

# **Search for Elastic and Inelastic Dark Matter Interactions in XENON1T and Light Detection for XENONnT**

DISSERTATION

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

vorgelegt der  
Mathematisch-naturwissenschaftlichen Fakultät  
der  
UNIVERSITÄT ZÜRICH

von  
ADAM MICHAEL BROWN  
aus  
dem Vereinigten Königreich

PROMOTIONSKOMMISSION

Prof. Dr. Laura Baudis (Vorsitz)  
Prof. Dr. Ben Kilminster  
Dr. Michelle Galloway

ZÜRICH, 2020

# Abstract

A wide range of astrophysical observations indicates that baryonic matter makes up only a small fraction of the total matter in the Universe. The rest, called dark matter, is over five times more abundant by mass, but its particle content is unknown. One of the most popular candidates is the weakly interacting massive particle (WIMP).

XENON1T was a dark matter detector based on a dual-phase xenon time projection chamber (TPC), located in the Gran Sasso National Laboratory (LNGS). The first part of this thesis describes searches for WIMPs scattering in XENON1T, via both elastic and inelastic interactions. A correction for the spatial dependence of the area of charge signals is developed. Event selection criteria are also shown, in particular a selection based on the fraction of the charge signal seen by the top PMT array. These corrections and criteria helped XENON1T set a world-leading upper limit on the cross-section of spin-independent elastic WIMP-nucleon scattering, reaching  $4.1 \times 10^{-47} \text{ cm}^2$  for  $30 \text{ GeV}/c^2$  WIMPs.

During inelastic scattering, the target nucleus is excited. Here, we consider scattering off  $^{129}\text{Xe}$ . The expected signal is a 39.6 keV photon from the nuclear de-excitation, detected together with the nuclear recoil from the WIMP interaction. By searching for such signals, XENON1T also sets a world-leading upper limit on the cross-section of inelastic scattering, with a minimum of  $3.3 \times 10^{-39} \text{ cm}^2$  for  $130 \text{ GeV}/c^2$  WIMPs.

The second part of this thesis focuses on the photomultiplier tubes (PMTs) for XENONnT, the successor to XENON1T, and their read-out. A total of 368 new PMTs were tested for XENONnT, in both liquid and gaseous xenon, to ensure their suitability for long-term, stable operation. Of these, 105 were tested at the University of Zurich. Particular attention was paid to signs of light emission and afterpulses.

Each PMT has a voltage divider, or base, to power it. The preparation and installation of these bases is presented, as well as the cabling need to connect them to the PMT power supplies and read-out electronics. Two different kinds of cable are used: 30 AWG Kapton-insulated wire for high voltage supplies and PTFE-insulated RG196A/U coaxial cables for PMT signals. Custom-design connectors are used to connect three sections of each cable in order to simplify their installation.

Finally, details of the preparation, installation and testing of the two PMT arrays are shown. These were assembled in a cleanroom at LNGS, where the PMT bases and cables were also attached. Every PMT was tested before installing the arrays in the TPC, which is now in its final position inside the cryostat. The cryostat has already been filled with gaseous xenon and at the time of writing is being cooled down ready for liquid to be filled.

# Acknowledgements

There are a great many people I would like to thank, and without whom I would never have been here writing this thesis. Although not all of them are mentioned here by name, their contributions are greatly appreciated.

Firstly, I would like to thank my supervisor, Prof. Laura Baudis. She gave me the chance to delve into the field of astroparticle physics by offering me a position in her group, and has accompanied and supported me throughout my time as a PhD student. That her office door was always open has not gone unnoticed.

I would also like to thank the other members of my PhD committee, Prof. Ben Kilminster and Dr Michelle Galloway, who have kept up with and supported my work throughout.

I am lucky to have been a member of a wonderful group at the University of Zurich – thank you all. I have found the atmosphere to be very friendly and collaborative and our regular discussions, whether about physics or not, were always a joy. I have spent quite some time in the Irchel Bar over the last four years, gathering inspiration and new ideas – thank you to Frédéric, Kevin and Roman who shared more time with me there than most. Thanks also to all the postdocs from the group who have supported me: Patricia, Shayne, Ale, Shingo and Alex. Chiara and Giovanni have shared many experiences with me, both in Zurich and LNGS. Finally, thanks to Yanina, who has been in the group since I joined. To all of you, our friendships are cherished and I'm sure you'll hear from me again!

Many members of the physics department have always been ready to help, from the two workshops, the administrative team and the other research groups. I would like to mention Daniel Florin of the electronics workshop in particular, whose work on the XENONnT bases and cabling was instrumental to the success of the experiment.

I would like to thank all members of the XENON collaboration for warmly welcoming me. I have enjoyed numerous meetings and spent quite some time at LNGS – to all those who were there with me, thank you for the good times spent together (even with an earthquake and a pandemic).

Thank you to my family for their never-ending support and encouragement, and for visiting me in Zurich many times. And finally, a special thanks to Charlie, for helping me constantly over the past years, putting up with my frequent absences to LNGS, and enduring the many hours of travel between Munich and Zurich.

# Preface

The work presented in this thesis was performed as part of an international collaboration, and thus not everything which I mention is my own work. In the following I give a short summary of my contribution to the work described in each chapter.

- Chapter 3, [Elastic WIMP scattering](#): The search for elastic WIMP scattering was a collaborative effort performed by a large analysis group. My major direct contributions were the S2 ( $x, y$ )-dependent correction (section 3.1.2), the S2 area fraction top cut (section 3.2.1) and the mis-identified S1  $^{83\text{m}}\text{Kr}$  cut (section 3.2.2).
- Chapter 4, [Inelastic WIMP scattering](#): All the analysis described in the chapter on inelastic WIMP scattering is my own work. Nevertheless, this work was based on many other collaboration members' contributions to understanding the detector response.
- Chapter 5, [PMT testing](#): The PMT tests performed at the University of Zurich and described here were performed by myself. I must give credit, however, to Y. Wei and J. Wulf for their work on the MarmotX testing facility and introducing me to it, and to S. Kazama and G. Volta, who worked with me for several of the tests and collaborated with me on the analysis of test data.
- Chapter 6, [XENONnT](#): I was responsible, with G. Volta, for the preparation of XENONnT's bases (section 6.2) and cables (section 6.3). I worked on the cold tests of the PMT array sector (section 6.4.1) as part of a group of several people, notably S. Lindemann. I was involved throughout the assembly of the arrays (section 6.4.2), which was performed by a team comprising the XENON PMT working group and other members of the collaboration.

# Contents

<b>1</b>	<b>Searching for dark matter</b>	<b>1</b>
1.1	Astrophysical evidence	1
1.1.1	Galaxy clusters	1
1.1.2	Galaxies	2
1.1.3	Cosmological effects	4
1.2	Particle candidates	5
1.2.1	The Standard Model	5
1.2.2	Weakly interacting massive particles	6
1.2.3	Beyond WIMPs	7
1.3	The hunt for dark matter	7
1.4	Direct dark matter detection	9
1.4.1	Elastic WIMP-nucleus scattering	9
1.4.2	Inelastic WIMP scattering	13
1.4.3	Annual modulation of direct dark matter signals	14
1.4.4	Selected experimental efforts	14
<b>2</b>	<b>Dark matter detection with liquid xenon</b>	<b>18</b>
2.1	Xenon as a dark matter target	19
2.2	Signals from dual-phase noble element TPCs	21
2.3	S1 and S2 signal detection	24
2.4	Challenges for dual-phase xenon TPCs	26
2.5	The XENON1T TPC	27
2.6	XENON1T: the rest	29
2.7	Calibrating XENON1T	31
<b>3</b>	<b>Elastic WIMP scattering</b>	<b>34</b>
3.1	Event corrections	36
3.1.1	Overview of detector effects which need correcting	36
3.1.2	S2 ( $x, y$ )-dependent correction	37
3.2	Quality cuts	42
3.2.1	S2 area fraction top cut	43
3.2.2	Misidentified S1 $^{83\text{m}}\text{Kr}$ cut	45
3.2.3	Overall cut performance	46
3.3	Signal and background modelling	47
3.4	Results	49

<b>4</b>	<b>Inelastic WIMP scattering</b>	<b>53</b>
4.1	Expected signal . . . . .	55
4.1.1	Uncertainties in the signal model . . . . .	57
4.2	Background modelling . . . . .	57
4.3	Statistical interpretation . . . . .	62
4.3.1	Binning structure optimisation . . . . .	63
4.3.2	Systematic uncertainties . . . . .	64
4.4	Results . . . . .	68
<b>5</b>	<b>PMT testing</b>	<b>72</b>
5.1	Photomultiplier tubes . . . . .	73
5.1.1	Photocathode . . . . .	73
5.1.2	Dynode design . . . . .	73
5.2	The MarmotX PMT evaluation facility . . . . .	75
5.3	The XENONnT PMT testing campaign . . . . .	76
5.3.1	Light emission . . . . .	77
5.3.2	Afterpulses . . . . .	79
5.4	Summary . . . . .	83
<b>6</b>	<b>XENONnT</b>	<b>85</b>
6.1	Overview of XENONnT upgrade . . . . .	85
6.1.1	The XENONnT TPC . . . . .	86
6.2	PMT Bases . . . . .	87
6.2.1	Production . . . . .	90
6.2.2	Cleaning . . . . .	91
6.3	Cables . . . . .	94
6.3.1	Cabling scheme . . . . .	94
6.3.2	Cable screening and procurement . . . . .	97
6.3.3	Connectors . . . . .	99
6.3.4	Cleaning and installation . . . . .	102
6.3.5	Installation of cabling . . . . .	104
6.4	PMT arrays . . . . .	105
6.4.1	Cold test . . . . .	107
6.4.2	Assembling the arrays . . . . .	109
6.5	First light in XENONnT . . . . .	114
<b>7</b>	<b>Concluding remarks</b>	<b>116</b>

# Chapter 1

## Searching for dark matter

It has now been almost a century that we have known that ‘normal’, luminous matter makes up only a small fraction of the total mass of the Universe. Indeed, we shall see that there is a huge variety of astrophysical and cosmological observations which can only be explained by the existence of ‘dark matter’, which has remained invisible until now except through its gravitational effects. Unsurprisingly, the quest to find out what this dark matter is and to determine its properties continues to receive considerable interest and experimental effort.

It is in the context of this worldwide effort searching for dark matter that this thesis is written. After a brief introduction to the topic we will explore the XENON1T experiment, built to directly detect dark matter interactions in its liquid xenon target. We will talk about searches for signatures of both elastic and inelastic scattering of dark matter using data collected during the  $\sim 2$  years that XENON1T was operational. Later we turn towards the future, XENONnT, concentrating on the testing and installation of its light detectors and their connections, and we see that the first light signals have already been recorded by those detectors.

### 1.1 Astrophysical evidence

#### 1.1.1 Galaxy clusters

Galaxy clusters have provided evidence for dark matter since the beginning: Fritz Zwicky first studied the Coma Cluster in 1933 [1]. He determined that there was not enough visible matter to explain the velocity dispersion in the cluster, although he proposed a far higher amount of dark matter (400 times the visible mass) than is accepted today.

Nowadays, most studies compare the distribution of baryonic matter to the overall mass distribution. Typically the baryonic matter distribution is determined using electromagnetic telescopes; usually this means detecting X-rays emitted from hot intracluster gas. The overall mass distribution can either be inferred from the X-ray measurements themselves [2], or measured directly using weak gravitational lensing measurements [3]. A particularly visually appealing example is ‘Pandora’s Cluster’, Abell 2744, seen in figure 1.1. This consists of several clusters which are actively merging, resulting in gas being stripped away from



Figure 1.1: Pandora's Cluster (Abell 2744), with the density of gas (observed using X-rays) in red and the density of mass as determined from gravitational lensing in blue. Credit: X-ray: NASA/CXC/ITA/INAF/J. Merten et al. Lensing: NASA/STScI; NAOJ/Subaru; ESO/VLT Optical: NASA/STScI/R. Dupke

the dark matter halos. Lensing measurements show that much of the mass is concentrated around different regions than the gas (which dominates the baryonic matter) [4]. Further examples include the well-known 'Bullet Cluster' [5], Abell 1689 [6] and Abell 520, dubbed a 'cosmic train wreck' due to its very messy and complicated structure [7, 8].

### 1.1.2 Galaxies

Rotation curves of galaxies provide some of the most intuitive evidence for the existence of dark matter. Indeed, it was when looking at motion within the Milky Way that Jacobus Kapteyn proposed searching for dark matter via its effect on galactic rotation, although he did not find any evidence at the time [9]. In 1932, Jan Oort reached a different conclusion and claimed that the total density in the Milky Way is 1.8 times the density of visible matter\* [10].

Modern efforts tend to compare *rotation curves*, showing the rotational velocity of stars and gas in galaxies as a function of radius from their centre, to the distribution of their visible matter. These curves are often much 'flatter' than expected, by which we mean that the velocity is constant as a function of radius, even well outside the discs containing most visible matter. In the absence of a dark matter *halo* extending out further than the visible matter,

---

\*Oort claimed evidence for 'dark matter', by which he meant 'nebulous' or 'meteoric' baryonic matter rather than a new particle, being present in the galactic disc. Our modern understanding is a little different: the dark matter forms a halo and more recent measurements find no evidence for dark matter being present in the disc itself [11].






Figure 1.2: Rotation curve of NGC 6503. The contributions from the galaxy's gas (dotted line), disk (dashed line) and halo (dot-dashed line) are shown, as is their sum (solid line). Figure from [13].

we would instead expect this to fall proportional to  $1/\sqrt{r}$ . A flat curve suggests a dark matter density proportional to  $1/r^2$  at large radii. Prominent examples of such works include a study of the Andromeda Galaxy by Vera Rubin [12], who pioneered this technique, and [13], in which ten spiral galaxies are studied (see figure 1.2).

While a lot of interest is in the behaviour of galaxies at large radii, it can also be interesting to know the shape of a dark matter halo at the galactic centre. For this, studies of *low surface brightness* galaxies are useful; these are galaxies in which dark matter dominates at all radii. Most observations suggest a so-called *core* at the centre of the halo – an approximately constant density below a certain radius – whereas simulations have tended to prefer a *cusp*, or sharp density peak. This tension is commonly known as the *cusp-core problem*. Recent work has indicated possible solutions, either by including baryons in the simulations or by modifying certain assumptions about dark matter (for example, allowing significant interactions of dark matter with itself). A fairly recent review of the cusp-core problem and various possible resolutions can be found in [14].

Attempts to explain galactic rotation curves without requiring dark matter have proposed modifying gravity instead. Most popular is the theory of modified Newtonian dynamics (MOND) and relativistic extensions of it such as TeVeS [15]; see [16] for a review. However, while MOND can do a reasonable job of explaining rotation curves for most galaxies, it struggles to deal with the cosmological observations described in the next section [17]. Gravitational wave detectors have provided a new way to test such theories. Based on the event GW170817, many theories have been ruled because they predict different speeds for electromagnetic and gravitational radiation, contrary to observations [18]. Furthermore, a galaxy has recently been found which doesn't seem to contain dark matter [19]. This is very difficult to explain with modified gravity, whose effects should apply universally.

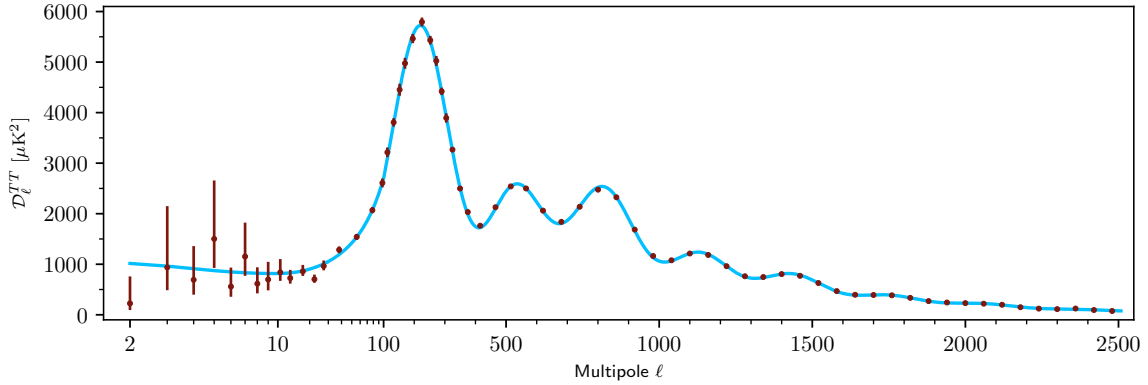


Figure 1.3: CMB temperature power spectrum as measured by Planck, with the best-fit model shown in blue. Figure from [22].

### 1.1.3 Cosmological effects

When talking about cosmological evidence for dark matter, most people immediately turn to the cosmic microwave background (CMB). Once the Universe had cooled to around 3000 K, neutral atoms were able to form from what were previously free electrons and protons – this time is known by the name *recombination*. For the first time, light was able to propagate with only a small probability of scattering off electrons, and has been travelling almost uninterrupted ever since. Nowadays, we can still see this radiation, which was produced as an almost perfect black-body spectrum at the temperature of the Universe at that time. Due to the expansion of the Universe, it has by now been redshifted to a temperature of around 2.7 K, with wavelengths in the microwave region.

By looking at temperature anisotropies in the CMB, we can gain an insight into fluctuations in the matter density at the time of recombination. Recent measurements of the CMB have been made with incredibly high precision. This makes a very effective quantitative test of dark matter theories possible and places strong constraints, in particular on its density [20].

Such comparisons are generally made by looking at the power spectrum of temperature fluctuations, as seen in figure 1.3. From such data, it is possible to find the *relic abundance*, or present-day density, of the three components: baryonic matter, dark matter and dark energy. The height of the third peak in the spectrum is particularly sensitive to the dark matter relic abundance. Current observations, most recently obtained by the Planck Collaboration [21], are explained extremely well by a  $\Lambda$ CDM model, standing for  $\Lambda$  cold dark matter. The  $\Lambda$  refers to dark energy; the cold indicates that the dark matter has non-relativistic speeds at the time of decoupling.

With such a model, Planck finds that the dark matter relic density is  $\Omega_c h^2 = 0.120 \pm 0.001$ , 5.4 times the density of baryonic matter. Here,  $\Omega_c$  is the dark matter density, as a fraction of the critical density, and  $h$  is the reduced Hubble constant (that is, the Hubble constant divided by  $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ).

## 1.2 Particle candidates

Given that the existence of dark matter is almost universally accepted, research focuses on the question of what the dark matter is. There is a huge variety of possible theories containing particles which could make up the dark matter in the Universe. However, some are more popular than others because of their simplicity and/or because they simultaneously solve other existing problems with the Standard Model of particle physics.

### 1.2.1 The Standard Model

The Standard Model of particle physics, which was developed during the second half of the 20th century, describes all currently known particles and their interactions. It contains twelve elementary *fermions*, all with spin  $1/2$ , which can be divided into four groups of particles with very similar properties. Within each group, the three particles differ only in their mass; these are referred to as *generations*. The fermions of the Standard Model are the three charged leptons: electron, muon and tau; their neutral partners: the three neutrinos; and six quarks: three up-like with charge  $+2/3$  (up, charm and top) and three down-like with charge  $-1/3$  (down, strange and bottom). According to the Standard Model, all fermions have mass except the neutrinos, which are massless. However, the discovery of neutrino oscillations around the turn of the millennium [23, 24] means we know now that at least two of the three neutrinos are massive.

As well as the fermions, there are a set of spin 1 *gauge bosons*. These mediate the interactions of the Standard Model: electromagnetism by the massless photon, the weak force by the massive  $W$  and  $Z$  bosons, and the strong force by eight massless gluons. Finally, the spin zero Higgs boson is required to explain why some of the gauge bosons (the  $W$  and  $Z$ ) have mass.

The Standard Model works extremely effectively at explaining particle physics observations. It correctly predicted the existence of the top quark, tau neutrino and Higgs boson, before they were discovered in 1995 [25], 2000 [26] and 2012 [27, 28], respectively. Predictions of the fine-structure constant, which describes the strength of electromagnetic interactions, agree with astonishing precision (better than one part in a billion) with experimental values [29].

Despite its many successes, there are some outstanding issues that are as yet unexplained. As well the fact that no Standard Model particle can make up the dark matter in the Universe, two particular problems which are relevant to the discussion here are the *hierarchy problem* and the *strong CP problem*. The hierarchy problem concerns itself with why the mass of the Higgs boson is so small. Loop-level corrections are expected to contribute an amount to its mass on the order of  $\Lambda_{\text{UV}}$ , an unknown energy scale up to which the Standard Model is valid and after which new physics exists. The only known cutoff for the Standard Model's validity is at the Planck mass,  $2.435 \times 10^{18} \text{ GeV}/c^2$ , many orders of magnitude higher than the mass of the Higgs boson,  $(125.18 \pm 0.16) \text{ GeV}/c^2$  [30]. To explain this difference, the various loop-level corrections would have to almost perfectly cancel out, although there is no reason why

they should.

The strong CP problem concerns the fact that although weak interactions are observed to violate charge-conjugation parity (CP) symmetry, no such observation has been made for interactions of the strong force. This is despite the fact that such CP violation is perfectly allowed in the Standard Model. The amount of CP violation is controlled by a single angle, which could reasonably take any value over a range of  $2\pi$ . Experimental results, such as limits on the electric dipole moment of a neutron [31], rule out strong CP violation at anything more than a very tiny level, requiring this angle to be extremely small. This is an therefore an example of a fine-tuning problem.

These problems and others are summarised concisely in [32], along with a discussion of their relevance for dark matter.

### 1.2.2 Weakly interacting massive particles

The most obvious place to start when seeking a particle which could be dark matter is surely what we already know. Of Standard Model particles, only the neutrino has the required properties: it is neutral, weakly interacting and stable. Unfortunately, neutrinos are simply too light to form a significant fraction of dark matter [33]. Neutrinos are also relativistic – they would form hot, rather than cold, dark matter.

Supersymmetry is a class of independently motivated theories, first conjectured as a way to solve the hierarchy problem in the Standard Model. In general, supersymmetric theories propose a symmetry between integer-spin bosons and half-integer-spin fermions. These contribute equal but opposite loop-level corrections to the Higgs mass, so that they cancel out exactly. In the simplest models, such as the minimal supersymmetric extension to the Standard Model (MSSM), the number of particles is doubled by introducing a *supersymmetric partner* for each Standard Model particle. A conserved parity, known as *R-parity*, is also introduced, which implies that supersymmetric particles can only be created or destroyed in pairs. As a result, whatever is the lightest supersymmetric particle (LSP) must be stable: it cannot decay to a supersymmetric particle since it has too little mass, and cannot decay to a Standard Model particle due to *R-parity*. This particle would be an excellent dark-matter candidate. One of the most popular is the lightest of the four *neutralinos*, a linear combination of the superpartners of the  $W^0$ , B and two Higgs bosons. For more complete treatments of supersymmetry in the context of dark matter, see reviews such as [34].

Neutrinos, LSPs, and other candidate dark matter particles share many properties and as a group are referred to as weakly interacting massive particles (WIMPs) [33]. These have become the most popular class of dark matter candidate, with a large number of experiments dedicated to hunting for them. The WIMP mass,  $m_\chi$ , should be between  $\mathcal{O}(1 \text{ GeV})$  and  $\mathcal{O}(100 \text{ TeV})$  [35].

To be a viable candidate, a particular model must have a working ‘production mechanism’. In other words, the theory must be able to predict the dark matter density which we observe

today. For WIMPs, the mechanism is known as freeze-out. In a nutshell, this says that at some point in the Universe’s past, all elementary particles were in thermal equilibrium with one another. As the Universe cooled and temperature,  $T$ , dropped below the mass of the particle ( $k_B T < m_\chi c^2$ ), WIMPs decoupled from other particles and began annihilating. This continued until the expansion of the Universe became greater than the annihilation rate, at which point annihilation stopped: freeze-out. For standard models, this happens when the temperature of the Universe is roughly  $m_\chi/20$  [36]. After freezing out, its density is only affected by the continued expansion of the Universe.

The relic abundance is smaller the greater the annihilation cross-section is, and is often quoted as  $\Omega_\chi h^2 = 3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1} / \langle \sigma_A v \rangle$  [34, 37], where  $\sigma_A$  is the dark matter annihilation cross-section,  $v$  is its velocity, and the average is over the velocity distribution. A new particle which interacts at the scale of the weak force has more-or-less the right  $\langle \sigma_A v \rangle$  to explain the observed relic abundance. This coincidence is known as the *WIMP miracle* is often taken to justify WIMPs as being the most promising class of candidate for dark matter.

### 1.2.3 Beyond WIMPs

WIMPs are only one of many possible dark matter candidates. One further example is the axion, a natural consequence of Peccei and Quinn’s solution to the strong CP problem [38–40]. Although the axion as proposed by them has been ruled out experimentally, other models remain possible [41]. Considering alternative production methods (other than freeze-out, described above) leads to more classes of dark matter candidates. For example, superweakly interacting massive particles (superWIMPs), could be produced if WIMPs which freeze out in the early Universe later decay into lighter particles [42]. Freeze-in is a proposal whereby dark matter interacts far more weakly than standard WIMPs, such that it is only slowly produced in the hot early Universe and never reaches equilibrium [43]. We will not go into details of these or other models here, and the remainder of this work focuses on the search for WIMPs. The interested reader is referred to reviews such as [33] for more.

## 1.3 The hunt for dark matter

A large number of experiments are ‘searching for’, or trying to ‘detect’ dark matter. These terms are somewhat loosely defined – have we not ‘found’ dark matter in that we know where it is with a reasonable precision? The term usually means attempting to discover non-gravitational interactions which would make it possible to identify the particles that constitute dark matter and their properties. Note that such interactions must exist if the dark matter was in thermal equilibrium in the early Universe. It is common to divide dark matter experiments into three classes: direct detection, indirect detection, and production. In discussing these three, we focus primarily on WIMPs.

Direct detection refers to the observation of signals resulting from dark matter interactions in a particle detector. Since we assume that the Milky Way is inside a dark matter halo,

dark matter particles should continuously be passing through the Earth. If we wait long enough and they are able to interact in some way other than just gravitationally, then one (or hopefully several) of them will eventually scatter in any target we construct. An observation of such scattering would be considered a direct detection of dark matter. The problem is that dark matter, by its nature, interacts only very rarely with normal matter, so we expect only occasional scattering. This means direct detection experiments must strive to minimise any backgrounds that could look like dark matter and disguise a signal, and, where possible, to find ways of distinguishing interactions of dark matter from those of ordinary particles.

The alternative way to ‘see’ dark matter is known as indirect detection. In practice, this means searching for the products of dark matter annihilation in the Universe. The obvious place to look is where there is a lot of dark matter, since the annihilation rate should increase with the square of the dark matter density. This usually means looking at the centre of galaxies, where the dark matter halo is presumably densest, and by far the easiest galaxy to use is the Milky Way. An excess of WIMPs could also be expected at the centre of large astronomical objects, where they could be captured by a combination of the object’s gravitational field and repeated energy loss due to scattering [44]. Experimentally, signatures of dark-matter annihilation can either be spatial or spectral excesses. Some experiments detect gamma-rays, either directly (for example the space-based Fermi-LAT [45]), or from the showers they produce when passing through the atmosphere (for example HESS [46, 47] and the up-coming Cherenkov Telescope Array (CTA) [48]). Others search for neutrinos, examples include Super-Kamiokande [49] and IceCube [50, 51]. Recent reviews of indirect searches for dark matter can be found in [52] and [53].

When talking about production, we normally mean at particle colliders such as the Large Hadron Collider (LHC). In general the signature of dark matter being produced would be some missing energy, since the dark matter particle is extremely unlikely to then interact with the detector. Attempting to probe the nature of dark matter using colliders carries some severe challenges. Probably one of the most significant is that even if a new particle is discovered in this way, it is difficult to prove that it is the same type which naturally occurs in the Universe – it may not even be stable. Furthermore, there are a large number of possible dark matter models and collider searches are inherently very model-dependent, looking for a specific signature present in a tiny fraction of events. It can be difficult to know where to look, and indeed a vast number of searches have been performed at the LHC (a review can be found in [54]). Direct detection experiments tend to be more sensitive to spin-independent scattering, while the two approaches have comparable sensitivity to spin-dependent interactions. For heavier WIMPs, with TeV-scale masses or higher, current colliders run out of energy and direct detection experiments dominate. On the other hand, lighter WIMPs, with masses below a few GeV, would deposit less energy in direct detection experiments’ targets and therefore be difficult to detect, but could still be produced in colliders. The two approaches are therefore somewhat complementary in terms of their strengths.

## 1.4 Direct dark matter detection

The final section of this chapter lays some of the groundwork needed for the remainder of this thesis. Before talking about specific experiments, we will briefly cover some of the theory underpinning direct dark matter searches. The results of this section only concern WIMPs. While we try to remain general, at times we do specifically concentrate on scattering off a xenon target, where it helps provide concreteness.

When trying to detect dark matter, it is possible to either look for its interactions with atomic nuclei or with the electrons around them. Since it carries no electric charge, we can generally expect scattering off the much more massive nuclei to be the most sensitive. However, due to kinematic considerations, this is only true for sufficiently heavy WIMPs. For masses lower than around  $1 \text{ GeV}/c^2$ , the average nuclear recoil energy is too small to be detected, but a much larger, detectable energy would be expected for electronic recoils [55]. Here we focus on the nuclear scattering of GeV to TeV scale WIMPs, to which experiments like XENON1T have the most sensitivity.

### 1.4.1 Elastic WIMP-nucleus scattering

In general, the total rate of WIMP-nucleus scattering can be written as

$$R = Nn\langle v\sigma \rangle \quad (1.1)$$

where  $N$  is the number of atoms in the target,  $n$  is the number density of WIMPs around Earth, and  $v\sigma$  is the product of the speed relative to the detector and cross-section of WIMP interactions, averaged over the distribution of WIMP velocities.

In most experiments the energy imparted on the recoiling nucleus,  $E_R$ , can be measured. We are therefore interested not only in the total rate, but in the differential rate as a function of energy recoil:

$$\frac{dR}{dE_R} = \frac{\rho_0}{m_\chi m_N} \int |\mathbf{v}| \frac{d\sigma}{dE_R} f(\mathbf{v}) d^3\mathbf{v}, \quad (1.2)$$

where  $m_\chi$  is the mass of the WIMP,  $m_N$  is the mass of the nucleus and  $\mathbf{v}$  is the WIMP's incident velocity.

A standard set of dark-matter-related astrophysical parameters are generally assumed to be valid for the purpose of publishing experimental results, known as the standard halo model [33, 56]. This makes it easier to compare results between experiments, since these parameters are common to all experiments and affect their constraints in essentially the same way. Dark matter in the Milky Way is thus assumed to have a flat rotation curve, or density  $\rho \propto r^{-2}$ , and a local density (in the Solar System) of  $\rho_0 = 0.3 \text{ GeV cm}^{-3}$ . Dark matter particles are further assumed to have a Maxwell-Boltzmann velocity distribution characterised by its most probable speed  $v_0 = 220 \text{ km s}^{-1}$ , but with a cutoff at the escape velocity of

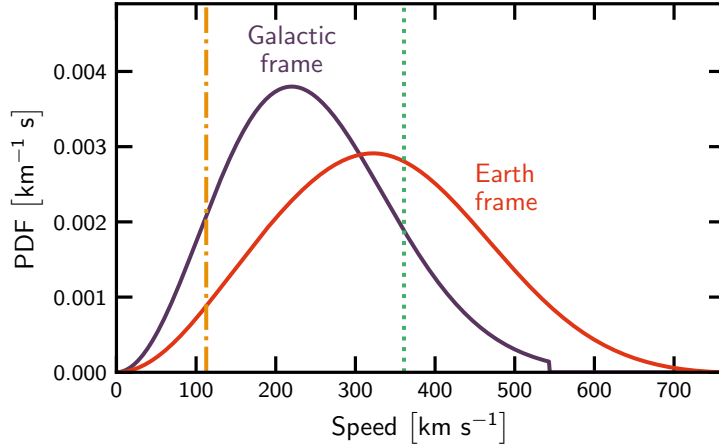


Figure 1.4: Distribution of WIMP speeds according to the standard halo model, in both the Galactic and local frame. The vertical orange dot-dashed line indicates the minimum impact speed needed to obtain a recoil of 7 keV (roughly the detection threshold in XENON1T [57]), and the green dotted line shows the minimum impact speed required for an inelastic recoil to be possible (with  $^{129}\text{Xe}$  as the target nucleus). Both cases are for a  $100 \text{ GeV}/c^2$  WIMP.

the Milky Way,  $v_{\text{esc}} = 544 \text{ km s}^{-1}$ :

$$f(\mathbf{v}) \propto \begin{cases} \exp\left(-\frac{v^2}{v_0^2}\right) & \text{if } |\mathbf{v}| < v_{\text{esc}} \\ 0 & \text{if } |\mathbf{v}| \geq v_{\text{esc}} \end{cases} \quad (1.3)$$

We must also take into account the velocity of the detector, which is orbiting in the Milky Way along with the rest of the Earth. Taking the Sun's orbit to be circular with a speed of  $220 \text{ km s}^{-1}$ , we get to the WIMP impact speed distribution shown in figure 1.4. This introduces an anisotropy in the velocity distribution (although most detectors do not measure the direction).

The standard halo model has significant attractions, notably its simplicity and widespread use. However, the precision of modern measurements has massively improved our understanding of the local dark matter halo. The Gaia satellite is a space telescope primarily designed to measure the position and motion of stars in the Milky Way [58]. Perhaps one of the most significant discoveries is of a second, anisotropic component to the halo, likely the result of the accretion of a large dwarf galaxy around 10 billion years ago [59]. Analysis of a combined dataset from the Sloan Digital Sky Survey and the second Gaia data release suggests there are significantly fewer dark matter particles in the high-speed tail of the distribution, above about  $550 \text{ km s}^{-1}$  [60]. Analysis of Gaia data also provides a handle on the local dark matter density, which has been estimated at  $(0.61 \pm 0.38) \text{ GeV cm}^{-3}$  [61]. In [62] an updated standard halo model is proposed, with a local density of  $(0.55 \pm 0.17) \text{ GeV cm}^{-3}$  amongst other changes.



Non-relativistic kinematics gives the recoil energy as a function of WIMP velocity and scattering angle [63]:

$$E_R = \frac{\mu^2 v^2}{m_N} (1 - \cos \theta), \quad (1.4)$$

in which  $\mu$  is the reduced mass of the WIMP and nucleus,  $\theta$  is the angle which the WIMP is scattered through, and  $v = |\mathbf{v}|$ . From this we can see that the maximum recoil energy, when  $\theta = 180^\circ$ , is  $E_{\max}(v) = 2\mu^2 v^2 / m_N$ .

Experiments usually consider two ways for dark matter to interact with nuclei: either spin-independently (SI) and spin-dependently (SD). The scattering cross-section is therefore the sum of these two, and we can write it as

$$\frac{d\sigma}{dE_R} = \frac{m_N}{2v^2\mu^2} (\sigma_{\text{SI}} F_{\text{SI}}^2(E_R) + \sigma_{\text{SD}} F_{\text{SD}}^2(E_R)), \quad (1.5)$$

where the  $\sigma$ s are the scattering cross-sections for each mode and the  $F(E_R)$  are the nuclear form factors, contributing to the energy dependence. We will discuss these components in turn, for both modes of interaction.

In the spin-independent case, dark matter interacts with all  $Z$  protons and  $(A - Z)$  neutrons present in the nucleus in the same way, regardless of their spin state. We can then write the total cross section in terms of the couplings to protons and neutrons,  $f_p$  and  $f_n$ :

$$\sigma_{\text{SI}} = 4 \frac{\hbar^2 c^2}{\pi} \mu^2 (Z f_p + (A - Z) f_n)^2. \quad (1.6)$$

Often these two couplings are assumed to be the same, and most experiments give their results in terms of a cross-section for scattering off a single nucleon  $\sigma_n$ , by writing

$$\sigma_{\text{SI}} = A^2 \frac{\mu^2}{\mu_n^2} \sigma_n, \quad (1.7)$$

in which  $\mu_n$  is the reduced mass of the WIMP and a single nucleon. This expression makes the  $A^2$ -dependence of the cross-section of spin-independent interactions clear, and an additional scaling with  $\mu^2$ . For heavy WIMPs ( $m_\chi \gg m_N$ ), the reduced mass is approximately the nuclear mass, and therefore also proportional to  $A$ ; in this case the overall dependence is roughly  $A^4$ . This better-than-quadratic increase in cross-section with the mass of the target nuclei means that there are significant advantages to be had using heavy atoms in direct detection experiments.

The form factor describes the energy-dependence, which for spin-independent scattering results from a loss of coherence at higher energies and can be fairly easily understood, at least qualitatively. When the momentum transfer is very small, the wavelength of the mediator particle is very large, so it will have the same value everywhere in the nucleus. This is what makes a coherent interaction possible and gives rise to the  $A^2$ -dependence. As the momentum increases, this wavelength shrinks and when it becomes comparable to the size of




Figure 1.5: Expected differential scattering rate for a  $100 \text{ GeV}/c^2$  WIMP in various common target materials. Heavier targets are better at the lowest recoil energies, but suffer more due to the form factor at increasing energies. Figure from [64].

the nucleus, the coherence gets lost. Since this is sooner for larger nuclei, the interaction rate drops off more quickly with increasing recoil energy, as seen in figure 1.5. At certain values of momentum transfer the contributions of all the nucleons perfectly cancel one another out, and the scattering rate drops to zero. A quantitative treatment can be found in reviews such as [63].

Spin-dependent scattering is a little more complex. In a naïve view, nuclei can only have a single unpaired nucleon of each type (proton and neutron), and the spin-dependent interactions with all others cancel out. Firstly, this means there is no  $A^2$ -dependence like in the case of spin-independent scattering. Secondly, it means that only some isotopes are useful for searching for such interactions.

In reality, chiral two-body currents complicate the picture and mean that there is sensitivity to spin-dependent interactions even in isotopes with an even number of protons/neutrons, albeit much weaker. They are also largely responsible for an energy-dependence in the form factor. It is generally expressed using a structure function  $S_A(E_R)$ :

$$F_{\text{SD}}^2(E_R) = \frac{S_A(E_R)}{S_A(0)}, \quad (1.8)$$

where  $A$  is either  $p$ , for scattering off protons, or  $n$ , for neutrons. Computing these structure function involves detailed nuclear physics calculations; those reported in [65] were used for the elastic scattering search in XENON1T. It is still the case that results are presented in terms of an interaction between a WIMP and a single nucleon to enable comparisons between experiments with different target materials. However, protons and neutrons are

treated separately in this case, giving [66]

$$\frac{d\sigma_{\text{SD}}}{dE_{\text{R}}} = \frac{E_{\text{R}}}{3\mu_p v^2} \frac{4\pi}{2J+1} S_p(E_{\text{R}}) \sigma_p. \quad (1.9)$$

The  $p$  subscripts in this expression make clear that it assumes WIMPs scatter only off protons, and  $\sigma_p$  is the cross-section for scattering off an isolated proton, but we could equally replace  $p$  with  $n$  and refer to neutrons.

