

Analysis and Modelling for CRESST-II

Andrew Brown
Christ Church, Oxford



A thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy at the University of
Oxford, Michaelmas Term 2011

Abstract

The dark matter search CRESST-II completed its most recent run, Run 32, in 2011. Compared to previous runs, the quantity of data taken in this run increased significantly. In this work, Oxrop, analysis software in use within the CRESST collaboration, is upgraded to analyse this new data. At the same time, Oxrop's internal structure is improved so that it can now handle data from detectors across different experiments consistently. This upgrade was performed with a view to developing Oxrop's candidacy for use with EURECA, a future dark matter experiment.

Oxrop is then used to model CRESST-II data.. First, light detector response to scintillation light produced in γ interactions in CRESST-II's target crystals is examined. A factor influencing detector efficiency is the time constant of scintillation light production, and this light detector examination is performed with a view to extracting the scintillation time constants of the target crystals. A simple model of light detector response of one exponential rise and two exponential decay times is initially considered. It is shown that this simple model does not closely match the light detector response to γ interactions in the crystal scintillator. Empirical extensions to this expected model are then made, allowing for additional decay times. These extensions allow the light detector response to crystal scintillator interactions to be well modelled, and allow estimates of the millikelvin γ scintillation time of CaWO_4 and ZnWO_4 . This model is then also applied to X-ray interactions directly in the light detectors. It is seen that, even with these model extensions, interactions directly in the light detector still show significant tension with the applied model. This implies that direct calibration of light detectors with X-rays is not possible without a further understanding of light detector response, or that future direct calibrations should be done with optical photons.

Position dependent effects in Run 32 calibration data are then studied. A phenomenon that has previously been considered as unrelated to position dependence, the anti-correlation effect between phonon and light detector signals, is shown to exhibit a position dependent effect in at least one light/phonon detector pair under study. Additionally, the collection efficiency of the light detector is shown to be related to the mean interaction position. Collection efficiency is found to reduce when mean interaction position is close to the cylindrical surfaces of CRESST's CaWO_4 target crystals. The magnitude of the difference in light collection efficiency between surface and bulk interactions is also seen to be correlated with high energy light detector resolution.

The WIMP-nucleon cross section limits resulting from the CRESST-II commissioning run (2007) are also reanalysed in this work. The original analysis of the commissioning run accounted only for tungsten recoils in the CaWO_4 crystals used in CRESST-II. Here, interactions from calcium and oxygen nuclei are also accounted for. The resulting WIMP-nucleon cross section limits were improved at light WIMP masses $\sim \mathcal{O}(10 \text{ GeV}/c^2)$. These limits show a mild tension with a recent dark matter analysis of Run 32, particularly for WIMP masses below $10 \text{ GeV}/c^2$. Possible causes of this tension are discussed.

Acknowledgements

There are a great many people without whom this work would have been impossible. Here is an insufficient attempt to thank those who made it possible for this thesis to exist.

First is Professor Hans Kraus. Hans has often taken the time out of his particularly busy research, tutoring and meeting schedule to offer me guidance in our frequent discussions about what work to do. Without his input, this thesis would have been very different and, simply, a much lesser piece of work.

Also thanks to Chris McCabe, his comments and input have always helpful and insightful, and without doubt improved my analysis. Vitaliy Mykhailik, Oxford's resident inorganic scintillator expert, also certainly deserves acknowledging for the advice he has given me. Richard McGowan, who I took over from, was vital in giving me my first grounding in CRESST analyses and Oxrop. Mike McCann, who taught me (most of?) what I was doing wrong in coding, and generally improved my research and my time researching throughout my D. Phil. My gratitude also to my other office mates: Alice Lynch, Phil Coulter, Xiaohe Zhang, James Imber, Stuart Ingleby and Matthew Malek, for their encouragement and help. Sam Henry also, though not technically an office-mate.

With enormous gratitude to my Mum and Dad. Though running off to Malaysia mid-way through my research, my parents have always been there to support me, and certainly I could not have written this thesis without them. To my brother James as well, though mostly for lending me his television. My friends: Iain Johnston, Nick Howarth, Kerrith Davies, Dragos Petrescu, and far too many more to name here. I have thoroughly enjoyed my time at Oxford, and it is certainly my friends that have meant this. Credit also to the girls, you know who you are.

I wish to thank Rafael Lang for his help getting me going on understanding CRESST detectors at the start of my research. Further, Jens Schmalzer's email discussions over the years have also been of great benefit to me, and I have appreciated his time in this. Within the CRESST collaboration, I had several happy times with many in the Tübingen crew, and my biggest regret of my time on CRESST is not having spent a few weeks working over in Tübingen during my D.Phil. Of course, this work would not have been possible without everyone in CRESST who actually physically made the experiment, and sat through the seemingly endless shifts in Gran Sasso. Everyone in the MPI, TUM and Tübingen groups involved in this has my sincerest gratitude.

I wish to thank STFC, who funded this work.

Dedicated to Bryn and Jack, my friends.

Contents

| | |
|---|-----------|
| Contents | 8 |
| Introduction | 1 |
| 1 Dark Matter | 2 |
| 1.1 The Dark Matter puzzle | 2 |
| 1.2 Evidence for Dark Matter | 3 |
| 1.2.1 Galaxy rotation curves | 3 |
| 1.2.2 Mass-to-light ratio of galactic clusters | 4 |
| 1.2.3 Colliding galactic clusters | 5 |
| 1.2.4 Cosmological evidence | 6 |
| 1.2.4.1 Cosmic Microwave Background and WMAP | 7 |
| 1.2.4.2 Big Bang Nucleosynthesis | 9 |
| 1.2.4.3 Structure formation | 10 |
| 1.3 Dark matter candidates | 11 |
| 1.3.1 Sterile neutrinos | 11 |
| 1.3.2 Axions | 12 |
| 1.3.3 MOND | 12 |
| 1.3.4 Standard model candidates | 13 |
| 1.4 WIMPs and Supersymmetry | 14 |
| 1.4.1 Motivation | 14 |
| 1.4.2 Abundance | 16 |
| 1.4.3 Detecting WIMPs | 17 |
| 1.5 Detection experiments | 19 |
| 1.6 Summary | 23 |
| 2 CRESST-II | 24 |
| 2.1 Overview | 24 |
| 2.2 Detectors | 26 |
| 2.2.1 Scintillators | 26 |
| 2.2.2 Detector principles | 29 |
| 2.2.3 Superconducting phase transition thermometers | 34 |
| 2.2.4 Measurement and stabilisation | 38 |

| | | |
|----------|---|-----------|
| 2.2.5 | Data acquisition..... | 40 |
| 2.2.6 | Calibration and background discrimination..... | 42 |
| 2.3 | Backgrounds and shielding..... | 45 |
| 2.3.1 | Cosmic Ray Muons..... | 46 |
| 2.3.2 | Shielding..... | 48 |
| 2.3.2.1 | Neutrons..... | 48 |
| 2.3.2.2 | Radon..... | 49 |
| 2.3.2.3 | Gamma and beta background..... | 50 |
| 2.3.2.4 | Electronic and vibrational noise..... | 52 |
| 2.4 | WIMP searches..... | 53 |
| 2.5 | Summary..... | 54 |
| 3 | Oxrop and Analysis..... | 55 |
| 3.1 | Overview..... | 55 |
| 3.2 | Software basis..... | 56 |
| 3.2.1 | ROOT..... | 56 |
| 3.2.2 | Oxrop functionality..... | 57 |
| 3.3 | Structural development..... | 62 |
| 3.3.1 | Oxrop 4..... | 63 |
| 3.3.2 | Oxrop 5..... | 67 |
| 3.4 | Analysis of the sub-1 MeV CRESST data..... | 72 |
| 3.4.1 | Energy estimates..... | 72 |
| 3.4.2 | Non-linear response corrections..... | 76 |
| 3.5 | Summary..... | 83 |
| 4 | Cuts for dark matter and detector response analyses..... | 84 |
| 4.1 | Introduction..... | 84 |
| 4.2 | Cut details..... | 85 |
| 4.2.1 | Stability cut..... | 85 |
| 4.2.2 | Rate..... | 88 |
| 4.2.3 | Trigger delay..... | 89 |
| 4.2.4 | Peak Position..... | 90 |
| 4.2.5 | DAQ range..... | 90 |
| 4.2.6 | Delta voltage..... | 91 |

| | | |
|----------|--|------------|
| 4.2.7 | Amplitude | 92 |
| 4.2.8 | Right – left of baseline..... | 92 |
| 4.2.9 | Peak Position – Peak Onset..... | 93 |
| 4.2.10 | RMS..... | 95 |
| 4.2.11 | Coincidence – Muon Veto..... | 97 |
| 4.2.12 | Coincidence - detector | 98 |
| 4.3 | Live-time | 99 |
| 4.4 | Total cut effects and summary | 99 |
| 5 | From light detector response to scintillation time..... | 102 |
| 5.1 | Introduction | 103 |
| 5.2 | Determining scintillation time | 104 |
| 5.3 | Standard-pulse treatment | 106 |
| 5.4 | Results..... | 107 |
| 5.5 | Extended model of light detector response..... | 110 |
| 5.6 | Extended model fit results | 111 |
| 5.7 | Direct calibration..... | 115 |
| 5.8 | Conclusions..... | 118 |
| 6 | Position dependence in CRESST target crystals..... | 120 |
| 6.1 | Introduction | 121 |
| 6.2 | Modelling detected signals..... | 121 |
| 6.3 | Position dependence and anti-correlation studies review | 125 |
| 6.4 | Anti-correlation measurement method | 128 |
| 6.4.1 | Thorium-232 calibration backgrounds..... | 132 |
| 6.4.2 | Cobalt-57 calibration backgrounds..... | 134 |
| 6.4.3 | Mean penetration depth estimate..... | 135 |
| 6.5 | Measurement | 137 |
| 6.5.1 | Cobalt-57 Calibration..... | 138 |
| 6.5.2 | Thorium-232 Calibration..... | 139 |
| 6.5.3 | Collection efficiency..... | 143 |
| 6.5.4 | Light detector resolutions | 146 |
| 6.6 | Conclusions and perspectives | 148 |

| | | |
|-------------------------|--|------------|
| 7 | Extending commissioning run limits to lower masses | 152 |
| 7.1 | Introduction | 152 |
| 7.2 | Including Oxygen and Calcium..... | 153 |
| 7.3 | Maximum gap | 158 |
| 7.4 | WIMP limits..... | 161 |
| 7.5 | Possible sources of systematic error in CRESST background analyses | 163 |
| 7.5.1 | Electron recoil leakage..... | 163 |
| 7.5.2 | Alpha quenching factor..... | 166 |
| 7.5.3 | Low energy behaviour of the lead recoil spectrum | 167 |
| 7.5.4 | Long decay or thermal relaxation events | 170 |
| 7.6 | Conclusions | 174 |
| 8 | Conclusions and the future | 175 |
| Appendices | | 178 |
| Appendix A | Oxrop UML diagrams | 178 |
| Appendix B | Calculating detector live time | 180 |
| Appendix C | Probability distributions integrals..... | 182 |
| Appendix D | Radial penetration depth estimates..... | 184 |
| Appendix E | Alternative light detector resolutions | 186 |
| References | | 187 |

Introduction

What is Dark Matter? On all astronomic scales, there is evidence that normal, baryonic matter constitutes only about a fifth of the total amount of matter in the Universe. The remainder is attributed to “Dark Matter”, an as yet unknown substance, existing in far larger quantities than the types of matter so far observed on Earth. Many possible candidates have been proposed for this unknown substance, with the Weakly Interacting Massive Particle, or WIMP, currently the most popular.

There are strong theoretical motivations for believing WIMPs represent a solution to the dark matter problem. But WIMPs are also particularly interesting for the experimentalist, simply because their existence can be directly verified here on the Earth. Searching for WIMPs through their interactions with ordinary matter targets has become a competitive field, with many technologies currently employed. This thesis focuses on data taken from one such WIMP search experiment, CRESST-II. CRESST-II is a cryogenic, scintillating dark matter search, which uses both light and phonon detectors to search for dark matter. CRESST-II is also a precursor experiment to EURECA, a future dark matter experiment. After introducing both dark matter and CRESST-II to the reader, this thesis first focuses on developments in Oxrop, software used to analyse CRESST-II data.

Later chapters then use Oxrop to model phenomenological aspects of CRESST-II data. Individual light detector behaviour is studied to see what information can be gleaned from how these detectors respond to particle interactions. Then, position dependence effects in CRESST-II’s target crystals are studied. In particular it is shown that detected light can vary depending upon where in a CRESST-II target crystal a particle interacts. Lastly, data from the CRESST-II commissioning run (2007) is reanalysed, using only published data, improving the constraints that CRESST-II can set on the existence of light mass WIMPs.

1 Dark Matter

Evidence for dark matter on galactic, galactic cluster and cosmic scales is reviewed. A brief summary of popular dark matter candidates is then presented, focusing on WIMPs. The expected rate of WIMP interactions in the example of a generic calcium tungstate detector is discussed.

1.1 The Dark Matter puzzle

The first indication of an excess of non-luminous, gravitationally interacting matter in the Universe was found by Oort in 1932 [1]. Oort measured the velocities of stars near the galactic plane. He found that some stars were moving fast enough to escape the galaxy given only the galaxy's luminous mass, implying a large unseen but gravitationally interacting component in the galaxy. One year later, Zwicky also found evidence for dark matter using a separate method [2]. Using the virial theorem to estimate the mass of the Coma galaxy cluster, Zwicky found the ratio of gravitating mass to luminous mass to be a factor 400. While this value was too high, with modern values closer to 180 [3], it still appears today that there is more gravitating mass in galactic clusters than interacts electromagnetically. A few years later, Babcock's estimates of the Andromeda galaxy indicated a mass-to-light ratio of Andromeda to be about 50. Babcock obtained this high number from the observation of a fast circular velocity in the outer parts of the galaxy. This was yet another indication of non-luminous matter in the Universe [4].

Since then, evidence for invisible, "dark matter" has been observed on astronomic scales ranging from the galactic up to the cosmological. This chapter summarises the main motivations for believing dark matter exists in the Universe. For further reading, see the reviews [5-7].

1.2 Evidence for Dark Matter

1.2.1 Galaxy rotation curves




Fig. 1.1 Galaxy Rotation curve NGC 3198. Expected behaviour if the matter distribution followed that of visible matter alone is shown by the “disk” curve. The observed velocity distribution is better modelled by contributions from both the visible disk and an invisible, dark matter halo. Halo contribution to the rotation velocity is shown, with the dark matter distribution modelled by $\rho(r) \propto 1/\left(1 + \left(\frac{R}{a}\right)^Y\right)$ [8].

The primary piece of evidence for dark matter on galactic scales is found in the rotation curves of spiral galaxies. Galactic rotation curves show the circular speed of rotation as a function of distance from the centre of the galaxy, exemplified in Fig. 1.1. Circular velocity is usually measured from the Doppler shift in the 21cm line of hydrogen, necessitating a galaxy to be viewed at-least partially side-on. To maintain stable orbits, the circular velocity of stars should behave as:

$$v(r) = \sqrt{\frac{GM(r)}{r}}, \quad (1.1)$$

with the mass enclosed within radius r given by:

$$M(r) = \int_0^r \rho(r') 4\pi r'^2 dr'. \quad (1.2)$$

If galactic mass were dominated by baryonic matter, the circular velocity of galaxies should fall off as $v(r) = \sqrt{\frac{GM_{Tot}}{r}} \propto \frac{1}{\sqrt{r}}$ beyond the central, visible matter disk. However, circular velocity remains approximately constant, implying a matter distribution beyond the central disk behaving approximately as $\rho(r) \propto \frac{1}{r^2}$. This

observation can then be understood if dark matter exists in a spherical halo in spiral galaxies, extending beyond the visible disk.

1.2.2 Mass-to-light ratio of galactic clusters

Evidence for dark matter can be seen when comparing the total galactic cluster mass to total emitted light. The total mass can be determined from several methods. The oldest, used by Zwicky, is to use the virial theorem:

$$2\langle T \rangle = -\langle V \rangle, \quad (1.3)$$

that is, for a gravitationally bound, dynamically relaxed system, twice the total, time averaged kinetic energy is equal to the negative of the potential energy. By measuring the Doppler shifts of spectral emission lines, cluster velocity dispersion, and thus cluster kinetic energy can be determined. This in turn allows for potential energy and mass to be estimated.

Another method is to use gravitational lensing. Mass curves space-time in its vicinity, and as light from distant sources passes through a cluster, light paths are bent along the geodesics of this distortion. Detailed images of the mass distribution in galactic clusters can then be constructed by measuring the distorted image(s) of light sources behind the cluster.

A third method is to measure the temperature of the ionised gas within the galactic cluster through X-ray emission. Assuming the gas is in hydrostatic equilibrium, the mass of the galactic cluster can then be deduced. Finally, a new method has been developed using the Sunyaev-Zel'dovich effect [9]. This is where CMB photons inverse Compton scatter off high speed electrons in a cluster's ionised, contributing to the anisotropy of the CMB.

These total matter estimates are then compared to baryonic matter estimates in the galactic cluster. Baryonic matter totals are themselves estimated using the

emitted light, correcting for absorption and the dimming of distant galaxies. Mass-to-light ratio estimates for galactic clusters are then typically 200 – 300 [5], implying a large contribution from unseen matter.

1.2.3 Colliding galactic clusters




Fig. 1.2a The Bullet Cluster. The matter distribution of the hot gas is shown, determined from X-ray emission, superimposed with the gravitational potential outline (green contours) determined from weak gravitational lensing. Gravitational centres are separated from the baryonic mass centres by about 8σ [10]. Fig. 1.2b Shows a similar image for the cluster MACS J0025.4-1222. An offset is seen between gravitational potential centres (red contours) and X-ray emission (yellow contours) by 4σ . White indicates galaxy position, determined from radio emission, which appears to follow gravitational centres [11].

Colliding galactic clusters provide some of the most convincing evidence for dark matter, particularly over its main competitor MOND (section 1.3.3). Two cases of this are the “Bullet Cluster”, and MACS J0025.4-1222, both shown in Fig. 1.2. In both cases, baryonic mass distribution is compared to the total mass distribution. $\sim 90\%$ of a cluster’s baryonic mass exists in a hot intra-cluster gas. The position of this gas, and thus the position of most of the baryonic mass, can be determined from X-ray emission. Gravitational lensing can then be used to separately determine the total mass distribution.

In colliding galactic clusters, hot gas within each cluster interacts through ram pressure. Dark matter and galaxies, on the other hand, are almost collision-less.

This behavioural difference can cause the position of the hot gas to become offset from the position of the dark matter and galactic content during a collision. Evidence for this separation, seen in Fig. 1.2, provides a strong indication of an additional un-seen matter component, dominating the total matter distribution.

However, in counter-point to this evidence, it should be noted that another colliding cluster, Abell 520, dubbed the “Train Wreck Cluster”, indicates understanding of cluster collisions is not complete. In Abell 520, gravitational centres exist in regions without galaxies. This is not normally expected in collision-less dark matter models, where the galaxy distribution is expected to closely follow the dark matter distribution [12].

1.2.4 Cosmological evidence

Considering cosmological evidence for dark matter, it is first useful to introduce a few important terms in which the density components of the Universe are often described. The Universe’s total matter and energy densities are usually expressed in terms of the critical density ρ_{crit} , itself given by:

$$\rho_{crit} \equiv \frac{3H_0^2}{8\pi G_N}, \quad (1.4)$$

where H_0 is the *Hubble parameter*, relating velocity to proper distance, and G_N is the gravitational constant. H_0 is often given in terms of $h \cdot 100$ km/s/Mpc, with a recent estimate of h being 0.702 ± 0.014 [13]. Individual contributions to the density content of the Universe are then expressed by:

$$\Omega_i = \frac{\rho_i}{\rho_{crit}} \quad (1.5)$$

for each matter, energy or dark energy contribution to the density of the Universe ρ_i . The total density of the Universe is the sum of all these components, $\rho_{Tot} = \sum_i \rho_i$. The relation of this total density ρ_{Tot} to the critical density ρ_{crit} determines the curvature of the Universe. If $\rho_{Tot} > \rho_{crit}$, the Universe will be closed, if $\rho_{Tot} < \rho_{crit}$, then it will be open, if $\rho_{Tot} = \rho_{crit}$, then it will be flat.

1.2.4.1 Cosmic Microwave Background and WMAP




Fig. 1.3 WMAP 7-year Cosmic Microwave Background temperature anisotropy map, with dipole and galactic foreground emission removed. From [14].

Wilkinson Microwave Anisotropy Probe (WMAP) is a satellite-borne experiment aimed at measuring the Cosmic Microwave Background (CMB). The CMB is an imprint of the big-bang, originating from the time when electrons and protons combined to form hydrogen. This was approximately 380,000 years after the big bang, when the temperature of the Universe was at $0.26 \text{ eV}/k_B$ [15]. This marks where the Universe becomes transparent to radiation, and is the point of last scattering for CMB photons. After this point, the scattering length for CMB photons is longer than the size of the Universe, and these photons can then still survive today.

The CMB today is a near perfect black-body spectrum, with a temperature of $2.725 \pm 0.002 \text{ K}$. The largest deviation from this spectrum is a dipole, at about 1 part in 1000, caused by the peculiar motion of the Earth relative to the CMB rest frame. Subtracting this, and the next largest contribution from emission by the Milky Way, an image of temperature anisotropies in the CMB is formed (Fig. 1.3) [16]. At the time of last scattering, CMB photons were coupled to baryons as a relativistic fluid in the Universe. In such fluids, density is closely related to

temperature. Anisotropies in the CMB are the result of density inhomogeneity in the Universe at the time of last scattering. These fluctuations are on the order $\sim 10^{-5}$, and this tiny variation itself a strong support of the principle that Universe is everywhere homogenous and isotropic, known as the cosmological principle. This temperature map can then be expressed in terms of its spherical harmonics, decomposing into an angular temperature spectrum (Fig. 1.4).




Fig. 1.4 WMAP 7-year Cosmic Microwave Background temperature spectrum. The red line indicates the best fit in the Λ CDM model. From [13].

The position of the first peak determines the total density of the Universe ρ_{Tot} . The denser the Universe, the larger the angular scale the first peak in the anisotropy spectrum is seen (towards the left hand side of Fig. 1.4). Several other processes contribute to the other structure seen in Fig. 1.4, and two effects are particularly relevant here. The first is baryon acoustic oscillations, where the competing effects of gravitation, tending to collapse the early Universe, and radiation pressure, tending to drive it apart, form oscillations. Larger baryonic content tends to increase compression in the ionized plasma of the early Universe with respect to rarefaction. This increases the heights of the odd numbered peak (e.g. first from the left of Fig. 1.4) against the even numbered peaks (e.g. second from the left of Fig. 1.4). Comparisons between the heights of the first and second peaks then indicates the fraction of the Universe in baryons.

Another important effect is the Sachs-Wolfe (SW) effect, the effect of gravitational red and blue shifting of the CMB by density perturbations in the Universe. The integrated SW effect is seen since the time of last scattering, and is related to the evolution of large scale structure in the Universe since that time, and is important to the large angular scale structure of the CMB. As a photon falls into a gravitational potential it blue-shifts, and as it leaves it red-shifts. There is only a net effect if the potential changed while the photon is in transit. This occurs only when the Universe is *not* matter dominated, when there is a significant dark energy component. Measurement of the integrated SW effect allows dark matter and dark energy components of the Universe to be resolved. The primary results of Cosmological interpretations of the CMB relevant for this work are:

The total density is in agreement with the critical density:

$$\Omega_{Tot} = 1.080^{+0.093}_{-0.0801}.$$

The baryon content is small:

$$\Omega_b = 0.0449^{+0.0028}_{-0.0028}.$$

The dark matter content is much larger than the baryon content:

$$\Omega_c = 0.228^{+0.027}_{-0.027}.$$

1.2.4.2 Big Bang Nucleosynthesis

Today, the relative amounts of ^2H , ^3He , and other light nuclei in the Universe depend upon the baryon density shortly after the Big Bang. In the very early Universe, protons and neutrons exist in equilibrium. As the Universe cools, the weak charged-current interaction, that maintains this equilibrium, reduces in rate. At about $T \sim 0.7$ MeV, the interaction rate becomes too low to maintain equilibrium. The neutron-proton ratio then freezes out and, after neutron β decay, leaves about one neutron for every seven protons. At temperatures lower still, $T \sim 100$ keV, photons are rarely energetic enough to dissociate deuterium, and the epoch of nucleo-synthesis begins.

The vast majority of neutrons remaining after freeze-out will end up in ${}^4\text{He}$. The fraction of ${}^4\text{He}$ in the Universe is then relatively insensitive to total baryon density. For deuterium however, a higher density of baryons leads to a faster rate of helium production and a lower fraction of deuterium in the Universe today. Processes that remove deuterium from the Universe, such as stellar nuclear fusion, are well known. Measurement of today's deuterium content then allow a reconstruction of the Universe's total baryonic matter fraction. Such measurements give $\Omega_b h^2 = 0.021 \pm 0.001$ from deuterium [17], compatible with the WMAP 7-year result of 0.02255 ± 0.0054 [13].

1.2.4.3 Structure formation

Structures - galaxies, galaxy clusters and super-clusters - are expected to form hierarchically, bottom-to-top. Small quantum fluctuations, magnified by inflation, are expected to collapse into larger structures by gravity. This process then determines the current large scale structure of the Universe. The final structure is dependent on the power spectrum of these density fluctuations. Measurements of large scale structure today can infer this power spectrum, and can be used to estimate the separate baryonic and non-baryonic components of the universe. These measurements support a total universal matter content much higher than the baryonic content, with $\Omega_m = 0.231 \pm 0.021$ and $\Omega_b/\Omega_m = 0.185 \pm 0.046$ measured by the 2dF Galaxy Redshift Survey [18].

Further, simulations of how structure forms have confirmed the need for cold dark matter. Structure formation models must account for several effects; by itself, gravity is relatively easily modelled, but these models must also take account of gas dynamics, radiative cooling, and other effects. To model structure evolution, it is then necessary to form so-called N-body simulations, such as in [19]. In such simulations, the presence of dark matter is needed to reproduce observed structure. Further this dark matter must be "cold" (non-relativistic),

else the dark matter particles have enough energy to escape galactic gravitational potentials, leading to a top-down formation of structure in the Universe.

1.3 Dark matter candidates

Now that evidence for dark matter is seen to exist at a variety of astronomical scales, it is natural to ask what this matter could actually be. At the same time as these astronomical observations, extensions to the standard model of particle physics have become necessary. The discovery that neutrinos have mass, the strong CP problem and the hierarchy problem, have all suggested new physics beyond the standard model. Extending the standard model of particle physics to solve these problems often provides suitable dark matter candidates. Today, the idea that dark matter is particulate, cold and non-baryonic, and that “dark energy” also exists, forms the Λ CDM model, often described as the *cosmological standard model*. Weakly Interacting Massive Particles (WIMPs) are the most popular candidate to be this dark matter, and calculations of expected abundance and the possibility of their detection is discussed later in this chapter. Before detailing WIMPs, several other popular dark matter particle candidates are introduced. While it will be shown that WIMPs are well motivated, other possibilities are examined to demonstrate that WIMPs are not the only possible constituent of dark matter in the Universe.

1.3.1 Sterile neutrinos

In the standard model, neutrinos are exactly massless. However, the observation of massive neutrinos allows for the existence of sterile, right handed neutrinos as an addition to the standard model [20]. In this scenario, neutrinos could acquire their mass through the “see-saw” mechanism by the introduction of electro-weak singlet(s) into the Standard Model Lagrangian [21]. Sterile neutrinos could function as hot, warm or cold dark matter. Limits on their existence come directly from neutrino oscillation data, and as dark matter candidates by a combination of

background & cluster X-ray measurements, which attempt to measure sterile neutrino decay, and Lyman- α forest small-scale structure formation measurements [7].

1.3.2 Axions

Without fine-tuning CP violation is expected in the strong sector of the standard model [22]. Strong CP violation gives rise to a significant neutron electric dipole moment. However, experimental limits show this dipole moment to be very small, with $|d| < 2.9 \times 10^{-26}$ [23]. This implies only a very small amount of CP violation in the strong sector. To solve this fine-tuning problem, the CP violating term in the QCD Lagrangian can be understood in terms of a scalar, Axion field [24,25]. The field gives rise to a massive, scalar, electrically neutral particle that would interact with matter very rarely through the strong and weak forces, similarly to a neutral pion. Axions are predicted to have a very small mass (less than about 10 meV). To be a cold dark matter candidate, Axions would then have to be produced non-thermally in the big bang. Detection experiments focus on searching for anomalous rotation of polarised light (such as PVLAS [26]) or the conversion of Axions into photons in the presence of a magnetic field (such as CAST [27] or ADMX[28])

1.3.3 MOND

Models exist which theorise no non-baryonic matter component to the Universe at all. These models are in general modifications of gravity or Newtonian physics, known collectively as Modified Newtonian Dynamics (MOND). MOND was introduced to explain galactic rotation curves without recourse to unseen matter [29]. The guiding principle was that at very small accelerations, force may not scale linearly with acceleration, such that:

$$\vec{F} = m\vec{a}\mu(a/a_0). \quad (1.6)$$

While variants of MOND deal with galaxy rotation curves very well [30], MOND has persistently had problems when dealing with the behaviour of galactic clusters [31]. The most recent problems have been with the Bullet Cluster. As noted in section 1.2.3, this colliding galactic cluster has its centre of visible, baryonic matter separated from its gravitational centres, ruling out many MOND theories which do not include non-visible matter [32].

1.3.4 Standard model candidates

Early attempts to solve the dark matter problem focused on unseen *baryonic* dark matter in the form of Massive Compact Halo Objects (MACHOs). Possible MACHO candidates are brown dwarfs or black holes, objects containing significant amounts of baryonic matter but not emitting much light. However, searches for MACHOs were consistent with null results [33], and given the cosmological result that $\Omega_b \ll \Omega_m$, the idea that MACHOs were responsible for the entirety of dark matter was abandoned.

Following the discovery that neutrinos were massive [34], neutrinos also became a possible candidates for dark matter, having the “undisputed virtue of being known to exist” [5,35]. However, neutrinos are very light, with mass constraints from a combination of WMAP-7 year, Sloan Digital Sky Survey, and Hubble telescope data giving $\sum_{i=1}^3 m_\nu < 0.41 \text{ eV}$ [15]. Neutrinos would then be relativistic, “hot dark matter”. Hot dark matter tends to destroy small scale structure, causing a top-down formation of structure, not supported by observations [36]. Using the relation of [13] that $\Omega_\nu h^2 = \sum \frac{m_\nu}{94 \text{ eV}}$, one sees that $\Omega_\nu h^2 < 0.0044$, limiting the neutrino mass contribution to less than about a fifth of the baryonic.

1.4 WIMPs and Supersymmetry

1.4.1 Motivation

A primary motivation for *supersymmetry* (SUSY) is the Hierarchy Problem: why is the Higgs mass so small? In the standard model, fundamental particles (quarks, leptons, weak gauge bosons) acquire their mass through interactions with the Higgs field. Spontaneous symmetry breaking in the electro-weak sector is then caused by this Higgs field, and leads to the expectation that a massive, scalar “Higgs boson” exists. The Hierarchy Problem is the problem that, without fine tuning, one would expect the mass of the Higgs boson to be of order the reduced Planck mass, $M_p = 2.4 \times 10^{18}$ GeV, due to loop corrections in the Higgs mass [37]. However, the Higgs mass is tightly constrained, with $114.4 \text{ GeV} < M_H < 185 \text{ GeV}$ by LEP [38], and now also limited to not be in the range $141 \text{ GeV} - 476 \text{ GeV}$ by a combination of ATLAS and CMS results [39].

The difference in scale between the non-fine-tuned expectation of the Higgs mass and the experimental expectation of the Higgs mass is the Hierarchy problem. To solve this problem, supersymmetry is introduced. In this theory, for each particle in the standard model, a super-symmetric partner differing in spin by $\frac{1}{2}$ exists, collectively called *sparticles*. Loop corrections to the Higgs mass arising from interactions with these sparticles then stabilise the Higgs mass.

To preserve the observation that lepton and baryon number are closely conserved, many SUSY models introduce a conserved quantity, R-parity. R-parity is given by $R = (-1)^{3(B-L)+2s}$, where B and L are the baryon and lepton numbers of a particle and s the particle spin. By convention, standard model particles have R-parity of +1, SUSY sparticles -1 . R-parity conservation implies that sparticles must be pair-produced in collider experiments, and that the Lightest SUSY Particle (LSP) must be stable. The characteristics of the LSP

depend upon the mechanism in which SUSY is broken. If SUSY was a perfect symmetry, sparticles would have identical masses to their standard model partners. As such identical mass sparticles have not been observed to date, SUSY must be a broken symmetry. The form of the LSP depends on the way in which SUSY is broken [37]. Many SUSY breaking mechanisms have the LSP as the neutralino, a linear composition of the neutral gauginos. The neutralino does not interact electromagnetically or strongly, only weakly and gravitationally, and is thus a WIMP dark matter candidate.




Fig. 1.5 Summary of exclusion limits set by CMS on the existence of super-symmetry, from [40].

SUSY searches at ATLAS and CMS have so far only reported null results, with a summary of ATLAS results shown in Fig. 1.5. However, the LHC is yet to fully probe the full SUSY parameter space. Yet, there is reason to believe that SUSY may be very difficult to find. In supersymmetry one expects a measurable electron-dipole moment. The current strong experimental limits on this dipole moment [40] rule out many super-symmetry models.

However, with respect to WIMPs, it is also important to note that the SUSY neutralino is not the only generic WIMP dark matter candidate. Particles of similar behaviours arise in Kaluza-Klein [41], little Higgs [42] and asymmetric dark matter [43] models. Despite the null-results of collider and electron dipole moment experiments so far limiting super-symmetry, the generic WIMP is still a highly motivated candidate due to simple estimates of WIMP abundance, discussed now.

1.4.2 Abundance

Thermal production of neutralinos in the big bang could leave us with dark matter today. In the early Universe, the temperature would have been high enough that the LSP, χ , would exist in thermal equilibrium. This equilibrium is maintained through the competing processes of pair-production and annihilation. While equilibrium holds, the number density of χ will decrease as the Universe expands and cools, until the rate of annihilation per particle:

$$\Gamma = \langle \sigma_{\chi\chi \rightarrow ff} v \rangle, \quad (1.7)$$

falls below H , the rate of universal expansion. $\langle \sigma_{\chi\chi \rightarrow ff} v \rangle$, is the velocity averaged total cross section for the annihilation of χ into standard model particles. At this point, χ “freezes-out”; χ number density still decreases, but only due to the effects of universal expansion. χ number density is then too dilute to maintain thermal equilibrium. χ density today can then be estimated from the Boltzmann equation, through the Friedman equations:

$$\frac{dn}{dt} = -3Hn - \langle \sigma_{\chi\chi \rightarrow ff,eff} v \rangle (n^2 - n_{equ}^2), \quad (1.8)$$

where the change in χ density per unit time, $\frac{dn}{dt}$, is in terms of the density decrease from expansion $-3Hn$, and n_{equ}^2 the equilibrium χ density. By considering χ density when the annihilation rate Γ is equal to the rate of expansion of the Universe H at freeze-out temperature T_f , it can be shown that [44]:

$$\Omega_\chi h^2 \cong \frac{n_\chi m_\chi}{\rho_{crit}} = \frac{3 \times 10^{-07} \text{ cm}^3 \text{ s}^{-1}}{\langle \sigma_{\chi\chi \rightarrow ff,eff} v \rangle}. \quad (1.9)$$

The density of χ is then approximately not related mass of χ , and is inversely proportional to the velocity averaged annihilation cross section $\langle\sigma_{\chi\chi\rightarrow ff,eff}v\rangle$. Using the weak interactions strength for $\langle\sigma_{\chi\chi\rightarrow ff,eff}v\rangle$ of $\mathcal{O}(10^{-0f}\text{cm}^3\text{s}^{-\text{m}})$ then approximately gives the correct dark matter abundance, typically taken to be compatible with a 50-1000 GeV/c² WIMP [6].

1.4.3 Detecting WIMPs

This section deals with how WIMPs may be practically observed on Earth, following [45]. The first quantity to be calculated is the differential rate of WIMP-nucleon recoils seen in the detector:

$$\frac{d\Gamma}{dE_R} = N_T \frac{\rho_\chi}{m_\chi} \int f(\vec{v}) \left(v \frac{d\sigma}{dE_R}(v, E_R) \right) dv, \quad (1.10)$$

where $\frac{d\Gamma}{dE_R}$ is the differential interaction rate with respect to the recoil energy. This is expressed in terms of the number of target nuclei, N_T , the WIMP number density $\frac{\rho_\chi}{m_\chi}$, and the velocity averaged differential cross section (Integral term). $f(\vec{v})$ is the WIMP velocity distribution, taken as the Maxwell-Boltzmann distribution in the isothermal WIMP halo model, truncated at the galactic escape velocity v_{esc} . The recoil energy, E_R is then:

$$E_R = \frac{m_{Red}^2}{m_n} v^2 (1 - \cos(\theta^*)) = E_R^{Max} \frac{(1 - \cos(\theta^*))}{2} = 2m_n q^2. \quad (1.11)$$

Eqn. 1.11 shows E_R as different expressions in terms of the centre of mass frame scattering angle θ^* , reduced mass $m_{Red} = \frac{m_\chi m_n}{m_\chi + m_n}$, the energy at maximum momentum transfer E_R^{Max} and momentum transfer, q . The cross section is then:

$$\frac{d\sigma}{dE_R}(v, E_R) = \frac{\sigma_0}{E_R^{Max}} F^2(q), \quad (1.12)$$

where σ_0 is the point-like cross section for coherent WIMP-nucleon interactions, and $F^2(q)$ is the square of the nuclear form factor, accounting for non-point like nuclear behaviour at non-zero momentum transfer. In this work, the Helm form factor as parameterised in [46] is used, shown for CaWO₄ in Fig. 1.6.

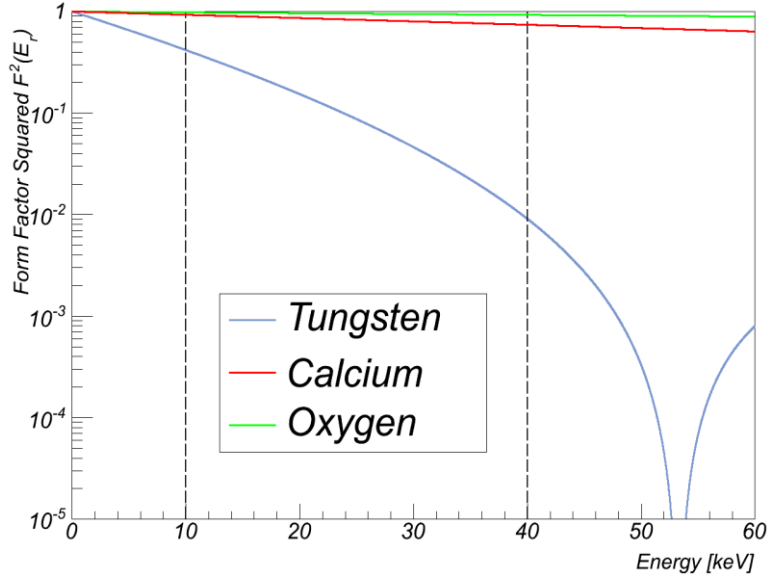


Fig. 1.6 Form Factors for calcium (red), oxygen (green) and tungsten (pale blue), as a function of recoiling nuclear energy, as described in [46]. The tungsten form factor drops towards a minimum at ~ 53 keV.

Eqn. 1.11 and 1.12 then allow Eqn. 1.10 to be re-expressed as:

$$\frac{d\Gamma}{dE_R} = N_T \frac{\rho_\chi}{m_\chi} \frac{m_\chi \sigma_0}{2m_{Red}^2} F^2(q) \int \frac{f(\vec{v})}{v^2} dv = N_T \frac{\rho_\chi}{m_\chi} \frac{m_\chi \sigma_0}{2m_{Red}^2} F^2(q) \mathbb{I}(v_{min}, v_\odot, v_{esc}), \quad (1.13)$$

where the integral \mathbb{I} is related to the velocity distribution function. \mathbb{I} is a function of the galactic escape velocity v_{esc} , the velocity of the solar system relative to the local system v_\odot , and the minimum velocity (v_{min}) of WIMP that can cause recoil energy E_R , ($\cos(\theta^*) = -1$ in Eqn. 1.11). The non-rotating WIMP halo model of [45] is used, with $v_{esc} = 650$ km/s, $v_\odot = 232$ km/s, $\rho_\chi = 0.3$ GeV/cm $^{-m}$.

The total cross section is the sum of a spin independent term σ_{SI} and a spin dependent term σ_{SD} . In the case of CaWO_4 , the majority of nuclei have no net spin, and only the spin-independent term is considered. The cross section is then expressed in a *per nucleon* basis, with target nucleon number A :

$$\sigma_{SI, WIMP-nucleon} = \left(\frac{1+m_\chi/m_N}{1+m_\chi/m_p} \right)^2 \frac{\sigma_0}{A^2}. \quad (1.14)$$

For a given WIMP mass, target nucleus and cross section, the interaction rate may then be calculated. This is done for CaWO_4 in Fig. 1.7. In real detectors, two more effects need to be accounted for. The first is detector energy resolution, which modifies Eqn. 1.13 for energy resolution ΔE as:

$$\frac{d\Gamma}{dE_R} = \int_0^\infty \frac{1}{\Delta E \sqrt{2\pi}} e^{-\frac{(E_R - E'_R)^2}{2\Delta E^2}} \frac{d\Gamma}{dE'_R} dE'_R. \quad (1.15)$$

The second effect is acceptance, the probability A_i that WIMPs scattering with recoil energy E_R from a given nuclei will be observed at a given energy. The observed rate is then:

$$\left. \frac{d\Gamma}{dE_R} \right|_{Obs} = \sum_{i=0}^N A_i \frac{d\Gamma_i}{dE_R}, \quad (1.16)$$

for all of the N possible target nuclei.

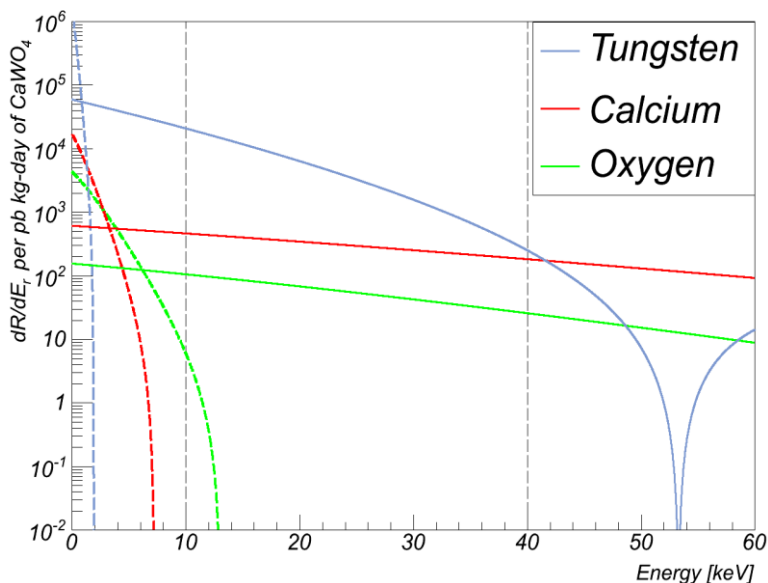


Fig. 1.7 Expected WIMP interaction rate $d\Gamma/dE_R$ for target nuclei Calcium (red), Oxygen (green) and Tungsten (light blue). Values are scaled to their relative constituent mass fractions in CaWO_4 , per kg-day of CaWO_4 and picobarn. The solid lines represent expected rates for 100 GeV/c^2 WIMPs, dashed lines 5 GeV/c^2 WIMPs. Approximate detector acceptance region limits are indicated by dashed black lines. Within the acceptance region, tungsten recoils dominate for 100 GeV/c^2 WIMPs, and oxygen for 5 GeV/c^2 WIMPs.

1.5 Detection experiments

Modern WIMP searches focus on direct detection, looking for rare WIMP-nucleon scattering, and indirect detection, evidence of WIMP annihilation in gravitating bodies (e.g. galaxy centres or the Sun). In both classes, many experimental collaborations have formed to hunt dark matter, and reviews of the dark matter detection field may be found in [47,48].

