First WIMP results of XENONnT and its signal reconstruction

Angevaare, J.R.

Publication date
2023
Document Version
Final published version

First WIMP results of XENONnT and its signal reconstruction

Joran Remco Angevaare

First WIMP results of XENONnT and its signal reconstruction

Joran Remco Angevaare

First WIMP results of XENONnT and its signal reconstruction

Joran Remco Angevaare

Cover created using the `fast-neural-style` neural network with the mosaic style from the `pytorch` example library. The original image shows the inside of the XENON1T experiment during its disassembly, seen through the holders of the light sensors at the bottom of the detector. Photo credit: XENON collaboration.

ISBN: 978-94-6419-809-6

Printed by Gildeprint

First WIMP results of XENONnT and its signal reconstruction

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op woensdag 7 juni 2023, te 16.00 uur

door Joran Remco Angevaare
geboren te Amsterdam

Promotiecommissie

Promotores:	prof. dr. M.P. Decowski	Universiteit van Amsterdam
	prof. dr. A.P. Colijn	Universiteit van Amsterdam
Overige leden:	dr. F. de Almeida Dias	Universiteit van Amsterdam
	prof. dr. G. Bertone	Universiteit van Amsterdam
	prof. dr. R. Fleischer	Vrije Universiteit Amsterdam
	dr. T.R. Pollmann	Universiteit van Amsterdam
	prof. dr. M. Schumann	University of Freiburg
	prof. dr. I.B. van Vulpen	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Publications

This work features two chapters (chapter 4 and 5) that are based on the two following publications. The author provided essential contributions to both, as detailed for each publication below.

Chapter 4: J. R. Angevaare, G. Bertone, A. P. Colijn, M. P. Decowski, and B. J. Kavanagh, *Complementarity of direct detection experiments in search of light Dark Matter*, JCAP 10 (2022) 004, reference [1].

All authors were involved with defining and refining the scope of the paper. JRA performed the required analyses, statistical inference, wrote the code [2] and the paper. All authors were essential to the content of the paper due to their helpful feedback. All figures were produced by JRA.

Chapter 5: The XENON collaboration: E. Aprile, J. Aalbers, K. Abe, F. Agostini, S. Ahmed Maouloud, L. Althueser, B. Andrieu, E. Angelino, J. R. Angevaare, V. C. Antochi, D. Antón Martin, F. Arneodo, L. Baudis, A. L. Baxter, L. Bellagamba, R. Biondi, A. Bismark, E. J. Brookes, A. Brown, S. Bruenner, G. Bruno, R. Budnik, T. K. Bui, C. Cai, J. M. R. Cardoso, D. Cichon, A. P. Cimental Chavez, D. Coderre, A. P. Colijn, J. Conrad, J. J. Cuenca-García, J. P. Cussonneau, V. D'Andrea, M. P. Decowski, P. Di Gangi, S. Di Pede, S. Diglio, K. Eitel, A. Elykov, S. Farrell, A. D. Ferella, C. Ferrari, H. Fischer, M. Flierman, W. Fulgione, C. Fuselli, P. Gaemers, R. Gaior, A. Gallo Rosso, M. Galloway, F. Gao, R. Glade-Beucke, L. Grandi, J. Grigat, M. Guida, R. Hammann, A. Higuera, C. Hils, L. Hoetzsch, N. F. Hood, J.

Howlett, M. Iacovacci, Y. Itow, J. Jakob, F. Joerg, A. Joy, N. Kato, M. Kara, P. Kavrigin, S. Kazama, M. Kobayashi, G. Koltman, A. Kopec, F. Kuger, H. Landsman, R. F. Lang, L. Levinson, I. Li, S. Li, S. Liang, S. Lindemann, M. Lindner, K. Liu, J. Loizeau, F. Lombardi, J. Long, J. A. M. Lopes, Y. Ma, C. Macolino, J. Mahlstedt, A. Mancuso, L. Manenti, F. Marignetti, T. Marrodán Undagoitia, K. Martens, J. Masbou, D. Masson, E. Masson, S. Mastroianni, M. Messina , K. Miuchi, K. Mizukoshi, A. Molinario, S. Moriyama, K. Morå, Y. Mosbacher, M. Murra, J. Müller, K. Ni, U. Oberlack, B. Paetsch, J. Palacio, R. Peres, C. Peters, J. Pienaar, M. Pierre, V. Pizzella, G. Plante, J. Qi, J. Qin, D. Ramírez García, A. Rocchetti, L. Sanchez, P. Sanchez-Lucas, J. M. F. dos Santos, I. Sarnoff, G. Sartorelli, J. Schreiner, D. Schulte, P. Schulte, H. Schulze Eißing, M. Schumann, L. Scotto Lavina, M. Selvi, F. Semeria, P. Shagin, S. Shi, E. Shockley, M. Silva, H. Simgen, A. Takeda, P.-L. Tan, A. Terliuk, D. Thers, F. Toschi, G. Trinchero, C. Tunnell, F. Tönnies, K. Valerius, G. Volta, C. Weinheimer, M. Weiss, D. Wenz, C. Wittweg, T. Wolf, D. Xu, Z. Xu, M. Yamashita, L. Yang, J. Ye, L. Yuan, G. Zavattini, S. Zerbo, M. Zhong, T. Zhu , *The Triggerless Data Acquisition System of the XENONnT Experiment*, arXiv preprint arXiv:2212.11032 (2022), reference [3].

Corresponding authors: J. R. Angevaare, A. Elykov, D. Masson, and S. Mastroianni.

The design and commissioning of the DAQ was performed by the XENON collaboration. JA (J. Aalbers), JRA, DC (D. Coderre), AE, PG (P. Gaemers), DM, and FT (F. Tönnies) developed the required concepts, software and tools. The paper was written by JRA, AE, DM and SM. DM wrote the "Introduction" (section 5.1) and "System Control & Oversight" (section 5.5), and JRA wrote the "Live Processing" (section 5.4) and "Conclusion" (section 5.7) sections. DM, AE and SM wrote the "Data Acquisition & Readout" (section 5.3) section, and AE and JRA wrote the "Performance" section (section 5.6). All authors contributed to the "From XENON1T to XENONnT" section (section 5.2). JRA performed the analyses of subsection 5.6.2 and subsection 5.6.4. JRA made Fig.'s 5.1, 5.3, 5.5, 5.7, AE Fig.'s 5.2.5 and 5.6, and DM Fig. 5.4. All authors participated in the refining of the paper by iteratively reviewing the other authors' sections. The feedback of the XENON collaboration during the internal review process improved the quality of the paper. JRA was essential to the internal review process within the XENON collaboration, and submitted this preprint version to arxiv and prepared it for JINST submission.

Contents

Summary
of "First WIMP results of XENONnT and its signal reconstruction"

As physicists, we are trying to solve the puzzle that is our Universe, a very tough challenge with only the few pieces we have so far. The pieces we have, that indicate ordinary, baryonic matter, only account for 16% of the total mass of the Universe. The remaining 84% should be some form of matter that has never been observed, yet. This mysterious matter component is called "Dark Matter", for it does not interact with light directly, making it hard to detect.

In this work we will show the first Dark Matter search results of the XENONnT experiment in the last chapter (Chapter 7). First, we must understand why Dark Matter exists (Chapter 1) and what it could be (Chapter 2). We also explain the XENONnT experiment (Chapter 3), and how XENONnT together with the SuperCDMS experiment could reconstruct Dark Matter (Chapter 4). One day, we hope to discover Dark Matter and reconstruct its properties as in Fig. 1. The digitization of the XENONnT data (Chapter 5) and reconstruction (Chapter 6) are essential to the DM search of the XENONnT detector (Chapter 7), which did not detect Dark Matter, but did set stringent limits (Fig. 2). Below, we provide a brief summary of each of the chapters.

Chapter 1: Why does Dark Matter exist?

The evidence of Dark Matter, the extra source of gravity, is seen on many scales, from galaxies such as our own to the oldest imprint of the Universe seen in the electromagnetic spectrum, the Cosmic Microwave Background. The Cosmic Microwave Background was emitted 380,000 years after the big bang at the time that the Universe became transparent to photons. The small fluctuations in the Cosmic Microwave Background allow precise modeling of the contributions of

baryonic matter and Dark Matter. These small fluctuations are enhanced by the Dark Matter content, since it works as a gravitational well when clumped together, in contrast to the hot plasma of ordinary matter, which would prevent clumping due to the pressure it experiences. Too much Dark Matter would result in too much clumping, and thereby too many overdensities in the Cosmic Microwave Background, too little Dark Matter would result in too few overdensities. The Λ_{CDM}-model parameterizes our current understanding of the evolution of the Universe, which fit to the latest PLANCK data shows us we have so far only observed 16% of the puzzle.

Chapter 2: Dark Matter

The exact nature of Dark Matter remains unclear, nevertheless the Weakly Interacting Massive Particle (WIMP)-model provides a well motivated candidate for the nature of Dark Matter. WIMPs are particles that would interact, besides gravity, via a mediator similar to those of the weak force. WIMPs would naturally be produced during a process called freeze-out and still be present in the Universe today. Among the many experimental techniques that strive to detect WIMPs, liquid xenon Time Projection Chambers are on the forefront of the Dark Matter search for WIMP masses $\gtrsim 5\,\mathrm{GeV}/c^2$. One of these experiments is XENONnT at Laboratori Nazionali del Gran Sasso in Italy, and it primarily aims to detect Dark Matter as an excess of scatters off xenon nuclei.

Chapter 3: XENONnT

If WIMPs exist, they will rarely interact with ordinary matter. How often this occurs is encoded in the cross section. Because this process is so rare, the XENONnT experiment has to reduce backgrounds as much as possible to have a chance of detecting this, yet unobserved, interaction. The Time Projection Chamber of the XENONnT experiment works by monitoring a large mass (8.6 t) of liquid xenon. If a particle scatters in the liquid xenon and deposits energy, it will create prompt scintillation light and ionize xenon atoms (illustrated with Fig. 3.2). The scintillation light is detected by photo-multiplier tubes (PMTs) as the first signal (S1). The freed electrons due to ionization are drifted and extracted from the liquid (into a gaseous xenon layer) by electric fields within the Time Projection Chamber. The extracted electrons create secondary scintillation in the gaseous xenon as they are accelerated, the secondary scintillation is detected as the second signal (S2). Based on the S1 and S2, the properties of the interaction are determined, the time between the S1 and the S2 being a

measure for how deep the interaction occurred in the Time Projection Chamber, while the pattern of the S2 signal on the PMTs at the top of the detector allows obtaining the xy-coordinate. The S1 and S2 size depend on the amount of energy deposited and their relative size discriminates between interactions with the xenon nucleus (nuclear recoils) and the atomic electrons (electronic recoils). These properties of the interaction are necessary to distinguish from sources of background. One of the key aspects of XENONnT that makes it an excellent Dark Matter detector is that it achieved very low backgrounds, thereby putting it in a prime position to discover rare interactions.

Chapter 4: Complementarity of light Dark Matter searches

In Chapter 4 we include Ref. [1] that focuses on finding "light" Dark Matter. If Dark Matter consists of light WIMPs with a mass of $\mathcal{O}\left(0.1 - 10\,\mathrm{GeV}/c^2\right)$, different experiments and detection techniques may be able to discover it. The nuclear recoil energies induced by this type of WIMPs may be too small to detect in XENONnT and fall below the energy threshold. Instead, they could be detected as electronic recoil signals induced by the so-called Migdal effect, caused by the displacement of the xenon nucleus with respect to its electron cloud. These electronic recoil signals would be rare, but result in detectable energy deposits as they yield a higher number of scintillation photons for the same energy. Additionally, an experiment like the Super Cryogenic Dark Matter Search (SuperCDMS) at SNOLAB in Canada, that monitors germanium and silicon crystals at cryogenic temperatures has a very low energy threshold and excellent energy resolution, making it ideal for detecting nuclear recoil signals from this type of light WIMPs. Furthermore, the Migdal effect may also extend the Dark Matter search of SuperCDMS to very low WIMP masses. Combining the data from XENONnT and SuperCDMS, using a combination of nuclear recoil and Migdal analyses can lead to a better reconstruction of the properties of the WIMP if it is detected. These experiments hope to one day detect Dark Matter, allowing them to constrain its properties such as mass (M_χ) and cross section ($\sigma_{\mathrm{S.I.}}$) as in Fig. 1. In Chapter 4, we compute the complementarity of XENONnT and SuperCDMS using the combination of nuclear recoil and Migdal analyses. We show that XENONnT reconstructs WIMPs best for masses $> 5.6\,\mathrm{GeV}/c^2$ and would significantly benefit from combining results with SuperCDMS for lower masses, allowing for more accurate reconstruction of the WIMP parameters. The greatest complimentary is found between SuperCDMS nuclear recoil and Migdal analyses (and to a lesser extent the XENONnT Migdal analysis) for masses in the range of $[0.2, 0.6]\,\mathrm{GeV}/c^2$.

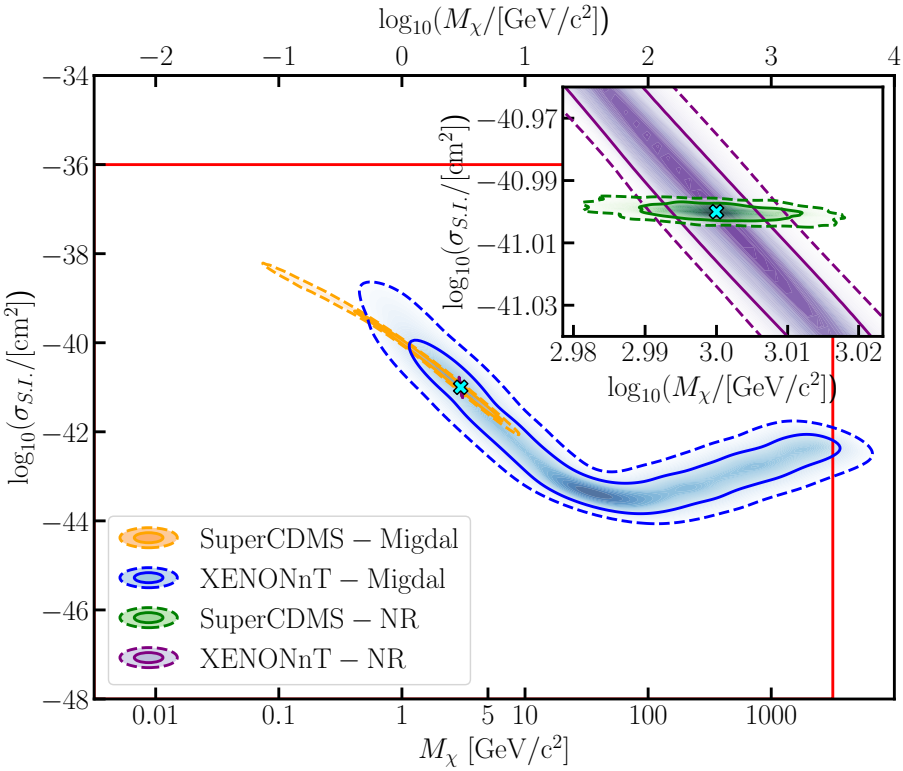

Figure 1: The reconstruction of simulated WIMPs with a mass $M_\chi = 3$ GeV/c^2 and coupling strength $\sigma_{\mathrm{S.I.}} = 10^{-41}$ cm^2 (cyan cross), using a nuclear recoil (NR) or Migdal analysis for SuperCDMS and XENONnT. The colored regions indicate where each analysis reconstructs the WIMP, the smaller the region, the better it is reconstructed, and brighter colors indicate higher probabilities. The SuperCDMS and XENONnT NR analyses both reconstruct the WIMPs (top left zoomed inset), but SuperCDMS does so more precisely. Combining the two analyses further increases the precision. The Migdal analyses have a lower precision to reconstruct these WIMP parameters, their precision is higher for lower masses (Fig. 4.7). This figure corresponds to Fig. 4.3 on page 63.

Chapter 5: The XENONnT DAQ

Chapter 5 includes Ref. [3], describing the design, commissioning, and performance of the XENONnT Data Acquisition system (DAQ). The DAQ is respon-

sible for digitizing the analogue signals from the PMTs. The XENONnT DAQ was designed around a triggerless readout of the Time Projection Chamber, where all individual signals from PMTs above a small digitization threshold are stored. The DAQ is able to handle the vastly different types of signals from fast S1s ($\sim$100 ns) to S2 signals of up to $\sim$100 µs. Additionally, the DAQ is designed to operate the three subdetectors of the Time Projection Chamber, Muon Veto, and the newly added Neutron Veto independently or as a single operational unit. When the DAQ operates as a single unit, the subdetectors are "linked" such that the clock signals are distributed to all the involved hardware and timestamps are synchronized between the three subdetectors.

The DAQ also performs the processing of the data from the PMTs. The processing is sufficiently fast and setup in such a manner that the data is fully processed (up to matched S1 and S2 signals) within $\mathcal{O}(10\,\mathrm{s})$. This data is stored in an online database which enables live monitoring of the performance of the detector.

Chapter 6: Signal reconstruction

The reconstruction by the PMTs, DAQ, and processing software is validated in Chapter 6 using simulated data. Two key concepts are directly related to the reconstruction: the energy threshold and the energy reconstruction. The first is inferred from the S1 detection efficiency where we characterize the probability, as a function of the signal size, of the S1 being reconstructed. The energy reconstruction is affected by the reconstruction "bias", which quantifies the mismatch between the true size of a signal and the reconstructed value. The reconstructed value is on average -2% to 1% of the true size of the signal depending on the size and type (S1 or S2) of the signal, predominantly due to the digitization threshold and the inclusion of PMT afterpulses. The resultant non-linearity of the energy scale is rectified by properly accounting for the reconstruction bias.

Chapter 7: First WIMP results with XENONnT

In the final chapter (Chapter 7), the WIMP results of XENONnT are discussed. No excess above the expected backgrounds is observed, yielding new, stringent exclusion limits on the properties of WIMP Dark Matter as in Fig. 2. XENONnT excludes the spin independent WIMP-nucleon cross section down to $\sigma_{n,\,\mathrm{S.I.}} = 2.2 \times 10^{-47}\,\mathrm{cm}^2$ at a WIMP mass of $26\,\mathrm{GeV}/c^2$.

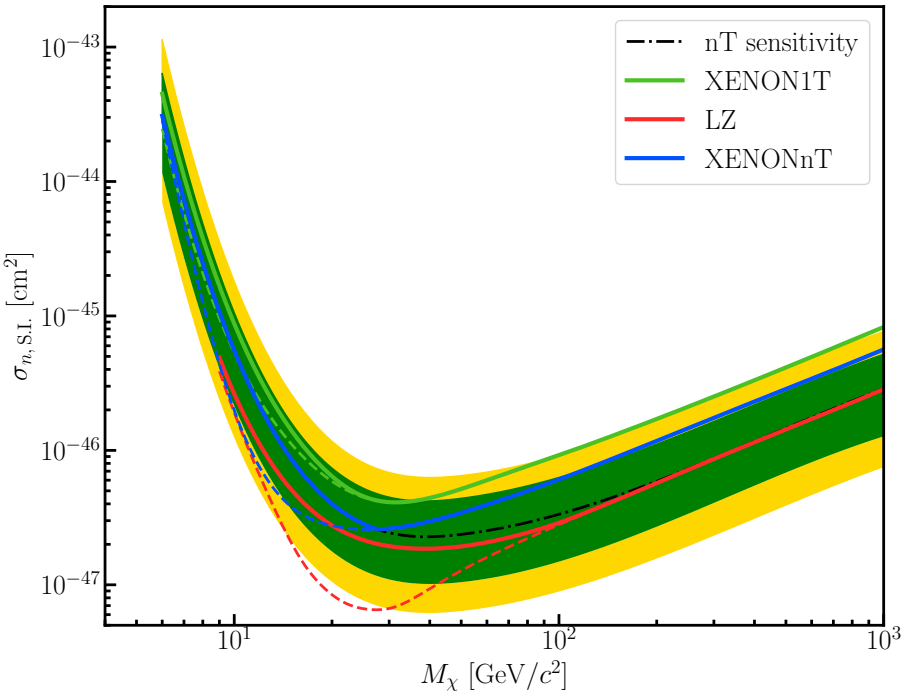

Figure 2: The XENONnT exclusion limit of WIMPs (blue line) with coupling strength $\sigma_{n,\mathrm{S.I.}}$ as a function of mass (M_χ), and the sensitivity (dash dotted black line, with 1σ and 2σ uncertainty bands in yellow and green respectively). Everything above the exclusion limit is excluded (at 90% confidence level), while WIMPs could still exist with masses and coupling strengths below the limit. The dashed lines exclude WIMPs down to the -2σ sensitivity while the solid lines only exclude to the median sensitivity. The limit of XENON1T (the predecessor of XENONnT) is indicated by the green line [4], and that of LZ (an experiment similar to XENONnT) by the red line [5], showing that the WIMP searches progresses to lower and lower coupling strengths. This figure corresponds to Fig. 7.6 on page 147 and is adapted from Ref. [6].

Chapter 1

Why does Dark Matter exist?

Our understanding of the Universe only extends to 16% of the mass of the Universe. The remaining 84% of the mass in the Universe is "Dark Matter", yet, this Dark Matter has never been observed. Before we can start the hunt for Dark Matter in this thesis, we first have to appreciate why we are joining this chase.

1.1 Introduction

A century ago, in 1922 Kapteyn first used the term "Dark Matter" to describe invisible matter that could only be inferred via gravitational effects [7]. Since then, (astro)physicists found overwhelming evidence for the existence of such a mass-component of the Universe. After many breakthroughs in (astro)particle physics where the Standard Model of particle physics (SM) is validated with good precision, there is no viable candidate for DM in the SM. A DM particle has yet escaped an unambiguous detection.

By definition DM should be the source of the additional gravity that is observed on astrophysical scales from galaxies and larger. In contrast to ordinary, baryonic, matter, DM should not exhibit strong self-interactions and is therefore distributed differently with respect to baryonic matter. In the absence of a detection, the properties of DM have not been established, as will be discussed in Chapter 2. Nevertheless, the existence of this additional source of gravity, is backed up by a plethora of experimental results.

We find evidence for DM from the dynamics of galaxies to the furthest observable signals of the Universe in the electromagnetic spectrum, the Cosmic Microwave Background (CMB). The list of evidence for DM is long and diverse, as we will illustrate with a short description of some of these experimental endeavors. A complete overview of evidence for DM can be found in any of Ref.'s [8–11].

A particularly striking piece of evidence for DM is the difference in the spatial distribution of baryonic- and DM matter demonstrated in colliding galaxy clusters. For example, the observation of 1E0657-588 (the Bullet Cluster) shows a clear separation of the hot gas in the cluster (observed in X-rays) and the main mass distribution reconstructed from gravitational lensing [12]. While ordinary matter has slowed down during the collision of the two progenitor galaxy clusters (as expected for baryonic matter), most of the mass did not, indicating a large, frictionless, DM component in the progenitor galaxy clusters. On even larger scales, we observe that DM also affects structure formation. Despite the Universe being very uniform, under- and over-densities exist, the latter resulting in structures such as galaxies. While ordinary matter experiences pressure and therefore an outward force when clumped together, the pressureless DM allows accumulation of mass more efficiently [13]. Additionally, the velocity of the DM plays a role in the structure formation. Relativistic DM would wash out structure due to its high velocity, while non-relativistic (cold) DM would enhance it. A large cold DM component needs to be considered [14–16] to successfully replicate the Universe and the structure in the distribution of galaxies in simulations. Otherwise, the pressure experienced by baryonic matter prevents sufficient accumulation of matter to match the observed structure.

In the following two sections, we elaborate on the evidence for DM in galaxies and from the CMB. In section 1.2, we discuss the DM distribution in galaxies as it is particularly relevant for the Milky Way when investigating the discovery power of Earth-based DM detectors, which is the subject of Chapter 4. The CMB allows us to determine the DM content of our Universe, which we therefore discuss in section 1.3.

1.2 Dark matter in galaxies

Kepler's law dictates that the velocity $v(R)$ of objects within a galaxy, spiraling around its center at radius R, depends on the mass enclosed within that radius

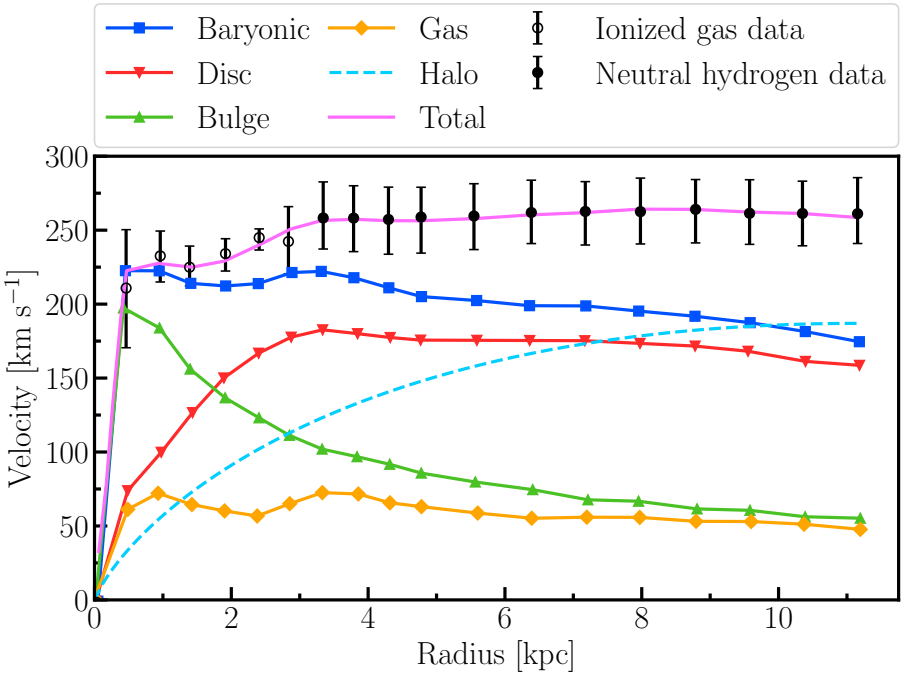

Figure 1.1: Decomposition of the rotation curve of NGC5005 [17]. The observed rotation velocity (ionized gas data and neutral hydrogen data) is not explained by the rotational velocity of baryonic matter only and requires a DM contribution ("Halo") to match the data. The total gas, stellar bulge, stellar disc, and dark matter halo model components are added in quadrature to achieve the best overall fit (solid pink line) to the observed rotation curve. Figure adapted from [17], source code available at [18].

$M(R)$,

$$v(R) = \sqrt{\frac{GM(R)}{R}}\,, \tag{1.1}$$

where G is Newton's constant. At greater R we therefore expect stars in a galaxy to orbit at a lower velocity, since the density of stars and gas decreases at larger galactic radii. However, in conflict with this expectation, observations show that the rotational velocity does not decrease as a function of R.

Fig. 1.1 compares the observed velocity distribution of galaxy NGC5005 for several components of this galaxy [17]. The total rotational velocity stays constant as a function of the radius ($R \gtrsim 3\,\text{kpc}$), which cannot be reconciled with Eq. (1.1) if only the observed mass of the galactic disk, bulge, and hot gas are included. Their total contribution results in the baryonic component in Fig. 1.1, which does not match the total observed velocity distribution, and the discrepancy becomes larger at larger radii. The observed velocity distribution can be explained by assuming the presence of a DM halo. These results for NGC5005 do not stand on their own, as the velocity distribution of many galaxies show the same behavior [7, 11, 19–21].

In order to match the observed velocity distribution of rotating galaxies, a DM halo with a Navarro–Frenk–White (NFW) mass profile can be used to match the data [22, 23]:

$$\rho(r) = \frac{\rho_0 R_s^3}{r\,(r+R_s)^2}\,, \tag{1.2}$$

where the halo density ρ is given as a function of radius r for the characteristic density ρ_0 and scale radius R_s. The characteristic density and scale radius are distinctive for each considered galaxy. The enclosed mass $M(R) = \int_0^R 4\pi r^2 \rho(r) dr$, together with Eq. (1.1) allows estimating the DM density at a given radius.

Highly relevant to Earth-based DM-searches, we also observe the discrepancy between the expected velocity distribution based on the baryonic matter and the observed velocity distribution in our galaxy. To calculate the interaction probability of DM within a detector on Earth, it is important to know the exact abundance of the DM. Recently, the second data release of the GAIA satellite [28, 29] helped to improve our understanding of the DM distribution and its structure within our Milky Way [25, 30, 31]. Ref. [25] found that the Milky Way can be modeled using a slightly adapted NFW halo model[1]. The velocity distribution is modeled comprising constitutions from the DM Halo, the galactic disc, the bulge and the circumgalactic medium (CGM). Fig. 1.2 shows the model of the rotation curve within the Milky Way [25] together with data inferred from the Gaia satellite [24]. From this model, the DM density at the location of the Sun is estimated to be $0.33(2)\,\text{GeV/cm}^3$. In addition to this result, other methods also allow determining the DM density and their results often have statistical uncertainties which are smaller than the scatter between them. Values for the DM density are found to be between $0.4 - 0.6\,\text{GeV/cm}^3$

[1]This model is adapted from the NFW model in Eq. (1.2) to account for the contraction of DM due to baryons settling in the center of the Milky Way [25].

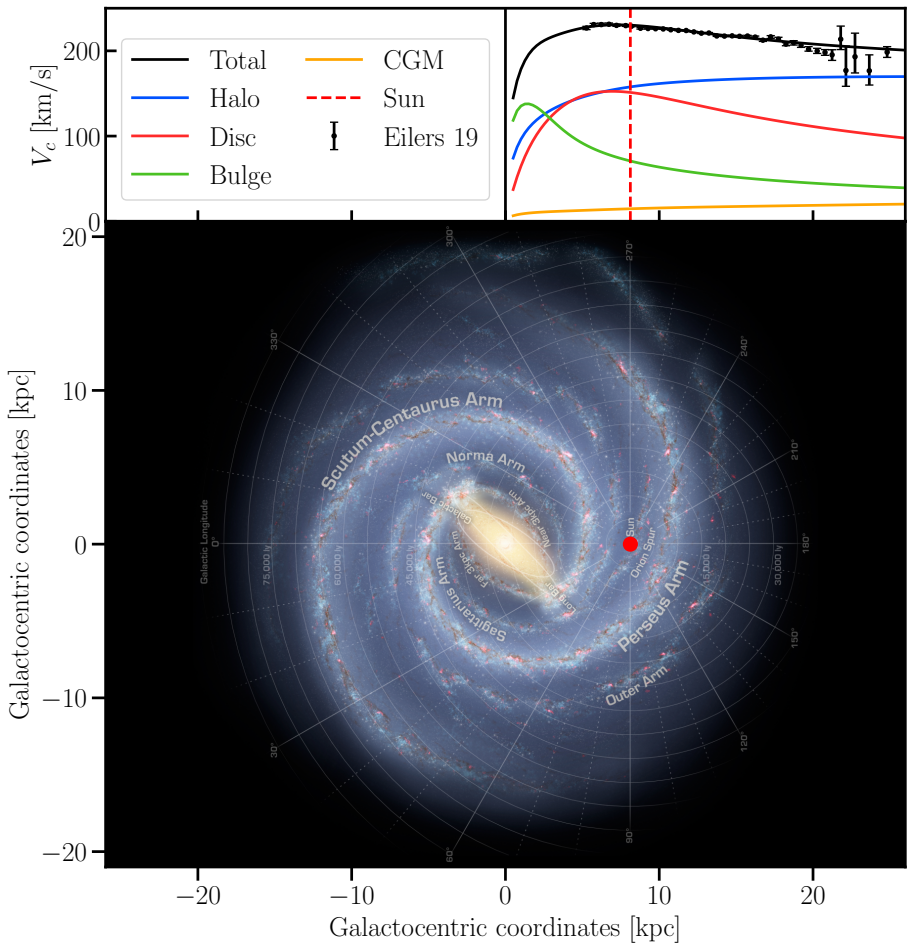

Figure 1.2: Rotation curve of the Milky Way (top panel) where the data [24] matches the vector sum of the modeled contributions [25], see text for details. The bottom panel shows an artist impression of our Milky Way [26]. The position of the Sun in the top panel (red line) illustrates the respective expected contributions at Earth, in the bottom panel the position of the Sun (red dot) is overlaid on the Milky Way. Source code available at [18], based on Ref.'s [26, 27].

or $0.3-0.5\,\mathrm{GeV/cm^3}$ depending on the type of analysis, Ref. [32] provides an overview of recent publications on the value of the DM density at the location of the Sun.

While the measurements of galaxy rotation curves are well studied, alternative models to DM exist that explain the galaxy rotation curves, as will be discussed in subsection 2.1.3. In fact, assuming sufficient amounts of ordinary matter that is very hard to detect would also describe these rotation curves. However, such models are strongly disfavored by cosmological evidence, as we will show. It is therefore presently believed that the most likely explanation for the galaxy rotation curves is the existence of DM-halos [33].

1.3 Cosmological evidence

We also see the imprint of DM in the temperature fluctuations of the Cosmic Microwave Background (CMB). When modeling the evolution of the Universe since the Big Bang, a DM contribution is required to match the observations [34].

1.3.1 The Λ_{CMD} model

The Universe is expanding at an increasing rate. The time dependent Hubble parameter, $H(t)$, characterizes the expansion of the Universe,

$$H(t) = \frac{\dot{a}(t)}{a(t)}\,, \tag{1.3}$$

where $a(t)$ is the scale factor of the expanding space-time.

The Einstein equation relates the energy content and the metric of the Universe,

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}\,, \tag{1.4}$$

where $G_{\mu\nu}$ is the Einstein tensor, Λ is the cosmological constant, $g_{\mu\nu}$ is the metric tensor, c the speed of light, and $T_{\mu\nu}$ the energy content of the Universe. The Einstein equation can be solved exactly with the Friedmann–Lemaître–Robertson–Walker metric describing a flat, homogeneous, isotropic and expanding Universe for which we can write the metric tensor

$$g_{\mu\nu} = \begin{pmatrix} -c & & & \\ & a^2(t) & & \\ & & a^2(t) & \\ & & & a^2(t) \end{pmatrix}. \tag{1.5}$$

Using this metric and the Einstein equation, we can relate $H(t)$ to the density of the Universe, $\rho(t)$,

$$H^2(t) = \left(\frac{\dot{a}(t)}{a(t)}\right)^2 = \frac{3\pi G}{3c^4}\rho(t)\,. \tag{1.6}$$

The evolution of $a(t)$ as a function of $\rho(t)$ depends on the composition of the Universe.

For non-relativistic matter, the energy density is the energy per particle per unit volume. The energy per particle is the rest-mass of the particle, which is invariant under the evolution of $a(t)$. As a consequence, the energy density of non-relativistic matter scales inversely proportional to the volume, and therefore $\propto a^{-3}(t)$ as $a(t)$ increases as a function of time. The energy density in photons, on the other hand, decreases with an additional power of $a^{-1}(t)$, as the energy per photon decreases with its wavelength. As $a(t)$ increases over time, the wavelength correspondingly increases, which lowers the photon energy. As such, the energy density of radiation is the product of the number density and the average energy per photon and scales as $\propto a^{-4}(t)$. Finally, observations indicate a third type of energy density in our Universe, that does not scale with $a(t)$ and is given by the cosmological constant Λ [35]. The contribution of Λ to the energy density of the Universe is called "Dark Energy". If we solve Eq. (1.6) and substitute that for the present day (t_0), $H(t_0) = H_0$ and $a(t_0) = 1$, we find [34]:

$$\rho(t) = \frac{3H_0^2c^4}{8\pi G}\left(\Omega_{\mathrm{R},\,0}\,a^{-4}(t) + \Omega_{\mathrm{M},\,0}\,a^{-3}(t) + \Omega_{\Lambda,\,0}\right) \tag{1.7}$$

$$= \rho_{\mathrm{cr},\,0}\left(\Omega_{\mathrm{R},\,0}\,a^{-4}(t) + \Omega_{\mathrm{M},\,0}\,a^{-3}(t) + \Omega_{\Lambda,\,0}\right)\,, \tag{1.8}$$

where the energy density is expressed in terms of the critical density $\rho_{\mathrm{cr},\,0}$, radiation density $\Omega_{\mathrm{R},\,0}$, matter density $\Omega_{\mathrm{M},\,0}$ and Dark Energy $\Omega_{\Lambda,\,0}$.

The matter density $\Omega_{\mathrm{M},\,0}$ is about 30% of the energy content of the Universe as measured by the Planck Collaboration [35]. However, the mass of baryonic matter only contributes to about 5% of this density. The remaining 25% should be DM such that the matter contribution can be split in a baryonic ($\Omega_{\mathrm{b},\,0}$) and cold DM part ($\Omega_{\mathrm{c},\,0}$):

$$\Omega_{\mathrm{M},\,0} = \Omega_{\mathrm{b},\,0} + \Omega_{\mathrm{c},\,0}\,. \tag{1.9}$$

This model describes the evolution of the Universe and is called the Λ_{CMD} model. This model can be tested with good precision with the data from CMB measurements, as discussed in the following section.

1.3.2 Cosmic Microwave Background radiation

The earliest signal in the electromagnetic spectrum is the Cosmic Microwave Background (CMB) radiation. The CMB was emitted when the Universe expanded to such an extent that the adiabatic decrease in temperature allowed the formation of neutral atoms. At this point, the Universe transitioned from a plasma state to becoming opaque to photons. We can still observe these photons today in the CMB.

The CMB-radiation describes an almost perfect blackbody spectrum with an average temperature of ~ 2.725 K. While the CMB is very isotropic, small anisotropies are observed in the temperature, the latest result come from the Planck collaboration [35]. The temperature differences result from over- and under densities in the early Universe and are a probe of the DM content in the Universe. While baryonic matter experiences pressure when clumped together, DM does not [13]. Both types of matter contribute to the total matter content ($\Omega_{\mathrm{M},0}$), but their effect on structure formation is different. Pressureless DM freely accumulates and enhances structures formation, on the other hand, baryonic matter experiences pressure, so it accumulates less efficiently. As such, the CMB is an excellent probe for the Λ_{CMD}-model or alternative models since the structures observed in the CMB strongly depend on the DM content in the Universe.

The Λ_{CMD} model describes the temperature anisotropies measured in the CMB accurately. By decomposing the temperature fluctuations in the CMB into spherical harmonics, the temperature fluctuations can be expressed as a function of the multipole moment ℓ, where increasing numbers of ℓ correspond to increasingly smaller angular scales. For example, $\ell = 0$ corresponds to the average CMB temperature, $\ell = 1$ is the dipole moment, $\ell = 2$ the quadruple moment and so forth. The power spectrum of the multipole expansion of the CMB provides a powerful tool to test Λ_{CMD} as the precise locations of the peaks in the power spectrum depend on the exact interplay of radiation, matter, and DM content.

Fig. 1.3 gives the multipole expansion of the CMB temperature anisotropy measurements (by the Planck collaboration [35]) together with the fit for Λ_{CMD}, which are in excellent agreement [35]. The fit of Λ_{CMD} yields the following contributions to the energy density of the present-day Universe,

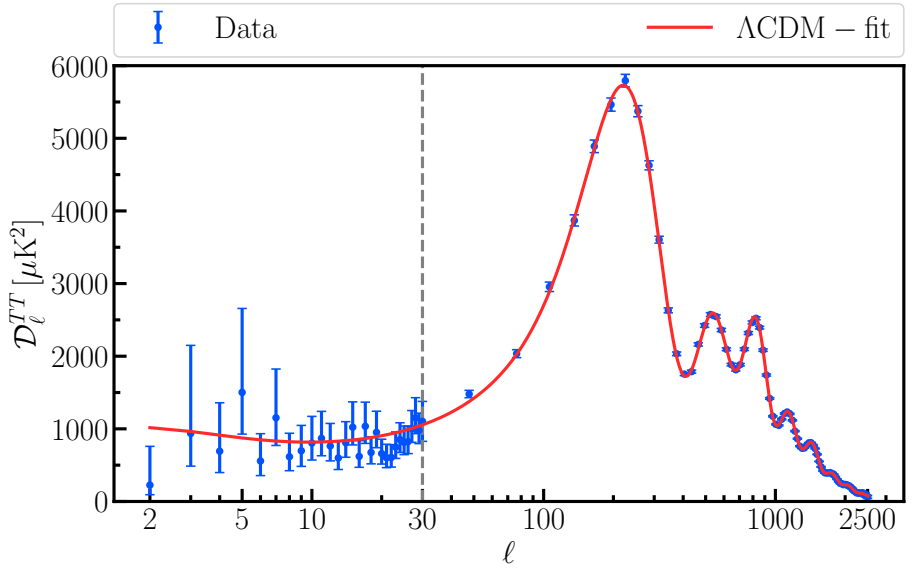

Figure 1.3: Power spectrum ($\mathcal{D}_\ell^{TT}$) of the CMB temperature anisotropy measurements (blue) as a function of the multipole moment ℓ and the Λ_{CMD} best fit (red). The data points above the vertical dashed line (gray) are rebinned. Figure adapted from [35], source code available at [18].

$$
\begin{aligned}
\Omega_{\Lambda,\,0} &= 0.691 \pm 0.006\,, \\
\Omega_{\mathrm{M},\,0} &= 0.309 \pm 0.006\,, \\
\Omega_{\mathrm{R},\,0} &\lesssim 10^{-4}\,, \\
\Omega_{\mathrm{b},\,0} &= 0.049 \pm 0.003\,, \\
\Omega_{\mathrm{c},\,0} &= 0.259 \pm 0.002\,.
\end{aligned}
\tag{1.10}
$$

These numbers show that the contribution of DM to the mass of the Universe ($\Omega_{\mathrm{c},\,0}$) to $\Omega_{\mathrm{M},\,0}$ is 84%.

In this chapter we have seen evidence for DM on many different scales. The results of the Planck collaboration have allowed us to estimate the total contribution of baryonic matter to the mass of the Universe, which is only one

sixth. We are not missing a piece of the puzzle, rather, we only have one piece while the puzzle is missing. In this thesis, we will discuss the search for this missing puzzle. The next chapter (Chapter 2) briefly discusses the properties of DM. We describe the XENONnT DM detection experiment in the following chapter (Chapter 3). In Chapter 4 we discuss how several search strategies and detector designs could be used to reconstruct the properties of DM when it is detected. Chapter 5 discusses how we acquire- and process the data of XENONnT, followed by an explanation of our signal reconstruction in software (Chapter 6). We conclude with the Science Run 0 WIMP results of XENONnT in Chapter 7.

Chapter 2

Dark Matter

Just as the evidence for Dark Matter (DM) is diverse, so too are the hypotheses of its constituents. Throughout this work, we will generally assume that DM consists of a type of particles referred to as Weakly Interacting Massive Particles (WIMPs), but alternative explanations also exist. If DM is explained by a single underlying cause - a particle - it must be long lived[1] because of its effects on the evolution of the Universe as we have seen in subsection 1.3.2. While it interacts with gravity, it does not have color or electrical charge[2]. Finally, DM is non-relativistic, and is therefore often referred to as cold DM, as relativistic particles would have washed out the structures we observe at large scales. One of the natural candidates in the Standard Model of particle physics (SM) would be neutrinos as they fulfill the above-mentioned requirements. However, the previously (subsection 1.3.1) discussed PLANCK result [35] shows that SM neutrinos cannot contribute sufficiently to the overall energy density to be DM because of their low mass. Other particles in the SM do not fulfill these requirements, so particle DM requires a beyond the SM theory. The existence of a yet unobserved particle has sparked a broad range of models that would predict a suitable candidate. While, in contrast, some theories predict that DM may be explained by something else than a particle.

To cover the range of models predicting DM, experiments are correspondingly diverse in their attempt to find it. These endeavors include telescopes, neutrino- and charged particle detectors, collider experiments, and direct DM

[1] Meaning that it must be stable on time scales similar to the age of the Universe $\sim$13.7 Gyr.

[2] More precisely, it must have a charge q of $q/e \lesssim 4 \times 10^{-7}$ for a DM-mass of $M_\chi \sim 1\,\mathrm{GeV/c^2}$ and not emit/absorb photons [36, 37].

detection experiments. In this chapter we focus on WIMPs and their detection, while also discussing different models and strategies to provide context.

2.1 Dark Matter candidates

2.1.1 Weakly Interacting Massive Particles

One of the historically most popular class of particles to constitute DM are Weakly Interacting Massive Particles (WIMPs) as they arise naturally from beyond the SM theories, such as super-symmetry. Additionally, a relatively simple mechanism of "freeze-out" produces the right relic abundance of DM as we will describe below. A WIMP is a DM candidate that has some interaction with SM-particles via a massive mediator and has a cross section and mass similar to the cross sections and masses typical for the weak interaction.

Freeze-out

The WIMP hypothesis allows obtaining the right relic abundance of DM via a relatively simple model of its creation. Less than one nanosecond after the big bang the Universe is still hot and dense and in a plasma phase. The WIMPs are in chemical and thermal equilibrium with the plasma. At that moment, the temperature of the Universe is so high that all particles in the plasma are relativistic, and the particle masses are negligible compared to their total energy. The number density of WIMPs only decreases with the expansion of the Universe. As the temperature of the Universe decreases due to its expansion, the chemical equilibrium is broken when the temperature becomes below the WIMP-mass. At that moment, SM particles have insufficient energy to create WIMPs, and therefore the number density of WIMPs decreases because DM self-annihilates. Later, the WIMP density becomes so low that also self-annihilation ceases. At this point, the WIMP density becomes "frozen-out" and the density only decreases due to the expanding Universe. The coincidence that a thermally averaged self-annihilation cross section of roughly $\langle \sigma v \rangle \approx 3 \times 10^{-26}$ cm^3/s and WIMP-mass of $\sim$100 GeV/c^2 gives the right relic abundance of DM is sometimes called the "WIMP-miracle" [8, 38].

The resultant DM density is directly related to the self-annihilation cross section. If the cross section is higher, the self-annihilation continues for longer, yielding a lower relic density. The assumption of a massive mediator (in the order of the weak scale) also dictates a lower bound on the mass of WIMPs (M_χ). If M_χ is too low, the thermal freeze-out process would start too early,

resulting in a too high DM relic abundance. This is called the Lee-Weinberg limit [39], which is a lower bound on $M_\chi \gtrsim 2$ GeV/c^2. Additionally, M_χ must be below $\mathcal{O}\left(100\,\text{TeV/c}^2\right)$ because of the unitarity of the total S-matrix of DM-particle - antiparticle scattering [40].

Significant portions of parameter space for the existence of WIMP-dark matter have been ruled out by direct detection experiments (section 2.3) leading to increased interest in other particle models, such as the axion(-like) particles discussed below. Yet, substantial portions of well-motivated parameter space remain unprobed [41] and the WIMP-search continues to be one of the prime candidates for DM searches [5, 42–47].

2.1.2 Axion and axion-like particles

Axions are a class of particles that were originally proposed to solve the apparent absence of Charge Parity (CP) violation in quantum chromodynamics [48, 49]. While a CP-violating term is allowed, it is not observed, a problem is known as the strong CP problem. The observation that the strong CP violation is absent or very small would be solvable if a new global symmetry were spontaneously broken, where the axion particle is the new pseudo-Goldstone boson. Often, a larger class of bosons is considered as DM candidates, known as axion-like particles (ALPs), which do not necessarily solve the strong CP problem, but their properties make them viable DM candidates [49, 50].

2.1.3 Alternative models

WIMPs, and, more recently, axions and ALPs, are contemporarily seen as the main DM candidates. However, because of the null results so far, the DM-community has developed more exotic models and names such as wimpzilla's, fimps, EWIMPS, asymmetric dark matter, and more. A complete overview of all these models is beyond the scope of this work, as we will primarily focus on WIMP-DM. For a summary of the many models, see for example [42, 51, 52].

In addition to particle-DM, other DM explanations have historically also been considered. One type of the earlier considered candidates are so called Massive Compact Halo Objects (MACHOs) [53, 54], dark and massive objects, such as planets or black holes. When these large objects pass in front of bright objects, they gravitationally lens the light from the bright object. The absence of observation of these lensing effects has largely ruled out MACHOs as the sole constituent of DM [55–57]. Yet, these kinds of searches remain important for DM detection, primarily as they might find a special class of black holes,

called primordial black holes. Primordial black holes are black holes that formed shortly after the Big Bang in the Universe. Although, primordial black holes are not able to explain the full DM puzzle [58], they would be an excellent seed for DM clustering. Moreover, the detection of a primordial black hole would be the nail on the coffin of the "classic" WIMP-theories, as these seeds would predict observable WIMP self-annihilation [59]. The absence of such self-annihilation signals means that either there are no stellar-mass primordial black holes - or - there are no WIMPs.

Another group of non-particle DM theories, commonly referred to as Modified Newtonian dynamics (MOND) [60–62], aim to explain the experimental evidence in terms of a modified theory of gravity [54]. This branch of theories started from the notion that the observations of the dynamics of stars and gas within galaxies could be reconciled with a modification of Newton's second law for low accelerations by modifying $F = ma$ to $F = ma^2/a_0$ for very small accelerations ($a \ll a_0 \sim 1.2 \times 10^{-10}$ m/s^2). This theory would alleviate the need of DM, but did lead to many issues to be embedded consistently in a realistic theoretical framework. Further models build on this (or start from a similar) postulate that a modification of Newtonian dynamics would solve the DM puzzle. While these models successfully explain one or two of the experimental results discussed in Chapter 1, they are not able to provide an explanation for all of them [54]; most prominently, their ability to match evidence of DM beyond the galaxy-rotation curves is low, such as the Bullet Cluster discussed in section 1.1.

2.2 Dark Matter particle detection

In the following sections we discuss three experimental approaches to detect DM.

- Indirect detection; self-annihilating DM may be indirectly observed through the production of SM-particles.
- Direct detection; DM may scatter off SM particles in Earth based detectors (section 2.3).
- Accelerator production; DM could be produced in high energy collisions at particle colliders such as the Large Hadron Collider (LHC).

The type of interactions searched for is illustrated in Fig. 2.1, where indirect detection (a) looks for the decay products of DM annihilation, direct detection

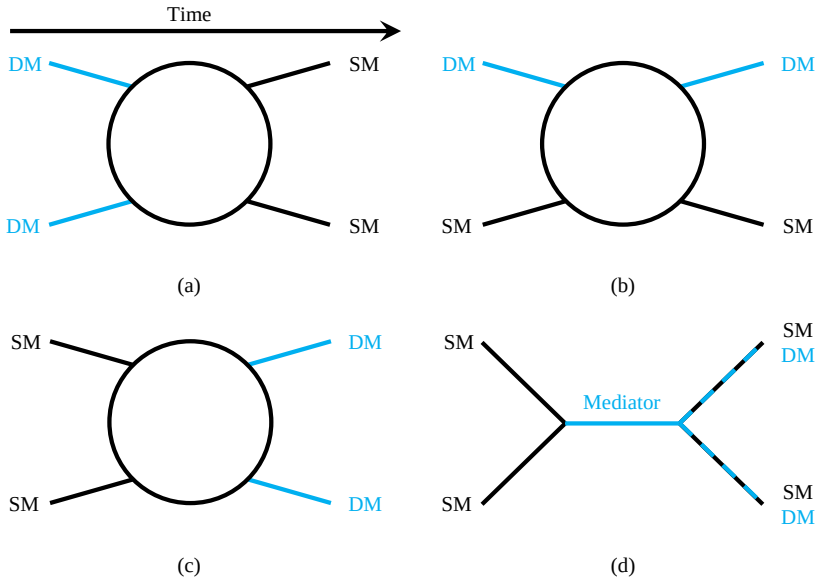

Figure 2.1: Particle DM detection techniques, relying on a form of interaction with SM particles. The first technique (a) depends on two DM particles annihilating to yield two SM particles. Many different interactions or higher order diagrams may play the role of the large circle indicated here, as the couplings between DM and SM particles are unknown. Direct detection (b) aims to detect DM through a scatter off it of a SM particle. At colliders, DM may be produced (c) by annihilating two SM particles, which is the inverse process of (a). Additionally, if there is a "dark mediator" (d) it could result in the annihilation of two SM particles that may decay either to two DM or two SM particles. Figure inspired by [63, 64], source file available at [18].

