Searching for Dark Matter with XENON100, Research and Development for XENON1T, and Modulating Radioactive Decay Rates

Dissertation

zur

Erlangung der naturwissenschaftlichen Doktorwürde
(Dr. sc. nat.)

vorgelegt der

Mathematisch-naturwissenschaftlichen Fakultät

der

Universität Zürich

von

Peter R. Barrow

von

Leukerbad VS

Promotionskomitee:

Prof. Dr. Laura Baudis (Vorsitz)
Prof. Dr. Ben Kilminster
Prof. Dr. Marc Schumann

Zürich, 2016

Abstract

Zusammenfassung

Contents

Listing of figures

Listing of tables

Acknowledgments

1

Introduction

Since its hypothesised existence in 1933 by Fritz Zwicky [4] dark matter remains one of the greatest known mysteries of the universe. Observations such as the rotation velocities of galaxies that show orbits of stars moving at far greater speeds that the gravitational force of luminous matter should allow, and the bullet cluster, where a distinct offset between the gravitational centre, and observable of mass is demonstrated. The ever mounting evidence points towards a yet undiscovered, and vast quantity of matter that drastically alters the gravitational properties of large scale structures in the universe. Cosmic microwave background (CMB) observations, most recently from the Plank experiment, predict that dark matter occupies 26.8% of the total energy density, and thus outweighs "normal" baryonic matter by over a factor of 5 at 4.9%(Plank Collaboration, 2013), with the remaining 68.3% due to dark energy. As such, the determination as to the nature of dark matter is of high importance to our understanding of the dynamical mechanics of the universe.

1.1 Cosmological Evidence for Dark Matter

Currently, dark matter has demonstrated solely gravitational effects on large scale objects such as galaxies or clusters of galaxies, as well as light. The first observations of this, performed by Fritz Zwicky, attempted to judge the total mass of the the Coma cluster

through redshift measurements, using the virial theorem, relating the total kinetic and gravitational potential energy as shown by:

$$\frac{M_{tot} \cdot v^2}{2} = \frac{G \cdot M_{tot}^2}{4 \cdot R_{tot}} \tag{1.1}$$

Where M_{tot} is the total mass of the system, v is the average velocity, G is the gravitational constant, and R_{tot} is the radius of the system. By arranging equation 1.1 to find the mass of the object, we arrive at:

$$M_{tot} = \frac{2 \cdot R_{tot} \cdot v^2}{G} \tag{1.2}$$

Measurements using the virial theorem resulted in a large amount of matter not accounted for from predictions by luminous galaxies. A study of 89 galaxy clusters predicted a mass-to-light ration of 230-250 [5].

Additional evidence towards the presence of dark matter also arose from a study performed by Rubin and Ford [6] who studied the rotation of the Andromeda nebula. What they found, contrary to conventional Newtonian dynamics, was that the rotation of objects at large distances from the galaxies core remained approximately constant as a function of radius. This observation is explained via the presence of a dark matter halo that encompasses each galaxy, allowing these object to orbit at such high velocities.

The potential gravitational influence of dark matter may also be observed in its manipulation of light. Massive cosmological objects are capable of observably altering the direction of light as shown in Figure 1.1

Observations of such gravitational lensing are capable of estimating the mass of the system. Comparing this value to the luminous matter observed leads to further evidence towards an unseen mass.

1.2 Predictions for particle Dark Matter

1.2.1 WIMPs

It may be possible for dark matter to be capable of interacting with baryonic matter. Such a particle must be consistent with current comsological observations. This gives rise to a new theoretical particle outside of the Standard Model known as a Weakly Interacting Massive Particle, or WIMP [7]. Such a particle is best motivated by theories involving supersym-

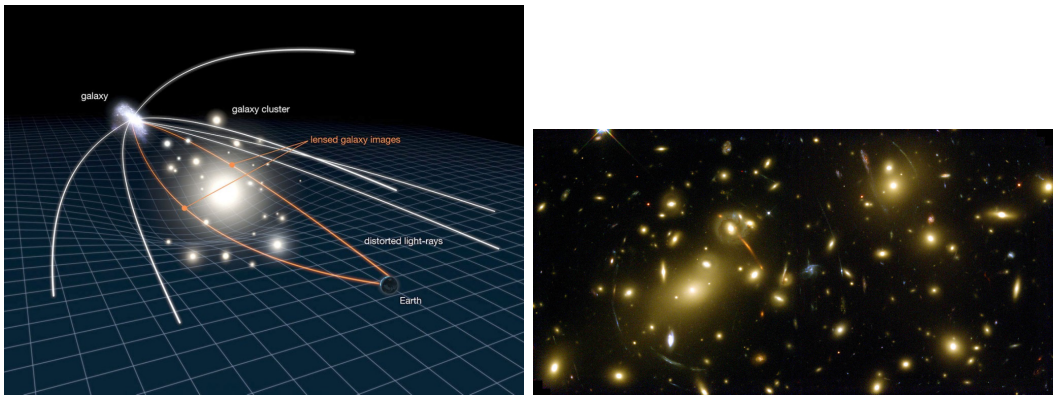

Figure 1.1: Left: Schematic representation of the process by which galaxy clusters gravitationally lens light. Right: NASA Hubble image taken of galaxy cluster Abel2218.

metry [8] or universal extra dimensions [9]. a WIMP would interact only very weakly with baryonic matter, on the weak-scale, thus accounting for the primarily gravitational nature of dark matter. This aspect also arises due to our persistent efforts to observe new particles. A strongly interacting dark matter candidate would have likely been discovered by previous experiment. A WIMP, however, must maintain some cross-section with baryonic matter in order for a direct observation to be possible. Additionally, a wimp must not interact with light, and thus be neutrally charged. Therefore, by design, a WIMP would be an ideal candidate for a particle outside the standard model in order to explain the multiple, and varied observations of dark matter.

1.2.2 Super-WIMPs

There is an increasingly overwhelming amount of gravitational evidence for the presence of dark matter within our universe. From the rotational velocity distribution of galactic objects, to gravitational lensing from large scale structures in our universe, as well as results from CMB measurements, the existence of yet undiscovered matter is increasingly apparent. As a result, the search for dark matter has sparked a wide vareity of experiments to search for it's presence. Furthermore, simulations of Cold Dark Matter (CDM) find a richer galactic structure than that observed within our own galaxy. As no such particle has yet been found and confirmed. This motivates the hunt for alternative signatures that may be possible to detect in existing experiments. One such model is the bosonic super-WIMP. These arise in a model of warm, bosonic dark matter with keV scale masses that couple electromagnetically to the standard model.

In this thesis I present the search for the presence of vector and pseudo scalar dark matter with the XENON100 detector. For vector superWIMPs, the absorption of the particle is analogous to that for the photo-electric effect. As a result, one can scale the cross-section accordingly shown in eq 1.3.

$$\frac{\sigma_{abs} v}{\sigma_{photo}(\omega = m_V)c} \simeq \frac{\alpha'}{\alpha} \tag{1.3}$$

Where σ_{abs} and v are the cross section of absorption for the vector boson and it's velocity respectively. σ_{photo} is the cross section of the photoelectric effect. α and α' are, respectively, the fine structure constant, and the vector boson equivalent.

For pseudo-scalar superWIMPs, the scaling of the cross section must be performed differently as shown in eq 1.4.

$$\frac{\sigma_{abs} v}{\sigma_{photo}(\omega = m_V)c} \simeq \frac{3m_a^2}{4\pi\alpha f_a^2} \tag{1.4}$$

Where m_a is the mass of the pseudo-scalar superWIMP, and f_a is a dimensional coupling constant.

Assuming that super-WIMPs are non-relativistic, and that they consititute the entirety of the dark matter halo within our galaxy with a local density of 0.3GeV/cm^3, the interaction rate in a direct detection experiment can be expressed as in eq. 1.5 and 1.6 for pseudo-scalar and vector super-WIMPs respectively.

$$R = \frac{1.29 \times 10^{19}}{A} g_{Ae}^2 m_A \sigma_{pe} [\frac{1}{kg} \frac{1}{day}] \tag{1.5}$$

Where $g_{Ae} = \frac{2m_e}{} f_a$ is the axio-electric coupling and A is the atomic number of the target atom

$$R = \frac{4 \times 10^{23}}{A} \frac{\alpha'}{\alpha} \frac{1}{m_A} \sigma_{pe} [\frac{1}{kg} \frac{1}{day}] \tag{1.6}$$

Due to the non-relativistic recoils that one would observe under such conditions, the recoil spectrum is expected to be a monoenergetic peak around the value of the super-WIMP mass.

From here, limits and predictions can be set as to the coupling constants expected from a particular mass of super-WIMP. Figure 1.2 displays constraints given by indirect calculations with respect to the theoretical prediction given by the abundance of dark matter.

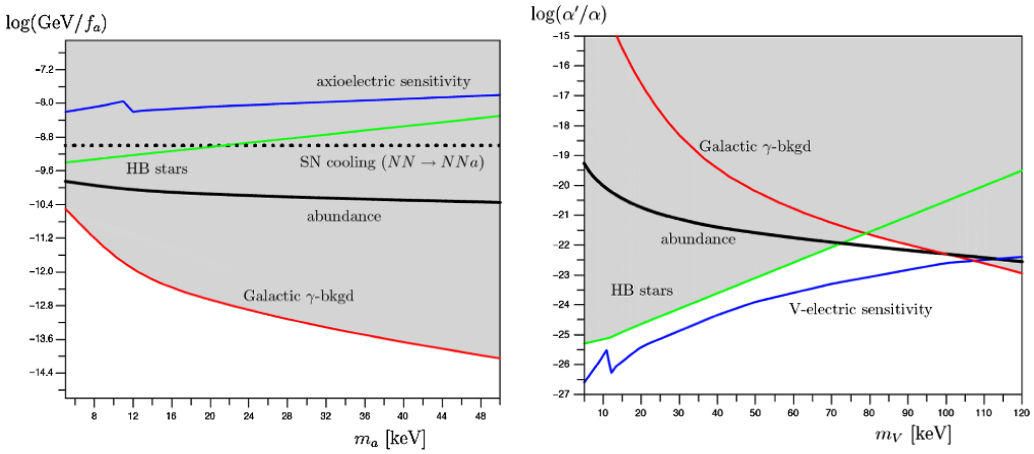

Figure 1.2: Direct-detection sensitivity to pseudoscalar (left) and vector (right) dark arising from the axioelectric cross section on Ge, assuming a fiducial sensitivity of the detector equivalent to a 1 pb cross section for a 100 GeV WIMP. We also show the constraints arising from the He-burning lifetime in HB stars, from SN cooling via a coupling to the neutron magnetic moment with $f_{aNN} = f_a$, and most significantly the monochromatic γ background from decays in the Galaxy. The gray shaded region is excluded by the latter indirect constraints. The thick black line corresponds to the parameters required to reproduce the required dark matter abundance from thermal production with $f_{abb} = 14 f_a$. Figure and text taken from [1]

1.3 Detection of Dark Matter

The process of observing a dark matter particle may be placed under three distinct classes of experiment: accelerator, indirect and direct detection. Accelerator experiments, such as the LHC at CERN, have the potential to create dark matter via their high energy collisions, and utilise missing momentum to obtain an observation. Indirect experiments intend to observe the presence of an excess of particles created via the collision, or annihilation of dark matter. Most notably, AMS-2 published results in 2013 attempting to observe an excess of positrons [10]. Direct detection experiments take advantage of the weakly interacting nature of WIMPs to observe effects of dark matter colliding with nuclei. Numerous direct detection experiments have been performed over the last decade alone, notably the XENON10 experiment [11], WARP [12] the XENON100 experiment published results in 2012 producing the lowest limits at the time for the cross section of spin-independent WIMP

interactions[13]. Only recently, in October 2013, the LUX experiment announced its results further lowering the limit of WIMP cross section to $7.6\text{x}10^{-46}\text{cm}^2$ [14]. Figure **??** displays the current status of the dark matter search. The area above each line has been excluded by the respective experiment. The closed lines represent claimed dark matter signals from DAMA [15], CoGeNT [16], and CRESST [17]. The rate at which experiments are producing new, and groundbreaking data for dark matter shows both the importance and urgency to understand it. As a result, the XENON1T experiment is currently under construction to push this sensitivity limit further.

1.3.1 Direct Detection

In the case of a detection of dark matter, direct detection experiments will provide the most conclusive evidence of its existence. These primarily rely on the interaction of dark matter with baryonic matter through some channel that communicates with both models.

1.3.2 Liquid Xenon Based Direct Detection

A promising sensitive medium in which to observe direct interaction of dark matter is Liquid xenon.

Particle interactions within the detector both excite and ionize the xenon atoms. De-excitation from this state produces a signal from the scintillation (S1). The ionized electrons are then drifted via an electric of 530 V/cm to the liquid-gas boundary where a 12 kV/cm extraction field produces proportional scinitillation in the gas (S2). Three-dimensional position reconstruction of the original interaction can then be performed from the information of these two signals. The time difference between the S1 and S2 signal allows the vertical position (z) to be determined, while the hit pattern on the PMTs caused by the S2 allows the radial position of the event (x,y) to be calculated. The conversion from hit pattern to position reconstruction is performed via machine learning through simulated events.

The primary observable difference between an electronic and nuclear recoil (ER, and NR respectively) is the quantity of ionized electrons from an interaction. As a result, the S2 signal for electronic recoils is greater than that for an equivalent nuclear recoil. This discrepency can be exploited to differentiate between these two types of events by using the ratio between the size of the S1 and S1 signal. In contect of the search for super-WIMPs, due to the coupling to the electromagnetic field, this ratio will be employed to eliminate contamination of data from nuclear recoils.

A schematic representation of the signal process can be observed in Figure **??**.

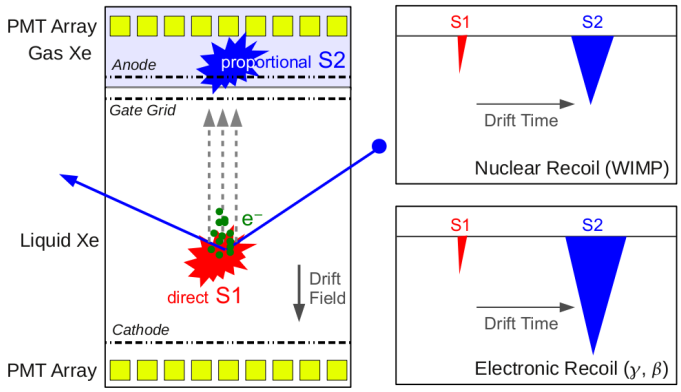

Figure 1.3: Schematic demonstrating the detector principle of a two-phase time-projection chamber.

The determination of the energy of an event is calculated via the S1 signal using the equation 1.7.

$$n^{exp}(E) = R(E) \times Q(E) \times f \times E \equiv L_Y \times E, \tag{1.7}$$

where the factor $f = 3.76$ PE/keV is the derived XENON100 light yield at 32.1 keV and zero field. The function n^{exp} is the expected number of photoelectrons as a function of Energy.

1.3.3 Indirect Detection

There is a possibility in which there is no interaction that may cause dark matter to directly interact with baryonic matter, or that the local abundance of dark matter may not be great enough to currently detect it . In this case, it may still be possible to observe dark matter via its interactions with itself, producing standard model particles which can then be observed. These direct detection experiments intend to search for the remnants of these collisions with the intent of finding an excess of particles of a particular energy as a result of collisions with dark matter not directly made within the experiment.

1.3.4 Production in Colliders

If dark matter were to have a channel connecting it to the standard model, it stands to reason that it would be therefore possible to create dark matter particles in high energy colliders. In this case, the newly formed dark matter particle would escape the detector without interaction. The missing energy from such an event can then be re-constructed in order to potentially determine the mass of the particle.

1.4 Modulating Radioactive Decay Rates

In 1902 the radioactive decay law was discovered by Rutherford and Soddy, shown in equation 1.8.

$$N(t) = N_0 \exp^{-\lambda t} \tag{1.8}$$

Where N_0 and $N(t)$ is the number of active nuclei at $t = 0$ and at time t respectively. λ is the nuclear decay constant. This law, since its discovery, has been used to determine the half life of radioactive decays.

Recently, however, precise measurements of the rate of radioactive decays have provided interesting results, giving hints that this exponential decay law may not describe the true nature of radioactive decays, and that it they could be influenced via external sources.

1.4.1 Modulation Observations from PTB

The Physikalisch-Technische Bundesanstalt (PTB) has made precise measurement of various radioactive isotopes since 1989. Each measurement has been performed within their 4πpressurised chamber, with the ionisation current being read out via the Townsend balance-of-charge method [18]. This method was upgraded in 1998 by a Keithley electrometer. As a result, all measurements are separated into two datasets, those taken via the Townsend balance-of-charge mathoed between 1990-1995 and the Kiethely electrometer measurements between 1999-2008. Measurements were taken with the following isotopes: ^{8}5Kr, ^{9}0Sr, 108Ag, ^{1}33Ba, ^{1}37Cs , ^{1}52Eu, ^{1}54Eu while using ^{2}26Ra as a reference source.

In order to determine the systematic errors in the measurements, the residuals were plotted and defined as:

$$R = A(t)_{\text{expected}} - A(t)_{\text{measured}} = -N_0 \cdot \exp^{\lambda t} - N(t)_{\text{measured}} \tag{1.9}$$

Where $A(t)$ is the activity at time t, N_0 is the initial number of active nucleii, and λ is the decay constant.

The ^{2}26Ra reference source was measured over a period of 10 years [2], the results of the residuals of this measurement are shown in Figure 1.4 (left). Results displayed an annual modulation of magnitude A_{mod}=0.15% with the maximum positive deviation occurring during February denoting a phase shift $\phi_{\text{mod}} = 0.144$ months. To determine the validity

of the modulation, the residual ratio between $^{1}54$Eu and $^{2}26$Ra was taken and is shown in Figure 1.4 (right). Plotting this ratio as a function of time visibly eliminates any previously observed modulation. This then implies that the phase and magnitude of the modulation is identical for both sources.

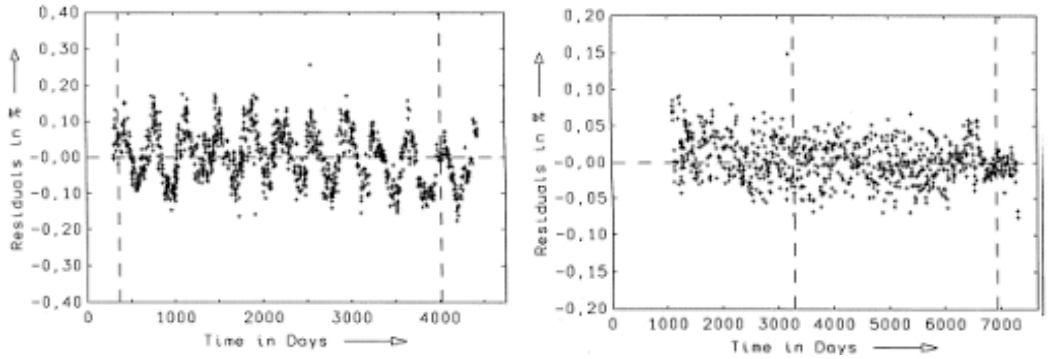

Figure 1.4: "The modulation is seen in the individual data, but disappears when one takes the ratio of these two sources. A data point is an average of the ratio of sample to reference source currents, containing about 30 individual sample measurements taken over 3 days and corrected for background." Figure and text from [2]

To determine the source of the modulation, the residuals of $^{8}5Kr$, $^{1}08$Ag, $^{1}52$Eu and $^{1}54Eu$ were also studied. The results of the observed residuals from $^{1}54Eu$ are shown in Figure 1.5

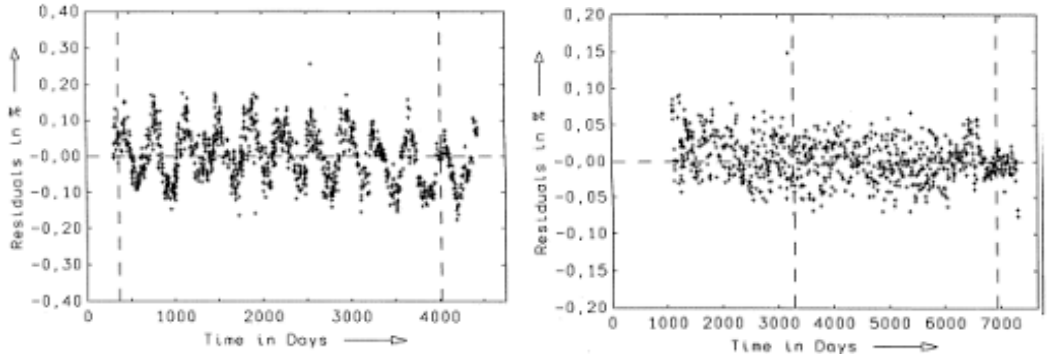

Figure 1.5: "The modulation is seen in the individual data, but disappears when one takes the ratio of these two sources. A data point is an average of the ratio of sample to reference source currents, containing about 30 individual sample measurements taken over 3 days and corrected for background." Figure and text from [2]

1.4.2 Influence of Solar Flares

In order to describe the modulations observed in other experiments, it has been postulated that a change in the flux of cosmic particles may have yet unknown interactions with radioactive sources, causing them to alter their decay rate during times of high flux. Such a change in the influx of particles may be observed during periods in which a solar flare is directed towards Earth. A solar flare is capable of releasing up to 6×10^{25} J of energy over a span of minutes, ejecting a plasma consisting primarily of electrons and protons into the solar system. Not only may a change in the flux of solar particles be observed over this short time frame, but the Sun itself exhibits an activity cycle that spans 11 years, over which time the flux of solar particles may alter.

A study by Jenkins and Fischbach in December 2006 [3] monitored the decay rate of $1\mu Ci$ sample of ^{5}4Mn using an NaI(Ti) scinitllator as a detector. During thier measurement two solar flares occurred, detectord by satellites due to the decrease in high energy protons and X-rays. They observed a decrease in measured event rate from their expected exponential curve. The results of their observations are shown in Figure 1.6

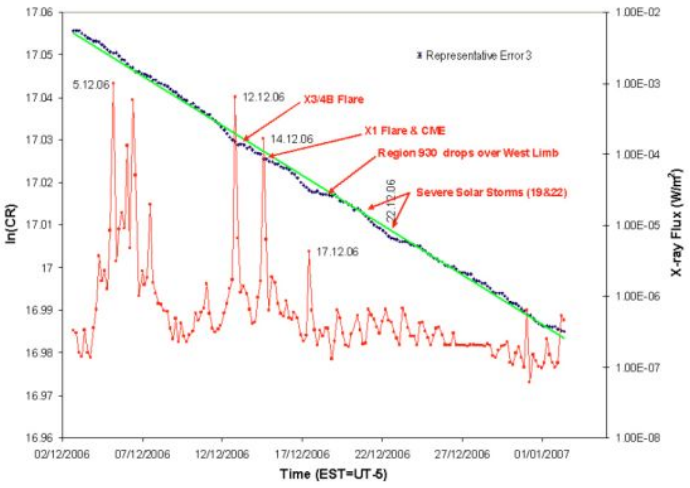

Figure 1.6: "December 2006 54Mn data, and GOES-11 x-ray data, both plotted on a logarithmic scale. For 54Mn, each point represents the natural logarithm of the number of counts 2.5×10^7 in the subsequent 4 hour period, and has a Poissonion statistical error shown by the indicated error bar. For the GOES-11 x-ray data, each point is the solar x-ray fluxin W/m2 summed over the same real time intervals as the corresponding decay data. The solid line is a fit to the 54Mn data. The dates for other solar events are also shown by arrows." Figure and text from Jenkins et al., (2008) [3]

2

XENON100

Installed at the Laboratori Nazionali del Gran Sasso (LNGS) XENON100 is the second iteration of the XENON project, after XENON10, with the goal to detect WIMPs via the use of a dual-phase time-projection chamber (TPC). Results from XENON 100 released in 2012 provided limits on the WIMP-nucleon cross section at $2 \times 10^{-45}\, cm^2$ at 55 GeV. The lowest limit at the time of publishing.

The primary purpose of the XENON100 detector is to search for the presence of Weakly Interacting Massive Particles (WIMPs) via the elastic scattering off nuclei in liquid xenon (LXe). Located, and operating in the Laboratori Nazionali del Gran Sasso (LNGS), the detector itself consisted of a cylindrical TPC 30 cm in height and 30 cm in diamter containing 62 kg of LXe. The TPC is then covered above and below the sensitive volume with 242 low-radioactivity Photo-Multiplier Tubes (PMTs).

[More details on XENON100 to be added]

This section will provide details to the work performed for XENON100. This will include the analysis towards defining a cut for Run12, and the search for bosonic dark matter in the form of super-WIMPs.

2.1 Noisy PMTs from Run12

Within each run from XENON100, a particular set of PMTs are observed with systematically higher rates from the average. These PMTs are thought to be registers such events due to a higher rate of noise as a result of a change in the conditions, or the deterioration of the PMTs with time. The specific detectors that display such a behaviour are known to change between runs, and thus defining which PMTs are responsible is of the upmost importance due to the potential for incorrectly recognising a signal as a result of an increased coincidence due to noise.

2.1.1 Creating a Manual Cut

As has been performed in previous runs, a manually defined list of noisy PMTs must be determined, and the coincidence level must then be raised based on the number of these PMTs that observed a given signal. As one would expect that signal across PMTs in an array would be approximately even, displaying the number of S1 events observed for each PMT over calibration data may give an insight into detectors that are experiencing high noise. Figure 4.4 applies a cut to AmBe data requiring an S1 signal to be below 3pe. This requirement is to focus the analysis towards the energy region in which noise is most likely.

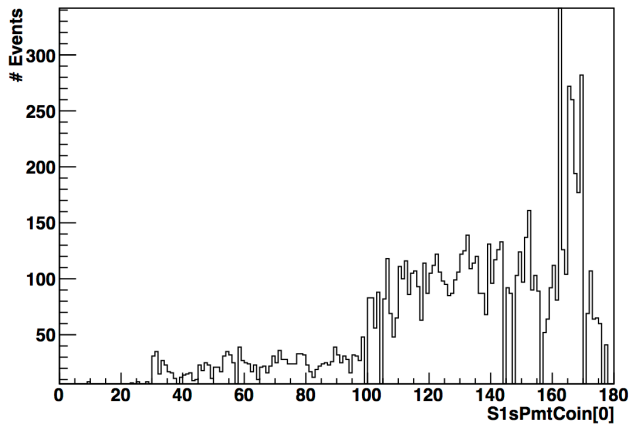

Figure 2.1: PMT numbers registering events for the greatest s1 signal over 5 AmBe files using run 12 cuts excluding coincidence, entropy, and veto requirements. Events surviving a cut on S1 width and below 3pe have been selected

From Figure 4.4, the largest peak rises well above the shown plot, and can be attributed to PMT 162. The following, simple, coincidence cut can be placed in an attempt to remove these events. If peaks remain after this, we will introduce a further cut to eliminate noisy

PMTs, as has been accomplished in run 10. One cannot assume that the same noisy PMTs in run 10 are also present in run 12, therefore a new cut to establish the list of noisy PMTs for this run must be made, named noisypmts12. The cuts are defines as such:

$$Xs1coin0 = (S1sCoin[0] > 1) \tag{2.1}$$

$$Xs1coin3 = (S1sCoin[0] > (1 + noisypmts12)) \tag{2.2}$$

Applying the Xs1coin0 cut results in the distribution observed in figure 2.2.

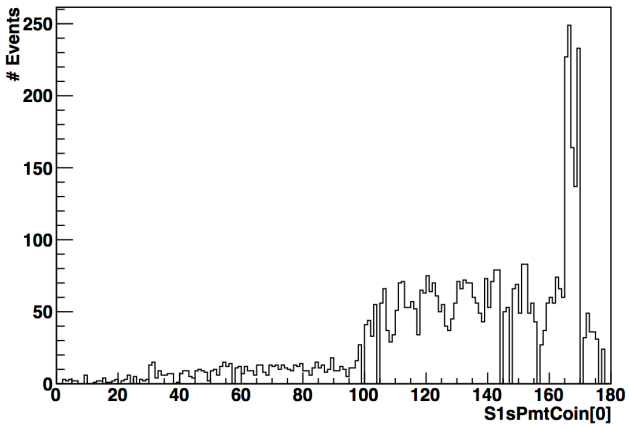

Figure 2.2: PMT numbers registering events for the greatest s1 signal over 5 AmBe files using run 12 cuts excluding coincidence, entropy, and veto requirements. Events surviving a cut on S1 width and below 3pe have been selected

From Figure 2.2, one can observe that the peak observed in Figure 4.4 at PMT 162 has been eliminated. There remains, however, an excess of events from PMTs 165 to 169. To determine which of these PMTs are contributing the most, or all of the peak seen in Figure 2.2, the cut noisypmts12 was implemented on one pmt at a time in the form:

$$Xs1coin3 = S1sCoin[0] > (1 + ((S1s[0][PMT] > 0.35))) \tag{2.3}$$

Zooming into the peak from PMTs 164-169, the implementation of this cut for each of these PMTs is as follows in Figure **??**.

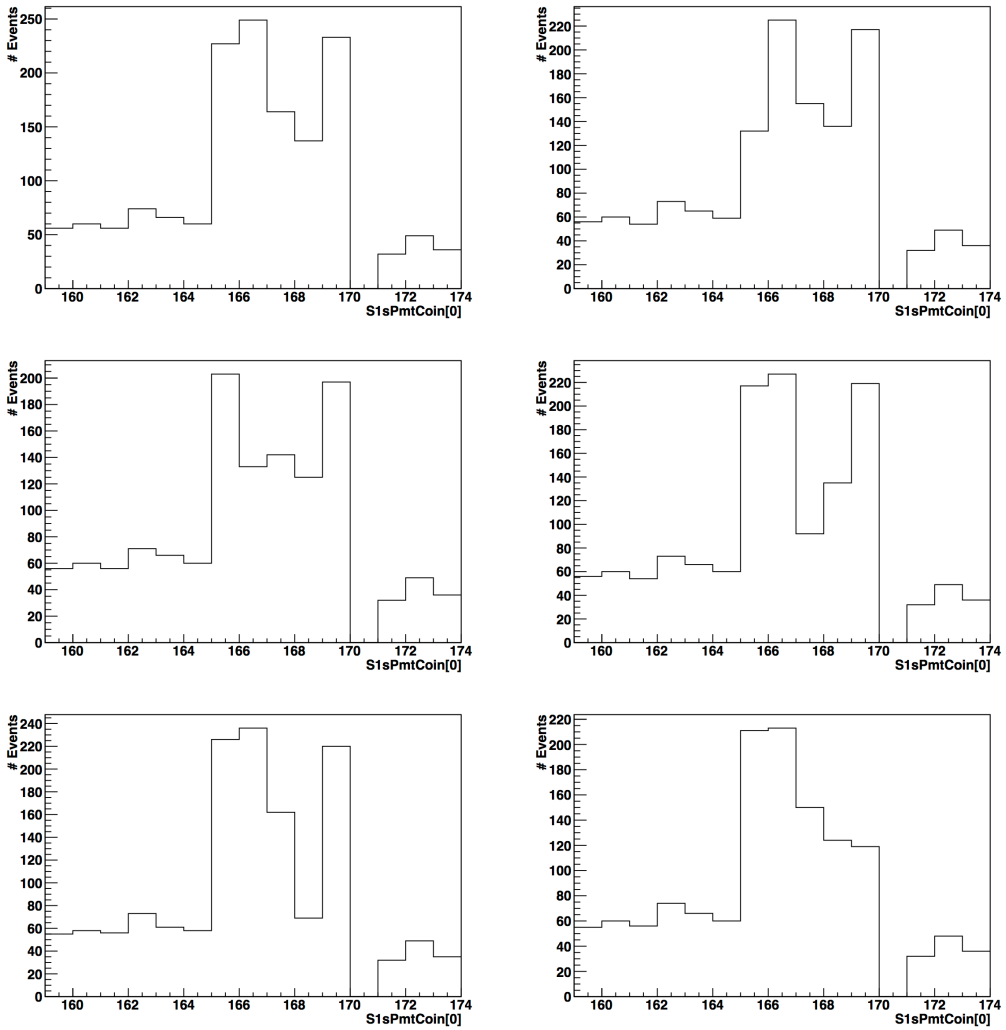

Figure 2.3: Results of adding an additional coincidence cut to PMTs 165-169 individually.

From Figure 2.3, no single noise cut is able to completely eliminate more than the peak of PMT being cut. It therefore brings to reason, that all PMTs should be implimented into "noisypmts12" as follows:

$$\begin{aligned} Xs1coin3 = S1sCoin[0] > (1 + ((S1s[0][165] > 0.35) + (S1s[0][166] > 0.35) \\ +(S1s[0][167] > 0.35) + (S1s[0][168] > 0.35) + (S1s[0][169] > 0.35))) \end{aligned} \tag{2.4}$$

The results of which are shown in Figure 2.4. From this figure, it is clear that the peaks shown in the previous plots have been eliminated. Upon closer inspection however, the cuts may have been too stringent. A compromise must therefore be found.

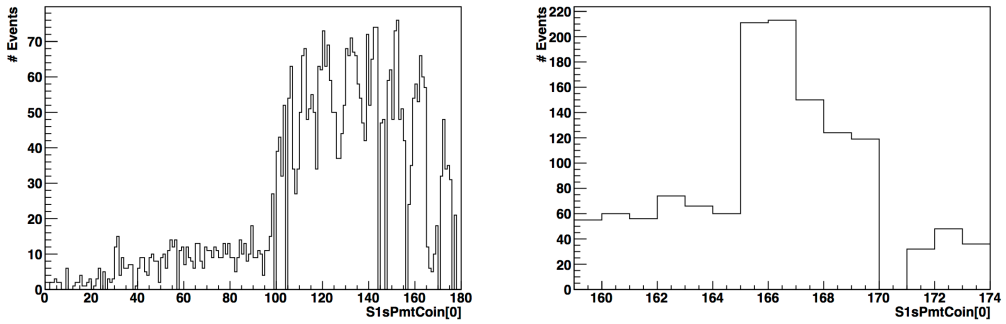

Figure 2.4: Implementation of an additional coincidence cut for all PMTs from 165-169.

As this cut is too strict on events it is worth looking at which PMTs contribute most to the peak, and to attempt to eliminate those. From figure 3, it can be seen that PMTs 165, 166, and 169 count almost double the number of counts as 167 and 168. A NoisyPmts12 cut involving just these three PMTs is shown below in Figure 2.5.

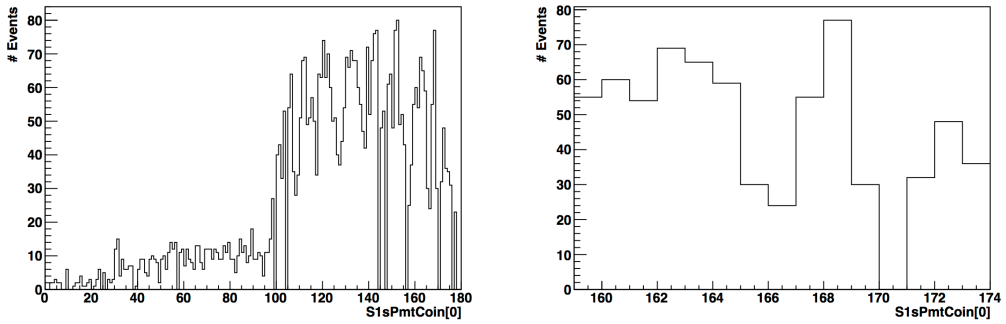

Figure 2.5: Implementation of an additional coincidence cut for all PMTs from 165-169.