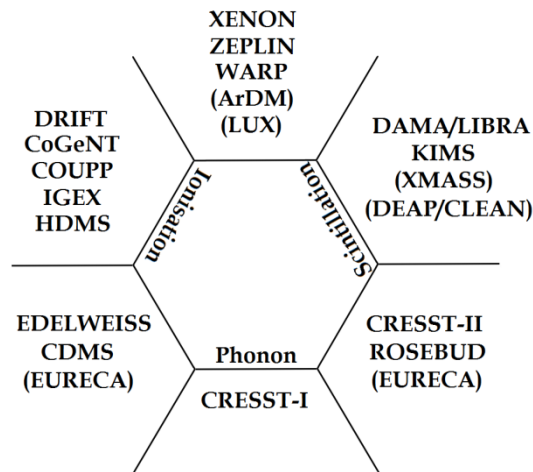


Fig. 1.8 Direct dark matter detector experiments, arranged by detection channel. Direct detection dark matter experiments are based around phonon, ionisation or scintillation detection technologies (or a combination of two of these). This list is non-exhaustive, showing several experiments that have produced WIMP limits or favoured regions [49-64] as well as some future experiments [65-69].

Current direct detection experiments limit the spin-independent WIMP-nucleon cross section at $100 \text{ GeV}/c^2$ to be no higher than 10^{-8} pb [49]. With this cross section, a target mass of (for example) 10 kg of CaWO_4 would see only $\mathcal{O}(10)$ events per year. Expected WIMP interaction rates are then very low. For experiments to be sensitive to these low rates, detectors must be shielded from natural α , β and γ , cosmic ray and neutron backgrounds. It is never possible to completely shield against these backgrounds, and so most dark matter search experiments have the capability to *discriminate* between background radiation and dark matter interactions. Broadly, direct detection dark matter experiments can be split between those that look for signal in phonon, ionization or scintillation light channels (or a combination of two of these), as in Fig. 1.8.

As of late 2011, three dark matter search experiments claim a signal compatible with elastic, spin-independent WIMP scattering: DAMA [51], CoGeNT [63] and an analysis of the new run of CRESST-II [53]. DAMA searches for an annual modulation by looking at its scintillation signal count rate, using sodium iodide target crystals. As the Earth rotates around the Sun, so too will the WIMP flux the Earth is exposed to vary; the Earth should see a larger flux of WIMPs close to the

2nd of June, where the Earth's motion around the Sun is most closely aligned with the Sun's motion around the galaxy. DAMA have reported a positive annual modulation signal when examining thirteen annual cycles [51]. CoGeNT also looks for an annual modulation, looking in the ionisation channel using a Germanium target. CoGeNT also see an annual modulation, though the resulting WIMP signature of [63] does not agree with DAMA when following standard assumptions about the dark matter halo (Fig. 1.9).




Fig. 1.9 Comparison of the favoured regions and 90% CL WIMP-nucleon cross section limits on elastic, spin-independent WIMP-nucleon cross sections from various experiments. The favoured regions from DAMA [51], CoGeNT [63] and CRESST-II (2011) [53] are shown. Along-side are WIMP cross section limits from CDMS II [57], CDMS II (low energy) [70], XENON100 [49], XENON10 (S2 only) [64], ZEPLIN-III [50] and EDELWEISS-II [56]. The extended CRESST-II limits are derived from the CRESST-II commissioning run, as outlined in chapter 7. With gratitude to F. Pröbst for providing the confidence contours of CRESST-II (2011), and similarly to C. McCabe for providing the DAMA contours.

Separately, CRESST-II looks for dark matter using a combination of both phonon and light signals, using scintillating calcium tungstate crystal targets, looking for a signal in excess of background. The CRESST-II collaboration has recently

published results from the new CRESST run claiming a signal compatible with WIMPs in [53]. The signal region claimed, however, does not appear to coincide with the CoGeNT or DAMA signals (Fig. 1.9).

Several other experiments partially or completely exclude the WIMP interpretations of DAMA, CoGeNT and the recent analysis of CRESST-II data. XENON100 [49], XENON10 [64] and ZEPLIN-III [50] look for WIMP interactions in two-phase cryogenic xenon detectors. These experiments all look for WIMP interactions in liquid xenon targets, using both scintillation and ionisation signals. By comparing these signals, most radioactive backgrounds can be distinguished from possible WIMP interactions signals, allowing these experiments to set limits on the existence of dark matter (Fig. 1.9).

CDMS-II [57], and EDELWEISS-II [56] look for WIMP interactions through the phonon and ionisation channels in Germanium crystal targets. By comparing signals in these two channels, background γ and β interactions in these experiments can be discriminated against possible WIMP interactions. CDMS-II and EDELWEISS-II also set limits in some tension with the WIMP interpretations of DAMA, CoGeNT and the analysis of recent data by CRESST-II. While there are several other experiments of interest, finally it is noted that the results of chapter 7 of this work show that the CRESST-II commissioning run (2007) are in mild tension with the WIMP interpretations of DAMA [51] and results of the most recent run of CRESST-II (2011) [53]. Thus, direct WIMP dark matter searches continue to be a competitive field, with no agreed upon conclusion yet.

1.6 Summary

Dark matter has remained an unsolved mystery for nearly 80 years. During this time, evidence for its existence has continued to mount up on all astronomical scales. While few possible dark matter candidates have been fully ruled out, WIMPs have become a popular solution to the dark matter problem. There are several compelling reasons to believe that super-symmetric WIMPs form dark matter. WIMPs solve the particle physics Hierarchy problem. At the same time, WIMPs give the observed dark matter abundance today simply by assuming dark matters annihilates with rates expected due to weak coupling. WIMPs can also be practically observed on Earth, and several dark matter experiments are on-going to hunt for these WIMPs. In the next chapter, we focus on one of these experiments, CRESST-II.

2 CRESST-II

CRESST-II is a direct dark matter search designed to detect WIMPs elastically scattering off one of several scintillating crystal dark matter targets. CRESST-II's detector operation, signal evolution and discrimination capabilities are outlined in this chapter. Shielding against natural backgrounds is also described.

2.1 Overview

CRESST is short for the Cryogenic Rare Event Search with Superconducting Thermometers. The original version of the experiment, CRESST-I, used Al_2O_3 (Sapphire) phonon detectors only. CRESST-I stopped taking data in 2002. Since then, scintillating calcium tungstate (CaWO_4), and now also zinc tungstate (ZnWO_4) target crystals have been used. Now, both phonon and light produced by particle interactions in these target crystals are measured by independent *phonon detectors* and *light detectors*. This second version of the experiment was dubbed "CRESST-II". For brevity, the "II" is dropped in the rest of this work.

The aim of CRESST is to see WIMPs in scintillating crystal targets. These targets are seen in the centre of a schematic diagram of CRESST in Fig. 2.1. CRESST detectors are calorimeters that measure the heat changes induced by particle interactions. This is directly in phonon detectors, and indirectly by the production of heat from the absorption of scintillation light in light detectors.

The expected nuclear recoil energies caused by WIMP interactions are very small, of $\mathcal{O}(10 \text{ keV})$. To detect these small energies CRESST is held at cryogenic temperatures, both to reduce detector heat capacity and allow the use of sensitive super-conducting phase transition thermometers ("SPT"s). This requires the use of a cryostat, seen at the top of this figure.




Fig. 2.1 Schematic cross section of CRESST. The detectors are located in the small volume at the centre of the figure, with the crystals (green), mounted on a copper detector carousel. Working horizontally outward, the thermal shields within the copper cold box are seen. Next is 14cm of high purity copper (orange), followed by 20cm of low background lead (grey). Surrounding this is the radon box (thin black line), and then the muon veto system (dark blue). Lastly a 45cm thick polyethylene neutron moderator surrounds the experiment (yellow). Water, also used as a neutron moderator is shown in light blue. At the top of the figure is the S1000 dilution refrigerator, separated from the detector carousel by a 1.3m long copper “cold finger”. Image taken from [71].

However, WIMPs are not the only interactions that will be seen in CRESST. Background interactions can also be seen from natural radioactive decays near to or in detectors, or from or induced by cosmic rays. To reduce these backgrounds requires several different levels of shielding, which constitute the majority of Fig. 2.1. In describing CRESST, first the operation of detectors is focused on. This is then followed by descriptions of the various backgrounds that can be seen in CRESST, and the shielding used to reduce these backgrounds.

2.2 Detectors

2.2.1 Scintillators

The targets for dark matter interactions in CRESST are scintillating CaWO_4 and ZnWO_4 crystals. When a particle (α, β, γ , neutron, WIMP, cosmic ray) interacts within the crystal, both phonons (heat) and light are produced. These scintillators produce different amounts of light when excited by different types of radiation – β / γ vs. α vs. neutron vs. WIMP – as a fraction of the energy of that radiation. Such scintillators can then be used to discriminate between different types of radiation [72].

The selection of a scintillator is performed with not only discrimination power in mind. For dark matter applications, the spin-independent, WIMP-nucleon cross section scales approximately with the mass of the target nucleus squared. This makes heavy target nuclei desirable. Further, a significant source of background events is often the crystals themselves; it is difficult to produce many inorganic scintillators with good radio-purity. Taking these factors into account, CRESST uses mostly CaWO_4 and also ZnWO_4 target crystals. In both cases, tungsten fulfils the heavy nuclear target criteria, and both crystals can be made with good radio-purity.

CaWO_4 and ZnWO_4 both fall within the broad class of crystalline, inorganic crystal scintillators which have similar scintillation mechanisms described here following [73]. Photons lose their energy when traversing the crystal through the photo-electric effect, Compton scattering and electron-positron pair-production (above 1.02 MeV). Charged particles lose energy through the ionisation of atoms in the crystal and bremsstrahlung. Neutrons lose energy through elastic scattering which can then cause knock-on electrons, or inelastic scattering which also leads to γ particle emission. WIMPs are hypothesised with several modes of

interaction, but are most frequently assumed to scatter elastically off crystal nuclei with similar nuclear recoil characteristics to neutron elastic scattering.

As photons lose energy through the photo-electric effect, Compton scattering and pair production, and charged particles lose energy from ionisation, their energy is distributed among multiple secondary particles within the crystal. These lower energy particles can then couple with electrons in the crystal lattice. This coupling can excite crystal electrons from the valence band into the conduction band creating an electron-hole pair or an *exciton*. If the exciton has less energy than the ionisation energy, the exciton will be trapped and cool through phonon emission until the electron reaches the bottom of the conduction band and the hole the top of the valence band¹. At this point, excitons can interact with a *luminescent centre*, where they can produce light (below). If the exciton had more energy than the ionisation energy, then the exciton is initially free within the crystal lattice, eventually becoming trapped by a defect or recombining at a luminescent centre, where they can also produce scintillation light.

All inorganic crystal scintillators contain luminescent centres. Scintillators have energy levels of radiative emission at luminescent centres between the valence and conduction bands in the crystal, so as to avoid reabsorption of any emitted light. In the case of CaWO_4 and ZnWO_4 , the luminescent centres are the $(\text{WO}_4)^{2-}$ and $(\text{WO}_6)^{2-}$ molecular complexes respectively. The actual mechanism of exciton coupling to luminescent centres varies between crystal scintillators. In CaWO_4 and ZnWO_4 , excitons interact through charge transfer, further described in [74].

The scintillation emission spectra of CaWO_4 and ZnWO_4 are shown in Fig. 2.2. The scintillation efficiencies of CaWO_4 and ZnWO_4 for γ interactions are $\sim 8.3\%$ and $\sim 10.0\%$ at cryogenic temperatures, respectively [75].

¹ In fact, electron-hole pairs can also be produced with *less* energy than the crystal band gap, due to the presence of crystal defects or impurities.




Fig. 2.2 Emission spectra of CaWO_4 (top) and ZnWO_4 (bottom) when excited by different photon sources at 9K, from [76]. Photon excitation energies are: 4.4 eV (1), 4.5 eV (2), 7.1 eV (3), 31 eV (4) and 3.2 keV (5). Scintillation photons are of optical frequency for both CaWO_4 and ZnWO_4 . Insets in both figures show the change in scintillation light intensity with temperature when using a 3.2 keV photon source. Scintillation light output in both cases increases with decreasing temperature.

Scintillation emission is generally characterised by a sum of several (two to three) exponential decays, such that scintillation emission may be described by [75]:

$$f(t) = y_0 + \sum_{i=1} A_i e^{-t/\tau_i} , \quad (2.1)$$

with y_0 modelling background and A_i the weight of each component. CaWO_4 and ZnWO_4 can both be described by one long decay time $\mathcal{O}(100 \mu\text{s})$ and one or more short $\mathcal{O}(1 - 10 \mu\text{s})$ decay times. The relative weighting of the long decay time $A_i \tau_i$ component dominating at low temperatures, such as those of CRESST. [75]. Evolution of the long decay component in both crystals is show in Fig. 2.3.




Fig. 2.3 Evolution of decay time constants of CaWO_4 (filled) and ZnWO_4 (empty) with temperature, from [76]. Detector response was modelled as two or three exponential decays, with decay times shown. At lower temperatures, the dominant long decay time component tends towards slower decay times.

2.2.2 Detector principles

As noted, different types of radiation (α , β , γ , neutron, WIMP, cosmic ray) interact through different mechanisms. These different mechanisms lead to different fractions of deposited energy ending up in emitted scintillation light, a characteristic that is exploited to make background-discriminating dark matter detectors. To perform discrimination, both scintillation light and total deposited energy are measured through separate cryogenic calorimeters. The calorimeters are called light detectors (top of Fig. 2.4) and phonon detectors (middle of Fig. 2.4), and are used for measuring light and heat respectively. Only a few percent of the energy of interactions in the target crystal is converted into scintillation light [75], with remaining energy going into phonons. The approximation is then made that detection of scintillation light can be used solely for discrimination purposes, and detection of phonons alone can provide an energy estimate.




Fig. 2.4 Diagram of a CRESST detector module. The phonon detector (crystal target and thermometer) and light detector (light absorber and thermometer) are shown. Each thermometer is thermally connected to the heat bath. Modules are surrounded by a scintillating, reflecting foil. Image from [71].




Fig. 2.5 Image of an opened CRESST module. On the left is the light detector, with the small tungsten thermometer just visible on the surface. On the right in the phonon detector, with the thermometer evaporated onto the crystal surface. Note the scale – CRESST detectors are quite small, with the target crystal measuring 4cm in diameter and 4cm in height. Image from [71].

The phonon detector consists of the target crystal with an attached phonon thermometer. In most phonon detectors, the thermometer is a thin film of tungsten in its α phase, evaporated directly onto the crystal surface. However, several crystals are of a more recent design, where the phonon thermometer is first evaporated onto a small *carrier* crystal that is then glued onto the larger *target* crystal [77]. The thermometer in all cases is a super conducting phase transition (SPT) thermometer.

Light detectors have two main components - a light absorber a tungsten SPT thermometer. In most cases, the light absorber is a silicon-on-sapphire wafer, 4cm in diameter. The exceptions are where the light absorber is made of silicon only. Light detectors operate similarly to the phonon detectors. Incident energy (scintillation light) is converted to phonons by the absorber. The temperature change of these phonons are then measured by an associated SPT. Real light detectors are more complicated, with additional components such as phonon collectors to improve phonon collection efficiency. Detailed descriptions of CRESST light detectors can be found in [78-80].

Light detectors are paired with phonon detectors. They are placed facing the crystal face opposite to the phonon thermometer. The detector pair is then placed inside of reflective, scintillating housing. On the inner surface of the housing is a layer of 3M VM2002 foil. Using values in [81,82], the foil reflects about 95% of incident light, or possibly as high as 98% from [79]. The foil is used to improve light collection efficiency. The foil is also scintillating; any ionizing radiation exciting the foil will produce scintillation light. This effect can be used to discriminate surface alpha-decays from signal events, using a method further described in [83].




Fig. 2.6 Diagram of the CRESST Carousel. Up to thirty three detector modules can be installed on the copper CRESST Carousel. Diagram taken from [83].

The combined phonon/light detector pairing inside the reflective housing is then named a detector *module*. Thirty three such modules can be installed in the CRESST set up, up to three in each of the twelve towers that form the detector *carousel* (Fig. 2.6). Module positions are labelled by their tower number, between 01 and 12, and their position, bottom (B), middle (M) and top (T).

In the most recent CRESST run, “Run 32”, eighteen modules were installed in an arrangement shown in Fig. 2.7. Of these eighteen, ten were fully active: 01T, 04B, 06B, 06T, 08B, 08M, 09B, 09M, 10B and 10T. Of the ten fully active modules, nine are CaWO_4 modules, and one is a ZnWO_4 module (08B). The remaining eight modules are almost all inactive due to non-functional light detectors. This is due to the cryostat not reaching a low enough temperature for the light detector thermometers to enter their transition. The exception is 01M, where the phonon detector “VK32” was not functional. Table 2.1 lists the positions and some basic parameters of each detector module.

| Module | Detector | Detector type | Material | Resolution @ 122 keV [keV] | Baseline Resolution [keV] | Detector trigger rate [mHz] | Module trigger rate [mHz] |
|--------|----------|---------------|-------------------|----------------------------|---------------------------|-----------------------------|---------------------------|
| 01T | VK33 | P (glued) | CaWO ₄ | 0.92 | 0.20 | 101.1 | 107.1 |
| | Franz | L | SOS | 5.43 | 0.70 | 93.4 | |
| 04B | Verena | P | CaWO ₄ | 0.77 | 0.21 | 109.0 | 233.9 |
| | Burkhard | L (*) | SOS | 5.76 | 1.10 | 187.8 | |
| | Q | L (*) | SOS | 6.61 | 0.79 | 78.5 | |
| 06B | Maja | P (glued) | CaWO ₄ | 1.05 | 0.21 | 201.3 | 229.7 |
| | Hans | L (*) | SOS | 6.84 | 0.68 | 146.4 | |
| 06T | Sabine | P | CaWO ₄ | 0.76 | 0.19 | 80.2 | 192.8 |
| | Josef | L | SOS | 6.62 | 1.35 | 70.0 | |
| 08B | ZnWO4-4 | P (glued) | ZnWO ₄ | 3.95 | 1.31 | 78.6 | 130.6 |
| | Ulrich | L (*) | Silicon | 5.93 | 1.61 | 104.4 | |
| 08M | Wibke | P | CaWO ₄ | 1.36 | 0.54 | 87.9 | 107.8 |
| | X | L (*) | SOS | 6.90 | 0.52 | 106.3 | |
| 09B | K07 | P | CaWO ₄ | 0.75 | 0.12 | 86.4 | 88.3 |
| | David | L | SOS | - | - | 65.6 | |
| 09M | Daisy | P | CaWO ₄ | 0.53 | 0.11 | 73.4 | 74.4 |
| | SOS08 | L | SOS | 6.18 | 1.74 | 55.2 | |
| 10B | Rita | P (glued) | CaWO ₄ | 0.74 | 0.32 | 85.2 | 158.1 |
| | Steven | L (*) | SOS | 4.24 | 0.42 | 156.4 | |

Table 2.1 Detector properties table for operational detectors in Run 32 from which data is used anywhere in this thesis. Detector types are split between Phonon (P) and Light (L) detectors. Some phonon detectors are glued detectors, as noted by "(glue)" (see text). Additionally, some light detectors have been outfitted with Fe-55 X-ray sources (noted by "(*)"). Phonon detectors are made using either calcium tungstate (CaWO₄) or zinc tungstate (ZnWO₄) as the target crystal. Light detectors are either of the silicon-on-sapphire (SOS) or pure silicon design (see [79] for details). Detector resolutions are taken from a Co-57 γ calibration data, whereas background trigger rates are taken from a 14 day "training sets" (See chapters 3 and 4). The light detector "David" was not active during Co-57 calibration, and thus this resolution could not be measured.

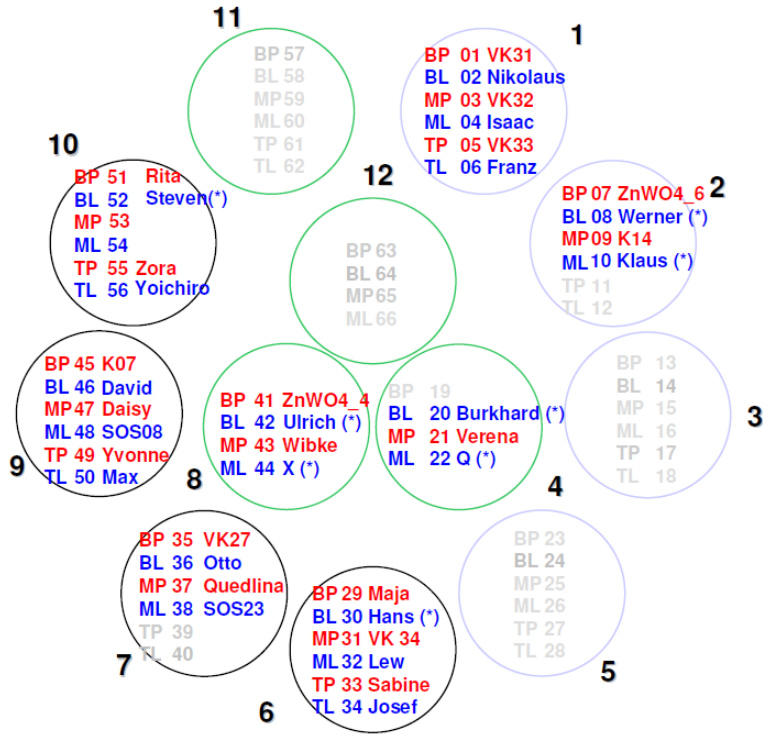


Fig. 2.7 Detector positions in CRESST's Run 32. Phonon detectors are shown in red, light detectors in blue. The first letter in the two letter code before each detector name represents the location Bottom, Middle or Top, the second letter whether the detector is a Phonon detector or Light detector. Light detectors marked with a star (*) were outfitted with an Fe-55 X-ray source on the opposite side to the phonon detector. One module, 04B, with phonon detector Verena and light detectors Burkhard and Q, has two light detectors. Burkhard is positioned above the phonon detector and Q below. The modules 02B (ZnWO₄_6 and Werner) and 08B (ZnWO₄_4 and Ulrich) use zinc tungstate based phonon detectors. All other modules use calcium tungstate based phonon detectors.

2.2.3 Superconducting phase transition thermometers

When a particle interacts in a detector, it deposits energy ΔE . This energy deposition will cause an approximate temperature change ΔT in the absorber:

$$\Delta T = \frac{\Delta E}{C}, \quad (2.2)$$

for heat capacity C . Energies of interest are of $\mathcal{O}(10 \text{ keV})$, and thus heat capacity C must be very small to allow measurable changes in temperature ΔT . At low temperatures, the heat capacity of non-conducting crystals follow the Debye model of crystal lattice heat capacity $C \propto (T/\theta_D)^3$, with Debye temperature θ_D . Thus, to reach low heat capacities, detectors must be cooled. In CRESST's case, detectors are cooled to very low temperatures of $\mathcal{O}(10 \text{ mK})$.

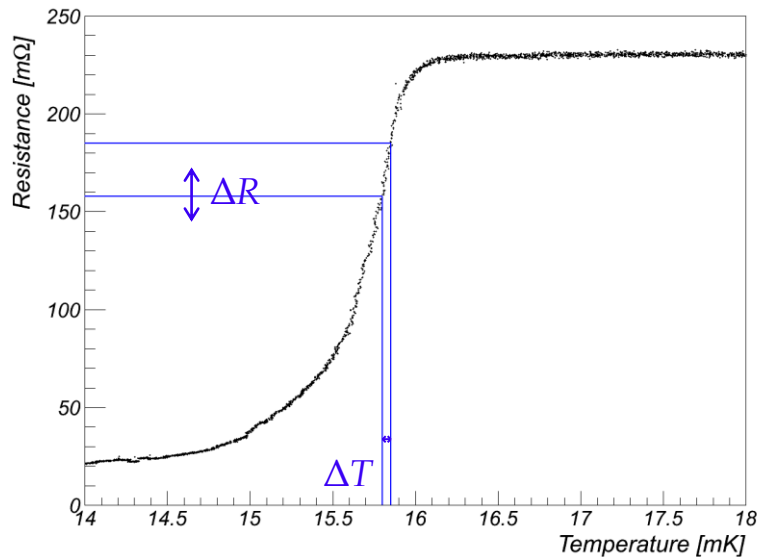


Fig. 2.8 Illustration of the operational principle of an SPT. Before the deposition of energy by a particle, the detector is held at a constant temperature – the “operating point”. Deposition of energy then causes an increase in temperature ΔT . This small temperature change causes a large resistance change, ΔR .

As noted, temperature changes are measured using SPT thermometers. Between particle interactions, thermometers are held at the same point within their transition between normally conducting and superconducting states. Energy deposition by incident particles causes a small change in temperature. This small temperature change then causes a large change in resistance. Measurement of this resistance change will be described in the next section.

This description is of course simplistic. A more detailed description is now given, summarising that in [84]. When radiation deposits energy in the absorber, high frequency a-thermal phonons are created. The absorber in the phonon detector is the target crystal. In the light detector, it is the silicon of the light absorber. For ionising radiation on time scales 1 ns, these phonons decay to acoustic phonons of about half the Debye frequency of the absorber. This frequency is > 1 THz [84,85] for all absorbers in question, so these phonons are still a-thermal. Particle induced nuclear recoils (neutrons, WIMPs) produce a range of a-thermal phonons, with frequencies up to the Debye frequency. In both cases, these phonons decay to frequencies $\mathcal{O}(100$ GHz) within a few tens of microseconds. These phonons, still a-thermal, do not decay further on time scale \sim ms.

Also on time scales of tens of microseconds, these athermal phonons will have travelled to fill the absorber uniformly after a few surface reflections. These then form a bath of athermal phonons connected to the thermometer. The fate of these phonons is then to be absorbed by the electrons in the thermometer, or decay in the absorber to thermal phonons, or escape into the heat bath undetected. If these phonons enter the tungsten film, they interact strongly with the electrons in the thermometer, thermalizing and raising the thermometer's electronic temperature. This gives rise to a fast time constant change in temperature of the electrons in the SPT. Part of this energy then escapes to the heat bath, and part is re-radiated to the absorber as thermal phonons. The eventual transfer of thermal phonons in the absorber back into the thermometer gives rise to a second, slower thermal time constant.

The resulting time dependence of detector response can then be modelled quantitatively. This is done by considering the coupled power inputs to the absorber and thermometer, assuming an instantaneous deposition of energy, first done in [84].

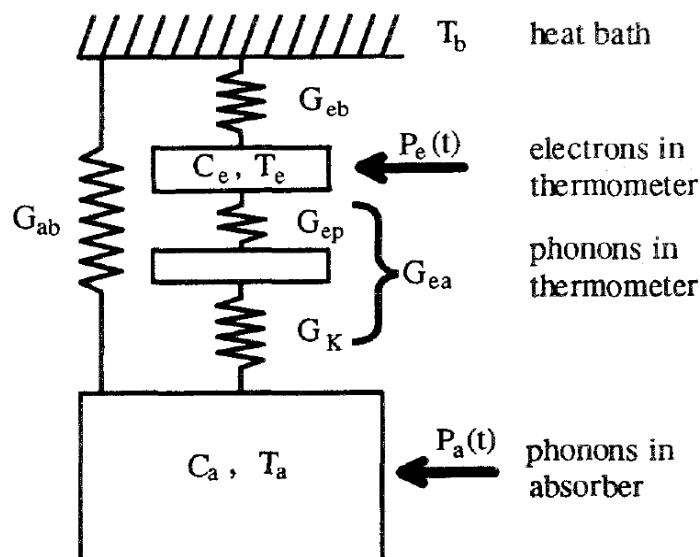


Fig. 2.9 Model of phonon transport in CRESST detectors, taken from [84]. Both the phonon system of the absorber (bottom) and electron system of the thermometer (upper-middle) are connected to a heat bath (top) held at temperature T_b , with thermal conductances G_{ab} and G_{eb} respectively. The thermometer electron and absorber phonon systems see power inputs $P_e(t)$ and $P_a(t)$, and have heat capacities $C_e(t)$ and $C_a(t)$.

The electron system of the thermometer and the phonon system of the absorber are described by temperatures and heat capacities T_e, C_e and T_a, C_a respectively (Fig. 2.9). Electron and phonon systems of the thermometer are treated as separate systems. At low temperatures and electron-phonon thermal conductance between them is expected to be $G_{ep} \propto T^5$ [84]. The phonon systems of the thermometer and absorber are coupled with Kapitza thermal conductance $G_K \propto T^3$ [86]. As thermometer phonon system heat capacity is negligible, coupling between the absorber phonon system and the thermometer electron system is effectively:

$$G_{ea} = \left(\frac{1}{G_{ep}} + \frac{1}{G_K} \right)^{-1} \quad (2.3)$$

The electron system of the thermometer is coupled to the heat bath by the gold bond wire with electron-bath thermal conductance $G_{eb} \propto T$. Lastly, the absorber is connected to the heat bath through the clamps, with Kapitza, absorber-bath thermal conductance $G_{ab} \propto T^3$. Energy depositions by particles give rise to high-frequency, non-thermal phonons. Power input into the detector system by these non-thermal phonons is then modelled as split into two fractions¹. A fraction $(1 - \epsilon)$ of non-thermal phonons is modelled to thermalize in the absorber, giving a power input to the absorber $P_a(t)$. The remaining fraction ϵ thermalizes in the thermometer, providing a power input $P_e(t)$ directly to thermometer electrons. $P_e(t)$ and $P_a(t)$ are then:

$$P_e(t) = \Xi(t)e^{-t/\tau_n} \text{ and } P_a(t) = \frac{(1-\epsilon)}{\epsilon} P_e(t), \quad (2.4)$$

introducing the non-thermal phonon population decay time, τ_n , arising from thermalisation in both absorber and thermometer, and $\Xi(t)$ the step function. By neglecting thermal conductance along the thermometer (i.e. assuming G_{eb} is solely modelled by thermal conductance of the gold bond wire), the thermal model above can be expressed by two coupled equations:

¹ This model assumes thermalisation of non-thermal phonons is frequency independent. [84] also notes that phonons should be treated in terms of a range of frequencies, however, it is assumed that describing these phonons using only two classes is sufficient to model detector response.

$$C_e \frac{dT_e}{dt} + (T_e - T_a)G_{ea} + (T_e - T_b)G_{eb} = P_e(t), \quad (2.5a)$$

$$C_a \frac{dT_a}{dt} + (T_a - T_e)G_{ea} + (T_a - T_b)G_{ab} = P_a(t), \quad (2.5b)$$

which leads to a model of the thermometer temperature change ΔT_e of:

$$\Delta T_e = \Xi(t) \left[A_n \left(e^{-t/\tau_n} - e^{-t/\tau_r} \right) + A_t \left(e^{-t/\tau_t} - e^{-t/\tau_n} \right) \right], \quad (2.6)$$

with two components: an athermal component A_n and a thermal component A_t . This equation also introduced the thermal phonon lifetime τ_t and the detector rise time τ_r . $\Xi(t)$ models the instantaneous deposition of energy. An important result is that this *shape*, as described by the time constants τ_n , τ_t , τ_r and the ratio A_t/A_n , is expected to be independent of both the type of interacting particle and the energy of the interaction. Different energy interactions are expected to cause a simple linear scaling of maximum temperature change ΔT_M with interaction energy, as in Eqn. 2.2.

2.2.4 Measurement and stabilisation

Temperature changes are expected to be linear with energy in the thermometer, but what of resistance changes? Expected temperature changes for low energies are very small, of $\mathcal{O}(10 \text{ nK})$. Most SPT thermometers have transitions of width $\mathcal{O}(1 \text{ mK})$, and almost always have *linear regions* over which thermometers can be operated¹. In these regions, resistance changes are approximately linearly related to temperature changes. Changes in resistance of a film are measured using a Superconducting Quantum Interference Device (SQUID), in a circuit illustrated in Fig. 2.10.

From Fig. 2.10, the current through the SQUID input coil, given a thermometer resistance R_T and two reference resistors of resistance $R_R / 2$, is then:

$$I_S = I_0 \frac{R_T}{R_T + R_R}, \quad (2.7)$$

¹ There is a certain freedom in selecting the precise operating point (position within the transition) in which to operate the thermometer. There is also freedom of choice in bias current (I_0 in Fig. 2.10). Both parameters affect response linearity, and are optimised with linearity (among other factors) in mind.

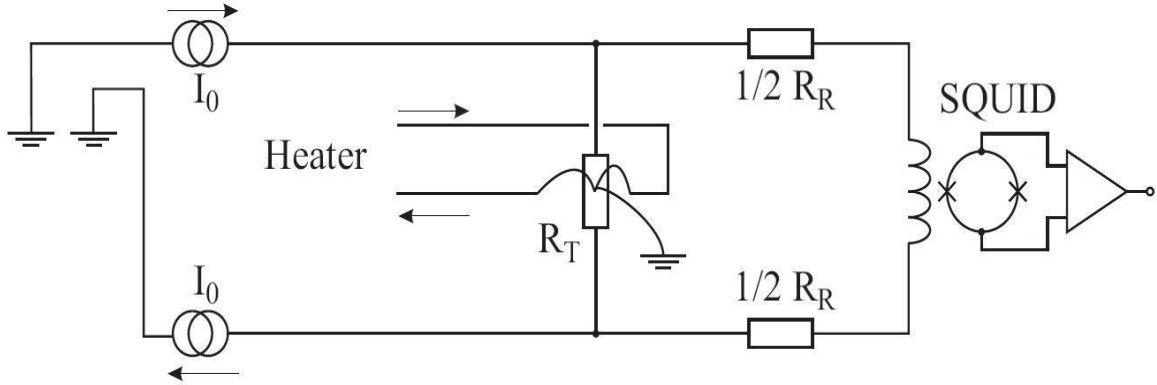


Fig. 2.10 Illustration of the thermometer circuit. In the outer, thermometer circuit, two reference resistors of equal resistance are placed in series with the SQUID input coil. Two are used instead of one to minimise cross-talk. Changes in thermometer resistance then induce changes in current through the SQUID input coil. Thermometer stability is maintained with the inner heating circuit. Image taken from [87].

for small changes in temperature and resistance $\Delta T \propto \Delta R_T \ll R_T$. The change in current through the SQUID input coil is then to first order:

$$\Delta I_S = I_0 \frac{R_R}{(R_T + R_R)^2} \Delta R_T \propto I_0 \frac{R_R}{(R_T + R_R)^2} \Delta T. \quad (2.8)$$

Small changes in temperature then cause approximately linear changes in current through the SQUID input coil. SQUIDs are operated in flux locked mode. Changes in current through the input coil are then directly related to changes in SQUID output voltage, to within an integer number of flux quantum $h/2e$.

Detectors are stabilised through a separate heater circuit, shown on Fig. 2.10. To operate all detectors, the base temperature of the carousel must be lower than the operating point of the detector with the lowest transition temperature. To maintain a specific operating temperature, each detector then has an attached heater structure. Heat is provided by gold wire that remains resistive at millikelvin temperatures. Across this heater is a constant heating bias, maintaining thermometer temperature. Additionally, heating *control pulses* are sent every 3s to measure operating point position. A control pulse drives a thermometer out of its linear region, and the resulting change in voltage, the *control pulse height*, is measured. This pulse height forms a measure of current operating position in the transition, and is held fixed by changing the heating

bias. By doing so, thermometers are maintained at the same operating point in their transition throughout an entire run.

Separately to these control pulses, heating *test pulses* are sent every 30s. This is to probe non-linear variation in detector response with injected energy. The combined, applied voltage of the test pulse and heating bias is first passed through a voltage square rooter (Eqn. 2.9). The energy input into the detector then scales linearly with the applied voltage (Eqn. 2.10):

$$I_{Au}(t) \propto \sqrt{V_{Bias} + \alpha V_{Pulse}(t)}, \quad P(t) = R_{Au} I_{Au}^2(t) \propto (V_{Bias} + \alpha V_{Pulse}(t)), \quad (2.9)$$

$$E = \int P(t)dt \propto (\int V_{Bias}dt + \int \alpha V_{Pulse}(t)dt), \quad \Delta E = \int \alpha V_{Pulse}(t)dt. \quad (2.10)$$

As applied test voltage is linearly related to energy, measuring detector response to test-pulses of different voltages allows the detector response - energy relation to be probed in detail, described further in section 3.4.

2.2.5 Data acquisition




Fig. 2.11 Simplified diagram of DAQ components relevant to particle pulse recording. SQUID signal is split into two outputs. One output is filtered and amplified, then fed into the trigger unit. The other is fed directly to the transient recorder. Readout of the transient recorder is instanced by the trigger unit.

Output from the SQUID is then split in two. One output is filtered and amplified, and then fed into a *trigger unit*. The other output fed is fed into a *transient recorder*. The memory buffer of the transient recorder is recorded to disk when a trigger is received from the trigger unit (Fig. 2.11).

Detectors trigger when the change in voltage on the trigger unit output exceeds a manually set *trigger threshold*. At this point, the transient recorder begins post trigger sampling for all other detectors on a module for the next 245.76ms. From the time of the first trigger, other detectors on the same transient recorder are allowed to trigger for a further 122.88ms, at which point the triggers on all detectors are deactivated for the remaining 122.88ms. There then follows an 8 ms period where the memory of the transient recorder is read out, or 16ms if two modules triggered on the same transient recorder. The trigger remains deactivated for a further two pre-trigger lengths, $2 \times 122.88\text{ms}$ for particle pulses, or one pre-trigger length 122.88ms for test pulses. This additional deactivation period is to ensure adequate baseline restoration (below) for subsequent recordings. The transient recorder read-out associated with a single detector and trigger then forms a *record*. This record constitutes the 122.8ms period before the trigger, the *pre-trigger* region, and the 245.76ms region following a trigger, the *post-trigger* region, for a total record length of 327.68ms. The pre-trigger region describes the state of the detector before a particle has interacted. The post-trigger region contains the response of a detector to a particle interaction. The interval between voltage readings in a record is $40\mu\text{s}$, and records then constitute 8192 voltage readings. A typical record is shown in Fig. 2.12. Several modules can share the same transient recorder. Transient recorders are not capable of recording each module with respect to its own trigger time. Thus, if several modules on the same transient recorder trigger, the onset of a pulse will occasionally be shifted relative to the pre and post trigger regions.

Briefly, it is noted that information from other sources than from the cryogenic (phonon and light) detectors described above is also recorded in CRESST. One source is the muon veto. When the muon veto triggers, the pulse height of the triggering panel(s) and a sum of all panels is also stored. This information can then be used to remove events seen in the cryogenic detectors that were coincident with muons from a dark matter analysis.




Fig. 2.12 Typical record of a 122.06 keV γ interaction in a phonon detector. Records consist of 8192 voltage readings separated in time by $40\mu\text{s}$, for a full record length of 327.68ms. Time shown is with respect to the point at which the first detector on the transient recorder triggered; times before zero define the pre-trigger region, after zero the post trigger region.

Information associated with control pulses, the heating pulses used to stabilise detectors, is also written out. Most importantly, this information includes the pulse height of these control pulses, itself a measure of the operating point of detectors. As control pulses are sent every few seconds, control pulse heights are then used to ensure detectors were at their correct operating point throughout the entire run. Lastly, environmental monitor information, such as air pressure and carousel temperature, is also recorded throughout a run. This allows for, for example, the course of cryostat cooling to be tracked in offline analyses.

2.2.6 Calibration and background discrimination

From the preceding sections, it is expected that detected changes in SQUID output voltage and deposited energy are linearly related. To set the constant of this energy scaling relation, and external Co-57 source emitting γ rays at 122.06 keV (and other energies) is introduced between the cold box and the copper shield. Phonon detector calibration is to direct interactions in the phonon detector, whereas light detector calibration is to the light detected due to 122.06 keV γ interactions in the phonon detector. Given that the light detector is then not

calibrated directly, light detector energy units are then often expressed in keV “electron equivalent” (keVee). Light detectors are typically calibrated such that one unit of energy deposited in the phonon detector by a γ interaction leads to a one unit keVee response. The light signal divided by phonon signal, the *light yield* $L_\gamma(E)$, is thus calibrated to unity. The full calibration process also attempts to remove non-linearity in detector response, discussed in detail in section 3.4.

With the energy scaling relations set by the detectors response to γ events, background discrimination can be discussed. The use of inorganic crystal scintillators for detecting dark matter was first discussed in [88]. The use of such crystals relies on the known difference between the amount of scintillation light produced by β or γ interactions when compared to neutron (or expected dark matter) interactions in such scintillators.

β and γ events, such as the 122.06 keV calibration γ rays, produce the largest amount of light of all interactions occurring in the target crystal. These interactions cause recoiling electrons in the scintillator target, which generate large amounts of light. The amount of light produced per unit energy by γ and β interactions is approximately constant with energy [89].

All other interactions (e.g. α , neutron), produce less light than γ or β interactions per unit energy, and are said to be *quenched* with respect to them. For example, fast neutrons tend to elastically scatter off nuclei in CaWO_4 . The resulting nuclear recoil produces less light than recoiling electrons.

This effect is illustrated in Fig. 2.13. The detector has been exposed to an AmBe neutron source, again placed between the copper cold box and the copper shield. The neutron source produced neutrons at energies ranging up to 11 MeV, which are primarily seen in CRESST through their elastic scattering off Oxygen, and to a lesser extent calcium and tungsten. Elastic neutron scattering off oxygen in

calcium tungstate produces almost a factor of 10 times less scintillation light than equivalent energy β or γ interactions. Elastic neutron scattering interactions then form the lower gradient band in Fig. 2.13.




Fig. 2.13 Typical AmBe neutron calibration, with light signal on the y-axis vs. phonon signal on the x-axis. The diagonal band consists of β and γ events. The peak in the middle of the γ / β band is caused by Pb-210 β decays, which mostly give 46.5 keV γ rays. Neutrons form the lower gradient band, with light output comparatively reduced by a factor of ~ 10 . Energy normalisation is with respect Co-57 calibration events, accounting for transition non-linearity (see section 3.4). Energies are thus calibrated with 122.06 keV γ scintillation events in the target crystal. “keVee” is the electron-equivalent energy of the light signal, as energies are normalised to the amount of scintillation light seen from the target crystal and not absolute energies (section 3.4).

WIMPs are expected to scatter primarily off tungsten, and the resulting tungsten recoil produces even less light. The reduction in light of nuclear recoils with respect to electron recoils is quantified by a quenching factor Q_i for each species i of recoiling nucleus. The expected amount of light from a nuclear recoil is:

$$L_{Q_i}(E) = Q_i L_\gamma(E) , \quad (2.11)$$

where $L_\gamma(E)$ is the light expected for β or γ interactions of the same energy. Nuclear recoils are then seen in a *quenched band*, with expected light yield defined by Eqn. 2.11. Important quenching factors are shown in Table 2.2:

| Element | Quenching factor Q_i [%] |
|----------|----------------------------|
| Oxygen | $11.09^{+0.909}_{-0.909}$ |
| Calcium | $6.38^{+0.619}_{-0.653}$ |
| Tungsten | $3.91^{+0.478}_{-0.430}$ |

Table 2.2 Selected quenching factors in CaWO_4 , taken from [80].

While quenching factors in Table 2.2 are the most recent, quenching factor measurements have varied between studies. For example, Q_O is measured to be $10.4^{+0.5}_{-0.5}\%$ in [53], $7.8^{+0.3}_{-0.3}\%$ in [90], $7.1^{+0.1}_{-0.1}\%$ in [91], and $Q_W < 3.0\%$ (2σ) in [90] and $2.5^{+0.2}_{-0.2}\%$ in [91]. Typically a neutron calibration is performed in each run to ensure that the oxygen quenched factor can be fixed for detectors in that run. However, due to the large cross section for oxygen scattering, it is not possible to perform such in-situ measurements for the calcium and tungsten bands with reasonable accuracy. These quenching factors must then be taken from dedicated measurements.

WIMPs tend to scatter off tungsten. As Q_W is very low, events in this band have negligible light signals compared to energy. While WIMPs are generally expected to scatter preferentially off tungsten in calcium tungstate, threshold effects (chapter 7) can mean that calcium and oxygen recoils are significant when performing dark matter searches for light WIMPs of mass $\mathcal{O}(5 \text{ GeV}/c^2)$.

2.3 Backgrounds and shielding

As “rare event search” implies, CRESST looks for interactions with exceedingly low rates. Indeed, the last chapter showed that the expected number of WIMP interactions that can be seen in CaWO_4 detectors is less than 10 events per 10 kg-years. A major challenge in any dark matter experiment then is to remove any background events that have even a small probability of mimicking a dark matter signal. A variety of such natural background sources exist, both terrestrial and extra-terrestrial in origin. While it is possible to discriminate many of these from dark matter interactions, it is preferable to reduce their occurrence altogether.

2.3.1 Cosmic Ray Muons

Cosmic ray muons are the most numerous charged particle caused by cosmic rays seen at sea level [92]. Both primary (extra-terrestrially sourced) and secondary muons are seen at sea level. Most are secondary, caused by the decay of charged mesons produced when high energy cosmic ray nuclei (mostly protons) interact with the atmosphere. Muons are a very dangerous background to direct dark matter searches. While muons can interact directly in detectors, a more dangerous background occurs when muons interact close to the detector producing neutrons. Muons can induce these neutrons through spallation interactions. This is dominantly where a real γ produced in an e-m shower photo-disintegrates a nucleus, and also where muon scatters in-elastically off a nucleus via virtual photon exchange. These processes can both produce free neutrons. Other processes, such as secondary neutron production, also contribute to the total neutron flux induced by muons on the earth [93].