(b) looks for a scatter of DM off a SM particle, and finally collider production may create DM particles (c), or DM mediators (d) as explained below.

Indirect detection of DM searches for signatures of self-annihilating DM which may either be produced directly or via loop diagrams. Self-annihilation is one of the ingredients of the WIMP-paradigm, see section 2.1.1. However, as also explained earlier in this section, WIMPs have by now long frozen out, so DM annihilation should be a rare process. For DM-annihilation to take place, an

increased DM density would be required, such as at a galactic center (Eq. (1.2)), near (primordial) black holes (subsection 2.1.3), dwarf galaxies, or even celestial objects. Such self-annihilation could result (among other particles) in γ-ray emission which could reach telescopes and point back to the source. In recent years, Fermi-LAT detected an excess in γ-rays of $\mathcal{O}(10\,\mathrm{GeV})$, which was at first considered to be a potential signature of DM self-annihilation [65]. However, it has also been interpreted as an excess in the form of un-modeled backgrounds or point-like sources, such as a population of millisecond pulsars [66]. As discussed in Ref. [66], the latter explanation is currently favored over a DM interpretation of the excess. In addition to γ-rays, the decay products of DM self-annihilation may also be observed via X-rays, antiprotons, positrons, neutrinos, or other particles, see for example Ref. [67].

Experiments at particle colliders could detect DM through several mechanisms. When DM is produced at, for example, the Large Hadron Collider (LHC) an experiment (like ATLAS [68] or CMS [69]) may detect it as missing transverse momentum. This missing momentum could be the result of an interaction like in Fig. 2.1 (c / d), where two DM particles are created that can then escape the detector undetected. The missing transverse momentum would be reconstructed as a non-zero sum of transverse momenta of outgoing particles during an event. Alternatively, collider experiments search for mediators between SM and DM-particles (Fig. 2.1d). If a mediator from the "dark sector" would be created, this could be observed as a resonance at the mass of the mediator.

One of the advantages of accelerator DM searches over other methods is that the searches are not affected by uncertainties in astrophysical parameters, such as the local DM density (section 1.2). On the other hand, if a new particle is found by missing transverse momentum, it does not necessarily have to be a DM-particle. Any particle that is sufficiently long-lived to escape the experiment can result in missing transverse momentum. It does not mean that it is stable on the timescale of the Universe as required for DM, only that it is stable for the time it traverses the experiment. Furthermore, a disadvantage of the collider-based DM-searches is that the possible signatures and backgrounds greatly depend on the specific DM-model assumed. For an overview of DM searches at the LHC see Ref. [70].

2.3 Direct detection

Earth-based DM detection experiments aim to discover DM by measuring a scatter of a DM particle off a target material. This scatter can be off an atom,

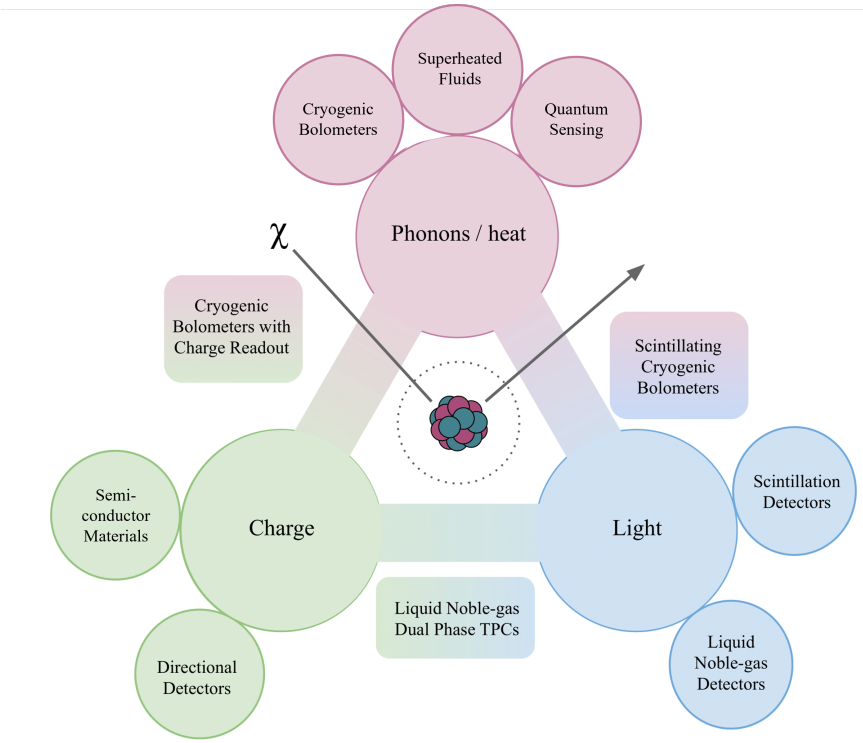

Figure 2.2: Signals of energy deposited by DM (χ) scattering off a given target such as an atom, nucleus, and/or electron. The large circles (charge, light, phonons/ heat) describe the possible signals that arise from scatters, while the smaller circles illustrate the techniques to detect these signals. The rectangular boxes are detection techniques that are able to detect two of these signals. For example, the blue-green box represents experiments like XENONnT (next chapter), LZ [5], and Panda-X [71] and the pink-green box experiments like Edelweiss [72], and SuperCMDS SNOLAB (Chapter 4). A relatively new, but interesting technique is using quantum sensing techniques like super conducting nanowires to detect minuscule energy deposits [73–75]. Figure adapted from [63].

nucleus, and/or electron depending on the experiment. It is usually performed in a controlled environment where the target is monitored to search for an excess

of interactions above the background of SM processes. A scatter of DM would deposit a (small) amount of energy in the target material. The general strategy is to detect this energy deposit via heat, ionization, or scintillation in the target material while shielding the detector from as many sources of background as possible. In Fig. 2.2, several direct detection approaches are illustrated.

For most direct detection experiments, the main backgrounds (γ, β, α-radiation or neutrinos) manifest themselves as recoils to the atomic electrons. However, DM may couple to electrons as well as to atomic nuclei (although scattering off electrons is strongly suppressed for typical $\mathcal{O}\left(10 - 100\,\mathrm{GeV}/c^2\right)$ WIMP-masses due to kinematic considerations). As these backgrounds and DM may scatter in different ways off the target, distinguishing between scatters off the nucleus (nuclear recoils) and bound electrons (electronic recoils) often provides a powerful tool to reduce these backgrounds. Other backgrounds require different techniques to be distinguished from DM. For example, neutrons can scatter off the nucleus in the same way as DM with the notable difference that neutrons may scatter multiple times, both inside and outside the detector. Neutron backgrounds can be reduced by detecting and actively vetoing such multiple scatters. Neutrinos can be another source of nuclear recoils, which will be very hard to distinguish from WIMP scatters off the nucleus. They could only be distinguished from DM if the directionality of the incoming particle could be reconstructed, on which subject interesting R&D is ongoing [74, 76]. The neutrinos are not a dominant background yet, but future detectors are expected to be limited by this background [77].

To discriminate between electronic recoil (ER) and nuclear recoil (NR), experiments often use more than one of the detection techniques depicted in Fig. 2.2. For example, a dual phase xenon Time Projection Chamber (TPC) such as XENONnT [44] uses both scintillation and ionization signals, while a cryogenic bolometer using semiconductors like SuperCDMS SNOLAB [45] employs both charge and phonon detectors. Both experiments are extensively discussed in Chapter 3 and Chapter 4. It must be noted that not all experiments require two detection methods to discriminate between NR and ER. Notably, liquid argon DM detection experiments like DEAP [78] can distinguish ER and NR interactions based on the shape of the scintillation signal, as will be discussed in the context of xenon in section 3.1.

2.3.1 Dark Matter-scattering

Assuming that WIMPs are the particles that make up the DM halo (in our Milky Way) as in section 1.2, they can scatter off a target material and deposit

energy, like in Fig. 2.1(b). As discussed above, we limit ourselves to nuclear recoils that result from these energy deposits. For an Earth based detector, the differential nuclear recoil rate $dR/dE_{\rm nr}$ (measured in number of events per energy × mass × time) is given by:

$$\frac{dR}{dE_{\rm nr}}\left(E_{\rm nr}\right) = \frac{\rho_0}{M_\chi M_N} \int_{v_{\rm min}}^{v_{\rm max}} d^3\vec{v} v F\left(\vec{v}+\vec{v}_e\right) \frac{d\sigma_{\chi-N}}{dE_{\rm nr}}(v, E_{\rm nr}, A)\,, \tag{2.1}$$

where $E_{\rm nr}$ is the nuclear recoil energy, typically in the order of $\mathcal{O}(1-10)$ $\mathrm{keV_{nr}}$[3] for a ~ 10 GeV/c^2 WIMP, ρ_0 the local DM density, M_N the mass of the target nucleus, M_χ the WIMP mass, $\sigma_{\chi-N}$ the WIMP nucleus cross section, $\vec{v}$ the WIMP velocity in the detector's rest frame, $\vec{v}_e$ the Earth's velocity with respect to the galactic rest frame, and $F(\vec{v})$ the WIMP velocity distribution in the galactic rest frame. We make explicit that the differential WIMP-nucleus cross section $d\sigma_{\chi-N}/dE_{\rm nr}$ depends on v and $E_{\rm nr}$, as well as the atomic number of the target material A, as we will explain below. The integral is performed from $v_{\rm min}$, the minimum WIMP velocity required to generate an NR of energy $E_{\rm nr}$, to $v_{\rm max}$ which is given by the sum of the DM escape velocity ($v_{\rm esc}$) and v_e. For a non-relativistic collision - the right regime for cold-DM WIMPs where $v/c \sim 10^{-3}$ - the nuclear recoil energy is kinematically fixed as,

$$E_{\rm nr}\left(M_\chi, c, A, \theta\right) = \frac{\mu_N^2 v^2\left(1-\cos\theta\right)}{M_N}\,, \tag{2.2}$$

where μ_N is the reduced mass $\mu_N = M_\chi M_N/(M_\chi + M_N)$ and θ the scattering angle in the center-of-mass frame. For $\theta = \pi$, we obtain the lowest velocity for which a recoil of $E_{\rm nr}$ is created,

$$v_{\rm min}\left(E_{\rm nr}, M_\chi, A\right) = \sqrt{\frac{M_N E_{\rm nr}}{2\mu_N^2}}\,. \tag{2.3}$$

The particle physics is encapsulated in the differential WIMP-nucleus cross section. This cross section can both contain spin-dependent (S.D.) and spin-independent (S.I.) terms,

$$\frac{d\sigma_{\chi-N}}{dE_{\rm nr}}(v, E_{\rm nr}, A) = \frac{M_N}{2v^2\mu_N^2}\left(\sigma_0^{\rm S.I.}\mathcal{F}_{\rm S.I.}^2(E_{\rm nr}) + \sigma_0^{\rm S.D.}\mathcal{F}_{\rm S.D.}^2(E_{\rm nr})\right), \tag{2.4}$$

[3]Where the subscript in $\mathrm{keV_{nr}}$ indicates that this is the nuclear recoil energy as discussed in section 3.1.

where $\sigma_0^{\text{S.I.}}$ and $\sigma_0^{\text{S.D.}}$ are the spin-(in)dependent cross sections at zero momentum transfer, and $\mathcal{F}_{\text{S.I.}}$ and $\mathcal{F}_{\text{S.D.}}$ the nuclear form factors that encapsulate the energy dependence of the interaction and thereby the nuclear physics. For a low energy interaction, a particle will coherently interact with all the nucleons in the nucleus. An interaction is coherent if the transferred momentum (and the corresponding de Broglie wavelength) is small relative to the size of the xenon nucleus. But at high momentum transfer, the coherence of the interaction is lost.

The WIMP may couple differently to protons and neutrons (with strengths f_p and f_n respectively), such that the spin-independent zero momentum transfer cross section may be written as,

$$\sigma_0^{\text{S.I.}} = \frac{4\mu_N^2}{\pi} \left(Z f_p + (A - Z) f_n\right)^2 , \tag{2.5}$$

But often, it is approximated that $f_n \sim f_p$ such that we may obtain,

$$\sigma_0^{\text{S.I.}} \approx \sigma_{n,\,\text{S.I.}} \frac{\mu_N^2}{\mu_n^2} A^2 , \tag{2.6}$$

where $\sigma_{n,\,\text{S.I.}} = \frac{4}{\pi}\mu_n^2 f_n^2$ is the spin-independent WIMP-nucleon cross section denoted as simply $\sigma_{\text{S.I.}}$, and μ_n the DM-nucleon reduced mass. In contrast to $\sigma_{\chi-N}$, which scales as A^2, $\sigma_{\text{S.I.}}$ does not depend on the target nucleus, and is therefore often used to compare between experiments.

Similarly, the spin-dependent WIMP-proton cross section can be obtained,

$$\sigma_0^{\text{S.D.}} \approx \sigma_{p,\,\text{S.D.}} \frac{4\pi\mu_N^2}{3\mu_p^2(2J+1)} S_{A,\,p}\left(E_{\text{nr}} = 0\right) , \tag{2.7}$$

where J is the total nuclear spin, $S_{A,p}$ is the structure function which is proportional to $\mathcal{F}^2_{\text{S.D.}}$, and μ_p the DM-proton reduced mass [79]. The spin-dependent WIMP-neutron cross section can be obtained by replacing $p \to n$ in Eq. (2.7), although $S_{A,p}$ and $S_{A,n}$ are notably different. Additionally, not all isotopes of the target material contribute, for example for xenon, only ^{129}Xe (spin 3/2 and natural abundance of 26%) and ^{131}Xe (spin 1/2 and natural abundance of 21%) contribute significantly [79]. We will focus primarily on spin-independent coupling of WIMPs to SM as XENONnT is most sensitive to such interactions.

2.3.2 Astrophysical inputs

The velocity distribution $F(\vec{v})$ is usually assumed to described by the Standard Halo Model (SHM) [80], where it is a Maxwellian velocity distribution. In the

SHM, the local $F(\vec{v})$ is given as a function of the galactic velocity $\vec{v}_g$, where $\vec{v}_g = \vec{v} + \vec{v}_e$ is the sum of the velocity $\vec{v}$ and Earth's velocity of $\vec{v}_e$ with respect to the galactic rest frame,

$$F\left(\vec{v}_g\right) \propto \frac{1}{v_0^3} e^{-\left(|\vec{v}_g|/v_0\right)^2}\,, \tag{2.8}$$

where v_0 is the most probable speed of the dark matter particles also known as the local standard of rest. However, gravitationally bound WIMPs cannot have arbitrarily large velocities, since they would escape the Milky Way if the velocity is too high. To account for this, Eq. (2.8) is cut off for velocities greater than the escape velocity v_{esc} in the galactic rest frame,

$$F\left(\vec{v}_g\right) \propto \begin{cases} \dfrac{1}{v_0^3} e^{-\left(|\vec{v}_g|/v_0\right)^2} & |\vec{v}_g| < v_{\text{esc}} + |\vec{v}_e| \\ 0 & |\vec{v}_g| > v_{\text{esc}} + |\vec{v}_e|\,, \end{cases} \tag{2.9}$$

where we explicitly added $\vec{v}_e$ to compare $\vec{v}_g$ in Earth's rest frame. Earth's velocity relative to the galactic rest frame relates to the velocity with respect to the local standard of rest, the peculiar velocity ($\vec{v}_{\text{pec}}$) of the Sun with respect to v_0 and Earth's velocity ($\vec{v}_{\text{Earth-Sun}}$) via

$$\vec{v}_e = \vec{v}_0 + \vec{v}_{\text{pec}} + \vec{v}_{\text{Earth-Sun}}\,. \tag{2.10}$$

In the case of non-directional detectors, we can simplify Eq. (2.9) using the DM speed distribution $f(v) = 4\pi v^2 F(v)$ and ignoring annual modulation effects due to the Earth's orbit around the Sun. The values of ρ_0, v_0, $\vec{v}_{\text{pec}}$, $\vec{v}_{\text{Earth-Sun}}$ and v_{esc} determine how many WIMPs signals can be detected in a given detector. From now on, we drop the explicit vector notation for the velocities and use the speeds instead, as we focus on non-directional experiments.

For comparing similar experiments like XENON1T [4] and LZ [81] identical benchmark values are used for the calculation of the WIMP induced recoil rates. However, until recently [31], these values were often based on outdated observations. In Tab. 2.1 we compare the values of the SHM that were used until recently (2021) [4] to modern observations. Ref. [31] aims to homogenize the assumed parameters for DM experiments to use modern values, which the community is starting to use [5]. For $\vec{v}_{\text{pec}}$ and $v_{\text{Earth-Sun}}$ there is no (significant) difference between the values that were historically considered and are currently advised in Tab. 2.1. Therefore, we from now on use the values as advised for these two parameters without further discussing them.

Parameter [unit]	Historical	Modern	Advised [31]
v_0 [km/s]	220 [82–84]	238.0 ± 1.5 [85, 86]	238.0
v_{esc} [km/s]	544 [87]	528^{+24}_{-25} [88]	544
ρ_0 [GeV/cm^3]	0.3 [89]	$0.3 - 0.6$ [32]	0.3
$v_{\text{Earth-Sun}}$ [km/s]	29.79 [89]	29.79 [89, 90]	29.8
$\vec{v}_{\text{pec}}$ [km/s]	(11, 12, 7)	$\left(11.1^{+0.7}_{-0.7}, 12.2^{+0.5}_{-0.5}, 7.3^{+0.4}_{-0.4}\right)$ [91]	$(11.1, 12.2, 7.3)$

Table 2.1: Values of parameters in the SHM. The second column shows the parameters as historically used. The third column shows more recent observations of the respective parameters. The fourth column shows the recommended values as motivated by [31]. As also listed in Ref. [31], multiple groups found different values of v_{esc} based on the same data [28, 29], therefore the recommended value of v_{esc} does not follow the modern value. Similarly, ρ_0 is not updated to a more modern value as significant differences exists between its observations [32] and it only serves as an overall scaling factor of the expected WIMP-rate, see Eq. (2.1). The peculiar velocity ($\vec{v}_{\text{pec}}$) has three vector components as (v_r, v_ϕ, v_θ) with r pointing radially inward, ϕ in the direction of the Milky Way's rotation, and θ the orthogonal angle.

2.3.3 Expected recoil rates

Combining Eq. (2.1) and Eq. (2.9) allows us to calculate the expected nuclear recoil rates for a given WIMP-model and a given experiment. In this section we compare the effects of changing the values as in Tab. 2.1 and the expected recoil rates for different target materials. To this end, we will illustrate three things:

1. the effect of changing astrophysical parameters on the velocity distribution, Eq. (2.9),
2. how that affects the lower (v_{min}) and upper (v_{esc}) integration bounds for the nuclear recoil rate, Eq. (2.1),
3. the effects of astrophysical parameters on the exclusion limit for a liquid xenon DM detector like LZ [5].

In Fig. 2.3, we illustrate the change in $f(v)$ under different assumptions of the parameters in the SHM. A change in v_0 from 220 km/s to 238 km/s, shifts $f(v)$ to higher speeds. In contrast, the change of v_{esc} from 544 km/s to 528 km/s decreases the integrated phase only at the tail of $f(v)$. This affects different experiments differently, as $f(v)$ is integrated over in Equation 2.1. To

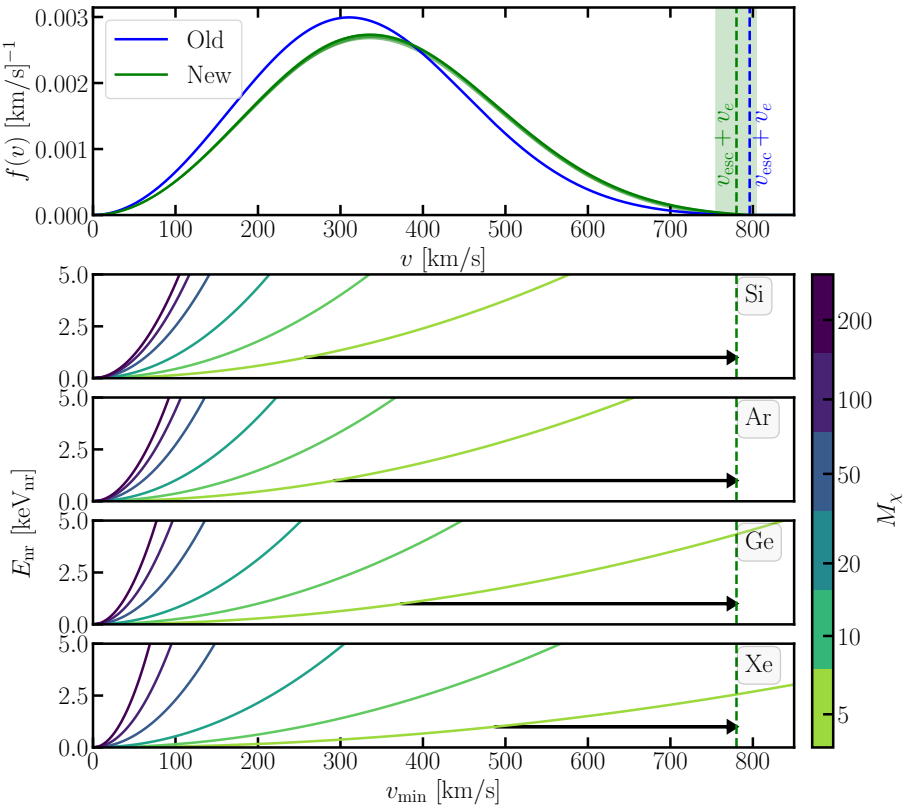

Figure 2.3: Local velocity distribution $f(v)$ for different SHM parameters (top panel). "Old" parameters refer to $v_{\rm esc}$ of 544 km/s and v_0 of 220 km/s, "New" refers to 528 km/s and 238 km/s, respectively. Other parameters are the same (following [31], rightmost column Tab. 2.1). The effect on the nuclear recoil rates for several WIMP-masses (M_χ) can be understood by plotting the lower integration bound of Eq. (2.1), $v_{\rm min}$ for several targets (bottom panels) for a given nuclear recoil energy ($E_{\rm nr}$). For example, the arrow indicates the parameter space $v_{\rm min}$ to $v_{\rm esc}$ for a 5 GeV/c^2 WIMP resulting in a 1 keV$_{\rm nr}$ recoil for each target material. Source code available at [18].

illustrate this, we show the lower integration bound, $v_{\rm min}$ for several target

materials and masses. The upper integration bound is given by $v_{\rm esc} + v_e$. As an example, for a 5 GeV/c^2 WIMP, inducing a 1 keV$_{\rm nr}$ recoil, the $v_{\rm min}$ is roughly 490 km/s for xenon (Xe) and 260 km/s for silicon (Si). This has a large effect on the expected recoil rates at these energies, as for Si, a much larger range of $f(v)$ will be integrated than for Xe. We also see that the shift from the old values to the new values has a larger effect for expected recoil rates of low-mass WIMPs than for high-mass WIMPs where most of the phase space of $f(v)$ will be integrated - irrespective of the exact shape of $f(v)$.

In Fig. 2.4, we evaluate differential nuclear recoil rates (Eq. (2.1)) for spin-independent coupling of WIMPs for the same materials as in Fig. 2.3. For this figure we use $\rho_0 = 0.3$ GeV/c^2, $v_0 = 238$ km/s and $v_{\rm esc} = 544$ km/s (the advised column in Tab. 2.1). Fig. 2.4 shows how the dependence on A, Z and M_N propagate to the nuclear recoil rate for a WIMP with a spin-independent cross section of $\sigma_n = 10^{-47}$cm^2. Heavier elements give more nuclear recoils because of the A^2 dependence in Eq. (2.6). For smaller WIMP-masses, the heavier target elements yield smaller rates at $\gtrsim \mathcal{O}(1)$ keV$_{\rm nr}$ as $v_{\rm min}$ becomes so large that little phase space up to $v_{\rm esc}$ remains. Although not considered in Fig. 2.4, detectors employing semiconductor materials like Si and Ge may additionally benefit from lower energy thresholds compared to Xe and Ar-based detectors, which typically have energy thresholds of $\gtrsim \mathcal{O}(1)$ keV$_{\rm nr}$. On the other hand, it is relatively easier to accumulate a large target mass (multi-ton scale) and to mitigate backgrounds for Xe and Ar detectors. Precisely for this trade-off, we compare a Ge (and Si) based detector (SuperCDMS) and a Xe based detector (XENONnT) in depth in Chapter 4.

Using Fig. 2.3 and Fig. 2.4 we can now probe how the parameters in Tab. 2.1 affect the final results of a DM experiment like LZ [5]. At the time of writing - no DM has been found. Therefore, the results of experiments are in the form of exclusion limits, where portions of WIMP cross section and mass parameters space are excluded. To calculate a limit, it is required to integrate Eq. (2.1) to yield the total rate for a given WIMP-model. For this, a spectrum as in Fig. 2.4 needs to be convolved with the detector response to account for the detector resolution and threshold. To get the total number of events, the rate is multiplied by the exposure, which is the product of the target mass and the observation time. The expected signal and background for a given WIMP-model, allows the calculation of the median sensitivity. The median sensitivity is where a limit would be put 50% of the time if an experiment would be repeated many times.

Below, we will show the effect of changing the astrophysical parameters on the sensitivity by using the combined efficiency of LZ [5] but neglecting the

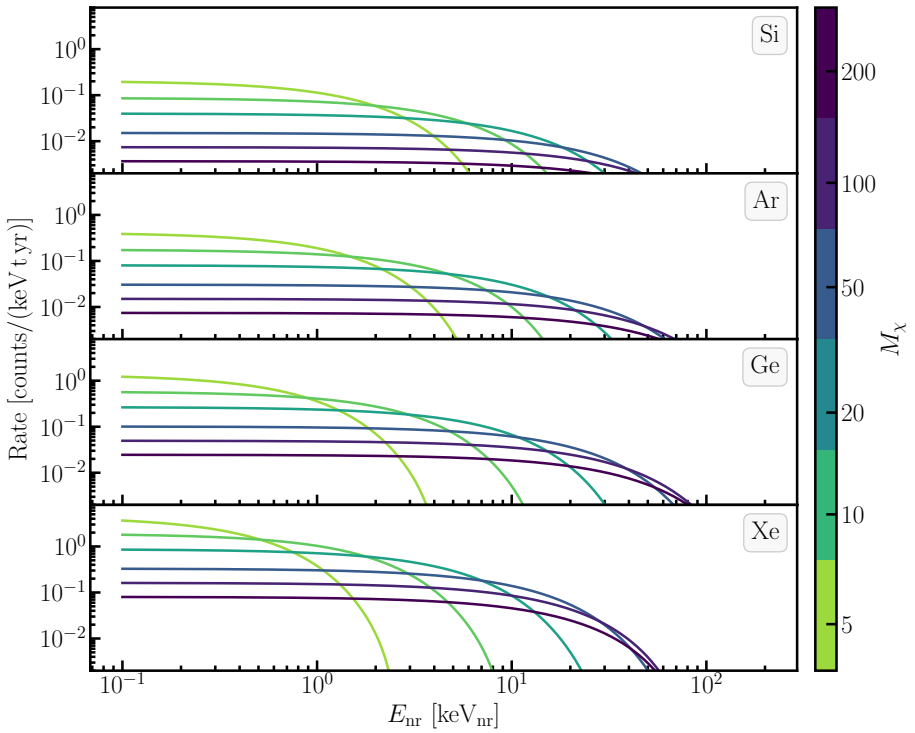

Figure 2.4: Expected nuclear recoil rates for WIMPs with a spin-independent cross section of $\sigma_n = 10^{-47}$ cm^2 and given masses. These recoil spectra do not account for detector effects like energy resolution and thresholds. For these spectra, we use $\rho_0 = 0.3$ GeV/c^2, $v_0 = 238$ km/s and $v_{\mathrm{esc}} = 544$ km/s. Source code available at [18].

energy resolution of the experiment. The combined efficiency of LZ is shown in Fig. 2.5 and incorporates the efficiencies from the reconstruction (similar to section 6.4), data-quality cuts, and the energy Region of interest (ROI). The ROI is motivated by the expected extent of the nuclear recoil spectrum (Fig. 2.4) and therefore limited to roughly $\lesssim$60 keV$_{\mathrm{nr}}$.

In Fig. 2.6 we compare the sensitivity from LZ [5] to several hypothetical sensitivities with varying assumptions for the SHM. The published sensitivity considers the full background expectation and detector response in a likelihood

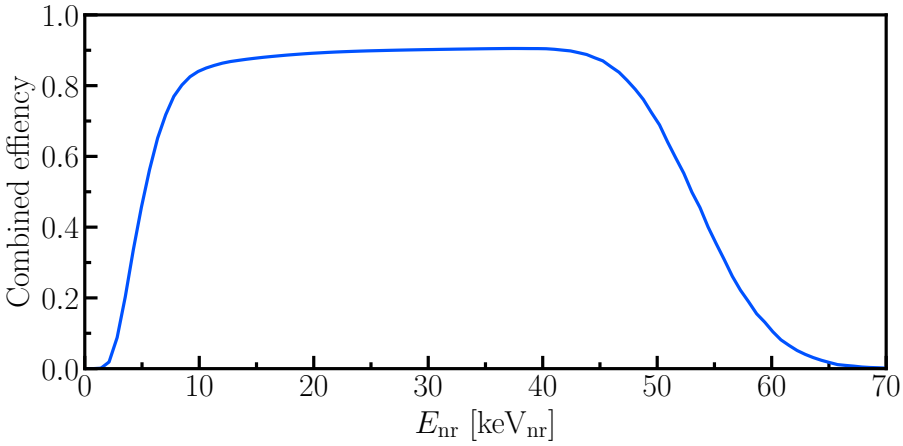

Figure 2.5: Combined efficiency of the LZ experiment. The combined efficiency incorporates the reconstruction efficiency (the most prominent contribution to the efficiency loss at low energies), data-quality cuts (intermediate energies), and ROI (cutoff at $\sim 60\,\mathrm{keV_{nr}}$). Figure based on public data from [5], source code available at [18].

analysis. It shows where, at 90% confidence level (C.L.), a limit would be put for the expected background. The yellow and green bands show the position of the sensitivity in upward/downward fluctuations at 1 σ and 2 σ respectively. In this figure we have not included the achieved limit [5] and only compare sensitivities.

To reproduce the sensitivity of LZ under varying halo parameters would require access to the full likelihood, which is not publicly available. Instead, a simpler, yet powerful method can be used to mimic the best-case scenario to obtain a sensitivity - the no-background assumption. In case of no-background, we can set a limit at each WIMP-mass at 90% C.L. by assuming that we have no background and measure no signal. This translates to evaluating the probability of finding 0 events in a Poisson counting experiment. This translates to solving:

$$P(n_{\mathrm{measured}}) = \mathrm{Poisson}_{\mathrm{cdf}}(n_{\mathrm{measured}},\, n_{\mathrm{expected}}) = 0.1\,, \tag{2.11}$$

which yields $n_{\mathrm{expected}} = 2.3$ for $n_{\mathrm{measured}} = 0$. In other words, for a DM model which expects 2.3 WIMP events, we would measure 0 events only 10% of the

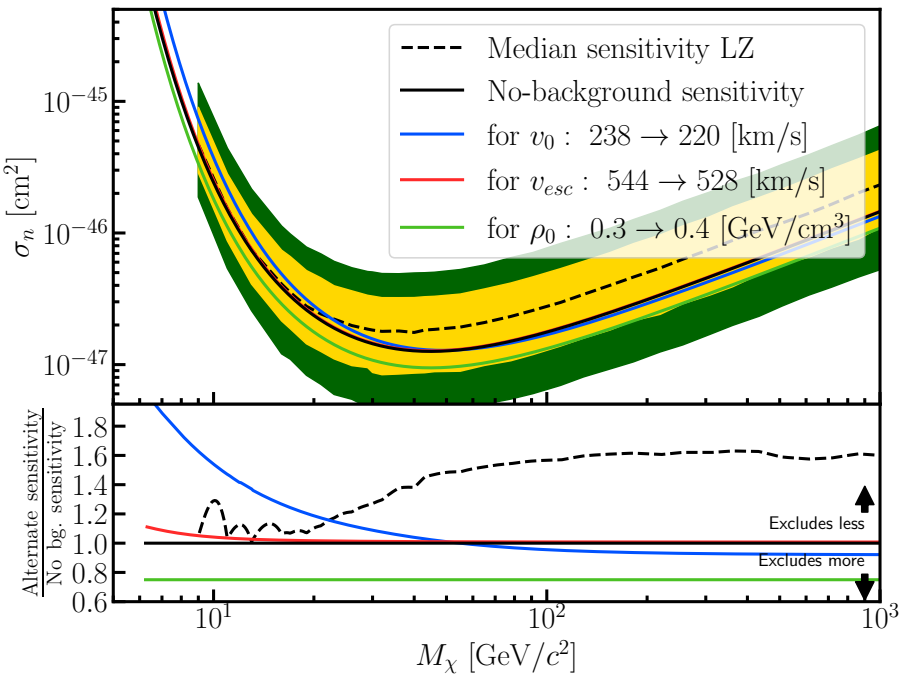

Figure 2.6: LZ sensitivity (dashed black line) with 1σ and 2σ uncertainty (yellow and green bands) [5]. Solid lines indicate the no-background sensitivities for several alternate SHM values. For the no-background sensitivity we used $\rho_0 = 0.3$ GeV/c^2, $v_0 = 238$ km/s and $v_{\text{esc}} = 544$ km/s. For the other sensitivities, we altered one of the parameters in the SHM. To compare different sensitivities, the bottom panel shows the alternate sensitivity divided by the no-background sensitivity. When v_0 is changed from 238 → 220 km/s (i.e., historical benchmark value in Tab. 2.1), it mostly affects the sensitivity at low WIMP-masses. If v_{esc} is changed from 544 → 528 km/s, which is the change from the historical benchmark value to the modern observation, little is changed in the sensitivity. The effect of changing $\rho_0 = 0.3 \rightarrow 0.4$ GeV/c^2 yields a proportionally lower exclusion sensitivity as the DM density is proportional to the cross section. Source code available at [18].

time, and therefore give 90% confidence that there must be 2.3 or fewer events

when no events are measured. Using this fact, we can find the σ_n for each M_χ where the number of events is 2.3. We find the 2.3 events by integrating the product of the WIMP-spectrum (Fig. 2.4) and the combined efficiency (Fig. 2.5). This yields the no-background sensitivity in Fig. 2.6. Despite neglecting the energy resolution, background expectations, and the statistically more powerful tools used in [5], the official sensitivity and no-background sensitivity agrees within $\sim 60\%$. The difference between the sensitivity and no-background sensitivity for $M_\chi \gtrsim 25$ GeV/c^2 is mostly governed by expected backgrounds in the LZ likelihood which are completely neglected in the no-background sensitivity.

Varying the astrophysical parameters in Tab. 2.1 has different effects on the sensitivity. Changing ρ_0 simply scales the sensitivity since it is proportional to the cross section. An increased value of ρ_0 by 25%, yields an equal decrease of the sensitivity by 25%. For this reason - in the absence of a detection - the exact value of ρ_0 does not matter so much if all experiments use the same value when comparing. Changing v_{esc} has a minute impact on the sensitivity at low WIMP-masses, as the available phase space for integration from v_{min} to v_{esc} is large (see Fig. 2.3). The modification of v_0 from 238 $\rightarrow$ 220 km/s has a large impact at low WIMP-masses with up to a factor ~ 2 for $M_\chi = 5$ GeV/c^2.

In this chapter, we have seen that there are many DM candidates, and detection techniques which thus far yielded null results. We also showed how many WIMP events are expected for a given experiment and how astrophysical parameters play a role in the search for WIMPs. In the next chapter (Chapter 3) we will focus on the XENONnT detector, primarily for detecting WIMPs.

Chapter 3

XENONnT

In this chapter, we discuss how XENONnT, a dual-phase xenon Time Projection Chamber (TPC), is designed to discover the Weakly Interacting Massive Particle (WIMP) dark matter particle. First, we explain how particles in general interact with liquid xenon and the detection principles of a dual-phase xenon TPC. Then we describe the XENONnT experiment and the design choices that, for example, aim to reduce backgrounds.

3.1 Energy deposits in xenon

When a particle scatters off a xenon atom (either off the xenon nucleus or off an electron) a fraction of its energy is transferred. The recoiling particle (nucleus or electron) will traverse the surrounding xenon, thereby creating a track of ionized and excited xenon atoms. If the scatter is off an electron, it creates an electronic recoil (ER) while a scatter off a nucleus results in a nuclear recoil (NR). As the nucleus is much heavier than an electron, NRs create shorter tracks with a higher energy density than ERs.

An energy deposit (ER or NR) excites and ionizes the surrounding xenon, which will result in scintillation, ionization, and heat. In an experiment like XENONnT, the scintillation and ionization signals are measured, while energy transferred to heat goes undetected. Below, we describe how several processes lead to ionization signals (freed electrons), scintillation signals (UV-photons), or are lost as heat. This is schematically shown in Fig. 3.1.

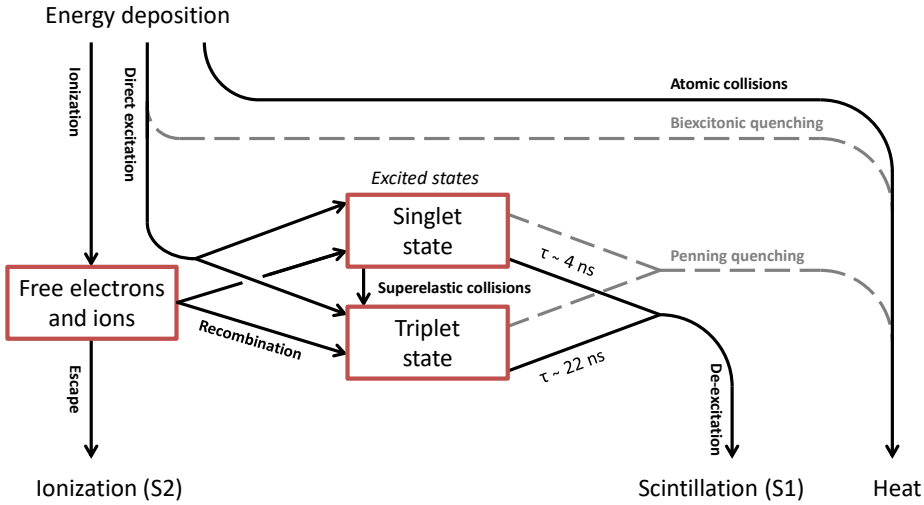

Figure 3.1: Energy depositions in liquid xenon that through several processes result in ionization, scintillation, and heat as described in the text. Figure from Ref. [92].

Excited and ionized xenon atoms emit scintillation light when de-exciting and recombining. Due to the meta-stable molecular states that emit the scintillation light, the wavelength of the scintillation light is different from atomic transitions, making the xenon transparent to its own scintillation light. Direct excitation (Fig. 3.1) yields scintillation light as follows. An excited xenon atom Xe^* quickly combines with a ground state Xe atom to form an excited dimer (excimer) Xe_2^*. The Xe_2^* is created in either a spin singlet or triplet state [93]. The excimer decays to two ground state Xe atoms by the emission of a 7 eV photon ($h\nu$) which corresponds to a wavelength of 178 nm:

$$Xe^* + Xe \rightarrow Xe_2^* \,, \tag{3.1}$$

$$Xe_2^* \rightarrow 2Xe + h\nu \,. \tag{3.2}$$

The decay times for the spin singlet and triplet states are 4 ns and 22 ns respectively [92]. The scintillation light of 178 nm is detected by photo-multiplier tubes (PMTs) that are custom designed by Hamamatsu. The XENONnT PMTs have a typical timing resolution of 9.2 ± 0.5 ns [94], while the Data Acquisition system (DAQ) digitizes the PMT signal every 10 ns. This prohibits a clear dis-

tinction between the different decay times, and thus the separation between the singlet and triplet state is marginal. Argon, on the other hand, has a similar scintillation mechanism but with vastly different decay times. The excited argon singlet and triplet states decay after ∼7 ns and ∼1600 ns [93]. This enables argon detectors to discriminate between ER and NR based on the pulse shape of the scintillation light, since ERs and NRs result in a different singlet to triplet ratio.

A second contribution to the prompt scintillation light is the recombination of ionized Xe^+ with a xenon atom to form Xe_2^+. The Xe_2^+ and a free electron create a doubly excited atom Xe^{**} and a ground state atom. The doubly excited xenon atom decays to Xe_2^*, which yields an UV-photon as in Eq. (3.2), and heat:

$$Xe^+ + Xe \rightarrow Xe_2^+ , \tag{3.3}$$

$$Xe_2^+ + e^- \rightarrow Xe^{**} + Xe , \tag{3.4}$$

$$Xe^{**} \rightarrow Xe^* + \text{heat} . \tag{3.5}$$

Those electrons that are freed during ionization are measured separately by applying a drift-field (which suppresses recombination) as will be discussed in the next section. When this drift-field is applied, the Xe_2^+ ions that do not recombine are drifted in the opposite direction from the electrons.

The third channel (Fig. 3.1) by which energy is deposited is through heat in atomic collisions. Elastic atomic collisions can result in a significant portion of the energy being lost in the form of heat. This form of energy loss is most relevant for NRs, as electrons only have 4.2×10^{-6} of the mass of a xenon atom, and so the elastic energy loss of ERs is insignificant.

Two additional processes (indicated in gray in Fig. 3.1) result in heat. First, biexcitonic quenching is when excitons Xe^* recombine before forming an excimer:

$$Xe^* + Xe^* \rightarrow Xe^+ + Xe + e^- . \tag{3.6}$$

Second, penning quenching [92] occurs when two excimers (Xe_2^*) collide, in which process one excimer dissociates to the ground state:

$$Xe_2^* + Xe_2^* \rightarrow Xe_2^* + 2Xe + \text{heat} . \tag{3.7}$$

These processes are important for NRs that deposit their energy in a relatively small volume. As a result, an NR will deposit more energy as heat, causing a smaller ionization and scintillation signal compared to an ER signal with the same energy.

The last transition in Fig. 3.1 is that from the singlet state to the triplet state of an excimer caused by "supereleastic" collisions of the singlet state with a free electron. This process is more likely to occur if recombination is slow because it requires a free electron, such as for low energy transfer tracks [92].

The relative strength of scintillation signals and ionization signals is different for NR compared to ER. This difference is employed to separate NR from ER in XENONnT. Despite this powerful tool - misidentification of ERs as NRs, known as "ER-leakage", occurs (section 7.3), which warrants also decreasing ER backgrounds. The energy deposited in the form of ionization is measured by extracting the electrons as we will explain shortly in section 3.2, while the energy in scintillation light is measured in the form of created UV-photons. The exact interplay of the processes described above is energy-dependent as well as dependent on the drift field that is applied. Furthermore, the effects are partially degenerate and therefore experiments use semi-empirical models like NEST [95, 96] to describe the detector response to energy deposits.

3.1.1 Electronic recoil

Electronic recoils (ERs) are the group of interactions where a particle collision transfers energy to an atomic electron. The following particles/radiation scattering off electrons are relevant for the XENONnT experiment: γ-rays, X-rays, β-radiation, α-radiation, neutrinos, and WIMPs (although very kinematically suppressed, see section 2.3). For XENONnT [97], the dominant sources of background are radioactive impurities. For low energies, $\mathcal{O}(< 140\,\mathrm{keV})$ where they might be reconstructed as NRs, these are specifically, the β-emitter ^{214}Pb, the second-order weak processes of the $2\nu\beta\beta$ decay of ^{136}Xe, the double-electron capture decay of ^{124}Xe, and to a lesser extent ^{85}Kr, $^{83\mathrm{m}}$Kr, ^{133}Xe, and solar ν. The scattering processes differ for the physical processes that are ER but they yield similar signatures in XENONnT.

For ERs, we can neglect elastic energy loss and biexcitonic quenching, such that we obtain the energy E of the interaction:

$$E = W(n_e + n_{\mathrm{ph}})\,, \tag{3.8}$$

where the energy is proportional to the sum of the number of electrons (n_e) and to the number of photons (n_{ph}) via the average energy required to produce a quantum (of charge or light) $W = 13.7$ eV/quanta. Energies obtained via Eq. (3.8) are called ER-equivalent energies and use the notation $\mathrm{keV_{ee}}$.

3.1.2 Nuclear recoil

In XENONnT coherent elastic neutrino nucleus scattering (CEνNS), WIMPs, and neutrons can cause nuclear recoils (NRs). CEνNS is the scattering of Solar, atmospheric, or supernova neutrinos off a xenon-nucleus via a Z-boson exchange. At the time of writing, CEνNS has not been detected (yet) in XENONnT, but it has been by other experiments [98–100]. CEνNS interactions are expected to be the ultimate background for XENONnT's successor XLZD (formerly Darwin) [42] as its signature is the same as for WIMPs. Therefore, the cross sections where CEνNS are expected to put a constraint onto the parameter space that can be probed by direct detection experiments and is referred to as the "neutrino floor" or "neutrino fog" [101].

Neutrons, as a background for XENONnT, are produced from spallation by cosmic muons and radioactivity. They are distinguished from WIMPs if they are detected to scatter multiple times either within, or outside the liquid xenon (LXe) volume, which is likely since the mean free path of a neutron in LXe ($3\,\mathrm{g/cm^3}$) is $\sim$10 cm (for a neutron with a kinetic energy of $1 - 10^4$ keV). In contrast, the cross section of WIMPs are so low that the chance of having a WIMP double scatter is effectively zero. Furthermore, neutrons provide the only NR calibration source available for XENONnT.

Similar to ER-equivalent energies, NR-equivalent energies are obtained using an adapted version of Eq. (3.8) and are expressed in units of $\mathrm{keV_{nr}}$. This is the topic of section 4.6.1 and NR-equivalent energies are obtained by dividing the right-hand side of Eq. (3.8) by an energy-dependent Lindhard factor $L \sim 0.2 - 0.3$ (for $E_{\mathrm{nr}} \sim 5 - 100\,\mathrm{keV_{nr}}$) to take the energy losses in terms of heat for NRs into account (Eq. 4.15). Since $L < 1$, measuring the number of photons and electrons results in a larger value of E for NRs. That is equivalent to stating that for the same ER and NR energies, more photons and electrons are generated for ERs than for NRs because of quenching.

3.2 Dual-phase Time Projection Chamber

At the heart of the XENONnT experiment is a dual-phase Time Projection Chamber (TPC). The goal of the TPC is to extract the ionization and scintillation signals from the liquid xenon (LXe) as discussed in the previous sections. The working principle of a TPC is explained in Fig. 3.2, which we will discuss together with its implementation in the XENONnT experiment [44].

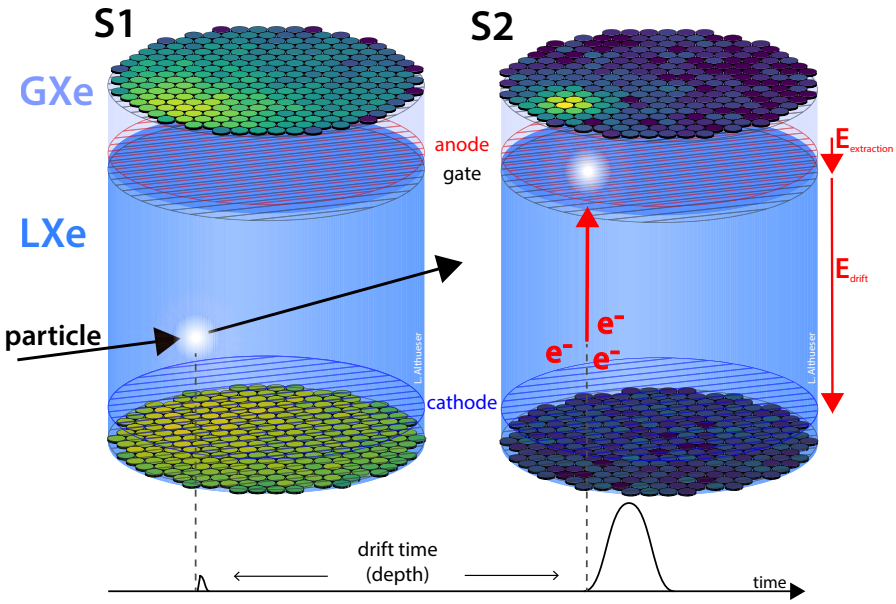

Figure 3.2: Working principle of a liquid xenon (LXe) TPC. The LXe target is continuously monitored by two arrays of PMTs at the top and the bottom of the detector. When a particle interacts within the LXe, it will create a first prompt scintillation signal (an S1). Due to the ionization induced by the interaction, electrons are liberated, and a cloud of freed electrons is drifted upwards by an applied electric drift field E_{drift}. Near the liquid-gas interface, a stronger extraction field $E_{\text{extraction}}$ is applied, that accelerates the electrons which collide with the gaseous xenon atoms, in order to produce secondary proportional scintillation light (the S2). The time difference between the S1 and the S2 is called the drift time, a measure for the depth of the interaction. This, together with the xy-reconstruction based on the hit pattern on the PMTs allows for full 3D reconstruction of the interaction position. Credits L. Althüser.

The XENONnT TPC measures 1.3 m in diameter and 1.5 m in height. Inside the TPC, the main volume is filled with LXe with a layer of gaseous xenon (GXe) on top. The LXe serves as the target where the scintillation and ionization signals originate. The scintillation light is observed by PMTs, both at the top

and bottom of the TPC, this signal is called the *S1*. To measure the ionization, liberated electrons are drifted to the top of the TPC by the application of a drift field (E_{drift}) between the cathode and gate electrodes. Just above the gate is the liquid-gas interface, above which another electrode, the anode, creates a stronger field ($E_{\text{extraction}}$). The $E_{\text{extraction}}$ extracts the electrons from the liquid and accelerates the electrons. The accelerated electrons gain enough energy to excite xenon atoms in the gas, and the resulting in secondary scintillation light is also detected by the PMTs. This second signal is called the *S2*.

With the S1 and S2, the depth of an interaction is reconstructed based on the time difference between the S1 and S2 signals for a given E_{drift}. The time between the signals is determined by the depth of an interaction since the electrons need to be drifted to the liquid-gas interface (up to 2.2 ms for XENONnT where $E_{\text{drift}} = 23$ V/cm), while the S1 is emitted almost instantaneously (with decay times of 4 and 22 ns for singlet and triplet states). Therefore, shallow interactions have a short time difference between S1 and S2, while interactions deep in the detector have a larger time difference between the two. The amount of the S2 signal observed per PMT in the top array (the hit pattern) allows reconstructing the lateral position of the interaction. Combined, the xyz-position of the interaction are reconstructed - which is essential for background reduction and correcting for position dependent effects.

3.3 The experiment

The XENONnT experiment (Fig. 3.3) is located in Laboratori Nazionali del Gran Sasso (LNGS) in central Italy and is the latest of a series of successively larger experiments operated by the XENON collaboration. In order to be sensitive for unexplored WIMP parameter space, the experiment needs to achieve:

- As large as possible target mass, since more target mass means a higher probability of seeing an interaction.
- As low as possible backgrounds to be sensitive to the few signals that at first would be observed for DM.

