

A SPARK IN THE DARK

SCINTILLATION TIME DEPENDENCE AND
NEUTRON-INDUCED SIGNALS
IN DUAL-PHASE XENON TPCS

Erik Hogenbirk

A SPARK IN THE DARK

SCINTILLATION TIME DEPENDENCE AND
NEUTRON-INDUCED SIGNALS
IN DUAL-PHASE XENON TPCS

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus

prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties
ingestelde commissie, in het openbaar te verdedigen
in de Aula der Universiteit
op vrijdag 8 maart 2019, te 13:00 uur

door

Erik Hogenbirk

geboren te Alkmaar.

Promotiecommissie:

Promotores: Prof. dr. A.P. Colijn

Universiteit van Amsterdam/

Universiteit Utrecht*

Prof. dr. F.L. Linde

Universiteit van Amsterdam

Overige Leden: Prof. dr. M.P. Decowski

Universiteit van Amsterdam

Prof. dr. ir. P.J. de Jong

Universiteit van Amsterdam

Prof. dr. E.L.M.P. Laenen

Universiteit van Amsterdam

dr. S. Ando

Universiteit van Amsterdam

Prof. dr. A.R. Pellegrino

Rijksuniversiteit Groningen

Prof. dr. M.H.M. Merk

Vrije Universiteit Amsterdam

dr. R.F. Lang

Purdue University

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

*Als professor verbonden aan de laatstgenoemde instelling (affiliated as professor with the latter institution).



Netherlands Organisation
for Scientific Research



This work is part of the research programme of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organisation for Scientific Research (NWO). The work is carried out at the National Institute for Subatomic Physics (Nikhef) in Amsterdam, The Netherlands.

Cover: Hubble Space Telescope composite image of M82. Image credit: NASA, ESA and STScI, www.hubblesite.org/image/1876/news_release/2006-14.

Copyright © 2019, Erik Hogenbirk

Distributed under Creative Commons (Attribution-NonCommercial 4.0 International (CC BY-NC 4.0))

ISBN 978-94-028-1367-8

Printed by IJskamp

PUBLICATIONS

Chapter 3: E. Hogenbirk, J. Aalbers, M. Bader, P. A. Breur, A. Brown, M. P. Decowski, C. Tunnell, R. Walet, A. P. Colijn, *Commissioning of a dual-phase xenon TPC at Nikhef*, Nuclear Instruments and Methods in Physics Research Section A, 840 (2016), pp. 87-96, reference [1].

All authors contributed to the development of the measurement setup and the TPC. RW was essential to the technical design of the TPC and its planning. Data taking was done by EH and MB, with on-call shifts done by JA, PAB, AB, MPD and APC. The data acquisition setup and analysis code were implemented with help from JA and CT. Data analysis was done by EH, partly following up on the analysis done by MB. The single-electron based gain analysis presented in section 3.5 was developed by EH. All authors were essential to the analysis and the production of the paper by their helpful discussions and criticism. All text and all images were produced by EH.

Chapter 5: E. Hogenbirk, J. Aalbers, P. A. Breur, M.P. Decowski, K. van Teutem, A. P. Colijn, *Precision measurements of the scintillation pulse shape for low-energy recoils in liquid xenon*, Journal of Instrumentation 13, Po5016, reference [2].

Data taking was done by EH and KvT, with on-call shifts done by JA, PAB, MPD and APC. The pulse shape fitting framework was developed by JA and heavily extended by EH. The basic Monte Carlo code used in section 5.5 was developed by JA with the pulse shape information added by EH. All further analysis was performed by EH. All authors were essential to the analysis and the production of the paper by their helpful discussions and criticism. All text and all images were produced by EH.

Chapter 6: E. Hogenbirk, M. P. Decowski, K. McEwan, A. P. Colijn, *Field dependence of 511 keV gamma-induced signals in a dual-phase liquid xenon TPC*, Journal of Instrumentation **13**, P10031, reference [3].

Data taking was done by EH and KME. All analysis was performed by EH under the guidance of MPD and APC. All text and all images were produced by EH.

Chapter 7: R. F. Lang, J. Piernaar, E. Hogenbirk, D. Masson, R. Nolte, A. Zimbal, S. Röttger, M. L. Benabderahmane, G. Bruno, *Characterization of a deuterium-deuterium plasma fusion neutron generator*, Nuclear Instruments and Methods in Physics Research Section A, 879 (2018), pp. 31-38, reference [4].

The project that led to this publication was led by RFL. The measurements at PTB were performed by JP, EH, MLB and GB under supervision of RN and SR. The measurements at Purdue were performed by JP and DM. The Monte Carlo simulations were done by JP. The neutron energy spectrum analysis (section 7.5) was spearheaded by EH with contributions from JP and AZ. Most of this text and all images in this section were produced by EH with corrections provided by the other authors. JP and DM performed the main part of the analysis in section 7.6. All authors contributed to the writing and revision process.

CONTENTS

Publications	v
Contents	vii
Summary	1
1 DARK MATTER	13
1.1 Introduction	13
1.2 The Λ_{CDM} model	14
1.3 Dark matter in galaxies	23
1.4 Dark matter particle detection	23
1.5 Direct detection	25
2 PARTICLE DETECTION IN XENON	33
2.1 Detection principle	33
2.2 Properties of xenon	35
2.3 Light and charge generation	36
2.4 Signal corrections and interaction reconstruction	40
2.5 Backgrounds	47
3 COMMISSIONING OF A DUAL-PHASE XENON TPC AT NIKHEF	51
3.1 Abstract	53
3.2 Introduction	54
3.3 The XAMS setup	54
3.4 Data reduction and results	58
3.5 Single-electron S ₂ -signals	66
3.6 Summary	75
3.7 Acknowledgments	76
4 CALIBRATION OF XAMS	77
4.1 Gamma-ray interactions	77
4.2 Data taking and processing	79
4.3 Data selection	82
4.4 Position calibration	88

4.5	S ₁ and S ₂ corrections	89
4.6	Residual gas analysis	97
5	PRECISION MEASUREMENTS OF THE SCINTILLATION PULSE SHAPE FOR LOW-ENERGY RECOILS IN LIQUID XENON	101
5.1	Abstract	107
5.2	Introduction	108
5.3	The process of scintillation	109
5.4	Measurements	113
5.5	Monte Carlo model	118
5.6	Results	123
5.7	Application to pulse shape discrimination	129
5.8	Conclusions	132
5.9	Acknowledgments	133
6	FIELD DEPENDENCE OF 511 KEV GAMMA-INDUCED SIGNALS IN A DUAL-PHASE LIQUID XENON TPC	135
6.1	Abstract	139
6.2	Introduction	140
6.3	Measurements	141
6.4	Results	144
6.5	Conclusions	153
6.6	Acknowledgments	155
7	CHARACTERIZATION OF A DEUTERIUM-DEUTERIUM PLASMA FUSION NEUTRON GENERATOR	157
7.1	Abstract	159
7.2	Introduction	160
7.3	Setup	160
7.4	Monte Carlo Simulation	164
7.5	Neutron Energy Spectrum	166
7.6	Neutron Flux	173
7.7	Angular Emission of Neutrons	177
7.8	Conclusions	179

8 NUCLEAR RECOIL CALIBRATION AND DARK MATTER SEARCH WITH XENON1T	181
8.1 Neutron generator installation and data acquisition	182
8.2 Data selection and band fit	183
8.3 Dark matter search results	185
9 OUTLOOK	189
Populairwetenschappelijke samenvatting	193
Acknowledgments	203
A APPENDIX	205
A.1 Residual gas analysis	205
A.2 Piping and instrumentation diagram	206
A.3 Recombination model	207
A.4 Abbreviations	209
Bibliography	211

Contents

SUMMARY

(OF: A SPARK IN THE DARK)

CHAPTER 1: EVIDENCE FOR DARK MATTER

We have never seen the most common material in the Universe. From many independent astronomical and cosmological observations, we know that 85 % of all matter must be *dark matter*, a mysterious form of matter that is not part of the Standard Model.¹ The nature of the particle responsible for dark matter remains one of the outstanding problems in modern fundamental physics.

One of the most precise measurements of the dark matter abundance comes from the observation of the cosmic microwave background (CMB) radiation. In the early Universe, density fluctuations were amplified by dark matter: it attracts gravitationally, but lacks the repulsive pressure of ordinary matter. The density fluctuations cause temperature fluctuations in the measured CMB map, which clearly bears the distinct fingerprint of the large amount of dark matter in the Universe.

There is one thing that cosmology and astronomy cannot do: tell us what *particle* dark matter consists of. Since dark matter is inferred from its gravitational interactions only, single particles cannot be observed.² The particle constraints from cosmology are therefore weak, so that many models exist that can solve the dark matter problem. One of these is the WIMP (Weakly Interacting Massive Particle) model. The WIMP has a mass of about 100 GeV and an interaction cross section of the order of the weak scale. One of its merits is that there is a very natural production mechanism in the early Universe, that predicts the right amount of dark matter that we observe. It is furthermore heavy enough to be ‘cold’, i. e. slow-moving dark matter, which is in agreement with simulations of large-scale structure formation. If WIMPs are supersymmetric particles, they could also solve the hierarchy problem in particle physics.

One of the ways WIMPs may be observed is through direct detection. WIMPs are moving through the Galaxy with typical speeds of about 200 km/s, and added to

¹ In terms of energy density, there is even more *dark energy*, something even more puzzling than dark matter, but this is a discussion for another day.

² Unless the dark matter ‘particles’ are extremely heavy, such as in the case of primordial black holes.

Summary

that comes the speed due to the motion of the Sun around the center of the Galaxy. This speed is high enough for WIMPs to produce a detectable signal if they collide with an atomic nucleus, transferring up to tens of keVs to the recoiling nucleus. The challenge of direct detection dark matter searches is to build large detectors capable of reaching this low energy threshold.

CHAPTER 2: DARK MATTER DETECTION WITH LIQUID XENON

In recent years, there has been a rapid increase in the sensitivity to low dark matter cross sections, mainly as a result of the success of dual-phase time projection chambers (TPCs) using liquid and gaseous xenon. Figure 1 schematically shows how dual-phase TPCs work. The cylindrical detector volume is partly filled with liquid xenon, cooled to -90°C , with a layer of gaseous xenon at the top. The top and bottom of the TPC are lined with sensitive light detectors, photomultiplier tubes (PMTs). When a particle interacts, it transfers (part of) its energy to a xenon nucleus or an electron, which creates a short track of excited xenon atoms (Xe^*), xenon ions (Xe^+) and electrons (e) (see panel 1a). The excited xenon atoms produce scintillation light that is detected by the PMTs as a short pulse, called S_1 . Some of the free electrons generated in the interaction are also measured. They are pulled up towards the gas layer by an electric field, generated by a high voltage applied on the cathode. At the liquid-to-gas interface, there is a stronger field from the anode, which pulls the electrons out of the liquid into the gas. The electrons gain energy in the gas and produce excited xenon atoms. This causes the emission of a second flash of light in the gas phase, which is detected by the same PMTs. This signal is called S_2 . The result from one initial interaction is therefore two measured signals, with a delay in between (see top panel in figure 1).

By combining information of the two signals, many properties of the interaction can be derived. The first is the position of the interaction. Since the velocity of the electrons in the liquid xenon is constant, the time between the S_1 and the S_2 (drift time) is directly proportional to the depth of the interaction. The other two coordinates in the plane are derived from the light distribution of the S_2 across the top PMT array. A second property is the type of the recoiling particle, which can either be a xenon nucleus (nuclear recoil, NR) or an electron (electronic recoil, ER). Most of the backgrounds, such as those caused by beta and gamma radiation, will cause ER signals, while WIMPs are expected to only give observable interactions from NRs. Compared to ERs, NRs cause more excitations and fewer ionizations, resulting in a relatively large S_1 signal and small S_2 signal. This means that the recoil

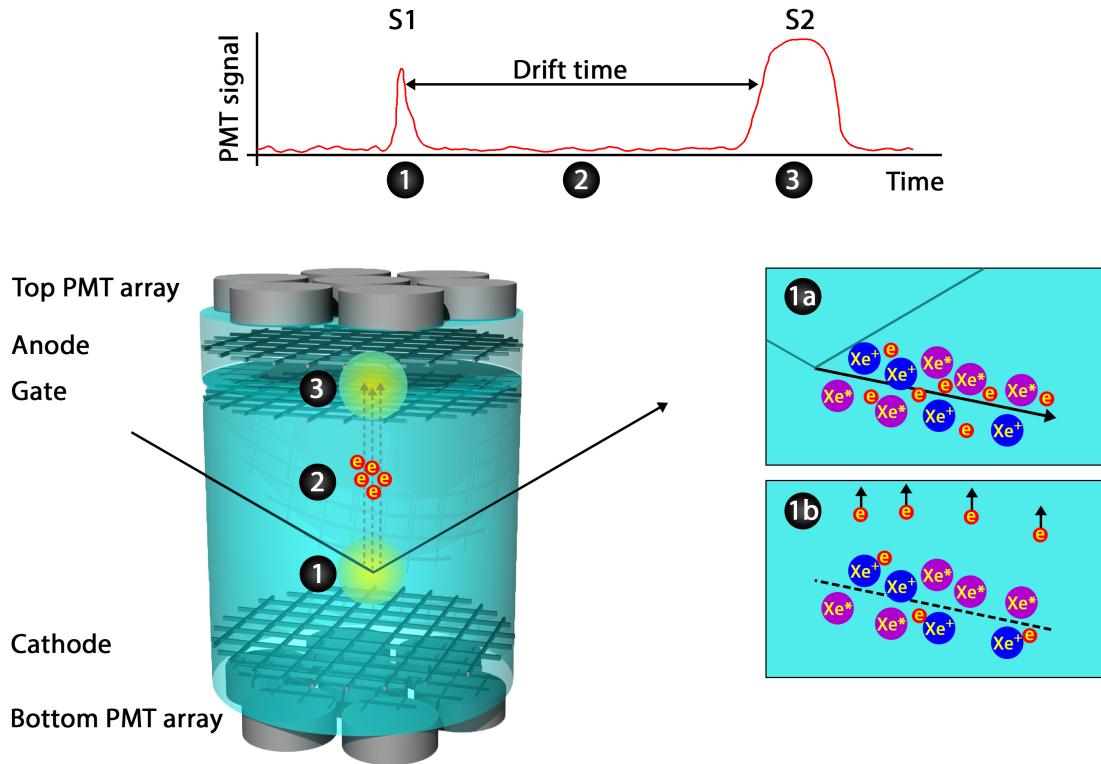


Figure 1: Schematic showing the operational principle of dual-phase TPCs. A particle interacts in the liquid xenon, which causes excitations and ionizations (1a). The excitations decay and emit scintillation light (1), detected by the PMTs as a short pulse (S_1). The electrons are pulled upwards by an electric field (1b) and travel through the liquid (2) until they arrive at the liquid-to-gas interface. A stronger field pulls them out of the liquid, causing secondary scintillation in the gas (3). This is the S_2 signal.

type can be determined experimentally from the ratio S_2/S_1 . The combination of the interaction position and the recoil type gives the excellent background rejection required for dark matter searches.

There are several corrections and calibrations required for xenon TPCs. For instance, the size of the S_1 signal is position-dependent due to varying light detection efficiency. The S_2 signal size furthermore decreases with drift time due to electrons being trapped by impurities during their drift through the liquid xenon. Both of these are corrected with mono-energetic sources.

Another calibration is the energy calibration. Both the average S_1 size and the S_2 size increase with energy, so that they can both be used as a measure for the recoil energy. However, a superior energy resolution can be achieved by using a combination of both signals, because the S_1 and S_2 signal sizes are anti-correlated.

Summary

The electrons that are not extracted from the interaction site recombine with xenon ions, and in this process generate excited xenon atoms that contribute to the S₁ signal, rather than to the S₂ signal. This process is called recombination. By using a linear combination of S₁ and S₂, fluctuations in the signal size due to fluctuations in the fraction of electrons participating in recombination are canceled out. This anti-correlation can be resolved by using multiple mono-energetic sources, or with a combination of multiple field settings.

Finally, the S₂/S₁ ratio is mapped out using ER and NR sources with varying energies. XENON1T uses an internal ²²⁰Rn source for its ER calibration, giving beta radiation going down to low energies. For the NR calibration, high-energy neutrons are used. Since they are neutral, they do not interact with the electrons but rather with the atomic nucleus, just as WIMPs are expected to do. XENON1T uses both a radiogenic neutron source (²⁴¹AmBe, through an (α , n) reaction) and a deuterium-deuterium fusion neutron generator, which is described in detail in chapters 7 and 8.

CHAPTER 3: THE XAMS SETUP

As the field of dual-phase xenon TPCs has progressed, the effort of R&D into the operational principles of these detectors has been growing too. Some of the research groups in the large collaborations that operate dark matter detectors have built their own small-scale TPCs. The purpose of these is to test new experimental techniques, research the processes leading to signal generation and to give us a more fundamental understanding of the dual-phase technology. XAMS is one of such setups in Amsterdam, featuring a TPC containing about 430 grams of liquid xenon in its active volume. The 10 cm tall, 4.5 cm diameter cylindrical volume is viewed from the top and bottom by two two-inch PMTs. Most systems function just as the systems of XENON1T, including the cooling, gas purification, data acquisition and data processing. In fact, XAMS has been the first real test case of the data acquisition and processing software of the XENON1T experiment.

Chapter 3 contains the first publication of the XAMS experiment, which marks its introduction into the scientific community as a fully functional setup. Measurements with a ²²Na gamma ray source revealed many of the operational parameters and provided essential calibration of the setup, giving the position-dependent light yield correction, the correction for electron loss of the S₂ and the energy resolution. An unusual population of events was found, where the S₂ size was up to about 80 % larger than the usual S₂s. The strong temporal correlation of these events

leads us to the conclusion that these are caused by instabilities in the liquid level, which are in turn related to a xenon gas pump in the purification system. In later measurements, this effect could be mitigated by using a different gas flow setting.

Another result from the analysis is a new PMT calibration technique. Usually, pulsed LED light sources are used to calibrate the PMT response due to single photons. The approach used in chapter 3 is different: it uses single-photon signals found in small S₂ signals. These S₂s are caused by single electrons that are liberated within the TPC due to the UV scintillation light (photo-ionization). Since they consist of only a few photons spread out over a microsecond or more, the individual single-photon pulses can be resolved, which makes it possible to measure the PMT's single-photon response. The single-photon signals are small and therefore close to the electronic noise level, which means an amplitude-dependent acceptance corrections needs to be applied for this method. Nonetheless, there are several advantages: this method requires no special hardware or calibration measurements, and directly probes the response at the xenon scintillation light wavelength (which is difficult to do with LED light through fibers). This makes it an ideal method to study effects only occurring for the UV scintillation light. An example of this is the emission of two photoelectrons as a result of one photon, which causes signals twice as high as expected.

CHAPTER 4: CALIBRATION OF XAMS

Chapter 4 improves and extends upon the initial characterization of XAMS that is presented in chapter 3. The S₁ and S₂ correction factors are recalculated (with an improved method of selecting the mono-energetic peaks of ¹³⁷Cs and ²²Na), the energy reconstruction that combines the S₁ and S₂ signal is shown and the effect of diffusion is measured. It is shown that multiple S₂s are separated for $\Delta z > 3$ mm, and merged multiple S₂s are identified for $\Delta z \gtrsim 2$ mm. The electron lifetime reached values up to 0.81 ms, and the typical light detection efficiency of the S₁ signal is approximately 10 %. Chapter 4 is a prelude to the next two chapters, since the calibration is used as the starting point of more complex analyses. In addition, this chapter gives a more detailed description of some of the analysis shown in the next chapters.

CHAPTER 5: THE SCINTILLATION PULSE SHAPE

S

One of the distinguishing features of XAMS is its excellent timing, which makes it suitable to measure the pulse shape of the S₁ signal that happens on the timescale of a few nanoseconds. One of the reasons to investigate the scintillation pulse shape is because of its importance for pulse shape discrimination (PSD). This method exploits the difference in the pulse shape between ERs and NRs to improve the separation between these events, and thus reduce the (ER) background. The effectiveness of PSD depends on the precise pulse shape of ERs and NRs. For high-energy recoils, this has been measured and the difference between ERs and NRs is large enough for effective PSD. However, measurements of the pulse shape at the low energy relevant to direct detection dark matter experiments were lacking. The publication in this chapter fills the gap in our knowledge with measurements and analysis of the pulse shape at the energy and electric field relevant to dark matter detectors.

The scintillation pulse shape depends on the way that the scintillation light is produced. There are two main ways that this occurs: through direct excitation and recombination (see figure 2). In the case of direct excitation, the scintillation light is emitted only after the formation of an excited molecular state, called an excimer. While excimer formation is very rapid, the finite lifetime of the excimer causes a significant delay in the emission of scintillation light. There are two possible excimer states, the singlet and the triplet state (corresponding to the spin state of the excimer), with approximate lifetimes of 3 ns and 22 ns, respectively. The observed pulse shapes from these states follow an exponential function and are shown in figure 3.

Scintillation can also occur through electron-ion recombination. The mechanism of excimer formation and decay is the same as in the case of direct excitation, but this is preceded by the recombination process, which causes an additional delay. Depending on the electric field, energy and particle type, the recombination process can either be very fast, or cause a significant delay. Figure 3 shows the singlet and triplet recombination pulse shape in the latter case.

The total observed S₁ pulse shape is a combination of all four previously mentioned components: the direct singlet and triplet states, and the singlet and triplet states due to recombination. For any S₁ pulse, all of these are superimposed. A model of the scintillation pulse shape should in principle include all these components.

In practice, there are some difficulties with constructing a model like this. The summed pulse shape in figure 3 is rather featureless and depends on many pa-

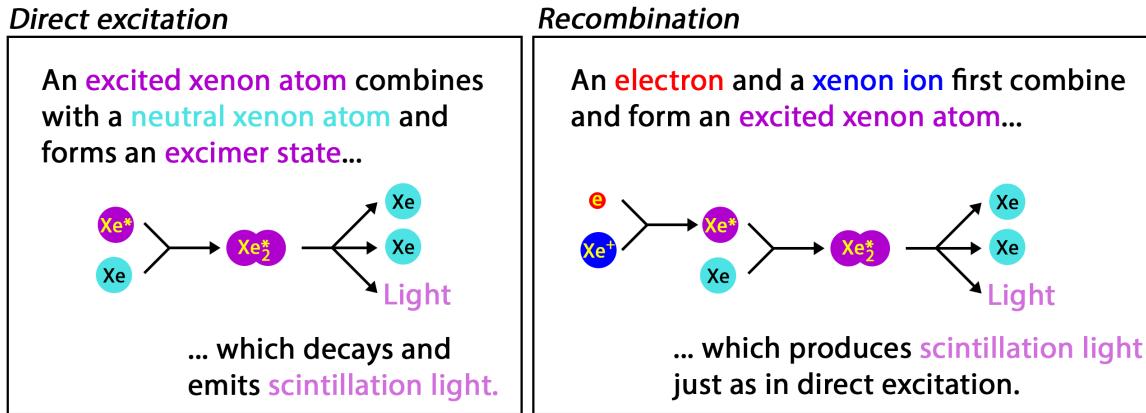


Figure 2: The processes leading to scintillation light. Direct excitation (left panel) and recombination (right panel) first produce a molecular excited state (excimer), which then decays and emits scintillation light.

rameters that are correlated when fitting a pulse shape model. In addition, the time behavior of recombination is not well known, and models are usually derived based on approximations that do not always hold. The model shown in figure 3, for instance, is based on high-energy electronic recoils at zero electric field. To counter these difficulties, an effective model is often used. In this model, the singlet and triplet times are allowed to vary to capture the slower tail of the observed pulse shape. We apply this model to data taken with XAMS, which includes ER and NR data going down to energies of a few keV, and at three different electric fields.

The results of the pulse shape measurements show a difference between ER and NR pulses that is smaller than expected, and decreases at the lowest energies. Essentially, the lower the energy gets and the higher the field gets, the smaller the difference between ER and NR pulse shapes. This makes the PSD performance worse than initially expected. Based on a pulse shape simulation, the increase in sensitivity using PSD corresponds to an effective increase in exposure of at most 6.8 % for large dark matter experiments, and only if the time resolution is improved. The minor increase in exposure likely does not justify the effort to improve the time resolution for the next generation of xenon-based dark matter detectors.

Apart from the gloomy conclusion about PSD, the measurements of chapter 5 tell us more: since the pulse shape depends on the dynamical behavior of electrons, it gives us information about the recombination process. For instance, the recombination time for low-energy ERs was previously assumed to be < 1 ns, but our results show an increase of the effective triplet time from 22 ns up to 25 ns, even when the field is relatively high. This suggests a recombination time that is at least approx-

Summary

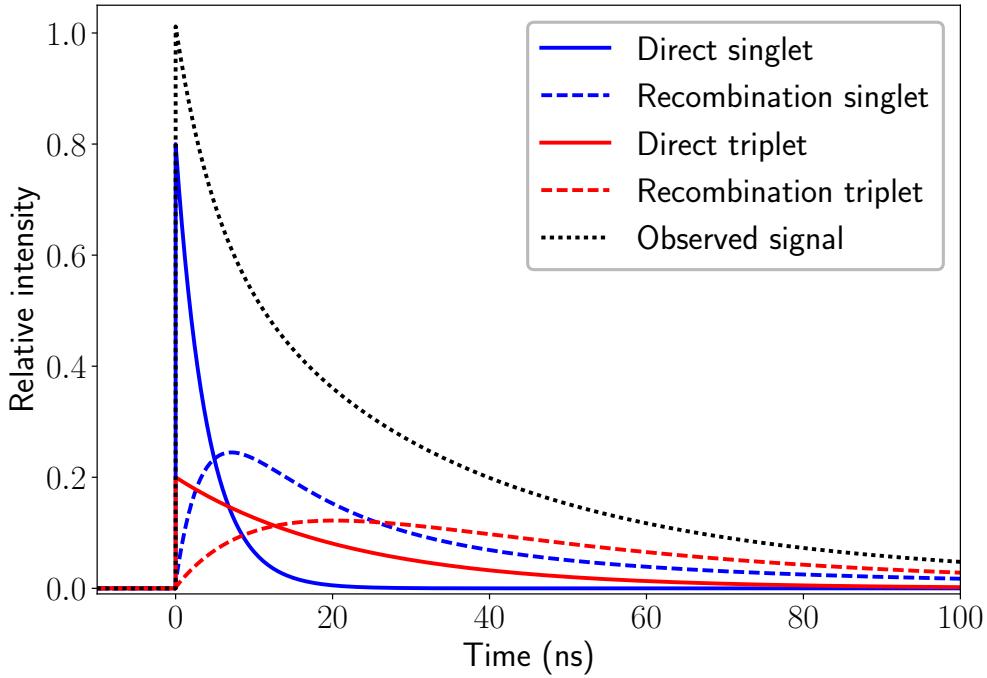


Figure 3: The measured scintillation pulse shape is built up from several components, coming from the two possible excimer states with different lifetimes (3 ns for the singlet state, 22 ns for the triplet state) and from the delaying effect of recombination.

imately 3 ns, and thus shows that the delay cannot be neglected. The dependence of pulse parameters on recoil energy further implies that there is a correlation with the linear energy transfer for both ERs and NRs. Research into the pulse shape thus opens a window into the physics happening on the microscopic distances and on the nanosecond scale.

CHAPTER 6: FIELD DEPENDENCE

In addition to pulse shape measurements, the XAMS setup has been used for measurements of the field dependence of the signals in dual-phase TPCs. As xenon dark matter detectors have become larger, the voltage V required to maintain the same electric field E rises according to $V = E \cdot L$, with L the detector drift length. Despite significant effort in high voltage engineering, the last dark matter experiments never reached their target field, but rather operated at lower fields than intended. The voltage was kept low to avoid discharges or electron emission from the cathode, which appear at higher voltage and make proper operation impossible. As a result, the electric field has decreased from approximately 500 V/cm down to

100 V/cm in the last years. If this trend persists, the field will become even lower for the next generation of TPCs.

At the current field of about 100 V/cm, we have reached a critical field strength, since many of the relevant properties of TPCs show a large rate of change for fields below this value. Because of this, small inhomogeneities in the field cause relatively large fluctuations in observed parameters. To counter this, one must do three things: design the field such that the inhomogeneity is minor, model the electric field, and correct for field dependence of the operational parameters of TPCs. The measurements and analysis of chapter 6 present a systematic study of recoils at fields ranging from 10 V/cm up to 500 V/cm and at zero field, going down to the low fields that are not commonly probed. We thus obtain a better knowledge of the low-field operation of dual-phase TPCs.

The results of this analysis show the variation of the drift velocity, electron lifetime, diffusion constant, and light and charge yields as a function of field. One of the lesser known quantities of these is the diffusion constant. This parameter indicates how much the electrons drift apart during their journey through the liquid xenon to the top of the TPC. Interestingly, at low fields, the diffusion constant rises rapidly. In combination with the lower drift velocity, this means that the electrons diffuse more due to a combination of a longer drift time (at fixed drift length) and more diffusion during this time. More diffusion means that the S₂ signals become wider, since the electrons arrive at the gas layer at different times. This might pose a problem for the rejection of events with multiple interaction sites. Usually, these events cause multiple S₂ signals, but if the S₂s become wider, they might overlap, so that a single interaction is reconstructed. Since multiple interaction sites are a clear indication of background rather than a WIMP signal, the misidentification of multiple scatters as single scatters causes an increase in background. On the other hand, the width of the S₂s is highly dependent on the local electric field in the TPC, which can be used to *reconstruct* the field given the dependence of the diffusion constant as a function of field. This method has already been applied in a preliminary study of the XENON1T field.

In addition to the aforementioned properties, the dependence of the scintillation pulse shape on field has also been determined. As the recoils measured here are 511 keV ER signals, this gives complementary information to the pulse shape measurements from chapter 5, which are performed at lower energy. The effective triplet time changes from 45 ns at low field down to 25 ns for the highest fields applied. The pulse shape in literature for high-energy (approximately 1 MeV) recoils is usually given as '45 ns decay for zero field and 27 ns with applied field'. While

the results from chapter 6 are consistent with this, it paints a more nuanced and detailed picture, showing the gradual change of pulse shape with increasing field.

CHAPTER 7: CHARACTERIZATION OF THE NEUTRON GENERATOR

As mentioned before, dark matter detectors need to be calibrated with sources of ER and NR events to enable the distinction between signal and background based on the recoil type. For XENON1T, a neutron generator is used for the NR calibration. This device allows a variable neutron flux, depending on the applied voltage and current. However, before its deployment as a calibration source for XENON1T, the device itself needs to be calibrated. The characterization of the neutron generator aims to answer three questions: how many neutrons are produced (what is the absolute flux), where are they going (what is the angular distribution), and what is their energy? All of this was determined after an extensive calibration campaign. Based on this, the paper shown in chapter 7 was published.

The angular distribution of the neutrons appears to be consistent with an isotropic distribution, if the internal geometry of the neutron generator is taken into account. The absolute flux was determined using measurements at different voltage and current settings, taking into account the internal geometry of the neutron generator and the experimental setup with a Monte Carlo simulation.

For the energy spectrum of the neutrons, measurements with a liquid scintillator detector were used. The observed energy deposition in the detector is not equal to the energy of the neutrons, because the energy deposition depends on the unknown scattering angle. This means that even for mono-energetic neutrons, a spectrum of observed recoil energies is possible. Rather than measuring the energy of the incident neutrons event-by-event, the observed recoil spectrum can be calculated from the incident neutron energy spectrum given the known detector response. In the analysis in chapter 7, there are two methods used to retrieve the neutron energy spectrum. First, the neutron energy spectrum at production in the neutron generator and at the liquid scintillator detector are calculated with a Monte Carlo simulation. Using the response function, the recoil energy spectrum is then calculated, which can be matched to the observed spectrum. Second, a method called deconvolution performs the inverse operation to convolution, so that the neutron energy spectrum at the detector can be computed. Both methods agree well. This thus gives us knowledge of the neutron energy spectrum. In the case of a plasma deuterium neutron generator, this is not equal to a single peak at the reaction neutron energy of 2.45 MeV, but is rather a spectrum with two peaks at 2.2 MeV and

2.7 MeV. This is because the fusion reaction does not occur in the same frame of reference as the lab frame, which gives an observed kinetic energy that depends on the reaction angle.

An unexpected result from the energy calibration of the neutron generator is the observation of neutrons with a much higher energy than the neutrons from deuterium-deuterium fusion. Using the Monte Carlo method described above, we reconstruct a neutron energy of roughly 14 MeV, consistent with neutrons from deuterium-tritium fusion. We attribute this to minor quantities of tritium produced in the deuterium-deuterium fusion reaction. Although the amount of tritium is small, the cross section of this fusion reaction is much higher, so that the contribution of these high-energy neutrons is around 3.5 %. This result illustrates the importance of a detailed calibration of neutron generators, since unknown effects like this can cause an overestimate of the signal yields in xenon if they remain unaccounted for.

CHAPTER 8 AND 9: NEUTRON CALIBRATION AND DARK MATTER SEARCH

After the characterization of the neutron generator described in chapter 7, it was deployed as a calibration source for XENON1T. The results of this calibration are shown in chapter 8. The calibration of the neutron generator gives the input into Monte Carlo models, which are then matched to data from XENON1T. This gives essential information that is used for the dark matter search of XENON1T. In particular, the S_2/S_1 ratio required for background rejection is mapped out with this data. The nonhomogeneous spatial distribution of events is furthermore used to validate the position reconstruction.

In the year-long science run of XENON1T, no significant excess of events was found. Unfortunately, this means that dark matter has yet to be discovered. With the null result, XENON1T tightens the constraints on dark matter, excluding parameter space in the plane of mass and cross section, as shown in figure 4. The XENON1T limits are the strongest limits to date.

Summary

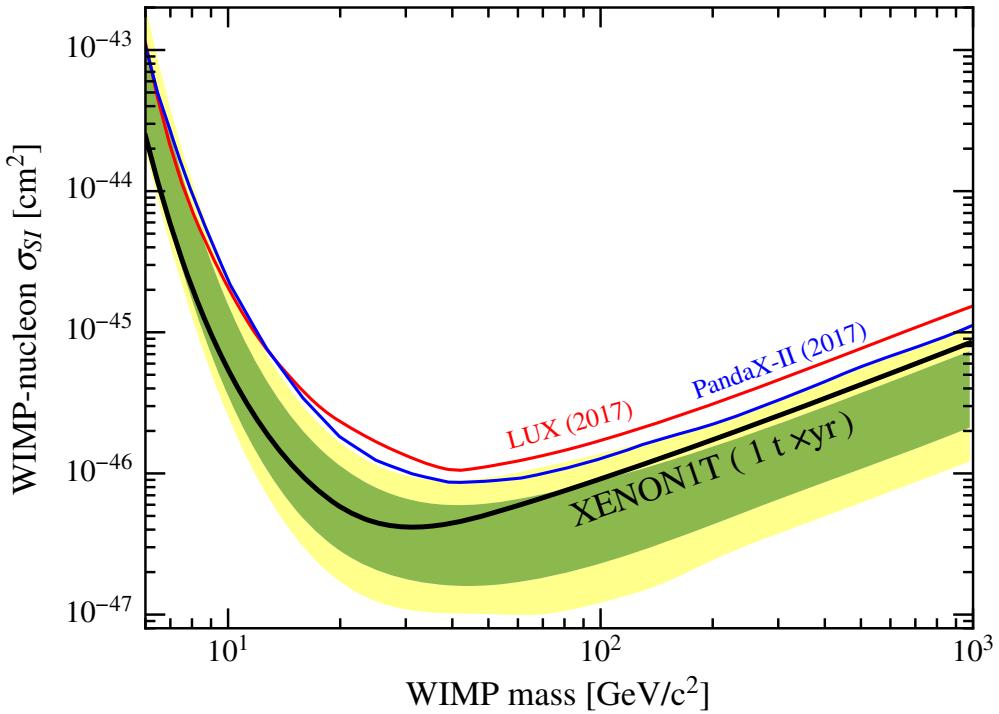


Figure 4: The obtained exclusion limit from the 1 tonne-year exposure of XENON1T (solid black line). The red and blue lines show the results from LUX and PandaX-II. The green and yellow regions are the 1σ and 2σ sensitivity bands, respectively. Adapted from [5].

The search for dark matter is not over yet; larger xenon-based experiments, such as LZ and XENONnT, are under construction and will likely start taking data in 2019. These experiments will be up to 10 times more sensitive to WIMPs than XENON1T due to the combination of a larger target mass and a lower background. It is impossible to say if these experiments will find dark matter. We just have to try hard and hope for the best. In any case, the next few years will be exciting, and I eagerly await the results from the next phase of xenon experiments, whether it is a null result or, finally, a detection.

DARK MATTER

1.1 INTRODUCTION

When it comes to cosmology, we live in an extraordinary age. Thanks to the great advances in theory and tremendous experimental efforts, which necessarily go hand in hand, our understanding of the Universe has progressed immensely. This has currently come to the level that questions that were once considered to be of a philosophical nature are creeping into the realm of fundamental physics. What is the world made of – which particles exist? Why does the Universe exist – what is the mechanism of the Big Bang? How old is time itself – what is the current age of the Universe? Amazingly, mankind begins to answer some of these questions quantitatively, if they have quantitative answers at all.

The current best model of the Big Bang is the so-called Λ_{CDM} model, which describes the evolution of the Universe under the assumption of the existence of inflation, cold (i. e. non-relativistic) dark matter and dark energy in the form of a cosmological constant Λ . This model is remarkably simple (only six parameters are needed), yet it seems to agree with several independent cosmological measurements. However, it requires adding things that we have never observed before: inflation, dark energy and dark matter. Of these three, dark matter may be the least elusive and readily observable. Moreover, the Λ_{CDM} model actually provides a good prescription for getting dark matter particles in the right abundance if dark matter particles are WIMPs: Weakly Interacting Massive Particles. These are produced in reactions in the hot plasma that existed when the Universe was less than a nanosecond old. This model also tells us that the particles must have a very small cross section, which could be the reason that they have evaded detection for so long.

In this chapter, we will first go over the Λ_{CDM} model, the reasons why it needs to introduce dark matter and the observations that support this claim. After that, we show the methods that are employed to look for individual dark matter particles. Finally, we focus on the ingredients needed for predicting the dark matter direct detection signatures and derive the expected rate and recoil energy spectrum.

1.2 THE Λ _{CDM} MODEL

According to the Big Bang model, the Universe has always been expanding due the expansion of space itself. The expansion rate is dictated by the energy content of the Universe, which is expressed by the Einstein equation as:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (1.1)$$

In this equation, $G_{\mu\nu}$ is the Einstein tensor, related to the curvature of spacetime, G is the gravitational constant and $T_{\mu\nu}$ describes the energy content of the Universe. From observations on large scales and early times (see section 1.2.1), the Universe appears to be very uniform and isotropic, meaning that at early times the expansion of space can be described by the metric:

$$g_{\mu\nu} = \text{diag}(-c, a^2(t), a^2(t), a^2(t)), \quad (1.2)$$

where $a(t)$ describes the scale of the Universe as a function of time with the current value $a_0 = 1$. The Einstein equation then becomes:

$$H^2(t) \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3c^4} \rho(t), \quad (1.3)$$

assuming a flat Universe, with $\rho(t)$ the density in the Universe as a function of time and $H(t)$ denoting the Hubble parameter with its current value H_0 . The evolution of a depends on the scaling of ρ with a (and therefore time). Based on this, the density of the Universe can be split into three parts:

$$\begin{aligned} \rho(t) &= \frac{3H_0^2 c^4}{8\pi G} \left(\Omega_{R,0} a^{-4} + \Omega_{M,0} a^{-3} + \Omega_{\Lambda,0} \right) \\ &= \rho_{cr,0} \left(\Omega_{R,0} a^{-4} + \Omega_{M,0} a^{-3} + \Omega_{\Lambda,0} \right) \end{aligned} \quad (1.4)$$

where we have defined the critical density ρ_{cr} and the subscript 0 denotes the present value. The three parts are the density of radiation (Ω_R), matter (Ω_M) and dark energy (Ω_Λ).

As it turns out, the value of Ω_M expected based on measurements that rely on properties of ordinary matter ('baryons'¹) is roughly 5 % of the critical density, which is too low compared to the overall matter density inferred from other measurements of around 30 %. The theory is saved by introducing dark matter that

¹ In a cosmological context, electrons are usually also taken as part of 'baryonic matter'.

takes up the remaining 25 %, or:

$$\Omega_M = \Omega_{DM} + \Omega_b. \quad (1.5)$$

It is worthwhile to pause for a second here and realize what we have proposed: according to this, all we have ever observed or studied amounts to one part in six of all the matter there is. This is an extraordinary claim that requires extraordinary evidence, but this extraordinary evidence is certainly there as given by the success of cosmology in describing the astrophysical observations in great detail, on scales going from the scale of dwarf galaxies all the way up to the entire Universe and matching the precision from, for instance, the measurements of the cosmic microwave background radiation. Moreover, the Big Bang theory rather naturally gives a production mechanism for a dark matter particle.

In figure 1.1, an overview of the formation of the Universe described by the Λ CDM model is given, with several ‘major events’ indicated as a function of time: inflation, freezeout of WIMPs (if they are the dark matter particle), nucleosynthesis, neutrino freezeout, the emission of the CMB and structure formation. These events happen on vastly different timescales, as indicated by the time axis in the middle. The colored bar indicates the energy component dominating the expansion of the Universe (see equation 1.4).

1.2.1 Cosmic microwave background radiation

One of the most remarkable achievements of the Λ CDM model is the detailed description of the cosmic microwave background (CMB) radiation, as recently measured precisely by the Planck satellite. This is the remnant radiation from the time when the Universe was approximately 380 000 years old. Around this time, the Universe underwent a transition from a hot plasma of ionized nuclei (mostly hydrogen) and free electrons to stable atoms, as a result of the reduced temperature in the expanding Universe. This means that the Universe became electrically neutral, so that photons could travel freely without scattering: the Universe became transparent. This is the first light that can be observed after the Big Bang.

The CMB radiation is emitted in the form of blackbody radiation, which is currently observable in the microwave ($\mathcal{O}(100\text{ GHz})$) band due to the redshift caused by the expansion of space. From the CMB, we can first of all conclude that the Universe was very homogeneous at the time of the CMB emission. Fits to the blackbody spectra show an average temperature of 2.7 K, with fluctuations only at the

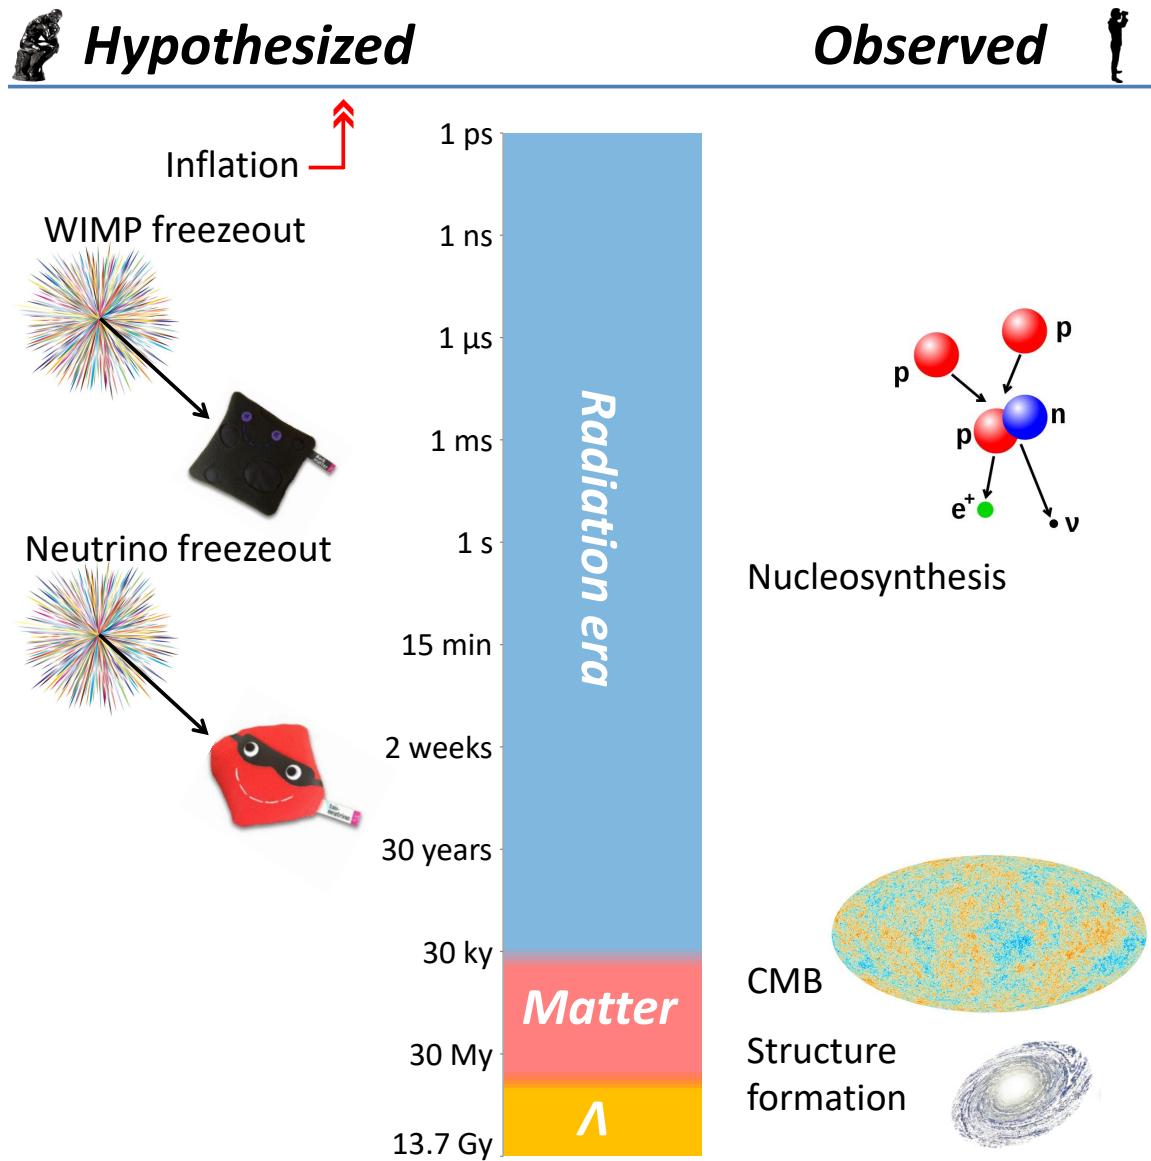


Figure 1.1: Rough timeline of the Universe according to the Λ _{CDM} model, starting at 1 ps going up to today. On the right, several (directly or indirectly) observed events are indicated at the time they occur: formation of light elements (in the first three minutes), emission of the cosmic microwave background (CMB) radiation at 380 000 years and the following formation of large-scale structure. On the left, hypothesized events are shown: freezeout of WIMPs (at times < 1 ns) and neutrinos (at times of approximately 1 s) and inflation, which occurs at $t \sim 1 \times 10^{-35}$ s, another 23 orders of magnitude before the earliest time indicated. The colored bar indicates the dominant energy density constituent that determines the expansion of the Universe, which changes due to the increasing scale factor a (see equation 1.4).

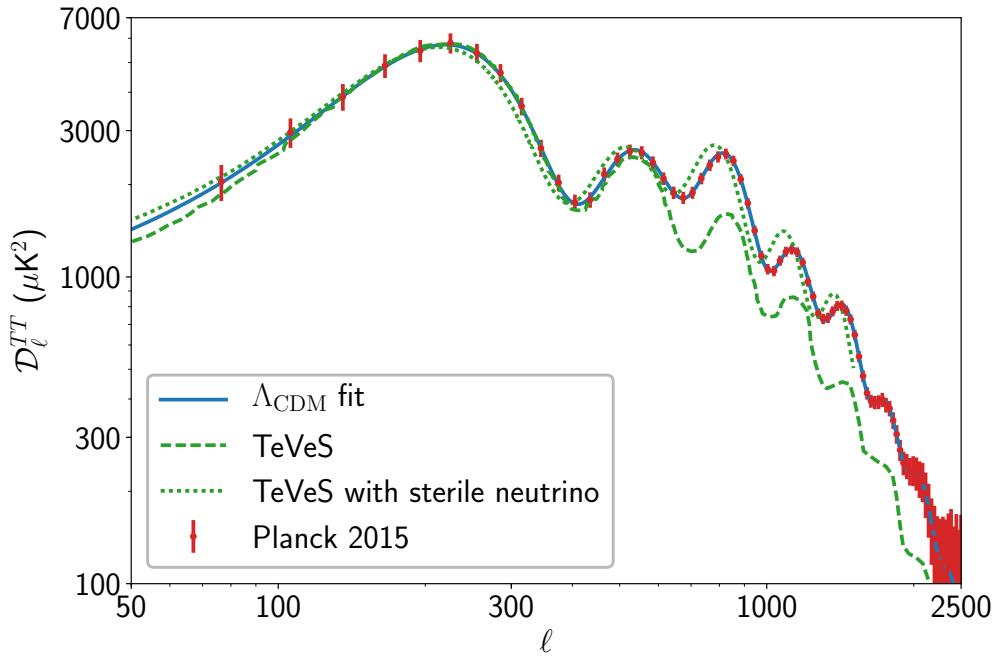


Figure 1.2: Power spectrum of the temperature anisotropies in the cosmic microwave background radiation, expressed as a function of multipole moment ℓ . The red points and error bars show the measurements from Planck and the blue line shows a fit from the Λ _{CDM} model [6]. The error bars are multiplied by five to increase visibility. The green dashed and dotted lines show the best-fit model predictions for the TeVeS model, an alternative model to dark matter that includes modified gravity [7, 8]. The model corresponding to the dotted line also includes a sterile neutrino component that acts as hot dark matter. Planck data taken from [9], source code available at [10].

10^{-5} level. This homogeneity is the main motivation for the inflation theory, since it requires that all regions of the sky must have been causally connected: there must have been some time when the Universe was small enough so that equilibrium could form. Inflation predicts a period of rapid expansion at $t \sim 1 \times 10^{-35}$ s, so that all regions in the sky start out homogeneously.

The fluctuations in the CMB temperature are small, but measurable, and there is a richness of information in the angular distribution of temperature anisotropies, since they correspond to density fluctuations. Temperature fluctuations of the CMB are usually analyzed by making a decomposition into spherical harmonics as a function of ℓ , where small values of ℓ correspond to large angular scales and vice versa. Figure 1.2 shows the measured results from the Planck satellite along with a fit based on the Λ _{CDM} model [6]. The power spectrum shows a distinct series of peaks, reflecting the angular size of pressure oscillations in the plasma at the

time of recombination (neutral atom formation). These oscillations are the result of the balance of the gravitational attractive force and the repulsive pressure, which increases when the density increases. Since dark matter is pressureless, its density changes the balance and thus the location and amplitude of the peaks [11]. This makes Planck's measurement of the CMB the most precise determination of the content of the Universe, giving (present-day) values of:

$$\begin{aligned}\Omega_\Lambda &= 0.6911 \pm 0.0062, \\ \Omega_{\text{DM}} &= 0.2589 \pm 0.0022, \\ \Omega_b &= 0.0486 \pm 0.0031.\end{aligned}\tag{1.6}$$

Note that due to the expansion of the Universe, the present-day value of Ω_R is very small ($< 1 \times 10^{-4}$). The measurement of the CMB and the excellent agreement with these measurements is one of the triumphs of the Λ_{CDM} model.

1.2.2 Structure formation

Although the early Universe is extremely uniform, small overdensities tend to grow due to gravitational attraction. This causes a clustering of the matter in the Universe that depends on its content and causes structure on the largest observed scales. When baryonic matter clusters, it exerts an outward force due to the increased pressure in the overdense region. Dark matter, on the other hand, is pressureless and amplifies overdensity fluctuations much more efficiently. In addition, structure formation depends on the *velocity* of the dark matter particles. If it is slow-moving, 'cold' dark matter, structure formation is efficient, while fast-moving, 'hot' dark matter washes out the structure. Roughly speaking, a Universe with baryonic matter only or with hot dark matter does not form significant large-scale structure, while a Universe with cold dark matter will.

The large-scale matter distribution of the Universe can be observed in several ways. Figure 1.3 shows measurements of the distribution of galaxies in a slice of the Universe, as observed in three different sky surveys (CfA2, 2dFGRS, SDSS, all shown in blue). This illustrates that there is a filament-like structure on the largest observable scales, sometimes called the 'cosmic web'. The red slices show simulated galaxy distributions, calculated from a part of the 'Millenium simulation' selected to match the general features at similar locations as observed in the surveys.

A quantitative comparison of the Λ_{CDM} cosmology to the large-scale structure observations gives a computational as well as an observational challenge. Due to

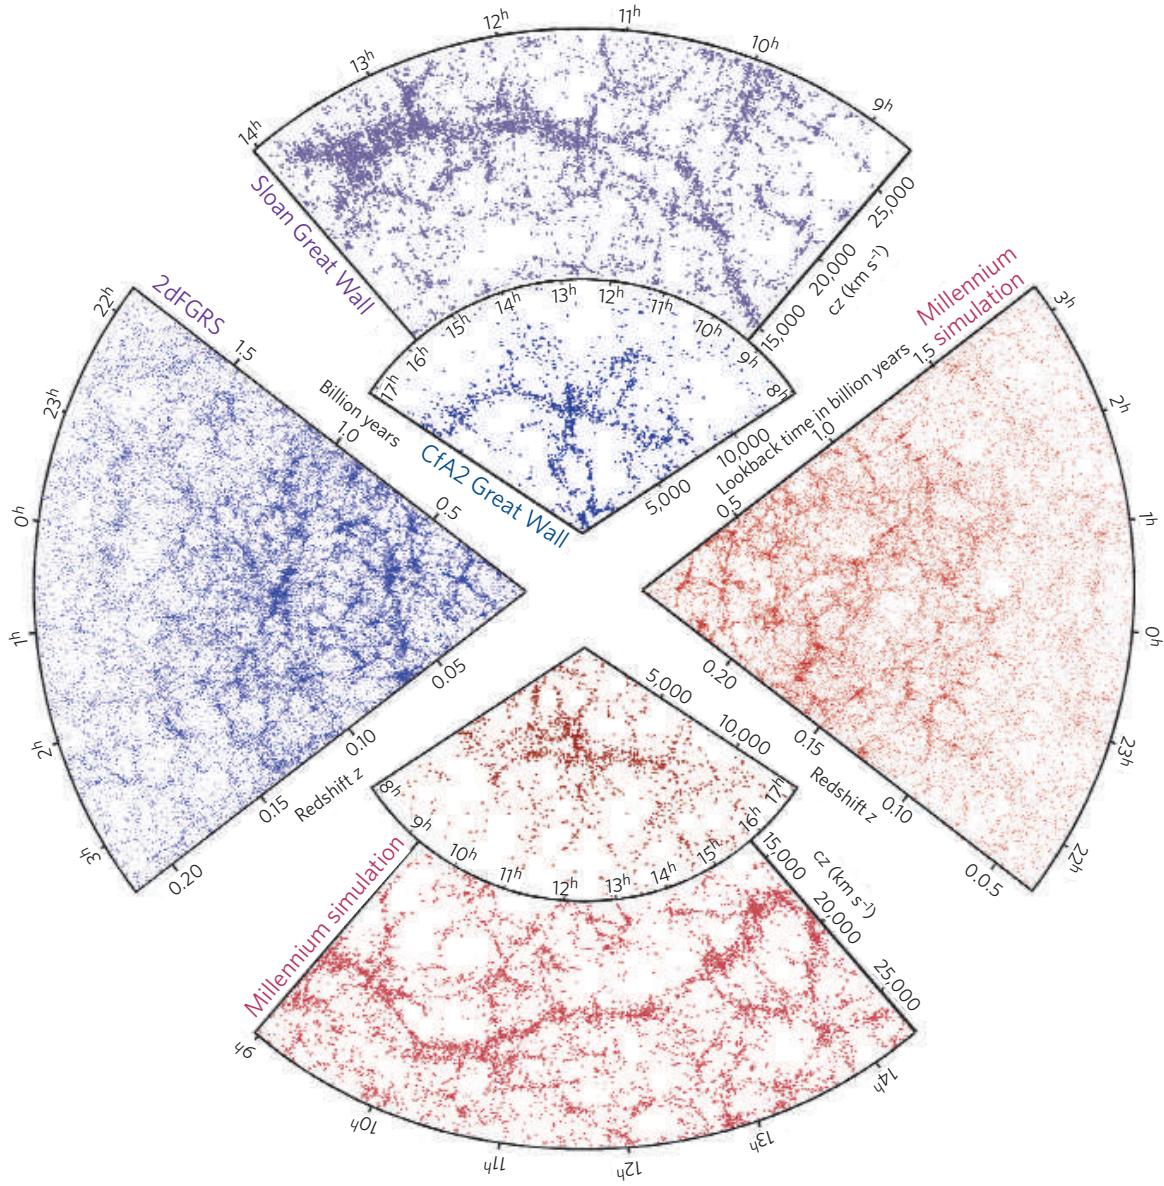


Figure 1.3: The galaxy distribution on large scales, from galaxy surveys in blue (top slice: SDSS [12], left slice: 2dFGRS [13], small top slice: CfA2 [14]) and simulated distributions (from [15]) in red on opposing sides. The observed distributions show a filament-like structure, the so-called ‘cosmic web’, on the largest scales, which are well reproduced by the Λ CDM simulations. Figure taken from [16].

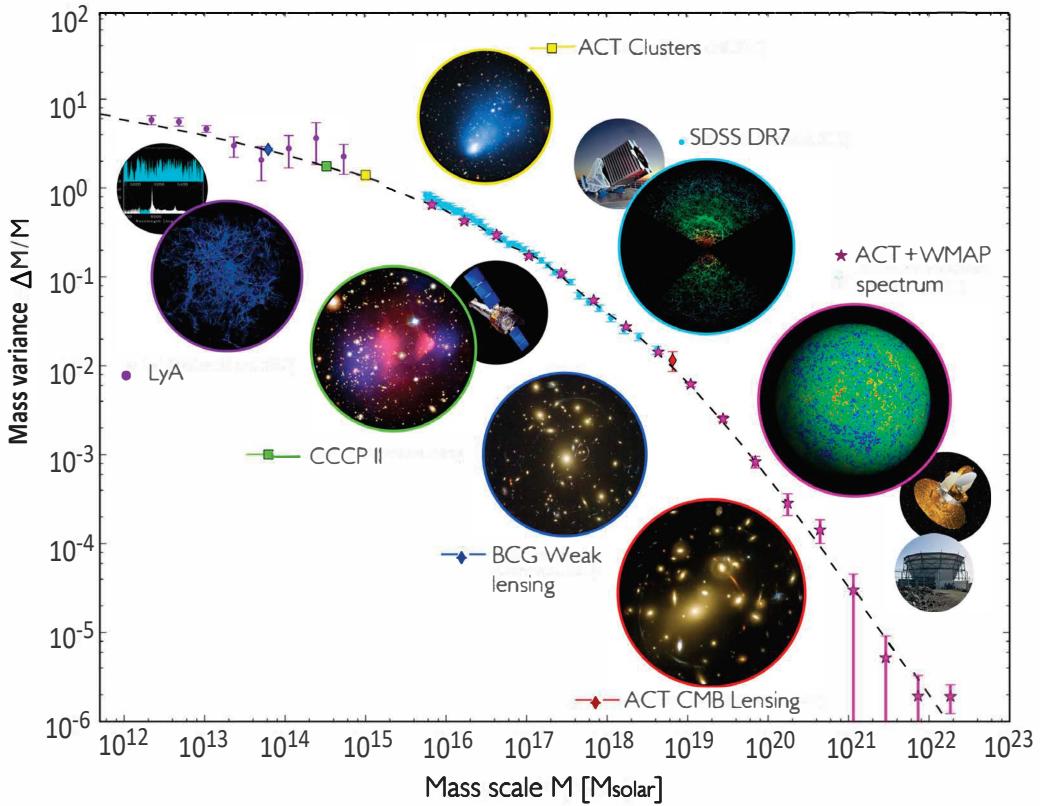


Figure 1.4: The mass variance $\Delta M/M$ as a function of mass scale, showing the amount of large-scale structure. The observations spanning the plot are from Lyman- α absorption lines (purple points [17]), weak lensing (dark blue point [18] and red point [19]), galaxy surveys (yellow point [20], green point [21] and light blue points [22]) and from the CMB power spectrum (purple stars [23, 24]). The dotted line shows the Λ CDM prediction, giving agreement over ten orders of magnitude in mass scale. Figure adapted from [25].

non-linear effects in clustering, the large-scale structure can only reliably be computed using full numerical simulations. On the observational side, the challenge is met by combining several methods to determine the matter anisotropy on scales spanning many orders of magnitude. In addition to the galaxy surveys, the Lyman- α absorption lines in distant quasars can be used to measure the hydrogen gas density as function of redshift and therefore distance. This has the advantage that it can be extended to larger scales and does not depend on the details of galaxy formation. Another method uses weak gravitational lensing: the distortion of background galaxy images due to the gravitational potential between source and observer. This has the advantage that it is sensitive to gravitational interactions and thus the total matter content, regardless if it is in the form of stars, gas or dark matter. The resulting mass variance (that is a measure for the amount of structure) from

these and several other methods is shown in figure 1.4. The dashed line, giving the Λ CDM prediction, agrees well over the full mass range and critically depends on the dark matter content and its velocity.

1.2.3 Freezeout

From the Λ CDM model, there is a clear observational motivation to introduce dark matter: it is essential to obtain agreement with various observations on scales spanning many orders of magnitude. So far, however, we have omitted any explanation about what dark matter *is*, except for that it has been around since the early Universe and that it cannot consist of relativistic particles. Surprisingly, it is relatively simple and natural to introduce particles that have a weak scale cross section and a mass $\mathcal{O}(100\text{ GeV})$.² This class of particles is called weakly interacting massive particles, or WIMPs.

The mechanism by which WIMPs are produced in the early Universe is usually called freezeout, or thermal production. When the Universe was less than a nano-second old, it consisted of a hot plasma of relativistic particles with such a high energy that the particle masses can be neglected and all creation and annihilation reactions are in equilibrium. This means that the number density n_χ of the dark matter particles χ decreases only with the expansion of the Universe, and the co-moving number density is constant (see figure 1.5). As the Universe expands, it cools and the dark matter annihilation reactions occur more frequently than their opposite reaction, resulting in a net decrease of the equilibrium number density. When the density becomes too low, the annihilation reaction becomes less likely so that the number density starts to depart from equilibrium and eventually remains constant: the density freezes out to a specific value. The remaining dark matter density in the Universe depends strongly on the point where the reactions start to fall out of equilibrium: if the cross section is higher, the reactions stay in equilibrium for a longer time and the net content of dark matter in the Universe is lower (and vice versa). This is shown by the three solid lines in figure 1.5, where three different cross sections yield three different present-day dark matter densities. The thick orange line for $\sigma = 2 \times 10^{-39} \text{ cm}^2$ shows a cross section that yields roughly the correct dark matter density. This is a cross section that is typical for weak force interactions, such as neutrino interactions, though it should be noted that the ‘weak’ of WIMPs does not necessarily imply that the weak force mediates the reaction, but

² Technically, the mass is measured in GeV/c^2 , rather than in GeV . Here and throughout we will use natural units, i. e. omit the $1/c^2$.

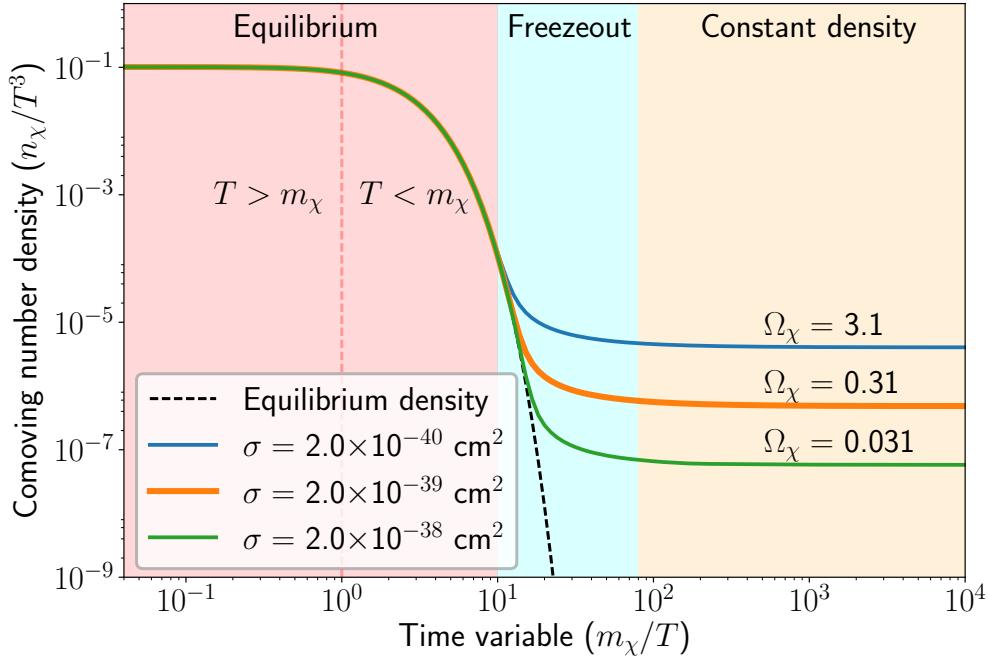


Figure 1.5: The number density n_χ of dark matter particles χ during the freezeout process, calculated following [26]. Since the process occurs during the radiation era (see figure 1.1), the scale factor a is proportional to T^{-1} so that m_χ/T can be used as a time parameter and n_χ/T^3 is the comoving number density, which compensates for the expansion of the Universe. The black dashed line shows the number density if the annihilation reactions are in equilibrium with the cosmic plasma. At early times (red region), the reaction rate is much greater than the expansion rate of the Universe, so that the density traces the equilibrium density. As the temperature drops below the WIMP mass (right to the dashed red line), the WIMP-producing reaction stops and the equilibrium density falls off. When the Universe expands far enough, the annihilation rate drops because of the decreased number density, and the dark matter number density falls out of chemical equilibrium, as indicated by the cyan region. The point where this occurs depends on the cross section of the annihilation reaction, as demonstrated by the solid lines for three different cross sections. At late times (yellow region), the number density has decreased enough to stop all reactions and the (comoving) density remains constant. The thick orange line at $\sigma = 2 \times 10^{-39} \text{ cm}^2$ gives roughly the correct dark matter density as observed today. Source code available at [10].

rather than the cross section is low.

1.3 DARK MATTER IN GALAXIES

We have seen in the previous sections that dark matter permeates the Universe in time and space, since moments after the big bang up to now and on scales ranging up to the entire Universe. The *smallest* scale on which the influence of dark matter has been observed is on the scale of galaxies. This is due to a discrepancy in the motion of stars and gas around the galactic center.

For an object at a distance R from the galactic center, the circular rotation velocity v_c is given by Kepler's law:

$$v_c(R) = \sqrt{\frac{GM(R)}{R}}, \quad (1.7)$$

with G Newton's constant and $M(R)$ the mass enclosed in a sphere of radius R from the galactic center. Since v_c is an astronomical observable, using emission lines shifted by the Doppler effect, one can indirectly measure the mass distribution in the galaxy.

Figure 1.6 shows v_c as a function of R , a *rotation curve*, for M33. Given the observed mass distribution from stars and gas, the expected rotation curve can be computed. At large distances, the observed baryonic mass density is low and equation 1.7 simplifies to $v_c \sim R^{-1/2}$, as seen from the green dash-dotted line that traces the gray background lines for $R \gtrsim 8$ kpc. However, the observed rotation curve shows an increasing trend, implying a large unseen mass component that extends beyond the visible galaxy: the dark matter halo. Observations such as these generally show a more or less flat rotation curve at large distances and have historically been of great importance for recognizing the dark matter problem. Moreover, it is of paramount importance for the search for dark matter particles, except for searches at colliders.

1.4 DARK MATTER PARTICLE DETECTION

As seen in the previous sections, there is overwhelming evidence that dark matter exists. Dark matter cannot consist of any of the particles in the standard model: they either interact too strongly, or, in the case of neutrinos, are not abundant enough to reproduce the dark matter relic density [28]. From cosmology only, there is little guidance as to what the properties of the particles are, giving rise to a

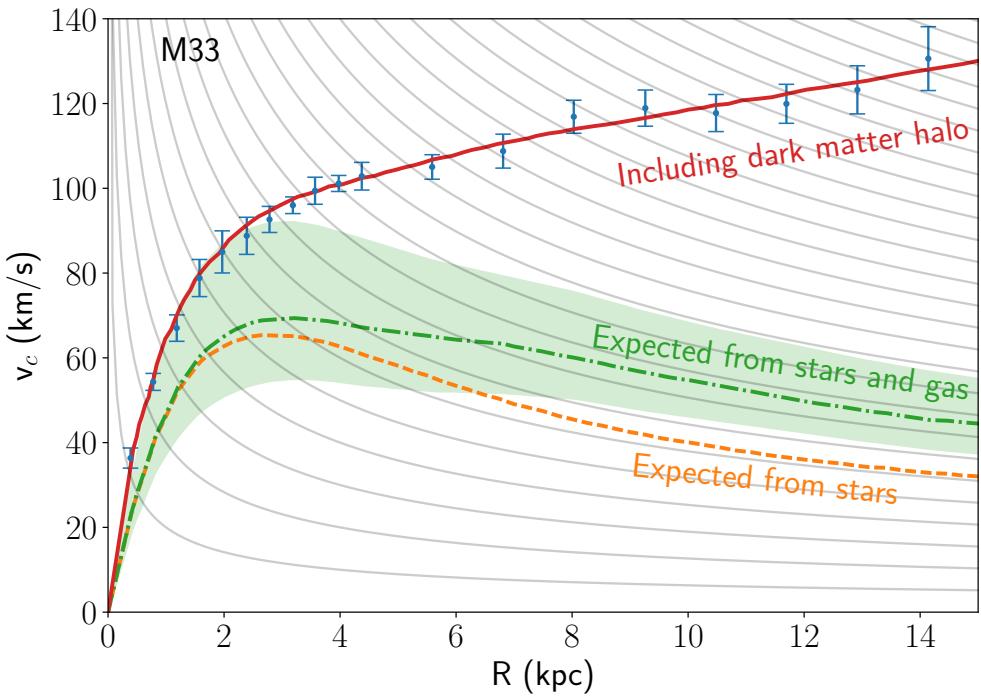


Figure 1.6: The rotation curve of M33, measured using the Doppler shift from the 21 cm radio emission of atomic hydrogen. The orange dashed line shows the rotation curve expected from the matter in stars, while the green dash-dotted line gives the total contribution from ordinary matter in the form of stars and gas with the green band showing the systematic uncertainty on the mass model. At large radii, the distributions trace the gray background lines proportional to $R^{-1/2}$. The measurements show no such decrease and disagree unless a large unseen mass component, the dark matter halo, is introduced (solid red line). Data taken from [27], source code available at [10].

large ‘zoo’ of hypothesized particles that evade the constraints on mass and cross section and usually simultaneously solve outstanding problems in particles physics. The most popular class of particles is that of WIMPs, largely due to the elegant production at the right abundance through freezeout (section 1.2.3) and the fact that it is introduced roughly at the scale that supersymmetry would be expected ($\mathcal{O}(100 \text{ GeV})$ or more). Here and throughout this work, we assume that dark matter particles are WIMPs; a good review of other candidates can be found in [28].

There are, in general, three ways to try to detect dark matter particles, each of which comes with its own advantages and disadvantages. If dark matter particles are light enough, they may be created in high-energy collisions at particle colliders such as the Large Hadron Collider (LHC). Since dark matter must be stable, it escapes the detector unseen, which is observed as missing energy and momentum

in the collision. Unfortunately, dark matter constraints set by collider experiments are highly model dependent, but given some assumptions the constraints are on par with or better than direct detection experiments, especially if the dark matter mass is less than half the Higgs mass (125 GeV) or if it interacts only through spin-dependent interactions [29–31].

The second way that WIMPs may be detected is indirectly, by looking for particles resulting from dark matter particle-antiparticle annihilations. The detected particles can be gamma radiation, neutrinos or charged antimatter, such as positrons. Since the reaction rate depends on the density of both initial particles (which can be the same, if the dark matter particle is its own antiparticle), the signal strength depends on the dark matter density squared, so that annihilation signals are expected to come from high density regions. Examples of targeted regions are the center of the Milky Way, the center of the Sun or dwarf spheroidal galaxies. Complicating factors of indirect detection are the heavy dependence of the dark matter density, astrophysical uncertainties and background sources. Recently, a tantalizing excess of gamma rays with energies of a few GeV from the Galactic center was found by the Fermi-LAT instrument, which neatly illustrates this issue: though the signal fits a dark matter signature, the Galactic center is very bright in gamma rays and uncertainties are large, so that the dark matter interpretation remains a point of vivid discussion [32–34].

