Royal Holloway University of London

Physics Department

Submitted for the degree of Doctor of Philosophy

An Analysis of a Data Set from a Prototype Dark Matter Detector and the Design, Simulation and Construction of a New Detector

Author:
Richard L. Eggleston

Supervisor:
Dr. Jocelyn Monroe

June 22, 2016

Declaration of Authorship

I Richard Eggleston hereby declare that this thesis and the work presented within is entirely my own. Where I have consulted the work of others, this is always clearly stated.

Signed: 

Date June 22, 2016

Abstract

Dark matter is considered one of the most significant outstanding problems in modern physics. Over the years a number of experimental techniques have been developed with the aim of making a direct detection. The Dark Matter Time Projection Chamber (DMTPC) collaboration uses gasesous CF_4 in a time projection chamber (TPC) with charge and imaging readout. The work presented in this thesis represents an analysis on an existing detector with the aim of understanding the backgrounds that are present and improving techniques to reject them from the data. Following this, the design, fabrication and commissioning of a new large-scale detector is described, with focus on reduction of backgrounds from component materials.

The process of analysing the data, results in an improvement to the background rejection methods. A reduction in the overall rate of events passing the selection criteria is observed, verifying the improvement. The upper limit on the spin-dependent dark matter-proton interaction cross-section is comparable to the previous analysis despite having lower total exposure. The value achieved is $\sigma_p^{SD} = 6.88 \times 10^{-33}$ cm^2 for a dark matter mass of 145 Gev/c^2. A background estimation is presented providing evidence that the remaining candidate events can be attributed to radon progeny recoils, present due to radon emanating from and plated out on the materials that comprise the detector. This finding is used to drive the design of the next detector. This is done by producing a metric which satisfies the desired fiducial volume whilst minimising the surface area contributions of materials. The final field cage design reduces the surface-area to fiducial-volume ratio by a factor of 9.5 compared to the previous detector prototype and a factor of 28 compared to the first prototoype (the detector used for the analysis of this thesis).

Acknowledgements

I would like to thank all the people that have supported me in the process of doing the work that is presented in this thesis as well as during the period of writing. Of particular note are: my supervisor, Dr. Jocelyn Monroe, for giving me the opportunity to work on the DMTPC experiment. Also the RCUK and STFC for providing the funds to allow me to do the work. Also, previous and current DMTPC collaborators, especially Cosmin Deaconu, Shawn Henderson, Jeremy Lopez, Asher Kaboth, Hidefumi Tomita, Ross Corliss, Michael Leyton and Natalia Guerrero. I would also like to thank the team of engineers who helped with both the design and fabrication of the parts used for this thesis, especially Richard Elsom and Paul Bamford of Royal Holloway, Andrew Gallant of MIT and Jim Kelsey, Chris Vidal and Erni Ilhoff of Bates Laboratory. Dr Chamkaur Ghag also offered support during my revisions, for which I am most appreciative. I am also grateful to Alistair Butcher and Gabriela Druitt, my colleagues at Royal Holloway, my family and my partner Jamie Elmer for their support and patience throughout the whole process.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This chapter is dedicated to giving an overview of the context of this thesis: what is dark matter, what can currently be said about it and what are the methods employed to detect it? In section 1.1 a brief historical background is given of the first hints that additional matter might be a necessary input into our physical theories. This is followed by a discussion of observed phenomena and the inferences that can be made. Next a selection of proposed particles are described that have been put forward as possible candidates for the dark matter constituent. Section 1.2 outlines the methods of dark matter detection, describing how the observed rate of interactions within a detector is used to produce an upper limit on the strength of the dark matter interaction. Section 1.3 describes the theoretical aspects of producing the result.

Chapter 2 provides an overview of DMTPC, giving the experimental considerations and techniques employed. Included is the fundamental composition of the detectors, what a dark matter interaction looks like in a detector and how the resulting signal is readout and recorded. Section 2.2 gives an overview of the types of backgrounds that can occur in the experiment. Section 2.3 gives a description of the detectors that have been built by the collaboration. Section 4.2 describes the simulation of the detector output.

Having provided a background to the theory and experimental techniques of dark matter detection and an outline of the DMTPC experimental status, chapter 3 goes on to describe the analysis of an underground DMTPC data set. The chapter contains descriptions of the detector calibration and the data-selection cuts used to discriminate between events that are signal-like and those that are due to backgrounds. The analysis process is described and then the results of the new analysis are presented, including comparison to previous studies. Subsequently, an investigation into the radiological background contribution is made.

In chapter 4 the focus moves away from data calibration and analysis and shifts toward detector design. As will be seen, the results of the analysis of chapter 3 have fed into the design requirements for the internals of the next generation detector. The chapter deals with the design, simulation and mechanical prototyping of this new configuration. Chapter 5 discusses the current status of the new design and the progress that has been made. Chapter 6 then summarises the work presented.

1.1 Dark Matter: Evidence Overview

Before discussing the evidence for dark matter, it is pertinent to place it in the context of what we can already say about this vast and magnificent Universe we call home. This section starts by providing a brief summary of our current understanding of cosmology.

1.1.1 The Λ-CDM Model

The concordance model of the evolution of the Universe is the Λ-CDM model. This is a cosmological model that includes the Big Bang, inflation, a cosmological constant (Λ), dark energy[1] and cold dark matter (CDM). The full details of the model are

[1] An energy which is thought to be causing the Universe expansion to accelerate

outside the scope of this work, so only a summary is provided (for more detailed discussions see [1, 2, 3]). The key principles that go into the model are:

1. The Copernican Principle: we are not in a special place within the Universe.

2. The Priniciple of Equivalence: inertial and gravitational mass are identical.

3. The Cosmological Principle: the Universe is homogeneous and isotropic.

The first of these derives from the move from a geo-centric to a helio-centric view of the Solar system. It is then a natural extension to assume that our Solar system and galaxy are not located in some special or unique part of the Universe. The second is a fundamental principle of general relativity, which is assumed true for this model[2]. The final principle comes from observations of regions of the sky, made by various telescopes, most significantly the Hubble telescope, and the cosmic microwave background (discussed later), that show the average density of galaxies is uniform across the sky (homogeneity) and that no one location should observe the Universe to be any different to any other (isotropy).

Combining the equations of gravitational fields from general relativity with the Cosmological Principle produces the Friedmann equations, which describe the evolution of the Universe (equations 1.1-1.3).

$$\frac{\ddot{a}}{a} = \frac{4\pi}{3}(\rho + 3p), \tag{1.1}$$

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi}{3}\rho - ka^{-2}, \tag{1.2}$$

$$\dot{\rho} = -3\frac{\dot{a}}{a}(\rho + p), \tag{1.3}$$

[2]General relativity is a very well tested theory of gravity however there is as yet no quantum theory for gravity, it is therefore considered that general relativity may not be the full story.

where a is the scale factor, defined as $a(t) = \frac{x_t}{x_0}$, with x being the distance measured. a is a time dependent variable and describes the relative expansion of the Universe, with t being the time under consideration and 0 being the the value at the present time. The time-dependence is necessary due to the observed expansion of the Universe which was deduced by comparing the relative velocities, or red-shifts, of distant objects and their distances [4]. The other variables in the Friedmann equations are; ρ, the density of the Universe; p, the pressure and k, a parameter which represents the curvature of space-time. There are three types of curvature the universe could take; positive, negative or zero-valued. The curvature can be determined by studying fluctuations in the cosmic microwave background, as its value will affect their apperent size. Figure 1-1 shows this effect. The current measurements from Planck report a value of < 0.005, indicating the geometry of our Universe is almost, if not completely, flat [5]. The value it takes has implications for the evolution of the universe, as is shown in figure 1-2.

The Friedmann equations help us to understand the dynamics of the Universe and the evolution of matter, from its beginning to its final state. This is achieved by combining them with the equation of state, the relation between pressure and density, for the different types of mass-energy components of the Universe. This equation is typically quite simple and it is sufficient to use equation 1.4 with values for the proportionality constant σ listed in table 1.1 (reproduced from reference [1]).

$$p = \sigma\rho \tag{1.4}$$

From here the Λ-CDM model is constructed. The density parameter can be separated into contributions from radiation and matter, as such: $\rho = \rho_m + \rho_r$. There are then three epochs that can be considered: the radiation-dominated, the matter dominated and the vacuum dominated. The first is the very early Universe when the energies were so high that particle annihilation and creation where constantly taking

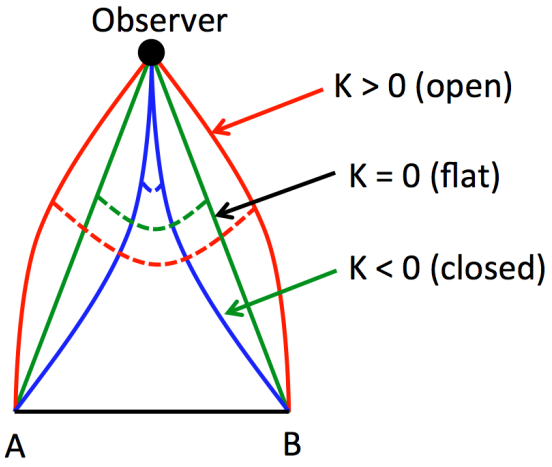

Figure 1-1: A diagram to demonstrate how the perceived angular size of a distant object is affected by the curvature of the Universe. The line AB represents the object being observed. The solid coloured lines represent the path taken by a photon under different geometrical conditions. The dashed lines show the change in angular size.

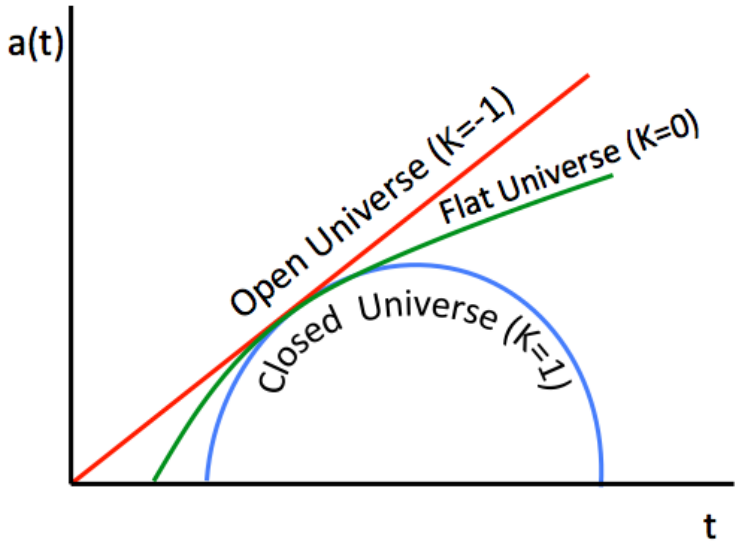

Figure 1-2: A representation of the impact of the curvature parameter on the evolution of the Universe. The abscissa is time, the ordinate is the scale factor $a(t)$. The red line represents an open Universe, the green line represents a flat Universe. The blue line shows a closed Universe. An open Universe will continue to expand indefinitely, a closed Universe will eventually collapse in on itself. A flat Universe will also continue to expand indefinitely. This is due to the presence of dark energy.

Table 1.1: Proportionality constants for different energy forms

σ	Form of Energy Density
0	Nonrelativistic matter (dust)
$+\frac{1}{3}$	Relativistic matter (e.g. neutrinos) Radiation (photons, gravitational waves)
+1	Free massless scalar fields Shear energy
$-\frac{1}{3}$	Curvature energy Cosmic string networks
$-\frac{2}{3}$	Domain walls (related to changes in the vacuum configuration, causing regions of differing lowest-energy states for a field ϕ)
-1	Cosmological constant

place. During this time the Universe expanded and cooled to the point where the temperature was less than twice that of the mass of the electron, meaning that more electron and positrons were being produced than were being annihilated. This caused an abundance of electrons and positrons (in comparison to the nucleon-anti-nucleon abundance) and the weak force was the dominant mediator. This led to the radiation-matter equality epoch, which occured at around 1,500 years after the big bang. Then, about 70,000 years after the big bang, matter began to dominate. 380,000 years after the big bang, due to the continuing expansion, the density reduced sufficiently to allow for fewer and fewer interactions of photons with the surrounding matter. This is the point at which the cosmic microwave background was produced - more of which is discussed in section 1.1.5.

In this model, 85% of the mass content of the Universe is currently missing [5]. The remainder of this section gives an overview of the evidence that has accumulated for the existence of a non-baryonic, particulate component of the Universe.

1.1.2 Galactic Motion

One of the earliest observations of a mass discrepancy was made by Jon Oort in 1932 [6]. Oort made use of the Jeans Equations, which relate the motions of the stars within a galaxy to the gravitational force, to measure the velocity of stars within the Milky Way. After doing this Oort identified that there was insufficient luminous mass to account for the motions he observed and thus proposed that invisible matter could account for the discrepancy - he described this invisible matter as 'dark matter'. A year later, Fritz Zwicky studied the Coma galaxy cluster [7]. Zwicky used a different technique for his measurements, making use of the Virial Theorem, which relates the total average kinetic energy of a system to the total potential energy. He too found that the observable mass could not account for the motions of the galaxies within the cluster. Zwicky also mentioned the effect of gravitational lensing as a useful method for obtaining a mass measurement, though he did not implement its use in this case.

1.1.3 Gravitational Lensing

Gravitational lensing is based on Einstein's General Theory of Relativity and the principle of least action. The result is that light from a distant source will be bent when passing a strong gravitational field between the source and observer. The gravitational field produced by the distribution of matter is referred to as a lens. This effect can be categorised into strong and weak lensing. Strong is when the lens is strong enough to create distortions in the observed image of objects behind the lens. Weak lensing is when the lens is only sufficiently strong to distort the objects surrounding the lens. Both can be used for obtaining a mass measurement. The most well-known evidence for dark matter, which uses weak lensing measurements, is the Bullet cluster (figure 1-3). This image shows the mass distribution following the collision of two galaxy clusters with two coloured regions on each side. The colours have been added to highlight the different types of matter: pink is the luminous and blue the

non-luminous. The luminous matter is deduced via x-ray emission from hot gas and the dark matter is inferred from the mass distribution, calculated via gravitational lensing, as shown in figure 1-4.

These two figures demonstrate that the luminous (pink) matter has interacted, as can be seen by the co-mingling of matter in the central region of each image. There is also a distinct cone-like shape visible just off the centre, to the right, which illustrates a drag-like effect as the luminous matter regions collide. The blue regions, however are well separated suggesting that there was no disturbance caused as the galaxies crossed paths, thus if the matter that constitutes these regions interacts at all, it can only be very weakly. The indication from this is that there is a large component of massive, non-luminous, weakly-interacting matter that was surrounding each galaxy, prior to the collision.

Figure 1-3: The Bullet cluster with artificially coloured regions [8, 9]. Blue highlights non-luminous matter, pink highlights the luminous.

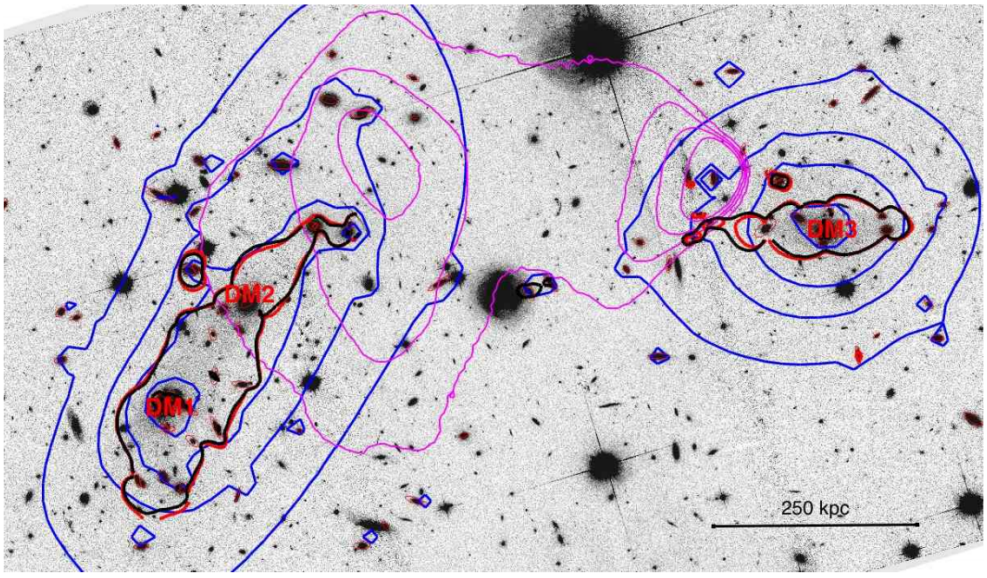

Figure 1-4: The Bullet cluster mass distributions as deduced by gravitational lensing [10]. The blue and pink contours are iso-mass lines, connecting points with the same mass value.

1.1.4 Rotational Velocity Curves

The next major study to suggest missing mass used galactic rotational velocity curves and was performed by Vera Rubin [11]. These curves show the angular velocity of stars as a function of distance from the galactic centre. The expectation, given the distribution of visible baryonic mass, is that the further a star is from the galactic centre, the slower its velocity. This expectation arises from equating the gravitational force to the centripetal force, leading to an expression for the angular velocity. Equation 1.5 shows the comparison and equation 1.6 shows the result.

$$\frac{GMm}{r^2} = m\frac{v^2}{r} \tag{1.5}$$

$$v = \left(\frac{GM}{r}\right)^{1/2}, \tag{1.6}$$

where the left hand side of 1.5 is the gravitational force, the right hand side is the centripetal force, G is the gravitational constant, M is the sum of the mass encompassed by the boundary of the stars orbit, m is the mass of the star, r is the orbital radius and v is the angular velocity. In systems such as the Solar System, where most of the mass is concentrated at the centre of the orbital radius, the velocity is seen to follow the inverse proportionality to $r^{1/2}$, resulting in slower velocities at larger radii. However, as can be seen in the example given in figure 1-5, the observations of stars in galaxies show the velocity to be approximately flat out to large distances. The conclusion from this is that there is an additional contribution to the inter-orbital mass[3]. In Figure 1-5, the markers are the measured values, with error bars, and the dashed lines show the contributory components.

[3]There are theories that attempt to explain this without the contribution of additional matter, known as Modified Newtonian Dynamics theories, whereby gravity is modified at the scales relevant to the problem. However, these theories struggle to explain the mass distributions seen in galactic collisions, as shown in section 1.1.3.

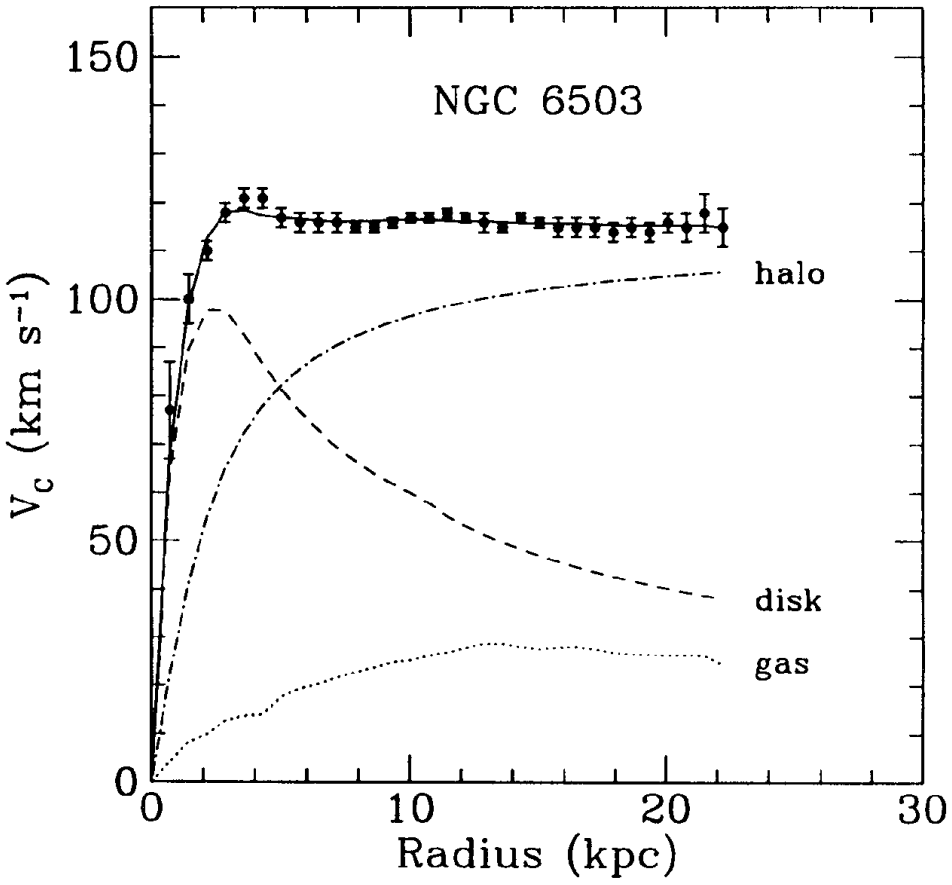

Figure 1-5: An example rotational velocity curve showing contributing elements [12]. The dark matter halo component is that required to make up the difference between the observed values and the amount attributable to luminous matter.

1.1.5 The Cosmic Microwave Background and Acoustic Oscillations

The cosmic microwave background (CMB) is an imprint of the temperature and density fluctuations of the Universe at the point of photon-matter decoupling, around 380,000 years after the big bang. These fluctuations are caused by variations in the mass-energy density of the early Universe, which result in regions of varying gravitational potential. This means that, at the point of decoupling, if a region compresses, a photon can be red-shifted. This results from the photon leaving an area of greater density than that which it entered.

Another effect that arises from the density variations of the early universe is the phenomenon of acoustic oscillations. The density variations cause pressure waves that have a propagation dependence on the speed of sound. Since the compression and rarefaction is dependent upon the gravitational attraction, which is in turn dependent on the amount of mass, these variations can be used to understand the abundance of dark matter.

By plotting the average correlation of variations in the CMB, as a function of angular separation in the sky, the power spectrum of the acoustic oscillations can be studied. The equation describing the spectrum is given by equation 1.7 (referenced from [3]).

$$\begin{aligned} C(\theta) &= \left\langle \left(\frac{\delta T(\bar{n})}{T}\right)\left(\frac{\delta T(\bar{m})}{T}\right)\right\rangle \\ &= (1/4\pi)\Sigma(2l+1)C_l P_l(cos\theta) \end{aligned}$$

where T is the average temperature over the whole sky, δT is the difference from the average at points m and n, θ is the angle separating the two points and P_l is the l^{th} Legendre polynomial with coefficients C_l. The average is taken over all points

separated by θ.

The most recent measurement of the CMB and the associated power spectrum comes from the Planck experiment and are shown in figures 1-6 and 1-7, both are referenced from [13]. The peaks in the power spectrum are called the acoustic peaks and represent different modes of the spectrum. This plot can be used to deduce various cosmological parameters by varying the input parameters to a theoretical model and comparing. By doing this for the Λ-CDM model, a value for the baryonic density Ω_b of the Universe is determined to be 0.049 [5]. This value indicates that only 5% of the mass-energy density is composed of the ordinary kind of matter with which we are familiar and it will be useful when considering the nature of dark matter in section 1.1.8.

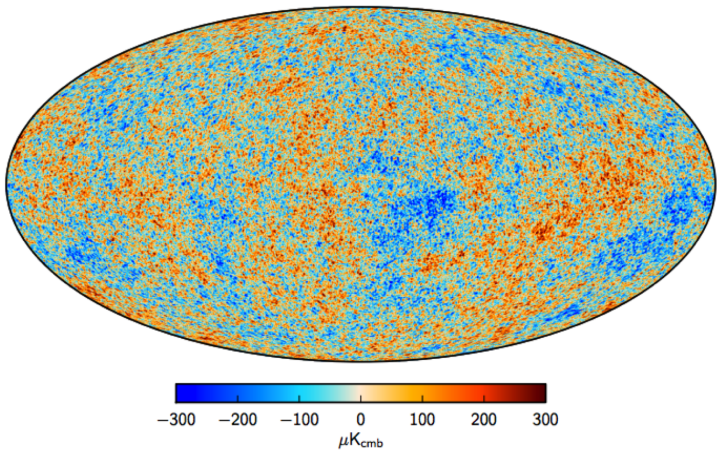

Figure 1-6: Temperature variations in the CMB across the sky, units are micro-Kelvin.

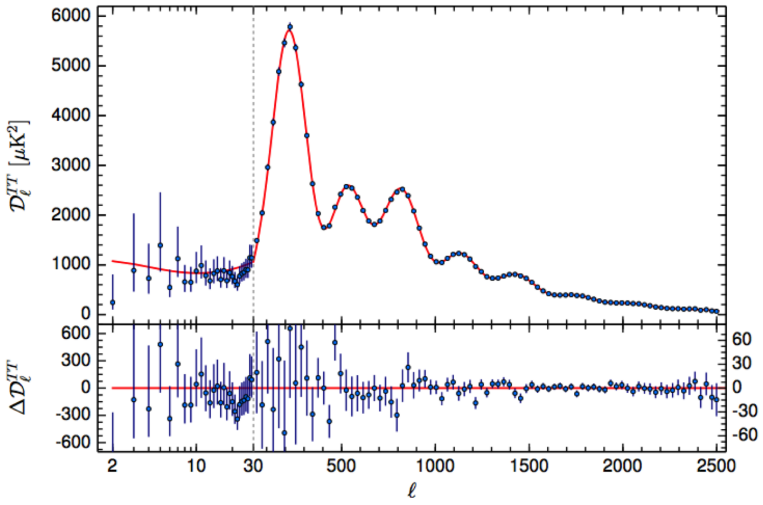

Figure 1-7: The angular power spectrum for the temperature, as described by equation 1.7. The abscissa is l from equation 1.7, representing the Legendre polynomial coefficient and the ordinate is the correlation value, normalised by $l(l+1)/(2\pi)$. The top panel is the power spectrum, showing data points with $\pm$ 1 σ error bars and the red line is a fit to the data using the Λ-CDM model. The lower panel shows the residuals of the data to the fit. The peaks are described in the main text.

1.1.6 Large Scale Structure and Simulations

The early Universe, as described by the Λ-CDM model, is homogeneous and isotropic. This would seem to be in conflict with the large scale structure and vast regions of seemingly empty space that comprise the Universe today. This apparent contradiction can be well resolved by the CDM component of the model. Since dark matter interacts only via the gravitational force, it can cause matter to accumulate and produce regions of high and low density. In order to test this theory, a simulation can be produced to model the evolution of the Universe with varying parameters related to the model. A number of simulations have been produced that do this, though possibly the most comprehensive of these is the DEUS Consortium which produced the first ever full Universe simulation in this way [14].

These simulations are useful as they allow us to test very complex theories which would be difficult, if not impossible, to solve analytically. Comparing the results with observation provides a means of evaluating the validity of the theory. Through these simulations it has been deduced that the dominant contribution of dark mattter must be cold (non-relativistic) at the point of radiation-matter decoupling. If this is not the case, then the large scale structure does not accumulate and galaxy formation would not have occured. A comparison of simulations using cold, warm and hot dark matter can be seen in figure 1-8.

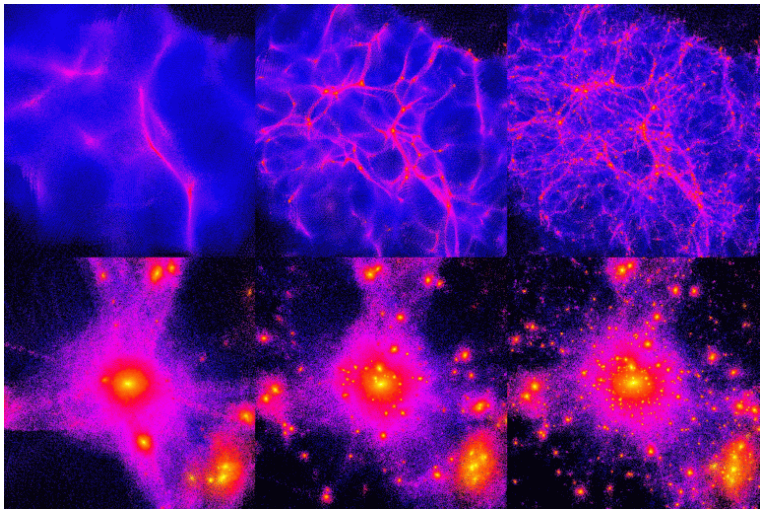

Figure 1-8: A comparison of the output from a simulation using cold, warm and hot dark matter[15]. Going from top left to bottom right, shows the decreasing temperature resulting in increased clustering.

1.1.7 The Origin of Dark Matter

With evidence for the existence of dark matter now examined, this section describes the process of dark matter creation in the early Universe. The model used for this is Big Bang Nucleosynthesis (BBN). BBN is premised on the concept that the statistical equilibrium of particle annihilation and creation breaks down, a process called freeze-out. Equation 1.7 is the Boltzmann Equation and it describes the process.

$$\frac{dn_\chi}{dt} + 3Hn_\chi = - < \sigma_a v > [(n_\chi)^2 - (n_\chi^{eq})^2], \tag{1.7}$$

where n_χ is the number density of dark matter particles (χ), t is time, H is the Hubble rate of expansion, defined as $\dot{a}/a$, σ_a is the annihilation rate, v the particle velocity (these values appear in brackets, representing the thermal average) and n_χ^{eq} is the number density in thermal equilibrium. The left hand side of equation 1.7 describes the time evolution of the number density, incorporating the expansion of the Universe, and the right hand side contains the physical processes. The first term in the square brackets is for the annihilation of particles and the second term is for creation. This

equation shows that as the cross-section for the annihilation interaction increases, the abundance of particles at the point of freeze-out, the relic abundance, decreases. Equation 1.8 shows the fractional value of the present day density, relative to the critical density, as derived in reference [16].

$$\Omega_\chi h^2 \sim 3 \times 10^{-27} cm^{-3} s^{-1} < \sigma_a v >^{-1}, \tag{1.8}$$

where $\Omega_\chi = \rho_\chi / \rho_c$, with ρ_c the critical density[4] and h represents the uncertainty in the Hubble expansion rate and can take a value between 0.5 and 1.

1.1.8 Matter Type

The dominant contribution to dark matter in the universe is most likely non-baryonic in nature. This arises when we consider the total baryonic contribution to the matter in the universe, the baryonic density. As in equation 1.8, this value is usually quoted as a fractional amount of the critical density and has the currently favoured value of $\Omega_m = 0.049$, deduced using the Λ-CDM model and the Planck measurements of the CMB [17]. Next, the amount of visible baryonic matter is calculated by studying a region of space and making use of the mass to luminosity ratio, which tells us how much mass there is based on the amount of light being produced. It is reasonable to do this as observations indicate that the Universe is homogeneous and isotropic. From this we get a value of $\Omega_{vis} \approx 0.01$. At first it would seem that the dark matter could be baryonic and make up the unseen amount. However, the galactic rotational curves mentioned above, imply that there is about 10 times more dark matter than ordinary, which would bring the combined value to $\Omega_{vis} + \Omega_{dark} \approx \Omega_{vis}(1+10) = 0.11$. This is greater than the expected value of 0.049 and as such, dark matter is considered to be largely non-baryonic.

[4]That required for a flat Universe

1.1.9 Dark Matter Candidates

There are many possible candidates for the missing mass. Baryonic candidates include Massive Compact Halo Objects (MACHOs) e.g. neutron stars, brown dwarfs and black holes. Non-baryonic candidates are particulate in nature and are the focus of this section. It is possible that these candidates are composite in nature or there may be a number of particles comprising a dark-sector, however, for simplicity, a single particle hypothesis is frequently used and is the current purvue for DMTPC. There exists a plentiful array of proposed particles that have the potential to play the role of dark matter in our Universe. Neutrinos are an example that have one of the required characteristics, that they interact only very weakly with ordinary matter: this must be the case otherwise the interaction would already have been observed. And since the verification that neutrino oscillations do occur [18] and thus have mass, they would seem to be like the particle being sought. However, as discussed above, the dark matter required for galaxy formation must be cold and this is not the case for cosmic neutrinos. As such they do not meet the criteria. The theories that describe candidate dark matter particles are numerous and complex so only a brief outline is given below.

1.1.9.1 Axions

These were originally proposed to solve the strong CP-violation problem [19]. The nature of this problems lies in the lack of observation of charge-parity violation in the strong sector. Quantum Chromodynamics (QCD), the theory which describes the strong sector, has no explicit reason for charge-parity to be inviolate. As such, this is considered a fine tuning problem in particle physics, since a phenomenon of the theory is not observed in experiment.

The axion, which is as yet unobserved, comes from the breaking of a symmetry in the QCD lagrangian. Much like the breaking of the electro-weak symmetry pro-

duces the Higgs boson, the breaking of this Peccei-Quinn symmetry (named after its founders) also produces a particle: the axion. Current constraints on the axion properties allow it to be a potential candidate for the dark matter dominating the Universe.

1.1.9.2 Kaluza-Klein Particles

Other candidates which have some popularity in the physics community are Kaluza-Klein particles. The theory behind these particles was first proposed in the early part of the nineteenth century [20] and unites gravity and electromagnetism by the inclusion of a fifth dimension (or more, as in the later extra-dimensional String Theory). Unlike Supersymmetry, Kaluza-Klein theory does not introduce additional particles but instead allows the existing Standard Model particles to propagate into the extra dimension(s). These extra dimensions are compactified, which results in the quantisation of a particles momentum. This quantisation then produces a set of modes, known as Kaluza-Klein states. A stable dark matter candidate can be produced from these states by imposing translational invariance along the extra dimension. When this is done, a discrete symmetry, Kaluza-Klein parity, produces a particle called the Lightest Kaluza-Klein Particle (LKP) which is the dark matter candidate.

1.1.9.3 WIMPs and SUSY

Weakly Interacting Massive Particles (WIMPs) is a generic term to describe particles that only interact very weakly and are a relic from the hot plasma of the early Universe. This type of particle is so-grouped due to the seemingly miraculous result that when the strength of the weak force is used for the annihilation cross section of equation 1.8, then the value that comes out is very close to that required for dark matter. A good candidate for such a particle is the neutralino which is a particle that comes from Supersymmetry (SUSY).

SUSY is a proposed solution to the Hierarchy Problem in particle physics, which is related to renormalisation: wherein quantum loop corrections in mass and coupling calculations are considered for the effective value. When this is done it would appear that finely tuned cancellations are required to match experimental results. SUSY resolves this problem by proposing that there is a symmetry between fermions and bosons, such that for every boson, a supersymmetric particle exists with the same properties except that it is a fermion and vice-versa. When this is done, the loop contributions cancel each other naturally which makes SUSY a very elegant solution, it is therefore widely popular.

SUSY contains many variables, as such modifications have been made to make them more manageable. One of these modifications is the Minimal Supersymmetric model, another is the Constrained Minimal Supersymmetric Model. Common SUSY models is an additional particle property dubbed R-parity. It is only by conservation of R-parity that a Supersymmetric particle can survive, otherwise at the end of a decay chain, only standard model particles would be left. The neutralino is one of the Lightest Stable Particles (LSPs) at the end of such a decay chain. Other possibilities include the Higgsino and the Bino.

1.1.10 Local Dark Matter

The model describing the shape and structure of the dark matter that surrounds our galaxy, that is predominately used by the direct detection community for comparison of results, is called the Standard Halo Model (SHM). This model posits a large sphere surrounding a galaxy, with a Maxwell-Boltzmann velocity distribution (eqn. 1.9) [21]. This distribution is used on the assumption that the dark matter behaves like an ideal gas and is collisionless.

$$f(\vec{v}, \vec{v_E}) = \mathrm{e}^{\frac{-(\vec{v}+\vec{v_E})^2}{v_0^2}} \tag{1.9}$$

The variables and values that are typically used for the SHM are listed in table 1.2.

Using these parameters and the observed rate of events seen in a detector, a limit can be set on the interaction[5] cross-section as a function of mass, using the prescription outlined in section 1.2. There are other halo models but this is the prevailing one within the dark matter detection community. The rationale for this is that the consistency of model used between experiments allows for easier comparison of results.

Table 1.2: SHM values used as standard by direct detection experiments. ρ_D is the local dark matter density, v_E is the Earths velocity, v_{esc} is the galactic escape velocity and v_0 is the velocity dispersion of the halo.

Parameter	Value
ρ_D	0.3 GeV cm^{-3}
v_E	244 km s^{-1}
v_{esc}	544 km s^{-1}
v_0	230 km s^{-1}

[5]Between dark matter and detector target.

1.1.11 Current Status of Dark Matter Experiments

There are many dark matter experiments that employ different detection techniques and are searching for different signatures - see section 1.2. Figure 1-9 shows the latest results for the spin-dependent cross-section limit. The upper limit on the cross-section is plotted as a function of dark matter mass. The lowest bounds on the spin-dependent WIMP-proton cross section were set by PICASSO for masses 4 GeV/c^2 to 8 GeV/c^2 [22], by SIMPLE for masses 8 GeV/c^2 to 40 GeV/c^2 [23] and by COUPP for masses above 40 GeV/c^2 [24]. PICASSO and COUPP then combined to form PICO, who published their first result in 2015, for a 2L volume detector [25] and have published new limits for 2016 [26]. The improvement in sensitivity can be clearly seen in figure 1-9. Figure 1-10 shows the plot for the spin-independent case, for this the LUX collaboration have produced the most stringent limit to date.

In both plots there is shown a region of claimed discovery, made by the DAMA/LIBRA collaboration [27]. This experiment exploits the modulation of the expected dark matter signal (see section1.2.2). The DAMA/LIBRA result is a contraversial one since subsequent experiments with greater sensitivity have excluded this region.

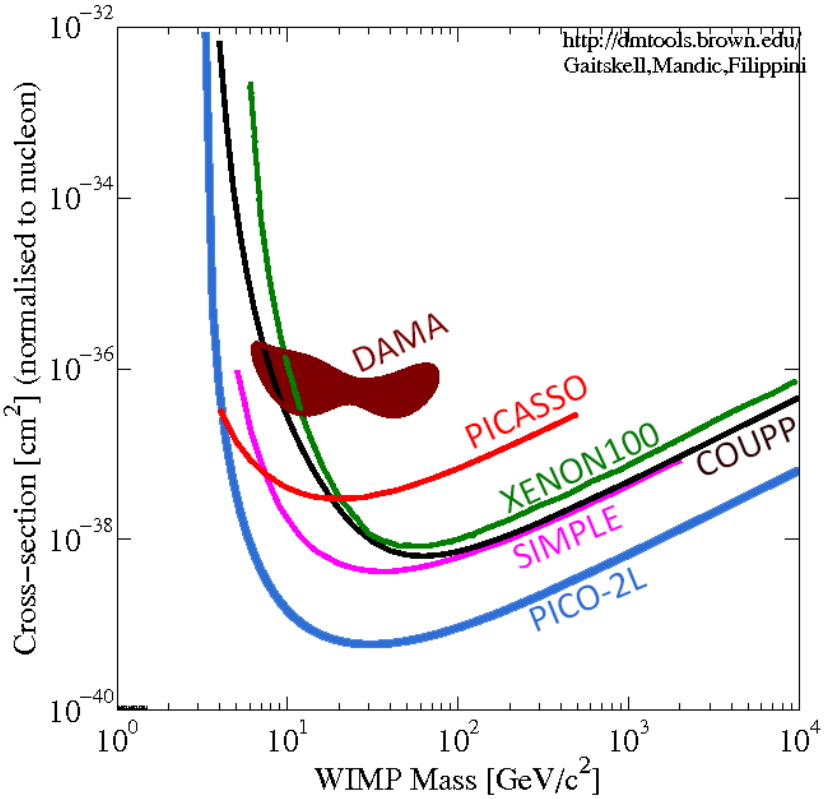

Figure 1-9: Plot showing the exclusion limit cross-section as a function of mass for the spin-dependent interaction between a WIMP particle and a proton. The experiments shown are PICASSO [22], XENON100 [28], SIMPLE [29], COUPP [30], PICO-2L [26]. Also shown is the claimed discovery region made by DAMA [27].

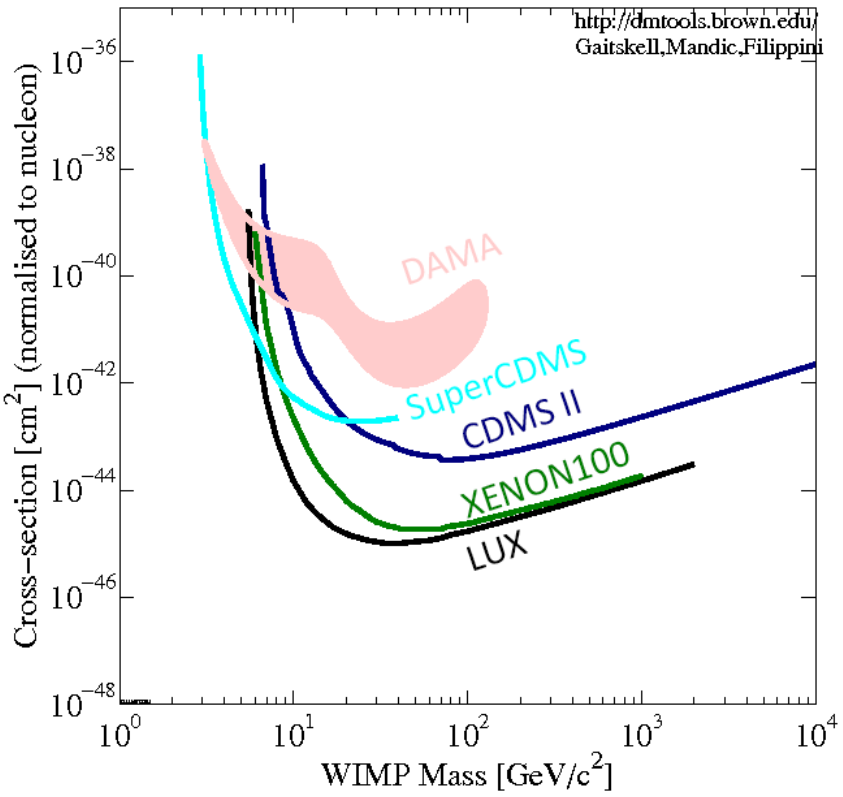

Figure 1-10: Plot showing the exlusion limit cross-section as a function of mass for the spin-independent interaction between a WIMP particle and a nucleon. The experiments shown are SuperCDMS [31], CDMS II [32], XENON100 [28] and LUX [33]. Also shown is the claimed discovery region made by DAMA [27].

1.2 Dark Matter Detection

There are three channels considered for detecting a dark matter signature, each of which is shown in figure 1-11. DMTPC is a directional, direct detection experiment. Direct detection experiments, aim to detect the recoil of a dark matter particle off the nucleus of a target material. This is considered a direct detection since the interaction would have occured within the detector in question and is a detection of currently extant dark matter. Indirect detection exploits the annihilation process of a dark matter with its anti-particle, with the aim of detecting the standard model particles produced in such an interaction. For example, PAMELA [34] are searching for an excess of anti-protons from the direction of the Sun[6], see [34] for more on this technique. The production method, as the name implies, is the method of generating dark matter particles via the collision of standard model particles at high energy. Since the dark matter interaction strength is weak and it does not interact electromagnetically, its presence must be inferred by missing energy from the summation of energy in the calorimeter detectors. So far there have been no observation of any supersymmetric particles[7]. See, for example, [35] for an up to date discussion on this area. A consideration with dark matter production discovery is that there is no guarantee that the particle detected is the same as those which make up the dark matter haloes surrounding galaxies. Finally, directional detectors aim to reconstruct the direction of the recoiling nucleus resulting from a direct interaction, within a detector volume. The reasons for this are described in this section, alongside the interaction process and the technology employed for detection.

[6]The sun is a preferred direction due to the high mass and thus strong gravitational pull, causing dark matter particles to coalesce.

[7]The current favoured dark matter candidate, see section 1.1.9.3.

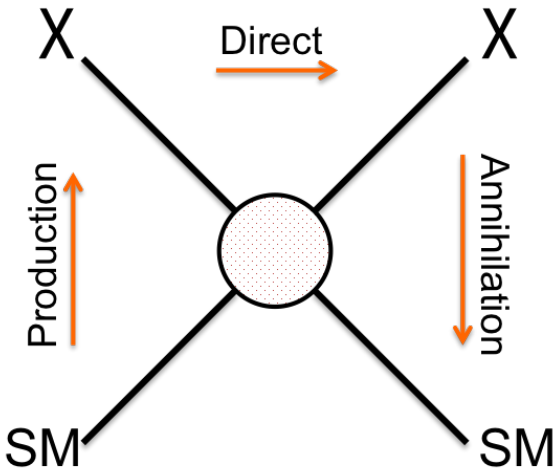

Figure 1-11: A diagram representing the methods that can be employed to detect a dark matter interaction. Going upwards is production, right to left is direct and downwards is indirect detection. Each of these is described in the main text. The circle in the middle is to indicate that the interaction mechanism is as yet unknown.

1.2.1 Direct Detection

The aim of a direct detection experiment is to observe the interaction between a dark matter particle and the nucleus of a medium. The low velocities of dark matter result in an elastic scatter off the target nucleus and a transfer of momentum causes the nucleus to recoil. An experiment can be designed to detect the recoil based on the effects it will have on the surrounding medium. At the low energies involved, only the scalar (spin-independent, or SI) and axial-vector (spin-dependent, or SD) components of the interaction are considered. The two types are compared in equations 1.10 (SI) and 1.11 (SD). In both cases it is typical to make a simplification, whereby the couplings of the dark matter particle to the proton and neutron are considered equal. Once this is done the SI interaction reduces to have an A^2 dependence, where A is the atomic mass number. This makes the spin-independent cross-section larger for higher mass targets. The spin-dependent cross-section does not simplify in this way, owing to the dependence on the net spin contributions to the nucleus for the proton and neutron ($\langle S_p \rangle$ and $\langle S_n \rangle$). A spin-dependent measurement is best suited to targets

with individually unpaired neutrons or protons.

$$\sigma_{SI}^{N} \propto \frac{\mu^2}{\pi}|Zf_p + (A-Z)f_n|^2 \tag{1.10}$$

$$\sigma_{SD}^{N} \propto \frac{4\mu^2}{\pi}|Za_p \langle S_p \rangle + (A-Z)a_n \langle S_n \rangle|^2 \frac{J+1}{J} \tag{1.11}$$

In equation 1.10 f_p and f_n are the spin-independent couplings to the proton and neutron respectively. In equation 1.11 a_p and a_n are the spin-dependent couplings to the proton and neutron. In both, A is the atomic mass number, Z is the atomic number, $\langle S_p \rangle$ and $\langle S_n \rangle$ are the spin expectation values for the proton and neutron, μ is the reduced mass for the dark matter particle and the target and J is the total angular momentum.

Different techniques and targets are employed to detect these interactions. Some detectors use liquid as the target material, for example DEAP [36], where the scintillation light produced from an interaction is detected, and PICASSO [22] where the acoustics from bubble nucleation in a super-heated liquid are recorded. Solid targets are also used, as in the case of SuperCDMS [31], where scintillation light and phonons are detected. And others make use of gas, such as DRIFT [37] and DMTPC. One of the benefits of a gaseous target is the potential for reconstructing the direction of the recoil, as is shown in section 1.2.1.

Due to the weakness of the interaction, only very few events are expected to take place for a given mass and target material. Based on the expected rate, equation 1.15, found in section 1.3, figure 1-12 shows the expected rate as a function of recoil energy for a 100 GeV/c^2 WIMP mass and a fluorine target. Figure 1-12 highlights the increase in the expected number of events at low energy. This is a driving force behind building detectors which are capable of detecting low energy events.

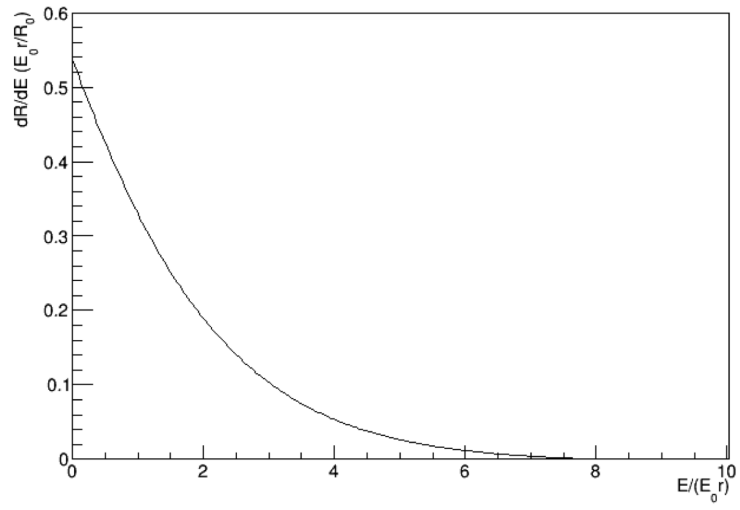

Figure 1-12: The expected rate of events for a fluorine target and WIMP mass of 100 GeV/c^2, using SHM parameters of table 1.2 and assuming an interaction cross-section of one pico-barn. The abscissa is the normalised recoil energy of the nucleus and the ordinate is the normalised count rate. In the normalisation values, E_0 is the most probable incident kinetic energy of the WIMP, r is a kinematic factor and R_0 is the total rate. See section 1.3 for more on this topic.

1.2.2 Directional Detection

Owing to the motion of the Solar System through the galactic plane and the orbit of Earth around the Sun, there is an apparent WIMP wind experienced on Earth, see figure 1-13. The effect of this is to give a preferred direction from which WIMP events are expected to originate. With the interaction being elastic, the recoiling nucleus will preserve the direction of the particle that has interacted with it. Other factors, such as energy, come into play as to whether the direction can be extracted from the detector signals. Section 1.2.4 details the interaction process within the detector further. These factors combined can be exploited as a strong indication that a reconstructed nuclear recoil is from a dark matter interaction [38].

A modulation to the mean direction of the wind (as seen in the detector) is expected to be seen: daily for the Earths rotation about its axis and annually for Earth's orbit around the Sun, see figure 1-14. The signal is strongest for the diurnal modulation and the detection of such a signal would be very convincing as being due

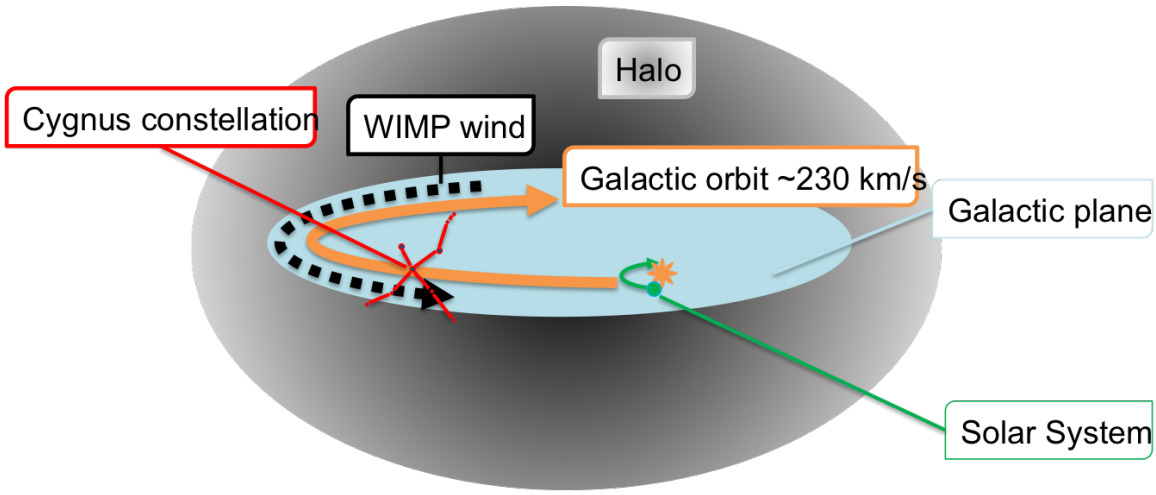

Figure 1-13: Diagram showing the motion of the Solar System through the dark matter halo.

to dark matter interactions.

Perhaps the most significant reason for pursuing a directional dark matter signal is what is referred to as the neutrino floor [39, 40]. With direct detection experiments achieving greater and greater sensitivity and, as yet, finding no irrefutable signal, the coherent scattering cross-section of neutrinos off nuclei will soon be detectable. This represents an irreducible background in standard direct detection experiments. However, the dominant source of neutrinos to this background originate from the sun. This, combined with the expected WIMP wind, the use of directionality as an additional observable provides a means of discriminating between the two event types.

Another benefit to reconstructing the direction is the increased ability to discriminate a dark matter signal from an isotropic background. It has been shown that reconstruction of the sense of a two-dimensional track provides two orders of magnitude improvement over the number of events required to confirm anisotropy [41]. This is compared with only a single order of magnitude improvement by going from two-dimensional axial to three-dimensional axial. Figure 1-15 shows this pictorially.

To be able to reconstruct the direction of the recoiling nucleus, the path of the nucleus must be discernible by the detector being used. This is the reason behind the

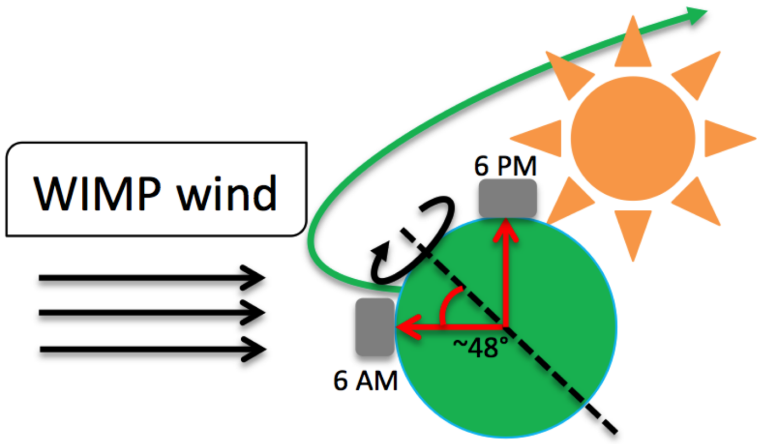

Figure 1-14: Diagram showing the Earth's rotation and the effect on the directional component of the WIMP wind. 48° is approximately the location that would maximise this effect.

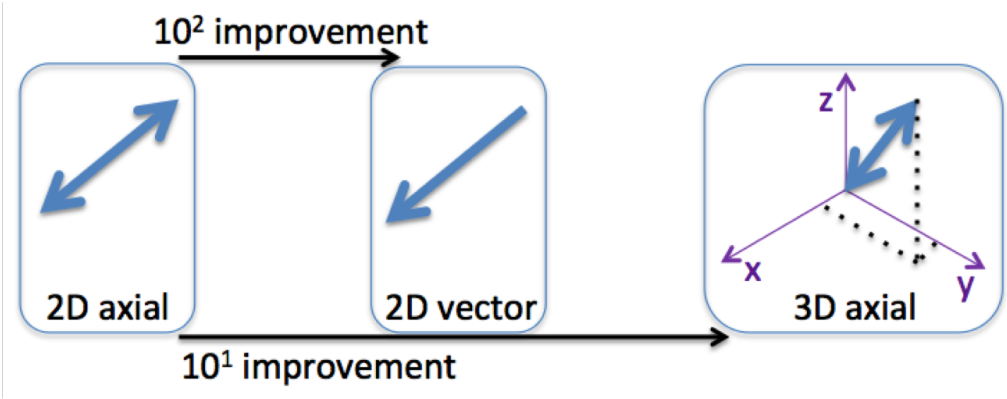

Figure 1-15: A diagram to show the improvement in the ability to confirm an anisotropic signal. Going from two-dimensional axial to vector gains two orders of magnitude improvement compared with only one order of magnitude improvement going from two-dimensional axial to two-dimensional vector.

use of gaseous detectors and is explored further in section 1.2.5. In order to obtain the sense of the track, the stopping power as a function of energy is employed. Figure 1-16 gives a diagrammatic representation of this to show how, at the low energies of the nuclear recoils under consideration, which can range from of a few eV to 100s of keV, the energy deposition is largerst at the start of the track. An example track with this energy deposition profile is given in section 2.1.2, figure 2-4.