### 1.4.2 Inelastic WIMP scattering

It would also be possible for WIMPs to inelastically scatter off xenon nuclei, leaving them in an excited nuclear state. In addition to the nuclear recoil like for elastic scattering, such an interaction would result in an electronic recoil due to photon emission during the subsequent de-excitation.

The kinematics are slightly more complicated than we found before for elastic scattering. There is now a minimum, as well as a maximum, recoil energy. These are given by [67]

$$E_{\text{min/max}} = \frac{\mu^2 v^2}{2m_N} \left( 1 \mp \sqrt{1 - \frac{2E^*}{\mu v^2}} \right)^2, \quad (1.10)$$

where  $E^*$  is the energy of the excited state. We can also see from here that there is a minimum WIMP velocity in order to induce an inelastic recoil:

$$v > \sqrt{\frac{2E^*}{\mu}}. \quad (1.11)$$

Taking the example of inelastic scattering off  $^{129}\text{Xe}$ , which is the topic of chapter 4, the lowest excited state has an energy 39.6 keV above the ground state. This means a 100 GeV WIMP would need an impact velocity of at least  $361 \text{ km s}^{-1}$ . By comparing this to the  $\sim 270 \text{ km s}^{-1}$  which characterises the WIMP velocity distribution in the Milky Way, we see that inelastic scattering can only probe the upper tail of WIMP velocities. This heavily suppresses the inelastic channel relative to the elastic one, especially for heavier WIMPs and for isotopes with higher-lying first excited states [68]. However, in a (hypothetical) detector with a relatively high energy threshold of  $\mathcal{O}(10 \text{ keV})$  the rate of observable inelastic scattering events may dominate, since the de-excitation photon can help to push the detected energy above the threshold [67].

Somewhat related to this is so-called *inelastic dark matter* [69]. This is a modification to the standard WIMP model in which the dark matter particle itself is excited during interactions. Such a model was first conjectured as a way to resolve the tension between the annual modulation signal reported by DAMA [70] and direct constraints on the cross-section of WIMP interactions, from CDMS [71]. The kinematics of inelastic dark matter interactions

would be similar to those described here, but no de-excitation photon would be expected. The minimum WIMP velocity, which depends on the mass of the target nucleus, could ‘hide’ inelastic dark matter from experiments based on lighter nuclei (such as the germanium in CDMS) while still allowing others to detect it (such as DAMA with its NaI target). The topic of this thesis, in chapter 4, is the inelastic scattering of standard WIMPs, where the nucleus is excited, not models in which the dark matter particle is excited.

### 1.4.3 Annual modulation of direct dark matter signals

In section 1.4.1 we briefly mentioned that the Earth’s motion through the Milky Way affects the rate and energy spectrum of WIMP interactions in a direct detection experiment. This motion can be broken down into two components: the Solar System’s motion through the Milk Way, and the Earth’s orbit around the Sun. The latter is much slower, and only becomes interesting when looking at the time behaviour of dark matter signals. When the Earth is moving in the same direction as the Solar System, their velocities add up. This leads to a higher WIMP flux and a higher average WIMP velocity relative to the Earth, meaning dark matter interactions are more likely to be possible. Depending on the minimum WIMP velocity required for an interaction, this leads to an annual modulation in the rate of  $\mathcal{O}(1\%)$ – $\mathcal{O}(10\%)$  [72].

### 1.4.4 Selected experimental efforts

To conclude this brief introduction to direct dark matter detection, we will see a short overview of some of the most important experimental strategies, focusing on those which are sensitive to WIMPs’ interactions. The major experiments in the direct dark matter detection world tend to be one of two types: noble liquid or solid crystalline detectors.

Noble liquid detectors make use of either argon or xenon as the target material. We will explore the way these work and why they are effective in the next chapter, but with a focus on dual-phase xenon TPCs. The choice between xenon and argon comes down to a few key differences. Up to now xenon experiments have (in most regions of the WIMP parameter space) outperformed their argon counterparts, primarily because xenon does not contain any radioactive isotopes, and its higher atomic mass number means a greater scattering rate for the same WIMP-nucleon cross-section. A major disadvantage of xenon is its low natural abundance of 86 parts per billion, which means it is very expensive (although larger argon experiments must use argon depleted in its radioactive isotope  $^{39}\text{Ar}$ , which makes using argon expensive as well). Argon also has the advantage that extremely efficient pulse-shape discrimination is possible between electronic and nuclear recoils.

As well as the choice of target material, there are two different detector technologies in use. The simpler of the two is the single-phase detector, where the usually-spherical, liquid target is surrounded by photon detectors with almost  $4\pi$  coverage. An event’s position can be determined from a combination of the light pattern and arrival time at different detectors,

with an accuracy of  $\mathcal{O}(1 \text{ cm})$ . Single-phase experiments can use pulse-shape discrimination to identify backgrounds. For argon this is very effective, rejecting all but a small fraction,  $\sim 10^{-8}$ , of electronic recoil events [73, 74]. For xenon, on the other hand, the discrimination is substantially weaker (rejecting around 85%–95% of electronic recoils [75]). Prominent examples of single-phase detectors include DEAP-3600 [76], using argon, and XMASS [77], using xenon.

The second technology is the dual-phase time projection chamber (TPC). These are generally cylindrical detectors, filled almost, but not quite, to the top with liquid xenon, with photodetectors generally placed at the top and bottom. The addition of a gas layer, just below the upper layer of photodetectors, makes it possible to detect both light and ionisation signals (for a detailed description of how TPCs work see the next chapter). Importantly for xenon, these allow electronic recoils to be identified from the charge-to-light ratio, albeit far less powerfully than pulse-shape-discrimination in argon. Position reconstruction is slightly more precise than for single-phase detectors, especially in the vertical direction (along the axis of the cylinder). Major dual-phase xenon TPCs include LUX [78], PandaX-II [79] and XENON1T [80], while the argon TPC world is dominated by DarkSide-50 [74].

Noble liquid detectors all suffer from being unable to measure the majority of the energy deposited in a dark matter interaction, which is lost as heat. This shows up most clearly in their mass threshold, since lower mass WIMPs are unable to deposit as much energy and often produce no or little scintillation or ionisation, and only heat. Crystalline detectors, where heat depositions can be recorded from the temperature increase of the crystal following an interaction, are at an advantage here. By simultaneously measuring two forms of energy (heat and either scintillation or ionisation), it is possible to discriminate nuclear and electronic recoils much more effectively than in dual-phase noble TPCs. These detectors use germanium, silicon or  $\text{CaWO}_4$  crystals at cryogenic temperatures (tens of mK), where they have a very small heat capacity. Examples include SuperCDMS [81–83], CDEX [84] and EDELWEISS [85], all measuring heat and charge in germanium/silicon crystals; and CRESST-III, measuring heat and scintillation light in  $\text{CaWO}_4$  [86].

While the two categories of experiment described above dominate current efforts, there are a variety of alternative techniques in use. Bubble chambers contain a superheated liquid in which a WIMP interaction can deposit enough energy to lead to a localised phase transition, leading to a bubble which continues to grow [87]. This is generally a refrigerant such as  $\text{C}_3\text{F}_8$ , which was used by the PICO collaboration to set world leading limits on the spin-dependent WIMP-proton scattering cross-section [88]. Bubble chambers excel at this channel, since they can quite flexibly choose the target compound. They often focus on compounds containing fluorine, whose only stable isotope  $^{19}\text{F}$  is the most sensitive towards such spin-dependent scattering off protons [65].

Alternatively, by using a compound with a higher relative molecular mass, it is possible to optimise bubble chambers' sensitivity to spin-independent scattering, as PICO did with  $\text{CF}_3\text{I}$  [89]. One of the main drawbacks of such bubble chambers is the long deadtime after events, since the liquid must be compressed to return it to its superheated liquid state with no

bubbles. A related technology is the superheated droplet detector, as used by PICASSO [90]. Instead of a single superheated liquid volume, these contain many superheated droplets in a polymer structure. Because only one droplet is affected by each interaction, multiple interactions can be recorded before needing to re-liquefy. In both types of experiment, piezoelectric sensors are used to detect the acoustic signals accompanying bubble formation.

A fairly recent development is the use of silicon charge-coupled devices (CCDs) to search for dark matter. Holes produced by particles interacting in the bulk silicon are drifted by an electric field to the surface where an image is created. This can contain tens of megapixels, giving precise two-dimensional position information. Measuring the amount that holes have diffused by gives information about the third dimension. CCD-based detectors have particularly low energy thresholds, making them very effective tools for searching for light dark matter. The most prominent such experiment is DAMIC, with seven 6 g CCDs [91].

Some experiments are designed to search for the annual modulation which is expected in the rate of dark matter interactions, as described in section 1.4.3. Perhaps the most widely known of these is DAMA/LIBRA, using an array of NaI(Tl) scintillator crystals, with a total of around 250 kg of sensitive mass [92]. After collecting data over 20 full annual cycles (seven of which with the smaller 115.5 kg DAMA/NaI experiment) DAMA/LIBRA claim  $12.9\sigma$  evidence of a modulating dark matter signal [93]. This is in strong tension with limits set by other experiments, which have searched directly for an absolute rate of WIMP interactions rather than a modulation signature, although it is true that the DAMA/LIBRA signal is to some degree more model-independent. The COSINE-100 [94, 95] and ANAIS-112 [96] experiments, both containing around 110 kg NaI(Tl) scintillator, will directly test the DAMA/LIBRA claim using the same technology.

Many future detectors are planned, to build on the successes of the efforts mentioned above. Dual-phase xenon TPCs are destined to get larger, with three new multi-ton detectors being commissioned at the moment and hoping to start taking data this year: PandaX-4T [97], LUX-ZEPLIN [98] and XENONnT [99], which is the subject of the latter part of this thesis. Further in the future, DARWIN will push the target mass to 40 t [100]. The major current argon collaborations are joining forces to work towards the future experiments DarkSide-20k [101] (20 t argon) and eventually Argo (300 t). Crystal-based detectors will also get larger; the EURECA collaboration, combining the expertise of EDELWEISS and CRESST, and SuperCDMS SNOLAB are both planning to operate detectors with targets on the scale of 100 kg [102, 103]. A further phase of EURECA is also planned, with a 1 t target.

A selection of upper limits on the cross-section of elastic, spin-independent WIMP-nucleon scattering, set by a variety of experiments we have discussed in this section, can be seen in figure 1.6.

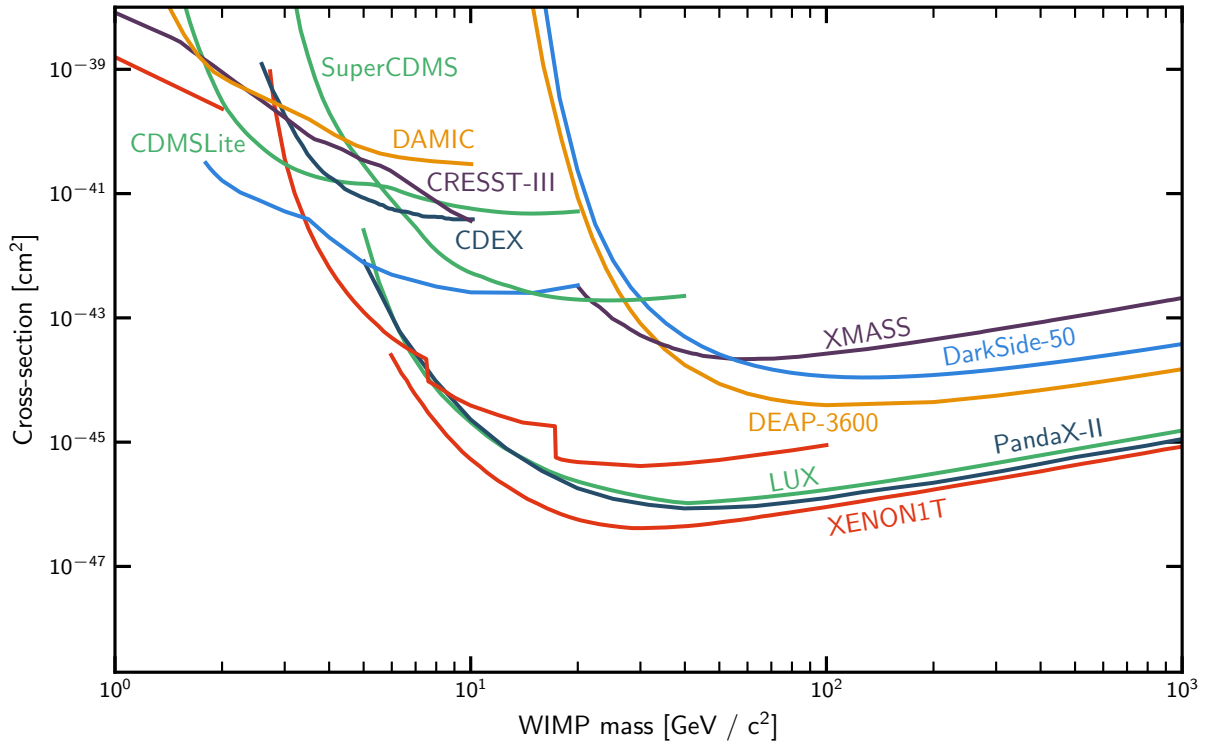


Figure 1.6: Selected experimental results on the cross-section of spin-independent WIMP-nucleon scattering from experiments mentioned in the text. All results are 90% confidence upper limits. The limits shown here were reported in [84] (CDEX), [86] (CRESST-III), [91] (DAMIC), [104, 105] (DarkSide-50), [106] (PandaX-II), [107] (DEAP-3600), [108] (LUX), [109, 110] (SuperCDMS), [57, 111, 112] (XENON1T) and [113] (XMASS).

# Chapter 2

## Dark matter detection with liquid xenon

Having seen an overview of the dark matter detection landscape in the last chapter, we now focus on its search using liquid xenon. This will be the topic of the remainder of this thesis.

In particular, we will explore certain aspects of the XENON dark matter programme. XENON is an international collaboration searching for dark matter interacting in Earth-based liquid-xenon detectors. The collaboration is undertaking a science programme consisting of multiple evolutions of the same concept: the dual-phase xenon time projection chamber (TPC). With each evolution the detector becomes larger, while the rate of backgrounds is reduced. From fairly humble (at least by modern standards) beginnings of XENON10 [114], with its 14 kg instrumented xenon, followed by XENON100 [115] (62 kg instrumented), the collaboration has come a long way to operate XENON1T [80], with 2 t instrumented. The next step in the journey will be XENONnT [99], with 5.9 t. Liquid xenon which will continue to be used to search for dark matter, in the DARWIN observatory (40 t instrumented), for several years to come [100].

This chapter describes the design and operation of dual-phase xenon TPCs in the context of XENON1T, and sets the stage for the following two chapters about the analysis of XENON1T data. Chapter 6 provides an introduction to the next phase of the program, which is under commissioning as this thesis is being written: XENONnT.

At the heart of XENON1T and similar experiments is a dual-phase TPC, immersed in a xenon-filled cryostat. The TPC is where dark matter interactions could be detected. It is only one part of the experiment, however, whose other subsystems include the cooling needed for cryogenic operation, purification of the xenon, and data acquisition. The cryostat is located inside a water tank, which has two functions: firstly to act as a passive shield against external radiation from the lab itself and the surrounding rock, and secondly as an active muon veto.



Table 2.1: Selected properties of xenon.

Property	Value	Unit	Notes / reference
Atomic number	54		
Standard atomic weight	131.290	$\text{g mol}^{-1}$	[116]
Boiling point	165.11	K	At 1 atm
Pressure	1.92	bar	At 177 K
Triple point	161.4	K	[117]
Density	2.86	$\text{g cm}^{-3}$	At 177 K, [118]
Most common isotopes	$^{129}\text{Xe}$ (26.4%), $^{132}\text{Xe}$ (26.9%), $^{131}\text{Xe}$ (21.2%) $^{134}\text{Xe}$ (10.4%), $^{136}\text{Xe}$ (8.9%)		
Fano factor	$\sim 0.05$		[119]
Scintillation wavelength	175	nm	At 168 K, [120]

## 2.1 Xenon as a dark matter target

To begin with, let us consider the detection medium itself: liquid xenon. Xenon is an excellent choice of target material for detecting rare particle interactions, such as those that would be expected from dark matter. It is the fifth noble gas in the periodic table, and the heaviest noble gas with stable isotopes. Table 2.1 contains a few useful basic properties of xenon.

Xenon was given its name by William Ramsay, who discovered it in 1898 together with Morris Travers [121]. The name comes from the Greek word for stranger, ξένος, and refers to the very low natural abundance of xenon in air – less than one part in ten million [122] – which results in its thousands of euros per kilogram price tag.

In its liquid form, xenon has a fairly high density ( $2.86 \text{ g cm}^{-3}$  at the 177 K operating temperature of XENON1T) [118]. This means that it has a high stopping power for external radiation such as gamma rays (see figure 2.1), and also means that even detectors with a large target mass are relatively compact.

Compared to other cryogenic fluids, xenon is liquified at a relatively high temperature, between around 170 K and 180 K depending on the operation pressure (see the phase diagram in figure 2.2). This makes handling it straightforward and enables simple and efficient cooling using liquid nitrogen (although experiments such as XENON1T tend to use active cooling with pulse tube refrigerators for greater control and stability).

One of the most important properties for dark matter searches is the large atomic mass number  $A \approx 131$ . Dark matter, at least when scattering spin-independently, can interact coherently with all the nucleons together, as we saw in chapter 1. Neglecting other factors, this gives a ten-fold enhancement in scattering rates compared to argon, for example.

When reconstructing the energy of particle interactions, it is useful to know the average energy required to produce a single ion-electron pair  $W_i = (15.6 \pm 0.3) \text{ eV}$  in xenon [126],




Figure 2.1: Length scales for photon and electron interactions in liquid xenon as a function of energy. For photons, the attenuation length is shown, using data from [123]. For electrons, the inverse stopping power is shown, using data from [124].




Figure 2.2: Phase diagram for xenon, in the region interesting for liquid-xenon-based experiments. The red dotted lines indicate the operating conditions of XENON1T. Figure data from [125].

and that to produce a single photon,  $W_{\text{ph}} = (13.8 \pm 0.9) \text{ eV}$  [127]. For practical purposes, it is enough to consider a single value of the average energy per quantum produced – often  $W = 13.7 \text{ eV}$  is used [128]. Ignoring the difference between scintillation and ionisation does not result in loss of generality modelling the response. A real detector will always have different efficiencies for detecting each type of signal, which are able to ‘swallow up’ the difference.

The Fano factor,  $F$ , for liquid xenon is rather small. This relates the intrinsic energy resolution (due to fluctuations in the amount of scintillation and ionisation produced) to the energy of an interaction, compared to what would be expected for a purely statistical process:

$$\sigma(E) = \sqrt{FEW}. \quad (2.1)$$

Precise measurements of  $F$  are difficult, partly because any experiment has other factors affecting the resolution, and partly due to the uncertainty on  $W$ , but it is around 0.05 [119].

Finally, and in contrast to alternative noble gases – notably argon – there are no relevant long-lived radioactive isotopes of xenon which contribute significantly to the dark matter search background. The isotope  $^{127}\text{Xe}$ , decaying by electron capture, has a half-life of 36.3 days. All other isotopes are either shorter-lived or stable, with the exception of two whose very slow decays mean they are not a concern\*. These are  $^{124}\text{Xe}$ , decaying through double electron capture with a half life of  $(1.8 \pm 0.05) \times 10^{22}$  years [129], and  $^{136}\text{Xe}$ , decaying through double beta-decay with a half life of  $(2.165 \pm 0.061) \times 10^{21}$  years [130]. Indeed, as we will see, these two isotopes each provide an interesting experimental opportunity, with the former having been observed for the first time by XENON1T.

## 2.2 Signals from dual-phase noble element TPCs

In this and following sections we delve into the design and operation of XENON1T’s TPC. First, we look at the physics which describes interactions in the liquid xenon and their immediate aftermath. Then, in the next section, we talk about how the light and charge signals produced are detected. Finally, we see details of the XENON1T TPC’s design specifically, as well as the other systems which form the experiment and how they are connected.

In the following we occasionally distinguish between two types of interaction: *electronic recoils* and *nuclear recoils*. As the name suggests, the former involves particles interacting with the electrons in xenon atoms – this usually means the particle is an electron or photon interacting electromagnetically. The latter involves interactions with nuclei and in practical terms only those of electrically neutral particles. The name is a slightly inaccurate description of the process, however, because in the low-energy nuclear recoils of interest to dark matter experiments the electrons are (at least mostly) carried along with the nucleus after

---

\*While these isotopes contribute only a negligible background to searches for low-energy elastic WIMP scattering,  $^{124}\text{Xe}$  is relevant for the search for inelastic scattering described in chapter 4;  $^{136}\text{Xe}$  will become relevant for future detectors such as DARWIN [100].

the interaction. In that sense the term ‘atomic recoil’ could arguably be considered more appropriate. Only recently have experiments begun to pay attention to the possibility that the electrons do not immediately follow the nucleus, leading to possible ionisation of the target xenon atom with a corresponding electronic recoil signal: the *Migdal effect* [112, 131].

Energy depositions in the liquid xenon can initially result in a mixture of ionisation, excitation and heat production. Of these, only the first two are detectable by most experiments. One of the unusual properties of noble gases which makes them particularly useful in particle detectors is that they are transparent to their own scintillation light. Instead of de-exciting directly, excited xenon atoms (excitons,  $\text{Xe}^*$ ) rapidly interact with neighbouring xenon atoms to form excited dimers (excimers,  $\text{Xe}_2^*$ ). It is during the decay of these excimers to separate xenon atoms that the characteristic scintillation light is released, with an average wavelength of 175 nm in the vacuum ultraviolet range.

The excimers can form in either a spin-singlet or a spin-triplet state. In the latter case direct decay to the ground state is forbidden, leading to a longer decay time:  $\sim 27$  ns compared to  $\sim 2$  ns for the singlet [132]. Because the ratio of singlet to triplet state excimers produced depends on the type of interaction (nuclear or electronic recoil), this difference could in principle be used to distinguish between them. However, both of these decays are so fast that they cannot be distinguished experimentally, given the temporal resolution of most light sensors and digitisers. In contrast, the triplet state in liquid argon has a lifetime of around a microsecond, and argon-based experiments have shown excellent discrimination is possible based on this timing information, with typically  $\sim 10^{-8}$  leakage fraction of electronic recoils at  $\sim 90\%$  acceptance of nuclear recoils [73, 74].

Turning to the ionisation signal, an externally applied electric field drifts the freed electrons away from the interaction point, to be detected at the top of the TPC (see the next section). The atomic xenon ions quickly combine with a neighbouring xenon atom to form a diatomic ion,  $\text{Xe}_2^+$ . Note that the positive charge is also drifted, in the opposite direction to the free electrons. It is not the ions themselves that drift (their drift speed of  $\mathcal{O}(0.1 \text{ mm s}^{-1})$  is too slow [133]), but holes which are exchanged between xenon atoms (giving a charge drift speed of  $\mathcal{O}(1 \text{ mm s}^{-1})$ ) [134].

A rich set of microphysics describes the relative yields of photons and electrons. Here we describe only the main concepts and avoid lengthy quantitative treatment. For the interested reader, an example of such is that used for the Noble Element Simulation Technique (NEST) model [128, 135, 136], which is widely used in the field of dual-phase xenon TPCs and for dark matter searches in particular. We will talk about the various steps involved in production of scintillation light and electron release in sequence; these are illustrated in figure 2.3.

We already said that energy deposited in the xenon can initially result in ionisation, excitation and heat. Electronic recoils result in negligible heat, whereas this is very important for low energy nuclear recoils. The fraction of energy, often called the quenching factor  $L$ , left for ionisation and excitation after heat loss is described using Lindhard’s treatment and increases as a function of the recoil energy [137, 138].

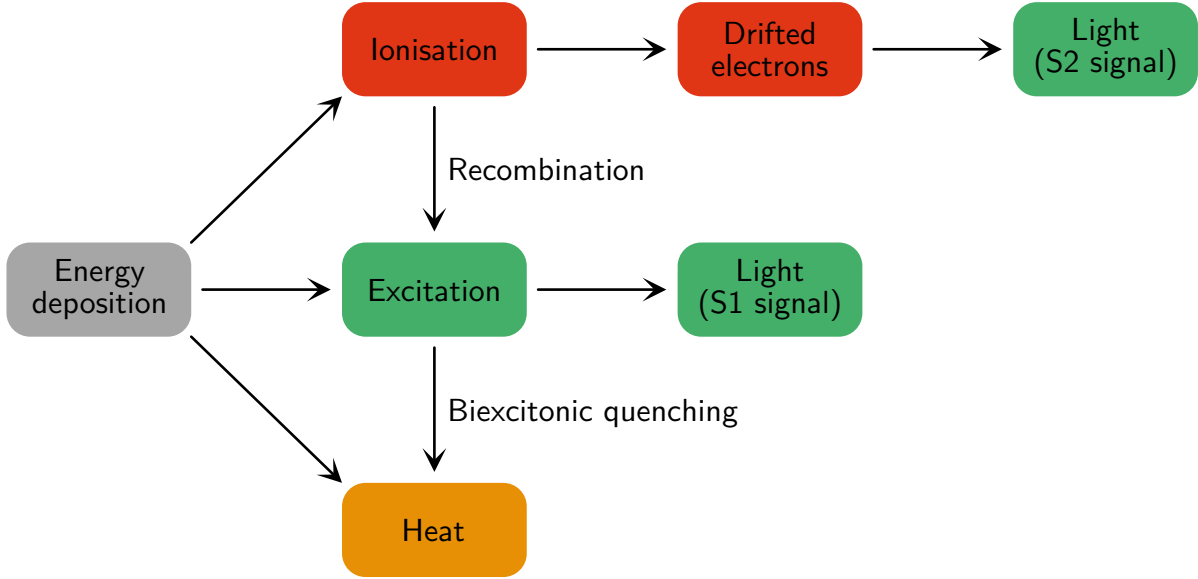



Figure 2.3: Main steps involved in the conversion from deposited energy to S1 and S2 signals in a TPC.

As discussed, the remaining energy  $E_0L$ , where  $E_0$  is the amount of energy deposited, is split between ionisation and scintillation. The initial division between the two is described by the mean exciton-to-ion ratio  $\langle N_{\text{ex}}/N_{\text{i}} \rangle$ . In general this ratio can depend on both the electric field and the deposited energy. That said, for electronic recoils assuming a constant value is sufficient. This has been theoretically predicted to be 0.06 [126] and measured to be as high as  $0.20 \pm 0.13$  [139], albeit with considerable uncertainty. For nuclear recoils, there is a significant energy dependence: using the framework in [136], this ratio is predicted to be  $0.72 \pm 0.06$  for a 5 keV nuclear recoil and  $1.00 \pm 0.07$  at 20 keV, with XENON1T's electric field of 81 kV/cm.

Two more processes can alter the division of energy between the three forms (and introduce energy-dependence also for electronic recoils). The first is the *recombination* of electrons and ions to produce an excimer. This is a process where a diatomic xenon ion meets a free electron, neutralising it and producing an excimer. The amount of recombination depends on the type and energy of the interaction, since this affects the track density and thus the resulting charge distribution. Recombination is also heavily field dependent: at zero field most electrons would eventually find their way to an ion and recombine; as the drift field is increased it rapidly becomes easier to remove them before this happens.

The second, *Biexcitonic quenching*, also known as Penning quenching, is the process of two excitons combining to form an electron/ion pair and an (unexcited) xenon atom:



Thus, either a single photon or electron is produced, depending on whether there is recombination. This reduces the number of quanta produced by one compared to the normal process.




Figure 2.4: Charge (red) and light (blue) yields, as used for XENON1T analysis. The upper two lines are for electronic recoils and the lower two for nuclear recoils. The varying start and end points result from different energy thresholds for each type of recoil, due to the quenching of nuclear recoils. Adapted from [140].

Figure 2.4 shows the modelled charge and light yields as a function of recoil energy, as determined from XENON1T calibration data. These models, taking into account all the effects above, are described in detail in [140] and based on NEST’s treatment, with some adaptation to ensure a good description of XENON1T’s response. The significant quenching for nuclear recoils is clearly visible by the much smaller yields (for both types of quanta).

## 2.3 S1 and S2 signal detection

The general operating principle of a TPC is shown in figure 2.5. The vacuum ultraviolet scintillation photons can be detected directly by specially designed photomultiplier tubes (PMTs). This light signal, originating at the interaction site, is referred to as the S1, since it is the first of two signals. In XENON1T two arrays of PMTs perform this task: one at the top and one at the bottom of the TPC. In total 248 Hamamatsu R11410-21 3” PMTs are used. Chapter 5 goes much deeper into the details of these PMTs, in the context of XENONnT.

Electrons which escape recombination are drifted upwards in the TPC, under an externally applied electric field. In the XENON1T TPC, this field was  $120 \text{ V cm}^{-1}$  and  $81 \text{ V cm}^{-1}$  for the two main periods of data taking, respectively. Two electrode grids placed above and below the liquid-gas interface, with a separation of 5 mm in XENON1T, produce a higher *extraction field* of  $8.1 \text{ kV cm}^{-1}$ , so called because it enables the ‘extraction’ of electrons into the gas phase. As they are accelerated upwards in the gas, these electrons gain sufficient energy between collisions with xenon atoms to excite those atoms and produce a second light signal, known as an S2. The size of the S2 is proportional to the number of electrons extracted and is amplified substantially by the multiple collisions in the gas phase, so that it is much larger than the S1 – in XENON1T around 28 photoelectrons are detected per extracted electron [141].




Figure 2.5: Operating principle of a TPC, showing that an interacting particle ( $X$ ) creates direct scintillation light ( $S1$ ) at the interaction site, and a second light signal ( $S2$ ) via proportional amplification after drifting the electrons to the gas phase.

It is in this two-signal event response that the strength of a dual-phase TPC lies. The energy deposited in an event can be determined from the number of quanta produced. The location of the deposition can be determined in all three dimensions. The depth ( $z$ ) can be found from the time delay between the detection of the  $S1$  and the  $S2$  light, knowing the average speed at which electrons drift in the TPC. This speed, in the range of  $\text{mm } \mu\text{s}^{-1}$ , depends on the field strength [142–144]. However, the dependence is only weak for relatively large fields and eventually saturates. The  $S2$ 's light pattern gives us the horizontal ( $x, y$ ) location: since the electrons are drifted vertically and the  $S2$  light is produced close to the top PMT array, it is very localised on that array.

Position resolution (in the transverse direction) can be affected by several factors. Transverse diffusion of the drifting electron cloud is generally unimportant for the ( $x, y$ ) resolution: typical diffusion coefficients around  $50 \text{ cm}^2\text{s}^{-1}$  result in  $\mathcal{O}(\text{mm})$  size electron clouds [143, 145, 146] (but the centre can still be determined with greater precision). The granularity of the PMT array is comparable to the gap between the amplification region and the array itself, which determines the size of  $S2$ 's light patterns. Random fluctuations play a significant role for small  $S2$ s, up to around 1000 pe, while for larger signals inaccuracies in the optical simulations needed to construct position reconstruction algorithms limit the resolution [141].

The transverse position resolution is generally on the order of a centimetre [141, 147]. On the vertical axis, resolution is a few millimetres and, due to longitudinal diffusion of the electrons while drifting, depends on the depth of an interaction. This resolution is enough to perform fiducialisation for dark matter searches, but not sufficient to resolve particle tracks. This means that the direction from which a particle arrived can't be determined. Because of its potential usefulness, some interest has been shown in possible technologies providing directional information (for high-pressure gaseous xenon detectors, rather than liquid) [148, 149].

As mentioned in the previous section, the charge and light yields are different for electronic and nuclear recoils. It is therefore possible to discriminate between these two types of events. This makes efficient rejection of background events possible, which as we will see predominantly induce electronic recoils.

For an electronic recoil, where quenching is irrelevant, the energy of an interaction can be estimated from a linear combination of the number of photons,  $n_\gamma$ , and the number of electrons,  $n_e$ . This is equivalent to a combination of the S1 and S2 signals:

$$E = W(n_\gamma + n_e) = W \left( \frac{S1}{g_1} + \frac{S2}{g_2} \right). \quad (2.2)$$

The  $g_1$  and  $g_2$  here define the proportionality between the number of photons/electrons produced and the number of photoelectrons detected for each signal. For S1s, this is simply the efficiency of detecting photons; for S2s, it includes the secondary amplification. We have made the common simplification of considering only a single value for the average energy needed to create a quantum,  $W$ . Recall that while the energy per scintillation photon and per ionisation electron are not quite the same, this simplification can be made without loss of precision due to the degeneracy with the parameters  $g_1$  and  $g_2$ .

## 2.4 Challenges for dual-phase xenon TPCs

While electrons are being drifted upwards in the TPC towards the liquid surface, some may be captured and not reach the top. This is possible if they collide with electronegative impurities in the xenon, such as oxygen. It is important to maximise the proportion which reach the surface and are extracted, both to improve resolution by reducing (relative) fluctuations and to lower the detection threshold for low energy events in which few electrons are released to start with. This is why an efficient xenon purification system (described in section 2.6) is essential. Quantitatively speaking, we usually talk about purity in terms of the *electron lifetime*: the average time before electrons are captured by an impurity. The electron lifetime is inversely proportional to the concentration of impurities in the xenon.

In a similar vein, a high efficiency of detecting the scintillation photons produced in interactions is critical. In fact, it is the S1 signal's threshold which dominates the overall event detection threshold for XENON1T, since the light yield becomes very small for events with  $\mathcal{O}(1 \text{ keV})$  energy [141]. Efficiency losses are mostly dominated by the *quantum efficiency* of the PMTs. The PMTs used in XENON1T do have a high quantum efficiency for xenon scintillation light in relative terms; nevertheless this is only around 35% [150]. Since it is pretty much fixed and difficult to improve, it is only realistically possible to make gains by increasing the chance of photons reaching the PMTs in the first place, commonly referred to as the *light collection efficiency*. This mandates the use of polished PTFE for all inward-facing surfaces of the TPC, which has a very high reflectance (around 97%) for vacuum ultraviolet light [151], as well as high-transparency electrodes.

We already discussed the importance of the electron lifetime so that electrons reach the liquid



surface. Of equal importance is that they are drifted vertically upwards, so that information about the position of the interaction is retained. This means that the electric field inside the TPC must be highly uniform and vertical, which is only possible with a large number of *field-shaping rings* arranged outside the TPC. An even worse side-effect of a non-uniform field would be charges being trapped on the PTFE walls if they are drifted outwards. This could bias the recorded S2 size and potentially lead to long-term charge accumulation and worsening electric field. The impact of these effects can be mitigated to some degree using a variety of position-dependent event corrections, which are the subject of section 3.1.

Finally, recall that one of the properties of xenon which make it attractive as a detection medium is the lack of long-lived unstable isotopes. While this is true, other noble gases often present in xenon do have naturally occurring radioactive isotopes. An important example is krypton, containing a small fraction of  $^{85}\text{Kr}$ . In addition,  $^{222}\text{Rn}$  is continuously released by most materials, and discussed in the next section. Unlike most other impurities, noble gases cannot be removed by the purification systems generally used to remove electronegative species. Specialised techniques such as distillation are required instead\*.

## 2.5 The XENON1T TPC

The XENON1T TPC [80] is almost cylindrical in shape, 97 cm tall and 96 cm in diameter. This leads to a total of 2.0 t of liquid xenon inside the TPC, in which particle interactions can be recorded. It is held inside a double-walled vacuum-insulated cryostat with an internal diameter of 1.1 m. An additional 1.2 t of liquid xenon fills the gap between the TPC and inner surface of the cryostat, acting as a passive, non-instrumented shield against external radiation.

In general there are two technologies used to control the level of the liquid surface in dual-phase TPCs. The first, used by LUX [78] amongst others, is a ‘weir’. In this case the level is higher inside the TPC than outside, and spills out at a well-defined position. The level is maintained by constantly filling liquid xenon inside the TPC. The second is a ‘diving bell’, as used by XENON1T [80]. Here, an open-bottomed cylinder of stainless steel covers the top of the TPC. By flushing gaseous xenon into the bell at a slight over-pressure, the liquid level inside is forced down to the required position. This means the level is higher outside the TPC; indeed one of the advantages of such a mechanism is that the outer xenon level can be raised above the top of the bell, to provide  $4\pi$  coverage of passive xenon shielding (although this is not done in XENON1T, leading to a higher background rate at the top of the TPC).

As mentioned above, the two light signals are detected at the top and bottom of the TPC by two arrays containing a total of 248 3” Hamamatsu R11410-21 PMTs. For now it is sufficient to know that when a photon reaches a PMT, it can produce a *photoelectron* (pe) which is

---

\*This problem is even more severe for experiments based on argon time projection chambers. Argon naturally contains  $^{39}\text{Ar}$  at a concentration of  $(8.0 \pm 0.6) \times 10^{-16}$  g/g, which beta-decays with a 268-year half-life. This has prompted the construction of a 350 m tall distillation tower in a coal mine in Sardinia, to improve the isotopic purity of argon for the DarkSide argon-based dark matter experiment [152, 153].




Figure 2.6: Illustration of the XENON1T TPC design, with the main components labelled. Figure from [80].

amplified into a measurable charge signal or *pulse*. A more detailed treatment of the photomultipliers can be found in chapter 5. Much of the TPC’s design is to maximise the chance of photons from the S1 and S2 reaching one of these PMTs. The walls are lined with interlocking diamond-polished PTFE panels, which have excellent reflectivity for xenon scintillation light. The gaps between PMTs in their array are also covered with polished PTFE.

To maintain the drift and extraction fields, a total of five electrodes are used along with 74 field-shaping rings around the outside of the PTFE reflectors. These are visible in figure 2.6. As well as the cathode and the *gate* electrode, which define the drift field, and the anode which defines the extraction field together with the gate, two screening meshes are located below the cathode and above the anode. These are biased at  $-1.5$  kV, similar to the PMT operating voltage, and are intended to prevent the TPC’s electric fields from leaking into the PMT arrays and disrupting their operation.

Because of the need to minimise the background rate of interactions in the TPC, both nuclear and electronic recoils, it is important to use materials with a very low radioactivity. Essentially all materials used in or near the XENON1T TPC were screened for radioactivity, and wherever possible the materials with the lowest radioactivity were selected. Two techniques are used for such screening. The first, gamma-ray spectroscopy, uses germanium crystals in low-background, underground laboratories [154–156] to directly detect gamma-rays from radioactive decay in materials. The exact parts to be used can be measured and, since the decay radiation is observed directly, there are few systematic errors in the measurement. The second, inductively coupled plasma mass spectrometry (ICP-MS), measures the mass number of atoms in a material, allowing the identification of radioactive isotopes [157]. ICP-MS requires only a very small sample, but is destructive and is not always able

to reach the same sensitivity as gamma-ray spectroscopy. For more details of the material screening programme for XENON1T, see [158].

