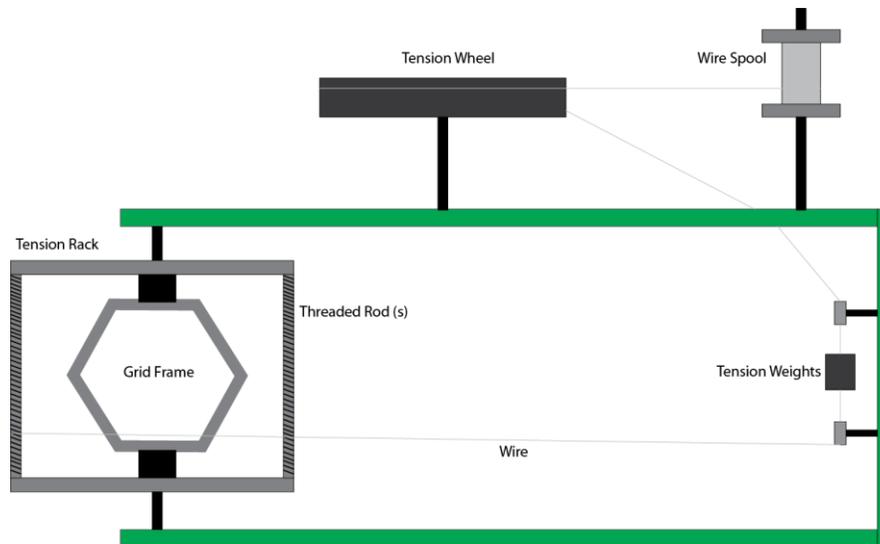


Figure 2.11: Illustration of the wire grid loom in an overhead view. Two operators are needed in order to use the loom, one to feed the wire from the spool, and the other to wind the wire around the loom. The process takes anywhere from 2-4 hours.

The wire wrapping process requires two operators, although it is entirely possible for this process to be automated. One operator draws the wire from the spool while simultaneously moving the tension wheel, which separates the low tension side of the spool with the high tension side of the loom. Without the tension wheel if the wire is drawn while under tension then several “kinks” (places where the wire is permanently bent) form on the wire which reduces strength performance. After the wire passes the tension wheel it is then passed through a sequence of pulleys to a dip where weights are held by a small rotatable pulley. As long as the wires are roughly parallel then the tension in the wire should be exactly half of that on the weights.

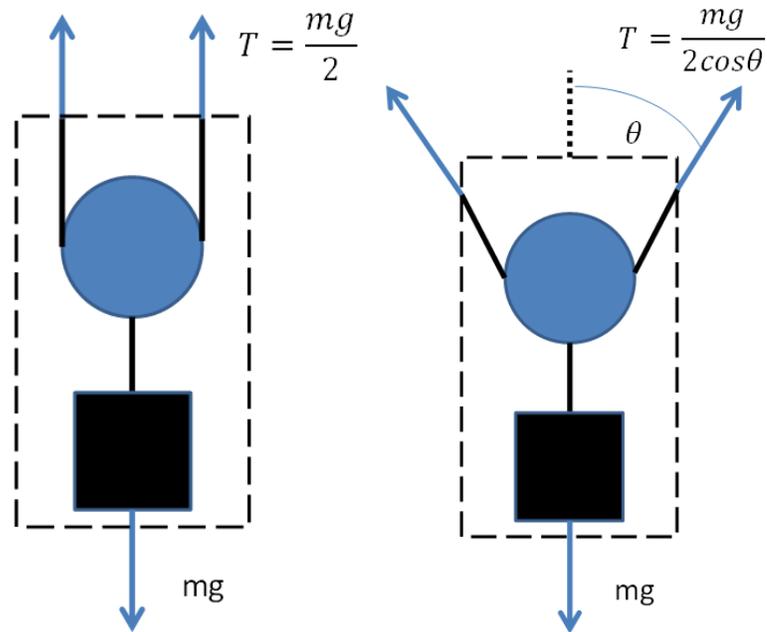


Figure 2.12: Left: Free-body diagram of the pulley system with near vertical wires. As long as the wires are vertical then the tension in the wire should be exactly half of the weight system. Right: Free-body diagram of the pulley system with splayed wires. The tension in the wire actually increases due to the wires being pushed apart.

Thus, the wire spool operator’s job is to keep the weight system low, because the lower it is the more closely the approximation in Figure 2.12 is true. If the weights get too high the wires will form a v-shape and then tension will increase due to the extra “cost” of stretching the wire horizontally.

The other operator’s job is to rotate the wire loom itself. The wires are guided by threaded rods that are on either side of the loom which have a 1-mm pitch. This sets the minimum distance between each wire, although larger pitches are possible. As the mass is being lowered by the wire spool operator, the loom operator rotates the loom and places the wire on the threaded guide rod. This raises the mass so the two operations must be done in sync to keep the tension weights at an optimal location.

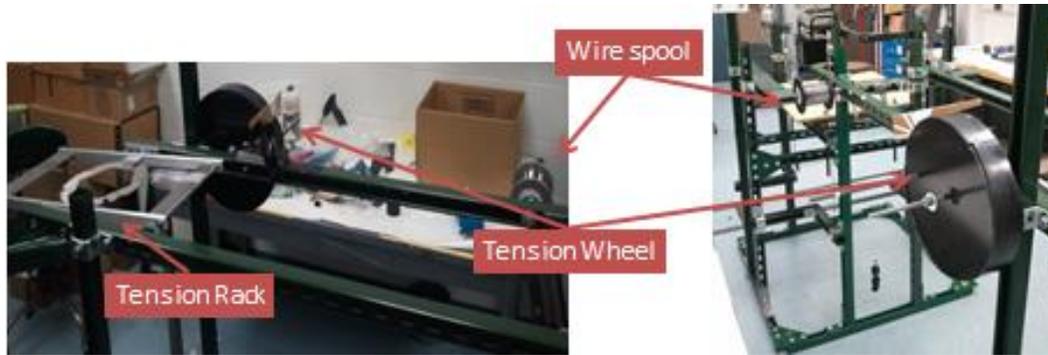


Figure 2.13: Annotated view of the wire stretcher. On the right the weights are visible and are between the two green strut bars.

Tin-Silver solder was chosen for mounting the wires to the frames. This was for strength purposes and also to minimize adding electronegative impurities into the liquid. Most of the grids would be in the liquid xenon, which means that any electronegative impurities that an epoxy would have would be filtered into the liquid, since xenon's polarizability can detach impurities [54]. For soldering there is a risk of contaminating the xenon with gamma-ray impurities, since there are trace radioactive impurities in solder [51]. For this reason Tin-Silver solder was chosen over the traditional lead-based solders. While any solder means that there will be higher background levels in the detector, the actual signal performance will be unaffected, provided that the event rate is higher than the background rate.

Monel was chosen as the material for both the frames for the wires to be mounted on and the wires themselves. This material was chosen for its corrosion resistance, steel-like hardness, and high solderability [55]. There are 5 frames in total that were fabricated: a cathode grid, an anode grid, a gate grid, and two shield grids. The wire itself was AWG 40 spring-hardened Monel, which is about $80\ \mu\text{m}$ in diameter. Each wire had about 250g of tension from the weight system, which was picked based on the yield strength of Monel [55] as well as the operating limitations of the wire loom system. Since each frame had about 150 wires this means that the axial tension on the frame was about 367.5 kN, or about 83 lbs.

Based on the simulations of the electric field, the pitch was chosen to be 1mm for all of the “primary” grids (cathode, anode, and gate), which is an optical transparency of 92%. By extrapolating the simulations done in Section 2.3.1, the field variation should be about 1% at a distance of 0.5-1mm away from the grid surface. For the shield grids that protect the upper and lower PMTs a wire pitch was chosen to be 3mm, which is an optical transparency of 96%. The field uniformity in that area is not as important however optical transparency still is, so the wire pitch can be more as long as there are no field leakages [56].

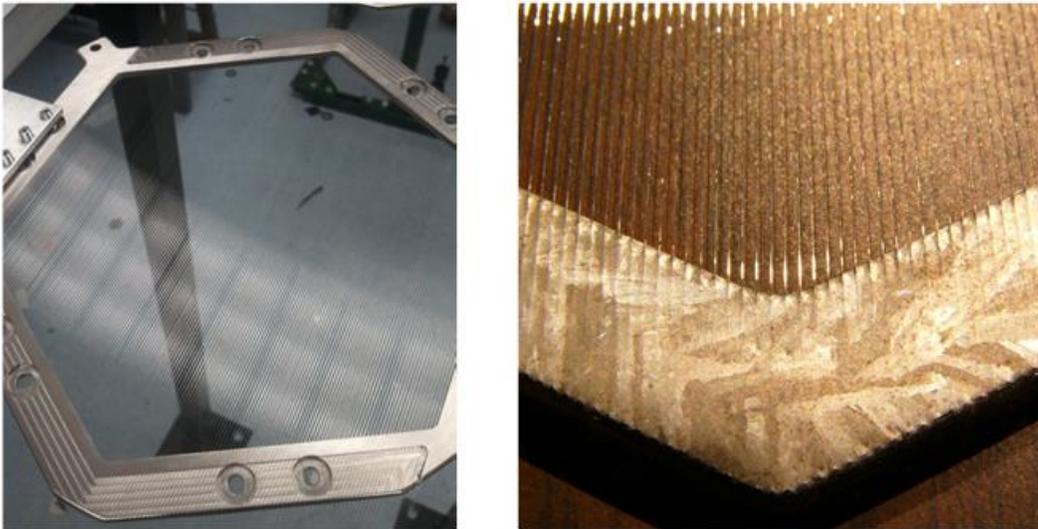


Figure 2.14: Left: an example of a wound frame. The wires that are visible are due to the low angle scattering of the above florescent lights in the room, otherwise the completed wires themselves are difficult to see. Right: close-up of the soldering operation after the wires have been wound.

In practice the loom system was able to fabricate two grids at a time. Identical grids were fabricated and the one with the best soldering operation and wire uniformity was picked. Since the failure rate was 30-50% multiple grids were made.

The main mode of failure was typically at the soldering step. The high heat combined with high tensions meant that failures during that step were likely. In particular, since the frames are hexagonal there was a gap of around 4cm where the hexagonal points met, which would cause frequent wire

breakages. For these grids if one wire breaks then the whole grid has to be made again. Besides the soldering step, there was one grid failure that occurred during the multi-phase cleaning operation. There were no grid failures during detector construction.

The detector had 2 major runs, one of which lasted for over year. Upon disassembly of the detector none of the grids had failed during operation, nor had any wire lost any of the tension in the grids due to solder creep or thermal cycling.

2.4 High Voltage Design Components

2.4.1 Introduction

One of the design goals of PIXeY was to be able to test and observe various phenomena at a variety of drift and proportional scintillation fields. This includes being able to test high value and low value fields. The design goal of the anode was +10 kV and the design goal of the cathode was -30 kV. The electric field components must be able to withstand high voltages to achieve those values. Design attention was put into the various connector components so that there would be no sparking or glowing at various connection surfaces at the high electric fields of interest.

2.4.2 The Anode Hardware

The feedthrough for the anode voltage is introduced through the gas and terminates in the gas. These terminations were considered, since the electric field breakdown in the gas is in general much less than it is in the liquid. Also the anode grid itself must be properly isolated from the rest of the system. There are a number of places that the anode grid can spark to, many of which would cause serious damage within the detector, including the lower gate grid, the upper shield grid, and the capacitive level meters that determine liquid height and liquid tilt.

A spacer had to be designed in order to properly insulate the anode from the rest of the detector, while still establishing the electric field in the proportional scintillation region. This was constructed from a solid piece of Teflon that was hallowed out into a well shape. The gate grid was below the anode insulator and the anode rested inside the well. This could not be constructed from multiple components because it was found that electrical breakdown can easily travel between the cracks in these parts, thus limiting the electric field that could be established in the system.

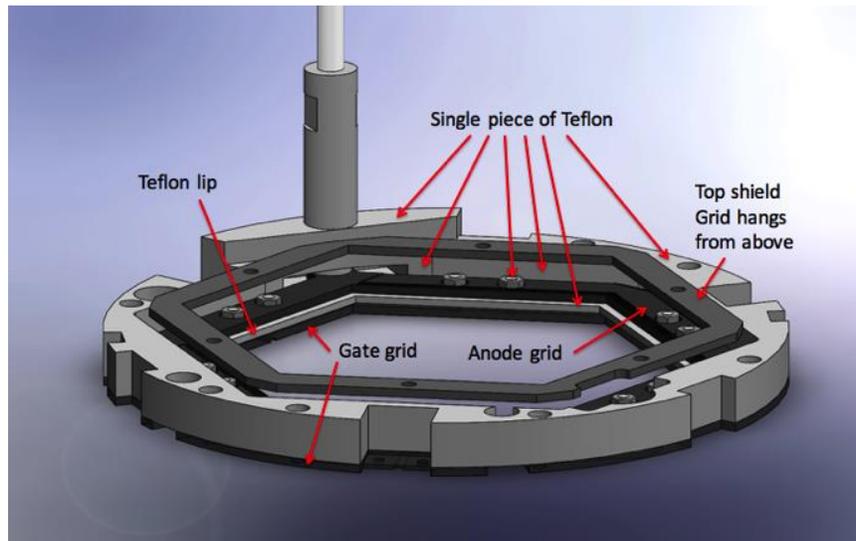


Figure 2.15: Annotated Solidworks figure of the anode spacer. Note the “floating” shield grid above it since it is not attached to the anode spacer in any way to avoid electric breakdown traveling across the surfaces.

This part was a significant technical challenge in order to fabricate. Intermediate parts had to be designed in order to fabricate certain features on the anode spacer, such as a swing chopper to cut the inset for the anode connector and custom dyes in order to thread the threaded posts that the anode attaches to. Fabrication took approximately 1 week or 40 man-hours.



Figure 2.16: Close-up view of the recess created where the anode voltage connects to the frame. The connector is threaded Teflon with a brass “bullet” that connects to the anode wire. Some of the connector components are visible in the left picture. In the left and the right the threaded posts that attach the anode to the spacer are shown. These were held in place by delicate Teflon nuts, all of which were very susceptible to stripping.

During installation the anode was attached to the spacer as a single piece. Once the anode was properly in place the frame had to be connected to its power supply. This was done with a “bullet” capture system, which has been used in other detectors [57]. Notches are cut into a hollow brass shape and then splayed apart. This makes an over fit with a matching hole that it is forced into. Since the metal has some elasticity the bullet creates an electrical contact with the hole. This receiver hole is then attached to the frame itself via a small threaded rod. All sharp edges in this connection are internal to the connector which minimizes the concentration of the electric fields.

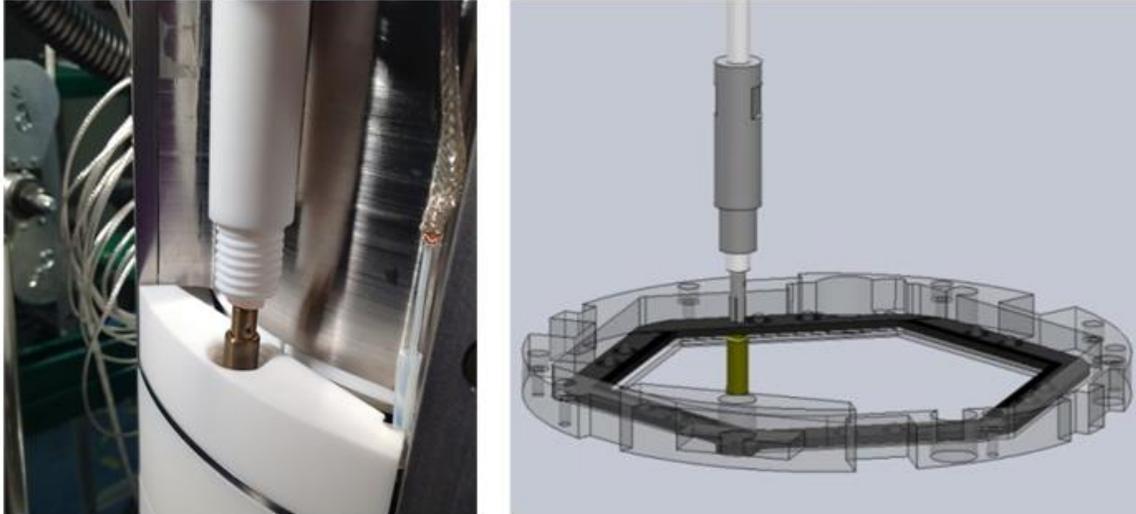


Figure 2.17: Left: picture of the HV “bullet” with Teflon sheath. Right: Schematic drawing of the principle of operation. The “fingers” of the bullet are visible which create the electrical contact between parts.

The first round of testing was done using only the frames in a test chamber. The main concern was the sparking of the anode grid to either the gate grid below it or the shield grid above it. The field cage parts were assembled and placed into a box that was overfilled with CO₂ gas. There was a small plastic window installed on the box so that a potential sparking location could be seen. The field was increased until electrical breakdown was observed. For the first rounds of testing the electrical breakdown occurred at around 10 kV. Since these tests were in CO₂ the values would have to be converted to xenon gas values using Paschen curves [58]. Frequently the sparking occurred on sharp features on the anode grid, which was expected. Eventually after several modifications the grid frames were able to withstand 18 kV. This corresponds to about 9 kV in cold xenon gas.

The electric fields were tested at the end of an operating run (so as not to damage any internal hardware during testing) in cold xenon gas. The highest steady value achieved was 10.8 kV on the anode, which is slightly above the design goal of 10 kV. During normal operation, the anode was set to 5.3 kV.

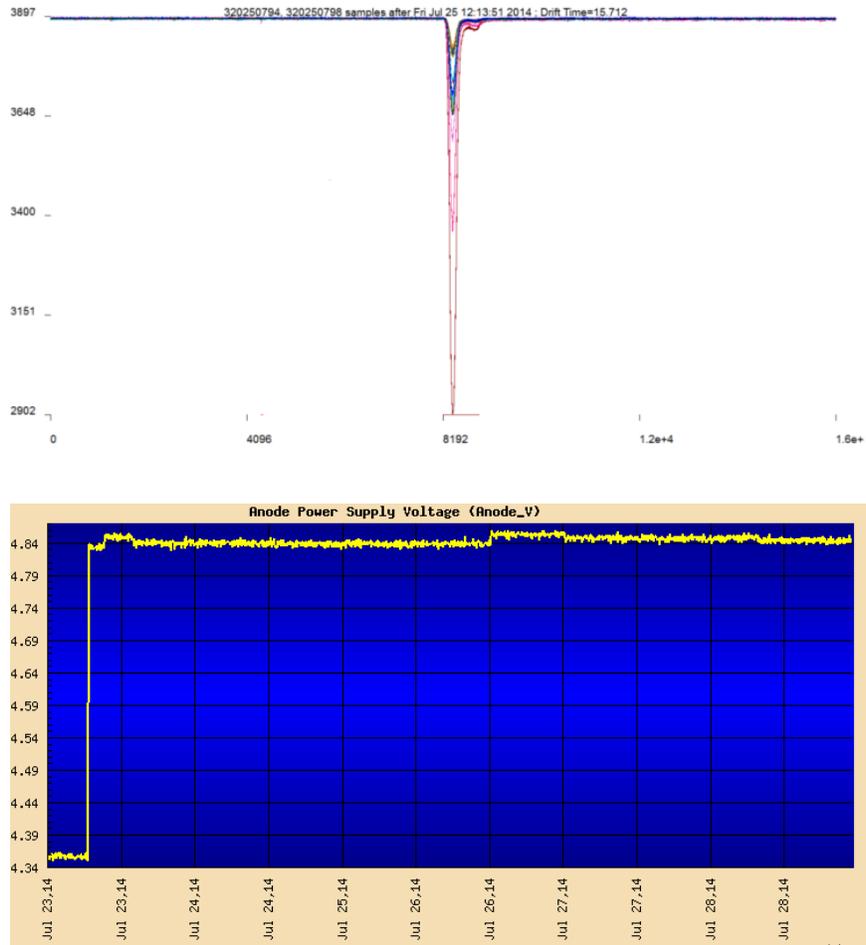


Figure 2.18: Above: a waveform during the operation of the anode at 10.8 kV, showing normal operation with no breakdowns observed. Below: the readout of the anode voltage box. Due to a readout error the readout is roughly $\frac{1}{2}$ of the actual voltage. There are some small changes in baseline voltage but overall the voltage was stable for several days.

2.4.3 The Cathode Hardware

The Cathode Hardware, while being important in establishing the fields in the drift region, was of less importance in testing. The connector interface for the cathode was in the liquid, thus sparking is much less of an issue since the liquid is much more resistive to electric field breakdown than gas.

For the cathode a “boot” connector was designed to interface with the cathode grid. This “boot” was attached to the frame by tapped holes that were in the cathode grid. The cathode wire attached to the “boot” by the bullet method described in the anode section.

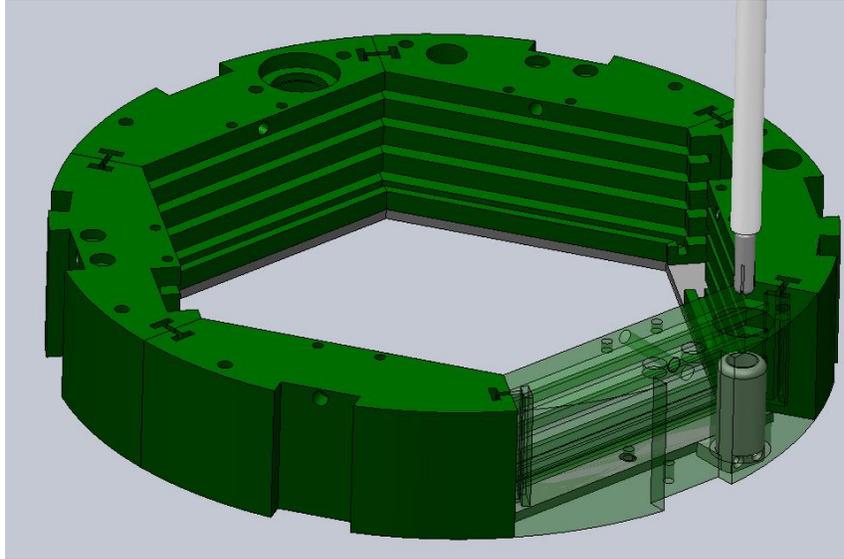


Figure 2.19: Schematic of the “boot” design of the cathode connector. Also shown is the bullet design that was also utilized for the anode connector.

The cathode was tested during actual operation at the end of an operational cycle. It was found that the Cathode was able to go to -14 kV before electrical breakdown. This is considerably lower than the design goal of -30 kV, however it was determined that such a high drift field was probably not worth testing due to the diminishing returns of charge-light recombination at increasing drift fields. For operational purposes the highest operational value of the cathode voltage was -10 kV.

2.5 Experimental Setup

2.5.1 Signal Chain

Signals induced on the PMTs are decoupled from the power at the base, with separate SHV power lines providing power to the PMTs. The signal lines are first fed to a voltage amplifier which has a gain value of 8, chosen to be low enough to preserve shaping. The signal then goes to a 12-bit CAEN ADC system which reads out the pulses, with one card reading the information from the top array and the other card reading the information from the bottom array. All 14 PMTs are then optically connected to a Linux machine via the VME crate interface. Custom data acquisition software (PIXeYCalc) acquires and stores the waveform information, which will be explained in further detail in Section 3.

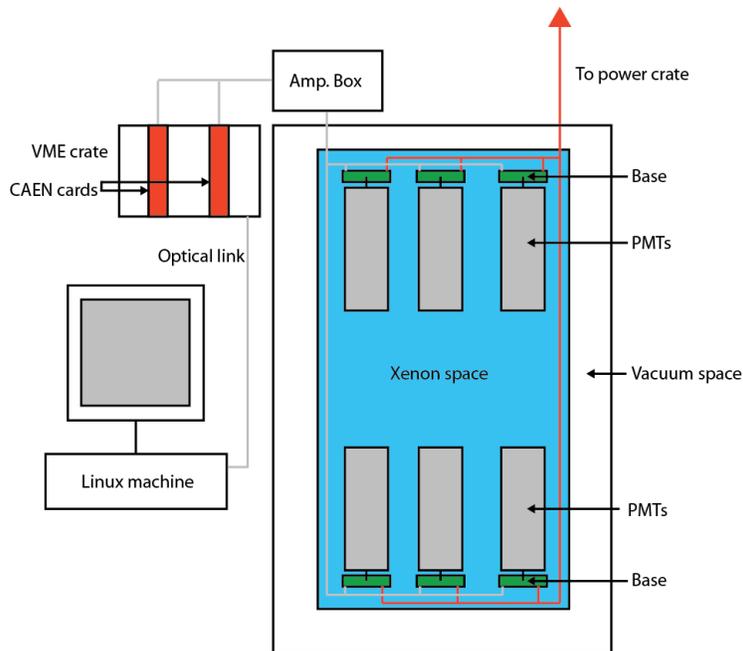


Figure 2.20: Annotated schematic of PMT signal and power chain in PIXeY.

2.5.2 Data Acquisition and Triggering

The primary trigger condition was a twofold width and height condition, in order to trigger on pulses that are S2-like. The trigger time was symmetric with the event window, so that S1s with the

longest drift time can be captured and also so that multiple scatter events could be captured. The full waveform was a readout of all 14 PMTs (whether they saw a signal or not) for 60 μ s each.

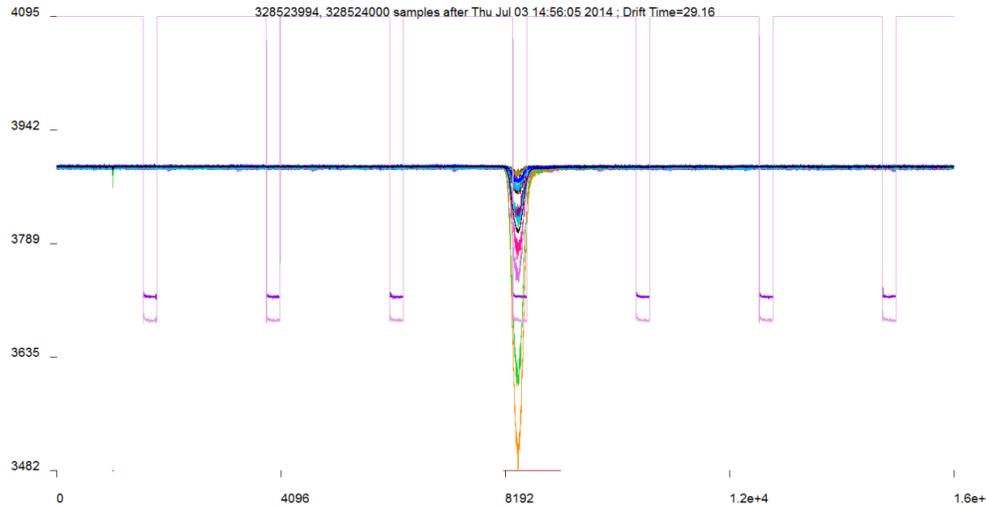


Figure 2.21: Full waveform view with pulse generator pulse visible. The pulse generator doubles as the moderator for event rates as well as aligns the separate card information in time.

For most runs we are rate-limited by data transfer from DAQ to disk. This was most likely due to the single optical connection that was used to connect the two cards to the Linux computer. As a result, we use an additional coincidence with a pulse generator to maintain a controlled acquisition rate of about 100Hz. As shown in Figure 2.21, this pulse is also read out by the DAQ system in order to match the two cards in time as well.

2.5.3 Procedure

A range of different external sealed radioactive sources were used to irradiate the detector with gamma-rays of different energies. Additionally, an internal $^{83\text{m}}\text{Kr}$ source was used to calibrate the full volume of the detector with good precision. Each source was run for 20-40 hours, in order to acquire enough data for the quality cut process. The data was saved to 1 GB binary files which were then later compressed to 250 MB. A typical source run had between 2000-4000 files. These files were then taken to

a workstation where the bulk of the raw information was processed into reduced quantities with in-house built software which is described in further detail in Section 3. This would output a reduced data file which could then later be analyzed or optimized further.

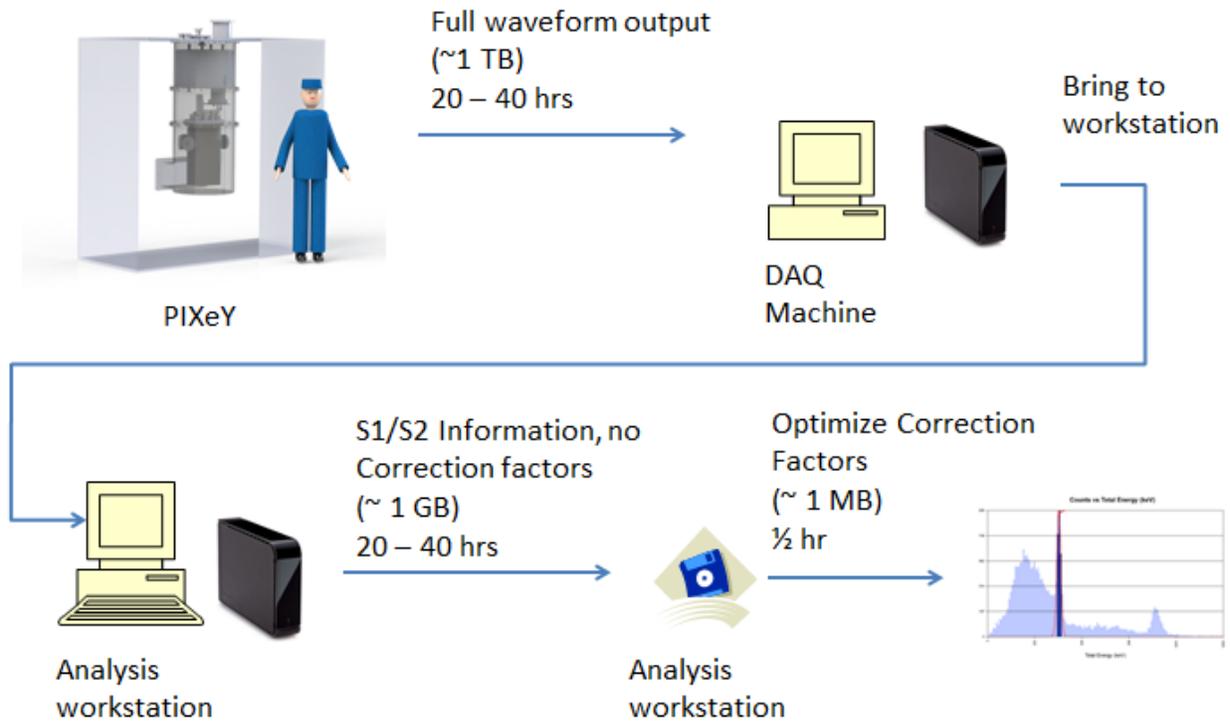


Figure 2.22: Process flow diagram that shows the steps taken from raw waveform information to optimized spectra.

3 Analysis and Acquisition Code – PIXeYCalc

3.1 Introduction

Quality software must be written in order to ensure the successful operation of the detector. This includes: acquiring and storing the waveform information coming out of the DAQ; analyzing waveform information; displaying and presenting results; optimizing data to find results; and finally building empirical and theoretical models that not only fit the data but also further the understanding of detector physics and basic science. In short useful software must be written at every step of the way for any successful project.

The analysis and acquisition code PIXeYCalc (Particle Identification in Xenon at Yale Calculator) was a project to build code in the traditional sense as a software product. It was designed to be used on multiple platforms (Windows, Linux); it could do some analysis in-situ to the data acquisition; and was written in C++ with an object oriented paradigm. The code was designed to be scalable and easily reusable so that future projects could build on it and add the features that they need. It has 90 unique classes, over 100,000 lines of code, and is based on the University of Michigan UMI project [59].

Building on this, the key component to the PIXeYCalc system was the data analysis code, which was designed both to do in-situ partial analysis as well as deep analytical studies. The data analysis code was designed to be modular so that separate components can be worked on and changed if need be, as well as to better facilitate iterative calculations. The code is split up between two main blocks: The first pass and second pass block.

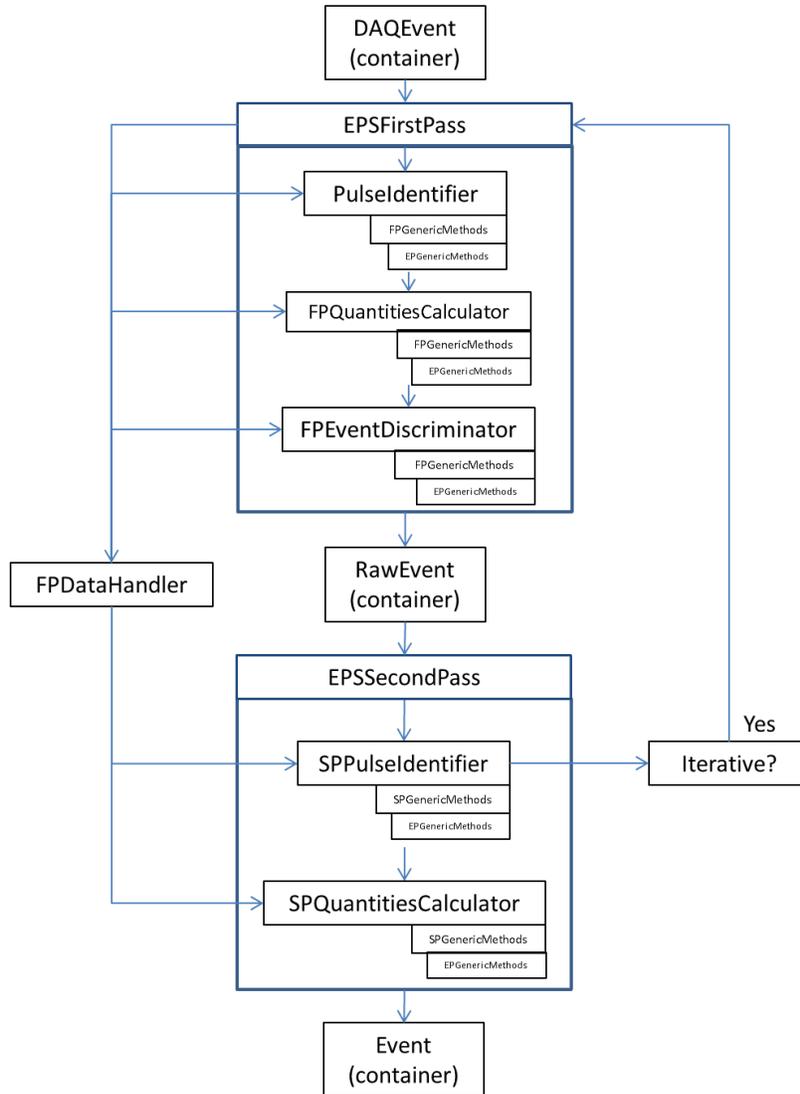


Figure 3.1: Analysis process flow chart of PIXeYCalc. The class FPDDataHandler stores all intermediate calculations and memory management the analysis classes. The iterative block is for analysis that needs to be redone based on information in later classes.

The first pass takes in information that is read out directly from the DAQ. This storage container is called DAQEvent and it contains the raw waveforms and other timing and DAQ information. The first stage of the processing is to identify all of the pulses, which is done in PulseRegionIdentifier. Then the reduced quantities such as pulse areas, heights, and widths need to be calculated, which is done in

FPQuantitiesCalculator. Next if there are any events that should be filtered FPEventDiscriminator checks if there were bad calculations or any errors that happened in the previous classes, and filters those events.

After first pass processing is completed the information is stored into RawEvent, which is an inherited DAQEvent container class (so no information is lost). RawEvent then passes its information to the second pass classes for processing. Like the first pass processing above, the information goes through a pulse region identifier. Although this time instead of identifying any pulse SPPulseRegionIdentifier identifies S1 and S2 pulses, which is explained in more detail in Section 3.4. Once this information is obtained the final scientific observables like S1 energies, S2 energies, drift times, positions, and quality metrics are calculated in SPQuantitiesCalculator. This final information is stored in the Event container class (inherited from RawEvent), which is the class that interfaces with the rest of the software program, such as building spectra or viewing waveforms.

3.2 Baseline Determination

The methodology of PIXeYCalc was to test multiple parallel algorithms and pick the one that performed the best. As such, there were two methods of determining the baseline that were investigated, which were called “first n points” and “non-pulse samples”. Both of these methods have their advantages and disadvantages, which were studied and evaluated.

3.2.1 First n-Points

This is a relatively simple method which takes the first n-points of a waveform, assuming it contains no pulse information, and does an average of those points to determine a baseline. The standard deviation of the baseline was also calculated to ensure that there was no pulse information in the baseline,

which was usually set to about 3-5 samples. Since pulses originating in the first 100-200 points was exceedingly rare this exception was rarely thrown.

This method has a number of advantages. First off is that the baseline is calculated independently of any other information. The baseline algorithm doesn't need to know a threshold or any other number that it has to be fed to; it simply can calculate the baseline based on the raw waveform information. This proved to be useful because the PMT base design caused the baseline to fluctuate between each event due to the large capacitor on the last dynode stage. Furthermore, this meant that any sort of threshold for pulse finding can be set *relative* to the baseline, rather than having to know a fixed number—which again is problematic due to the moving baseline.

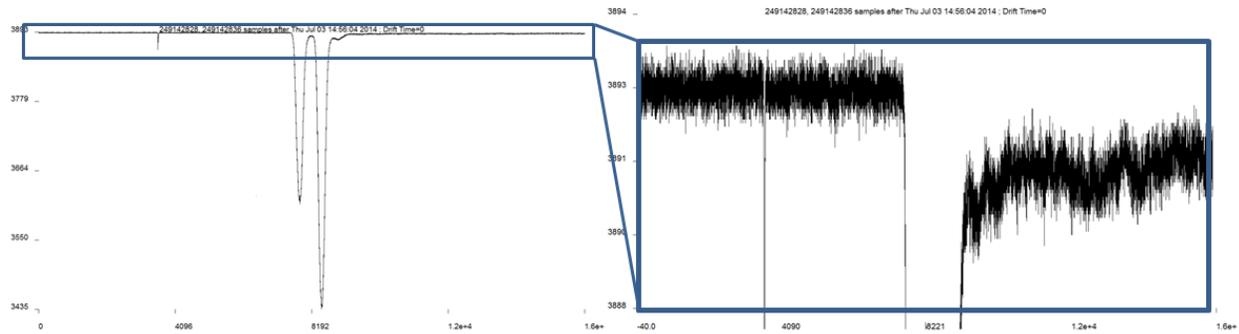


Figure 3.2: Example waveform of Th228 at 2 ADC*samples/phe gain. Left: The full waveform view which illustrates the difficulty in visually noticing the shift. Right: A zoomed in view of the baseline. The baseline shift appears after the large S2 from a Thorium event is visible. The size of the shift is typically between 2-7 samples, and the shift depends on the size of the S2.

Another issue is that large S2s would cause an immediate baseline shift after the S2 pulse, typically on the order of 2-7 samples. Since the baseline is a fixed measurement and is typically smaller than the new baseline after the large S2 pulse this would cause many baseline regions to be erroneously flagged as pulse regions. This problem was solved by imposing a height threshold on these erroneously flagged ROIs (region of interest). By setting the height requirement to be just 5-10 samples

these erroneous regions never get identified as an S1 or S2, and thus does they did not factor into any other calculations.

3.2.2 All Non-Pulse Samples

The second, more technically complicated but computationally elegant technique was using all points that were not in pulse regions. This method requires a fixed threshold to be used for all of the channels and thus the algorithm already has an assumption about where the baseline should be. Another way of thinking about this algorithm is that it is a software version of zero suppression. However, for most zero suppression algorithms to work well it is assumed that the baseline will be a constant throughout data acquisition. Since the baseline in PIXeY could shift during acquisition it meant this version of baseline determination was problematic, and thus was not used.

3.2.3 Other Considerations and Future Improvements

Most computers when translating a floating or continuous value to an integral value simply do a floor function, or cut off the decimal values. So, numbers like 3.2, 3.4, and 3.6 all become 3. The conversion of the analog PMT signal to a digitized computer information signal can introduce this error, so this type of rounding is not easily changeable, if at all. One consideration on fixing this was to shift the baseline by $\frac{1}{2}$ sample to compensate for any round-off error that may have occurred during the digitization process.

As shown in Figure 3.3, early tests of this adjustment showed a 0-1% improvement in the energy resolution of the S2 for a ^{137}Cs test run, although the improvement was less than $1\text{-}\sigma$ from the error. It wasn't until later analysis confirmed that the baseline adjustments did not improve the energy resolution in PIXeY, which was apparent in the combined signal. Combined energy spectra were compared that had and didn't have the baseline adjustment. The combined energy spectra with the baseline adjustments were measurably worse, so the adjustment of this type was abandoned. This effect is shown in Figure 3.4.

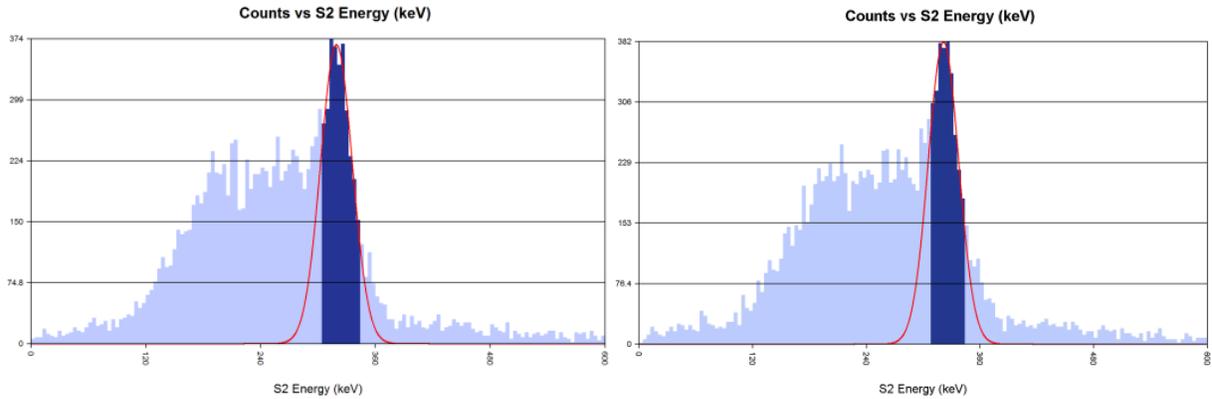


Figure 3.3: Comparison of a Cs137 spectrum with the left having no baseline adjustment and the right having a baseline adjustment. The measured S2 energy resolutions for these spectra are $5.18 \pm 0.36\%$ and $5.00 \pm 0.35\%$, so the improvement was less than 1-sigma.

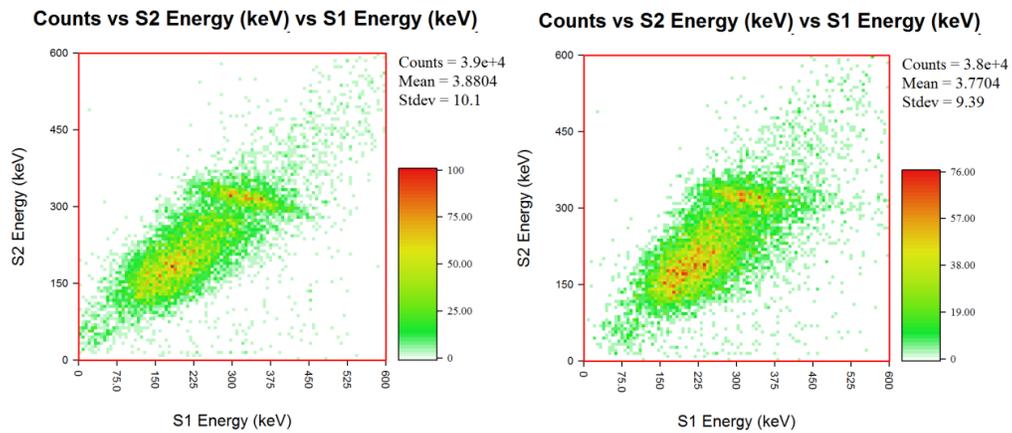


Figure 3.4: Measured S1 S2 pixel plot of a Ba133 spectrum. The left figure is without baseline adjustment and the right figure is with the baseline adjustment. The baseline adjustment seems to negatively affect the region in which the S2 energy is low or high and where the S1 energy is high, going from 2.3 ± 0.16 to $3.6 \pm 0.25 \sigma/E$ Combined.

One final thing to consider is the possibility of doing an interpolated baseline for the S2 pulses. The basic idea is this: there is a measurable change in the baseline after a particularly large S2 pulse, as shown in Figure 3.2. However, the baseline that's used to calculate the S2 pulse is from before the shift, so it's not quite the "true" baseline of the S2 pulse. The baseline both before and after an S2 pulse can be

calculated and then linearly interpolated, which would act as a moving baseline for the S2 pulse area calculation. This idea is illustrated in Figure 3.5.

