

MAX-PLANCK-INSTITUT
FÜR PHYSIK

TECHNISCHE UNIVERSITÄT MÜNCHEN
FAKULTÄT FÜR PHYSIK

MAX-PLANCK-INSTITUT FÜR PHYSIK
WERNER-HEISENBERG-INSTITUT

Low-Threshold Cryogenic Detectors for Low-Mass Dark Matter Search and Coherent Neutrino Scattering

Johannes Felix Martin Rothe

Vollständiger Abdruck der von der Fakultät für Physik der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigten Dissertation.

Vorsitzender: Prof. Dr. Andreas Weiler

Prüfende der Dissertation:

1. Prof. Dr. Stefan Schönert
2. Prof. Dr. Susanne Mertens

Die Dissertation wurde am 22.10.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 14.01.2021 angenommen.

Abstract

The study of rare processes is crucial to our modern understanding of physics. This work is concerned with a particular class of rare process: nuclear recoils in the previously inaccessible energy range of 10-100 eV. More than 25 years of searching for dark matter with the CRESST experiment has yielded a mature, very sensitive technology for rare event search. This work has focused on the optimization of these cryogenic detectors, with a detailed analysis of the trade-off between target mass and energy threshold. For physics cases with rising rates towards low energies, such as direct detection of low-mass dark matter and coherent elastic neutrino-nucleus scattering ($CE\nu NS$), increased sensitivity can be achieved with smaller detectors. This approach is fundamental to the strategy of CRESST-III. In addition, it opens a way to a cryogenic detection of reactor neutrinos. This new concept developed here helped initiate the NUCLEUS experiment.

A study of the impact of magnetic fields on transition edge sensors (TES) is described, showing sensitivity and stability of detectors can be improved by canceling background fields. With the conception and installation of a magnetic field compensation system at LNGS, this work contributed to the stable operation of CRESST detectors with energy thresholds far below 100 eV. In above-ground measurements, different TES designs were characterised on multiple crystal materials, laying the foundation for the next generation of CRESST detectors with further improved energy resolution.

Sensitivity calculations for $CE\nu NS$ detection are presented, which are the basis for a new experimental effort to detect coherent neutrino-nucleus scattering at a power reactor. The feasibility of a measurement is shown and the driving parameters for the experiment are identified. Detailed studies of sensitivity to beyond-the-Standard-Model physics scenarios are performed that motivate the experimental effort. The first cryogenic detector prototypes for NUCLEUS were commissioned in this work. The most important requirement, a nuclear recoil energy threshold in the 10-20 eV range, has been achieved. This resulted in new leading constraints on very light dark matter. The functionality of a novel detector-holder arrangement that provides cryogenic anticoincidence vetos against surface and penetrating backgrounds has been demonstrated. This is a promising technique to mitigate limiting low-energy backgrounds in the future.

This thesis is organized as follows. In chapter 1, the physics cases are presented. A historical perspective and an overview of experimental approaches to dark matter direct detection and $CE\nu NS$ are given. The possible implications of a detection on the Standard Model of particle physics is described. While a direct detection of dark matter particle would itself be a sign of physics beyond the Standard Model, $CE\nu NS$ is a well-predicted process within the Standard Model. A precision measurement can uncover deviations caused by non-standard neutrino properties. Chapter 2 introduces cryogenic calorimeters, technical implementations and the particular solution studied in the rest of this work: detectors using tungsten TES sensitive to athermal phonon signals. An important scaling relation is derived: reducing the target mass in an athermal phonon

detector improves its energy resolution following a power-law. The energy threshold can be lowered, with important implications for CRESST and NUCLEUS detectors.

Chapter 3 introduces the CRESST experiment located at the underground laboratory of Laboratori Nazionali del Gran Sasso, a leading effort to detect low-mass dark matter through its potential rare interactions with atomic nuclei.

Chapter 4 describes the first systematic study of the impact of magnetic fields on the operation of sensors in CRESST. Background fields as low as a few μT are shown to adversely affect the superconducting transition of the TES. Within this work, an active field compensation system was developed and implemented at the CRESST setup. Comparison studies with operating cryogenic detectors show an increase in detector gain up to a factor of 2 through the suppression of the Earth magnetic field, and improved stability by compensating magnetic transients.

Chapter 5 deals with the TES design for the ongoing experimental campaign CRESST-III. Different sensor sizes are studied on Al_2O_3 , CaWO_4 and Si crystals developed for use in CRESST. Phonon thermalization times are shown to be reduced with larger-area TES structures, suggesting an improved energy resolution achievable by smaller designs. A fundamental limitation to downsizing TES films is explored, caused by thermal diffusion in extended metallic heat links.

Chapter 6 introduces the NUCLEUS experiment, a new approach attempting the first observation of reactor neutrinos with cryogenic detectors. The basic premises of the experiment, required energy threshold and target masses, were developed within this work and are described in detail. The new approach to background suppression using cryogenic anticoincidence vetos is explained. An “inner veto” acts as instrumented holder and 4π surface veto. A massive “outer veto” tags multiple-scattering backgrounds.

Chapter 7 encompasses detailed sensitivity studies performed for the NUCLEUS experiment. A study of different possible neutrino sources identifies a power reactor as the preferred site. A benchmark scenario is developed, and the conditions under which a $\text{CE}\nu\text{NS}$ observation is possible are defined in terms of reactor distance, energy threshold and background level. Various physics scenarios beyond the Standard Model of particle physics are explored, and the potential of NUCLEUS to detect statistically significant deviation from standard $\text{CE}\nu\text{NS}$ is quantified.

Chapter 8 describes the characterization measurements performed in this work with the first NUCLEUS detector elements. The NUCLEUS-1g prototype contains a single target detector and the cryogenic vetos. The first operation of a target detector showed a world-leading nuclear recoil energy threshold of 19.7 eV. New constraints on light dark matter down to masses of $140 \text{ MeV}/c^2$ were derived from the measurement. Over several experimental campaigns, the target detector was operated in the newly designed holder, and the cryogenic vetos were demonstrated for NUCLEUS. Below ~ 200 eV recoil energy, a large number of background events currently prevents sensitivity to $\text{CE}\nu\text{NS}$. The prototype NUCLEUS detector commissioned in this work will be used to characterize this background and develop methods to suppress it.

Zusammenfassung

Die Untersuchung seltener Ereignisse trägt entscheidend zu unserem modernen Verständnis der Physik bei. Diese Arbeit befasst sich mit einer bestimmten Sorte seltener Ereignisse: Kernrückstöße im bisher unzugänglichen Energiebereich von 10-100 eV. Über 25 Jahre der Suche nach Dunkler Materie mit dem CRESST-Experiment haben eine ausgereifte, sehr empfindliche Technologie zur Suche nach seltenen Ereignissen hervorgebracht. Diese Arbeit konzentriert sich auf die Verbesserung dieser kryogenen Detektoren, mit einer detaillierten Analyse des Trade-Offs zwischen Targetmasse und Energieschwelle. Bei Physik-Anwendungen mit zu kleinerer Energie ansteigenden Raten, wie die direkte Suche nach leichter Dunkler Materie und kohärent-elastischer Neutrino-Kern-Streuung ($CE\nu NS$), kann verbesserte Empfindlichkeit mit kleineren Detektoren erreicht werden. Dieser Ansatz ist fundamental für die Strategie von CRESST-III. Zusätzlich eröffnet er eine Möglichkeit, Reaktor-neutrinos mit kryogenen Detektoren nachzuweisen. Dieses neue Konzept das in dieser Arbeit entwickelt wurde, half, das NUCLEUS-Experiment zu initiieren.

Eine Studie des Einflusses von Magnetfeldern auf Übergangskantensensoren (TES) wird beschrieben. Sie zeigt, dass Empfindlichkeit und Stabilität der Detektoren durch Kompensation des Hintergrundfeldes verbessert werden können. Durch den Entwurf und die Inbetriebnahme eines Magnetfeldkompensationssystems am LNGS trug diese Arbeit zum stabilen Betrieb der CRESST-Detektoren mit Energieschwellen weit unterhalb von 100 eV bei. In Messungen an der Erdoberfläche wurden verschiedene TES-Entwürfe auf unterschiedlichen Kristall-Materialien charakterisiert. Dies stellt die Grundlage für die nächste Generation von CRESST-Detektoren mit weiter verbesserter Energieauflösung dar.

Sensitivitätsstudien für die Detektion von $CE\nu NS$ werden präsentiert, aus denen ein neues experimentelles Programm entstand, um kohärente Neutrino-Kern-Streuung an einem Kernkraftwerk zu detektieren. Die Machbarkeit der Messung wird gezeigt und die kritischen Parameter für das Experiment werden bestimmt. Ausführliche Studien der Empfindlichkeit auf Physik-Szenarien jenseits des Standardmodells werden durchgeführt, die das experimentelle Programm motivieren. Die ersten Prototypen für kryogene Detektoren für NUCLEUS wurden in dieser Arbeit in Betrieb genommen. Die wichtigste Voraussetzung, eine Energieschwelle für Kernrückstöße im Bereich von 10-20 eV, wurde erfüllt. Dies ergab neue führende Ausschlussgrenzen für sehr leichte Dunkle Materie. Die Einsetzbarkeit einer neuartigen Detektor-Halter-Vorrichtung wurde gezeigt. Sie ermöglicht den Einsatz kryogener Antikoinzidenzvetos gegen Oberflächen- und durchdringende Untergründe. Dieser technische Ansatz ist vielversprechend für die zukünftige Abschwächung begrenzender niederenergetischer Untergründe.

Diese Arbeit ist wie folgt gegliedert. In Kapitel 1 werden die physikalischen Anwendungsgebiete vorgestellt. Eine historische Einordnung und ein Überblick der experimentellen Ansätze zum direkten Nachweis Dunkler Materie und von $CE\nu NS$ wird gegeben. Die Bedeutung einer Entdeckung auf das Standardmodell der Teilchenphysik wird be-

schrieben. Während der direkte Nachweis eines Teilchens der Dunklen Materie selbst ein Zeichen von Physik jenseits des Standardmodells wäre, ist $CE\nu NS$ ein genau vorhergesagter Prozess innerhalb des Standardmodells. Eine Präzisionsmessung kann Abweichungen aufdecken, die von nicht dem Standardmodell entsprechenden Eigenschaften von Neutrinos hervorgerufen würden. Kapitel 2 führt kryogene Kalorimeter ein, technische Umsetzungen und im Besonderen die Lösung, mit der sich der Rest der Arbeit befasst: Detektoren mit Wolfram-TES, die für Signale nicht-thermischer Phononen empfindlich sind. Eine wichtige Skalierungs-Beziehung wird hergeleitet: eine Verringerung der Targetmasse in einem Detektor für nicht-thermische Phononen verbessert die Energieauflösung nach einem Potenzgesetz. Die Energieschwelle kann dann gesenkt werden, mit wichtigen Konsequenzen für CRESST- und NUCLEUS-Detektoren.

Kapitel 3 führt das CRESST-Experiment am Untergrundlabor Laboratori Nazionali del Gran Sasso ein, ein führendes Programm zur Suche nach leichter Dunkler Materie mittels möglicher seltener Wechselwirkungen mit Atomkernen.

Kapitel 4 beschreibt die erste systematische Untersuchung der Auswirkungen von Magnetfeldern auf den Detektorbetrieb in CRESST. Hintergrundfelder von nur wenigen μT beeinflussen den supraleitenden Übergang der TES negativ. In dieser Arbeit wurde ein aktives Feldkompensationssystem entworfen und am CRESST-Aufbau in Betrieb genommen. Vergleichsstudien mit laufenden Detektoren zeigen eine bis zu einem Faktor zwei höhere Verstärkung durch die Unterdrückung des Erdmagnetfelds, und verbesserte Stabilität durch die Kompensation magnetischer Transienten.

Kapitel 5 behandelt das Design der TES für die laufende Phase des Experiments CRESST-III. Verschieden große Sensoren wurden auf Al_2O_3 -, $CaWO_4$ - und Si-Kristallen untersucht, die für den Einsatz in CRESST entwickelt wurden. Es wird gezeigt, dass die Thermalisierungszeit nicht-thermischer Phononen durch die größeren TES-Strukturen verkürzt wird. Dies deutet darauf hin, dass mit kleineren TES-Flächen eine bessere Energieauflösung erreichbar ist. Eine grundlegende Grenze für die Verkleinerung der TES-Filme wird untersucht, die sich aus thermischer Diffusion in ausgedehnten metallischen Wärmeverbindungen ergibt.

Kapitel 6 führt das NUCLEUS-Experiment ein, das mit einem neuen Ansatz auf den ersten Nachweis von Reaktor-neutrinos mit kryogenen Detektoren hinarbeitet. Die Grundvoraussetzungen dieses Experimentes, die nötige Energieschwelle und Targetmassen, wurden in dieser Arbeit hergeleitet und werden im Detail beschrieben. Der neue Ansatz zur Unterdrückung von Untergrund mittels kryogener Antikoinzidenzveto wird erklärt. Ein "inneres Veto" wird als instrumentierter Halter und 4π -Oberflächen veto verwendet. Ein massives "äußeres Veto" markiert Untergrund durch Mehrfachstreuungen.

Kapitel 7 umfasst detaillierte Sensitivitätsstudien, die für das NUCLEUS-Experiment durchgeführt wurden. Eine Untersuchung verschiedener möglicher Neutrinoquellen identifiziert ein Kernkraftwerk als bevorzugten Standort. Ein Vergleichs-Szenario wird entwickelt, und die Bedingungen, unter denen eine Beobachtung von $CE\nu NS$ möglich ist werden als Funktion des Reaktorabstands, der Energieschwelle und des Untergrundniveaus eingegrenzt. verschiedene Physik-Szenarien jenseits des Standardmodells werden erkundet, und das Potential von NUCLEUS, statistisch signifikante Abweichungen von standard- $CE\nu NS$ nachzuweisen, wird quantifiziert.

Kapitel 8 beschreibt die Charakterisations-Messungen, die in dieser Arbeit mit den ersten Detektor-Elementen für NUCLEUS durchgeführt wurden. Der NUCLEUS-1g Prototyp enthält einen einzelnen Target-Detektor und die kryogenen Vetos. Die erste In-

betriebsnahme eines Target-Detektor zeigte eine weltweit führende Energieschwelle für Kernrückstöße von 19.7 eV. Neue Ausschlussgrenzen auf leichte Dunkle Materie mit Massen bis zu $140 \text{ MeV}/c^2$ wurden aus der Messung abgeleitet. Über mehrere Messkampagnen wurde ein Target-Detektor im neu entworfenen Halter betrieben, und die kryogenen Vetos wurden für NUCLEUS demonstriert. Unterhalb von $\sim 200 \text{ eV}$ in Kernrückstoß-Energie verhindert momentan eine große Anzahl von Untergründereignissen Empfindlichkeit auf $\text{CE}\nu\text{NS}$. Der NUCLEUS Prototyp-Detektor, der in dieser Arbeit in Betrieb genommen wurde, wird verwendet werden um diesen Untergrund zu charakterisieren und Methoden zu entwickeln, ihn zu unterdrücken.

Contents

Abstract	i
Zusammenfassung	iii
1. Rare-event searches at low energies	1
1.1. Direct detection of dark matter	1
1.1.1. The dark matter problem	1
1.1.1.1. Gravitational evidence for dark matter	1
1.1.1.2. Particle candidates and alternatives	5
1.1.2. Experimental approaches	10
1.1.2.1. Collider Searches	11
1.1.2.2. Indirect detection	13
1.1.2.3. Direct detection	14
1.1.2.4. Comparability and complementarity	18
1.1.3. Technology for direct detection	19
1.1.3.1. Semiconducting/scintillating crystals	19
1.1.3.2. Cryogenic detectors	20
1.1.3.3. Noble liquid detectors	20
1.1.3.4. Bubble chambers	21
1.1.3.5. Directional detectors	21
1.2. Coherent elastic neutrino-nucleus scattering	22
1.2.1. Prediction and significance	22
1.2.2. Cross-section in the Standard Model	23
1.2.3. Sensitivity to New Physics	26
1.2.3.1. Weak mixing angle	26
1.2.3.2. Non-standard interactions	27
1.2.3.3. Electromagnetic form-factors	28
1.2.3.4. New light mediators	32
1.2.4. Artificial neutrino sources for CE ν NS	33
1.2.4.1. Stopped-pion sources	33
1.2.4.2. Nuclear reactors	34
1.2.4.3. Radiogenic sources	35
1.2.5. CE ν NS experiments	35
1.2.5.1. The COHERENT experiment	36
1.2.5.2. Upcoming CE ν NS experiments at nuclear reactors	38
2. Cryogenic detectors for rare event searches	41
2.1. Cryogenic calorimeters	41
2.2. Cryogenic sensor technologies for fundamental physics	42
2.2.1. Neutron transmutation doped germanium thermistors	42
2.2.2. Metallic magnetic microcalorimeters	43

2.2.3.	Microwave kinetic inductance detectors	44
2.2.4.	Transition edge sensors	45
2.3.	Principles of TES macrocalorimeters	46
2.3.1.	Thermal detector model	46
2.3.2.	SQUID readout	50
2.3.3.	Temperature stabilization	51
2.3.4.	Scaling of energy resolution with target mass	52
3.	The CRESST experiment	55
3.1.	The CRESST Facility	55
3.2.	The phonon-light technique	57
3.3.	The CRESST-III low-mass dark matter search	58
3.3.1.	Detector modules in CRESST-III	58
3.3.2.	CRESST-III low-mass dark matter results	60
3.3.3.	Exploration of low-energy background	61
4.	Active magnetic field compensation for CRESST	63
4.1.	Magnetic fields and superconducting films	63
4.2.	TES response to transverse fields	64
4.2.1.	Measurement principle	64
4.2.2.	Observation of a large-area sputtered tungsten film	65
4.2.3.	Observation of a structured evaporated tungsten TES	67
4.3.	Fields and coils at CRESST	71
4.3.1.	The magnetic environment at LNGS	71
4.3.2.	Magnetic field compensation setup	72
4.4.	Static field compensation	76
4.4.1.	Identifying the zero-field configuration	76
4.4.2.	Transition measurement	81
4.5.	Active magnetic field control	84
4.5.1.	Magnetometer noise and sensitivities	84
4.5.2.	Field compensation algorithm	87
4.5.2.1.	Minimizing magnetic field at magnetometer positions	87
4.5.2.2.	Minimizing magnetic field at an extrapolated location	87
4.5.3.	Perturbation test	88
4.5.4.	Zero-field detector performance	90
4.6.	Outlook	94
5.	TES design for CRESST-III	95
5.1.	Scaling of Transition Edge Sensors	95
5.1.1.	TES on CaWO_4 at LNGS	97
5.1.2.	TES on Si and Al_2O_3 at MPP	99
5.1.3.	Discussion	101
5.2.	Thermal diffusion as a limit to weak link designs	103
5.2.1.	Diffusion in an extended thermal link	103
5.2.2.	Expected effect on different TES designs	105
5.2.3.	Implications for future TES designs	106
5.2.4.	Proposed experimental study	107

6. The NUCLEUS experiment	109
6.1. Ultra-low threshold detectors for CE ν NS	109
6.1.1. Neutrino flux model and recoil spectra	109
6.1.2. Gram-scale cryogenic detectors	111
6.1.3. Science phases of NUCLEUS	113
6.1.3.1. NUCLEUS-10g	113
6.1.3.2. NUCLEUS-1kg	115
6.2. A fiducial-volume cryogenic detector	116
6.2.1. Target array	116
6.2.2. Cryogenic inner veto	117
6.2.3. Cryogenic outer veto	119
6.3. The NUCLEUS experimental site	120
6.3.1. The Very-Near-Site at Chooz	120
6.3.2. Radiogenic background	121
6.3.3. Cryogenic infrastructure	123
6.3.4. Passive shielding	123
6.3.5. Muon anticoincidence veto	124
7. Sensitivity studies for NUCLEUS	127
7.1. Comparison of Neutrino Sources for NUCLEUS	128
7.1.1. Stopped-pion source	128
7.1.2. Research reactor	129
7.1.3. Radiogenic neutrino source	131
7.2. CE ν NS Observation Potential at a Power Reactor	134
7.2.1. Benchmark case and variations	134
7.2.2. Multi-target approach: background of unknown shape	136
7.2.3. Correlation with reactor power	139
7.3. Studying New Physics at VNS	142
7.3.1. Achievable precision at VNS	143
7.3.2. Weak mixing angle	144
7.3.3. Non-standard Neutrino Interactions	145
7.3.4. Neutrino electromagnetic form factors	150
7.3.5. New Light Mediators	153
8. The NUCLEUS-1g prototype	159
8.1. Prototype components	159
8.1.1. Target detectors	160
8.1.2. Inner Veto	161
8.1.3. Holding Structure	162
8.1.4. Outer Veto	163
8.2. Experimental setup	164
8.2.1. Cryogenic infrastructure	164
8.2.2. Data Acquisition	164
8.3. Measurements	165
8.3.1. Prototype Run 1: first target detector	165
8.3.1.1. Detector response and template pulse	167
8.3.1.2. Calibration and truncated template fit	168

8.3.1.3.	Optimum filter method and theoretical energy resolution	169
8.3.1.4.	Combination of energy reconstruction methods	172
8.3.1.5.	Energy resolution from simulated events	172
8.3.1.6.	Threshold from trigger efficiency simulation	173
8.3.1.7.	Estimate of noise trigger spectrum	175
8.3.1.8.	Stream dataset: Dark Matter Analysis	178
8.3.2.	Prototype Run 2: background measurement in silicon holder . . .	182
8.3.3.	Prototype Run 3: first outer veto operation	188
8.3.4.	Prototype Run 4: target and active inner veto	191
8.3.5.	Conclusions from Prototype Runs	198
9.	Conclusion and Outlook	201
A.	Appendix to active magnetic field compensation for CRESST	205
A.1.	Fields of rectangular coils	205
A.2.	Determining TES orientations	205
A.3.	Simulation of compensation algorithm	208
A.4.	Configuration of field compensation at CRESST	210
B.	Statistical methods for sensitivity studies	215
B.1.	Likelihood-Ratio-Test	215
B.2.	Unbinned likelihood	215
B.3.	Binned likelihood	216
B.4.	Toy Monte Carlo	217
B.5.	Asimov data set	217
B.6.	Nuisance parameters	219

1. Rare-event searches at low energies

1.1. Direct detection of dark matter

1.1.1. The dark matter problem

Arguably the most important “known unknown” in fundamental physics today concerns the nature of Dark Matter. The consequences of its gravitational interaction have been mapped out by astronomers over a wide range of length-scales with increasing precision during the past century. Dark matter has come to be accepted as a necessary ingredient in our understanding of the universe and its development. At the same time, a bewildering variety of theoretical models compete to explain its origin, fundamental nature, and the relation of dark matter to visible matter. The experimental quest for non-gravitational signatures of dark matter has ramped up over the past decades and turned into a global endeavour joining fields from astronomy (with light, particles and gravitational waves) to high-energy physics, quantum sensing and low-radioactivity underground physics. Not a single detection, but many measurements following different methodologies and employing a variety of technologies, are required to bring the endeavour of “understanding dark matter” to fruition.

This section provides an overview of the different observations of dark matter in the universe through gravitational effects, and outlines physical models of both particle dark matter and alternatives.

The next section explores the complementary methodologies of hunting for a non-gravitational signature of particle dark matter. The chapter concludes with an overview of the plethora of detector technologies employed in direct searches for dark matter.

1.1.1.1. Gravitational evidence for dark matter

Often, Fritz Zwicky is credited with the “invention” of dark matter, in his 1933 study of the motion of galaxies in the Coma cluster [1]. In fact, a number of earlier works estimating the total mass of the Galaxy from stellar motion [2, 3, 4, 5, 6] introduce the idea of dark, invisible or unseen matter. The journey from the first order-of-magnitude estimates of “missing mass” to a precise cosmological model including non-baryonic dark matter took many decades with decisive inputs from many pioneers. For a review of dark matter history see [7]. At this point, there is strong astronomical evidence that 85% of the matter in the universe (26% of the mass-energy density) is in the form of cold dark matter. In the following, this evidence is presented from large to small length-scales.

Cosmic microwave background and large-scale structure The cosmic microwave background (CMB) was first observed by Penzias and Wilson in 1965 [9] and interpreted by Dicke and others [10] as the afterglow of the Big Bang. It consists of photons streaming freely through the cosmos since 380.000 years after the Big Bang, when the primordial plasma became transparent. Tiny temperature anisotropies (on the level of 10^{-5}) carry




Figure 1.1.: Temperature power spectrum of the cosmic microwave background, for different dark matter density parameters $\Omega_{cdm} = 0.11 \dots 0.43$. The relative heights of the first few acoustic peaks allow to independently constrain dark matter and baryon density in the early universe. Figure taken from [8].

information of density fluctuations in the primordial plasma, which in turn depend on its composition.

The observed power spectrum (oscillation strength as a function of angular scale, see Fig. 1.1) can only be explained by assuming the presence of non-baryonic cold dark matter (CDM) in addition to normal matter and radiation. Normal matter (mostly electrons and protons, “baryons” in cosmological language) is charged, causing internal pressure, which leads to sound waves propagating in the plasma before release of the CMB. Dark matter density fluctuations simply grow by gravitational collapse and do not oscillate. This allows independent measurements of the baryon density parameter $\Omega_b h^2 = 0.02237 \pm 0.00015$ and CDM density parameter $\Omega_{cdm} h^2 = 0.1200 \pm 0.0012$ (with the reduced Hubble parameter $h = H_0/100$ km/s/Mpc) [11]. Density parameters $\Omega_x = \rho_x/\rho_c$ relate a mass density to the cosmological critical density $\rho_c = 3H_0^2/8\pi G \approx 8.4 \cdot 10^{-27}$ kg/m³ needed to close the universe.

Complementary information can be derived from large-scale structure. The observed CMB anisotropies by themselves are not large enough to explain structure formation in the universe since decoupling. One has to take into account the density perturbations in dark matter, which are not supported by pressure in the plasma and undergo gravitational collapse already before decoupling. Baryons subsequently fall into the gravitational potential of the pre-formed dark matter structures, which accelerates the growth process. Furthermore, the galaxy distribution today carries direct information about processes in the early universe. Mapping out the three-dimensional positions of ten-thousands of galaxies, evidence for a correlation over a distance around 150 Mpc has been found [12]. The length scale of this “baryon acoustic peak” corresponds to the distance traveled by the baryon sound wave in the early universe (before decoupling of matter and radiation), away from the original dark matter density distribution. Together with the CMB

anisotropies, a consistent and precisely constrained understanding of the early universe emerges, including roughly five times more dark matter than normal matter.




Figure 1.2.: The Bullet Cluster (1E 0657-57), in a composite image of optical (yellow) and X-ray (magenta) images and a mass map (blue) determined by weak gravitational lensing. Credit: X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al.. Figure taken from [13].

Galaxy clusters The influence of dark matter in galaxy clusters has been observed in a variety of ways. Fritz Zwicky studied the redshift of eight galaxies within the Coma Cluster [1] and noted a large deviation among the galaxies, corresponding to radial velocity differences above 2000 km/s. Using the virial theorem (assuming dynamical equilibrium) he converted this to an estimate of total mass of the galaxy cluster. Comparing to a visible mass derived by estimating mass and number of the constituting galaxies, he concluded the cluster dynamics to be dominated by unseen “dark matter”.

Modern observations of galaxy clusters allow much more refined conclusions. Maps of X-ray emission from the hot intracluster gas allow separately measuring total, gas and stellar mass for many individual galaxy clusters [14], finding consistently baryon fractions of around 1/6.

The effect of gravitational lensing, i.e. deflection of light passing near massive objects, can be used to map out the gravitational mass distribution in a galaxy cluster [15]. In the case of strong lensing, multiple images of a fortunately aligned background object appear around the foreground galaxy cluster. The positions and distortions of these images are very sensitive to the mass distribution of the lensing cluster. In weak lensing, no lucky alignment is needed: through the detection of systematic patterns in the slight distortions of hundreds of background galaxies again the mass distribution of the foreground galaxy

cluster can be inferred.

The most visually striking evidence for dark matter comes from a combination of such measurements. The Bullet Cluster [16] (see Fig. 1.2) consists of a pair of galaxy clusters about 150 million years after colliding with each other. The distributions of total mass, gas, and stars can be mapped out separately (by lensing, X-ray and optical imaging respectively). The gas in the two clusters, dominating the baryonic mass, interacted strongly and was slowed down. The well-spaced galaxies and stars within them continued moving largely unperturbed by the merger and are now far ahead of the gas clouds. The lensing map reveals that the dominant, unseen mass concentrations of the two clusters ended up near the galaxies, showing that dark matter is collisionless. This observation places limits on the self-interaction of dark matter particles.