As well as these bulk radioactivity measurements, components with a significant surface area are screened for *radon emanation*. This is the release of  $^{222}\text{Rn}$  from their surface into the surrounding xenon. The  $^{222}\text{Rn}$  is produced from the alpha-decay of  $^{226}\text{Ra}$ , part of the decay chain of  $^{238}\text{U}$ , which is commonly known as the *uranium series*. Most detector materials contain trace amounts of naturally occurring  $^{238}\text{U}$ ; indeed, this one of the main targets of the bulk screening campaign discussed above. Since it is gaseous,  $^{222}\text{Rn}$  can diffuse to the surface of whichever material it is produced in and escape into the surrounding medium. A dedicated screening procedure is used to measure the  $^{222}\text{Rn}$  emanation rates of components [159]. They are filled or flushed with a carrier gas, either helium or nitrogen. The radon is then trapped by passing the helium through activated carbon at liquid nitrogen temperature (radon freezes at  $-71\text{ }^\circ\text{C}$ ). Afterwards, the radon is released from the carbon trap and filled into a proportional counter to measure its activity.

The TPC, in its cryostat, is contained inside a water tank with 9.6 m diameter and 10.2 m height. The primary function of the water is to act as a shield against external radiation from the laboratory and surrounding rock. It is also instrumented and used as a muon veto. XENON1T is located in the Gran Sasso National Laboratory (LNGS) in central Italy, where it is covered by around 1300 m of rock. This is important in order to reduce the flux of cosmic muons, which has been measured as  $(3.41 \pm 0.01) \times 10^{-4} \text{ m}^{-2}\text{s}^{-1}$  [160].

## 2.6 XENON1T: the rest

Up to now we have only discussed the XENON1T TPC. While this can be considered the ‘heart’ of the experiment, it is far from the only component. Other systems which are closely related to the TPC include the data acquisition system, xenon purification, distillation and cryogenics systems, and the xenon recovery and storage facility. The active muon veto, using the water tank, is important for background reduction.

Most of these ancillary systems are located in the three-storey service building, adjacent to the water tank. The top floor contains the cryogenic and xenon purification systems. In order to maintain a high electron lifetime, xenon from the cryostat is continuously recirculated and purified. Liquid xenon is extracted from the bottom of the cryostat through a heat exchanger, used to cool returning xenon gas and massively reduce the cooling power needed [161]. Two loops, operated independently in parallel, pump the warm, gaseous xenon through a SAES PS4-MT50-R hot zirconium getter. This removes many impurities from the xenon by forming chemical bonds between them and the zirconium. Noble gases are not affected.

Three independent cold heads are used to cool and liquify the xenon, contained in the cooling towers on the top floor. Two of these use Iwatani PC-150 pulse-tube refrigerators (PTRs) to provide the cooling power. Each can manage a load of around 250 W at the operating temperature of  $-96\text{ }^\circ\text{C}$ , enough that one is sufficient for operation with one being spare. The




Figure 2.7: The XENON1T watertank and service building. Credit: the XENON collaboration.

third cold head is a backup in case of power loss and is cooled by liquid nitrogen. Enough liquid nitrogen is stored on-site to keep the TPC cold for at least two days without PTR operation.

A vacuum-insulated pipe connects the cryostat with the top floor of the service building, carrying the necessary liquid and gas xenon piping as well as cables for signals and voltage supply from and to the detector.

The bottom floor of the service building houses the xenon storage and recovery system (ReStoX) [162]. This is a vacuum-insulated, liquid-nitrogen-cooled spherical vessel, able to contain up to 7.6 t xenon in gaseous, liquid or solid form. By keeping ReStoX at a cold temperature, quick recovery of liquid xenon from the cryostat is possible in case of emergency through dedicated vacuum-insulated lines.

Also on the bottom floor is the krypton distillation column [163]. This roughly 5 m column was used to reduce the concentration of krypton to  $(0.66 \pm 0.11)$  ppt\* (mol/mol) [57], compared to a concentration on the ppb<sup>†</sup> to ppm<sup>‡</sup> level in commercially available xenon. An extremely low concentration of krypton is important because natural krypton contains  $15 \pm 4$  ppt  $^{85}\text{Kr}$  [164], which undergoes beta-decay with a 10.7 year half-life and a 687 keV endpoint energy. Were it not for the distillation column,  $^{85}\text{Kr}$  would be the dominant electronic recoil

---

\*parts-per-trillion, 1 ppt =  $10^{-12}$

†parts-per-billion, 1 ppb =  $10^{-9}$

‡parts-per-million, 1 ppm =  $10^{-6}$

background and result in a background rate more than 1000 times higher than that observed in XENON1T.

Finally, the middle floor of the service building contains the data acquisition system, as well as a working space for underground operators. PMT signals and power supplies are carried on cables running through the pipe to a breakout chamber on the top floor. From here they are connected to Phillips 776 ten-fold amplifiers and high-voltage supply modules, respectively, on the middle floor. After amplification, signals are recorded with a 100 MHz sampling rate using CAEN V1724 digitisers. Every pulse reaching the threshold, equivalent to around 0.3 photoelectrons, is digitised and saved. Six independent computers are used to read the digitised data which is buffered locally before being stored remotely for further processing. The system can handle a maximum rate of  $3000 \text{ MB s}^{-1}$ , which is equivalent to an event rate of approximately 100 Hz. Further details of the data acquisition system and trigger logic can be found in [165].

The final system to be discussed here is the active muon veto [166], used to identify the interactions of muon-induced neutrons. The water tank in which the cryostat is suspended is instrumented with 84 Hamamatsu R5912ASSY 8" PMTs. These are used to detect the Čerenkov light produced when charged particles pass through the water. Reflective foil on all surfaces of the water tank ensures a high photon detection efficiency despite a relatively small area coverage by the PMTs. Muons directly passing through the water tank can thus be detected with an efficiency of 99.5%. It is also possible for muons to induce neutrons without themselves passing through the water. In this case they can be tagged if particles from their hadronic shower enter the tank, with an efficiency of 70%.

## 2.7 Calibrating XENON1T

Older liquid xenon TPCs were calibrated using a variety of gamma-ray sources positioned near the detector [115, 147]. The size of XENON1T means that gamma-rays originating outside the cryostat are unlikely to make it beyond the outer few centimetres of the TPC. Furthermore, the response to gamma particles is not identical to that to beta particles (the predominant background), as we saw earlier. The response to electronic recoils is therefore determined using an internal beta-decaying source [167]. A chamber in the gas recirculation system containing a source of  $^{228}\text{Th}$  can be flushed with xenon gas which is on its way to the TPC. This decays to  $^{220}\text{Rn}$ , which can mix with the xenon and be carried into the detector.

The  $^{220}\text{Rn}$  alpha-decays with a half-life of 56 s via  $^{216}\text{Po}$  to  $^{212}\text{Pb}$ , after which a mixture of beta and alpha decays take it to the stable isotope  $^{208}\text{Pb}$ . This latter half of the decay chain, with a half-life of 10.6 hours determined by the  $^{212}\text{Pb}$ , provides low energy electronic recoils which mimic the main background for WIMP searches.

The nuclear recoil response is determined using neutron calibrations. Two independent sources of neutrons are used. The first, used throughout the lifetime of XENON1T, is an  $^{241}\text{AmBe}$  source providing a range of neutron energies with a mean of about 4.2 MeV [168].

This is housed in a tungsten collimator, and can be lowered through a hatch in the roof of the water tank to a position near the cryostat, using a system of belts. The second is a deuterium-deuterium plasma fusion neutron generator, which provides neutrons mostly with energies in one of two peaks at 2.2 and 2.7 MeV [169].

In addition to the nuclear recoils induced by neutrons scattering elastically of xenon nuclei, the neutron sources produce a variety of mono-energetic gamma lines. These include lines from the decay of activated isomers at 164 keV ( $^{131m}\text{Xe}$ ) and 236 keV ( $^{129m}\text{Xe}$ ). Inelastic neutron scattering produces further gamma lines at 39.6 keV ( $^{129}\text{Xe}$ ) and 80.2 keV ( $^{131}\text{Xe}$ ), albeit with an additional but small nuclear recoil component – the importance of these lines will become clear in chapter 4.

A final calibration source used for XENON1T is  $^{83m}\text{Kr}$ , a second internal source first used by LUX [170] after being tested in smaller research detectors at the University of Zurich and Yale University [171, 172]. This is produced from its parent  $^{83}\text{Rb}$ , which decays with a half-life of 86.2 days. As with  $^{220}\text{Rn}$ , a chamber containing the parent is flushed with xenon which is then returned to the TPC, carrying the  $^{83m}\text{Kr}$  with it. The metastable isomer  $^{83m}\text{Kr}$  decays to its stable ground state with a half-life of 1.83 hours through two transitions, the first releasing 32.2 keV and the second releasing 9.4 keV. The intermediate stage has a half-life of 157 ns, long enough that the two S1 signals are often separated (although the S2s are usually merged). This two stage decay means that it is easy to identify  $^{83m}\text{Kr}$  events while rejecting background with a very high efficiency, since it is unlikely for two decays to coincidentally be so close in time.

The utility of  $^{83m}\text{Kr}$  calibrations will become apparent in the next chapter. Because its decays always deposit the same energy and are easy to select efficiently, they are useful for identifying position-dependence in the detector’s response.

To conclude this section, let us briefly look at how the gain parameters  $g_1$  and  $g_2$  can be determined. Recall that these are required to relate the energy of an interaction to its charge and light signals. There are at least two common and related methods used for this purpose. The first requires a calibration with a single mono-energetic line. When plotting the distribution of such interactions’ S1 and S2 areas, they generally form a two-dimensional Gaussian, as illustrated in figure 2.8 (left). Fluctuations along the major axis correspond to a clear anti-correlation between the amounts of charge and light produced. Since energy must be conserved, if less scintillation light is produced, more ionisation must take place, and vice versa. The minor axis essentially represents the energy resolution of the detector.

By fitting this response and extracting the major axis it is easy to find  $g_1$  and  $g_2$ : rearranging equation 2.2 shows that the line through it is described by:

$$S2 = g_2 \left( \frac{E}{W} - \frac{S1}{g_1} \right). \quad (2.3)$$

The alternative is to look at several different mono-energetic peaks. When plotting their *charge and light yields*, defined as the amount of charge/light (in detected photoelectrons) per unit energy deposited, these should all fall on a single line, with the  $x$  and  $y$  inter-




Figure 2.8: Determination of the photon detection efficiency and charge amplification factor from mono-energetic calibration lines. Left: illustration of the anti-correlation of charge and light signals for a single mono-energetic line. Right: the light and charge yields of several mono-energetic lines; the intercepts give the two values of interest. Figure adapted from internal XENON1T data.

cepts being  $g_1/W$  and  $g_2/W$ , respectively. This latter method is less affected by detector effects, since the central point of the two-dimensional Gaussians is less likely to be biased by, for example, non-uniform light and charge detection efficiencies. This is the technique which is used by XENON1T, with the data points used for its calibration shown in figure 2.8 (right). The values extracted from fitting these points are  $g_1 = (0.142 \pm 0.002)$  pe/photon and  $g_2 = (11.4 \pm 0.2)$  pe/electron, where the latter considers only the part of S2s observed by the bottom PMT array (as we will see in the next chapter, this is less affected by position-dependent fluctuations).

# Chapter 3

## Elastic WIMP scattering

We have seen an introduction to the XENON1T experiment and explored the technologies used by it, in particular in the time projection chamber. We now turn our attention to the analysis of data which was collected by XENON1T, working towards a constraint on the cross-section of interactions between WIMPs and nuclei. In this chapter we concentrate on elastic interactions, and on the way we will go into some detail about aspects of the lower-level data analysis needed to convert flashes of light detected by PMTs into events with physical values assigned, such as their energy and location. In the next, we turn to inelastic WIMP interactions, in which the target xenon nucleus is excited.

Elastic scattering is the primary focus of many experiments' searches for WIMPs. Limits on the cross-section of spin-independent, elastic WIMP-nucleus scattering in particular could be considered the canonical results obtained by most experiments seeking the direct detection of dark matter. Experiments such as XENON1T are designed specifically with the search for such signals in mind; their ability to discriminate between nuclear recoils, expected from such elastic WIMP interactions, and electronic recoils, expected from most background interactions, boosts their sensitivity to them. By searching for spin-independent interactions, experiments using fairly heavy target nuclei such as xenon gain an additional advantage: due to the ability of WIMPs to interact coherently, the scattering rate increases proportional to  $A^2\mu^2$ , where  $A$  is the number of nucleons and  $\mu$  is the WIMP-nucleus reduced mass, as we saw in section 1.4.1.

Very roughly outlined, the data analysis pipeline used by XENON1T is as follows. First, individual *hits* from PMTs are grouped together according to their detection times, to form *peaks* (borrowing notation from [141], where more details can be found). These peaks should correspond to individual S1 or S2 signals, and are classified as such based on their rise-time. To be classified as a physical signal, S1s must contain hits from at least three PMTs within 100 ns; S2s must contain hits from at least four (with no timing requirement).

The peaks are grouped into *events*, which are defined as the collection of all signals in a window  $2\ \mu\text{s}$  long, centred around a trigger. A single physical interactions in the TPC should lead to an individual event, containing the S1 and S2, but events can also sometimes contain signals from multiple interactions or collections of signals arising from other sources (more on this later).






Figure 3.1: Example of a XENON1T waveform. This example is taken from a  $^{83\text{m}}\text{Kr}$  calibration; the two separate S1s are clearly visible (in purple) while the S2s are merged (in red).

Basic properties of these events can then be computed, mostly information about the signals within them (for example, the S1's and S2's area and rise time or the time gap between them, which is used to compute the depth of the event). The contribution of individual PMTs to each signal is also used to compute other properties; the most obvious is the  $(x, y)$  position of the event, which is extracted from the distribution of S2 light on the top PMT array using a neural network. Technically, the steps up to this point are performed by the PYTHON-based software package PAX (standing for 'Processor for Analyzing XENON') [173] – this part of the analysis chain is often referred to as data processing. An example waveform is shown in figure 3.1.

In latter stages of analysis, we first transform these event properties into more useful variables. Many of the event properties that were computed need correcting, to counteract known systematic effects, as detailed in the first section of this chapter. They can also be combined to calculate more physical variables, for example the energy of an event is estimated using a combination of the S1 and S2 area, as we saw in section 2.7. Now we can try to filter out some of the noise by selecting only those events which are consistent with what would be expected of a physical interaction (see section 3.2). Finally, the 'highest' level of analysis aims to place constraints on fundamental physics by interpreting the events identified up to now. This can be rather complex in its own right, and in this chapter on the search for elastic WIMP scattering is kept fairly concise; in the next one about inelastic scattering there are more details. Here, the focus is on my personal contributions to the analysis: the S2  $(x, y)$ -dependent correction, S2 area fraction top cut and misidentified S1  $^{83\text{m}}\text{Kr}$  cut.

## 3.1 Event corrections

### 3.1.1 Overview of detector effects which need correcting

One of the reasons dual-phase time projection chambers are so widely used is their uniform response. Wherever a particle interacts in the TPC, more or less the same signal is observed on average. However, this uniformity is not perfect: there are small variations depending on the position of the interaction. Most analyses therefore work with variables which have been corrected to account for these non-uniformities. These corrections are the subject of this section, in which we focus in particular on the interaction position and the S1 and S2 areas.

The first effect we consider is due to the fact that the electric field is not perfectly uniform. In particular at the corners of the TPC, external fields can ‘leak’ in and create distortions. This means that especially for events from the lower, outer regions, electrons tend to drift inwards as they move up towards the liquid surface. As the S2 light used for reconstructing their position is produced at the top, bottom events’ interaction locations would be biased inwards. To counter this effect, a correction is applied to convert observed positions (which are the position that the S2 is produced, smeared according to the position resolution) into interaction positions.

The light collection efficiency, defined as the probability of a photon, which has been produced in a particle interaction, reaching a PMT photocathode, is also not uniform throughout the TPC’s volume. Events which occur in the middle of the TPC and just above the cathode have the highest light collection efficiency. To correct for this effect, S1s’ areas are multiplied by a correction factor depending on the position of the interaction. This correction factor is quantified by measuring the size of the S1s in  $^{83\text{m}}\text{Kr}$  calibration data, in which all events should produce the same number of photons on average. The general trend is that the efficiency decreases moving upwards in the TPC and moving closer to the PTFE wall. The reason for this is the increased probability of reflection off the wall, or at the top, off the liquid surface.

Since S2s are all produced at the top of the TPC, near the liquid-gas interface, we can divide systematic effects into two. Some may occur before S2 production, while the electrons are being drifted through the TPC, and should only depend on the depth of the interaction  $z$ . Others affect the S2 light production itself, and therefore depend only on the  $x$  and  $y$  coordinates where the S2 is observed. These two effects are considered separately, and referred to in the following as the *electron lifetime correction* and the *S2 ( $x, y$ )-dependent correction*.

The electron lifetime correction takes into account the fact that as electrons drift through the liquid xenon, some can be captured by electronegative impurities which they encounter. Since these impurities should be homogeneously distributed in the xenon, we expect a constant probability for an electron to be captured per unit distance drifted in the xenon. This results in the electron cloud becoming exponentially smaller as it drifts, and can be characterised by the average ‘lifetime’ of an electron  $\tau_e$ , where of course by lifetime we don’t refer

to its decay, but rather the average time before capture.

The S2 ( $x, y$ )-dependent correction primarily considers the effects of non-uniformity in the extraction field, and PMT-related effects. We characterise these as a multiplicative correction to the S2 area, depending on the position ( $x_{\text{obs}}, y_{\text{obs}}$ ). The subscript ‘obs’ is used in this chapter to refer specifically to the position at which an event is observed, as opposed to where the interaction itself took place, or in other words the position before the field-distortion correction is applied. Further details about this correction can be found in section 3.1.2, below. Putting these effects together, the size of an observed S2 is related to the number of electrons produced in an interaction,  $n_e$ , according to:

$$S2 = g_2 \exp\left(-\frac{\tau}{\tau_e}\right) c_{xy}(x_{\text{obs}}, y_{\text{obs}}) n_e, \quad (3.1)$$

where  $c_{xy}$  is a function describing the ( $x, y$ )-dependence and  $g_2$  is the nominal S2 amplification factor, as defined in section 2.3.

Other corrections are important for different analysis regions but not needed for the results presented in this thesis. In particular, a desaturation correction is used for analyses looking at high energy events, in the MeV region. At these energies, nonlinearities from the PMTs, their voltage-divider circuits and digitisers become important. Details of the correction can be found in [174], and in chapters 5 and 6 we touch on the origins of such saturation effects, in the context of XENONnT.

Of the various types of calibration data taken,  $^{83\text{m}}\text{Kr}$  is particularly suitable for characterising systematic effects which need correcting, thanks to the ease of tagging events as described in the previous chapter. In general, events with the two S1s reconstructed separately and the two S2s merged are used. A typical set of selection criteria, as used for the S2 ( $x, y$ )-dependent correction described below, would be to require that the two S1s are separated by between 0.6 and 2  $\mu\text{s}$  (long enough to mean that the S1s are well separated but short enough to avoid most random coincidences of two events) and that there is a good match between the first S1’s light pattern and its reconstructed position\*.

### 3.1.2 S2 ( $x, y$ )-dependent correction

We can divide the sources of ( $x, y$ )-dependent effects on the S2 area into two categories. The first is those that occur during the production of S2 light: these can be related to the extraction efficiency and the amplification factor of S2s. The other occurs during S2s’ detection, and is related to the S2 light collection efficiency. In practice, this is largely uniform except for drops in regions where PMTs are turned off, so it introduces more localised variations. It only introduces a significant effect for the part of S2s seen by the top PMT array (which we refer to as ‘top S2s’,  $S2_t$ ) since these are localised on few PMTs. In contrast, S2 light is uniformly

---

\*This is quantified using the  $\chi^2$  goodness-of-fit between the expected pattern for an event at that position and the observed pattern.




Figure 3.2: The analytic fit to the bottom (left) and top (right) S2 ( $x, y$ ) variation, shown as contours relative to the maximum. The data used for the fit is shown behind, with colouring representing the area of each event's S2 as detected on the relevant PMT array.

spread out over the whole array by the time it reaches the bottom of the TPC, and, as we will see, bottom S2s ( $S2_b$ ) display no such small-scale variations.

For the purposes of correcting S2s' area, these two scales of variation are considered one at a time, and the S2 light seen by each PMT array is treated separately.

First, we correct for the large scale variation due to S2 production effects. This is assumed to be caused by variations in the strength of the extraction field: the electric field in the gap between the liquid surface and the anode, and in the length of this gap. These variations arise from 'sagging' of the electrodes due to their mutual attraction when charged, or residual tilting of the TPC which means that the electrodes are not perfectly level.

The average area of an S2 detected at position  $(x_{\text{obs}}, y_{\text{obs}})$  due to this variation is assumed to follow a quadratic dependence

$$S2 \propto 1 - b \frac{(x_{\text{obs}} - x_0)^2 + (y_{\text{obs}} - y_0)^2}{r_0^2}, \quad (3.2)$$

where  $r_0$  is the radius of the TPC, 48 cm. The position  $(x_0, y_0)$  is where observed S2s have the greatest area, and  $b$  represents the size of the non-uniformity.

Figure 3.2 shows the size of  $^{83\text{m}}\text{Kr}$  events' S2s (after correcting for electron lifetime only) on the bottom and top array, respectively, and the functional fit to them. It is already clear




Figure 3.3: Variations in the mean S2 on the bottom (left) and top (right) PMT arrays, after applying the electron lifetime correction and correcting for the large-scale variations described analytically. The colour represents the mean S2 area in each bin, as a fraction of the largest bin.

that the average bottom S2 displays no significant small-scale structure, in contrast to the top. Note that in the case of the top array, the quadratic is fit with different parameters even though the effects of electrode sagging are expected to affect both signals equally. Some of the PMT effects have biased the fit: the large number of turned-off PMTs in the upper-right corner means that S2s are generally lower there. This bias doesn't affect the production of a correction, since it will be properly accounted for in the next step, but it does mean that any physical interpretation of the effects described here should be done using bottom S2s only.

In the second step, the remaining small-scale variations are considered by binning the TPC and calculating the mean area in each bin, after correcting using the quadratic formula above. A radial and azimuthal binning structure is used, as seen in figure 3.3. The reader may be wondering why it wouldn't be possible to jump straight to the second stage of the correction and take into account all variations using the binned mean S2. The reason is that most of the variation has almost\* radial dependence; the small scale effects are also a smaller magnitude. In order to get the best correction, we would therefore tend to have bins which are quite small in radius, but which would then have to be quite large in arc length in order to have enough statistics in them. By first taking out this effect, it is possible to choose squarish bins that are better able to represent the small-scale variations. In addition, the radial binning structure would tend to favour a correction depending on the radius from the centre of the TPC, i.e. with  $(x_0, y_0) = (0, 0)$ .

The benefit of this correction can be measured by the spread in area of S2s from  $^{83m}\text{Kr}$  events. With no correction, the standard deviation is 17.0%, after the analytic correction only this becomes 15.7%, and with the full correction 15.0%. Note that the relative effect on the energy

\*It is not quite a radial dependence, since the maximum position  $(x_0, y_0)$  is a few centimetres from the centre of the TPC.

resolution, while not quantified here, will be even better: the S2 resolution contains a constant spread due to the anti-correlation of S1s and S2s, which is removed when calculating events' energy.

In order to simplify the application of the S2  $(x, y)$ -dependent correction, a combined, up-sampled correction map was produced: the binned part of the correction was interpolated on a finer grid, always using a distance-weighted mean of the nearest four bin centres, and combined with the analytic part. As a final remark on the technical implementation, note that the above technique is only useful for measuring relative differences in S2. An overall scaling factor is arbitrary, and would simply result in a different value of the constant  $g_2$  in equation 3.1. For XENON1T we chose to have the correction map normalised such that the average correction factor was unity.

### Physical effects causing S2 variation

In the following we will delve more deeply into the physical effects which are behind the  $(x, y)$ -dependence described above. As mentioned above, this can be divided into three physical origins, whose effects should multiplicatively combine:

$$c_{xy} \propto \epsilon_{\text{ex}} G_{\text{S2}} \epsilon_{\text{PMT}}, \quad (3.3)$$

where  $\epsilon_{\text{ex}}$  is the relative efficiency with which electrons are extracted from the liquid to the gaseous phase,  $G_{\text{S2}}$  is the relative S2 amplification factor, and  $\epsilon_{\text{PMT}}$  is the relative photon detection efficiency. The last term accounts for PMT effects (particularly turned-off PMTs), and is only relevant for top S2 signals.

To study the first two terms, we can look at the variation in the size of single electron signals. A background of single electrons is present throughout the operation of XENON1T; there are a variety of origins including photo-ionisation of metal surfaces in the TPC, delayed extraction of electrons in an S2 signal, and slow release of electrons from impurities such as oxygen [175–177]. For this study, single electrons signals present before the main S1 in  $^{83\text{m}}\text{Kr}$  event windows were used. In principle any data could be used for this, but at the time this study was performed background data was blinded. By searching only the part of each event before the main interaction's S1, any effects due to that interaction should be removed.

The advantage of looking at single electron events is that they are effectively a sample where electrons were extracted with 100% efficiency – a single electron can either be extracted or not, and we only see the events where it is. The  $(x, y)$ -dependence of their area therefore only includes the latter two terms of equation 3.3. We can extract the dependence of  $\epsilon_{\text{ex}}$  from the mean single electron area,  $S_{2\text{SE}}$ , according to:

$$\epsilon_{\text{ex}}(x_{\text{obs}}, y_{\text{obs}}) \propto \frac{c_{xy}(x_{\text{obs}}, y_{\text{obs}})}{S_{2\text{SE}}(x_{\text{obs}}, y_{\text{obs}})}. \quad (3.4)$$

In fact, we will use the number of hits in the single electron peak,  $n_{\text{hits}}$ , instead of its area,  $S_{2\text{SE}}$ , since it is less affected by fluctuations. With a mean of only 22.0 hits per single electron

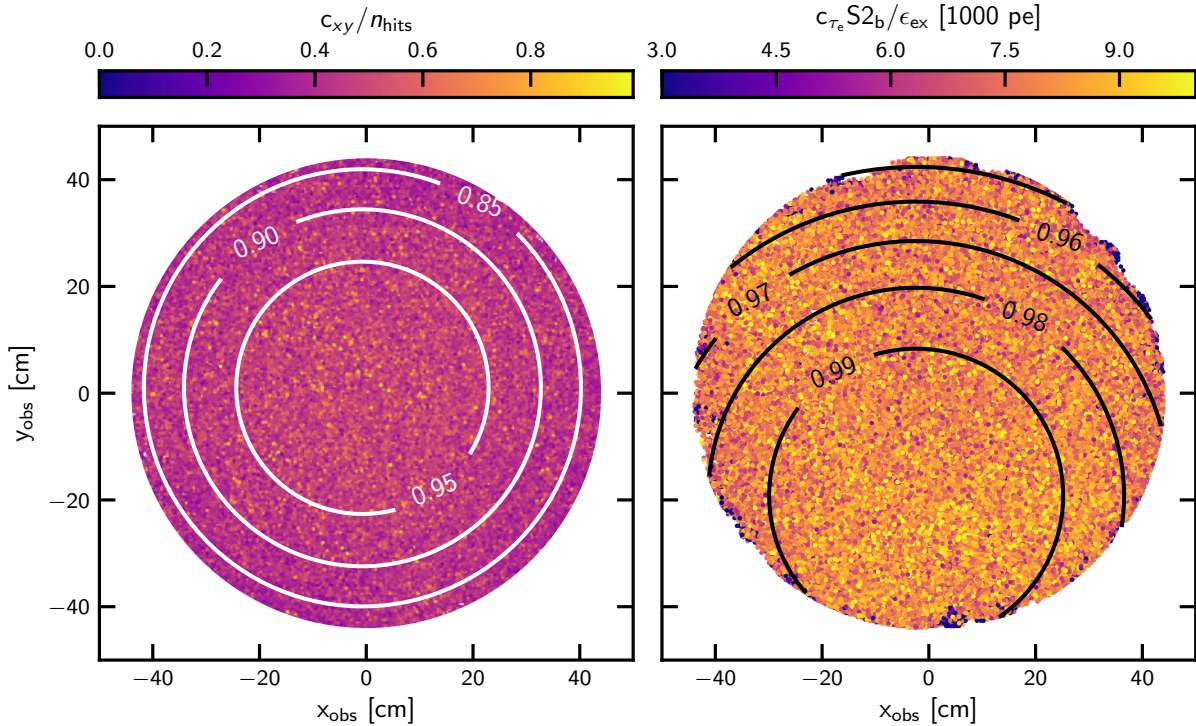


Figure 3.4: Breakdown of the large-scale S2 ( $x$ ,  $y$ ) dependence into variation of electron extraction efficiency (left) and amplification factor (right). The data shown by colour in the latter case is the bottom S2 area divided by the estimated extraction efficiency at that point, from the fit shown on the left.

shared across all the PMTs and over around 200 ns, it is very unlikely for any PMT to detect more than one hit [178].

This is seen in figure 3.4 (left). There are two interesting things to see here: firstly, the size of the effect is similar to the overall variation that we saw in figure 3.2; secondly, the centre-point of the behaviour is very close to the centre-point of the TPC. This is what we may have expected: the extraction efficiency should depend purely on the electric field at the liquid-gas interface. This is influenced much more by sagging of the electrodes than tilting, which would not change their separation.

We are now also in a position to look at the variation of the S2 amplification factor  $G_{S2}$ . For this we divide the area of bottom S2s from  $^{83m}\text{Kr}$  events by the extraction efficiency found above, leaving this S2 amplification as the only variation. This can be seen in figure 3.4 (right), where it is clear that the effect is much smaller (up to around 5% compared to 20% for the full variation) and very off-centre. Again, this might be expected: tilting will result in a larger region being available for S2 amplification and thus a higher gain.

Finally, we can try to understand the origin of the PMT-related variations. Figure 3.5 shows the ratio between the top and bottom S2 signals of  $^{83m}\text{Kr}$  events, which should be proportional to  $\epsilon_{\text{PMT}}$ . This is fit with a simple empirical function, starting with a uniform efficiency




Figure 3.5: Description of the S2 variation by turned off PMTs. The data plotted is the ratio of events' top and bottom S2 areas. The contours are a simple model, with parameters fit to this data, as described in detail in the main text.

and subtracting a two-dimensional Gaussian for every turned-off PMT:

$$\epsilon_{\text{PMT}} = 1 - \sum_{i \in O} \alpha \exp\left(-\frac{-(x - x_i)^2 + (y - y_i)^2}{2r_{\text{out}}^2}\right) - \sum_{i \in I} \beta \exp\left(-\frac{-(x - x_i)^2 + (y - y_i)^2}{2r_{\text{in}}^2}\right). \quad (3.5)$$

Here  $O$  is the set of PMTs in the outer ring of the top array which are turned off and  $I$  is the set of all other turned-off PMTs. The outer PMTs are considered separately because reflections off the TPC wall can affect the region which they are sensitive to. The parameters  $\alpha$ ,  $\beta$ ,  $r_{\text{out}}$  and  $r_{\text{in}}$  are free in the fit.

Visually, this is a good fit to the variations seen in XENON1T data. Residual variations are up to around 5%, and may be explained by a more complex model; this is however beyond the scope of this work. We can nevertheless conclude that these smaller-scale variations are related to the reduced efficiency of collecting photons where PMTs are turned off.

## 3.2 Quality cuts

Not all events algorithmically identified are genuine particle interactions within the TPC. They may come from random coincidences of spurious S1-like and S2-like signals, such as PMT dark counts or the single electron background. In addition, some events are reconstructed badly. Typical examples include events very close to the wall, where the electrons can be trapped on the PTFE surface, resulting in a smaller-than-expected S2. For events that occur near an area with several turned-off PMTs, the position may be incorrectly reconstructed due to the limited knowledge of their light pattern. These events are removed from the search data using a variety of quality criteria, which are the subject of this section.



Broadly speaking, most quality criteria can be identified as belonging to one of three categories. The first is time-based, and rejects events which occur close in time to a muon interaction in the water tank or when the DAQ is ‘busy’ and unable to fully record data (in which case part of the event may not have been fully recorded). Other criteria identify unusual properties of either the S1 or the S2, and reject those events. Some prominent examples are the S1-based and S2-based *single scatter* cuts, which reject events where a second S1-like or S2-like peak is present and there may be some confusion over which should be paired; cuts on the fraction of each peak seen by the top PMT array; and the *S2 width* cut, which rejects events with an unusually long S2 peak.

Data quality cuts must be designed to have a high *acceptance*, the rate of good events being kept, which should ideally be one. It is also important to know the acceptance well, in order to know how many signal events one would expect for a given interaction cross-section. They should also have a reasonably high *rejection power*, their ability to remove bad events. It is, however, less important to know the rejection power very precisely, provided any significant populations of bad events which are not removed by any cut can be satisfactorily modelled.

Here we will explore two examples: the so-called S2 area fraction top cut and a cut targeting  $^{83m}\text{Kr}$  events with a misidentified S1. A detailed discussion of the full set of quality cuts is beyond the scope of this work; a more comprehensive treatment can be found in [141].

### 3.2.1 S2 area fraction top cut

Various cuts are designed to remove events which are physical but either badly reconstructed or for some reason cannot be dark matter events. An example of these is the *S2 area fraction top (AFT) cut*. This is based upon the corrected fraction of the S2 seen by the top PMT array, or cS2 AFT:

$$\text{cS2 AFT} = \frac{\text{cS2}_t}{\text{cS2}}. \quad (3.6)$$

This parameter is particularly suitable for identifying *gas events*: interactions which occur in the gaseous xenon, generally above the anode. Such interactions can produce both S1 and S2 signals, like liquid events do. Any electrons produced are accelerated downwards towards the anode, and can produce an S2 by secondary amplification in this region. The electric field above the anode is different from the usual extraction field and the distance over which S2 amplification takes place depends on the interaction position. The xenon microphysics describing the ratio of electrons to photons produced is also different in gas and liquid. Therefore, the S2/S1 ratio can vary greatly between events and is not well modelled, unlike for usual liquid events. It is important to identify these events and remove them from the science data used to search for dark matter, to prevent them appearing as spurious WIMP signals.

The S2 light from gas events is produced in the region above the anode, and potentially starting from near the top screening mesh. It is therefore more likely to be detected by the top array, so the cS2 AFT will in general be larger than for normal events.




Figure 3.6: Calibration data used to define the S2 area fraction top cut, and the cut definition. The blue shading represents the density of calibration events (from both  $^{220}\text{Rn}$  and  $^{241}\text{AmBe}$  calibration data), and the vertical histograms show the distribution of cS2 AFT within a series of cS2 slices. The 0.5th and 99.5th percentiles of these distributions are shown as black crosses; the cut is defined using a fit to these points as shown in red.

Normal S2s, produced in the narrow gap between the liquid surface and anode, will have a single average value of cS2 AFT. The exact cS2 AFT for a particular event depends on statistical fluctuations, and can be expected to follow a binomial distribution. We can therefore expect the fluctuations in the cS2 AFT to be proportional to  $1/\sqrt{S2}$ . Note that we expect the dependence to be on S2 rather than cS2, since it is the number of photons produced in the gas which matter for these statistical fluctuations, not the number of electrons produced at the interaction site.

Earlier versions of this cut were based on the S2 AFT (rather than cS2 AFT). However, the uncorrected values vary depending on whether the S2 is produced near PMTs which are turned off; by using the corrected top and bottom S2 signals, these variations are removed. The electron lifetime correction affects both top and bottom S2 equally, so is irrelevant when considering the ratio.

The S2 AFT cut is designed to accept 99% of real liquid interactions, independent of their S2 area. Events are accepted or not based on their location in the (S2, cS2 AFT) parameter space. The expected distribution is determined from  $^{220}\text{Rn}$  calibration data. After dividing the data into slices based on S2 area, the cS2 AFT in each is fit with a Gaussian distribution. The 0.5th and 99.5th percentiles of these distributions are shown in figure 3.6 as a function of S2 area, on top of the data used for the fit. These percentile points are in turn fitted to give an analytic expression for the cut boundary; the final selection criterion is

$$0.629 - 1.89 \cdot S2 < \text{cS2 AFT} < 0.634 + 1.57 \cdot S2. \quad (3.7)$$




Figure 3.7: Data for estimating the power of the S2 area fraction top cut at rejecting gas events. The purple points above the dashed line represent gas events removed by the cut; the red crosses are those five which are not removed. The green points below the dashed line are assumed to belong to the population of good events. The cut definition is shown as black lines.

Given its definition based on quantiles of the cS2 AFT distribution recorded in calibration data, this cut's acceptance should be 99%. We can try to estimate the efficiency at which the cut rejects gas events by selecting a sample rich in such events. For this we can use another cut, which also primarily targets gas events: the S2 width cut. This is based on events' *width*, defined as the shortest time window which contains at least 50% of the S2's area. As mentioned above, the S2 amplification takes place over a longer (and variable) distance for gas events, so their width is generally unusually large. By reversing the S2 width cut, we obtain a sample high in gas events; this is shown in figure 3.7. After discounting those events which are clearly not from the gas region (their cS2 AFT is far away from the range generally seen for gas events), 1303 remain. Five of these are not removed by the S2 area fraction top cut. These five may or may not be gas events, since they are consistent with either population. We can thus place a lower limit on the rejection efficiency at 99.6%.

### 3.2.2 Misidentified S1 $^{83\text{m}}\text{Kr}$ cut

Not all selection criteria are built to search for good events by selecting those with properties that behave as expected. Some are more directly targeted at known categories of bad events, to remove a specific contamination in the data. Here, we briefly look at one example, built to remove certain events arising from  $^{83\text{m}}\text{Kr}$  in background data.

As we have already seen,  $^{83\text{m}}\text{Kr}$  is periodically injected into the liquid xenon and used as a calibration source. It is produced from the decay of its parent isotope,  $^{83}\text{Rb}$ , which has a half-life of 86 days. During the operation of XENON1T a small amount of  $^{83}\text{Rb}$  leaked inside the main xenon volume, leading to a permanent  $^{83\text{m}}\text{Kr}$  contamination. In general these events show up as normal electronic recoils, but other behaviour is also possible due to the two-stage decay. One less common possibility is that the two S2 signals are merged (as

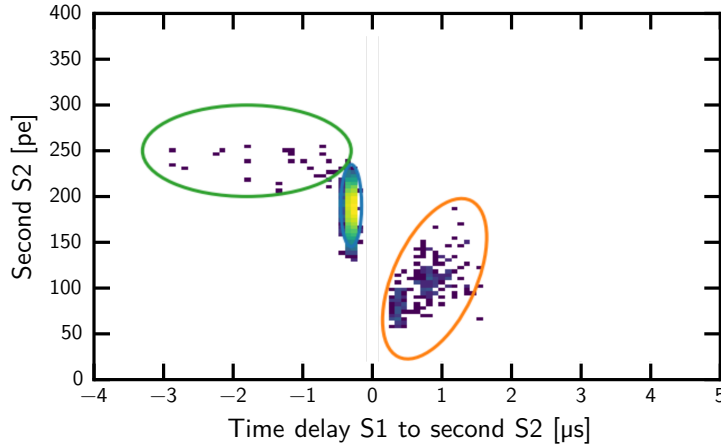


Figure 3.8: Events in  $^{83\text{m}}\text{Kr}$  calibration with the second-largest S2 near the main S1. The population circled in blue is where the 32.1 keV decay's S1 is identified as an S2; for the population in orange it is the 9.4 keV decay's S1 which is misidentified. The small population circled in green is probably events where two merged S1s are identified as an S2 and an afterpulse is identified as an S1.

happens most of the time) but the two S1s are not, and one of the S1s is misidentified as an S2. Since the false S2 will be rather small, these events are difficult to remove using the usual single-scatter selection criteria. On the other hand, because the S1 and S2 are effectively created from different energy deposits the events can fall somewhere outside the ER band, and therefore not be considered by any background model.