From Figure 2.5, the Noisy PMT cut on pmts 165 166 and 169, seems to have completely eliminated the peak of events seen previously. Other combinations of these PMTs did not effectively remove the excess of events, therefore the current version of the cut is as follows:

$$noisypmts12 = ((S1s[0][165] > 0.35) + (S1s[0][166] > 0.35) + (S1s[0][169] > 0.35)) \quad (2.5)$$

2.1.2 Alternative Noise Cuts

There may be a more efficient way to eliminate noisy effects. Rather than the more "brute force" method of specifically selecting noisy pmts, by using such entropy cuts, or just by increasing the coincidence level, it might be possible to have a more general cut that can be applied to all data and runs. The following section investigates the possibility to find such a cut that efficiently eliminates all noise events, without the need to select specific noisy PMTs.

In the figure below, the cuts of Xs1coin0 are compared to the addition of cuts such as S1scoin[0]>2, Xentropy0, and the Xs1coin3 as defined in the previous section.

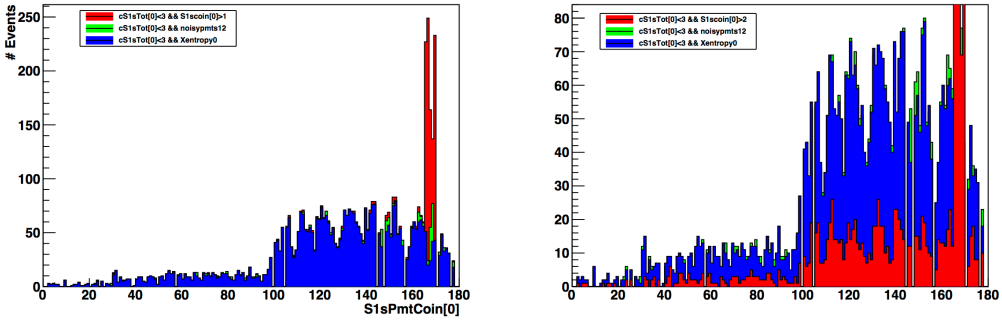

Figure 2.6: Number of events against PMT number demonstrating the effect of the cuts S1scoin[0]>1, Xentropy0,S1scoin[0]>2 and noisypmts12 on the reduction of noisy events in AmBe data.

Figure 2.6 demonstrates that the addition of the noisypmts12 cut in Xs1coin3 has a very similar effect to the entropy cut. Further more, increase the global coincidence level to S1scoin[0]>2, does not eliminate the excess of events, and is clearly too strict on PMTs that do not appear to be as noisy.

In conclusion, Xs1coin3 as defined in the previous section with noisypmts12, as well as Xentropy0 seem to be effective in remove the observed excess of events. Analysis into the individual events cut by these parameters must be investigate to determine the efficiency in which both cuts perform. At first glance, however, it seems as though Xentropy0 is a concise and effective cut to eliminate noise without the need for a specific list of noisy pmts.

2.1.3 Cross-check with Dark Matter Data

In order to determine the effectiveness of the Xs1Coin3 and Xentropy0 cuts, it is necessary to determine how many events that are selected are also noise cuts. The 5 AmBe files that have been used in the analysis above provide statistics that are too high to reasonably

sift through all the selected, or non-selected events. Dark Matter data, therefore, must be analysed as the few remaining selected events can be checked individual to determine the efficiency with which each cut performs. Detailed below is the analysis comparing the Xs1coin3 and Xentropy0 cuts.

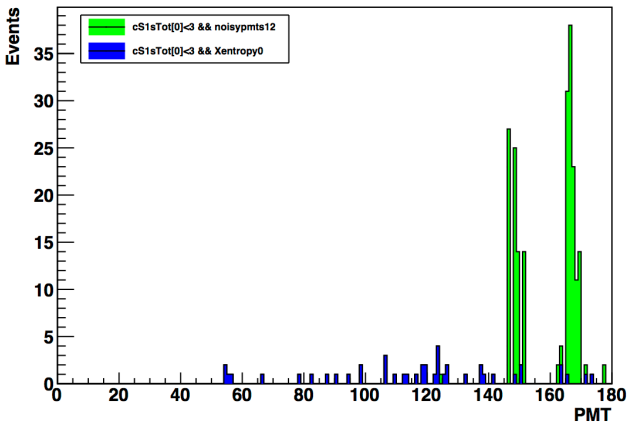

Figure 2.7: Comparison plot of the NoisyPmts12 and entropy cut for dark matter data with events with S1<3pe.

Figure 2.7 shows a comparison on dark matter data for events S1<3pe between the Xs1coin3 and Xentropy cut, as described in the previous section. It can clearly be seen that an excess of events not seen in the AmBe data, is clearly present in the dark matter data with Xs1coin3. The entropy cut, on the other hand, does not show any obvious excess of noisy events.

In an attempt to reduce the excess of events, NoisyPMTs12 was expanded to include the PMTS: 146, 148, 165, 166, 167, 169. The results of this comparison are shown below. The Xs1coin3 cut is now as follows:

$$\begin{aligned} noisypmts12 = ((S1s[0][146] > 0.35) + (S1s[0][148] > 0.35) + (S1s[0][165] > 0.35) \\ +(S1s[0][166] > 0.35) + (S1s[0][167] > 0.35) + (S1s[0][169] > 0.35)) \end{aligned} \tag{2.6}$$

From Figure 2.8, it can be observed that even with 6 noisy PMTs included in Xs1coin3, excesses of events still remain, and further noisy PMTs are therefore still required to be added. At this point, however, it is becoming clear that the number of PMTs required to cut out noisy events via Xs1coin3 is unviable in comparison to the cut on entropy. It is

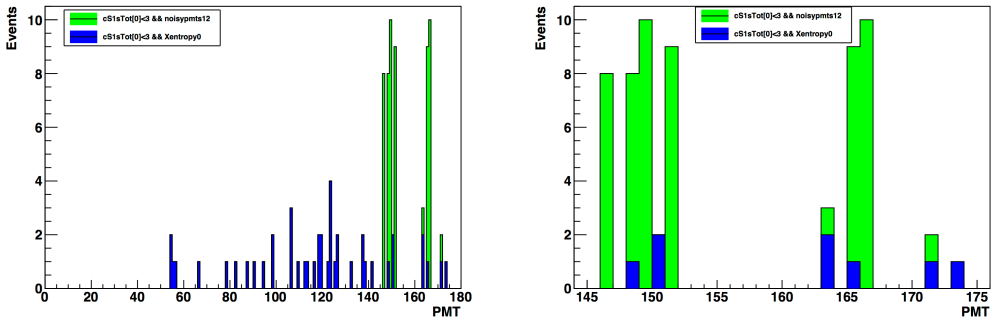

Figure 2.8: Comparison of number of events observed in dark matter data between the implimentation of an entropy cut, or a manual cut , increasing the coincidence on PMTs146, 148, 165, 166 and 167. The events for all pmts are show (left) along with a zoom into the region of excess (right).

then necessary to determine which PMTs should be added to NoisyPMTs12 to eliminate all excesses of events. By the edition of PMTs indexes 149 and 152, we arrive at Figure 2.9.

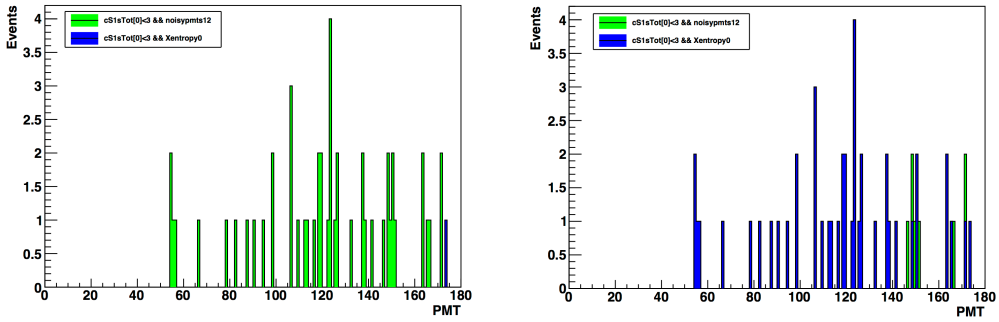

Figure 2.9: Comparison of number of events observed in dark matter data between the implimentation of an entropy cut, or a manual cut , increasing the coincidence on PMTs146, 148, 149, 152, 165, 166 and 167.

$$\begin{aligned} noisypmts12 = ((S1s[0][146] > 0.35) + (S1s[0][148] > 0.35) + (S1s[0][149] > 0.35) \\ +(S1s[0][152] > 0.35) + (S1s[0][165] > 0.35) + (S1s[0][166] > 0.35) \\ +(S1s[0][167] > 0.35) + (S1s[0][169] > 0.35)) \end{aligned} \quad (2.7)$$

To conclude the analysis, AmBe data has been shown to be ineffective at accurately determining the presence of noisy events. The high statistics mean that small excesses in noise are not identifiable. Furthermore, a cut on the entropy can eliminate observed excesses

in events efficiently on its own, without any other coincidence requirement. It is therefore more efficient, and flexible to apply a cut on entropy above increasing the coincidence on individual noisy PMTs. In further support in the use of the entropy cut, future runs for both XENON100, and in the future, XENON1T, will no longer require an individual list of noisy PMTs in order to eliminate noise from events.

2.2 The Search for Bosonic Super-WIMPs

In the context of searching for dark matter with XENON100, bosonic super-WIMPs allow the detector to search for potential signals in previously unexplored territory. Due to the axio-electric effect, bosonic super-WIMPs couple electronically to the target atom, and thus signals produced from the absorption of such a particle are found within the electronic recoil band. Furthermore, due to the unbounded masses on super-WIMPs, XENON100 is able to probe all potential energies sensitive to the detector. In contrast, previous searches for dark matter have involved WIMPs, signals of which are expected within the nuclear recoil band, and axions, that although predicted to interact electronically, the analysis is limited to energies less than 30keV in XENON100. The ability to probe the electronic recoil band at higher energies than previously searched exemplifies the necessity and importance of such an analysis.

2.2.1 Predictions of Interaction Rates in XENON100

Using the equations for the rate of both pseudo-scalar and vector super-WIMPs it is possible to predict the potential number of interaction within XENON100 for a range of coupling constants in our detector. The exposure chosen for this prediction was that for the dark matter run, 34kg fiducial volume and 225 live days. The results of which are shown in Figures 2.10 for pseudoscalar and vector super-WIMPs.

The structures observed at 35 keV, as well as at low energies are a result of an increase in the cross section of the photo-electric effect due to the excitation of higher energy levels.

2.2.2 Determining a Preliminary Limit

As this analysis will be performed in the electronic recoil band, a system to select events within the band must be made within the region of interest before further analysis can continue. To begin, it is important to understand the results performed in the initial analysis by kevin.In order to do this one must redefine the electronic recoil band from scratch using the cuts applied for run 10 for the electronic band flattening. The following cuts were applied to Co and Th data using the minimum bias 2 files.

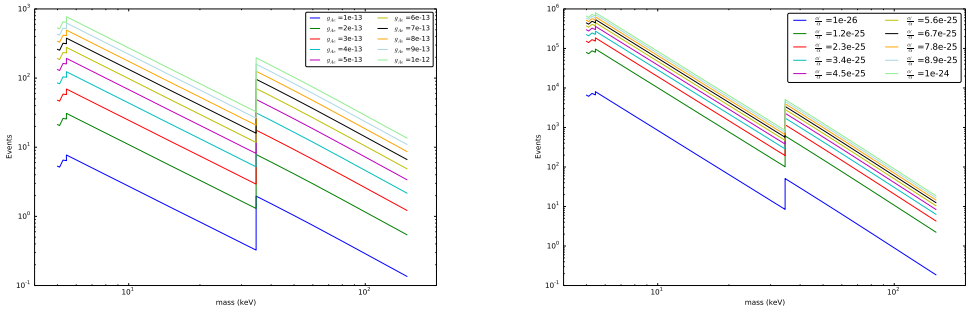

Figure 2.10: Predicted number of events over a range of coupling constants for pseudo-scalar (left) and vector (right) super-WIMP interactions within a xenon based direct detection experiment with 34kg fiducial volume, and 225 live days of data.

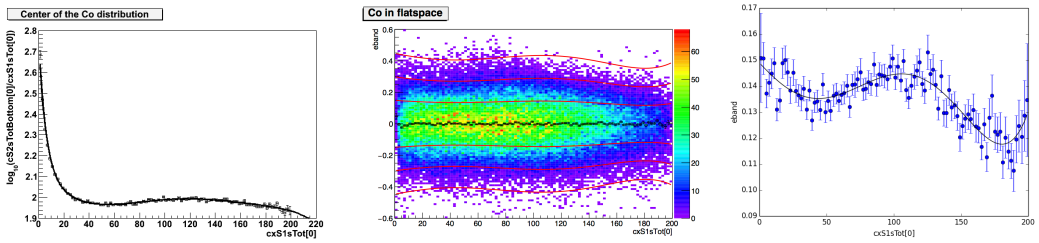

Figure 2.11: Left: Fit of the eletronic recoil band prior to flattening. Centre: Flattened spectrum of the electronic recoil band passing all initial cuts. Right: 1 sigma band of the electronic recoil band in flattened space.

From the remaining events, the s1/s2 ratio was taken as a function of s1 between 0-200pe in bins of 2pe.The resulting histogram for each bin was used to define the centre, and 1 sigma regions of the electronic recoild band. The spectrum was then flattened, and the band was defined as the 3 sigma region on this flattened band, as per kevin's definition.

The results of the electronic band definition can be seen in the figures below.

In order to perform accurate background subtraction, the known background levels must be understood. Using the cuts described in the background analysis from Kevin's note Here, the cuts that were used were as follows:

The electronic recoil band has been selected to be the 3 sigma margin as described above.

The background rates are shown in the figure below.

The average rate of approximately 0.05 DRU is in agreement with previous measurements. As is the peak at 165keV which is a results of metastable xenon after AmBe calibration. One can remove the files in which these decays occur to allow a flat BG, but this would reduce the sensitivity in all regions. The high rate at low energies is yet to be explained,

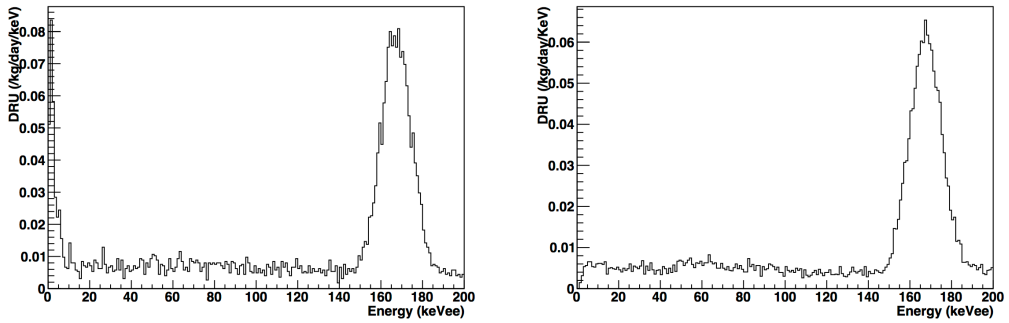

Figure 2.12: Left: Electronic Background rate given 225 days and a fiducial volume of 10kg using kevins background cuts. Right: Electronic background rate given 225 days of dark matter data and a 34kg fiducial volume with kevin's analysis cuts.

but may be a results of strange acceptance effects at low energies. The peak that can be observed in both plots is due to the decay of metastable xenon 131 that releases a gamma at 163.9keV.

Once the electronic recoil band has been adequately defined, and the cuts with their acceptances have been determined to be appropriate, One can begin to look at the remaining events cuts within the XENON100 run 10 dark matter data. The plot below shows the spectrum of events in the combined energy scale between 0-150keV with all the applied cuts from kevin defined above.

As a first order approximation, one can assume this rate to be flat at 0.005 DRU. From this rate, a limit can be calculated using the equations above. The resultant upper limit for vector and pseudoscalar bosonic superWIMPs can be seen in the figure below.

The approximated limits from Figure 2.13 are in the same order as the results published by XMASS-I. The potential for improvement, however, lies in the greater sensitive range of energies that XENON100 is able to probe.

For a complete analysis, however, new cuts must be independently implemented, and their acceptances determined in order to produce a final limit. Furthermore any new cuts that are decided to be implemented must ensure a consistent efficiency over the energy ranges required for the analysis. As many pre-defined cuts for the XENON100 axion and WIMP analyses focus on much lower energies, it may be necessary to define custom cuts for bosonic superWIMPs.

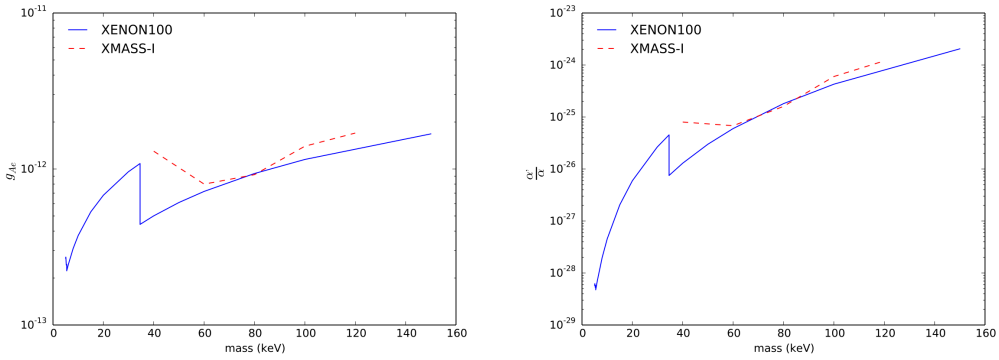

Figure 2.13: Limit of pseudoscalar (left) and vector (right) bosonic superWIMPS given from 225 of dark matter data, a fiducial volume of 34kg and assuming a constant upper limit of the rate of 0.005 events/kg/day. The red, dashed line represents the latest limits released by XMASS.

2.2.3 Cut Acceptances at High Energies

The selection of the appropriate data, and it's acceptance in the signal region is imperative for the accurate determination of a signal, or limit. In order to reach an appropriate selection of cuts, we have analysed Co and Th calibration data. Such a dataset has been used, due to the high statistics within the electronic recoil band, and as a result allows an accurate calculation of the acceptance of any applied cuts.

For the purposes of consistency, all cuts have been defined in S1 space, between 0-400 pe. Initial quality cuts applied to the data prior to the analysis demand the presence of at least one S2, fewer than 400 pe, that at least 2 PMTs observed the largest S1, that the largest S2 contains at least 150 pe. In addition, there are cuts on the requirement for single scatters, as well as an initial 48kg fiducual volume requirement.

Acceptances are calculated via the N-1 method. That is, for each bin in energy, the ratio is calculated between the total number of events with all cuts without the cut in question, normalised by the total number of events with all cuts. For the base cuts, only these were taken into account for the acceptance, and the appropriate acceptances are shown in Figure [add figure] below. All further acceptances will be calculated the the base cuts taken as a constant.

The acceptances shown in the following sections were calculated using a subset of Co and Th data, using 30 files from each. The base cuts shown below were then applied to this data to give a base selection of events.

(NbS2Peaks > 0) ,(cS1sTot[0] < 400.) , Xs1coin0, Xs2peaks2 , Xs2single3, Xs1single4, X48kg0 In order to determine the validity of these base cuts at higher energies, the N-1 acceptances of these cuts were calculated on 10 Co and 10 Th files. Since the cuts NbS2Peaks>0, cS1sTot[0]<400, Xs2peaks2 and X48kg0 are not at risk of being poorly defined at higher energies, they will not be studied here.

For the remaining cuts, the acceptances are shown in Figure 2.14.

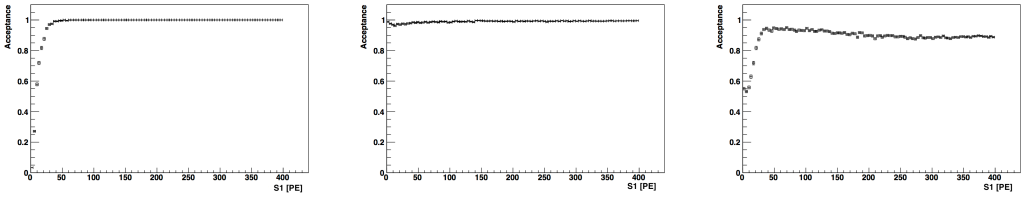

Figure 2.14: Acceptance of base cuts to be used for the superWIMP analysis as a sunction of S1. From left to right: xs1coin0, xs1single4, xs2single3.

The following cuts have been selected to be chosen for the superWIMP analysis due to their inclusion into the axion analysis. The cuts are listed below alongside their motivation as to why they would be necessary for the analysis.

- **Xsignalnoise3, Xsignalnoise4, Xentropy0**: Checks the quality of events against noise. Useful for eliminated uphysical events
- **Xs1width0**: Eliminated events with widths too large for single s1 such as double scatter.
- **xs2width0**: Eliminated events with high s2 width cause by such events as double scatters.
- **Xs2top0**: Removes the hotspots from run 10.
- **Xhighlog1**:Eliminated events where the S1 and S2 are not part of the same event.
- **XPL013_97**: Determines physicality of events based on the PMT pattern likelihood.
- **Xlownoise0**: Eliminated events with a very high S2 asymmetry, such as gas events.
- **Xs1coin2**: Demands a higher coincidence level for each noisy PMT that observed an event.

- **Xchisquares0**: Looks at the chi square of the position reconstruction between various methods, and eliminates though with high chi2 to eliminate events such as double scatters.
- **Xposrec1**: Looks at the difference in position reconstruction between various methods, and eliminates those with a high difference, to eliminate double scatters.

The following analysis will focus on a selection of these cuts that have not been verified for high energies required for the superWIMP analysis.

The purpose of Xs2chisquared0 purpose is to use the reduced chi squared of the S2 signal in order to remove events such as double scatters, that have a high chi squared. This is performed by dividing the result for the chi squared by the number of PMTs on the top array that observed the event, in order avoid energy dependence.

The cut is defined as follows

$$Xs2chisquared0 = ((S2sPosNn[0][3]/(S2sCoinTop[0] - 1)) < 7) \tag{2.8}$$

Where S2sPosNn[0][3] is the χ^2 value of the largest S2 and S2sCoinTop[0] refers to the number of PMTs in the top array that observed the same S2. The normalisation factor shown is known as the reduced chi^2 of the event. The acceptance of this cut with respect to the base cuts defined above is shown in Figure 2.15

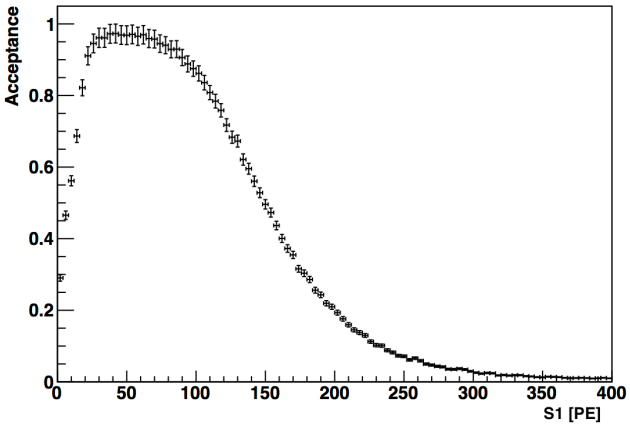

Figure 2.15: Acceptance of Xs2chisquared1 against S1 for the superWIMP analysis.

Looking into the distribution of the χ^2 and minimised χ^2 in Figure 2.16 we can see that there seems to be a second population of events at (S2sPosNn[0][3]/(S2sCoinTop[0]-1))>10 which would be removed from the Xs2chisquare0 cut.

To obtain a better idea as to what this population of events may be, the minimised χ^2 was plotted against the size of the largest S2 signal as seen by both the top and bottom PMTs separately. The result of which are shown in Figure 2.17.

This distribution shows that the saturation of PMTs causes the reduced chi squared to shift to higher values at higher energies. The dependence of the minimised chi square on energy, is therefore clearly not flat as the cut requires. Instead, due to this shifting of higher

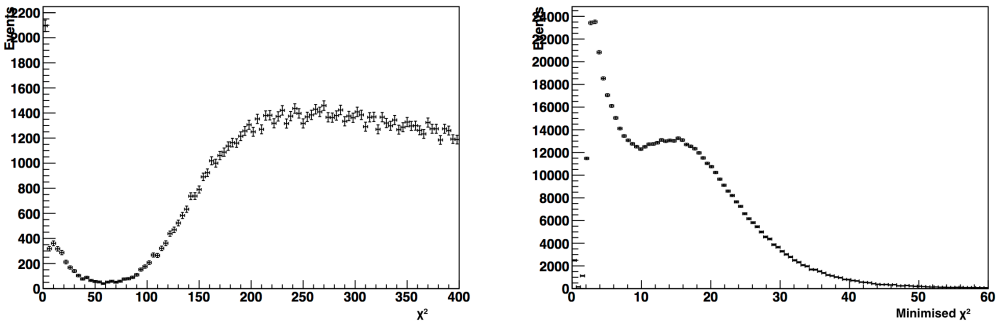

Figure 2.16: Distribution of events with basic cuts for χ^2 (left) and minimised χ^2 of the largest S2 signal in each event.

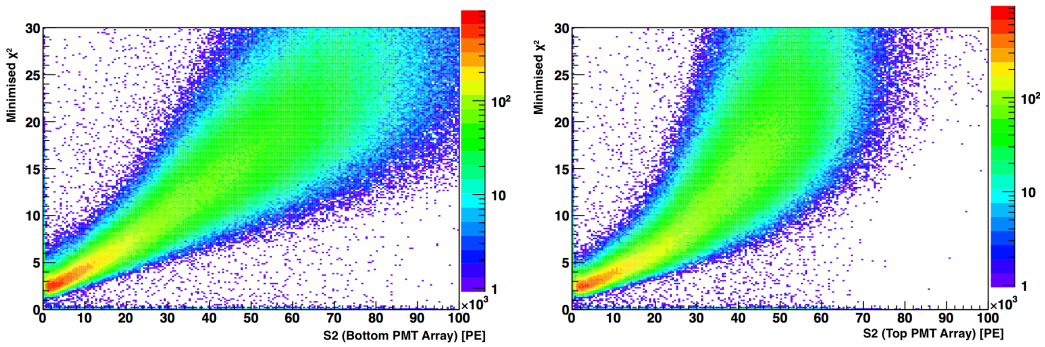

Figure 2.17: Distribution of the minimised χ^2 of the largest S2 of each event against the size of the S2 signal as seen by the top (left) and bottom (right) PMT arrays.

energy events to higher chi square values, a flat cut removes a higher fraction of events as a function of energy, as was predicted by Figure 2.15. A zoom into the lower energy region in Figure 2.18 shows how the motivation behind a flat cut works at these energies

From Figure 2.18, it becomes clear that the cut has only been specifically designed for low energies. To observe the effect of this cut on the electronic recoil band, the ratio of S2/S1 against S1 is plotted before and after applying the Xs2chisquared0 cut. The differences in the two distributions can be observed in Figure 2.19.

The excess of events at around 350pe seen in the left plot shown above are suspected to be from lead, and backscatters causing a peak at approximately the same energies. We can attempt to confirm this by looking at the distribution of events in S2 vs S1. By zooming into the higher energy region of this distribution, one can observe in Figure 2.20 the same excess as in Figure 2.19 (right), also distributed accordingly in S2.

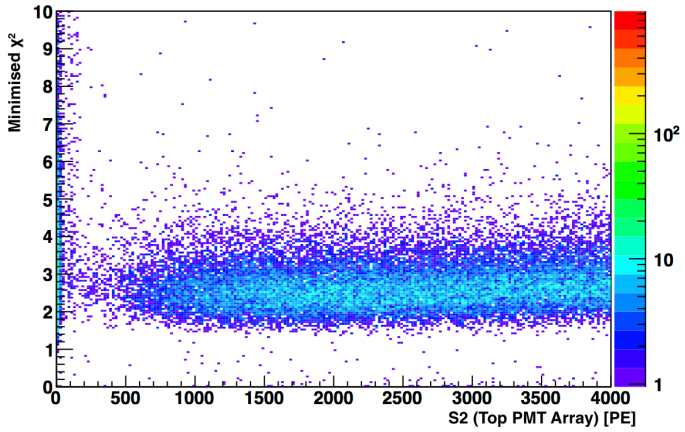

Figure 2.18: Distribution of the minimised χ^2 of the largest S2 of each event against the size of the S2 signal as seen by the top PMT arrays

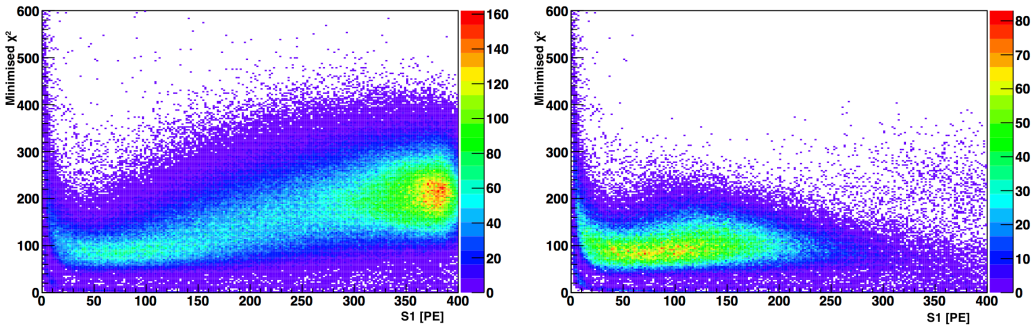

Figure 2.19: s2/s1 against s1 for events selected with only the base cuts (left) and with Xs2chisquare0 (right)

Overall, it is clear the the right hand plot in Figure 2.19 has had good events removed from the analysis. Coupled with the strong energy dependence of the minimised χ^2 on the S2 signal from figure 3, it can be clearly shown that a flat cut here is not suitable for our analysis. Since our predicted acceptance for this cut would be 99.5%, one can argue that the removal of this cut from the analysis would be more beneficial than potentially removing 0.5% of bad events from the analysis. Furthermore, other cuts in the analysis may also include the ability to remove double scatter, and could fill in the place of this cut. In this case, one must also removed the position reconstruction cut Xposrec1, as this cut is connected with Xchisquare0.

The Xs2width8 cut removes events that have widths that are too large for the size of events such as double scatters that are only recognised by the processor as having 1 S2. The

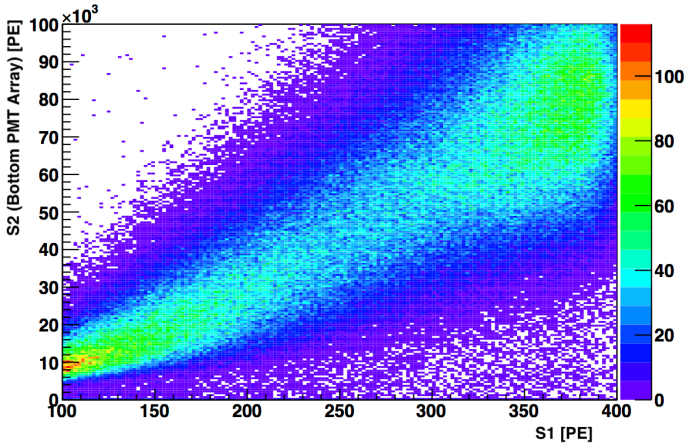

Figure 2.20: Distribution of S2 vs S1 with only a basic selection of cuts. The figure focuses on the events with S1>100, to focus on the high energy excess observed at around 350pe in S1.

elimination of potential double scatters is crucial for the analysis, and bosonic superWIMPs are only expected to interact once within the detector, if at all. Although Xs2width8 was used for the axion analysis, it had since been updated for higher energies and for electronic recoil studies. One must therefore re-run the acceptance calculation with the new cut, Xs2widther3 is shown in Figure 2.22

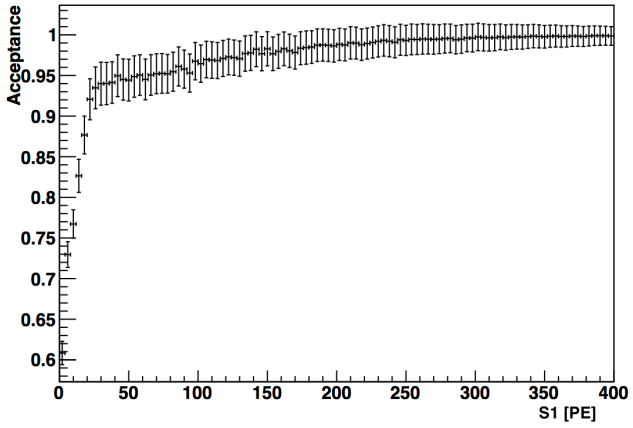

Figure 2.21: Acceptance from base cuts of Xs2widther3.

This cut, however, was defined for use on XENON100's run 12 data. In order to determine the viability of this cut for run 10, this definition is overlapped onto run 10 data. The result is shown in Figure **??**

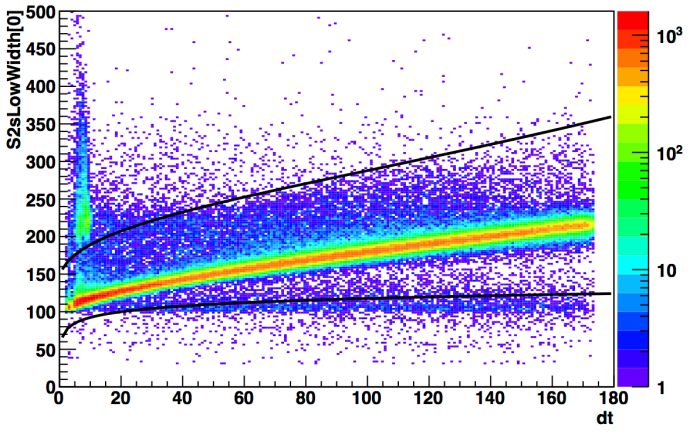

Figure 2.22: Acceptance from base cuts of Xs2widther 3.

With the implementation of these cuts, it is now possible to define the region of the electronic recoil band, as a definition of the region of interest to be used for the search for super-WIMPs. This is performed by taking the ratio of the largest S1 and S2 signal for each event, as a function of S1. A function is then defined to the center of this band, which then allows the spectrum to be flattened. This new flattened space results in the distribution observed in Figure 2.23.

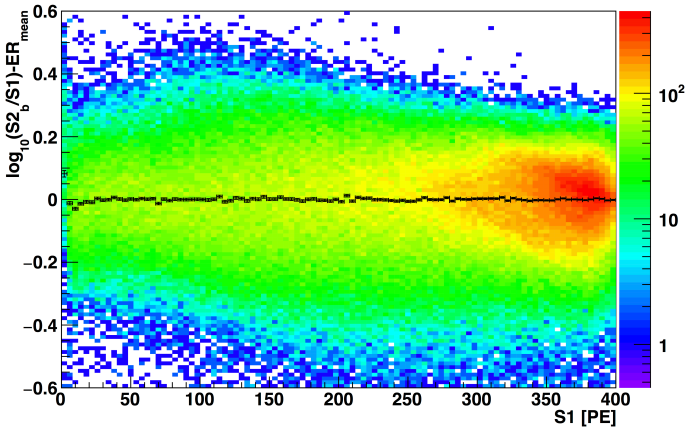

Figure 2.23: Distribution of Co and Th calibration data as a function of S1 in flattened electronic recoil band space

To constrain the region of interest, the 2 σ distribution from the mean is taken, and a flat approximation over the total energy range it taken between -0.3 and 0.3 in this flattened

space. The total acceptance, including that on the electronic recoil band is shown in Figure 2.24.