The third method of detecting WIMPs is by looking for an interaction with ordinary matter. In such a collision, kinetic energy from a WIMP is transferred to normal matter, whose recoil energy can then be detected. As will be shown in section 1.5, the energy transfer is usually minor, so that $\mathcal{O}(\text{keV})$ energy thresholds are required. In addition, exceptional shielding from and rejection of backgrounds from ambient radiation are required. This simultaneously sets the ultimate sensitivity for direct detection experiments, since neutrinos will eventually give a background that is impossible to beat, except possibly by using the as yet unobtainable directional information [35].

1.5 DIRECT DETECTION

The aim of direct detection experiments is to detect the energy transfer of a dark matter particle that collides with a detector on Earth. The scattering rate and the recoil energy spectrum are set by a range of parameters; the dark matter density, mass and velocity distribution, the type of interaction and the target used.

1.5.1 Spatial and velocity distribution

The spatial and velocity distribution of dark matter in the Galaxy are important for direct detection experiments for two reasons. Firstly, the *local* dark matter density determines the number of particles that can scatter in the detector and is therefore proportional to the scattering rate. Secondly, the velocity distribution determines the energy that is available in the interaction and therefore influences the energy spectrum. The flattening of the rotation curves, as mentioned in section 1.3, gives some guidance on the spatial distribution, since this occurs at large radii where the baryonic matter contribution to the total mass becomes subdominant. This implies that the dark matter must be spatially distributed as $\rho(r) \sim r^{-2}$, since equation 1.7 then yields the observed constant circular velocity through:

$$M(R) = \int_{r < R} \rho(\vec{r}) d^3r = 4\pi A \int_0^R \frac{r^2}{r^2} dr = 2\pi A R \Rightarrow v_c = \sqrt{4\pi A G}, \quad (1.8)$$

where A is a normalization constant and baryonic matter is neglected. A description of the dark matter halo that satisfies this spatial distribution is obtained by assuming thermodynamic equilibrium, implying that the dark matter phase space density is described by:

$$f(\vec{r}, \vec{v}) \sim \exp \frac{-E(\vec{r}, \vec{v})}{k_b T}, \quad (1.9)$$

where T is the dark matter temperature, k_b is Boltzmann's constant and the energy E is the sum of kinetic and gravitational energy:

$$E(\vec{r}, \vec{v}) = \frac{1}{2} m_\chi |\vec{v}|^2 + m_\chi \Phi(\vec{r}), \quad (1.10)$$

with m_χ the dark matter particle mass and Φ the gravitational potential. This can be combined with Poisson's equation for gravity:

$$\nabla^2 \Phi = 4\pi G \rho(\vec{r}), \quad (1.11)$$

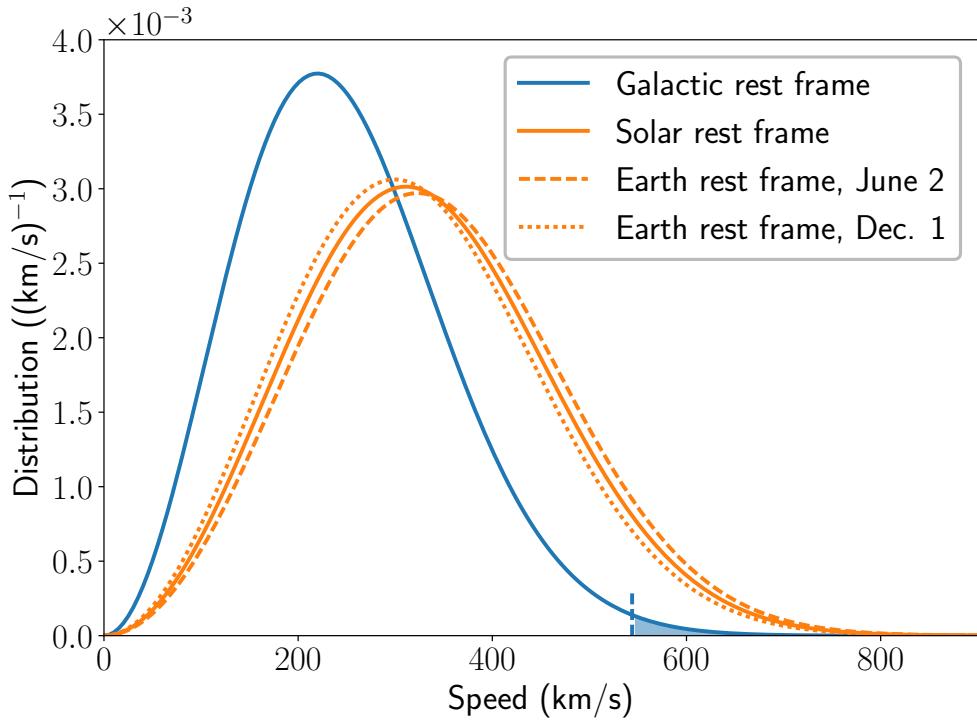


Figure 1.7: The normalized dark matter speed distribution in different rest frames for an isothermal halo model. The parameters assumed are $v_0 = 220 \text{ km/s}$, $v_\odot = 232 \text{ km/s}$, increased by 15 km/s in June and decreased in December due to the Earth's motion around the Sun. The effect of the finite escape velocity (544 km/s, indicated by the dashed blue vertical line) is neglected. Source code available at [10].

with the density ρ obtained by $\int f(\vec{r}, \vec{v}) d^3v$. This yields a differential equation that is solved by the spatial and velocity distribution:³

$$\rho(\vec{r}) = \frac{k_b T}{2\pi G m_\chi} \frac{1}{|\vec{r}|^2} \text{ and} \quad (1.12)$$

$$f(\vec{v}) = \left(\frac{m_\chi}{2\pi k_b T} \right)^{3/2} \exp \frac{-m_\chi |\vec{v}|^2}{2k_b T}, \quad (1.13)$$

where the velocity distribution has been normalized (for the full derivation, see [36, 37].) This is called the *isothermal* density profile, since the temperature is the same for all positions in the halo. To satisfy the relation from equation 1.8, the exponential factor in equation 1.13 must equal $-v^2/v_0^2$, with v_0 equal to the rotation velocity far away from the Galactic center, $v_0 \equiv \lim_{R \rightarrow \infty} v_c(R) \approx 220 \text{ km/s}$.

³ Note that these functions are three-dimensional distributions, i. e. $f(\vec{v}) d^3v$ and $\rho(\vec{r}) d^3r$. The density in terms of the magnitude of \vec{v} and \vec{r} yields a factor $4\pi v^2$ and $4\pi r^2$ respectively, due to the spherical integration.

The velocity distribution in equation 1.13 is not the final distribution for (Earth-borne) direct detection experiments, since the Earth is moving with respect to the Galactic rest frame. This breaks the spherical symmetry and gives a net increase of the average speed. The velocity of the Earth is composed of two parts: a large component due to the motion of the Solar System in the Galaxy and a smaller, yearly-modulating addition due to the motion of the Earth around the Sun. Figure 1.7 shows the effect on the speed distribution of dark matter particles. The net motion of the Solar System is key in directional detection experiments: a larger flux and higher energy collisions are expected to come from the direction of the Solar System in the Galactic rest frame. The yearly modulation can be used as a signature of dark matter interactions [38], though the effect is small ($\sim 3\%$) and requires excellent understanding and control of all radioactive backgrounds.

Although the dark matter distribution in the Milky Way remains a topic of study, the isothermal halo model has become a standard within the direct detection community since these distributions are required to calculate event rates and relate parameter space exclusion limits across experiments. Values commonly used are $v_0 = 220 \text{ km/s}$, $\rho_0 \equiv \rho(R = R_\odot) = 0.3 \text{ GeV/cm}^3$ and $v_\odot = 232 \text{ km/s}$ (see, for instance, [39]). In addition, the velocity distribution in the Galactic rest frame is usually truncated at the escape velocity of 544 km/s. In the distribution with the aforementioned parameters, this cuts only 0.66 % of the distribution, so that the curves are qualitatively the same, though the effect on the maximum allowed energy deposition will affect the high end of the recoil energy spectrum.

1.5.2 Scattering rate

Given the dark matter density and velocity distribution, we can formulate the expected scattering rate per unit mass in a detector as a classical scattering experiment:

$$R = \frac{N_A}{A_T} \sigma \Phi = \frac{N_A}{A_T} \sigma n v = \frac{N_A}{A_T} \frac{\rho}{m_\chi} \sigma v, \quad (1.14)$$

with N_A Avogadro's number, A_T the atomic mass of the target, Φ the flux and n the dark matter number density ρ/m_χ . The complicating factors are twofold: firstly, as described in section 1.5.1, the dark matter velocity is a distribution rather than a fixed value, and secondly, the cross section in general depends on the dark matter velocity and the transferred momentum, and therefore the recoil energy E_R . This

transforms equation 1.14 into:

$$\frac{dR}{dE_R} = \frac{N_A \rho_0}{A_T m_\chi} \int v \cdot f(v) \cdot \frac{d\sigma}{dE_R}(E_R, v) dv. \quad (1.15)$$

where we use the *speed* distribution $f(v)$ in the detector frame and we have changed to differential rate and cross section. If we take an isotropic scattering cross section and if we take the speed distribution as the Maxwell-Boltzmann distribution in the Galactic rest frame with infinite escape velocity (the blue curve in figure 1.7), it can be shown that the recoil spectrum has an exponential form [40]:

$$\frac{dR}{dE_R} = \frac{R_0}{E_0 r} \exp\left(\frac{-E_R}{E_0 r}\right), \quad (1.16)$$

where R_0 is given by 1.14 with $v = \langle v \rangle = \frac{2}{\sqrt{\pi}} v_0$, $E_0 = \frac{1}{2} m_\chi v_0^2$ and r is a kinematic factor that yields the maximum fraction of energy that can be transferred in a collision between two particles of mass m_χ and m_N (the mass of the nucleus):

$$r = \frac{4m_\chi m_N}{(m_\chi + m_N)^2}. \quad (1.17)$$

The qualitative effect of the spectrum of including the Earth's velocity is a boost of the energy spectrum, as the average velocity is higher, and including the escape velocity introduces a cutoff in the spectrum at $E_R = E_{R,max} = r \frac{1}{2} m_\chi v_{esc}^2$.

The scattering cross section σ in equations 1.14 and 1.15 is by construction the WIMP-nucleus cross section. However, the particle physics quantity of interest is the WIMP-*nucleon* cross section. This cross section is usually split into two parts: one that couples to the spin of the nucleus, called the spin-dependent (SD) cross section, and one that couples to the nucleons themselves, the spin-independent (SI) cross section. In the case of the spin-independent cross section, the nucleon scattering amplitudes add up coherently if the transferred momentum is low compared to the nucleus size. As an example, a 50 GeV WIMP colliding onto a xenon target gives $\langle E_R \rangle = E_0 r = 11$ keV with a momentum transfer of 53 MeV and a corresponding de Broglie wavelength h/p of 23 fm, which is of the same order as the xenon nucleus diameter of approximately 12 fm so that partial coherence is expected. This effect may be described by a nuclear form factor $F(q)$ that depends on the momentum transfer q and modifies the cross section according to $\sigma(q) = \sigma_0 F^2(q)$.

Given the type of interaction (spin-dependent or spin-independent), the energy spectrum can be calculated for each target used, taking into account the details of the speed distribution and the form factor. There are two remaining unknowns: the

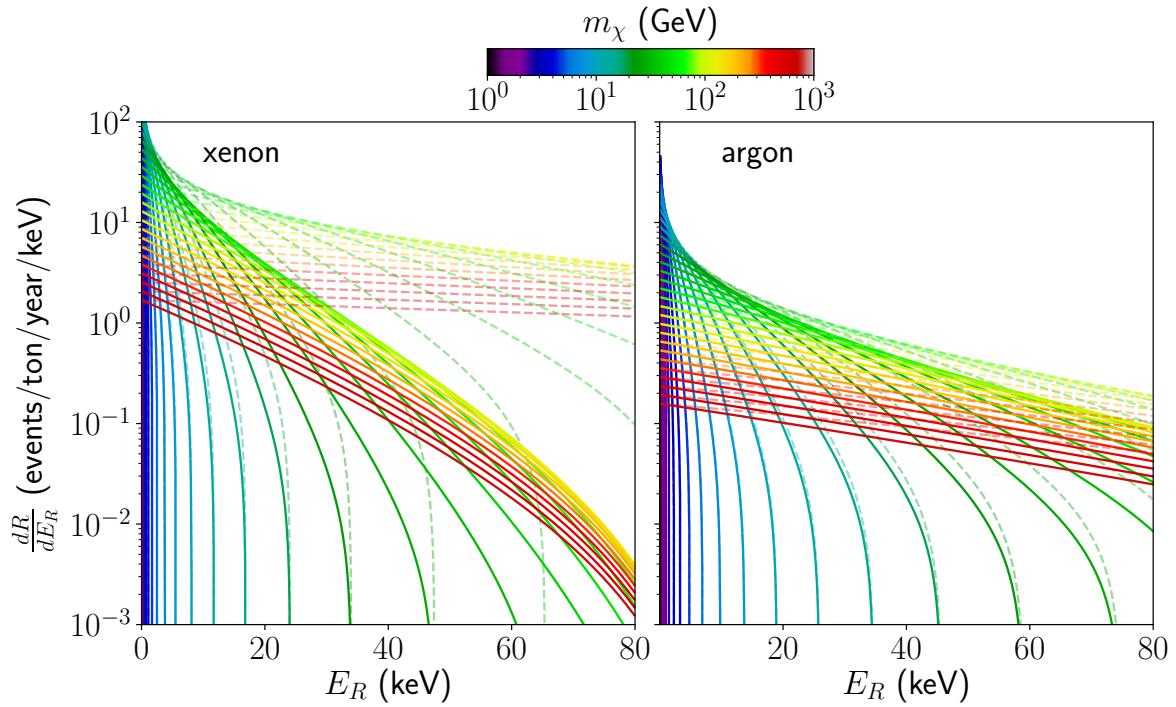


Figure 1.8: Expected recoil energy spectra for various WIMP masses, for xenon (left) and argon (right) targets. All curves are computed for a fixed spin-independent WIMP-nucleon cross section of $1 \times 10^{-45} \text{ cm}^2$. The dashed curves neglect the nuclear form factor and are a good approximation if the recoil energy is low or for light-mass targets. For high-mass WIMPs, the spectra extend to higher energies up to a maximum set by the escape velocity. The rate is in general higher for xenon than for argon due to the A^2 -dependence of the rate. For high WIMP masses, the number density is low and a lower rate is expected. The calculations are based on derivations found in [40] and [41] and source code is available at [10].

WIMP-nucleon cross section, that scales the overall interaction rate, and the WIMP mass, that in combination with the choice of target material changes the spectral shape as well as the total rate. Figure 1.8 shows the expected recoil spectra for various WIMP masses at a cross section $\sigma_{SI} = 1 \times 10^{-45} \text{ cm}^2$ for two common targets: xenon ($A = 131$) and argon ($A = 40$). The dashed lines indicate the spectrum neglecting the form factor contribution.

The spectra are smoothly decreasing functions of recoil energy that are well approximated by the exponential distribution (equation 1.16) for low recoil energies and relatively high WIMP masses. If the mass is very low, the energy transfer is low due to two reasons: the kinetic energy of the dark matter particles is proportional to m_χ , and the kinematic fraction r is low when the target mass is very different

from the dark matter mass. As the mass increases, the spectrum shifts to higher energies up to the maximum energy due to the finite escape velocity. At higher energies, the momentum transfer increases as well, lowering the rate due to the loss of coherence, as seen by the difference between the dashed and solid curves. This is especially noticeable for high mass targets, as the nuclear radius scales as $r_n \sim A^{1/3}$ and q scales as $(E_R A)^{1/2}$, so that the effect is smaller for argon than for xenon at equal recoil energy. In addition, an overall decrease of all spectra is observed, coming from the $1/m_\chi$ dependence in equation 1.14 that has its physical origin in the decreased number density at equal mass density. When xenon and argon are compared, it is clear that xenon target gives a higher rate, especially at low recoil energies. This is due to the coherent scattering contribution that scales the cross section with A^2 , and this (partly) motivates the choice of heavy nuclei for direct detection experiments. Since spin-dependent cross sections scale with nuclear angular momentum and therefore have no A^2 -like scaling, the limits are usually less stringent so that the main focus of many direction experiments is on the spin-independent cross section. Nevertheless, a spin-dependent cross section limit can be set using xenon, as it contains odd-numbered isotopes that have a nonzero nuclear spin.

PARTICLE DETECTION IN XENON

As shown in chapter 1, xenon is an excellent target for direct detection dark matter searches due to its large mass number and because the nuclear mass is close to the expected WIMP mass of $\mathcal{O}(100\text{ GeV})$. In addition, there are several practical aspects that make it possible to build large-scale, low-background experiments using xenon as a detector medium. Since the development of dual-phase liquid xenon time projection chambers (TPCs), the sensitivity to WIMP recoils has roughly doubled every year for the last decade: a stunning achievement that shows the remarkable success of this type of detector [5, 42].

The liquid xenon TPC that currently sets the most stringent limit to the WIMP-nucleon interaction cross section is XENON1T, an experiment with a multi-ton liquid xenon target housed at Laboratori Nazionali del Gran Sasso (LNGS) in an underground laboratory [5]. As an R&D effort for the XENON1T experiment, its successors XENONnT and DARWIN [43] and for liquid xenon particle detection in general, a small-scale dual-phase TPC, called XAMS, was constructed at Nikhef in Amsterdam. Both of these setups are the topic of this thesis. Although XENON1T and XAMS differ by over 3 orders of magnitude in target mass and radio-purity, their operating principles are notably similar. In this chapter, we will introduce the operating principle of liquid xenon TPCs and some of the backgrounds currently limiting WIMP direct detection searches. *There is more than meets the eye...*

2.1 DETECTION PRINCIPLE

In liquid xenon TPCs, there are two distinct signatures for each interaction: the first signal, S_1 , is due to scintillation light, while the second signal, S_2 , is due to liberated electrons. A diagram of the operating principle of XENON1T is shown in figure 2.1. When a particle interaction occurs in liquid xenon, energy is transferred to either the xenon nucleus or one of the atomic electrons, resulting in a short track of excitations and ionizations due to the recoiling particle. The excitations (and some of the ionizations, see section 2.3) cause scintillation light, which is detected by photomultiplier tube (PMT) arrays on the bottom and the top of the detector. An

electric field is applied across the liquid xenon volume using grid wires or meshes, pulling the ionized electrons up towards the xenon gas layer at the top of the TPC. At the liquid-to-gas interface, the electrons are extracted into the gaseous xenon by a stronger electric field, causing the secondary scintillation signal S_2 . Both signals cause light in the UV range at 178 nm, which is directly measured by the PMT arrays. Typically, the S_2 signal is much larger, since each electron extracted from the liquid generates $\mathcal{O}(100)$ scintillation photons.

2

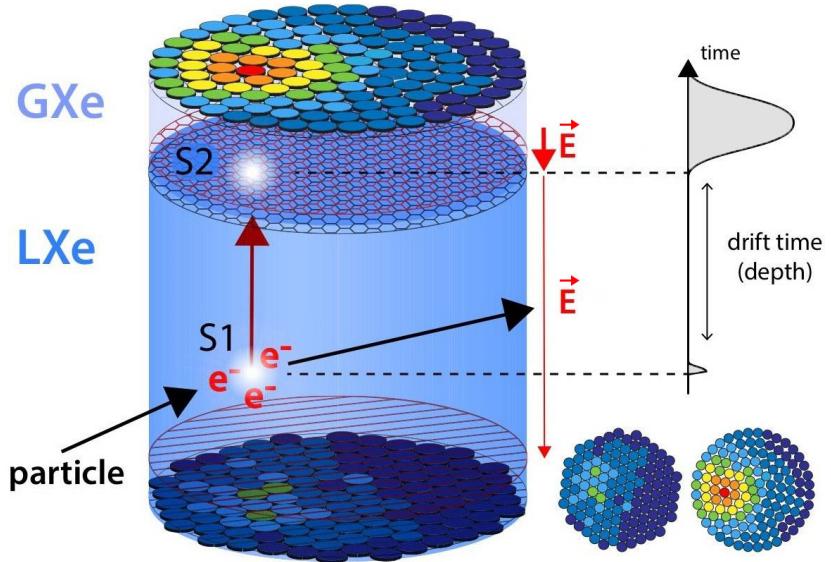


Figure 2.1: Schematic of an interaction occurring in the XENON1T time projection chamber (TPC). A particle interacts in the liquid xenon (LXe) and transfers its energy to a xenon nucleus or an electron. The recoiling particle generates excitations and ionized electrons. The excitations generate direct scintillation light, which is detected by photomultiplier (PMT) arrays at the top and the bottom of the TPC. This signal is called S_1 . The electrons are drifted towards a gas layer (GXe) at the top of the TPC, where they are extracted into the gas and produce secondary scintillation. This signal, S_2 , is detected by the same PMT arrays. The resulting PMT sum signal waveform is indicated by the trace on the right. The x - and y -positions are derived from the S_2 light distribution in the top array, indicated by the colored patterns.

The two signals provide a wealth of information about the original interaction. Firstly, the three-dimensional position can be reconstructed from each S_1 - S_2 pair. The depth of the interaction (the z -coordinate) is proportional to the drift time, which is defined as the time delay between the S_1 and the S_2 , because of the constant drift velocity of the electrons. This typically generates a delay of several μs up to almost 1 ms (depending on the TPC size and applied field), as the drift velocity is

of the order of $1 \text{ mm}/\mu\text{s}$. The position in the horizontal plane (x - and y -positions) is inferred from the S_2 light distribution across the top PMT array, which is highly dependent on the interaction position as the secondary scintillation emission occurs close to that array. In addition to the interaction position, the relative S_2 and S_1 size makes particle identification possible. WIMP recoils are expected to interact with a xenon nucleus, causing a *nuclear recoil* signal, while most of the background interactions cause electronic recoils with a larger S_2/S_1 -ratio.¹ Finally, if there are multiple interaction positions at significantly different z -positions, there will be multiple S_2 -signals with different drift times. Since the probability for a WIMP to scatter more than once is extremely low, these events are indicative of backgrounds and can be rejected.

2.2 PROPERTIES OF XENON

The use of liquid xenon for dark matter detection experiments has several advantages. Table 2.1 lists some of the properties that are relevant to the use of liquid xenon as a detector medium. Xenon is a heavy element, with a mean atomic weight of 131.3, giving a large cross section due to coherent scattering (see section 1.5.2). Its relatively high boiling point means that a moderate cryogenic installation is required to liquefy the xenon. Liquid xenon detectors are therefore housed in a vacuum-insulated double-walled vessel. Since xenon has a high density, only a small detector volume needs to be instrumented to reach a relatively high detector mass.

Liquid xenon is an excellent scintillator, with a maximum scintillation yield of 73 photons/keV; higher than most commonly used scintillators (such as NaI(Tl) with 38 photons/keV). The high yield enables a low energy threshold (1.4 keV for XENON1T [5]) that is required due to the exponentially falling energy spectrum of WIMP recoils. The scintillation wavelength of liquid xenon is in the UV range at 178 nm, directly detectable by special UV-sensitive PMTs.

Large liquid xenon detectors benefit from self-shielding due to the low path length of background radiation in the medium. Table 2.1 lists the range of gamma and neutron radiation in liquid xenon for several energies. As the typical gamma-ray path length is low with respect to the detector size, most of these background events are located within a few cm of the edge of the detector volume. The in-

¹ In principle, electronic recoils are not excluded, however, for elastic WIMP recoils, the maximum energy transfer that is kinematically allowed is too small to be detectable (see equation 1.17). Dark matter other than WIMPs or other types of dark matter interactions may give electronic recoils or higher energy nuclear recoils, see, for instance [44–46].

Table 2.1: Selected properties of xenon.

	Property	Value	Reference
General	Atomic number	54	[47]
	Atomic mass	131.3	[47]
	Boiling point (at 1 bar)	–108 °C	[47]
	Liquid density (at boiling point)	3.057 g/cm ³	[47]
Scintillation	Maximum yield	73 photons/keV	[48]
	Wavelength	178 nm	[47]
Path lengths	100 keV gamma	0.18 cm	[49]
	1 MeV gamma	5.9 cm	[49]
	1 MeV neutron	10.6 cm	[50]
	10 MeV neutron	15.9 cm	[50]

2
nermost volume, called the *fiducial volume*, is therefore a region of very low background. As the position for each event in liquid xenon TPCs is known, most of the gamma-ray background and some of the neutron background can be removed by a selection on interaction position.

2.3 LIGHT AND CHARGE GENERATION

When a particle interacts in liquid xenon, it transfers some of its energy to the recoiling particle, which can be either a xenon nucleus or an electron around the nucleus. The first type is called a nuclear recoil (NR) signal. In this interaction, the xenon atom usually partially ionizes as a result of the transferred energy. The second interaction type, in which an electron is liberated from the xenon atom, is called an electronic recoil (ER). Although energetic electrons caused by beta decay of materials within the xenon are technically not recoiling particles, the experimental signature is the same and the signal is also called ER. Both ERs and NRs result in an energetic particle that moves though the high density liquid xenon and thus quickly loses its energy. This creates a track of excitations and ionizations that are at the root of the observable signals, S₁ and S₂, respectively. In the case of an NR, some of the energy is lost to atomic motion through atomic collisions. This mechanism is called nuclear quenching. The three main energy dissipation channels are shown by the lines originating in the top left in figure 2.2.

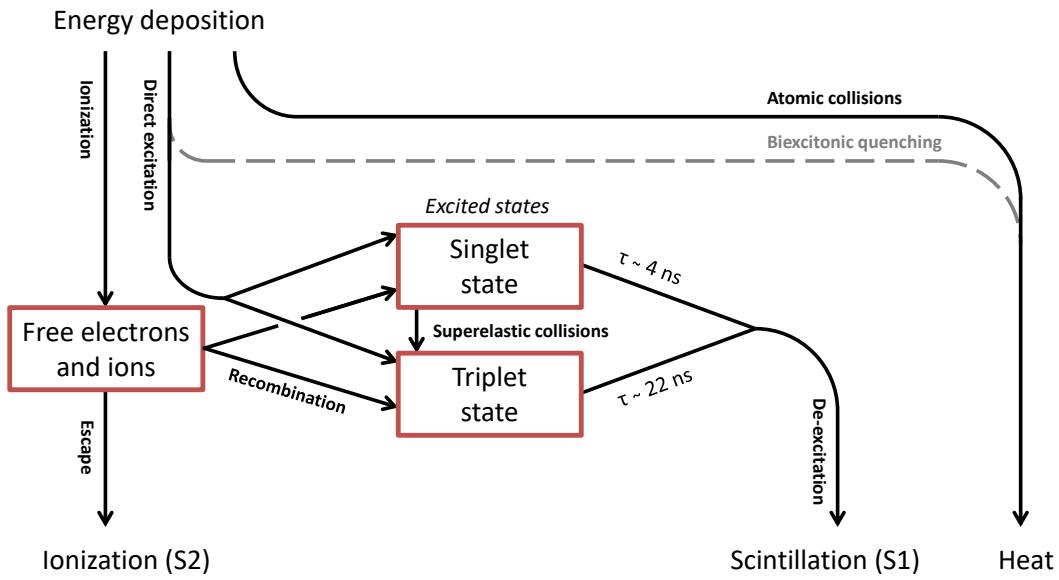


Figure 2.2: After an energy deposition in liquid xenon, different processes lead to ionization, scintillation and heat. The quenching process indicated by the gray dashed line is only relevant for high ionization density recoils. Figure adapted from [2] (chapter 5).

2.3.1 Direct luminescence

There are two paths that can cause scintillation light: direct excitation or recombination. In the first of these, excited xenon atoms are directly formed by the recoiling particle in the original track. In less than 1 ps after the excitation, the excited xenon atom (exciton) forms a bound state with a stable xenon atom through the self-trapping process $\text{Xe}^* + \text{Xe} \rightarrow \text{Xe}_2^*$. The bound dimer state Xe_2^* is called an excimer. The excimer can be in two spin states: a singlet or a triplet state, referring to the combined spin state of the electron and the angular momentum due to the molecular orbit. The singlet and triplet states have decay times of approximately 4 ns and 22 ns, respectively [51]. When the state decays, the molecule dissociates and the energy is emitted in the form of a 178 nm ultraviolet (UV) photon. Since the dimer state is at a lower energy level than the excitation energy of an individual xenon atom, the medium is transparent to the UV light so that the scintillation light can be detected over significant distances.

2.3.2 Recombination luminescence

The second process that can cause scintillation light is through recombination of electrons and ions. Free xenon ions form charged dimer states Xe_2^+ which can form doubly-excited excitons Xe^{**} when they recombine with free electrons according to the reaction $\text{Xe}_2^+ + e^- \rightarrow \text{Xe}^{**} + \text{Xe}$. This then relaxes to the first excited state Xe^* and forms a excimer state similar to the direct excitation process. A sizable fraction of the scintillation light is generated through recombination; for $\mathcal{O}(\text{MeV})$ ERs, approximately 70 % of all light can be attributed to recombination luminescence [52–54]. When an electric field is applied, some of the electrons that could participate in recombination are extracted from the interaction site, so that the light yield decreases.

2.3.3 Particle type and ionization density dependence

In liquid xenon TPCs, the main discrimination between ER and NR events is due to the difference in their light and charge yield: NR events give a smaller S2 signal for similar S1 sizes. This is mostly due to a difference in the energy distribution of the primary energy channels shown in figure 2.2, which is characterized by the exciton to ion ratio $N_{\text{ex}}/N_{\text{ion}}$. For ERs, $N_{\text{ex}}/N_{\text{ion}} \approx 0.13$ [55, 56], while NR events have $N_{\text{ex}}/N_{\text{ion}} \approx 1$ [48, 57]. Additional differences in the scintillation and ionization signals are caused by a difference in the density of ionizations and excitations in the track. This can be characterized with the linear energy transfer (LET), a measure of the energy transferred to the medium per cm along the path. ERs have a low LET and give a track with a typical length of 10 to 100 μm . NRs, on the other hand, result in a very dense core of excitations and ionizations with a surrounding penumbra of lower excitation density caused by delta rays. Because of the high ionization density, the electrostatic force from the ions at the interaction site makes it difficult to extract the electrons, while only moderate electric fields are required for ERs.

For very high LET events, such as caused by alpha decays or fission fragments, a reduced scintillation yield has been observed. This is likely due to *biexcitonic quenching*, in which two excitons produce one ionization according to the reaction $\text{Xe}^* + \text{Xe}^* \rightarrow \text{Xe}^+ + e^- + \text{Xe}$. In the absence of this process, two observable quanta are produced, while only one quantum (either an electron or a photon, if recombination occurs) is produced when this reaction occurs. Note that quenching of excimer states does not appear to be the major quenching mechanism, which would be evidenced by a shorter observed singlet and triplet lifetimes if quenching occurs.

2.3.4 Scintillation pulse shape

Most of the aforementioned reaction steps that lead to the formation of excimers occur on short timescales compared to the excimer decay times. However, the recombination process is a notable exception to this, with timescales up to $\mathcal{O}(10\text{ ns})$. This is highly dependent on the ionization density and the electric field, as it depends on the dynamical behavior of the free electrons. Recombination causes a delay in the formation of excimer states that is reflected in the scintillation pulse shape. The delaying effect is most notable for high-energy (MeV) ERs at zero field, where the ionization density is relatively low. In this case, a single decay time of 45 ns is observed that obscures the double-exponential scintillation intensity time dependence, which would be expected based on the two excimer states with lifetimes of 4 ns and 22 ns. If the ionization density is higher or if an electric field is applied, the effect on the pulse shape from recombination delay is less profound, so that it may be described with increased effective singlet and triplet lifetimes that are longer than the true lifetimes. For high ionization density tracks, recombination is fast enough so no delay is observable, resulting in the same measured lifetimes for all particles.

Apart from the delayed luminescence caused by recombination, the scintillation pulse shape is fully characterized by the relative population of the singlet and triplet states. This may be described by the singlet fraction f_s ; the fraction of all scintillation light resulting from the decay of the singlet state. The singlet fraction depends on several factors, such as the particle type, recoil energy and applied electric field. The dependence of particle type and recoil energy is partly caused by a difference in the primary energy distribution into excitations and ionizations, as they both form excimers, but not necessarily in the same singlet to triplet ratio. In addition, both processes may have their own dependence on field and ionization density. This results in a nontrivial behavior of f_s .

In liquid xenon, it has been shown that the singlet fraction increases with LET. Interestingly, this is opposite to the observed effect in organic scintillators. For organic scintillators, the probability of recombination with the same ion that the electron was released from (geminete recombination) increases for low LET tracks. This type of recombination predominantly produces a singlet state ($f_s = 1$) as the electron spin is conserved in the process [58]. On the other hand, if the spin is randomly oriented, the probability of producing a singlet state is dictated by the multiplicity of the singlet and triplet states, so that the triplet state is three times more likely than the singlet state and $f_s = 0.25$. This would be expected for recombination in

high LET tracks, as the electrons can recombine with many ions in its vicinity. In contrast to these expectations, f_s increases with LET in liquid xenon. A suggested reason for this is the occurrence of superelastic collisions between excimers and electrons that change the singlet state to a triplet state [59]. Since this requires a free electron, this process is more likely if recombination is slow, which is the case for low LET tracks. Note that the reverse process is less likely due to the slightly lower energy level of the triplet state.

2.3.5 Pulse shape discrimination

For zero electric field, the broad pulse shape of ERs due to the slow recombination and low singlet fraction can be used to discriminate ER and NR signals. This was used in early single-phase liquid xenon dark matter detectors [60, 61]. However, liquid xenon TPCs allow for discrimination based on S_2/S_1 by applying an electric field that simultaneously diminishes the difference between ER and NR scintillation traces. Moreover, the S_2/S_1 method has been shown to be much more powerful, especially at low energies. In chapter 5, the pulse shape dependence for ER and NR signals at low energies with applied fields is investigated further. At low energies, the pulse shapes of ER and NR signals appear to be more similar than at high energies; an effect that can be intuitively understood as the LET becomes more similar. The effect of applied field on the ER pulse shape is described in chapter 6.

2.4 SIGNAL CORRECTIONS AND INTERACTION RECONSTRUCTION

The S_1 and S_2 signal generated for each event in liquid xenon TPCs are used to derive the physical properties of the interaction, most notably the interaction position, recoil energy and recoil type (ER or NR). The reconstruction of these properties requires detailed understanding of the TPC. There are several corrections that need to be applied, which are based on calibration data, predominantly from mono-energetic gamma-ray sources. For small TPCs, an external source (such as ^{137}Cs or ^{22}Na) can be used. However, for TPCs such as XENON1T, the penetrating power of such sources is too low to be used for calibration of most of the detector, so that internal calibration sources must be used. For most calibrations of XENON1T, an internal $^{83\text{m}}\text{Kr}$ source is used. This decays emitting a mono-energetic gamma ray of 32.2 keV and a second gamma, delayed with a half-life of 155 ns, at 9.4 keV [50]. This provides a very typical event fingerprint of two mono-energetic emissions, and thus gives a very clean calibration signal. Other internal sources are radon-induced

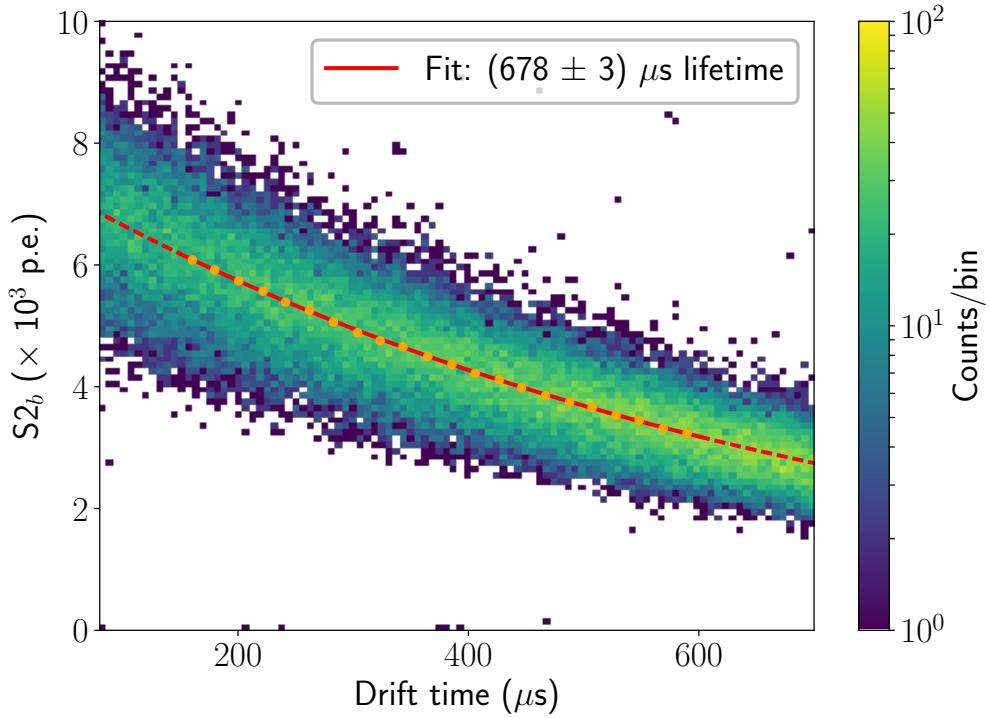


Figure 2.3: The S_2 size, measured in the bottom PMT array, for selected events from a ^{83m}Kr source in XENON1T. The S_2 decreases exponentially with increasing drift time due to electron absorption. This is fit in the range of 150 to 600 μs , yielding an electron lifetime of $(678 \pm 3) \mu\text{s}$.

alpha decays or the decay of neutron-activated radioactive xenon isotopes and inelastic transitions, giving gamma rays of 39.6 keV (^{129}Xe), 80.2 keV (^{131}Xe), 163.9 keV (^{131m}Xe) and 236.1 keV (^{129m}Xe). In this section, the most critical corrections and calibrations for XENON1T are discussed. The procedure used for XAMS differs on certain points and is further explained in chapters 3 and 4.

2.4.1 Electron lifetime

A first correction of the S_2 -signal is due to charge loss during the electron drift from the interaction position up to the top of the TPC. As electrons move through the liquid, there is a finite probability that they attach to impurities in the xenon and are absorbed. This is characterized by the electron lifetime τ_e , which is the mean time before an electron is absorbed. Since the S_2 size is proportional to the number

of electrons that arrive at the top of the TPC, the S_2 is expressed as:

$$S_2(t_d) = S_{20} \exp\left(-\frac{t_d}{\tau_e}\right), \quad (2.1)$$

with t_d the drift time and S_{20} the size of the S_2 if all electrons survived. The applied correction is simply the inverse of this function, so that the corrected S_2 is equal to the S_2 size in the limit of infinite electron lifetime. Figure 2.3 shows the S_2 size as a function of drift time for selected ^{83m}Kr events. The event distribution traces the exponential fit ($\tau_e = (678 \pm 3) \mu\text{s}$), indicated by the red line.

2.4.2 Light detection efficiency

For both the S_1 and the S_2 , the light yield depends on the interaction position due to light detection efficiency (LDE) variation across the TPC. This is due to a variation of optical effects that is difficult to accurately model. Among other factors, the LDE depends on light propagation in the liquid xenon, reflection efficiency at the walls, the angular dependence of the transmission probability into the PMT window and the quantum efficiencies of the PMTs. Since modeling all these effects is a serious challenge, usually a data-derived LDE map is constructed based on a mono-energetic calibration source. The corrected S_1 is then computed by dividing the S_1 by the LDE value at the interaction position, normalized to the volume average within the TPC. For the S_2 , the light is always emitted in the narrow region between the liquid level and anode mesh, so that the correction map is two-dimensional. Since the S_2 scintillation occurs very close to the top PMT array, the LDE variation in this array is highly dependent on accurate reconstruction of the interaction position. For this reason, only the S_2 size in the bottom PMT array (S_{2b}) is used for energy reconstruction. Other than LDE effects, variation of the S_2 size as function of x - and y -position occurs because of changes in the secondary scintillation gain, mostly caused by a difference in the gas gap between the gate and anode mesh. This is a result of the net effect of gravitational sagging of the mesh and the attractive electrostatic force resulting from the high voltage applied. In the case of XENON1T, this results in an S_2 that is 32 % larger in the center than at the edge [5].

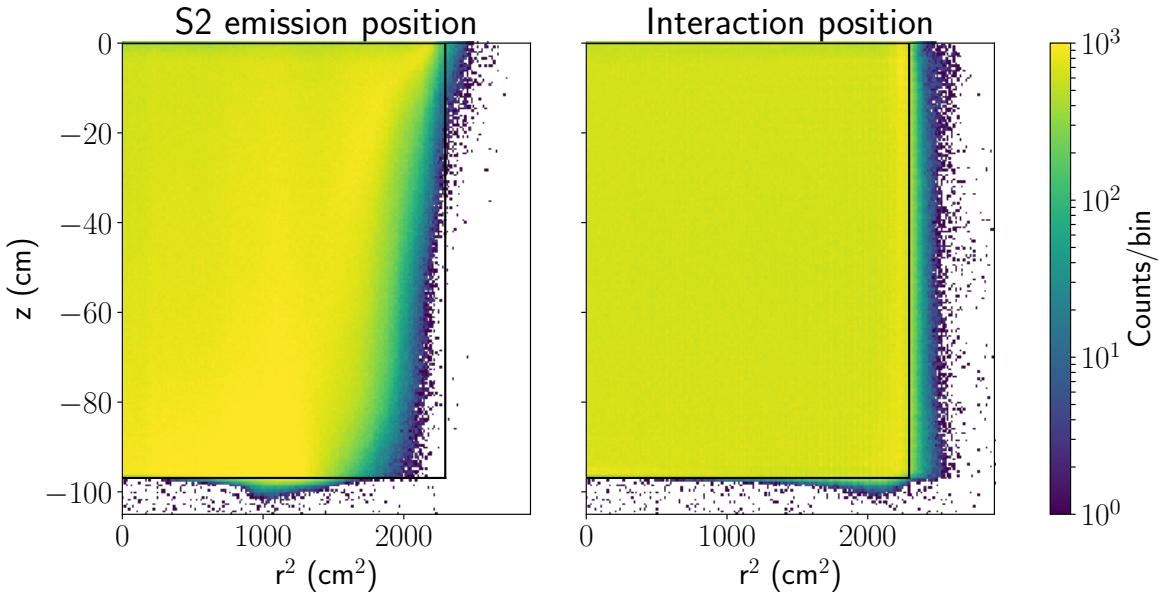


Figure 2.4: Reconstructed x and y interaction position distribution at the $S2$ emission position (left figure) and the corrected interaction position distribution (right figure) for ^{83m}Kr events. Although a uniform distribution is expected for this source, a significant distortion is observed in the uncorrected distribution due to an inhomogeneous electric field. The correction is applied in r and z .

2.4.3 Field inhomogeneity

In section 2.1, we noted that the x - and y -positions are reconstructed from the light distribution pattern across the PMTs in the top array. However, this relies on the implicit assumption that the electrons arriving at the gate mesh are at exactly the same x - and y -position as the original interaction position, that is, that the electrons are drifted straight up. Electrons that drift through the liquid xenon follow electric field lines, so that the drift of the electrons depends on the electric field in the TPC. The field is shaped by rings along the entire length of the TPC, supplied with the appropriate voltages by a resistor divider chain to maximize field homogeneity. Nevertheless, some drift field inhomogeneity remains, distorting the reconstructed interaction positions. This is shown in the left panel of figure 2.4, where a density histogram of the reconstructed position of ^{83m}Kr events is shown. Since this is an internal source, a homogeneous distribution in (r^2, z) is expected, but significant distortion is found predominantly at deep events (high drift time). The drift field inhomogeneity is partly attributed to a lower cathode voltage than the design voltage (8 kV for the distribution shown in figure 2.4, design voltage

50 kV [62]) resulting in a sub-optimal voltage distribution across the field shaping rings. Furthermore, there is evidence that the Teflon walls of the TPC slowly accumulate a negative charge, an effect that was more prominently observed by the LUX collaboration [63]. The supposed cause of this is that the UV scintillation light causes electron-hole pairs in the Teflon. Since the holes have a higher mobility than electrons in Teflon, the holes are preferentially removed by the electric field and a negative charge remains. A distinct signature of this effect is that the field homogeneity ameliorates with time, as charge is building.

2

The position distortion from the field is corrected based on ^{83m}Kr calibration data, where events are required to be spread homogeneously over the TPC's volume.² Based on this, each position is corrected in r . The z -coordinate is corrected for the longer drift over the skewed trajectory. The resulting position distribution is shown in the right panel of figure 2.4. The position correction is validated by confirming agreement with simulation for other uniform and non-uniform sources.

2.4.4 Energy reconstruction

The observable light and charge signals S_1 and S_2 come from the energy transferred to excited dimer states and electron-ion pairs, respectively. The number of photons and electrons may vary due to recombination efficiency or exciton fraction variation. However, as the average energy to produce an excimer state or ionized state is equal, the energy can be reconstructed according to:

$$E = W (n_\gamma + n_{e^-}), \quad (2.2)$$

with the W -value of (13.7 ± 0.2) eV [48]. This can be expressed in the experimentally measured light and charge signals:

$$E = W \left(\frac{cS1}{g_1} + \frac{cS2_b}{g_2} \right), \quad (2.3)$$

where $cS1$ and $cS2_b$ denote the corrected signals and we have defined the primary and secondary scintillation gain g_1 and g_2 in p.e./ γ and p.e./ e^- , respectively. Conceptually, g_1 denotes the probability for a scintillation photon to release a photo-electron in one of the PMTs, while g_2 is the average number of photoelectrons for each electron reaching the liquid-to-gas interface, considering the combined effects

² For the correction procedure outlined in this section, only homogeneity in r is required. This assumption is justified by the relatively high convection velocity, which is expected to give good mixing within the lifetime of ^{83m}Kr (1.83 h).

of extraction efficiency, light production in the gas gap and the light measurement efficiency, so that it is typically greater than one. The values of g_1 and g_2 can be measured by using different calibration sources or different electric fields, so that the signal ratio between S_1 and S_2 is effectively varied. Rewriting equation 2.3, we have:

$$\left(\frac{cS2_b}{E} \right) = \frac{g_2}{W} - \frac{g_2}{g_1} \left(\frac{cS1}{E} \right). \quad (2.4)$$

If multiple values of $cS2_b/E$ and $cS1/E$ are known, g_1 and g_2 can be extracted from a linear fit, setting the intercept g_2/W and slope $-g_2/g_1$. Figure 2.5 shows the energy calibration of XENON1T, used for the dark matter search results of the full exposure [5]. It should be noted that any quenching is neglected in this calibration procedure, which is a good approximation for electronic recoil interactions at these energies.

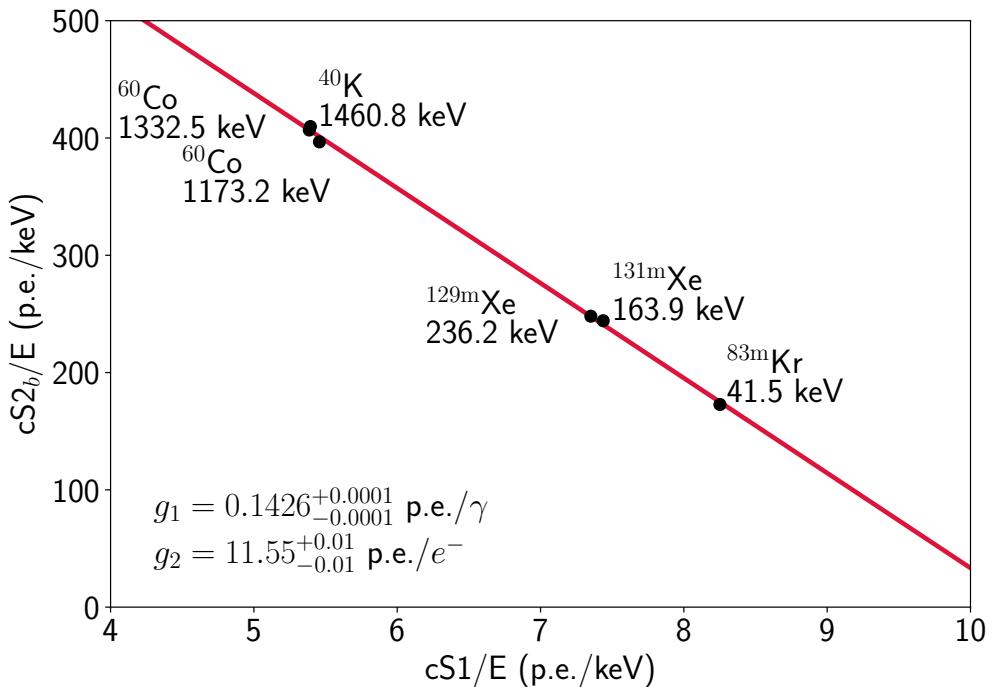


Figure 2.5: Energy calibration of XENON1T. The charge and light yield are anti-correlated: a decrease in charge yield corresponds to an increase in light yield (see equation 2.4). By using various sources with different light to charge ratios, the values of g_1 and g_2 are determined from a linear fit, shown by the red line.

For nuclear recoils, the energy is still computed using g_1 and g_2 derived from the electronic recoil calibration, even though the reconstructed energy does not match

the recoil energy because of nuclear quenching. This difference is usually denoted by using subscripts in the energy units; the ER equivalent energy is denoted keV_{ee} (electron equivalent), while the recoil energy is denoted as keV_{nr} (nuclear recoil). The conversion between these is calculated using a model of nuclear quenching [64].

2.4.5 Interaction type identification

The major background rejection in liquid xenon TPCs (apart from fiducialization) comes from the ability to distinguish nuclear and electronic recoil interactions from their different light and charge yields. In general, nuclear recoils cause a smaller S_2 with respect to the S_1 signal. The distinction between the two interaction types is thus made by defining regions in the space of cS_1 and cS_{2b} , called the ER and NR bands. Since the position and shape of the ER and NR bands depend on detector parameters, such as the electric field and LDE, the ER and NR response need to be calibrated using ER and NR collisions that span the energy range relevant to dark matter searches. For large-scale TPCs, the ER response is most conveniently calibrated using internal beta decay sources, giving a wide range of energies. For the ER calibration of XENON1T, a ^{220}Rn source is used, which has the advantage of good solubility in xenon and a relatively short lifetime (10.6 hours for the longest living isotope in the decay chain). The beta decaying isotope in the ^{220}Rn decay chain is ^{212}Pb with a Q -value of 570 keV. Alternatively, tritiated methane (CH_3T) may also be used [65]. This source lacks the high-energy alpha decays and thus allow for higher rate calibration, but needs to be efficiently extracted from the xenon using cryogenic distillation or getter purification.

For the NR calibration, neutrons with MeV energies can be used as they give energy depositions up to tens of keVs, corresponding to the approximate range of dark matter recoil interactions (see section 1.5.2). Neutron sources that may be used are (α, n) source mixtures (such as $^{241}\text{AmBe}$), spontaneous fission sources (such as ^{252}Cf) [66] or neutrons emitted in a fusion process in a neutron generator.³ In XENON1T, a $^{241}\text{AmBe}$ source was used for the NR calibration of the first science results [39], while a deuterium-deuterium fusion neutron generator was used for the most recent result [5]. Prior to its use as a calibration source, the neutron generator was itself calibrated in an extensive measurement campaign at the

³ In addition to (α, n) source mixtures, (γ, n) mixtures are also a possibility (^{88}YBe , $^{205}\text{BiBe}$) [67]. However, as these give relatively low neutron energies, these are mostly useful for extremely low energy calibrations.

Table 2.2: *Backgrounds causing low-energy recoils in XENON1T [68].*

	Edge (external)	Uniform (internal)
Electronic recoil	γ -emission from materials	Radon daughters ^{85}Kr Neutrino-electron scattering ^{136}Xe double beta decay
Nuclear recoil	Radiogenic neutrons Muon-induced neutrons	Coherent neutrino-nucleus scattering
Other	Wall leakage	Accidental coincidence

Physikalisch-Technische Bundesanstalt (PTB) in Germany and at Purdue University in the United States. The results are described in chapter 7, and are published in [4]. The NR calibration using the neutron generator is further described in chapter 8.

2.5 BACKGROUNDS

In direct dark matter detection experiments, the major challenge is background rejection. Since rare event searches like this require almost no background for optimal sensitivity, exceptional care has to be taken to reduce all backgrounds to extremely low values. This challenge becomes stronger for larger detectors, as the number of events increases with the detector mass. The background is firstly reduced to a minimum by material selection and xenon purification. Secondly, most background events can be rejected by either their interaction position (fiducialization) or their S_2/S_1 signal ratio (ER/NR discrimination). Table 2.2 lists the most relevant backgrounds for XENON1T’s dark matter search, organized by their position distribution and recoil type.

2.5.1 *Electronic recoil backgrounds*

Although the largest background rate, measured in the number of events in the entire detector volume, comes from external gamma ray recoils from slightly radioactive materials, it is one of the minor backgrounds in the full analysis. This is because these interactions are easily distinguishable from dark matter interactions: they occur mostly at the edges of the detector, and are of the ER type. The most relevant ER backgrounds are internal backgrounds. The dominant ER background in both of the XENON1T dark matter search results comes from ^{214}Pb , a β^- -emitting

isotope that comes from ^{222}Rn that is continuously emitted from the internal detector materials and distributed homogeneously throughout the detector volume. Background from ^{85}Kr comes from the krypton contamination in xenon when it is procured, and is reduced by cryogenic distillation of xenon, reaching sub-ppt levels within months in XENON1T [69]. Neutrino-electron scattering (mostly from the pp fusion process in the Sun) causes an ER background at very low energy. This background cannot be mitigated, except by improving ER/NR discrimination. Finally, natural xenon contains ^{136}Xe : a long-living radioactive isotope decaying via double beta decay. However, its impact on the low-energy background is small due to its high decay energy of 2.46 MeV.

2.5.2 Nuclear recoil backgrounds

Nuclear recoil backgrounds have the same experimental signature as WIMP recoils and thus cannot be discriminated from signal interactions. These backgrounds, mainly from neutrons, must therefore be limited to a minimum. Neutrons with energy of $\mathcal{O}(\text{MeV})$ can come from spontaneous fission or (α, n) reactions within the detector materials, mostly from ^{238}U contamination. These backgrounds are external, so can be cut by fiducialization, although neutrons are more penetrating in liquid xenon than gamma radiation. Neutrons are also produced in cosmic ray interactions. As XENON1T is located in an underground lab, the cosmic ray flux is greatly reduced, leaving only a very small fraction of muons (and a negligible flux for any other cosmic ray particle). An instrumented water tank around XENON1T furthermore gives a veto signal for almost any crossing muon by its Cherenkov emission. This reduces the muon-induced neutron background to a negligible level. An ultimate background for liquid xenon TPCs comes from coherent neutrino-nucleus scattering: it is impossible to mitigate, fiducialize or discriminate events from this process. The expected dark matter limit for liquid xenon based detectors bounded by this background is usually called the *neutrino floor*, and it is expected that this limit is reached within the next decade [43].

2.5.3 Other backgrounds

In addition to backgrounds classified as ER and NR based on their relative S_1 and S_2 signal sizes, there are detector effects that give signals that are incompatible with ER and NR. These are indicated by the bottom row in table 2.2. The ‘wall leakage’ background comes from radon daughter nuclei that plate out on the TPC walls.

When they then decay, their observed S₂ size is reduced due to electron attachment to the TPC walls. This causes these ER events to shift to lower S₂ sizes, overlapping with the NR band, but also extending below it. Finally, energy depositions in charge- or light-insensitive regions cause only S₁ or S₂ signals, respectively, which may be accidentally paired and mimic a real interaction. This causes an additional background that does not correlate with one physical event and typically exhibits a low S₂.

COMMISSIONING OF A DUAL-PHASE XENON TPC AT NIKHEF

The rise of liquid xenon TPCs in the field of dark matter direct detection has sparked great interest in liquid xenon particle detection technology. Although large-scale dark matter detectors can be used for R&D purposes after their dark matter search has concluded, many of the advances in liquid xenon technology or the understanding of the physics of signal generation would not have been possible without the aid of small-scale liquid xenon TPC setups. These operate using exactly the same technology as large-scale dark matter detectors, but their small size and lower cost makes it easier to conduct research that would otherwise be too risky or time-consuming. In addition, small-scale detectors can be optimized for specific measurements, for instance, by maximizing the time resolution, light yield or electric field.

In this context, an R&D dual-phase TPC setup, called XAMS, was constructed at Nikhef, Amsterdam. The paper shown in this chapter marks the introduction of XAMS into the field of liquid xenon TPC research. In addition, a PMT gain calibration method based on single photoelectron signals in single electron S2 signals is introduced, similar to a method employed by the LUX collaboration [70]. This has the advantage that it can be applied without the need for special data-taking and that it directly probes the PMT response to UV light, so that effects such as double photoelectron emission may be measured [71]. This new method means that this paper not only introduces a new TPC, but simultaneously shows its first scientifically relevant result to the community.

In the paper shown in this chapter, we introduce technical details of the XAMS setup that is further used in chapters 4, 5 and 6. Chapter 4 shows the analysis steps that are necessary for a basic calibration that is used in chapters 5 and 6. In chapter 5, the scintillation pulse shape at low energies is measured for electronic and nuclear recoils. The effect of the electric field on several properties is explored in chapter 6.

COMMISSIONING OF A DUAL-PHASE XENON TPC AT NIKHEF

3

E. Hogenbirk, J. Aalbers, M. Bader, P.A. Breur, A. Brown, M.P. Decowski, C. Tunnell, R. Walet, A.P. Colijn

Nuclear Instruments and Methods in Physics Research Section A, 840 (2016), pp. 87-96

Accepted 23 September 2016

Available online at

<https://dx.doi.org/10.1016/j.nima.2016.09.052>

<https://arxiv.org/abs/1602.01974>

Reference [1]

3.1 ABSTRACT

A dual-phase xenon time-projection chamber was built at Nikhef in Amsterdam as a direct dark matter detection R&D facility. In this paper, the setup is presented and the first results from a calibration with a ^{22}Na gamma-ray source are presented. The results show an average light yield of (5.6 ± 0.3) photoelectrons/keV (calculated to 122 keV and zero field) and an electron lifetime of (429 ± 26) μs . The best energy resolution σ_E/E is (5.8 ± 0.2) % at an energy of 511 keV. This was achieved using a combination of the scintillation and the ionization signals. A photomultiplier tube gain calibration technique, based on the electroluminescence signals occurring from isolated electrons, is presented and its advantages and limitations are discussed.

3.2 INTRODUCTION

There is considerable evidence from astrophysical observations that there is more mass in the universe than can be accounted for with only standard model particles [72–74]. The most popular theory that explains this discrepancy introduces dark matter particles called WIMPs [75]. In past years, the sensitivity of direct dark matter search experiments has increased by orders of magnitude, lead by the development of large dual-phase xenon time-projection chambers (TPCs) [42, 47, 76]. In the context of dark matter research, a small-scale liquid xenon TPC, called XAMS (Xenon Amsterdam), has been designed, built and commissioned at Nikhef in Amsterdam. The setup described in this work is similar to small-scale dual-phase xenon setups, such as described in [77–79].

Dual-phase TPCs detect a particle interaction using two distinct signals. The first comes from excitations and recombined electron-ion pairs. Bound excited states of two atoms form, and subsequent decays of these excitons causes scintillation light that is detected by photomultiplier tubes (PMTs). This signal is called S_1 . The second signal is caused by ionization electrons that do not recombine with ions. These are drifted up by an electric field and extracted by a second, stronger field into the gas phase, where secondary scintillation (S_2) is caused and measured by the same PMTs. The drift time between these signals is proportional to the interaction depth (z). In addition to this, the ratio of S_2/S_1 provides a powerful discrimination between electronic and nuclear recoils. In large-scale TPCs, such as XENON100 [80] and LUX [81], the light distribution of the S_2 in the PMTs gives the coordinates in the plane of the PMTs, so that a three-dimensional resolution is obtained.

This article has the following structure. In section 3.3, the XAMS setup and the TPC are introduced. Section 3.4 discusses the data processing and gives results based on the main S_1 - and S_2 -signals. In section 3.5, S_2 -signals from single electrons are analyzed and a PMT calibration technique based on these signals is presented. In section 3.6, we give a summary of the analyses in these sections.

3.3 THE XAMS SETUP

3.3.1 *The XAMS TPC*

The XAMS TPC features a cylindrically-shaped active volume of 154 cm^3 , which holds 434 g of liquid xenon at a temperature of $-90\text{ }^\circ\text{C}$, as shown in Fig. 3.1 [82–84].

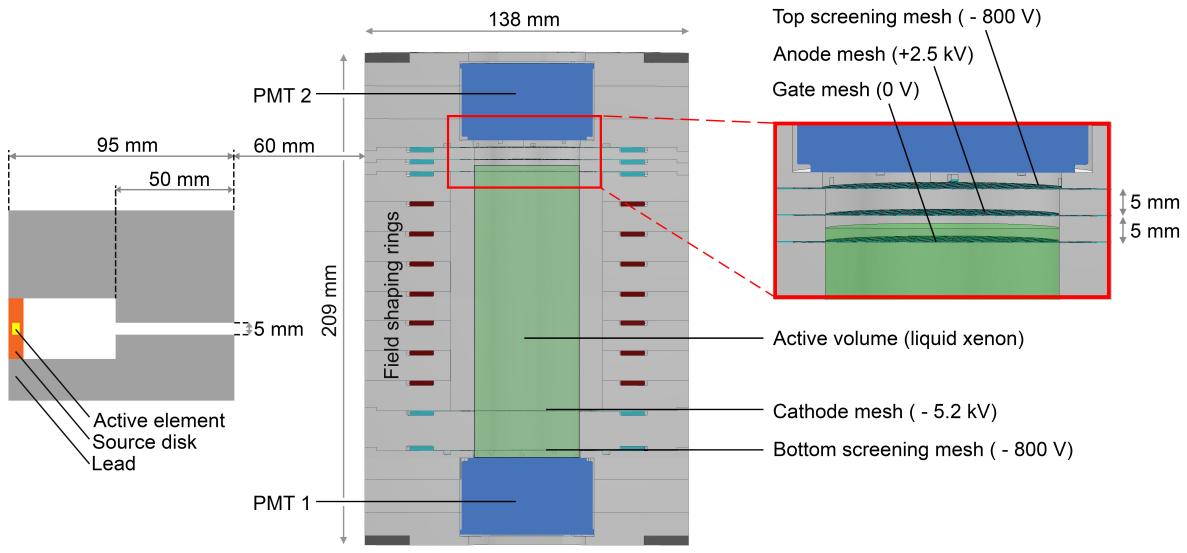


Figure 3.1: Cross-section of the XAMS TPC and the source in the collimator (drawn to scale). All elements of the TPC are contained in a cylindrical PTFE structure made from stackable disks, as indicated by the gray color. The electric field is defined by five meshes and seven copper rings, serving to homogenize the drift field. The top and bottom screening meshes are held at the cathode potential of the PMTs. The active volume is defined by the cylindrical volume between the gate and cathode mesh, measuring 100 mm (height) \times 44 mm (diameter). The ^{22}Na gamma source (described in section 3.3.3) is mounted in a collimator that is made of two lead blocks with cylindrical holes. These are positioned on the outside of the outer vessel (vessels not shown) and allow for a beam size of 11 mm at the closest edge of the active volume. A two-inch NaI(Tl)-detector (not depicted) is used for triggering, and is positioned 100 mm to the left edge of the collimator. The z-position of the collimator is adjustable.

There are two PMTs, one at the top and one at the bottom, that view the active volume and record the S₁- and the S₂-signals. Five meshes define the electric field: the drift field of 0.52 kV/cm is between the cathode and the grounded gate mesh, whereas the extraction field is between the gate and anode mesh, where a voltage of 2.5 kV is applied over 5 mm. The meshes were made by chemical etching of a 150 μm thick stainless steel sheet. They have a square pattern with a pitch of 2.45 mm and a wire thickness of 150 μm , giving a head-on optical transparency of 88 %. The drift field is shaped by a series of copper rings connected to a resistor chain between the cathode and gate mesh. Two additional meshes shield the PMTs from the TPC's electric fields. The distance between the cathode and the gate mesh, which defines the maximum drift length, is 100 mm.

The PMTs are circular two-inch UV-sensitive low-temperature Hamamatsu PMTs of type R6041-406. The low transit-time spread of 0.75 ns in combination with the fast 500 MSa/s digitizer type CAEN V1730D makes XAMS well-suited for fast-timing applications, such as pulse-shape discrimination studies.

3.3.2 Cryogenics and gas system

For the successful operation of a dual-phase xenon TPC, a cryogenic cooling system is required in combination with a purification and storage system. The piping and instrumentation diagram of the XAMS setup is included in A.2.

The cryogenic part of the system consists of double-walled stainless steel vessels. The insulation volume between the vessels is continuously pumped out during normal operation, and pressures of 3×10^{-7} mbar are reached. In addition, aluminum-coated Mylar foil is inserted in the insulation volume to shield from radiative heat transfer. The cooling is provided by an Iwatani PDCo8 pulse tube refrigerator (PTR), which gives an effective cooling power of (22 ± 2) W at -90°C . We apply the cooling to a copper cold finger, where the xenon condenses and droplets fall down into a funnel leading into the TPC. A resistive heating band wrapped around the cold finger enables us to regulate the temperature. The current to the heating band is controlled by a PID controller based on the temperature read by a Pt100 temperature sensor at the cold finger.

A cooling power failure may result in a rising pressure in the TPC. A burst valve with a pressure limit of ~ 4.0 bar is connected to the inner volume to ensure no higher pressure can build. We provide emergency cooling with a pressurized liquid nitrogen dewar, with the flow controlled by a solenoid valve that is switched by a pressure sensor. In addition, text and email warning messages are automatically sent in case of abnormal behavior of the system. The pressure sensor, the solenoid valve and the computer that sends the messages are powered by an uninterruptible power supply.

The required xenon purity level is achieved by continuous circulation through a high-temperature SAES MonoTorr PS3-MT3-R-2 getter with a maximum flow rate of 5 standard liters per minute. We use a heat exchanger at the cryogenic part of the system to achieve this flow rate with only modest cooling power. In section 3.4.2, we show that we achieved an impurity level of (1.2 ± 0.1) ppb (oxygen-equivalent).

We use an EMP MX-808ST-S diaphragm pump to establish the flow in the recirculation circuit. The flow is controlled with a needle valve and measured with a thermal mass flow meter. For the measurements described in this work, no buffer

volumes were installed at the inlet or outlet of the pump, causing oscillatory behavior in the flow. The presumed effect on the measurements is described in detail in section 3.4.3. We recognize this as a design flaw, which we have since adjusted by installing gas bottles as buffer volumes in the system.

The liquid level in the TPC is monitored by a stainless steel cylindrical capacitive level meter, which is read out by a custom-programmed Arduino board. The flow control of the needle valve is used to set the liquid level, as we noticed that the liquid level decreased as we increased the flow rate. We assume that this effect is due to a changing thermal equilibrium in the heat exchanger, where a nonnegligible amount of liquid xenon is kept.

The total xenon content in the XAMS setup is roughly 6 kg, most of which surrounds the PTFE structure of the TPC. The time required to fill the TPC, limited by the maximum cooling power of the PTR, is roughly 10 hours. We perform recuperation by immersing gas bottles into liquid nitrogen dewars and allowing gas to deposit on the walls of the cylinder. The time for a full recuperation is roughly 8 hours.

3.3.3 Trigger and DAQ

We use a ^{22}Na gamma source with an activity of $(368 \pm 11) \text{ kBq}$ to perform our studies. The source is mounted in a lead collimator (see Fig. 3.1) on the outside of the insulation vacuum vessel, with an opening angle of 2.9° such that the beam has a width of 11 mm at the closest edge of the active volume. The direction of the beam is horizontal, giving lateral irradiation of the TPC. We change the z-position of the collimator by varying the height of the platform on which the collimator is mounted. To reach the active volume, the gamma rays have to cross the walls of the inner and outer vessels, a thin layer of liquid xenon and the PTFE holding structure of the TPC, so that the total material traversed is 6 mm of stainless steel, 2 mm of liquid xenon and 46 mm of PTFE, respectively.

^{22}Na decays by positron emission (branching ratio 90.4 %) or electron capture (branching ratio 9.6 %). The decay is almost always followed by the emission of a 1274 keV gamma ray from its ^{22}Ne daughter. In the case of positron emission, two additional back-to-back gamma rays of 511 keV are produced from positron annihilation. By using thallium-doped sodium iodide ($\text{NaI}(\text{Tl})$) as a coincidence detector that measures one of the 511 keV gamma rays, the other 511 keV gamma ray going directly toward the active volume is tagged. This increases the fraction of events where all the energy is absorbed, since the number of events where gamma

rays enter the active volume after Compton scattering on the material surrounding the detector is reduced.

The trigger is based on a threefold coincidence of the two PMTs in the TPC and the external NaI(Tl) detector. If the trigger condition is satisfied, all three channels are digitized by a CAEN V1730D digitizer board. This board has 8 channels that are digitized with a time resolution of 2 ns and a voltage resolution of 14 bits, distributed over a dynamic range of 2 V. We choose an event window of 163 μ s: more than twice as long as the maximum drift time of 60 μ s. We place the trigger position in the middle, such that an (accidental) trigger on an S₂-signal will always contain the S₁ in the same window. A cut in post-processing ensures that there was a true coincidence with the S₁ and the external NaI(Tl) (and not, for example, a coincidence with the S₂-signal and an uncorrelated interaction in the NaI(Tl) crystal).

The simple coincidence means that all three channels must exceed the threshold *at the same time*; no coincidence window was used. The time offset between the two PMTs in the TPC is negligible, however, the start of the peak of the NaI(Tl) detector output was shown to occur (22 ± 6) ns later than that of the PMT signals. The trigger condition was therefore satisfied only if both PMT signals were still above threshold at this time after the peak amplitude. In the case of high energy recoils, the pulses are sufficiently large and this causes no problems. However, for low energy recoils we observe a low trigger efficiency, which we identify in the comparison to Monte Carlo simulation in section 3.4.3 at energies below 150 keV.

3.4 DATA REDUCTION AND RESULTS

3.4.1 Peak finding, clustering, identification

The data of each event consist of the waveforms of the two PMTs with a duration of 163 μ s (Fig. 3.2). The data processor, which is the same software developed for XENON1T [85], analyzes the waveforms in each individual PMT channel by looking for significant excursions above the baseline. These are called *hits*. In XENON100 and XENON1T, *zero-length encoding* is used: the data consists of small chunks of data around a significant excursion from the baseline, so that the baseline is suppressed and the data volume is reduced [86]. In order to be compatible with this structure, we apply a software zero-length encoding with a very low threshold. The hitfinder threshold is dynamically determined as 4.5 times the standard deviation of the noise in the first 40 digitizer samples (80 ns) of the zero-length

encoded chunk containing the hits. The hits from both channels are then clustered into *peaks* based on the gap between the edges of the hits: if this exceeds 450 ns, the hits are clustered into separate peaks. The area and the width of the peak are computed based on the summed gain-corrected waveform properties. The width metric uses the range containing 50 % of the peak area with 25 % on either side. The peak position is defined as the amplitude-weighted mean time of the samples in the peak.

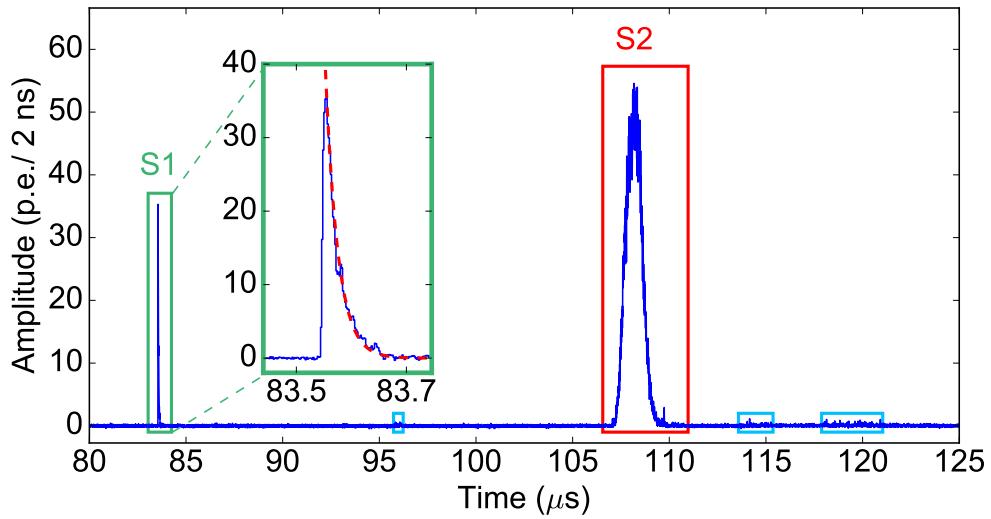


Figure 3.2: Typical gamma-ray-induced sum-signal of the two PMT channels, showing the S1 (green box) due to prompt scintillation and the S2 (red box), delayed by the drift time. The inset shows a detailed view of the S1-signal and an exponential fit to the falling slope of the S1, with a decay time of (22.8 ± 0.1) ns. Additional peaks are found (in the blue boxes), mostly happening after the S2. Details of this kind of signal can be found in Fig. 3.7. The data processing software finds the hits in each channel, clusters them into peaks, determines peak properties and classifies each peak based on the width and area.