Fortunately it is easy to select such events, which must have a characteristically sized second S2, just before the S1. Figure 3.8 shows  $^{83\text{m}}\text{Kr}$  calibration data, after applying all other quality cuts, with the second S2's area plotted against the time of this second S2 relative to the S1. Two populations due to mis-classified S1s are visible, where either the 32.1 or 9.4 keV decay's S1 is classified as S2. It is the former which are most concerning, since the disparity between the S1 (from only the 9.4 keV part) and the S2 (from the full 41.5 keV) is largest. We can tag such events with essentially perfect efficiency by looking for a second S2 up to 3  $\mu\text{s}$  before the S1 and with an area of at least 100 photoelectrons.

Since this signature is extremely unlikely to randomly occur in good events, a cut removing all such tagged events will have a very good efficiency. We can estimate this by looking in calibration data, for events with a second S2 fulfilling the area criterion (at least 100 photoelectrons) before the main S1. In  $^{220}\text{Rn}$  calibration data, 0.038% of all events have such a second S2 in the 1 ms before the S1. Given that we expect these to be random in time, the acceptance of the cut looking only at a 3  $\mu\text{s}$  window should be very close to one,  $1 - 1.1 \times 10^{-6}$ .

### 3.2.3 Overall cut performance

To estimate the overall acceptance rate of dark-matter-like interactions, the cumulative effect of all the quality criteria applied must be considered. Most criteria have only negligible correlation and their acceptances can simply be multiplied, but four S1-related cuts are cor-




Figure 3.9: Overall cut efficiency as a function of S1 (left) and S2 (right). Purple: the acceptance of quality cuts which select good S1s and S2s individually (including the S2 area fraction top cut). Red: adds cuts which consider correlations between the S1 and S2 signal (for example ensuring that the S1 light pattern is consistent with the pattern reconstructed from the S2). Green: includes also cuts designed to select events with only a single interaction. Orange: adds the effect of some cuts designed to deal with effects specific to the high rate of interactions in  $^{220}\text{Rn}$  calibration data, which were also applied to background data for consistency. Figure adapted from [141].

related. The combined acceptance of these is determined by applying them all together on  $^{220}\text{Rn}$  calibration data, with all other criteria already applied. The fraction of events which are accepted is then taken as a conservative\* measure of the combined acceptance. The overall acceptance, as a function of S1 and S2 size, is shown in figure 3.9.

### 3.3 Signal and background modelling

We have seen an overview of the signal corrections and selection criteria which can be regarded as the low-level part of the analysis chain, with a few more detailed examples. We now move on to the interpretation of XENON1T data in terms of a limit on elastic WIMP scattering. An overview is given in the following, more details can be found in [57, 140].

The data used to search for WIMP elastic scattering were collected in two separate science runs. The first, SR0, started on the 22nd November 2016 and was halted due to disruption caused by the series of four magnitude 5.2–5.7 earthquakes which struck the region on the 18th January 2017. The second, SR1, ran from the 2nd February 2017 to the 8th February 2018. The only significant difference between the two is in the drift field in the TPC:  $120 \text{ V cm}^{-1}$

\*At least in the scenario where a limit is set (as opposed to a discovery made), this is conservative. An underestimation of the fraction of dark matter interactions accepted would mean that a given limit on number of interactions observed corresponds to an overestimation in the actual number of interactions, so a weaker upper limit on the interaction cross-section.

for SR0 and  $81 \text{ V cm}^{-1}$  for SR1. A third science run, SR2, which began immediately after SR1 and finished on the 3rd December 2018, is not used for this search.

In the combined SR0+SR1 dataset a total of 302.1 days of data were collected. A small fraction of the total is *dead time*, when any possible dark matter events would either not be recorded or not be considered as good events due to external influences. This can be introduced by the data acquisition system being unavailable, activity in the muon veto or periods of delayed electron extraction which follow large energy depositions. After subtracting the dead time, 32.1 and 246.7 days of data remain in SR0 and SR1, respectively, to search for WIMP events.

The dataset was *blinded* while the event selection criteria were developed and the expected signal and background shapes were determined. Events which were below the  $-2\sigma$  quantile of the expected electronic recoil band, based on preliminary models, were hidden from analysts. This is a common technique for such searches, and is intended to reduce bias by preventing selection criteria or models being fine-tuned based on individual events. The data were only unblinded once the full analysis chain was fixed.

For the purposes of searching for WIMPs, only a  $(1.30 \pm 0.01) \text{ t}$  sub-volume of the TPC was used, in order to avoid regions near its boundary where the background rate is higher. This *fiducial volume* is bounded at 8.0 cm below the liquid surface, above which a small residual population of gas events exists; 2.9 cm above the cathode, below which the electric field is more non-uniform leading to biases in position reconstruction and the charge-to-light ratio; and 5.2 cm inside the TPC wall, outside which there are a very large number of surface events (see below). In addition, the top and bottom of the fiducial volume vary as a function of radius, such that there is no more than 10% increase in the background rate at its upper and lower boundaries.

The shape of the electronic and nuclear recoil bands is modelled following the treatment used by NEST [136], as described in more detail in section 2.2. These models of the two bands are fit simultaneously to all  $^{220}\text{Rn}$ ,  $^{241}\text{AmBe}$  and neutron generator calibration data taken during both science runs (with the field-dependence being considered as part of the model). The resulting models allow a 99.7% rejection of electronic recoil events in a region between the  $-2\sigma$  quantile and median of a  $200 \text{ GeV}/c^2$  WIMP's expected signal in  $(cS_1, cS_2)_b$  space.

Seven sources of backgrounds are taken into account for the search. Backgrounds which induce electronic recoils include both intrinsic sources, that are mixed into the xenon itself, and those originating from natural radioactivity of the detector's materials. The rate of the latter is reduced by only searching within a smaller fiducial volume inside the TPC, thanks to the effective shielding of the xenon itself.

The important intrinsic sources of radioactivity are  $^{85}\text{Kr}$  and the daughters of  $^{222}\text{Rn}$ . Since both of these are noble elements they are not efficiently removed by the getter. The former is present in air (from which xenon is distilled), mostly due to emissions from nuclear re-processing [179]. Its concentration was reduced to subdominant levels by cryogenic distillation [163]. The latter is the result of radon emanation, as described in section 2.5. The

decay chain of  $^{222}\text{Rn}$  includes  $^{214}\text{Pb}$ , whose beta-decay dominates the intrinsic background spectrum at low energies. In the energy range 1.4 keV to 10.6 keV, which corresponds to the region used for this analysis, these both have an almost-flat spectrum due to their beta-decays. Their combined rate is  $82^{+5}_{-3}(\text{sys}) \pm 3(\text{stat})$  events/(t yr keV).

Nuclear recoils can be induced by solar neutrinos via coherent elastic neutrino-nucleus scattering (CE $\nu$ NS) [180] or by neutrons of both radioactive and cosmological origin. The most important of these is radiogenic neutrons, expected to contribute  $1.43 \pm 0.66$  events in the 1 t yr exposure. Their rate is constrained using both a GEANT4-based simulation [181, 182] and by searching for events in which one neutron scatters multiple times within the TPC. Cosmogenic neutrons are those induced by cosmic muons. They can be identified with a very high efficiency by the muon veto system, so they only contribute a negligible amount to the total NR background and are not considered.

Two final backgrounds, arising from detector effects, are also considered. The first is accidental coincidences of lone S1s or lone S2s, that happen to look like genuine particle-induced events. Samples from dedicated searches for lone S1s and lone S2s in the two science runs are paired randomly to create a model for such events. The second is so-called *surface events*. These are decays which occur on the surface of the PTFE walls of the TPC. Some of the electrons produced can be trapped on the wall, so the S2 is smaller than expected. These events are mostly from the decay of  $^{210}\text{Pb}$  and its daughters. These can be found on the TPC walls due to *radon plate-out*, a process whereby charged daughters of  $^{222}\text{Rn}$  become incorporated into the PTFE [183, 184].

Table 3.1 lists the sources of background mentioned above together with the expected number of events for each, calculated using the likelihood’s best fit. The number of events are listed both for the full analysis region, as seen in figure 3.10, and for a nuclear recoil ‘reference region’, defined as the region between the  $-2\sigma$  quantile and the median of the 200 GeV/ $c^2$  WIMP model in (cS1, cS2<sub>b</sub>) space. This region is intended to illustrate the signal-like content in the dataset, but is not directly used to constrain the cross-section.

## 3.4 Results

An unbinned extended likelihood is used to provide a statistical interpretation of the search data. Models used for this likelihood are produced in three dimensions: cS1, cS2<sub>b</sub> and radius. In addition a central region is identified which contains half of the fiducial volume, and events are categorised as being inside or outside this region. The central region is chosen because of its substantially smaller neutron rate ( $0.63 \text{ t}^{-1} \text{ yr}^{-1}$  compared to  $1.43 \text{ t}^{-1} \text{ yr}^{-1}$  in the full 1.3 t volume).

The same data and set of background models can be used to search for WIMPs scattering with a variety of physics channels. As mentioned in the introduction to this chapter, the canonical route is to search for spin-independent scattering. Signal models are computed for a variety of WIMP masses between 6 and 1000 GeV/ $c^2$ .

Table 3.1: Expected number of events in the full dataset from each background and from a 200 GeV/c<sup>2</sup> WIMP, based on the profiled likelihood.

Source	Events	
	Total	Reference region
Electronic recoils	$627 \pm 18$	$1.62 \pm 0.30$
Neutrons	$1.43 \pm 0.66$	$0.77 \pm 0.35$
CE $\nu$ NS	$0.05 \pm 0.01$	$0.03 \pm 0.01$
Accidental coincidence	$0.47^{+0.27}_{-0.00}$	$0.10^{+0.06}_{-0.00}$
Surface events	$106 \pm 8$	$4.84 \pm 0.40$
Total background	$735 \pm 20$	$7.36 \pm 0.61$
100 GeV/c <sup>2</sup> WIMP	3.56	1.70
Observed	739	14




Figure 3.10: Data used to search for elastic WIMP scattering. Here, the expected region for ER events is shown in red and that for surface events in blue. The expected signal from a 100 GeV/c<sup>2</sup> WIMP is shown in green. Figure adapted from [57]. The black dots show the locations of events within the 1.3 t fiducial volume after unblinding the data.



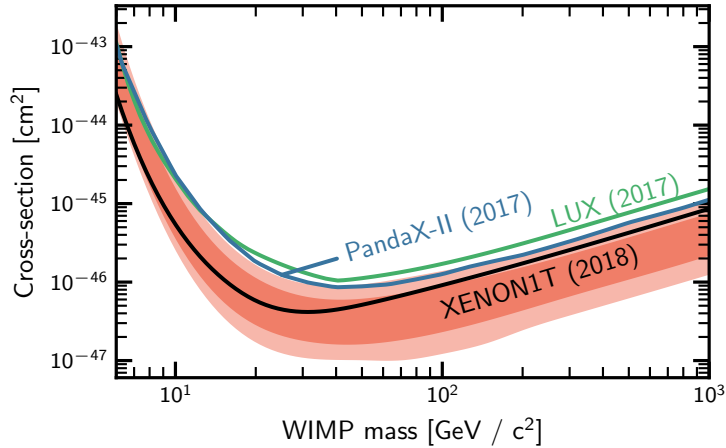


Figure 3.11: Upper limit on the cross-section of spin-independent elastic WIMP scattering off nucleons, obtained with XENON1T data. The shaded regions are the  $1\sigma$  and  $2\sigma$  range of expected values, respectively. Limits published by the LUX [108] and PandaX-II [106] collaborations are also shown for comparison. Figure adapted from [57].

To place constraints on the interaction cross-section for each of these masses, a profile likelihood-ratio test is used, as described in [185]. The number of events observed in the reference region is consistent with the background expectation to a level of  $2.1\sigma$ . Based on the likelihood analysis, this slight excess is not consistent with a significant WIMP signal. A limit is therefore placed on the WIMP-nucleon scattering cross-section at 90% confidence level, as shown in figure 3.11. The most stringent limit is set at a cross-section of  $4.1 \times 10^{-47} \text{ cm}^2$ , for  $30 \text{ GeV}/c^2$  WIMPs.

A similar technique can be used to set limits on spin-dependent WIMP interactions. The resulting limits, at 90% confidence level, can be seen in figure 3.12.

In this chapter we have seen an overview of the analysis techniques which allow dual-phase xenon time projection chambers to achieve excellent background rejection and reach unbeatable sensitivities to WIMP scattering. The limit set by XENON1T for the spin-independent elastic WIMP-nucleon scattering cross-section is, at the time of writing, the most stringent for WIMPs heavier than  $6 \text{ GeV}/c^2$ . As we will see in chapter 6, increasing the size of the experiment while introducing new technology to reduce some backgrounds and veto others, will allow future experiments such as XENONnT to reach even greater sensitivity. This may allow the discovery of dark matter interactions, and if not will certainly make it possible to push this upper limit further down.

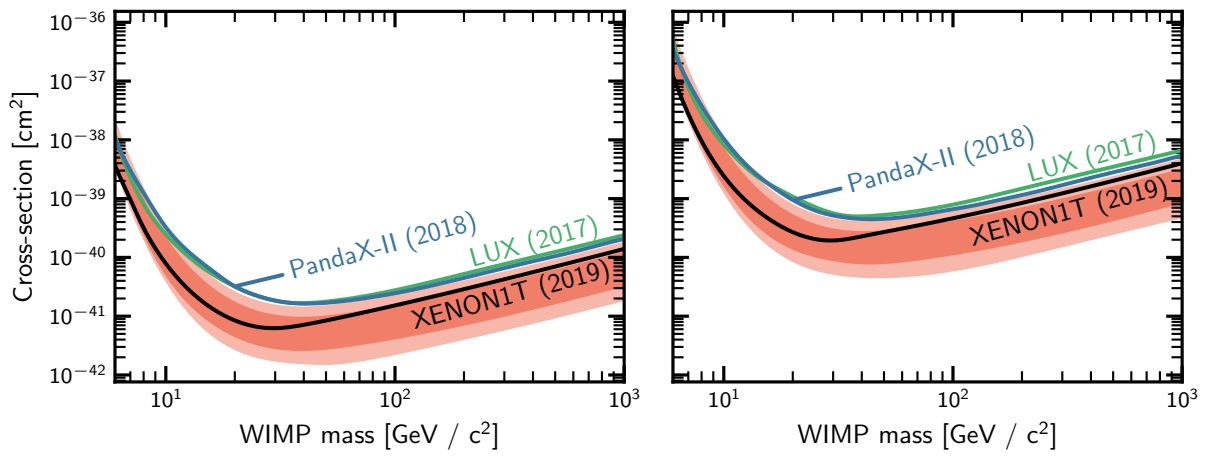


Figure 3.12: Upper limits on the cross-section of spin-dependent elastic WIMP scattering off neutrons (left) and protons (right), obtained with XENON1T data. The shaded regions are the  $1\sigma$  and  $2\sigma$  range of expected values, respectively. Limits published by the LUX [186] and PandaX-II [187] collaborations are also shown for comparison. Figure adapted from [66].

# Chapter 4

## Inelastic WIMP scattering

In the previous chapter we studied the search for the elastic scattering of WIMPs off xenon nuclei, resulting in an energy deposit in the form of a nuclear recoil. It is also possible for WIMPs to scatter inelastically, meaning that some of the energy of the interaction is used to excite the target nucleus [68, 188]. The nucleus would subsequently return to the ground state, emitting a photon in the process. The signature of such an interaction would therefore consist of both the nuclear recoil itself and an electronic recoil arising from absorption of the de-excitation photon.

Two isotopes of xenon could be interesting targets for inelastic WIMP scattering. These are  $^{129}\text{Xe}$ , which has a  $3/2+$  excited state 39.6 keV above the  $1/2+$  ground state, and  $^{131}\text{Xe}$ , with a  $1/2+$  excited state 80.2 keV above its ground state. Both isotopes form a significant fraction of natural xenon, with abundances of 26% and 21%, respectively. The structure functions and resulting recoil spectra for inelastic scattering off both of these two isotopes were calculated in [67], and are used for this analysis. A comparison of the recoil spectra for elastic and inelastic scattering, and of the inelastic spectra for a variety of WIMP masses, can be seen in figure 4.1. Here, we concentrate only on scattering off  $^{129}\text{Xe}$ , since the substantially higher energy of the first excited state suppresses the rate for  $^{131}\text{Xe}$  by around an order of magnitude.

Because some of the energy transferred is used to excite the target nucleus, imposing a higher threshold on WIMP velocity for inelastic scattering, the expected event rate is smaller at a given cross-section than elastic scattering. However, the subsequent de-excitation of the nucleus also means that all events will be above the detection threshold. This is in stark contrast to elastic scattering where the differential event rate increases towards arbitrarily small recoil energies.

Where inelastic scattering becomes especially interesting is in the case of a detection. The inelastic event rate is expected to be many orders of magnitude smaller than the elastic rate for spin-independent interactions, but much closer in the spin-dependent case. This means a detection of inelastic WIMP scattering would be strong evidence of a spin-dependent interaction [67, 189]. Although here we only consider standard spin-dependent scattering, Arcadi et al. recently showed that effective field theories contain certain operators for which inelastic WIMP scattering would provide even better sensitivity than elastic scattering [190].




Figure 4.1: Left: comparison of recoil spectra for the elastic (dotted) and inelastic (solid) scattering of  $100 \text{ GeV}/c^2$  WIMPs off  $^{129}\text{Xe}$  (purple) and  $^{131}\text{Xe}$  (red) nuclei. Right: recoil spectra for  $30 \text{ GeV}/c^2$  (red),  $100 \text{ GeV}/c^2$  (purple) and  $10 \text{ TeV}/c^2$  (green) WIMPs scattering inelastically off  $^{129}\text{Xe}$  nuclei. In each case the spectra are calculated using the results of [67], assuming a scattering cross-section of  $10^{-40} \text{ cm}^2$ .

We saw in chapter 1 that the velocity threshold is different for elastic and inelastic WIMP scattering. Because of this, we can probe different parts of the local velocity distribution using each type of interaction. A detection of both types of interaction would provide a way to constrain the shape of that distribution, using a single Earth-bound detector, by comparing the rates of each [67].

To search for inelastic WIMP scattering, data collected by XENON1T during SR1 is used. The same 1.3 t fiducial volume is used as for the search for elastic WIMP scattering, but unlike there the 24 hours of data after every  $^{83\text{m}}\text{Kr}$  calibration are excluded. This is because the background from  $^{83\text{m}}\text{Kr}$  is very close to the signal expected from inelastic scattering, and removing those 24 hour periods results in roughly a factor three reduction of the average  $^{83\text{m}}\text{Kr}$  background rate. As a result, the exposure is lower than for the elastic scattering search, at 0.89 t yr.

A slightly different set of selection criteria was developed for use at energies of several tens of kilo-electron-volts, compared to the set used for the elastic scattering search. The same selection criteria were used for the original search for  $^{124}\text{Xe}$   $2\nu\text{ECEC}$  [129], with the only difference being that an additional criterion was used in that case to ensure that events were inside the expected region for electronic recoils. A summary of these can be found in table 4.1. Within the region of interest for this search, the combined acceptance of all cuts for good events is  $(93.6 \pm 0.7) \%$ .

Like the elastic scattering search, this analysis was performed using blinded data. Events with a  $\text{cS1}$  between 250 pe and 550 pe, and a  $\text{cS2}_b$  between 5000 pe and 13000 pe were hidden by the blinding procedure; meaning the central  $2\sigma$  region of the expected signal model

Table 4.1: Overview of selection criteria used for the search for inelastic WIMP scattering.

Criterion	Description
cS2 <sub>b</sub> threshold	cS2 <sub>b</sub> must be at least 1000 pe.
S2 Width	Depth-dependent constraint on duration of S2 signal.
Krypton mis-ID S1	Misidentified S1 <sup>83m</sup> Kr cut (section 3.2.2).
S2 area fraction top	Constrains S2 fraction seen by top PMTs (section 3.2.1).
Flash	Rejects events occurring soon after a PMT flash.
S2 single scatter	Limit on the area of the second-largest S2-like peak.
S1 single scatter	Only one S1-like peak with a possible matching S2.
S2 tails	Rejects events in single-electron trains after large S2s [191].
Pre-S2 junk	Rejects noisy events with spurious signals before S2.
S1 area fraction top	Depth-dependent constraint on S1 fraction seen by top PMTs.
S1 pattern likelihood	Goodness-of-fit constraint for pattern of S1 light.
DAQ veto	Rejects events which are only partially recorded.
Fiducial volume	Events must be within 1.3 t fiducial volume.

was hidden for all WIMP masses. The expected signal response and all backgrounds were characterised without looking at the blinded events, in order to minimise unintentional bias.

## 4.1 Expected signal

In order to model the response of XENON1T to an inelastic WIMP-nucleus interaction, we must consider both the nuclear recoil itself and the effect of its de-excitation.

For this analysis we rely on two assumptions about the interplay between these two energy deposits. Firstly, we consider them independently: we compute the expected light signal from each and assume that the total light signal is simply the sum, and the same for the charge signal. This is important, since the light and charge responses of liquid xenon are not linear (their yields vary with energy). The mean absorption length of a 39.6 keV photon is 150  $\mu\text{m}$  [123]. On the other hand, ionisation electrons thermalise over a characteristic distance of 4.6  $\mu\text{m}$  in liquid xenon [192]. It is therefore safe to assume that the two electron clouds are well separated and don't affect each other's recombination probabilities.

Secondly, we assume that although the ER and NR signals are produced independently, they occur so close in both space and time that they are experimentally indistinguishable from a single energy deposition. The 150  $\mu\text{m}$  characteristic separation mentioned above is significantly less than the  $\mathcal{O}(1\text{ mm})$  position resolution of XENON1T [141]. The 0.97 ns half-life of the excited state is smaller than the PMT sampling period (10 ns) and the decay time involved in the scintillation process (3 ns or 27 ns for the singlet and triplet excited state, respectively). This means the two signals can safely be considered to be simultaneous.




Figure 4.2: Best-fit model for inelastic scattering of neutrons off  $^{129}\text{Xe}$  (red lines, showing the  $1\sigma$ ,  $2\sigma$  and  $3\sigma$  contours). This is compared to data from calibration with the neutron generator (2D histogram indicated by shading within the  $2\sigma$  contour and individual events shown in green outside).

As mentioned, the nuclear recoil spectra are used as calculated in [67]. As in the search for elastic scattering, the expected detector response to the nuclear recoil is calculated using a NEST-like model [136], with the parameters constrained by a Bayesian fit to XENON1T calibration data [140]. The only significant difference between this analysis and searches for elastic scattering is that the threshold criteria, such as the three-fold PMT coincidence requirement for S1s, are not applied. Since the nuclear recoil will later be combined with a significant electronic recoil, it will always pass the detector threshold.

In general, the  $(cS1, cS2_b)$  distribution for events from a mono-energetic electronic recoil peak is described by a two-dimensional Gaussian function [174]. The five parameters of that function (the mean  $cS1$ ,  $\mu_1$ , and the mean  $cS2_b$ ,  $\mu_2$ ; the width in each direction,  $\sigma_1$  and  $\sigma_2$ ; and the correlation between  $cS1$  and  $cS2_b$ ,  $\rho$ ) are determined using the neutron generator calibration data.

The complete model for neutrons inelastically scattering off  $^{129}\text{Xe}$  is fitted to this calibration data by maximising the binned likelihood. The resulting best-fit model is shown in figure 4.2, together with the calibration data. An additional component is included in the fit, representing the ER background. This is assumed to have the same distribution as  $^{220}\text{Rn}$  calibration events, and the normalisation is left free to be fitted. The expected distribution of likelihoods was computed by generating Monte-Carlo datasets from the best fit model. By comparing the likelihood for the real data to this distribution, it is possible to compute a p-value for the quality of the fit, which is 0.22.

Finally, to determine the response model for WIMPs' inelastic scattering, the electronic recoil from the neutron fit is convolved with the expected nuclear recoil for the inelastic scattering of a WIMP. Figure 4.3 shows the result for a variety of WIMP masses.

### 4.1.1 Uncertainties in the signal model

The final signal model depends on the value of eighteen parameters, listed in table 4.2. Five of these are related to the electronic recoil part of the signal. The other thirteen are physical parameters used to describe the nuclear recoil part.

An uncertainty in each of these is evaluated, as discussed in the following. Signal models are also computed with each parameter being varied individually. The propagation of these uncertainties into the search for inelastic WIMP scattering is discussed later, in section 4.3.2. For the nuclear recoil parameters, the width of the posterior probability distribution for each is taken as the uncertainty. For the electronic recoil parameters, a confidence interval is computed by using the likelihood described above.

## 4.2 Background modelling

Four backgrounds are considered in the region of interest, from the decays of  $^{214}\text{Pb}$ ,  $^{83\text{m}}\text{Kr}$ ,  $^{124}\text{Xe}$  and  $^{125}\text{I}$ . The way that these are modelled is described in the following.

As for the elastic scattering search, a significant background for inelastic WIMP scattering is the beta-decay of  $^{214}\text{Pb}$ . Both the energy spectrum and the  $\text{cS2}_b/\text{cS1}$  distribution of these must be modelled. The decay spectrum is determined from detailed GEANT4 simulations of the detector. Data outside the blinded region can be used to verify the accuracy of this model, as seen in figure 4.4. Three models are needed to describe this data satisfactorily:  $^{214}\text{Pb}$ , discussed here,  $^{133}\text{Xe}$ , which is an unstable isotope of xenon produced by neutron capture during neutron calibrations, and  $^{83\text{m}}\text{Kr}$ . With these three models, the p-value determined from the  $\chi^2$  of the best-fit is 0.57.

A small hint of an additional component is visible just at the edge of the blinding region, around 65 keV, but the significance is not large so it is not included here. As we will see below, this is probably caused by events from the decays of  $^{124}\text{Xe}$  and  $^{125}\text{I}$ . For energies below 71 keV,  $^{133}\text{Xe}$  makes up less than 1% of the total rate. It is therefore not considered as a background for the search for inelastic WIMP scattering. The upper energy limit of the region of interest used for the search is defined by this energy, and  $^{133}\text{Xe}$  is not included in the likelihood fit.

The  $\text{cS2}_b/\text{cS1}$  distribution, at a given energy, is determined using data from  $^{220}\text{Rn}$  calibrations. Within the region of interest for this analysis, we see only pure beta-decays of  $^{212}\text{Pb}$  with a similar spectrum to that of the background  $^{214}\text{Pb}$ . To ensure that the  $\text{cS2}_b/\text{cS1}$  response is the same as in background data, we again look at energies either side of the blinded region. At higher energies than the blinded region, the additional component,  $^{133}\text{Xe}$ , affects this ratio. Its beta-decay occurs only to excited states of  $^{133}\text{Cs}$ , and so electrons from the decay are always detected together with a photon of at least 81 keV. By excluding data within two months of a neutron calibration (with either an  $^{241}\text{AmBe}$  source or the neutron generator), it is possible to remove events from the decay of  $^{133}\text{Xe}$ , which has a half-life of 5.2 days.




Figure 4.3: Expected signals from the inelastic scattering of 30 GeV/c<sup>2</sup> (red), 100 GeV/c<sup>2</sup> (purple) and 10 TeV/c<sup>2</sup> (green) WIMPs off <sup>129</sup>Xe nuclei. The contours correspond to the 1 $\sigma$  and 2 $\sigma$  regions.

Table 4.2: The eighteen parameters affecting the expected signal model for inelastic WIMP scattering off <sup>129</sup>Xe.

Symbol	Constraint	Description
$\mu_1$	$317.1 \pm 1.3$	Mean cS1, in pe
$\mu_2$	$7318 \pm 42$	Mean cS2 <sub>b</sub> , in pe
$\sigma_1$	$30.1 \pm 1.1$	Width in cS1, in pe
$\sigma_2$	$1060 \pm 34$	Width in cS2 <sub>b</sub> , in pe
$\rho$	$0.476 \pm 0.044$	Correlation between cS1 and cS2 <sub>b</sub>
$\alpha$	$1.240 \pm 0.079$	NR response parameters [136, 140]
$\zeta$	$0.047 \pm 0.009$	
$\beta$	$239 \pm 28$	
$\gamma$	$0.0139 \pm 0.0007$	
$\delta$	$0.062 \pm 0.006$	
$\kappa$	$0.139 \pm 0.003$	
$\eta$	$3.3 \pm 0.7$	
$\lambda$	$1.14 \pm 0.45$	
$W$	$13.7 \pm 0.2$	Work function, in eV
$g_1$	$0.142 \pm 0.002$	Mean cS1 gain, in pe/photon
$g_1$	$11.4 \pm 0.2$	Mean cS2 <sub>b</sub> gain, in pe/electron
$p_{\text{dpe}}$	0.18–0.24	Probability of detecting two pe per incident photon
$\delta\tau_e$	$0.00 \pm 0.02$	Electron lifetime – fractional deviation from best-fit



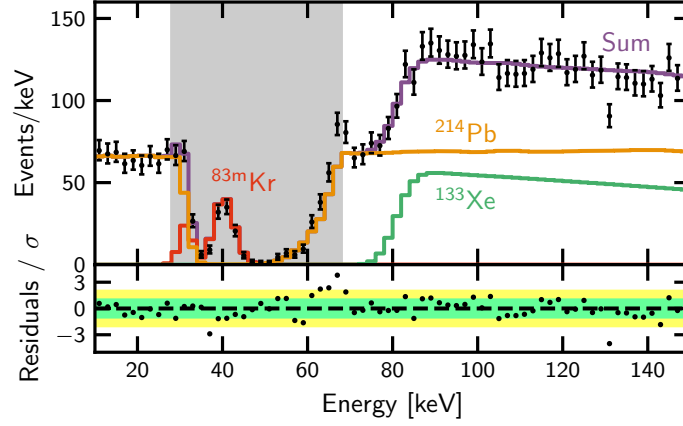


Figure 4.4: Energy spectrum of data outside the blinded region, compared to models. The  $^{214}\text{Pb}$  (orange) and  $^{133}\text{Xe}$  (green) models are from Monte-Carlo simulations using Geant4. The  $^{83\text{m}}\text{Kr}$  (red) model is from calibration data. The sum is shown in purple. The grey shaded region indicates bins where the blinding cut has an effect, and these bins are not included in the fit.

In contrast, at energies below 71 keV these events are not present and we can use the full exposure. The comparison is performed in two bins between 15 keV and 25 keV, and three between 70 keV and 85 keV. The distribution within each of these bins can be found in figure 4.5. In each bin, a Kolmogorov-Smirnov test is used to compute a p-value for the quality of match between the  $\text{cS2}_b/\text{cS1}$  distributions for  $^{220}\text{Rn}$  calibration and background events. We can then use Fisher's method to combine these into a single p-value, which is 0.13.

Having established that the simulated data and  $^{220}\text{Rn}$  calibration data describe the energy spectrum and  $\text{cS2}_b/\text{cS1}$  distribution, respectively, well, we can combine these to compute a model for  $^{214}\text{Pb}$  events. The number of events expected in each bin are given by

$$n_{ij} = \tau \frac{f_i^{\text{MC},220\text{Rn}}}{f_j^{\text{MC},220\text{Rn}}} f_{ij}^{\text{cal},220\text{Rn}}, \quad (4.1)$$

where  $f_{ij}^{\text{MC/cal},X}$  is the fraction of the simulated/calibration  $X$  events in energy bin  $i$  and  $\text{cS2}_b/\text{cS1}$  bin  $j$ .

In the previous chapter, we mentioned that there is a contamination of the calibration source  $^{83\text{m}}\text{Kr}$  present throughout the XENON1T science data. While before we were concerned about removing mis-reconstructed events, we are now concerned about the main peak of properly reconstructed decays. The total energy deposited in the two stages of  $^{83\text{m}}\text{Kr}$  decay is 41.5 keV, only marginally greater than the 39.6 keV energy of the first excited state of  $^{129}\text{Xe}$ . It is therefore important to properly account for this background and ensure that it is not incorrectly interpreted as a dark matter signal. Fortunately, the abundant  $^{83\text{m}}\text{Kr}$  calibration data itself can be used to construct a high-statistics model.

The two remaining backgrounds both contribute mono-energetic peaks to the region of interest, and are considered in the same way. The first of these is from the decay of  $^{124}\text{Xe}$




Figure 4.5: Comparison between the  $cS2_b/cS1$  distribution in  $^{220}\text{Rn}$  calibration data (blue shaded region) and in background data (black data points and error bars), for a variety of energy ranges.

to  $^{124}\text{Te}$  via double electron capture with the emission of two neutrinos ( $2\nu\text{ECEC}$ ). The difference in nuclear binding energies is carried away by the two neutrinos and not detected. Instead, the detectable signal is a cascade of Auger electrons and X-rays produced when the two vacancies are filled. This decay was first measured in [129]. The second is the decay of  $^{125}\text{I}$  to  $^{125}\text{Te}$  via electron capture (EC). Again, the detectable signal is produced when the vacancy in an inner electron shell is filled.

In 80.1% of  $^{125}\text{I}$  decays, it captures the electron from the K shell (referred to here as EC-K), while in 15.6% of cases it is from the L shell (EC-L) [193]. The remaining 4.3% is made up of capture from higher energy shells, but is not important here since the signal falls outside the region of interest for this study. Similarly, when  $^{124}\text{Xe}$  decays, each electron can be captured from either the K shell or the L shell. In 76% of cases both electrons are captured from the K shell ( $2\nu\text{ECEC-KK}$ ); 23% of the time one is captured from the K shell and the other from the L shell ( $2\nu\text{ECEC-KL}$ ) [194]. Capture of both from the L shell, or from higher shells, happens in only 1% of cases, and is not considered here. These modes of decay are summarised in

Table 4.3: Summary of electron capture and double electron capture signals which create a background for the search for inelastic WIMP scattering. Capture from higher energy shells, including  $2\nu\text{ECEC-LL}$ , is not included here since it produces signals outside the region of interest.

Decay	$2\nu\text{ECEC-KK} / \text{EC-K}$		$2\nu\text{ECEC-KL} / \text{EC-L}$	
	Energy [keV]	Branching ratio	Energy [keV]	Branching ratio
$^{124}\text{Xe } 2\nu\text{ECEC}$	64.3	76%	36.7	23%
$^{125}\text{I } \text{EC}$	67.3	80%	40.4	16%

table 4.3, together with the energies released when filling the vacancies.

Since these signals contain a mixture of Auger electrons and X-rays, modelling the observed light and charge signals for them is complex. No attempt is made to predict the  $cS2_b/cS1$  distribution of events from the two higher energy signals, from  $2\nu\text{ECEC-KK}$  and  $\text{EC-K}$ . For this reason, at energies above 55 keV the analysis bins are divided only by energy and cover the full range of  $cS2_b/cS1$ . Since these peaks only overlap with the high-energy tail of the expected signal, and only for higher WIMP masses, this only has a limited effect on the sensitivity of the search.

On the other hand, the two lower-energy peaks, from  $2\nu\text{ECEC-KL}$  and  $\text{EC-L}$ , lie very close in energy to the expected dark matter signal. It is therefore important to model the distribution of  $cS2_b/cS1$  for their events. Fortunately, we can make use of the nearby  $^{83m}\text{Kr}$  signal to help with this. To do so, the contributions of Auger electrons and X-rays to the signal observed for each decay is first computed using RELAX [195]. NEST [128] is then used to simulate the light and charge signals for each. This is done by performing simulations assuming that the full energy deposit is in the form of first electrons and then X-rays, and taking the weighted mean of these two signals according to the results from RELAX. NEST is also used to simulate the signal from  $^{83m}\text{Kr}$ . Finally, the observed distribution of  $(cS1, cS2_b)$  for  $^{83m}\text{Kr}$  events is morphed by using the ratios of the expected light and charge signal sizes from  $2\nu\text{ECEC/EC}$  and from  $^{83m}\text{Kr}$ .

In order to estimate the uncertainty arising from this template morphing procedure, the same technique is used to predict the mean  $cS2_b/cS1$  of a pure 41.5 keV beta interaction. This is compared to  $^{222}\text{Rn}$  calibration data, selecting events between 41 keV and 42 keV. The discrepancy is converted into an uncertainty on the charge yield and light yield predicted by NEST. Since the number of quanta is fixed for a given energy, the uncertainties on each of these should be perfectly anti-correlated, and are estimated at 0.66 quanta/keV.

### 4.3 Statistical interpretation

The search for inelastic WIMP scattering is performed using a frequentist approach based on the binned likelihood

$$\ln \mathcal{L}_{\text{science}} = \sum_{i=1}^n d_i \ln f_i - f_i, \quad (4.2)$$

where  $d_i$  is the number of events observed in bin  $i$ , and  $f_i$  is the expected number of events. This is calculated by summing the expectations from the four background sources and WIMP inelastic scattering:

$$f_i = \epsilon_{\text{eff}} \mu_{\text{WIMP}} f_{\text{WIMP},i} + \mu_{\text{Rn}} w_i A_i + \sum_{S \in \{\text{Kr}, \text{Xe}, \text{I}\}} \mu_S f_{S,i} \quad (4.3)$$

where  $\mu_S$  is the number of events from the source  $S$  and  $f_{S,i}$  is the fraction of those events in bin  $i$ . For the  $^{222}\text{Rn}$  background,  $A_i$  is the fraction of events in  $^{220}\text{Rn}$  calibration data which are in bin  $i$  and  $w_i$  is the per-bin weighting needed to transform these into the number of expected  $^{222}\text{Rn}$  events in that bin (essentially the first fraction from equation 4.1). We will encounter  $\epsilon_{\text{eff}}$ , the effective efficiency parameter, shortly.

To produce limits on the scattering cross-section, we use the log likelihood ratio

$$t(\mu_{\text{WIMP}}) = -2 \ln \frac{\mathcal{L}(\mu_{\text{WIMP}}, \hat{\theta})}{\mathcal{L}(\hat{\mu}_{\text{WIMP}}, \hat{\theta})}, \quad (4.4)$$

where  $\mu_{\text{WIMP}}$  and  $\hat{\theta}$  are the values of the WIMP scattering rate parameter  $\mu_{\text{WIMP}}$  and all other parameters  $\theta$  (the *nuisance parameters*) which together maximise the likelihood, and  $\hat{\mu}_{\text{WIMP}}$  and  $\hat{\theta}$  are the values of the nuisance parameters which maximise it for a given value of  $\mu_{\text{WIMP}}$ .