In CRESST specifically, neutrons are a background to dark matter searches because they have a high cross section for elastic scattering off oxygen nuclei (and, dependent on energy, other nuclei) in the CaWO_4 and ZnWO_4 target crystals. Such recoils are very difficult to distinguish in the acceptance region from tungsten nuclear recoils that are expected from WIMPs heavier than $\sim 20 \text{ GeV}/c^2$. For lighter WIMPs $\mathcal{O}(5 \text{ GeV}/c^2)$, neutrons are impossible to distinguish from elastic WIMP scattering. This is because light WIMPs scatter almost exclusively off oxygen in the acceptance region. Thus, for both light and heavy WIMPs, it is vital to remove all possible neutron backgrounds.

By far the best defence against muons is to reduce this muon flux. This is most easily done by placing experiments deep underground. Muon flux decreases rapidly with depth, as can be seen in

Fig. 2.14. For this reason, all direct dark matter detection experiments are located far beneath the Earth's surface. CRESST itself is located in the Laboratori Nazionali del Gran Sasso (LNGS) deep underground laboratory, which has a minimum 1400m rock overburden. The cosmic ray muon flux on the earth's surface is $\approx 1 \text{ cm}^{-2} \text{ min}^{-1}$ [92]. The muon flux at Gran Sasso is reduced by ~ 6 orders of magnitude from this surface flux.

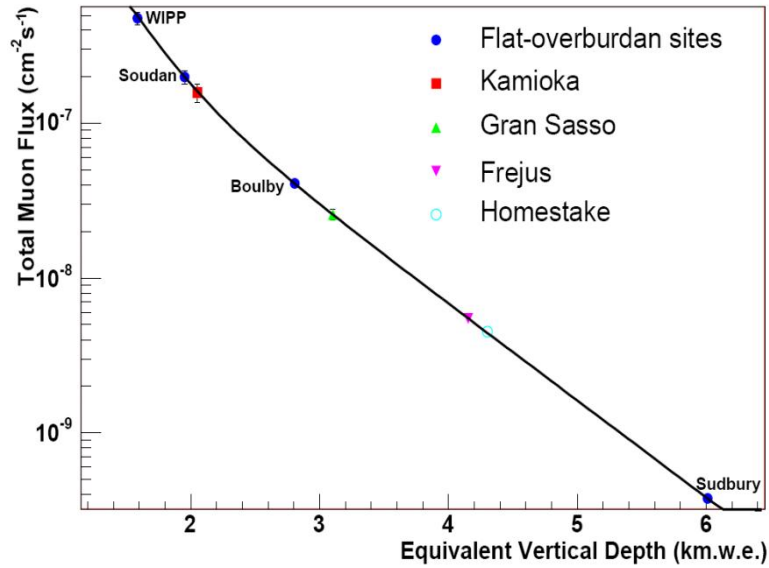


Fig. 2.14 The variation in cosmic ray muon flux with flat-overburden depth (water equivalent), from [94]. Deeper sites tend to see a drastically reduced cosmic ray muon flux, and then also tend to see a reduced muon induced neutron flux. Equivalent vertical depth for sites without a flat-overburden, such as LNGS, account for the variation in depth with incident angle.

Even at this depth however, cosmic ray muons can still induce neutrons in the surrounding rock or the experiment itself. Thus a muon veto system is installed around CRESST, located within the polyethylene shield (Fig. 2.1). The muon veto consists of twenty plastic scintillating panels surrounding the experimental volume. The total solid angle coverage of these panels is 98.7%. Each panel is read out through a single photo-multiplier tube. When a muon veto panel triggers due to an incident muon (or γ), information associated with that trigger, such as PMT pulse height, is recorded. Further details are found in [95].

2.3.2 Shielding

Muons are not the only source of backgrounds that can be harmful to a dark matter search. The need for further shielding against such backgrounds has led to the current CRESST design seen in Fig. 2.1. In describing the experimental layout, each possible source of background is considered. The resulting shielding used to defend against each background source is then described.

2.3.2.1 Neutrons

With the muon flux reduced, the largest source of neutrons is then from (α, n) interactions off light elements and the spontaneous fission of U-238 [96]. CRESST is located in Hall A of LNGS. The rock surrounding Hall A has a relatively high concentration of natural radioactive isotopes Th-232 and U-238 [97]. However, the neutron flux is expected to mostly originate from spontaneous fission and (α, n) interactions in the concrete that surrounds the LNGS site [96]. The neutron flux is ~ 20 neutrons $\text{m}^{-2}\text{h}^{-1}$ above 1 keV in LNGS, from both simulation and experiment [96,98]. To shield CRESST from these energetic neutrons, a polyethylene neutron shield is installed around the experimental volume. This shield acts as a moderator, thermalising the fast neutrons down to thermal energies $\mathcal{O}(25 \text{ meV})$. Neutrons at these energies are no longer able to trigger CRESST detectors, drastically reducing the effective neutron count rate. This shielding is highly effective at removing neutrons from the concrete surrounding the experimental set up, as seen from simulation results in Fig. 2.15.

However, even with the neutron shield installed, neutrons can still be seen in the 10-40 keV approximate signal region of CRESST detectors, as seen in Fig. 2.15. These signals can be from fast neutrons produced by muons in interactions both inside and outside of the neutron shielding. Neutrons can also originate from radioactive decays *within* the experimental volume, such as from spontaneous

fission inside the lead shield [100]. The integrated neutron count rate expected in CRESST detectors between 10 and 40 keV can be taken from the simulations in Fig. 2.15 (b, c, d and e combined), and is approximately 0.01 counts per kg-day. Some of these neutrons can be cut by the coincidence cut in section 4.2.12.




Fig. 2.15 Simulations of the expected neutron backgrounds in CRESST detectors from [99]. a shows the expected neutron count rate in a single CaWO_4 detector without neutron shielding originating from the concrete in the surrounding halls, b with 50cm of neutron shielding. c shows the expected flux from spontaneous fission of U-238 in the lead shield, d the flux from muon induced neutrons originating outside the experimental volume, e muon induced neutrons in the experimental set up itself.

2.3.2.2 Radon

Rn-222 is part of the U-238 decay chain. Decays from U-238 chain contaminants in the rock surrounding LNGS halls can then lead to radon gas entering the hall itself. This can be through diffusion, or from first dissolving in the ground water. Measurements of radon contamination in Hall A vary from 20 Bq m^{-3} [101] to about 200 Bq m^{-3} [102]. Were radon gas allowed near to detectors, it would lead to a variety of dangerous backgrounds; a full list can be found using [103]. To prevent radon from coming near the detectors, CRESST is housed in a radon box, seen in Fig. 2.1. The radon box is continuously flushed with pure nitrogen gas, preventing radon from coming close to the experimental volume.

2.3.2.3 Gamma and beta background

CRESST is located in Hall A of LNGS. As mentioned, the rock in Hall A, when compared to Hall B and C, is surrounded by relatively high background rock (Table 2.3). The total γ flux in Hall A, below 3 MeV, was measured recently at $0.3 \text{ cm}^{-2} \text{ s}^{-1}$ [104], lower than previous estimates of $\sim 1 \text{ cm}^{-2} \text{ s}^{-1}$ [105]. The first, outer shield against this external radiation is 20cm of low background lead. Lead is a high density material with high atomic weight nuclei, and thus efficiently stops external γ rays.

| Rock location | Contamination (Bq/kg) | | |
|---------------|-----------------------|------------------|-----------------|
| | ^{238}U | ^{232}U | ^{40}K |
| Hall A | 116±12 | 12±0.4 | 307±8 |
| Hall B | 7.1±1.6 | 0.34±0.11 | 7±1.7 |
| Hall C | 11±2.3 | 0.37±0.13 | 4±1.9 |

Table 2.3 Contamination of the rocks surrounding LNGS, from [97]. Hall A is surrounded by relatively radioactive rock when compared with Halls B and C.

However, lead itself is not radio-clean, containing the unstable isotope Pb-210. This isotope is a natural contaminant, originating as a daughter from the U-238 decay chain. The β decay of Pb-210, and the subsequent decay of its daughters, can also lead to low energy backgrounds in the detectors. While reduced background lead is used in this shielding, an additional 14cm copper radiation shield is also installed. Copper can be produced with high radio-purity, and is then used closer to the detectors. Radio-clean copper is used for the inner radiation shield, the temperature shields around the detectors, and the detector carousel on which detectors are mounted.

It is worth noting that the dilution refrigerator is not radio-clean. The S1000 $^3\text{He}/^4\text{He}$ dilution refrigerator used to keep detectors cold is commercially available, and not designed for low background applications. Thus, the refrigerator itself is kept physically separated from the detectors (Fig. 2.1). In

between are layers of extremely low background lead, in an effort to efficiently stop backgrounds originating from the refrigerator.

With this shielding, the most significant γ and β source is from within and very close to the target crystals themselves [83]. Energy spectrums of the two crystals in operation in the 2007 CRESST-II commissioning run are shown in Fig. 2.16.




Fig. 2.16 Detector activity of modules active in CRESST-II commissioning run, from [83]. Fig. 2.16a represents 21.2 kg-days of data for module "Verena/SOS21", Fig. 2.16b for module "Zora/SOS21", taken over 22.8 kg-days. With gratitude to R. Lang for providing the raw data points of the shown data.

The expected background rates that could arise in the signal acceptance region of the commissioning run can be estimated by using the information in Fig. 2.16 combined with detector resolution functions of commissioning run detectors in

Table 7.1. Under the assumptions that: the event rate in Fig. 2.16 is completely dominated by γ/β interactions at all energies; that γ/β band leakage may be modelled by a Gaussian function with parameters given in Table 7.1; and that the signal acceptance region is the same as that given in the CRESST-II commissioning run paper [106]; estimates of γ/β band leakage can be made. These estimates are ~ 0.1 per kg-day for module Verena/SOS21 and ~ 2 per kg-day for module Zora/SOS23.

However, it should be noted that, unlike the dark matter limits derived in chapter 7, these background leakage estimates are not stable to errors in light detector resolution estimation. If γ/β band resolution was systematically higher than that shown in Table 7.1 by 10%, but acceptance region definitions remained the same, γ/β band leakage estimates using this method would be a factor ~ 2 higher for both modules. Comparing the resolution functions extracted from the acceptance regions shown in the commissioning run paper [106] (Table 7.1) and those in an independent analysis in [83] (Appendix E) highlights this concern when estimating γ/β band leakage. Position dependence is also not addressed with this simple calculation, a phenomena further discussed in chapter 6. Other problems with estimating γ/β band leakage in this way are also discussed in section 7.5.1.

2.3.2.4 Electronic and vibrational noise

One must also be cautious of the more mundane external influences of electromagnetic interference (EMI) and vibrations. To guard against EMI, the experiment, including all shielding, is housed inside a faraday cage. For vibrations, the entire cryostat itself is also placed on air dampers. This helps remove the effects of mechanical vibrations induced by, for example, earthquakes and passing traffic.

2.4 WIMP searches

A key tenet of a dark matter search, or indeed any experiment aimed at discovering new physics, is the *blind analysis*. A blind analysis is one where the experimentalist defines all datasets, cuts, and analysis methods without access to the final data set from which new physics could be established. The purpose of this blinding is to prevent expectation bias. If allowed to modify an analysis while simultaneously allowed access to the final results, the experimentalist can, consciously or unconsciously, change the final results towards an expected result by tuning analysis parameters. Several examples of this bias are suggested in [107]. The systematic bias introduced by allowing non-blind selection of analysis parameters is unquantifiable, and is thus avoided wherever possible.

Cuts must be used to clean raw data for final analysis. This is to discard spurious non-particle and non-WIMP particle interactions that can trigger recordings during a WIMP search. For CRESST, a blind analysis is pursued by defining cuts using a *training* data set – a short period of background data on which cuts may be trained, but no dark matter information can be extracted. In this work, cuts in chapter 4 were originally designed to analyse background data in the most recent CRESST run using the first 13 days of background data as a training set¹, and this data would then not have been used for a full dark matter analysis. These cuts were then extended to allow the detector response analyses of chapters 5 and 6. However, while these cuts were originally designed for a dark matter analysis, background data from the most recent CRESST run (2011) are not used to produce WIMP limits in this thesis. This decision was made because several backgrounds contaminated the expected signal region in the 2011 run, such as lead and alpha recoils, not evident in earlier runs. That these backgrounds exist can be seen in an analysis of this same data by the CRESST collaboration in [53].

¹ One module, K07/David, (09B) was not operational during the first part of Run 32 as the experiment was not yet cold enough for the light detector to enter its superconducting transition. The training set for this module was then selected to be of the same length as the other modules.

To account for these backgrounds required models to be made of expected spectral shape of these backgrounds in the dark matter signal region. This was performed in [53] using *side-band analyses*; the spectral behaviour of backgrounds in areas outside the dark matter signal region was extrapolated to infer the behaviour of those same backgrounds within the dark matter signal region¹. However, such side-band modelling and extrapolation can be subject to several possible systematic errors, many of which are outlined in section 7.5 of this work. Instead, dark matter limits in this work are derived with recourse to information from the publication [106] concerning an earlier (2007) run of CRESST. These limits, described in chapter 7, use the same data set, and thus the same cuts, as originally published in that work.

2.5 Summary

CRESST is a dark matter search using up to 33, 4cm by 4cm CaWO_4 or ZnWO_4 scintillating target crystals, where interactions in the crystal both deposit heat and produce light. CRESST can see interactions of $\mathcal{O}(10 \text{ keV})$ in energy in these target crystals due to the use of SPT thermometers held at cryogenic temperatures. By using scintillating targets, CRESST is able to discriminate between electron recoils and nuclear recoils. This is due to different amounts of light produced by each kind of interaction, relative to the heat of the interaction. By simultaneously measuring the light and heat produced using dedicated phonon and light detectors, background interactions can be distinguished from nuclear recoils. None-the-less, background reduction is desirable, and CRESST's location and design attempts to reduce backgrounds from cosmic ray muons, neutrons, radon, γ and β interactions.

¹ It should be noted that in [53] the magnitude of these background predictions is also affected by information within the signal region. This can be seen in Table 4 of [53] where a different implied WIMP parameters in the signal region (M1 or M2) causes predictions of e.g. neutron background events to change.

3 Oxrop and Analysis

This chapter introduces the software used to analyse CRESST data, Oxrop. Oxrop is a data analysis program used to make the transition from raw data taken by an experiment to finished, calibrated yield plots used in analysis. Recent upgrades to this software are detailed. How this software is used to determine energies in CRESST data below ~ 1 MeV, used later in this work, is then described.

3.1 Overview

CRESST aims to detect rare interactions, a few amongst tens of millions of background events. These background events then provide large data volumes. Software used to analyse background data must be able to reliably cope with such large volumes while at the same time clearly outlining results at each analysis stage. For this purpose, the software package Oxrop has been developed. Oxrop is written in C++ and is based on the data analysis framework ROOT. Oxrop is an interactive analysis program, and has been used amongst the CRESST collaboration since 2006. It contains all the tools necessary to generate CRESST yield plots ready for a dark-matter search.

Oxrop follows the principles of Object Oriented design and allows data produced by several different experiments to be read, processed and stored for analysis. Whilst initially developed for CRESST, Oxrop 5 can now handle data from Institute of Nuclear physics of Lyon (IPNL), EDELWEISS, the data acquisition system used within the Oxford group (OxDAQ). Oxrop may also be used in the analysis of EURECA, a next-generation dark matter search experiment. The purpose of this chapter is to outline upgrades to Oxrop performed by the author, and the use of Oxrop analysis of the low energy (< 1 MeV) region of CRESST data.

3.2 Software basis

Oxrop consists of nearly 100,000 lines of code. Whilst it would have been entirely possible to have designed Oxrop without reference to classes outside the C++ standard library, it would have been inefficient to re-implement functions that are not particular to the analyses it will be used for. One of the advantages of Object oriented programming is the ability to reuse code through the mechanism of inheritance, the transparent specification of generalized classes designed and written by external sources. Within particle physics, the ROOT package is both extensive and commonly used across multiple experiments, and provides a strong basis for any more specific analysis software.

3.2.1 ROOT

ROOT [108] is a data analysis framework developed at CERN. ROOT was designed with the aim of providing a data analysis structure suitable for use with the large amount of experimental data that was to be collected at the Large Hadron Collider. While CRESST is not related to the LHC, many functional features of ROOT, such as file operations and graphical interfacing, remain relevant. ROOT also incorporates a C++ interpreter, CINT. The interpreter allows for bespoke analyses through macro processing without modification of Oxrop's source code. This feature provides for simpler software maintenance and development. The internal workings of Oxrop may change significantly between versions, but the functionality presented to the user for use in macros remains almost unchanged.

3.2.2 Oxrop functionality

The majority of classes in Oxrop both inherit from and utilize ROOT classes in their functions. As such, there is consistency in data management operations such as file access, object naming and list or array operations. Oxrop has been designed for use with CRESST with its own specific analysis requirements, and appropriate tools have been developed. A brief description of the most important of these tools is given here.

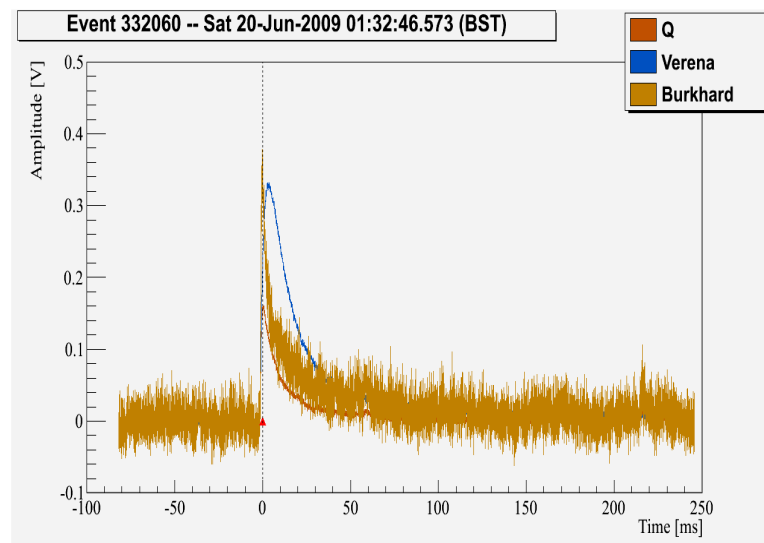


Fig. 3.1 Example of Oxrop's event display. A 59 keV electron recoil-like interaction in detector module 04B during 57-Co calibration. "Verena" is the phonon detector, with "Burkhard" and "Q" the associated light detectors. The red triangle indicates a coincident event in a separate module. The coincident event is a 77 keV β/γ like interaction in nearby module 08M. This display then highlights the possibility that a 136 keV γ has interacted in 08M, producing a Tungsten $K\alpha 1$ X-ray, which has escaped and interacted in 04B.

Oxrop can be used to visualise recorded data through the *event display*. The event display allows records stored on disk to be displayed. A trigger causes all the detectors in a module to be read out, with each detector having a separate record. All detector records associated with this triggering module are then grouped in *events*, and the event display allows these records to be compared simultaneously. Events close to this event in time on other transient recorders can also be seen through the event display.

When loading files, Oxrop examines all records to generate several *parameters*. Oxrop also gathers further parameters that were written with the record by the data acquisition. These parameters are useful for a first analysis of a record, and the following is a list of important parameters that are stored at this point:

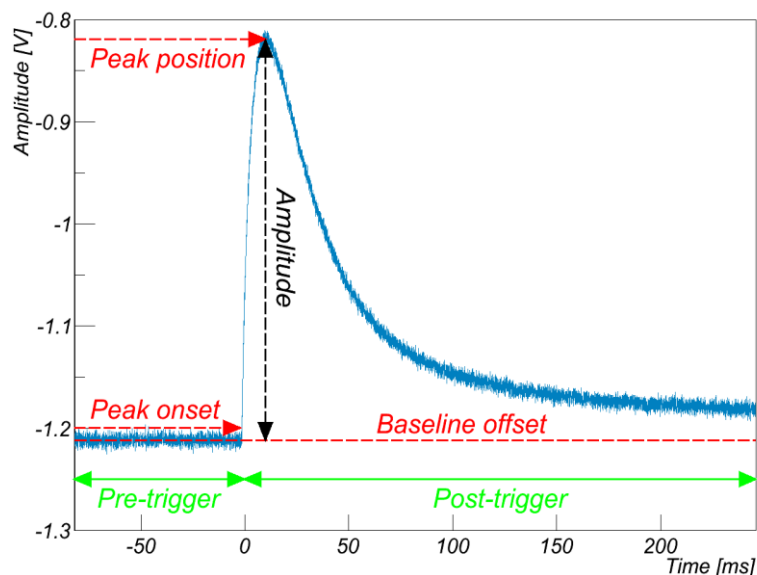


Fig. 3.2 Typical pulse record with some important parameters labelled. Calculated parameters are shown in red. Pre and post trigger regions are defined by the data acquisition unit, and are shown in green. The y-axis shows the raw voltage output as seen by the transient recorder. A template fit parameter (section 3.4) is seen in black.

Calculated parameters:

Pulse height: The maximum change in voltage between points before and after a trigger. This can be used as a first estimate of the energy of an interaction to within a scaling factor (section 2.2.3).

Time stamp: The time in microseconds of the beginning of the post-trigger region. This is calculated from the 10 MHz DAQ clock and the computer timestamps of the recording session.

Baseline voltage: The average voltage in the first 95% of the pre-trigger region. This identifies the voltage level before an event occurred.

Peak position: The time from the beginning of a record to the maximum voltage. A 2ms moving average is used to find this peak position.

Peak onset: This is the time from the beginning of a record that the pulse began. This is found moving back from the peak position to the point where the voltage returns to the baseline voltage, plus the noise on the baseline.

Minimum / Maximum voltage: The lowest and highest voltages seen in a record.

Right – left of baseline: The average voltage in the last 2ms, minus the average voltage in the first 2ms of a record. This is the voltage offset across a record.

Delta voltage: The maximum drop in voltage between adjacent and next-to-adjacent voltage readings. This drop is divided by the RMS of the pre-trigger region, forming a measure of maximum voltage drop relative to noise.

DAQ parameters:

Trigger delay: The trigger time, relative to the first trigger on a transient recorder. Detectors in a module trigger independently, thus trigger delay can be non-zero for at least one detector in a module. Further, there is often more than one module on a single transient recorder, which can cause relative trigger delays between modules.

Delay firing T.P. (control pulses): The delay in firing a control pulse within a control pulse recording. Control pulse firing can be delayed by pile-up, where a time-coincident particle interaction automatically delays the actual sending of control pulse heating. The delay between the time a control pulse recording started and heating pulse is actually sent is the firing delay. If this firing delay is large, then a control pulse might not have been properly recorded.

Oxrop can visualise parameters through histograms and graphs, incorporating the ROOT histogram and graph functionality through a GUI. Cuts on parameters can also be defined through the GUI, either visually or numerically.




Fig. 3.3a example of histogram of pulse height for a 122.06 keV Co-57 Calibration of the module 10B, Rita/Steven. 3.3b Shows an example graph of calibrated energies for the same data.

Pulse height is approximately linearly related to the energy deposited in a detector. The energy deposited in the light detector by γ interactions is also approximately linearly related to the energy deposited in the phonon detector, seen in Fig. 3.3b (discussed in chapter 6).

It can be necessary to mathematically transform one parameter (or more) into other, new parameter. This then allows amplitude energy relations, or light / phonon energy relations, to be drawn up. Incorporated into the GUI is a mathematical operation (MOP) function, allowing relationships to be constructed between existing parameters to create new parameter values (Fig. 3.4).

As can be seen in Fig. 3.1, noise is evident in recorded data. By examining the frequency spectrum of this noise, sources of electromagnetic disturbances, such as ground loops, or common noise modes across multiple detectors, can be identified. For this purpose, Oxrop has a Fourier transformation mechanism that

allows time-domain data output from the DAQ to be transformed into the frequency domain to determine what noise components are present in the data. Additional packages allow the generation of estimated noise simulations from a sample of real noise data [109]. These simulated “empty baselines”, so called as they lack any pulse data, can then be used to study noise effects on resolution.

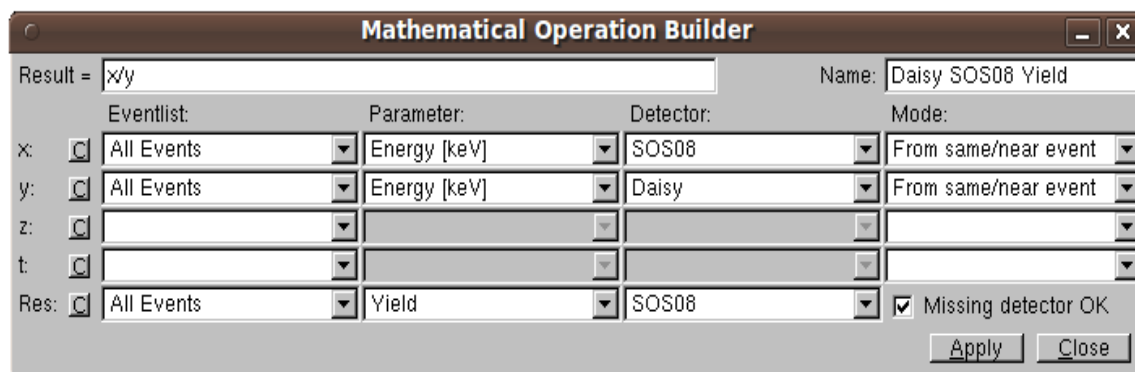


Fig. 3.4 Example of a mathematical operation. The “light yield” of the module 09M is defined by the light detector energy divided by the phonon energy.




Fig. 3.5a Example of a noise only recording for the light detector “Josef” and Fig. 3.5b corresponding Fourier transform from noise only recorded calibration data of the most recent CRESST Run.

All particle pulses in a particular detector are expected to show the same generic pulse shape in recorded data. This fact can be used to make more accurate estimates of pulse amplitude than a simple pulse height found from first examination of a record. Two methods have been incorporated into Oxrop to fit these generic pulse shapes to recorded data. The first is a template fit. An

expected pulse profile is generated by averaging over hundreds of standard interactions. This expected shape is then fit to recorded samples, reducing noise effects on the estimation of the amplitude of an interaction, when compared to the pulse height. This template fit is used in CRESST low energy analyses, up to ~ 1 MeV, further described in section 3.4. The second fit mechanism utilises a more generalised multi-parameter fit. Eqn. 2.6 is used to model recorded pulse shape. Transition non-linearity effects are then modelled by up to a 15th order polynomial. By accounting for transition non-linearity, energies of up to ~ 10 MeV can be reconstructed. Further details are found in [110].

A further important process in Oxrop relates to how the test pulses (section 2.2.4) are used to correct for non-linear effects in the amplitude – energy relation, given that the transition between normally conducting and superconducting transitions in SPT thermometers is not fully linear in resistance against temperature. This process is also outlined in section 3.4.2.

3.3 Structural development

One of the key issues within any package aimed at analysing CRESST data is how to process, store and represent data from the different DAQ systems which are obtained over periods longer than a year. These different data types are from cryogenic particle data from the phonon and light detectors, muon veto data, control pulse data, and environmental monitoring data, with the total number of events recorded of the order hundreds of millions. The importance of considering the practical limitations of the computers used to analyse this data cannot be underestimated. Without due weight to memory and processing optimization, any analysis program will fail due to the weight of data it is required to process.

3.3.1 Oxrop 4

Oxrop 4 was in use without significant issue in the CRESST commissioning run (“Run 30”) in 2007. In that run, 3 data modules (6 detectors) recorded background data for a little over 7 months. “Run 31” (2008) and the most recent run, “Run 32” (2009-2011) saw significant jumps in the quantity of information required for processing. This was both in terms of the number of detectors and the length of recording.

In Run 32, 28 cryogenic (phonon or light) detectors took data, and the experimental run lasted nearly 2 years. One detector in particular caused additional data storage issues. The detector, “K14”, had significant U-238 decay chain contamination, providing a total particle count rate of 0.52 Hz, responsible for over 30% of all particle event triggers in the first year’s data set. Additionally, Run 32 saw time-varying noise (see section 4.2.2). The magnitude of baseline noise could at times exceed trigger threshold, causing rapid triggering of the detector readout. Combined, these factors meant that the first year of Run 32 data represented over 80 million cryogenic detector records, an increase of a factor of 20 over the requirements of Run 30 (approximately 4 million cryogenic records). Raw pulse data have always remained on disk unless required. However, parameters associated with both events and pulses in Oxrop 4 were located entirely in memory, leading to a significant increase in memory requirements over Run 30. Run 30 required approximately 6.5 gigabytes of memory for the entire run; the large increase in data volume for Run 32 meant that it became impractical for Oxrop to be used for analysing the entirety of Run 32 in a single Oxrop session.

A further issue with Oxrop 4, with regards to extending its use in data from other current experiments, such as EDELWEISS, and future experiments, such as

EURECA, was that all event types were not stored in the same fashion. Oxrop 4 stored events originating from four distinct sources:

- (1) Triggers in the cryogenic phonon and light detectors.
- (2) Control events that stabilise cryogenic detectors
- (3) Muon veto recordings.
- (4) Monitor data. Environmental variables such as temperature or air pressure are periodically recorded in CRESST.

Information from these sources was stored through different methods in Oxrop. Before describing Oxrop's structure further, it is useful to first outline several key C++ classes in Oxrop as they relate to data storage, particularly for recorded data from cryogenic (phonon and light) detectors.

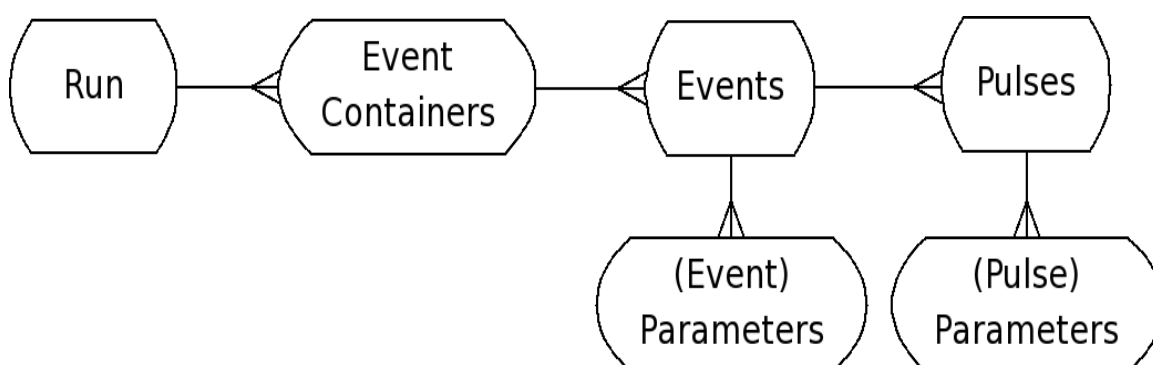


Fig. 3.6 Relational diagram showing how cryogenic event data is stored in Oxrop. UML diagrams, indicating the implementation of these relations, are shown in Fig. A.1 and Fig. A.2 in Appendix A.

Pulse: Contains information associated with a single record (8192 voltage readings each separated by 40 μ s in a single detector) for a single trigger. Pulses contain a detector identifier and the position of the recorded data within a file on the disk. In Oxrop 4, parameter values associated with that single record (e.g. pulse height) are stored in the pulse.

Event: A collection of pulses associated with a single trigger in a single transient recorder. In Oxrop 4, parameter values associated with the trigger itself are also stored within an event, such as the test pulse voltage (zero for particle events).

Event Container: A grouping of all the events within a single readout session for a specific type of event. This could be for example all cryogenic detector events or all muon veto events. Descriptions of the recording data format, and the start and stop time of the readout session are also kept in this object.

Run: A global object containing references to every object within the analysis session. This includes for example event containers, event lists and histograms.

The information above describes in general terms how data from cryogenic data is represented in memory. However, as mentioned, parameter values associated with each of the four classes of events (cryogenic, control, muon, monitor) are stored differently.




Fig. 3.7 Diagram describing cryogenic event parameter storage. An event with two pulses is shown. Events contain an array allowing sixteen parameters. An event also contains a pointer to a pulse. The pulse then has a similar array allowing up to sixty four parameters, and can point to another pulse.

Both pulse and event parameter values for cryogenic (phonon and light) detectors are stored in fixed size arrays within the “Event” and “Pulse” objects respectively. One or more pulses can be associated with a single cryogenic event, representing the one or more detectors recorded in a single trigger on a single transient recorder. Pulses are stored in a singly-linked list within the event.

However, control event information was stored with a different method. Each control pulse event contained an array of structs, with a struct existing for each detector associated with the event. Each struct then contained a fixed array of parameter information associated with each detector in the event.

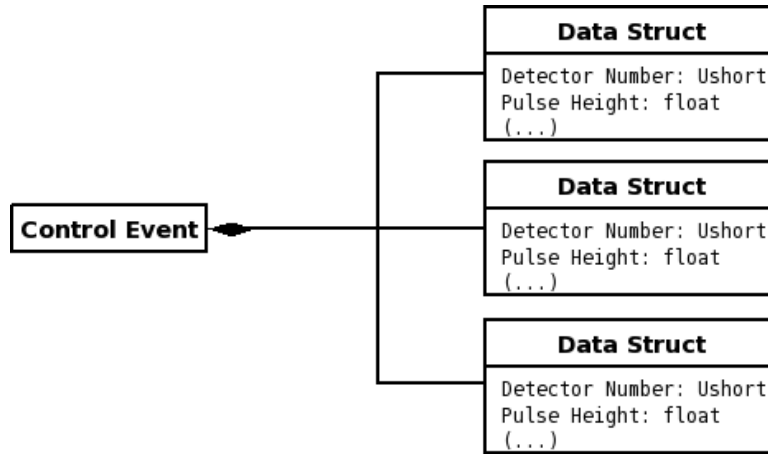


Fig. 3.8 Diagram of control pulse parameter storage. A control pulse containing information for three detectors is shown. No parameter data is held in the control event class itself, except for the time of the event. The control event points to an array of structs. Each struct contains five parameters (two shown), that store detector specific information.

Veto events were again dealt with differently. The CRESST muon veto consists of 20 scintillating plastic panels, each with a single PMT. If several panels trigger, then then the pulse heights on all the triggering panels are recorded (with an additional “sum”). This leads to events with a varying data volume. To store this data from muon veto system triggers, a different class-object would be instantiated depending on the number of panels triggered, labelled “2”, “3”, “5”, “10” and “20”. Panel data was held in fixed sized arrays within these objects, and memory size depended upon the instanced class.




Fig. 3.9 Diagram describing muon veto event parameter storage in muon events. Different classes are instances depending on the number of panels that triggered in a muon event. Two events are shown, one requiring space for two panels of data, one requiring space for five panels of data.

Lastly environmental monitor storage is also stored. Each monitor event consisted of up to forty parameters describing the environment of the experiment. Storage is similar to cryogenic event information, except that there are no “pulses” associated with monitor events.

This variety of data storing methods made it difficult to extend Oxrop for use in other experiments: simply, which method should be used to represent data from a new experiment? An additional problem with this structure is the method by which parameters were accessed. While the method of storing parameters was different for each event type, ultimately a wrapper class is used so that users accessed parameters in the same way for all kinds of events. Necessarily, this wrapper presented the user with the ability to add and remove user defined parameters for all event types. Structurally however, only cryogenic events could actually add such additional parameters.

These issues combined meant that Oxrop could no longer be used for the main purpose its users wished to use it - to process and analyse data for an entire run interactively. By far the largest problem was memory usage. Attempts to analyse Run 32 required many tens of gigabytes of memory to analyse even small sections, with the majority of memory being used to store parameter data. Extensive structural changes were then required to reduce memory, with three further primary aims. First, interactive analysis capability needed to be maintained, as this is the primary way Oxrop is used within the collaboration. Second, data storage should be consistent across all event types. Lastly, while it would be necessary to break backward compatibility, existing analysis macros should require minimal changes for use between software versions.

3.3.2 Oxrop 5

Whilst originally designed under the Object Oriented (OO) model, Oxrop's purpose has much in common with that of Relational Database Management Systems (RDBMS). Under such systems, elements which would, under a purely OO design, be considered data members of a specific class, are not necessarily stored within the class itself. Such elements are stored in separate tables, with interfaces provided to access data within that table. To allow Oxrop to handle the

increase in data volume required a translation of the existing structure to one more resembling those of relational databases. This required a complete transformation of the existing structure, such that the “event” and “pulse” objects no longer stored the parametric data. Events and pulses instead now provide an interface to tables that hold this data elsewhere. The tables themselves are located within the event containers. Full UML diagrams of the most important features changed between Oxrop 4 and 5 can be found in Appendix A.

ROOT Trees were considered extensively for storing data. Trees are designed to hold the large amount of data produced by CERN experiments. Superficially, “events” holding many “tracks” bear a strong similarity to CRESST, with events containing several pulses. Trees have the added bonus of compression, fast access times and extensive support in the Particle Physics community.

Oxrop however is an interactive analysis program which *can* allow variables to be updated individually, such as when a mathematical operation or a standard pulse fit is carried out on small event lists spanning many event containers. Modifying a parameter value within a Tree requires the saving of (at least) the entirety of the parameter store in that Tree, the smallest practical size in Oxrop being an event container. To both use trees and maintain interactive speed, all capability of updating individual parameters would have to be removed. This would have both been time consuming and necessitated a significant breaking of backward compatibility. Instead, a disk data management system was designed specifically for Oxrop's purposes.

No parameter is now stored directly within events or pulses. With the exception of the microsecond event time, used to identify an event's parent container, parameter data is now stored entirely in the event containers. The new method of storing data can essentially be viewed as parameter data being stored in tables. The number of rows is fixed by the number of events (for event parameters) or

pulses (for pulses parameters) that were written to disk in a single recording session. Each parameter then defines a column in this table. The number of entries in each column is then defined by the number of events or pulses in a recording session, appropriately. When asked for a parameter value, each event, containing its unique event time, interrogates the global run object to find its event container, and is then able to return the relevant value from the container, as depicted in Fig. 3.10.

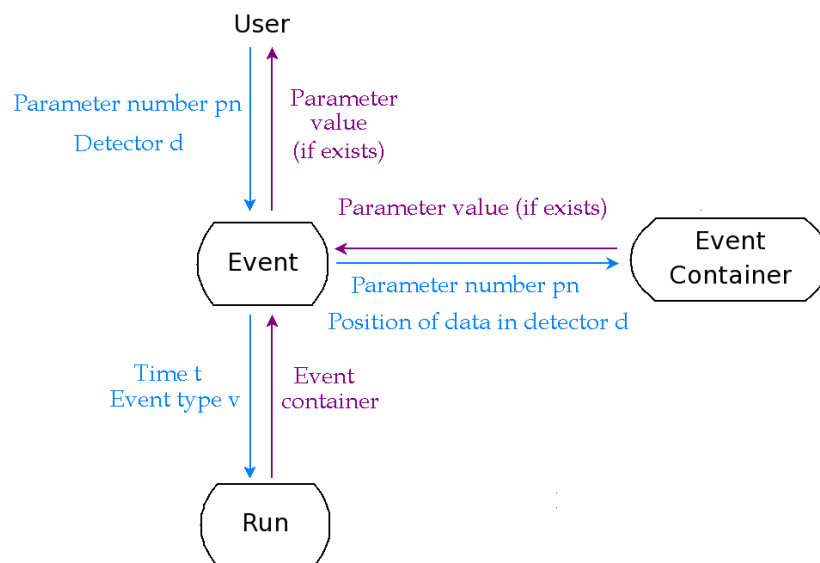


Fig. 3.10 Simplified data flow diagram for parameter access in Oxrop 5. Blue indicates information sent in a request, purple information received from a request. A user requests the value of a parameter numbered pn , in a detector numbered d from an event. The event finds its corresponding container from the global "Run" object, using its event time and type. The event then interrogates the event container to see if it contains information about detector number d with parameter number pn , returning this value to the user.

Events have now been split into two broad categories: events which have sub-detectors, and events which don't. Sub-detector data is defined as information associated with an event which can have a multiplicity larger than 1. This includes pulse parameters of cryogenic events, panel data of Muon Veto events¹ and detector parameters for control pulses. All event types now store data consistently. Both in memory and to the user, data is accessed identically.

¹ Or the PMT signals (pulse height and timing) for the EDELWEISS muon veto, two for each panel.

Events still exist within memory, containing their event time and a unique identifier. These values are used to find their container and position within that container respectively. Parameter data associated with that event are now stored in the event container. Sub-detector data are roughly equivalent to “Pulse” data in Oxrop 4. Events containing sub-detectors also contain the position in the event container where sub-detector parameter data starts, and how many sub-detectors are in that event. Pulses no-longer exist as stand-alone entities. Parameter data previously associated with pulses is now accessed through the parent event. However, for backward compatibility with pre-existing macros, a pulse-like object is available to provide an interface to the parameter table.

Previously, it was required that the entire pulse should exist in memory. The major advantage of storing parameter data in a table-like format is that this table can be kept on disk. In Oxrop 5, this table can be written to disk, with a small section (or *page*) kept in memory. The optimum size of the page was determined considering two competing factors. The first is random access time, such as for event lists with few events that span many event containers, which favours small pages. The second is sequential access time, for initial processing or fitting to all data, which favours large pages.

Using a paging system also means disk storage is optional. When analysing small sections of a run, or where one wishes an analysis to be portable, it is often more convenient to store parameters entirely in memory instead. If a user wishes to keep parameters in memory, this option is provided. The page size is simply expanded to include all data in an event container.

How information related to DAQ, Heater, Bias, Analysis and SQUID settings for each detector is stored via Infos has also been reworked. Previously, only the cryogenic event container could contain information specific to detectors. Muon veto event containers could contain no information about the veto-panels

themselves. Control-event containers referenced the corresponding cryo-event container for details about their own detectors, breaking the concept of data encapsulation. In Oxrop 5, detector specific information is grouped together in the corresponding event container. This change is particularly relevant for the muon veto. In CRESST, each of the 20 muon veto panels contains a single recording PMT. However, in EDELWEISS, each panel gives two signals, one for each end of a single panel. Without a method for storing information about the detector positions within panels, it was not possible to group detectors by panel. In the new structure this is possible, a point of strong importance in future use of Oxrop for EURECA.

These structural changes have several advantages. As seen in Table 3.1, storing parameters on disk has cut memory requirements dramatically. The speed of generating histograms has worsened slightly, as, despite paging, accessing data on disk is slower than accessing data stored in memory. When entirely located in memory, Oxrop 5 takes up less space, operates faster in saving and loading, and is marginally faster at data analysis actions such as drawing histograms.

In conclusion, Oxrop has undergone major changes, aimed at improving its performance and structure. The amount of memory required for operation has decreased drastically, with relatively small impact on operational speed. Parameter storage has now been sufficiently generalised, allowing new detector types to be added following a consistent protocol. When the data structure for EURECA is decided, it will doubtless present new challenges, however, Oxrop is now much better placed to meet such requirements.

| | Oxrop 4.9 | Oxrop 5 | Oxrop 5 (Disk Mode) |
|---|------------|-----------|------------------------|
| Total Memory Loading files (no GUI) | 11145.1 MB | 4978.6 MB | 2289.6 MB |
| Total Memory Interactive Use (with GUI) | 11239.4 MB | 3548.2 MB | 856.3 MB |
| Save Time | 290s±90s | 57s±21s | 27±3s |
| Load Time | 330s±60s | 32s±16s | 17s±3s |
| Histogram for “All Events” of “Time Since Start [h]” | 5.2s±0.1s | 3.0s±0.1s | 7.0s±0.2s |

Table 3.1 Benchmarks for software versions of Oxrop for a section at the end of Run 32, between 19:22 on the 26th of November, 2010 and 07:44 1st of March, 2011, 84 days and 11 hours of recording time. Section was analysed using the 64-bit, Oxford Particle Physics linux cluster “pplxint6”. The reduction in memory between loading files without a GUI and opening a saved session with a GUI is seen in Oxrop 5, is due to the extra memory space required to load files.

3.4 Analysis of the sub-1 MeV CRESST data

The first aim of most analyses of CRESST data is to determine the energy of every interaction observed in both the light and phonon detectors. This process will ultimately allow for discrimination between interactions of different types, such as between electronic and nuclear recoils. This process also allows for a characterization of low energy radioactive backgrounds in the data set. The following details energy reconstruction applicable below ~ 1 MeV, where detector response is well within the range of the heating test pulses (below)¹.

3.4.1 Energy estimates

First, consider a detector response which is within the linear region of the superconducting phase transition. In this region, measured amplitude scales approximately linearly with energy. Thus, the first quantity requiring measurement is the pulse amplitude.

¹ The limit of the test pulses varies between detectors, but rarely extends beyond 2 MeV.