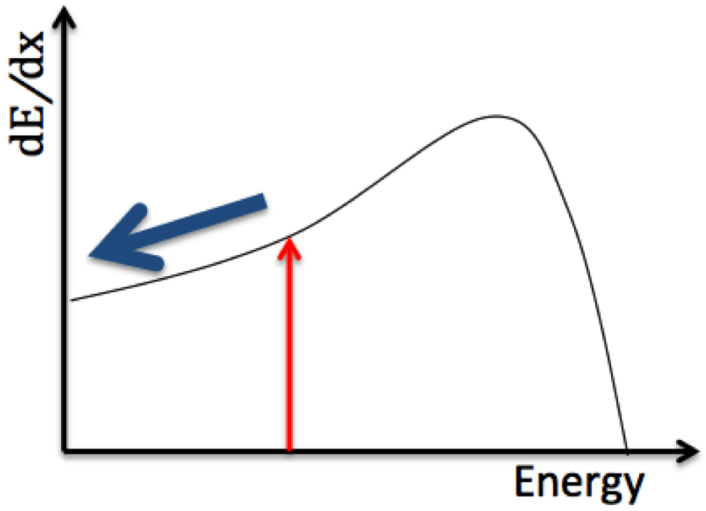

Figure 1-16: A representation a Bragg curve, which plots the stopping power as a function of energy. The red arrow points to the energy region of interest. The black arrow shows how the drop in stopping power, as the track loses energy, can be used to reconstruct the track sense.

1.2.3 The Time Projection Chamber

The time projection chamber (TPC) configuration used by DMTPC consists of two volumes, each of which is placed under an electric field. A diagrammatic representation is shown in figure 1-17. The first, larger, volume is called the drift region and the second, smaller, volume is called the amplification region. The drift region is so-called as primary electrons from an ionised track that occurs in this volume are transported, or 'drifted', towards the amplification region, due to the presence of the electric field. The amplification region is made up of two electrodes closely spaced. The electric field in the amplification region is much stronger than the drift region, resulting in multiplication of the initial electrons, as described in section 1.2.5. This amplification of the primary electrons results in a signal that is detectable.

The uniformity of the electric field is an important factor as it can cause deflection of an electron cloud[8] or cause the cloud to not make it to the amplification region at

[8]Generated by ionisation of the gas, see section 1.2.4

all. In the ideal case the field would be perfectly uniform throughout the drift-region volume. Explicitly this means the field magnitude would be the same at all points within the drift region and would have no off-axis components. In the case of the diagram, this would require the field to point fully in the z-direction and to have no radial component at any position. To better understand the uniformity it is useful to think of the electrode caps of the drift region as a capacitor. In the absence of the field rings, collectively called the field cage, the field will have strong edge effects, as seen in figure 1-18.

To achieve field uniformity, the potential difference between the two plates can be gradated through the use of a sequence of conductors and resistors. The finer the resolution of gradation, the more uniform a field will be created, however too fine a gradation results in a large amount of material surrounding the drift region. Each ring of the field cage is connected to the next by a resistor, providing a smooth voltage drop across the drift region, which acts to maintain a level of field uniformity.

As will be discussed further in chapter 3, the materials of the detector are considered to be one of the largest sources of background events, and so there becomes a trade off between electric field uniformity and background reduction. This trade off was studied as a topic of this thesis, with the aim of maintaining a level of electric field uniformity that would provide the desired fiducial volume, as determined by track survival from the top of the drift region to the amplification region, whilst keeping material usage to a minimum, see chapter 4. The reduction in materials goes alongside using materials that are low in backround.

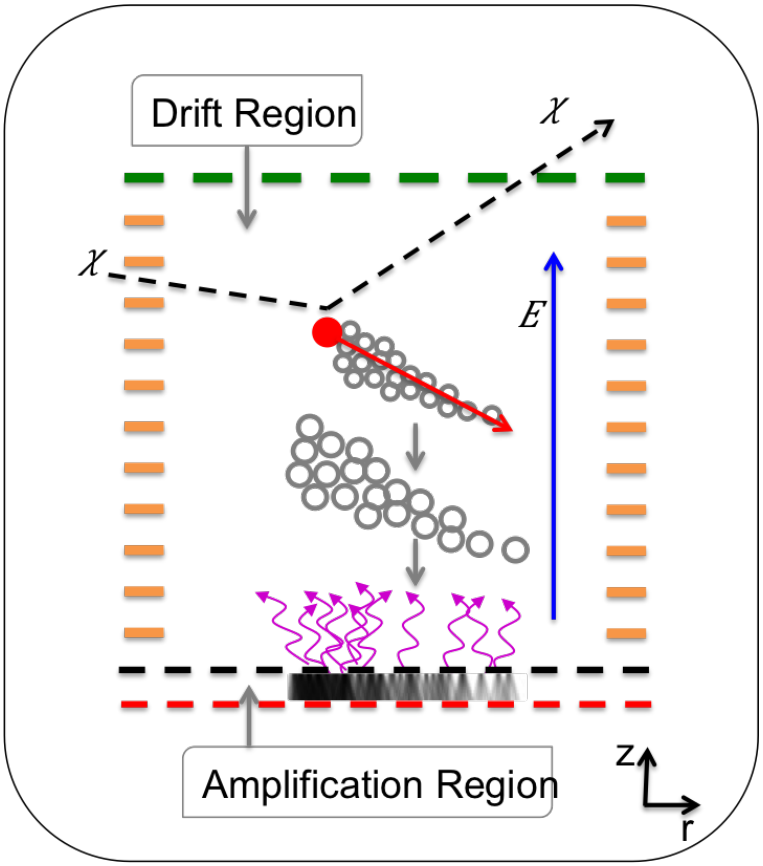

Figure 1-17: A diagram of a typical TPC arrangement used in the DMTPC detectors. The dashed lines represent the TPC electrodes: green is the cathode, black is ground and red is an anode. The orange bars represent the field-shaping rings. The blue arrow shows the direction of the electric field. An example event is also shown, as is discussed in section 1.2.4. For this, χ is an incident dark matter particle and the dashed black line traces out its path through the detector. The red dot represents a nucleus of the target, with the red arrow tracing out its recoil from the interaction. The grey circles represent the electron cloud generated from the ionisation induced by the nuclear recoil. Also shown is the diffusion of this cloud as the electrons are drifted to the amplification region. The black shading in the amplification region represents the avalanche of electrons produced during amplification. And the purple wavy arrows represent the scintillation photons produced.

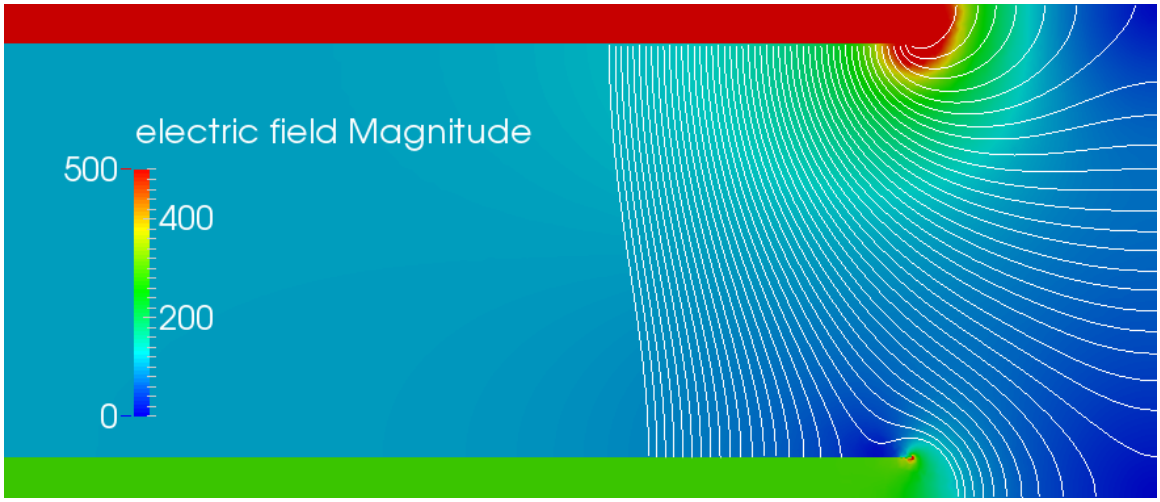

Figure 1-18: The field and field lines of a capacitor, demonstrating edge effects. The white field lines trace out the path of a particle originating from the uppermost electrode, showing the distortion and loss of particles from the drift region volume.

1.2.4 The Interaction Process

Assuming our theories are correct and that dark matter behaves as expected, the generation of a signal will occur in the following manner. As shown in figure 1-17, a dark matter particle enters the detector, undergoes elastic scattering with a nucleus of the target. The dark matter transfers some of its energy to the nucleus, causing it to recoil. As the nucleus recoils it can lose its energy through heat, scintillation and ionisation. The DMTPC detectors are currently only sensitive to the ionisation channel therefore this fractional energy loss must be accounted for. Equation 1.12 is based on a power law described by Lindhard [42] and gives a quenching factor used to describe the fraction of energy that goes into ionisation. To distinguish between the energy of the recoil and the ionisation energy, the notation keVnr and keVee are used, where nr denotes nuclear recoil and ee denotes electron-equivalent.

$$\begin{aligned} E_I &= E_R/0.76879 \quad (E_R > 200keV) \\ &= a + bE_R - cE_R^2 + dE^3 - eE^4 \quad (E_R < 200keV) \end{aligned} \tag{1.12}$$

where the subscripts I and R stand for ionisation and recoil respectively and the values for CF_4 are: $a = 4.46388$ keV, $b = 1.65231$ keV^{-1}, $c = 2.86719 \times 10^{-3}$ keV^{-2}, $d = 5.8344 \times 10^{-6}$ keV^{-3} and $e = 4.22845 \times 10^{-9}$ keV^{-4}.

The process of ionisation incorporates complicated microphysics which is simplified by a parameterisation: the work function W. This value gives the mean energy needed to free an electron from the target atom or molecule and is used to estimate the number of electrons liberated for a particular ionisation energy. For CF_4 the value of the work function has been measured by the DMTPC collaboration to be 33.8 $\pm$ 0.4 eV [43].

The ionised electrons can undergo attachment to a molecule or ion. This loss results in a reduction in the detectable signal. CF_4 has low electron affinity, however impurities that are present can have high affinity. Oxygen, for example, has a high electron affinity and so it is imperative that the detector contains as little as possible. This is achieved by ensuring there are no leaks in the detector and minimising the materials used, from which oxygen can outgas.

The cloud of electrons produced by the ionisation undergoes a process called diffusion, whereby the spatial extent of the electron cloud will increase. The amount of diffusion is dependent on the drift length the track travels, the strength of the electric field and the nature of the gas itself and occurs in both the drift region and the amplification region, along the transverse and longitudinal directions. Only the two-dimensional projection of a track is read out by DMTPC (see section 2.1.2) and so the transverse component is the parameter of interest. Any further reference to diffusion will be considering the transverse component only. The diffusion is modelled as a Gaussion with standard deviation shown in equation 1.13

$$\sigma^2 = \sigma_0^2 + 2\frac{D}{\mu}\frac{zL}{V}, \tag{1.13}$$

where σ_0 is a constant, which is governed by the detector resolution. D is the diffusion

constant and is gas-dependent, μ is the mobility of electrons, z is the drift length, L is the distance from cathode to ground and V is the cathode potential.

Figure 1-19 and 1-20 show the ratio of the diffusion to the electron mobility and the electron drift velocity, respectively, both as a function of the reduced field[9]. The electron drift velocity is an important parameter as it relates to the time an electron takes to travel to the amplification region which affects diffusion and signal collection time. From figures 1-19 and 1-20 it can be seen that there is a sweet spot at $\sim$10 Td where the drift velocity peaks. Much higher than this and the transverse diffusion begins to considerably increase. The values used for the diffusion graph are combined from reference [44] (15 Td and above) and reference [45], a diffusion study performed by the DMTPC collaboration. Values for the drift velocity are all from reference [44].

Once the electrons arrive in the amplification region, they undergo acceleration due to the increased electric field. This acceleration causes the primary electrons to ionise the gas further, resulting in an avalanche effect and scintillation photons to be released. Avalanching is an exponential process and is described by equation 1.14.

$$n(x) = n(0)\mathrm{e}^{\alpha x} \tag{1.14}$$

where n is the number of electrons, $n(0)$ is the initial number of electrons, x is the distance travelled and α is the first Townsend coefficient. α describes the number of ion pairs generated and is dependent upon the electric field strength. The large number of electrons also means a large number of ions are produced in the amplification process. These ions de-excite and produce scintillation photons. The number of photons produced in CF_4 has been measured by the DMTPC collaboration to be 0.34 per electron in the avalanche [46].

[9]The reduced field is the electric field divided by the number density and is given in units of Td (Townsend units), 1 Td is equivalent to 1×10^{-17} V/cm. This is a useful parameter for comparison of parameters since the field and density can both vary in different experiments.

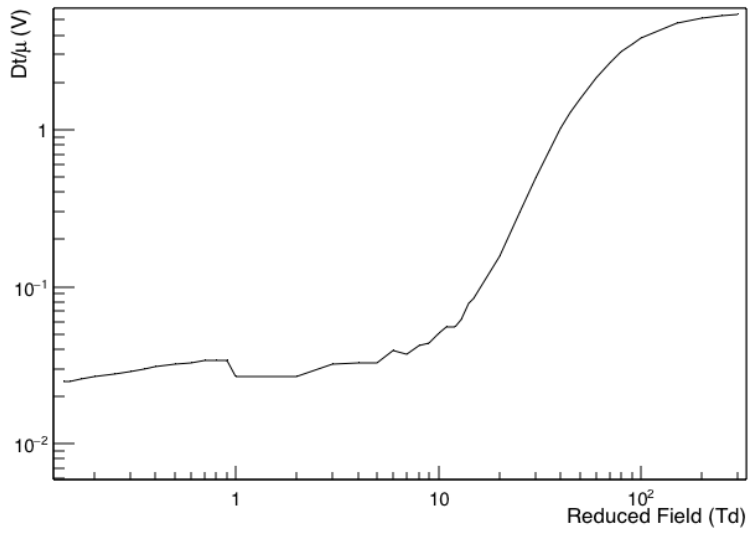

Figure 1-19: The ratio of transverse electron diffusion to electron mobility as a function of reduced field.

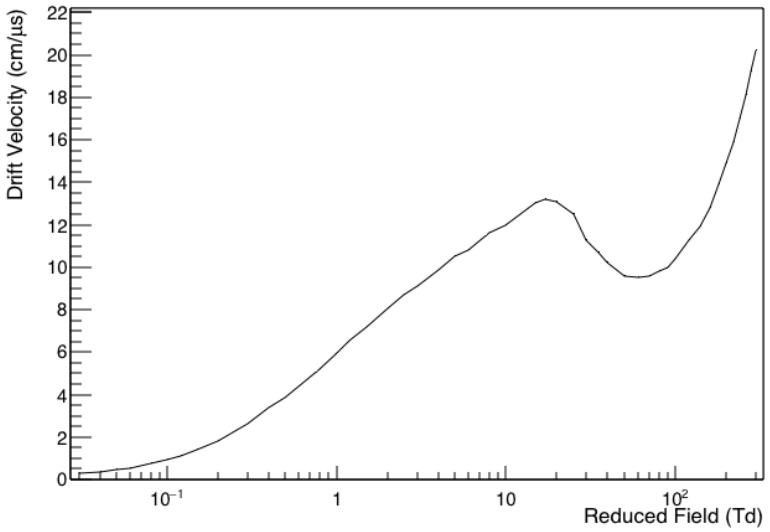

Figure 1-20: The electron drift velocity as a function of reduced field.

1.2.5 The Gas

DMTPC uses low pressure gas as the target material since the mean free path of a recoiling nucleus is much larger than that for a liquid or solid. This means that the recoiling nucleus will travel further before losing all of its energy, resulting in a longer path length. Increasing the track length improves the chances of correctly reconstructing the recoil direction.

CF_4 is the chosen gas for DMTPC as the fluorine atoms each have a single unpaired proton in their nucleus. This makes CF_4 a good candidate for detecting the WIMP-proton spin-dependent interaction. The spin-independent interaction can still occur with the carbon atoms of the molecule and the detector does not distinguish between the two types, however, due to the four-to-one ratio of fluorine to carbon, all interactions are considered to be spin-dependent. Other factors that make CF_4 a good choice are:

1. High electron drift velocity, $\mathcal{O}(10)$ cm/μs : fast signal collection
2. High scintillation at wavelengths that correspond with the peak quantum efficiency of CCD[10] chips, $\sim 70\%$ at ~ 625 nm : optimised light signal collection
3. Low electron diffusion, $\mathcal{O}(1)$ mm at 12 Td[11] : well defined tracks
4. Low to zero electron attachment : little to no loss of signal electrons

A drawback to the use of a low pressure gas is the small amount of mass. As shown in equation 1.16, the event rate is determined as a function of exposure, which is reported in kg-days.

[10] Charge-coupled device
[11] Td = Townsend Units, equivalent to 1×10^{-17} V cm.

1.2.6 Backgrounds

There are a number of backgrounds that can be seen in dark matter detectors, some of which are related to the detection method employed and others which are due to background radiation. The detector-specific backgrounds, for DMTPC, are reserved for discussion in section 2.2, alongside the detector descriptions. The background due to radiation is discussed in this section. Of prime concern to experiments such as DMTPC are the decay products of the uranium-238 and thorium-232 decay chains, shown in figures 1-21 and 1-22. All raw materials contain some abundance of these elements so it is important to understand their contribution to background radiation. Radon from the decay chain emanates from the materials as a gas and the progeny of the radon decays are born charged. Being charged, these progeny can plate out on the detector materials or attach to dust, which can also attach to the detector materials. The subsequent decays can produce nuclear recoils, which mimic those of a dark matter recoil. It is thus important to both reduce (ideally eliminate) radon in the detector and be able to distinguish between the two types of event.

Neutrons can also be produced within the detector through spontaneous fission and the (α-n) interaction. Since neutrons are electrically neutral and are therefore effectively dark matter themselves, they can cause nuclear recoils which also mimic the signal we are trying to detect.

One such means of rejecting radon related events, called radon-progeny recoils (RPRs), is fiducialisation. Where radon can plate out on materials and thus the components that make up a detector, they are most often likely to occur at the surface of the detector component. Fiducialisation is the method of selecting a region of the detector volume and only considering events that occur in that space. By fiducialising away from the component surfaces, the radon background can be reduced. The other solution is to use as little material as possible in the design of the detector and be rigorous in the cleaning of any materials which are used, to remove any particulate

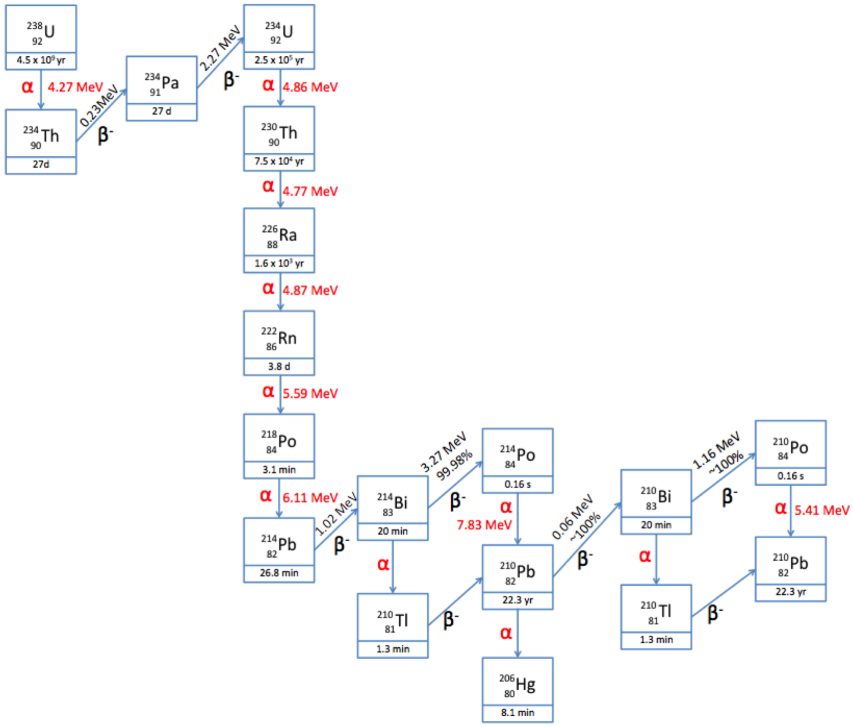

Figure 1-21: The uranium-238 decay chain. Alphas and their associated energy are shown in red, betas and their associated energies are in black.

matter which might have radon attached. Finally, the use of a pure target material, with as little radioactive contamination as possible, is another means of reducing the radon background.

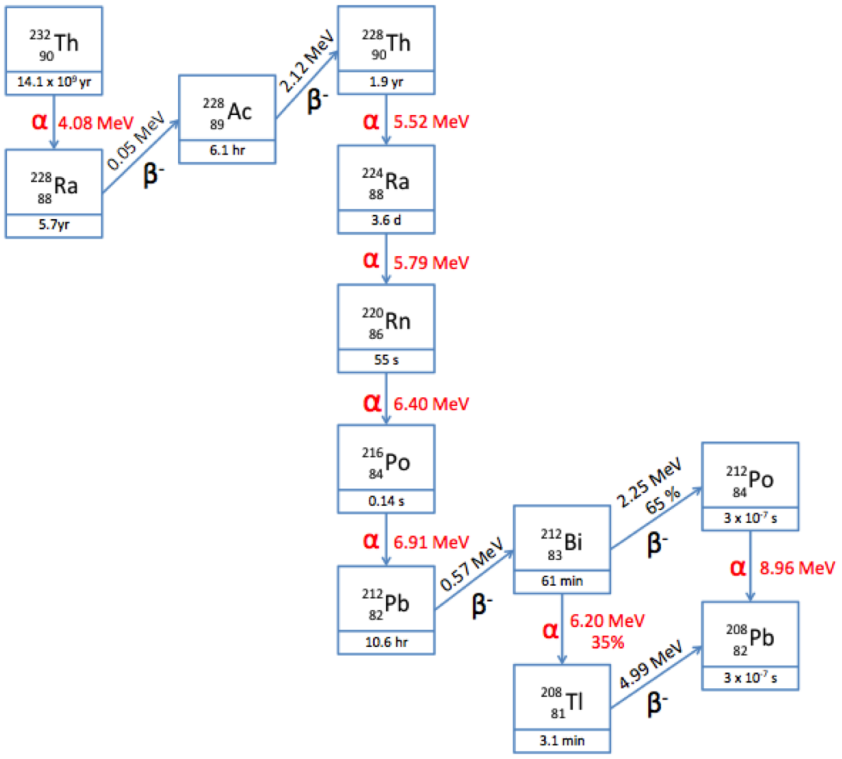

Figure 1-22: The thorium-232 decay chain. Alphas and their associated energy are shown in red, betas and their associated energies are in black.

1.3 Dark Matter Interaction Rates in Direct Detection Experiments

Direct detection experiments are aimed at detecting dark matter that is present within the galaxy. This is achieved by observing the after-effects of an interaction between a dark matter particle and a target material: CF_4 in the case of DMTPC. Section 1.2.1 discusses the experimental aspects of direct detection. This section outlines the theory that goes into setting a limit on the dark matter interaction cross section. It closely follows the work of Lewin and Smith [21] but focuses on the implementation used in the analysis presented in this thesis, considering spin-dependence and CF_4 as the target. Though DMTPC is a directional experiment, directional information was not used in the analysis. This is due to the directional reconstruction and analysis being at an early stage of development, see [47] for more on the current status.

When publishing results, dark matter experiments report a limit on the dark-matter-nucleon interaction cross-section as a function of dark matter mass. In order to deduce the values used for such a plot, such as the one seen in figure 1-9 in 1.1.11, the detectors predominantly operate as counting experiments, accumulating a rate of events which are passed through a series of selection criteria, hereafter called cuts, to remove background events (as will be shown in later chapters). Once this has been done, any remaining events that have passed all cuts and cannot be otherwise explained, are treated as candidate signals. The rate of these signal-like events is measured as a function of exposure and is reported in units of inverse kg-days. To see how this count rate is converted to an interaction cross-section limit, a discussion of the relevant components is outlined below.

1.3.1 Event Rate

The simplest form for the differential energy spectrum of dark matter is

$$\frac{dR}{dE_R} = \frac{R_0}{E_0 r} \mathrm{e}^{-\frac{E_R}{E_0 r}} \tag{1.15}$$

where R is the event rate per dark matter mass, E_R is the recoil energy, R_0 is the total event rate at zero-momentum transfer , E_0 is the most probable incident kinetic energy for a dark matter particle of mass M_D and r is a kinematic factor equal to $4\frac{M_D M_T}{(M_D+M_T)^2}$, where M_T is the mass of the target. Equation 1.15, however, assumes that the detector is stationary in the galaxy. Since the earth is spinning on its axis, around the sun and around the galaxy, the spectral function must be modified to reflect this. The differential rate must also be updated to account for the form factor of the nucleus with which the dark matter interacts and to account for the efficiency of the detector used. Finally, the contributions from spin dependence must be considered.

1.3.2 Rate Modifications

Equation 1.15 can be modified to incorporate the additional factors mentioned above as per equation 1.16. Here the notation of [21] is adopted.

$$\left.\frac{dR}{dE_R}\right|_{observed} = R_0 S(E_R) F^2(E_R) I \tau_m(\epsilon(E_R)) F_U \tag{1.16}$$

where $S(E_R)$ accounts for the Earth's motion through the galaxy, $F^2(E_R)$ contains the nuclear form factor corrections, $\tau_m(\epsilon(E_R))$ is the efficiency-weighted exposure, I covers the spin-dependence, F_U is the fraction of usable mass of the target material. Each component is discussed below.

The Earth's movement through the halo results in a variation of its velocity with respect to the dark matter halo. This can be accounted for by a simple modification to the basic rate spectrum:

$$\frac{dR}{dE_R} = c_1 \frac{R_0}{E_0 r} \exp^{-c_2 \frac{E_R}{E_0 r}} \tag{1.17}$$

where c_1 and c_2 are constants. The values of c_1 and c_2 vary to reflect the date under consideration, ranging from 0.738 to 0.771 for c_1 and 0.534 to 0.592 for c_2, however it is sufficient to use an averaged value for each [21]. The values used are $c_1 = 0.751$ and $C_2 = 0.561$.

The nuclear form factor, F^2, accounts for the non-spherical shape and constituent nature of the nucleus. This becomes an important factor when the momentum transfer, q, is large enough to make the wavelength h/q comparable or smaller than the nuclear radius. The result is a reduction in the effective cross-section. When this factor is introduced, the cross-section becomes separated into a zero-momentum transfer component and the correction factor F, as in 1.18.

$$\sigma(qr_n) = \sigma_0 F^2(qr_n) \tag{1.18}$$

where $q = \sqrt{2M_T E_R}$ is the momentum transfer, σ_0 is the zero-momentum transfer cross section and r_n is the nuclear radius. In the case of a spin-dependent calculation, the modification is made through use of the first spherical Bessel function, $j_0(x) = \frac{\sin(x)}{x}$. In this case $x = qr_n$. This construction is based on the thin-shell model which represents a single-nucleon on the outer shell of the nucleus.

When considering the spin-dependence correction, any spin-paired nucleons can be thought of as contributing a net scattering amplitude of zero. A better representation is the odd-group calculation which considers the net contribution of nucleons with spin the same as the unpaired nucleons. This is done by using the magnetic moment of the nucleus in question and the nucleons themselves. The result is a multiplicative factor:

$$I = C^2 S_{odd}^2 \frac{j+1}{j} \tag{1.19}$$

where C takes into account the spin content of the constituent quarks (equation 1.20),

S_{odd} contains the magnetic moment calculation (equation 1.21) and j is the usual total angular momentum quantum number.

$$C = \Sigma_q T_q^3 \Delta q \quad (q = u, d, s) \tag{1.20}$$

$$S_{odd} = \frac{\mu_{mag} - g_N^l j}{g_N^s - g_N^l} \tag{1.21}$$

In equation 1.20 T_q^3 is the third component of isotopic spin for each quark and Δq is the spin contribution of the quark to the spin of the nucleon. In equation 1.21 μ_{mag} is the nuclear magnetic moment, g_N^l and g_N^s are the dimensionless magnetic moments of the nucleon N for angular momentum l and spin s respectively.

In order to be able to compare results from experiments that use different target materials, the result is normalised to represent the spin-dependent cross-section for the nucleon target via equation 1.22:

$$F_N = \left(\frac{\mu_p}{\mu_T}\right)^2 \frac{[\lambda^2 j(j+1)]_p}{[\lambda^2 j(j+1)]_T} \left(\frac{C_{Wp}}{C_{WN}}\right)^2 \tag{1.22}$$

where μ_p is the reduced mass for the proton, with mass m_p and dark matter mass M_D. μ_T is the reduced mass for the target with M_D, λ accounts for the proton being a composite particle, C is the WIMP-nucleon spin factor for the proton (p) and the nucleon under consideration (N) - in the case of CF_4 this is the proton. The values used in this thesis for these parameters are m_p = 938.272 046(21) MeV/c^2 [48], C_{WP}^2 = 0.46 [21].

Since there is an efficiency attributed to the detector response at different energies, this must be used as a weighting for the exposure used in the rate calculation. Owing to the energy dependence, this must be incorporated into the full integrated equation such that it is evaluated correctly for the energy range under consideration. The efficiency of the detector is estimated through the use of simulated events, more on

this can be found in section 3.3.6.

The parameter F_U of equation 1.16 represents the usable fraction of the total fiducial mass of the detector. This is to take into account composite targets and in the case of CF_4 represents the fractional amount of fluorine to the total CF_4 molecule as in equation 1.23.

$$F_U = \frac{4 \times 19(F)}{4 \times 19(F) + 12(C)}. \tag{1.23}$$

1.3.3 Setting a Limit

With all of the modifications accounted for, the event rate per dark matter mass R can be calculated. This is compared to the expected rate, an example of which is given in figure 1-12, in order to set a limit on the strength of the interaction cross-section. For the purposes of the analysis presented in chapter 3 the Feldman Cousins Technique was employed [49].

The Feldman Cousins method is a conservative way of assigning a confidence interval to a measured value x. It is conservative in that the 90 % confidence limit that is reported can in fact be slightly higher than 90 %. To demonstrate the method I will use the prescription outlined in [49] where a Poissonian process with background is considered. In this example the probability distribution follows equation 1.24.

$$P(n|\mu) = \frac{(\mu + b)^n \exp^{-(\mu+b)}}{n!} \tag{1.24}$$

where b is the background mean, n is the number of observed events and μ is the expected mean. In order to calculate the range of expected mean values μ is found such that equation 1.25 is satisfied.

$$P(\mu \in [\mu_1, \mu_2]) \geq \alpha \tag{1.25}$$

where μ_1 and μ_2 are the lower and upper limit of the range of μ i.e. the set of values that μ could be in, and α is the confidence level, more often than not chosen to be 0.1. This means that the value μ will be contained within the set of values μ_1 to μ_2 in 90 % of the experiments that are performed to measure μ. This equation describes the coverage of the experiment. If the probability $P(\mu \in [\mu_1, \mu_2])$ were in fact less than the quoted α, then this is under-coverage. If it is greater than α, this is over-coverage. Both of these scenarios are undesirable, the result should be quoted correctly. However, over-coverage is the better of the two since under-coverage implies the result is better than it actually is.

To determine this range of values for a Poissonian process with background events the probability value from equation 1.24 for a range of μ values is calculated at multiple n values. i.e. multiple calculations must be done to deduce the probability of getting a value n given that the mean is the value being considered. The probability is also then calculated for each n under consideration for different μ values and the best result is chosen giving μ_{best}.

With an upper limit on the observed rate of events deduced, this can then be translated into an upper limit on the cross section interaction via equation 1.26 [21].

$$\sigma_0 = R_0 \frac{M_D M_T}{\rho_D v_0} \tag{1.26}$$

where σ_0 is the zero-momentum transfer cross-section, M_D is the mass of the dark matter particle, M_T is the target mass, ρ_D is the dark matter density and v_0 is the mean dark matter velocity. By substitution of R_0 from equation 1.17, the spin-dependent cross-section can be estimated.

Chapter 2

DMTPC

This chapter provides an overview of the DMTPC experimental status and how the fundamentals outlined in the first chapter have been implemented. Section 2.1 discusses the readout channels used by the DMTPC detectors. Section 2.2 describes the DMTPC-specific backgrounds. Section 2.3 gives a description of the detectors that have been used, including a first look at the latest iteration and the backgrounds specific to the DMTPC experiment. Section 4.2 deals with how the readout from the detectors is simulated, for the purpose of understanding how the detectors are expected to behave. Section 3.2 outlines the calibration procedure which describes how the detectors are characterised and prepared for the data taking process. Section 3.4 details the data processing, wherein potential signal events are extracted from the recorded data and the method used for producing a cross-section limit.

2.1 Readout Channels

2.1.1 Charge Readout

The electrons generated in the amplification region produce two signals that can be read out. The first is the collection of the electrons on the amplification region

anode, which produces a current. A charge-sensitive preamplifier is used to integrate the current and output a voltage that is proportional to the number of electrons. The amplifier used is a Cremat 113 which has a gain of 1.3 mV/pico-Coulomb and a 1 ns response time. The lower energy threshold for this is in the order of a few keV. This provides a measure for the energy of the recoiling nucleus. The second uses a fast amplifier on the ground electrode to read out the current pulse induced by the ions and the image charge of the electrons. Due to the large difference between the drift velocities for electrons and ions, in CF_4 $\mathcal{O}(\text{cm } \mu\text{s}^{-1})$ and $\mathcal{O}(\text{mm } \mu\text{s}^{-1})$ respectively, the pulse generated shows a sharp peak followed by a wide distribution. The fast pulse readout can be used as a means of event identification.

To show the type of pulses observed, the examples from [50] are shown, see the reference for more on the subject. For that study, the end of an alpha track was used to mimic nuclear recoils. This is representative of nuclear recoils since at the end of the alpha track the energy is low enough such that the energy depositions and stopping powers are comparable. Figure 2-1 shows a diagram of how this method was employed, with the alpha outside of the drift region, pointing at an angle into it. SRIM [51] calculations report that over 99% of the alpha energy goes into ionisation, as such the units given for these pulses is given in keV_α, this is similar to the $\text{kev}_e e$ unit, as described in section 1.2.4.

Pulses with the peak distribution described and shown in figure 2-2 a) are generated from these end-of-alpha track recoils. Since the track is short, the electrons enter the amplification region almost simultaneously. The pulse of figure 2-2 b) is from an electronic recoil. Electronic recoils, such as those from compton-scattered photons, produce longer tracks. This is a consequence of multiple scattering within the drift region and the low mass of the electron causing significant deflection from the initial scattering direction. The longer track results in a pile of sharp peaks which produces the more spread out distribution.

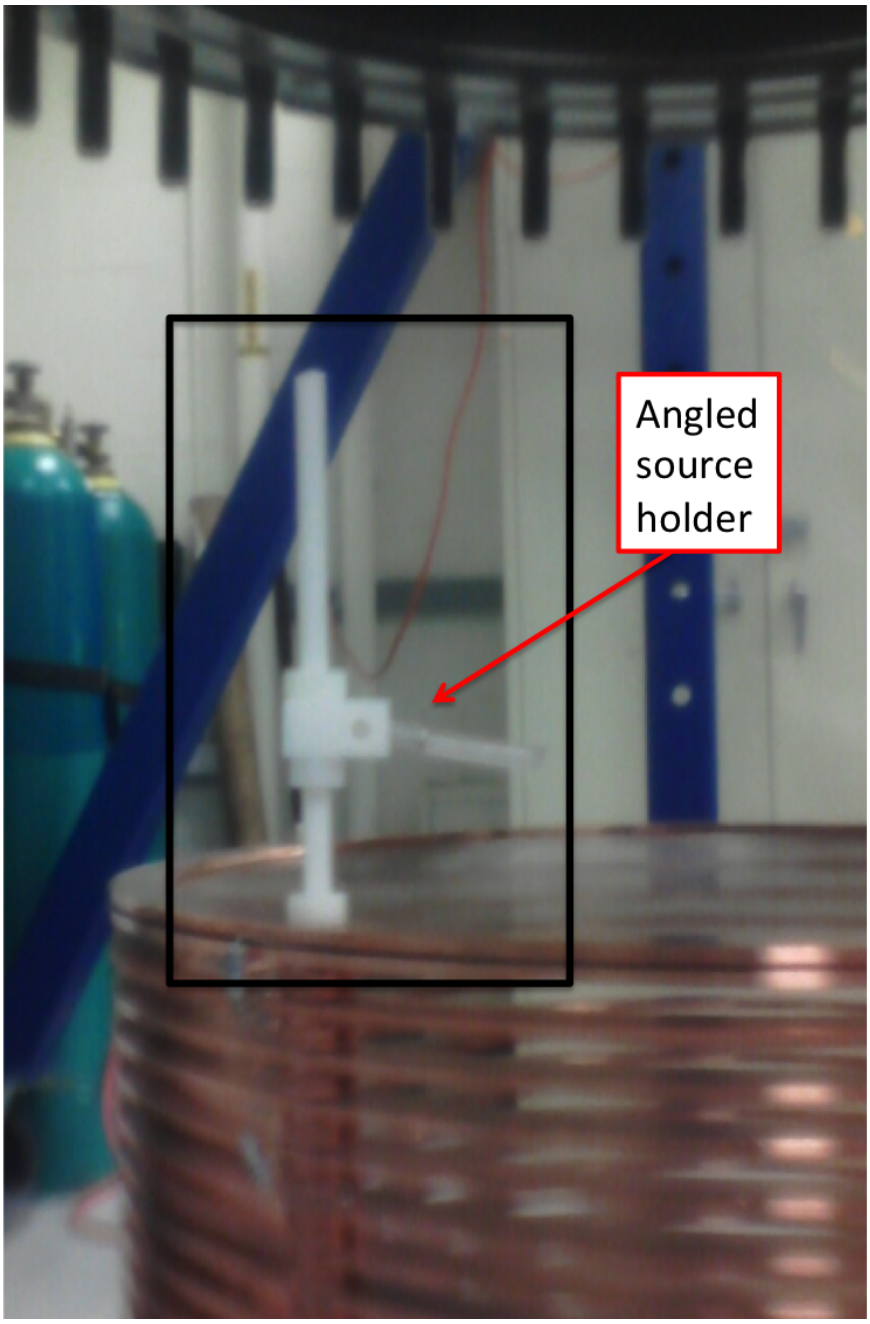

Figure 2-1: This photograph shows the angled placement of an alpha source. The alphas are aimed into the drift region such that only the end part of the track enters the drift region.

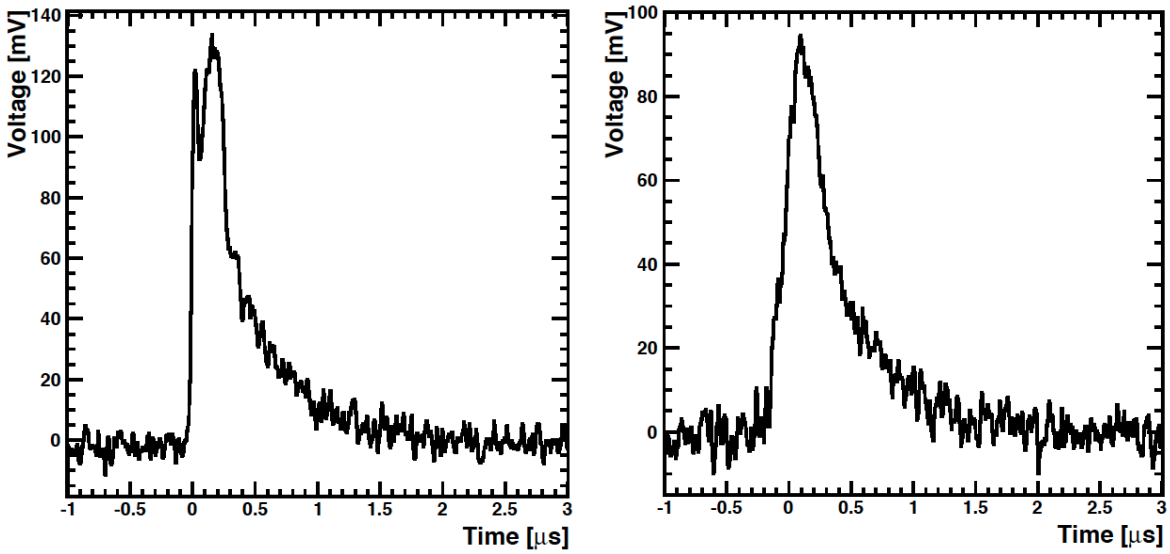

(a) Nuclear recoil event

(b) Electron recoil event

Figure 2-2: An example of a pulse for a) a nuclear-recoil-like event, generated by the tail end of an alpha track. The equivalent energy is 75 keV$_\alpha$. b) shows an electron recoil event with an energy of 60 keV$_\alpha$. Both are readout from the TPC ground plane. The ordinate is voltage in mV and the abscissa is time in μs. A double peak structure can be seen in a) and not in b). This can be used to identify event types.

2.1.2 Image Readout

CCDs are used to image the scintillation photons that are produced at the amplification stage - see section 1.2.4. An example of an alpha track can be seen in figure 2-3 and a candidate nuclear recoil in figure 2-4. The reconstructed parameters for the alpha track are labelled as in-image as the source is placed outside of the imaging area, so the start of the track is not imaged.

An important aspect of the CCD readout is that the image is a two-dimensional projection of the full three-dimensional track. A representation of this is shown in figure 2-5. From the image, parameters that describe the track such as: energy, range, width and sense[1], are reconstructed using the analysis software described in section 3.4. Figure 2-6 shows a close up of candidate track with some of the parameters labelled. For each data set, before data taking begins, a series of one hundred images

[1]The sense of the track is the two-dimensional direction of the track.

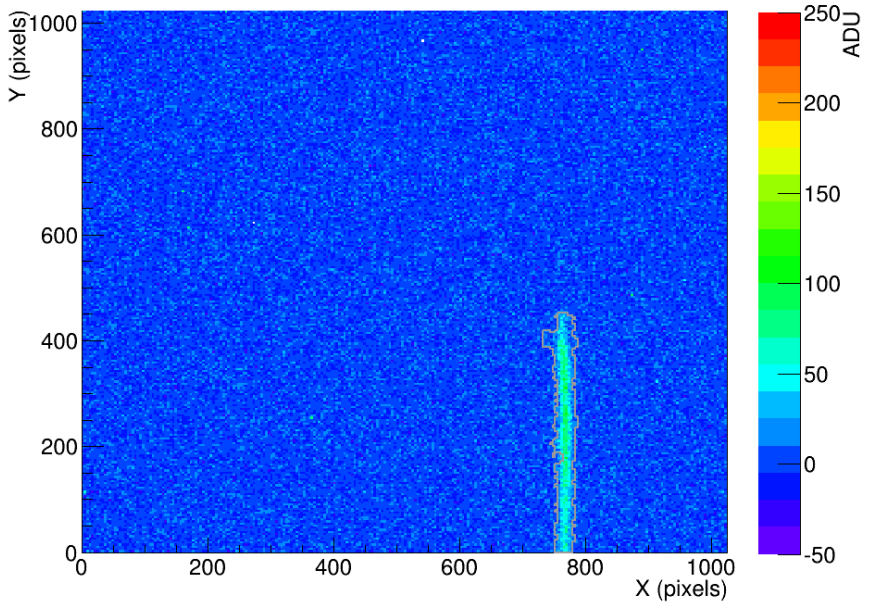

Figure 2-3: An example of an alpha track as seen in the CCD. The abscissa and ordinate are in pixels, the applicate is in analog-to-digital units (ADU). The in-image, reconstructed energy and range for this track are 2.6 MeV_{α} and 77 mm.

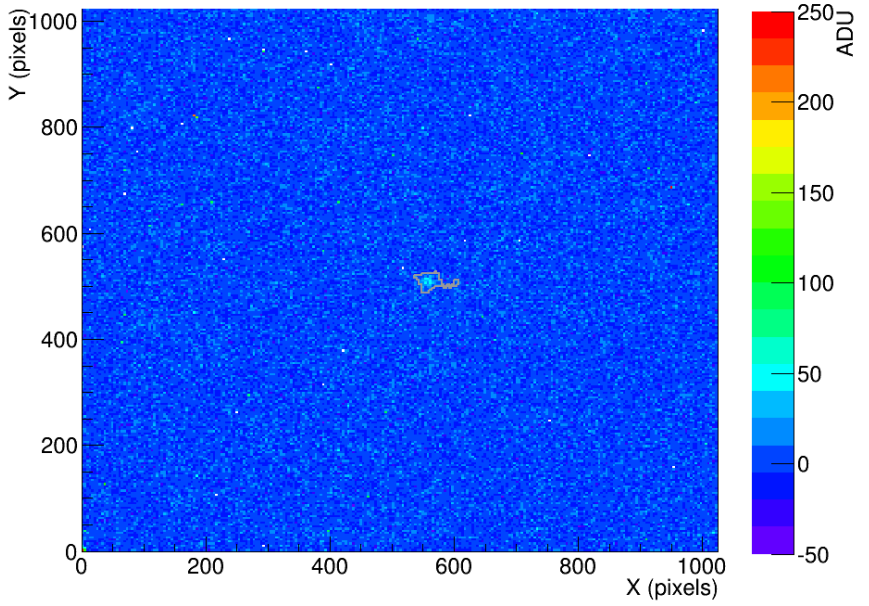

Figure 2-4: An example of a candidate nuclear recoil track as seen in the CCD. The abscissa and ordinate are in pixels, the applicate is in ADU. The reconstructed energy and range for this track are 91 keV_{ee} and 6.34 mm.

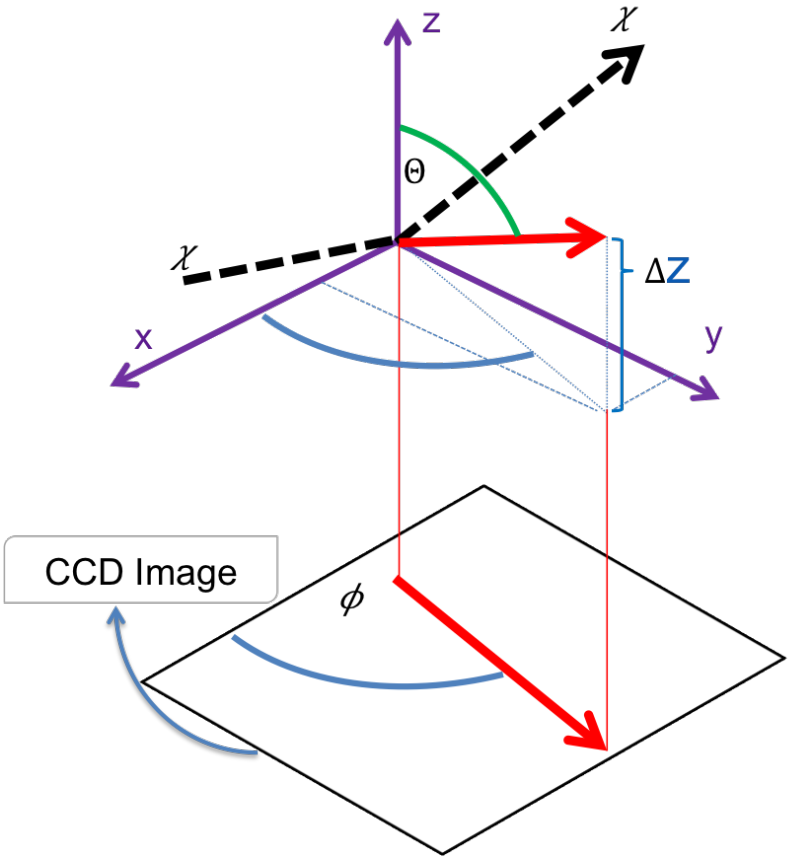

Figure 2-5: A diagram of the two-dimensional projection of a three-dimensional track. The coordinates shown in this diagram define the DMTPC co-ordinate system. χ is an incoming dark matter particle, the red arrow represents a recoiling nucleus, Φ is the azimuthal angle and Θ is the polar angle.

is taken, with the shutter closed. These images are summed and averaged in order to give a guide to the baseline value of the pixel intensities across the chip. This averaged image, called the bias-frame, is saved with the data for bias-subtraction, again, described in section 3.4.

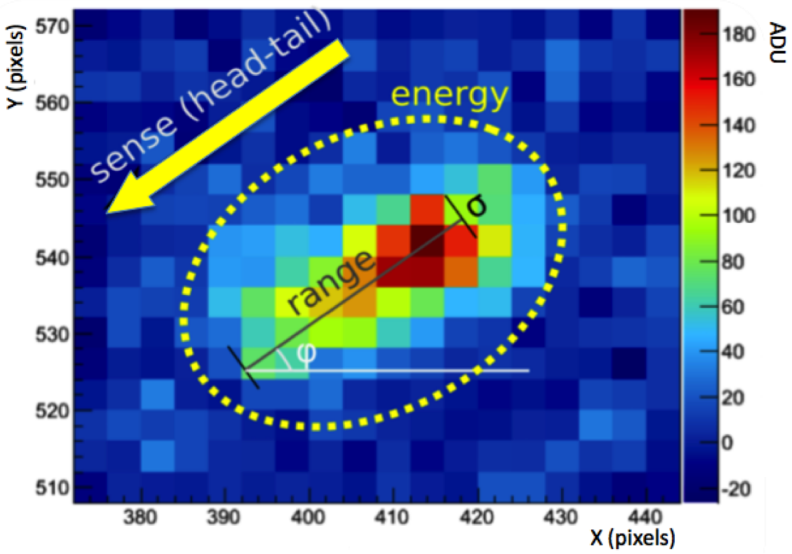

Figure 2-6: Close-up image of a reconstructed track with parameter labelling.

2.2 Backgrounds

The backgrounds specific to DMTPC can be placed into two categories: physical and non-physical. The physical are those caused by background radiation, such as those described in section 1.2.6, and sparks in the amplification region. The non-physical are a consequence of artefacts that occur in the CCD image, for example due to hot or cold pixels[2]. However, a physical background can also cause an artefact in the camera, namely if a cosmic ray passes through the CCD chip, it can leave a trace. Below is a list of backgrounds currently accounted for:

1. "Worms" : This is the term given to tracks that are reconstructed from the CCD image but have no correlation to a physical event within the fiducial volume of the detector. These include the aforementioned cosmic ray events, hot and cold pixels, as well as noise fluctuations in the image.

2. Sparks : Due to an effect called Paschen breakdown, at low pressure and high voltage, sparking can occur in the amplification region. The Paschen breakdown

[2]These are caused by defective or dead pixels in the CCD chip

voltage is given by equation 2.1 and is shown in figure 2-7.

$$V_b = \frac{Bpd}{ln(Apd)}, \tag{2.1}$$

where V_b is the breakdown voltage, p is the pressure, d the separation distance and A and B are constants related to the gas. For CF_4 they are 11 and 213 respectively [53]. The The bright flash of the spark can saturate part of the CCD image and the brief loss of voltage results in a temporary drop in gain.

3. RBI : RBI stands for Residual Bulk Image. This is a well-known phenomenon in CCD imaging [54], whereby long wavelength light can cause trapped charge in the CCD substrate-epitaxial interface. The trapped charge slowly diffuses over time causing a glowing effect in the affected region. These can result in a reconstructed event which have reconstructed parameters indistinguishable from nuclear recoils. However, due to the slow diffusion of charge, an RBI will show up in subsequent images in the same location, providing the possibility for identification. Sparks are the dominant cause of RBIs.

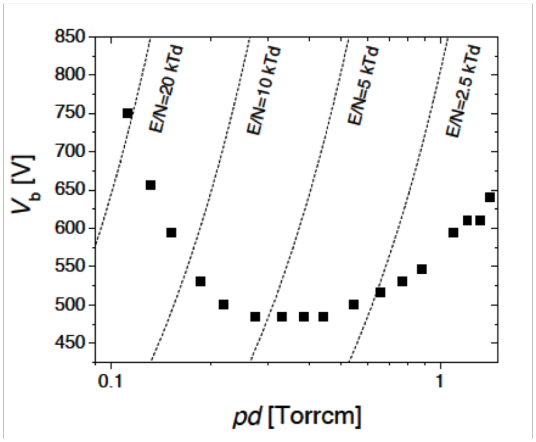

Figure 2-7: The Paschen curve for CF_4 [52]. The curve shows the voltage at which breakdown, or sparking, will occur as a function of the pressure $\times$ electrode separation distance (pd).

2.3 Detectors

This section describes the three main detectors used by DMTPC. This serves to demonstrate the progression in design and how each has fed into the next. The detectors all consist of a vacuum chamber in which is housed one or more TPCs, as described in section 1.2.3 and shown in figure 1-17. Due to the use of imaging as a readout, described in section 2.1.2, it is necessary for the cathode, anode and ground electrodes of DMTPC field cages to allow the transmission of light. To achieve this, a steel-wire woven mesh is attached to a metal frame. The mesh for the amplification region has 100 lines-per-inch, resulting in 79 % open area, the cathode uses mesh with 50 lines-per-inch and has 88 % open area. The purpose of the more densely woven mesh is to allow the mesh to be stretched and glued to its frame at tension, typically $\mathcal{O}(20)$ N/cm. The tension is required because of the stronger electric field of the amplification region, to resist electrostatic attraction, and also to keep the plane as flat as possible and minimise any non-uniformity in the gain. To ensure the electrodes of the amplification region remain separated, a small amount of material is placed in-between in the form of thin dielectric strips, referred to as spacers.

2.3.1 Ten-Litre: 10L

The 10L detector consists of two back-to-back TPCs, each consisting of nineteen stainless steel rings of 1 mm thickness and spaced at 1 cm intervals. The total height of the drift region is 19.7 cm and the ring inner diameter is 27 cm. The rings are connected by 1 MΩ resistors and with a -5 kV voltage applied to the cathode, the electric field strength is 254 V/cm. The rings are kept in place by eight acrylic rods spaced evenly around the ring and are kept apart by acrylic spacers that sit on the rods. The cathode is made from a stainless steel mesh, with 88% open area, glued onto the top ring of the field cage structure. The ground of the amplification region is also made from stainless steel mesh with a 79% open area. The anode consists of a copper plate on a piece of fiberglass epoxy laminate.

The CCDs used by the 10L are the Apogee Alta U6 which uses a Kodak KAF-1001 sensor chip with a pixel size of 24×24 μm and a 1024×1024 array. The pixels are binned by hardware on the chip as this improves the signal-to-noise ratio (SNR), resulting in a 256×256 array. The physical area that is imaged by the cameras differs for each TPC, owing to each being at a slightly different distance from the amplification plane. Due to the orientation of the detector they have been labelled the top and bottom camera see figure 2-8.

The top camera images an area of 140.29×140.29 mm and the bottom images 176.13×176.13 mm, resulting in a fiducial volume of $\sim 1 \times 10^{-2}$ m^3. The fast amplifier used for the 10L is custom made and has a response time of 40 ns. This fast response allows for the peak separation, observed for nuclear recoils and thus short tracks, as detailed in section 2.1.1. The integrated signal is obtained using a Cremat 113 amplifier which has a slower response time of 1μs. With two readouts on each TPC this results in four signals overall which are connected to an Alazar ATS 860 digitizer.

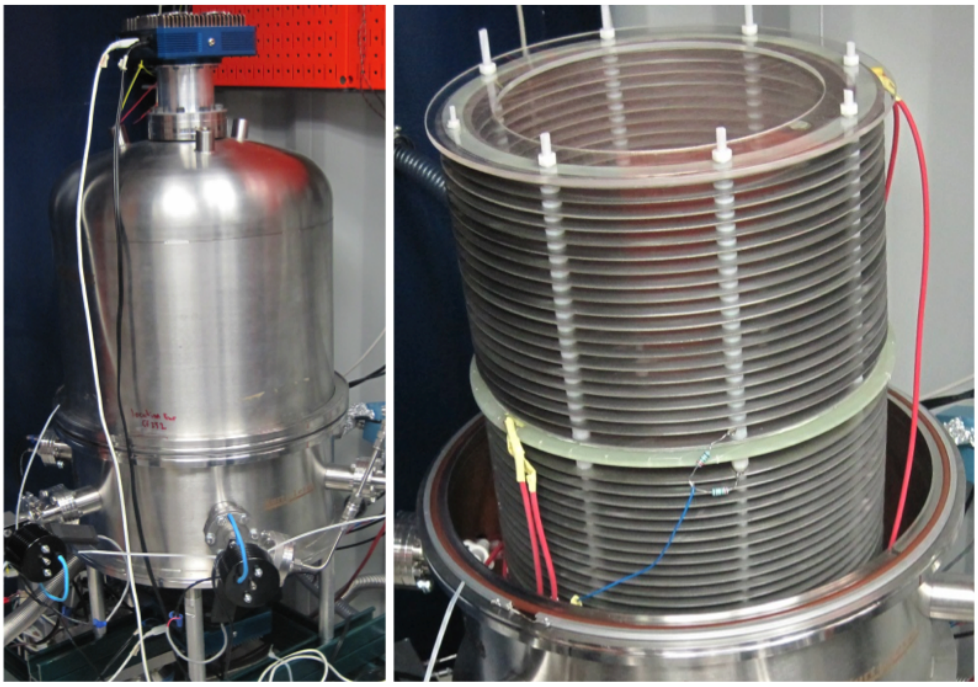

Figure 2-8: The 10L detector. Left: the enclosed detector. Right: the internal TPC structure.