At the center of XENONnT is the TPC as discussed above, which is housed in a double-walled stainless-steel cryostat. The inner cryostat is filled with 8.6 t of xenon, of which 5.9 t is within the TPC. The remainder is outside the TPC and mainly serves as a passive shield from external radiation, such as from the cryostat itself. The walls of the TPC are made of Polytetrafluoroethylene

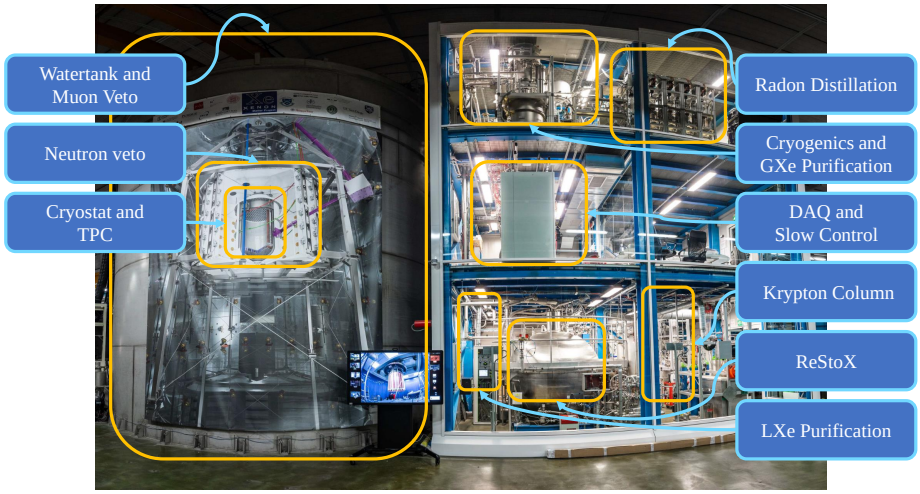

Figure 3.3: Picture of the XENONnT experiment underground at LNGS. The left shows the water tank on which a banner hangs that shows the TPC, Neutron Veto (NV) and Muon Veto (MV). The right half shows the adjacent support building, where the support systems are housed as described in the text. Credits picture H. Schulze Eißing.

(PTFE) to reflect the UV-photons. The XENONnT TPC is instrumented with 494 Hamamatsu R11410-21 3-inch PMTs (253 in the top array, 241 in the bottom array). These PMTS are selected based on their low radioactivity [102]) and high quantum yield ($\sim$ 34% [44]) for the wavelength of the xenon scintillation light. During the design and assembly of XENONnT the uttermost care was taken to select only those materials with the lowest contaminants for the construction of the TPC and cryostat [102].

Apart from the careful selection of materials for the XENONnT TPC, many more steps were taken to reduce backgrounds. The experiment is located in the underground laboratory of LNGS under 1500 m of rock (3600 m water equivalent) to shield the experiment from cosmic rays. Below, we will discuss the systems of XENONnT to achieve low background, stable operation, and signal recording, see Fig. 3.3. We emphasize the importance of those systems that are otherwise rarely mentioned throughout this work.

Muon Veto (MV) The XENONnT cryostat is housed within the Muon Veto (MV), which consists of a water tank that is filled with ultrapure water and is instrumented with 84 Hamamatsu R5912ASSY 8-inch PMTs. The MV is shown on the left-hand side of Fig. 3.3 on which a banner hangs that depicts the inside of the water tank (including the cryostat). The MV serves both as a passive shielding against radioactivity from the rock as well as an active veto against cosmic muons by detecting the Cerenkov light emitted by muons travelling through water at a speed faster than the speed of light in water. The MV has a tagging efficiency of 99.5% for muons and 70% for secondary particles from muon interactions in the surrounding rock [103, 104][1].

Neutron Veto (NV): In addition to the MV, a second veto system, the Neutron Veto (NV) is situated between the cryostat and the optically separated MV (as illustrated on the banner of the water tank Fig. 3.3). As the name suggests, the NV aims to veto neutrons, that scatter both inside and outside the TPC. The NV does this by detecting the 2.2 MeV capture gamma after thermalization of neutrons through the detection of the Compton scatters of the gamma. The NV is instrumented with 120 Hamamatsu R5912-100-10 8-inch PMTs. The NV has an estimated tagging efficiency of neutrons of 68% [97]. In future operations, the water inside the NV will be loaded with gadolinium (Gd), which increases capture probability[2]. A thermal neutron captured by Gd will result in a 8 MeV gamma cascade, which will increase the tagging efficiency to 87% [44].

Support building: Next to the water tank is the support building (Fig. 3.3, right) which houses the systems required to run the experiment. Much of the infrastructure is re-used from the predecessor of XENONnT, XENON1T. These systems are discussed below.

Cryogenics: The cryogenic system liquifies the xenon in the experiment and keeps it stably at a temperature of −96 °C. To get xenon to and from the TPC an “umbilical" runs from the top of the cryostat. Additionally, an LXe-pipe runs from the bottom of the LXe purification.

[1] Note that these numbers were from XENON1T, before the installation of the optically separated Neutron Veto (NV). Hence, if the MV is taken alone in XENONnT, the expected efficiency is lower, but together with the NV it should at least equal the cited XENON1T efficiency.

[2] The probability to capture a neutron in less than 20 cm is 91% for doping with 0.2% Gd relative mass concentration compared to 9% for H [44].

GXe and LXe purification Drifting electrons in the LXe can be trapped by electronegative impurities like H_2O and O_2 dissolved in the LXe. These diminish the number of electrons that arrive at the liquid-gas interface, and thereby the size of the S2 signal. The xenon purity is measured by the so-called "electron lifetime", which is the typical time a freed electron is able to drift before being "trapped" by some impurity in the xenon.

One of the upgrades from XENONnT with respect to XENON1T, is the addition of the liquid xenon purification which allows purification of vast quantities of xenon [105]. In XENON1T, the purification was only performed via GXe purification, which serves as an ancillary method for XENONnT. For the GXe purification, liquid xenon has to be converted to the gas phase, is then passed through zirconium getters, and has to be re-liquified before being recirculated to the TPC. The GXe purification is able to process 510 kg/d, while LXe purification processes 8.6 t/d, i.e., the entire xenon mass of XENONnT every day. GXe purification does have a $\sim 67\%$ higher efficiency of removing electronegative impurities than LXe purification, but despite this, the LXe purification processes the xenon much more often, resulting in a greatly improved electron lifetime.

Using the LXe purification, much better xenon purities are achieved in XENONnT. In XENON1T, the electron lifetime was 380 µs–650 µs, while in XENONnT this number is >10 ms. Without the improvements in purification, the XENONnT experiment would (with an electron lifetime of 650 µs) have had severe problems with reconstructing events from the bottom of the TPC as their electrons have to drift as much as 2.2 ms to the liquid-gas interface.

Radon and krypton distillation: Two additional systems, the radon and krypton distillation, specifically target the backgrounds that dissolve in xenon. The first is the radon distillation system, which is used to reduce the otherwise dominant source of background. ^{222}Rn constantly emanates from materials like the stainless steel used for the cryostat, cables, PTFE, and other detector materials, requiring continuous radon distillation. The ^{222}Rn has a long decay chain with many radioactive daughter isotopes. Especially ^{214}Pb causes a β-spectrum extending down to the energy threshold of the experiment [97], where it poses a background for beyond the SM searches [97]. Predominantly at low ($< \mathcal{O}(10)$ keV) energies, it can also form a background for WIMP-searches, as low energy ERs can be misreconstructed as NRs. The radon distillation was significantly upgraded from XENON1T, where already an all-time low ER background of 76 ± 2 events/(t × yr × keV) [106] was achieved. In XENONnT

the ER background was further reduced to 15.8 ± 1.3 events/(t × yr × keV) [97], which is mainly driven by advancements in the radon distillation and due to the radio purity screening [102]. In Chapter 7, the results of the low ER search are shown in Fig. 7.3 clearly indicating the reduction of ER background with respect to XENON1T. The XENONnT radon distillation system is able to process a large fraction of the xenon, with a flow of 1.7 t/d through the distillation system.

The other system is responsible for distilling out anthropogenic ^{85}Kr that is present in trace amounts in air, and thus xenon as it is obtained from air. When newly purchased xenon is brought underground, it is only distilled once since - unless in the case of an air leak - no new contamination of ^{85}Kr takes place.

ReStox The steel container shown in the bottom left of the support structure in Fig. 3.3 is ReStox I. ReStox I is a safety and storage system that allows recuperation of 7.6 t of xenon and storing it at cryogenic or room temperatures. ReStox II (not shown in this figure) is responsible for the recuperation of the remainder of the xenon in case of an emergency. It can store 10 t of xenon at room temperature (more than the total XENONnT inventory). In case of an emergency, 500 kg/h can be recuperated into ReStox II.

Slow Control and DAQ The Slow Control and the DAQ are two separate systems, both housed in the DAQ-room of the experiment. Slow Control monitors and manages the safety of other systems and sends out alarms and messages in case one of the parameters goes over a certain limit. It also provides a human operational interface for operations and monitoring of the systems mentioned above. The DAQ is the system responsible for writing the signals as measured by the PMTs to disk and providing the tools to monitor the detector conditions live. The DAQ will be elaborated on extensively in Chapter 5.

In this section, we have described the XENONnT experiment. We have showed the infrastructure to operate one of the least radioactive environments on earth in the hunt for dark matter. Xenon makes an excellent DM detection material for its high atomic mass, its scalability to the multi tonne scale, its high density (3 g/cm^3), and transparency to its own scintillation light. Furthermore, advancement in purification techniques yield very clean environments.

Dual-phase TPCs currently are the most sensitive detection technique for $\gtrsim$5 GeV/c^2 [4, 5, 107] WIMPs. Lighter WIMPs will - if they exist - be detected by a very different kind of experiment, as will be discussed in the next chapter.

Chapter 4

Complementarity of light Dark Matter searches

The detection of Weakly Interacting Massive Particle (WIMP) Dark Matter (DM) would be a huge step forward in the understanding of the universe on cosmological scales. At the same time, it would give rise to many questions as to the nature of the WIMP particle. Furthermore, any initial WIMP detection claim would have to be followed up by a confirmation by another experiment, preferably by an experiment that is based on a different technology. Combining the results of multiple experiments allows constraining the reconstructed parameters of the WIMP particle and the distribution of DM itself.

Recently, the scope of searches for WIMP particles has broadened (to lower WIMP-masses) as a "classical", heavier, WIMP particle remains undetected. The experimental efforts also diversify, such as recent searches for WIMP interactions with a detection material through the hypothesized Migdal effect.

In this chapter, we show the results as published in Ref. [1], by simulating WIMP induced nuclear recoil (NR) and Migdal signals in the SuperCDMS SNOLAB and the XENONnT experiments. We probe their complementarity for reconstructing the WIMP parameters.

Complementarity of direct detection experiments in search of light Dark Matter

Published in: Journal of Cosmology and Astroparticle Physics, 10 (2022) 004.
Publication date: Oct 3, 2022

Available online at:

- https://doi.org/10.1088/1475-7516/2022/10/004
- https://arxiv.org/abs/2204.01580

Reference [1], source code available at [2].

Abstract

Dark Matter experiments searching for Weakly interacting massive particles (WIMPs) primarily use nuclear recoils (NRs) in their attempt to detect WIMPs. Migdal-induced electronic recoils (ERs) provide additional sensitivity to light Dark Matter with $\mathcal{O}(\mathrm{GeV}/c^2)$ masses. In this work, we use Bayesian inference to find the parameter space where future detectors like XENONnT and SuperCDMS SNOLAB will be able to detect WIMP Dark Matter through NRs, Migdal-induced ERs or a combination thereof. We identify regions where each detector is best at constraining the Dark Matter mass and spin independent cross-section and infer where two or more detection configurations are complementary to constraining these Dark Matter parameters through a combined analysis.

4.1 Introduction

Many Dark Matter direct detection experiments aim to observe Dark Matter (DM) through an excess of nuclear recoils (NRs) caused by Weakly Interacting Massive Particles (WIMPs) scattering off nuclei from a target material [108–111]. For light Dark Matter, this is not always the most sensitive method of detection. For example, the dual-phase liquid xenon experiment XENON1T has reached world-leading sensitivities for a broad range of WIMP-masses using NRs [4] but sensitivity drops quickly for WIMP masses $\lesssim 5$ GeV/c^2 as the kinetic energy of the WIMP is not sufficient to generate a detectable recoil. The lower energy NRs for lighter WIMP-masses typically produce fewer photons and

the signal drops below the detection threshold. In contrast, cryogenic semiconductor experiments like the Super Cryogenic Dark Matter Search at Sudbury Neutrino Observatory Lab (SuperCDMS) [45] are much better suited for detecting such light DM, due to a combination of a lighter target element, a low energy threshold, and an excellent energy resolution.

The Migdal effect [112–117] is a rare, inelastic scattering process that allows the transfer of more energy to the target than with an ordinary NR. When an NR causes displacement of the nucleus with respect to the electrons of the atom, the resulting perturbation to the electric field experienced by the electrons may cause ionization or excitation of the atom. As such the Migdal manifests itself as an NR causing an electronic recoil (ER). While it has not been experimentally confirmed, it offers the possibility for experiments to extend their DM search region to lower WIMP masses [72, 118–122] since NRs that fall below the NR energy threshold of an experiment may result in detectable ERs.

This paper demonstrates the capability of experiments like XENONnT [44] (the upgrade of XENON1T) and SuperCDMS to reconstruct light Dark Matter, through a combination of NR and Migdal searches. Furthermore, we show how the combination of the two experiments would further improve the reconstruction of the DM properties. We benchmark the sensitivity of a given detection channel by simulating low mass WIMP signals. We then use Bayesian inference to reconstruct the simulated WIMP mass and cross section. By combining the likelihoods of the two experiments, we study their complementarity.

Ref.'s [123, 124] have previously demonstrated how experiments employing different target materials such as germanium, xenon and argon could complement each other when using an NR search to reconstruct the Dark Matter mass and cross section. Additionally, the effect of uncertainties of astrophysical parameters on the reconstruction was investigated (see for example Refs. [125, 126]). In this work, we will take into account more recent detector characteristics specifically aimed at detecting light Dark Matter through NRs or Migdal analyses.

In the following section (section 4.2), we review the theory of the NR and Migdal processes. The methods section (section 4.3) discusses the XENONnT and SuperCDMS detectors, after which the statistical inference framework is introduced. In the results section (section 4.4) we show the posterior distributions for several benchmarks of interest which we then generalize by exploring the parameter space for WIMP-masses between $0.1 - 10$ GeV/c^2 and we conclude by summarizing the results (section 4.5).

4.2 Theory

4.2.1 Nuclear recoils

The elastic recoil spectrum caused by a WIMP of mass M_χ scattering off a target nucleus $N(A, Z)$ with mass M_N is described by the differential recoil rate [123]:

$$\frac{dR}{dE_{\rm nr}}\left(E_{\rm nr}\right) = \frac{\rho_0}{M_\chi M_N} \int_{v_{\rm min}}^{v_{\rm max}} d^3\vec{v} v F\left(\vec{v} + \vec{v}_e\right) \frac{d\sigma_{\chi-N}}{dE_{\rm nr}}(v, E_{\rm nr}, A)\,, \tag{4.1}$$

where $E_{\rm nr}$ is the nuclear recoil energy, $\vec{v}$ is the WIMP velocity in the detector's rest frame for a Dark Matter model with local Dark Matter density ρ_0, $\vec{v}_e$ is the Earth's velocity with respect to the galactic rest frame, $F(\vec{v})$ the WIMP velocity distribution in the galactic rest frame and $\sigma_{\chi-N}$ is the WIMP-nucleus cross section. We will use the same formulation of $\sigma_{\chi-N}$ as in Ref. [123], and only take the spin-independent WIMP-nucleus cross section ($\sigma_{\rm S.I.}$) into account. The upper integration limit $v_{\rm max}$ is given by the sum of the Dark Matter escape velocity $v_{\rm esc}$ and $\vec{v}_e$. The lower integration limit $v_{\rm min}$ is the minimum WIMP velocity required to generate an NR of energy $E_{\rm nr}$. The value of $v_{\rm min}$ is kinematically constrained and dependent on the target material and recoil energy,

$$v_{\rm min}\left(E_{\rm nr}, M_\chi, A\right) = \sqrt{\frac{M_N E_{\rm nr}}{2\mu_N^2}}\,, \tag{4.2}$$

where $\mu_N = \frac{M_\chi M_N}{M_\chi + M_N}$ is the reduced mass and A the atomic mass number of $N(A, Z)$. From Eq. (4.1) we see that for a given recoil rate, a degeneracy exists between $\sigma_{\chi-N}$ and M_χ. However, since $v_{\rm min}$ also depends on M_χ, this degeneracy may be broken. Only when $M_\chi \gg M_N$, Eq. (4.2) becomes effectively independent of M_χ, at which point Eq. (4.1) becomes degenerate for the cross section and WIMP-mass.

In the case of non-directional detectors like XENONnT and SuperCDMS, we can simplify Eq. (4.1) using the Dark Matter speed distribution $f(v) = 4\pi v^2 F(v)$ and ignoring annual modulation effects due to the Earth's orbit around the Sun,

$$\frac{dR}{dE_{\rm nr}}\left(E_{\rm nr}\right) = \frac{\rho_0}{M_\chi M_N} \int_{v_{\rm min}}^{v_{\rm esc}} dv\; v f\left(|\vec{v} + \vec{v}_e|\right) \frac{d\sigma_{\chi-N}}{dE_{\rm nr}}(v, E_{\rm nr}, A)\,. \tag{4.3}$$

Earth's velocity relative to the galactic rest frame $\vec{v}_e$ relates to the velocity with respect to the local standard of rest ($\vec{v}_{\rm lsr}$), the peculiar velocity ($\vec{v}_{\rm pec}$) of the Sun with respect to $\vec{v}_{\rm lsr}$ and Earth's velocity ($\vec{v}_{\text{Earth-Sun}}$) via

$$\vec{v}_e = \vec{v}_{\rm lsr} + \vec{v}_{\rm pec} + \vec{v}_{\text{Earth-Sun}} \simeq \vec{v}_{\rm lsr} = \vec{v}_0 \,, \tag{4.4}$$

where we have approximated $\vec{v}_e \simeq \vec{v}_{\rm lsr}$ which will be referred to as $\vec{v}_0$ throughout this work [127].

We use a Maxwellian velocity distribution for the Dark Matter velocity distribution $F(v)$, also referred to as the Standard Halo Model [80]. For the astrophysical parameters we assume $v_0 = 233$ km/s, $v_{\rm esc} = 528$ km/s and $\rho_0 = 0.55$ GeV/cm^3 [30]. This Dark Matter density ρ_0 is different from the 0.3 GeV/cm^3 usually assumed for direct detection Dark Matter experiments [4, 89, 128] which is adopted by convention as its value is directly proportional to the recoil rate as in Eq. (4.1) and can therefore be easily scaled. Ref. [32] provides an overview of recent publications on ρ_0 where ranges of $0.4 - 0.6$ and $0.3 - 0.5$ GeV/cm^3 are quoted depending on the type of analysis. Using Eqs. (4.1-4.4), the differential NR rate can be computed for a given target material and a set of astrophysical parameters.

4.2.2 Migdal

For lower mass WIMPs, fewer NR energies exceed the energy threshold. However, low-energy recoil interactions may be detected through the so-called Migdal effect. Although it is usually assumed that the electrons after an NR interaction always accompany the nucleus, it actually takes some time for the electrons to catch up, resulting in ionization and excitation of the recoil atom [115]. These effects can lead to detectable energy deposits in a detector similar to the energy depositions caused by ERs. The differential recoil rates are calculated for several materials assuming isolated atoms in Ref. [115]. For semiconductors, the calculation of the Migdal-induced rates needs to go beyond this isolated atom approximation as was done in Ref. [117].

In the isolated atom approximation of Ref. [115], the differential rate for Migdal-induced signals combines the standard NR recoil energy distribution with the electronic band structure of the target atoms. The differential Migdal rate is described by the convolution of the NR differential rate with the probability of ionization,

$$\frac{dR}{dE_{\rm er}} \simeq \int dE_{\rm nr} dv \frac{d^2R}{dE_{\rm nr} dv} (E_{\rm nr}) \times \sum_{n,l} \frac{d}{dE_{\rm er}} P^c_{q_e} (n, l \to E_{\rm er} - E_{n,l}) \,, \tag{4.5}$$

where $P^c_{q_e}$ is the probability for an atomic electron with quantum numbers (n, l) and corresponding energy $E_{n,l}$ to be emitted with a kinetic energy of $E_{\rm er} - E_{n,l}$. The values of $P^c_{q_e}$ are taken from Ref [115].

Ref. [117] includes a derivation of the Migdal-induced rates in semiconductors for WIMP-nucleus scattering. Because of the smaller gap for electron excitations, the Migdal rates are found to be higher than for the isolated atom approximation. The differential electronic recoil rate is

$$\frac{dR}{dE_{\rm er}} \simeq \frac{\rho_0}{M_\chi M_N} \frac{4\alpha Z^2}{3\pi^2 E_{\rm er}^4 M_N} \int dk\, k^2 {\rm Im}\left(\frac{-1}{\epsilon(k, E_{\rm er})}\right) \int_{v_{\rm min}}^{v_{\rm max}} d^3\vec{v} v F\left(\vec{v} + \vec{v}_e\right) \int dE_{\rm nr}\, E_{\rm nr} \frac{d\sigma_{qe}}{dE_{\rm nr}}\,, \quad (4.6)$$

where α is the fine structure constant, $\frac{d\sigma_{qe}}{dE_{\rm nr}}$ the quasi-elastic cross section from [117], ${\rm Im}(-\epsilon^{-1}(k, E_{\rm er}))$ the energy loss function with ϵ the momentum and frequency dependent longitudinal dielectric function, and k is the momentum associated with the electronic excitation.

Using the Migdal effect, the NRs that fall below the energy threshold of experiments may still be indirectly detected as ERs. In other words, there is the possibility to detect NRs that are below the threshold through the associated ERs, thereby allowing detectors to be sensitive to smaller WIMP masses that would otherwise be undetectable.

4.3 Methods

We consider two experiments: XENONnT and SuperCDMS. These detectors are both sensitive to $\mathcal{O}\left({\rm GeV}/c^2\right)$ mass WIMPs, but with significant differences: SuperCDMS has a high quantum yield with a relatively modest target mass, while XENONnT combines a lower light and charge yield with a multi-tonne target mass.

In the remainder of this section, we describe the methods we use for modeling the detectors, calculating the signal spectra, and inferring projected constraints on the DM parameters. The detector characteristics which are used are summarised in Tab. 4.1. Example NR and Migdal spectra for the experiments are shown in Fig. 4.1. We use `pymultinest` to sample from the posterior distribution of the spin-independent WIMP-nucleon cross section and WIMP mass $(\sigma_{\rm S.I.}, M_\chi)$, assuming the benchmark points and priors given in Tab. 4.2. The

Experiment	XENONnT	SuperCDMS			
		Ge HV	Si HV	Ge iZIP	Si iZIP
NR and Migdal (ER)					
Target mass (kg)	4×10^3	11	2.4	14	1.2
Live time	100%	80%	80%	80%	80%
Run time (yr)	5	5	5	5	5
Exposure (kg · year)	20×10^3	44	9.6	56	4.8
k-parameter for Eq. (4.14)	0.1735	0.162	0.161	0.162	0.161
NR					
E_{range} (keVnr)	[0, 5]	[0, 5]	[0, 5]	[0, 5]	[0, 5]
Cut- and detection-eff.	0.83	$0.85 \cdot 0.85$	$0.85 \cdot 0.85$	$0.85 \cdot 0.75$	$0.85 \cdot 0.75$
Energy resolution for $\sigma_{\text{ph, nr}}$ (HV) / $\sigma_{Q,\text{nr}}$ (iZIP)	Eq. (4.8)	Eq. (4.22) 10 eV	Eq. (4.22) 5 eV	Eq. (4.23) 100 eV	Eq. (4.23) 110 eV
BG. $\left(\frac{\text{counts}}{\text{kg·keV·year}}\right)$	2.2×10^{-6}	27	300	3.3×10^{-3}	2.9×10^{-3}
E_{thr} (keV_{nr})	1.6	0.040	0.078	0.272	0.166
Migdal (ER)					
E_{range} (keV_{ee})	[0, 5]	[0, 0.5]	[0, 0.5]	[0, 0.5]	[0, 0.5]
Cut- and detection-eff.	0.82	$0.5 \cdot 0.85$	$0.675 \cdot 0.85$	$0.5 \cdot 0.75$	$0.675 \cdot 0.75$
Energy resolution	Eq. (4.7)	0.4 eV_{ee}	0.15 eV_{ee}	19 eV_{ee}	7 eV_{ee}
BG. $\left(\frac{\text{counts}}{\text{kg·keV·year}}\right)$	12.3×10^{-3}	27	300	22	370
E_{thr} (keVee)	1.0	0.004	0.003	0.14	0.05

Table 4.1: The assumed detector characteristics of XENONnT and SuperCDMS. SuperCDMS consists of various detector target materials (Si, Ge) and designs (HV, iZIP). The first set of detector parameters (top part of the table) are independent of the type of analysis (NR or Migdal). For the NR and Migdal searches, the respective values are listed separately in the middle and bottom of the table.

results of these benchmark points are further generalized in the Results section (section 4.4).

For both experiments we assume a five-year run time which the experiments aim to acquire on similar timescales [44, 45]. The product of a combined cut- and detection- efficiency, run time, live time and target mass yields the effective exposure ϵ_{eff}.

Below, we describe the detector characteristics which are used for the recoil rate calculations, summarized in Tab. 4.1. In the following sections, we use the Lindhard theory [129] to convert between NR energies (E_{nr}) and electronic equivalent energies (E_{ee}) as explained in Appendix 4.6.1. For both the NR and Migdal search, we require the cut- and detection-efficiency, energy resolution, background rate, and energy thresholds for the calculation of the spectra. As the Migdal effect manifests itself as an ER signal, some parameters are different from the NR search, such as the expected background in case the detector has the ability to distinguish NRs and ERs. Other parameters like target mass and exposure are independent of the type of search. We conclude this section with a description of the Bayesian framework we use for the analysis.

4.3.1 XENONnT

XENONnT is the upgrade of XENON1T with a larger target mass and lower background expectation [44]. For the NR and Migdal detection channels, we assume a 4 tonne active target mass and continuous data taking (live time of 100%), yielding a total of 20 tonne year exposure.

XENONnT measures both prompt scintillation light (S1) and ionization signals (S2). Since NRs with the same energy cause relatively smaller ionization signals, XENONnT is able to distinguish between ERs and NRs. Most of the background events in XENONnT are from radioactive contaminants like radon and krypton causing ERs within the active target volume. The background rate for the NR search can therefore be reduced because of the ER/NR discrimination. We assume a background rate of $2.2 \cdot 10^{-3}$ (12.6) $\text{keV}^{-1}\text{t}^{-1}\text{yr}^{-1}$ for the NR (Migdal) search [44]. We will first discuss the parameters relevant for the Migdal search followed by those for the NR search.

For the Migdal search, the detector ER energy resolution (σ in keV_{er}) is assumed to be the same as for XENON1T [106] which is given by the empirical formula:

$$\sigma_{\text{er}}(E_{\text{er}}) = 0.31\,\text{keV}_{\text{er}}\sqrt{\frac{E_{\text{er}}}{\text{keV}_{\text{er}}}} + 0.0037\,E_{\text{er}}\,. \tag{4.7}$$

The ER detection energy threshold relevant for the Migdal search ($E_{\mathrm{thr,\,er}}$) is assumed to equal 1.0 $\mathrm{keV_{er}}$ [106]. This energy threshold is dictated by the requirement of reconstructing the S1 of an interaction [118]. While lower thresholds are achieved in S2-only analyses, these can only lead to exclusion of Dark Matter models as not all backgrounds can be adequately modelled [130]. Therefore, these lower thresholds are not used here.

The Migdal recoil energies are limited to the interval of [0, 5] $\mathrm{keV_{er}}$. While Ref. [115] assumes target materials to consist of isolated atoms, XENONnT uses liquid xenon as the target material. To account for this difference and in order to be conservative, the contribution to the differential recoil rate from the $5,1$ shell is neglected. We do take the $5,0$ shell into account which contributes $\lesssim 2\%$ to the total rate for the masses considered in this work. Furthermore, the innermost electrons are considered too tightly bound to the nucleus to contribute significantly [72, 115, 118]. Finally, we assume a combined detection and cut efficiency of 83% (82%) for NR (Migdal) [44].

For the NR search, we use the Lindhard factor L (explained in section 4.6.1) in Eq. (4.15) to convert E_{nr} to E_{ee} and treat the energy resolution (Eq. (4.7)) as the uncertainty on the value of the detected energy:

$$\sigma_{\mathrm{nr}}(E_{\mathrm{nr}}) = \frac{\mathrm{d}E_{\mathrm{nr}}}{\mathrm{d}E_{\mathrm{er}}}\sigma_{\mathrm{er}}(E_{\mathrm{ee}}) = \frac{\mathrm{d}E_{\mathrm{nr}}}{\mathrm{d}E_{\mathrm{er}}}\sigma_{\mathrm{er}}\left(L(E_{\mathrm{nr}})\cdot E_{\mathrm{nr}}\right)\,, \tag{4.8}$$

to obtain the NR energy resolution σ_{nr}. A value of $k = 0.1735$ [131] is used for XENONnT in Eq. (4.14). We assume an analysis optimized for low energy events. We set an energy threshold $E_{\mathrm{thr,\ nr}}$ of 1.6 $\mathrm{keV_{nr}}$, which has been achieved in XENON1T with the dedicated low energy NR search for coherent elastic scattering of solar neutrinos [132]. The energy range of interest is set to [0, 5] $\mathrm{keV_{nr}}$.

4.3.2 SuperCDMS

The SuperCDMS experiment [45] has two detector designs each using germanium and silicon as target material. The so-called HV detector only utilizes phonon sensors, whereas the iZIP detector uses both phonon and ionization sensors, thereby allowing ER/NR discrimination. Since the HV detectors are not able to distinguish between ER and NR, most of the detector parameters are the same for the Migdal (ER) and NR search. For the iZIP detectors some detector parameters differ for the two types of searches because of the ER/NR discrimination.

The HV detectors have better phonon energy resolution compared to the iZIP detectors, which results in a better sensitivity for WIMP masses $\lesssim$ 5 GeV/c^2 as lower WIMP masses cause lower recoil energies. The iZIP detectors have better sensitivity for higher masses. We model each of the target materials for each of the detector designs, yielding four different configurations. The detector parameters are listed in Tab. 4.1.

The background in each detector is directly obtained from Table V. in Ref. [45]. The backgrounds of the HV detector (NR and Migdal search) are given by the ER backgrounds dominated by ^{3}H and ^{32}Si decays. The iZIP detector background for Midgal is also given by the ER background whereas the NR search background, which is mostly due to coherent neutrinos, is significantly lower due to the NR/ER discrimination.

The energy-scales, -resolution and -thresholds for the four detector configurations for both NR and Migdal are summarized in Appendix 4.6.1. Their respective values are listed in Tab. 4.1. For the NR search, we use a [0, 5] $\mathrm{keV_{nr}}$ energy range. As the electronic recoil energies for the Migdal search are typically at low energy, we focus on the energy range of [0, 0.5] $\mathrm{keV_{er}}$.

4.3.3 Recoil rates

In order to evaluate the recoil spectra, we evaluate Eq. (4.1) or Eq. (4.5) using the `wimprates`-framework [133] and Eq. (4.6) using the `darkelf`-framework [117, 134]. For evaluating the energy loss function in Eq. (4.6), we use the GWAP method for $E_{\mathrm{er}} < 60$ eV and Lindhard method for larger energies as no data for the GPAW [134] method is available at energies $E_{\mathrm{er}} \gtrsim 75$ eV and the methods agree well for recoils above 60 eV. To calculate the recoil rates, we assume the astrophysical parameters as per the Standard Halo Model. We will limit ourselves to WIMPs that couple to the target nucleus through spin-independent interactions.

We add a flat background spectrum to the NR or Migdal recoil spectrum prior to convolving the spectrum with the detector resolution σ, resulting in the detectable energy spectrum

$$\frac{d\tilde{R}}{dE_R} = \int dE' \frac{dR}{dE_R}(E') \frac{e^{-\frac{(E-E')^2}{2\sigma^2(E')}}}{\sqrt{2\pi}\sigma(E')} \,. \tag{4.9}$$

The number of expected events N_i in a given energy bin is obtained by integrating Eq. (4.9) times the effective exposure (ϵ_{eff}) between the bin edges

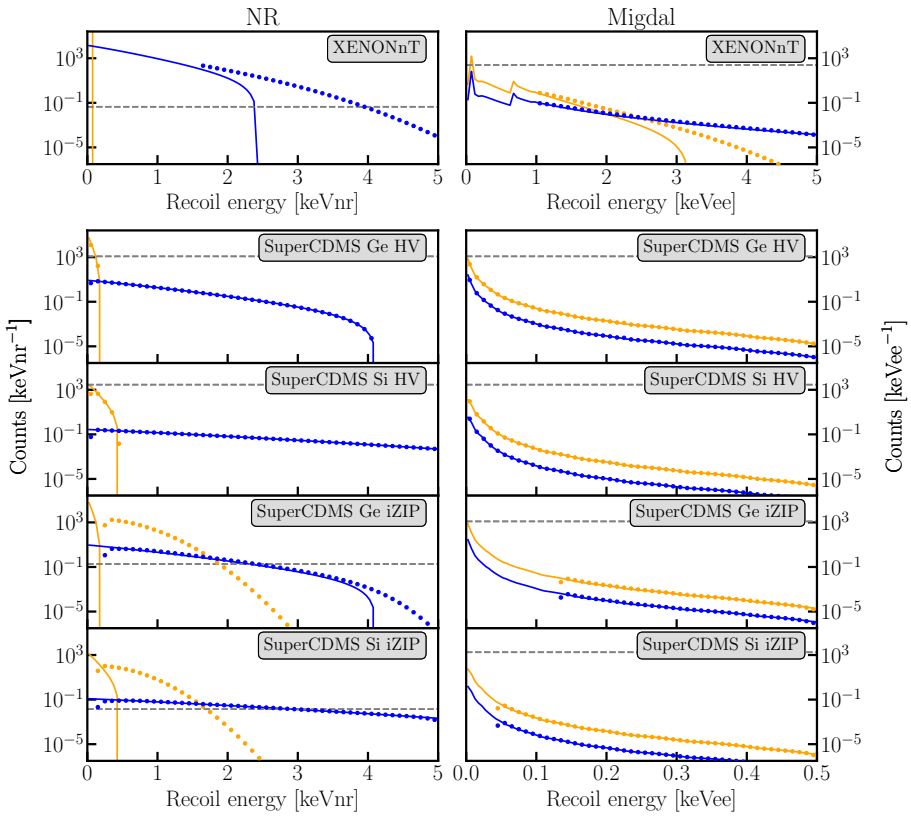

Figure 4.1: The differential recoil rate (solid line) results in the detectable spectrum (dots) when the detector energy threshold and detector resolution are taking into account, and the spectrum is binned in 50 energy bins. The background rates for the given exposures are shown separately (dashed gray lines). The left column shows the NR spectra and the right column the ER spectra as a result of the Migdal effect. For all NR searches, the energy range is restricted to [0, 5] $\mathrm{keV_{nr}}$, while for Migdal the SuperCDMS searches use a smaller energy interval of [0, 0.5] $\mathrm{keV_{er}}$ compared to XENONnT ([0, 5] $\mathrm{keV_{er}}$). Recoil spectra for WIMP DM with $M_\chi = 5$ GeV/c^2 and $\sigma_{\mathrm{S.I.}} = 10^{-45}$ cm^2 (blue) and $M_\chi = 1$ GeV/c^2 and $\sigma_{\mathrm{S.I.}} = 10^{-42}$ cm^2 (orange) for the exposures listed in Tab. 4.1. In the XENONnT-NR panel, the recoil rate for $M_\chi = 1$ GeV/c^2 falls off exponentially well below the energy threshold of 1.6 $\mathrm{keV_{nr}}$ and the detectable spectrum is ~ 0 counts $\mathrm{keV_{nr}}^{-1}$. For example for the XENONnT detector, especially with $M_\chi = 1$ GeV/c^2, the top panels show why the Migdal effect can help experiments extend their search region, since even though the spectrum drops steeply below the NR energy threshold, the Migdal spectrum extends sufficiently beyond the detector energy threshold of 1.0 $\mathrm{keV_{ee}}$ to higher ER energies.

$E^i_{\rm min}$, $E^i_{\rm max}$,

$$N_i = \int_{E^i_{\rm min}}^{E^i_{\rm max}} dE_R \epsilon_{\rm eff} \frac{d\tilde{R}}{dE_R} \,. \tag{4.10}$$

Fig. 4.1 shows the spectra obtained for NR and Migdal before- and after-including detector effects as well as the background rates for each detector. We approximate the spectrum by a 50-bin spectrum which allows for reasonably fast computation of spectra.

We model the Migdal spectra and NR spectra independent from each other. In a real detector when DM would be observed through the Migdal effect, the direct NRs may also be observed. This is especially relevant for detectors where there is no NR/ER discrimination as the Migdal and NR contribution could not be disentangled. Since we want to investigate the ability of detectors to detect DM through either Migdal or NR, we take their resultant spectra separately into account as if only one or the other would be observed.

4.3.4 Statistical inference

We follow a Bayesian approach [135] to extract the parameters of interest (M_χ and $\sigma_{\rm S.I.}$) similar to the method described in Ref. [123]. The total likelihood $\mathcal{L}$ is the product of the likelihood for each detector which is given by the product of the Poisson probability of each of the energy bins

$$\mathcal{L}\left(\Theta\right) = \prod_{j}^{\rm detectors} \left(\prod_{i}^{\rm bins} \frac{\hat{N}_{ij}(\Theta)^{N_i}}{N_i!} e^{-\hat{N}_{ij}(\Theta)} \right) , \tag{4.11}$$

where N_i is the number of counts in each energy bin (i) and $\hat{N}_{ij}(\Theta)$ is the expected counts for a given detector (j) at the set of parameters Θ, where Θ contains the DM parameters of interest,

$$\Theta = \{M_\chi, \sigma_{\rm S.I.}\} \,. \tag{4.12}$$

To infer the posterior distribution, the likelihood $\mathcal{L}(\Theta)$ is multiplied by the prior $p(\Theta)$ for given parameters Θ. We choose a flat prior in log-space for the mass and cross section as their true value is unknown and the aim is to reconstruct these parameters. Given the very steep rise in sensitivities for SuperCDMS and XENONnT in the mass range considered here, a large prior range was chosen for the masses of interest. Each of the prior ranges was set around the central value for the three benchmark points of interest, as in Tab. 4.2.

M_χ (GeV/c^2)	$\sigma_{\mathrm{S.I.}}$ (cm^2)	prior-range $\log_{10}\left(M_\chi/\left(\mathrm{GeV}/c^2\right)\right)$	prior-range $\log_{10}\left(\sigma_{\mathrm{S.I.}}/\mathrm{cm}^2\right)$
5	10^{-45}	$\log_{10}(5)-2.5, \log_{10}(5)+3.5$	$-52, -40$
3	10^{-41}	$\log_{10}(3)-2.5, \log_{10}(3)+3.5$	$-48, -36$
0.5	10^{-38}	$\log_{10}(0.5)-2.5, \log_{10}(0.5)+3.5$	$-45, -33$

Table 4.2: Benchmark points and corresponding prior ranges. For both the WIMP mass cross sections a flat prior is assumed in log-space. As the relevant cross sections greatly differ for the three WIMP masses, the prior ranges are scaled accordingly.

The likelihood for SuperCDMS at Θ is given by the product of the likelihood of the Ge HV, Si HV, Ge iZIP and Si iZIP detectors. When combining the results of XENONnT and SuperCDMS, all five detectors are taken into account in the product over the detectors in Eq. (4.11).

To sample the posterior distribution several sampling methods are implemented in Ref. [2] such as `emcee` [136], `nestle` [137] and `pymultinest` [138]. Since the results are independent of the sampling method and `pymultinest` proved the fastest, it is used here. The `pymultinest`-package is a pythonic interface to the `multinest` algorithm [139, 140].

Using the `pymultinest` sampler, 1000 “live points" are generated that populate the prior volume. The live points iteratively probe the prior volume to obtain the posterior, see Ref. [140]. A tolerance of 0.5 is used as a stopping criterion. The samples are weighted to represent the posterior distribution density.

4.4 Results and discussion

For a given set of Dark Matter parameters Θ, a benchmark recoil spectrum is calculated for each of the detectors. We obtain the posterior distribution density using `pymultinest` to investigate how a binned Poisson likelihood analysis would be able to reconstruct the set DM parameters. This section compares the ability of SuperCDMS and XENONnT to correctly reconstruct Θ using either an NR or Migdal search.

SuperCDMS and XENONnT have different characteristics (Tab. 4.1) and their ability to reconstruct the benchmark value depends strongly on the assumed DM parameters. We give results for the three benchmark points in

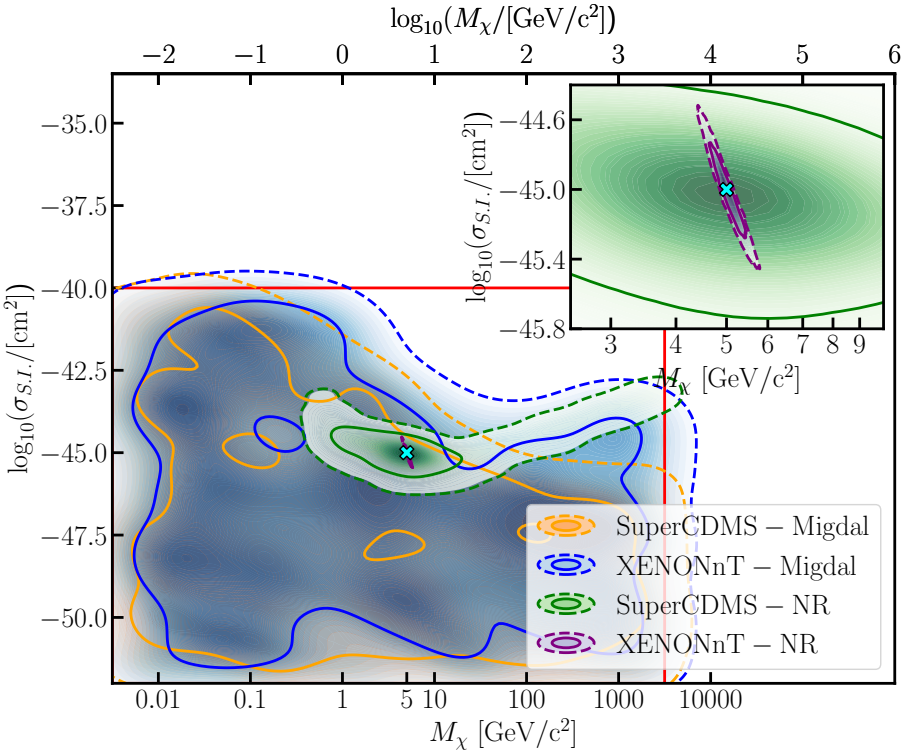

Figure 4.2: Posterior distribution densities reconstructed for a WIMP with $M_\chi = 5$ GeV/c^2 and $\sigma_{\mathrm{S.I.}} = 10^{-45}$ cm^2 in the four detector configurations. The 68% and 95% CIs are illustrated with the solid and dashed lines, respectively. Whereas the NR searches are able to reconstruct the set benchmark (cyan), the Migdal searches are not. The inset shows the posterior distribution densities XENONnT-NR and SuperCDMS-NR, where the 68% CI for the former is much smaller than that of the latter. The XENONnT-Migdal and SuperCDMS-Migdal reconstructed posteriors fill the prior volume (indicated by the red box), consistent with no signal.

Tab. 4.2 which lie close to the detection threshold of XENONnT. Next, we generalize this for other masses and cross-sections to find the complementarity of the four detector configurations.

4.4.1 5 GeV/c^2

We first simulate a benchmark Dark Matter model for WIMPs with $M_\chi = 5$ GeV/c^2 and $\sigma_{\text{S.I.}} = 10^{-45}$ cm^2. Fig. 4.2 shows the inferred posterior distribution for these Dark Matter parameters, which XENONnT NR-search (XENONnT-NR) reconstructs since the benchmark value is in the center of the posterior distribution density. Also, the SuperCDMS NR-search (SuperCDMS-NR) gives the Dark Matter parameters albeit with a larger 68% credibility interval (CI), while at large M_χ the 95% CI contour lines do not close due to a mass-cross section degeneracy as mentioned in the Theory section (section 4.2). The difference between XENONnT-NR and SuperCDMS-NR can be understood from Fig. 4.1: the number of expected events for XENONnT-NR for $M_\chi = 5$ GeV/c^2 is higher while the background is relatively lower than for SuperCDMS-NR, leading to a tighter 68% CI for XENONnT-NR.

The XENONnT Migdal-search (XENONnT-Migdal) and SuperCDMS Migdal-search (Super-CDMS-Migdal) are not able to reconstruct the benchmark point. For these detector configurations, the prior volume is filled where the signal would be consistent with no signal, since the expected recoil rates in Fig. 4.1 are relatively low and backgrounds generally higher compared to the NR searches (Tab. 4.1). When the cross section and WIMP mass are both higher, a sizable Migdal signal is expected. Therefore, the prior volume in the upper right corner of Fig. 4.2 is not filled by the posterior distributions of XENONnT-Migdal and SuperCDMS-Migdal.

We quantify how well the benchmark is reconstructed by calculating the fraction of the prior volume filled by the posterior volume in log-space of the enclosed 68% CI:

$$\phi = \frac{\log_{10}\left(\frac{M_\chi^{\text{enc. 68\%}}}{\text{GeV}/c^2}\right) \cdot \log_{10}\left(\frac{\sigma_{\text{S.I.}}^{\text{enc. 68\%}}}{\text{cm}^2}\right)}{\text{prior-volume}}, \tag{4.13}$$

which is the surface enclosed by the solid lines in Fig. 4.2 divided by the surface within the red box. The 68% CI is obtained using a bi-variate Gaussian kernel density estimator based on code from Ref. [141]. Values of $\phi \sim \mathcal{O}(0.1 - 1)$ indicate low power to reconstruct a benchmark model since the posterior volume is of similar size as the prior volume, the lower ϕ, the better the benchmark is reconstructed as the parameters are better constrained.

Evaluating ϕ for the results in Fig. 4.2 yields $\phi_{\text{XENONnT-NR}} = 6.1 \times 10^{-5}$ while $\phi_{\text{SuperCDMS-NR}} = 8.1 \times 10^{-3}$, showing that the XENONnT-NR

search yields $\mathcal{O}(10^2)$ times tighter constraints on the reconstructed parameters. For the Migdal searches ϕ is large ($\phi_{\text{XENONnT-Migdal}} = 3.9 \times 10^{-1}$) and ($\phi_{\text{SuperCDMS-Migdal}} = 3.5 \times 10^{-1}$). As the 95% CI do not close before the prior boundaries, these numbers only indicate that neither XENONnT-Migdal nor SuperCDMS-Migdal is able to reconstruct the DM parameters.

4.4.2 3 GeV/c^2

We simulate a WIMP of $M_\chi = 3$ GeV/c^2 and $\sigma_{\text{S.I.}} = 10^{-41}$ cm^2 near the detection threshold of XENONnT. At this mass and cross section, XENONnT-NR and SuperCDMS-NR both reconstruct a tight posterior distribution as in Fig. 4.3. As this cross section is higher than what was considered for 5 GeV/c^2, SuperCDMS-Migdal and XENONnT-Migdal are also able to reconstruct a broad posterior distribution which, for XENONnT-Migdal, has non-closing contour lines due to the mass-cross section degeneracy also observed for SuperCDMS-NR in Fig. 4.2.

We study the complementarity of XENONnT-NR and SuperCDMS-NR in Fig. 4.4. Whereas the reconstructed 68% CI for XENONnT-NR has a relatively large spread in $\sigma_{\text{S.I.}}$, SuperCDMS-NR has a large spread in M_χ. The likelihood of XENONnT-NR changes rapidly as function of M_χ since the drop in the recoil spectrum occurs close to the energy threshold for these WIMP masses. As a result, the likelihood constrains M_χ around this mass relatively well. In contrast, the uncertainty of SuperCDMS-NR is mostly in M_χ since a shift in the spectral shape as function of M_χ has a relatively smaller effect for SuperCDMS-NR on the number of events above threshold. Since $\sigma_{\text{S.I.}}$ is proportional to the number of events observed it is therefore relatively well constrained for SuperCDMS-NR.

When the likelihoods of the NR searches are combined, the 68% CI is reduced. Quantitatively, one can see this from $\phi_{\text{XENONnT-NR}} = 2.8 \times 10^{-6}$ and $\phi_{\text{SuperCDMS-NR}} = 1.1 \times 10^{-7}$ while the combination of the two gives $\phi_{\text{XENONnT-NR+SuperCDMS-NR}} = 5.1 \times 10^{-8}$. This corresponds to a reduction of ϕ by a factor of 54 (2.1) when the likelihoods of these detector configurations are combined, compared to XENONnT-NR (SuperCDMS-NR) alone. Both Migdal searches also constrain the posterior distribution, $\phi_{\text{SuperCDMS-Migdal}} = 8.8 \times 10^{-4}$ and $\phi_{\text{XENONnT-Migdal}} = 2.3 \times 10^{-2}$. However, since the 68% CI of SuperCDMS-Migdal and XENONnT-Migdal fully enclose the 68% CI of the XENONnT-NR search, their combination with the NR searches does not result in a lower value of ϕ.

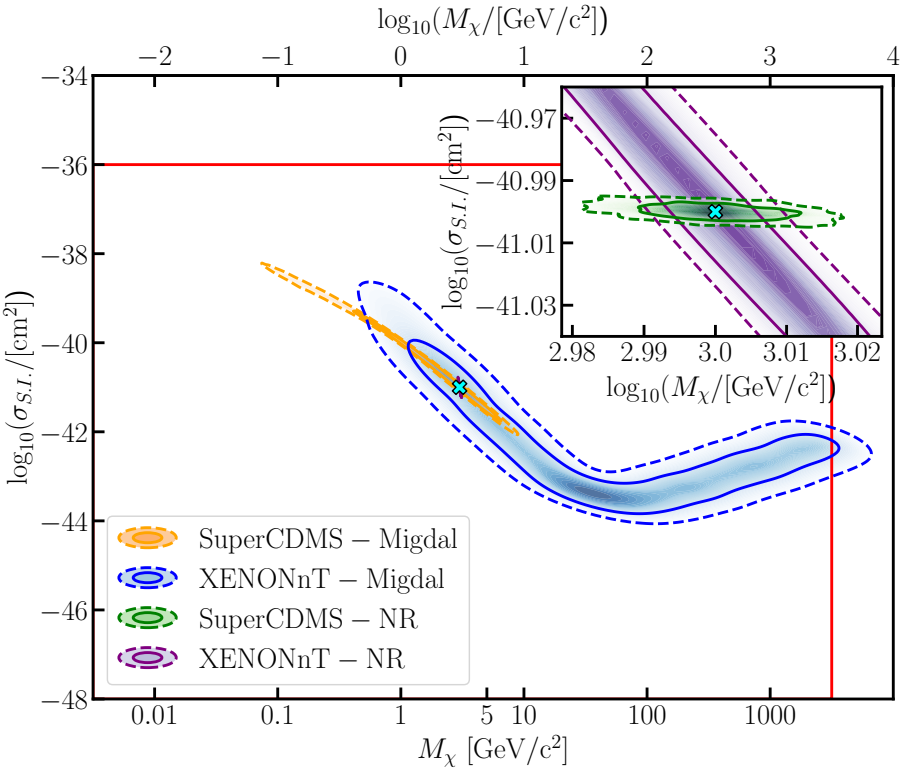

Figure 4.3: Posterior distributions reconstructed for a WIMP with M_χ = 3 GeV/c^2 and $\sigma_{\mathrm{S.I.}}$ = 10^{-41} cm^2 in the four detector configurations. SuperCDMS-NR and XENONnT-NR both reconstruct the benchmark point (cyan) even though the shapes of the posterior differ. Furthermore, the SuperCDMS-Migdal is also able to constrain the DM parameters with larger 68% and 95% CIs. The posterior for XENONnT-Migdal has non-closing contour lines as it extends to the boundary of the prior range as in Tab. 4.2.