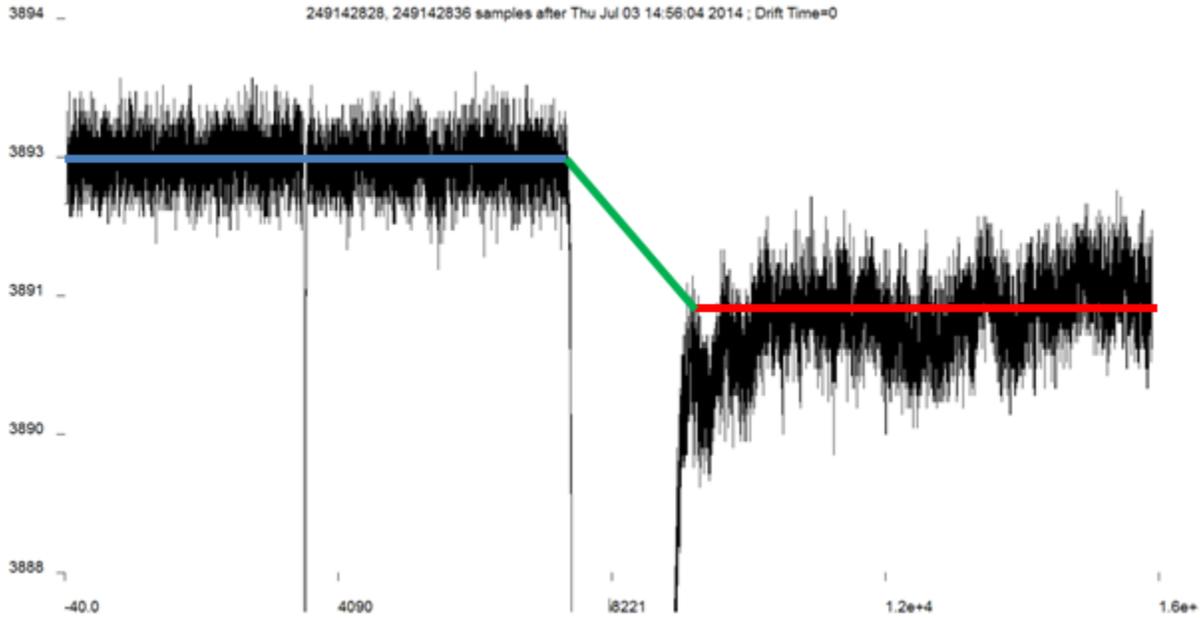


Figure 3.5: Visualization of using an interpolating algorithm for a moving baseline for calculating the S2 pulse area. The blue line is the baseline before the S2 pulse, the red line is the baseline after the S2 pulse and the green line would be the interpolated baseline that the algorithm would use to calculate the S2 area.

The structure of the analysis code is iterative (as shown in section 3.1) so this would be an iterative step since a region would first have to be identified before this modified baseline calculation could be done. Also there would be difficulty when there are multiple S2 pulses to consider, since many a time during a Compton scatter event there aren't many samples that return to baseline, if at all.

An interpolated baseline would attempt to fix any sort of asymmetry in the peak region. This is because large S2 pulses produce a larger baseline shift, so that biases a peak region to the right hand side of the peak. This would be unusual since most biases are to the left from incomplete signal collection. As

an example in Figure 3.5 the triangle area is about 2,000 ADC*samples, but the pulse area is about 458,000 ADC*samples, so this correction is about 0.4% of the pulse area. Since there wasn't any large right hand asymmetry observed in the S2 and combined peaks this algorithm was never implemented, but it remains as a consideration if any future experiment has a similar issue with a moving baseline and would need an implementation fix, or when sub-percent errors in a detector start to become the dominant error levels.

3.3 Pulse Finding

Since no two experiments are generally the same most pulse finding software is written from scratch. For PIXeY many different pulse-finding techniques were implemented then tested on simulated data. The metrics were based on pulse resolution and pulse thresholding. There were three different pulse finding techniques that were studied: pulse finding using the sum waveform with a voltage or height threshold, pulse finding using each individual PMT signal with a voltage or height threshold, and pulse finding using the sum waveform by doing a sliding dual box integrated signal threshold.

3.3.1 Sum Waveform

This method involves summing up the all of the PMTs and then performing a threshold analysis on the signals and identifying all of the various pulses. Of all the methods it is by far the simplest, since it's essentially only looking at one channel at a time to do pulse finding. The pulse regions of interest are then fed to the individual PMT channels where all of the signal metrics are found (area, width, height, etc.).

3.3.2 Waveform by Waveform

The second method was the waveform by waveform implementation. This was similar to the sum waveform as above but instead looking at each PMT signal individually to perform pulse finding analysis. The main motivation would be to find better ROI bounds that more accurately reflect the region of interest for each signal, rather than simply copying the information from the sum signal.

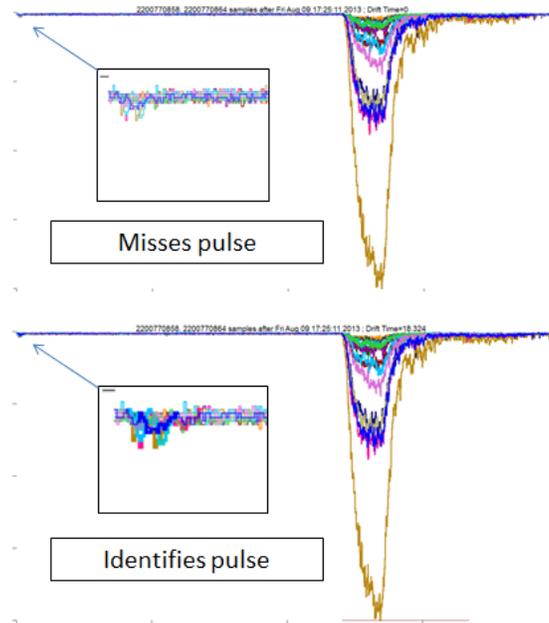


Figure 3.6: Top: pulse finding using sum waveform. Bottom: pulse finding using waveform by waveform. Since the waveform by waveform pulse-finding focuses only on each individual PMT it is able to find much smaller pulses that the sum waveform may miss.

This technique, while potentially useful, had a number of technical setbacks. It turns pulse-finding from a 1D problem to a 14D problem, since now each PMT has its own information on the pulse regions of interest. Different channels won't agree on the locations of all pulses and thus there has to be some higher order logic that can group pulses together into a cluster that represents the same event, which multiplies the algorithmic complexity. This algorithm, while possible to construct, became very computationally expensive and time consuming and was ultimately not used due to complexity of implementation.

3.3.3 Double Box Method (Sum)

The final method that was tested was area threshold analysis using a double sliding box. This technique integrates the signal in a box window and thus its threshold is set on area, not height. The sizes of the boxes are based on the widths of the S1 and S2 pulses, since their approximate widths are known beforehand. The S2 size window is scanned first finding any regions of interest that would have a size that is S2-like. The second window scans the regions not already identified by the large box to check if there are any smaller S1-like pulses. This operation, while theoretically able to be performed on the individual PMT waveforms, was done on the sum instead for simplicity.

3.3.4 Threshold Testing

The different pulse finding techniques were tested; however the waveform by waveform method was not investigated further due to complexity. The area threshold method is most often used when trying to find small pulses, which is important for WIMP dark matter searches [3]. It was necessary to investigate if this method would be useful for this application space, since it was assumed that the area threshold method was the superior pulse finding technique.

The performances of these methods were tested using simulated data, which was generated as a continuous pulse area spectrum. This spectrum was convolved with simple background noise representing error in the baseline of 2-4 samples of signal deviation. This simulated data used in the different pulse-finding algorithms to determine if there was a significant difference in threshold response.

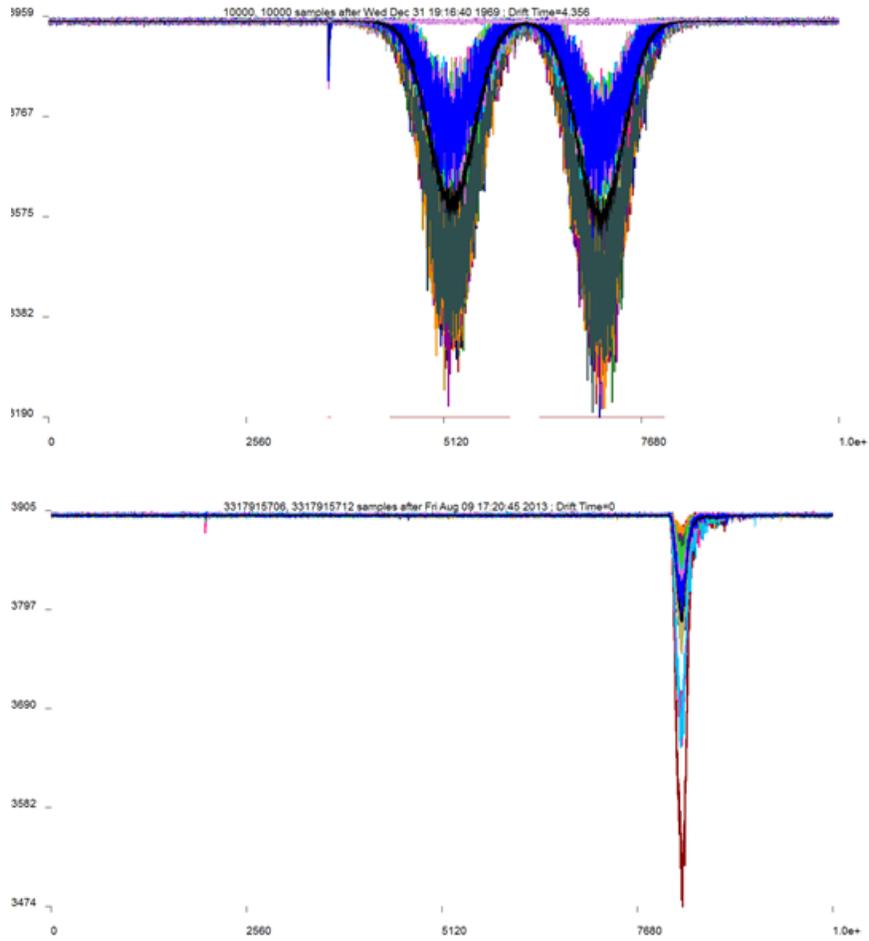


Figure 3.7: Comparing simulated data to real S1 S2 data for the purposes of threshold determination. Top: simulated S1 S2 double S2 Compton Scatter. Bottom: real data of Cs137 test.

The results were that the area pulse finding method and voltage pulse finding method have nearly identical low energy thresholds, which for a gain of 5 mV-ns/phe was about 1 keV. It shows that the implementation of more complex pulse finding algorithms doesn't necessarily increase performance, at least for the purposes of PIXeY data. For other systems in which the accuracy of small pulses is more important the pulse area method could be used. However for the sake of making a system that is easy to understand for our purposes the voltage threshold method for pulse finding was used.

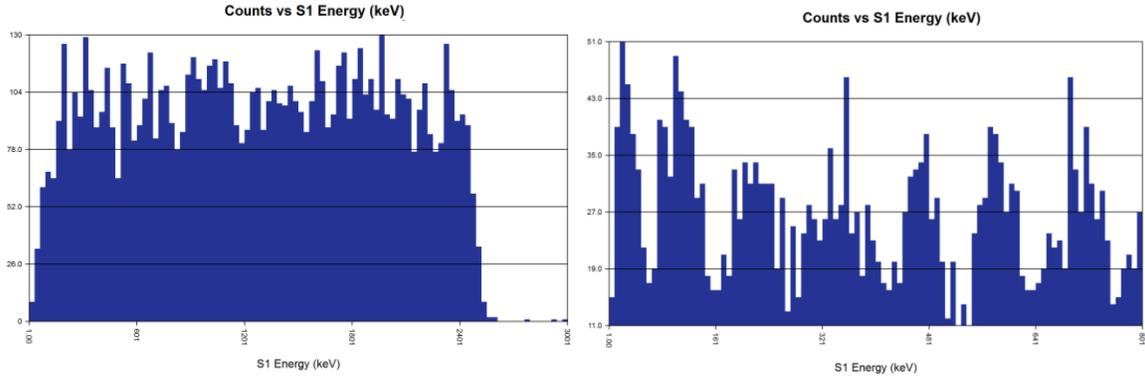


Figure 3.8: Comparison of voltage threshold method and area threshold method. Left: voltage threshold method. Right: area threshold method. Note: area threshold method is zoomed in to the first fifth of the voltage threshold method picture. The peaks and troughs of the data are due to artifacts in the generated data.

Finally, it was found that if the sum waveform method was used on the bottom PMTs only, then this could boost the sensitivity to small signals without the increase in complexity. This would allow some improvement that was shown in the waveform by waveform method but still keeping the relatively simple sum waveform pulse finder. The reason for this has to do with light collection profiles in the detector.

It was found for PIXeY that on average about 80% of the S1 light is entirely contained in the bottom PMT array, which is mainly due to total internal reflection [60]. Also for the S2, about 55% of the S2 light on average is in the bottom PMT array as well. This is due to the fact that geometrically, PIXeY has a relatively short depth and large width (5cm and 15cm), so its aspect ratio is about 3. Most other detectors have much larger depths compared to their widths, such as LUX which has an aspect ratio of 1.02 [61]. With a larger aspect ratio PIXeY captures most of its S2 light in the bottom array, because of total internal reflection. This means that on average adding the top array to pulse-finding adds more noise than it adds signal for calculating the pulse regions of interest. The study of the light collection profiles in PIXeY is done in Section 5.2.2.

3.4 Identifying S1 and S2

Many different methods of classifying S1 and S2 pulses have been developed by previous experiments, mostly aimed at very low energy pulses for WIMP dark matter studies [62]. All parts of the analysis chain can affect the final dark matter threshold, so these experiments typically want to push the ability to see small signals. For this experiment since the gain on the PMT was freely adjusted and the lowest energy probed was 42 keV this consideration was less important. Thus for this analysis upper and lower bounds on the full width at half max (FWHM) were set in order to determine if a pulse was S1 or S2-like.

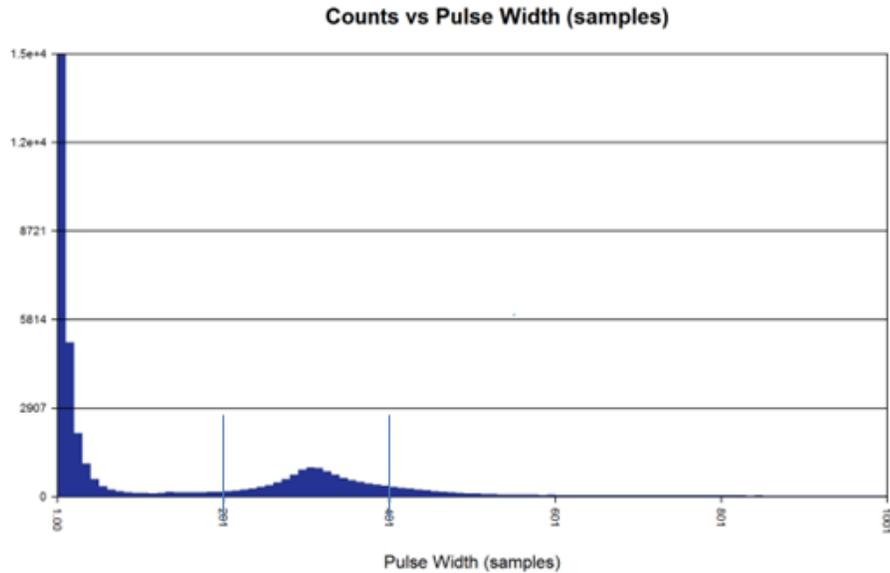


Figure 3.9: Width spectra of all of the pulses. The blue bars indicate the region of acceptance for S2 pulses. The spike on the left represents small pulses and S1 like pulses. Care must be taken to move the window if the anode voltage is changed since the S2 pulse width is anode dependent.

The FWHM for each pulse was calculated for each individual waveform, and were compared to upper and lower bounds for S1 and S2-like pulses. Physically, the S1 bounds are well defined by the triplet and singlet lifetime state [12], and the S2 width is well-defined by the width of the gas gap in the

proportional scintillation region. If 4 PMTs had a pulse that was in the specified bounds, then it was categorized as an S1 or an S2. Finally, a small height threshold was placed on the S2-like pulses to filter out erroneous baseline measurements, which was explained in Section 3.2.1.

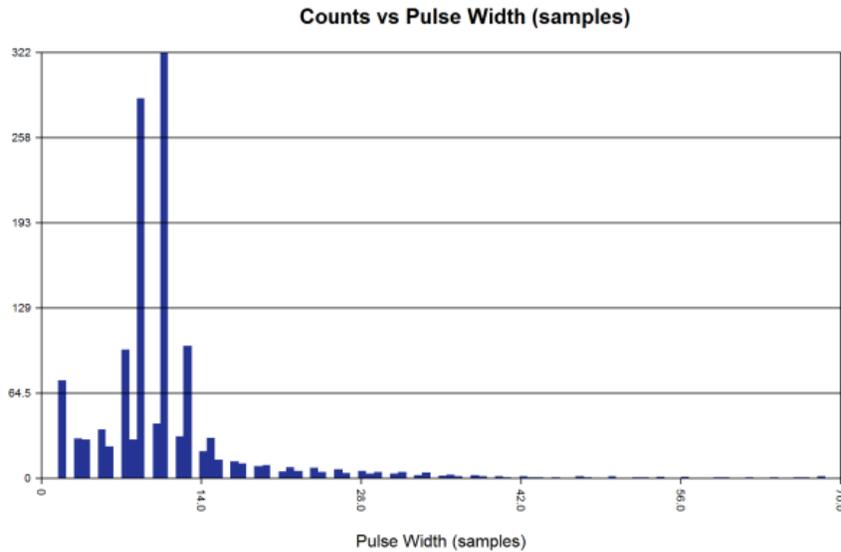


Figure 3.10: Zoomed in plot showing the S1 width spectrum that is used to make selection of S1 events. Some bins are not filled due to a rounding error that moves values from one bin to another.

For many experiments it can be difficult to find a lower bound for the S1 pulse verification, so very small pulses are difficult to filter. As shown in Figure 3.10, a lower bound can be found for the S1 spectrum as well as an upper bound, which may be because of the 4-ns timing resolution and the fact that the PMT gains were tuned to maximally resolve S1 pulses. The lower and upper bounds should not change with anode field, but should change slightly with cathode field due to suppression of recombination, which has a lifetime of about 120 ns in zero-field [63]. This effect however is small enough though that a large enough upper bound can be since there are so few events in the tail of the spectrum. Thus in general the bounds for S1 were set to be between 4-50 samples, or 16-200 ns.

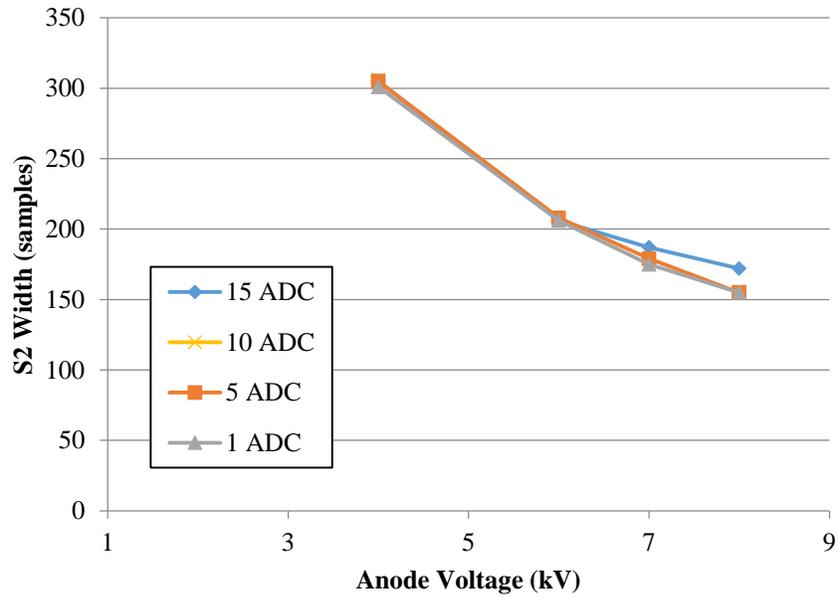


Figure 3.11: S2 width peaks (in samples) as a function of Anode voltage using the injected Kr-83m source. The spread at the high anode voltage may be due to some saturation occurring at the high anode voltage region.

The S2 acceptance window has to be adjusted depending on the value of the Anode voltage. The anode voltage increases the electric field in the gas, which causes the drifting electrons in the gas to drift faster. Because the S2 can be thought of as the time integral of the electron path signal, the width of this S2 pulse is shortened, which is shown in Figure 3.11. Eventually for the rest of the energy resolution studies the anode was kept at a fixed voltage of 5 kV, which is explained in Section 4.9. The electric field in the gas was calculated to be 8.3 kV/cm, which is explained in Appendix A.3.

3.5 Pulse Declustering

During normal operation, there were some waveform events that would categorize pulses incorrectly, specifically Compton scatter events. These types of events produce more than one S2 per one S1 of the event. Often times these S2 pulses are very close to each other in time. What can happen is that

the pulses never really return to baseline before the second pulse happens. As such, it's relatively easy for two pulses to be grouped together and counted as the same S2 pulse in a pulse finding algorithm. And while this does not alter the sum area calculation for most of this work multiple scatter events were excluded, so it was desirable to have the correct classification algorithm in order to filter these events.

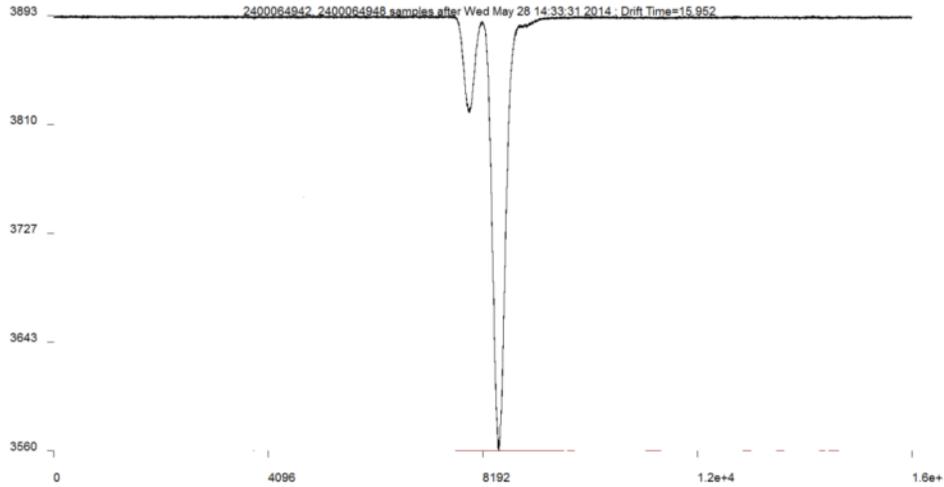


Figure 3.12: Waveform with a Compton scatter S2. The red bars on the bottom indicate the ROIs that the program classifies as S2 pulses. Notice that one large solid bar is on the two S2 pulses, whereas it should be two. Also this picture points out another issue with the pulse-finder which is spurious events after a baseline shift being labeled as S2. While the pulse declustering was much more complicated to work out the spurious post-S2 events were filtered with a simple height check requirement, which is explained in more detail in Section 3.2.1.

In order to fix this the following algorithm was implemented: any identified region went through a “pulse declustering” check. This part of the code looked at all ROIs that were flagged as being S2-like². The derivative of the waveform in the ROI was computed. This was done through several steps: first the original waveform was smoothed using an n-period average smoothing, then the slope of the curve was computed using an n-point secant interpolation, and then finally the slope curve is also smoothed using n-period moving average as well. Typical values for n were between 3-5 samples, and were tuned based on the specificity of the S2-like pulses.

²One of the computational advantages of this was that the code had knowledge that it was most likely an S2 at this point, previous implementations of this algorithm were used at the pulse finding step and thus implemented this algorithm on every pulse region, which significantly slowed computation time.

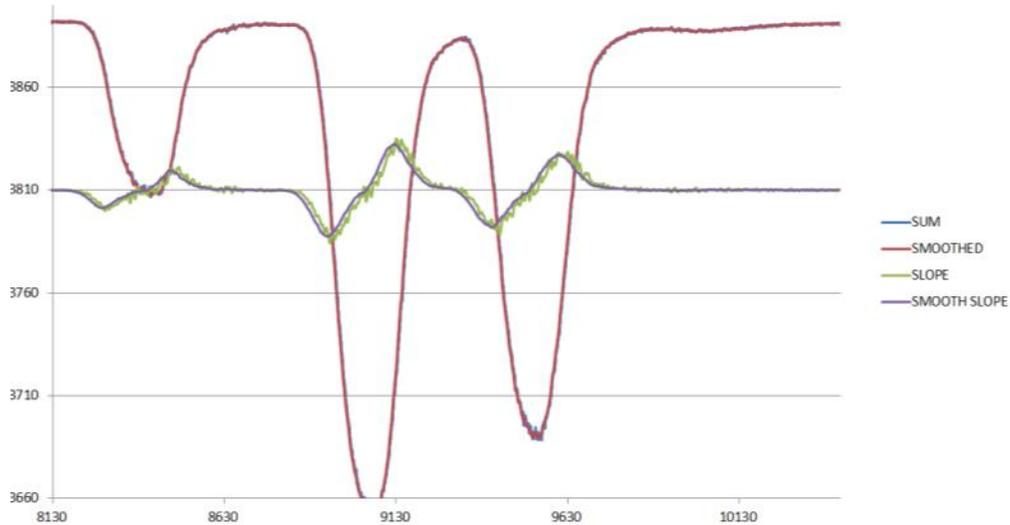


Figure 3.13: Slope deconvolution of a multiple S2 pulse event. The waveform goes through various stages of smoothing, differentiation, and smoothing again before analysis is done on the derivative of the waveform.

Once the derivative was calculated the code looked for zero crossings in the data. Specifically, it looked for places that went from positive to negative or negative to positive in some smooth transitional way. An assumption is made on the order of extrema and thus all pulses can be independently found. Once they were found in this way the S2 ROI was then split up according to the number of minima's found by using information from the maxima as references as to where to break up the ROI into multiple regions.

Once the separate regions of interest were identified, the modified information was sent back to the first pass analysis step (after the pulse finding step) to have all of the areas, times, heights, and other waveform metrics recalculated. This was part of the iterative step shown in Section 3.1 and was one of the main reasons for an iterative data analyzer.

4 Detector Diagnostics

4.1 Introduction

The successful operation of any detector requires multiple measurements and calibration studies that measure the health of the detector before scientific studies can be performed. These include calibrating and normalizing PMT gain, testing electric fields in the detector, and determining and adjusting the liquid height for optimal operation. Work was done in each of these areas to ensure that the energy resolution measurements were made with a detector in optimal condition.

4.2 PMT Calibrations

The first step in preparing any detector system is making sure all sensors are calibrated properly, which for PIXeY would be the PMTs which detect both the S1 and S2 light. The gain curves of all PMTs were mapped so that a variety of gains could be used to fully optimize the energy resolution. Internally installed LEDs that were on the top and bottom PMT block were pulsed while the system was filled with cold gas for calibration purposes.

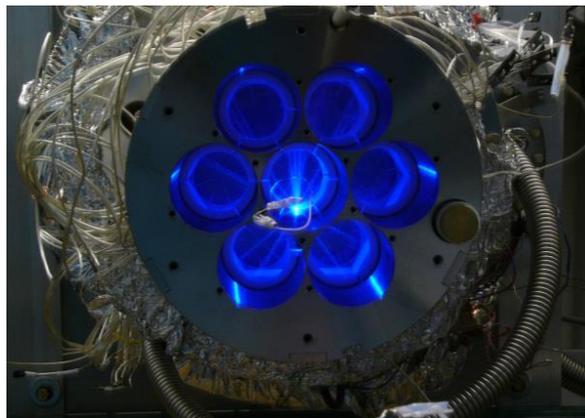


Figure 4.1: Image of bottom of detector being illuminated by the top LED. The LEDs were installed on Teflon spacers.

The standard way to calibrate the gain of a PMT is to measure the single photoelectron peak, which is the peak associated with the emission of a single photoelectron off of the photocathode of the PMT. Using this signal size can be converted from mV-ns or ADC*samples (the units used in the DAQ) to photoelectrons, or phe. This is found by turning on the PMT and looking for dark counts: stray light and/or electrons may make an excitation on the photocathode which will cause a single electron to be released. By looking at the spectrum of events there should be a peak associated with the single photoelectron as well as an exponential tail that corresponds to integrated background noise. The exponential noise and single phe peak are the most prominent features in the signal, although there are many other features present [57].

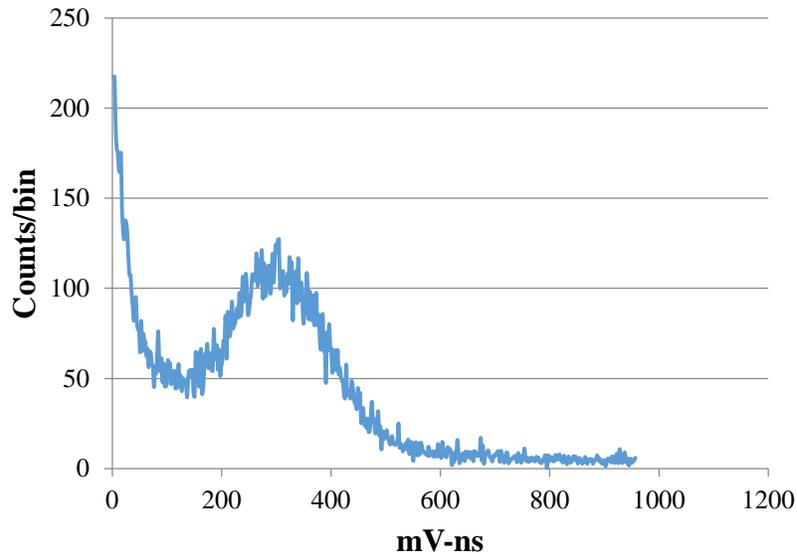


Figure 4.2: Example of a single photoelectron spectrum. The continuous slope towards zero area is background and baseline integration, with the peak being the first single photoelectron peak.

Usually the gain has to be sufficiently high in order to observe such an effect. For detectors where PMT gains need to be high (such as in WIMP dark matter detectors) this means that the gains can be calibrated directly at the operating PMT voltages. However for PIXeY most of the energies studied are

much higher than this, which meant that PMT voltages had to be kept low to avoid saturation. Thus it became increasingly difficult to find the single phe peak in the data itself.

The single photoelectron measurement could not be done at voltages lower than 1600 V for all of the PMTs, so the signal response for voltages ranging from 1000-V to 1600-V were mapped with an LED pulse of finite size. With this the parameters of the gain curve could be found, and then the relative parameter values needed to be calibrated to an absolute measurement. The single photoelectron was found for 1800V, 1600V, and 1400V where possible and those values were used to position the relative gain curve to an absolute scale. These gain points as well as a fit for them are shown on Figure 4.5.

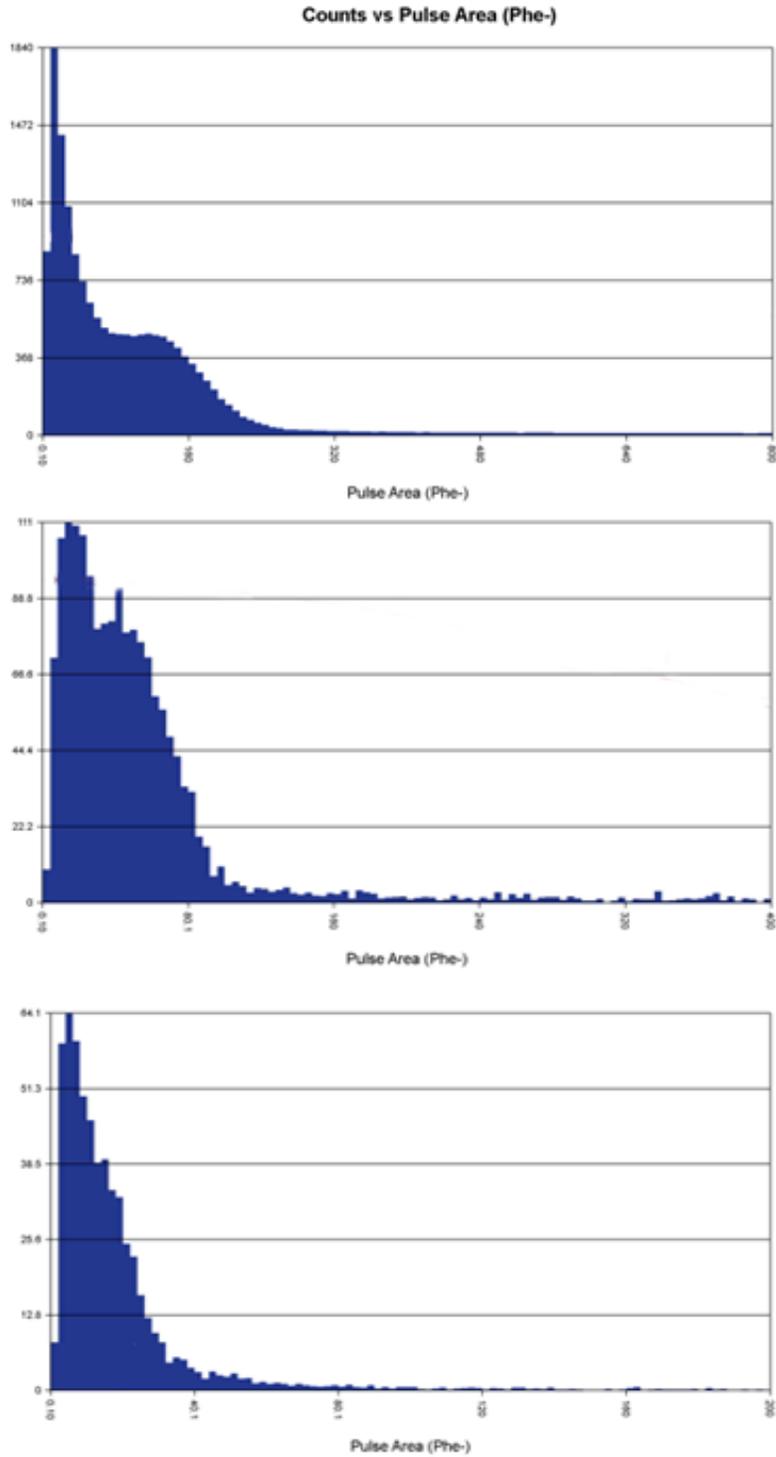


Figure 4.3: Comparison of the single photoelectron changing as a function of the top center PMT voltage. Top: 1800 V Middle: 1600 V Bottom: 1400 V. As the voltage gets lower it gets harder to accurately fit the single photoelectron peak.

Calibrating the gains based on a finite pulse size can be explained in the following way. Assume that the gain of the PMT can be modeled as being proportional to the voltage to the power of some constant. That is:

$$Gain = A * V^B \quad (4.1)$$

where A and B are constants that are the fit parameters needed for calibration. Now suppose that some reference values are known:

$$Gain_{Ref} = A * V_{Ref}^B \quad (4.2)$$

We can divide these two equations and rewrite in terms of gain:

$$Gain = Gain_{Ref} * A * \left(\frac{V}{V_{Ref}} \right)^B \quad (4.3)$$

with the constant A technically cancels but we restore it here to take care of any asymmetry. We use reference values because typically the unnormalized value of A is extremely small (order of 10^{-38}), which could propagate rounding error to other calculations. Now, for a finite sized pulse it's simply the intrinsic light content or "size" of the pulse times the gain:

$$\begin{aligned} Signal &= Size * Gain \\ Signal_{Ref} &= Size * Gain_{ref} \end{aligned} \quad (4.4)$$

with the same sized pulse being tested at different gains. This means that the ratio between the signals is equal:

$$\frac{Signal}{Signal_{Ref}} = \frac{Gain}{Gain_{Ref}} \quad (4.5)$$

But the ratio on the right hand side is known, which can be re-written in terms of voltages:

$$\frac{Signal}{Signal_{Ref}} = \left(\frac{V}{V_{Ref}} \right)^B \quad (4.6)$$

Thus the constant B can be calculated from the signal without knowing the absolute calibration of the gain. So the strategy will be to find the value of B for all of the PMTs and then use a few single phe measurements at high gain to provide an absolute calibration for the measurements taken at finite pulse sizes.

Finally, the gain can be calculated based on the finite signal and the reference values:

$$Gain = Gain_{Ref} * \frac{Signal}{Signal_{Ref}} \quad (4.7)$$

provided that the reference here for the gain and the signal is the same.

Additionally since there were two LEDs in the system it was possible to calculate the PMT gains based on several different methods, such as using the top or bottom LED as the absolute calibration measurement, or just doing a simple extrapolation based on the known single phe values. This was done to see if there were any systematic differences between the measurements using the top LED, bottom LED, or a gain extrapolation.

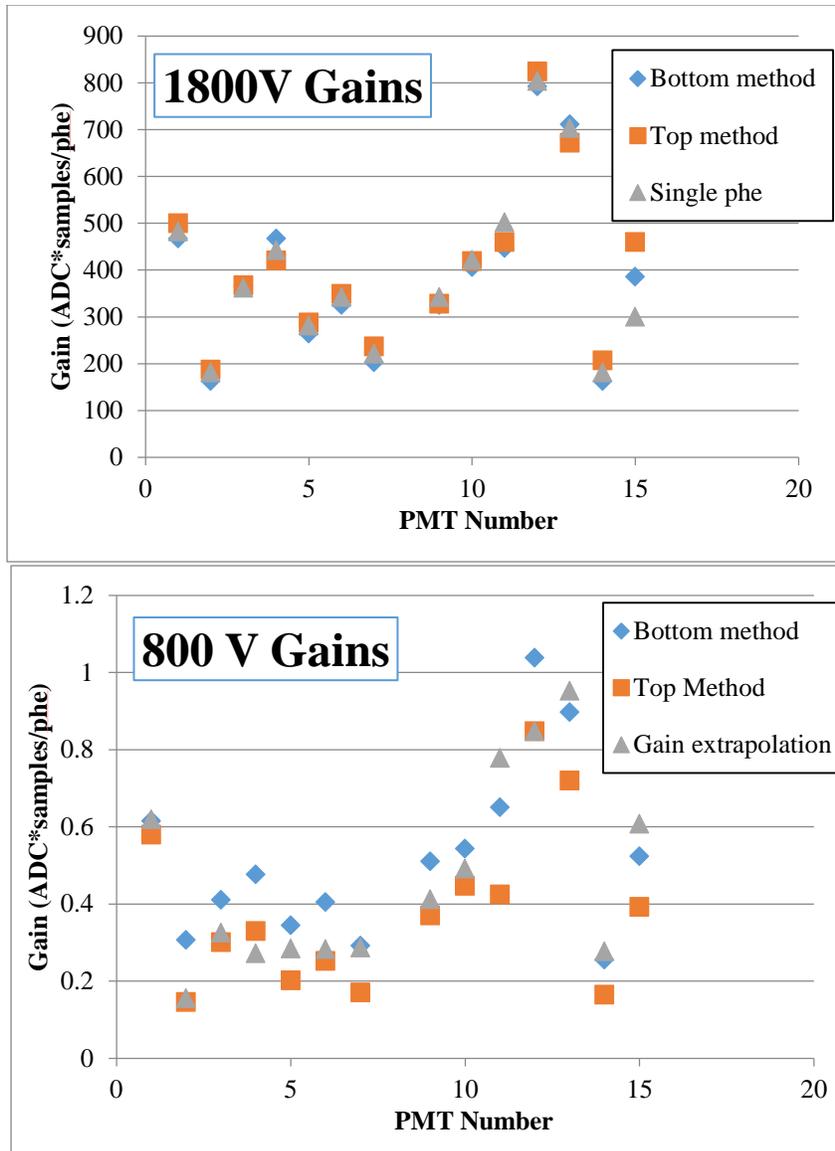


Figure 4.4: Comparing different methods of determining gain, at a high PMT voltage of 1800 Volts (top) and a low PMT voltage of 800V (bottom) for all 14 PMTs.

For the 1800 V gains, the average variation is 4.7%, with the worst variation of 9.8%, which occurs on PMT 15. For the 800 V gains the average variation is 25% with a worst variation of 41% on PMT 11. This shows that in the region where the single phe is known well (1800V) there was low systematic uncertainty since all of the values have a much smaller variance. However, the single phe was not directly calculable at 800 V and thus has a much higher systematic uncertainty. It's because of this

level of systematic uncertainty at low PMT voltages that eventually inspired the data optimization methods explain in section 5.3.3. Nonetheless, all of this info was combined to extrapolate the final gain curves that were used for PIXeY.

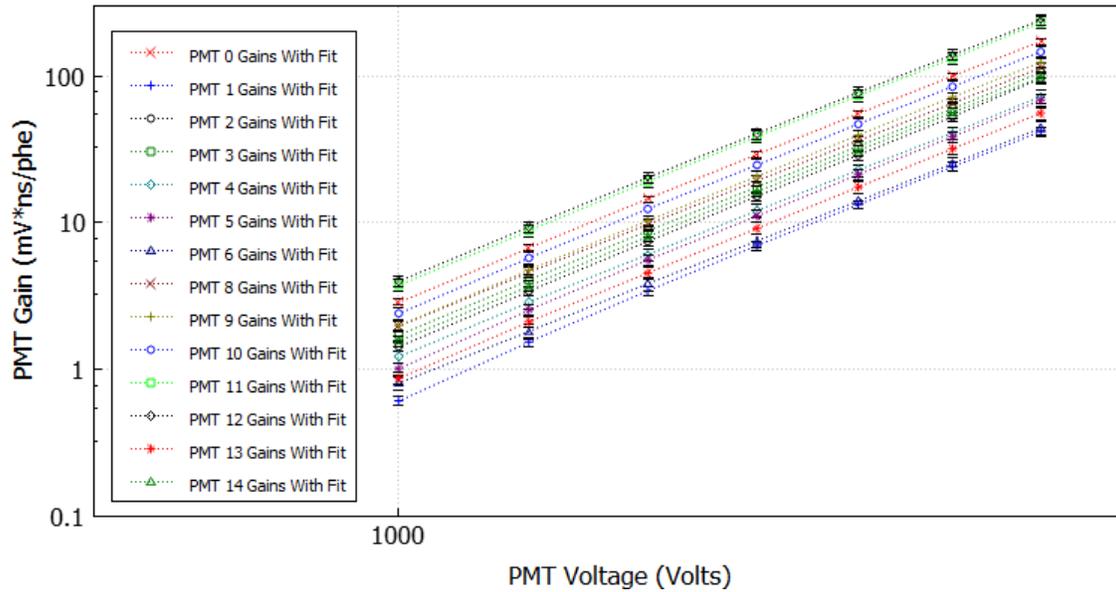


Figure 4.5: Calibration of all 14 PMTs with calibrated fit line.

4.3 LED Saturation Tests

4.3.1 Introduction

It is desirable to know the saturation levels of the PMTs before any sources were used. The PMTs are wired to a set of capacitors that control the voltages on the various stages of the dynodes inside the PMT; if any of those capacitors were to deplete during operation then this would cause the PMT response to go down. This is known as saturation because PMT response is flattening out or “saturating” and is no longer behaving linearly. This is generally bad because the underlying assumption while using these detectors is that the response is linear; if at some point that assumption breaks down then this is one of the primary places that could lead to worsening of energy resolution.

A variety of tests studied to probe the extent of the saturation using the installed LEDs that were used for PMT calibration. Since there were two LEDs in the system they could be pulsed separately with different sizes and frequencies to see how the two signals compared.

4.3.2 Pulse Linearity Studies

One such test was called the simultaneous test, which was to see if pulses added linearly, which would be expected if there is no saturation present. The procedure of this test was to measure the size of two pulses by themselves then add them together at the same time to see if the sum of the signal is equal to the sum of each signal individually. It is here that we define the fractional signal difference as:

$$\delta_{fractional} = (S_{BLED} + S_{TLED} - S_{BLED||TLED}) / (S_{BLED} + S_{TLED}) \quad (4.8)$$

Where S_{BLED} is the signal size from the bottom LED, S_{TLED} is the signal from the top LED, and $S_{BLED||TLED}$ is the signal from the two run simultaneously. Thus the fractional difference should be 0 if the signals are adding as expected and +1 if the simultaneous signal is extremely small. The interesting thing about this test is that we did not expect to see any values go negative, which is how this metric was designed. However when the tests were performed it was observed that this fractional difference could go negative. This is most likely an “enhancement effect” in which the response actually increases slightly before starting to saturate. This was an effect that was observed during PMT testing.