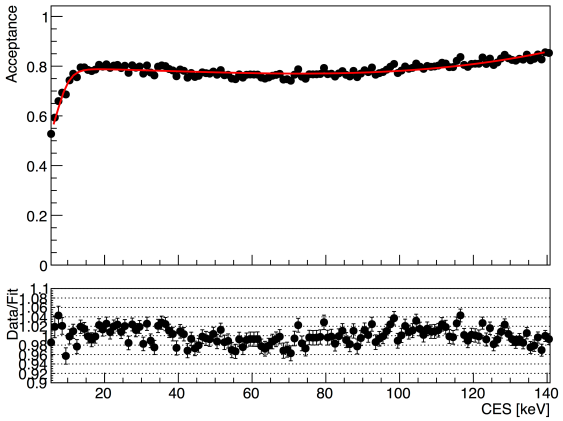

Figure 2.24: Acceptance of all applied cuts.

2.2.4 The Background Model and Signal region

A required aspect for the production of a profile likelihood limit is an accurate model describing the behaviour of the background in the RoI in the absence of a signal. To achieve this, Co and Th calibration data is used, as the electronic recoils produced from this data is expected to follow the same shape as the background, while the high statistics masks any potential signal that may arise from superWIMPs. As both Co and Th datasets are taken independently, it is imperative that, when observed separately, the shape of each are comparable. If, for any reason, the distribution of the Co and Th data do not follow the same trend, it is not possible to determine the appropriate background model, and thus cannot be used for the analysis.

Figure 2.27 compares the Co and Th data, plotted independently of one another as a function of CES.

As can be observed in Figure 2.27, there is a significant deviation from the ratio of 1.4 in the energy region below 5keV. In order to determine our signal region, the method by which the limit is intended to be determined must first be know. For the profile likelihood analysis, the signal region for a given mass of superWIMP is blinded in the region determined by the resolution of XENON100. The region +-2 sigma from the mass in question is then used to perform the profile likelihood analysis. From this, this signal region can then be determined based on the maximum and minimal energies from which the Co and Th data

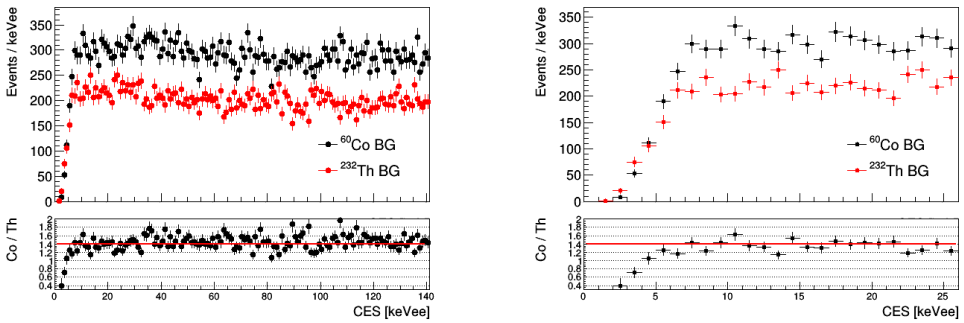

Figure 2.25: Left: Comparison between Co and Th data. Right: zoom into the low energy region.

can accurately predict the shape of the background. The minimal energy is defined at 5keV by the discrepancy between the Co and Th data, and the maximum energy is defined by the peak produced from Xe131, and lies at 140keV, as can be seen in the DM spectrum shown in Figure **??**.

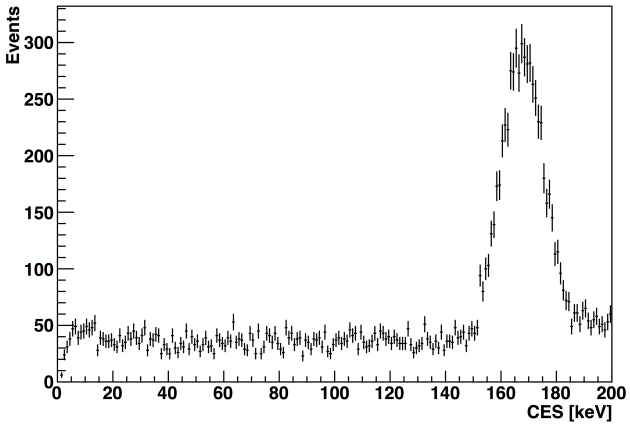

Figure 2.26: Dark matter spectrum of run 10.

Finding the energy at which the 2 sigma of the resolution does not fall below the minimal threshold, or the highest threshold then gives us the energy region we can probe. This amounts to a sensitive region between 8-125keV.

With the sensitive energies determined, the BG model can then be created. For this analysis, a modified fermi function was used to fit the data. The full equation is shown in eq. 2.9

$$BG = \frac{553.012}{1 + e^{-x - \frac{-4.69899}{1.03531}}} + -23.1283 + -.0144532 \times x^2 + 8.87476 \times 10^{-5} \times x^3 \quad (2.9)$$

The resultant fit to the Co and Th data can be observed in Figure **??**, and is the distribution to be used when forming a profile likelihood limit.

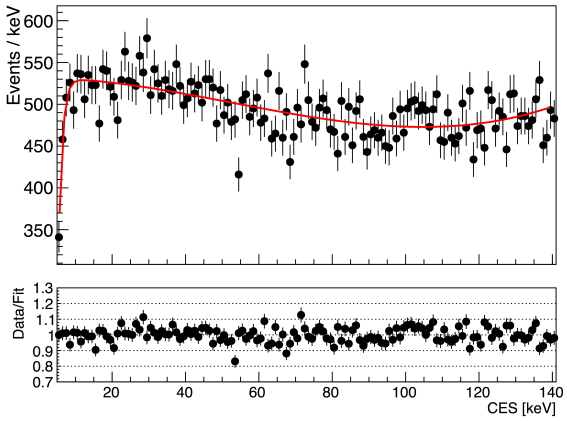

Figure 2.27: Dark matter spectrum of run 10.

2.2.5 Producing a Profile Likelihood Limit

To constrain limits on g_{aee} and $\frac{\alpha'}{\alpha}$ a profile likelihood analysis, as described [here] and [here] is employed. The full likelihood function is given by

$$\mathcal{L}((g_{Aee}, \alpha'/\alpha), N_b) = Poiss(N|N_s + N_b) \prod_{i=1}^{N} \frac{N_s f_s(CES_i) + N_b f_b(CES_i)}{Ns + Nb} \quad (2.10)$$

Where g_{Aee} and α'/α are the coupling constants for pseudo-scalar and vector super-WIMPs respectively, N is the total number of observed events N_s is the number of signal events, N_b is the number of background events, $f_b(CES)$ is background model as shown in the above sections, and $f_s(CES)$ is the signal model derived from a monochromatic peak smeared by the energy resolution of XENON100 given a rate determined by theory.

The parameters of interest are g_{aee} or $\frac{\alpha'}{\alpha}$ whereas N_b and N^{exp} are considered as nuisance parameters.

For each considered mass of bosonic superWIMP, the region $\pm 2\sigma$ as determined by the energy resolution is blinded. The events remaining in the dark matter spectrum are then used to scale the background model as shown in Figure **??**, and is used to make a prediction on the number of background events in the region of interest. The region is then unblinded and the profile likelihood calculation is performed. As this 2σ window must not exceed the sensitive energy region between 5-140 keV, the range of bosonic superWIMP masses is restricted to those between 5-125 keV.

Using the models and methods described above, a profile-likelihood limit can be set on the coupling constants for each type of superWIMP, and compared to the current limit produced by XMASS. The results are shown in the Figure 2.28, along with the 1 and 2 sigma bands of the sensitivity, and the results from XMASS-I.

As shown in figure 7, XENON100, as a result of its higher background, is unable to reduce the limit shown by XMASS for masses between 60-120. XENON100, is however, able to probe much lower masses, and extend the signal region slightly up to 125keV. This result is therefore able to prove that XENON100 has probed new parameter space for bosonic superWIMPs in extending the energy range from that previously obtained by XMASS-I.

To determine the validity of implementation of the profile likelihood limit, cross checks can be studied. One would be a study to determine if the analysis used is able to identify a signal if one were to exist. This was performed by injecting a Gaussian signal of standard deviation equal to the energy resolution of the detector to a simulated background from the background model. This was tested with a pseudo-scalar signal with a coupling constant of $g_{aee} = 1 \times 10^{-}12$. The extracted coupling constant from the profile likelihood analysis can then be subtracted from the input value to determine the accuracy of the analysis code. The result of this cross-check performed over the full range of super-WIMP masses sensitive to XENON100 can be observed in Figure **??**.

In addition to determining the ability to extract a signal, alternative methods of performing the profile-likelihood calculation can be tested. In particular, in the absence of a signal, one may take the entire dark matter data as background, and perform an analysis as such. This removed the necessity to blind the 2σ region around the tested super-WIMP mass, and thus have greater statistics to determine a limit. As a result, the statistical fluctuations observed in the limits shown in Figure 2.28 may be reduced. A plot comparing these two methods with respect to the results from XMASS-I can be observed in Figure 2.30. Although there are small discrepancies between the two methods, in which the single signal

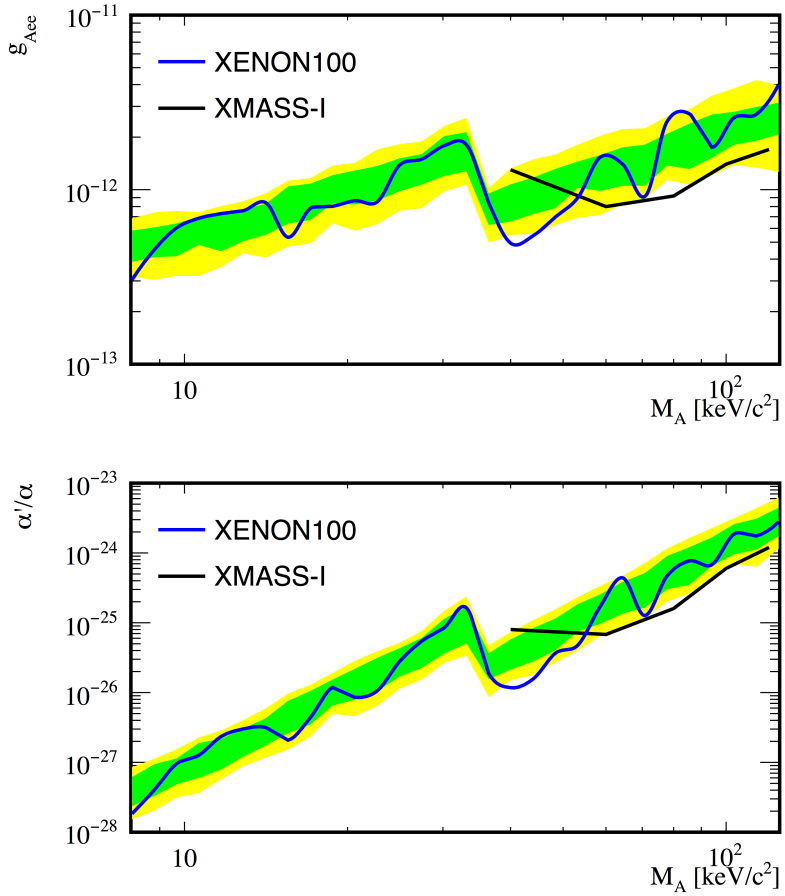

Figure 2.28: Profile likelihood limits for superWIMPs in XENON100 for peudo-scalar (top) and vector (bottom) superWIMPs in comparison to the latest results from XMASS. The green and yellow bands denote the 1 and 2 sigma regions from the sensitivity.

region method provides a slightly better limit in the center of the sensitive mass range, the differences are statistically consistent with one another. As a result, both methods are valid, and the statistical method holds in both cases.

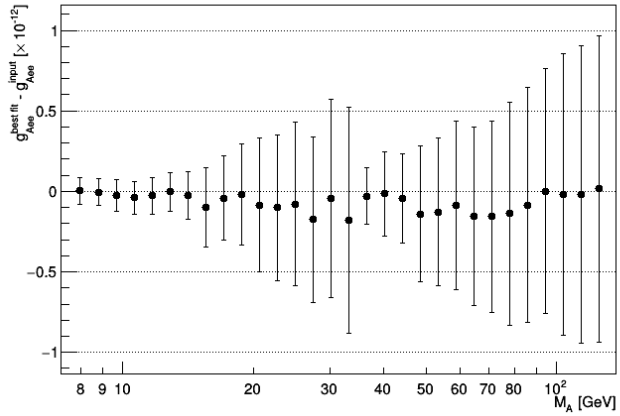

Figure 2.29: Results from injecting a pseudo-scalar super-WIMP signal with $g_{aee} = 1 \times 10^{-}12$. The applied signal is subtracted from that measured from the profile-likelihood method and shown as a function of the super-WIMP mass.

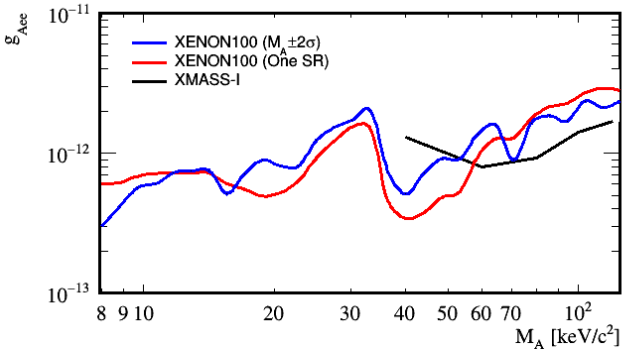

Figure 2.30: Comparison between a profile-likelihood limit calculated by blinding the 2σ signal region defined by the energy resolution of XENON100 (blue curve), and a method by which the entire background is used within the calculation (red curve). These results have been performed for the pseudo-scalar super-WIMP, and are placed in context with the latest results from XMASS.

2.3 Conclusion

In summary of the work performed on the XENON100 experiment, as detailed in the above chapter, an analysis was performed to determine which PMTs are the lead cause of noise in run 12. By observing the number of events for each detector within AmBe calibration data, and requiring events less than 3pe, one can study the levels of noise by looking for PMTs with an excess of events. After implementing increasing levels of cuts on the coincidence level of observed S1 signals, a list of high-noise PMTs could be established. Namely, events in which an S1 is observed on PMTs 165, 166 or 169 must have the coincidence level for a signal increased to account for the elevated probability of noise in these detectors. It was then studied if there was an alternative, less manual cut that may be implemented in place of this cut on noisy PMTs. It was found that by using a cut on the entropy of events, one could effectively remove noisy events reliably without the need to a constant re-definition of a noisy cut for each run. These two systems were then tested on dark matter data, whereby an increased level of noise is more significant, and it was found that both the noisy PMT cut, and the entropy cut performed equally well. Due to the flexible nature of the entropy method, however, it was decided that this should be used in replacement for future runs.

An analysis to search for the presence of bosonic superWIMPs in XENON100 was also performed. These theorised particles, in the form of a pseudo-scalar and vector type, are predicted to deposit a monochromatic peak within the detector at the approximate rest mass of the particle via the axio-electric effect. The resultant signal is therefore expected within the electronic recoil band. As such, a necessary analysis into providing a comprehensive list of cuts designed for the electronic recoil band, for electronic recoils up to 200keV was performed. This concluded with a list of cuts, effective over the entire energy range of interest, and developed specifically for the electronic recoil band that were to be used for the remainder of the analysis. The model to the acceptance of the final cuts was then made using a modified fermi function.

To provide a profile likelihood limit for bosonic superWIMPs, the background model and the sensitive energy regions were studied. The background was modelled based on Co and Th calibration data. Due to a discrepancy in shape between the Co and Th spectra, energies below 5keV could not be considered for this analysis. Furthermore, as a result of the peak resulting from the decay of Xe131, an upper limit on the sensitive energy was places at 140keV. Placing a limit at a given energy via profile likelihood is performed by blinding a region $\pm 2\sigma$ of the resolution of XENON100 from the energy of interest. As a result, the range of energies that may be probed for this analysis that do not exceed the upper and

lower limits of energy when this blinded region is taken into account is between 8-125 keV. The Co and Th data is then modelled between 5-140 keV, again using the modified fermi function as to be consistent with the acceptance model.

With the cuts, acceptance, and background model defined, the profile likelihood calibration was performed. g_{Aee} and α'/α are the coupling constants for pseudo-scalar and vector superWIMPs respectively, and represent the parameters of interest that limits will be set upon. The number of background and the number of signal events are then implemented as nuisance parameters. The final limit lies above the current best limits from XMASS-I in the energy region between 60-120 keV. The remaining energies are new limits for the bosonic super-WIMP model.

3

XENON1T

Using a total of 3.1 tonnes of liquid Xenon (LXe) in a cylindrical cryostat, shown in Figure **??** (top), XENON1T intends to expand upon the success of XENON100 by producing a detector two orders of magnitude more sensitive than its predecessor, placing a predicted limit on the spin-independent cross section at $2\text{x}10^{-47}$ cm^2 at 30 GeV. This would then be the lowest limit set for WIMPs.

3.1 The XENON1T TPC

At the heart of XENON1T, encompassing the sensitive volume and allowing S2 signals to be detected lies the TPC. A schematic displaying the basic dimensions and components of the TPC can be observed in Figure **??**

Each component of the TPC has been rigorously developed with low radioactive materials, and then screened to quantify the activity of every component. Capping the TPC, and laying above the top PMT array is a stainless-steel diving bell. This structure, 1 m in diameter, and 0.6 m in height, is responsible for maintaining the liquid level of XENON1T TPC. A full, detailed description of the workings, and testing of this component can be seen in section 3.5.

Below the upper PMT array lies the uppermost part of the electric field cage. The purpose of this cage is to apply a uniform electric field throughout the TPC, to ensure the

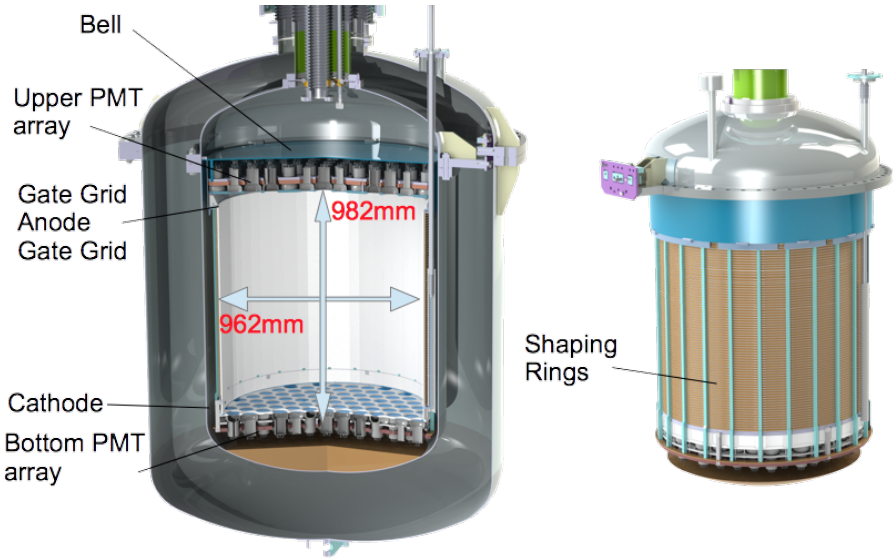

Figure 3.1: Schematic representation of the XENON1T experiment, with the basic dimensions and component that comprise the time-projection chamber.

reliable drifting of electrons produced from the initial interaction. The desired electric field goal is 1 kV/cm, that will result in an electric drift velocity of approximately $2\,mm/\mu s$.

The top of the cage consists of 3 hexagonal meshes. Two are named as gate meshes, and placed on ground voltage. The top gate mesh, also known as the screening mesh, consists of wires $178\,\mu m$ in diameter with a cell opening size of 5.7 mm. The lower gate mesh contains wires of diameter $127\,\mu m$, and a cell opening of 3.5 mm. The anode lies in between the two gate meshes, is made of $178\,\mu m$ diameter wires, with a cell opening of 3.5 mm. Enclosing the bottom of the field cage are two more meshes. These are the cathode, and bottom screening mesh. Both consist of parallel wires, $216\,\mu m$ in diameter and with a pitch of 7.5 mm. A summary of the information of all of these meshes can be observed in Table 3.1.

Connecting the cathode to the anode, through a chain of resistors are the shaping rings. Arguably the most important structure in the field cage for the electric field uniformity, the shaping rings apply a radial electric field to counter that produced from the anode and cathode. Their design and spacing has been specifically chosen based on numerous simulation, some which will be detailed in the sections below. 74 of these, made from low-activity OFHC copper are evenly spaced between the cathode and upper meshes. Between the shaping rings, two redundant chains of $5\,G\Omega$ resistors are placed. A resistance twice this is then used to connect the cathode and first shaping ring.

TPC Electrode	**Type**	**Material**	**Wire Diameter** [μm]	**Pitch/Cell opening**	**Transparency**
Top Screen	hex etched	stainless steel	178	5.7	95.3
Anode	hex etched	stainless steel	178	3.6	90.4
Gate	hex etched	stainless steel	127	3.5	93.1
Cathode	parallel wires	Au-plated stainless steel	216	7.5	97.1
Bottom Screen	hex etched	Au-plated stainless steel	216	7.5	97.1

Table 3.1: Specifications of the five TPC electrodes which were optimized for S1 light collection and high voltage stability. Both wire electrodes feature one single wire installed perpendicular to all others to minimize elastic deformations of the frames.

On the inside of the shaping rings, 48 high quality PTFE reflection panels are positioned, being supported by 24 PTFE pillars, both running the length of the TPC. The purpose of the reflector panels is to incrase the light collection efficiency of events within the sensitive volume of the detector.

3.2 Field Simulation of the XENON1T TPC

This section will detail the electric field simulations performed to improve upon, and increase the uniformity of the the electric field within the XENON1T TPC by altering the design of the field cage.

3.2.1 COMSOL and the Finite Element Method

The finite element method is a numerical technique for finding a solution in a domain with boundary values. Therefore we have to discretize the whole computational domain into N finite elements to obtain the electric potential in the inner region. Common elements are tetrahedral in three dimensions and triangles in two dimensions. The electric potential inside one triangle can be described by the three corner potentials of the triangle: $\phi 1, \phi 2, \phi 3$.

For the case we have only one element the corner potentials are well known. The electric potential inside the triangle is assumed with a linear function by the following equation:

$$\phi_{tri}(x, y) = c_1 + c_2 \cdot x + c_3 \cdot y \tag{3.1}$$

Where x and y sare the coordinates and c_1, c_2 and c_3 are coefficients. The coefficients c_1, $c_2$and c_3 can be expressed by the the corner potentials of the triangle. In this way we obtain a linear equation system for the electric potential inside the triangle:

$$\begin{bmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \end{bmatrix} = \begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ 1 & x_3 & y_3 \end{bmatrix} \times \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} \tag{3.2}$$

After solving for the coefficients we get the potential in one element in terms of the corner potentials. Given by the following formula

$$\phi_{tri}(x, y) = N_1\phi_1 + N_2\phi_2 + N_3\phi_3 \tag{3.3}$$

Where N_1, N_2 and N_3 are defined interpolation functions from the solution of 1.2 If we have more than one triangle, we have to set up a global matrix over an additional condition.

3.2.2 Optimising the Simulation

In the context of performing the electric field simulations for the XENON1T TPC, there are three primary methods that must be explored to determine the most viable solution. Those are a flat 2D simulation whereby only the cross section of the TPC is taken and simulated upon, a 2D axisymmetric model where the geometry is assumed to be cylindrically symmetric, and thus a 3D simulation can be assumed from a 2D cross-section, and finally a full 3D representation of the chamber. Each of these methods is a compromise between simulation and accuracy, and it is therefore of great importance to quantify the accuracy of each method with respect to the full 3D simulation.

The geometry for these simulations is an approximation based on a general layout of a TPC containing 2 PMTs and is not a representation of an existing TPC. The model cross-section is shown in Figure 1 and is identical for both the 2-dimensional axisymmetric and

3-dimensional models. Although both models are provided with the same cross section, the differing geometries require a different coordinate system. In these simulations the 'z' coordinate of both models are identical, and the radial component of the two dimensional simulation corresponds to the y-componend of the three dimensional model. That is to say that the direction of the mesh wires propagates perpendicularly to the y axis in this coordinate system.

The importance in determining differences between these models lies in the effects of the simulated mesh layouts. An axisymmetric model allows a greater simulated resolution, however each mesh is modelled as concentric circles. A full 3 dimensional model provides a much greater geometric accuracy, however computing memory restricts the simulation to a lower resolution and longer computation time.

Each model contains the same object dimensions, and positions are defined from the middle of the central chamber. Wires have been given a 0.8mm diameter with a 3mm mesh spacing and a nominal drift field of 1kV/m. The cathode is placed at -3kV, and the anode at 5kV and are placed at z=-15mm and z=20mm respectively. Four shaping rings 5mm in length and 2mm in width are added, and equally spaced within the length of the central TPC chamber. The two PMTs are positioned at either end of the model are placed at a 1kV potential, and the entire model is then ecompassed via a grounded cryostat vessel. All materials, including liquid and gaseous xenon, copper, PTFE, and structural steel for the surrounding structure have been accurately implemented into the simulation. The 2D cross-section which is then applied to every version of the geometry is shown in shown in Figure 3.2

Plots comparing the radial and normal components of the electric field [V/m] between then two models are shown below. Results are taken as a function of both TPC length and radius. The average normal field within the 2 dimensional model was measured to be 1.01×10^5 V/m while the full 3 dimensional model simulated an average field of 0.99×10^5 V/m, a percentage difference of 2%. In an ideal model, the expected average field within the central TPC chamber would be 1×10^5 V/m by design. The electric field for the full simulation of the 2D axisymmetric and 3D model can be seen in Figure 3.3

Figure 3.4 illustrates the differences between the 3 model types along the vertical z-axis of the TPC at two positions along the radial coordinate. Comparison between the geometry of the two 2D simulation show that the axisymmetric model requires half of the full cross section to simulate, the 2D model requires the full cross section. The primary advantage in the 2D simulation is that any specific cross section which may differ from cylindrical symmetry may be taken, however, only the axisymmetric model is simulated in three di-

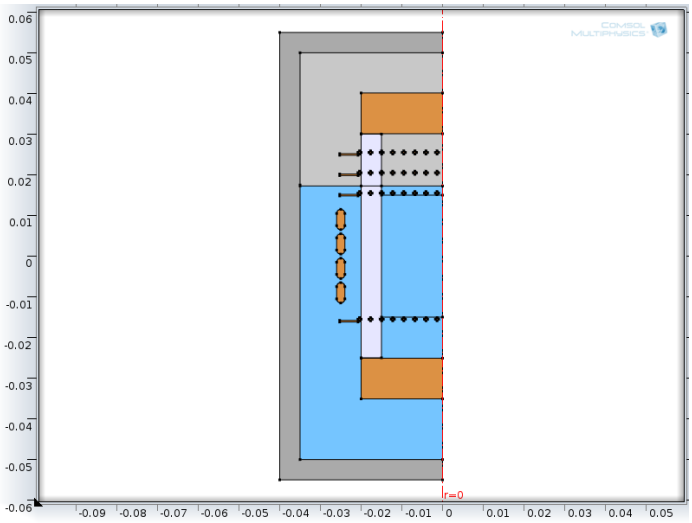

Figure 3.2: 2-dimensional cross section of a small, fictional TPC.

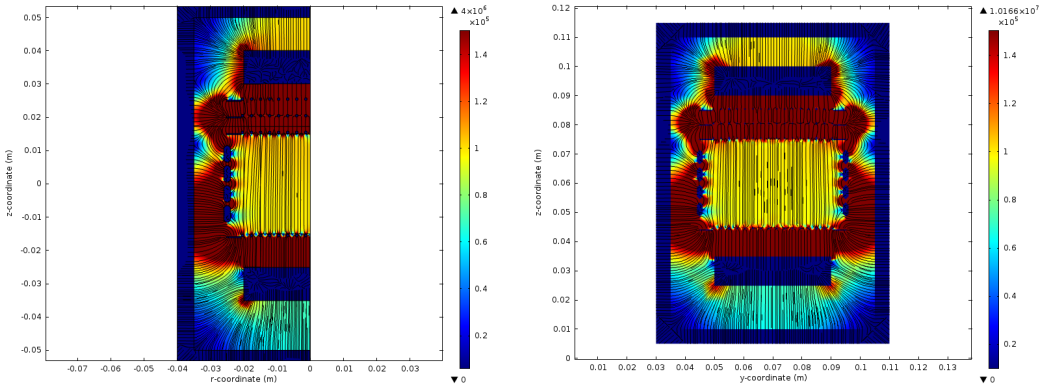

Figure 3.3: Electric field cross section simulation for a 2D axisymmetric model (left) and 3D (right) along the r/y-z plane.

mensions, and thus the purpose of this investigation is to determine how these differences affect electric field results. It may be observed in this figure that the 2D- axisymmetric and 3D models are in good agreement, within 0.1% of each other along the centre of the drift field. Although the 2D simulation shows a similar trend to the alternative simulations, a consistent discrepancy in the slope is observed, leading to the impression of a more uniform field.

The electric field along the radial axis for the three simulation types can be observed in Figures 3.5 and 3.6 as a function of the normal electric field, and radial electric field respectively. Both the 2D and 2D axisymmetric simulations demonstrate an agreement

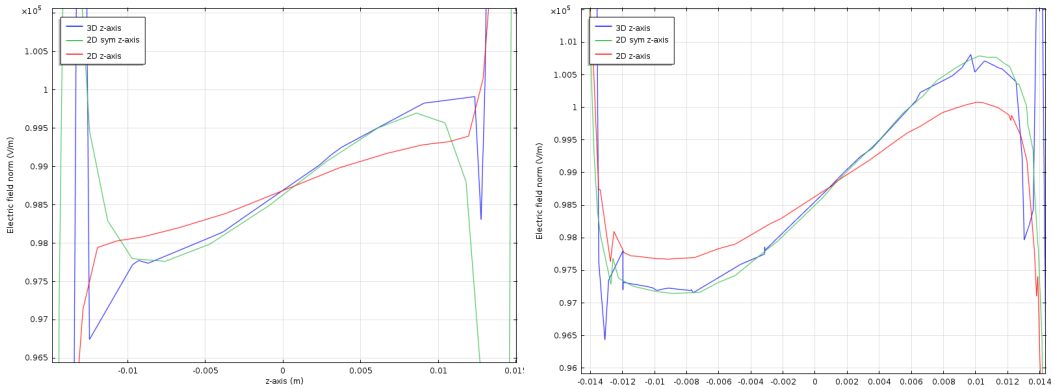

Figure 3.4: Normal electric field value for a 2D, 2D axisymmetric 3D simulation between the cathode and ground mesh at r=0 (left) and r=10mm (right)

within 1% of the normal electric field value, however the 2D-axisymmetric simulation is observed to have a more consistent gradient with respect to the 3D model.

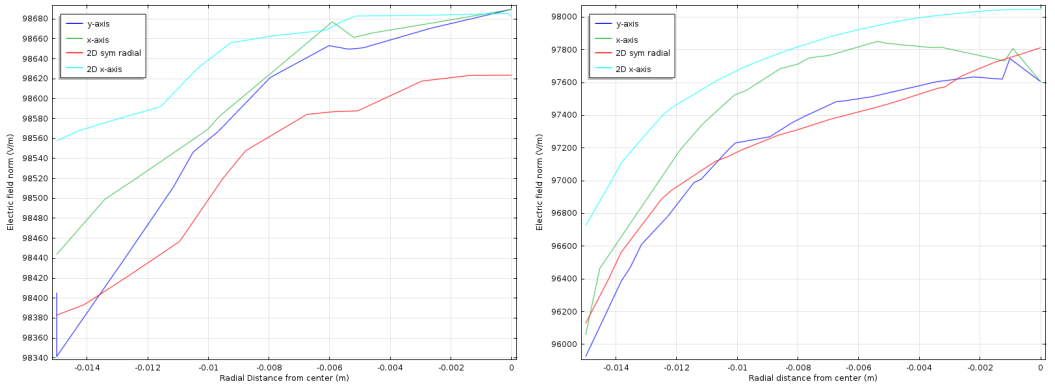

Figure 3.5: Normal electric field value for a 2D, 2D axisymmetric 3D simulation between the cathode and ground mesh at r=0 (left) and r=10mm (right)

To conclude, each simulation type demonstrated in this section agree well with one another within 2% of the full normal electric field. The full 2D simulation however, produces results of a more uniform electric field in almost all cases analysed above. As an alternative to the full 3-dimensional simulation, however, a greater agreement may be observed with an axisymmetric model than a full 2D-cross section. A a result, to improve upon simulation time, and computing power required, the 2D axisymmetric model will be employed for all future simulations.

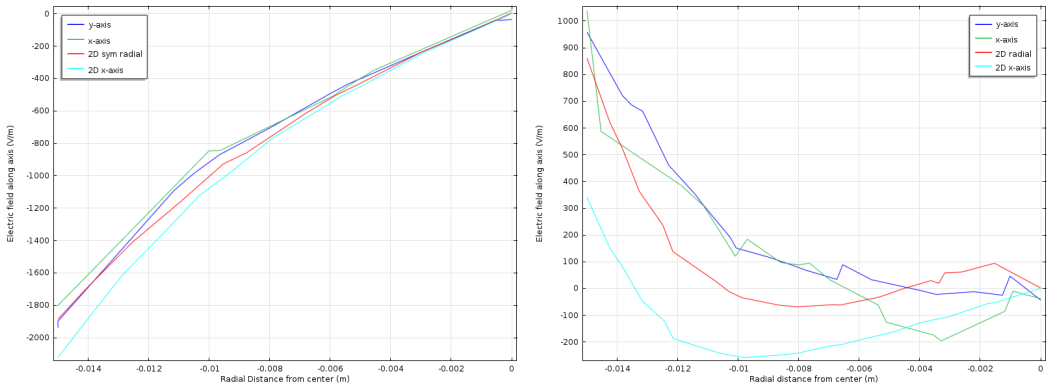

Figure 3.6: Radial electric field value for a 2D, 2D axisymmetric 3D along the radius of central chamber at z=0 (left) and z=-15mm (right)

3.2.3 Variable Parameters of the Simulation

3.2.4 Creating a Standard for Simulation Comparison

For long-term comparison between electric field simulation models that may performed via different methods or institutes within the collaboration, a universal standard must be developed to define regions of interest within the TPC that must be analysed in order for a fair contrast to be made with previous results.

In order to determine how such a standard must be defined, one must first recognise the areas within the TPC that are of the greatest interest, and importance. Immediately, the field at the center of the TPC, both extending vertically and horizontally must be considered. The uniformity along these lines allows for a clear understanding of the uniformity of the electric field along the entire length and width of the TPC.

As the purpose of these simulations is to develop a design to produce the most uniform electric field, regions within the detector must also be define where the value of this electric field is expected to deviate most from the mean. To this extent, three lines, and two areas of the TPC are considered. Those are the regions just below the gate, just above the cathode, as well as close the the side walls of the TPC. Further to this the regions at the top and bottom of the chamber where the gate and cathode meet with the walls must also be taken into account. In order to remove small-scale effect caused by the positioning of the electric field shaping rings, or the wires of the gate and anode, the lines to be standardised are placed horizontally 50 mm below the gate, 50 mm above the cathode, and similarly 50 mm away from the walls of the TPC. The regions enclosed within the intersections of these lines at the top and bottom of the chamber will also be standardised as an area of interest.