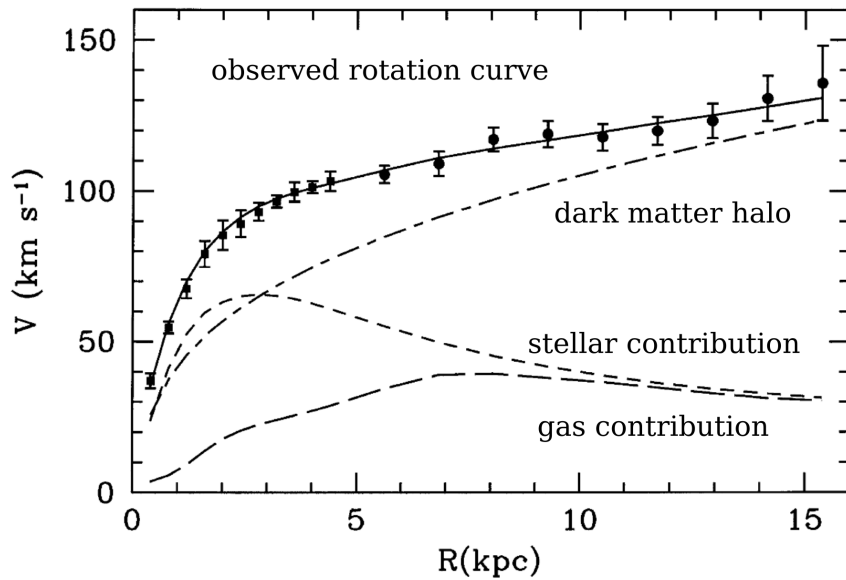


Figure 1.3.: Rotation curve of the Triangulum Galaxy M33, adapted from [17], measured by 21 cm spectroscopy. The short-dashed, long-dashed and dot-dashed lines show the modeled contributions from stars, gas, and the dark matter halo, respectively. The fact that the rotation curve keeps rising far beyond the visible stellar disk is explained by the extended dark matter halo surrounding the galaxy.

Galactic rotation curves The rotation curves of galaxies played an especially important role in convincing the astronomical community of the existence of unobserved missing mass. Starting in the early days of extragalactic astronomy in the beginning of the 20th century, the rotation of galaxies was used to infer their mass-to-light ratio and thus make statements about their composition. Several decades of technological advancement was needed to bring the field to the point where clear discrepancies with the assumption of spatially constant mass-to-light ratio could be claimed in the 70es.

Among the many increasingly precise observations, the most famous examples of “flat rotation curves” come from Rubin, Ford and Thonnard’s [18] optical spectroscopic observations, and Albert Bosma’s [19] 21-cm observations of atomic hydrogen clouds ex-

tending out to radii beyond the visible stellar disks. Assuming Newton’s law of gravity, this shows that the mass distribution in the observed galaxies falls off much more slowly than the visible stellar distribution. In 1974, two influential papers [20, 21] concluded that large amounts of unseen mass is present around galaxies. The modern explanation is that all galaxies are embedded in extended, nearly spherical dark matter halos. Fig. 1.3 shows an example of a rotation curve, measured in 21 cm emission for the nearby galaxy M33 [17]. A fit to a mass model allows to decompose the rotation curve into contributions from stars, interstellar gas, and the dark matter halo. At large radii, the gravitational influence of the dark matter dominates.

Local stellar motion Studying the motion of stars near the sun is arguably the first method for constraining the amount of unseen matter in the universe, dating back to the works of Kapteyn, Poincaré, Öpik, Jeans and Oort [22, 3, 4, 5, 6]. At the same time, it is the only direct way to estimate the dark matter density in the solar system: gravitational effect on bodies inside the solar systems are unobservably small. The local dark matter density is an important parameter for direct detection experiments attempting to observe scattering of dark matter in an Earth-bound detector.

Two basic approaches are possible to constrain the dark matter density near the sun: local methods analysing the dynamics of tracer stars in the solar vicinity, and global methods modeling the whole Galaxy. In general, global methods can reach smaller uncertainties, but have to make more assumptions on the overall shape of the Galactic halo. Nearly a century of progress in local dark matter measurements is described in the review [23]. Modern estimates fall in the range of $\rho_{\text{dm}} = 0.2 - 0.5 \text{ GeV}/c^2/\text{cm}^3$, in the particle-physics units used by direct-detection experiments¹. A wealth of stellar kinematic data has become available from Galactic surveys such as the Hipparcos astrometry mission [24] and the Sloan Digital Sky Survey [25]. Most recently, the long-awaited ESA/Gaia mission [26] precisely measured proper motion and radial velocity for hundreds of millions of stars in the Galaxy. This revolutionary dataset has so far not resulted in model-independent precision measurements of the local dark matter density. The more complete information requires correspondingly more complex models for the Galaxy. This has resulted in similarly uncertain and discrepant density estimates as in the past, though relying on fewer assumptions on general properties of the Galaxy’s dark matter halo [27].

1.1.1.2. Particle candidates and alternatives

The wealth of gravitational evidence for missing mass is nowadays often taken as the most concrete indication for physics beyond the Standard Model, in the form of a new elementary particle. There are compelling theoretical models explaining all the different dark matter observations by adding a single new ingredient to standard Big Bang cosmology. Still, one should keep in mind that the particle nature of dark matter is as yet unproven. This section first discusses some alternative ideas, before describing our understanding of particle dark matter, some candidates and their properties.

¹to an astronomer, this is on the order of $10^{-2} \text{ M}_{\odot}\text{pc}^{-3}$, showing that the solar system does not contain enough dark matter to observe gravitational effects from it. The solar system contains much more normal matter than dark matter. Dark matter dominates only on larger scales.

Baryonic dark matter A very old scenario is dark matter in the form of Massive Compact Halo Objects (MACHOs), such as brown dwarfs, orphaned planets or black holes. This is an elegant hypothesis in the sense that it does not require new elementary particles. To avoid conflicts with the modern cosmological measurements of baryon and dark matter densities, one can postulate MACHOs in the form of Primordial Black Holes (PBH). These could be formed before Big Bang nucleosynthesis through a mechanism beyond the Standard Model, and thus do not contribute to the baryon budget of the early universe. This type of dark matter can populate a wide mass range. Its most important astronomical feature is that MACHO dark matter would be much “lumpier” on small scales than smoothly distributed particle dark matter. This allows a variety of techniques to place limits on MACHO dark matter.

In gravitational microlensing, a characteristic flare of a distant star is produced when a MACHO passes in front. By monitoring tens of millions of stars in the Magellanic clouds, limits on the number density of MACHOs have been placed, showing that MACHOs below $\sim 30 M_{\odot}$ cannot make up the galaxies dark matter halo [28, 29].

Heavier MACHOs can be constrained by their tendency to disrupt wide binary stars and destroy star clusters by dynamic heating. Studies of such systems allow to exclude dark matter composed entirely of MACHOs up to arbitrarily high masses [30].

Modified Gravity The evidence for dark matter comes from vastly different astronomical systems and using many independent strategies of inference. Still, they all ultimately concern the gravitational interaction of dark matter. It is therefore conceivable that a new, more complete theory of gravity can solve the “dark matter problem” without the need of dark matter. Such a modified gravity theory has to explain all the different discrepancies attributed to dark matter in diverse astrophysical contexts, while obeying all observational constraints confirming the validity of general relativity.

In 1983, Milgrom proposed a theory of Modified Newtonian Dynamics (MOND) [31]. Its basic postulate is the suppression of Newtonian forces below a tiny acceleration scale $a_0 \sim 1.2 \cdot 10^{-10} \text{ m/s}^2$. This simple modification is very successful in explaining the observed rotation curves of galaxies without missing mass, making MOND the most popular alternative to dark matter at the time.

By itself, MOND is not a complete theory and violates conservation of energy, momentum and angular momentum. Intense theoretical efforts have yielded relativistic theories of gravity which include MOND-like behaviour at low energies and reproduce the observed strength of gravitational lensing [32].

Work is continuing to find a particular theory of gravity that can also explain the CMB power spectrum and even lensing in the Bullet Cluster [33, 34].

The recent coincident observation of gravitational and electromagnetic waves from binary neutron star merger GW170817 [35] confirmed with high precision that gravitational waves travel with the speed of light. In addition, it showed that photons and gravitational waves are subject to an identical time delay due to intervening gravitational potentials (a special case of the weak equivalence principle). This presents an additional challenge for some classes of modified gravity theories [36].

Particle dark matter: known properties In spite of the hypothetical nature of a dark matter particle, a number of precise statements can be made about particle dark matter in general. In [37], the authors give a list of conditions a theoretical candidate particle

has to fulfill in order to be considered a good dark matter candidate. Mutatis mutandis, this list can be used as a summary of our knowledge on particle dark matter in general.

- **Relic density:** one of the most precise measurements about dark matter is its abundance in the early universe, measured in the CMB power spectrum to percent-level precision. Big Bang cosmology consistently explains the number density and behaviour of the particle species of the Standard Model, from Big Bang nucleosynthesis in the first few minutes to recombination after 380.000 years. An extended Standard Model containing extra particles should likewise reproduce the observed particle densities, along with the correct dark matter density. The generic initial conditions of the Big Bang have to lead to a natural production of the observed amount of dark matter. Any plausible dark matter model contains, along with the candidate particle, a production mechanism compatible with Big Bang cosmology. This production mechanism can be used to classify dark matter particle candidates into several categories, discussed below. In the same vein, as the gravitational influence of dark matter is still observed today, over 13 Gyr later, dark matter must be made of stable particles (or at least with a decay time much longer than the age of the universe). Big Bang nucleosynthesis is likewise an impressive probe of early-universe dark matter properties [38]. The observed abundances of light elements precisely match the predictions of the model and tell us the baryon-to-photon ratio and expansion speed of the universe minutes after the Big Bang. The model is very sensitive to the number of relativistic particle species in the primordial plasma, constraining some dark matter scenarios.
- **Density and velocity distribution in the Galaxy:** from the rotation curve of the Milky Way Galaxy, we broadly know the mass, shape and size of its dark matter halo. In the simplest model, this allows to extract the DM velocity distribution as an isotropic Maxwellian with velocity dispersion tied to the depth of the galactic potential well. This can be extracted from the observed rotation curve, the circular speed $v_c(R)$ as a function of galactic radius. The so-called Standard Halo Model gives a velocity distribution

$$f(\vec{v}) \propto \exp(-|v|^2/v_c(R)^2) \quad (1.1)$$

with a local circular speed of $v_c(R_0) = 220$ km/s [39]. Importantly, the particles of the halo must be gravitationally bound to the Galaxy, i.e. there is a cut-off velocity above which dark matter particles escape. The galactic escape speed is estimated by the RAVE survey [40] to be 544 km/s. The Standard Halo Model is frequently used by experiments to give comparable exclusion limits. Debate and further work is ongoing to explore its degree of accuracy. More realistic velocity distributions are studied in elaborate N-body simulations of galactic formation and evolution including dark matter and baryons, for reviews see [41, 42]. Moderate deviations from Maxwellian distributions are sometimes found, which will become important for deeper understanding of a high-statistics direct dark matter signal in the future. There is considerable uncertainty on the realistic density profile in the Galactic center (cusp vs. core problem [43, 44]), but this has little effect on the halo at the solar radius. More important for direct detection would be the presence of a “dark stream”, i.e. a non-equilibrated faster moving dark matter

component at the location of the sun. Detectability and impact of such streams on direct searches are discussed e.g. in [45].

- **Neutral and at most weakly interacting:** the simplest way to explain the invisibility of dark matter is to assume it is made of an electrically neutral particle. Charged particles couple to photons and thus emit light, and undergo Coulomb scattering and even chemical interactions with ordinary matter. All of these phenomena are highly constrained by observations, leaving little room for exotic charged dark matter models. A similar argument can be made regarding the strong force: color-charged dark matter would self-interact (similar to neutron-neutron scattering), dissipating energy in the halo and accumulating in the galactic center. Therefore, other than by gravity, dark matter is only allowed weak interactions with ordinary matter. “Weak” may stand for the electroweak force of the Standard Model, or some new interaction of similar or lesser strength.




Figure 1.4.: Cosmic large-scale structure simulated in the “Millennium Simulation”, image from [46]. Left: dark matter density (brightness) and velocity dispersion (color) 13.6 Gyr after the Big Bang. Right: projected mass density of galaxies. The clumps and filaments in the distribution of galaxies follow the underlying network of the dark matter distribution. “Warm dark matter” made of relativistic particle would free-stream to a greater extent and fail to produce the observed fine structures.

- **Cold dark matter:** the favored cosmological model, Λ CDM, explicitly names “cold dark matter” (as well as a cosmological constant Λ for accelerated expansion). Cold is an abbreviation for the more unwieldy term “non-relativistic during structure formation”. The growth of large-scale structure in the cosmos can be traced in observations of the “Lyman- α forest”: absorption features in the short-wavelength flank of hydrogen emission from a distant quasar trace the density of intervening hydrogen gas as a function of redshift (i.e. cosmic expansion history). The detailed structure of these absorption features follow the evolution of matter into clumps, filaments and subsequently clusters and galaxies. This data can be used to constrain the free-streaming scale of dark matter, which would wash out small-scale features as long as dark matter stays relativistic [47]. Coupled with unprecedentedly resolved computer simulations of cosmic structure formation (de-

scribed e.g. in [46], see Figure 1.4), “cold dark matter” emerges as the preferred scenario. This implies a lower mass bound on thermally produced dark matter particles around several keV. For this reason, for example the Standard Model neutrinos with masses below few tenths of an eV are excluded from being the dominant form of dark matter.




Figure 1.5.: Allowed mass ranges for some types of DM candidates, from [48].

Candidate particles There is a wide variety of frameworks describing potential dark matter particles which respect the above properties. These range from general mechanisms to fully fleshed-out theories. Figure 1.5 shows the staggering mass range over which these particles could occur. As the zoo of candidate dark matter particles is so diverse, the overview given here does not attempt to be complete. The list below should be taken as an outline of some of the most popular candidates.

Certainly the best-studied dark matter candidate is the **Weakly Interacting Massive Particle** (WIMP) produced by thermal freeze-out in the early universe. The energy density close to the Big Bang is high enough to produce all possible particles, which frequently undergo creation and annihilation reactions. Under these conditions, Standard Model particles and WIMPs are in thermal equilibrium through their common annihilation products. As the universe expands, the number densities of particles drops. While normal matter keeps interacting and cooling, at a certain point WIMPs with their weaker interactions no longer get produced and annihilated frequently enough to stay in equilibrium: they decouple and freeze out. The known relic density of dark matter tells us the thermally-averaged annihilation cross-section at which this happens: $\langle\sigma v\rangle \approx 3 \cdot 10^{-26} \text{ cm}^3\text{s}^{-1}$ [49]. The main motivation of the WIMP concept is that this number is the expected value for a “weak-scale” particle (i.e. mass from $1 - 10^5 \text{ GeV}$) with weak interactions ($\sigma \approx 10^{-36} \text{ cm}^2 = 1 \text{ pb}$). Such particles are believed to exist for reasons related to the internal consistency of the Standard Model of Particle Physics (the gauge hierarchy problem). Cosmology tells us that they would naturally be produced at the correct dark matter abundance in the early universe. This remarkable coincidence is known as the “WIMP miracle”. WIMPs with their weak interactions would in addition be detectable non-gravitationally. All this has motivated large experimental efforts in WIMP searches over the past decades. The term WIMP is used generically for a class of particles, for reviews see [50, 51]. The most prominent WIMP candidate historically has been the lightest supersymmetric particle (in many scenarios the neutralino, a Majorana fermion typically around 100 GeV mass). In many theories, conservation of a new quantum number produces a new stable particle as a DM candidate.

In the **asymmetric dark matter** [52] scenario, the dark matter particle is not its own antiparticle. Instead, a slight asymmetry during production (analogous to the baryon-antibaryon asymmetry) leads to a relic density after dark matter annihilation. Most generically, the particle mass relative to that of protons is related to the primordial

abundances, predicting $m_{dm} \approx 5 \text{ GeV}/c^2$.

If dark matter interaction cross-sections are far below the electroweak scale, dark matter particles never reach thermal equilibrium with Standard Model particles in the early universe. This scenario is called **Feebly Interacting Massive Particle (FIMP)** [53]. Instead of “freezing out” (annihilation stopping with expansion of the universe), these particles can “freeze in”: non-thermal production from SM particles is interrupted by the expansion of the universe as the correct relic density is reached.

The **Strongly Interacting Massive Particle (SIMP)** paradigm [54] describes dark matter with sizable self-interactions, albeit below the limits derived from the Bullet Cluster. Such dark matter models produce the relic density via number-changing $3 \rightarrow 2$ annihilation processes, leading to a lower mass estimate $\sim 100 \text{ MeV}$. Dark matter with self-interactions seems to be favored by tensions in structure formation below scales of few tens of kpc, known as the “missing satellite problem” [55, 56], “cusp vs. core problem” [57] and the “too big to fail problem” [58, 59].

Another dark matter candidate addressing structure formation problems are **sterile neutrinos** [60]. A wide range of mass-scales is possible for these particles “just outside the Standard Model”, with different motivations and cosmological implications. As a dark matter candidate, sterile neutrinos are viable only in the mass range of a few hundred eV to some tens of keV. From below, this range is limited by the Tremaine-Gunn bound [61] derived from the Pauli exclusion principle and the necessity to amass enough fermions in galactic centers to form a sufficiently heavy halo. On the high side, sterile neutrinos would rapidly decay into normal neutrinos and monoenergetic photons, which are not observed. Sterile neutrino dark matter can be produced thermally or athermally, by mixing with SM neutrinos, decay of heavier particles or new gauge interactions. They can be cold or slightly warm dark matter, or not have a thermal spectrum at all.

Very light dark matter can still be cold if it is not produced thermally, e.g. by the misalignment mechanism [62, 63, 64]. In such a scenario, a new field would be brought to a random initial state in the early universe (for example by inflation). Such fields would subsequently start oscillating around their minimum. The coherent field oscillations can behave just like cold dark matter. This production mechanism applies to **QCD axions** [65, 66, 67], **axion-like particles (ALPs)** [68] and **hidden photons** [69], collectively known as WISPs (Weakly Interacting Slim Particles) [70]. As bosons, their densities are not limited by Pauli exclusion. The mass of bosonic dark matter must only be larger than 10^{-22} eV , where the wavelength becomes comparable to the size of galactic dark matter halos.

1.1.2. Experimental approaches

Assuming a non-gravitational interaction between dark matter and Standard Model particles, the symmetries of quantum field theory open up three avenues to search for signatures of this interaction. This is often illustrated with a Feynman diagram (see Figure 1.6) showing how the interaction can result in three fundamentally different observable processes: dark matter scattering of SM particles in **direct detection**, dark matter annihilating into SM particles in **indirect detection**, and production of dark matter in SM particle collisions at **collider searches**.

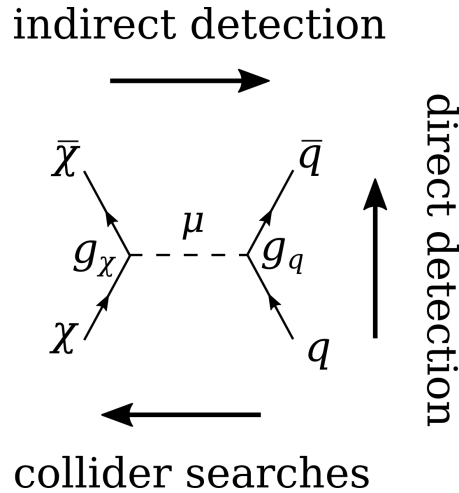


Figure 1.6.: Three avenues for dark matter search, illustrated in a schematical Feynman graph to be read in three orientations. Assuming some fundamental interaction exists between dark matter particles χ and Standard Model particles q (here coupled by a mediator μ), three different processes are allowed. Indirect detection searches for dark matter particles annihilating into SM particles. Collider searches aim at producing dark matter particles in collisions of SM particles. Direct detection strives to observe dark matter particles scattering elastically off SM particles.

1.1.2.1. Collider Searches

At high-energy particle colliders, dark matter particles can be produced if the dark matter mass is within the range of the machine energy. The signature of dark matter production is inconspicuous. The chance of a dark matter particle to interact again on the way out of an experiment is negligible. Thus, dark matter is invisible at a collider experiment. Therefore the observable of such an event consists only of an imbalance of the momenta of reconstructed particles, known as missing transverse momentum (colloquially: missing energy).

At the Large Hadron Collider, many kinds of dark matter searches are actively pursued (for a review see [72]). Many of them are searches for specific extensions of the Standard Model which also include a dark matter candidate, such as supersymmetry. These complete models often predict very specific final states, such as a number of high-momentum SM particles associated with missing energy. Conversely, there is a nearly uncountable number of possible searches for such very specific models, motivating more general approaches.

More agnostic dark matter searches target dark matter particle pair production. These events do not produce a visible signature, unless an SM particle is additionally emitted from the initial or final state. One therefore looks for these subdominant events, for examples called mono-jet, mono-V (vector boson: γ , W^\pm , Z_0) or mono-Higgs searches. Backgrounds include detector inefficiencies, wrongly reconstructed particles, and Standard Model processes such as a Z-boson decaying invisibly to neutrinos. The expected missing energy spectrum from this particular background can be estimated from the corresponding leptonic Z-decays. Dark matter coupling predominantly to heavy particles




Figure 1.7.: ATLAS limits on a simplified model of WIMP pair production from an axial-vector mediator [71]. Left: WIMP mass versus mediator mass. Right: comparison with spin-dependent direct detection limits. To translate the simplified model observables to the WIMP-proton cross-section, model-dependent choices have to be made for the mediator couplings to quarks and dark matter and the mediator decay width.

are searched in top quark pair production with missing energy. Dark matter lighter than half the Higgs boson mass could cause invisible Higgs decays, which are studied using a number of different signatures.

Given a search channel (i.e. event class potentially caused by dark matter) and an understanding of expected backgrounds (from SM processes and detector effects mimicking the signal), one still needs to specify a dark matter model to make a statement on the confirmed or excluded dark matter properties. Due to the unknown nature of dark matter and the vast theoretical possibilities, a compromise between generality and realism is needed.

The minimalist strategy, known as Effective Field Theory approach [73, 74, 75], adds only a single DM particle to the Standard Model, and an effective contact interaction mediated by an unspecified particle heavier than the LHC energy scale. Here, simplicity is given preference over internal consistency and motivation by general theoretical principles.

One step more detailed are so-called simplified models [76, 77], in which also the mediator of the DM-SM interaction is kinematically accessible at LHC. Particularly interesting classes of simplified models have been identified [78, 79] and studied at both ATLAS and CMS [71, 80]. These models can be required to be normalizable, and necessarily have more free parameters than the EFT approach.

For example, an s-channel simplified model is specified by the interaction structure (axial-vector, vector, scalar or pseudoscalar) and five parameters (dark matter and mediator masses, couplings of the mediator to dark matter and SM particles, and width of the mediator). Commonly, the couplings are fixed at some benchmark values, and the width is assumed to be minimal (mediator does not decay into other light particles). The sensitivity of a collider search is then presented in the remaining mass-mass plane

(see Figure 1.7).

The idea of the simplified model approach is to capture the relevant phenomenology of many classes of underlying complete theories, so that experimental constraints on simplified models can subsequently be translated into constraints on them. The detailed mapping process can be challenging, and at the same time questions remain on realism (respecting unitarity and gauge-invariance) and generality (ignoring relations between parameters present in complete models) of some simplified models.

1.1.2.2. Indirect detection

The field of indirect detection searches for standard model particles in the universe generated by annihilation or decay of particle dark matter. A recent review can be found in [81]. The messengers of dark matter can take a number of different forms, such as photons, neutrinos, electrons, positrons, protons, anti-protons and heavier nuclei or antinuclei. The energy of the messenger is related to the unknown dark matter mass. Searches range from keV X-ray lines to cosmic rays and neutrinos of the highest energies (PeV-EeV).

Thermal relics are a popular model to probe by indirect detection, as the flux of annihilation products is directly proportional to the velocity-averaged cross-section, which is known for thermal production in the early universe. Also sterile neutrinos leave a conspicuous astronomical fingerprint, an X-ray line at the neutrino mass-energy, through their radiative decay into SM neutrinos. Another model that lends itself to indirect testing is superheavy dark matter [82], whose particle would be extremely rare but could decay or annihilate into ultra-high energy cosmic rays.

Gamma rays are observed by the Fermi satellite [83] (20 MeV-300 GeV) and at even higher energies by ground-based Cherenkov telescope arrays (such as HESS, VERITAS, MAGIC and the planned CTA [84]). Neutrinos are detected in the IceCube, ANTARES and BAIKAL-GVD neutrino telescopes, as well in the water Cherenkov detector Super-Kamiokande. Cosmic rays are detected on-orbit by the PAMELA and AMS experiments, and on the ground in the Pierre Auger Observatory, an array of water Cherenkov detectors spanning 3000 km².

As annihilation and decay rates of dark matter are largest where dark matter density is highest, indirect searches are typically aimed at dense regions of the universe. Examples are the Galactic center, the Milky Way halo, satellite or other nearby galaxies, and galaxy clusters. Neutrino telescopes can in addition target the Sun and the Earth, which may gravitationally trap dark matter. Cosmic ray experiments can only study the local flux without information on origin (due to deflection of charged particles by galactic magnetic fields).

A number of tantalizing hints have been found in the past decade, such as an excess of 1-10 GeV photons from the Galactic center [85], and a rising positron-to-electron ratio above 10 GeV in the local cosmic ray flux [86]. While both signals can be explained by dark matter annihilation, it is difficult to exclude astrophysical origins, such as unresolved millisecond pulsars to explain the Galactic excess, and nearby pulsars for the positron fraction [87]. Better data from upcoming generations of experiments may be able to distinguish the different proposed scenarios.




Figure 1.8.: Nuclear recoil energy spectra for spin-independent interaction of dark matter particles with a mass of (1, 10, 100) GeV/c^2 with a CaWO_4 target. Dotted lines show the contribution from oxygen recoils, dashed lines show tungsten recoils. The differential rate $\frac{dR}{dT}$ is normalized to an interaction cross-section of $\sigma_n = 1$ pb.

1.1.2.3. Direct detection

Direct detection experiments try to observe energy depositions in a target caused by interactions of dark matter particles traversing the Earth. The principle of this approach was first discussed in 1984 [88], and quickly followed by the first experimental constraint [89]. For generic WIMP models, the most prominent signature of dark matter interactions are coherent elastic nuclear recoils. The energy transfer T in non-relativistic two-body scattering is [39]:

$$T = \frac{q^2}{2m_N} = \frac{\mu^2 v^2}{m_N} \cdot (1 - \cos \theta) \quad (1.2)$$

with momentum transfer q , reduced mass $\mu = \frac{m_\chi m_N}{m_\chi + m_N}$, nuclear mass m_N , dark matter mass m_χ , relative velocity v , and center-of-mass scattering angle θ . The first fundamental challenge of direct detection results directly from this: with typical values $m_N = 100 \text{ GeV}/c^2$ and $v = 200 \text{ km/s}$, the maximum recoil energy (in backscattering $\theta = \pi$) for $m_\chi = 100 \text{ GeV}/c^2$ is only 22 keV. For $m_\chi = 10 \text{ GeV}/c^2$, this reduces to 0.7 keV. Therefore, any detector used in direct detection must be sensitive to keV-scale energy depositions.

Similarly, the interaction rate R can be estimated as the product of target number n , dark matter particle flux Φ and interaction cross-section σ :

$$R = n \cdot \Phi \cdot \sigma = \frac{M}{m_N} \cdot \frac{\rho_\chi \cdot \langle v \rangle}{m_\chi} \cdot \sigma \quad (1.3)$$

With target mass $M = 1$ kg, dark matter density $\rho_\chi = 0.3$ GeV/c²/cm³, average incident velocity $\langle v \rangle = 200$ km/s and cross-section $\sigma = 1$ pb = 10^{-36} cm² one finds $R = 0.09$ /day, a single event every eleven days. This shows the second fundamental challenge of direct detection: a successful experiment must combine a large target with an extremely low background rate. The quoted numbers serve for illustration only: pb cross-sections are excluded over a wide mass-range, much lower rates are currently being probed.

Most direct detection experiments precisely measure the size of energy depositions in the target. The expected energy spectrum (differential rate) of dark matter recoils can be calculated as [90]:

$$\frac{dR}{dT}(T) = \frac{\rho_\chi M}{m_N m_\chi} \int_{v_{min}}^{v_{esc}} dv \quad v \cdot f(v) \cdot \frac{d\sigma}{dT}. \quad (1.4)$$

Here, $\frac{d\sigma}{dT}$ is the differential interaction cross-section between dark matter particles and target nuclei, and $f(v)$ denotes the velocity distribution of dark matter particles in the laboratory frame. The velocity integral has particular limits: particles slower than $v_{min} = \sqrt{m_N T / 2\mu^2}$ cannot induce recoils of energy T and particles with a speed higher than v_{esc} are absent as they are not bound to the galactic potential well.

The differential cross-section is an unknown function and takes many forms in different dark matter models. A useful framework to classify the options is non-relativistic effective field theory [91]: independent from the fundamental theory describing dark matter, by Galilei invariance only a limited number of operators can describe the non-relativistic scattering of dark matter particles. In many models, one of only two operators dominates the interaction. They are known as “spin-independent” (SI) and “spin-dependent” (SD) interactions. In the SI case (induced by scalar or vector four-fermion interactions), the dark matter particle interacts identically and coherently with protons and neutrons in the target nucleus. This leads to an enhancement of the cross-section compared to DM-nucleon scattering:

$$\frac{d\sigma}{dT}_{SI} = \frac{m_N}{2v^2} \cdot \frac{\sigma_n}{\mu_n^2} \cdot A^2 \cdot F^2(T) \quad (1.5)$$

σ_n is the underlying DM-nucleon cross-section and μ_n the DM-nucleon reduced mass. The form-factor $F^2(T)$ describes the loss of coherence at large transferred momenta: it is the Fourier transform of the nuclear mass density, equal to unity at $T = 0$ and falling as the wavelength of the exchanged momentum drops below the nuclear size.

In the SD case (induced by axial-vector four-fermi interactions), the dark matter couples only to the net spin of the target nucleus. The A^2 enhancement factor is absent, and only nuclei with non-zero spin (such as ¹H, ⁶Li, ⁷Li, ¹⁹F, ⁷³Ge, ¹²⁹Xe, ¹³¹Xe, ¹⁸³W) have any sensitivity. Instead of a single DM-nucleon cross-section, different couplings $a_{n,p}$ to neutron and proton spins are possible and separately constrained.