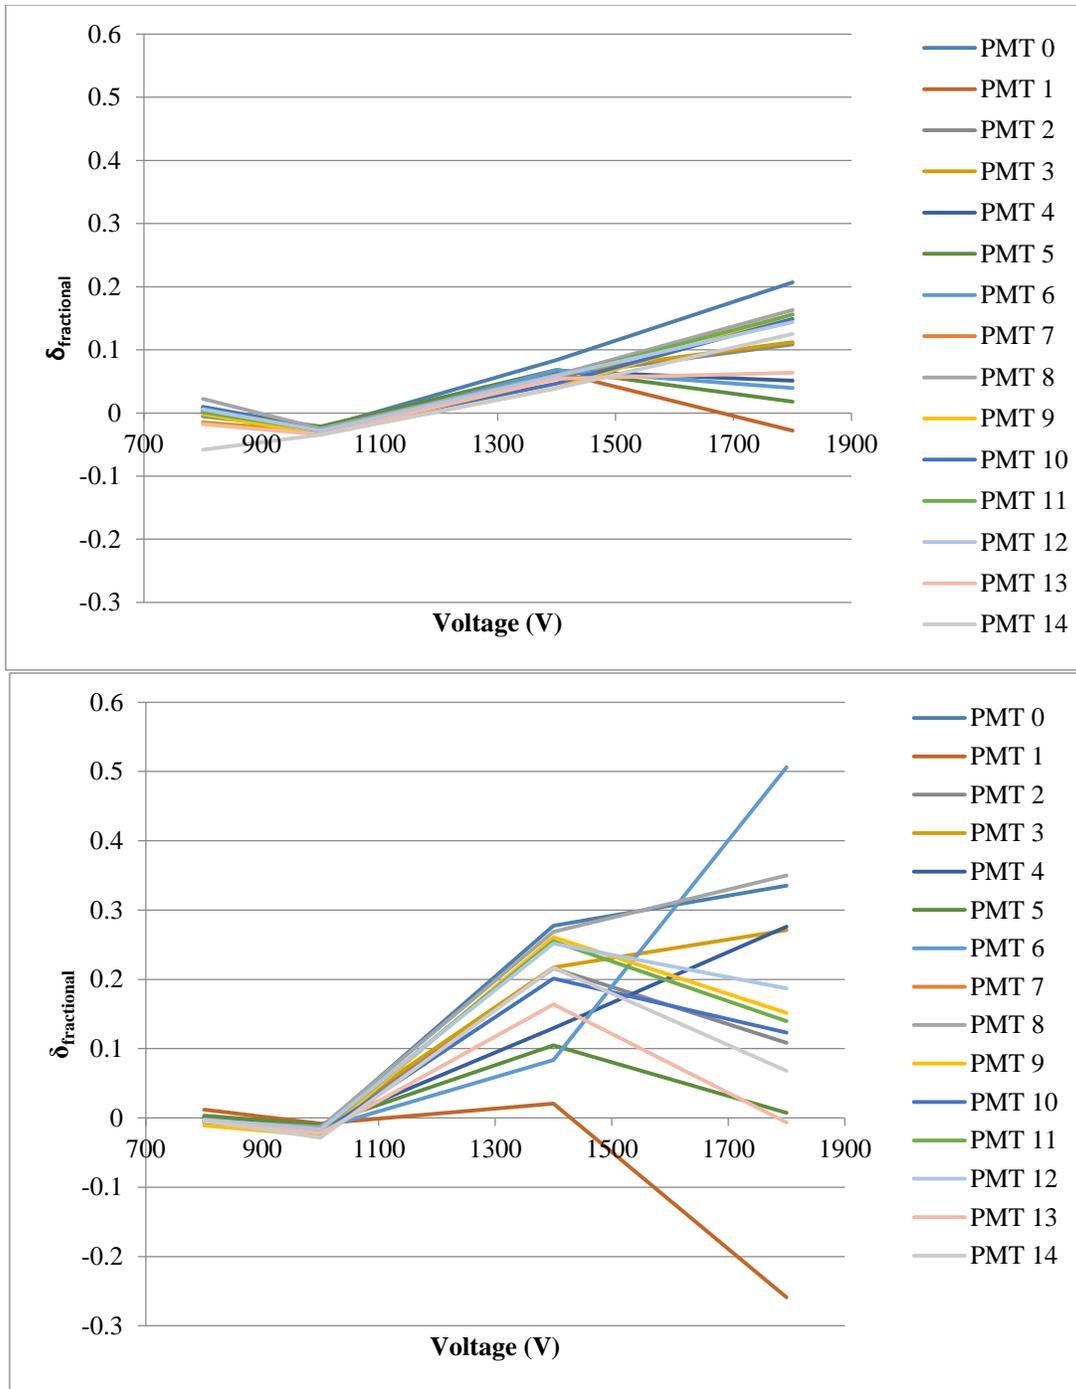


Figure 4.6: Top: simultaneous test with LED settings of 2.6V and 2 μs window on the top LED and 3.40V and 2 μs window on the bottom LED, both pulsed at 100 Hz. Bottom: simultaneous test with PMT settings of LED settings of 2.65V and 2 μs window on the top LED and 3.50V and 2 μs window on the bottom LED, both pulsed at 100 Hz. As expected the larger pulses have higher non-linearity, and thus have a larger fractional difference.

As shown in Figure 4.6, the results demonstrate the limitation on the PMTs with respect to increasing voltage. Also noted previously there was also the interesting observation of a small “enhancement” effect that occurs before normal saturation, which was about $-0.04 \delta_{fractional}$. This was seen as a clue to see if saturation is about to take place. There is a slight enhancement effect for both test pulse configurations at around 1,000 V but then the PMTs begin to saturate soon after. The max saturation is only about 20% with the average being about 10% for the top configuration in Figure 4.6 because it has less of an overall amount phe, as shown in Figure 4.7. For the bottom configuration, the saturation is as high as 60% with the average about 20% although the deviation was about 8% for the bottom configuration and only about 3% for the top configuration. This could be due to the total amount of light that each PMT saw, which was fixed since the LEDs did not move in their relative position. It is unknown why PMT1 in the bottom configuration had a large negative enhancement at high PMT voltages.

As shown in Figure 4.17, looking at the sum total of the signals the response was decreasing as a function of increasing voltage. The bottom configuration peaks at 1000 V at 700k phe and then decreases to about 14% of its signal size. This can only be explained by saturation of the PMTs which for all runs began at 1000 V with a signal size greater than 500k phe. Not only that, but the individual pulses themselves were shown to be saturating. So even if this linearity tests were not done, we can still see the saturation of a pulse by just tuning the voltage and looking for the single response. This is what inspired the later test mentioned in section 4.3.4. Based on the total phes of this test this meant that saturation would occur at around 500,000 phe total signal size if the voltage on the PMTs was on average higher than 1,000V.

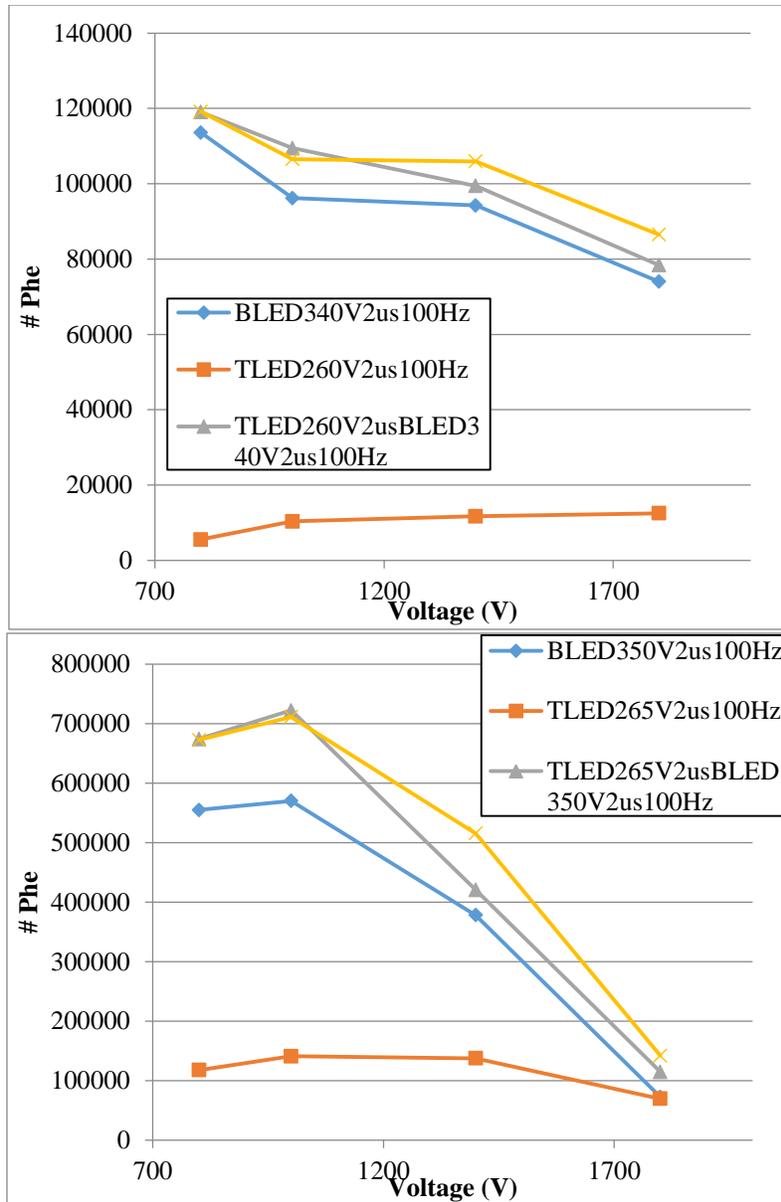


Figure 4.7: Total Phe of the two simultaneous tests as a function of voltage. It is clear at these levels that each large pulse is self-saturating, since the LED parameters were not changed during each data point.

4.3.3 Frequency-Based Saturation

The second test that was done using the LEDs was called the stagger test, which was to see if pulse area is also affected by the presence of other pulses at varying frequencies and at different times.

The test is designed to approximate backgrounds of various rate frequencies to see how the PMTs respond and if they respond linearly. Once again we define a fractional difference for the stagger tests as:

$$\delta_{fractional} = (S_{TLED} - S_{TLED|BLED})/S_{TLED} \quad (4.9)$$

with the signal sums having their same meanings as above. So here once again 0 would mean that there is no difference between the two different situations, and +1 is if there is heavy saturation on the signal when both pulses are present.

As shown in Figure 4.8, at 1000 Volts there is evidence of large enhancement effects at higher frequencies, as much as 40% enhancement with an average of about 30% at 10 kHz and 1000V. Even for moderately sized pulses and smaller voltages like 800 V, there still 10% change at 1 kHz. These signals were normalized at 100 Hz, which is why these graphs have 0 fractional differences at that value.

This type of saturation was termed frequency-dependent saturation, and the previous type was called self or size-dependent saturation. Basically, it is the non-linear response that occurs in the presence of some background events at a certain frequency. This frequency-dependent saturation would only lead to a broadening of the energy resolution if the background or source rate were constantly changing. However the background and source rate for the detector was approximately constant, with rates usually between 100-300 Hz, which means this would matter if measurements with a source frequency of over 10 kHz were studied. In experiments with active interrogation event pile-up and frequency-based saturation effect is more important due to the sheer high number of events. Even for sources that were in the range of 2-20 μCi , the raw event rate in the detector was never higher than 2 kHz.

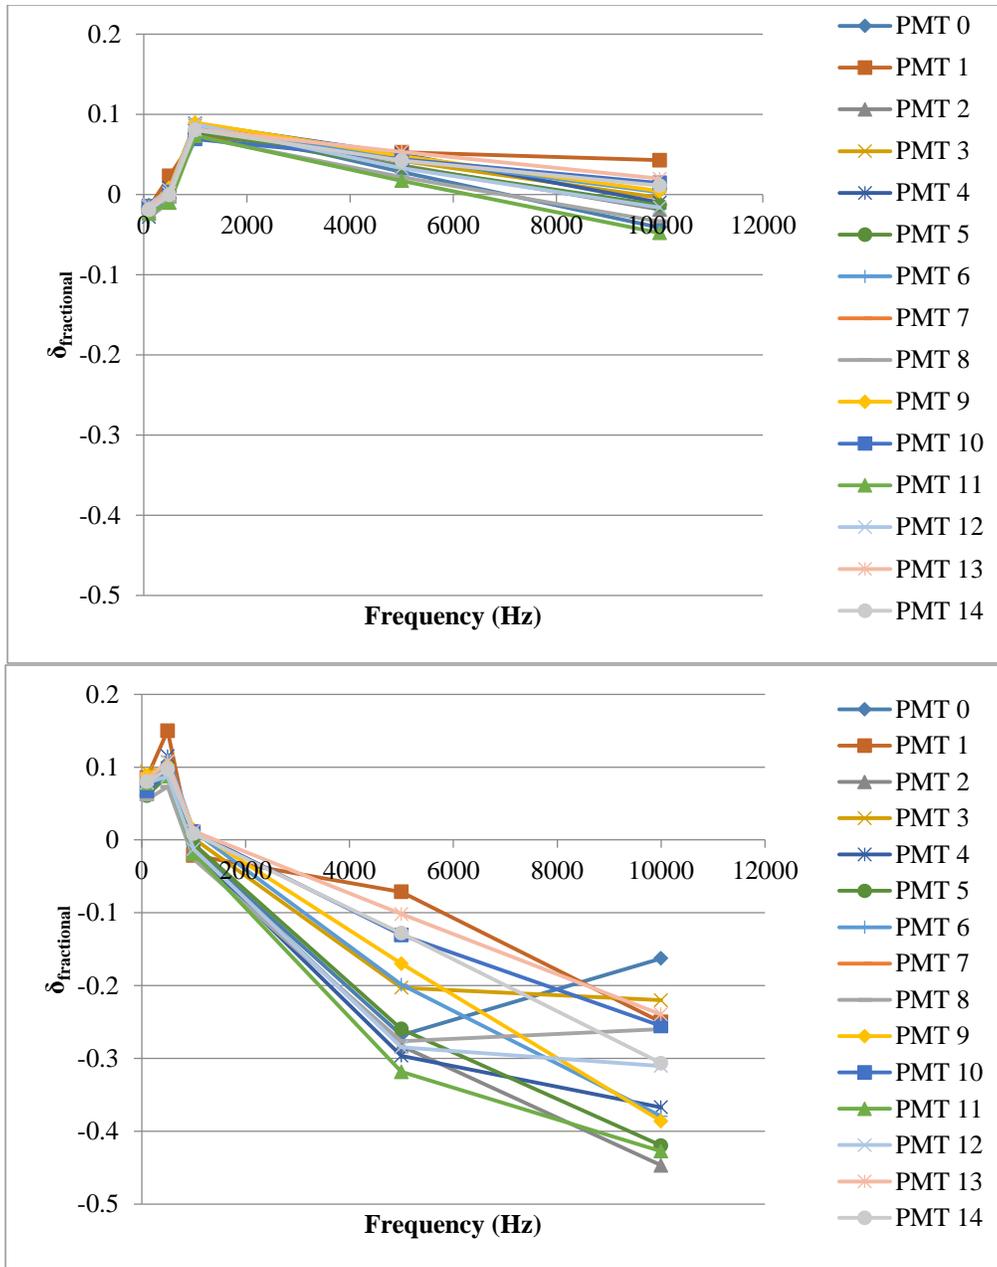


Figure 4.8: Top: stagger test with PMTs at 800V and 2.75V and a 1 μs window on the top LED at 100 Hz and 3.40V 2 μs window for the BLED with varying frequency with PMTs at 800V (top) and 1000V (bottom). The bottom LED corresponds to a pulse size of roughly 500k phe. The interesting thing to note here is the large negative fractional difference (enhancement) at increasing frequencies.

4.3.4 Large Pulse Test

At this point there was a good idea of what S2 pulse sizes can be expected for various sources at various gains. The goal of this round of testing was to see if there was any saturation present at the detector metrics that we had been using with source tests.

The methodology for this testing was much simpler than the previous ones: the voltage and timing of the LED pulses were tuned by matching the phe sizes of sources used already in PIXeY. That LED would then be pulsed at various frequencies to check for frequency-dependent saturation and see if it was likely that it was present in the data. The reference points used were for ^{83m}Kr -like pulses, ^{137}Cs -like pulses, and ^{228}Th -like pulses, which had phe sizes of approximately 50,000 phe, 500,000 phe, and 2,000,000 phe respectively.

As shown in Figure 4.9, for the LED tests meant to approximate a ^{83m}Kr pulse none of them seemed to saturate at any frequency. In fact, all PMTs observed at both 5 and 10 mV-ns/phe gain seemed to have a slight enhancement that was anywhere from 2-10%. While this meant that the PMTs do not saturate in this configuration it also means that they are nearing saturation. As such it meant that the PMT voltages can't really go any higher for a ^{83m}Kr study without expecting some level of saturation at 5 or 10 mV-ns/phe.

As shown in Figure 4.10, for the ^{137}Cs test the results showed that for some PMTs there can be saturation for a ^{137}Cs S2-like pulse, and some of them won't. For some PMTs an enhancement was observed that was as much as 15% and for other PMTs saturation was observed that was nearly 50%. When some PMTs saturate and others don't it's undesirable because it means that the energy resolution will most likely get worse. This is because, for the LED, it is at a fixed location. However for real ^{137}Cs events in the detector they will be in different locations, so the PMT response will greatly depend on where the event occurred in the detector. Since this can never be properly corrected this will almost always mean that the resolution will worsen.

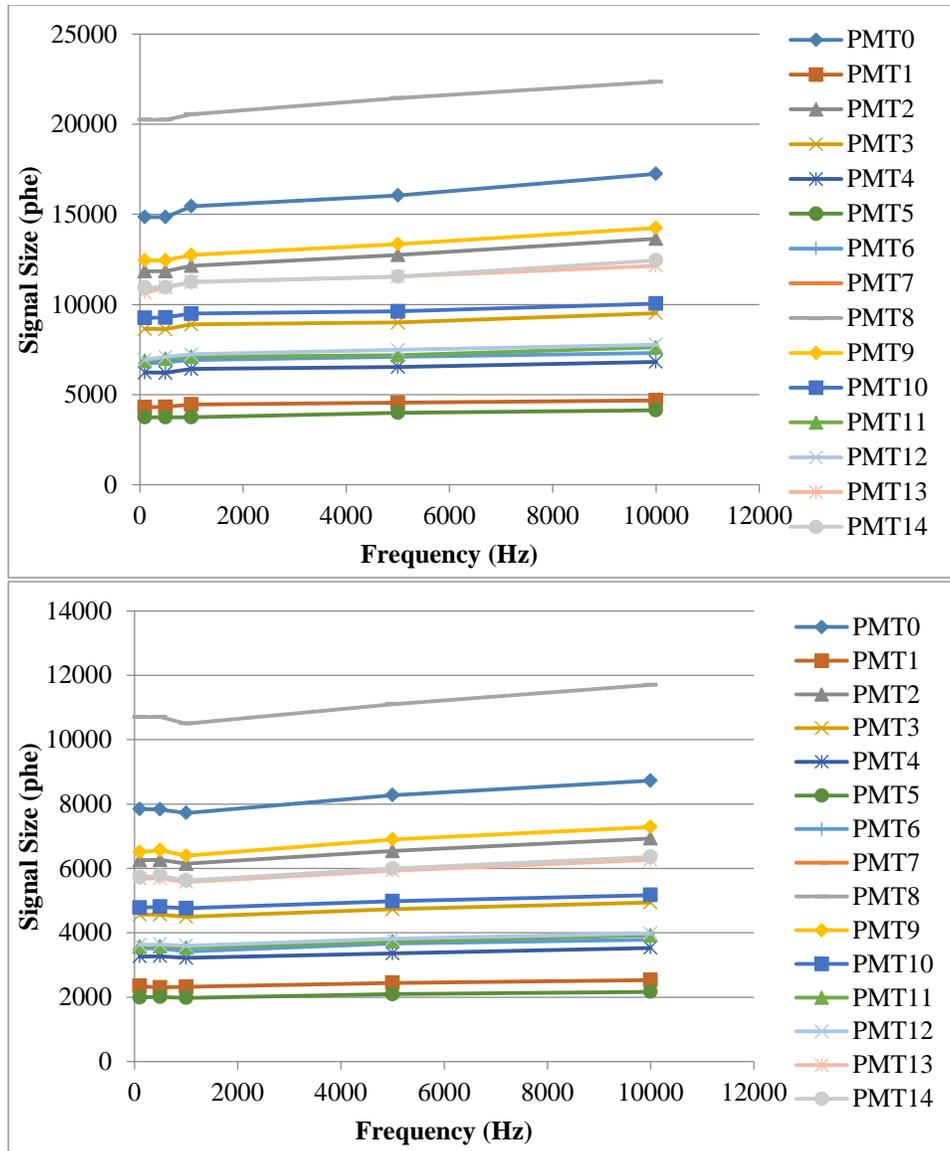


Figure 4.9: tuned LED that is meant to emulate the S2 size of a ^{83m}Kr -like pulse. Figure is pulse size in phe as a function of frequency. Top: PMT gains set to 10 mV-ns/phe Bottom: PMT gains set to 5 mV-ns/phe.

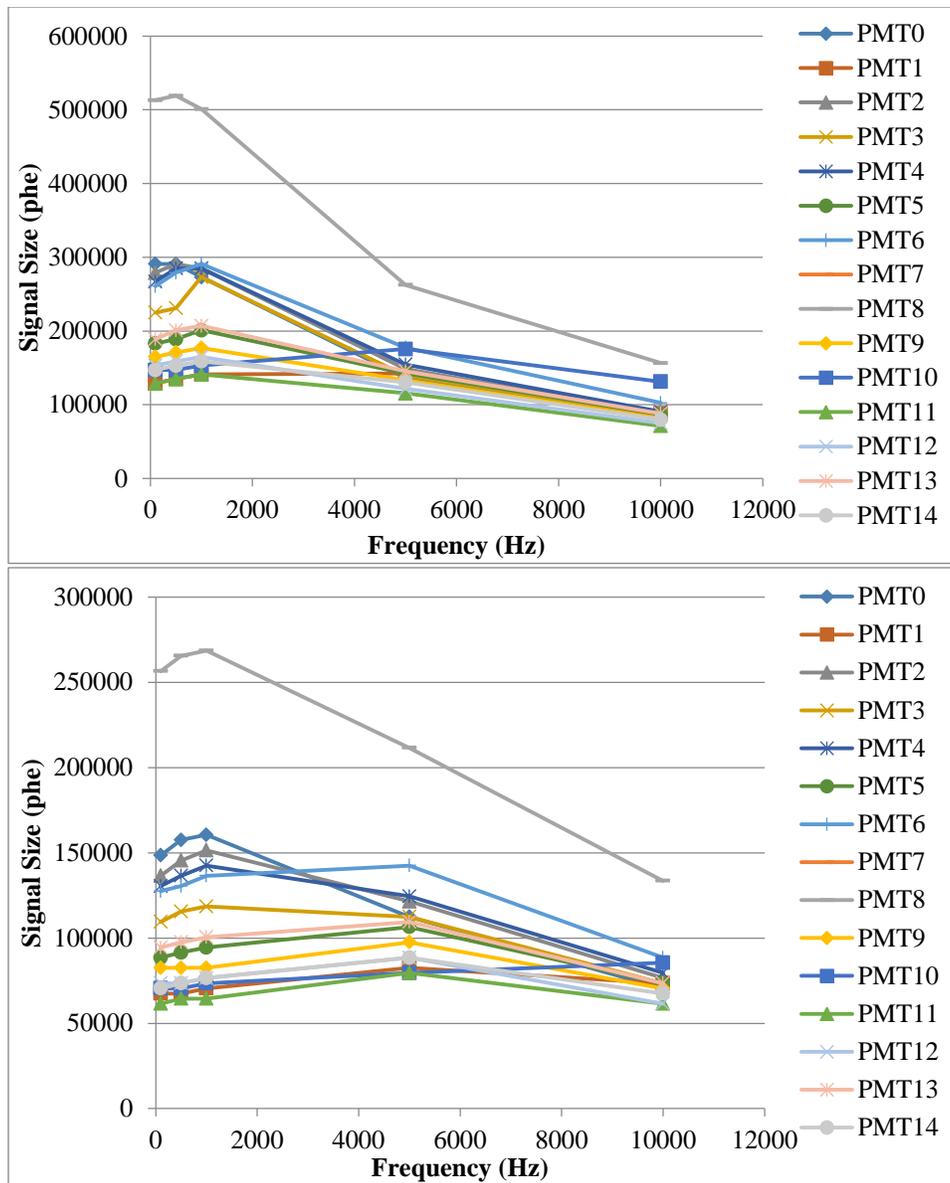


Figure 4.10: tuned LED that is meant to emulate the S2 size of a ^{137}Cs -like pulse. Figure is pulse size in phe as a function of frequency. Top: PMT gains set to 10 mV-ns/phe Bottom: PMT gains set to 5 mV-ns/phe.

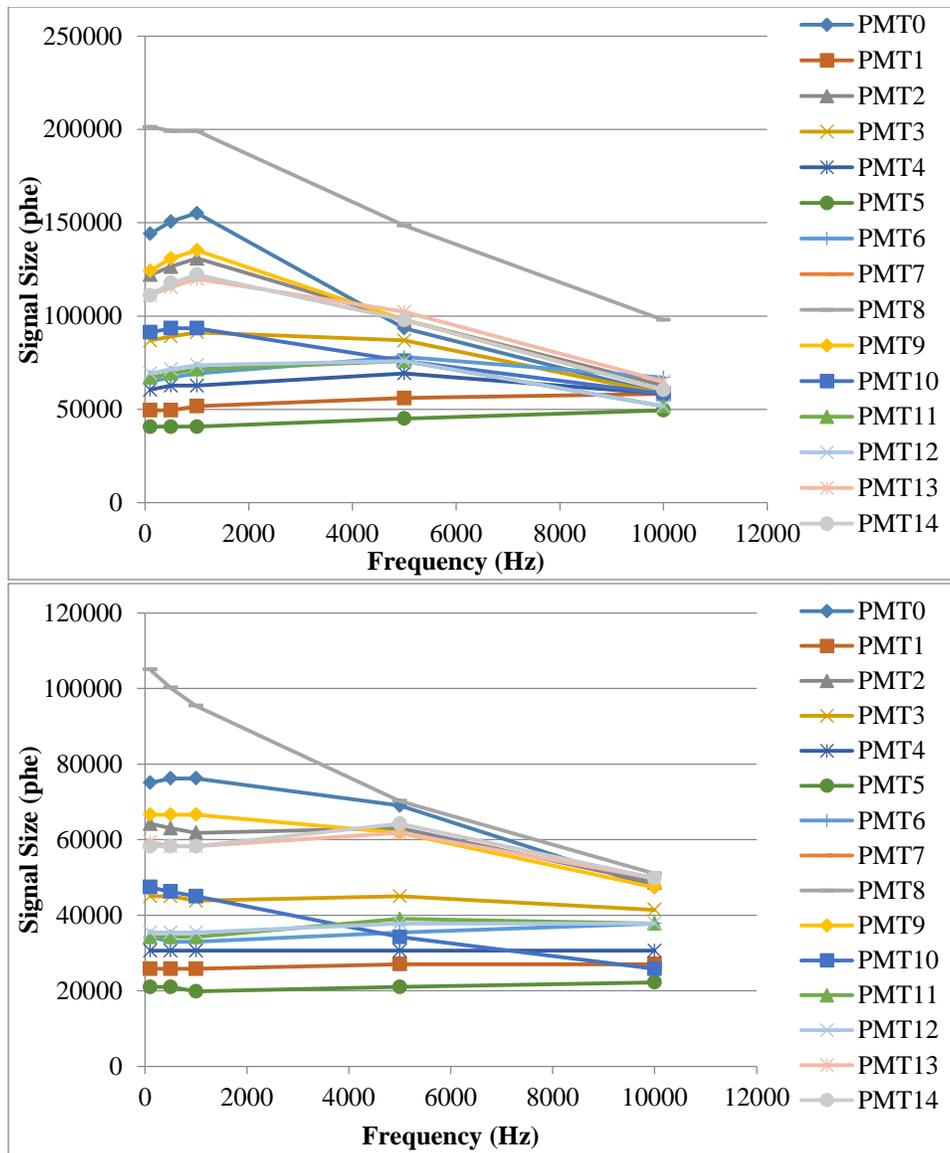


Figure 4.11: tuned LED that is meant to emulate the S2 size of a ~ 2 MeV like pulse. Figure is pulse size in phe as a function of frequency. Top: PMT gains set to 1 mV-ns/phe Bottom: PMT gains set to 0.5 mV-ns/phe.

As shown in Figure 4.11, for the 2 MeV-like pulse it was clear that a large portion of the PMTs had experienced 20-40% saturation, mainly the PMTs that were closest to the light source. Even with the PMT gains being turned down by a factor of 10 all PMTs with initial signal sizes of 50-100k phe had saturation starting at 2 kHz.

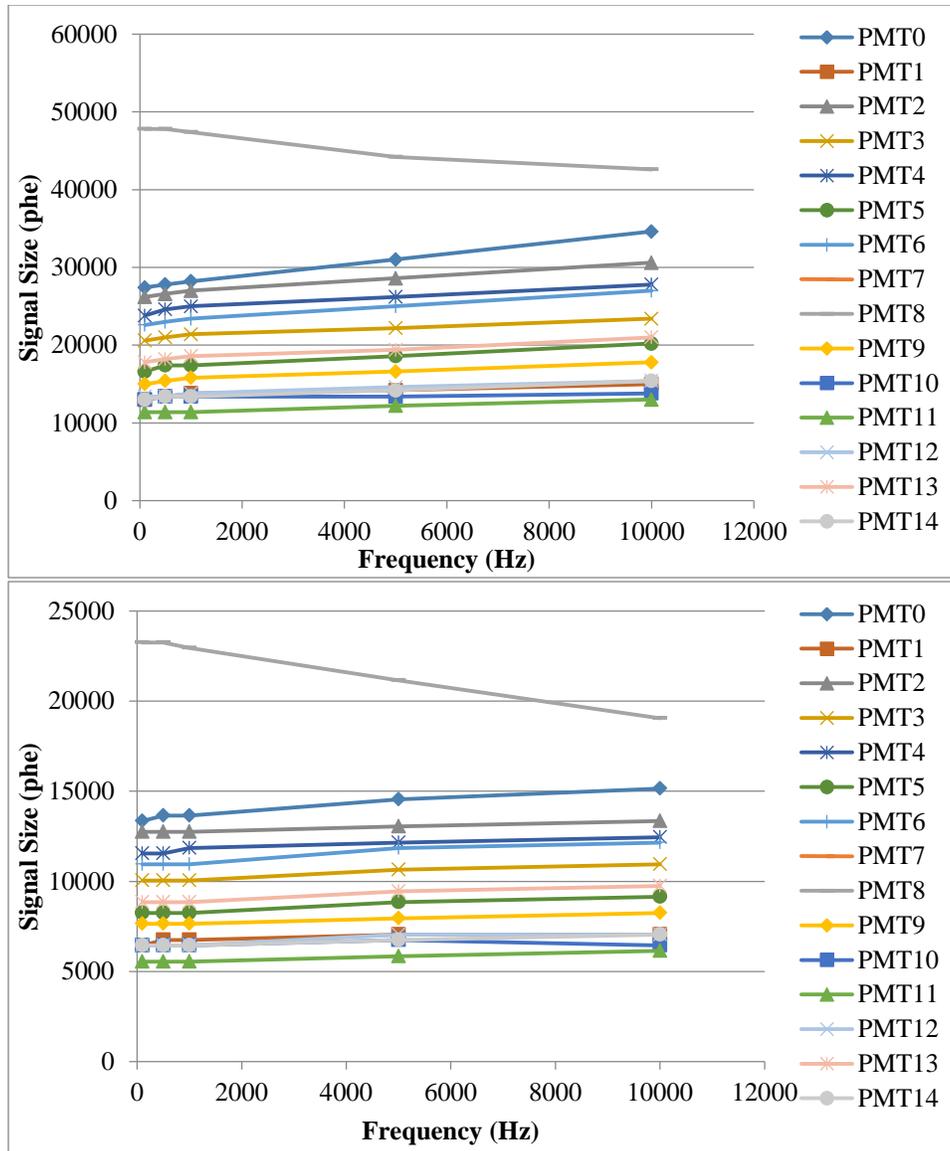


Figure 4.12: tuned LED that is meant to emulate the S2 size of a ^{137}Cs -like pulse. Figure is pulse size in phe as a function of frequency. Top: PMT gains set to 1 mV-ns/phe Bottom: PMT gains set to 0.5 mV-ns/phe.

As shown in Figure 4.12, the test was re-done for the ^{137}Cs S2-like pulse but using gains that were a factor of 5 smaller than the previous. The interesting and disappointing thing here was that almost every PMT seemed to be fine, observing some slight enhancement in the 5-10% range, except for PMT2. PMT2 was most likely the PMT that was nearest to the LED that was being used. This means that for a ^{137}Cs S2-

like pulse there could still be some saturation, even that these low PMT levels, because of the proximity to that one PMT. This eventually led to the realization that because of PIXeYs unique flat geometry the PMT response across different PMTs is highly stratified. This means that one PMT can see a large amount of light while another close by can see virtually none. This sort of effect is common amongst larger dark matter detectors because of their sheer size. However, in those detectors since they are so large the total amount of light is much more spread out among the PMTs that do see a signal.

The stratification of the light was difficult to solve because of the dynamic range of the S1 and S2 signals, if the gains were high enough to see the S1 then the S2s would almost certainly saturate. If the gains were low enough such that the S2s were not saturating then the S1 would either not be resolved or would have such SNR issues that the resolution would be extremely poor. For instance in Figure 4.12 an S1 in that configuration would not be resolvable.

Previous to this discovery best combined energy resolution obtainable in the system was around 2.3%, which was the limit between the S1s and S2s dynamic. Any attempt to adjust any parameter, be it the anode voltage, cathode voltage, or PMT voltage would cause the energy resolution to get worse. It was clear then that the first round of data that was collected with ^{137}Cs , although promising, had PMT saturation which limited the overall energy resolution. Unfortunately errors of this type may not be obvious in the S2 energy resolution signal, but in the combined energy signal, which will be explained in more detail in Section 6.6.

In addition even gains at the very lowest end of our calibration which was 0.5 mV-ns/phe seemed to be saturating at the ~2 MeV like pulse. For gains this low the PMT voltages were between 700-900V, which is at the limit of their operation. This hindered the study of ^{228}Th which has an energy of 2.6 MeV. This dynamic range problem described was only solved with the addition of an amplifier to the signal

chain³. This allowed the high gain that was needed without the saturation that limits the dynamic range of the S1 and S2 signals.

4.4 Measurement of the liquid level

During operation it is desirable to have the liquid level be as flat as possible in order to normalize the S2 signal response. Although the greatest attempt is made to keep the field cage area as level as possible during installation by using a bubble level there still may be some level mismatch that will cause an overall tilt in the detector. Any sort of tilt will make the proportional scintillation region not uniform across the detector, which changes the S2 response. Although with precise determination of good position reconstruction it is possible to correct for a spatial anomaly it is nonetheless desirable to try to minimize such correction factors as possible.

Thus an algorithm was developed that can calculate the level of the detector using any type of dataset. Initially the algorithm was used in conjunction with a ⁵⁷Co source however it was discovered that any source, even a background source, could be used for a level measurement. The algorithm is based on the following: assume the electron drift speed is uniform throughout the gas. Then the gas gap is simply:

$$\text{Gas gap} = (\text{Drift time}) * (\text{Measured Width}) \quad (4.10)$$

where the measured width is measured as the FWHM of the S2 pulses, and the drift time is calculated from the following relation:

$$\text{Drift Time} = A * (\text{Electric Field in Gas})^B \quad (4.11)$$

³There were many datasets taken in PIXeY that in the end were not useful due to saturation. Looking back in hindsight it seems obvious what the problem was and how to fix it but at the time it was simply hard to see. We needed an outside tool (an amplifier) to solve the problem but sometimes people get stuck and think that whatever they have is good enough.

where A and B are parameters that relate to the type of material and the pressure of the xenon which have empirical values of $A = 0.856437545$ and $B = 0.91793104$ [64]. This presumes the electric field in the gas is known, which can be calculated from the following formula:

$$Electric\ Field\ in\ Gas = \frac{V_{anode} * \epsilon_{liquid}}{\epsilon_{gas} * d_{total} + (\epsilon_{liquid} - \epsilon_{gas}) * d_{gas}} \quad (4.12)$$

where V_{anode} is the voltage on the anode, ϵ_{gas} is the permittivity of xenon gas, ϵ_{liquid} is the permittivity of xenon liquid, d_{total} is the total gap between the gate and the anode, and d_{gas} is the gas gap. This equation is derived in Appendix A.2. It is an iterative calculation since it assumes some liquid level guess and the new liquid level is calculated. This process is repeated until the difference between the two values converges to a difference that is less than 1 μm , although the algorithm can be run to greater convergence as well. All that is needed externally is the voltage on the anode and the rest of the calculation is done based on the S2 PMT response and the xy position, which is calculated based on a weighted sum of the PMT signal on the top array and thus the level can be compared at different locations within the detector.

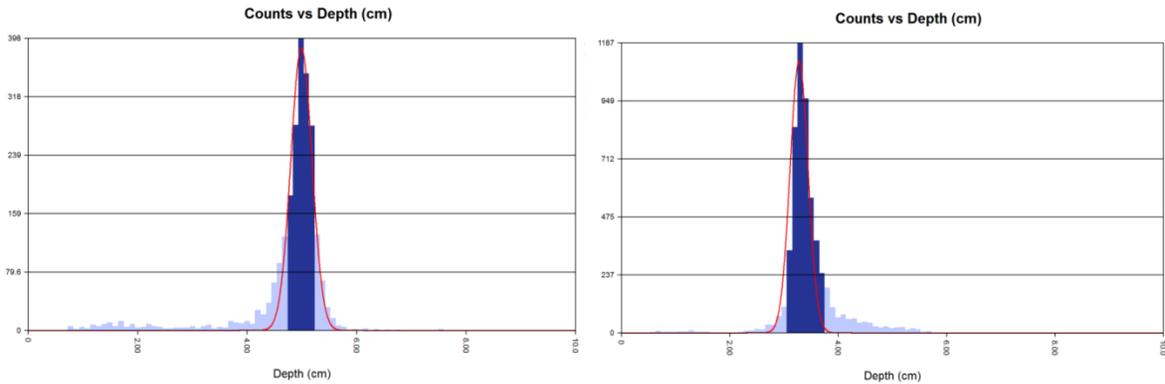


Figure 4.13: Comparison of the height spectrum during a filling operation. As more xenon was being added to the detector the height went from 5mm to 3.8mm, showing the liquid level change in real time.

The detector was leveled based on four leveling screws that were at the top, which was done first by the tilt sensors and then later much more accurately by the level determination algorithm. Since there were four leveling screws instead of three the process was iterative with some educated guessing. With

this algorithm we were able to get a deviation of less than 40 μm , which is about 1% of the total gas gap in the detector. These numbers were obtained from the very edges of the detector, in most data runs some kind of radial quality cut is used which lessens the overall impact of such a correction factor. This correction is a sub-percent correction factor so once energy resolutions go below 1% this sort of correction factor will become important and such a tilt adjustment would have to be better finely tuned.

4.5 Lifetime Measurements

Another important measurement to ensure healthy operation of the detector is to measure the electron lifetime in the liquid. The lifetime of the electron is how long a free electron lasts until being captured by an impurity. These electronegative impurities, like oxygen, are attracted to the free electrons and then once attached move so much slower and are not detected.

The traditional method for calculation of the lifetime is to generate a pixel plot of the S2/S1 ratio signal as a function of drift time. The ratio between the S1 and S2 signals will cancel out any energy drift time dependence, so no quality cuts are needed. However if there is any sort of geometric response in the S1 or S2 signals as a function of depth (and they generally won't be the same) then this will cause error with this measurement since the lifetime will be convoluted with the geometric response of the S2 and S1 signals. In general it was found that the S1 response has a linear response as a function of depth, which means that to maximize the usage of this technique it should be used *after* geometric corrections have been made.

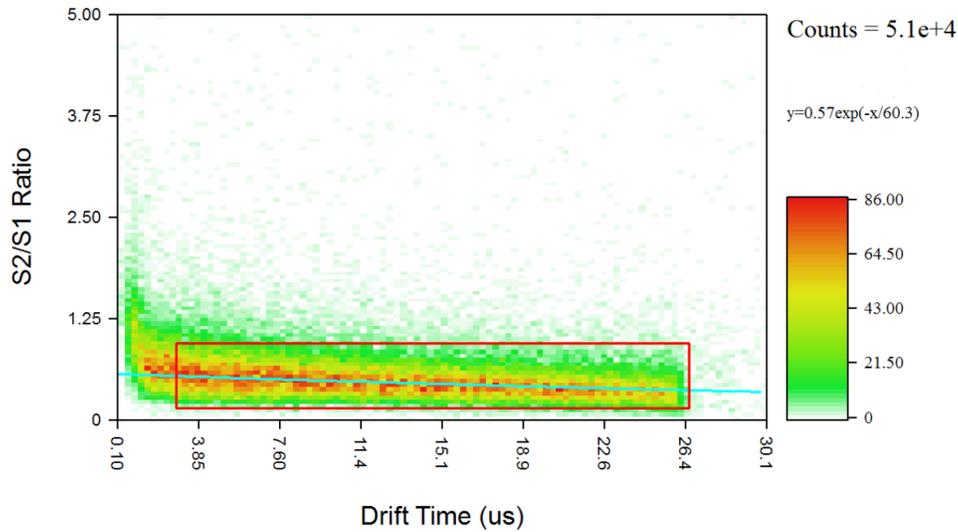


Figure 4.14: Determination of the lifetime by comparing the S2/S1 ratio with the drift time of the events. This figure was generated using ^{137}Cs at 2,000 V/cm drift field. The value for this dataset was about 60 us.

An alternative method is to find the S2 peak location at discrete drift time steps and plot the overall data, as shown in Figure 4.15. For instance a source that is placed externally to the detector can be used, and the peak response at different drift time windows can be found. Since a peak is being measured and not just all spectra this method does not suffer from any peak energy dependence. It also is less sensitive to geometric fluctuations since only the S2 is used to calculate the lifetime instead of the S1 and S2 signals. However, it is typically more difficult to do since it requires finding an energy peak at many discrete locations.

This is why it is important to have a continuous circulation and purification system installed in a LXeTPC. Such a system will be described in greater detail in an upcoming hardware paper, but suffice to say the design goal of such a circulation and purification system was to remove electronegative impurities in the range of a few parts per billion of impurities. Interestingly enough, this level is too low to measure commercially, so the electron lifetime becomes in a way the best way to measure the impurities in the system.

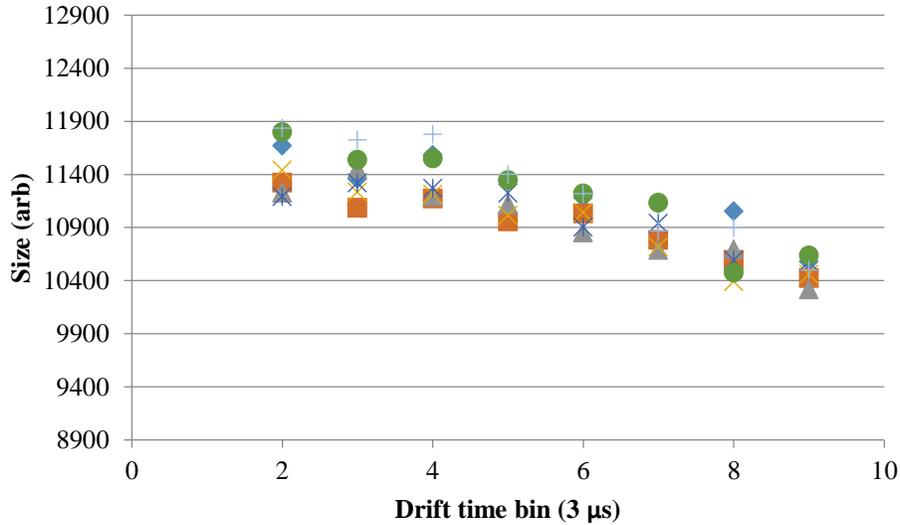


Figure 4.15: S2 size data in normalized counts per bin. The cut was for $r < 1.75$ cm, each theta bin is separate, and it was a Cs137 source with 1,000 V/cm drift field. The loss in peak size is due to the electron lifetime; in this figure the lifetime was calculated to be ~ 80 μ s.