Region of interest	**Purpose**
Radially central, extending TPC in the z direction	Provides information as to the uniformity across the direction of drift.
Line centered along the z direction, extending radially across the TPC	radial drift of electrons, useful for the field map can be determined.
Line 50 mm from the side walls, extending along the z dimension along the length of the TPC	Measurement of uniformity along sensitive region near the shaping rings.
Line 50 mm above the cathode extending radially	Measurement of the uniformity in the region close to the cathode.
Line 50 mm below the gate grid extending radially	Measurement of the uniformity in the region close the the gate.
Region around the intersection between the TPC wall and cathode	High electric field due to the presence of the first shaping ring, and the cathode
Region around the intersection between the TPC wall and the gate	High electric field due to the presence of the last shaping ring, and the grounded gate.

Table 3.2: Summary of Standardised lines and regions for electric field simulations.

A summary of the standardised analytical lines and regions can be observed in Table 3.2 below.

3.2.5 The Geometrical and Electric potential Models

Using a Using a 2-dimensional cylindrical symmetry it is possible to simulate the entire TPC by building half the cross-section of the full model. The base of the central chamber of the TPC is defined to be at z=0, and the radial coordinate is centered on the full model. All model descriptions below will assume this generalisation.

The default model of XENON1T used is shown in the Figure 3.7. A more detailed view of the anode and cathode regions is then shown in Figure 3.8. The outer dimensions extend from z=-230mm to z=1300mm and r=0mm to r=540mm with the central TPC chamber positioned from z=0mm to z=983.9mm, the cathode mesh has been placed just below and a ground mesh directly above this. Five copper meshes are located at coordinates z=-50mm z=-1mm z=984mm, z=989mm, z=994mm. The three uppermost meshes are fixated via a the structure shown in Figure 2 which is encompassed in PTFE. Mesh spacing is 5mm for

all but the anode and the ground mesh directly above the central chamber. These have a spacing of 2.5mm and are phase shifted by 180 degrees to each other. All mesh wires have a diameter of 0.25mm. The liquid-gas interface is located at z=987mm, between the anode and ground mesh. PTFE is placed along the outer vertical perimeter of the central TPC chamber with a thickness of 5mm, located 1mm from the chamber wall and resting directly against the shaping rings. The shaping rings themselves are 20mm in length and 5mm in width with rounded edges. 39 of these rings are spaced 5mm apart equally between the cathode and first ground mesh. The PMTs, for the purposes of the simulation, are modeled as two solid copper objects 50mm in height and 4875mm in radius positioned at z=1075mm and z=-65mm. Figure 3 shows the cathode holding structure. The current design measure 11mm in with and 15mm in height, including the 5mm wide protrusion. It has been positioned to be as close as possible to the radius of the cathode rings.

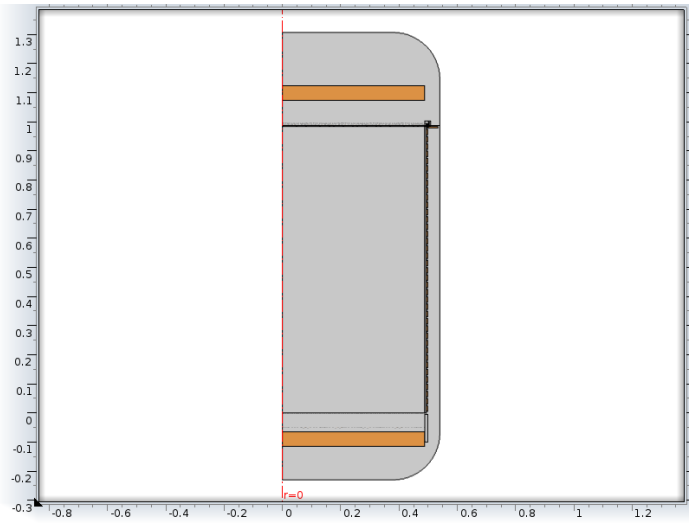

Figure 3.7: Cross-section of base XENON1T simulation model.

The geometry as shown in the figures above will be known as the version 0 of the geometry. All other versions of the geometry will be based on this.

The following is a full description of each geometry to be tested to determine the most optimal setup:

- **Version 1**: The primary difference to version 0 lies within the distribution and size of the shaping rings. They have been altered such that both the spacing between the rings, and the size of the rings themselves have been halved. They now measure 10mm in length, 5mm in width and have a spacing of 2.5mm. With the decreased

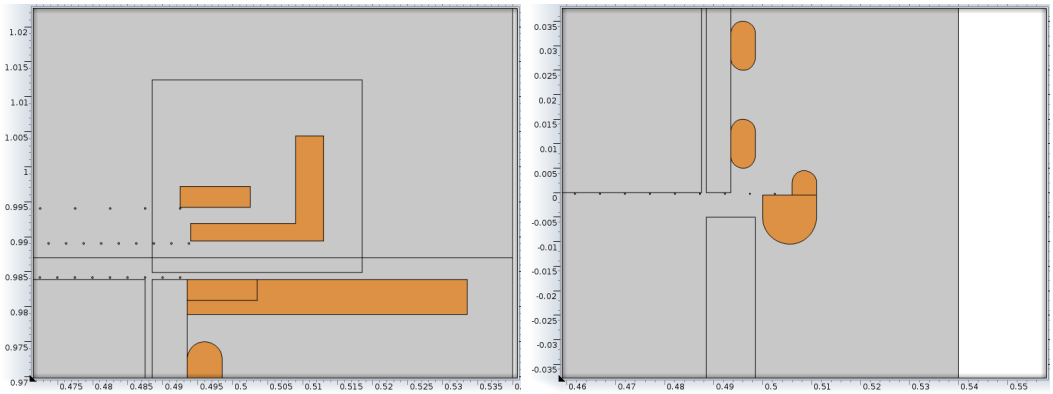

Figure 3.8: Detailed view of the COMSOL simulation geometry in the region of the anode (left) and cathode (right).

spacing and size of the shaping rings, 78 are able to be placed in this version where 39 were in Version 0.

- **Version 2**: Shaping rings have been altered to be 10mm in length as in version 1. The ring positioning, however, has remained unchanged from version 0 thus. The number of rings, therefore, remain unchanged and the ring spacing has become 15mm.

- **Version 3**: Shaping rings have been altered to be 10mm in height, with spacing between them of 5mm. The total number of shaping rings is 65, evenly spaced between the cathode and ground mesh. Variations on this geometry are defined by the spacing between shaping rings. For example, simulation 3-0 10mm refers to geometry version 3 using a 10mm spacing between shaping rings.

- **Version 4**: The upper and lower edges of the shaping rings have been altered from their semi-circular shape to a more "flat" configuration with rounded edges instead. Shaping ring size stands at 10mm, but spacing between them is at 2mm. That is to say, a center to center distance of 12mm with 81 shaping rings evenly spaces between the cathode and ground mesh.

There are two models of the electric potential that may be coupled with each of the geometrical simulations. These are described below

- **Version 0**: The base model of the electric field simulation contains the general outline from which other configurations may be based. The outer structure of the model is grounded, so too are the lowest mesh as well as the two meshes above and below the

anode, which is placed at 5kV. For the holding structure shown in Figure 2, all copper objects are placed at the potential of the mesh to which they connect. The cathode is given a -100kV potential, from which the electric potential of the shaping rings are defined. The step in potential between the shaping ring is based on the number of rings in the geometry and the cathode potential. For example, for the geometrical version 0, there are 39 shaping rings before the ground mesh. The ring directly above the cathode (assuming this is at -100kV) would therefore be placed at -97.5V, then -95kV etc such that the ring before the ground mesh is -2.5kV. This is to simulate uniform resistors between shaping rings.

- **Version 1**: This setup is identical to Version 0, except for the potential between the cathode and the first shaping ring. This particular configuration places a virtual resistor that is larger than in Version 0 solely between the cathode and first shaping ring. Subsequent potential differences between the rings are evenly distributed to the ground mesh. This is simulated by considering a cathode at a higher or lower voltage than the actual (-100kV), and distributing the potentials in an identical fashion as outlined in Version 0. This can be defined by stating the virtual difference of the cathode from the actual electric potential.

3.2.6 Notation for Comparing Simulations

The purpose of this investigation is to accurately determine the distribution and uniformity of the electric field within XENON1T using the electrostatics module of COMSOL. Alterations to the geometry and electric potential have been implemented to observe further effects. Results of electric field simulations using the COMSOL electrostatics module are outlined below. The labelling system used is to differ between the multiple versions of geometry and electric potential layouts. The first number in this labelling system represents the version number of the geometrical configuration, a "-" then divides between the version number of the electric field configuration. For example, results of simulation 0-1 would be the results using geometry version 0 and electric potential version 1. This notation can be altered with additional information specific to a geometrical or electric field configuration, such as specifying the shaping ring distance, size, or potential difference between the cathode and first dynode. This notation is consistently used throughout the following sections.

3.2.7 Determining the Optimal TPC Design

In order to methodically determine the Optimal TPC design for XENON1T, each geometrical and electric field version combination must be explored. Each comparison between combinations must be analysed to determine the improved model which will then become the new standard from which future simulations must be compared. In order to begin the process of optimisation, the base model, version 0-0 is simulated in order to find the base electric field uniformity.

Figure 3.9 shows the normal electric field distribution within XENON1T with simulation version 0-0. This simulation is to be the basis by which future simulations will be improved upon. Field lines have been added and distributed evenly over the model. The colour scale has been altered to produce greater detail within the drift field region. The greatest deviations from the mean electric field value of 99.9kV/m are found at the upper and lower corners of the central TPC chamber in the region where the side walls intersect with the cathode and gate. This is due to the presence of the shaping rings that have are placed on a different potential than the surrounding field cage.

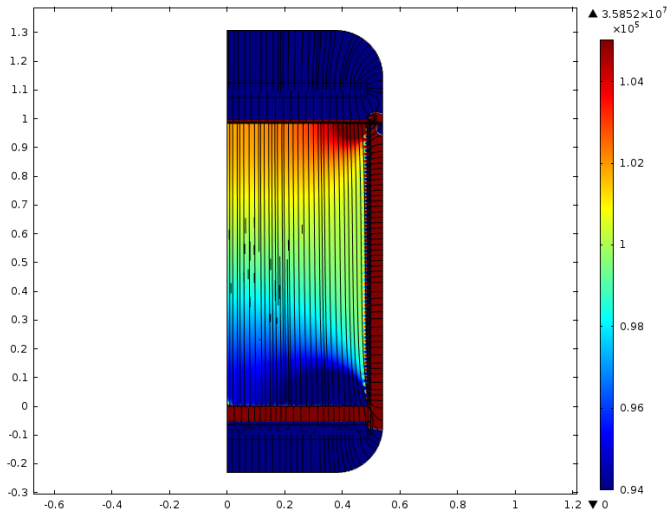

Figure 3.9: Normal electric field simulation of the base design for XENON1T.

By adding a virtual resistor to the space between the cathode and first shaping ring equivalent of the 2750 V potential difference results in the comparison shown in Figure 3.10 between the base 0-0 simulation and 0-1 with this resistor.

From a visual comparison of the two it is clear that the introduction of a higher potential between the cathode and first resistor results in a greatly improved uniformity within the

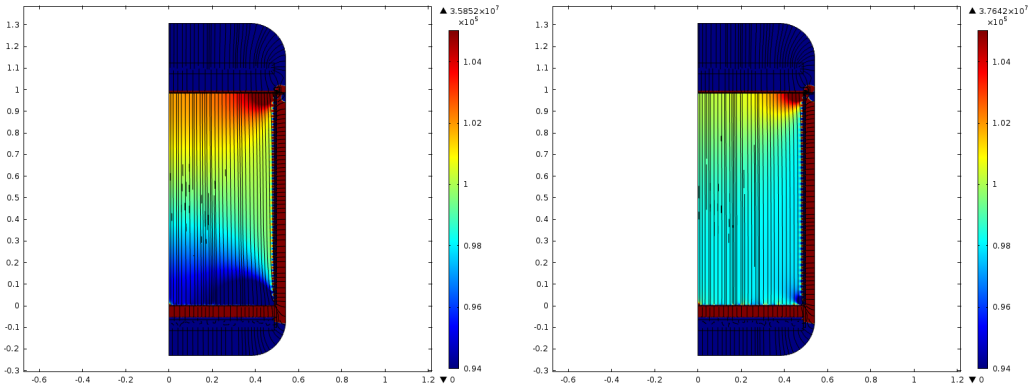

Figure 3.10: Normal electric field simulation of the base design for XENON1T (left), and the same geometrical design with an added electric potential of 2750 V.

sensitive volume of the chamber. The 0-1 simulation must therefore be the new base with which to compare with future simulations. Prior to this, however, one must determine which of the geometrical versions of the TPC result in the most uniform field, after which one may apply version 1 of the electric field.

Results below demonstrate the effect of smaller shaping rings using the same potential distribution between them. Smaller shaping rings can be seen to provide a smaller variation by a factor depending on the position of the analytical line. As may be expected from introducing a larger gap between shaping rings, the leak-through of the electric field is greater than in other geometrical configurations at approximately 25mm from the central TPC chamber wall. This is apposed to 10mm from version 0-0. Geometry version 3-0, using 10mm shaping rings with 5mm spacings, are also plotted alongside the alternative geometries. It was also decided to vary the distance between 10mm shaping rings to observe the effects not only of the size of the rings, but also the spacing. For this reason, an additional geometry entitled "3-0 10mm" has been created using 49 10mm shaping rings of separation 10mm.

The average normal electric field for 0-0, 2-0, 3-0 and 3-0 10mm over the entire central chamber is 99820 V/m, 99780 V/m, 99850 V/m, and 99793 V/m respectively. A comparison of the electric field norm can be observed in Figure 3.11.

Figure 3.12 shows a comparison between the simulations 3-0 and 3-0 10 mm by observing the differences along the z dimension at both at the center of the TPC, and then again 50 mm from the edge of the side walls. From this figure, it is clear that simulation 2-0 is able to maintain a more uniform electric field along the full length of the TPC. This

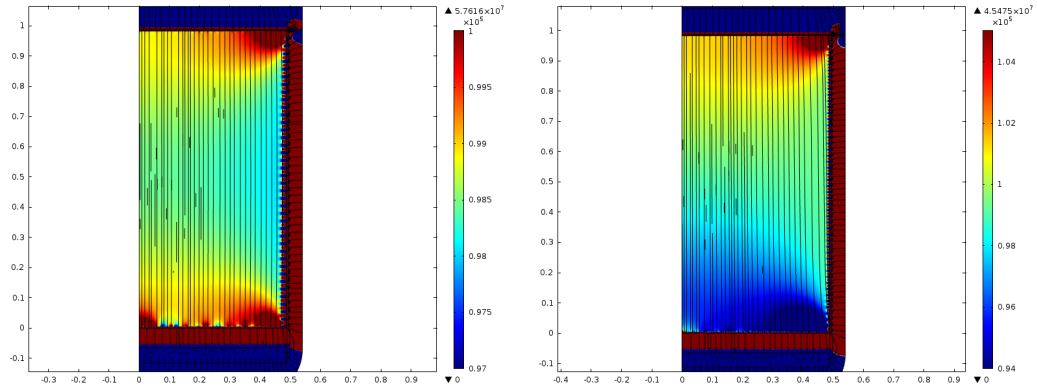

Figure 3.11: Cross section displaying the normal electric field within XENON1T for verision 2-0 (left) and 3-0 (right).

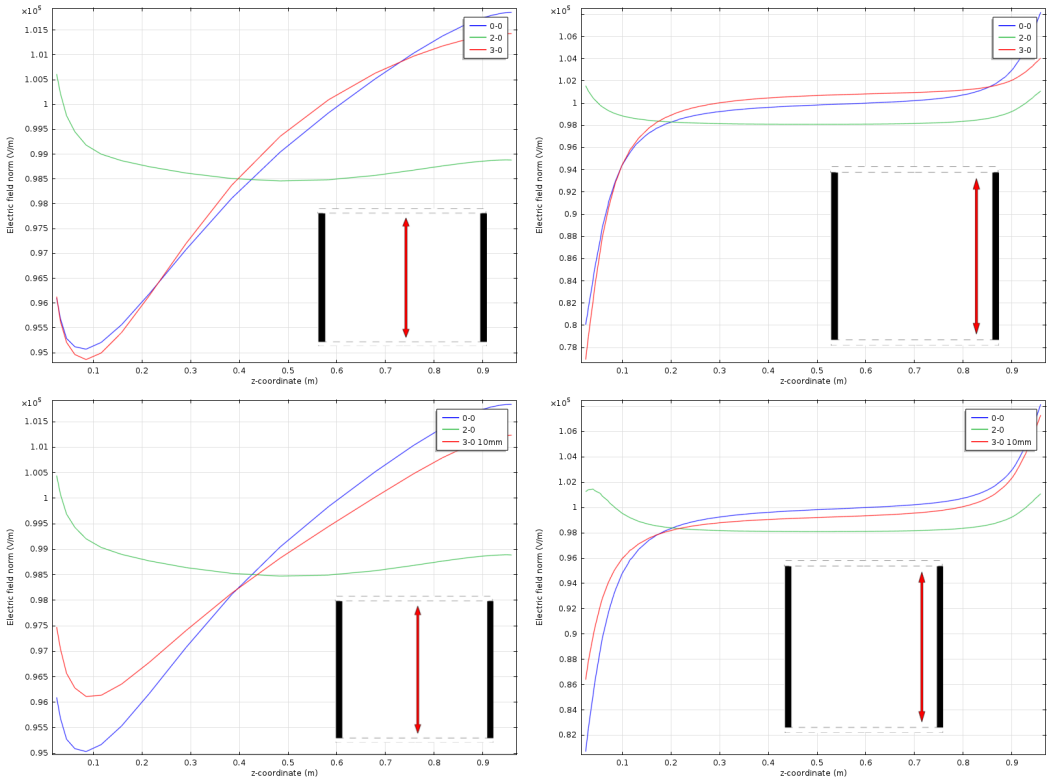

Figure 3.12: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right).

conclusion is supported in Figures 3.13 and 3.14 that investigate the uniformity of the normal and radial electric field along the radial dimension of the TPC respectively.

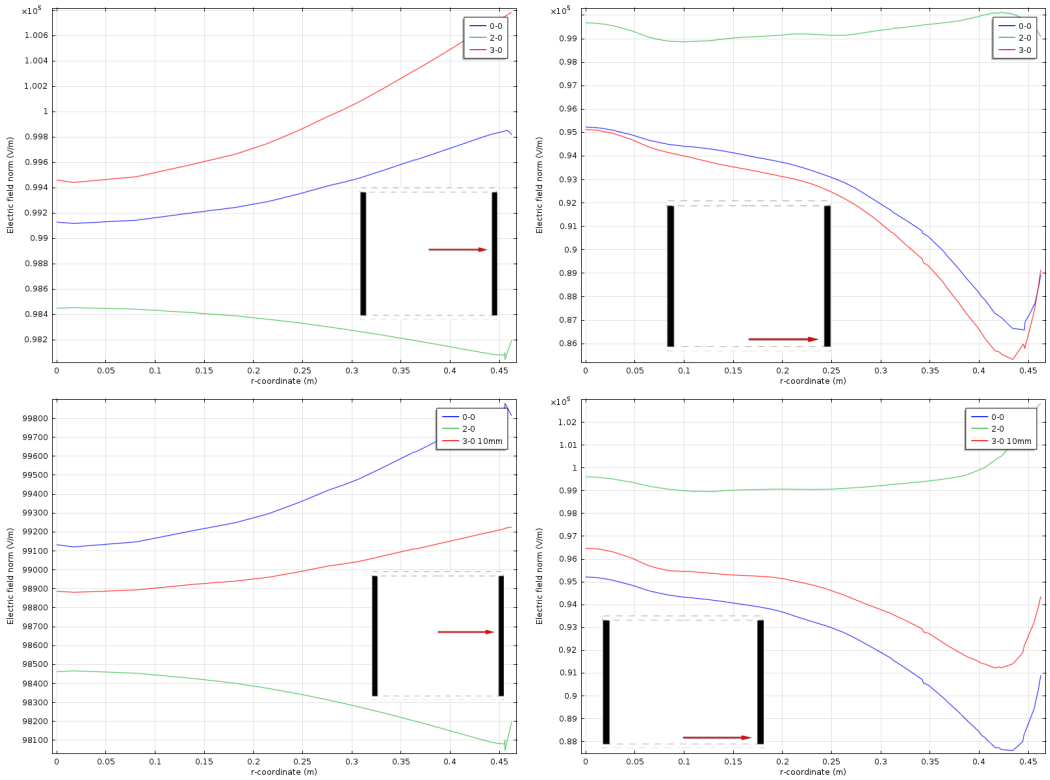

Figure 3.13: Normal electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right).

Configuration 2-0 clearly demonstrates an improvement over alternative geometries. The plots shown above, however, do not demonstrate the electric field "leak through" caused by the greater physical space between shaping rings. Figure 3.15 below displays the electric field 10mm from the PTFE wall along the vertical axis of the TPC, where the large variations in electric field can be observed as a cause of this greater spacing in configuration 2-0.

The following is an investigation into the effects of creating a greater potential gap between the cathode and first shaping ring. The geometry chosen for this investigation consists of 10 mm shaping rings spaced evenly by 10 mm between the cathode and gate mesh. The difference in potential between the cathode and first shaping ring for configuration 3-1 is 1500 V greater than that of 3-0. The average electric field within the TPC for simulation 3-1 10 mm has been calculated at 99762 V/m.

Figure 3.16 displays the electric field within the TPC, comparing the results of the smulation of version 3-0 10 mm and 3-1 10 mm, with a focus on the regions of high electric field around the top and bottom corners of the TPC where the side walls meet the gate and

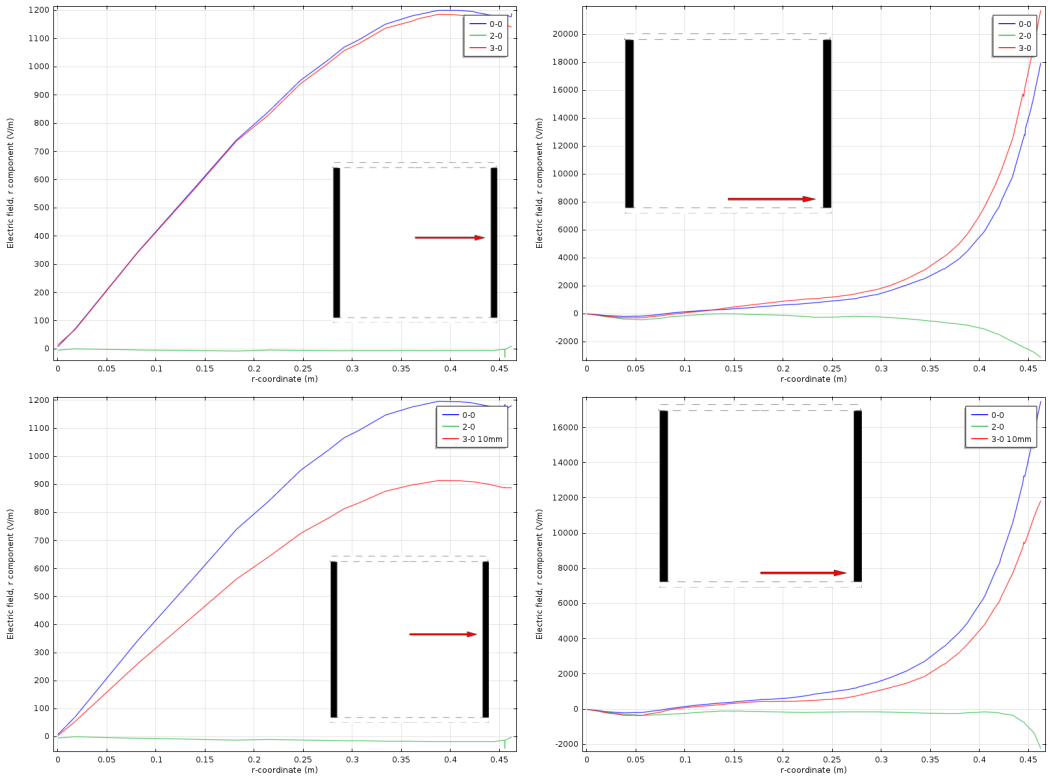

Figure 3.14: Radial component of the electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right).

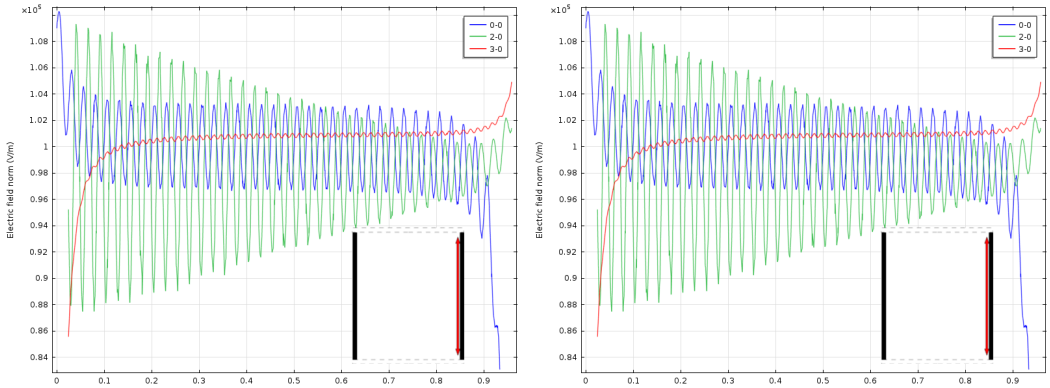

Figure 3.15: Normal electric field along the vertical, z-axis 10mm from the PTFE surface of the TPC chamber.

cathode. By eye, one can observe a drastic improvement in the uniformity of the electric field within the entirety of the TPC in this figure. By looking into the regions of high electric field in more detail, where the primary difference between the two models lie, the

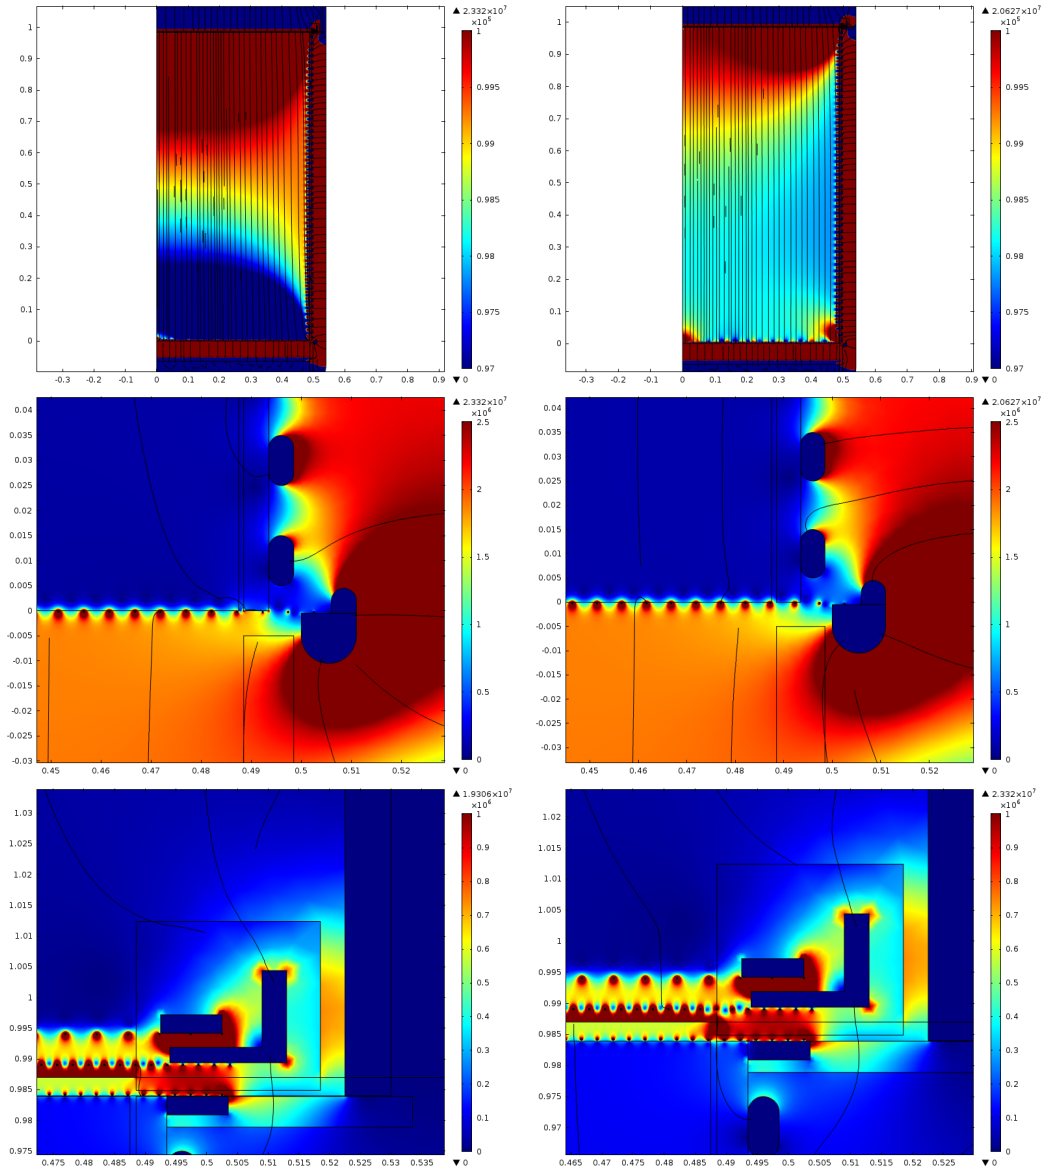

Figure 3.16: Cross section displaying the normal electric field within XENON1T for verision 3-0 10mm (left) and 3-1 10mm (right). The colour scale displayes the electric field in V/m. The simulated half cross section is shown (top) along with the cathode area (center) and anode (bottom).

specific differences in these regions can be observed. In order to quantify the improvement to the uniformity of the electric field with simulation 3-1 10mm, however, the standardised lines must be investigated and compared with simulation 3-0 10mm.

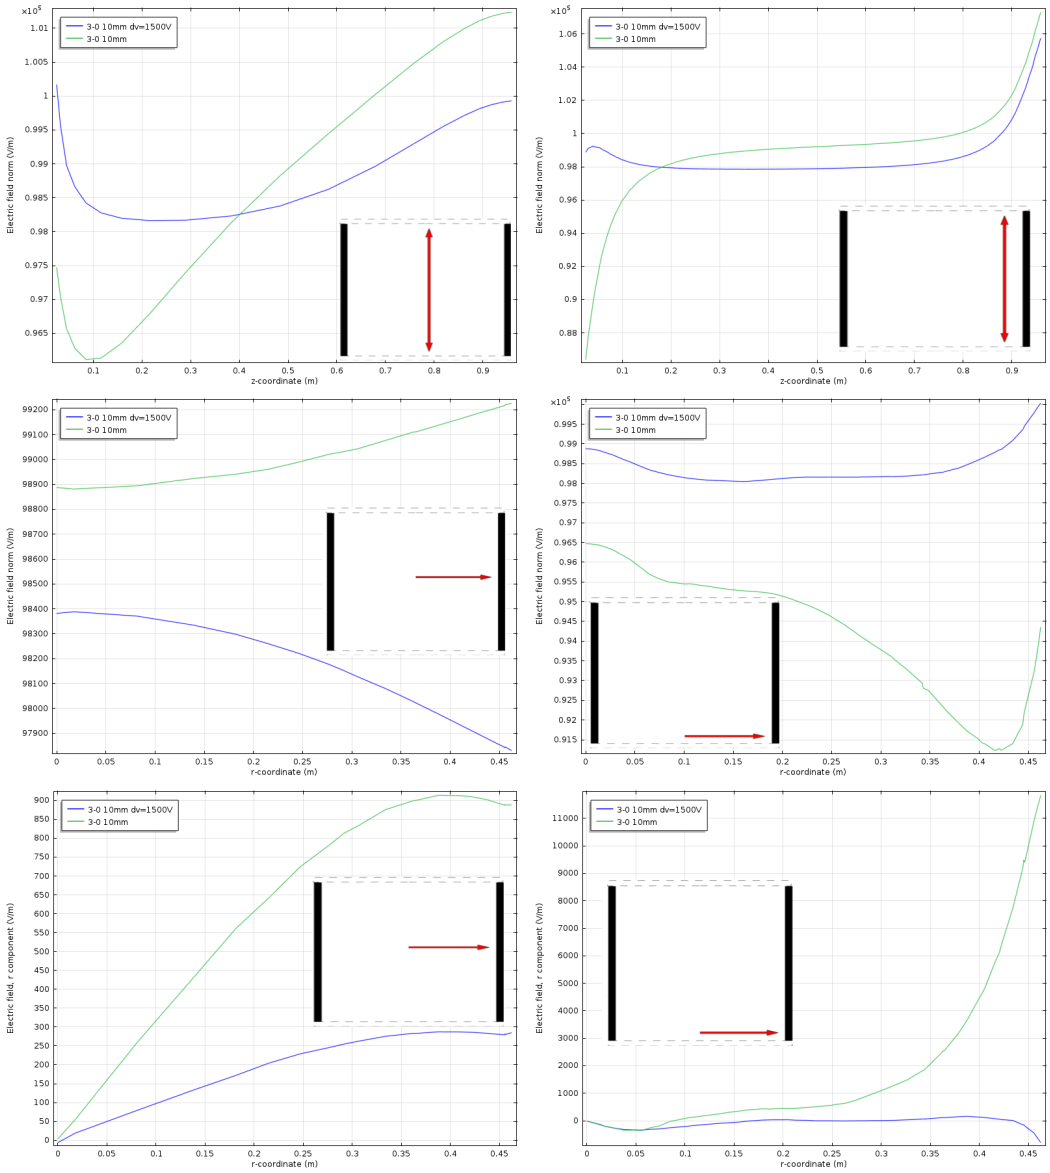

Figure 3.17: Top: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right). Middle: Normal electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right). Bottom: Radial component of the electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right)

Comparing the normal electric field for simulation 3-1 10mm and 3-0 10mm along the z-axis of the TPC as in Figure 3.17 results in a clear improvement with the implementation of electric field version 1. Although an improvement is observed at the edges of the TPC,

the improvement is most striking at the center, where the variation in the normal electric field has dropped from $\approx 4\%$ to ≈ 2 %.

Following the comparison of simulations 3-0 10 mm and 3-1 10 mm along the z-axis, the radial axis must also be investigated in order to determine if the same improvements observed in Figure 3.17 hold along both directions. In order to determine any potential drift of electrons along the radial component of the TPC, the electrical field along that direction is also simulated.

A re-design of the shaping rings is then considered, altering the shape to 8mm spaced evenly with a center to center distance of 10mm. Comparing this new design to the previous best simulation 3-1 10mm yields the following results shown in Figure 3.18.