Using the likelihood ratio as a test statistic to set limits means that we can switch from reporting only upper limits, when there is insufficient evidence of WIMP scattering, to two-sided intervals, once sufficient evidence exists. The correct coverage will also be preserved [185]. This means that, whatever the true value of  $\mu_{\text{WIMP}}$ , 90% of experiments like XENON1T would report a limit or an interval containing the true value. This property is typically lost when using other methods of constructing limits, when switching from upper limits to two-sided intervals. A decision to only report upper limits if the evidence for a signal is less than  $3\sigma$ , however, leads to *over-coverage*, especially for small WIMP scattering cross-sections: more than 90% of intervals would contain the true value.

In the remainder of this section we will first explore the motivation for the binning structure used to evaluate the likelihood, which is seen in figure 4.6. Then we discuss which systematic uncertainties are important for the result and how these are incorporated.




Figure 4.6: Region of interest showing the four background models considered –  $^{222}\text{Rn}$  (orange),  $^{83\text{m}}\text{Kr}$  (green),  $^{124}\text{Xe}$  (pink) and  $^{125}\text{I}$  (blue) – and the expected signal from a  $100 \text{ GeV}/c^2$  WIMP (red). The boundaries of the bins used to evaluate the likelihood are indicated by thin black lines. The  $^{124}\text{Xe}$   $2\nu\text{ECEC-KK}$  and  $^{125}\text{I}$  EC-K peaks are visualised assuming the same  $cS2_b/cS1$  distribution as for  $^{222}\text{Rn}$ , although their true shape is unknown (and unimportant) in this dimension.

### 4.3.1 Binning structure optimisation

For the best sensitivity, the bins used to evaluate the likelihood should generally be as small as possible. However, several constraints prevent us from using arbitrarily small sizes. The most obvious is that as the number of bins increases the complexity of the likelihood increases, and therefore so does the computational time needed to set limits.

A second requirement is that bins should not be too empty. The main reason for this is that when bins are sufficiently full, the likelihood ratio gains certain useful properties. In particular, we make use of Wilks’s theorem [196]. This states that in the limit of large statistics, the log likelihood ratio  $t(\mu)$  is  $\chi^2$ -distributed with a number of degrees of freedom equal to the number of dimensions in  $\mu$  – in this case one – when  $\mu$  has its true value. As the bins become smaller it also becomes more difficult to incorporate systematic uncertainties into the likelihood. We will see that for the  $^{222}\text{Rn}$  background this is done on a per-bin basis. If the bins were much smaller, the per-bin uncertainty would become relatively larger and eventually could ‘wash out’ all other information. For these reasons, all bins are required to have at least five expected events.

Let us first consider the ‘main part’ of the region of interest, where there are regularly spaced bins 2 keV in width and with a  $cS2_b/cS1$  range of 5. This size is optimised by comparing the sensitivity to WIMPs as the bin size is gradually increased, as seen in figure 4.7. For this study the entire region of interest was divided into rectangular bins. The optimised bin size results




Figure 4.7: Relative sensitivity loss as the size of bins is gradually increased. In each case the sensitivity is compared to that obtained with 56 1-keV-wide energy bins between 15 keV and 71 keV, and 50 bins between 0 and 50 in  $cS_{2b}/cS_1$ . In purple:  $28 \times 25$  bins, in red:  $28 \times 10$  bins, in green:  $11 \times 10$  bins, and in orange:  $11 \times 5$  bins. The black dotted line shows a projection onto energy, for comparison (i.e.  $56 \times 1$  bins).

in roughly an 8% loss of sensitivity compared to the baseline (1 keV wide and a  $cS_{2b}/cS_1$  range of 1), which is considered acceptable.

The remaining structure in figure 4.6 can be explained as follows. At high energies (above 71 keV), bins are divided only by energy, due to the lack of knowledge about the charge and light yields for the  $^{124}\text{Xe}$   $2\nu\text{ECEC-KK}$  and  $^{125}\text{I}$  EC-K peaks, as described in the previous section. At low energies (below 15 keV) only a single large bin is used. This is because there is only one background present in that region ( $^{222}\text{Rn}$ ). Instead of being included in the main science likelihood, this bin is used to add a constraint to the  $^{222}\text{Rn}$  rate. This simplifies sensitivity estimation and provides a degree of protection against bugs (by ensuring the rate parameter remains within the central region of its constraint) without otherwise affecting the result. Finally, at particularly low and high values of the ratio  $cS_{2b}/cS_1$  (below 15 and above 35), larger bins are used. This is to ensure the five-event criterion mentioned above is met.

### 4.3.2 Systematic uncertainties

In this section we discuss the implementation of systematic uncertainties on each of the models in the likelihood. Most of the uncertainties which are considered as nuisance parameters in the fit, which are summarised in table 4.4.

Table 4.4: Overview of parameters in the likelihood. Where a Gaussian constraint is included in the likelihood for that parameter, the mean and standard deviation of the constraint term are given. The constraint on the effective efficiency  $\epsilon_{\text{eff}}$  varies between  $0.936 \pm 0.033$  for 30 GeV/ $c^2$  WIMPs and  $0.936 \pm 0.014$  for 10 TeV/ $c^2$  WIMPs.

Symbol	Constraint	Description
$\mu_{\text{Rn}}$	$1.00 \pm 0.03$	$^{222}\text{Rn}$ rate, relative to expectation from low-energy bin
$\mu_{\text{Kr}}$	–	$^{83\text{m}}\text{Kr}$ rate
$\mu_{\text{Xe}}$	–	$^{124}\text{Xe}$ $2\nu\text{ECEC}$ rate
$\mu_{\text{I}}$	–	$^{125}\text{I}$ EC rate
$\mu_{\text{WIMP}}$	–	WIMP scattering cross-section, in units of $10^{-40} \text{ cm}^2$
$E_{\text{Xe,KK}}$	$64.3 \pm 0.6$	Mean energy of the $^{124}\text{Xe}$ $2\nu\text{ECEC-KK}$ peak, in keV
$\sigma_{\text{Xe,KK}}$	$2.6 \pm 0.3$	Energy resolution for the $^{124}\text{Xe}$ $2\nu\text{ECEC-KK}$ peak, in keV
$E_{\text{I,K}}$	$67.3 \pm 0.5$	Mean energy of the $^{125}\text{I}$ EC-K peak, in keV
$\sigma_{\text{I,K}}$	$2.8 \pm 0.5$	Energy resolution for the $^{125}\text{I}$ EC-K peak, in keV
$\Delta Q_y^{\text{EC}}$	$0.00 \pm 0.66$	Charge yield for KL and K relative to prediction, in quanta/keV
$\epsilon_{\text{eff}}$	Yes	Effective efficiency
$\mu_2$	$7318 \pm 42$	Mean $\text{cS2}_b$ of inelastic ER, in photoelectrons

### Data-driven backgrounds

Two of the backgrounds – the ER continuum due to  $^{222}\text{Rn}$  daughters and the  $^{83\text{m}}\text{Kr}$  peak – are modelled using calibration data, in the former case combined with simulated data. Since these calibrations have finite statistics, there is an uncertainty on the number of events expected in each bin. The proper way to treat these uncertainties is to say that there is a true, unknown model, which predicts the number of events in each bin of both the calibration and the science data. This model should be determined by maximising the combined likelihood  $\ln \mathcal{L} = \ln \mathcal{L}_{\text{science}} + \ln \mathcal{L}_{\text{calibration}}$ .

This procedure, commonly referred to as a *combined fit*, is also used for XENON1T searches for elastic WIMP scattering [140]. However, in that case an unbinned, not a binned, likelihood is used and the uncertainties are not on the number of events in each bin but a set of physical parameters which describe the model. The problem here is that the combined fit would include an additional 64 parameters: one per bin. Maximisation of the likelihood with so many parameters would be very challenging computationally, if not unfeasible. A solution is provided by the method described in [197], which replaces this 64-dimensional fit with 64 independent equations.

Turning now to the calibration datasets themselves, further simplification is possible. Firstly, the uncertainty in each bin arising from random fluctuations in the  $^{83\text{m}}\text{Kr}$  calibration is never more than 0.4% of that from  $^{220}\text{Rn}$  calibration, so we can neglect any uncertainty in the  $^{83\text{m}}\text{Kr}$  model. Secondly, at any given energy there are between 700 and 800 times as many simulated  $^{222}\text{Rn}$  events as there are events in the  $^{220}\text{Rn}$  calibration data. It is therefore safe to ignore

uncertainties arising from the simulated data, and consider only the  $^{220}\text{Rn}$  calibration's statistics. We are now at a point where there are only two datasets to include in the combined fit: the science data itself and the  $^{220}\text{Rn}$  calibration. This allows a further simplification of the method from [197], since the 64 equations become analytic when there is only one calibration source. They then manifest themselves as an analytic 'adjustment' to the per-bin number of expected events in the  $^{222}\text{Rn}$  model, based on the distribution of events in the science data. This adjustment is performed before maximising the likelihood.

### Electron capture and double electron capture

The two backgrounds from  $^{124}\text{Xe}$   $2\nu\text{ECEC-KK}$  and  $^{125}\text{I}$   $\text{EC-K}$  have uncertainties coming from how well the energy scale and resolution of the detector are known. Writing these again explicitly now, they enter the likelihood as

$$\mathcal{L}_X = \text{Gauss}(E_X | \tilde{E}_X, \delta E_X) \cdot \text{Gauss}(\sigma_X | \tilde{\sigma}_X, \delta\sigma_X), \quad (4.5)$$

where  $X$  is either  $2\nu\text{ECEC-KK}$  or  $\text{EC-K}$ ,  $E_X$  is the mean energy of the corresponding peak in the energy spectrum and  $\sigma_X$  is its width, or energy resolution. Throughout this section, the notation  $\tilde{x}$  refers to the best fit value of  $x$ , from an independent measurement and  $\delta x$  to its uncertainty. The two decays' expected mean energy and peak resolutions are as reported in [129].

The dominant uncertainty in the models from  $^{124}\text{Xe}$   $2\nu\text{ECEC-KL}$  and  $^{125}\text{I}$   $\text{EC-L}$  is on the predicted charge and light yields for both their decays and that of  $^{83\text{m}}\text{Kr}$ . As we discussed in section 4.2, these impact the morphing procedure used to generate their expected distributions. Given that signals from the two decays are generated via a very similar process, the uncertainties are assumed to be correlated between the two models. Therefore a single extra parameter,  $\Delta Q_y^{\text{EC}}$  is introduced to the likelihood, along with a constraint term:

$$\mathcal{L}_{Q_y} = \text{Gauss}(\Delta Q_y | 0, 0.66). \quad (4.6)$$

This parameter describes an adjustment to the charge yield of the  $2\nu\text{ECEC-KL}$  and  $\text{EC-L}$  peaks, compared to the nominal value, with a simultaneous adjustment of  $-\Delta Q_y^{\text{EC}}$  to the light yield of each in order to preserve the total number of quanta. The origin of the value 0.66 quanta/keV is described in section 4.2.

### Signal model

As we saw in section 4.1, the shape of the expected dark matter signal depends on eighteen parameters. Including all of these in the likelihood fit would not be viable, both in terms of computation time and the possibility of converging to a false (local) minimum. Instead, we take each independently in turn, and investigate its effect on the parameter of interest: the inelastic cross-section.






Figure 4.8: Equivalent rate uncertainty for each parameter affecting the shape of the signal model, for 30 GeV/c<sup>2</sup> (red) and 100 GeV/c<sup>2</sup> (purple) WIMPs. The first five parameters describe the electronic recoil model;  $\alpha$ – $\lambda$  describe the liquid xenon response to nuclear recoils and are defined in [140] and [136]; the remainder are detection-related: the work function,  $W$ , and the S1 and S2 gains,  $g_1$  and  $g_2$ , were described in chapter 2,  $p_{\text{dpe}}$  is the probability that a PMT emits two photoelectrons for a single incident photon, and  $\delta\tau_e$  is the 2% uncertainty on the electron lifetime [140].

A collection of 100 simulated datasets, with no signal present, is produced by drawing events from the expected background distributions. Each signal model parameter is then taken in turn and three alternative signal models are computed, with the parameter at the best fit and varied up and down by its uncertainty. A limit can then be set on the cross-section using each of these models, for all of the simulated datasets. We can therefore find an equivalent uncertainty on the signal rate  $r_j$  for each parameter of its model, computed as the median effect of a one-sigma variation:

$$r_j = \left| \text{Median} \left( \frac{l_{+1}^j - l_{-1}^j}{2l_0^j} \right) \right|, \quad (4.7)$$

where  $l_{0/\pm 1}^j$  is the limit set for a signal model in which parameter  $j$  is set at its nominal value or  $\pm 1\sigma$ . Figure 4.8 shows the equivalent effect of the eighteen signal-model parameters for two different WIMP masses.

For low-mass WIMPs, the mean,  $\mu_2$ , of the cS2<sub>b</sub> of the inelastic signal’s electronic recoil is dominant, with a  $1\sigma$  change resulting in a median of around a 5% difference in the limit. This makes some sense: for low masses the nuclear recoil only makes a minor contribution to the total signal (see figure 4.3), which is very similar to a pure 39.6 keV gamma-ray. The mean S2 of events from this gamma-ray strongly affects how well it can be distinguished from the nearby <sup>83m</sup>Kr peak and the <sup>222</sup>Rn background. Due to its importance,  $\mu_2$  is included directly in the likelihood fit with its constraint from the fit to neutron generator data as an additional term:

$$\mathcal{L}_{\mu_2} = \text{Gauss}(\mu_2 | \tilde{\mu}_2, \delta\mu_2). \quad (4.8)$$

None of the other parameters reach this level of importance, and we therefore convert them into an ‘effective rate uncertainty’. This parameter is essentially an additional scaling factor in the conversion from expected number of interactions in the TPC to expected number of events in the search data. Therefore, it has the same effect as an uncertainty on the efficiency of the selection criteria, and we combine the two to produce the effective efficiency,  $\epsilon_{\text{eff}}$ , which we encountered in equation 4.2, with the constraint

$$\mathcal{L}_\epsilon(\epsilon_{\text{eff}}) = \text{Gauss}\left(\epsilon_{\text{eff}} \mid \tilde{\epsilon}, \sqrt{\delta\epsilon^2 + \sum r_j^2}\right), \quad (4.9)$$

where  $\tilde{\epsilon}$  and  $\delta\epsilon$  are the estimated efficiency of the selection criteria and its uncertainty, respectively. The sum is over the seventeen remaining signal model parameters after excluding the mean ER cS2<sub>b</sub>.

## 4.4 Results

Before unblinding the full dataset and computing a limit, three predefined checks were performed. These were designed to identify any mistakes or mis-modelling while there was still time to fix them without unnecessarily biasing the result.

The first check is to unblind only a selection of bins where no WIMP signal is expected. These are chosen to be any bins in which only 1% or less of the total expected number of events is from WIMPs, if the cross-section is at the level of the 90% confidence upper limit set by XENON100 [198]. The  $\chi^2$  quality of fit of the data in these bins to the expectation is 21.5. This is compared to the 95th percentile of the expected distribution of  $\chi^2$  values, obtained using simulated datasets, which is 32.8. It is important to calculate the 95th percentile using simulated data instead of simply using the standard  $\chi^2$  distribution, because this simple test doesn’t take into account variations of the nuisance parameters.

The second check involves maximising the likelihood using the full science dataset, but without looking at either the events themselves or the best-fit rate. Instead, we check how much each of the parameters which has a constraint term in the likelihood has been ‘pulled’. This is measured as its best-fit deviation from the nominal value, divided by the standard deviation in the Gaussian constraint term:  $(\hat{x} - \tilde{x})/\delta x$ . Figure 4.9 shows the pull for each relevant parameter. It is clear that no parameter has been pulled very much from its nominal value (less than one standard deviation in every case). It is therefore unlikely that the result is impacted by one of these being incorrectly modelled.

The final check is intended to protect against severe mis-modelling of the 2 $\nu$ ECEC-KL and EC-L peaks. Since these are present at slightly lower energies than the expected signal, we look at the region between 29 keV and 35 keV. This contains a significant fraction of the two background peaks but no part of the central 2 $\sigma$  region of the expected signal. Again we compare the  $\chi^2$  quality of fit, 9.9, against the 95th percentile of the expected distribution, in this case 18.2.




Figure 4.9: Parameters' best-fit values, using the full unblind dataset, relative to their nominal values, measured in number of standard deviations.

The unblinded data is shown in figure 4.10, compared to the background-only expectation: the expected number of events when the signal rate is set to zero and all nuisance parameters fit to the data. Using this data, an upper limit is placed on the cross-section for inelastic WIMP- $^{129}\text{Xe}$  scattering, at 90% confidence level. This can be seen in figure 4.11, where it is also compared to the upper limit reported by the XMASS collaboration in [199]. The upper limit from this analysis is the most sensitive for WIMPs heavier than  $100 \text{ GeV}/c^2$ , and reaches  $3.3 \times 10^{-39} \text{ cm}^2$  for a mass of  $130 \text{ GeV}/c^2$ . Because of the lower expected event rate and the less efficient discrimination between electronic recoil backgrounds and the expected signal, the upper limit is not as strong as for the search for spin-dependent elastic scattering, where the strongest limit of  $6.3 \times 10^{-42} \text{ cm}^2$  is set for  $30 \text{ GeV}/c^2$  WIMP-neutron scattering.

Also visible in figure 4.11 is the expected range of upper limits. This is determined by simulating datasets from the best-fit models, with no signal. The illustrated  $1\sigma$  and  $2\sigma$  bands correspond to the central 68% and 95% ranges of limits computed for these datasets.

Especially around  $50 \text{ GeV}/c^2$ , there is a maximum of a roughly  $2\sigma$  upwards fluctuation of the limit. However, there is no significant ( $3\sigma$ ) evidence of an excess and we therefore only report an upper limit. That no particularly large excesses are visible in figure 4.10 suggests that this limit is produced by a genuine signal-like-nature of the dataset rather than a mis-modelling of one or more of the backgrounds. Figure 4.12 suggests the same; here, the bins are arranged according to their sensitivity to WIMP scattering, quantified by the expected signal content divided by the square root of the expected background.

In summary, the search for spin-dependent inelastic WIMP scattering using a  $0.89 \text{ t yr}$  exposure of XENON1T is the most sensitive search for this process to date, and the upper limit is the strongest for a wide range of WIMP masses.




Figure 4.10: Observed events in the full 0.89 tyr exposure shown as black error bars. The background expectation, from the likelihood best-fit, is shown in grey and the expected signal for a  $50 \text{ GeV}/c^2$  WIMP is shown in red, normalised to a cross-section of  $10^{-37} \text{ cm}^2$ . The five histograms between 29 keV and 55 keV correspond to the  $cS_{2_b}/cS_1$  bins used for the likelihood, which can be seen in figure 4.6.

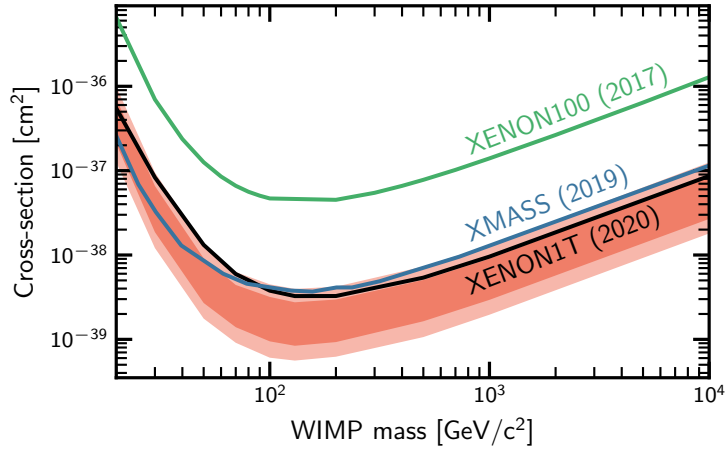


Figure 4.11: Limit on the cross-section of inelastic WIMP scattering off  $^{129}\text{Xe}$ , set using XENON1T data (black line) and the  $1\sigma$  and  $2\sigma$  range of expected limits (light and dark red shading, respectively). The limit is compared to those reported by the XMASS collaboration [199] (blue) and XENON100 [198] (green).




Figure 4.12: Observed events after grouping the bins according to their sensitivity to inelastic scattering of  $50 \text{ GeV}/c^2$  WIMPs. This is quantified using the ratio between the expected number of signal events and the square root of the expected number of background events. The colours have the same meaning as figure 4.10, but the signal is included in the fit and depicted with its best-fit rate. This grouping shows the origin of the  $2\sigma$  upwards fluctuation of the limit, as compared to the expected range: there is a roughly  $1\sigma$  excess in the final bin, and roughly  $2\sigma$  in the penultimate bin.

# Chapter 5

## PMT testing

The ‘eyes’ of XENONnT are its PMTs. These are what detect the tiny flashes of scintillation light – often only a few photons in the S1 – expected from a dark matter interaction. PMTs are a relatively old and simple technology, first invented in 1935 [200]. Nevertheless, more modern alternatives such as silicon photomultipliers have not yet managed to achieve comparable performance for low-rate experiments, although there is considerable research interest [144, 201, 202].

The PMTs used for both XENON1T and XENONnT are the R11410, manufactured by Hamamatsu. This 3” model was developed specifically for use in liquid xenon experiments, with an operating temperature range of  $-110\text{ }^{\circ}\text{C}$  to  $50\text{ }^{\circ}\text{C}$  and maximum pressure of 3 bar. Previous studies have shown stable operation in the conditions typically encountered in dual-phase xenon TPCs [203].

The R11410 is available in a variety of versions, which are the result of iterative improvements, after collaboration between both the LUX and XENON collaborations and Hamamatsu [204, 205]. The specific model used by XENONnT is the R11410-21, which uses extremely low radioactivity materials. In particular, the radioactivity of the ceramic stem, used to insulate the various dynode connections, was minimised. The PMT body is constructed from Kovar, an iron-nickel-cobalt alloy which has a high magnetic permeability. This effectively shields external magnetic fields which could otherwise affect the PMT’s operation [206].

Of the 494 PMTs needed to instrument XENONnT’s TPC, only 153 were used in XENON1T. Due to a relatively high rate of failure (79 of the 248 PMTs) during the operation of XENON1T, the new PMTs for XENONnT were tested much more thoroughly before installing them. Every photomultiplier was operated at cryogenic temperatures in both gaseous and liquid xenon. This was intended to simulate the conditions for PMTs used in both arrays of XENONnT – the top PMT array being above the liquid-gas interface. While cold, a variety of test data was collected for every PMT. The aim of this process was firstly to identify severely problematic PMTs which could not be considered for use in the detector, secondly to search for warning signs that could point to problems developing during future operation, and thirdly to provide a ranking of which of the ‘acceptable’ PMTs would be most preferred for eventual use.

The testing was carried out at three institutes, the Max Planck Institute for Nuclear Physics in Heidelberg, Stockholm University and the University of Zurich, the last two of which have facilities for cryogenic xenon testing. In this chapter, after an introduction to PMTs, we focus on the procedure used in Zurich, where in total 105 PMTs were tested, and the results obtained there. At the end of the chapter we look at these results in the context of the overall testing campaign and its implications for XENONnT.

## 5.1 Photomultiplier tubes

Conceptually, a photomultiplier tube can be thought of in two parts. The first is the photocathode, where incident photons are converted into electrons via the photoelectric effect. The second is an amplification stage, where the electron signal is multiplied by a factor of several million in order to output a measurable charge. We will treat each of these in turn.

### 5.1.1 Photocathode

The most important property of a photocathode is its quantum efficiency: the probability that when a photon reaches it, a photoelectron is released. This tends to prefer materials with a low work function, although with a lower work function there will also be a greater rate of thermal electron emission. Operating PMTs at low temperature reduces the latter. For the R11410, a bialkali (antimony-potassium-caesium) photocathode was developed by Hamamatsu [207]. This has a particularly high quantum efficiency for the low vacuum ultraviolet (VUV) wavelength of xenon scintillation light: 175 nm. According to the specification agreed between the XENON collaboration and Hamamatsu this is a minimum of 28%, but we find a much higher average value of 34% for the PMTs purchased for XENONnT.

In general, a single photoelectron is emitted per incident photon. It is also possible, however, for two photoelectrons to be emitted. This is known as double photoelectron emission, and happens for just over one fifth of emitted photoelectrons in the case of the R11410 [208].

A second property of a good photocathode is low resistance. In order to ensure a linear PMT response, all components must have low resistance. This is most important in the last stages of the dynode chain, as we will see later, but also relevant for the photocathode. Unlike the dynodes, the alkali metals used for this are chosen primarily for their work function and spectral response, not their resistivity. By adding a small amount of bismuth Hamamatsu was able to decrease the resistivity of the R11410-21's photocathode material.

### 5.1.2 Dynode design

In order to produce a measurable charge from a single photoelectron, repeated stages of electron acceleration which induces *secondary electron emission* are used. This happens in the dynode chain, a series of usually around ten electrodes. Each dynode is held at around

## Chapter 5 PMT testing

Table 5.1: Dynode voltages recommended by Hamamatsu for operation of an R11410-21 PMT, assuming a photocathode (K) voltage of  $-1500$  V and a grounded anode (P).

Dynode voltage [V]													
K	1	2	3	4	5	6	7	8	9	10	11	12	P
-1500	-1176	-1054	-892	-811	-730	-649	-568	-486	-405	-324	-243	-122	0
Difference from previous dynode [V]													
	+324	+122	+162	+81	+81	+81	+81	+81	+81	+81	+81	+122	+122

100 V higher potential than the last. Electrons are accelerated into every dynode in turn. Here, secondary emission of additional electrons (proportional to the number of incident electrons) results in a multiplication of the signal. The final electrode, the anode, is designed to collect the charge signal, which by then has been amplified by a factor of typically several million.

The relative voltages between the various dynodes are optimised by the manufacturer. The recommended voltages for the R11410 PMTs used by XENONnT are shown in table 5.1. The general trend here is similar for many PMTs: the first few and last few dynodes are separated by a greater potential difference than most. The reason for this can easily be explained qualitatively, as found in [209], for example, and summarised here. At each dynode stage, some of the electrons might escape and either be lost completely or skip straight to the following stage. If this happens to the photoelectron at the first dynode, no signal will be produced at the anode. This potential signal loss is represented by the *collection efficiency*, which is around 90% for the R11410 [206]. Even in the second and third stages, it is especially important for the electrons to reach every dynode. There are fewer electrons propagating here, so random effects are not averaged out so effectively and fluctuations impact the overall resolution more severely.

The last few stages also benefit from a higher voltage difference, for a different reason. These stages are where the number of electrons is largest. Space-charge effects, or the effect on the propagating electrons of their own electric field, are most severe. By increasing the external electric field, the relative effect of these can be minimised. This helps to ensure a constant amplification factor for a wide range of signal sizes, or in other words a *linear response*.

Linearity is important for accurate event reconstruction, especially for determining events' energies. As we saw in chapter 2, the energy deposited in an interaction can be found from a linear combination of the S1 and S2 signals. This computation becomes substantially more involved if the signal measurement itself is non-linear. For analysing XENON1T's data, a 'desaturation correction' was used for high energy events (although the limited dynamic range of the digitisers was the dominant effect on linearity, not the PMTs). However, this involved essentially discarding much of the information from the most illuminated photomultipliers and replacing it with a template [174]. This loss of information inherently leads to a lower energy resolution than might otherwise be achievable.






Figure 5.1: The MarmotX facility. Left: a rendering of ten 3" PMTs within their support structure. The pairwise arrangement, with two PMTs always facing each other, can be seen. Right: the PMT support structure seen within the vacuum-insulated cryostat, and connections to the PTR and cable feedthroughs on top. Credit: A. James.

## 5.2 The MarmotX PMT evaluation facility

In order to carry out repeatable and relevant qualification tests, it is important to have a facility with stable and accurate temperature control, efficient cool-down & warm-up and xenon filling & recovery procedures, and the ability to record PMT waveforms with and without a calibration LED. In Zurich the MarmotX facility, seen in figure 5.1, provides such an environment. This section provides an overview of the facility itself, while the next gives details of the tests. A detailed description of MarmotX can also be found in [210].

The core of MarmotX is a vacuum-insulated cryostat with space for ten 3" PMTs, arranged in two layers of five. Cooling is provided by an Iwatani PDC08 pulse tube refrigerator (PTR), providing 24 W of cooling power at 164 K. Since the cooling power of the PTR cannot be controlled, two heaters are attached to the cold head. A CyroCon 32 proportional-integral-derivative (PID) temperature controller supplies power to the heaters in order to maintain a stable cold-head temperature. An independent backup system uses liquid nitrogen to cool two copper blocks. These are attached to two of the tubes used for feedthroughs into the inner vessel. A solenoid valve opens when a threshold pressure, typically 2.5 bar, is reached inside the cryostat, allowing liquid nitrogen to flow to the copper blocks. Since the flow is not controlled, this system is only used in case of emergency such as power failure. A final safety mechanism is a spring-loaded pressure-release valve which activates at 3 bar, the maximum safe pressure for R11410 PMTs. If the pressure rises above this level, xenon is vented into the room until an acceptably low pressure has been reached.

The internal environment is monitored by eight Pt100 temperature sensors at a variety of depths and a vacuum pressure gauge. These are read out by custom software running on linux. During operation the cold head temperature is generally chosen such that the internal pressure is  $\sim 2.3$  bar. This is close to the pressure experienced by PMTs in the bottom array of XENONnT.

An Iseg EDS-F-130n 16-channel high voltage module provides up to  $-3$  kV to power the PMTs. PMT waveforms are recorded at a sampling rate of 100 MHz by two CAEN V1724 digitisers, after amplification by a factor of ten with a Philips 776 amplifier. An additional copy of the amplified signal is fed into a CAEN N845 discriminator, whose output passes to a CAEN V260 scaler. The discriminator threshold is set to 11 mV, roughly half the amplitude of a single photoelectron signal from a PMT operated with a gain of  $5 \times 10^6$ . The scaler is used to gain an accurate measurement of the total rate of signals from each PMT. A blue LED, within the cryostat, is used in conjunction with a pulse generator for calibrating PMTs.

The two layers of PMTs are arranged in pairs, one above and one below. Within a pair, the two PMTs are facing one another (window-to-window) and separated by only a small gap. The PTFE plate between the two layers, which acts as both a reflector and spacer, has circular holes to ensure that paired PMTs can see one another. This is intended to ensure that light emitted by one PMT can be seen by its partner, as detailed in section 5.3.1.

### 5.3 The XENONnT PMT testing campaign

All the new PMTs (those which were not used for XENON1T) were tested in liquid xenon, either at Stockholm University or the University of Zurich. A subset of those tested in Stockholm were also tested at the Max Planck Institute for Nuclear Physics in Heidelberg, where a cold nitrogen/argon facility is used instead. The focus of the Heidelberg tests was slightly different, with the possibility to characterise properties such as the transit-time spread. A manuscript providing an overview of the full testing campaign at all three centres is being prepared; here we concentrate only on the tests performed with MarmotX at the University of Zurich.

During each test cycle, the MarmotX cryostat, with ten PMTs to be tested, was cooled down twice. The first time, the inner cryostat vessel was filled with sufficient liquid xenon to completely cover both layers of photomultipliers (19 kg). The second time, only a small amount of liquid xenon was filled so that the PMTs were in a gaseous xenon atmosphere very close to the phase transition temperature. These two conditions should mimic those for PMTs in the bottom and top arrays, respectively, when they are being operated in the XENONnT dual-phase TPC.

In both cases the cryostat was held at a constant (cold) temperature for one week before warming up and extracting the xenon. During that time a variety of diagnostic tests were carried out, which are discussed in the remainder of this section. The aim of those tests was to identify which photomultipliers are suitable to be used in XENONnT, avoiding any that

may cause problems for stable operation.

### 5.3.1 Light emission

Light emission is a process whereby a PMT can produce individual photons, at a wavelength such that the PMTs themselves are sensitive to it. Previous studies have shown that this can be observed in R11410 PMTs [80, 211, 212]. The origin of light emission is not conclusively understood, but Hamamatsu have suggested that it may be the result of charge-up of the PMT's ceramic stem [213]. The light can either be seen by the same PMT in which it is produced (effectively a form of dark count) or by another one. Some degree of light emission is present in all PMTs but the rate varies considerably. If the rate is too high, multiple PMTs may observe these small signals in coincidence and they could appear like a scintillation signal. This could produce an additional background to searches for rare, low-energy interactions.

We identify two categories of light emission. *Micro light emission* is a relatively small rate of single photon emission. This happens continuously while a PMT is turned on and affects all PMTs, although to differing degrees. The rate of micro light emission depends on both the temperature and the voltage that a PMT is operated at [211].

Certain PMTs exhibit a second kind of behaviour, whereby they emit light at a much greater intensity but not continuously. This can be in the form of *flashes*, as identified in [150]. These can be identified by sudden increases in the trigger rate of all PMTs (but especially the one which is flashing), followed by an exponential decrease lasting several minutes or hours. The origin of flashes is not fully known, but it is thought to involve discharges within the PMT [212]. Alternatively, PMTs can have *intermittent light emission*, meaning there are extended periods (several hours or days) during which they emit high intensity light, but without a characteristic exponential drop-off. Both cases are often induced by bright illumination – in XENON1T this was often due to a muon interacting in the TPC [80] – but can also occur apparently spontaneously.

We performed a dedicated test to detect micro light emission, described in the following. Imagine we want to test a certain PMT (labelled here as PMT A) for light emission. We measure the trigger rate in the PMT which is facing it (PMT B), as we adjust the power supply voltage of PMT A. If we see that the the trigger rate of PMT B decreases while we turn down the voltage supplied to PMT A, we can conclude that the difference in trigger rate is due to light emitted by PMT A. To quantify the amount of micro light emission, we record the difference in PMT B's trigger rate when PMT A's supply voltage is increased from 900 V to 1500 V. It is difficult to convert this into an absolute measure, since it depends on the quantum efficiency of PMT B as well as geometrical effects. Nevertheless, this is sufficient for comparing PMTs and identifying those with unusually high emission rates. Previous testing campaigns defined the amount of light emission slightly differently, as the rate difference between 0 V and 1500 V [150]. However, we observed a slight increase below around 900 V, presumably due to sparking between the two PMTs, and therefore use this measure instead.




Figure 5.2: Left: distribution of PMTs' micro light emission rates, defined as the difference in rate (seen by the opposite PMT) between 0.9 kV and 1.5 kV operating voltage. Right: example of an individual measurement, showing the opposite PMT's trigger rate as a function of the operating voltage.

The micro light emission distribution for all PMTs tested in Zurich is shown in figure 5.2, together with an example of the measurement for a particular PMT. The worst PMTs, with rates of several hundred Hz or more, were not used in the XENONnT PMT arrays.

Strong light emission can spontaneously begin at any time. For that reason, the trigger rate of all PMTs is continuously monitored while they are being tested, for a total of at least two weeks, and checked for unexplained increases. It can also be induced, and we therefore also perform a 'stress test' twice for each PMT: once in liquid and once in gaseous xenon. This is designed to simulate the most intense light that PMTs experienced in XENON1T, which is during  $^{220}\text{Rn}$  calibrations. The stress test consists of two twelve-hour long periods. During the first, the PMTs are illuminated with short (i.e. S1-like) light pulses creating signals of a few hundred photoelectron signals. The second twelve hours is intended to represent S2 signals, and the PMTs are illuminated with  $2\ \mu\text{s}$  pulses creating around  $10^4$  photoelectrons.

Figure 5.3 shows an example of intermittent light emission, occurring during a stress test. In this case, one of the PMTs begins emitting light around ten hours after the start of the stress test, which is seen by the opposite PMT at a high rate. Towards the end of the test, all PMTs trip. This is presumed to be due to sudden bright illumination, most likely from the light-emitting PMT. All PMTs are turned back on at the end of the stress test, and light emission continues. In order to verify which is the light emitting PMT, it and the PMT opposite it are each turned off in turn. When the opposite PMT is turned off, only its trigger rate decreases. On the other hand, when the light emitting PMT is turned off, both trigger rates decrease significantly.

In total we observed four such cases of intermittent light emission out of the 105 PMTs tested. For three of these it was possible to identify a period of LED illumination which first triggered the light emission. In the other case, the light emission appeared to begin spontaneously, but became much worse following a stress test. Of the four, two displayed sufficiently strong light

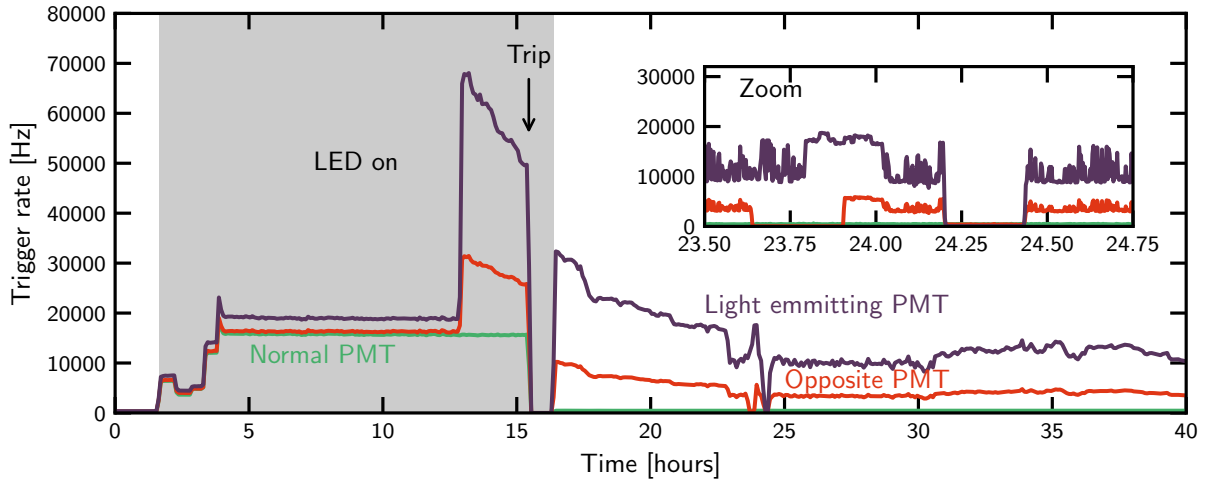


Figure 5.3: Intermittent strong light emission, beginning during a stress test in liquid xenon. The light emitting PMT (purple) shows the highest trigger rate, and significant light is also seen by the PMT facing it (red), whereas another separate PMT (green) sees only little light. The grey shaded region indicates the time that the LED was turned on (pulsed). This includes a short period of lower-intensity illumination, for afterpulse characterisation, followed by the higher-intensity stress test. The inset shows a zoom into the time when the light emitting PMT and the one opposite it were turned off in turn.

emission that they were not considered for XENONnT.

### 5.3.2 Afterpulses

It is well known that PMTs can produce additional signals shortly after the usual pulse from incident photons. These are referred to as *afterpulses*. A variety of types of afterpulse, with distinct origins, have been proposed. These include elastic scattering of electrons off the first dynode [209] and electron-induced luminescence within the dynode chain [214], but both of these produce only fairly small afterpulses soon after the main peak and they are unimportant for us. Here we concentrate on afterpulses produced by ionisation of residual gases in the PMT vacuum.