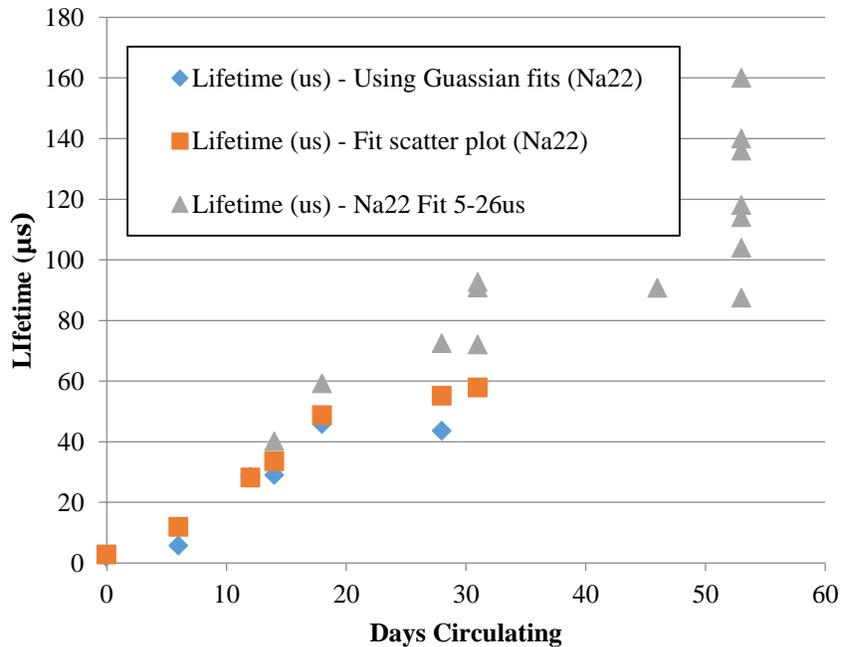


Figure 4.16: Measurements of the lifetime as a function of circulation time for a variety of different lifetime calculation algorithms. The values increase with time which implies that impurities are being removed, although the variance increases as well. The graph stops at ranges at around 100-160 μ s, which is where the values leveled off at 123 ± 23 μ s at 53 days.

The values of the lifetime calculated after over 60 days of circulation ended up being in the range of 100-200 us, which corresponds to an impurity level of 8.5-4.3 parts per billion [65]. This ends up being over 3x the drift length in the PIXeY for field configurations between 1-2 kV/cm.

4.6 S1 vs. Depth Dependence

As mentioned in the previous section the amount of S1 light captured is known to vary with depth. Physically this is a geometric effect and it has to do with the reflection of the S1 light within the detector. This can be corrected by using a similar method employed in the previous section but comparing the S1 signal size instead of the S2 signal size.

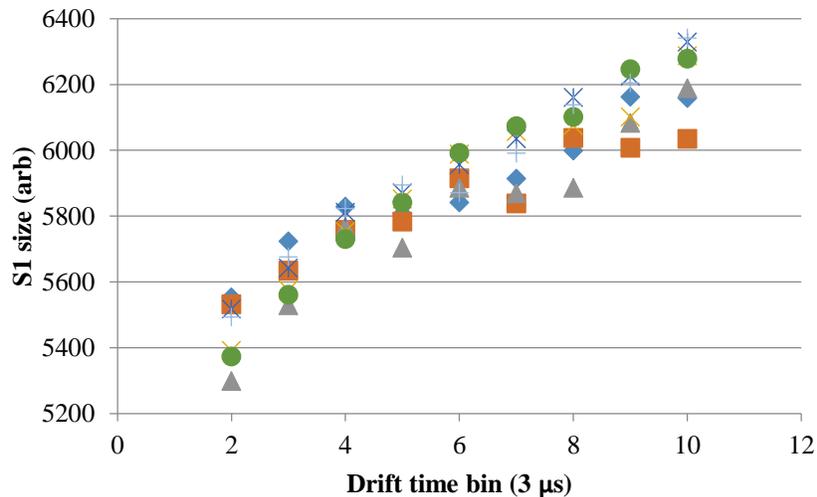


Figure 4.17: S1 signal size as a function of Drift time in normalized counts/bin. The increase in signal size is a geometric effect that is most likely due to the proximity of the PMTs. Data shown with a ^{137}Cs source with 250 V/cm drift field.

The results show that the S1 increases size as a function of depth because events that are closer to the bottom of the detector are closer to the bottom PMT array and thus see proportionately more light.

Since most of the light is internally reflected at the liquid surface this means that events higher up in the detector scatter more than events lower in the detector. This gives the light more chance to be absorbed by any other surface. The data is fitted and then corrected with linear or quadratic fits, with typical values in the range of 15-31 keV/ μs and 3-0.5 keV/ μs^2 .

4.7 Quality Factors

There were a number of quality measures that were experimented with to try and determine if energy resolution could be improved. And while many of these were not used in energy resolution measurements they lead to a discovery that was worth the time invested. They all attempt to measure the quality in which the data analysis algorithm is finding and evaluating data and all of these factors have been used in previous experiments [11].

4.7.1 Q-Factor

The Q-Factor or Quality Factor is defined as the total area of all identified pulse regions divided by the total area of the trace, or simply:

$$Q_{factor} = \frac{Area_{pulseregions}}{Area_{trace}} \quad (4.13)$$

For an ideal pulse finding algorithm, this factor should be 1 since it should find all the regions that are above the baseline. It's important to note that there was no zero suppression done in any dataset so this factor also measures the baseline algorithm performance, or more technically the convolution of the baseline algorithm and the pulse finder algorithm.

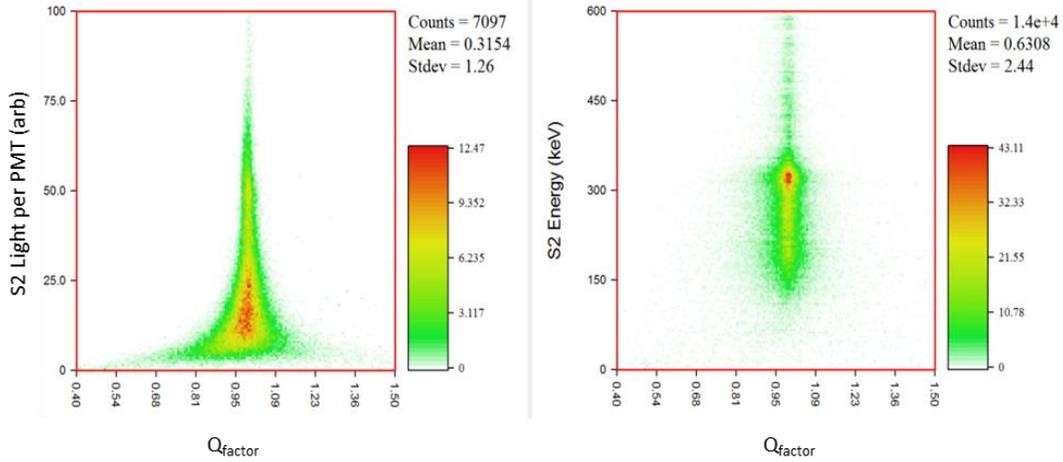


Figure 4.18: Left: pixel plot of individual pulse areas as a function of Q-factor. At higher pulse areas the Q-factor is between 95-100% but is anywhere between 54-100% at pulse areas less than 10 keV. Right: total pulse area as a function of Q-factor. Regardless of total area most Q-factors are between 95-100%. Some Q-factors are more than 100% which is due to baseline shifting.

The results seemed to reveal more about the baseline calculation than the pulse finding calculation. In particular, sometimes the baseline can shift as explained in Section 3.2.1 which makes the Q-factor be more than 100%. By cutting these events however it didn't seem to have any noticeable improvement on energy resolution. The most probable explanation was that there were very few events with poor Q-factor, and after position quality cuts the amount that they could broaden the energy resolution was negligible.

4.7.2 R-Factor

The R-Factor was named as such because R comes after Q. The R-Factor, or “rality factor”, is a measure of the performance of the S1 and S2 identifier. It measures the total area identified as being S1 or S2 and divides that by the total area found in all pulse regions. More specifically:

$$R_{factor} = \frac{Area_{S1S2}}{Area_{pulseregion}} \quad (4.14)$$

Once again this factor should be 1 (or very close to 1) when all of the S1s and S2s have been properly identified.

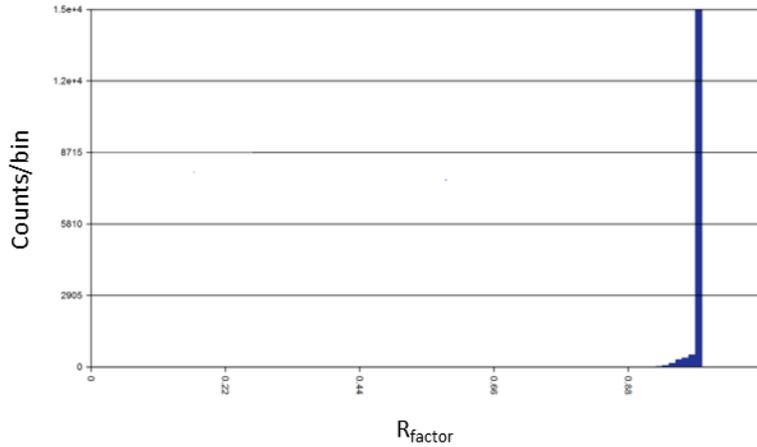


Figure 4.19: Measurement of the R-factor for a Cs137 spectrum. 99.6% of the data is in the 95-100% data range.

In practice since the S2 areas are so much bigger than the S1 areas the R-factor only measures how well S2s were being found. And since they are at least 50 times bigger than S1s this quality factor contains no real useful information. As expected almost all events reported 1 for the R-Factor and a negligibly small amount reported non-unity values of the R-factor, which were even smaller subset of events that was cut compared the Q-factor and thus was not necessary.

It should be noted that in previous work [11] both the Q and R factors were merged into a single quality measure. The goal of these studies was to determine which element of these factors was useful, if any. As it turns out any combination of these two factors will ultimately measure the baseline and pulse finding algorithm, because of how close to one the R-factor is compared to the Q-factor.

4.7.3 TB-Factor

We conclude the study of factors with one that is not a measure of a pulse algorithm. The TB-Factor stands for “top bottom” factor, and it’s a measure of the ratio of the S2 light seen in the top array versus the total amount of S2 light. This was used in previous experiments as a way to filter “spurious” S2 events for Dark Matter searches [11]. Here the investigation was to see if there was any improvement in energy resolution.

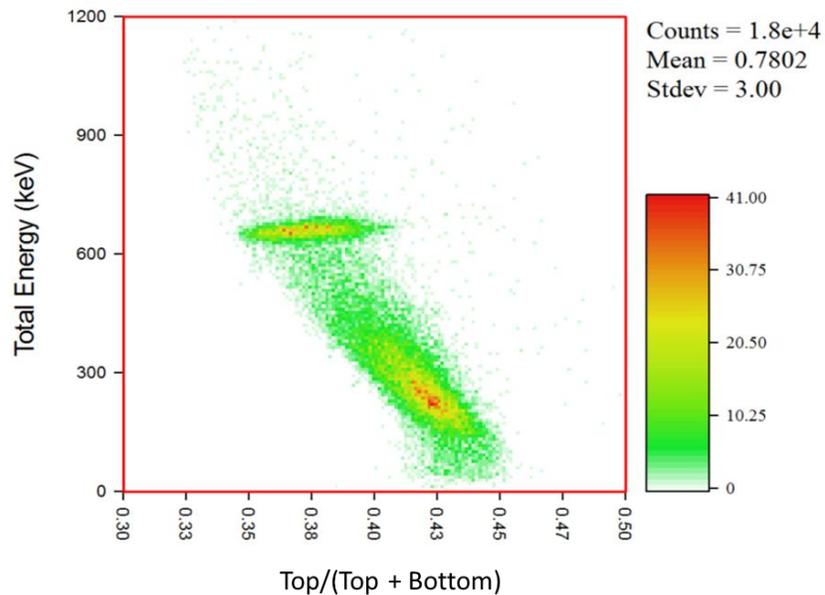


Figure 4.20: Total light energy as a function of the Top/(Top + Bottom) ratio. Note the correlation between energy and top/bottom ratio, with the peak positon moving by about 3%.

Once again it was found that such a correction factor did not noticeably improve the energy resolution. What was instead found was a strange correlation between the TB-factor and the total amount of light in the event. Note that such a correlation probably makes no physical sense: the ratio of light from the top array should be a geometric effect; it should depend on how the detector is constructed. The total amount of light should not really change how that area is collected, especially since S2 sizes are in the range of 100-1,000k phe, which is about 3-30k electrons based on the calculations in Section 6.5.1. These amounts are dominated by Poisson statistics and are too high to be susceptible to shot noise.

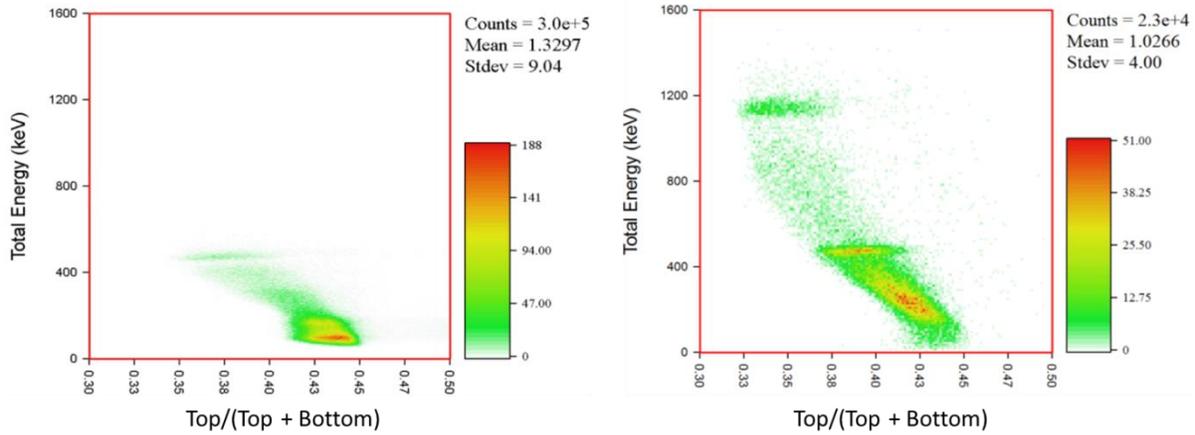


Figure 4.21: Comparison of TB-Factors as a function of energy for different isotopes and gains. Left: ^{228}Th data at 1 mV-ns/phe gain. Right: ^{22}Na data at 5 mV-ns/phe gain. Note the trends seem similar even though the energy content is different.

It was then realized that this may not be a physical effect but rather a manifestation of self-saturation. If there was no saturation, then a pixel plot of the total light collected as a function of the TB-factor should be constant, it should look like a pillar. Now suppose that one and only one of the PMTs is saturating. What this means is that fractionally less light will appear in that PMT. While this doesn't change the total amount of light by much it does change the ratio of light, since the ratio will be more sensitive than the total. This means that instead of being a pillar, the spectrum should shift over to the left as the total amount increases.

Tests were done in order to verify this, which was to look at different isotopes at different gains. For this ^{228}Th was chosen at 1 mV-ns/phe gain and ^{22}Na was chosen at 5 mV-ns/phe gain. The idea is that the 511 keV line in ^{22}Na is roughly 5 times as small as the 2.6 MeV line in ^{228}Th . So, if it was an energy/geometric effect then the difference in gains shouldn't matter and the scaling should be identical however it was a gain effect then the scaling would be different.

As shown in Figure 4.21, what was found was that the effect scaled just as one would expect if this was a PMT effect, not an energy one. This developed into an in-situ method of determining whether or

not saturation was present in data. This was more useful than the LED tests, since it can connect with actual data in a running detector that doesn't have to be abstracted by considering the data rates of background or other considerations.

4.8 Saturation Plots

4.8.1 Introduction

A "saturation plot" is a pixel mapping of the area of the S2-like pulse compared to the ratio of the top to total signal size, originally inspired by the TB-factor studies in Section 4.7.3. If there is no saturation then the pixel plot should appear as a column: there should be no energy dependence in the ratio between the top and bottom signals. However if a PMT channel is saturating at a certain size then this will show up as a decrease in the ratio of the top signal. If a PMT is saturating then its area response goes down, and if all of the other areas are constant then that leads to an overall decrease in the ratio. This effect is shown in Figure 4.21.

4.8.2 Saturation per PMT

In addition to identifying saturation between the top and bottom arrays it was later found that saturation plots could be made of individual PMTs to check and see which PMT was saturating. This was also a good diagnostic measure to figure out how exactly saturation was happening, and which PMT was doing it. Mathematically speaking the ratio of the amount of light seen in one PMT with the total amount of light seen in all channels is computed, that is:

$$\delta_{PMT1} = \frac{Signal_{PMT1}}{Signal_{PMT1} + Signal_{PMT2} + \dots + Signal_{PMT14}} \quad (4.15)$$

Equation (4.15) is for δ_{PMT1} ; although similar equations can be constructed for the other PMTs. This should be able to calculate the amount of saturation in each individual PMT, provided that the total amount of saturation is small.

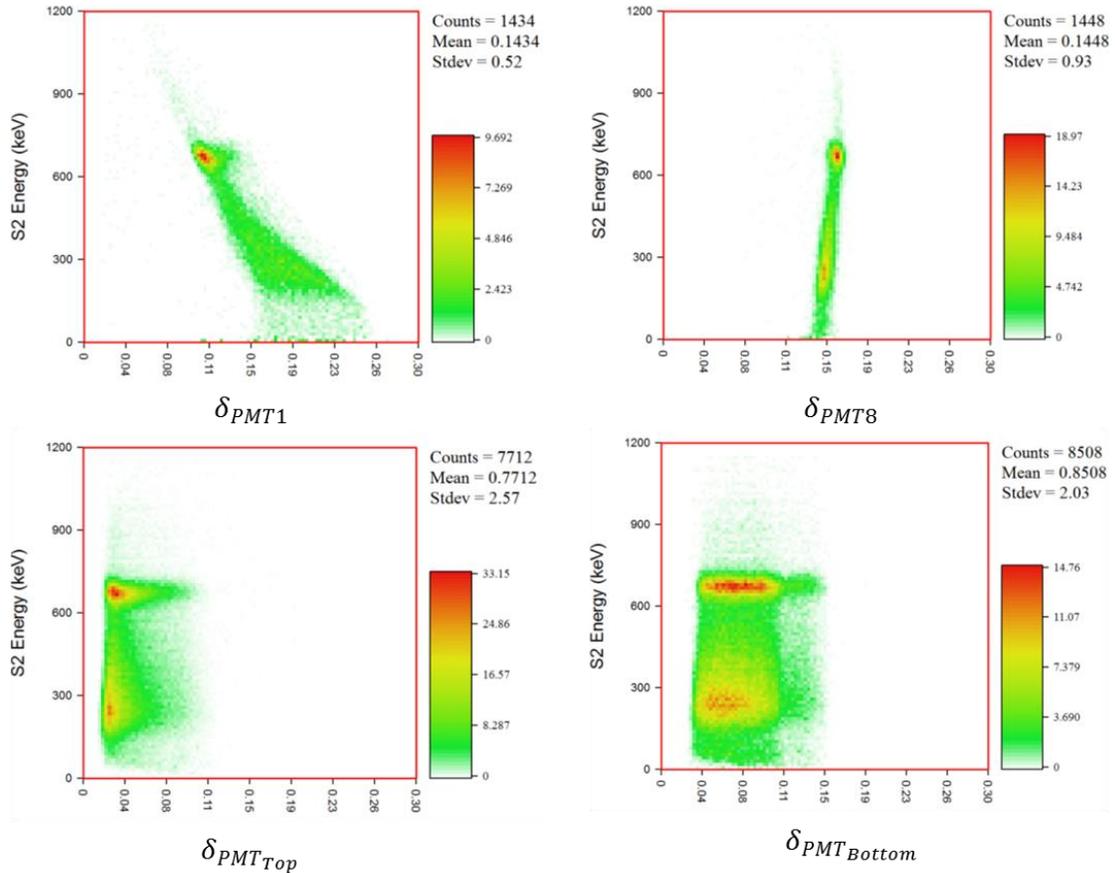


Figure 4.22: Comparison of the various pixel plots of total area as a function of the percentage of light seen in that PMT for a ^{137}Cs source at 1,000 V/cm drift field and 10 mV-ns/phe gain. Top Left: top center PMT Top Right: Bottom center PMT Bottom Left: Top off-center PMTs. Bottom Right: Bottom off-center PMTs.

What was found was that the top center PMT was the most likely culprit for saturation. The error had been that the PMT gains were calibrated such that there shouldn't have been any saturation on average. However, as shown in Figure 4.22 the top center PMT sees between 15-20% of all the light

when it is unsaturated, which is more than twice the amount of light that one would expect if all PMTs saw the same amount of light, which would be about 7%.

In addition to the total amount of light getting smaller the deviation of light gets smaller as well, which paradoxically means that saturation can improve energy resolution. This added another confounding variable to the historical problem because it meant that fiducializing down would improve more than it should, since only one PMT would be saturating. Ultimately what would happen is that the best energy resolution would go down to the deviation in the saturated PMT since it was a systematic effect. This value was still too high to be considered a good energy resolution value, which also was puzzling for energy resolution measurements.

This test was more successful compared to the LED testing because of the relative static nature of the LED. Each LED was fixed and mounted to one spot in the detector. This means that the light collection out of each PMT is exactly the same for each pulse, which makes detecting saturation difficult since any one PMT is seeing exactly the same amount of light and thus the reduced amount may just be the normal operating mode of the PMT. But when the event happens in different locations throughout the detector and at different energies it's much easier to cross-reference all of the events, and check if the response of one PMT is aberrant or not.

4.8.3 Determining Optimal PMT gains

Next the optimal PMT gains for each source run were found, which is necessary due to the dynamic range difference in size between the S1 and S2 pulses. If the gain is too high then the S2 signals will saturate and degrade energy resolution performance, if the gain is too low then the S1 energy resolution will be poor due to small pulse size. For each source run a saturation plot was made to make sure that the large S2-like pulses were not saturating.

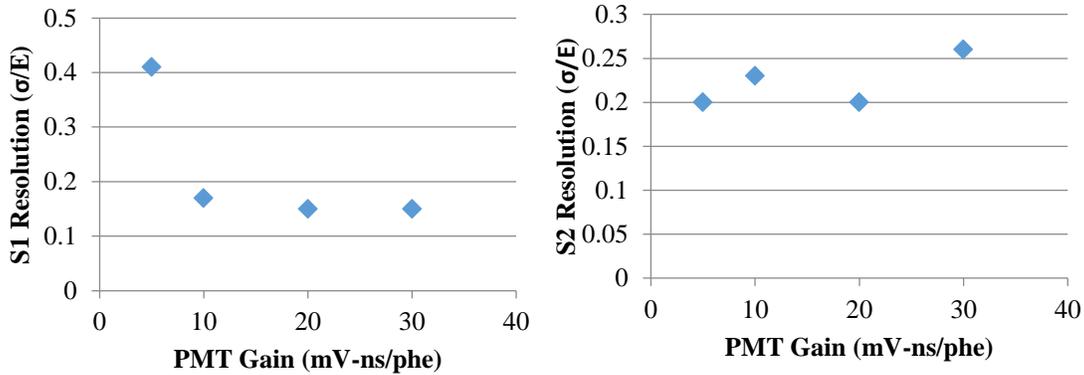


Figure 4.23: Comparison of energy resolution values as a function of PMT gain for ^{83m}Kr source. The S2 gains (right) remain relatively flat with a variation of about 3% but the S1 gains (left) seem to improve with higher gains, most likely due to signal-to-noise reasons. This sort of evidence set the trend for wanting to have as high PMT gains as possible due to desire for optimizing the S1 energy resolution, which propagates to the combined energy resolution.

With the addition of the amplifier and the saturation plots the optimal PMT gains were able to be determined. This was done by increasing the PMT gain until saturation was observed, then decreasing the PMT gain so that there was no saturation but the S1s could still clearly be resolved. Based on this method the gains used for the sources were found and the results are in Table 4.1.

Source	Gain(mV-ns/phe)
^{83m}Kr	16
^{57}Co	16
^{133}Ba	16
^{22}Na	4
^{137}Cs	10,5(center)
^{228}Th	2

Table 4.1: PMT gains used for all optimization source runs.

For the lower energies (^{83m}Kr , ^{57}Co , and ^{133}Ba) the gain was found by making sure that the ^{133}Ba was not saturated while at the same time making sure that the ^{83m}Kr S1s can be resolved without having SNR issues. The ^{22}Na was tuned for the 1274 keV line rather than the 511 keV. This may explain why the 511 keV resolution is not as good as current work with other detectors [66]. This will be explained in

further detail in Section 6.4. For the ^{137}Cs it was attempted to try and turn down the gain in the center tube because this usually saturates due to the light collection profiles, which is shown in detail in Section 5.2.2. Finally, the ^{228}Th gain was chosen based on half of the ^{22}Na gain, although some saturation was observed in the ^{228}Th spectra, especially at higher field values, which is shown in Figure 6.14.

4.9 Tuning Gas Field Region

4.9.1 Introduction

Of the two field regions that are in the detector, the region that controls the amount of proportional scintillation needs care in tuning. Without any field the charge signal could not be read out, and if the field was too large then the signal could easily saturate the PMTs or even cause railing in the DAQ. It is for these reasons that precise tuning of the anode voltage that determines the electric field in the gas must be done in such a way as to optimize the S2 signal without saturation. Once the anode voltage has been tuned properly the energy resolution is not expected to change due to the large electron statistics in the S2 pulse region [2]. In this section the methodology that was used in determining the level of anode voltage that was used for the energy resolution studies was studied.

4.9.2 Anode Studies with $^{83\text{m}}\text{Kr}$

There were many different experiments that were done to test the viable ranges of the drift and proportional scintillation regions. They involved searches as a function of various quality cuts and cross-referencing voltages on the Cathode and Anode. One experiment was carried out by using injected $^{83\text{m}}\text{Kr}$ into the liquid xenon. This source is different in that unlike the rest of the sources, which are sealed external sources, $^{83\text{m}}\text{Kr}$ gets injected into the circulation system where it spreads through the detector. The active isotope has a half-life of 1.83 hours [67], so it is completely removed after several hours after injection.

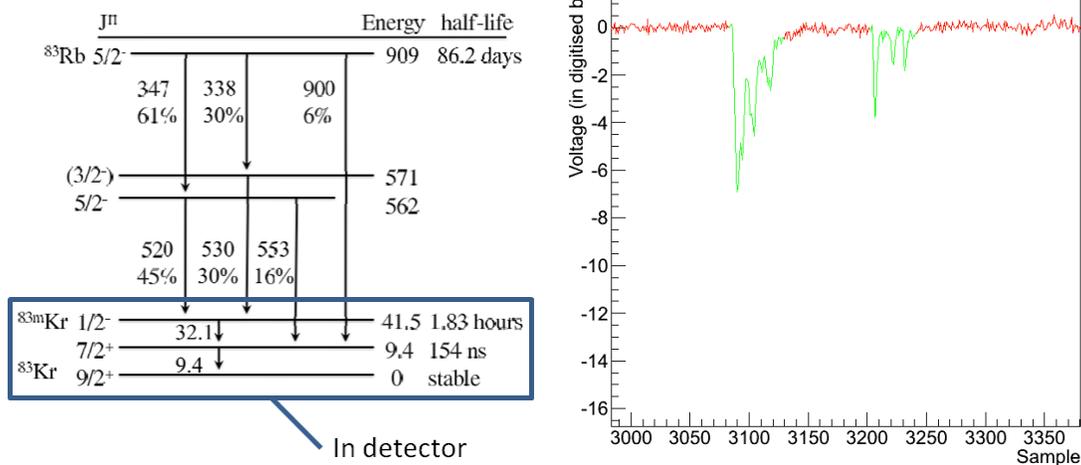


Figure 4.24: Left: Energy level diagram which shows the production of the $^{83\text{m}}\text{Kr}$ values. The decays have a 152 ns lifetime between them which is readable by the 4 ns PIXeY DAQ system. Right: observation of the double S1 pulse signals in PIXeY data.

$^{83\text{m}}\text{Kr}$ is used over other sources for detector calibrations because it tends to spread itself uniformly throughout the detector. This means that detector effects such as a position dependence and field dependence can be studied with the assumption that the source is uniformly distributed. Based on experiments the $^{83\text{m}}\text{Kr}$ signal would appear around 5 min. after injection and then 10-20 minutes later it would be completely distributed throughout the detector. In addition because $^{83\text{m}}\text{Kr}$ has two separate beta decays it has a characteristic “double S1” signal that can be used to eliminate virtually all backgrounds in the energy region.

For this study there were 4 values of the Cathode studied and 4 values of the anode studied, which were 1 kV, 2 kV, 3 kV, 4 kV for the cathode and 4 kV, 6 kV, 7 kV, 8 kV for the anode. This corresponds to drift fields in the liquid of 200 V/cm, 400 V/cm, 600 V/cm, and 800 V/cm, and 6.4 kV/cm, 9.6 kV/cm, 11.2 kV/cm, and 12.8 kV/cm in the gas. Thus this was a 4x4 parameter run space so 16 runs were done to test not only the ideal gas region field but to also see if there were any cross-correlations between the drift and proportional scintillation fields with respect to S2 energy resolution.

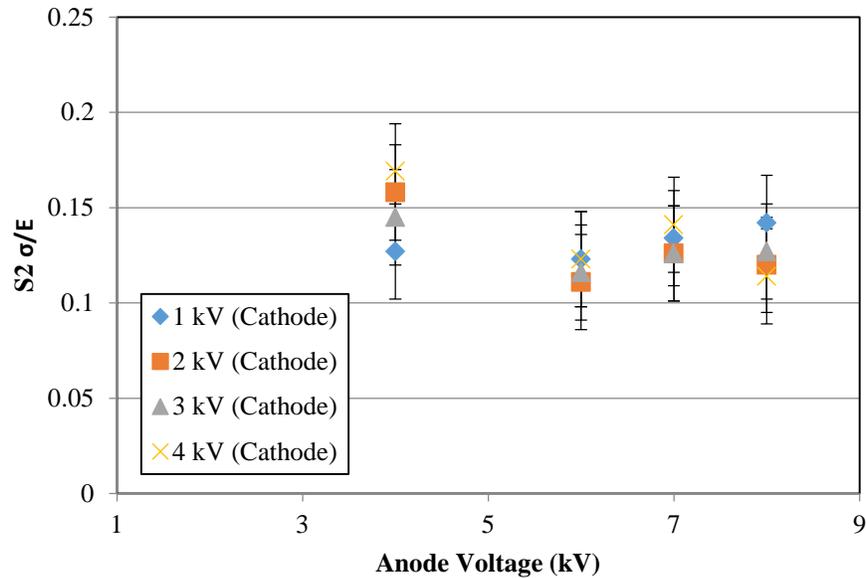


Figure 4.25: An experiment using ^{83m}Kr on how to set the optimal field configurations. Multiple values of the anode and cathode were tested. Taking into account the uncertainties in the measurements and assuming the resolution is constant the energy resolution values in S2 from 6-8 kV was $0.138 \pm 0.009 \sigma/E$, which is a 6% variation of the local value.

The results of this test were that, combined with the uncertainties of calculating the energy resolution values, that the energy resolution in the data points between 6 and 8 kV on the anode were flat with an energy resolution of $0.138 \pm 0.009 \sigma/E$, which is a 6% variation. There is some improvement going from 4 kV and higher, but there was no further improvement observed. This was expected, since the fluctuations due the gas field is supposed to be dominated by electron extraction efficiency and once the statistics are high enough there shouldn't be much difference in the energy resolution due to the anode.

In addition, there was no observed cross-correlation between the cathode and anode values. Although energy resolution is expected to improve due to increased drift field this was not observed during this test run, which may have been due to the fact that there weren't many quality cuts done to

these datasets other than basic edge fiducialization. Still, these results suggest that there is no reason to go higher in anode voltage than 6 kV in terms of energy resolution optimization.

4.9.3 Anode studies with ^{137}Cs

The last test that was done was to study electron extraction efficiency, so that the gas field can be maximally efficient across many different source runs. Since there are many parameters that go into the running of the detector it was desired that the anode voltage should constant across all resolution studies. If the anode voltage becomes a parameter that is continually tweaked and studied during energy resolution testing then there would be too much data to study. So if an ideal value could be used it would be one more common parameter and one less dimension in cross-comparisons.

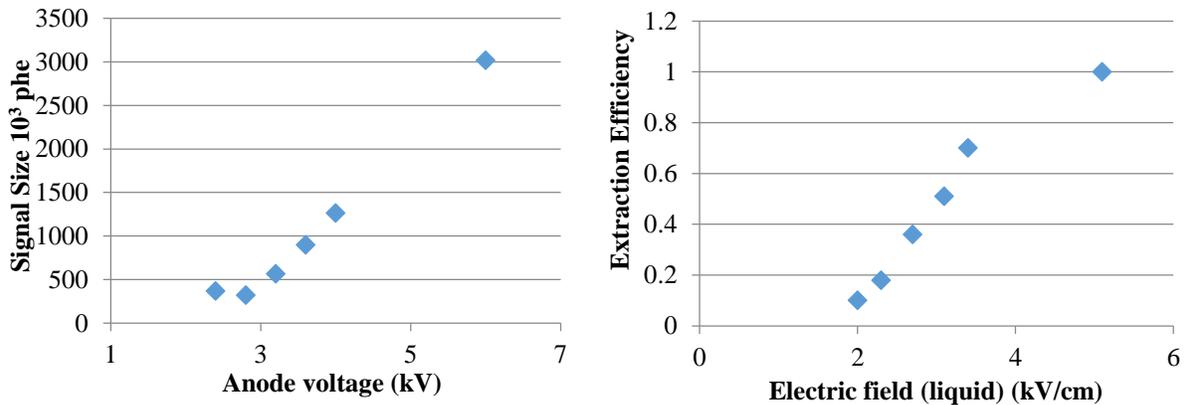


Figure 4.26: Left: Size of a ^{137}Cs signal in 10^3 phe as a function of anode voltage for 250 V/cm on the cathode. Using this and assuming 100% extraction efficiency at 6 kV a crude extraction efficiency plot can be calculated. Right: Extraction efficiency as a function of extraction field, which is calculated from Appendix A.2.

For this experiment a ^{137}Cs sealed source was used and placed externally to the detector. Several values of the anode voltage were used and the drift field was kept constant at 250 V/cm. The values used for the anode voltage were 2.4 kV, 2.8 kV, 3.2 kV, 3.6 kV, 4 kV, and 6 kV, which corresponds to electric fields in the liquid of 2 kV/cm, 2.3 kV/cm, 2.7 kV/cm, 3.1 kV/cm, 3.4 kV/cm, 3.4 kV/cm, and 5.1 kV/cm.

In this test it was assumed that the extraction efficiency had reached 100% when at 6 kV, which was based on the results of the flat energy resolution and linear signal size response from Section 4.9.2.

With the results from the $^{83\text{m}}\text{Kr}$ and ^{137}Cs test it was decided that the best value for the anode voltage was about 5 kV. This corresponds to a drift field in gas of 8.3 kV/cm, which corresponds to an extraction efficiency of about 78-88% [68]. With this value the extraction efficiency is high enough such that it will not be a major factor in S2 energy resolution while still being a reasonable enough size that the PMTs won't saturate if the DAQ doesn't rail, even at the highest of source measurements.

4.10 Position-Based Quality Cuts

4.10.1 Introduction

In order to optimize the energy resolution in the detector one must make a number of cuts to the dataset in order to factor out extraneous events, which is mainly done through the event position. The z-component of the event location can be inferred through the drift time of the event (S1/S2 separation), while the xy position was estimated using a weighted sum of the top PMT array, which was chosen due to the proximity of the S2 events.

In addition for this energy resolution study only single scatters were used, which can be easily localized in the detector in the small signal acceptance regions. It is possible with better knowledge of the position reconstruction that multiple scatter events can be added as well into the energy resolution study. This can be done for future work of this type.

4.10.2 Planar Quality Cuts

The xy positions for each event were calculated using a centroid reconstruction from the S2 signal top PMT array. Quality cuts here exclude events on the edge of the detector, which can be from

background events or unwanted Compton scatters and typically are in highly non-uniform electric field regions, and thus have a significantly different PMT response than the response in the bulk. Furthermore, a large amount of background is known to be at the edge of the detector due to impurities from the walls of the detector as well as Compton scattering from external detector materials that have a partial energy deposit near the edge of the detector.

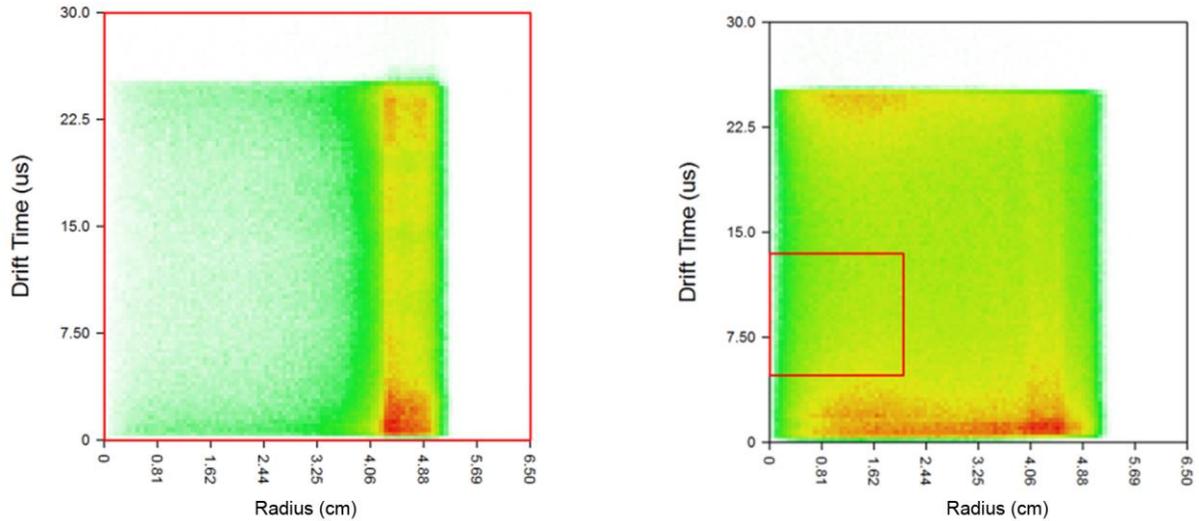


Figure 4.27: Fiducial region of ^{137}Cs spectrum. Right: Fiducial region of ^{228}Th spectrum. The red box on the ^{228}Th spectrum shows the quality region used for parameter optimization.

As shown in Figure 4.27, the profile in the detector can vary depending on the source being used; however there are common characteristics. There are large background regions at the edge of the detector that have hard cutoffs at $r \sim 5\text{cm}$, $r \sim 4\text{cm}$, and $r \sim 3.5\text{ cm}$, which most likely correspond to the various hexagonal edges in the detector. The outermost one is the hard edge, which gives us an idea of the distortion in geometry; it should roughly be 7.5 cm at the farthest edge. Going further in we would expect to see an “edge” that corresponds to being perpendicular to the plane of one of the faces of the detector, which in these pictures is most likely the 4 cm – 3.5 cm cutoff lines.

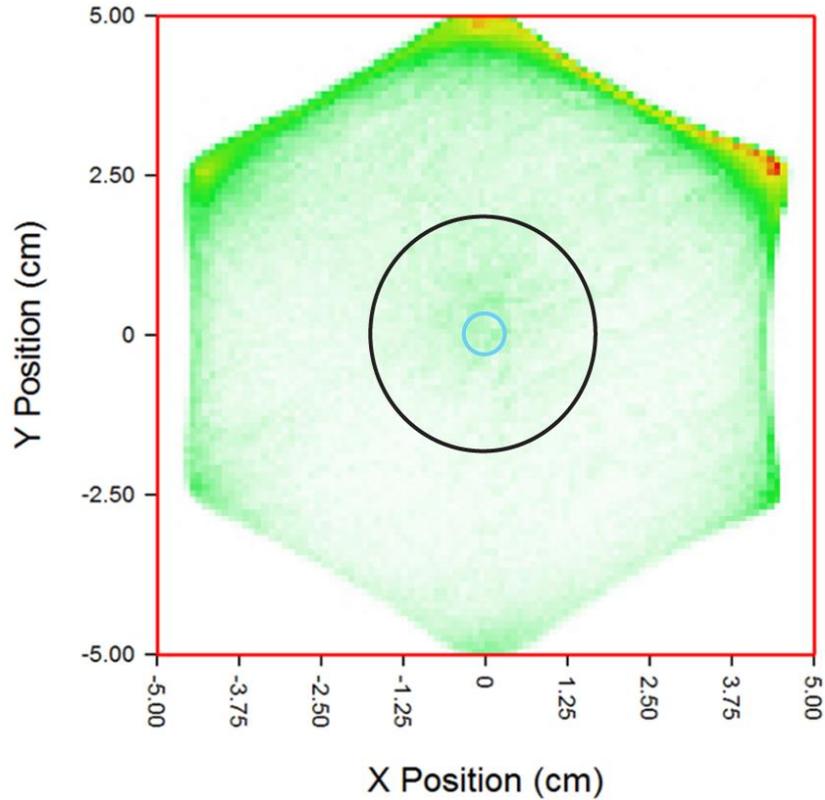


Figure 4.28: XY scatter plot showing the optimization region in black as well as the smaller region for ^{22}Na , ^{137}Cs , and ^{228}Th in blue. The source was on top relative to this image which is why there are more background and edge counts on that side.

For the ^{137}Cs quality region there aren't many events that make it to the central region past $r < 3.25\text{cm}$. This is most likely because of the stopping power of liquid xenon for the 662 keV gamma-rays, which for this energy is $0.4\text{ keV}/\mu\text{m}$ [69]. This is a problem because it means that a large portion of the acquired events must be excluded because they are not in regions where the electric fields are uniform. Also shown are the background events from the top and bottom of the detector, which are most likely impurities and radioactivity from the gate and cathode surface events as well as liquid surface events at the top of the liquid surface. For both top and bottom regions the skin depth is about 2-3 μs .

For the ^{228}Th in Figure 4.27 we see that the events are much more evenly distributed throughout the detector, although the hard cutoffs that were in ^{137}Cs are visible as well. This adds evidence to the fact

that those are the correct edges, because they show up in about the same places but for a different source that has different penetrating characteristics and were also at different PMT gains. In particular, since the ^{228}Th is higher energy than ^{137}Cs (2.6 MeV compared to 662 keV) it is much more likely to make deposits into the central regions. Thus the size of PIXeY starts to be useful at higher energies, since for smaller R&D detectors it is usually not possible or very difficult to resolve a high energy due to the relatively few amounts of full energy deposits that happen. In addition the skin of high count events at the top and bottom goes deeper at almost 4 μs , which is also most likely due to the penetrating power of ^{228}Th .

Thus based on these studies the events with a calculated centriodal radius $r > \sim 1.75$ cm were cut. This was to make sure that all of the edge events from the detector were excluded as well as minimizing the need for any position dependent corrections, since larger quality regions would require a more precise correction by position which would require better knowledge of the event positions. The $r < 1.75$ cm region the region used for the optimization process described in Section 5. After optimization, a stronger quality cut was used for the ^{22}Na , ^{137}Cs , and ^{228}Th runs due to their higher concentration of central events of $r < \sim 0.3$ cm, which enhanced the energy resolution results.

4.10.3 Depth Quality Cut

The z-position quality cut is done separately since it can be found based on the relative distance between the S1 and S2 signals. The events at both the top and bottom of the detector need to be excluded as shown in Section 4.10.2. The events at the top of the detector occur near the liquid surface, which have a higher amount of impurities, non-uniform electric fields, and also contain events that occur above the gate grid but below the liquid level surface. The events at the bottom of the detector have non-uniform fields near the cathode surface as well as event leakage into the space between the cathode and bottom shield grid.

The minimum value was chosen so as to cut out events that occur between the gate and anode grid. Since the liquid level is above the gate grid, there is a small liquid section that has energy

depositions. Those events do not have the same drift field and furthermore the field in that region is known to be highly irregular. Thus, the characteristics of that region of the detector are completely different and are thus cut.