These assumptions on the interaction of dark matter with nuclei allow calculating an expected recoil spectrum as a function of just two numbers (besides the parameters of the Standard Halo Model): dark matter mass and cross-section. Some exemplary recoil spectra for the spin-independent case are shown in Figure 1.8. Together with an experimentally observed spectrum, the modelled recoil spectra for different sets of dark matter mass and cross-section allow placing limits on the parameter space as well as finding mass and cross-section corresponding to an observed signal. Various statistical tools are used to this end, such as Maximum Likelihood fits [94] or the Yellin Optimum




Figure 1.9.: Limits on spin-independent interactions of dark matter particles with nuclei, from [92]. For references to individual experimental results, see [92] and references therein.

Interval method [95]. Selecting the DM-nucleon cross-section as variable of interest as shown above for the SI case allows comparing the sensitivities of experiments with different nuclear targets. This comparison is only possible under the outlined assumptions regarding the DM-nucleus coupling and the halo velocity distribution. Figure 1.9 shows the limits on spin-independent DM interactions from the first release of CRESST-III data [92]. Similar plots for spin-dependent couplings to protons and neutrons from [93] are shown in Figure 1.10. In an analogous way, limits on scattering of light dark matter with electrons can be derived for experiments with sensitivity to single charges (see Figure 1.11 from [8]).

There is a well-predicted astrophysical background to direct searches for nuclear recoils: coherent elastic scattering of solar and atmospheric neutrinos. The $\text{CE}\nu\text{NS}$ process from Standard Model neutrinos (discussed in section 1.2) shares many features with the hypothetical dark matter scattering and produced an indistinguishable experimental signature. It therefore acts as an irreducible background affecting the sensitivity of future direct search experiments [96]. The parameter space in which a dark matter observation will be complicated by detection of $\text{CE}\nu\text{NS}$ events is shaded gray in Figure 1.10, as calculated in [97].

Several features still distinguish a dark matter signal from the $\text{CE}\nu\text{NS}$ background. One such is the “annual modulation”, for a recent review see [98]. The velocity of the Earth with respect to the Galactic DM halo changes over the year, due to the Earth’s rotation around the Sun which is added to the solar motion around the galactic center.




Figure 1.10.: Limits on spin-dependent interactions of dark matter particles with nuclei, from [93]. For references to individual experimental results, see [93] and references therein.

The effect is quite small, with only about half of Earth’s 30 km/s orbital velocity aligned with the solar motion of around 220 km/s. This can nonetheless have important effects on the high-energy tail of DM scattering events, so that the events detected above an energy threshold may be strongly modulated. This modulation, if detected with high confidence, is an indicator of the Galactic origin of the signal. Many backgrounds also exhibit an annual modulation (atmospheric muons, neutrinos and cosmogenic neutrons due to seasonal density changes in the atmosphere [99], solar neutrinos due to the shifting Sun-Earth distance [100]) albeit with a different amplitude and phase. The only signal claim from a modulation experiment, the long-standing result of the DAMA collaboration [101], remains controversial as it is incompatible with many other experiments in the standard interpretation.

Another distinct feature of a dark matter signal is its directionality, for a review see [102]. Due to the solar motion around the Galactic center, dark matter particles in the Earth reference frame appear to be streaming in a “dark matter wind” from the direction of the constellation Cygnus. (The annual modulation discussed above corresponds to a small variation in direction and average speed of this wind.) In consequence, nuclear recoils induced by dark matter show a forward-backward asymmetry pointing away from Cygnus. Detectors with the ability to reconstruct nuclear recoil directions thus have the power to statistically disentangle a Galactic signal from omnidirectional environmental




Figure 1.11.: Limits on scattering of MeV/c^2 -scale dark matter scattering with electrons, from [8]. For references to individual experimental results, see [8] and references therein.

background or even solar neutrino background. In addition, the directional dark matter signal may carry information on details of the dark matter coupling and the local velocity distribution.

1.1.2.4. Comparability and complementarity

In principle, all three ways to observe non-gravitational effects of dark matter are a logical consequence of any underlying interaction. In practice, the three approaches are far from equivalent. To start with, indirect detection by annihilation requires two dark matter particles to produce an interaction. Direct detection requires one, and a collider search needs none at all. As a consequence, direct and indirect event rates scale with dark matter density ρ and ρ^2 respectively. This implies astrophysical uncertainties on indirect and to a lesser extent on direct searches. While the sensitivity of direct searches depends on the local dark matter density, an indirect search relies on an understanding of the dense astrophysical region it probes (e.g. the Galactic center). The sensitivity of a collider search is independent of the dark matter density and scales instead with the collider integrated luminosity.

Also, the three approaches probe vastly different energy scales: in a collider, the momentum transfer is set by the center-of-mass energy delivered by the machine. The transferred momentum μ in dark matter annihilation is about twice the dark matter mass. In direct detection, dark matter undergoes non-relativistic scattering with $\mu \sim \mathcal{O}(1 - 100 \text{ MeV})$. The strength of the interaction of dark with normal matter at each of the energy scales depends on the specific dark matter model.

The kinematics of the process is also different in the three cases. Indirect detection measures the thermally-averaged annihilation cross-section that may be responsible for the relic density in case of a thermal production mechanism. Collider searches can probe

dark matter masses up to the machine energy, heavier dark matter is inaccessible. For direct detection, heavy dark matter causes larger recoil energies, while light dark matter becomes invisible due to the detector energy threshold.

Production at a collider is the only technique where interactions of the new particle with the different SM partners can be systematically probed. This produces unique insight into the physics of the dark matter particle. Conversely, a collider does not allow to ascertain the “dark matter nature”, i.e. cosmological stability and abundance in the universe.

These examples show that there is no approach preferable a-priori for detecting dark matter. A single clear detection in either channel will not mark the end of the quest to understand the nature of dark matter. It will be only the beginning.

1.1.3. Technology for direct detection

As introduced above in section 1.1.2.3, competitive dark matter detectors must fulfill a number of requirements. They need a low energy threshold to be sensitive to keV-scale energy depositions. The detector target should be scalable to large masses to constrain increasingly rare interactions. At the same time, the detector must be constructed out of low-radioactivity materials to keep background levels low. Further methods of event discrimination are desirable to separate dark matter interactions from residual background events. Various schemes are employed, such as veto detectors (against muons or muon-induced interactions), two-channel detection (allowing nuclear- vs. electron recoil discrimination), fiducialisation (to distinguish events in the bulk of the detector from interactions near contaminated surfaces) or information on the direction of the interacting particle.

In solid, liquid and gaseous detector materials, nuclear recoils lead to various excitations that can be detected with suitable sensor technology. The most common ones are electric charges (electrons and holes) registered in charge-sensitive amplifiers and photons that can be counted using photomultiplier tubes or solid-state photodetectors. Only a small fraction of the nuclear recoil energy creates light or charge signals. For electron recoils this fraction is larger (the nuclear recoil signal is said to be “quenched”), which causes complications in calibration of the nuclear recoil energy scale. The ionization or light output depends on the ionization density of the primary event, which is different for electron and nuclear recoils. In either case, most of the deposited energy goes into thermal motion in the detector material. For solid-state targets, this corresponds to lattice vibrations (phonons), which become detectable at cryogenic temperatures.

A plethora of technological implementations using these detection channels is applied in ongoing direct dark matter searches. The following is intended as an overview of technologies, not an exhaustive list.

1.1.3.1. Semiconducting/scintillating crystals

Semiconductors with their eV-scale bandgaps allow creation of many electron-hole pairs in low-energy nuclear recoils, and a correspondingly good energy resolution. Germanium crystals in particular have a long history in direct detection, with the historically first limit obtained with a 0.72 kg high-purity germanium crystal with an energy threshold of 15 keV for nuclear recoils [89]. Modern experiments use point-contact germanium diodes and reach thresholds of 160 eV for electron recoils [103]. This corresponds to a nuclear

recoil threshold around 0.8-1.3 keV due to quenching, which is subject to considerable uncertainty at these low energies [104]. The size of individual charge detectors is limited to kg-scale by electronic noise related to the detector capacity. Tonne-scale arrays are challenging to construct due to the complexity of the individual detector. Silicon CCD detectors [105] offer high charge resolution, a lighter target nucleus, and imaging possibilities for background identification. Non-destructive readout techniques [106] allow detection of single charges. Still, large arrays are necessary to even reach a kg-scale target mass.

Scintillation detectors overcome this challenge using a simpler detector design. In a scintillator, charge recombination after an energy deposition creates photons which escape the transparent material. The photons can be detected by coupling scintillator crystals to low-background photon detectors. A famous example of this technique is the DAMA experiment [101] with its long-standing modulation signal claim, operating a 250 kg array of NaI scintillators.

1.1.3.2. Cryogenic detectors

A challenge for both semiconductors and scintillators is their inability to distinguish nuclear from electron recoils. In addition to raising doubts about the nature of a possible signal, this causes uncertainties in the energy scale, which is usually calibrated by an electron-recoil signal and has to be converted using uncertain quenching factors. Cryogenic detectors address this challenge by measuring heat signals, which are nearly independent of the interacting particle type. By using cryogenic semiconductor or scintillator crystals, a heat and a light or charge signal can be recorded simultaneously for each energy deposition. With an event-by-event measurement of light or charge yield, electron and nuclear recoil events can be distinguished. Thermal or athermal phonon signals can be detected using Neutron Transmutation Doped germanium thermistors (NTDs) [107] or tungsten thin-film Transition Edge Sensors (TESs) [108, 109]. Ongoing projects using two-channel cryogenic detectors are EDELWEISS [110] (NTDs and charge readout), SuperCDMS SNOLAB [111] (W-TESs and charge readout) and CRESST-III [92] (W-TES on both heat and light channels). Event-by-event particle identification, ultra-low energy thresholds (30 eV in CRESST [92]) and a wide range of possible target materials are advantages of cryogenic detectors. A disadvantage is the cumbersome cryogenic operation as well as the smallness and complexity of the detector elements posing difficulty to scale to large target masses. In consequence, cryogenic detectors lead the search for light dark matter (below few GeV/c^2 , where expected recoils are more frequent but less energetic), but cannot compete with other approaches at higher masses.

1.1.3.3. Noble liquid detectors

Experiments with a noble liquid target (usually xenon or argon) are currently the most sensitive to the classic WIMP of 10-1000 GeV/c^2 mass. Dual phase time-projection chambers combine the readout of light and charge in PMT arrays facing the liquid target. Prompt scintillation light is directly registered in the PMTs, charges are drifted to the liquid surface in a moderate electric field, and then amplified in the gas phase by a much higher applied field. This leads to secondary scintillation which is detected mostly in the top PMT array. This approach allows both particle identification and event localisation in the detector, which is a powerful handle on surface and external backgrounds. The liquid

target can be continuously circulated and purified, allowing unprecedented background levels combined with the self-shielding effect of the target (82 events/tonne/year/keV_{ee} for electron recoils in XENON-1T [112]). The technology allows scaling to large masses. Over the past decade, a series of detectors using liquid xenon has reached tonne-scale targets. The next-generation experiments XENON-nT [113] (5.9 t target mass) and LUX-ZEPLIN [114] (7 t target mass) are expected to observe a few CE ν NS background events, touching for the first time the neutrino floor. Argon-based TPCs are currently smaller (46 kg in DarkSide-50 [115]), but aiming to scale even faster (23 t in DarkSide-20k [116]). Argon features a lighter nucleus, and an additional handle on nuclear/electron recoil discrimination using the pulse-shape (fast and slow scintillation component in liquid argon). Challenges include a lower light yield compared to xenon, and the presence of the radioactive isotope ^{39}Ar which has to be removed from natural argon.

1.1.3.4. Bubble chambers

A completely different type of excitation is exploited by bubble chambers: a liquid is brought into a superheated state by heating and compression. A sufficiently dense energy deposition can cause nucleation of a bubble, which can be detected acoustically and photographically. The bubble is then removed by compressing and expanding the fluid. Such a detector does not measure the deposited energy, but counts only events above a thermodynamical threshold. Operation parameters can be tuned such that electron recoils do not nucleate bubbles due to their lower ionization density. Backgrounds from α decay can be discriminated by their acoustical signature. Neutron backgrounds can be estimated from multiple-scattering events. An advantage of bubble chambers is the variety of possible liquid targets, often containing ^{19}F for SI or iodine for SD sensitivity. A detector with 50 kg target mass of C_3F_8 and 3.3 keV energy threshold currently sets the most stringent limit on spin-dependent DM-proton coupling [117].

1.1.3.5. Directional detectors

A variety of technologies has been proposed for directional detection of dark matter, see the review in [102]. The most advanced experiments are low-pressure gaseous TPCs, profiting from long track lengths in the low-density medium. The most sensitive experiment is DRIFT [118], operating a m³-scale TPC filled with a $\text{CS}_2/\text{CF}_4/\text{O}_2$ mixture at 41 torr pressure. The total target mass corresponds to 140 g of F nuclei. Track directionality was measured down to an energy threshold of several tens of keV for electron recoils. Although the directional signal provides a smoking-gun evidence for the dark matter nature of a signal, target masses have to be increased by many orders of magnitude for these experiments to become competitive with energy-only direct searches.

1.2. Coherent elastic neutrino-nucleus scattering

The hypothetical spin-independent, coherent interaction of dark matter particles with nuclei has a fascinating sibling within the standard model: the low-energy weak neutral current interaction of neutrinos with nuclei. It is a long-standing prediction of the standard model, which has been observed experimentally only recently. The study of this process opens up many connections to different areas of physics.

1.2.1. Prediction and significance

The elusive neutrino is often introduced with the vivid example that they are sufficiently weakly interacting to easily penetrate a hundred light-years of steel [119] or even a thousand light-years of lead [120]. This may be true in most contexts, but in special circumstances neutrinos can behave in stranger ways. At very high energies, neutrino interaction cross-sections grow large: at $E_\nu > 40$ TeV neutrino absorption in the Earth has been measured using the IceCube detector [121]. At very low neutrino energies $E_\nu \sim 10$ MeV, the trapping of neutrinos in dense material during stellar collapse (with a mean free path of only few kilometers) is an important mechanism for supernova explosions [122, 123]. The process responsible for this second phenomenon, coherent elastic neutrino-nucleus scattering, and its detection on Earth, is the subject of this chapter.

The idea of coherent elastic neutrino-nucleus scattering was formulated by D. Freedman in 1974 [124], soon after experimental results from the Gargamelle experiment at CERN [125, 126] first established the existence of a weak neutral current. Freedman called his suggestion to search experimentally for this interaction “an act of hubris” [124]. This proved not to be an understatement, given the more than 40 years of technological development that were necessary before the first detection of $\text{CE}\nu\text{NS}^2$ was made by the COHERENT collaboration in 2017 [127].

Already in the first paper ever about $\text{CE}\nu\text{NS}$, Freedman pointed out boon and bane of its detection: while the cross-section is enhanced by $\sim N^2$ through the coherent action of all nucleons, the only observable signature is a low-energy nuclear recoil³. These critical features make the research of $\text{CE}\nu\text{NS}$ both challenging and compelling.

There is a common origin for both the enhanced interaction rate and the difficulty of observing its signature: simply put, the nucleus as a neutrino target is much larger than an elementary particle (increasing the rate), but also much heavier (reducing the maximal recoil energy). A quantitative estimate of this can be done in a back-of-the-envelope manner. The exchanged momentum q between neutrino and nucleus is associated with a length-scale h/q which must be a few times larger than the nucleus for the interaction to receive full coherent enhancement. For nuclear radii R around 5 fm, this limits momentum exchange to order $h/R \sim 100$ MeV/c. The exchanged momentum is related to the observable recoil energy T by $q^2 = 2MT$. For a nuclear mass of 100 GeV/c², this translates to recoil energies of only tens of keV. Using the relation between neutrino energy

²The acronym established for coherent elastic neutrino-nucleus scattering established by the COHERENT collaboration is pronounced “sevens” and provides a unique identifier after previous use of a number of different abbreviations.

³Initially, an enhancement $\propto A^2$ was expected. Due to the structure of the weak neutral current and the value of the weak mixing angle, the proton weak charge is strongly suppressed with respect to the neutron weak charge, so that the enhancement instead follows N^2 .

E_ν and maximum recoil T_{max} , $T_{max} = 2E_\nu^2/M$, we find that neutrinos predominantly producing coherently enhanced interactions should have an energy of tens of MeV or below.

In exchange for this challenge on the detector sensitivity, an experiment gains 3-4 orders of magnitude in interaction rate per target atom, allowing for a vastly reduced target size compared to traditional neutrino detectors. Conversely, a large CE ν NS experiment can reach high neutrino statistics enabling precision measurements.

Since its first prediction, there has been cross-fertilization between the fields of direct detection of dark matter and the search for CE ν NS. Traditional dark matter searches look for coherent nuclear scattering of a particle with unknown mass and coupling, but identical experimental signature as CE ν NS: a low-energy nuclear recoil. With detector technology improving and a first detection becoming plausible over the past decade, many efforts have started around the world with the goal of a CE ν NS observation.

With higher neutrino energies such as provided by stopped-pion sources, there is partial coherence in the nuclear interaction. This allows to directly measure the neutron distribution in the target nuclei for the first time. This has implications for other physics fields, such as atomic parity violation and the neutron star equation-of-state.

At lower neutrino energies, e.g. provided by the beta-decay processes at nuclear reactors, there is full coherence and thus no dependence of the cross-section on nuclear parameters. This allows precision measurements of the CE ν NS cross-section and neutrino properties. A study of such an experiment using cryogenic detectors in the main topic of this chapter.

At even smaller neutrino energies, coherent scattering of solar neutrinos can be observed. This poses an irreducible background to dark matter searches and warrants a precise independent measurement of CE ν NS cross-sections, to account for this background.

Similarly, large low-threshold dark matter experiments become sensitive to CE ν NS induced by a galactic supernova neutrino burst. Such a detection proceeding through neutral-current CE ν NS provides important information on all neutrino flavors, complementary to existing large neutrino experiments sensitive to charged current interactions.

Beyond all these physics applications, the study of CE ν NS is motivated by a unique technical possibility. The enhanced cross-section of CE ν NS enables small detectors. The best example is the 14 kg CsI crystal employed by COHERENT in the first observation. This allows for the first time transportable neutrino detectors that operate outside of underground laboratories. This technology may find application in the monitoring of nuclear reactors.

1.2.2. Cross-section in the Standard Model

The purpose of this section is to give sufficient detail on the cross-section of CE ν NS to allow the detailed calculations and sensitivity studies in chapter 7.

With a modern understanding of electroweak theory, a more detailed SM expectation for the CE ν NS differential cross-section can be specified [128]:

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{2\pi} \left\{ (G_V + G_A)^2 + (G_V - G_A)^2 \left(1 - \frac{T}{E_\nu}\right)^2 - (G_V^2 - G_A^2) \frac{MT}{E_\nu^2} \right\} \quad (1.6)$$

where T is the kinetic energy of the nuclear recoil and M the nuclear mass. The vector

and axial nuclear weak charges $G_{V,A}$ are related to the nucleus' constituents (through neutron and proton numbers) and their distributions (through the vector and axial form factors $F_{V,A}$):

$$G_V = [g_V^p \cdot Z + g_V^n \cdot N] \cdot F_V(Q^2) \quad (1.7)$$

$$G_A = [g_A^p \cdot (Z_\uparrow - Z_\downarrow) + g_A^n \cdot (N_\uparrow - N_\downarrow)] \cdot F_A(Q^2) \quad (1.8)$$

The coupling constants are the weak vector and axial neutrino-nucleon couplings $g_{V,A}^{n,p}$. The axial current couples to the imbalance of spin up/down nucleons, while the vector current couples to the total nucleon number. From this we can conclude that G_A vanishes for even-even nuclei. For other nuclei the ratio of vector to axial charge is on the order of A , so for large nuclei the axial contribution can be neglected to first order. This assumption leads to a simplified cross-section:

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{2\pi} G_V^2 \left\{ 1 + \left(1 - \frac{T}{E_\nu}\right)^2 - \frac{MT}{E_\nu^2} \right\} \quad (1.9)$$

The kinematics of the remaining terms leads to a further simplification. The form-factor enforces the coherence requirement, i.e. only neutrino energies below some tens of MeV contribute, therefore we have $E_\nu \ll M$ in any target nucleus. Since the maximum recoil energy obeys

$$T_{max} = \frac{2E_\nu^2}{Mc^2 + 2E_\nu} \approx \frac{2E_\nu^2}{Mc^2} \ll E_\nu \quad (1.10)$$

the middle term related to $(G_V - G_A)^2$ in Eq. 1.9 is very close to 1. Finally we can write

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{\pi} G_V^2 \left\{ 1 - \frac{MT}{2E_\nu^2} \right\} \quad (1.11)$$

or, switching to experimental units and making the dependence on the nuclear target explicit:

$$\frac{d\sigma}{dT} = \frac{1}{\pi} \frac{G_F^2}{(\hbar c)^4} \cdot Mc^2 \cdot [g_V^p \cdot Z + g_V^n \cdot N]^2 \cdot F_V^2(T) \cdot \left\{ 1 - \frac{Mc^2 T}{2E_\nu^2} \right\} \quad (1.12)$$

The 2018 Review of Particle Physics [129] gives the constants

$$\frac{G_F}{(\hbar c)^3} = 1.1663787(6) \cdot 10^{-5} \text{ GeV}^{-2} \quad (1.13)$$

$$\hbar c = 197.3269788(12) \text{ MeV} \cdot \text{fm} \quad (1.14)$$

leading to the numerical value

$$\frac{G_F^2}{(\hbar c)^4} \approx 5.297270 \cdot 10^{-60} \text{ m}^2/\text{eV}^2 \quad (1.15)$$

with an uncertainty below 1 ppm. The unusual choice of units is convenient for later multiplication with a neutrino flux. The neutrino-nucleon couplings can be related to the SM neutrino-quark couplings [129] of

$$g_L^u = +\frac{1}{2} - \frac{2}{3} \sin^2 \theta_W \quad (1.16)$$

$$g_R^u = -\frac{2}{3} \sin^2 \theta_W \quad (1.17)$$

$$g_L^d = -\frac{1}{2} + \frac{1}{3} \sin^2 \theta_W \quad (1.18)$$

$$g_R^d = +\frac{1}{3} \sin^2 \theta_W, \quad (1.19)$$

which are given by combinations of the quarks' weak isospin and their electric charge rotated by the weak mixing angle θ_W . These couplings can be combined into vector and axial combinations $g_{V,A} = g_L \pm g_R$ and summed over the quark contents of the nucleons:

$$g_V^p = 2g_L^u + 2g_R^u + g_L^d + g_R^d = +\frac{1}{2} - 2 \sin^2 \theta_W \quad (1.20)$$

$$g_V^n = g_L^u + g_R^u + 2g_L^d + 2g_R^d = -\frac{1}{2} \quad (1.21)$$

$$g_A^p = 2g_L^u - 2g_R^u + g_L^d - g_R^d = +\frac{1}{2} \quad (1.22)$$

$$g_A^n = g_L^u - g_R^u + 2g_L^d - 2g_R^d = -\frac{1}{2}. \quad (1.23)$$

These values, following directly from the underlying symmetry of the electroweak field theory, are the exact description for the weak neutral current at tree level. For a comparison with experiment high-order terms in the form of radiative correction have to be included. Table 10.3 in [129] gives values for the effective low-energy neutrino-quark couplings:

$$g_L^u = +0.3458 \quad (1.24)$$

$$g_R^u = -0.4288 \quad (1.25)$$

$$g_L^d = -0.1552 \quad (1.26)$$

$$g_R^d = +0.0777, \quad (1.27)$$

which can be translated to the neutrino-nucleon couplings

$$g_V^p = +0.0301 \quad (1.28)$$

$$g_V^n = -0.5116 \quad (1.29)$$

$$g_A^p = +0.4955 \quad (1.30)$$

$$g_A^n = -0.5120. \quad (1.31)$$

The dominant contribution to CE ν NS is rooted in the vector coupling to neutrons g_V^n . The table shows the big relative suppression of g_V^p due to the structure of the weak interaction, namely $\sin^2 \theta_W \approx 1/4$. As discussed above, the contributions of $g_A^{n,p}$ are suppressed due to the small axial nuclear charge.

With these parameters, the last remaining unknown in Eqn. 1.9 is $F_V(T)$, the weak vector nuclear form factor. It is more commonly written as a function of exchanged momentum $q^2 = 2MT$ and quantifies (roughly speaking) the overlap of the neutrino

wavefunction with the target nucleus. Explicitly, it is defined as the Fourier transform of the nuclear weak charge density. As the neutron coupling is much larger than the proton coupling, it is related to the nuclear neutron distribution, which is difficult to access by independent measurement. A popular parametrisation was introduced by Helm in 1956 [130], and discussed in greater detail in [39]. To model the effect of the nuclear skin, the density is taken to be a hard sphere of homogeneous density, convoluted with a Gaussian of thickness s . The nuclear radius is given by $r_n^2 = c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2$, where $a = 0.52$ fm, $s = 0.9$ fm, and $c = A^{\frac{1}{3}} \cdot 1.23$ fm $- 0.6$ fm. The Fourier transform can be calculated analytically:

$$F(qr_n) = 3 \cdot \frac{j_1(qr_n)}{qr_n} \cdot e^{-(qs)^2/2}, \quad (1.32)$$

where $j_1(x)$ is the first spherical Bessel function. In this way, a simple approximation to the form factor can be found for any nucleus. In practice, the form factor is one as long as $qr_n \ll 1$. At higher nuclear recoil energies, the roll-off of the form factor describes the increasing loss of coherence in the neutrino-nuclear scattering.

1.2.3. Sensitivity to New Physics

CE ν NS is a unique low-energy probe of the weak neutral current with the potential for precision measurements. Deviations from the Standard Model arise naturally in many extensions.

1.2.3.1. Weak mixing angle

Looking back at the couplings in Eqn. 1.20, a CE ν NS measurement can be viewed as constraining the weak mixing angle at low momentum transfers. As an alternative to incorporating the radiative corrections into the neutrino-quark couplings, the weak mixing angle can be said to run, i.e. to change value according to the exchanged momentum of the process in question. At low energies (below the weak scale), it picks up correction terms from the particle content of the SM. New physics, e.g. related to the muon g-2 anomaly, can result in a deviation of the SM expectation of the weak mixing angle at low energies. Thus, although the weak mixing angle is a standard model quantity, it provides sensitivity to BSM phenomena.

Fundamentally, the weak mixing angle is a parameter of electro-weak symmetry-breaking, describing the emergence of electric charge from the higher symmetries of weak isospin and weak hypercharge. In [129], it is defined via the gauge couplings of the electroweak symmetries SU(2) (weak isospin) g and U(1) (weak hypercharge) g' , as $\theta_W = \tan^{-1}(g'/g)$. Empirically, its cosine appears at tree level as the ratio of the W and Z boson masses. At lower energies, it has been measured in neutrino-lepton scattering [131] (neutrino/antineutrino ratio), deep-inelastic neutrino scattering [132, 133] (charged- to neutral-current ratio), parity-violating electron scattering [134, 135] and atomic parity violation [136].

The scale-dependence of the weak mixing angle, governed by a renormalization group equation (RGE), is a firm prediction of the SM [137]. The coefficients of the RGE change whenever the energy scale crosses a particle mass, where the particle decouples and is integrated out. Thus the running of the weak mixing angle depends on the ensemble of weakly interacting particles. Precision measurements of the weak mixing

angle at different exchanged momentum scales therefore contribute to a SM test and are sensitive to BSM models such as Z' extensions of the weak sector.

In the case of $CE\nu NS$, the weak mixing angle enters through the proton contribution to the nuclear weak charge (i.e. the small proton vector coupling to the Z -boson). The dependence can be made explicit by inserting the tree-level neutrino-nucleon couplings in the cross-section formula of Eqn. 1.12:

$$\sigma \propto (Z \cdot (4 \sin^2 \theta_W - 1) + N)^2 \quad (1.33)$$

Due to the strong suppression $(4 \sin^2 \theta_W - 1) \approx 5\%$, mostly the neutrons contribute to the large $CE\nu NS$ cross-section. Measuring precisely the suppression of the proton contribution allows constraining the weak mixing angle.

1.2.3.2. Non-standard interactions

Non-standard neutrino interactions (NSI) in general describe deviations from SM predictions in the couplings of neutrinos to other particles. Usually, contact interactions assuming a heavy mediator are studied, and only neutrino couplings to first-generation fermions (u, d, e) produce detectable effects.

Historically, neutral current NSI were first introduced by Wolfenstein to show they influence neutrino oscillation physics through matter effects (and induce oscillations also for massless neutrinos) [138]. In the age of precision oscillation experiments, NSI are studied for their potential to throw off the global interpretation of oscillation data by influencing all experiments in production, propagation and detection [139]. More generally, NSI arise in many extensions of the SM (such as theories of neutrino mass) and can therefore be used to probe BSM physics in an effective field theory approach, without assuming a specific model. NSI can arise at mass-dimension 6 from operators that produce (by $SU(2)_L$ symmetry) equally strong charged lepton interactions [140, 141]. These are strongly bounded by observation of charged lepton processes. To avoid these constraints, one has to consider dimension-8 operators involving the Higgs doublet, for which the $SU(2)_L$ relation may be broken, such that no effects on charged lepton physics are induced at low energy. In general, dimensional arguments lead to the expectation $\varepsilon \propto m_W^2/m_X^2$ for the amplitude of NSI at low energy ε and the NSI-generating new physics mass-scale m_X [142]. With a W-boson mass of $m_W \approx 80$ GeV and a 1(10) TeV NSI scale, one naturally expects $\varepsilon \propto 10^{-2}(10^{-4})$. For review articles see [142, 143]. Experimentally, the study of $CE\nu NS$ provides a new window for high-statistics neutrino measurements and promises unique sensitivity to particular NSI parameters.