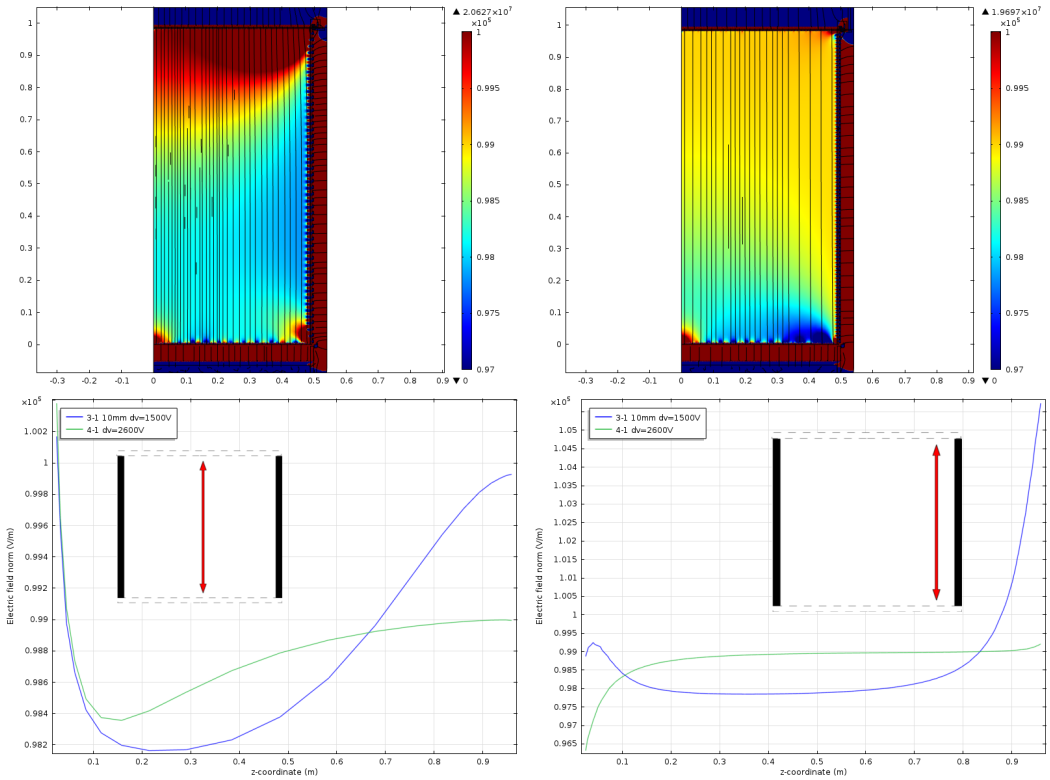

Figure 3.18: Top: Cross section displaying the normal electric field within XENON1T for verision 3-0 10mm (left) and 3-1 10mm (right). The colour scale displayes the electric field in V/m. The simulated half cross section is shown (top) along with the cathode area (center) and anode (bottom). Bottom: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right).

From Figure 3.18, a factor 2 improvement in the uniformity of the normal electric field for simulation 4-1 may be observed along z-axis in the center of the TPC. To judge the

radial distribution of this improvement, the simulations were performed across the remaining standard analytical lines, the results of which are shown in Figure 3.19. These additional comparisons between 3-1 and 401 comprehensively demonstrate that simulation 4-1 improves the uniformity along almost all directions by a factor of 2. Furthermore, when observing the variations of the electric field nearby the PTFE panels, 4-1 is able to completely eliminate the non-uniformities previously observed in simulation 3-1.

Concluding the simulations provided above. The geometrical and electric field configurations that maximising the uniformity the greatest along the majority of the TPC is simulation 4-1. Average normal electric field variation for this simulation is maintained at 1-2% across the entire TPC.

The following compares the variations in the electric field between shaping rings spaced 2mm apart with respect to 3mm. A stronger potential difference has been placed between the cathode and first shaping ring of 2600V for the 2mm and 2800 for the 3mm. These values have been chosen to optimise the uniformity under each case to provide a more direct comparison.

Using the optimised version of the TPC design, one can now begin to study the dependence of field uniformity on the radial positioning of the shaping rings. Any correlation between these two parameters is of great importance due to the structural variations that may occur during the constructions of the shaping rings. As such, one must study the change in uniformity of the electric field in the case in which the shaping rings are moved of order 1 mm, the typical production error stated for the shaping rings.

Prior to performing any simulations, the scenarios to be compared must first be determined. The first case is that in which all the field shaping rings are appropriately constructed, however the cathode is 1 mm too large in radius. The second is the inverse of this case, in that all the shaping rings are systematically positioned 1 mm greater in radius, and the cathode is correctly positioned. The third scenario is one in which the cathode and shaping rings are randomly determined to either be correctly positioned or 1 mm larger in radius. This last study is a worst-case scenario in the event in which the shaping rings have a 50% probability to be incorrectly constructed with a 1 mm larger radius.

The first simulations were preformed looking at the uniformity of the field along the z-direction of the TPC in two positions, both in the center, and half way towards the edge of the detector at r=250 mm. The results of these simulations are shown in Figure 3.22, and show no significant deviations of the field between each other.

The normal electric field must then also be compared along the radial component of the TPC. To compare the deviation of each scenario along the entire length of the TPC,

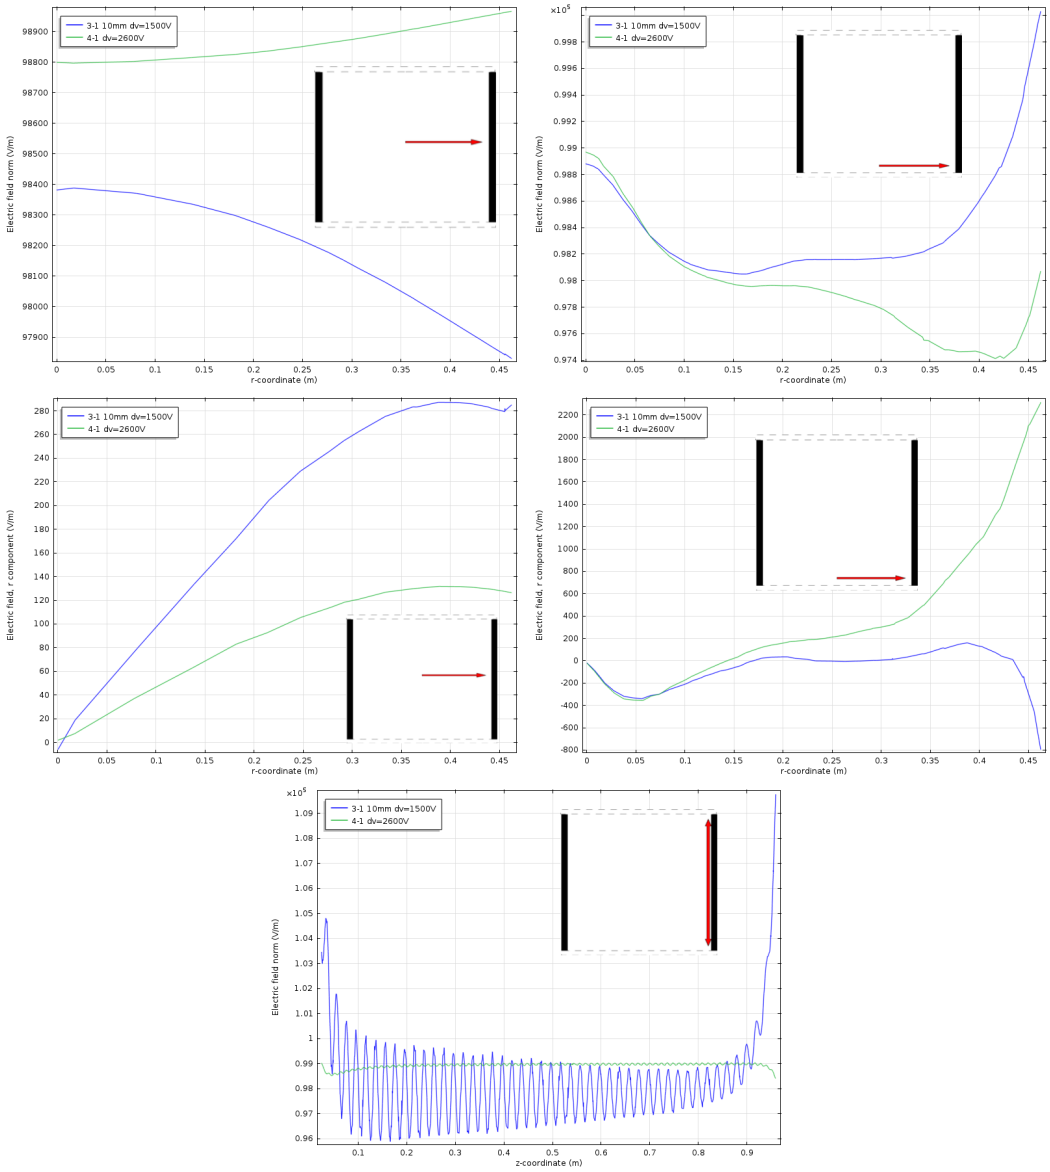

Figure 3.19: Top: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right). Middle: Radial component of the electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right). Bottom: Normal electric field along the vertical, z-axis 10mm from the PTFE surface of the TPC chamber.

four Z-coordinates have been selected, Z =100 mm, 200 mm, 500 mm, and 600 mm. The results can be shown in Figure 3.23. From this figure one can observe a large deviation from uniformity in the case in which the shaping rings are expanded radially at random

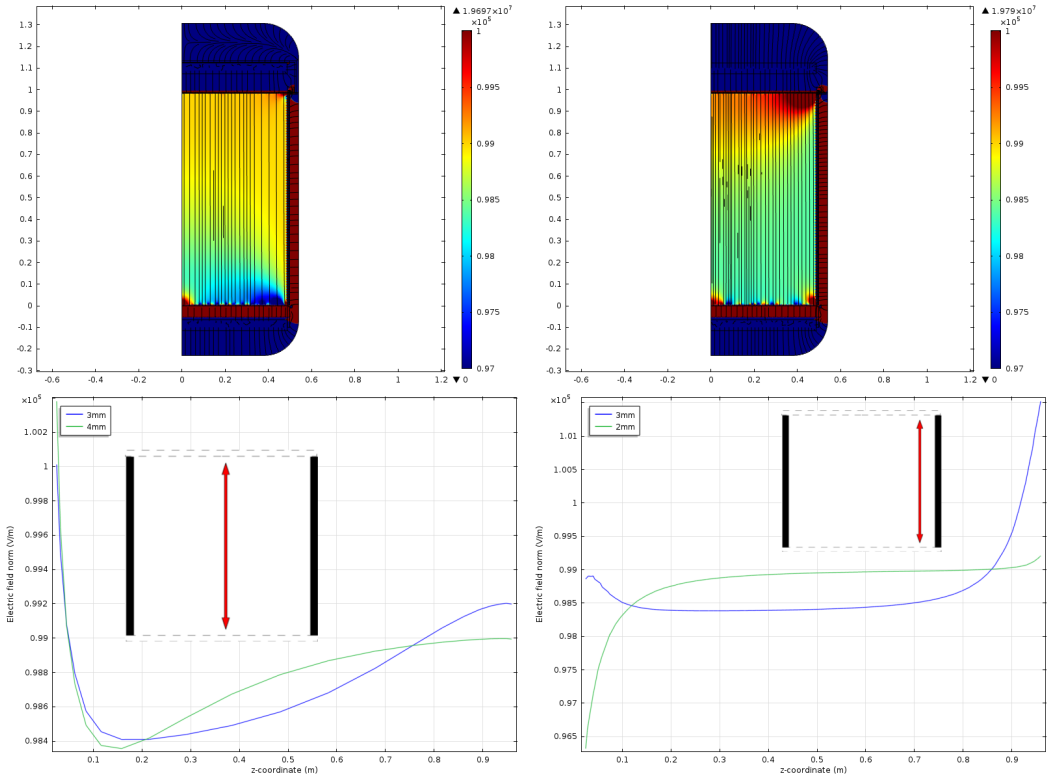

Figure 3.20: Top: Cross section displaying the normal electric field within XENON1T for verision 3-0 10mm (left) and 3-1 10mm (right). Bottom: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right)

in the region between r=450-480 mm, at the edges of the TPC. By observing how this deviation evolves as a function of Z within these plots, it can be see that this is a result of the positioning of the analysis line with respect to the Z-position of the shaping rings. As these values of radius are not to be within the fiducial volume of the detector, variations in this region are not vital to the accurate operation of the TPC.

From Figures 3.22 and 3.23, one can conclude that uniformity is held along all simulations along the entirety of the TPC. The largest deviation is shown in the range between r=450-480 mm in the case in which the shaping rings are distributed randomly. To determine the average deviation along the length of the TPC, the normal electric field at r=460 mm along the z-direction of the detector, along with the 2-D colour map of the random shaping ring distribution is simulated and shown in Figure 3.24. From this figure, the average deviation is shown to be <0.2%.

The results shown above demonstrate the a deviation of 1mm on shaping rings near the cathode, have effects <0.2% on field uniformity, with the greatest deviations occurring

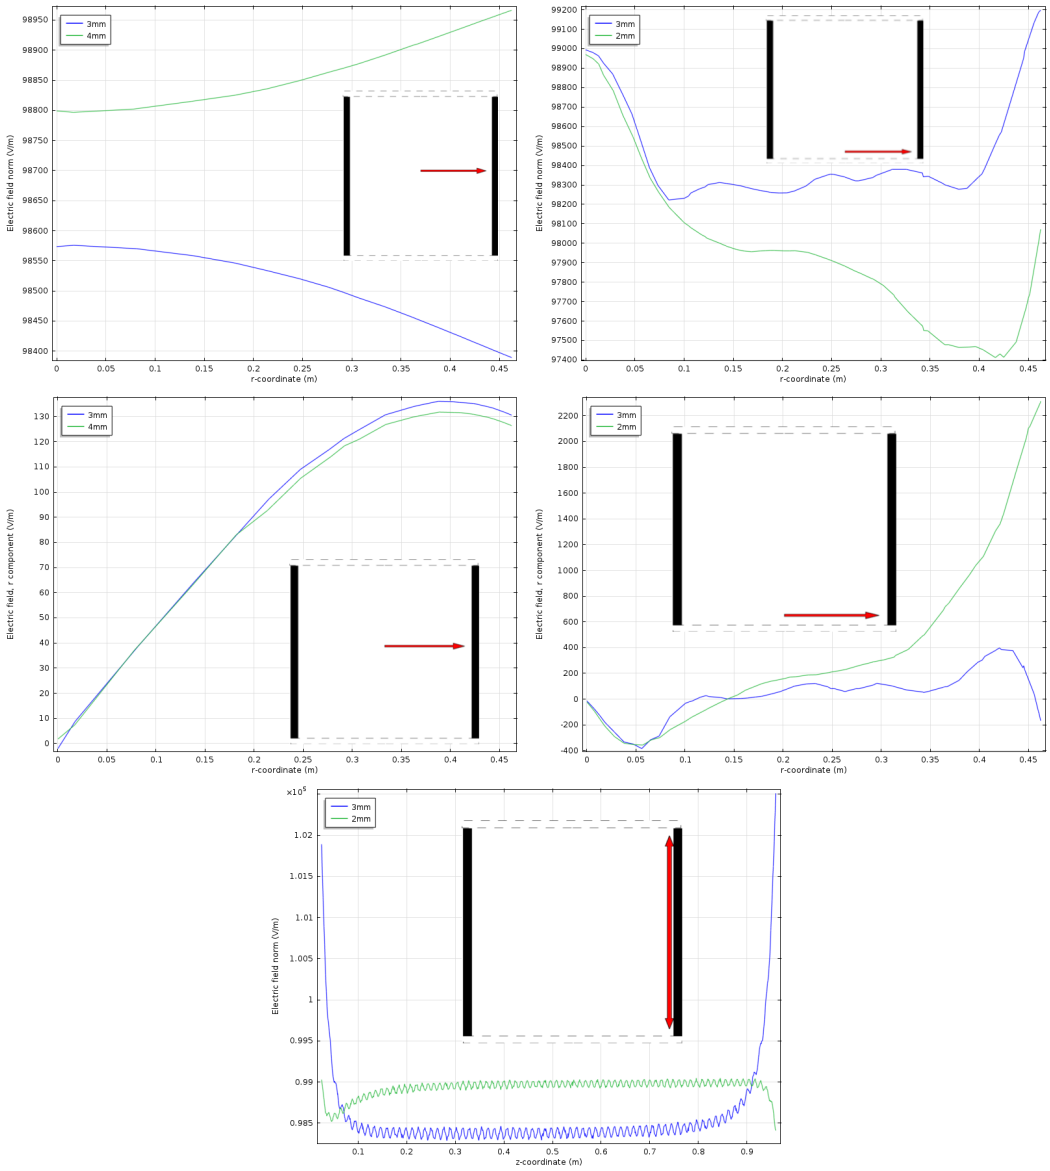

Figure 3.21: Top: Normal electric field along the z axis at radial positions r=0 (left) and r=440mm, 50mm from the edge of the central chamber (right). Middle: Radial component of the electric field along the radial axis from the center of the TPC chamber (left) and z=50mm (right). Bottom: Normal electric field along the vertical, z-axis 10mm from the PTFE surface of the TPC chamber.

within 30mm of the TPC wall. This effect is small enough to be managed easily, and it is possible that the volume will be cut out of the fiducial mass. It can therefore be concluded

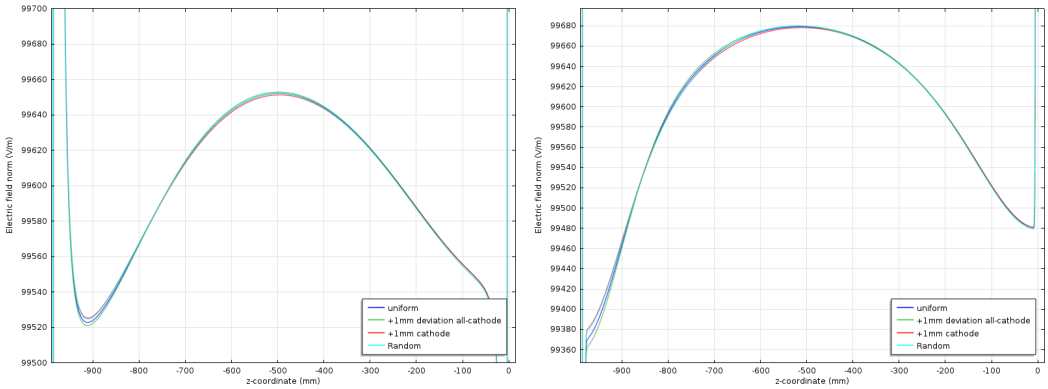

Figure 3.22: Electric field norm along Z direction of XENON1T at the centre of the TPC (left), and at r=250mm (right). The blue line indicates no change from the XENON1T model, the green line indicates and increase of 1mm in radius of all shaping rings but the cathode, and the red line is a result of only increasing the cathode radius by 1mm.

that a production error of 1mm on the radius of the shaping ring will not have any significant effect on the electric field uniformity.

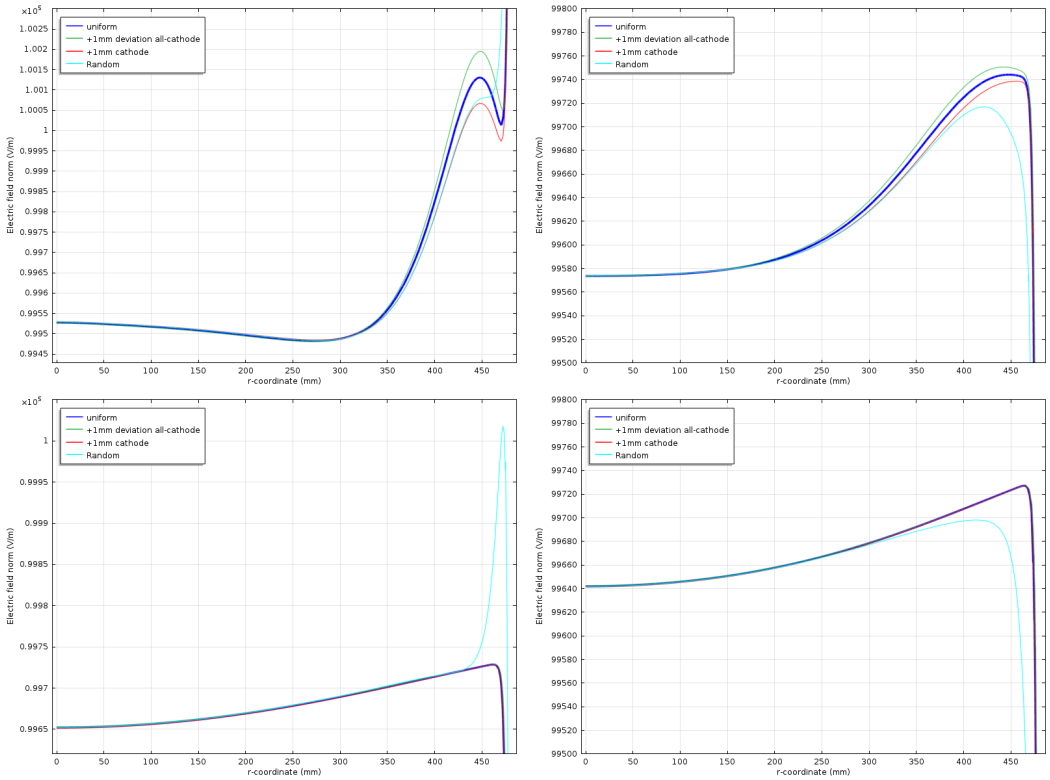

Figure 3.23: Electric field norm along radial direction of XENON1T at z=100 mm (top left), z=200 mm (top right), z=500 mm (left), and at z=100mm (right). The blue line indicates no change from the XENON1T model, the green line indicates and increase of 1mm in radius of all shaping rings but the cathode, and the red line is a result of only increasing the cathode radius by 1mm.

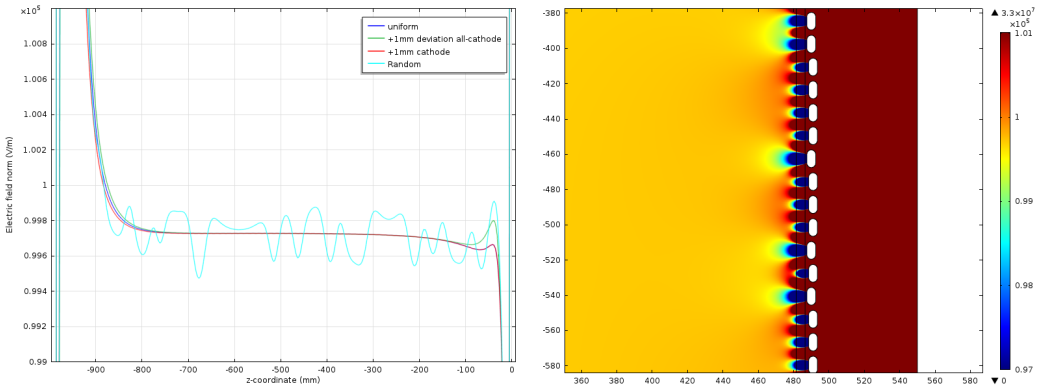

Figure 3.24: Electric field norm along Z direction of XENON1T at the centre of the r=460mm (left) along with the 2-D field results of the scenario in which shaping rings are placed 1 mm greater in radius at random. The blue line indicates no change from the XENON1T model, the green line indicates and increase of 1mm in radius of all shaping rings but the cathode, the red line is a result of only increasing the cathode radius by 1mm, and the cyan indicates the results for when shaping rings are placed 1 mm greater in radius at random.

3.2.8 Summary of Simulation Results.

As can be observed in the summary of the results in Table 3.3, after performing electric field simulations on numerous geometrical and electric field configuration. Currently simulation 4-1, utilising 81 10mm shaping rings spaced 2mm apart with rounded edges as opposed to semi-circular ones, provides the most uniform field with an average variation of approximately 1kV and an average field of 99755 V/m.

Simulation	**average variation from 100kV/m (%)**	**Leak-through (approx) (mm)**	**Reduced Fiducial Volume (approx, cm^3)**	**Additional Comments**
0-0	5-10	10	30	Default Geometry
0-1	2-5	10	30	Potential difference of first shaping ring varies on geometry
1-1	1-3	5	15	Shaping rings too close together for practical use
2-0	1-2	20	60	High Leak-through, low variance
3-0	5-10	10	30	Very small differences to 0-0
3-0 10mm	2-5	10	30	Compromise between 2-0 and 3-0
3-1 10mm	1-3	10	30	1500V dV between cathode and 1st shaping ring
4-1(2mm spacing)	1-2	5	15	2600V dV Most uniform field currently obtainable, if physically viable
41 (3mm spacing)	1-2	5	15	2800V dV, uniform field and structurally stable.

Table 3.3: Summary of results from each electric field simulation for the XENON1T TPC.

3.3 Cryogenic Testing of a TPC Segment

This field cage mockup is consists of 75 10mm shaping ring sections spaced 3mm apart from each other and placed within two teflon comb structures. The total lenth of the setup is 985mm.

Figure 3.25: Overall view of the experimental setup to test the field cage segment

Half of the segment was placed within a dewar filled with liquid nitrogen for approximately 5 minutes, until the mockup had reached thermal equilibrium, using the setup displayed in the images below.

After cooling, it was noted that the overall length of the segment had been reduced to 975mm. Deformations along the remaining axis were also observed, however the deviations of these deformations were determined to not be significant enough to be of concern . After the segment had warmed to room temperature, the overall length returned to 985mm, with not structural weaknesses observed. It can therefore be concluded that, during the extreme cooling scenarios implemented upon this TPC segment, due to the lack significant structural

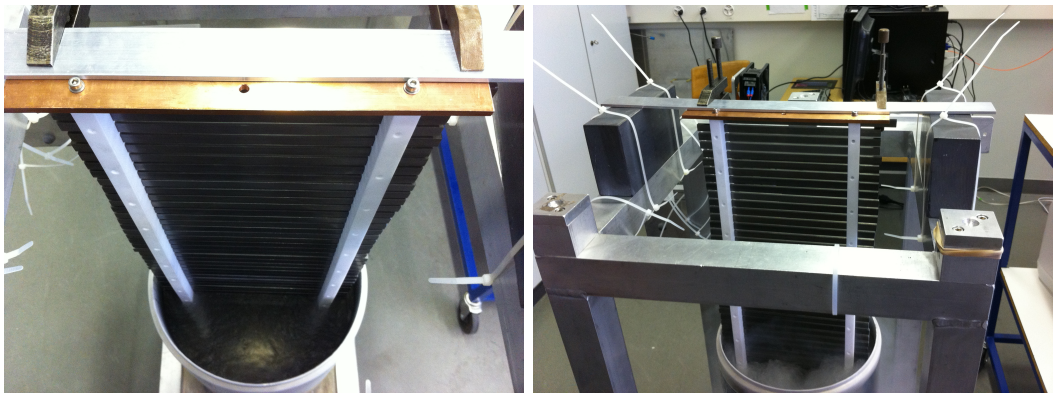

Figure 3.26: Image of the mounting structure for during the cryogenic tests of the TPC segment.

deformities observed, the current design of TPC will be able to withstand operational conditions within XENON1T.

Figure 3.27: Images displaying deformations along 2 axis after cooling

3.4 Electric Field simulations of Photo-Multiplier Tubes

3.4.1 Light Sensors for XENON1T

The ability for XENON1T to observe single photo-electrons (PE) is due to the precise photo-detectors used, known as Photo-Multiplier Tubes (PMTs). Optimised over several iterations for low-radioactivity and high quantum-efficiency, 248 tubes are currently operated in XENON1T. They consist of photo-cathode, 3" in diameter coated in a tri-alkali solution designed to be sensitive to xenon scintillation light. The electron released from the result of an interaction between a photon and the photo-cathode is accelerated via an electric field applied to the PMT due to the approximately 1.5kV provided to the voltage divider that applies the field across the photo-cathode and all subsequent dynodes. The accelerated electron then interacts with the first dynode, exciting further electrons. This process is repeated across all 12 dynodes until the resultant electrons are extracted from the anode as current. The desired gain, the ratio between initial photo-electrons and the final signal electrons, is approximately 1×10^6, allowing for the observation of signals as small as one electron.

3.4.2 Simulation Geometry

The entire structure of the PMT is not necessary in order to produce an accurate electric field between the Cathode and first dynode. Therefore the geometry used to produce the following simulation results includes the entire body of the PMT, the grid with the surrounding plate, and the first dynode. The dimensions of each of these structures have been carefully measured and implemented into the simulation. An Overview of the geometry can be seen in Figure 3.28 below, and will be used for all subsequent electric field, and particle drift simulations.

3.4.3 Electric Field Simulations

For the electric field simulation, the entire body of the PMT as well as the cathode placed at the upper limit of the body is placed at a voltage of -1600V. The grid, the plate surrounding the grid as well as the first dynode are placed at the same voltage of -1254V.

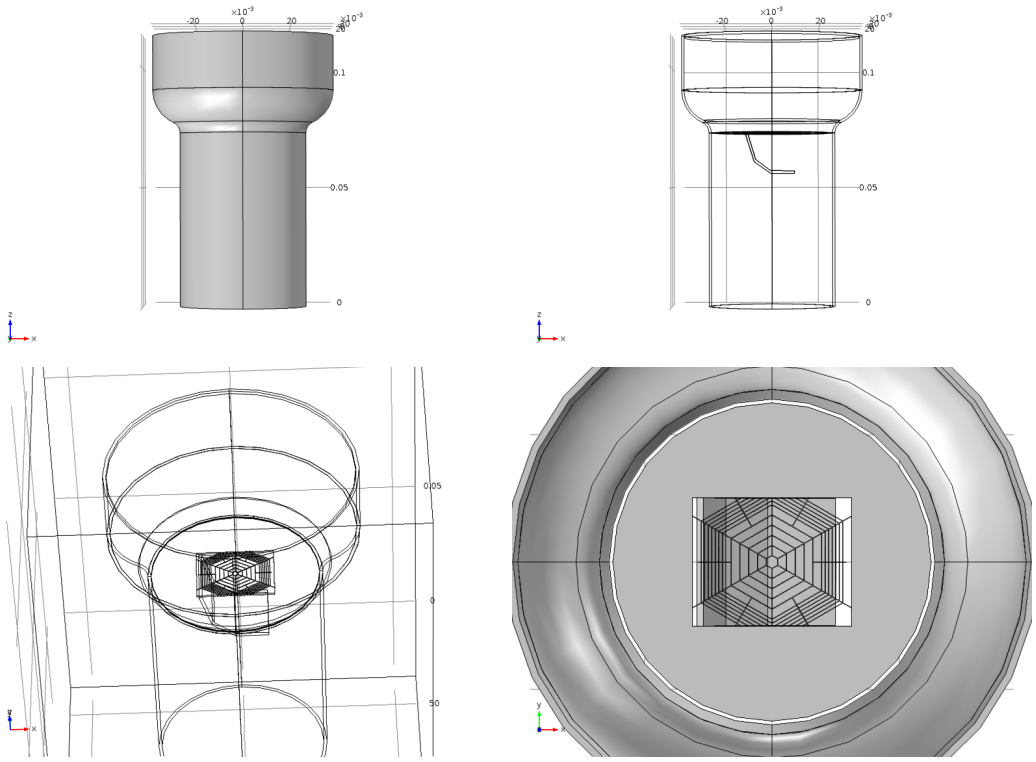

Figure 3.28: Geometric setup used in simulations of the XENON1T photomultiplier tubes.

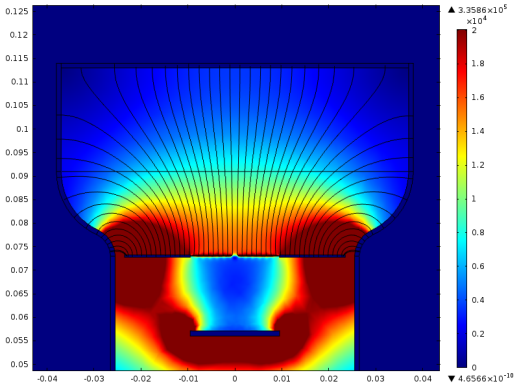

Figure 3.29: 2D cross section of the electric field map used in simulation of the XENON1T photomultiplier tubes.

3.4.4 Electron Timing

To test the validity of the simulation, electrons were placed randomly on the cathode and given a nominal velocity in the Z direction of -10m/s to negate edge effects of the simulation.

The resultant electron paths are shown in Figure 3.30 below. It is clear from the figure that electrons, regardless of position created on the cathode are focussed efficiently onto the first dynode.

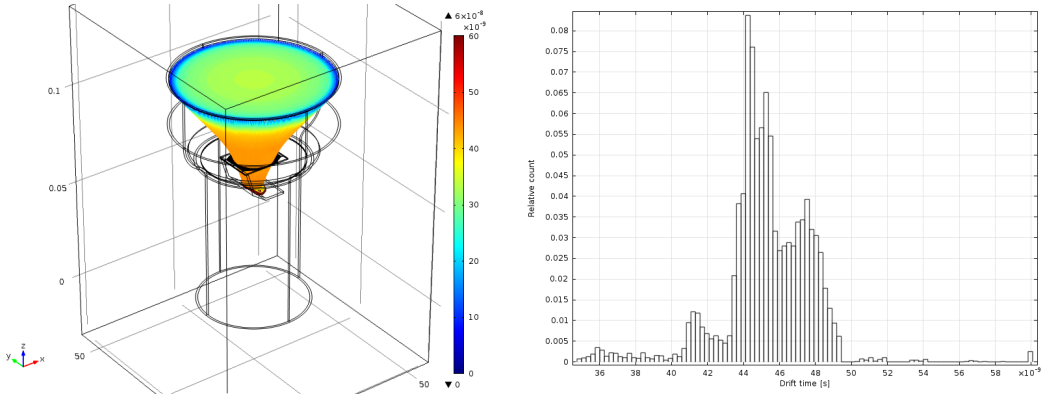

Figure 3.30: Electrons released at the cathode with an initial velocity directed downward at 10m/s. The 3D representation of each particle path is shown (left) along with the distribution of arrival times (right). The colour axis represents time in seconds.

From Hamamatsu, we know that electrons should arrive at around 46ns with a spread of 9ns. Figure 3 demonstrates that both of these are consistent in the simulation. 10,000 electrons are initially released from the cathode, of these 9500 arrive at the cathode, giving an efficiency of 95%. This is higher than the predicted (90%) from Hamamatsu, but may be due to a difference in wire thickness between simulation and reality.

3.4.5 After Pulse Simulations

For afterpulse simulations, a volume was defined within the upper cylindrical section of the PMT. Particles can then be created randomly within this volume and, starting from rest, be accelerated due to the field shown in figure 2. The results of this analysis are shown in Figure 3.31. The figure shows a consistent arrival time for protons colliding with the cathode at around 0.2-0.3 μs. This is consistent with the expected value of 0.27 μs.

Using the same method as described previously for Figure 3.31, after pulses due to xenon 130 were analysed. A single ionised atom of xenon using the same initial starting positions as shown in figure 4 are placed at rest within the electric field. Results of the analysis are shown in Figure 3.32. The expected drift time for xenon according to the potential after pulse timing peak at around 2.8 us. The drift times shown in Figure 3.32 show a much broader distribution than expected with most take between 2-3 us.

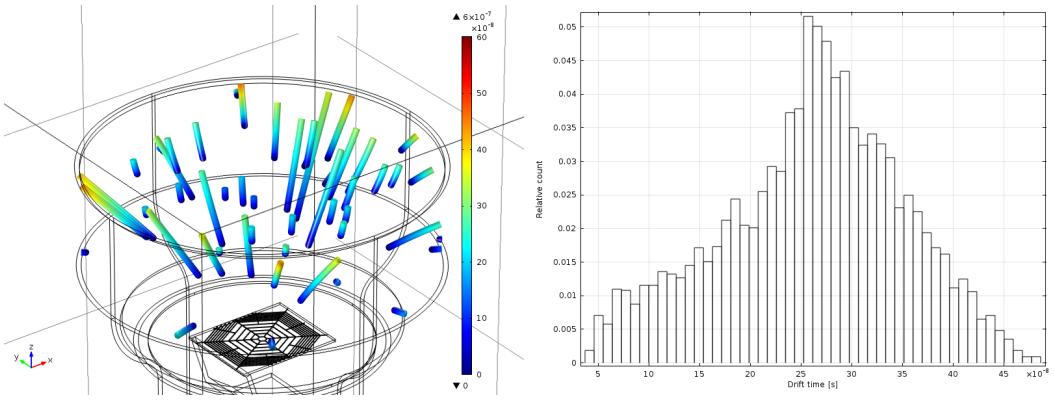

Figure 3.31: Protons released at rest randomly within the volume of the PMT above the mesh. The 3D representation of each particle path is shown (left) along with the distribution of arrival times (right). The colour axis represents time in seconds.