4.4.3 0.5 GeV/c^2

When considering a lower mass WIMP of M_χ = 0.5 GeV/c^2 and $\sigma_{\mathrm{S.I.}}$ = 10^{-38} cm^2 the situations changes. The spectra in Fig. 4.1 are shifted to lower energies and for XENONnT-NR, the spectrum (before taking the detector ef-

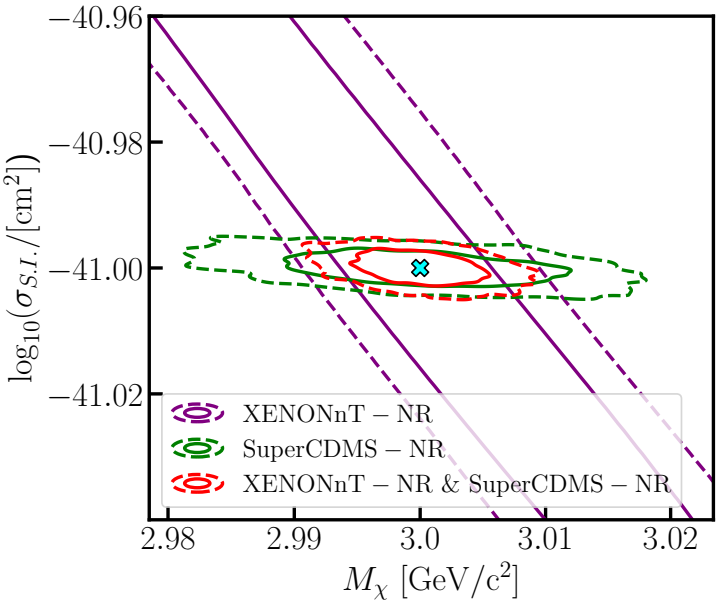

Figure 4.4: Overlaid posterior distributions reconstructed for a WIMP with $M_\chi = 3$ GeV/c^2 and $\sigma_{\mathrm{S.I.}} = 10^{-41}$ cm^2 for SuperCDMS-NR (green), XENONnT-NR (purple) and the combined result for SuperCDMS-NR and XENONnT-NR (red). The 68% CI (solid) and 95 % CI (dashed) contour lines are shown. The two experiment are complementary to each other since a combination of the two experiments yields a substantially tighter 68% CI as explained in the text.

fects into account) drops steeply below the energy threshold, leading to close to no events in the detector. At this cross section, the recoil rate for XENONnT-Migdal becomes sufficient to constrain the DM parameters. Fig. 4.5 shows the posterior distributions for the four detector configurations.

The SuperCDMS-Migdal search is able to reconstruct these DM parameters best, resulting in $\phi_{\text{SuperCDMS-Migdal}} = 6.0\times10^{-5}$. The NR search of SuperCDMS also constrains the DM parameters, achieving $\phi_{\text{SuperCDMS-NR}} = 2.2\times10^{-4}$. The XENONnT-NR search becomes insensitive as fewer signals are above the energy threshold ($\phi_{\text{XENONnT-NR}} = 2.3 \times 10^{-1}$), the posterior distribution function fills the prior volume up to ~ 3 GeV/c^2, where NRs are starting to be just above

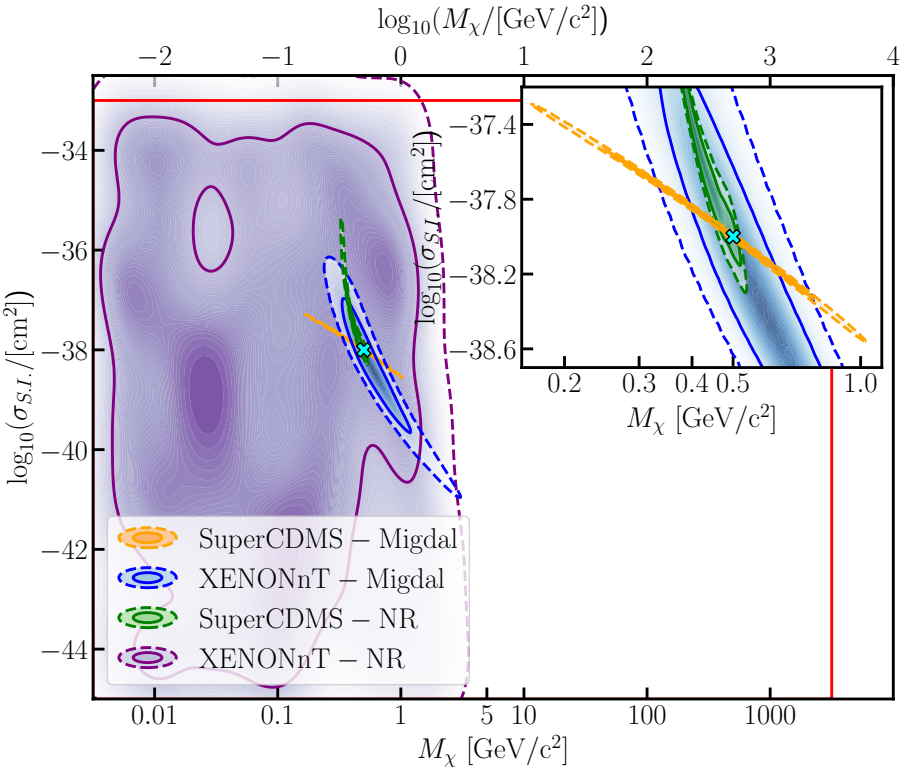

Figure 4.5: The posterior distributions reconstructed for a WIMP with $M_\chi = 0.5$ GeV/c^2 and $\sigma_{\mathrm{S.I.}} = 10^{-38}$ cm^2. SuperCDMS-NR and SuperCDMS-Migdal reconstruct the benchmark point (cyan) as the 68% CI (solid) and 95 % CI (dashed) center around the set benchmark. Whereas XENONnT-NR does not reconstruct the benchmark, the Migdal search does. Due to the few detected recoils and relatively large background for XENONnT-Migdal, the credibility interval is significantly larger than for SuperCDMS-NR or SuperCDMS-Migdal.

the detection energy threshold. In contrast, for such a cross section and mass, the XENONnT-Migdal search is able to constrain the posterior distribution ($\phi_{\text{XENONnT-Migdal}} = 2.3 \times 10^{-3}$). With the considered M_χ being close to the energy threshold of SuperCDMS-NR, the 68% CI of SuperCDMS-NR extends

to lower masses and higher cross sections with respect to the benchmark point since a higher mass would result in many more events. In contrast, the 68% CI of XENONnT-Migdal is quite broad due to the limited number of events at this cross section and mass, while being less affected by the energy threshold. Since the 68% CI of SuperCDMS-NR and XENONnT-Migdal cover different portions of the prior volume the combination of the two has a much lower ($\phi_{\text{SuperCDMS-NR+XENONnT-Migdal}} = 3.4 \times 10^{-5}$), which is a factor of 6 lower than for SuperCDMS-NR and a factor of 69 compared to XENONnT-Migdal. Even better results are achieved with the combination of SuperCDMS-NR and SuperCDMS-Migdal, where $\phi_{\text{SuperCDMS-NR+SuperCDMS-Migdal}} = 8.1 \times 10^{-8}$, which corresponds to a reduction of 7×10^2 for SuperCDMS-Migdal and 3×10^3 for SuperCDMS-NR.

4.4.4 Masses between 0.1-10 GeV/c^2

In order to generalize the results as in the sections above, we investigate how the following combined analyses would reconstruct Dark Matter parameters at several WIMP-masses and cross sections:

- A combined *NR* analysis using XENONnT-NR and SuperCDMS-NR,
- A combined *Migdal* analysis using XENONnT-Migdal and SuperCDMS-Migdal,
- A combination of *All* analyses; being XENONnT-NR, XENONnT-Migdal, SuperCDMS-NR and SuperCDMS-Migdal.

For each of these analyses, we evaluate ϕ for a scan of points in M_χ-$\sigma_{\text{S.I.}}$ space. We will refer to these values as ϕ_{NR}, ϕ_{Migdal}, and ϕ_{All} respectively. This allows us to split the contributions of an NR/Migdal analysis to a fully combined search.

We perform a grid scan of M_χ in the range of [0.1, 10] GeV/c^2 and $\sigma_{\text{S.I.}}$ in the range of [10^{-47}, 10^{-28}] GeV/c^2. The points are equally spaced in log space for $\sigma_{\text{S.I.}}$ and M_χ. In order to find the parameters resulting in equal ϕ for the combination of all detector configurations, the prior range is fixed to [10^{-2}, 10^2] GeV/c^2 for M_χ and to [10^{-53}, 10^{-27}] cm^2 for $\sigma_{\text{S.I.}}$. This prior volume is 24% larger than the priors considered in the previous section (Tab. 4.2), which would therefore yield equally smaller values of ϕ for properly reconstructed benchmarks because of the denominator in Eq. (4.13). Additionally, the number of live points considered here is only 300 in order to save computation time and the values of ϕ obtained proved to be similar for 1000 live points.

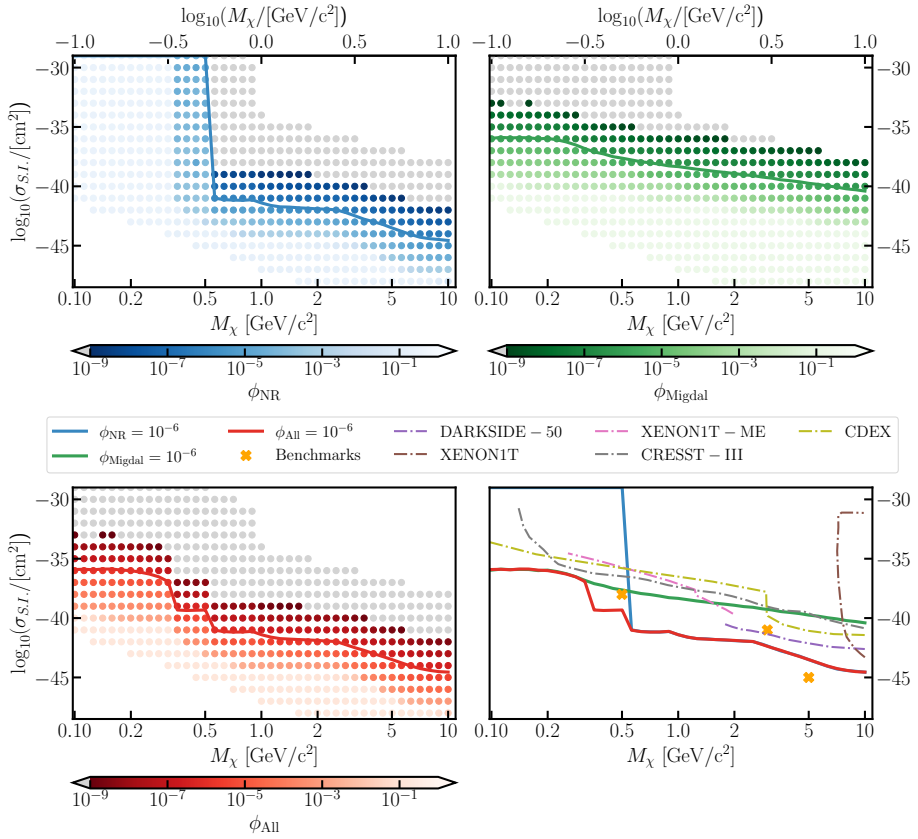

Figure 4.6: Values of ϕ for the combined likelihood using the NR (top left), Migdal (top right), or all (bottom right) experiments, where smaller values of ϕ indicate a tighter 68 % CI. For each of these results, ϕ was interpolated to obtain points where $\phi = 10^{-6}$ (solid lines) which are shown again in the comparison panel (bottom right). This panel also shows the current experimental exclusion 90% CL limits of XENON1T Migdal (ME) [118], XENON1T [4], CRESST [46], CDEX [47], and DarkSide [43]. The benchmark points from Tab. 4.2 are plotted as the orange crosses for reference. While it is tempting to interpret the lines of $\phi = 10^{-6}$ as exclusion limits, this is not correct as elaborated on in the text. The results for each of the m1sses of ϕ_{All} is interpolated to find the corresponding $\sigma_{\mathrm{S.I.}}$ where $\phi = 10^{-6}$ which are the points used in Fig. 4.7. Points where $\phi < 10^{-9}$ are excluded from the color-scales and all set to gray; these points are all well above the current exclusion limits. Points where $\phi \sim \mathcal{O}(10^{-1} - 10^{0})$ correspond to Dark Matter parameters that cannot be reconstructed with the 68 % CI being of similar size as the prior volume.

Fig. 4.6 shows the results of the grid scan for M_χ and $\sigma_{\text{S.I.}}$ for the three combinations of analyses. Whereas the NR analysis (top left panel) constrains the Dark Matter parameters well for $M_\chi \gtrsim 0.5$ GeV/c^2 since ϕ_{NR} is small, it does not have constraining power below this WIMP-mass. The Migdal analyses (top right panel) do have constraining power at these lower WIMP-masses. Compared to the NR analysis, the Migdal analysis achieves similar values of ϕ above $M_\chi \gtrsim 0.5$ GeV/c^2 only at larger $\sigma_{\text{S.I.}}$, meaning that the NR analyses constrain the DM parameters more stringently.

Generally, for small M_χ and $\sigma_{\text{S.I.}}$, $\phi \sim \mathcal{O}(1)$, the combined analyses do not allow constraining the set Dark Matter parameters. For large M_χ and $\sigma_{\text{S.I.}}$, ϕ becomes small as the Dark Matter parameters are reconstructed with good precision.[1]

The combination of all analyses is shown in the bottom left panel, where the contributions of the NR and Migdal analyses are apparent. For $M_\chi \gtrsim 0.5$ GeV/c^2, the combined result follows the result for NR, while it is dominated by the Migdal result for $M_\chi \lesssim 0.3$ GeV/c^2.

To illustrate this further Fig. 4.6 shows for each of the three combinations the value where $\phi = 10^{-6}$. While there is nothing particularly special to the value of $\phi = 10^{-6}$, it corresponds to values of $(M_\chi,\ \sigma_{\text{S.I.}})$ that are close to and below the current 90% confidence level (CL) exclusion limits as illustrated in the bottom right panel of Fig. 4.6. Although it is tempting to interpret the lines where $\phi = 10^{-6}$ in this panel as exclusion limits, they are very different. Exclusion limits are obtained by doing a one-dimensional fit for a fixed mass and show the (frequentist) 90% CL upper limit, while in contrast the lines of $\phi = 10^{-6}$ show where a two dimensional fit would be able to reconstruct the WIMP mass and cross section simultaneously with good precision.

To extract points where $\phi = 10^{-6}$, we interpolate for each mass in Fig. 4.6 to find the corresponding $\sigma_{\text{S.I.}}$. We extract where $\phi = 10^{-6}$ in order to obtain $(M_\chi,\ \sigma_{\text{S.I.}})$-points that are not excluded by experiments at the time of writing [4, 43, 46, 47, 72, 118]. For ϕ_{All} and ϕ_{NR} a jump occurs at $M_\chi \sim 0.5$ GeV/c^2 as this is near the detection threshold of SuperCDMS-NR; for ϕ_{All} this is where the transition starts from NR to Migdal being the largest contribution to the total likelihood.

For the $(M_\chi,\ \sigma_{\text{S.I.}})$-points where $\phi_{\text{All}} = 10^{-6}$, ϕ is also calculated for each of the four separate detector configurations to find the detector configuration contributing most to the likelihood. If ϕ_{All} is lower than the ϕ of individual detector

[1] A significant portion of this parameter space is already excluded by direct detection experiments [4, 43, 46, 47, 72, 118].

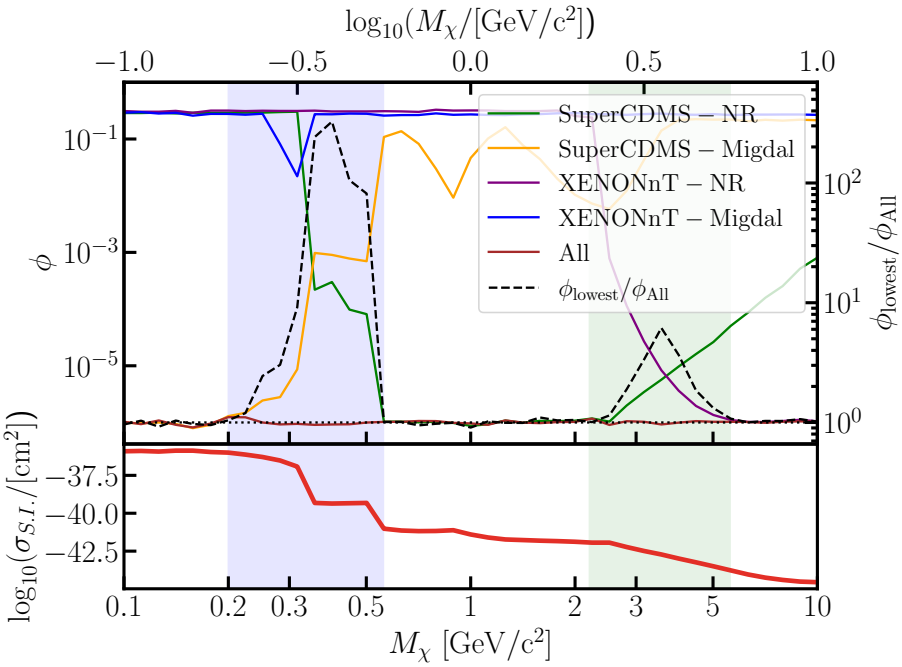

Figure 4.7: Parameter ϕ for the four individual detector configurations and ϕ_{All} (top panel) for the interpolated points from Fig. 4.6. Due to the interpolation, $\phi_{\mathrm{All}} \sim 10^{-6}$ (the horizontal dotted line). The right axis (top panel) shows $\phi_{\mathrm{lowest}}/\phi_{\mathrm{All}}$, the ratio of the lowest ϕ of one of the detector configurations and ϕ_{All}. If $\phi_{\mathrm{lowest}}/\phi_{\mathrm{All}} \sim 1$, the combined likelihood is dominated by the likelihood from one detector configuration as that constrains the parameters well. If $\phi_{\mathrm{lowest}}/\phi_{\mathrm{All}} \gg 1$, this means that the combination of detector configurations is better at constraining the overall likelihood than the individual detector configurations. Two mass ranges with high complementarity are shaded and are discussed in the text. The bottom panel shows the cross section for the masses considered, these correspond to $\phi_{\mathrm{All}} = 10^{-6}$ extracted from the lower left panel of Fig. 4.6.

configurations, this means that the detector configurations are complementary to each other, as in Fig. 4.4.

Fig. 4.7 evaluates ϕ for the individual detector configurations at the points where $\phi_{\text{All}} = 10^{-6}$ in Fig. 4.6. We increase the number of live points back to 1000 from the 300 in considered in Fig. 4.6. Each of the detectors has a mass-range for which it is the most constraining. The contribution of XENONnT-NR to the combined likelihood is largest for $M_\chi \gtrsim 4$ GeV/c^2 since $\phi_{\text{All}} \sim \phi_{\text{XENONnT-NR}}$. Similarly, SuperCDMS-NR is most constraining for $M_\chi \sim [0.5,\ 2.2]$ GeV/c^2, SuperCDMS-Migdal for $M_\chi \lesssim 0.3$ GeV/c^2. We see that the contribution to the combined likelihood from XENONnT-Migdal is small, only achieving values of $\phi_{\text{XENONnT-Migdal}} \sim \mathcal{O}(10^{-2} - 10^{-1})$ since either XENONnT-NR, SuperCDMS-NR or SuperCDMS-Migdal observes higher rates at the DM parameters considered here.

At several intermediate masses we find that the combination of detector configurations yields smaller ϕ values than the individual detectors. For example, between [2.2, 5.6] GeV/c^2, the combination of XENONnT-NR and SuperCDMS-NR yields a smaller value of ϕ. The value of ϕ_{All} is lower than the individual ϕ for the detector configurations of SuperCDMS-NR, SuperCDMS-Migdal and XENONnT-Migdal in the mass range between $\sim [0.2,\ 0.6]$ GeV/c^2 as all three (mostly SuperCDMS-NR and SuperCDMS-Migdal) are constraining the likelihood. In this mass range, a combined analysis will enhance the ability to reconstruct the DM parameters as the ϕ_{All} is $\mathcal{O}\left(10^1 - 10^2\right)$ smaller than the smallest ϕ for these WIMP masses.

4.5 Conclusion

We have investigated the potential of two future detectors, XENONnT and SuperCDMS, to discover light WIMP Dark Matter using an NR or Migdal search or combination thereof. Using a Bayesian framework to probe the Poisson likelihood, the posterior distributions of benchmark points were obtained for WIMP masses of 5, 3 and 0.5 GeV/c^2 and cross section of 10^{-45}, 10^{-41} and 10^{-38} cm^2 respectively. For 5 GeV/c^2 (Fig. 4.2), XENONnT-NR constrained the Dark Matter parameters most, whereas for 0.5 GeV/c^2 (Fig. 4.5) this was done by SuperCDMS-Migdal. At an intermediate mass of 3 GeV/c^2 (Fig. 4.3) the parameter ϕ reduces for the posterior of the combined likelihood by a factor of 54 for XENONnT-NR and 2.1 for SuperCDMS-NR (Fig. 4.4).

More generally, we probed a large parameter space in (M_χ, $\sigma_{\text{S.I.}}$) to find the set of DM parameters where a combined inference of the NR, Migdal, all combined-analyses would be able to reconstruct those DM parameters to an equally sized 68% CI (Fig. 4.6). Using those points, we observed several re-

gions in which one of the detection configurations was outperforming the other detector configurations (Fig. 4.7). Near the detection threshold of XENONnT-NR ($\sim [2.2,\ 5.6]$ GeV/c^2), the combination with SuperCDMS-NR helps in reconstructing the DM parameters. The largest complementarity can be found for SuperCDMS-NR, SuperCDMS-Migdal, and to a lesser extent, XENONnT-Migdal in the mass range between $\sim [0.2,\ 0.6]$ GeV/c^2.

In future work, several effects may be worth exploring. One of the most important parameters for XENONnT is the energy threshold. Experiments are cautious with claiming discoveries near detection thresholds as threshold effects are difficult to model fully. An interesting study would be to take the value of the energy threshold into account as a nuisance parameter in Eq. (4.12). Similarly, as was done previously in Ref. [123], it is worth doing the same for the astrophysical DM parameters. While this has been well-studied for NR searches, their effect on Migdal searches have not been investigated. Finally, the Earth shielding effect [142] should be taken into account when discussing the ability to detect strongly interacting Dark Matter, either at the very small or very large WIMP-masses where large cross sections are not excluded by experimental results.

We have demonstrated the complementarity of two planned Dark Matter direct detection experiments to observe light Dark Matter through a combination of Migdal and standard NR searches. These results highlight in particular that over certain WIMP mass ranges the combination of standard NR and Migdal searches can lead to tighter constraints on the Dark Matter parameters than from either analysis alone.

4.6 Appendix

4.6.1 Energy scales

In this appendix we review several details required for converting the energy scales relevant for the detectors in this work.

Lindhard quenching

The two detectors of interest (SuperCDMS-SNOLAB and XENONnT) both use ionization signals caused by interactions to characterize the type of interaction (ER or NR) within the target volume. In xenon, germanium and silicon, an ER of a given energy will result in more detectable ionization energy than an NR of the same energy due to nuclear quenching [129, 131]. We adopt the following

notation for the ER recoil energy $E_{\rm er}$ and the NR recoil energy $E_{\rm nr}$. In order to compare NR and ER energies it is often useful to calculate how much ionization energy a nuclear recoil would have deposited if the recoil was an electronic recoil: the electronic equivalent energy ($E_{\rm ee}$). Using the Lindhard factor L [129, 131],

$$\begin{aligned} L(E_{\rm nr}) &= \frac{k\ g(\epsilon)}{1+kg(\epsilon)}, \\ g(\epsilon) &= 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon, \\ \epsilon &= 11.5\frac{E_{\rm nr}}{\rm keVnr}\, Z^{-7/3}\,, \end{aligned} \tag{4.14}$$

we can convert $E_{\rm nr}$ to $E_{\rm ee}$:

$$E_{\rm ee} = L(E_{\rm nr}) \cdot E_{\rm nr}\,. \tag{4.15}$$

Here, k is a detector specific parameter and Z the atomic number of the target material. From Eq. (4.14), we can directly see that $L < 1$. The Lindhard factor is used to convert $E_{\rm nr}$ into $E_{\rm ee}$ and vice versa in the methods section (section 4.3).

Following [143], we rewrite Eq. (4.14) to take the atomic binding energy into account for semiconductor materials:

$$\begin{aligned} L(E_{\rm nr}) &= \frac{k\ g(\epsilon')}{1+kg(\epsilon')} = \frac{\epsilon' - \bar{\nu}(\epsilon')}{\epsilon'}\,, \\ \bar{v}(\epsilon') &= \bar{v}_L + C_0{\epsilon'}^{\frac{1}{2}} + C_1 + u\,, \\ \bar{v}_L(\epsilon') &= \frac{\epsilon'}{1+kg(\epsilon')}\,, \\ u &= 11.5\frac{E_{\rm nr}}{\rm keVnr}\, Z^{-7/3}U\,, \\ \epsilon' &= \epsilon - u\,, \end{aligned} \tag{4.16}$$

where U is the energy lost to disruption of atomic bonding, C_0 and C_1 are material specific parameters. For $C_0 = C_1 = 0$ and $U = 0$ keV, Eq. (4.16) reduces to Eq. (4.14). We use the best fit parameters as obtained in Ref. [143]. For Si we take $C_0 = 9.1{\times}10^{-3}$, $C_1 = 3.3{\times}10^{-5}$ and $U = 0.15$ keV. For Ge, we take $C_0 = 3.0{\times}10^{-4}$, $C_1 = 6.2{\times}10^{-6}$ and $U = 0.02$ keV. We assume a value of k of 0.162 for Ge and 0.161 for Si [143] in Eq. (4.16).

SuperCDMS energy-resolution and -threshold

In this appendix, the two relevant energy scales for SuperCDMS are discussed as well as how the values for Tab. 4.1 for the energy-resolution and -threshold are obtained.

There are two energy scales in the SuperCDMS experiment that relate to the ER/NR recoil energy scales [45], namely the phonon energy E_{ph} and the ionization energy E_Q, where the latter is given by:[2]

$$E_{Q,\,\mathrm{nr}} = y(E_{\mathrm{nr}}) \cdot E_{\mathrm{nr}}\,, \tag{4.17}$$

where $y(E_{\mathrm{nr}})$ is the ionization yield, which is set to be equal to $L(E_{\mathrm{nr}})$ for large enough E_{nr}. For ERs, where $y = 1$, we can explicitly rewrite this as:

$$E_{Q,\,\mathrm{er}} = E_{\mathrm{er}}\,. \tag{4.18}$$

Additionally, the phonon energy scale is given by:

$$\begin{aligned} E_{\mathrm{ph,\,nr}} &= E_{\mathrm{nr}} + E_{\mathrm{Luke,\,nr}} \\ &= E_{\mathrm{nr}} \left(1 + \frac{y(E_{\mathrm{nr}}) e\Delta V}{\delta}\right)\,, \end{aligned} \tag{4.19}$$

$$\begin{aligned} E_{\mathrm{ph,\,er}} &= E_{\mathrm{er}} + E_{\mathrm{Luke,\,er}} \\ &= E_{\mathrm{er}} \left(1 + \frac{e\Delta V}{\delta}\right)\,, \end{aligned} \tag{4.20}$$

where the E_{Luke}-term is the signal generated through the Luke-Neganov effect [45], δ is the average energy required to make an electron-hole pair (3.0 eV for Ge and 3.82 eV for Si) and $e\Delta V$ is the work done to move one charge through a crystal, which depends on the bias voltage applied to the detector. The value of $e\Delta V/\delta$ depends on the detector design and is 1.6 (Ge) or 2.7 (Si) for IZIP, and 26 (Ge) or 33 (Si) for HV. As such a relatively modest E_{er} can correspond to a large E_{ph}.

For Migdal, the recoil spectrum is computed in E_{er}. However, in Ref. [45], the resolution and energy thresholds are given in E_{ph}. We need to convert the energy threshold by inverting Eq. (4.20) and substituting the E_{ph} (from Table VIII in Ref. [45]).

[2] Here, we are only considering “bulk events" that have a correction factor $\eta = 1$ in Equations 3 and 4 of Ref. [45].

Similar to the energy threshold, the energy resolution is given in the phonon resolution $\sigma_{\rm ph}$. This resolution is in the order $5-50$ eV. We relate the phonon resolution σ_{ph} to the ER resolution $\sigma_{\rm er}$ using Eq. (4.20):

$$\sigma_{\rm er} = \sigma_{\rm ph} / \left(1 + \frac{e\Delta V}{\delta}\right) . \tag{4.21}$$

For the NR search in SuperCDMS we need to have the relevant energy resolutions and thresholds By inverting Eq. (4.19), we can obtain the values listed for the NR energy threshold in Ref. [45], which are directly used in Tab. 4.1. For the NR case, we need to distinguish between the ionization resolution relevant for the iZIP detectors and the phonon resolution, relevant for the HV detectors. As such, if we treat $\sigma_{\rm ph,nr}$ as the uncertainty on $E_{\rm ph,nr}$, we can propagate the resolution $\sigma_{\rm ph,nr}$ to $\sigma_{\rm nr}$ as:

$$\sigma_{\rm nr} = \frac{{\rm d}E_{\rm nr}}{{\rm d}E_{\rm ph,nr}} \sigma_{\rm ph,\ nr} , \tag{4.22}$$

and resolution of $\sigma_{Q,\,\rm nr}$ to $\sigma_{\rm nr}$ as:

$$\sigma_{\rm nr} = \frac{{\rm d}E_{\rm nr}}{{\rm d}E_{Q,\,\rm nr}} \sigma_{Q,\,\rm nr} , \tag{4.23}$$

where Eq. (4.22) applies to the HV detectors and Eq. (4.23) to the iZIP detectors. We solve Eqs. (4.22-4.23) numerically. From Eqs. (4.22-4.23), we see that the energy resolution $\sigma_{\rm nr}$ has an energy dependence through the ionization yield $y(E_{\rm nr})$ even though $\sigma_{\rm ph,\ nr}$ and $\sigma_{\rm Q,\ nr}$ are assumed to be energy independent.

Chapter 5

The XENONnT DAQ

As experiments grow in size, so do their number of photosensors and requirements of Data Acquisition systems (DAQs). At the same time, experiments aim to further lower their energy thresholds and improve their understanding of the detector response at the lowest energies (see Chapter 4 and section 6.4). To facilitate the understanding of very small signals, the XENONnT DAQ was designed around a triggerless paradigm where all signals from individual photomultiplier tubes (PMTs) that pass a digitization threshold are stored. Additionally, as experiments grow in sensitivity, their need to distinguish neutron induced backgrounds increases, to which end the XENONnT detector has a new neutron veto (see section 3.3). The DAQ is configured to digitize the data of the Time Projection Chamber (TPC), neutron veto, and muon veto either by running as a single or as independent operational subsystems. Finally, the increased data rate - due to more PMTs and very low drift fields (subsection 6.3.1) - requires fast processing, which in the case of XENONnT is so fast that it is able to provide fully reconstructed data to analysts within $\mathcal{O}(10\,\mathrm{s})$, thereby allowing live monitoring of the data. In this chapter, we discuss the design, commissioning, and performance during Science Run 0 (SR0) of the XENONnT DAQ as published in Ref. [3].

The Triggerless Data Acquisition System of the XENONnT Experiment

Pre-published in: arXiv preprint arXiv:2212.11032 (2022)
Publication date: Dec 21, 2022

Available online at:

- https://doi.org/10.48550/arXiv.2212.11032
- https://arxiv.org/abs/2212.11032

Reference [3].

Abstract

The XENONnT detector uses the latest and largest liquid xenon-based time projection chamber (TPC) operated by the XENON Collaboration, aimed at detecting Weakly Interacting Massive Particles and conducting other rare event searches. The XENONnT data acquisition (DAQ) system constitutes an upgraded and expanded version of the XENON1T DAQ system. For its operation, it relies predominantly on commercially available hardware accompanied by open-source and custom-developed software. The three constituent subsystems of the XENONnT detector, the TPC (main detector), muon veto, and the newly introduced neutron veto, are integrated into a single DAQ, and can be operated both independently and as a unified system. In total, the DAQ digitizes the signals of 698 photomultiplier tubes (PMTs), of which 253 from the top PMT array of the TPC are digitized twice, at $\times 10$ and $\times 0.5$ gain. The DAQ for the most part is a triggerless system, reading out and storing every signal that exceeds the digitization thresholds. Custom-developed software is used to process the acquired data, making it available within $\mathcal{O}(10\,\mathrm{s})$ for live data quality monitoring and online analyses. The entire system with all the three subsystems was successfully commissioned and has been operating continuously, comfortably withstanding readout rates that exceed $\sim 500\,\mathrm{MB/s}$ during calibration. Livetime during normal operation exceeds 99% and is $\sim 90\%$ during most high-rate calibrations. The combined DAQ system has collected more than 2 PB of both calibration and science data during the commissioning of XENONnT and the first science run.

5.1 Introduction

A variety of experiments use time projection chambers (TPCs) filled with liquid noble elements (usually xenon or argon) in the search for Weakly Interacting Massive Particle (WIMP) dark matter and rare radioactive decays [71, 144–147]. While the details of each detector differ, common design features include arrays of photosensors at the ends of the drift region and accompanying readout systems.

Interactions in a dual-phase TPC are observed via two processes: scintillation and ionization. When a particle interacts with either the electrons or nucleus of a target atom, prompt scintillation light and liberated electrons are produced, resulting in two signals. Two arrays of photosensors, photomultiplier tubes (PMTs) in XENONnT, are located above and below the cylindrical drift region to capture these signals. The detected scintillation light is referred to as the "S1" signal, while the electrons are drifted under an external electric field towards the liquid-gas interface. When the electrons reach this interface, a much stronger electric field extracts them from the liquid and causes electroluminescence in the gas, producing additional proportional scintillation that is detected and referred to as the "S2" signal. The time between the S1 and S2 signals, which is the drift time of electrons in the liquid phase, as well as the pattern of illumination on the top PMT array caused by the S2, are used to reconstruct the interaction vertex in the detector. The S2 is typically much larger than the S1, and the relative sizes of these two signals are used to discriminate between electronic recoil (ER) and nuclear recoil (NR) interactions.

The XENON collaboration has operated a series of increasingly larger dual-phase xenon TPCs at the INFN Laboratori Nazionali del Gran Sasso (LNGS) in central Italy for nearly two decades, probing WIMP-nucleon cross-sections down to $4.1 \times 10^{-47}\,\mathrm{cm}^2$ (for a $30\,\mathrm{GeV/c}^2$ WIMP) [4]. The latest is the TPC of the XENONnT detector, containing 5.9 t in its active target mass and is expected to be sensitive to spin-independent WIMP-nucleon cross-sections down to $1.4 \times 10^{-48}\,\mathrm{cm}^2$ (for a $50\,\mathrm{GeV/c}^2$ WIMP) [148].

5.2 From XENON1T to XENONnT

The upgrade from XENON1T to XENONnT saw the TPC increase in size from ∼1 m in diameter and length to ∼1.3 m and ∼1.5 m, respectively, to accommodate a larger target mass. This increase in TPC size was accompanied by a corresponding increase in the number of PMTs to 494, 253 in the top array in the

gas phase and 241 in the bottom array in the liquid below the target. This constitutes a two-fold increase from XENON1T where the TPC was instrumented with 248 PMTs.

As detectors continue to grow in size, the maximum drift time (the drift length) of electrons in the TPC grows in accordance. A lower drift field of 23 V/cm in the first science run (SR0) of XENONnT [149] compared to 81 V/cm–120 V/cm in XENON1T [4] means a further increase in the drift length. The need to store and read out one continuous drift length of data, which can exceed 2 ms, is alleviated by the firmware used by most of the readout hardware. This digital pulse processing with dynamic acquisition windows (DPP-DAW) firmware was developed in collaboration with CAEN [150] for XENON1T [151], and an updated version was used for XENONnT. It affords many useful techniques such as baseline suppression or zero length encoding (ZLE), dynamically-sized acquisition windows that automatically extend as long as the input is above the digitization threshold, and the independent and continuous readout of each channel. The digitization thresholds are set relative to dynamically-calculated baselines. However, increased drift time leads to an increased temporal width of S2 signals as the freed electrons diffuse over a larger amount of time. This increased temporal width directly increases data rates as signals remain above threshold for longer. Thus, new challenges arise as drift times increase and S2s become longer.

A new active neutron veto sub-detector was built to suppress the NR background from radiogenic neutrons generated through spontaneous fission and alpha-neutron reactions, as these mimic WIMP-induced signatures. It is made of an octagonal structure (3 m-high and 4 m-wide) placed inside the water tank around the cryostat that houses the TPC and is optically separated from the existing muon veto [148]. To improve the neutron detection efficiency, the water will be loaded with gadolinium (Gd). A total of 120 Hamamatsu 8" high quantum efficiency PMTs with low-radioactivity windows are placed along the lateral walls. Neutrons that leave the TPC volume are moderated by the water around the cryostat before being captured Gd or H. A gamma-ray cascade with total energy of about 8 MeV is generated for capture of Gd and a single 2.2 MeV gamma is emitted in the case of capture on H. The gammas in water, mainly through Compton scattering, are converted into electrons and ultimately into Cherenkov photons. Monte Carlo studies indicate that the neutron veto is expected to reduce the total NR background by a factor of six [148]. Dedicated hardware was implemented to manage the neutron veto data readout. The water Cherenkov muon veto surrounding the cryostat, instrumented with 84 PMTs, is otherwise largely unchanged from XENON1T [151, 152].

5.2.1 General DAQ upgrades

The XENONnT data acquisition (DAQ) system is an evolution of that which was successfully used for XENON1T [151]. Many aspects of the system have received modifications and improvements based on the XENON1T system and the experience of operating it. An overview of the new system design is shown in Fig. 5.1.

One challenge of the general DAQ design is the range of sizes and shapes of signals the system must handle. The muon veto and neutron veto are Cherenkov detectors, registering photon signals over an interval of at most $\mathcal{O}(1\,\mu s)$. The TPC, in contrast, must record both S1 and S2 signals. S1s can be very small, down to a single photon, and are very fast, lasting up to $\mathcal{O}(100\,ns)$. S2s are much larger, potentially millions of photons, and can have temporal widths exceeding $100\,\mu s$. Representative signals in all three subsystems are shown in Fig. 5.2.5.

While XENON1T used the Phillips Scientific 776 amplifiers with dual $\times 10$-gain outputs, XENONnT uses custom dual-gain ($\times 10$ and $\times 0.5$) amplifiers developed at the University of Zurich [153]. The low-gain signals from the top PMT array are digitized separately from the high-gain signals to try to improve energy and position reconstruction involving large signals that otherwise saturate the input stage of the digitizers. The low-gain signals from the bottom array are summed together and used by the high-energy veto, discussed in subsection 5.3.3. This results in the XENONnT TPC effectively having three times the number of PMT readout channels of the XENON1T TPC (747 compared to 248), as the number of PMTs is doubled and half are read out twice.

5.2.2 Triggerless data streams

The XENON1T TPC utilized a triggerless readout and a central buffer built from a MongoDB database [154]. Triggering software ran live over this database and determined "events" which were written to disk for later analysis, the remainder of the data was deleted. While this paradigm was successful at realizing a very low effective trigger threshold, the estimation of some backgrounds was more difficult due to the forced truncation of events after a certain maximum duration. Additionally, this database would not scale effectively to match the increased load foreseen by the demands of a larger system. The solution is to forego the software trigger and save all the data, leaving the determination of events to much later in the data processing pipeline.

The removal of all triggers except the per-channel digitization threshold does not lead to significantly increased storage requirements. The primary driver

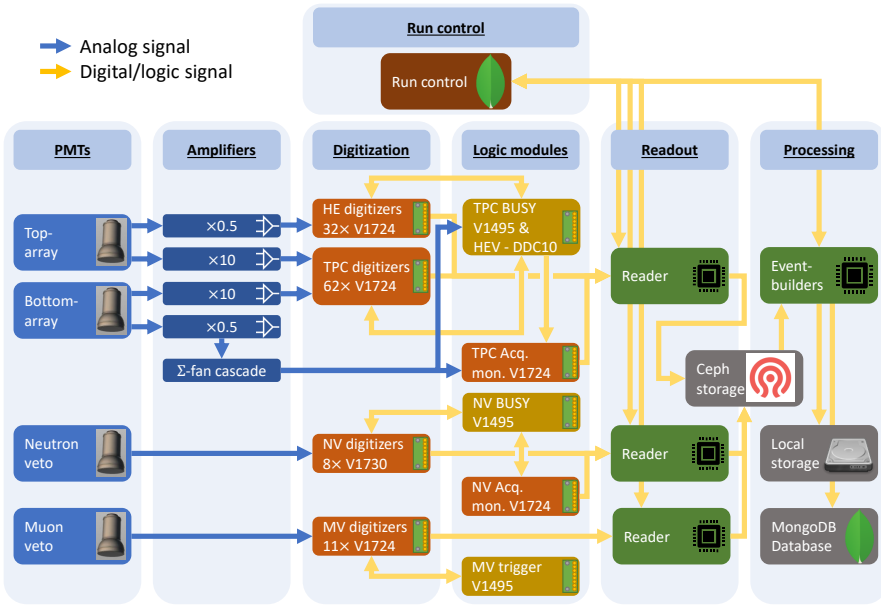

Figure 5.1: The XENONnT DAQ layout at several stages. The TPC PMT signals are amplified before being digitized by CAEN V1724 modules. The neutron veto (NV) and muon veto (MV) PMT signals are digitized by the associated V1730S and V1724 digitizers. The top array is digitized twice at ×10 and ×0.5 gain, the signals from the latter constitute the high energy system. The ×0.5 gain signals of the bottom array are fed into the sum-signal fan cascade (Σ-fan cascade). The busy (V1495) and acquisition monitor (V1724) modules handle the busy logic and monitor the system performance. The reader servers read out the digitizers and write the data to a common (Ceph) storage disk. The eventbuilder servers do the processing of the reader data which get written to local storage where it can be distributed to other storage sites. A portion of the processed data is also written to the MongoDB database where it can be accessed for online data monitoring. The control of the DAQ is achieved through the website that communicates commands (run control) to the readers/eventbuilders.

of data rate is not PMT dark counts or other small signals seen by only a few channels, rather it is large S2 signals that can remain above threshold for

$\mathcal{O}(10-100\,\mu s)$, and are seen by a large number of channels. For instance, in typical conditions during SR0, S2s from single electrons account for 30% of all reconstructed signals but only 2% of the data volume. In contrast, for very large S2s these values are reversed, accounting for 2% of the reconstructed signals but 30% of the data volume. Any trigger would be configured to save the large S2s, so additionally saving everything else (mostly S2s) does not represent a significant increase on the requirements of long-term storage. To support this, work was done studying data formatting and data compression, and a storage format was chosen that compresses more efficiently than the XENON1T storage format. Further, the removal of the software trigger eliminates the requirement for a database that can act as a base for fast triggering software, so the readout processes write data directly to high-speed disks in a continuous stream.

5.2.3 Fast data processing & immediate data availability

In addition to the hardware upgrades from XENON1T to XENONnT, the readout and processing software was also upgraded. To handle the continuous data stream of roughly three times the number of channels, the processing framework *PAX* [155] was replaced by the generic framework *strax* [156] and implemented for XENONnT in *straxen* [157]. Strax and straxen are referred to as strax for simplicity. Strax is written in python and was initially based on a re-write of PAX with a different memory model. It uses packages from the scipy stack [158, 159], just-in-time compilation (numba) [160], and a tabular data format to allow for fast processing by exploiting autovectorization. While PAX achieved processing speeds of $\mathcal{O}(100\,\mathrm{kB/s/core})$, strax can process data at rates of $\mathcal{O}(10-100\,\mathrm{MB/s/core})$. Strax achieves its highest per core processing speeds when running on only a few cores but also allows parallelization to tens of cores, albeit at lower per core performance. For the data rates observed during SR0, including the associated calibration periods where a higher rate was expected, the processing time was much lower than the data collection time. The strax framework will be further elaborated on in section 5.4.

Strax does complete online reconstruction of all the data within $\mathcal{O}(10\,s)$ after the PMTs detect light. This allows the detector performance and stability to be monitored with high-level data without the need for selections or triggers. To enable remote, online access to the data while it is being processed at the DAQ, several types of data are uploaded to the MongoDB database in a dedicated collection. This database is accessible from outside LNGS, such that it can be retrieved from anywhere. During normal operation, these data are available online within $\mathcal{O}(30\,s)$. With online data access the performance can be monitored

using fully reconstructed data. This is especially useful for stability checks, as well as detailed feedback on operations with rapidly changing conditions such as calibrations or changing field configurations. These online data are blinded, and only when purposely unblinded and reprocessed the science results [149] are obtained.

5.2.4 Neutron veto DAQ

The goal of the neutron veto in XENONnT is to detect the capture process of those neutrons responsible for NR background events, which can mimic the interaction of a WIMP. A neutron detection tagging efficiency greater than 85% is desired [148]. Since the expected Cherenkov signal in the case of neutron capture by H is of about 20 PE in total, it is important to have a high detection efficiency for each photon. To achieve such a high efficiency in a trigger-based DAQ architecture it would have been necessary to reduce the number of coincident PMTs that form the trigger. This, in turn, would have lead to an increased number of triggers and acquired data, making it challenging for the DAQ readout.

Therefore, the neutron veto DAQ is designed around a triggerless data collection scheme like the TPC. Its ability to provide both the pulse shape and the timestamp of each PMT signal supports data collection with fully independent channels without the use of a global trigger, typically based on channel multiplicity. As will be described in section 5.4, the event building is done in software after data acquisition, where timestamps and coincidences between PMT signals are used to define events. This architecture based on a readout system of independent channels allows the acquisition of all the PMT signals above the digitization threshold and the lowering of the energy threshold.

PMT characteristics such as dark rate, afterpulsing and timing resolution are essential for the choice of front-end electronics. In particular, the dark rate puts a limit on the detection of very small signals, and can be used to estimate the accidental coincidence rate with a defined number of PMTs within a specific time window. Operating with a threshold of 0.5 photoelectrons (PE), the measured PMT dark rate during detector commissioning was about 0.96 kHz, generating an accidental coincidence rate that exceeded 4 kHz for a 2-fold coincidence between two random neutron veto PMTs. In addition, the materials in the sub-detector itself (PMTs and stainless steel structure) induce events that mimic NR signals in the neutron veto of $\mathcal{O}(100\,\mathrm{Hz})$.

To efficiently tag neutron events, the electronics must be able to acquire signals ranging between 0.5 PE and $\mathcal{O}(100\,\mathrm{PE})$, requiring a wide dynamic range.

The small signals last about $100-200\,\text{ns}$, resulting mainly from the dark rate, and define the lower limit for the neutron veto data throughput. In contrast, gamma and beta particles from materials radioactive decays with a typical rate of $\mathcal{O}(100\,\text{Hz})$ exhibit waveforms that last up to $10\,\mu\text{s}$ (considering signals and associated afterpulses) in many channels, requiring a much higher data collection rate. Therefore, the readout electronics must be able to handle extremely different time acquisition windows with the presence of sharp peaks of data rate.

In addition, the fast response (few ns) of PMTs used by the neutron veto to acquire single photoelectrons requires a fast waveform digitizer for signal sampling. This high time resolution is necessary to efficiently separate neutrons produced close to and far from the TPC cryostat, where the former are the primary target.

5.2.5 Three integrated DAQ subsystems

One requirement for the XENONnT DAQ system was that the three DAQ subsystems (TPC, muon veto, and neutron veto) should be able to operate both independently and as one combined system. While both the TPC and muon veto in XENON1T used the same $50\,\text{MHz}$ clock signal, there was no synchronization of the start signals issued to the two readout systems, so there was some variation in reconstructed timestamps between the two detectors, and data from the two subsystems were analyzed separately. The trigger signal of the muon veto was recorded in one of the TPC's digitizers, but this did not provide the equivalent timestamp in the muon veto's data stream, thus viewing the corresponding event as observed in the muon veto required analysts to perform additional steps.

To ameliorate this, the XENONnT DAQ was designed to allow for the start signal from one subsystem to be issued directly to one or both of the others, essentially combining them into one and ensuring that timestamps recorded in one can be directly compared to those from another. Subsystems can be combined or "linked" together as determined by the requirements of the data being taken, or can operate independently. Subsystems operating in linked modes are controlled as a single operational unit, and the data they record are combined at the readout level and processed together to facilitate handling and analysis.

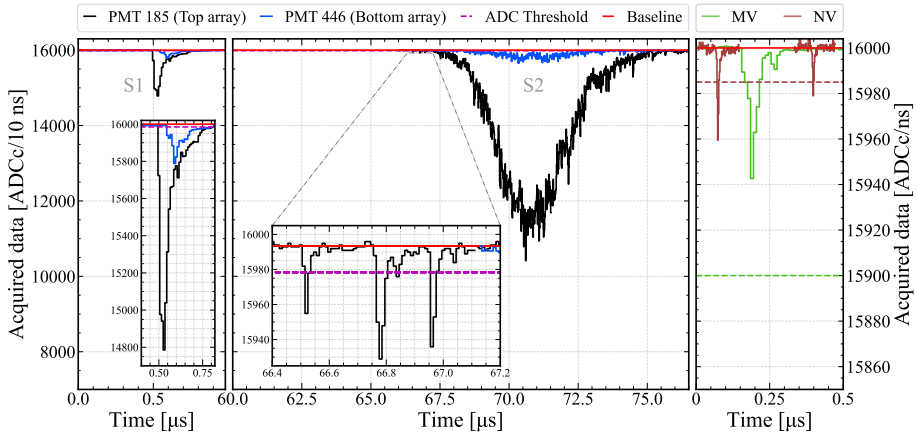

Figure 5.2: An illustration of the variety of signals read out by the XENONnT DAQ in ADC counts (ADCc). The left and middle panels showcase raw signals from a selection of two TPC PMTs in a single event. The inset in the left panel zooms in on the S1 to emphasize its narrow width and short risetime, contrasting with the wide S2 in the middle panel. An inset in the middle panel zooms in on minute signals in the leading edge of the S2 waveform. In both panels, the black lines correspond to signals from a PMT in the top array, and the blue from the bottom array. The red and purple dashed lines represent the baseline and digitizer threshold, respectively. The rightmost panel shows signals that were recorded by individual channels in the muon veto (green) and neutron veto (dark red) DAQ subsystems. These signals are not correlated with the ones showcased for the TPC channels. Typical muon veto and neutron veto thresholds are depicted with dashed lines of matching color. The higher sample rate of the digitizers employed by the triggerless neutron veto subsystem is clearly visible in comparison to the triggered muon veto. In this case, a relative baseline is shown in red for illustration purposes only. As explained in section 5.3.2, this muon veto signal is read out, despite not being above threshold, because of the hardware coincidence trigger during this time interval.

5.3 Data Acquisition & Readout

The data acquisition is organized in two broad schemes following the system used in XENON1T. At the greatest scope are so-called "science runs" which represent a period of months or even years with a targeted science objective

where the detector conditions are held constant. For the daily operation of the experiment, the organizational unit called "runs" is used, where each run represents a continuous period of a few minutes up to a few hours using a set of configuration options that remain constant for the duration of the run.

The three DAQ subsystems rely predominantly on commercially-available analog front-end electronic modules supported by custom hardware. The firmware and software include both custom components and some provided by CAEN. All commercially available CAEN products are marked with their model number throughout this work, the reader is referred to the company's website for more details and manuals [150]. All DAQ hardware is installed within eight racks located in the DAQ room on the first floor of the XENON service building in Hall B of LNGS. An air conditioning system provides cooling for the electronics, maintains a constant air temperature in the room, and reduces the collection of dust.

5.3.1 Analog electronics

The PMTs of the TPC are powered by an array of multi-channel CAEN A7030LN, A1536LN and A1535LN high-voltage supplies. Five CAEN A7435SP high-voltage (HV) boards supply power to the neutron veto PMTs, and four CAEN A1535SP boards power the PMTs of the muon veto. The above HV boards are housed in separate CAEN SY4527 Universal Multichannel Power Supply System crates. HV boards are known to produce high-frequency switching noise. Hence, before being supplied to TPC PMTs the output from each HV board is passed through a custom filter box. Within each filter box, every HV channel line goes through a low-pass filter, removing electronic noise with frequencies greater than $\sim$250 kHz.