Considering only neutral-current NSI contact interactions (relevant to $CE\nu NS$ and not constrained directly by charged-lepton physics), the additional terms in the Lagrangian can be written (following [140]):

$$\mathcal{L}^{NSI} = -\varepsilon_{\alpha\beta}^{fP} 2\sqrt{2}G_F(\bar{\nu}_\alpha\gamma_\rho L\nu_\beta)(\bar{f}\gamma^\rho P f) \quad (1.34)$$

where $\alpha, \beta = e, \mu, \tau$ denote neutrino flavors, $f = e, u, d$ is a first-generation fermion, and $P = L, R$ is a projection operator selecting chirality. The parameter $\varepsilon_{\alpha\beta}^{fP}$ gives the strength of the NSI interaction relative to the standard weak interaction. The operators can be classified as “flavor-changing” ($\alpha \neq \beta$) and flavor-conserving “non-universal” ($\alpha = \beta$).

Relevant to CE ν NS of (unoscillated) reactor neutrinos are only the couplings involving quarks (u,d) and an electron neutrino. Existing constraints on NSI parameters vary greatly in precision [140]. Monojet searches at the LHC constrain all neutrino-quark NSI at a level of 10% – 20% [141]. Hadronic scattering of electron neutrinos has been directly measured in the CHARM experiment [144] and $\mathcal{O}(50\%)$ constraints can be set on flavor-diagonal and flavor-changing electron-neutrino quark interactions. For $\varepsilon_{\mu e}^{qP}$, much stronger constraints ($\mathcal{O}(10^{-4})$) come from limits on $\mu \rightarrow e$ conversion on nuclei ($\varepsilon_{\mu e}^{qP}$ contributes at higher order via loop contributions).

As discussed in section 1.2.2, the axial weak nuclear charge is suppressed, so that CE ν NS is predominantly sensitive to combinations $\varepsilon_{\alpha\beta}^{qV} = \varepsilon_{\alpha\beta}^{qL} + \varepsilon_{\alpha\beta}^{qR}$ and insensitive to $\varepsilon_{\alpha\beta}^{qA} = \varepsilon_{\alpha\beta}^{qL} - \varepsilon_{\alpha\beta}^{qR}$. A future precise CE ν NS measurement using reactor neutrinos can give relevant constraints to the parameters ε_{ee}^{qV} and $\varepsilon_{\tau e}^{qV}$, for $q = u, d$.

1.2.3.3. Electromagnetic form-factors

With increasing precision of a CE ν NS measurement, observations of electromagnetic neutrino form factors come into view. The neutrino charge radius, induced by W-charged-lepton-loops, is predicted by the SM and in reach for planned CE ν NS experiments, but not particularly dependent on new physics [145, 146]. A more interesting (and challenging) case is the neutrino magnetic dipole moment (MDM). It predicts an additive electromagnetic component to the scattering cross-section (essentially Rutherford scattering) with a sharp rise in the recoil spectrum at extremely low energies. An observation of MDM provides deeper insight into neutrino mass generation mechanisms and the Majorana nature of neutrinos, as there are different allowed ranges for Dirac and Majorana neutrinos [147, 146].

Theoretical motivation In the SM, the neutrino is massless, electrically neutral and has exactly one nonzero electromagnetic form factor: a charge radius induced by electroweak radiative corrections. To accommodate the neutrino masses observed in oscillation experiments, the SM must be extended to a more fundamental theory. Depending on the structure of this underlying theory, e.g. Dirac or Majorana character of the neutrino, the neutrino may acquire additional electromagnetic interactions from loop effects induced by new physics.

Electromagnetic neutrino interactions are therefore an interesting sector where further non-standard behaviour of neutrinos may be observed. Electromagnetic neutrino properties can be probed in the laboratory through neutrino scattering. Stringent constraints are also derived from astrophysical observations, where electromagnetic neutrino interactions can lead to observable effects. A detailed review is given in [146].

Here, the discussion is focused on the charge radius (non-zero in the SM) and the magnetic dipole moment (linked to neutrino mass generation). The most straightforward electromagnetic form factor, the neutrino electric charge, is strongly constrained by the neutrality of matter and charge conservation in beta decay ($q_{\nu_e}/e < 10^{-21}$) [148].

The neutrino charge radius in the SM has been the subject of considerable debate before its renormalizability and nature as a physical observable were shown [149]. The association of the neutrino with a charge distribution is related to vertex corrections involving W/charged-lepton loops. The charge radius is the length scale associated with

this charge distribution, or more precisely the second moment of the charge form factor:

$$\langle r_\nu^2 \rangle = 6 \left. \frac{df_Q(q^2)}{dq^2} \right|_{q^2=0} \quad (1.35)$$

The result of the one-loop SM calculation including γ -Z-mixing and vector boson box diagrams is [149]:

$$\langle r_{\nu_l}^2 \rangle = -\frac{1}{2\sqrt{2}\pi} \cdot \frac{G_F}{\hbar c} \left[3 - 2 \ln \left(\frac{m_l^2}{m_W^2} \right) \right] \quad (1.36)$$

for neutrino flavor and associated lepton l . In the literature, there exist some conflicting definitions involving a factor 2 and a different sign. The sign flip is a convention ($\langle r_\nu^2 \rangle$ is equal and opposite for neutrinos and antineutrinos), while the factor of 2 is related to the normalization of the interaction vertex (essentially assuming an additional SM-forbidden anapole moment of the same strength) [145].

Following the definitions and formula above, for the electron neutrino one finds numerically $\langle r_{\nu_e}^2 \rangle = -0.83 \cdot 10^{-32} \text{ cm}^2$ [150].

The magnetic dipole moment interaction induces a neutrino helicity flip, and is therefore directly connected to the neutrino mass term in the simplest extensions of the SM. For massless neutrinos, no right-handed state exists and the magnetic dipole moment interaction strictly vanishes. For massive Dirac neutrinos, each mass state obtains a magnetic dipole moment proportional to its mass [146]:

$$\mu_{\nu,i} = \frac{3e}{8\sqrt{2}\pi^2} \cdot \frac{G_{Fc}}{\hbar^2} \cdot m_i = \frac{3}{4\pi} \cdot \frac{m_e c^2}{V} \cdot \frac{m_i c^2}{V} \cdot \mu_B \approx 3.2 \cdot 10^{-19} \cdot \mu_B \cdot \left(\frac{m_i}{\text{eV}} \right), \quad (1.37)$$

around 8 orders of magnitude below experimental constraints. There also exist transition magnetic dipole moments between the mass states, but they are suppressed by at least another factor of 10^{-4} . In the case of massive Majorana neutrinos, only two helicity states exist (compared to 2 particle- and 2 antiparticle states for Dirac neutrinos). This reduces also the allowed electromagnetic form factors, for example prohibiting diagonal magnetic dipole moments. The transition dipole moments are a-priori expected in the same range as for Dirac neutrinos. Due to the different structure of the mixing matrix (with extra phases compared to the Dirac case), the transition magnetic dipole moments of Majorana neutrinos could be enhanced in ways depending on the detailed physics model [146].

Turning the statement around, the model-independent theoretical limits on magnetic dipole moments of Dirac neutrinos are orders of magnitude stronger than those for Majorana neutrinos, which are around the level of astrophysical and laboratory limits [147]. Therefore a discovery of a neutrino magnetic dipole moment near current experimental limits would demonstrate the Majorana nature of neutrinos. This is a strong motivation for experimental searches for neutrino magnetic dipole moments.

Astrophysical constraints Neutrinos play important roles in many astrophysical situations, such as energy production in stellar fusion and energy transport in core-collapse supernovae. Observations on the behaviour of these astrophysical systems therefore allows precise deductions on the underlying neutrino properties [151].

An important example is the plasmon decay of photons into neutrino pairs in the interior of stars. In vacuum, such a decay is kinematically forbidden for the massless

photon (in the same way as conventional e^+e^- pair-production). Within the stellar plasma, the dispersion relation of photons is modified. This allows anomalous energy loss from photons into neutrinos, given some electromagnetic neutrino coupling. The observed life-time of the Sun can be used to place limits on electromagnetic neutrino interactions. More stringent limits on anomalous energy loss come from measurements of stellar core masses in globular clusters at helium ignition. They lead to constraints of $\mu_\nu < 3 \cdot 10^{-12} \mu_B$ and $q_\nu < 2 \cdot 10^{-14} e$ [151]. This limit applies to all neutrino states light enough to be produced in stars ($\lesssim 5$ keV) and also to transition moments, equally for Dirac and Majorana neutrinos.

Interesting inferences on neutrino properties can also be made from the observed neutrino signal from supernova 1987A. In a supernova core during collapse and explosion, matter is dense enough to efficiently trap neutrinos, which then move in a diffusive fashion through the core and contribute to the energy transport. A magnetic dipole interaction for Dirac neutrinos would flip neutrino helicities, making the neutrinos right-handed and nearly sterile with respect to the standard weak interaction. This is an escape mechanism for neutrinos from the supernova core, which would shorten the duration of the detectable neutrino burst. This leads to a similar constraint $\mu_\nu(\text{Dirac}) < 3 \cdot 10^{-12} \mu_B$. The escaping high-energy neutrinos can subsequently spin-flip back in galactic magnetic fields and cause additional high-energy events in terrestrial neutrino detectors. The absence of such signal leads to a stronger constraint $\mu_\nu(\text{Dirac}) < 1 \cdot 10^{-12} \mu_B$. Also the neutrino charge has a detectable effect on a supernova burst: charged neutrinos are deflected by galactic magnetic fields leading to an energy-dependent time-delay. The absence of such a dispersion leads to a charge limit $q_\nu < 3 \cdot 10^{-17} e$.

Weaker limits on electromagnetic form factors can also be derived from Big Bang Nucleosynthesis and radiative neutrino decay.

In summary, strong constraints on neutrino interactions can be inferred from observations of the cosmos. Exciting improvements could be possible with the observation of the next galactic core-collapse supernova. On the other hand, limits e.g. from plasmon decay are difficult to improve upon. This leaves an opportunity for laboratory limits to catch up with astrophysical constraints.

Laboratory limits and scattering phenomenology The best existing laboratory limits on electromagnetic neutrino form factors are obtained from elastic neutrino scattering on electrons. Among them, the charge radius has a special role in that its interaction is indistinguishable from the standard weak interaction, and therefore the scattering amplitudes interfere. A measurement of the charge radius is therefore equivalent to a precise determination of the total elastic cross-section. The best experimental limit for the electron neutrino charge radius is obtained by TEXONO [152], observing scattering of reactor $\bar{\nu}_e$ in a 187 kg CsI scintillator array above 3 MeV energy. With a measurement precision of 24% on the scattering cross-section, the charge-radius-constraint is $-4.2 \cdot 10^{-32} \text{ cm}^2 < \langle r_{\bar{\nu}_e}^2 \rangle < 6.6 \cdot 10^{-32} \text{ cm}^2$ (the numerical value is scaled up by factor of two with respect to [152] due to a different normalization of $\langle r_{\bar{\nu}_e}^2 \rangle$). This is one order-of-magnitude away from the SM predicted value.

For the case of CE ν NS, several limits have been derived from COHERENT data [153, 150, 154, 155] with slightly different treatment, e.g. $-62 \cdot 10^{-32} \text{ cm}^2 < \langle r_{\nu_e}^2 \rangle < 12 \cdot 10^{-32} \text{ cm}^2$ [153]. COHERENT constrains also $\langle r_{\nu_\mu}^2 \rangle$, which is not accessible to reactor experiments.

neutrino coupling	cross-section		ratio
	ν -electron	ν -nucleus	
Q_W	$\frac{G_F^2 m_e}{\pi} \cdot \frac{5}{4}$	$\frac{G_F^2 M}{\pi} \cdot \frac{N^2}{4}$	$\frac{N^2}{5} \frac{M}{m_e}$
μ_ν	$\frac{\pi \alpha^2}{m_e^2} \cdot \frac{\mu_\nu^2}{\mu_B^2} \cdot \frac{1}{T}$	$\frac{\pi \alpha^2 Z^2}{m_e^2} \cdot \frac{\mu_\nu^2}{\mu_B^2} \cdot \frac{1}{T}$	Z^2
q_ν	$\frac{2\pi \alpha^2}{m_e} \cdot q_\nu^2 \cdot \frac{1}{T^2}$	$\frac{2\pi \alpha^2 Z^2}{M} \cdot q_\nu^2 \cdot \frac{1}{T^2}$	$Z^2 \frac{m_e}{M}$

Table 1.1.: Electron antineutrino differential cross-sections $\frac{d\sigma}{dT}$ [159, 150, 160] to lowest order in T (i.e. $T \ll E_\nu \ll M$), for different couplings (weak charge Q_W , magnetic dipole moment μ_ν , electric charge q_ν) and targets (electrons, nuclei). To make a leading-order comparison, nuclear form-factors are neglected as well as axial weak charge and magnetic dipole moment of the nucleus. Weak couplings are simplified by setting $\sin^2 \theta_W \approx 1/4$.

In neutrino-nucleus scattering, the charge-radius contribution to scattering couples to the nuclear charge and interferes with the standard weak amplitude, essentially modifying the neutrino-proton coupling g_V^p for CE ν NS. It can be expressed with the following substitution [156]:

$$g_V^p \rightarrow g_V^p - \frac{2 m_W^2 c^4 \sin^2 \theta_W}{3 (\hbar c)^2} \langle r_{\nu e}^2 \rangle \quad (1.38)$$

This SM contribution to CE ν NS is therefore exactly degenerate with a deviation in the weak mixing angle, or non-universal NSI obeying $\varepsilon_{ee}^{uV} = -2\varepsilon_{ee}^{dV}$ (which cancels the effect on the neutron coupling). Treating the charge radius as a contribution to the weak mixing angle as measured in CE ν NS, the relative change in $\sin^2 \theta_W$ becomes

$$\sin^2 \theta_W \rightarrow \sin^2 \theta_W \cdot \left(1 + \frac{1}{3} \frac{m_W^2 c^4 \langle r_{\nu e}^2 \rangle}{(\hbar c)^2} \right) \quad (1.39)$$

which translates to a 4.4% reduction.

When analysing experimental data, one is left with the choice to assume the uncorrected SM value of the weak mixing angle and use the observed value as a measurement of the charge radius contribution, or else to subsume the charge radius into the expectation and use the observation as a test of this SM prediction.

The other electromagnetic parameters induce a neutrino helicity change and therefore do not interfere with the weak interaction. Also, dipole moment and charge interactions are strongly forward-focused for typical neutrino energies, resulting in a concentration of observable recoils at low energies (compared to the weak interaction). The best limits correspondingly come from small-target low-threshold experiments. The reactor experiment GEMMA [157, 158] using high-purity germanium spectrometer with a threshold of 2.8 keV does not observe the weak process, but places leading laboratory limits on dipole moment and charge interactions. This is explained by the different behaviour of the weak and EM cross-sections at low energies, as shown in Table 1.1. While the weak differential cross-section becomes constant at small recoil energies T , the dipole cross-section shows a $1/T$ divergence (charge even $1/T^2$). This lets the electromagnetic cross-sections rise far above the weak cross-section, which stays below backgrounds at low energies.

For neutrino-nucleus scattering, the situation is different. The weak process is enhanced roughly by a factor $N^2/5 \cdot M/m_e$ (see Table 1.1) with respect to ν -e scattering.

The factor $N^2/5$ is the coherent enhancement of the total cross-section, while the factor M/m_e results from the compression of the differential cross-section below a much-reduced kinematic endpoint. For μ_ν (q_ν), the coherent enhancement by Z^2 is comparable, but the interactions are suppressed by the target mass (squared), canceling the kinematic compression.

Thus the dipole moment sensitivity of a CE ν NS experiment is coherently enhanced with respect to electron scattering and there is an additional advantage $\propto 1/T$ from ultralow thresholds. On the other hand, CE ν NS acts as an irreducible background which is strongly suppressed for ν -e scattering. For this reason, a light nuclear target can be advantageous in spite of lower Z^2 , as the reduction of the CE ν NS cross-section is stronger than that of the dipole interaction. This is true in the optimistic case in which the sensitivity to a dipole moment is limited dominantly by a clear and strong CE ν NS signal and not by some neutrino-unrelated background.

For the direct measurement of neutrino charge, neutrino-nucleus scattering will not be competitive with neutrino-electron scattering. The rate in neutrino-nucleus scattering is comparable to the one in ν -e scattering (enhanced by Z^2 , suppressed by m_e/M). Observing with a small ultralow-threshold detector does not promise enhanced sensitivity, especially because of the presence of CE ν NS at low energies. Another advantage of neutrino-electron scattering in this case is the much higher typical recoil energies, which make the suppression of other backgrounds an easier task.

1.2.3.4. New light mediators

Non-standard interactions allow constraining the impact of new physics on a CE ν NS signal without making any assumptions to its detailed nature, i.e. particle content and coupling structure. This is possible as long as the new physics effects are limited to contact interactions, i.e. when all BSM particles interacting with neutrinos are much heavier than the exchanged momentum in the observed CE ν NS interaction. This in turn implies a restricted phenomenology of NSI: only the normalization of the CE ν NS recoil spectrum can be affected by such heavy new physics.

Treating the new physics with one more level of detail, a new mediator (in this case between neutrinos and quarks) with a mass not much heavier than the momentum exchange can be explicitly considered. In the simplest cases, the new mediator can be a scalar, or a vector particle. One should keep in mind that such simplified models are incomplete by construction and implicitly require larger new physics sectors to avoid unitarity bounds or other conceptual problems.

The phenomenology of new light mediators has been discussed in the context of LHC DM searches [161], direct detection experiments [162] and CE ν NS in COHERENT as well as other upcoming experiments [163, 164, 165].

Strong constraints exist on light mediators from a variety of laboratory, astrophysical and cosmological constraints. Important bounds on light scalars from supernova physics, Big-Bang Nucleosynthesis and the Cosmic Microwave Background are discussed in [163]. They depend on the type of Lagrangian considered (lepton-number violating or not). Many laboratory bounds on scalar and vector mediators rely on their coupling to charged leptons [164]. CE ν NS can therefore provide independent bounds on mediators coupling only to neutrinos and quarks. This argument does not apply to inelastic hadronic neutrino scattering and deep-inelastic scattering experiments, for which

however, no limits on new light mediators have been derived [164].

1.2.4. Artificial neutrino sources for CE ν NS

Although laboratory observation of CE ν NS has been discussed in the context of naturally occurring neutrino fluxes, either from the sun or during supernova bursts, there are a number of artificial neutrino sources which allow greater experimental control and are therefore deemed more useful to an experiment aiming to study the CE ν NS process.

1.2.4.1. Stopped-pion sources

Many accelerator facilities around the world produce neutrinos as a by-product of their normal operation. Sufficiently energetic proton beams impinging on a fixed target cause neutrino emission through pion decay at rest. The charged pions produced in hadronic interactions are stopped within the target and decay to monoenergetic muons (29.9 MeV) in the two-body reaction $\pi^+ \rightarrow \mu^+ \nu_\mu$. The emitted muons are also stopped within the target and subsequently decay at rest (after a few lifetimes of 2.2 μ s) and produce continuous spectra of ν_e and $\bar{\nu}_\mu$ in the three-body reaction $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ [166, 167]. An advantage of this type of neutrino source is the well-predicted energy distribution of the neutrino flux, shown in Figure 1.12.




Figure 1.12.: Neutrino emission from pion decay at rest, normalised to one neutrino per flavor as parametrised in [168]. The monoenergetic line of ν_μ results from pion decay. The continuous ν_e - and $\bar{\nu}_\mu$ -emission comes from the subsequent muon decay at rest, and is delayed by the muon lifetime (2.2 μ s mean).

The energy range of neutrinos emitted from pion decay at rest is ideally suited for the detection of CE ν NS. For an experiment, one would like the largest neutrino energies (resulting in a higher cross-section by E_ν^2 and higher-energy nuclear recoils) while still profiting from the coherent enhancement of the cross-section by N^2 . The neutrino energy should thus be high enough that typical exchanged momenta are below but close to the scale where coherence is lost. For a simple estimate of the required neutrino energy, we can set the wavelength associated to the maximum momentum exchange to

the diameter of the target nucleus: $q_{\max} = h/(2R)$. Kinematics gives the maximum momentum exchange as $q_{\max} = 2E_\nu/c$. The diameter of heavy target nuclei is around 10 fm. With these numbers the highest neutrino energy profiting from coherence is around $E_\nu = hc/4R \approx 60$ MeV. Stopped-pion sources in that sense hit the sweet spot for CE ν NS detection. With neutrino energies up to about 53 MeV, the nuclear recoils extend up to 56 keV in recoil energy (assuming a target nucleus with $A = 100$).

An important class of stopped-pion sources are spallation neutron sources, such as SNS at Oak Ridge National Laboratory, where CE ν NS was first observed by the COHERENT experiment [127]. At SNS, a 1 GeV proton beam strikes a mercury target for neutron production, with a 60 Hz repetition rate of 1 μ s beam spills. $5 \cdot 10^{20}$ protons on target per day result in an average neutrino emission of $5 \cdot 10^{16}/s$ in each flavor (ν_μ , ν_e , $\bar{\nu}_\mu$). Because of the time structure of the proton beam, the peak instantaneous neutrino flux is over four orders of magnitude larger, so that fast detectors profit greatly from the timing information to subtract steady-state backgrounds. As the beam spills at SNS are shorter than the muon lifetime, detectors with sub-microsecond time resolution can even distinguish beam-related backgrounds from the delayed neutrino signal from muon decay. All CE ν NS experiments at neutron sources depend on the availability of an experimental site free from the “main product” of the facility, as neutrons are a dangerous background to CE ν NS searches. In addition, a location with some overburden against cosmic radiation is desirable.

1.2.4.2. Nuclear reactors

Nuclear reactors are strong artificial continuous sources of electron antineutrinos. They emit an average of 6.7 $\bar{\nu}_e$ per fission event at a fission rate of $3.1 \cdot 10^{19} \text{ s}^{-1}$ per GW of thermal power [169]. This translates to an antineutrino rate of $8.7 \cdot 10^{20} \text{ s}^{-1}$ emitted isotropically from a reactor with 4 GW thermal power. The typical energy of the produced antineutrinos is a few MeV, with a high energy tail extending up to about 13 MeV [170]. 75% of the flux falls below 1.8 MeV [169], where detection through the inelastic inverse beta decay reaction $\bar{\nu}_e + p \rightarrow n + e^+$ is impossible. Modeling of reactor antineutrino fluxes is discussed in more detail in section 6.1.1.

Experimentally, the high flux ensures large event numbers, even after accounting for the reduced cross-section due to the lower neutrino energy. Interaction rates can rise above 10^3 counts/(keV·kg·day) at 40 m distance from a power reactor, as shown in section 6.1.1. The big challenge of reactor-neutrino CE ν NS is the recoil energy: the highest recoil energy for a typical neutrino energy of 3 MeV is only 180 eV (assuming a target nucleus with $A = 100$). Ultra-low threshold detector technology is therefore a prerequisite for a detection.

Reactor neutrinos arrive continuously and without time-tag, a fast subtraction of reactor-unrelated backgrounds is therefore not possible. Reactor-off periods, which are often scheduled for several weeks per year, can be beneficial for this purpose. This depends on the reactor schedules and the number and proximity of reactor cores at the experimental site. A detailed study of this is given in section 7.2.3.

In selecting an experimental site at a nuclear reactor, there is a trade-off between sufficient distance from the reactor to avoid reactor-correlated neutrons (a dangerous background) and enough proximity for a high neutrino flux. This is discussed more deeply in section 7.2.1.

A special case are research reactors, typically with a thermal power of tens of megawatt, orders of magnitude weaker than power reactors. Still, the research-oriented nature of such reactors may allow finding a suitable experimental site. Potential benefits of such a location, apart from good access to the experimental site, typically include more frequent and prolonged off-periods compared to power reactors. A study of such a scenario is shown in section 7.1.2.

1.2.4.3. Radiogenic sources

A class of artificial neutrino sources used successfully in past experiments are radiogenic sources, i.e. carefully purified samples of a beta-decaying isotope producing (anti-)neutrinos. Prominent examples include the ^{51}Cr source used for calibration in the solar neutrino experiment GALLEX [171], and the planned ^{144}Ce source for the cancelled neutrino oscillation experiment SOX [172].

Compared to a nuclear reactor, radiogenic sources offer the possibility to place the source in close proximity to the detector and locate the experimental setup underground for control of external backgrounds. Also source-correlated backgrounds can be addressed by experimental design of the source shielding and stringent requirements on isotope purity. The challenges are a much lower neutrino energy (750 keV and 430 keV for ^{51}Cr , 3 MeV endpoint for ^{144}Pr produced from ^{144}Ce) and comparatively small neutrino fluxes.

Radiogenic neutrino sources generally produce unavoidable gamma radiation from beta decay to excited states, which has to be shielded in a high-Z source container. In addition, dangerous backgrounds to CE ν NS measurements can arise from impurities in the source. The chemical purity of the source thus has to be strictly controlled to avoid e.g. neutron production on a level prohibitive to the neutrino experiment.

Due to the low neutrino energies, ultra-low detection thresholds even beyond the range required for reactor experiments are necessary. Large target masses are also required, as the lower neutrino flux cannot be compensated by an arbitrarily long measurement time, due to the finite lifetime of the radiogenic source. In a recent study, the CE ν NS potential of a re-activation of the ^{51}Cr source used in GALLEX is investigated [173]. In section 7.1.3, a hypothetical CE ν NS experiment using a SOX-like source is studied.

1.2.5. CE ν NS experiments

Through their common event signature, a low-energy nuclear recoil, the processes of dark matter direct detection and CE ν NS are intimately related. This is apparent in the history of the detection efforts: a widely noted suggested technology for the observation of CE ν NS “for neutrino physics and astronomy” by Drukier and Stodolsky in 1984 [174] was quickly followed by the insight (by Goodman and Witten, 1985 [88]) that such a technology would also be sensitive to dark matter from the galactic halo. While the specific approach envisioned by Drukier and Stodolsky, based on metastable superconducting grains, did not lead to practical experiments, the first direct dark matter limit was derived a few years later using a low-background germanium spectrometer [89]. Since these early days, technology developed in the pursuit of CE ν NS has often been applied to direct dark matter search, and vice versa. The dramatic improvement in detector technology evidenced by the tightening in dark matter exclusion limits since the 1990s has led to numerous attempts to observe CE ν NS.

Shared requirements for dark matter- and $\text{CE}\nu\text{NS}$ -experiments are a low energy threshold, large number of target nuclei, and low intrinsic background rate. Large nuclei are favored to take advantage of the coherent enhancement. In addition, $\text{CE}\nu\text{NS}$ detectors must operate in close proximity to a neutrino source, which may be inherently a high-background environment.

While only the COHERENT collaboration has achieved a significant observation of $\text{CE}\nu\text{NS}$, there is now a host of experiments in various stages of planning, commissioning and taking data. These experiments employ different low-threshold detector technologies using various target materials. To reap the full science potential of $\text{CE}\nu\text{NS}$, it is mandatory for the community to make many independent measurements ideally at distinct neutrino sources using a wide variety of nuclei. At the moment, COHERENT remains the only collaboration working on pion decay-at-rest neutrinos, while other approaches target nuclear reactors as neutrino sources.

1.2.5.1. The COHERENT experiment

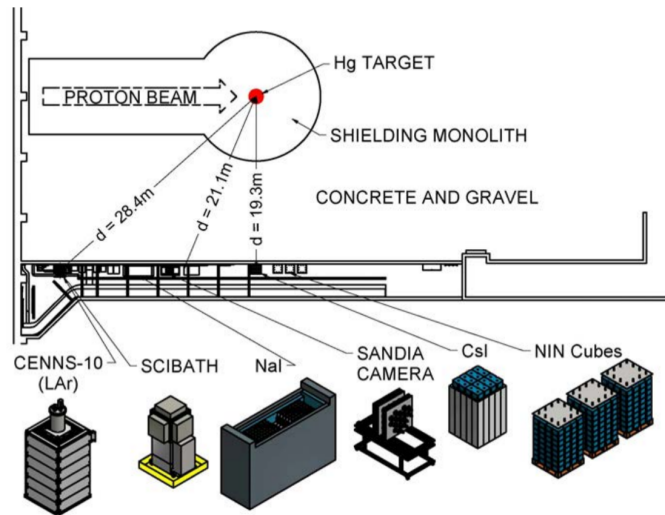


Figure 1.13.: COHERENT detectors deployed at SNS in Oak Ridge, Tennessee. The “neutrino alley” has an overburden of 8 m.w.e. and more than 12 m shielding (concrete, gravel) against neutrons from the spallation target. The experimental location is densely populated with detectors using different target materials for $\text{CE}\nu\text{NS}$ and other processes such as neutrino-induced neutrons. Figure from [127].