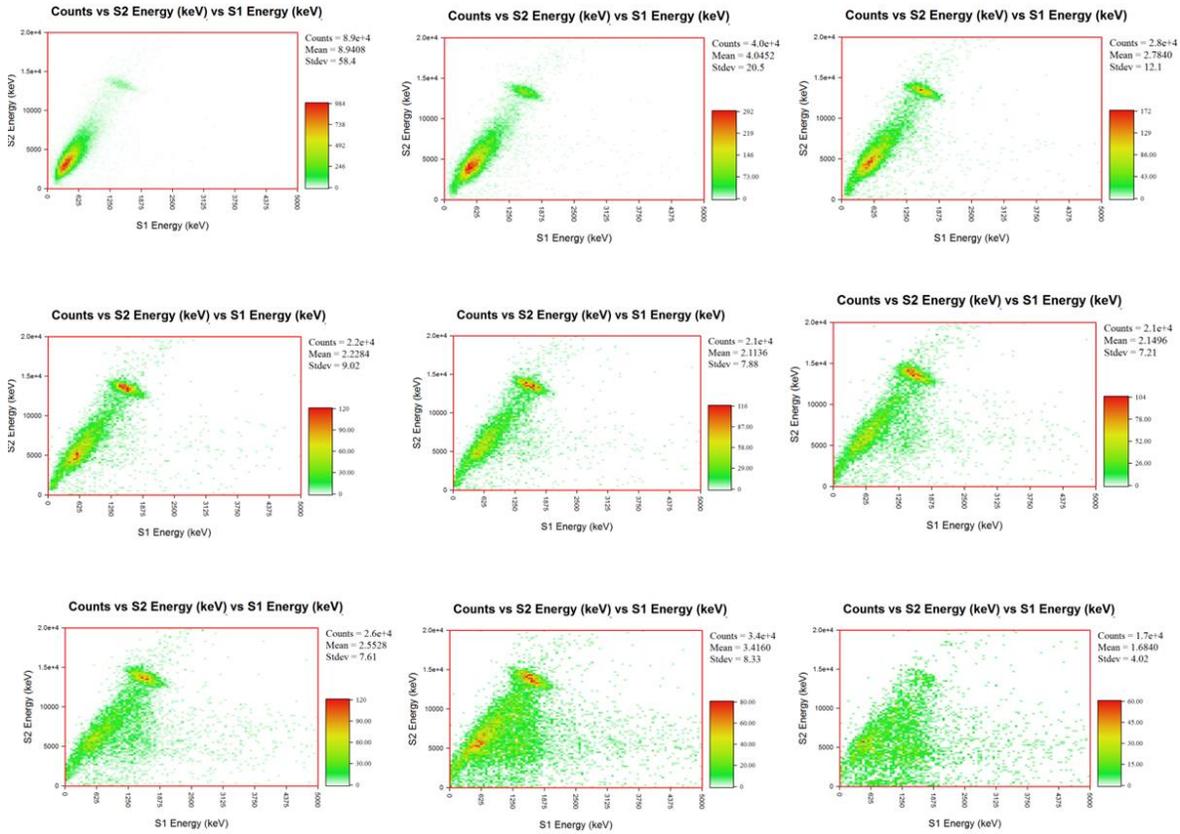


Figure 4.29: Visualization of the degradation of the energy resolution and signal shape as a function of drift time in windows of $3\mu\text{s}$ each. The first window contains mainly events that happen between the gate and anode surface. The smearing effect is due to events Compton scattering into the “dead zone” of the detector, which is below the cathode but above the bottom shield grid. Secondary scatters cannot be detected in this region. The software identifies it as a single scatter and includes it in the spectrum. The smearing has a soft vertical cutoff at the peak energy, here 662 keV.

For the lower bound of the z-component leakage events are cut. Leakage events are Compton scatter events where one or more scatters scatter in the region below the cathode grid. With these events some of the S2 energy is lost because the charge drifts and is lost on the cathode. Based on this 15 μs was chosen as the value with minimal to no leakage while also balancing the number of events. Simply put,

the stricter the fiducial cut the better the energy resolution however this must also be balanced by the number of acquired events⁴.

Thus for this work, the best region was found a few mm below the gate grid, near the top of the detector and thus at drift times between 3 μs – 15 μs . This allowed for a healthy size acquisition reading while still excluding the majority of the waste events at the top and bottom of the detector.

⁴ Since our data acquisition is already on the order of 1-2 days any finer fiducal cut could make acquisition 2-5times longer. Any longer than a few days and we would have to consider long-term integrated effects like lifetime or PMT signal drift on the waveform datasets.

5 Spectrum Optimization – PIXeYOptimizer

5.1 Introduction

Many parameters go into the calculation for the combined energy resolution, all of which must be precisely known. However, since many of these values are correlated with the signal size the relative energy resolution (σ/E) has less systematic error associated with it. This allows us to focus on parameters which would change the relative energy resolution.

The parameters that can potentially change the relative energy resolution are: relative weights on the PMTs for the S1 and S2 energy calculations, the S2 electron lifetime, the S1 and S2 planar position dependence, the S1 depth dependence, and the S1 and S2 anti-correlation angle. So maximally there are seven unique sets of parameter spaces which impact the relative energy resolution which must be precisely determined. Typically there is an independent measurement that determines the value of the parameter, which is prone to measurement error. This measurement error can negatively impact the energy resolution.

Instead of measuring the parameters independently a program was designed that checks against a large array of possible values for each parameter to see if one minimizes the energy resolution. This resembles an unsupervised machine learning algorithm that “learns” the best values based on a training set of data. The measurements aren’t being made directly, the parameters are cross-checked with energy resolution, and the value that minimizes the energy resolution is chosen.

Investigations and considerations were made on how exactly to design a parameter optimization program. This was done in order to match the algorithm with the detector response. For instance, when considering all of the parameter spaces, one must first decide if there is to be any special ordering for those parameter spaces. It was decided that any sort of PMT weighting optimization should be done first, since this could change the value of any geometric dependent parameter. Once the PMT weightings have

been calculated the geometric parameters for the S1 and S2 peaks can be calculated. Finally, after all geometric parameters have been calculated the anti-correlation angle can be optimized, which is the last parameter.

5.2 Exploratory Studies

5.2.1 Introduction

The methodology of the PMT optimization process was to only consider the relative weightings on the PMTs. For instance, suppose we have a system of just two PMTs. For such a system only a single weighting factor needs to be applied to each of them in order to minimize energy resolution:

$$S = p(PMT_1) + (1 - p)(PMT_2) \quad (5.1)$$

where S is simply the weighted sum of the PMT signal—the justification for this is explained in Appendix A.4. Since the goal is improved energy resolution and not absolute signal size, the set of all possible values of p is simply $p = [0, 1]$.

The most straight-forward generalization would be to use $N-1$ weighting factors for N PMTs. However, we can take advantage of the S1 and S2 light collection profiles in order to design an optimization process that is appropriate for the detector response. Information about the system can be used to design the weighting optimization. Therefore studies were done based on the light collection profiles of the system with the goal of aiding in the design of the optimization algorithm.

5.2.2 Light Collection Profile Investigation

Studies were done to determine how the light is divided in the detector, which was done for the optimization region in the detector described in Section 4.10. The fiducial region for this study was $r < 1.75$ cm and a drift time window of [3-15] μ s.

The central PMTs were investigated first, since it was surmised that they would see the most light. This was indeed the case, for unsaturated S2 signals, with the top-center PMT seeing anywhere between 15-25% of the light. As explained in Section 4.7.3, this was an inexpediently large amount of light that a single PMT saw, which is why so many early datasets had saturation issues. The bottom center PMT saw anywhere between 14-18% of the light. The interesting thing here is that since the variation in light collection for this PMT is so small the relative weighting on this PMT is not as much of a factor. Since its light value is relatively constant any weighting factor added to this will simply move the mean signal size but it won't change the σ/E value.

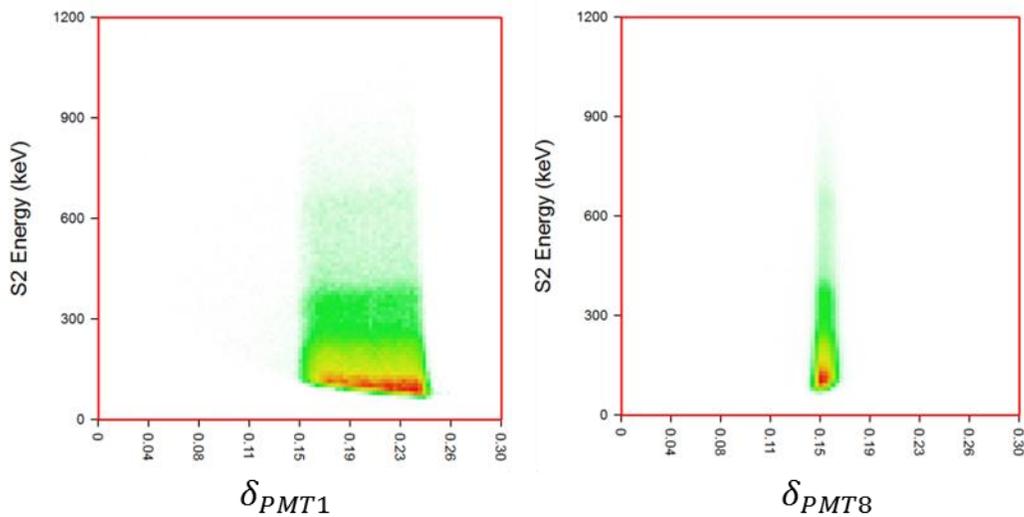


Figure 5.1: Top and Bottom center PMT light collection profiles for ^{228}Th data at 250 V/cm drift and PMT gain mV-ns/phe with an amplifier with a gain of 8.

For the off-center PMTs they see much less light on average than the center PMTs. Interestingly though is that they almost all see a nonzero amount of light. This is most likely because the quality cut used is the central core of the detector so the outer ring of the PMTs have some level of S2 light. This makes it clear that all relative weighting factors are important because each PMT sees some light. For the top off-center PMTs they see anywhere from 2-10% of the light and for the bottom off-center PMTs they see anywhere from 4-15% of the light, as shown in Figure 5.2.

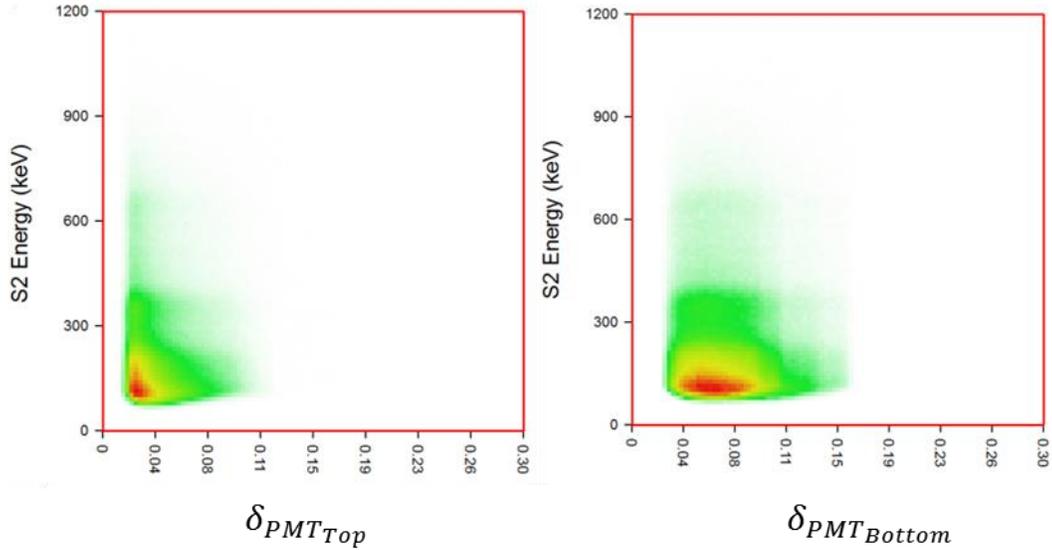


Figure 5.2: Top and Bottom off-center PMT light collection profiles for ^{228}Th data at 250 V/cm drift and PMT gain 0.25 mV-ns/phe with 8 gain Amplifier.

The unsaturated S1 light profiles differ from the S2 profiles. The top center PMT sees only about 1-9% of the light and the bottom center PMT sees 15-30% of the light, as shown in Figure 5.3. This is almost certainly due to the total internal reflection of the liquid xenon surface, which means that the majority of the light is captured by the bottom array [60]. The top off-center PMTs see anywhere from 0-5% of the light and the bottom off-center PMTs see anywhere from 4-18% of the light. Most likely this means that the optimization of the weighting factors for the top PMT array for the S1 light will be less important since the top PMT array only sees about 20% of the S1 light on average, and the majority of the light being on the bottom center PMT.

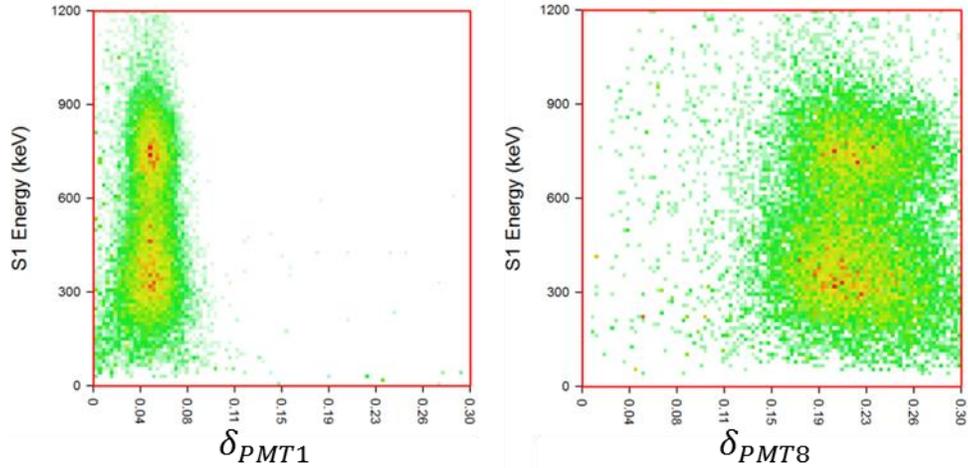


Figure 5.3: S1 light collection profiles in ^{137}Cs for the top and bottom center PMT.

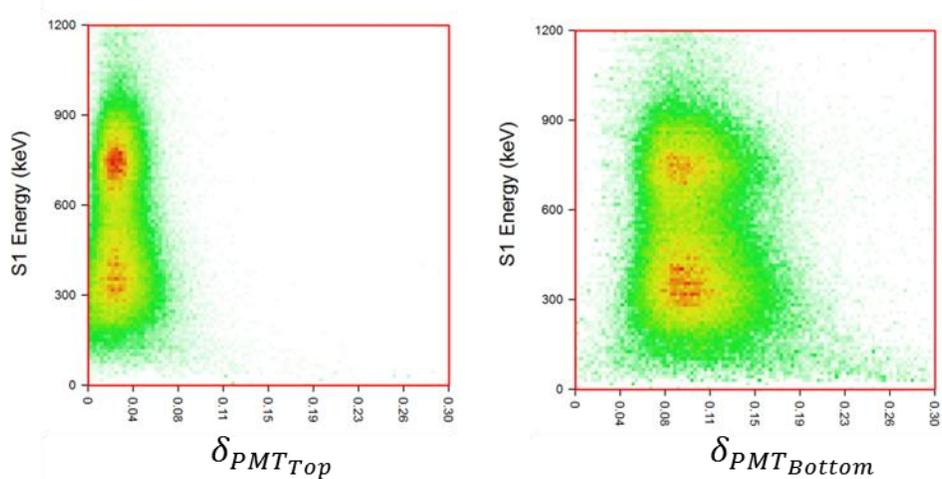


Figure 5.4: S1 light collection profiles in ^{137}Cs for the top and bottom off-center PMT.

With this knowledge of the light collection profiles the peaks in these arrays were studied. Specifically if only a fraction of each array is used is it still possible to find an energy peak that can be used in the optimization process. Optimization can be done on the entire PMT array however that means the algorithm will be less sensitive to changes in the weighting factors since the errors tend to average out across the different PMTs. Thus, the methodology is to isolate as many PMTs as possible before

optimization so that the optimization process will be more sensitive to each individual weighting parameter.

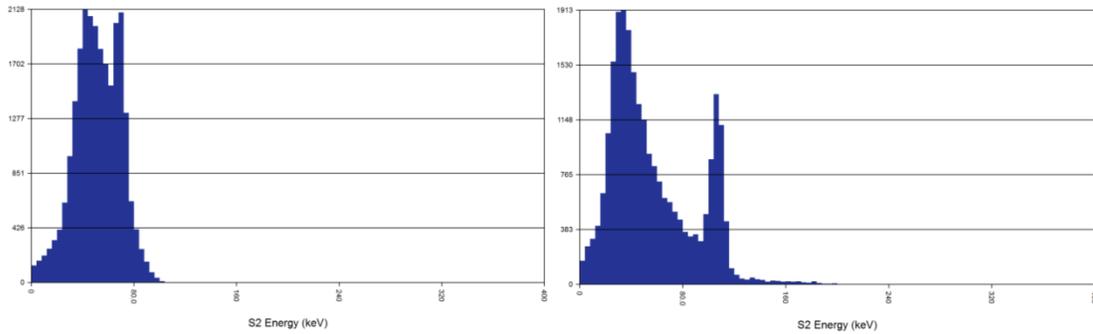


Figure 5.5: S2 Top center and bottom center. Peaks are clearly visible in both spectra.

For the top and bottom PMTs alone it was found that peaks could be found in those PMTs, as shown in Figure 5.5. This means that optimization could be done relative to that PMT and the peak response can be studied by changing parameters.

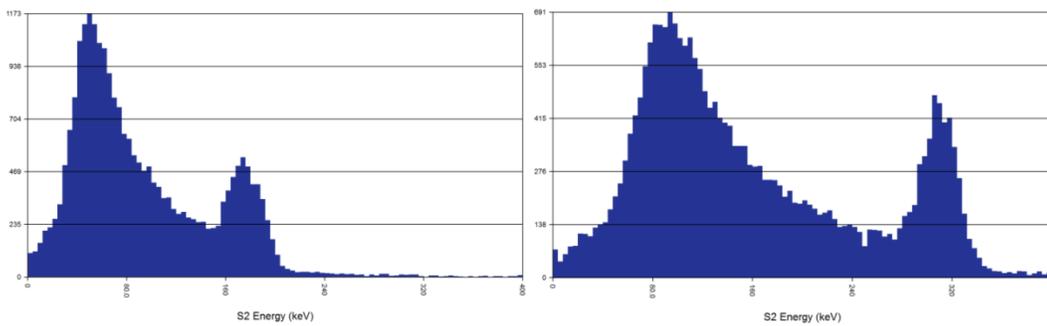


Figure 5.6: S2 Top off-center and bottom off-center. Peaks are visible only if the entire off-center array is considered

For the top and bottom off-center PMTs peaks were able to be resolved provided that the entire off-center array was used, as shown in Figure 5.6. This informed the design process of the optimizer,

which meant that for optimizing the off-center PMTs it can be made relative to that entire array, but not less than that. This splits the PMT arrays into 4 sub regions: the top PMT, the bottom PMT, the top PMT ring array, and the bottom PMT ring array.

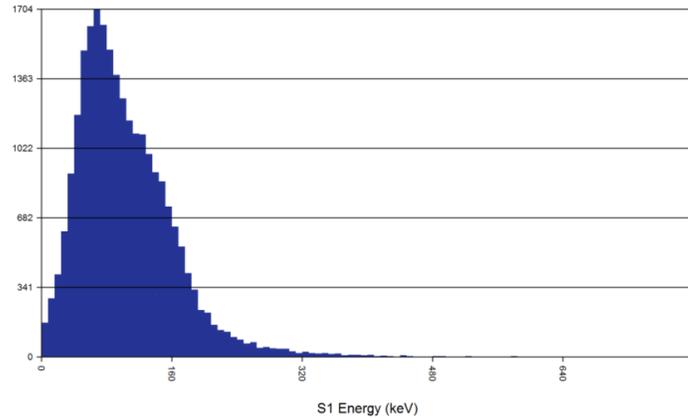


Figure 5.7: S1 Top PMT array. No clear peak is visible even with the central PMT added in. A slight shoulder may be visible.

The same mythology was used when looking at the S1 arrays. The discretization of the PMT array was studied so that the optimization process could be more sensitive to small changes in the relative weighting factors. For the S1 signal however since only 20% of the light is in the top array no peak was observed, even when including the center PMT. Thus it would not be useful to try to optimize the top off-center PMTs for the S1 signal. Even so, an optimization is still run comparing the center PMT to the off-center array since the top center PMT still sees a nonzero amount of light. Thus, the only time the top PMT array is used for the optimization algorithm is when it is compared to the bottom PMT array and when the center is weighted against the ring. Otherwise the process is identical to the S2 PMT array profiles, with peaks visible in the bottom off-center ring and the bottom center PMT as shown in Figure 5.8.

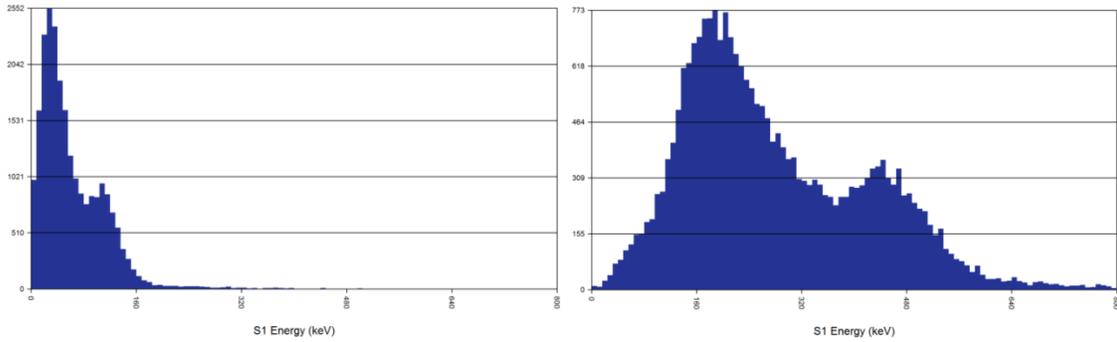


Figure 5.8: S1 bottom center and bottom off-center PMTs. Peaks are visible in these spectra and once again for the ring it is only true if the entire array is used.

5.2.3 Light Collection Profile Simulations

With the knowledge of the light collection profiles a program was written in order to estimate the broadening of the energy resolution due to improperly weighted PMTs, which can be due to anything from variations in the quantum efficiency between PMTs to incorrect PMT calibrations or PMT calibration fluctuations.

A Monte Carlo simulation was made that simulated “detectors”. This was a two-layer iteration program with the first being a detector configuration and the second being an individual energy event. Each detector had a random but fixed set of incorrect PMT weightings based on a Gaussian probability distribution with known error percentages. The second layer of the iteration was to reconstruct an energy event based on the known light collection profile of the detector. The procedure for getting the results is as follows: the initial energy value is randomly varied based on some pre-determined “ideal” energy resolution. This number is then broken up into the components of the 14 PMTs based on known light-collection patterns, with randomized amounts based on the variances in the distribution. Finally, the energy is reconstructed by adding all of these parts together with the “incorrect” weighting factors based on each detector simulation. 10,000 events were done for 10,000 different detectors, and four PMT weight errors were used.

The simulation was run for three error configurations: 10%, 20%, and 50% all with Gaussian errors, which means that the values could be broadened or reduced by the error amount according to a Gaussian distribution. In addition, the simulation was run with two different intrinsic energy resolutions: 1% and 4%. This was an attempt to normalize the results based on signal size; only the ratio to the Gaussian width and mean should matter, so the simulation was run with two different energy resolutions in mind.

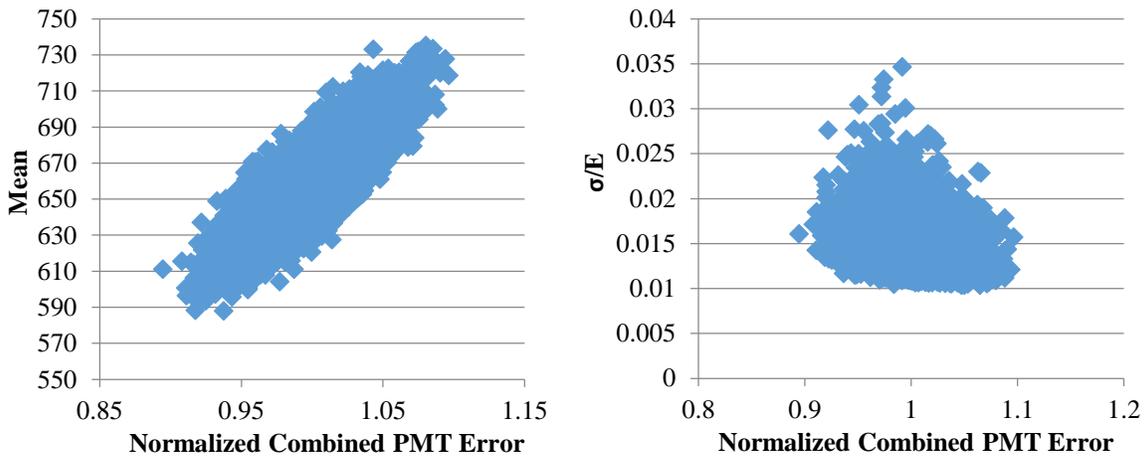


Figure 5.9: Top: Overall signal mean of a 662 keV simulated spectrum as a function of the normalized combined PMT error. The signal mean is correlated with this due to the fact that if the PMT errors are slightly above 1 then the signal will most likely appear as being slightly larger. Bottom: σ/E as a function of normalized combined PMT error. 0.01 is the initial energy resolution before convolution which is why the plot asymptotes to it.

The results were that the better the PMTs were calibrated to the smaller the improvement would be in tweaking those values to get them correctly. In addition no set of PMT weightings could get better than the initial energy resolution, which is a consistency check that this method cannot produce energy resolution values that are unphysical. Most graphs resemble Poisson distributions, since this resembles a counting operation. As shown in Figure 5.10, there is a mean deviation of energy resolution from 1% to 1.4% so that most often times a 10% error from PMT calibration will result in a 40% worsening of the energy resolution for an intrinsic resolution of 1%.

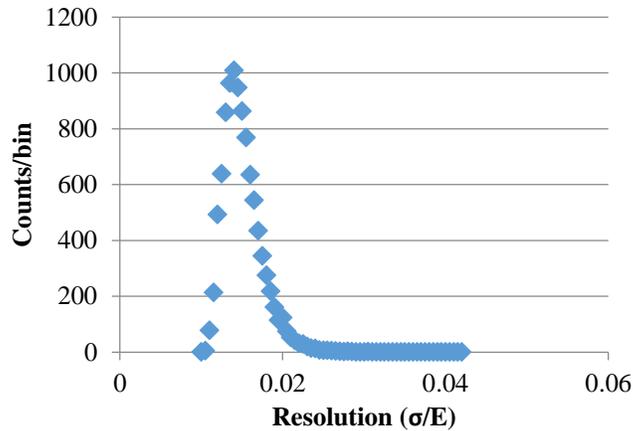


Figure 5.10: Energy resolution spectrum with an initial energy resolution of 1% with 10% PMT weighting error. The resulting spectrum is a Poisson distribution which shifts the mean energy resolution to $1.4 \pm 0.2\%$.

When the weighting errors are larger the average mean error broadens as well. As shown in Figure 5.11, the spectrum gets shifted to 2.2% and 5.3% for 20% and 50% weighting error respectively. This means that if the PMT calibration process is poor then it's more likely that energy resolution values will suffer, especially if the intrinsic value is low like 1%. However, when the intrinsic energy resolution is not as good, then the effect of incorrect PMT weighting error is not as pronounced. As shown in Figure 5.12, the intrinsic values do not shift as much relative to the initial baseline value. In fact, at the worst PMT calibration error of 50% the shift was from 1% to 1.5%, which was almost as bad as the 10% calibration error when the intrinsic was 1%.

It seems that when the intrinsic energy resolutions are good, then any PMT weighting error is more likely to have a large impact on the spread of the energy resolution. In all cases when the intrinsic was set to 1%, the relative spread was much larger than when the intrinsic was set to 4%. These fluctuations start to matter when looking at the combined resolution, because a fluctuation in the PMT weighting error, even if it originates in the S1 or S2 signal, can propagate to the combined energy resolution signal, which can limit energy resolution performance, since the combined energy is the best resolution.

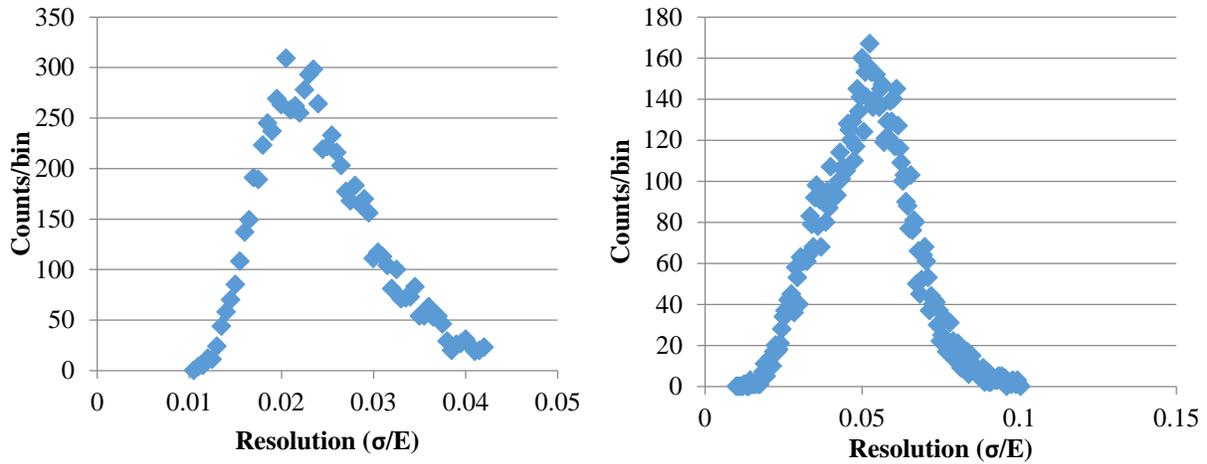


Figure 5.11: Top: Initial energy resolution of 1% with 20% PMT weighting error. The spectrum is shifted to $2.2 \pm 0.6\%$. Bottom: Bottom: 1% initial energy resolution with 50% PMT weighting error. The spectrum is shifted to $5.3 \pm 1.2\%$.

This means that in the S1 and S2 signals a PMT calibration error will manifest itself more closely to the simulation with 4% intrinsic energy resolution, that is it will be difficult to detect. However, since many other fluctuations cancel out at the level of the combined energy resolution the calibration fluctuations will manifest in the combined, but not be as easily detectable in the S1 or S2 signals alone. Thus care must be taken to optimize the S1 and S2 signals because a small improvement in one of those signals can become a big improvement in the combined energy resolution.

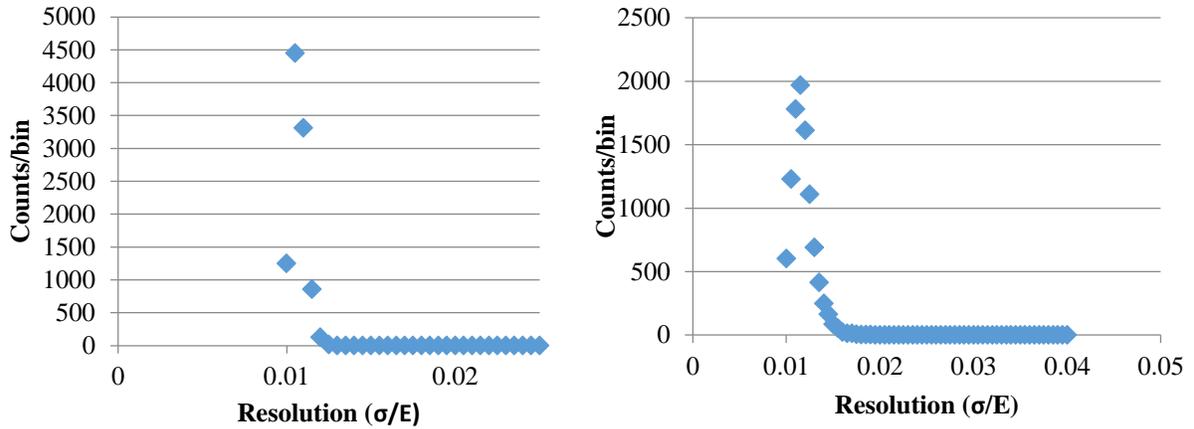


Figure 5.12: Initial energy resolution of 4% with 10% and 20% PMT error renormalized to 1%. The spectrum is shifted to 1.07% and $1.1 \pm 0.1\%$. The shift is less when the intrinsic energy resolutions aren't as good but they are still shifted by a non-trivial amount.

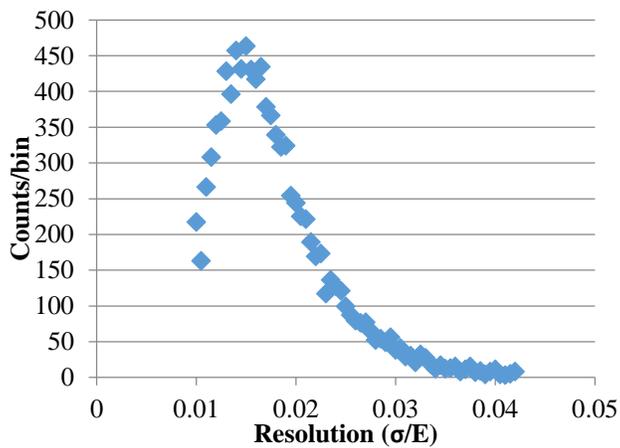


Figure 5.13: Initial energy resolution of 4% with 50% PMT error renormalized to 1%, which shifts the spectrum to $1.5 \pm 0.4\%$.

5.3 Program Design

5.3.1 PIXeYIterator

The first version of an optimization algorithm was to take a generic set of parameters that have a generic range and do a cross-referencing matrix to every parameter set and then bin and do Gaussian fits

to the resultant peaks. That is to cross-reference all possible sets of parameters to get a picture of the entire space of energy resolution values. This process, although complete, is very computationally intensive. Just considering the PMT weights the computation time alone would be extremely long. As an example with 14 PMTs each with 10 values for the gain the program would have to do 10^{14} iterations. Even if a single iteration—which includes passing in the values, applying the correcting factors, and then doing the auto-bin procedure—is only 1 μ s then one single pass of this program would take 31 years, which is after optimization with a map-reduced algorithm. Even if this can be passed to a cluster running multiple cores it would still take years to complete. Iterative computation time goes like $O(n^p)$, with p being the number of parameters, so for any true cross-referencing computation program it would take a very large amount of time.

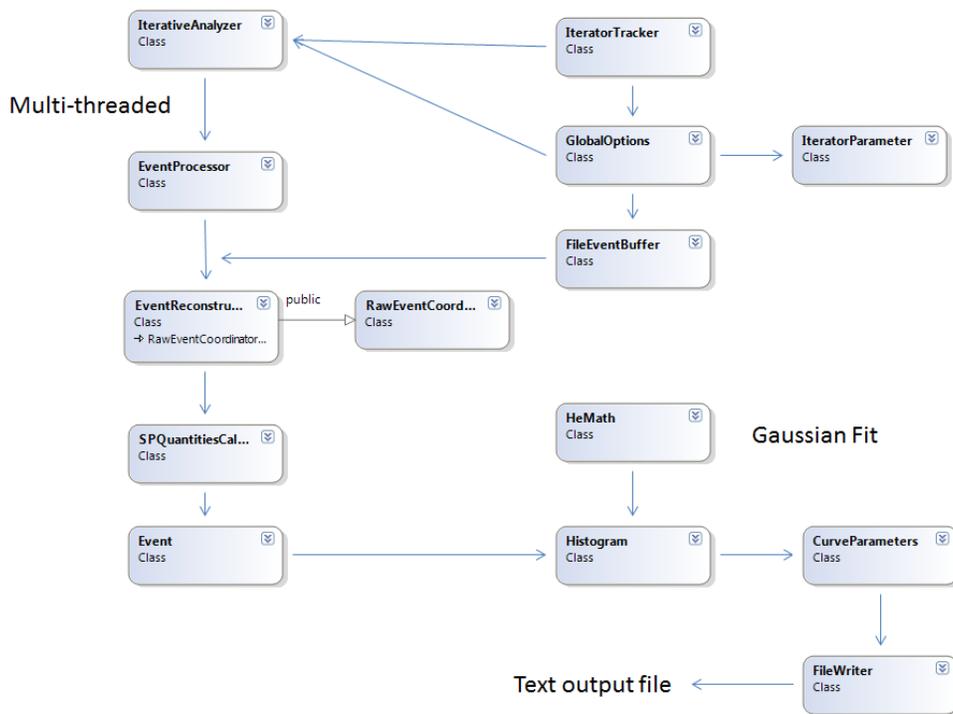


Figure 5.14: Class description and general call list of the iterative optimizer. The IterativeAnalyzer classes uses a map reduced multi-threaded implementation in order to improve computation time.

5.3.2 Optimization Algorithm

The iterator process ended up forming the core of a larger analysis program⁵. Instead of cross-referencing a multi-parameter dataset a single parameter is checked at a time. An energy resolution spectrum that is based on that parameter is constructed and a minimum is found. The parameter that minimizes that energy spectrum is saved and then the next parameter is checked. This process continues until all parameters have been optimized, at which point a final optimized spectra is run to obtain the S1, S2, and combined energy resolutions.

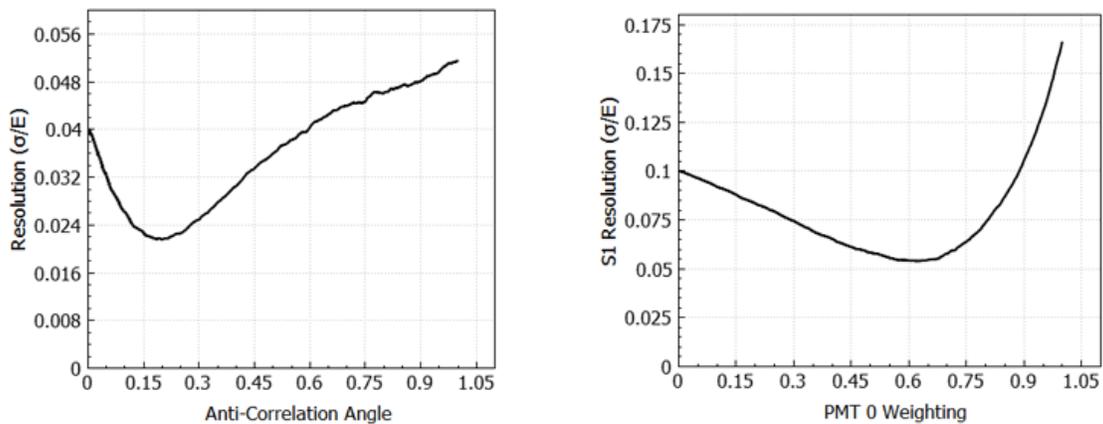


Figure 5.15: Examples of some optimization processes as they run through the various values. In each of them, the minimum is chosen as it reflects the best value for that parameter. Left: iteration through the anti-correlation angle. Right: iteration through PMT 0’s weighting. 0.5 should be the “ideal” value, since it means that it should be exactly half and half weighting between it and the other values. Here it is shown to be shifted to the right at a value of about 0.69, which indicates that more relative weighting should be put on this PMT.

Each optimization process followed a simple algorithm. In the case of the weighting PMT factors all probabilities were sampled and the resulting energy resolutions were computed and a minimum can be found. For each non-complete weighting factor a range will be empirically determined until a minimum with a clear point can be found. The minimum of this curve was used in the next subsequent optimization process. This process continued until all weighting factors were calculated.

⁵ Since the code was made to be scalable from the start this sort of broadening of scope was easier to do from a software engineering perspective. It’s a good lesson for building software engineering projects to be scalable from the start, since it’s usually not known ahead of time what portion of a project may be picked up and advanced on.

5.3.3 Overview of the Optimization Process

It has been empirically observed that the bulk of the light is collected by the top center and bottom center PMT (~20% and ~15% respectively), mainly due to the fact that the data that was analyzed here has a strong central fiducial cut in order to minimize backgrounds interfering with the signal. Thus, the optimization process was subdivided into several different PMT regions: The top PMT ring, the top center PMT, the bottom PMT ring, the bottom center PMT, and the top to bottom PMT arrays.

The weighting factors are:

$$\begin{aligned}
 S_{off\ center} &= p_1(PMT_1) + (1 - p_1)(PMT_2 + PMT_3 + \dots) \\
 S_{center} &= p_{center}(PMT_{center}) + (1 - p_{center}) \sum_{ring} PMT_{ring} \\
 S &= p \sum PMT_{top} + (1 - p) \sum PMT_{bottom}
 \end{aligned} \tag{5.2}$$

Thus there are 15 optimization processes per signal type. However based on the light collection studies a few reductions can be made: the bottom center PMT for S2 does not need to be optimized because the light that it sees is uniform to within a few percent. In addition, only the top to top off-center weighting will be used for the top PMT array for S1, which reduces the total number of iterating factors. Otherwise, they will be done according to ring, ring-to-center, and then top-to-bottom. The parameters will be convolved and normalized when appropriate. This combined with the lifetime, z-dependent light collection parameters, and anti-correlation angle make 28 optimization parameters in total.

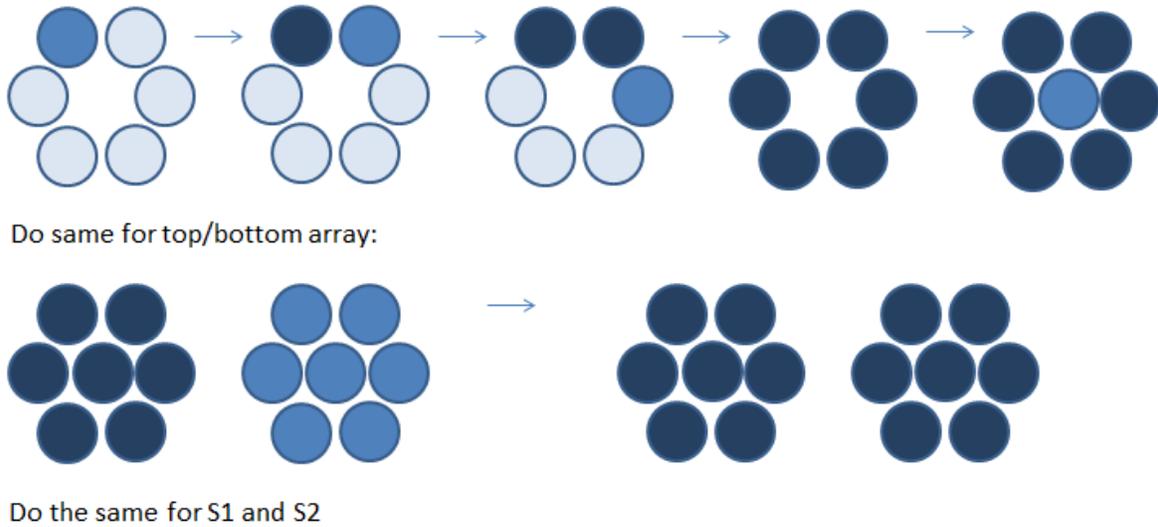


Figure 5.16: Visualization of the PMT iteration process. The process is based on light collection profiles and simulations done in previous sections.

After all of the PMTs weights have been obtained the optimizer will then go through all of the “non-complete” optimization parameters, the lifetime correction, the S1 vs. z correction, any other position corrections, and then finally the anti-correlation angle. These will had test sets that were known before hand to likely contain the minima, but not be so coarse as to limit the precision of the optimizer. For instance, based on lifetime studies in Section 4.5 the lifetime was in the range of 100-200 μs , so a search window that’s [50, 500] μs would be suitable. The evolution of these processes is visualized and annotated in Figure 5.17.