2.3.2 Four Shooter: 4Sh

The 4Sh is so called due to its having four cameras imaging the amplification region. For this detector there is a single TPC, measuring 26.7 cm tall, composed of twenty-six copper rings. The diameter of the anode is 29.21 cm resulting in a fiducial volume of $\sim 2 \times 10^{-2}$ m^3. This is a two-fold increase in the volume over the 10L. The rings are connected via 1 MΩ resistors and are spaced using a stack of copper and acrylic - this is to reduce the amount of acrylic in the detector and thus reduce the amount of Radon present. The meshes used are the same as for the 10L as is the Copper-plated fiberglass epoxy laminate for the anode.

The four CCDs are the same Apogee Alta U6s as used for the 10L. They are arranged such that they image a quarter of the full amplification region each, an example image can be seen in figure 2-10.

The benefit to this arrangement is that the full amplification region is imaged, not just a central region as per the 10L. This is useful since the field-shaping rings

Figure 2-9: The 4Sh detector. In this image the TPC is exposed and can be identified by the copper structure in the lower section. The upper steel part is the bell-jar of the chamber and during operation sits over the top of the TPC, enclosing the volume.

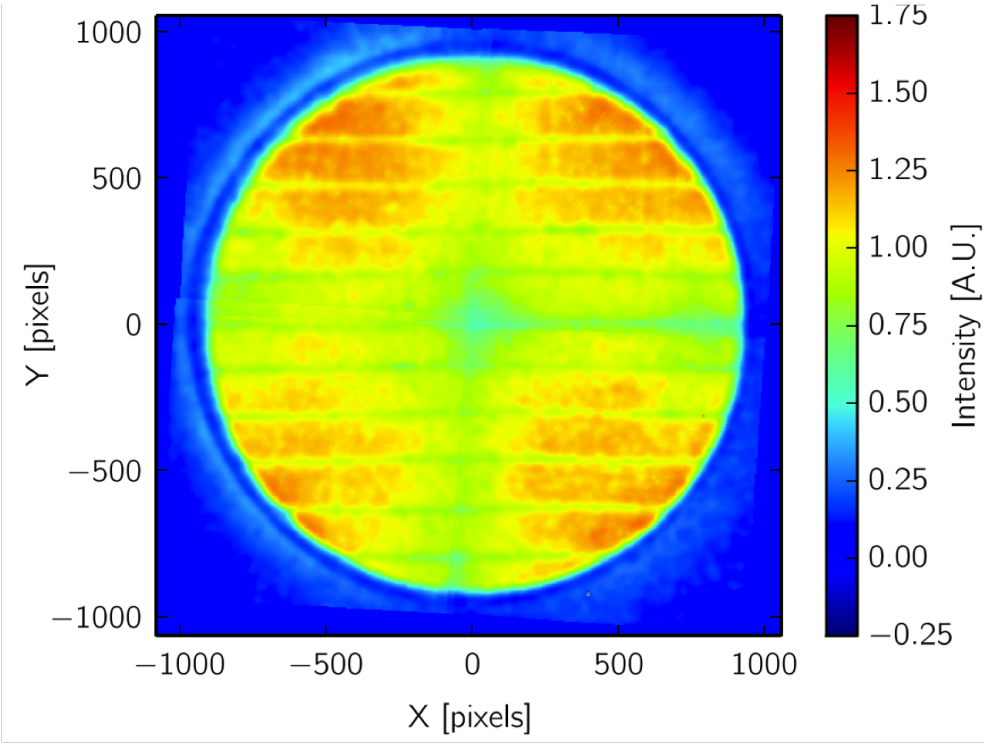

Figure 2-10: The coverage of the CCDs for the 4Sh detector. The central circular region shows the amplification region and plots the gain variations. The horizontal structures that can be seen are due to dielectric thread used to separate the ground mesh from the anode. These threads cause a drop in gain. This image is composed of four CCD images stitched together. The scale conversion is 0.1604 ± 0.0004 pix/mm [55].

are a dominant source of alpha backgrounds and an alpha originating from them will be visible in the image.

2.3.3 One Meter Cubed: $1m^3$

The $1m^3$ detector is the fifth DMTPC prototype to be built and consists of a fifty-fold increase in the fiducial volume (from the four-shooter). At this scale it is not practical to have a vertical detector due to the weight of the bell jar that would be required, as well as the height of the room needed to enable removal of the bell jar. It is also necessary to divide the drift region up, much like the 10L detector, as increasing the height (or in this case length) of the drift region, results in increased diffusion from tracks that start near the cathode, as seen by equation 1.13. The change in detector orientation also requires a new mechanical support system for the field cage to be designed. A detailed description of the design for the $1m^3$ TPC can be found in chapter 4. The imaging for the $1m^3$ is a hybridisation of the 10L and the 4Sh. There are two sets of back-to-back TPC constructions within the vacuum chamber, each with a triple mesh configuration for the amplification region, see figure 2-11. The triple mesh is connected such that the central mesh is at voltage and the two outer meshes are at ground. This provides two back-to-back amplification regions that can be imaged from a single side. For one of the TPC sets, four Fingerlake Proline 09000s image in the same manner as the 4Sh. The other triple-mesh is imaged using a single camera, a Spectral 1100S with a Fairchild 486 sensor. Though only a single camera is used, the full amplification region is still imaged.

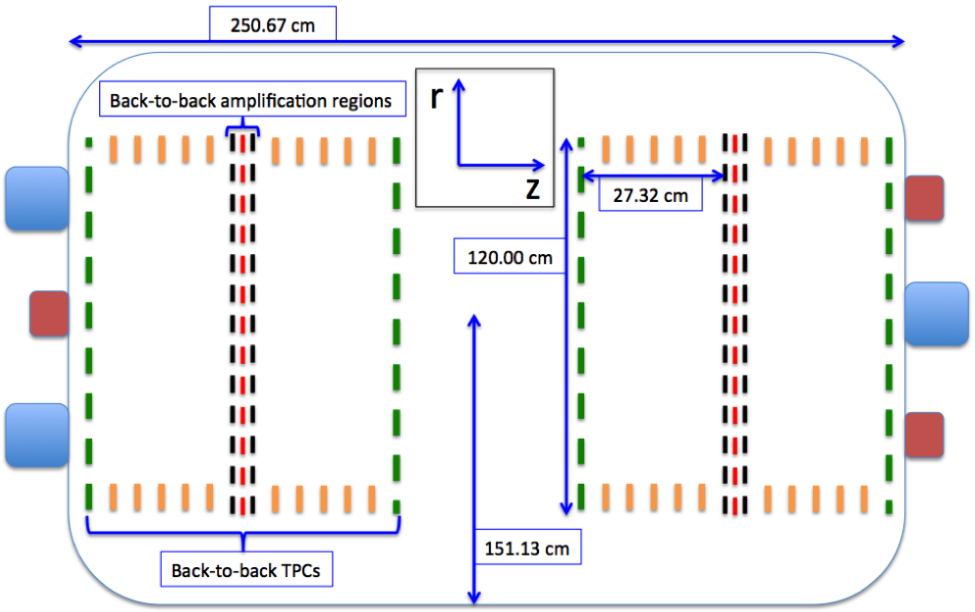

Figure 2-11: Diagram representing the $1m^3$ detector configuration, not to scale. This shows two back-to-back TPC configurations, with colour coding the same as figure 1-17, totalling four field cages. The box surrounding the TPCs represents the vacuum chamber. On the outer edge of the chamber, the blue boxes represent CCD cameras and the red boxes represent PMTs (photo-multiplier tubes).

2.4 Simulation

A simulation is used to generate the kind of output expected to be seen in the detectors. These simulations are used in the analysis procedure at various stages including the Multivariate Analysis (section 3.3.4), the calibration procedures (section 3.2) and for deducing cut efficiencies (section 3.3.6). On commencement of this project a CCD readout simulation was in place but no charge readout simulation had been done. The charge is a very powerful means of discriminating backgrounds so a quick simulation was used and is described below.

2.4.1 CCD Readout Simulation

The simulation is a basic one and uses the assumption that the energy deposited occurs along a straight line. This is not entirely the case, particularly at low energies.

New simulations have been prepared [47] that work to overcome this simplification but were not used for this thesis as they have been prepared as part of a new software architecture that is still in the process of completion. The premise of the simulation used is to propagate particles based on the stopping power tables generated by the SRIM package [51]. At each incremental step in range, dx, an energy deposition is calculated as $\Delta E = \frac{dE}{dx}dx$. This is then multiplied by the system gain and used to produce a histogram as seen in the CCD cameras. In order to best represent the data collected, the tracks generated from this process can be 'mixed' with real images, taken while the detector is in gas-refilling mode. This process is referred to as cosmixing, since the majority of artefacts are of cosmic orgin. The images taken during this time will not contain any tracks that occur in the detector, since the TPC voltages will be switched off. They will contain, though, any CCD artefacts. Mixing this cosmic data with simulated events gives a more realist representation of the detector image output.

2.4.2 Charge Readout Simulation

In order to make an estimation of the expected charge readout for a particular generated track, a charge-light calibration, referenced in section 3.2.6 is used to convert the integrated cluster of the CCD simulation into a voltage value. On top of this is, it is necessary to add a noise value. For this, a source-free data set was used (see section 3.1). So that only the noise and no signal was measured, only the bins that come before the start of a pulse are included in the estimation. Two simple cuts were applied to remove anomalous pulses that were observed, caused by sparks, the recovery of the circuitry post-spark and pulses longer than the scope time window or 'trailing' pulses. Figure 2-12 shows an example of a trailing pulse. The first cut removes values that are outside the scope range of -0.1 V and 0.1 V, removing railing events. To remove trailing pulses, a cut was made on the first bin entry value, to be

greater than -0.7 V but less than 0.7 V. These values come from visual inspection of a selection of typical nuclear recoil pulses, see figure 2-13.

In the resulting distributions, figures 2-14 and 2-15, there can still be seen populations below zero that do not agree with the expected distribution, that of a Gaussian centred around zero. Despite attempts to remove these values, they have persisted. A Gaussian was fitted to the distribution and the resulting mean and standard deviation used as a measure of the charge readout noise. This fit, since it includes the additional low population, is a conservative estimate of the noise distribution and is therefore reasonable to use. The values obtained for the fit are given in table 2.1. The fitted values were then used to produce a random variable with a Gaussian distribution from which the expected charge was assigned a noise value.

Table 2.1: Charge noise fit parameters.

TPC	Mean	Standard Deviation
Top	1.72e-4	2.053e-3
Bottom	1.29e-3	2.873e-3

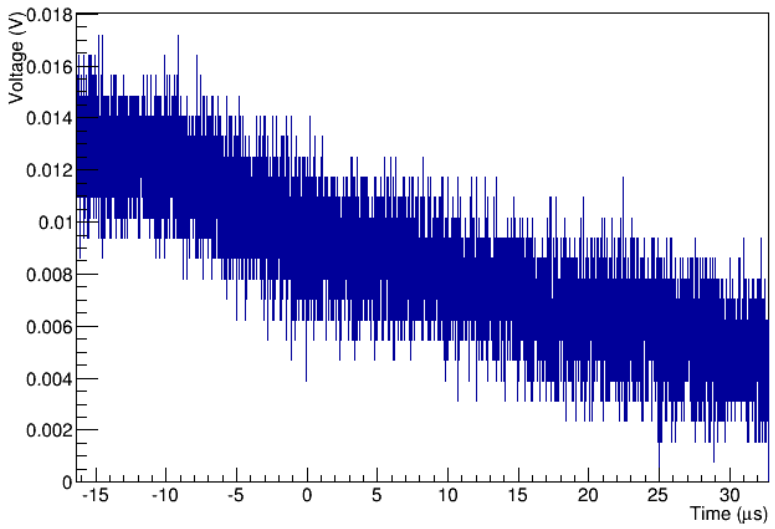

Figure 2-12: Example trailing pulse: the tail end of a pulse which is longer than the read-out window.

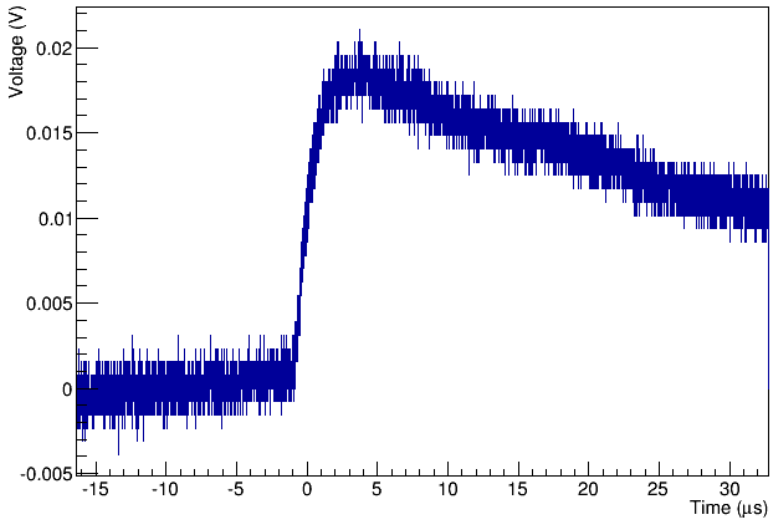

Figure 2-13: Example of a typical nuclear recoil pulse. This is the matching pulse for the recoil of figure 2-4 with reconstructed energy of 91 keVee.

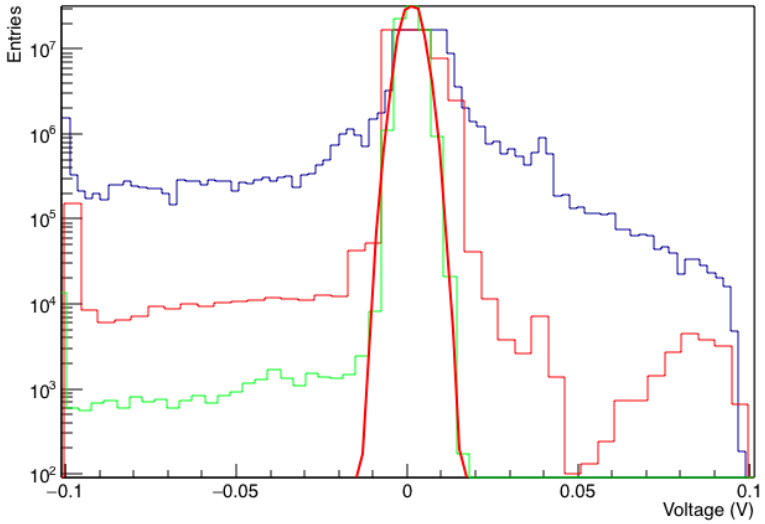

Figure 2-14: Histogram of the pre-pulse bin values after application of cuts for the bottom camera. Blue shows the distribution before any cuts, red shows the effect of the first cut made and the green is the final histogram used, after the trailing pulse cut was applied. The red curve overlaying the green distribution is a Gaussian fit.

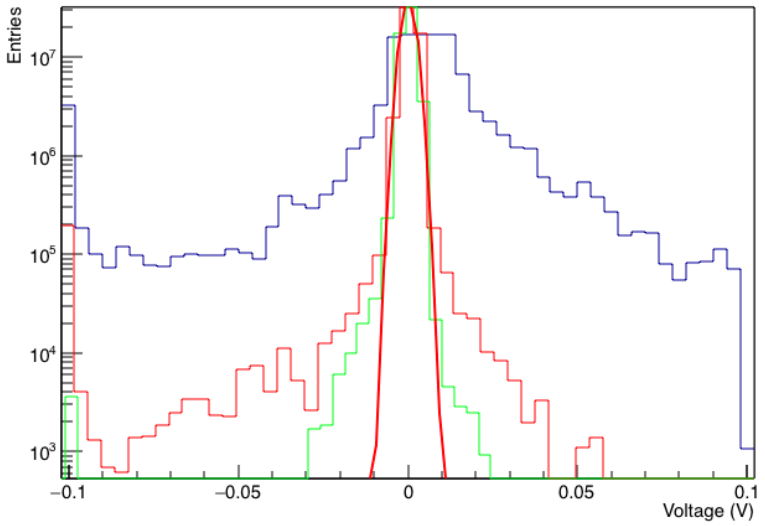

Figure 2-15: Histogram of the pre-pulse bin values after application of cuts for the top camera. Blue shows the distribution before any cuts, red shows the effect of the first cut made and the green is the final histogram used, after the trailing pulse cut was applied. The red curve overlaying the green distribution is a Gaussian fit.

Chapter 3

Analysis of Data Acquired Underground

This chapter details the analysis of a data set acquired using the 10L detector whilst located underground at the Waste Isolation Pilot Plant (WIPP), a salt mine which has been converted to a storage facility for defense nuclear waste. WIPP is located in new Mexico and the mine has a depth of 650 metres (about 2000 metres water-equivalent). Underground operation is a desired aspect of dark matter detection as the layers of rock and earth provide a reduction in the cosmic ray background. This analysis represents the third to be done on data acquired with the 10L detector. The first was a surface run, the purpose for which was to act as a comparison to underground operation as well as verification of suitable operability, prior to underground deployment. The results for this have been published in reference [56]. The second was done on data taken after the detector was placed in the mine. The result of this analysis can be found in reference [57]. It found that the rate of events passing all cuts had in fact increased, compared to overground operation. This is completely counter to the expectation and the reasons suspected were an increase in RPRs and RBIs. The RPR rate is likely to have gone up owing to the fact that the vacuum

chamber was opened between each data set, resulting in the detector internals being exposed to atmosphere. This gives an opportunity for any airborne radon progeny to plate out onto the TPC components. The RBI rate is likely to have increased since the amplification anode voltage was increased from 680 V to 700 V. This voltage increase will produce a higher spark rate and since sparks are the cause of RBIs, it would not be surprising that the rate of RBIs would also increase. This third analysis represents an opportunity to alter the operating parameters and further study the observed backgrounds to the experiment and attempt to remove them.

3.1 The Data

The detector parameters used for each of the three runs mentioned are outlined in table 3.1[1]. The first underground run is referenced as WR5 (WIMP-Run 5) and the second as WR6. The drift voltage was the same for all three and was -5 kV. The anode voltage was increased in WR5, relative to the surface run, in order to increase the gain, but was reduced again in WR6 to reduce sparking. The CCD[2] exposure was increased to five seconds in WR6 to reduce the fractional amount of dead time for a run - time during which the CCD image is read-out. The shutter was turned off in WR5 in case of mechanical failure but was used again in WR6 and monitored. The shutter is used to reduce the CCD exposure to sparks and to eliminate 'cut-off'[3] tracks. To ensure the mechanical failure of the shutter had not occurred during WR6, the mean of the images being taken was checked across multiple runs, during acquisition. An increase in the mean can clearly be seen during the data-taking period and a decreased mean is seen during cosmic data taking, figure 3-1. Cosmic data are

[1] See appendix for the exposure calculation.

[2] Charge-Coupled Device

[3] A cut-off track is an event that occurs when the shutter is not used. Since the CCD is still exposed during readout, if any light is produced at this time, from scintillation or sparks, the pixels still to be readout are exposed. If the illumination spans a region where some pixels have been read and others have not, the image, or a track in the image, has a hard edge and appears to be cut.

acquired when the detector is being refilled with gas and consists of CCD images taken while the shutter is closed. These data are used in the calibration process, to provide a true representation of the noise present in the images and has been termed 'cosmic' since many of the artefacts seen are the result of cosmic ray interactions occurring in the CCD chip. Due to the detector being located underground there were time and accessibility constraints at the time of data acquisition. Therefore the total exposure for WR6 is less than for previous runs. Calibration sources were also only allowed to be placed by WIPP personell, imposing further constraints, as such the data obtained for calibration were limited. The data used for calibration is listed in table 3.2. The gain map uses gamma sources to uniformly illuminate the detector in order to highlight any gain variations. The range calibration uses a neutron source, which also illuminates the detector uniformly. The energy calibration uses an alpha source, attenuated such that the rate is $\mathcal{O}(1)$ Hz. Reducing the rate allows for the acquisition of images with single alpha tracks present, which is necessary for correct determination of the energy conversion factor.

As well as a change in the operating parameters, the bottom camera was re-focussed, as it was found to be not fully focussed after initial deployment. Additionally, during the data taking period for WR6, the bottom camera developed additional noise in the image. Noise of this type has been observed in previous cameras that have been operated underground, it has been attributed to the salt that is present

Table 3.1: A table showing the differences between datasets.

Parameter	Surface Run	WR5	WR6
Pressure (Torr)	75	60	60
Anode voltage (V)	680	700	690
Exposure (sec)	1	1	5
Shutter	On	Off	On
Total Exposure (inc. parasitic) (gm-days)	35.7	33.13	24.16

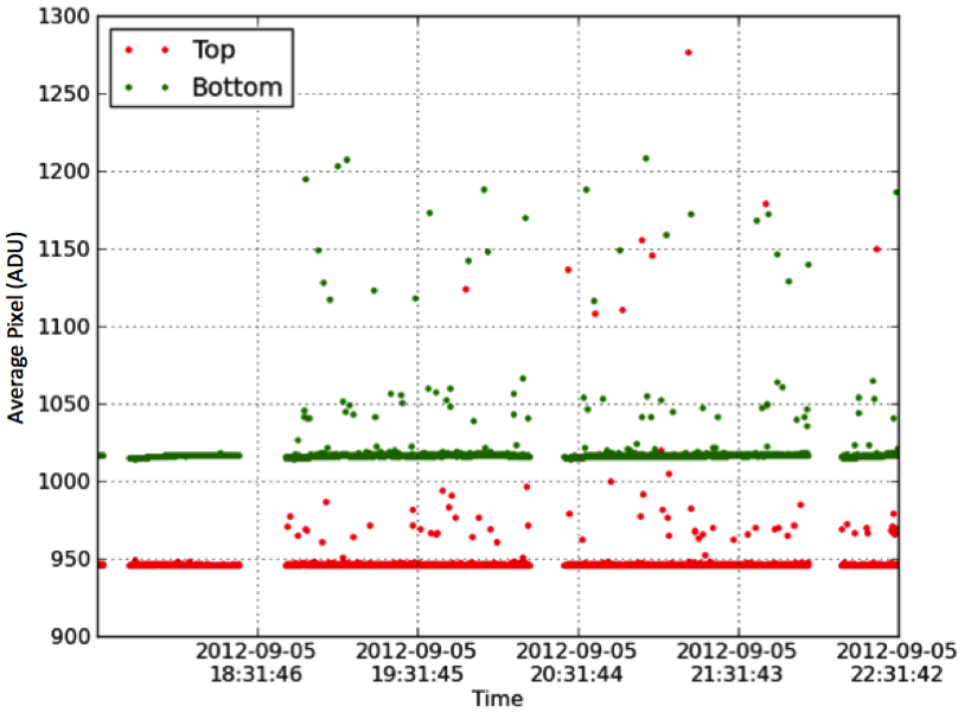

Figure 3-1: Average pixel value versus time. A clear difference can be seen between the population below 18:31 and above 18:31. Below 18:31 is when the shutter was closed and is fairly flat in average pixel vs. time. After this point is when the shutter was opened and there can be seen a number of images with higher average pixel values. However, as most images are empty, there is still a population at the low average pixel value.

Table 3.2: A table of the sources used for detector calibration. The two gain map sources were used concurrently. The energy is most relevant for the energy calibration, specific values for the sources used for each camera are given in section 3.2.4.

Calibration	Source	Decay Product	Energies
Gain map	^{137}Cs	Gamma	662 keV
Gain Map	^{57}Co	Gammas	122 and 136 keV
Range	^{252}Cf	Neutron	2.3 MeV (average)
Energy	^{241}Am	Alpha	5.5 MeV

in the WIPP environment affecting the read out electronics. To verify the data were still appropriate for use, the mean of the cosmic data images was plotted. If this shows a Gaussian distribution, as relied upon by the cluster-finding algorithm, then it is suitable. As can be seen in figure 3-2, this was satisfied. The effect of the noise on the analysis is accounted for by the use of the image mixing procedure outlined in 2.4.1 and the energy threshold, used in the data cuts, is adjusted to reflect the accuracy of the reconstructed energy.

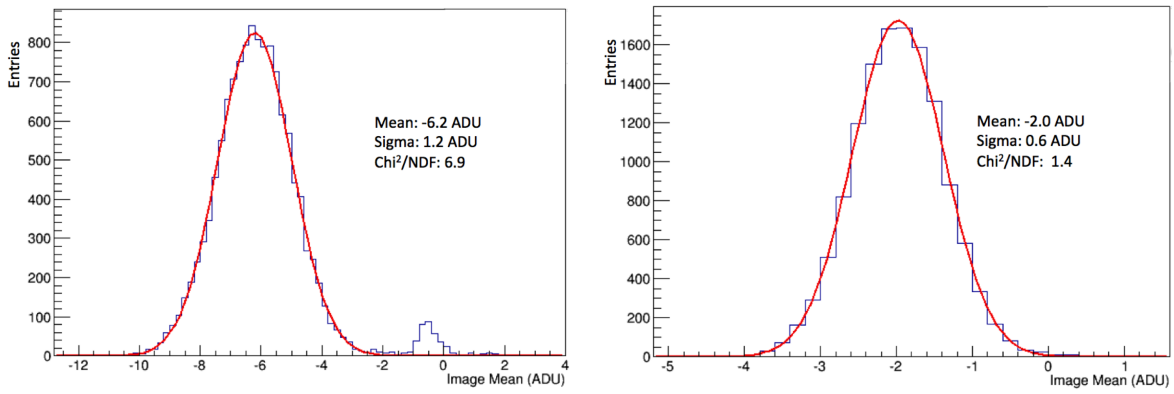

(a) Bottom camera

(b) Top camera

Figure 3-2: Image mean distributions for WR6 cosmic data.The source of the bump in the bottom camera is unidentified.

3.2 Calibration

A calibration procedure is performed to characterise the detector and to be able to convert from read-out values to physical values. There are six parts to the calibration procedure which are given below.

1. Gain map: a plot of the pixel-to-pixel gain variations in the CCD readout.
2. Diffusion: adjusting the readout simulation to represent data.
3. Range conversion: a CCD conversion factor for pixels to mm.
4. Energy conversion: a conversion factor for the CCD to convert ADU[4] to keV.
5. Energy and range corrections: fits to correct for bias introduced by the track reconstruction algorithm.
6. Charge-light conversion: a fit to to compare energy deposition in the charge readout with that of the CCD readout.

Some of the calibration procedures use simulation to compare with what is seen in the detector. In order for the simulation to be as representative of the actual data as

[4]Anologue to Digital Units.

possible, the generated tracks are mixed with real noise from the CCD, using cosmic data. To do this properly the generated image is given a zero bias and read noise and the number of images generated is the same as that of the data that is added. This ensures that the image bias subtraction operates in the same manner as the data sets acquired for analysis. The neutron data used for the range calibration are cleaned by applying a nominal set of CCD cuts: if the image contains a spark; or if a track is an RBI. For the charge-light match, two charge cuts are also made: one on the mesh pulse rise time, requiring: $10 < t_{rise}(25\% - 75\%) < 22$ ns and the other on the integrated charge peak voltage value: $0.01 < \mathrm{V} < 0.18$. The rise time values are guided by a previous study on the charge readout [50]. These, as well as the other cuts made in the analysis, are mentioned in the following discussion however their full description is reserved for the discussion on the data analysis of chapter 3.

3.2.1 Gain Map

A gain map describes the pixel-to-pixel variations in the system gain as seen in the CCD image. To make a gain map the detector is uniformly illuminated with two photon sources, ^{57}Co which decays to ^{57}Fe, producing photons of 122 and 136 keV and ^{137}Cs which decays to ^{137m}Ba, producing a photon of 662 keV. These photons produce a uniform distribution of electron cascades across the detector. The electrons then illuminate the amplification region through scintillation of the gas. In order to measure a signal above the pedestal[5] of the cameras, data are required with the detector both active and inactive: voltages on and off respectively. The voltage-off data provides a value of the readout baseline - the pedestal. Multiple CCD exposures are taken in both configurations then the images are summed together and averaged. The averaging of the data is performed on a pixel-by-pixel basis as sparks and RBIs must be removed. The voltage-off image is then subtracted from the voltage-on

[5]The pedestal consists of the DC offset of the circuit, leakage current and random values induced by variations in material and construction processes.

allowing the source-induced signal to become apparent.

A gain map had already been produced for the 10L detector which was used in the two analyses that preceded the one presented here [58]. On inspection it was found that some of the pixel entries were zero-valued. To correct this, the mean of the surrounding pixels was substituted. The resulting gain maps are shown in figure 3-3.

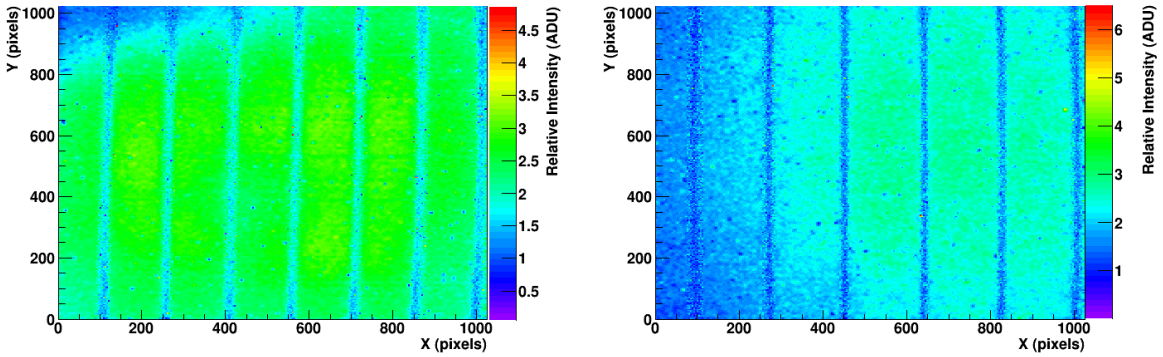

(a) Bottom camera gain map

(b) Top camera gain map

Figure 3-3: Relative gain maps for each camera. The vertical lines are due to a drop in gain at the spacer locations. The left image is the bottom camera and the right image is the top.

3.2.2 Range Calibration

The range calibration is done after the gain map and determines the vixel[6] size and thus the image size, which is used in the generation of simulation. As seen in the gain map, figure 3-3, there is a drop in the gain in the vicinity of the spacers, this can be used to compared the in-image separation to the measured separation of 2.5 ± 0.1 cm [56] to deduce a pixels-to-mm conversion factor. Previously this was achieved using summed images acquired with a gamma source present. When summed, the spacers that are used in the amplification region become apparent due to the reduced gain. Due to the aforementioned constraints, no new gamma source data were acquired for WR6. However, neutron data were obtained and so these data were used instead. The signal to noise ratio was fairly low, however the dip in gain around the spacer locations is apparent (figures 3-4 and 3-6).

To get the calibration factor, the summed histogram is split into four horizontal sections which are then projected onto the abscissa. The location of the low points - corresponding to the spacers - are then selected, resulting in four vertical points for

[6] a vixel is the equivalent space spanned by a CCD pixel at the focal point.

each spacer.[7]. The selected points are then plotted with an ordinate value equal to the median of the section it was taken from: these values are 128, 384, 640 and 896 pixels. These points are then fitted to a line such that any diagonality of the spacer, relative to the vertical of the image, can be accounted for. The final step is to select a number of points along the fitted lines to calculate the parallel separation between any two spacers. This is done by calculating the angle of the line to the vertical and multiplying the horizontal distance by the cosine of this angle - see figure 3-10. The results from this procedure are in table 3.3.

Table 3.3: Range calibration factors.

Camera	Range Calibration (mm/pix)
081264	0.1367 ± 0.0042
100439	0.1718 ± 0.0037

[7]In previous analyses the spacer location has been selected by fitting the spacer regions to a Gaussian and using the minimum. Here, however, the points were selected by eye. As detailed in chapter 3 the bottom camera was refocussed thus a different value to that previously deduced would be expected. The top camera, however, remained unchanged. As such this method would be expected to produce the same calibration factor. This was the case and so the method was verified and implemented for the bottom camera.

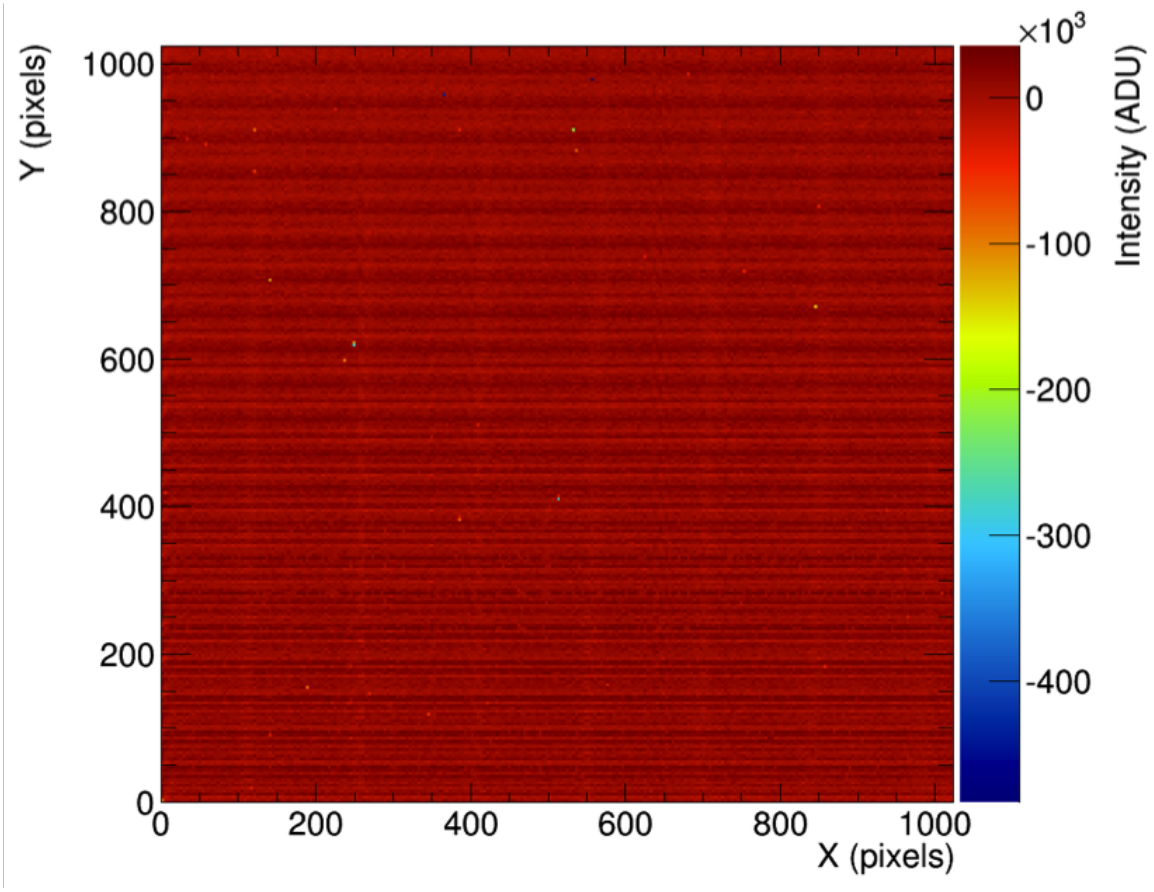

Figure 3-4: The resulting image from summing neutron data for the bottom camera. The horizontal white bars are due to noise in this camera. The spacers can be seen as vertical white lines.

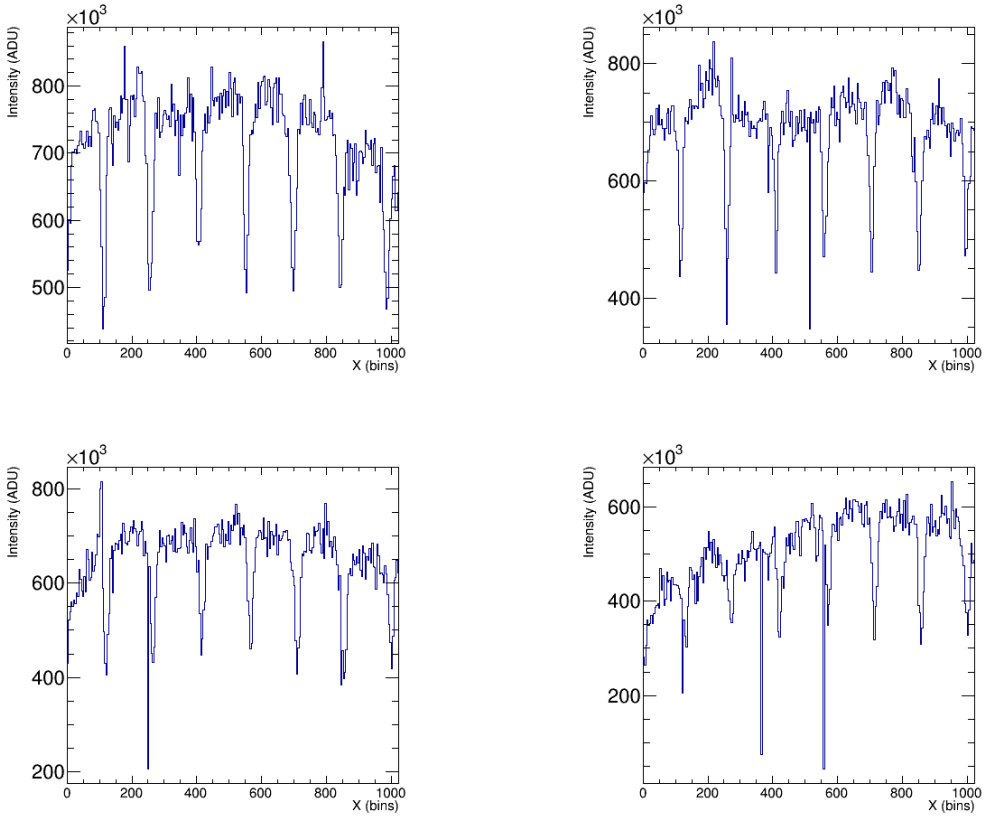

Figure 3-5: The summed neutron data for the bottom camera, projected onto the abscissa at four points along the ordinate. The abscissa is in pixels and the ordinate is in ADU. The spacers are located where the ADU count dips.

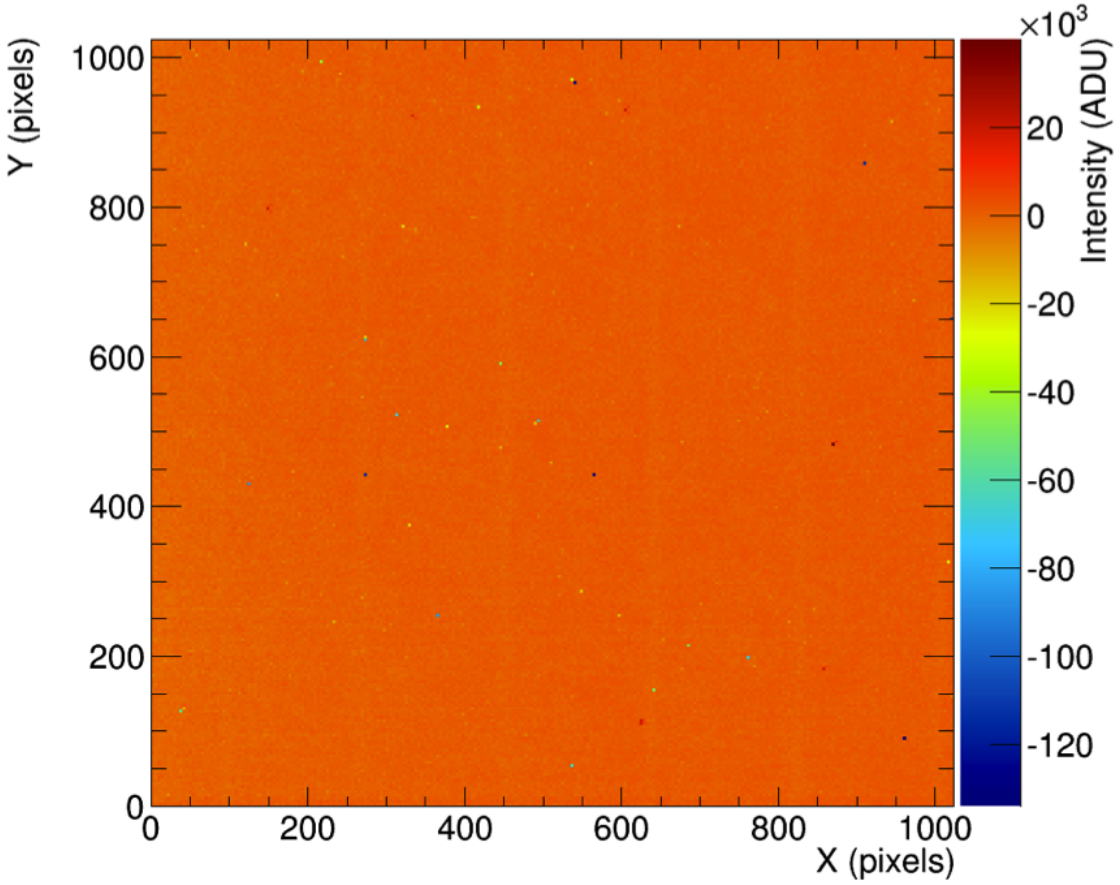

Figure 3-6: The resulting image from summing neutron data for the top camera. The abscissa and ordinate are both in pixels, the applicate is in ADU. There are no horizontal bars present in this image as the top camera did not develop additional read out noise. The vertical dips due to the spacers are faint but visible.

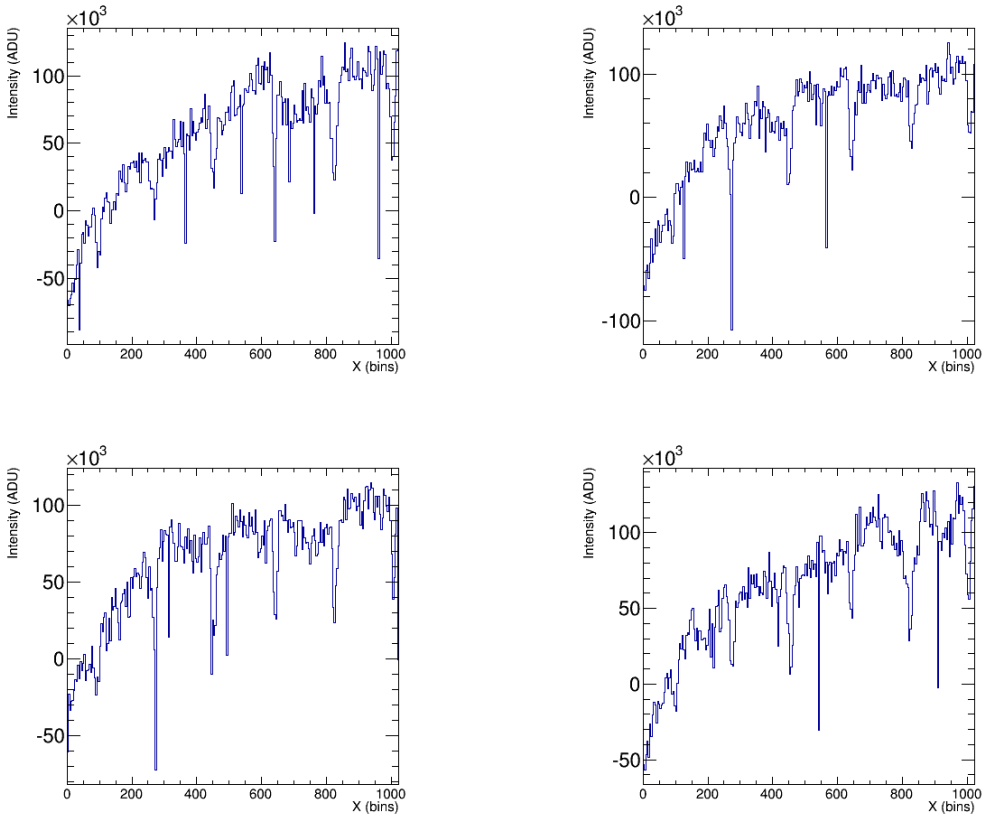

Figure 3-7: The summed neutron data for the top camera, projected onto the abscissa at four points along the ordinate. The abscissa is in pixels and the ordinate is in ADU. The drop in values on the left hand side of each histogram are due to gain variations across the image. This can also be seen in the gain map of image 3-3.

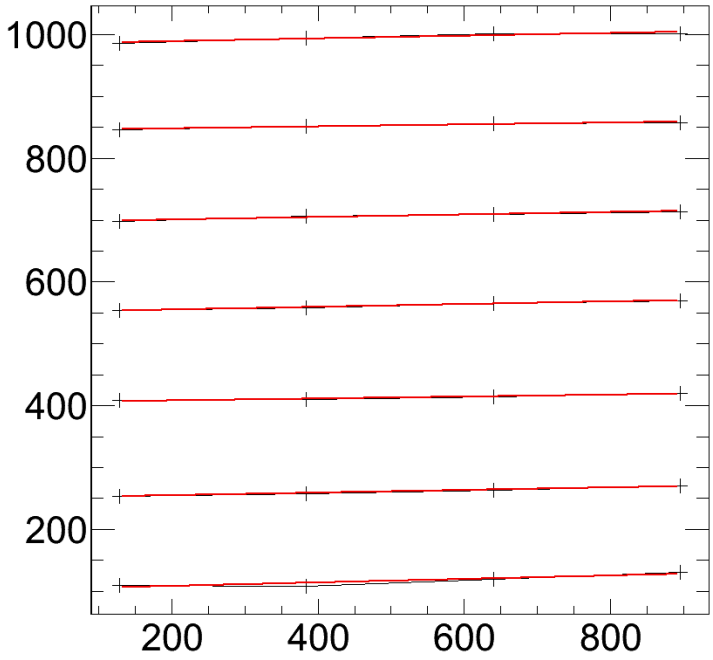

Figure 3-8: Fitted lines to spacer point locations for the bottom camera. The image is rotated to account for the handling of fits in Root. The black cross markers mark the location of the dips. The red line is the linear fit.

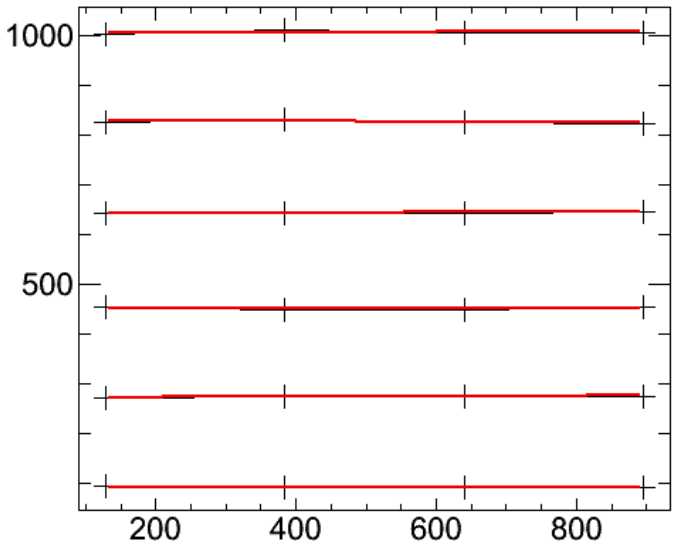

Figure 3-9: Fitted lines to spacer point locations for the top camera in the same manner as figure 3-8.

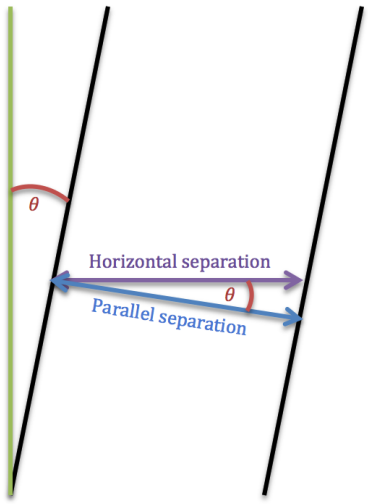

Figure 3-10: Diagram describing the spacer parallel separation calculation.

3.2.3 Diffusion

The result of the diffusion process outlined in section 1.2.4 is to cause the energy deposition in the CCD to be more spread out than it is in the original track, affecting the spatial and stopping power resolution. These are both crucial as the range of a track is a variable used for background elimination and the stopping power is used to calibrate the energy and determine the sense of a track. As such, before further calibration can be performed, the simulation diffusion must be tuned to match that seen in the detector. To achieve this the alpha source data are used.

The constant of equation 1.13 is dependent on the image pixel-binning and the camera resolution. It is a variable in the CCD simulation and is adjusted until the transverse width of generated tracks agree with the data, within error. A Gaussian is fitted to each and the fit parameters are used for comparison. Figures 3-11, 3-12 and table 3.4 show the results of the diffusion study.

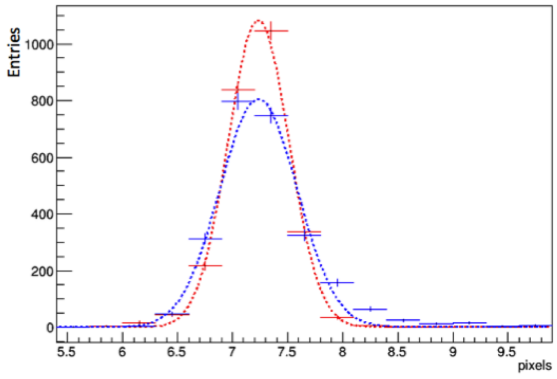

Figure 3-11: Track transverse width for top camera: blue = data, red = simulation.

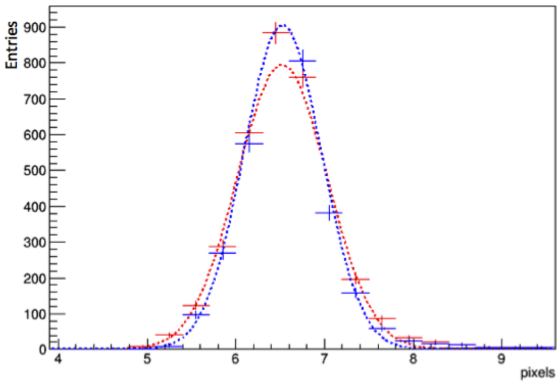

Figure 3-12: Track transverse width for bottom camera: blue = data, red = simulation.

Table 3.4: Track transverse widths resulting from diffusion study.

Camera	Data Transverse Width (σ_0) (pix)	Simulation Transverse Width (σ_0) (pix)
081264	7.235 +/- 0.023	7.236 +/- 0.012
100439	6.527 +/- 0.018	6.522 +/- 0.022

3.2.4 Energy Calibration

In order to deduce the conversion factor from CCD ADU to an energy in keV, the tracks from an ^{241}Am alpha source, of known energy, and Monte Carlo simulations are compared. The energies of the alphas are 4.44 MeV for the top camera source and 4.51 MeV for the bottom camera source. These values are lower than the 5.486 MeV of ^{241}Am alphas due to being attenuated by a covering of unknown thickness and composition. This is required such that the rate of alpha emittance is $\mathcal{O}$(1Hz), to allow single alpha imaging [59]. The energy deposition as a function of range, or stopping power, of each is plotted as a function of the track range and the simulation is adjusted until it aligns with the data, figure 3-13. Once aligned, the ratio of the two histograms is taken, providing the relative gain. This is then multiplied by the gain value set in the simulation to provide the conversion factor, figure 3-14. To ensure a correct comparison between the data and simulation, two selection criteria are used: that the track is straight and that there is only one track in the image. These criteria ensure that the stopping power, plotted as a function of range, is an accurate representation. In order to plot the stopping power, for each selected track a projection is taken, parallel to the direction of motion, of the ADU count along the length of the track (figure 3-13). These are then summed and averaged.

This procedure uses high energy alphas but the experiment is concerned with low energy nuclear recoils. In order to account for this, only the low energy part of the alpha track is used. To deduce the region to use, the path length for a particle with incident energy of 50 keV and an expected gain in the region of 10 ADU/keV is obtained using SRIM [51]. The reported length is approximately 1 mm, which covers two bins in either camera. Therefore the dE/dx range of interest is 250 ADU per bin. For comparison, with the conversion factors shown in table 3.3, the 500 (560) pixel range of these alphas is equivalent to 86.0 (76.7) mm for the bottom (top) camera. Values higher than this are included in a fit to determine the gain value

to use. Another difference between these two interactions is that the alpha deposits almost all of its energy into ionisation whereas the nuclear recoils have a quenching factor - see section 1.2.4. This should not impact the calibration as it is and ADU to keVee conversion that is produced, the quenching factor can then be applied to obtain the energy in keVnr. The result of this procedure is a conversion factor of 9.74 $\pm$ 0.14 and 10.38 $\pm$ 0.14 ADU/keVee for the top and bottom camera respectively.

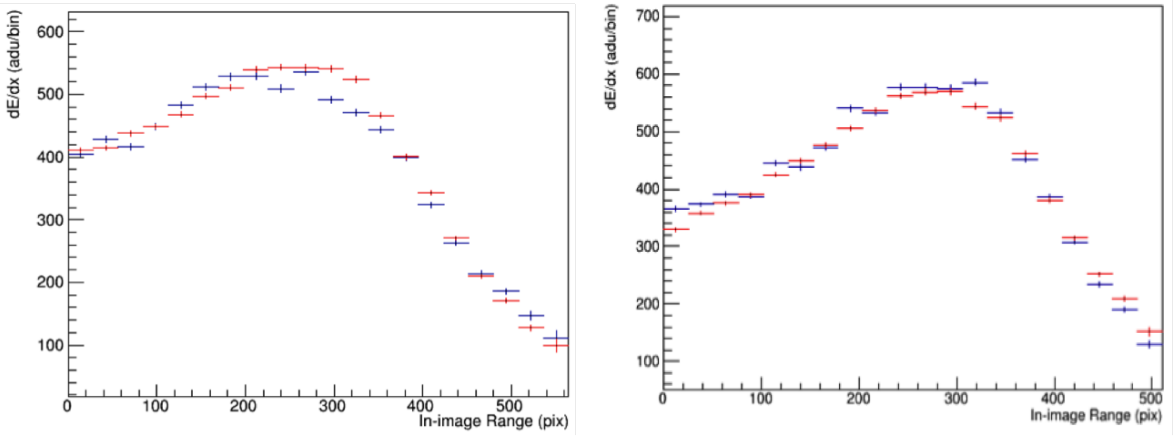

(a) a (b) b

Figure 3-13: Track deposited energy for a) top camera, b) bottom camera. Red is simulation, blue is data. The abscissa is in-image range as the source does not lie directly at the image edge.

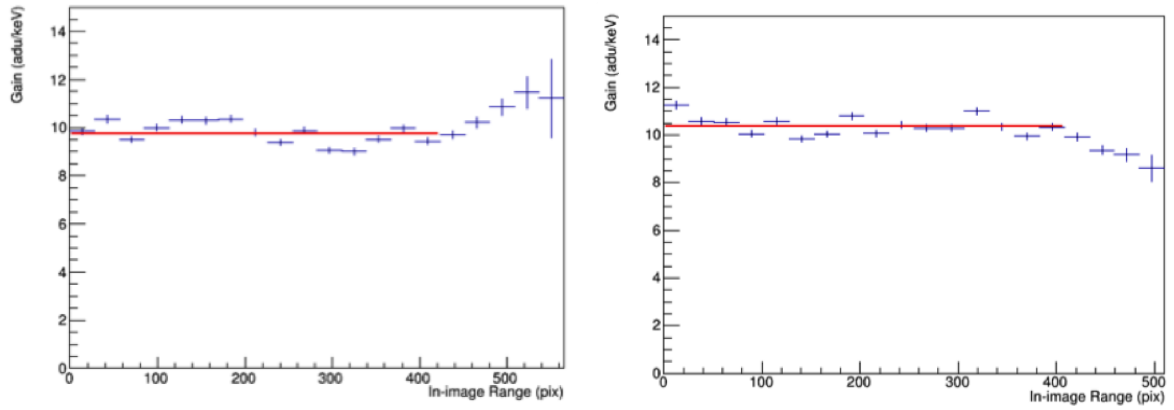

Figure 3-14: The ratio of simulation to data for histograms in figure 3-13 for a) top camera b) bottom camera. The red line is a fit to the bin values, used to produce a gain conversion factor. The features at high range values show some stopping-power dependence to the gain. These bins are excluded from the fit as they fall outside the range of interest for low energy nuclear recoils.