The COHERENT collaboration operates a suite of detectors in close proximity to the spallation target at the Spallation Neutron Source (SNS) at Oak Ridge National Lab in Tennessee, USA. The potential of stopped-pion sources for $\text{CE}\nu\text{NS}$ has been discussed for over a decade [167]. The advantages of such a source are the highest neutrino energies profiting from coherent enhancement combined with a time-structure suitable for background subtraction. Even weighed against a small overall flux compared to nuclear reactors, these advantages allowed the first observation of $\text{CE}\nu\text{NS}$. An important requirement for $\text{CE}\nu\text{NS}$ at a neutron source is an experimental location well-shielded

from beam-related neutrons. COHERENT succeeded after numerous background measurement campaigns to identify an experimental location without direct connection to neutron beam facilities. Figure 1.13 shows a map of “neutrino alley”, a basement corridor repurposed as an experimental space. A variety of target materials using different detector technologies is deployed, with the goal of demonstrating the N^2 -dependence of the $\text{CE}\nu\text{NS}$ cross-section.




Figure 1.14.: COHERENT dataset from the first $\text{CE}\nu\text{NS}$ observation [127]. Shown are differences in photon counts in the 14.6 kg CsI scintillator, between the 12 μs following a beam trigger and the 12 μs before. These differences are displayed as a function of photon number (A) and arrival time (B). “Beam OFF” signifies data accumulated over 153.5 live-days of SNS inactivity and “Beam ON” is from 308.1 live-days with neutrino production. An excess matching the SM prediction (shown as coloured histograms in the right panels) is observed at a significance of 6.7σ .

The first COHERENT detector, deployed in 2015, is the 14.6 kg CsI scintillator which subsequently made the first observation of $\text{CE}\nu\text{NS}$ [127] (announced in August 2017). The dataset from 153.5 live-days without beam and 308.1 live-days with beam is shown in Figure 1.14. Compared to a SM prediction of (174 ± 49) $\text{CE}\nu\text{NS}$ events, (134 ± 22) counts were observed. Adding the prediction and observation uncertainties in quadrature, this can be interpreted as a cross-section measurement with a precision of 33%.

This important milestone for $\text{CE}\nu\text{NS}$ physics triggered a large number of papers devoted to interpretations and derivations of bounds on new physics from this dataset.

Since the first observation using CsI, the COHERENT collaboration published results from a three-months engineering run of a 29 kg single-phase liquid argon detector, CENNS-10 [175]. With a threshold of 80 keV (NR) this detector was not able to observe $\text{CE}\nu\text{NS}$ in that run, but provided important measurements on beam-related prompt and delayed neutron backgrounds. In January 2020, results from the first production run of CENNS-10, with an improved threshold of about 20 keV (NR), were presented at a Fermilab seminar [176] and published subsequently as a preprint [177]. With $159 \pm 43(\text{stat.}) \pm 14(\text{syst.})$ observed events over a period of 17 months (July 2017 - December 2018), an observation of $\text{CE}\nu\text{NS}$ with 3.5σ significance was made, limited by statistics. Additional data-taking with CENNS-10 is ongoing at the time of writing.

Future planned detectors of the COHERENT collaboration include a tonne-scale liquid

argon detector CENNS-750 (potentially using underground argon to avoid backgrounds from ^{39}Ar), 16 kg of high-purity germanium detectors and a tonne-scale NaI scintillator array. Presently, all cross-section measurements at SNS suffer from a 10% uncertainty in neutrino production. As a measure to reduce this uncertainty, deployment of a tonne-scale heavy-water detector is planned. This could reduce the neutrino flux uncertainty by measuring the well-predicted charged-current reaction between electron neutrinos and deuterium (known to $\sim 3\%$).

1.2.5.2. Upcoming $\text{CE}\nu\text{NS}$ experiments at nuclear reactors

Although the first observation of $\text{CE}\nu\text{NS}$ at a spallation source demonstrates the advantages of this experimental concept, the high neutrino fluxes available only at nuclear reactors promise large statistics for more precise $\text{CE}\nu\text{NS}$ measurements. Several experimental efforts are tackling the challenge of detecting nuclear recoils from neutrinos with roughly one order of magnitude lower energy.

TEXONO The TEXONO collaboration formed in 1997 is an established player in low-energy reactor neutrino physics [178]. They set up a neutrino physics laboratory at Kuo-Sheng Nuclear Power Station and operate detectors in 28 m distance to the reactor core. Neutrino-electron scattering has been studied with a CsI scintillator array, resulting in the current best limit on the electron neutrino charge radius [152]. Leading searches for the magnetic dipole moment of the neutrino using low-threshold high-purity germanium detectors have been performed [179]. With an achieved energy threshold of around 300 eVee, $\text{CE}\nu\text{NS}$ has so far not been observed at Kuo-Sheng Nuclear Power Station. Efforts to lower backgrounds and energy threshold are ongoing to push for a $\text{CE}\nu\text{NS}$ detection [180]. The collaboration is a founding member in the CDEX effort applying germanium technology for dark matter searches, an example of the synergies between $\text{CE}\nu\text{NS}$ and direct detection.

CONUS CONUS is a comparatively young experimental effort at Max-Planck-Institut für Kernphysik in Heidelberg, Germany [181]. Four 1 kg low-threshold point-contact germanium spectrometers equipped with electrical cryocoolers are deployed at 17 m distance from a reactor core at Kernkraftwerk Brokdorf in northern Germany. The experiment uses a high-efficiency muon veto and a compact passive shielding evolved from the GIOVE setup for low-background gamma spectrometry [182]. A favourable background level of $\sim 1 /(\text{keV kg day})$ was demonstrated with the CONUS shielding in an underground laboratory in Heidelberg. A challenge for this kind of experiment is the lack of experimental data on the “quenching factor” in germanium at $\text{CE}\nu\text{NS}$ -relevant energies. This conversion factor is needed to account for the reduced charge production by nuclear recoils compared to electron recoils (as induced by the gamma sources used in calibration). Extensive studies on reactor correlated neutron backgrounds were performed [183]. Muon-induced neutron backgrounds from within the shielding are expected to dominate, while neutrons from the reactor cores are efficiently shielded. The experiment has been taking data since April 2018. Data from the first exposure of two months (114 kg days reactor-off and 112 kg days reactor-on) has been presented and shows an excess of about 1 count/kgday at a significance of 2.3σ [184]. Sensitivity improvements with more statistics are expected in the near future. A second phase of

CONUS with 100 kg target mass is planned for precision CE ν NS studies.

CONNIE The CONNIE collaboration operates sensitive CCD detectors 30 m from the Angra 2 nuclear reactor at Almirante Alvaro Alberto Nuclear Power Plant in Brazil. An interesting feature of these detectors is their natural imaging capability: muons and Compton electrons can be easily distinguished from diffusion-limited hits, such as produced by X-rays or nuclear recoils. A disadvantage is their slow readout, which for example prevents the operation of anti-coincidence vetos. In an engineering run in 2014-15, a prototype array of four 1 g Si CCDs achieved a noise of 7 eVee. A 48 g array of 8 CCDs installed in 2016 [185] recorded 3.7 kg day of data with a nuclear recoil threshold of 1 keV. While no excess in reactor-on data with respect to reactor-off was observed and sensitivity is a factor 40 above the predicted CE ν NS rate, the dataset was used to set leading limits on light scalar and vector mediators contributing to the neutrino scattering cross-section [186]. This highlights the potential of studying reactor neutrinos with low-threshold detectors. The CONNIE collaboration is working on more accurate determinations of the quenching factor in silicon and a detector upgrade employing higher-efficiency Skipper CCDs. The current analysis is based on a subtraction of reactor-off and -on data, limited by reactor-off statistics. A full background model is in development that will allow using the full recorded statistics.

MINER The US experiment MINER [187] plans to deploy low-threshold cryogenic germanium and silicon detectors at a research reactor at Texas A&M University. The detector technology is derived from the CDMS dark matter experiment, using Neganov-Luke amplification to measure drifting charges via their heat signal. The reactor power is comparatively small at 1 MW, but the detectors are placed at a distance of only 2 m, resulting in a neutrino flux comparable to a site 40 m from a 4 GW power reactor. In addition, the reactor core is movable, potentially providing interesting sensitivity to sterile neutrinos [188]. Neutron and gamma background measurement campaigns are reported in [187]. According to [180], a payload of 10 kg is planned which will allow a CE ν NS observation within one month. The cryogenic infrastructure for the experiment has been commissioned and first data-taking is expected to be ongoing.

ν GeN Kalinin Nuclear Power Plant (KNPP) in Russia has been hosting a number of neutrino experiments since almost 20 years [189]. Building on experience from GEMMA [157], a germanium spectrometer placed at a baseline of 14 m (which sets the current best limit on neutrino magnetic dipole moment via $\bar{\nu}_e$ -e scattering), a new experiment called ν -GeN targets detection of CE ν NS using four 400 g Ge diodes [190]. The experiment uses the shielding developed for GEMMA, which allowed a background level at the reactor site of 2.5 /keV kg day at few keV energy. The ν GeN detectors were characterised at the underground laboratory of Modane and showed an energy resolution of 220-270 eV at 10.4 keV. At KNPP, ν GeN is targeting a recoil energy threshold of 350 eV, which will allow detection of tens of CE ν NS events per day.

RED-100 The dual-phase liquid xenon detector RED-100 [191] with 160 kg in the active volume has been under development since 2002. Deployment at SNS has been studied as well a location at KNPP [192]. With a threshold of 3 photoelectrons, more

than a thousand CE ν NS events per day are expected at the reactor distance of 19 m at KNPP.

In an engineering run of RED-100 performed above ground in February 2019 [193] the single-electron count rate was found to be ~ 250 kHz. This implies a higher threshold of 5 photoelectrons has to be used to suppress accidental coincidences. Even under these circumstances and using a more pessimistic quenching factor, around 6 CE ν NS events are expected per day, which makes CE ν NS detection at KNPP feasible. Efforts to reduce the accidental coincidence background include a more effective muon windowing, higher fields for better photon yield (improving spatial reconstruction) and better purification.

RICOCHE The concept for the cryogenic experiment RICOCHET has been under development for almost 10 years, starting as an idea in 2011 to validate the Reactor Antineutrino Anomaly at an intense ^{37}Ar source with 10^4 Si bolometers of 50 g each equipped with Mo/Au TES [194]. Later in 2012 it was proposed using modified CDMS Ge and Si detectors (employing W-TES with a lower T_c) for a first CE ν NS observation at a research reactor [195]. Some years later an experiment employing 10 kg of Zn and Ge bolometers with 100 eV recoil threshold at Chooz Nuclear Power Plant was studied [196]. A neutrino laboratory at the “Near Site” of Double Chooz (~ 400 m) was considered to host the experiment. In [165], the physics reach of Zn and Ge targets (500 g, 50 eV recoil threshold for phase 1, 5 kg, 10 eV recoil threshold for phase 2) at the Double Chooz Near Site or a nearer location is studied in greater detail. The goal is to develop single-charge-sensitive readout using HEMT amplifiers [197], achieving simultaneously 10 eV heat and 20 eVee charge resolution for background rejection at lowest energies.

2. Cryogenic detectors for rare event searches

2.1. Cryogenic calorimeters

It is a basic consequence of the third law of thermodynamics that all heat capacities must vanish at zero temperature¹. This allows for cryogenic calorimetry: the precise measurement of energies via temperature changes observed at very low temperature. In the basic calorimeter equation:

$$\Delta T = \frac{\Delta E}{C}, \quad (2.1)$$

the heat capacity C appears in the denominator, so that a small C ensures a large temperature rise ΔT for a given energy deposition ΔE .

In an insulating crystal at low temperature, the heat capacity is determined by the crystal lattice contribution which is well described by the Debye model [198]:

$$C = 9Nk_B \cdot \left(\frac{T}{T_D}\right)^3 \cdot \int_0^{T/T_D} dx \cdot \frac{x^4 e^x}{(e^x - 1)^2} \approx \frac{12\pi^4}{5} \cdot Nk_B \cdot \left(\frac{T}{T_D}\right)^3 \quad (2.2)$$

where the second transformation uses $T \ll T_D$. The Debye temperature T_D is a material parameter related to speed of sound (i.e. phonon density of states) and volume density of the lattice. Far below the Debye temperature, most of the phonon modes of the crystal are energetically inaccessible, which dramatically reduces the heat capacity. For example, for a 24 g Al_2O_3 crystal ($M = 102$ g/mol, $T_D = 1042$ K [199]), the heat capacity at $T = 15$ mK is $C = 1.4 \cdot 10^{-12}$ J/K. To put this in particle physics units, the temperature rise per unit of deposited energy is:

$$\frac{\Delta T}{\Delta E} = C^{-1} = 7.3 \cdot 10^{11} \text{ K/J} \approx 120 \text{ } \mu\text{K/keV} \quad (2.3)$$

This demonstrates how, thanks to the freeze-out of phonon modes at low temperature, a keV-scale energy deposition on one of the $\mathcal{O}(10^{23})$ atoms in a macroscopic crystal can lead to a significant temperature increase (around 1% in this example)². This makes cryogenic detectors an excellent choice for low-energy rare-event searches, where ultimate sensitivity to small energy depositions is needed.

This chapter gives an overview of different implementations of cryogenic detectors in fundamental physics, before expanding upon the general principles of the specific type

¹The third law states that the entropy of any system becomes constant as it approaches absolute zero. Entropy is linked to heat capacity by the relation $dS = \delta Q/T = C(T) \cdot dT/T$. For dS to vanish as $T \rightarrow 0$, also $C(T)$ must vanish.

²The simplistic discussion here assumes thermal equilibrium within the crystal, which is the case for some (slower) implementations of cryogenic calorimetry. In the macroscopic detectors discussed in this work, the situation is more complex (as discussed in section 2.3.1) but leads to the same general picture.

of cryogenic detector considered in this work: a macroscopic calorimeter equipped with a Transition Edge Sensor (TES).

2.2. Cryogenic sensor technologies for fundamental physics

The key technology in cryogenic calorimetry is the sensor which performs the temperature measurement, i.e. converts a temperature change to an electrically readable quantity (current, voltage change). It should be precise (high resolution), fast (high bandwidth), work at very low temperatures (small readout power), and preferably consist of simple elements that can be scaled up to large detector arrays. Depending on the application, several different technical solutions for cryogenic temperature sensors have emerged with different advantages and challenges.

2.2.1. Neutron transmutation doped germanium thermistors

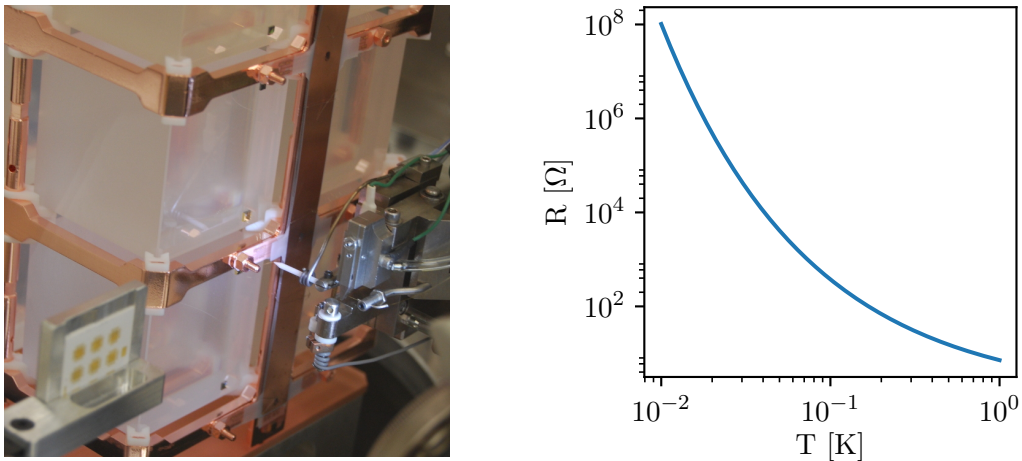


Figure 2.1.: Neutron transmutation doped germanium thermistors as low-temperature sensors. Left: cryogenic detector developed for the CUORE experiment [200]. The NTD sensor is the dark square glued to the large absorber crystal. Right: temperature dependence of NTD resistance, following equation 2.4, with typical parameters $R_0 = 1.15 \Omega$, $T_0 = 3.35 \text{ K}$ [201].

Semiconductors with an extremely homogeneous dopant concentration exhibit a sharp temperature dependence of resistivity at low temperature. Under these conditions, conduction happens via tunneling between states localized at dopant sites, through the mechanism of variable range hopping [202]. The temperature dependence of the resistance is:

$$R(T) = R_0 \cdot e^{\sqrt{T_0/T}} \quad (2.4)$$

Technically, semiconductor samples with the required homogeneous doping can be produced using neutron transmutation doping of germanium [107, 203] (NTD Ge thermistors). In an intense thermal neutron flux, some isotopes of germanium are converted

by neutron capture to donor (As, Se) or acceptor (Ga) atoms. In this process, the initially homogeneous distribution of isotopes in the crystal is transferred to a homogeneous distribution of dopants.

Figure 2.1 (left) shows a detector developed for the CUORE experiment [204] pursuing neutrinoless double-beta decay research [205]. In Figure 2.1 (right), the NTD resistance as a function of temperature is shown with parameters typical for CUORE.

NTDs are high-impedance ($> \text{M}\Omega$) sensors that are current-biased (typically $\mathcal{O}(\text{nA})$) and read out with voltage amplifiers. An advantage is the comparative simplicity of the sensor and required electronics, allowing arrays of $\mathcal{O}(10^3)$ channels. A disadvantage (in some applications) is their slow response time ($\sim\text{ms}$) and relatively large heat capacity (from the NTDs $\sim\text{mm}^3$ volume).

2.2.2. Metallic magnetic microcalorimeters

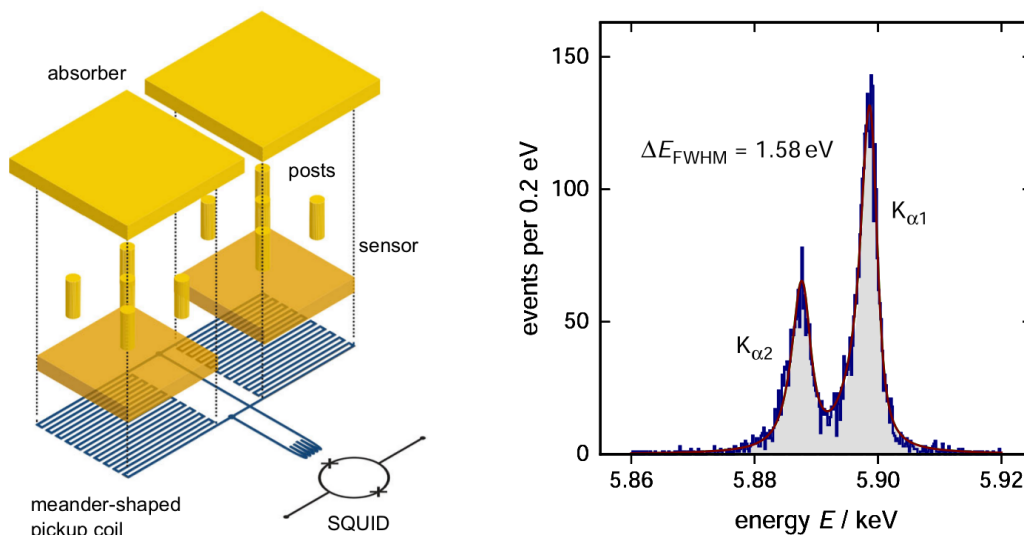


Figure 2.2.: Metallic magnetic microcalorimeters (MMCs). Left: principle schematic, showing the paramagnetic sensor (orange) connected to a metallic photon absorber (yellow). The SQUID magnetometer measures the induction signal caused by a temperature change in the sensor. Picture from [206]. Right: K_{α} X-ray lines of ^{55}Mn observed with an MMC [206]. The energy resolution of 1.58 eV (FWHM) allows separating the $K_{\alpha 1}$ and $K_{\alpha 2}$ sub-components.

Metallic magnetic microcalorimeters use the temperature-dependent magnetization of a paramagnetic material in a weak magnetic field. At low temperature, this follows the Curie law $M \propto T^{-1}$. Modern devices use dilute paramagnetic alloys, mostly gold doped with a few 100 ppm of erbium [207]. Figure 2.2 shows a conceptual drawing of an MMC X-ray calorimeter, and the eV-scale energy resolution obtained with it for ^{55}Mn X-ray lines. Advantages of this technology include the high energy resolution, good detector linearity (up to tens of keV) and fast rise times $\mathcal{O}(100 \text{ ns})$. A challenge is the complex cryogenic readout based on SQUID magnetometers, which require development of multiplexing schemes to scale up to large arrays. The ECHo experiment [208] is developing a 100 pixel (eventually many kilopixel) MMC array to determine the effective electron

neutrino mass via the endpoint of ^{163}Ho beta-decay. The AMoRE experiment [209] uses MMCs coupled to massive crystals (CaMoO_4 , few 100 g each) to search for neutrinoless double-beta decay of ^{100}Mo .

2.2.3. Microwave kinetic inductance detectors




Figure 2.3.: Left: MKID-based large-area cryogenic light detector developed within the CALDER project (photograph from [210], device schematic from [211]). Right: response of the microwave resonance of an MKID to increasing load, from [212].

Microwave kinetic inductance detectors [213] (MKIDs) take advantage of the high mobility of charge carriers in superconductors. In principle any conductor is characterized, along with its geometric inductance, by a kinetic inductance which arises not from the magnetic field energy of a flowing current, but from the kinetic energy of the charge carriers. Due to the short scattering times of electrons in normal metals, this effect is usually observable only at very high frequencies. Using a superconducting material, it is possible to construct devices where the kinetic inductance dominates over the geometric one at GHz frequencies. This allows fabricating microwave resonators whose frequency and quality factor is sensitive to the Cooper pair density in the material. In the conceptually simplest case, the lumped-element kinetic inductance detector [214], the resonator is formed by a meander-shaped kinetic inductor in parallel to an interdigitated capacitor. Figure 2.3 (left) shows a schematic along with a four-pixel device. Figure 2.3 (right) shows the transmission of a single MKID as a function of frequency. Quasiparticles excited by deposited energy in the superconducting meander change the kinetic inductance, leading to a frequency shift of the resonance and a reduced quality factor. This is measurable as an amplitude and phase shift of a microwave probe tone coupled to the MKID via a transmission line. The most important advantage of MKIDs (combined with high-bandwidth cryogenic microwave amplifiers) is the capability to probe a large number of separate resonators with a single transmission line. In addition, MKIDs are relatively easy to fabricate, requiring only a single superconducting material and no superconducting readout electronics. As a consequence, large MKID arrays are developed and deployed as astronomical instrumentation. A recent example is a 10 kilopixel array demonstrated on-sky in a near-infrared integral field spectrograph [215]. The CALDER project [216] aims to port this technology from astronomy to particle physics, for use as large-area cryogenic light detector for background suppres-

sion in neutrinoless double-beta decay. Figure 2.3 (left) shows a four-pixel CALDER test device that achieved 200 eV baseline energy resolution. This was improved to 26 eV using an Al/Ti/Al trilayer process [216].

2.2.4. Transition edge sensors

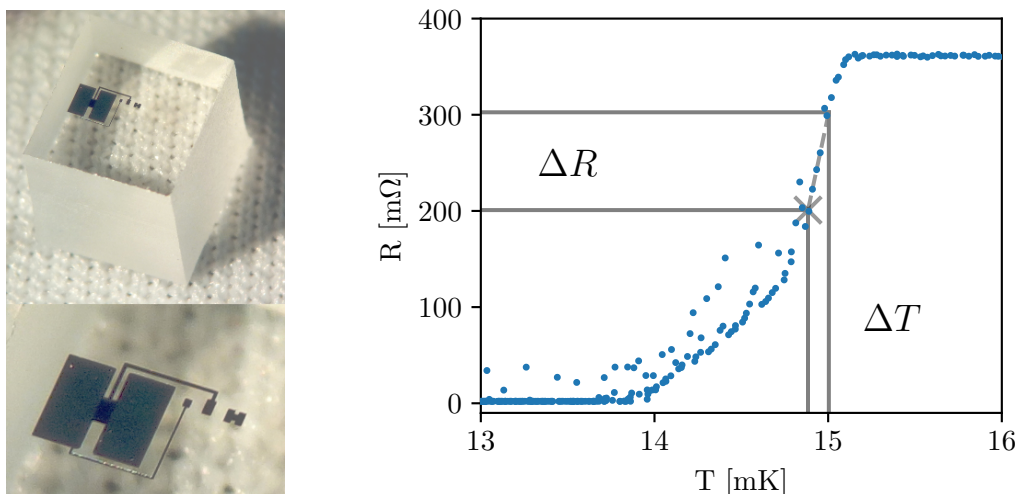


Figure 2.4.: TES on a macroscopic absorber. Left: 5 mm Al_2O_3 particle detector equipped with a tungsten TES. The bottom zoom shows the tungsten (black), aluminum (grey) and gold layers forming the sensor. Right: TES resistance as a function of temperature. Near the top of the transition, a linear range can be identified in which a temperature rise $\Delta T = 120 \mu\text{K}$ is transformed into a resistance change of $\Delta R = 102 \text{ m}\Omega$. The corresponding operation point of the TES is marked with a cross. Outlying data points above the curve are due to particle events during the recording of the transition.

Among the most sensitive cryogenic thermometers are Transition Edge Sensors (TESs) [217]. Such devices, originally also known as Superconducting Phase Transition (SPT) thermometers [109], use the steep resistance change occurring in a thin metal film as it passes from the normal-conducting into the superconducting state.

TES microcalorimeter arrays approach the energy resolving power of dispersive optical instruments, but retain the large collection area and high efficiency of calorimeters. This led to a rapidly increasing use of these instruments for X-ray and gamma-ray detectors over the last two decades (for a review see [218]). Applications of TES microcalorimeter arrays include material analysis and X-ray astrophysics. In these fields, TES are often made from alloys (e.g. Al-Mn) or bilayers (e.g. Mo-Cu, Mo-Au, Ti-Au) with transition temperatures between 50 mK and few 100 mK (depending on the cryogenic technology).

For rare-event searches, much larger absorbers are used, since the physics sensitivity scales with the target mass. To optimize the energy resolution, lowest possible temperatures are desirable. For the detectors discussed in this work, the material of choice is tungsten. Its bulk transition temperature of 15 mK (in the α crystal structure) is well

matched to the lowest temperatures continuously sustainable using He^3/He^4 dilution refrigerators (5-10 mK for optimized setups, with a record of 1.75 mK [219]).

At MPP Munich, tungsten thin films of $\mathcal{O}(\text{mm}^2)$ area and 200 nm thickness are produced by electron beam evaporation (with an alternative magnetron sputtering process under development [220]). The sample temperature is held at 400 – 600 °C during evaporation. This degree of freedom allows tuning the film transition temperature in the range 10 – 25 mK. Fig. 2.4 shows a detector equipped with a tungsten TES and its superconducting transition. Between 14 and 15 mK, the normal resistance of 360 m Ω vanishes completely. During operation, the TES is heated to a fixed temperature, the operation point (an example is marked by a cross in Figure 2.4). In the limited linear range above the operation point, a temperature rise is directly mapped to a resistance increase. The sensitivity of a TES is usually described by the dimensionless parameter

$$\alpha = \frac{dR}{dT} \cdot \frac{T_0}{R_0}. \quad (2.5)$$

In the given example, $\alpha \approx 63$. The large values of α realized in TES devices explain their outstanding sensitivity.

TES are low-impedance devices that require SQUID amplifiers for readout. Often, also the fabrication of thin films with tightly controlled parameters is a technical challenge. This makes large-scale application cumbersome. The unparalleled energy resolution motivates the ongoing efforts applying TES-based macrocalorimeters for rare-event searches.

2.3. Principles of TES macrocalorimeters

2.3.1. Thermal detector model




Figure 2.5.: Model of heat capacities and thermal conductances relevant for signal evolution in TES macrocalorimeters [109] (drawing adapted from [221]). The labelled quantities are defined in the text.

The simple presentation of cryogenic calorimetry in section 2.1 assumes thermal equilibrium and makes unrealistic assumptions such as neglecting the presence of the thermometer. To understand the signal formation in TES macrocalorimeters, processes in

the system after an energy deposition have to be studied in some additional detail. The discussion here summarizes the arguments and conclusions in [109].

The relevant heat capacities and thermal conductances in a macroscopic cryogenic calorimeter are shown in Figure 2.5. Besides the heat capacity of the absorber crystal C_c , also the thermometer phonons and electrons are systems with independent temperatures. As the thermometer is metallic, its electron heat capacity C_e is proportional to temperature and much larger than its phonon heat capacity, which can usually be neglected. In fact, the electron specific heat capacity of the thermometer is much larger than the specific heat capacity of the crystal. The thermometer does not dominate the overall heat capacity only because of its small size. In typical detectors, C_e and C_c are of a similar magnitude. Both heat capacities are coupled to the thermal bath at constant temperature T_b . For the crystal, this coupling G_{cb} is accidental, arising from the mechanical connections of the crystal. It is often small and can be neglected. For the TES electrons, the coupling G_{eb} is part of the detector design and implemented as a thin gold stripe wire-bonded to the detector holder. It sets a timescale in detector relaxation and is sized deliberately. There is an indirect thermal coupling between the crystal phonons and TES electrons, G_{ec} . It arises as a series combination of the Kapitza coupling G_K between TES and crystal phonons, and the electron-phonon coupling G_{ep} . At the extremely low operating temperature (~ 10 - 20 mK), this parameter becomes highly suppressed (electron-phonon decoupling). This presents a challenge to cryogenic calorimetry: the only quantity measured in TES readout is the thermometer electron temperature T_e . If T_e decouples from the crystal temperature T_c , the detector becomes blind to the temperature rise from particle events in the crystal. In fact the main sensitivity of TES macrocalorimeters at such extremely low temperatures does not rely on thermal phonon transmission, but on a prompt athermal signal.