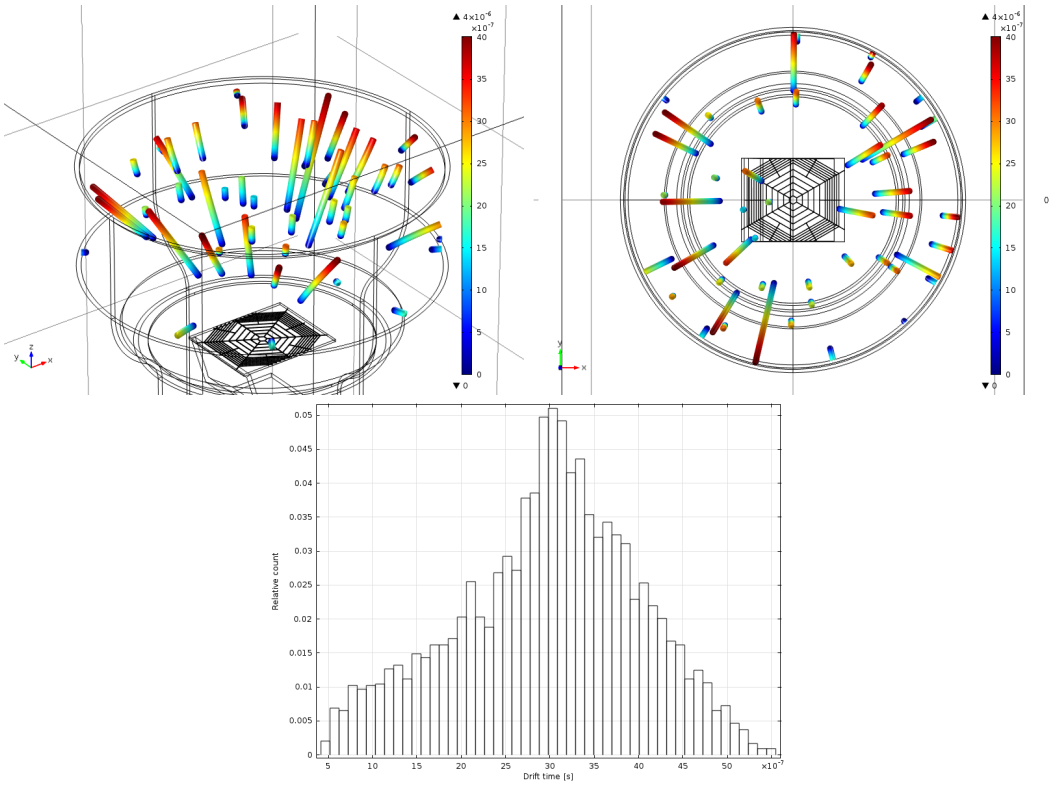

Figure 3.32: Xe130 released at rest randomly within the volume of the PMT above the mesh.. The 3D representation of each particle path is shown (left) along with the distribution of arrival times (right). The colour axis represents time in seconds.

Multiple other ions were also tested and the drift time distributions are shown in the Figure 3.34 below.

Ion	Simulated Drift Time [μs]	Experimental Drift Time [μs]	Error with respect to Experiment [%]
He	0.53	0.52	1.9
CH4	1.05	1.00	5.0
CO	1.40	1.32	6.1
Ar/CO2	1.65	1.58	4.4

Table 3.4: Summary of results from each after pulse simulation, compared with the experimentally observed value.

The drift times of ions simulated in Figure 3.34 and compared in the table above all correspond to the peaks given by experiment within 0.1us. This corresponds to a fractional error with respect to experimental results ranging up to 6.1%.

Figure **??** below demonstrates a comparison between all simulated and expected experimental peaks. A fit is also displayed on this plot, the gradient of which has a value of 1.04.

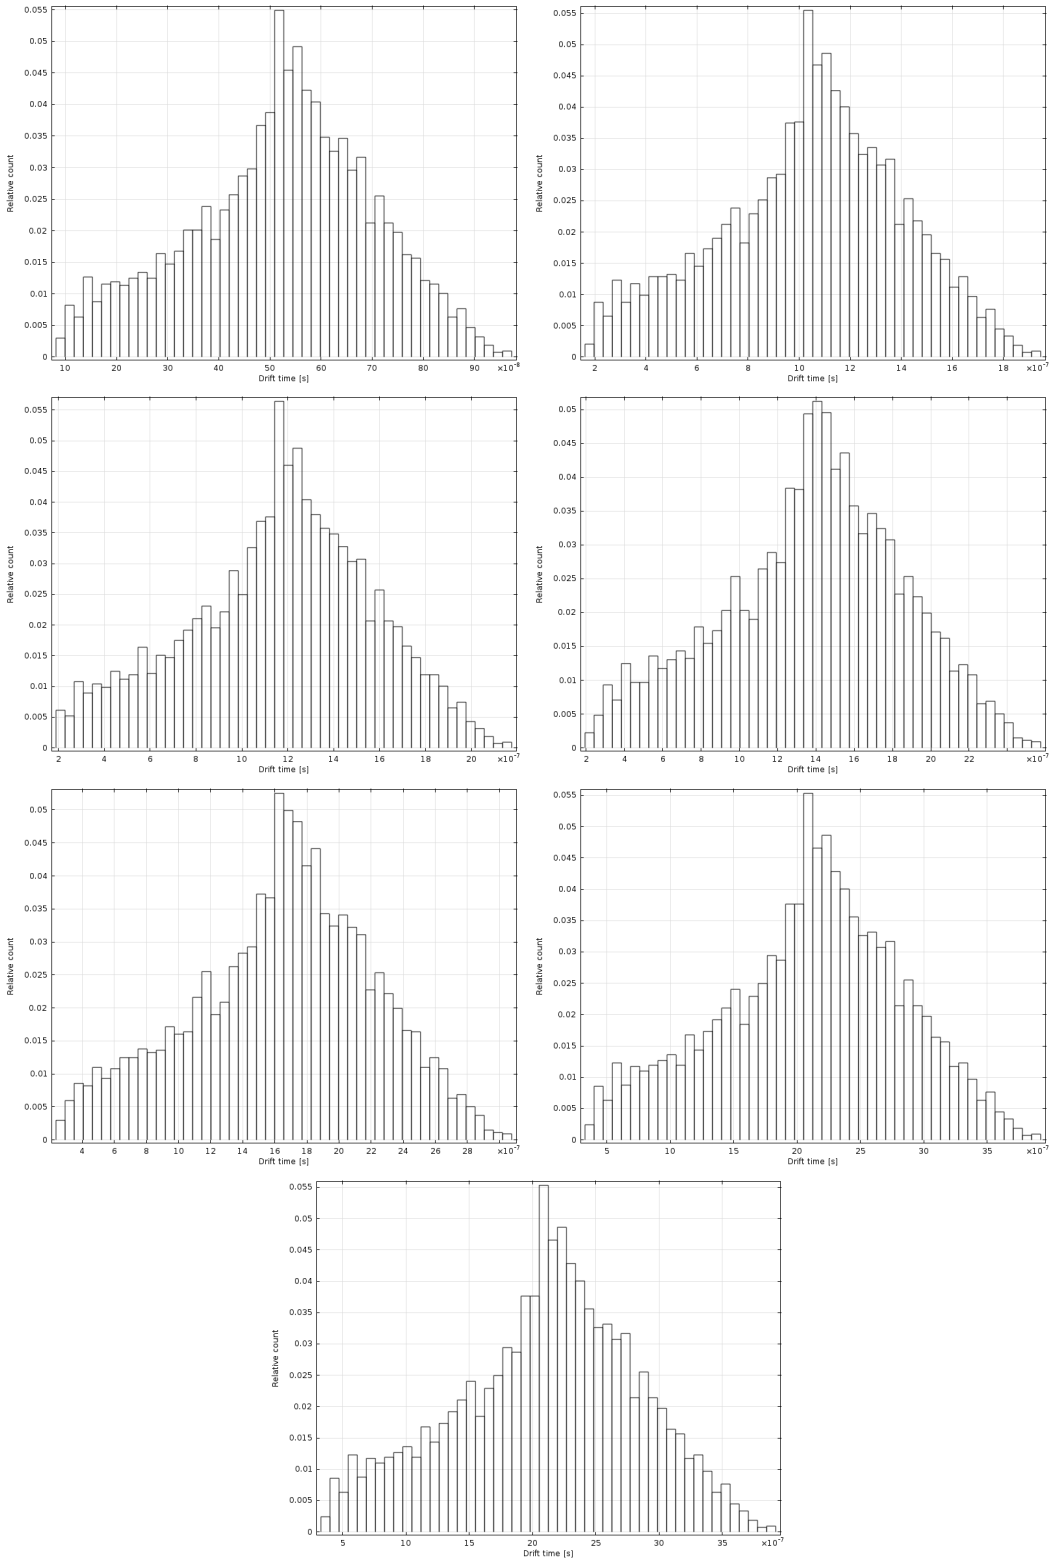

Figure 3.33: Time from release to cathode for various ion released at rest randomly within the volume of the PMT above the mesh. Reading left to right, He, CH4, Ne, CO, Ar, Xe130++, Xe130

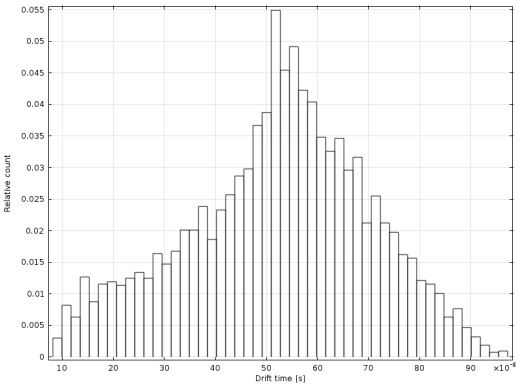

Figure 3.34: Graphical comparison between experimental and simulated afterpulse drift times for the following ions listed from shortest to longest drift times: He, CH4, Ne, CO, Ar, Xe130++, Xe130

If we assume an electron from rest at the cathode, it accelerates though 350V before it arrives at the first dynode. Using this, if an initially standing electron were accelerated to the first dynode at deposited all of its energy to a CH4 mass of ion, it would have an upper limit of around 65,000 m/s. This is, of course, not a realistic case, and therefore a velocity of the order 10,000m/s (approximately 1/6 of maximum, around 60eV) seemed more appropriate. This proportional alteration in vertical velocity has been applied to all ions based on their mass. The direction of the initial impulse to the 10,000 ions was given directly towards the grid to ensure highest probability of transmission to the cathode. The time distribution of ions reaching the cathode is shown in the figure below. Ions have been released from the two surfaces on the first dynode that electrons were shown to hit in Figure 3.30. Figure 3.35 provides a visual insight into the paths of ions (Ne in this case) released from the first dynode.

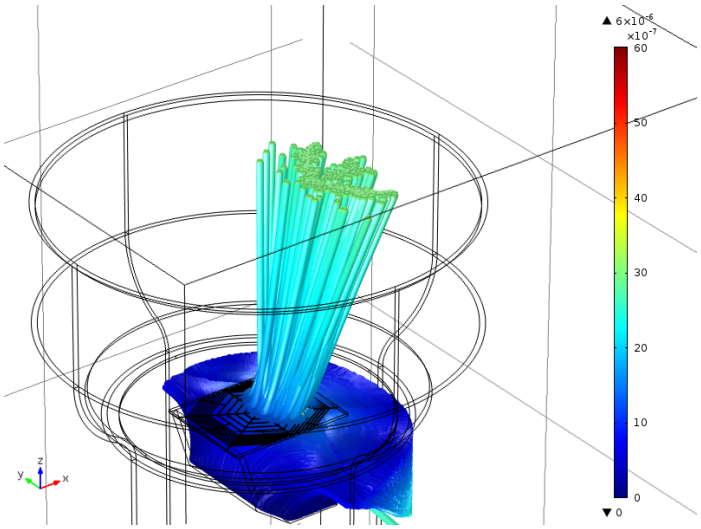

Figure 3.35: Particle paths of Ne, starting from the first dynode with an initial velocity of 9,600m/s towards the grid.

Of the 10,000 initial ions, it is observed that only approximately 5% made to the cathode in this scenario. The results showed that the middle surface of the three on the first dynode was providing all of the counts in Figure 3.36. All particles released from the horizontally orientated section were accelerated towards the PMT wall. None of the peaks present in Figure 3.36 correlate to their respective expected after pulse arrival times for this release velocity. Please note the small time window present in the figure. It is possible to alter this release velocity to purposefully produce an expected arrival time. Figure 3.37 shows

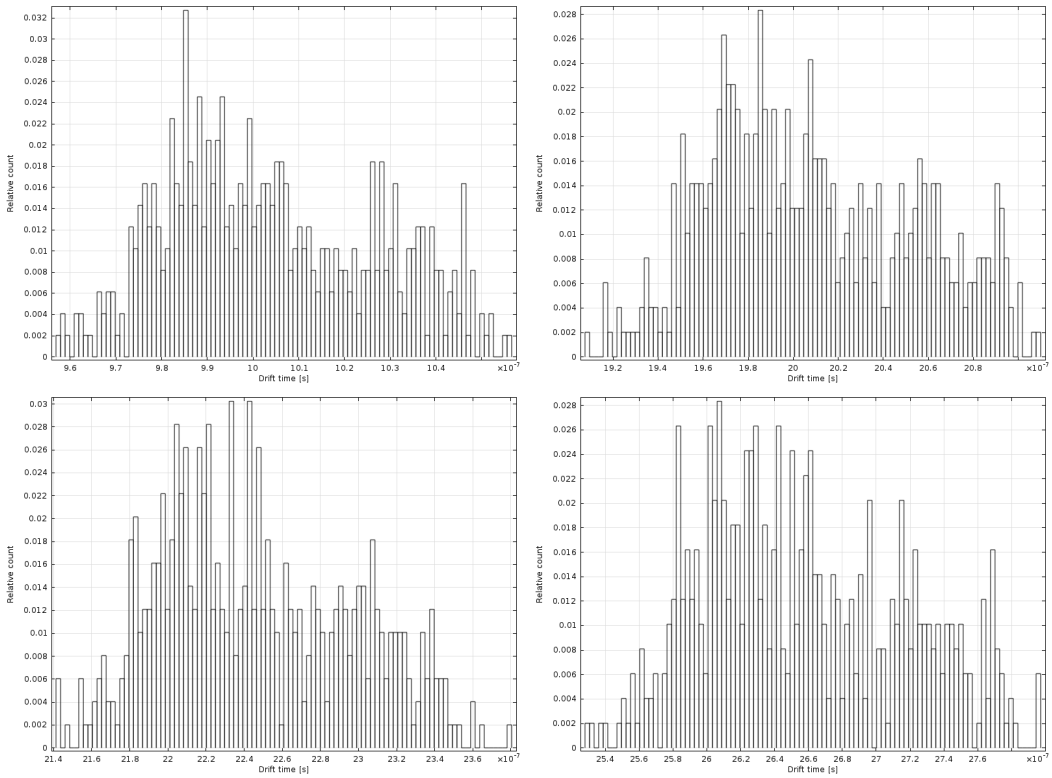

Figure 3.36: Drift time distribution of ions starting from the first dynode with an initial energy of 60eV towards the grid. Ions reading left to right: He, CH4, Ne, CO.

He released at around 60% of the total energy it could received from the electron. This fraction holds true for all other ions.

In this scenario, for He, a 3.5x increase in release velocity halved the arrival time. Although the peak is highly focussed and, by definition, the expected drift time value. The specific requirements for this scenario to occur are unlikely and almost unphysical.

3.5 Control of the Liquid Xenon level

XENON100 uses a bell system to control and stabilise the liquid level inside the Time-Projection Chamber (TPC).[**?**]. This system requires a constant flow of gas being directed into a diving bell to maintain the required overpressure for the desired level of liquid xenon (LXe). Due to the success of the bell in XENON100, XENON1T also uses the diving bell system to control the liquid level. As such, it becomes imperative to consider the implications of increasing the diameter of the bell from the 376.6 mm in XENON100, to the 1056.0 mm in XENON1T. Effects such as liquefaction from the interior surface of the

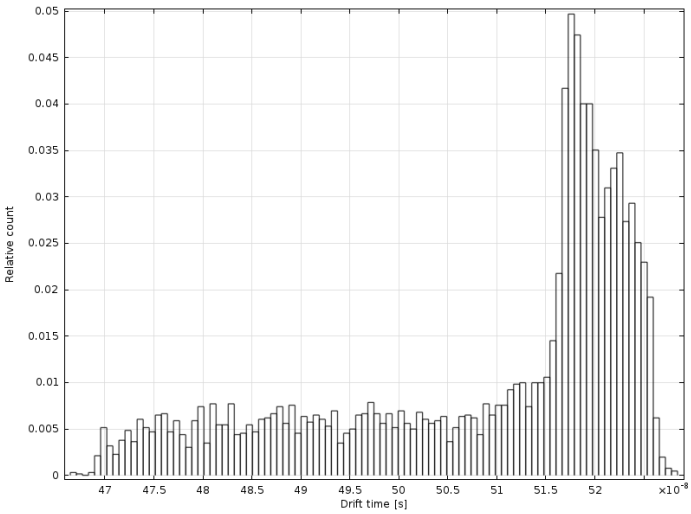

Figure 3.37: Drift time distribution of He, given 207eV of energy to reproduce the expected arrival time

bell could have ramifications for the heating power required. On the contrary, heat input goes into the bell could require larger cooling power for the experiment. We designed and operated a precise 204.0 *mm* bell to test such parameters in order to extrapolate any observed effects to XENON1T. In addition, in order to accurately determine the scaling factor, a smaller 161.0 *mm* diameter bell was subjected to identical tests.

3.5.1 Importance on Liquid Level for Liquid Xenon Experiments

3.5.2 Methods of Control

3.5.3 The Bell Experimental Setup

Figure 3.38 shows the general setup of the experiment. With the use of a recirculation pump, xenon gas is extracted from the main chamber, and pumped into the bell, with the option of opening a bypass back to the main chamber in the case in which the flow must be controlled via the flow controller. The flow meter is placed such that the total flow of gas in the system can be measured, and the flow controller such that only the flow into the bell is recorded.

The main chamber itself is a cylinder with 250.0 mm diameter and 295.0 mm height, surrounded by a vacuum insulation to keep the temperature inside it. Three screws extend downwards from the upper plate of the inner chamber that secure to top plate of the bell.

Three pipes extend from the chamber to the recirculation system. One to extract gas and is extended to the bottom of the inner chamber via a PTFE tube. There are then two inputs, one of which connects directly into the bell through a flexible metallic pipe to a VCR connection, and another that acts as a bypass for excess flow that outputs directly into the chamber.

In addition, a power supply is connected to the heater inside the bell which consists of a 50 Ω resistor placed within a PTFE shell to avoid disturbance of the LXe surface from bubbles. The heater is placed underneath the level of the side walls of the bell, to ensure that it is permanently submerged in LXe, even when the bell is emptied. This is to increase the efficiency with which it may control the liquid level.

Two sizes of bell were used for the purposes of comparison differing only in their diameter of 200.0 mm and 150. mm that shall be referred to as the large and small bell respectively. These bells are made from aluminium and have side walls of 65.0 mm from the outside as can also be seen in Figure 3.38.

Two level meters are implemented to observe the liquid level. They consist of two parallel steel plates, separated by 1.0 mm and maintained in place via a PEAK holder. At the top of each level meter, small wings are placed for the attachment of the signal cables. One of length 120.0 mm, and one of 40 mm to be placed outside and inside the bell respectively. Using information from both of these level meters, it is possible to determine the absolute height of the liquid level, as well as the height difference between the inside and outside of the bell. The expected change in capacitance for each level meter, using LXe is 2.38 pF for the short level meter, and 7.15 pF for the long level meter. For simplicity, the capacitance measured has no reference for which to determine the capacitance in pF, as such the output value is given in arbitrary units. The level meter can then be calibrated by observing the value at known point along the level meter.

3.5.4 Initial Testing and Results

In initial tests an issue arose with the level meters whereby the capacitance would suddenly and sporadically drop as shown in figure 3.39. The cause of this is unknown, but is suspected to be a bad connection between the capacitors. Despite this, the capacitance level can still be clearly read, and filling continued to ensure that the level meters were indeed working, and that the tests could continue as normal.

After one day, approximately 5.5 hours of filling, and 2.6KG of Xe added the results were as follows. None of the previous sporadic behaviour can be observed here. Level meter 4 shows a strong indication that it is registering a signal from the rising level of LXe. The

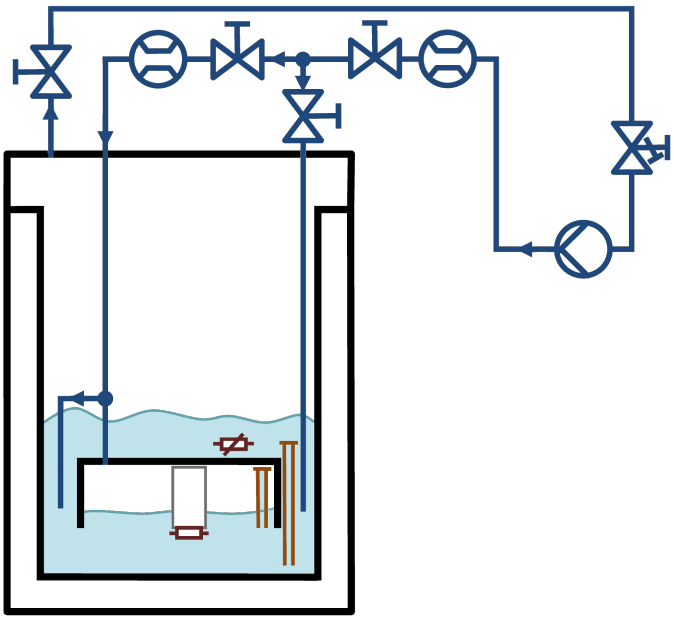

Figure 3.38: Schematic of experiment setup

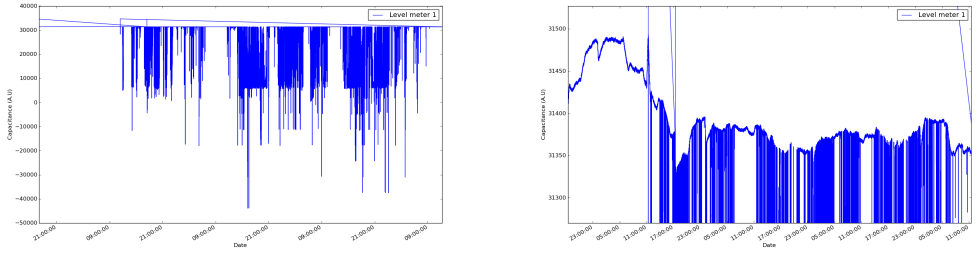

Figure 3.39: Plots of level meter 1 during cooldown. The left plot shows all data, while the right provides a zoom on the y-axis over all the data. All level meters showed exact same behaviour.

interior level meters do not show any such sign, but at the moment, it may be due to the level not being high enough. If these continue to not show anything, calibration can still be performed on level meter 4.

It is anticipated, based on technical drawings, that the interior level meters should see a signal after approximately 3-4kg of Xe have been added. Below are the plots of the second day of filling. An addition 2.75kg of Xe has been added, making a total of 5.35kg

LM4 started to show odd behaviour towards the end of the filling day, and continues to do so after filling has stopped.

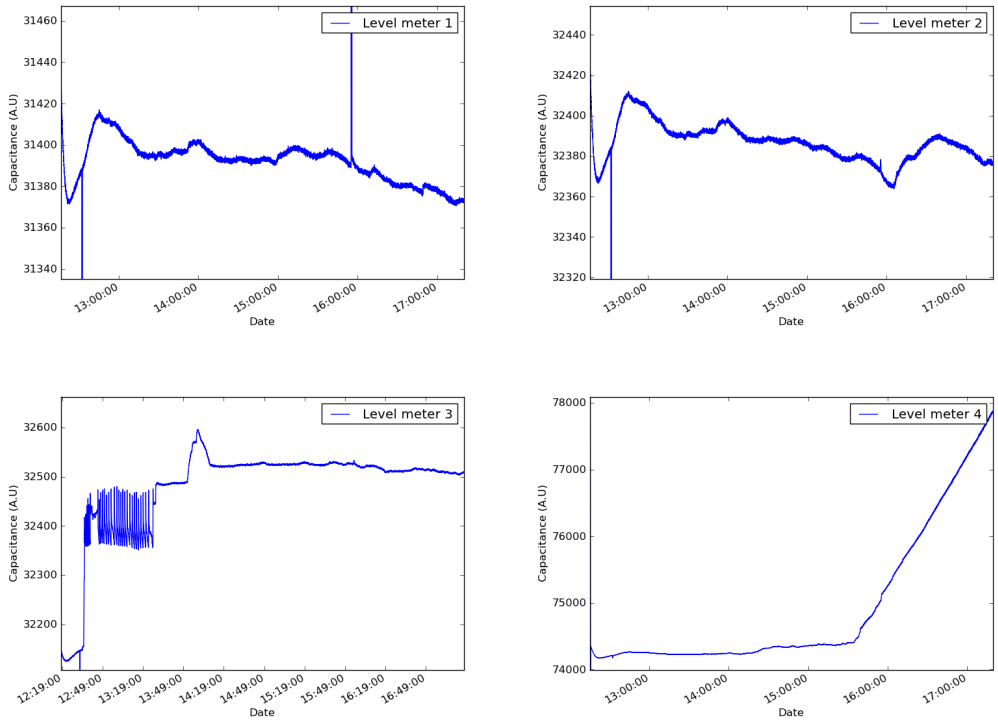

Figure 3.40: Filling data after 5.5 hours of adding Xe, or 2.6kg.

Below is a complete plot of all data from each level meter, since first starting to fill.

Using data from Figure 3.43, calibration can be performed on the interior LMs by analysing the filling up to the point to which they reach one of the fillers. The fillers are spaced one cm apart, starting 1cm above the copper ring at the bottom of the capacitors. Figure 9 shows a step increase in gradient up to its maximum that can be observed at approximately 15:00. This is taken to be the point at which the LXe level rises above this copper ring, and therefore the step up to the first spacer seen at around 15:30 will be the distance between the copper ring and the first spacer, 1cm.

A similar analysis can be done with LM4, the exterior capacitor. Due to the pausing in filling overnight, however, the change in capacitance must be taken into account over both days of filling.

The horizontal step is a little Teflon ring around the inner conductor which is 2,5mm high. So we were filling LXe at 64um/min. Then, using the second delitmited region, the UTI increased by 2060 arbitrary units in 58 minutes. At 64um/min, this gives about 1.8um/A.U.

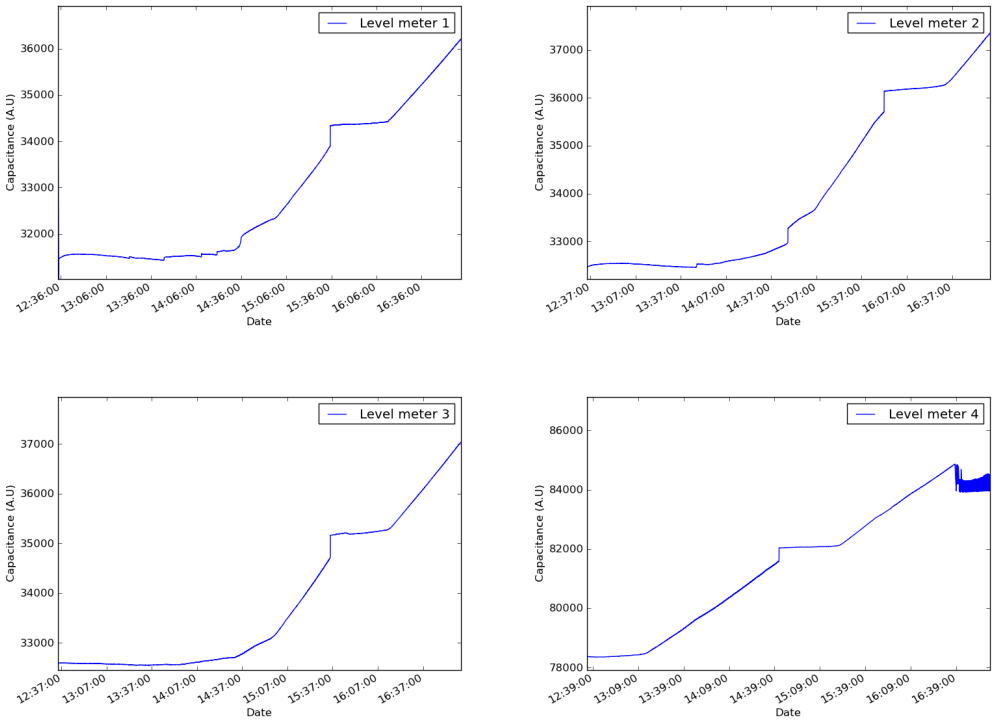

Figure 3.41: Filling data after 5.35kg of Xe

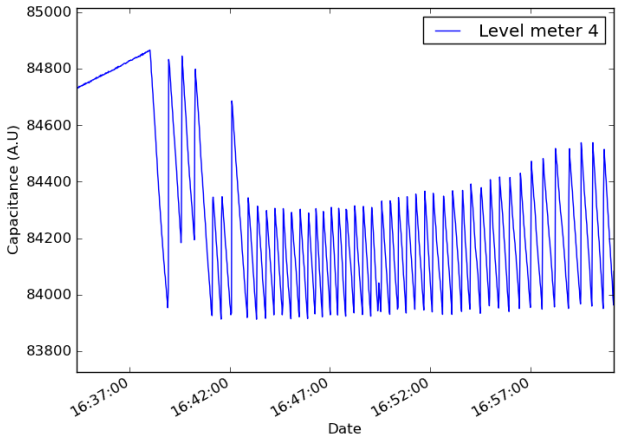

Figure 3.42: Pulsing behaviour of level meter 4

The fact that the flow meter is placed before the gas is split to the chamber and the bell, we cannot now the flow through one, or each, of the different pathways. This leaves the possibility of exploring stability.

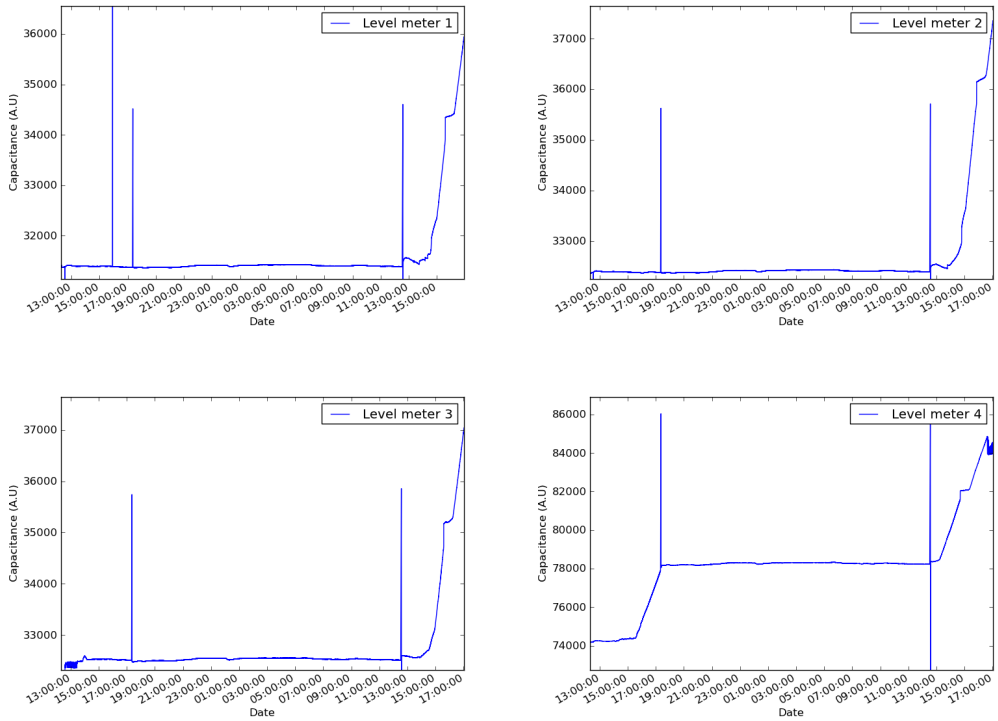

Figure 3.43: Full dataset of the filling procedure

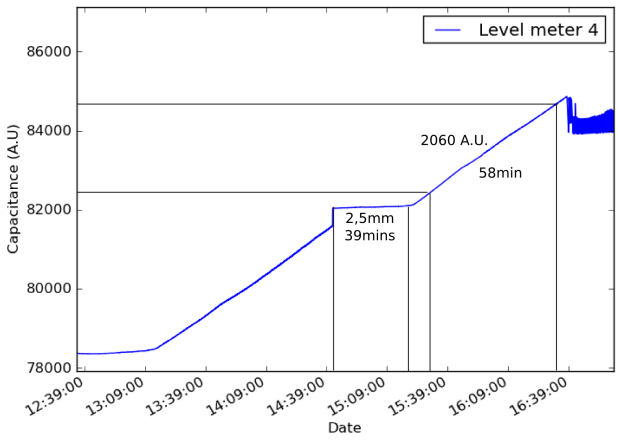

Figure 3.44: Pulsing behaviour of level meter 4

In the previous figure, the bigger spike at 13:28 would represent a variation of 2.12mm. But if we look at the more stable zone between 11:30 and 13:00, we can assume to control

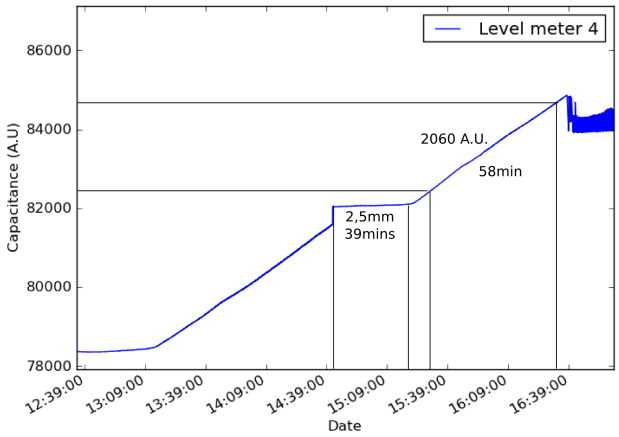

Figure 3.45: Example of calibration method for level meter 4

the level within 1.04mm. This value is the height difference of the larger peaks. Also keep in mind that this signal has not been denoised at all. The high frequency structure is surely not physical variations in the liquid level. Take these numbers as upper limits.