3.2.5 Range and Energy Correction

In order to account for any bias that can be introduced in the reconstruction of tracks, a correction is applied, again using simulation generated tracks. This correction accounts for effects such as focussing of the camera and the cluster reconstruction algorithm. The process is the same for both the range and the energy and involves comparing the known values for simulated events with the values reconstructed by the analysis software.

With the correct calibration parameters, as deduced above, simulated neutron events are generated in the energy region of 0 to 500 keVnr and at random and isotropic x, y and z positions. Once the events have been passed through the image cleaning and track finding software, the reconstructed values and the truth values are plotted against each other in a two dimensional histogram. To ensure the comparison is valid, the x and y locations of the true cluster are compared with the reconstructed cluster. Once plotted, the difference of x bin center to y bin centers is plotted for each energy (range) bin and the 1 (1.5) sigma outliers removed. Finally the profile of the 2D histogram is taken and fitted to a line, as shown in figures 3-15 and 3-16. These corrections are only applied at the point of cutting on range and energy. This means all other cuts are calibrated and applied based on the original reconstructed energy - avoiding the need for this correction to be applied multiple times. The reconstructed values can be seen to worsen at lower energies and ranges. This is a demonstration of the lower limit of recronstruction which is accounted for by a reconstruction efficiency which is shown in section 3.3.6.

Table 3.5: Reconstruction-bias fits for for energy.

Camera	Energy Fit
081264	$E_{Truth} = (-1.96 \pm 0.17) \times 10^1 + (1.0944 \pm 0.0010) E_{rec}$
100439	$E_{Truth} = (-2.94 \pm 0.18) \times 10^1 + (1.1055 \pm 0.0010) E_{rec}$

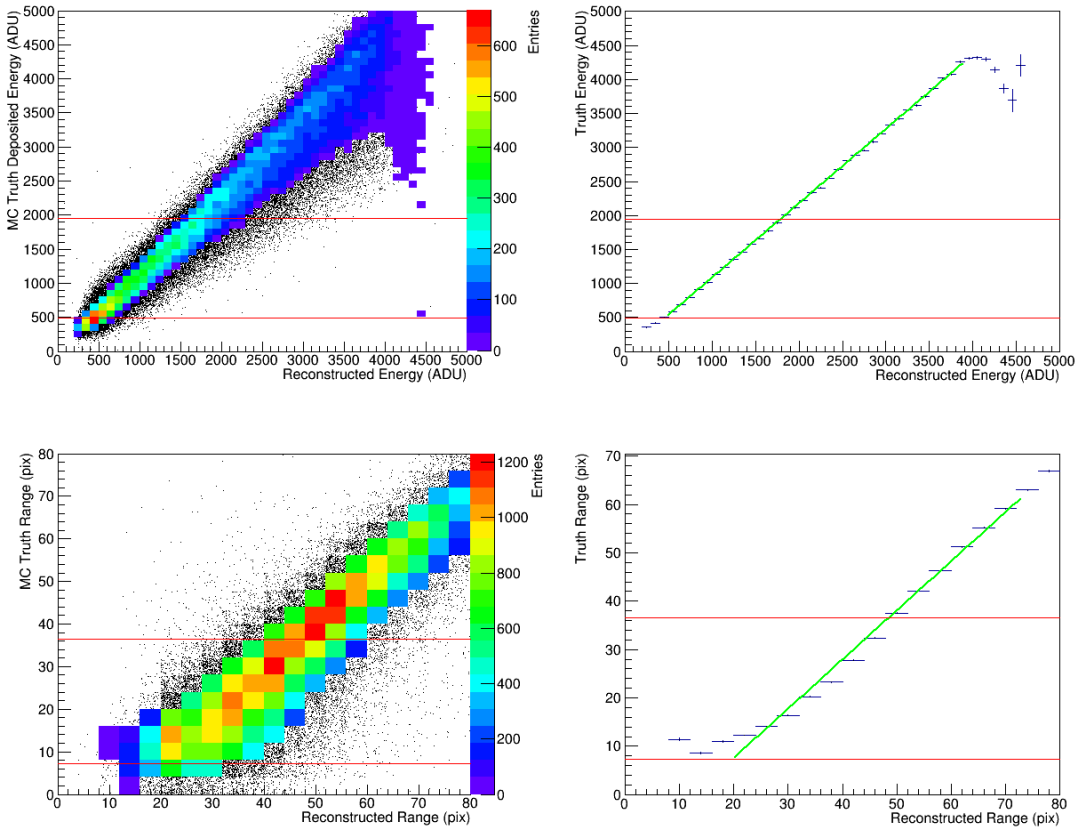

Figure 3-15: Correction histograms for the top camera. The left hand side shows the 2D histograms and the right shows their profile. The top histograms are for the energy correction and the bottom two are for the range. The black dot markers are the full distribution, the coloured boxes are the distributions with outliers removed. The red lines show the typical range of values for nuclear recoils with energies of 50 and 200 keVee. At the lower energies, the linear relationship between truth and reconstructed values tapers off. This shows the reduced ability to accurately reconstruct at such low energies.

Table 3.6: Reconstruction-bias fits for range.

Camera	Range Fit
081264	$R_{Truth} = (-12.962 \pm 0.075) + (1.0183 \pm 0.0016) R_{rec}$
100439	$R_{Truth} = (-7.342 \pm 0.098) + (1.0155 \pm 0.0025) R_{rec}$

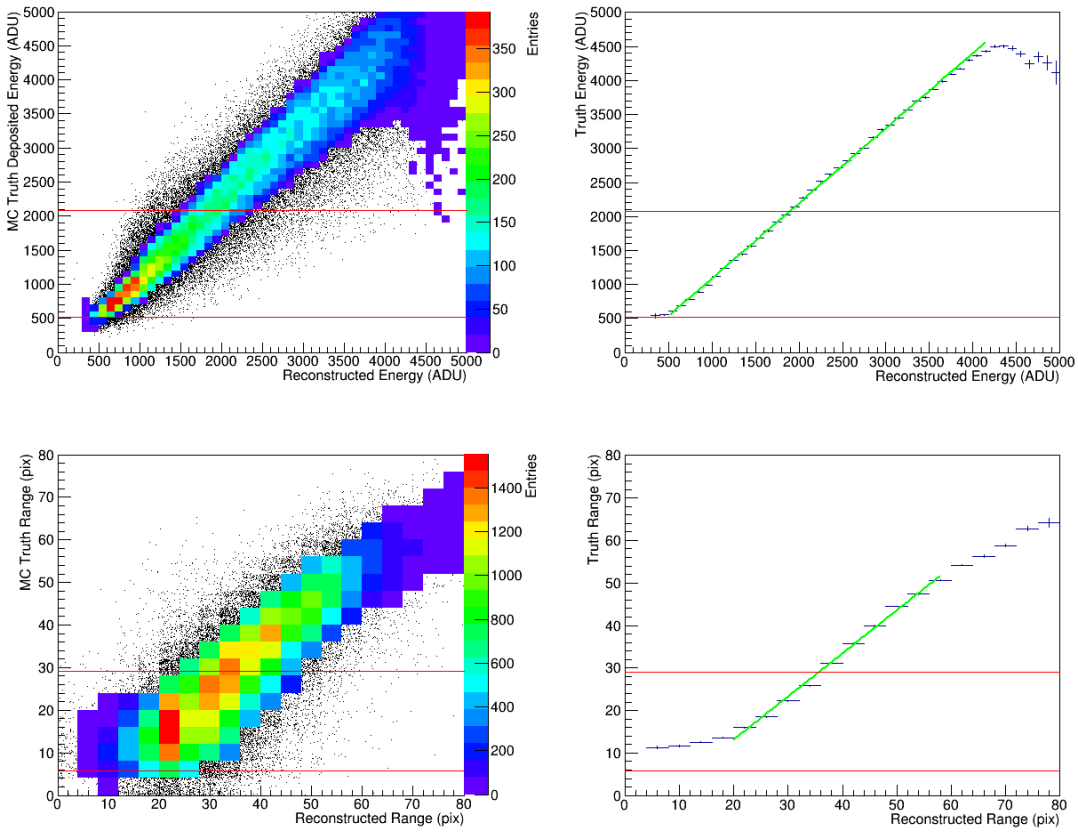

Figure 3-16: Correction histograms for the bottom camera. The left hand side shows the 2D histograms and the right shows their profile. The top histograms are for the energy correction and the bottom two are for the range. The black dot markers are the full distribution, the coloured boxes are the distributions with outliers removed. The red lines show the typical range of values for nuclear recoils with energies of 50 and 200 keVee. At the lower energies, the linear relationship between truth and reconstructed values tapers off. This shows the reduced ability to accurately reconstruct at such low energies.

3.2.6 Charge-Light Match

Prior to this analysis, the charge readout data had not been used for the 10L detector. For this analysis it was decided to use a simple charge-light energy match. This is a powerful method for discriminating between real events and CCD artefacts as a real event will have a charge event with equivalent energy to that of the CCD. The method used is to plot the peak voltage value from the anode of the amplification region versus the energy gotten from the CCD in a two-dimensional histogram. The profile of this histogram is then taken, whereby a two-dimensional histogram is effectively flattened to one-dimensional, with the ordinate values equal to the mean of the distribution. This profile is then used to produce a linear fit which is used to compare results in the data and reject events that do not have a matching CCD and charge-readout energy, within error.

To calibrate this cut, neutron source data are used, after being cleaned with some basic data quality cuts. For the CCD readout, alpha events and poorly reconstructed tracks are rejected. It is also necessary to place some cuts on the charge readout as there are often multiple pulses recorded over the duration of a CCD exposure and there is possible cross-talk between the back-to-back amplification regions. Here the work done by a previous collaborator was drawn upon as a guide [50].

There are two parts to the charge readout - the anode integrated charge and the mesh fast-response pulse. The mesh pulse rise time, in the region of 25%-75% of the full rise, is found to be useful in selecting nuclear recoils from electronic recoils [50]. The cut value for this is deduced by plotting all triggers and selecting the region of nuclear recoils, see images in figure 3-17. The sources of all the populations seen have not been deduced and was not seen in [50]. They are possibly due to different track orientations. For example, an electronic track travelling in the x-y plane will be read out in a short time, since the extent in the drift direction is small, however, a track travelling in the z direction will be read out over a longer time period, for the

opposite reason. Both will have low energy but will vary in pulse rise time. Some events are also likely due to cross-talk, where a signal in one amplification region induces a signal in the other. This effect would not have been seen in [50] as there was only one amplification region.

The events of interest lie in the region of approximately $0.01 <$ anode (V) < 0.18 and $10\times 10^{-6} <$ rise time (s) $< 16 \times 10^{-6}$, corresponding with the region for nuclear recoils observed, simulated using an alpha source, in [50]. This allows a large portion of the neutrons through whilst rejecting most other pulses. Then, to remove the possibility that a value used is due to cross talk from the other amplification region, a simple check was made to be sure that the voltage peak is greater than the greatest peak value from the other side.

With all these cuts in place, the charge value of the remaining events are plotted against the CCD energy in ADU, figure 3-18. Performing the calibration on the raw voltage and ADU value, rather than as a function of recoil energy, avoids the need for conversions, which can introduce additional errors. A fit is then applied to the profile of the resulting histogram in the limited region away from the vertical band. The fit implements the Least Trimmed Squares Robust method [60], whereby a percentage of the data is chosen to perform the fit, selecting those that minimise the sum of squared residuals. The percent of data to be used was set to 70 %. The reason for this value being chosen was that the fitted line at this percentage had the closest-to-zero offset value, as would be expected: no energy in the CCD should correspond to no energy in the charge readout. The reason for using this technique is that there is the possibility for multiple charge events to pass the selection criteria and there is no current technique for knowing which is the matching pulse. It is therefore reasonable to use only a percentage of the passing data. Using this fit, a conservative cut was made on events lying more than one sigma away on the linear offset, figure 3-18.

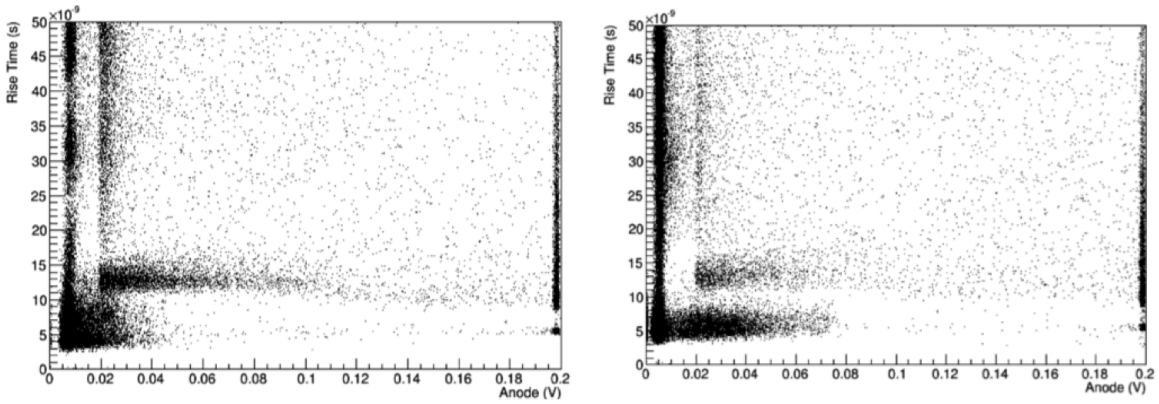

Figure 3-17: Ground-mesh pulse rise-time vs. anode peak voltage plotted for all charge data. The two L-shaped structures seen in both histograms are attributed to cross-talk between the back to back amplification regions. The voltage window is configured for low energy events and so the vertical band at 0.2 V is due to high energy events. The vertical band at low voltages is attributed to electron recoils that have a longer extent and so can take more time to deposit their energy. The events of interest, neutron recoils, lie in the region of approximately $0.01 <$ anode (V) < 0.18 and $10\times10^{-6} <$ rise time (s) $< 16\times10^{-6}$.

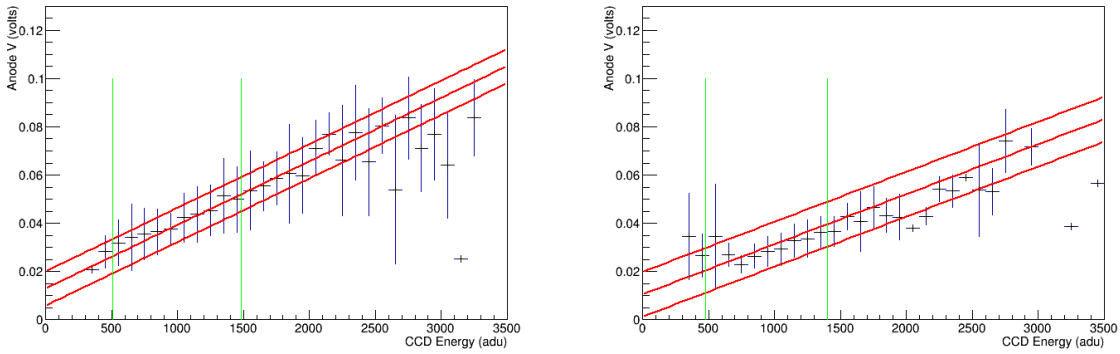

Figure 3-18: Charge-light calibration fits. The red lines indicate the fit with $\pm$ 1σ. The green lines indicate the 80-200 keVnr energy region.

Table 3.7: Linear fit results for charge readout voltage as a function of CCD energy in ADU.

Camera	Charge-Light Fit
081264	Charge (V) = 2.073×10^{-5}E$_{ccd}$(ADU) + 0.010
100439	Charge (V) = $2.064\times10^{-5}E_{ccd}$(ADU) + 0.013

3.3 Cuts

As mentioned, part of the analysis process is to pass the data through a number of cuts that are used to remove background events, discussed in section 2.2, whilst preserving signal-like events. This section discusses the cuts made and how they were configured for the underground analysis. The differences in detector configuration for WR6 compared to WR5 require that the cuts be re-tuned, such that they are tailored to the data set. The additional noise (bottom camera only), along with the increased exposure time, required that the cuts related to camera artefacts be updated. The increased exposure time affects the spark cut. Also, the analysis of the data of WR5 resulted in a higher rate of passing events than found in the surface run. This was considered to be due to RBIs and as such the RBI cut was further developed. First the cuts are discussed then their efficiencies are presented at the end of the section.

3.3.1 Spark Cuts

When a spark occurs in the amplification region, a reduction in the gain is observed, which is due to the recovery time of the high voltage supply. For the 10L detector, this lasts for approximately five seconds. Due to this, the energy of any tracks that are recorded in this time period cannot be relied upon. As such, any data that were taken during or following a spark was discarded. For the one second exposure of WR5, this is achieved by ignoring the five images that immediately follow a spark. For the five second exposure, it is only necessary to reject one image. There is a simple cut on the charge that requires that there are fewer than twenty-five recorded pulses, or triggers, contained in the event. This stems from the observation that, when sparking, many more pulses are recorded, thus events with a large number of triggers are likely to be due to sparks.

Two image cuts are applied to remove sparks. The first is the ratio of the image-

mean to the previous image-mean. The second is the difference between the image mean and the mean of the overscan region[8]. To determine the values used for rejection, a selection of sparks were identified by eye and the values compared to non-spark images. Figures 3-19 and 3-20 show the results of this.

The cut value for each of these is listed in table 3.8. The values were deduced by maximising the signal efficiency whilst maintaining a low background efficiency. The efficiencies for each cut is alsogiven in table 3.8. An additional consideration for sparking is whether it may have occurred just prior to a data set being taken. With no image to act as a reference, the cut values cannot be deduced. As such, the initial image of each run is discarded as a conservative measure.

[8]The overscan region is an array of pixels which are covered and thus shielded from illumination. These pixels are therefore useful as a baseline comparison.

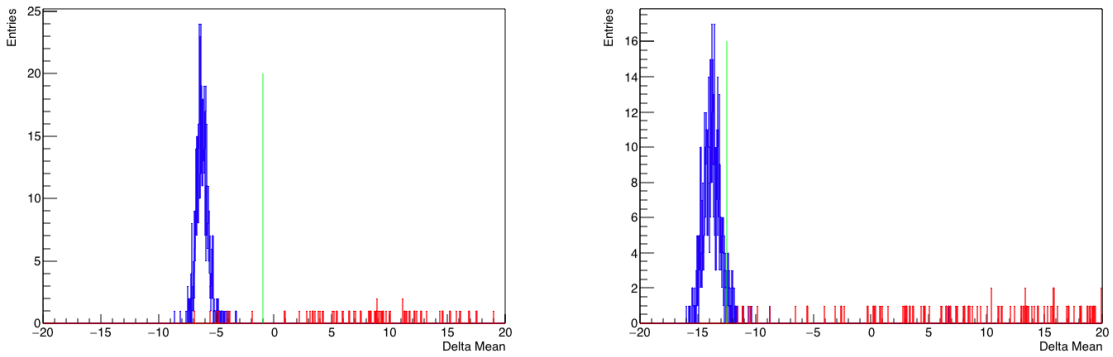

Figure 3-19: Delta-mean variable. Left is the top camera, right is the bottom camera. Red: sparks, blue: non-sparks. The green line shows where the cut was placed, all images with the variable greater than this value were discarded.

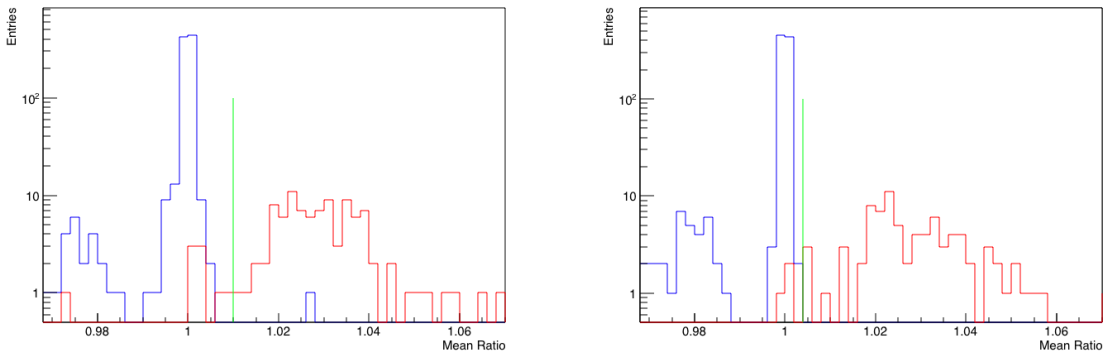

Figure 3-20: Mean-ratio. Left is the top camera, right is the bottom camera. Red: sparks, blue: non-sparks. The green line shows where the cut was placed.

Table 3.8: Spark cut values. Efficiencies are listed for sparks first then non-sparks. The matching efficiencies for the top camera are not erroneous, they have been verified.

Camera	Mean Ratio	Efficiencies %	Delta Mean	Efficiencies %
Top	1.010	6.4 / 99.9	-1.0	6.4 / 99.9
Bottom	1.004	5.0 / 99.9	-12.5	0.7 / 94.0

3.3.2 Charge Cut

The charge cut is a very basic one and relies solely on the number of charge triggers acquired during the exposure time of an event. The rejection of events that contain a high number of triggers is an additional means of removing images that might contain a spark. This is because sparks have been seen to cause a high number of recorded readouts - owing to the high intensity electrical discharge induced. The value of this cut was unchanged from the previous analyses and was set to accept events with twenty-five triggers or fewer. Twenty-five might seem like a large number of triggers, however the charge readout has a lower energy threshold, $\mathcal{O}$(keV), due to few to no losses of the generated signal and due to the nature of the readout electronics, a pulse is recorded for both TPCs regardless of the origin of the event that caused the trigger.

3.3.3 Energy Match

As described in section 3.2.6, a physical event occurring within the fiducial volume of the detector would be expected to produce a signal in both the CCD and charge readouts (assuming the energy is above the detection threshold $\sim$ 50 keV). This can be used to remove a significant fraction of the backgrounds that are recorded. To do this, the CCD integrated ADU count is evaluated against the fit shown in figure 3-18. The event is accepted if the evaluated value lies within the ± 1 σ acceptance region, also shown in figure 3-18.

3.3.4 Worm Cut

The worm cut is based on a multivariate analysis (MVA), the training for which makes use of the ROOT implementation of TMVA [61]. There are many different types of MVA and the scope of the topic is too broad for inclusion here, however the underlying premise is to use a set of parameters from known background and signal events to produce a single parameter, the distribution of this parameter, for each event type, can then be used to discriminate between the two. One of the simplest implementations of this is a linear discriminant, in which each parameter is given a weight. The sum of the weighted variables then provides the discriminating parameter. The process requires training on the known samples to produce the two distributions which are most separated, determined by their mean values. The more separated the distributions, the stronger the power of discrimination. The following variables were used for WR5 and WR6: energy, range, cluster RMS, maximum pixel value, number of neighbours to maximum pixel with energy above cluster-finding threshold, number of pixels in the cluster and number of pixels in the reduced cluster. The background training data come from the cosmic data sets, however, with the new image cleaning method and smaller data set of WR6, only very few worm events were found, therefore a charge parameter was added. The addition of the charge parameter is useful in this situation as it is a powerful discriminant between signal and background, since only physical events occurring within the fiducial volume should produce both. See section 2.4.2 for further information on how the charge assignment was performed. The signal events used for the MVA are simulated WIMPs, generated using the expected energy distribution, with an energy range of 0-600 keVnr and all the relevant chamber/camera properties for the top and bottom TPCs e.g. gain, image width and pressure. To get the best results, three different MVA methods were tested: a linear discriminant, boosted decision tree (BDT) and rule fit, see [61] for more on these methods. The response value chosen to set as the cut value was

decided by that which reported zero percent background efficiency. The best method was found to be different for each camera, which would ordinarily be undesirable as we would expect them to exhibit similar behaviour. However, as has been seen, the noise in the bottom camera differs to that of the top, so it is not unreasonable to expect this response. The results of the training can be seen in figures 3-21 to 3-32.

Table 3.9 summarises the cut values and resultant signal efficiencies for each method used. The method chosen was that which maximised the signal efficiency. For the top camera this was the linear discriminant, for the bottom camera this was the BDT. Referring to the response plot for the BDT discriminant used on the bottom camera, figure 3-27, there is a large spike in the training background data response at around -0.3. This might indicate overtraining, whereby the response has been over-tuned to fit the training data, however it is the tested background event responses that are used to set the cut value. The distribution of the tested events cover the trained region and extend further, ensuring that the cut value chosen is representative and conservative.

In WR5 a linear discriminant was used for botoh cameras and had 1.5 % and 0.5 % background efficiency and 94 % and 88.7 % signal efficiency (bottom and top respectively). The cut values and efficiencies for WR6 are in table 3.9.

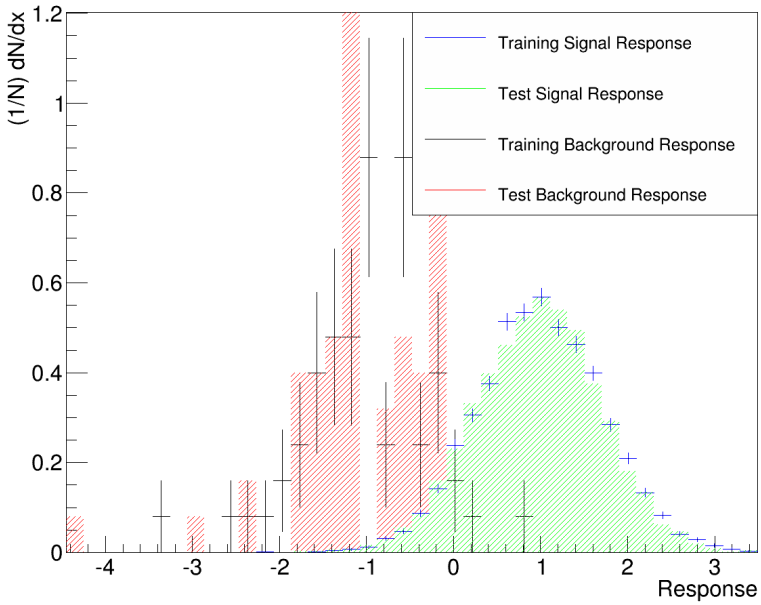

Figure 3-21: Linear response plot for the top camera.

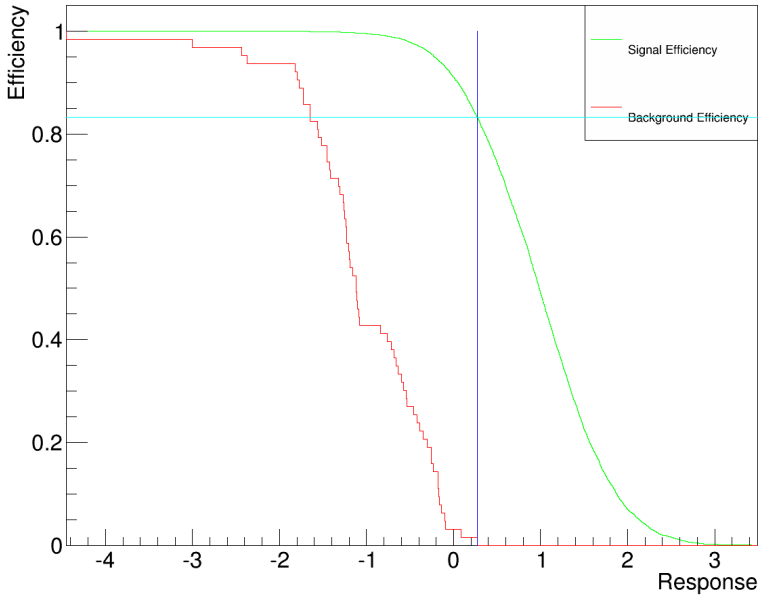

Figure 3-22: Efficiency as a function of linear response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

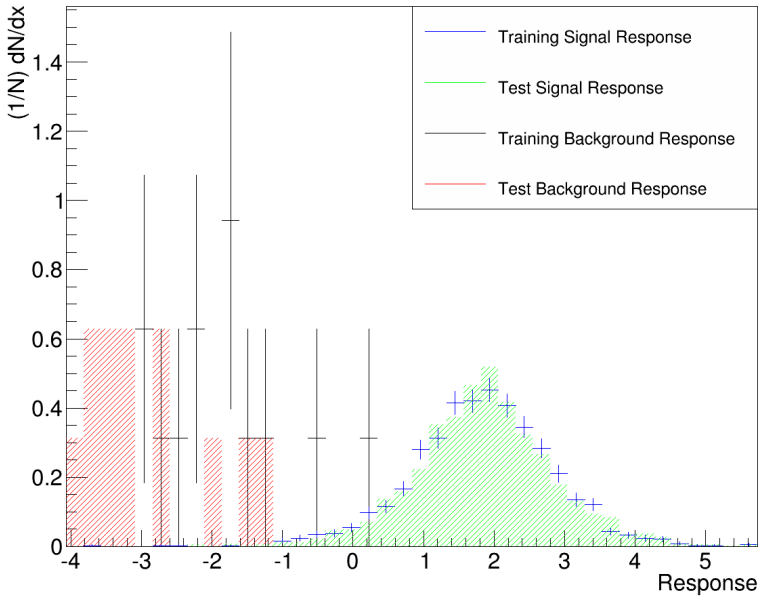

Figure 3-23: Linear response plot for the bottom camera.

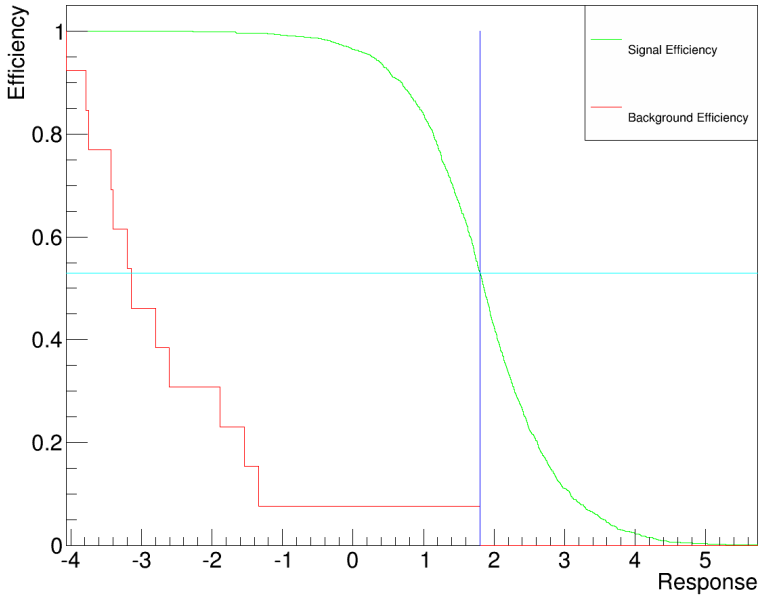

Figure 3-24: Efficiency as a function of linear response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

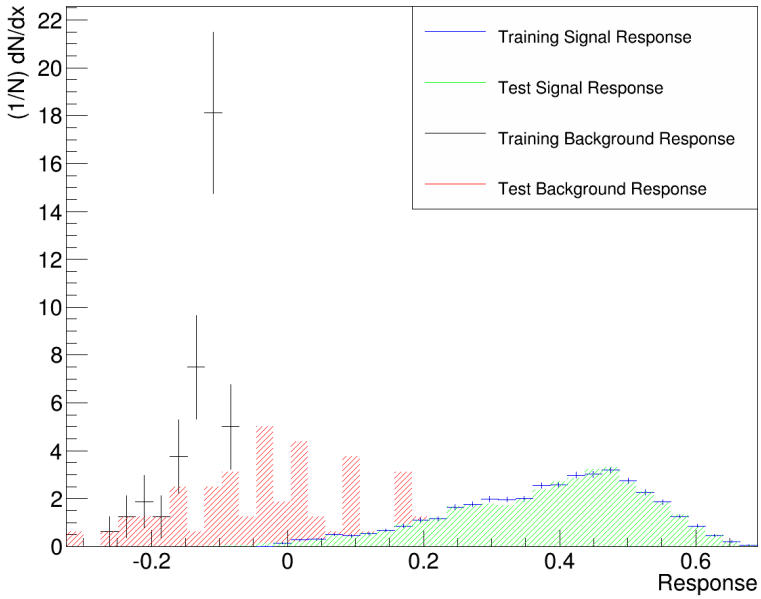

Figure 3-25: BDT response plot for the top camera.

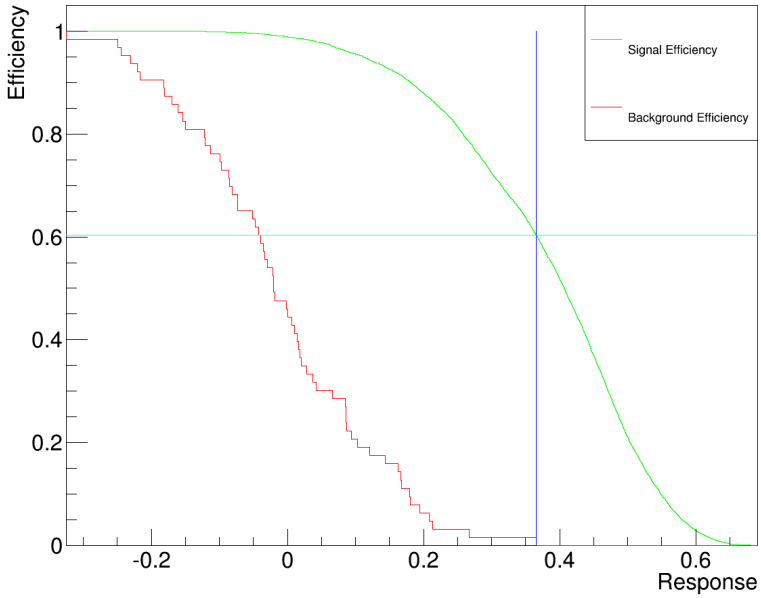

Figure 3-26: Efficiency as a function of BDT response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

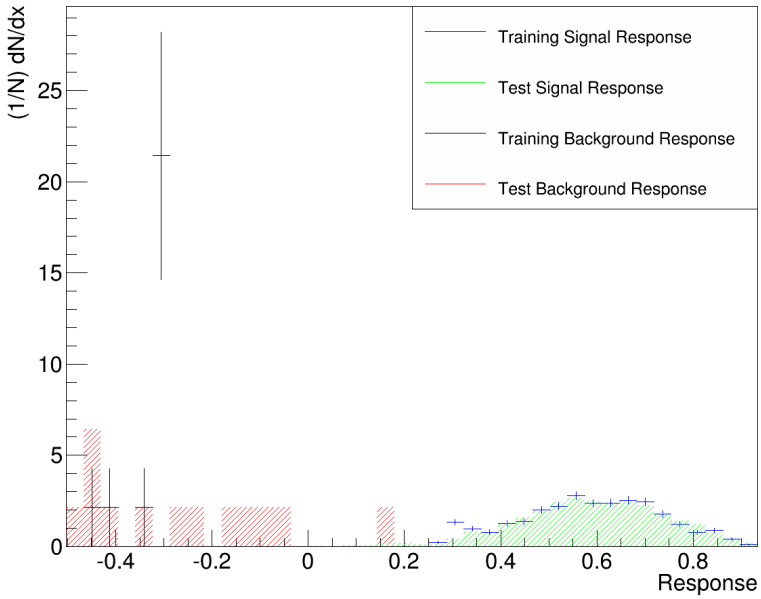

Figure 3-27: BDT response plot for the bottom camera.

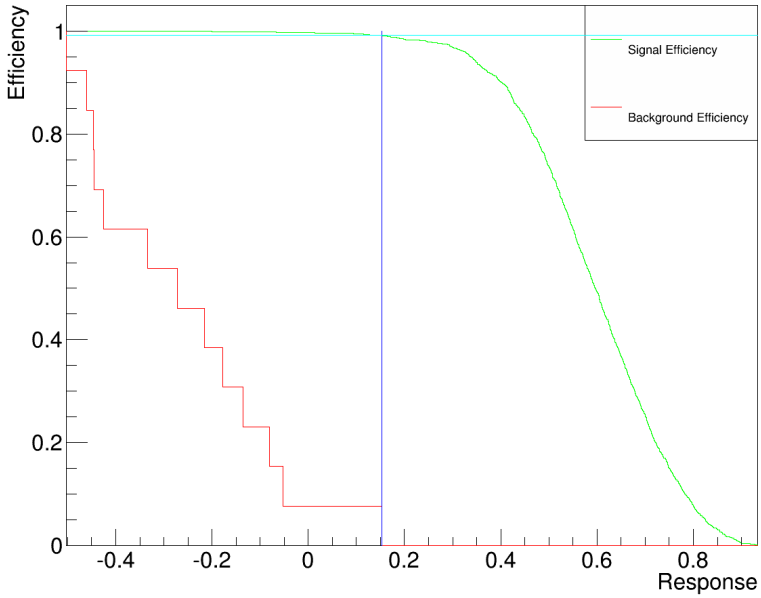

Figure 3-28: Efficiency as a function of BDT response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

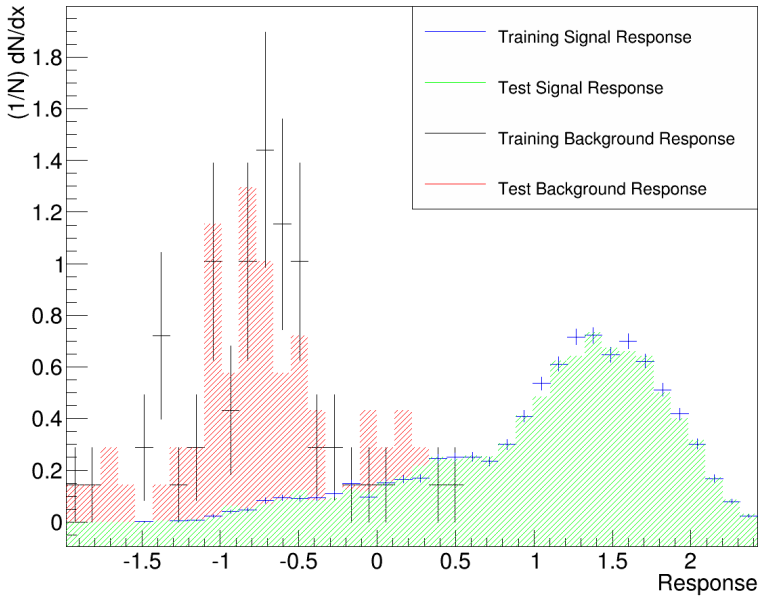

Figure 3-29: Rule fit response plot for the top camera.

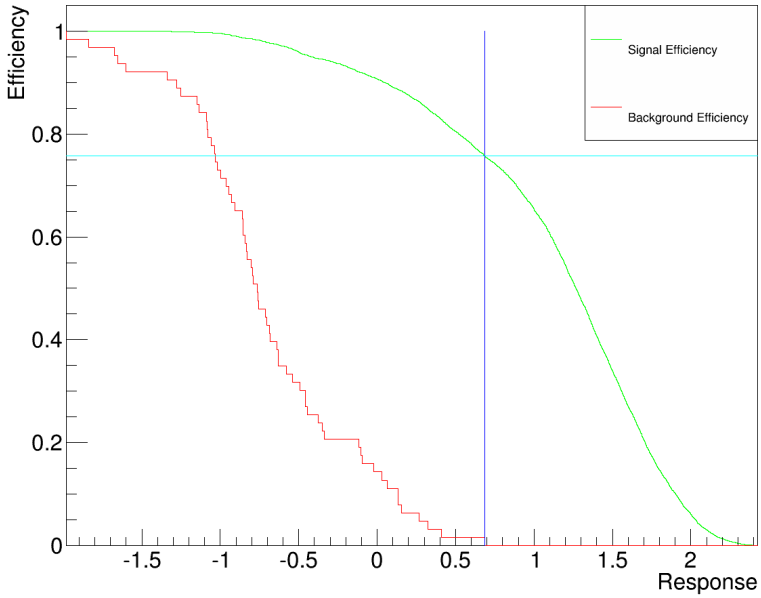

Figure 3-30: Efficiency as a function of rule fit response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

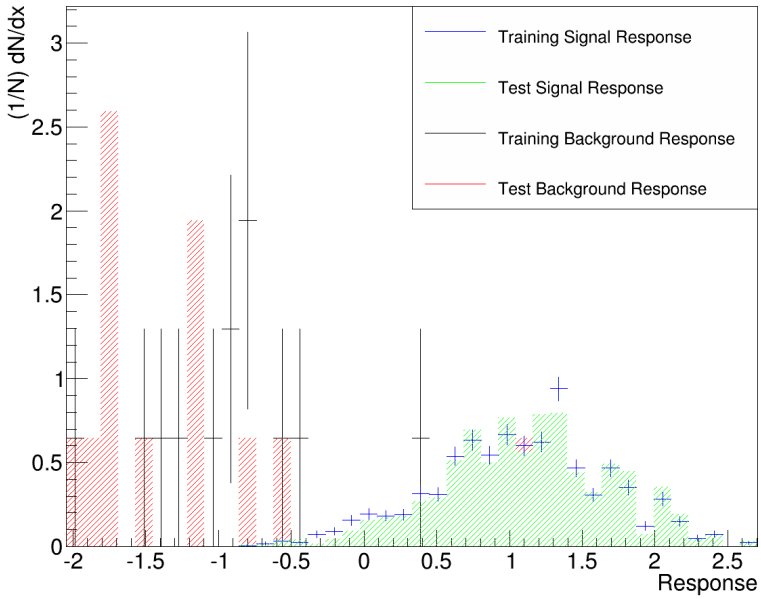

Figure 3-31: Rule fit response plot for the bottom camera.

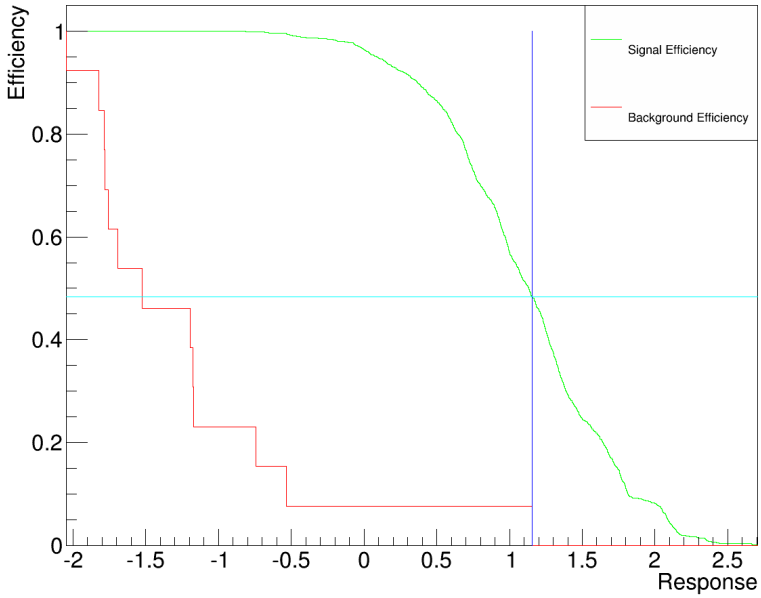

Figure 3-32: Efficiency as a function of rule fit response value. The vertical blue line indicates the zero background efficiency response value. The horizontal cyan line indicates the corresponding signal efficiency.

Table 3.9: Worm MVA cut values and signal efficiencies. The chosen methods are in bold.

Camera	Method	Cut Value	Signal Efficiency (%)
Top	**Linear**	**0.28**	**83**
Bottom	Linear	1.8	53
Top	BDT	0.37	60
Bottom	**BDT**	**0.16**	**99**
Top	Rule Fit	0.69	76
Bottom	Rule Fit	1.16	48

3.3.5 RBI Cut

The original cut used in WR5 looks at all tracks from a run and checks if there are more than two clusters which overlap in position and are within four bins of a saturated pixel. If they do, they are discarded. To improve on this for WR6, the overlap check was extended to look at the previous run as well. To determine how far back into the run to search, the intensity decay of saturated pixels was studied. For this, the intensity of any pixels tagged as having saturated during the run was plotted as a function of the number of events. This was also done for non-saturated pixels for a comparison to the baseline. The distributions for these can be seen in figure 3-33

In the bottom camera the intensity does not quite make it back to the baseline and in the top camera it overshoots. For the bottom camera, this is likely a consequence of the additional noise. For the top camera this may be an indication that the cleaning method is too severe. The value used for the cut is the point at which the intensity crosses the baseline for the top camera. Since both cameras are the same make and model, the RBI behaviour would be expected to be very similar and unaffected by the in-image noise. The noise is also suspected to be read-out related and not chip-related. The event number at which this crossover occurs is 375. It is clear to see that if a spark occurs near the end of a run, the RBI could persist into the subsequent run. Therefore, the new cut ensures that when looking back 375 events for overlapping

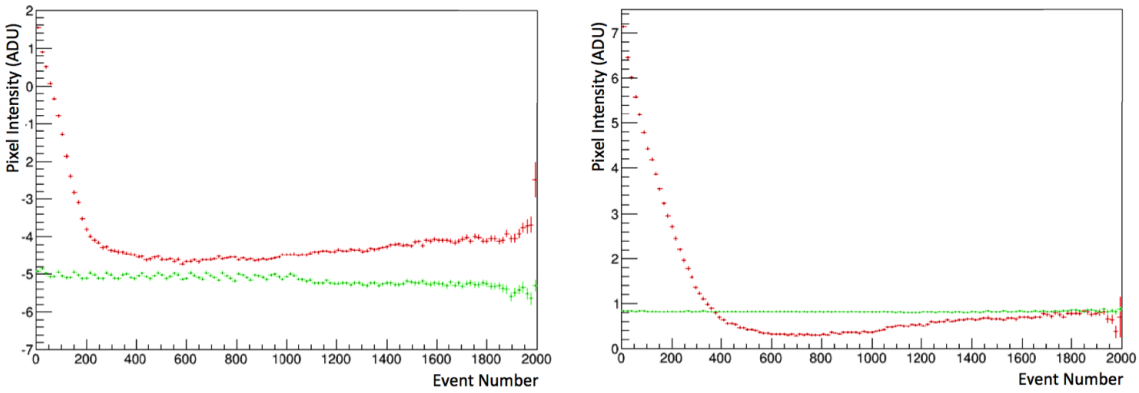

Figure 3-33: Pixel intensity dropoff vs. event number, where each data set contains one thousand events. Red shows the saturated pixels and green the non-saturated. Left is for the bottom camera and right is the top.

clusters, the previous run events are also considered.

3.3.6 Efficiencies

The efficiency of each cut is calculated by using simulation-generated nuclear recoils, with the expected WIMP distribution and energy range. The exception to this is the spark trigger cut (charge cut), which uses neutron source data. The number of events that pass each cut, out of the total number of events generated, is plotted as a function of energy. The results of these are shown in figures 3-34 and 3-35. Most of these cuts have been discussed, those that have not are: region/edge - removes tracks at the image edge as the energy of these tracks cannot be reliably reconstructed; spacer cut - removes tracks that are near spacers (which separate the ground mesh from the anode plate in the amplification region) for the same reason as the region/edge cut; long tracks - removes events that contain a long track ($>$ 5 mm) to remove RPR events; reconstruction - accounts for the ability of the cluster-finding algorithm to reconstruct tracks. The final efficiency histogram is then used with the exposure of the run to produce an exposure weighted efficiency plot, as seen in figure 3-36. The reconstruction efficiency is calculated using a separate procedure and so is seen in the plot as an independent distribution, but the final overall efficiency is the accumulation of all cuts. There is a hard edge for the bottom camera charge efficiency and a gap for the top camera charge efficiency. Thiese arise due to there being no data available at the energies in question. To produce the efficiency-weighted exposure plot of figure 3-36, the fiducial mass of each camera is multiplied by each bin of the energy-dependent efficiency. This is then used in the limit setting procedure as outlined in section 3.4.2.

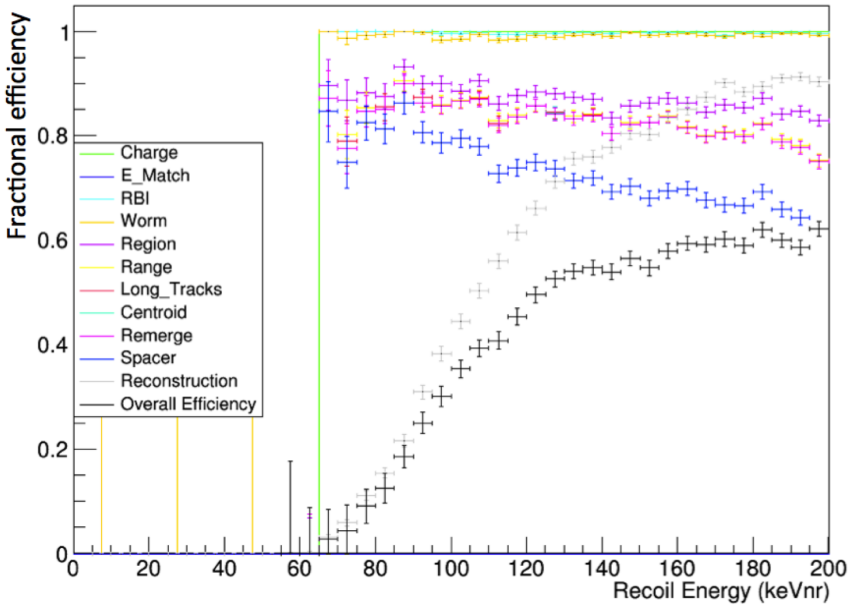

Figure 3-34: Cut efficiencies for the bottom camera.

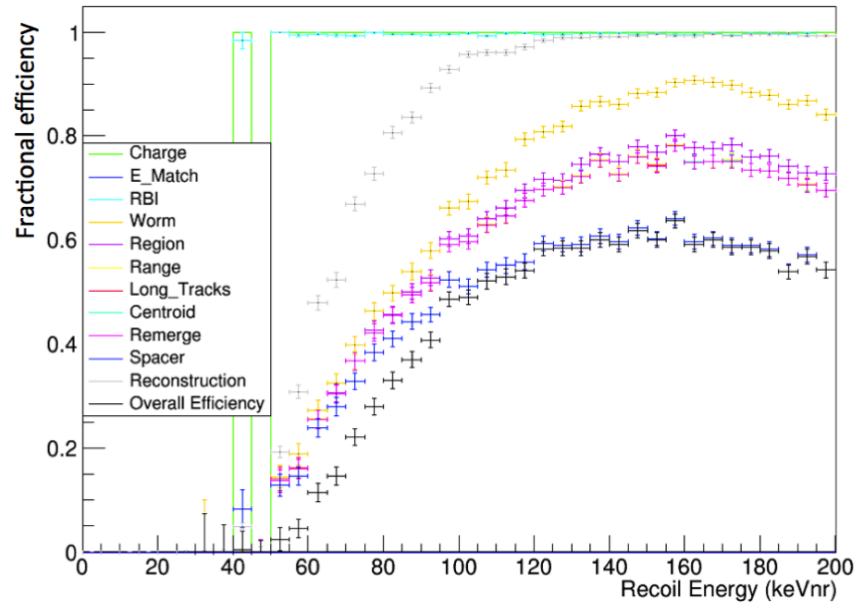

Figure 3-35: Cut efficiencies for the top camera.

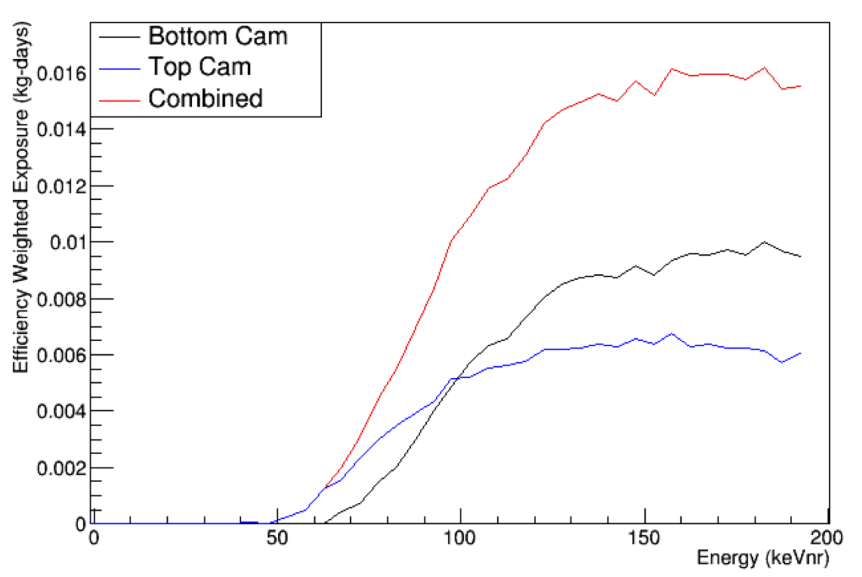

Figure 3-36: Efficiency-weighted exposure as a function of energy. Colours indicate a different camera or their combination

3.4 Analysis

The DMTPC experiment is operated in observation mode, whereby a certain exposure time is used during which the CCD shutter is open and all pulses in the charge readout are stored. Because of this a first pass analysis is performed which both cleans the data and tries to detect potential signals. In the case of the CCD images, this consists of searching the obtained images and constructing clusters of pixels that would represent the two-dimensional projected energy of an ionisation track. This section describes this pre-processing

3.4.1 cleanSkim

cleanSkim is the name given to the image cleaning and cluster finding software. It is well detailed in references [55, 57]. Firstly, the bias images undergo a cleaning process, replacing the value of high intensity pixels with the mean of the surrounding pixels. These high-intensity pixels are caused by fluctuations in the CCD and are defined as being 10% greater than the image mean. This cleaning is performed three times for

the bias frame. Then, each of the event images is checked against the spark cut criteria (see chapter 3). If the image is identified as having a spark, any saturated pixels are stored in an array for later referencing (see section 3.3.5) and no cluster finding is performed. If the image passes the spark criteria, it too is cleaned, using the median value of surrounding pixels. The median is used here, instead of the average, since it is more robust against hot pixels that might be in the surrounding pixels. Crucially, for event images, it is necessary to check the surrounding pixels for increased intensity, otherwise all candidate tracks will be cleaned away. If the surrounding pixels have a higher intensity, the pixel remains unmodified. With the cleaning completed, the mean value of the bias image is subtracted from each pixel of the event image, to centre the mean around zero. Finally the cluster finding algorithm is run to build up clusters of pixels which become the candidate events of the data. The cluster finding is based on a seeding algorithm which is run on a low-pass-filtered version of the cleaned image. The seeding algorithm works in the following way. Firstly, the highest intensity pixel of the image is selected as the seed. Then the surrounding pixels are added to the cluster if they are above the threshold value. The threshold value is decreased by 75 % on each iteration of surrounding pixels. A reduced cluster is also generated, using the pixels of the filtered cluster, but using only those that are 2 σ above the noise of the unfiltered image. These clusters become the candidate events for which a selection of relevant parameters are deduced, for example:

1. Energy: this is the integral of all pixels in the cluster
2. Range: deduced by the separation between the furthermost two pixels
3. Centroid: calculated by weighting of the pixel intensities
4. X, Y location: the location in the image of the track centroid
5. Phi: the angle of the track made with respect to the x-axis of the image

The charge waveform data are also processed at this stage, determining values such as the pulse height and baseline value. The analysis of the charge data is still in production, see [50, 62] for further reading. Due to this the charge data are used in a limited capacity in the analysis of this thesis. It is used in the generation of a simulated charge signal, as shown in section 4.2. It is also used in the calibration process, for producing an energy match between light seen in the CCD image and the charge read-out, see section 3.2.6.