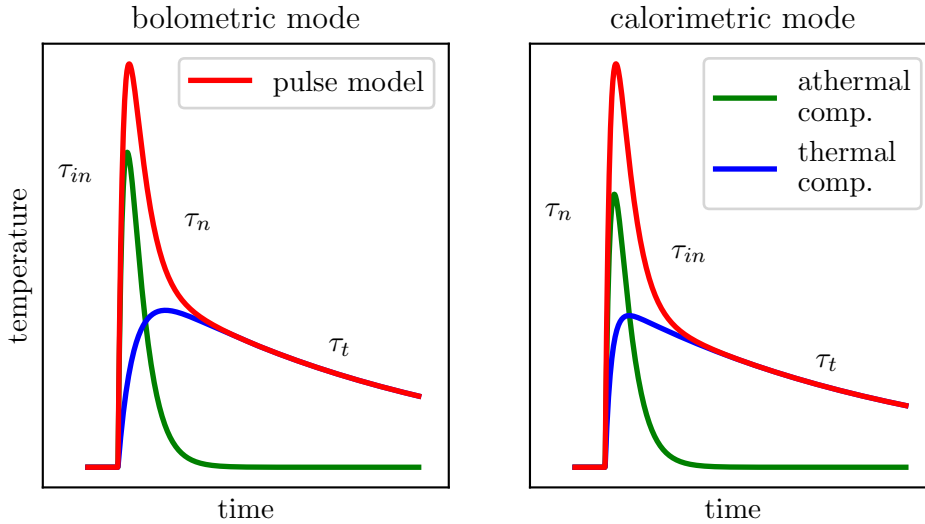


Figure 2.6.: Pulse models following equation 2.11, for the bolometric ($\tau_{in} < \tau_n < \tau_t$, left) and calorimetric ($\tau_n < \tau_{in} < \tau_t$, right) operation regimes. The relevant time-scales are marked on the pulses in both cases. The pulse model (red) is a sum of the athermal component (green) and the thermal component (blue).

In a particle event in the crystal, energy first reaches the crystal lattice in the form of high-frequency optical phonon modes. These immediately decay (within nanoseconds) into acoustic phonons which then move through the crystal. Acoustic phonons are subject to anharmonic decay, a one-to-two process in which the energy is distributed over a population of daughter phonons. As the anharmonic decay rate is strongly frequency-dependent ($\propto \omega^5$), this process is initially very fast, but essentially stops once phonons reach frequencies of some hundreds of GHz. These phonons free-stream ballistically through the crystal, carrying energies (corresponding to few K temperature) far above those of the thermal phonon background in the crystal. These athermal phonons eventually (after some ms) thermalize via inelastic processes on the crystal surfaces. A fraction of the crystal surface is covered by the TES film, where athermal phonons are efficiently absorbed and split cooper-pairs, raising T_e . This is depicted by the two power inputs $P_e(t)$ and $P_c(t)$ in Figure 2.5. The direct input of power into TES electrons by athermal phonons bypasses the bottleneck of G_{ec} and causes the leading signal contribution in the presence of electron-phonon decoupling. With athermal phonon lifetime τ_n and film-thermalization fraction ε , the power inputs can be written:

$$P_e(t) = \Theta(t) \cdot \varepsilon \cdot \Delta E / \tau_n \cdot e^{-t/\tau_n} \quad (2.6)$$

$$P_c(t) = \Theta(t) \cdot (1 - \varepsilon) \cdot \Delta E / \tau_n \cdot e^{-t/\tau_n}, \quad (2.7)$$

where ΔE is the time-integrated phononic energy input. τ_n and ε can be expressed through the competition of film thermalization with a time constant τ_f , and crystal surface thermalization with a time constant τ_c :

$$\tau_n = \left(\tau_f^{-1} + \tau_c^{-1} \right)^{-1} \quad (2.8)$$

$$\varepsilon = \frac{\tau_f^{-1}}{\tau_f^{-1} + \tau_c^{-1}} = \frac{\tau_n}{\tau_f} = \frac{\tau_c}{\tau_c + \tau_f} \quad (2.9)$$

The crystal thermalisation time τ_c depends on the crystal material, size and shape (essentially sound speed and mean-free-path between surface scatterings) and the surface properties (thermalization probability). The film thermalisation time τ_f can be expressed

$$\tau_f = \frac{1}{\bar{\eta}} \cdot \frac{2 \cdot V_a}{A_t \langle v_{\perp} \alpha \rangle} \quad (2.10)$$

as a function of absorber volume V_a , TES area A_t , and phonon absorption fraction $\bar{\eta}$. $\langle v_{\perp} \alpha \rangle$ is the mode-averaged product of phonon transmission probability α and group velocity perpendicular to the absorber-thermometer interface v_{\perp} . $\bar{\eta}$ can be calculated analytically by assuming that longitudinal phonons are fully absorbed while transverse modes do not interact with the TES electron system.

The coupled system of equations represented by the schematic in Figure 2.5 exhibits two internal time constants given by combinations of the various heat capacities and conductivities. The duration of the energy input τ_n adds a third time constant. This leads to the following form for the time evolution of TES electron temperature:

$$T_e(t) = \Theta(t) \cdot \left\{ A_n \cdot \left(e^{-t/\tau_{in}} - e^{-t/\tau_n} \right) + A_{th} \cdot \left(e^{-t/\tau_{th}} - e^{-t/\tau_n} \right) \right\}. \quad (2.11)$$

This pulse-shape can be understood as a superposition of two exponential pulses with finite rise and decay times, shown in Figure 2.6. The amplitudes A_n , A_{th} and time

constants τ_{in} , τ_t in general have a complex dependence on the system parameters. In an idealized limiting case, they can be understood in a simple picture. Neglecting G_{cb} and assuming $G_{eb} \gg G_{ec}$, the full expressions derived in [109] simplify to:

$$\tau_{in} = C_e/G_{eb} \quad (2.12)$$

$$\tau_t = C_c/G_{ec}. \quad (2.13)$$

τ_{in} is the relaxation timescale of the thermometer through the thermal link G_{eb} while the much longer τ_t is the relaxation of the absorber crystal through the bottleneck G_{ec} . Also when relaxing the assumptions, τ_{in} describes the relaxation of the TES temperature (initially heated due to the direct power input $P_e(t)$) into a temporary equilibrium between the warm crystal and the cold bath, and τ_t is the slower relaxation of the entire system back (TES and crystal) to bath temperature. The slowest timescale τ_t describes thermal relaxation after the end of athermal phonon energy input. This motivates naming the slower exponential component proportional to A_t the ‘‘thermal component’’, while the faster component proportional to A_n is the ‘‘athermal component’’. Due to electron-phonon decoupling at 15 mK, the athermal component dominates the overall amplitude. Two regimes can be distinguished depending on the order of the timescales τ_n , τ_{in} , τ_t , which is settled by detector design.

Bolometric mode: $\tau_{in} < \tau_n < \tau_t$ The TES temperature reaches its temporary equilibrium before the decay of the athermal energy input. The temperature rise is then proportional to the time-dependent input power, or athermal phonon flux. In allusion to the bolometer, an instrument measuring a flux of energy, this is called the ‘‘bolometric mode’’. The pulse rise-time is then given by τ_{in} , the fast decay time is τ_n . For the pulse amplitude one finds

$$A_n = \varepsilon \cdot \frac{\Delta E}{\tau_n G_{eb}} \quad (2.14)$$

Calorimetric mode: $\tau_n < \tau_{in} < \tau_t$ The thermometer begins to equilibrate only after the end of the athermal energy input. The initial temperature rise is therefore proportional to the deposited energy. Like a calorimeter measuring total particle energy, in the ‘‘calorimetric mode’’ the TES integrates over the athermal phonon pulse caused by a particle interaction. τ_n appears as the pulse rise-time, while τ_{in} is the fast decay time. For the pulse amplitude one finds

$$A_n = \varepsilon \cdot \frac{\Delta E}{C_e} \quad (2.15)$$

Bolometric mode was used in CRESST-II phonon detectors, to speed up relaxation times in the large 4 cm target crystals (where τ_n can be $\mathcal{O}(10 \text{ ms})$). The light channel was designed for calorimetric operation. In CRESST-III, all detectors operate in calorimetric mode. In the NUCLEUS-1g prototype, the outer veto is designed as a bolometer, while all smaller detectors (target, inner vetos) are calorimeters.

Phonon collectors Phonon collectors, or quasi-particle traps, are a possibility to enhance the signal collection for a detector in calorimetric mode. They are formed by a superconducting film with a larger gap (higher T_c) than the TES, deposited next to

the TES with some overlap. In the detectors discussed in this work, phonon collectors made of electron-beam evaporated aluminum are used and serve a secondary role as superconducting bias contact pads. The idea is that athermal phonons are energetic enough to break Cooper pairs also in the aluminum, which is far below its T_c . The broken Cooper pairs form quasi-particles which quickly relax to the edge of the superconducting band gap. Some energy is lost in this process due to the emission of subgap phonons which can only be reabsorbed in the W film. The quasiparticles then diffuse through the superconducting film and can recombine, which leads to a finite lifetime. If a quasiparticle encounters the edge of the tungsten film, the lower-bandgap material traps it. The energy of the quasiparticle will be thermalised in the TES electron system. Phonon collectors are a powerful way to enhance the athermal signal, as they increase the effective collection area of the TES without adding to its heat capacity.

2.3.2. SQUID readout




Figure 2.7.: SQUID readout circuit for one TES channel in CRESST-II, from [222]. The bias current is split over the TES and a reference resistor R_R . The current flowing in the shunt branch is coupled into the SQUID via an input coil.

Thin-film TES are low-impedance devices with resistances of a fraction of an Ω , that need to be measured to $< \text{m}\Omega$ precision. Due to heating constraints, readout currents have to be kept below few tens of μA . This limits the choice of amplifier technology to essentially a single option: Superconducting Quantum Interference Devices (SQUIDs).

SQUIDs are formed by a superconducting loop broken by two insulating Josephson tunnel junctions. For a detailed discussion of these devices see [223]. The current-voltage characteristic of a SQUID is modulated by the magnetic flux through the loop, with a periodicity of one magnetic flux quantum $\Phi_0 = h/(2e)$. A SQUID is thus an extremely sensitive magnetometer. By coupling a current into the SQUID loop using an input coil, a SQUID becomes a current meter with current noise levels as low as a few $\text{pA}/\sqrt{\text{Hz}}$. To linearize the SQUID response, a feedback coil cancels the flux signal from the detector. This holds the SQUID at a fixed working point. Figure 2.7 shows the circuit allowing a resistance measurement with a SQUID. A constant bias current is applied to a parallel combination of the resistance of interest (the TES) and a reference resistor (shunt). The fraction of the current flowing through the shunt branch is monitored by the SQUID current sensor.




Figure 2.8.: Heater pulses used for temperature stabilization and monitoring of the detector energy scale. The injected amplitudes are 0.1 V, 0.3 V, 0.7 V, 1.0 V, 2.0 V, 4.5 V, 5.0 V, 8.0 V. Left: heater pulses of various amplitudes that are continuously injected every few seconds. Right: reconstructed amplitudes of heater pulses throughout a recorded file.

2.3.3. Temperature stabilization

To ensure optimal gain and gain stability in a TES detector, it is important to control the TES temperature (or better the operating resistance) very precisely over the measurement duration. Some TES detectors use electrothermal feedback, caused by the bias current joule heating decreasing with increasing resistance in typical readout circuits. This effect is strong when the cryostat base temperature is far below the TES transition temperature, as significant joule heating is necessary to bring the TES into transition in this situation. In the detectors discussed here, the transitions are relatively close to the cryostat base temperature and electrothermal feedback is weak. The TES temperature is therefore stabilized via an actively controlled resistive heater. It is fabricated as a ~ 50 nm gold layer, and a DC heater current as well as fast pulses can be applied. The general stabilization scheme is depicted in Figure 2.8: pulses of varying amplitude are periodically injected in the calorimeter. Saturated heater pulses, named “control pulses”, serve as a direct measure of the TES position in the transition (or its operating resistance relative to the normal conducting resistance). The recorded control pulse amplitudes are the input for a PID control loop on the DC heater current, stabilizing the detector against slow temperature drifts (on a time scale of tens of seconds).

The smaller heater pulses, named “test pulses”, are spread from near threshold up to detector saturation. They serve to verify both linearity and time stability of the detector energy scale. The relative injected energy in the heater pulses is known a priori, and the reconstructed test pulse amplitudes over a dataset (Figure 2.8 right) can be used for an energy scale linearization in post-processing.

2.3.4. Scaling of energy resolution with target mass

The energy resolution of a cryogenic detector is determined by many parameters. Some of them can be manipulated by the designer, others are material quantities or technical limitations that are hard to improve upon. Examples are choice of sensor technology, target material, readout electronics, sizes of the target and thermometer, thin-film layouts, mechanical and thermal parameters of the detector environment. In addition to the energy resolution, many other figures-of-merit change in the optimization (such as pulse time constants or detector bandwidth, readout power, minimal process size required) which may pose challenges to fabrication, handling and operation. Therefore, optimizing cryogenic detector energy thresholds is a multi-parameter process, and (given the timescales of fabrication and testing) inherently slow. Here we explore a simple shortcut to detector optimization, previously derived in [221, 224], which offers an interesting trade-off between energy resolution and target mass. This scaling relation between target mass and energy resolution has been used to increase the sensitivity of CRESST-III to low-mass dark matter, and its further extrapolation leads to the gram-scale cryogenic detectors developed for CE ν NS detection within NUCLEUS. In that sense, this scaling relation is the foundation for the two experiments discussed in the rest of this work.

In this section, the change of cryogenic detector energy resolution as a function of the size of the target is considered. We use a strongly simplified approach, looking only at the temperature rise achieved in the thermometer for a given energy deposition. This glosses over many aspects, such as the level and frequency dependence of thermometer noise and the signal power spectrum, which may also change with the target size. The simplifications are primarily justified by a good fit of our result with past observations from detectors.

In a thermal detector, the thermometer is assumed to be in equilibrium with the target. Ideally, the target heat capacity dominates, so that the temperature rise $\Delta T = \Delta E/C_a$ is given by the deposited energy and absorber heat capacity. The absorber heat capacity scales with volume or mass of the target. Assuming the energy resolution to be proportional to the target mass, we find $\sigma_E \propto m$.

For a cryogenic detector primarily sensitive to athermal phonons, basic considerations yield a different relation. In these detectors, electron-phonon decoupling at low temperatures result in a weak thermal link between absorber and thermometer. Therefore the temperature rise in the absorber is mainly caused by ballistic athermal phonons, which can be absorbed directly in the thermometer electron system.

A particle interaction, be it electron or nuclear recoil, results in a number of displaced nuclei near the interaction site (plus electronic excitations not relevant to our thermometer on the crystal surface). The displaced nuclei can be expressed as a linear combination of phonon states with frequencies up to the Debye frequency of the lattice. These high-frequency phonons quickly experience anharmonic decay $\propto \omega^5$. The strong frequency dependence of this 1-to-2 phonon process effectively means that after a few generations of phonon splitting (with on average, frequencies divided in half and decay times increased by factors of 32), the phonons “freeze out” and start free-streaming ballistically through the crystal. Depending on lattice parameters, these phonons have mean frequencies corresponding to temperatures of several K, i.e. they are far out of thermal equilibrium with the crystal. This population of athermal phonons is typically

thermalized on crystal surfaces over a timescale of tens or hundreds of milliseconds. This marks the epoch of energy collection by the thermometer.

In the picture of a homogeneous ensemble of athermal phonons scattering off the crystal surfaces, energy collection becomes a competition between phonon absorption in the thermometer and thermalization on the crystal surfaces. The fraction of athermal phonons absorbed rather than thermalized is then $\varepsilon = (\kappa_f A_f)/(\kappa_c A_c)$, where $A_{f,c}$ denotes thermometer film and absorber crystal surface areas and the κ are material parameters independent of geometry. For the temperature rise, one gets:

$$\Delta T = \varepsilon \cdot \frac{\Delta E}{C_f} = \frac{\kappa_f A_f}{\kappa_c A_c} \cdot \frac{\Delta E}{C_f} \propto \frac{1}{A_c}. \quad (2.16)$$

The only geometrical factor remaining is the crystal surface area A_c . The thermometer film surface A_f cancels, as the film heat capacity C_f is proportional to it. A larger thermometer collects more athermal phonons, but its increased heat capacity leads to an unchanged temperature rise. (It is assumed that the thermometer film thickness is fixed by fabrication constraints.) As the dependence of target surface area and mass on linear dimension d are $A_c \propto d^2$ and $m \propto d^3$, the final result for the scaling of energy resolution in athermal detectors is:

$$\sigma_E \propto m^{2/3} \quad (2.17)$$

The derivation considers only the scaling of thermometer temperature rise. The result reflects the fact that the signal is produced in the bulk (volume) of the target, but detected on its surface. It has to be stressed that many possibly relevant factors are ignored here, which amount to the capability of maintaining an equally performing thermometer independent of the target size.

Taking the simple scaling law at face value, we can derive predictions from past detectors. Figure 2.9 depicts masses and energy thresholds of detectors operated in various phases of the CRESST experiment. CaWO_4 detectors correspond to green symbols, while Al_2O_3 detectors are shown in blue. The solid prediction lines follow $m^{2/3}$, while the dashed line, computed for thin plates of constant thickness [221], approaches m^1 . This difference originates in the mean free path of a phonon between two surface scatterings. In the same way, the relative factor ~ 0.77 between Al_2O_3 and CaWO_4 comes from different sound speeds and transmission probabilities at the target-thermometer interface. The single free parameter, the common normalization of all prediction lines was fitted to the CRESST-I and CRESST-II data points (blue cross, green triangles). Both the energy thresholds reached by CRESST-II light detectors (blue circles) and CRESST-III phonon detectors (green diamonds) scatter around the expected range. This shows the predictive power of the scaling model for different target materials, sizes and geometries.

For CRESST-III, the scaling relation predicts energy thresholds below the 100 eV design goal for the 24 g target crystals. For NUCLEUS (as discussed in section 6.1.2) energy thresholds in the 10-20 eV regime are required. The scaling relation suggests 4-5 gram detector targets. For NUCLEUS, a target size of 5 mm were chosen (for reasons discussed in section 6.1.2), leading to a predicted threshold of 4.4 eV (Al_2O_3) and 7.5 eV (CaWO_4 , red stars in Figure 2.9). This should be understood as a ballpark estimate only, or as a performance goal to be achieved after careful optimization.




Figure 2.9.: Masses and energy thresholds of CRESST detectors, with scaling model. Solid lines follow $m^{2/3}$ for omnidirectional scaling, dashed line shows m^1 for disks of constant thickness. The normalization of all prediction lines was fitted to the data points above 200 g. Blue cross from CRESST-I (4 cm)³ sapphire [225], green triangles from CRESST-II [226, 227] (4 cm CaWO₄ cylinders), blue circles from CRESST-II light detectors (4 cm sapphire disks) [221], green diamonds from CRESST-III [228] (2 cm × 2 cm × 1 cm CaWO₄ cuboids). Red stars are predictions for thresholds of (5 mm)³ Al₂O₃ and CaWO₄ cubes (4.4 eV and 7.5 eV).

All CRESST detectors at LNGS were selected through preliminary measurements at MPP, and operated in a highly optimized experimental setup regarding stability of the environment and electronics noise conditions. The scaling relation anchored at CRESST detectors thus serves as a guideline for what performances are potentially achievable, and has been used for high-level detector design in both CRESST and NUCLEUS.

3. The CRESST experiment

The Cryogenic Rare Event Search with Superconducting Thermometers (CRESST) experiment has been pursuing the direct search for dark matter since the late 1990s. The experiment is located at the underground laboratory of Laboratori Nazionali del Gran Sasso (LNGS) in Abruzzo, Italy. The first generation of the experiment, CRESST-I [229], used Al_2O_3 crystals with only one readout channel (phonons). For the second generation experiment CRESST-II, the scintillating target material CaWO_4 was introduced and a cryogenic light channel added, allowing for discrimination of electronic recoil background. The current experimental stage CRESST-III is focusing on low mass dark matter by improving the nuclear recoil energy threshold to unprecedented sensitivity.

This chapter describes the deep-underground experimental facility and the working principles of CRESST detectors. The results and status of the ongoing data-taking campaign are outlined.

3.1. The CRESST Facility




Figure 3.1.: Flux of cosmic muons as a function of depth in major underground laboratories [230].

As mentioned in the previous chapter, direct detection experiments must be exquisitely shielded from environmental backgrounds as well as sensitive to small energy depositions. The CRESST facility at LNGS is carefully designed to provide a low-background environment for the experiment, as well as the ultra-low temperatures required for highest




Figure 3.2.: The CRESST installation at LNGS. Left: overview of the underground laboratory and location of the CRESST installation. Right: cut-away view of the CRESST setup at LNGS, showing the cryostat and multi-layer shielding surrounding the detectors. Figure reproduced from [231].

sensitivity.

The first step for virtually all dark matter searches is hiding from cosmic-ray background. The Earth's surface is constantly showered with particles created in high-energy collisions in the atmosphere. The most penetrating component are muons, which arrive at a rate of about $1 / \text{cm}^2 / \text{min}$ at sea level. The solution is to go deep underground. Figure 3.1 shows the muon flux as a function of depth, measured in several underground laboratories. LNGS is covered by ~ 1300 m of limestone (equivalent to ~ 3800 meters of water), and the muon flux is reduced by six orders of magnitude, to about one per m^2 per hour [232].

The layout of the underground lab and the location of CRESST is shown in Figure 3.2 (left). The experiment is contained in a three-story building, providing access for detector mounting at the ground floor, cryogenic operations at the second floor, and auxiliary lab space including a clean room at the top floor. The experimental setup itself occupies the first two floors and is shown in a cross-section view in Figure 3.2 (right). The upper center contains a wet dilution refrigerator (Oxford Instruments Kelvinox 400) providing a base temperature below 6 mK. This temperature is transferred to the detectors located at the bottom center via a system of copper cold fingers. The detectors are displaced from the coldest point of the cryostat to allow for material control and shielding. The cryostat itself cannot be made sufficiently radiopure to directly house the detectors.

The shielding of CRESST consists of an onion-shell structure of layers suppressing backgrounds from different sources. An outer layer of 40 cm polyethylene thermalises

neutrons produced by natural radioactivity in the rock or spallation by the remaining cosmic muons. A layer of plastic scintillators acts as a veto against direct hits by cosmic muons on the experiment. Inside of this, an airtight box flushed with nitrogen (the “radon box”) guards against deposition of radioactive radon decay products on the surfaces of the experiment. A 20 cm layer of lead, partially installed inside the cryostat at low temperature, covers the sides and the top of the experiment. This efficiently suppresses background from environmental gamma rays due to the high density and atomic number of lead. The radioactive isotope ^{210}Pb itself causes background, so an inner layer of 14 cm high-purity copper constitutes the innermost massive shielding layer. An additional 3.5 cm of polyethylene installed inside the cryostat copper vessel blocks neutron background generated in the massive lead shield.

An integral part of the experiment is its “nervous system”, the sensitive electronics that make it possible to read out signals from the cryogenic detectors. The CRESST setup is equipped with 66 DC-SQUID amplifiers mounted inside the liquid helium reservoir [222]. The circuit shown in Figure 2.7 is implemented with bias lines filtered at 4 K and shunt resistors mounted at the mixing chamber plate (< 10 mK).

3.2. The phonon-light technique




Figure 3.3.: Neutron calibration data from Run 34’s Detector A displayed in the light-yield vs. nuclear recoil energy plane [92]. The light yield is the ratio of detected light to phonon energy, and normalized to 1 for e^-/γ events of the calibration source. The blue lines show the 90% band for e^-/γ events, red and green are the nuclear recoil bands for oxygen and tungsten, respectively. The calcium band is in between and not shown for clarity. All bands are heavily populated by the AmBe neutron source. Powerful discrimination of electron and nuclear recoils is possible above ~ 1 keV recoil energy.

The ultimate goal of any dark matter search is to be background-free in the region of interest (ROI). This can be achieved in two ways: background suppression (e.g. by shielding) and background identification (i.e. gathering information and more narrowly defining the ROI). The massive shielding constructed for CRESST is not sufficient to

become background free. An important residual background comes from β and γ radiation from radioactive impurities in the cryostat, shielding and the detector materials itself. For detailed studies on radioactive backgrounds in CRESST-II see [233, 234]. The continuous energy spectra of these backgrounds extend to low energies, causing events in the dark matter ROI. Fortunately, there is a way to distinguish e^-/γ interactions from dark matter scattering. The former create electron recoils with long track length inside the crystal compared to short-range nuclear recoils. In consequence, electron recoils cause a lower ionization density than nuclear recoils. In CRESST, this can be used by employing scintillating crystal targets: highly ionizing interactions cause less light output per energy, as scintillation centers in the crystal become saturated (light quenching). Measuring independently the phonon energy and scintillation light of a particle interaction therefore allows to distinguish electron from nuclear recoils. An example from CRESST-III is shown in Figure 3.3: the light yield, defined as the ratio of measured light to phonon energy, strongly differs for electron recoils (blue) and nuclear recoils (red, green). The CaWO_4 material chosen as the scintillator in CRESST was identified in a compromise between radiopurity, light output and phonon properties [235, 236]. The light detector in CRESST uses the same technology and readout as the phonon channel: it is implemented as a thin silicon-on-sapphire wafer equipped with a TES, held independently and facing the main absorber crystal inside a reflective housing [237].

3.3. The CRESST-III low-mass dark matter search

After the introduction of the phonon-light technique in CRESST-II, the strategy of CRESST-III is to push the nuclear recoil threshold as low as possible. This allows unprecedented sensitivity to low-mass dark matter (below a few GeV/c^2). In this region of parameter space experiments with a small target mass like CRESST can make important contributions. The “classic WIMP” parameter space above $\sim 10 \text{ GeV}/c^2$ is more effectively explored by tonne-scale experiments with higher thresholds. The main change from CRESST-II to CRESST-III is a reduction of the target mass in each module, from 300 g to 24 g. Reducing the target mass increases thermal and athermal signals for a given energy deposition, allowing to lower the energy threshold (see also the discussion in section 2.3.4).

3.3.1. Detector modules in CRESST-III

The detector modules deployed in CRESST-III are the product of lessons learned in the previous generation of the experiment. The best sensitivities to dark matter in CRESST-II were reached by the modules Lise [227] (which reached the lowest energy threshold) and TUM40 [226] (which combined low internal radioactivity with a fully scintillating housing). In addition, a fully active surface veto was developed in the so-called “beaker modules” [239]. In the first measurement campaign of CRESST-III, Run 34, ten identical standard modules (labeled A-J) and three upgraded beaker modules were operated from May 2016 to February 2018.

The standard modules are shown in Figure 3.4. The concept of the standard module implements the “stick design” of TUM-40 for the smaller ($2 \text{ cm} \times 2 \text{ cm} \times 1 \text{ cm}$) target crystals. The fully scintillating housing avoids an important problem of previous CRESST-II detectors: α background on the detector-facing surfaces can lose most of its




Figure 3.4.: The standard detector module of CRESST-III. Left: schematic drawing with labeled individual constituents (adapted from [92]). Red squares represent TESs on instrumented components. Right: photograph of an opened detector module. Phonon and light detector TES are facing the viewer.




Figure 3.5.: The beaker module of CRESST-III. Left: silicon beaker suspended in the bottom half of the copper holder. Center: schematic cross-section, adapted from [238]. Right: top view on carrier disk and ring with TES.

MeV-scale energy in the surface before escaping and depositing the rest of its energy in the detector. This background can mimic nuclear recoil signals due to the partially undetected energy. Imperfect modeling of this surface background is thought to be responsible for the excess observed in previous CRESST-II detectors [240, 241]. By making all surfaces facing the detector scintillating, surface backgrounds emit an additional light signal and are clearly distinguished from real nuclear recoils. Compared to CRESST-II stick modules, the CRESST-III fully scintillating housing features an upgrade: instrumented holding sticks (iSticks). The stick is a dielectric crystal in mechanical contact with the main absorber. After energy depositions in the stick, phonon transmission into the main absorber is possible. By equipping the sticks with their own TES, such holder-events can be vetoed.

The beaker modules (shown in Figure 3.5) feature a large-area light detector in the form of a milled silicon beaker. This serves a dual purpose: strongly enhanced light-collection efficiency, and an active surface veto [238]. The weak point of beaker modules

in CRESST-II was the composite design: the top of the module consists of a large disk featuring the TES (“carrier”), glued to the main absorber with epoxy glue spots. The purpose of this layout is to close the 4π -veto by distinguishing carrier events and main absorber events via their different pulse shape (the method is applied to CRESST-II composite detectors in [242]). This beaker module design suffered from low-energy dark counts attributed to thermal stress in the glue spot. The upgraded beaker design addresses this challenge through the use of a third channel: the carrier ring. While the carrier disk is strongly coupled to the target using a large interface of silicon oil, the ring is connected with small glue spots, providing mechanical support and closing all line-of-sight from the main absorber to the outside. The pulse height ratio between carrier and ring then allows to distinguish events originating in the holder from main absorber events.

3.3.2. CRESST-III low-mass dark matter results




Figure 3.6.: Energy thresholds achieved in the standard modules of Run 34, the first measurement campaign of CRESST-III [228].