As seen in Fig. 3.2, the main signals are the S1- (highly peaked signal at 84 μs) and the S2-signal (the broad signal at 108 μs). After the S2, some peaks with low area and high width are found (shown in blue boxes). These signals are due to secondary emission of electrons caused by photo-ionization of S2 UV photons and drifted up to produce an additional, much smaller S2. These signals will be discussed in section 3.5. All peaks are classified as either 'S1', 'S2' or 'other' based on their width and area.

Table 3.1: Data selection cuts and the number and fraction of events surviving each cut. The cuts are applied successively.

Cut	Events	Fraction
No cuts	215 831	100.0 %
At least one S ₁	205 417	95.2 %
Only one S ₁	166 353	77.1 %
At least one S ₂	158 586	73.5 %
Only one S ₂	115 005	53.3 %
Coincidence S ₁ and NaI(Tl)	105 062	48.7 %
NaI(Tl) < 600 keV	101 381	47.0 %
Total	101 381	47.0 %

3.4.2 S₁ and S₂ corrections

After data processing, the following selection criteria are applied to the events. First of all, only events with a single S₁ and S₂ are kept. This cut rejects pileup events, double scatter events (which cause two S₂-signals), or events where no S₂ is generated (for instance, where the interaction occurs below the cathode mesh). Events where the S₁ is not in coincidence (difference of peak center position less than 200 ns) with a signal in the NaI(Tl)-crystal are also cut. In addition, the energy deposition in the NaI(Tl)-crystal is required to be less than 600 keV, so that the triggers on 1274 keV gamma ray are cut. We impose no lower bound other than the trigger threshold on the NaI(Tl) energy, so that we keep events where the 511 keV Compton scattered in the NaI(Tl)-crystal. A summary of all the cuts and the number of events surviving each successive cut is given in table 3.1. The fraction of events surviving all cuts for the analysis presented here is 47.0 %. Most events cut are due to multiple S₁s or S₂s.

For both the sum-signals of the S₁- and the S₂-signals, the area of the peak is proportional to the recoil energy. However, the response to a mono-energetic energy deposition is not uniform throughout the TPC, requiring spatial corrections. Since the XAMS TPC has only two PMTs, the position in the x,y-plane cannot be determined, but the z-coordinate is calculated based on the drift time that is defined by the difference of the weighted mean times of the S₂ and the S₁.

A z-dependent scale factor is applied to the S₁-signal to eliminate differences in light detection efficiency (LDE). The amount of light detected by the PMTs for different interaction positions depends on optical properties of the TPC, such as the

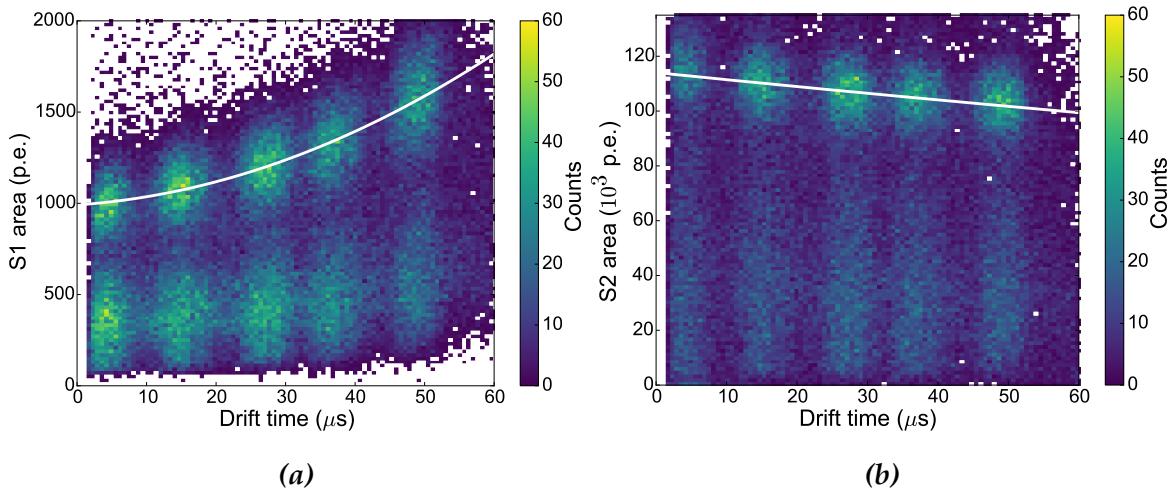


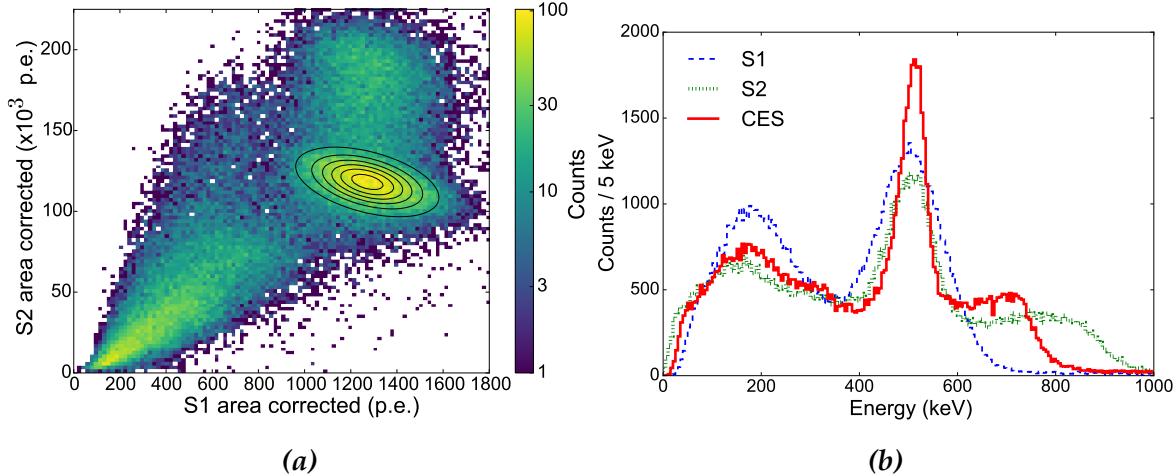
Figure 3.3: Density plot of the area of the sum-signal of the S1 (a) and the S2 (b) signal for different z-positions in the TPC, corresponding to different drift times. The data shown in these figures were taken with the collimated source pointing at five different positions in the TPC. The thick white lines are fits to the photo-peak. For the S1, an overall increase is found due to LDE effects, a second degree polynomial fit gives the correction. For the S2, an exponential fit provides a correction for loss of electrons during the drift time.

reflection properties of the walls of the TPC, optical transparency of the meshes and reflection on the liquid-to-gas interface. The secondary scintillation light of the S2-signal is always produced in the small region between the liquid-to-gas interface, so no z-correction for LDE has to be applied. However, the number of electrons that create the S2-signal decreases with increasing drift time due to attachment of electrons to impurities in the xenon. Assuming n_0 electrons are produced at the interaction position, the number of electrons n_e left after a drift time t_d can be calculated with

$$n_e = n_0 \exp(-t_d/\tau_e), \quad (3.1)$$

where τ_e is the *electron lifetime*, which is an indirect measure of the purity of the xenon.

Five datasets were taken with a different z-position of the collimator. Fig. 3.3 shows the area of the sum-signal of the S1- and the S2-signal for all datasets, each containing a prominent peak at high energy and a broad shoulder for lower energies. The former is attributed to the full absorption peak (mostly due to photoelectric absorption, or multiple scatter events where the S2s are too close together to be separated), whereas the latter is due to Compton-scatter events.



3 *Figure 3.4: (a): Density plot of the area of the S1 and the area of the S2 in the same event for a 511 keV gamma-ray source. The ellipse shows the anti-correlation between the S1- and the S2-signal at the expected photo-peak. A superior energy resolution is found by fitting the photo-peak and projecting along the short axis of the ellipse. The shoulder at low energy is due to Compton-scatter events; the second peak at higher S2 area than the photopeak is discussed in the text. (b): The spectra using the S1, S2 and the combined signal. The energy resolution at 511 keV improves from $(14.5 \pm 0.2)\%$ and $(10.8 \pm 0.4)\%$, respectively, for the S1- and the S2-signal alone to $(5.8 \pm 0.2)\%$ for the combined spectrum.*

For the S1, uncertainties on optical parameters limit the use of a detailed LDE model. We therefore use a data-driven approach, modeling the correction function as a second degree polynomial. We determine this function in two steps. We first fit a Gaussian function to the photopeak for several slices in drift time, and then fit the photopeak position as a function of drift time with a second-degree polynomial. This polynomial function, shown by the white line in Fig. 3.3a, provides the correction factor for the LDE. The average value of the fit function, which gives the volume-averaged light yield for 511 keV gamma rays, is $(1.29 \pm 0.07) \times 10^3$ p.e., or (2.5 ± 0.1) p.e./keV in this configuration. This is equivalent to (5.6 ± 0.3) p.e./keV at zero field and 122 keV using data from NEST [87], which is comparable to TPCs like XENON100 (4.3 p.e./keV) and LUX (8.8 p.e./keV) [80, 88]. An overall increase of LDE with drift time is found, since most of the scintillation light is detected by the bottom PMT.

For the S2, the correction function is expected to be an exponential (see equation 3.1). The electron lifetime as determined from the fit is $(429 \pm 26)\mu\text{s}$, similar to the average lifetime of 514 μs during the year-long science run of XENON100 [89], and was achieved in only 7 days of continuous purification with the high-

temperature getter. We observed that the electron lifetime rapidly increases in the first 6 days, but levels off after this [83]. For XAMS, this electron lifetime means that even at the maximum drift time, only 13 % of the S₂-signal is lost. Using the values in [90], this electron lifetime corresponds to an impurity level of (1.2 ± 0.1) ppb (oxygen-equivalent). We kept the recirculation flow rate constant over the full duration of all measurements described here (one day).

3.4.3 Energy calibration

After the corrections for the S₁- and the S₂-signal have been applied, the absorbed energy can be determined. Both signals provide a measurement of the deposited energy, since the area of the S₁-signal is proportional to the number of photons produced in the interaction and the area of the S₂-signal is proportional to the number of electrons produced. The total energy deposited in these events is always identical: the ionization and scintillation signals are therefore anti-correlated. In Fig. 3.4a, this anti-correlation is clearly visible as the ellipse with a downward slope. The best energy resolution is achieved by using a projection along the short axis of the ellipse, which is known as the combined energy scale (CES) [91, 92]. We use the same projection for all energies, which is a good approximation for energies greater than roughly 100 keV [80]. In Fig. 3.4b, the spectra obtained from the S₁, the S₂ and the CES are shown. The energy resolutions, as defined by σ_E/E for a Gaussian fit, are (14.5 ± 0.2) %, (10.8 ± 0.4) % and (5.8 ± 0.2) %, respectively.

High-S₂ population

In addition to the photopeak, a second peak at the same S₁-area but larger S₂-area was found, see Fig. 3.4a. We also find this effect for the Compton-scatter events, and throughout all datasets. The appearance of the high-S₂ events is highly correlated in time, with a frequency of (0.110 ± 0.006) Hz, i.e., roughly a 9 s period (see Fig. 3.6a).

The cause of a varying S₂ size can be related to only few parameters. Since the S₁-signal is unaffected, the PMT gain, cathode voltage, DAQ problems or processing errors can be excluded. Possible detector parameters changing the S₂ size, but not the S₁ size, are the xenon purity, the anode voltage, and the liquid level. The anode voltage was not monitored by the slow control system, but the display showed a stability of better than 1 V. Unfortunately, we cannot correlate the detector parameters monitored by the slow control to the time behavior found in the high-S₂ population appearance. This is because the variables were read out only every two

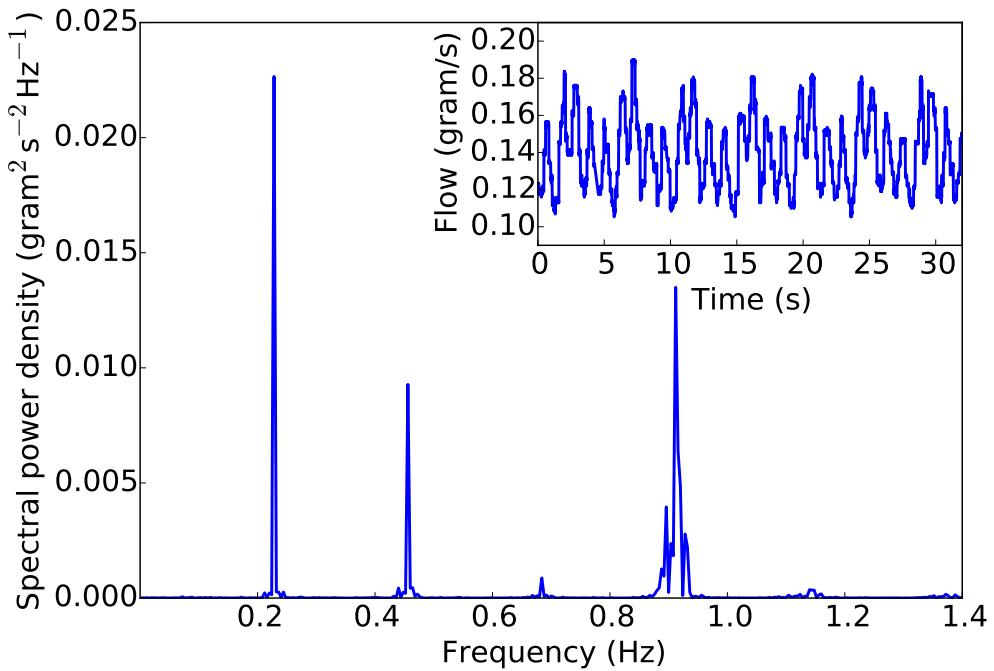


Figure 3.5: Power spectrum of the flow rate as measured in a test where nitrogen gas was pumped through the detector volume. Shown in the inset is the flow as a function of time for the first 30 seconds of this measurement. A clear peak at a frequency of (0.228 ± 0.004) Hz is visible, along with several harmonics of this frequency.

minutes; a decision that was taken because the readout of temperature sensors in the TPC caused noise in the PMT signals by electronic pick-up.

A plausible explanation found is a time-varying liquid level in the TPC. The S_2 size is highly dependent on this, so that only a small change in liquid level can still give significant effects. Alternatively, there could be ripples on the liquid surface, appearing every 9 s. One of the mechanisms that could cause either a changed liquid level or ripples on the surface is related to the recirculation flow. During the measurements, we observed that the gas flow rate in the recirculation system was constantly varying. To investigate this effect further, we did a test where nitrogen gas was pumped through the system. We observed a highly periodic behavior of the flow rate, with a period depending on recirculation speed. Fig. 3.5 shows the flow rate for a mass flow similar to the flow used during the measurements with liquid xenon. The typical frequencies found in these tests are higher than the (0.110 ± 0.006) Hz found in the data, but it should be noted that the systems with liquid xenon and with nitrogen gas are not equivalent, and that the frequency found in the nitrogen gas tests depends on the pressure and the recirculation flow. The reason for this oscillatory behavior is related to the absence of a buffer volume

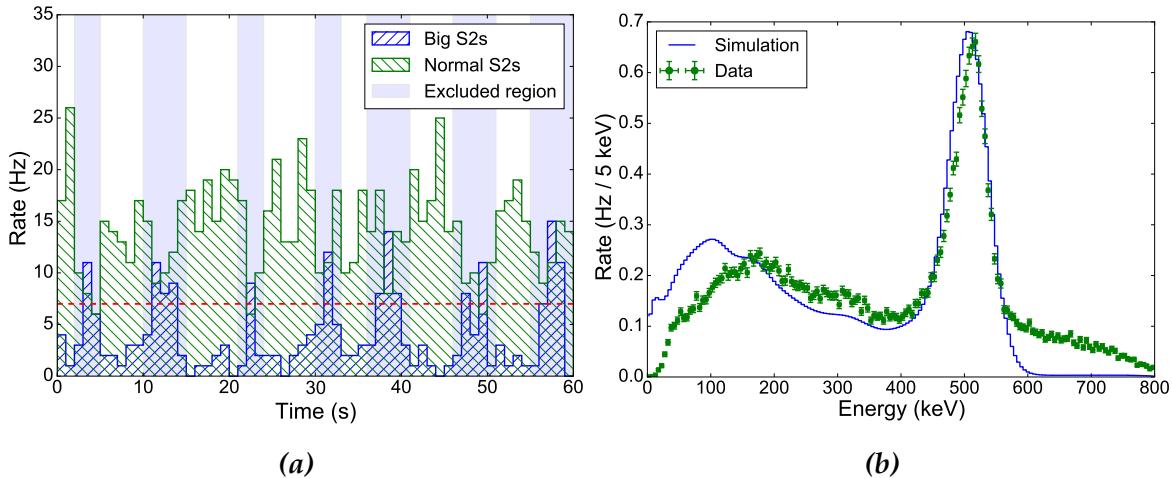


Figure 3.6: (a): The rate of events with an S_2 larger than 150 000 p.e. (blue) and the rate of other events (green) for the first 60 seconds of the dataset. A clear time-correlation is visible. The shaded regions show the times that are cut. (b): The CES spectrum from data after applying the time cut (green points) compared to a smeared spectrum from a Monte Carlo simulation (blue line). The data points are scaled to reflect the rate before any cuts. The mismatch between simulation and data at low energy is due to a decreased trigger efficiency, as described in section 3.3.3. At high energy, this is due to the partly cut high- S_2 population described in section 3.4.3.

at the recirculation pump. Buffer volumes were installed, and subsequent tests showed a significant increase in the stability of the flow rate. Future measurements with liquid xenon will show if the effect is related to the instability of the flow rate.

As illustrated in Fig. 3.6a, we can use a time cut to remove a large fraction of the events with a large S_2 -signal. Whenever more than six events with an S_2 size larger than 150 000 p.e. are found within one second, the events from one second before to one second after this bin are cut. This removes 41.1 % of all events passing previous cuts.

Comparison to simulation

The resulting spectrum was compared to a GEANT4 [93] Monte Carlo simulation, where the energy deposition was registered when there is a simultaneous energy deposition in the NaI(Tl) crystal and the liquid xenon active volume. The result was then smeared with an energy resolution function according to

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}}, \quad (3.2)$$

where a is fixed by the requirement that $\sigma_E/E = 5.8\%$ at 511 keV. The comparison is shown in Fig. 3.6b, where the green points are from data and the blue line is from simulation. The data points are scaled to the total rate observed before any cuts of 26.6 Hz, which agrees well with the rate of 25.9 Hz from simulation. The contribution at S2 sizes larger than 600 keV is still visible.

At energies below ~ 150 keV, the simulation predicts a higher rate than observed in measurement. This difference is due to a timing offset between the NaI(Tl) and the S1-signal, which causes a trigger on the falling edge of the S1 instead of on its peak amplitude and deteriorates the trigger efficiency for low energy recoils, as described in section 3.3.3.

3.5 SINGLE-ELECTRON S2-SIGNALS

A distinct signal that is found in dual-phase xenon TPCs is that of S2-signals produced by single electrons [94–96]. The scintillation light of xenon, at 178 nm, can liberate electrons in the TPC. In general, the electrons come from impurities in the xenon that have a low ionization potential, such as O^- ions, or from exposed metallic surfaces. If these electrons are somewhere in the active volume, they will in turn drift upward and create very small S2-signals. Since the main S2-signal is the dominant source of UV photons in the TPC, it causes the large majority of single-electron peaks.

Fig. 3.7 shows an example of a single-electron signal found in the data. As described in section 3.4.1, the waveform is cut into small sections analogous to zero-length encoded data, so that the hitfinder threshold is dynamically updated based on the local noise level. The blue and green parts of the waveform show the hits that are found, when a threshold of 4.5 times the standard deviation of the baseline noise is crossed (indicated by the dashed lines). The width of the signal is around 1 μ s, comparable to normal S2-signals.

Fig. 3.8 shows the time distribution of the peak position relative to the position of the S2 for all candidate single-electron S2-signals, namely all peaks that are not classified as S1 or S2 and have a coincidence of both PMTs. A large fraction of the peaks occurs between 0 μ s to 60 μ s (as defined by the maximum drift time). We observe a clear increase at 60 μ s, which is due to the S2 light impinging on the cathode mesh, where electrons are liberated relatively easily due to the low ionization potential of the iron in the stainless steel. Before the S2 ($\Delta t < 0$), as well as after the full drift length ($\Delta t > 60 \mu$ s), there is a nonzero contribution, which is partly due to noise hits clustered into peaks, but partly shows the same properties

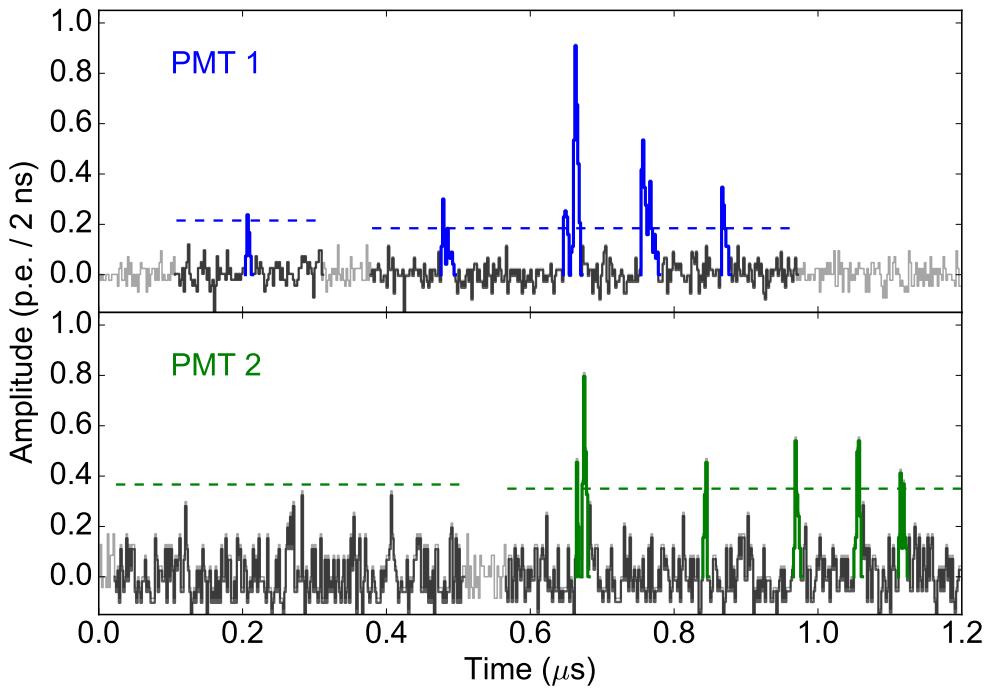


Figure 3.7: Example of a single-electron S₂-signal, shown for both PMTs separately. The data is cut into small pieces based on the crossing of a very low threshold. This is indicated by the dark gray part of the waveforms. The noise level is determined on the first 40 samples of these pieces, yielding a hitfinder threshold of 4.5 times the standard deviation of the noise (indicated by the dashed lines). When the waveform crossed this threshold, a hit is found, indicated by the blue and green waveforms.

as the single-electron signals in the drift region. These peaks can be caused by a delayed extraction phenomenon, as discussed in [95] and [96].

Single-electron S₂s can be effectively used as ‘calibration sources’: the detector response to just one electron can be probed in this way. This enables the direct determination of various parameters, such as the secondary scintillation gain. In addition, these signals can be used for a PMT gain calibration, since they consist of single-photoelectron hits.

3.5.1 Gain calibration

The PMT gain is defined as the average number of electrons at the anode responding to one electron emerging from the photocathode. Often PMT calibrations are done using external pulsed light sources. Although such calibration provides a direct and usually accurate gain calibration, it requires an interruption during dark

matter data taking. In addition, a dedicated LED calibration system and calibration measurements are necessary. Finally, the LED calibration is usually performed at a higher wavelength than the xenon scintillation light of 178 nm, because it is technically challenging to guide UV light through an optical fiber. This makes it impossible to probe effects like double-photoelectron emission, which occurs only at short wavelengths [71].

In this section, we discuss a method to use single-electron peaks for PMT gain calibration. These consist of well-separated single-photon hits and are abundant in all data, so they can be used to measure the PMT gain continuously. The LUX collaboration already uses single-electron signals as part of their gain calibration, which operates on different principles [70].

Hit data selection

Single-electron S2s typically have the same width as ordinary S2-signals (about 1 μ s wide), but consist of a small number of photoelectrons. For example, for XENON100, these signals consists of roughly 20 photoelectrons [96].

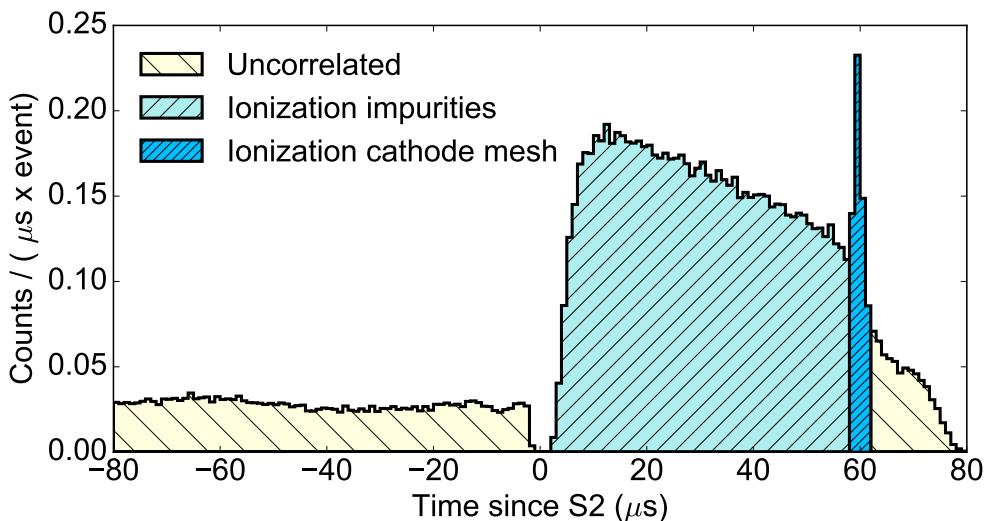


Figure 3.8: Distribution of peak positions relative to the position of the S2. A large fraction of the peaks occurs between 0 μ s to 61 μ s, which is expected for single-electron signals that are caused by the S2 light. The large peak at 60 μ s is due to the cathode mesh at this drift length: electrons are easily liberated from the stainless steel.

Since the hits in single-electron S2s are spread out over a relatively long duration, the PMT hits of the detected photons can be found individually (see Fig. 3.7). This means that the single-electron S2s provide a source of single photoelectron hits,

which can be used for an *in-situ* gain calibration.

We apply cuts on the event, peak and hit level to select proper single-photon hits in proper single-electron signals. Events are selected by the same criteria as in section 3.4. For the peaks, defined as clusters of hits, we introduce the following cuts. Since single-electron S2s are primarily caused by S2 photons, only peaks that occur within 5 μ s to 60 μ s after the S2 are selected. The lower bound ensures no fragments from accidentally split S2s are selected, and the upper bound cuts peaks beyond the maximum drift time. Both PMTs are required to contribute to the peak, to suppress peaks consisting of noise and dark current hits. In order to reduce contamination from common-mode noise clustering, we apply a cut on the average area of the hits in a peak: since noise hits on average have a smaller area than particle-induced PMT hits, a cluster of noise hits will have a low average area.

Finally, at the hit level, we do not consider hits with an extremely small width, indicative of noise hits rather than real PMT hits. The width parameter used here is the sum absolute deviation (SAD), given by

$$\text{SAD} = \sum_i \frac{A_i}{A_{\text{tot}}} |t_i - t_c|, \quad (3.3)$$

where i runs over all samples in the hit, t_c is the amplitude-weighted mean time of the hit, and A denotes the area. This parameter takes continuous values greater than or equal to zero, which has the advantage of discriminating different widths even if this is at the same order as the sampling time. Typical values for a single-photoelectron hit are about 3.5 ns for the XAMS PMTs. We cut hits with an SAD less than 0.5 ns, which mostly consist of hits that are just one sample wide (such that $\text{SAD} = 0$ exactly) or where the hit is two samples wide but the area is dominated by just one sample.

With the above selection of hits we proceed to calibrate the gain of each PMT. For gain calibration, the parameter of interest is the *area* of the hits (given by $\int V dt$), which is proportional to the number of electrons n_e at the PMT output according to

$$n_e = Q/q_e = \frac{1}{q_e} \int I dt = \frac{1}{q_e R} \int V dt, \quad (3.4)$$

where Q is the charge at anode, q_e is the charge of the electron and where R denotes the termination resistance of the digitizer (50 Ω). Because of this relation, the area of a hit can be expressed as number of electrons equivalent area. If the hit results from one photoelectron, the average number of electrons at the output is simply

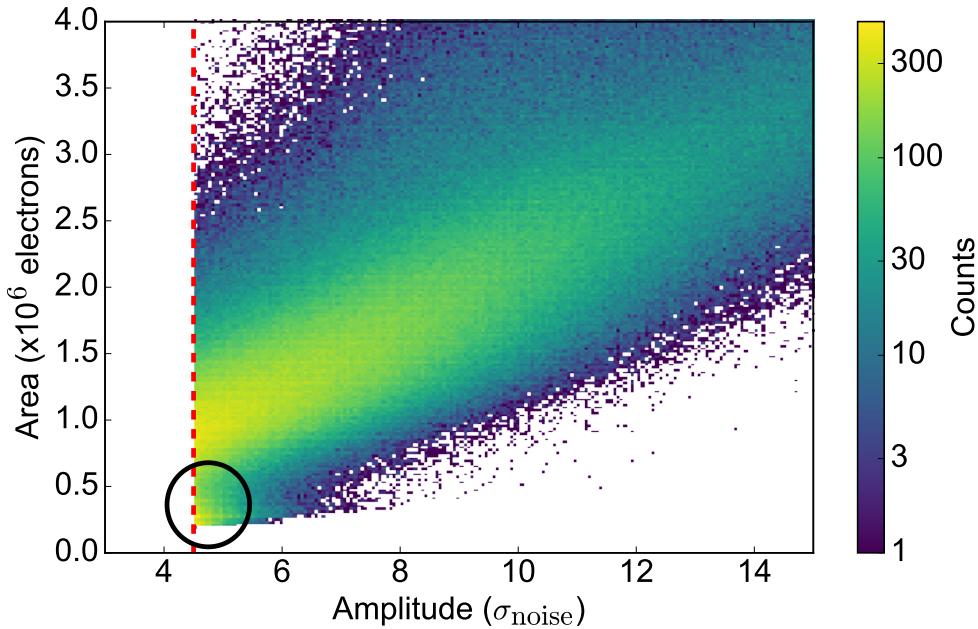


Figure 3.9: The distribution of selected single-electron hits in amplitude and area. There is a clear correlation between the hit area and amplitude. The distribution is cut at $4.5\sigma_{\text{noise}}$ as defined by the hitfinding threshold. A correction for hits below this threshold is calculated and used in the analysis. At small area and low amplitude, a second band coming from noise hits can be seen (indicated by the black circle).

$\overline{n_e} = G$, and the gain G can be computed.

Acceptance correction

The hitfinding algorithm preferentially detects high-area hits, since these are more likely to exceed the hitfinding amplitude threshold. Fig. 3.9 shows the correlation between the amplitude, measured in units of σ_{noise} , and the area of the hit. The distribution is sharply cut at $4.5\sigma_{\text{noise}}$, the hitfinder threshold. This was chosen to limit the contribution of noise hits, which are visible in the bottom left corner of Fig. 3.9.

To correct for this loss of hits below the threshold, we must estimate the acceptance ϵ of the hitfinder, i.e. the fraction of photon hits found by the hitfinder, as a function of the hit area. This function can be estimated by studying the amplitude distribution for a sample of hits with similar area (equal up to a difference of 0.1×10^6 electrons). This distribution is cut at the hitfinder threshold, but since it is well-described by a Gaussian distribution above the threshold, we will assume that it follows a Gaussian function also below this threshold. By evaluating the fraction

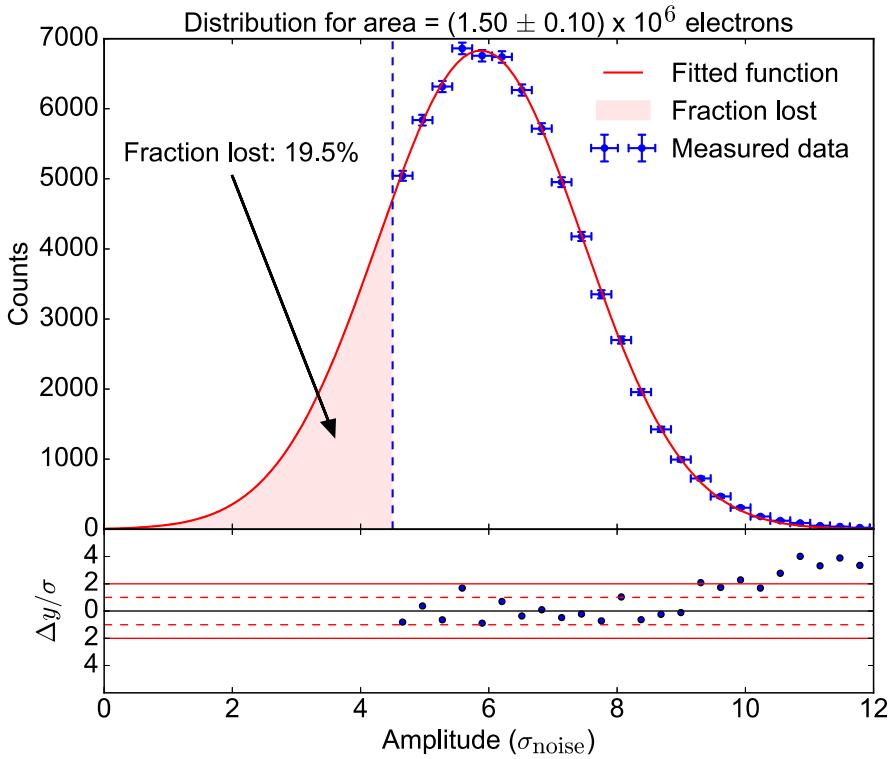


Figure 3.10: Example of a fit used to determine the acceptance for an area slice. The fraction of the hits that is not found is inferred from the area under the fit in the region of the function that extends below the threshold. In this example, this fraction is 19.5 %, so the acceptance is 80.5 %. The bottom panel shows the deviation from this fit in units of the error on the points. The dashed and solid lines indicate a deviation of 1σ and 2σ , respectively. Up to a threshold of $9\sigma_{\text{noise}}$, the distribution is well described by the fit.

of the area under the fitted distribution below the threshold, the acceptance can be calculated, as illustrated in Fig. 3.10.

The distribution of hits and Gaussian fits are shown for all area slices in Fig. 3.11. For low-area hits, only a tail of the Gaussian distribution exceeds the hitfinding threshold of $4.5\sigma_{\text{noise}}$, making it difficult to fit the distribution. We instead infer the parameters μ and σ by extrapolation. Since the shape of PMT hits is to a good approximation independent of the area, the mean μ is extended linearly to zero [82]. We assume that the standard deviation σ is constant at low area, since this should be dominated by baseline noise on the highest bin and is therefore independent of the hit area.

With the amplitude distribution specified by $\mu(A)$ and $\sigma(A)$, the acceptance for every area and for different hitfinding thresholds can be computed. For a hitfinding

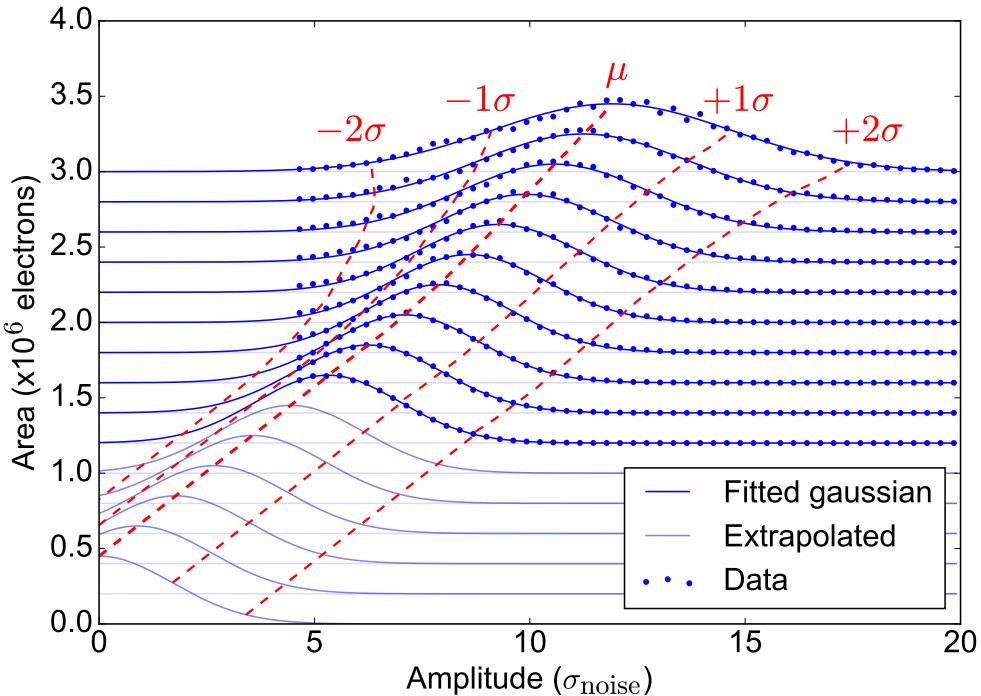


Figure 3.11: Stacked hit amplitude histograms for each area slice (blue points), together with Gaussian fits (blue lines). The data as well as the fits are scaled such that the maximum amplitude of all distributions is the same. Red dashed lines indicate the mean and standard deviations of the fitted Gaussians. For low area slices, the amplitude distribution is estimated by extrapolating the mean and standard deviations found in higher-area slices as described in the text.

threshold of $n_{tr}\sigma$, the acceptance ϵ as a function of peak area A is

$$\epsilon(A) = \int_{n_{tr}\sigma}^{\infty} g(x; \mu(A), \sigma(A)) dx \quad (3.5)$$

where g is a normalized Gaussian distribution and x denotes the amplitude. To correct the area spectrum, it is divided by $\epsilon(A)$.

Although a hitfinder threshold that is as low as possible is desired for determining the acceptance, it is not necessarily ideal for determining the gain. This is because low-amplitude noise hits are too dominant in the area spectrum, so that the gain will be underestimated. We therefore use a higher amplitude threshold to remove any possible bias due to noise hits, which we compensate by a corresponding change in the acceptance function. In Fig. 3.12, the uncorrected and corrected area spectrum for a hitfinding threshold of $6.5\sigma_{\text{noise}}$ is shown, together with the acceptance function for this threshold. A clear peak is visible, which is fit in the

area around the peak to determine the gain. The fit is limited to a part around the maximum; at low area, the noise contribution becomes dominant, while at high area the contribution from two-photoelectron hits cannot be excluded. Similar features are found in other PMT calibrations, such as in [97].

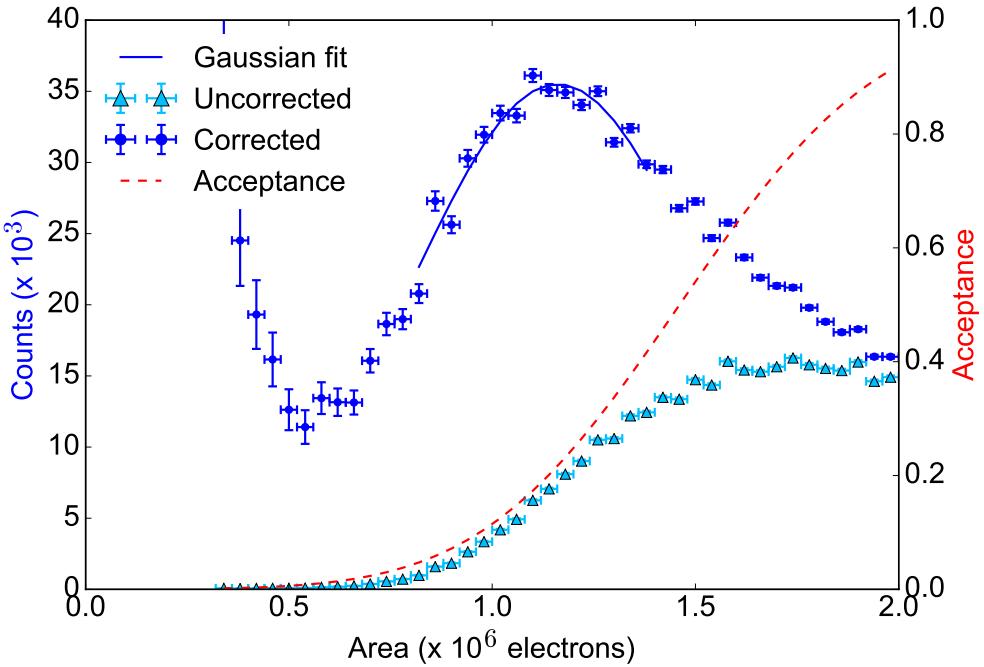


Figure 3.12: The area distribution for the hits with (dark blue points, left scale) and without (light blue points, left scale) correction from the acceptance function (red dashed line, right scale) as determined from the model described in section 3.5.1. A Gaussian fit (blue solid line) is used to determine the gain. At low area, noise hits give a large contribution to the corrected spectrum, since they are highly amplified by the acceptance correction. Points below 0.3×10^6 electrons area, where the acceptance drops below 0.1 %, are omitted from the plot.

In Fig. 3.13, the determined gain is plotted as a function of the hitfinding threshold. For low thresholds, the noise contribution becomes too pronounced to properly fit the spectrum. For higher thresholds, the gain that is determined converges to a final value, which we infer to be the true gain of the PMT. We allow a range around this value from uncertainty on the convergence of the final points, which we estimate to be 5 % for PMT 1 and 10 % for PMT 2. The PMT gains found in this analysis were $(1.30 \pm 0.07) \times 10^6$ for PMT 1 and $(0.71 \pm 0.07) \times 10^6$ for PMT 2; both close to the typical gain of 1.0×10^6 quoted by Hamamatsu for this type of PMT [98]. The error bars in Fig. 3.13 originate from systematic errors on the acceptance function, which we calculate by perturbing the fit parameters $\mu(A)$ and $\sigma(A)$ in

equation 3.5.

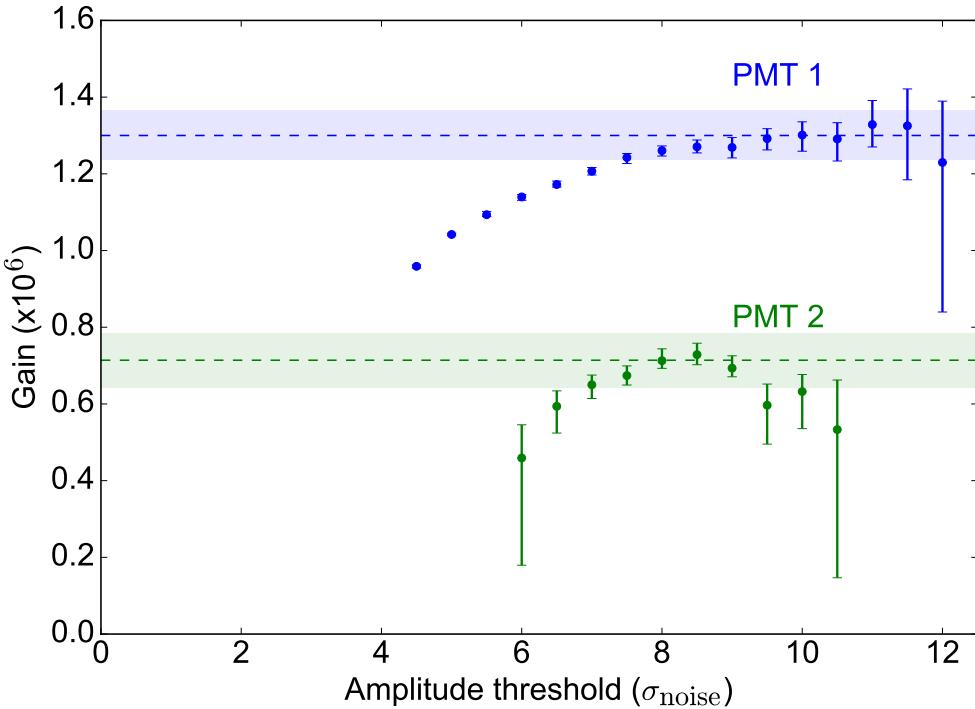


Figure 3.13: The gain for different hitfinding thresholds (in units of standard deviation of the noise), for the two PMTs in XAMS. As the hitfinding thresholds rises far enough above the low-amplitude noise, the gain that is determined levels off to one value, which is the gain of the PMT. The error bars are calculated by allowing varying values for $\mu(A)$ and $\sigma(A)$ for the acceptance function. At high thresholds, the result of the fit is very sensitive to variations in the acceptance function, causing the large error bars. The systematical errors are estimated to be 5 % for PMT 1 and 10 % for PMT 2, as shown with the bands, based on the convergence of the final points.

3.5.2 Discussion

The method used in this work relies on modeling the PMT hits. In particular, the amplitude of hits of a given area is assumed to be normally distributed. For large-area hits, this assumption can be verified since most hits are above the hitfinding threshold, and the distribution follows a Gaussian distribution to a high degree. For smaller areas, a significant part of the distribution is inaccessible and needs to be inferred from the visible part of the distribution (as in Fig. 3.10). Moreover, the parameters μ and σ of the distribution are extrapolated from larger areas where fits

can be made (Fig. 3.11). The validity of the extrapolation can break down at small area, although this will not affect the gain determination if the approximation is valid sufficiently far below the average single-photoelectron hit area.

In the case of low PMT gains, the separation of noise and true PMT hits becomes a serious issue. This is the case for PMT 2 in our analysis, where at high thresholds the errors increase and we eventually fail to fit the distribution because of low statistics. Moreover, we cannot confirm if convergence is reached before this effect starts to dominate. We therefore estimate the systematic errors to be higher for PMT 2 than for PMT 1. It should be noted that these limitations becomes less important if the PMT gain is higher, so that the PMT hits are more separated from the noise.

A PMT calibration requires a source of single photoelectron hits. For single-electron S₂s, a few photoelectron signals are seen in the PMT channels over a time window of typically 1 μ s. There is a finite probability of having two or more PMT signals clustered together into one hit, so that there is a contribution of two-photoelectron hits in the data. The importance of this effect could be different for other TPCs, as it depends on several parameters such as the transient time spread of the PMTs, the sampling time of the ADCs, the width of the S₂ and the anode voltage. These effects will thus need to be studied further if this method is to be used for other TPCs.

Compared to the normal PMT calibration with LED pulsed light, there are some definite advantages to using single-electron S₂s. One of these is they are usually readily available and easily identified in ordinary (energy) calibration or dark matter data; no extra dedicated calibration runs are required. This means that the drift of PMT gains can be monitored on timescales far shorter than with ordinary PMT calibration runs. A second advantage is that the response to the scintillation light is directly probed. The scintillation light of xenon has a wavelength of 178 nm, but since this is technically challenging to provide for a calibration, higher wavelengths are used. For example, XENON100 uses an LED at 470 nm [81]. The method described here makes it possible to study, for example, the possibility of two-photoelectron emission due to one scintillation photon at the photocathode.

3.6 SUMMARY

In this work, the first data of the XAMS TPC were presented. An energy resolution of $(5.8 \pm 0.2)\%$ was achieved at 511 keV. The electron lifetime was found to be $(429 \pm 26)\mu\text{s}$, which is sufficient for this TPC, after only 7 days of purification. An

average light yield of (5.6 ± 0.3) photoelectrons/keV (recalculated to zero field and 122 keV) was found, which is comparable to TPCs like XENON100 and LUX.

A new PMT calibration method based on single-electron S2-signals was explored. Since single-electron S2-signals are very abundant in dual-phase xenon TPCs, this method of PMT calibration can give an important independent cross-check of the normal PMT calibration, with the advantage of superior time resolution and no need for dedicated PMT calibration data.

3.7 ACKNOWLEDGMENTS

This work is part of the research program of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organization for Scientific Research (NWO). We gratefully acknowledge the technical support from to the mechanical, electrical and computing departments at Nikhef.

4

CALIBRATION OF XAMS

In chapter 2, the calibration and correction procedures for liquid xenon TPCs were briefly outlined with a focus on XENON1T. This chapter describes the calibration procedure for XAMS, the setup introduced in chapter 3. The basic analyses presented here provide the light detection efficiency (LDE) correction used for the S_1 , the electron lifetime correction for S_2 , the energy calibration through g_1 and g_2 , the position reconstruction calibration and the double scatter rejection efficiency. These properties provide the corrected S_1 and S_2 and the reconstructed energy, which are used by the more complex analyses that are shown in chapters 5 and 6.

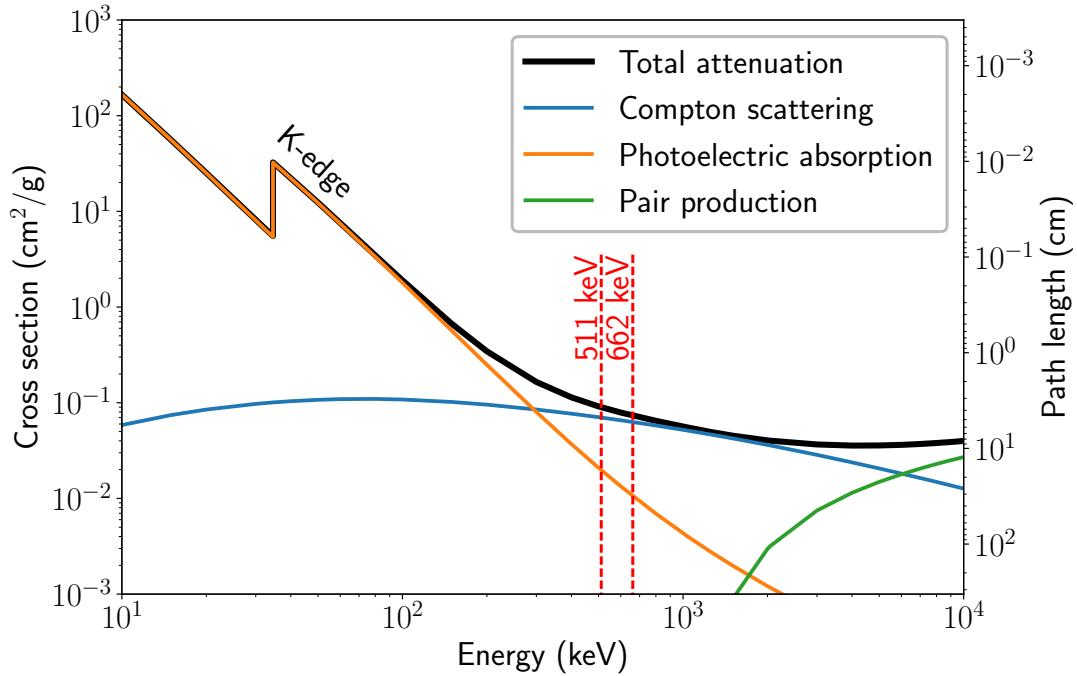
The calibration of the LDE, electron lifetime and g_1 and g_2 require a source of mono-energetic interactions. Since XAMS is a relatively small TPC, external gamma-ray sources can be used for this. The sources used for the measurements described in this section are ^{137}Cs and ^{22}Na , at gamma-ray energies of 662 keV and 511 keV, respectively.¹

4.1 GAMMA-RAY INTERACTIONS

There are three relevant types of gamma-ray interactions with matter: photoelectric absorption, Compton scattering and pair production.² The interaction cross section σ for these processes in xenon as a function of energy is shown in figure 4.1. In the case of photoelectric absorption, the photon is absorbed by an interaction with the atom, which emits an electron from one of the atomic shells. The ejected electron obtains a kinetic energy equal to the gamma-ray energy minus the binding energy of the shell that it was emitted from. The resulting vacancy is most often filled by one of the electrons in the higher energy atomic shells, resulting in a cascade of characteristic X-rays. In some cases (12.5% in xenon), the shell is filled by

¹ In addition to 511 keV gamma rays, ^{22}Na also emits a 1274 keV gamma ray in its decay. However, the trigger setup used for ^{22}Na (described in section 4.2) limits the number of 1274 keV passing the trigger selection.

² A fourth process is coherent or Rayleigh scattering, which changes the direction of the gamma ray but not its energy. Since there is no observable recoil in this interaction, we omit this from the discussion here.



4

Figure 4.1: Cross sections and corresponding path length of gamma-ray interactions in xenon as a function of the gamma ray energy (colored lines) and the sum of all cross sections (black line). The dominant process changes with energy; photoelectric absorption is the main process up to approximately 300 keV, followed by Compton scattering and finally pair production, at energies above 6 MeV. The energies of the gamma rays from ^{22}Na (511 keV) and ^{137}Cs (662 keV) are indicated by the red dashed lines. Data taken from [49]; source code available at [10].

the Auger effect, where the energy difference from an electronic transition is transferred to one of the outer electrons [99]. Since the range of X-rays (less than one mm) and Auger electrons ($\mathcal{O}(10 \mu\text{m})$) at these energies is small with respect to the typical detector position resolution, the net result of a photoelectric absorption is a full energy absorption at essentially one position. A notable exception to this is if the interaction occurs close to the edge of the sensitive volume, so that one of the X-rays may escape the detector unobserved.

The second interaction process is Compton scattering, which becomes the dominant process in xenon at energies above approximately 300 keV. In this case, the photon scatters off an electron and is deflected at an angle. Only part of the initial photon energy is transferred to the electron, depending on the deflection angle. The scattered photon moves away after the interaction and may interact again or exit the detector. As the photon loses its energy by Compton scatters, photoelectric

absorption becomes more likely, so that the photon may deposit all its energy in the detector in multiple interactions. At the energy of the ^{22}Na (511 keV) and ^{137}Cs (662 keV) gamma rays, Compton scattering is the dominant process, at cross sections which are respectively 3.5 and 5.5 times higher than those for photoelectric absorption.

For high gamma-ray energies, pair production may occur. In this process, a photon converts to a electron-positron pair under the influence of the nuclear electric field (or, to a lesser extent, the electronic electric field). The positron annihilates with an electron in the interaction medium (after losing its kinetic energy) and yields a pair of 511 keV gamma rays. Since two massive particles with $m_e = 511$ keV are produced, pair production can only occur if the incident gamma-ray energy exceeds 1022 keV. For the sources used in this analysis, the only gamma ray that satisfies this criterion is the 1274 keV gamma ray from ^{22}Na , although the pair production cross section is only 0.5 % of the total cross section.

4.2 DATA TAKING AND PROCESSING

A schematic of the XAMS data acquisition (DAQ) chain is shown in figure 4.2. The analog waveforms from the two PMT channels are fed to two channels of a CAEN V1730D digitizer. On-board trigger logic issues a start-readout if both signals simultaneously cross a pre-set threshold. The trigger level was set at 7.0 mV and the coincidence time was set to 120 ns. This trigger setting is suited for both the short S_1 signals as well as the longer S_2 signals, so that an event can trigger the DAQ even if the S_1 is very small or if there is no S_1 in the event at all. When a trigger is generated, the full waveform information for both channels is sent to the DAQ PC via optical link. The digitizer has a circular memory buffer, which continuously digitizes the signal channels, so that the trigger position can be placed in the middle of the event window. The window of 163.8 μs is chosen such that the maximum drift time (roughly 60 μs) is comfortably contained in the event, regardless of which signal triggered the readout.

For the ^{22}Na data, an external NaI(Tl) detector was used in a setup similar to the one described in section 3. This exploits the two simultaneous back-to-back 511 keV gamma rays that result from the annihilation of the positron produced in the decay of ^{22}Na (with a branching ratio of 90.4 %). The trigger was in this case set to a three-fold coincidence requirement of the two internal PMTs and the external detector. This setup enables the suppression of background radiation, so that only radiation from the ^{22}Na source is left in the data sample. Since a 1274 keV gamma

ray is emitted (almost) simultaneously with the positron in the decay of ^{22}Na , it is possible that a trigger is generated when this gamma ray interacts in the TPC if one of the 511 keV gamma rays interacts in the NaI(Tl) detector. However, since the NaI(Tl) detector, source and TPC are in line-of-sight and the 511 keV gamma rays are emitted in opposite directions, the 511 keV gamma rays are much more likely to generate a trigger.

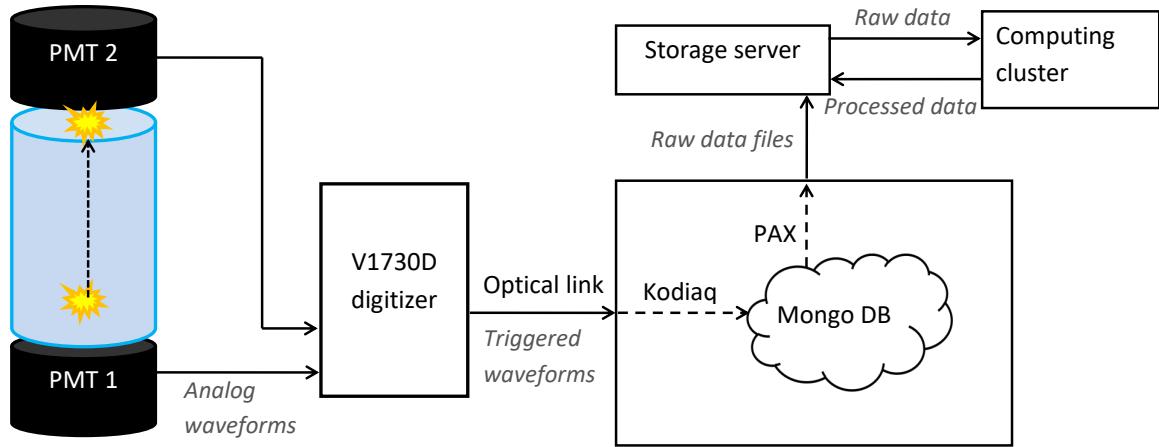


Figure 4.2: Schematic of the DAQ process for the XAMS measurements. The gray cursive labels indicate the type of data at the corresponding arrows. The PMT data is saved by a V1730D digitizer, which also handles triggering based on a coincidence of both PMT channels crossing a threshold. The full waveform information is sent to a DAQ PC (bottom right panel), which runs Kodiaq to store the data in a temporary Mongo database. PAX reads events from this database, performs software zero-length encoding (see text) and stores raw data files to the Nikhef server. An automatic processing script sends a batch job to the local computing cluster, which runs an instance of PAX that analyzes the waveforms and outputs the processed data into a new file.

The readout PC runs several processes simultaneously during data taking. The software that handles the data as it comes in from the digitizer is a standalone version of Kodiaq [100]; the same readout software as used for the XENON1T experiment. The data is written to a Mongo [101] database. An instance of the XENON1T data processing code PAX [85] reads the waveforms from the database and builds raw data files from the waveforms. It simultaneously applies a software equivalent of zero-length encoding (ZLE), which is a data reduction algorithm installed on the V1724 digitizers used in XENON1T [86]. This means that all samples sufficiently close to the baseline are removed from the waveform. If there is an excursion more than 5 or 6 ADC counts from the baseline, only this data and 50 samples before and

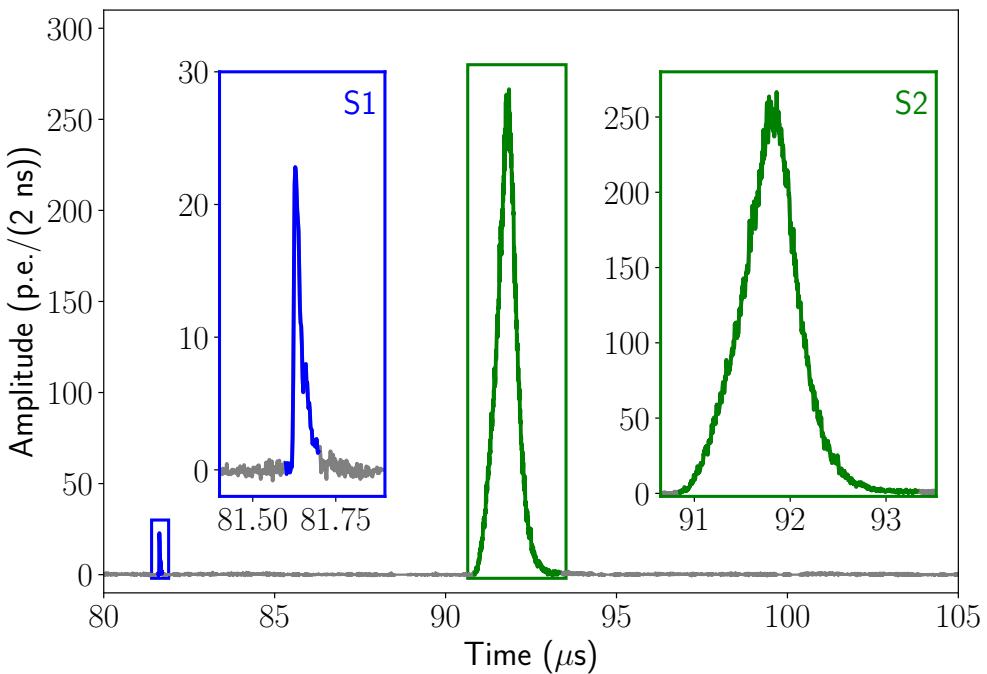


Figure 4.3: A typical sum-waveform for an event coming from the ^{137}Cs calibration. The first narrow peak, shown in the blue inset, is the S_1 signal. The second, much wider signal shown in the green inset is the S_2 -signal. The S_1 and S_2 area are approximately 330 p.e. and 88×10^3 p.e. and the widths are 25 ns and 395 ns, respectively.

after this are retained. By applying the software ZLE, the data volume is greatly reduced and low-frequency noise in the baseline can be accounted for (as there is a baseline correction on the start of each stretch of data that passes the ZLE).

The DAQ PC also runs a database that keeps track of all datasets taken so far. When a run finishes and all data files are built, the status of that run is set to ‘pending’, which triggers a script running on the Nikhef server. This builds a batch job that is sent to the local computer cluster, so that the data is processed by PAX and can be analyzed.

Figure 4.3 shows a waveform from the ^{137}Cs measurement. PAX finds significant excursions from the local baseline, called *hits*, by comparing to a threshold set at 7 times the standard deviation of the local baseline (computed on the first 40 samples of the ZLE chunk). The boundaries of the peaks are set at 10 ns left and right of the points where the 7σ threshold is crossed. The hits are then clustered into *peaks* based on their separation; if the hits are less than 250 ns from each other, they are supposed to belong to the same peak. A second step cuts the peaks into separate peaks if a shallow enough local minimum is found in the peak, so that even peaks

that are close to each other can be separated (such as events where two scatters occur, creating two S₂s, or events where the drift time is very short). All relevant properties of the peaks, such as the area, height and width are then computed.

The peak classification is based on the peak area and the width, as measured by the time span containing the central 50 % of the area of the peak. If this is more than 100 ns and the area exceeds 8 p.e., the peak is classified as an S₂; if the width is less than 60 ns, the peak is identified as an S₁.³ An additional requirement is that both PMT channels have hits contributing to the peaks. All other coincident peaks are classified as ‘unknown’. Based on the width distribution of the S₁s and S₂s, the acceptance loss due to classification is estimated to be negligible. Only 0.4 % of all events contain an ‘unknown’ peak, and most of them are consistent with split S₂ tails or single-electron S₂ signals. After classification, the S₁s and S₂s are paired to form an interaction. The main interaction, containing the main S₁ and main S₂ used for analysis, are defined as the largest S₂ in the event with at least one S₁ occurring before it, and the largest S₁ before this S₂. Based on the information of these peaks, the relevant properties of the interaction that occurred can be derived.

4

4.3 DATA SELECTION

In general, most analyses require a single interaction in the sensitive volume of the TPC, which generates exactly one S₁ and one S₂ following it. However, there are several event topologies that contaminate the single-scatter signal and must be rejected from the analysis. Table 4.1 lists the experimental signatures of several physical processes that occur in the TPC.

Table 4.1: Event properties for various interaction possibilities in the TPC.

Signature	Physical process
No S ₁ , no S ₂	Electronic noise
No S ₁	Interaction in light-insensitive region or merged S ₁ -S ₂ pair
No S ₂	Interaction in charge-insensitive region
Multiple S ₁ s	Pileup
Single S ₁ , multiple S ₂ s	Resolved multiple scatters
Single S ₁ and single wide S ₂	Clustered S ₂ s from multiple scatters
Single S ₁ , single S ₂	Single scatter (or unresolved multiple scatter)

³ For low drift fields, such as in used in chapter 6, the upper bound for the S₁ width is set to 80 ns to capture wide S₁s caused by a slower recombination process.

Events without an S₁ or an S₂ are usually triggered by bipolar noise induced by the Pt100 readout system. An event containing only an S₂ is indicative of interactions that occur in a light-insensitive region in the TPC or with an S₁ too close to the S₂ to be separated. Most of the events with no S₁ have S₂ properties similar to the S₂ properties found for events with very short drift times that occur in the liquid layer between gate and anode. Events with no S₂ are more common and are produced in the charge-insensitive region below the cathode, where the direction of the electric field is away from the gas layer. If there are multiple uncorrelated interactions in the detector volume within the duration of a waveform, two S₁s are found in the waveform (except in very exceptional cases). In the ¹³⁷Cs calibration, 2.1 % of all events are cut by this requirement, consistent with the expected pileup fraction of 2.2 % at the trigger rate of 137 Hz.

4.3.1 Multiple scatters

In the ¹³⁷Cs and ²²Na calibration datasets, multiple scatter events are quite common. This is because at the energies of the gamma rays from these decays (511 keV and 662 keV, respectively), Compton scattering is the dominant interaction process, as discussed in section 4.1. If multiple interactions occurs, their S₁s coincide as the delay between these interactions is typically less than 1 ns in XAMS (given the speed of light of 30 cm/ns). In contrast, the S₂s are delayed by a time equal to their drift time, which depends on the z-coordinate of the interaction position. If the interaction position separation Δz is sufficiently large, the drift time difference is enough so that an event with multiple S₂s is observed. In the ¹³⁷Cs calibration, roughly 20 % of all events passing previous cuts are classified as multiple scatters. Figure 4.4 shows the difference of the z-coordinates of the two largest S₂ signals in events with multiple S₂s.

The distribution of separation between interactions Δz depends on the detector geometry and the range of the gamma rays between the scatters, which depends on the gamma-ray energy (see figure 4.1). To properly account for this, a Monte Carlo simulation is required. However, some simplifying assumptions allow for the formulation of a basic expected distribution. Supposing a constant effective mean free path R_{eff} for all gamma rays between the scatters, the one-dimensional probability function for the distance between the scatters r is:

$$p_r(r; R_{\text{eff}}) = \frac{1}{R_{\text{eff}}} \exp\left(-\frac{r}{R_{\text{eff}}}\right). \quad (4.1)$$

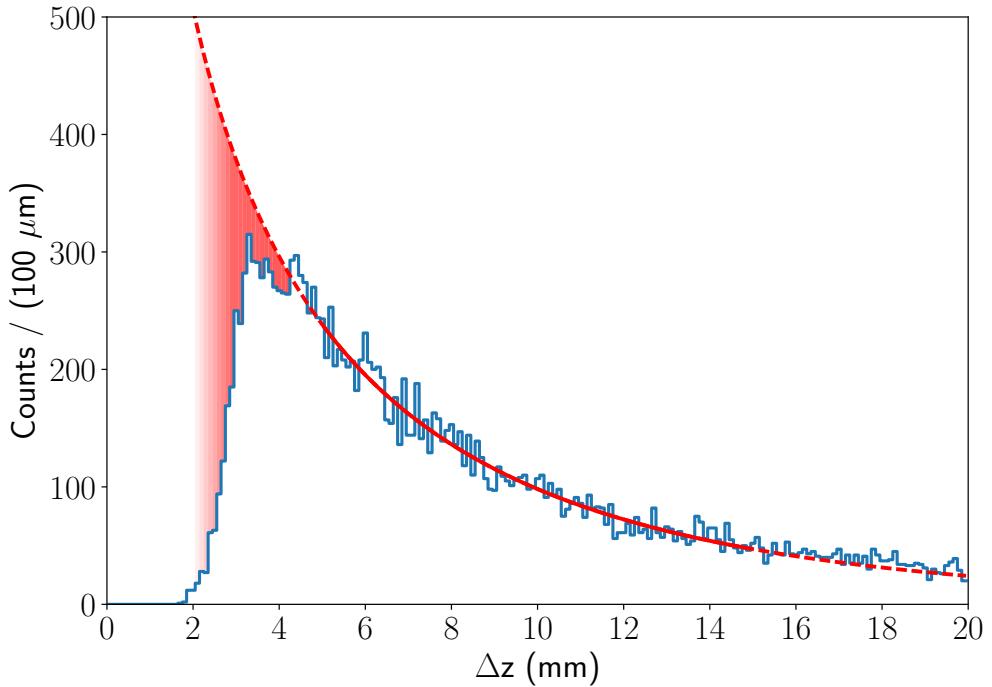


Figure 4.4: The difference in reconstructed z -position for the events containing at least two S2s. A decreasing trend is observed, coming from the finite path length of the gamma ray after a Compton scatter, in combination with a geometric effect. At low separation distances (< 4 mm), the distribution starts to fall as the S2s are too close in time and are thus merged. A fit of equation 4.3 is shown by the red curve; this gives an effective path length of (11.0 ± 0.4) mm. At $\Delta z = 3.0$ mm, 50 % of the events are correctly split and 50 % are left in the single scatter sample, based on the extrapolation of the fit curve. Some of the events that are too close to be separated are cut from the single scatter sample by the S2 width cut (see section 4.3.2). Based on the number of events cut and the fit extrapolation, the events removed by this cut would fill the red shaded region that extends down to 2.0 mm. We therefore estimate that only double scatter events with $\Delta z < 2.0$ mm or true single scatter events are left in the sample after both cuts.

For three dimensions, the probability density function must be divided by $4\pi r^2$ due to the spherical integration normalization. If we define \vec{r} as the separating vector between the scatters, we obtain:

$$p(\vec{r}; R_{\text{eff}}) = \frac{A}{r^2} \exp\left(-\frac{r}{R_{\text{eff}}}\right) = \frac{A}{x^2 + y^2 + z^2} \exp\left(-\frac{\sqrt{x^2 + y^2 + z^2}}{R_{\text{eff}}}\right), \quad (4.2)$$

with $A = 1/(4\pi R_{\text{eff}})$ a normalization factor. The distribution of the z -distance between the scatters can now be computed through:

$$\begin{aligned} p_z(z; R_{\text{eff}}) &= \iint p(\vec{r}; R_{\text{eff}}) dx dy \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{A}{x^2 + y^2 + z^2} \exp\left(-\frac{\sqrt{x^2 + y^2 + z^2}}{R_{\text{eff}}}\right) dx dy, \end{aligned} \quad (4.3)$$

which has no closed-form expression but may be calculated numerically. Note that in addition to the assumption of constant mean free path R_{eff} we implicitly assume here that the detector is large with respect to R_{eff} (since the integral boundaries are infinite). Figure 4.4 shows a fit of equation 4.3 to the observed multiple scatter separation between 5 and 15 mm. The best-fit effective path length is $R_{\text{eff}} = (11.0 \pm 0.4)$ mm, roughly corresponding to the path length of 200 keV gamma rays of 9.8 mm [49]. This is a reasonable value for the typical energy of ^{137}Cs gamma rays after an initial Compton scatter. We note that the fit describes the distribution surprisingly well, in spite of the rough assumption of one effective path length which in fact depends on the energy rather strongly (see figure 4.1).

At low separation distances, the observed Δz -distribution in figure 4.4 starts to deviate from the expected distribution. This is due to the clustering of multiple S2s; if the distance between two scatters is too low, the difference in drift time is too low and the S2s occur nearly simultaneously. This causes the waveforms to overlap. However, some of these events may be distinguished using the increased width that results from the delay between the S2s.