3.4.2 Limit Setting

With the preprocessing complete and the events passed through the selection cuts. A list of events passing each cut can then be generated. The final number of events passing is then used to set a limit on the WIMP-proton interaction cross-section as a function of WIMP mass, as described in section 1.3.3. As part of this project, the software that sets the limit was improved upon. Previously, an energy-averaged value for the detector efficiency and nuclear form-factor were used, deduced for a single WIMP mass value. This was modified such that both values are calculated at each energy and mass under consideration. Additionally, the combination of two experiments was updated. This is relevant to DMTPC as we have multiple cameras which can be treated as the combination of multiple experimental results. Originally, the rates from each camera were simply combined and the spectrum-averaged efficiency used. This neglects the differing, energy-dependent, efficiency of each camera. The original calculation of the cross-section was based on equation 3.1

$$\sigma_{SD} = \frac{\sigma_0 F_C}{\bar{\epsilon}\bar{F}^2 F_U} \tag{3.1}$$

where the parameters represent those given in section 1.3.3 and $\bar{F}^2$ and $\bar{\epsilon}$ are the *spectrum-averaged* nuclear form factor and detector efficiency respectively (equations

3.2 and 3.3).

$$\bar{F^2} = \frac{\int_{E_{min}}^{E_{max}} \frac{dR}{dE} F^2(E)}{\int_{E_{min}}^{E_{max}} \frac{dR}{dE}}, \tag{3.2}$$

$$\bar{\epsilon} = \frac{\int_{E_{min}}^{E_{max}} \frac{dR}{dE} \times \epsilon(E) dE}{\int_{E_{min}}^{E_{max}} \frac{dR}{dE} dE}, \tag{3.3}$$

This helps to reduce calculation but also reduces the accuracy of the result: since to average over an energy range, the energy distribution is required, thus a WIMP mass must be assumed. This is clearly inaccurate since the limit result is reported as a function of WIMP mass, as mentioned in section 1.3.3. To improve on this, the calculation of equation 1.26 is used, with the combined efficiency-weighted exposure.

In order to test the code produces the expected result, the limit for the DMTPC 10L surface run was computed and compared and found to be well reproduced, with a minimum cross section found to be 2.7×10^{-33} cm^2 at a WIMP mass of 150 GeV/c^2, compared to the previous value of 2.0×10^{-33} cm^2 at 115 GeV/c^2. To check the combining of results, the code was run on data from EDELWEISS [63] and CDMS [64] and the two combined. This was useful as the two data sets have already been combined as shown in figure 3-39 [65], so a direct comparison to a published result could be made. The only difference, in this case, is that the results are for the spin-independent cross-section, but this does not affect the method of combining results. A comparison of results can be seen in table 3.10.

The computed results do vary from the published results, however these exper-

Table 3.10: A comparison of published dark matter cross-section exclusion-limits to those computed with the updated limit-setting code.

Experiment	M_D GeV/c^2	Published σ (cm^2)	Calculated σ (cm^2)	Difference (%)
CDMS	70	3.8×10^{-44}	4.7×10^{-44}	21
EDELWEISS	85	4.4×10^{-44}	6.0×10^{-44}	34
Combined	90	3.3×10^{-44} at	5.0×10^{-44}	52

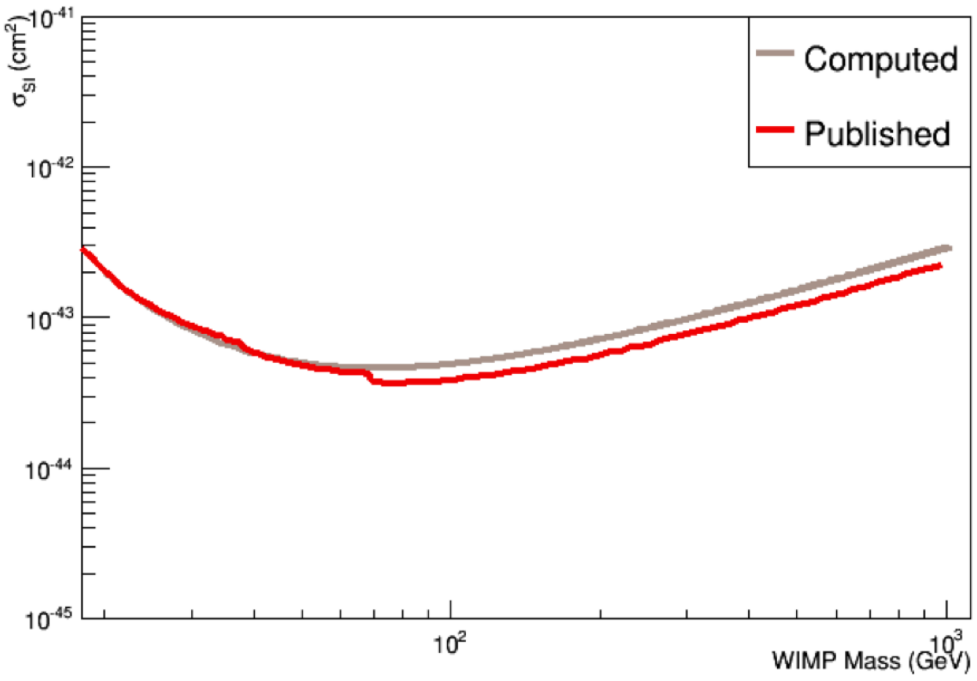

Figure 3-37: CDMS published result compared to calculated.

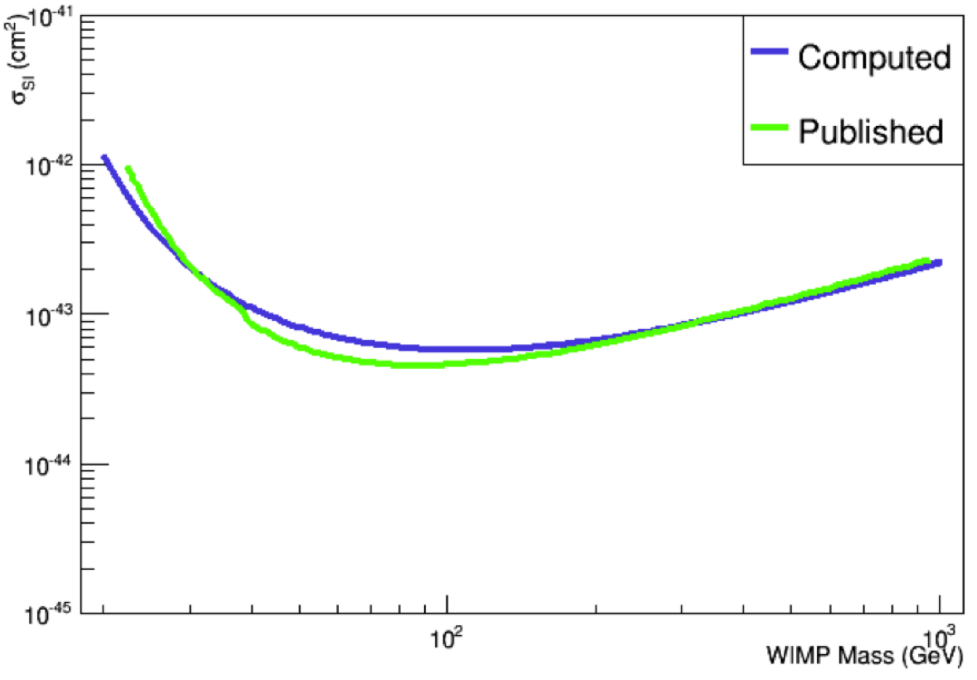

Figure 3-38: EDELWEISS published result compared to calculated.

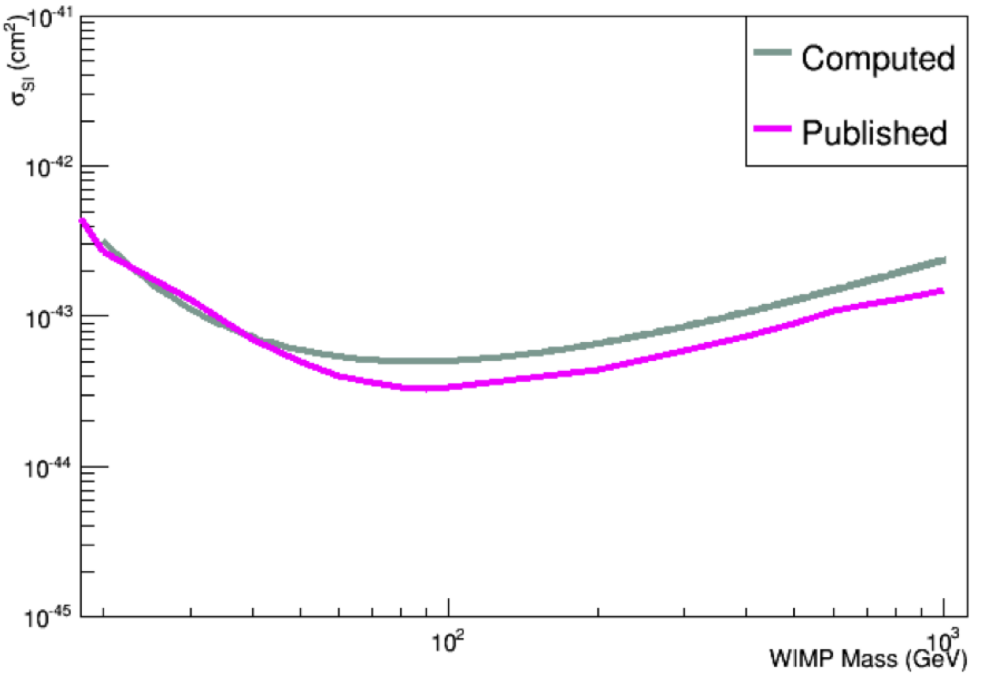

Figure 3-39: CDMS and EDELWEISS combined published result compared to calculated.

iments calculate the limit using the optimal interval method [66]. In simple terms, this method chooses the energy range of interest based on that which minimises the number of events that have passed the selection criteria. This means the energy of the events is taken into consideration. This was not used for this calculation and is the most reasonable cause for the variation. The two individual experiments agree well for many M_D, for the CDMS result, the dip in the published result is likely due to the low energy of the passing events. This dip also drives the difference between the computed and published combined limit at higher mass ranges. Despite these discrepancies the agreement was deemed sufficient to implement the use of the new prescription.

3.5 Results

3.5.1 Rates

Table 3.11 shows the number and rate of events and the reduction in these as each cut is applied. For comparison, the rates table for WR5 is also presented in 3-40.

Though it is not straightforward to do a direct comparison of the event rejection, due to the changes in the detector configuration, the fact that the chamber was opened in between runs and the change to the noise in the bottom camera, the final rate figures are useful to compare. For WR5 each camera had a final rate of events $\mathcal{O}(10^{-4})$ Hz. The final rates for WR6 are less than half this value, indicating an improvement in the background rejection. Not listed in the rates for WR5 is the percentage of non-spark images, which was 87.4 % for the bottom camera and 84.3 % for the top. The corresponding values for WR6 are 86.7 % and 91.5 %, indicating that the lower voltage has reduced the spark rate.

	Camera					
	081264		100439		Total	
Cut	Number	Rate (Hz)	Number	Rate (Hz)	Number	Rate (Hz)
Non-Spark Images	1666084		1607601			
All Tracks	221586	0.133	189195	0.118	410781	0.251
Non-RBI	148966	0.0894	181921	0.113	330887	0.203
Non-Worm	17456	0.0100	25654	0.0160	43110	0.0260
Fiducial	1326	0.000796	1343	0.000835	2669	0.00163
50 < E < 300 (keV)	806	0.000484	702	0.000437	1508	0.000921
Non-Cutoff	446	0.000268	390	0.000243	836	0.000510
Remerge	434	0.000260	383	0.000238	817	0.000499
Non-Edge Spark	433	0.000260	380	0.000236	813	0.000496
Non-Spacer	368	0.000221	268	0.000167	636	0.000386
80 < E < 200 (keV)	205	0.000123	162	0.000101	367	0.000224

Figure 3-40: Equivalent rates table for the previous run, reproduced from [57].

Table 3.11: Table of event rates on application of selection cuts for all source free data of WR6.

No. Images	165000			
Cut	Bottom Cam	Rate (Hz)	Top Cam	Rate (Hz)
Total	165000		165000	
Non-Spark Images	143124	86.74 (%)	150955	91.49 (%)
Spark Images	21876	13.26 (%)	14045	8.51 (%)
Tracks	69260	9.7×10^{-2}	152368	2.0×10^{-1}
Charge	65856	9.2×10^{-2}	145495	1.9×10^{-1}
E Match	3665	5.1×10^{-3}	74209	9.8×10^{-2}
RBI	2330	3.2×10^{-3}	3620	4.8×10^{-3}
Worm	2073	2.9×10^{-3}	756	1.0×10^{-3}
Edge/Region	280	3.9×10^{-4}	112	1.5×10^{-4}
Range	106	1.5×10^{-4}	103	1.4×10^{-4}
Long Tracks	52	7.3×10^{-5}	82	1.1×10^{-4}
Centroid	52	7.3×10^{-5}	82	1.1×10^{-5}
Remerge	52	7.3×10^{-5}	82	1.1×10^{-5}
Spacer	45	6.3×10^{-5}	72	9.5×10^{-5}
Energy	19	2.7×10^{-5}	30	4.0×10^{-5}

3.5.2 Cross-Section Limit

A limit on the spin-dependent WIMP-proton interaction is determined using the methods of Lewin and Smith [21] and Feldman and Cousins [49] as outlined in sections 1.3.3 and 3.4.2. The results of the calculation are presented in figures 3-41 and 3-42. Figure 3-41 demonstrates that there is still some work to be done before DMTPC can produce a leading limit. Figure 3-42 compares the analyses that have been performed on the 10L detector to date. The figure shows that, despite WR6 having a lower total exposure time, the minimum cross-section limit is still comparable to that of WR5.

The effect of the higher energy threshold for the bottom camera can be seen in the low-mass region. The bottom camera, however, performs better at the high mass end, due to having fewer remaining candidate events. However, the two combine to produce a more stringent limit at an intermediate mass value. The minimum cross-section and corresponding mass values are given in table 3.12.

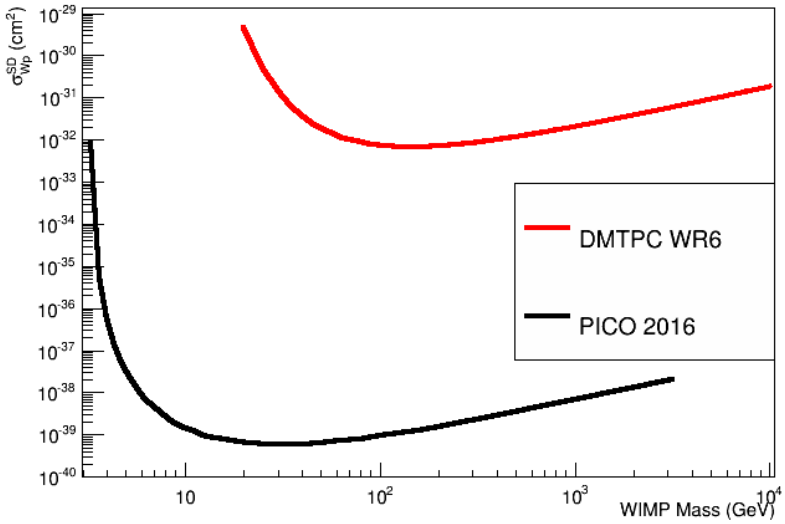

Figure 3-41: Calculated upper limit for WR6 with cameras combined. Also shown is the current leading limit from PICO [26].

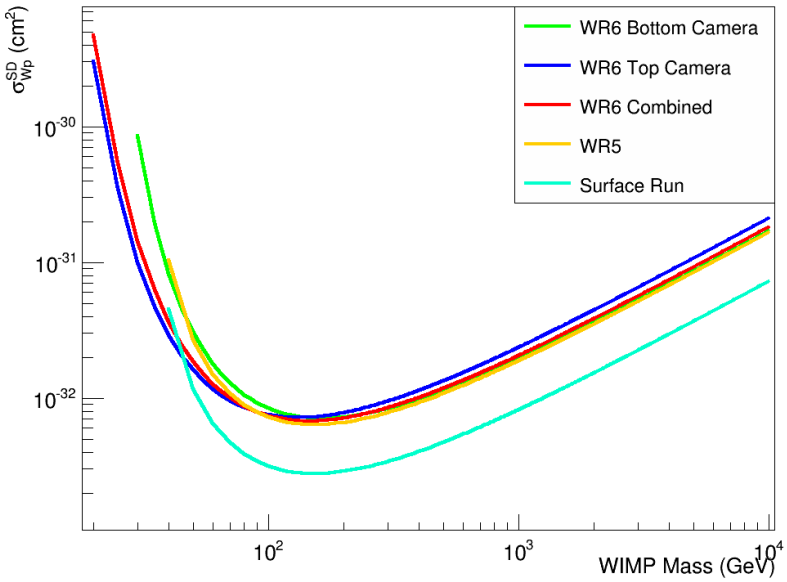

Figure 3-42: Calculated upper limits for the DMTPC results, including from previous analyses.

Table 3.12: Table of the exclusion limit on σ_p^{SD} and corresponding M_D for this analysis.

Experiment	Minimum σ_p^{SD} (cm^2)	M_D (GeV/c^2)
WR6 Top	7.2×10^{-33}	130
WR6 Bottom	7.1×10^{-33}	165
WR6 Combined	6.8×10^{-33}	145
WR5	6.4×10^{-33}	150
Surface	2.77×10^{-33}	150
PICO 2016	5.9×10^{-40}	31.62

3.5.3 Background Estimation

An estimate of the neutron background contribution due to the uranium and thorium contamination of materials is provided. The materials used were stainless steel, for the field rings, and acetal for the support structure. Due to the unavailability of information on acetal, the values for acrylic are used instead. Table 3.13 shows the contamination values that were used for the calculation. No radioassays of the materials used were performed, conservative values for the uranium and thorium content of materials are taken from those provided by the UKDMC [67]. The values for the spontaneous fission plus (α-n) interaction cross-sections were produced using SOURCES and were kindly provided by Dr. Chamkaur Ghag.

Table 3.14 shows the mass contributions for each material and the resulting neutron rate. Since each camera has a different live time, owing to differing number of sparks and RBIs[9], the expected background rate must be calculated independently. To do this, half the total mass is attributed to each TPC and the total neutron yield

Table 3.13: Radio-impurities of materials and the associated neutron yield due to spontaneous fission (SF) and (α-n) interactions. Units are parts-per-billion (ppb) and neutrons per second per gram per ppb.

Material	^{238}U (ppb)	Neutrons (n/s/g/ppb)	^{232}Th (ppb)	Neutrons (n/s/g/ppb)
Steel	5	1.85×10^{-11}	5	5.77×10^{-12}
Acrylic	10	2.65×10^{-11}	4	5.05×10^{-12}

[9]A pixel-time equivalent reduction is made to the live time, see the exposure calculation appendix.

calculated based on the live time for each camera. Table 3.15 shows the resulting values.

The other main contributors to nuclear-recoil like events are RPRs. An estimate for the RPR contribution is deduced using radon emanation values obtained from the DEAP/CLEAN radon resource [68]. For most materials multiple values are quoted, the values used are conservative and correspond to the highest value listed. Again, the values for acrylic are substituted for acetal. Table 3.16 shows the radon emanation values used and table 3.17 shows the estimated amount of radon produced during the analysis live time. The value is calculated for each camera individually.

The values presented in these background estimations are all conservative as they account for emanation into the entire detector and not into the fiducial volume. They are useful in demonstrating that the radon content of materials is a far larger contributor to backgrounds than neutrons leading to the conclusion that the remaining events from the analysis are most likely due to RPRs.

Table 3.14: The calculated mass and total neutron rates, combining uranium and thorium contributions, for the 10L detector materials, alongside total number of neutron events expected during the analysis live time (per camera).

Material	Mass (gms)	Neutron Rate (n/s)
Steel	7,440	9.1×10^{-7}
Acrylic	288	8.2×10^{-8}

Table 3.15: Total number of neutrons expected for each camera.

Camera	Neutron Yield
Top	0.37
Bottom	0.35

Table 3.16: Radon emanation rates for the materials of the 10L detector.

Material	Radon Emanation (per m^2 per hr)
Steel	0.036
Acrylic	8
Copper	0.0072

Table 3.17: Estimated amount of radon produced during the analysis live-time.

Material	Top Cam	Bottom Cam
Steel	0.3	0.3
Acrylic	170.0	159.0
Copper	0.1	0.1
Total	170.4	159.4

3.6 Conclusions

From table 3.11 and figure 3-40, it can be seen that the rates of events passing all cuts in WR6 has been reduced by an order of magnitude, compared to WR5. There have been many changes made to the detector operation and analysis procedure so it is difficult to say exactly which of these changes dominate the improvement. Undoubtedly it is the combination of modifications that have achieved this result. Looking at the calculated limit for WR6, it can be seen that the top camera has the best result at low mass-values, this is due to the lower energy threshold not being affected by additional noise in the way the bottom camera was. However, the bottom camera has a better upper limit at higher masses, this is due to there being fewer events overall. This is understandable as the noise causing the lower energy events to be washed out by the noise. When combining the two cameras, a lower limit, at an intermediate mass-value, is achieved.

The limit for WR6 is very close to WR5, however it has not beaten it. Reasons for the limit to be higher are that the detector had to be reopened in order to refocus the bottom camera - exposing the internals to radon-contaminated air - and the reduced overall exposure time that was taken. Additionally, the read noise problem developed by the bottom camera developed, raising the energy threshold. Considering the live time is 73 percent of WR5, that the limit is so close shows considerable improvement overall.

This work has been useful in improving the analysis process in preparation for the new $1m^3$ detector, by close examination and improvement of the cuts and updates to the limit setting procedure. It has also shown that there are background events that must still be excluded. Since the fraction of non-spark images has increased, the number of RBIs will be decreased. The rejection of RBIs has also been modified to work across data sets, meaning the number passing cuts will be reduced. With the improvement in rejection of these backgrounds and the low contribution of neutrons,

as shown, the most likely candidate for the background events are the decay products of radon. This is a driving force behind the design of the internal components of the new detector, as will be seen in chapter 4.

Chapter 4

Field Cage Design

4.1 Introduction

This chapter is concerned with the design of the field cage for the $1m^3$ detector. As described in chapter 3, though the rate of events passing the background rejection criteria was reduced in comparison to the previous underground analysis, there are still a number of events that pass. The most probable source of these recoil candidate backgrounds is radon progeny, as shown in section 3.5.3. Not only does radon emanate from materials but environmetal radon also plates out on detector surfaces. The design, therefore, of the $1m^3$ aims to minimise the amount of material used for the field cage and its support structure. Of particular significance to the experimental sensitivity, and the goal of this chapter, is to minimise the surface area to fiducial volume ratio.

Another matter of critical importance to the design is constraining the height of the drift region. This value was set to minimise diffusion of the ionised electrons and was based on DMTPC studies of diffusion [45, 55]. From these studies, a track of $\sim$1 mm in length[1] was determined to undergo a transverse spread of 1 mm at a starting drift height of 25 cm. It is clear that the ratio of the track range to track width will

[1]For a 50 keVee recoil generated at 75 Torr and a reduced field of 10 Td.

impact on the ability to reconstruct the sense of a track. 25 cm was therefore chosen as the maximum height that would ensure this ratio does not go below unity. A study has been done by the group to quantify and optimise this metric and can be found in reference [47].

In order to optimise the design, the field cage was modelled with different geometries with the aim of achieving a configuration that would minimise the surface area of the material used, whilst maintaining a fiducial volume of 1m^3 - a design goal of the original concept. The final design uses field shaping rings with the long axis parallel to the drift direction, which is contrary to previous iterations, and also rounds the edges of the rectangular profile. This rounding not only reduces surface area but reduces the number of sharp edges from which arcing can occur. This design is also used by ARGONTUBE [69] and EXO [70] and is preferable to the previous design, where the long axis was normal to the z-direction, as the rings take up less space in the radial direction, allowing a smaller inner diameter for a particular outer diameter and thus a greater fiducial volume. See figures 4-1 and 4-2 for diagrammitic representations and photos of the ring orientations. In these figures are shown the variables a and b for the rectangular shape. The width of the rectangle is the a variable and b is the length.

The design process for this project is unique as it used a simulation-based approach to optimise the design, using a new track-survival metric. The electric field of previous detectors has been studied using a field-uniformity metric of $\frac{E\perp}{|E|}$. The new metric traces the path of tracks, within the electric field, with the aim of eliminating RPRs entering the fiducial volume. By tracing the path of a particle that starts at the surface of a field ring and drifts to the cathode, where it can subsequently decay, the fiducial radius can be set so that those particles do not enter the fiducial volume. This is the first instance of such design-optimisation and is discussed in further detail throughout the chapter.

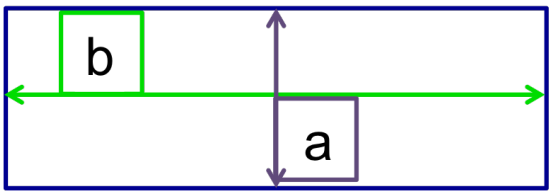

(a) Original ring cross-sectional profile.

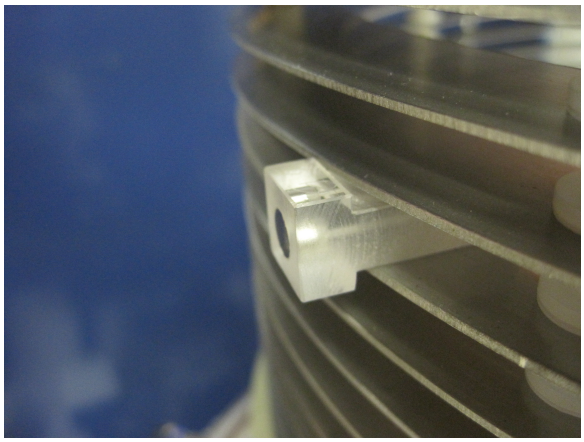

(b) Original design field cage stack.

Figure 4-1: In a) is shown the cross-sectional profile of all previous TPC ring designs (not to scale). The flat face was always pointing upwards, closely spaced and with the normal parallel to the drift direction, as can be seen in (b). The small acrylic piece in (b) is a source holder, used to produce alphas in the detector for calibration (as described in chapter 2 sections 3.2.3 and 3.2.4.

As mentioned in section 2.3.3, the $1m^3$ detector is orientated with the drift direction in the horizontal plane, in contrast to previous DMTPC prototypes which have drifted in the vertical direction. This required a new support structure which was also designed and produced as part of this project. For this the new ring orientation was also beneficial: instead of threading rods through the rings (as seen in the 10L design, figure 2-8) then having to connect these to the chamber walls, the rings themselves can be clamped onto.

The outline of this chapter is as follows: section 4.2 describes the simulation technique employed, including the software that was used and the choices made for generating a result. Section 4.2.5 gives the results of the initial simulations and describes the reasons for the choice made for the final design, with some comparison to other results and why they were not chosen. Section 4.3 outlines the design of the cathode and the components used to support the field cage, amplification region and cathode within the vacuum chamber. Not included in this project was the design and fabrication of the amplification region. The current design for this has the meshes

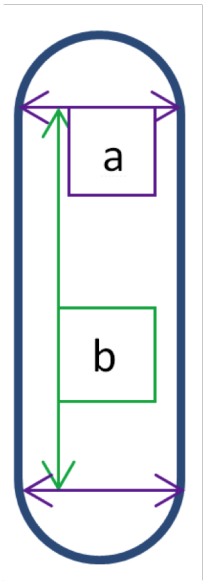

(a) New ring cross-sectional profile.

(b) New design field cage stack.

Figure 4-2: In (a) is shown the new ring cross-sectional profile (not to scale). It can be seen in (b) that the profile has been rotated such that the flat face of the ring is directed inwards and is now normal to the drift direction. In (b) can also be seen the much larger ring spacing that was achievable in this new configuration. The large white block that surrounds the rings consists of individual ring support housings inside of which is a resistor connecting each ring to the next. The white material behind the copper rings is part of the support structure. Both of these will be discussed further on in the text.

segmented into quadrants and each quadrant held in place by a cruciform-shaped frame. Section 4.4 then describes an additional simulation that was performed after the amplification region was designed and built. This was in order to model the impact that the additional plastic of the cruciform, that crosses the amplification plane and extends further than the field rings into the drift region, has on the drifting of electrons within the drift region and the impact on the overall fiducial volume.

Section 4.5 details the mechanical prototyping of the new design concept, including the building of the test-stand detector into an operational state. Section 4.6 gives a brief account of the optimisation of the field-cage support structure design that was aided by the assistance of the mechanical engineers of Bates Laboratory. Section 4.7 presents the overall result of the field cage design in reference to the total surface area of all components and the fiducial volume, specifically the ratio of the two values. The same values are given for the previous iterations of detector design to demonstrate the significant improvement that has been made.

4.2 Simulation

4.2.1 Finite Element Analysis

To simulate the detector geometry, a finite element analysis (FEA) was performed. FEA is an approach used to solve complicated physical systems for which an analytical solution would be very difficult, if not impossible to find. It can be used for many applied parameters such as stress, strain, fluid flow or, as in this case, electrostatics. In order to do this, the system is discretised by breaking the geometry down into small elements and performing the relevant calculation for each element. The composition of these geometric elements is referred to as the mesh and the intersection of each element of the mesh is referred to as a node.

The geometry of the field cage is rotationally symmetric along the central axis

thus a 2D slice is sufficient to describe the problem. This is beneficial also, as the computer memory requirements are kept low and errors due to use of an unstructured mesh[2] are minimised. The process used to generate and compare different geometries was:

1. Initial concepts generated in CAD[3] using Autodesk Inventor
2. Geometry for simulation generated in Gmsh[4]
3. Mesh produced by Gmsh and converted using Elmer[5]
4. Electric field map calculated by Elmer
5. Paraview[6] used to plot electric field lines[7]
6. Track survival deduced using ROOT [74]

4.2.2 Geometry

The geometry itself is fairly simple, however, it is necessary to generate a completely separate simulation for each configuration, since once a mesh is created, it cannot be easily modified. This is because the mesh file is a text document containing a list of node co-ordinates and it is a non-trivial task to modify these to represent a new geometry. So discrete values for the variable parameters are chosen and compared. To assist with the listing of these variable parameters, the detector axes are reproduced again in figure 4-3.

[2]An unstructured mesh consists of an irregular pattern of, typically, triangles or tetrahedra. This allows greater flexibility in modelling more complicated geometries. A structured mesh uses parallelotopes, which produce a more regular pattern but is more restrictive in geometrical application.

[3]Computer aided design

[4]A three-dimensional finite element mesh generator in c++ [71].

[5]Multiphysical simulation software [72].

[6]Data analysis and visualisation software [73].

[7]The field lines describe the path that a massless charged particle would take in the presence of the field.

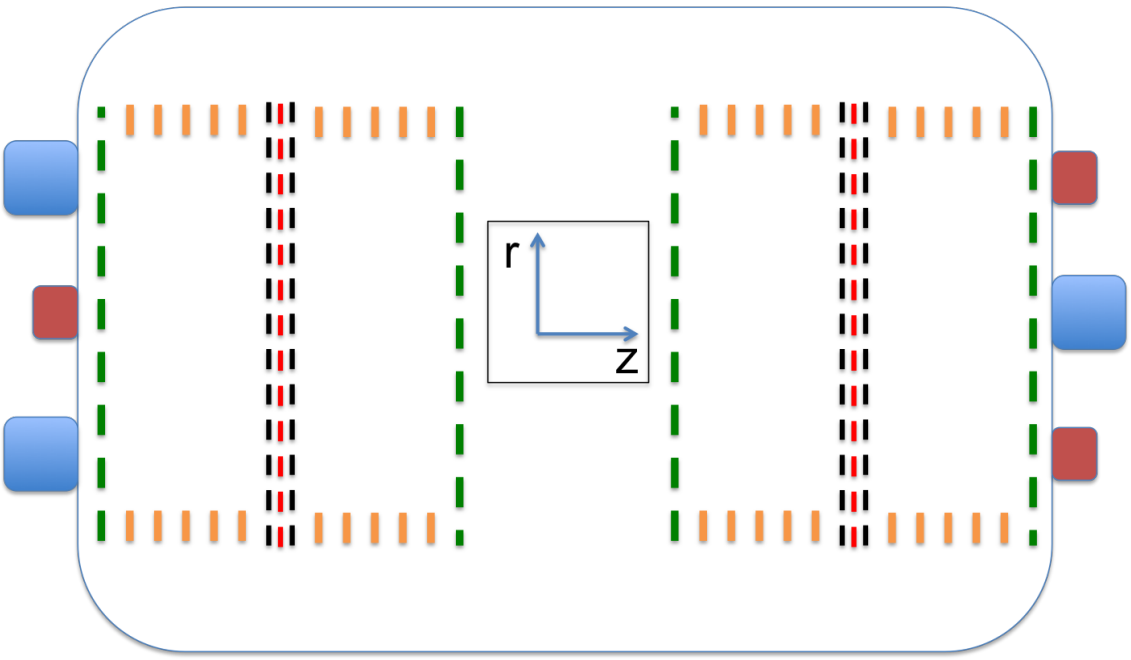

Figure 4-3: Diagram of the $1m^3$ detector with axis reference, as previously seen in chapter 2.

The list of variable parameters is then:

1. Ring thickness in the radial direction (this relates to variable a for the new configuration, figure 4-2 and to variable b for the old, figure 4-1)

2. Ring height in the z-direction (this relates to variable b for for the new configuration, figure 4-2 and variable a for the old, figure 4-1)

3. Ring spacing - the distance between each ring in z

4. Ring cross-sectional profile

It should be noted that in the first round of simulations, the rounded rectangle geometry was generated in a way that added the curvature after setting the length (b) to the desired value. As such, the total height of the rounded ring, for these simulations, must include the rounded ends. For simplicity the roundness was kept uniform, such that the radius of curvature was half the width of the profile, i.e.

variable a. This makes the total height of the rounded ring profile $b + a$. This is crucial to keep in mind when comparing the results of different ring geometries as the ring heights will be different, but the b-values will be the same. This was later updated such that the total height of each geometry is directly comparable.

The number of rings required for each TPC was calculated by using 25 cm as the minimum drift height, but this was allowed to increase as a consequence of the discreteness of the number of rings and ring-spacing of each simulation, as will be seen in the results section. The different spacings that were tried were: $\frac{1}{4}$", $\frac{1}{2}$", 1" and 1 $\frac{1}{2}$". For comparison, the ring spacings of the 10L and 4Sh are roughly $\frac{2}{5}$" and $\frac{1}{4}$" respectively. Many different ring heights and thicknesses were tried, initially based on available metal sizes from two vendors: OnlineMetals.com [75] and McMaster-Carr [76] and ranged from several mm to several cm. There were three different ring cross-section profiles tried: round, rectangular and rounded rectangular (as seen in diagrams 4-1 and 4-2). The outer diameter of the ring was fixed and chosen such that there would be a sufficient stand-off distance from the vacuum chamber wall to minimise the possibility of sparking to occur. This was deduced via the Paschen Law, using the minimum expected operating pressure of 20 Torr and a maximum expected operating voltage of 8 kV (as set by the maximum output of the high-voltage supply to be used). The distance that satisfies the extremes of operation is 15.5 cm. With the vacuum chamber inner diameter being 151 cm, this gives a ring outer diameter of 120 cm.

The geometry is generated by creating a Gmsh function for each different ring type, which is iterated to be drawn at the appropriate spatial locations in the field cage stack. Then the boundary of the vacuum chamber and cathode, ground and anode electrodes are drawn. The electrodes in the real field cage are made from a mesh of woven steel however, due to the small radius of the wire of 15 μm, and their close proximity of 457 μm, it is impractical to model, owing to the resolution of mesh

that would be required, thus a solid plane was used instead. In this configuration, an element of the mesh covers about 100 wires of the mesh. It is therefore more practical to use a solid plane than place a wire at every 0.5 cm (the mesh resolution used). Given that the electric field of a wire is inversely proportional to the wire radius, at 15 μm these effects are only noticeable very close to the plane of the woven mesh. The final geometry configuration can be seen in figure 4-4.

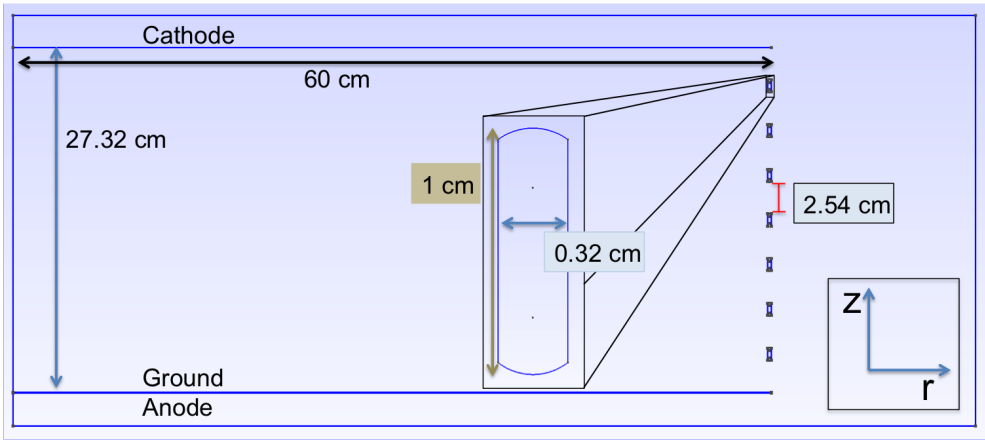

Figure 4-4: This is the geometry of the final model as viewed in the Gmsh viewer. The vertical is the z-axis in relation to the detector, the horizontal is the radius. The horizontal bars are: cathode (top) and amplification region (bottom). The boxes aligned vertically on the right hand side are the profiles of the rounded rectangular rings - a close up is provided to show the final geometry as per the material delivered (see section 4.7). All these are bounded by a chamber geometry, the blue line that encloses all previously mentioned components.

4.2.3 Mesh

The mesh is generated using quadrilaterals and, as mentioned, was generated at a resolution of 0.5 cm, figure 4-5 shows an example mesh. This resolution was chosen after several simulations were generated to test the process. 0.5 cm was found to be the finest resolution that was computable in a reasonable amount of time and with reasonable computer memory allocation to allow many simulations to be quickly generated for comparison. This resolution factor does not restrict each quadrilateral to being 0.5 cm in length on each side, but requires that the geometrical elements that compose the model are discretised by 0.5 cm increments. The mesh is then generated on this basis and elements may be produced that are slightly smaller than this to allow for the structures that are in place. A quadrilateral mesh element was chosen over triangular due to the geometry of the structure being simulated - the potential gradient is expected to be very close to uniform in the central drift region, thus the use of quadrilaterals produces almost parallel lines in this region, which would be

expected to minimise errors in the computation per element. As can be seen in figure 4-6, where a triangular mesh was tested, the nodes are quite randomly distributed and do not reflect the expected uniformity of the simulation. Though no qualitative comparison was made between the two types, the uniformity of the quadrilaterals was chosen as best representing the geometry being modelled.

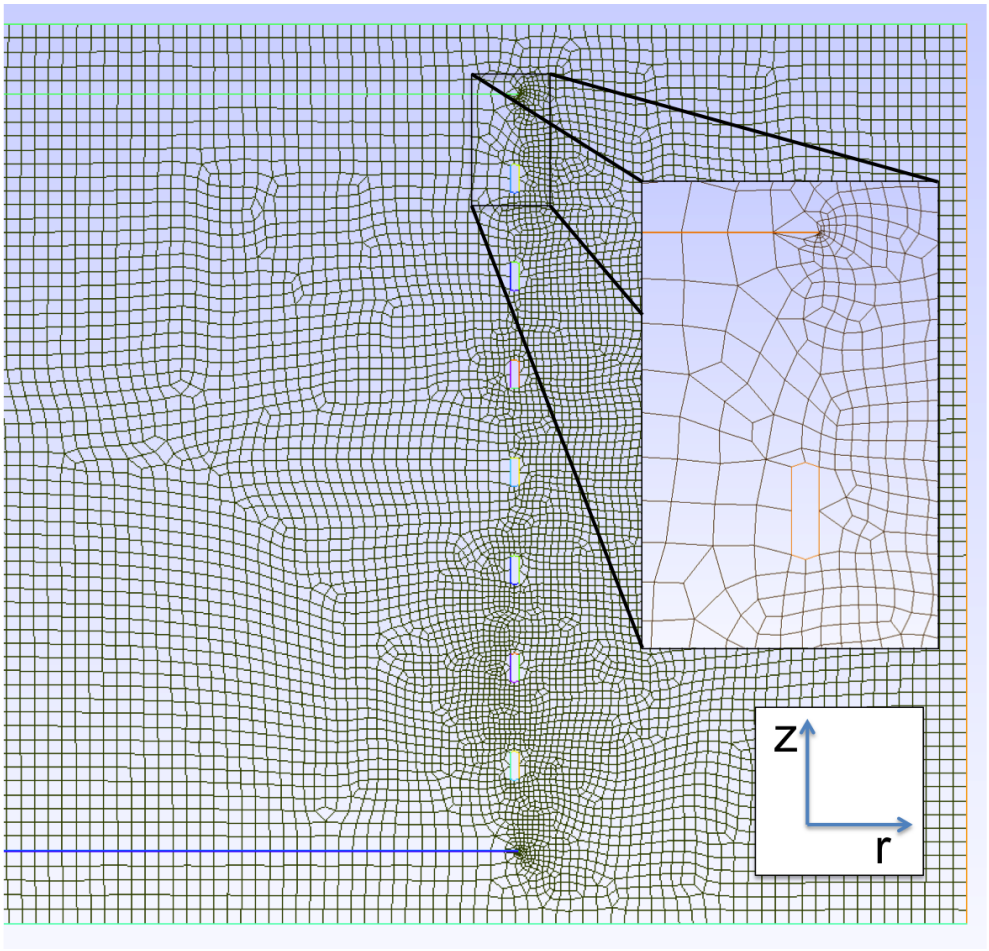

Figure 4-5: An example of a quadrilateral mesh used in the simulation process. The expanded area shows how the geometrical elements are subdivided by the set resolution of 0.5 cm, the mesh is then generated from the divided starting nodes. The coloured lines are the geometrical components and the black lines represent the elements of the generated mesh. The relative uniformity of the elements can also be seen in the more central drift region on the left hand side of the image.

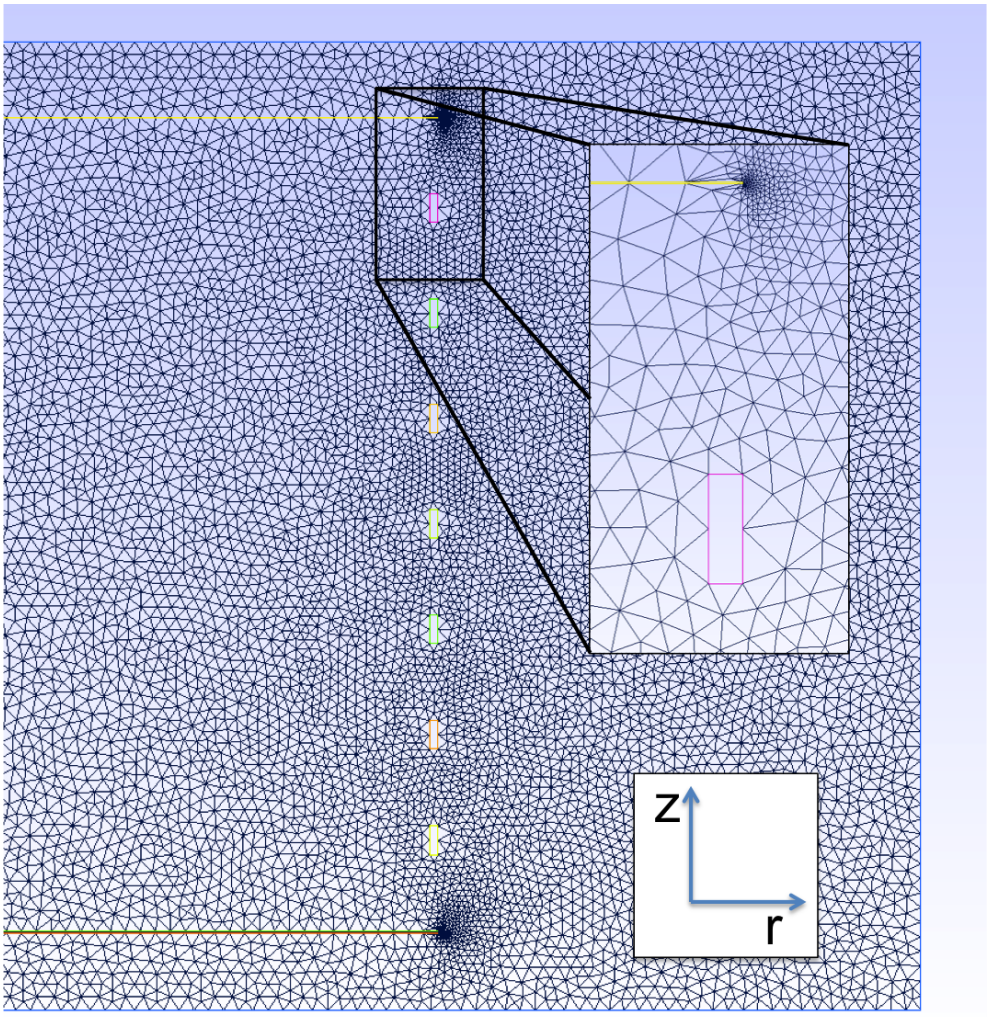

Figure 4-6: A mesh generated using triangular elements. The more random distribution of elements can be seen in comparison to figure 4-5.

4.2.4 Post Processing

An example of the electric field map that is generated by Elmer can be seen in figure 4-7. The program ParaView is used for the visual representation and to generate field lines within the field map. It should be noted that no circuitry was simulated i.e. the field rings were not connected by resistors and the voltage drop deduced from the connections, rather the voltage was divided evenly across the rings and the appropriate value applied to each. Two sets of field lines are generated; one at the surfaces of a bounding box for each of the rings that make up the field cage, shown in figure 4-8, and another just below the plane of the cathode shown in figure 4-10. The first set of field lines is used to estimate the track survival probability of the expected radon decay progeny to drift from the ring surface to the cathode, where they can stick and subsequently decay. These decays can mimic the WIMP signal, if the resultant track successfully drifts to the amplification region. Therefore, by ensuring the radius to which these field lines survive is outside of the required fiducial radius, any events of this type should be minimised, if not eliminated. The radius to which these field lines are propagated is denoted r_1. In the example of figure 4-8 the value of r_1 would be reported as 57.6 cm. Figure 4-9 shows the second set of field lines, those generated just below the cathode. These are used to deduce the radial position at which tracks generated from this location, i.e. with the longest drift length, will successfully drift all the way to the amplification region and be readout by the readout channels. It has been seen in the four-shooter, that some tracks that start at the top of the drift region have not made it to the amplification region, thus reducing the effective fiducial radius. Ensuring the field lines from this location extend to the ground plane of the amplification region ensures a track will enter the amplification region and be detected. The maximum radius at which field lines extend to the ground plane, and thus survive to the amplification region is denoted r_2. This can be visualised by figure 4-10: the white line that drifts outside of the drift region,

that is mentioned in the caption, would set the limit on r_2. The value reported would be the originating r co-ordinate value of the previous, surviving, field-line.

The two values, r_1 and r_2 can be combined to set a fiducial radius for which nuclear recoils produced within this radius, at high z, will successfully drift to the amplification region and be readout, but radon progeny originating from any of the field ring surfaces, will not propagate to a point on the cathode at which the subsequent decay products can successfully drift to the amplification region - mimicking the expected signal. In this simplified construction, gas physics is not taken into account and the field lines represent the path taken by a massless point particle. The most significant effect missing due to this simplification is that of track diffusion, as discussed in section 1.2.4. In order to account for diffusion, the radial value for the track survival metric is reduced by 3σ of the typical track diffusion, where 1σ is $\sim$1 mm. This value is calculated using equation 1.13, using a constant value approximated from the results of the analysis shown in section 3.2.3 of 0.6 cm, a typical drift height of 25 cm, a voltage of 2000 V and pressure of 30 Torr[8], which results in a D/μ value of 4.3×10^{-2} and the initial z-position of the track being at the top of the field cage, 25 cm.

[8]The operating parameters for the $1m^3$ detector.

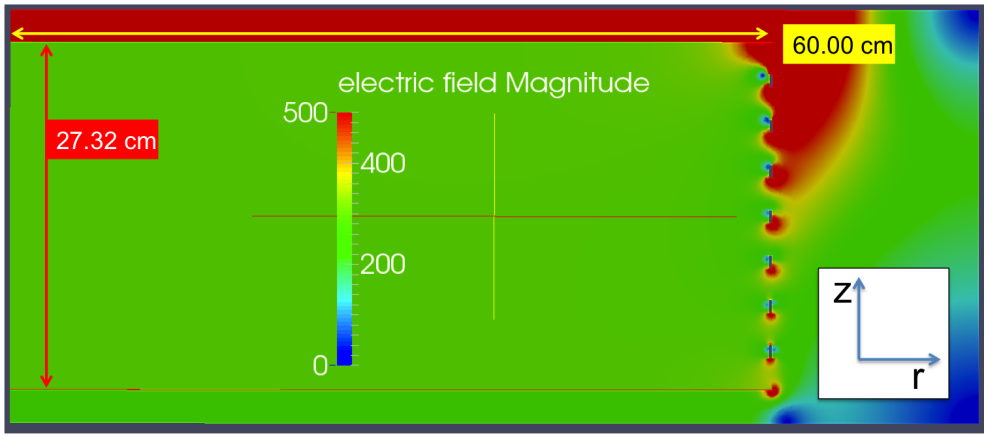

Figure 4-7: A field map generated by ElmerFEM, visualised in Paraview. The electric field magnitude is for the total electric field: $|\vec{E}| = \sqrt{E_r^2 + E_z^2}$ and is given in V/cm. Using figure 4-5 as a reference for the geometry, a clear uniform electric field can be seen well within the drift region, whilst outside, the field varies greatly.

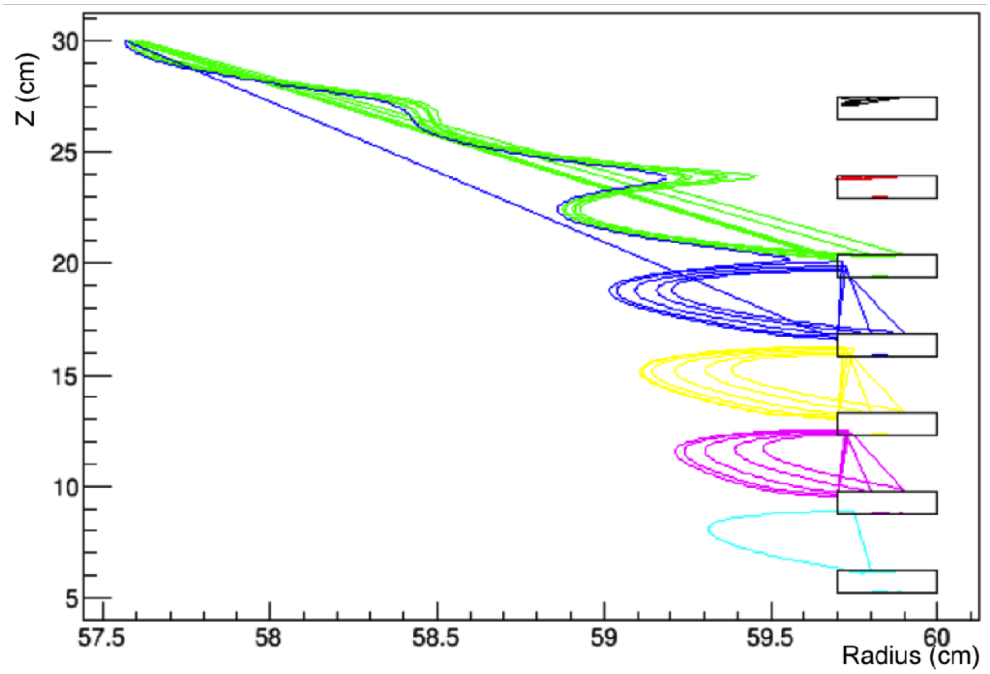

Figure 4-8: Field lines generated on bounding box of ring surface. The black boxes are the rectangular bounding box of the field cage rings and the starting point for the electric field lines. The different colours used in this image are used for visualisation purposes only - to be able to see which ring any field lines, that survive to the cathode (z=30 cm), belong to. The straight lines that can be seen are an artefact of the drawing technique.

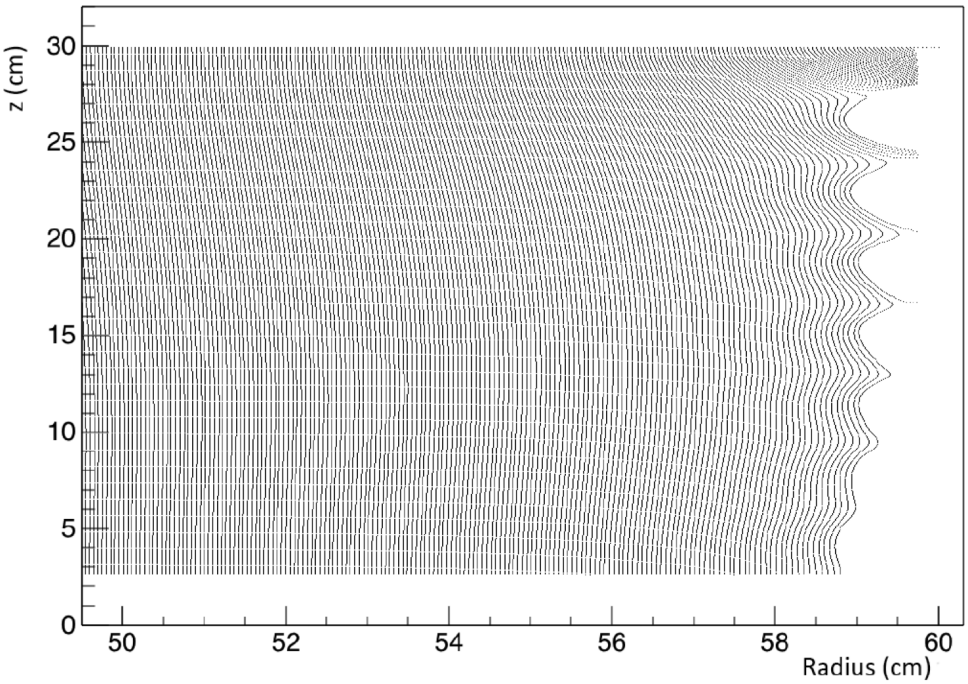

Figure 4-9: Field lines generated from just under the cathode. The abscissa is the radius and the ordinate is the z-direction in the detector. The lines propagate from the cathode and terminate at the ground electrode. This image shows clearly the field lines that terminate outside of the field cage, those in the top right corner of the plot. It also shows that the field lines follow a wavy pattern nearer the field rings. This is discussed further in section 4.2.5.

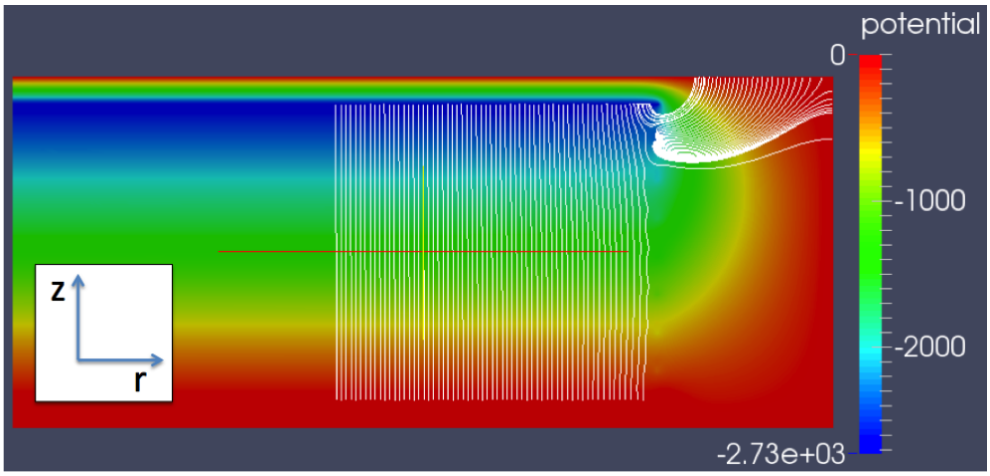

Figure 4-10: Electric field lines generated in ParaView. The white lines in this figure are the electric field lines. They are overlain on a map of the electric potential. The field rings are less discernible in this image due to the potential value being mapped. They lie in the same position as in figure 4-7. The cathode is located just above the horizontal white line in the upper part of the image. The field lines in the top right corner of the image are outside of the field cage. The field line just below this region, that starts inside the drift region and ends on the right hand side boundary, demonstrates how a nuclear recoil produced near the cathode, at this radius, would drift outside of the field cage.