Before installation, all TPC signal cables were grouped based on the location of their corresponding PMTs in the arrays and assigned to specific hardware modules. The resulting cable map ensures an equal distribution of the data load on readout electronics and was used as a guide throughout the hardware installation process. Signals from both PMT arrays are passed through custom amplifiers as mentioned in subsection 5.2.1. As illustrated in Fig. 5.1, both the high- and low-gain signals from the top PMT array are propagated to dedicated groups of digitizers. However, only high-gain signals from the bottom PMT array are passed to digitizers. The low-gain signals from the bottom array are summed up using a cascade of linear fan-in/fan-out modules and are used by the high energy veto (HEV). Lastly, each DAQ subsystem hosts a range of logic modules that are used for distributing initialization and trigger signals.

5.3.2 Digital electronics

For time synchronization across its subsystems the XENONnT DAQ relies on a CAEN DT4700 clock generator module. Its 50 MHz low-voltage differential signaling (LVDS) outputs are propagated via seven shielded custom-manufactured cables to the first digitizer in each VME crate. Shielding the clock-carrying cables reduces the amount of external noise that can be injected into the cables, improving the stability of the clock signal. The propagated signals are then distributed within each VME crate by shorter clock cables from digitizer to digitizer, ensuring the temporal synchronization of the entire DAQ system. Time offsets in these clock chains were manually calibrated out, securing synchronization well below the digitizer temporal resolution.

Additionally, a GPS timing module [161] is used to distribute a 0.1 Hz trigger signal to dedicated digitizers in each DAQ subsystem. Each trigger is associated with a GPS timestamp (accurate to $\sim$10 ns), providing another layer of time synchronization within the DAQ. The same signal can also be used for absolute time synchronization with other experiments.

TPC

The core of the TPC readout is formed by 95 CAEN V1724 digitizers running the DPP-DAW firmware, an updated version of what was used in XENON1T [151]. The V1724 is an 8-channel board featuring a sample rate of 100 MHz, a dynamic input range of 2250 mV (input impedance 50 Ω) with 14-bits of resolution, and an input bandwidth of 40 MHz. Of these digitizers, 62 read the 494 high-gain signals from the top and bottom arrays, 32 read the 253 low-gain signals from the top PMT array, and 1 acts as the TPC's Acquisition Monitor detailed in subsection 5.3.4. The boards are distributed across five VME crates and connected to readout servers via daisy-chained optical links. Most optical links contain the maximum of 8 digitizers, while the acquisition monitor is read out via its own dedicated optical link to ensure it never goes busy.

One CAEN V2718 crate control module is used as a synchronizing module to produce the sync/start/stop digital input (S-IN) signal that begins and ends the acquisition. This signal is distributed to all digitizers via logic fans, with the signals all reaching their respective digitizers with a spread of $<$4 ns. Additionally, this module provides gate logic signals that control the propagation of the S-IN signal to the muon and neutron veto digitizers that are activated during linked-mode operation. A periodic external trigger signal can be generated by this module, which is distributed both to all digitizers and an external LED

pulser used to calibrate the response of the PMTs. The third type of module is a general-purpose CAEN V1495 board running custom firmware, which manages the busy subsystem detailed in subsection 5.3.3.

Finally, two NIM crates hold the logic fan modules used to distribute the S-IN and trigger/veto signals to all TPC digitizers, as well as a gate module, a NIM-TTL level converter, and a delay generator. These latter two are used to connect and synchronize the TPC DAQ with the LED calibration system.

Muon veto

The muon veto readout is unchanged from XENON1T as described in [151], though some additional connections were made between this subsystem and those of the TPC and neutron veto. Eleven V1724 digitizers with the default ZLE firmware form the readout system, although the zero length encoding features are not used. Three optical fibers are used to read out these digitizers. A V2718 module provides the S-IN signal for these digitizers during unlinked operation. A CAEN V976 unit serves as a logic fan to distribute both this S-IN signal and that of the TPC during linked operation to all muon veto digitizers. A V1495 board acts as a programmable trigger unit, allowing the user to specify both the number of participating channels and the coincidence window necessary to generate a hardware trigger.

Neutron veto

The 120 PMTs of the neutron veto are connected to the readout electronics and HV system located in the DAQ room by means of 30 m coaxial cables with separate grounding for signal and high voltage cable lines. A custom-made patch panel mounted on the back side of the neutron veto rack gathers HV lines in one section and signal lines in another. Signal lines are directly connected to the front-end electronics via a panel feedthrough. HV lines are low-pass filtered to reduce high frequency noise ($\gtrsim$ MHz) and connected to the CAEN A7435SP HV boards.

To take advantage of the fast response of the PMTs and to efficiently reconstruct the fast component of Cherenkov photons in the neutron veto subdetector, eight CAEN V1730S new generation digitizers are used to acquire PMT signals. Each V1730S board is a VME 6U module housing a 16-channel 14-bit 500 MHz flash ADC. The input dynamic range can be set to either 2 V or 0.5 V on single ended MCX coaxial connectors. During commissioning and SR0, the 2 V dynamic range was used. The input section is 50 Ω-coupled and

feeds a programmable gain amplifier to select the suitable analog range. In case all the buffer memory is filled, a busy condition occurs and a logic module inhibits the data acquisition for all the boards (as described in subsection 5.3.3). The V1730S digitizers are operated with the DPP-DAW firmware like the TPC digitizers. An exemplary neutron veto waveform is shown in Fig. 5.2.5.

The V1730S module is also able to work with a common global trigger, either coming from the digital input external trigger (TRG-IN) input or a coincidence trigger. In particular, the external trigger mode is used by the neutron veto system during calibration. A V2718 board hosted in the VME crate generates several control signals (mainly the start-of-acquisition and calibration signals) that are subsequently distributed to the digitizers via logic fan-in/fan-out modules. Two additional boards are hosted in the neutron veto crate: a V1495 to manage the V1730S busy signals and provide the veto signal, and a V1724 digitizer that serves as an acquisition monitor.

The neutron veto digitizers are connected to a readout server via two optical links; one daisy-chains the V1730S digitizers while the second is for the V1724 acquisition monitor. In order to synchronize all the digitizers in the neutron veto DAQ and limit the clock uncertainties to below $\sim$1 ns, an external common clock reference feeds all the modules. The V1724 digitizer receives the common 50 MHz clock signal (see subsection 5.3.2), which is then upconverted to 62.5 MHz via a phase locked loop device and propagated through the V1730S boards. Lastly, several auxiliary electronic modules are hosted in a NIM crate, managing the distribution of calibration triggers and run start signals.

5.3.3 Busy & high-energy veto

The V1724 and V1730S digitizers have a limited on-board memory buffer for storing data between digitizing and readout, amounting to 1 MB/channel and 10.24 MB/channel, respectively. If incoming data accumulates in the digitizer's memory buffer faster than it is read out, the buffer will become full and the digitizer will no longer be able to acquire new signals, rendering it *busy*. To ensure the integrity of individual events the triggerless TPC and neutron veto DAQ subsystems employ a hardware-based *busy veto*. Each subsystem hosts a general-purpose V1495 board, equipped with field-programmable gate array (FPGA) firmware developed in-house. When a digitizer enters the busy state it emits an LVDS signal via a pair of connectors on its front panel. These signals are propagated via ribbon cables from each digitizer in each subsystem to its respective V1495 module.

Whenever the V1495 recognizes that any of the digitizers emits an LVDS busy signal, it outputs a veto NIM signal for a fixed duration. This veto signal is distributed to all the digitizers within the subsystem, inhibiting data acquisition for 1 ms or until none of the digitizers are busy. Within the FPGA firmware, busy intervals are assigned with start and stop NIM signals. These are also output from the V1495 board and propagated to the relevant acquisition monitor digitizer, as explained in subsection 5.3.4. The TPC V1495 board has a more advanced version of this firmware. Besides being responsible for the busy veto, it is also capable of generating an artificial periodic veto, with user controlled duration and frequency. During detector commissioning, the water tank was empty and the detector was not shielded from radiation in the experiment hall. This periodic hardware-induced veto allowed the TPC DAQ to handle high background rates in addition to taking ^{83m}Kr calibration data. Additionally, the V1495 module performs several other important functions. In the TPC DAQ it collects and propagates the HEV signal, and assigns it with start and stop NIM signals that are read by the TPC acquisition monitor. In both the neutron veto and TPC subsystems the V1495 board is also responsible for the propagation of the LED trigger to the digitizers during LED calibration.

The TPC DAQ also employs a hardware veto to reduce the load on the system during acquisition of high-rate calibration data. This HEV was developed based on a commercially-available multipurpose digital pulse processor DDC-10 from SkuTek [162]. It hosts a variety of chips and daughter cards on a Black-VME S6 motherboard, including a Spartan-6 FPGA and a 100 MHz, 10-channel, 14-bit ADC. Mounted on the FPGA is custom-developed firmware whose main goal is to identify and veto high-energy S2 signals. The HEV digitizes the analog sum signal from the low-gain TPC channels of the bottom array, determining the risetime, width and integral of acquired signals. If any identified S2 exceeds predefined threshold parameters, the HEV issues a 3 ms veto NIM signal. To provide the HEV module with enough time to make the veto decision, data readout is delayed within the TPC digitizers by 10 µs. The veto signal generated by the HEV is propagated to the TPC V1495 module, from where it is distributed to each TPC digitizer.

Drift field conditions in the TPC during SR0 produced broad S2 signals with widths greater than $\mathcal{O}(10\,\mu s)$, which were found to have the largest contribution to the DAQ rate. It is difficult to identify and characterize the shape parameters of such signals within the 10 µs time limit. Hence, the HEV firmware has an additional operation mode, whose purpose is to veto low-amplitude high-width S2 signals that might last for $\mathcal{O}(10\,\mu s)$. In this mode, if the HEV is not able to determine the width and the risetime of the signal within $\mathcal{O}(5\,\mu s)$ and the signal's

amplitude is still above the HEV threshold it will consider the signal to be a high-width S2, and will issue a veto. The aforementioned HEV operation modes can be utilized separately or run in parallel. A schematic view of the hardware-based veto systems described above is shown in Fig. 5.1. The operation of the HEV results in raw data reduction at the readout stage of up to 40%, depending on the utilized HEV settings. Throughout SR0 the HEV was utilized during AmBe and ^{220}Rn calibration data taking.

5.3.4 Acquisition monitors

The TPC and the neutron veto DAQ subsystems each host a dedicated V1724 digitizer, whose aim is to collect information about the status of the DAQ itself and the operation of its hardware veto modules. Both the TPC and neutron veto acquisition monitors receive the start and stop NIM signals from their respective V1495 veto modules, which indicate the boundaries of busy veto intervals. Additionally, these digitizers also acquire the 0.1 Hz NIM synchronization signal from the GPS module [161]. Uniquely for the TPC, its acquisition monitor also digitizes the same analog sum waveform signal that is seen by the HEV module. To prevent the TPC acquisition monitor from ever going busy a relatively high threshold of 100 ADCc (ADC counts) or 14 mV is set on this channel. Furthermore, acquisition monitors are also excluded from the busy veto distribution scheme.

Acquisition monitor data are read out identically to the rest of the digitizers and incorporated into the overall data processing chain. These data are then used to diagnose the performance of the busy and HEV systems, and to determine the deadtime they induce. The measured deadtime under several operational modes is discussed in subsection 5.6.2. Moreover, the same data are used as a basis for a data quality cut. The cut removes any events that could be misreconstructed due to missing information as a result of their proximity to a busy or a HEV veto interval. The cut decreases the livetime and is accounted for in the exposure rather than the cut acceptance [149].

5.3.5 Servers & software

Five server computers are responsible for the readout of all the digitizers, three for the TPC and one each for the muon veto and neutron veto. Two additional servers provide backup capacity. The TPC readout servers each have four 960 GB write-intensive solid state drives which are configured together as a Ceph cluster [163] to form a single high-speed buffer disk with approximately

10 TB of capacity that is accessible from all servers within the DAQ network. While replication is possible using Ceph, it is not necessary for a short-term buffer disk, so the configuration is equivalent to RAID0 (data striping) to provide the highest access speeds. This buffer disk can sustain simultaneous read and write operations from multiple sources at rates exceeding 1 GB/s. Data are stored on the Ceph buffer from the start of acquisition until the live processing for that run has successfully concluded, which is typically one or two hours, so very little data are lost in the event of disk failure. The combined data rate from the three subsystems during science data taking is approximately 40 MB/s, so the disk can potentially buffer data for a considerable amount of time in case of issues in the live processing.

Each readout server is equipped with at least one CAEN A3818 PCIe interface card. Each A3818 supports up to 4 optical fibers, with each fiber capable of daisy-chaining up to 8 digitizers and supporting a maximum data throughput of 80 MB/s–90 MB/s. Digitizers and optical links are distributed to approximately balance the load on each of the readout servers.

The readout servers all run the *redax* software package [164], which copies data from the digitizers and transforms it from the digitizer-native format into one compatible with the strax data processing package [156, 157]. Data are read from digitizers in block transfers via the CAENVMElib C++ library, where each optical fiber is read out exclusively by a dedicated readout thread. A round-robin technique is used, where each board on an optical fiber is successively polled. When a digitizer has data available for readout, block transfers are performed until all data stored on that digitizer have been read into the server's memory. The readout threads then transfer data asynchronously to processing threads, where the binary format transformation is performed. Each processing thread periodically compresses its buffered output data and writes it to the Ceph buffer in fixed-time intervals called chunks following the chunking paradigm in strax as described in section 5.4. Chunks are labeled with the name of the readout process, the chunk number, and also the ID of the thread that wrote that chunk, which acts as a unique identifier. Additionally, redax is responsible for programming the digitizers in preparation for each run via configurations it obtains from a central database. Redax also writes status snapshots to this database once per second, including quantities such as the current state of that instance of redax, the amount of data currently buffered in memory, and the data rate for each channel of each digitizer being read out.

Six additional servers, called the *eventbuilders*, are responsible for the live processing. In case of high data rates or unlinked operation, when the three DAQ subsystems run independently, multiple hosts can process data simultane-

ously. For low data rates only a single host is required for the processing. Three eventbuilders are PRIMERGY RX2540 M4 Fujitsu servers with two Intel® Xeon® Gold 6128 CPUs at 3.40 GHz and 202 GB of RAM each. Additionally, there are three backup PRIMERGY RX2540 M1 Fujitsu servers with two Intel® Xeon® E5-2660 v3 CPUs at 2.60 GHz and 135 GB of RAM each. These backup servers were also used in XENON1T. Two of these three mainly serve as extra redundancy, while the third acts as a general purpose machine with access to the latest data. This machine, for example, automatically produces online monitor plots and handles requests for retrieving the same (as explained in subsection 5.4.4).

5.4 Live Processing

The data stream of raw data from the digitizers is fully processed onsite at LNGS. The triggerless data stream is handled by the data stream processor, strax [156, 157]. Using live processing and online data storage, data can be accessed while their collection is still ongoing.

5.4.1 Data stream versus discrete events

The triggerless design of the XENONnT DAQ results in a continuous data stream. For processing as well as storage purposes, handling discrete time intervals of data is advantageous as it allows for parallelization. To this end, the digitizer data which are read out by redax [164] are partitioned in 5 s–20 s time intervals called chunks. Each chunk is accompanied with an overlap region of ∼0.5 s to the previous and following chunk. As such, the overlap region is being saved twice, once with the previous chunk and once with the following. These overlap regions, called the pre- and post-chunk, are processed together with a chunk to ensure that each process has access to sufficient data to do the reconstruction. This is important for reconstructing S1/S2 signals (peaks) whose data might otherwise be split into consecutive chunks. Strax searches for time regions where there are no data for 1 µs within the pre- and post-chunk and discards the data before (pre-chunk) or after (post-chunk) this time region. This discarded time region will instead be processed together with the previous or next chunk. If no time interval of 1 µs is found within the overlap region, artificial deadtime would be inserted, which was never required for the entire SR0 dataset including calibration data. Due to this temporal separation between chunks, single chunks of low-level data are handled independently, allowing for

processing in parallel. For high level data, such as events, the processing is based on stateful algorithms which are therefore single threaded and may rearrange chunk boundaries.

5.4.2 Strax(en) data format

The processing at LNGS is handled by the eventbuilders and uses the publicly-available strax framework [156, 157]. Strax is a purely python-based streaming processor. Autovectorization, just-in-time compilation, and a tabular data format make the processing fast. The tabular data format is achieved by fixing the shape of the data fields in software. At the level of PMT traces (Fig. 5.2.5), this is achieved by splitting one variable-length PMT trace into a sufficient number of fixed-length intervals. The data are organized in a hierarchical structure of "datatypes". At higher level datatypes, like S1/S2-peaks, the summed waveform of all PMTs is down-sampled to a fixed number of samples.

There are several steps in the processing, from PMT-traces as in Fig. 5.2.5 at the lowest level, to a fully reconstructed S1 and S2 pair originating from one physical interaction within the TPC at high level. The PMT-traces are stored as *raw-records* as they are the lowest level (raw) datatype that is stored long term. The S1s and S2s are saved as *peaks* level data which can be grouped in time to form *events*. The time scales of the typical objects in raw-records, peaks and events differ by orders of magnitude. For example, Fig. 5.2.5 shows that raw-records can be of $\mathcal{O}(1\,\mu s)$ and an S2 peak of $\mathcal{O}(30\,\mu s)$ The duration of an event is set to be at least as long as the drift length (2.2 ms). Correspondingly, for higher level data, the number of items and the data size decreases by orders of magnitude. For instance, an hour of data may amount to 250 GB of raw-records, 5 GB of peaks, and only 30 MB of events.

The different levels of data processing are organized in software modules called *plugins*, each producing one or more datatypes which can serve as the input data for subsequent (higher level) plugins. This structure allows for a modular design and a flexible processing framework. When a chunk of data is processed, it is transferred between processing threads to any higher level plugins requiring it as input. Using this structure, the versioning of the data is handled per datatype by tracking the dependency chain. This has the benefit for reprocessing that, for example, a new or modified plugin at event level can only require event level input data to (re)compute and does not affect any lower level datatypes. During processing, auxiliary information on several quantities required for data processing, like PMT gains, are queried from a dedicated collection within the MongoDB database. This collection is frequently updated

with the latest values to ensure the data are processed with as up to date corrections and detector variables as possible.

5.4.3 Online processing

The processing by strax includes several stages, a full description of all its aspects is beyond the scope of this paper. Instead, some of the aspects are briefly discussed to illustrate that the full reconstruction of all the data is done live. The lowest level data, the raw-records, are all written to disk without further processing, allowing to always go back to the unprocessed data. After the raw-records level, PMT-traces are baseline-subtracted, inverted and integrated. For the TPC, time intervals are obtained wherein photon hits are extracted from the PMT traces to build peak sub-clusters, *peaklets*. The peaklets are classified and re-clustered according to their type to obtain peaks; S1-peaks are assumed to consist of only one peaklet, while S2-peaks can consist of many. Using this two-step clustering, strax is able to deal with the very short S1 signals while also being able to reconstruct the longer S2 signals as single peaks. Three different neural networks are applied on the peak level data for xy-position reconstruction based on the PMT hit pattern, which allows for cross-validation of their results. Events are built on the basis of a large S2 peak (the "triggering" peak). The triggering peak should be >100 PE and there should be fewer than 8 other peaks with at most 50% of the area of the triggering peak in a 10 ms window around the triggering peak. An event is the time region from 2.45 ms before and 0.25 ms after the triggering peak. This time region is set to be longer than the drift length of 2.2 ms (in SR0) and all peaks within the time region are considered part of the event. This is effectively the event trigger, which is set as a high-level configuration in the processing chain, in stark contrast to XENON1T [151], where an event window was fixed once at the DAQ and all other data was discarded. As a result, the event trigger is easily re-optimized in a high-level analysis.

Processing of muon veto and neutron veto data is also performed within strax using dedicated veto plugins which are applied similarly to both types of veto data. These plugins reconstruct veto-events based on the number of PMT hits. Additionally, a software coincidence trigger for the neutron veto reduces the data stream at a low data level. This software trigger is not used for the muon veto because it has a hardware coincidence trigger.

The program *bootstrax* is responsible for processing on the eventbuilders and is optimized per host machine to provide the maximum performance under a wide variety of data rates. As soon as a new run is issued by the dedicated

program (the dispatcher, see subsection 5.5.3), bootstrax looks for newly written chunks on the Ceph buffer disk. Bootstrax marks a set of data ready for uploading into long-term storage after completing the processing.

5.4.4 Online monitoring

The live processing on the eventbuilders is able to keep up with the data rates observed during SR0, including all the calibration periods. This opens up possibilities to use fully reconstructed data to monitor the state of the detector while data collection is ongoing. To this end, several datatypes are uploaded while data are being collected. This includes acquisition monitor data, all the fully reconstructed events and selections of data from the muon veto, neutron veto, and a selection of the peaks data from the TPC.

Redax buffers at least two chunks in memory, which first have to be written to the Ceph buffer disk before that data can be processed. Several chunks are usually combined in memory during processing before writing it to disk to reduce the number of small files in long-term storage. However, when chunks of processed data are uploaded to the MongoDB database, there is no such limitation and a chunk of processed data is therefore immediately uploaded after processing. This usually results in the data being available in the database $\mathcal{O}(30\,\mathrm{s})$ after light has been detected by PMTs.

Status overview plots to monitor the detector conditions and data quality are made with the open-source infrastructure [156, 157] and additional XENONnT software. As an example, Fig. 5.3 shows two unrelated changes in detector conditions close in time: a period of intermittent light emission ("flash") of PMT313 [165] (apparent from panel A) and the start of a calibration period with $^{83\mathrm{m}}$Kr (most clearly visible in panel C). This figure can be produced continuously to see changing detector conditions live. Additionally, each hour a plot is automatically produced and sent to the XENONnT-Slack [167] workspace which is used as the common chat room for the entire experiment. On Slack, one can also easily request periods of time to make this plot for, which is handled automatically by one of the backup eventbuilders. Alternatively, the data can be directly retrieved from the MongoDB database or via strax to do custom analysis, for example to create SuperNova Early Warning System (SNEWS) [168] warnings.

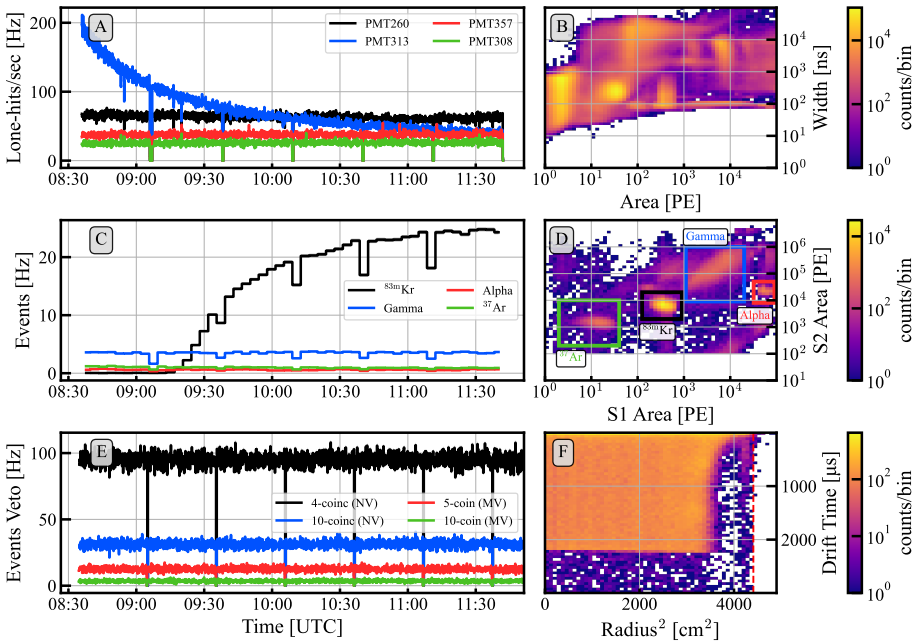

Figure 5.3: Online monitor plot for monitoring the detector status. Panel A shows the per PMT *lone-hit* rate, which are pulses that are seen in one PMT without any pulses in other PMTs within a short time interval. PMT313 (blue) just "flashed" [165], and is slowly returning to a rate comparable to the other PMTs. Panel B shows the area of peaks versus the range of 50 percent decile (known as the width) of the sum-waveform of a peak. This parameter space is useful for identifying peak populations, e.g., the peaks from ^{83m}Kr S1 are visible in the range 80-800 PE at roughly 100 ns width. Panel C shows the evolution of reconstructed events, which are roughly selected on their S1 and S2 area as shown in figure panel D. Since this was the start of a calibration period with ^{83m}Kr following an ^{37}Ar calibration, the event rate of ^{83m}Kr is increasing in panel C while the ^{37}Ar remnants are being removed via online distillation to a negligible level [166]. Panel E shows the evolution of the number of veto events in the veto-systems over time. Panel F shows the reconstructed event positions throughout the TPC, where at larger drift times (deeper into the detector), the events are reconstructed inward due to an inhomogeneous drift field inside the detector [104]. The drops in the rate to 0 Hz (panels A and E) mark the periods where the DAQ is switching from one run to another. As the data in panel C are re-binned, the run transitions manifest in $\mathcal{O}$ (10%) drops in the rate.

5.4.5 Data storage infrastructure

During the commissioning of XENONnT and the first science run (SR0), the DAQ collected >2 PB of uncompressed data. To reduce the required amount of long term storage, aggressive compression algorithms like *bz2* (in case of low data rates $\lesssim$65 MB/s) and *zstd* (in case of higher data rates) are used to compress the low-level data. The bz2 and zstd algorithms compress the raw data by factors of 5 and 4, respectively. While this increases CPU usage on the eventbuilders compared to the faster compression algorithms used for high level data, such as *blosc*, CPU usage is usually not the constraining factor for the eventbuilders.

The eventbuilders write to their own hard disks configured in a RAID5 configuration for performance and redundancy resulting in 22 TB of storage per eventbuilder. These hard disks, shown as local storage in Fig. 5.1, can be accessed by other hosts within the LNGS network. The data are uploaded from these disks into long-term storage as soon as bootstrax marks them ready for upload in the MongoDB database.

5.5 System Control & Oversight

Control and oversight of the DAQ and its associated subsystems are handled via databases, a user-interface website, and a software controller that coordinates the readout processes. Two additional servers are used in these roles, one in the LNGS surface server room and one underground with the other DAQ servers in the DAQ room. The surface server hosts the necessary databases for the DAQ, and also acts as a secure gateway through which experts can remotely access the DAQ subnet underground. The underground server hosts the DAQ website and the software controller.

5.5.1 Databases

Two databases are used for system control, monitoring, and interprocess communication, both implemented in the NoSQL-based MongoDB [154]. These are referred to as the "DAQ" and "Runs" databases. Each database is subdivided into "collections", analogous to tables in an SQL-based database, each containing "documents" which are analogous to rows. Unlike SQL-based databases, documents in one collection are not required to have the same schema, which allows for considerable flexibility.

The Runs database is a three-node replica set of servers located in LNGS, University of Chicago, and Rice University, the latter two being the primary XENONnT analysis facilities. This ensures that analysts have access to the database in the event of transient network disruption and protects against data loss due to hardware failure. One collection in this database contains metadata for each discrete run, which includes quantities such as the run start and end times, which of the three detectors were assigned to this run, the full readout configuration of each detector, and a listing of all datatypes for this run and their storage locations. Another collection contains data produced for online monitoring as described in subsection 5.4.4.

The DAQ database, in contrast, is expressly for the operation of the readout and not required for analysis. All data in this database are either only stored temporarily, or change very infrequently and can be restored from periodic backups in the event of data loss. This database, therefore, is neither replicated nor directly accessible outside of LNGS. Several collections in this database contain the regular status snapshots used to monitor the various components of the DAQ. These collections include the status snapshots of the redax readout processes, the health and performance of all the servers and NIM/VME crates in the DAQ system, and the status of the live processing. These collections are configured with time to live (TTL) indexes to only store data for 3 days, primarily to ensure that queries against these collections remain fast, and also because the information contained can either be reconstructed from the processed data or is additionally written to disk for long-term availability. Other collections store all available readout configurations, the system operational goal state set by the website, commands being issued to the readout processes, and important logging messages from the various DAQ processes.

5.5.2 User interface website

To facilitate easy use of the DAQ for the day-to-day operation of the experiment, a front-end website was developed using NodeJS [169]. A variety of pages allows users to view the current readout performance, set the system's operation goal state, and monitor the data rates from each readout channel to identify potential problems in the detector, such as localized regions of sustained electron emission, also known as "hotspots". The *status* page displays instantaneous data rates for each readout process in the entire DAQ system, information about the current activity of each eventbuilder, and the current status of the Ceph high-speed buffer disk. Additionally, a plot displays the recent data rate for each readout process, which allows a user to identify transient behavior in the system that may

not be clear from the instantaneous rates alone. This page is shown in Fig. 5.4. The *control* page allows users to specify the operational goal state of all three subsystems, such as selecting an operational mode and a desired run duration. The *monitor* page displays the instantaneous data rate for each channel in the TPC in a convenient layout mirroring that of the physical locations of PMTs in the TPC (shown as the inset in Fig. 5.4). In this view, a hotspot will appear as a localized region, typically three adjacent PMTs, with a data rate significantly higher than other nearby PMTs. Additionally, a plotting function is provided to allow for a direct comparison of recent rates between different channels. Another page provides a user interface to the Runs database, where metadata about each run can be viewed. Other pages allow experts to modify and create preset operational modes and configurations, monitor the status of all the servers in the DAQ network, and interface with the dispatcher (described below).

Finally, an application programming interface (API) is provided to enable programmatic control of the DAQ and access to part of the DAQ database. This is used, for instance, by Slow Control to perform the periodic automatic calibration of the PMTs via a pulsed LED. Slow Control continuously queries the API, and notifies experts if any aspects of the performance of the system deviate from what is expected, if disks are full, or a hotspot is suspected based on the per channel data rate.

5.5.3 Readout coordination software

To oversee and coordinate the readout processes, a program called the *dispatcher* was developed. The primary responsibility of the dispatcher is to convert the desired operational goal state as specified on the website into direct commands issued to the various readout processes. To do this, the dispatcher retrieves the most recent status snapshot of each readout process. These are aggregated together to determine an overall status for each subsystem, such as if the subsystem is idle, running, in a transitional state, or if processes are not responding. This aggregated status is then compared to the desired operational goal state of each subsystem, and commands are issued to the readout processes to make the former match the latter. For example, if a user wants the readout to begin with a certain operational mode, the dispatcher will ensure that the necessary processes are capable of starting, issue commands to begin the digitizer-programming sequence, wait until all necessary processes report the successful completion of this sequence, and then issue the start command. When an active run reaches the desired length as specified by the website, the readout is stopped, and the cycle is repeated. In rare cases where a readout process or digitizer stops responding

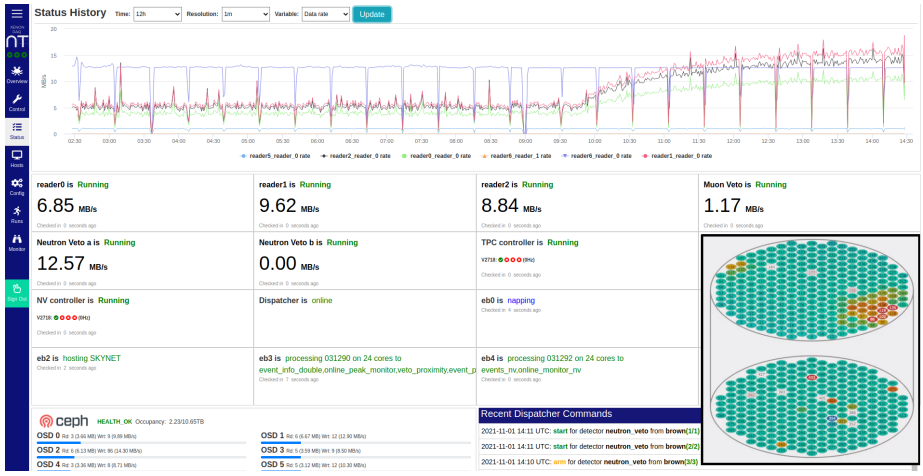

Figure 5.4: The Status page on the DAQ interface website. The plot shows the data rates for the past 12 hours, and the status cards show the instantaneous statuses of all readout and live processing elements in the DAQ system. An increase in the rate due to the start of regular detector calibration with ^{83m}Kr is clearly visible. A navigation bar on the left provides convenient links to other pages in the website. The inset in the bottom right shows the per-channel rate as displayed on the monitor page.

properly, the dispatcher will automatically kill and restart delinquent readout processes and power-cycle VME crates as necessary to restore normal behavior. Additionally, if such action is necessary during linked-mode operation and restarting a process or VME crate fails to rectify the situation, the dispatcher will unlink the detectors so the readout of the detectors that are responding normally can continue. Experts receive notifications whenever automatic actions such as these are taken. At the start of every run, the dispatcher creates an entry in the Runs database containing a copy of the readout configuration and other metadata necessary for the live processing. At the end of a run, the corresponding entry is updated with additional quantities such as the end time of the run and the average data rates for all contributing detectors.

5.6 Performance

In the first two years of operation, the TPC subsystem collected more than 1280 TB of data, while the muon veto and neutron veto collected 28 TB and

680 TB, respectively. The performance of the overall DAQ system can be measured in several ways. Most obvious are the maximum data rate the system can maintain and the livetime with which the system operates, but other criteria such as inter-detector synchronization and noise levels are also important. An additional key performance metric is the speed of the live processing, as it is crucial that the data are made available for transfer off-site at least as fast as it is recorded.

5.6.1 Noise levels

Externally triggered, short runs were taken weekly throughout SR0 to assess the noise levels in all TPC channels, using fixed windows of approximately 1 ms duration. The mean RMS noise level was found to be stable for each PMT array with values of 0.23 mV and 0.34 mV for the high-gain channels in the top and bottom arrays, respectively, and 0.16 mV for the low-gain channels in the top array. The installed filter boxes are effective in suppressing electronic noise >250 kHz, which is related to HV power supplies. However, as expected, the filter boxes have a negligible effect on the low-frequency noise peak at 24.41 kHz, which is correlated with intrinsic noise produced by CAEN digitizers. Channels from the bottom PMT array on average exhibited an RMS noise level which was ∼1 ADCc higher when compared to top PMT array channels. This effect could be attributed to either a different resistor type and assembly procedure that was employed for filter boxes used for the bottom PMT array, or different and noisier HV power supplies. These low levels of noise support low digitization thresholds for the PMTs. Over 98% of TPC channels have thresholds set at 2 mV and only 1 PMT has a threshold higher than 3.4 mV, giving an average acceptance to single PE > 90%. For the neutron veto, 109 PMTs (91%) have thresholds set at 1.8 mV (∼ 0.3 PE) and the other 11 at 2.4 mV (∼ 0.4 PE). The average noise for the neutron veto PMTs is 0.3 mV (∼0.05 PE). Lastly, for the muon veto the average RMS noise is 0.18 mV, the thresholds are set for all channels at 3 mV (∼ 1 PE). All above voltages were estimated for an input impedance of 50 Ω.

5.6.2 Livetime

Throughout SR0 and the commissioning of XENONnT all three DAQ subsystems operated stably, collecting in total more than 200 days of commissioning and science data, and close to 100 days of various calibration data. The deadtime fraction induced by the operation of the busy veto is 2×10^{-5} for the

majority of SR0 science data (which is typically $\lesssim$25 MB/s), as illustrated in Fig. 5.5. The average deadtime fraction for all SR0 science data is 3×10^{-4}. Furthermore, during high-rate ^{220}Rn and AmBe calibration periods the deadtime resulting from the combined operation of the busy and HEV on average amounts to $\sim$ 10%. The above deadtime values describe only the intrinsic deadtime produced by the operation of the busy and HEV modules, and not the data reduction caused by the analysis cut described in subsection 5.3.4. Lastly, it should be noted that the busy veto-induced deadtime of the NV was found to be negligible.

The highest sustained data rate in SR0 was $\sim$500 MB/s during an AmBe calibration (population near 500 MB/s in the bottom panel of Fig. 5.5). The DAQ was designed to withstand higher data rates, but performing high rate calibrations in SR0 was inhibited by the long event duration which leads to pileup of events where they start overlapping. A few chunks were digitized with data rates of up to 600 MB/s.

5.6.3 Time synchronization

Throughout SR0 data-taking the three DAQ subsystems operated in linked mode with full temporal synchronization. However, as was described above, to facilitate the operation of the HEV the TPC data are delayed within the V1724 digitizers by 10 µs, as compared to the TPC acquisition monitor and the other sub-detectors. The time synchronization across all sub-detectors was verified using the 0.1 Hz signal generated by the GPS. It was supplied to a dedicated TPC digitizer channel, as well as to its acquisition monitor. Additionally, this signal was acquired by the acquisition monitor of the neutron veto and a digitizer in the muon veto. A comparison between the timestamps of these signals was used to measure the average temporal difference between the sub-detectors. The average time difference between the TPC and neutron veto was measured to be 10 157 ns, while the time difference between the TPC and muon veto was found to be 5283 ns. The variation in the measured delay time is related to the trigger formation time used by the muon veto. After SR0 these constant offsets are subtracted out during readout, resulting in time synchronization with a precision of $\sim$10 ns, which is comparable to the sampling time of the digitizers.

To illustrate the inter-detector synchronization a muon event passing through all XENONnT sub-detectors is presented in Fig. 5.6. The signals obtained in each detector are aligned by accounting in software for the time differences described above. In addition to the sub-detector signals, also the analog sum waveform that is acquired by the acquisition monitor of the TPC

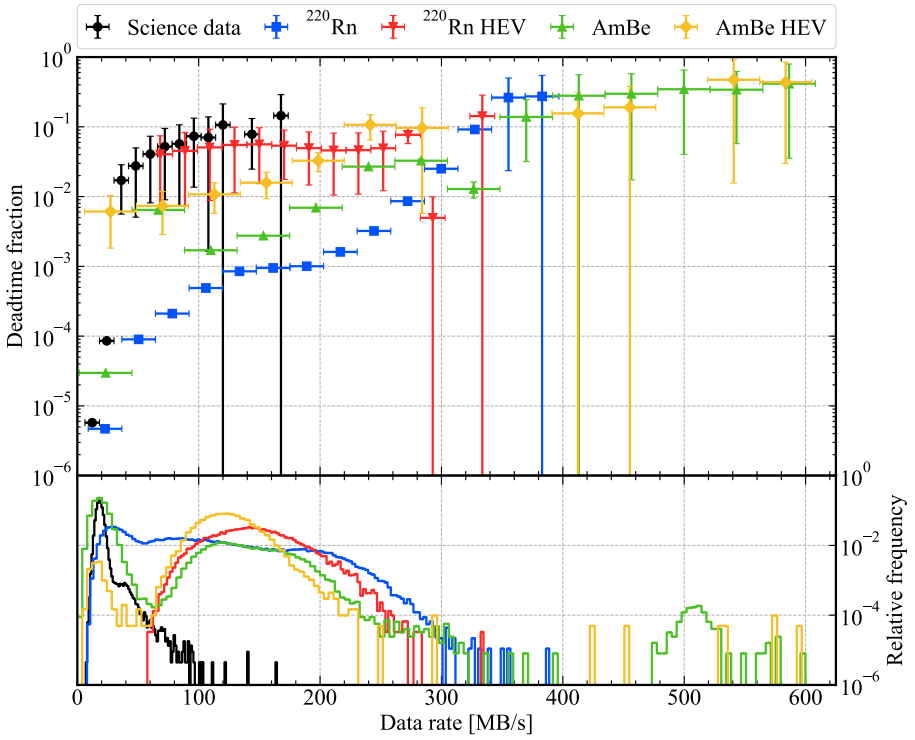

Figure 5.5: Mean total deadtime per chunk (top panel), and relative frequency of chunks (bottom panel) as a function of data rate per chunk for several run modes. Each chunk is a time interval of 5 s to 20 s. AmBe calibrations are performed by keeping the source at several positions with respect to the TPC, leading to distinct populations in the bottom panel. The High Energy Veto (HEV) run modes (AmBe HEV and ^{220}Rn HEV) have higher deadtime fractions as a result of the inserted deadtime by the HEV, see subsection 5.3.3. For typical (98%) science data ($\lesssim$25 MB/s) the deadtime fraction is 2×10^{-5}. In science data, higher rate data points are caused by short periods in time following a muon traversing the TPC, leading to high data rates and deadtime fractions of $\mathcal{O}(1\%)$. ^{220}Rn has a lower deadtime fraction for $\gtrsim$25 MB/s than science data, since these higher rates are caused by a higher S2 rate, rather than muons. Above $\sim$250 MB/s the data quality deteriorates due to the onset of pileup.

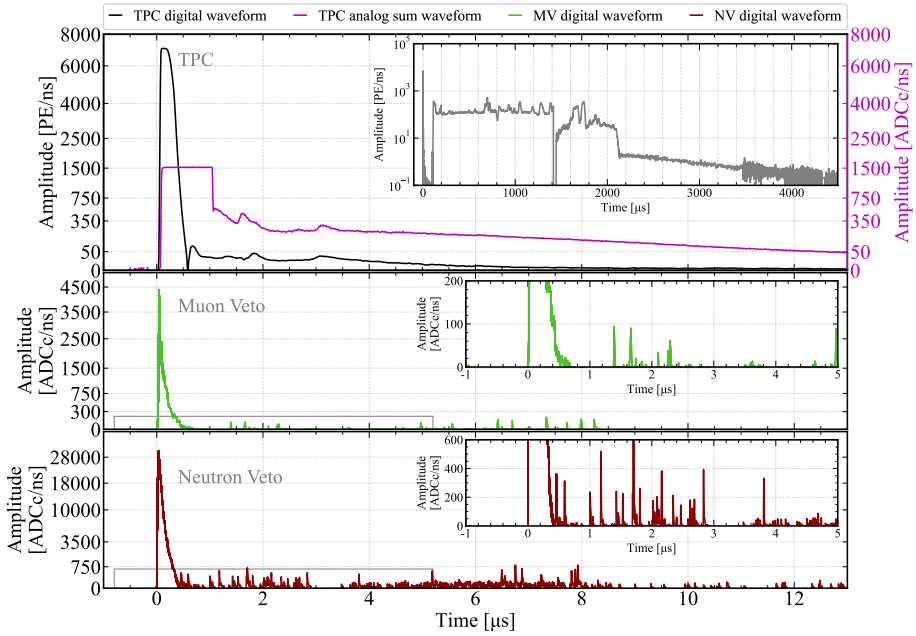

Figure 5.6: An illustration of a muon event recorded by all three DAQ subsystems. The top panel shows a zoomed-in view of the beginning of the muon waveform that was recorded by the TPC digitizers (black). An analog sum waveform of the bottom PMT array that was recorded by the TPC's acquisition monitor digitizer is also shown in the same panel (magenta). The inset in the top panel shows the entire muon waveform duration as seen by the TPC digitizers. The drop at ∼1400 µs is caused by a baseline fluctuation. The middle panel shows the same muon event recorded by the muon veto, while the bottom panel shows the data recorded by the neutron veto. Insets in both middle and bottom panels show a zoomed-in view of the respective waveforms. All three waveforms were aligned based on the TPC signal (black).

is shown in the top panel. As seen in the inset of the top panel of Fig. 5.6, the prompt S1 from the relativistic muon's interactions is followed by a sustained S2, which lasts for the full drift time of the detector. This extended S2 indicates a vertically traversing muon interacting along most of the drift column to produce ionization electrons. The subsequent long tail is formed from photoion-

ization electrons that follow the S2 for multiple milliseconds. The muon track in the muon veto and neutron veto sub-detectors corresponds to the Cherenkov photons detected by the PMTs of these detectors, followed by PMT afterpulses lasting up to ∼10 µs. The general shape of the waveform in all three panels is the same, indicating clearly that the same event is shown. Lastly, the GPS synchronization signal was also used to evaluate the clock drift of the DT4700 clock module, yielding approximately 2 µs over a period of 10 s, or 0.2 ppm.

5.6.4 Live processing performance

The DAQ and its software were designed to run under high data rates during detector calibrations using a combination of internal and external radioactive sources to quantify the detector performance. However, due to a limited voltage on the TPC cathode during SR0 compared to the design value, high rate calibrations could not be performed as events quickly piled up because of the long drift time of 2.2 ms. As a consequence, the system did not have to work under persistent high data rates and live processing was always able to process the data faster than it could be collected. During SR0, the data rate never exceeded ∼50 MB/s for extended periods of time. To quantify the performance of the eventbuilders in high data rate conditions, pre-SR0 commissioning data that were taken during a high rate ^{83m}Kr calibration are used.

Fig. 5.7 shows bootstrax cumulative processing time for different datatypes. Here, raw-records is the lowest level datatype, followed by peaks and finally events. The results for higher level datatypes include the time to compute the lower level datatypes as well. There are some additional datatypes of intermediate data in between, which one can find in the straxen documentation [170], several of which were briefly discussed in section 5.4. The total processing time comprises the time of starting bootstrax for a given run, decompressing the redax data, processing the data until the specified datatypes, compressing and writing all of the processed data to disk.

Fig. 5.7 shows that for data rates below 250 MB/s, the processing time is shorter than the collection time and a single eventbuilder can manage the entire data stream regardless of the datatype considered. For higher rates, the processing up to the events or peaks datatype is not fast enough to keep the processing live as each chunk would be processed slightly later than it is acquired. At these data rates, the finite RAM of the servers and increased disk read/write operations prevent processing at the same rate as at lower data rates, since processing each new chunk on a separate core starts requiring more memory than available on the host. The break-even line for one eventbuilder in Fig. 5.7

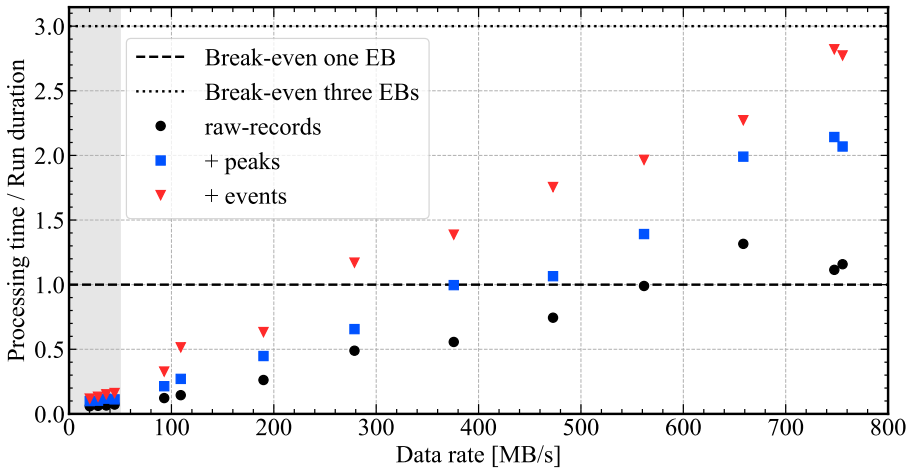

Figure 5.7: Live processing time as a function of the raw data rate for several target datatypes. The raw-records datatype is the lowest level datatype, followed by peaks and finally events (for simplicity this is called events, even though the benchmarks were obtained for the event-basics datatype [170]). When a higher level datatype is computed, all the lower level datatypes are also produced, so processing event also includes raw-records and peaks. All SR0 science data in this figure are 50 MB/s (gray band). Pre-SR0 commissioning data were used for the high rate data points where the DAQ was operated with a fractional livetime mode (discussed in subsection 5.3.3) during a high rate ^{83m}Kr-calibration. For data rates $\lesssim$250 MB/s the live processing keeps up with one eventbuilder (EB) server as the processing time is lower than the acquisition time for any datatype. For any data rate in this plot the points are below the break-even line of the three eventbuilders, meaning that live processing at the DAQ could keep up with the readout.

lies around ∼250 MB/s for events, ∼400 MB/s for peaks, and ∼550 MB/s for raw-records. Additionally, the work is divided among three eventbuilders (with two additional as backup) and the combined eventbuilders can keep up with much higher data rates.

The total rate of the neutron veto and muon veto subsystems was found to be relatively constant and amounts to 10 MB/s–20 MB/s and ∼1 MB/s, respec-

tively. It takes about 80 s to process a 1800 s run for the neutron veto data, and $\lesssim$40 s for a 1800 s run of muon veto data.

5.7 Conclusion

The XENON collaboration has designed and commissioned the triggerless XENONnT DAQ. By forgoing a trigger and relying instead on fast software to handle the continuous data stream, all data exceeding the digitization threshold is written to disk. The TPC, muon veto and neutron veto subsystems that constitute the DAQ can be operated independently, or as one linked system sharing the same 50 MHz clock signal. The increased number of PMTs and the double digitization of the top PMT array leads to roughly three times the number of channels with respect to XENON1T for the TPC. While the triggered muon veto-subsystem remains virtually unchanged, the new neutron veto-subsystem was successfully built to enable tagging of neutron events, one of the main backgrounds for the XENONnT WIMP search. A 500 MHz sampling rate enables the required characterization of neutron signals.

The DAQ is able to operate at the highest data rates observed during the first science run of XENONnT (SR0) of ∼500 MB/s with the potential to go higher. The deadtime fraction is as low as 3×10^{-4} for science data and $\lesssim 10\%$ for calibration data at high rates of <350 MB/s.

Using online processing, high level data are directly made available to monitor the detector. This enables analysts to have immediate feedback on changing detector conditions with fully processed data. The online processing is able to handle all data rates observed during SR0, where each of the three dedicated servers is able to process the data with a rate of up to ∼250 MB/s. The maximum observed data rates during SR0 were limited by the low drift-field conditions. However, the DAQ was designed and is capable of dealing with data throughput rates greater than 750 MB/s.

During commissioning and SR0, the XENONnT DAQ has collected more than 2 PB of both science and calibration data, and it will continue to operate in subsequent science runs. The successful operation of the XENONnT DAQ and the implementation of the triggerless readout paradigm provides a solid basis for the development of DAQ systems for the next generation of liquid xenon dark matter experiments [42, 77].

Chapter 6

Signal reconstruction

The UV-photons of the primary scintillation (S1) and the secondary scintillation (S2) signals are detected by the two PMTs arrays (as discussed in Chapter 3). The events are reconstructed from these signals by identifying and matching S1 and S2 signals from the same interaction. This chapter describes the reconstruction of the S1 and S2 signals starting from digitized PMT waveforms. Using simulations, we validate and quantify the reconstruction performance.

6.1 Detector response

The response to single UV-photons is the same for both S1 and S2 signals, and a detailed characterization of the Single Photoelectron (SPE) response of the PMTs is at the basis of understanding the detector. The incident photon liberates a photoelectron from the photocathode that gets amplified by multiple dynodes inside the PMT (see [171] for a full description on PMTs). Each dynode stage amplifies the signal, so that the last dynode (the anode) receives as much as 2×10^6 (the set gain for XENONnT PMTs) electrons for an SPE. This creates the pulse that is digitized by the Data Acquisition system (DAQ). The gain of the PMTs is determined primarily by the voltage applied to the PMT. A higher voltage increases the secondary emission of electrons at each dynode. Additionally, the response of a PMT to UV-photons depends on several effects:

- The probability of having an under-amplified signal, caused by effects like a photon striking the first dynode instead of the photo-electrode.

- The amount of noise from the PMT, high voltage supply, signal cables, and digitization hardware involved in the readout of that PMT.

- Double Photoelectron Emission (DPE), where a single incident photon leads to two photoelectrons. The probability of this effect depends on the energy of the photon [172] and is 22% for the 178 nm xenon scintillation light and XENONnT PMTs. The LED data does not show this high DPE probability as the wavelength of the LED light is much longer (470 nm).

- The energy resolution of the PMT. The energy resolution is primarily set by the gain, as more quanta are incident on the anode, the resolution improves, and the PMTs are therefore set to a high gain of 2×10^6 [173].