4.3.2 S2 width

The S2 width increases with the drift time due to longitudinal diffusion of electrons during the drift from the interaction position to the gas region. The standard deviation σ of the S2 time distribution depends on the interaction position z according

to [102]:

$$\sigma_{S2} = \sqrt{\sigma_{\text{diff}}^2 + \sigma_0^2} = \sqrt{\frac{2D_L t_d(z)}{v_d^2} + \sigma_0^2}, \quad (4.4)$$

with D_L the diffusion coefficient, and t_d and v_d the drift time and drift velocity, respectively, while σ_0 is the basic width of an average single electron S2 signal (determined by the amplification gap) [57]. Equation 4.4 is an approximation that holds if the S2 shape is sufficiently Gaussian, which is the case for long drift times as diffusion smears the S2. We therefore fit only for drift times $> 30 \mu\text{s}$, selecting the bottom half of the detector. Furthermore, since equation 4.4 describes the standard deviation of the S2, an additional factor is required to convert to the width parameter W used in analysis: the range containing the center 50 % of the area. Again assuming a Gaussian distribution for the S2, these are related according to $W = 1.349\sigma$, so that equation 4.4 becomes [44]:

$$W(t_d) = \sqrt{\frac{3.640 \cdot D_L t_d(z)}{v_d^2} + W_0^2}. \quad (4.5)$$

Figure 4.5 shows the two-dimensional distribution of the S2 width as a function of drift time. The increase of the S2 width as a result of diffusion is clearly visible, as is a nonzero contribution of events with a higher width than expected from the model, caused by merged S2s. The data is sliced in bins of $2 \mu\text{s}$ and the median is calculated for each slice, yielding the blue points in the figure. The S2 width model of equation 4.5 is fit to all datapoints between 30 and $60 \mu\text{s}$. This gives a diffusion constant of $(15.5 \pm 0.6) \text{ cm}^2\text{s}^{-1}$ at a drift field of 500 V/cm , consistent with the diffusion constant found for XENON100 of $(14.1 \pm 1.5) \text{ cm}^2\text{s}^{-1}$ at a comparable field strength of 530 V/cm [44]. The uncertainty quoted incorporates the systematic uncertainty induced by changing the bin width or making changes to the fit range, as well as the statistical uncertainty.

Given the fitted width model, the deviation from the expected width can be calculated for each event. A histogram of this distribution is shown in figure 4.6. There is a population with widths considerably larger than expected from diffusion, originating from merged S2s. We apply a cut at 95 % of the population, giving a threshold of $2.5 \times 10^2 \text{ ns}$, which is indicated by the orange line in figure 4.5.

By extrapolating the fit to the distribution in figure 4.4, an estimate of the minimum distance required to reject a multiple scatter event based on the width can be obtained. For this, we assume that (i) all events cut by the S2 width cut are due

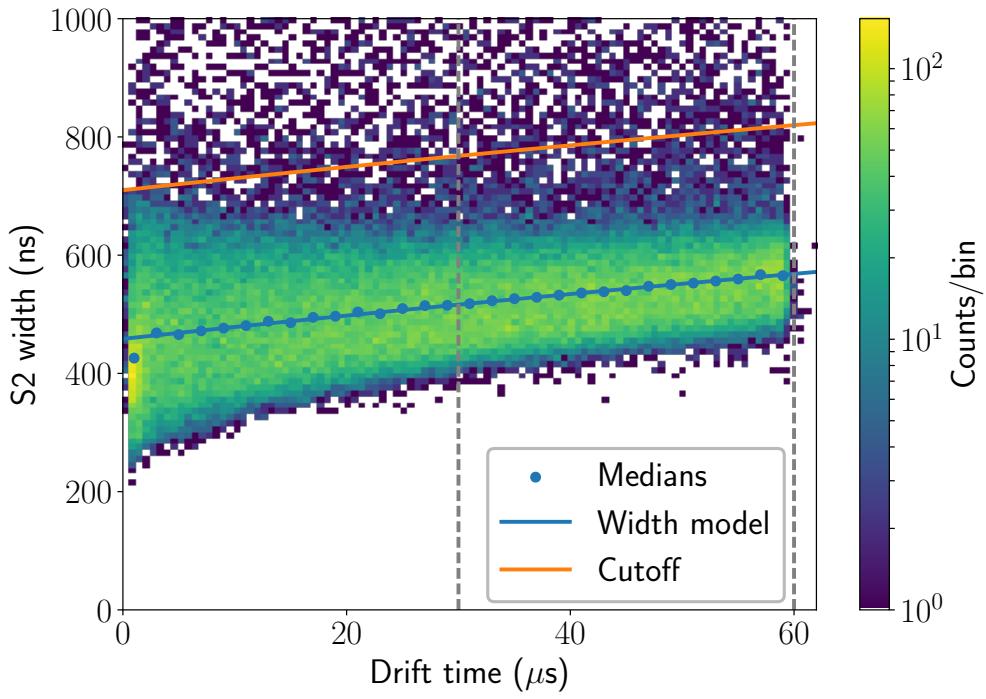
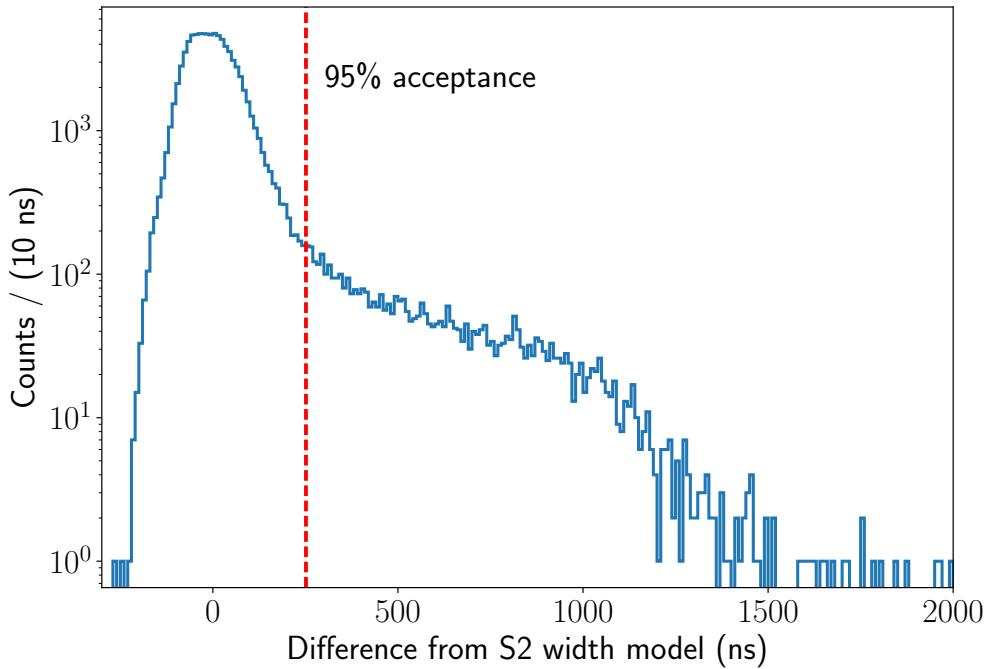


Figure 4.5: The S_2 width, measured as the time containing the middle 50 % of the area of the S_2 , as a function of drift time. The blue points indicate the median value for drift time slices of $2\ \mu\text{s}$ wide. The width is increasing due to electron diffusion, described by the width model in equation 4.5. A fit of the blue data points is shown by the blue line, yielding a diffusion constant of $(15.5 \pm 0.6)\ \text{cm}^2\ \text{s}^{-1}$. The fit is performed only for values between 30 and 60 μs (indicated by the gray dashed lines) as the width model approximation breaks down at low drift times. The orange line indicates the upper bound for the S_2 width: S_2 s that are too wide are likely caused by merged S_2 s from multiple scatter events.

to clustered multiple scatters and (ii) their separation Δz is too small to separate the S_2 s (i. e. below 4 mm), but still relatively high, as indicated by their large width. This means that the events would fill the space of ‘missing events’ indicated by the red shaded region in figure 4.4. The minimum separation Δz required for multiple scatters events to be cut by either having multiple S_2 s or a wide S_2 obtained in this way is 2.0 mm. Based on the extrapolation below 2.0 mm, roughly 20 % of the events left after all cuts are actually multiple scatter events, while all other events are true single scatter events.



4

Figure 4.6: Distribution of the difference of S_2 width and the width described by the diffusion model (shown in figure 4.5). There is a contribution of events with a large width, coming from merged S_2 s from double scatters that are too close to be resolved. A cut at the 95 % quantile, indicated by the red dashed line at 2.5×10^2 ns, is applied to remove these events.

4.4 POSITION CALIBRATION

The z -position of an interaction in the TPC can be determined by its drift time. However, even events occurring close to the gate mesh (which we define as $z = 0$) will have a nonzero drift time. This is due to the finite drift time in the shallow layer of liquid between the gate and the anode mesh and due to the calculation of the drift time that uses the difference in the center time of the S_1 and S_2 peak. If the drift field, and thus the drift velocity, are assumed to be constant throughout the drift region between the cathode and the gate mesh, the mesh positions in drift time can be used as calibration points to determine the drift time offset and the drift velocity. A further examination of this assumption is provided in chapter 6.

As the field strength near the gate mesh is much higher than in the drift field region, free electrons are pulled away and contribute to the S_2 signal rather than to the S_1 signal through recombination. This can be observed in the left panel of figure 4.7, where the ratio S_2/S_1 shows a clear decrease between 1 and 2 μ s drift.

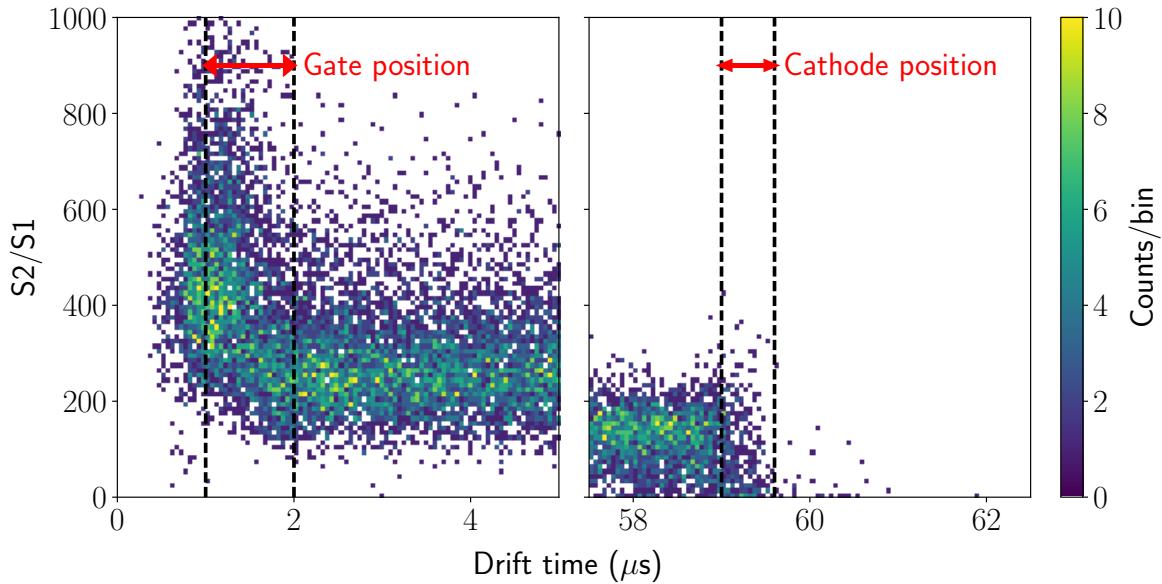


Figure 4.7: The ratio of S_2/S_1 for 5 μs of drift time, at the gate (left panel) and cathode mesh (right panel) position. There is a clear jump coming from the changing field conditions around the meshes, changing the recombination efficiency. The drift times associated to these meshes are $(1.5 \pm 0.5) \mu s$ and $(59.3 \pm 0.3) \mu s$, respectively. These ranges are indicated by the dashed lines.

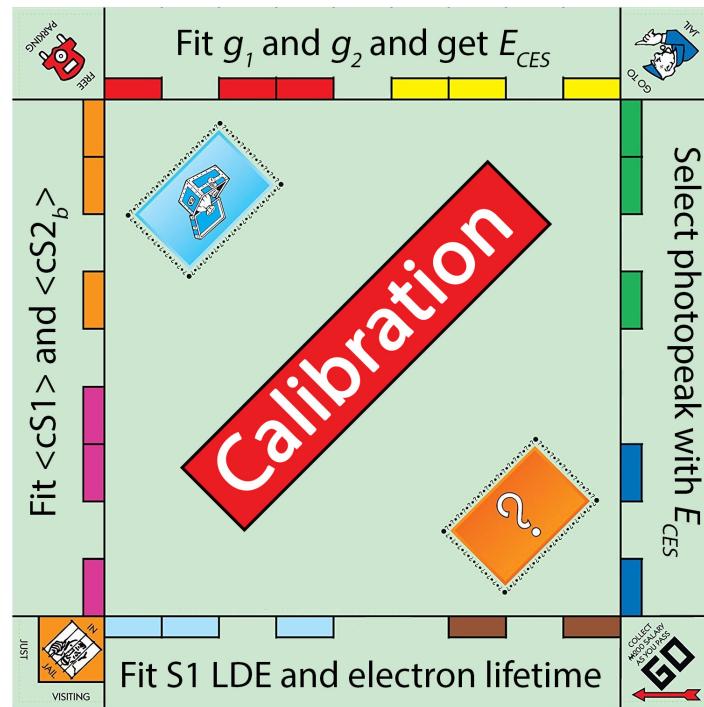
4

Based on this, we determine the drift time corresponding to the gate position as $(1.5 \pm 0.5) \mu s$. For the cathode mesh, a similar effect can be observed in the right panel of figure 4.7. The cathode position is also the position of maximum drift time, so that few events surpass the drift time corresponding to the cathode position. The cathode drift time is set at $(59.3 \pm 0.3) \mu s$, so that we arrive at an average drift velocity of $(1.73 \pm 0.04) \text{ mm}/\mu s$ over the full length of $(100.0 \pm 2.0) \text{ mm}$, at a drift field of 500 V/cm. Based on simulations of the electric field shown in chapter 6, the inhomogeneity of the field causes a variation of the drift velocity in the drift region of $\pm 1.5\%$, or $\pm 0.03 \text{ mm}/\mu s$. Given the gate drift time offset and the drift velocity, the interaction z -position can be computed for each event.

4.5 S₁ AND S₂ CORRECTIONS

4.5.1 General outline

Some of the corrections that are required for liquid xenon TPCs are computed from events with constant recoil energy. For the S_1 , the light detection efficiency (LDE)



4

Figure 4.8: The iterative correction and energy calibration procedure. After an initial rough photopeak selection (at GO), the S_1 as a function of z and the S_2 as a function of drift time are used to determine the LDE and electron lifetime corrections. This makes it possible to compute the average values of cS_1 and cS_{2b} , which allows the determination of g_1 and g_2 . The combined energy scale, calculated using these values, makes a better selection of the photopeak possible, so that the corrections can be refined. This process is repeated until all values converge.

varies with position as the probability of the scintillation photons reaching either one of the PMTs depends on the optical properties of the TPC. For the S_2 , all the light is created in the same narrow region between the liquid level and the gate mesh, but the number of electrons depends on the drift time as the electrons can attach to impurities in the liquid xenon during the drift to the surface. Both of these corrections were discussed in chapter 2 in the context of XENON1T, for which the distinct event structure of ^{83m}Kr is used to select a single energy. The LDE and electron lifetime calibration of XAMS use events from mono-energetic gamma rays. However, background gamma radiation and Compton scatter events require that the mono-energetic photopeak first be selected from the data.

The best selection (after the single-scatter cuts outlined in section 4.3) is performed by the combined energy scale that uses the combination of S_1 and S_2 and

is computed using:

$$E_{CES} = W \left(\frac{cS1}{g_1} + \frac{cS2_b}{g_2} \right), \quad (4.6)$$

where $cS1$ and $cS2_b$ are the corrected signals and g_1 and g_2 are the primary and secondary scintillation gain as defined in section 2.4.4. However, it is apparent from equation 4.6 that calculating E_{CES} requires the knowledge of the corrected S_1 and S_2 , which is precisely the goal of the selection. In addition, the values of g_1 and g_2 are also obtained using multiple known mono-energetic values of $cS1$ and $cS2_b$ (see section 2.4.4). We thus arrive at a situation where the corrections are best calculated if the corrections are already known. This deadlock is broken by adopting an iterative approach, which is schematically shown in figure 4.8. Initially, the photopeak selection is performed on simple by-eye cuts in S_1 and S_2 . S_1 and S_2 corrections are performed based on this data selection, after which the mean values of $cS1$ and $cS2$ can be determined using a Gaussian fit. This is done for both ^{137}Cs (at 662 keV) and ^{22}Na (at 511 keV). When this is done, the values of g_1 and g_2 can be determined, and a better photopeak selection can be obtained using the combined energy parameter and the preliminary S_1 and S_2 corrections. The correction factors and central values for $cS1$ and $cS2$ are then recomputed, after which g_1 and g_2 are redetermined, and the analysis is performed until all fitted values converge within their statistical errors. For the sake of readability, in this section we will show the corrections and fits only for the selection used in the last step.

4.5.2 Light yield calibration

Figure 4.9 shows the S_1 area as a function of the z -coordinate in the TPC for the 662 keV photopeak of ^{137}Cs . In addition to the photopeak selection, outlier events in $cS1$ are removed to improve the fit stability. Outlier events are defined as being more than 2.5σ from the mean μ , where both μ and σ are determined from a Gaussian fit of $cS1$. This cut is indicated by the gray dashed curves in the figure. The resulting data shows a clear z -dependence of the LDE as a result of the light propagation properties in the TPC and its geometry. Since there is no *a priori* expectation of the functional form, we fit the data with a second-degree polynomial in the range between -9.5 cm and -0.5 cm (indicated by the gray dashed vertical lines), excluding 5 mm of the volume near the gate and cathode mesh to remove data which may be influenced by the field distortion. The functional form of the fit formula f

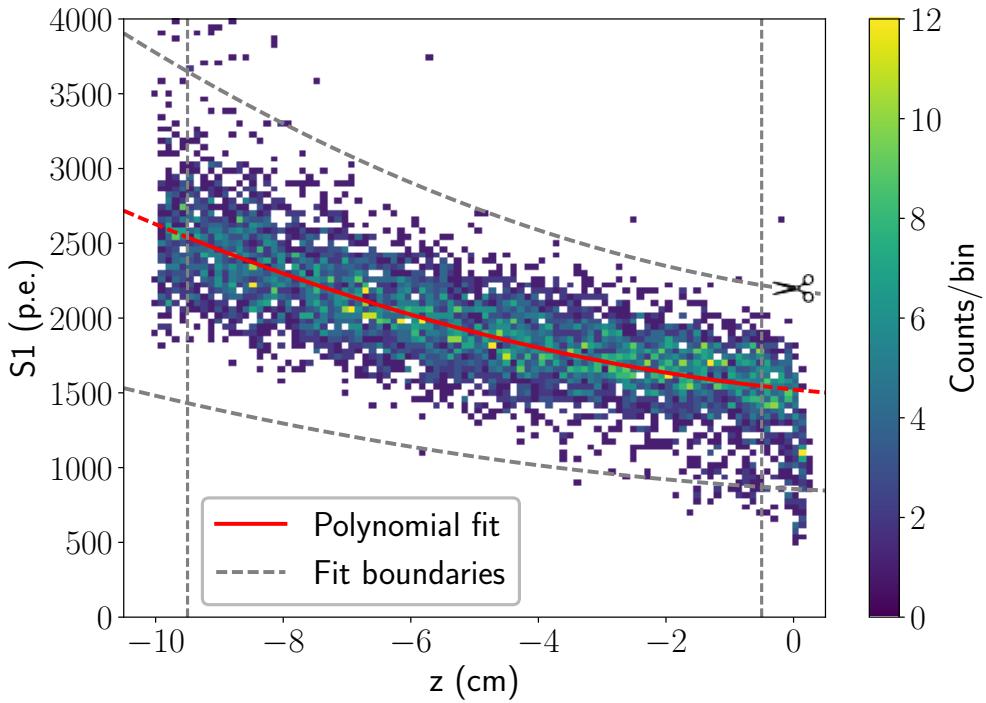


Figure 4.9: The S_1 as a function of reconstructed z -coordinate (with respect to the gate mesh) for the photopeak selection described in the text. A second degree polynomial function is fit to the data to correct for the position dependence of the light detection efficiency. The fit range is indicated by the dashed lines, and is limited in order to avoid edge effects close to the cathode and gate meshes.

is thus:

$$f(z) = c_2 \cdot z^2 + c_1 \cdot z + c_0, \quad (4.7)$$

and the correction factor C_{S1} is normalized to the volume average, so that:

$$C_{S1}(z) = \frac{\langle f(z) \rangle}{f(z)} = \frac{\int f(z') dz'}{L} \frac{1}{f(z)}, \quad (4.8)$$

where the integral runs over the full length L of the active region of the TPC, thus scaling the S_1 to the average response in the TPC. The corrected S_1 can now easily be computed with $cS1(z) = C_{S1}(z) \cdot S1$.

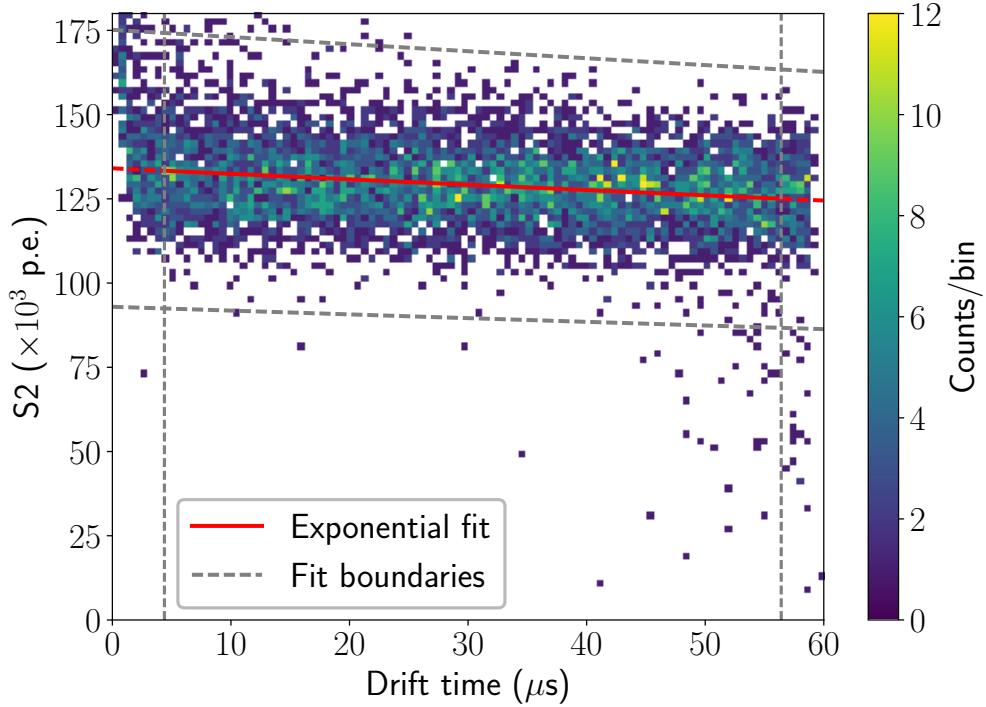


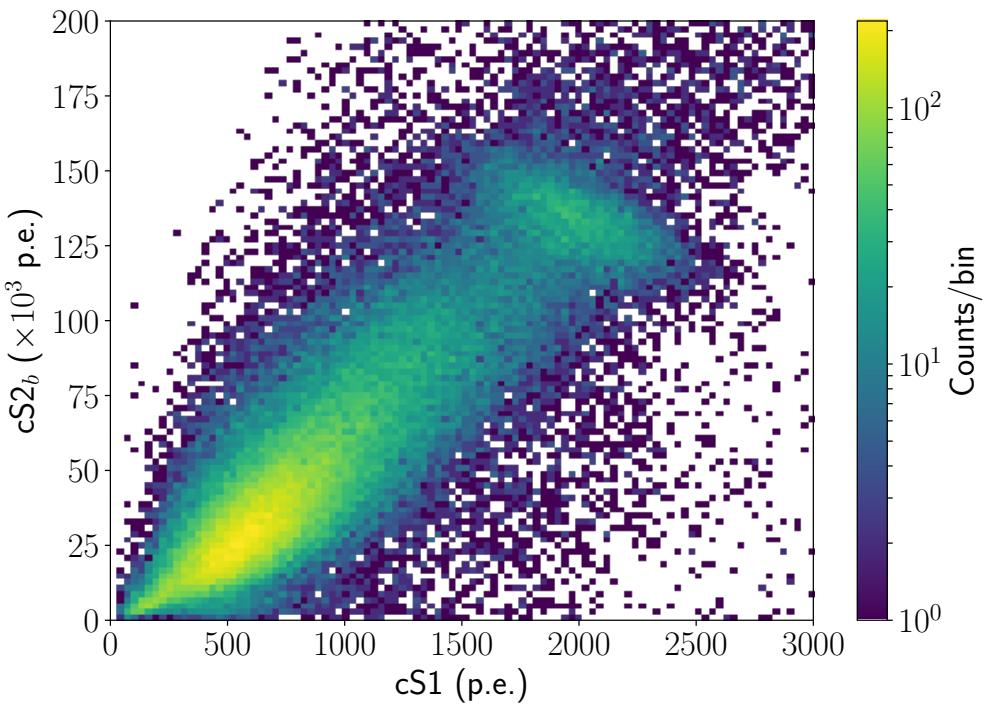
Figure 4.10: The S_2 as a function of drift time for the photopeak selection described in the text. A net decrease is found due to electron attachment to impurities in the liquid. This is described by an exponential function with an electron lifetime $\tau_e = (0.81 \pm 0.04) \text{ ms}$, shown by the red curve.

4.5.3 Electron lifetime

The data selection method for the S_2 electron lifetime correction is similar to the one used for the S_1 correction. Again, the photopeak selection is determined from a fit to the CES spectrum, after which outliers are removed using a Gaussian fit to the cS_2 spectrum and cutting at 2.5σ from the mean. The resulting data and exponential fit are shown in figure 4.10. The fit gives an electron lifetime of $(0.81 \pm 0.04) \text{ ms}$. All events are corrected for electron loss using the inverse factor $\exp(t_d/\tau)$.

4.5.4 Energy calibration

After all the aforementioned cuts and the corrections for S_1 and S_2 , the distribution of corrected S_1 and S_2 is shown in figure 4.11. The photopeak around 2000 p.e. in S_1 is clearly visible, with a Compton scatter contribution extending down to the origin.



4

Figure 4.11: The distribution of events in corrected S_1 - and S_2 -space. For the S_2 , the area in the bottom PMT is used to avoid base saturation. A clear photopeak is seen at 2000 p.e. in cS_1 . Extending to lower values of cS_1 and cS_{2b} are Compton scatters and low-energy background events.

Rather than using the corrected value of the full S_2 , only the S_2 in the bottom PMT (referred to as S_{2b}) is used. This is due to saturation observed in the top PMT for large S_2 sizes, as evidenced from a decreasing fraction of the S_2 observed in this PMT for large S_2 s.

The values of cS_1 and cS_{2b} are both related to the energy of the interaction, as more energy on average yields an increase in the light and charge signals. However, at fixed energy the two signals are anti-correlated due to recombination luminescence variation: a decrease in the charge signal S_2 correlates with an increase in the light signal S_1 due to the increased fraction of electrons participating in recombination. This anti-correlation is seen in figure 4.11 by the downward sloping ellipse of the photopeak. The anti-correlation may be exploited by the combined energy scale, consisting of a linear combination of S_1 and S_2 .

As noted in section 2.4.4, the calibration of the combined energy scale corresponds to determining the values of g_1 and g_2 in equation 4.6 by variation of the scintillation to ionization ratio. This is accomplished by one of two methods: measuring multiple sources at various energies or changing the drift field. In the case

of XENON1T, a variety of sources is used (see figure 2.5). For XAMS, only ^{137}Cs and ^{22}Na data was taken, giving relatively large uncertainties on the values of g_1 and g_2 . In chapter 6, a variation of drift field is used for the energy calibration; the values obtained there are in agreement with the values found here.

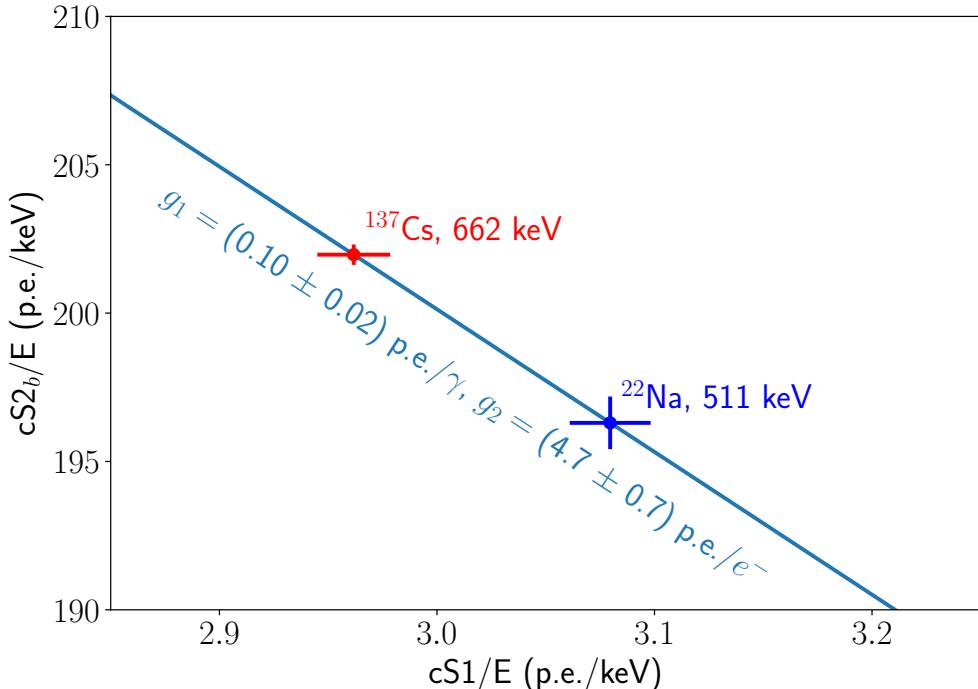
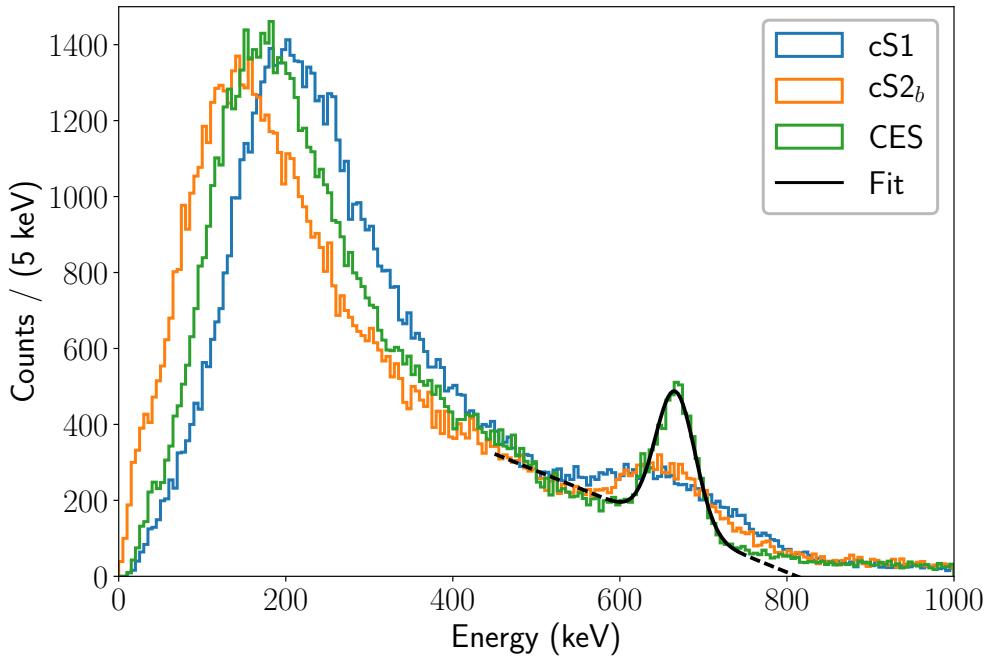


Figure 4.12: ‘Doke plot’ for the two calibration sources used. The blue line corresponds to $g_1 = (0.10 \pm 0.02) \text{ p.e.}/\gamma$ and $g_2 = (4.7 \pm 0.7) \text{ p.e.}/e^-$.

Figure 4.12 shows the energy calibration of XAMS. The two points correspond to the photopeak position in $cS1$ (x coordinate) and $cS2_b$ (y coordinate) for the two sources, divided by the photopeak energy. The photopeak position is determined by an unbinned Gaussian fit of the $cS1$ and $cS2_b$ spectrum after a photopeak selection, similar to the selection used for the LDE and electron lifetime corrections. The uncertainty, as shown by the error bars in the figure, is a combination of the statistical error of the fit and a systematic error that was estimated by varying the fit range and photopeak selection range. From these two points, the values of g_1 and g_2 can be determined using equation 2.4. The uncertainties on g_1 and g_2 are determined by perturbing the data points in figure 4.12 based on the uncertainties in both $cS1$ and $cS2_b$, and calculating the corresponding spread in g_1 and g_2 . This finally yields $g_1 = 0.10 \pm 0.02 \text{ p.e.}/\gamma$ and $g_2 = 4.7 \pm 0.7 \text{ p.e.}/e^-$.

In addition to the combined energy scale, the values of g_1 and g_2 indicate characteristic values that indicate the performance of the TPC. This is especially important



4

Figure 4.13: The distributions of energies for ^{137}Cs estimated from the corrected S_1 (blue histogram), S_2 (orange histogram) and the combined energy scale (green histogram). A significant increase in the energy resolution is achieved when using the combined energy scale. The black curve is an unbinned Gaussian fit with linear background to the CES spectrum, and is used for the photopeak selection.

for low energy events, as the signals are very small. The value of g_1 is related to the average probability of detection of a scintillation photon, as cS_1 is normalized to the average value of S_1 in the TPC. From the values found here, this is approximately 10% in XAMS; a reasonable value expected from the combined effects of the quantum efficiency of the PMTs (usually around 25 %) and light absorption. It should be noted that double photoelectron emission has also been observed with PMTs at the xenon scintillation wavelength [71]; this means that a light detection probability slightly less than 10% is possible. The value of g_2 is related to the electron extraction efficiency, the number of photons per extracted electron generated in the top gas layer and the light detection efficiency of the bottom PMT for light generated between the gate and anode mesh. Since this is a combination of multiple variables, an interpretation is less straightforward in this case. It should furthermore be noted that the extraction efficiency and number of photons created depends on the local electric field, which in turn depends on the position of the liquid to gas interface and the applied anode voltage. If these parameters change significantly, a correction or recalibration of g_2 may be required to retain an accurate energy calibration.

Given the values of g_1 and g_2 , we can calculate the energy spectrum based on equation 4.6. The combined energy scale spectrum of the ^{137}Cs measurement is shown in figure 4.13, along with scaled spectra of cS1 and cS2_b for comparison. There is a clear increase of the energy resolution with respect to using cS1 and cS2_b individually. A Gaussian fit with a linear background component is also shown in the figure; this gives an energy resolution of 5.0 % at 662 keV.

4.6 RESIDUAL GAS ANALYSIS

The operation of xenon TPCs requires a high level of purity from electronegative contaminants. XAMS is fitted with a residual gas analyzer (RGA), so that we can investigate which contaminants are still present in the residual gas when the TPC volume is vacuum pumped prior to or after operation with xenon.

The RGA fitted onto XAMS is a Pfeiffer QME 200 quadrupole mass analyzer [103]. The residual gas content is measured in a process of three steps [104]. First, the gas is ionized with a filament, after which the positively charged gas fragments are accelerated using a potential of typically around 100 V. The ions then enter a quadrupole mass selector. This system consists of four rods with a time-varying potential, which is arranged in a pattern such that only fragments with a specific mass-to-charge ratio m/q are able to pass the mass selector, while other fragments are ejected to the sides. After the mass selector, an ion detector measures the total charge, which is proportional to the number of ions that are measured. By varying the potential applied to the mass selector, a range of masses can be scanned, which is from 0 to 200 amu (atomic mass unit) in the case of the QME 200.

Figure 4.14 shows the first part of an RGA scan of XAMS, taken at a pressure of 2.2×10^{-6} mbar that was achieved after 20 days of continuous pumping. Prior to pumping, the detector was opened for maintenance and exposed to air. After a value of m/q of 48 amu/e, no more noteworthy peaks are found except for a peak structure at approximately 130 and 65 amu/e, corresponding to singly and doubly ionized xenon, respectively. These will be discussed in the following section.

In the figure, the main peaks are labeled with the ion fragment causing it. Table A.1 in appendix A.1 lists peaks and their most probable source. Most ion fragments are caused by constituents of air (N_2 , O_2 , Ar, Ne) or the water in the air. There are some fragments of carbohydrates, which may be attributed to the ethanol that is used to clean the surfaces of the TPC and the vessel, or plastics in the cables and TPC components. We have performed an RGA scan before and after any time that the detector volume was opened, and confirmed that there were no

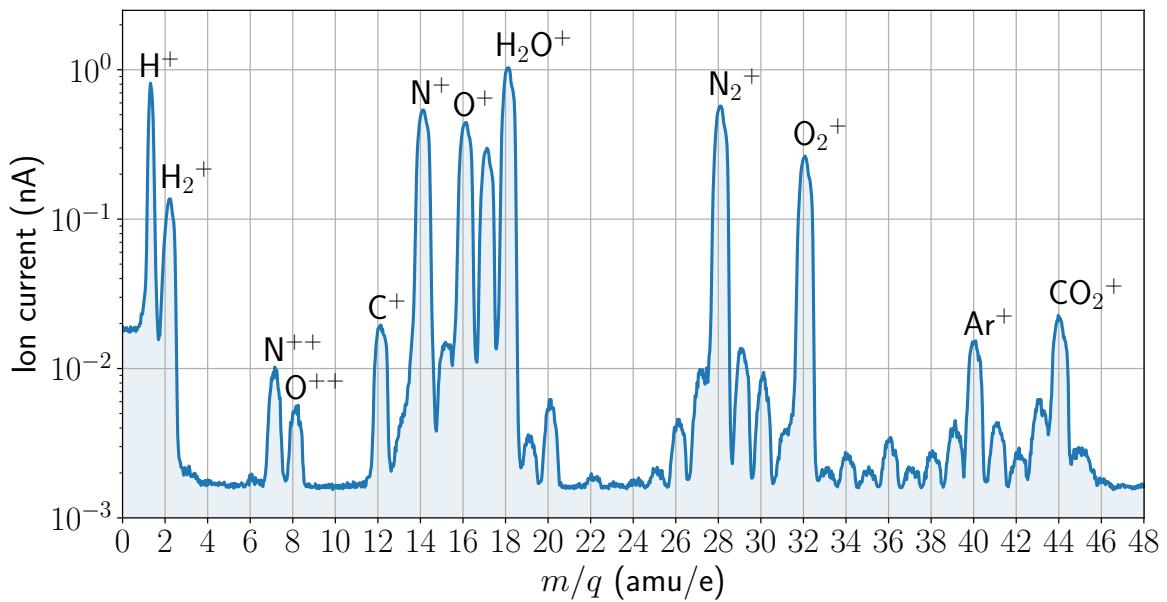


Figure 4.14: The first part of an RGA scan taken to investigate the residual contaminants of the vacuum prior to filling XAMS with xenon. An overview of the peaks, along with the corresponding ion fragment, is given in table A.1. The peaks found here are due to constituents of air, water or ethanol.

4

unexpected peaks in the mass scan or that they were transient.⁴

In addition to the peaks at low mass, a peak structure was found at the mass of xenon isotopes. This is shown in figure 4.15. Since there had been xenon in the system before the scan discussed here was taken, a xenon component can be expected due to outgassing of surfaces of the detector. A total of 9 peaks can be distinguished in figure 4.15, corresponding to all stable isotopes of xenon (124, 126, 128 to 132, 134 and 136). The figure also shows a Gaussian fit to all peaks. For the peak structure between 128 and 132, all peaks were fit simultaneously, with the peak centers constrained at integer isotope masses and with the same standard deviation σ for all Gaussians. The peaks at 124, 126, 134 and 136 were fit with a single Gaussian function, with σ constrained at the same value found in the fit of the large peak structure. The fit range was set at ± 0.5 amu around each peak.

If we assume that the RGA detection efficiency is the same for all isotopes, the area under each peak is proportional to the abundance of the corresponding xenon isotope. The area under each peak is approximately equal to the area under

⁴ For instance, two new peaks at 69 and 83 amu were observed one day after the installation of an optical fiber into the TPC, but they disappeared after two weeks of pumping. This could be due to the TorrSeal glue that was still setting at that time.

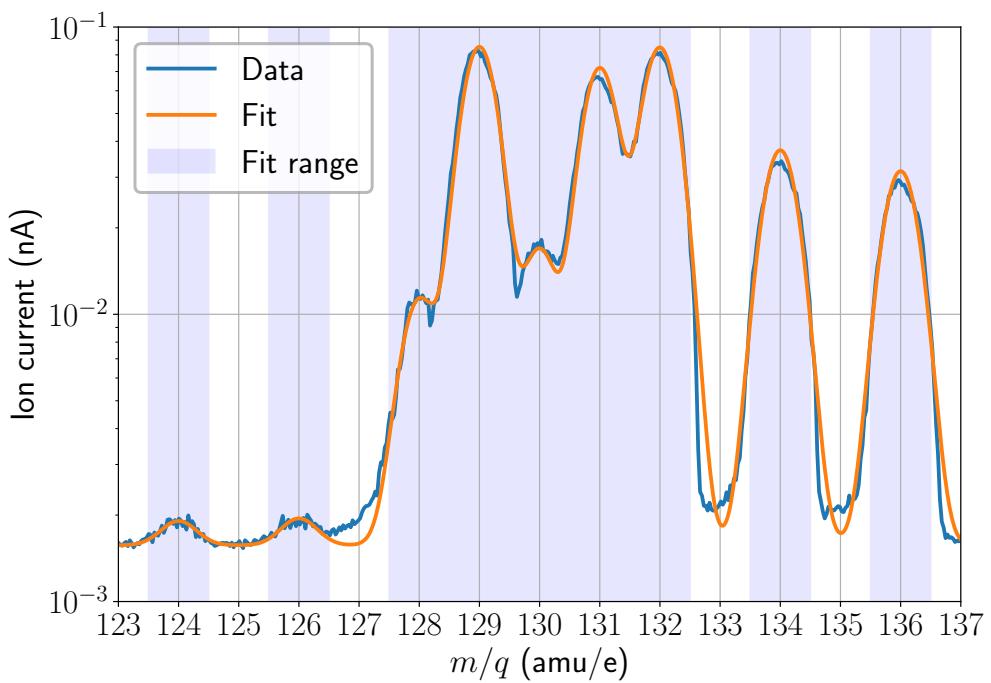


Figure 4.15: The part of an RGA scan showing the xenon isotopes found in the residual gas (blue line). Xenon is found due to emanation of surfaces that have trapped xenon during measurements of the TPC. The orange line shows a Gaussian fit to all peaks. The structure between 128 and 132 is fit by a simultaneous fit of 5 Gaussian functions. For each fit, the shaded area shows the fit range, which is 0.5 amu on either side of the peak. The resulting isotope fractions are shown in figure 4.16.

the Gaussian function, although it can be observed that the Gaussian fit does not fully describe each peak; the peaks in the RGA scan seem to have wider tails than pure Gaussian functions (which can be seen, for instance, in the region between 126.5 and 127.5 amu). This introduces a systematic error in the peak area, which is estimated at 5 % from the area of the tails in the regions between the fits. The resulting abundance for all xenon isotopes is shown in figure 4.16. There is good overall agreement with literature values for atmospheric xenon, except for the isotopes 126, 128 and 130 which show an abundance higher than found in literature. This discrepancy may be explained by the large peak at 129 amu, which has wider tails than given by the Gaussian function. This means that part of the contribution that is attributed to the area at 126, 128 and 130 is actually part of the large peak at 129.

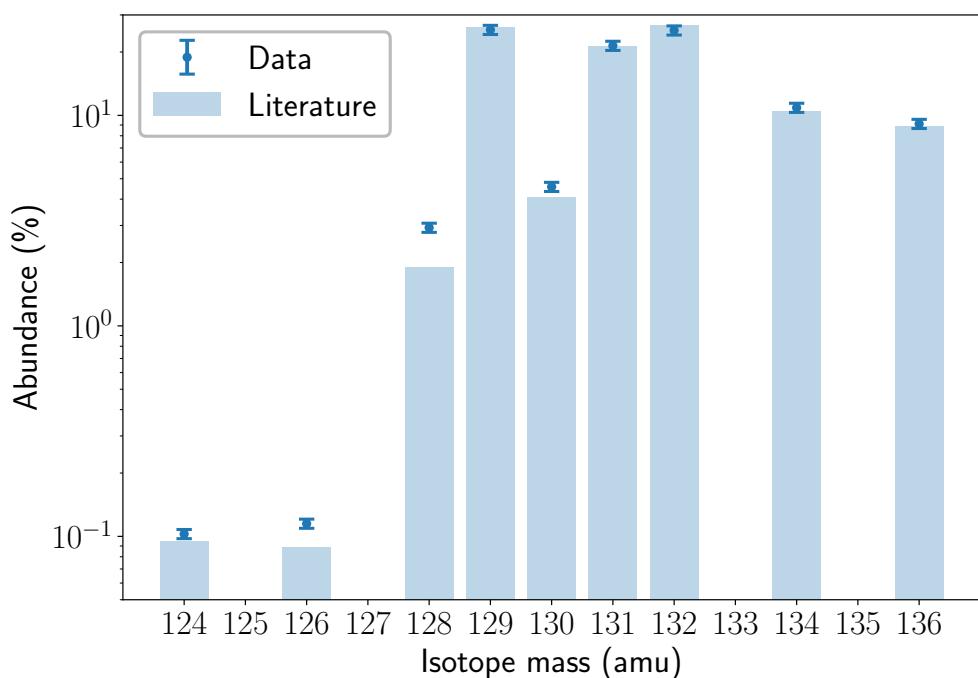


Figure 4.16: The isotope abundances as determined from the fit to the RGA spectrum shown in figure 4.15. All stable isotopes are found. The abundances are in agreement with values found in literature, except for the isotopes 126, 128 and 130, which is likely due to spillover of the large peak at 129 (see text).

PRECISION MEASUREMENTS OF THE SCINTILLATION PULSE SHAPE FOR LOW-ENERGY RECOILS IN LIQUID XENON

Particle detection in liquid xenon is often described in a simplified way: when a particle interacts, it deposits (part of) its energy and the recoiling particle creates excitations and ionizations, which form the S₁ and S₂ signal, respectively. This type of description omits many details of the microscopic physics, but it is usually sufficient for most purposes. Although the S₁ and S₂ yields depend on the particle type, electric field and recoil energy, they are calibrated with neutron and gamma sources. These calibrations, in combination with effective models fit to data (i. e. NEST), directly predict which signals are to be expected from WIMPs and which from backgrounds. In principle, anyone working on a liquid xenon TPC can thus be blissfully unaware of the detailed physical processes. However, a good understanding of the microphysics may help the improvement of future detectors, and is an interesting field of study by its own right.

It is certainly not an easy task to find out what processes are happening in the liquid xenon when a particle interacts. Typical ionization track lengths range from the sub-micrometer scale for nuclear recoils (NRs) up to several micrometers for electronic recoils (ERs) (and up to a few mm for very highly energetic ERs), making it very hard to directly observe the physical extent of the track [44]. Moreover, the finite thermalization radius of ionized electrons and electron diffusion obscures the initial distribution of the ionized electrons. The time scales for the relevant processes are short: typically on the nanosecond scale, but even faster for some of the processes, such as excimer formation, which occurs within picoseconds [51].

Since it is practically impossible to observe the microphysical processes directly, one has to rely on the indirect observables, as illustrated in figure 5.1. The variables are limited to the electric field, the recoil energy and the type of particle (i. e. ER, NR, alpha decay, fission fragments). The latter two variables change the structure

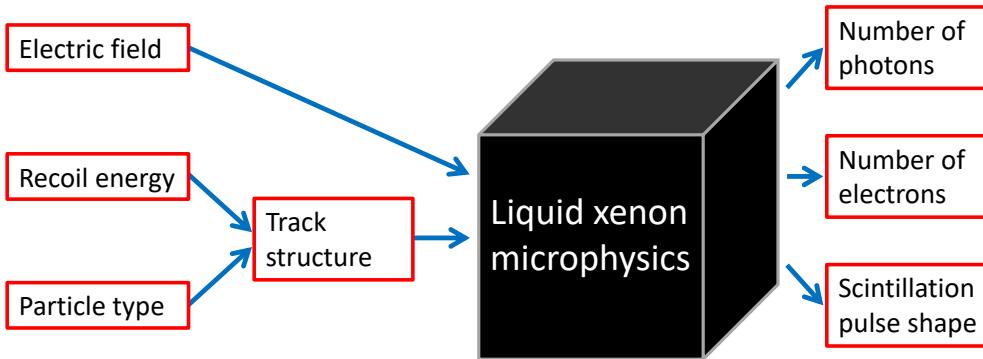


Figure 5.1: Diagram showing the parameters that influence the scintillation and ionization processes (left side) and the accessible observables for liquid xenon TPCs (right side). The scintillation pulse shape provides information that is essential to the knowledge of liquid xenon microphysics.

of the track: the physical position of excitons, electrons and ions, their ratio and the density. Two of the observables from the interaction are the number of photons and the number of electrons, which are simply observed as the size of the S₁ and S₂ signals. From just these signals and their fluctuations, many properties of the interaction can already be derived (see, for example, [48]). However, the scintillation pulse shape provides an important independent handle on the processes and may give valuable input into theoretical models.

Scintillation in liquid xenon comes from the decay of excimer states, which may either be in the spin singlet or triplet state. These have different decay times (approximately 4 ns and 22 ns) and can thus be distinguished by the scintillation pulse shape. The relative occurrence of these states depends on all variables on the left-hand side of figure 5.1. An additional factor in the pulse shape is recombination: if electrons and ions combine, they form excimers in one of the spin states. The delay in light emission due to recombination is observable in some cases; most notably for electronic recoils at high energy and low (or zero) field. The systematic study of the scintillation time profile thus opens a window into the microphysics, in particular recombination, that would otherwise be inaccessible. The paper that is presented in this chapter gives exactly this, showing the energy dependence of the pulse shape for ER and NR at high field, low field (comparable to the XENON1T field) and at zero field (for ER only). The analysis uses a pulse shape simulation that is matched to data, so that the underlying theoretical time dependence of the scintillation is decoupled from detector effects, such as the finite time resolution, pulse smearing due to misalignment and the shape of single-photon pulses of the

PMTs. This greatly improves upon earlier measurements that were taken at zero field only [105–107] or at vastly different energies than those relevant for dark matter searches [51, 53]. In chapter 6, the pulse shape for 511 keV electronic recoils as a function of electric field is investigated.

There is also considerable interest in the scintillation pulse shape for a more practical reason: the prospect of better ER/NR discrimination using pulse shape discrimination (PSD). This method exploits the difference in pulse shapes between ERs and NRs and has been proven to work in liquid argon. Since the difference between the singlet and triplet lifetimes of the scintillation from liquid argon (approximately 6 ns and 1.5 μ s [108]) is large, the discrimination can be based on the relative occurrence of the singlet and triplet states, usually parametrized by the *prompt fraction*, the fraction of all light observed in the first 120 ns of the pulse.¹ Figure 5.2 shows an example of a simulated ER and NR pulse at 50 keV_{ee} and zero field, based on measurements in [108]. There is more light coming from the singlet state for NRs than for ERs, giving a high prompt fraction.

Compared to liquid argon, it is more challenging to apply PSD for scintillation light from liquid xenon. This is chiefly due to the shorter timescales and the smaller difference in decay times for the singlet and triplet state (3 ns and 22 ns for xenon, 6 ns and 1.5 μ s for argon). In addition, the number of available quanta at the energies of interest in dark matter searches is low. Still, measurements at high recoil energy have shown a large difference in the singlet and triplet state population for ER and NR. The effect of recombination further broadens the ER pulse significantly, giving an effective lifetime of 45 ns for recoils at approximately 1 MeV at zero field. This effect further increases the difference of the prompt fraction between the short NR pulse and the wide ER pulse, although it diminishes with applied field.

Recent attempts have shown that PSD is possible in liquid xenon at low recoil energies [54, 60, 79, 109–111]. These mainly focused on obtaining optimum PSD capabilities for the detector that is used, so that the conclusions that are drawn are detector-dependent. In the paper presented in this chapter, we give a description of the pulse shapes that is decoupled from the detector effects, so that a theoretical best-case PSD performance can be calculated if the detector time resolution is known. We find that the pulse shapes for ER and NR at low energies are more similar than expected from high-energy recoils, and that the difference is further reduced with applied electric field, so that the PSD discrimination is not as strong as previously expected. As a definite example, we simulate a XENON1T-like detector with a theoretical time resolution up to 2 ns and compute the performance of

¹ Other window sizes can also be used; the range is usually optimized for PSD performance.

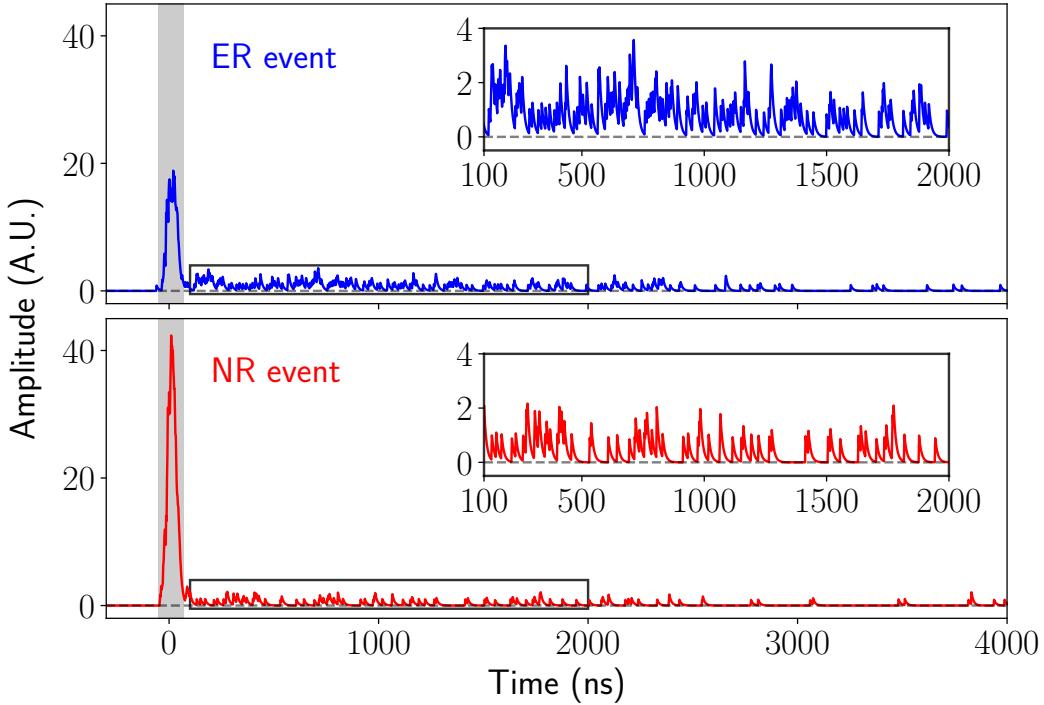


Figure 5.2: Simulated waveforms for a 50 keV_{ee} ER (top panel) and NR (bottom panel) in liquid argon, based on measurements in [108], illustrating the pulse shape difference between an ER and an NR with the same total signal area. A relatively large fraction of the scintillation light comes from the singlet state for NRs, as evidenced by the large initial peak and lower contribution at late times (see insets). The discrimination parameter that is usually used is the relative fraction of the light observed in the first part of the pulse, which is approximately 0.30 and 0.65 for ERs and NRs at this energy, respectively. The shaded region indicates a typical window (120 ns wide) that is used to compute the prompt fraction. Source code available at [10].

the S₂/S₁ method combined with PSD, using a maximum likelihood method for the discrimination. For the regular WIMP search, the improvement is very small, so that increasing efforts to measure the pulse shape better for future detectors is likely unnecessary.

Paradoxically, many of the features that make a TPC more sensitive to very low dark matter interaction cross sections make it worse at PSD. One basic consideration is the TPC size: for large TPCs, the time for a scintillation photon to traverse the TPC is non-negligible with respect to typical scintillation times. For instance, the time to traverse the drift length of XENON1T is approximately 6 ns, compared to only 0.6 ns for XAMS. If there is uncertainty in optical path lengths, for instance, due to reflections, this can cause significant uncertainties in the original pulse shape.

Large TPCs also usually feature large PMTs (giving more transit time spread), long cables and many PMT channels with different time offsets. Finally, the digitizer time resolution chosen for large TPCs is usually 10 ns, which reduces the data rate and volume and has the added benefit of concentrating most of the PMT signal area into one bin, making peakfinding and data reduction by zero-length encoding easier. In spite of these issues, the LUX collaboration has shown PSD ability in the LUX detector in a paper that appeared on the arXiv while the paper in this chapter was in preparation [112]. The LUX analysis also aptly demonstrates the difficulty of using a large TPC for the pulse shape analysis, requiring a detailed light propagation simulation and a photon time reconstruction that is six times smaller than the time bin size. This also results in large uncertainties on the reconstructed pulse shape parameters. Even so, the combined uncertainties of the LUX measurement and the data presented here are not large enough to explain the discrepancy that is found between these measurements (see figure 5.11), so that this remains an open issue. The conclusions from the LUX analysis about the PSD performance appear to agree with the PSD simulation shown in this chapter.

In conclusion, PSD appears to give only small improvements to the NR/ER discrimination ability at a large cost, either in electronics (fast digitizers, large data volume), manpower, or both. This is partly due to the similarity between the pulse shapes at low energy for ERs and NRs and partly due to the limited photon statistics at low energy. A lower drift field, which may be used by future liquid xenon TPCs, could possibly change this, as recombination for ERs becomes slower (see chapter 6) and the number of photons becomes higher. Dark matter models that predict high energy nuclear recoils, such as inelastic dark matter [113], could also benefit from PSD ability. However, the current data in combination with the aforementioned considerations for large TPCs suggests that PSD will likely not be useful for the regular WIMP search using elastic nuclear recoils.

PRECISION MEASUREMENTS OF THE SCINTILLATION PULSE SHAPE FOR LOW-ENERGY RECOILS IN LIQUID XENON

E. Hogenbirk, J. Aalbers, P.A. Breur, M.P. Decowski, K. van Teutem, A.P. Colijn

Journal of Instrumentation **13**, Po5016

Accepted 4 May 2018

Available online at

<https://iopscience.iop.org/article/10.1088/1748-0221/13/05/P05016>

<https://arxiv.org/abs/1803.07935>

Reference [2]

5

5.1 ABSTRACT

We present measurements of the scintillation pulse shape in liquid xenon for nuclear recoils (NR) and electronic recoils (ER) at electric fields of 0 to 0.5 kV/cm for energies <15 keV and <70 keV electron-equivalent, respectively. The average pulse shapes are well-described by an effective model with two exponential decay components, where both decay times are fit parameters. We find significant broadening of the pulse for ER due to delayed luminescence from the recombination process. In addition to the effective model, we fit a model describing the recombination luminescence for ER at zero field and obtain good agreement. We estimate the best performance of a combined S_2/S_1 and pulse shape ER/NR discrimination and show that even with 2 ns time resolution, the improvement over S_2/S_1

discrimination alone is marginal, so that pulse shape discrimination will likely not be useful for future dual-phase liquid xenon experiments looking for elastic dark matter recoil interactions.

5.2 INTRODUCTION

Over the past decade, the field of direct dark matter detection has been led by dual-phase xenon time projection chambers (TPCs) [39, 63, 114]. These experiments are particularly well-suited to investigate weakly interacting massive particles (WIMPs), the leading dark matter candidate. The main challenge for these experiments is to reduce radioactive backgrounds to the level of a few events per year, so that an (almost) background-free environment can be achieved. A crucial part of background rejection is the ability to distinguish electronic recoil (ER) events, caused by background gamma or beta radiation, from nuclear recoil (NR) events, which are expected from WIMP interactions. In dual-phase xenon TPCs, this is usually achieved by using the ratio of the secondary scintillation (S_2) to the direct scintillation signal (S_1), which is a powerful discriminant. Alternatively, there is considerable interest in using the scintillation pulse shape as a particle identification method. This method, called pulse shape discrimination (PSD), has been applied in, e.g., liquid organic scintillators [115] and in liquid argon, where it provides the main discrimination against the sizable ^{39}Ar ER background [116, 117]. In xenon, the scintillation pulse shape has been shown to depend on ionization density (and therefore particle type and recoil energy) [51] and electric field [53, 54], so that pulse shape discrimination is in principle possible. However, due to the short timescales ($\sim 4\text{ ns}$ and $\sim 22\text{ ns}$ for the two exponential decay times for alpha particles and fission fragments [51]) compared to typical photomultiplier tube (PMT) responses, pulse shape discrimination is challenging in liquid xenon detectors. The method has nevertheless been successfully applied as a particle identification method, either in combination with the S_2/S_1 ratio [79, 109], or by itself, i.e., in single-phase detectors [54, 60, 110, 111]. In general, a single pulse shape parameter or an effective distribution (such as a single exponential) is used to describe the pulse shape. Recently, the LUX collaboration measured the pulse shape by using a sum of two exponential functions with a fixed time constant [109]. This distribution correctly describes the decay of the two excited states, but neglects any time delay in their production, caused primarily by recombination of electron-ion pairs. This delay is non-negligible for $\mathcal{O}(\text{MeV})$ ERs, but the influence of recombination on the pulse shape was assumed to be minor based on extrapolations from

measurements at high energy, as measurements at low energy are lacking [118].

This paper aims to directly determine the scintillation pulse shape for $ER < 70 \text{ keV}_{ee}$ and $NR < 15 \text{ keV}_{ee}$ at electric fields up to 0.5 kV/cm using measurements from XAMS, a dual-phase TPC setup described in [1]. All measured energies are reported relative to gamma-ray (ER) calibrations, which we make explicit by using the unit keV_{ee} (electron-equivalent) for both ER and NR. We use a sum of two exponential distributions to describe the pulse shapes, where all parameters are allowed to vary to capture the effect of recombination. Additionally, we use a recombination model to describe the pulse shape for ER at zero field. Using the pulse shape model extracted from the data, we calculate the theoretical improvement in discrimination when combining PSD and S_2/S_1 discrimination in future large-scale TPCs searching for elastic WIMP recoils.

5.3 THE PROCESS OF SCINTILLATION

When ionizing radiation energy is deposited in liquid xenon, the energy is transferred to atomic excitation, ionization and heat (see figure 5.3). Excited xenon atoms combine with stable xenon atoms to form excited xenon dimers (Xe_2^*), which can be in either a spin singlet or triplet state. The lifetimes of these states, τ_s and τ_t , are measured at 2 to 4 ns and ~ 22 ns, respectively [51]. The intensity of the scintillation light from the decay of one of these states is characterized by

$$I_d(t, \tau) = \frac{1}{\tau} \exp\left(\frac{-t}{\tau}\right), \quad (5.1)$$

with τ the lifetime of the state. If we define the fraction of photons from the singlet state as f_s , we can write the time dependence of the direct scintillation as

$$I(t, \tau_s, \tau_t, f_s) = f_s I_d(t, \tau_s) + (1 - f_s) I_d(t, \tau_t). \quad (5.2)$$

In addition to the direct excitation, excited dimer states are formed by electron-ion recombination. In this process, ionized electrons recombine with xenon ions to form doubly excited xenon atoms Xe^{**} , which finally form xenon dimer states.

The scintillation process in liquid xenon depends on the ionization density, which to first order is proportional to the linear energy transfer (LET). This has been studied extensively by using different LET tracks caused by (in order of increasing LET) high-energy ($\mathcal{O}(\text{MeV})$) ERs [52, 53, 119], low energy ($< 100 \text{ keV}$) ERs [105–107, 109], nuclear recoils [105, 106, 109], alpha particles [51, 120] and fission frag-

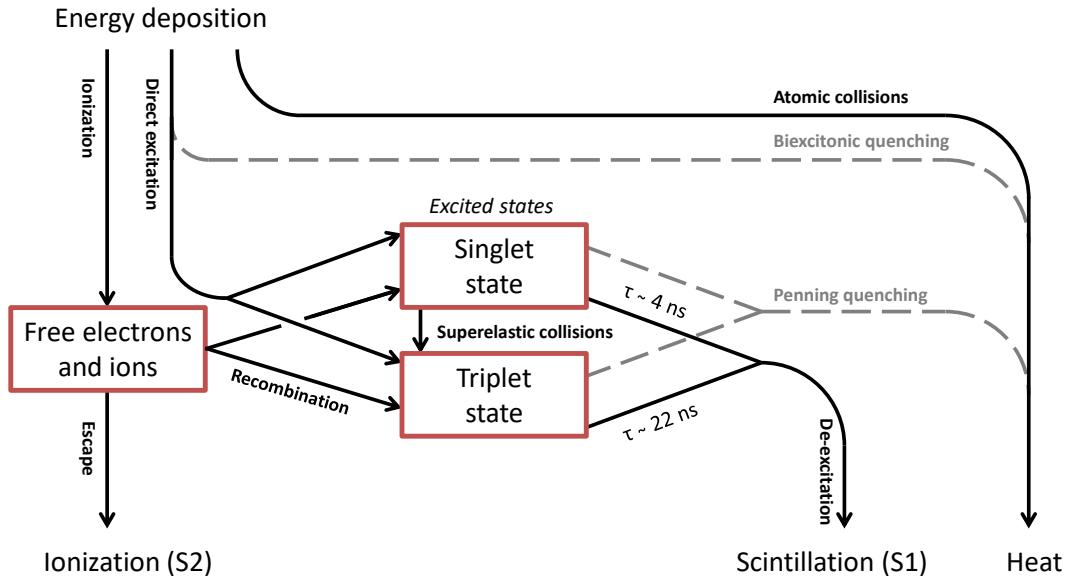


Figure 5.3: After an energy deposition in liquid xenon, different processes lead to ionization, scintillation and heat. The processes indicated by the gray dashed lines are most relevant for recoils resulting in high ionization density.

5

ments [51]. First of all, the timescale of recombination increases for low LET tracks. For high-energy ERs at zero field, the timescale is comparable to the decay times, so that the scintillation profile consists of the cumulative effect of dimer states forming and decaying.² In this case, a decay curve different from the distribution given in eq. (5.2) has been found, with a single apparent decay constant of (33 ± 1) ns [52, 119] or ~ 45 ns [53] and significant broadening of the peak intensity. Secondly, the scintillation yield decreases with LET, as there are quenching mechanisms that become more relevant for high ionization density tracks (indicated by the gray dashed lines in figure 5.3). In biexcitonic quenching, two excitons are quenched before the excited molecular states are formed. In Penning quenching, two excited molecular states collide to produce one excited state and one of the states dissociates to the ground state. For nuclear recoils and fission fragments, the yield decreases further due to elastic energy transfer to xenon atoms, resulting in energy dissipation through atomic motion (heat) instead of scintillation. Finally, the fraction of photons coming from the singlet state f_s increases with LET. Measured values are 0.05 for $\mathcal{O}(\text{MeV})$ electrons [52], 0.04 [109] and 0.045 to 0.145 [107] for low-energy (< 100 keV) ERs, 0.21 for NRs [109], 0.31 ± 0.05 [51] and ~ 0.6 [120] for alpha particles and 0.62 ± 0.08 for fission fragments [51]. It has been suggested [51] that this is

² As commonly reasoned (or so this is commonly supposed), a recombination efficiency for other recoils (not electronic) rarely deviates significantly.

due to superelastic collisions, where a singlet excited state collides with a free electron and forms a triplet state. As this requires a free electron, this process is more likely to occur if recombination is slow, which is the case for low energy transfer tracks.

The applied electric field influences the pulse shape due to its effect on free electrons. If an electric field is applied, the ionization electrons are drifted away from the interaction site, separating the electrons from the ions that are pulled to the opposite side (but at a much lower drift speed [47]), so that the recombination luminescence is suppressed. It should be noted, however, that even at high field of 4 kV/cm an apparent lifetime of (27 ± 1) ns was found for electronic recoils [52], which disagrees with the ~ 22 ns found for alpha and fission fragment recoils.

5.3.1 Recombination model

The time dependence of recombination luminescence in the absence of an electric field can be described following a model by Kubota et al. [53]. This model assumes a region of uniform ionization density, where a fraction of the electrons η escapes at $t = 0$ and the timescale of recombination is governed by the recombination time T_R . The derivation of the scintillation intensity time dependence is given in Appendix A.3 and yields

$$I_r(t, \tau, T_R, \eta) = A \exp\left(\frac{-t}{\tau}\right) \times \int_0^t n_e(t') n_{ions}(t') \exp\left(\frac{t'}{\tau}\right) dt', \quad (5.3)$$

with A a normalization factor, τ the decay time of the excimer state, T_R the recombination time and η the probability of electrons to escape the recombination region. The dependence on T_R and η is implicit through the solution of differential equations describing the electron and ion densities n_e and n_{ions} . For a general scintillation pulse, the time dependence in its most general form will be a combination of four terms, corresponding to the two states in the direct decay from excitations and the same two states formed by recombination:

$$I(t, \tau_s, \tau_t, T_R, f_s, f_R, \eta, f_s^R) = (1 - f_R) \cdot (f_s I_d(t, \tau_s) + (1 - f_s) I_d(t, \tau_t)) + f_R \cdot (f_s^R I_r(t, \tau_s, T_R, \eta) + (1 - f_s^R) I_r(t, \tau_t, T_R, \eta)), \quad (5.4)$$

where the integral of I , I_d and I_r is normalized to 1 and f_R , f_s and f_s^R are the fraction of photons from recombination, the fraction of direct scintillation photons coming from the singlet state and the fraction of recombination photons coming

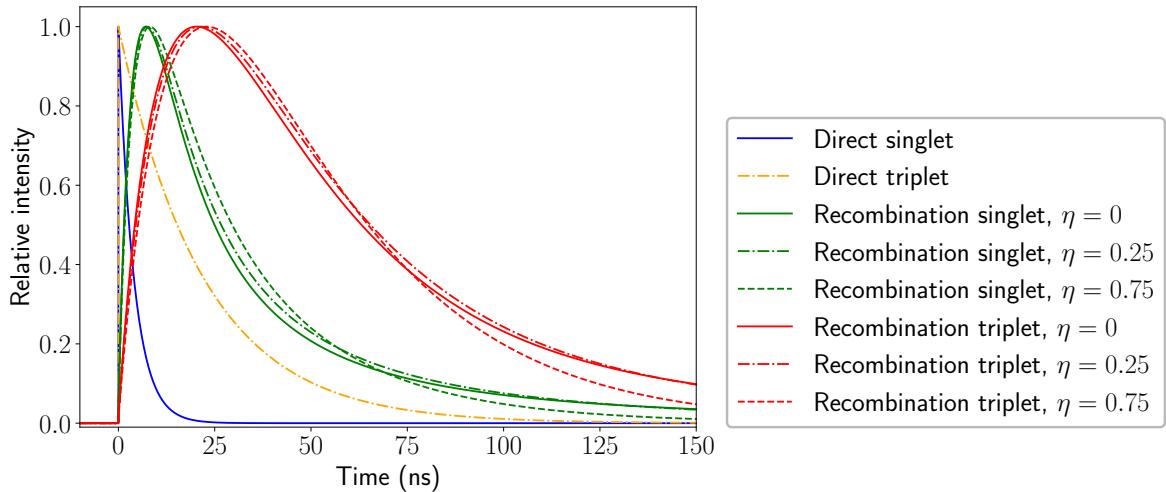


Figure 5.4: Example curves of the scintillation time dependence components in eq. (5.4). The time constants used are $\tau_s = 4 \text{ ns}$, $\tau_t = 22 \text{ ns}$ and $T_R = 25 \text{ ns}$. All curves are scaled such that the maximum is 1.

from the singlet state, respectively. Figure 5.4 illustrates the time dependence of the scintillation intensity for the components of eq. (5.4) for a specific choice of parameters.

5.3.2 Effective model

To avoid the complexity of including a recombination model and the uncertainty of the physical mechanisms, we introduce the effect of recombination by altering the lifetimes of the singlet and triplet states. In this case, we follow the model of eq. (5.2), with the difference that the lifetimes are interpreted as effective lifetimes, i.e.,

$$I(t, \tau_s^{\text{eff}}, \tau_t^{\text{eff}}, f_s) = f_s I_d(t, \tau_s^{\text{eff}}) + (1 - f_s) I_d(t, \tau_t^{\text{eff}}). \quad (5.5)$$

While eq. (5.4) provides a model that may be motivated better physically, it comes at the cost of four extra parameters. Therefore, the effective lifetime model can be chosen over the full model if only a working description of the pulse shape is required, or if the contribution from recombination is negligible. In this paper, we use the model in eq. (5.5) as an effective model and apply the model from eq. (5.4) to the data at zero field only.