Of the ten standard detector modules (A-J) operated in Run 34, three (A, E, J) performed well with all channels (phonon, light, iStick) fully operational. Three more modules (B, C, D) performed well on the phonon and light channels but had some non-working iSticks. The remaining four modules were discarded due to their insufficient energy resolution or readout problems on phonon or light channel. All three beaker modules were not fully operational due to separate issues with their electronics channels (not related to detector design).

In CRESST-III, the analysis profits from a continuous DAQ [243], which allows trigger optimization after the completion of data-taking, and use of the optimum filter technique already at trigger level. With this technique, four of the iStick modules reached or exceeded the goal of 100 eV nuclear recoil threshold [228] (see Figure 3.6). The best-performing detector A reached 30.1 eV nuclear recoil energy threshold and was used to place new limits on low-mass dark matter parameter space [92]. The corresponding limit curve (solid red) is drawn in Figure 1.9. The energy spectrum observed by detector A is reproduced in Figure 3.7. All events triggered in the phonon channel are shown in grey.




Figure 3.7.: Energy spectrum observed by detector A in Run 34 [92]. Grey: all events triggered in the phonon channel. Red: accepted events passing a light-yield cut designed against e/γ backgrounds. Cosmogenic activation of ^{182}W produces lines at 2.6 keV and 11.27 keV.

The red events are selected for their light-yield compatible with nuclear recoils. This step is designed to remove e/γ background. Between threshold and ~ 200 eV, a large number of events with a falling exponential spectrum are observed. This background is limiting the sensitivity of the CRESST experiment to unexplored low-mass dark matter parameter space.

The data analysis of Run 34 is described in greatest detail in two recent PhD theses [244, 245]. The former focuses on detector A, detailed cross-checks on the characteristics of the low-energy background, and classification of events by origin in the module. The latter features analyses of the modules A, J, B and E, showing that the background is present in all of them, but with incompatible rates and spectral shapes. This excludes a common astrophysical origin of the observed events.

3.3.3. Exploration of low-energy background

With the sensitivity of the experiment limited by the presence of an unknown low-energy background, the following measurement campaign, Run 35, was devoted to studies of its possible origins. Eleven modified standard modules were operated from 11/2018 to 10/2019 along with three beaker modules. The standard modules feature modifications such as polished surfaces, copper clamps instead of CaWO_4 holding sticks, removal of scintillating foil or new crystal materials: commercial CaWO_4 , CaWO_4 grown at TU Munich and two modules featuring an Al_2O_3 target. Two of the three beaker modules could be operated successfully. At the time of writing, data analysis of Run 35 is ongoing and initial cryogenic operation of Run 36 is well underway.

4. Active magnetic field compensation for CRESST

4.1. Magnetic fields and superconducting films

Superconductivity is named after the complete loss of electrical resistance some metals exhibit at low temperature, as first discovered by Kamerlingh Onnes in mercury in 1911 [246]. The connection of superconductivity to magnetism was shown in 1933 with the observation of the Meissner-Ochsenfeld effect [247]. This phenomenon of flux expulsion from a superconducting sample proved crucial in understanding the underlying mechanism of superconductivity as a macroscopic electron condensate. The interrelation of superconductivity and magnetic fields thus has a long history. Also in the context of TES devices, their sensitivity to external magnetic fields is a long-standing subject of study [248].

While the transition of an ideal bulk superconductor is perfectly sharp, real thin films such as TES show a finite transition width. Many effects play a role in the gradual onset of resistance, such as critical temperature non-uniformities, current densities approaching critical values, magnetic fields induced by currents or external magnetic fields [249]. One mechanism involving magnetic fields is the motion of flux tubes (in type II superconductors) or small normal regions (in type I superconductors). These can be created either by an external field or thermal excitation close to the transition temperature. A bias current across the TES creates a Lorentz force acting on the flux-carrying region. Any resulting motion of magnetic flux perpendicular to the current flow induces voltage parallel to the current flow, thus creating resistance. This simplistic picture illustrates the connected roles of magnetic fields, applied current and resistance near a TES' transition temperature. Only flux penetrating the TES plays a role in suppressing superconductivity, so that only the field component normal to the thin film is expected to contribute.

Effects of magnetic field on TES transition slope and high-frequency noise has been reported for TES microcalorimeters [250]. While many experiments and test facilities for TES devices use careful magnetic shielding (e.g. [251]), the CRESST experiment and test facilities at MPP have been built without magnetic field suppression¹. A passive magnetic shielding, integrated in the experiment, would remove magnetic field effects but also the possibility to study them. Also, constraints in material radiopurity and cryogenic design limit the implementation of such a shielding in the running experiment. Using instead an active field compensation, consisting of coils at room temperature and outside the low-radioactivity shielding, the impact of magnetic fields can be systematically studied. In addition, a time-dependent field cancellation can be implemented. This chapter describes an exploration of magnetic field effects on TES performed at

¹The SQUID sensors in the readout circuit, which are highly sensitive magnetometers, are enclosed in niobium or cryoperm shields to avoid flux pickup.

the cryogenic characterization setup at MPP Munich. Next, the implementation of an active field compensation for CRESST at LNGS is presented, with studies to show its benefits to experimental sensitivity both by static compensation of the background field and active suppression of magnetic transients.

4.2. TES response to transverse fields

4.2.1. Measurement principle

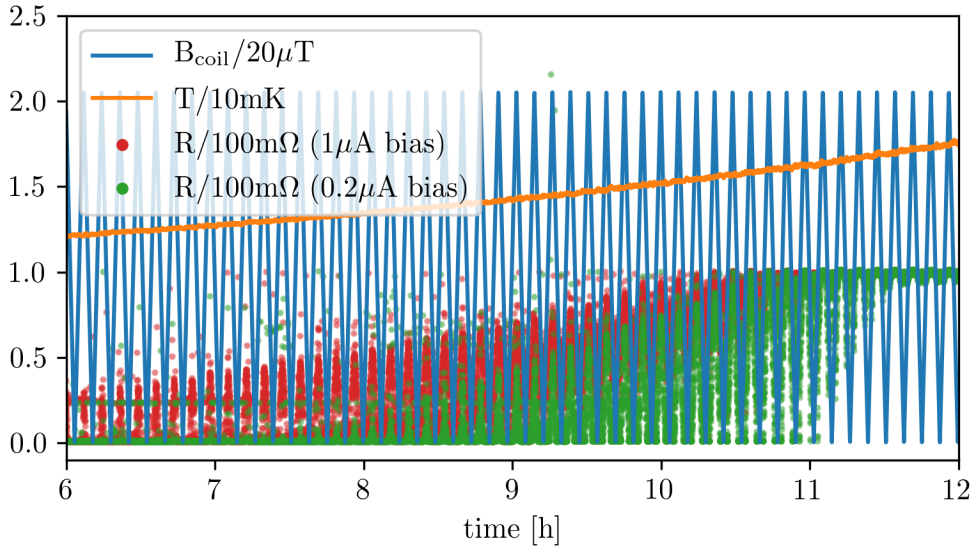


Figure 4.1.: Transition measurement of a TES (WI-594/2M) varying temperature, bias current and magnetic field simultaneously. The plot shows mixing chamber temperature and applied magnetic field as a function of time since start of the measurement, and the resulting TES resistance values at two different bias currents.

To study the effect of magnetic fields on the TES produced at MPP, one setup in the cryogenic test lab at MPP was equipped with a Helmholtz coil made from copper wire. The coil consists of four windings per pair on the outside of the cryostat dewar (46 cm diameter, $0.2\ \Omega$ total resistance), generating $16\ \mu\text{T}/\text{A}$ at the central location, with good homogeneity expected across the sample space (e.g. 2% field strength variation over a central cylinder of 15 cm height and diameter). A current of only a few ampere is needed to compensate magnetic fields on the order of the Earth magnetic field, so a simple lab power supply with analog programming was used.

The variables of interest affecting the TES resistance are the bath temperature, bias current and applied magnetic field. The resistance is measured by sequentially sending symmetric current steps of different sizes to the bias circuit, and recording the SQUID response. The switch from $-I_B$ to $+I_B$ allows to remove the unknown offset of the SQUID amplifier. Simultaneously, the magnetic field is ramped up and down continuously with a time scale of around 5 minutes. The bath temperature is raised slowly by the mixing chamber heater to ensure thermalization of the sample to the mixing




Figure 4.2.: Transition measurement of a TES (WI-594/2M) varying temperature, bias current and magnetic field simultaneously (same as Figure 4.1). Horizontal axis zoomed in to show a small number of field sweeps.

chamber. After sweeping the TES temperature in a sufficient span to observe the whole transition, the recorded data can be selected and displayed in various ways. Figure 4.1 shows a three-dimensional transition sweep recorded using this strategy, as a function of time. Figure 4.2 zooms in on a smaller number of magnetic field-sweeps over part of the transition. This method was used to study the magnetic response of a number of tungsten films of different origin and layout.

4.2.2. Observation of a large-area sputtered tungsten film

One of the samples characterised as a function of temperature, current and magnetic field is a large (6 mm × 8 mm) rectangular tungsten film produced by magnetron sputtering at Max Planck Institute for Plasma Physics in Garching [220]. The goal of the cryogenic measurement was primarily the determination of transition temperature (not operation as a detector). Therefore the TES was measured as a large film, without phonon collectors or a weak thermal link. The 20 mm × 10 mm × 0.4 mm silicon substrate was strongly coupled to the thermal bath via gold wirebonds and a large spot of vacuum grease.

The result of the measurement is displayed in Figure 4.3. The three panels correspond to different bias currents, the TES resistance is color-coded. The emerging picture can be summarized in a few key observations:

- the transition temperature is highest at an applied field of $B_0 = 28 \mu\text{T}$. This is expected for zero transverse field at the TES, i.e. when the applied field cancels the background field.
- for other applied field values, the transition temperature is shifted down by about 1 mK for every $20 \mu\text{T}$ deviation from B_0 .




Figure 4.3.: Resistance of TES IPP_runI.005 as a function of mixing chamber temperature and applied magnetic field. Resistance is color-coded between 0 and 100 mΩ. The panels show the different bias currents (left to right): 0.1 μA, 0.55 μA, 1.0 μA.

- the dataset is symmetric around a horizontal line at 28 μT, reflecting the fact that the TES does not distinguish between positive and negative transverse field values.
- the width of the transition is < 1 mK and is independent of magnetic field and bias currents within the observed ranges.

The local perpendicular field at the TES location, visible as the vertical offset of the symmetric pattern, is not in line with a predicted 43.7 μT vertical Earth magnetic field component expected at MPP, from the current World Magnetic Model [252, 253]. The deviation could be explained by the magnetic field of the lab building or by local distortions from magnetic materials inside the setup. For the purpose of studying and mitigating the effects of magnetic fields on TES, the best estimate of the local transverse field comes from the TES itself.

A shift of the transition temperature as a function of (transverse) magnetic field is predicted by the temperature-dependence of the critical field of a superconductor. The empirical relation [223]

$$H_c(T)/H_c(0) = 1 - (T/T_c)^2 \quad (4.1)$$

for ideal type-I-superconductors predicts a value of $-\frac{1}{2} \cdot T_c/B_c(0)$ for the slope of the transition temperature as a function of magnetic field near zero field. For tungsten with $T_c = 15$ mK and $\mu_0 \cdot H_c(0) = 120$ μT [254], one therefore expects a shift of transition temperature by 1 mK per 16 μT of applied field.

This expectation is close to the observed magnetic-field-induced transition shift. The large-area sputtered tungsten film thus behaves as expected for an ideal (bulk) type-I-superconductor and shows a simple, current-independent response to a transverse magnetic field.

4.2.3. Observation of a structured evaporated tungsten TES




Figure 4.4.: Resistance of TES WI-594/2M as a function of mixing chamber temperature and applied magnetic field. Resistance is color-coded between 0 and 100 m Ω . The panels show the different bias currents (top left to bottom right): 0.2 μA , 0.6 μA , 1.0 μA , 1.4 μA .

A similar dataset was recorded with a completed TES (WI-594/2M), fabricated by electron-beam evaporation at MPP. The corresponding detector is the 5 mm sapphire cube that will be discussed later in section 8.3.4. The TES consists of a $72 \mu\text{m} \times 144 \mu\text{m}$ tungsten rectangle between two $210 \mu\text{m} \times 320 \mu\text{m}$ Al phonon collectors. A thin Au stripe ($20 \mu\text{m} \times 3 \text{mm} \times 50 \text{nm}$) provides a weak link to the thermal bath. The layout is depicted in Figure 8.3.

The full dataset is shown in Figure 4.4, where mixing chamber temperature is plotted against applied magnetic field, with the color of the markers encoding the TES resistance. Each panel displays the data recorded for one bias current. Although the general behavior is recognizable, more features appear compared to the large-area sputtered tungsten film:

- compensation of the local transverse field happens around 15 μT . The offset compared to TES IPP_runI_005 can be explained by changes in the local background field or in the setup in the three months that passed between the two measurements.

- the width of the transition is smallest when the applied field cancels the background field, and increases with local transverse field
- at higher bias currents, the transition appears at lower bath temperatures. This shift is larger for higher residual magnetic fields.
- at $1.4 \mu\text{A}$, the highest observed bias current, a periodic pattern in the transition temperature appears as a function of magnetic field, with a pronounced maximum around the compensated field value

Additional TES mounted in the same run in close proximity and oriented parallel to each other show the same symmetry-offset, i.e. they agree about the value of the local transverse field. In addition, a TES on another detector was mounted in an orthogonal orientation, so that the applied magnetic field vector falls in the film plane of the TES. On this sensor, no magnetic field dependence on the transition was observed, showing that in fact only the transverse field component affects the behavior of the superconducting phase transition.




Figure 4.5.: Transition curves of TES WI-594/2M at background magnetic field and compensated magnetic field, extracted by selecting data points taken around $0 \mu\text{T}$ (coil off) and $15.5 \mu\text{T}$ (field compensated). The solid lines are Bezier curves intended as guides for the eyes. Outlying data points above the transition curve can be caused by particle interactions spuriously increasing TES resistance during the bias current step. The transition at compensated magnetic field is much sharper than the corresponding “coil off”-transition, and stays sharper also for increasing currents, which highlights the interplay of transverse magnetic field and bias current.

From the full dataset, some simpler views can be extracted as slices. Figure 4.5 shows transition curves at constant magnetic field (i.e. horizontal slices of Fig. 4.4) for $0 \mu\text{T}$ and $15.5 \mu\text{T}$ (i.e. background field or “coil off” and compensated field or “coil on”). The

transition temperature is increased by 1-2 mK with compensated field, and the transition width is sharper (around 0.5 mK from 10% R_n to 90% R_n instead of > 1 mK with “coil off”). Furthermore, the transition width for compensated field does not increase with larger bias currents, contrary to the situation with background magnetic field.

The apparent current dependence of the “coil on” transition can be explained by the bias heating of the TES. The TES temperature is increased with respect to the (measured) bath temperature, thus the transition appears shifted to lower temperatures. This can be taken into account by changing the plotted variable from bath temperature T_b to film temperature T_f using

$$T_b = T_f + \frac{I_B^2}{g} \cdot \frac{R_s R_{TES}}{(R_s + R_{TES})^2} \quad (4.2)$$

The TES electron-bath thermal coupling g (formed by the thin gold stripe) can be estimated from the heat capacity of the TES electrons and the fast decay time of observed pulses, in this case one gets roughly 0.5 pW/K. Taking this number, bias-current self-heating explains well the current dependence of the “coil on” transition, while the “coil off” transition shows a current dependence beyond self-heating. In general, the picture in Figure 4.4 changes by less than a mK through self-heating. More importantly, self-heating is magnetic field-independent and thus doesn’t affect any of the statements above.




Figure 4.6.: Resistance of TES WI-594/2M for 1.4 μA bias as a function of mixing chamber temperature and applied magnetic field. An empirical Fraunhofer pattern (white markers) is overlaid on the data, showing qualitative agreement with the transition temperature-magnetic field dependence for a periodicity of 4 μT .

The periodic pattern observed at the highest bias current is reminiscent of a Fraunhofer single-slit interference pattern. Such a dependence (known as superconducting

quantum interference) is expected for the critical current in a spatially extended Josephson junction, as a function of magnetic field across the tunnel barrier [223]. In general, the magnetic-field dependence of the critical current is the Fourier transform of the current density across the junction area, averaged along the direction of current flow [255]. This relation has been used to study the current density distribution in exotic superconductors, to reveal localized conducting quantum states [256].

For the case of a tungsten TES close to transition temperature, the appearance of weak-link behavior can be explained by inhomogeneities in the transition temperature across the TES area. With large parts of the TES in the normalconducting state, most current is carried by narrow superconducting paths, which are driven close to transition and can therefore form localized tunnel barriers. Figure 4.6 shows a single-slit interference pattern superimposed on the 1.4 μA data, following the empirical formula:

$$T(B) = 13.5 \text{ mK} + 2.0 \text{ mK} \cdot \frac{\sin\left(\pi \cdot \frac{B-B_0}{B_1}\right)}{\pi \cdot \frac{B-B_0}{B_1}} - \frac{1 \text{ mK}}{12 \mu\text{T}} \cdot |B - B_0| \quad (4.3)$$

with a field offset $B_0 = 14.0 \mu\text{T}$ and periodicity $B_1 = 4.0 \mu\text{T}$. The periodicity, corresponding to one flux quantum per $(23 \mu\text{m})^2$, gives a scale for the size of the weak link area. Exploring the appearance of superconducting quantum interference patterns may be a way to characterize homogeneity of superconducting films for TES application and, while not directly related to interesting physics phenomena, a useful quality control measure.

The finished TES shows a stronger and more complex magnetic field response than the large-area tungsten film. Apart from the weak thermal coupling adding the phenomenon of self-heating and the fabrication method (evaporation vs sputtering) that may affect the film homogeneity, also the Al phonon collectors can contribute to this difference. As superconductors far below T_c they behave perfectly diamagnetic and deflect transverse magnetic flux through the action of persistent eddy currents. This displaces field lines onto the tungsten area (flux focusing) and may exacerbate the impact of a transverse magnetic field.

In summary, the intrinsic transition temperature and sharpness (in zero-field conditions) of CRESST TES is reduced significantly by the presence of background fields of tens of μT such as the Earth magnetic field. As detector performance depends on the stability and noise as well as transition temperature and slope, the magnetic field cannot be ignored when setting up and operating a detector. These observations demonstrate that magnetic fields need to be controlled at least on a μT level, to ensure optimal or even reproducible conditions for running a detector.

4.3. Fields and coils at CRESST

4.3.1. The magnetic environment at LNGS




Figure 4.7.: Fluxgate data recorded for one file during Run 34 in June 2017, along with trigger rates of two detectors. For clarity, offsets of (30.4 μT , 1.1 μT , 11.4 μT) are added to the magnetic field components. Four short periods of changed magnetic field can be identified, one of which coincides with a high-rate period in detector E-Ph.

In parallel with the magnetic field studies at MPP, the influence of magnetic fields at the CRESST setup at LNGS was studied in greater detail. The list begins with a singular event that helped focus the attention on magnetic fields as a limit to detector stability, and continues with more details learned subsequently.

CUORE lead shield The CUORE Experiment [257], direct neighbor of CRESST in Hall A at LNGS, moved its outer shielding (including 70 tonnes of lead) from storage to operational position in April 2017. CRESST detectors could not be operated at the pre-established thresholds due to excessive noise for about two weeks. The situation could be remedied by installation of mu-metal panels between CRESST and CUORE.

As a reaction, a three-axis fluxgate magnetometer (APS Model 520) was installed in proximity to the CRESST experiment on May 29, 2017. With its help, a number of

smaller sources of regular external magnetic disturbances could be identified. Figure 4.7 shows the magnetic field components measured by the fluxgate over a two-day period between cryostat refills, along with the trigger rates of two detectors.

Faraday cage door During cryostat refills and other operations in the CRESST building, the Faraday cage door at the middle floor is opened. This can be identified on the fluxgate traces. Usually data taking is interrupted while people are present at the setup.

Experimental hall doors The large gates of Hall A are opened for truck access, i.e. for weekly liquid nitrogen deliveries. The magnetic disturbance associated with this coincides in Figure 4.7 with noise periods in detector E-Ph. The latter can be due to the vibrations of truck movement rather than the magnetic influence.

Parked vehicles lab personnel and companies supporting the infrastructure have been observed parking cars next to the CRESST building, leading to a magnetic disturbance of similar magnitude than the hall doors. Without 24/7 surveillance of the surroundings of the experiment, it is difficult to identify all such disturbances and remove the corresponding periods from data.

Hall crane A reproducible observation of likely magnetic influence from an external source comes from the positioning of the Hall A crane running along the ceiling (passing over the CRESST building). On March 14, 2019, the crane was moved from above the CRESST experiment to a parking position at the end of the experimental hall. Figure 4.8 shows the magnetic field change recorded by the fluxgate and the test-pulse pattern of some detectors over a larger time span including the crane move. The fluxgate recorded a magnetic field change of the order of $0.5 \mu\text{T}$ in each component. Some detectors are unaffected, while others show a permanent change in the testpulse pattern. This excludes disturbance by vibration or other features of the operating crane, and can be understood as follows: the magnetic field change influences the transition shape and temperature. Any change in the control pulse height due to this is corrected by the control loop through the detector heater, effectively changing the TES temperature to achieve the same operation resistance in the new transition. The difference of the transition shape (slope, curvature) around this new operation point with respect to the old one results in changing amplitudes for non-saturated test pulses. As the test pulses map out the non-linear response of the detector, it is not surprising to see some test pulse amplitudes more strongly affected than others. The change in the testpulse pattern (amplitudes and possibly resolutions) is an indication that detector energy resolution may have changed. This underlines that precise magnetic field control is necessary for stable detector operation.

4.3.2. Magnetic field compensation setup

The observations of magnetic field effects on TES made at MPP and LNGS point towards a measurable improvement in detector sensitivity and stability for CRESST, achievable by canceling magnetic fields near the detectors.

One effective method of magnetic field control is the use of passive magnetic shielding, for example mu-metal, cryoperm or superconducting shields. A combination with

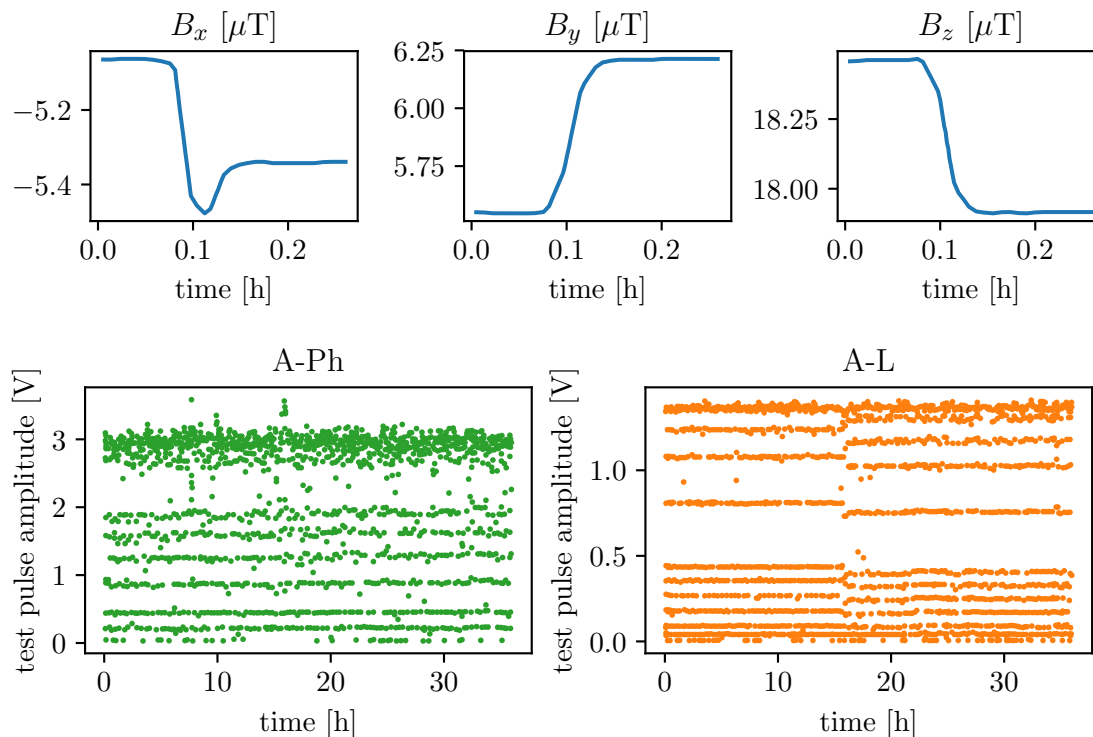


Figure 4.8.: Effect of the Hall A crane on CRESST detectors. The crane, initially parked above the CRESST experiment, was moved to the far end of the experimental hall. Top row: magnetic field readings of the fluxgate sensor during the crane move. Bottom row: detector test pulse patterns from the corresponding file. The crane move occurs around hour 16. Some detectors are unaffected (A-Ph, left), while others (A-L, right) show a lasting change in testpulse amplitudes, i.e. a distortion of the energy scale.

an active field compensation using electromagnets can be useful to avoid flux frozen into the superconductors [251]. CRESST is facing the special situation of having to retrofit a running cryogenic experiment with a magnetic field compensation system. There are constraints on available space, power dissipation and radiopurity of the introduced materials. The running schedule of the experiment has to be adapted to get access for modifications. As a first step, passive shielding was therefore avoided and instead a combination of air-coils and magnetometers were installed around the running setup, outside the lead/copper shielding and the air-tight radon box. This has the advantages of avoiding a modification of the cryostat and reduced cleanliness demands. The modification can take place with the experiment running, and an active compensation allows to study magnetic field effects in greater detail before fixing the design of a final magnetic compensation.

Compensation coils At the beginning of Run 35 in May 2018, the fluxgate magnetometer was moved to a closer location directly under the detector carousel, as shown in Figure 4.9. At the same time, three pairs of rectangular coils were added between

the radon box and muon veto surrounding the CRESST setup. The rectangular coils represent an approximation to the ideal Helmholtz geometry using the available space. The coils have to fit around the radon box with a square cross-section of $1.4 \text{ m} \times 1.4 \text{ m}$. The vertical positions of the z-coils can be freely chosen, so that a square Helmholtz configuration can be achieved. The coils for generating the horizontal field components have a fixed distance and maximum lateral extent. Since the distance is too large for a Helmholtz configuration, the most homogeneous field inside the cryostat is reached by making the coils as large as possible (symmetric around the target region). For the given situation, this results in z-coils with $1.4 \text{ m} \times 1.4 \text{ m}$ cross-section and 76 cm spacing, and x,y-coils with $1.4 \text{ m} \times 1.5 \text{ m}$ cross-section and 1.4 m spacing.

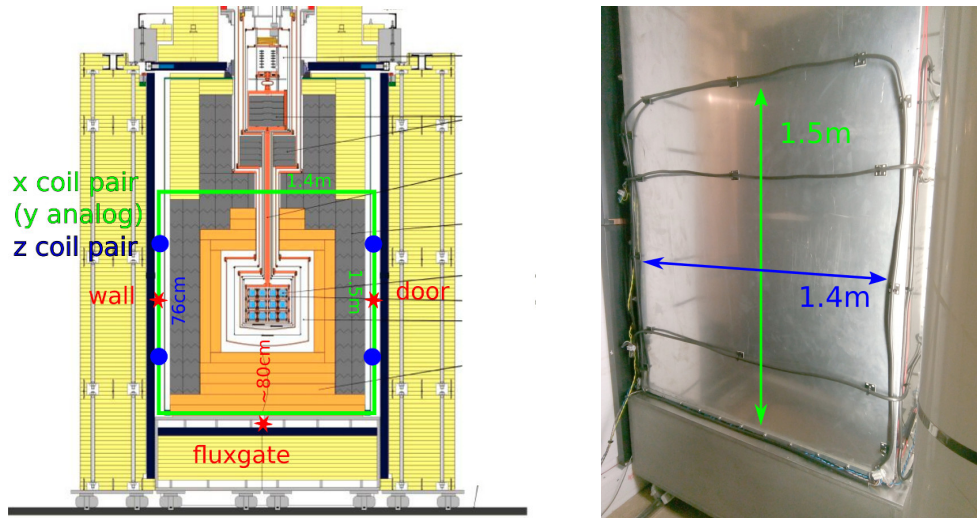


Figure 4.9.: Left: schematic cross-section of the CRESST setup showing the locations of the APS Model 520 fluxgate magnetometer and the compensation coils since the beginning of Run 35 in May 2018. Additional magnetometer positions (“wall”, “door”) are indicated in red. Right: photograph of the setup with opened polyethylene shielding, showing the radon box with compensation coils.

Calculation of field distributions for rectangular coils are outlined in Appendix A.1. The results are shown in Figures 4.10 and 4.11. The x- and y-coils achieve a homogeneity of 3.5% over the control region (a $30 \text{ cm} \times 30 \text{ cm}$ cylinder), while the z-coil achieves homogeneity better than 0.3% .

The “coils” are formed by a single loop of standard 5-pole 16 A cable, with the poles connected in series to form five windings. They are expected to generate $3.3 \mu\text{T}/\text{A}$ (x,y) and $5.7 \mu\text{T}/\text{A}$ (z) at the central location. A current of 10 A per pair is assumed to be sufficient to generate the necessary range of B-values to study the compensation of background fields. As each coil pair has a resistance of about 0.5Ω , no high voltages are required for operation. The coils are driven by analog-programmable DC power supplies (Delta Elektronika ES 015-10), with an output range of 0-15 V and 0-10 A. Given that the power supplies are unipolar, the correct polarity has to be manually assured. As long as the dominant background field is the nearly constant Earth field, no sign changes on the currents are expected to become necessary.