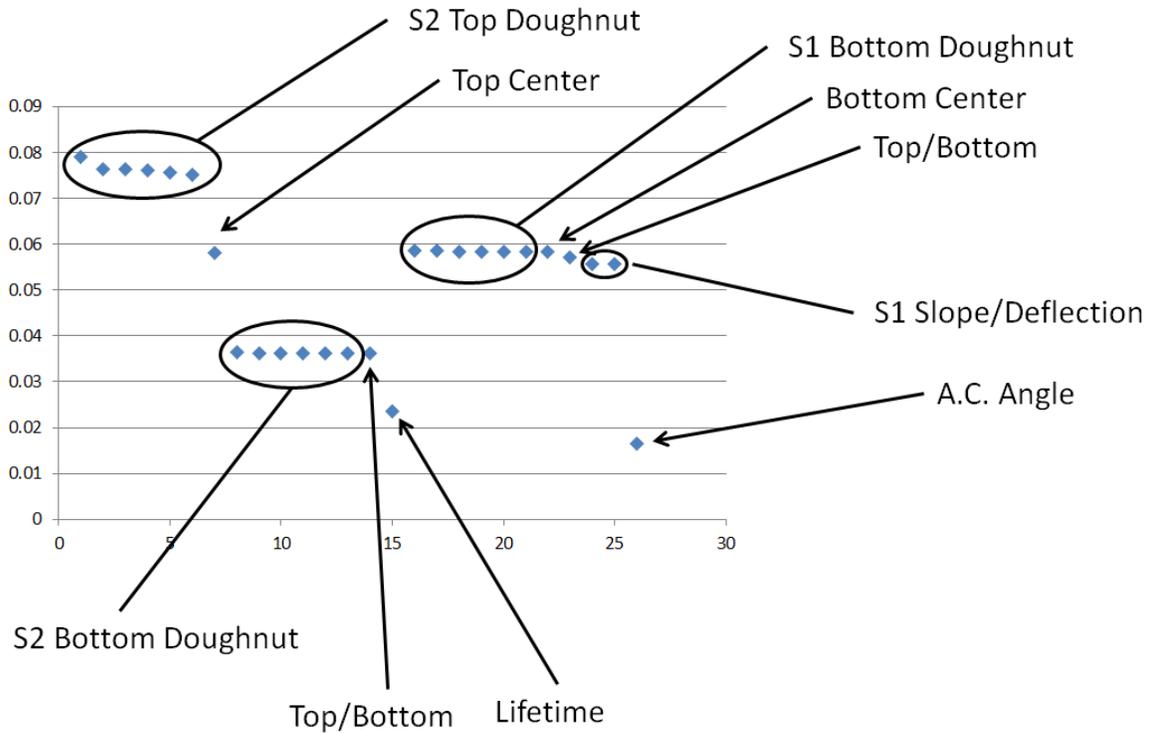


Figure 5.17: Annotated view of each step in the optimization process as a function of energy resolution without xy position corrections.

5.3.4 Auto-binning and Peak Fitting Studies

During the optimization process, data is received as an array of values that have not been binned. The program has to bin and fit the data for each parameter value so that a minimum can be found before going on to the next parameter. A maximum likelihood method of finding the mean and standard deviation of the signal was found per iteration. These values approximate the true mean and standard deviation and this information is then used to find a custom window that is made for each histogram. Each histogram window is centered on this average and the lower and upper bounds are $3\text{-}\sigma$ from the guessed standard deviation. This forms the window in which 200 bins get placed and so an accurate description of the shape can be drawn. Each spectrum had between 5-10k events in total so 200 bins was appropriate for the window.

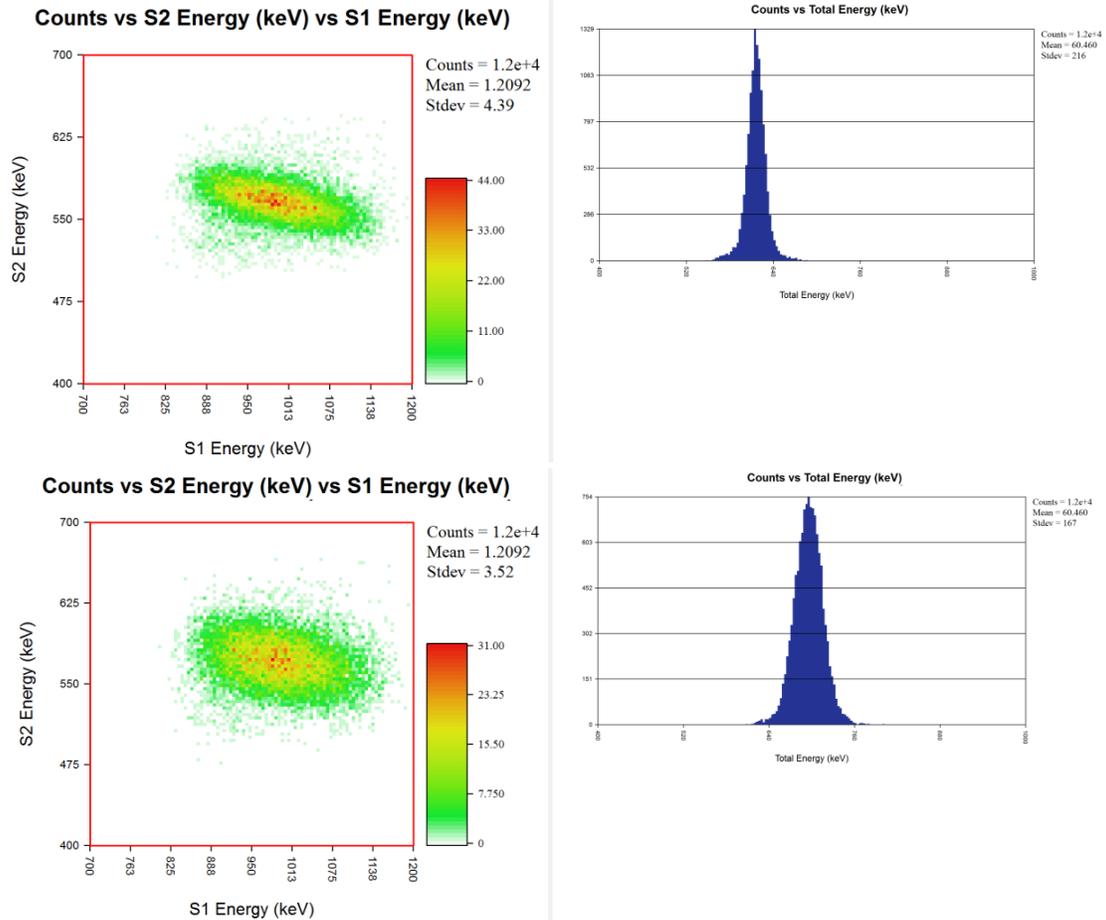


Figure 5.18: Comparison of different fitting routines used during the optimization process. The Gaussian fitting routine produces the better spectrum, perhaps due to the fact that it incorporates more of the spectrum for optimization.

Two different fitting methods were compared to be used in the optimization algorithm. One was a Gaussian fitting routine and the other was a FWHM (Full Width at Half Max) routine. What was found was that the Gaussian fitting routine had an almost 60% improvement in the final combined energy resolution compared with the FWHM one, as shown in Figure 5.18. Interestingly, the FWHM reported better equivalent resolutions at each step of optimization until the final one. This is most likely because the FWHM routine only analyzes the top half of the signal, whereas the Gaussian fit uses the entire spectrum to compare its value. As such, the Gaussian fitting routine was chosen for the optimization algorithm.

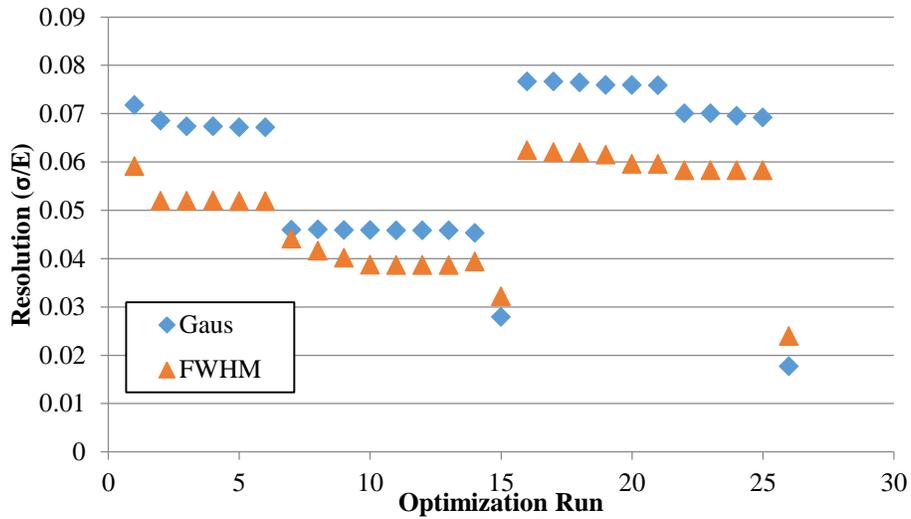


Figure 5.19: Optimization processes and the gradual improvement at each step. Two different fitting methods are compared, one based on a Gaussian fit to the data, and the other a FWHM fit. The Gaussian fit was preferred since the final result almost always had a better improvement, here with a nearly 0.5% difference.

5.4 Results

5.4.1 Visual Optimization Examples

Every source dataset saw a measurable improvement in the energy resolutions after the optimization process. The S1, S2, and combined energy resolution spectra all saw remarkable improvements in their values, many of which are visually apparent as shown in Figure 5.20 and Figure 5.21.

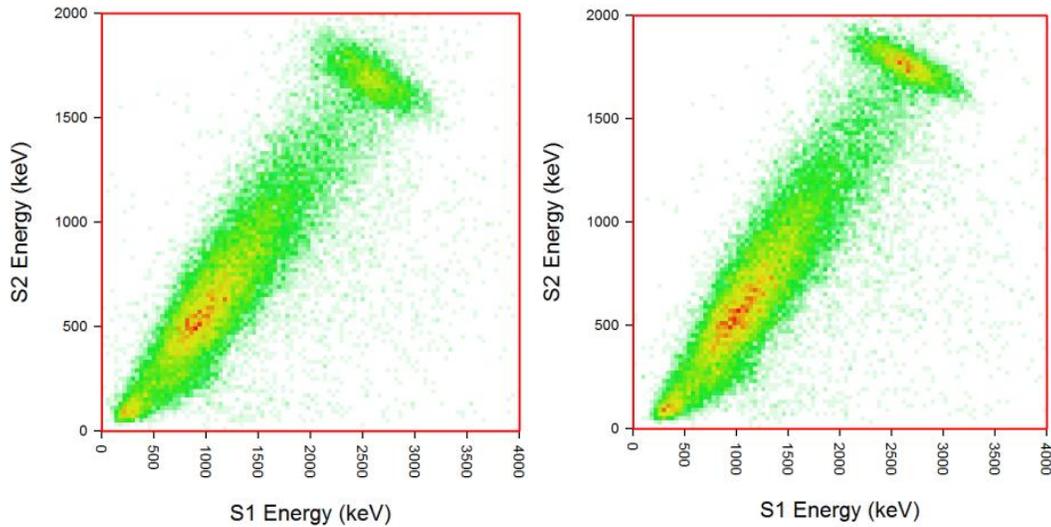


Figure 5.20: Comparison of unoptimized and optimized signal on ^{137}Cs data at 1,000 V/cm. The S1 values go from 0.092 to 0.069, the S2 values go from 0.046 to 0.026, and the combined values go from 0.027 to 0.014. All energy resolution values are in σ/E .

When comparing the before and after spectra for ^{137}Cs as shown in Figure 5.20, the optimizer succeeds in the most important aspect which is pushing events towards the center anti-correlation region. In the after plot there is a clustering of events (shown in red) that push the events into that region, whereas previously the peak events are much more dispersed throughout the anti-correlation region. Interestingly enough there doesn't seem to be a large effect in the Compton continuum spectrum, as that remains nearly unchanged. This is most likely due to the fact that since these events are background then there are no optimization regions for the background to be improved from. It also is most likely the case that certain background events just "moved" to different locations in the same background, thus the appearance is unaltered.

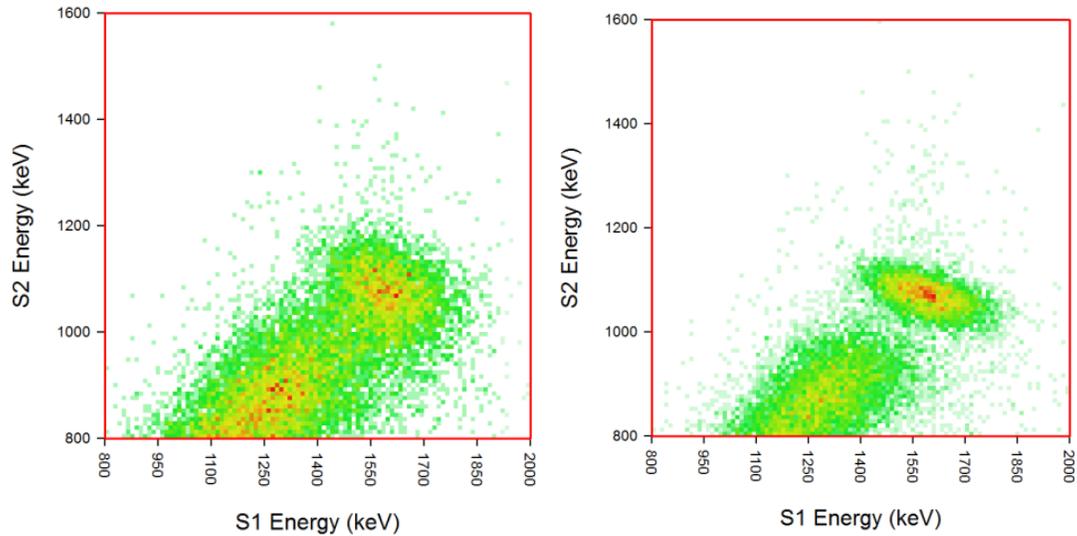


Figure 5.21: Comparison of unoptimized and optimized signal on ^{228}Th data at 1,000 V/cm. The S1 values go from 0.060 to 0.047, the S2 values go from 0.042 to 0.016, and the combined values go from 0.029 to 0.014. All energy resolution values are in σ/E .

When comparing the before and after spectra for ^{228}Th as shown in Figure 5.21, there is a much more dramatic improvement in the visualization of the peak region. There are definite visual improvements in the S1 resolution but the S2 and combined are sharper. As was the case in ^{137}Cs the events are pushed toward the center of the anti-correlation region, which is the main form of this improvement. Also what's interesting as well is that the Compton continuum in these datasets also seems to improve. This suggests that there is a focusing effect on the background which pushes it more towards the central region, which must mean that the effect was not as pronounced in the ^{137}Cs optimization.

5.4.2 Concluding Remarks

While the program was successful in achieving the design goals of energy resolution improvement there remains a few pathways to explore. In particular, there were not many xy corrections that were used in the generic optimization process, mainly because of the relatively low position resolution. A few simple radial corrections were used and attempted, which is studied in Section 6.2. The

main reason why it wasn't widely used was because most of the data had a heavy quality cut on the radius, with the final quality cut being just 3% of the total volume of the detector. This was a "brute force" method of getting rid of planar based energy resolution fluctuations. There are hints that with precise knowledge of the event position more sophisticated correction factors can be used on the planar position, and future work with position sensing grids described in Section 7 will test that.

6 Energy Resolution Measurements

6.1 Introduction

With the detector properly calibrated and analysis and optimization software designed and implemented several energy resolution measurements were performed. All of the energy resolution measurements were made to probe both detector physics as well as to gain a better understanding of the underlying xenon physics. Here detector physics refers to all of the experimental and detector elements, the engineering aspects of the detector. These are PMT calibrations, grid responses to electric and light stimuli, variations in signal response due to imperfect detectors. These are worthy of study but are more engineering driven rather than science driven.

On the other hand is the xenon physics, which is more interesting from a basic science perspective. These are recombination fluctuations and the response of the xenon atoms to gamma-rays. These are in general harder to probe since they are often convolved with detector effects. Thus in general it is desirable to minimize the effects of detector physics so that the xenon physics can be studied, because that is more about the fundamental nature of the system.

Source	Energy(keV)	Gain(mV-ns/phe)	Drift Field (kV/cm)	Gas Field (kV/cm)
^{83m}Kr	41.5	16	1,2	8.3
^{57}Co	122	16	1,2	8.3
^{133}Ba	334	16	1,2	8.3
^{22}Na	511	4	0.25,0.5,1,2	8.3
^{137}Cs	662	10, 5(top center)	0.25,0.5,1,2	8.3
^{22}Na	1274	4	0.25,0.5,1,2	8.3
^{228}Th	2614	2	0.25,0.5,1,2	8.3

Table 6.1: Tables of parameter values used for source runs. The gains shown include the amplifier gain as well as the PMT gain.

As shown in Table 6.1, a range of energies were studied from 41.5 keV (S2 equivalent) to 2.6 MeV. The gains were tuned to maximize S1 energy resolution while keeping the S2 signal unsaturated, as explained in Section 4.8.3. The drift fields that were studied were 1,000 V/cm and 2,000 V/cm for energies up to 334 keV and 250 V/cm, 500 V/cm, 1,000 V/cm, and 2,000 V/cm for energies above 334 keV, which was done to map the energy resolution as a function of drift field (Section 6.3), as well as determine the detector parameters g_1 and g_2 (Section 6.5.1). Finally, the gas field was tuned to maximize extraction efficiency at 75-85% while also preventing DAQ saturation or PMT saturation (Section 4.9).

All of the measurements made with the goal of studying fundamental xenon physics will be discussed. Detector physics and performance diagnostics will also be discussed since that is important for research and development but the end results and goals are most focused on gaining new insights into xenon physics.

6.2 Position Dependence

6.2.1 Introduction

The xy position dependence in the energy resolution signals needed to be measured and studied. The location of the energy peak may vary due to the location of the event, which is most likely due to imperfections in detector design such as grid deflection from electric fields, imperfectly calibrated PMTs, and any other geometric light collection response effect. This blurs the energy resolution due to the systematic shift in peak location. It is desirable to try to eliminate this effect as much as possible for optimal energy resolution results, which has been done in other detectors [66].

6.2.2 No Quality Cuts

The position dependence of a ^{137}Cs source at 1,000 V/cm was examined. This was to look at the aggregate effect of the position dependence on the signal across the whole detector and to see if there is

an improvement with an xy position correction. So, the optimization algorithm described in Section 5 was run with a quadratic radial correction done to the data without any cut in the xy plane, and a drift time window of [3, 15] μs was still used. This position-dependent correction factor was done in addition to all of the other correction factors described previously.

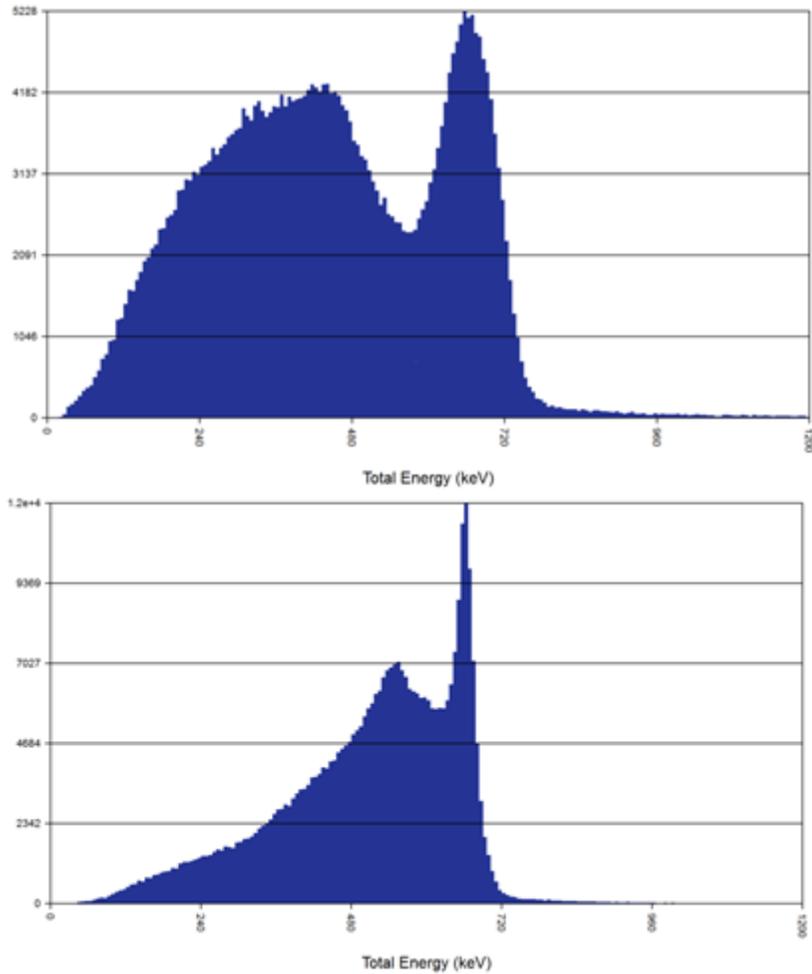


Figure 6.1: Comparison of the combined energy spectrum with no correction factors (other than optimized anti-correlation angle) on top and all correction factors (including quadratic radial correction) on bottom. The resolutions are $4.76 \pm 0.33\% \sigma/E$ and $1.73 \pm 0.12\% \sigma/E$ respectively. By far the radial correction is the factor that optimizes this resolution the most.

Adding a radial correction to the full dataset with no xy corrections greatly improved the energy resolution signal from $4.76 \pm 0.33\% \sigma/E$ to $1.73 \pm 0.012\% \sigma/E$, a 275% improvement, as shown in Figure

6.1. However the best value for this dataset was $1.16 \pm 0.08\% \sigma/E$, which had a heavy radial cut in addition to data optimization. This may be due to the complexity of what a “true” correction would look like, since the correction here is only a simple quadratic fit to the radius, there was no theta correction attempted. This is due to the poor position resolution performance in the detector, which was on the order of 1cm. For better correction factors to correct the spatial variation of the signal within the detector better energy resolution is needed.

Another odd feature of the optimized spectrum was the distortion of the Compton continuum before the peak. Before optimization the top spectrum shown in Figure 6.1 has a Compton peak and Compton edge, as expected. However after the correction there is a long tail towards the lower energies, which seems unphysical. A comparison of the energy to the radial position of the detector was studied.

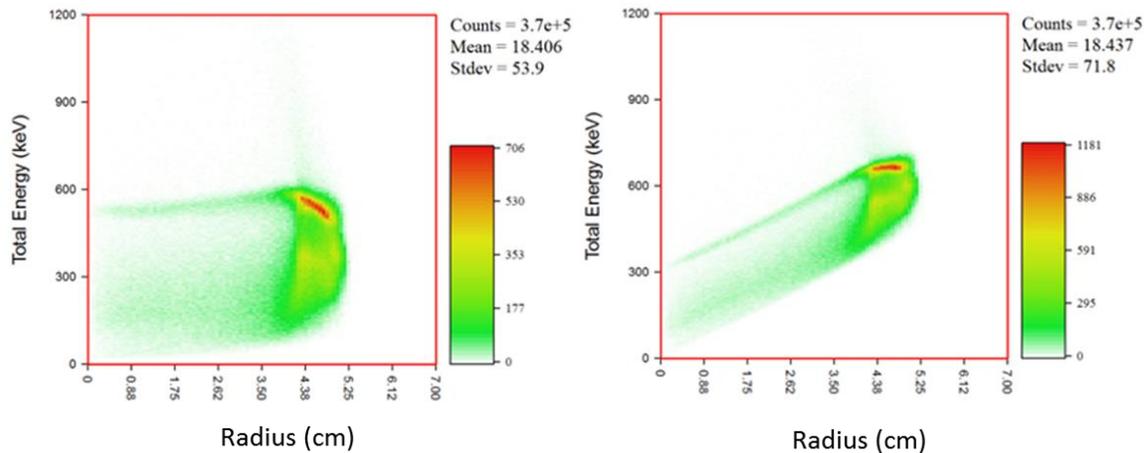


Figure 6.2: hummingbird plots that compare before and after radial corrections of the combined energy spectrum. The eye of the hummingbird is the full energy deposition of the ^{137}Cs 662-keV peak, which decreases towards the edge. This is either a geometric or could be a saturation effect. The optimized spectrum flattens out this area because it by far has the most counts. The beak of the hummingbird is what is usually used to characterize the detector, since that is the inner quality cut range.

The results were that the optimization algorithm had flattened out the region at the very edge of the detector, because that area has by far more counts than in the central region, as shown in Figure 6.2. This is because the optimization algorithm is an unsupervised machine learning algorithm. It doesn't rely

on any assumptions about what the spectrum should look like and instead is trying to minimize the energy resolution in the peak region. In this case the learning algorithm found that there was much more events in a tight region and chose to ignore all of the other events. It's an important lesson on the limits of such learning algorithms: typically most can only find a local minimum when they find convergence, and it's very difficult to design an algorithm than can find a global minimum for a suitably complex system.

6.2.3 Reduced Datasets

For the reduced datasets there was a tight initial quality cut done for the optimization window and then the higher energy sources were constrained even further for their energy resolution values, because there is a remaining signal dependence both on size and energy resolution as a function of radial cut. XY corrections on the S1 and S2 signal were not done for this dataset, but improvements with energy resolution by reducing the quality region occur even when those corrections are performed [66]. A correction factor based on the events' radius was attempted, which improved the energy resolution the optimization region but not in the smaller region, there was no improvement in the best combined energy resolution. In future work a xy position correction could be done on this data or data of similar types both to try and see if there are any energy resolution enhancements but also more importantly to expand the quality cut region for efficiency. It is thought that a position reconstruction improvement would be better for better position resolution, which may be tested in the next version of PIXeY which will have single wire grids for a xy readout of each with position resolution of $\sim 1\text{mm}$.

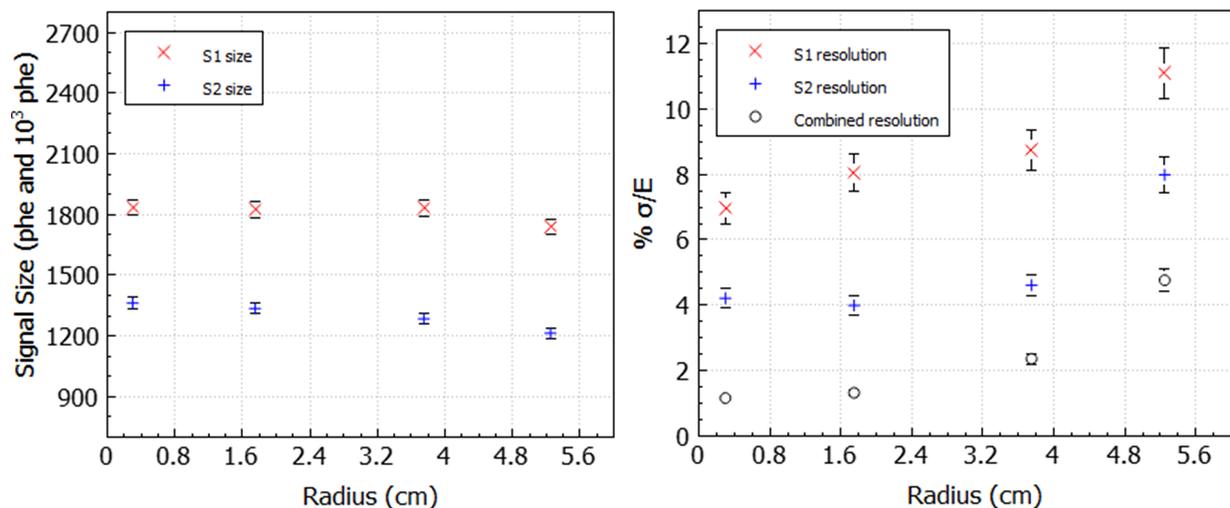


Figure 6.3 :Left: Signal size of ^{137}Cs at 1000 V/cm as a function of radial cut. The S1 y-axis is in phe and the S2 y-axis is in 10^3 phe. Assuming that the signal sizes are constant the S1 and S2 values are 1961 ± 43 and 1329 ± 58 phe and 10^3 phe respectively, which is about a 2.2 and 4.4 percent variation in signal size. Right: Energy resolution of ^{137}Cs at 1000 V/cm as a function of radial cut. It should be noted that all of these optimization points were from a separate optimization run, and there was no position correction attempted.

The best energy resolution values were those with the smallest acceptance regions. For the ^{22}Na , ^{137}Cs , and ^{228}Th the tighter acceptance region was done. This additional quality cut improved the resolutions of all four lines, as shown in Figure 6.4. For $^{83\text{m}}\text{Kr}$ the extra cut did not significantly improve the resolutions. For ^{57}Co there were no usable events in the center, in fact for all ^{57}Co results the outside edge of the detector had to be used to find a peak due to the poor absorption into the center, which is why the energy resolution of ^{57}Co is so much poorer than other values. And finally for ^{133}Ba there weren't enough events in the central region of the detector for a tighter quality cut like ^{57}Co . Thus for energies studied lower than 511 keV There were enough events to optimize the values but there were not enough events to do a tighter radial quality cut.

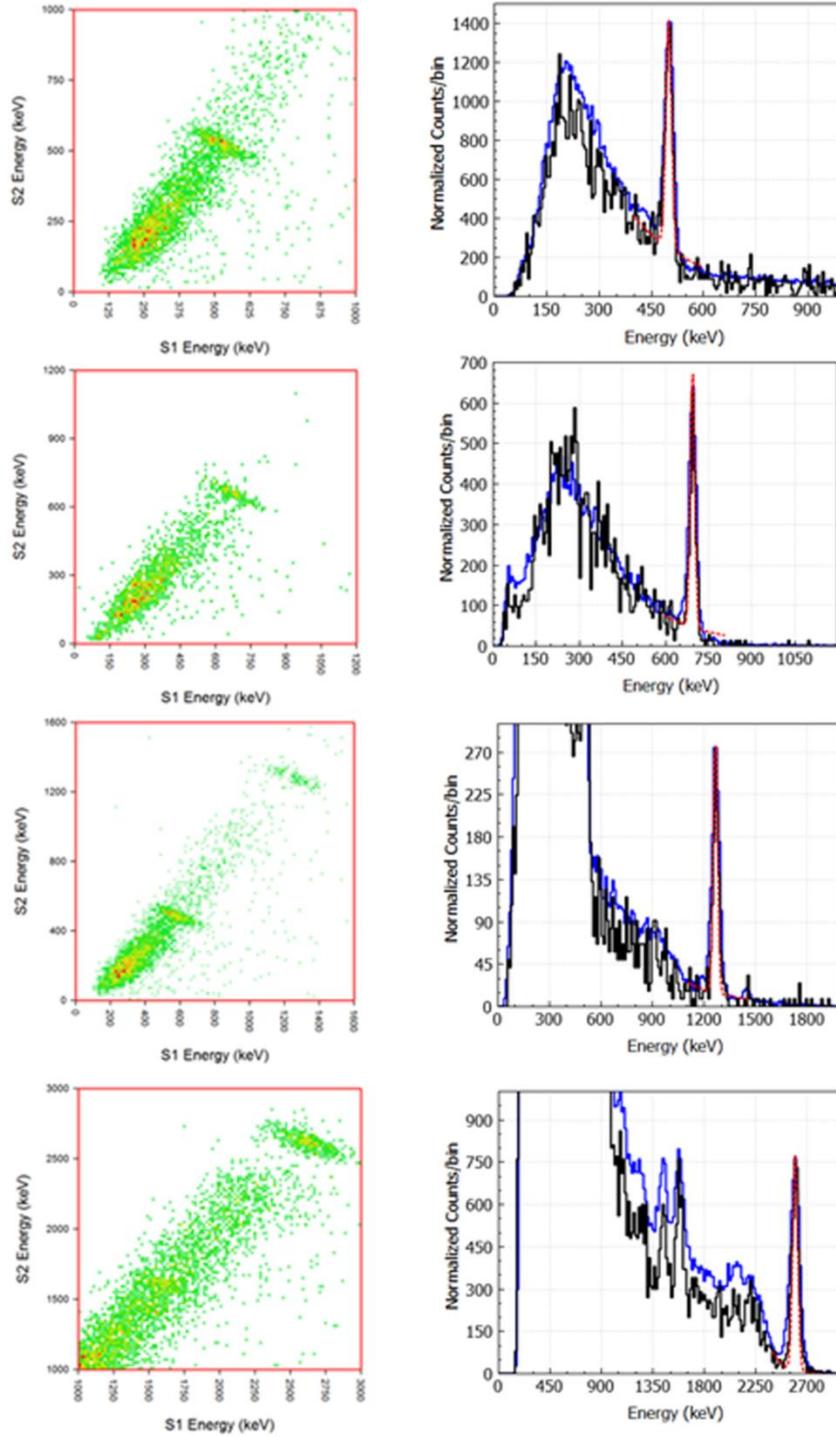


Figure 6.4: Top to Bottom: ^{22}Na 511 keV line with a resolution of $1.71 \pm 0.12 \% \sigma/E$, ^{137}Cs 662 keV line with a resolution of $1.16 \pm 0.08 \% \sigma/E$, ^{22}Na 1274 keV line with a resolution of $1.01 \pm 0.07 \% \sigma/E$, and ^{228}Th 2614 keV line with a resolution of $0.85 \pm 0.06 \% \sigma/E$. Left: S1 S2 pixel plot of the spectrum. Right: Energy spectrum with 1.75 radial cut in blue and 0.3 radial cut in black. The red fit is for the 0.3 radial cut. All data shown here is for 1000 V/cm drift field.

6.3 Electric Field Dependence

This section is a report on the energy resolution performance and how it relates to applied drift field. Data was taken with a variety of different fields and the S1, S2, and combined energy resolutions were studied.

6.3.1 Reduced Datasets

There is a standard prediction to the energy resolution performances by increasing the strength of the drift field. If the drift field is increased this increases the amount of escaped charge in the proportional scintillation signal, which should improve the S2 energy resolution since an increase in the amount of charge should improve statistical-based fluctuations. Conversely, The S1 signal size should be decreasing with field, since less charge is recombining to add to the S1. This implies that the energy resolution should be poorer at higher drift fields due to the relative increase in the number of statistical fluctuations. Finally, since the S2 pulse has the better energy resolution then the combined energy resolution should improve as the drift field is increased.

As shown in Figure 6.5, the S2 energy resolution improved as the drift field was increased for all sources. In addition, the combined energy resolution appears to improve as the drift field is increased, and the values flatten out and even in some cases get slightly worse at high drift fields. For this work the best combined energy resolutions were at 1,000 V/cm.

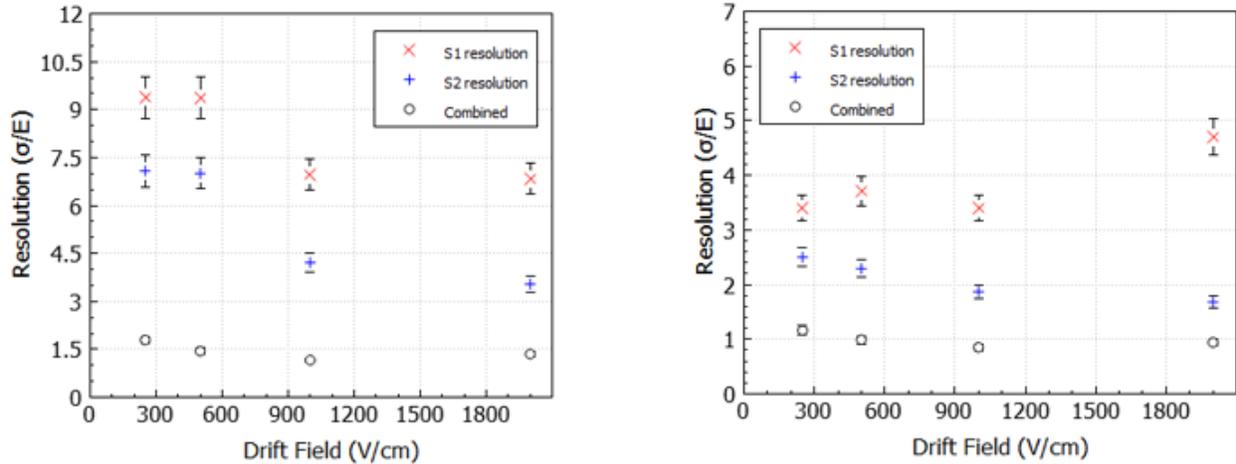


Figure 6.5: Left: S1, S2, and Combined energy resolution as a function of drift field for the ^{137}Cs runs. The S1 energy resolution appears to improve as drift field is increased, which may be due to a decrease in recombination fluctuations at higher drift field. This effect is observed for the ^{22}Na (511 and 1274 keV) source and ^{137}Cs source Right: S1, S2 and Combined energy resolution as a function of drift field for ^{228}Th source runs. The S1 energy resolution is mostly flat with an increase at 2000 V/cm which may be due to signal size issues.

But the interesting thing is that, for the ^{22}Na and ^{137}Cs sources, the S1 energy resolution appeared to improve with increasing drift field, changing by as much as 50%. This may be because the S1 energy resolution is flat with drift field or due to the fact the fluctuations themselves are decreasing. This can be true if the relative field strength in the micro region of the event in the detector is proportionally stronger than the local field produced by the freed charges and not just bigger in an absolute sense. With only three lines showing this it cannot be concluded, and it could simply be some unknown systematic effect.

6.4 Comparison with Other work

The best energy resolution results for all source runs and fields are shown on Figure 6.6. In terms of energy resolution values PIXeY has the best values for energies at and above 662 keV to date, which is

shown on Figure 6.7. The set of data that was used for this comparison was all of the results obtained at 1000 V/cm drift field, which had the best overall energy resolution values.

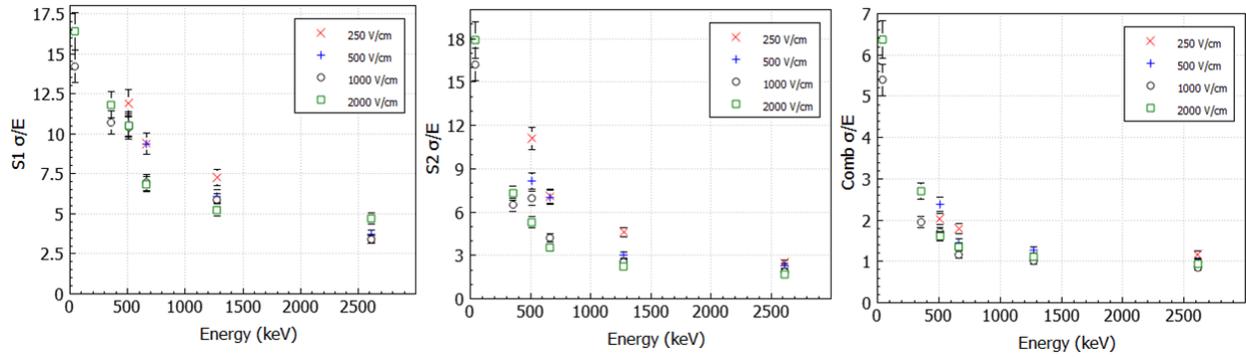


Figure 6.6: All energy resolution results from S1 pulses (left), S2 pulses (center), and combined (right).

As shown in Figure 6.7 our values are higher at lower energies due to the larger size of our detector and the inability to get events into the center of the detector. The ^{57}Co point was not used for this fit since its energy resolution value was outside the normal quality region. The fit shown in Figure 6.7 is A/\sqrt{E} fit with $A = 1.11 \pm 0.03\%$.

One possible issue with the 511 keV line for this work was that the gain for that source was tuned for 1274 keV rather than 511 keV. It's possible that that energy resolution value could be improved if the gain of that source was higher, perhaps a separate run could have been done for 511 keV. This is due to fluctuations in the S1 signal due to SNR problems, which is explained in Section 6.6.

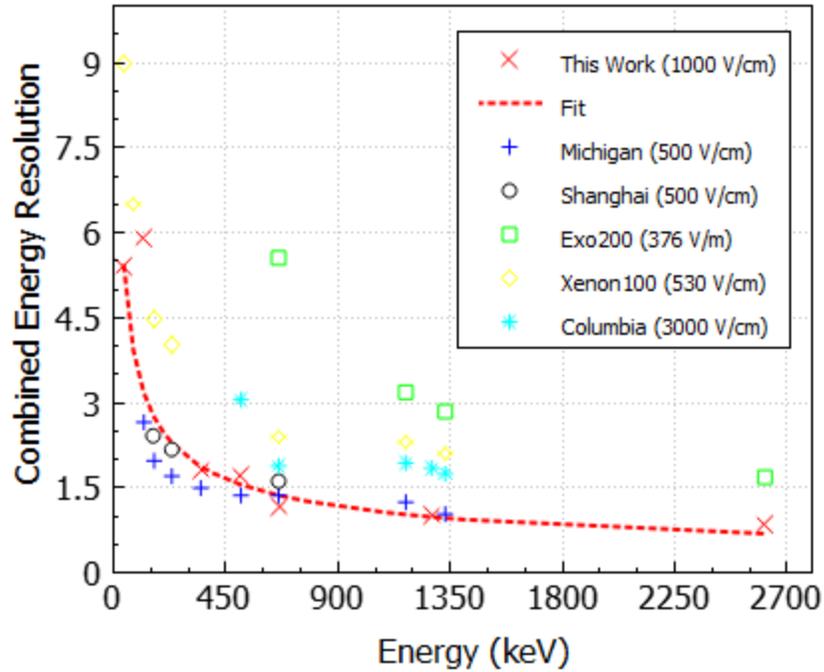


Figure 6.7: All combined energy resolution compared with previous work with specified drift fields shown [66] [70] [71] [72] [11]. The ^{57}Co data was not used for this fit since it was outside the quality range from the other data points.

6.5 Light and Charge Yields

6.5.1 Calculation of g_1 and g_2

In order to model the results in terms of physical quanta such as photons and charge, the light and electron yield in the liquid xenon must be obtained. To calculate this, multiple source spectra were compared by varying the drift field in the detector which changes the ratio of charge and light in the detector response. By plotting how the ratio of the photoelectrons obtained by the source energy is changing with drift field, the parameters g_1 and g_2 can be extrapolated, which convert the phe signal into light and charge quanta respectively. This is called a “Doke” plot [73], and for enough values a line should be obtained and the parameters can be extracted. Since different drift fields for the same source run are used a large amount of systematic uncertainties can be controlled for—like PMT gains.

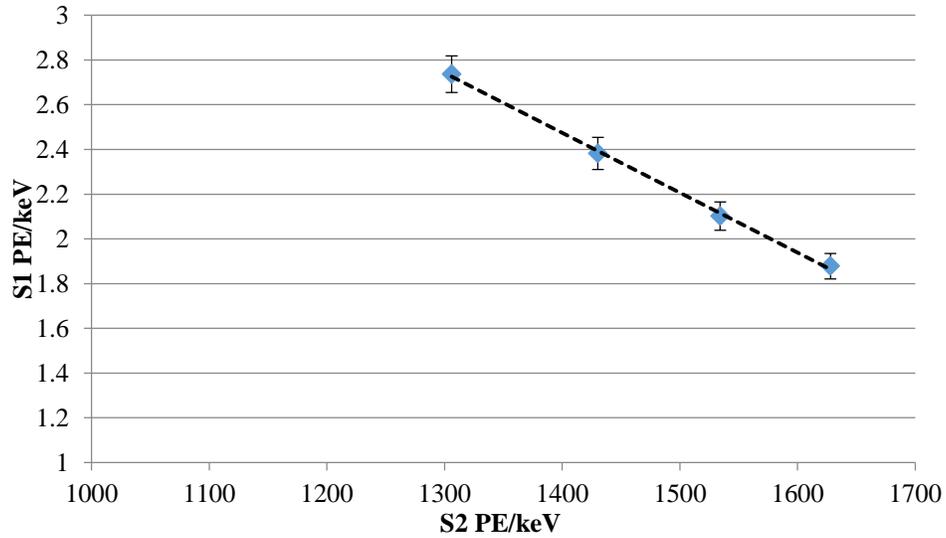


Figure 6.8: An example of a Doke plot made by a single source (^{137}Cs) and by varying the drift fields. Since only the drift field is varying across these sets most sources of error are correlated and do not affect the goodness of fit, but it will affect the values of the constants. Since the energy of the source is known the photoelectron response can be normalized by the energy in keV.

Physically g_1 is the probability that a photon in the field region becomes a photoelectron in the PMTs. This means that it is a combination of several correction factors including the probability of creating a photon in the liquid, the light collection properties of the detector, and the quantum efficiency of the PMTs. Based on past work this value should be around 7-8% [62] and we expected similar results from our detector. The parameter g_2 is the number of photoelectrons produced in a PMT from the drift of a single extracted electron. In addition to all of the factors above (with probability of creating a photon changed to probability of creating an electron) it also includes drift region specific values like the extraction efficiency from the liquid to the gas as well as the number of photons that a single drifting electron makes. A value from a similar detector for g_2 is about 31 phe/e [62] and this expected to be consistent with PIXeY as well.

Extrapolating the constants in the best fit curve $y = mx + b$ in Figure 6.8 gives the equations for g_1 and g_2 to be

$$g_1 = b * W \quad (6.1)$$

$$g_2 = \frac{b * W}{m} \quad (6.2)$$

where W is the work function in liquid xenon, here assumed to be 13.6 eV/quanta [56]. With these two measured values the S1 and S2 yield can be converted into number of photons and ions respectively. We can directly observe the charge and light fluctuations and compare those values from results obtained in simulations to not only see how well PIXeY operates compared to other liquid xenon detectors, but also to gain insight into the detector operation itself, and see if there are any systematic effects or variations in the parameters themselves as well as the calculated quanta.