5.4 MEASUREMENTS

5.4.1 Data acquisition and calibration

The data was collected with the XAMS setup, a dual-phase liquid xenon TPC. The active volume of the TPC has a cylindrical shape with a diameter of 44 mm and a total effective length (as defined by the distance between the cathode and gate meshes) of 100 mm, giving a total xenon content of 154 cm³ (434 g at -90°C). The walls of the TPC are made of Teflon to maximize the light detection efficiency. The active volume is viewed from below and above by two 2-inch circular PMTs (Hamamatsu type R6041-406 [98]). For all the measurements described here, a voltage of 3 kV is applied between the anode mesh and the gate mesh (5 mm separation), extracting the electrons from the liquid into the gas and providing the proportional scintillation. The waveform data from both PMTs was collected by a CAEN V1730D digitizer with a time resolution of 2 ns. The high sampling rate, in combination with the relatively small geometry and the low transit time spread of the PMTs (~ 0.75 ns [98]) makes the setup well-suited for pulse shape measurements. Further details of the setup can be found in [1, 82].

Figure 5.5 shows the examples of the summed PMT waveforms for an NR and ER event. The trigger was set at a level of 7 mV, corresponding to 3.0 ± 0.2 times the average single photoelectron (SPE) amplitude for both channels, and the coincidence time window was set to the maximum of the on-board trigger of 120 ns. This allowed to trigger on small S₂s, which is required for low-energy recoils. Data processing followed the method outlined in [1]. Peaks were classified based on their width, which is defined as the time range containing the center 50 % of the area in the peak. The peaks were classified as S₁ if the width was less than 60 ns. For S₂s, we required a width of at least 100 ns, and an area of at least 100 times the average single photoelectron area. In addition, for both S₁ and S₂ we required both PMTs to contribute to the signal.

An optical fiber connected to a blue pulsed LED was installed, allowing in-situ PMT calibration to SPE. The gain of the PMTs was used to convert the area of the PMT signals to the number of photoelectron-equivalent, or p.e. (photo-electrons). We use the photopeak in ¹³⁷Cs and ²²Na calibration data to correct the S₁-signal for the z-dependent light detection efficiency (cS₁) and the S₂ for charge loss during electron drift (cS₂). For the S₁, we fitted a second-order polynomial to the S₁ as a function of z. All S₁ values were then scaled according to this fit and normalized to the volume average. The S₂ correction was determined by an exponential

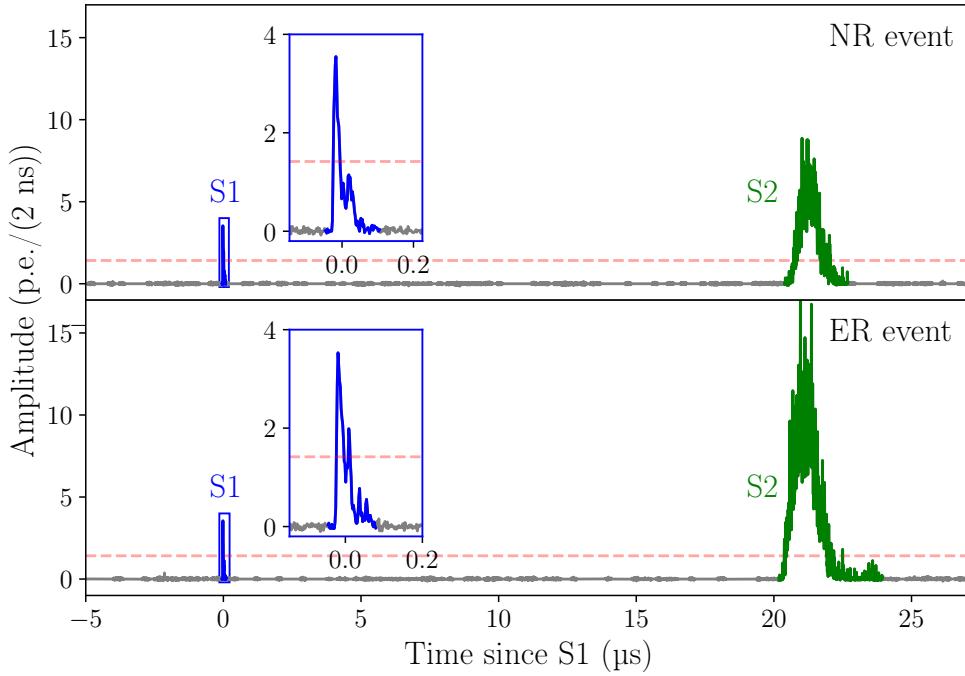


Figure 5.5: Examples of the summed waveforms for an NR event (top panel) and an ER event (bottom panel) with comparable drift time. The inset shows a zoom of the S1 signal. The dashed line shows the sum of the trigger levels for the two channels (note that a coincidence of both channels is also required). The energy, reconstructed using a combination of S1 and S2, is 9.5 keV_{ee} and 11.3 keV_{ee} for the NR and ER event, respectively.

fit to the photopeak as a function of drift time, where the normalization was to the point where the drift time is zero. The electron lifetime was determined to be (0.81 ± 0.04) ms. We calculate the recoil energy from

$$E = W \left(\frac{cS1}{g1} + \frac{cS2_b}{g2} \right), \quad (5.6)$$

assuming a W -value of (13.7 ± 0.2) eV and using only the S2 in the bottom PMT ($cS2_b$). The photon gain $g_1 = (0.10 \pm 0.02)$ p.e./photon and electron gain $g_2 = (4.7 \pm 0.7)$ p.e./e⁻ were obtained from the photopeak positions of ¹³⁷Cs and ²²Na, following the method described in [48].

For measurements with long timescales, we found that the S2 gain slowly decreased, while the S1-signal was unaffected. The gain decreased more rapidly during high-rate calibrations, and reset after cycling the anode power or temporarily raising the liquid level above the anode. After ruling out faulty cabling or equip-

ment, we attributed this to charge-up of floating Teflon dust particles, which we found in the detector after the measurements. These were probably produced while drilling a hole in the TPC wall to install the optical fiber. We corrected for the S_2 decrease by fitting the electronic recoil band in half-hour time slices; the correction is typically 15 % and at most 30 % (after 3 hours).

To collect the electronic recoil data, we acquired data with ^{137}Cs (662 keV gamma), ^{22}Na (positron decay, where we trigger on the 511 keV back-to-back gammas using an external NaI(Tl) detector) and background radiation with 0.5 kV/cm field, and with ^{137}Cs only with 0.1 kV/cm and zero field. We combined all three datasets (^{137}Cs , ^{22}Na and background) at 0.5 kV/cm into one ER dataset to increase statistics. For the nuclear recoil data, a $^{241}\text{AmBe}$ neutron source with a neutron intensity of 1.3×10^3 n/s was used. A lead shield of 25 mm was used to attenuate the gamma radiation from ^{241}Am . In addition, a cylindrical lead shield with a thickness of 12 mm was placed outside the outer cryostat vessel to reduce the influence of background radiation for all measurements.

5.4.2 Data selection

For all data taken at nonzero drift fields we require exactly one S_1 and one S_2 , reducing pileup and double scatter events, respectively. To further reduce the number of double scatter events, we remove events with a wider S_2 signal than expected based on a diffusion model fit to calibration data, yielding a diffusion constant of $(36.9 \pm 0.6) \text{ cm}^2/\text{s}$ and $(15.5 \pm 0.6) \text{ cm}^2/\text{s}$ for a drift field of 0.1 kV/cm and 0.5 kV/cm, respectively. These cuts combined give an estimated double scatter rejection resolution of 2 mm in depth. Furthermore, we cut events occurring within 1 ms of a previous event. This removes events that triggered on a tail of single-electron S_2 s originating from photoionizations caused by large events. Finally, we cut events within the top and bottom 5 mm, giving an effective mass of 390 g of liquid xenon in the TPC.

Only an S_1 -signal is generated for the data taken with zero field. Nevertheless, the z-coordinate of the interaction can be estimated using the distribution of light across the two PMTs, which we related to the z-position using calibration data taken at nonzero drift field. We cut on the z-coordinate to remove energy depositions above the gate, in the gas layer or below the cathode. In addition, we used the z-coordinate to correct for the position-dependent light detection efficiency described in section 5.4.1. This method breaks down for small S_1 s ($\lesssim 40$ p.e.) because of statistical fluctuations in the light distribution. We therefore set a lower energy

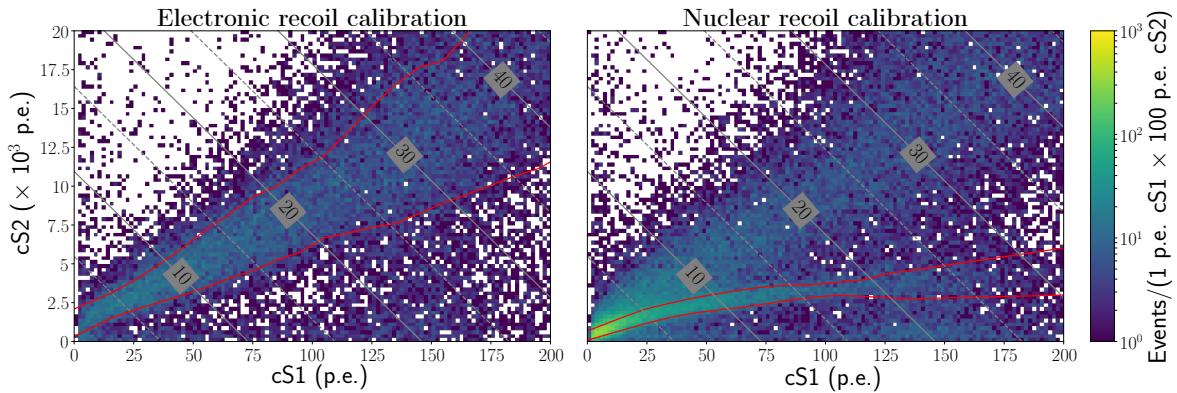


Figure 5.6: Event distributions in $(cS_1 - cS_2)$ -space for electronic recoil (ER, left panel) and nuclear recoil (NR, right panel) calibrations. Note that the density color scale is logarithmic. The red lines indicate the fitted ER and NR 1σ bands. The gray lines indicate the reconstructed recoil energy, with the labels in keV_{ee} . In the NR calibration, the ER band is of a similar magnitude as observed in background data taking.

limit of 10 keV_{ee} . The S_1 classification was relaxed to include all peaks where both channels contribute. Events with more than one S_1 -peak were still discarded, and the holdoff of 1 ms from the previous event was also still applied to ensure a stable baseline condition.

5.4.3 Electronic and nuclear recoil selection

The electronic and nuclear recoil events were selected based on the charge-to-light ratio cS_2/cS_1 . Figure 5.6 shows the event distribution for the combination of all ER data (left panel) and the NR data (right panel) for a field strength of 0.5 kV/cm . In the nuclear recoil calibration, a distinct NR band appears below the ER band. As the rate of nuclear recoils is relatively low, an ER band is still present in the nuclear recoil calibration dataset due to background radiation. We observe some leakage of events below the bands, which we attribute to events at the edges of the detector, where there is incomplete charge collection and therefore a reduced S_2 size. This is consistent with the lower S_2 width observed for these events, coming from the shorter gas gap due to the capillary effect of liquid xenon.

The ER and NR bands are fit by slicing the data in cS_1 and performing an unbinned Gaussian fit to the distribution of cS_2 in each slice. The fit range is determined iteratively; an initial fit is done, after which the fit range is adjusted to $\mu \pm 1\sigma$ of this fit. This procedure is repeated until convergence is reached. The resulting

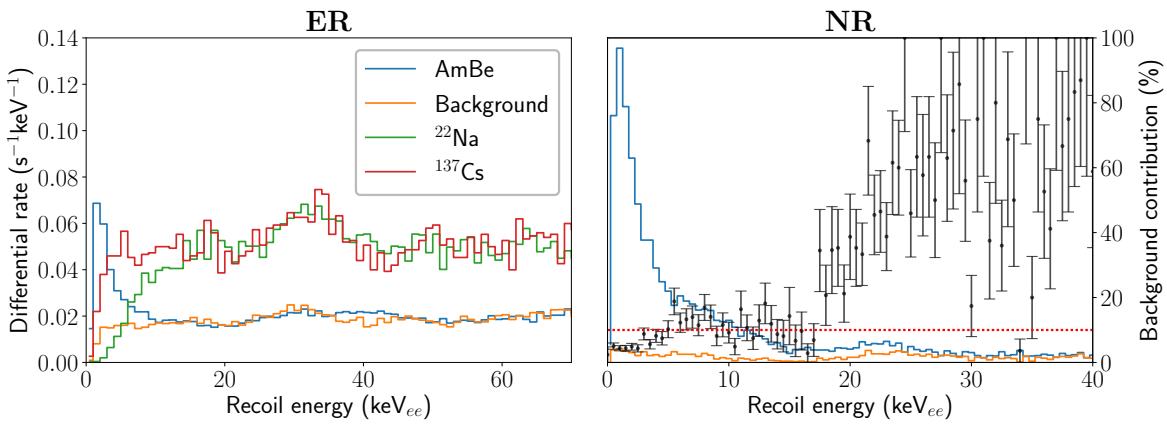


Figure 5.7: The energy spectrum for the ER (left panel) and NR (right panel) selection shown in figure 5.6 for different sources. **Left panel:** The AmBe and background ER rate are consistent above 10 keV_{ee} , where the bands are separated, except for a small excess at $E = (40 \pm 2) \text{ keV}_{ee}$ from the inelastic transition of ^{129}Xe . A broad peak at $(31 \pm 2) \text{ keV}_{ee}$, caused by xenon X-rays, is visible in all ER spectra. The low-energy rolloff of the ^{22}Na spectrum comes from the lower efficiency of the S1-only trigger. **Right panel:** Energy spectrum for all events within the NR selection in AmBe and background data. The gray points and error bars indicate the ratio between background and AmBe data, which gives the ER contribution in the NR dataset. These points have been smoothed by a three-point moving average filter to reduce the statistical fluctuations at energies $\gtrsim 25 \text{ keV}_{ee}$. At energies below 15 keV_{ee} , the background contribution is at roughly 10 %, as indicated by the dotted red line. The broad peak observed in neutron and background data between 15 and 35 keV_{ee} is due to leakage of the $(31 \pm 2) \text{ keV}_{ee}$ xenon X-ray line below the ER band.

ER and NR bands are indicated in the left and right panel of figure 5.6, respectively. ER and NR data selection is consequently performed by cutting all the events outside the $\pm 1\sigma$ lines indicated in figure 5.6.

Figure 5.7 (left) shows the ER spectra for three different sources and background data. The rate of events in the ER band in the AmBe dataset is consistent with the background rate (which we can confirm above 10 keV_{ee} where the separation between the bands is good enough), except for a small excess of events at $E = (40 \pm 2) \text{ keV}_{ee}$. This is attributed to the neutron-induced inelastic transition of ^{129}Xe , which gives a gamma ray at 39.6 keV. The appearance of this excess also confirms our calibration of g_1 and g_2 at the low energies of interest. Another broad structure was found in all the datasets at $(31 \pm 2) \text{ keV}_{ee}$. We believe this is due to an X-ray energy transition from xenon at 29.8 keV [121]. If a photoelectric absorption interaction occurs in an insensitive region of the detector, an X-ray can result and penetrate

into the sensitive volume. This hypothesis is strengthened by the observation that the excess increases in intensity with higher ER rate. Also, we find significantly more leakage to low S_2 values at this energy, indicating that the events are concentrated at the edges of the detector, where there is incomplete charge collection. This is consistent with the low range (~ 0.4 mm) of X-rays at this energy [49]. For the background and ^{137}Cs data (which have the same trigger), we evaluate the low-energy acceptance by extending the observed flat spectrum between 5 and 20 keV_{ee} to lower energies. This gives an acceptance of 0.7 ± 0.1 at 2 keV_{ee} and reaching >0.9 at 3 keV_{ee} . For the ^{22}Na data taking, the trigger was set to a triple coincidence with an external NaI(Tl) detector. This requires a trigger on the S_1 -signal rather than the much larger S_2 -signal and causes a clear reduction in the trigger efficiency at low energies, as evidenced by the rate decrease at low energies.

The NR energy spectra for the $^{241}\text{AmBe}$ and background datasets are shown in figure 5.7 (right). The gray points shows the relative contribution of background events within the NR selection. At low energies, the ER and NR bands overlap, so that the background rate in this region increases. However, since the majority of NR events occurs at low energy (due to the effect of kinematics and the energy spectrum of the $^{241}\text{AmBe}$ neutrons), the percentage of background events is low. At higher energies, the background comes from events that are below the ER band due to incomplete charge collection. The background rate increases at ~ 17 keV_{ee} due to the earlier described line feature with significant leakage below the ER band. We therefore limit the analysis of NR data to < 15 keV_{ee} . We limit all analysis of electronic recoil data to 70 keV_{ee} (roughly corresponding to 400 p.e. cS_1 and 300 000 p.e. in cS_2), as ADC saturation starts appearing above this energy.

5.5 MONTE CARLO MODEL

5.5.1 Single photoelectron pulse model

Our S_1 simulation uses a single photoelectron pulse shape model based on LED calibration measurements using an oscilloscope (Keysight DSO-S 254A) with a time resolution of 0.1 ns. Figure 5.8 shows the average of single photoelectron waveforms, aligned on the point where 10 % of the area is reached. We rebin the model to 2 ns and apply a 250 MHz low-pass frequency filter to correct for the lower time resolution and analog bandwidth of the CAEN V1730D digitizer used in the main measurements. The resulting pulse model is shown in figure 5.8. We compare the rebinned model to LED calibration data taken at cryogenic temperatures with

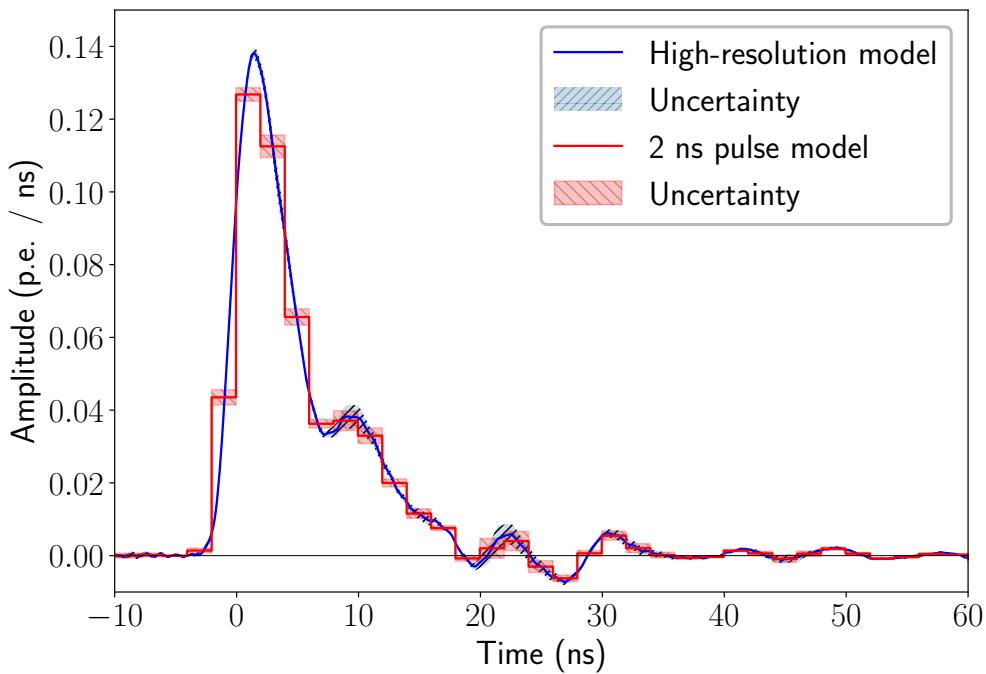


Figure 5.8: Single photoelectron pulse shape model, measured in high resolution (blue, smooth line), and rebinned in 2 ns bins (red histogram) as used in the model. The specific form of the single photoelectron waveform depends on the position of the high-resolution pulse within the 2 ns bins. The uncertainty on the coarse single photoelectron waveform is taken as the difference between the shape obtained from the rebinned high-resolution calibration and the calibration taken in liquid xenon. This uncertainty is indicated by the bands around the waveforms.

the V1730D digitizer by computing the average single photoelectron pulse shapes, aligned to the maximum in the sample. The difference between these is taken as a systematic uncertainty in the main analysis.

The single photoelectron pulse shape has a 5 ns width (FWHM) and shows a \sim 90 MHz oscillation (*ringing*) consistent with standing waves in the \sim 2 m long cables between the PMTs and the DAQ, presumably due to an impedance mismatch between the PMT bases and the DAQ. The bottom PMT shows a lower oscillation frequency consistent with its longer cable length.

5.5.2 Average pulse shape

Since our analysis focuses on low-energy recoils, the number of photons in each individual pulse is low. We therefore use the average pulse shape to fit the measured

data to the scintillation model. The average pulse shape model is constructed in the following way. A large sample of photon production times with a distribution according to the scintillation model (eq. (5.4) or eq. (5.5)) is generated. The times are smeared with a normal distribution with mean 0 and a standard deviation which we define as the detector time resolution σ_{det} . The PMT channel for each photon is randomly sampled from a binomial distribution according to the observed light distribution in measured data (with probability of the photon being in the top PMT of 0.28 ± 0.13). The area of each single-photon pulse is sampled from the measured area distribution in the single photoelectron gain calibration for the PMT that it was observed by. The array of photons is then split into S1 clusters, where the areas of the S1s are taken from the area distribution observed in measured data (within the energy selection). The S1 pulse shapes are calculated as the sum of the single photon pulse shapes, where the normalization of each single-photon pulse is provided by the sampled area and the time shift is given by the photon arrival time. We use the high time resolution model described in section 5.5.1, rebinned to 0.2 ns for ease of computation, as the single-photon pulse model. The resulting pulse shape is then shifted by a random integer up to 10 of 0.2 ns bins to simulate the random alignment of the pulse within the digitizer time bins and rebinned to 2 ns. The average pulse shape is computed as the average of all normalized pulse shapes, aligned on the point where 10 % of the area is found.

5.5.3 Uncertainties

We compute the goodness of fit $\chi^2 / n_{\text{d.o.f.}}$ from

$$\frac{\chi^2}{n_{\text{d.o.f.}}} = \frac{1}{n_s - 1} \sum_i \frac{(y_{\text{data},i} - y_{\text{model},i})^2}{\sigma_i^2}, \quad (5.7)$$

where i runs over all 2 ns samples in the pulse that are within the fit range, n_s is the number of samples in the waveform (ranging from 58 to 76 as the fit range is changed, see section 5.5.4) and y and σ are the average waveform amplitude and its uncertainty. The average waveform uncertainty we compute is built up of several parts. The dominant contribution comes from the single photoelectron pulse shape model described in section 5.5.1. We compute the corresponding uncertainty on the average pulse shape by adding the uncertainty on the single photoelectron pulse shape for all photons in quadrature per S1 waveform. A second contribution of the uncertainty is the statistical uncertainty due to the finite number of S1 waveforms in the data sample. We compute this by resampling the data waveforms 250 times

and recomputing the average waveform, where we define the uncertainty as the per-sample standard deviation. A third source of uncertainty comes from the distribution of photons across both PMTs, which varies stochastically in the data with changing light detection efficiencies at different interaction positions. We therefore recompute the average waveform with light fraction in the top PMT of 0.28 ± 0.13 (corresponding to 1σ) and take the variation on the average waveform as a systematic uncertainty. Finally, we add a constant uncertainty of 10^{-4} in units of fraction of amplitude, corresponding to approximately 0.1 - 0.2 % of the maximum pulse amplitude, to account for any unmodeled uncertainties (such as the influence from randomly distributed noise hits or dark counts in the waveform). This uncertainty is subdominant in all parts of the waveform, but stabilizes the fit if all the other uncertainties fluctuate downward (causing very large values of $\chi^2/n_{\text{d.o.f.}}$ because of the $1/\sigma^2$ dependence in eq. (5.7)). All these uncertainties are added in quadrature to yield the per-sample uncertainty on the average waveform.

5.5.4 Fit procedure

We produce the average pulse shape for various parameter combinations of both scintillation models (eq. (5.4) or eq. (5.5)) on a parameter grid. For the exponential fit, the fit parameters are f_s , τ_s^{eff} , τ_t^{eff} and the detector resolution σ_{det} and the grid spacing is 0.005, 0.25 ns, 0.25 ns and 0.125 ns, respectively. For the recombination model, some of the parameters are fixed (see section 5.6.2), leaving T_R , η , f_R and f_s^R with a grid spacing of 1 ns, 0.05, 0.025 and 0.05, respectively. The best-fit parameters are taken as the point where the minimum value of $\chi^2/n_{\text{d.o.f.}}$ is reached. Figure 5.9 shows an example for one of the waveform fits.

We validate the fit procedure by generating a sample of waveforms with known parameter values and then fitting this in the same way as done with measured waveforms. This is done for each energy bin in each dataset with the same number of S1 waveforms as observed in measured data. In this way, we simultaneously evaluate the statistical uncertainty, which we define as the standard deviation of the parameter around the fit point.

There are also systematic uncertainties coming from the freedom in the definition of the fit range and possible errors in the estimate of the uncertainty of the average waveform. We estimate the effect of the fit range by varying the left and right boundary within -14 ns to -6 ns and 110 ns to 140 ns (with respect to the 10 % alignment point). To quantify the uncertainty of the fit parameters coming from the uncertainty estimate on the average waveform, we vary the constant uncertainty

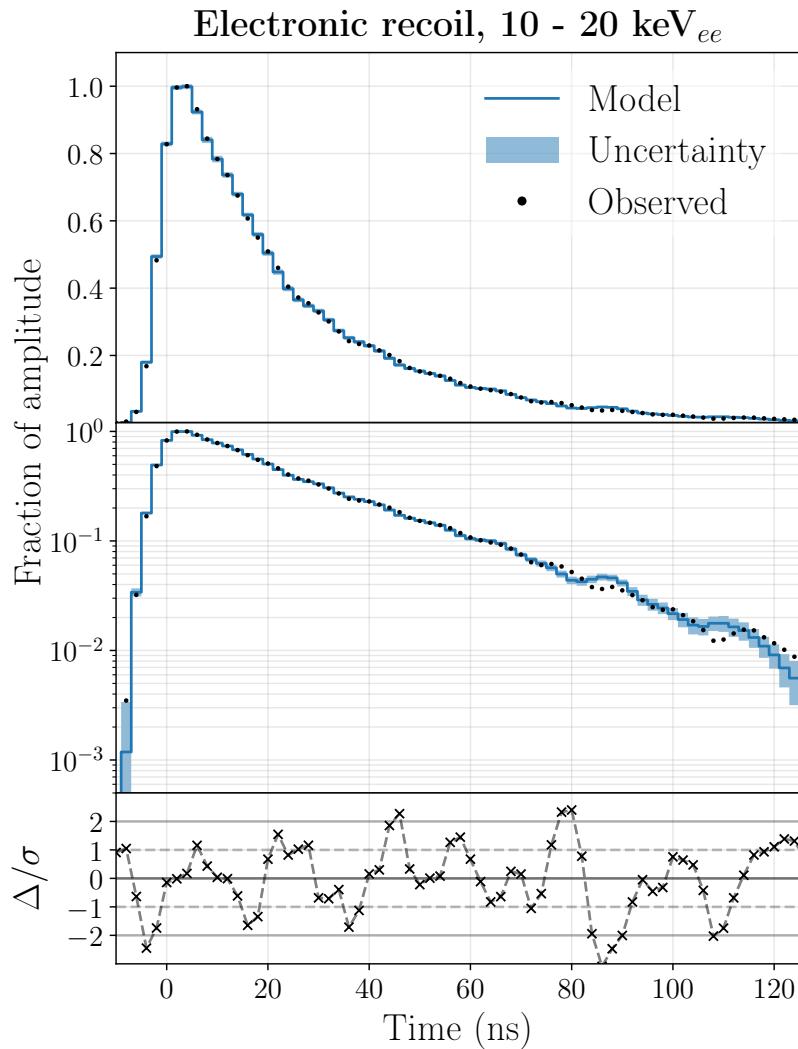


Figure 5.9: Example of an average waveform shape for electronic recoil data between 10 and 20 keV_{ee}, taken at 0.5 kV/cm, shown on a linear scale (top panel) and logarithmic scale (middle panel). Measured data is indicated by the black dots, while the blue line and the shaded blue band show the best-fit simulated pulse shape and its uncertainty. The bottom panel shows the residuals, normalized by the uncertainty. An oscillatory behavior is seen for both the data and the model, coming from the ringing observed in the single photoelectron waveforms (see figure 5.8). The value of $\chi^2/n_{\text{d.o.f.}}$ for this fit is 1.44.

described in section 5.5.3 from 0 to 2×10^{-4} . These three parameters are all varied simultaneously and the best-fit point is recomputed. We take the uncertainty on the best-fit point as the standard deviation of the parameters.

The detector resolution σ_{det} is a quantity that depends on the setup, hence it should be a constant parameter independent of the measurement. Nevertheless, we allowed this quantity to vary for all exponential fits. Figure 5.10 shows a com-

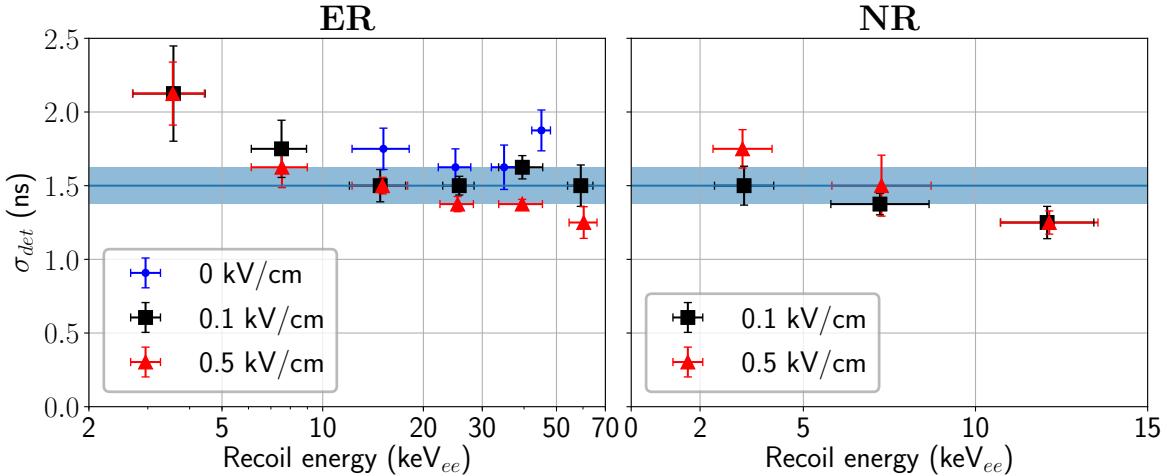


Figure 5.10: The detector resolution σ_{det} for all ER data (left panel) and NR data (right panel). The values shown here were obtained by applying an exponential fit (eq. 5.5) with a Gaussian smearing of the photon times with standard deviation σ_{det} . As the detector resolution is a detector quantity, we fix the value of this parameter to 1.5 ns for all fits. We estimate the uncertainty to be 0.1 ns, and vary σ_{det} with one step in the parameter grid (± 0.125 ns, as indicated by the blue band) to evaluate the effect on the other fit parameters.

pilation of the values of σ_{det} . From this, we estimate the detector resolution is (1.5 ± 0.1) ns. This is in agreement with the expected resolution given the PMT transit time spread of ~ 0.75 ns [98] and the influence of photon travel time due to reflections. As the length of the TPC is approximately 10 cm, we expect typical photon path differences to be of the same order, giving an additional spread of ~ 0.5 ns.³ For the rest of the analysis, the detector resolution is fixed at 1.5 ns. We determine any uncertainty on the other fit parameters by varying σ_{det} by 0.125 ns, corresponding to one step in the parameter grid spacing of σ_{det} .

5.6 RESULTS

5.6.1 Double-exponential fit

The data taken consists of five measurement series; two for the two different fields for NR data and three for the ER data that includes data taken at zero field. The

³ Note the effective speed of light in liquid xenon is lower due to the index of refraction of 1.69 ± 0.02 [122].

data is binned in energy as obtained from the combined energy scale for all the points with nonzero field and from the cS1 in the case of zero field. The value and the error of the energy for each bin is set to the mean and standard deviation of the energy in the bin. Figure 5.11 shows the fit values for a double-exponential fit on ER (left panels) and NR (right panels) data, at 0.5 kV/cm and 0.1 kV/cm and including zero field for ER. The bottom scales show the measured recoil energy, while the top scales show the corresponding electronic LET (obtained using the ESTAR [123] and the SRIM [124] models for ER and NR, respectively), as given in [78]. The recoil energy for nuclear recoils was calculated using the quenching factor parametrized in [125]. The shaded bands in the figures show results from Akimov et al. (blue band) [105] and Ueshima (green hatched band) [106], both using data at zero field and a single exponential fit, measurements from XMASS (orange band) [107] at zero field and results from LUX (red bands) [109], where the light shaded bands indicate the systematic error due to uncertainties in the optical model. The field for the LUX data was 0.41 kV/cm, comparable to 0.5 kV/cm as used for our high-field measurements.

The top two panels show the effective triplet lifetime τ_t^{eff} as a function of energy for ER and NR. For all datasets, the apparent lifetime increases with energy, though the increase in the NR datasets is mild and could at least partly be explained by the $\mathcal{O}(10\%)$ ER contamination described in section 5.4.2. The values found converge to (22 ± 1) ns for low energy nuclear recoils at both field strengths, in agreement with (zero-field) measurements of low-energy nuclear recoils reported by Akimov et al., but disfavoring values found by Ueshima and LUX. For ER, there is an increase of the effective triplet lifetime that becomes stronger for low field strengths and high energy due to the non-negligible recombination time. This is in disagreement with the analysis followed by LUX, where a recombination time <0.7 ns was assumed based on an extrapolation of an empirical formula to low energies [118]. In contrast, at a field of 0.5 kV/cm, we find an increase from roughly 22 to 25 ns, giving a recombination time of at least ~ 3 ns at 50 keV_{ee}. The high-energy limit of 25 ns seems to correspond reasonably well with the value of (27 ± 1) ns, which was found at a higher field strength (4 kV/cm) and a higher energy (from a ²⁰⁷Bi source, typically MeV electrons). For zero field, the apparent triplet lifetime component is significantly longer and seems to increase to beyond the observed maximum of (33 ± 1) ns. It is therefore unclear if the high-energy limit found here favor the values found in literature of (34 ± 2) ns [52] and (33 ± 1) ns [119] or ~ 45 ns [53]. We find good agreement with the values found by Akimov et al. [105] and XMASS [107], but significant disagreement with measurements from Ueshima [106].

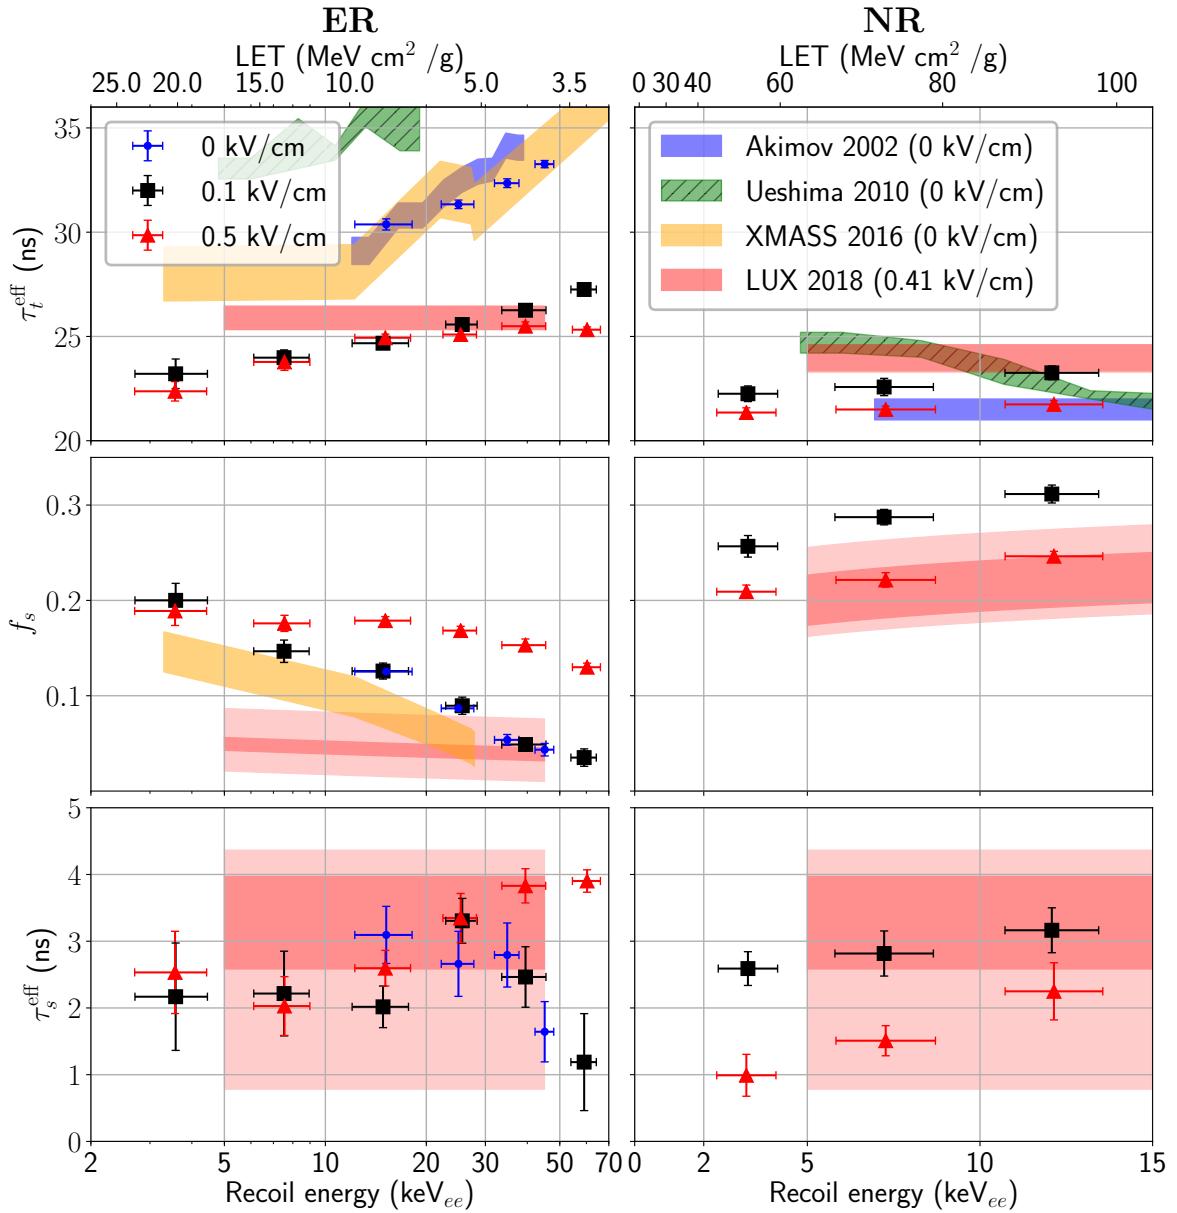


Figure 5.11: The best-fit parameters using a double-exponential model (eq. (5.5)) for electronic recoils (left row) and nuclear recoils (right row) as a function of recoil energy. The electronic linear energy transfer (LET) is indicated on the top scales. The black squares and the red triangles show the data at low field (0.1 kV/cm) and high field (0.5 kV/cm), the blue circles are for zero field data (ER only). The error bars indicate the statistical and systematic uncertainty, described in section 5.5.4, added in quadrature. The triplet lifetimes as measured by Akimov *et al.* [105], Ueshima [106] and the triplet lifetimes and singlet fractions measured by XMASS [107] are indicated by the blue, hatched green and orange bands, respectively. Note that these measurements were performed at zero field, and fit using a single exponential distribution for Akimov *et al.* and Ueshima. The LUX measurements (0.41 kV/cm) are indicated by the red bands, with the light shaded band indicating the systematic uncertainty due to the optical simulation [109].

The second row in figure 5.11 shows the singlet fraction as a function of energy. The ER and NR datasets are showing opposing trends; for higher energies, the singlet fraction rises for NR and decreases for ER events. This suggests that f_s is correlated with the LET (shown on the top scales), which decreases as a function of recoil energy for ER but increases for NR. This is in agreement with previous measurements that found a higher value of f_s for particle species that have a higher ionization density. For high fields, the influence is reduced, pointing to the recombination process as a physical origin of this difference. We find reasonable agreement with the zero-field measurements from XMASS. The LUX measurements showed compatibility with a linear or an exponential model. We find better agreement with the exponential trend and show these bands for comparison. While the NR data confirms LUX's result, the values of f_s found for ER show a strong disagreement; the values found by LUX of ~ 0.04 are much lower than the high field points around 0.16 for this energy range and electric field.

The bottom row in figure 5.11 shows the apparent singlet lifetime τ_s^{eff} . For this parameter, the fit uncertainties are much larger than for the first two parameters. This is partly due to a correlation between τ_s^{eff} and σ_{det} , which can be understood from the waveform model: both have the net effect of smearing out the strongly peaked part of the average waveform. In addition, the short timescales with respect to the digitizer time resolution makes the fit value depend strongly on only a few samples around the maximum, making the fit sensitive to mismodeling of the waveform shape. Given these caveats, we find decent agreement with the singlet lifetimes in the range of 2.2 ns to 4.3 ns found in literature [51, 109].

5.6.2 Recombination model

Rather than using the simple exponential model, we can use the full recombination model introduced in section 5.3.1. This is a seven-parameter model, given by eq. (5.4), so that an unconstrained fit would be challenging due to correlation between the parameters. However, since the scintillation model is split into a part that is due to direct excitation (dependent on τ_s , τ_t and f_s) and a part due to recombination luminescence (dependent on T_R , f_s^R and η), we can use the parameters of τ_s , τ_t and f_s that are found for the cases where the influence of recombination is negligible, such as for nuclear recoils at high field. We therefore fix the values of τ_s , τ_t and f_s to 1.5 ns, 21.5 ns and 0.23, respectively. The detector resolution σ_{det} is kept fixed at 1.5 ns. This leaves four parameters: the recombination time T_R , the electron escape probability η , the fraction of photons due to recombination f_R and the

fraction of recombination photons coming from the singlet state f_s^R . The resulting four-parameter fit was performed for the zero field ER data.

Figure 5.12 shows the resulting best-fit values. All the uncertainties are determined in the same way as described in section 5.6.1. The uncertainties of the recombination time T_R , the escape fraction η and the singlet recombination fraction f_s^R are substantial due to the observed correlation of these parameters. This makes it difficult to draw quantitative conclusions for these parameters. The recombination time T_R and the recombination singlet fraction f_s^R at high energies seems to be consistent with the values found by Kubota et al. of (15 ± 2) ns and 0.44 ± 0.11 [53], though it should be noted that this was fit for high-energy electrons and using a different value for η . Rather than the values $\eta < 0.05$ found for high-energy electronic recoils, we find that the escape probability is significantly higher at the energies studied here (>0.4 for $E < 40$ keV $_{ee}$).

The most notable result from the fits, is the fraction of scintillation light coming from recombination f_R . This value may be compared to the measurements of luminescence quenching with a sufficiently high electric field so that all electrons are pulled away from the interaction site, which indicate a value of around 0.7 as indicated by the colored bands in figure 5.12 [52–54]. At high energies, where these measurements were performed, the results obtained here show excellent agreement; the recombination fraction appears to asymptotically approach this value. It should be noted that the method described here is completely independent of the measurements that rely on extracting the electrons from the interaction site: the same result is found without applying any electric field.

In general, if the results of Kubota et al. are compared to the low-energy equivalent presented here, the conclusion can be drawn that the recombination efficiency decreases at low energies. This is evidenced by the decrease in f_R with respect to the constant value of 0.7 and the increase in η towards low energies. This conclusion is consistent with recent measurements that indicate a decreasing scintillation yield and an increasing charge yield for low-energy electronic recoils [126]. Furthermore, we note that the apparent decrease of f_s at high energy as observed in the double-exponential fit (section 5.6.1) is modeled without including superelastic collisions. Instead, this is accounted for in the following way. At high energy, f_R increases, so that a larger fraction of the luminescence comes from recombination. Since the recombination timescale is long with respect to the singlet lifetime, most of the recombination luminescence is fitted by the triplet component in the double-exponential model, reducing the value of f_s where recombination is dominant.

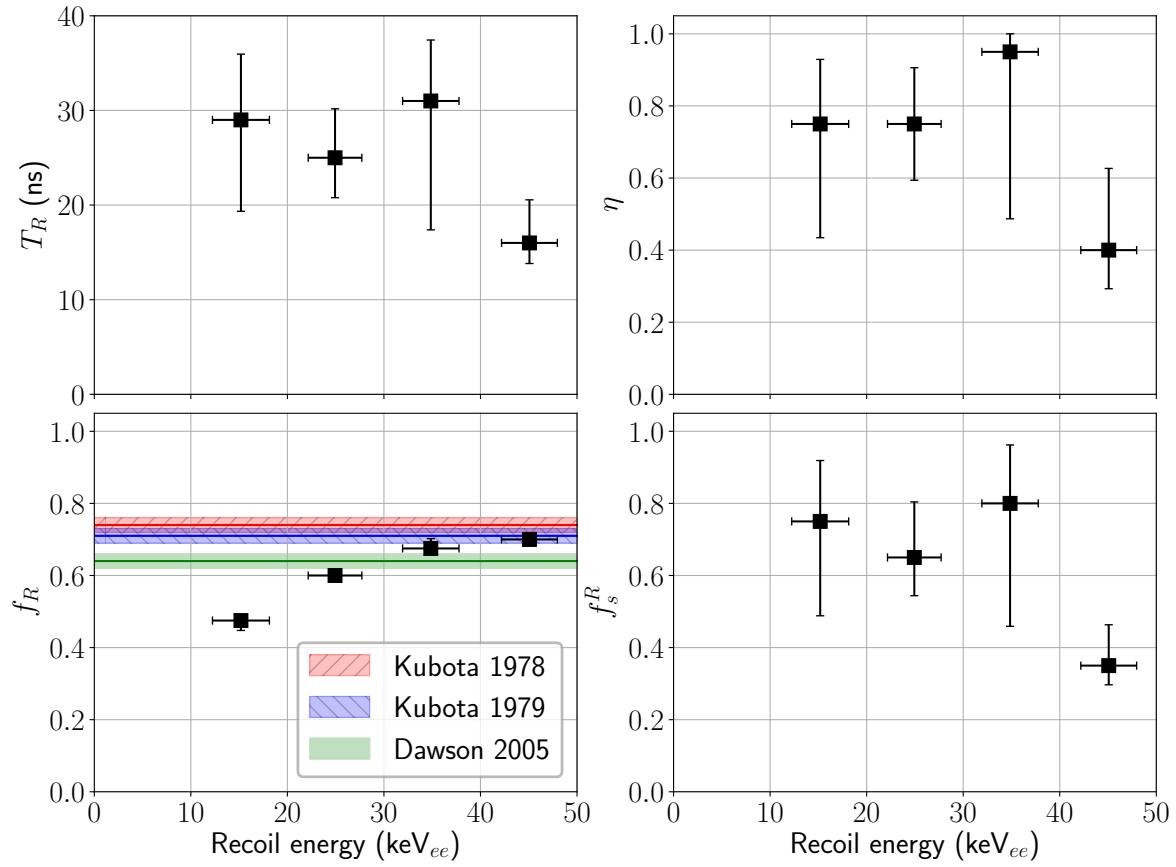


Figure 5.12: The best-fit parameters for the model including recombination introduced in section 5.3.1, described by eq. (5.4). The fit is applied to the ER data taken at zero field. The panels show the fit parameters for the recombination time T_R , the electron escape probability η , the recombination fraction f_R and the recombination singlet fraction f_s^R . The parameters τ_t , τ_s , f_s and the detector resolution σ_{det} are fixed at 21.5 ns, 1.5 ns, 0.23 and 1.5 ns, respectively. The recombination fraction f_R increases with energy up to values consistent with measurements performed at high energies, shown by the colored bands [52–54].

In this section, we have shown that a relatively simple recombination model can be used to explain the pulse shape difference observed at zero field. However, it should be noted that the true physical process can be more complex and may not be fully captured by the model. In particular, the ad-hoc parameter η could be replaced by a more detailed description of the dynamical behavior of electrons. It is furthermore possible that the quenching processes and superelastic collisions described in section 5.3 give a non-negligible influence on the pulse shape. Finally, the results here depend on assumed values of τ_s , τ_t and f_s , which were estimated from figure 5.11.

5.7 APPLICATION TO PULSE SHAPE DISCRIMINATION

As liquid xenon TPCs are currently leading the field of direct dark matter discrimination, there has been considerable interest in using pulse shape discrimination as a method of reducing the electronic recoil background in large-scale TPCs [79, 109, 110]. The theoretical performance of PSD in large-scale dual-phase detectors was previously studied by Kwong et al. [79], who found that using PSD was not performing well enough to be productive in large-scale TPCs looking for dark matter elastic recoils. This was concluded based on the comparison of the PSD performance with the regular S_2/S_1 discrimination. In the LUX detector, a combined discriminator decreased the ER contamination from $(0.4 \pm 0.1)\%$ to $(0.3 \pm 0.1)\%$ at a fixed NR acceptance of 50% [109]. In this section, we explore the theoretical performance of a combined discrimination in XENON1T. We use simulated data from the Laidbox code [127] combined with Blueice [128] to compute the likelihood function.

Currently, the method used by XENON1T to determine WIMP-nucleon interaction cross section limits uses the likelihood ratio computed in two dimensions, S_1 and S_2 [39]. This may be extended to a three-dimensional parameter space, using a PSD parameter as the third parameter. Given the description of the pulse shape parameters at low energies shown in section 5.6.1, we can calculate the performance of such a combined discrimination in an idealized experiment.

We produce Monte Carlo samples of a flat electronic recoil background component and that of a WIMP recoil. For each Monte Carlo generated event, we sample photon production times from the distribution described by eq. (5.5) with τ_s^{eff} , τ_t^{eff} and f_s taken from an interpolation of the best-fit values displayed in figure 5.11. As the drift field of XENON1T is close to 0.1 kV/cm, we take the values found for the low drift field. The times are then smeared with an adjustable detector resolution σ_{det} . The PSD parameter used is the logarithm of the likelihood ratio that is calculated from the photon probability density functions $f_{\text{ER}}(t)$ for electronic recoils and $f_{\text{NR}}(t)$ for nuclear recoils. The likelihood ratio \mathcal{L} is computed as,

$$\mathcal{L} = \frac{\prod_{t \in T} f_{\text{NR}}(t; E, \sigma)}{\prod_{t \in T} f_{\text{ER}}(t; E, \sigma)}, \quad (5.8)$$

where T is a set of photon detection times and where the energy E is the reconstructed ER equivalent energy (the dependence on E is through the dependence of the pulse shape parameters). The numerator and the denominator give the likelihood of the times in T to be drawn from the distributions for nuclear recoils

and electronic recoils, respectively, so that a high value is classified as an ‘NR-like’ event, while a low value is likely to be an ER event. We maximize both likelihoods with respect to the start time of the pulse. The discrimination efficiency depends on the scintillation and ionization gains g_1 and g_2 ; the values taken here are 0.144 and 11.5, respectively, as found for the first results of XENON1T. We simulated WIMP spectra for 20 GeV, 50 GeV and 500 GeV and a time resolution of 10, 5 and 2 ns, roughly representing the current time resolution, an improved time resolution and a significantly improved time resolution. It should be noted that the estimate of 10 ns is optimistic, given the typical PMT transit time spread of 9 ns [129] and the digitizer resolution of 10 ns [62]. For all the simulated datasets, over two million events were simulated.

We evaluate the ER and WIMP separation efficiency by constructing the receiver operating characteristic (ROC) curve, which gives the WIMP acceptance fraction as a function of ER acceptance. A three-dimensional histogram of the parameters S_1 , S_2 and \mathcal{L} is constructed for the WIMP signal and the ER background data. For each bin, the signal-to-background ratio is evaluated based on the number of WIMP and background events in the bin. Given the minimum signal-to-background ratio allowed, the ER and WIMP acceptance may be calculated as the fraction of all events in the bins with a higher signal-to-background ratio. The ROC curves for S_2/S_1 or PSD only are evaluated similarly, with two- and one-dimensional histograms. The allowed range for all events is 3 to 70 p.e. for S_1 and 50 to 8000 p.e. for S_2 , corresponding to the analysis range used for XENON1T’s first results. All acceptances are expressed as a fraction of events after passing the S_1 and S_2 selection criteria.

Figure 5.13 shows four ROC curves for a 500 GeV WIMP. The solid black line shows the discrimination achieved by using only the position of the events in (S_1, S_2) -space, neglecting all pulse time information. The dashed blue and dotted green lines show the improvement that can be achieved by including the pulse time information. As these lines barely improve over the baseline scenario, the improvement by including PSD is marginal. At an ER background rejection of 99.5% (indicated by the red dashed line), which is a typical value used for liquid-xenon based dark matter searches, including PSD increases the WIMP acceptance from 54.0% to 57.7% even for a time resolution of 2 ns, giving a relative exposure increase of 6.8%. At lower WIMP masses, this drops to even lower values (3.9% increase at 50 GeV, 0.6% at 20 GeV), as may be expected due to the decreased performance of PSD at the low energies that are more relevant for low WIMP masses. The thick blue line shows the ROC curve for PSD only. The shaded yellow portion of the graph is bounded by the line $y = x$, which would be the ROC curve from a

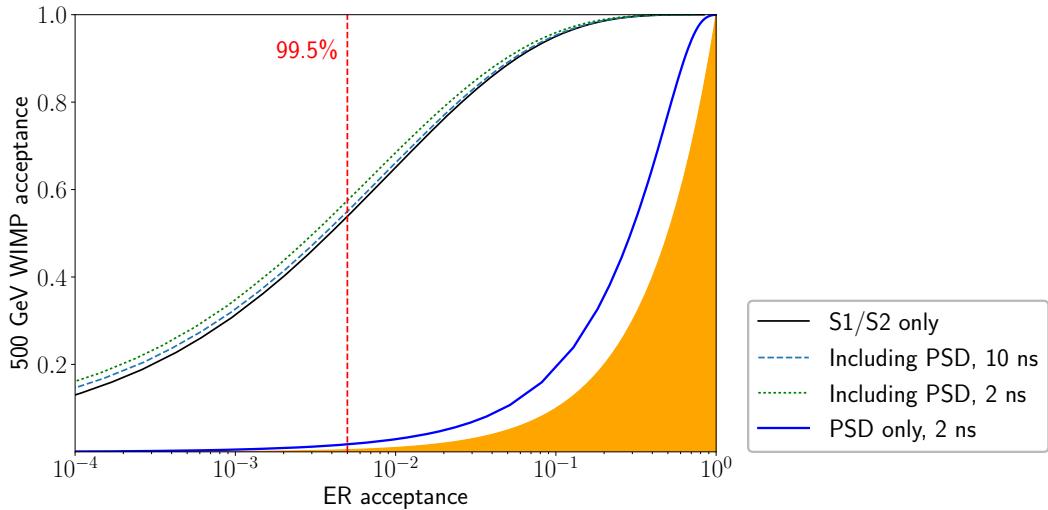


Figure 5.13: Simulated acceptances of events from a 500 GeV WIMP recoil as a function of the acceptance of a flat ER background for four different discrimination methods. The solid black line shows the results from using only the information of the S_1 and the S_2 size, while the dashed blue and dotted green lines are the performance when this method is combined with PSD, given a detector timing resolution of 10 or 2 ns. The thick solid blue line shows the performance using PSD alone, while the yellow shaded area is bounded by the line $y = x$ that gives the acceptance of a random cut. At an ER rejection of 99.5 %, the WIMP acceptance rises by 6.8 % if PSD is included, given a time resolution of 2 ns. For lower WIMP masses, the PSD performance decreases even more.

fully random cut on the data.

The LUX collaboration has used the pulse shape information to decrease the ER background leakage in the LUX detector from $(0.4 \pm 0.1) \%$ using only the S_2/S_1 ratio to $(0.3 \pm 0.1) \%$ using a combination of S_2/S_1 and PSD [109]. This number is the average leakage over the range of 10 to 50 p.e. at a fixed NR acceptance of $\sim 50 \%$. Although a full simulation of the LUX detector requires a detailed understanding of the LUX detector that goes beyond the scope of this work, we determine the background leakage measured in the same way in our simulation, setting the time resolution to LUX's value of 3.8 ns. This gives a background decrease from 0.47 % using S_2/S_1 only to 0.40 % including PSD, in reasonable agreement with the measured values.

It should be noted that the performance of PSD as given in figure 5.13 depends on several assumptions. Firstly, the scintillation yield was set to the XENON1T value. If the light yield increases, there are more photons available to increase the statistical power of the discrimination in eq. (5.8). This will also ameliorate the

S_2/S_1 discrimination. Secondly, the pulse shape parameters depend on the applied field. If the drift field is lower, the ER pulse shape broadens due to the longer effective triplet lifetime and lower singlet fraction (see figure 5.11), while the NR pulse shape appears to shorten (as seen by the increase in f_s for the lower field data), so that the pulse shapes become more distinct. Given the recent trend of large-scale detectors to operate at a lower drift field, this may become more relevant for next generation detectors. Finally, the simulation assumes an ER background in (S_1, S_2) -space, neglecting anomalous backgrounds formed by accidental coincidence of S_1 - and S_2 -signals, wall events that have reduced S_2 sizes, or unknown causes. These backgrounds may escape the S_2/S_1 discrimination, so that PSD could be used as an independent method to determine their origin or suppress them.

5.8 CONCLUSIONS

We have measured the scintillation pulse shape for electronic and nuclear recoils in liquid xenon down to low energies at different field strengths. By using a Monte Carlo based pulse shape description, we have probed the effective singlet and triplet lifetimes and the fraction of light coming from the singlet state. The high-energy limits of the results are compatible with previous measurements taken at higher energies. The measurements for electronic recoils are in tension with the measurements from LUX, as we found significantly higher values for the singlet fraction and a non-negligible recombination timescale for electronic recoils. A model of the time dependence for the recombination luminescence was taken from Kubota et al. [53] and applied to the electronic recoil data taken at zero field. Although correlation between the fit parameters makes it difficult to draw solid conclusions, we show that a good description of the pulse shape may be found by using a model that accounts for the recombination luminescence time dependence. Moreover, the fitted fraction of light from recombination shows excellent agreement with independent measurements. Given the measured pulse shapes, the theoretical WIMP/background discrimination performance of a large-scale dual-phase xenon detector was computed for a method combining the S_2/S_1 method and PSD in an optimal way. It is shown that even with a very fine time resolution and a relatively high-mass WIMP, the exposure increase is marginal, so that PSD will likely not be used as a main method for discriminating elastic WIMP recoil interactions.

5.9 ACKNOWLEDGMENTS

This work is part of the research program with project number FOM VP139, which is financed by the Netherlands Organisation for Scientific Research (NWO). We gratefully acknowledge the technical support from to the mechanical, electrical and computing departments at Nikhef. This work was carried out on the Dutch national e-infrastructure with the support of SURF Foundation.

6

FIELD DEPENDENCE OF 511 KEV GAMMA-INDUCED SIGNALS IN A DUAL-PHASE LIQUID XENON TPC

At the time that the technology of dual-phase liquid xenon TPCs was originally developed, it was not immediately obvious that it could be used for dark matter searches or that it would work at all. There are a number of boundary conditions that need to be fulfilled in order to build large, efficient detectors suitable for dark matter searches. For instance, the drift velocity and electron lifetime need to be sufficient so that the charge signal is measurable over long drift lengths. Already in the late 1960s, it was found that the electron mobility (which is related to the drift velocity) is high in liquid xenon, and later measurements indicated that good electron lifetimes could be achieved, so that “liquid xenon will be useful as a detector medium in the so-called time projection chamber” [130]. In a more systematic study, it was shown that sufficient lifetimes could be achieved with a modest purification setup [131]. Measurements of single-phase [132] and dual-phase [77] TPCs demonstrated the required performance. Finally, the discrimination of nuclear and electronic recoils and low energy threshold were demonstrated, both of which are of paramount importance for dark matter detection [78]. Moreover, it was concluded that “only modest fields are needed for the background-discriminating measurement of nuclear recoils”.

Today, it appears we may similarly need to re-evaluate the basic properties that are at the foundation of the principle of TPC operation, as maintaining the field has proven to be a difficult task as the TPC sizes grew. While the design drift field has been constant at approximately 1 kV/cm (for XENON10 [133], XENON100 [80], XENON1T [134], and 2 kV/cm for LUX [81]), the *voltage* applied to the cathode seems to be roughly constant for all TPCs, broadly hovering around 15 kV, as can be seen in figure 6.1. At higher voltages, spurious electron emission, photon emission or (micro)discharges occur, so that the TPC cannot be used for dark matter detection [135, 136]. As a consequence of the longer drift lengths for larger TPCs, the field has thus become lower. For future detectors, indicated by the dashed lines

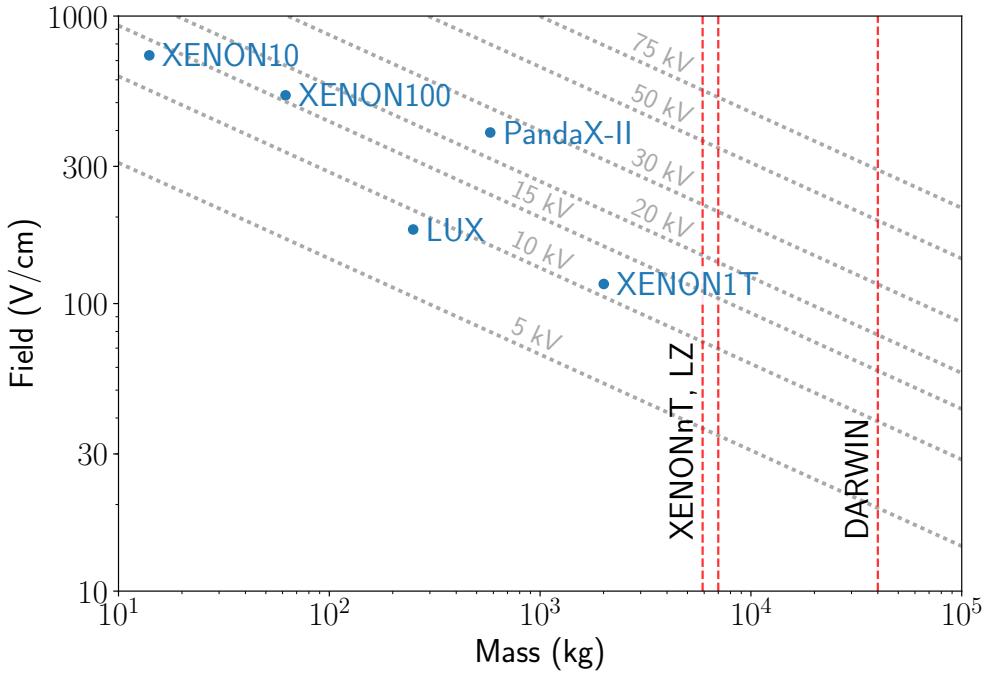


Figure 6.1: The drift field as a function of the instrumented xenon mass for several liquid xenon TPCs. The red dashed lines show the mass of planned future TPCs. The gray dashed lines show the voltage required to reach the indicated field, assuming a cylindrical TPC geometry with $L = 2R$ and $\rho = 3.0 \text{ g/cm}^3$. While the design field of most indicated TPCs was 1 kV/cm or above, the field has been decreasing with mass with cathode voltages around 15 kV for all TPCs. If this trend continues for future TPCs, fields below 100 V/cm will be used. Source code available at [10].

in figure 6.1, the longer drift length means that either improved high voltage engineering is required, or that the drift field will drop to values below 100 V/cm .

The paper presented in this chapter provides a systematic study of some of the important parameters for liquid xenon TPCs as a function of field, using 511 keV gamma-ray recoils from a ^{22}Na source. Interestingly, it seems that many of the variables of interest, such as the drift velocity, electron lifetime and charge-to-light ratio start to show a large rate of change at fields below roughly 100 V/cm : exactly the — relatively uncharted — domain that future TPCs may be venturing into. In particular, the drift velocity decreases and the diffusion constant rises at low fields, so that S_2 signals broaden due to the combined effect of more diffusion during a longer drift time. This may cause difficulties in multiple scatter event rejection, as wide S_2 s are more likely to overlap. The effect of the changing charge-to-light ratio could deteriorate NR/ER discrimination, which is currently an active field

of study [137, 138]. Finally, low fields significantly broaden the scintillation pulse shape for 511 keV ERs. Although this is in principle advantageous for pulse shape discrimination purposes, it should be noted that the pulse shape also depends on energy (see chapter 5), so that these results may not quantitatively carry over to low energy ERs.

The large rate of change of the parameters at low fields means that inhomogeneities in the field will cause relatively large changes in the parameters. This can be countered with several methods: (i) homogeneous drift field design, optimized for realistic cathode voltages; (ii) good modeling of drift field and (iii) detailed measurements of properties as a function of drift field. The paper presented here aims to improve these measurements, aiming for stable future TPC operation even if the drift fields reach low values.

FIELD DEPENDENCE OF 511 KEV GAMMA-INDUCED SIGNALS IN A DUAL-PHASE LIQUID XENON TPC

E. Hogenbirk, M. P. Decowski, K. McEwan, A. P. Colijn

Journal of Instrumentation **13**, P10031

Accepted 28 September 2018

Available online at

<http://iopscience.iop.org/article/10.1088/1748-0221/13/10/P10031>

<https://arxiv.org/abs/1807.07121>

Reference [3]

6

6.1 ABSTRACT

We present measurements of light and charge signals in a dual-phase time projection chamber at electric fields varying from 10 V/cm up to 500 V/cm and at zero field using 511 keV gamma rays from a ^{22}Na source. We determine the drift velocity, electron lifetime, diffusion constant, and light and charge yields at 511 keV as a function of the electric field. In addition, we fit the scintillation pulse shape to an effective exponential model, showing a decay time of 43.5 ns at low field that decreases to 25 ns at high fields.

6.2 INTRODUCTION

The search for weakly interacting massive particles by direct detection has seen sensitivity improvements of orders of magnitude in recent years, chiefly due to the employment of dual-phase liquid xenon time projection chambers (TPCs). In this type of detector, two signals following an energy deposition in the liquid xenon are registered, the first due to the scintillation light, called S_1 , and the second due to the ionized electrons, called S_2 . For the S_2 , the electrons liberated in the liquid xenon drift up towards a gas layer under the influence of an electric field, where a stronger electric field extracts them from the liquid and produces proportional scintillation. The light from both signals is detected using photomultiplier tubes (PMTs).

The magnitude of the applied drift field influences the performance of TPCs, as the dynamical behavior of the free electrons is changed. Firstly, for lower fields, the electrons are more likely to recombine with xenon ions, forming excitons that decay and contribute to the scintillation signal instead of the ionization signal. In effect, this changes the ratio of S_2/S_1 as well as the scintillation pulse shape due to the time delay in exciton formation. Secondly, the drift field determines the electron drift velocity, which rises steeply at fields up to ~ 100 V/cm and saturates for higher fields. Thirdly, longitudinal electron diffusion depends on the applied field, so that S_2 signals become wider at lower fields for the combination of two reasons: a lower drift velocity (and thus more time for electron diffusion for a given drift length) and a higher diffusion constant. Finally, the electron lifetime, the average time before an electron is absorbed in the liquid xenon, is assumed to depend on the electric field [47, 90], although direct measurements of this are scarce [131, 139, 140].

As the size of liquid xenon TPCs has increased, it has become apparent that the applied drift fields have become progressively lower. Longer drift lengths require higher cathode voltages, providing an engineering challenge. For example, the XENON10, XENON100, LUX and XENON1T TPCs have operated at maximum average fields of 730 V/cm [141, 142], 530 V/cm [80], 181 V/cm [88] and 117 V/cm [39], respectively. Planned future TPCs, such as XENONnT [68], LZ [143] and DARWIN [43] will feature even longer drift lengths, requiring vastly improved high voltage engineering or operating at fields lower than 100 V/cm. Some of the aforementioned low drift field effects, such as the light and charge yields, are well-measured and described by the simulation toolkit NEST [87, 144], while other effects have not been systematically measured yet.

This work describes measurements with XAMS [1], a dual-phase liquid xenon

TPC, operating at fields between approximately 10 V/cm and up to 500 V/cm, as well as measurements at zero field. We use data from 511 keV gamma-ray recoils from a ^{22}Na positron annihilation source. The drift velocity, electron lifetime, diffusion constant, and light and charge yields are determined for all measured fields and are compared to NEST where possible. In addition, we fit the scintillation pulse shape using a model with two exponential decay components [2], where we allow the effective triplet lifetime component to vary.

6.3 MEASUREMENTS

6.3.1 *Data acquisition and processing*

The active volume of the XAMS TPC has a cylindrical geometry, with the distance of the gate to cathode mesh of (100 ± 2) mm and a diameter of 44 mm (at -90°C). S_1 and S_2 photons are detected with two 2-inch R6041-406 Hamamatsu PMTs [98] that view the active volume from above and below. The high voltage required for the PMT bias, anode and cathode are provided by a CAEN DT1471ET power supply. During all measurements, the PMT voltages were 750 V and 700 V for the bottom and top PMT, respectively, resulting in a gain of 2.69×10^5 and 2.90×10^5 for the bottom and top PMT, respectively. This relatively low value was set to avoid signal saturation. The extraction field was supplied by biasing the anode with 3.5 kV for all measurements, except for the zero field measurement.

The gamma-ray source is a (172 ± 5) kBq ^{22}Na source, that decays by positron emission and subsequent 1274 keV gamma ray with a branching ratio of 90.4 %. The positron rapidly annihilates and produces two back-to-back 511 keV gamma rays. We use an external NaI(Tl) detector to trigger on one of the three gamma rays, with one of the other gamma rays interacting in the TPC. If a three-fold coincidence of both PMTs in the TPC and the NaI(Tl)-detector is found within 120 ns, full waveform signals are recorded for all channels. We changed the waveform duration from 164 μs (at high fields) up to 328 μs to account for the reduced drift velocity at low fields. For all except the lowest field measured, the event window was long enough to capture the entire drift length, with the S_1 in the middle of the waveform. The source was mounted in a lead collimator setup (similar as described in [1]) aimed at the bottom part of the TPC. The z -position distribution of the interactions follows a Gaussian distribution at (-58.6 ± 16.6) mm ($\mu \pm \sigma$) (see section 6.4.1). For the measurements described here, we varied the electric field strength by applying different voltages on the cathode. A total of 10 voltage settings were used with voltages rang-

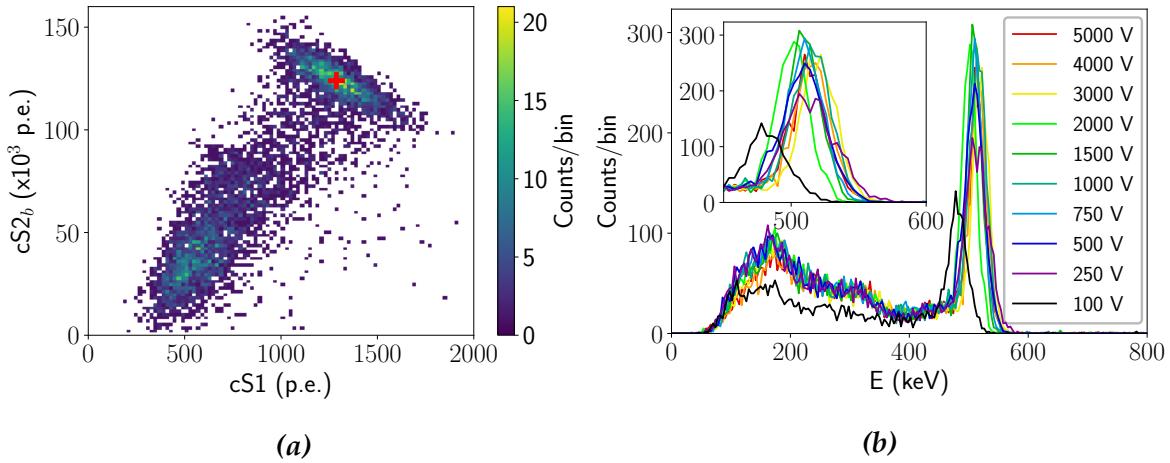


Figure 6.2: (a): The distribution of the S_1 and S_2 signals after all cuts and corrections for light detection efficiency (for S_1) and electron lifetime (for S_2) for the described ^{22}Na source. The red cross shows the photopeak position as determined from Gaussian fits (see section 6.4.4). The population at lower values of cS_1 and cS_{2b} is due to Compton scatter events, which give a smaller energy deposition. No events are observed at very low values of S_1 due to the trigger threshold. (b): The energy spectrum, reconstructed using a linear combination of S_1 and S_2 , for all data. The inset shows a detailed view of the photopeak at 511 keV.

ing from 100 V to 5000 V, giving fields up to 500 V/cm. In addition, we measured at zero applied voltage, where only the scintillation signal is produced.