To measure the SPE-spectrum, periodic LED calibrations were performed during Science Run 0 (SR0) [174]. These externally triggered runs digitize the signals of the PMTs at the time of the emission of a low intensity LED pulse. A low intensity LED pulse is used to prevent multiple photons arriving at the PMTs at the same time, with an estimated probability of 0.3% [174]. The LED light is used to measure the gain and characterize the response of the PMT to an SPE. The LED data for PMT 129 is shown in Fig. 6.1. A model from Ref. [175] is fitted to the data that describes the components from noise, 1 PE, and several PE. Apart from the Gaussian components that describe noise, 1 PE, and several PE, contributions from the PMT response like dark counts and under-amplified signals are modeled by modified exponential Gaussians. The modified Gaussian terms provide an empirical model that allows obtaining a good fit ($\chi^2/n.d.f. \sim 1$) to the data. The SPE spectrum[1] is obtained from the LED data as the 1 PE contributions plus the 0 PE modified Gaussian exponential [174].

The digitization thresholds discussed in Chapter 5 are determined using the SPE-spectra. In an idealized situation, a photon would always yield a 1 PE signal (i.e., probability 1 at 1 PE, 0 everywhere else), however, due to PMT effects (e.g., a finite energy resolution, the photon skipping the photoelectrode and hitting the first dynode or noise), a portion of the signal is under-amplified, so there is a probability of a photon only creating a <1 PE signal. This is particularly important, since this can lead to signals that are below the digitization-trigger (hitfinder) thresholds in the DAQ (straxen). The digitization threshold reduces

[1] For the modeling in WFsim [176], the LED data is first subtracted with the noise spectrum that is obtained in dedicated runs, such that the subtracted LED spectrum does not include the data from the 0 PE peak as in Fig. 6.1. Following the same procedure as described here, the WFsim SPE-model is obtained.

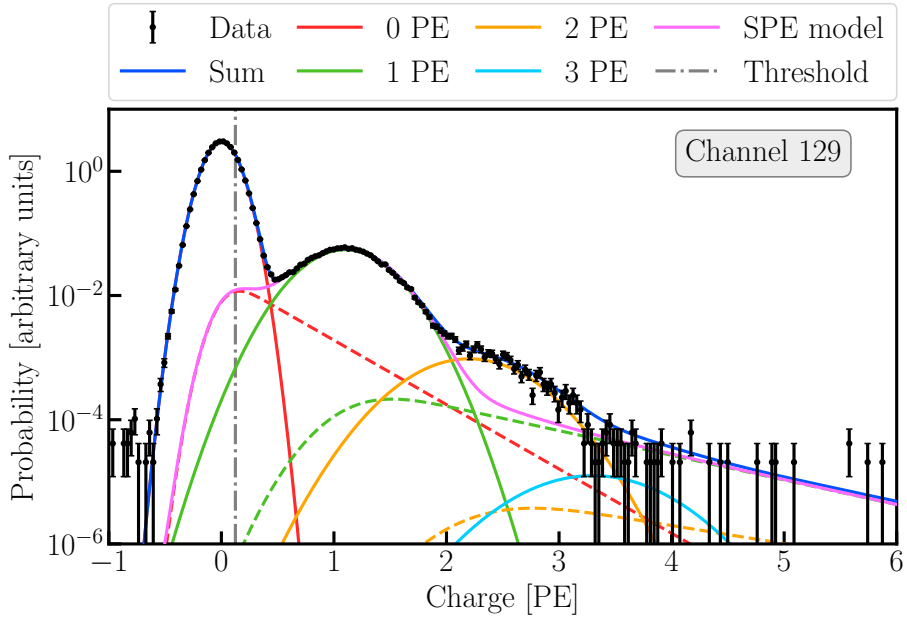

Figure 6.1: LED data for PMT 129 to determine the Single Photoelectron (SPE) spectrum [174]. The SPE model consists of several contributions, as explained in the text. The exponential modified Gaussians (dashed lines) are shown for the 0, 1 and 2 PE contribution. The lowest part of the LED data is due to noise and the pedestal (the response of the PMT to no signal). Bipolar noise may be both above and below baseline, hence if only noise is observed after an LED signal, the area may be negative. Using the SPE shape, the digitization thresholds discussed in Chapter 5 were optimized. The 2 and 3 PE contributions are also added for illustration purposes — although the statistics for fitting in their respective ranges is small and their contribution to the total spectrum minute.

the data rate, the higher the threshold, the less noise is digitized, at the expanse of not digitizing a portion of the SPE signals. The SPE acceptance is the fraction of the SPE response that is above the digitization threshold. Based on SPE-spectra, the average single photoelectron acceptance is 91%.

The spread on the arrival times of S1 photons at the PMT arrays are dominated by the distance of the interaction to the PMT arrays, and by diffusion for S2 signals, as will be explained in the following sections. In addition to S1s and S2s, there are several other signals that affect the reconstruction. Afterpulses are due to the ionization of residual gas particles inside PMTs that are attracted by the photocathode and cause emission of one or more additional electrons [165]. Delayed electron extraction - which are electrons trapped at the gas liquid interface - are an additional source of S2-like signals [173]. Finally, PMTs exhibit thermal dark counts [177].

6.1.1 Signal modeling

The signal formation and detection are simulated by the WFsim [176] package, which combines results from the generic NEST [95, 96] package and XENONnT specific detector modeling. It handles simulations of optical properties like the light collection efficiency, photon arrival times, electron propagation through the LXe, electron extraction and other detector effects. WFsim conveniently casts simulated data into a format native to strax(en) [156, 157] such that it can be analyzed the same way as the real data.

The overall hierarchy of signals that are used for the event reconstruction is shown in Fig. 6.2. At the lowest level is the SPE response of a PMT. The signals of S1s and S2s are the sum of these responses, and how many photoelectrons are detected for each is determined by the energy deposited in the LXe and the type of energy deposit. WFsim is able to model the most important types of interactions in the TPC; S1s, S2s, their respective after pulses, and noise are all part of the simulation stack. Some secondary effects, like PMT dark counts and, the perpendicular wires (subsection 6.3.1) are not included.

6.1.2 Reconstruction

The reconstruction chain works similar to the signal hierarchy in Fig. 6.2, and starts at the per-PMT waveforms. The algorithm searches the PMT waveforms for signals induced by photoelectrons, and the time intervals of these signals are called “hits”. These per-PMT hits are sequentially grouped with neighboring hits (from any PMT) within 700 ns in time to form clusters. Isolated hits (“lone hits”), which have no neighboring hits in this time window, are mostly due to afterpulses or dark counts, and are handled and stored separately. The clustered groups of hits are iteratively split to smaller clusters based on their timing information and the summed waveform of all hits in the cluster. These smaller

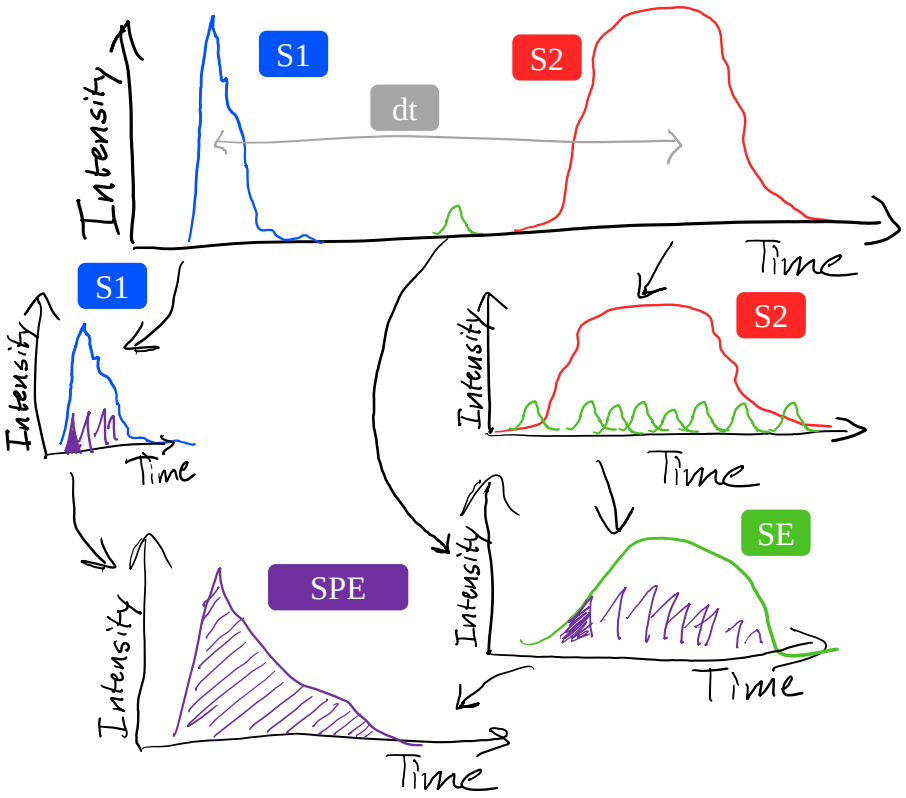

Figure 6.2: Signal hierarchy in XENONnT. At the highest level are "events", the matched S1 and S2 signals. These are separated by drift time *dt*. The S1 is a peak consisting of discrete photons that, when incident on a PMT, each produce an SPE response. The S2 in contrast, is the scintillation light caused by the sum of (single) electrons (SEs, see text). The SEs in turn - like S1s - are detected as the sum of SPEs.

clusters, "peaklets", are sequentially classified as S1 or S2 peaklets, based on their waveform shapes, as will be discussed in Fig. 6.3. S2s peaklets are merged to form S2 "peaks" based on their area and timing information to adjacent S2 peaklets. If any lone hits fall within the duration of the S2, they are included in the S2 peak.

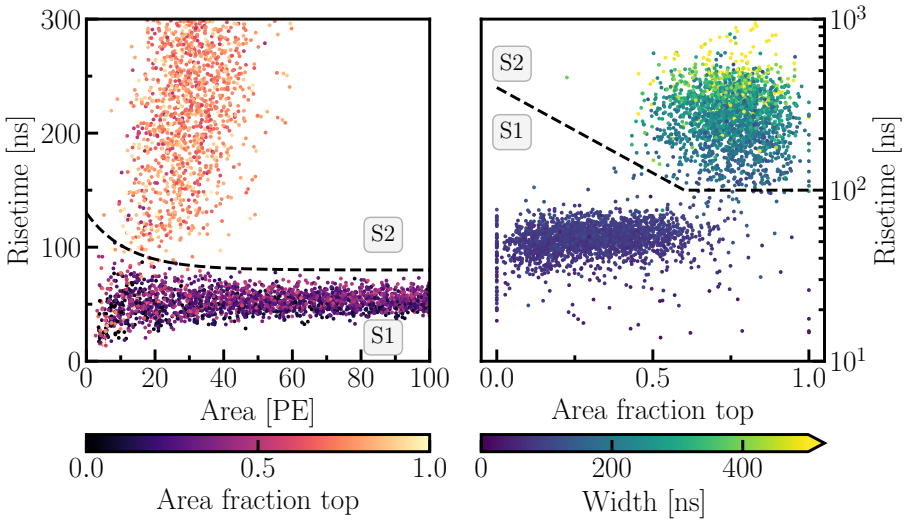

Figure 6.3: Classification boundaries in straxen [157] between S1s and S2s using the peak characteristics as explained in the text, showing the S1/S2 classification boundaries in area versus risetime (left) and Area Fraction Top versus risetime (right). The data shown are S1 peaks from a ^{220}Rn calibration, and isolated single electrons (the smallest S2s, see next section). Peaks are only classified as S1s if they are below the dashed lines. S1 signals are typically much faster (i.e., have a shorter risetime) than secondary scintillation, this property is used in the left panel to discriminate between the two signals. S1s signals are expected to be measured predominantly by the bottom PMT array, while secondary scintillation light is produced in the GXe and therefore detected most by the top array.

The classification boundaries between S1/S2s are shown in Fig. 6.3. These boundaries are encoded in the reconstruction software straxen [157]. The boundaries utilize several characteristics of the peak(let)s:

- The *area* of the peak, which is the total charge (gain corrected, in PE) measured by all PMTs during the peak.

- The *risetime* is defined as the time between the 10% and the 50% area quantiles of the sum waveform[2]. The 10% and 50% area quantiles are obtained from the time intervals since the start of the peak, wherein 10% and 50% of the total charge of the peak is achieved respectively.
- The *width* of the peak is the time range of the central 50% of the area of the peak resides.
- The *tight coincidence* (TC) is the number of different PMTs that have a hit within 100 ns around the area-weighted center time of the peak.
- The *area fraction top* (AFT) is the fraction of the total area seen in the top PMT array.

In the following sections, we will use these properties to discuss and quantify the performance of the reconstruction using simulations.

6.2 Scintillation signals

In this section, we will briefly describe the shape of the S1 signals. The performance of the reconstruction is then validated with the simulation framework. Finally, we show that this validation is in agreement with an independent data-driven method.

6.2.1 S1 Shape

The S1 photons arrive at different times at a photocathode depending on their generation timing (such as the singlet/triplet ratio) and propagation. The propagation is mainly dominated by the distance of an interaction to the PMT-arrays, which includes reflections off the PTFE-walls of the TPC. The TPC is 1.3 m in diameter and 1.5 m long, which takes the UV-photons 7 ns and 9 ns to traverse. An additional spread on the timing information of the S1 signal is caused by the different paths an electron can have in the PMT depending on where the photocathode was hit by a photon. This induces a 9.2 ± 0.5 ns [94] time spread on the arrival times of the signal (the transit time spread) for the Hamamatsu R11410-21 PMTs. Finally, the signals from the PMT are digitized with a sampling rate of 100 MHz.

[2] In the literature a slightly different definition is often used that refers to the time it takes to rise from 10% to 90% of the signal.

Short description	Long description
`found`	The peak was matched one-to-one between truth and the outcome (type agrees, no other peaks in time range). The area, widths, etc. can still be quite different.
`misid as XX`	The peak is present in the outcome, but it is misidentified as `XX`.
`merged`	The peak is merged with another peak in the outcome, the new 'super-peak' has the same type.
`merged to XX`	As `merged`, but the 'super-peak' has type `XX`.
`split`	The peak is split in the other list, but more than one fragment has the same type as the parent.
`chopped`	As `split`, but one or several fragments are unclassified, exactly one has the correct type.
`split and unclassified`	As `split`, but all fragments are unclassified in the outcome.
`split and misid`	As `split`, but at least one fragment has a different peak type.

Table 6.1: Possible matching outcomes for simulated data. Each peak in the simulated data is matched to its instruction, and the reconstructed type (S1 or S2) is compared to the instruction. Additionally, a peak is considered accepted if it is found, chopped, split, or split and misid and its area is sufficiently large, as described in the text.

6.2.2 S1 Efficiency

The performance of the reconstruction is quantified by evaluating the efficiency (the fraction of events correctly reconstructed), and bias (how well does the reconstructed area agree with the simulation input). The efficiency determines the energy threshold, and the bias affects the energy scales that are reconstructed from data.

We simulate events homogeneously distributed in the TPC with a flat energy spectrum. The most important parameters for this study are the shape

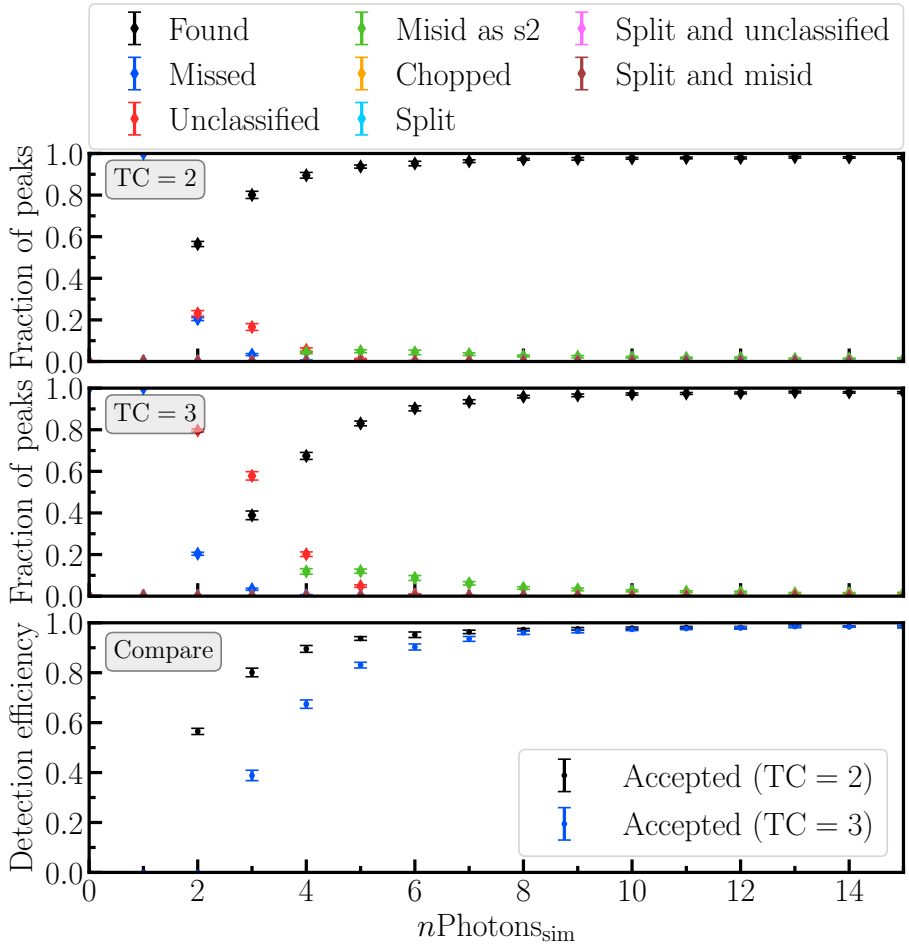

Figure 6.4: Simulated S1 efficiencies, where the top panel shows the matching outcomes for a twofold Tight Coincidence (TC) requirement, the middle panel shows the same for a threefold TC requirement, and the bottom panel shows the accepted fraction of the top two panels.

characteristics of the signals (as discussed in subsection 6.1.2), since the classification boundaries depend on these characteristics. The simulation outputs

of S1s (and S2s) have been shown to match in these parameters, as shown in Ref.'s [178, 179], and we will cross-validate the following simulation-based results with a data driven method (subsection 6.2.3).

For each peak in the simulated dataset, we match the reconstructed peak with the original instruction ("truth") tracked by WFsim. We categorize the possible outcomes in Tab. 6.1. Fig. 6.4 shows the acceptance of S1s as a function of number of simulated UV-photons incident on the PMT ($n\text{Photons}_{\text{sim}}$, also see Fig. 6.7). We consider the S1 peak "accepted" if[3]:

- A single peak in the truth matches to a single peak in the reconstructed peaks.
- An S1 peak is also considered "accepted" if it is "split", "split and misid", or "chopped" due to an afterpulse shortly after the peak. If the reconstruction artificially reconstructs an S2, it will never be considered accepted. For a peak to be accepted despite being split or chopped, its reconstructed area has to be at least 80% of the area of the original peak. This additional inclusion of these matching outcomes is particularly useful for e.g., the energy threshold determination (section 6.4), since S1s in the data analysis are also taken into account irrespective if an afterpulse is merged with the S1. If dedicated cuts would remove S1s that have an afterpulse merged to the S1, this would be accounted for in the respective cut efficiency. The inclusion of these outcomes is relevant for $\sim 2\%$ of S1s with $n\text{Photons}_{\text{sim}} \gtrsim 30$ and negligible for smaller signals.

If the signal is missed, the peak is detected by less than two PMTs, such that it is considered a lone hit. When the peak remains unclassified, it means that the properties of the S1 are not consistent with the S1 classification boundaries in straxen, relating most often to the TC requirement, the risetime, or to a lesser extent the AFT (Fig. 6.3), nor with the S2 requirement of at least 4 PMTs contributing to the peak.

The main loss mechanisms of S1 efficiency are the requirements we impose on the classification of an S1 (Fig. 6.3). The first requirement is that there must be at least two or three hits in distinct PMTs within a TC window of 100 ns around the center time of the peak. This is a powerful tool to reduce accidental coincidence events that would arise from random pairing of dark counts or noise. For the first DM search [6] and the low energy ER search [97],

[3]Note that "found" and accepted are used for different purposes. As a peak may be split, but still consider accepted, for example if a following afterpulse was not merged with a peak it will be "split" or "chopped" but still considered accepted.

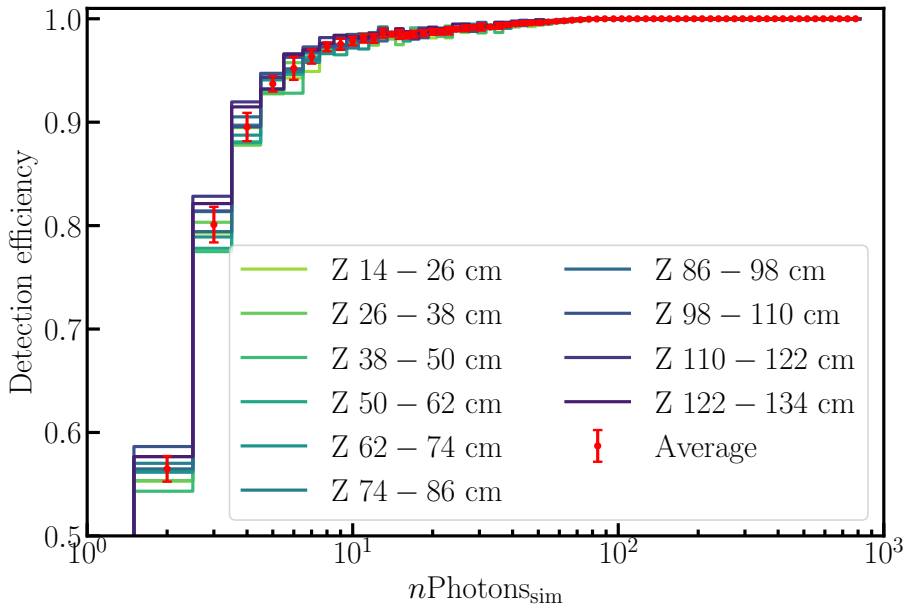

Figure 6.5: S1 acceptance as a function $n\mathrm{Photons}_{\mathrm{sim}}$ for different interaction depths. The efficiency is depth-dependent (especially for few photons). Therefore, we calculate the efficiency (e.g. Fig. 6.4) by taking ten z-bins (in the FV). The mean efficiency is calculated by averaging over these ten z-bins. The uncertainty is taken as the standard deviation of these z-bins (for a given number of photons). The obtained standard deviation is at least the statistical error and the systematic error from the z-dependence.

a TC requirement of three PMT hits (threefold coincidence) was required. In XENON1T, dedicated low energy analyses like for CEνNS [132] lowered the threshold to a twofold coincidence. The effect of the TC requirement is mostly at low energies, where there is a relatively high probability that there are too few photons within this time window. The second requirement that is preventing us from classifying S1s correctly are the risetime and AFT boundaries on S1 classification aimed at discriminating between S1s and Single Electron (SE) shown in Fig. 6.3. As will be discussed in the next section, S2s signals are the

scintillation light of multiple electrons[4]. The misclassification observed at all energies and mostly results in high risetime S1s to be misclassified as S2s.

The shape of the S1 signal is depth-dependent, as there is a larger arrival time spread of the UV-photons that are further removed from the PMT arrays. As the shape of the S1 signal changes with depth, so does the detection efficiency, as shown for the simulations in Fig. 6.5. S1s from the upper part of the LXe volume on average are wider in time due to the higher number of reflections inside the TPC. In the modeling of the detector response, a z-dependent S1 detection efficiency could account for this. However, as this is computationally expensive, we instead model the z-dependence as a systematic uncertainty. We use ten bins within the Fiducial Volume (FV) and take the median and the 1σ quantiles of each of those bins as the efficiency and uncertainty on the efficiency.

6.2.3 Data driven S1 efficiency

The simulation driven method discussed above was cross-validated using a data driven method [180], which results in Fig. 6.6. The data driven method of finding the fraction of accepted S1 signals relies on the principle that the S1s are the sum of several SPE signals, which can be re-sampled to form a "fake" S1 of fewer number of photoelectrons. By feeding these re-sampled S1s into the reconstruction chain, the reconstruction efficiency can be evaluated. By selecting mono-energetic S1 signals from ^{37}Ar and $^{83\mathrm{m}}$Kr, pure S1 samples are obtained that serve as parents for the generation of fake S1s. The data driven study estimated the number of UV-photons from the reconstructed hits by accounting for DPE and the loss of SPE signals due to the digitizer threshold. While simulation driven efficiency estimates are prone to mismodeling in simulation, this data driven method does not rely on simulation. However, the selection of parent peaks can induce bias. For example, as the number of hits in the fake peak approaches the number of hits in the parent peak, the outcome is inadvertently classified as an S1, as the parent was classified as an S1. This effect was accounted for by an increased systematic uncertainty as detailed in [180].

6.2.4 S1 bias

Apart from the reconstruction efficiency, another important metric of the performance is the reconstruction "bias" that quantifies how correct the reconstructed

[4]S2s can be seen as the sum of SE signals, and they have larger area's, risetimes, and widths, but similar AFT-values.

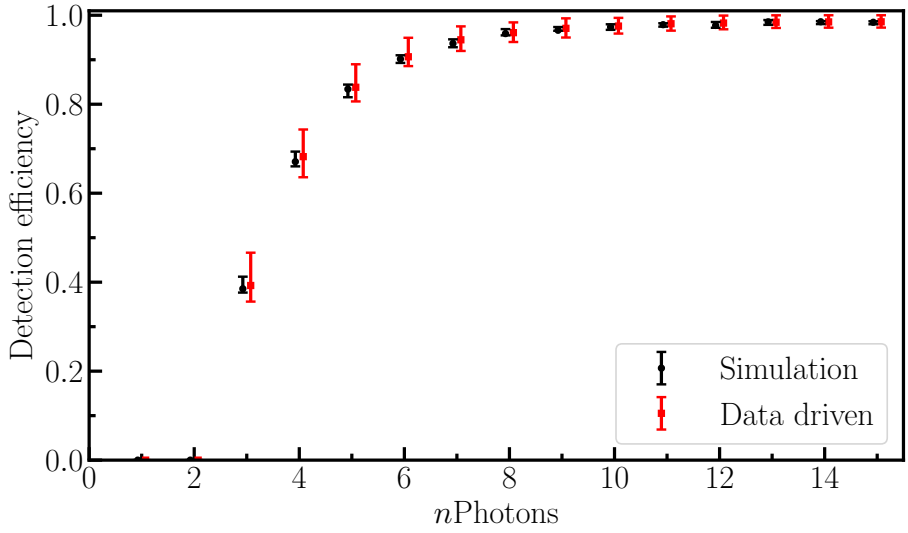

Figure 6.6: S1 acceptance using simulations and the data driven method to evaluate the detection efficiency for a threefold tight coincidence showing excellent agreement. The points of the two approaches are slightly offset for illustration purposes.

areas are, for those signals that are correctly digitized ("accepted" in the previous section). We use the same simulation as in the previous section. As before, the area (A) of a peak is the sum of the gain corrected measured charge in the PMTs, expressed in the number of photoelectrons. Ideally, we express the bias in terms of the difference between the input (A_{sim}) and output reconstructed area (A_{rec}),

$$\phi = \frac{A_{\mathrm{rec}} - A_{\mathrm{sim}}}{A_{\mathrm{sim}}}\,, \tag{6.1}$$

where A_{rec} is associated to the simulated peak after processing it with the processing framework. For some applications - like the inference of the detector response - it is sometimes more convenient to describe the bias as the difference in the number of quanta. For this purpose, we also express the bias in terms of the simulated number of freed photoelectrons ($n\mathrm{PE}_{\mathrm{sim}}$) which includes any DPE photoelectrons, as schematically shown in Fig. 6.7. For this definition, the

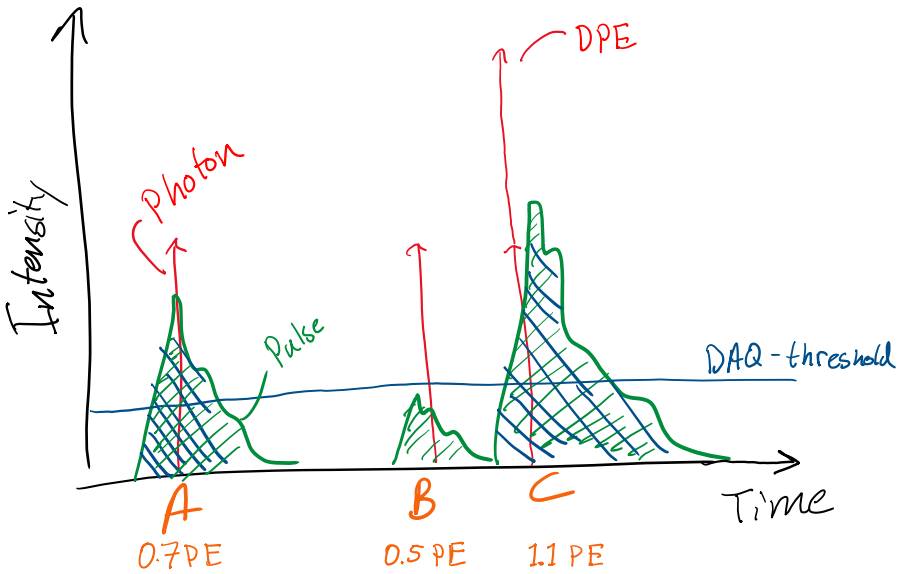

Figure 6.7: Schematic of signals in WFsim, showing the detection of three photons (labeled ABC, red upward errors). The third photon (C) causes DPE (indicated by the two vertical red arrows). As such, there are 3 photons ($n\text{Photons}_{\text{sim}}$) and 4 photoelectrons ($n\text{PE}_{\text{sim}}$). However, the photon signals must also be registered as a hit. The digitizable waveform (i.e., the three pulses/hits) is shown in green, the DAQ digitization threshold in blue. Signals A and C are above the threshold and will be digitized (hence their areas are also filled blue). Pulse B does not pass the digitization threshold and will therefore not be stored. As a consequence, while the total area (green hatches in pulses) of the pulses is 2.3 PE, the area above threshold is lower, 1.8 PE.

bias takes the similar form of:

$$\psi = \frac{A_{\text{rec}} - n\text{PE}_{\text{sim}}}{n\text{PE}_{\text{sim}}} \,. \tag{6.2}$$

Notice that ψ (Eq. 6.2) will yield a larger spread than ϕ (Eq. 6.1), since the first definition does not include any smearing due to the SPE-shape. We will use ϕ for the results in Fig. 6.9 and Fig. 6.15, which are used in the signal response modeling (in terms of quanta), and ψ for the results in Fig. 6.10 and Fig. 6.16 as these results form the basis of the bias correction discussed in section 6.5.

Fig. 6.8 shows several effects that affect the median bias. If a single photoelectron signal would always result in 1 PE reconstructed, the mean bias (and spread on the bias) would be zero. However, several effects impede this hypothetical situation:

- The detector has a finite resolution, the better the resolution, the smaller the spread on the bias.
- Photoelectron signals can be under-amplified (which we parameterize in the SPE) yielding a too small signal.
- The DAQ digitization threshold (often called self-trigger) prevents minute signals from being registered. In software, the hitfinder threshold works similarly.
- Noise can distort a signal, which can result in a positive as well as negative bias.
- Afterpulses, when merged with their progenitor peaks, will yield a positive bias.
- The reconstruction in software may reconstruct too small or too large signals, for example, if a portion of the signal is wrongfully not considered part of the peak.

The realized median S1 bias at low energies is shown in Fig. 6.9. The median bias is -2% for these small signals due to the digitization threshold. For the lowest two bins, the median bias is $\sim 0\%$, which is a selection effect of only showing the accepted S1 signals. For two (three) photons, only those signals that are sufficiently large will correctly be reconstructed because of the two- (three-) fold requirement. In this energy range, the median bias is $\sim -2\%$, albeit with a large spread due to noise and small statistics (indicated by the quantile ranges).

Fig. 6.10 shows the results for three simulations that were performed to reproduce the sources of bias as in Fig. 6.8:

- An idealized simulation where the SPE shape is replaced by a delta function at 1 PE. Additionally, the transit time spread of the PMTs was set to be 0 ns, and no extra time smearing on the arrival times is applied apart from the digitization effects imposed by the 100 MHz sampling rate. Additionally, afterpulses and noise are turned off in the simulation.

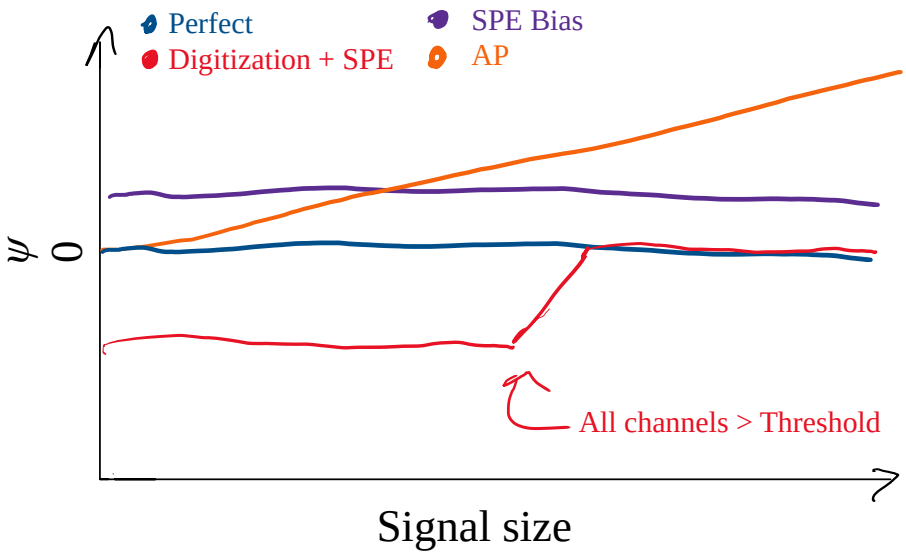

Figure 6.8: Median bias as a function of signal size following Eq. (6.2). A hypothetical "perfect" situation would result in 0% median bias for all energies. The combination of under-amplification and the digitization/hitfinder threshold induces negative bias for small signals. The inclusion of afterpulses (AP) into peaks creates a positive bias, as larger signals remain above threshold for longer, giving them a higher probability to include afterpulses. Mismodeling of the SPE shape can yield a constant offset in bias when using ψ Eq. (6.2) if the average SPE is mismodeled and not equal to 1 PE. Alternatively, for ϕ, Eq. (6.1), the same qualitative energy dependence is expected, except for the SPE bias as explained in the text.

- An afterpulse free simulation, where the SPE shape is set to a realistic model as well as including noise, without the inclusion of afterpulses.
- A realistic simulation, including realistic SPE shapes, noise and afterpulses.

For large signals, the median reconstruction bias is $\sim 0\%$ if no afterpulses are included. For the lowest bins, the threshold effects mentioned in Fig. 6.8 leads to a median bias of $\sim 0\%$. The inclusion of afterpulses does change the result as

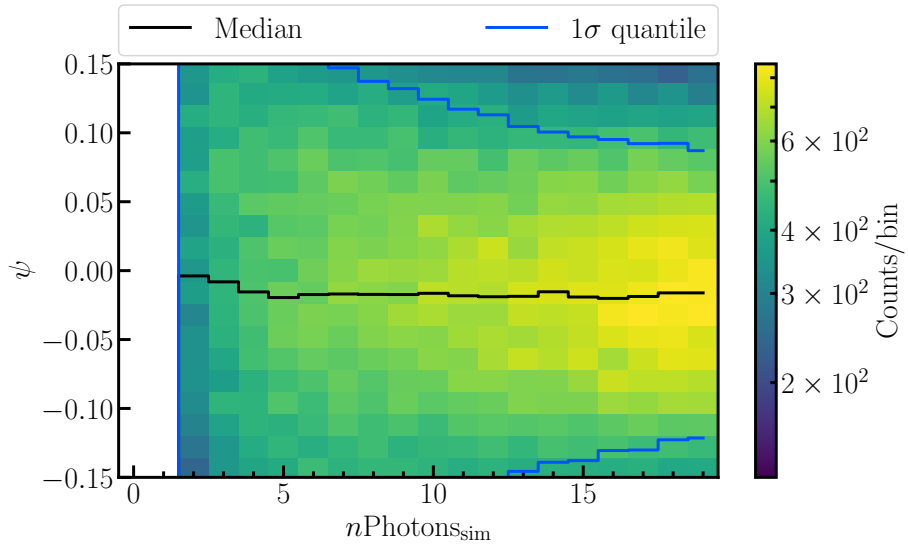

Figure 6.9: S1 bias as in Eq. (6.2) as a function of $n\text{Photons}_{\text{sim}}$ for a twofold TC requirement. At these small energies, the bias is $\sim 0\%$ for $n\text{Photons}_{\text{sim}} < 3$ and close to -2% elsewhere. The spread (1σ) of the bias as indicated by the shaded lines is mostly due to noise and small statistics, but the SPE resolution also plays a role, as the recorded area for single photons follow the SPE area distribution.

expected, leading to larger positive biases for larger peaks, as more afterpulses are included into the main peaks. The median bias is $-2\,\%$ to $0.5\,\%$ for realistic simulations, depending on the energy. The negative bias is mostly attributed to the under-amplification in combination with the digitization threshold. This study also includes the bias determination of S2s (Fig. 6.16), and these biases will be accounted for in the energy reconstruction section 6.5 At higher energies, the positive bias is attributed to the onset inclusion of afterpulses.

6.3 Ionization signals

6.3.1 S2 Shape

The ionization signals are measured by extracting electrons from the LXe and creating secondary scintillation light in the GXe. We can characterize these

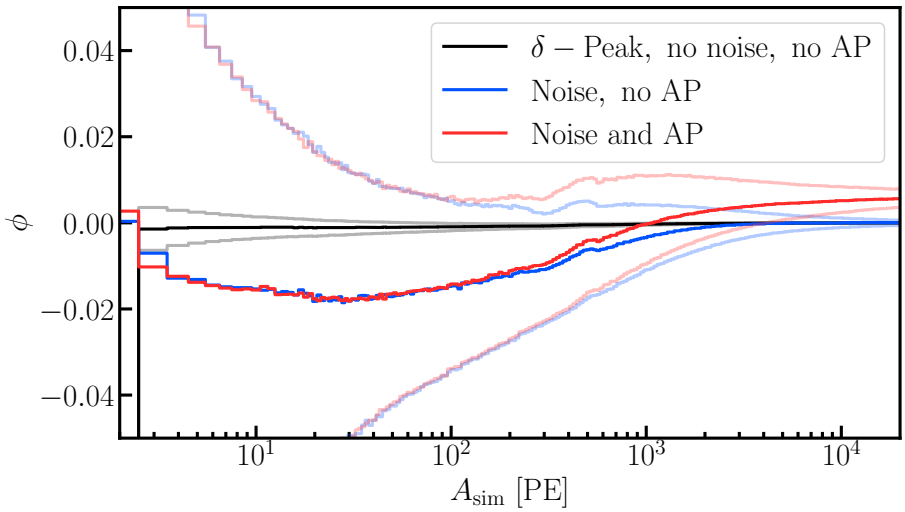

Figure 6.10: Median S1 bias according to Eq. (6.1), where several contributions of bias are isolated by simulating progressively more realistic datasets. The 1σ quantiles are shown as the shaded lines for each dataset. As we compare A_{sim} to A_{rec}, the spread in this bias curve is smaller compared to Fig. 6.9 since A_{sim} accounts for the area sampled from the SPE. The most idealized dataset simulates data with a delta-function as the SPE shape and no transit-time spread on the PMT response, without noise and afterpulses (AP). This yields an absolute median bias of $< 0.15\%$. The more realistic simulations show the trend as quantitatively explained in Fig. 6.8, and at the lowest energies, we see the selection effect as in Fig. 6.9 (only peaks with large areas are reconstructed, resulting in a positive bias).

signals by examining the response to single electrons (SEs). SE are plentiful in XENONnT and are mostly caused by photo ionization of the initial scintillation light of metal surfaces or electronegative impurities [173, 181, 182]. The scintillation light emitted by GXe is very similar to the scintillation light (and process) of LXe. The singlet and triplet scintillation times in GXe are 5.9 ± 5.5 ns and, 100 ± 8 ns, respectively [183].

Liberated electrons in the LXe are drifted upwards by the drift field between the cathode and gate. The extraction field between the gate and the anode then extracts the electrons from the LXe and accelerates them towards the anode.

As the electron is accelerated towards the anode, it excites the xenon gas [184]. The scintillation yield Y is proportional to:

$$Y \propto \left(\frac{E_g}{p} - b\right) p\, d_g\,, \tag{6.3}$$

where E_g is the electric field strength in the gas, p is the gas pressure, d_g is the gas gap length and b is an empirical offset of $\sim$1.0 kV/(bar cm) [185]. During SR0 the electric field was $E_g = 2.9\,\mathrm{kV/cm}$ and the gas gap was $d_g = 3\,\mathrm{mm}$. If the electron is sufficiently far removed from the anode mesh, the electric potential can be approximated as the potential between two parallel plates with two media inside it. When an electron comes closer to the anode wire, it experiences a stronger field as the distance r to the anode wire decreases. At a characteristic distance, the electric field starts being dominated by the proximity to one wire (instead of the anode mesh), and the electric field correspondingly increases with $1/r$ (as the electric field description near an infinitely long wire), and thus so does the scintillation light yield, as Eq. (6.3). We will show how these two regimes are observed in the SE shape in Fig. 6.12, by calculating the average SE shape for xy positions in the detector, as discussed below.

Fig. 6.11 shows the average single electron width throughout the TPC during SR0. The SEs are selected by requiring isolated signals (i.e., not relying on straxens classification boundaries) that have no other signals 2.2 ms before or after the SE with an area greater or equal to 25% of the SE. We also apply some loose cuts to prevent multiple electron signals by requiring that the width is 30 ns to 2000 ns, the area 5 PE to 100 PE, the number of hits contributing to the peak are 5 hits to 50 hits, at least four channels contributing, and that the signal has a risetime of >50 ns.

The width of the SEs in Fig. 6.12 shows several effects:

- Two perpendicular wires in the anode and cathode that were added to prevent wire sagging of these electrodes [44]. The perpendicular wires are visible as the two parallel lines across the TPC. Near these wires, two effects result in higher widths; the reduced sagging increases the width of the SE signal as the electron is drifted through the gas longer, and the infinite long wire approximation (and corresponding high electric field) does not suffice to describe the electric field near the perpendicular wire as it crosses the (non-perpendicular) wires spaced 5 mm apart.

- A region at $x \approx 10$ and $y \approx -15$ cm, where a region of clearly lower SE widths $\sim$200 ns is observed. This region is the "hotspot" region (also

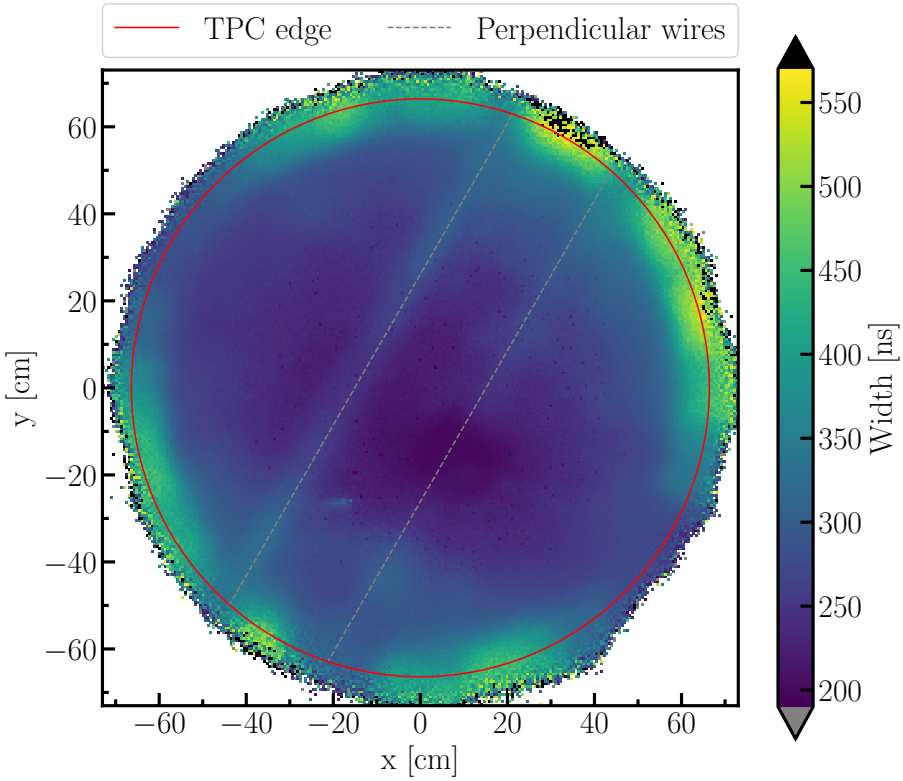

Figure 6.11: Average SE width throughout the detector.

mentioned in section 5.5). Naively, the lowered width could indicate a lower d_g. However, the temporal stability of the spot argues against such a hypothesis, as the average SE width within this region fluctuates up and down over time. The conundrum of the physical explanation of the hotspot remains unresolved at the time of writing.

- An overall sagging of wires, such that the center $x, y = 0, 0$ has a smaller gas gap. As the wires come closer to the suspension points at the edge of the TPC, the sagging naturally decreases. The sagging results in smaller SE widths for signals closer to the center of the TPC. This was also ob-

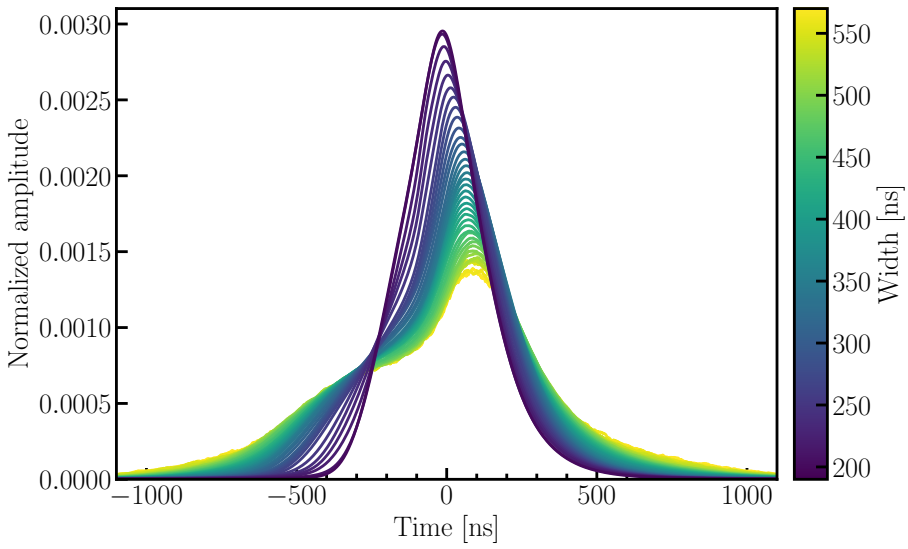

Figure 6.12: Average SE shape grouped by their widths as in Fig. 6.11 throughout the detector as described in the text.

served in XENON1T [184]. A tilt of the anode with respect to the liquid-gas interface, as observed in XENON1T, is not observed.

We show the average SE waveform for bins in xy of equal width (Fig. 6.11). Each xy bin is categorized as being of a certain width, and which for each width the average waveform is obtained by aligning the waveforms on the center of gravity of the waveform. The two regimes of E_g, described by a parallel plate approximation and by an infinitely long wire, are apparent, especially for the higher width bins. The waveform first plateaus, characteristic for the parallel plate electric field, before rising again when the electric field increases close to an anode wire. Due to the relatively low extraction field of 2.9 kV/cm and the aforementioned gas gap, signals from a single electron are roughly twice as wide as in XENON1T [184].

The S2 signal shape is, in addition to the response of the detector to individual electrons, also dictated by the collective behavior of the electrons. As electrons are drifted upwards to the gas-liquid interface, they diffuse longitudinally and transversely to the direction of the electric drift field. For the times at

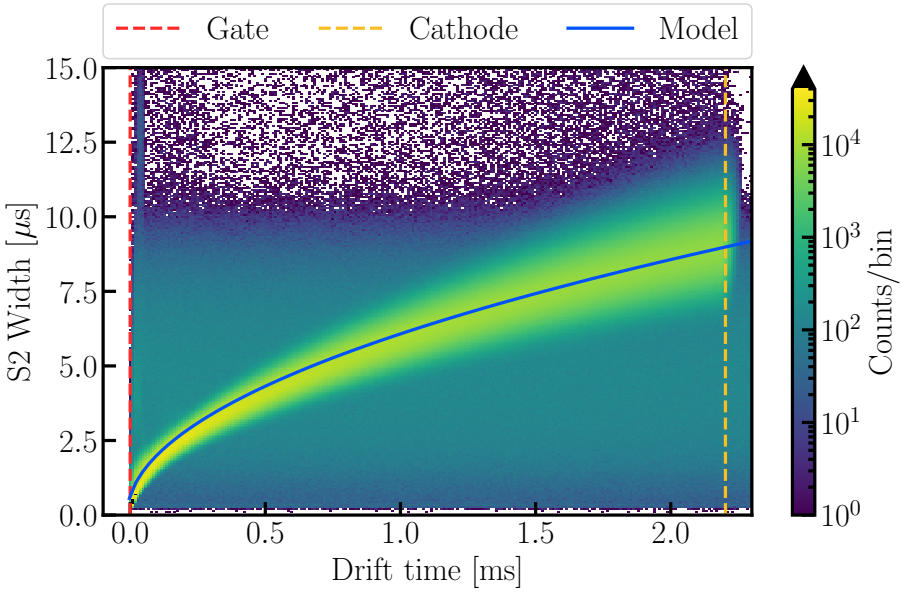

Figure 6.13: The S2 width as a function of drift time (time difference between the S1 and S2). Due to diffusion, the S2 width increases for longer drift times. For SR0, the diffusion model was fitted to be $D = 45.7\,\mathrm{cm/s^2}$, $v = 0.675\,\mathrm{mm/\mu s}$, and $\sigma_0 = 375\,\mathrm{ns}$ [186]. Only minimal cuts are applied to this data, and accidental coincidence signals (such as S1s from below the cathode paired with delayed electrons) show up as the band of events that do not follow the diffusion model.

which the S2 is measured (the shape) the longitudinal drift is most important. For a point-like source drifting for time t under standard Brownian motion, the distribution broadens:

$$\Delta t_\sigma = \sqrt{\frac{2Dt}{v}}\,, \tag{6.4}$$

where v is the drift velocity, D the diffusion constant and Δt_σ the spread of the signal. To account for the non-zero width of the single electrons, we can adapt Eq. (6.4):

$$\Delta t_\sigma = \sqrt{\frac{2Dt}{v} + \sigma_0^2}\,, \tag{6.5}$$

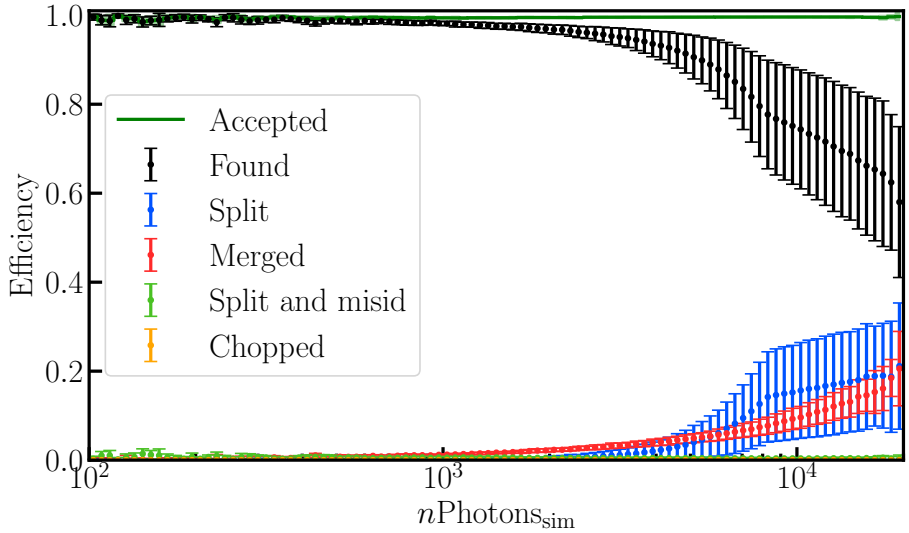

Figure 6.14: S2 acceptance, as the size of the S2s increases, so does the probability that the S2 is either split - or merged to an afterpulse or delayed electron. The "accepted" fraction of the peaks accounts for these effects.

where σ_0 is the broadening of the S2 due to the width of the secondary scintillation. Finally, we have to convert the Gaussian width $\Delta t_\sigma \sim 1.349 \Delta t_{50\%}$ where $\Delta t_{50\%}$ is the time range that contains 50% of the area of the peak:

$$\Delta t_{50\%} = \sqrt{\frac{3.64 D t}{v} + 1.82\sigma_0^2}\,. \tag{6.6}$$

Fig. 6.13 shows the diffusion for $^{83\text{m}}$Kr data with only minimal quality cuts applied, and only considering the region that is at least 4.45 cm removed from the position of the perpendicular wires, as the S2 shape is different near these wires.