If a photoelectron encounters an atom of gas between being emitted from the photocathode and reaching the first dynode, it can ionise it. The positively charged ion will then drift back towards the photocathode and upon impact releases several electrons. For the electric potential present in PMTs such as the R11410 ( $V \propto s^2$ , where  $s$  is distance from the cathode), the time take for an ion to reach the cathode is almost independent of where it is produced [150], and is given by:

$$t = 1.134 V^{1/2} \mu\text{s cm}^{-1} \sqrt{\frac{L^2 M}{V_0 Q}}, \quad (5.1)$$

where  $L$  is the distance between the photocathode and first dynode,  $V_0$  is the voltage difference between the photocathode and first dynode and  $M/Q$  is the mass-to-charge ratio of the

ion, measured in atomic mass units per elementary charge.

In themselves, afterpulses pose only a minor problem for the experiment, provided that they are properly accounted for when modelling the detector's response. However, during the operation of XENON1T, it was observed that PMTs with especially large afterpulse rates are more susceptible to flashing than others [80]. In the more severe cases this required the PMT to be turned off. Furthermore, the afterpulse rate can increase over time if xenon is able to leak inside the PMT. This was also observed in XENON1T, and indeed most PMTs which had to be turned off showed evidence of a leak. A major aim of the testing campaign was therefore to identify any leaky PMTs.

Given that xenon has such a low natural abundance in air, there is no reason for xenon contamination to be present after the PMTs have been manufactured. We can safely assume that any xenon-related afterpulse, even at a small rate, indicates that xenon leaked into the PMT during the test. This idea, of using afterpulses to check for leaks, has been used before by both XENON1T and PandaX-I [150, 212].

In order to characterise a PMT's afterpulsing, we illuminate it with a reasonably large pulse of light from the LED. The larger the pulse, the better the sensitivity, since each photoelectron in the main signal can independently produce an afterpulse. All PMT signals in a  $5\ \mu\text{s}$  window after the main pulse are recorded.

The afterpulse characteristics of a PMT are often represented as a two-dimensional histogram with the time delay between the main LED signal and the afterpulse on the  $x$ -axis and the area of the afterpulse, in photoelectrons, on the  $y$ -axis. An example, for a leaky PMT, can be found in figure 5.4. Three categories of afterpulses can be identified, which are labelled A1, A2 and A3 in [150]. There are small afterpulses with a charge of  $\mathcal{O}(1\ \text{pe})$  present at all times (A2). At small delay times the tail of a second population (A1) of small afterpulses is visible. These are possibly due to back-scattering of electrons off the first dynode or photoluminescence, as mentioned above. Most afterpulses of this latter type are not visible, since their delay time is so small that they overlap with the main signal.

The third category (A3) is most interesting for this study, and consists of substantially larger afterpulses. These are found at characteristic delay times in the microsecond range, and are due to ionisation of residual gases. The upper  $x$ -axis of figure 5.4 indicates the mass-to-charge ratio of the ions which would be expected to cause afterpulses at each time delay. A xenon afterpulse is clearly visible at a delay time of  $\sim 2.7\ \mu\text{s}$ .

In order to search for and quantify the size of xenon afterpulses, we first select only those signals with an area of at least 2 pe. By doing this, most of the population of small afterpulses can be removed while keeping most ion-induced signals.

A histogram of the remaining afterpulses' time delays is then produced (as in figure 5.5, for the same example PMT as before). The region where a xenon afterpulse would be expected is fit with a Gaussian peak on top of a linear background. This doesn't perfectly describe the data in every case, since there is a complex background of afterpulses at all time delays, whose shape varies from PMT to PMT. However, it is sufficient for identifying leaky PMTs,




Figure 5.4: Afterpulses' size and time-delay, for an example leaky PMT (left) and a good PMT (right). The three different populations of afterpulse can be seen. On the left, a large xenon afterpulse is clearly visible with a delay of roughly  $2.7 \mu\text{s}$ . The upper x-axis shows the mass-to-charge ratio of ions expected to produce afterpulses at that delay time.

and gives a reasonable picture of their severity.

It is then possible to compute the afterpulse rate, defined as the fraction of photoelectrons in the main signal which induce an afterpulse. If  $n_{\text{Xe}}$  events contain a xenon afterpulse, out of a total of  $N$ , and the average number of photoelectrons in the main pulse is  $p$ , Poisson statistics tells us that the afterpulse rate is given by:

$$r = -\frac{1}{p} \ln \left( 1 - \frac{n_{\text{Xe}}}{N} \right). \quad (5.2)$$

For the example shown in figure 5.5 this gives a rate of 0.46%. Of the 105 PMTs tested in Zurich, xenon afterpulses were found in two of them; the other had a rate of 0.45%.

A significant improvement has been made to the rigour with which we searched for afterpulses, compared to the testing campaign used before constructing XENON1T. At that time only a subset of all the PMTs were tested in liquid xenon. A larger LED pulse is used now of around 100 pe, compared to only around 1 pe during the XENON1T tests. In addition, we have increased the total number of LED pulses from  $10^5$  to  $10^6$ . With these changes we estimate that we are sensitive to afterpulse rates down to around  $10^{-5}$  %.

While testing, we also search for argon ion afterpulses. It is fairly common to find afterpulses due to residual argon gas, which is used during the production process. Noble gases cannot be removed by the PMT's built-in getter, so any argon inside after production will remain in the PMT's vacuum. We have previously observed that PMT's with unusually large argon afterpulse rates are more likely to be leaky, although the causal relationship is unknown. To determine the argon afterpulse rate, the same method is used as described above for xenon. The resulting distribution of rates for the 105 PMTs tested is shown in figure 5.6. Those PMTs with a measurable argon afterpulse were assigned 'penalty points' (see below) according to its rate. They were therefore either not used or placed in less critical positions within the XENONnT TPC.

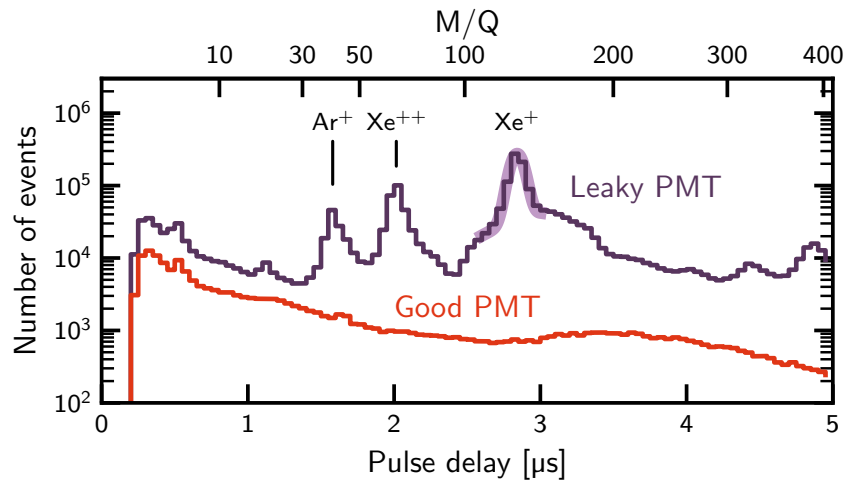


Figure 5.5: Example afterpulse time-delay distributions, for the same leaky PMT (purple) and good PMT (red) as in figure 5.4. The best-fit model for the peak due to  $\text{Xe}^+$  ions is shown as a thick, light purple line. Prominent peaks are labelled according to their expected delay times.




Figure 5.6: Distribution of fitted argon afterpulse rates of all PMTs tested in Zurich.



Table 5.2: Summary of main problems identified in the new PMTs for XENONnT during the testing campaign.

	Tested	Leak identified	Strong light emission	Other problem
In total	368	13	11	2
Tested at UZH	105	2	2	1

## 5.4 Summary

In total 368 new PMTs were tested for XENONnT, 105 of which at the University of Zurich. Each PMT was tested over a period of at least one week in liquid xenon followed by another week in gaseous xenon. This mimicked the conditions for the bottom and top PMT arrays in a dual-phase XENON TPC.

During the test, particular attention was paid to two known potential problems: light emission and afterpulsing, especially that related to xenon atoms in the vacuum which indicates leaks. Table 5.2 summarises the PMTs which ‘failed’ the test – those with severe enough problems to immediately be removed from consideration for use in the XENONnT TPC.

The PMTs were tested much more thoroughly than they were for XENON1T, when only a selection were tested in xenon. We estimate that we are sensitive to xenon-related afterpulse rates on the order of  $10^{-5}$  % or larger. This would be sufficient to detect the smallest identified leaks in XENON1T PMTs, over the course of the two weeks in xenon. Although we cannot rule out the possibility that new leaks opened between the tests and installation in XENON1T, this gives us confidence that the overall fraction of leaky PMTs will be substantially smaller in XENONnT.

Two PMTs were rejected because we were unable to power them on, due to an extremely high leakage current. These PMT was presumably damaged between being produced and arriving at their respective testing facilities.

For the remaining 342 PMTs, a system of ‘penalty points’ was developed to describe their performance overall. Each problem was assigned a certain number of points, including micro light emission, low-rate intermittent light emission and larger than average argon afterpulse rates. In addition, points were given to PMTs which were manufactured in batches containing several other PMTs which had developed serious problems. After removing those with serious problems, 197 PMTs from XENON1T were also included in this penalty point system (including 32 which were spare). The overall distribution of penalty points is shown in figure 5.7. In total, 497 PMTs had at most 2.0 penalty points; 494 of these were selected to be used for the XENONnT PMT arrays. The number of penalty points was also considered when arranging the PMTs within the two arrays. In the next chapter, we will see details of the construction and installation of the PMT arrays and their associated cabling.




Figure 5.7: Distribution of penalty points awarded to all PMTs, including those left over from XENON1T as well as those acquired for XENONnT, whose testing is described in this chapter.

# Chapter 6

## XENONnT

The next stage in the XENON programme is XENONnT, an evolution of XENON1T, which is being commissioned as this thesis is being written. When XENON1T was designed, attention was already being paid to a possible future upgrade, with a significantly larger amount of xenon. The total amount of xenon will be 8.4 tonnes, or 2.6 times as much as before, and the increase in size of the TPC itself is even greater, from 2.0 t to 5.9 t of instrumented xenon. In addition to the much larger TPC, a new neutron veto and radon distillation system will enable improved nuclear recoil background rejection and decrease the electronic recoil background rate, respectively.

In this chapter we will first see an overview of XENONnT, concentrating on what is new and what is needed to make sense of the rest of this chapter. More details can be found in [99]. Then we will move on to some details of the PMT arrays and cabling, starting with an overview of the bases' design, before discussing their production and preparation for installation in XENONnT. Later in the chapter is information about the cabling which connects these bases to the data acquisition system, in the service building adjacent to XENON's water tank. Finally we will see some details of the assembly and installation of the PMT arrays and related cabling, now undergoing commissioning at LNGS. Here we attempt to include sufficiently detailed information to be useful as documentation of the cabling while avoiding unnecessary minutiae.

### 6.1 Overview of XENONnT upgrade

The following presents only a summary of the most important aspects of the XENONnT upgrade. By increasing the sensitive mass by almost a factor of three, the fiducial volume used for analysis may be almost four times as large. Much of the infrastructure from XENON1T is being reused, thus accelerating and simplifying the upgrade: the cryogenic and purification systems are being kept with a few important upgrades, the water tank and muon veto are unchanged (but see details of the neutron veto below), and the diameter of the outer cryostat is the same so that the old vessel can be reused with a simple extension.

A new neutron veto system in XENONnT will make it possible to reduce the nuclear recoil background significantly, providing an 87% efficiency for tagging neutrons which have

scattered once in the TPC. The neutron veto is based on the same technology as the existing muon veto, with 0.2% gadolinium ions being added to the water. An inner region of the water tank will be optically isolated using expanded PTFE (ePTFE) reflectors, and instrumented with 120 eight-inch photomultipliers to detect the Čerenkov light which eventually follows neutron capture by gadolinium [215].

As already mentioned, the cryogenic and purification systems from XENON1T are largely being inherited, but some significant improvements have been made. These include new magnetically-operated piston pumps to recirculate the xenon, which introduce significantly less radon into the system [216]. A dedicated radon distillation column will make it possible to remove radon from xenon being recirculated into the TPC with a high efficiency, effectively mitigating all sources of radon outside the cryostat itself [217]. With this distillation column, it should be possible to reduce the  $^{222}\text{Rn}$  activity to around  $1 \mu\text{Bq kg}^{-1}$ .

A new liquid xenon purification system, using a filter to remove oxygen and other electronegative impurities from liquid xenon without needing to boil and re-condense it. This enables a much faster recirculation rate, meaning the xenon is cleaned more often. As a result, XENONnT should have a longer electron lifetime of around  $1000 \mu\text{s}$  (compared to around  $650 \mu\text{s}$  in XENON1T).

A full simulation-based study of the sensitivity of XENONnT can be found in [99], considering the upgrades mentioned above. Combined with the lower target krypton concentration of 0.1 ppt, the ER background rate in XENONnT is expected to be reduced by around a factor six, to  $(13.1 \pm 0.6) \text{keV}^{-1} \text{t}^{-1} \text{yr}^{-1}$ . This will make it sensitive to elastic scattering cross-sections more than an order of magnitude better than XENON1T, with an expected sensitivity, to  $50 \text{GeV c}^{-2}$  WIMPs, of  $1.4 \times 10^{-48} \text{cm}^2$ .

### 6.1.1 The XENONnT TPC

An illustration of XENONnT's TPC can be seen in figure 6.1. The walls of the TPC are made up of a total of 48 overlapping PTFE panels, with diamond-polished surfaces to maximise their reflectivity to scintillation light. The sensitive region of the TPC, when cold, is 149 cm in height (between the two electrodes defining the drift region) and 133 cm in diameter (measured between two opposite panels of the wall, the exact value varies since the TPC has a 48-sided, rather than a circular, cross-section).

In total the TPC contains five electrodes: three defining the drift and extraction fields and two 'screening electrodes' at the top and bottom which prevent the electric fields from leaking into the PMT arrays. These are all built from parallel stainless steel wires with a diameter of  $216 \mu\text{m}$ , except the cathode with a diameter of  $304 \mu\text{m}$ . The gate and anode have additional wires (two and four, respectively) running perpendicularly across the others. These are designed to reduce any deformation of the electrodes due to the electrostatic forces between them. This minimises the position-dependent field dependence, which is what makes an S2 ( $x, y$ )-dependent correction like we saw in chapter 3 necessary.



Figure 6.1: Illustration of the XENONnT TPC. Figure from [99].

Outside the PTFE walls, copper field shaping rings ensure a good uniformity of the drift field. There are two sets of these in XENONnT. The inner set consists of 71 wires with a diameter of 2 mm, which touch the outer surface of the PTFE. These are intended to mitigate charge build-up on the PTFE, by providing a short path to a conductor. The other set contains 64 rings, with a rounded rectangular 5 mm by 15 mm cross-section. These help with field uniformity near the boundaries of the TPC, by preventing electric fields from leaking inside the wires.

The significantly larger diameter than the 96-cm-wide XENON1T TPC leads to a requirement of 494 PMTs, 246 more than were used before. We saw details of the PMTs and their qualification tests in the previous chapter. We will first discuss the connections which provide the high voltages needed to operate them and a way to read out their signals, and later in this chapter we will see details about the construction and installation of the PMT arrays.

## 6.2 PMT Bases

The PMTs are powered and read out through their *bases*, small printed circuit boards connected directly to the PMTs. These are essentially voltage dividers, which supply the various voltages needed for the PMT's dynodes from a single externally applied high voltage. The bases are also the first step in the signal readout chain, forming the bridge between the anode of the PMT and the cables carrying signals to the data acquisition system.

The design of the voltage divider circuit has been inherited from XENON1T [218], where it

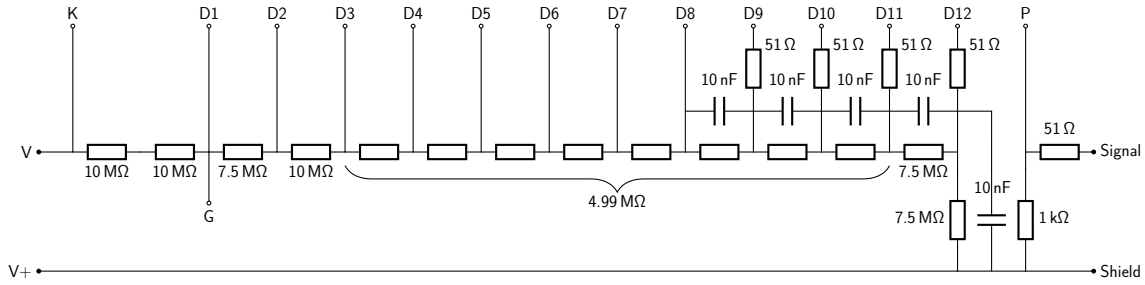


Figure 6.2: Base voltage divider circuit used in XENON1T and XENONnT, as developed in the framework of [218].

worked successfully during the lifetime of the experiment and no unforeseen problems were identified. The circuit diagram used is shown in figure 6.2.

The most important factor in the base design is the relative voltages of each dynode, which directly impact the way electrons are propagated through the dynode chain. Since the base, at its core, is a voltage divider, this is determined by the ratio of resistances between the dynodes. These values are recommended by the manufacturer in order to optimise the PMTs' performance.

A second, somewhat more intricate consideration is linearity. We discussed the importance of a linear PMT output in the previous chapter, and here we consider the influence of the base design on this.

In the most simple case, where a base consists purely of a chain of resistors forming a voltage divider, it begins to saturate when the current flowing between the dynodes is comparable to the *base current*: the steady current flowing through the voltage divider. The base cannot then provide enough current to the dynodes without their voltage being affected. It is worth noting that by reducing a PMT's operating voltage, and thus its gain, saturation is also reduced since the base current is proportional to the voltage, while the gain increases more quickly.

The simplest way to improve linearity is to decrease the total resistance of the base, thereby increasing the current which flows through it. This has the side effect, however, of increasing the power dissipated in the resistor chain. Particularly when the base is operated in a cryogenic environment, the power dissipation must be kept fairly small, so there is a trade-off to be found between linearity and heat production. The XENONnT bases dissipate 0.024 W each (when operated at 1500 V), or 12 W in total. For comparison, the total heat load in XENONnT is predicted to be around 275 W.

Other tricks are possible, in order to improve base linearity while keeping a low base current. These include the use of transistors or Zener diodes to stabilise the last few dynodes' voltages, having separate power supplies maintain the voltage on the final dynodes, or even replacing the voltage divider with a Cockroft-Walton generator. These techniques and their implications are discussed in detail in [209, 214]. None are optimal for a low-background setting, since all involve introducing extra material – active electronic components or sub-

stantial additional cabling – which can significantly affect the radioactive background in the detector.

A final technique involves adding decoupling capacitors in parallel to the last resistors in the chain. In a TPC the PMTs are illuminated only by brief pulses of light, with relatively long dark periods in between. Therefore, although the peak output current can be high compared to the base current (during a pulse), the average output current is much lower. The decoupling capacitors can provide the higher peak current and limit saturation, and are recharged slowly between pulses. The number of photoelectrons per signal for which saturation will begin,  $n_{\text{sat}}$ , then depends on the capacitance  $C$ , the voltage  $V_{\text{D12-P}}$  between the final dynode and anode and the gain PMT's gain  $G$  according to [218]:

$$n_{\text{sat}} \sim \frac{CV_{\text{D12-P}}}{G \cdot e}.$$

As can be seen in figure 6.2, the last few dynodes of the XENONnT bases are connected by 10 nF capacitors in parallel with the resistor. This means we can expect signals with areas on the order of a million photoelectrons or more to saturate, assuming a gain of around 5 million at a 1500 V operating voltage. Indeed, in XENON1T data, base saturation was observed to begin at around this point [174]. The dominant effect was saturation of the analogue-to-digital converters (ADCs) used to digitise signals, which happens for pulses with a charge greater than around  $10^4$  pe. New dual gain amplifiers installed in XENONnT, with a  $\times 0.5$  channel in addition to the  $\times 10$  already used, will increase this threshold by a factor of 20, meaning that base saturation will take on a greater importance.

A few more features are important in the base circuit design. Large enough pulses can result in a phenomenon known as *ringing*, high-frequency oscillations in the falling part of a pulse. This can be mitigated by using damping resistors for the last few dynodes [209, 214]. Signals from the PMTs are carried by several-metre long coaxial cables to the data acquisition system. To avoid reflections, the output impedance of the base must be matched to the cable impedance (which must in turn be matched to the input impedance of the amplifiers at the other end). This leads to the 50  $\Omega$  resistor at the signal output of the base.

Of course, only the relative voltages between a PMT's dynodes are important, the absolute voltage range makes no difference to its performance. This means that the PMT can be operated with either the anode or the photocathode connected to ground\*. We use the former option, which has the advantage that signals from the PMT, which originate from the anode, are fluctuations around ground and easily amplified and digitised. The ground connection is also generally less susceptible to high-frequency noise which is often present in the output of high-voltage power supplies, so by connecting the anode to ground, signals are recorded on top of a less noisy baseline. A corollary of the anode being grounded is that the photocathode is at a large negative voltage, and therefore so is the exterior of the PMT. This is not a problem, but does mean that the PMT must be electrically insulated from any surrounding conductive structures.

---

\*It would also be possible to operate a PMT with neither anode nor photocathode at ground, but this would be unlikely to be useful.




Figure 6.3: Left: A complete base, before attaching cabling. Right: schematic of a base PCB, showing traces on the front (red), back (blue), and connections between the two layers (green).

### 6.2.1 Production

We have seen a description and explanation of the base design, which has been kept the same as that used for XENON1T. Where possible, the materials used to produce the bases were also the same. In the following we discuss briefly these materials, their screening for radioactivity levels and the production of the bases.

By far the largest mass of material per base (not including cables) is the substrate upon which the printed circuit board (PCB) is produced. Standard PCB substrate materials, such as FR-4 (a grade of fibreglass), are not suitable for operation in cryogenic environments and are not generally radioactively pure. Two alternative materials were considered for XENON1T: CuFlon (copper-plated PTFE, produced by Polyflon) and Cirlex (a laminate produced from Kapton film). Although CuFlon was found to be somewhat cleaner [158], the PCB traces were found to easily separate from the substrate and break during cooling [218].

The PCBs were produced by Fralock in the USA, who are also the exclusive manufacturer of Cirlex. The Cirlex was screened for radioactivity both before and after printing the traces, using the GeMPI gamma-ray spectrometer [155].

Where possible, identical electronic components were used to produce the bases as were used for XENON1T. Enough spare of each component was procured for it to be screened with ICP-MS (see section 2.5). Comparing the screening results to those of the components used in XENON1T showed similar levels of radioactivity, as seen in table 6.1.

Components were soldered to the PCBs by Elfab in Switzerland. All materials were supplied by us, including the solder paste: Genma NP303-CQS-2, which had been screened by ICP-MS. Figure 6.3 shows a finished base, with components but before adding the coaxial and Kapton cabling (see section 6.3 for details about the cabling itself). A total of 567 such bases were produced, accounting for some spare above the 494 needed.

The fully assembled bases were screened again, using the GeMSE facility [156], after being



superficially cleaned by immersion in and wiping with ethanol\* and bagged. The screening results are compatible with the sum of the individual electronic components' contributions (the PCB itself is negligible in comparison). The base radioactivity is comparable to or lower than the PMTs for all isotopes considered: for example an activity of  $(650 \pm 25) \mu\text{Bq}$  compared to  $(600 \pm 100) \mu\text{Bq}$  for  $^{226}\text{Ra}$ , and  $(287 \pm 45) \mu\text{Bq}$  compared to  $(1200 \pm 200) \mu\text{Bq}$  for  $^{40}\text{K}$ . Individual screening results for each component can be found in table 6.1. Detailed simulations based on these results show that the bases are expected to contribute 5.7% of the total electronic recoil background rate and 5.5% of the nuclear recoil background [99]. The PMTs themselves contribute 46.1% and 23.6% of the two backgrounds, respectively.

After production each base was tested under the application of high voltage at 1500 V to ensure correct operation. We found that 66 of the 567 produced failed this test. The majority of failures were due a broken connection in one of the three *vias*, which connect the two layers on either side of the PCB. These were fixed by the electronics workshop of the University of Zurich physics institute, by soldering a small length of wire into the hole, ensuring a good connection. After the repair, bases were re-tested and found to work without problem. More rarely a component was missing from the final base; in this case a spare base was used instead. A number corresponding to its position in the PMT array was laser-engraved into each base, allowing its individual test results to be traceable even after installation.

The signal and HV cables were also soldered to the bases by Elfab, using Stannol Flowtin KS115 solder wire, also screened by ICP-MS. Before attaching the cables the bases were arranged on acrylic support structures, which replicate the placement of PMTs in the two arrays. Each acrylic structure contained all bases corresponding to one PMT array sector, a subsection of the array, as we will see in section 6.3.3; these acrylic structures can be seen in figure 6.4. By doing this it is was possible to arrange the cables in the final configuration and simplify the assembly of the arrays at LNGS. The acrylic structures doubled as protection while transporting and cleaning the bases.

## 6.2.2 Cleaning

As with all components the bases were cleaned inside the cleanroom at LNGS before the assembly of the PMT arrays. Due to the variety of materials used, a fairly simple procedure was followed in order to prevent chemical damage to the bases. They were first immersed in a soap bath for 15 minutes (see table 6.6 for details), before being rinsed thoroughly with deionised water. Whereas most components are subjected to ultrasound during their soap cleaning, there was concern that this could damage the fragile bonding of the PCB traces to the Circlex base substrate. However, it was deemed desirable to perform the full ultrasonic cleaning on the cables, which are soldered directly to the bases, especially due to their large surface area. Therefore the bases were cleaned twice: first with only the cables immersed in

---

\*It is worth mentioning here that while cleaning the bases before screening, a small amount of brown residue was noticed on the sides. At the time not much was thought of this, and it was assumed that it would be removed during the later, more thorough cleaning, but it was discovered again after the PMT arrays had already been built – see section 6.4.2 for more details.

Table 6.1: Measured radioactivity of the components used to produce bases. Some gaps are present due to the varying measurement and analysis techniques used. Results for components used in XENON1T bases are also shown, in blue, as are average values for a Hamamatsu R11410-21 PMT (as measured for XENON1T) for reference [158]. The first highlighted row, “Total (no PCB)”, shows the totals expected per base (except the PCB itself), based on the measurements of individual components. The second highlighted row, “Total with PCB”, shows measurements of the finished bases, including components.

Component	Radioactivity per component [ $\mu\text{Bq}$ ]							
	$^{238}\text{U}$	$^{226}\text{Ra}$	$^{228}\text{Ra}$	$^{228}\text{Th}$	$^{40}\text{K}$	$^{60}\text{Co}$	$^{137}\text{Cs}$	$^{210}\text{Pb}$
10 M $\Omega$	< 30	4.7(4)	5.3(6)	4.7(4)	14(3)	< 0.116	< 1.135	70(6)
(1T)	< 40	5.4(4)	7.0(7)	5.7(6)	3.3(6)	< 0.1	< 0.3	
7.5 M $\Omega$	21(2)	2.6(3)	1.0(2)	0.9(1)	13.0(2)	< 0.12	< 0.11	272(20)
(1T)	1.0(2)	0.18(2)	0.14(3)	0.13(2)	1.2(2)	< 0.03	< 0.02	
4.99 M $\Omega$	29(6)	1.4(1)	0.6(1)	0.6(1)	11(2)	< 0.068	< 0.072	138(11)
(1T)	< 40	2.1(5)	< 2.8	< 1.1	21(5)	< 0.5	< 0.8	
1 k $\Omega$	0.7(1)	0.26(3)	0.23(3)	0.17(2)	0.33(17)	< 0.017	< 0.012	8.2(7)
(1T)	< 145	3(1)	< 3.9	< 4.0	< 9.4	< 0.8	< 1.5	
51 $\Omega$	1.7(3)	0.40(6)	0.53(7)	0.53(6)	< 1.14	< 0.02	< 0.034	27.4(2.7)
(1T)	< 120	< 1.5	3(2)	< 2.6	< 10	< 0.9	< 0.8	
10 nF	16(4)	119(6)	22(2)	13(1)	< 7.8	< 0.4	< 0.098	142(20)
(1T)	< 105	63(3)	26(3)	6(2)	< 13.7	< 0.4	< 1.2	
Socket	87(29)	0.7(4)	1.9(6)	1.8(3)	< 15.6	< 0.32	0.39(0.18)	< 0.006
(1T)	100(20)	0.26(9)	0.9(2)	1.5(2)	3.1(8)	< 0.08	< 0.09	
Total (no PCB)	1700(400) + < 90	640(30)	165(14)	116(7)	169(18) + < 279	< 8.1	6(3) + < 4.9	2990(150)
PCB	19.8(1.3)	1.4(0.6)	0.89(0.17)	7(3)	< 0.49	< 0.24		
Total with PCB	1463(140)	650(25)	142(12)	53(3)	287(45)	< 2.94	< 2.01	
PMT	8000(2000)	600(100)	700(200)	600(100)	1200(200)	840(90)		




Figure 6.4: Acrylic structures used to hold bases before assembling the PMT arrays. Top-left: bases in a single sector's structure, before cables were added; bottom-left: acrylic laid out in the pattern of the bottom array; right: one sector's acrylic structure after adding cables.

Table 6.2: Components used to manufacture XENONnT bases.

ID	Component	Quantity	Manufacturer	Product
1	10 M $\Omega$	3	Rohm Semiconductor	KTR03EZPF1005
2	7.5 MOhm	3	Vishay Dale	CRCW08057M50FKEA
3	4.99 MOhm	8	Vishay Dale	CRCW08054M99FKEA
4	1 kOhm	1	Yageo	RC0402FR-071KL
5	51 Ohm	5	Yageo	RC0402FR-0751RL
6	10 nF	5	Kemet	C0603C103K2RACTUSMD
7	Socket	15	TE Connectivity	2-331677-6




Figure 6.5: Bases suspended above an ultrasonic bath while cables are cleaned with ultrasound in a soap solution. The second bath, in the background, is used to rinse them with deionised water afterwards.

the soap solution and the bases held above the surface (seen in figure 6.5), with the ultrasound turned on; then with the bases and cables immersed but no ultrasound.

## 6.3 Cables

Each base has four electrical connections: the HV supply and return, signal output and a ground connection. Given the 494 PMTs, this means almost 2000 base connections are needed in the cryostat. The reliable operation of such a large number of connections requires careful planning of the cabling scheme. The need for ultra-low radioactivity is a further challenge; the large mass of cables, totalling around 60 kg and more than 16 km, must not cause too much of an increase in background for the detector.

### 6.3.1 Cabling scheme

This section provides an overview of the cabling needs of XENONnT and how they are met on a conceptual level. We distinguish between two types of cable, needed to perform different roles. Signals are transmitted using coaxial cables, where the signal being carried on a central conductor is isolated from electromagnetic noise by a conductive shield surrounding it. The high voltage power supply to the PMTs can be transmitted over cables with a single conductor, which must however have a sufficiently high voltage rating. A schematic view of

Table 6.3: Summary of total cable lengths installed in XENONnT. The numbers for cryostat refer to cables between the bases and the connector at the cryostat end of the pipe. Numbers for the new/old pipe refer to the entire length of cable between the two connectors, although some of the total length is with the cryostat and breakout chamber. All lengths are nominal; the production accuracy was  $\pm 1\%$ .

	Coaxial		High voltage	
	Number	Length [m]	Number	Length [m]
Cryostat	531	1522.3	1030	1820.9
Old pipe	432	4147.2	468	4492.8
New pipe	110	1056.0	104	998.4
Old breakout chamber	432	108.0	472	118.0
New breakout chamber	144	36.0	104	26
Air	576	3312.0	576	230.4

the cabling can be found in figure 6.6, with an explanation provided in the following. Table 6.3 summarises the total lengths of cables used in each region of the experiment.

Each PMT requires a supply voltage of around 1.5 kV. It is desirable to be able to individually control the voltage supplied to any PMT. This allows the gains to be chosen individually, enabling optimisation of signal acceptance and energy resolution. It is also essential to be able to completely turn off certain PMTs if they develop problems; during the lifetime of XENON1T several were turned off due to after-pulsing or flashing [80]. Therefore, each PMT must have its own independent power supply cable.

As discussed in section 6.2, the PMTs are operated such that the anode is connected to ground and the photocathode is supplied with a negative high voltage. This simplifies the signal readout electronics. The shielding of the coaxial signal cables should be fixed to the same source of ground as the amplifiers to minimise noise. In order to avoid creating a ground loop, the high voltage supply\* has a floating ground. This is connected to the signal cable's shielding, and therefore to the common ground, through the base. A resistor (located outside the cryostat) between the two is intended to prevent voltage fluctuations from the power supply unit from inducing noise in the signal cables.

To facilitate installation each connection is split into three sections, as can be seen in figure 6.6. The first runs from the base to a connector at the top of the cryostat. These cables are one of two lengths: 3.7 m from the bottom PMTs and 2.2 m from the top. The signal and high voltage supply cables are the same length for a given PMT. The second 9.6 m section passes through one of the two cable pipes. These pipes pass through the water tank and terminate inside the service building. There are two for historical reasons: during the construction of XENON1T enough provision was made in the pipe to operate 432 PMTs [219]. However,

\*The high voltage supply is provided by CAEN A7030LN boards for PMTs in the bottom array and CAEN A1536LN boards for those in the top array. There is no reason that each array uses a different model, the A7030 boards are simply newer.




Figure 6.6: Cabling scheme, showing connections between photomultipliers and the data acquisition system. The connections to the two PMT arrays are indicated, as are the two cable pipes (the older one on top and the newer one underneath) and the two breakout chambers. Figure not to scale.

the final XENONnT design, with an even-larger TPC than envisioned, requires 494. The extra cabling needed is provided for in a new, second pipe while the old one remains in place unchanged. The old pipe also contains much of the tubing needed for the cryogenic system, to feed liquid xenon into the cryostat and extract gaseous xenon into the purification system. Both pipes are initially vertical as they leave the cryostat, then bend and run diagonally upwards to the top floor of the service building. They are vacuum insulated; although at the top the xenon inside is at room temperature, at the bottom, where they meet the cryostat, it is close to the temperature of liquid xenon. The end of each pipe is connected to one of the two breakout chambers, where cabling and other connections are carried out of the xenon atmosphere into the surrounding air.

These breakout chambers are vacuum-tight vessels with several flanges available for vacuum feedthroughs. The third section of cables, enclosed in a *potted feedthrough*, are what carry the connections outside the xenon atmosphere. The potted feedthroughs, produced by Reliable Hermetic Seals, are produced by sealing the cables in an empty flange with a black epoxy\*. A total of eight coaxial cable feedthroughs are needed, each containing 72 cables potted in a DN 63 CF flange. All eight were purchased new for XENONnT since the old feedthroughs had a different length of cable on the air-side of the flange, and were terminated with LEMO 00 rather than SMB connectors as are used now. The Kapton wires are potted in DN 40 CF flanges, with up to 104 in each. Since the design is unchanged, the four used for XENON1T are kept and only an additional three were produced, of which one with 104 cables is used in the new pipe's breakout chamber while two with 80 cables each are added to the old pipe's chamber.

\*The exact composition of this epoxy is a trade secret. It is designed for use with vacuum on the inside, and good experience has been had using it with xenon.




Figure 6.7: Return collector. Left: photo of a return collector block with all wires soldered in place. Right: illustration of how the Kapton-insulated wires are fixed to the copper block, passing through one hole and being soldered in the second.

Outside the breakout chambers, the high voltage cables are plugged into sturdier multi-conductor cables from the high voltage supply modules; the coaxial cables are routed directly to the amplifiers. The difference in length between the internal signal cables for top and bottom PMTs is corrected for here in order to ensure that signals from both arrays have an equal travel time of 90 ns before reaching the amplifiers. Cables corresponding to top PMTs are 1.5 m longer on the air-side than those for the bottom array: 6.75 m instead of 5.25 m, of which 25 cm is on the inside of the feedthrough in each case.

Because the ground is common to all bases, it makes sense to economise on the length of cabling used by combining grounds within the cryostat. A 50 cm wire is used to carry ground from the base to a ‘return collector’. This is a copper block with two rows of 26 holes into which the 24 return wires from each PMT array sector’s bases are permanently soldered. The remaining two pairs of holes are used for shared wires connecting to the high voltage supply modules. The two wires are connected in parallel to provide redundancy, so that a single wire breaking cannot render an entire sector of 24 PMTs unusable.

### 6.3.2 Cable screening and procurement

The cables used must be operable in cryogenic conditions and compact enough to fit into the small space available, particularly around the edge of the TPC. It is also important to limit their contribution to the radioactive background in the detector. Since the mass of cables near the TPC is relatively small, yet the total surface area within the xenon environment is rather large, the most important factor is their rate of radon emanation. Whereas most other radioactive impurities will decay within the cable and their radiation products are unlikely to be detected, radon produced in the decay chain of naturally occurring  $^{238}\text{U}$  can drift inside the xenon and end up inside the TPC.

For the PMT signals, RG196A/U coaxial cables were used. These were selected for XENON1T due to their low radioactivity and ease of handling [219]. The now defunct US military standard RG196A/U specifies the materials and dimensions of the conductor, dielectric and

Table 6.4: Summary of the  $^{222}\text{Rn}$  emanation rates of various cables tested for use in XENONnT.

ID	Manufacturer	Insulation	Emanation [ $\mu\text{Bq}/\text{m}$ ]
1	Huber+Suhner <sup>1</sup>	PFA	$0.5 \pm 0.2$
2	Habia <sup>2</sup>	PTFE	$0.34 \pm 0.12$
3	Habia	PTFE	$12.6 \pm 1.8^3$
4A	Huber+Suhner	PFA	$< 2.0$
4B			$< 0.7$
4C			$1.7 \pm 0.8$
5	Pasternack	PTFE	$0.5 \pm 0.2$
6	Accu-Glass	Kapton	$0.55 \pm 0.19$
7	Accu-Glass	Kapton	$0.41 \pm 0.15$

<sup>1</sup> Measured for XENON1T, never used.

<sup>2</sup> Used in XENON1T.

<sup>3</sup> After wiping with alcohol (not standard procedure) this was reduced to  $(10.3 \pm 0.9) \mu\text{Bq}/\text{m}$ .

shielding [220]. The material used for the external insulation depends on the manufacturer; we tested and used cables with both PTFE and PFA insulation. RG196A/U cables are about 2 mm in diameter, with the conductor consisting of seven strands of 0.102 mm diameter silver-plated copper-clad steel wire. The dielectric is PTFE, making it perfect for use in low-radioactivity cryogenic environments.