6

6.3.2 Event selection

Valid events are selected using three cuts. Only events with precisely one S_1 and S_2 are selected, removing a large fraction of events containing pileup signals and multiple scatter events. Events with high S_2 widths are cut, as they are indicative of merged S_2 signals coming from unresolved multiple scatters. All events are further required to contain an S_1 -signal coincident with a signal in the NaI(Tl) detector. The final fraction of events passing all cuts is $(57 \pm 2)\%$ for all measurements except for the lowest field, where the event window was not long enough to capture events from the bottom of the TPC resulting in only 40 % of the events passing cuts. For some of the analyses described below, only the events with a z -coordinate within the collimated beam ((-58.6 ± 16.6) mm) are selected. This cut accepts $(64 \pm 1)\%$ of all events passing previous cuts. For the lowest field measurement, the selection is limited to $z \geq -60$ mm, giving an acceptance of 49 %.

An example of the energy distribution in corrected S_1 (cS_1) and the corrected S_2 using only the bottom PMT (cS_{2b}) is shown in figure 6.2a for a cathode voltage of 5000 V. This distribution is plotted after applying all the aforementioned cuts and after correcting for the z -dependent light detection efficiency for S_1 and for the electron lifetime in S_2 (see section 6.4.2). The 511 keV photopeak is found at the position of the red cross, showing a downwards sloping ellipse due to the anti-correlation of S_1 and S_2 . An additional distribution due to Compton scatter events extends down to lower energies. At low S_1 values, no more events are found due to the trigger cutoff for low amplitude S_1 s. For lower applied cathode voltages, the S_1 increases in favor of the S_2 signal, yielding a similar plot as figure 6.2a but shifted down in S_2 and up in S_1 .

Since the S_1 and S_2 signals are anti-correlated, a superior energy resolution can be achieved with a linear combination of the two signals. This is called the combined energy scale (see section 6.4.4). In figure 6.2b, the combined energy spectrum is shown for all data. All measurements except for the one taken at 100 V show a very similar energy spectrum (the discrepancy of the 100 V datapoint is discussed in section 6.4.4) with a clear photopeak at 511 keV. The energy resolution achieved at this energy is $(2.8 \pm 0.5) \%$, where the uncertainty indicates the standard deviation across different measurements.

6.3.3 Electric field simulation

The drift field in the TPC is calculated using the Comsol Multiphysics package [145]. The model, described in detail in [82], uses the cylindrical symmetry of the TPC and includes the geometry and electrical properties of the Teflon structure, the meshes and the vessel holding the TPC. Figure 6.3 shows the resulting electric field as a function of the z -coordinate for all used voltages. This gives the field averaged over r^2 (volume average), as the dependence on r is relatively minor. There are two regions of high field distortion: at the top, caused by the high anode voltage ‘leaking’ through the grounded gate mesh at the top of the TPC ($z = 0$), and at the bottom, caused by the high cathode voltage. The region where the collimated beam is located, indicated by the dashed lines, is a region of relatively low field distortion. We calculate the average field and its uncertainty by taking the volume-averaged field and its standard deviation in the relevant regions. Very close to the meshes, the field model breaks down due to the assumed cylindrical symmetry which is incorrect for the square mesh structure. We therefore restrict the average field computation to between -5 mm and -95 mm.

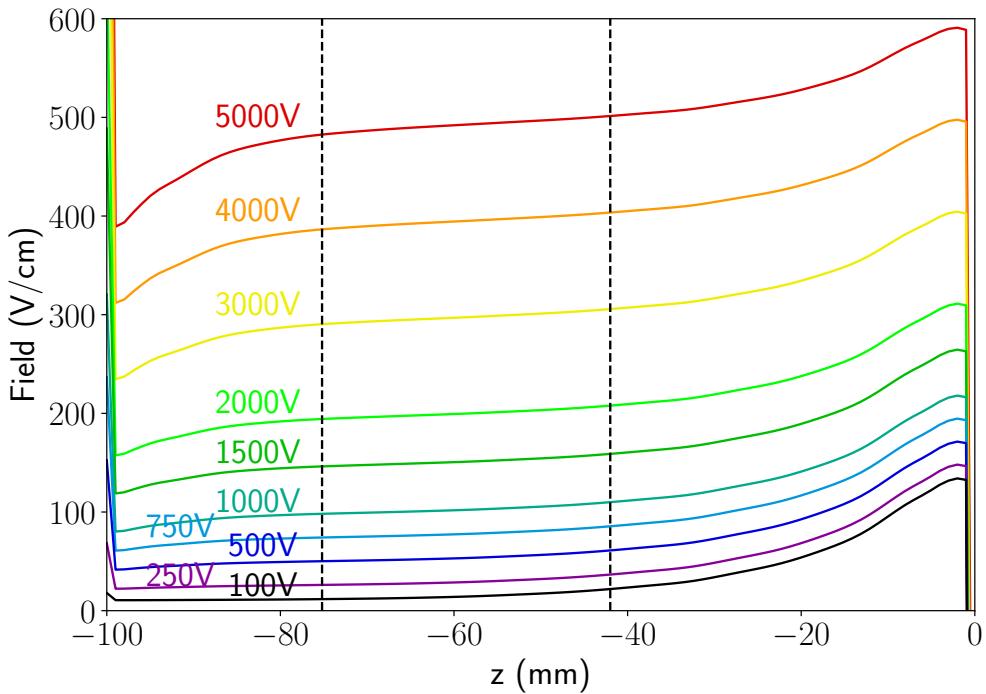


Figure 6.3: Electric field as a function of the z -coordinate for various applied cathode voltages, calculated using Comsol Multiphysics [145]. The gate mesh is located at $z = 0$, and the cathode is at -100 mm. The dotted lines indicate the region containing the collimated beam ($\mu \pm 1\sigma$).

6.4 RESULTS

6.4.1 Drift velocity and z -coordinate

The drift velocity as a function of field is determined from the maximum observed drift time, which corresponds to the position of the cathode. We subtract (1.0 ± 0.5) μ s from this time to account for the finite drift time between the gate mesh and the liquid level, which we observe in the data by the changing S_2/S_1 ratio as a result of the high field between the gate mesh and anode mesh. The uncertainty of 0.5 μ s on the drift time and the estimated uncertainty of 2 mm of the distance from the cathode to gate are propagated to the drift velocity uncertainty. We reconstruct the z -coordinate by assuming a constant drift velocity over the full drift region. Although the calculated electric fields show non-uniformity in the drift region and therefore a non-uniform drift velocity, for high enough fields the drift velocity saturates so that the drift velocity is relatively constant over the full range. The relative inhomogeneity in drift velocity in the range from -5 mm to -95 mm

is 1.5 % for an applied voltage of 5000 V and increases to 7 % for 750 V, calculated assuming drift velocities from NEST. Correcting for this effect would require a parameterization of $v_d(E)$ and an iterative approach, and was deemed unnecessary for this analysis.

For voltages of 500 V and below, we observe a significant change in the z -positions reconstructed using the drift time with respect to the higher field measurements. We attribute this to field inhomogeneity and the large gradient of the drift velocity as function of field below 100 V/cm. For these measurements, we reconstruct the z -coordinate by using the fraction of S1 light observed in the top PMT f_t . We calibrate f_t as a function of z using the highest field measurement, where field distortion is minimal. For the three lowest field measurements, we determine the drift velocity by taking the average drift times at two points with f_t -derived z -coordinates. We estimate that using these coordinates introduces an uncertainty of 2 mm in the drift distance and take this as the uncertainty on the drift velocity. The resulting drift velocities are indicated in figure 6.6(a) and show good agreement with NEST.

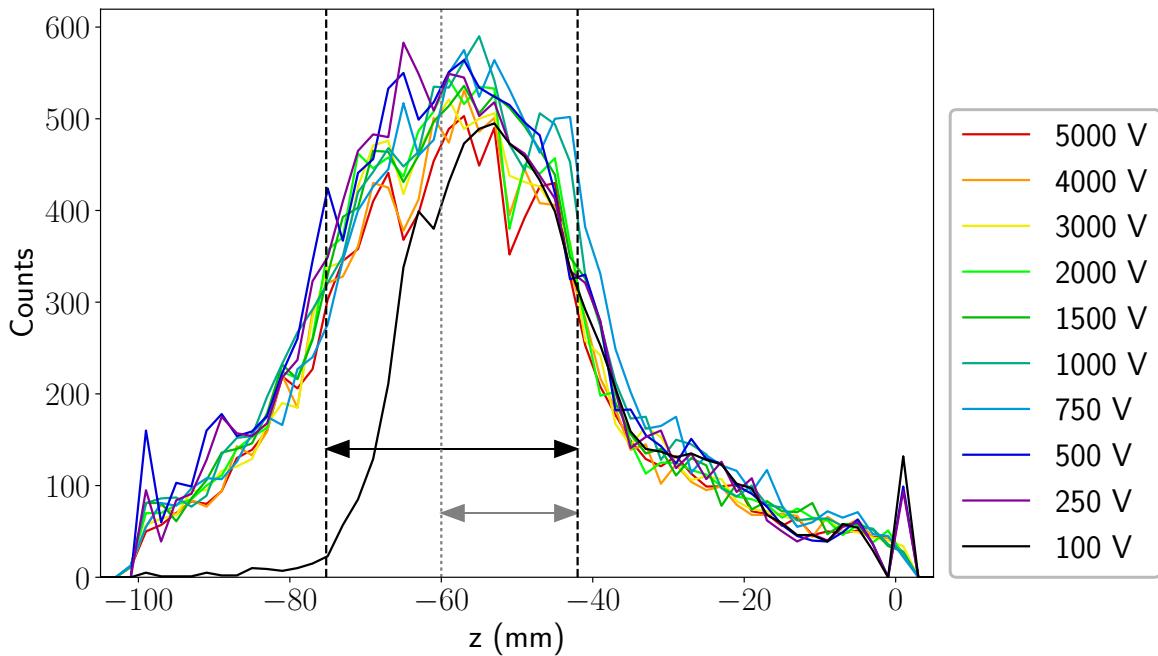


Figure 6.4: The z -distribution for all measurements, reconstructed using the drift time for measurements above 500 V and using the S1 light distribution for 500 V and below. The distributions show the same trend around the position of the collimated beam, except for the measurement at the lowest field due to an insufficiently long event window setting. The black arrow indicates the selection range used in most of the analyses. For the lowest field, the region indicated by the gray arrow is used instead.

The final position distribution for all measurements is shown in figure 6.4. The distributions show the same beam profile due to the collimator setup, except for the measurement at 100 V that lacks events for low values of z since the event window was too short to capture these events. The dashed black lines indicate the boundaries of the position cut used in several of the following analyses. For the 100 V measurement, the dotted gray line is used as a lower bound. The spiked feature at 0 and -100 mm for measurements at 500 V and below are due to interpolation artifacts in the f_t -derived z -coordinate.

The S_1 -signal is corrected for z -dependent light detection efficiency (LDE). We use a previously determined LDE map (as described in [2]) for this correction, where the corrected S_1 is normalized to the volume-averaged S_1 . We confirm that this LDE map is consistent with current conditions using a measurement where the ^{22}Na source was uncollimated, so that the events are spread out over the full detector volume.

6.4.2 Electron lifetime

The attachment of free electrons to impurities in the liquid xenon causes a decrease of electrons with increasing drift time, which follows an exponential distribution characterized by the electron lifetime τ according to

$$n_e(t_d) = n_{e,0} \exp\left(-\frac{t_d}{\tau}\right). \quad (6.1)$$

We fit this function for events with a z -coordinate within (58.6 ± 16.6) mm of the collimated beam position. To determine the electron lifetime, we select events within the full absorption peak, corresponding to a 511 keV energy deposition, so that the mean initial number of electrons $\langle n_{e,0} \rangle$ is constant. This is done iteratively. Full absorption events are first selected using a Gaussian fit to the corrected S_1 spectrum. After S_2 correction, the energy can be calculated using a combination of the S_1 and S_2 signals (see section 6.4.4). The reconstructed energy is then used to refine the photopeak selection, after which the S_2 correction is recomputed. This process is repeated until the correction on the electron lifetime is small with respect to the fit uncertainty.

The electron lifetime as determined from a direct exponential fit is biased due to field inhomogeneity within the fit range. Figure 6.3 shows that high z -coordinates correspond to a relatively higher field than events occurring deeper down in the TPC. This implies that for high z -coordinates, more electrons are extracted from

the interaction site, leading to a steeper observed fall of $n_e(t_d)$. We calculate the magnitude of this effect by simulating the observed effective electron lifetime as a function of the true electron lifetime, assuming the field dependence of the S2 size as described by NEST. We thus obtain a (higher) corrected electron lifetime and use this lifetime for the S2 correction. In figure 6.6(b), the directly fitted and corrected electron lifetimes are shown with blue circles and green squares, respectively. The indicated uncertainties correspond to the statistical uncertainty from the exponential fit.

The electron lifetime depends on the attachment cross-section of impurities in the liquid xenon, which depends on the electric field [47, 146]. This cross-section may either increase or decrease with applied field, depending on the type of impurity. For example, the lifetime increases with field for O₂ and H₂O, but decreases for N₂O and CO₂. In this case, we find a maximum electron lifetime at fields between 100 and 200 V/cm, decreasing at lower and higher fields. However, it is unknown what impurity dominates absorption in our TPC, so that we are unable to extend this to a more general statement for other TPCs.

6.4.3 Diffusion

The width of an S2 increases for interactions occurring deeper in the TPC. The standard deviation of the S2 time profile σ_{S2} can be described by

$$\sigma_{S2} = \sqrt{\frac{2Dt_d(z)}{v_d^2} + \sigma_0^2}, \quad (6.2)$$

with D the diffusion constant, t_d the drift time and v_d the drift velocity. The zero-drift S2 standard deviation σ_0 comes from the S2 width due to the time electrons emit light within the gas gap. For the measurements with applied voltages above 500 V, we fit this function to the observed distribution of S2 width as a function of drift time. An example of this fit (for 1000 V cathode voltage) is shown in figure 6.5. The fit range, indicated by the black dotted lines, corresponds to the z -range shown in figure 6.4. We divide the range into 10 slices and compute the median width in each slice. This is indicated by the red points. The solid red curve shows a fit of equation 6.2 to the medians. The dotted orange curve traces the fit plus a constant offset of 150 ns; all events above this line are likely due to merged S2s from multiple scatter events and are cut in the analysis.

For the three lowest fields, the aforementioned drift field inhomogeneity causes

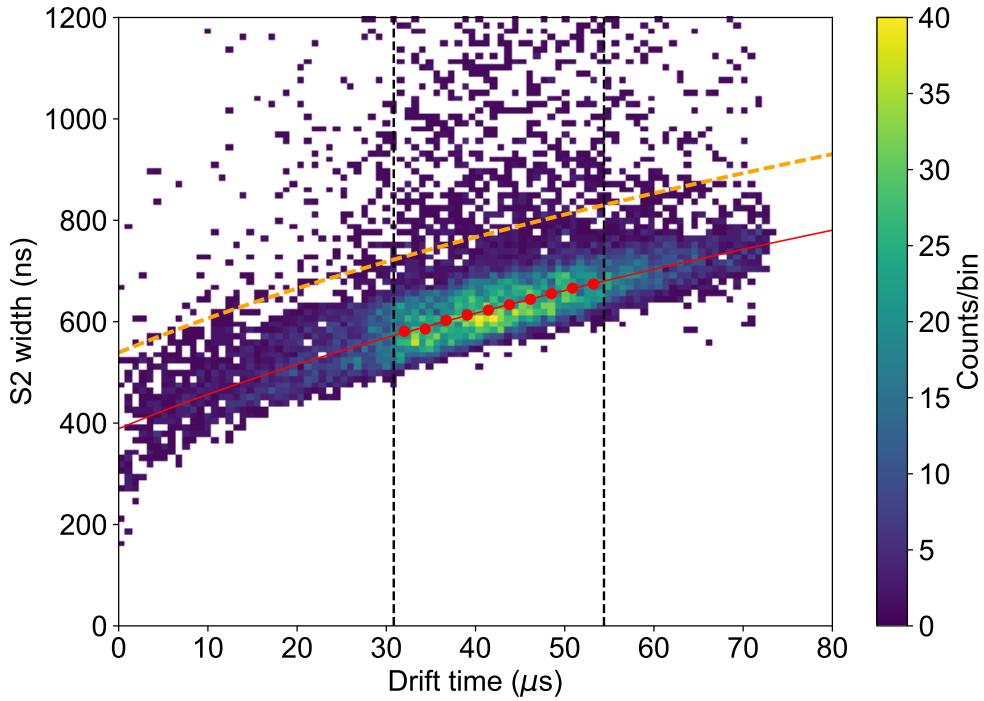


Figure 6.5: The S_2 width as a function of drift time for the measurement taken at a cathode voltage of 1000 V. The increase in width is caused by diffusion and traces the square root fit that is shown by the red solid line. This is fit to the median S_2 width in drift time slices, as indicated by the red points.

6

a significant deviation of this behavior due to the dependence of D on the field, and thus on z . To account for this, we instead use

$$\frac{d\sigma_{S_2}^2}{dt_d} = \frac{2D(z(E))}{v_d^2}. \quad (6.3)$$

The left hand side is calculated by taking the difference between $\sigma_{S_2}^2$ in drift time slices, divided by their separation in drift time. This has the advantage that the diffusion constant can be probed if it changes with z , however, it is highly sensitive to uncertainties in σ_{S_2} . This causes larger uncertainties of D for the lowest field values. We determine the uncertainty on D calculated in this way from the standard deviation of $\Delta\sigma_{S_2}^2$ for several neighboring drift time slices, while the uncertainty of D calculated with the direct fit of equation 6.2 is the statistical uncertainty on the fit parameter. The diffusion constant as a function of field is shown in figure 6.6(c). The drift velocity used to compute D is identical to the velocity shown in panel (a).

For the measurements where we use the direct fit (above 500 V), we estimate the impact of drift field inhomogeneity on the fitted value of D by a simulation of the S_2

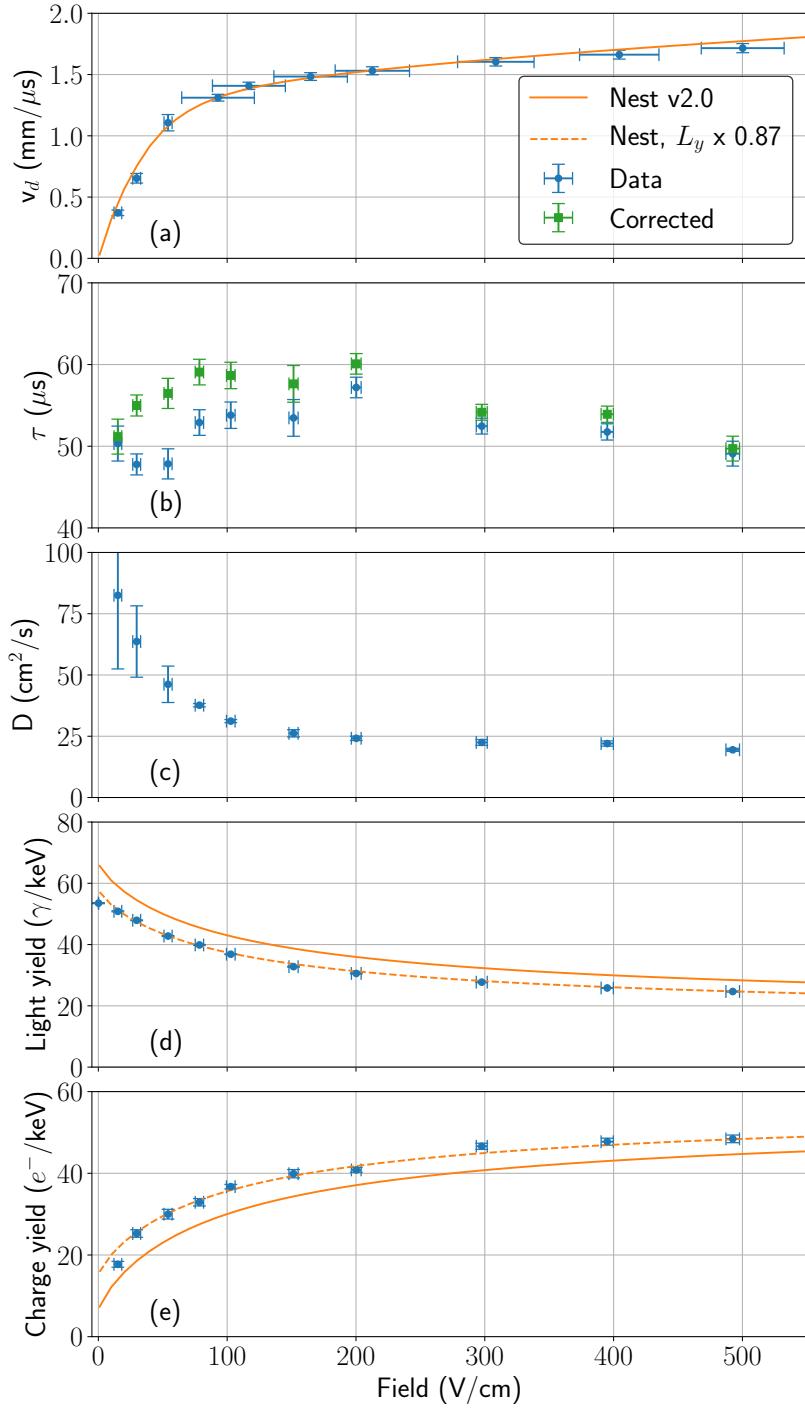


Figure 6.6: Various measured field-dependent properties compared to NEST predictions where available (orange lines) [87, 144]. The panels show (a) the electron drift velocity, (b) electron lifetime, (c) diffusion constant, (d) light yield and (e) charge yield. Blue points indicate the directly measured or fitted values, green points show values after corrections for field non-uniformity. The dashed lines correspond to the NEST light yield decreased by 13 % and give a better fit to the measured light and charge yield. The uncertainties in the field is evaluated as the standard deviation of the field within the z-selection. For the drift velocity (a), the full z-range is used, giving a relatively large uncertainty. Datapoints from this figure are available as ancillary files [147].

width as a function of drift time. This simulation uses the electric field as function of z (see figure 6.3), the field dependence of v_d (from NEST) and an interpolation of D as a function of field. We then fit the simulation with equation 6.2, neglecting the effects of field inhomogeneity, and compare the result with the simulated value of D . For the measurements above 500 V, the difference is less than 1 %; well below the uncertainty on D . We therefore conclude that the correction would be minor and that the field inhomogeneity can be neglected for these measurements.

6.4.4 Light and charge yields

After S_1 and S_2 corrections, we determine the corrected S_1 and S_2 (cS_1 and cS_{2b}), using only the bottom PMT for the S_2) corresponding to the full absorption peak at 511 keV. The cS_1 and cS_{2b} are anti-correlated, as a decrease in the extracted electrons contributes to the scintillation through the recombination process. This is described by

$$E = W \left(\frac{cS_1}{g_1} + \frac{cS_{2b}}{g_2} \right), \quad (6.4)$$

where $W = 13.7$ eV [48] and g_1 and g_2 are the photon and electron gain, respectively. The values of g_1 and g_2 are detector-dependent and can be determined from the linearity of cS_{2b}/E as a function of cS_1/E . This is shown in figure 6.7. The photopeak values of cS_1 and cS_{2b} (for example, as indicated by the red cross in figure 6.2a) are determined from individual Gaussian fits to cS_1 and cS_{2b} of events in the photopeak (selected with the same procedure as outlined in section 6.4.2). The cS_1 uncertainty is taken as the statistical uncertainty from the Gaussian fit, while the cS_{2b} uncertainty is dominated by the electron lifetime uncertainty. The measurement at 100 V shows a significant deviation from the linear behavior and is excluded from the fit. We attribute this to an imperfect integration of the full S_2 signal due to its extreme width and low area, resulting in a low amplitude waveform that is only partially integrated by the peakfinding algorithm. The obtained values from the fit are $g_1 = 0.102 \pm 0.003$ p.e./ γ and $g_2 = 5.01 \pm 0.13$ p.e./ e^- , in agreement with values found in [2].

Given the values of g_1 and g_2 , the S_1 and S_2 yields can be recalculated to absolute yields (in photons and electrons per keV, respectively), through

$$L_y = \frac{\langle cS_1 \rangle}{g_1 \cdot E}; \quad Q_y = \frac{\langle cS_{2b} \rangle}{g_2 \cdot E}. \quad (6.5)$$

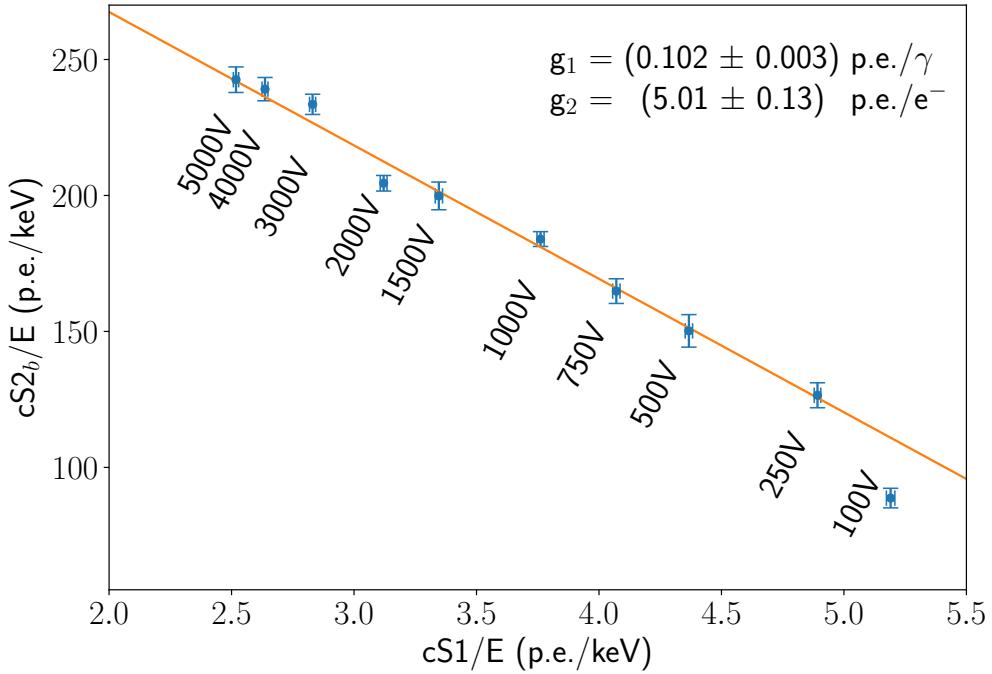


Figure 6.7: The S_2 and S_1 yield, changing from high field (top left) to low field (bottom right) measurements. The best-fit values for g_1 and g_2 give the curve indicated by the orange line. The measurement at the lowest field is excluded from the fit.

The yields depend on the incident particle type, energy and applied field. The yield for various fields for a gamma interaction at an energy of 511 keV is extracted from NEST and compared against our observations. This is shown in figure 6.6(d) and 6.6(e). While the field dependence of light and charge yields is well captured by the NEST description, our measurements favor a lower light yield and higher charge yield at 511 keV for all fields. Note that the light and charge yield are anti-correlated as the total number of quanta E/W is fixed (equation 6.4). A best-fit description is found for a 13 % decrease in light yield, as shown by the dashed curves in the figures. Changing the values of g_1 and g_2 similarly shifts the light and charge yields, however, since the uncertainty on these parameters is 3 %, this can only partly cause the observed discrepancy.

6.4.5 Scintillation pulse shape

In the absence of recombination luminescence, the scintillation pulse shape is described by a double exponential distribution due to the existence of two exciton states: the singlet state with lifetime τ_s and the triplet state with lifetime τ_t . Recom-

bination luminescence has the net effect of broadening the pulse due to a delay in the formation of exciton states, and possibly a preferred formation of the triplet state for the recombination process. This is most notable for low ionization density recoils and low electric fields [53]. Rather than constructing a detailed model of recombination, the scintillation pulse shape is usually described by an effective model, using a single exponential distribution [54, 105] or absorbing the delay due to recombination into an effective lifetime τ_t^{eff} [107, 109]. The normalized photon emission time distribution then becomes

$$I(t; \tau_s, \tau_t^{\text{eff}}, f_s) = f_s \left(\frac{1}{\tau_s} \exp \frac{-t}{\tau_s} \right) + (1 - f_s) \left(\frac{1}{\tau_t^{\text{eff}}} \exp \frac{-t}{\tau_t^{\text{eff}}} \right), \quad (6.6)$$

with f_s the fraction of light observed from the singlet state.

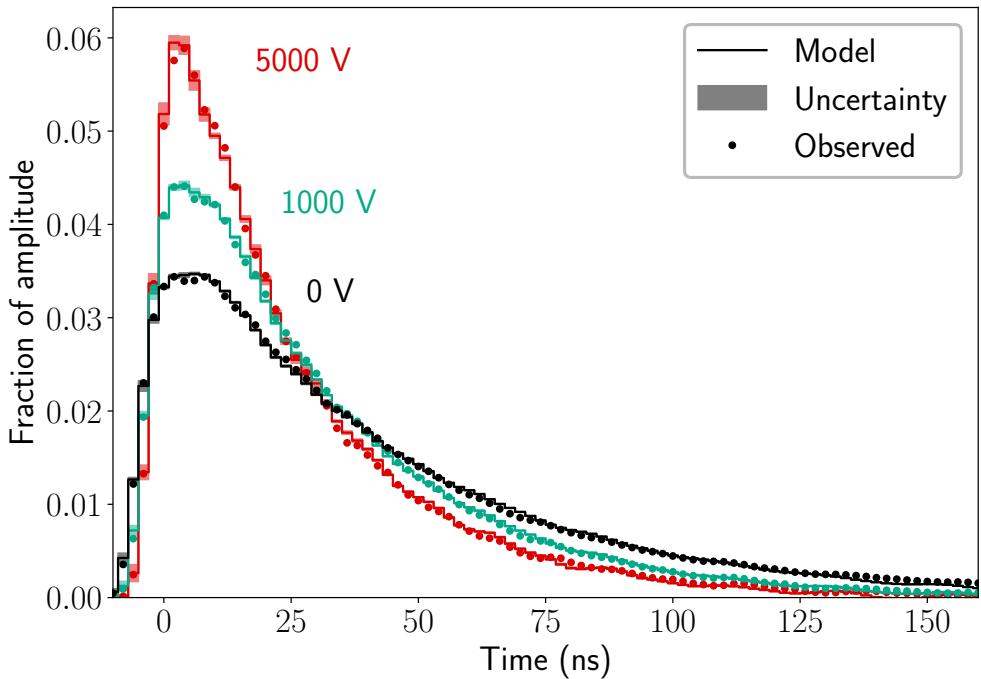


Figure 6.8: The average scintillation pulse shape for three of the cathode voltages. Measured values are indicated by the points, while the solid line and the shaded area show the model description and its uncertainty, respectively.

We use the model from equation 6.6 and fit this to the average pulse shape observed for all measured fields. Only S1s from events occurring within the full absorption peak at 511 keV and within the collimated beam position are used for this analysis. The fitting procedure is the same as outlined in [2] and uses simulated individual pulse shapes to account for the effect of pulse alignment on the average

pulse shape, as well as a high time resolution single photoelectron pulse shape model. In addition to the parameters in equation 6.6, an addition parameter σ_{det} comes from the detector time resolution, which is assumed to smear all photon detection times by a Gaussian distribution. This gives a total of four parameters: τ_s , σ_{det} , τ_t^{eff} and f_s . We fix the singlet lifetime and the detector time resolution to values found in [2] ($\tau_s = (2.0 \pm 1.0)$ ns and the detector time resolution $\sigma_{\text{det}} = (1.5 \pm 0.5)$ ns). Both are varied within their uncertainty and the effect on the best-fit values of τ_t^{eff} and f_s is taken as a systematic uncertainty on these values.

Figure 6.8 shows the average normalized waveform for three of the measured voltages. All the waveforms are aligned such that 10 % of the area is at $t < 0$. The points indicate the average of the measured S1 pulse shapes. The histograms show the best-fit model and the shaded regions around it indicate the uncertainty on the waveform that comes from the uncertainty on σ_{det} and τ_s . The curves clearly show the effect of the field on the scintillation pulse shape; much wider curves are found for low fields.

The best-fit values of τ_t^{eff} and f_s as a function of field are shown in figure 6.9. At low field, the pulse is characterized by a long effective triplet lifetime of approximately 43.5 ns and only a very small contribution of the singlet fraction (0.03 ± 0.01 below 100 V/cm), in agreement with the ~ 45 ns observed in [53] for $\mathcal{O}(\text{MeV})$ electron recoils at zero field. We note that the lifetime at zero field is slightly below the lifetime at the lowest nonzero field. This difference may be due to multiple scatter events that are not cut in the zero field data, as there are no S2s. In this case, the observed scintillation light comes from multiple energy depositions, so that effectively a lower energy than 511 keV is probed. Since the effective triplet lifetime decreases with decreasing energy for electronic recoils, this may cause an artificially reduced effective triplet lifetime. At high fields, the singlet fraction increases up to 0.13 ± 0.02 , while the lifetime drops to 25 ns. The data suggest that even at high field strengths, the effective triplet lifetime is still appreciably higher than the actual triplet lifetime of (22 ± 1) ns measured using high ionization density tracks [51]. This could imply that either the recombination process is still significant even at high field strengths, or that the low ionization density associated with 511 keV gamma recoils causes a delay in the direct process of excimer formation.

6.5 CONCLUSIONS

We have measured the dependence on electric field of various quantities of interest to dual-phase liquid xenon TPCs. For this, we used a setup with a collimated

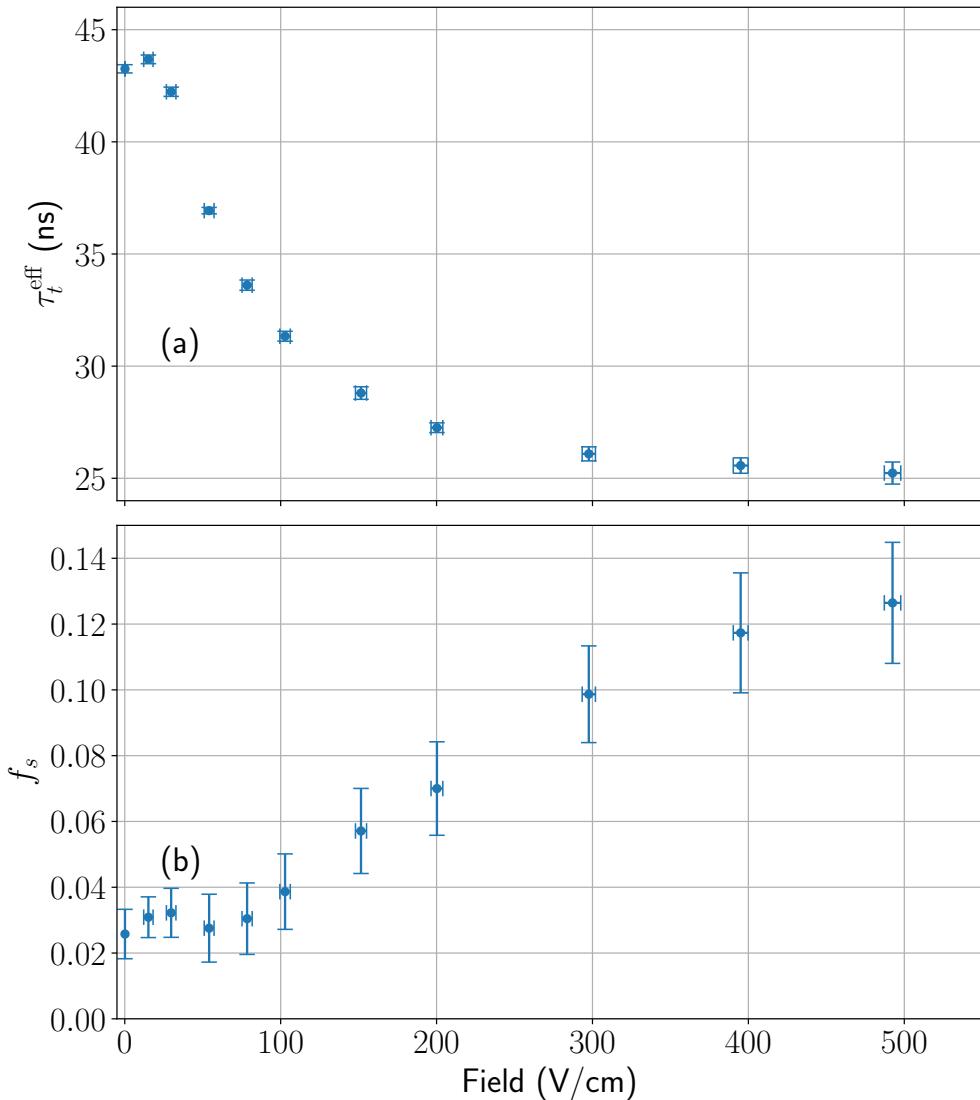


Figure 6.9: Best-fit values of the effective triplet lifetime τ_t^{eff} (a) and the singlet fraction f_s (b) as a function of field, fit to equation 6.6. The values of the singlet lifetime and the detector time resolution are varied within $\tau_s = (2.0 \pm 1.0) \text{ ns}$ and $\sigma_{\text{det}} = (1.5 \pm 0.5) \text{ ns}$, respectively. The pulse shapes at low field are characterized by a long decay time of up to 43.5 ns, and a very low singlet fraction consistent with 0.03 below 100 V/cm. At higher fields, the singlet fraction increases and the effective triplet lifetime drops to 25 ns. Datapoints from this figure are available as ancillary files [147].

^{22}Na source, triggering on 511 keV gamma recoils. We use an electric field simulation, and whenever necessary the field inhomogeneity is taken into account for the analysis. We measured the drift velocity, electron lifetime, diffusion constant and light and charge yields for fields ranging from 10 V/cm up to 500 V/cm. We find a

minor field dependence of electron lifetime, with a maximum value at 100 V/cm to 200 V/cm, although this may depend on the type of electron-attachment impurity. The diffusion constant is shown to remain relatively constant for fields higher than 100 V/cm, but rises steeply for lower fields. The light and charge yield dependence on field is well captured by NEST, although a systematic light yield decrease of 13 % at 511 keV is suggested by our data.

We fit the average scintillation pulse shape to a model containing two exponential decays, where the field-induced change of the recombination luminescence time dependence is fit by changing the effective triplet lifetime and singlet fraction. The effective triplet lifetime reaches values up to 43.5 ns for low fields and converges to 25 ns at high fields, while the singlet fraction increases from 0.03 ± 0.01 to 0.13 ± 0.02 .

We note that the strong field dependence of the diffusion constant and the drift velocity for low fields may provide a challenge for large dual-phase TPCs if a field below 100 V/cm is used. The combination of a high diffusion constant and low drift velocity makes the S₂s very wide, causing overlapping S₂ signals for multiple scatter events if the z-separation for the interaction positions is not sufficiently large. This could cause a significant increase in multiple scatter backgrounds into the single scatter signal sample, thus limiting the sensitivity to dark matter interactions.

6.6 ACKNOWLEDGMENTS

6

We gratefully acknowledge support through the FOM VP139 program, financed by the Netherlands Organisation for Scientific Research (NWO). We thank all members of the Nikhef Dark Matter group for shift duties during data-taking and for useful discussions. We are grateful for the technical support from the mechanical, electronics and computing departments at Nikhef. This work was carried out on the Dutch National e-Infrastructure with the support of the SURF Cooperative.

CHARACTERIZATION OF A DEUTERIUM-DEUTERIUM PLASMA FUSION NEUTRON GENERATOR

Perhaps the most important and challenging part of direct dark matter detection, or any rare event search, is in the reduction and rejection of backgrounds. For TPCs like XENON1T, exceptional care is given to material selection to prevent any radioactive decays occurring in materials close to the detector volume. In analysis, the background rejection strategy is largely based on two methods: fiducialization and discrimination of electronic recoil (ER) and nuclear recoil (NR) events (see chapter 2). Fiducialization means that events in the outermost region of the detector are rejected, as they are most likely to occur from radiation entering from the outside. The ER/NR discrimination further limits the accepted events to only those that show an experimental signature consistent with energy transfer to the xenon nucleus, rather than to an electron (which happens for gamma and beta radiation). This is done based on the position of nuclear recoil events in S_1 - S_2 space, which is usually called the *NR band*.

The response of liquid xenon TPCs to NRs is calibrated with neutron radiation, as these are neutral particles and thus interact with the xenon nucleus rather than the electrons. The energy transfer in a neutron-nucleus collision in xenon is very inefficient due to the large difference between the neutron and the xenon nucleus mass (131 amu), resulting in a maximum energy transfer of about 3 % of the kinetic energy of the neutron. This means that for the energy range of interest for WIMP-xenon collisions (up to tens of keV), neutrons with an energy of the MeV scale (*fast* neutrons) are required. Several sources of fast neutrons can be used. While previous experiments mostly relied on (α, n) sources, such as $^{241}\text{AmBe}$ [148, 149], recent TPCs have shifted to using deuterium-deuterium (D-D) fusion neutron generators for the NR calibration [39, 150]. These have several advantages over the (α, n) sources: (i) the neutron rate is adjustable and no radiation is produced when the device is off, (ii) no high-energy gamma radiation produced in the process that might introduce ER contamination into the ER band and (iii) the neutron energy spectrum

is approximately mono-energetic. The latter is a great advantage for the calibration of the energy dependence of the signal yields as a function of recoil energy, which is required for a signal model of WIMP-xenon interactions.

In D-D fusion, the outgoing neutron carries a kinetic energy of 2.45 MeV in the inertial reference frame of the fusion reaction. In the laboratory frame, this means that the neutron energy depends on the outgoing neutron angle with respect to the collision angle of the deuterium nuclei in the fusion process. In beam-on-target neutron generators, such as used by LUX, the energy is fixed by choosing a fixed angle relative to the device [151, 152]. The recent analysis by LUX also determined the recoil angle in xenon by using multiple interactions, so that the recoil energy can be computed and a direct calibration of the charge yield is possible [150]. For the D-D plasma fusion generator, the fusion process occurs within the plasma so that the interaction angle is randomly oriented [153]. This means that the spectrum is isotropic and contains a finite spread due to the kinematics. For any analysis using the initial neutron spectrum, this has to be taken into account.

The paper shown in this chapter presents the analysis that comprises a full calibration of the neutron generator used in XENON1T. The neutron energy spectrum is determined in section 7.5. In addition to the D-D neutrons, a non-negligible amount of high-energy neutrons from deuterium-tritium fusion were found. The angular distribution and the neutron flux as a function of applied voltage and current are determined in section 7.6. The combination of all analyses shown in this paper yields a full characterization of a D-D neutron generator that has never been performed before, and thus contributes to the understanding of plasma fusion neutron generators, and provides essential information required for a full understanding of the nuclear recoil response of XENON1T.

CHARACTERIZATION OF A DEUTERIUM-DEUTERIUM PLASMA FUSION NEUTRON GENERATOR

R.F. Lang, J. Pienaar, E. Hogenbirk, D. Masson, R. Nolte, A. Zimbal, S. Röttger,
M.L. Benabderrahmane, G. Bruno

Nuclear Instruments and Methods in Physics Research Section A, 879 (2018), pp. 31-38

Accepted 1 October 2017

Available online at

<https://doi.org/10.1016/j.nima.2017.10.001>

<https://arxiv.org/abs/1705.04741>

Reference [4]

7.1 ABSTRACT

We characterize the neutron output of a deuterium-deuterium plasma fusion neutron generator, model 35-DD-W-S, manufactured by NSD/Gradel-Fusion. The measured energy spectrum is found to be dominated by neutron peaks at 2.2 MeV and 2.7 MeV. A detailed GEANT4 simulation accurately reproduces the measured energy spectrum and confirms our understanding of the fusion process in this generator. Additionally, a contribution of 14.1 MeV neutrons from deuterium-tritium fusion is found at a level of 3.5 %, from tritium produced in previous deuterium-deuterium reactions. We have measured both the absolute neutron flux as well as its relative variation on the operational parameters of the generator. We find the

flux to be proportional to voltage $V^{3.32 \pm 0.14}$ and current $I^{0.97 \pm 0.01}$. Further, we have measured the angular dependence of the neutron emission with respect to the polar angle. We conclude that it is well described by isotropic production of neutrons within the cathode field cage.

7.2 INTRODUCTION

Neutron generators are a convenient, commercially available source of neutrons widely used in science and engineering. They can easily achieve a tuneable neutron flux of 10^6 n/s with some generators operating above the 10^{10} n/s range, they pose no or only minimal safety concerns when turned off, and they are available in a variety of configurations. The latest advances in the field of compact sealed-tube neutron generators toward the development of smaller, lighter and less expensive systems further extend their applicability.

Two main reactions are exploited in such generators: deuterium-tritium fusion yielding 14.1 MeV neutrons, and deuterium-deuterium fusion yielding 2.45 MeV neutrons in the center-of-mass frame. Two operating principles are commonly employed to induce fusion. One is to accelerate a beam of deuterium ions onto a solid state target which contains either deuterium or tritium. Another principle is the fusion of ions in a plasma in the presence of a high voltage potential. Indeed, there are detailed discussions of the characteristics of deuterium-tritium generators [154], deuterium-deuterium generators [153, 155, 156] as well as neutron generators in general [157, 158]. However, to our knowledge, the measurements reported here represent the first complete characterization of a deuterium-deuterium plasma fusion generator, including the determination of absolute neutron yield, neutron energy spectrum and emission anisotropy.

7

7.3 SETUP

7.3.1 Neutron Generator

The neutron generator characterized in this work is a model 35-DD-W-S deuterium-deuterium plasma fusion neutron generator manufactured by NSD/Gradel-Fusion. This generator produces 2.45 MeV neutrons based on the fusion of deuterium



The generator is capable of delivering neutron fluxes up to 10^7 n/s. Given our particular application of this generator in the field of direct dark matter detectors [68], it was modified to enable stable operation even at fluxes as low as 10 n/s. The neutron generator has a cylindrical shape with a length of 940 mm and a diameter of 138 mm. It has a standalone high voltage power supply module, a slow control program to monitor system parameters, and a water cooling loop.

The working principle of the neutron generator is based on inertial electrostatic confinement (IEC). The generator has a fusion chamber filled with deuterium gas. The deuterium pressure is reduced to a level that allows for plasma ignition by glow discharge (Paschen's law). The primary source of neutrons is considered to be the plasma in the volume surrounded by the cathode. Deuterium gas in the fusion chamber is ionized and the resulting ions are accelerated toward the cathode field cage. Once inside the field cage, the ionized gas is confined by applying a high voltage ranging between 10 – 100 kV. When the necessary conditions to overcome the Coulomb barrier are met, fusion occurs, emitting approximately mono-energetic 2.45 MeV neutrons.

7.3.2 Liquid Scintillators

We use two liquid scintillator detectors, a 3" \times 3" EJ301 cell and a 2" \times 2" NE213 cell, to measure the neutron energy spectrum and relative flux. The scintillator used in these detectors is identical, apart from the manufacturer [159]. Liquid scintillators are very popular for fast neutron detection as they can easily be shaped into the desired size and geometry of a given application. The process of elastic scattering by neutrons off the protons found in the hydrocarbon molecules produces prompt scintillation that offers excellent timing performance. Such nuclear recoils exhibit greater ionization density rates than electronic recoils that are induced by various backgrounds. Consequently, the ionization tracks of nuclear recoils produce higher yields of delayed fluorescence, resulting in scintillation pulses that decay more slowly than those of electronic recoils of comparable energy. The different pulse shapes that arise from electronic and nuclear recoils in liquid scintillators can thus be exploited using pulse shape discrimination methods [160–163]. Additionally, if the detector response to neutrons of specific energies is well known, these detectors can be used to reconstruct the energy spectrum of the incident neutron flux [164, 165].

The EJ301 detector cell was optically coupled to a ETEL 9821KB photomultiplier (PMT) which was operated at a voltage of 1700 V. The anode signal of the

PMT was acquired using a CAEN DT5751 digitizer, which samples at 1 GHz with 10 bit resolution and has a 1 V dynamic range. The NE213 detector cell was coupled to an XP2020 PMT via a short lucite light guide. The PMT was operated at a voltage of -1950 V. Standard nuclear electronics modules were used for analogue signal processing. A signal proportional to the total amount of scintillation light produced in the detector cell (pulse height) was derived from the ninth PMT dynode. A second signal related to the decay time of the light pulse (pulse shape) was derived from the PMT anode using the zero-crossing technique. The two signals were digitized using pulse-height sensitive ADCs and a PC based multi-parameter data acquisition system. Deuterium-deuterium IEC fusion devices are known to produce bremsstrahlung in addition to neutrons [166]. Therefore we operated the liquid scintillator detectors with lead shielding.

The gain of the PMT in the NE213 detector was constantly measured and adjusted using a feedback loop. This was constructed with an integrated LED, that gave constant light pulses at a rate of 65 Hz, and a voltage added to the bias voltage based on the PMT's response to the LED pulse. This setup allowed for a constant gain during the long measurements required for determining the energy spectrum (section 7.5).

7.3.3 *De Pangher Long Counter*

We use a De Pangher Long Counter to measure the absolute neutron flux. This detector is made of concentric layers of polyethylene and borated polyethylene. The borated polyethylene shields the detector from neutrons entering the detector through its side wall, whereas incident neutrons from the source, which enter through the planar surface, are thermalized by the polyethylene. At the center of the Long Counter is a tube filled with BF_3 that is enriched in ^{11}B . Thermalized neutrons that are scattered into this volume can be captured, producing an alpha through the $^{11}\text{B}(\text{n}, \alpha)^7\text{Li}$ reaction. The detection of the emitted alpha particle produces a unique signal in the detector that is used to tag neutrons.

The Long Counter has the ability to measure neutron fluxes with a response that is almost independent of neutron energy in the range from a few keV to almost 10 MeV. Their high sensitivity allows them to be used to measure low neutron fluxes. Additionally, they are directional, which is necessary to control the impact of in-scattered neutrons [167].

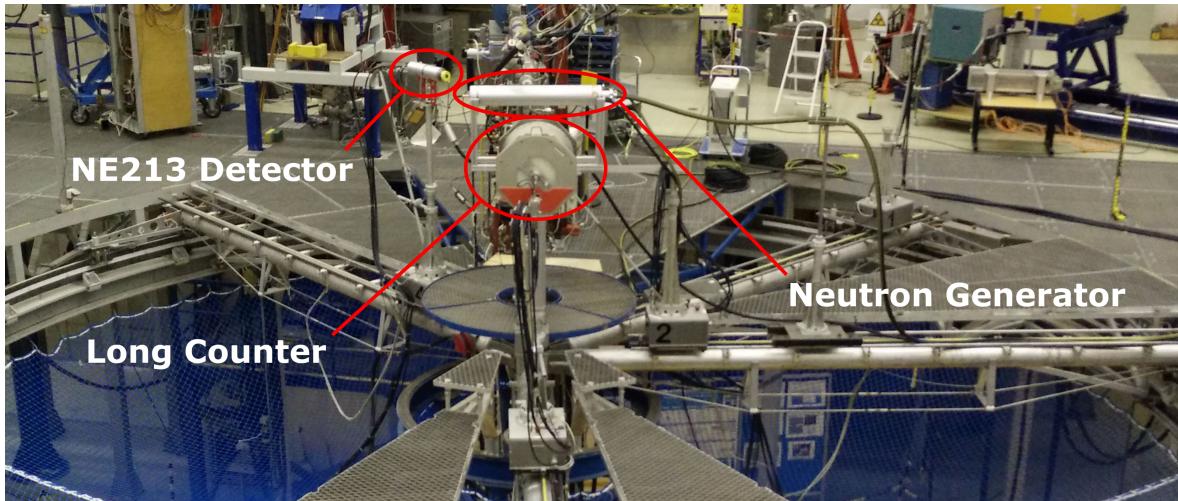


Figure 7.1: The experimental setup at PTB. The neutron generator is located at the center of the circle with the Long Counter at a radial distance of 1569 mm. On the left hand side, the NE213 detector can be seen at a distance of 2069 mm and at polar angle 230° .

7.3.4 Experimental Setup

Measurements of the neutron flux were performed both at Purdue University and at the Physikalisch-Technische Bundesanstalt (PTB). Fig. 7.1 shows the experimental setup at PTB. The neutron generator was placed with its horizontal axis at the center of the experimental facility. We define the polar angle such that the direction perpendicular to the axis of the neutron generator is 90° , as shown in Fig. 7.2. All distances in this section are measured from the center of the neutron generator to the front face of the detector in question.

To measure the neutron energy spectrum, a NE213 detector was placed at a radial distance of 865 mm from the generator, facing it at a 90° polar angle. These measurements were taken for a duration of 6.0 hours while the neutron generator was operated at 50 kV and 2.5 mA. To be able to perform background discrimination and subtraction, data were also collected for 21.1 hours while the neutron generator was off.

Additionally, data were taken with a shadow cone placed between the neutron generator and the Long Counter, where the Long Counter was placed at two different radii. These measurements were used to verify that the Monte Carlo simulations of the experimental setup correctly accounts for the number of neutrons scattered off the air into the Long Counter.

For the measurements of the angular emission (section 7.7), the Long Counter was placed on a radial arm at a distance of 1569 mm from the neutron generator

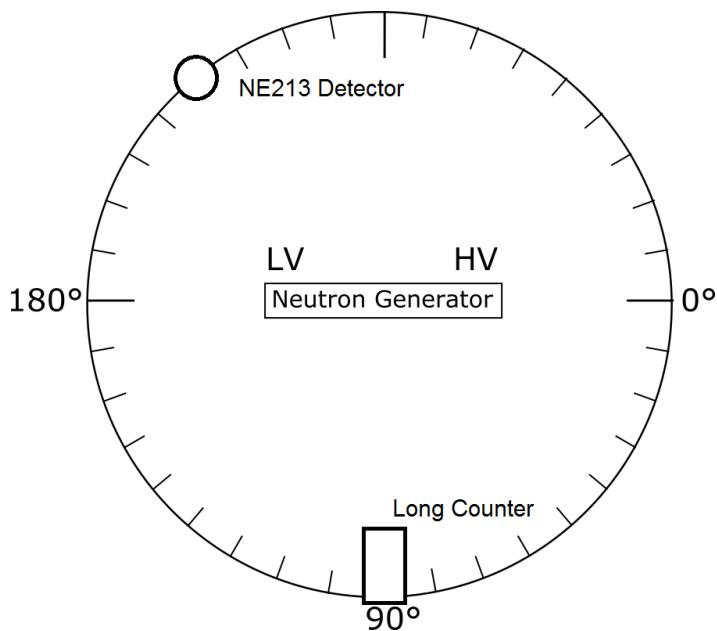


Figure 7.2: The coordinate system used at PTB. The schematic shows the neutron generator from above, with its high voltage (HV) end taken to be at 0° polar angle and its low voltage (LV) end taken to be 180° polar angle. Also shown are the Long Counter at 90° and the NE213 detector at 230° , as the detectors were placed when Fig. 7.1 was taken.

but at various polar angles. A $3'' \times 3''$ EJ301 detector was placed at a distance of 2000 mm at a fixed angle of 230° in order to have a permanent measurement of the stability of the generator during the angular scan.

The functional dependence of the neutron flux on the applied voltage and current were determined using three EJ301 liquid scintillator cells at Purdue in various orientations. The response of these detectors to 2.45 MeV neutrons has previously been characterized [163]. Furthermore, the functional dependence and absolute flux were measured at PTB with the Long Counter at a radial distance of 1569 mm and a polar angle of 90° .

7.4 MONTE CARLO SIMULATION

We have developed a detailed Monte Carlo simulation of the experimental setup in order to assist with the interpretation of the obtained data. The simulation was developed using the GEANT4 toolkit [93]. Technical drawings of the neutron generator and its interior, as well as the three different neutron detectors, were used to create a complete description of the major components that can produce significant scattering of neutrons.

7.4.1 Physics List in GEANT4

This simulation made use of version 9.4-patch02 of the GEANT4 toolkit. Since the energy of the neutrons of interest is below 20 MeV, we use the High Precision physics list, with G4NDL 3.14. This list contains cross-sections down to thermal energies in order to accurately describe the elastic, inelastic and capture processes of neutrons in the Long Counter. The radioactive α , β^+ , β^- or electron capture decays are simulated using G4RadioactiveDecay. Information about the half-lives, nuclear level structure, decay branching ratios, and the energies of decay processes are taken from the Evaluated Nuclear Structure Data Files (ENSDF) [168].

The tracking of particles in GEANT4 is divided into spatial steps. The length of these steps is automatically set depending on the energy and type of each particle, as well as the material it is propagating through. For each interaction, we record the position, deposited energy, particle type, initial energy of the particle, and the process responsible for the energy loss.

7.4.2 Neutron Generator Model

The description of the neutron generator in the GEANT4 toolkit is reproduced from technical drawings and information provided by the manufacturer, NSD/- Gradel Fusion. The important internal components included in the simulation are the reaction chamber and cathode, high voltage feedthrough, getter pump, and the water cooling system. The deuterium gas conditions inside the fusion chamber were modelled using the known pressure of the deuterium gas. We assumed homogeneous neutron production within the cathode volume, producing an isotropic neutron flux. This assumption is consistent with the measured angular neutron flux (section 7.7) and the energy spectrum measured at two different polar angles (section 7.5).

The energy spectrum used as an input in the GEANT4 simulation was calculated using the known energy-dependent differential cross-section of deuterium-deuterium fusion. The characteristic deuteron energy was described by a Gaussian distribution with a mean of 30 keV and sigma of 3 keV. While this is consistent with the applied high voltages of 40 kV and 50 kV, at which measurements of the energy spectrum were obtained, assumed average kinetic energies between 30 keV and 50 keV also fit the data.

The energy of the incident particle in a fusion event is determined by randomly sampling from the aforementioned Gaussian distribution. Angles ($\theta \in [0, \pi]$) were

randomly sampled to describe the emission angle of the neutron relative to the direction of the incident particle. The energies of the neutrons were determined from scattering kinematics given θ . Differential production cross-section information for neutrons from deuterium-deuterium fusion on thin targets in the center-of-mass frame were taken from [169]. Using the parametrization suggested therein, the differential production cross section can be described by

$$\frac{d\sigma}{d\Omega}(\theta) = \frac{d\sigma}{d\Omega}(0^\circ) \sum_i A_i P_i(\theta), \quad (7.2)$$

where $\frac{d\sigma}{d\Omega}(0^\circ)$ is the differential production cross section at 0° , and A_i are the recommended Legendre coefficients in the center-of-mass frame for a particle with incident energy E_d . Using Equation 7.2, we calculate the production cross-section of neutrons in the lab frame.

The resulting lab frame neutron energies are used to produce an energy histogram, where each entry is assigned a weighting factor determined by the production cross-section of the fusion event that produced a neutron of that energy. During Monte Carlo studies, neutron energies are produced by randomly sampling from this distribution.

7.4.3 Liquid Scintillator Detector Model

The two liquid scintillators EJ301 and NE213 were modeled as simple cylinders of the appropriate dimensions. The material properties of both liquid scintillators were taken to be those of EJ301 [159].

7

7.4.4 Long Counter Model

The De Pangher Long Counter simulation geometry was also reproduced from technical drawings. The concentric layers of polyethylene and borated polyethylene were implemented, along with the central volume of enriched BF_3 . In the Monte Carlo calculation, the emission of an α -particle following neutron capture on ^{11}B in the central tube is assumed to represent a neutron detection event.

7.5 NEUTRON ENERGY SPECTRUM

We use the $2'' \times 2''$ cylindrical NE213 detector to determine the energy spectrum of the neutrons produced by the generator. Since this detector only measures the

energy of the recoiling proton, there is no unique measurement of the incident neutron energy. Nevertheless two methods can be used to determine the incident neutron energy spectrum. Both require knowledge of the response function of the scintillator, which is simply a matrix that gives the distribution of the amount of scintillation light for a given incident neutron energy. We use a Monte Carlo-derived response function, which is based on measurements in monoenergetic and broad ns-pulsed neutron fields at PTB. The latter allows one to apply time-of-flight methods to select specific neutron energies [164, 170, 171].

If the initial neutron energy spectrum is known, the observed recoil spectrum can be calculated from the convolution of this response function with the incident spectrum, as we show in section 7.5.3. Alternatively, the observed recoil spectrum can be deconvoluted with the help of the response function in order to extract the incident neutron spectrum, as we do in section 7.5.4. The results that we obtain from both methods are in good agreement with each other.

7.5.1 Data Selection

The two parameters available for each event in the NE213 scintillator are the pulse height, which increases with the nuclear recoil energy, and the Zero Crossing Method pulse shape discrimination parameter [172], which allows one to distinguish nuclear and electronic recoil events. Fig. 7.3 shows the histograms for both background data and data taken with the neutron generator turned on. The pulse height is converted to recoil energy (electronic recoil equivalent) by Monte Carlo-matching of a ^{207}Bi gamma calibration spectrum. For proton recoil events, the non-linear relation between proton recoil energy and the resulting signal amplitude, known as the light output function, is considered in the neutron response function data [173]. A linear energy scale is assumed, which we verified to be correct at six energies ranging up to an electron-equivalent energy of 1546 keV_{ee} using the Compton edges of the calibration sources ^{207}Bi , ^{22}Na and ^{137}Cs .

We select neutron-induced events by cutting at the 99 % electronic recoil background rejection line in the pulse shape discrimination parameter space, as indicated in Fig. 7.3. We subtract the background rate in the selected region by computing the number of events that pass the cut in the background data set and using the appropriate scaling for live-time. We apply an energy threshold of 120 keV_{ee} electron-equivalent energy in order to stay above energies where characteristic X-ray emission from the lead surrounding the NE213 detector starts to dominate.

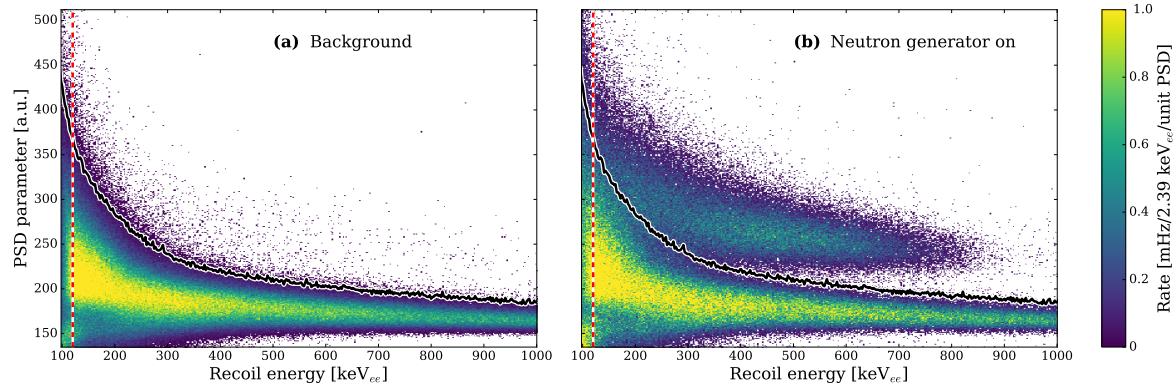


Figure 7.3: (a) Data taken with the NE213 scintillator from a 21-hour background run and (b) a 6-hour run with the neutron generator turned on, showing the pulse shape discrimination parameter as a function of recoil energy. A clear second neutron-induced band can be seen in the right panel. These recoils are selected using a 99 % background rejection cut in the pulse shape parameter, as indicated by the solid (black) line. The vertical dashed (red) line indicates the analysis threshold of 120 keV_{ee} electron-equivalent energy.

7.5.2 Cut Acceptance

At low recoil energies, the pulse shape discrimination becomes less efficient. Due to 99 % background rejection criterion, for lower pulse heights an increasingly bigger fraction of neutron events does not fall in the selected region of neutron events. We calculate the fraction of neutron events passing this cut (the *acceptance*) as a function of recoil energy as follows: for each energy bin, we subtract the live-time normalized background data (Fig. 7.3a) from the neutron generator data (Fig. 7.3b) and fit a Gaussian to the resulting nuclear recoil distribution. The neutron acceptance, shown in Fig. 7.4, is then taken as the area fraction of the Gaussian above the background rejection line. Also shown in the Fig. 7.4 is a smoothed interpolation that we use in the calculations that follow. We calculate the uncertainty on the acceptance by changing the fit parameters of the Gaussian within their uncertainty and re-computing the resulting acceptance.

7.5.3 Convolution

The initial neutron energy spectrum in the neutron generator fusion region, calculated as described in section 7.4.2, is shown in Fig. 7.5. As can be seen, a plasma fusion generator as used here does not produce a truly monoenergetic neutron spectrum. Due to the dependence of the fusion cross section on the neutron emis-

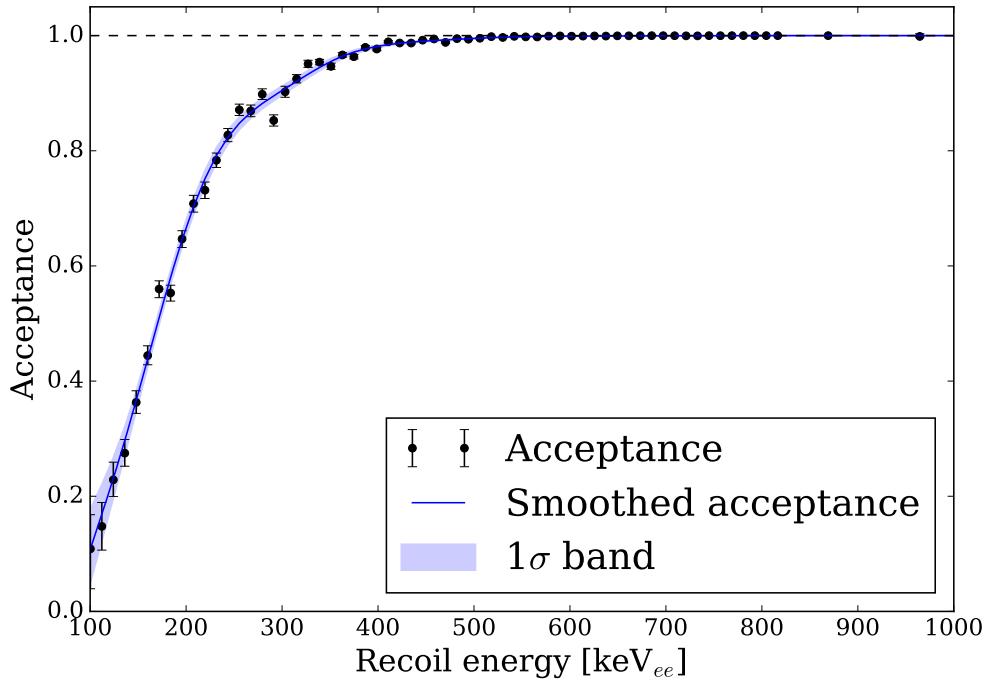


Figure 7.4: The acceptance of the pulse shape cut (shown in Fig. 7.3) as a function of energy. The Gaussian fraction above the pulse shape cut is shown by the (green) data points. The two points at the highest recoil energies are calculated with a wider bin width due to low statistics in the neutron band at these energies. The (blue) solid line is a lowpass-filtered interpolation of those data points. The 1σ uncertainty band on this acceptance is also shown (shaded blue).

sion angle relative to the momentum of the incident deuteron, the spectrum shows two peaks at 2.22 MeV and 2.72 MeV in the lab frame, corresponding to emission angles of 180° and 0° , respectively.

Neutrons from this spectrum are propagated from the fusion region using GEANT4. The resulting neutron energy spectrum at the NE213 detector position is also shown in Fig. 7.5, where we normalized the spectra to the energy range from 2.2 MeV to 2.7 MeV for ease of comparison. The most significant impact to the neutron energy spectrum comes from scattering in the cooling water that surrounds the neutron generator fusion chamber. This causes the long tail towards lower energies and gives rise to the asymmetric peak structure.

We calculate the expected pulse height spectrum observed by the NE213 detector from the incident Monte Carlo-derived neutron energy spectrum. To this end, we apply the response function described earlier and use the data-derived acceptance function (Fig. 7.4) to correct for acceptance losses at low pulse height energies.

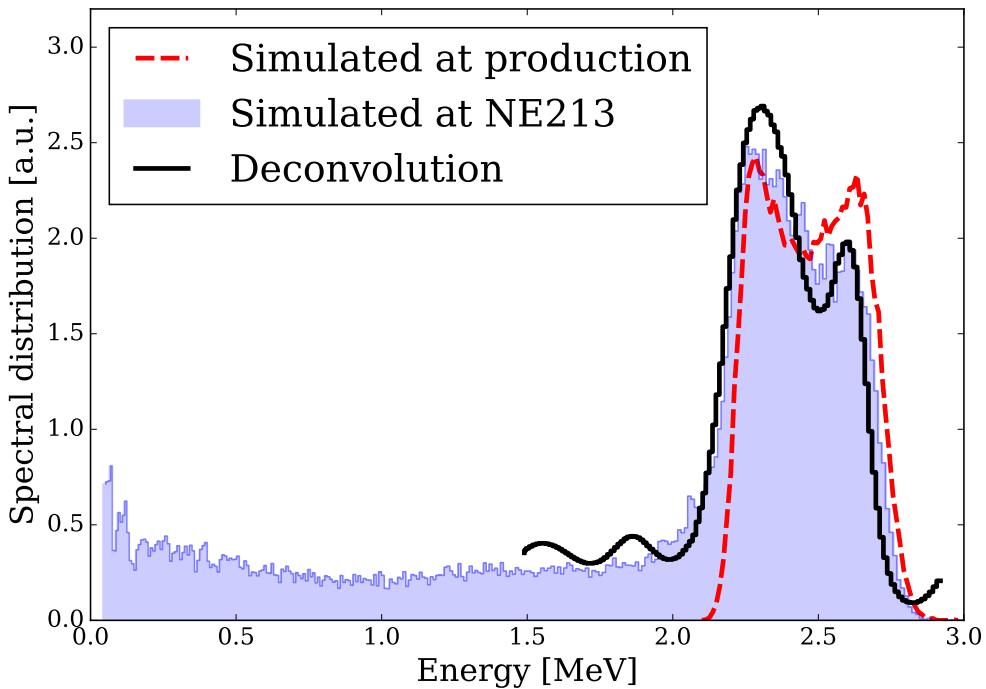


Figure 7.5: Shown are the simulated neutron spectrum in the neutron generator fusion region (red dashed line) and the resulting simulated spectrum at the NE213 detector (blue shaded histogram). Also shown is the result of the deconvolution of our data (black solid line), which is discussed in section 7.5.4. All spectra are normalized to the region 2.2 MeV to 2.7 MeV for ease of comparison.

We normalize the expected pulse height spectrum to the observed data using a χ^2 -minimization in the energy range $(120 - 900)$ keV_{ee} electron-equivalent energy, where the upper limit of the considered energy range is placed at the point where signal and background rates become comparable. The result of this normalization is shown in Fig. 7.6 with its 1σ uncertainty, and the residuals. We find good agreement between the expected and observed distributions, validating the assumed energy spectrum of produced neutrons shown in Fig. 7.5. A closer inspection of the residuals of the minimization indicates that our nuclear recoil acceptance may be slightly underestimated below 350 keV_{ee}.