The following test is to observe if the bell can filled or emptied completely. The two limit points to this measurement are at the top: there is one dead centimeter because of the construction and the attachment of the levelmeters. At the bottom: the exhaust pipe. So we can never empty below the exhaust pipe and detect filling above the detection limit. This gives a maximum height differential of 45mm. So by creating a large differential pressure between the inside and the outside of the bell, we could produce this:

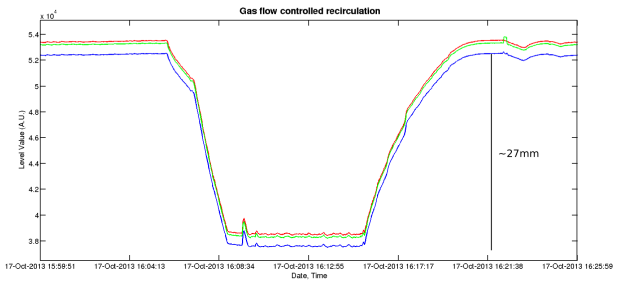

Figure 3.46: Emptying the bell with the use of recirculation flow.

We believe the slower oscillations at the bottom, when the bell is 'empty' to be bubbles escaping from the exhaust pipe. The difference of the highest to the lowest point would

give a total displacement of 27mm. The fact that we could not completely fill the entire bell is also apparent in the leveling off, and not a hard cutoff, after refilling.

So this gives us a definite zero-point calibration but does not confirm our height-arbitrary units calibration.

Operation of a resistive heater in the Liquid can be used to control the height of xenon. This operation produced much less actionable data. The operation of the heater interfered with the functioning of the level meters which frequently dropped out or maxed out. We were able to again empty the bell though.

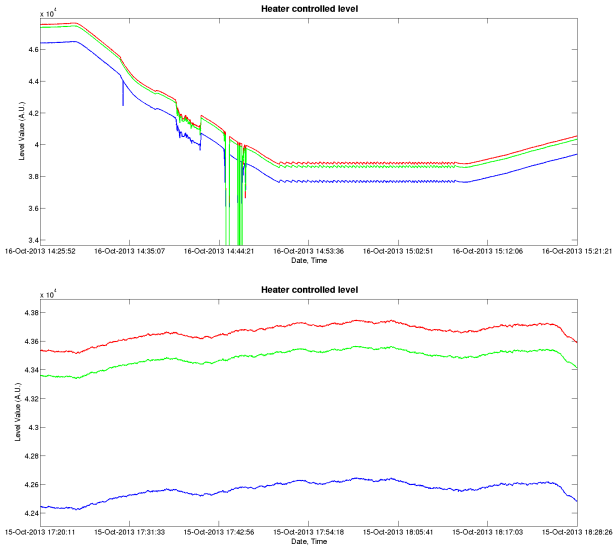

Figure 3.47: Controlling the level of liquid xenon with the use a of a heater in the form of a resistor placed within the bell itself. The upper plot displays the full time range of the tests, while the bottom figure shows a zoomed region where the level was controlled.

One can observe the droupouts in the level readout around 14:45. But this confirms our zero point calibration value from the gas controlled section. It also seems to confirm the bubbling from the exhaust pipe.

An improvement is seen over the gas circulation control. With 300mW of power through the resistor over a period of about 1 hour, the liquid level stays within 0.5mm.

In conclusion to these stability tests, stable operation was extremely delicate. This is specially true for the heater, just too much power and the bell empties, just not enough and it fills. This would necessitate a PID control system. With that in mind, the heater would be easier to have under PID control since only one value (level) controls one operable (heating). Controlling two valves simultaneously would be harder to implement. The combination of

level meter design and UTI readout proved to be very unreliable. A simpler plate capacitor readout using an LRC circuit analyzer could greatly enhance the robustness and readout ease.

Using the new parallel plate capacitors we filled 7.5kg LXe into MarmotXl before running into issues with the recirculation and were forced to recuperate. The new system for reading out the level meters included connecting each side of both capacitors to an arduino that sent a signal through one plate, and waited for a response from the other. The time taken for the signal to complete the loop is related to the capaciatance, and thus could be calculated.

Initially the level meters were tested by inserting paper approximately 2cm up between the plates. The ability to recognise the change in capacitance here would ensure the minimum required sensitivity for the bell tests.

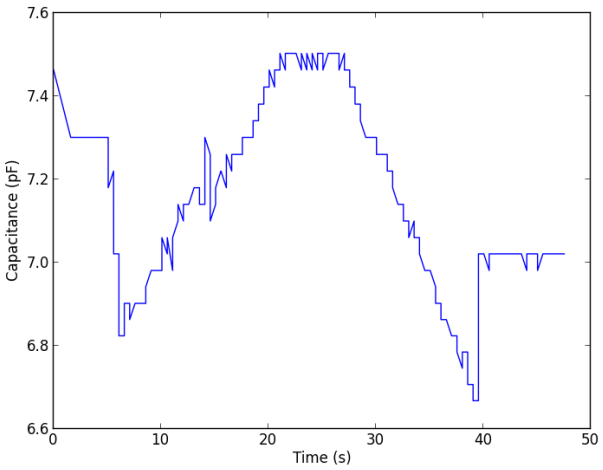

Figure 3.48: Results of inserting paper 2cm between the plates of the capacitor.

During cooling, data was taken to observe any differences in capacitance. None were found over this time, but the noise level was higher than previously seen.

Filling was then started. Data was taken only during the hours when which xenon was being addes to the chamber. No change in capacitance was observed with over 7.5Kg of LXe in the detector. These results concluded that potentially the arduino was unable to accurately read out through the feedthroughs of the detector.

The arduino assumes an internal capacitance within the system. In order to obtain an accurate measurement, this capacitance must be determined via calibration. By altering this internal capacitance value to only account for that within the arduino alone, the value of both level meters is reduced. The idea was to produce a lower noise level in the hope

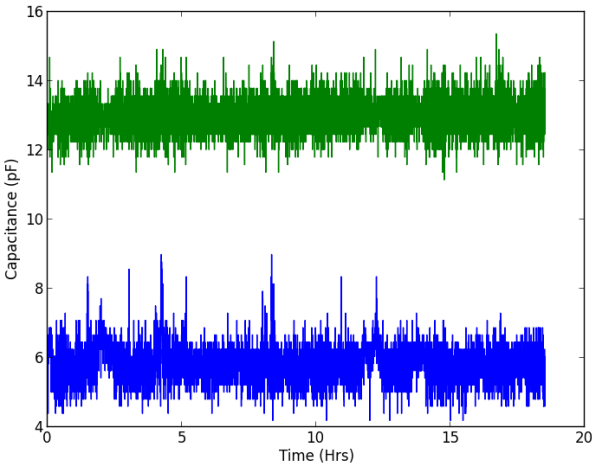

Figure 3.49: Level meter results of an overnight cooling from approx 300K to 170K.

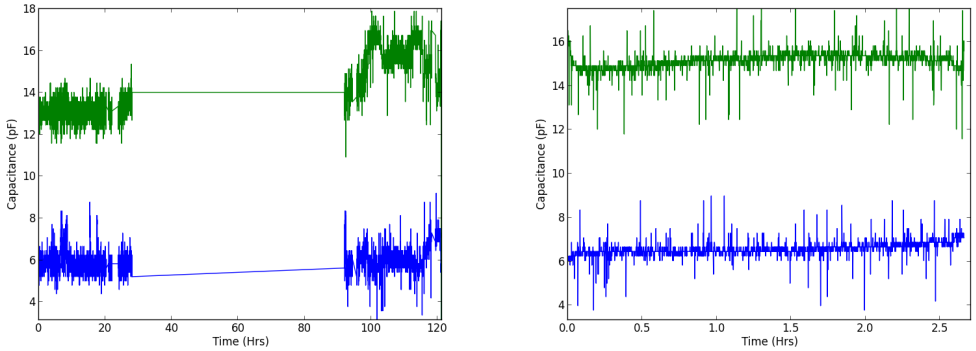

Figure 3.50: All capacitor data from filling 7.5Kg of LXe (left), and a zoom of a specific filling section for details on noise (right) .Each data point is one averaged over 50 seconds of data collection, with data being collected at a rate of 2 Hz

that we would then be able to observe the level more effectively. Unfortunately, this was not the case as can be seen in the figure below.

We observed less high noise data points using this internal capacitance level. We then attempted to push xenon inside the bell to determine if a change in level could be observed.

The DAQ system was then changed to readout through a UTI, using labview to read and store the data. The pump and readout issues that were present before now seem to have been eradicated. With 6kg of LXe inside the chamber, the recirculation was tested with a unregulated flow (controller set point=20slpm) through both the controller only, and with the bypass open, to observe the readings of the level meters. The results are shown below.

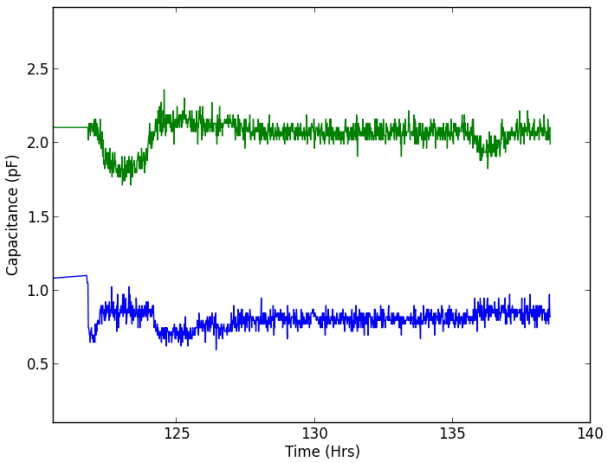

Figure 3.51: Capacitance of both level meters after altering the value for the internal capacitance of the system to reduce noise. Each data point is one averaged over 50 seconds of data collection, with data being collected at a rate of 2 Hz

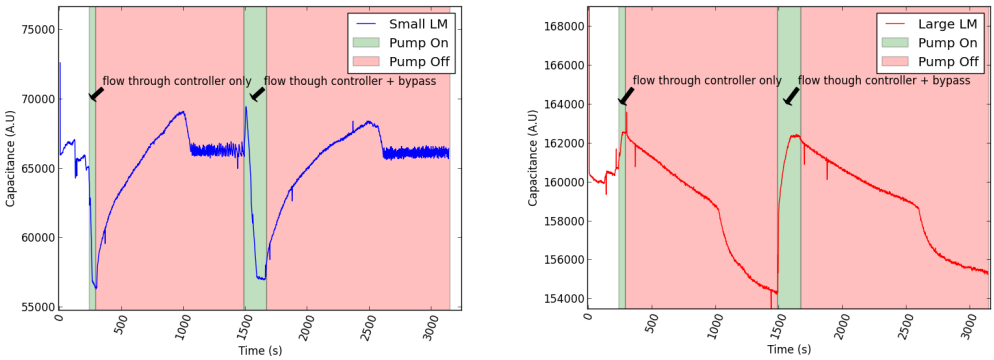

Figure 3.52: Level meter data with recirculation through the main chamber. The system was tested with flow both entirely through the flow controller and with the bypass open. Flow controller set point=20slpm (max).

Results show a marked decrease in the level within the bell, and an increase outside the bell while the recirculation pump is on, as expected. Furthermore the rate of increase/decrease was reduced with the bypass open, also in agreement with expectations. The level meter inside the bell shows some irregularities at a certain point, this is suspected to be an issue when the level reaches the connections on top of the capacitor.

Rough calibrations estimates show the variation from an empty bell to full produce a capacative difference in the small level meter by 69000-57000=12000 A.U. Over the 4cm

length, this corresponds to 3000 A.U/cm. For the large level meter such a calibration cannot be made, as we are unable to determine the height difference experiences by the capacitor. From the previous estimate, however, one can say that the approximate change in level outside the bell corresponds to (163000-154000)/3000=3.0cm.

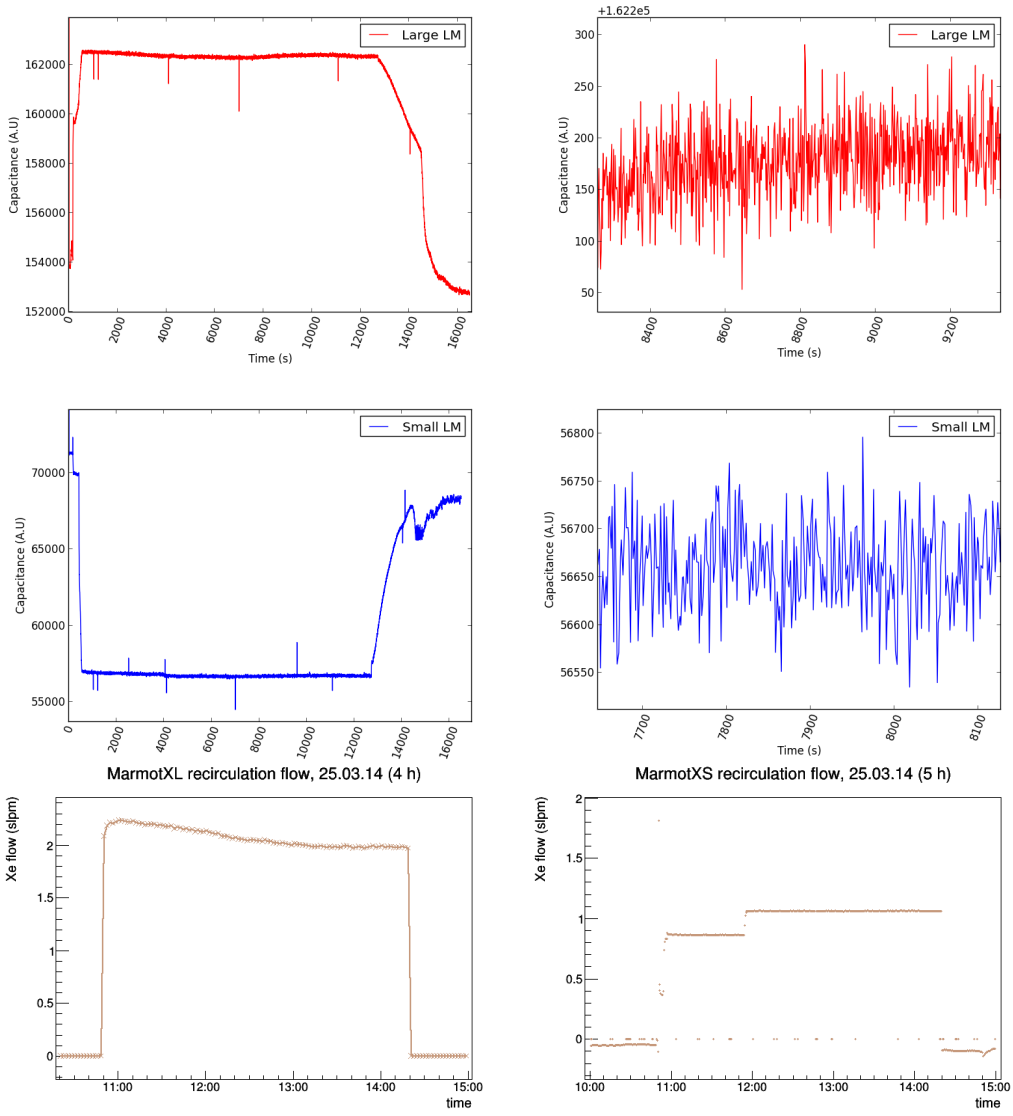

Figure 3.53: Recirculation test to determine stability of liquid level within the bell. Results from the exterior, long level meter are shown above. The left plots show the entirety of the test data, whereas the right plots show a zoomed in section during the recirculation to demonstrate the capacitor sensitivity variation. The central plots denote results from the large level meter, while the lower plots are from the small level meter. The lower two plots denote display the recirculation flow (left) and flow into the bell (right).

Using the earlier calibration estimate of 3000 A.U/cm, one can deduce from the fluctuations in the small level meter are approximately 150/3000=0.05cm→0.5mm. I would be possible to improve this vis averaging over several data points.

Averaging over 4 data points (approx 4 seconds per data point) we obtain the following distribution when zoomed.

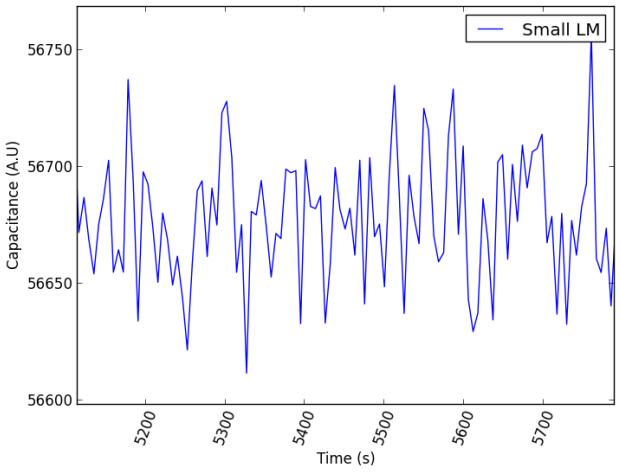

Figure 3.54: Small level meter distribution when averaged over 4 data points.

Using the standard deviation of the distribution of these points, we can determine the error in A.U for this data over the displayed range is 27. This corresponds to an error in distance of 27/3000=0.009cm, approx 0.1mm, a factor 5 improvement over the non-averaged data.

The Bell was next filled further with xenon to ensure a resting level above that height of the bell. Future results are performed with a total of 8.5kg of xenon.

The "saw tooth" like structure observed in the top right plot indicates a limit where the flow is not high enough to reach the bottom of the bell until the pressure rises, once the bottom has been reached, gas is released from the exterior pipe and thus the pressure inside the bell drops, raising the level.

3.5.5 The Level Meters and Their Characterisation

To ensure accurate calibration of the level meters, one must determine how the level meter values scale with capacitance. This study was performed by methodically placing the level meters into oil with steps of 10 mm, and observing the change in the output capacitance.

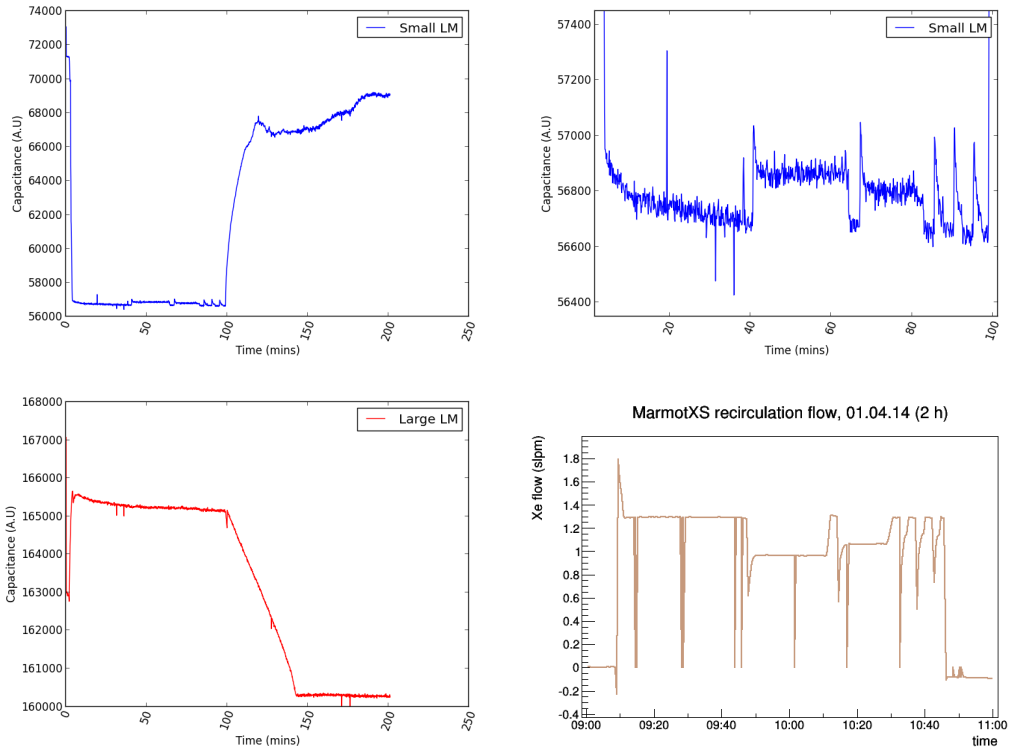

Figure 3.55: Results of recirculation test to determine minimum flow required to reach the bottom of the bell. This data was obtained at a flow of 1.9slpm. Data is averaged over 4 points, each data point is therefore separated by approximately 4 seconds.

Oil was chosen as the medium in which to perform these tests due to its high dielectric constant $\varepsilon = 2.2$ comparable to that of LXe.

The results can be observed in Figure 3.56. This Figure demonstrates that the level meters accurately return the liquid level with a linear response with capacitance. This characteristic is then used for all future calibration of the devices and is applied to the tests detailed below.

3.5.6 The Readout and DAQ

3.5.7 COMSOL Simulations on Heat Input

In the model, shown below in Figure 3.57, the bell itself is surrounded by a Xenon-like material placed at 165K. This therefore applied a constant temperature to the entirety of the bell at this value. The upper plane of the three screws are connected to a plate (not shown for clarity of model) that is at 200K.

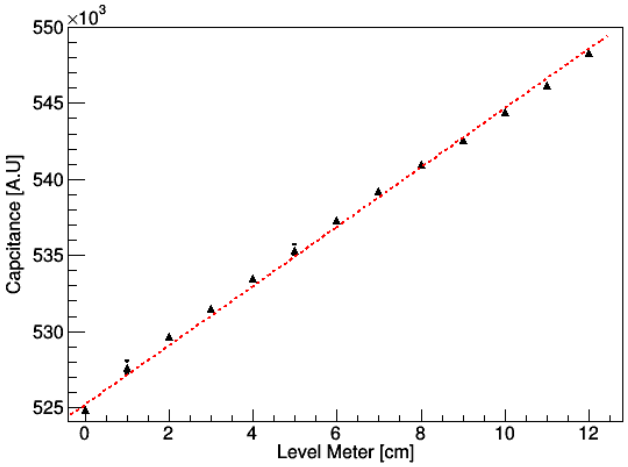

Figure 3.56: Testing the linearity of the long level meter with paraffin oil of $\varepsilon_r = 2.2$. The horizontal axis denotes the depth of oil the level meter was submerged into.

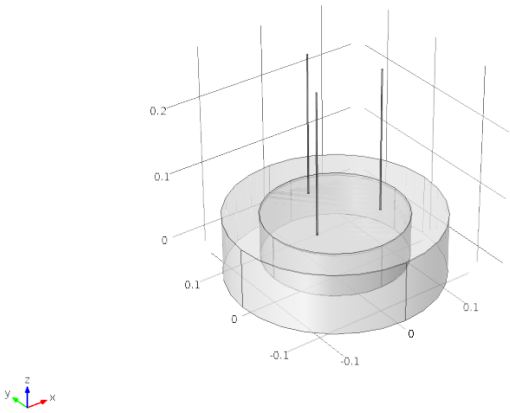

Figure 3.57: Model used for thermal simulations of the bell.

n order to determine at what point in time thermal equilibrium is achieved through the length of the screws, the temperature along the screw at each time step (1s for 1000s) was plotted to find the point at which the gradient is linear. The results of which are shown below. The upper most line indicates t=1s, the one below t=2s etc. This plot demonstrates that equilibrium can be achieved within 10s, and therefore the analysis can be comfortably performed at t=100s for security.

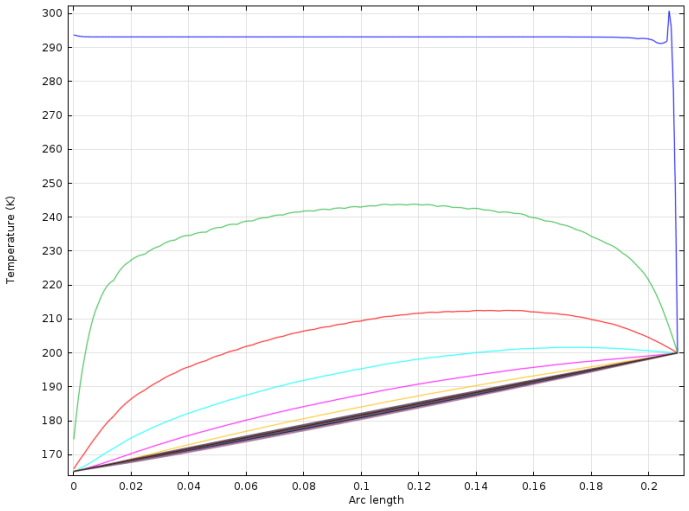

!!![h]

Figure 3.58: Temperature along length of one screw. The specific lines indicate a timestep.

At t=100s, the energy flux through the top of the screw can be plotted as a colour map, and shown below.

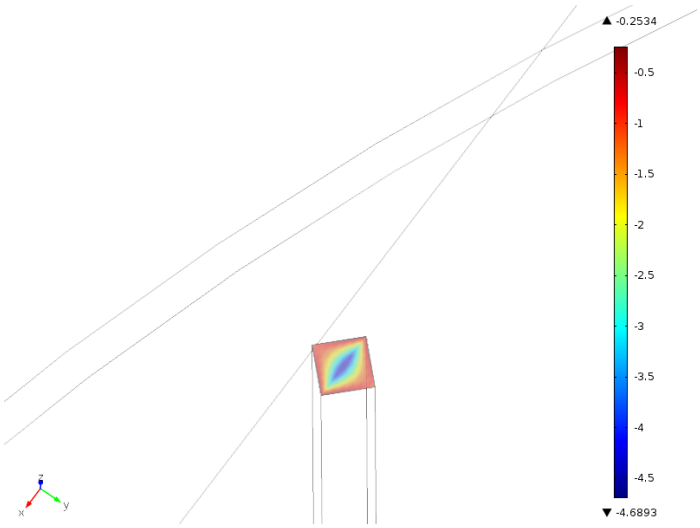

Figure 3.59: Colour map of the energy flux through the upper section of the screw in W/m

Taking an average value of $2.5\,Wm^{-}2$, and correcting for the cross-section of each screw, the total flux for all three can be calculated to be approximately $30\,\mu W$. This result is negligible to the predicted heat influx into the system of 1 W, and therefore can be ignored as a source of heat input.

3.5.8 Decoupling of Recirculation flow and Bell Input Flow

One potential advantage of using a bell to control the liquid level is the possibility to decouple the recirculation flow to that required to control the liquid level. To test this, two scenarios were implemented. Firstly, the bypass valve was closed, restricting all flow to enter the bell. This valve was then opened, and the resulting change in level within the bell was observed. Secondly, the flow was regulated into the bell with the use of the flow controller. The overall recirculation flow was then reduced, until this flow was lowered below that of the set point on the flow controller. The level both inside and outside of the bell was recorded over this period of time, and the stability of the liquid level was observed. Results of the decoupling of the recirculation flow from the flow into the bell can be observed in Figure 3.60. The fluctuation on the noise from the short level meter is shown in Figure 3.61. The stability of the liquid level shown in Figure 3.61 demonstrates that the recirculation flow can be altered accurately while maintaining a liquid level within 0.1 mm from the initial value.

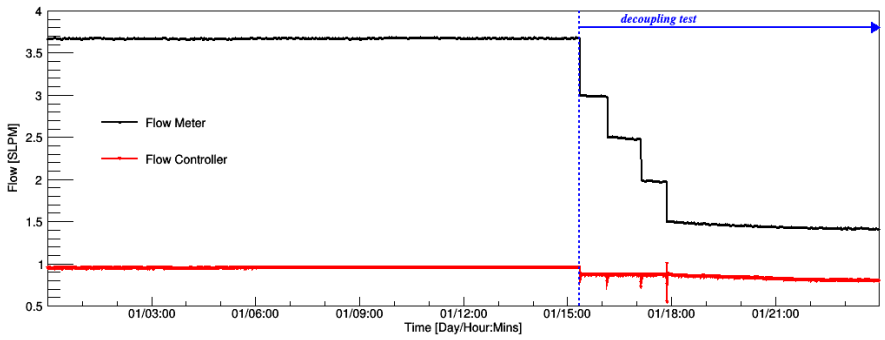

Figure 3.60: The overall recirculation flow (black curve) and the flow into the bell (red curve) against time for the period in which the recirculation flow was systematically reduced. The flow of xenon into the bell was maintained via the use of the flow controller.

3.5.9 Control Liquid Level With the Use of a Heater

3.5.10 Liquefaction From the Top Plate of the Bell

A major concern for large scale experiments is the heating power required to maintain a certain liquid level. This is dependent on the height difference between inside and outside of the bell, as well as the loss of heat into the bell resulting in liquefaction of xenon. The

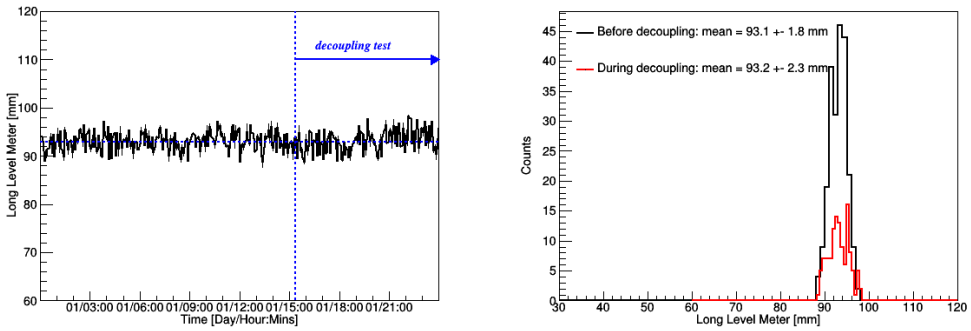

Figure 3.61: The performance of long level meter during the decoupling tests. Left: The calibrated value of the long level meter during the stable period prior to altering the recirculation flow. Right: The distribution of the measured liquid level position of the same period of time.

accurate determination of the amount of liquefaction within the bell and its origin will aid in the prediction of heating power required to maintain a set liquid level.

A candidate for the source of liquefaction is the top plate of the bell. l With a large surface area, covered by LXe, the temperature gradient between both sides of this plate can cause condensation. To test this hypothesis, the flow required to reach a liquid level within the bell was observed under two scenarios. One in which the liquid level outside of the bell covered the top plate, and one where the liquid level only in contact with the side walls. A schematic representation of these to scenarios can be observed in Figure 3.62

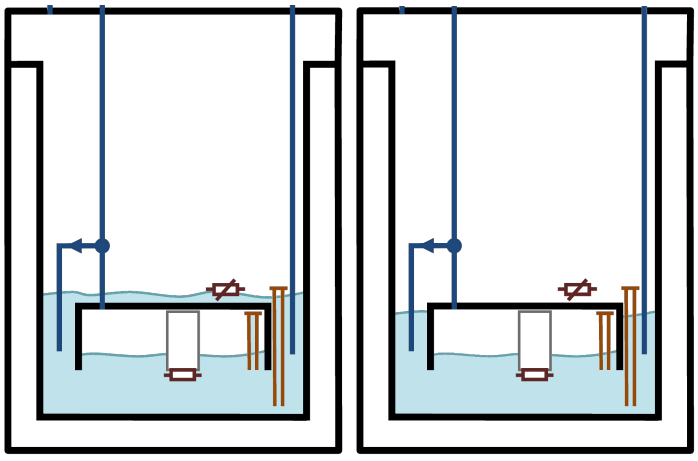

Figure 3.62: Schematic representation of scenarios in which the liquid level outside the bell covers the top plate of the bell (left) and in which the liquid level is below this point (right).

To achieve this, the liquid level was observed during the time of filling, and halted at the point at which the liquid level outside of the bell covered the top plate. The first tests were then performed before recuperating as little xenon as possible for the liquid level to fall below this point. The data from the flow meter and flow controller over the course of both tests can be observed in Figure 3.63. The unstable behaviour observed after the first test with negative flow is attributed to the recuperation process. The tests were then compared, and the results can be observed in Figure 3.64, whereby no observable difference is found. It can therefore be concluded from these results, that the primary source of liquefaction does not arise from the top plate of the bell.

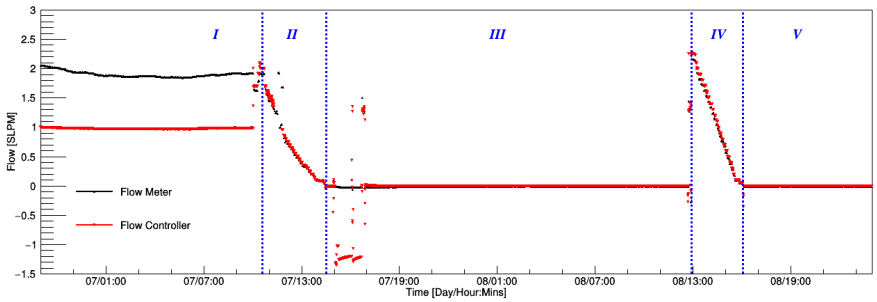

Figure 3.63: The overall recirculation flow (black curve) and the flow into the bell (red curve) against time for the period in which the liquefaction inside the bell was tested. The vertical dashed lines denote the transition of different periods of the tests. I: Equalisation of the liquid level inside and outside of the bell. II: Varying the flow into the bell while LXe lies above the top plate of the bell. III: Recuperation of LXe to being the liquid level outside of the bell below the top plate, and equalisation of the liquid level. IV: Varying the flow into the bell while LXe lies below the top plate of the bell. V: Post-test stabilisation.

3.5.11 Determining the Heat Input into The Bell

With the use of a heater to control the liquid level, one can estimate the heat input into bell inside our system. When no power is supplied to the heater, the natural heat input to the bell well maintain a non-zero liquid level difference between the inside and outside of the bell, as is shown in Figure 3.65. By then altering the power to the heater, one can determine the dependence of heating power to level difference. The process by which the liquid level difference within the bell was determined is shown in Figure 3.66.

Initially the pressures inside and outside of the bell were equalised by opening the appropriate valves through the recirculation system for a period of approximately one hour. This allows the liquid level within the bell to rise to that of the outside level. At this point, all valves are closed and the system is allowed to reach a state of equilibrium as can be seen

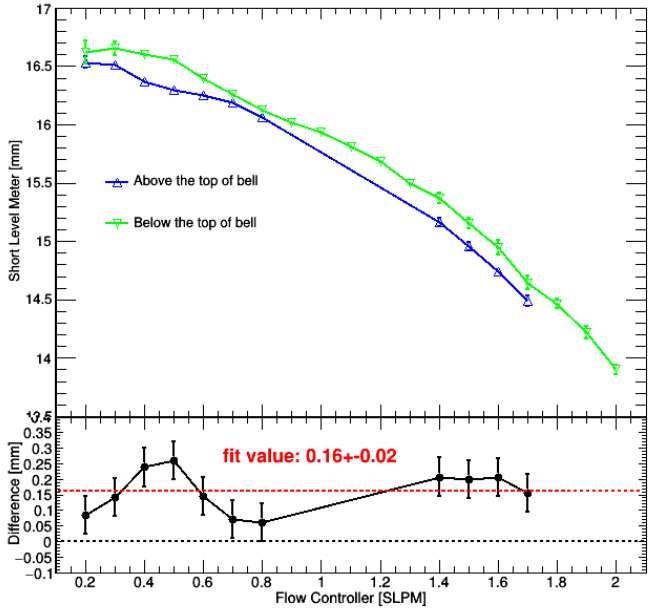

Figure 3.64: The liquid level inside the bell as a function of the flow into the bell for the cases in which liquid xenon was filled above the top plate of the bell (blue), and when it was below (green).

in Figure 3.65. Once the liquid level is seen to be stable, the heater is then activated, and the voltage is raised in steps of 0.5 V until the bell is completely emptied. The liquid level inside and outside of the bell is recorded over this period, and the height difference of the LXe inside and outside of the bell as a function of power supplied to the heater is plotted. A least chi^2 fit to the data is then performed, and extrapolated to the point of zero height difference. The power input to the bell is then determined. For the large bell, this was found to be 0.91±0.02 W. as show in Figure 3.67.