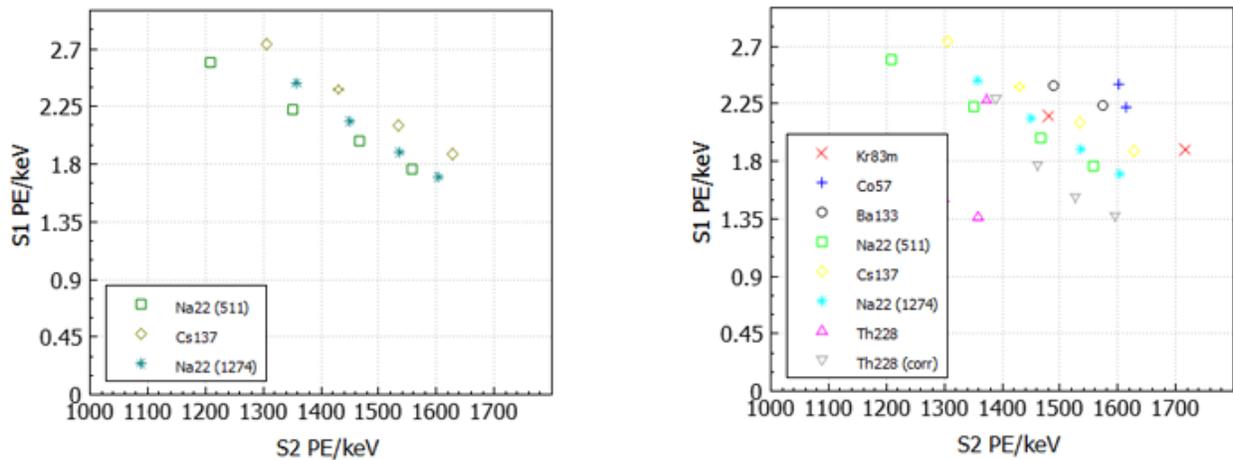


Figure 6.9: Left: Doke plot of all data taken for this run. The ^{228}Th run most likely had saturation which explains the unusual shape. This can be corrected by the light and charge yields. Right: Doke plot of values used to calculate g_1 and g_2 .

Shown in Figure 6.9 on the right are all source runs. It is apparent here that there was some saturation effect on the ^{228}Th data, which is why the points are non-linear. When the Thorium points are corrected for signal response we see that the points converge to the line at less than 10%. For the ^{57}Co points they were not in same fiducial region as other sources because of its inability to penetrate deeper into the detector. And finally the ^{133}Ba and $^{83\text{m}}\text{Kr}$ points seem to have a different shape and slope entirely,

which may be due to the systematic gain difference between those points and the higher energy points, since they were at lower gains.

For this work the calculation of g_1 and g_2 was done on the ^{22}Na 511 keV and 1274 keV and ^{137}Cs 662 keV points. The ^{228}Th data was not used because of the saturation present in the dataset. Using these 3 sources with 4 different drift values the value for g_1 and g_2 was found to be 0.083 ± 0.006 and 30.7 ± 0.9 respectively.

6.5.2 N_γ , N_e , and N_q

In order to calculate the signal response in terms on number of quanta it is necessary to convert the signal, this is in terms of photoelectrons, to number of quanta. To calculate that we use:

$$N_\gamma = \frac{S1(phe)}{g_1} \quad (6.3)$$

$$N_e = \frac{S2(phe)}{g_2} \quad (6.4)$$

where g_1 is the average photon detection efficiency and g_2 is the product of the electron extraction efficiency and the electron transmission efficiency described in the previous section.

The values of g_1 and g_2 can now be used to calculate the light and charge yields of the detector. The light and charge yields are normalized by the source energy to line match the sources on the same graph. The sources that had the most similar parameters and PMT gains have a variation of less than 8% in S1 and S2, and the smaller energies that were not as optimized and had different PMT gains had a variation of 18% in S1 and 9% in S2. This is shown in Figure 6.10. As expected, the fractional charge yield increases with drift field and the fractional light yield decreases with drift field.

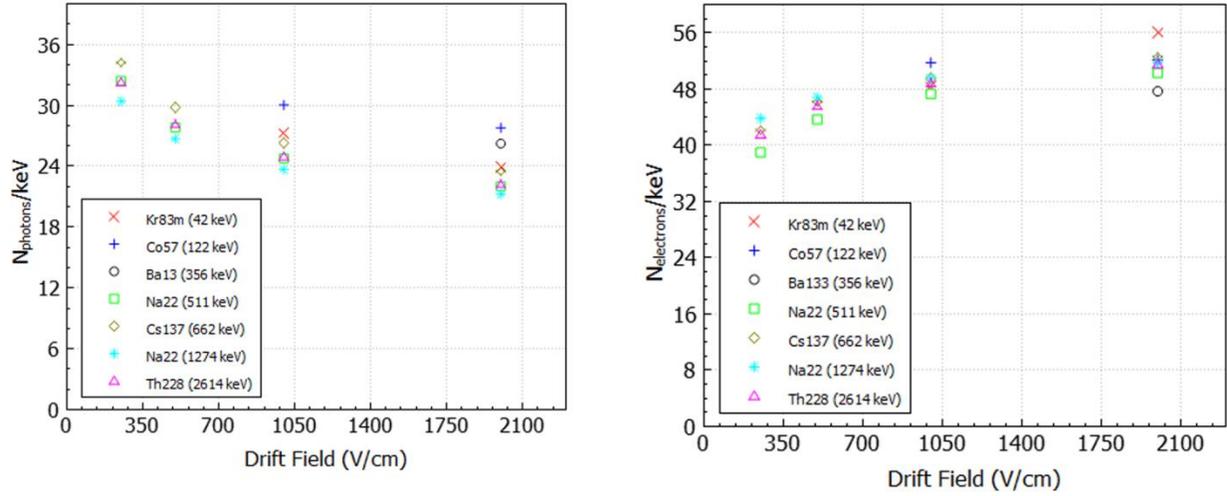


Figure 6.10: Left: Fractional light yields for all data taken in this run. Right: Fractional charge yields for all data taken in this run.

Next the total quantum as a function of energy was constructed. This equation is simply:

$$N_q = N_\gamma + N_e \quad (6.5)$$

This is useful not only to check the signal response of the detector, but also as a calibration between N_q and energy. As shown in Figure 6.11, the total number of quanta is linear with energy, except for the ^{228}Th points. It is here that the saturation of ^{228}Th datasets was discovered. The behavior is expected since the amount of saturation increased with increasing drift field and the 2,000 V/cm ^{228}Th run was saturated by about 15%. Using the extrapolated quanta the Thorium yield was corrected.

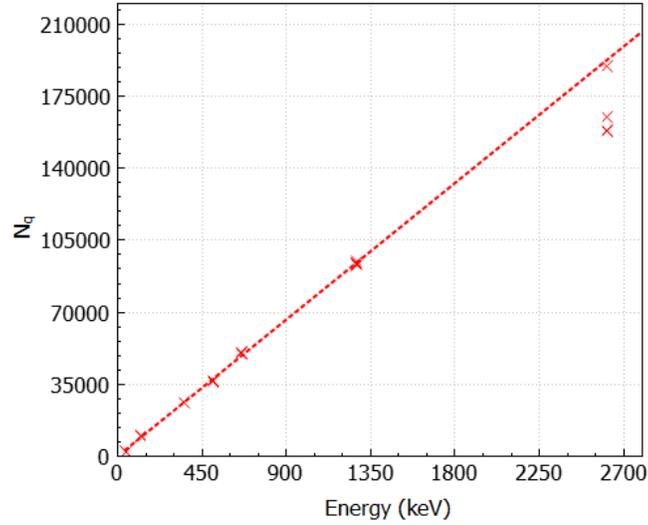


Figure 6.11: Total yield as a function of energy. At the end the Thorium quanta yields are reduced which is most likely due to S2 saturation.

6.5.3 Mean Recombination Probability

Next, the mean recombination probably was compared with source energy. The mean recombination probability is the average amount of ion electron recombination that occurs for a given drift field and energy. It is known to have an energy dependence and has been calculated previously for energies up to ^{137}Cs [62]. The mean recombination probability is defined as:

$$R = \frac{\frac{N_\gamma}{N_e} - \alpha}{\frac{N_\gamma}{N_e} + 1} \quad (6.6)$$

with α is defined as the ratio of excitons to ions. For this work we assume a value of 0.2, it is not expected to be energy dependent [73] [74].

As shown in Figure 6.12 the value appears to increase with decreasing energy at energies less than ~ 300 keV and then flattens out as energy is increased. The asymptotic behavior is field dependent with 2,000 V/cm leveling out to ~ 0.16 and 250 V/cm leveling out to ~ 0.32 . So even though the energy and therefore N_{ions} is increasing the average amount of recombination is the same. This may be due to the

fact that at some energy the electron tracks get to be so big that electrons from one side of the track cannot recombine with electrons on the other side of the track, which makes the function flat instead of changing with energy as previously expected.

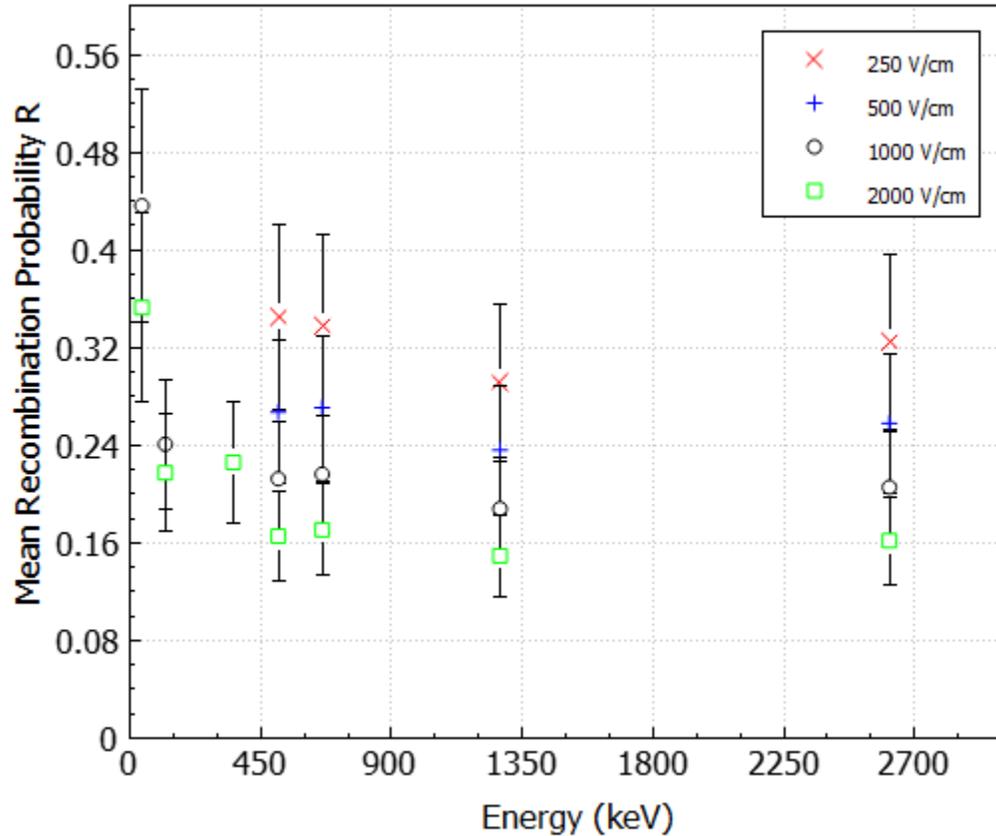


Figure 6.12: Mean recombination probability R as a function of energy for all source runs. At values at and above 511 keV the mean recombination probability can be modeled as a constant with values of 0.325 ± 0.072 , 0.258 ± 0.057 , 0.205 ± 0.045 , and 0.162 ± 0.036 for drift field values of 250 V/cm, 500 V/cm, 1,000 V/cm, and 2,000 V/cm respectively.

6.6 Detector and Recombination Fluctuations

Using the light yields, charge yields, S1 resolution, S2 resolution, and combined resolution it is possible to calculate values that relate to the systematic effects of the detector as well as offer insights to

energy resolution performance. One such quantity is the increase in energy resolution due to charge/light recombination. It can be shown that this can be calculated as [62]:

$$\sigma_R^2 = \frac{1}{2} \left(\frac{\sigma_{S1}^2}{g_1^2} + \frac{\sigma_{S2}^2}{g_2^2} - \frac{\sigma_E^2}{W^2} \right) \quad (6.7)$$

Using this equation and assuming that recombination fluctuations appear in both the S1 and S2 signals, the fluctuations in the S1 and S2 signals with the recombination fluctuations subtracted out can be calculated:

$$\sigma_{N\gamma_{Det}}^2 = \frac{\sigma_{S1}^2}{g_1^2} - \sigma_R^2 \quad (6.8)$$

$$\sigma_{Ne_{Det}}^2 = \frac{\sigma_{S2}^2}{g_2^2} - \sigma_R^2 \quad (6.9)$$

These are the "detector" fluctuations for the S1 and S2 signals and they are composed of random and systematic errors. If we plug Eqn. (6.8) and Eqn. (6.9) into (6.7) the S1 and S2 detector fluctuations make up the combined energy resolution:

$$\frac{\sigma_E^2}{W^2} = \sigma_{N\gamma_{Det}}^2 + \sigma_{Ne_{Det}}^2 \quad (6.10)$$

This means that any systematic effect that is present in the S1 or S2 signals will be present in the combined energy resolution, since it doesn't cancel out like the recombination fluctuations. These are mainly important as diagnostic tools of detector performance. It highlights an issue of the combined energy resolution, which is that there may be errors that are small in the S1 or S2 channels that are difficult to detect in the S1 or S2 channels themselves but they become apparent after calculations and are present in the combined energy resolution.

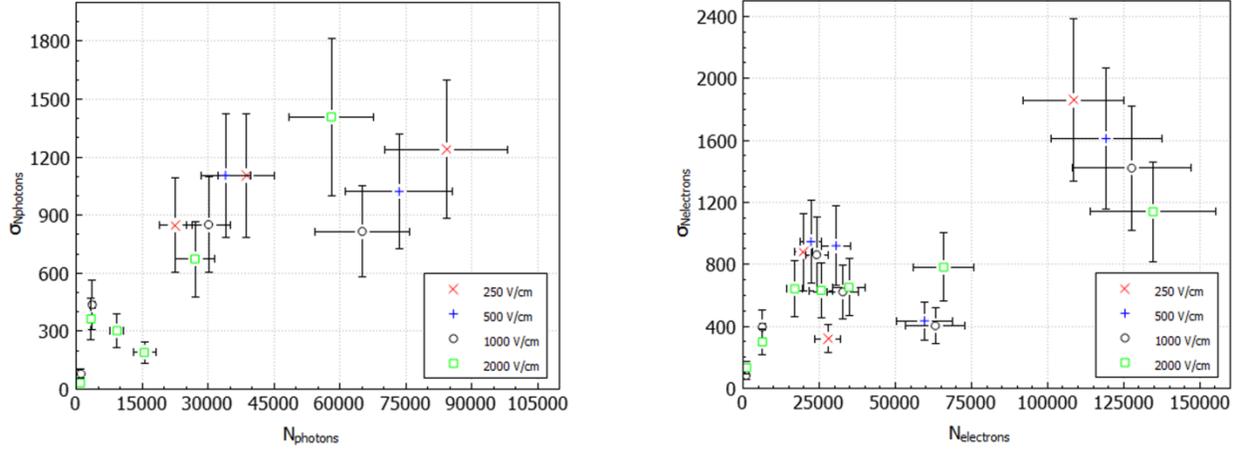


Figure 6.13: Left: S1 Detector fluctuations as a function of photon quanta. Right: S2 detector fluctuations as a function of electron quanta.

Shown in Figure 6.13 are the results of the detector fluctuations for the S1 and S2 pulses. The S1 fluctuations appear to follow a curve that increases as the square root of energy, which is expected based on statistical fluctuations. This curve is described by:

$$\sigma_{\gamma_{Det}}^2 = \left(4.79 \pm 0.26 * \sqrt{N_{\gamma}}\right)^2 \quad (6.11)$$

For this fit $\chi_{red}^2 = 1.69$, which means that this is a slight underfit to the data since $\chi_{red}^2 > 1$. Also since $\sigma_{\gamma_{Det}}^2$ is modeled as a statistical process there isn't any large systematic effects in the S1 channel that would contribute to the worsening of the combined energy resolution.

A linear fit was also attempted for $\sigma_{\gamma_{Det}}^2$, which is described as:

$$\sigma_{\gamma_{Det}}^2 = \left(0.0188 \pm 0.0017 * N_{\gamma}\right)^2 \quad (6.12)$$

For this fit $\chi_{red}^2 = 2.99$, which is again a slight underfit to the data and fits slightly worse than Equation (6.11). Nonetheless, for both fits $\chi_{red}^2 \cong 1$ so it is not definite as to which fit is valid.

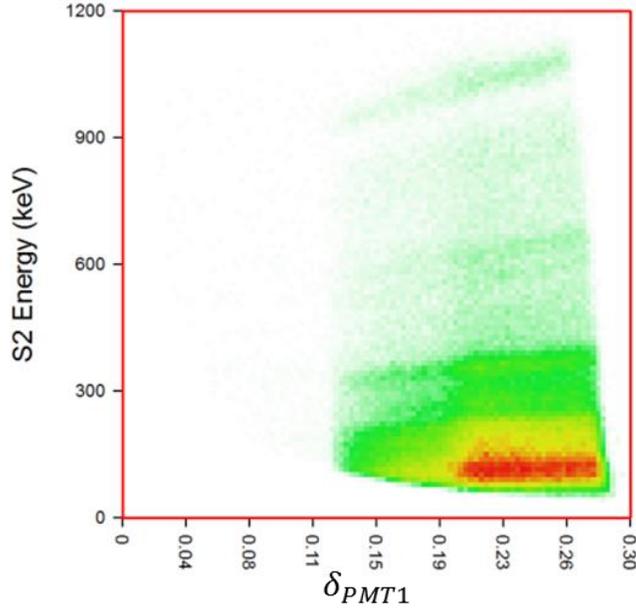


Figure 6.14: Light fraction plots of the top center PMT in the ^{228}Th data set with 2,000 V/cm showing that there indeed was some saturation in the peak signal.

For the S2 signal this also appears to be the case for the majority of the data. However, for the ^{228}Th points on the far right we see that the fluctuations once again increase linearly with number of quanta, which is most likely due to saturation. This is shown in the saturation plot in Figure 6.14. If we ignore the Thorium data points then the curve is modeled as:

$$\sigma_{e_{Det}}^2 = (4.21 \pm 0.39 * \sqrt{N_e})^2 \quad (6.13)$$

For this fit $\chi_{red}^2 = 0.884$, which means that this was a reasonable fit to the data since $\chi_{red}^2 \cong 1$. The Thorium data points are about 14% higher than the curve suggests for 250 and 500 V/cm which is comparable to the saturation amount. It's possible that the intrinsic energy resolution for Thorium could be even better than what's reported here for those field values; by Eqn. (6.13) this would improve the combined energy resolution by about 7%, so a value of 0.79 may be possible that without any saturation in the S2 signal. It is also possible that there is still some systematic electric field dependence in $\sigma_{e_{Det}}$, which would explain why the ^{228}Th values are decreasing with increasing electric field.

A linear fit was also attempted for $\sigma_{e_{Det}}^2$, with the results as:

$$\sigma_{e_{Det}}^2 = (0.0193 \pm 0.0031 * N_e)^2 \quad (6.14)$$

For this fit $\chi_{red}^2 = 4.56$, which means that it is a slight overfit to the data. As with $\sigma_{\gamma_{Det}}^2$ the $\sqrt{N_e}$ model fits slightly better but both are reasonable models to the data.

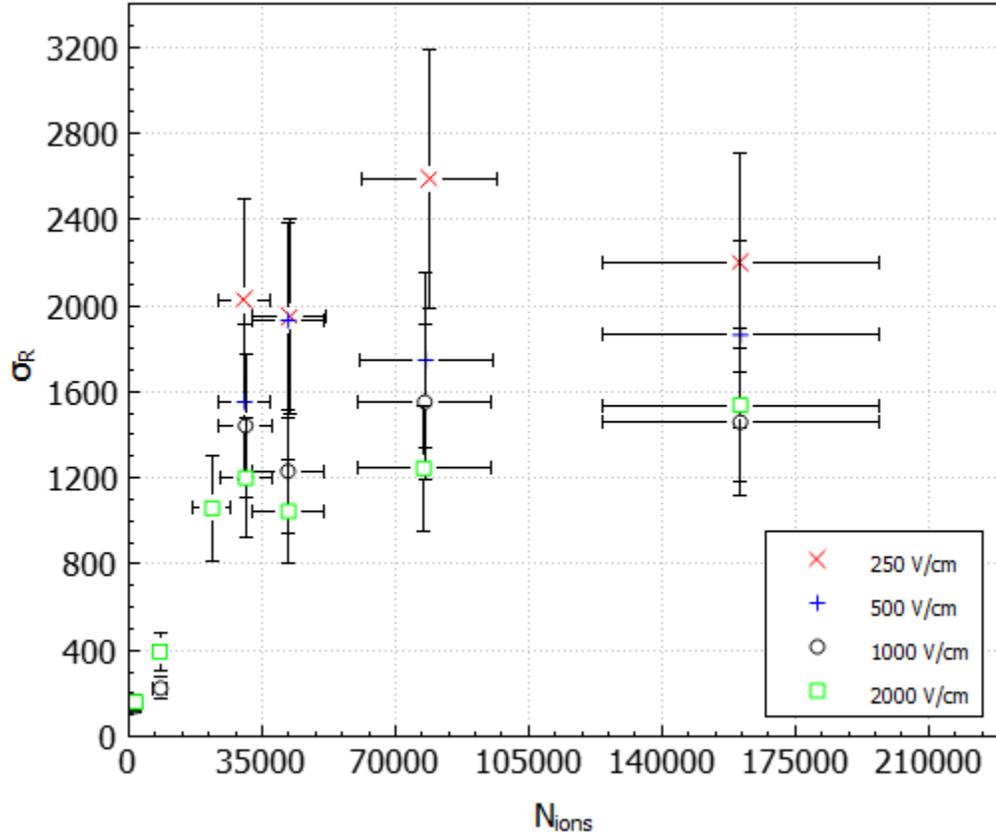


Figure 6.15: Recombination fluctuations as a function of ion quanta. Recombination fluctuations increase with number of quanta and then flatten out at higher quanta. The plateau appears to have a field dependence and flatten out at values of $2.36 \cdot 10^4$, $1.83 \cdot 10^4$, $1.54 \cdot 10^4$, and $1.39 \cdot 10^4$ for 250, 500, 1000, and 2000 V/cm respectively.

Going back to Eqn. (6.7) the recombination fluctuations can be calculated. This quantity offers insights about xenon detector physics, Figure 6.13 is more about systematic diagnostic tools. For this, the recombination fluctuations are plotted as a function of ion quanta, which is defined as:

$$N_{ions} = \frac{N_q}{1 + \alpha} \quad (6.15)$$

For this work, we assume that $\alpha = 0.2$ and is constant, which means that N_{ions} is simply proportional to N_q .

The recombination fluctuations appear to level off at higher ion quanta with the value of the plateau being field dependent. These values are $2.36 \cdot 10^4$, $1.83 \cdot 10^4$, $1.54 \cdot 10^4$, and $1.39 \cdot 10^4$ for 250, 500, 1000, and 2000 V/cm respectively. Previous work has shown that recombination fluctuations are linear with ion quanta [62]. This appears to be in mild agreement since our results are at higher ion quanta. For ion quanta less than 40,000 our data roughly observes a linear relationship, however further values appear to level off. This matches the mean recombination probability which was also leveling off with the plateau being field dependent as shown in Figure 6.15.

Since the behaviors of mean recombination probability and recombination fluctuations are similar they were combined and compared to total ion quanta, which is shown in Figure 6.16.

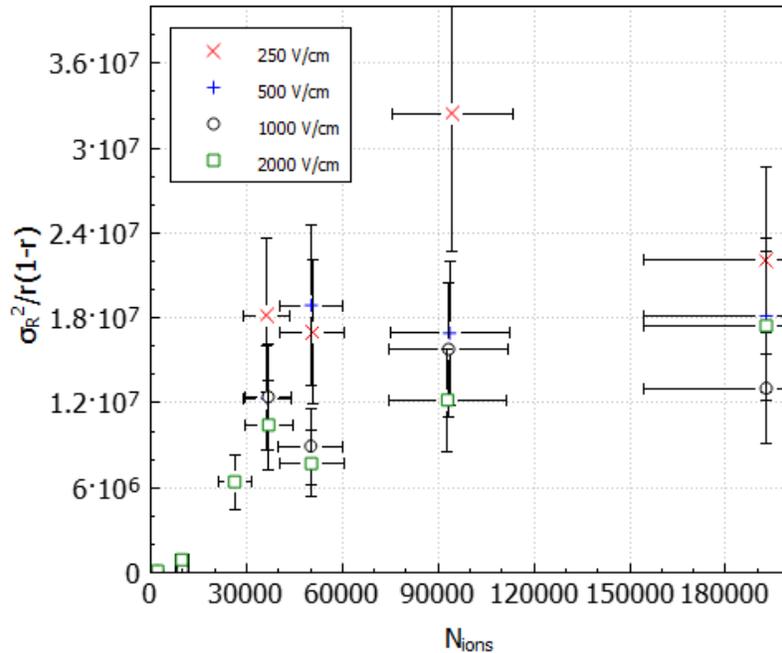


Figure 6.16: The ratio of recombination variance to the binomial square of the mean recombination probability as a function of N_{ions} .

The data points were modeled as:

$$\frac{\sigma_R^2}{r(1-r)} = \frac{C}{\sqrt{d_{field}}} * (1 - e^{-\left(\frac{N_{ions}}{N_{\tau}}\right)^2}) \quad (6.16)$$

Where d_{field} is the drift field in the detector in V/cm. The values of the modeled constants are $N_{\tau} = 3.24 \pm 0.82 * 10^4$ and $C = 4.21 \pm 1.26 * 10^7$, using Figure 6.11 this corresponds to an energy of $E_{\tau} = 527 \pm 137$ keV. In terms of performance, $\chi_{red}^2 = 8.09$ which is a underfit to the data, most likely due to the fact that there were only 4-6 data points per drift field.

The recombination fluctuations follow a binomial variance with an extra prefactor that is a function of N_{ions} . In the limit of high N_{ions} , the recombination fluctuations are proportional to a binomial variance with a constant value of $\frac{C}{\sqrt{d_{field}}}$. In the limit of low N_{ions} , the function simplifies to a binomial variance of N_{ions}^2 with an extra amplification factor 'A' that is equal to $\frac{C}{\sqrt{d_{field} * N_{\tau}^2}}$. Previous work calculated this value as 0.022 for 170 V/cm drift field [62], which is 0.031 ± 0.012 for this work.

This implies that recombination fluctuations follow a binomial variance with an energy-dependent ion function. This could mean that as the event gets longer for larger energies the recombination fluctuations are stabilizing due to the lack of recombination from far ends of the track. Studies have been performed that measure recombination saturation due to increasing drift field [75] [76], this is the first study to measure recombination saturation due to increasing energy.

7 PIXeYCI – A Liquid Xenon Compton Imager

7.1 Introduction

PIXeY has demonstrated both the method and the hardware means in which one can achieve good energy resolution. With this detector the intrinsic energy resolution of liquid xenon has been studied, and limiting factors of energy resolution have been proposed. In addition, PIXeY was used to study electron extraction efficiency, gamma/neutron discrimination, and low energy thresholds with ^{37}Ar . Now that those studies have concluded, PIXeY will be used to study Compton imaging as described in Section 1.2.3. This will require a complete redesign of the detector internals, mainly for new hardware to capture the xy position of each event.

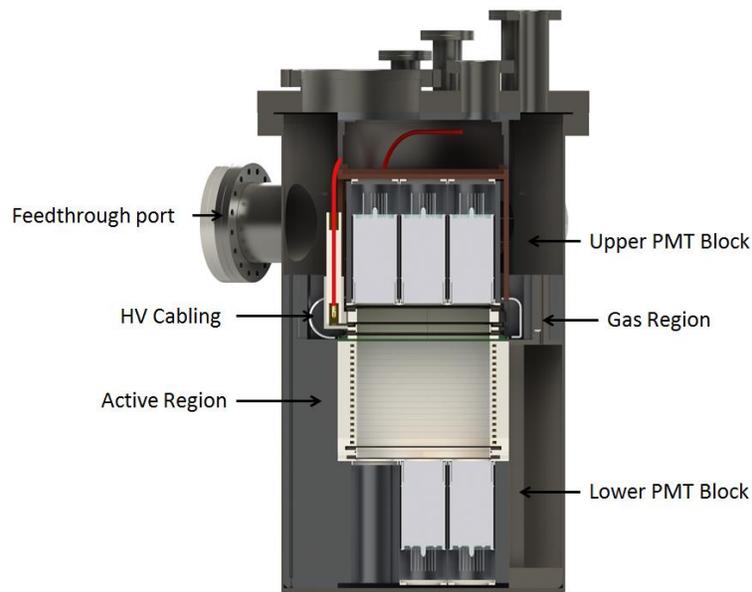


Figure 7.1: Schematic Concept of the PIXeYCI detector being enclosed in its inner can. The detector has a similar overall structure of top block, active region, and middle block. Also visible is the cabling for the anode voltage (shown in red), which follows a similar encapsulated design as PIXeY.

PIXeYCI (Particle Identification in Xenon at Yale Compton Imager) has a similar 3 body structure as PIXEY, but is an array of 3x3 PMTs on top and bottom and has a longer drift region of about 13cm and an active region of 3.226 l. The geometric redesign is to facilitate more Compton scatter events in the active region. PIXeYCI will also be a dual phase detector with charge detection via proportional scintillation. In addition it will also have two wire readout grids that will detect the charge drifting through the active region by way of image charges. A large portion of the redesign effort is to accommodate the hardware necessary for a dual planar xy wire readout system.

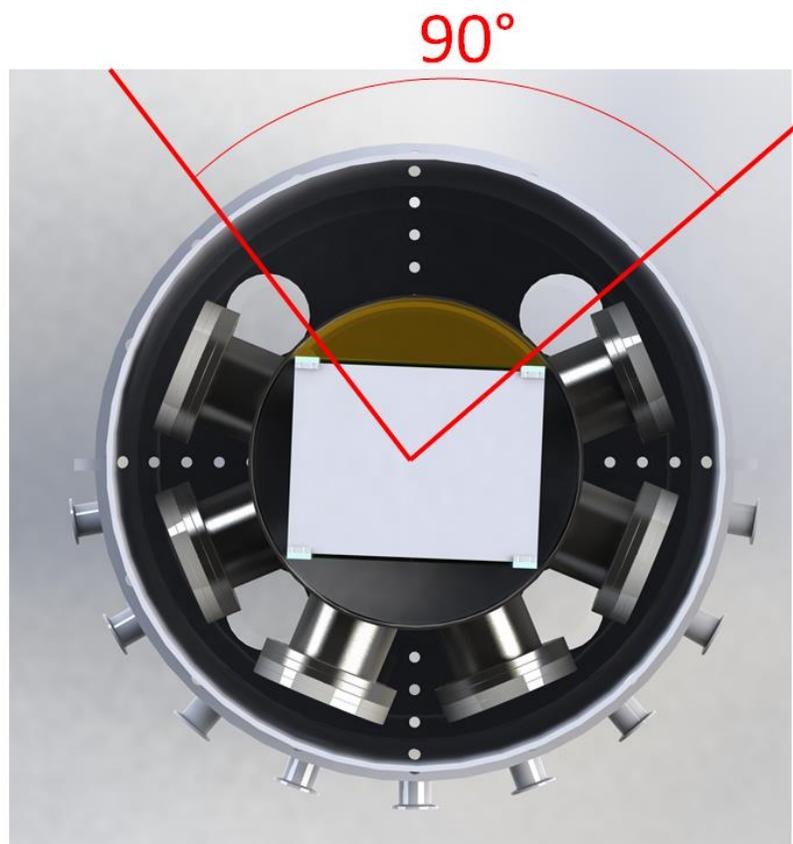


Figure 7.2: Overhead view of the PIXeYCI detector, which will have an approximately 90° viewing window for source study.

PIXeYCI's hardware configuration was designed in part to reduce incomplete scatter backgrounds within the detector, which was one of the largest background sources in PIXeY. These incomplete scatter

events deposit some of their energy outside of the active region, and some within the active region. Only events that deposit all of their energy in the active region are useful. The incomplete scatter background is shown in Figure 6.4, which is the large band before each energy peak. For external sources this background is unavoidable, since the gamma-rays are scattering off of all of the external hardware that is necessary to run the detector. To help avoid this vacuum flange ports are arranged on one side with an approximately 90° viewing window for source study, as shown in Figure 7.2. Essentially, if there is less material in the mean free path of the external source gamma-ray, then there is less potential for scattering before penetrating the active field region.

7.2 Wire Readout

7.2.1 Introduction

One of the challenges in getting a Compton imager to work for any detector is the accurate position reconstruction that is needed in order to do image reconstruction. Through advanced algorithms, many groups claim mm-level resolution with the use of only the PMTs [77], but those algorithms rely on a relatively large number of PMTs and symmetric assumptions that cannot be made for a small detector like PIXeYCI.

Instead, PIXeYCI will use an orthogonal array of charge sensitive wires. These wires will detect the charge as it passes by the wires, creating an image charge on the wire that is amplified by a charge sensitive amplifier. The idea was initially used in the LXeGrit detector mentioned in Section 1.2.3, and PIXeYCI will use a similar design. So in addition to the PMT information there will also be two planes of charge sensitive wires that will be read out in coincidence.

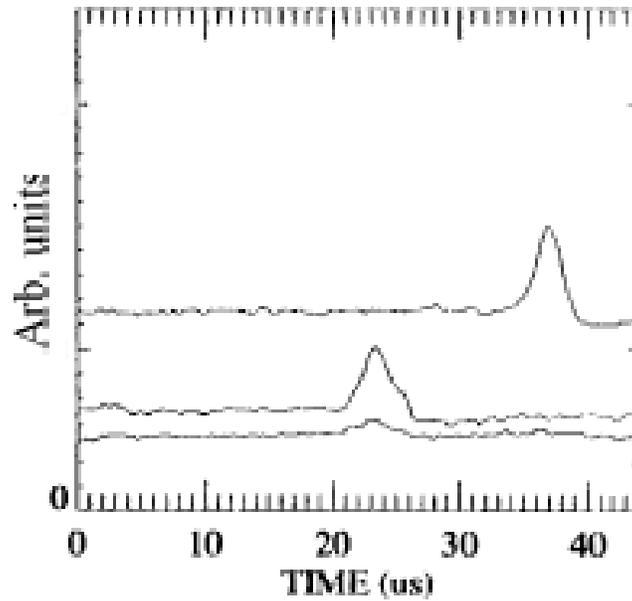


Figure 7.3: Figure: Example pulses from the LXeGrit detector [78]. The pulses indicate areas where charge passed by. Sometimes there may be pulses on adjacent wires, and it's possible to interpolate between those two wires to improve the position resolution.

As shown in Figure 7.3, a typical pulse from the wire readout system will have a FWHM of about $2.5 \mu\text{s}$. These pulses will be matched with the S2 pulses from the PMT readout and matched in time to determine which S2 pulse goes with which wire readout pulse. Since the electron cloud will physically pass between two sets of wires, there should be two matched pulses that indicate their detection. In Figure 7.3 it's clear that there are two pulses on adjacent wires on the lower center wires. By doing a weighted sum it's possible to get ~ 1.5 - 1 mm position resolution with a wire pitch of 3 mm , with the possibility of even better position resolution with more advanced algorithm techniques such as maximum likelihood and closest approach.

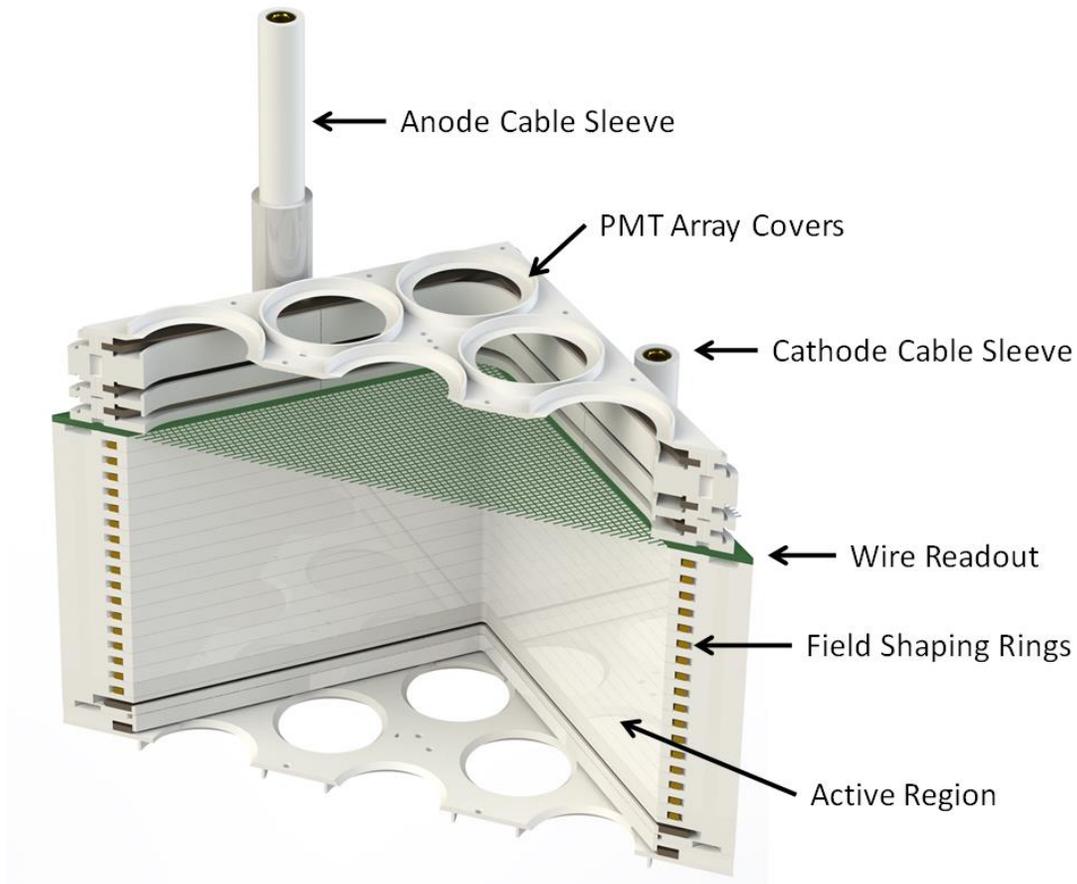


Figure 7.4: Illustration of the wire readout system and how it fits in with the rest of the detector. The wire readout system will be just below the gas region and detect the charge as it passes by the wires.

As shown in Figure 7.4, PIXeYCI's geometry is a 15x15x13 cm cube, and the pitch of the wires for the readout will be 3mm. This means that there will be 50 wires per array, so 100 wires in total. Each of those wires will require its own charge sensitive amplifier, be digitized, and require power. Thus there will be significant hardware engineering challenges in fitting and routing 100 signal and power lines into and out of the detector. Cross-talk and pickup noise will certainly be an issue in the system, and those issues will be diagnosed and remedied during construction.

7.2.2 Amplifier Design

One of the ways to avoid cross-talk and unwanted pickup noise is to have at least part of the amplification to happen as close to the wires as possible. Before the amplifier each wire can be thought of as a little radio antenna and so one would ideally want to make that antenna as small as possible. This can be accomplished by having an amplification step near the wires, and thus somewhere in the xenon space. Such an amplifier must operate in a liquid or gas xenon environment, work at temperatures of 175K, have low noise, and low cross-talk between the other amplifiers. Based on simulations done in Maxwell, the amount of image charge that actually gets to the wire could be as little as 10%, which would be if the electron cloud were exactly in the middle of two wires. Of course, it also means that a roughly equal 10% would be on the other wire as well. Nonetheless this means that proper, reduced noise amplifiers need to be designed and fabricated.

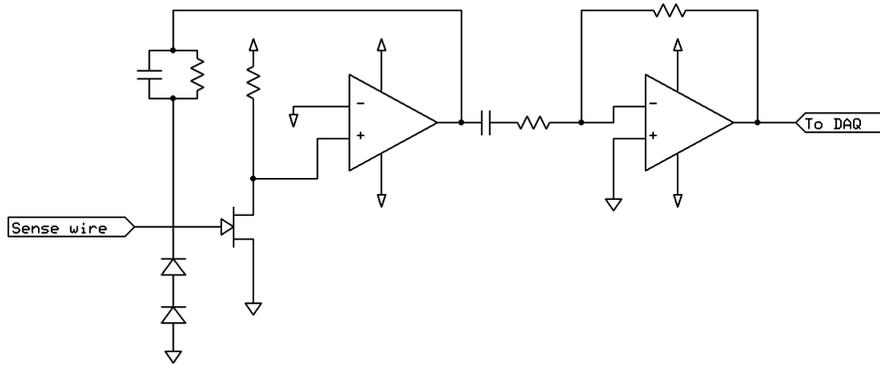


Figure 7.5: Figure: Two-stage amplification idea based on GERDA's conceptual design [79]. There is a JFET stage of amplification and a feedback stage that acts as a current amplifier. The second stage is voltage amplification and can be outside of the xenon environment. Also visible are diodes that protect the DAQ from high reflective voltages.

The amplification for the charge sensitive wires will be done in three stages. The first stage is a charge amplification using a BF862 JFET produced by Phillips with a power bias on the source. This will be the main entry point to the feedback loop which is modulated by a 0.3 pF capacitor. The amplifier is an

AD8651, an all CMOS amplifier that can work at liquid nitrogen temperatures (77K). Thus the amplifier with the JFET in feedback will produce an output voltage that is proportional to the input charge. The second gain stage will be voltage amplification with a gain around 10, which will preserve the shape of the amplifier pulse. It is desired to retain the shape of the pulse for advanced algorithms to determine the electron cloud location.

7.2.3 Test Boards

Boards have been designed and constructed that will test the amplifier concept, and specifically to tweak the values of the various input capacitors and resistors to check for performance, mainly concerning with stability. As with feedback amplifiers the stability of the amplification is important to test because errors in the manufacturing process that cause deviations in the reported values of resistance and capacitance or thermal cycling and cooling could cause oscillations.

The boards themselves are multilayered precision etched boards that were designed with a proprietary program called PCB Artist [80]. It has a top masking layer to shield conductive lines, then a conductive layer that has common footprints for various parts and also acts as a ground plane. Next is the resistive wafer itself which acts as the main structure of the board. Finally the whole layering system is repeated on the backside so components or connections can be placed on the back. The two ground planes are connected via thru holes.

The amplifier needs to have a small form factor because of the number of amplifiers needed as well as the space that can fit them. Even in the xenon gas, space will be limited due to the design constraints of the inner can. Nonetheless it should be possible to fit all of the electronic components in as little as 2x4 cm. As shown in Figure 7.6, the components on the left are to prepare the signal coming from the sense wires, and the right set of components are all stages of amplification. The first pin is the input preparation; the second pin is the JFET bias line with various buffering resistors and capacitors. The fourth pin is a test pin used to test the amplifier with a function generator. The fifth pin is the power line

to the amplifier and the eighth pin is the signal output pin. The second, sixth, and seventh pins were not used.

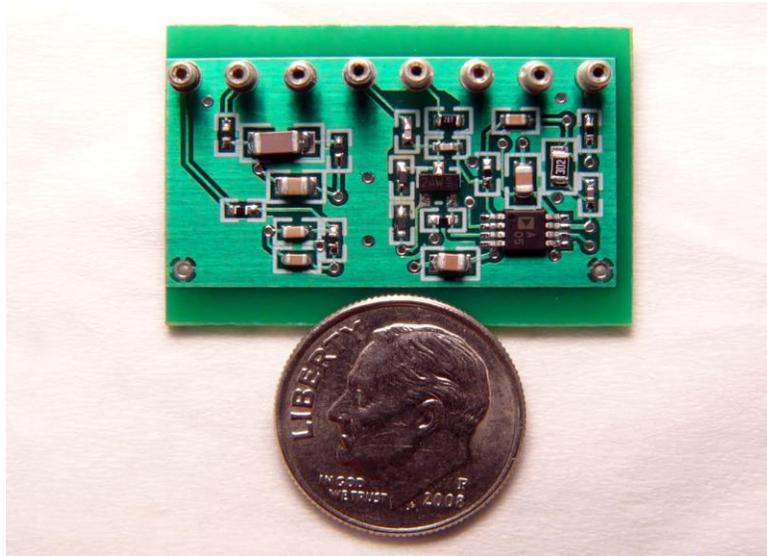


Figure 7.6: Image of an early design of the charge sensitive amplifier with a dime for scale. The parts were a mix of 1mm, 3mm, and 5mm size standards and were manually placed with tweezers and a magnifying glass. The components were held in place with a sticky liquid solder flux paste and were soldered with a compact precision hot air gun.