In addition to the data selection method outlined in section 7.5.1, we repeated the analysis with no selection on the pulse shape discrimination parameter. The pulse height spectrum from the neutron generator is in this case computed by subtracting the background spectrum from the spectrum taken with the neutron generator turned on. Since this method requires no acceptance correction, it can be used as a cross-check of the cut-based analysis, however, with larger statistical

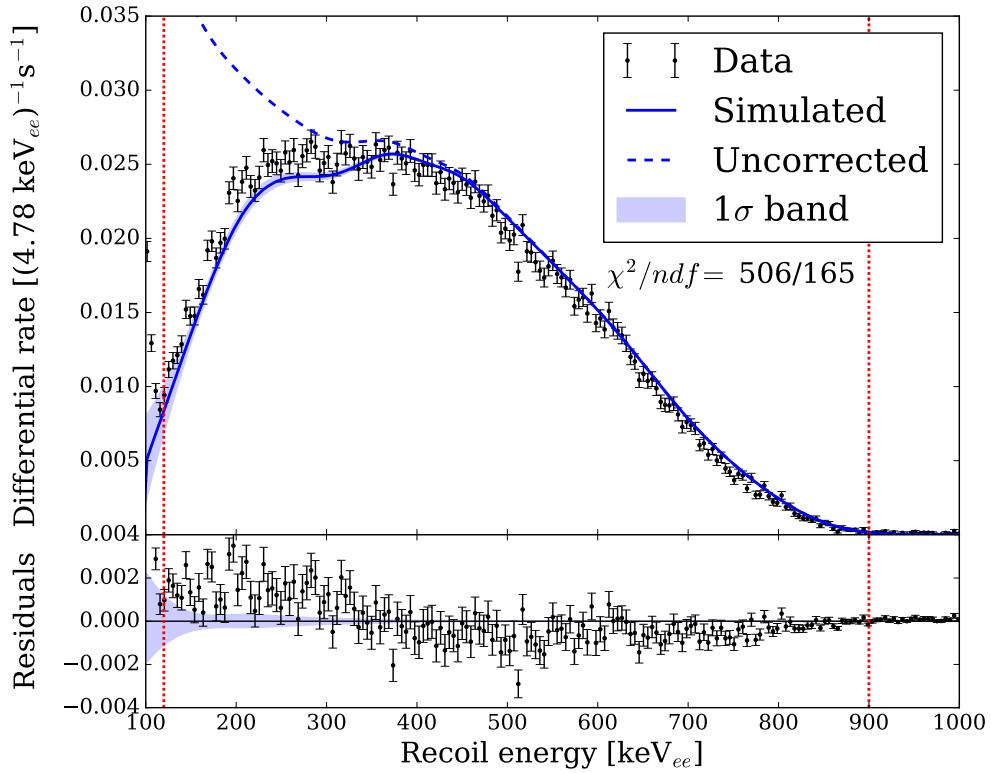


Figure 7.6: The observed background-rejected energy spectrum (green data points) together with the normalized expected distribution from simulation (solid blue line). The energy range considered for the normalization is indicated (red dotted lines). The reduced χ^2 of the fit is 3.07. An indication of the effect of the acceptance correction can also be seen (dashed line). The (light blue) band around the simulated distribution indicates the uncertainties from the acceptance function. The bottom panel shows the residuals between data and simulation after normalization.

uncertainty from the increased number of bins. Any gamma radiation caused by neutrons or Bremsstrahlung is ignored in this approach. By comparing the gamma-induced band for the neutron generator and background runs, we estimate that this contributes $\lesssim 10\%$ to the number of events between 300 and 800 keV_{ee}. We find that the data selected by this simple background subtraction is consistent with the cut-based analysis if this effect is considered.

7.5.4 Deconvolution

We determine the neutron spectrum through a deconvolution of the observed nuclear recoil pulse height spectrum. For this analysis, we use the same data selection,

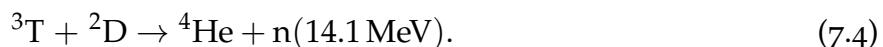
background treatment and acceptance correction as in section 7.5.3, but restrict the energy range to $350 \text{ keV}_{\text{ee}} - 950 \text{ keV}_{\text{ee}}$ electron-equivalent energy. The main purpose of the deconvolution is an independent confirmation of the expected line shape of the 2.45 MeV neutron peak resulting from reaction kinematics and Monte Carlo modeling of the setup. The recorded pulse height events below $350 \text{ keV}_{\text{ee}}$ do not contribute to this information but cause instabilities in the deconvolution process due to non-perfect neutron-gamma separation and systematic limitations in the precise determination of the response function in this energy range.

We use a combination of the GRAVEL [174] and MAXED [175] deconvolution codes, with GRAVEL providing the starting values that are then used for further refinement using the MAXED code. The allowed neutron energy range for this analysis is 1.49 MeV to 2.92 MeV, corresponding to the light output of the selected pulse height range.

The neutron spectrum obtained from this deconvolution is shown in Fig. 7.5. In agreement with the Monte Carlo simulation, this deconvolved spectrum shows a double-peaked structure around 2.4 MeV, rather than a purely monoenergetic neutron emission. At low energies ($< 2.0 \text{ MeV}$), an oscillatory feature appears. However, only a small fraction of the low-energy neutrons can induce a signal above our analysis threshold of $350 \text{ keV}_{\text{ee}}$. Consequently, small changes in the analysis (such as the considered energy range or data selection criteria) result in significant changes to the spectral shape below 2.0 MeV. We thus do not consider this neutron energy range any further. A small contribution is seen at the highest allowed energies ($> 2.8 \text{ MeV}$). As discussed in the next section, we attribute this contribution to the presence of high-energy neutrons from deuterium-tritium fusion.

7.5.5 High-Energy Neutrons

Upon inspection of the NE213 nuclear recoil data at high ($> 1000 \text{ keV}_{\text{ee}}$ electron-equivalent energy) energies, we observed a secondary population at energies above those expected from 2.45 MeV deuterium-deuterium neutrons. We attribute these high-energy recoils to deuterium-tritium fusion, produced in the neutron generator as a result of the reactions



Reaction (7.3) is equally likely to occur at 50 keV as the main neutron producing reaction (Eq 7.1) [176]. Since the deuterium-tritium fusion reaction has a cross section that is more than two orders of magnitude higher than that of deuterium-deuterium fusion, even a small tritium contamination can give a non-negligible amount of 14.1 MeV neutrons in the spectrum.

To test this hypothesis, we rebin the recoil energy spectrum as shown in Fig. 7.7. Due to the low statistics, the deconvolution code is prone to reconstruction artifacts. We therefore only use the convolution method, with 14.1 MeV neutrons as the initial spectrum originating in the neutron generator fusion volume. Again, scatters in surrounding materials were taken into account by propagating the neutron spectrum to the detector in GEANT4. As the separation between the electronic and nuclear recoil bands in the NE213 detector is excellent at these high recoil energies ($>1000 \text{ keV}_{ee}$), the acceptance is taken to be unity. The resulting simulated recoil spectrum is shown in Fig. 7.7, scaled to data. We find an excellent agreement between data and simulation, confirming our hypothesis of deuterium-tritium fusion taking place in the generator.

For a quantitative analysis, we calculate the ratio of the deuterium-tritium neutron flux to the flux of neutrons integrated between 2.0 and 2.8 MeV in the deconvoluted energy spectrum, taking into account the energy-dependent response of the NE213 detector. We arrive at a total ratio of $(5.5 \pm 0.3)\%$. This is consistent with the yield expected from tritium produced during the operation of the neutron generator prior to the data taking presented here. To determine the ratio of deuterium-tritium neutrons to the flux of all neutrons from deuterium-deuterium fusion, we calculate the fraction of neutrons produced in the GEANT4 MC simulation with energies below 2.0 MeV. In total 37.34 % of the simulated neutrons have incident energies below 2.0 MeV when they reach the detector. This is in agreement with the difference in flux measured by the NE213 detector and the Long Counter, discussed at the end of section 7.6. We conclude that the total ratio of deuterium-tritium neutrons is $(3.5 \pm 0.2)\%$.

7.6 NEUTRON FLUX

7.6.1 Relative Dependence

Measurements were taken at Purdue University to measure the functional dependence of the neutron flux on the operational parameters of the neutron generator, namely the high voltage V and the current I applied to the cathode. High voltages

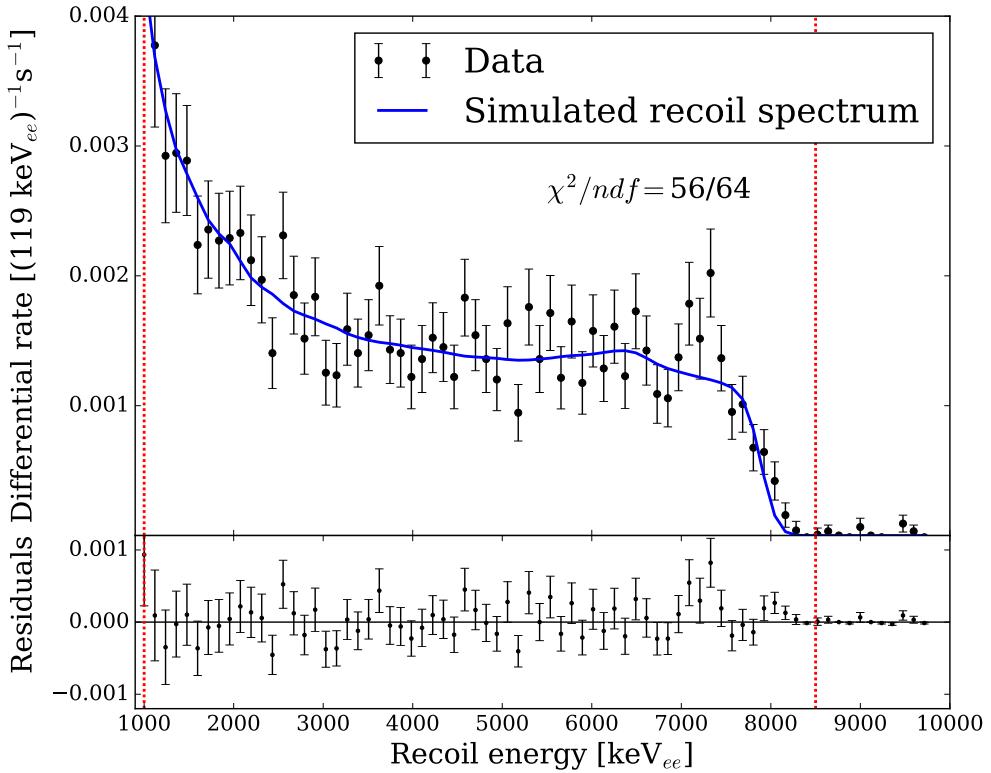


Figure 7.7: The observed spectrum at high recoil energies (green data points) together with the normalized distribution expected from 14.1 MeV neutrons from deuterium-tritium fusion (solid blue line). Again, the energy range considered for the normalization is indicated (red dotted lines) and residuals from the normalization are shown in the bottom panel. The reduced χ^2 of the fit is 0.875.

7

were set to between 30 kV and 52 kV, and currents were set to between 0.5 mA and 1.1 mA. Measurements of the high voltage, current, and getter pump temperature were collected in one minute intervals. Three EJ301 organic liquid scintillator detectors were used to measure the neutron flux. Because this experiment was conducted in a small lab, the resulting backscattering of neutrons from the walls prevents us from using this data to determine the absolute neutron flux. This does not however prevent measurement of the relative functional dependence on V and I .

A total of 107 datasets with a combined live-time of 55.3 hours were collected with the detectors level with the neutron-producing region of the generator. We have previously studied and improved upon the pulse shape discrimination using the EJ301 detectors. That work allows us to reduce the recoil energy threshold for this analysis to 50 keV_{ee} electron-equivalent energy at 99.5 % electronic recoil rejection, using a Laplace transform-based pulse shape discrimination parame-

ter [163].

We obtain a rate of nuclear recoils passing this rejection cut and fit a function of the form

$$F(V, I) = aV^b I^c. \quad (7.5)$$

As the neutron flux depends strongly on the applied high voltage, which varied during a run by up to several hundred volts, the voltage measurements are averaged together for each run via the expression

$$\langle V \rangle = \left(\frac{1}{n} \sum_{i=0}^n V_i^p \right)^{1/p} \quad (7.6)$$

The exponent p is determined to be 3.33 through an iterative process where a value is chosen, the functional dependence is calculated from the data, and the resulting voltage exponent used to re-average the voltage measurements. This was repeated until the values converged. Typical variations in the applied current were of order μA . Since the relative variations are much smaller and the flux depends only weakly on current, no averaging process was required for the current.

To estimate systematic uncertainties in the values of b and c , we investigate variations of the selection criteria and also compare the results from the three different detectors. We find the results to be robust against changes in the 50 keV_{ee} energy threshold requirement as well as against variations in the exponent p used when averaging high voltage measurements.

Due to scattering of neutrons from the walls, the value found for a is not an accurate measurement of the overall scaling of the function, but we find $b = 3.32 \pm 0.14$ and $c = 0.97 \pm 0.01$. Fig. 7.8 shows the normalized measured neutron rates as a function of cathode high voltage, at fixed current.

7.6.2 Absolute Flux

To measure the absolute neutron flux, a similar measurement was performed at the PTB, where the large experimental hall resulted in significantly less environmental scattering of neutrons. The Long Counter was used as the detector, placed at a distance of 1.569 m from the neutron generator, as measured to the front face of the detector. A total of 10 datasets were collected at different high voltage and currents settings over a wider range than measured at Purdue. The same fitting procedure described above was applied, resulting in $b = 3.31 \pm 0.08$, and $c = 1.00 \pm 0.02$,

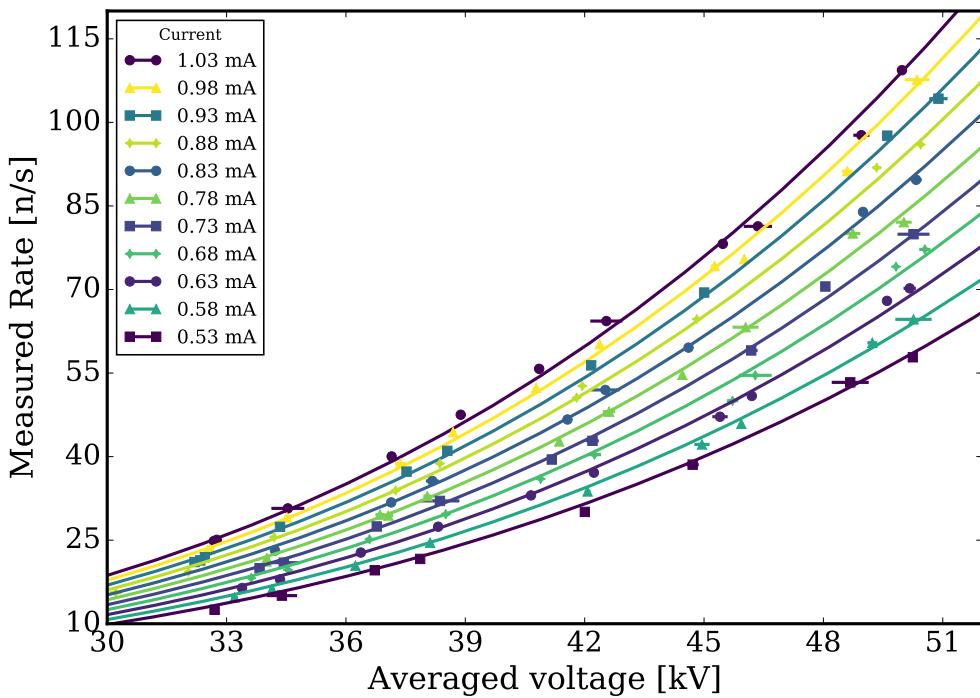


Figure 7.8: Measurements of the relative neutron flux produced by the neutron generator at Purdue University for various values of the operational parameters (cathode high voltage and current). The current at which each measurement was taken is represented by the color of the data point. The flux from the functional fit was found to be proportional to $(\text{Voltage})^{3.32}$ and $(\text{Current})^{0.97}$. The solid lines represent the predicted neutron flux as a function of cathode high voltage, at a fixed current.

consistent with the Purdue measurements.

We use the energy-dependent Long Counter [177] response, taking into account neutron contributions from both deuterium-deuterium and deuterium-tritium fusion, to correct the measured neutron flux for the space angle covered by the detector. Additionally we correct for the residual environmental scattering from the experimental hall. Furthermore, we also take into account the anisotropic neutron emission presented in the next section. We thus find for the absolute deuterium-deuterium neutron output outside of the neutron generator in the full 4π space angle (see equation 7.5) $a = (8.0 \pm 0.6) \times 10^{-2} \text{ s}^{-1}$ with V in kV and I in mA. Given the determined values of a , b and c , the measurements of neutron flux performed here span a range from 7.8×10^3 to $1.16 \times 10^5 \text{ s}^{-1}$. We expect that the functional dependence of the rate on the operational parameters is valid for neutron fluxes higher than those measured in this work.

As a cross check for consistency, we compare the measured neutron flux in the

NE213 detector and the Long Counter. For the Long Counter we use data collected at a polar angle of 90° and an effective radial distance of 1663 mm. Here we have included the distance to the effective center of the Long Counter, which is 94 mm behind the front face of the detector. The neutron generator was set to 2.0 mA and 50 kV, resulting in a measured neutron flux in the Long Counter of $(0.247 \pm 0.022) \text{ cm}^{-2}\text{s}^{-1}$. We compare this to data collected with the NE213 detector at an effective radial distance of 892 mm (which includes the distance to the effective center of the NE213 detector of 27 mm) and a polar angle of 90° . This yielded a measured neutron flux of $(0.643 \pm 0.064) \text{ cm}^{-2}\text{s}^{-1}$ between 2.0 and 2.9 MeV. Taking into consideration the full neutron energy spectrum (discussed in section 7.5.5) the total flux is $(1.06 \pm 0.11) \text{ cm}^{-2}\text{s}^{-1}$. This value is then corrected for the difference in radial distances between the two detectors and scaled for the different current settings using equation 7.5. The comparable neutron flux from the NE213 detector is thus determined to be $(0.246 \pm 0.025) \text{ cm}^{-2}\text{s}^{-1}$.

7.7 ANGULAR EMISSION OF NEUTRONS

The internal geometry of the neutron generator cylinder is azimuthally symmetric, but not along its axis. We therefore assume that the flux of neutrons is independent of azimuthal angle, and measure the neutron flux as a function of polar angle (compare Fig. 7.2). An angular scan was performed in steps of 10° ranging from 10° to 180° using the Long Counter. For these measurements, the neutron generator was operated at a high voltage of 50 kV and a current of 2 mA, and we collected approximately 1000 neutron counts for each measurement.

Data taken with both the Long Counter and the EJ301 detector at PTB were corrected for the neutron flux that was produced, given the average high voltage and current, as discussed previously. After correction, the measured number of neutrons in the EJ301 detector (which was held at a constant angle throughout) was found to have a standard deviation of 7.5 % across the angular scan measurements. The neutron flux at the location of the Long Counter, shown in Fig. 7.9, was calculated from the count rate using the energy-dependent response of the Long Counter [177].

A fourth order polynomial, of the form $F(\cos \theta) = A + B \cos^2 \theta + C \cos^4 \theta$, is fitted to the Long Counter measurements in order to parametrize the angular dependence of the neutron flux. The resulting fit parameters are $A = (0.237 \pm 0.006) \text{ cm}^{-2}\text{s}^{-1}$, $B = (-0.051 \pm 0.028) \text{ cm}^{-2}\text{s}^{-1}$ and $C = (-0.130 \pm -0.025) \text{ cm}^{-2}\text{s}^{-1}$.

We simulate the expected angular dependence using the detailed neutron gene-

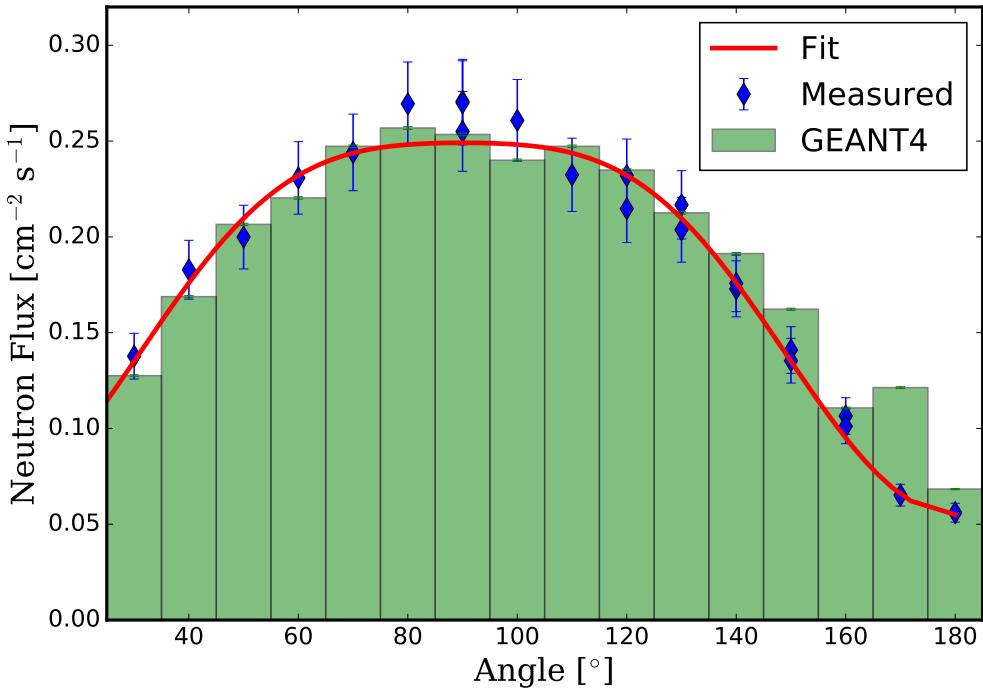


Figure 7.9: Measured neutron flux as function of polar angle. Data taken with the Long Counter is shown (blue diamonds) together with the angular neutron flux dependence predicted by a detailed GEANT4 simulation of the neutron generator (green bars). A fourth order polynomial fit to the data is shown as well (red line), parametrizing the measured dependence.

rator setup in GEANT4. The Long Counter is placed at the correct distance and angle of each respective measurement. The EJ301 detector is also present in the simulation to account for any scattering off of the EJ301 cell into the Long Counter as the latter is placed closer to the EJ301 cell for large polar angles. At each angle we simulate 10^8 neutrons distributed homogeneously throughout the cathode and emitted isotropically with the energy spectrum previously calculated. The GEANT4 Monte Carlo data sets are expressed in terms of a neutron flux, using the energy-dependent response function of the Long Counter derived from GEANT4. The resulting simulated angular dependence is also shown in Fig. 7.9.

The simulation results agree well with data. We conclude that the measured angular emission spectrum is well described by isotropic production of neutrons within the volume of the cathode field cage, with any angular variations being a result of the interior geometry and composition of the neutron generator.

7.8 CONCLUSIONS

We have performed the first characterization of the neutron flux produced by a deuterium-deuterium plasma fusion neutron generator. Our interest lies in the application of this generator as a nuclear recoil calibration source for a sensitive dark matter scattering experiment [68]. For this application, accurate knowledge of both the energy spectrum and the absolute flux are mandatory.

We found that the energy spectrum is not strictly monoenergetic, but contains two peaks at 2.22 MeV and 2.72 MeV. These are understood to be caused by the dependence of the neutron energy on the neutron emission angle relative to the momentum of the incident deuteron. Running this generator produces small quantities of tritium that resulted in a measurable flux of 14.1 MeV neutrons due to deuterium-tritium fusion in the plasma. These deviations from an ideal, monoenergetic spectrum will have to be taken into account in any measurement using deuterium-deuterium generators that aims to use the energy information of the incident neutron.

We have also characterized the dependence of the absolute neutron flux on the applied high voltage and current, as well as the dependence of the emitted neutron flux on polar angle. Monte Carlo simulations showed that the angular distribution of the neutron flux is affected by transmission of neutrons through the generator housing. The measured angular distribution is consistent with an isotropic neutron source inside the generator. Taken together, knowledge of this parametrization will allow us to break the degeneracy in calibration between decreasing acceptances of a detector and varying flux output from the neutron source [178].

We have developed a Monte Carlo simulation of the neutron generator characterized in this work, which accurately predicts the emitted neutron flux. The simulation is able to reproduce both the measured angular emission spectrum and energy of emitted neutrons. Thus we have a predictive model of the behavior of the neutron generator as a nuclear recoil calibration source for other applications.

7

ACKNOWLEDGEMENTS

We thank the manufacturer NSD/Gradel fusion for many useful discussions and for providing technical drawings of the neutron generator. This work is supported by grants #PHY-1206061, #PHY-1209979, #PHY-1412965 and #PHY-1650021 from the National Science Foundation (NSF) and carried out under a cooperation agreement between Purdue University and PTB. This work is supported by the research pro-

gram of the Foundation for Fundamental Research on Matter (FOM), which is part of the Netherlands Organization for Scientific Research (NWO) (grant number FOM VP139). JP is supported by scholarship #SFH13071722071 from the National Research Foundation (NRF).

NUCLEAR RECOIL CALIBRATION AND DARK MATTER SEARCH WITH XENON1T

The discrimination of electronic recoil (ER) and nuclear recoil (NR) interaction types is essential to the dark matter search of liquid xenon TPCs, as the most prominent background contribution comes from sources that give ER signals (see section 2.5). The interaction discrimination is based on the values of cS_1 and cS_{2b} . Compared to ERs, NR events typically generate a relatively small S_2 , so that the ER and NR distributions in (cS_1, cS_{2b}) are distinct. These distributions are called the ER and NR band. The band positions and their spread are measured and fitted with calibration sources. The process of measurement and subsequent band fit are called ER and NR calibration.

The ER band calibration of XENON1T is performed with ^{220}Rn , giving internal beta radiation spread out throughout the detector volume. For the NR calibration, the neutron generator that was described in chapter 7 can be used. Due to unforeseen technical issues, the neutron generator became operational well after data taking for the first science result of XENON1T finished, so that a $^{241}\text{AmBe}$ source was used instead [39]. For the second, updated result, a re-calibration of the NR band was required due to the lower drift field [5]. For this result, which comprises the most stringent dark matter exclusion limit to date, a combination of neutron generator and $^{241}\text{AmBe}$ data was used as the input of the NR band fit. The results from this fit are used in the likelihood models that yield the dark matter exclusion limit.

In this chapter, the data taken with the neutron generator is shown and compared to the ER data from the ^{220}Rn calibration. The results of the 1 tonne-year dark matter search of XENON1T are then discussed. For details of the operating principle of the neutron generator, see chapter 7.

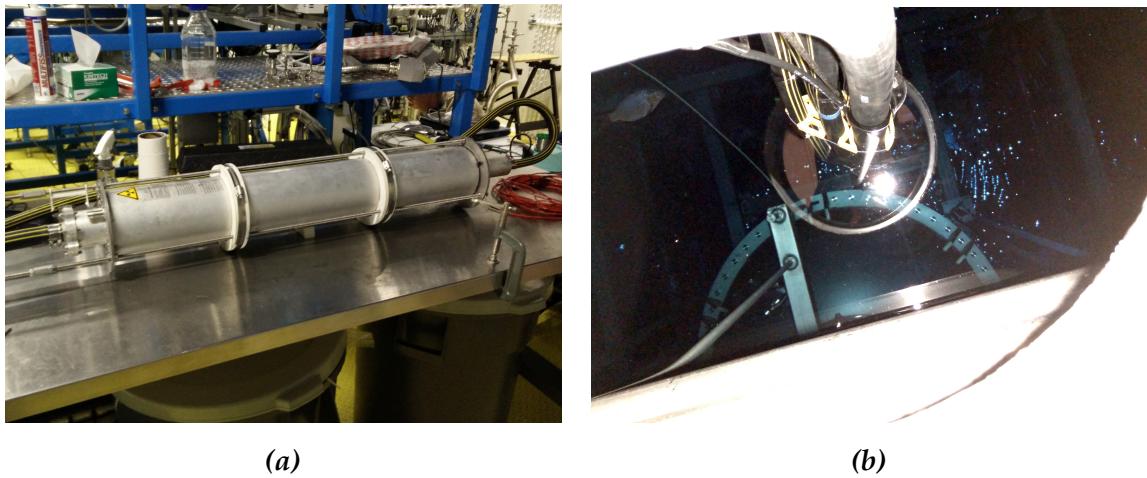


Figure 8.1: The neutron generator as used for the NR calibration of XENON1T. Figure 8.1a shows the neutron generator prior to deployment in the water tank, with all cables and hoses mounted into the holding structure. Figure 8.1b shows the neutron generator as it is lowered into the water tank, next to the outer vessel containing the XENON1T TPC.

8.1 NEUTRON GENERATOR INSTALLATION AND DATA ACQUISITION

All neutron sources that are commonly used for the NR calibration of liquid xenon TPCs are external sources: the neutrons are generated outside the detector and have to cross the detector materials before they can interact in the active volume. In the case of XENON1T, the TPC is surrounded by two nested steel vessels, which are suspended in a large tank filled with deionized water. The water tank provides an important shielding to gamma and neutron radiation from the outside, and simultaneously acts as a veto detector for cosmic muons that may generate neutrons in reactions with the material around the TPC. However, the water surrounding the outer steel vessel simultaneously provides a challenge for the neutron calibration. Since fast neutrons have a short mean free path in water (less than 6 cm for 2.45 MeV neutrons [50]), the source is moved close to the vessel. This means that the active part of the neutron generator is suspended in the water tank. Figure 8.1 shows how this is achieved. There are two coolant hoses and two cables attached to the neutron generator, which are clamped into a holding structure made of steel and plastic parts (see figure 8.1a). This structure is lowered into the water tank, where it hangs from a steel wire that is attached to a winch mounted on a flange on the top of the tank (see figure 8.1b). In this way, the neutron generator can be maneuvered close to the outer vessel of XENON1T. The distance to the water tank is monitored by a laser proximity sensor mounted on the holding structure of the neutron generator.

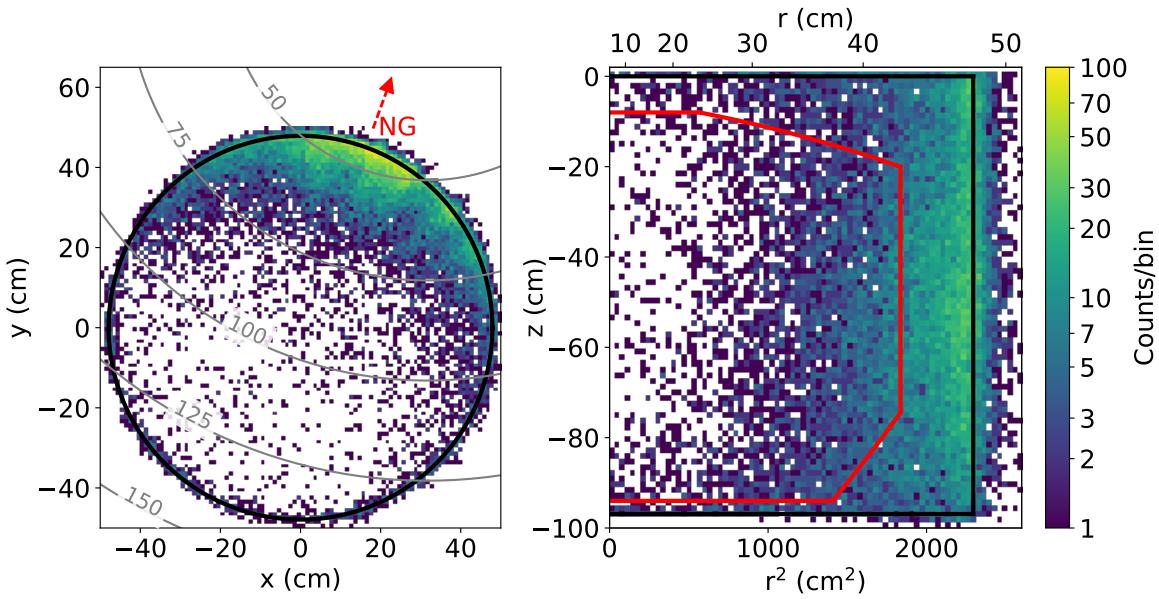


Figure 8.2: Density distribution of the reconstructed interaction position in (x, y) (left panel) and (r^2, z) (right panel) of all events in the neutron generator calibration datasets, with all cuts except the fiducial volume cut applied. The red arrow in the left panel points towards the location of the neutron generator and the labeled gray lines give the distance towards it in cm (in the plane of the neutron producing part). Most of the events are found close to the neutron generator, consistent with the expectation based on Monte Carlo simulation. The black line in both panels indicates the edge of the TPC volume. In the right panel, the red line indicates the fiducial volume; events that are contained within this boundary are used for the NR band fit.

8.2 DATA SELECTION AND BAND FIT

The first NR calibration data with the neutron generator was taken in May 2017 and consisted of over 50 hours of exposure time. Figure 8.2 shows the reconstructed interaction position of the events passing all cuts except for the fiducial volume cut. This includes data quality cuts, single scatter cuts and a low recoil energy selection. The position of the bulk of the interactions is found close to the neutron generator (as seen from the gray lines of equal distance) and the obtained position distribution is consistent with Monte Carlo simulation. The black line in both figures indicates the edge of the TPC volume. The red line in the right panel of figure 8.2 indicates the 1.3-tonne fiducial volume cut used for the dark matter exposure [5]. This cut is required to remove background events coming from detector materials and events that suffer from incomplete charge collection near the detector wall. For the NR

band fit, the same fiducial volume was used, so that only events contained within this volume are selected. The limited range of neutrons in xenon means that most events are close to the wall, so that most neutron events are also cut by the fiducial volume cut.

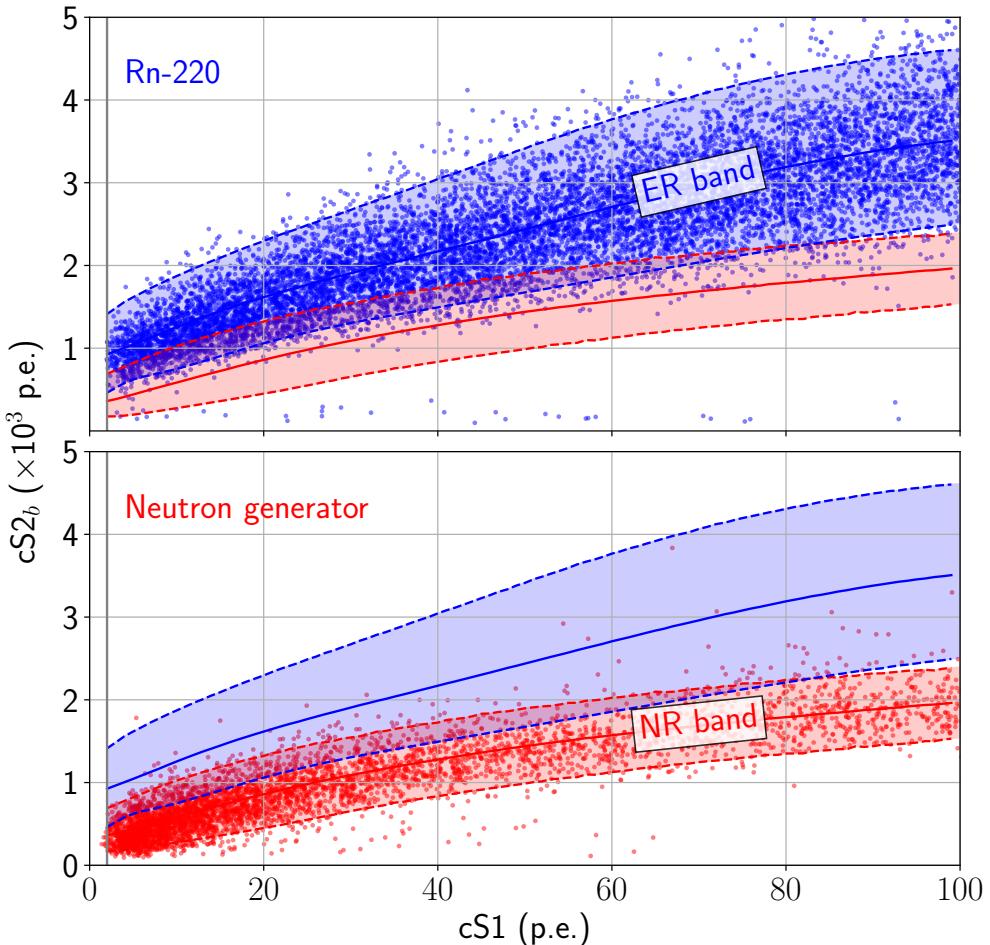


Figure 8.3: The position of events passing all cuts in (cS_1, cS_{2b}) parameter space, for ^{220}Rn (ER, top panel) and neutron generator data (NR, bottom panel). The solid lines indicate the mean of the ER (blue) and NR (red) band and the dashed lines show the mean $\pm 2\sigma$. The band position is determined in a simultaneous fit that incorporates detector effects as well as the microphysics.

The values of cS_1 and cS_{2b} for all selected events are plotted in the bottom panel of figure 8.3. The ER calibration dataset uses the internal ^{220}Rn source and is shown in the top panel. The observed event distribution depends on many parameters, of which some are related to the microphysics and others are related to properties of the detector (such as the position reconstruction resolution and the double photo-electron emission probability) and the analysis (such as the peakfinding efficiency

and the total cut acceptance). Since some of the parameters are shared between the two calibration types, the ER and NR bands are fitted simultaneously using a Bayesian fitting framework. The NR light and charge yields are parametrized with a model taken from NEST [125]. The resulting band positions are indicated with the colored lines in figure 8.3: the solid lines show the mean band position and the dashed lines show the position $\pm 2\sigma$. The obtained separation between the bands enables a rejection of 99.7% of the ER events at an NR acceptance of approximately 50% in the energy range of interest for the dark matter search (between 3 and 70 p.e. in cS1) [5].

8.3 DARK MATTER SEARCH RESULTS

The dark matter search data of XENON1T is split into two parts. The first data taking period ('Science run 0') consisted of 34.2 live days, and the resulting analysis set the world's most stringent spin-independent exclusion limit above a WIMP mass of 10 GeV [39]. After an earthquake, the drift field was reduced due to the appearance of single-electron S2 signals at the cathode voltage used for the initial science run. The changing conditions prompted the definition of a second science run ('Science run 1'). The second science run consisted of 246.7 live days accumulated over roughly a year of data taking. During this time, the detector was taking data almost continuously, including regular PMT gain calibrations, calibration with $^{83\text{m}}\text{Kr}$ and ^{220}Rn and neutron calibrations taken with $^{241}\text{AmBe}$ and the neutron generator. The analysis combining the two science runs is published in [5].

The second result improves upon the initial one not only due to the significantly longer live-time, but also due to several improvements to the analysis. The field inhomogeneity correction and the modeling of events coming from the wall ('surface' events, see section 2.5.3) made it possible to increase the fiducial mass from 1.0 t to 1.3 t. The combination of the data from both science runs and the increased fiducial mass meant that the total exposure (live-time times fiducial mass) for the first time crossed the value of 1 tonne-year, truly ushering in the era of tonne-scale detectors.

Another improvement to the analysis has been to the background and signal models. These models describe the likelihood for signal and various background components as a function of the analysis dimensions. In the first analysis, the modeled dimensions were cS1 and cS2_b. However, some backgrounds exhibit a spatial dependence, so that a better signal-background distinction can be obtained by including the spatial coordinates into the likelihood. Since the surface background has a strong dependence on the radial dimension, the radius R was taken as an

additional model dimension in the new analysis. Furthermore, the Z dimension was later included to take backgrounds with a strong Z-dependence into account; in particular, neutrons that scatter once below the cathode and once more in the fiducial volume (so-called ‘neutron-X’ events) are expected near the bottom part of the fiducial volume. Since including a fourth analysis dimension into the fit models is computationally highly expensive, this dimension was included in only two bins, so that an inner volume and outer volume were defined. This division was triggered by a neutron-like event found near the bottom of the detector after unblinding. All signal and background models are thus described in terms of cS_1 , cS_{2b} , R and Z .

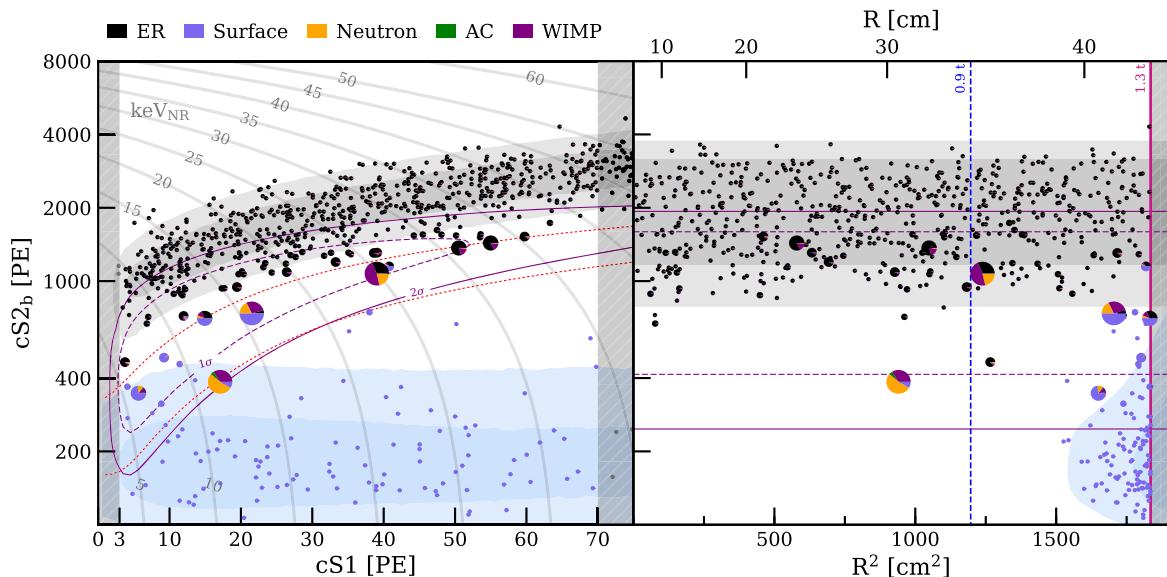


Figure 8.4: The events in the 1 tonne-year dark matter search exposure in the analysis space (cS_1 , cS_{2b}) (left panel) and (R^2 , cS_{2b}) (right panel). The pie charts indicate the relative probabilities of signal and several background components of the best-fit model (200 GeV, $\sigma_{SI} = 4.7 \times 10^{-47} \text{ cm}^2$), indicated by the color code in the top left. Events with a relatively large signal probability are shown with larger pie charts. The black and blue regions indicate the distribution of the ER and surface background, respectively. The signal contours are indicated by the purple dashed and solid lines. Taken from [5].

Figure 8.4 shows the events passing all cuts from the 1 tonne-year exposure in (cS_1 , cS_{2b}) (left panel) and (R^2 , cS_{2b}) (right panel). Both panels show the same events, but in different slices of the four-dimensional analysis space. The observed data is fit with a model containing signal and various background contributions. This model describes the likelihood for each event as a function of the analysis dimensions. The best-fit model is then obtained by maximizing the combined like-

lihood of the observed data, and contains a nonzero signal component for most WIMP masses. In figure 8.4, the best-fit model for a WIMP mass of 200 GeV and $\sigma_{SI} = 4.7 \times 10^{-47} \text{ cm}^2$ is shown in various ways. First, each of the events is shown in the figure by a pie chart that indicates the relative contribution of each likelihood component for this event. The color code is given in the top left of the figure. Events with low signal probability are drawn as small pie charts (and therefore resemble points), while events that are more signal-like are drawn larger. The shaded gray and blue bands indicate the 1σ and 2σ regions of the ER and surface background, respectively. The best-fit WIMP model signal contours are indicated by the dashed (1σ) and solid (2σ) purple lines.

The WIMP mass is an unknown factor that changes the shape of the recoil spectrum and therefore influences the likelihood function. The likelihood therefore has to be recomputed as a function of WIMP mass and cross section. From the likelihood as a function of cross section, either a two-sided confidence can be reconstructed (in case of a detection) or a one-sided exclusion limit can be reported. A 90 % confidence level is used, meaning that the experiment would include (not exclude) the true (mass, cross-section) in its confidence interval in 90 % of the cases, regardless of what it is [179]. In the case of no signal, the exclusion limit can be simulated, giving an expected limit and its spread. The median value of the expected limit is quoted as the sensitivity and its spread is called the sensitivity band. Prior to unblinding, the decision was made that a 3σ deviation above the sensitivity would be reported as a two-sided confidence interval.

Figure 8.5 shows the obtained exclusion limit for the analysis in [5]. The solid black line is the obtained limit. Over the full mass range displayed in the figure, the black line surpasses the limits of the previous experimental results, which means that this analysis is the most stringent to date for WIMP masses above 6 GeV. The green and yellow regions show the 1σ and 2σ sensitivity bands, respectively. The limit is fully contained within the 2σ sensitivity band, so that no statistically significant excess of events can be claimed.

The inset of figure 8.5 shows the limits and sensitivity bands normalized with respect to the median sensitivity of the analysis from [5]. The red and blue regions indicate the 1σ sensitivity band from LUX and PandaX-II. From this figure, it is evident that the difference between the LUX and PandaX-II *limits* (red and blue lines) and the current best limit (black line) at high mass is relatively minor, but the difference between the median *sensitivity* (shown by the blue dashed line for PandaX-II and the middle of the red band for LUX) and the current best median sensitivity (dashed black line at $y = 1$) is much larger: approximately a factor 7. This effect is

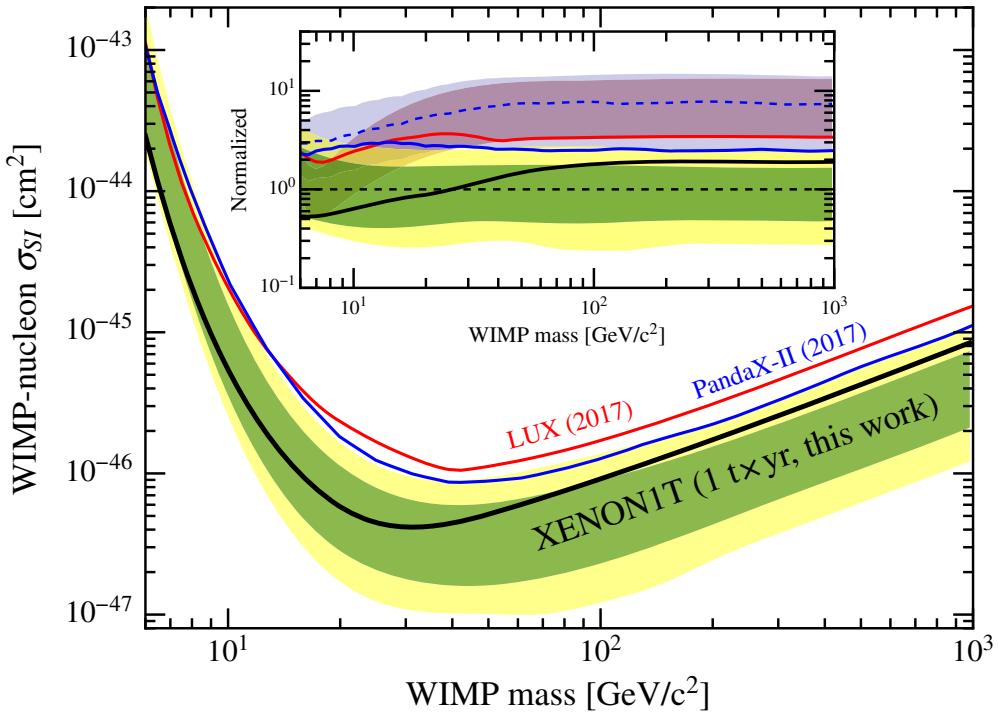


Figure 8.5: The obtained exclusion limit from the 1 tonne-year exposure of XENON1T (solid black line). The red and blue lines show the results from LUX [63] and PandaX-II [114]. The green and yellow regions are the 1σ and 2σ sensitivity bands, respectively. The inset shows the limit normalized to the median sensitivity of this analysis, including the 1σ sensitivity band of the XENON1T, LUX and PandaX-II (green, red and blue bands, respectively). Taken from [5].

due to the large statistical uncertainty on the limits that comes from background fluctuations common to rare-event searches such as these. Simply put, the last two experiments seemed to have been ‘lucky’, observing an underfluctuation of the expected background that sets the limit at the bottom part of the sensitivity band, while a slight overfluctuation is observed for the XENON1T result. It is interesting to note that this is the first time since 2011 that a world-leading experiment reports a limit that is higher than the sensitivity band $+1\sigma$ [180], with most experiments reporting limits mostly tracing the bottom [39, 63, 89, 114, 149, 181] or middle [88, 182] of the sensitivity band.¹ This appears to suggest a systematic overestimate of the background, which is a conservative approach for dark matter detection but artificially lowers the limit in the case of a null result.

¹ These results are not all independent, since updates of results from the same experiment usually include the previous data. Still, the outcome of all these experiments would make for an interesting meta-statistical exercise.

OUTLOOK

Since the development of the technology of dual-phase xenon TPCs, there has been a rapid improvement in the sensitivity to dark matter interactions. Over the past 10 years, the limits have roughly decreased by a factor of 2 each year. This is illustrated in figure 9.1, which shows the minimum of the WIMP-nucleon exclusion limit of various experiments.

The field of dual-phase xenon TPCs has been and continues to be rapidly expanding, with many advances made over the years. All these are essential to the experimental effort, as the increase of the detector size needs to be accompanied by a corresponding decrease in background. In addition, there has been development in the understanding of the low-energy signals in liquid xenon. For example, the light and charge yields have been measured to great detail [126, 138, 150, 183]. The worldwide data is fit in the NEST simulation framework, which has become accepted and widely used [144]. New calibration techniques have been developed and continue to be researched [65, 184, 185], and the understanding of other properties of interest to dual-phase xenon TPCs is also advancing [71, 78, 96]. The publications contained in this thesis aim to contribute to this expanding body of knowledge by the systematic measurement of the scintillation pulse shape dependence on particle type, energy and field (chapter 5, [2]) and the variation of many relevant properties with applied field (chapter 6, [3]).

While the improvement rate of xenon TPCs is remarkable, the consistent null results in the search for dark matter are agonizing, and are prompting different alternative analyses, focusing on other signal types, WIMP masses or other particles. Examples of these are the search for annual modulation [186, 187], axions or axion-like particles [46, 188] and inelastic dark matter [189]. Other analyses focus on lowering the energy threshold by using only the S₂ signal [44, 190], or by looking for signals from Bremsstrahlung or the Migdal effect [191]. The drive for exploring as much of the data as possible usually goes beyond the standard WIMP hypothesis. However, this is fully justified as the parameter space for dark matter is unprecedented in its vastness and theoretical models are likewise abundant, and in the absence of signal we should look at the data in any way that we can [192].

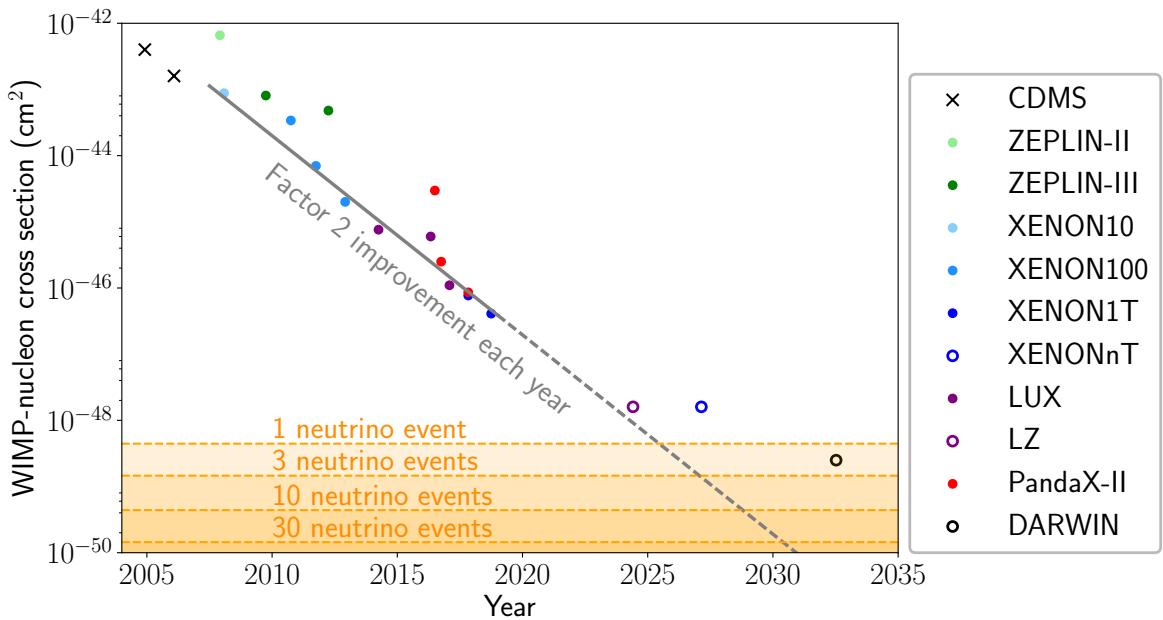


Figure 9.1: Minimum of the WIMP-nucleon exclusion limit over time. The filled circles show the limits obtained by dual-phase xenon TPCs, which have taken over from CDMS (black crosses) as the most sensitive experiment since XENON10. The limits shown are from CDMS [193, 194], ZEPLIN-II [195], ZEPLIN-III [196, 197], XENON10 [141], XENON100 [89, 180, 198], XENON1T [5, 39], LUX [63, 88, 182] and PandaX-II [114, 149, 181]. The limits roughly trace the gray line, which follows a factor 2 improvement each year. This illustrates the stunning improvement of the dual-phase liquid xenon technology in the past 10 years. The open circles show the minimum of the projected final limit for XENONnT [68], LZ [199] and DARWIN [43].¹ For DARWIN, neutrino backgrounds start to become a limiting factor, slowing down the improvement of the limit. The dashed horizontal lines show the number of expected neutrino events leaking into the parameter space for a 50 GeV WIMP [35]. Source code available at [10].

The next few years bear the promise of a new generation of larger TPCs with even lower backgrounds. The XENONnT [68] and LZ [143] experiments, both of which are currently under construction, will feature a similar detector mass and are expected to increase the sensitivity approximately by 2020.² The increased size requires a decrease in backgrounds, mainly from internal radon contamination and from radiogenic neutrons. The radon concentration can be decreased by several

¹ Assuming a fiducial mass of 4 t, 5.6 t and 30 t for XENONnT, LZ and DARWIN, assuming data taking starts 2020, mid-2020 and 2023, respectively, and assuming a livetime fraction of 70 % for all experiments.

² The planned successor of PandaX-II, PandaX-4T [200, 201], will have a detector mass slightly below that of XENONnT and LZ, and will likely only be competitive if it starts data taking earlier.

methods, such as the selection of even more radio-pure detector materials and by cryogenic distillation. The neutron background is mitigated by specialized detectors surrounding the cryostat, using a liquid scintillator subdetector in the case of LZ and gadolinium-doped water in the case of XENONnT. The development of these techniques is currently in full swing. After the next ‘round’ of detectors, there will likely be a global effort for the construction of DARWIN, a 40 t TPC that is intended to reach the neutrino floor and will thus become background-limited due to neutrino interactions (although the planned successor of PandaX-4T, PandaX-30T, is intended to reach a similar sensitivity [202]). Because of this irreducible background, this would be the moment of truth for liquid-xenon based detectors: either dark matter will be found in the next two generations of detectors, or one has to look for it in different ways.

The theory of dark matter is a firmly established theory with many astronomical observations that bear its fingerprint. It is in our nature as mankind to seek knowledge to understand how nature works and what the Universe is made of, and the question of dark matter is one of the clearest examples of a manifestation of this question. Dual-phase xenon TPCs will continue to contribute to the hunt for the dark matter particle and either yield a detection or tighten the constraints on what dark matter can or cannot be.

POPULAIRWETENSCHAPPELIJKE SAMENVATTING

(VAN: A SPARK IN THE DARK)

“*Waar is alles van gemaakt?*” Hoewel deze vraag zo simpel lijkt, is het beantwoorden ervan enorm moeilijk. We kunnen deze vraag op twee manieren interpreteren: als een vraag over de grootste afstanden (waar is het heelal van gemaakt?) en als een vraag over de kleinste afstanden (waaruit zijn objecten opgebouwd?). Als we, ten eerste, gaan kijken naar de kleinste afstanden, dan blijkt dat alles opgebouwd is uit dezelfde bouwstenen. Objecten zijn gemaakt van moleculen. Als we deze opbreken, dan komen we uit op atomen, wat letterlijk ‘ondeelbaar’ betekent. Ook deze deeltjes zijn, in tegenstelling tot hun naam, op te delen in nog kleinere deeltjes: elektronen, protonen en neutronen. De protonen en neutronen zijn dan weer opgebouwd uit quarks.

Voor zover we weten zijn we met de quarks en elektronen aangekomen op de kleinste schaal. Deze en andere, meer exotische deeltjes, noemen we elementaire deeltjes. De deeltjesfysica heeft momenteel een theorie ontwikkeld, het Standaardmodel, die deze deeltjes en hun interacties met elkaar extreem gedetailleerd beschrijft. Zelfs de meest precieze metingen, zoals die worden gedaan bij de Large Hadron Collider in Genève, zijn volledig in overeenstemming met de voorspelling hoe deeltjes zich gedragen volgens het Standaardmodel.

Om de vraag “*Waar is alles van gemaakt?*” te beantwoorden, kunnen we ook kijken naar de grootste afstandsschalen, namelijk die van astronomie en kosmologie. Kosmologen hebben hun eigen modellen om de ontwikkeling van het heelal en astronomische waarnemingen te voorspellen. De afgelopen tijd is het hun gelukt om tot een precisieniveau te komen dat bijna even sterk is als dat van de deeltjesfysica. Er is inmiddels een kosmologische versie van het Standaardmodel ontwikkeld, die de Oerknal en de ontwikkeling van het universum zeer gedetailleerd beschrijft. Er is alleen één probleem: alles wijst erop dat er materie moet zijn die niet zichtbaar is. Dit wordt donkere materie genoemd. Zonder deze materie zouden de kosmologische modellen in duigen vallen en kunnen we onze astronomische waarnemingen niet verklaren. Donkere materie wordt donker genoemd omdat het niets met licht doet: het produceert, reflecteert en absorbeert het licht niet. We kunnen het daarom

ook niet zien met onze telescopen. We kunnen het bestaan van donkere materie alleen afleiden doordat het interactie met de zwaartekracht heeft, en daardoor andere materie aantrekt en het pad van lichtstralen afbuigt. Maar liefst 85% van alle materie in het universum is donkere materie, en we hebben geen idee wat het is.

DE ONGRIJPBARE DEELTJES

Als we de combinatie maken van kosmologie en deeltjesfysica, dan wordt onze beginvraag dus specieker: *“Uit wat voor deeltjes bestaat donkere materie?”* Geen enkel deeltje in het Standaardmodel past bij de beschrijving van de astronomen: zwaar, zonder interactie met licht, en met hooguit een zeer zwakke interactie met andere deeltjes. We moeten dus een deeltje aan het befaamde Standaardmodel toevoegen. Het Standaardmodel is misschien heel goed in het beschrijven van alle normale deeltjes, maar blijkbaar is dit slechts 15 % van alle deeltjes in het universum.

Gelukkig waren er al een paar andere problemen met het Standaardmodel, en om deze op te lossen wilden deeltjesfysici toch al nieuwe deeltjes toevoegen. Eén type deeltjes wat de donkere materie zou kunnen zijn is de WIMP: een Weakly Interacting Massive Particle. WIMP’s hebben precies de goede eigenschappen om zowel problemen in de deeltjesfysica op te lossen, als de kosmologische waarnemingen te verklaren. Dit deeltje is ongeveer 100 keer zo zwaar als een proton, en heeft weliswaar een interactie met normale deeltjes, maar alleen erg zwak.

De enige manier om erachter te komen of WIMP’s bestaan is proberen om ze te meten. Eén van de methodes om dat te doen is om te zoeken naar een botsing tussen een WIMP en een atoom hier op aarde. We hoeven geen moeite te doen om WIMP’s naar de aarde te brengen: als ze bestaan, dan moeten ze overal zijn en vliegen ze door de Melkweg heen met een enorme snelheid. Aangezien ze zo’n zwakke interactie met normale materie hebben gaan ze door praktisch alles heen: voor WIMP’s lijkt het net alsof de aarde doorzichtig is. Als ze wél een keer botsen met een atoom, wat erg onwaarschijnlijk maar niet onmogelijk is, dan geven ze een klein beetje energie aan dat atoom. In een donkere materiedetector kan deze botsing gedetecteerd worden. Veel energie is het niet: WIMP’s zijn weliswaar snel, maar ook erg klein. We hebben dus een gevoelige detector nodig om WIMP’s te kunnen meten.

ZOEKEN NAAR ENKELE DEELTJES

De meest gevoelige WIMP-detector ter wereld is het XENON1T-experiment (zie figuur 1). XENON1T bevat ongeveer 2 ton aan xenonatomen: zo veel mogelijk, om de kans dat WIMP's botsen met het xenon zo groot mogelijk te maken. Vóór XENON1T waren er al eerdere, kleinere experimenten, die allemaal niet in staat waren om donkere materie te vinden. Daarom zijn de xenonexperimenten in de loop van de tijd steeds groter geworden, van een paar kilo tot meer dan een ton in iets meer dan tien jaar tijd.



Figuur 1: Het XENON1T-experiment. Aan de linkerkant is een watertank te zien, waarin het hart van XENON1T zich bevindt: een vat gevuld met vloeibaar xenon. Aan de rechterkant staat een gebouw met instrumenten die nodig zijn het experiment te laten werken: een zuiveringsinstallatie, data-apparatuur en een tank voor xenon-opslag.

Uniek aan xenon is het ontstaan van meetbare signalen als gevolg van de energie die WIMP's overdragen. Helaas zijn WIMP's niet de enige deeltjes die energie kunnen overdragen aan xenon: deeltjes die het gevolg zijn van radioactief verval kunnen dit ook doen, met als gevolg een vergelijkbaar signaal als dat van WIMP's. Straling is overal aanwezig, vooral doordat veel doodgewone materialen een klein beetje radioactieve stoffen bevatten. Verder komt er ook kosmische straling uit de ruimte, die zo doordringend is dat die door dikke lagen afscherming heen kan komen. De grootste uitdaging van donkere materiedetectoren is niet om donkere

materie te detecteren, maar om niet continu achtergrondstraling te detecteren, en om het te kunnen onderscheiden van het signaal van donkere materie.

In het XENON1T-experiment worden er verschillende technieken gebruikt om de achtergrond te minimaliseren. Ten eerste staat het experiment in een diep ondergronds lab in Gran Sasso (Italië), waar 1,4 km aan gesteente een natuurlijke afscherming biedt tegen de doordringende kosmische straling. Verder bevindt het xenon zich in het middelpunt van een 10 meter grote watertank. Deze tank biedt afscherming tegen straling, maar is ook zelf een detector, zodat we weten wanneer er kosmische straling binnenkomt. Als er tegelijkertijd een interactie wordt gemeten in het water en in het xenon, dan weten we dat deze signalen worden veroorzaakt door kosmische straling, en niet door donkere materie. Tenslotte worden alle materialen die in of om het experiment heen komen uiterst zorgvuldig geselecteerd, zodat ze zo min mogelijk radioactieve verontreinigingen bevatten.

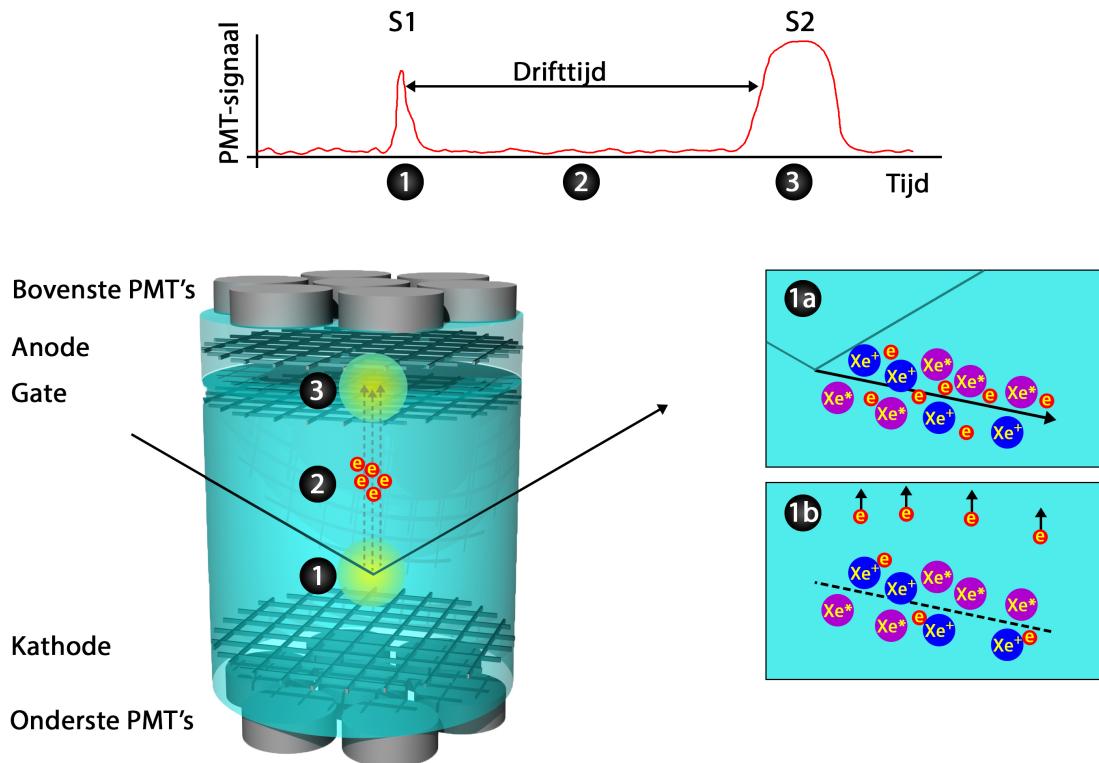
DE WERKING VAN TPC'S

Zelfs na alle moeite om achtergrond te minimaliseren komt er nog steeds straling de detector binnen. In het geval van XENON1T zorgt dit voor ongeveer 5 gemeenten interacties per seconde, terwijl XENON1T is ontworpen om een paar WIMP-botsingen per *jaar* te meten. Om de aanwezigheid van WIMP's aan te kunnen tonen, moeten we de achtergrondsignalen, die ondanks de maatregelen passeren, onderscheiden van de signalen van WIMP's. Dit is een enorme opgave: in totaal worden er meer dan 100 miljoen achtergrondinteracties per jaar gemeten, en we moeten dit terug zien te dringen tot bijna niets.

Om uit te leggen hoe dit in zijn werk gaat moeten we eerst wat dieper ingaan op hoe XENON1T werkt. In figuur 2 is een schematische weergave van de XENON1T-detector te zien. Deze bestaat uit een cilinder die deels gevuld is met vloeibaar xenon, gekoeld tot -90°C . Het overgebleven gedeelte aan de bovenkant van de cilinder bevat gasvormig xenon. Aan de boven- en onderkant is de cilinder bedekt met gevoelige lichtdetectoren: PMT's (photomultiplier tubes, ofwel fotoversterkingsbuisen).

Als een deeltje botst in het vloeibaar xenon, dan geeft het zijn energie af aan de kern van een xenonatoom, of aan een elektron om de atoomkern heen. Door de botsing kaatst de atoomkern of het elektron weg met hoge snelheid. Daarbij maakt het vrije elektronen (e), xenon-ionen (Xe^+) en aangeslagen xenonatomen (Xe^*) (zie 1a in figuur 2). De aangeslagen xenonatomen vallen kort daarna weer terug naar normale xenonatomen, en zenden daarbij lichtdeeltjes (fotonen) uit. Het gevolg

is een korte lichtflits, die door de PMT's gedetecteerd kan worden. Het gemeten signaal noemen we 'S1', signaal 1.



Figuur 2: Schematische weergave van de werking van het XENON1T-experiment.

De elektronen die in de botsing vrijgemaakt worden zorgen voor een tweede signaal. Met behulp van twee elektrodes aan de onder- en bovenkant van de vloeibaar xenon (kathode en gate) wordt een elektrisch veld opgewekt, zodat de elektronen omhoog getrokken worden (1b en 2). Wanneer de elektronen aankomen bij de gaslaag, worden ze door het nog sterkere veld van een derde elektrode in het gas (de anode) aangetrokken. Er wordt zo hard aan de elektronen getrokken dat ze uit de vloeistof komen en veel energie krijgen terwijl ze in het gas zijn. Hierdoor ontstaat opnieuw een lichtflits, die weer door de PMT's aan de bovenkant en onderkant wordt gemeten (3). Dit signaal heet S2.

Tussen de S1 en de S2 bestaat een tijdverschil, omdat het even duurt voordat de elektronen de bovenkant bereiken. De tijd die de elektronen onderweg zijn noemen we de drifttijd. Hoe dieper in de detector de interactie was, hoe langer de drifttijd. Hierdoor kunnen we met behulp van de tijd van de twee signalen de diepte van de interactie bepalen. Detectoren die op deze manier de diepte achterhalen noemen we TPC's, of tijdprojectiekamers (time projection chambers). De drifttijd alleen geeft

een ééndimensionale positie, maar we kunnen ook de andere twee coördinaten achterhalen. De PMT's die dichtbij het punt zijn waar het licht van de S_2 wordt uitgezonden zullen namelijk meer van het licht zien dan de andere PMT's. De combinatie van deze informatie betekent dat we precies weten waar in het volume van de detector de interactie plaats heeft gevonden.

ACHTERGROND ONDERDRUKKING

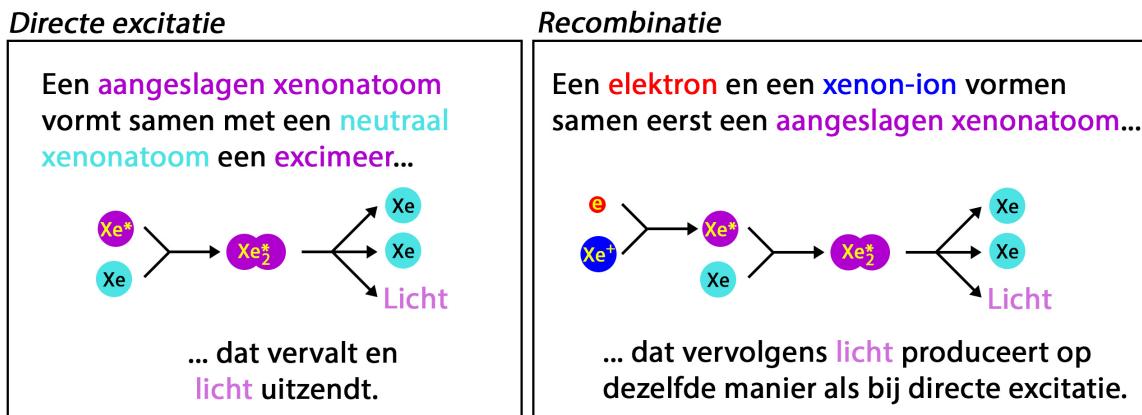
Met behulp van de positie van de interactie kan een groot deel van de achtergrondinteracties worden onderscheiden van WIMP-signalen. De meeste gemeten achtergrondinteracties zijn afkomstig van radioactieve stoffen in de materialen van de detector. De straling komt meestal niet ver het vloeibaar xenon in. We kunnen alle interacties aan de randen daarom weggooien.

De tweede methode om achtergrondstraling te onderdrukken is door de verschillende deeltjes die wegkaatsen te onderscheiden. Voor WIMP's verwachten we alleen botsingen met de atoomkern, en niet met de elektronen, terwijl de meeste achtergronden juist komen door botsing van straling op elektronen. Een botsing met een atoomkern noemen we een atoombotsing en een botsing met een elektron een elektronenbotsing. Voor allebei de interactietypes krijgen we zowel een S_1 als een S_2 , maar niet in dezelfde *verhouding*. Een atoombotsing maakt meer aangeslagen xenonatomen en minder elektronen dan een elektronenbotsing, zodat er verhoudingsgewijs een grotere S_1 en een kleinere S_2 ontstaat. We kunnen dus het onderscheid tussen een atoombotsing en een elektronenbotsing maken door de verhouding S_2/S_1 te bekijken. Deze methode is zo goed, dat we er 99,7 % van de achtergrond mee weg kunnen halen.

ACHTERGROND ONDERSCHEIDEN MET DE PULSVORM

Hoewel het onderscheid tussen atoombotsingen en elektronenbotsingen met de S_2/S_1 -verhouding erg goed is, is het nog steeds niet perfect. Daarom is men altijd op zoek naar manieren om de achtergrondsignalen nog beter van WIMP-signalen te onderscheiden. Eén van de manieren waarop dit mogelijk kan, is door heel precies naar de vorm van het S_1 -signaal te kijken. Deze is voor WIMP's en achtergrond namelijk verschillend.