6.3.2 S2 reconstruction

Fig. 6.14 shows the efficiency of the S2 reconstruction. One notable difference with the S1 efficiency is the difference in the amount of split and chopped S2 peaks, especially for S2s that originate from interactions deeper in the detector

as they diffuse for a longer time. Similar to S1s, the reason for allowing an S2 to be split over multiple peaks is that S2s are easily split due to their discrete electron nature. We consider an S2 "accepted" if the peak is found, split, split and misid, merged[5], or chopped, and the peak has at least 85% of the area of the original peak. Fig. 6.14 shows that following these requirements, the S2 acceptance is $\sim 100\%$ over the entire energy range from 100 PE[6] to 2×10^4 PE (corresponding to $\sim 3 - 10^3$ electrons).

Like for S1s, the inclusion of afterpulses leads to a positive median bias at higher energies (see Fig. 6.15 and Fig. 6.10). At lower energies ($n\text{Photons}_{\text{sim}} \lesssim 10^4$, $\text{A}_{\text{sim}} \lesssim 10^4$ PE), the median bias is negative due to the digitization threshold. Since also small S2s are always digitized and the acceptance is $\sim 100\%$ for the entire energy range, there is no apparent selection effect in the bias for the lowest energy S2s (as there was for the S1 bias). The results are otherwise similar to the S1 results, Fig. 6.15 serves as the input for the detector response modeling, and we use a simulation similar to Fig. 6.10 to discuss the energy response of the detector (section 6.5).

6.4 Energy threshold

The S1 detection efficiency directly relates to the energy threshold of the detector for analyses requiring both S1 and S2 signals. The energy threshold for XENONnT, based on the work in the previous section, is shown in Fig. 6.17. In XENON1T [106], an excess of ER events was observed, most prominently in the energy region of 2 keV to 3 keV, which is close to the energy threshold of 1 keV. In XENONnT a similar search was performed to validate these results [97], a comparison between the XENON1T data and XENONnT data will be shown in the next chapter (Fig. 7.3). It was vital to understand the response of the detector near the energy threshold, where the S1 detection efficiency determined in subsection 6.2.2 was used as a direct input. Similarly, for future studies of WIMPs, CEνNS, and WIMPs through the Migdal effect, understanding the energy threshold will be required as the signals for these analyses are expected at low energies, driving the push for lowering the energy threshold [4, 118, 132].

The energy threshold (Fig. 6.17) is determined by converting the detection efficiency expressed in the number of photons detected to the "combined" energy scale Eq. (3.8). For future searches, work is ongoing to lower the Tight

[5] Where merged S2s are S2s "merged" with a following delayed extraction electron signal

[6] Well below the S2 thresholds of 500 PE for the the low energy ER analysis and 200 PE for the WIMP analysis (Chapter 7).

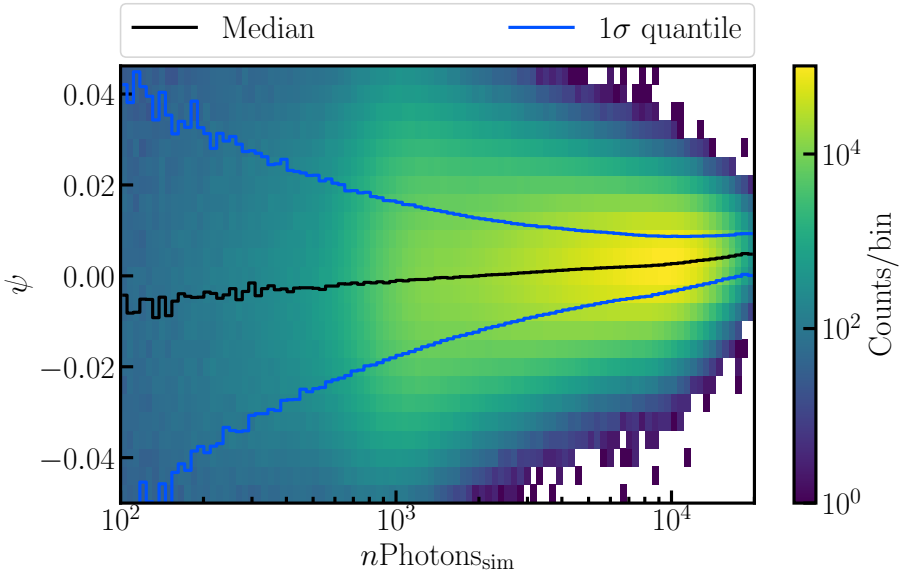

Figure 6.15: S2 bias (Eq. 6.2) as function of of $n\text{Photons}_{\text{sim}}$.

Coincidence (TC) requirement to twofold. Using the results in subsection 6.2.2, we can compare the energy threshold that will be obtained using this lowered S1 detection requirement in Fig. 6.17. The energy threshold is defined as the energy where 10% of the signals are detected [118], which corresponds to $1.0\,\text{keV}_{\text{er}}$ and $0.7\,\text{keV}_{\text{er}}$ for a three- and twofold coincidence threshold, respectively. Lowering the TC requirement will come with additional stringent cuts to remove the extra Accidental Coincidence (AC) background.

6.5 Light and charge yield

The energy of an event is reconstructed based on the number of photons n_{ph} and the number of electrons n_e. To convert the S1 and S2 signals to the energy of the interaction, there are two conversion factors $g_1 = S1/n_{ph}$ and $g_2 = S2/n_e$ that relate the average S1 and S2 area to number of photons and number of electrons. The values of g_1 and g_2 depend on the light and charge yield and are detector specific parameters. To calibrate the detector with a linear energy

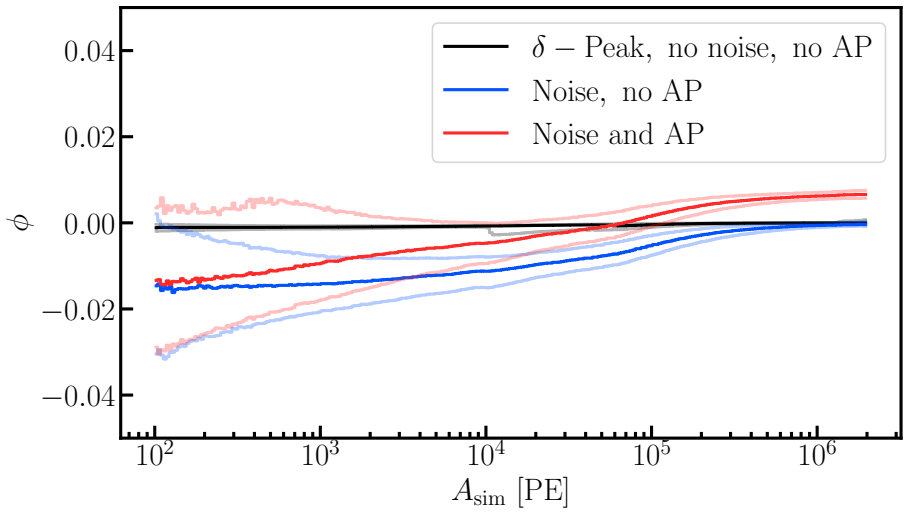

Figure 6.16: Realized median S2 bias and $\pm 1\sigma$ quantiles (shaded lines), see caption of Fig. 6.10 for more details.

response of both S1 and S2 signals, it is convenient to rewrite Eq. (3.8) to:

$$\frac{S2}{E} = \frac{g_2}{W} - \frac{g_2}{g_1}\frac{S1}{E}\,. \tag{6.7}$$

By plotting the $S1/E$ and $S2/E$ (the "Doke" plot [187]), a linear relation should appear that allows us to fit g_1 and g_2 simultaneously.

However, during SR0, the Doke plot (using real data) showed significant non-linearity. We attribute this to the reconstruction bias as we have seen in Fig. 6.10 and Fig. 6.16. The reconstructed data may be biased up to a few percent depending on the energy of the interaction.

Using many calibration sources, the values of g_1 and g_2 for XENONnT were inferred, but the data points deviated significantly from a linear energy scale. Below, we will illustrate how a linear energy scale can seem non-linear due to the several effects that bias has on the energy reconstruction. To this end, we simulated mono-energetic sources that are also used in data to infer the values of g_1 and g_2. This is a slight oversimplification - especially for a source like $^{83\mathrm{m}}$Kr (which decays in two steps to ^{83}Kr via an excited state which relaxes after 154 ns to the ground state) but as the goal is to evaluate Eq. (6.7) for

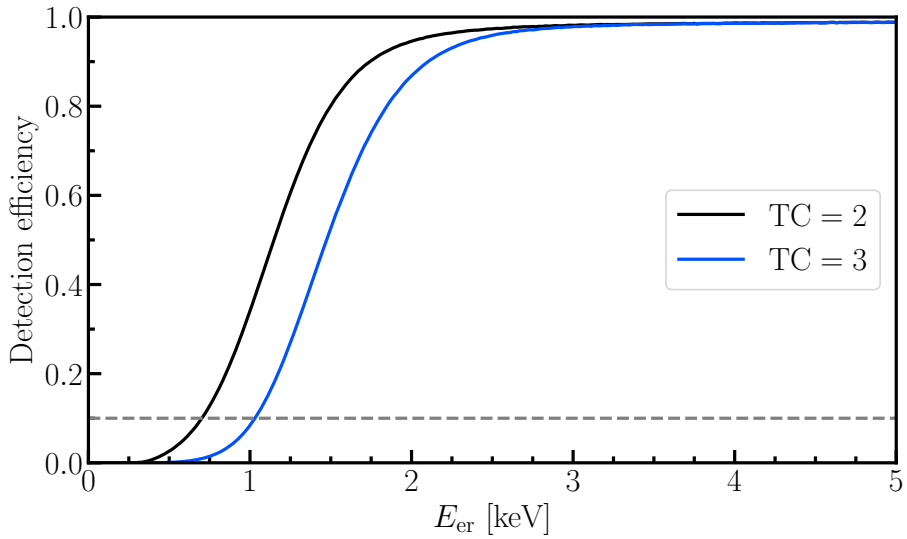

Figure 6.17: The detection efficiency for different TC requirements. From these curves, the energy threshold (the energy where the detection efficiency starts being greater than 10%) are determined to be 1.0 $\mathrm{keV_{er}}$ and 0.7 $\mathrm{keV_{er}}$ for a three- and twofold coincidence threshold.

several calibration sources it is a justifiable simplification. We computed the median bias at these energies using Eq. (6.1). Additionally, we compare the areas in cS1/cS2 space which are position-corrected S1 and S2 areas (cS1 and cS2) to compensate for position-dependent detector effects like light collection efficiency and the electron lifetime.

Fig. 6.18 shows where each of the energy deposits end up in the Doke plot. We rewrite Equation 6.7 in terms of $Q_{ee} = cS2/E$ and $L_y = cS1/E$ such that we find that Q_{ee} should be linear with L_y:

$$Q_{ee} = -\frac{g_2}{g_1} L_y + \frac{g_2}{W} \,. \tag{6.8}$$

By correcting the values of Q_{ee} and L_y with the median bias-values for S1s and S2s at their respective energies, we can regain linearity as shown on the right-hand side of Fig. 6.18.

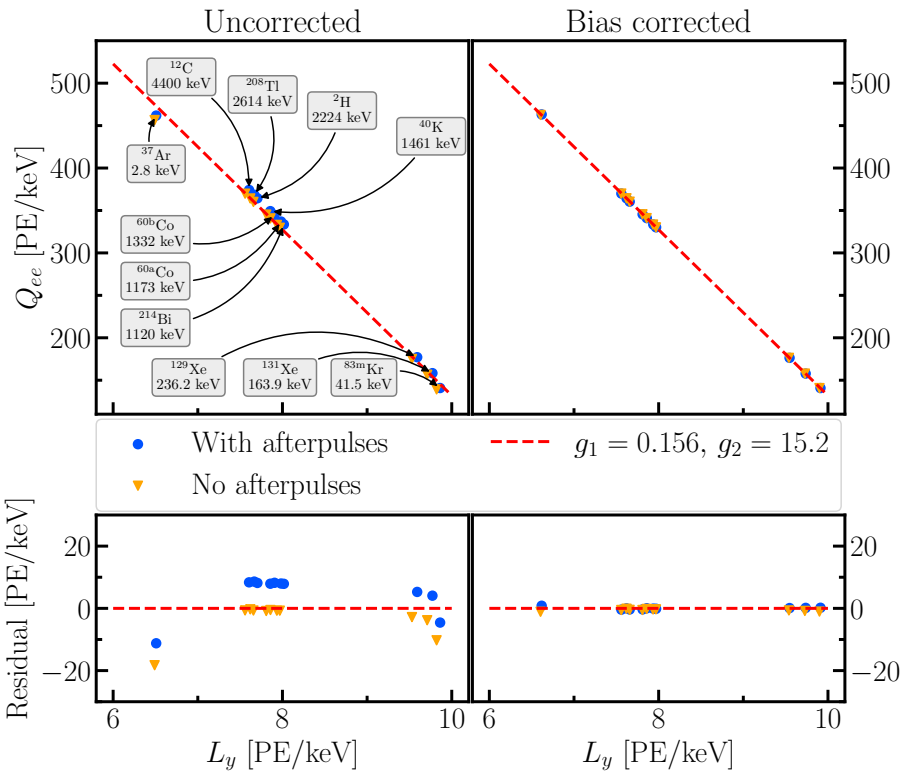

Figure 6.18: Doke plot using simulated data. The data points (with bias included) are shown on the top left plot, indicating a non-linear energy scale. The right panels show the same data but shifted by their respective bias values using the dedicated simulation. We obtain the values of $g_1 = 0.156$, $g_2 = 15.2$ from a fit to the data using the "With afterpulses" corrected data points. The non-linearity is even clearer on the bottom left plot, that shows the residuals of the fit. The non-linearity does not get worse (or better) if we include afterpulses (bottom panels), due to the fact that the DAQ digitization threshold has the largest effect on the bias (and the largest non-linear contribution), see Fig. 6.10 and Fig. 6.16.

These results were also used to correct the real data - in order to extract g_1 and g_2, as shown in Fig. 6.19. The simulation (Fig. 6.18) and data (Fig. 6.19)

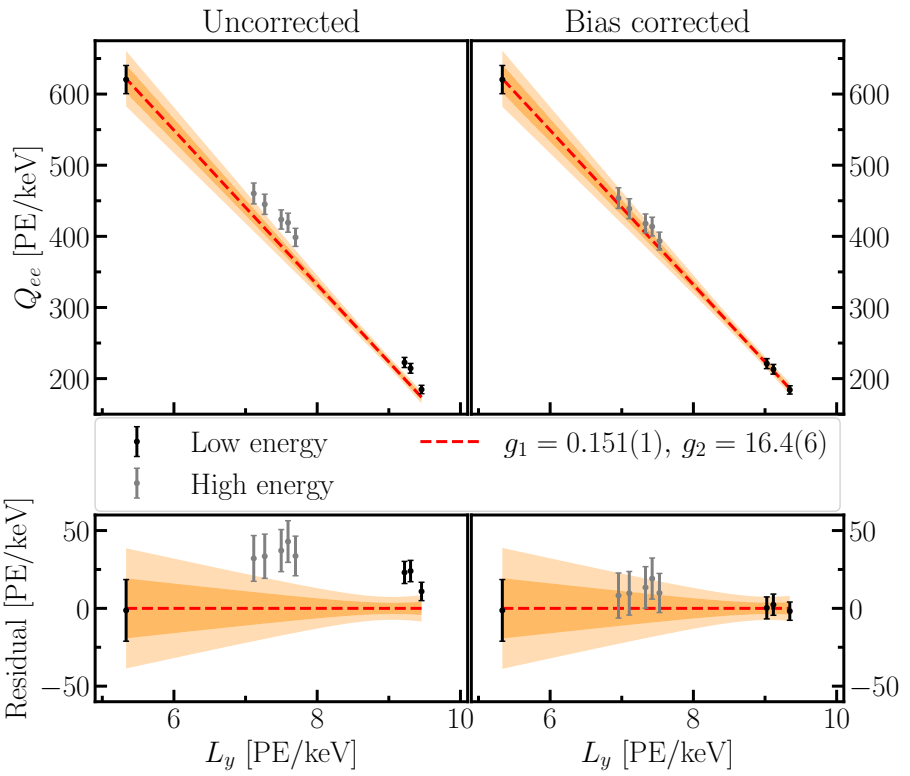

Figure 6.19: Realized Doke plot using real data when not corrected for reconstruction bias (top left) and using the median bias values for the sources as obtained in this chapter (top right). The low energy data points from the top right plot are fitted [188, 189], and shown for reference on the top left plot too. The residuals in the bottom panel also show the residuals with respect to the fit from the top right panel.

show that while the data and simulation match for the peak-shape parameters [178, 179], the charge- and light yield of these sources show differences of up to 30% (comparing the axes of the top left plot), and unsurprisingly, the g_1, g_2 values from the data are different from the simulation. Yet - as we have seen in Fig. 6.10 and Fig. 6.16, the change in the median bias is gradual, and

a 30% offset in the area would yield similar values for the median bias. To account for this, a systematic error was added to the results obtained from the bias corrected real data [188]. Additionally, for the low ER [149] and WIMP search, the primary focus was the low energy part of the spectrum. Therefore, the fit shown in Fig. 6.19 included the data points from the low energy data points of ^{37}Ar, ^{83m}Kr, ^{131}Xe and ^{129}Xe.

In this chapter, we have validated the reconstruction chain for XENONnT, as used during SR0. We have shown how two key analyses, that of the energy threshold and that of the energy reconstruction, are directly related to the reconstruction performance. In the next, final section, we will briefly discuss the WIMP results of XENONnT.

Chapter 7

First WIMP results with XENONnT

This chapter reports on the latest results of XENONnT in the search for WIMP Dark Matter (DM). It was performed in parallel to the low energy electronic recoil (ER) search [149] and focuses on detecting nuclear recoil (NR) signals that would be induced by Weakly Interacting Massive Particles (WIMPs).

7.1 Science run 0

The data taking took place during the first science run (SR0) of XENONnT, between May 3, 2021, and December 10, 2021. SR0 started and ended with calibration periods, the science data of SR0 was collected between July 6 and November 10, yielding a total of 97.1 d of exposure during this 127 d, see Fig. 7.1. The Data Acquisition system (DAQ) busy signals reduced the effective exposure by 0.05%[1]. Additionally accounting for the Neutron Veto (NV) and Muon Veto (MV) veto periods leads a total 2% reduction of the effective exposure, such that the total effective exposure is 95.1 d. The average uptime (when we collected science data, corrected for busy/veto time intervals) in the period between July 6 and November 10 is 75%, see Fig. 7.1. The contributors to a decreased uptime

[1]The actual DAQ-deadtime (when the DAQ does not record data) is smaller than the 0.05% exposure loss (when events cannot be completely reconstructed). Events that are close (< 2.2 ms, the drift length) to a busy signal are discarded, even when the DAQ was recording data as the event can be only partially reconstructed, as discussed in Chapter 5.

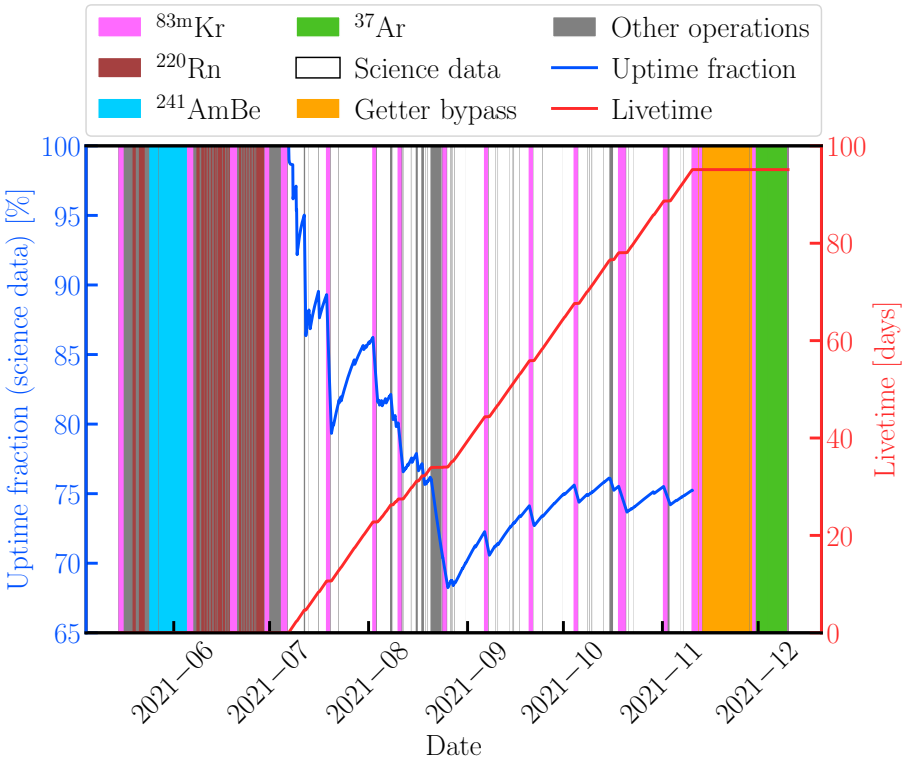

Figure 7.1: The cumulative average uptime fraction (left) and the accumulated livetime (right) of SR0 during the science data taking. After SR0 a special mode "getter bypass" aimed to enhance the tritium concentration in the detector as it was a possible explanation of the XENON1T low ER energy excess (discussed below), see Ref. [149] for more details.

are mostly the "other operations" (e.g., LED calibration, non-SR0 electrode configurations, maintenance), the weekly ^{83m}Kr calibrations (discussed below), and the $\sim$ 3% time where the DAQ is transitioning from one run to another (one minute restart time for thirty minute long runs)[2].

[2]This is decreased $\sim 12\times$ after SR0 by increasing the run duration to 6 h.

During SR0, the gain of the 494 PMTs was set to be $\sim 2 \times 10^6$ and regular gain calibrations ensured that any time dependencies in the gains, caused by effects like temperature changes, prolonged exposure to high illumination, and voltage changes, were calibrated out. This yielded an estimated gain stability of 3% after correction. In total, 17 PMTs were excluded (in hardware or software) because of stability issues, high noise, high afterpulse rates, vacuum degradation or light emission.

The stability of the detector was monitored by weekly ^{83m}Kr calibrations. ^{83m}Kr decays with a half-life of 1.83 h predominantly to ^{83}Kr under the emission of two conversion electrons with energies of 32.1 keV and 9.4 keV. The intermediate state of the excited ^{83}Kr after the first decay has a half-life of 154 ns. ^{83m}Kr mixes homogeneously through the liquid xenon (LXe) and is therefore suited to calibrate the full LXe volume. ^{83m}Kr was used to calibrate and monitor the light- and charge-yield of the detector. Additionally, the ^{83m}Kr data was used to measure the field distortion effects on the position of the S2. Due to an inhomogeneous drift field, S2s signals from deeper in the detector are reconstructed radially inwards, as the electrons drift to smaller r. This effect is shown in Fig. 5.3 panel F, which shows the data before the field distortion correction.

The inhomogeneous drift field is in part the result of the relatively low high voltage applied to the electrodes (see Fig. 3.2, where each PMT array is accompanied by a protective screening mesh to screen the PMTs from high electric fields). A short-circuit occurred just before SR0 between the cathode and the bottom screening mesh, limiting the cathode voltage to -2.75 kV, while the gate was set to 0.3 kV, and the anode to 4.9 kV, resulting in a drift field of 23 V/cm and an extraction field of 2.9 kV/cm. The voltage applied to the cathode (and bottom screening mesh) was limited to this value, as a too low voltage might damage the bottom PMTs. The additional rings that surround the PTFE walls of the TPC were set to 650 V to reduce the inhomogeneity of the drift field. Due to the limited voltage on the cathode electrode, the gate electrode had to be set to a higher voltage than anticipated. The drift field was therefore $\sim 10\times$ smaller than the design goal of 200 V/cm [148] and the realized extraction efficiency was $53 \pm 3\%$ [149] versus the design goal of 96%. Due to the low extraction field, the average single electron gain was 31.2 ± 1.0 PE/electron [149].

7.2 Electronic recoils

The data of SR0 was used both for a low energy ER search [149] and WIMP search [6]. The response to ER signals was calibrated using a combination

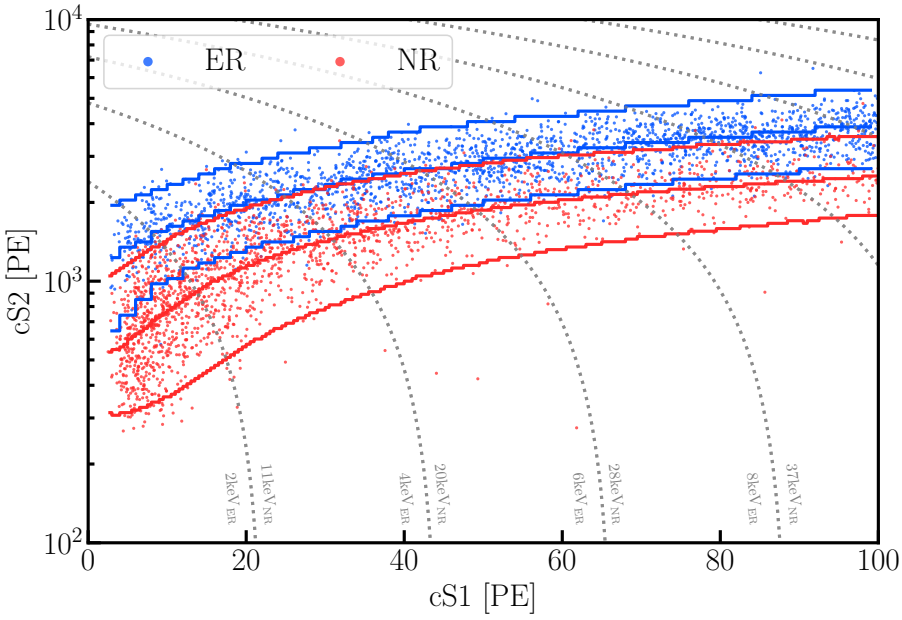

Figure 7.2: The ER (blue) and NR (red) response of the detector calibrated with ^{220}Rn and ^{241}AmBe respectively. The red and blue lines show the median and the $\pm 2\sigma$-quantiles for the ER and NR calibration data. Dashed gray lines show contours of equal energy for ER and NR signals.

of ^{37}Ar [190] (decaying through the emission of two mono-energetic lines at 2.82 keV and 0.27 keV) and ^{220}Rn [191] calibration sources. A daughter isotope of ^{220}Rn, ^{214}Pb emits a β-spectrum at low energies that makes it well suited to calibrate the detector at low energies (the Region of interest (ROI) being $[1, 30]\,\mathrm{keV_{er}}$ [149]). The ER response in cS1-cS2 space, the ER-band, is shown in Fig. 7.2. The XENON1T low energy ER search [106] had shown an excess which could be interpreted as tritium[3], but also as beyond the SM physics. Due to the lowered ER background, XENONnT was much more sensitive in the relevant energy region. However, the excess as in XENON1T was not observed in XENONnT [149], indicating that the cause of the XENON1T excess was not beyond the SM physics. Figure Fig. 7.3 compares the data (and background-

[3]That emits a β-spectrum with a Q value of 18.6 keV and has a half-life of 12.3 yr.

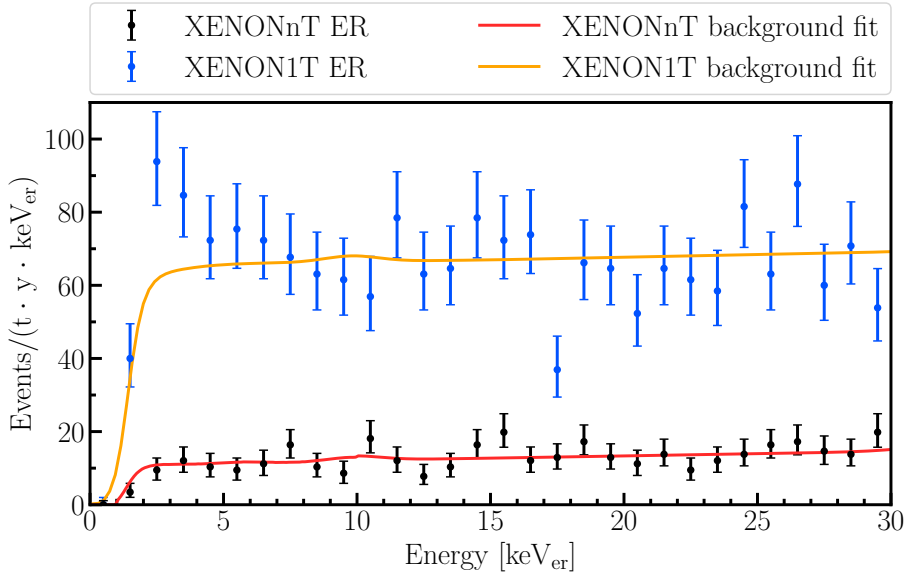

Figure 7.3: ER background fits for XENONnT [149] and XENON1T [106]. The XENONnT background is $\sim 5\times$ lower than that of XENON1T. No excess - as observed in XENON1T - is observed in the XENONnT data.

only fit) for both experiments, clearly showing the reduction of ER background and no excess in the low energy region of XENONnT.

7.3 Nuclear recoils

The WIMP search looks for WIMP induced NR scatters, for which there are fewer sources of background than for ER. The data below the -2σ quantile in cS2 of the ER band was blinded for the WIMP search. The NR band was calibrated using an ^{241}AmBe external calibration source. The ^{241}AmBe source emits neutrons via an alpha-capture reaction $^{9}\mathrm{Be}(\alpha, n)^{12}\mathrm{C}$, where the α decay of $^{241}\mathrm{Am}\rightarrow^{237}\mathrm{Np}+\alpha$ causes the ^{9}Be to emit a neutron $^{9}\mathrm{Be}+\alpha\rightarrow^{12}\mathrm{C}+n$. The energy of the emitted neutrons follows a spectrum ($\lesssim 10\,\mathrm{MeV}$) with many peaks, that produce $\mathcal{O}(10\,\mathrm{keV_{nr}})$ NR scatters in the LXe. The NR sample was obtained by selecting events in the TPC based on a coincidence of detecting a corresponding γ in the neutron veto. The alpha-capture reaction $^{9}\mathrm{Be}(\alpha, n)^{12}\mathrm{C}$

results in 60% of the cases in an excited state of ^{12}C that subsequently emits a 4.4 MeV γ. Tagging TPC events that occur in coincidence with the detection of a 4.4 MeV γ by the neutron veto allows to obtain a very pure n sample [192]. As explained in section 3.1, NR signals have a lower cS2/cS1 ratio, as is visible in Fig. 7.2. However, the lower extraction efficiency causes the ER/NR bands to overlap more that anticipated [148]. This impacts the WIMP search as it decreases the separation power between the ER and NR signals, as events in the NR band have a higher probability of being ER background signals[4]. This diminishes the discovery power for WIMPs (like all sources of background, see subsection 2.3.3). Fortunately, the ER background is also very low (as in Fig. 7.3), partially mitigating this issue. To compare, the LZ collaboration achieved a better ER/NR separation due to a higher extraction field [5]. With the probability of the events below the NR-median band being NRs of 99.75% for LZ [193] and 99.5% for XENONnT.

The WIMP search uses only well reconstructed NR events by applying quality cuts [149], and removing double scatter events since WIMPs are not expected to scatter twice in the detector (section 2.3). The TPC was partitioned into two sections based on the proximity in xy (<4.45 cm) to a perpendicular wire, and the two regions (near-wire or off-wire) were taken into account as separate likelihood terms in the statistical analysis of the data. As we have seen in subsection 6.3.1, the shape of the S2 signal is different for the near-wire region. Therefore, the events in this region are treated differently with respect to the rest of the TPC, most prominently, by S2 quality cuts like the S2-width. The statistical analysis of the data was performed on the central 4.18 ± 0.13 t fiducial mass, with the fiducial volume defined by the shape as in Fig. 7.4. Several sources of background are accounted for in the statistical analysis of the data: accidental coincidences, radiogenic sources, wall events, CEνNS, and ER-leakage (ER events that "leak" into the NR band). The distribution of the unblinded dataset in cS1, cS2 space and the background models are shown in Fig. 7.5.

7.4 Statistical analysis

The statistical analysis was performed in a blinded manner to reduce human induced bias. The statistical analysis uses an extended unbinned likelihood in

[4] And, strictly speaking, vice versa, but the NR backgrounds are $\mathcal{O}(1 - 10\,\text{events}/(\text{t} \cdot \text{y}))$, compared to $\sim 460 \pm 40\,\text{events}/(\text{t} \cdot \text{yr})$ in the $[1, 30]$ keV$_{\text{er}}$ ROI (from the background rate of $15.8 \pm 1.3\,\text{events}/(\text{t} \cdot \text{yr} \cdot \text{keV})$ [149]).

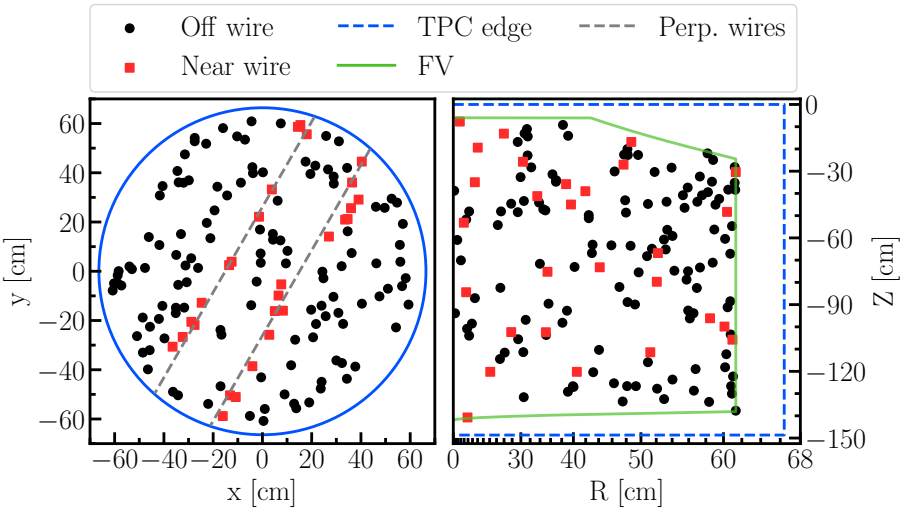

Figure 7.4: Spatial distribution in reconstructed xy (left) and reconstructed depth Z vs. reconstructed radius R (right) of events after unblinding within the fiducial volume (FV). Events are subdivided into "off-wire" and "near-wire" based on their reconstructed proximity to the perpendicular wires (gray). Figure adapted from [6].

cS1, cS2, and R. The likelihood of a WIMP signal given the data was evaluated for each WIMP mass. Ahead of the unblinding, the following criteria were decided: if the best fit WIMP would have a $> 2\sigma$ significance, a two-sided interval would be reported. If the significance would be $< 2\sigma$, a one-sided 90% confidence upper limit of the cross section (as discussed in subsection 2.3.3) would be reported. After unblinding, the dataset consisted of 152 events, of which 16 were in the blinded region. A higher ER leakage fraction is observed for $\gtrsim 50\,\mathrm{PE}$ as shown in Fig. 7.5. Additionally, a bug was discovered after unblinding in the neutron veto processing chain that meant that the neutron background prediction turned out to be too low. A post unblinding change accounted for this increased neutron background prediction. This change scaled up the neutron prediction by using neutron veto tagged events. No significant excess of events above the background model was found, and a new exclusion limit was set. In Ref. [6], the limits are presented using a conservative approach to prevent setting limits below the median sensitivity. Previously, limits were

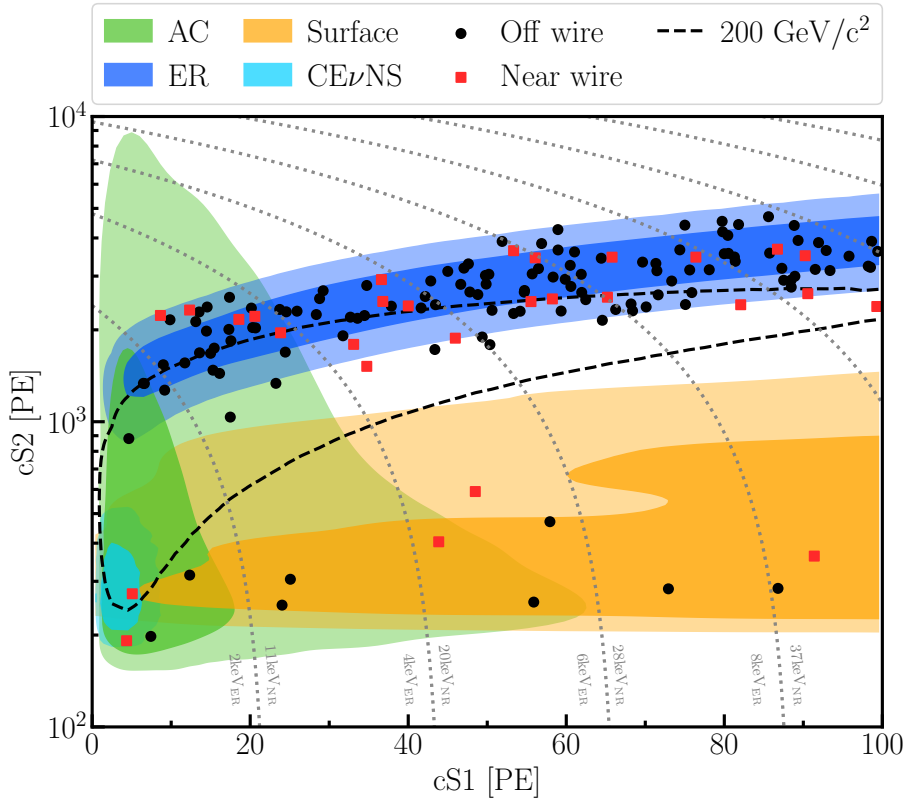

Figure 7.5: Events after unblinding and signal models shown as 1σ (dark) and 2σ (light) contours. The dashed line indicates where the 2σ signal region of a $200\,\mathrm{GeV}/c^2$ WIMP would be expected.

presented down to the -2σ sensitivity band [5, 31]. The exclusion limit at 90% confidence level is shown in Fig. 7.6, with the lowest upper limit of $\sigma_{n,\mathrm{S.I.}} = 2.6\times10^{-47}\,\mathrm{cm}^2$ at a WIMP mass of $28\,\mathrm{GeV}/c^2$. The new exclusion limit is lower than the XENON1T [4] result. The result of LZ [5], that operated a similar detector observed a very large under fluctuation of their background, yielding a very stringent limit that is lower for all masses $M_\chi \gtrsim 9\,\mathrm{GeV}/c^2$ (which is the lowest WIMP mass for which they reported their result).

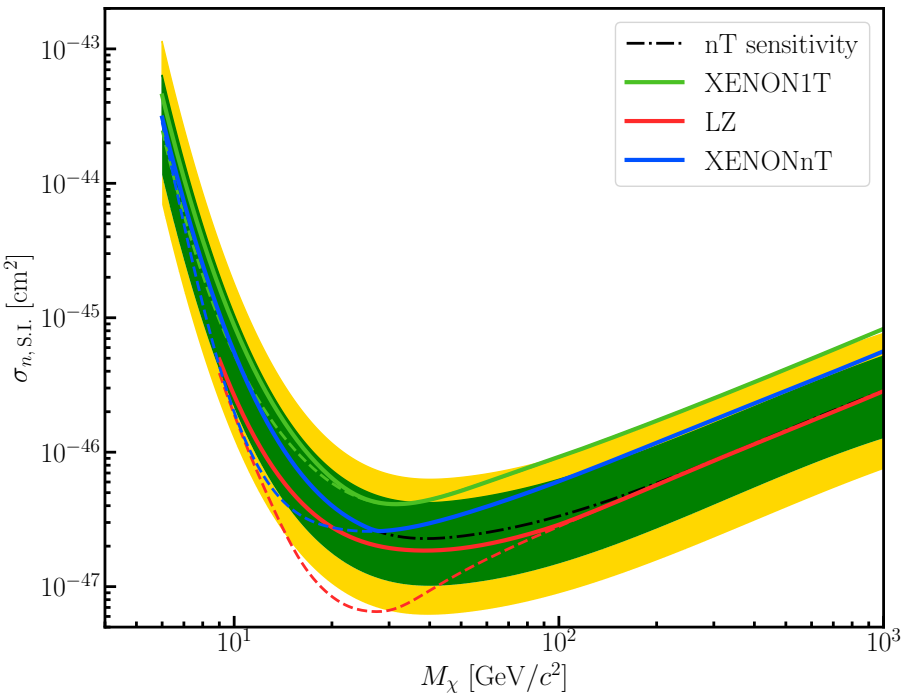

Figure 7.6: Upper exclusion limit of the spin independent WIMP-nucleon cross section as a function of WIMP mass at 90% confidence level (blue line), and sensitivity (dash dotted black line, green 1σ band, and yellow 2σ band). The solid lines show the exlusion using the conservative limit setting as in Ref. [6] (as described in the text), the dashed lines follow the earlier conventions [31]. The LZ result is indicated by the red lines [5] and XENON1T by the green lines [4]. Figure adapted from Ref. [6].

7.5 Outlook

After finishing SR0, the experiment continues to take data in Science Run 1 (SR1). During SR1 the electron lifetime continues to be >10 ms, and the ER backgrounds due to the decay chain of ^{222}Rn are further reduced down to <1 µBq/kg. Additionally, the neutron veto is being prepared for doping with Gd-Sulfate to further increase the neutron veto tagging efficiency (section 3.3).

After SR1, there may be a time to improve the electric fields by opening up the detector and replacing the faulty electrodes to fix the low drift and extraction fields to improve the ER/NR separation, remove the inadequately modeled perpendicular wires, reduce the large drift that causes long events and high data rates, and hopefully resolve the hotspot occurrences.

Acknowledgements

It has been a true privilege to do a PhD in the Nikhef Dark Matter group and be able to work on a huge range of topics with inspiring colleagues all around me.

Without my promotors, Auke-Pieter and Patrick, I would never have been able to spend these four thrilling years in the XENON experiment. You gave me the freedom to work on whichever topic I considered interesting and were always able to provide spot-on feedback. For me, you were the best promotors I could have wished for, striking the right balance between making me fully responsible for my own work and yet providing guidance when needed. Auke-Pieter, above all, you taught me how to be exact but also put things in context, both in physics where we are fond of scrutinizing 0.1% effects while the uncertainties on our limit plots are measured in orders of magnitude, as well as pointing out the true important topics of in person collaboration meetings. Patrick, your patient and forward-looking supervision always pointed me in the right direction. As just one example, shortly after receiving my first login credentials to the computing cluster, you already sent me to an expert-software workshop in Chicago, which really paid off later.

I'm grateful to all the members of my committee: Flavia de Almeida Dias, Gianfranco Bertone, Robert Fleischer, Tina Pollman, Marc Schumann, and Ivo van Vulpen for reading my thesis and being my opponents during my defense, where you will undoubtedly ask excellent questions. Thank you, Robert, for also being my C3 member. Our meetings were always a great moment of reflection, providing both great support as well as a roadmap for the period ahead.

To do scientific research is to stand on the shoulders of giants, and I have been fortunate with the giants whose work I could build on. Thank you Jelle, Sander, Erik, and Dan for being these giants and having been able to work with you for the brief period at the start of my PhD. This time was a true slingshot into the Dark Matter hunt. Jelle, in my view you are one of the greatest role

models of what a scientist should be, and you are always happy to discuss physics, statistics, or coding, while also giving the most constructive feedback that truly motivates the people around you. Sander, your enthusiasm and great mentorship during my master-project led me to this PhD (after some detours), for which I'm very grateful. You also really put me on track getting involved with the XENON1T data analysis. Erik, even though our overlap as PhDs was very short, you were able to teach me a lot in those few months, and I've greatly benefitted from your advice while I was still getting started with my PhD.

The most valuable resource in academia is the people that work in it, and Nikhef is a prime example of that, especially the Dark Matter group. Peter, you were there every stage during my PhD, always happy to discuss the detector physics, the latest XENON gossip, or the date of the next "berkfest". Thank you Maricke, for always being eager to take on more interesting topics (how you are able to work on so many, different, projects always amazes me), and for running the DAQ group while also making everyone tune in to the weekly DAQ meeting. Carlo, you have just started as a PhD, but you are already producing great work for the experiment. Serena, Stefan, Alvaro, Tina, Maxime, Kelly, Marjolein, James, and Emily, thank you for being such great colleagues with whom I shared so many interesting on- and off-topic discussions. Barbara, it was wonderful to supervise your master's project, and I'm very happy you chose to continue as a PhD at the Weizmann Institute.

Thank you, Darryl, for all the fun we had during the commissioning of the XENONnT DAQ. We spent many late nights, weekends, 4 a.m. calls, and even more chats building a truly amazing system. We had to commission the DAQ during the pandemic with a timeline that ended before the PMTs were even connected, with zero downtime allowed. When I look back on my PhD, these moments where we pulled everything together with a monumental dose of sarcasm are among the most valued. Alexey, Dan, Peter, Alvaro, Andrea, and Florian, it has been a blast developing the XENONnT DAQ system. Being awarded "the XENON working group" of the year by the rest of the experiment was an unexpected acknowledgement for a system which people seem to only think about on those very rare occasions that it does not work, and I'm proud of what we have built together. It is now time to let our baby mature, and it will be taken excellent care of by Maricke, Carlo, Robin, Valerio, and Alessandro.

Daniel, the commissioning of the XENONnT detector would not have been the same without you. During the covid pandemic we were pulling the software framework together, understanding our new detector, and designing the future steps along the way, which was so much fun to do in collaboration with you.

Thank you, Gianfranco and Bradley, for writing the complementarity paper together. You really took the time to share ideas and improve every time we had a meeting. Your feedback and help on writing the paper are among the best that I received during my PhD. I'm proud of the result and grateful of how much I learned in the process.

The XENON-family has been a great and motivating atmosphere, which in my opinion really floats on the work of truly inspiring individuals that are not afraid to work around the clock to make the experiment work at its prime. Thank you, Daniel, Diego, Yossi, Lanqing, Joey, Giovanni, Peter, Jelle, Luca Scotto, Evan, Jörn, Micha, Francesco, Riccardo, and all other XENON members for all our constructive, productive, and heated discussions on how to make things run smoothly.

I like to thank Nikhef, for providing such a great environment to do science. I had the pleasure of being on the works council together with Jan, Rasa, Sander, Dimitri, Marieke, Peter, Anton, and Jesse, which really showed me that there are so many people involved in making Nikhef such a great place to work, guided by Stan, Pieter, and Arjen.

Floris, Nino, Jaïr, Alik, Sjaak, Amir, Stan, and Pim, you were always there when it was time for non-physics. You were able to get me away from the keyboard when I was slowly being absorbed by my work. Your quick-wittedness on anything but physics is always refreshing after a period of work. No one can make me laugh as you guys.

Tycho and Fabian, thanks for always being there, I know I can always count on you. I'm lucky to have two great brothers who I can have so much fun with and rely on.

Els and Rick, you truly are the two pillars on which I have always been able to build. Without your infinite and unconditional support I could never have dreamt of pursuing a PhD in the first place, nor write it, and certainly not finish it. Even when the going was tough, you have been, are, and will be at the foundation of what I do.

Femke, you experienced me during all the phases of my PhD, from being called from bed to do midnight debugging, getting papers published, to the tough period of thesis writing, and you never failed to support me lovingly. During those stressful times of finalizing my thesis, we also decided to buy a house, and I don't fully understand how, but you made it work. I cannot put in words how much you meant and mean to me.

Finally, I would like to thank everyone which I had the privilege to work with during my PhD. The great colleagues made these four truly fantastic years!

Nederlandse Samenvatting

van "First WIMP results of XENONnT and its signal reconstruction" (Eerste WIMP-resultaten van XENONnT en signaalreconstructie ervan)

Als natuurkundigen proberen wij het universum volledig te begrijpen, wat een lastige puzzel is met de paar puzzelstukjes die we hebben. Deze stukjes beschrijven namelijk normale, baryonische, materie en geven ons maar 16% van de totale massa van het universum. De overige 84% zou een soort materie moeten zijn die we tot nu toe nog nooit hebben gezien. Deze mysterieuze vorm van materie wordt "donkere materie" genoemd omdat het niet met licht wisselwerkt (het is volledig transparant en zendt geen enkele vorm van licht of warmte uit), waardoor het heel lastig is om donkere materie te detecteren.

In deze thesis laten we de eerste donkere materie resultaten van het XENONnT experiment zien in het laatste hoofdstuk (Hoofdstuk 7). Echter zullen we eerst moeten begrijpen waarom donkere materie bestaat (Hoofdstuk 1) en wat het precies zou kunnen zijn (Hoofdstuk 2). We zullen het XENONnT experiment uitleggen (Hoofdstuk 3), en hoe XENONnT samen met het SuperCDMS experiment donkere materie zou kunnen detecteren (Hoofdstuk 4). We hopen ooit donkere materie te detecteren en daarmee de eigenschappen ervan te kunnen reconstrueren. Dit zou worden gedaan zoals is weergegeven in Fig. S1 waarin we ons op gesimuleerde data baseren. Het vastleggen en digitaliseren van de XENONnT data (Hoofdstuk 5) en reconstructie (Hoofdstuk 6) zijn essentieel voor de donkere materie zoektocht van de XENONnT detector (Hoofdstuk 7). Hoewel XENONnT (nog) geen donkere materie heeft ontdekt, sluit het wel veel parameterruimte van donkere materie uit (Fig. S2). Er volgt nu een korte samenvatting van ieder hoofdstuk.

Hoofdstuk 1: Why does Dark Matter exist? (Waarom bestaat donkere materie?)

We zien het bewijs voor donkere materie, de extra bron van zwaartekracht, op vele afstandsschalen, van sterrenstelsels zoals ons eigen Melkwegstelsel tot de kosmische achtergrondstraling, de oudste signalen die we kunnen zien in het elektromagnetische spectrum. De kosmische achtergrondstraling werd 380.000 jaar na de oerknal uitgezonden op het moment dat het universum doorzichtig werd voor lichtdeeltjes (fotonen). De kleine verschillen in de kosmische achtergrondstraling stellen ons in staat om heel exact de bijdrage van normale materie en donkere materie te modelleren. Deze kleine verschillen worden versterkt door donkere materie omdat samengepakte donkere materie door zwaartekracht meer materie zal aantrekken, terwijl normale materie juist een grote druk zou ervaren als het wordt samengepakt. Te veel donkere materie zou ervoor zorgen dat er veel plekken zouden zijn waar materie samenkomt, terwijl te weinig donkere materie zou leiden tot te weinig van zulke regio's met een hogere dichtheid. We proberen de evolutie van het universum te beschrijven met het Λ_{CDM}-model. Met de laatste data van de PLANCK satelliet, laat Λ_{CDM} ons zien dat we maar 16% van de puzzelstukjes in handen hebben (de bijdrage van normale materie aan de totale hoeveelheid materie in het universum). Donkere materie moet dus niet alleen bestaan, maar het moet $> 5\times$ meer bijdragen aan de massa van het universum dan normale materie.