Fig. 3.11 Diagram of different estimates of pulse height pulse in the light detector “X” at 28.4 keV, with a template fit measured amplitude (black solid line). Amplitude is over estimated by a simple maximum voltage search (red dashed line), yet underestimated by a 2ms moving average (blue dashed line). The template fit amplitude is seen with the black dashed line.

A first estimate of this amplitude can be taken from the magnitude of difference between the maximum voltage seen within a record, and the average of the voltage that preceded the response onset (the “Baseline Voltage”). However, as seen in Fig. 3.11, this simple estimate can be skewed by noise. The absolute maximum measured voltage in the pulse tends to increase with increasing noise, introducing a systematic, positive shift in this amplitude estimate. To reduce this effect, a moving average can be used, also seen in Fig. 3.11. This moving average pulse height is calculated by Oxrop from a first examination of all records, and can be used as the first line of an analysis.

However, a more accurate amplitude estimate makes use of the fact that all particle interactions have the same pulse shape in a given detector (section 2.2.3). An amplitude estimate can then be found by fitting the expected response to the recorded response. Fitting over the entire pulse significantly reduces the effect of noise seen in a record. Fits can then be defined by three, free parameters - time-shift, amplitude and baseline voltage.




Fig. 3.12a Shows a typical 122.06 keV response in module 10B, with phonon detector “Rita” (Blue) and the light detector “Steven” (Green). Fig. 3.12b shows the normalized standard pulses generated by averaging over 1,837 similar responses. By averaging over many pulses, a template response shape is derived with negligible noise with respect to a recorded pulse.

To perform this fit, the expected pulse shape must be obtained. This is found by averaging the detector response to many interactions of the same energy. Pulses of the same energy are selected via a cut on pulse height. Ideally, one would use pulses at energies near the dark-matter signal region (below 40 keV) for calibration and standard pulse building. However, photons at these energies from external sources do not penetrate the copper shielding of the cryostat between the target crystal and source. Thus it is necessary to use higher energy interactions for a calibration. Standard pulses are then first built from the 122.06 keV γ line of a Co-57 source¹. A typical standard pulse, derived from these interactions, is shown in 3.12b. After the standard pulse is built, it is fit to the list of events that were used to generate it. Contaminants with, for example, a poor goodness of fit can then be removed, and the standard pulse is finally rebuilt. It should be noted that it is only necessary for $\sim\mathcal{O}(10)$ detector responses to be used to generate a standard pulse [83]. However, many thousands are generally available. This standard pulse is then fit to all observed pulses within a detector.

¹ In practice, a further iteration using 46.5 keV γ Pb-210 intrinsic contaminant events is used for subsequent chapters. This is to ensure pulses are built from the linear region of transition response – see section 3.4.2.

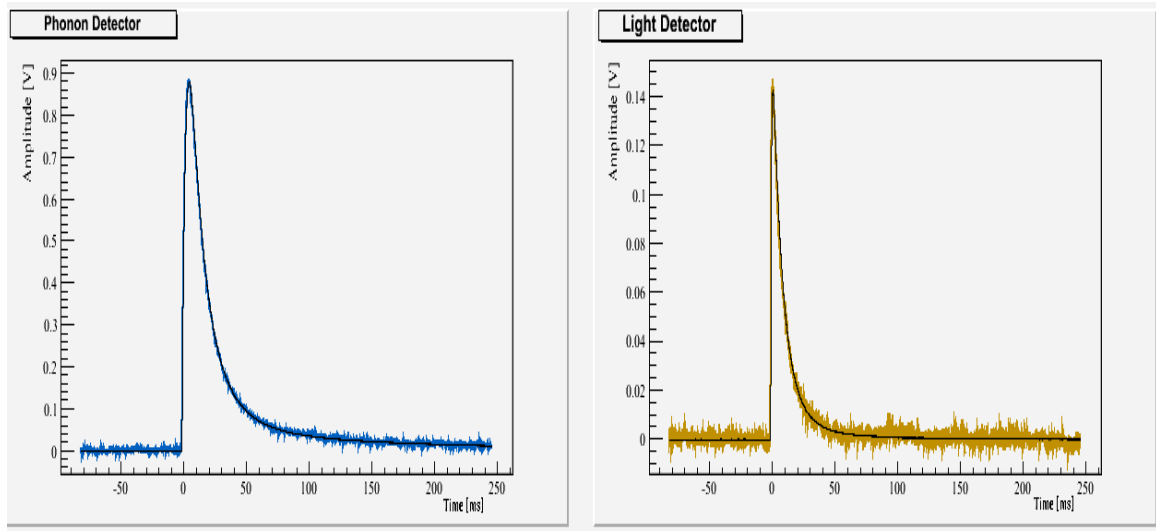


Fig. 3.13 A correlated fit for a 122 keV gamma interaction in module 06T ("Sabine"/"Josef"). A common time-shift is used to minimize the effects of noise on measurements of pulse onset and amplitude.

Every detector in a module has a standard pulse generated from the same group of calibration interactions. These standard pulses are then fit to recorded events simultaneously between light and phonon detector pairs through a *correlated* standard pulse fit. When fitting to an individual detector, a standard pulse is fit to an observed record with three, free parameters:

Template fit parameters:

Amplitude: The height of a template pulse fit to a recorded pulse. This value is approximately linearly related to the energy.

Baseline offset: The average voltage across a record after subtracting the fit to a detector's response to an interaction.

Shift: The position in time of a pulse's onset, relative the standard pulse onset.

Residual goodness-of-fit or *RMS* is also stored. In a correlated fit, the time shift is a common parameter between detectors. An example is shown in Fig. 3.13. Correlated fits are conducted to allow accurate reconstruction of the energy for

events with negligible amplitude in one of the detectors. This would be the case for WIMP signals, where the signal in the light detector is expected to be small compared to that of the phonon detector. In such cases with individual fits, the onset of the interaction in the detector with a small amplitude response can be misjudged due to noise. By correlating the fit, the large amplitude phonon detector fit aligns the light detector fit. True amplitude is then more accurately recovered. This effect is illustrated in Fig. 3.14.




Fig. 3.14 Reconstructed energies from a correlated fit (red) and uncorrelated fit (blue) to the light detector “Hans” in simulated events for module 06B. Events are simulated with 15.00 keV energy in Maja and 0.586 keV (electron equivalent) energy in Hans, modelling tungsten recoil events at this energy. Without correlation, fitting is strongly affected by noise; pulse maximum is often overestimated simply as the maximum within the noise. With correlation, fits reconstruct light detector energy with a reduced systematic shift from the true mean (dashed black line). A small noise induced difference remains; the mean of the correlated fit (dashed fuchsia line) is $0.14 \pm 0.01 \text{ keV}$ higher than the true mean. Simulations performed using the events simulator from [109].

3.4.2 Non-linear response corrections

A standard pulse fit, as outlined above, is only a valid representation of energy when detector response falls entirely within the linear region of a detector. This typically relates to $\lesssim 200$ keV interactions. For higher energy interactions, detector response exceeds this linear region. There is then a deviation from a linear relationship between energy and measured amplitude.

Transition non-linearity is first evident when looking at the behaviour of the standard pulse fit residual goodness of fit RMS against amplitude. When detector response exceeds specific voltages, labelled “truncation voltages”, goodness-of-fit begins to degrade. Beyond this point, detector response is distorted by the shape of the SPT thermometer phase transition. This effect is seen in Fig. 3.15.




Fig. 3.15 RMS of fit against pulses height for interactions in the phonon detector “Rita”. Above 0.6V, detector response becomes increasingly non-linear, and an untruncated fit (black) becomes a steadily poorer approximation of the true pulse shape. The choice of fit-truncation voltage should be lower than this point to ensure minimal non-linearity is introduced in fitting. In this case, fit-truncation voltage is set at 0.5V (black dashed line). Higher amplitude pulses then see an improvement in fit RMS (red).

To reduce this effect, fits are limited to only include the region where detector response is within the linear region; that is, below the truncation voltage. This is exemplified in Fig. 3.16. Truncation voltages are generally equivalent to the amplitude of particle pulses at ~ 100 keV. Above these points, events enter the non-linear region of detector response.

It is then also necessary to ask if the Co-57 122.06 keV γ interactions used to generate the standard pulse access this non-linear region. This may be checked by fitting the standard pulse to energies up to the calibration energy, and examining the residual root-mean-squared (RMS) of the fit, as in Fig. 3.17.




Fig. 3.16 Truncated fit for “Rita”. Pulse samples above 0.5V are ignored, and only the response in the linear region is taken into account when fitting.




Fig. 3.17a shows the RMS profile of fits with a template constructed of 122 keV Co-57 γ interactions for the phonon detector “VK33”, and Fig. 3.17b shows the same but with a template of 46.5keV Pb-210 γ interactions, with the fit truncated near to 80 keV. By using a standard pulse built from Pb-210 interactions, the energy estimate of the 54.9 keV escape peak, originating from 122 keV γ rays interacting with a tungsten K β 1 X-ray escaping from the crystal, is improved from $52.80 \pm 0.07 \text{ keV}$ to $53.42 \pm 0.06 \text{ keV}$.

In Fig. 3.17, the detector VK33 is shown to have a non-linear RMS profile when using a standard pulse generated from Co-57 122.06 keV γ interactions. Two fit RMS minima are seen. One is at the energy that the template was constructed from, and one at zero energy, where there is negligible pulse amplitude. At intermediate energies the template shape differs from that of the recorded pulse, increasing RMS. This pattern is symptomatic of standard pulses generated from

events at energies beyond a detector's linear region, and causes a systematic shift in energy determination (Fig. 3.17 caption).

Standard pulses must then be built from lower amplitude responses. A useful set of interactions arises from Pb-210 contamination. Pb-210 causes a significant number of events in all CaWO₄ modules, and to a lesser extent in ZnWO₄ modules. Pb-210 decays could be either within the crystal or on the surface. While there may be a discrepancy in observed amplitude between contaminant positions (chapter 6), pulse shape is not expected to be affected. 84% of the time, Pb-210 decays lead to a 46.5 keV γ ray with a 17.0keV end-point β . After calibrating with the 122.06 keV Co-57 source, background events with energies very close in energy to the 46.5 keV γ ray peak can be selected. These events have only a very small β contribution, and are almost monoenergetic. These events can be used to build a standard pulse at lower energy than 122.06 keV Co-57 γ events. Fig. 3.17b shows the RMS profile for fitting a 46.5 keV γ standard pulse to VK33. The systematic shift in energy associated with non-linearity is reduced, though not eliminated. VK33 is an extreme case – most detectors do not show this effect clearly. However, in this work, standard pulses from 46.5 keV Pb-210 events are used for all detectors. This is to minimise non-linear effects when fitting to events up to 911.2 keV in chapter 6. Energy normalization (below) is still performed using Co-57, 122.06 keV γ events.

However, even with the truncated fit, non-linearity may not be fully eliminated. The amplitude-energy relation is then more closely examined to increase accuracy in energy reconstruction. For this purpose, artificial heating pulses are introduced into detectors. Described in section 2.2.4, each thermometer is fitted with a resistive gold wire. This wire acts as a heater that stabilizes thermometers within their superconducting phase transition. Incorporated into the heating circuit is voltage square-rooter. The square rooter exists such that the relationship between applied test pulse voltage V_{Bias} and deposited energy is linear.

The time evolution of applied test pulse voltage is designed such that it mimics particle pulse responses. A range of test pulse voltages is applied, large enough to probe a detector's non-linear regime. As with particle interactions, a standard pulse is constructed of test pulses, which is then fit to all recorded test pulses. This template test pulse is then fit to all recorded test pulses, allowing fit amplitudes to be related to the input voltage, as in Fig. 3.18.

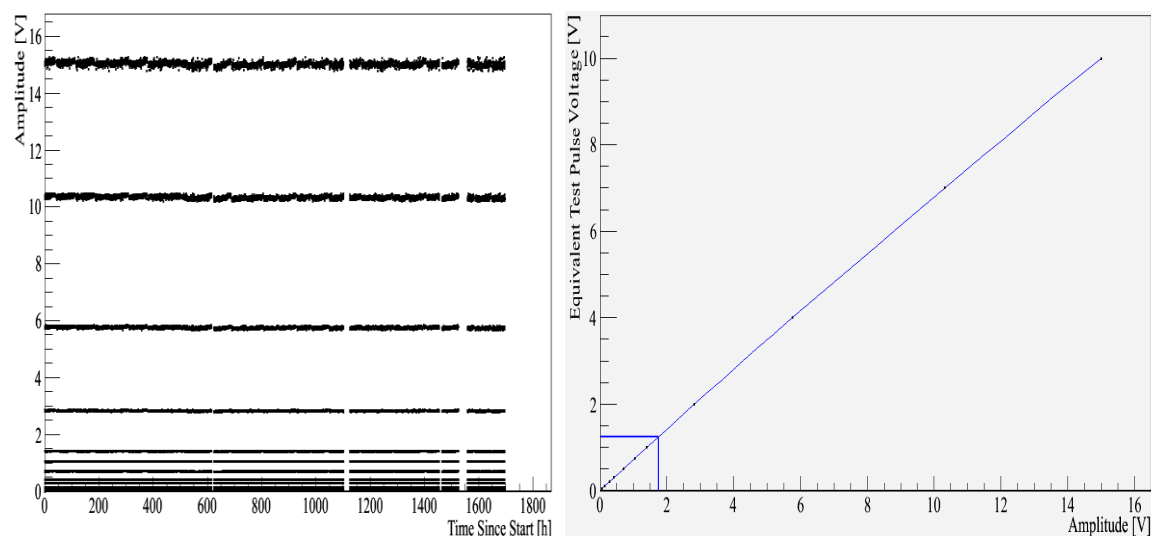


Fig. 3.18a Shows the amplitude of all test pulse voltages in the phonon detector “Sabine” during a section of background data taking. Detector response to test pulse injection is almost constant for each Test Pulse Amplitude. This allows a relation between measured amplitude and test pulse amplitude to be drawn up, as seen in Fig. 3.18b. Though a 5th order polynomial has been used to express this relationship, the use of a truncated fit has meant that the translation is almost linear.

Given this amplitude-voltage relation, particle amplitudes can then be related to particle test pulse voltages. The amplitudes of fits to particle interactions are generally between the fitted amplitudes of the discrete voltage set of test pulses. Thus, the equivalent test pulse voltage is found by interpolating between the test-pulse amplitudes at each test pulse voltage, a process known as “CPE”.

The actual energy input by test-pulses to detectors at each test pulse voltage is unknown. However, energy input is known to scale linearly with voltage input (Eqn. 2.10). Thus, test pulse voltage is related to energy by a scaling factor. This scaling factor is determined using Co-57, 122.06 keV γ calibration events,

independently for all detectors in a module. Energies are then related to γ responses. This is significant for the light detector, as electron and nuclear recoils generate significantly different amounts of light. Light detector energies are then often quoted in “keVee” – electron recoil equivalent energy. On the other hand, the phonon detector is used to estimate total energy deposition. As the phonon signal is not expected to vary greatly with interaction type, phonon readings tend to be quoted in “keV”.




Fig. 3.19 shows the 3V test pulse for Sabine, from the same background data set used in Fig. 3.19. Gaps between data points are when helium and nitrogen coolant is refilled, and the detector carousel warms significantly. Between refills, the carousel slowly cools, and detector heater demand increases, slightly altering detector response over the period. The deviation is smaller than a few percent in all detectors in the low energy region but is none-the-less corrected for by measuring test pulse amplitude at all times. Blue shows a smoothed cubic spline, used to model the variation in measured amplitude.

The relationship between amplitude and energy can vary over time. Variations in cryostat temperature can cause the heating bias voltage, applied to keep detectors at their operating point, to vary. This heating variation can slightly alter the shape of a thermometer's transition. To correct for this effect, the relation between test pulse amplitude and applied voltage is measured periodically throughout recording, following [110]. This allows for any time variation, such as that shown in Fig. 3.19, to be accounted for.

Finally, an underlying assumption in this process is that non-linearity corrections in the amplitude-energy relation are the same for test pulses and particle pulses. This assumption could be incorrect in the case of, for example, any non-linear relationship in cross-talk between the heater and readout electronics, or a highly non-linear pulse shape at large test pulse voltages. However, this assertion of linearity can be justified from reconstructing the energies of known lines (Table 3.2). Almost all detectors show systematic energy errors of at most two percent using Co-57 data alone. The linearization process has then meant that reconstructed energies are very close to the actual interaction energies.

| Line (Energy) | Detector (Module) | | | |
|--|-------------------|-------------------|-------------------|-------------------|
| | VK33 (01T) | Verena (04B) | Maja (06B) | Sabine (06T) |
| Cu, K α (8.0 keV) | (Too weak) | (Too weak) | (Too weak) | 7.90 \pm 0.11 |
| 122.1 Co-57 γ - W,K β 1 (54.9 keV) | 53.42 \pm 0.06 | 55.04 \pm 0.10 | 54.55 \pm 0.18 | 54.13 \pm 0.07 |
| Co-57 γ (136.5 keV) | 136.53 \pm 0.04 | 136.26 \pm 0.05 | 136.64 \pm 0.17 | 136.39 \pm 0.02 |
| Co-57 γ (692.0 keV) | 714.1 \pm 0.9 | 669.4 \pm 0.3 | 689.9 \pm 0.8 | 679.6 \pm 0.3 |

| Line (Energy) | Detector | | | |
|--|-------------------|-------------------|-------------------|-------------------|
| | Wibke (08M) | K07 (09B) | Daisy (09M) | Rita (10B) |
| Cu, K α (8.0 keV) | (Too weak) | (Too weak) | (Too weak) | 8.08 \pm 0.31 |
| 122.1 Co-57 γ - W,K β 1 (54.9 keV) | (See caption) | 54.77 \pm 0.11 | 54.77 \pm 0.07 | 54.56 \pm 0.07 |
| Co-57 γ (136.5 keV) | 136.46 \pm 0.18 | 136.40 \pm 0.04 | 136.38 \pm 0.03 | 136.47 \pm 0.05 |
| Co-57 γ (692.0 keV) | 669.6 \pm 3.8 | 677.4 \pm 0.7 | 690.7 \pm 1.0 | 686.6 \pm 0.2 |

Table 3.2 Table of selected lines seen in Co-57 data as seen by the phonon detectors. Relatively good agreement with known values is seen across all lines, with the largest systematic error seen in Verena of 3.4%. The 122.06 γ with escaping W K β 1 X-ray was chosen due to its proximity upper bound of the signal region \sim 10-40 keV, relatively intensity, and that it is separable from other lines. Wibke could not have its linearity tested as easily; a combination of relatively poor phonon resolution and a weak 122 γ W, K β 1 line meant the position of this line could not be accurately determined.

3.5 Summary

Oxrop, analysis software used within the CRESST collaboration, has been upgraded. This has been to reduce memory requirements, while maintaining interactive speed. Further, with a view to Oxrop's possible use with EURECA, Oxrop's internal structure has been simplified, allowing for easier analysis of new experimental data.

How Oxrop is used to determine energies in the low energy, sub MeV region of CRESST data has been outlined. Energies are reconstructed by measuring the amplitude of detector responses, exploiting the expectation that all particle induced events will have the same pulse shape in a given detector. A significant consideration is in removing non-linear effects caused by detector transition shape. Non-linear effects are reduced by the use of "truncated fits" and linearization using test-pulses. Even with these efforts, non-linear systematic errors remain, but these errors are at the few percent level below the calibration energy of 122.06 keV and above up to 692 keV.

4 Cuts for dark matter and detector response analyses

Cuts developed with a view to cleaning cryogenic data obtained from CRESST for the detector response analyses of chapters 5 and 6 are outlined. Specific contaminants seen in CRESST data are identified, and the corresponding cuts explained. These cuts are used in detector response chapters 5 and 6. While originally designed to also allow a dark matter analysis, these cuts are not used for dark matter limit setting in chapter 7. The limits in chapter 7 are extensions from the published CRESST commissioning run analysis in [106]. The cuts, acceptances and livetimes in chapter 7 are thus the same as those in that analysis.

4.1 Introduction

In the last chapter, it was shown how to get from raw data to a list of events calibrated in energy. However, a further step is required to perform analysis. The list of events must be cleaned to ensure all events were well reconstructed, particle interactions in the target crystal. The purpose of this chapter is to show how to move from the set of all records seen from cryogenic detectors, to a final set of events ready for analysis.

Readout of the transient recorder is initiated when the SQUID signal in the trigger unit (Fig. 2.11) from any detector on a module exceeds a pre-set trigger threshold. When this occurs, all detectors on a module are recorded. This trigger ensures all WIMP and target crystal, particle interactions signals are recorded during detector live-time. However, a number of noise, non-particle and partly recorded interactions can also cause a trigger and be recorded. These contaminants need to be removed before detector response analyses such as those in chapters 5 and 6, requiring cuts.

The cuts used in this work fall broadly into four categories: stability cuts (4.2.1, 4.2.2), contaminant cuts (4.2.3-4.2.8), pulse shape cuts (4.2.9, 4.2.10) and coincidence cuts (4.2.11, 4.2.12). Cuts here were developed considering those described in theses [83], [109], and [111]. Cuts in this chapter can be used for a dark matter analysis, though are not in this thesis. Efficiency discussions, relevant to a dark matter analysis, are found in Appendix B.

4.2 Cut details

4.2.1 Stability cut

SPT thermometers are used in CRESST due to their high sensitivity to small changes in temperature. A drawback of this sensitivity is a limited range of operation. Small changes in operating point of tens of nanokelvin can impact detector response. It is thus critical that all detectors in a module are at a constant operating point over an entire run. Temperature stability is ensured by periodic injection of heating control pulses. Control pulses heat thermometers out from within the linear region of their superconducting transition. The resulting voltage change is measured as the “control pulse height”. Negative feedback with respect to the heating bias keeps control pulse height at a fixed value for the duration of a run. However, external traumas can cause a thermometer to fall outside of its transition for short periods. Such traumas can be from surrounding works or earthquakes [112]. As operating point stability is necessary for all detectors, a cut is made on control pulse height, to ensure that every detector in a module is always at its pre-set operating point.

This cut is performed under the following protocol. First, control pulses are themselves cut to ensure that only real control pulses are considered. Control pulses can be misrecorded due to pile-up; while a “control pulse” is recorded, a particle interaction can delay the actual sending of a heating pulse. This can cause non-control events to be recorded in place of control events. By considering the

delay of the firing of a control pulse with respect to the recording window, such spurious events can then be removed. After cuts, the control pulse height distribution is then found (Fig. 4.1). The central peak is well described by a Gaussian distribution. However, the tails of this distribution are not so. Further, it is more likely for lower pulse heights with respect to the most probable control pulse height to occur than higher ones. This asymmetry can be seen in Fig. 4.1, and is partially due to pile-up of control and particle pulses, generally leading to reduced pulse height estimates.




Fig. 4.1 Distribution of control pulse heights in the first 14 recording days of the light detector “Franz” in module 01T. Fitting to the central region with a Gaussian distribution is shown by the red dashed line. Control pulses have been cut to ensure only valid control pulses are considered. Cut limits, using the method outlined here, are indicated by dotted blue lines. The total distribution is non-Gaussian and asymmetric, with more events below the lower cut limit than above the upper limit.

Operating point stability cut limits are made noting the non-Gaussian nature of the pulse height distribution. First, the median and variance of the control pulse height distribution is determined using the Median Absolute Deviation (MAD) method. MAD provides a robust mechanism for describing a distribution which reduces the effect of outliers caused by disturbances. This median and deviation is then used to define an approximately Gaussian “central region”, over which the mean and RMS are determined. Control pulse height cut limits are set such that pulse heights outside of $\pm 3\sigma$ from the mean of this central region are considered “unstable”.

The distribution of unstable control pulse heights in time provides for a cut. However, it is seen that in this time distribution, there are many occasions with only a single outlying control pulse height (Fig. 4.2). This is where only one unstable control pulse height is seen without another adjacent in time. Such single outliers are not in general indicative of instability. These occur most frequently due to pile-up, where a particle interaction occurs shortly before a control pulse. Cuts are then only placed when two or more unstable pulse heights are seen adjacent in time. Time limits of this cut are set noting that nothing can be said about detector stability between control pulses. The cut is then between the last good and the next good control pulse around control pulse heights tagged as unstable. Any particle events in-between are removed due to possible instability. This is shown diagrammatically in Fig. 4.2, and is performed for each detector individually. One last consideration is for long (>30s) periods where no control pulse was recorded. If a gap between control pulses is larger than this time, then this time period is also cut.




Fig. 4.2 Control pulse heights in time for detector "Franz" in module 01T. The same cut limits as shown in Fig. 4.1 are applied. Several single outliers are seen, but only periods with multiple outliers are cut.

4.2.2 Rate

In the most recent run, noise on detectors was not constant in time. The baseline noise on detectors could increase radically for short periods, causing “excess noise” events and “microphonics” events. The cause of “microphonics” was attributed to mechanical vibrations induced by the cryostat [113]. An example of a microphonics-induced event is shown in Fig. 4.3. Both types of event are dubbed “noise-only” records, and occur when noise volume exceeds detector trigger threshold. Particle interactions in these high noise periods could not be reliably attributed to dark matter interactions, thus these periods must be cut.




Fig. 4.3 Microphonics event in phonon detector “Maja” on module 06B. During periods of rapid triggering, it is possible for large levels of low frequency noise to be seen on the phonon detector. Neglecting non-linear corrections, the energy-pulse height conversion in Maja is approximately 0.032V per 10 keV.

If noise is constant in time, the recording rate in real time will follow a Poisson distribution. Time-varying noise generally caused short bursts of rapid recording, from a few minutes to a few hours. In these periods, recording is much more frequent than the expected Poisson interaction rate, as seen in Fig. 4.4. This behaviour allows for another time cut. Events are split into short time bins, and periods are excluded when the number of pulses recorded per unit time significantly exceeds the Poisson expectation value.

The rate cut is however not 100% effective; bursts of noise can be for a very short time, or only enter the very beginning or end of a real time bin. Noise-only events can then escape this cut. These remaining events are however efficiently removed by the RMS cut (section 4.2.10), as the additional noise causes excessive degradation of goodness-of-fit when fitting with a standard pulse.




Fig. 4.4 All particle events in the first 100 hours of recording for module 06B, in 0.1h time bins. Short periods of rapid recording are seen, where the phonon detector “Maja” has triggered repeatedly. In these periods, most recordings are noise only. Cut limit on the number of events per 0.1 h is shown (dashed line).

4.2.3 Trigger delay

It is possible for interactions to occur in the light detector itself, not in the scintillating crystal target. Such direct light detector interactions are dubbed *light-only* interactions. This occurs particularly frequently in detectors that have been fitted with a dedicated Fe-55 X-ray source, intended to calibrate the light detectors directly. For a light-only event, negligible response is seen in the phonon detector, and the phonon detector does not trigger in general. As these events have significant light signals and negligible phonon signals, they are very different from dark matter signals (significant phonon, negligible light signal). These events can then be efficiently cut by ensuring that the phonon detector triggered, using the parameter “trigger delay”.

4.2.4 Peak Position

Fig. 4.5a and Fig. 4.5b show two types of recording not suitable for analysis. The first occurs when a high amplitude interaction has occurred shortly before a recording. This can be either due to the thermometer not having yet returned to its operating point following an event on a previous recording, or due to an event occurring during a period when the trigger was inactive. In either case, the detector re-triggers immediately upon trigger reactivation. As this recording does not include the onset of the pulse, energy cannot be determined.

A similar contaminant is shown in Fig. 4.5b. A particle interaction occurs before the trigger is enabled, but starts after the beginning of the pre-trigger region. While the amplitude of such pulses can be re-constructed, the baseline region before the pulse onset is too short to ensure the detector returned properly to its baseline position before recording.

The “peak position” in a pulse is the position of the maximum in a record. In both these event classes, the peak position is well within the pre-trigger region. A peak position cut then removes these events. Cut limits are set far from the distribution of normal particle events, which generally have peaks shortly after the pre-trigger region 81.92 ms into a record.

4.2.5 DAQ range

It is possible for a record to reach the limits of the DAQ system. This can be caused by, for example, by sharp downwards noise fluctuations (see below). Events reaching the DAQ limits on any record in an event cannot be properly reconstructed, and so are cut.




Fig. 4.5a Typical falling baseline on phonon detector “Rita” in module 10B. An interaction had occurred shortly before the record, and the corresponding detector response occurs over the course of the record. Fig. 4.5b shows an interaction occurring at the start of a record, on the same detector and module as Fig. 4.5a. The operating parameters before pulse onset cannot be reliably determined from the record in Fig. 4.5b.

4.2.6 Delta voltage

A contaminant peculiar to CRESST recordings is the “delta-spike”, an example seen in Fig. 4.6. Such events mostly occur simultaneously across all detectors, and briefly ($\sim 100\mu\text{s}$) take all detectors towards, but are limited by, the lower DAQ limit of -10V . The source of these disturbances is unknown. However the source of these noise events must be after the SQUID feedback electronics, as such events never cause a SQUID reset (section 4.2.8) despite being capable of taking the recording to -10V , below the limits of the SQUID feedback system.

Given that these noise events occur only *mostly* on all detectors, coincidence with other detectors cannot be relied upon to completely remove these events. A parameter was designed in [109] to remove these contaminants, named the “delta voltage”. This parameter records the largest voltage drop between adjacent and next-to-adjacent points in a record as a ratio to the baseline noise, the *delta-voltage* parameter. Records with very large values of delta voltage are then cut.




Fig. 4.6 Typical delta-spike. Almost every cryogenic detector (coloured lines) triggers at the same time, leading to a recording in almost all detectors. Every detector is driven down in voltage, limited by the DAQ limit of -10V. Stated time is with respect to the recording time on module 01T.

4.2.7 Amplitude

High energy (> 1 MeV) interactions, such as most α decays, lead to large amplitude detector responses. Responses to these interactions can exceed the amplitude range of the test pulses (section 3.4). The amplitude where this begins is detector dependent, but is generally well beyond 1 MeV. As detector linearity is un-probed in this region, attempts to reconstruct the equivalent test pulse voltage in the CPE process can result in nonsensical results. Thus, if the measured amplitude is larger than the amplitude range of test pulses on any detector in a module, the event is cut.

4.2.8 Right – left of baseline

If the current through the SQUID input coil changes too rapidly, this change can exceed the slew-rate of the corresponding SQUID, causing a quantum flux loss. The voltage equivalent value of a flux quantum depends upon detector, and is generally ~ 0.85 V. An example is shown in Fig. 4.7a. The result is that the voltage readout at the end of a record is significantly below the beginning. A quantum flux loss may also lead to a detector being below the limits of the SQUID

feedback system, generally slightly above the -10V DAQ limit of detectors. This will cause a SQUID reset, which leads to an increase in SQUID output voltage equal to ten flux quanta, $\sim 8.5\text{V}$. This voltage increase leads voltage at the end of a record to be very much higher than at the beginning (Fig. 4.7b).

A cut is then made on the “Right – left of baseline” parameter, the gradient across the record. This cut is not energy independent. Higher amplitude pulses, such as those resulting from alpha interactions in detectors with relatively long decay times do not return to the baseline by the end of recording. Cut limits are then set allowing for a linear change in right-minus-left of baseline with amplitude.




Fig. 4.7a Flux quantum loss in the phonon detector “Sabine” on module 06T. A flux quantum is the equivalent of an 0.85V change on Sabine’s SQUID voltage, and almost this drop is seen across the record. Fig. 4.7b shows a typical SQUID reset. The difference in voltage from the start to the end of the record is an integer (10) number of flux quanta.

4.2.9 Peak Position – Peak Onset

Three of the phonon detectors on complete modules included in this work are “glued” detectors, in modules 06B, 08B and 10B. In these detectors, a small carrier crystal with a tungsten thermometer evaporated onto it is glued to the larger crystal [77]. However, a side-effect of this change is that interactions can occur in the carrier crystal itself. Such interactions have a distorted phonon detector pulse shape, with faster rise and decay times, as shown in Fig. 4.8.




Fig. 4.8 Carrier crystal hit on the glued phonon detector “Rita” in module 10B. Crystal hits (blue) have much faster rise and decay times than standard pulse templates built from bulk crystal hits (black).

On all phonon detectors, interactions in the phonon thermometer cause similar pulses with fast rise and decay times. Events occurring either in a carrier crystal or in the phonon thermometer can see a quite reduced light signal with respect to the measured phonon signal, and thus need to be cut in a dark matter analysis.

Cuts are made noting that there is a large difference in the rise time of these pulses. Carrier crystal interactions have very fast, sharp increases in voltage at the point of interaction. The difference between the peak position and peak onset is a measure of this rise time, and this parameter is shown in Fig. 4.9. From Fig. 4.9, carrier crystal hits form a population of events separable from bulk crystal hits in this parameter. Cut limits are then set to remove all carrier crystal or direct thermometer hits, yet minimise the removal of bulk hits. While this cut has been designed to remove a specific contaminant, the peak position - peak onset is the first of two cuts that can also be labelled “pulse shape” cuts.




Fig. 4.9 Peak position – peak onset vs. energy for the first 14 recorded days of phonon detector “Rita”, module 10B. Carrier crystal hits form a distribution in peak-position – peak onset (below dashed line) separable from bulk hits (above dashed line).

4.2.10 RMS

The RMS of a pulse is a measure of the total discrepancy of a record in a standard pulse template fit. Cuts on this discrepancy can be used to remove several contaminants not dealt with by other cuts. An example is *pile-up*, shown in Fig. 4.10. Non-particle like interactions can also be removed with this cut. These can be noise-induced triggers that have escaped the rate cut in section 4.2.2.

Particle induced, non-bulk crystal interactions can also be cut. Pulses such as *excess light* events can also be removed by this cut. Excess light events appear as a contaminant where more light is seen than expected for electron recoils. The cause of these events is conjectured to be backscattered electrons from the crystal surface [83], causing some part of an interaction to be directly in the light absorber. These events show a slightly different pulse shape (faster rise time) compared to normal events, and thus will also be cut by an RMS cut. Lastly, long decay-time events described in chapter 7, can also be removed by this cut.




Fig. 4.10 Pile-up on phonon detector “Maja” (module 06B). Two particle interactions have occurred in the same record. Significant excursion from the expected response of a single particle like interaction is seen.

Fit RMS is not energy independent in general. Despite the use of the truncated fit described in section 3.4, there remains a small energy dependence on fit RMS on several detectors (Fig. 4.11). This energy dependence is here described by a simple linear increase in RMS with fitted amplitude in cuts. RMS cut limits here are set more tightly on the phonon detector than the light detector, as in [111]. This is because the short duration, excess noise that was the cause of the “Rate” cut more strongly affected phonon detectors than light detectors. Events can then have loud noise amplitude on the phonon detector but little noise amplitude on the light detector. Such events can be reconstructed with significant phonon and zero light amplitude, naively mimicking a WIMP signal. Phonon RMS cuts are then tighter, to fully remove excess-noise events that escape the rate-cut.

Some detectors show a strong variation in noise over the course of data taking. A weakness in the RMS cut here and in other works is that it is generally set once for the entire run – thus, changes in noise can change the efficiency of this cut at removing certain contaminants. Further, the efficiency of the RMS cut at removing some contaminants does not appear to be energy independent. This last effect is discussed briefly in chapter 7.




Fig. 4.11 RMS cut on phonon detector Maja in module 06B. Data shown is for the first 13 live days of the most recent CRESST run.

4.2.11 Coincidence - Muon Veto

The muon veto was designed to remove muon related events in or close to the experiment. The muon veto however triggers on both muons and incident γ rays. However, muons can cause multiple panel triggers, whereas single γ rays will not¹. Muons are also expected to give a larger pulse height seen in the veto than γ rays. The original intention was then to use the muon veto to detect through-going muons by looking for triggers in two panels, and with a minimum pulse height to reject radioactively induced triggers.

However, the muon veto is not fully efficient [53]. The full solid angle of possible muon incidence into the cryostat is not covered (1.3% gap). It is also possible for muons to hit the corners or sides of a panel, leading to a lower pulse height than one might expect for muons. Cut parameters are then set such that if any panel triggers within a time window of ± 10 ms of a veto trigger an event is cut. Window limits are set for compatibility with the analysis of [111]. No minimum pulse height requirement is set. In [95] the muon tagging efficiency was predicted to be 98% with strict muon cut definitions. As all muon veto panel triggers are used

¹ Though multiple γ rays could cause multiple panel triggers.

here, the muon tagging efficiency here will be at least this, though many accidental γ ray coincidences will be included. The trigger rate of the muon veto is 2.9 Hz in the training data set, thus the efficiency loss introduced by this cut is 5.8%, from both muons and accidental coincidences.

4.2.12 Coincidence - detector

Naturally occurring radioactive decays can cause triggers in multiple CRESST modules. This can be due to multiple Compton scattering of photons, or multiple particles produced in a decay, such as when two γ rays are produced in β decay of Lu-176, a trace contaminant in CRESST detectors [83]. These separate particles can interact in separate CRESST detectors, causing a coincident trigger.

More importantly for dark matter searches, neutrons have a non-zero probability of multiply scattering in different CRESST modules. Possible neutron sources have been simulated in [100], and multiple scattering of neutrons in CRESST data has been studied in [53]. These neutrons come from two main sources: radioactive decays (s.f. or (α, n)) or induced by muons. Muon induced neutrons will only be seen if the muon veto cut fails. This could happen if the muon passed through the hole in the top of the muon veto, or if highly energetic neutrons are created outside the neutron shield, and thus also the muon veto [100]. Such neutrons can be energetic enough to penetrate the neutron shield and reach detectors, and can also induce secondary neutrons in the lead and copper shields.

Both sources of neutrons were shown to have a non-zero probability of scattering in multiple detector modules in [53]. WIMPs, on the other hand, have a very small cross section with matter. The probability of a single WIMP multiply scattering in CRESST is vanishingly small. Events occurring across multiple modules simultaneously then cannot be from a WIMP interaction alone. This behaviour allows for a cut, enacted under the following protocol. First, the onset

of a pulse is used to estimate the interaction point in microseconds for every recorded event in every active CRESST module. Events are then checked for coincidence with any other detector module in a short time-window. If events are coincident, they are cut. The time window limits are selected to be $\pm 5\text{ms}$, again to be compatible with [111].

4.3 Live-time

Detector *live-time* is relevant to a dark matter analysis. It is the time during which a dark matter interaction could have been recorded. This is reduced from the total recording time by a combination of detector trigger inactivity and cut reduction. Different cuts can have different effects on live-time. Cuts removing spurious, non-particle events causing a trigger only remove the duration of that event's recording (e.g. the trigger delay cut). Cuts that infer periods of detector instability (e.g. the stability cut) remove all live time over the unstable period. More general cuts, such as the "RMS" cut, require further assumptions about the kind of events being cut. However, no analysis of data from the new run, Run 32, is used to produce new dark matter limits in this thesis; cuts here are not used for the extended dark matter analysis of the commissioning run in chapter 7. The interested reader is then directed to Appendix B for a further discussion of live-time effects by cuts.

4.4 Total cut effects and summary

The full data set generated by the data acquisition system of CRESST's cryogenic phonon and light detectors contains records that are pathological in a dark matter or other analysis of CRESST data. Many records cannot be reliably attributed to particle interactions in the target crystal; this can occur due to operating point instability, noise, SQUID flux losses / resets, and non-target crystal particle interactions.

| Cut name | 01T | % | 04B | % | 06B | % | 06T | % | 08B | % |
|----------------------|-------|------|--------|------|--------|------|-------|------|-------|------|
| Total | 41975 | 100 | 154443 | 100 | 147437 | 100 | 14474 | 100 | 66200 | 100 |
| Stability | 36435 | 86.8 | 135820 | 87.9 | 115976 | 78.7 | 13279 | 91.7 | 57565 | 87.0 |
| Rate | 32982 | 78.6 | 110571 | 71.6 | 107421 | 72.9 | 13279 | 91.7 | 55350 | 83.6 |
| Trig. delay | 31656 | 75.4 | 12602 | 8.2 | 86477 | 58.7 | 12504 | 86.4 | 14436 | 21.8 |
| Peak pos. | 26442 | 63.0 | 10991 | 7.1 | 44896 | 30.5 | 10719 | 74.1 | 10767 | 16.3 |
| DAQ range | 26290 | 62.6 | 10958 | 7.1 | 44850 | 30.4 | 10633 | 73.5 | 10753 | 16.2 |
| Delta volt. | 25900 | 61.7 | 10864 | 7.0 | 44158 | 30.0 | 10447 | 72.2 | 10577 | 16.0 |
| Amplitude | 22713 | 54.1 | 7652 | 5.0 | 21986 | 14.9 | 8593 | 59.4 | 9634 | 14.6 |
| R.-L. base. | 22406 | 53.4 | 7268 | 4.7 | 21638 | 14.7 | 7708 | 53.3 | 9584 | 14.5 |
| P. pos.- ons. | 22368 | 53.3 | 7256 | 4.7 | 17552 | 11.9 | 7701 | 53.2 | 7716 | 11.7 |
| RMS | 20787 | 49.5 | 6740 | 4.4 | 17017 | 11.5 | 7482 | 51.7 | 6587 | 10.0 |
| Muon veto | 19667 | 46.9 | 6399 | 4.1 | 16028 | 10.9 | 7064 | 48.8 | 6275 | 9.5 |
| Coincidence | 18930 | 45.1 | 5936 | 3.8 | 15344 | 10.4 | 6744 | 46.6 | 5731 | 8.7 |

| Cut name | 08M | % | 09B | % | 09M | % | 10B | % | Total | % |
|---------------------|-------|------|-------|------|-------|------|-------|------|--------|------|
| Total | 39092 | 100 | 25427 | 100 | 12214 | 100 | 84211 | 100 | 585473 | 100 |
| Stability | 36315 | 92.9 | 22625 | 89.0 | 11258 | 92.2 | 76122 | 90.4 | 505395 | 86.3 |
| Rate | 36315 | 92.9 | 22605 | 88.9 | 11258 | 92.2 | 66491 | 79.0 | 456272 | 77.9 |
| Trig. delay | 26897 | 68.8 | 21083 | 82.9 | 10698 | 87.6 | 17900 | 21.3 | 234253 | 40.0 |
| Peak pos. | 23501 | 60.1 | 16851 | 66.3 | 8436 | 69.1 | 11857 | 14.1 | 164460 | 28.1 |
| DAQ range | 23422 | 59.9 | 16841 | 66.2 | 8353 | 68.4 | 11857 | 14.1 | 163957 | 28.0 |
| Delta volt. | 23001 | 58.8 | 16768 | 65.9 | 7890 | 64.6 | 11502 | 13.7 | 161107 | 27.5 |
| Amplitude | 14965 | 38.3 | 11642 | 45.8 | 7247 | 59.3 | 8378 | 9.9 | 112810 | 19.3 |
| R.-L. base. | 13948 | 35.7 | 11290 | 44.4 | 6999 | 57.3 | 8320 | 9.9 | 109161 | 18.6 |
| P. pos.-ons. | 13943 | 35.7 | 11286 | 44.4 | 6977 | 57.1 | 8217 | 9.8 | 103016 | 17.6 |
| RMS | 13511 | 34.6 | 10757 | 42.3 | 6864 | 56.2 | 7973 | 9.5 | 97718 | 16.7 |
| Muon veto | 12802 | 32.7 | 9681 | 38.1 | 6504 | 53.3 | 7510 | 8.9 | 91930 | 15.7 |
| Coincidence | 12087 | 30.9 | 9331 | 36.7 | 6071 | 49.7 | 7183 | 8.5 | 87357 | 14.9 |

Table 4.3 Reduction of events across all modules with applied cuts in the first 14 live days of Run 32.

Further, particle interaction characteristics cannot be well reconstructed if interactions occur close in time to another interaction, the effect of pile-up. Pathological records caused by these effects have been identified, and the cuts used to separate these events have been outlined. This has been with a view to preparing data for analyses in subsequent chapters.

Additionally, it is noted that both the muon veto and the multiply modular nature of CRESST detectors can be used to isolate non-WIMP events. WIMPs would not interact with either the muon veto or across multiple detectors. This behaviour can be used to cut events occurring in either the muon veto or other detectors as not having originated from a WIMP alone. Table 4.3 then indicates the number of events removed by each cut. The method to calculate acceptances of these cuts has been left to Appendix B.

5 From light detector response to scintillation time

Relatively poorly known constants affecting CRESST light detector response are the scintillation times of CaWO_4 and ZnWO_4 . Previously, parametric modeling of CRESST light detector response has been used to infer these scintillation times, as well as to directly calibrate the light detectors themselves. This modelling has interpreted CRESST light detector response to instantaneous power inputs in terms of a well-motivated model, with one exponential detector rise time and two exponential detector decay times, derived from Eqn. 2.6.