Initially samples of cables produced by Habia, as used for XENON1T, were screened for radon emanation. These were found to emanate radon at more than 30 times the rate of the XENON1T sample. Although wiping with alcohol (which is not part of the usual procedure for measuring materials' emanation rates) reduced the rate by around 20%, it remained far higher than desirable. Several alternative manufacturers were identified, and radio emanation rates of cables from Pasternack and Huber+Suhner were measured. Huber+Suhner produced the total length of cable required in five batches; these were combined into three samples (4A, 4B and 4C) for radon emanation measurements. The results can be found in table 6.4. Although not as clean as the cables which were used in XENON1T, both have a much lower emanation rate than the new Habia cables. The Pasternack cable was somewhat cleaner, but the radon emanation rate of the Huber+Suhner cable was considered acceptable and this cable was chosen as the entire length had been purchased prior to screening.

Having split the cable into three batches for screening, it was possible to optimise the usage of each batch to minimise the overall radon budget in XENONnT. Inside the cryostat is the most sensitive region. Radon emanated here is close to the TPC so has a good chance of ending up inside. In contrast, xenon is continuously extracted from the two breakout chambers, creating an overall flow along both pipes away from the cryostat. The extracted xenon is distilled, so radon emanated in either of the pipes or breakout chambers is unlikely to end up inside the TPC.



The cleanest cables were therefore used in the cryostat. In order of cleanliness these are the XENON1T cables, from which only the cables for the old bottom array are long enough to be used for the new top array; the Huber+Suhner screening batch 4B; and part of batch 4A. The remainder of batch 4A and batch 4C were used for the new pipe. Table 6.5 provides more details of the usage of each batch. Habia RG196 cable was used to produce the new potted cable feedthroughs. Since each cable extends only 25 cm on the inside of the flange, these contain only 2% of the total xenon-exposed coaxial cable length. Their contribution to the cable-related radon level in the TPC will be smaller still, thanks to the efficient extraction and distillation of xenon from the breakout chambers where the feedthroughs are installed.

High voltage supplies for the PMTs are carried over 30 AWG (0.255 mm diameter) single-core gold-plated copper wires with Kapton insulation, manufactured by Accu-Glass. These cables, with an outer diameter of 0.5 mm, are rated to be used with voltages up to 2 kV, more than the 1750 V limit for XENONnT PMTs. The radon emanation of these cables is around a factor four smaller than the coaxial cables per unit length. Details can be found in table 6.4; the cable with screening ID 6 was used in the new cable pipe and screening ID 7 was used in the cryostat.

### 6.3.3 Connectors

The three sections of each cable are joined at the top of the cryostat, near the pipes' entrances, and in the two breakout chambers. Commercial connectors are used: D-subminiature pins and sockets for the Kapton wire and MMCX, or micro-miniature coaxial, connectors for the coaxial cable. In both cases, the male connector is used on the section of cable closer to the PMTs, while the female is used on the section closer to the data acquisition system.

To provide some structure, the PMTs are divided into groups of 24, which are referred to in the following as *PMT array sectors*. The division of each array into sectors is illustrated in figure 6.8. Custom-design blocks provide strain relief to avoid cables coming loose from their connectors and to ensure reliable connection between the male and female parts [219]. These blocks contain one PMT array sector's worth of cables each – 24 coaxial cables or 26 Kapton wires, including the two grounds. Dividing into 24 leaves over 13 PMTs on the top (sector 5) and one on the bottom. These are combined in terms of connectors, and the shared ground for sector 5 also supplies the single bottom PMT. Since the copper block where the grounds are combined is located near sector 5, at the top of the TPC, this is the only PMT with a longer (3.7 m) Kapton cable connecting the base and ground collector block.

The coaxial connector blocks consist of five layers of ridged PTFE, between which four layers of coaxial cables are sandwiched. Two stainless steel bolts clamp together the five pieces of PTFE. When closed, the holes through which the coaxial cables pass are too small for the connectors themselves to fit, holding them in place. Finally, two copper rods are used to firmly hold together the male and female connectors' blocks. These are threaded into the female side and attached to the male side using a small copper nut, at two opposite corners of the block.

Table 6.5: Usage of each batch of RG196A/U cable. The uses were chosen to optimise radioactive background levels in the TPC: the XENON1T cable was used where possible, and the cleanest of the new cable was prioritised for being in the cryostat, where emanated  $^{222}\text{Rn}$  is more likely to enter the TPC.

Screening ID	Manufacturer, Batch	Length [m]	Emanation [ $\mu\text{Bq/m}$ ]	Used in:	
				New pipe	Cryostat
2	Habia, XENON1T cable				120 $\times$ 2.2 m = 264 m (Top)
4A	Huber & Suhner 1003215228	1048	< 2.0	46 $\times$ 9.6 m = 441.6 m	52 $\times$ 3.7 m = 192.4 m (Bottom) 56 m <sup>1</sup> (Sensors)
4B	Huber & Suhner 1003215184	610	< 0.7		133 $\times$ 2.2 m = 292.6 m (Top) 54 $\times$ 3.7 m = 199.8 m (Bottom) 9 m <sup>1</sup> (Sensors)
	Huber & Suhner 1003215196	604			135 $\times$ 3.7 m = 499.5 m (Bottom) 9 m <sup>1</sup> (Sensors)
4C	Huber & Suhner 1003215162	459	1.7 $\pm$ 0.8	28 $\times$ 9.6 m = 268.8 m	
	Huber & Suhner 1003215259	576		36 $\times$ 9.6 m = 345.6 m	

<sup>1</sup> This length is estimated. The three batches making up screening samples 4A and 4B were used in roughly this ratio. The cables for each sensor were individually cut to length by hand depending on the position of the sensor.




Figure 6.8: Division of the two PMT arrays into sectors. Both arrays are shown as viewed from above. Adapted from figure by G. Volta.




Figure 6.9: Left: female Kapton-wire connector block (before cleaning). Center: both types of connector, with the PEEK cable ties visible. Right: coaxial connector block after installation in the breakout chamber. The arrangement of the 24 cables is visible, as are the two copper rods holding the male and female parts together.

The blocks for Kapton wire have a slightly different design; here two plates of PTFE are used; each with 26 holes, one per wire. They sandwich a small ridge in the D-subminiature pins and sockets, securing the small connectors reliably. Because of this design, the wires must be passed through one of the plates before crimping the connectors themselves, which has implications for the cleaning procedure used, as discussed below. The two plates are held together by small stainless steel bolts. Finally, a D-shaped piece of copper, attached to the wire-side of the connector block, acts as a strain relief. When XENON1T was constructed small lengths of copper wire were used to secure the wires to these D-shapes. These were found to break easily when repeatedly bent, and small loose pieces of them were later found while installing the TPC. We therefore switched to PEEK cable ties for XENONnT; these double as labels since they are engraved with the connector's unique identifier. Nevertheless, copper wire remains in place on the connectors of wires in the old pipe; removing these and replacing them would have been more risky than leaving them, both from the point of view of copper dust and damage to the connectors themselves.

As mentioned above, the cable ties used to secure Kapton-insulated wires to their strain reliefs double as labels to uniquely identify each connector block. These are laser-engraved with a code in the format pTnnG, where p, indicating the position of the connector, is either 1 (at the cryostat end of the cable pipe) or 2 (in one of the breakout chambers); T, standing for the type of cable, is either C (coaxial cable) or H (for high-voltage Kapton-insulated wire); nn is the number of the connector block from 00 to 24; and G is the connector gender, either M or F. In the case of coaxial cables, these cable ties are simply fixed around the bundle of cables corresponding to each connector block. The cables which are left over from XENON1T have an old style of labels, with only the PMT array sector number nn stamped on a copper plate, secured to the cables with copper wire. These were left in place during the upgrade and not replaced.

The connectors used to terminate those cables which are potted in feedthroughs are different, in order to interface with the data acquisition system. The coaxial cables are terminated with female SMB connectors, which can be plugged into the amplifiers. The high voltage cabling uses 52-pin connectors produced by Radiall; since each connector is used for a single PMT array sector only 26 of the 52 pins are connected.

### 6.3.4 Cleaning and installation

A detailed cleaning plan was developed in collaboration with the XENON cleanliness working group, taking into account the peculiarities of the various commercial components. Most parts for XENONnT are cleaned with a material-specific recipe; the general cleaning procedure involves first a soap bath to remove greasy surface contamination, followed by acid cleaning which effectively removes contaminants – particularly oxidation products – near the surface. While these 'standard' material-specific procedures could be used for the small copper and PTFE connector-related components, the generally commercial, composite parts needed greater care to be taken to avoid damaging them. For a short summary of the cleaning procedures used for all cable-related parts see table 6.6, or find more details in the following.

Table 6.6: Summary of cleaning procedures used for base- and cable-related parts. Concentrations of solutions used are 5% (EC-60), 1% ( $\text{H}_2\text{SO}_4$ ), 3% ( $\text{H}_2\text{O}_2$ ), 1% (citric acid), 6 molar ( $\text{HNO}_3$ ). US: in ultrasound bath. Bases are cleaned together with the cables which are soldered directly to them.

Part	Degreasing	Surface treatment
Cables	EC-60, 15 min @ 40 °C, US	–
Copper parts	Acetone, > 30 min	1. $\text{H}_2\text{SO}_4/\text{H}_2\text{O}_2$ , 10 min 2. Citric acid, 10 min
PTFE parts	EC-60, 15 min @ 40 °C, US	$\text{HNO}_3$ , 10 min
Feedthroughs	EC-60, 15 min @ 40 °C, US	–
Bases	2× EC-60, 15 min @ 40 °C: 1. No US, bases+cables 2. US, cables only	–

Loose cables (those intended for the new pipe), were cleaned using a soap solution and ultrasonic bath. While coaxial cables were cleaned with just the MMCX connectors but no PTFE connector blocks, the D-subminiature connectors on Kapton wires were already in their PTFE blocks. In the latter case, a sacrificial copper strain-relief was used during the cleaning process and for transport, and later replaced with a properly-cleaned alternative. We used Elma EC-60 soap, which is an acid-based soap intended for use in medical and dental settings, in combination with ultrasound. EC-60 is one of a few soaps which have been screened for radioactivity and considered suitable for use, provided parts are rinsed thoroughly afterwards. The cables were cleaned in soap for 15 minutes, at 40 °C and with ultrasound. This was followed by thorough rinsing and a further ultrasonic bath, also at 40 °C, in pure deionised water.

The feedthroughs were, like the loose cables, cleaned in a soap-solution ultrasound bath and rinsed with deionised water. The manufacturer, RHSeals, confirmed that it was safe to clean them using ultrasound and that this would not damage the potting material. The entire coaxial cable feedthrough assemblies, including both the internal and long external sections of cable, could be submersed for cleaning. On the other hand, we wanted to avoid getting the Radiall connectors used on the air-side of the Kapton wire feedthroughs wet. Due to their design, water which gets inside the connector may not dry easily or, worse, soap residue could end up being stuck inside. We therefore wrapped the air-side part of the assemblies in plastic, to keep any dirt in and water out, and cleaned only the internal section of the wires in soap. The ‘dirty’ air-side section was unwrapped only after closing the flange, so that it did not re-contaminate the clean part. Since it is not in contact with any xenon, this section itself does not need to be clean.

Copper is one of the most critical materials to get right when it comes to cleaning since it is easily oxidised in air. We first degreased the copper parts by immersing them in acetone

for at least half an hour. Tests showed that acetone performed better than EC-60, in that parts degreased with acetone were visibly less oxidised after the entire cleaning procedure than those degreased with soap. After degreasing, the copper was pickled for 10 minutes in a solution containing 1% sulphuric acid and 3% hydrogen peroxide. Pickling removes surface impurities, such as oxidation products and other contaminants. After pickling, passivation using a ten minute bath in 1% citric acid prevents further oxidation from damaging the part [221].

As a general rule, parts are cleaned inside the clean room where they are going to be installed, so as not to expose them to dust after cleaning. Sometimes this is not possible due to logistical constraints, in which case a couple of workarounds can be used. The first is to clean the parts in an alternative clean room and seal them in a plastic bag, or plastic wrapping, until they are needed. The alternative is to perform the main cleaning in a non-dust-free environment, and then wipe them with ethanol or deionised water to remove dust once they have been brought into the clean room for installation.

### 6.3.5 Installation of cabling

As we've seen, the central section of each cable is routed through a several-metre-long pipe, to bridge the space between the cryostat and service building. Most cables were already installed for the construction of XENON1T, and are reused now in their original pipe. To accommodate the full number of additional PMTs, an extra pipe was needed containing 110 coaxial and 104 high voltage cables, including some spares. Due to the very limited space in the water tank, we inserted the cables into the pipe before it was installed. This is the procedure which was successfully used for the old pipe [219], and the preparation of the new pipe is described here.

The cables had been cleaned using the procedure described above, but not in a clean room (although care had been taken to minimise their exposure to dust, by promptly bagging them). They were brought into the clean room bagged, and wiped once more with ethanol to remove any dust. Having laid out the cables on a long strip of plastic film, the remaining work needed for the connectors was done: the coaxial connectors were assembled fully at this point, identifying the two ends of one cable by using a multimeter. The Kapton wire connectors were already mostly assembled before cleaning; only the copper strain relief was exchanged for one which had undergone the proper copper cleaning procedure. Once all the connectors were prepared, the plastic strip was folded around the cable and heat-sealed along its length to form a long sleeve. This sleeve had a diameter of around five centimetres, small enough to fit easily through the flange at the cryostat end of the pipe.

In order to insert the cables into the pipe itself, we first pushed through a fairly stiff steel wire, as a pilot. The wire was tied to all the connector blocks individually, taking care that they were not overly 'bunched up', to make it easier to fit the collection of all connectors through the pipe. We then pulled the wire from the breakout-chamber end of the pipe, while the cables entered from the cryostat end. As the cables entered the pipe the plastic sleeve

was removed, a small section at a time (although this operation was done in a clean room, it was necessary to lay the cables on the ground and the pipe was imperfectly clean, so we attempted to keep the cables wrapped as much as possible). After inserting the cables, we tied the wire which was used to pull them through the pipe to a special plug, which tightly fits into the breakout-chamber end of the pipe. Since this is the higher end of the pipe after installation, tying the cables here prevents them from sliding down the pipe under their own weight. The pipe was installed in this form, with most of the extra cable length hanging from the end of the pipe, inside the cryostat.

The feedthroughs were added to the two breakout chambers after they had been installed together with the cable pipes in their final position. Each breakout chamber is a cylinder, with the flanges for feedthroughs in the curved wall and a large flange at each end allowing easy access. We first installed the feedthroughs themselves in the new, smaller breakout chamber. The feedthroughs had been wrapped while still inside the clean room, and were only unwrapped immediately before their installation. Once all had been installed, the large access flange was opened and cables connected. A similar process was followed for the old chamber, although in that case the old coaxial cable feedthroughs were removed first. During the entire procedure, for both chambers, nitrogen gas was flushed through in order to prevent dust from getting inside.

Once the feedthroughs had been installed, we tested the connections of all cables in the pipe and potted feedthroughs. For this testing we used a multimeter, with one person in the service building touching each connector in turn, and one person in the water tank touching the other end of that cable. A few small problems, described in the following, were found during this testing; these were fixed immediately. In a few connectors which were left over from XENON1T, the ordering didn't match on the inside and outside. This was resolved by changing the order of the connectors at the cryostat end. More seriously, the connection was completely broken in some cases. Of the signal cables, one had no connection: cable 18-9 (PMT array sector 18, connector 9). To enable the operation of this PMT we moved a connector from the incomplete block of spares to its position. Therefore the PMT at position 18-9 is connected to the cable 22-10. Two of the Kapton wire blocks also had issues, both were installed before the operation of XENON1T but not used until now. In one case, block number 13, a single connector was not working. Since PMT array sector 5 only contains 14 PMTs we swapped blocks 13 and 5 such that the broken cable is unused. In the other case, block number 17, several connectors didn't work. This was swapped with block 12, the spare.

## 6.4 PMT arrays

In chapter 5 we saw details of the photomultiplier tubes used to detect light in XENONnT. Here we will look at how they are incorporated into the TPC. We begin with a brief discussion of the structures which hold the PMT arrays in place: their design and the considerations which led to it. Then we will see the story of constructing the arrays, how they were installed as part of the TPC, and the first steps of commissioning the new detector at LNGS.




Figure 6.10: The bottom PMT array being rotated prior to assembly of the TPC. The honeycomb pattern is clearly visible. Notice that six PMTs' cutouts are incomplete. These PMTs overlap with the walls of the TPC, and the filled-in section is not facing inside the TPC.



There are two arrays of PMTs in the XENONnT TPC, holding 253 photomultipliers at the top and 241 at the bottom. Looking at each array as if from the centre of the TPC, you see a polished PTFE plate with cutouts for each PMT, arranged in a hexagonal, honeycomb-like structure as seen in figure 6.10. The PTFE is important in order to reflect light which is not incident on a photomultiplier and thus to maximise the overall collection efficiency of light. Further back, hidden by the PTFE, a copper plate provides support to the PMT array: holding up the PMTs in the top array and preventing those in the bottom from floating upwards in the liquid xenon. The PMTs themselves are fastened using circular clips, which securely hold their stem.

A particular challenge for the design of the arrays is the different thermal expansion coefficients of PTFE and copper. During cooling, the copper plate will shrink less than the PTFE: 0.2% and  $\sim 1.5\%$  compared to a room temperature of 25 °C, respectively [222, 223] (see below for more details about the contraction of PTFE). As a result, the PMTs cannot be clamped directly onto the copper, else they would move within the cutouts of the reflector during cooling. To ensure that their windows always line up perfectly with the cutouts, the movement of the PMTs must instead be defined by the shrinkage of PTFE. This is achieved by fixing the clamps, at the back of each PMT, to the reflector by means of copper rods. The effect of this is that each PMT is sandwiched between the reflector and its clamp, and thus tightly held in position. The rods pass through slots in the copper plate, allowing them to move radially inwards when the PTFE contracts. A shoulder in the copper rod on one side, and a PEEK spacer on the other, fix the vertical position of the PTFE reflector relative to the copper plate, as seen in figure 6.11, without restricting this radial movement. Only at the centre of the array is there a single rod which has only a hole and no extended slot, this is the point around which the entire array contracts or expands.

### 6.4.1 Cold test

This concept of the array design was tested before producing the full arrays, to ensure that the sliding design worked without hiccups. A sector of the full array was prepared for the tests, which were performed at the University of Zurich. The test sector, as seen in figure 6.12 was approximately one sixth of the full top PMT array, including the central PMT of the array. While suspended from an overhead crane with three chains, the sector was lowered into a large dewar, while being filmed in stop-motion from above. Inside the dewar, a shallow layer of liquid nitrogen provided cooling; the PMT array sector itself was lowered into the cold gas above this.

For the first tests, a further PTFE sheet was used on the ‘back’ side of the copper (top in figure 6.12), instead of washers. This sheet was kept thin – 2 mm – to reduce the total material in the array structure. There was a tendency for whichever side of the sector was facing down, towards the liquid nitrogen, to cool more quickly. After several attempts trying different techniques, we were able to counteract this effect fairly efficiently by spraying cold gaseous boil-off nitrogen onto the upper side. Nevertheless, the temperature difference between the two PTFE sheets reached around 20 °C during the cooling, and this resulted in




Figure 6.11: PMT clamp attachment to array structure, showing the pillars used. The slot in the copper support plate extends into/out of the page in this view. The PMT clamp itself is screwed into the top of the PEEK spacer. Figure to scale.




Figure 6.12: Final PMT array sector after performing the cooling test. The PTFE reflector can be seen at the bottom, with a Pt100 temperature sensor inserted. The black plastic parts are mockup PMTs, and their PTFE clamps can be seen at the top, which are attached to the PEEK spacers, and in turn to the copper rods which pass through the copper support plate.

visible bending of the copper pillars. Although this didn't seem to affect the performance of the design, we observed that the thin PTFE sheet warped significantly, presumably due to the different contraction rates, and was left with some plastic deformation upon warming up the sector. We supposed that this sheet was too thin to provide any useful stability, and that just the reflector and clamps were enough to ensure uniform shrinking. A final test was therefore made, with the thin sheet replaced by small PTFE washers (as were already present on the other side of the copper plate); this is the design seen in figure 6.12. With this new design the contraction and expansion were smooth when cooling down and warming up the array sector.

During some of the tests we measured the amount by which the PTFE contracted during cooling. The generally accepted contraction from 25 °C to -96 °C is around 1.5%, but varies substantially depending on the production process used and the residual stresses present in the material [223]. We observed the contraction of the PTFE reflector relative to an aluminium bar, attached near the apex of the sector. The contraction of the bar itself was accounted for by comparing its length, when cold, to a second identical but warm bar. We measured a contraction of  $(1.1 \pm 0.2)\%$ , where the uncertainty, due to the difficulty of observing the movement of the PTFE relative to the bar, was estimated by eye.

#### 6.4.2 Assembling the arrays

This section describes the procedure for assembling and testing the two PMT arrays, with details of cleaning steps being one of the most important topics. We concentrate on the PMT-related assembly steps of the two arrays, and will not go into details of the array structure's assembly itself. For logistical reasons the two array structures were ready at quite different times, and the PMTs were therefore also inserted in two separate periods: the top array was finished October 2019 and the bottom array in December 2019.

A careful protocol was followed in order to bring the photomultipliers into the cleanroom. They were first wiped briefly with ethanol, to remove excessive grease, and blown with nitrogen, to remove dust. This was done close to the door of the anteroom, and they were brought into the anteroom in batches of around six shortly after the external cleaning. In the anteroom they were cleaned again by more thoroughly wiping with ethanol, as is standard for tools and materials being brought inside. They were then transferred to the main part of the cleanroom, again in small batches.

The final cleaning was done inside the cleanroom, using a combination of 'dipping' and 'swooshing' the PMTs in an ethanol bath, and carefully wiping them. The aim was that the bath would be able to clean areas which are very difficult to wipe thoroughly, such as around the stem of the PMT and between the pins. A small gap, or slot, between the window and body, is another region that is almost impossible to wipe properly. We used analysis-grade (> 99.9%) absolute ethanol, and replaced the bath's contents after approximately every 50 photomultipliers with clean ethanol to avoid gradually accumulating contaminants. After the ethanol cleaning, compressed nitrogen was again used to remove any residual dust. This was




Figure 6.13: The bottom PMT array, seen in the cleanroom before being installed in the TPC array. From this angle the copper rods which connect the reflector plate and PMT clamps are clearly visible.



Figure 6.14: The top PMT array, hanging in the bell before the field cage was added.



Figure 6.15: PMTs being cleaned by dipping in an ethanol bath and wiping.

a last-resort measure, because while the nitrogen can effectively move dust around, doing so in a cleanroom means the dust will simply be moved onto another surface – the floor or walls if lucky, but another PMT or the array structure if unlucky. Occasionally, we noticed deposits on the PMT body after it had been cleaned. These were eventually determined to come from our gloves, when it was held while still wet with ethanol. An attempt was made to minimise the amount of contact with wet PMTs, and after the final nitrogen blowing all PMTs were inspected and cleaned again if deposits were found.

The PMT arrays were assembled gradually, inserting the PMTs sector-by-sector. This allowed us to attach the PMT bases as we progressed, and in particular to attach the bases at the interior of the array while the outer PMTs were not in place. Once an entire sector's PMTs were in place, the bases for that sector were attached, removing them from the acrylic structures and placing them in the same arrangement on the PMTs. After attaching the bases of one sector, the last parts of the connectors were added: the coaxial connector blocks were assembled and the copper parts replaced in the high voltage connector blocks.

At this point, the electric connections to each PMT were tested, simply by verifying connectivity between the relevant point on a base (anode, cathode, ground) and its corresponding connector. In a few cases a connection was broken, always because of a loose cable on the base, and these were repaired in-situ. While soldering inside the cleanroom is far from ideal, no alternatives were deemed feasible. We took precautions to try to avoid unnecessary contamination, by covering the entire array with plastic, leaving just a small hole for the base being repaired, and cleaning the cleanroom after performing the repairs.

After assembling the entire array, a further dust-removing step was followed. A length of 5 mm diameter plastic tubing was connected to a pump (outside the cleanroom) to form a



Figure 6.16: Deposits on a PMT after cleaning and drying. These were found to be released from cleanroom gloves when in contact with ethanol, and easily removed by wiping the PMT once more.

kind of fairly gentle vacuum cleaner. Using this we removed individual visible dust particles which we found on the PMT array. These were mostly copper or PEEK shavings, which we supposed to be from parts' threads, particularly internal threads. After completing the assembly we also noticed some brown residues on the cut surfaces of PMT bases, presumed to be left over from production of the bases. It is possible that they were not removed during the cleaning procedure due to the lack of ultrasound. An ethanol-soaked wipe was used to clean the surfaces of all bases, this was effective at removing the residue.

Once each entire PMT array had been completed it was tested in a light-tight box. The same table which had been used to construct the array acted as the base of the box, while a second piece covered the array. Black plastic foil taped around the bottom of the box and to the table provided additional protection against stray light. The cables were routed through four holes with  $\sim 10$  cm diameter, with plastic elbows and additional black plastic foil to mitigate light intrusion. The black box, after closing, can be seen in figure 6.18. Each PMT was ramped up to 1300 V in turn to ensure that dark count signals could be seen. During the testing one PMT was identified which did not work and the problem diagnosed to be a broken base connection. This was fixed and the PMT re-tested.

After the arrays had been installed in the TPC it was transported underground as a single unit and raised almost into its final position, below the dome of the inner cryostat. The cables were routed correctly around the exterior of the field cage, and connected on top of the bell. Trays running from the top to bottom of the field cage allowed cables from the bottom array to be neatly routed, and all cables were fastened on top of the bell using PEEK cable ties and mounts. While an attempt was made to maintain order on top of the bell, with the huge number of cables in a limited space, and additional constraints from the fragile optical fibres, inevitably it was difficult to keep the cabling tidy. It is unlikely this mess will cause serious




Figure 6.17: Complete bottom PMT array, showing the reverse (cable side).




Figure 6.18: Black box containing the top PMT array, ready to be tested, and detail of the elbow through which signal and high voltage cables are accessible.

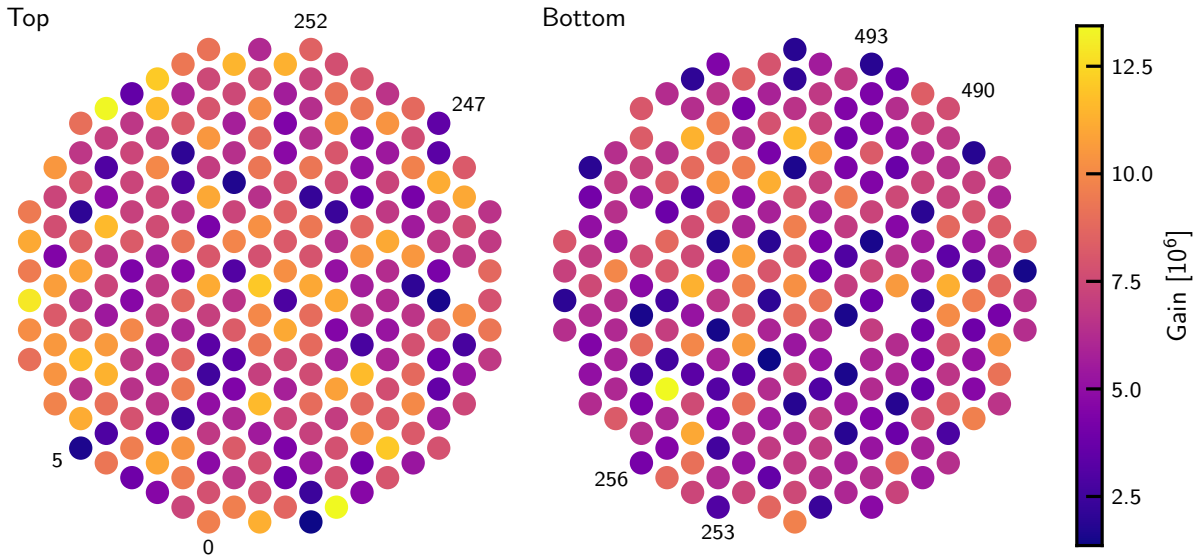


Figure 6.19: Gains of the XENONnT PMTs at 1500 V operating voltage, as measured on the 24th July 2020. The five gaps are from the five PMTs with issues, as described in the text.

problems for XENONnT, but for any future larger experiments with even more channels, more thought will likely need to be given to how to keep cabling tidy and out of the way.

## 6.5 First light in XENONnT

The inner cryostat vessel was closed on the 20th March 2020. This was followed by a period used for diagnostic measurements, including measurements of the radon emanation rate, and some time under vacuum conditions to allow outgassing of the materials inside. As of the beginning of July 2020, all but five of the PMTs are working well. Of these five, three have connection problems, one has a large afterpulse rate due to air having leaked inside during assembly, and the fifth reaches higher-than-usual base currents, maybe due to light emission. During this time, regular PMT calibration data have been taken, in particular to characterise their gains and afterpulse characteristics. An example of the results from a recent PMT gain calibration are shown in figure 6.19.

The first xenon was filled into the cryostat on the 3rd July 2020. With this xenon inside, it is possible to detect scintillation events. One event collected early on is shown in figure 6.20. Since there is no drift field in the TPC to extract electrons, only an S1 can be seen. The tail is longer than usual for an S1, due to the recombination of electrons which would have been drifted away in the presence of an electric field.

As thesis is being submitted, the XENONnT cryostat is being cooled down, ready to be filled with liquid xenon. This will happen in the coming months, while the remaining work on the new neutron veto is carried out. This will be followed by an extended period of calibration before the first science data, to be used to search for dark matter, is taken around the end of



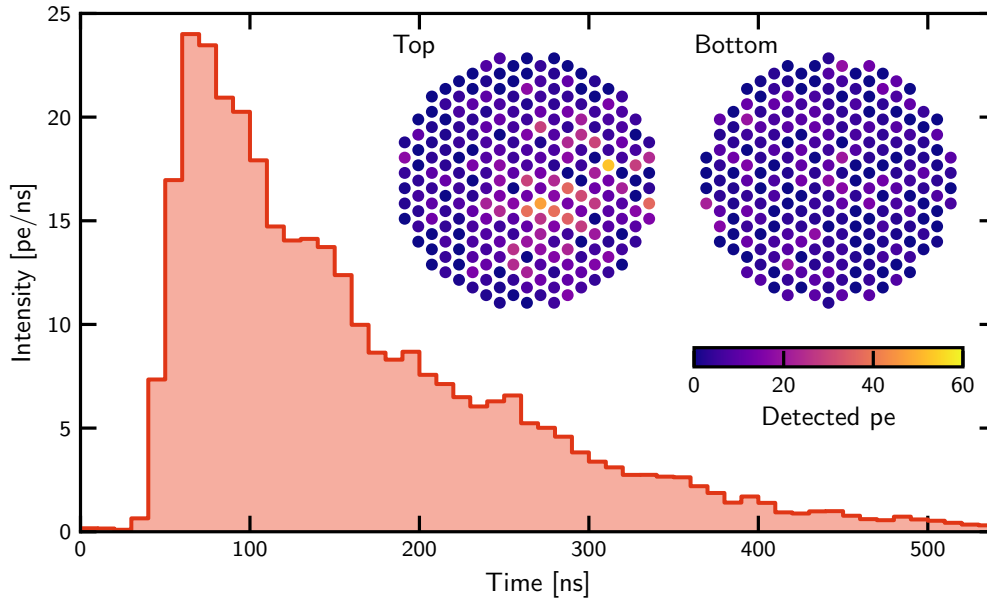


Figure 6.20: One of the first scintillation events detected in XENON1T, with gaseous xenon in the cryostat. This interaction was recorded on the 9th July 2020 at 17:55:16.

this year. The first results from XENONnT will then surely be eagerly awaited by many in the community.

# Chapter 7

## Concluding remarks

There is widespread evidence, from a large variety of astronomical and cosmological sources, for the existence of a hitherto unseen form of matter. This is known as dark matter, and is five times more abundant than baryonic matter. Evidence suggests that this dark matter consists of new particles, which are not part of the Standard Model of particle physics. One of the most popular candidates for these new particles are WIMPS, or weakly interacting massive particles. As we have seen, dual-phase xenon TPCs are an ideal technology with which to search for WIMP interactions.

The leading experiment using such technology is XENON1T, which was operating until 2018 in LNGS, Italy. Inside its TPC, 2.0 t of liquid xenon acts as a target for dark matter interactions. Light signals produced during particle interactions are detected by 248 PMTs, arranged in two arrays at the top and bottom of the TPC.

Analysing the data collected by such an experiment is a rather complicated, multi-stage process. It is vitally important to effectively characterise every particle interaction, in order to know which of these may be due to dark matter. This necessitates the use of corrections, to account for known biases in the reconstruction of events' properties. We saw a detailed example in chapter 3: the  $S2(x, y)$ -dependent correction, designed to remove the effects of a non-uniform extraction field. We must also take care to remove poorly reconstructed events from the dataset. This is the job of a range of selection criteria, which identify suspicious events. We saw the example of the  $S2$  area fraction top cut, which can identify events due to interactions in the gaseous phase, as well as a cut designed to target a particular case of misidentified  $^{83\text{m}}\text{Kr}$ -decay events.

Having applied these corrections and selection criteria, we can look for dark matter interactions in a variety of ways. In this thesis we concentrated on searches for the elastic and inelastic scattering of WIMPs. For both of these, we can to some degree distinguish between WIMP interactions and background events based on the reconstructed charge and light signals. That said, the discrimination is much better for elastic scattering interactions, and XENON1T is substantially more sensitive to these. Using data from a 1 tyr exposure of XENON1T, a world-leading upper limit of  $4.1 \times 10^{-47} \text{ cm}^2$  was placed on the cross-section of elastic, spin-independent WIMP interactions, for a mass of  $30 \text{ GeV}/c^2$ , at a 90% confidence level. Inelastic scattering, while not providing sensitivity to quite such small cross-sections,

gives us a complementary way to look for dark matter. Its detection would make it possible to distinguish between the spin-dependent and spin-independent nature of a WIMP's interaction. Furthermore, the interplay between both the elastic and inelastic channels allows us to say more about dark matter than by looking for either alone. A search for inelastic scattering was performed using a slightly smaller exposure than for elastic scattering, 0.9 t yr. No significant evidence for dark matter interactions was found, and an upper limit was placed on the cross-section reaching  $3.3 \times 10^{-39} \text{ cm}^2$  for 130 GeV/c<sup>2</sup> WIMPs, also at a 90% confidence level. This limit is also world leading above 100 GeV/c<sup>2</sup>.

One of the many advantages of dual-phase time projection chambers is their scalability. This is illustrated very nicely by the example of the XENON programme, which started with the commissioning of XENON10 in 2006, containing 14 kg of instrumented xenon. As this thesis is being written, the first data is flowing from XENONnT, the successor to XENON1T. One of the most important changes is a significant increase in the size of the TPC, now designed for 5.9 t of liquid xenon.

Because the TPC is larger, more photomultipliers are needed to fully instrument it: 494 in total. While 153 of these were reused from XENON1T, the remainder were produced new. All of these were tested thoroughly in both liquid and gas xenon, mimicking the conditions they will experience in XENONnT, at either the University of Zurich or Stockholm University. Particular attention was paid to possible leaks, which can allow secondary signals known as afterpulses to develop. If the rate of afterpulses becomes too large, PMTs start to show other problems, including light emission. In the worst cases, they must be turned off. In order to identify leaky PMTs, we searched for precisely those afterpulses which are caused by xenon contamination of the PMT's vacuum. These can be identified by their characteristic delay after the main pulse, of around 2.7  $\mu\text{s}$ . We identified leaks in 13 of the 368 PMTs tested.

A second focus of the testing campaign was on light emission by the PMTs. This can be in the form of continuous micro light emission, or can be intermittent. If signals from light emission and other sources happen to coincide in multiple PMTs, they can look like a dark matter signal. This additional background weakens the sensitivity to genuine dark matter interactions. Light emission was identified in 11 of the tested PMTs. The PMTs identified to have leaks, light emission, or other serious problems were not used for the XENONnT TPC. As a result of this thorough testing campaign, we can be confident that the PMTs will work effectively and reliably throughout the lifetime of XENONnT.

As well as the light detectors themselves, we have seen how their signals are carried away from the innards of XENON1T. Each PMT has its own base, which divides up the single voltage supplied to it into 14 levels, for its 12 dynodes, photocathode and anode. These come in the form of a small printed circuit board, which is attached directly to the PMT. The bases are optimised for their low radioactivity and are therefore constructed from components which have been screened for radioactivity. Kapton-insulated wires connect these bases to the PMTs' power supplies, while PTFE-insulated RG196 coaxial cables carry signals to be amplified and digitised. Custom connectors join together the three sections of each of these, which must be split to enable their installation. A detailed cleaning procedure, adapted to

each of the components, was developed.

Finally, we have seen how the PMT arrays were assembled. As with most components of a low-background experiment like XENONnT, much of the effort here was to ensure a high level of cleanliness. Therefore, the assembly was performed in a cleanroom and a careful protocol established for cleaning the PMTs themselves. These were inserted into the PMT arrays and their bases (with cables) were attached in the cleanroom. After assembling each of the two arrays, every PMT was tested in situ. They were then installed as part of the TPC, which was moved to its final position and connected inside the cryostat.

XENONnT is currently filled with gaseous xenon and in the process of being cooled down, ready to begin filling it with the total 8.4 t liquid xenon. The first calibration and commissioning data is already being taken.

Searching for the direct detection of dark matter is a very exciting and fast-moving field. In the coming months, both XENONnT and LUX-ZEPLIN expect to begin collecting science data, with which to search for WIMPs and other rare interactions. Even larger experiments are planned farther down the line: DARWIN, with 40 t xenon in its TPC, should be commissioned in 2026. Whether or not one of these experiments makes the discovery of dark matter that we have all been waiting for, there are sure to be many interesting challenges, solutions and results along the way.

# Bibliography

- [1] F. Zwicky, ‘Die Rotverschiebung von extragalaktischen Nebeln’, [Helv. Phys. Acta 6, \[Gen. Rel. Grav.41,207\(2009\)\]](#), 110–127 (1933).
- [2] H. Katayama and K. Hayashida, ‘X-ray study of the dark matter distribution in clusters of galaxies with Chandra’, [Adv. Space Res. 34](#), 2519 (2004).
- [3] R. Massey, T. Kitching and J. Richard, ‘The dark matter of gravitational lensing’, [Rept. Prog. Phys. 73](#), 086901 (2010).
- [4] J. Merten et al., ‘Creation of cosmic structure in the complex galaxy cluster merger Abell 2744’, [Mon. Not. Roy. Astron. Soc. 417](#), 333–347 (2011).
- [5] D. Clowe, A. Gonzalez and M. Markevitch, ‘Weak lensing mass reconstruction of the interacting cluster 1E0657-558: Direct evidence for the existence of dark matter’, [Astrophys. J. 604](#), 596–603 (2004).
- [6] A. N. Taylor, S. Dye, T. J. Broadhurst, N. Benitez and E. van Kampen, ‘Gravitational lens magnification and the mass of Abell 1689’, [Astrophys. J. 501](#), 539 (1998).
- [7] A. Mahdavi, H. Hoekstra, A. Babul, D. Balam and P. Capak, ‘A dark core in Abell 520’, [Astrophys. J. 668](#), 806–814 (2007).
- [8] M. Girardi, R. Barrena, W. Boschin and E. Ellingson, ‘Cluster Abell 520: a perspective based on member galaxies. A cluster forming at the crossing of three filaments?’, [Astron. Astrophys. 491](#), 379 (2008).
- [9] J. Kapteyn, ‘First attempt at a theory of the arrangement and motion of the sidereal system’, [Astrophys. J. 55](#), 302–328 (1922).
- [10] J. H. Oort, ‘The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems’, [B. Astron. I. Neth. 6](#), 249 (1932).
- [11] J. Holmberg and C. Flynn, ‘The local density of matter mapped by Hipparcos’, [Mon. Not. Roy. Astron. Soc. 313](#), 209–216 (2000).
- [12] V. C. Rubin and J. Ford W. Kent, ‘Rotation of the Andromeda Nebula from a spectroscopic survey of emission regions’, [Astrophys. J. 159](#), 379–403 (1970).
- [13] K. Begeman, A. Broeils and R. Sanders, ‘Extended rotation curves of spiral galaxies: dark haloes and modified dynamics’, [Mon. Not. Roy. Astron. Soc. 249](#), 523 (1991).
- [14] J. S. Bullock and M. Boylan-Kolchin, ‘Small-scale challenges to the  $\Lambda$ CDM paradigm’, [Ann. Rev. Astron. Astrophys. 55](#), 343–387 (2017).