4.2.5 Simulation Results

To achieve the $1m^3$ fiducial volume from four field cages with a nominal height of 25 cm, the minimum fiducial radius that could be tolerated is 56.5 cm. However, at focus the cameras are able to view to a radius of 57 cm. Allowing for the 3σ error on diffusion of 3 mm, this puts an upper limit on the allowable value for r_1 and r_2 of 57.3 cm. It is also important to keep the total surface area of the construction down, to minimise the total contribution to backgrounds and to minimise the weight for mechanical considerations. Table 4.1 gives some examples of the results produced from this study. This is a small sample of the full range of results, however it shows that the number of field rings can be reduced whilst maintaining the minimum fiducial radius and keeping the total surface area low. The final configuration is given in the table, the last line in set (3). Only set (1) contains a round ring profile result as by this stage it was apparent this profile would not be suitable, thus the following

sets compare the rectangle and rounded rectangle designs only. It was also decided at this stage to not use the standard available copper sizes so arbitrary heights, in units of 1 cm were modelled. It will be noted that the r_2 value for the chosen model is in fact below the stated 57.30 cm requirement. However, as it is only 1 mm away from the desired value it was deemed acceptable. One might conclude that the non-rounded rectangular profile should have instead been chosen, owing to the r_1 and r_2 values, however notice the increase in surface area this produces. Recall also that it is beneficial to have rounded edges in order to reduce the chance of charge build-up on pointed surfaces and thus reduce the likelihood of arcing.
For quick reference, the variables of table 4.1 are:

1. *Round* is a ring with a circular cross-section profile
2. *Rnd rect* is a ring with a rounded rectangle cross-section profile
3. r_1 = minimum inner radius reached by field lines from rings.
4. $r_2 = r_{start}$ for where the electric field lines propagate successfully to the amplification region, with inclusion of 3σ diffusion.

From the selection of models generated, the final configuration chosen, which satisfied the criteria, has the following parameters;

1. Number of rings, per TPC: 7
2. Ring height (z): 1 cm
3. Ring thickness (r): 0.32 cm
4. Ring spacing: 2.54 cm
5. Profile: rounded rectangle
6. Drift height: 27.32 cm

Table 4.1: An example set of results from the simulation study: a small set from of order 300 simulations that were generated. In the first set (1) are the results for a round, rectangular and rounded rectangular ring. The last two are at the same height and thickness (0.2 cm), the rounded is shorter and wider due to the geometrical value being set via the variable b, as described in the main text. You can see that all three of these results would satisfy the stated requirements i.e. r_1 and r_2 are both ≤ 57.30 cm. However they have large total surface area, due to the increased number of rings required by the small, 1.27 cm spacing. In set (2) are the same models but spaced at 2.54 cm, twice the previous distance. These start to fail the radial requirements, but the total surface area is lowered quite considerably. Set (3) shows the results for a ring with a height of 1 cm and a width of 0.32 cm at 1.27 cm and 2.54 cm spacing. Here it can be seen that the requirements for r_1 and r_2 are met at the spacing of 2.54 cm, contrary to set (1). The 2.54 cm spacing rounded rectangle result shown here is the one chosen for the final design. To see that additional variations did not improve the situation, set (4) shows the 1 cm height ring at 3.81 cm spacing and set (5) shows a 2 cm height ring at 2.54 cm spacing. Set (4) does not satisfy the requirements of r_1 and r_2 and though those in set (5) do, the surface area is considerably increased.

Type	Ring IR	Ring H	Spacing	Drift	N Rings	r_1	r_2	SA_{Tot} (m^2)
(1)								
Round	58.73	1.27	1.27	26.67	10	57.64	57.30	5.95
Rect	59.60	1.77	1.27	25.57	8	57.95	57.72	5.21
Rnd rect	59.60	1.77	1.27	25.17	8	57.90	57.73	4.76
(2)								
Round	58.73	1.27	2.54	25.40	6	57.21	56.90	3.57
Rect	59.60	1.77	2.54	28.38	6	57.26	57.64	3.91
Rnd rect	59.60	1.77	2.54	27.98	6	56.62	57.74	3.57
(3)								
Rect	59.68	1.00	1.27	26.24	11	58.46	58.10	4.37
Rnd rect	59.68	1.00	1.27	26.24	11	58.31	58.00	3.91
Rect	59.68	1.00	2.54	27.32	7	60.00	57.30	2.78
Rnd rect	59.68	1.00	2.54	27.32	7	57.52	57.20	2.49
(4)								
Rect	59.68	1.00	3.81	27.86	5	60.00	56.40	1.99
Rnd rect	59.68	1.00	3.81	27.86	5	56.71	56.4	1.78
(5)								
Rect	59.68	2.00	2.54	25.24	5	57.32	57.00	3.49
Rnd rect	59.68	2.00	2.54	25.24	5	57.30	57.00	3.12

7. Total surface area (four field cages): 2.49 m^2

8. Single ring weight: 0.96 kg

9. Fiducial volume = 1.1 m^3

The final vendor chosen for the ring material provided a slightly different thickness to that modelled (0.3 cm as opposed to 0.32 cm) and on delivery the metal had a slightly different profile, shown in figure 4-11 . The modifications were made to the simulation and no significant impact was observed: the new model had an r_1 value of 57.54 cm, only 0.2 mm different from the original (in fact an improvement, though very minimal) and an r_2 value of 57.2 cm - identical to the original simulation. The fiducial volume quoted in the list above is given for the initial minimum radius of 57.3 cm. However, with the drift height set at 27.32 cm, to achieve a fiducial volume of 1 m^3, the fiducial radius need only be 54 cm. This becomes useful when considering the tolerance on the build of the field cage rings. As seen in figure 4-9, the field lines trace out a wavy pattern near to the field rings. Though this effect was not used in the metric for determining the optimum geometry, figure 4-12 shows the maximum amount of radial drift, plotted as a function of the starting radius, for the chosen configuration. For a radius of 54 cm, the maximum deflection is just over 5 mm. The effect drops to 1 mm at $\sim$ 41 cm. This effect is discussed further in section 4.4. Figure 4-13 shows the variation in radial drift (for initial radius of 54 cm) and surface area as a function of ring spacing. The figure demonstrates that though the radial drift is reduced when the rings are more closely spaced, the total surface area contribution of the rings increases significantly. This additional surface area will not only contribute to the radon content within the detector but also to the impurities in the gas due to outgassing. Reduced spacing also requires more rings per field cage, increasing fabrication time, cost and complexity.

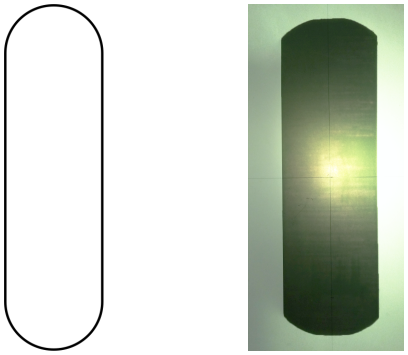

(a) (b)

Figure 4-11: Comparison of the original simulated cross-section profile (a) and the profile of the copper as delivered (b). Slightly sharp edges can be seen in the delivered copper, which could act as a point of charge build-up and thus arcing, however the rings are spaced far from each other and the chamber wall so this does not present a problem. This was verified at the stage of commisioning, section 5.2, where the field cage was operated at voltages above those that were required for operation.

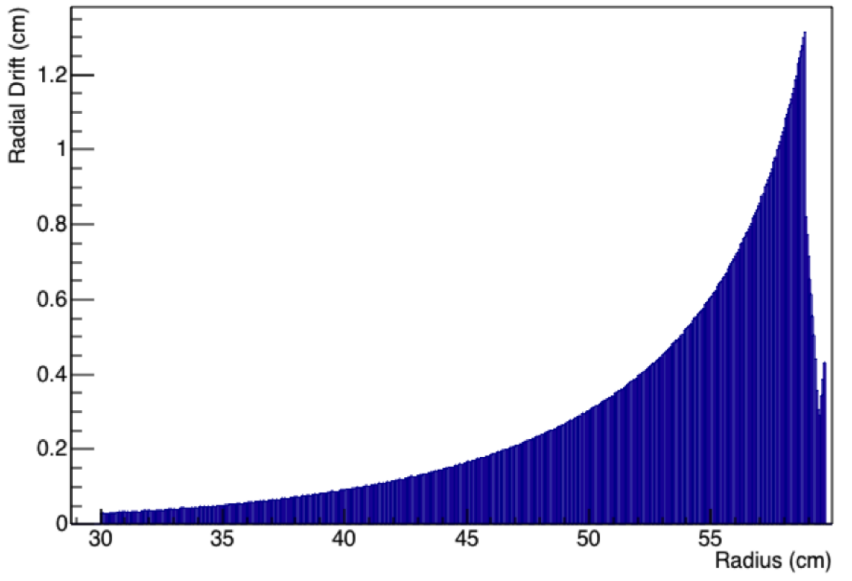

Figure 4-12: A plot of the maximum radial deflection of a field line as a function of the starting radius. The feature on the far right occurs at the location of the field ring inner radius.

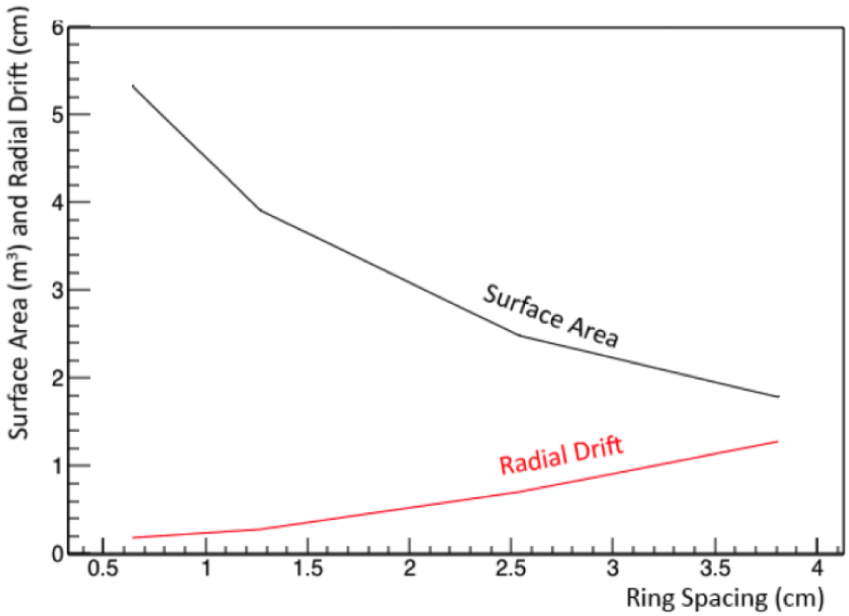

Figure 4-13: A plot of the field line deflection, for an initial radius of 54 cm, and the total surface area of four TPCs, as a function of the field ring spacing. A spacing of 1.5 inches does not satisfy the requirements of the survival metric. Though the radial deflection at 0.5 inch is less than at 1 inch, the total surface area is greatly increased. The surface area goes up by nearly 60 percent but the deflection is reduced by only about 40 percent.

4.3 Cathode and Support Structure

4.3.1 Cathode Design and Electrode Supports

Although the amplification region was not part of this project, the cathode was. As were the support structure to hold the cathode and amplification regions, which interface with the field cage support structure to make the whole TPC.

When designing the cathode a restriction was found in the form of the width of commercially available meshes, being a maximum of 48" (1.22 m). Though this is just over the 1.2 m of the field cage design of section 4.2.2, it leaves no room for placement in the stretching apparatus. The stretcher shown in figure 4-14 also has a limited aperture for placing the frame onto which the mesh was to be glued. With these constraints, the maximum achievable outer diameter of the frame was 111.76 cm. This would leave a gap in the radial direction, between the cathode and closest field ring. The effect of this gap would be a distortion of the electric field around this region. In order to make this up to the required 120 cm, metal flaps were designed to screw onto the outer edges, see figure 4-15. The design also used the amplification region design as a guide, specifically the hole placements for the support pieces, in order to allow correct alignment of both parts.

For the support pieces, the aim, once more, is to keep materials used to a minimum yet maintain sufficient mechanical support. The designs were kept as simple as possible and are shown in figure 4-16. The design was based on the premise of each support piece covering at least four of the aforementioned holes that are present in the circumference of both the amplification region and the cathode. The material used is Delrin which is a brand name for acetyl-resin.

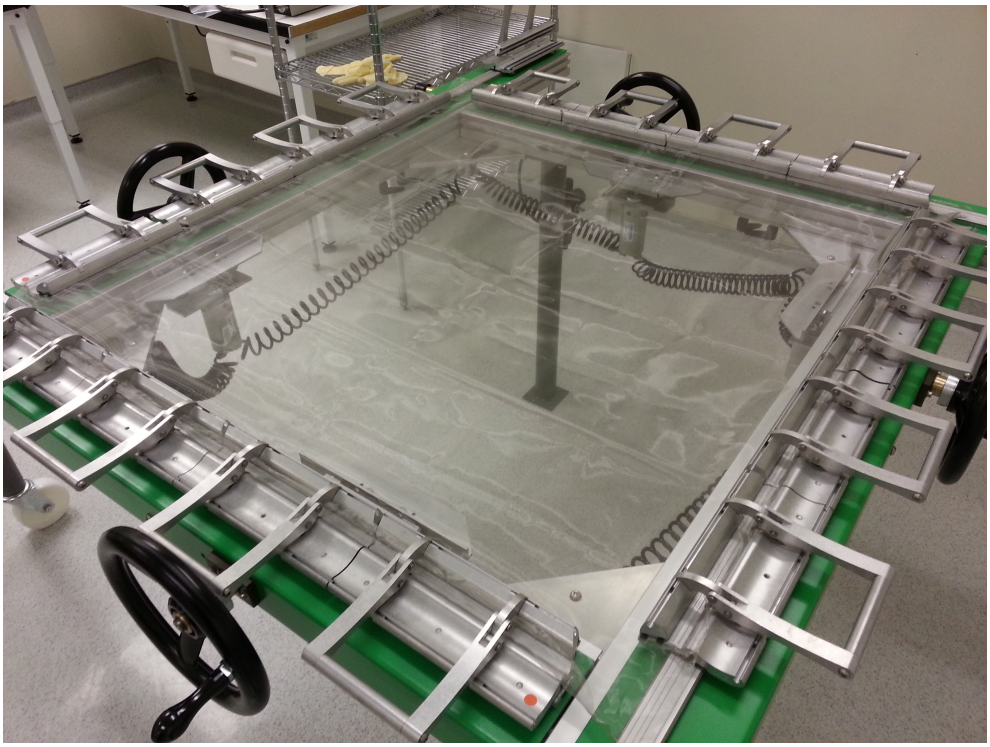

Figure 4-14: The mesh stretcher.

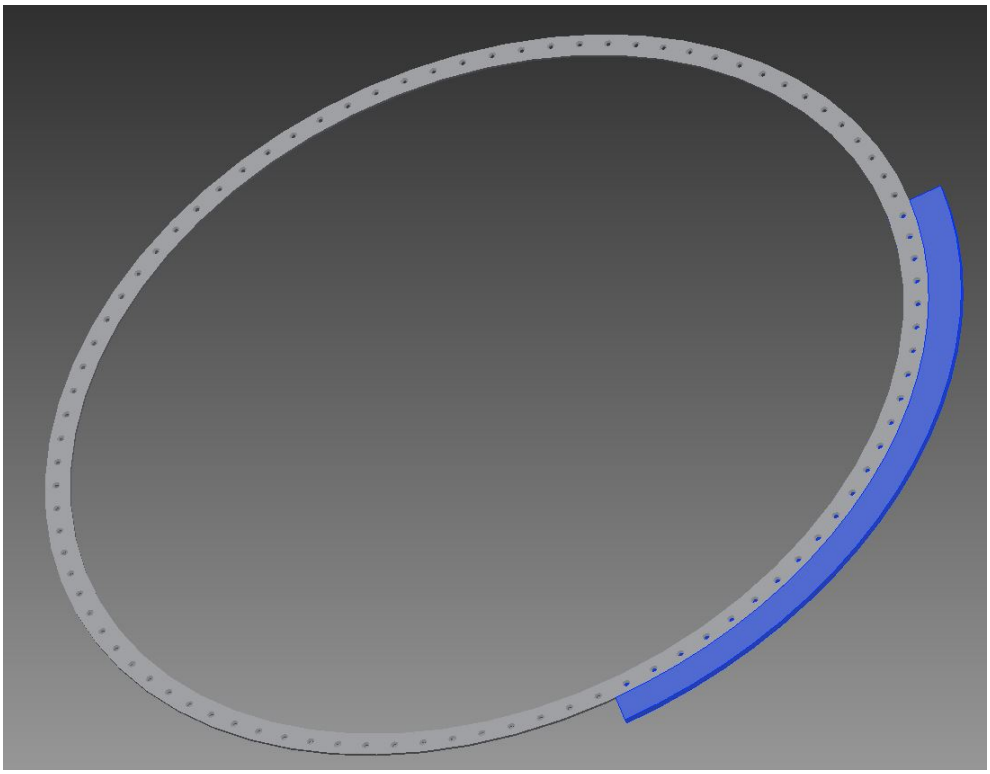

Figure 4-15: The cathode.

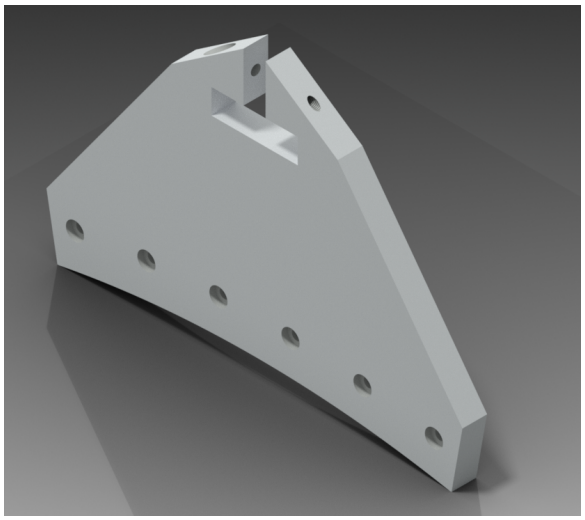

(a) Cathode top

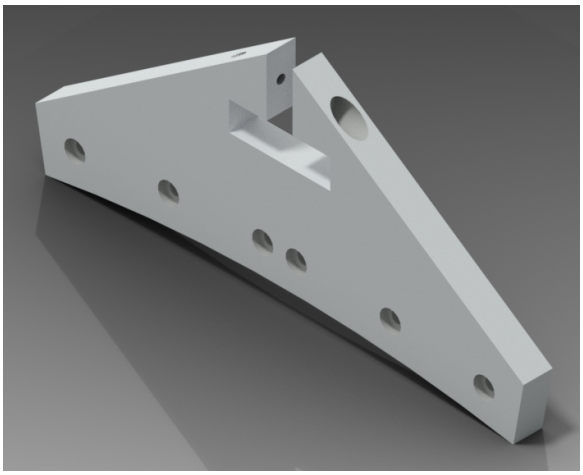

(b) Amplification top

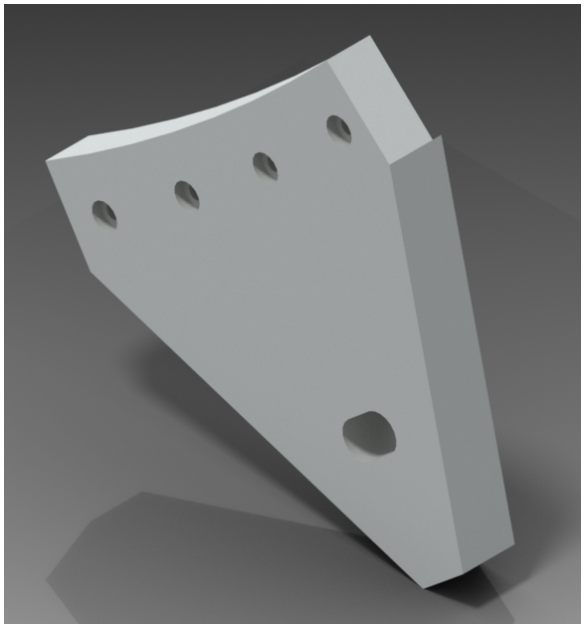

(c) Cathode side

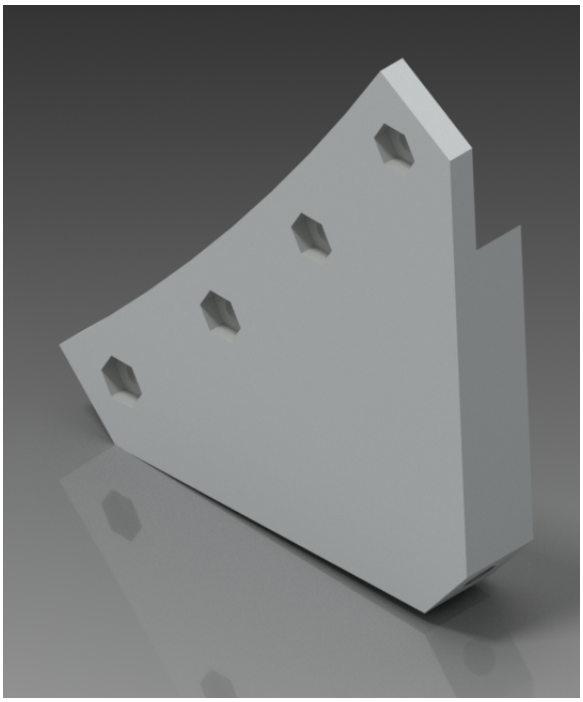

(d) Amplification side

Figure 4-16: CAD images of the Delrin support pieces for the cathode and amplification region. Each piece shows the hole pattern that was the starting point for the design. The two top support pieces (a and b) show the T-bar connection to the top slide rail, required for installation into the vacuum chamber. The cathode side support has an oval hole which allows for the rod which extends the length of the field cage. The side support for the amplification region also requires this hole, it was added during fabrication.

4.3.2 Field Cage Support

As mentioned, the change in TPC orientation and ring design required that a new support system be designed. For the ring support piece, inspiration was drawn from the EXO TPC [70] and the comb style housing encompassing their field rings. The initial design was a single long piece of Delrin in which the rings could be encapsulated. This included a recess for the rounded shape of the ring and a plate, with an extruded rounded shape, to secure the ring in place. Figure 4-17 shows the 3D CAD model for this design. The thick piece in the image is the outside piece and contains the recesses. The thin piece is the securing plate.

In order to hold the rings in the horizontal configuration, a supporting piece is required to extend to the wall of the vacuum chamber. Three rails were already installed into the vacuum chamber, one placed at the top and one on each side at 135°. The design was thus built around this existing mounting framework. The chosen design for the support rod was a clevis system, which consists of a central body into which a rod is screwed at each end, but with opposite handed threads. The central body is called the clevis and by turning it, the length of the rod and clevis construction can be increased or decreased. This allows for overcoming of tolerances in the fabrication and also allows a tension to be applied if required. Figure 4-18 shows a CAD image of this design.

These initial designs were prototyped in the DaRCO detector at Royal Holloway, as will be seen in section 4.5, and were very successful. However the attachment to the rail system was required to be slightly different in DaRCO compared to the $1m^3$ due to the orientation and style of the rails attached to the walls of the chamber. In DaRCO there are two extruded pieces in the horizontal plane, in line with the central axis, and a similar extruded piece at the very bottom of the chamber (see figure 4-19 (a)). These allow for a sliding mechanism to be used directly over the rail. The design for this was quite simple and involved a metal housing that fits over the extrusion

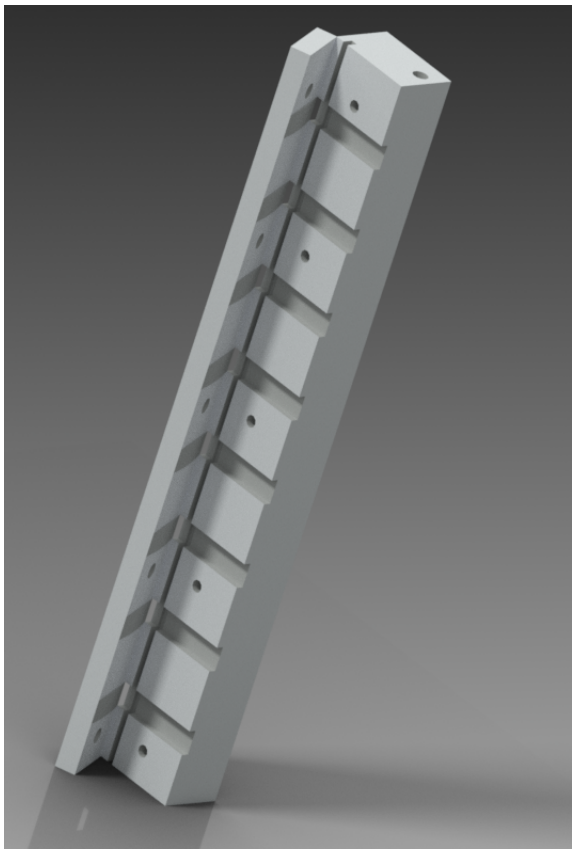

Figure 4-17: CAD image of the comb-style ring support piece. A recess for each ring can be seen along the length of the thicker part. The thinner, enclosing, piece shows the rounded extrusion. Along the length of the thick piece runs a channel in which the resistors would be housed.

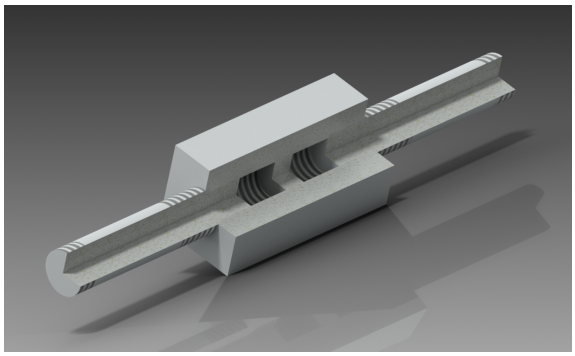

Figure 4-18: CAD image of the clevis support mechanism. This image shows a three-quarter cut-out view to reveal the inner mechanism.

along which threaded holes are placed to allow the clevis rod to be screwed into (see figure 4-19 (b).)

The rails of the $1m^3$ are placed at 135 °s from the top central point and there is a rail at the top of the chamber as opposed to the bottom. They also differ in that they are horizontally flat and have screw holes along their length to enable a support mechanism to be attached. Both rail systems were designed and in place prior to the support structure designs being initiated.

Figures 4-20 to 4-22 show the initial conceptual ideas that were put forward. Common to each design was the angle of the support rod to the chamber wall. This was made normal to the radial direction for uniform support. Figure 4-20 shows a pillow block design, which is a commercially available unit composed of a steel outer shell with a Delrin inner housing, allowing for a rod to be smoothly slid along the length. Unfortunately these were not feasible due to cost constraints. The design in figure 4-21 was inspired by the pillow block and was intended to be made entirely of Delrin. A T-shaped aperture was proposed for which the the support rods would have the complimentary shape, such that they would slide along the length of the chamber. Due to the complicated shape of this design, machining would have been difficult. Therefore the next iteration (figure 4-22) broke this down into a two-piece component with a bar connecting to the chamber rail and the sliding section slotted in behind it. With these concepts in place, the designs were passed to a group of engineers for optimisation based on their experience of machining feasibility and detector construction. The results of which will be shown in section 4.6, after the prototyping of these initial designs has been demonstrated (section 4.5).

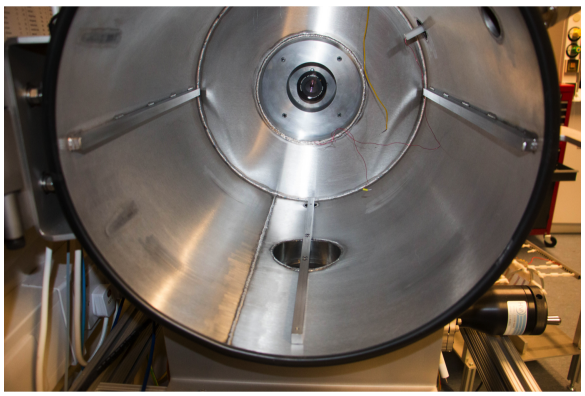

(a) DaRCO rails

(b) Slide mechanism

Figure 4-19: In (a) is a front view of the inside of DaRCO. The side and lower rails can be seen attached to the walls of the chamber. (b) shows the slide rail mechanism that was employed to support the field cage.

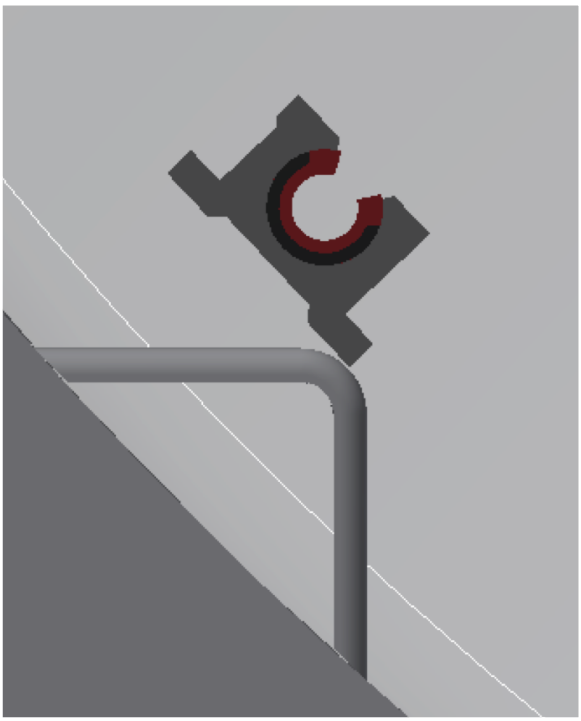

(a) Pillow block

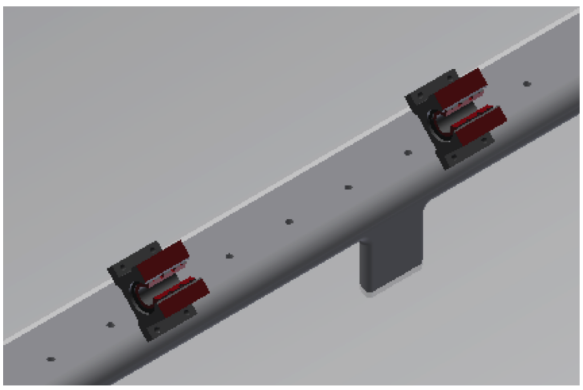

(b) Pillow block

Figure 4-20: CAD images of the pillow block rail design (a) front on and (b) from an angle. The blocks appear to be free floating - had the design been chosen a support wedge would have been put in place to hold them at the angle shown.

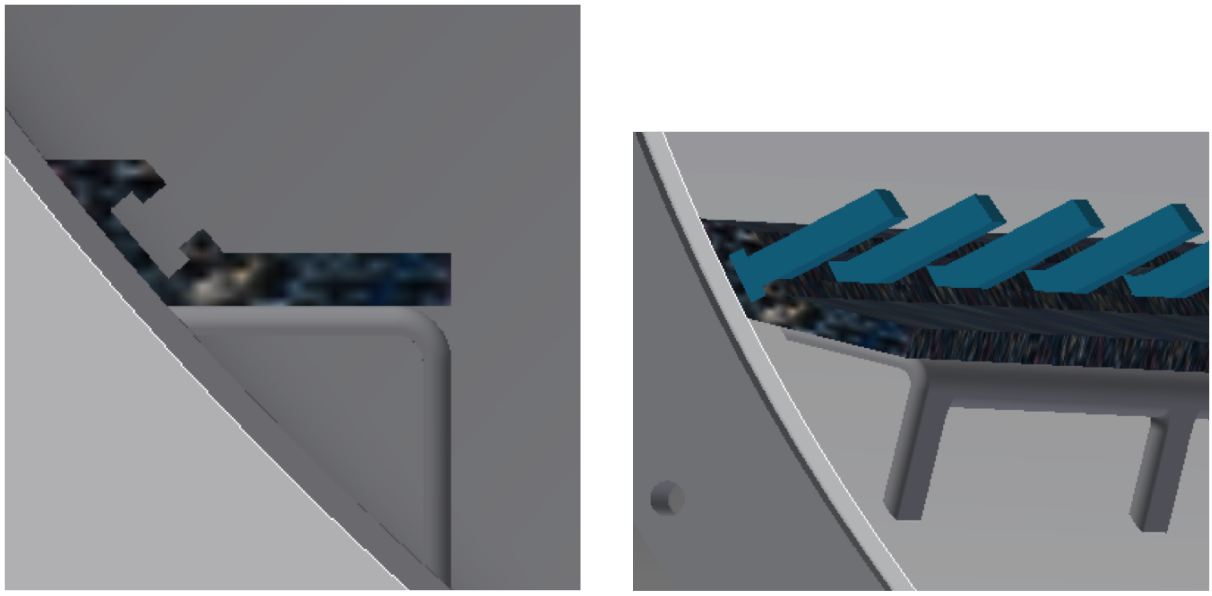

(a) Delrin T-bar (b) Delrin two-piece

Figure 4-21: CAD image of the T-bar rail design (a) front on and (b) at an angle and with support rods included. The colours are artificial to make the parts clearer to see against the grey of the chamber.

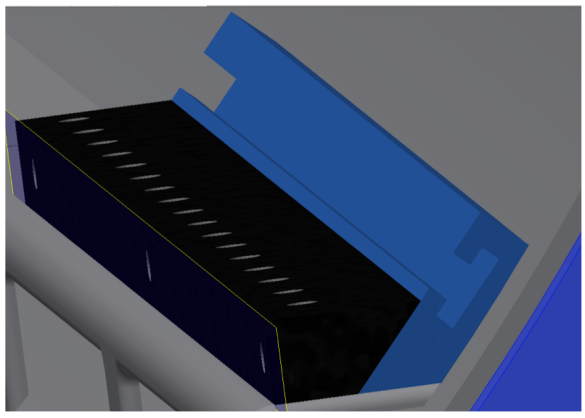

Figure 4-22: CAD image of the T-bar rail design modified to be more feasible to machine. The colours are artificial, to highlight the components. The black piece screws onto the existing rail within the chamber. The blue piece then slides behind it.

4.4 Post Analysis

For the $1m^3$ detector, the amplification region was segmented into four quadrants. The reason for this is that at this scale, an individual amplification region would have very high capacitance, owing to the linear proportionality of capacitance with electrode area: $C = \epsilon A/D$, where C is the capacitance, ϵ is the dielectric constant of the material between the electrodes, A is the surface are of the electrodes and D is the electrode separation distance. A high capacitance would affect the fast pulse readout electronics and diminish the ability to distinguish between electronic and nuclear recoils (see section 2.1.1). To support the individual quadrants, a Delrin cruciform frame was designed. Due to the need for the frame to provide sufficient support, the width and thickness was required to be considerable and will thus have an impact on the field within the TPC - the dimensions are shown in figure 4-25.

To quantify the effect and to see the impact on the achievable fiducial volume, a three-dimensional model was constructed with the relevant geometry. To keep the file size and computation time down, this was done for just a single ring. A three-dimensional model is required in this case due to the geometry of the cruciform - an alternative would be to take two-dimensional slices at selected angles, however as only a single model was required to be studied, a three-dimensional version is not unreasonable to build. Figure 4-26 shows the Gmsh geometry of the cruciform and figure 4-27 shows two slices of the electric field map in the r-z plane. In figure 4-27 the left hand side intersects the middle of a bar of the cruciform shape and the right hand side is at 45°s to the bar, to demonstrate the impact of the plastic at each location. Figure 4-27 shows a close-up view of each slice, near the edge of the field cage.

Before considering the impact of the cruciform, this model can be used to guide the fabrication tolerances for the physical build. Figure 4-23 shows the field uniformity, U_F, as defined by equation 4.1, at the edge of the 45° slice. Figure 4-24 plots the value at the location of the white line of figure 4-23, where the non-uniformity has

the greatest extent into the drift region. A one percent non-uniformity is reached at a radius of 57.3 cm, this is approximately one ring spacing away from the ring radius. By setting the tolerance to half a ring spacing, or to be conservative and explicit, 1 cm, the extent of the non-uniformity should not impact on the fiducial region.

$$U_F = \frac{E_x^2 + E_y^2}{|\bar{E}|^2} \tag{4.1}$$

With a full 3D model, the generated field map can be imported into Garfield[9] [77]. Garfield is a piece of software that is specifically designed to simulate drift chambers and can simulate the transport of electrons through a particular gas at a particular pressure in a given field configuration. By using Garfield, the impact of the cruciform on the path of a particle can be seen. To do this, electrons were generated in a line at the top of the field cage and drifted to the ground electrode. Figure 4-31 shows the results of two cases. The first with the electrons generated along a line perpendicular to one of the arms of the cross part of the cruciform shape. The second is with the electrons generated from the centre to the outer edge of the field cage at a 45° angle of a quadrant. The location of these lines are represented in figures 4-29 and 4-30. The extent of the material covering the ground electrode in a single quadrant is as follows: cruciform bar: 2 cm; outer radial edge: 4.66 cm; centre corner: 2.83 cm. Figure 4-31 shows how the electrons survive right up to the edge of the plastic. As a safety margin, the 100% survival point is considered to be 0.5 cm away from the edge, allowing for the 0.5 cm resolution of the mesh. Using these values, the reduced fiducial volume is calculated to be 0.892 m^3, figure 4-33. Figure 4-31 also shows a plot of the radial deflection of the drifted electrons, with the end radius plotted as a function of the start radius. Figure 4-32 uses this plot to show the difference in the electron starting radius to its end radius. Though this does not give information

[9]Garfield is capable of using the output from 2D models, however it is not currently configured to accept the 2D output of Elmer

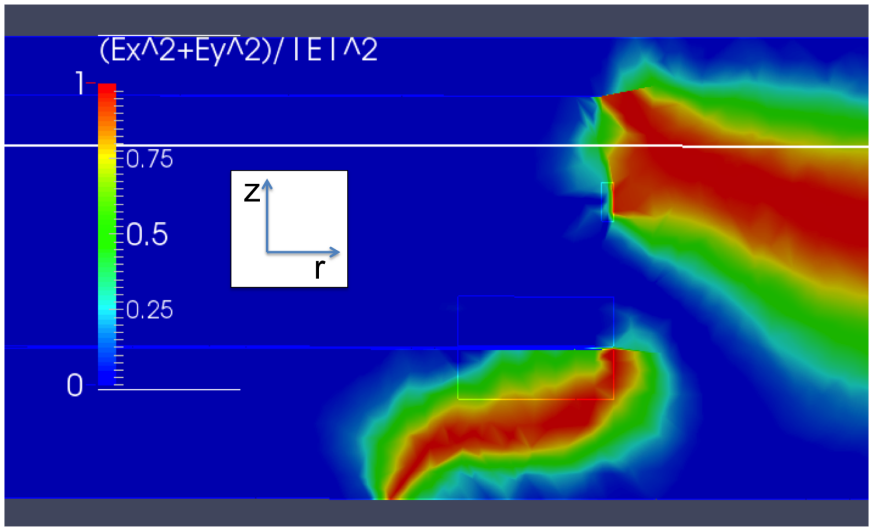

Figure 4-23: The field uniformity, as defined in equation 4.1, at a 45° slice of the 3D model . The scale of the uniformity is in a percentage. The location of greatest non-uniformity lies between the ring and the cathode.

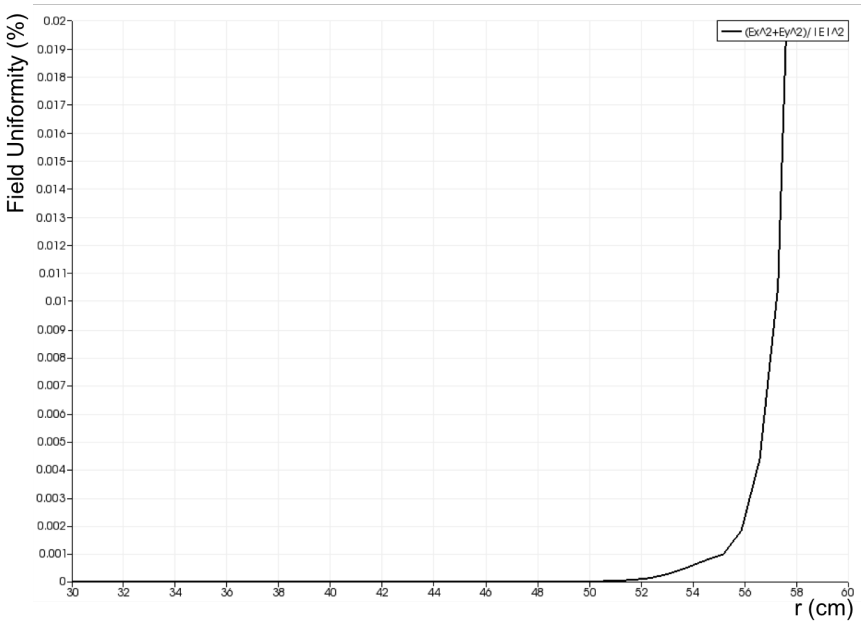

Figure 4-24: A plot of the field uniformity for values over the horizontal white line of figure 4-23. The abscissa is the radial dimension, in units of cm, the ordinate is the uniformity as a percentage. The point of one percent uniformity can be seen at $r = 57.3$ cm.

about the diffusion of an electron cloud, it does indicate that the radial deflection near the field rings is slightly less than that indicated in figure 4-12: at 54 cm, the difference has a maximum of 4 mm. It also shows that the median value for the radial drift is lower at 2 mm.

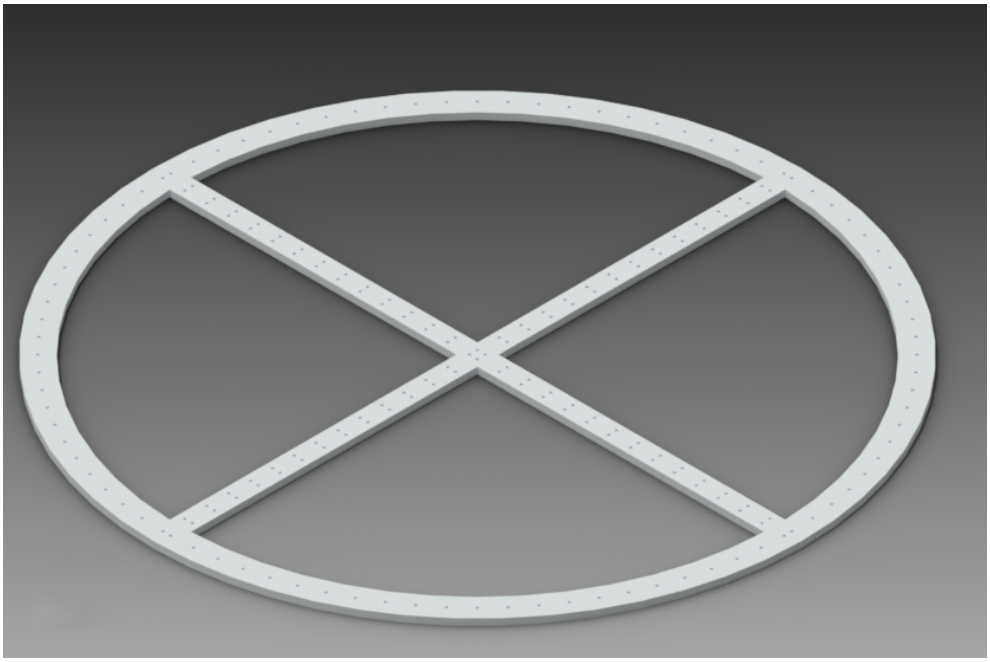

Figure 4-25: CAD image of the cruciform structure designed to support amplification region quadrants. The thickness of the outer circular part is 4.66 cm and the thickness of the bars of the cruciform shape is 4.00 cm.

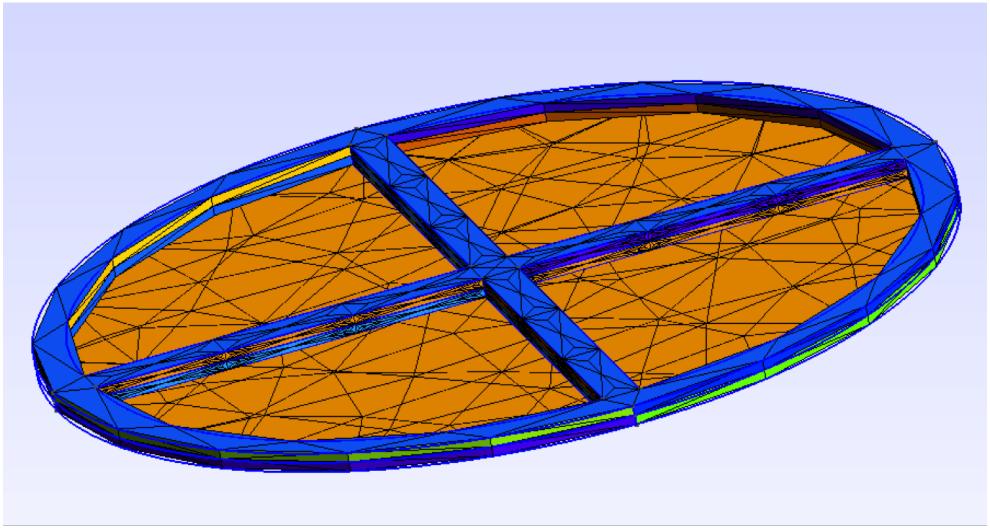

Figure 4-26: Gmsh geometry showing the cruciform frame used for the segmented amplification region. This was modelled at the full dimensions given in the drawing of figure 4-25. The plane of the orange quadrants is the radial direction and the normal to this plane is the z-direction.

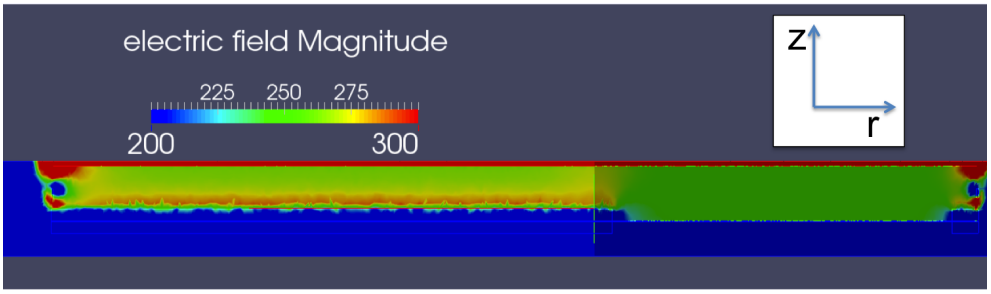

Figure 4-27: Two slices of the 3D model. The left hand side, which is slightly brighter, shows a head-on view that intersects the cruciform structure. The right hand side shows a 45° angle slice, showing the impact of the cruciform at this point. These are shown in close up in figure 4-28. The field magnitude is given in V/cm.

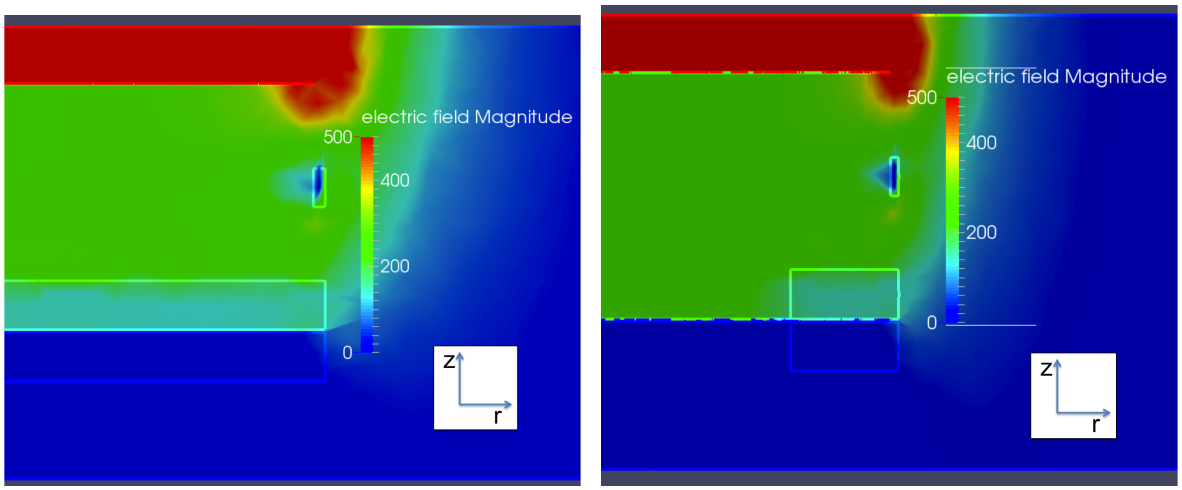

(a) (b)

Figure 4-28: In (a) is a close up of the left hand side of figure 4-27, where the cruciform is intersected at the midpoint. Note that the image has been mirrored along the vertical, for easier comparison to the opposite side. The highlighted lines represent the outlines of the geometry. The impact of the cruciform can clearly be seen by the light blue shading in the highlighted box in the drift region. In (b) is a close up of the right hand side of figure 4-27, where the field cage is intersected at a 45° angle. The light blue shading within the highlighted cruciform shape shows the variation in the electric field magnitude caused by the plastic. The impact on the inner drift region away from the bar is clearly much diminished. In both figures, the field magnitude is given in V/cm.

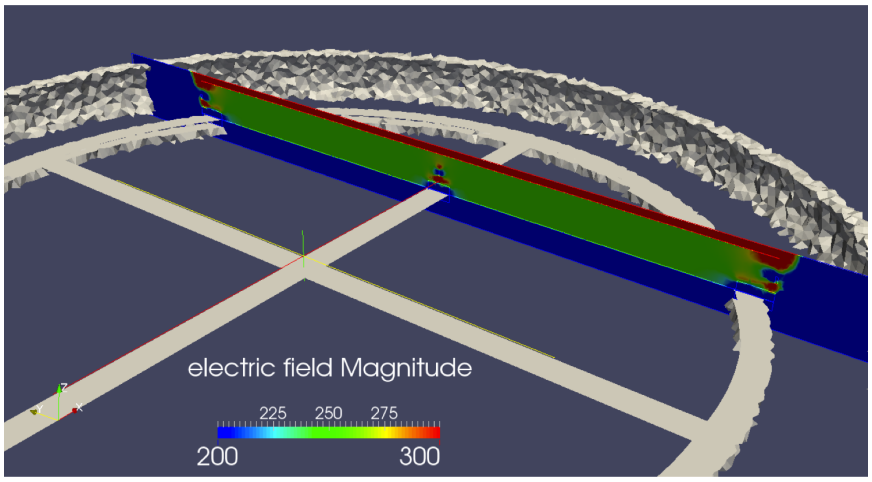

Figure 4-29: This figure demonstrates the location of the line along which electrons were generated in Garfield, for the first case mentioned in the main text. The off-white parts of the geometry show the outline of the cruciform and the outer edges of the enclosing chamber. The view is looking slightly above and down at the chamber. The coloured plane is the location at which the line of electrons was generated. The colours represent the electric field in this plane and the scale is given in V/cm.

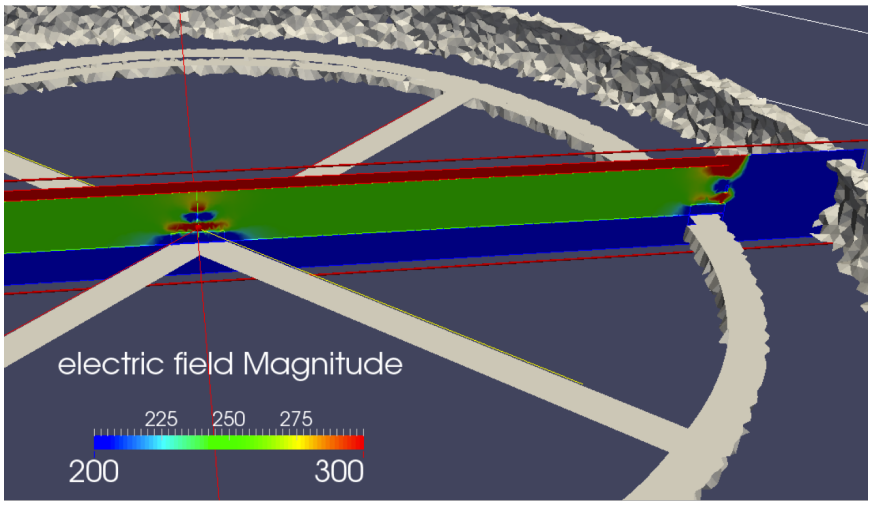

Figure 4-30: This Figure is the same form of representation as for figure 4-29, but for the second case given in the text. Here the line of electrons was generated at a 45° angle from one of the cruciform arms, to show the impact of the cruciform in a radial sense. The colours of the plane represent the electric field and the scale is given in V/cm.

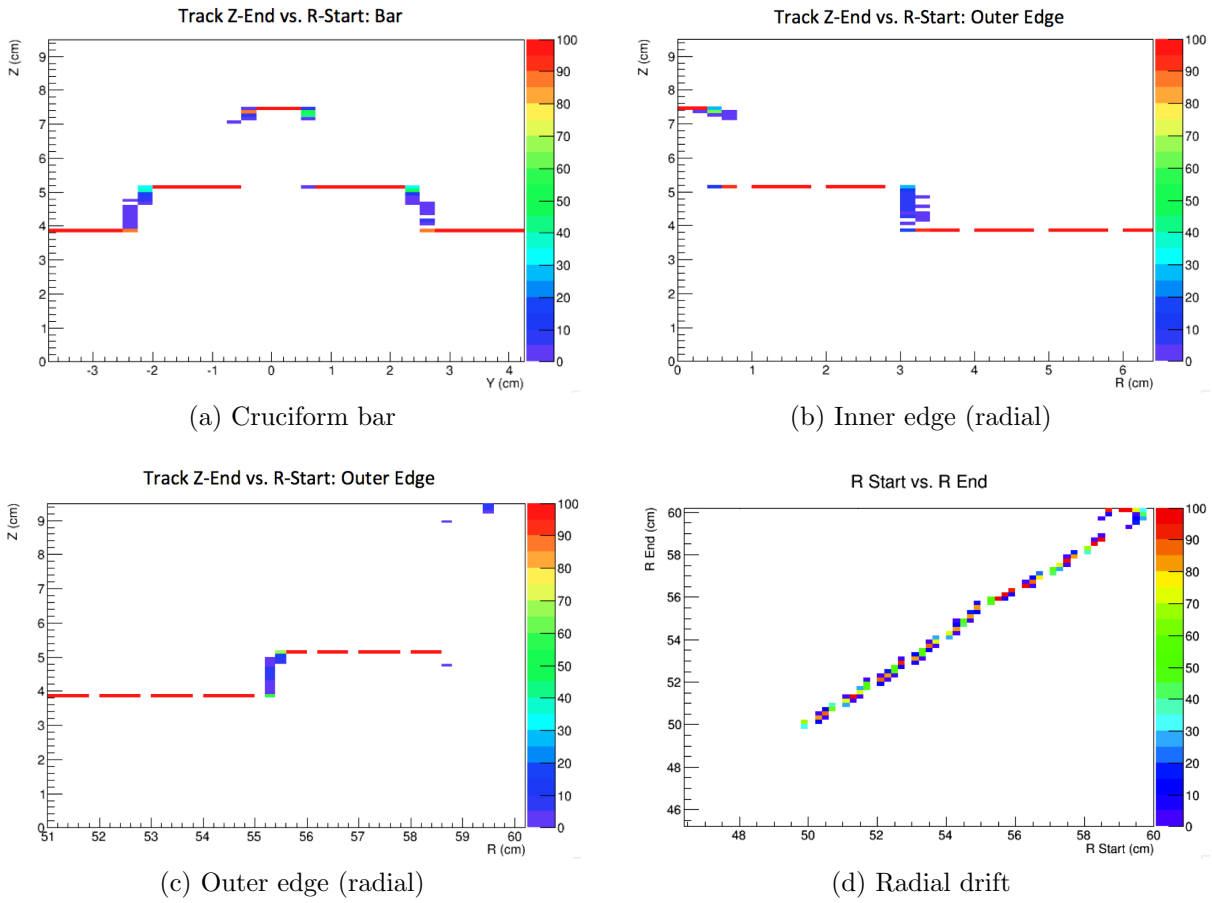

(a) Cruciform bar

(b) Inner edge (radial)

(c) Outer edge (radial)

(d) Radial drift

Figure 4-31: Garfield drifting results. In figures (a) to (c), the abscissa is the radial location in the detector, the ordinate is the height in the drift region - the ground electrode of the amplification region is located just below 4 cm. The applicate represents the number of electrons that drifted to that location. For example, in (a) at $y = 3$ cm, all electrons generated at this radial location, just below the cathode, successfully drifted to the ground electrode and will thus be amplified and readout as signal. Whereas at $y =$ 2.5 cm fewer electrons survived. The shape seen in (a) (excluding the applicate values between 7 and 8 cm - explained shortly) roughly outlines the cruciform bar. Figure (a) shows these results for the line of electrons generated across the bar of the cruciform, relating to figure 4-28 (a). The central region has a higher applicate value (between 7 and 8 cm) as this model also tested segmenting of the ring for additional support. This was not used and does not affect the result of the simulation. Figure (b) shows these results across the radial inner edge for the 45° line and (c) is for across the central corner. (b) and (c) both relate to (b) of figure 4-28. Figure (d) shows the radial location at the end of drifting versus the starting point. The applicate is again the electron count.