Here, this model is applied to CRESST light detector responses in the most recent CRESST run, and is shown to be in significant tension with observed response for both scintillation events and direct interactions in the light detector. Modelling CRESST light detector response is then improved by empirical extensions to the one rise, two decay model of CRESST light detector response, allowing for extra decay times. With these extensions, light detector response to scintillation events can be modelled approximately to within estimated statistical error. Using these extensions a scintillation time of $600\mu\text{s}\pm 100\mu\text{s}(\text{sys})$ can be obtained for γ interactions in CaWO_4 , with systematic error estimated from variation in measurements on the nine light detectors analysed. For ZnWO_4 , only a single light detector could be analysed. γ scintillation time measurements for ZnWO_4 were degenerate between values of $80\mu\text{s}\pm 40\mu\text{s}(\text{stat})$ or $480\mu\text{s}\pm 40\mu\text{s}(\text{stat})$.

However, even using these extensions, direct interactions of 5.9 keV X-rays in the light detector cannot be modelled to within statistical error. This implies that direct calibration of light detectors with X-rays cannot be performed for CRESST detectors. However, as detector response to optical, scintillation photons can be well modelled, this suggests that future runs should attempt to directly calibrate light detectors with optical photons from an artificial source, as light detector response to optical photons should be similar to scintillation events.

5.1 Introduction

Calcium tungstate (CaWO_4) exists in the scheelite structure, with low temperature light emission attributed to charge transfer within the tetrahedral $(\text{WO}_4)^{2-}$ anion. On the other hand Zinc tungstate (ZnWO_4) exists in the wolframite structure, with low temperature emission attributed to charge transfer within the octahedral $(\text{WO}_6)^{2-}$ anion [74]. These processes lead to scintillation decay times of $8.9 \pm 0.2 \mu\text{s}$ and $25.7 \pm 0.3 \mu\text{s}$ at room temperature, and $90 \pm 20 \mu\text{s}$ and $110 \pm 15 \mu\text{s}$ at 9K, for CaWO_4 and ZnWO_4 , respectively [75]. However CRESST operates CaWO_4 and ZnWO_4 scintillating target crystals at much lower temperatures, close to 10 mK. At this temperature fewer measurements for the scintillation time of CaWO_4 and ZnWO_4 exist. Currently, millikelvin γ scintillation times for CaWO_4 have been measured to be $400 \pm 100 \mu\text{s}$ at 23mK [114] (using a similar method to section 5.4) and $500 \pm 60 \mu\text{s}$ at 20mK [115]. ZnWO_4 scintillation times are even less well known, measured to be $\approx 200 \mu\text{s}$ at 400 mK for α interactions [116]. Independent of this work, the same data set as used here was used to estimate the millikelvin scintillation times of γ interactions in CaWO_4 and ZnWO_4 in [80], with a different method. In [80] these times were estimated to be $\approx 400 \mu\text{s}$ - $500 \mu\text{s}$ and $\approx 110 \mu\text{s}$ respectively.

That these scintillators have a finite scintillation time means that energy deposition in the light detectors should not be modelled as instantaneous (Eqn. 2.6). Indeed, as these decays are quite long, of the order half a millisecond, this finite scintillation time significantly changes light detector response from the instantaneous energy deposition model (see below).

Separately, the actual amount of light detected in CRESST light detectors, relative to the energy of an interaction, is not found through the calibration process in section 3.4. Calibration normalizes detector response with respect to the amount of scintillation light seen for 122.06 keV γ interactions from an external source.

Thus, calibrations are only sensitive to the *relative* amount of energy in scintillation light reaching the light detector. It would however be useful to know the energy calibration in terms of the actual energy input to the light detector. This would allow determination of the fraction of energy in scintillating crystal interactions actually detected in the light detector. If crystal scintillation efficiency is known independently, then the light collection efficiency could also be determined. To attempt to determine these direct calibrations Fe-55 radioactive sources were placed inside many CRESST detector modules (Table 2.1). The sources are placed so that these X-rays can interact directly with the light detector. Fe-55 decays lead to a cascade of low energy (< 10 keV) X-rays. In principle, these interactions allow a direct calibration of the light detector.

This chapter details analysis work conducted to attempt to understand both the scintillation time of CaWO_4 and ZnWO_4 , and absolute calibrations of the light detectors. Working from a model of light detector scintillation response originally developed in [114], attempts are made by the author to model light detector response in the most recent CRESST run, Run 32. When these attempts fail, further extensions are made by the author to model detector response.

5.2 Determining scintillation time

First, the method to determine scintillation time in CRESST detectors is examined. This model is derived from that given by Eqn. 2.6, which is slightly rearranged here as Eqn. 5.1. This equation consists of one exponential rise time, τ_r , and two exponential decay times τ_0 and τ_1 :

$$\Delta T = \Xi(t)A \left[-e^{-t/\tau_r} + \left(1 - \frac{A_1}{A}\right) e^{-t/\tau_0} + \frac{A_1}{A} e^{-t/\tau_1} \right]. \quad (5.1)$$

These three time constants are referred to as *detector time constants*. This model has been the basis¹ for all derived time constants for work on both phonon and light detectors for CRESST to date (e.g. [78,80,83,84,109,110]). This model

¹ Either Eqn. 2.6, or the model in Eqn. 5.5 resulting from non-instantaneous power inputs.

assumes that power input function is instantaneous, shown explicitly as:

$$\Delta T = \int_{-\infty}^t E \delta(s - 0) A \left[-e^{-(t-s)/\tau_r} + (1 - f_1)e^{-(t-s)/\tau_0} + f_1 e^{-(t-s)/\tau_1} \right] ds, \quad (5.2)$$

with $\frac{A_1}{A} \equiv f_1$. However, scintillation of CaWO_4 and ZnWO_4 is a non-instantaneous process. Scintillation light output by these crystals can be modelled by an exponential decay, introducing a time constant τ_s representing the long *scintillation time* noted in section 2.2.1. Light detectors then see a non-instantaneous power input function. The power input function in Eqn. 5.2 is then replaced with this exponential decay, such that:

$$\Delta T = \int_{-\infty}^t P_A \Xi(s - 0) e^{-s/\tau_s} A \left[-e^{-(t-s)/\tau_r} + (1 - f_1)e^{-(t-s)/\tau_0} + f_1 e^{-(t-s)/\tau_1} \right] ds, \quad (5.3)$$

where P_A is normalized so that:

$$\int_{-\infty}^{\infty} P_A \Xi(s - 0) e^{-s/\tau_s} ds = P_0 \therefore P_A = P_0/\tau_s. \quad (5.4)$$

The equivalent relation for temperature change in time then becomes:

$$\Delta T = P_0 A \left[-\frac{e^{-t/\tau_r} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_r)} + (1 - f_1) \frac{e^{-t/\tau_0} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_0)} + f_1 \frac{e^{-t/\tau_1} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_1)} \right]. \quad (5.5)$$

Typical time constants for light detectors in previous work, obtained by fitting Eqn. 5.2, are $\tau_r \sim 500\mu\text{s}$, $\tau_0 \sim 5\text{ms}$, $\tau_1 \sim 25\text{ms}$, from [83] and [109]. Typical pulse shapes are then shown diagrammatically in Fig. 5.1.

As noted from [114] and [115] the scintillation time is expected to be $\sim 500\mu\text{s}$. The only detector time constant expected to possibly be of similar magnitude to the scintillation time is then the rise time. In such cases, with $\tau_0, \tau_1 \ll \tau_r$, Eqn. 5.5 becomes:

$$\Delta T \approx P_0 A \left[-\frac{e^{-t/\tau_r} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_r)} + (1 - f_1)e^{-t/\tau_0} + f_1 e^{-t/\tau_1} \right]. \quad (5.6)$$

In this case, τ_s and τ_r behave similarly; if they are switched, the first term is the same to within a scaling factor. This near-degeneracy is highlighted in Fig. 5.1.

Before fitting Eqn. 5.5 to measured detector response, special preparation of the data is discussed.




Fig. 5.1 Example pulse shapes showing different scintillation and rise times. All example pulses have the same athermal and thermal decay times, 5ms and 25ms respectively (with $f_1=0.8$). To illustrate the near degeneracy of τ_s and τ_r , the dashed blue line shows the $\tau_s = 100\mu\text{s}$, $\tau_r = 500\mu\text{s}$ example pulse scaled in amplitude by 107.3%. This is to illustrate this pulse's similarity to the $\tau_s = 500\mu\text{s}$, $\tau_r = 100\mu\text{s}$ pulse (green).

5.3 Standard-pulse treatment

It is expected that interactions of the same type and energy lead to identical detector responses (section 2.2.3). This behavior was used to construct template pulses in section 3.4. However, particular care is taken here to ensure accurate reconstruction of the rising part of a pulse. It is important that each pulse summed to generate the standard pulse template must have interacted at the same point in the record, to within the resolution ($40\mu\text{s}$) of recorded data. Due to noise fluctuations, triggering might not occur at identical times after the physical interaction of particles of the same energy. Noise effects can then shift records relative to each-other. If these shifts are not accounted for, the form of the standard pulse can be slightly distorted, an effect shown in Fig. 5.2b.

To remove this effect, template standard pulses are first constructed as outlined in section 3.4. The template is then re-fit to the records from which it was created, to measure the shift of each recorded pulse relative to the template. Outlying pulses are then cut, and the template pulse rebuilt. The process is repeated until the measured shift for all records, relative to the template, is zero.




Fig. 5.2a Fitted pulse onsets for the list of template events used to build a standard pulse, as noted in section 3.4, for light detector “Josef”. In Fig. 5.2b a standard pulse generated from events with zero shift relative to the most probable shift position (black) is shown relative to one generated from all pulses (blue). The main difference is that the two standard pulses are shifted slightly with respect to each other. However, noting the relative shifts, pulse shape differences are apparent very near the standard pulse onset near -2ms.

All standard pulses generated here are from 46.5 keV γ interactions caused by Pb-210 decays. This contaminant is intrinsic to all crystals, and thus such standard pulses can be built for all modules. Necessarily, the scintillation time constant determined is then of γ interactions. After finalising the standard pulse, it is noted that detectors display a small baseline gradient caused by pile-up (see section 4.2.10). The magnitude of this gradient across the record was less than 0.5% of pulse amplitude in all cases, and generally smaller than 0.1%. This effect was modelled by allowing a small linear variation in the baseline voltage.

5.4 Results

Standard pulses were constructed for nine light detectors which had a corresponding phonon detector in the most recent CRESST run. A first fit of Eqn. 5.5 is performed on the standard pulses. Fitting was performed using a numerical fitter based on generic ROOT classes and corroborated with a Levenberg-Marquardt fitter, adapted from that in [110]. Results from fitting were identical when using the two fitters, given in Table 5.1.

| Detector | Module | Material | τ_s [μ s] | τ_r [μ s] | Goodness of Fit |
|----------|--------|-------------------|------------------------|------------------------|--------------------|
| Franz | 01T | CaWO ₄ | 340 | 340 | 5.27 |
| Q | 04B | CaWO ₄ | 340 | 340 | 8.59 |
| Burkhard | 04B | CaWO ₄ | 300 | 300 | 4.81 |
| Hans | 06B | CaWO ₄ | 610 | 130 | 2.01 |
| Josef | 06T | CaWO ₄ | 630 | 130 | 1.31 |
| X | 08M | CaWO ₄ | 670 | 160 | 1.09 |
| David | 09M | CaWO ₄ | 380 | 360 | 1.09 |
| Steven | 10B | CaWO ₄ | 280 | 270 | 48.1 |
| Yoichiro | 10T | CaWO ₄ | 520 | 150 | 1.55 |
| Ulrich | 08B | ZnWO ₄ | 240 | 240 | 6.11 |

Table 5.1 Results from fitting to standard pulses, using Eqn. 5.5 for each light detector with one scintillation time, one rise time, and two decay times. Note the strong correlation between scintillation-time and goodness of fit χ^2_{\min}/ν . As fits degrade, both scintillation time and rise-time are systematically skewed towards equality in fitting. Given the significant model discrepancy, these results are recorded here for reference only, and are not intended for use in estimating the scintillation time or CaWO₄ or ZnWO₄.

To estimate fit χ^2 , noise remaining in the pre-trigger region of standard-pulses is used (section 4.1). In this region, only a noise-filled baseline is expected. This noise is used as an error estimate in the voltage coordinate across the entire record. The goodness of fit is then calculated with degrees of freedom ν given by the 8192 data points, reduced by the number of fit parameters. The resulting chi-squared distribution is a gamma distribution with mean ν and variance 2ν . The best fit is χ^2_{\min} with goodness-of-fit χ^2_{\min}/ν . Statistical errors on parameters can be found where $\chi^2_{\min} \leq \chi^2 < \chi^2_{\min} + \Delta\chi^2$, with $\Delta\chi^2$ approximately $\sqrt{2\nu}$ for large ν .

It is seen that, while some detectors were modelled well by Eqn. 5.5, many detectors displayed distinct, non-model-like behavior, showing very poor goodness-of-fit values. Fig. 5.3 shows a comparison between the two extreme cases of well-modelled (“X”) and poorly-modelled (“Steven”) signals. As the applied model was not in general a good representation of the data, the results displayed in Table 5.1 are supplied for reference purposes only. Errors have not been stated in Table 5.1 (though will be in Table 5.2). This is because given the poor fit results seen in many detectors, quoting such statistical errors would be meaningless without an understanding of additional, systematic errors.




Fig. 5.3 Fits to the light detectors “X” (left) and “Steven” (right). X sees the model (black) fit the measured response (red) approximately to within noise, with the residual of the fit shown in blue. Steven shows considerable errors, much larger than the noise, and fit residual is dominated by systematic discrepancies.

Given the degraded fit quality evident in Table 5.1, it is necessary to model the additional behavior observed in the majority of light detectors. Before doing so, certain features evident in the fitting process need to be discussed. In all fits, scintillation time and detector rise time measurements were strongly correlated. This is true both of relatively well modelled (e.g. “X”) and poorly modelled detector responses (e.g. “Steven”). In general, fitting identified two minima, with τ_s and τ_r reversible to give each minima. This result is explained by again noting that if τ_r and τ_s are switched in Eqn. 5.6, the resulting pulse shape description is very similar. To select which minima to quote in Table 5.2 as “scintillation time”, an appeal is made to previous results. In [115] and [116], mK scintillation times for CaWO_4 and ZnWO_4 are $\sim 500\mu\text{s}$ and $\sim 200\mu\text{s}$ respectively. The closest fit results to these times are selected for table 5.1. The alternative possibility can be realized by switching rise and scintillation times in Table 5.1 (and Table 5.2).

5.5 Extended model of light detector response

The model of CRESST detector response incorporating two decay times and a single rise time of Eqn. 5.5 does not appear to well represent the observed response on many light detectors. The basis for this simple model was described in section 2.2.3. Here, several other possible features that can affect detector response are discussed. The first possible cause of a distortion from the one rise, two decay time model can come from the phonon collector. Phonon collectors are incorporated into light detector design to improve phonon collection efficiency. The phonon collector could however introduce a further time constant into the light detector signal model, though evidence for this extra decay time has previously not been seen [80]. Another possible distortion could come from the underlying assumption that thermalisation of phonons is roughly frequency independent. In the case of light detectors examined here, a deviation from the one rise, two decay time behavior is seen, which could possibly be explained by a variation in lifetimes of a-thermal phonons based on frequency. Lastly, any spatial non-uniformity in phonon thermalisation has not been considered, and could also cause a variation in a-thermal phonon lifetimes.

If any of these effects is significant, it will result in additional detector decay times than in the model of Eqn. 5.5. Thus there is some motivation for adding these additional decay times, and this model is empirically extended as:

$$\Delta T = P_0 A \left[-\frac{e^{-t/\tau_r} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_r)} + (1 - \sum_{i=1}^N f_i) \frac{e^{-t/\tau_0} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_0)} + \sum_{i=1}^N f_i \frac{e^{-t/\tau_i} - e^{-t/\tau_s}}{(1 - \tau_s/\tau_i)} \right], \quad (5.7)$$

where $\frac{A_i}{A} \equiv f_i$. Eqn. 5.5 is reproduced by setting $N = 1$. The first natural extension is with $N = 2$, where the additional time constant could be explained as due to the influence of the phonon-collector. Here, N is increased simply *until the model represents the data*. This is performed under the following protocol:

- 1) Best fit parameters for the $N = 1$ case are determined across the range of τ_r, τ_s range $50\mu\text{s}$ to $1000\mu\text{s}$.

- 2) Using these results as starting parameters, fit with $N_{n+1} = N_n + 1$. Initialise the fit with a small (1%) admixture of the additional decay component. Find best χ^2 across the range of τ_r, τ_s range $50\mu\text{s}$ to $1000\mu\text{s}$.
- 3) If χ^2 is improved beyond estimated statistical error, increment N and repeat the process. Otherwise, the best fit and number of decay times needed has been determined as N_n .

This method has then departed from the simple yet theoretically well motivated model derived from Eqn. 2.6. However, it remains the case that the original model does not represent the actual behaviour of light detectors within error. In the next section it is seen that even with the simple extension of $N = 2$ in Eqn. 5.7, many detectors were still not well described. It is then desirable to see if light detector behavior *can* be modelled by this simple extension. Undoubtedly, these extensions need stronger theoretical motivation. Scintillation time estimates in the next section must be taken in this light.

5.6 Extended model fit results

Extended model fit results are seen in Table 5.2. As one would expect, adding additional decay times has improved quality of fit in all cases, especially in detectors that had a poor goodness of fit when fitting the original model.

Fig. 5.4 shows an example of this improvement.

Fitting Eqn. 5.7 to each detector's response again yielded two fit minima, with scintillation time and rise time reversible to give each minima. Measurements of scintillation time are then degenerate between the times estimated for the rise time and the scintillation time. Given that scintillation time for calcium tungstate has been previously measured to be $\sim 500\mu\text{s}$, the minima closest to this value was used to choose the minima to represent in Table 5.2, and the second minima, at values $\sim 150\mu\text{s}$ in calcium tungstate detectors, is listed as the rise time.

For zinc tungstate, the choice is more difficult. Only a single measurement of $\approx 200\mu\text{s}$ for α interactions exists separately to data used in this work, the selection cannot be made clearly. This is emphasized in the “Ulrich” entry of Table 5.2.

| Detector | Module | Material | τ_s [μs] | Error [μs] | τ_r [μs] | Goodness of Fit | N |
|----------|--------|-------------------|-------------------------------|----------------------------|-------------------------------|--------------------|-----|
| Franz | 01T | CaWO ₄ | 560 | +50/-50 | 230 | 0.98 | 3 |
| Q | 04B | CaWO ₄ | 560 | +140/-160 | 160 | 0.83 | 4 |
| Burkhard | 04B | CaWO ₄ | 510 | +220/-90 | 220 | 1.00 | 4 |
| Hans | 06B | CaWO ₄ | 630 | +10/-20 | 110 | 1.54 | 2 |
| Josef | 06T | CaWO ₄ | 640 | +10/-20 | 120 | 1.18 | 2 |
| X | 08M | CaWO ₄ | 700 | +30/-20 | 150 | 0.91 | 3 |
| David | 09M | CaWO ₄ | 650 | +140/-280 | 160 | 0.95 | 2 |
| Steven | 10B | CaWO ₄ | 500 | +70/-30 | 100 | 0.94 | 4 |
| Yoichiro | 10T | CaWO ₄ | 550 | +20/-20 | 130 | 0.75 | 2 |
| Ulrich | 08B | ZnWO ₄ | 80(480) | +40/-40 | 480(80) | 0.61 | 4 |

Table 5.2 Scintillation time results from fitting Eqn. 5.7 to standard pulses for each light detector, using one scintillation time, one rise time, and multiple decay times. As with Table 5.1, fit minima exist switching τ_s and τ_r ; ordering is chosen with the closest match to existing values in the literature. This effect is particularly important in the zinc tungstate module, as the two minima are similarly close to the expectation of $\mathcal{O}(200\mu\text{s})$.

Table 5.2 shows that, in terms of goodness-of-fit, all detectors can now be reasonably well modelled by the addition of extra decay times. This should be compared to the case of direct light detector interactions (next section), which can-not. Some detectors display over-fitting (e.g. “Yoichiro” and “Ulrich”). Additionally, some detectors (“Hans”) still appear to fit poorly. This is conjectured to be caused by noise in the pre-trigger region of standard pulses not perfectly representing noise over the whole standard pulse. This can be caused by the cut of section 5.3, intended to remove noise-induced trigger position dependence. This cut is affected both by small baseline gradients and noise close to the standard pulse onset, and could cause asymmetry in noise magnitude across a standard pulse.

A few detectors were well described by $N = 2$. For these detectors, the additional time constant can then be understood as due to the influence of the phonon

collector. However, other detectors required several additional decay times to fully describe measured response. For example, “Steven” was seen to require $N - 1 = 3$ additional decay time constants to model detector response to approximately within noise error, shown in Fig. 5.4.




Fig. 5.4 Comparison of fit residual to “Steven” with $N=1$ (left), as in Fig. 5.3 and $N=4$ (right), with goodness-of-fit values 48.1 and 0.94 respectively. Adding additional decay times allows detector response to be modelled almost entirely within noise.

Scintillation time estimates for CaWO_4 in Table 5.2. do not agree with each-other within statistical error. This effect was true also of applying the original model of Eqn. 5.5 (Table 5.1), and in the independent measurements of the same detectors in [80]. Several possible causes for these discrepancies are now considered.

First, it is possible that an excess of contaminant events in the standard pulse, such as the excess light events noted in section 4.2.10, could contribute a possible systematic shift between measurements. To investigate this possibility, it is noted that such contaminants are expected to be more common at lower energies [111]. Standard pulses were then built from the lower energy interactions of 24.5 keV γ rays caused by Ac-227 beta decay, for detectors, “X”, “Q” and “Burkhard”. However, fitting displayed identical scintillation time results to the pulses constructed from 46.5 keV γ ray interactions.

Another possible explanation is simply that CaWO_4 shows a range of scintillation time constants at mK temperatures, depending on the individual crystal manufacture process. In this case, given that the crystals under observation in this data set are from separate manufacturers and production processes, one would expect this deviation to be largest on detectors of separate manufacturers and smallest on those of the same manufacturer. Fig. 5.4 shows scintillation time versus manufacturer. While the data set is small, no such manufacturer dependence is evident in the data set.

A further possible explanation for the CaWO_4 result is a small variation in detector preparation after receipt from the manufacturer. All CRESST crystals undergo annealing at 800°C for 48 hours in a pure oxygen flow. Most CaWO_4 crystals in this data set have then subsequently had a tungsten thermometer evaporated onto them. The exception are crystals in modules 06B and 10B which are “glued” detectors, with light detectors “Hans” and “Steven” respectively.




Fig. 5.5 Scintillation times determined from Table 5.2, against manufacturer. The rough average measurement of $600\mu\text{s}$ is shown by the dashed black line. No strong manufacturer dependence is evident.

All of these processes are known to effect scintillation properties in terms of light yield [117], and different crystals show different final light outputs [118]. One possible explanation is that different crystals display different scintillation time

constants too. However, given the empirical nature of the applied model, these variations are attributed to a purely systematic measurement error. All scintillation time measurements for CaWO_4 fall approximately within the range $\approx 600\mu\text{s} \pm 100\mu\text{s}(\text{sys})$, with systematic errors estimated from the distribution of the measurements from each of the nine light detectors analysed. This is compared to $400 \pm 100\mu\text{s}$ at 23mK [114] and $500 \pm 60\mu\text{s}$ at 20mK [115]. Measurements here then appear high, though not significantly so. Independently, from the same data and a different method, scintillation time was determined to be $\approx 400\mu\text{s}$ - $500\mu\text{s}$ in [80].

For ZnWO_4 , only a single light detector could be analysed. The minimum of the fit was degenerate between scintillation time measurements of $80\mu\text{s} \pm 40\mu\text{s}$ and $480\mu\text{s} \pm 40\mu\text{s}$. This compares to $\approx 200\mu\text{s}$ for α events at 400mK from [116], and independent measurements of $\approx 110\mu\text{s}$ for the same detector from [80], using a different method. Unfortunately, as only one ZnWO_4 module is fully operational, it is not possible to see if the same variation exists in ZnWO_4 measurements as in CaWO_4 measurements.

The light detector “Ulrich” used for this measurement is also the only silicon light detector, with all others being of silicon-on-sapphire design. These two designs tend to have different detector time constants [119]. In particular, the rise time is expected to be much shorter in silicon detectors. The rise time of the detector is determined by the decay lifetime of non-thermal phonons in the detector. This decay is mostly caused by inelastic surface reflections, which occur more frequently in the thinner silicon detectors [119].

5.7 Direct calibration

Several light detectors are fitted with Fe-55 calibration sources. Fe-55 undergoes electron capture to Mn-55 with a half-life of 2.73 years. This decay leads to an X-ray cascade, with primary energies 5.90 keV from the Mn-K α 1, 5.89 keV K α 2,

and also 6.49-6.54 keV $K\beta'2$, at relative intensities 0.1656, 0.0845, 0.034 respectively [120]. Other lines exist; for example, the presence of Aluminium foil around the Fe-55 source should give an Al fluorescence line at 1.49 keV $K\alpha$. These other lines however are often too weak to provide useful measurements.

Six active light detectors have been outfitted with Fe-55 sources: Burkhard, Q, Hans, Ulrich, X and Steven, causing count rates in each of $\mathcal{O}(10$ mHz). The sources are placed on the side of the light detector opposite the phonon detector, in order not to trigger the phonon detector. These X-rays could not affect WIMP searches even if they were to reach the phonon detector. The X-ray energies are too far below the 10 keV minimum threshold for WIMP acceptance.

The previous section showed that all light detectors required additional decay times over Eqn. 5.5 to have their response to target crystal interactions modelled within noise. It is then desirable to see if the same is also true of direct interactions in the light detector from Fe-55 X-rays. As the light detector Steven showed the largest deviation from the model of Eqn. 5.5 (Table 5.1), yet could ultimately be well modelled by extensions in Eqn. 5.7, it is focused on here. A standard pulse is built from the 5.9 keV Mn-55 $K\alpha$ peak of Fe-55 direct interactions in the light detector. The only change in pulse shape expected between direct interactions and scintillation interactions is that the scintillation time is effectively infinitesimal for direct interactions. First, the pulse description yielded from applying Eqn. 5.7 to the γ scintillation standard pulse was applied to the direct interaction standard pulse with $\tau_s = 0$, seen in Fig. 5.6.

It is clear that signal evolution of direct light detector interactions behaves very differently to the response to scintillator hits. However, to directly calibrate the light detector, one must relate detector response caused by target crystal interactions to detector response caused by direct light detector interactions. The question is: can any meaningful information be obtained that can relate direct

light detector interactions to crystal scintillation interaction? As a second attempt, Eqn. 5.7 is fit to the 5.9 keV $K\alpha$ peak of Steven with all parameters free (except scintillation time, left at zero). The results are shown in Fig. 5.7a.




Fig. 5.6 Fit results of the scintillation pulse events (black) are applied to the data from direct hits of the light detector (red), setting $\tau_s = 0$. This model clearly very poorly describes data,




Fig. 5.7 Fits to the direct hits 5.9 keV X-rays of Rita/Steven with Eqn. 5.7. Fig. 5.7a shows a fit (black) with $N = 4$ to the measured pulse (red) with no attempt to model for transition linearity. Fig. 5.7b shows a fit accounting for the influence of transition non-linearity. Neither fit models data well, with goodness of fit values 262 and 154, respectively.

Again, fits could not reproduce the observed response (goodness-of-fit 262). However, it should be noted that, in practice, the 5.9 keV X-ray line from Fe-55 decays lies within the non-linear region of that detector (as described in section

3.4). Without taking transition non-linearity into account, any deviation in detector response could be attributed to transition non-linearity effects. To account for this non-linearity, it was shown in [110] that it is possible to fit to high amplitude, scintillation pulses to determine the evolution of pulse shape behaviour in the non-linear region. This was also performed here, with the results shown in Fig. 5.7b. When accounting for non-linear adjustments to pulse shape, fit results do improve (goodness of fit 154). However, light detector response can still be not modelled closely to within estimated noise error.

5.8 Conclusions

The behaviour of light detectors in the most recent CRESST run was analysed. The model of light detector behaviour of one rising and two falling exponential decay times, convolved with an exponentially decaying scintillation light output, was applied. It was seen that this model did not well describe light detector response for either scintillation events or direct light detector hits to within estimated noise error.

The possibility that additional decay times exist was seen to be motivated, particularly due to the presence of the phonon collector in CRESST's light detectors. Empirical modifications to the expected behaviour were then made, introducing additional exponential decay times. Adding these decay times did allow scintillation events to be modelled close to estimated error. Fits results in calcium tungstate detectors were degenerate between two minima at $\sim 150\mu\text{s}$ and $\sim 600\mu\text{s}$. Previous results that the scintillation time of calcium tungstate is $500\pm 60\mu\text{s}$ at 20mK [115] were used to select the $\sim 600\mu\text{s}$ minima. All resulting γ scintillation time estimates then fell within the range $\approx 600\mu\text{s}\pm 100\mu\text{s}$. However measurements did not agree with each-other within estimated statistical error. A single measurement of scintillation time for zinc tungstate was also made, with fit results degenerate between scintillation times of $80\mu\text{s}\pm 40\mu\text{s}$ or $480\mu\text{s}\pm 40\mu\text{s}$,

where the errors are statistical. This compares to previous measurements that millikelvin scintillation times for α interactions are $\approx 200\mu\text{s}$ for alpha events at 400mK from [116], and independent measurements of $\approx 110\mu\text{s}$ for the same detector from [80].

These same extensions, allowing for extra decay times, were also applied to the response of direct, 5.9 keV X-ray interactions in the light detector. These extensions however did not allow behaviour of direct hits of the light detector to be modelled within estimated error. Accounting for thermometer transition non-linearity also did not remedy this poor description of detector behaviour. Thus, it does not appear possible to perform a direct calibration of CRESST light detectors with the Fe-55 X-ray sources currently in CRESST. However, given that response to optical, scintillation photons can be well modelled, this suggests that in future runs, optical photons, of similar frequency to the scintillation light emission of the target crystals, should be used to directly calibrate the light detector.

6 Position dependence in CRESST target crystals

CRESST's cylindrical target crystals have finite size, 4 cm in diameter and height. Position dependence refers to the possibility that the signal seen in CRESST's light and phonon detectors could be related to interaction position in these crystals, particularly as a function of radius. Given that external γ sources of energies ranging from 122.06 keV to 911.2 keV were used for calibration in the most recent CRESST run, the variation in penetration depths of these γ rays in target crystals with energy can be used to probe behaviour at different positions in the crystal.

Anti-correlation between signals from phonon and light detectors is discussed in this context. It is shown that the gradient of this anti-correlation is constant at all energies without position dependence, but can change if position dependence is introduced. Anti-correlation gradients are then measured, accounting for various background phenomena that can distort measurements, and it is shown that changes in this anti-correlation gradient are evident in at least one module in calibration data from the most recent CRESST run.

Position dependent effects on total light collection efficiency – the amount of light produced that actually reaches the light detector – are also considered. Light collection efficiency is shown to vary radially in target crystals, increasing radially towards the crystal centre with interaction position. The magnitude of the change in light collection efficiency from the surface of the crystal to the bulk of the crystal is then compared to light detector resolution. A correlation is seen between the magnitude of surface-to-bulk changes in light collection efficiency and the (high energy) light detector resolution of detectors. This is explained simply as larger variations in light collection efficiency naturally lead to a worse resolution of light detectors.

6.1 Introduction

Position dependence expresses the possibility that the light or phonon signals measured for a given energy and type (e.g. α , β , γ) of interaction could be a function of position within that crystal. In CRESST analyses to date (e.g. [53,106]) this effect is only considered as possibly influencing the light detector resolution, as part of the resolution term that linearly scales with energy (σ_2 in Eqn. 7.1), and thus signal acceptance. However, position dependence has been predicted [121] and measured [83] to cause further effects in CRESST data. This work expands on these observations, further reinforcing the evidence that position dependence causes a variety of effects in CRESST detectors.

6.2 Modelling detected signals

To consider position dependent effects, one first needs to describe the signals detected in the phonon and light detector relative to the energy of an event. First, detected light is considered. The average amount of light produced in a CRESST crystal for a γ interaction can be expressed as:

$$L_{Prod} = EK_{\gamma}(E), \quad (6.1)$$

where $K(E)$ is the *scintillation efficiency* at energy E . The scintillation efficiency of γ events in CaWO_4 is measured to be $K_{\gamma} \sim 8.3\%$ at cryogenic temperatures [75]. Variations in scintillation efficiency are limited to no more than a few parts in a hundred beyond ≈ 100 keV [89]. $K_{\gamma}(E) \approx \text{const} \equiv K$ then serves as a good first approximation. Explicit energy dependence is dropped in the following equations, but the possibility of a small variation in K_{γ} with E is considered in section 6.4.3. Emission is in the optical frequency range with a peak emission at 430nm, at which frequency the refractive indices of CaWO_4 are $n_0 = 1.96$ and $n_0 = 1.98$ [122].

We now wish to express the signal actually seen in the light detector in response to these optical photons, relating this to the light produced L_{prod} . This must account for several effects. The first is the fraction of light produced that is then reabsorbed by the crystal, B . Visible scintillation light has a finite attenuation length in CaWO_4 crystals of 10-30cm (from the review in [121]), so light travelling through the 4cm diameter, 4cm height CRESST crystals has the possibility of being reabsorbed. B took values between $\sim 65\%$ and $\sim 85\%$ depending on position, from simulations of polished CRESST crystals in [121]. Light reabsorbed in the crystal will create phonons, and these phonons will subsequently be detected by the phonon detector, as discussed shortly. $(1 - B)$ is then the fraction of light produced that escapes the crystal. Light escaping the crystal will either reach the light detector, or be absorbed by surrounding surfaces. These surfaces are the reflecting foil that surrounds detectors, and the clamps that hold target crystals in place. The magnitude of scintillation light absorption by surrounding surfaces is quantified by a fraction C . This is the fraction of light escaping the crystal not absorbed in the light detector¹. C represents a “loss channel”, where scintillation light energy is completely undetected. $(1 - B)(1 - C)$ is then the *light collection efficiency*: the fraction of light produced that reaches the light detector.

Finally, one must also account for an overall scale factor between light absorbed by the light detector and the practically measured signal. This scale factor is affected by the detector efficiency, as not all energy in light reaching the light detector will cause a measurable signal. Further, one must account for calibration normalization effects. All detectors are calibrated using scintillation light from 122.06 keV γ interactions in the target crystal, not directly to a known amount of energy deposited in the light detector. A scaling factor ϵ is introduced to account for both these effects. Combining all these factors leads to an expression for the signal seen in the light detector, as related to the produced light:

¹ The meaning of C is subtly different from c in [121]. C here is the fraction of light absorbed in the surrounding surfaces relative to the total amount of light escaping the crystal, whereas c in [121] was relative to the total amount of light produced.

$$L_{Det} = \epsilon(1 - B)(1 - C).EK . \quad (6.2)$$

We now turn to the signal seen in the phonon detector. This is related to the sum of the energy that did not contribute to generating scintillation light, $E(1 - K)$, plus any additional energy that went into scintillation light that was later reabsorbed by the crystal, $B.EK$. As with the light detector, the phonon signal model must also account for a finite efficiency and calibration normalization effects. Calibration normalisation effects for the phonon detector are subtle. Calibration normalises phonon detector response to the energy deposited in the phonon detector. However, only the majority of interaction energy goes into the phonon detector, not the total. Calibration has then actually normalized phonon detector response only to the amount of energy in phonons, not the total energy. A scale factor α is then used to model both phonon detector efficiency and calibration normalization. The signal seen in the phonon detector is then:

$$P_{Det} = \alpha(E(1 - K) + B.EK) . \quad (6.3)$$

With models of the detected phonon and light signals now outlined, *anti-correlation* between phonon and light signals can be described. This can be seen if the scintillation efficiency K , the fraction of scintillation photons produced reabsorbed in the crystal B , and the fraction of scintillation photons escaping the crystal absorbed in the surrounding surfaces C , can all vary slightly between interactions of the same energy. This is modelled such that, for a given interaction, $K = \bar{K} + \Delta K$, $B = \bar{B} + \Delta B$ and $C = \bar{C} + \Delta C$, with ΔK , ΔB and ΔC small relative to the expectation values \bar{K} , \bar{B} and \bar{C} . To first order, the resulting anti-correlation gradient is then:

$$\frac{\Delta L_{Det}}{\Delta P_{Det}} = \frac{-\epsilon(1-\bar{C})}{\alpha} + \frac{\epsilon\bar{K}(1-\bar{B})}{\alpha} \left(\frac{\Delta C}{\Delta K(1-B) - \bar{K}\Delta B} \right) . \quad (6.4)$$

Initially, these slight variations are considered in the case of *no* position dependence. The number of photons produced is finite, 28.7 photons per keV in CaWO_4 at milliKelvin temperatures [75]. The scintillation efficiency K , the fraction of produced light not escaping the crystal B , and the fraction of light escaping the crystal absorbed by surrounding surfaces C , are then described by

independent Poisson variables. K , B and C are then uncorrelated, and the second term in 6.4 vanishes when summed over many events. This leaves an average anti-correlation gradient $\frac{-\epsilon}{\alpha}(1 - \bar{C})$. An important result is that this gradient is independent of both energy and interaction position, and should be the same for all interaction types (γ/β , α , neutron). This behaviour was confirmed in measurements of cadmium tungstate [123].

We now consider the effects that a position dependent variation in crystal reabsorption B and surface absorption C might have. Position dependence could mean that changes in B and C could then be *correlated*. Such co-variance would mean that the second term in Eqn. 6.4 no-longer averages to zero; anti-correlation gradient can change, and this change can be related to position. The position dependent case of this model then allows for a variation in the anti-correlation gradient, where the position independent case does not. Further, *if* anti-correlation gradient is related to position, then it will *also* be related to energy. This is because the second term in Eqn. 6.4 will be affected by the relative magnitude of position dependent, correlated changes contributing to ΔB and ΔC , which won't vary with energy, and the magnitudes of position independent changes contributing to ΔB , ΔC and ΔK (Poisson variation), which will.

Separately from correlated changes between C and B , position dependence can also affect the average total light signal through the light collection efficiency term $(1 - B)(1 - C)$ in Eqn. 6.2. This is the case if photons produced in different regions of the crystal have different probabilities of actually reaching the light detector, versus being absorbed in either the crystal or surrounding material. This possibility is predicted in simulations of [121] and supported by measurements in [83].

Before discussing measurements of these effects, previous studies on both position dependence and anti-correlation in CRESST and CRESST-like detectors

are briefly reviewed. Though predicted in the closing chapter of [121], both effects have until now been treated separately in CRESST, and so it is useful to place this work in reference to existing literature.

6.3 Position dependence and anti-correlation studies review

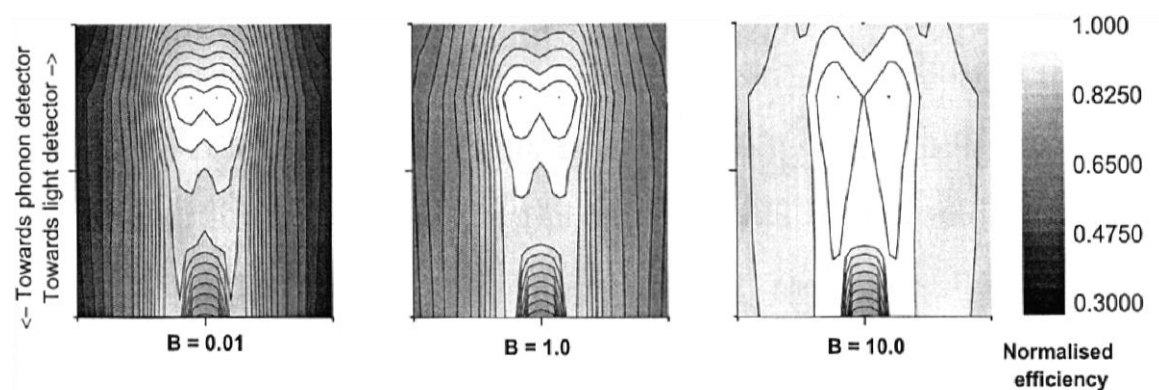


Fig. 6.1 Simulated light collection efficiency with position in polished surface crystals, from [121]. The x-axis represents position along the 4cm diameter of the crystal, the y-axis along the 4cm crystal height. “B” in [121] is the scattering-to-absorption ratio for scintillation photons in a crystal. Higher B values in [121] implied a larger ‘relative chance to photons “survive” an attenuating interaction’. Each different diagram represents a different possible scattering-to-absorption ratio (an unknown) supposed for CaWO_4 . In all cases, radial position dependence is seen, and light collection efficiency drops towards the surface. The small patch of reduced light at the bottom and centre of each diagram is caused by the phonon tungsten thermometer absorbing scintillation light directly. Diagrams correspond to changes in light collection efficiency $(1 - B)(1 - C)$ in this work.

Position dependent effects in detected light were predicted in simulations of CRESST detector response with polished target crystal surfaces [121]. The results are shown in Fig. 6.1. In these simulations, light collection efficiency decreased significantly radially towards the crystal surface. The principle cause of this effect was light-trapping, where scintillation light is totally internally reflected by the cylindrical surfaces of the CaWO_4 crystal, and is then re-absorbed by the crystal.

However, roughening of the top surface of CRESST crystals has been shown to reduce this light-trapping effect significantly [117], and all CRESST crystals are now of this type. In [117], position dependence was examined by looking for a non-Gaussian bulk-behaviour of electron recoils in the light detector. Compared to polished surface crystals, position dependence was reduced such that it was

no-longer observable in surface roughened crystals through this method. However, radial position dependence was first observed in surface roughened CRESST crystals using a different technique in [83]. In [83], a comparison was made in the light yield seen for 122.06 keV γ interactions from Co-57 calibrations at different source positions. [83] made use of the small penetration depth of 122.06 keV γ interactions in CaWO_4 (0.8mm). Interactions with the calibration source directly below the crystal are approximately radially isotropic in the crystal, whereas those with the source radially separated from the crystal will not be. 122.06 keV γ interactions with the calibration source positioned below the crystal produced significantly more detected light on average than those with the calibration source positioned radially with respect to the crystal. A decrease in light towards the crystal edge of $\sim\mathcal{O}(10\%)$ was then inferred. [83] also found a similar reduction in light towards the crystal centre when comparing the behaviour of background interactions to 122.06 keV γ calibration interactions.

Anti-correlation measurements have yet to be considered in terms of position dependence in measurements of real crystals. For solid, inorganic scintillators, either in CRESST or very similar set-ups, the anti-correlation effect itself was examined previously in [109,124] for (surface-roughened) CRESST calcium tungstate crystals, and in [123] for cadmium tungstate crystals in a similar set-up. Fit results of [109] are shown in Fig. 6.2. In these works, scintillation light reabsorption by target crystals and surface losses were not considered. The driving force of anti-correlation was everywhere assumed to be the Poisson variation of the number of scintillation photons produced through K . By comparison, this interpretation leads Eqn. 6.2 - 6.4 to become:

$$L_{Det} = \epsilon KE , \quad (6.5)$$

$$P_{Det} = \alpha E(1 - K), \quad (6.6)$$

$$\frac{\Delta L_{Det}}{\Delta P_{Det}} = \frac{-\epsilon}{\alpha} . \quad (6.7)$$




Fig. 6.2 Anti-correlation measurements of [109]. A bivariate Gaussian distribution was fit to event distribution of the 122.06 keV Co-57 calibration peak in the CRESST commissioning run. Shown are the data (left) and results (right) are the results for fits to 09T in the commissioning run. The anti-correlation gradient was then related to the scintillation constant K assuming no losses or reabsorption.

Comparing Eqn. 6.7 and Eqn. 6.4, the two interpretations lead to very similar expectations for the anti-correlation gradient. Indeed, in the absence of a loss channel, such that $C = \Delta C = 0$, the equations are identical - though the resulting formalism differs between considering the fraction of energy in produced photons K and the fraction of energy leaving the crystal $K(1 - B)$. The important distinguishing feature between these representations is that the anti-correlation gradient is a constant when both losses and position dependence are neglected. Eqn. 6.5 - 6.7 can also be rearranged to give $K = 1 / \left(1 - \frac{\Delta L_{Det}}{\Delta P_{Det}} \frac{P_{Det}}{L_{Det}} \right)$. Thus, in these interpretations, *if* loss channel effects can be neglected, it is possible to reconstruct the scintillation efficiency K by measuring the anti-correlation gradient¹. In [109], scintillation constant K was then reconstructed to be 5.2% and 2.0% for the two CaWO_4 crystals under study, somewhat reduced with respect to 8.3% measured in [75]. Under the assumption that losses were correctly neglected in [109], this reduction could be explained by crystal reabsorption of produced scintillation light.