Hoofdstuk 2: Dark Matter (Donkere materie)

Hoewel de exacte aard van donkere materie een mysterie blijft, is het "Weakly Interacting Massive Particle" (zwak-wisselwerkend zwaar deeltje) ookwel WIMP - model een goed onderbouwde verklaring voor wat donkere materie zou kunnen zijn. WIMP's zijn deeltjes die naast zwaartekracht ook met een andere kracht, vergelijkbaar met de zwakke kernkracht, met normale materie wisselwerken. WIMP's zouden worden geproduceerd via een mechanisme dat "uitvriezen" wordt genoemd en zouden nog steeds aanwezig zijn in ons universum vandaag de dag. Er zijn veel technieken ontwikkeld om WIMP's te ontdekken en tijdprojectiekamers op basis van xenon leiden de zoektocht naar donkere materie voor WIMP-massa's $\gtrsim 5\,\mathrm{GeV}/c^2$. Een van deze experimenten is XENONnT bij het Laboratori Nazionali del Gran Sasso in Italië. XENONnT zoekt voornamelijk naar donkere materie door te zoeken naar een overschot van botsingen met xenonatoomkernen.

Hoofdstuk 3: XENONnT

Als donkere materie bestaat zal het haast nooit met normale materie botsen. Hoe vaak dat toch gebeurt drukken we uit in termen van de werkzame doorsnede (cross section). Het XENONnT experiment moet zo weinig mogelijk achtergronden hebben om een kans te maken om deze nog nooit gedetecteerde soort botsingen te zien. Het XENONnT experiment monitort een massa van 8.6 ton vloeibare xenon. Als een deeltje tegen het xenon botst, zal het een deel van zijn energie overdragen en een lichtflits veroorzaken (door scintillatie van het xenon) en xenonatomen ioniseren (zoals weergegeven in Fig. 3.2). Het scintillatielicht wordt gedetecteerd door middel van fotomultiplicatoren (lichtsensoren, photo-multiplier tubes, PMTs) als een eerste signaal (S1). De elektronen die zijn losgemaakt tijdens de ionisatie worden door een elektrisch veld naar de bovenkant van de detector geleid, waar ze vervolgens door een sterker elektrisch veld uit de vloeibare xenon getrokken worden en worden versneld in een laag xenongas. Deze versnelde elektronen zullen ook scintillatie veroorzaken in het xenongas en daarmee het tweede signaal (S2) veroorzaken. Op basis van de S1 en S2 kunnen de eigenschappen van de botsing worden achterhaald, zo geeft de tijd tussen de S1 en S2 aan hoe diep de botsing gebeurde en geeft het patroon van de S2 op de fotomultiplicatoren de informatie die nodig is om te herleiden waar in het horizontale vlak de botsing plaatsvond. De grootte van de S1 en S2 signalen hangt af van de hoeveelheid energie die er is overgedragen tijdens de botsing en de relatieve grootte van de S1 vergeleken met de S2 stelt ons in staat om onderscheid te maken tussen botsingen met de xenonatoomkern of met de elektronen van de xenonatomen. Deze eigenschappen zijn nodig om onderscheid te kunnen maken tussen een donkerematerie botsing of een achtergrondsignaal. Een van de eigenschappen dat XENONnT een erg goede donkerematerie detector maakt, is dat het zeer weinig achtergrondsignalen heeft, waardoor het in een uitstekende positie is om zeldzame processen te ontdekken.

Hoofdstuk 4: Complementarity of light Dark Matter searches (Complementariteit van zoektochten naar lichte donkere materie)

In Hoofdstuk 4 hebben wij Ref. [1] opgenomen dat zich richt op het vinden van "lichte" donkere materie. Als donkere materie bestaat uit WIMP's met een massa van $\mathcal{O}\left(0.1 - 10\,\mathrm{GeV}/c^2\right)$ kunnen verschillende experimenten en detectietechnieken het ontdekken. De energieoverdracht van een botsing van een WIMP en een xenonatoomkern kan te klein zijn om te detecteren in XENONnT.

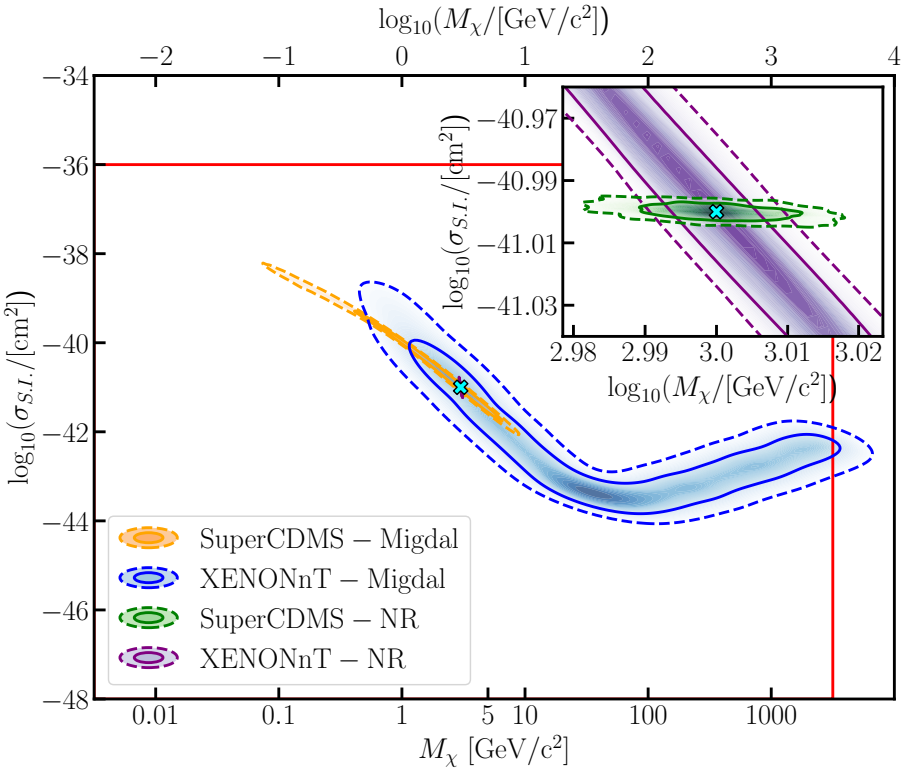

Figuur S1: De reconstructie van gesimuleerde WIMP's met een massa $M_\chi = 3$ GeV/c^2 en werkzame doorsnede $\sigma_{\mathrm{S.I.}} = 10^{-41}$ cm^2 (cyane kruis) met SuperCDMS en XENONnT die ieder een atoomkernbotsing- (nuclear recoil, NR) of een Migdalanalyse toepassen. De gekleurde gebieden geven aan waar de analyses de WIMP-eigenschappen zouden reconstrueren, hoe kleiner het gebied, hoe beter de reconstructie is, waarbij de heldere kleuren hogere waarschijnlijkheden aangeven. De SuperCDMS en XENONnT NR-analyses kunnen beiden de WIMP-eigenschappen goed reconstrueren (zie inzet in de rechterbovenhoek) waarbij SuperCDMS de hoogste nauwkeurigheid heeft. Wanneer de analyses van SuperCDMS en XENONnT worden gecombineerd kunnen de eigenschappen nog preciezer gemeten worden. De Migdalanalyses hebben een lagere nauwkeurigheid voor deze WIMP-massa, maar voor lagere massa's wordt hun precisie relatief beter (Fig. 4.7). Dit figuur komt overeen met Fig. 4.3 op pagina 63.

In plaats daarvan zou deze kunnen worden gedetecteerd als energieverlies van de atomaire elektronen door het "Migdal"-effect, waarbij de xenonatoomkern verplaatst wordt ten opzichte van de atomaire elektronen door een overdracht van energie op de atoomkern. Deze vorm van energieverlies van atomaire elektronen zou zeldzaam zijn maar detecteerbare energieoverdrachten veroorzaken, die een hoger aantal scintilatiefotonen veroorzaken vergeleken met dezelfde energieoverdracht van directe atoomkernbotsingen. Verder heeft een experiment zoals de Super Cryogenic Dark Matter Search (SuperCDMS) bij SNOLAB in Canada, dat germanium en silicium kristallen bij cryogene temperaturen monitort, een erg lage energiedrempelwaarde en een zeer goede energieresolutie, waardoor het ideaal is voor het detecteren van atoomkernbotsingen van dit type lichte donkere materie. Bovendien kan het Migdaleffect ook de donkeremateriezoektocht van SuperCDMS naar zeer lage WIMP-massa's uitbreiden. Het combineren van de data van XENONnT en SuperCMDS met een combinatie van atoomkernbotsing- en Migdalanalyses leidt tot een nauwkeurigere reconstructie van de WIMP-eigenschappen als deze gedetecteerd worden. Deze experimenten hopen ooit donkere materie te meten waardoor ze de eigenschappen ervan kunnen reconstrueren zoals de massa (M_χ) en werkzame doorsnede ($\sigma_{\text{S.I.}}$) zoals in Fig. S1. In Hoofdstuk 4, berekenen we de complementariteit van deze twee experimenten en zoektochten. We laten zien dat XENONnT WIMP's het beste reconstrueert voor massa's $> 5.6\,\text{GeV}/c^2$ en baat zou hebben van het combineren van de resultaten met SuperCDMS voor lagere massa's om daarmee een preciezere reconstructie van de WIMP-eigenschappen te bereiken. De grootste complementariteit is gevonden tussen de atoomkernbotsing- en Migdalanalyse van SuperCDMS (en in mindere mate de XENONnT Migdalanalyses) voor massa's tussen $[0.2, 0.6]\,\text{GeV}/c^2$.

Hoofdstuk 5: The XENONnT DAQ (De XENONnT DAQ)

Hoofdstuk 5 bevat Ref. [3] waarin we het ontwerp, de ingebruikname, en de prestaties van het XENONnT data-acquisitie systeem (DAQ) beschrijven. De DAQ is verantwoordelijk voor het digitaliseren van de analoge signalen van de fotomultiplicatoren. De DAQ is ontworpen zonder het gebruik van een zogenoemde "trigger" op de signalen van de tijdprojectiekamer, en alle signalen van individuele fotomultiplicatoren worden gedigitaliseerd als ze boven een kleine drempelwaarde uitkomen. De DAQ kan omgaan met de grote verscheidenheid aan signalen, van korte S1s ($\sim$100 ns) tot lange S2 signalen van tot $\sim$100 µs. Bovendien is de DAQ ontworpen om de drie subdetectoren: de tijdprojectiekamer, de muonenveto en de nieuwe neutronenveto gezamelijk of onafhankelijk

aan te sturen. Als de subdetectoren als een gezamenlijke eenheid worden aangestuurd worden ze "geketend" waarbij de kloksignalen worden verdeeld over alle betrokken hardware zodat de tijden en klokken gesynchroniseerd zijn tussen de subdetectoren.

De DAQ doet ook de dataprocessing van de data van de fotomultiplicatoren. De dataprocessing is dusdanig snel en zo opgezet dat de data volledig gereconstrueerd is in $\mathcal{O}(10\,\mathrm{s})$ tot data waarbij de S1 en S2 signalen van een botsing met elkaar gekoppeld zijn. Deze data wordt vervolgens opgeslagen in een online database zodat deze live kan worden bekeken en de detector live kan worden gemonitord.

Hoofdstuk 6: Signal reconstruction (Signaalreconstructie)

De reconstructie van botsingen die wordt gedaan door de fotomultiplicatoren, de DAQ en de processingsoftware is gevalideerd in Hoofdstuk 6 met gesimuleerde data. Twee sleutelconcepten zijn direct verbonden met de reconstructie; de energiedrempelwaarde en de energiereconstructie. De eerste is afgeleid van de S1 detectie-efficiëntie wat de waarschijnlijkheid geeft dat een S1 wordt gereconstrueerd als functie van de signaalgrootte. De energiereconstructie wordt beïnvloed door de reconstructie-onzuiverheid (bias), oftewel het verschil tussen de werkelijke signaalgrootte en de gereconstrueerde signaalgrootte. De gereconstrueerde signaalgrootte is gemiddeld -2% tot 1% van de werkelijke signaalgrootte afhankelijk van de grootte en het type (S1 of S2) van het signaal. De voornaamste oorzaken voor dit effect zijn de drempelwaarde van de digitalisatie en het meenemen van napulsen van de fotomultiplicatoren. De ogenschijnlijke niet-lineaire energierespons die hierdoor wordt veroorzaakt corrigeren we achteraf.

Hoofdstuk 7: First WIMP results with XENONnT (De eerste WIMP-resultaten met XENONnT)

In het laatste hoofdstuk (Hoofdstuk 7) bespreken we de WIMP resultaten van XENONnT. Er is geen overschot aan interacties boven de verwachte achtergrond waargenomen en nieuwe uitsluitingslimieten zijn gezet op de WIMP-eigenschappen zoals in Fig. S2. XENONnT sluit de spinonafhankelijke WIMP-nucleon werkzame doorsnede uit tot $\sigma_{n,\,\mathrm{S.I.}} = 2.2 \times 10^{-47}\,\mathrm{cm}^2$ bij een WIMP-massa van $26\,\mathrm{GeV}/c^2$.

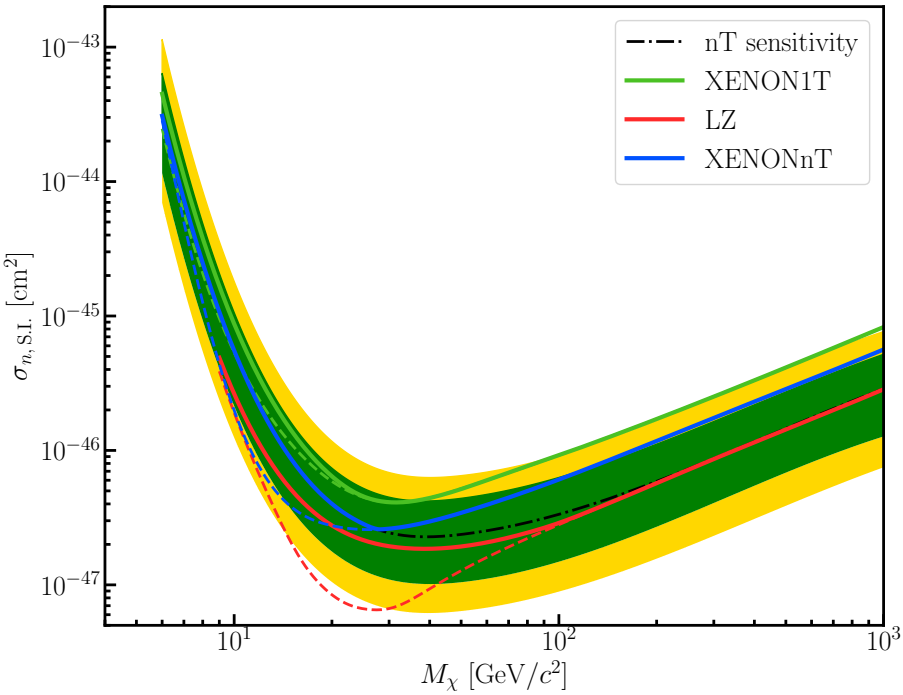

Figuur S2: De uitsluitingslimieten van XENONnT voor WIMPs (blauwe lijn) met werkzame doorsnede $\sigma_{n,\mathrm{S.I.}}$ als functie van massa (M_χ), en de sensitiviteit (zwarte streepjes-stippellijn) met 1σ en 2σ onzekerheidsbanden in geel en groen respectievelijk. Alles boven de uitsluitingslimiet is uitgesloten (bij 90% betrouwbaarheidsniveau), terwijl WIMP's nog zouden kunnen bestaan als de massa en werkzame doorsnede onder de limiet ligt. De gestreepte lijnen sluiten WIMP's uit tot de -2σ onzekerheidsband terwijl de ononderbroken lijnen WIMP's uitsluiten tot de mediane sensitiviteit. De limiet van XENON1T (XENONnT's voorganger) is weergegeven met de groene lijn [4], en die van LZ (een experiment vergelijkbaar met XENONnT) met de rode lijn [5]. De WIMP-zoektocht vindt plaats bij lagere en lagere werkzame doorsnedes. Dit figuur komt overeen met Fig. 7.6 op pagina 147 en is aangepast van Ref. [6].

Glossary

AC Accidental Coincidence. 133

AFT Area Fraction Top. 114, 115, 118–120

ALP axion-like particle. 19

CEνNS coherent elastic neutrino nucleus scattering. 39, 119, 132, 144

CMB Cosmic Microwave Background. 1, 2, 8, 12–15

CP Charge Parity. 19

DAQ Data Acquisition system. 4, 5, 36, 75, 109, 110, 122, 123, 136, 139, 140

DM Dark Matter. 1–3, 5, 7–10, 12–28, 30, 32–34, 41, 45, 47, 118, 139

DPE Double Photoelectron Emission. 110, 120–122

ER electronic recoil. 3, 24, 35, 37–39, 44, 45, 118, 132, 138–145, 147, 148

FV Fiducial Volume. 119, 120

GXe gaseous xenon. 2, 40, 44, 114, 125, 126

LHC Large Hadron Collider. 20, 22

LNGS Laboratori Nazionali del Gran Sasso. 2, 41, 154

LXe liquid xenon. 2, 39, 40, 43, 44, 112, 120, 125, 126, 141, 143

MACHO Massive Compact Halo Object. 19

MOND Modified Newtonian dynamics. 20

MV Muon Veto. 5, 42, 43, 139

NR nuclear recoil. 3, 4, 24, 25, 35, 37–39, 44, 47, 139, 142–144, 148, 156

NV Neutron Veto. 5, 42, 43, 139

PMT photo-multiplier tube. 2, 3, 5, 36, 40–43, 45, 75, 109–115, 118–121, 123, 126, 141, 155

PTFE Polytetrafluoroethylene. 41, 42, 44, 115, 141

ROI Region of interest. 31, 142

SE Single Electron. 113, 119, 120, 126–129

SHM Standard Halo Model. 26–29, 31, 33

SM Standard Model of particle physics. 7, 17, 18, 20–22, 24, 26, 44, 142

SPE Single Photoelectron. 109–113, 120, 122–126

SR0 Science Run 0. 16, 75, 110, 127, 134, 139–141

TC Tight Coincidence. 115, 117–119, 125, 132, 133, 135

TPC Time Projection Chamber. iv, 2, 3, 5, 24, 35, 39–45, 75, 112, 115, 116, 120, 127, 128, 141, 143, 144

WIMP Weakly Interacting Massive Particle. iii, 2–6, 17–21, 24–35, 38, 39, 41, 44, 45, 47, 132, 138, 139, 141, 143–147, 154, 158, 159

Bibliography

[1] J.R. Angevaare, G. Bertone, A.P. Colijn, M.P. Decowski and B.J. Kavanagh, Complementarity of direct detection experiments in search of light dark matter, JCAP **10** (2022) 004 [2204.01580].

[2] J.R. Angevaare, Joranangevaare/dddm: v4.0.0, July, 2022. 10.5281/zenodo.6922328.

[3] XENON collaboration, The triggerless data acquisition system of the xenonnt experiment, arXiv pre-print (2022) [2212.11032].

[4] XENON collaboration, Dark matter search results from a one ton-year exposure of xenon1t, Phys. Rev. Lett. **121** (2018) 111302 [1805.12562].

[5] LZ collaboration, First dark matter search results from the lux-zeplin (lz) experiment, arXiv pre-print (2022) [2207.03764].

[6] XENON collaboration, First dark matter search with nuclear recoils from the xenonnt experiment, arXiv pre-print (2023) [2303.14729].

[7] J. Kapteyn, First attempt at a theory of the arrangement and motion of the sidereal system, The Astrophysical Journal **55** (1922) 302.

[8] G. Bertone, J. Silk, B. Moore, J. Diemand, J. Bullock, M. Kaplinghat et al., Particle dark matter: observations, models and searches, Cambridge University Press (2010).

[9] G. D'Amico, M. Kamionkowski and K. Sigurdson, Dark matter astrophysics, in Dark Matter and Dark Energy: A Challenge for Modern Cosmology, S. Matarrese, M. Colpi, V. Gorini and U. Moschella, eds., (Dordrecht), pp. 241–272, Springer Netherlands (2011), DOI.

[10] M. Roos, Dark matter: The evidence from astronomy, astrophysics and cosmology, 2010. 10.48550/ARXIV.1001.0316.

[11] M. Lisanti, Lectures on dark matter physics, in New Frontiers in Fields and Strings, WORLD SCIENTIFIC, nov, 2016, DOI.

[12] D. Clowe, M. Bradač, A.H. Gonzalez, M. Markevitch, S.W. Randall, C. Jones et al., A direct empirical proof of the existence of dark matter, The Astrophysical Journal **648** (2006) L109.

[13] A.L. Serra and M.J.L.D. Romero, Measuring the dark matter equation of state, Monthly Notices of the Royal Astronomical Society: Letters **415** (2011) L74.

[14] V. Springel, C.S. Frenk and S.D.M. White, The large-scale structure of the universe, Nature **440** (2006) 1137.

[15] C. Knobel, An introduction into the theory of cosmological structure formation, arXiv pre-print (2012) [1208.5931].

[16] A.L. Coil, The large-scale structure of the universe, in Planets, Stars and Stellar Systems, pp. 387–421, Springer Netherlands (2013), DOI.

[17] E.E. Richards, L. van Zee, K.L. Barnes, S. Staudaher, D.A. Dale, T.T. Braun et al., Baryonic distributions in the dark matter halo of NGC 5005, Monthly Notices of the Royal Astronomical Society **449** (2015) 3981.

[18] J.R. Angevaare, "Joranangevaare/thesis_plots: v1.1.1." https://doi.org/10.5281/zenodo.7568784, Jan., 2023. 10.5281/zenodo.7568784.

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

[20] F. Zwicky, Die rotverschiebung von extragalaktischen nebeln, Helvetica Physica Acta **6** (1933) 110.

[21] E. Corbelli and P. Salucci, The extended rotation curve and the dark matter halo of m33, Monthly Notices of the Royal Astronomical Society **311** (2000) 441.

[22] J.F. Navarro, C.S. Frenk and S.D.M. White, A universal density profile from hierarchical clustering, **490** (1997) 493 [astro-ph/9611107].

[23] J.F. Navarro, C.S. Frenk and S.D.M. White, The structure of cold dark matter halos, **462** (1996) 563 [astro-ph/9508025].

[24] A.-C. Eilers, D.W. Hogg, H.-W. Rix and M.K. Ness, The circular velocity curve of the milky way from 5 to 25 kpc, The Astrophysical Journal **871** (2019) 120.

[25] M. Cautun, A. Benítez-Llambay, A.J. Deason, C.S. Frenk, A. Fattahi, F.A. Gómez et al., The milky way total mass profile as inferred from gaia dr2, **494** (2020) 4291 [1911.04557].

[26] H. Leung. https://milkyway-plot.readthedocs.io/en/latest/, 2022.

[27] M. Cautun and T.M. Callingham, Milky way mass profile, Apr., 2020. 10.5281/zenodo.3740067.

[28] Gaia Collaboration, T. Prusti, J.H.J. de Bruijne, A.G.A. Brown, A. Vallenari, C. Babusiaux et al., The gaia mission, **595** (2016) A1 [1609.04153].

[29] Gaia Collaboration, A.G.A. Brown, A. Vallenari, T. Prusti, J.H.J. de Bruijne, C. Babusiaux et al., Gaia data release 2. summary of the contents and survey properties, **616** (2018) A1 [1804.09365].

[30] N.W. Evans, C.A. O'Hare and C. McCabe, Refinement of the standard halo model for dark matter searches in light of the gaia sausage, Physical Review D **99** (2019) 023012.

[31] D. Baxter et al., Recommended conventions for reporting results from direct dark matter searches, Eur. Phys. J. C **81** (2021) 907 [2105.00599].

[32] P.F. de Salas and A. Widmark, Dark matter local density determination: recent observations and future prospects, Reports on Progress in Physics **84** (2021) 104901.

[33] J. de Swart, G. Bertone and J. van Dongen, How dark matter came to matter, Nature Astron. **1** (2017) 0059 [1703.00013].

[34] S. Dodelson, Modern Cosmology, Elsevier (2003).

[35] Planck Collaboration, Ade, P. A. R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J. et al., Planck 2015 results - xiii. cosmological parameters, A&A **594** (2016) A13.

[36] S.D. McDermott, H.-B. Yu and K.M. Zurek, Turning off the lights: How dark is dark matter?, Phys. Rev. D **83** (2011) 063509.

[37] A. Caputo, L. Sberna, M. Frías, D. Blas, P. Pani, L. Shao et al., Constraints on millicharged dark matter and axionlike particles from timing of radio waves, Phys. Rev. D **100** (2019) 063515.

[38] G. Steigman, B. Dasgupta and J.F. Beacom, Precise relic wimp abundance and its impact on searches for dark matter annihilation, Phys. Rev. D **86** (2012) 023506.

[39] E.W. Kolb and K.A. Olive, Lee-weinberg bound reexamined, Phys. Rev. D **33** (1986) 1202.

[40] K. Griest and M. Kamionkowski, Unitarity limits on the mass and radius of dark-matter particles, Phys. Rev. Lett. **64** (1990) 615.

[41] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre et al., The waning of the wimp? a review of models, searches, and constraints, Eur. Phys. J. C **78** (2018) 203 [1703.07364].

[42] J. Aalbers et al., A next-generation liquid xenon observatory for dark matter and neutrino physics, J. Phys. G **50** (2023) 013001 [2203.02309].

[43] DarkSide collaboration, Low-mass dark matter search with the darkside-50 experiment, Phys. Rev. Lett. **121** (2018) 081307.

[44] E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, L. Althueser, F. Amaro et al., Projected wimp sensitivity of the xenonnt dark matter experiment, Journal of Cosmology and Astroparticle Physics **2020** (2020) 031.

[45] R. Agnese, A. Anderson, T. Aramaki, I. Arnquist, W. Baker, D. Barker et al., Projected sensitivity of the supercdms snolab experiment, Physical Review D **95** (2017) 082002.

[46] T. Emken, R. Essig, C. Kouvaris and M. Sholapurkar, Direct detection of strongly interacting sub-GeV dark matter via electron recoils, Journal of Cosmology and Astroparticle Physics **2019** (2019) 070.

[47] CDEX collaboration, Studies of the earth shielding effect to direct dark matter searches at the china jinping underground laboratory, Phys. Rev. D **105** (2022) 052005 [2111.11243].

[48] R.D. Peccei and H.R. Quinn, CP conservation in the presence of pseudoparticles, Phys. Rev. Lett. **38** (1977) 1440.

[49] D.J.E. Marsh, Axion cosmology, Phys. Rept. **643** (2016) 1 [1510.07633].

[50] Particle Data Group collaboration, Review of particle physics, Phys. Rev. D **98** (2018) 030001.

[51] H. Baer, K.-Y. Choi, J.E. Kim and L. Roszkowski, Dark matter production in the early universe: Beyond the thermal wimp paradigm, Physics Reports **555** (2015) 1.

[52] A. Drlica-Wagner, Y.-Y. Mao, S. Adhikari, R. Armstrong, A. Banerjee, N. Banik et al., Probing the fundamental nature of dark matter with the large synoptic survey telescope, arXiv pre-print (2019) arXiv:1902.01055 [1902.01055].

[53] K. GRIEST, The search for the dark matter: WIMPs and MACHOs, Annals of the New York Academy of Sciences **688** (1993) 390.

[54] G. Bertone and D. Hooper, History of dark matter, Rev. Mod. Phys. **90** (2018) 045002.

[55] K. Griest, A.M. Cieplak and M.J. Lehner, New limits on primordial black hole dark matter from an analysis of kepler source microlensing data, Phys. Rev. Lett. **111** (2013) 181302.

[56] P. Tisserand, L. Le Guillou, C. Afonso, J.N. Albert, J. Andersen, R. Ansari et al., Limits on the macho content of the galactic halo from the eros-2 survey of the magellanic clouds, **469** (2007) 387 [astro-ph/0607207].

[57] C. Alcock, R.A. Allsman, D. Alves, R. Ansari, É. Aubourg, T.S. Axelrod et al., EROS and MACHO combined limits on planetary-mass dark matter in the galactic halo, The Astrophysical Journal **499** (1998) L9.

[58] M. Sasaki, T. Suyama, T. Tanaka and S. Yokoyama, Primordial black holes—perspectives in gravitational wave astronomy, Classical and Quantum Gravity **35** (2018) 063001.

[59] G. Bertone, A.M. Coogan, D. Gaggero, B.J. Kavanagh and C. Weniger, Primordial black holes as silver bullets for new physics at the weak scale, Phys. Rev. D **100** (2019) 123013.

[60] M. Milgrom, A modification of the newtonian dynamics as a possible alternative to the hidden mass hypothesis., **270** (1983) 365.

[61] M. Milgrom, A modification of the newtonian dynamics - implications for galaxies., **270** (1983) 371.

[62] J.D. Bekenstein, Relativistic gravitation theory for the modified newtonian dynamics paradigm, Phys. Rev. D **70** (2004) 083509.

[63] A. Elykov, "Background modeling & data acquisition for xenon detectors." https://doi.org/10.6094/UNIFR/225492, 2021. 10.6094/UNIFR/225492.

[64] O. Buchmueller, C. Doglioni and L.T. Wang, Search for dark matter at colliders, Nature Phys. **13** (2017) 217 [1912.12739].

[65] D. Hooper and L. Goodenough, Dark matter annihilation in the galactic center as seen by the fermi gamma ray space telescope, Physics Letters B **697** (2011) 412.

[66] O. Macias, C. Gordon, R.M. Crocker, B. Coleman, D. Paterson, S. Horiuchi et al., Galactic bulge preferred over dark matter for the galactic centre gamma-ray excess, Nature Astron. **2** (2018) 387 [1611.06644].

[67] J.M. Gaskins, A review of indirect searches for particle dark matter, Contemporary Physics **57** (2016) 496 [https://doi.org/10.1080/00107514.2016.1175160].

[68] ATLAS collaboration, The ATLAS experiment at the CERN large hadron collider, Journal of Instrumentation **3** (2008) S08003.

[69] CMS collaboration, The CMS experiment at the CERN LHC, Journal of Instrumentation **3** (2008) S08004.

[70] F. Kahlhoefer, Review of lhc dark matter searches, Int. J. Mod. Phys. A **32** (2017) 1730006 [1702.02430].

[71] PandaX-4T collaboration, Dark matter search results from the pandax-4t commissioning run, Phys. Rev. Lett. **127** (2021) 261802 [2107.13438].

[72] EDELWEISS collaboration, Searching for low-mass dark matter particles with a massive ge bolometer operated above ground, Phys. Rev. D **99** (2019) 082003.

[73] Y. Hochberg, Y. Kahn, N. Kurinsky, B.V. Lehmann, T.C. Yu and K.K. Berggren, Determining dark-matter–electron scattering rates from the dielectric function, Phys. Rev. Lett. **127** (2021) 151802 [2101.08263].

[74] Y. Hochberg, B.V. Lehmann, I. Charaev, J. Chiles, M. Colangelo, S.W. Nam et al., New constraints on dark matter from superconducting nanowires, arXiv pre-print (2021) [2110.01586].

[75] Y. Hochberg, E.D. Kramer, N. Kurinsky and B.V. Lehmann, Directional detection of light dark matter in superconductors, arXiv pre-print (2021) [2109.04473].

[76] F. Mayet et al., A review of the discovery reach of directional dark matter detection, Phys. Rept. **627** (2016) 1 [1602.03781].

[77] DARWIN collaboration, Darwin: towards the ultimate dark matter detector, Journal of Cosmology and Astroparticle Physics **2016** (2016) 017.

[78] DEAP-3600 collaboration, First results from the deap-3600 dark matter search with argon at snolab, Phys. Rev. Lett. **121** (2018) 071801.

[79] XENON collaboration, Constraining the spin-dependent wimp-nucleon cross sections with xenon1t, Phys. Rev. Lett. **122** (2019) 141301.

[80] A.M. Green, Astrophysical uncertainties on direct detection experiments, Mod. Phys. Lett. A **27** (2012) 1230004 [1112.0524].

[81] B.J. Mount et al., Lux-zeplin (lz) technical design report, arXiv pre-print (2017) [1703.09144].

[82] L.E. Strigari, Galactic searches for dark matter, Phys. Rept. **531** (2013) 1 [1211.7090].

[83] A.M. Green, Astrophysical uncertainties on the local dark matter distribution and direct detection experiments, J. Phys. G **44** (2017) 084001 [1703.10102].

[84] S.E. Koposov, H.-W. Rix and D.W. Hogg, Constraining the milky way potential with a 6-d phase-space map of the gd-1 stellar stream, Astrophys. J. **712** (2010) 260 [0907.1085].

[85] J. Bland-Hawthorn and O. Gerhard, The galaxy in context: Structural, kinematic, and integrated properties, Annual Review of Astronomy and Astrophysics **54** (2016) 529 [https://doi.org/10.1146/annurev-astro-081915-023441].

[86] GRAVITY Collaboration, Abuter, R., Amorim, A., Bauböck, M., Berger, J. P., Bonnet, H. et al., Improved gravity astrometric accuracy from modeling optical aberrations, A&A **647** (2021) A59.

[87] M.C. Smith et al., The rave survey: Constraining the local galactic escape speed, Mon. Not. Roy. Astron. Soc. **379** (2007) 755 [astro-ph/0611671].

[88] A.J. Deason, A. Fattahi, V. Belokurov, N.W. Evans, R.J.J. Grand, F. Marinacci et al., The local high-velocity tail and the galactic escape speed, Monthly Notices of the Royal Astronomical Society **485** (2019) 3514.

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

[90] C. McCabe, The earth's velocity for direct detection experiments, JCAP **02** (2014) 027 [1312.1355].

[91] R. Schoenrich, J. Binney and W. Dehnen, Local kinematics and the local standard of rest, Mon. Not. Roy. Astron. Soc. **403** (2010) 1829 [0912.3693].

[92] E. Hogenbirk, J. Aalbers, P. Breur, M. Decowski, K. van Teutem and A. Colijn, Precision measurements of the scintillation pulse shape for low-energy recoils in liquid xenon, Journal of Instrumentation **13** (2018) P05016.

[93] V. Chepel and H. Araújo, Liquid noble gas detectors for low energy particle physics, Journal of Instrumentation **8** (2013) R04001.

[94] V. Antochi, L. Baudis, J. Bollig, A. Brown, R. Budnik, D. Cichon et al., Improved quality tests of r11410-21 photomultiplier tubes for the xenonnt experiment, Journal of Instrumentation **16** (2021) P08033.

[95] M. Szydagis, J. Balajthy, G. Block, J. Brodsky, J. Cutter, S. Farrell et al., "Noble element simulation technique." https://doi.org/10.5281/zenodo.7061832, Sept., 2022. 10.5281/zenodo.7061832.

[96] S. Farrell, C. Tunnell, J.R. Angevaare, G.R. Rischbieter, N. Carrara and mszydagis, "Nestcollaboration/nestpy: Update to nestpy internals, adding in basic larnest bindings.." https://doi.org/10.5281/zenodo.7062199, Sept., 2022. 10.5281/zenodo.7062199.

[97] XENON collaboration, Search for new physics in electronic recoil data from xenonnt, Phys. Rev. Lett. **129** (2022) 161805.

[98] D. Akimov, J.B. Albert, P. An, C. Awe, P.S. Barbeau, B. Becker et al., Observation of coherent elastic neutrino-nucleus scattering, Science **357** (2017) 1123 [https://www.science.org/doi/pdf/10.1126/science.aao0990].

[99] COHERENT collaboration, Measurement of the coherent elastic neutrino-nucleus scattering cross section on csi by coherent, Phys. Rev. Lett. **129** (2022) 081801.

[100] J. Colaresi, J.I. Collar, T.W. Hossbach, C.M. Lewis and K.M. Yocum, Measurement of coherent elastic neutrino-nucleus scattering from reactor antineutrinos, Phys. Rev. Lett. **129** (2022) 211802 [2202.09672].

[101] D.S. Akerib et al., Snowmass2021 cosmic frontier dark matter direct detection to the neutrino fog, in 2022 Snowmass Summer Study, 3, 2022 [2203.08084].

[102] XENON collaboration, Material radiopurity control in the xenonnt experiment, Eur. Phys. J. C **82** (2022) 599 [2112.05629].

[103] XENON collaboration, Conceptual design and simulation of a water cherenkov muon veto for the xenon1t experiment, Journal of Instrumentation **9** (2014) P11006.

[104] XENON collaboration, Xenon1t dark matter data analysis: Signal reconstruction, calibration, and event selection, Phys. Rev. D **100** (2019) 052014.

[105] G. Plante, E. Aprile, J. Howlett and Y. Zhang, Liquid-phase purification for multi-tonne xenon detectors, Eur. Phys. J. C **82** (2022) 860 [2205.07336].

[106] XENON collaboration, Excess electronic recoil events in xenon1t, Phys. Rev. D **102** (2020) 072004 [2006.09721].

[107] XENON collaboration, Light dark matter search with ionization signals in xenon1t, Phys. Rev. Lett. **123** (2019) 251801 [1907.11485].

[108] A. Drukier and L. Stodolsky, Principles and applications of a neutral current detector for neutrino physics and astronomy, Phys. Rev. D **30** (1984) 2295.

[109] M.W. Goodman and E. Witten, Detectability of certain dark matter candidates, Phys. Rev. D **31** (1985) 3059.

[110] A.K. Drukier, K. Freese and D.N. Spergel, Detecting cold dark matter candidates, Phys. Rev. D **33** (1986) 3495.

[111] J. Billard, M. Boulay, S. Cebrián, L. Covi, G. Fiorillo, A. Green et al., Direct detection of dark matter – APPEC committee report, arXiv (2021) [2104.07634].

[112] J.D. Vergados and H. Ejiri, The role of ionization electrons in direct neutralino detection, Phys. Lett. B **606** (2005) 313 [hep-ph/0401151].

[113] C.C. Moustakidis, J.D. Vergados and H. Ejiri, Direct dark matter detection by observing electrons produced in neutralino-nucleus collisions, Nucl. Phys. B **727** (2005) 406 [hep-ph/0507123].

[114] R. Bernabei et al., On electromagnetic contributions in wimp quests, Int. J. Mod. Phys. A **22** (2007) 3155 [0706.1421].

[115] M. Ibe, W. Nakano, Y. Shoji and K. Suzuki, Migdal effect in dark matter direct detection experiments, Journal of High Energy Physics **2018** (2018) 1.

[116] M.J. Dolan, F. Kahlhoefer and C. McCabe, Directly detecting sub-gev dark matter with electrons from nuclear scattering, Phys. Rev. Lett. **121** (2018) 101801 [1711.09906].

[117] S. Knapen, J. Kozaczuk and T. Lin, Migdal effect in semiconductors, Phys. Rev. Lett. **127** (2021) 081805 [2011.09496].

[118] XENON collaboration, Search for light dark matter interactions enhanced by the migdal effect or bremsstrahlung in xenon1t, Phys. Rev. Lett. **123** (2019) 241803 [1907.12771].

[119] CDEX collaboration, Constraints on spin-independent nucleus scattering with sub-gev weakly interacting massive particle dark matter from the cdex-1b experiment at the china jinping underground laboratory, Phys. Rev. Lett. **123** (2019) 161301 [1905.00354].

[120] SuperCDMS collaboration, A search for low-mass dark matter via bremsstrahlung radiation and the migdal effect in supercdms, arXiv pre-print (2022) [2203.02594].

[121] EDELWEISS collaboration, Search for sub-gev dark matter via migdal effect with an edelweiss germanium detector with nbsi tes sensors, arXiv pre-print (2022) arXiv:2203.03993 [2203.03993].

[122] D. Akerib, S. Alsum, H. Araújo, X. Bai, J. Balajthy, P. Beltrame et al., Results of a search for sub-GeV dark matter using 2013 LUX data, Physical Review Letters **122** (2019) .

[123] M. Pato, L. Baudis, G. Bertone, R. Ruiz de Austri, L.E. Strigari and R. Trotta, Complementarity of dark matter direct detection targets, Phys. Rev. D **83** (2011) 083505 [1012.3458].

[124] A.H.G. Peter, V. Gluscevic, A.M. Green, B.J. Kavanagh and S.K. Lee, Wimp physics with ensembles of direct-detection experiments, Phys. Dark Univ. **5-6** (2014) 45 [1310.7039].

[125] P.J. Fox, J. Liu and N. Weiner, Integrating out astrophysical uncertainties, Phys. Rev. D **83** (2011) 103514 [1011.1915].

[126] B.J. Kavanagh and A.M. Green, Model independent determination of the dark matter mass from direct detection experiments, Phys. Rev. Lett. **111** (2013) 031302 [1303.6868].

[127] C. McCabe, The earth's velocity for direct detection experiments, Journal of Cosmology and Astroparticle Physics **2014** (2014) 027 [1312.1355].

[128] SuperCDMS collaboration, New results from the search for low-mass weakly interacting massive particles with the cdms low ionization threshold experiment, Phys. Rev. Lett. **116** (2016) 071301.

[129] J. Lindhard, V. Nielsen, M. Scharff and P. Thomsen, Integral equations governing radiation effects, Mat. Fys. Medd. Dan. Vid. Selsk **33** (1963) 1.

[130] E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, L. Althueser, F. Amaro et al., Light dark matter search with ionization signals in xenon1t, Physical Review Letters **123** (2019) 251801.

[131] LUX-Collaboration, D.S. Akerib, S. Alsum, H.M. Araújo, X. Bai, A.J. Bailey et al., Low-energy (0.7-74 kev) nuclear recoil calibration of the lux dark matter experiment using d-d neutron scattering kinematics, arXiv pre-print (2016) [1608.05381].

[132] XENON collaboration, Search for coherent elastic scattering of solar ^{8}B neutrinos in the xenon1t dark matter experiment, Phys. Rev. Lett. **126** (2021) 091301.

[133] J. Aalbers, B. Pelssers, J.R. Angevaare and K.D. Morå, "Jelleaalbers/wimprates: v0.4.1." https://doi.org/10.5281/zenodo.7041453, Sept., 2022. 10.5281/zenodo.7041453.

[134] S. Knapen, J. Kozaczuk and T. Lin, python package for dark matter scattering in dielectric targets, Physical Review D **105** (2022) .

[135] T. Bayes, Rev., An essay toward solving a problem in the doctrine of chances, Phil. Trans. Roy. Soc. Lond. **53** (1764) 370.

[136] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, emcee: the mcmc hammer, Publications of the Astronomical Society of the Pacific **125** (2013) 306.

[137] K. Barbary, kbarbary / nestle : v0.2.0, Nov., 2016.

[138] J. Buchner, A. Georgakakis, K. Nandra, L. Hsu, C. Rangel, M. Brightman et al., X-ray spectral modelling of the agn obscuring region in the cdfs: Bayesian model selection and catalogue, Astronomy & Astrophysics **564** (2014) A125.

[139] F. Feroz, M.P. Hobson, E. Cameron and A.N. Pettitt, Importance nested sampling and the multinest algorithm, Open J. Astrophys. **2** (2019) 10 [1306.2144].

[140] F. Feroz, M.P. Hobson and M. Bridges, Multinest: an efficient and robust bayesian inference tool for cosmology and particle physics, Mon. Not. Roy. Astron. Soc. **398** (2009) 1601 [0809.3437].

[141] M.L. Waskom, seaborn: statistical data visualization, Journal of Open Source Software **6** (2021) 3021.

[142] B.J. Kavanagh, Earth scattering of superheavy dark matter: Updated constraints from detectors old and new, Phys. Rev. D **97** (2018) 123013 [1712.04901].

[143] Y. Sarkis, A. Aguilar-Arevalo and J.C. D'Olivo, Study of the ionization efficiency for nuclear recoils in pure crystals, Physical Review D **101** (2020) .

[144] XENON collaboration, The XENON1T dark matter experiment, Eur. Phys. J. C **77** (2017) 881.

[145] LZ collaboration, The lux-zeplin (lz) experiment, Nucl. Instrum. Meth. A **953** (2020) 163047 [1910.09124].

[146] nEXO collaboration, nexo: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity, J. Phys. G **49** (2022) 015104 [2106.16243].

[147] DarkSide collaboration, Darkside-50 532-day dark matter search with low-radioactivity argon, Phys. Rev. D **98** (2018) 102006 [1802.07198].

[148] XENON collaboration, Projected wimp sensitivity of the xenonnt dark matter experiment, JCAP **11** (2020) 031 [2007.08796].

[149] XENON collaboration, Search for new physics in electronic recoil data from xenonnt, Phys. Rev. Lett. **129** (2022) 161805.

[150] "Caen." https://www.caen.it/.

[151] XENON collaboration, The XENON1T data acquisition system, JINST **14** (2019) P07016.

[152] XENON collaboration, Conceptual design and simulation of a water cherenkov muon veto for the xenon1t experiment, JINST **9** (2014) P11006 [1406.2374].

[153] J. Wulf, Direct Dark Matter Search with XENON1T and Developments for Multi-Ton Liquid Xenon Detectors, Ph.D. thesis, Universität Zürich, 2018.

[154] "Mongodb." https://mongodb.com.

[155] J. Aalbers, C. Tunnell, B. Pelssers, A. Buß, F. Gao, Q. Lin et al., "PAX data processor." https://doi.org/10.5281/zenodo.4290785, Mar., 2020. 10.5281/zenodo.4290785.

[156] J. Aalbers, J.R. Angevaare, D. Wenz, C. Tunnell, pyup.io bot, Y. Mosbacher et al., "Axfoundation/strax: v1.4.0." https://doi.org/10.5281/zenodo.7216048, Oct., 2022. 10.5281/zenodo.7216048.

[157] J.R. Angevaare, J. Aalbers, D. Wenz, E. Shockley, P. Gaemers, A. Higuera et al., "Xenonnt/straxen: v2.0.3." https://doi.org/10.5281/zenodo.7308136, Nov., 2022. 10.5281/zenodo.7308136.

[158] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau et al., Array programming with NumPy, Nature **585** (2020) 357.

[159] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau et al., SciPy 1.0: Fundamental algorithms for scientific computing in python, Nature Methods **17** (2020) 261.

[160] S.K. Lam, A. Pitrou and S. Seibert, Numba: A llvm-based python jit compiler, in Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, pp. 1–6, 2015.

[161] M. De Deo et al., Accurate GPS-based timestamp facility for Gran Sasso National Laboratory, JINST **14** (2019) P04001.

[162] SkuTek Instrumentation, Digital Pulse Processor for Nuclear Physics. https://www.skutek.com/ddc10.htm.

[163] "Ceph." https://ceph.io.

[164] D. Masson, D. Coderre, J.R. Angevaare, A. Elykov, S.D. Pede, P. Gaemers et al., "Axfoundation/redax: Version 2.3.0." https://doi.org/10.5281/zenodo.5882717, Jan., 2022. 10.5281/zenodo.5882717.

[165] V.C. Antochi et al., Improved quality tests of r11410-21 photomultiplier tubes for the xenonnt experiment, JINST **16** (2021) P08033 [2104.15051].

[166] XENON collaboration, Application and modeling of an online distillation method to reduce krypton and argon in xenon1t, PTEP **2022** (2022) 053H01 [2112.12231].

[167] "Slack." https://slack.com/.

[168] S. Al Kharusi, S.Y. BenZvi, J.S. Bobowski, W. Bonivento, V. Brdar, T. Brunner et al., SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy, New Journal of Physics **23** (2021) 031201.

[169] "Nodejs." https://nodejs.dev.

[170] "Straxen documentation." https://straxen.readthedocs.io/.

[171] G.F. Knoll, Radiation Detection and Measurement, 3rd ed., John Wiley and Sons, New York, 3rd edition ed. (2000).

[172] C.H. Faham, V.M. Gehman, A. Currie, A. Dobi, P. Sorensen and R.J. Gaitskell, Measurements of wavelength-dependent double photoelectron emission from single photons in vuv-sensitive photomultiplier tubes, JINST **10** (2015) P09010 [1506.08748].

[173] XENON collaboration, Emission of single and few electrons in xenon1t and limits on light dark matter, Phys. Rev. D **106** (2022) 022001 [2112.12116].

[174] G. Volta, "Single photoelectron shape for science run 0." https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:giovo:spe_spectra_sr0_new, June, 2022.

[175] M.V. Diwan, Statistics of the charge spectrum of photo-multipliers and methods for absolute calibration, JINST **15** (2020) P02001 [1909.05373].

[176] P. Gaemers, T. Zhu, J.R. Angevaare, D.R. García, K. Mizukoshi, J. Aalbers et al., "XENONnT/WFSim: v1.0.2." https://doi.org/10.5281/zenodo.7216324, Oct., 2022. 10.5281/zenodo.7216324.

[177] P. Barrow et al., Qualification tests of the r11410-21 photomultiplier tubes for the xenon1t detector, JINST **12** (2017) P01024 [1609.01654].

[178] L. Yuan, "Xenonnt sr0 wfsim-data s1 matching summary." https://xe1t-wiki.lngs.infn.it/doku.php?id=lanqing:wfsim_data_peak_level_s1matching, March, 2022.

[179] J. Qi, "Wfsim single electron improvements." https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:jianyang:wfsim_data_se_matching, March, 2022.

[180] L. Yuan and E. Angelino, "Xenonnt sr0 data-driven s1 efficiency summary." https://xe1t-wiki.lngs.infn.it/doku.php?id=lanqing:data_driven_s1_sr0, May, 2022.

[181] XENON100 collaboration, Observation and applications of single-electron charge signals in the xenon100 experiment, J. Phys. G **41** (2014) 035201 [1311.1088].

[182] B. Edwards et al., Measurement of single electron emission in two-phase xenon, Astropart. Phys. **30** (2008) 54 [0708.0768].

[183] J. Mock, N. Barry, K. Kazkaz, M. Szydagis, M. Tripathi, S. Uvarov et al., Modeling pulse characteristics in xenon with nest, JINST **9** (2014) T04002 [1310.1117].

[184] J. Aalbers, Dark matter search with XENON1T, Ph.D. thesis, Universiteit van Amsterdam, 2018.

[185] XENON100 collaboration, The xenon100 dark matter experiment, Astropart. Phys. **35** (2012) 573 [1107.2155].

[186] V. D'Andrea, R. Peres and S. Kazama, "Optimization of s2 width cut for xenonnt sr0 lower unblinding." https://xe1t-wiki.lngs.infn.it/doku.php?id=dandrea:s2widthcutsr0update20220618, June, 2022.

[187] T. Doke, A. Hitachi, J. Kikuchi, K. Masuda, H. Okada and E. Shibamura, Absolute scintillation yields in liquid argon and xenon for various particles, Jap. J. Appl. Phys. **41** (2002) 1538.

[188] K. Morå, "A conservative upper bound on g1g2 error." https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:analysis:ntsciencerun0:g1g2_conservative_error, June, 2022.

[189] H. Schulze Eißing, M. Guida and Y. Ma, "g1 g2 analysis for sr0." https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:analysis:ntsciencerun0:g1g2_update, May, 2022.

[190] XENON collaboration, Low-energy calibration of xenon1t with an internal ^{37}ar source, arXiv pre-print (2022) [2211.14191].

[191] XENON collaboration, Results from a calibration of xenon100 using a source of dissolved radon-220, Phys. Rev. D **95** (2017) 072008 [1611.03585].

[192] J. Scherzinger, J. Annand, G. Davatz, K. Fissum, U. Gendotti, R. Hall-Wilton et al., Tagging fast neutrons from an 241am/9be source, Applied Radiation and Isotopes **98** (2015) 74.

[193] D. Woodward, First results from the lux-zeplin (lz) dark matter experiment, 14th Conference on the intersections of Particle and Nuclear Physics, LZ collaboration, August, 2022.