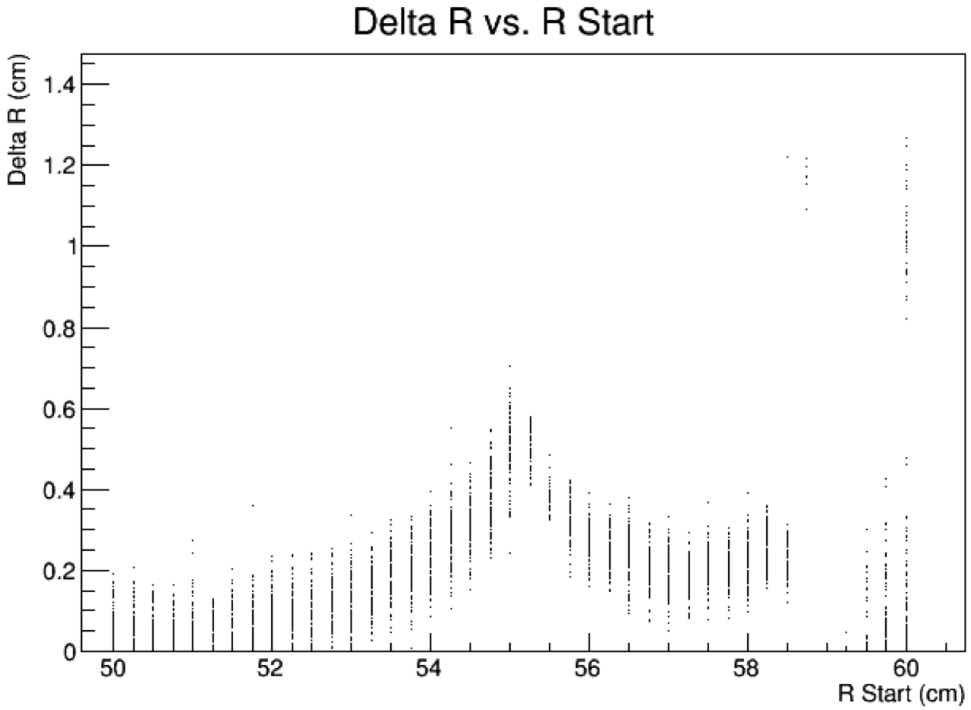

Figure 4-32: The difference in electron radial start position to end position, plotted as a function of the starting radius. The deflection at 54 cm is 4 mm. The peak at 55 cm is where the plastic of the cruciform starts to impact the result.

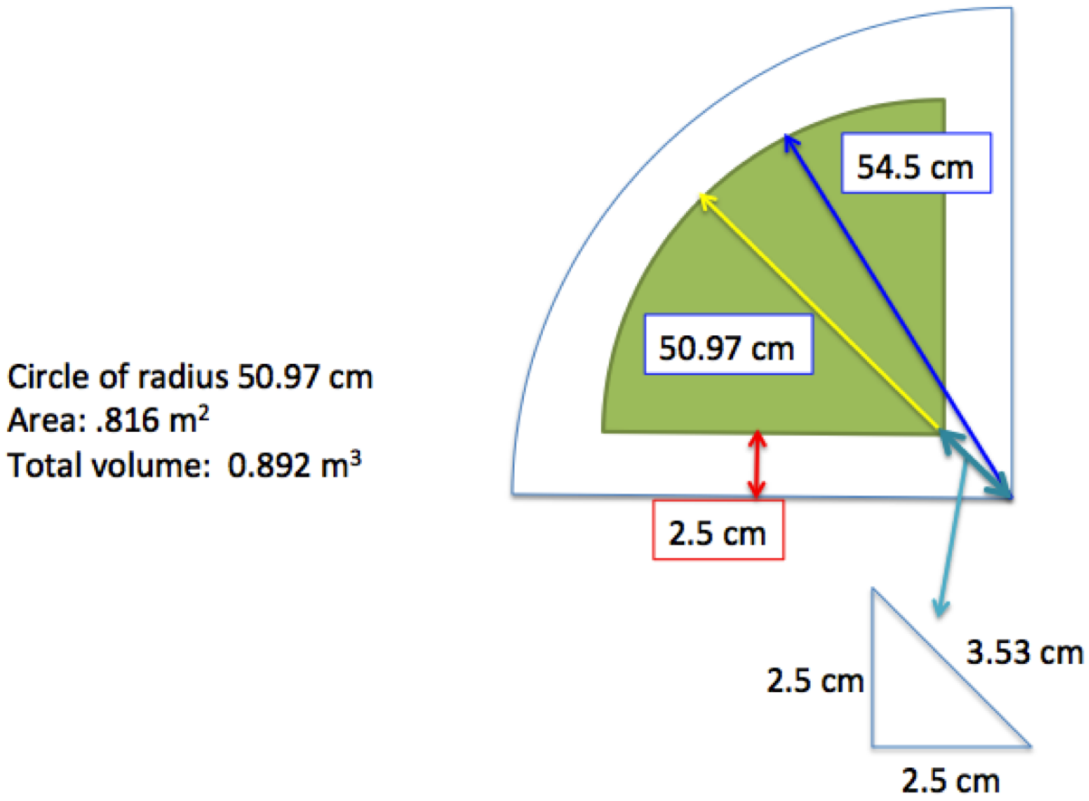

Figure 4-33: The fiducial volume calculation considering the cruciform support frame for the amplification region. The green quadrant represents the area for which electrons survive in the post-analysis simulation. The blue quadrant represents a quarter of a field cage ring.

4.5 DaRCO

DaRCO: Dark matter Research Chamber Optimisation is the name given to a prototype detector that was built at Royal Holloway and which, as the name suggests, was used as a place to test the new field cage design for the large scale detector. It started as an empty vacuum chamber on a stand and over time was built up to contain a working field cage with CF4 gas plumbed in, slow control electronics and a CCD camera connected.

4.5.1 Purpose and Comparison to Other Detectors

With the design requirements for the new $1m^3$ detector being very different to its predecessors, DaRCO was used as a means to test new ideas. The main difference is the chamber orientation - all previous detectors have had the TPC upright from amplification region to cathode. This makes it quite straightforward to simply stand the TPC in the chamber, with no support to the walls required. In the new horizontal configuration, wall supports are essential in order to keep the field rings of the TPC out of contact with the chamber walls. The reason for this is that the rings are held at high voltage and without sufficient stand off distance from the chamber walls, sparks can occur. For this reason also, the support system must be made of a dielectric material, in this case Delrin was used.

4.5.2 Field Cage and Amplification Region

Initially the field cage rings were not expected to change and so a field cage was made with the original design. At this stage, however, a triple mesh design for the amplification region had already been considered. In previous detectors the anode of the amplification region was made up of a copper plated disc of G10[10]. The idea

[10] An epoxy glass

for the $1m^3$ was to use the back to back TPC concept of the 10L detector, but to make the anode out of mesh, allowing both sides to be imaged with a single camera. Therefore, a triple mesh was made, based on an existing design in the collaboration. At this point, the new support system design was also introduced. Inspiration was drawn from existing tension mechanisms and the clevis system was chosen as a viable option. In this system, two oppositely threaded rods are screwed into a central 'clevis' connector. The opposing threads cause the overall length of the clevis-rod system to contract or extend as the clevis is turned. The test stand piece for this design can be seen in fig. 4-34.

When the suggestion of using a rotated ring geometry was made a prototype was fabricated and can be seen in fig. 4-35.

As can be seen from the two images, the prototype constructions were successful and gave the group confidence in going forward with the new design. It was shown that the TPC could be securely held in place using the clevis system and the rotated ring design allowed for a much simpler support mechanism.

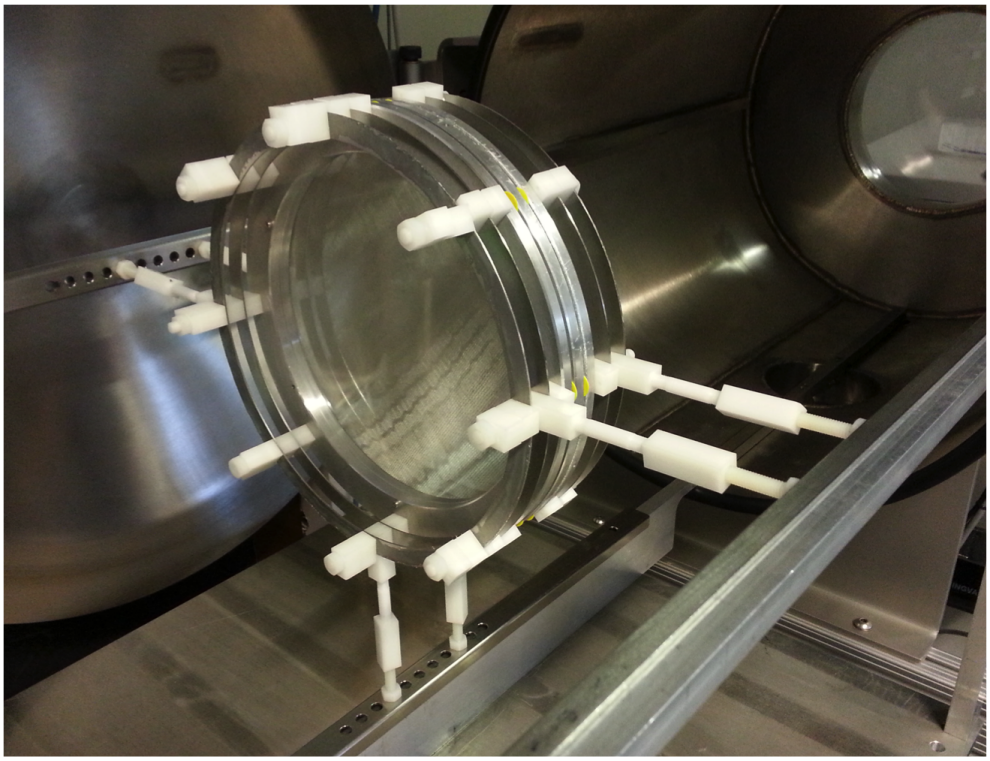

Figure 4-34: Prototype field cage based on the original ring design. In this construction, two back-to-back TPCs were built. A triple mesh amplification region can be seen as the thicker, lighter, central region of the construction. This image also shows the clevis system being used to support the field cage.

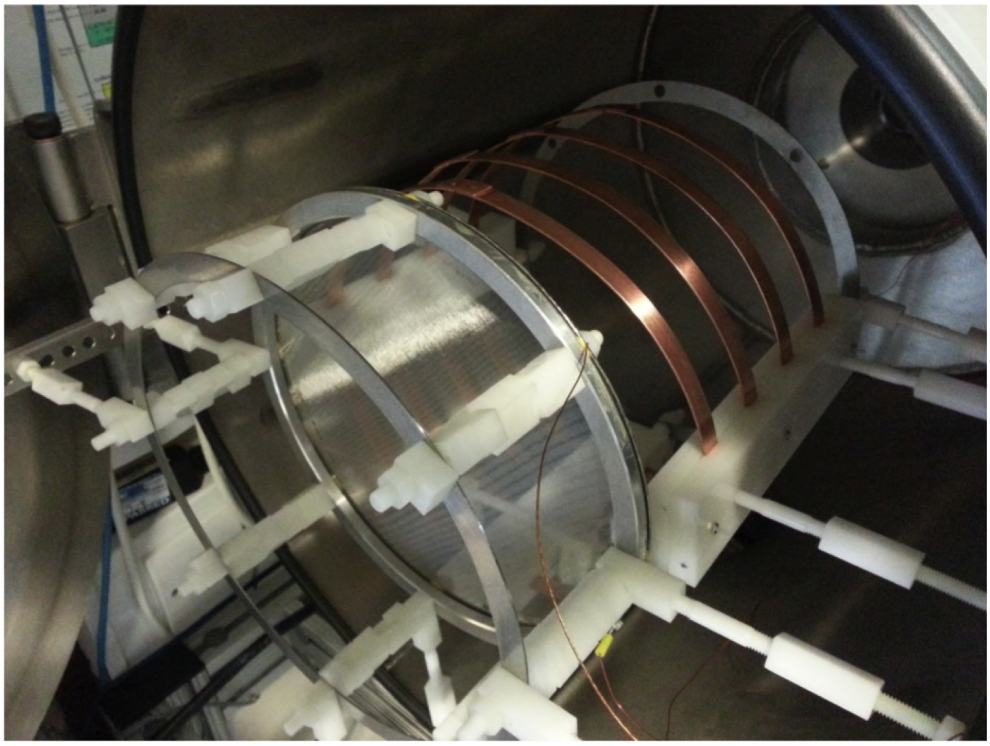

Figure 4-35: Prototype field cage based on the new ring design. For this a single TPC was made. The copper rings show the new construction, including the first iteration of the ring housing mechanism.

4.5.3 Construction

The parts purchased for the slow control were the same as those for the 10L detector, to allow easier porting of the existing software for operation. A National Instruments PCI-6221 card, connected to the BNC-2090A breakout box was used to read in and control the voltages on the high voltage supply and a Synaccess Netbooter was used to automate the valve operation. Two gauges were also connected to the PCI card, an ionisation gauge for the sub-Torr pressures and a convection gauge for the higher pressure range. The slow control was configured to readout the pressure from the gauges and to switch to reporting the correct gauge readout dependent upon the pressure range. The voltage supplies were connected and were configured to some degree, however the full slow control was not fully implemented as there were difficulties in reliably ramping up and down. Additionally, multiple computer crashes meant that full integration was not possible at the time. The high voltage was also connected to the chamber, allowing voltage to be applied to the amplification region and the cathode. A camera was also placed onto the window at the end of the chamber allowing for images to be acquired.

4.5.4 Operation

DaRCO was tested using both the original and the new prototype field cage. Both were successfully connected electrically, with 1 MΩ resistors between rings to provide the smooth voltage drop to the amplification region. The original prototype was operated at a drift field of $\sim$ 200 V/cm and the new prototype was operated at $\sim$ 100 V/cm. The reason the new prototype was at a lower drift field is because of the increased drift length. The amplification region was also successfully connected and high voltages - comparable to those used for the 10L - were applied to the anode, with the expected level of sparking.

Figures 4-36 and 4-37 show events seen in CCD images. As DaRCO was intended

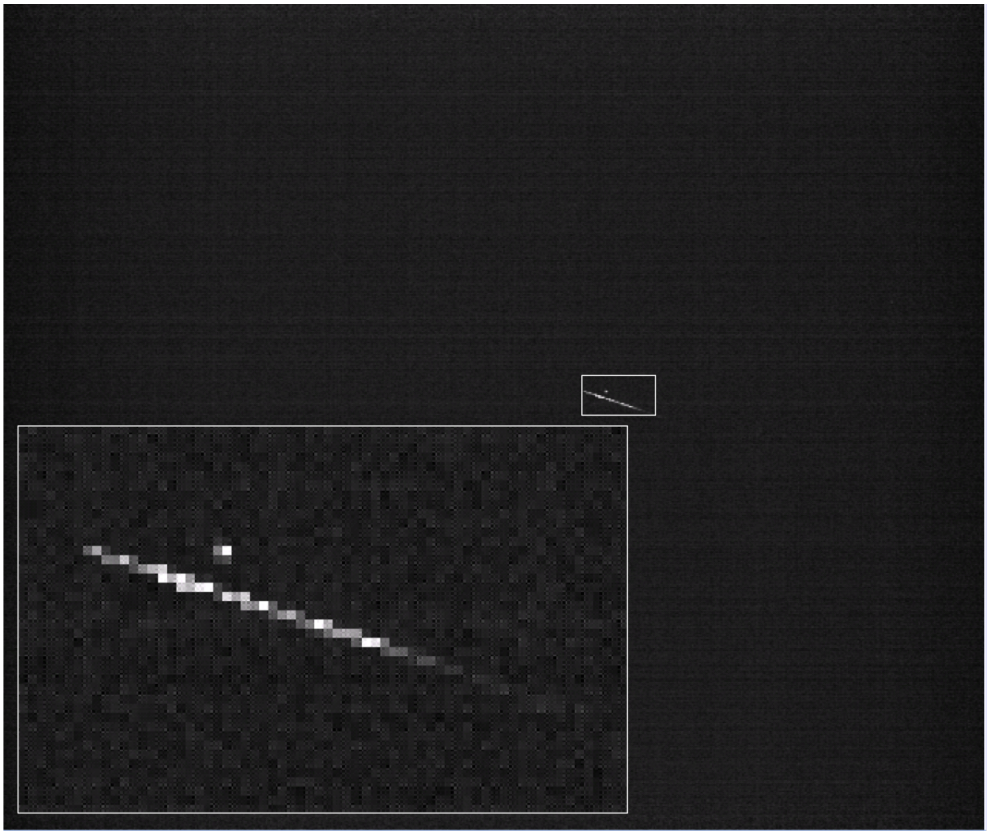

Figure 4-36: A background event seen with the original field cage design.

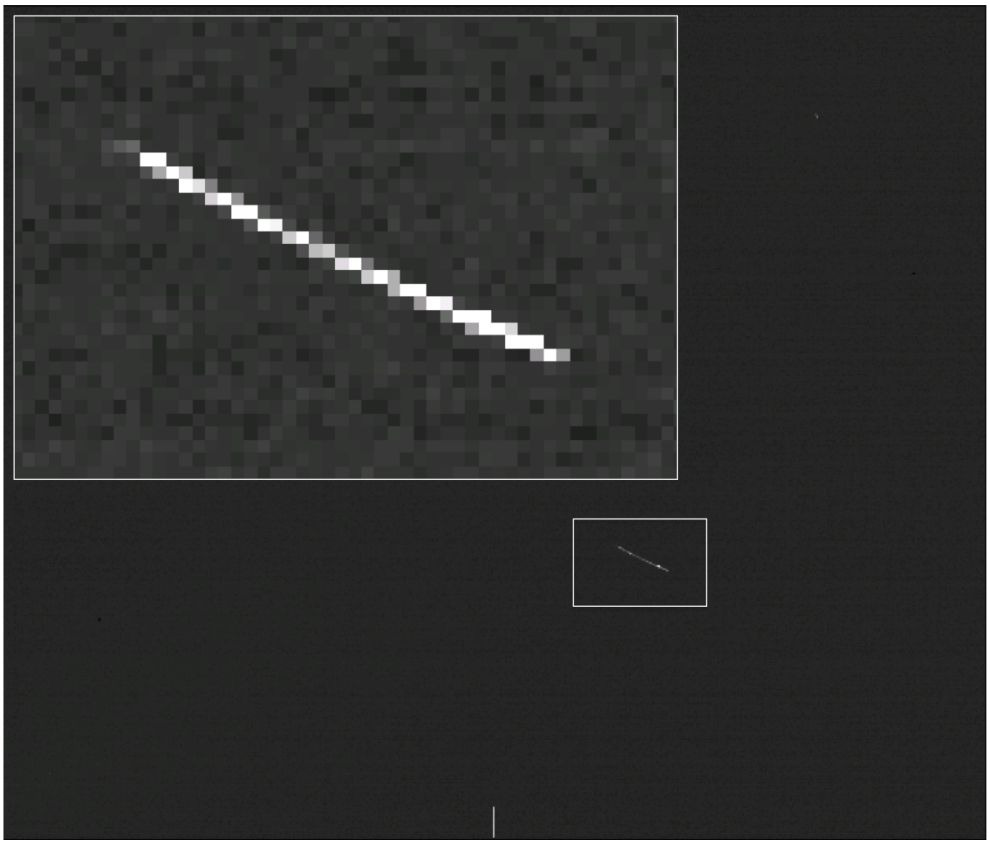

Figure 4-37: A background event seen with the new field cage prototype.

to be used predominantly as a mechanical prototype, no analysis was performed. DaRCO was a useful and successful test for the mechanical design of the new style field cage proposed for the $1m^3$ detector. It also provided the opportunity to work on constructing a gaseous detector of this type, giving experience of operation with gas and the plumbing thereof, communication with hardware instrumentation and remote control via a computer.

4.6 Optimised Field Cage Design for the $1m^3$ Detector

After successful mechanical prototyping of the new field cage assembly, the design was refined with the assistance of the engineers at Bates Laboratory. The end result was a more modular construction in which each ring has its own separate holder and all are connected by a rod extending along the length of the field-cage in the z-direction. With a thread at each end, a nut can be screwed onto each end of this rod, securing the construction. The modular design improves the construction process by allowing for one ring section to be made at a time, rather than having to construct the field cage as one entire piece. By this stage the resistors had been selected. The type used are made by Ohmite, model RX-1M. The resistors have a value of 3 MΩ and a housing length of 47.63 mm. These resistors were chosen for their high accuracy (1%), stability and zero outgassing, making them ideal for use in vacuum and low pressure gas. To fit these between the 25.4 mm spacing, each ring support piece has a diagonal channel to fully house the resistor and its wires. The design for the slide rails was a hybridisation of the pillow block and T-bar design. Each rail was designed to screw into the metal rail of the chamber and has a cylindrical, key-hole like aperture along the length, see images 4-39, 4-40. The round aperture was chosen to minimise friction when sliding. Initially the rail was designed such that the sliding component

would lay horizontal and would behave as an alignment mechanism only. This was done to avoid over-constraining the system. However it became necessary to modify this and revert to the normal-to-radial orientation after an initial prototype build, more of which is discussed in section 5.1.

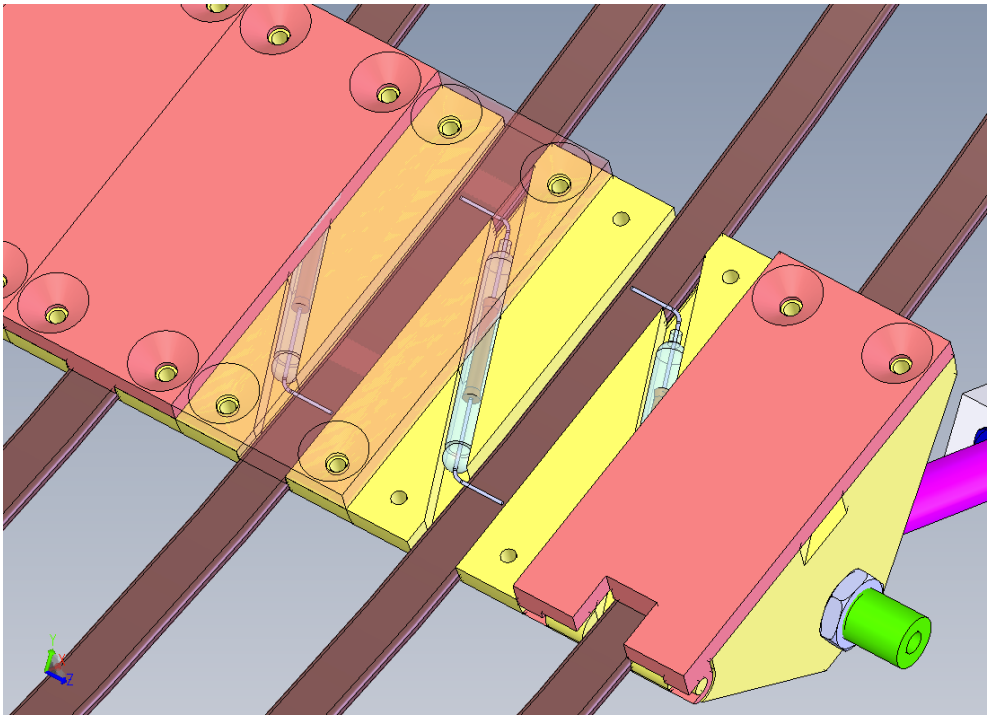

Figure 4-38: CAD image of the side support piece with a diagonal channel to house a resistor. The yellow piece is the outer piece with the ring recess. The red piece is a cap to hold the ring in place. The green piece is the securing rod. For scale, the red block has a 5 cm square cross-section.

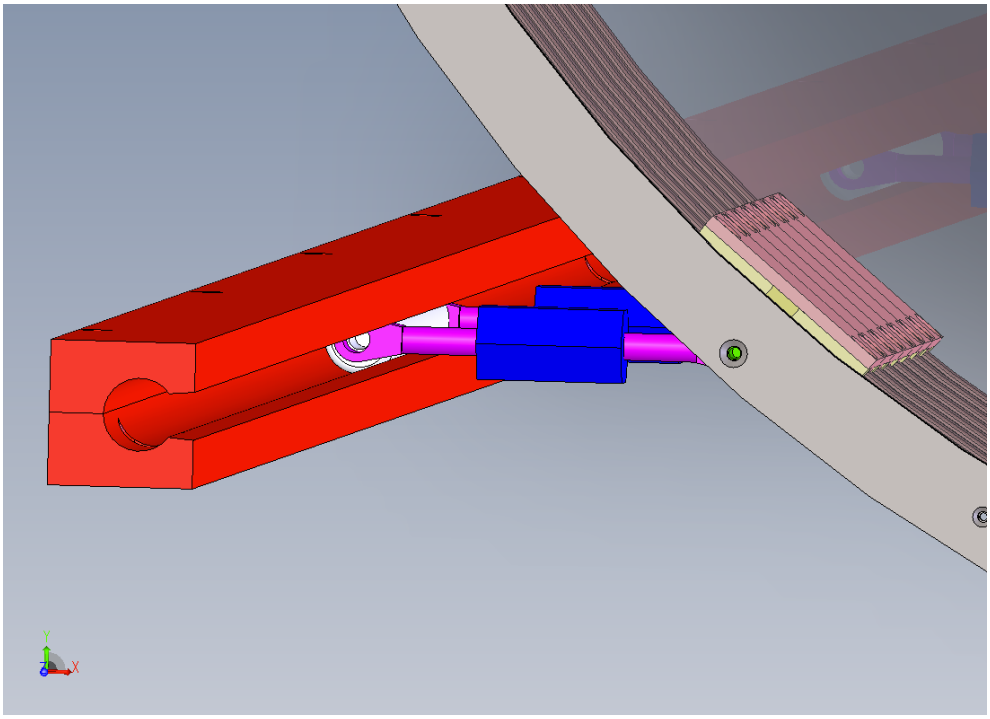

Figure 4-39: CAD image of the side rail design. The red piece connects to the vacuum chamber rail, the pink and blue pieces make up the clevis mechanism and the white piece is a bar that slides along the length of the rail.

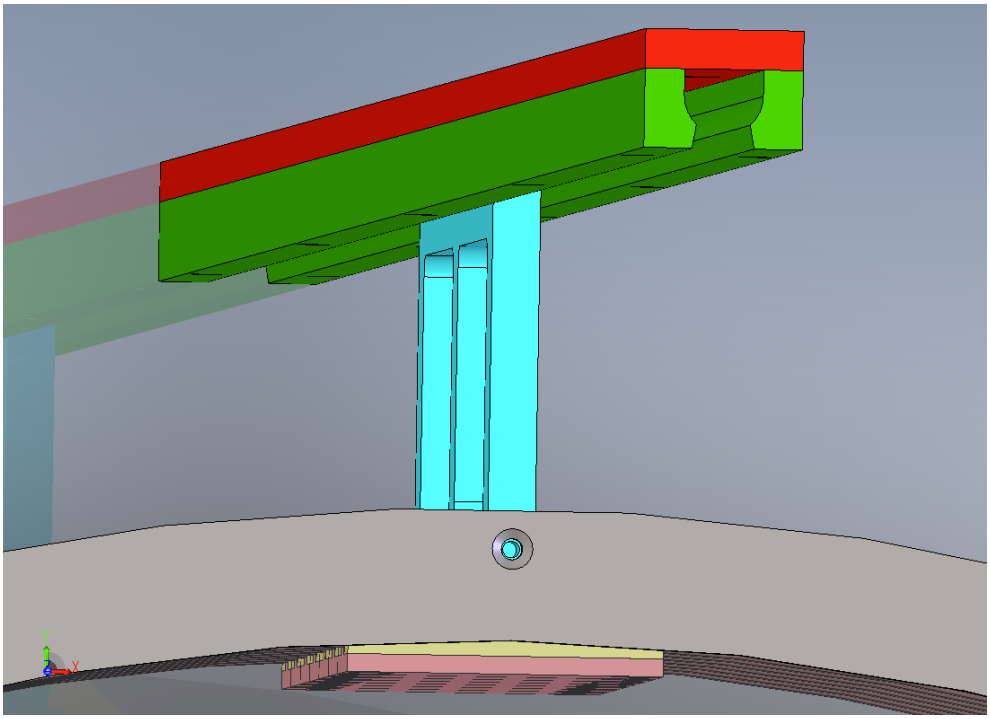

Figure 4-40: CAD image of the top rail design.

4.7 Results

With the final configuration chosen and a full 3D tracking simulation performed, a comparison of total surface area to fiducial volume is made for this and the previous detectors. First shown is the calculation including all support parts, except the amplification region support frame - as this was not part of the original design and simulation. The contribution of the frame is then included. As will be seen, a considerable improvement has been achieved in each case. Firstly, table 4.2 breaks down the surface area contributions to the field cage and its support structure. Then table 4.3 does the same for the four shooter detector. Finally table 4.4 lists the contributing surface areas for components of the support cage and the total field cage. The largest contributor to the surface area from the support cage is the ring holding pieces. These have been listed separately to highlight this. Following the breakdowns is a comparison of the surface area to fiducial volumes of each detector, in table 4.5. In this table, it is clear to see that the new field cage design has been very successful in improving the total sources of surface backgrounds in relation to the fiducial volume which is being observed. Without the quadrant this represents a 28 fold improvement over the 10L and a 10 fold improvement over the four shooter. Including the frame reduces this to a 20 fold improvement over the 10L and a 7 fold improvement over the four shooter. For comparison, the improvement in surface area to volume from the 10L to the four shooter was only 3 fold.

Table 4.2: A list of the surface area contributions for the field cage and support structure of the 10L detector.

Parts	Material	Surface Area (m^2)
Field rings	Copper	1.832
Field ring spacers	Plastic	0.204
Ground mesh	Steel	0.048
Cathode mesh	Steel	0.028
Anodes	Copper	0.115
Anode plate edges	G10	0.009
Total		2.235

Table 4.3: A list of the surface area contributions for the field cage and support structure of the 4Sh detector.

Parts	Material	Surface Area (m^2)
Field rings	Copper	1.056
Field ring spacers	Plastic and copper	0.036
Ground mesh	Steel	0.040
Cathode mesh	Steel	0.028
Anodes	Copper	0.115
Anode plate edges	G10	0.095
Total		1.36

Table 4.4: A list of the surface area for the field cage and support structure. The values are obtained using the CAD components and the Autodesk Inventor software. Only the exposed faces are considered.

Parts	Material	Surface Area (m^2)
All support pieces	Delrin	1.592
Field rings	Copper	2.556
Anode and Ground mesh	Steel	3.324
Cathode mesh	Steel	0.785
Initial total		8.255
Quadrant frame	Delrin	1.492
Overal total		9.747

Table 4.5: A list of the surface areas, volumes and ratios thereof for the three main DMTPC detectors. The new design, though it has a large overall surface area, has reduced the surface area to fiducial volume ratio significantly.

Detector	Surface Area (m^2)	Fiducial Volume	Surface Area/Volume
10L	2.2356	0.00999	224
4Sh	1.36	0.01789	76
$1m^3$ no qudarant	8.255	1.09168	8
$1m^3$ with quadrant	9.747	0.892	11

4.8 Summary

This chapter has demonstrated a new simulation-based design, built on the premise of minimising background events due to RPRs emanating from the surface of the field cage materials, whilst maintaining the required fiducial volume. This was then followed by the outlining of the additional components designed as part of this project, enabling the full TPC to be constructed, first externally in the construction frame and then finally in the vacuum chamber itself. A full version of the CAD for the detector, containing a single TPC, can be seen in figure 4-41 which shows all the parts that were designed, fabricated and installed successfully - excluding the amplification region (shown), which was managed by another collaborator (Dr. Michael Leyton). As demonstrated, the new design has satisfied the stated requirements, reducing the surface area to fiducial volume ratio by a factor of 9.5 in comparison to the previous iteration and a factor of 28 to the 10L detector, used for the analysis presented in this thesis. It has also been successfully operated, as will be outlined in chapter 5.

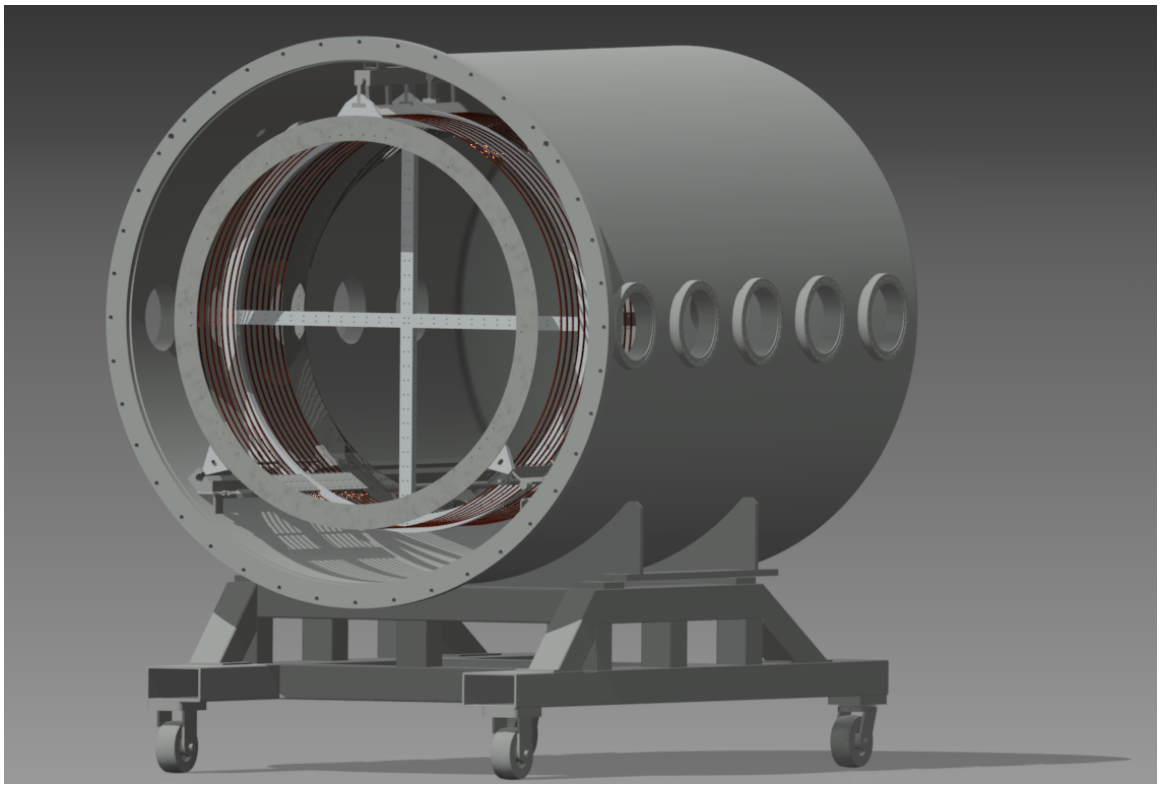

Figure 4-41: A CAD image showing the full TPC design within the vacuum chamber.

Chapter 5

1m^3 Commissioning

This chapter describes the construction of the newly designed TPC and the testing thereof.

5.1 Prototype

Before going ahead with the full four-cage construction, a prototype consisting of a single TPC was fabricated to verify the design and to determine any modifications that might be required. The first part to be built was the copper-ring field cage followed by the addition of the amplification region and cathode. Based on the results of the analysis presented in chapter 3 and the conclusion that the surface area of the detector components is a big contributor to the background, the construction process was closely controlled. Apart from the initial process outlined in section 5.1.1, all parts were put together in either a clean tent[1], as in the case of the field cage, or a clean-room. This helps to minimise the exposure of the components to dust, which can have radon stuck to it. It also minimises dirt which can outgas at low pressures, causing the time required to evacuate the chamber to take longer and causing possible impurities of the gas during detector operation. The clean tent air was measured

[1]Clean Air Products, model Series 577 class 1,000.

using a Durridge RAD7 radon detector, to have an average radon content of 3.59 $\pm$ 0.6 Bq/m^3. The dust level of the clean tent was also measured, using a Dylos DC1100. Both small ($> 1\ \mu m$) and large ($> 5\ \mu$ m) were both measured to the limit of the Dylos device of just under 3,000 particles per m^3. Protective clothing was also worn: nitrile gloves, hair net, shoe covers and overall, and the parts themselves were cleaned as outlined in section 5.1.2.

5.1.1 Preparation

The copper used for the field rings was OFE-OK oxygen free copper. The reason for this is that oxygen can outgas from the copper during operation. Oxygen is electronegative, thus any electrons produced from an ionising recoil could attach to the oxygen, causing a reduction in the signal. The initial prototyping done in DaRCO was made with spare copper from the workshop, the hardness for which is unknown. The maximum hardness that could be provided for the 1m^3 was H00, which is the second softest available. This was due to the decision to purchase pre-rolled copper that is provided on a spool at the desired dimensions. By having the copper pre-cut and pre-rolled, the ring fabrication process is much simplified, since cutting and rolling of hard copper into a round shape can easily warp the material. The only vendor found that could provide OFE-OK oxygen free copper supplied in this way, could only do so by the use of the H00 hardness. In this form the copper was easily cut to the appropriate length and the two ends braised to make a ring. However, a fabricated ring was incapable of supporting its own weight, due to the softness, causing it to sag and permanently deform as shown in figure 5-1. To overcome this the copper was work hardened. Different methods were tested in a qualitative study. First the strips were fed through a metal roller, then hammered using a rubber mallet and finally hammered with an air-hammer. Variations of multiple hits at low pressure to few hits at high pressure were tested. The roller was not useable as it introduced

considerable warping to the material. The rubber mallet did harden the copper to a reasonable amount but was very labour and time intensive so was not practical for the long term. Sufficient hardening was achieved by using the air-hammer at a pressure of 100 PSI and hammering several times on the outer surface of the ring. Figure 5-2 shows the stand and hammer used for this hardening process. The copper was hardened to the point where it could be held from the top and though it still sagged, it did not cause the metal to bend and remain deformed. During the construction phase it was actually found that the copper softness was advantageous, allowing subtle manipulation to accommodate the tolerances of the support structure, particularly the rails in the chamber.

(a) Saggingof the soft copper ring.

(b) Deformation

Figure 5-1: Images showing the softness of the copper (a) and the permanent deformation induced (b).

Figure 5-2: The frame used to harden the soft copper rings.

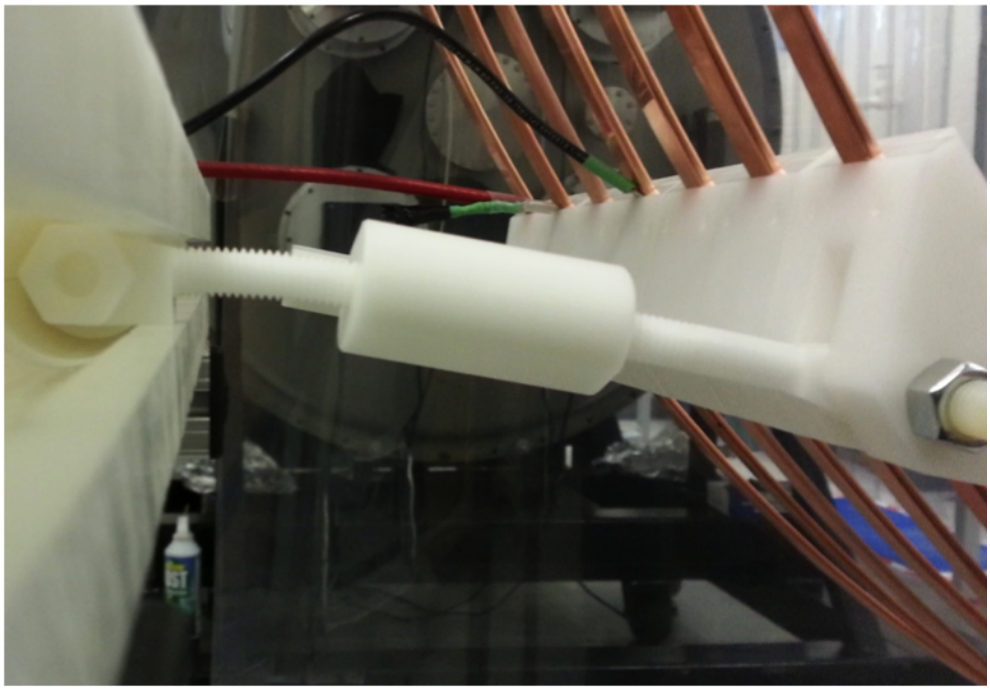

Figure 5-3: This image shows the sagging side support rods caused by the rings inability to support its own weight.

5.1.2 Cleaning

Before putting all the parts together they were each meticulously cleaned using a similar procedure outlined in [78] and [55]. This is required to remove both the dirt present from the machining of the parts, which can outgas, and any dust that may have settled on the parts. The procedure is listed below.

1. Scrub with soapy water
2. Wipe with acetone
3. Wipe with isopropanol
4. Ultrasonic cleaning in RadiacWash for one hour
5. Wipe with isopropanol
6. Wipe with de-ionised water

7. All wipes were done using lint-free KimWipes

Once the cleaning was completed the parts were wrapped in aluminium foil and stored in the laboratory, ready to be used in construction. At the end of the construction process all parts were given a final wipe with isopropanol prior to being transferred into the vacuum chamber. The chamber walls, up to the point reachable, and the inner door surface were also scrubbed with isopropanol. This was the procedure for the prototyping stage, to remove machine oil and dust. The plan for the future is to electro-polish the chamber and the field rings to remove radon plated out on the surfaces.

5.1.3 Construction

To build the field cage outside of the chamber, a frame which reproduced the internal rails of the chamber was required. This was designed and then built out of 80/20 frame material (figure 5-4). The A-frame legs of the design were found to take up too much room in the laboratory. The parts were therefore replaced, with no loss to stability. Figure 5-9 shows the CAD image and first build of the frame. Figure 5-9 contains the construction frame in its final configuration.

Once the hardening process had been completed for the seven rings of the prototype field cage they were added one-by-one to the build in the construction frame. On the first iteration of the build an additional problem arose from the sagging of the rings. Since they were unable to support their own weight, when hung from the top Delrin rail, there was significant bending of the side alignment rods, see figure 5-3. This was a concern not only in the immediate, but for long term usage, so a modification to the side rails was made, in the form of a wedge. This repositioned the side rods to extend from the normal of the ring surface, providing an additional upwards force as in the original proposed concepts. This simple fix turned out to be very effective (figures 5-5 and 5-6).

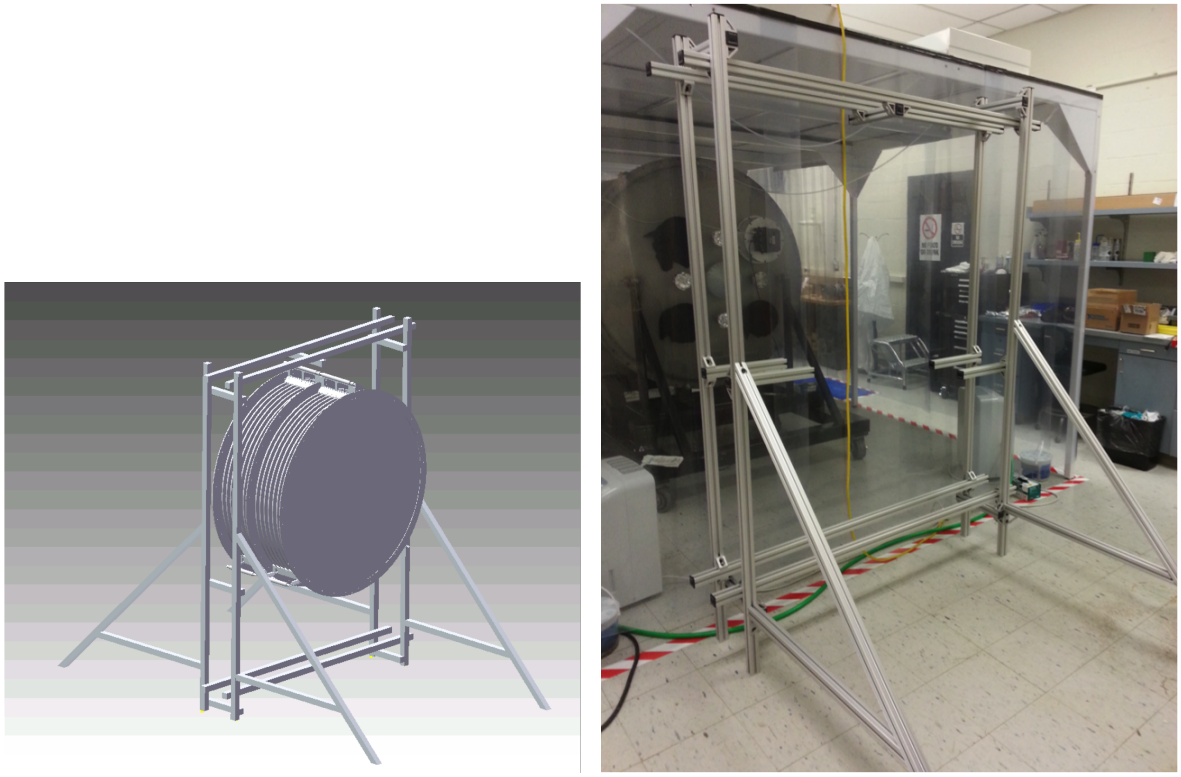

(a) (b)

Figure 5-4: (a) shows a CAD image of the construction frame for building the field cage outside of the detector. (b) shows the frame as first built.

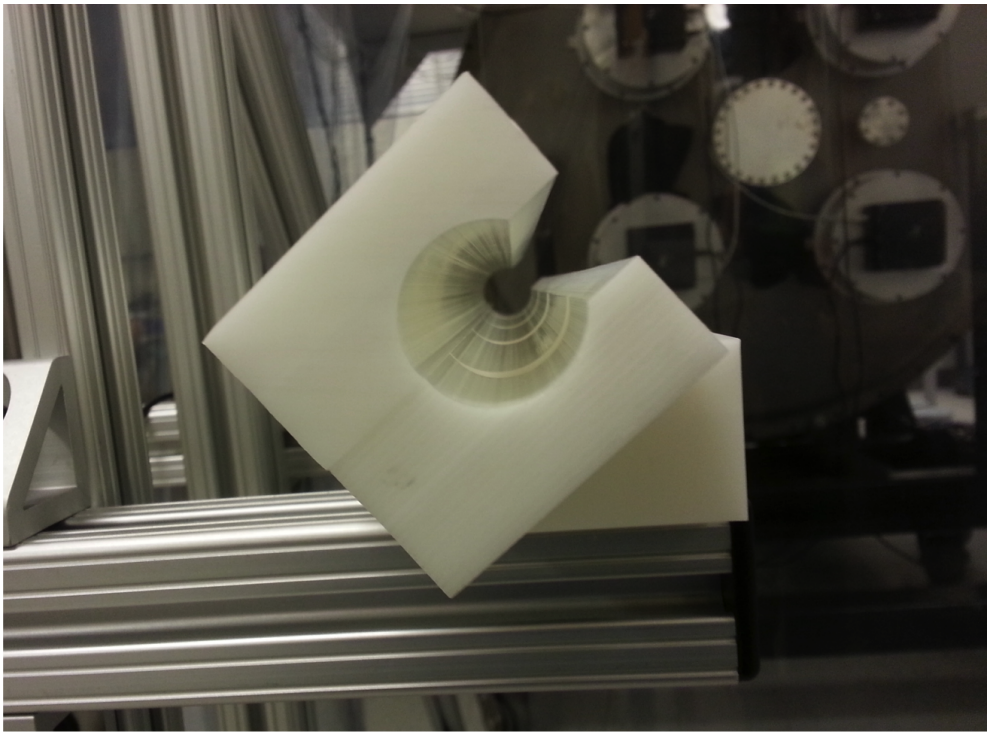

Figure 5-5: The construction frame shown with the added support wedge. This improved the support of the rings and enabled the correct circular shape to be produced.

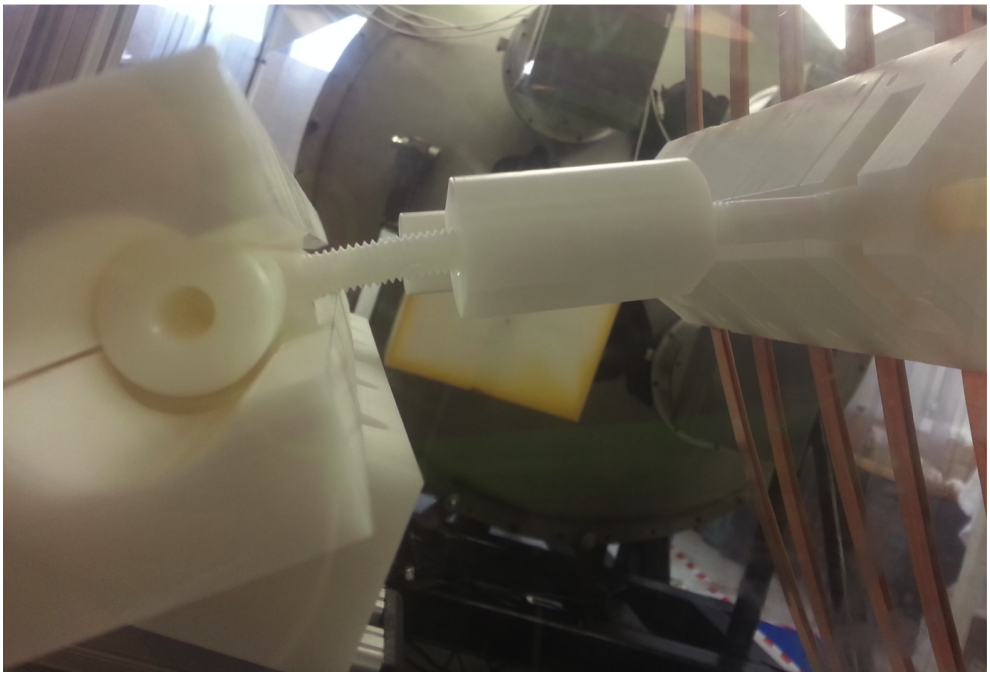

Figure 5-6: This image shows the effect of the wedge to the side support, compare to figure 5-3.

The next step in the process was to connect the rings using the resistors. This was achieved by using spare ring-support housings as placeholders, removing the front cap of the housing and then putting the resistor in place. To ensure good electrical conduct between the ring and the resistor, small conductive adhesive tabs were used[2]. Figure 5-7 shows the steps in the process.

Once in place, a great deal of care was taken to ensure the alignment of components was as uniform as possible. This involved the use of a spirit-level and tape-measure, used to check that the support structure was level and the spacing of the rings was even. Due to the softness of the metal, some subtle manipulation was also required to achieve satisfactory roundness of each ring. Figure 5-8 shows some examples of the alignment procedure. When the alignment and shape of the field cage was completed, the construction frame was wheeled up to the front of the chamber to begin the transfer process, figure 5-9. On the first attempt this was not a smooth process, but with some adjustments to the frame and some practice it soon became relatively easy to slide the field cage from the construction frame and into the vacuum chamber. After the placement, the manipulation and alignment process was repeated as an additional check. Figure 5-10 shows the prototype in place in the chamber. To quantify the alignment and uniformity of the construction, measurements of the ring spacing were taken at several locations around the ring: the approximate mid-point between supporting rails, corresponding with the least supported point, and either side of the supporting structures. The spacing was measured to be 2.54 ± 0.50 cm. The co-axial alignment and roundness of the ring was tested by measuring the distance of the outer ring edge to the chamber wall. This was measured at the same locations and was found to be 15.5 ± 0.5 cm. The diameter across the vertical, horizontal and both diagonal positions was measured as a test for ring roundness. At each location the value was measured to be 120.0 ± 0.5 cm. These values are in line with the

[2]Agar Scientific Leit adhesive tabs

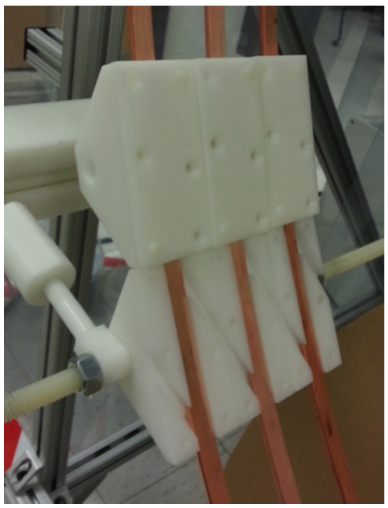

(a) Placeholders

(b) Stickers

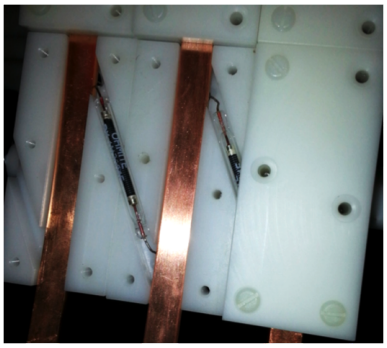

(c) In place

Figure 5-7: (a) the method for keeping the rings in place whilst placing resistors in their channels. (b) a picture showing the placement of the resistors and the adhesive carbon tabs used. (c) the resistors in place with rings.

expected values and well within the 1 cm tolerance prescribed in section 4.4.

Figure 5-8: These pictures show the alignment of the field cage rings.