¹ Neglecting losses in Eqn. 6.2-6.4, one instead gets $K(1 - B) = 1 / \left(1 - \frac{\Delta L_{Det}}{\Delta P_{Det}} \frac{P_{Det}}{L_{Det}} \right)$. Equivalently, one can then reconstruct the fraction of energy escaping the crystal, $K(1 - B)$

Briefly it is noted that variations in crystal re-absorption and surrounding surface absorption are not the only possible origins of position dependence. Position dependent effects would also arise if scintillation efficiency K varied with position. However, measurements of K have shown that such a variation is small, and scintillation efficiency is almost independent of position [121]. Position dependent effects might also occur if the light detector efficiency ϵ was dependent on the position of the absorbed scintillation photons. Measurements of light detector behaviour however do not support this possibility, and position dependence on the light detector itself 'can be neglected' [80].

6.4 Anti-correlation measurement method

To see if there is any position dependence in anti-correlation, measurements of anti-correlation gradients at different positions must be performed. The following details a method to perform such measurements. The principle idea is to make use of the fact that photons of different energy have different penetration depths in CaWO_4 . Photons differing in energy from an external source then act as probes of behaviour at different mean depths within the crystal. The method used to measure anti-correlation gradients is itself broadly derived from similar approaches in [109] and [123]. In both these works, it was assumed that calibration peaks could be modelled by two independent Gaussian variables in a frame rotated with respect to the phonon-light signal frame. In this work, calibration energy distortions and background effects are also accounted for.

Cobalt-57 and Thorium-232 radioactive sources, introduced in the most recent CRESST run, produce γ rays at a variety of energies up to several MeV. However, above ~ 1 MeV, systematic errors on energy estimate become significant, effectively providing an upper bound in energy for this work. Also, below 100 keV, γ rays are only weakly penetrating. Sources are limited to locations outside the copper cold-box (Fig. 2.1), and source γ rays much below 100 keV do not

readily penetrate this copper cold-box. Of interest to this work then are γ energies ranging from the 122.06 keV emission of Co-57, up to the 911.1 keV emission of Ac-228 in the Th-232 decay chain.

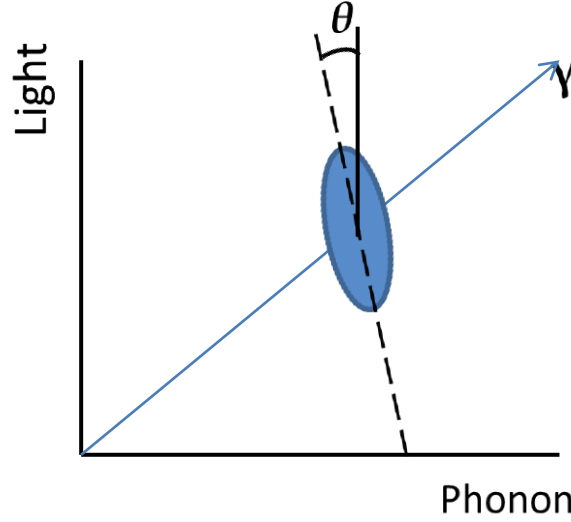


Fig. 6.3 Diagram of anti-correlation effect. The blue shaded region represents the probability distribution of detector response for events from γ rays of identical energy. The x - y basis represents the original phonon and light detectors measurements. The dashed lines represents a direction of the basis in which the event distribution may be modelled by two independent variables.

First, the simple case where all γ interactions are of equal energy is examined. The probability distribution of events can then be modelled as two, independent Gaussian variables, in the orthogonal variables x' and y' :

$$\rho(x', y') = \frac{1}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x'-\bar{x}')^2}{2\sigma_{x'}^2}} e^{-\frac{(y'-\bar{y}')^2}{2\sigma_{y'}^2}}. \quad (6.8)$$

These variables have co-ordinates transformed through a unitary transformation S in angle ϑ with respect to the phonon signal (x) and light signal (y). ϑ is then related to the anti-correlation gradient by $\frac{\Delta L_{Det}}{\Delta P_{Det}} = -\tan(\vartheta)$. Anti-correlation gradients can then be measured by transforming coordinate systems and fitting the distribution of Eqn. 6.8, with the transformation angle between axes a free parameter. The mean measurements of each independent variable \bar{x}' and \bar{y}' are then related to the mean reading in the phonon (\bar{x}) and light (\bar{y}) detectors by:

$$\begin{pmatrix} \bar{x}' \\ \bar{y}' \end{pmatrix} = S \begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix} = \begin{pmatrix} \cos(\vartheta) & -\sin(\vartheta) \\ \sin(\vartheta) & \cos(\vartheta) \end{pmatrix} \begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix}, \quad (6.9)$$

\bar{x} is then the mean phonon detector reading, related to interaction energy by:

$$\bar{x} = \frac{1 - K_\gamma(E)(1 - B)}{1 - K_{cal}(1 - B_{cal})} T_x E, \quad (6.10)$$

where possible phonon systematic error has been introduced with the variable T_x . In this equation, only \bar{x} is measurable directly, with E known. To discuss the meaning of other variables, an approximation must be made. It is noted that as $K_\gamma \approx 0.083$, it is quite small. $K_\gamma(E)(1 - B)$, the amount of energy in scintillation light that escapes the crystal, will then also be small¹. The approximation is made that $\bar{x} \approx T_x E$. $|T_x - 1|$ then estimates the systematic error of the phonon detector as a function of energy, due to un-corrected detector transition non-linearity.

Considering the average signal seen in the light detector, Eqn. 6.2 leads to:

$$\bar{y} = \frac{(1 - B)(1 - C)}{(1 - B_{cal})(1 - C_{cal})} \frac{K_\gamma(E)}{K_{cal}} T_y E, \quad (6.11)$$

where T_y is the systematic energy error on light signal, and $\frac{K_\gamma(E)}{K_{cal}}$ represents possible non-linear energy dependence in the scintillation efficiency (from Eqn. 6.1). As with the phonon signal, only \bar{y} is measurable, with E known. Interpreting measurements in terms of Eqn. 6.11 then poses even more of a challenge, as changes in collection efficiency $(1 - B)(1 - C)$ are *not* necessarily small. To proceed, the following assumptions are made:

- 1) $K_\gamma(E)$ is constant with respect to energy, $K_\gamma(E) \equiv K$.
- 2) Light detector systematic T_y errors are at worst as bad as the systematic errors T_x in the worst phonon detector.

Assumption 1) implies that all crystals produce the same amount of light with respect to incident energy. This assumption is supported by measurements in [89], which limit the variation in light production of CaWO_4 with energy to a couple of percent above 100 keV. Relaxation of this assumption will be examined in section 6.4.3.

¹ It is later shown that the light collection efficiency $(1 - B)(1 - C)$ varies by at most $\sim 18\%$ in the crystal showing the largest variation. Though it is not possible to disentangle this variation between changes of crystal absorption B , and surrounding material absorption C , assuming this entire variation is in $(1 - B)$ means $1 - K(1 - B)$ varies by at most $\sim 1.5\%$. This then compares to a maximum systematic error assumed in the phonon detector $|T_x - 1|$ of $\sim 4.2\%$. The assumption that systematic error dominates then appears reasonable.

The justification for assumption 2) is less clear. Detector systematic error occurs primarily from transition non-linearity. If one allows for a possible position dependence in light collection efficiency, these errors are only measurable on the phonon detector using $\bar{x} \approx T_x E$. To justify this assumption, the non-linear energy corrections introduced by the CPE process (section 3.4) are considered. These corrections were designed to remove most of the effects of this transition non-linearity, which can be quite large for the phonon detector but are rarely so for the light detector. Any remaining systematic errors might then well be expected to be larger for the phonon detector than the light detector. An over-estimate of possible systematic error would yield conservative results. Thus light detector systematic error with energy is modelled by behaviour of phonon detector showing the largest systematic errors at a given energy.

With these assumptions, measurements of \bar{y}/E can then be related to $\frac{(1-B)(1-C)}{(1-B_{cal})(1-C_{cal})}$. This is the fraction of light reaching the light detector at a given energy, relative to the fraction seen in calibration - the *relative light collection efficiency*. \bar{y} is practically measured as $\bar{y} = M_S E$. Ultimately, M_S will be composed of contributions from systematic error, possible non-linear behaviour of emitted light, and finally possible position dependent variation in light collection efficiency $(1 - B)(1 - C)$. Anti-correlation gradient measurements here have then naturally led to measurements of changes in light collection efficiency. However, this treatment has only described the simple case, where the probability distribution of events can be modelled by a bivariate Gaussian distribution. Real measurements must also account for background and source distortions (e.g. Compton scattering). Methods to account for these effects are outlined below. All fits in this chapter use the Maximum Likelihood method with the ROOFIT package [125].

6.4.1 Thorium-232 calibration backgrounds

Decays in the Thorium chain, from Th-232 down to a final, stable daughter Pb-208, produce a variety of γ rays that can penetrate the copper cold box. The most prominent of these, as seen in CRESST modules, are indicated in Fig. 6.4.




Fig. 6.4 ^{232}Th calibration spectrum as seen in the phonon detector of module 10B. The energies that give rise to each major peak in the spectra are noted, in keV, along with the decay(s) that principally contributes.

Here, only strong γ lines, relatively isolated from other possible lines are chosen for study. These choices are indicated in Table 6.1. Even after selection, most chosen lines have small, additional contributions from other possible γ emission branches in the Th-232 decay chain. A full list of all decay energies in this decay chain may be found in [103]. These small contributions are also allowed in fitting, but in general negligibly affect fits. Table 6.1 also details the energies at which these small contributions are expected to appear. Th-232 calibration sources then provide γ rays energies used in this work between 238.63 keV and 911.21 keV.

During calibration, β interactions and Compton scattered γ rays from higher energy decays are seen close in energy to those of interest here. These lead to a broad background contribution at energies above and below the lines of interest, exemplified in Fig. 6.5.

| Principal energy (keV) | Source | Secondary energy (keV) | Source |
|------------------------|-----------------------|------------------------|------------------------|
| 238.63 | Pb-212, β decay | 240.97 | Ra-224, α decay |
| 338.32 | Ac-228, β decay | 340.97 | Ac-228, β decay |
| 510.77 | Tl-208, β decay | 508.96 | Ac-228, β decay |
| 583.19 | Tl-208, β decay | 583.41 | Ac-228, β decay |
| 911.21 | Ac-228, β decay | | |

Table 6.1 Emission energies of the Th-232 source that are examined in this work. The principal energy is of the γ ray causing the majority of events in a fit region, the secondary energy the next largest γ contributor.




Fig. 6.5 Density plot of events from the 238.63 keV γ line in module 10B, principally from Pb-212 β decay. On the left is the observed event probability distribution, on the right the best fit using the outlined method. Of note: the small excess above (in energy) the central distribution, from 240.97 keV γ rays from Ra-224 β decay; the broad background at all energies; a systematic error of 0.4% in the phonon detector. Binning is by 0.5 keV in the phonon detector and 2 keV in the light detector. Fits themselves are all un-binned.

This broad background can form a significant contaminant to fits near Th-232 peaks examined here. A model must then be built to account for these interactions. Here, this background is assumed to represent an approximately flat contribution to event density over the narrow fit energy range. The contribution to the event density distribution is then modelled as:

$$\rho_{Back}(x', y') = A_B \int_0^\infty \frac{1}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x' - \bar{x}'_B(E))^2}{2\sigma_{x'}^2}} e^{-\frac{(y' - \bar{y}'_B(E))^2}{2\sigma_{y'}^2}} dE, \quad (6.12)$$

for background contribution fraction A_B . Eqn. 6.12 can be integrated with an analytic solution. This integration is performed in Appendix C, with the result:

$$\frac{A_B}{\sqrt{2\pi}\sigma_T} e^{-\frac{(x's_{y'} - y's_{x'})^2}{2\sigma_T^2}} \frac{1}{2} \operatorname{erfc}\left(\frac{-p'_B\sigma_T}{\sqrt{2}\sigma_x\sigma_y}\right). \quad (6.13)$$

Where $p'_B = \frac{(\sigma_{y'}^2 x' s_{x'} + \sigma_{x'}^2 y' s_{y'})}{\sigma_T^2}$. This background distribution is then summed with the expected source probability distribution. Two possible backgrounds are allowed, for surface and internal interactions, each allowed a different mean light signal per unit energy. The effect of Compton scattered photons from the peak of interest contaminating the data set (see next section) is minimal at these energies.

6.4.2 Cobalt-57 calibration backgrounds




Fig. 6.6 Density plot for the 122.06 keV Co-57 calibration line in the module 06B. The left figure shows the observed distribution, the right the best fit. A clear “tail” of Compton scattered photons, with energies below 122.06 keV, is evident. Binning is by 0.5 keV in the phonon and 5 keV in the light detector.

Co-57 undergoes β decay to Fe-57, mostly emitting a 122.06 keV or 136.47 keV γ ray¹. At these energies, attenuation by the copper cold box between source and detector is significant. The average penetration depth of 122.06 keV γ rays is ~ 3 mm in copper. This leads to a factor ~ 50 attenuation by the minimum 12 mm of copper separating source and detector. Importantly, this copper also causes a significant Compton scattered photon background to be seen in Co-57 calibration data, from the two energies of interest. This is the origin of the “tail” of events extending down in energy from 122.06 keV, seen in Fig. 6.6 (and Fig. 6.2). To account for these Compton scattered photons, a similar procedure to modelling

¹ Other lines seen in Co-57 calibrations are seen too weakly in detectors to be of use here.

the background in the Th-232 calibration is followed. Compton scattered photons are phenomenologically modelled by an exponential decay, such that:

$$\rho_{Comp}(x', y') = A_C \int_0^{E_{Cal}} \frac{N e^{-\frac{(E_{Cal}-E)}{\lambda}}}{\lambda} \frac{1}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x'-\bar{x}'(E))^2}{2\sigma_{x'}^2}} e^{-\frac{(y'-\bar{y}'(E))^2}{2\sigma_{y'}^2}} dE, \quad (6.14)$$

where N is a normalization factor ~ 1 and A_C the contribution of Compton scattered photons to the overall signal. The onset of the exponential decay is then the energy of the interaction leading to the Compton scattered photons, extending down in energy. The free parameter λ represents the decay constant of the empirical Compton scattered photon model. The integral also has an analytic solution (Appendix C), with the result:

$$\frac{A_C}{\sqrt{2\pi}\sigma_T} e^{-\frac{(x'_{y'}-y'_{x'})^2}{2\sigma_T^2}} \frac{N}{\lambda} e^{-\frac{(E_{Cal}\sigma_T^2 - (\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'}))}{\lambda\sigma_T^2}} \frac{\sigma_{y'}^2 \sigma_{x'}^2}{e^{2\sigma_T^2 \lambda^2}} \frac{\operatorname{erfc}\left(\frac{-(E_{Cal}-p'_C)\sigma_T}{\sqrt{2}\sigma_x\sigma_y}\right) - \operatorname{erfc}\left(\frac{p'_C\sigma_T}{\sqrt{2}\sigma_x\sigma_y}\right)}{2} dE. \quad (6.15)$$

Where $p'_C = \frac{(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'} + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda})}{\sigma_T^2}$. The full model of Co-57 calibration event distributions then includes the bivariate Gaussian distribution in x' and y' for each γ line, the empirical Compton scattered photon model for the calibration emission line at 122.06 keV and 136.47 keV, and a constant background component representing background γ/β rays.

6.4.3 Mean penetration depth estimate

Photon penetration lengths in elemental matter are well known, allowing accurate calculation of the penetration lengths in CaWO_4 , shown in Fig. 6.7. However, geometric effects must be accounted for to relate these lengths to interaction positions in the target crystals. Ideally, these geometric effects would be taken into account with a Monte-Carlo study. However no such study was available during this work. Thus, first order estimates are made using an analytic approach.

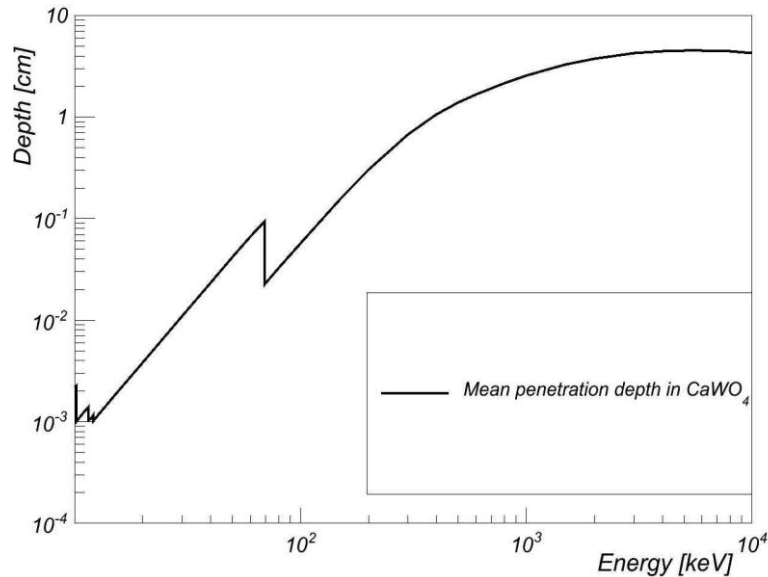


Fig. 6.7 Penetration length in CaWO_4 of photons, derived from [126]. This line is calculated from the known attenuation of gamma rays in the elemental constituents of CaWO_4 . This diagram then represents the mean interaction depth for photons in an infinite, planar detector exposed to an infinite, planar source. The sharp drops are at X-ray absorption edges, for example the K-edge of tungsten at 69.5 keV.

Target crystals measure 4cm in diameter and height. γ penetration lengths of interest here then range from 0.8mm at 122.06 keV, to 2.4cm at 911.2 keV. Calibration sources are placed approximately radially with respect to the crystal cylindrical axis during data taking. Given that radial position dependence is expected, it is then necessary to approximately relate these penetration lengths to mean radial interaction depths in the crystal. Difficulty exists in that during calibration, source position, relative to a given detector, was not recorded. Source position must then be estimated, providing a source of systematic error.

First order geometrical correction factors, neglecting the effects of blocking by other detectors, are then estimated (Appendix D). These calculations give rise to the mean radial depth interaction estimates in Fig. 6.8. A note of caution in that the radial depths in Fig. 6.8 are of the *mean* interaction position, not the depth of each *individual* interaction, and have neglected Compton scattering. The point of calculating these mean depths is simply to show that different energy calibrations probe radial length scales of different magnitudes. 122.06 keV γ rays from Co-57 sample behaviour approximately in the first millimetre from the cylindrical

surface, whereas 911.2 keV γ rays from Th-232 the bulk behaviour throughout the entire crystal. Thus, the two measurements at different energies are referred to here as *surface* and *bulk* measurements.




Fig. 6.8 Mean γ calibration radial interaction depth estimates. The penetration lengths of Fig. 6.7, (black) are used to estimate the mean interaction depth towards the azimuthal axis from the cylindrical crystal surface (cyan). This accounts for geometric factors of detector size and position, but neglects blocking (Appendix D). Vertical errors are not statistical, but represent source position uncertainty. At high energy, penetration depth is longer than the crystal radius, and interactions are approximately uniform in the crystal. Mean interaction position is then close to the average radius, $\frac{2}{3} R$, giving average depth $\sim R - \frac{2}{3} R = 0.67\text{cm}$ (horizontal dotted line). At low energies, the dominant error is caused by the small fraction of γ interactions that, due to vertical position error, might enter the crystal through top or bottom surfaces.

6.5 Measurement

The modules chosen for measurement are 01T, 04B, 06B, 06T, 08M, 09M and 10B. Module 04B has two light detectors, “Q” and “Burkhard” (labelled “Q” and “B” respectively). Anti-correlation in module 04B is considered between the phonon detector and each light detector separately. Other working modules not used here are: 08B, the active ZnWO_4 module, which had poor phonon detector resolution making separation of γ peaks difficult; 09B, which did not have an active light detector during the Co-57 calibration; and 10T, which also had poor phonon detector resolution. When measuring anti-correlation gradient, fitting is performed close to the position of each γ line. Fit regions are typically $\pm 1\text{-}3\%$ in

the phonon detector from the peak maxima, varied to account for phonon detector resolution. Results are stable within statistical error with respect to small changes in fit region choice.

6.5.1 Cobalt-57 Calibration

Co-57 calibration data was taken at the start of the most recent run, between the 18th and 25th of June, 2009. The Co-57 source was introduced to set the energy-amplitude relations for phonon and light detectors. With the source at a given position, many detectors are not adequately calibrated. Simply, it is possible that not enough γ rays reach a given detector from the source, due to distance or possible obstruction by other detectors. The source was then moved on the 23rd of June from position RBR to position BGB to ensure that all detectors were adequately calibrated. Movement of the source provides the first possible chance to see position dependent effects. If anti-correlation gradient differs between source positions, this would be direct evidence for position dependence. However, position dependence is expected to be principally a radial effect. As source movement was approximately in the crystal radial plane, this will not strongly affect average radial penetration depth, with maximal change corresponding to the vertical systematic error in Fig. 6.8. Thus, little change in anti-correlation gradient is expected.

However, tests are still performed to check for this small change. The calibration data set is split in two, between periods at each source position. Anti-correlation gradient measurements are then taken for each, shown in Table 6.2. A χ^2 significance test is then performed. The null-hypothesis of this test is that anti-correlation gradient is independent of source position. The question is posed: is each pair of individual detector measurements consistent with a single anti-correlation gradient, $\bar{\vartheta}_i = \left(\frac{\vartheta_{i,RBR}}{\sigma_{i,RBR}^2} + \frac{\vartheta_{i,BGB}}{\sigma_{i,BGB}^2} \right) / \left(\frac{1}{\sigma_{i,RBR}^2} + \frac{1}{\sigma_{i,BGB}^2} \right)$? Six detector pairs have readings at both positions, allowing six independent tests of this hypothesis.

Individual χ^2 tests are then combined as a total, global χ^2 to give a single confidence probability. The total χ^2 value is 10.47 with six degrees of freedom, implying results are consistent with the null-hypothesis with 10.6% confidence.

| Energy (keV) | Position | Anti-correlation angle ($\times 10^{-2}$ radians $\pm(\text{stat}) \pm(\text{sys})$) | | | |
|--------------|----------|--|----------------------------|----------------------------|----------------------------|
| | | 01T | 04B (Q) | 04B (B) | 06B |
| 122.06 | RBR | 6.51 \pm 0.24 \pm 0.02 | 4.53 \pm 0.29 \pm 0.01 | 4.84 \pm 0.33 \pm 0.01 | 4.99 \pm 0.29 \pm 0.01 |
| 122.06 | BGB | 5.24 \pm 0.52 \pm 0.02 | 4.88 \pm 0.29 \pm 0.01 | 5.11 \pm 0.34 \pm 0.01 | 7.44 \pm 1.68 \pm 0.02 |
| 136.47 | RBR | 5.44 \pm 0.54 \pm 0.02 | 5.03 \pm 0.55 \pm 0.02 | 4.98 \pm 0.55 \pm 0.02 | |
| 136.47 | BGB | 5.63 \pm 0.86 \pm 0.02 | 4.43 \pm 0.71 \pm 0.01 | 4.96 \pm 0.69 \pm 0.02 | 4.65 \pm 0.49 \pm 0.01 |

| Energy (keV) | Position | Anti-correlation angle ($\times 10^{-2}$ radians $\pm(\text{stat}) \pm(\text{sys})$) | | | |
|--------------|----------|--|----------------------------|----------------------------|----------------------------|
| | | 06T | 08M | 09M | 10B |
| 122.06 | RBR | 5.44 \pm 0.54 \pm 0.02 | 5.63 \pm 0.48 \pm 0.02 | | |
| 122.06 | BGB | 4.60 \pm 0.35 \pm 0.01 | 4.77 \pm 0.83 \pm 0.01 | 4.18 \pm 0.13 \pm 0.01 | 8.79 \pm 0.14 \pm 0.03 |
| 136.47 | RBR | 4.15 \pm 0.67 \pm 0.01 | 5.28 \pm 0.81 \pm 0.02 | | |
| 136.47 | BGB | 5.63 \pm 0.86 \pm 0.02 | | 4.14 \pm 0.24 \pm 0.01 | 8.92 \pm 0.29 \pm 0.03 |

Table 6.2 Anti-correlation gradients at 122.06 keV and 136.47 keV. RBR and BGB indicate source position. Gaps exist where source-strength was too weak with respect to background. The 122.06:136.47 keV γ production ratio is 100:12 for Co-57 before attenuation. As penetration depths are similar, 136.47 keV measurements are then more easily overwhelmed by background. These measurements have been corrected for the measured systematic error in each phonon detector T_x , though this change is $< 0.3\%$ at these energies. Additional possible systematic error in the light detector, estimated from behaviour of the worst phonon detector, is also noted. Systematic error is much less than statistical everywhere.

Movement of the calibration source has then had only a weakly significant effect on anti-correlation gradient using only the 122.06 keV Co-57 γ measurements. Measurements of the 136.47 keV line have larger associated errors. Position dependence effects on anti-correlation can then not be shown using Co-57 data alone. Data from the Th-232 calibration must then also be considered.

6.5.2 Thorium-232 Calibration

A Th-232 calibration source was introduced at the end of the most recent CRESST run, between the 1st and 15th of March, 2011. Anti-correlation gradients are then measured at each of the energies noted in Table 6.1, with results shown in Table 6.3. A slight complication was that one module, 09M, showed highly non-linear light-detector behaviour. This was shown as the fit RMS (see section 4.2.10) became highly energy (amplitude) dependent. In such situations, standard pulse

fits fail to accurately assess detector amplitude. Anti-correlation measurements could thus not be taken at high energy.

| Energy (keV) | Anti-correlation angle ($\times 10^{-2}$ radians $\pm(\text{stat}) \pm(\text{sys})$) | | | |
|-----------------|--|----------------------------|----------------------------|----------------------------|
| | 01T | 04B (Q) | 04B (B) | 06B |
| 238.63 | 6.01 \pm 1.12 \pm 0.06 | 5.64 \pm 0.88 \pm 0.06 | 5.35 \pm 0.81 \pm 0.06 | 4.14 \pm 0.74 \pm 0.04 |
| 338.32 | | | | 5.28 \pm 1.18 \pm 0.04 |
| 510.77 | | 4.53 \pm 1.06 \pm 0.10 | 5.29 \pm 0.60 \pm 0.12 | 3.01 \pm 0.83 \pm 0.07 |
| 583.19 | 5.98 \pm 0.99 \pm 0.16 | 4.49 \pm 0.52 \pm 0.12 | 4.25 \pm 0.30 \pm 0.11 | 3.91 \pm 0.48 \pm 0.10 |
| 911.21 | 6.45 \pm 1.59 \pm 0.27 | 5.25 \pm 0.39 \pm 0.22 | 4.22 \pm 0.32 \pm 0.18 | 3.48 \pm 0.69 \pm 0.15 |

| Energy (keV) | Anti-correlation angle ($\times 10^{-2}$ radians $\pm(\text{stat}) \pm(\text{sys})$) | | | |
|-----------------|--|----------------------------|----------------------------|----------------------------|
| | 06T | 08M | 09M | 10B |
| 238.63 | 3.92 \pm 0.29 \pm 0.04 | 6.36 \pm 2.22 \pm 0.07 | 3.72 \pm 0.16 \pm 0.04 | 7.91 \pm 0.42 \pm 0.08 |
| 338.32 | 4.30 \pm 0.57 \pm 0.04 | | | 7.43 \pm 0.42 \pm 0.08 |
| 510.77 | 4.69 \pm 0.31 \pm 0.10 | | | 5.62 \pm 0.50 \pm 0.12 |
| 583.19 | 4.40 \pm 0.15 \pm 0.12 | | | 5.86 \pm 0.23 \pm 0.16 |
| 911.21 | 4.26 \pm 0.18 \pm 0.18 | 4.22 \pm 0.92 \pm 0.18 | | 5.06 \pm 0.27 \pm 0.22 |

Table 6.3 Anti-correlation gradients for various energies in the Th-232 calibration. Gaps exist where source-strength was too weak to make measurements. Light detector response in module 09M became non-linear at relatively low energies, and higher energy measurements were not taken (see text). All measurements are corrected for the measured systematic error in each phonon detector T_x . Additional systematic error in the light detector is estimated from the worst phonon detector systematic error.

A test is performed to see if these anti-correlation gradients are constant across energy (and thus, penetration depth). The null hypothesis of this test is that anti-correlation gradient is independent of energy and penetration depth in each detector. Measurements for an individual module, from both Co-57 and Th-232 calibrations, are then compared to average value for that module $\bar{\vartheta}_i = \frac{\sum \frac{\vartheta_{i,E}}{\sigma_{i,E}^2}}{\sum \frac{1}{\sigma_{i,E}^2}}$, across all energies. The resulting χ^2 values are shown in Table 6.4.

It is clear that these measurements are not compatible with the null-hypothesis, considering all modules. However one module in particular, 10B, has dominated total results. This module should then be examined further, to see whether this result is expected or anomalous.

| Module | χ^2 | Degrees of freedom | Confidence probability |
|---------|----------|--------------------|------------------------|
| 01T | 7.43 | 6 | 28.2% |
| 04B (Q) | 3.76 | 7 | 80.7% |
| 04B (B) | 7.65 | 7 | 36.5% |
| 06B | 13.41 | 7 | 6.3% |
| 06T | 9.93 | 8 | 27.0% |
| 08M | 2.48 | 4 | 64.7% |
| 09M | 5.14 | 2 | 7.6% |
| 10B | 189.7 | 6 | < 0.1% |
| Global | 239.5 | 47 | < 0.1% |

Table 6.4 χ^2 of un-changing anti-correlation gradient assumption. The χ^2 is taken with respect to the mean of all results for a single detector. “Degrees of freedom” represents the number of measurements for a given module, minus one for the mean estimate. Each module represents an independent test of the null-hypothesis, and results are combined to give a global χ^2 test. Module 10B, dominates total results.

Two effects should be noted. First, module 10B shows superior light detector resolution at all energies (see section 6.4.4). This is likely a result of the phonon detector in 10B being a “glued” detector. As noted earlier in this work, the majority of phonon detectors in CRESST have their phonon thermometer directly evaporated onto their surface. Unfortunately, the evaporation process tends to degrade the scintillation characteristics of the crystal, and less light is seen from a crystal after it has been subject to thermometer evaporation. To avoid this problem, the phonon thermometer is first evaporated onto a small carrier crystal, which is then glued onto the cylindrical crystal. It should be noted that this modification is also true of module 06B. The second effect is that this module was well calibrated. Module 10B saw several tens of thousands of events in both Co-57 and Th-232 calibrations. This implies anti-correlation gradient will show smaller statistical errors for this module when compared to other modules. Both these factors lead to the impression that module 10B should show reduced errors over measurements from other modules.

The combined Co-57 and Th-232 anti-correlation gradient measurements for module 10B are shown in Fig. 6.9. Anti-correlation gradient magnitude decreases with increasing energy, and thus penetration depth. This could then be

understood in terms of Eqn. 6.4. At small penetration depths \sim one millimetre and low energies, changes in the fraction of light escaping the crystal that is absorbed in surrounding surfaces C could be highly correlated with changes in crystal re-absorption B , with respect to Poisson error. Across the bulk of the crystal at high energy, this correlation is less pronounced, reducing anti-correlation gradient magnitude. Alternatively, changes in C and B could be highly anti-correlated in the bulk at high energy (with respect to Poisson error), and less so on the surface at low energy. In short, changes in C and B might be more correlated (or less anti-correlated) near the surface at low energy than in the bulk at high energy, with respect to Poisson error.

While there does appear to be evidence here for position dependent effects in anti-correlation gradient, the result is unsatisfactory. Only one detector module showed clear evidence for this effect, 10B. Excluding module 10B, results would have been consistent with no position dependence effects. We then look now at position dependence effects in terms of light collection efficiency.




Fig. 6.9 Anti-correlation gradient of module 10B against energy, from Table 6.2 and Table 6.3. A reduction in anti-correlation gradient magnitude is observed with increasing energy. Statistical and estimated systematic errors have been summed in quadrature. Labels correspond to the average radial penetration depths implied from section 6.3.4.

6.5.3 Collection efficiency

In section 6.2, light detector collection efficiency was modelled to be the product of the fraction of light escaping the crystal, $(1 - B)$, and the fraction of that escaped light not absorbed in the surrounding surfaces $(1 - C)$. Mean light detector response to given energy interactions was then modelled in Eqn. 6.11. This model accounted for relative changes in collection efficiency, $\frac{(1-B)(1-C)}{(1-B_{Cal})(1-C_{Cal})}$, systematic error in the light detector T_γ , and possible changes in produced light $\frac{K_\gamma(E)}{K_{Cal}}$. Measured light signal divided by energy \bar{y}/E against energy is shown in Fig 6.10 for all light detectors.




Fig. 6.10 Variations in the light signal divided by the energy in each light detector, as a function of energy. The errors shown are the squared sum of both statistical and systematic errors. Statistical errors on the mean light yield are very small, and estimated systematic errors dominate everywhere.

From Fig. 6.10 it can be seen the calibration normalization of one unit of light signal per unit energy at 122.06 keV well represents all mean light yield readings. Most detectors are better than 1% calibrated, though almost all mean light signals are higher than normalization. At first glance one would have suspected precisely zero error in the normalization process. However, the calibration method, outlined in section 3.4, does not account for Compton scattered photons. These scattered photons will cause the mean energy of γ rays used for calibration to be slightly less than the expected 122.06 keV. Accounting for Compton

scattered photons here, slightly larger energies than calibration are measured. Only module 06T has worse than 1% error at $1.4 \pm 0.2\%$. This detector module appears to show the most position dependence in terms of light collection efficiency (below), which could explain this additional discrepancy.

The assumption that $K(E) \equiv K$ from section 6.3.1 is now used to estimate the magnitude of position dependence effects. First, the mean light signal measurements of Fig. 6.10 are corrected for the noted small discrepancy from normalization at 122.06 keV. These measurements are then plotted against the mean interaction depth estimates from section 6.3.4 in Fig. 6.8. Fig. 6.11 is then a plot of relative collection efficiency, $(1 - B)(1 - C)/(1 - B_{cal})(1 - C_{cal})$ against mean radial interaction depth.

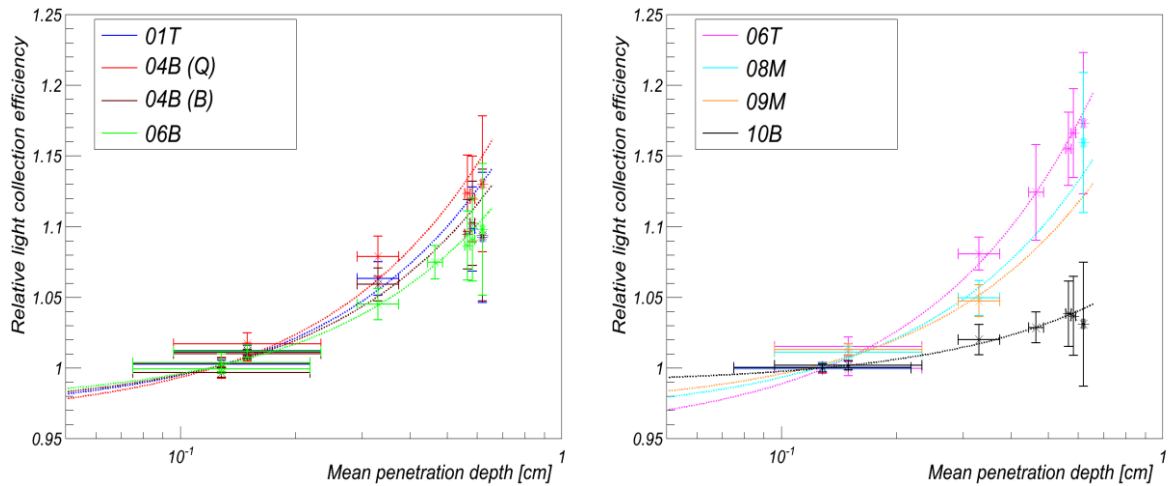


Fig. 6.11 Variations in light collection efficiency as a function of mean radial interaction depth, under the assumption $K(E) \equiv K$. x-axis values are estimates of mean radial interaction depth from section 6.3.4. Relative light collection efficiencies are corrected from Fig. 6.10 by the small discrepancy from calibration at 122.06 keV. Trend lines, linear in mean penetration depth, have been added as guides for the eye.

Under the assumption of linearity in light production $K(E) \equiv K$, light detector efficiency varies considerably with mean penetration depth across nearly all detectors. An increase in observed light is seen towards the azimuthal axis of nearly every crystal; collection efficiency is better across the bulk of the crystal than near the edge. This behaviour broadly matches the expected behaviour from simulations in [121] and measurements in [83].

This positional variation is not identical across all crystals. Minimal variation is seen in module 10B¹, in which a $3.1\pm 4.4\%$ variation between surface behaviour (122.06 keV measurements) and bulk behaviour (911.21 keV measurements) is observed. Maximal variation is seen on module 06T, which sees a $17.3\pm 5.1\%$ variation between surface and bulk behaviour. The optical environment surrounding each crystal is identical. Thus this variation between crystals is conjectured to be caused by a variation in optical transport properties within crystals. Scattering of optical photons in CaWO₄ is expected to be dominated by impurities and lattice deformations [121], and variations in these characteristics could vary between crystals.

However, before committing to this explanation, could this variation be explained by a simple non-linear variation in scintillation efficiency, without position dependence? To answer this question, the assumption in 6.3.1 of constant scintillation efficiency $K_\gamma(E) = K$ is examined. All crystals under examination are of the same chemical makeup. Thus, if there is a change in scintillation efficiency with energy, all crystals should show the same variation. Limits on scintillation efficiency changes with energy can then be derived from the module showing the least variation in received light, module 10B. Despite evidence in section 6.4.2 that 10B is affected by position dependence, limits on changes in K with energy are estimated assuming *no* position dependence in this module. Scintillation efficiency is then modelled noting the Birk's function trend of slightly increasing light with energy in CaWO₄ [127]. The phenomenological model of $K_\gamma(E) = \frac{K_1 E}{1 + K_2 e^{-EK_3}}$, presented in [111] for calcium tungstate, is then applied to the measurements of module 10B shown in Fig. 6.10. Applying this model sets a maximum possible variation in scintillation efficiency between

¹ That this detector sees the least variation in light collection is counter-intuitive, given the results in section 6.4.2, showing position dependence most clearly in module 10B. However, measurements of 6.4.2 were sensitive to *correlated changes* in parameters B and C . Measurements here relate to changes in total light collection efficiency $(1 - B)(1 - C)$, which is sensitive to correlated changes in B and C only to second order.

122.06 keV and 911.21 keV at $4.0 \pm 0.4\%$. Minimum variations in light collection efficiency using this model are then constructed for all other modules in Fig. 6.12.

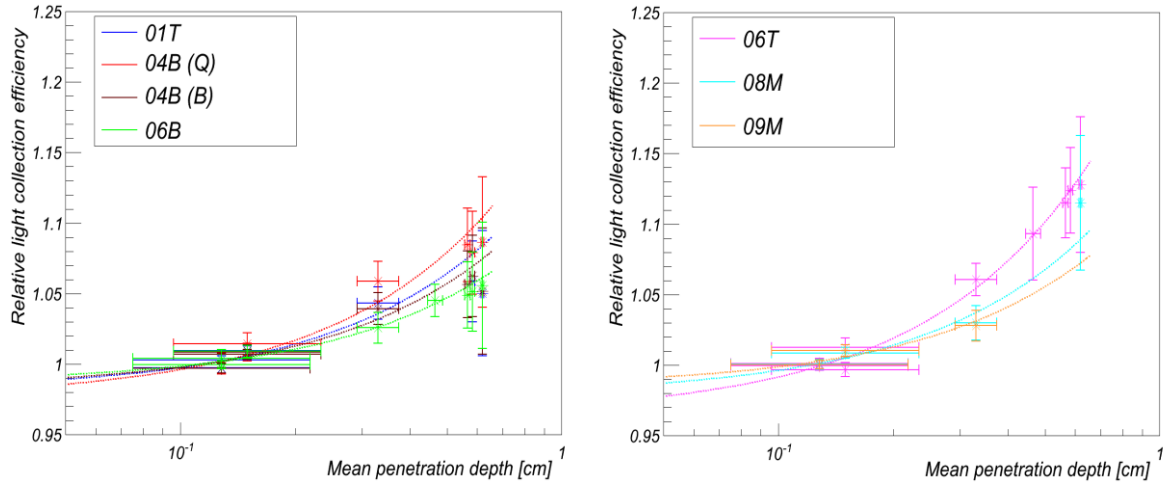


Fig. 6.12 Relative light collection efficiency, against mean radial interaction depth estimates. This figure is related to Fig. 6.11, except that the behaviour of $K(E)$ is allowed to be maximally non-linear, limited by measurements of module 10B. Again, trend lines, linear in mean penetration depth, have been added.

From Fig. 6.12, even assuming maximally non-linear behaviour in K , light collection efficiency still appears to increase from the edge towards the bulk of all remaining crystals. The strongest effect is still in module 06T. Allowing for maximally non-linear behaviour in K , this module then sees $12.8 \pm 4.8\%$ more light for bulk interactions compared to surface interactions, again using 911.2 keV and 122.06 keV measurements. While estimates of the magnitude of light collection efficiency changes have reduced over the constant scintillation efficiency assumption, they are still significant. The original assumption that scintillation efficiency is constant is returned to for the next section.

6.5.4 Light detector resolutions

Finally, light detector resolution, as related to position dependence, is considered. Light detector resolution may be calculated from fit results. The resolution estimates in the x' and y' directions are modelled as independent. Thus, projections of their errors will add in quadrature, and light detector resolution can then be found by:

$$\sigma_y = \sqrt{\cos^2(\vartheta)\sigma_{y'}^2 + \sin^2(\vartheta)\sigma_{x'}^2}. \quad (6.14)$$

This light detector resolution is then divided by the mean detected light output, giving σ_y/\bar{y} , or the 1σ resolution as a fraction of signal magnitude. Errors on σ_y are found through Monte-Carlo, knowing the correlated errors on ϑ , $\sigma_{x'}$ and $\sigma_{y'}$. Position dependence measurements are then estimated using the values of $(1 - B)(1 - C)/(1 - B_{\text{Cal}})(1 - C_{\text{Cal}})$ at 911.21 keV from Fig. 6.11. This resolution in the light detector at 911.21 keV is then compared to these position dependence magnitudes in Fig. 6.13.




Fig. 6.13 Light detector resolution as related to position dependence magnitudes from Fig. 6.11. Some correlation is seen between the light detector resolution at high energy and position dependence magnitude. “04B (Q)” and “04B (B)”, are the detectors on the two light detector module.

Fig 6.13 is then a comparison of the surface-to-bulk light collection efficiency difference and high energy light detector resolution. The Pearson correlation coefficient is calculated between these variables. Under the assumption of Gaussian error distribution in both variables, the most probable correlation coefficient is $0.69^{+0.20}_{-0.69t}$ (FWHM), with errors calculated using Monte-Carlo. It then appears that there is some correlation between the magnitude of radial position dependence and high energy light detector resolutions. This can be understood as larger variations in light collection efficiency naturally lead to a wider variation in the amount of light measured for events of a given energy, and thus a worse resolution.

Module 10B has both the best light detector resolution and the least position dependence in terms of light collection efficiency. The absolute best resolution a light detector can reach is limited by the Poisson error on the number of photons reaching the light detector. This can be estimated at 911.2 keV. If the amount of energy in light that reaches the light detector is 1-2.5% [80]¹, and is in 3eV photons², the Poisson error will be $\sim 1.1-1.8\%$. Thus the 1σ resolution of 2.0% at 911.21 keV in module 10B is of comparable order to the Poisson error. This relatively good resolution is then attributed to a lack of position dependence.

Unexpectedly, the two light detectors of module 04B do not appear anomalous in terms of resolution or position dependence magnitude. Given the presence of a second light absorber in this detector module, one might have expected these detectors to individually have worse light collection efficiencies, and thus worse resolutions, than lone light detectors in other modules. However, both light detectors fall within the bulk of measurements in high energy resolution and positional changes in light collection efficiency. This implies adding a further light detector to a single module does not appear to have significantly worsened the performance of each light detector individually.

6.6 Conclusions and perspectives

Position dependent effects have been examined in CRESST detector modules. The underlying principle has been to use γ rays from external calibration sources of different energies as different depth probes of target crystal behaviour, making use of the differing penetration depths of different energy γ rays in matter.

¹ Another estimate comes from simulation; given the $\sim 8.3\%$ scintillation efficiency of [75] and the $\sim 20\%$ average light collection efficiency in polished crystal simulations of [121], one reaches an approximate amount of energy in light reaching the light detector of $\sim 1.7\%$.

² Calcium tungstate peak emission is at 420nm [80]. This is the most probable wavelength, the additional Poisson error due to variation in emission between $\sim 380\text{nm}$ and $\sim 480\text{nm}$ has been neglected here.