The tests were then repeated in an identical fashion for the small bell. The valves were closed, and the system allowed to reach equilibrium before the tests were performed. The voltage to the heater was then increased in steps of 0.5V and the liquid level inside and outside the bell was recorded. The level meter data from the short and long level meters, as well as the hight difference between the inside and outside of the bell can be observed in Figure 3.68. Using the information from the difference in liquid level, and the power supplied to the heater, Figure 3.69 displays the relationship between liquid level height difference and heater power. A fit to this data results that the heat input to the small bell is 0.76±0.01 W.

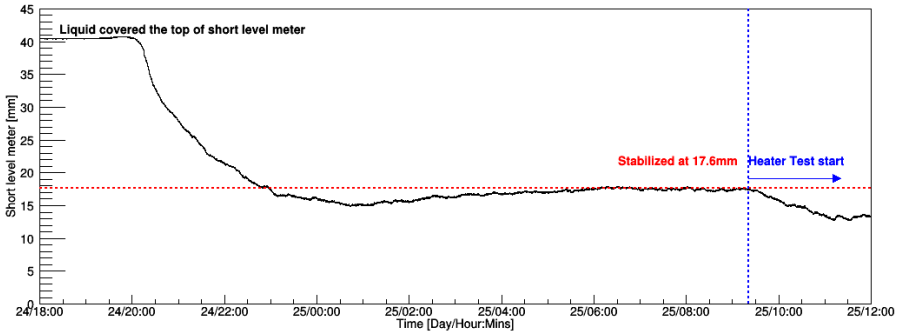

Figure 3.65: The level of liquid xenon inside the bell as a function of time starting just prior to closing all valves into, and out of the experimental chamber. The drop in level is caused by heat input into the bell until the liquefaction and heat input of the system stabilise.

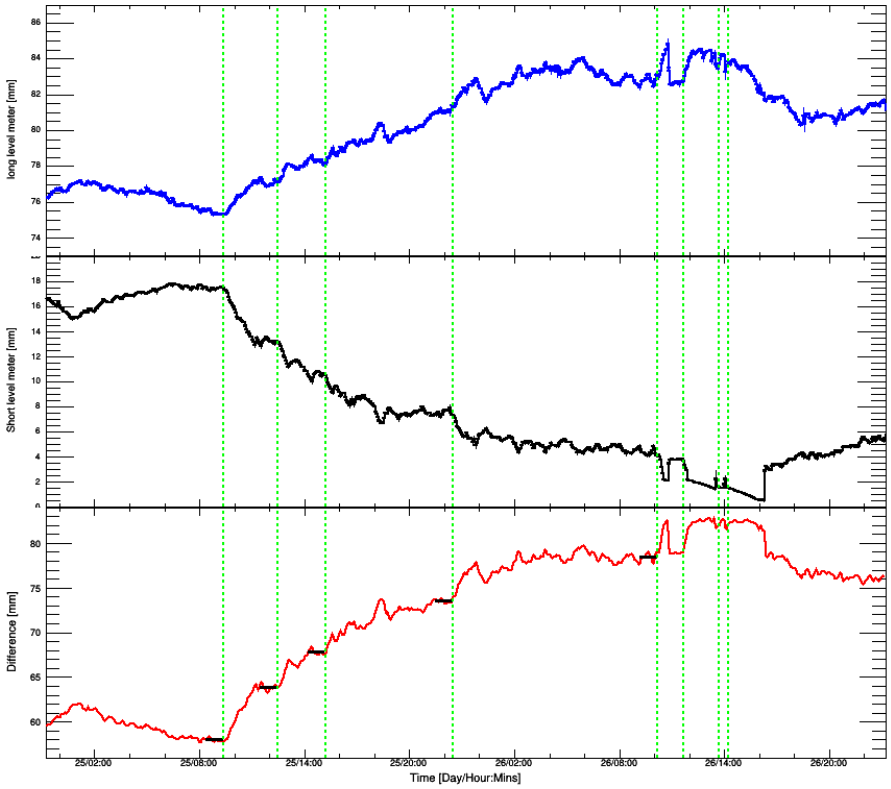

Figure 3.66: The performance of short and long level meter, and liquid level difference between inside and outside bell during the heat input test with the big bell. The vertical lines denote the time at which the power to the heater placed within the bell was changed.

This value can then be compared to that from the large bell, and extrapolated to XENON1T. Taking the results from the large bell, and extrapolating based on the area

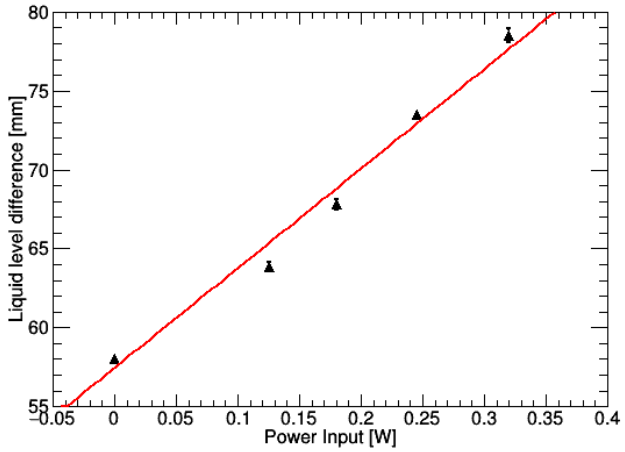

Figure 3.67: Liquid level difference between the inside and outside of the large bell against the power supplied to the heater within the bell to determine the external heat input with small bell.

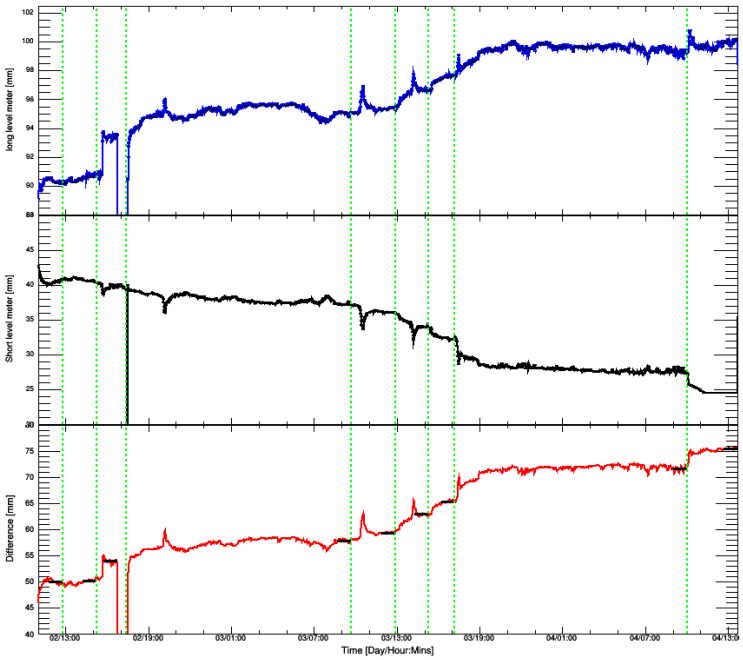

Figure 3.68: The performance of short and long level meter during heat input test with small bell.

results in a predicted heat input of 22.5±0.5 W. By using the results acquired from the small bell, one can linearly extrapolate the heat inputs up to the size of XENON1T. Results of such an extrapolation are presented in Figure 3.70 and predict a heat input for XENON1T of 11.3±1.4 W.

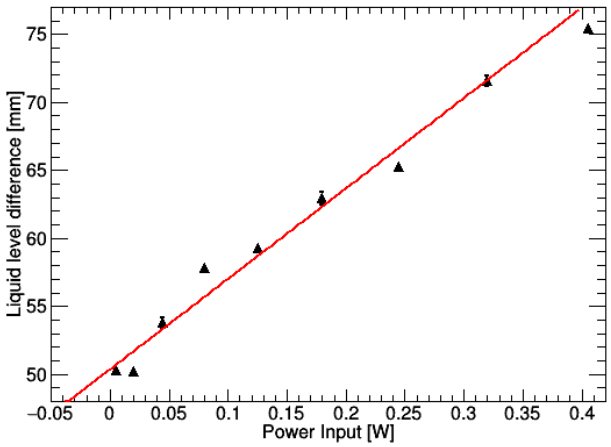

Figure 3.69: Liquid level difference between the inside and outside of the small bell against the power supplied to the heater within the bell to determine the external heat input with small bell.

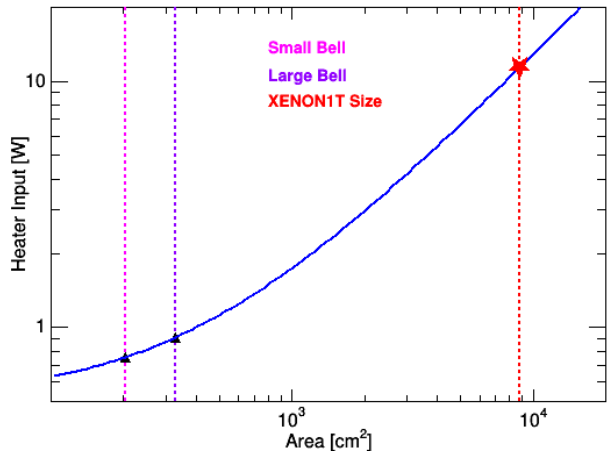

Figure 3.70: Extrapolation of heat input into the bell using results from two sizes of bell to XENON1T. A predicted value of 11.3±1.4 W is found from the fit.

3.6 Conclusion

This section will summarise the work performed in the research and development of systems for XENON1T.

Electric field simulations of the XENON1T TPC were performed with the intent to determine a design that would provide the most uniform field throughout the sensitive volume of the detector. It was first determined that, by comparing a full 3D simulations with a 2D

axisymmetric and full 2D simulation, that the 2D axisymmetric simulation would provide accurate results with respect to a full 3D simulation while maintaining the computing efficiency provided by a 2D one. Standardised analytical lines were then established in order to be able to compare results efficiently between one another. By altering the geometry, and spacing of the shaping rings as well as the potential difference between the cathode and first shaping ring, iterating between simulations resulted in an optimised TPC design. A further study into the effect on the electric field uniformity as a function of shaping ring radii compared how the electric field changed when some of the shaping rings were moved outwards by 1 mm. This was to

4

The Modulation Experiment

4.1 Detector Principle

The primary purpose of the Modulation experiment is to bring light to potential modulation of radioactive decay rates observed by other experiments. In order to accomplish this, the experiment must be able to obtain a sensitivity of order 0.1% on the amplitude of a modulation. Furthermore, to rule out the possibility of external, and environmental parameters affecting the results, measurements must be made of these and studies must be performed to determine the potential effects these could have on the observed rate. For the detector in Zurich, three radioactive sources will be measured, and compared to the background. These sources are Cs137, Co60, and Ti44.

4.1.1 Experimental Setup

Each of the three sources along with the background will each be monitored by two NaI(Tl) detectors placed opposite of each other, providing almost 4 coverage. A schematic of a detector pair is shown in Figure **??**. The sets of detectors will be enclosed in an "inner box." To exclude local variation, this inner box will be air tightly sealed, filled with radio pure nitrogen and kept at a steady temperature. In order to exclude as much natural radioactivity as possible, the radon level will be measured by using a RAD7 radon monitor. Sen- sors will also measure the pressure, temperature, humid- ity and magnetic field at a rate of 1 Hz. To lower the cosmic ray contribution, natural radioactivity and the gamma

rays of neighboring sources, the detector sets will be enclosed by a lead shield with a width of 5 cm.

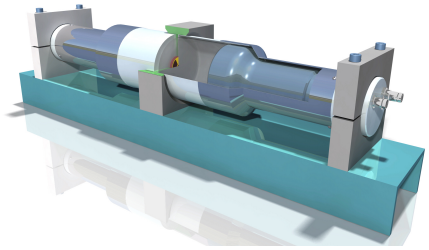

Figure 4.1: Schematic representation of one pair of detectors used for the modulation experiment.

An outer box will house the necessary electronics, including two CAEN DT5533P high voltage supplies, a NI-PXIe DAQ, a PID heater controller, a gas flow controller, and a RAID server. The external connections include 110 V power, ethernet and in the future, gas in- flow. The ethernet cable will be used to export processed data for faster analysis. A safety system will cut off the main power if one of the temperature monitors inside the inner or outer box measures a problem with the temperature. These safety measures are to make sure the heaters and power supplies will not damage the whole setup in case of a malfunction. A diagram displaying the inner, and outer box, as well as the electronics is shown in Figure 4.2

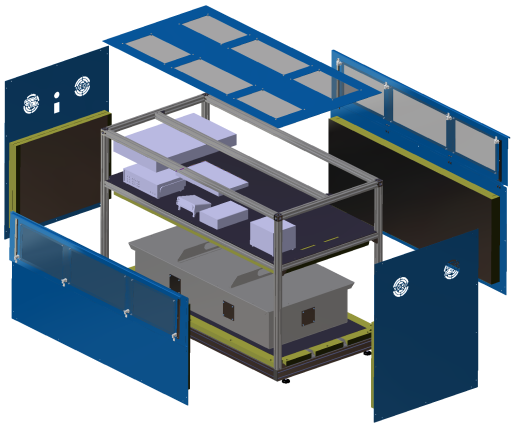

Figure 4.2: Technical disgram illustrating the inner, and outer box, as well as the positioning of the electronics with respect to them.

A variety of environmental parameters are read out and can be monitored individually in order to measure the influence of changes on the apparent measured decay rates. In

particular, temperature, pressure, humidity and the magnetic field are recorded with I2C sensors that are read out continuously through an Arduino board that is coupled directly to the DAQ via USB. We are able to control the temperature of the detectors to 0.1 C through the use of a heating pad and PID controller within the Styrofoam-insulated detector containment. This containment is kept at a slight overpressure in order to control the radon level, which will be monitored with a Durridge RAD7 monitor that is accurate to ±5%. Use of a Drystik unit allows the RAD7 monitor to be run continuously for months without the need to replace any desiccant under conditions of moderate humidity. Finally, we can monitor the line voltage and control the high voltage powering each detector directly through the multichannel high voltage power supply. All the afore-mentioned parameters are monitored and stored on a second to second basis, and are known as the slow control.

4.1.2 NaI Detectors

Eight Psionix 3" detectors are employed for use in the modulation experiment. These consist of a NaI scintillating crystal placed in front of a Photo-Multiplier Tube (PMT). An LED is placed inside the crystal itself for the purposes of stability. Connections to the detector consist of an SHV high voltage connector, BNC for signal, and LEMO for the operation of the LED.

Upon an interaction of a γ within the NaI crystal, scintillation light then interacts with the photocathode of the PMT whereby from the photo-electric effect, an electron is released. The high voltage which is applied to the detectors is split amongst each dynode within the PMT. The resulting electric field accelerates the electron into the first dynode, where it is multiplied and the resulting electrons are then accelerated to the next dynode. This process is repeated along each dynode until the signal is extracted at the anode. A schematic displaying the basic workings of the Nai detectors can be seen in Figure **??**

4.2 Processed Data Compression Studies

The theory behind the following analysis is to produce a set of fake data via mc simulations, and then to determine which of the following ways would be most appropriate to analyse fast and slow data to be able to create cuts based on the varying timestamps between them.

A gaussian curve placed on a uniform background is performed by a monte-carlo (MC) simulation. This was then then sampled at a rate determined by the temperature. The temperature data was taken at a rate of 1Hz and varied between 18-22 degrees in a sinusoidal function of period 100s. The temperature then determined the fast data acquisition rate from anywhere between 98-102Hz linearly correlated to temperature. Timestamps were

then produced via a MC simulation based on the current DAQ rate. Over 1,000,000 fast data data points were collected to provide statistical significance to the analysis.

Three separate techniques were implemented to store the fake data using root files:

1. Fast and slow data was stored in two separate trees in the same root file
2. Both data sets were saved in two root files.
3. All the data was saved into the same tree with the last slow data value used for each fast data point.

Method 1 intends to collect all the data into the same file, while maintaining the flexibility provided by separating the fast and slow data into different trees. Method 2 again separates the data, except into different files with the purpose of observing the efficiency in which root is able to compile both sets of data, and being able to compare this compression with method one. Method 3 combines all data, assigning a slow control value to every fast data event based on the slow control point taken. Although this method clearly saves a greater amount of data, it is thought that the compression from root may be capable of mitigating the difference between methods one and two.

To fill the gaps between timestamps on slow data and fast data for the former two methods, a root function was created with the slow data, producing a linear fit between each data point. This fit was then utilized to give an approximate value to the slow data at any given fast data timestamp, and thus was able to be analysed effectively. The file sizes of each method, along with the time take to produce three histograms using cuts on fast data from the slow data.

For the first test fast data was taken at an average rate of 100Hz, and slow data at a constant rate of 1Hz over a simulated 10,000s. This analysis was performed over 10,000 simulated seconds (2:45 hours). Methods 1 and 2 took approximately 45 seconds to produce the three histograms that were used during the time test. Method 3 took approximately 1, giving a significant time advantage to the analysis.

Method	File Size [MB]	Time Test [s]
1	6.014	45
2	5.969+0.053=6.022	45
3	6.126	1

In order to more accurately simulate the data. The DAQ rate for fast data was increased to a simulated average 10,000Hz, and humidity, pressure, and magnetic field were introduced

as parameters into the slow data, which was taken at a constant rate of 1/60Hz. The test was run for a simulated 600s. Using this improved simulation, the results were as follows:

Method	File Size [MB]	Time Test [s]
1	31.4	20
2	31.4+0.008≈31.4	20
3	32.2	2

It can be seen that with this simulated configuration, method 3 (last slow data value taken for each fast data) the file size comes to 2.5% larger than method 1 or 2. The ratio of analytical times of method 1 and 2, to method 3 for both simulations has improved due to the slower rate at which slow data had been taken with the smaller data acquisition period.

The simulation has been improved to include a simulated channel number (unsigned int), test pulse (boolean), pulse integral, base line and one other pulse parameter (all floats). The results of the simulation are below. Simulated run time was 10 minutes.

Method	File Size [MB]	Time Test [s]
1	108.8	-
2	108.8+0.008≈108.8	-
3	109.4	-

In this simulation, file size increase of method 3 in comparison to methods 1 and 2 has been reduced to 0.6%.

Using data from the tests described above, it can be concluded that by storing slow data for each fast data acquisition provides the best compromise between disk space and analysis performance. That is to say that benefits gained in terms of ease and time outweigh the small increase, as low as 0.6%, in file size in comparison to separating fast and slow data.

4.3 Characterisation of Detectors

4.3.1 Energy Calibration

The following calibration data was taken with detectors at 600V prior to gain calibration. Three sources were used for the calibration, Co57, Cs137, and Na22 providing 4 peaks at 122 keV, 662 keV, 511 keV, and 1274 keV. The spectrum of the integral of each waveform was then produced, and the respective peaks identified and fit using a Gaussian. Knowing the value of the energy for each peak from literature, the relationship between the integral and the predicted energy can then be produced. The results are shown in figure **??**.

All detectors are shown to have a linear response correlation between waveform integral and energy, as would be expected. A summary detailed the resultant fits to each detector

4.3.2 Voltage Calibration

This section details the steps performed to determine the optimal voltage to use for each NaI detector, Each calibration has been performed using the respective source that the detector is to be used for, such that we can ensure the calibration is optimised for its working conditions.

The voltage for each detector was varies around its recommended value. For each voltage a dataset of 500,000 events was taken, and the full absorption peak for the source was then identified from the integral spectrum and fit with a simple Gaussian. The mean of this Gaussian was then uses to scale the energy into keV, for the sake of comparison. The resolution was then calculated by FWHM/mean of the source peak. In the cases in which the source has multiple peaks, only one is chosen for the sake of the calibration. For $^{6}0$Co the 1173 keV peak is selected, and for $^{4}4$Ti, the electron-positron annihilation peak at 511 keV is chosen as the low energy peaks from the source itself are eliminated by the threshold for low voltages.

The method by which the appropriate voltage for each detector is determined is based on the shape of the resolution vs voltage curve. As voltage is increased, the resolution is expected to improve to a finite point. This plateau indicated the optimal resolution for the detector. In order to ensure that any drift in the performance of the detectors, or in the voltage supplied, a point an appropriate distance from the beginning of this plateau is used for the recommended voltage.

Measurements were performed by starting from when signals could first be observed, and then increasing the voltage in steps of 20 V until the voltage reached was 200 V above that specified by the manufacturer.

The figure below displays the results of the calibration for each individual detector, along with the selected voltage. The individual spectra for each calibration can be observed in the appendix.

With the voltages for each detector determined, it is now possible to apply energy thresholds from which the presence of a signal will be determined. This was performed by systematically increasing the threshold from the minimum value until such point at which the noise peak was removed.

Detector	Calibrated Voltage [V]	Threshold [A.U]
SBL256	680	800
SBL257	720	850
SBL258	760	950
SBL259	760	950
SBL260	700	900
SBL261	700	850
SBL262	840	900
SBL263	710	1000

Table 4.1: Summary of the voltage calibration and threshold determination tests.

4.4 Conclusions

To summarise the work performed, the appropriate format of data compression was studied. This entailed the comparison of three methods of data storage for the fast and slow data. Those were to save the datasets in the same file, but in different trees, into entirely separate files, or to save all data into the same root tree, using the last slow data point to assign a value to all fast data events. After producing fake data for 3 different fast and slow data rates, and parameters, it was determined that ROOT is capable of compressing data adequately such that storing a slow control value increased the file size by 0.6%, which is outweighed by the benefits gained during the analysis. As a result, this was the storage method decided upon.

The detectors were characterised to determine the appropriate voltage on which to power them, and their energy calibration. To determine the correct voltage, the resolution of a source was measured for each detector around the recommended voltage given by the producer. The curve of resolution vs voltage drops to a plateau at higher voltages, and thus a value suitably within the region of the plateau was selected as the operating voltage of each detector.

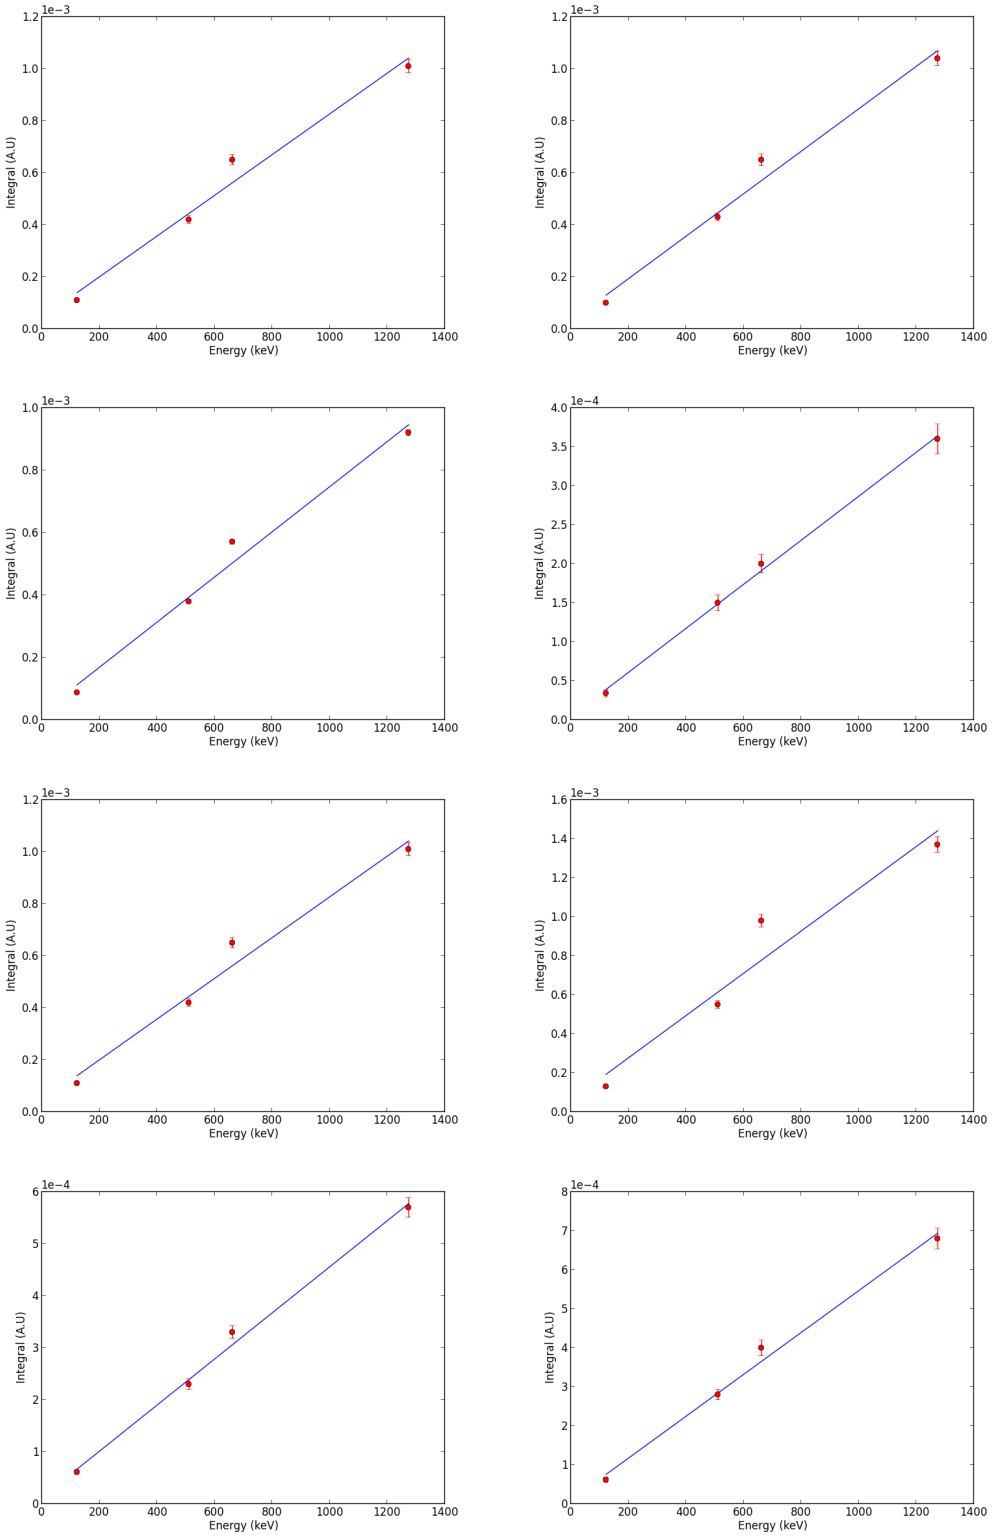

Figure 4.3: Reading left to right, the energy calibration results for NaI detectors SBL256-263. The position of each peak is calculated via the mean of the fit Gaussian for the sources of Co57, Cs137, and Na22 providing 4 peaks at 122 keV, 662 keV, 511 keV, and 1274 keV. The integral is calculated via the summation of each bin in every waveform.

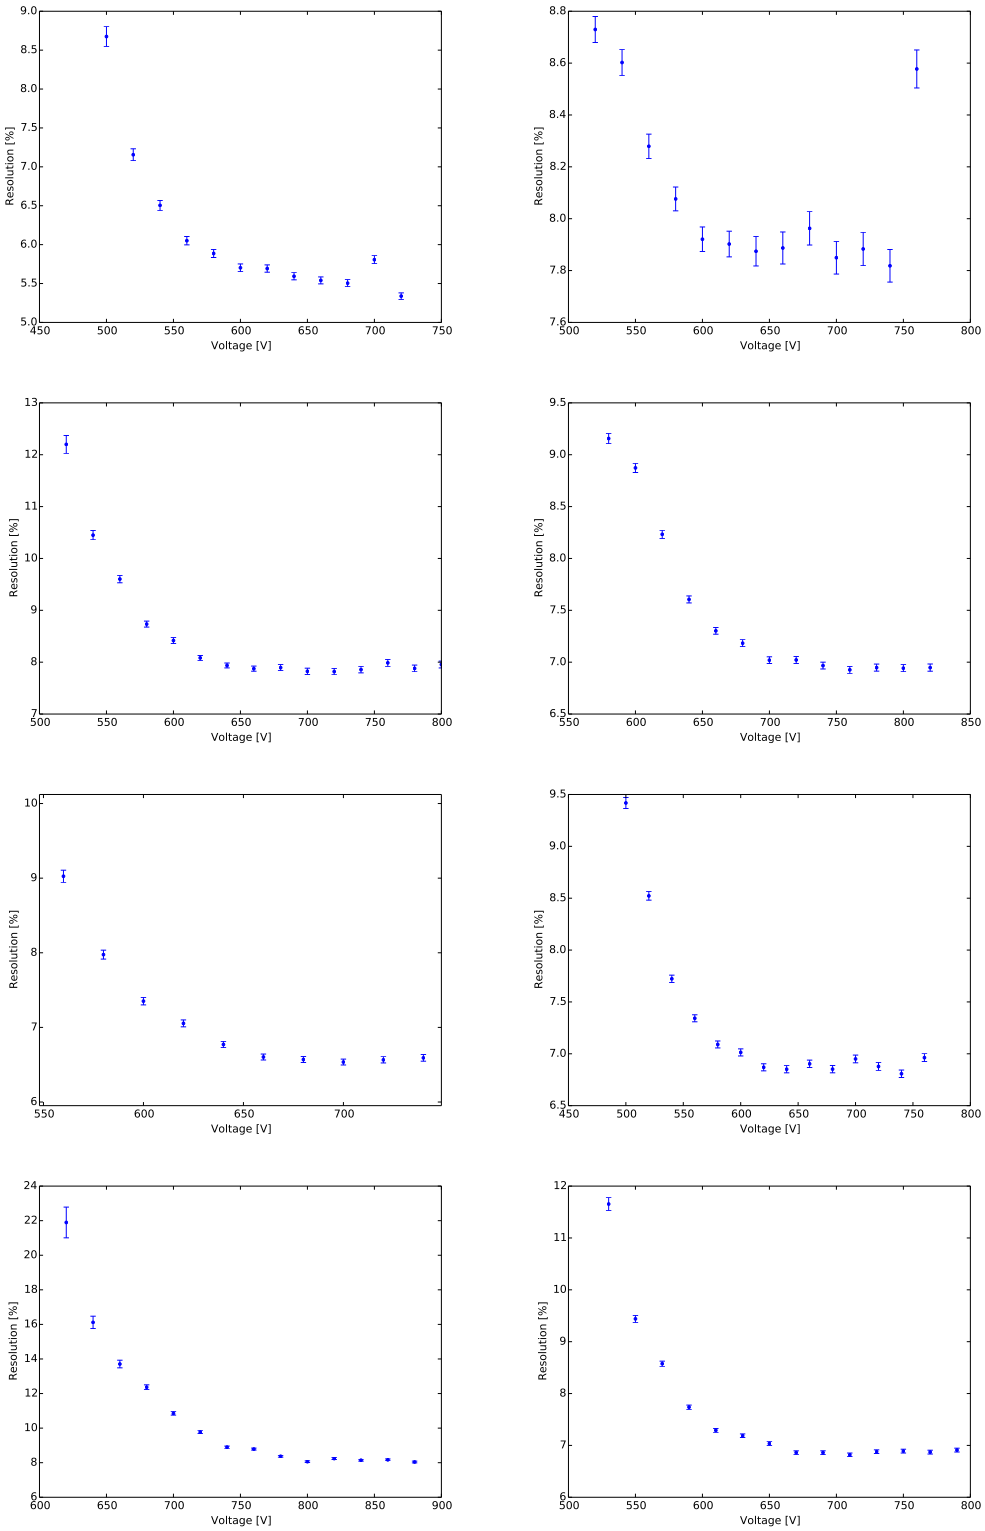

Figure 4.4: Reading left to right, the voltage calibrations for NaI detectors SBL256-263. The resolution is given by the ratio between the peak height and the FWHM of the corresponding peak used. The vertical red line denotes the selected voltage for each detector.

References

[1] Maxim Pospelov, Adam Ritz, and Mikhail Voloshin. Bosonic super-wimps as kev-scale dark matter. *Physical Review D*, 78(11):115012, 2008.

[2] Helmut Siegert, Heinrich Schrader, and Ulrich Schötzig. Half-life measurements of europium radionuclides and the long-term stability of detectors. *Applied Radiation and Isotopes*, 49(9):1397–1401, 1998.

[3] Jere H Jenkins and Ephraim Fischbach. Perturbation of nuclear decay rates during the solar flare of 2006 december 13. *Astroparticle Physics*, 31(6):407–411, 2009.

[4] Fritz Zwicky. Die rotverschiebung von extragalaktischen nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.

[5] Marisa Girardi, Giuliano Giuricin, Fabio Mardirossian, Marino Mezzetti, and Walter Boschin. Optical mass estimates of galaxy clusters. *The Astrophysical Journal*, 505(1):74, 1998.

[6] Vera C Rubin and W Kent Ford Jr. Rotation of the andromeda nebula from a spectroscopic survey of emission regions. *The Astrophysical Journal*, 159:379, 1970.

[7] Gary Steigman and Michael S Turner. Cosmological constraints on the properties of weakly interacting massive particles. *Nuclear Physics B*, 253:375–386, 1985.

[8] G Jungman, M Kamionkowski, and K Griest. Supersymmetric dark matter, phys. rept. 267 (1996) 195–373. *arXiv preprint hep-ph/9506380*, 2.

[9] Hsin-Chia Cheng, Jonathan L Feng, and Konstantin T Matchev. Kaluza-klein dark matter. *Physical review letters*, 89(21):211301, 2002.

[10] M Aguilar, G Alberti, B Alpat, A Alvino, G Ambrosi, K Andeen, H Anderhub, L Arruda, P Azzarello, A Bachlechner, et al. First result from the alpha magnetic spectrometer on the international space station: Precision measurement of the positron fraction in primary cosmic rays of 0.5–350 gev. *Physical Review Letters*, 110(14):141102, 2013.

[11] E Aprile, J Angle, F Arneodo, L Baudis, A Bernstein, A Bolozdynya, P Brusov, LCC Coelho, CE Dahl, L DeViveiros, et al. Design and performance of the xenon10 dark matter experiment. *arXiv preprint arXiv:1001.2834*, 2010.

[12] P Benetti, R Acciarri, F Adamo, B Baibussinov, M Baldo-Ceolin, M Belluco, F Calaprice, E Calligarich, M Cambiaghi, F Carbonara, et al. First results from a dark matter search with liquid argon at 87k in the gran sasso underground laboratory. *Astroparticle Physics*, 28(6):495–507, 2008.

[13] E Aprile, M Alfonsi, K Arisaka, F Arneodo, C Balan, L Baudis, B Bauermeister, A Behrens, P Beltrame, K Bokeloh, et al. Dark matter results from 225 live days of xenon100 data. *Physical review letters*, 109(18):181301, 2012.

[14] DS Akerib, HM Araujo, X Bai, AJ Bailey, J Balajthy, S Bedikian, E Bernard, A Bernstein, A Bolozdynya, A Bradley, et al. First results from the lux dark matter experiment at the sanford underground research facility. *arXiv preprint arXiv:1310.8214*, 2013.

[15] Christopher Savage, Katherine Freese, Paolo Gondolo, and Douglas Spolyar. Compatibility of dama. *LIBRA dark matter detection with other searches [JCAP 0904: 010*, 2009.

[16] CE Aalseth et al. Cogent collaboration et al. *arXiv preprint arXiv:1002.4703*, 2011.

[17] G Angloher, M Bauer, I Bavykina, A Bento, C Bucci, C Ciemniak, G Deuter, F von Feilitzsch, D Hauff, P Huff, et al. Results from 730 kg days of the cresst-ii dark matter search. *The European Physical Journal C*, 72(4):1–22, 2012.

[18] Samuel B Garfinkel. Semiautomatic townsend balance system. *Review of Scientific Instruments*, 30(6):439–442, 1959.