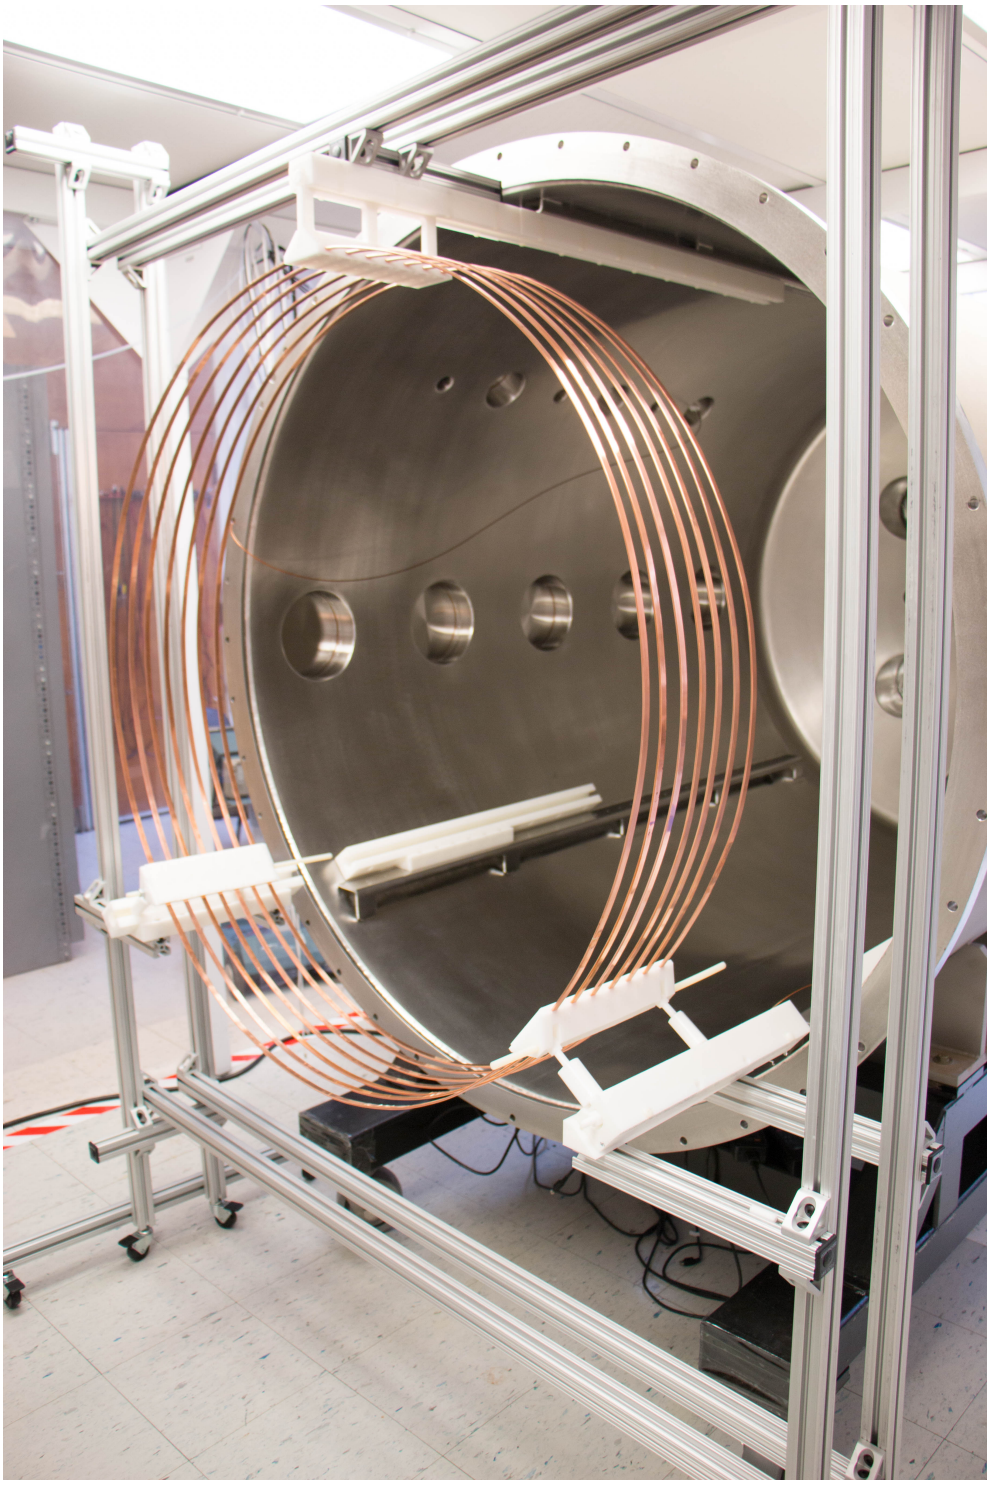

Figure 5-9: The completed field cage prototype being prepared for transfer to the vacuum chamber. This image also shows the construction frame with the A-frame legs replaced with shorter, lower support legs.

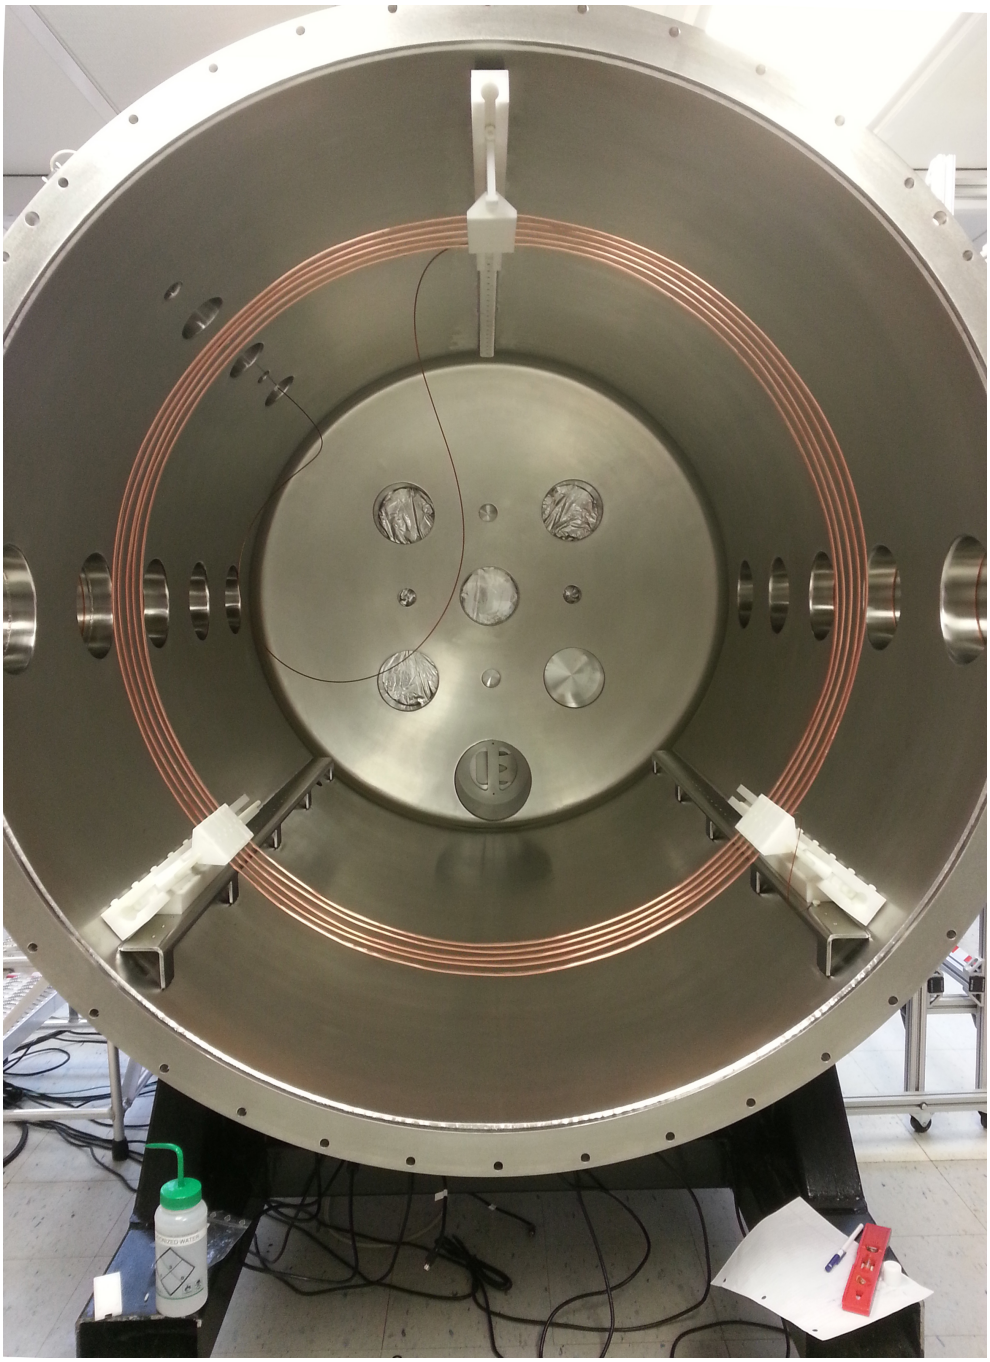

Figure 5-10: The field cage successfully placed inside the vacuum chamber.

5.2 Field Cage Prototype Commissioning

5.2.1 Field Cage Voltages

The voltage required to be applied to the cathode of the TPC is that for which a reduced field value of ∼10 Td is achieved and thus diffusion is kept low, as described in section 1.2.4. To test the voltage the chamber was evacuated to $\mathcal{O}(10^{-}6)$ before filling with CF_4 to the pressure being tested. The voltage on the supply was slowly ramped up until the supply tripped from a current draw. The value before which the trip occurred is considered the achievable voltage. This was tested at four different pressures of CF_4 and the achievable voltage was well above that required in each case. See table 5.1 for the results.

With the field-cage operating stably and exceeding the requirements, the remainder of the TPC was built. These parts, the cathode and the amplification region, were installed by a colleague due to the timing of travel and parts fabrication.

With all parts of the TPC installed and operational, some initial commissioning data were acquired, figure 5-11 shows the first alpha track observed. This was taken with no source present so is a background track. It demonstrates that the TPC is operational and the cameras are in focus. At this stage, with the cluster reconstruction algorithm and energy calibration still being undertaken, the energy and range of the track are unknown.

Table 5.1: Table of achieved and required voltages for the $1m^3$ prototype field cage.

CF_4 Pressure (Torr)	Required Voltage (V)	Achieved Voltage (V)
12	1100	3200
20	1900	5400
30	2800	6400
40	3600	8000

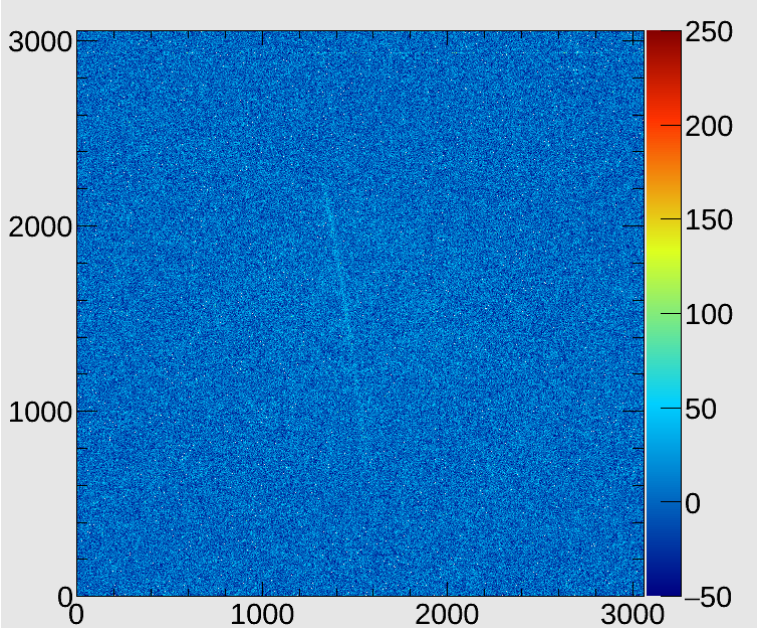

Figure 5-11: The first alpha track to be observed in the new detector.

5.2.2 System Gain Map

The next step in detector commissioning was to produce a gain map of the amplification region, as described in section 3.2.1. The source used in this instance was ^{55}Fe, which produces two photons, both ~ 6 keV. At 30 Torr of CF_4, the attenuation length[3] for a photon of this energy is approximately 2 m. This is calculated using the attenuation equation, 5.1, solved for x, as in equation 5.2. The attenuation coefficient, μ for CF_4 at 30 Torr is 34.22 cm^2/g [79] and the the density of CF_4 at 30 Torr is 1.47×10^{-4} g/cm^3. The placement of the field cage in the chamber is determined by the focussing of the cameras, which places the amplification region at approximately 84 cm away from the chamber door. It was therefore reasonable to place sources in the spare ports attached to the door, as shown in fig. 5-12. To do this, a source holder was designed and fabricated. Figure 5-13 shows a CAD image of the source holder.

[3]Defined as the length at which the intensity has dropped to $1/e$.

$$I = I_0 \mathrm{e}^{-\mu\rho x}, \tag{5.1}$$

$$x = \ln\left(\frac{I_0}{I}\right)\frac{1}{\rho\mu}, \tag{5.2}$$

where I is the intensity at x, I_0 is the initial intensity, μ is the mass attenuation coefficienct and ρ is the density of material. For equation 5.2, I_0/I is substituted for e, to get a value of $x = 2$ m.

At this point the detector is still in the commissioning stage and optimum operating conditions are still being studied. Specifically the maximum gain achievable, whilst maintaining a reasonably low spark rate ($\mathcal{O}(0.1)$Hz) is still to be deduced. As such there are no calibrated energy or range conversion factors available for the CCD clusters, so energies are presented as intensity in ADU and range is given pixels. For this test, the data were acquired for a single camera of the four-shooter side of the detector, at a range of anode voltages. Different voltages will result in differing absolute gains, however it is the relative gain that is required and so it is not unreasonable to use different voltages. The exposure length of the images is thirty seconds. A long exposure minimises read-noise relative to signal, as it is time-independent. Dark-noise[4], however, does increase with exposure, so a very long exposure is undesirable. Figure 5-14 shows a test of the pixel-intensity distribution standard deviation as a function of exposure. This was produced by acquiring data, with the camera shutter closed, at each exposure length, then plotting the pixel intensity distribution for each run. A Gaussian fit was then applied and the resulting standard deviation was used.

The data were checked for quality prior to being used. The bias frames of each run were checked and some were found to contain RBIs. This is indicative of significant sparking occurring prior to the run. These runs were excluded from the data. Next,

[4]Dark noise is induced by thermal fluctuations of the electrons in the CCD chip

Figure 5-12: ^{55}Fe sources placed in the vacuum chamber door. The sources are the black and white circles located in-between the CCD windows.

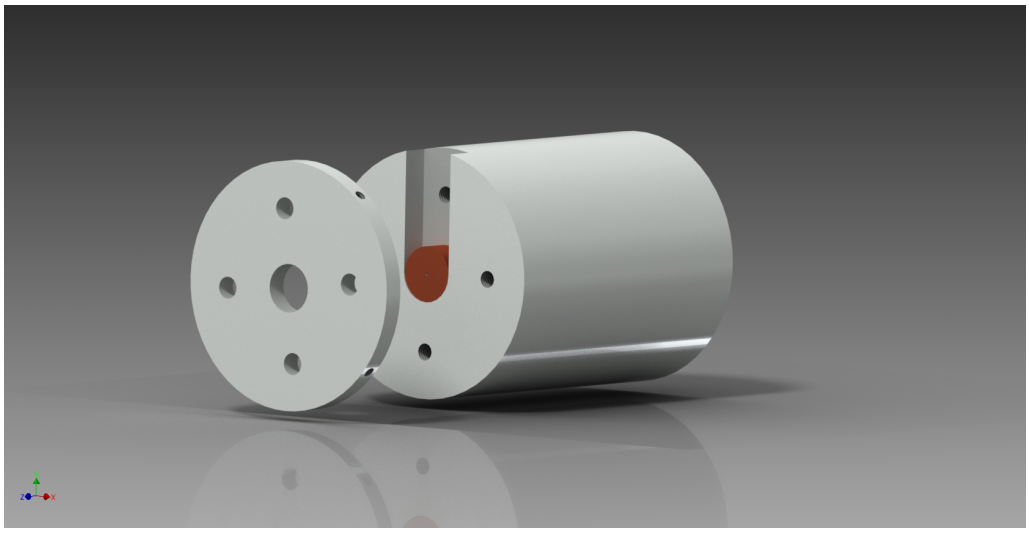

Figure 5-13: A CAD image of the source holder design. The long cylinder, which has a slot allowing for the easy removal of the source disc (shown in dark orange), is designed to sit in a port on the chamber door. The smaller cylinder screws onto the front, securing the source disc in place, but has an aperture so as to not attenuate the source.

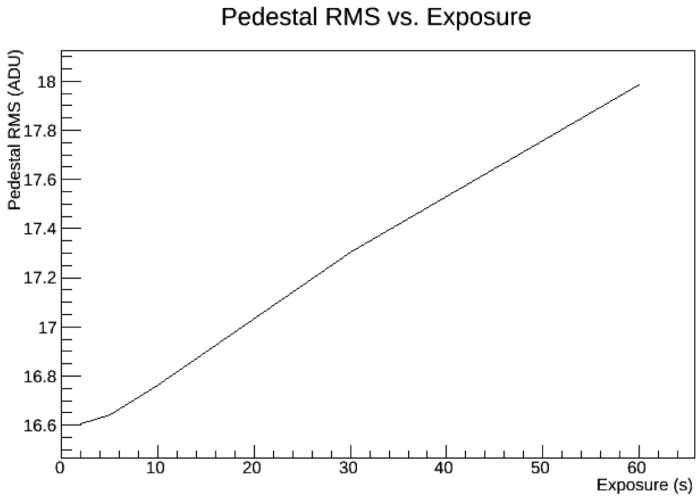

Figure 5-14: Standard deviation of pixel intensity distribution as a function of exposure length.

the image means were plotted as a function of run number, shown in figure 5-15. The image mean is defined as the sum of all the pixel intensities divided by the total number of pixels.

It is apparent that the first few runs of the voltage-off data show anomalous behaviour before settling down. This is possibly due to the camera not being properly cooled to the operating temperature of -20 °C. This cooling is required to minimise the dark noise. There are three runs of the voltage-on data, figure 5-15 (b), that have a raised mean in the left hand side of the histogram, between runs 310 and 320. These come just before and just after the region of data that was omitted due to RBIs being present in the bias frames. This is attributed to excessive sparking and so these runs were also excluded. The right hand side of the histogram shows data taken at varying voltages so a difference in mean is not unreasonable, however the first run shown here has a substantially lower mean and a large spread, therefore this run was also omitted from the analysis.

After run-level data quality cuts, event-level data quality cuts were employed because in the remaining runs there are still a number of sparks present. It is crucial to remove these from the summed image due to the illumination that they cause. The induced RBIs from these sparks must also be taken into account as they too act as artificial illumination. To remove the sparks, the same mean-ratio method, described in section 3.3.1 was modified to take the ratio between the event image mean and the bias-frame mean. The reason for this modification was that many sparks were found to be in multiple sequential images, thus taking the ratio with the previous image would not act as a good indicator. Figure 5-16 shows the distribution of the mean-ratio variable. The sparks that have been identified by scanning the data by eye are shown in red and all other images are in blue. A clear double population can be seen. There is a distribution centred around one, as would be expected for images with little to no illumination. There is also a distribution which has a mean

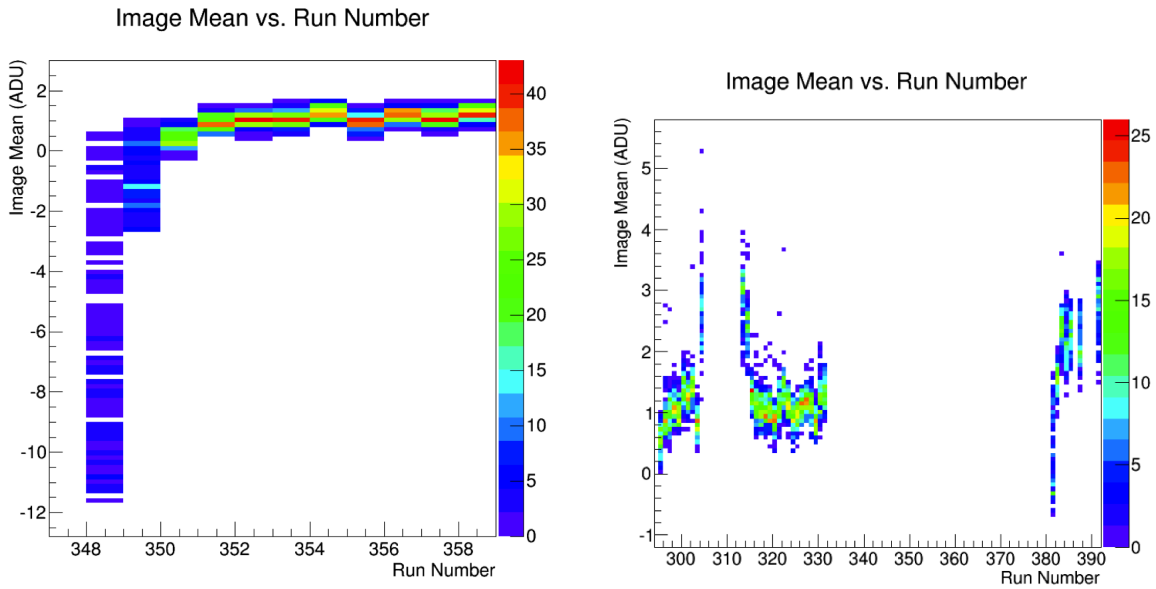

(a) Voltages off

(b) Voltages on

Figure 5-15: Image means plotted versus run number. For voltages off (a) and voltages on (b). In (a) there is a clear discrepancy in the value between the initial few runs (348-350) and the runs that follow. As such, these were not used in the analysis. In (b) a peak can be seen either side of the gap around run 310. There was sparking during this period so the runs were excluded. Around run 381, there is a dip in the mean from the subsequent runs. These runs were also discarded.

greater than one, as would be expected for images with increased illumination. To use these data to remove the sparks from the data set, a Gaussian fit was applied to the population, centred around one. The fit produced a mean of 1.00007 and σ of 0.00012. Images greater than three sigma away from this mean, with a value of 1.00043, were rejected as sparks.

The spark cut was applied after processing with cleanSkim[5], meaning that images with spark-induced illumination produced reconstructed clusters. Figure 5-19 shows an example image. This prompted the use of an additional cut on reconstructed clusters with a high intensity and range value[6]. To deduce the cut value, the intensity and range of reconstructed tracks was plotted - figures 5-17 and 5-18. Due to the size and intensity of the sparks, the values will be large. The values used for rejection were: intensity $\leq$ 20,000 ADU and range $\leq$ 300 pixels.

To exclude any RBIs after sparks, the pixels of reconstructed tracks that have been excluded as sparks, were checked for being above a threshold value. If they were, these pixels were then excluded from the remainder of the run. To determine the threshold value to use, all bin values of all reconstructed clusters were plotted, figure 5-20. Though it is not an event-dependent variable, the bin value was plotted as a function of event number. Figure 5-21 shows the one-dimensional distribution of bin values. The threshold value was chosen to be 5000 ADU, where the tail of the low end diminishes.

Using these methods, each bin of the summed image histogram is filled a different number of times. It is therefore necessary to keep track of the filling to ensure the correct averaging of the final image. For this a separate histogram was made to count the number of times each bin was used. This histogram was then used to average each bin individually. Figure 5-22 shows the resulting pedestal-subtracted summed-average

[5]This is due to cleanSkim being configured to base the cut on the original method

[6]At this stage the cluster finding algorithm has yet to be configured, as such the reconstructed clusters are only used at this stage as a supplementary spark removal tool. Though this should also remove background alpha tracks

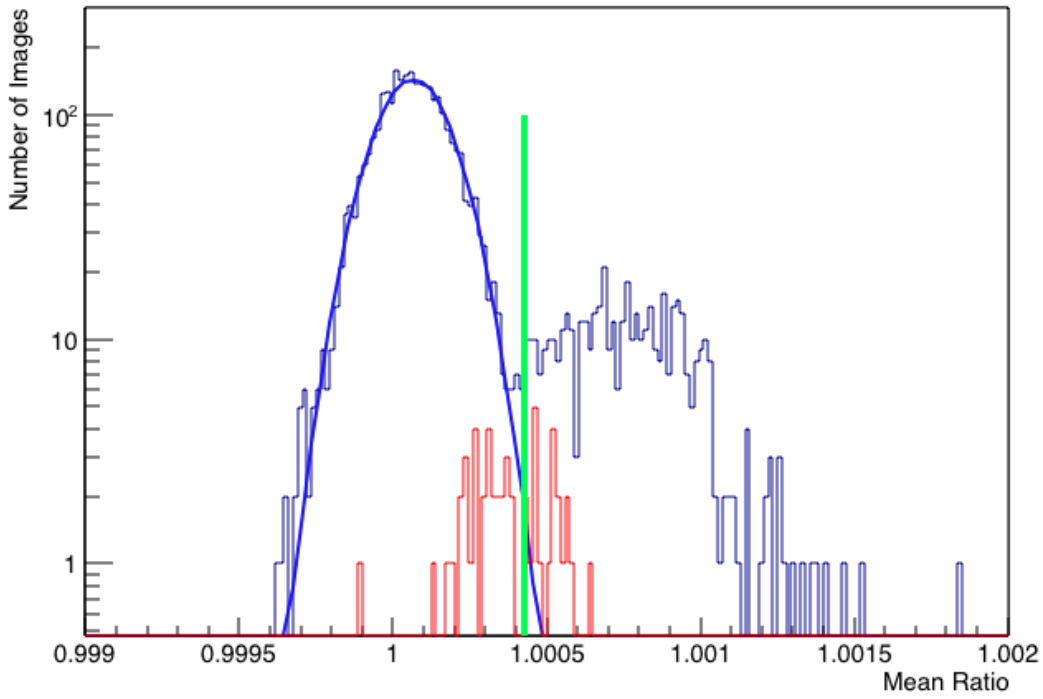

Figure 5-16: Ratio of image means to bias frame means. The abscissa is the ratio of image mean to bias frame mean, the ordinate is the number of images. The blue line is a Gaussian fit to the population centred around one - the expected mean for a non-spark image. The red data points are known spark images. The green line indicates the cut value, values below the line are accepted. The population with higher image means could be attributed to unidentified sparks or the images might contain alpha tracks.

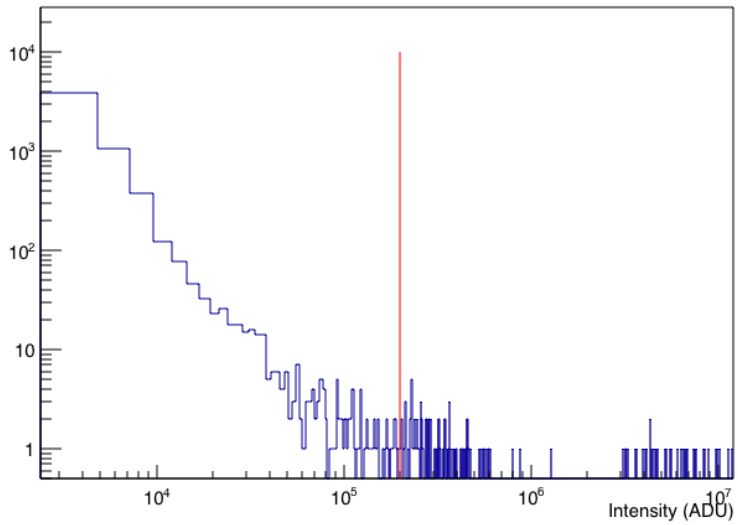

Figure 5-17: Intensity distribution for reconstructed clusters. The red line shows where the cut was made, clusters with intensity above this value were discarded as sparks.

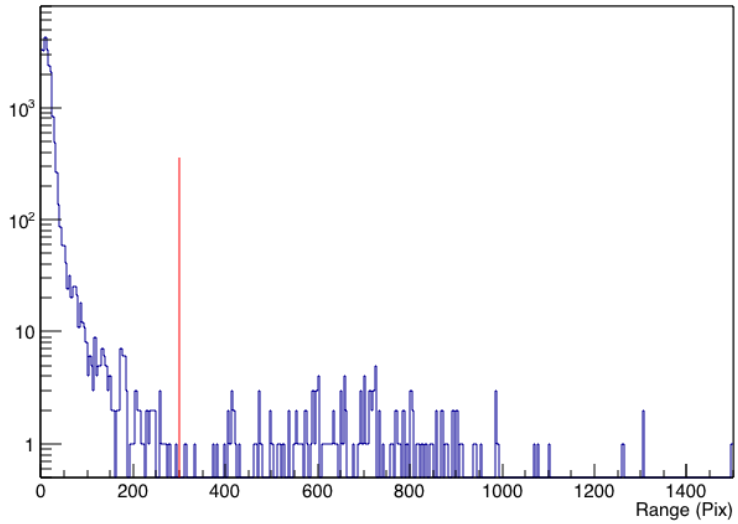

Figure 5-18: Range distributions for reconstructed clusters. The red line shows the where the cut was made, clusters with range above this value were discarded as sparks.

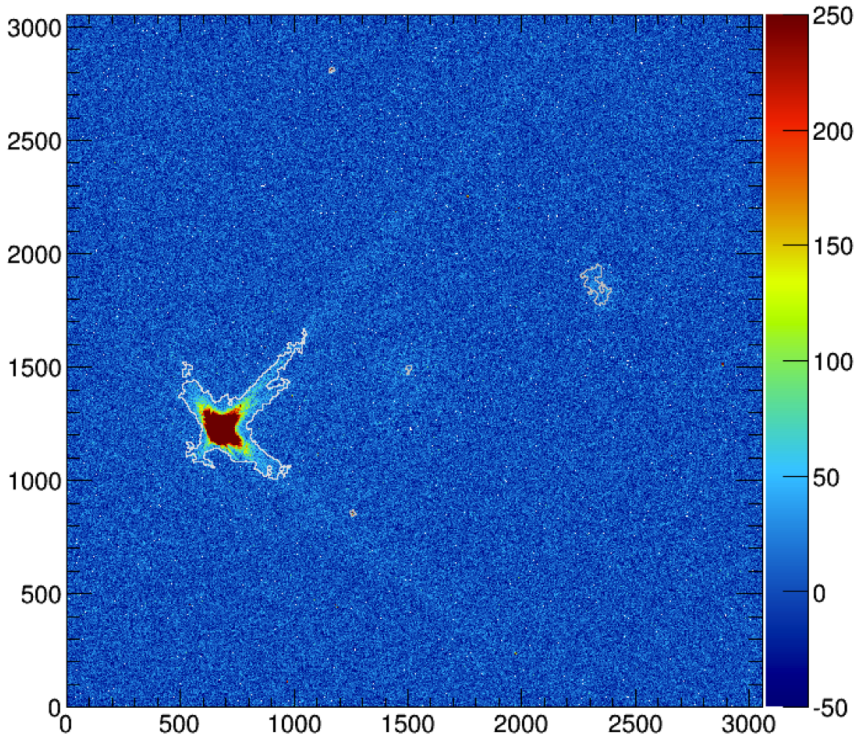

Figure 5-19: Image showing a spark reconstructed as a cluster. The abscissa and ordinate are in pixels, the applicate is ADU.

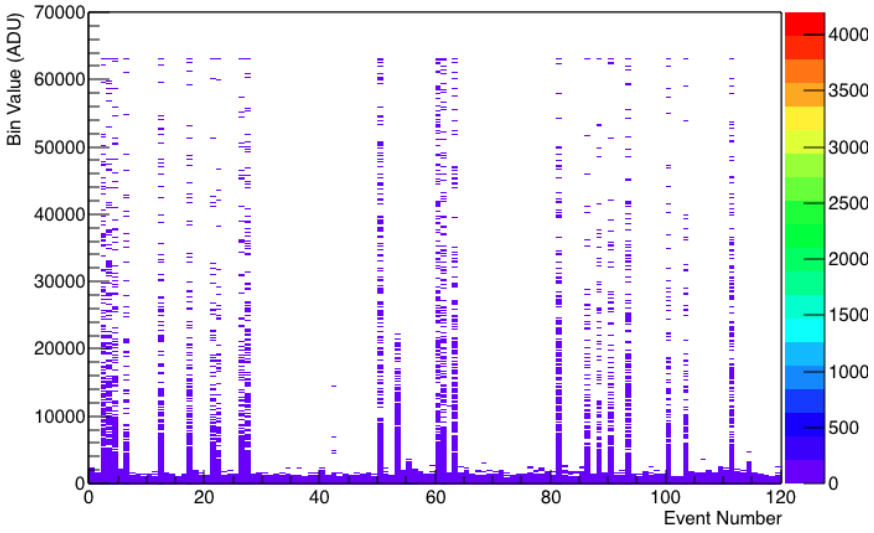

Figure 5-20: Cluster bin-values plotted as a function of event number.

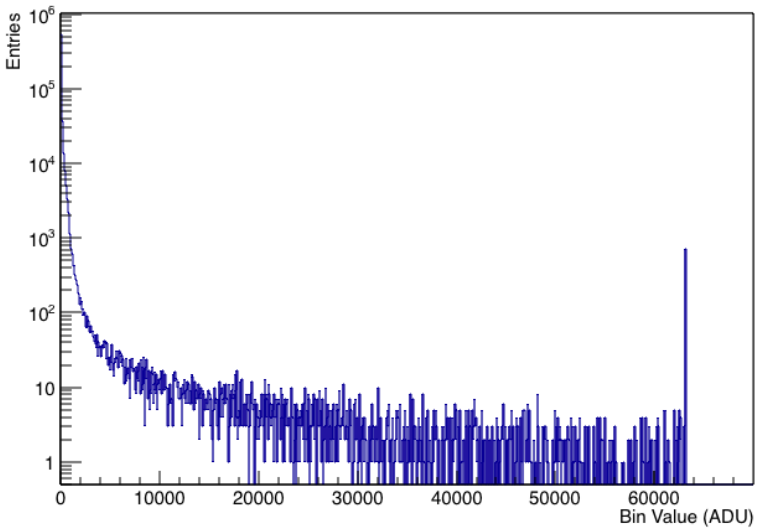

Figure 5-21: Cluster bin-values distribution.

histogram. If, over the total time of data-acquisition, the detector was illuminated sufficienty, it would be expected that the amplification region would be visible in this image. To demonstrate the impact of the cuts described above, figure 5-23 shows the bin-counting histogram. Figure 5-24 shows the resulting gain map, which is the inverse of figure 5-22. The inverse of the summed image is used for the gain map due to the nature of the calculation made for the correciton in cleanSkim, which is to multiply an unadjusted bin value by the gain map value. With the gain map as the inverse of the sum, this acts to normalise the value by the gain variation.

It is apparent that insufficient data were acquired to allow for the illumination of the amplification region to be seen above background. Owing to the imaging configuration, as described in section 2.3.3, the image would be expected to contain a distinctly highlighted region as outlined by the quadrant imaged by the camera. As an example, figure 5-25 shows an image of a purposely made light leak, highlighting the amplification region.

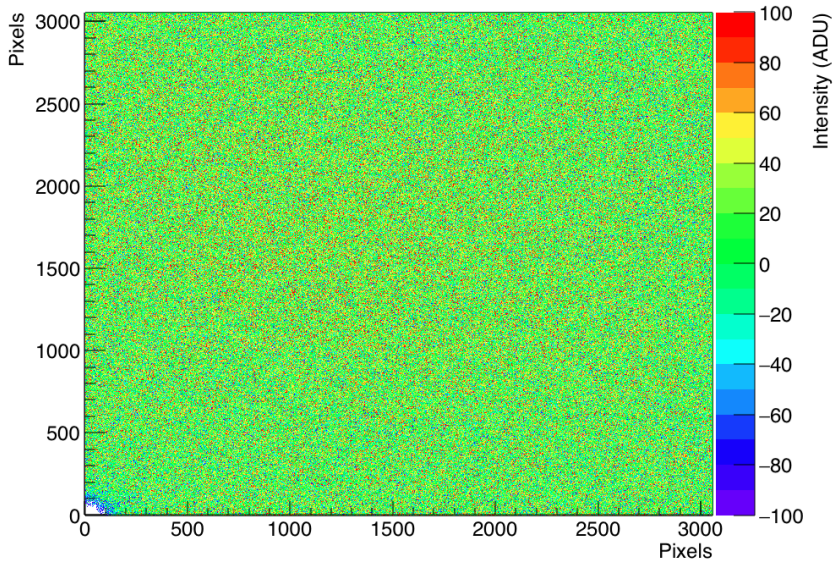

Figure 5-22: The pedestal-subtracted summed-averaged histogram of ^{55}Fe source data. The feature in the bottom left cornera, around (0,0), is due to a light leak in the camera housing. This camera is old and this type of feature has been seen before. It does not impact this result since the amplification region does not overlap with this location (see figure 5-25). No other features are discernible indicating either insufficent data were acquired or that the detector gain was too low.

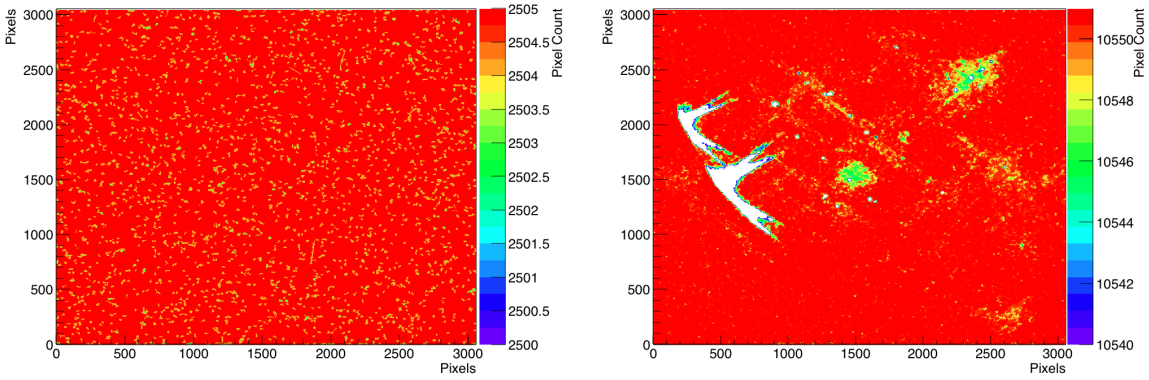

Figure 5-23: Histogram of stored number of entries added to the summed image for each bin. (a) voltage-off data. b) voltage on data. The abscissa and ordinate are in pixels, the applicate is ADU. The white regions in (b) contain values below the lower threshold of the applicate. The threshold value of the applicate was chosen to highlight the other regions, coloured yellow and green, which are not visible when the full scale is used.

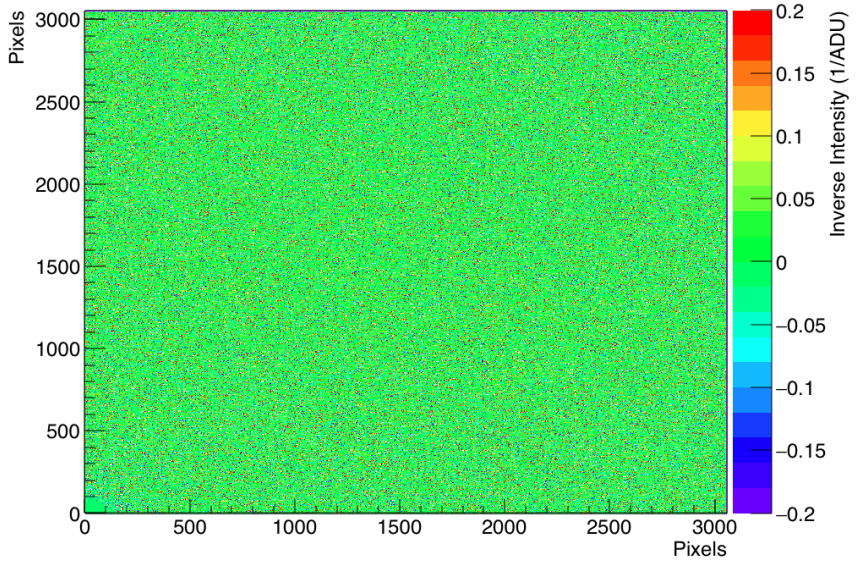

Figure 5-24: The resulting gain map for the $1m^3$. No variance can be seen at this scale, other than that induced by random fluctuations.

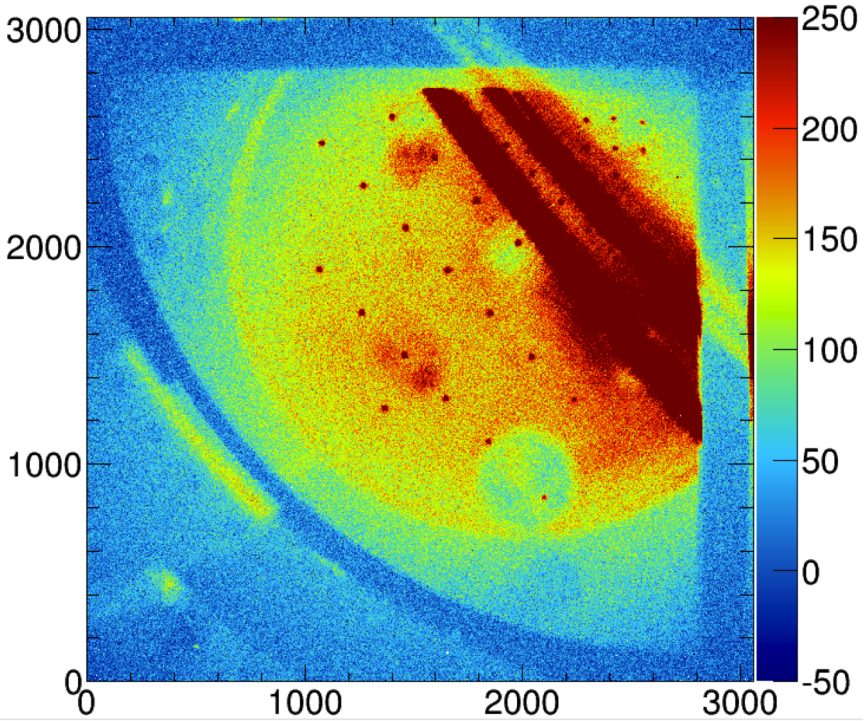

Figure 5-25: Image showing the amplification region quadrant as highlighted by an intended light leak The dark blue shaded region outlines the amplification region quadrant, the large yellow circle shows the back of the chamber, the dark red bars originate from the light leak - thus have the highest intensity - and the dark red dots are the spacers used to separate the amplification region meshes.

5.3 Conclusions

The field cage build was very successful, with a method deduced to harden the soft copper and a simple fix employed to provide additional support to the ring sides. The field cage has been operated at voltages that far exceed those required and has been used in the full TPC construction to produce observable data. The gain map currently is not highlighting any gain variation across the amplification region and this is likely a consequence of insufficient statistics. This can be resolved by acquiring additional data as well as increasing the amplification region anode voltage, and so the system gain.

Chapter 6

Summary

After an introduction to the origins of dark matter and a description of the methods employed to detect it, the main body of this thesis described two main projects - a software-based analysis on an underground dataset (chapter 3) and a hardware-based design and construction process (chapters 4 and 5). In the analysis of chapter 3, the rate of background events passing all event-discrimination cuts was reduced from previous analyses, using the same detector. This was achieved through re-focussing one of the detector cameras, updating the calibration values, improving and re-tuning the event-discrimination cuts and updating the limit-setting code. The limit on the spin-dependent WIMP-proton cross-section was comparable to that of the previous analysis, despite having reduced total exposure and the additional noise of the bottom camera. One of the conclusions of this analysis is that there are still a significant number of background events with nuclear-recoil like event parameters. Since radon, plated out onto the materials that comprise the TPC, is expected to be the dominant contributor to these events, the need to reduce the total surface area of detector internals became a priority in designing the next generation detector.

Chapters 4 and 5 then took the goals of reducing the TPC surface area and producing a fiducial volume of 1 m^3 to design, simulate, construct and test this

new TPC. Chapter 4 outlined the design ideas and showed how the simulation was constructed including the use of a new metric for determining the optimum model. The metric was constructed using the modelling of the transport of charged particles from both the top of the field cage, along the radial direction, and from the bounding box of the field cage rings. From these, two values were determined (r_1 and r_2) which were used with a minimum fiducial radius, at which charged particles originating from the top of the field cage would survive to the amplification region and be detected, but those originating from the field cage rings would not drift to the cathode and be within the fiducial region. This enabled a design to be selected that satisfied both of the design requirements.

With the final field cage configuration chosen, the design of the cathode and TPC support structure was discussed. This led to a final simulation to model the impact of the amplification region support frame, the result of which was a reduction in the fiducial volume from 1.1 m^3 to 0.892 m^3 ($\sim$ 20 %). This was followed by a demonstration of the mechanical prototyping and subsequent refinement of the support structure components. The final part of chapter 4 gave a comparison of the surface area to fiducial volume ratio for each of the three main DMTPC detectors given in section 2.3. This showed a significant improvement to the overall value and also validated the new metric employed for optimisation.

Finally, chapter 5 described the process of building and testing the TPC. The construction process was outlined, with reference to the support frame used to build the field cage outside of the vacuum chamber and modifications made to the field cage side support mechanism. When in place the shaping and alignment of the field cage was performed then measured and found to be well within the tolerances required. The field cage was then tested at different pressures to determine the maximum operating voltage. The result of this was successful operation well outside of the values needed to obtain a reduced field that satisfies the diffusion and drift velocity

requirements. The last part of the chapter discussed the generation of a system gain map by uniformly illuminating the drift region with an ^{55}Fe source. In this instance no illumination was seen, this was possibly due to insufficient data or too severe data quality cuts.

Overall, this thesis demonstrates how the results obtained from the 10L detector drove the design of the new TPC for the 1m^3 detector and how the new design represents a significant improvement over previous iterations.

Appendix A

Exposure Time Calculation

Total live time is the number of images acquired multiplied by the exposure length, which is then averaged over the two cameras. From this is subtracted the time-equivalence of images cut due to sparks and residual bulk images (RBIs). For example, in wr5, a spark image plus the following 5 images are discarded. Any bins that have been saturated by the spark are also discarded for the remainder of the run. Equation A.1 gives the result for this in terms of time.

$$T_{cut} = (N_b * (N_{cut} + 1) + N_{sat} * N_{events}) * S_p \tag{A.1}$$

Where T_{cut} is the time to cut, N_b is the total number of bins in a single image, in this case $256 \times 256 = 65536$, N_{cut} is the number of images cut post-spark - this is exposure dependent and is discussed further in section 3.3.1. The +1 is to account for the spark image itself, N_{sat} is the number of pixels saturated by the spark, N_{events} is the number of events following the spark - this part of the equation deals with the RBIs - and S_p is the seconds per pixel which is equal to the exposure time divided by the number of pixels. It is clearly necessary to do this separately for each camera as they will have different rates of sparking. In order to then get a total exposure time, the T_{cut} value for each camera is subtracted from the total live time (averaged). These

numbers are then multiplied by the fiducial volume of each TPC, which is different due to the different vixel sizes. Finally, these two numbers are added together.

The parasitic time is deduced by iterating through all of the runs and getting the total time taken from the start of the run to the end and subtracting the number of runs multiplied by the exposure.

Bibliography

[1] E. V. Linder. *First Principles of Cosmology.* Addison-Wesley, 1997.

[2] L Bergström and A Goober. *Cosmology and Particle Astrophysics.* Praxis Publishing, second edition, 2006.

[3] D. Perkins. *Particle Astrophysics*, volume 1. Oxford University Press, 2006.

[4] E. Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *PNAS*, 15(3):168–173, 1929.

[5] Planck Collaboration (P.A.R. Ade (Cardiff U.) et al.). Planck 2015 results. xiii. cosmological parameters. arXiv:1502.01589, February 2015.

[6] J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bulletin of the Astrophysical Institutes of the Netherlands*, 6:249, August 1932.

[7] F. Zwicky. On the masses of nebulae and of clusters of nebulae. *Astrophys. J.*, 86:217, October 1937.

[8] M. Markevitch. Chandra observation of the most interesting cluster in the Universe. arXiv:astro-ph/0511345.

[9] D. Clowe, S. W. Randall, and M. Markevitch. Catching a bullet: direct evidence for the existence of dark matter. arXiv:astro-ph/0611496.

[10] D. Paraficz et al. The bullet cluster revisited. arXiv: 1209.0384.

[11] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /r = 4kpc/ to UGC 2885 /r = 122 kpc/. *Astrophysical Journal*, 238:471–487, June 1980.

[12] K. Freese. Review of observationa evidence for dark matter in the universe and in upcoming searches for dark stars. arXiv:12090384v1.

[13] R Adam et al. Planck 2015 results. i. overview of products and scientific results. arXiv:1502.01582, 2015.

[14] J. M. Alimi et al. Deus full observable λcdm universe simulation: the numerical challenge. arXiv:1206.2838.

[15] Center for Theoretical Astrophysics and University of Zurich Computing. Comparison of cold, warm and hot dark matter simulations. http://www.ctac.uzh.ch/gallery/index.html.

[16] G. Jungman, Kamionkowski, and K. Kriest. Supersymmetric dark matter. *Phys. Rept.*, 267:195–373, 1996.

[17] Planck Collaboration (P.A.R. Ade (Cardiff U.) et al.). Planck 2015 results. xvi. isotropy and statistics of the CMB. arXiv:1506.07135v1.

[18] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81(8):1562–1567, August 1998.

[19] R. D. Peccei. The strong CP problem and axions. *Lect. Notes. Phys. 741*, pages 3–17, 2008.

[20] T. Kaluza. Zum unitätsproblem in der physik. *Sitzungsber. D. Berl. Akad.*, pages 966–972, December 1921.

[21] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors and corrections for dark matter experiments based on elastic nuclear recoil. *Astropart. Phys.*, 6:87–112, 1996.

[22] S. Archambault et al. Constraints on low-mass WIMP interactions on ^{19}F from PICASSO. *Phys. Lett. B*, 711:153–161, March 2012.

[23] M. Felizardo et al. The SIMPLE phase II dark matter search. *Phys. Rev. D.*, 89(072013), April 2014.

[24] E. Behnke et al. Erratum: First dark matter search results from a 4-kg CF_3I bubble chamber operated in a deep underground site. *Phys. Rev. D.*, 90(079902), October 2014.

[25] C. Amole et al. Dark matter search results from the PICO-2L C_3F_8 bubble chamber. *Phys. Rev. Lett.*, 114(231302), 2015.

[26] C. Amole et al. Improved dark matter search results from pico-2l run-2. arxiv.org/abs/1601.03729v1.

[27] R. Bernabei et al. New results from DAMA/LIBRA. *Eur. Phys. J. C*, pages 39–49, 2010.

[28] E. Aprile et al. Limits on spin-dependent WIMP-nucleon cross sections from 225 live days of XENON100 data. *Phys. Rev. Lett.*, 111(021301), July 2013.

[29] M. Felizardo et al. Final analysis and results of the phase II SIMPLE dark matter search. *Phys. Rev. Lett.*, 108(201302), 2012.

[30] E. Behnke et al. First dark matter search results from a 4-kg CF_3I bubble chamber operated in a deep underground site. *Phys. Rev. D.*, 86(052001), 2012.

[31] R. Agnese et al. Search for low mass WIMPs with SuperCDMS. *Phys. Rev. Lett.*, 112(24):241302, 2014.

[32] R. Agnese et al. Search for low mass WIMPs with SuperCDMS. *Phys. Rev. Lett.*, 112(24):241302, 2014.

[33] D. S. Akerib et al. First results from the LUX dark matter experiment at the sanford underground research facility. *Phys. Rev. Lett.*, 112(091303), 2014.

[34] P. Adriani et al. The cosmic-ray positron energy spectrum measured by PAMELA. arXiv:1308.0133v2, 2013.

[35] A. O. Keith. Supersymmetric dark matter after run i at the lhc: From a tev to a pev. arXiv:1510.06412.

[36] M. G. Boulay and A. Hime. Technique for direct detection of weakly interacting massive particles using scintillation time discrimination in liquid argon. *Astropart. Phys.*, 25(3):179–182, April 2006.

[37] E. Daw et al. The DRIFT directional dark matter experiments. In *Proceedings of the 3rd International conference on Directional Detection of Dark Matter (CYGNUS 2011)*, June 2011.

[38] D. N. Spergel. Motion of the earth and the detection of weakly interacting massive particles. *Phys. Rev. D.*, 37(6):1353–1355, March 1988.

[39] J. Monroe and P. Fisher. Neutrino backgrounds to dark matter searches. *Phys. Rev. D.*, 76(033007), 2007.

[40] J. Billard et al. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Phys. Rev. D.*, 89(023524), 2014.

[41] B. Morgan and A. Green. Directional statistics for WIMP direct detection II: 2-d read-out. *Phys. Rev. D.*, 72(12):123501, Dec 2005.

[42] J. Lindhard et al. Range concepts and heavy ion ranges. *Mat. Fys. Medd. K. Dan. Vidensk. Selsk*, 33(14):1–42. http://www.sdu.dk/media/bibpdf/Bind

[43] I. C. Wolfe. Measurement of work function in CF_4 gas, June 2010.

[44] L. G. Christophorou, J. K. Olthoff, and M. V. V. S. Rao. Electron interactions with CF_4. *NIST*, 1996.

[45] T. Caldwell et al. Transport properties of electrons in CF_4. arXiv:0905.2549v1.

[46] A. Kaboth et al. A measurement of photon production in electron avalanches in CF_4. *Nucl. Instrum. Meth A*, 592:63–72, 2008.

[47] C. Deaconu. *A Model of the Directional Sensitivity of A Model of the Directional Sensitivity of Low-Pressure CF_4 Dark Matter Detectors*. PhD thesis, MIT, 2015.

[48] K. A. Olive and others (Particle Data Group). The review of particle physics. *Chin. Phys. C*, 38(090001), 2014.

[49] G. J. Feldman and R. D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D.*, 57(7):3873–3889, Apr 1998.

[50] J.P. Lopez, D. Dujmic, S. Ahlen, J.B.R. Battat, C. Deaconu, et al. Background rejection in the DMTPC dark matter. *Nucl. Instrum. Meth*, A(696):121–128, 2012.

[51] SRIM: www.srim.org.

[52] N. Skoro et al. Low pressure breakdown and voltage characteristics of DC-discharge in CF_4. In *28th ICPIG*, 2007.

[53] M. A. Lieberman and A. J. Lichtenberg. *Principles of Plasma Discharges and Materials Processing*. John Wiley and Sons, second edition, 2005.

[54] J. R. Janesick. *Scientific Charged Coupled Devices*. SPIE Press, 2001.

[55] S. W. Henderson. *An Assessment of the Sensitivity of a Low Pressure Time Projection Chamber to the Direction of An Assessment of the Sensitivity of a Low Pressure Time Projection Chamber to the Direction of WIMP-Induced Nuclear Recoils*. PhD thesis, MIT, 2013.

[56] S. Ahlen et al. First dark matter search results from a surface run of the 10-l DMTPC directional dark matter detector. *Phys. Lett.*, B(695):124–129, 2011.

[57] A. C. Kaboth. *Detecting the Invisible Universe with Neutrinos Detecting the Invisible Universe with Neutrinos and Dark Matter*. PhD thesis, MIT, 2012.

[58] A. Lee. Point by point gain calibration. Technical report, MIT, September 2009.

[59] H. Yegoryan. Study of alpha background in a dark matter detector, June 2010.

[60] P. J. Rousseeuw and K. Driessen. Computing LTS regression for large data sets. *Data Mining and Knowledge Discovery*, 12:29–45, 2006.

[61] TMVA - multivariate toolkit: tmva.sourceforge.net.

[62] J.P. Lopez. *First Results from a 20-Liter Prototype Dark First Results from a 20-Liter Prototype Dark Matter Detector with Directional Sensitivity*. PhD thesis, MIT, 2014.

[63] E. Armangaud et al. Final results of the EDELWEISS-II wimp search using a 4 kg array of cryogenic germanium detectors with interleaved electrodes. *Phys. Lett. B*, 702:329–335, 2011.

[64] Z. Ahmed et al. Dark matter search results from the CDMS II experiment. *Science*, 327:1619–1621, March 2010.

[65] Z. Ahmed et al. Combined limits on WIMPs from the CDMS and EDELWEISS experiments. *Phys. Rev. D.*, 84(011102), 2011.

[66] S. Yellin. Finding an upper limit in the presence of unknown background. *Phys. Rev. D.*, 66(032005), 2002.

[67] UKDMC radioactivity test results.

[68] DEAP/CLEAN radon wiki.

[69] A. Ereditato et al. Design and operation of ARGONTUBE: a 5m long drift liquid argon TPC. *JINST*, 8:07002, 2013.

[70] M. Auger et al. The EXO-200 detector, part 1: Detector design and construction. *JINST*, 7, 2012.

[71] C. Geuzaine and J. F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.

[72] Elmer: www.csc.fi/web/elmer.

[73] Paraview: www.paraview.org.

[74] ROOT: root.cern.ch.

[75] Online Metals: www.onlinemetals.com.

[76] McMaster-Carr: www.mcmaster.com.

[77] Garfield: garfieldpp.web.cern.ch.

[78] S. W. Henderson. Cleaning procedures. Technical report, MIT, January 2013.

[79] V.K. Shen, D.W. Siderius, and W.P. Krekelberg. Nist standard reference simulation website, nist standard reference database, http://www.nist.gov/mml/csd/informatics˙research/srsw.cfm.