It was shown that when considering anti-correlation between the phonon and light detectors, position dependence can change the gradient of this anti-correlation, whereas such changes are not seen in the absence of position dependence. It was then shown that a variation in anti-correlation gradient is observed with mean interaction depth, between depths of $\sim 1\text{mm}$ and \sim homogeneously throughout the crystal, in at least one module. This change was explained in terms of a changing correlation between the amount of scintillation light produced by the crystal that is then reabsorbed, and the amount of light escaping the crystal that is absorbed in surrounding surfaces.

The average amount of light collected by light detectors was also shown to be related to the mean radial interaction depth of γ rays in nearly all light detectors, a further position dependent effect. Variations in the amount of light detected between interactions in the first millimetre from the cylindrical crystal surface (122.06 keV γ rays), and in the bulk of the crystal (911.2 keV γ rays), ranged between $3.1\pm 4.4\%$ and $17.3\pm 5.1\%$ between different light detectors. This variation between detectors is conjectured to be caused by different optical transport properties of scintillation light between crystals, itself caused by a variation in impurities and lattice deformations. These measurements of light collection efficiency changes assumed that scintillation efficiency does not vary between 122.06 keV and 911.21 keV. If scintillation efficiency is allowed to vary identically across all modules, position dependent effects on light collection efficiency are reduced, but still observable in nearly all detectors.

These light collection efficiency changes were seen to at least strongly affect module 10B, yet measurably affected this same module in terms of changes to anti-correlation gradient. The two measurements are of different effects. Anti-correlation gradient is sensitive to correlations in changes in crystal and surrounding surface absorption, whereas light collection efficiency measurements are sensitive to the magnitudes of crystal and surrounding surface

absorption. Thus, while counter-intuitive, this result is understandable in the position dependent framework outlined here.

High energy light detector resolution was found to show some correlation with the magnitude of position dependence in terms of light collection efficiency. This effect is attributed to variations in light collection efficiency naturally causing worse light detector resolutions. The two-light detector module, 04B, was seen to behave similarly to other modules in terms of both position dependence and detector resolution. This implies that additional light detectors may be added in future runs without significantly worsening the (high energy) resolution of individual light detectors resolution. As signals in the two light detectors can be combined, using two light detectors in a single module could improve overall module performance. However, the addition of extra light detectors should be weighed against the fact that six of nineteen light detectors installed did not take data in background files of the most recent CRESST run. Thus additional, independent light detectors increase the chance of having a non-operational detector in a module, and this should factor should be considered before using an extra light detector in all modules.

That module 10B shows superior high energy light detector resolution, and the least variation in light collection efficiency, can be attributed to the fact that module 10B is a glued detector. The gluing process was introduced to avoid evaporating thermometers directly onto the bulk crystal surface. This is because the evaporation process reduces scintillation light emission from a crystal. As 10B shows superior light detector resolution and low variations in light collection efficiency when compared to standard modules, this fact implies that making all future modules glued detectors should be considered. While the other glued detector, module 06B, did not show such as significant improvement in terms of position dependence variation, it still appears to be superior to other modules in terms of high energy light detector resolution.

The approximation was made in this work that the light detector systematic error is no worse than the worst phonon detector systematic error. However, with an accurate estimate of light detector systematic error, future work in this area could improve accuracy. This systematic error can be determined in future runs by direct light detector calibration at multiple energies. However, the last chapter showed this was not possible with the Fe-55 X-ray sources installed in the most recent run. As this chapter has showed that position dependence measurements in CRESST crystals can be performed in-situ, this further motivates the idea that light detectors should be calibrated in future runs using a different method, perhaps with optical photons.

Lastly, that source position relative to individual detectors was unavailable posed a challenge in this work. Source position was then estimated from schematic diagrams of CRESST, and this allowed a first approximation of mean γ interaction depth in target crystal. However, if source location becomes available in the future, results from this analysis could then be compared closely to simulation; crystal optical properties, such as optical photon scattering length, might then be inferable from such comparisons.

7 Extending commissioning run limits to lower masses

Motivated by the recent interest in light WIMPs of mass $\sim\mathcal{O}(10 \text{ GeV}/c^2)$, an extension of the elastic, spin-independent WIMP-nucleon cross section limits resulting from the CRESST-II commissioning run (2007) is presented. Previously, these data were used to set cross section limits for WIMP masses of $1000 \text{ GeV}/c^2$ down to $\sim 17 \text{ GeV}/c^2$, using tungsten recoils, in 47.9 kg-days of exposure of calcium tungstate. Here, the overlap of the oxygen and calcium bands with the acceptance region of the commissioning run data set is reconstructed, using previously published quenching factors. The resulting elastic WIMP cross section limits, accounting for the additional exposure of oxygen and calcium, are presented down to $5 \text{ GeV}/c^2$.

7.1 Introduction

In the CRESST-II commissioning run of 2007 [106], limits on elastic, spin-independent, WIMP-nucleon cross sections for WIMP masses from $1000 \text{ GeV}/c^2$ down to $\sim 17 \text{ GeV}/c^2$ were presented. However, light WIMPs, with a mass $\sim\mathcal{O}(10 \text{ GeV}/c^2)$ have been suggested as a possible interpretation to the experimental results of DAMA, CoGeNT and CRESST-II (2011) [51,53,63]. At the same time, several other experiments [49,50,57,64,70] partially or completely exclude the light WIMP interpretations of [51,53,63]. Therefore, it is of interest to see how data from the commissioning run compare to these results when extended to examine such light WIMP scenarios.

The author is not part of the CRESST-II (2011) results interpretation in [53] which considers a positive WIMP interpretation of the most recent CRESST run. The last part of this chapter constitutes the author's comments on possible systematic errors in this interpretation. The limits derived by the author in this chapter from the earlier commissioning run (2007) were published in [128].

7.2 Including Oxygen and Calcium

The CRESST-II commissioning run data set [106] consists of 47.9 kg-days exposure of CaWO_4 , taken between the 27th of March and the 23rd of July, 2007. The data were obtained from two independent modules, “Zora/SOS23” and “Verena/SOS21”, collecting 23.8 and 24.1 kg-days respectively [129]. As explained in chapter 2, the phonon detector is used to measure energy and the light detector to distinguish between event types, as different types of interaction in CaWO_4 produce different amounts of light, relative to the energy deposited. γ and β interactions, causing electron recoils, produce more light than nuclear recoils. The amount of light produced per unit energy of electron recoils, L_γ/E , was calibrated to those of gamma interactions at 122.06 keV and is normalised to a value of one. Nuclear recoils see a reduction in light compared to electron recoils, quantified by a quenching factor, Q_i , dependent upon the species i of recoiling nucleus.




Fig. 7.1 Acceptance region and nuclear recoil band diagram for Zora/SOS23. The black line indicates the light yield (L_γ/E) limit for the acceptance region in [106], and the vertical dashed lines the acceptance region's lower and upper energy thresholds. 1σ regions for observing recoiling oxygen (green), calcium (red) and tungsten (blue) nuclei are indicated, using light detector resolution parameters in Table 7.1 and quenching factors from [80]. The 90% acceptance line for tungsten from the commissioning run crosses the 1σ region for expected tungsten recoils due to the change in quenching factor measurements between runs. A similar diagram may be drawn for Verena/SOS21.

The limits of [106] resulted from assuming all recoils in the acceptance region, the region in which WIMP nucleon interactions are searched for, were from tungsten alone. However, due to the effects of a finite light detector resolution, recoils from both calcium and oxygen may also be seen within the same acceptance region. This effect is illustrated in Fig. 7.1. The additional exposure provided by the parts of calcium and oxygen bands that fall within the acceptance region can strengthen cross section limits for light WIMPs, with mass $\sim \mathcal{O}(10 \text{ GeV}/c^2)$. In [106], the acceptance region was chosen so that tungsten recoils would have been seen with minimal electron recoil band overlap. In energy, this was between 10 and 40 keV. In light yield, the upper limit was set so that 90% of tungsten recoils would occur in the acceptance region. Here, this same acceptance region is used, so that we do not introduce non-blind elements into the analysis.

To consider oxygen and calcium recoils in this region, the fraction of each nuclear species' recoils that fall within the acceptance region must be estimated. This requires several pieces of information. The first is the light detector resolution of the observed electron recoil band, as a function of energy. This resolution is expressed as:

$$\sigma_{\gamma}^2 = (\sigma_0)^2 + (\sigma_1)^2 E + (\sigma_2)^2 E^2 , \quad (7.1)$$

where E is the energy in the phonon detector. The resolution of the electron recoil band depends on three terms: σ_0 , reflecting electronic noise; σ_1 , related to the Poisson distribution of the expected number of detected photons; and σ_2 , incorporating position dependence and other possible effects seen in CRESST-II light detectors.

| Parameter | Zora/SOS23 | Verena/SOS21 |
|----------------------------------|------------|--------------|
| σ_0 (keV) | 0.784 | 1.508 |
| σ_1 (keV ^{1/2}) | 1.064 | 0.610 |
| σ_2 | 0.192 | 0.154 |
| P_0 (keV) | 0.56 | 0.11 |
| P_1 | 0.0040 | 0.0065 |

Table 7.1 Resolution values used in this work. σ_0 , σ_1 and σ_2 are light detector resolution parameters determined from figures in [106]. P_0 and P_1 are the phonon energy resolution parameters from [83].

We now wish to estimate the resolution of a quenched band. At energy E , events in the quenched band produce an average amount of light $L_{Q_i}(E) = Q_i L_\gamma(E)$. The resolution of a quenched band is assumed to be equal to the resolution of the electron recoil band at energy E' , where $L_{Q_i}(E) = L_\gamma(E')$. The exact light detector resolutions used in [106] are unavailable. However, these resolutions may be obtained by fitting to the acceptance region figures in [106]. To estimate light detector resolutions, oxygen and tungsten acceptance regions were modelled by:

$$Acc(Q_i, E) = \frac{Q_i L_\gamma(E) + N_{Sig} \sigma_{Q_i}(E)}{E}, \quad (7.2)$$

where Q_i is the quenching factor of the considered nucleus, and $N_{Sig} \approx 1.28$ is the number of standard deviations allowing 90% of quenched recoils to be seen in the acceptance region. These equations were fitted simultaneously to the tungsten and oxygen nuclear recoil acceptance regions in Figure 8 of [106]. The tungsten quenching factor is set at 2.5% and L_γ/E is taken to be one.

The oxygen quenching factor in the commissioning run was determined from a neutron calibration. The results of that neutron calibration are shown here as Fig. 7.2, together with expected the 90% acceptance line for oxygen nuclear recoils. The quenching factor for oxygen was given as ~11.1% in the text of [106]. It is stated in [106] that using the value 11.1%, one finds agreement with the expected 90% of nuclear recoils being within this acceptance region.




Fig. 7.2 AmBe neutron calibration of Verena/SOS21 in the commissioning run, from [106]. The expected 90% acceptance line for oxygen recoils, using a quenching factor of 11.1% is shown in red. In [106] it is stated that there is a good agreement between this expectation that 90% of oxygen recoils are seen below this line and measurement.

Here, as a test of fit method accuracy, the quenching factor is left free when fitting to both modules' acceptance region figures that are shown in [106]. The quenching factor was then found to be 11.0%. The obtained light detector resolutions are given in Table 7.1. Separately, the energy resolution of the phonon detector was modelled in [83] by $\Delta E = P_0 + P_1 E$, with energy resolution parameters also shown in Table 7.1.

The next piece of information is the quenching factor of each target nucleus. For this, the measurements in [80] are used, with $11.09^{+0.909}_{-0.908}\%$ for oxygen¹, $6.38^{+0.619}_{-0.653}\%$ for calcium and $3.91^{+0.478}_{-0.430}\%$ for tungsten. It should be noted that the more recent measurements of light output for tungsten recoils in CaWO_4 in [80] are higher than the 2.5% used in [106]. This means that the amount of light for tungsten recoils is on average higher than that which was expected in [106], causing less than the expected 90% of all tungsten recoils to fall within the acceptance region, as can be seen in Fig. 7.3 for Zora/SOS23.

¹ An oxygen quenching factor of $10.4^{+0.5}_{-0.5}\%$ was used in [53]. Limits calculated using this quenching factor are slightly stronger than those presented here, a result of a larger fraction of oxygen recoils being observable within the acceptance region.




Fig. 7.3 Estimates of the fraction of recoils of a given target nucleus that fall within the acceptance region for Zora/SOS23 in [106], as a function of energy. Oxygen is shown in green, calcium in red and tungsten in blue, with black showing the initial assumption of 90 of tungsten recoils falling within the acceptance region. Vertical dashed lines at 10 and 40 keV indicate acceptance region limits in recoil energy.

One last piece of information would be required for a complete description of detected light. This is the small deviation of observed light in the electron recoil band from the normalisation of one unit of light per unit energy. Two effects can cause this deviation: the dependence of light yield on energy in inorganic scintillators [130], and an overall calibration error. Such adjustments as used in the analysis of [106], are unavailable, although it is stated in [106] that the light yield is always near the normalisation of one. Here we use the approximation that the mean electron recoil light yield is one everywhere. In an independent analysis of the commissioning run data [83], the electron recoil band behaviour and light detector resolution were parameterised, with results repeated in Appendix E. As a check on these results, the WIMP cross section limits resulting from these light detector parameters was also considered. They are consistent with those presented here to within a few percent.

With this information, the fraction of recoiling nuclei from each constituent of CaWO_4 that falls within the acceptance region of [106] can be estimated, as shown in Fig. 7.3 for Zora/SOS23. With these acceptance fractions, the interaction

rate of WIMPs with oxygen and calcium in the commissioning run acceptance region may now be calculated. Here a method analogous to that in [106] is followed. The elastic, spin-independent WIMP-nucleon interaction rates are calculated following [45], using the Helm form factor parameterisation as suggested in [46]. The total rate expected from all target nuclei is then:

$$\frac{dR_{Tot}}{dE_R} = A_W \frac{dR_W}{dE_R} + A_{Ca} \frac{dR_{Ca}}{dE_R} + A_O \frac{dR_O}{dE_R}, \quad (7.3)$$

for the fractions A_i of each species' nuclear recoils that may be seen in the acceptance region. This rate is convolved with the observed phonon energy resolution, ΔE , as described in [46].

7.3 Maximum gap

With the spin-independent expected rate calculated, a given event distribution allows a calculation of WIMP-nucleon cross section limits. These limits set the maximum cross section that could give rise to the observed event distribution, to a given degree of confidence, across possible WIMP masses. A number of methods exist to make this step. The first is the Poisson method, which sets limits using the expected rate in the entire signal region and the simple observed count rate. If events were seen in the acceptance region, this method naturally sets the weakest cross section limits by assuming all events were equally likely to be signal events.

However, this simple limit can be improved using knowledge of the expected spectrum of WIMP-nucleon recoils. The rate of nuclear recoils induced by the coherent, elastic scattering of WIMPs is expected to follow an approximately exponential decay in energy [46]. If backgrounds are expected, then one can compare the observed spectrum of events to the expected WIMP spectrum of events to set tighter constraints. Particularly if background rates and spectrums are well known, limits can be set by *subtracting* these expected backgrounds, a method pursued in the most recent CRESST analysis of [53].

However, an alternative is to note that CRESST may have an effectively *unknown* background source (or sources), leading to events observed within the signal region. One can then compare the observed event distribution to the expected, accounting for a completely unknown background contamination. Such a method is pursued here in light of the difficulties in estimating backgrounds in CRESST detectors discussed in section 7.5.




Fig. 7.4 Diagram of Maximum Gap method, derived from a similar figure in [131]. Hypothetical event energies are indicated in blue. The maximum gap is shown by the purple shaded region. The total expected rate is the integral between the two vertical dashed lines, which represent the lower and upper thresholds.

The Maximum Gap method [131] was developed to set WIMP cross section limits in the unknown background scenario. This method can be used where expected signal is a function of cross section and a single further parameter. This parameter can be recoil energy E , such that $\frac{d\Gamma}{dE} \equiv \frac{d\Gamma}{dE}(\sigma, E)$. Given a set of N events ordered in recoil energy, the WIMP-nucleon cross sections that could give rise to this energy distribution of observed events can be found. This is done by computing the “gaps” between observed event energies, where each gap Γ_i is the expected number of events between two adjacent events:

$$\Gamma_i(\sigma) = \int_{E_i}^{E_{i+1}} \frac{d\Gamma}{dE_R}(\sigma_i) dE_R, \quad (7.4)$$

where $\frac{d\Gamma}{dE_R}(\sigma_i) = \frac{d\Gamma}{dE_R}(\sigma_0) \frac{\sigma_i}{\sigma_0}$. This integral is shown diagrammatically in Fig. 7.4.

The lower and upper thresholds are treated identically to events, labelled E_0 and

E_{N+1} respectively. The maximum gap Γ_M is then the largest of these Γ_i . This maximum gap is the maximum expected interaction rate between two event energies. The total expected rate between lower and upper thresholds is given by the integral of Eqn. 7.4, with integration limits E_0 and E_{N+1} . For a given WIMP-nucleon cross section σ_i , Γ_M and Γ_T can be related to C_0 , the probability that WIMPs at this cross section could give rise to the observed event distribution:

$$C_0(\Gamma_M, \Gamma_T) = \sum_{k=0}^m \frac{(k\Gamma_M - \Gamma_T)^k e^{-k\Gamma_M}}{k!} \left(1 + \frac{k}{\Gamma_T - k\Gamma_M}\right), \quad (7.5)$$

where m is the largest integer that satisfies $m \leq \frac{\Gamma_T}{\Gamma_M}$. Differences in light detector resolution between modules will cause slightly different overlaps between nuclear recoil bands and the acceptance region between modules. Each module will then show a slightly different expected interaction rate at each energy. Modules are then added in a “simple merger”, normalised to the cumulative expected event rate [132]. Finally, it is noted the Maximum Gap is not the only limit setting method available, with several methods reviewed for CRESST in [95]. In particular, several methods also make use of the event distribution in the discrimination parameter (light yield). These methods tend to set WIMP-nucleon cross section constraints at least as strict as the Maximum Gap. However, the Maximum Gap method is used here, to remain consistent with [106].

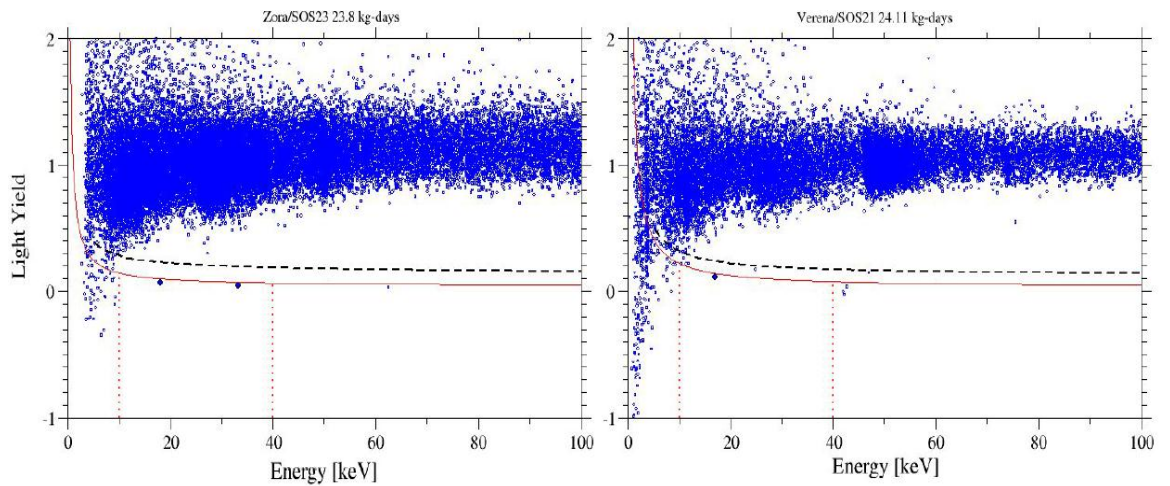


Fig. 7.5 Event distribution of the commissioning run dark matter data set. Blue points with light yield ~ 1 represents events reconstructed in the electron recoil band. The red (black-dashed) line represents the expected 90% acceptance lines for tungsten (oxygen) events. Three events are seen in the dark matter signal region of the commissioning run, defined as between 10 keV and 40 keV and below the expected 90% tungsten acceptance line. Figure taken from [106].

Three events are observed in [106], at 16.89 keV, 18.03 keV and 33.09 keV, as shown in Fig. 7.5. When the maximum gap is applied to commissioning run data, the maximum gap is the integrated rate between the lower energy threshold (10 keV) and the first event (16.89 keV) for all WIMP masses.

7.4 WIMP limits




Fig. 7.6 90% confidence limits on elastic, spin-independent, WIMP-nucleon cross sections from data in [106], considering all possible nuclear recoils within the acceptance region. The green dashed-dotted, red dashed and light-blue dotted lines result from considering WIMP interactions with oxygen, calcium and tungsten individually, and the solid blue line all nuclei. The black dashed line indicates the tungsten only limit using the same quenching factor as in the commissioning run. This limit is a few percent stronger than the limit using the updated quenching factor, and the two lines appear almost identical in this figure. The WIMP halo properties used are $\rho_{DM} = 0.3 \text{ GeV/cm}^3$, $v_{esc} = 544 \text{ km/s}$, $v_0 = 220 \text{ km/s}$ and $v_{sun} = 232 \text{ km/s}$. Resolutions from Table 7.1 and quenching factors from [80] were used to derive these limits.

Fig. 7.6 shows how the extended 90% confidence limits, calculated using the Maximum Gap, improves sensitivity to low mass WIMPs with commissioning run data. Interactions with WIMPs heavier than $\sim 17 \text{ GeV}/c^2$ in the acceptance region are dominated by tungsten recoils, and the cross section limit is well modelled by considering interactions from tungsten alone. At lower masses, calcium, then oxygen recoils become dominant, such that below $\sim 7 \text{ GeV}/c^2$, nearly all WIMP-nucleon interactions in the acceptance region are with oxygen nuclei. Considering all possible nuclear recoils then provides a significant strengthening of cross section limits at low masses compared to tungsten alone.




Fig. 7.7 Total 90% confidence limits on elastic, spin-independent WIMP-nucleon cross sections from a reanalysis of the commissioning run data in this work (solid blue). For comparison, the favoured regions from DAMA, from [51], CoGeNT [63] and CRESST-II (2011) [53] are shown. With gratitude to F. Pröbst and C. McCabe for providing the confidence contours of CRESST-II (2011), and DAMA respectively.

Fig. 7.7, a comparison of the total limits in this work is made to the elastic WIMP interpretation of other experiments. These limits, in relation to existing WIMP-nucleon cross section limits, were shown in Fig. 1.9. The CoGeNT, CRESST-II (2011) and DAMA results were already in tension with the results of XENON100, XENON10 (S2 only), CDMS II, CDMS II (low energy), ZEPLIN-III. The extended commissioning run limits introduce further tension with DAMA and a mild tension with CRESST-II (2011), particularly below 10 GeV/c^2 . As the commissioning run and CRESST-II (2011) results are with the same target nuclei with similar energy thresholds, it is difficult to reduce this mild tension by choosing different astrophysical parameters or particle physics models.

However, it should be noted that the CRESST-II commissioning run and CRESST-II (2011) run do not use the same acceptance region definitions. In the commissioning run, the acceptance region was 10-40 keV in energy, and chosen to include 90% of tungsten recoils. This same acceptance region has been used here, to ensure additional non-blind elements have not been introduced by this analysis. The analysis of the CRESST-II (2011) run uses a different acceptance

region definition. Lower energy thresholds are generally higher, at ~ 12 keV for most detectors. Higher thresholds tend to decrease sensitivity to WIMPs in a per-kg basis. [53] also includes regions of higher light yield, allowing for more calcium and oxygen recoils than in the commissioning run. This tends to increase the number of WIMPs that can be seen in the acceptance region, in a per kg-day basis. However, a significant difference between the *analyses* of the commissioning run [106] and of the 2011 run [53] is in terms of backgrounds. Between runs, the design of clamps in direct contact with the target crystals has changed, which as noted in [53] introduced additional backgrounds into the CRESST-II (2011) data set. To account for these additional backgrounds, the results of [53] have been background-subtracted, whereas those of this chapter have not. Given these differences, what follows is a discussion of *possible* sources of systematic error in background estimates of [53], with a view to their investigation for future CRESST dark matter searches.

7.5 Possible sources of systematic error in CRESST background analyses

7.5.1 Electron recoil leakage

The definition of the acceptance region in terms of light yield has changed in [53] from that used in the commissioning run [106]. Particularly, it includes regions of high light yield, to include slightly more oxygen and calcium recoils (at a given energy) in [53] than [106]. As described in chapter 2, the light yield expected for electron recoils (γ or β) is significantly higher than the light yield expected for nuclear recoils. Electron recoil leakage refers to the possibility that, given large numbers of γ and β interactions in background data, a small percentage of electron recoils leaks into the acceptance region. As electron recoils give more light on average than nuclear recoils, electron recoil leakage is a high light yield contaminant in the dark matter acceptance region of [53]. In [53], this leakage was modelled by the “tails” of a Gaussian distribution of light yield. The tails are the statistically disfavoured region of the assumed Gaussian distribution of events; in

this case, the low probability that electron recoils are seen with very low light yield. The mean and deviation of this distribution was determined from observed background γ and β interactions¹.

There are several possible systematic effects in this model that can affect electron recoil leakage estimates. The first is time-varying noise. As noted in section 4.2.2, noise on many detectors was not constant in time. Changing noise was noted to occasionally cause rapid triggering of the detectors. However, extended periods of higher noise might not cause repeated triggering, yet can still alter light detector resolution, through the noise term σ_0 in Eqn. 7.1. Fig. 7.8 shows an extreme case of time-varying noise for the detector “David” on module 09B in the most recent run.

Time-varying noise on light detectors implies a time-varying resolution. If only the average light detector resolution over the entire period is used, electron recoil leakage can be misestimated. This is simply because when taking the average of many Gaussian distributions of different widths, measurements are dominated by the bulk of this distribution, and the tails can be modelled poorly.

However, this effect can be checked for in future work. Light detector resolution might be better represented by several time-separated noise periods. Each period can then be allowed to have a different light detector resolution. Performing this separation will mitigate any time varying systematic effects. Such a separation was performed for commissioning run data in the theses [83] and [109], though not with a view to calculating electron recoil leakage. Each accounted for time-varying noise using different methods. Repeating either method, and comparing the resulting electron recoil leakage estimates, would reduce time-varying noise induced uncertainty.

¹ σ_0 , the noise term in the light detector resolution from Equ. 7.1, is however measured using the test pulses.




Fig. 7.8 FWHM of noise seen on recorded “empty baselines” for light detector “David” in module 09B of the most recent CRESST run. Empty baselines are noise recordings, recorded by an artificial trigger sent periodically to the trigger unit. Data taken is from a sub-set of files used in [53]. Only basic data quality cuts are applied to give a complete picture of noise variations - importantly, the “RMS” cut is not used, which would distort diagrams of noise vs. time

However, two other factors, more difficult to account for, could contribute to a non-Gaussian electron recoil leakage. The first is the difference between γ and β scintillation efficiency observed in [83]. This was measured such that γ interactions saw $13.9 \pm 0.6\%$ less light than β interactions at 74 keV. Light detector resolutions are measured from background interactions, assuming γ and β events have the same scintillation efficiency. However, if a discrepancy between scintillation efficiencies exists, this will distort light detector resolution estimates. Further, several γ peaks are seen at acceptance region energies, such as at 9.3 keV and 24.5 keV (total γ energies) from Ac-227 β decay. Near these energies, electron recoil leakage should vary if γ and β scintillation efficiencies are different.

Lastly, position dependence, discussed in chapter 6, could cause non-Gaussian electron recoil band leakage. In chapter 6, the amount of light reaching the light detector was seen to drop significantly radially towards the surface, confirming simulations in [121] and measurements in [83]. Though not studied here, detected light was predicted to also drop significantly for interactions very close to the tungsten thermometer in [121] (as can be seen at the bottom of Fig. 6.1).

If less light reaches the light detector, events will be reconstructed with a reduced light yield, and be more likely to leak into nuclear recoil bands. This effect could also distort electron recoil leakage from a Gaussian estimation. To account for this, it would be necessary to match observed position dependence to simulation.

7.5.2 Alpha quenching factor




Fig. 7.9 Dark matter data of Ch51/Ch52 (Module 10B, Rita/Steven in this work), from [53]. The acceptance region is shown in orange. The band of expected α recoils in yellow, oxygen recoils in purple, tungsten recoils in grey, and the reference region for lead recoils (section 7.5.3) shown in green. Black points are all events, encircled red events are those falling within the acceptance region.

α interactions are expected in the acceptance region of [53]. α decays typically have a Q value of several MeV. However, α events can be seen at low energy in CRESST due to surface effects. If an α particle loses energy on the way to the target crystal, or only deposits a portion of its energy in the target crystal, α interactions could be seen in the ~ 12 -40 keV acceptance region of [53].

A factor affecting calculation of the number of α events expected in the acceptance region is the α quenching factor, Q_α . Changes in Q_α will change the overlap of the alpha band with the acceptance region of [53], and also the number of events in the α “reference region” used in [53]. In [53], Q_α is stated to be $\sim 22\%$, shown in Fig. 7.9. However, several values in the literature exist for Q_α . The value

17.5±0.1% is found in [83] and [133]. 19% is used in [134]. 21% is used in [135]. 28% can be found in [78]. Further, these values are at a variety of energies, and the possibility that Q_α changes with energy, such that $Q_\alpha = Q_\alpha(E)$, has yet to be studied. It is then important that the low energy α quenching factor be determined accurately.

7.5.3 Low energy behaviour of the lead recoil spectrum




Fig. 7.10a Observed Pb-206 recoil spectrum in [53]. An exponential plus a constant has been used to model this spectrum in the 40-90 keV “reference region” (red line). Fig. 7.10b Simulated Pb-206 recoil energy spectrum using the SRIM package in [53]. Each curve represents a different depth profile of the decays in the surface in which the Pb-206 nuclei originate. The red curve is scaled from the same fit in Fig. 7.10a.

α decays near surfaces of CRESST target crystals can lead to recoiling heavy nuclei seen in the acceptance region of [53]. This background is principally caused by Po-210 α decays, where Po-210 is a daughter of radioactive radon gas (Rn-222). Radon gas can be adsorbed onto surfaces that surround CRESST target crystals. This adsorption is particularly dangerous if the surface is non-scintillating, such as in the silver that coats the bronze clamps in contact with CRESST crystals. In these cases, Po-210 decays can lead to a recoiling Pb-206 nucleus depositing energy in CRESST detectors. Recoiling Pb-206 nuclei produce very little light, with $Q_{Pb} = 2.35^{+0.383}_{-0.310}\%$ from [80]. Thus, Pb-206 recoils form a low light yield contaminant in the acceptance region of [53]. The full recoil energy of a Pb-206 nucleus emitted in a Po-210 α decay is 103 keV. However, if the decay is within a surface that surrounds a target crystal, some of this recoil energy can be

lost before the Pb-206 interacts in the target crystal. Lead nuclei can then deposit energies $\sim 12\text{-}40$ keV, a background in CRESST dark matter searches.

In [53], the energy spectrum of Pb-206 recoils was simulated using the Stopping and Range of Ions in Matter (SRIM) package [136]. 206-Pb recoils were simulated to originate in the silver that coats the clamps in direct contact with target crystals in the most recent run. Results of that simulation are shown in Fig. 7.10b. A variable in simulations is the implantation profile of Po-210 in the silver surface, that is, the depth-distribution of Po-210 decays. Simulations were performed with a variety of profiles, with the closest match to a Pb-206 recoil “reference region”, 40-90keV, used for background estimates. Using simulations with this implantation profile, Pb-206 recoils were then extrapolated to form an approximately flat background contribution in energy in the acceptance region, $\sim 12\text{-}40$ keV. However the actual implantation profile of 210-Po in the silver clamps is not known independently from these simulations. Further, no experimental evidence independently confirms the behaviour of this energy spectrum below 40 keV. Two experiments have however been performed in *similar* conditions to the recent CRESST run, and these are now discussed.




Fig. 7.11 Energy spectrum of events in the nuclear recoil band of module “Daisy/BE13” in CRESST Run 27, taken from [83]. A significant number of 206-Pb recoil events are seen at 103 keV, despite only 12.31 kg-days of data being represented in this figure. Noting low statistics, there appears to be a minimum in this spectrum at $\sim 20\text{-}60$ keV, but below this, the number of events in the nuclear recoil band appears to rise.

The first relevant experiment is CRESST Run 27 (2003). The most important difference in module design between Run 27 and the most recent CRESST run (Run 32) is in the reflecting foil surrounding each module. In Run 27, a non-scintillating silver foil surrounded detectors. In the most recent run a polymeric, scintillating foil is used. As Run 27 used non-scintillating foil, Po-210 decays in the foil also caused recoiling Pb-206 nuclei to be observed in the acceptance region. The results from an analysis of Run 27 in [83] are repeated in Fig. 7.12. Superficially, these results show an increase of events in the nuclear recoil band at low energy from a minimum between approximately 20 keV and 60 keV. However, statistics are small, with only 11 nuclear recoil events in the 12-40 keV region in Fig. 7.12. Further, it should be noted that no neutron shield was present in Run 27. Thus it is possible to attribute any low energy excess of events to neutrons. Gamma leakage was also not considered in [83]. It is then not possible to conclusively show the low-energy behaviour of Pb-206 recoils in Run 27.




Fig. 7.12 Po-210 surface calibration from [137]. Shown is the energy spectrum of the observed Pb-206 recoils in the nuclear recoil band. The energy and light yield behaviour is different to that in [53] (see text). Attempts have been made in this figure to correct for the anti-correlation effect discussed in chapter 6, which leads to the small difference between “with phonon quenching” and “without phonon quenching”.

After Run 27, a dedicated experiment was performed to examine the spectrum of Pb-206 recoil events in [137]. This experiment used a chemically produced Po-210 source. Po-210 was deposited onto a polished silver disk using a radium solution.

This Po-210 source was then placed inside a CRESST module, facing the cylindrical crystal surface. Fig. 7.12 shows the results of this Po-210 surface calibration. In this experiment, it was seen that the energy spectrum of Pb-206 recoils did not become flat below 40 keV from a peak at 103 keV. Instead, the spectrum reduced to a minimum at $\sim 40\text{-}60$ keV, then increased at lower energies. SRIM simulations of this resulting distribution were also performed in [137]. However, the behaviour of the mean light yield behaviour of Pb-206 recoils could not be explained in the SRIM simulations with silver only. To explain the light yield behaviour, a hydrocarbon polymer was inserted into simulation. This was however, not known to be on the surface of the source, but was added simply to match simulation to observed light yield behaviour [133].

The depth profile of Po-210 in the source and the geometry of the source position are different between [137] and [53]. Further, it is clear that the energy behaviour in the 40-90 keV region of [137] (Fig. 7.12) does not closely match that in [53] (Fig. 7.10). However, it is useful to note that the SRIM simulations of [137] did not match observed event characteristics without adding further materials to the simulation (in that case, a hydrocarbon polymer). It is thus vital that, if this background remains in future dark matter searches, its low energy behaviour should be experimentally determined, independently from simulation.

7.5.4 Long decay or thermal relaxation events

A background to the acceptance regions of the commissioning run was dubbed *long decay-time events* [83] or alternatively *thermal relaxation events* [109]. These events showed significant reconstructed phonon detector amplitude, but negligible light amplitude. These were thus low light yield contaminant to the commissioning run. These events were identified as being separable from particle events, showing a variety of long decay times, longer than decay times of particle pulses, with an example from [83] shown in Fig. 7.13 (left).

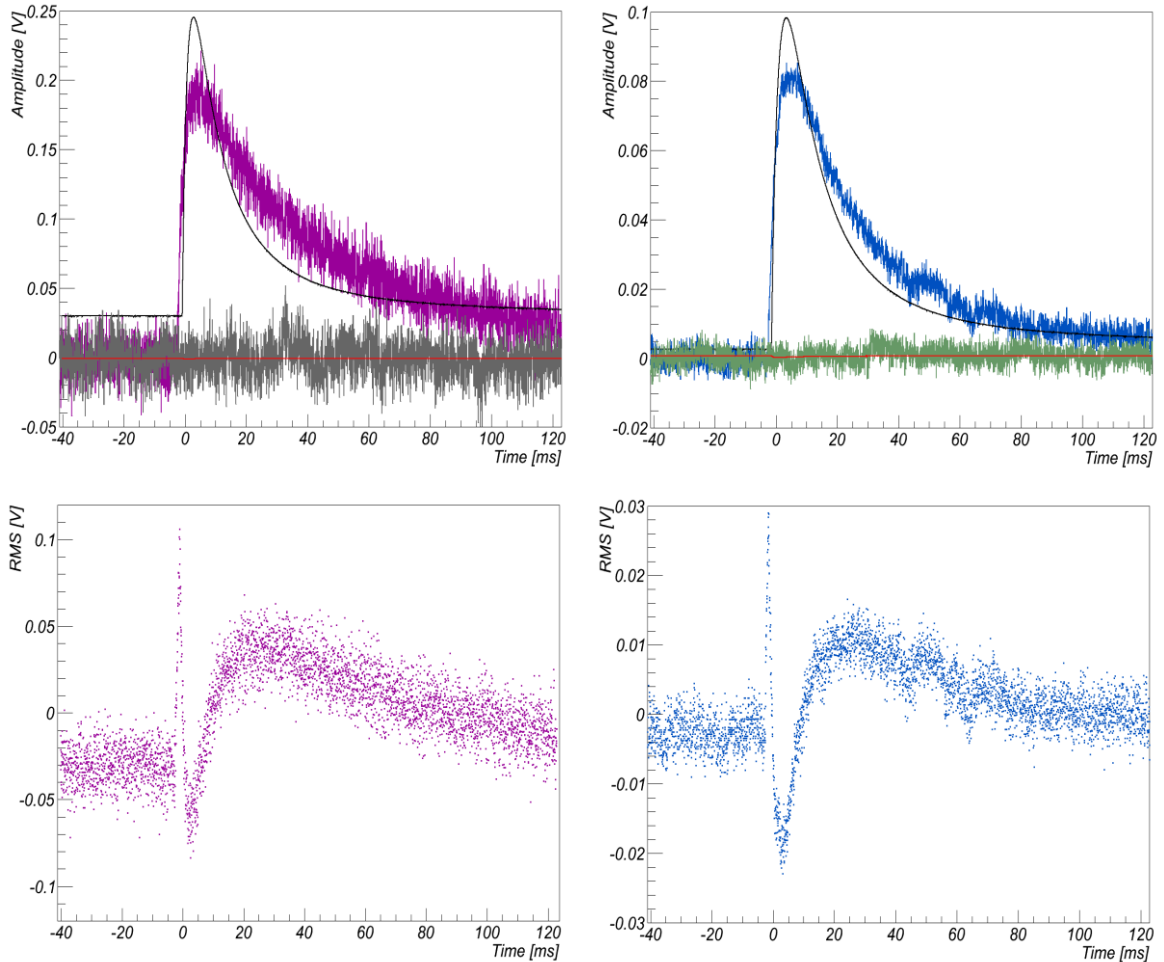


Fig. 7.13 (left) Long decay-time event in the commissioning run. Top left shows a standard pulse fit (black)¹ to the measured phonon detector response on “Zora” (magenta). This event survived the RMS cuts of [83], yet as can be seen in residual of the fit in the phonon detector (lower left), has a systematic pulse shape difference from particle pulses, with a longer decay time. Negligible amplitude is seen in the light detector “SOS23” fit (red) to data (grey). The event was then a contaminant to the acceptance region of [83], seen at 28.6 keV, passing RMS cuts. However, this event was identified as having a long decay time using the proposed “decay-time” cut in section 9.5.6 of [83]. Fig. 7.13 (right) shows a candidate long decay-time event in module 04B in the most recent CRESST run (Run 32). Top right shows the fitting of a particle standard pulse to response the phonon detector “Verena” (blue). A similar systematic pulse shape difference is observable in the RMS profile of the phonon detector (bottom right). Again, a negligible amplitude signal is reconstructed for the light detector “Q” (green) in the top right figure. 04B is the two light detector module. The other light detector (“Burkhard”) also sees a negligible amplitude response in this event. This event just fails the RMS cuts on the phonon detector with measured energy 17.0 keV. The event occurred at 13:04 (GMT), 27/12/2010.

¹ Technical note: Record length changed between the commissioning run (163.84ms) and the most recent run (327.68ms). Both fits in this figure are with a *flat* baseline assumption, whereas all other fits in this thesis allow a *linear* change. Given the shorter record length, allowing a linear baseline change in the commissioning run significantly distorts fits of long-decay time event with respect to the most recent run, as only half as much baseline was recorded in the commissioning run. To make first comparisons between runs, linear fitting is then disallowed to make this figure. The cut efficiency Fig. 7.14 below however allows for a linear baseline in fitting, as this is done at all other points in this work. Given recording length difference, only the (-40.96 ms, +122.88ms) part of a record is shown in Fig. 7.13

In [83], these events were attributed to the scintillating foil surrounding clamps in the commissioning run. This foil has since been removed, and long-decay time events were then not expected in the most recent CRESST run. However, examination of background files reveals the existence of a small population of these events on some modules in the most recent run, such as 04B. These events have a variety of decay times, longer than particle events, and thus often fail the RMS cut in section 4.2.10. An example of an event with a long decay-time in the most recent data failing the RMS cut in this work is shown in Fig. 7.13 (right).

That this population of long-decay time events appears in the new run is surprising. Many of these events appear to be cut by the RMS cut. However, the efficiency of the RMS cut in removing long decay-times events will not be independent of reconstructed energy. Long-decay time events are reconstructed with a variety of energies, at least up to several tens of keV. Qualitatively, if an event with the same pulse shape occurred with lower amplitude than that in Fig. 7.13 (right), the bare magnitude of pulse shape discrepancy between standard pulse shape and recorded pulse shape will be lower. It is then possible for the RMS cut to fail, such that these events could remain after cuts.

A first quantitative estimate of RMS cut efficiency can be obtained through simulation. The fact that the event of Fig. 7.13 (right) failed RMS cuts, with reconstructed energy 17.0 keV, is used to tag this event as a long-decay time event. Fitting Eqn. 2.6 to this event then yields a parametric description of this long-decay time event's pulse shape. The event simulator of [109] was then used to simulate long decay time events with this pulse shape, at reconstructed energies between 0 keV and 40 keV. Fig. 7.14 shows the efficiency of the RMS cut in removing events of this shape, against reconstructed energy.




Fig. 7.14 Cut efficiency versus reconstructed energy for events of the same pulse shape as that in Fig. 7.13 (right). Events are simulated with reconstructed energies of 0-40 keV using the event simulator of [109].

In simulation, the RMS cut in module 04B begins to fail in removing this class of events below ~ 15 keV, and is completely ineffective by 12 keV. Events of this class show negligible light signals. Thus, with an acceptance region of 12.9 keV to 40 keV for 04B as in [53], if events of the same pulse shape as in Fig. 7.13 (right) exist at energies 12.9 keV to ~ 15 keV, they can contaminate dark matter analyses using RMS cuts in this work.

However, this cut efficiency estimate is conservative – an event that just failed RMS cuts (at 17.0 keV) was used as a template, and it is then not surprising that the RMS cut becomes ineffective just below this energy. The amplitude and pulse shape distribution of these events is unknown, and further, fit method and RMS cut parameters used in this work likely differ slightly to those used in [53]. Thus, Fig. 7.14 is not a direct comparison to the cut efficiency for this class of events that will occur in the analysis of [53]. The conclusion of this section is simply to note that long-decay time events still exist in CRESST data, even after changes in clamp design. This class of events appears to be a possible background in dark matter searches, using cuts in this thesis. Any contamination of the acceptance region will depend on the distribution of these events in amplitude and pulse shape, as well as RMS cut choice. Thus, further investigation of this background is warranted in future CRESST dark matter searches.

7.6 Conclusions

The WIMP cross section limits for the 47.9 kg-days exposure of CaWO_4 in the CRESST-II commissioning run [106] have been extended down to a WIMP mass of $5 \text{ GeV}/c^2$. This analysis has accounted for possible oxygen and calcium recoils within the commissioning run acceptance region, using light and phonon detector resolutions in Table 7.1 and quenching factors from [80]. The improvement of cross section limits at light masses occurs because recoiling oxygen and calcium nuclei dominate over tungsten recoiling nuclei for light WIMPs within the commissioning run acceptance region.

Extending the commissioning run limits results in mild tension with the recent CRESST-II (2011) [53] and DAMA [51] results. The analyses of the 2011 run and the CRESST-II commissioning run use different methods; the analysis in [53] is background subtracted, and those of this work are not. *Possible* sources of systematic error in the background estimates of [53] have then been discussed, and several possible sources of systematic error in calculating the excess of signal over background in [53] have been identified. None of these can be conclusively shown to account for the excess of signal over background in [53], however all deserve study in future CRESST-II dark matter searches.