Het licht van het S_1 -signaal wordt niet precies in één keer geproduceerd, maar er is een korte vertraging. De lengte van de vertraging is afhankelijk van hoe het licht precies geproduceerd wordt. Hiervoor zijn in totaal vier mogelijkheden, die

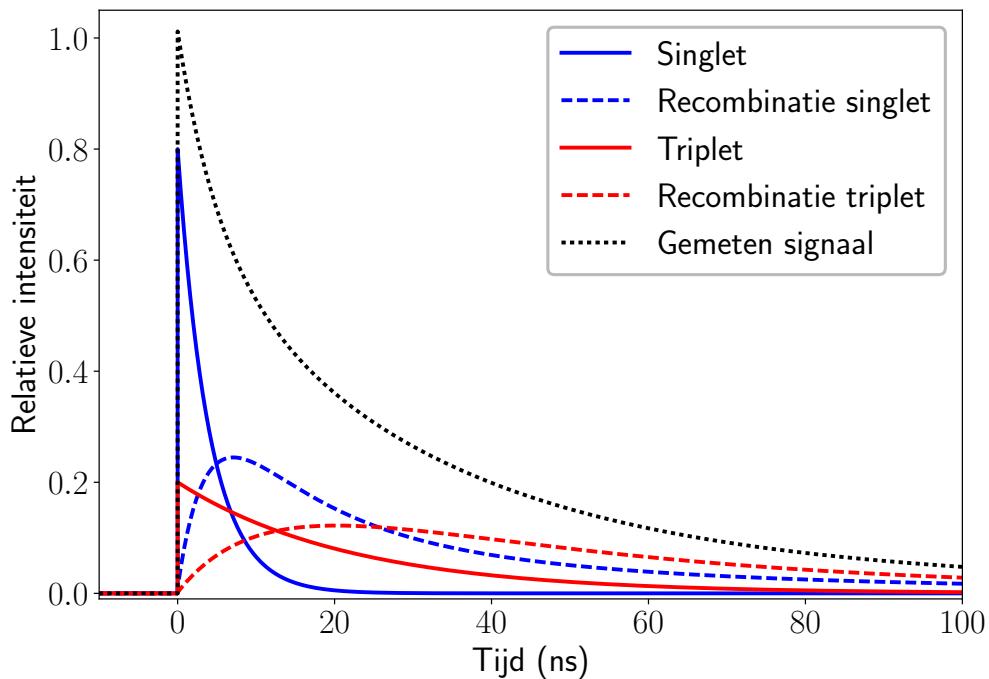


Figuur 3: De twee manieren waarop licht wordt geproduceerd bij een interactie in het xenon. Bij directe excitatie maken de aangeslagen xenonatomen een instabiel molecuul (excimeer), dat kort daarna licht uitzendt. Bij recombinatie werkt dit op dezelfde manier, maar wordt het voorafgegaan door het samenkommen van elektronen en xenon-ionen, met als gevolg een langere vertraging.

elk hun eigen pulsvorm hebben. Deze pulsvormen zijn aangegeven in figuur 4. Het uiteindelijk gemeten S1-signaal is een combinatie van deze vier pulsvormen.

In de botsing worden aangeslagen xenonatomen gemaakt, die voor ze licht uitzenden instabiele moleculen, genaamd *excimeren*, produceren (zie figuur 3). De productie van de excimeren gaat extreem snel, maar het verval niet: hier is een korte vertraging. Er bestaan twee soorten excimeer-toestanden, de singlet en de triplet, die een gemiddelde vervaltijd van respectievelijk 3 ns en 22 ns (nanoseconden, een miljardste van een seconde) hebben. Hierdoor zijn er twee mogelijke pulsvormen, die allebei het gevolg zijn van de aangeslagen xenonatomen die direct door de botsing worden gemaakt. Dit proces wordt ook wel directe excitatie genoemd.

Er kan ook nog een extra vertraging optreden door aangeslagen xenonatomen die niet direct door de botsing worden geproduceerd. Dit wordt veroorzaakt door elektronen die niet omhoog worden getrokken door het elektrische veld, maar achterblijven bij de xenon-ionen (zie 1b in figuur 2). Als een elektron en een xenon-ion weer bij elkaar komen, wordt er een aangeslagen xenonatoom gemaakt. Dit proces heet recombinatie (zie figuur 3). Het aangeslagen xenonatoom zal, net als in het geval van directe excitatie, een excimeer maken in de singlet- of triplet-toestand. Er is hier sprake van twee vertragende effecten voordat het licht wordt uitgezonden: eerst de productie, en daarna het verval van de excimeren. Ook hier zijn weer twee mogelijke pulsvormen als gevolg van de verschillende vervaltijden van de twee toestanden. De pulsvormen door recombinatie zijn weergegeven in figuur 4.

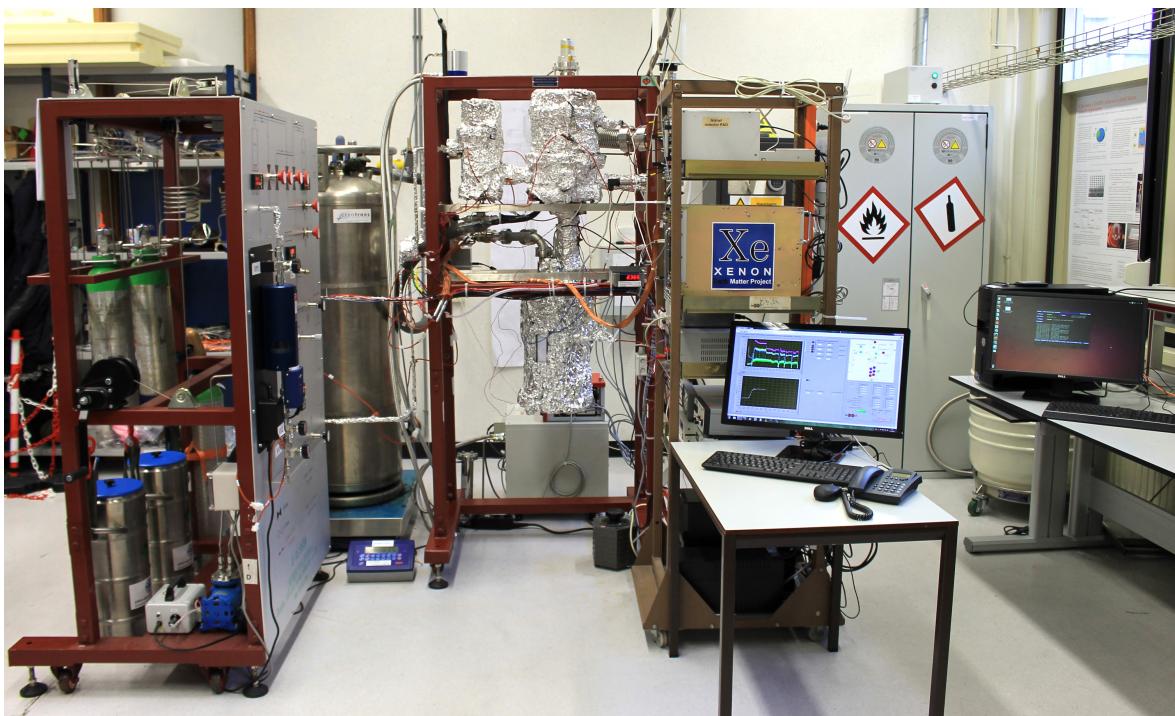


Figuur 4: De gemeten S_1 -pulsvorm is opgebouwd uit vier componenten: twee door directe excitatie en twee door recombinatie.

EEN KLEINE DETECTOR

In Amsterdam hebben we een opstelling die uiterst geschikt is voor metingen van de pulsvorm. Deze opstelling heet XAMS (zie figuur 5). XAMS werkt op precies dezelfde manier als XENON1T: er is een koelinstallatie, een gaszuiveringssysteem, elektronica en een dataverwerkingssysteem. Al deze systemen lijken op die van XENON1T. Naast het formaat (0,5 kg in plaats van 2 ton) is het enige wat aanzienlijk verschilt de achtergrondafscherming, die we niet nodig hebben omdat XAMS niet is gebouwd om donkere materie te zoeken. In plaats daarvan gebruiken we radioactieve bronnen en onderzoeken we de eigenschappen van de interacties van de straling.

Er is één eigenschap waar XAMS echt in uitblinkt: de tijdresolutie. De PMT-signalen worden normaal gesproken in tijdstappen van 10 ns gemeten, maar XAMS is met 2 ns vijf keer zo snel. Bij deze extreem korte tijden is zelfs de snelheid van het licht een factor om rekening mee te houden. Het licht doet er bijvoorbeeld ongeveer 6 ns over om van de ene kant van de XENON1T-detector naar de andere kant te komen. Dit is ook één van de redenen dat een kleine detector in dit geval beter is: het duurt minder van 1 ns om door de XAMS-detector heen te komen. Deze en



Figuur 5: De XAMS-opstelling: een kleine xenon TPC in Amsterdam, gemaakt voor onderzoek dat detectoren als XENON1T helpt in hun zoektocht naar donkere materie.

andere eigenschappen maken XAMS erg geschikt voor het meten van de pulsvorm van de S₁.

METINGEN VAN DE PULSVORM

Tot nu toe waren de pulsvormen van atoombotsingen en elektronenbotsingen nog niet goed genoeg gemeten. Daarom was niet bekend of de pulsvorm veel kan helpen bij de achtergrondonderdrukking in donkere materiedetectoren. In hoofdstuk 5 staan gedetailleerde metingen van de pulsvorm beschreven, die zijn gedaan met behulp van XAMS. De analyse daarin laat precies zien hoe de pulsvorm eruit ziet, afhankelijk van de energie, het elektrisch veld en het type botsing (elektronenbotsing of atoombotsing). Helaas is de conclusie dat het onderscheid niet gemakkelijk te maken is: de pulsvormen van elektronen- en atoombotsingen lijken simpelweg te veel op elkaar. Het is daarom waarschijnlijk niet de moeite waard om veel te investeren in het beter meten van de pulsvorm in donkere materiedetectoren. Behalve deze deprimerende conclusie kunnen we de metingen in dit hoofdstuk ook gebruiken om conclusies te trekken en modellen te maken over de processen die bij elkaar zorgen voor het uitzenden van het licht van de S₁. We concluderen bijvoorbeeld dat

het vertragende effect van recombinatie afhankelijk is van het type botsing en het veld, wat ons informatie geeft over wat de elektronen aan het doen zijn in de eerst paar nanoseconden nadat ze zijn bevrijd van hun atomen. Het verband tussen de pulsvorm en het elektrisch veld wordt verder beschreven in hoofdstuk 6.

STEEDS BETERE DETECTOREN

Xenon TPC's zijn ontzettend goede detectoren voor de zoektocht naar donkere materie. Dit is deels dankzij de detectoren die alsmaar groter worden. XENON1T bevat al ruim 2 ton aan vloeibaar xenon, maar de volgende fase xenonexperimenten zal een massa van 6 tot 7 ton aan vloeibaar xenon bevatten, en de geplande opvolger daarvan maar liefst 40 ton. Alleen steeds grotere detectoren bouwen om de gevoeligheid te verhogen is niet voldoende. We zullen ook steeds beter moeten begrijpen hoe deze detectoren werken. Een belangrijke reden hiervoor is dat de grotere detectoren ook een lagere achtergrond moeten hebben (per kilo xenon). XAMS draagt daarin bij met metingen zoals de eerder beschreven pulsvorm-metingen.

In de meting van meer dan een jaar van XENON1T is er, helaas, geen aanwijzing gevonden voor WIMP-interacties. Dit resultaat betekent dat we de duimschroeven van theoretische modellen van donkere materie kunnen aandraaien: we weten nog steeds niet welk deeltje donkere materie is, maar we weten wel dat de interactie zo zwak is dat XENON1T het niet kan detecteren. Om donkere materie wél te detecteren moeten we blijkbaar nog meer ons best doen: grotere detectoren bouwen, de achtergrond minimaliseren, en nog beter begrijpen hoe xenon werkt als detectormateriaal. Xenondetectoren worden momenteel beter op al deze fronten. De volgende experimenten, XENONnT en LZ, zullen 10 keer gevoeliger zijn en starten hun zoektocht mogelijk al in 2019. Misschien is de detectie van donkere materie dus al dichtbij. Anderzijds is het mogelijk dat de WIMP-theorie helemaal niet klopt, of dat WIMP's wel bestaan maar zo'n zwakke interactie met normale materie hebben dat we ze zelfs niet met de allergrootste detectoren kunnen waarnemen. Als we ooit de kans willen krijgen om te achterhalen wat donkere materie is, dan zullen we het gewoon moeten proberen en hopen dat we geluk hebben. Het is de moeite waard. Ik ben benieuwd waar de andere 85 % van het heelal uit bestaat — jij toch ook?

ACKNOWLEDGMENTS

The past four years (or even a bit more, counting my Master's project) have been a great thrill, and there are some people I would like to thank for everything. First of all, thanks to Auke-Pieter and Frank for being my promotors. Auke-Pieter, I have had a lot of fun discussing kitesurfing (and sometimes physics) during our many spontaneous coffee breaks. I think that your relentless optimism and my realistic pessimism was the perfect combination to improve XAMS and its measurements. Frank, your keen eye for detail has improved many of the plots and sentences in this thesis. I would also like to Patrick and Auke-Pieter for creating the positive environment that makes the Nikhef dark matter group thrive. It's great to feel valued, either by simple compliments or with actual champagne (although I think you still owe me a bottle for the field paper ;)). I consider myself very lucky that I have ended up under your supervision: you are the best.

I am grateful to all my committee members: Patrick, Paul, Eric, Shin'ichiro, Marcel, Rafael and Antonio. Thank you for taking the time to read my thesis and prepare (hopefully not too difficult) questions. A special thanks to Marcel for being member in my C3-committee and always being there, just in case.

The PhD would not have been so much fun without the kind fellow PhD students in the dark matter group: Sander, Jelle, Peter, Andrea and Joran. The usual 'wanna see a plot?' was always a source of fun, fierce and fruitful discussion about physics, aesthetics and details. Sander, it was great to share so many details and frustrations that come with finishing a PhD together. I'd gladly start the PaleoSolutions company with you if it weren't for your impending academic career, and I wish you all the best in the US. Jelle, the combination of your knowledge of details of xenon physics with crazy ideas is amazing, and it surprises me a little that we did not destroy XAMS with one of these ideas. In fact, your rapid-fire coding is the thing that made the XAMS DAQ work at all. Thanks also to the postdocs Chris and Andrew for helping out whenever needed. I loved our lunchtime discussions and your strong opinions about physics organization really made me think. I forgive both of you for killing me in Diplomacy :).

A great thanks to the Master students that I spent many lab-hours with: Maria, Kiefer, Kate and Alvaro. Maria, I am glad you chose to do a PhD after your Master, and the xenon group should envy the gravitational waves group for having you.

Acknowledgments

Alvaro, I trust you will take good care of XAMS: you're the perfect one for it, and Patrick and Auke-Pieter would be fools not to hire you.

Thanks to Rolf for helping me throughout my Master project. Despite all your good advice, you inspired me to start the PhD project, and I have not regretted it. I've had so much fun working on XAMS with you, and I am glad to know you continue to inspire the next generation of students.

There have been many people in the dark matter group who helped out with on-call shifts. Thanks to Matteo, Rolf, Faan, Chris, Sander, Jelle, Frank, Andrew, Froukje, Auke-Pieter, Patrick, Maria, Arjen, Kiefer, Mitch and Jasper for their time. I am glad that we never had to rely on our emergency procedures.

I enjoyed having the opportunity to cooperate with people all over the globe for the neutron generator paper. Waiting for a thousand neutron counts for several hours wasn't so bad having Jacques, Gianmarco and Lotfi around. The supervision and help we received at PTB and during the following analysis was indispensable and helped make the paper a success. I greatly value the close collaboration with Rafael, Jacques and Darryl on the analysis and the later installation of the stubborn neutron generator.

I am thankful that I have had the experience of being in a world-class scientific collaboration and to work on the amazing XENON1T detector. I would like to thank everyone in the XENON collaboration for the good times and wish you all the best with building XENONnT. I sincerely hope that XENON1T will turn out to be the detector that *almost* found dark matter, and not like one of the 'aether-hunting' machines.

Building and maintaining an experiment at Nikhef is possible only because of the help of all the excellent technical staff. I would like to thank Berend, Nico, Wim Gotink, Dimitri, Rob Walet and Jan Koopstra for their help, and all the others in the machine shop, electronics and computing department.

A big thank you to all my family and friends who managed to keep me in the 'real' world, and even showed interest in dark matter and fundamental physics. Thanks to the help of Danielle, Willem and Johan, the Dutch summary should be readable to almost anyone (and that really says something). Danielle and Maria, thanks for being my paranympths. Danielle, there might soon be a little less physics in your life, but you won't get rid of me any time soon.

Thanks to everyone – including those who I forgot to mention – for making these years insightful, challenging and so much fun. I am grateful to have had the privilege of working on fundamental physics. It's been a great ride. Thanks to all of you.

A

APPENDIX

A.1 RESIDUAL GAS ANALYSIS

Table A.1: Overview of the main peaks found in the RGA scan shown in figure 4.14, together with their most probable ion fragment and its source. Peak identification was performed using the library in the Balzers Quadstar Utility program.

<i>m/e</i>	Fragment	Source
1	H^+	Water
2	H_2^+	Water
7	N^{2+}	Air
8	O^{2+}	Air
12	C^+	Hydrocarbon, CO_2 (air)
14	N^+	Air
15	$\text{CH}_3^+, {}^{15}\text{N}^+$	Hydrocarbon, air
16	O_2^+	Air
17	OH^+	Water
18	H_2O^+	Water
19	$\text{H}_2{}^{17}\text{O}^+$	Water
20	$\text{H}_2{}^{18}\text{O}^+, \text{Ar}^{2+}$	Water, air
26	C_2H_2^+	Hydrocarbon/ethanol
28	N_2^+	Air
29	C_2H_4^+	Hydrocarbon/ethanol
30	$\text{C}_2\text{H}_5^+, \text{NO}^+$	Hydrocarbon/ethanol, air
31	CH_3O^+	Ethanol
32	O_2^+	Air
36-45	C_3H_x^+	Hydrocarbon
40	Ar^+	Air
44	CO_2^+ or N_2O^+	Air
45	$\text{C}_2\text{H}_5\text{O}^+$	Ethanol

A.2 PIPING AND INSTRUMENTATION DIAGRAM

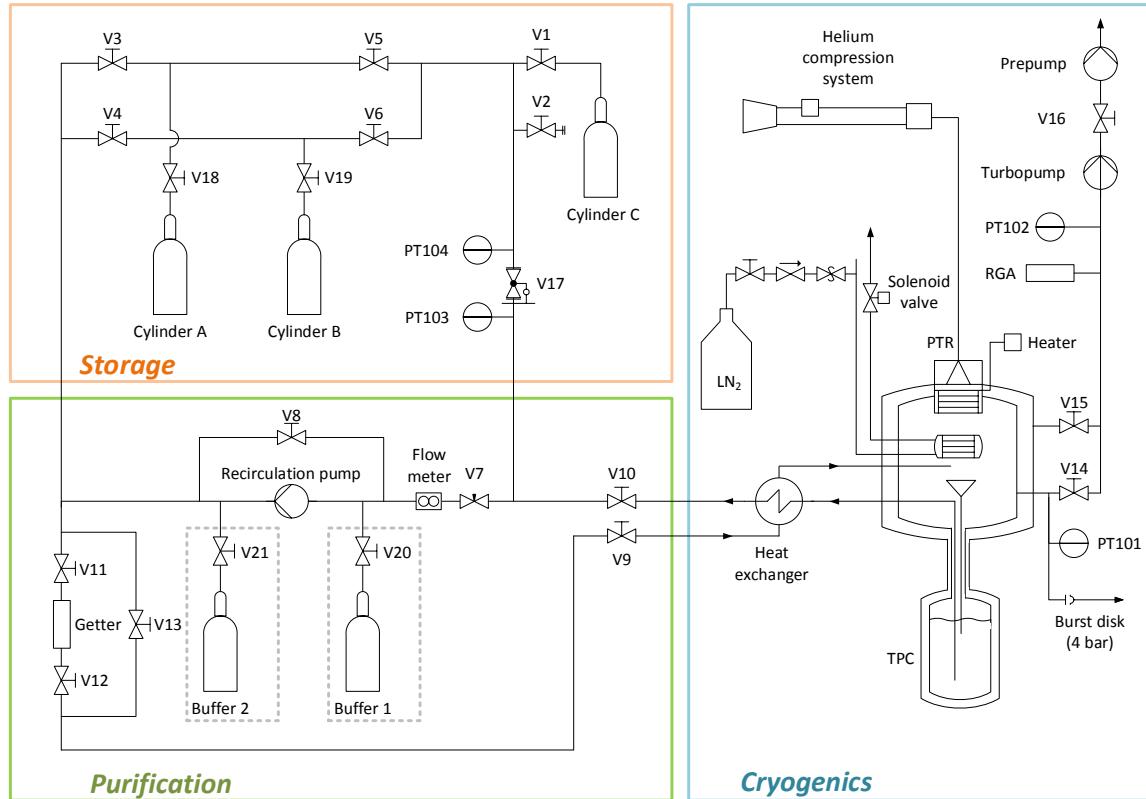


Figure A.1: Piping and instrumentation diagram of the setup. In the box on the right, the cryogenic part of the setup is shown. The main cooling is provided by the pulse tube refrigerator (PTR). The liquefied xenon drops down into a funnel, which leads into the detector volume. The temperature at the cold finger is controlled by adjusting the current to a resistive heating band. In case of a cooling failure, automatic emergency cooling is provided by a pressurized liquid nitrogen system. The bottom left box shows the purification system. Liquid xenon is extracted from the TPC and evaporates in the heat exchanger, so that gaseous xenon is be pumped through a getter. The flow speed is regulated with a needle valve (V7) and measured with a mass flow meter. The buffer volumes, shown in the dashed gray boxes, were added in the system after the measurements. When the detector is not running, we store the xenon as a pressurized gas in the storage system, shown in the top left. A pressure regulator (V17) serves to set a low pressure in the detector volume. Cylinders A and B can be submerged in liquid nitrogen dewars, causing xenon deposition on in the cylinders. This is used to recuperate the xenon from the setup.

A.3 RECOMBINATION MODEL

In this model, we consider a region with volume V where there is a uniform density of free electrons n_e and ions n_{ions} caused by the initial ionizing particle. Both of these densities are time-dependent, as electrons and ions recombine to form excited states. We describe the number density of the recombination electrons n_e by the differential equation

$$\frac{dn_e}{dt} = \frac{dn_{ions}}{dt} = -\alpha n_e(t) n_{ions}(t). \quad (\text{A.1})$$

If the free electrons and ions are completely confined to the volume V , the initial values of n_e and n_{ions} as well as their derivatives are equal. We can then define $n(t) \equiv n_e(t) = n_{ions}(t)$ and $n(0) = n_0$, giving the solution

$$n(t) = \frac{n_0}{1 + \frac{t}{T_R}}, \quad (\text{A.2})$$

where T_R is the recombination time constant given by $(n_0\alpha)^{-1}$. This case corresponds to full recombination, as for $t \rightarrow \infty$ all electrons and ions have recombined ($n \rightarrow 0$). However, since the electrons are initially energetic with respect to the energy associated with the Coulomb binding energy, there is a finite probability η that these electrons escape the region V where recombination occurs shortly after the primary interaction. In this case, the boundary conditions change to $n_{ions}(0) = n_0$ and $n_e(0) = (1 - \eta)n_0$, and the equation has to be solved numerically. If we assume that the recombination electrons and ions form an excited molecular state in either the triplet or singlet state, the corresponding number of excited states N with lifetime τ is given by

$$\begin{aligned} \frac{dN}{dt} &= -\frac{N(t)}{\tau} + -V \frac{dn_e}{dt} \\ &= -\frac{N(t)}{\tau} + V\alpha n_e(t) n_{ions}(t), \end{aligned} \quad (\text{A.3})$$

with the two terms on the right hand side corresponding to the decay and the production of the excited states. This equation has the solution:

$$N(t) = V\alpha \exp\left(\frac{-t}{\tau}\right) \int_0^t n_e(t') n_{ions}(t') \exp\left(\frac{t'}{\tau}\right) dt'. \quad (\text{A.4})$$

The luminescence time dependence $I(t)$ is equal to $\frac{1}{\tau}N(t)$, which gives:

$$I(t) = \frac{V\alpha}{\tau} \exp\left(\frac{-t}{\tau}\right) \times \int_0^t n_e(t') n_{ions}(t') \exp\left(\frac{t'}{\tau}\right) dt'. \quad (\text{A.5})$$

We define the normalized distribution I_r as

$$I_r(t, \tau, T_R, \eta) = A \exp\left(\frac{-t}{\tau}\right) \times \int_0^t n_e(t') n_{ions}(t') \exp\left(\frac{t'}{\tau}\right) dt', \quad (\text{A.6})$$

with A a normalization factor.

A.4 ABBREVIATIONS

CDM	(Λ_{CDM}) Cold dark matter
CMB	Cosmic microwave background
D-D	Deuterium-deuterium
ER	Electronic recoil
LDE	Light detection efficiency
NEST	Noble element simulation technique
NG	Neutron generator
NR	Nuclear recoil
PMT	Photomultiplier tube
PTFE	Polytetrafluoroethylene (Teflon)
PTR	Pulse tube refrigerator
RGA	Residual gas analyzer
TPC	Time projection chamber
WIMP	Weakly interacting massive particle
ZLE	Zero-length encoding

BIBLIOGRAPHY

- [1] E. Hogenbirk et al. „Commissioning of a dual-phase xenon TPC at Nikhef.” In: *Nuclear Instruments and Methods in Physics Research Section A* 840 (2016), pp. 87–96. doi: [10.1016/j.nima.2016.09.052](https://doi.org/10.1016/j.nima.2016.09.052).
- [2] E. Hogenbirk et al. „Precision measurements of the scintillation pulse shape for low-energy recoils in liquid xenon.” In: *Journal of Instrumentation* 13.05 (2018), Po5016. doi: [10.1088/1748-0221/13/05/P05016](https://doi.org/10.1088/1748-0221/13/05/P05016).
- [3] E. Hogenbirk et al. „Field dependence of electronic recoil signals in a dual-phase liquid xenon time projection chamber.” In: *Journal of Instrumentation* 13.10 (2018), P10031. doi: [10.1088/1748-0221/13/10/P10031](https://doi.org/10.1088/1748-0221/13/10/P10031).
- [4] R. F. Lang et al. „Characterization of a deuterium-deuterium plasma fusion neutron generator.” In: *Nuclear Instruments and Methods in Physics Research Section A* 879 (2018), pp. 31–38. ISSN: 0168-9002. doi: [10.1016/j.nima.2017.10.001](https://doi.org/10.1016/j.nima.2017.10.001).
- [5] E. Aprile et al., the XENON collaboration. „Dark Matter Search Results from a One Ton-Year Exposure of XENON1T.” In: *Physical Review Letters* 121 (11 2018), p. 111302. doi: [10.1103/PhysRevLett.121.111302](https://doi.org/10.1103/PhysRevLett.121.111302).
- [6] P. A. R. Ade et al., the Planck collaboration. „Planck 2015 results - XIII. Cosmological parameters.” In: *Astronomy & Astrophysics* 594 (2016), A13. doi: [10.1051/0004-6361/201525830](https://doi.org/10.1051/0004-6361/201525830).
- [7] C. Skordis. „The tensor-vector-scalar theory and its cosmology.” In: *Classical and Quantum Gravity* 26.14 (2009), p. 143001. doi: [10.1088/0264-9381/26/14/143001](https://doi.org/10.1088/0264-9381/26/14/143001).
- [8] X. Xu, B. Wang, and P. Zhang. „Testing the tensor-vector-scalar theory with the latest cosmological observations.” In: *Physical Review D* 92 (8 2015), p. 083505. doi: [10.1103/PhysRevD.92.083505](https://doi.org/10.1103/PhysRevD.92.083505).
- [9] *Planck Legacy Archive*. Version 2.14, April 2018. URL: <https://pla.esac.esa.int/pla/>.
- [10] E. Hogenbirk. *DMPlots v1.1.0*. Nov. 2018. doi: [10.5281/zenodo.1479669](https://doi.org/10.5281/zenodo.1479669).
- [11] W. Hu and S. Dodelson. „Cosmic Microwave Background Anisotropies.” In: *Annual Review of Astronomy and Astrophysics* 40.1 (2002), pp. 171–216. doi: [10.1146/annurev.astro.40.060401.093926](https://doi.org/10.1146/annurev.astro.40.060401.093926).
- [12] D. G. York et al., the SDSS collaboration. „The Sloan Digital Sky Survey: Technical Summary.” In: *The Astronomical Journal* 120.3 (2000), p. 1579. doi: [10.1086/301513](https://doi.org/10.1086/301513).
- [13] M. Colless et al. „The 2dF Galaxy Redshift Survey: spectra and redshifts.” In: *Monthly Notices of the Royal Astronomical Society* 328.4 (2001), pp. 1039–1063. doi: [10.1046/j.1365-8711.2001.04902.x](https://doi.org/10.1046/j.1365-8711.2001.04902.x).
- [14] M. J. Geller and J. P. Huchra. „Mapping the Universe.” In: *Science* 246.4932 (1989). doi: [10.1126/science.246.4932.897](https://doi.org/10.1126/science.246.4932.897).
- [15] V. Springel et al. „Simulations of the formation, evolution and clustering of galaxies and quasars.” In: *Nature* 435 (2005), pp. 629–636. doi: [10.1038/nature03597](https://doi.org/10.1038/nature03597).

- [16] V. Springel, C. S. Frenk, and S. D. M. White. „The large-scale structure of the Universe.” In: *Nature* 440 (2006), pp. 1137–1144. doi: 10.1038/nature04805.
- [17] P. McDonald et al. „The Ly α Forest Power Spectrum from the Sloan Digital Sky Survey.” In: *The Astrophysical Journal Supplement Series* 163.1 (2006), p. 80. doi: 10.1086/444361.
- [18] J. L. Tinker et al. „Cosmological Constraints from Galaxy Clustering and the Mass-to-number Ratio of Galaxy Clusters.” In: *The Astrophysical Journal* 745.1 (2012), p. 16. doi: 10.1088/0004-637X/745/1/16.
- [19] S. Das et al. „The Atacama Cosmology Telescope: A Measurement of the Cosmic Microwave Background Power Spectrum at 148 and 218 GHz from the 2008 Southern Survey.” In: *The Astrophysical Journal* 729.1 (2011), p. 62. doi: 10.1088/0004-637X/729/1/62.
- [20] N. Sehgal et al. „The Atacama Cosmology Telescope: Cosmology from Galaxy Clusters Detected via the Sunyaev-Zel’dovich Effect.” In: *The Astrophysical Journal* 732.1 (2011), p. 44. doi: 10.1088/0004-637X/732/1/44.
- [21] A. Vikhlinin et al. „Chandra Cluster Cosmology Project III: Cosmological Parameter Constraints.” In: *The Astrophysical Journal* 692.2 (2009), p. 1060. doi: 10.1088/0004-637X/692/2/1060.
- [22] B. A. Reid et al. „Cosmological constraints from the clustering of the Sloan Digital Sky Survey DR7 luminous red galaxies.” In: *Monthly Notices of the Royal Astronomical Society* 404.1 (2010), pp. 60–85. doi: 10.1111/j.1365-2966.2010.16276.x.
- [23] R. Hlozek et al. „The Atacama Cosmology Telescope: A Measurement of the Primordial Power Spectrum.” In: *The Astrophysical Journal* 749.1 (2012), p. 90. doi: 10.1088/0004-637X/749/1/90.
- [24] D. Larson et al. „Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Power Spectra and WMAP-derived Parameters.” In: *The Astrophysical Journal Supplement Series* 192.2 (2011), p. 16. doi: 10.1088/0067-0049/192/2/16.
- [25] J. R. Primack. *Cosmological Structure Formation*. 2015. url: <https://arxiv.org/abs/1505.02821>.
- [26] S. Dodelson. *Modern Cosmology*. Academic Press, 2003.
- [27] E. Corbelli and P. Salucci. „The extended rotation curve and the dark matter halo of M33.” In: *Monthly Notices of the Royal Astronomical Society* 311.2 (2000), pp. 441–447. doi: 10.1046/j.1365-8711.2000.03075.x.
- [28] G. Bertone, D. Hooper, and J. Silk. „Particle dark matter: evidence, candidates and constraints.” In: *Physics Reports* 405.5 (2005), pp. 279–390. ISSN: 0370-1573. doi: 10.1016/j.physrep.2004.08.031.
- [29] A. Boveia et al. „Recommendations on presenting LHC searches for missing transverse energy signals using simplified s-channel models of dark matter.” In: 2016. url: <https://arxiv.org/abs/1603.04156>.
- [30] V. Khachatryan et al., the CMS collaboration. „Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7, 8$, and 13 TeV.” In: *Journal of High Energy Physics* 2017.2 (2017), p. 135. ISSN: 1029-8479. doi: 10.1007/JHEP02(2017)135.

- [31] G. Aad et al., the ATLAS collaboration. „Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector.” In: *The European Physical Journal C* 75.7 (2015), p. 299. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3517-3.
- [32] M. Ackermann et al., the Fermi-LAT collaboration. „The Fermi Galactic Center GeV Excess and Implications for Dark Matter.” In: *The Astrophysical Journal* 840.1 (2017), p. 43. DOI: 10.3847/1538-4357/aa6cab.
- [33] M. Ajello et al., the Fermi-LAT collaboration. *Characterizing the population of pulsars in the inner galaxy with the Fermi Large Area Telescope*. URL: <https://arxiv.org/abs/1705.00009>.
- [34] R. Bartels et al. „Comment on “Characterizing the population of pulsars in the Galactic bulge with the Fermi large area telescope”.” In: *Physics of the Dark Universe* 20 (2018), pp. 88–94. ISSN: 2212-6864. DOI: 10.1016/j.dark.2018.04.004.
- [35] J. Billard, E. Figueroa-Feliciano, and L. Strigari. „Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments.” In: *Physical Review D* 89 (2 2014), p. 023524. DOI: 10.1103/PhysRevD.89.023524.
- [36] M. Dijkstra. *Structure of Dark Matter Halos*. <http://www.uio.no/studier/emner/matnat/astro/AST4320/h14/timeplan/lecture11note.pdf>, retrieved on May 14, 2018.
- [37] J. Binney and S. Tremaine. *Galactic Dynamics*. Princeton University Press, 2008.
- [38] K. Freese, J. Frieman, and A. Gould. „Signal modulation in cold-dark-matter detection.” In: *Physical Review D* 37 (12 1988), pp. 3388–3405. DOI: 10.1103/PhysRevD.37.3388.
- [39] E. Aprile et al., the XENON collaboration. „First Dark Matter Search Results from the XENON1T Experiment.” In: *Physical Review Letters* 119 (18 2017), p. 181301. DOI: 10.1103/PhysRevLett.119.181301.
- [40] J. D. Lewin and P. F. Smith. „Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil.” In: *Astroparticle Physics* 6.1 (1996), pp. 87–112. ISSN: 0927-6505. DOI: 10.1016/S0927-6505(96)00047-3.
- [41] A. P. Colijn. *Experimental Quantum Physics: Dark Matter*. 2016. URL: <https://github.com/acolijn/Teaching/tree/master/DMUU/dictaat>.
- [42] L. Baudis. „WIMP dark matter direct-detection searches in noble gases.” In: *Physics of the Dark Universe* 4 (2014). DARK TAUP2013, pp. 50–59. ISSN: 2212-6864. DOI: 10.1016/j.dark.2014.07.001.
- [43] J. Aalbers et al., the DARWIN collaboration. „DARWIN: towards the ultimate dark matter detector.” In: *Journal of Cosmology and Astroparticle Physics* 2016.11 (2016), p. 017. DOI: 10.1088/1475-7516/2016/11/017.
- [44] J. Aalbers. „Dark Matter Search with XENON1T.” PhD thesis. University of Amsterdam, 2018. URL: https://www.nikhef.nl/pub/services/biblio/theses_pdf/thesis_J_Aalbers.pdf.
- [45] E. Aprile et al., the XENON collaboration. „Search for WIMP inelastic scattering off xenon nuclei with XENON100.” In: *Physical Review D* 96 (2 2017), p. 022008. DOI: 10.1103/PhysRevD.96.022008.

- [46] E. Aprile et al., the XENON collaboration. „First axion results from the XENON100 experiment.” In: *Physical Review D* 90 (6 2014), p. 062009. DOI: 10.1103/PhysRevD.90.062009.
- [47] E. Aprile and T. Doke. „Liquid xenon detectors for particle physics and astrophysics.” In: *Reviews of Modern Physics* 82 (3 2010), pp. 2053–2097. DOI: 10.1103/RevModPhys.82.2053.
- [48] C. E. Dahl. „The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter.” PhD thesis. Princeton University, 2009. URL: <https://inspirehep.net/record/1374815/files/E.Dahlthesis.pdf>.
- [49] XCOM Photon Cross Sections Database, <https://www.nist.gov/pml/xcom-photon-cross-sections-database>, retrieved on July 5, 2018.
- [50] ENSDF database, Brookhaven National Laboratory, <https://www.nndc.bnl.gov/ensdf/>, retrieved on June 27, 2018.
- [51] A. Hitachi et al. „Effect of ionization density on the time dependence of luminescence from liquid argon and xenon.” In: *Physical Review B* 27 (9 1983), pp. 5279–5285. DOI: 10.1103/PhysRevB.27.5279.
- [52] S. Kubota, M. Hishida, and J. Raun. „Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton and xenon.” In: *Journal of Physics C: Solid State Physics* 11.12 (1978), p. 2645. DOI: 10.1088/0022-3719/11/12/024.
- [53] S. Kubota et al. „Dynamical behavior of free electrons in the recombination process in liquid argon, krypton, and xenon.” In: *Physical Review B* 20 (8 1979), pp. 3486–3496. DOI: 10.1103/PhysRevB.20.3486.
- [54] J. V. Dawson et al. „A study of the scintillation induced by alpha particles and gamma rays in liquid xenon in an electric field.” In: *Nuclear Instruments and Methods in Physics Research Section A* 545.3 (2005), pp. 690 –698. ISSN: 0168-9002. DOI: 10.1016/j.nima.2005.01.343.
- [55] T. Doke et al. „Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles.” In: *Japanese Journal of Applied Physics* 41.3R (2002), p. 1538. DOI: 10.1143/JJAP.41.1538.
- [56] T. Takahashi et al., the XMASS collaboration. „Average energy expended per ion pair in liquid xenon.” In: *Physical Review A* 12 (5 1975), pp. 1771–1775. DOI: 10.1103/PhysRevA.12.1771.
- [57] P. Sorensen and C. E. Dahl. „Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter.” In: *Physical Review D* 83 (6 2011), p. 063501. DOI: 10.1103/PhysRevD.83.063501.
- [58] J. L. Magee and J. J. Huang. „Triplet formation in ion recombination in spurs.” In: *The Journal of Physical Chemistry* 76.25 (1972), pp. 3801–3805. DOI: 10.1021/j100669a022.
- [59] A. Hitachi. „Properties of liquid xenon scintillation for dark matter searches.” In: *Astroparticle Physics* 24.3 (2005), pp. 247 –256. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2005.07.002.
- [60] G. J. Alner et al., the ZEPLIN collaboration. „First limits on nuclear recoil events from the ZEPLIN I galactic dark matter detector.” In: *Astroparticle Physics* 23.5 (2005), pp. 444 –462. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2005.02.004.

- [61] R. Bernabei et al. „The liquid xenon set-up of the DAMA experiment.” In: *Nuclear Instruments and Methods in Physics Research Section A* 482.3 (2002), pp. 728–743. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(01)01918-0.
- [62] E. Aprile et al., the XENON collaboration. „The XENON1T dark matter experiment.” In: *The European Physical Journal C* 77.12 (2017), p. 881. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5326-3.
- [63] D. S. Akerib et al., the LUX collaboration. „Results from a Search for Dark Matter in the Complete LUX Exposure.” In: *Physical Review Letters* 118 (2 2017), p. 021303. DOI: 10.1103/PhysRevLett.118.021303.
- [64] M. Szydagis et al. „Enhancement of NEST capabilities for simulating low-energy recoils in liquid xenon.” In: *Journal of Instrumentation* 8.10 (2013), p. C10003. DOI: 10.1088/1748-0221/8/10/C10003.
- [65] D. S. Akerib et al., the LUX collaboration. „Tritium calibration of the LUX dark matter experiment.” In: *Physical Review D* 93 (7 2016), p. 072009. DOI: 10.1103/PhysRevD.93.072009.
- [66] K. Ichimura. „XMASS detector calibration using a neutron source.” In: *Journal of Instrumentation* 11.02 (2016), p. C02034. DOI: 10.1088/1748-0221/11/02/C02034.
- [67] B. J. Mount et al., the LZ collaboration. *LUX-ZEPLIN (LZ) Technical Design Report*. 2017. URL: <https://arxiv.org/abs/1703.09144>.
- [68] E. Aprile et al., the XENON collaboration. „Physics reach of the XENON1T dark matter experiment.” In: *Journal of Cosmology and Astroparticle Physics* 2016.04 (2016), p. 027. DOI: 10.1088/1475-7516/2016/04/027.
- [69] S. Rosendahl et al. „A cryogenic distillation column for the XENON1T experiment.” In: *Journal of Physics: Conference Series* 564.1 (2014), p. 012006. DOI: 10.1088/1742-6596/564/1/012006.
- [70] D. S. Akerib et al., the LUX collaboration. *Improved WIMP scattering limits from the LUX experiment*. 2015. URL: <https://arxiv.org/abs/1512.03506v1>.
- [71] C. H. Faham et al. „Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes.” In: *Journal of Instrumentation* 10.09 (2015), P09010. DOI: 10.1088/1748-0221/10/09/P09010.
- [72] D. Clowe et al. „A direct empirical proof of the existence of dark matter.” In: *The Astrophysical Journal* 648 (2006), pp. L109–L113. DOI: 10.1086/508162.
- [73] P. A. R. Ade et al., the Planck collaboration. „Planck 2013 results. XVI. Cosmological parameters.” In: *Astronomy & Astrophysics* 571.A16 (2014), p. 66. DOI: 10.1051/0004-6361/201321591.
- [74] G. R. Blumenthal et al. „Formation of Galaxies and Large Scale Structure with Cold Dark Matter.” In: *Nature* 311 (1984), pp. 517–525. DOI: 10.1038/311517a0.
- [75] G. Jungman, M. Kamionkowski, and K. Griest. „Supersymmetric dark matter.” In: *Physics Reports* 267.5–6 (1996), pp. 195 –373. ISSN: 0370-1573. DOI: 10.1016/0370-1573(95)00058-5.
- [76] V. Chepel and H. Araújo. „Liquid noble gas detectors for low energy particle physics.” In: *Journal of Instrumentation* 8.04 (2013), R04001. DOI: 10.1088/1748-0221/8/04/R04001.

[77] E. Aprile et al. „Proportional light in a dual-phase xenon chamber.” In: *IEEE Transactions on Nuclear Science* 51.5 (2004), pp. 1986–1990. ISSN: 0018-9499. DOI: 10.1109/TNS.2004.832690.

[78] E. Aprile et al. „Simultaneous Measurement of Ionization and Scintillation from Nuclear Recoils in Liquid Xenon for a Dark Matter Experiment.” In: *Physical Review Letters* 97 (8 2006), p. 081302. DOI: 10.1103/PhysRevLett.97.081302.

[79] J. Kwong et al. „Scintillation pulse shape discrimination in a two-phase xenon time projection chamber.” In: *Nuclear Instruments and Methods in Physics Research Section A* 612.2 (2010), pp. 328 –333. ISSN: 0168-9002. DOI: 10.1016/j.nima.2009.10.106.

[80] E. Aprile et al., the XENON collaboration. „The XENON100 Dark Matter Experiment.” In: *Astroparticle Physics* 35 (2012), pp. 573–590. DOI: 10.1016/j.astropartphys.2012.01.003.

[81] D. S. Akerib et al., the LUX collaboration. „The Large Underground Xenon (LUX) Experiment.” In: *Nuclear Instruments and Methods A*704 (2013), pp. 111–126. DOI: 10.1016/j.nima.2012.11.135.

[82] E. Hogenbirk. „Development, commissioning and first results of XAMS: A dual-phase xenon time projection chamber at Nikhef.” MA thesis. Vrije Universiteit, 2014. URL: http://www.nikhef.nl/~ehogenbi/Erik_Hogenbirk-Thesis_LR.pdf.

[83] M. Bader. „First light from XAMS: Electronic recoil calibration of the Netherland’s first dual-phase xenon time projection chamber.” MA thesis. Utrecht University, 2015. URL: http://dspace.library.uu.nl/handle/1874/318246?_ga=1.29915965.361842327.1462893092.

[84] R. Schön. „XAMS – development of liquid xenon detector technology for dark matter searches.” PhD thesis. Universiteit van Amsterdam, 2015. URL: http://www.nikhef.nl/pub/services/biblio/theses_pdf/thesis_R_Schon.pdf.

[85] *PAX on GitHub*. retrieved on July 5, 2016. URL: <https://github.com/XENON1T/pax>.

[86] *Technical Information Manual V1724 Digitizer*. CAEN. URL: <http://www.caen.it/servlet/checkCaenManualFile?Id=12364>.

[87] M. Szydagis et al. „NEST: a comprehensive model for scintillation yield in liquid xenon.” In: *Journal of Instrumentation* 6.10 (2011), P10002. DOI: 10.1088/1748-0221/6/10/P10002.

[88] D. S. Akerib et al., the LUX collaboration. „First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility.” In: *Physical Review Letters* 112 (9 2014), p. 091303. DOI: 10.1103/PhysRevLett.112.091303.

[89] E. Aprile et al., the XENON collaboration. „Dark Matter Results from 225 Live Days of XENON100 Data.” In: *Physical Review Letters* 109 (18 2012), p. 181301. DOI: 10.1103/PhysRevLett.109.181301.

[90] E. Aprile et al. *Noble Gas Detectors*. pages 52–56. Wiley-VCH Verlag GmbH, 2006.

[91] T. Shutt et al. „Performance and fundamental processes at low energy in a two-phase liquid xenon dark matter detector.” In: *Nuclear Instruments and Methods in Physics Research Section A* 579.1 (2007), pp. 451 –453. ISSN: 0168-9002. DOI: 10.1016/j.nima.2007.04.104.

[92] E. Aprile et al. „Observation of anticorrelation between scintillation and ionization for MeV gamma rays in liquid xenon.” In: *Physical Review B* 76 (1 2007), p. 014115. DOI: 10.1103/PhysRevB.76.014115.

[93] S. Agostinelli et al., the GEANT4 collaboration. „GEANT4: A simulation toolkit.” In: *Nuclear Instruments and Methods in Physics Research Section A* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(03)01368-8.

[94] B. Edwards et al. „Measurement of single electron emission in two-phase xenon.” In: *Astroparticle Physics* 30.2 (2008), pp. 54–57. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2008.06.006.

[95] E. Santos et al., the ZEPLIN collaboration. „Single electron emission in two-phase xenon with application to the detection of coherent neutrino-nucleus scattering.” In: *Journal of High Energy Physics* 2011.12 (2011), pp. 1–20. ISSN: 1029-8479. DOI: 10.1007/JHEP12(2011)115.

[96] E. Aprile et al., the XENON collaboration. „Observation and applications of single-electron charge signals in the XENON100 experiment.” In: *Journal of Physics G: Nuclear and Particle Physics* 41.3 (2014), p. 035201. DOI: 10.1088/0954-3899/41/3/035201.

[97] A. Lyashenko. *Progress in characterization of the Photomultiplier Tubes for XENON1T Dark Matter Experiment*. 2015. URL: <https://arxiv.org/abs/1502.01000>.

[98] Hamamatsu. *Photomultiplier tubes and assemblies*. retrieved on October 30, 2018. URL: https://www.hamamatsu.com/resources/pdf/etd/High_energy_PMT TPMZ0003E.pdf.

[99] G. F. Knoll. *Radiation Detection and Measurement*. 3rd ed. John Wiley & Sons, Inc., 1999.

[100] Kodiaq on GitHub. retrieved on July 12, 2018. URL: <https://github.com/coderdj/kodiaq>.

[101] MongoDB. 2018. URL: <https://www.mongodb.com/>.

[102] P. Sorensen. „Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches.” In: *Nuclear Instruments and Methods in Physics Research Section A* 635.1 (2011), pp. 41–43. ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.01.089.

[103] Pfeiffer Vacuum. *Prisma 80 Quadrupole Mass Spectrometer manual*. retrieved on September 21, 2018. URL: <http://www.polvac.com/wp-content/uploads/Manuals/Balzers/qms-200.pdf>.

[104] F. W. McLafferty and F. Turček. *Interpretation of mass spectra*. 4th ed. University Science Books, 1993.

[105] D. Akimov et al. „Measurements of scintillation efficiency and pulse shape for low energy recoils in liquid xenon.” In: *Physics Letters B* 524.3 (2002), pp. 245–251. ISSN: 0370-2693. DOI: 10.1016/S0370-2693(01)01411-3.

[106] K. Ueshima. „Study of pulse shape discrimination and low background techniques for liquid xenon dark matter detectors.” PhD thesis. University of Tokyo, 2010. URL: http://www-sk.icrr.u-tokyo.ac.jp/xmass/_pdf/publist/ueshima_PhD.pdf.

[107] H. Takiya et al. „A measurement of the time profile of scintillation induced by low energy gamma-rays in liquid xenon with the XMASS-I detector.” In: *Nuclear Instruments and Methods in Physics Research Section A* 834 (2016), pp. 192–196. ISSN: 0168-9002. DOI: 10.1016/j.nima.2016.08.014.

[108] R. Acciarri et al. „Neutron to Gamma Pulse Shape Discrimination in Liquid Argon Detectors with High Quantum Efficiency Photomultiplier Tubes.” In: *Physics Procedia* 37 (2012). Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011), pp. 1113 –1121. ISSN: 1875-3892. DOI: 10.1016/j.phpro.2012.03.732.

[109] D. S. Akerib et al., the LUX collaboration. „Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector.” Feb. 2018. URL: <https://arxiv.org/abs/1802.06162v1>.

[110] K. Ueshima et al. „Scintillation-only based pulse shape discrimination for nuclear and electron recoils in liquid xenon.” In: *Nuclear Instruments and Methods in Physics Research Section A* 659.1 (2011), pp. 161 –168. ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.09.011.

[111] G. J. Davies et al. „Liquid xenon as a dark matter detector. Prospects for nuclear recoil discrimination by photon timing.” In: *Physics Letters B* 320.3 (1994), pp. 395 –399. ISSN: 0370-2693. DOI: 10.1016/0370-2693(94)90676-9.

[112] D. S. Akerib et al., the LUX collaboration. „Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector.” In: *Physical Review D* 97 (11 2018), p. 112002. DOI: 10.1103/PhysRevD.97.112002.

[113] D. Smith and N. Weiner. „Inelastic dark matter.” In: *Physical Review D* 64 (4 2001), p. 043502. DOI: 10.1103/PhysRevD.64.043502.

[114] X. Cui et al., the PandaX collaboration. „Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment.” In: *Physical Review Letters* 119 (18 2017), p. 181302. DOI: 10.1103/PhysRevLett.119.181302.

[115] M. Moszyński et al. „Study of n- γ discrimination with NE213 and BC501A liquid scintillators of different size.” In: *Nuclear Instruments and Methods in Physics Research Section A* 350.1 (1994), pp. 226 –234. ISSN: 0168-9002. DOI: 10.1016/0168-9002(94)91169-X.

[116] J. Calvo et al., the ArDM collaboration. *Backgrounds and pulse shape discrimination in the ArDM liquid argon TPC*. 2017. URL: <https://arxiv.org/abs/1712.01932>.

[117] P. A. Amaudruz et al., the DEAP collaboration. *Design and Construction of the DEAP-3600 Dark Matter Detector*. 2017. URL: <https://arxiv.org/abs/1712.01982>.

[118] J. Mock et al. „Modeling pulse characteristics in Xenon with NEST.” In: *Journal of Instrumentation* 9.04 (2014), T04002. DOI: 10.1088/1748-0221/9/04/T04002.

[119] J. W. Keto, R. E. Gleason, and F. K. Soley. „Exciton lifetimes in electron beam excited condensed phases of argon and xenon.” In: *Journal of Chemical Physics* 71 (1979), p. 2676. DOI: 10.1063/1.438625.

[120] S. Kubota et al. „Liquid and solid argon, krypton and xenon scintillators.” In: *Nuclear Instruments and Methods in Physics Research* 196.1 (1982), pp. 101 –105. ISSN: 0167-5087. DOI: 10.1016/0029-554X(82)90623-1.

[121] X-ray data booklet, Lawrence Berkeley National Laboratory, <http://xdb.lbl.gov/xdb.pdf>, page 37, retrieved on December 22, 2017.

[122] V. N. Solovov et al. „Measurement of the refractive index and attenuation length of liquid xenon for its scintillation light.” In: *Nuclear Instruments and Methods in Physics Research Section A* 516.2 (2004), pp. 462 –474. ISSN: 0168-9002. DOI: 10.1016/j.nima.2003.08.117.

- [123] ESTAR (stopping power and range tables for electrons), <https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>, retrieved on February 9, 2018.
- [124] SRIM (the stopping and range of ions in matter), <http://www.srim.org/>, retrieved on February 9, 2018.
- [125] B. Lenardo et al. „A Global Analysis of Light and Charge Yields in Liquid Xenon.” In: *IEEE Transactions on Nuclear Science* 62.6 (2015), pp. 3387 –3396. DOI: [10.1109/TNS.2015.2481322](https://doi.org/10.1109/TNS.2015.2481322).
- [126] L. W. Goetzke et al. „Measurement of light and charge yield of low-energy electronic recoils in liquid xenon.” In: *Physical Review D* 96 (10 2017), p. 103007. DOI: [10.1103/PhysRevD.96.103007](https://doi.org/10.1103/PhysRevD.96.103007).
- [127] Laidbox on GitHub, <https://github.com/XENON1T/laidbox>, retrieved on January 31, 2018.
- [128] Blueice on GitHub, <https://github.com/JelleAalbers/blueice>, retrieved on January 31, 2018.
- [129] P. Barrow et al. „Qualification tests of the R11410-21 photomultiplier tubes for the XENON1T detector.” In: *Journal of Instrumentation* 12.01 (2017), P01024. DOI: [10.1088/1748-0221/12/01/P01024](https://doi.org/10.1088/1748-0221/12/01/P01024).
- [130] T. Doke. „Recent developments of liquid xenon detectors.” In: *Nuclear Instruments and Methods in Physics Research* 196.1 (1982), pp. 87 –96. ISSN: 0167-5087. DOI: [10.1016/0029-554X\(82\)90621-8](https://doi.org/10.1016/0029-554X(82)90621-8).
- [131] E. Aprile, R. Mukherjee, and M. Suzuki. „Measurements of the lifetime of conduction electrons in liquid xenon.” In: *Nuclear Instruments and Methods in Physics Research Section A* 300.2 (1991), pp. 343 –350. ISSN: 0168-9002. DOI: [10.1016/0168-9002\(91\)90446-W](https://doi.org/10.1016/0168-9002(91)90446-W).
- [132] E. Aprile et al. „The electronics read out and data acquisition system for a liquid xenon time projection chamber as a balloon-borne Compton telescope.” In: *Nuclear Instruments and Methods in Physics Research Section A* 412.2 (1998), pp. 425 –436. ISSN: 0168-9002. DOI: [10.1016/S0168-9002\(98\)00480-X](https://doi.org/10.1016/S0168-9002(98)00480-X).
- [133] R. G. Gomez. „Characterization of the Xenon-10 dark matter detector with regard to electric field and light response.” MA thesis. Rice University, 2007. URL: <https://scholarship.rice.edu/bitstream/handle/1911/20507/1441821.PDF?sequence=1&isAllowed=y>.
- [134] E. Aprile et al., the XENON collaboration. *The XENON1T Dark Matter Search Experiment*. 2012. URL: <https://arxiv.org/abs/1206.6288>.
- [135] B. Rebel et al. „High voltage in noble liquids for high energy physics.” In: *Journal of Instrumentation* 9.08 (2014), T08004. DOI: [10.1088/1748-0221/9/08/T08004](https://doi.org/10.1088/1748-0221/9/08/T08004).
- [136] A. Tomás et al. „Study and mitigation of spurious electron emission from cathodic wires in noble liquid time projection chambers.” In: *Astroparticle Physics* 103 (2018), pp. 49 –61. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2018.07.001](https://doi.org/10.1016/j.astropartphys.2018.07.001).
- [137] V. Velan. *ER/NR Discrimination in Liquid Xenon with the LUX Experiment*. Identification of Dark Matter conference presentation. 2018. URL: https://indico.cern.ch/event/699961/contributions/3043414/attachments/1691959/2722381/VelanIDM2018_Final.pdf.

- [138] E. Aprile et al. *Simultaneous measurement of the light and charge response of liquid xenon to low-energy nuclear recoils at multiple electric fields*. 2018. URL: <https://arxiv.org/abs/1809.02072>.
- [139] A. Ferella et al. „Purity measurements in Liquid Xenon. Electron lifetime dependence on circulation time and rate.” URL: https://www.lngs.infn.it/images/REIS/Annual_Report/PREPRINT/preprint97.pdf.
- [140] A. D. Ferella. „Measuring electron lifetime and V_0 in Liquid Xenon.” In: *Nuclear Instruments and Methods in Physics Research Section A* 572.1 (2007), pp. 488–489. ISSN: 0168-9002. DOI: [10.1016/j.nima.2006.10.153](https://doi.org/10.1016/j.nima.2006.10.153).
- [141] J. Angle et al., the XENON collaboration. „First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory.” In: *Physical Review Letters* 100 (2 2008), p. 021303. DOI: [10.1103/PhysRevLett.100.021303](https://doi.org/10.1103/PhysRevLett.100.021303).
- [142] E. Aprile et al., the XENON collaboration. „Design and performance of the XENON10 dark matter experiment.” In: *Astroparticle Physics* 34.9 (2011), pp. 679–698. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2011.01.006](https://doi.org/10.1016/j.astropartphys.2011.01.006).
- [143] D. S. Akerib et al., the LZ collaboration. *LUX-ZEPLIN (LZ) Conceptual Design Report*. 2015. URL: <https://arxiv.org/abs/1509.02910>.
- [144] M. Szydagis et al. *Noble Element Simulation Technique v2.0*. July 2018. DOI: [10.5281/zenodo.1314669](https://doi.org/10.5281/zenodo.1314669).
- [145] COMSOL Multiphysics ® v. 5.1. <https://www.comsol.com/>. COMSOL AB, Stockholm, Sweden.
- [146] G. Bakale, U. Sowada, and W. F. Schmidt. „Effect of an electric field on electron attachment to sulfur hexafluoride, nitrous oxide, and molecular oxygen in liquid argon and xenon.” In: *The Journal of Physical Chemistry* 80.23 (1976), pp. 2556–2559. ISSN: 0022-3654. DOI: [10.1021/j100564a006](https://doi.org/10.1021/j100564a006).
- [147] E. Hogenbirk et al. *Field dependence of electronic recoil signals in a dual-phase liquid xenon time projection chamber*. 2018. URL: <https://arxiv.org/abs/1807.07121>.
- [148] P. Sorensen et al., the XENON collaboration. „The scintillation and ionization yield of liquid xenon for nuclear recoils.” In: *Nuclear Instruments and Methods in Physics Research Section A* 601.3 (2009), pp. 339–346. ISSN: 0168-9002. DOI: [10.1016/j.nima.2008.12.197](https://doi.org/10.1016/j.nima.2008.12.197).
- [149] A. Tan et al., the PandaX collaboration. „Dark Matter Results from First 98.7 Days of Data from the PandaX-II Experiment.” In: *Physical Review Letters* 117 (12 2016), p. 121303. DOI: [10.1103/PhysRevLett.117.121303](https://doi.org/10.1103/PhysRevLett.117.121303).
- [150] D. S. Akerib et al., the LUX collaboration. *Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics*. 2016. URL: <https://arxiv.org/abs/1608.05381>.
- [151] J. R. Verbus. „An Absolute Calibration of Sub-1 keV Nuclear Recoils in Liquid Xenon Using D-D Neutron Scattering Kinematics in the LUX Detector.” PhD thesis. Brown University, 2016. DOI: [10.7301/Z01G0JQ7](https://doi.org/10.7301/Z01G0JQ7).
- [152] J. R. Verbus et al. „Proposed low-energy absolute calibration of nuclear recoils in a dual-phase noble element TPC using D-D neutron scattering kinematics.” In: *Nuclear Instruments and Methods in Physics Research Section A* 851 (2017), pp. 68–81. ISSN: 0168-9002. DOI: [10.1016/j.nima.2017.01.053](https://doi.org/10.1016/j.nima.2017.01.053).

[153] G. H. Miley. „A portable neutron / tunable X-ray source based on inertial electrostatic confinement.” In: *Nuclear Instruments and Methods in Physics Research* 422 (1999), pp. 16–20. DOI: 10.1016/S0168-9002(98)01108-5.

[154] M. Guillame et al. „On the optimal generation of 14 MeV neutrons by means of tritiated titanium target.” In: *Nuclear Instruments and Methods* 92 (1971), p. 571. DOI: 10.1016/0029-554X(71)90116-9.

[155] G. H. Miley and J. Sved. „The IEC - A plasma-target-based neutron source.” In: *Applied Radiation and Isotopes* 48.10-12 (1997), pp. 1557–1561. ISSN: 09698043. DOI: 10.1016/S0969-8043(97)00257-1.

[156] G. H. Miley and J. Sved. „The IEC star-mode fusion neutron source for NAA - Status and next-step designs.” In: *Applied Radiation and Isotopes* 53.4-5 (2000), pp. 779–783. ISSN: 09698043. DOI: 10.1016/S0969-8043(00)00215-3.

[157] G. J. Csikai. *CRC Handbook of Fast Neutron generators*. CRC Press Inc., Boca Raton, FL, 1987.

[158] D. Chernikova et al. „Experimental and numerical investigations of radiation characteristics of Russian portable/compact pulsed neutron generators: ING-031, ING-07, ING-06 and ING-10-20-120.” In: *Nuclear Instruments and Methods in Physics Research Section A* 746 (2014), pp. 74–86. DOI: 10.1016/j.nima.2014.01.061.

[159] Eljen Technology. *Neutron/Gamma PSD Liquid Scintillator EJ-301, EJ-309*. 2016. URL: http://www.eljentechnology.com/images/products/data_sheets/EJ-301_EJ-309.pdf.

[160] F. D. Brooks. „A scintillation counter with neutron and gamma-ray discriminators.” In: *Nuclear Instruments and Methods* 4.3 (1959), pp. 151 –163. ISSN: 0029-554X. DOI: 10.1016/0029-554X(59)90067-9.

[161] F. T. Kuchnir and F. J. Lynch. „Time Dependence of Scintillations and the Effect on Pulse-Shape Discrimination.” In: *IEEE Transactions on Nuclear Science* 15.3 (1968), pp. 107–113. ISSN: 0018-9499. DOI: 10.1109/TNS.1968.4324923.

[162] L. J. Perkins and M. C. Scott. „The application of pulse shape discrimination in NE213 to neutron spectrometry.” In: *Nuclear Instruments and Methods* 166.3 (1979), pp. 451 –464. ISSN: 0029-554X. DOI: 10.1016/0029-554X(79)90534-2.

[163] R. F. Lang et al. „Improved pulse shape discrimination in EJ301 liquid scintillators.” In: *Nuclear Instruments and Methods in Physics Research Section A* 856 (2017), pp. 26 –31. ISSN: 0168-9002. DOI: 10.1016/j.nima.2017.02.090.

[164] H. Klein and S. Neumann. „Neutron and photon spectrometry with liquid scintillation detectors in mixed fields.” In: *Nuclear Instruments and Methods in Physics Research Section A* 476.1–2 (2002), pp. 132 –142. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(01)01410-3.

[165] V. V. Verbinski et al. „Calibration of an organic scintillator for neutron spectrometry.” In: *Nuclear Instruments and Methods* 65.1 (1968), pp. 8 –25. ISSN: 0029-554X. DOI: 10.1016/0029-554X(68)90003-7.

[166] N. Luo, M. Ragheb, and G. H. Miley. „Proton bremsstrahlung and its radiation effects in fusion reactors.” In: *Fusion Engineering and Design* 85.1 (2010), pp. 39–45. ISSN: 09203796. DOI: 10.1016/j.fusengdes.2009.05.005.

- [167] R. Nolte and D. J. Thomas. „Monoenergetic fast neutron reference fields: II. Field characterization.” In: *Metrologia* 48.6 (2011), S274. DOI: 10.1088/0026-1394/48/6/S05.
- [168] M. R. Bhat. „Evaluated Nuclear Structure Data File (ENSDF).” In: *Nuclear Data for Science and Technology: Proceedings of an International Conference, held at the Forschungszentrum Jülich, Fed. Rep. of Germany, 13–17 May 1991*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 817–821. ISBN: 978-3-642-58113-7. DOI: 10.1007/978-3-642-58113-7_227.
- [169] H. Liskien and A. Paulsen. „Neutron production cross sections and energies for the reactions T(p,n)³He, D(d,n)³He, and T(d,n)⁴He.” In: *Atomic Data and Nuclear Data Tables* 11.7 (1973), pp. 569 –619. ISSN: 0092-640X. DOI: 10.1016/S0092-640X(73)80081-6.
- [170] G. Dietze and H. Klein. „NRESP4 and NEFF4: Monte Carlo Codes for the Calculation of Neutron Response Functions and Detection Efficiencies for NE213 Scintillation Detectors.” In: *PTB Report PTB-ND-22* (1982). ISSN 0572-7170.
- [171] A. Zimbal et al. „High resolution neutron spectrometry with liquid scintillation detectors for fusion applications.” In: *Proceedings of Science* FNDA2006 (2006), p. 035.
- [172] T. K. Alexander and F. S. Goulding. „An amplitude-insensitive system that distinguishes pulses of different shapes.” In: *Nuclear Instruments and Methods* 13 (1961), pp. 244–246. ISSN: 0029554X. DOI: 10.1016/0029-554X(61)90198-7.
- [173] T. Novotny. „Photon spectrometry in mixed neutron-photon fields using NE213 liquid scintillation detectors.” In: *PTB-Report PTB-N-28* (1997). ISBN 3-89429-853-7.
- [174] M. Matzke. „Unfolding of Pulse Height Spectra: The HEPRO Program System.” In: *PTB-Report PTB-N-19* (1994). ISBN 3-89429-543-0.
- [175] M. Reginatto, P. Goldhagena, and S. Neumann. „Spectrum unfolding, sensitivity analysis and propagation of uncertainties with the maximum entropy deconvolution code MAXED.” In: *Nuclear Instruments and Methods in Physics Research Section A* 476 (2002), pp. 242 –246. DOI: 10.1016/S0168-9002(01)01439-5.
- [176] J. D. Huba. *NRL Plasma Formulary*. Naval Research Laboratory. 2013.
- [177] N. J. Roberts, H. Tagziria, and D. J. Thomas. „Determination of the effective centres of the NPL Long Counters.” In: *NPL Report DQL-RN 004* (2004). ISSN 1744-0629.
- [178] E. Aprile et al., the XENON collaboration. „Response of the XENON100 Dark Matter Detector to Nuclear Recoils.” In: *Physical Review D* D88 (2013), p. 012006. DOI: 10.1103/PhysRevD.88.012006.
- [179] J. Aalbers. Private communication, 2018.
- [180] E. Aprile et al., the XENON collaboration. „Dark Matter Results from 100 Live Days of XENON100 Data.” In: *Physical Review Letters* 107 (13 2011), p. 131302. DOI: 10.1103/PhysRevLett.107.131302.
- [181] A. Tan et al., the PandaX collaboration. „Dark matter search results from the commissioning run of PandaX-II.” In: *Physical Review D* 93 (12 2016), p. 122009. DOI: 10.1103/PhysRevD.93.122009.
- [182] D. S. Akerib et al., the LUX collaboration. „Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data.” In: *Physical Review Letters* 116 (16 2016), p. 161301. DOI: 10.1103/PhysRevLett.116.161301.

[183] D. S. Akerib et al., the LUX collaboration. „Ultralow energy calibration of LUX detector using ^{127}Xe electron capture.” In: *Physical Review D* 96 (11 2017), p. 112011. DOI: [10.1103/PhysRevD.96.112011](https://doi.org/10.1103/PhysRevD.96.112011).

[184] L. W. Kastens et al. „Calibration of a liquid xenon detector with $^{83}\text{Kr}^m$.” In: *Physical Review C* 80 (4 2009), p. 045809. DOI: [10.1103/PhysRevC.80.045809](https://doi.org/10.1103/PhysRevC.80.045809).

[185] E. M. Boulton et al. „Calibration of a two-phase xenon time projection chamber with a ^{37}Ar source.” In: *Journal of Instrumentation* 12.08 (2017), Po8004. DOI: [10.1088/1748-0221/12/08/P08004](https://doi.org/10.1088/1748-0221/12/08/P08004).

[186] E. Aprile et al., the XENON collaboration. „Search for Electronic Recoil Event Rate Modulation with 4 Years of XENON100 Data.” In: *Physical Review Letters* 118 (10 2017), p. 101101. DOI: [10.1103/PhysRevLett.118.101101](https://doi.org/10.1103/PhysRevLett.118.101101).

[187] M. Kobayashi et al., the XMASS collaboration. *Search for sub-GeV dark matter by annual modulation using XMASS-I detector*. 2018. URL: <https://arxiv.org/abs/1808.06177>.

[188] D. S. Akerib et al., the LUX collaboration. „First Searches for Axions and Axion-like Particles with the LUX Experiment.” In: *Physical Review Letters* 118 (26 2017), p. 261301. DOI: [10.1103/PhysRevLett.118.261301](https://doi.org/10.1103/PhysRevLett.118.261301).

[189] E. Aprile et al., the XENON collaboration. „Effective field theory search for high-energy nuclear recoils using the XENON100 dark matter detector.” In: *Physical Review D* 96 (4 2017), p. 042004. DOI: [10.1103/PhysRevD.96.042004](https://doi.org/10.1103/PhysRevD.96.042004).

[190] E. Aprile et al., the XENON collaboration. „Low-mass dark matter search using ionization signals in XENON100.” In: *Physical Review D* 94 (9 2016), p. 092001. DOI: [10.1103/PhysRevD.94.092001](https://doi.org/10.1103/PhysRevD.94.092001).

[191] M. J. Dolan, F. Kahlhoefer, and C. McCabe. „Directly Detecting Sub-GeV Dark Matter with Electrons from Nuclear Scattering.” In: *Physical Review Letters* 121 (10 2018), p. 101801. DOI: [10.1103/PhysRevLett.121.101801](https://doi.org/10.1103/PhysRevLett.121.101801).

[192] G. Bertone and T. M. P. Tait. „A new era in the search for dark matter.” In: *Nature* 562.7725 (2018), pp. 51–56. ISSN: 1476-4687. DOI: [10.1038/s41586-018-0542-z](https://doi.org/10.1038/s41586-018-0542-z).

[193] D. S. Akerib et al., the CDMS collaboration. „First Results from the Cryogenic Dark Matter Search in the Soudan Underground Laboratory.” In: *Physical Review Letters* 93 (21 2004), p. 211301. DOI: [10.1103/PhysRevLett.93.211301](https://doi.org/10.1103/PhysRevLett.93.211301).

[194] D. S. Akerib et al., the CDMS collaboration. „Limits on Spin-Independent Interactions of Weakly Interacting Massive Particles with Nucleons from the Two-Tower Run of the Cryogenic Dark Matter Search.” In: *Physical Review Letters* 96 (1 2006), p. 011302. DOI: [10.1103/PhysRevLett.96.011302](https://doi.org/10.1103/PhysRevLett.96.011302).

[195] G. J. Alner et al., the ZEPLIN collaboration. „First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two-phase xenon detector for dark matter detection.” In: *Astroparticle Physics* 28.3 (2007), pp. 287 –302. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2007.06.002](https://doi.org/10.1016/j.astropartphys.2007.06.002).

[196] V. N. Lebedenko et al. „Results from the first science run of the ZEPLIN-III dark matter search experiment.” In: *Physical Review D* 80 (5 2009), p. 052010. DOI: [10.1103/PhysRevD.80.052010](https://doi.org/10.1103/PhysRevD.80.052010).

[197] D. Akimov et al., the ZEPLIN collaboration. „WIMP-nucleon cross-section results from the second science run of ZEPLIN-III.” In: *Physics Letters B* 709.1 (2012), pp. 14 –20. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.01.064](https://doi.org/10.1016/j.physletb.2012.01.064).

- [198] E. Aprile et al., the XENON collaboration. „First Dark Matter Results from the XENON100 Experiment.” In: *Physical Review Letters* 105 (13 2010), p. 131302. doi: 10.1103/PhysRevLett.105.131302.
- [199] D. S. Akerib et al., the LZ collaboration. *Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment.* 2018. URL: <https://arxiv.org/abs/1802.06039>.
- [200] H. Zhang et al., the PandaX collaboration. *Dark matter direct search sensitivity of the PandaX-4T experiment.* 2018. URL: <https://arxiv.org/abs/1806.02229>.
- [201] L. Zhao and J. Liu. „PandaX: A deep underground dark matter search experiment in China using liquid xenon.” In: *Modern Physics Letters A* 33.30 (2018), p. 1830013. doi: 10.1142/S0217732318300136.
- [202] J. Liu, X. Chen, and X. Ji. „Current status of direct dark matter detection experiments.” In: *Nature Physics* 13 (2017), pp. 212–216. doi: 10.1038/nphys4039.