## Bibliography

- [15] J. D. Bekenstein, ‘Relativistic gravitation theory for the MOND paradigm’, *Phys. Rev. D* **70**, [Erratum: *Phys.Rev.D* **71**, 069901 (2005)], 083509 (2004).
- [16] B. Famaey and S. McGaugh, ‘Modified Newtonian dynamics (MOND): observational phenomenology and relativistic extensions’, *Living Rev. Rel.* **15**, 10 (2012).
- [17] B. Famaey and S. McGaugh, ‘Challenges for  $\Lambda$ CDM and MOND’, *J. Phys. Conf. Ser.* **437**, edited by L. Horwitz, L. Lusanna, T. Gill, M. Land and P. Salucci, 012001 (2013).
- [18] R. Sanders, ‘Does GW170817 falsify MOND?’, *Int. J. Mod. Phys. D* **27**, 14 (2018).
- [19] P. van Dokkum et al., ‘A galaxy lacking dark matter’, *Nature* **555**, 629–632 (2018).
- [20] W. Hu and S. Dodelson, ‘Cosmic microwave background anisotropies’, *Ann. Rev. Astron. Astrophys.* **40**, 171–216 (2002).
- [21] N. Aghanim et al. (Planck), ‘Planck 2018 results. VI. Cosmological parameters’, *Astron. Astrophys.* **641**, A6 (2020).
- [22] Y. Akrami et al. (Planck), ‘Planck 2018 results. I. Overview and the cosmological legacy of Planck’, *Astron. Astrophys.* **641**, A1 (2020).
- [23] Y. Fukuda et al. (Super-Kamiokande), ‘Evidence for oscillation of atmospheric neutrinos’, *Phys. Rev. Lett.* **81**, 1562–1567 (1998).
- [24] Q. Ahmad et al. (SNO), ‘Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory’, *Phys. Rev. Lett.* **89**, 011301 (2002).
- [25] F. Abe et al. (CDF), ‘Observation of top quark production in  $\bar{p}p$  collisions’, *Phys. Rev. Lett.* **74**, 2626–2631 (1995).
- [26] K. Kodama et al. (DONUT), ‘Observation of tau neutrino interactions’, *Phys. Lett. B* **504**, 218–224 (2001).
- [27] S. Chatrchyan et al. (CMS), ‘Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC’, *Phys. Lett. B* **716**, 30–61 (2012).
- [28] G. Aad et al. (ATLAS), ‘Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC’, *Phys. Lett. B* **716**, 1–29 (2012).
- [29] R. Bouchendira, P. Clade, S. Guellati-Khelifa, F. Nez and F. Biraben, ‘New determination of the fine structure constant and test of the quantum electrodynamics’, *Phys. Rev. Lett.* **106**, 080801 (2011).
- [30] M. Tanabashi et al. (Particle Data Group), ‘Review of particle physics’, *Phys. Rev. D* **98**, 030001 (2018).
- [31] C. Baker et al., ‘Improved experimental limit on the electric dipole moment of the neutron’, *Phys. Rev. Lett.* **97**, 131801 (2006).
- [32] J. L. Feng, ‘Dark matter candidates from particle physics and methods of detection’, *Ann. Rev. Astron. Astrophys.* **48**, 495–545 (2010).

## Bibliography

- [33] G. Bertone, D. Hooper and J. Silk, ‘Particle dark matter: evidence, candidates and constraints’, *Phys. Rept.* **405**, 279–390 (2005).
- [34] G. Jungman, M. Kamionkowski and K. Griest, ‘Supersymmetric dark matter’, *Phys. Rept.* **267**, 195–373 (1996).
- [35] L. Roszkowski, E. M. Sessolo and S. Trojanowski, ‘WIMP dark matter candidates and searches—current status and future prospects’, *Rept. Prog. Phys.* **81**, 066201 (2018).
- [36] J. Silk et al., *Particle dark matter: observations, models and searches*, edited by G. Bertone (Cambridge Univ. Press, Cambridge, 2010).
- [37] G. Steigman and M. S. Turner, ‘Cosmological constraints on the properties of weakly interacting massive particles’, *Nucl. Phys. B* **253**, 375–386 (1985).
- [38] R. Peccei and H. R. Quinn, ‘CP conservation in the presence of instantons’, *Phys. Rev. Lett.* **38**, 1440–1443 (1977).
- [39] S. Weinberg, ‘A new light boson?’, *Phys. Rev. Lett.* **40**, 223–226 (1978).
- [40] F. Wilczek, ‘Problem of strong  $P$  and  $T$  invariance in the presence of instantons’, *Phys. Rev. Lett.* **40**, 279–282 (1978).
- [41] L. Visinelli and P. Gondolo, ‘Dark matter axions revisited’, *Phys. Rev. D* **80**, 035024 (2009).
- [42] J. L. Feng, A. Rajaraman and F. Takayama, ‘Superweakly interacting massive particles’, *Phys. Rev. Lett.* **91**, 011302 (2003).
- [43] L. J. Hall, K. Jedamzik, J. March-Russell and S. M. West, ‘Freeze-in production of FIMP dark matter’, *J. High Energy Phys.* **03**, 080 (2010).
- [44] J. M. Gaskins, ‘A review of indirect searches for particle dark matter’, *Contemp. Phys.* **57**, 496–525 (2016).
- [45] W. Atwood et al. (Fermi-LAT), ‘The Large Area Telescope on the Fermi Gamma-ray Space Telescope mission’, *Astrophys. J.* **697**, 1071–1102 (2009).
- [46] A. Abramowski et al. (H.E.S.S.), ‘Search for a dark matter annihilation signal from the Galactic center halo with H.E.S.S.’, *Phys. Rev. Lett.* **106**, 161301 (2011).
- [47] *High Energy Stereoscopic System*, <https://www.mpi-hd.mpg.de/hfm/HESS>.
- [48] V. Lefranc, E. Moulin, P. Panci and J. Silk, ‘Prospects for annihilating dark matter in the inner Galactic halo by the Cherenkov Telescope Array’, *Phys. Rev. D* **91**, 122003 (2015).
- [49] S. Desai et al. (Super-Kamiokande), ‘Search for dark matter WIMPs using upward through-going muons in Super-Kamiokande’, *Phys. Rev. D* **70**, [Erratum: *Phys.Rev.D* **70**, 109901 (2004)], 083523 (2004).
- [50] M. Aartsen et al. (IceCube), ‘Search for annihilating dark matter in the Sun with 3 years of IceCube data’, *Eur. Phys. J. C* **77**, [Erratum: *Eur.Phys.J.C* **79**, 214 (2019)], 146 (2017).

## Bibliography

- [51] M. G. Aartsen et al. (IceCube), ‘Search for neutrinos from dark matter self-annihilations in the center of the Milky Way with 3 years of IceCube/DeepCore’, *Eur. Phys. J. C* **77**, 627 (2017).
- [52] J. Conrad and O. Reimer, ‘Indirect dark matter searches in gamma and cosmic rays’, *Nature Phys.* **13**, 224–231 (2017).
- [53] P. A. Zyla et al. (Particle Data Group), ‘Review of particle physics’, *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020).
- [54] F. Kahlhoefer, ‘Review of LHC dark matter searches’, *Int. J. Mod. Phys. A* **32**, 1730006 (2017).
- [55] R. Essig, J. Mardon and T. Volansky, ‘Direct detection of sub-GeV dark matter’, *Phys. Rev. D* **85**, 076007 (2012).
- [56] A. M. Green, ‘Astrophysical uncertainties on the local dark matter distribution and direct detection experiments’, *J. Phys. G* **44**, 084001 (2017).
- [57] E. Aprile et al. (XENON), ‘Dark matter search results from a one ton-year exposure of XENON1T’, *Phys. Rev. Lett.* **121**, 111302 (2018).
- [58] T. Prusti et al. (Gaia), ‘The Gaia mission’, *Astron. Astrophys.* **595**, A1 (2016).
- [59] V. Belokurov, D. Erkal, N. W. Evans, S. E. Koposov and A. J. Deason, ‘Co-formation of the disc and the stellar halo’, *Mon. Not. Roy. Astron. Soc.* **478**, 611–619 (2018).
- [60] L. Necib, M. Lisanti and V. Belokurov, ‘Inferred evidence for dark matter kinematic substructure with SDSS-Gaia’, (2018).
- [61] J. Buch, S. C. Leung and J. Fan, ‘Using Gaia DR2 to constrain local dark matter density and thin dark disk’, *J. Cosmol. Astropart. Phys.* **04**, 026 (2019).
- [62] N. W. Evans, C. A. O’Hare and C. McCabe, ‘Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage’, *Phys. Rev. D* **99**, 023012 (2019).
- [63] J. Lewin and P. Smith, ‘Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil’, *Astropart. Phys.* **6**, 87–112 (1996).
- [64] T. Marrodán Undagoitia and L. Rauch, ‘Dark matter direct-detection experiments’, *J. Phys. G* **43**, 013001 (2016).
- [65] P. Klos, J. Menéndez, D. Gazit and A. Schwenk, ‘Large-scale nuclear structure calculations for spin-dependent WIMP scattering with chiral effective field theory currents’, *Phys. Rev. D* **88**, [Erratum: *Phys.Rev.D* 89, 029901 (2014)], 083516 (2013).
- [66] E. Aprile et al. (XENON), ‘Constraining the spin-dependent WIMP-nucleon cross sections with XENON1T’, *Phys. Rev. Lett.* **122**, 141301 (2019).
- [67] L. Baudis et al., ‘Signatures of dark matter scattering inelastically off nuclei’, *Phys. Rev. D* **88**, 115014 (2013).
- [68] J. R. Ellis, R. A. Flores and J. D. Lewin, ‘Rates for inelastic nuclear excitation by dark matter particles’, *Phys. Lett. B* **212**, 375–380 (1988).



## Bibliography

- [69] D. Tucker-Smith and N. Weiner, ‘Inelastic dark matter’, *Phys. Rev. D* **64**, 043502 (2001).
- [70] R. Bernabei et al. (DAMA), ‘Search for WIMP annual modulation signature: Results from DAMA / NaI-3 and DAMA / NaI-4 and the global combined analysis’, *Phys. Lett. B* **480**, 23–31 (2000).
- [71] S. Golwala et al. (CDMS), ‘Exclusion limits on the WIMP nucleon cross-section from the Cryogenic Dark Matter Search’, *Nucl. Instrum. Meth. A* **444**, edited by P. de Korte and T. Peacock, 345–349 (2000).
- [72] K. Freese, M. Lisanti and C. Savage, ‘Colloquium: Annual modulation of dark matter’, *Rev. Mod. Phys.* **85**, 1561–1581 (2013).
- [73] P.-A. Amaudruz et al. (DEAP), ‘Measurement of the scintillation time spectra and pulse-shape discrimination of low-energy  $\beta$  and nuclear recoils in liquid argon with DEAP-1’, *Astropart. Phys.* **85**, 1–23 (2016).
- [74] P. Agnes et al. (DarkSide), ‘First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso’, *Phys. Lett. B* **743**, 456–466 (2015).
- [75] K. Abe et al. (XMASS), ‘A measurement of the scintillation decay time constant of nuclear recoils in liquid xenon with the XMASS-I detector’, *J. Instrum.* **13**, P12032 (2018).
- [76] P.-A. Amaudruz et al. (DEAP-3600), ‘Design and construction of the DEAP-3600 dark matter detector’, *Astropart. Phys.* **108**, 1–23 (2019).
- [77] K. Abe et al., ‘XMASS detector’, *Nucl. Instrum. Meth. A* **716**, 78–85 (2013).
- [78] D. Akerib et al. (LUX), ‘The Large Underground Xenon (LUX) experiment’, *Nucl. Instrum. Meth. A* **704**, 111–126 (2013).
- [79] A. Tan et al. (PandaX-II), ‘Dark matter results from first 98.7 days of data from the PandaX-II experiment’, *Phys. Rev. Lett.* **117**, 121303 (2016).
- [80] E. Aprile et al. (XENON), ‘The XENON1T dark matter experiment’, *Eur. Phys. J. C* **77**, 881 (2017).
- [81] R. Agnese et al. (SuperCDMS), ‘Search for low-mass weakly interacting massive particles using voltage-assisted calorimetric ionization detection in the SuperCDMS experiment’, *Phys. Rev. Lett.* **112**, 041302 (2014).
- [82] R. Agnese et al. (SuperCDMS), ‘Low-mass dark matter search with CDMSlite’, *Phys. Rev. D* **97**, 022002 (2018).
- [83] W. Rau (CDMS, SuperCDMS), ‘CDMS and SuperCDMS’, *J. Phys. Conf. Ser.* **375**, edited by L. Oberauer, G. Raffelt and R. Wagner, 012005 (2012).
- [84] H. Jiang et al. (CDEX), ‘Limits on light weakly interacting massive particles from the first 102.8 kg  $\times$  day data of the CDEX-10 experiment’, *Phys. Rev. Lett.* **120**, 241301 (2018).
- [85] E. Armengaud et al. (EDELWEISS), ‘Performance of the EDELWEISS-III experiment for direct dark matter searches’, *J. Instrum.* **12**, P08010 (2017).

## Bibliography

- [86] A. Abdelhameed et al. (CRESST), ‘First results from the CRESST-III low-mass dark matter program’, *Phys. Rev. D* **100**, 102002 (2019).
- [87] A. Pullia, ‘Searches for dark matter with superheated liquid techniques’, *Adv. High Energy Phys.* **2014**, 387493 (2014).
- [88] C. Amole et al. (PICO), ‘Dark matter search results from the complete exposure of the PICO-60 C<sub>3</sub>F<sub>8</sub> bubble chamber’, *Phys. Rev. D* **100**, 022001 (2019).
- [89] C. Amole et al. (PICO), ‘Dark matter search results from the PICO-60 CF<sub>3</sub>I bubble chamber’, *Phys. Rev. D* **93**, 052014 (2016).
- [90] E. Behnke et al., ‘Final results of the PICASSO dark matter search experiment’, *Astropart. Phys.* **90**, 85–92 (2017).
- [91] A. Aguilar-Arevalo et al. (DAMIC), ‘Results on low-mass weakly interacting massive particles from a 11 kg-day target exposure of DAMIC at SNOLAB’, (2020).
- [92] R. Bernabei et al. (DAMA), ‘The DAMA/LIBRA apparatus’, *Nucl. Instrum. Meth. A* **592**, 297–315 (2008).
- [93] R. Bernabei et al., ‘First model independent results from DAMA/LIBRA-Phase2’, *Universe* **4**, [At. Energ.19,307(2018)], 116 (2018).
- [94] G. Adhikari et al., ‘An experiment to search for dark-matter interactions using sodium iodide detectors’, *Nature* **564**, [Erratum: *Nature* 566, E2 (2019)], 83–86 (2018).
- [95] G. Adhikari et al. (COSINE-100), ‘Search for a dark matter-induced annual modulation signal in NaI(Tl) with the COSINE-100 experiment’, *Phys. Rev. Lett.* **123**, 031302 (2019).
- [96] J. Amaré et al., ‘First results on dark matter annual modulation from the ANAIS-112 experiment’, *Phys. Rev. Lett.* **123**, 031301 (2019).
- [97] H. Zhang et al. (PandaX), ‘Dark matter direct search sensitivity of the PandaX-4T experiment’, *Sci. China Phys. Mech. Astron.* **62**, 31011 (2019).
- [98] D. Akerib et al. (LZ), ‘The LUX-ZEPLIN (LZ) experiment’, *Nucl. Instrum. Meth. A* **953**, 163047 (2020).
- [99] E. Aprile et al. (XENON), ‘Projected WIMP sensitivity of the XENONnT dark matter experiment’, *arXiv:2007.08796* (2020).
- [100] J. Aalbers et al. (DARWIN), ‘DARWIN: towards the ultimate dark matter detector’, *J. Cosmol. Astropart. Phys.* **11**, 017 (2016).
- [101] C. Aalseth et al., ‘DarkSide-20k: A 20 tonne two-phase LAr TPC for direct dark matter detection at LNGS’, *Eur. Phys. J. Plus* **133**, 131 (2018).
- [102] G. Angloher et al. (EURECA), ‘EURECA Conceptual Design Report’, *Phys. Dark Univ.* **3**, 41–74 (2014).
- [103] R. Agnese et al. (SuperCDMS), ‘Projected Sensitivity of the SuperCDMS SNOLAB experiment’, *Phys. Rev. D* **95**, 082002 (2017).

## Bibliography

- [104] P. Agnes et al. (DarkSide), ‘DarkSide-50 532-day dark matter search with low-radioactivity argon’, *Phys. Rev. D* **98**, 102006 (2018).
- [105] P. Agnes et al. (DarkSide), ‘Low-mass dark matter search with the DarkSide-50 experiment’, *Phys. Rev. Lett.* **121**, 081307 (2018).
- [106] X. Cui et al. (PandaX-II), ‘Dark matter results from 54-ton-day exposure of PandaX-II experiment’, *Phys. Rev. Lett.* **119**, 181302 (2017).
- [107] R. Ajaj et al. (DEAP), ‘Search for dark matter with a 231-day exposure of liquid argon using DEAP-3600 at SNOLAB’, *Phys. Rev. D* **100**, 022004 (2019).
- [108] D. S. Akerib et al. (LUX), ‘Results from a search for dark matter in the complete LUX exposure’, *Phys. Rev. Lett.* **118**, 021303 (2017).
- [109] R. Agnese et al. (SuperCDMS), ‘Search for low-mass weakly interacting massive particles with SuperCDMS’, *Phys. Rev. Lett.* **112**, 241302 (2014).
- [110] R. Agnese et al. (SuperCDMS), ‘New results from the search for low-mass weakly interacting massive particles with the CDMS Low Ionization Threshold Experiment’, *Phys. Rev. Lett.* **116**, 071301 (2016).
- [111] E. Aprile et al. (XENON), ‘Light dark matter search with ionization signals in XENON1T’, *Phys. Rev. Lett.* **123**, 251801 (2019).
- [112] E. Aprile et al. (XENON), ‘Search for light dark matter interactions enhanced by the Migdal effect or bremsstrahlung in XENON1T’, *Phys. Rev. Lett.* **123**, 241803 (2019).
- [113] K. Abe et al. (XMASS), ‘A direct dark matter search in XMASS-I’, *Phys. Lett. B* **789**, 45–53 (2019).
- [114] E. Aprile et al. (XENON), ‘Design and performance of the XENON10 dark matter experiment’, *Astropart. Phys.* **34**, 679–698 (2011).
- [115] E. Aprile et al. (XENON100), ‘The XENON100 dark matter experiment’, *Astropart. Phys.* **35**, 573–590 (2012).
- [116] J. Meija et al., ‘Atomic weights of the elements 2013 (IUPAC technical report)’, *Pure Appl. Chem.* **88**, 265–291 (2016).
- [117] J. R. Rumble, ed., *Handbook of chemistry and physics* (CRC, 2020).
- [118] F. Theeuwes and R. J. Bearman, ‘The  $p, V, T$  behavior of dense fluids V: the vapor pressure and saturated liquid density of xenon’, *J. Chem. Thermodyn.* **2**, 507–512 (1970).
- [119] T. Doke, A. Hitachi, S. Kubota, A. Nakamoto and T. Takahashi, ‘Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon’, *Nucl. Instrum. Meth.* **134**, 353–357 (1976).
- [120] K. Fujii et al., ‘High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region’, *Nucl. Instrum. Meth. A* **795**, 293–297 (2015).
- [121] W. Ramsay, *Nobel lecture: the rare gases of the atmosphere*, <https://www.nobelprize.org/prizes/chemistry/1904/ramsay/lecture/>, Dec. 1904.
- [122] *The Merriam-Webster new book of word histories* (Merriam-Webster, 1991).

## Bibliography

- [123] M. Berger et al., *XCOM: photon cross section database (version 1.5)*, <http://physics.nist.gov/xcom>. National Institute of Standards and Technology, Gaithersburg, MD. 2010.
- [124] M. Z. M.J. Berger J.S. Coursey and J. Chang, *Stopping-power & range tables for electrons, protons, and helium ions*, [Online] Available: <https://www.nist.gov/pml/stopping-power-range-tables-electrons-protons-and-helium-ions> [2020, July 2]. National Institute of Standards and Technology, Gaithersburg, MD. 2017.
- [125] *Gas encyclopedia* (Air Liquide, 2017).
- [126] T. Takahashi et al., ‘Average energy expended per ion pair in liquid xenon’, *Phys. Rev. A* **12**, 1771–1775 (1975).
- [127] T. Doke et al., ‘Absolute scintillation yields in liquid argon and xenon for various particles’, *Jap. J. Appl. Phys.* **41**, 1538–1545 (2002).
- [128] M. Szydagis et al., ‘NEST: a comprehensive model for scintillation yield in liquid xenon’, *J. Instrum.* **6**, P10002 (2011).
- [129] E. Aprile et al. (XENON), ‘Observation of two-neutrino double electron capture in  $^{124}\text{Xe}$  with XENON1T’, *Nature* **568**, 532–535 (2019).
- [130] J. Albert et al. (EXO-200), ‘Improved measurement of the  $2\nu\beta\beta$  half-life of  $^{136}\text{Xe}$  with the EXO-200 detector’, *Phys. Rev. C* **89**, 015502 (2014).
- [131] M. Ibe, W. Nakano, Y. Shoji and K. Suzuki, ‘Migdal effect in dark matter direct detection experiments’, *J. High Energy Phys.* **03**, 194 (2018).
- [132] S. Kubota, M. Hishida and J. Raun, ‘Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton and xenon’, *J. Phys. C* **11**, 2645–2651 (1978).
- [133] W. F. Schmidt, O. Hilt, E. Illenberger and A. G. Khrapak, ‘The mobility of positive and negative ions in liquid xenon’, *Radiation Physics and Chemistry* **74**, Professor Robert Schiller’s 70th Birthday, 152–159 (2005).
- [134] O. Hilt and W. F. Schmidt, ‘Positive hole mobility in liquid xenon’, *Chem. Phys.* **183**, 147–153 (1994).
- [135] M. Szydagis, A. Fyhrie, D. Thorngren and M. Tripathi, ‘Enhancement of NEST capabilities for simulating low-energy recoils in liquid xenon’, *J. Instrum.* **8**, C10003 (2013).
- [136] B. Lenardo et al., ‘A global analysis of light and charge yields in liquid xenon’, *IEEE Trans. Nucl. Sci.* **62**, 3387–3396 (2015).
- [137] P. Sorensen and C. E. Dahl, ‘Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter’, *Phys. Rev. D* **83**, 063501 (2011).
- [138] J. Lindhard, V. Nielsen, M. Scharff and P. V. Thomsen, ‘Integral equations governing radiation effects’, *K. Dan. Vidensk. Selsk. Mat. Fys. Medd.* **33** (1963).
- [139] E. Aprile, K. Giboni, P. Majewski, K. Ni and M. Yamashita, ‘Observation of anti-correlation between scintillation and ionization for MeV gamma-rays in liquid xenon’, *Phys. Rev. B* **76**, 014115 (2007).

## Bibliography

- [140] E. Aprile et al. (XENON), ‘XENON1T dark matter data analysis: signal and background models, and statistical inference’, *Phys. Rev. D* **99**, 112009 (2019).
- [141] E. Aprile et al. (XENON), ‘XENON1T dark matter data analysis: signal reconstruction, calibration and event selection’, *Phys. Rev. D* **100**, 052014 (2019).
- [142] L. Miller, S. Howe and W. Spear, ‘Charge transport in solid and liquid Ar, Kr, and Xe’, *Phys. Rev.* **166**, 871–878 (1968).
- [143] J. B. Albert et al. (EXO-200), ‘Measurement of the drift velocity and transverse diffusion of electrons in liquid xenon with the EXO-200 detector’, *Phys. Rev. C* **95**, 025502 (2017).
- [144] L. Baudis et al., ‘The first dual-phase xenon TPC equipped with silicon photomultipliers and characterisation with  $^{37}\text{Ar}$ ’, *Eur. Phys. J. C* **80**, 477 (2020).
- [145] S. Kobayashi et al., ‘Ratio of transverse diffusion coefficient to mobility of electrons in high-pressure xenon’, *Jpn. J. Appl. Phys.* **43**, 5568–5572 (2004).
- [146] V. M. Atrazhev et al., ‘Electron transport coefficients in liquid xenon’, in *IEEE international conference on dielectric liquids, 2005* (2005), pp. 329–332.
- [147] D. Akerib et al. (LUX), ‘Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment’, *Phys. Rev. D* **97**, 102008 (2018).
- [148] F. Mayet et al., ‘A review of the discovery reach of directional dark matter detection’, *Phys. Rept.* **627**, 1–49 (2016).
- [149] G. Mohlabeng, K. Kong, J. Li, A. Para and J. Yoo, ‘Dark matter directionality revisited with a high pressure xenon gas detector’, *J. High Energy Phys.* **07**, 092 (2015).
- [150] P. Barrow et al., ‘Qualification tests of the R11410-21 photomultiplier tubes for the XENON1T detector’, *J. Instrum.* **12**, P01024 (2017).
- [151] F. Neves et al., ‘Measurement of the absolute reflectance of polytetrafluoroethylene (PTFE) immersed in liquid xenon’, *J. Instrum.* **12**, P01017 (2017).
- [152] S. Pandolfi, *A 350-metre-tall tower to purify argon*, (2017) <https://home.cern/news/news/engineering/350-metre-tall-tower-purify-argon> (visited on 02/07/2020).
- [153] M. Simeone, *The ARIA project: production of depleted argon for the DarkSide experiment*, May 2018.
- [154] L. Baudis et al., ‘Gator: a low-background counting facility at the Gran Sasso Underground Laboratory’, *J. Instrum.* **6**, P08010 (2011).
- [155] G. Heusser, M. Laubenstein and H. Nider, ‘Low-level germanium gamma-ray spectrometry at the  $\mu\text{Bq/kg}$  level and future developments towards higher sensitivity’, in *Radionuclides in the environment*, Vol. 8, edited by P. Povinec and J. Sanchez-Cabeza, Radioactivity in the Environment (Elsevier, 2006), pp. 495–510.
- [156] M. von Sivers, B. Hofmann, Å. Rosén and M. Schumann, ‘The GeMSE facility for low-background  $\gamma$ -ray spectrometry’, *J. Instrum.* **11**, P12017 (2016).

## Bibliography

- [157] S. Nisi, L. Copia, I. Dafinei and M. Di Vacri, ‘ICP-MS measurement of natural radioactivity at LNGS’, *Int. J. Mod. Phys. A* **32**, 1743003 (2017).
- [158] E. Aprile et al. (XENON), ‘Material radioassay and selection for the XENON1T dark matter experiment’, *Eur. Phys. J. C* **77**, 890 (2017).
- [159] G. Zuzel and H. Simgen, ‘High sensitivity radon emanation measurements’, *Appl. Radiat. Isot.* **67**, 889–893 (2009).
- [160] G. Bellini et al. (Borexino), ‘Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth’, *J. Cosmol. Astropart. Phys.* **05**, 015 (2012).
- [161] K. Giboni et al., ‘Xenon recirculation-purification with a heat exchanger’, *J. Instrum.* **6**, P03002 (2011).
- [162] *Facility and method for supplying liquid xenon*, 2618038 (European Patent Office, 2013).
- [163] E. Aprile et al. (XENON), ‘Removing krypton from xenon by cryogenic distillation to the ppq level’, *Eur. Phys. J. C* **77**, 275 (2017).
- [164] C. Y. Chen et al., ‘Ultrasensitive isotope trace analyses with a magneto-optical trap’, *Science* **286**, 1139–1141 (1999).
- [165] E. Aprile et al. (XENON), ‘The XENON1T data acquisition system’, *J. Instrum.* **14**, P07016 (2019).
- [166] E. Aprile et al. (XENON1T), ‘Conceptual design and simulation of a water Cherenkov muon veto for the XENON1T experiment’, *J. Instrum.* **9**, P11006 (2014).
- [167] R. F. Lang et al., ‘A  $^{220}\text{Rn}$  source for the calibration of low-background experiments’, *J. Instrum.* **11**, P04004 (2016).
- [168] J. Marsh, D. Thomas and M. Burke, ‘High resolution measurements of neutron energy spectra from AmBe and AmB neutron sources’, *Nucl. Instrum. Meth.* **366**, 340–348 (1995).
- [169] R. F. Lang et al., ‘Characterization of a deuterium–deuterium plasma fusion neutron generator’, *Nucl. Instrum. Meth. A* **879**, 31–38 (2018).
- [170] D. Akerib et al. (LUX), ‘ $^{83\text{m}}\text{Kr}$  calibration of the 2013 LUX dark matter search’, *Phys. Rev. D* **96**, 112009 (2017).
- [171] A. Manalaysay et al., ‘Spatially uniform calibration of a liquid xenon detector at low energies using  $^{83\text{m}}\text{Kr}$ ’, *Rev. Sci. Instrum.* **81**, 073303 (2010).
- [172] L. Kastens, S. Cahn, A. Manzur and D. McKinsey, ‘Calibration of a liquid xenon detector with  $^{83\text{m}}\text{Kr}$ ’, *Phys. Rev. C* **80**, 045809 (2009).
- [173] XENON Collaboration, *The pax data processor v6.8.0*, version 6.8.0, Mar. 2018.
- [174] E. Aprile et al. (XENON), ‘Energy resolution and linearity in the keV to MeV range measured in XENON1T’, *Eur. Phys. J. C* **80**, 785 (2020).
- [175] P. Sorensen and K. Kamdin, ‘Two distinct components of the delayed single electron noise in liquid xenon emission detectors’, *J. Instrum.* **13**, P02032 (2018).

## Bibliography

- [176] E. Aprile et al. (XENON100), ‘Observation and applications of single-electron charge signals in the XENON100 experiment’, *J. Phys. G* **41**, 035201 (2014).
- [177] B. Edwards et al., ‘Measurement of single electron emission in two-phase xenon’, *Astropart. Phys.* **30**, 54–57 (2008).
- [178] J. Aalbers, ‘Dark matter search with XENON1T’, PhD thesis (University of Amsterdam, 2018).
- [179] J. Ahlswede, S. Hebel, J. O. Ross, R. Schoetter and M. B. Kalinowski, ‘Update and improvement of the global krypton-85 emission inventory’, *Journal of Environmental Radioactivity* **115**, 34–42 (2013).
- [180] D. Akimov et al. (COHERENT), ‘Observation of coherent elastic neutrino-nucleus scattering’, *Science* **357**, 1123–1126 (2017).
- [181] S. Agostinelli et al. (GEANT4), ‘GEANT4: a simulation toolkit’, *Nucl. Instrum. Meth. A* **506**, 250–303 (2003).
- [182] E. Aprile et al. (XENON), ‘Physics reach of the XENON1T dark matter experiment’, *J. Cosmol. Astropart. Phys.* **1604**, 027 (2016).
- [183] M. Leung, ‘Surface contamination from radon progeny’, *AIP Conf. Proc.* **785**, edited by B. Cleveland, R. Ford and M. Chen, 184 (2005).
- [184] M. Stein et al., ‘Radon daughter plate-out measurements at SNOLAB for polyethylene and copper’, *Nucl. Instrum. Meth. A* **880**, 92–97 (2018).
- [185] G. J. Feldman and R. D. Cousins, ‘Unified approach to the classical statistical analysis of small signals’, *Phys. Rev. D* **57**, 3873–3889 (1998).
- [186] D. S. Akerib et al. (LUX), ‘Limits on spin-dependent WIMP-nucleon cross section obtained from the complete LUX exposure’, *Phys. Rev. Lett.* **118**, 251302 (2017).
- [187] J. Xia et al. (PandaX-II), ‘PandaX-II constraints on spin-dependent WIMP-nucleon effective interactions’, *Phys. Lett. B* **792**, 193–198 (2019).
- [188] M. W. Goodman and E. Witten, ‘Detectability of certain dark matter candidates’, *Phys. Rev. D* **31**, edited by M. Srednicki, 3059 (1985).
- [189] C. McCabe, ‘Prospects for dark matter detection with inelastic transitions of xenon’, *J. Cosmol. Astropart. Phys.* **1605**, 033 (2016).
- [190] G. Arcadi, C. Dring, C. Hasterok and S. Vogl, ‘Inelastic dark matter nucleus scattering’, *J. Cosmol. Astropart. Phys.* **1912**, 053 (2019).
- [191] D. Akerib et al. (LUX), ‘Investigation of background electron emission in the LUX detector’, (2020).
- [192] A. Mozumder, ‘Free-ion yield and electron-ion recombination rate in liquid xenon’, *Chem. Phys. Lett.* **245**, 359–363 (1995).
- [193] J. Katakura, ‘Nuclear data sheets for  $A = 125$ ’, *Nucl. Data Sheets* **112**, 495–705 (2011).
- [194] M. Doi and T. Kotani, ‘Neutrino emitting modes of double beta decay’, *Prog. Theor. Phys.* **87**, 1207–1231 (1992).

## Bibliography

- [195] D. E. Cullen, *Electron photon interaction cross sections*, <https://www-nds.iaea.org/epics>.
- [196] S. S. Wilks, 'The large-sample distribution of the likelihood ratio for testing composite hypotheses', *Ann. Math. Statist.* **9**, 60–62 (1938).
- [197] R. J. Barlow and C. Beeston, 'Fitting using finite Monte Carlo samples', *Comput. Phys. Commun.* **77**, 219–228 (1993).
- [198] E. Aprile et al. (XENON), 'Search for WIMP inelastic scattering off xenon nuclei with XENON100', *Phys. Rev. D* **96**, 022008 (2017).
- [199] T. Suzuki et al. (XMASS), 'Search for WIMP- $^{129}\text{Xe}$  inelastic scattering with particle identification in XMASS-I', *Astropart. Phys.* **110**, 1–7 (2019).
- [200] H. Iams and B. Salzberg, 'The secondary emission phototube', *Proc. Ire.* **23**, 55–64 (1935).
- [201] R. G. Wagner et al., 'The next generation of photo-detectors for particle astrophysics', *USDOE Office of Science Technical Report* (2009).
- [202] L. Baudis, M. Galloway, A. Kish, C. Marentini and J. Wulf, 'Characterisation of silicon photomultipliers for liquid xenon detectors', *J. Instrum.* **13**, P10022 (2018).
- [203] L. Baudis et al., 'Performance of the Hamamatsu R11410 photomultiplier tube in cryogenic xenon environments', *J. Instrum.* **8**, P04026 (2013).
- [204] D. Akerib et al., 'An ultra-low background PMT for liquid xenon detectors', *Nucl. Instrum. Meth. A* **703**, 1–6 (2013).
- [205] E. Aprile et al. (XENON), 'Lowering the radioactivity of the photomultiplier tubes for the XENON1T dark matter experiment', *Eur. Phys. J. C* **75**, 546 (2015).
- [206] K. Lung et al., 'Characterization of the Hamamatsu R11410-10 3-inch photomultiplier tube for liquid xenon dark matter direct detection experiments', *Nucl. Instrum. Meth. A* **696**, 32–39 (2012).
- [207] K. Nakamura, Y. Hamana, Y. Ishigami and T. Matsui, 'Latest bialkali photocathode with ultra high sensitivity', *Nucl. Instrum. Meth. A* **623**, edited by H. Iwasaki, T. K. Komatsubara and Y. Sugimoto, 276–278 (2010).
- [208] B. López Paredes et al., 'Response of photomultiplier tubes to xenon scintillation light', *Astropart. Phys.* **102**, 56–66 (2018).
- [209] *Photomultiplier tubes: basics and applications*, 4th ed. (Hamamatsu Photonics, 2017).
- [210] J. Wulf, 'Direct dark matter search with XENON1T and developments for multi-ton liquid xenon detectors', PhD thesis (Universität Zürich, 2018).
- [211] D. Akimov et al., 'Observation of light emission from Hamamatsu R11410-20 photomultiplier tubes', *Nucl. Instrum. Meth. A* **794**, 1–2 (2015).
- [212] S. Li et al., 'Performance of photosensors in the PandaX-I experiment', *J. Instrum.* **11**, T02005 (2016).
- [213] D. Akimov et al., 'Peculiarities of the Hamamatsu R11410-20 photomultiplier tubes', *Proc. Sci. PhotoDet2015*, 025 (2016).



## Bibliography

- [214] *Photomultiplier tubes: principles & applications* (Photonis, Brive, France, Sept. 2002).
- [215] L. Marti et al., ‘Evaluation of gadolinium’s action on water Cherenkov detector systems with EGADS’, *Nucl. Instrum. Meth. A* **959**, 163549 (2020).
- [216] E. Brown et al., ‘Magnetically-coupled piston pump for high-purity gas applications’, *Eur. Phys. J. C* **78**, 604 (2018).
- [217] E. Aprile et al., ‘ $^{222}\text{Rn}$  emanation measurements for the XENON1T experiment’, (2020).
- [218] A. Behrens, ‘Light detectors for the XENON100 and XENON1T dark matter search experiments’, PhD thesis (Universität Zürich, 2014).
- [219] G. Kessler, ‘Inelastic WIMP-nucleus interactions in XENON100 and cables and connectors for XENON1T’, PhD thesis (University of Zurich, 2016).
- [220] *RF transmission lines and fittings*, Vol. MIL-HDBK-216, Military Standardization Handbook (Department of Defense, Jan. 1962).
- [221] C. E. Aalseth et al., ‘Ultra-low-background copper production and clean fabrication’, in *Aip conf. proc.* Vol. 785, 1 (2005), pp. 170–176.
- [222] I. I. Lifanov and N. G. Sherstyukov, ‘Thermal expansion of copper in the temperature range  $-185$  to  $+300$  °C’, *Measurement Techniques* **11**, 1653–1659 (1968).
- [223] R. K. Kirby, ‘Thermal expansion of polytetrafluoroethylene (Teflon) from  $-190$  °C to  $+300$  °C’, *Journal of Research of the National Bureau of Standards* **57**, 91–94 (1956).