8 Conclusions and the future

CRESST-II completed its most recent run, Run 32 in 2011. Data taken in this new run presented new challenges in terms of analysis software requirements. Data volume increased substantially over previous runs, and it began to be difficult to use Oxrop, analysis software in use within the CRESST collaboration, to analyse this new data. At the same time, Oxrop needed to be extended to analyse the EDELWEISS experiment, which required additional modifications, and needed to be improved for future use with EURECA. To meet these challenges the structure of Oxrop has been redesigned (chapter 3). The new version of Oxrop, Oxrop 5, has a more intuitive, consistent structure, requires significantly less memory, and can be used for a greater variety of experiments. Cuts were then designed (chapter 4) to allow Oxrop to analyse CRESST data.

The unifying theme of this thesis was subsequently to use Oxrop to apply models to different aspects of CRESST analyses. First, light detector behaviour in Run 32 was analysed (chapter 5). It was seen that light detector behaviour is not well modelled by the simple expected model of one exponential rise time and two exponential decay times. This was both for light detector responses to γ interactions in the crystal scintillator, and to interactions directly in the light detectors themselves.

There are several reasons why the light detector might exhibit further decay times, for example due to the influence of the phonon collector. Empirical extensions to this simple model were made, allowing for additional decay times, and light detector response to crystal scintillator interactions was shown to be well modelled by these extensions. The millikelvin γ scintillation time of CaWO_4 of $\approx 600\mu\text{s} \pm 100\mu\text{s}$ (sys) was estimated using this model, with systematic error estimated by the distribution of the nine measurements from individual light detectors. Measurement of a single light detector associated with a ZnWO_4

module was also performed, and the γ scintillation time was inferred to be $80\mu\text{s}\pm 40\mu\text{s}$ (stat) or $480\mu\text{s}\pm 40\mu\text{s}$ (stat), with fit minima degenerate between the two possible scintillation times.

With model extensions, measured response to interactions directly in the light detector with 5.9 keV X-rays from an Fe-55 source still showed significant tension with the applied model. Even accounting for detector transition non-linearity, these direct interactions were poorly modelled. This suggests a new approach to direct calibration of light detectors should be made for future CRESST runs. This could perhaps be done using optical photons of the same frequency as scintillation light from the same target crystal, but from an artificial source, fed into the CRESST cryostat through an optical fibre.

Position dependence effects were then studied in calibration data of Run 32 (chapter 6). The work made use of the different penetrating depths of γ rays from external Co-57 and Th-232 sources at different energies in calcium tungstate. Different energy γ rays were described as sampling behaviour in different areas of the crystal target. Using this calibration data, it was seen that the fraction of light produced that is seen in the light detector, the light collection efficiency, measurably reduced when the mean position of γ interactions is closer to the cylindrical surface edge of CRESST target crystals than in the bulk. This behaviour was seen both if one assumes scintillation efficiency for γ interactions in CaWO_4 is independent of energy between 122.06 keV to 911.2 keV, or if one allows the efficiency of γ induced scintillation to vary according to measurements in this work. Using calibration data from multiple CRESST modules, it was also shown that the magnitude of the difference between surface and bulk collection efficiencies is correlated between with the high energy light detector resolution.

Anti-correlation between the phonon and light detectors was also considered in a position dependent framework. It was shown that without position dependence, anti-correlation gradient does not change with energy; however it can with position dependence. Given that a change was observed in at least one module under study, this shows for the first time that position dependence must play a role in anti-correlation between phonon and light detectors.

Lastly, data from the CRESST-II commissioning run was re-examined (chapter 7). The original analysis of the commissioning run assumed that all interactions in the acceptance region, the region in which WIMP interactions were searched for, were with the tungsten in the CaWO_4 crystals used in CRESST. In this work, this assumption was relaxed, modelling interactions from both calcium and oxygen nuclei in the same acceptance region as originally used in the commissioning run. This was done by reconstructing the behaviour of the calcium and oxygen bands, using light detector resolutions obtained by fitting to acceptance region figures in [106] and quenching factors from [80].

The resulting WIMP-nucleon cross section limits were improved at light WIMP masses $\sim \mathcal{O}(10 \text{ GeV}/c^2)$. These new limits introduce a mild tension with a recent dark matter analysis of CRESST-II [53] and DAMA [51]. Possible systematic errors in the recent dark matter analysis of CRESST-II in [53] were then discussed. No systematic error could be conclusively shown to account for the excess of signal over background in [53], however, several systematic error sources were noted that warrant further investigation.

Appendix B Calculating detector live time

This appendix discusses module live-time. No dark matter analysis of data from the new CRESST run (Run 32) is used in this thesis to make WIMP limits. However, this appendix is included for discussing the effects of cuts developed in chapter 4 in a full dark matter analysis. First, one must introduce the *dead-time*. This is the difference between this live-time and the total recording time. The data acquisition unit (DAQ) automatically increments the dead time when a module could not have seen a dark matter interaction. Briefly, DAQ dead time occurs for 3 reasons [111]:

- During the readout of a particle interaction. The readout time, plus an additional two pre-trigger lengths before trigger re-activation is added to the dead time.
- Readout of other modules on the same transient recorder. Other modules may only trigger in the first half of the post-trigger time in a record. Subsequent time before trigger re-activation is added to the dead time.
- The time of recording test or control pulses active on a module, plus readout time plus one pre-trigger length before trigger re-activation.

Cuts themselves then also introduce further reductions in detector live-time. The stability and rate cuts in chapter 4 both remove any live time accumulated during periods tagged as “unstable”. For the stability cut, the live time between the first stable control pulse before and after an unstable region is removed. For the rate cut, the time during a time-bin with a fast trigger rate is removed. Cuts in sections 4.2.3 – 4.2.9 are then all designed to remove specific contaminants. While the cut in section 4.2.9 is a pulse shape cut, effectively no significant fraction of background β or γ bulk crystal hits is removed by this cut, as illustrated in Fig. 4.9. All these cuts then decrement the live time simply by the live-time seen during the actual recording, during which a WIMP signal could not have been observed. Coincidence cuts in section 4.2.11 and 4.2.12 remove any detector live

time where a veto could have occurred. For the muon veto, this is the $\pm 10\text{ms}$ window around every muon event. For the coincidence cut, this is the $\pm 5\text{ms}$ window of every other cryogenic particle event recorded.

The last cut not dealt with so far is the RMS cut (section 4.2.10). There is considerable complexity on the true effects of the general RMS cut on real detector live time. On one hand, this cut behaves as *specific contaminant cut*, removing events such as long-decay time events, discussed in chapter 7, or excess-light events. In this case, the live time removed is only the live-time accumulated during recording. However, separately, this cut also serves to remove a fraction of the signal; pile-up, exemplified in Fig. 4.10, occurs randomly. Such events would be acceptable if their energy were accurately reconstructed by a standard pulse fit, which is not the case in general. In such cases, a fraction of the signal is removed, and thus the removed live time would be approximately the same as the fraction of events removed. Lastly, this cut removes noise-only triggers not removed by the rate cut. This can happen when noise increases for only short periods, or has not increased enough to cause continual triggering. In both cases, events recorded during the increased noise period will not be removed by the rate cut, but might always be removed by the RMS cut. In this case, the RMS cut is indicative of times where the detector is not able to detect any WIMP event. In these cases, the RMS cut functions effectively as a stability cut, and the live time removed should then be entire live-time between the last good and the next good event, as in the other stability cuts. From the perspective of setting WIMP-interaction limits, the most conservative approach is to treat the RMS cut under the last of these assumptions. From this assumption, the live time removed is, as with control events, the live time accumulated between the last event passing RMS cuts and the next good event. The second effect, that the RMS uses a fraction of the signal, is used as the basis for livetime removal in [111] and [109], though these other effects should be considered in future runs if e.g. time varying noise still exists within data.

Appendix C Probability distributions integrals

This appendix describes the integration of the background models in chapter 6, moving from Eqn. 6.12 to 6.13, and Eqn. 6.14 to 6.15 respectively.

$$\text{If: } \begin{pmatrix} \bar{x}'_B(E) \\ \bar{y}'_B(E) \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} \bar{x}_B(E) \\ \bar{y}_B(E) \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} T_x \\ M_B \end{pmatrix} E = \begin{pmatrix} S_{x'} \\ S_{y'} \end{pmatrix} E$$

Then from Eqn. 6.12, the “flat” background model is given by:

$$\begin{aligned} \rho_{\text{Back}}(x', y') &= A_B \int_0^\infty \frac{1}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x'-\bar{x}'_B(E))^2}{2\sigma_{x'}^2}} e^{-\frac{(y'-\bar{y}'_B(E))^2}{2\sigma_{y'}^2}} dE = \int_0^\infty \frac{A_B}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x'-S_{x'}E)^2}{2\sigma_{x'}^2}} e^{-\frac{(y'-S_{y'}E)^2}{2\sigma_{y'}^2}} dE \\ &= \frac{A_B e^{-\frac{(x')^2}{2\sigma_{x'}^2}} e^{-\frac{(y')^2}{2\sigma_{y'}^2}}}{2\pi\sigma_{x'}\sigma_{y'}} \int_0^\infty e^{-\frac{(\sigma_{y'}^2 S_{x'}^2 + \sigma_{x'}^2 S_{y'}^2)E^2 + 2(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'})E}{2\sigma_{x'}^2 \sigma_{y'}^2}} dE \end{aligned}$$

$$\text{Using } \sigma_{y'}^2 S_{x'}^2 + \sigma_{x'}^2 S_{y'}^2 = \sigma_T^2$$

$$\begin{aligned} &= \frac{A_B}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(\sigma_{y'}^2 (x')^2 + \sigma_{x'}^2 (y')^2)(\sigma_{y'}^2 S_{x'}^2 + \sigma_{x'}^2 S_{y'}^2) + (\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'})^2}{2\sigma_{x'}^2 \sigma_{y'}^2 \sigma_T^2}} \int_0^\infty e^{-\frac{-\sigma_T^2}{2\sigma_{x'}^2 \sigma_{y'}^2} \left(E - \frac{(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'})}{\sigma_T^2} \right)^2} dE \\ &= \frac{A_B}{2\pi\sigma_{x'}\sigma_{y'}} e^{-\frac{(x' S_{y'} - y' S_{x'})^2}{2\sigma_T^2}} \int_0^\infty e^{-\frac{-\sigma_T^2}{2\sigma_{x'}^2 \sigma_{y'}^2} \left(E - \frac{(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'})}{\sigma_T^2} \right)^2} dE \end{aligned}$$

$$\text{Using } \frac{(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'})}{\sigma_T^2} = p'_B \quad \text{and transforming } E = \frac{\sigma_x \sigma_y}{\sigma_T} Q + p'_B, \quad dE = dQ \frac{\sigma_x \sigma_y}{\sigma_T}$$

$$\begin{aligned} \rho_{\text{Back}}(x', y') &= \frac{A_B}{\sqrt{2\pi}\sigma_T} e^{-\frac{(x' S_{y'} - y' S_{x'})^2}{2\sigma_T^2}} \frac{1}{\sqrt{2\pi} \frac{-p'_B \sigma_T}{\sigma_x \sigma_y}} \int_{\frac{-p'_B \sigma_T}{\sigma_x \sigma_y}}^\infty e^{-\frac{Q^2}{2}} dQ \\ &= \frac{A_B}{\sqrt{2\pi}\sigma_T} e^{-\frac{(x' S_{y'} - y' S_{x'})^2}{2\sigma_T^2}} \frac{1}{2} \text{erfc} \left(\frac{-p'_B \sigma_T}{\sqrt{2}\sigma_x \sigma_y} \right) \end{aligned}$$

and from 6.14, the “exponential” background model is given by:

$$\rho_{Comp}(x', y') = A_C \int_0^{E_{Cal}} \frac{N e^{-\frac{(E_{Cal}-E)}{\lambda}}}{2\pi\lambda\sigma_{x'}\sigma_{y'}} e^{-\frac{(x'-S_{x'}E)^2}{2\sigma_{y'}^2}} e^{-\frac{(y'-S_{y'}E)^2}{2\sigma_{x'}^2}} dE$$

Where $\int_0^{E_{Cal}} \frac{N e^{-\frac{(E_{Cal}-E)}{\lambda}}}{\lambda} dE = 1 \therefore N = \left(1 - e^{-\frac{E_{Cal}}{\lambda}}\right)^{-1} \sim 1$ for $E_{Cal} \sim \mathcal{O}(100 \text{ keV})$ and $\lambda \sim \mathcal{O}(10 \text{ keV})$

$$\begin{aligned} \rho_{Comp}(x', y') &= \frac{A_C N e^{-\frac{(x')^2}{2\sigma_{x'}^2}} e^{-\frac{(y')^2}{2\sigma_{y'}^2}} e^{-\frac{E_{Cal}}{\lambda}}}{2\pi\lambda\sigma_{x'}\sigma_{y'}} \int_0^{E_{Cal}} e^{-\frac{-\sigma_T^2 E^2 + 2\left(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'} + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda}\right)E}{2\sigma_{x'}^2 \sigma_{y'}^2}} dE \\ &= \frac{A_C N}{2\pi\lambda\sigma_{x'}\sigma_{y'}} e^{-\frac{-(\sigma_{y'}^2 (x')^2 + \sigma_{x'}^2 (y')^2 + 2\sigma_{x'}^2 \sigma_{y'}^2 \frac{E_{Cal}}{\lambda}) (\sigma_{y'}^2 S_{x'}^2 + \sigma_{x'}^2 S_{y'}^2) + \left(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'} + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda}\right)^2}{2\sigma_{x'}^2 \sigma_{y'}^2 \sigma_T^2}} \int_0^{E_{Cal}} e^{-\frac{-\sigma_T^2}{2\sigma_{x'}^2 \sigma_{y'}^2} \left(E - \frac{\left(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'} + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda}\right)}{\sigma_T^2}\right)^2} dE \end{aligned}$$

Using $\frac{\left(\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'} + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda}\right)}{\sigma_T^2} = p'_C$ and transforming $E = \frac{\sigma_x \sigma_y}{\sigma_T} Q + p'_C$, $dE = dQ \frac{\sigma_x \sigma_y}{\sigma_T}$

$$= \frac{A_C N}{2\pi\lambda\sigma_T} e^{-\frac{-(x' S_{y'} - y' S_{x'})^2}{\lambda^2} \frac{E_{Cal} (\sigma_{y'}^2 S_{x'}^2 + \sigma_{x'}^2 S_{y'}^2) - (\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'}) + \frac{\sigma_{y'}^2 \sigma_{x'}^2}{\lambda^2}}{2\sigma_T^2}} \int_{\frac{-p'_C \sigma_T}{\sigma_x \sigma_y}}^{\frac{(E_{Cal} - p'_C) \sigma_T}{\sigma_x \sigma_y}} e^{-\frac{Q^2}{2}} dQ$$

$\rho_{Comp}(x', y')$

$$= \frac{A_C}{\sqrt{2\pi}\sigma_T} e^{-\frac{(x' S_{y'} - y' S_{x'})^2}{2\sigma_T^2}} \frac{N}{\lambda} e^{-\frac{-(E_{Cal} \sigma_T^2 - (\sigma_{y'}^2 x' S_{x'} + \sigma_{x'}^2 y' S_{y'}))}{\lambda \sigma_T^2}} \frac{\sigma_{y'}^2 \sigma_{x'}^2}{e^{2\sigma_T^2 \lambda^2}} \frac{\text{erfc}\left(\frac{-(E_{Cal} - p'_C) \sigma_T}{\sqrt{2}\sigma_x \sigma_y}\right) - \text{erfc}\left(\frac{p'_C \sigma_T}{\sqrt{2}\sigma_x \sigma_y}\right)}{2}$$

Appendix D Radial penetration depth estimates

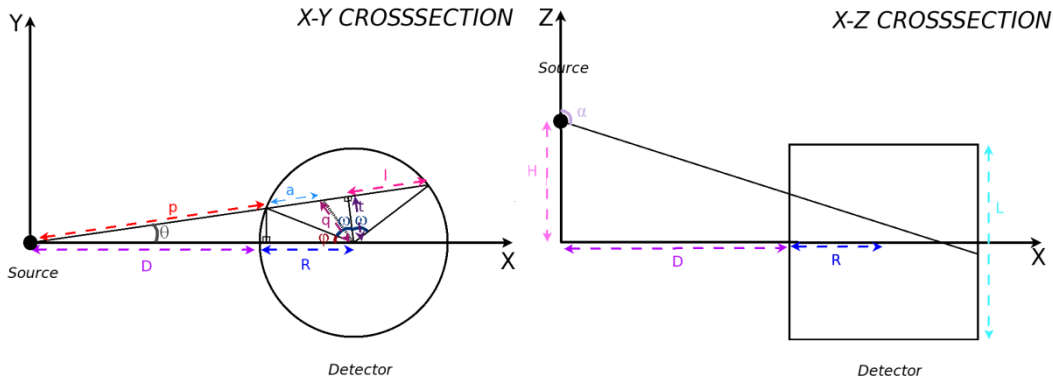


Fig. D.1 Diagram of the numerical integral calculation used to estimate penetration depths here.

This appendix describes how the position of a radioactive γ calibration source (D, H) , the crystal dimensions (R, L) , and the penetration depth of γ rays (λ) can be used to estimate the average radial penetration depth of γ rays in CRESST's target crystals $(R - \bar{q}$ in Fig. D.1). This numerical integral is used for the mean radial penetration depth estimates in chapter 6. D and H are estimated to be 20 cm to 40 cm and -10 cm to $+10$ cm from schematic diagrams of CRESST. First some useful angles are calculated, with $-\frac{\pi}{2} < \theta, \varphi < \frac{\pi}{2}$ everywhere (Fig. D.1).

From the triangle with side lengths $D + R, p$ and R :

$$\frac{R}{\sin(\theta)} = \frac{R+D}{\sin(\pi-\theta-\varphi)} = \frac{R+D}{\sin(\theta+\varphi)} \therefore \varphi = \sin^{-1}\left(\frac{R+D}{R}\sin(\theta)\right) - \theta.$$

From the triangle with side lengths $D + R, p + 2l$ and R :

$$\frac{R}{\sin(\theta)} = \frac{R+D}{\sin(\pi-\theta-(\varphi+2\omega))} \therefore \omega = \frac{\pi-\theta-\varphi-\sin^{-1}\left(\frac{R+D}{R}\sin(\theta)\right)}{2} = \frac{\pi}{2} - \sin^{-1}\left(\frac{R+D}{R}\sin(\theta)\right)$$

Knowing these angles the x-y projections: p , the source-crystal distance along θ ; t , distance of closest approach to the crystal axis in the x-y plane; and l , half the length of the crystal along θ , can then be calculated:

$$p = \left| \frac{R\sin(\varphi)}{\sin(\theta)} \right|$$

$$t = |R \cos(\omega)|$$

$$l = |R\sin(\omega)|$$

To begin, in spherical coordinates (r, θ, α) we then want to know a , the distance travelled from cylindrical crystal edge projected into the x-y plane. This will

allow us to calculate the distance to the azimuthal axis. This is calculated from distance in the x-y plane from the source (origin) minus the minimum source-crystal distance along θ , $a = |r \sin(\alpha)| - p$. The distance towards the azimuthal axis for (r, θ, α) is then $q = +\sqrt{t^2 + (l - a)^2}$.

Armed with this description of the azimuthal distance, the average interaction distance across the entire crystal can then be calculated.

Events are assumed to deposit all of their energy essentially at their interaction position (i.e. no scattering is assumed within the crystal). The source is assumed to be isotropic, such that the intensity of source photons is $\propto 1/r^2$, before taking account of absorption. To take account of absorption, the distance travelled in crystal up to this point, c , needs to be calculated. If photon enter the crystal through the cylindrical surface, this is $c = r - p/|\sin(\alpha)|$. However, if $H > \frac{L}{2}$ or $H < \frac{-L}{2}$, photons may also enter the crystal from the top or bottom surface. The distance travelled through the crystal will then be $c = r - \left| \frac{H-L/2.0}{\cos(\alpha)} \right|$ if $H > \frac{L}{2}$ and $c = r - \left| \frac{H+L/2.0}{\cos(\alpha)} \right|$ if $H < \frac{-L}{2}$, when α is such that a photon entered through the top surface. It should be noted that the fraction entering from the top or bottom surfaces will always be low, given that the source is always further from the crystal radially than any vertical offset. The average distance to the azimuthal axis can then finally be calculated as an integral over the crystal volume in spherical coordinates, with $dV = r^2 \sin(\alpha) d\alpha d\theta dr$. This value must be normalized to the total number of photons observed, such that:

$$\bar{q} = \frac{\iiint q \cdot \text{Intens}(r, \theta, \alpha) dV}{\iiint \text{Intens}(r, \theta, \alpha) dV} = \frac{\iiint q \frac{1}{r^2} e^{-c/\lambda} r^2 \sin(\alpha) d\alpha d\theta dr}{\iiint \frac{1}{r^2} e^{-c/\lambda} r^2 \sin(\alpha) d\alpha d\theta dr},$$

where the penetration lengths λ are calculated as described in section 6.3.4. This integral is then performed numerically, and the mean radial penetration depth is then $R - \bar{q}$.

Appendix E Alternative light detector resolutions

It was noted in chapter 7 that the resolutions and electron recoil band behaviours were also parameterised in [83]. This appendix repeats that parameterisation. In [83] the electron recoil band behaviour was modelled by:

$$L_{\gamma}(E) = \frac{L_1 E}{1 + e^{(-L_e E)}}$$

giving electron recoil band behaviour parameters and light detector resolutions for Zora/SOS23:

| Parameter | Zora/SOS23 |
|----------------------------------|-------------|
| l_1 | 1.068±0.03 |
| l_e (keV ⁻¹) | 0.180±0.007 |
| σ_0 (keV) | 1.1±0.3 |
| σ_1 (keV ^{1/2}) | 0.46±0.03 |
| σ_2 | 0.178±0.004 |

Table E.1 Electron recoil band behaviour parameters for Zora/SOS23 in the commissioning run, from [83].

Light detector resolution for Verena/SOS21 was modelled by three time-separated noise regions with differing light detector resolutions, labelled "High", "Medium" and "Low" noise regions. The electron recoil band behaviour and light detector resolution parameters were:

| Parameter | Verena/SOS21 | | |
|----------------------------------|--------------|-------------|--------------|
| | High Noise | Low Noise | Medium Noise |
| l_1 | 1.021±0.005 | 1.035±0.003 | 1.036±0.002 |
| l_e (keV ⁻¹) | 0.169±0.001 | 0.171±0.008 | 0.22±0.01 |
| σ_0 (keV) | 3.5±0.6 | 1.18±0.37 | 1.26±0.13 |
| σ_1 (keV ^{1/2}) | 1.00±0.15 | 0.72±0.07 | 0.76±0.04 |
| σ_2 | 0.03±0.08 | 0.09±0.01 | 0.106±0.006 |

Table E.2 Electron recoil band behaviour parameters for Verena/SOS21 in the commissioning run. Commissioning run has been split into 3 discrete time periods, labelled by noise volume. From [83].

The live-time distribution between regions is taken from Figure 9.22 of [83] at 20%, 32% and 48% respectively. However, when using these light detector resolutions, it should be noted that there are some differences between the exact cuts and data selection between [83] (47.5 kg-days) and [106] (47.9 kg-days).

References

- [1] J. Oort, *The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems*, Bulletin of the Astronomical Institutes of the Netherlands **6**, 249, 1932.
- [2] F. Zwicky, *Republication of: The redshift of extragalactic nebulae*, General Relativity and Gravitation **41** (1), 207-224, 2009.
- [3] S. M. Kent and J. E. Gun, *The dynamics of rich clusters of galaxies. I - The Coma cluster*, Astron. Journ. **87**, 945-971, 1982.
- [4] H. W. Babcock, *The rotation of the Andromeda Nebula*, Lick Observatory bulletin , University of California Press **48**, 41-51, 1939.
- [5] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, Physics Reports **405** (5-6), 279-390, 2005.
- [6] K. Garrett and G. Duda, *Dark Matter: A Primer*, Advances in Astronomy, 2011.
- [7] J. L. Feng, *Dark Matter Candidates from Particle Physics and Methods of Detection*, Annual Review of Astronomy and Astrophysics **48** (1), 495-545, 2010.
- [8] T. S. Vanalbadu et al., *Distribution of dark matter in the spiral galaxy NGC-3198*, Astrophysical Journal **295** (2), 305, 1985.
- [9] J. Carlstrom, G. Holder, and E. Reese, *Cosmology with the Sunyaev-Zel'dovich Effect*, Ann. Rev. Astron. Astrophys. **40**, 643-680, 2002.
- [10] D. Clowe et al., *A direct empirical proof of the existence of dark matter*, Astrophys. J. Lett. **648**, L109-L113, 2006.
- [11] M. Bradač et al., *Revealing the Properties of Dark Matter in the Merging Cluster MACS J0025.4-1222*, The Astrophysical Journal **687** (2), 959, 2008.
- [12] A. Mahdavi et al., *A Dark Core in Abell 520*, The Astrophysical Journal **668** (2), 806, 2007.
- [13] E. Komatsu et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation*, The Astrophysical Journal Supplement Series **192** (2), 18, 2010.
- [14] WMAP Science Team, <http://wmap.gsfc.nasa.gov/>, 2011.
- [15] S. Hannestad et al., *Neutrino and axion hot dark matter bounds after WMAP-7*, (8), 2010.
- [16] B. Gold et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Galactic Foreground Emission*, The Astrophysical Journal Supplement Series **192** (2), 15, 2010.
- [17] F. Iocco et al., *Primordial Nucleosynthesis: from precision cosmology to fundamental physics*, Physics Reports **472** (1-6), 1-76, 2009.
- [18] S. Cole et al., *The 2dF Galaxy Redshift Survey: Power-spectrum analysis of the final dataset and cosmological implications*, Monthly Notices of the Royal Astronomical Society **362** (2), 505-534, 2005.

- [19] A. P. Cooper et al., *Galactic stellar haloes in the CDM model*, Monthly Notices of the Royal Astronomical Society **406** (2), 744-766, 2010.
- [20] S. Dodelson and L. M. Widrow, *Sterile neutrinos as dark matter*, Physical Review Letters **72** (1), 17-20, 1994.
- [21] A. Kusenko, *Sterile neutrinos: The dark side of the light fermions*, Physics Reports **481** (1-2), 1-28, 2009.
- [22] O. Wantz and E. P. S. Shellard, *Axion cosmology revisited*, Physical Review D **82** (12), 123508, 2010.
- [23] C. A. Baker et al., *Improved Experimental Limit on the Electric Dipole Moment of the Neutron*, Physical Review Letters **97** (13), 131801, 2006.
- [24] G. Raffelt, *Axions*, Space Science Reviews **100** (1), 153-158, 2002.
- [25] R. D. Peccei and Helen R. Quinn, *CP Conservation in the Presence of Pseudoparticles*, Physical Review Letters **38** (25), 1440-1443, 1977.
- [26] F. Della Valle et al., *The PVLAS experiment and its results: probing the quantum vacuum with polarized light and magnetic fields*, Quantum Field Theory under the influence of external conditions (QFEXT09), 147-152, 2010.
- [27] E. Arik et al., *Probing eV-scale axions with CAST*, Journal of Cosmology and Astroparticle Physics **2009** (02), 008, 2009.
- [28] L. D. Duffy et al., *High resolution search for dark-matter axions*, Physical Review D **74** (1), 012006, 2006.
- [29] M. Milgrom, *A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis*, Astrophys. J. **270**, 365-370, 1983.
- [30] J. R. Brownstein and J. W. Moffat, *Galaxy Rotation Curves Without NonBaryonic Dark Matter*, The Astrophysical Journal **636** (2), 721, 2006.
- [31] A. Aguirre, J. Schaye, and E. Quataert, *Problems for Modified Newtonian Dynamics in Clusters and the Ly α Forest?*, The Astrophysical Journal **561** (2), 550, 2001.
- [32] M. Feix, C. Fedeli, and M. Bartelmann, *Asymmetric gravitational lenses in TeVeS and application to the bullet cluster*, Astron. Astrophys. **480** (2), 313-325, 2008.
- [33] C. Alcock et al., *The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations*, The Astrophysical Journal **542** (1), 281, 2000.
- [34] Y. Fukuda and Super-Kamiokande Collaboration, *Evidence for oscillation of atmospheric neutrinos*, Physical Review Letters **81** (8), 1562-1567, 1998.
- [35] L. Bergstrom, *Non-baryonic dark matter: observational evidence and detection methods*, Reports on Progress in Physics **63** (5), 793, 2000.
- [36] J. R. Bond, G. Efstathiou, and J. Silk, *Massive Neutrinos and the Large-Scale Structure of the Universe*, Physical Review Letters **45** (24), 1980-1984, 1980.
- [37] S. Martin, *A Supersymmetry Primer*, <http://arxiv.org/abs/hep-ph/9709356>, 2011.

- [38] G. Abbiendi, *Search for the Standard Model Higgs Boson at LEP*, Physics Letters B **565** (0), 61-75, 2003.
- [39] The ATLAS and CMS Collaborations, *Combined Standard Model Higgs boson searches with up to 2.3 fb⁻¹ of pp collision data at sqrt(s) = 7 TeV at the LHC*, ATLAS-CONF-2011-157, November 14, 2011.
- [40] P. de Jong, *SUSY Searches at ATLAS*, in *Hadron Collider Physics symposium (HCP-2011)* (Paris, France), Nov. 2011.
- [41] G. Servant and T. Tait, *Is the Lightest Kaluza-Klein Particle a Viable Dark Matter Candidate?*, Nuclear Physics B **650** (1-2), 391-419, 2002.
- [42] A. Birkedal-Hansen and J. Wacker, *Scalar Dark Matter From Theory Space*, Physical Review D **69** (6), 065022, 2004.
- [43] M. Graesser, I. Shoemaker, and L. Vecchi, *Asymmetric WIMP dark matter*, <http://arxiv.org/abs/1103.2771>, 2011.
- [44] G. Jungman, M. Kamionkowski, and K. Griest, *Supersymmetric Dark Matter*, Physics Reports **267** (5-6), 195-373, 1995.
- [45] F. Donato, N. Fornengo, and S. Scopel, *Effects of galactic dark halo rotation on WIMP direct detection*, Astroparticle Physics **9** (3), 247-260, 1998.
- [46] J. D. Lewin and P. F. Smith, *Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil*, Astroparticle Physics **6** (1), 87-112, 1996.
- [47] R. W. Schnee, *Introduction to dark matter experiments*, presented at the Theoretical Advanced Study Institute in Elementary Particle Physics, Singapore, 2009.
- [48] R. J. Gaitskell, *Direct detection of dark matter*, Annu. Rev. Nucl. Part. Sci. **54**, 315-359, 2004.
- [49] E. Aprile et al., *Dark Matter Results from 100 Live Days of XENON100 Data*, Physical Review Letters **107** (13), 131302, 2011.
- [50] D. Y Akimov et al., *WIMP-nucleon cross section results from the second science run of ZEPLIN-III*, <http://arxiv.org/abs/1110.4769>, 2011.
- [51] R. Bernabei et al., *New results from DAMA/LIBRA*, The European Physical Journal C - Particles and Fields **67** (1), 39-49, 2010.
- [52] Y. D. Kim et al., *Dark matter search experiment with CsI(Tl) crystal*, J. Korean Phys. Soc. **43**, S40-S44, 2003.
- [53] G. Angloher et al., *Results from 730 kg days of the CRESST-II Dark Matter Search*, <http://arxiv.org/abs/1109.0702>, 2011.
- [54] S. Cebrian et al., *First results of the ROSEBUD dark matter experiment*, Astroparticle Physics **15** (1), 79-85, 2001.
- [55] G. Angloher et al., *Limits on WIMP dark matter using sapphire cryogenic detectors*, Astroparticle Physics **18** (1), 43-55, 2002.

- [56] E. Armengaud et al., *Final results of the EDELWEISS-II WIMP search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes*, *Physics Letters B* **702** (5), 329-335, 2011.
- [57] Z. Ahmed et al., *Dark Matter Search Results from the CDMS II Experiment*, *Science* **327** (5973), 1619-1621, 2010.
- [58] M. Pipe, *Background reduction and spin-dependent limits using Drift - A directionally sensitive dark matter detector*, presented at the J. Phys., Conf. Ser., 2011.
- [59] W. J. Bolte et al., *A bubble chamber for dark matter detection - (the COUPP project status)*, presented at the Topics in Astroparticle and Underground Physics, Bristol, 2005.
- [60] I. G. Irastorza et al., *Status of IGEX dark matter search at Canfranc Underground Laboratory*, <http://arxiv.org/abs/astro-ph/0211535>, 2002.
- [61] B. Majorovits et al., *Status of dark matter search with the HDMS experiment and the GENINO project*, presented at the Sources and Detection of Dark Matter and Dark Energy in the Universe, Berlin, 2001.
- [62] P. Benetti et al., *First results from a dark matter search with liquid argon at 87 K in the Gran Sasso underground laboratory*, *Astroparticle Physics* **28** (6), 495-507, 2008.
- [63] C. E. Aalseth et al., *Search for an Annual Modulation in a P-type Point Contact Germanium Dark Matter Detector*, *Physical Review Letters* **107** (14), 141301, 2011.
- [64] J. Angle et al., *Search for Light Dark Matter in XENON10 Data*, *Physical Review Letters* **107** (5), 051301, 2011.
- [65] A. Marchionni et al., *ArDM: a ton-scale LAr detector for direct Dark Matter searches*, presented at the J. Phys., Conf. Ser., 2011.
- [66] D. N. McKinsey et al., *The LUX dark matter search*, presented at the J. Phys., Conf. Ser., 2010.
- [67] A. Minamino and XMASS Collaboration, *XMASS experiment, dark matter search with liquid xenon detector*, *Nuclear Instruments & Methods in Physics Research Section A* **623** (1), 448-450, 2010.
- [68] F. Giuliani, *Hunting the Dark Matter with DEAP/CLEAN*, *AIP Conf. Proc.* **1200**, 985-988, 2010.
- [69] H. Kraus et al., *EURECA - The future of cryogenic dark matter detection in Europe*, in *Cral-Ipnl: Dark Energy and Dark Matter: Observations, Experiments and Theories*, edited by T. Buchert, D. Stefano, and Y. Copin (E D P Sciences, Cedex A), Vol. 36, pp. 249-255, 2009.
- [70] Z. Ahmed et al., *Results from a Low-Energy Analysis of the CDMS II Germanium Data*, *Physical Review Letters* **106** (13), 131302, 2011.
- [71] The CRESST collaboration, <http://www.cresst.de>, 2011.
- [72] M. Altmann et al., *Results and plans of the CRESST dark matter search*, presented at the X International Symposium on Lepton and Photon Interactions at High Energies, 2001.

- [73] P. Lecoq et al., *Inorganic Scintillators for Detector Systems: Physical Principles and Crystal Engineering*, (Springer, 2006).
- [74] V. B. Mikhailik et al., *Luminescence of CaWO₄, CaMoO₄, and ZnWO₄ scintillating crystals under different excitations*, J. Appl. Phys. **97** (8), 2005.
- [75] V. B. Mikhailik and H. Kraus, *Performance of scintillation materials at cryogenic temperatures*, Phys. Status Solidi B-Basic Solid State Phys. **247** (7), 1583-1599, 2010.
- [76] V B Mikhailik and H Kraus, *Cryogenic scintillators in searches for extremely rare events*, Journal of Physics D **39**, 1181-1191, 2006.
- [77] Michael Kiefer et al., *Glued detectors for the CRESST-II experiment*, Optical Materials **31** (10), 1410-1414, 2008.
- [78] F. Petricca, *Dark Matter Search with Cryogenic Phonon-Light Detectors*, Doctoral Thesis, Technische Universität München, 2005.
- [79] E. Pantić, *Performance of Cryogenic Light detectors in the CRESST-II Dark Matter Search*, Doctoral Thesis, Technische Universität München, 2010.
- [80] P. Huff, *The Detector Parameters Determining the Sensitivity of the CRESST-II Experiment*, Doctoral Thesis, Technische Universität München, 2010.
- [81] M. Janecek and W. W. Moses, *Optical Reflectance Measurements for Commonly Used Reflectors*, IEEE Trans. Nucl. Sci. **55** (4), 2432-2437, 2008.
- [82] M. F. Weber et al., *Giant birefringent optics in multilayer polymer mirrors*, Science **287** (5462), 2451-2456, 2000.
- [83] R. Lang, *Search for Dark Matter with the CRESST Experiment*, Doctoral Thesis, Technische Universität München, 2008.
- [84] F. Pröbst et al., *Model for cryogenic particle detectors with superconducting phase transition thermometers*, Journal of Low Temperature Physics **100** (1), 69-104, 1995.
- [85] I. Bavykina, *Investigation of ZnWO₄ and CaMoO₄ as Target Materials for the CRESST-II Dark Matter Search*, Doctoral Thesis, Ludwig Maximilians Universität München, 2009.
- [86] NATO, *Nonequilibrium Superconductivity, Phonons, and Kapitza Boundaries*, (Plenum, New York, 1981).
- [87] S. Henry et al., *The 66-channel SQUID readout for CRESST II*, Journal of Instrumentation **2** (11), P11003, 2007.
- [88] L. Gonzalez-Mestres and D. Perret-Gallix, *Detection Of Low-Energy Solar Neutrinos And Galactic Dark Matter With Crystal Scintillators*, Nuclear Instruments and Methods in Physics Research Section A **279** (1-2), 382-387, 1989.
- [89] M. Moszyński et al., *Characterization of CaWO₄ scintillator at room and liquid nitrogen temperatures*, Nuclear Instruments and Methods in Physics Research Section A **553** (3), 578-591, 2005.
- [90] T. Jagemann et al., *Measurement of nuclear recoil quenching factors in CaWO₄*, Astroparticle Physics **26** (4-5), 269-281, 2006.

- [91] J. Ninkovic et al., *New Technique for the Measurement of the Scintillation Efficiency of Nuclear Recoils*, 2006.
- [92] K. Nakamura, *Review of Particle Physics*, *Journal of Physics G: Nuclear and Particle Physics* **37** (7A), 075021, 2010.
- [93] Y. F. Wang et al., *Predicting neutron production from cosmic-ray muons*, *Physical Review D* **64** (1), 2001.
- [94] D. M. Mei and A. Hime, *Muon-induced background study for underground laboratories*, *Physical Review D* **73** (5), 2006.
- [95] M. Kimmerle, *Data Analysis in the Direct Dark Matter Search Experiment CRESST and Calculation of the corresponding Limit on the Cross Section of Dark Matter*, Doctoral Thesis, Eberhard Karls Universität Tübingen, 2010.
- [96] H. Wulandari et al., *Neutron flux at the Gran Sasso underground laboratory revisited*, *Astroparticle Physics* **22** (3-4), 313-322, 2004.
- [97] C. Arpesella, *A low background counting facility at Laboratori Nazionali del Gran Sasso*, *Applied Radiation and Isotopes* **47** (9-10), 991-996, 1996.
- [98] P. Belli et al., *Deep underground neutron flux measurement with large BF₃ counters*, *Il Nuovo Cimento A (1971-1996)* **101** (6), 959-966, 1989.
- [99] H. Wulandari et al., *Neutron Background Studies for the CRESST Dark Matter Experiment*.
- [100] S. Scholl, *Neutron Background Simulation for the CRESST-II Experiment*, presented at the 12th International Conference on Topics in Astroparticle and Underground Physics, 2011.
- [101] A. Bassignani et al., *Review of long term radon studies at the gran sasso underground laboratory*, *Radiation Measurements* **25** (1-4), 557-560, 1995.
- [102] J. Kisiel, *Background measurements in the underground labs: Gran Sasso, Modane and Boulby*, presented at the ILIAS, 4th Annual Meeting, 2007.
- [103] R. B. Firestone and V. S. Shirley, *Table of isotopes (8th edition)*, *Nucl. Energy-J. Br. Nucl. Energy Soc.* **36** (6), 401-401, 1997.
- [104] M. Haffke et al., *Background Measurements in the Gran Sasso Underground Laboratory*, *Nuclear Instruments and Methods in Physics Research Section A* **643** (1), 36-41, 2011.
- [105] C. Arpesella, *Background measurements at Gran Sasso Laboratory*, *Nuclear Physics B - Proceedings Supplements* **28** (1), 420-424, 1992.
- [106] G. Angloher et al., *Commissioning run of the CRESST-II dark matter search*, *Astroparticle Physics* **31** (4), 270-276, 2009.
- [107] M. Jeng, *A selected history of expectation bias in physics*, *American Journ. Phys.* **74**, 578, 2005.
- [108] R. Brun and F. Rademakers, *ROOT - An object oriented data analysis framework*, *Nuclear Instruments & Methods in Physics Research Section A* **389** (1-2), 81-86, 1997.
- [109] R. McGowan, *Data Analysis and Results of the Upgraded CRESST Dark Matter Search*, Doctoral Thesis, University Of Oxford, 2008.

- [110] A. J. B. Tolhurst, *OxRop: Data Analysis Software for CRESST-II*, Doctoral Thesis, University Of Oxford, 2007.
- [111] J. M. Schmalzer, *The CRESST Dark Matter Search – New Analysis Methods and Recent Results*, Doctoral Thesis, Technische Universität München, 2010.
- [112] J. Imber, Doctoral Thesis, University Of Oxford, 2010.
- [113] P. Huff, *CRESST collaboration internal note*, 2010.
- [114] P. C. F. Di Stefano et al., *Textured silicon calorimetric light detector*, J. Appl. Phys. **94** (10), 6887-6891, 2003.
- [115] V. B. Mikhailik et al., *Scintillation studies of CaWO₄ in the millikelvin temperature range*, Physical Review B **75** (18), 184308, 2007.
- [116] V. Babin et al., *Decay kinetics of the green emission in tungstates and molybdates*, Radiation Measurements **38** (4-6), 533-537, 2004.
- [117] J. Ninković et al., *CaWO₄ crystals as scintillators for cryogenic dark matter search*, Nuclear Instruments and Methods in Physics Research Section A **537** (1-2), 339-343, 2005.
- [118] M. Kiefer, *CRESST collaboration internal note*, 2008.
- [119] A. Tanzke, *Optimization of Light Detectors for the CRESST Experiment*, Diploma Thesis, Technische Universität München, 2011.
- [120] International Atomic Energy Agency, *Update of X ray and gamma ray decay standards for detector calibration and other applications*, (Vienna). 2007.
- [121] D. Wahl, *Optimisation of light collection in inorganic scintillators for rare events searches*, Doctoral Thesis, University Of Oxford, 2005.
- [122] V. B. Mikhailik D. Wahl, H. Kraus, *The Monte-Carlo refractive index matching technique for determining the input parameters for simulation of the light collection in scintillating crystals*, Nucl. Instr. Meth. **570** (3), 529-535, 2006.
- [123] C. Arnaboldi et al., *CdWO₄ scintillating bolometer for Double Beta Decay: Light and Heat anticorrelation, light yield and quenching factors*, <http://arxiv.org/abs/1005.1239>, 2010.
- [124] A.J.B. Tolhurst, Probationer Research Student Report, University Of Oxford, 2003.
- [125] W. Verkerke, *RooFit*, presented at the Sos 08: In2p3 School of Statistics, Cedex A, 2010.
- [126] M. J. Berger et al., <http://www.nist.gov/physlab/data/xcom/index.cfm>, 2009.
- [127] R. F. Lang et al., *Scintillator Non-Proportionality and Gamma Quenching in CaWO₄*, <http://arxiv.org/abs/0910.4414>, 2009.
- [128] A. Brown et al., *Extending the CRESST-II commissioning run limits to lower masses*, <http://arxiv.org/abs/1109.2589>, 2011.
- [129] R. Lang and W. Seidel, *Search for Dark Matter with CRESST*, 2009.
- [130] R. B. Murray and A. Meyer, *Scintillation response of activated inorganic crystals to various charged particles*, Physical Review **122** (3), 815-&, 1961.
- [131] S. Yellin, *Finding an Upper Limit in the Presence of Unknown Background*, Physical Review D **66** (3), 032005, 2002.

- [132] S. Yellin, *Some ways of combining optimum interval upper limits*, 2011.
- [133] W. Westphal et al., *Characterization of the response of CaWO₄ on recoiling nuclei from surface alpha decays*, *Journal of Low Temperature Physics* **151** (3-4), 824-829, 2008.
- [134] C. Cozzini et al., *Detection of the Natural Alpha Decay of Tungsten*, 2004.
- [135] G. Angloher et al., *Limits on WIMP dark matter using scintillating CaWO₄ cryogenic detectors with active background suppression*, *Astroparticle Physics* **23** (3), 325-339, 2005.
- [136] J. F. Ziegler, M. D. Ziegler, and J. P. Biersack, *SRIM - The stopping and range of ions in matter (2010)*, *Nucl. Instrum. Methods Phys. Res. Sect. B-Beam Interact. Mater. Atoms* **268** (11-12), 1818-1823, 2010.
- [137] W. Westphal, *Development and Characterization of Cryogenic Detectors for the CRESST Experiment*, Doctoral Thesis, Technische Universität München, 2008.