7.2.4 Noise Analysis

A successful working amplifier must be able to distinguish its signal from the noise, which ideally should be minimized by design. Some amount of noise will be in the system however and it is critical to quantify this, since a system with too much noise relative to its signal would not be useful. With this in mind a noise analysis was done on each of the components assuming bias current sources at the resistors and current and voltage noise sources at the amplification stages.

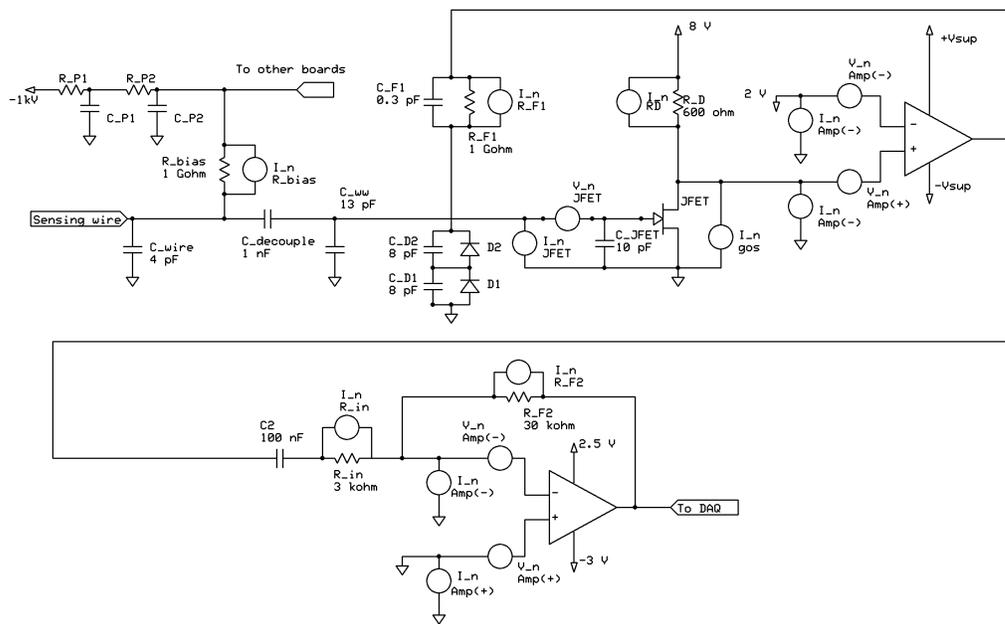


Figure 7.7: Annotated view of the current and voltage noise sources in the amplifier. Additional components, such as the ones that decouple the power lines and the other wires, needed to be modeled for a full noise analysis.

As shown in Figure 7.7, all components of the amplifier must be studied, including the various resistors and capacitors that buffer the power units to the amplifiers and JFET and decouple the signal and power lines from other wires. Each resistor has a bias current noise, and the amplifiers have a bias current noise on the inputs. There are capacitive noises on the diodes, the JFET, and the various wire components themselves. Finally, there are voltage bias noises on the amplifiers and on the JFET.

Noise source	Charge noise density referred to the input
Voltage noise at the JFET gate	$2.48 \times 10^{-20} \text{C}/\sqrt{\text{Hz}}$
Current noise at the JFET gate	$8.00 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Johnson noise of R_{bias}	$2.48 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Johnson noise of R_{F1}	$2.48 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Johnson noise of R_{D}	$3.31 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Johnson noise of g_{os}^{-1}	$4.27 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Total of noise sources excluding op-amps	$2.68 \times 10^{-20} \text{C}/\sqrt{\text{Hz}}$
Total noise from all CMOS op-amp inputs	$1.42 \times 10^{-20} \text{C}/\sqrt{\text{Hz}}$
Total noise from all BJT op-amp inputs	$8.00 \times 10^{-21} \text{C}/\sqrt{\text{Hz}}$
Total noise using CMOS op-amps	$3.03 \times 10^{-20} \text{C}/\sqrt{\text{Hz}}$
Total noise using BJT op-amps	$2.80 \times 10^{-20} \text{C}/\sqrt{\text{Hz}}$

Table 7.1: Values of noise sources for different components of the amplifier and comparing different types of amplifiers.

Taking everything into account the projected noise for this amplifier would be between 2.8-3.03 $\times 10^{-20} \text{C}/\sqrt{\text{Hz}}$. This is based on a summation on all of the line elements in Table 7.1, assuming the usage of a BJT op-amp instead of a CMOS op-amp. Since the pulse regions for the wire readout system are on the order of 2.5 μs at FWHM this means that our bandwidth region of interest is about 100 Hz – 100 kHz. Thus for this bandwidth region of interest this would be an integrated charge noise of about $1.92 \times 10^{17} \text{C}$, or 120 electrons.

Finally, with all of these components the power requirement would be 120 mW with JFET biasing and op-amp biasing of 5 V. For the entire system this would be 14.4 W of power, which if converted into heat would add 0.32 L-LN2/hr to the heat load of the system.

7.2.5 Simulated Gain Calculation

The gain at the first stage will be primarily dominated by the feedback capacitor which is 0.3 pF. One way to understand the gain of the amplifier is to calculate the signal response for the largest event

that would be encountered. Thus, the upper bound on the gain of the amplifier when taking into account the xenon physics information was studied.

For the simulation the largest gamma-ray of interest was chosen, ^{228}Th 2.6 MeV, and assumed a photoabsorption event. This puts all of its energy in one localized cloud of charge which has an estimated size of 100 μm [81]. A 2.6 MeV gamma-ray would make approximately 170,000 electrons, based on the work function of liquid xenon [82] and assuming that there will be a 50% light split at 1 kV/cm drift field [69]. Simulations were done in Maxwell with this estimated information of charge to see the type of image that is induced on a charge sensitive wire.

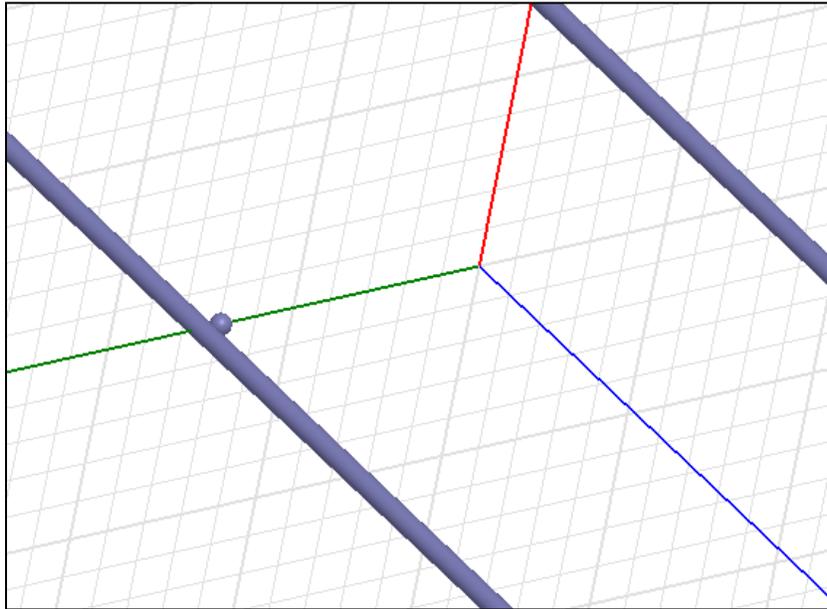


Figure 7.8: Maxwell simulation of a ball of charge drifting past a sense wire. The induced charge on the close wire was compared to the total charge of the ball.

The results of the simulation found that on average 80% of the charge in the ball is induced on the wires. If a 0.3 pF capacitor is used for the first gain stage, a gain of 28 for that stage, and finally a gain of 10 for the second stage this would make an overall gain of $3.33 \cdot 10^{13}$ V/C or 5.33 $\mu\text{V}/\text{electron}$. This

means that the maximum potential difference that a single electron would make on the output of the amplifiers is $5.33 \mu\text{V}$.

7.3 Projected Performance

Putting together the energy resolution results from current PIXeY, the simulations and expectations of the charge readout information and the geometry of the PIXeYCI simulations were done to determine the performance of the detector. Namely, based on given information of projected energy, wire readout, and drift time resolution the resolution of the Compton scattering angle was determined. A monoenergetic 662 keV gamma-ray point source was simulated in addition to background events at random locations within the search region from energies between 50 – 700 keV, which are at random locations in the event search view.

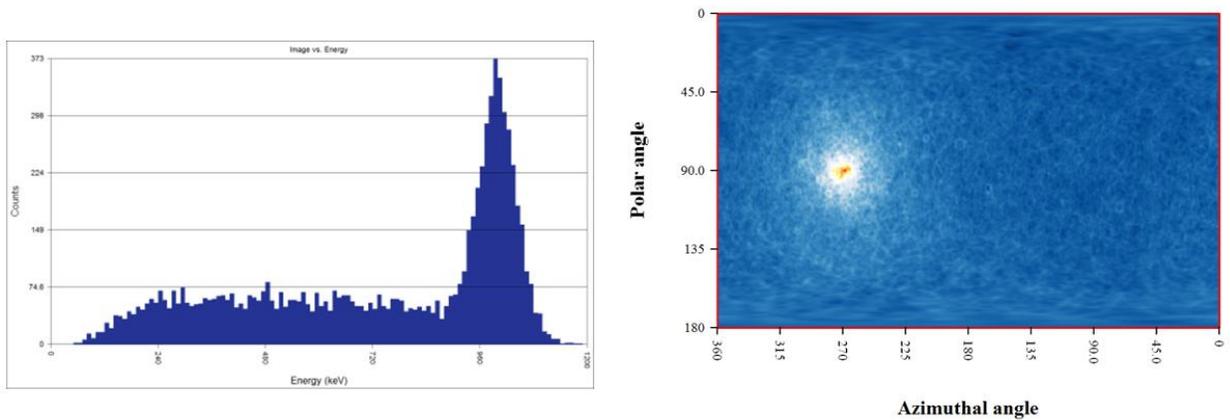


Figure 7.9: Projected performance of PIXeYCI with a monoenergetic point source of 662 keV assuming xy position resolution of 1.5mm and z resolution of 1mm. The projected performance on the Compton scattering angle is 13.5° FWHM.

Based on an S2 energy resolution of $4\% \sigma/E$, xy position resolution of 1.5 mm and z position resolution of 1mm, the resulting Compton scattering angle resolution is 13.5° FWHM. With these studies

it shows that this will be an exciting chapter in liquid xenon physics. The viability of such a material as a cheap, cost effective, and still good resolution material for Compton imaging technologies will be tested. The combination of competitive metrics with CZT and Ge with the bulk properties of liquid xenon mean that this detector material is sure to be used in a wide variety of new applications for the years to come.

Appendix

A.1. Theoretical Photon Yield Calculation

The following equation describes the photon yield in gaseous xenon as a function of electric field intensity, gas pressure, and grid distance [48]:

$$N_{ph} = 70 \left(\frac{E}{p} - 1.0 \right) xp \quad (\text{A.1})$$

where E is in kV/cm, p is in bar, and x is in cm. This equation accounts for motion in 1 dimension and a perfectly uniform electric field. What is needed then is an equation that takes into account variations in the electric field as well as 3D paths. I will start the analysis with a 1D case (the particles are constrained in one direction but can experience a varying field) since it's easier and then extrapolate to the 3D case.

A.1.1. One dimensional derivation

Suppose an electron is traveling in one dimension and experiences an electric field that varies with distance. Along the line that the electron travels, we can imagine cutting that path up into tiny segments, where the field is roughly constant in that segment. To find the total photon yield then, we simply add up all the contributions due to each tiny segment. Now suppose that there are just two segments that are simply connected to each other. The following equations would apply for each segment:

$$N_{ph}^{(1)} = 70p \left(\frac{E^{(1)}}{p} - 1.0 \right) (x_2 - x_1) \quad (\text{A.2})$$

$$N_{ph}^{(2)} = 70p \left(\frac{E^{(2)}}{p} - 1.0 \right) (x_3 - x_2) \quad (\text{A.3})$$

Here the electric field is constant within each segment but varies between segments. Now, without loss of generality we can assume that the total distance of each segment is the same, more formally:

$$\Delta x \equiv (x_2 - x_1) = (x_3 - x_2) \quad (\text{A.4})$$

Then:

$$N_{ph}^{(1)} = 70p \left(\frac{E^{(1)}}{p} - 1.0 \right) \Delta x \quad (\text{A.5})$$

$$N_{ph}^{(2)} = 70p \left(\frac{E^{(2)}}{p} - 1.0 \right) \Delta x \quad (\text{A.6})$$

To find the total photon yield due to each segment, we simply add the two equations together to get:

$$N_{ph}^{(1)} + N_{ph}^{(2)} = 70p \left(\frac{E^{(1)} + E^{(2)}}{p} - 2.0 \right) \Delta x \quad (\text{A.7})$$

From here it is easy to extrapolate to an arbitrary number of segments:

$$\sum_{i=1}^m N_{ph}^{(i)} = 70p \left(\frac{\sum_{i=1}^m E^{(i)} \Delta x}{p} - 1.0 * m \Delta x \right) \quad (\text{A.8})$$

which can be rewritten as:

$$\sum_{i=1}^m N_{ph}^{(i)} = 70 \left(\sum_{i=1}^m E^{(i)} \Delta x - 1.0 * p \sum_{i=1}^m \Delta x \right) \quad (\text{A.9})$$

To make this equation exact, we simply let the segment size go to zero and the number of segments go to infinity. Thus:

$$\lim_{m \rightarrow \infty} \lim_{\Delta x \rightarrow 0} \left\{ \sum_{i=1}^m N_{ph}^{(i)} = 70 \left(\sum_{i=1}^m E^{(i)} \Delta x - 1.0 * p \sum_{i=1}^m \Delta x \right) \right\}$$

$\xrightarrow{\text{yields}}$

(A.10)

$$N_{ph} = 70 \left[\int E(x) dx - 1.0 p * \int dx \right]$$

and evaluating these integrals one gets:

$$N_{ph} = 70[\Delta V - 1.0 * pd]$$
(A.11)

where ΔV is the total change in potential from the beginning point to the end and d is the total distance traveled by the electron.

A.1.2. Generalization to Three dimensions

We will follow the same analysis as in the previous section but here we will assume that the particle can take a path in three dimensions. But still in each segment the field is constant. Equation (A.5) and (A.6) becomes:

$$N_{ph}^{(1)} = 70p \left(\frac{\vec{E}^{(1)} \cdot \hat{s}^{(1)}}{p} - 1.0 \right) |\vec{x}_2 - \vec{x}_1|$$

$$N_{ph}^{(2)} = 70p \left(\frac{\vec{E}^{(2)} \cdot \hat{s}^{(2)}}{p} - 1.0 \right) |\vec{x}_3 - \vec{x}_2|$$
(A.12)

where the following replacements have been made:

$$\begin{aligned}
(x_2 - x_1) &\underset{3D}{\Rightarrow} |\vec{x}_2 - \vec{x}_1| \\
(x_3 - x_2) &\underset{3D}{\Rightarrow} |\vec{x}_3 - \vec{x}_2| \\
E^{(1)} &\underset{3D}{\Rightarrow} \vec{E}^{(1)} \cdot \hat{s}^{(1)} \\
E^{(2)} &\underset{3D}{\Rightarrow} \vec{E}^{(2)} \cdot \hat{s}^{(2)}
\end{aligned} \tag{A.13}$$

The first two replacements are simply transforming a distance in one dimension into a distance in three dimensions. The last two replacements are accounting for the fact that the field that we're interested in is the field along the direction of motion (so then \hat{s} is a unit vector that points along the direction of motion).

At this point we assume that the electron follows a smooth path that is continuous and continuously differentiable (a “nice” function) and that the path can be parameterized as follows:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x \\ f(x) \\ g(x) \end{pmatrix} \tag{A.14}$$

so that Equation (A.12) becomes:

$$\begin{aligned}
N_{ph}^{(1)} &= 70p \left(\frac{\vec{E}^{(1)} \cdot \hat{s}^{(1)}}{p} - 1.0 \right) \{ [x_2 - x_1]^2 + [f(x_2) - f(x_1)]^2 + [g(x_2) - g(x_1)]^2 \}^{1/2} \\
N_{ph}^{(2)} &= 70p \left(\frac{\vec{E}^{(2)} \cdot \hat{s}^{(2)}}{p} - 1.0 \right) \{ [x_3 - x_2]^2 + [f(x_3) - f(x_2)]^2 + [g(x_3) - g(x_2)]^2 \}^{1/2}
\end{aligned} \tag{A.15}$$

Once again without any loss of generality we can assume:

$$\Delta x \equiv (x_2 - x_1) = (x_3 - x_2) \tag{A.16}$$

so that Equation (A.15) becomes:

$$\begin{aligned}
N_{ph}^{(1)} &= 70p \left(\frac{\vec{E}^{(1)} \cdot \hat{s}^{(1)}}{p} - 1.0 \right) \left\{ 1 + \left[\frac{f(x_2) - f(x_1)}{\Delta x} \right]^2 + \left[\frac{g(x_2) - g(x_1)}{\Delta x} \right]^2 \right\}^{1/2} \Delta x \\
N_{ph}^{(2)} &= 70p \left(\frac{\vec{E}^{(2)} \cdot \hat{s}^{(2)}}{p} - 1.0 \right) \left\{ 1 + \left[\frac{f(x_3) - f(x_2)}{\Delta x} \right]^2 + \left[\frac{g(x_3) - g(x_2)}{\Delta x} \right]^2 \right\}^{1/2} \Delta x
\end{aligned} \tag{A.17}$$

Now the terms in the curly brackets are measurements of discrete slopes. Since these points are close to each other we can approximate that these slope values should be equal. And in the limit that Δx goes to zero this approximation becomes exact. Redefining the term in the curly brackets:

$$\begin{aligned}
\Delta s &\equiv \left\{ 1 + \left[\frac{f(x_2) - f(x_1)}{\Delta x} \right]^2 + \left[\frac{g(x_2) - g(x_1)}{\Delta x} \right]^2 \right\}^{1/2} \Delta x \\
&\cong \left\{ 1 + \left[\frac{f(x_3) - f(x_2)}{\Delta x} \right]^2 + \left[\frac{g(x_3) - g(x_2)}{\Delta x} \right]^2 \right\}^{1/2} \Delta x
\end{aligned} \tag{A.18}$$

Then Equation (A.17) becomes:

$$\begin{aligned}
N_{ph}^{(1)} &= 70p \left(\frac{\vec{E}^{(1)} \cdot \hat{s}^{(1)}}{p} - 1.0 \right) \Delta s \\
N_{ph}^{(2)} &= 70p \left(\frac{\vec{E}^{(2)} \cdot \hat{s}^{(2)}}{p} - 1.0 \right) \Delta s
\end{aligned} \tag{A.19}$$

These equations can be summed and generalized in the usual way:

$$\sum_{i=1}^m N_{ph}^{(i)} = 70 \left(\sum_{i=1}^m \vec{E}^{(i)} \cdot \hat{s}^{(i)} \Delta s - 1.0 * p \sum_{i=1}^m \Delta s \right) \tag{A.20}$$

and in the limit that Δs goes to zero and m goes to infinity:

$$\lim_{m \rightarrow \infty} \lim_{\Delta s \rightarrow 0} \left\{ \sum_{i=1}^m N_{ph}^{(i)} = 70 \left(\sum_{i=1}^m \vec{E}^{(i)} \cdot \hat{s}^{(i)} \Delta s - 1.0 * p \sum_{i=1}^m \Delta s \right) \right\}$$

$\xrightarrow{\text{yields}}$

$$N_{ph} = 70 \left[\int \vec{E}(\vec{x}) \cdot d\vec{s} - 1.0 * p \int ds \right]$$

(A.21)

which reduces to the familiar result:

$$N_{ph} = 70[\Delta V - 1.0 * ps]$$

(A.22)

with the only difference being that d has been replaced with s where s is the total length of the path that the electron takes in three dimensions.

A.2 Determination of the Electric Field in the Gas Gap

The liquid and gas will be modeled as capacitors in series and the capacitance in the gas region will be used to find the field.

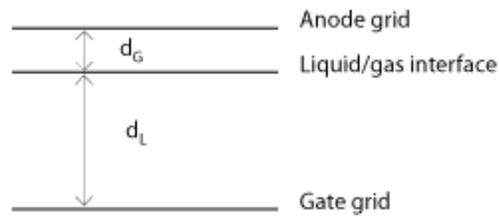


Figure A.1: Illustration of the liquid/gas interface with gas gap d_G and liquid gap d_L .

The capacitance of a parallel-plate capacitor is:

$$C = \epsilon_0 \epsilon_r \frac{A}{d} \quad (\text{A.23})$$

where C is the capacitance in Farads, ϵ_0 is the permittivity of free space, ϵ_r is the relative static permittivity, A is the area of the plate and d is the distance between plates. Since these are capacitors in series then the total adds in parallel, thus the total capacitance is:

$$C_{Total} = \left(\frac{1}{C_{gas}} + \frac{1}{C_{liquid}} \right)^{-1} \quad (\text{A.24})$$

Using the equation for a parallel-plate capacitor and labeling terms we get:

$$C_{Total} = \left(\frac{d_{gas}}{\epsilon_0 \epsilon_{gas} A} + \frac{d_{total} - d_{gas}}{\epsilon_0 \epsilon_{liquid} A} \right)^{-1} \quad (\text{A.25})$$

and after rearranging terms we get:

$$C_{Total} = \frac{\epsilon_0 \epsilon_{gas} \epsilon_{liquid} A}{\epsilon_{gas} d_{total} + (\epsilon_{liquid} - \epsilon_{gas}) d_{gas}} \quad (\text{A.26})$$

The total capacitance is related to the total charge and the total voltage drop. Since these are modeled as capacitors in series we assume that the total voltage drop is equal to the sum of the voltage drop in each section and that the total charge is the same on each section. This assumes that the field-defining charge has reached a steady-state condition:

$$\begin{aligned} V_{Total} &= V_{gas} + V_{liquid} \\ Q_{Total} &= Q_{gas} = Q_{liquid} \end{aligned} \quad (\text{A.27})$$

If the electric field in the gas is assumed to be a parallel-plate capacitor, then we can find the field based on the capacitance:

$$E_{gas} = \frac{V_{gas}}{d_{gas}} = \frac{Q}{C_{gas}d_{gas}} = \frac{Q}{\epsilon_0\epsilon_{gas}A} \quad (A.28)$$

The charge can be found from the total capacitance:

$$Q = V_{Total} * C_{Total} \quad (A.29)$$

so substituting this equation into equation (A.28) and equation (A.26) we get:

$$E_{gas} = \frac{V_{Total}}{\epsilon_0\epsilon_{gas}A} \left(\frac{\epsilon_0\epsilon_{gas}\epsilon_{liquid}A}{\epsilon_{gas}d_{total} + (\epsilon_{liquid} - \epsilon_{gas})d_{gas}} \right) \quad (A.30)$$

and finally after simplifying terms:

$$E_{gas} = \frac{\epsilon_{liquid}V_{Total}}{\epsilon_{gas}d_{total} + (\epsilon_{liquid} - \epsilon_{gas})d_{gas}} \quad (A.31)$$

A.3. Estimating the Electric Field on the Cathode Wire

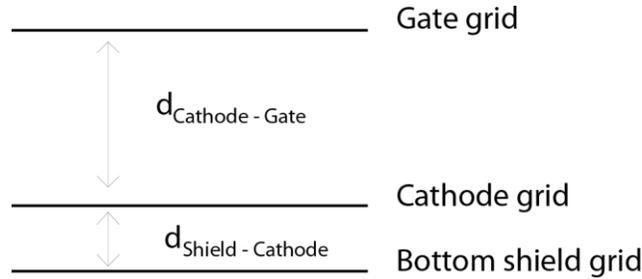


Figure A.2: The schematic of the grid configuration with Cathode Gate grid separation $d_{Cathode-Gate}$ and Cathode Shield grid separation $d_{Shield-Cathode}$.

First the charge on the surface of the entire cathode needs to be determined. Approximating all three surfaces as infinitely long thin sheets of charge the boundary condition for the middle layer becomes:

$$\frac{\sigma}{\epsilon_0} = E_{above} - E_{below} \quad (\text{A.32})$$

but these are just simple planar fields with $E = V/d$. So taking the proper voltage differences to get the fields:

$$\frac{\sigma}{\epsilon_0} = \frac{V_{Cathode} - V_{Gate}}{d_{Cathode-Gate}} - \frac{V_{Shield} - V_{Cathode}}{d_{Shield-Cathode}} \quad (\text{A.33})$$

where σ is the total charge per unit area on the cathode. To get the charge per unit length on the wire, divide the surface charge density by the wire pitch (this assumes that all charge is contained on the wire):

$$\lambda = \sigma \backslash w \quad (\text{A.34})$$

Where w is the pitch in units of *length/wire*. Thus the charge line density is:

$$\lambda = w * \epsilon_0 \left(\frac{V_{Cathode} - V_{Gate}}{d_{Cathode-Gate}} - \frac{V_{Shield} - V_{Cathode}}{d_{Shield-Cathode}} \right) \quad (\text{A.35})$$

Now that we have the charge density, we can use the electric field equation of an infinitely long wire to get the electric field at the wire. The electric field of an infinitely long wire is:

$$E_{wire}(r) = \frac{1}{2\pi\epsilon_0} \frac{\lambda}{r} \quad (\text{A.36})$$

Let the wire diameter be a . We are interested in the field at a :

$$E_{wire} = \frac{1}{2\pi\epsilon_0} \frac{\lambda}{a} \quad (\text{A.37})$$

Finally, substituting in the equation of λ and cancelling terms we get:

$$E_{wire} = \frac{w}{2\pi a} \left(\frac{V_{Cathode} - V_{Gate}}{d_{Cathode-Gate}} - \frac{V_{Shield} - V_{Cathode}}{d_{Shield-Cathode}} \right) \quad (A.38)$$

A.4. Uniqueness of Energy Resolution Minimization

In this section it will be demonstrated that the minimal energy resolution for a given system can be found by only considering a relative weighting factor between the PMTs, rather than the calculation of the value itself. To show this, let us consider a simple example of only 2 PMTs. The total signal would be:

$$S = PMT_1 + PMT_2 \quad (A.39)$$

Now suppose that each of these signals is off in some way from the correct values:

$$S = A_1 PMT_1 + A_2 PMT_2 \quad (A.40)$$

The values A_1 and A_2 represent some fixed erroneous weighting factor that can be any sort of systematic calibration error. The value $1-\sigma$ away from the curve will also be off by the same values:

$$S^{\sigma+} = A_1 PMT_1^{\sigma+} + A_2 PMT_2^{\sigma+} \quad (A.41)$$

The "energy resolution" of such a system would just be the ratio of these two values:

$$\frac{S^{\sigma+}}{S} = \frac{A_1 PMT_1^{\sigma+} + A_2 PMT_2^{\sigma+}}{A_1 PMT_1 + A_2 PMT_2} \quad (A.42)$$

For any non-zero non-negative values of A_1 and A_2 the ratio here will be larger than if they are simply unity. In practical terms this means that the energy resolution will be categorically worse for any additional incorrect weighting factor that is added to the sum of the PMT signal. This is proven in Appendix A.5. Now, suppose we add relative weighting factors to the two signals here:

$$\frac{S^{\sigma+}}{S} = \frac{wA_1PMT_1^{\sigma+} + (1-w)A_2PMT_2^{\sigma+}}{wA_1PMT_1 + (1-w)A_2PMT_2} \quad (\text{A.44})$$

with all possible values of w being from 0 to 1. Suppose we are able to find a w such that:

$$w = \frac{A_2}{A_1 + A_2} \quad (\text{A.45})$$

If we plug this into equation (A.44) we get:

$$\begin{aligned} \frac{S^{\sigma+}}{S} &= \frac{\frac{A_2}{A_1 + A_2}A_1PMT_1^{\sigma+} + (1 - \frac{A_2}{A_1 + A_2})A_2PMT_2^{\sigma+}}{\frac{A_2}{A_1 + A_2}A_1PMT_1 + (1 - \frac{A_2}{A_1 + A_2})A_2PMT_2} \\ &= \frac{\frac{A_1A_2}{A_1 + A_2} \frac{PMT_1^{\sigma+} + PMT_2^{\sigma+}}{PMT_1 + PMT_2}}{\frac{A_1A_2}{A_1 + A_2} \frac{PMT_1 + PMT_2}{PMT_1 + PMT_2}} = \frac{PMT_1^{\sigma+} + PMT_2^{\sigma+}}{PMT_1 + PMT_2} \end{aligned} \quad (\text{A.46})$$

Thus we were able to determine a weighting factor that minimizes the energy resolution without directly determining how off each PMT was in an absolute sense. Since the span of all weighting factors is between 0 and 1 the minimized resolution due to incorrect PMT weighting can be found by spanning the entire set. The advantage of this method is that it reduces the required number of parameters to find by one and since it is a bounded set the span of all values can be tested and such an optimization process is therefore only limited by the resolution of the optimization process.

A.5. Lemma for Energy Resolution Uniqueness

In this section it will be shown that changing the relative weighting of a PMT signal can only improve the energy resolution insofar as determining the "ideal" value found, changing the weighting of the PMTs cannot make the energy resolution any better than the intrinsic resolution.

Without any loss of generality, we can once again look at a simplified system of just 2 PMTs, with incorrect constant weighting factors added:

$$f(x, y|a, b, c, d) = \frac{ax + by}{cx + dy} \quad (\text{A.47})$$

Where x and y are the incorrect weighting factors and a , b , c , and d are the sigma and peak signals. Then the question becomes finding the values of these erroneous factors that minimize this ratio. Since we only have two parameters we want the minima of the two-dimensional space $f(x,y)$. Taking the derivatives:

$$\frac{\partial f}{\partial x} = \frac{(cx + dy)a - (ax + by)c}{(cx + dy)^2} \quad (\text{A.48})$$

$$\frac{\partial f}{\partial y} = \frac{(cx + dy)b - (ax + by)d}{(cx + dy)^2}$$

Set these derivatives to zero:

$$\frac{\partial f}{\partial x} = 0 \rightarrow (cx + dy)a - (ax + by)c = 0 \rightarrow acx - acx + y(ad - bc) = 0 \rightarrow y = 0 \quad (\text{A.49})$$

$$\frac{\partial f}{\partial y} = 0 \rightarrow (cx + dy)b - (ax + by)d = 0 \rightarrow bdy - bdy + x(bc - ad) = 0 \rightarrow x = 0$$

Thus we find that the minima for this system occurs at $x = 0, y = 0$. However the function is not well defined at these points. However we can conclude that since this ratio is less than 1 for all values of x,y then this function monotonically increases with increasing x,y . This means that we just have to find the smallest possible values without actually getting to zero.

As it turns out, the smallest possible values are when x,y are equal to each other, in that case the ratio reduces to:

$$\frac{a + b}{c + d} \quad (\text{A.50})$$

To show that this ratio is always smaller than Equation (A.47) we take two steps: first we show that it's true for all x,y values that are bigger than one, if that is true then:

$$\frac{a+b}{c+d} > \frac{ax+by}{cx+dy} \quad \forall x,y > 1 \quad (\text{A.51})$$

Is also true because by the derivative analysis above this function is monotonically increasing with x,y so any value of x,y is that is bigger than 1 will monotonically increase this function.

Next we simply show that this is also true for fractional values of x,y . Suppose x,y are fractions that can be written or approximately written as:

$$x = \frac{f_1}{f_2} \quad y = \frac{f_3}{f_4} \quad (\text{A.52})$$

where $f_1, f_2, f_3,$ and f_4 are all numbers that are larger than one. When we plug these in to the ratio we get:

$$\frac{a \frac{f_1}{f_2} + b \frac{f_3}{f_4}}{c \frac{f_1}{f_2} + d \frac{f_3}{f_4}} \quad (\text{A.53})$$

Simplifying terms:

$$\frac{f_2 f_4 a f_1 f_4 + b f_3 f_2}{f_2 f_4 c f_1 f_4 + d f_3 f_2} \quad (\text{A.54})$$

but $f_1, f_2, f_3,$ and f_4 are all numbers greater than 1 so therefore by the demonstration above this monotonically increases the function as well. Therefore, the only resolution that can be found by adding relative weights to the PMT signals is the intrinsic one and that value cannot be artificially lowered by adding weights to the PMTs.

Bibliography

- [1] Y. Ramachers et al., "Energy Resolution Improvement in Room-Temperature CZT Detectors," *JINST*, 2007.
- [2] E. Aprile and T. Doke, "Liquid Xenon Detectors for Particle Physics and Astrophysics," *Rev. Mod. Phys.*, pp. 82:2053-2097, 2009.
- [3] D. Akerib et al., "First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility," *Phys. Rev. Lett.* 112, 2014.
- [4] M. Auger et al., "Search for Neutrinoless Double-Beta Decay in ^{136}Xe with Exo-200," *Phys. Rev. Lett.*, p. 109, 2012.
- [5] A. Curioni et al., "A study of the LXeGrit Detection Efficiency for MeV Gamma-Rays During the 2000 Balloon Flight Campaign," *Nucl. Instr. Meth. A*, p. 576, 2007.
- [6] H. Zaklad et al., "A Liquid Xenon Radioisotope Camera," *IEEE Trans. Nucl. Sci.* 19, p. 206, 1972.
- [7] J. Nikkel et al., "Liquified Noble Gas (LNG) Detectors for Detection of Nuclear Materials," *JINST* 7 C03007, 2012.
- [8] D. Jordan et al., "Advanced Large-Area Plastic Scintillator Project (ALPS): Final Report," U.S. Department of Energy, 2007.
- [9] E. Siciliano et al., "Comparison of PVT and NaI(Tl) Scintillators for Vehicle Portal Monitor Applications," *Nucl. Inst. and Meth. in Phys. Res. A* 550, pp. 647-674, 2005.
- [10] C. M. Stahle et al., "Fabrication of CdZnTe Strip Detectors for Large Area Arrays," in *SPIE Conference Volume 3115*, 1997.
- [11] K. Ni, Development of a Liquid Xenon Time Projection Chamber for the XENON Dark Matter Search, New York, New York: Columbia University, 2006.
- [12] S. Kubota, M. Hishida and J. Raun, "Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton, and xenon," *J. Phys. C*, 11 (12), 1978.
- [13] J. Oort, "The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems," *BAIN*, p. 6:249, 1932.
- [14] F. Zwicky, "On the Masses of Nebulae and of Clusters of Nebulae," *API*, p. 86:217, 1937.
- [15] Q. Ahmad et al., "Direct evidence for neutrino flavor transformation from neutral current interactions

- in the Sudbury Neutrino Observatory," *Phys. Rev. Lett.* 89, 2002.
- [16] K. Abazajian and S. M. Koushiappas, "Constraints on sterile neutrino dark matter," *Phys. Rev. D* 74, 2006.
- [17] L. Bergstrom, "Non-baryonic dark matter: observational evidence and detection methods," *Rep. Prog. Phys.* 63, pp. 793-841, 200.
- [18] R. Bernabei et al., "First results from DAMA/LIBRA and the combined results with DAMA/NaI," *Eur. Phys. J. C* 56, pp. 333-355, 2008.
- [19] A. Giuliani and A. Poves, "Neutrinoless Double-Beta Decay," *Advances in High Energy Physics*, 2012.
- [20] J. Beringer et al., "Review of Particle Physics," *Physical Review D*, p. 86, 2012.
- [21] A. S. Barabash, "Experiment Double Beta Decay: Historical Review of 75 Years of Research," *Physics of Atomic Nuclei* 74, pp. 603-613, 2011.
- [22] Y. Fukuda et al., "Measurements of the Solar Neutrino Flux from Super-Kamiokande's First 300 Days," *Phys. Rev. Lett.* 81(6), p. 1158, 1998.
- [23] T. Araki et al., "Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion," *Phys. Rev. Lett.* 94(8), p. 081801, 2005.
- [24] 2015. [Online]. Available: http://www.etap.physik.uni-mainz.de/685_ENG_HTML.php.
- [25] M. Doi, T. Kotani and E. Takasugi, *Prog. of Theor. Phys. Suppl.* 83, pp. 1-175, 1985.
- [26] A. Gando et al., "Limit on Neutrinoless $\beta\beta$ Decay of Xe136 from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in Ge76," *Phys. Rev. Lett.* 110, 062502, 2013.
- [27] H. Klapdor-Kleingrothaus et al., "Evidence for neutrinoless double beta decay," *World Scientific*, 2006.
- [28] M. Agostini et al., "Measurement of the half-life of the two-neutrino double beta decay of Ge-76 with the GERDA experiment," *J. Phys. G; Nucl. Part. Phys.* 40, 2013.
- [29] A. H. Compton, "A Quantum Theory of the Scattering of X-Rays by Light Elements," *Physical Review* 21 (5), pp. 483-502, 1923.
- [30] A. Curioni, Laboratory and Balloon Flight Performance of the Liquid Xenon Gamma Ray Imaging Telescope (LXeGRIT), Columbia University, 2004.
- [31] 2015. [Online]. Available: https://en.wikipedia.org/wiki/Radiation_Portal_Monitor.

- [32] R. Dynes et al., "Evaluating Testing, Costs, and Benefits of Advanced Spectroscopic Portals for Screening Cargo at Ports of Entry," National Research Council, Washington, D.C., 2009.
- [33] K. Szymanska et al., "Resolution, efficiency and stability of HPGe detector operating in a magnetic field at various gamma-ray energies," *Nuc. Inst. and Meth. in Phys. A* 592, pp. 486-492, 2008.
- [34] T. Burr et al., "Radio-Isotope Identification Algorithms for NaI gamma spectra," *Algorithms* 2(1), pp. 339-360, 2009.
- [35] V. Y. Chepel, "A New Liquid Xenon Scintillation Detector for Position Emission Tomography," *Nuc. Tracks Radiat. Meas. Vol. 21 No. 1*, pp. 47-57, 1993.
- [36] J. Kurfess et al., *IEEE* 2, p. 1166, 2001.
- [37] T. Takahashi et al., *Phys. Rev. Lett. A*, pp. 12, 1771, 1975.
- [38] M. Szydagis et al., "NEST: A Comprehensive Model for Scintillation Yield in Liquid Xenon," *ArXiv e-prints*, 2011.
- [39] J. Wodin, "Phd. Thesis," Stanford University, 2007.
- [40] D. Cline et al., "A WIMP detector with two-phase xenon," *Ast. Phys.* 13 (4), p. 317, 2000.
- [41] L. Weissman et al., "Amplification and scintillation properties of oxygen-rich gas mixtures of optical-TPC applications," *JINST*, 2006.
- [42] M. Suzuki, "Mechanism of proportional scintillation in argon, krypton, and xenon," *Nucl. Inst. and Meth.* 164 (1), pp. 197-199, 1979.
- [43] E. Aprile et al., "Observation of Anti-correlation between Scintillation and Ionization for MeV Gamma-Rays in Liquid Xenon," *Phys. Rev. Lett.* 100, 2008.
- [44] S. Bricola et al., "Noble-gas liquid detectors: measurement of light diffusion and reflectivity on commonly adopted inner surface materials," in *Nuclear Physics B - Proceedings Supplements*, 2007.
- [45] "R8878 data sheet," Hamamatsu Photonics, 2008.
- [46] S. Biagi, MAGBOLTZ, program to compute gas transport parameters, 1997.
- [47] D. Y. Akimov et al., "The Zeplin-III dark matter detector: Instrument design, manufacture and commissioning," *Ast. Phys.* 27 (1), pp. 46-60, 2007.
- [48] A. Bolozyntsa, "Two-phase emission detectors and their applications," *Nucl. Inst. Meth. A* (422), p. 1999, 314-320.

- [49] [Online]. Available:
<http://www.ansys.com/staticassets/ANSYS/staticassets/resourcelibrary/brochure/ansys-maxwell-brochure.pdf>.
- [50] M. Ashby, "Chapter 12," in *Engineering Materials 2*, Oxford, Pergamon Press, 1992, p. 119.
- [51] M. Schwartz, *Soldering Understanding The Basics*, Materials Park, OH: ASM International, 2014.
- [52] J. Vuilleumier et al., "Search for neutrinoless double-beta decay in ^{136}Xe with a time projection chamber," *Phy. Rev. D* 48 (3), 1993.
- [53] R. Arnold et al., "Technical design and performance of the NEMO 3 detector," *Nucl. Inst. and Meth. in Phys. Res. A* 536, pp. 79-122, 2005.
- [54] G. Maroulis et al., "New basis for xenon and the interaction polarizability of two xenon atoms," *Chem. Phys. Lett.* 396, pp. 59-65, 2004.
- [55] [Online]. Available: <http://www.specialmetals.com/documents/Monel%20alloy%20400.pdf>.
- [56] C. Dahl, "The Physics of Background Discrimination in Liquid Xenon, And First Results from Xenon10 in The Hunt for WIMP Dark Matter," Phd Thesis, 2009.
- [57] C. H. Faham, "Prototype, Surface Commissioning and Photomultiplier Tube Characterization for the Large Underground Xenon (LUX) Direct Dark Matter Search Experiment," Brown University, Providence, RI, 2014.
- [58] V. Lisovski et al., "Low-pressure gas breakdown in uniform dc electric field," *J. Phys. D: Appl. Phys.* 33, pp. 2722-2730, 2000.
- [59] C. G. Wahl, J. M. Jaworski and Z. He, "UMImaging: A Software Package for Image Reconstruction From 3D-Position-Sensitive Gamma-Ray Detectors," *IEEE Transactions on Nuclear science*, Vol. 59, No. 4, pp. 1672-1680, 2012.
- [60] V. Solovov et al., *Nucl. Instrum. Methods Phys. Res., Sect. A* 516, p. 462, 2004.
- [61] D. Akerib et al., "The Large Underground Xenon (LUX) Experiment," *Nucl. Instrum. Meth. A* 704, pp. 111-126, 2013.
- [62] A. Dobi, Measurement of the electron recoil band of the LUX dark matter detector with a tritium calibration source, PhD Thesis, 2014.
- [63] E. Aprile, R. Mukherjee and M. Suzuki, *Nucl. Instr. Methods A* 300, p. 343, 1991.
- [64] H. Brooks, "Electron drift velocities in xenon," *J. Phys. D: Appl. Phys.* 15 L51, 1982.

- [65] G. Bakale, U. Sowadaand and W. Schmidt, *J. Phys. Chem.* 80, p. 2556, 1976.
- [66] S. Stephenson et al., "MiX: A Position Sensitive Dual-Phase Liquid Xenon Detector," *JInst*, 2015.
- [67] L. W. Kastens, S. Cahn, A. Manzur and M. D.N., "Calibration of a liquid xenon detector with ^{83}mKr ," *Physical Review C* 80, 045809, 2009.
- [68] E. Gushchin et al., *Sov. Phys. JETP*, 49 (5), pp. 856-858, 1979.
- [69] E. Aprile et al., "Simultaneous Measurement of Ionization and Scintillation from Nuclear Recoils in Liquid Xenon for a Dark Matter Experiment," *Phys. Ev. Lett.* 97, 2006.
- [70] Q. Lin et al., "High Resolution Gamma Ray Detection in a Dual Phase Xenon Time Projection Chamber," *JINST*, 2013.
- [71] M. Auger et al., "The EXO-200 Detector, part 1: detector design and construction," *JINST*, 2012.
- [72] E. Aprile and X. C. e. al, "The Xenon100 Dark Matter Experiment," *Astroparticle Physics*.
- [73] T. Doke et al., "Absolute scintillation yields in liquid argon and xenon for various particles," *Japanese Journal of Applied Physics*, 41(3R):1538, 2002.
- [74] M. Miyajima et al., "Average energy expended per ion pair in liquid argon," *Phys. Rev. A* 9:1438, 1974.
- [75] J. Thomas and D. A. Imel, "Recombination of electron-ion paris in liquid argon and liquid xenon," *Phys. Rev. A* 36, 2, 1987.
- [76] A. Mozumder, "Free-ion yield and electron-ion recombination rate in liquid xenon," *Chem. Phys. Lett.* 245, pp. 359-363, 1995.
- [77] V. e. a. Solovov, "Position Reconstruction in a Dual Phase Xenon Scintillation Detector," *IEEE Transaction on nuclear science*, Vol. 59, No. 6, pp. 3286-3293, 2012.
- [78] E. Aprile et al., "The electronics read out and data acquisition system for a liquid xenon time projection chamber as a balloon-borne Compton telescope," *Nuc. Instr. and Meth. in Phys. Res. A* 412, pp. 425-436, 1998.
- [79] S. Riboldi, "A low-noise charge sensitive preamplifier for Ge spectroscopy operating at cryogenic temperature in the GERDA experiment," in *Nuclear Science Symposium Conference Record*, 2010.
- [80] 2015. [Online]. Available: <http://www.4pcb.com/free-pcb-layout-software/>.
- [81] P. Sorensen, "Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches," *Nuc. Instr. and Meth. in Phys. A*, Vol. 635, pp. 41-43,

2011.

[82] M. Miyajima et al., *Phys. Rev. A*. *12*, pp. 1771-1775, 1975.