Figure 4.10.: Contour plots of magnetic field strength calculated for rectangular coil geometries for x and y directions. The plot is a central horizontal cut showing the field strength by the x -coil pair. Magnetic field strength is scaled to the nominal value at the center of the setup. The left panel shows the complete setup. Adjacent contours are spaced by 4% of the nominal field strength. The right panel zooms in on the control region, contours are spaced by 0.5%. The outline of the control region, a 30 cm \times 30 cm cylinder (corresponding to the carousel), is shown as a dashed line.

AMR magnetometers The coils were set up and first operated after the Run 35 data taking period, in June 2019. At this time, also additional magnetometers were installed at symmetric locations in the setup (“wall” and “door” in Figure 4.9). These magnetometers, based on anisotropic magneto-resistance (AMR) sensors and controlled by Arduino electronics, were designed as a simple unit allowing magnetic field mapping when deployed as an array. As opposed to the fluxgate, which allows an extremely accurate magnetic field measurements, the additional magnetometers were intended to provide additional, though less precise, magnetic field data from many locations. This is important since the magnetic field cannot be directly measured at the detector location in the cryostat and thus has to be obtained through a model. Field inhomogeneities on a scale larger than the setup can be interpolated using many magnetometers distributed in the vicinity. Naturally, a fluxgate array would be optimal. A less costly solution was pursued for this explorative phase of the magnetic-field study.

The AMR devices use the Honeywell HMC5883L three-axis magnetometer chip containing magnetoresistive elements measuring in a range of $-200 \mu\text{T} \dots +200 \mu\text{T}$ with a nominal precision of $0.2 \mu\text{T}$. The magnetometer readings are digitized on-chip in a 12-bit ADC, and sent to the Arduino via the I²C protocol. The Arduino then communicates with three identical MCP4725 12-bit DAC boards through a TCA9548A I2C multiplexer. This allows outputting the three magnetic field components simultaneously as analog voltages. Figure 4.12 shows an open view of one such magnetometer. The voltage regulator for the 9 V barrel input is on the left, followed by the Arduino Nano (with an external USB connection for programming, power supply and serial data readout used during development).




Figure 4.11.: Contour plots of magnetic field strength, analogous to Figure 4.10, here for z direction. The plot is a central vertical cut showing the field strength by the z -coil pair. Adjacent contours are spaced by 4% (left) and 0.1% (right) of the nominal field strength.

Five units were produced, allowing gradient measurements in all directions by deploying them e.g. north, south, east, west and below the detector carousel (the direction above is occupied by the cryostat and therefore inaccessible). In a first step, two magnetometers were deployed, characterized, and integrated in the compensation system. As the positions “wall” and “door” are located symmetrically opposite the detector carousel, a simple average of the two magnetometers extrapolates to the carousel location taking into account field inhomogeneities to linear order in all components.

4.4. Static field compensation

The magnetic compensation system was set up at LNGS and characterized during June 2019. The following section describes the characterization of the coils and background field without using the magnetometers. The setup of the active compensation using the magnetometers is described in the next section.

4.4.1. Identifying the zero-field configuration

The first exercise during setup of the field compensation is finding the current settings necessary to shield the static background field at the detector carousel. As seen in the test setup in Munich, the TES themselves can provide a clear signature of the magnetic field configuration. At CRESST, this becomes a three-dimensional problem, as the TES are not oriented along any coil axis, and furthermore are oriented in different directions. The normal field can therefore be compensated for all TES only by cancelling all magnetic field components separately.

A schematic of the distribution and orientation of detector modules in the CRESST carousel is depicted in Figure 4.13. The 11 detector modules A...L share the same basic layout (depicted in Figure 3.4) of a $20\text{ mm} \times 20\text{ mm} \times 10\text{ mm}$ target crystal equipped




Figure 4.12.: Arduino-controlled magnetometer based on the three-axis anisotropic magnetoresistive (AMR) sensor Honeywell HMC5883L. The sensor chip communicates with the Arduino Nano through an I2C interface. The analog output happens via three DAC modules that map the sensor range of $-200 \mu\text{T} \dots +200 \mu\text{T}$ to $0 \dots \sim 5 \text{ V}$.

with a TES reading the phonon-channel (Ph-TES), and a thin light detector carrying the LD-TES. The cuboid modules, with the exception of C and L, are arranged in a two-floor bar structure. The top-middle position is unoccupied, modules C and L are instead located elsewhere in the carousel. Detector module C is special in that its orientation is rotated by 90° so that the Ph-TES plane is horizontal. Module L was not considered in this study. The modules Gode 1,2 follow a different design using two independent phonon channels and a beaker light detector (shown in Figure 3.5). All 3 TES in a module (C “carrier” and R “ring” on the target crystal, L “light” on the beaker) are horizontal.

The TES can therefore be sorted into three groups by orientation: Gode 1,2 as well as C-Ph (horizontal), LD-TES (except C,L), Ph-TES (except C,L). Representatives from each groups were selected and their response to different magnetic field components studied. Similar to the experiment at MPP, different current steps were injected while the magnetic field components were slowly varied. The detector temperature was kept constant, as for purposes of field nulling only the symmetry of the response, or the minimum of resistance as a function of field, is of interest. The detector heaters were used to bring each TES close to its respective transition temperature while the carousel remained colder.

Before starting the systematic search of the nulling current configuration, some explorative trials were used to establish the necessary polarity on each coil pair. The initial guess from magnetometer readings were quickly confirmed by observing resistance trends in response to individual coil ramps. It was necessary to flip the polarity of the y-coil pair with respect to the coordinate axis, such that the positive coil control voltage creates a $-B_y$ magnetic field.

As the groups “LD-TES” and “Ph-TES” are not oriented along a coil axis, the horizontal null configuration has to be found in a 2d search. The vertical field component can be optimized independently, as all TES are either normal or parallel to this direction, and should therefore either be insensitive to a vertical field or insensitive to all other




Figure 4.13.: Schematic view of detector locations and TES orientations within the carousel for Run 35. The modules follow two general designs: the standard modules with a cuboid shape and two-channel phonon-light readout (11 modules named A...L), and the cylindrical three-channel beaker modules (two units named Gode 1,2). Module C is rotated by 90° so that its phonon TES lies horizontal. Module L is not considered here. All three TES in the beaker modules are oriented horizontally. Also marked are geographic north, the direction toward the CUORE experiment and the wall/door magnetometers, as well as the reference frame (x,y,z) derived from the coil pairs. The x-axis, which points along the experimental hall, lies at 128° azimuth, i.e. points roughly southeast, according to the OpenStreetMap service [258].

field directions. Figure 4.14 shows the response of selected TES to a sweep of B_z , with no current on the x,y-coil pairs. For the sweep, the z-coil pair was ramped from 0 A to 10 A at a rate of 1/15 A/s and back again. The resistance changes appear delayed with respect to the upwards and downwards field ramps. This can be attributed to eddy currents excited in the massive copper shielding of CRESST, which dissipate with a time constant of several seconds. By shifting the data from each TES separately in time (by 3.8 s for C-Ph, 5.2 s for Gode 1-C, 8 s for Gode 1-L and 12 s for A-Ph), the TES response from both ramps can be brought to a good match. To record a clear signature, each TES should ideally be heated into transition such that it is superconducting only for vanishing transverse magnetic field. The required coil settings for vanishing field are unknown initially, so that TES are instead heated to about half their normal resistance with no currents applied to the coils.

Figure 4.14 shows, as expected, that the sensitivity of the vertical TES A-Ph to B_z is much weaker than that of C-Ph, Gode 1-C and Gode 1-L. The three horizontal TES show a clear response to B_z with apparent minima. More unexpectedly, the minimal resistance of Gode 1-C appears at a $\sim 50\%$ larger applied field than the minimal resistance of C-Ph and Gode 1-L. Different possible reasons can be found to explain this observation:




Figure 4.14.: Response of selected TES to a sweep of vertical magnetic field. The observed size of applied bias steps is plotted against the control voltage on the z-coil pair current source. A fixed heating power was applied to each detector to bring it into its transition. The upwards ramp (x-markers) and downwards ramp (dot-markers) have to be shifted in time to coincide. 1 V of applied coil control voltage translates to 2 A coil current, resulting in $11.5 \mu\text{T}$ vertical field at the carousel.

1. Significant misalignment of some TES, then sensitive to other field components
2. Inhomogeneity of the applied field across the locations of the different TES
3. Inhomogeneity of the background field across the locations of the different TES
4. Intrinsic preference for different transverse fields at minimal resistance

A physical motivation for 4 is hard to imagine. Option 1 is the simplest explanation, but hardly possible since Gode 1-C and Gode 1-L are held together in the same module. Significant relative tilt of the detectors is excluded by the module geometry. Possibilities 2 and 3 appear unlikely given the spatial proximity of Gode 1-C and Gode 1-L, which are only about 4 cm apart in the vertical direction (the length of the beaker). The coil inhomogeneity should not exceed a few percent over the whole carousel. A probable scenario could be a small magnetic object near Gode 1-C, locally distorting the background field. As an ad-hoc compromise, a z-current of 2.0 A (control voltage 1.0 V, vertical field $11.5 \mu\text{T}$) was chosen.

After fixing the z-coil current, a two-dimensional sweep in x and y determines the full minimal field configuration. Figure 4.15 shows step sizes acquired from four TES as a function of applied x and y fields. A continuous ramp on the y-coil is superimposed with a slower ramp on the x-coil (i.e. the sweep progresses as a zig-zag line from left to right). A time lag is applied to the data to bring the resistance observed in successive rising and falling y-ramps to alignment. A low-resistance “valley” gives away the applied field configuration for which the field component normal to the TES plane is small.




Figure 4.15.: Response of differently oriented TES to a magnetic sweep in x- and y-directions. The plot axes display the x- and y-coil control voltage. The colors encode the range between minimal (blue) and maximal (red) resistance for each TES. The white markers show B_x - B_y combinations approximately canceling transverse fields for phonon (circles) and light (diamonds) detectors.

The direction of the valley in field-space is related to the geometric orientation of the TES films. This allows to measure the relative orientation of TES and local coil-field directions, a possibility explored in Appendix A.2. The crossing region of the valleys of differently oriented TES marks the compensated field configuration in x and y. Again it is observed that the two representatives of each TES orientation group, while agreeing in valley direction (i.e. orientation), do not agree precisely about the valley offset or normal field value at minimal resistance. Note that the width of the valley is arbitrary, depending on the chosen fixed detector heater value during the sweep. Compromises are again found in judging the valley center lines, indicated with white markers in Figure 4.15. The crossing point of the two lines sets the chosen minimal field configuration (determined on June 17, 2019): $(U_x, U_y, U_z) = (0.51, 0.44, 1.00)$ V, corresponding to an applied field of $(B_x, B_y, B_z) = (+3.4, -2.9, +11.4)$ μ T. Rotated into the coordinate system of the World Magnetic Model (components pointing north, east and down), we get a background field of $(+4.4, -0.9, +11.4)$ μ T. This is, like the vertical field in the

setup at MPP, incompatible with the World Magnetic Model [252, 253], which predicts $(+24.2, +1.5, +40.1)$ μT at the geographic location of LNGS. We have to assume that magnetic materials in the hall surrounding the experiment strongly distort and in this case shield much of the earth magnetic field.

4.4.2. Transition measurement




Figure 4.16.: Comparison of transitions of phonon TES recorded with coils off (darker markers) and coils on (lighter markers). The field-induced temperature shifts are 0.1 mK or smaller, matching the observation that transverse background fields for phonon TES were small due to their orientation. The transition temperatures of A-Ph and F-Ph are higher in zero-field, D-Ph is unaffected. I-Ph and J-Ph even show broader transitions with coils on. A possible trend between observed transition change and detector position hints at insufficient field homogeneity.

Having determined a current configuration to compensate the background magnetic field at the CRESST experimental site, its impact on detector operation can be visualized using a transition measurement at both field configurations. For a transition measurement at LNGS, all detector heaters are switched off and the carousel heater is used to change the temperature of all detectors with respect to the cryostat base temperature. Temperature is measured using a Speer-type carbon resistor (GS4) directly mounted on the carousel frame. The current applied to the heater is slowly increased over a timescale of several hours. Comparing an upwards and a downwards ramp the lag between TES temperature and GS4 temperature is confirmed to be below 0.05 mK. With the coil currents set to $I_x = 1.08$ A, $I_y = 0.59$ A, $I_z = 2.00$ A (the first null configuration later updated to 1.02 A/0.88 A/2.00 A), the transition measurement was repeated overnight.

Figures 4.16, 4.17 and 4.19 show the results grouped by TES orientation. The phonon detector transitions in Figure 4.16 are shifted by 0.1 mK or less.




Figure 4.17.: Comparison of transitions of light detector TES recorded with coils off (darker markers) and coils on (lighter markers). Detector H-L is suffering from a series resistance in the bias circuit. All transitions are shifted up between 0.5 mK (at the top) and over 1 mK (at the foot of the transition), and become steeper by around a factor of 2. Detectors A-L and F-L show nearly identical behavior in both magnetic configurations.

The fact that phonon TES are weakly affected by the coils can be understood through a detailed look at their orientation. With the compensating coil field pointing about 12° west of north (azimuth 348°), Figure 4.13 confirms that this vector happens to lie near the plane of the phonon TES. Equivalently, the valley in Figure 4.15 (upper row) passes very close to the origin of the plot, meaning “coils-off” produces a situation with small normal fields for TES oriented this way. Phonon TES in Run 35 happened to be mounted in the correct orientation to avoid strong magnetic effects by the background field.

Looking back to the disturbance by the crane in Figure 4.8, we can see that also its field change was roughly parallel to the phonon detectors. This explains their reduced sensitivity to this magnetic interference, purely by lucky orientation.

A trend of the phonon TES behaviour with detector position can be observed: from A-Ph and F-Ph at one end of the carousel (see Fig. 4.13), whose transition temperature increased, to I-Ph and J-Ph at the other end of the carousel, whose transition appears to get broader with field compensation. This points to a position-dependence in the quality of the compensation, i.e. either a local background field gradient, or insufficiently homogeneous compensation field.

Light detector TES, oriented orthogonally to the phonon TES, are affected more strongly by the background magnetic field. The transitions in Fig. 4.17 are consistently shifted up by around 1 mK and get steeper by about a factor 2. This shows the potential of improving sensitivity and stability using field compensation.

The beaker module, with three very different TES structures (see Fig. 4.18) sharing




Figure 4.18.: TES layouts used in the beaker module. Tungsten layers are shown in blue, aluminum in gray, gold in yellow. The carrier (C) uses a large-area bolometric TES, the ring (R) a smaller calorimetric design, while the beaker (light detector, L) is equipped with a very short TES (0.1 mm) including large (1 mm \times 1 mm) phonon collectors.

the same horizontal orientation, allows an interesting observation. The large-area carrier TES (Gode 1-C) with only narrow Al stripes for contacting, behaves similar to the sputtered TES in section 4.2.2. The transition is shifted by 0.4 mK, but the shape stays the same. The ring TES (Gode 1-R) with smaller tungsten film and larger Al phonon collectors, shows a noticeable sharpening along with the temperature shift. The beaker TES (Gode 1-L) shows the most dramatic magnetic field effect of all observed TES: the transition temperature shifts by more than 2.5 mK at the foot, and a large kink in the transition vanishes in zero-field conditions. This trend among TES layouts fits with the expectation from flux focusing by the superconducting Al films. By Meissner effect, persistent eddy currents prevent any magnetic flux from penetrating the Al films, and instead displace the field lines into neighboring regions. The field in the tungsten gap between the Al phonon collectors is therefore enhanced by a factor related to the relative width of aluminum and tungsten.




Figure 4.19.: Comparison of transitions of module Gode-1 recorded with coils off (darker markers) and coils on (lighter markers). Widely different behaviour is observed, from a transition purely shifted by 0.4 mK (C) over a larger shift and markedly increased sharpness (R) to a shift of over 2 mK and feature in the transition removed by field cancellation (L). This can be explained by flux focusing by the different Al structures in the TES layouts.

4.5. Active magnetic field control

4.5.1. Magnetometer noise and sensitivities

Static compensation is the first step in controlling magnetic fields. To stabilize the magnetic field over long times, canceling transients and drifts, a network of magnetometers is needed to update the desired current configuration. Mathematically, the active field compensation can be described as a problem in linear algebra. The magnetic field components \vec{B} measured by the magnetometers behave as

$$\vec{B} = \vec{B}_0 + M \cdot \vec{I} \quad (4.4)$$

given a set of background field components \vec{B}_0 and coil currents \vec{I} . The matrix M describes the field sensitivity of each magnetometer channel to current in each of the actuator coils. It uniquely characterizes the compensation system and can be determined from a geometric model of the setup, or empirically by driving current ramps in each coil. A linear regression for each magnetometer channel and each coil current gives the coefficients of M .

The data from such a series of current ramps is shown in Figure 4.20. Each column corresponds to a current ramp on one coil pair. Each of the rows shows the responses of one three-channel magnetometer, except for the last row showing the input voltage on the three coils.

The dataset can be used to characterise magnetometer resolution. The fluxgate uses

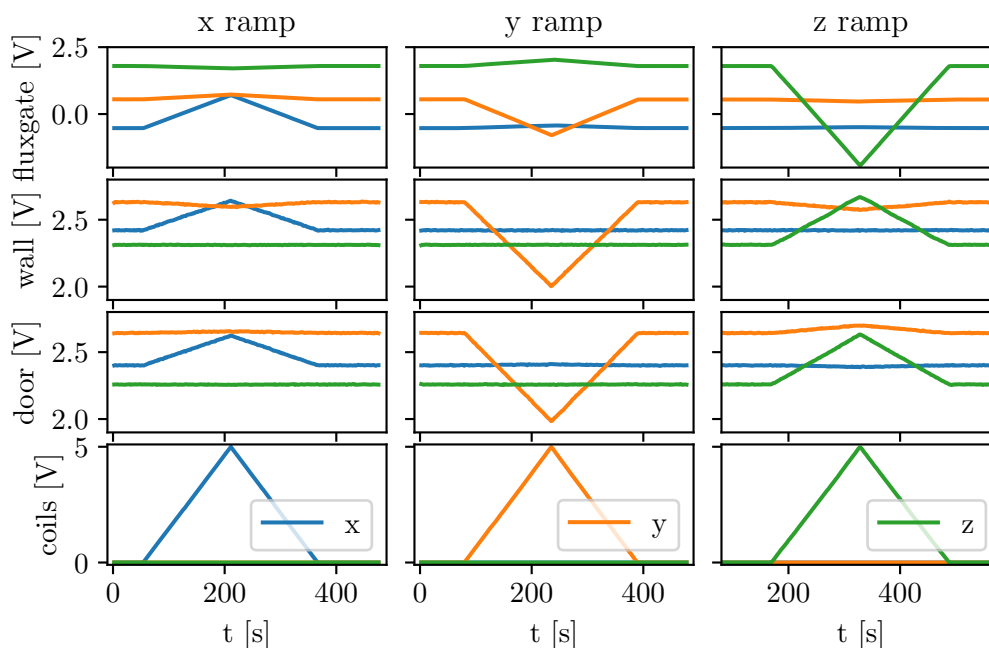


Figure 4.20.: Magnetometer responses to current ramps on the three coil pairs (left column: x, middle column: y, right column: z). The response slopes of the nine magnetometer channels resulting from the three coil ramps can be reduced to a 9×3 matrix describing the compensation system.

bipolar output with a scaling of $10 \mu\text{T}/\text{V}$ and a noise level of 0.4-0.5 mV rms. The AMR magnetometers map the sensitive range of $-200 \mu\text{T} \dots +200 \mu\text{T}$ into the output range $0 \dots 5 \text{ V}$, giving a scaling of $80 \mu\text{T}/\text{V}$. The observed noise level is higher at about 1-2 mV rms. The simply constructed AMR magnetometers thus operate with their specified sensitivity of $0.2 \mu\text{T}$ rms resolution. The sophisticated (and more costly) fluxgate probe (rms resolution $< 5 \text{ nT}$) easily beats their performance by more than an order of magnitude.

As intended, each coil predominantly affects the corresponding field component in the magnetometers. The y-coil has been inverted (positive current generates negative fields) to counteract the background field with the unipolar supply. The z-channel of the fluxgate probe is inverted relative to the coordinate axis². The smaller “off-diagonal” magnetometer responses show the deviations from perfect positioning and alignment of coils and magnetometers.

A linear regression on all combinations of magnetometer output voltage versus coil control voltage leads to the desired response matrix, which is displayed in Table 4.1 (left). The entries are the ratios of two voltages and thus dimensionless. To cross-check the result, we can compute an expectation matrix from a geometrically simplified model of the arrangement, shown in Table 4.1 (right)³. The expected sensitivities are listed in physical units in the last column. In general, they are largest for coils whose axis aligns

²The probe orientation in its confined position does not match the coordinate axes. x-, y- and z- outputs are internally remapped, leading to an apparently left-handed probe

³The field of a rectangular coil can be analytically calculated, see Appendix A.1.

	x-coil	y-coil	z-coil	x-coil	y-coil	z-coil	in $\mu\text{T}/\text{A}$
fg-x	246.5	21.1	7.3	279.3	0	0	1.40
fg-y	35.3	-268.6	-15.1	0	-279.3	0	-1.40
fg-z	-18.3	48.5	-744.7	0	0	-688.8	-3.44
wall-x	44.2	-0.5	-0.1	41.0	0	0	1.64
wall-y	-7.4	-126.0	-11.3	0	-109.9	0	-4.39
wall-z	-0.2	0.1	72.7	0	0	61.7	2.47
door-x	44.3	1.4	-2.5	41.0	0	0	1.64
door-y	2.4	-132.8	11.8	0	-109.9	0	-4.39
door-z	-0.3	-0.3	74.8	0	0	61.7	2.47

Table 4.1.: Response matrix of the compensation system (left) derived from the data shown in Figure 4.20, compared to a geometrical model expectation (right). The units are mV/V, that is mV of magnetometer response per V of coil control voltage. In the last column, the diagonal elements of the expected matrix are given in physical units. For comparison, the effect of the coils at the detector carousel is expected to be 3.33/3.33/5.74 $\mu\text{T}/\text{A}$.

with the magnetometers (y for wall/door, z for fluxgate) and decrease with distance. At equivalent locations, the z-coil produces a stronger field than x and y due to its narrower coil spacing (Helmholtz configuration).

The measured response matrix generally follows the expectation, but a number of deviations appear.

Sensitivities The fluxgate is less sensitive to B_x and B_y , but more sensitive to B_z , on a 10%-level. This is difficult to explain via the uncertain position of the fluxgate alone, and seems to indicate that the ideal rectangular geometry does not describe the coils at this level, or that local field distortions caused by the surrounding material become important. The wall/door magnetometers agree to each other on the 5%-level, but overshoot the expected sensitivity by 15% on average. A detailed study produced a partial answer: due to limitations of the internal voltage regulator, the actual output voltage range is compressed to about 0...4.75 V. This can explain about 1/3 of the effect.

Alignment The fluxgate is quite well aligned along z, but not as good in x,y. This is expected as the probe sits on a horizontal plane without a good orientation reference. The largest “misalignment” of wall/door is in the y-response to the z-coil, which is nearly opposite in wall and door. This points to a consistent deviation from the ideal geometry, like placing both magnetometers slightly out from the radon box and below the symmetric height.

Conclusion In general, the observed magnetometer responses agree with a simple model at the 10%-level. The deviations are taken into account by the minimization algorithm through the empirical response matrix. As the desired field configuration is determined by the TES and defined through a current configuration at a certain time, the actual position, orientation, offset and field calibration of the magnetometers plays a minor role

for the quality of the compensated field configuration. The decisive parameter, as shown below, is the magnetic field resolution.

4.5.2. Field compensation algorithm

4.5.2.1. Minimizing magnetic field at magnetometer positions

If the goal of the compensation is to reduce as much as possible the average field measured by the magnetometers, the desired compensation current can be found by a linear least-squares minimization of Equation 4.4:

$$\vec{I}_{min} = - (M^T M)^{-1} M^T \vec{B}_0. \quad (4.5)$$

Applying the current configuration I_{min} guarantees the smallest possible error term $|B|^2 = |B_0 + M \cdot I_{min}|^2$. It remains the responsibility of the designer to ensure that the configuration of coils and magnetometers is such that a small error term corresponds to a favorable 3D magnetic field configuration. The method can be simply extended to approach some non-zero target field configuration \vec{B}_{tar} given by desired values at the magnetometer locations:

$$\vec{I}_{min} = - (M^T M)^{-1} M^T (\vec{B}_0 - \vec{B}_{tar}). \quad (4.6)$$

4.5.2.2. Minimizing magnetic field at an extrapolated location

In CRESST, the field has to be controlled inside the cryostat volume, while the field is measured at room temperature, with magnetometers displaced from the cryostat. One must rely on the calculated magnetic field created by the coils in the control location (i.e. the detector carousel), which can be summarized in a matrix M_c . Similarly, the field value at the control location \vec{B}_c has to be inferred from the magnetometer data. This can be done simply with an average of the measured field components (assuming a homogeneous background field), but depending on the number and distribution of magnetometers also field gradients and higher orders could be modeled. To stay general, this extrapolation operation is denoted $(\dots)_{red}$, since it is a reduction of the available magnetometer information to a single value (for each component) at the control location. The estimate of the field there becomes

$$\vec{B}_c = (\vec{B}_0)_{red} = (\vec{B} - M\vec{I})_{red}. \quad (4.7)$$

The absolute square of this estimate is minimized using the linear least-squares formula, using the control-location response matrix M_c :

$$\vec{I}_{min} = - (M_c^T M_c)^{-1} M_c^T (\vec{B} - M\vec{I})_{red}. \quad (4.8)$$

The inputs for the algorithm are the measured field configuration \vec{B} , the magnetometer response matrix M , the set of applied currents \vec{I} and a modeled response matrix for a virtual magnetometer at the control location M_c . A non-zero target field can be included as in the direct minimization above. In plain language, the contribution caused by the coils is first subtracted from the measured field configuration, this field configuration is

then averaged or extrapolated to the control location. Finally the necessary canceling field is calculated from the modeled response of a virtual magnetometer located there.

A simulation study of this algorithm can be found in Appendix A.3. Further details on its implementation as process `flux` in the CRESST control software are shown in Appendix A.4. There, it is demonstrated that only the fluxgate has the resolution needed to accurately track magnetic transients affecting the detectors. Fluxgate readings are averaged over 5 s and a “damped approach” to new current configurations is added. These measures suppress the effect of magnetometer noise fluctuations on the compensation currents.

4.5.3. Perturbation test




Figure 4.21.: Fluxgate response to an excitation of 5 A on the perturbation loop. Two trials with different compensation settings are overlaid for comparison. Dotted lines: compensation passive, i.e. fixed currents on the coils, showing only the perturbation field. Solid lines: compensation active. The coil action is visible as the slow opposite ramp after each fast field change. The reduced amplitude of the excitation at the fluxgate position corresponds to a full compensation at the detector carousel.

To validate the effectivity of the active compensation, a test with a dedicated magnetic perturber was performed. It consists of a single wire loop surrounding the CRESST setup at a distance of several meters (running along the inner walls of the Faraday cage). In order to produce fields in all three components, the loop is tilted, reaching floor level on one side and passing over the entrance door on the other side. When powering the perturbation loop with 5 A, the fluxgate showed a reading changed by

($+0.45, +0.26, -1.60$) μT . This is comparable in magnitude to the disturbance caused by the hall crane.

Figure 4.21 displays the fluxgate response during two repetitions of the perturbation test. The dotted lines show the magnetic field recorded with the process `flux` switched off, i.e. fixed coil currents. The solid lines are from a perturbation of the same amplitude but somewhat shorter duration, with the compensation active. The difference in field amplitude is caused by the partial canceling effect of the coil fields at the fluxgate location. The additional currents applied by the compensation were ($-0.15, +0.09, -0.31$) A, the fluxgate reading changed by ($-0.17, -0.11, +1.05$) μT . This can be compared with the expectation from Table 4.1, which is ($-0.18, -0.12, +1.18$) μT for these currents. The expected field induced at the carousel is ($-0.50, -0.30, +1.76$) μT , which is within 15% of the field change measured at the fluxgate.




Figure 4.22.: Test pulse pattern recorded for detector J-Ph during the comparison test with the excitation coil. The amplitude of the largest test pulse gives the operation point of the detector, the different test pulse amplitudes show visually the stability of the energy scale. The influence of the excitation coil on the stability of J-Ph is strongly reduced when the compensation is active.

As magnetic fields cannot be measured in the CRESST cryostat, the effectivity of the compensation is best judged by observing a running cryodetector. Figure 4.22 shows the test pulse pattern of detector J-Ph during both repetitions of the disturbance test, which occurred within one hour from each other. The test pulse amplitudes show the detector energy scale over time, as introduced in section 2.3.3. During the first disturbance test with passive compensation (time range marked in red), the control pulse amplitude on J-Ph changed from about 6 V to less than 4 V, indicating a large displacement of

the transition. Subsequently, the heater control loop reacted by reducing the heating, essentially adjusting the detector temperature to the new reduced transition temperature. The amplitudes of the smaller testpulses do not return to their original values, indicating that a different slope of the magnetically shifted transition. This creates a distortion in the energy scale. After switching the disturbance off, the heater returned the detector to its original operation point within 10 minutes. After the situation stabilized, the field compensation was put in active mode and the perturbation switched on again (time range marked in green). During the ramping of the coils to the new current configuration, a few control and test pulses fall outside the pattern. Once the magnetic disturbance is fully compensated, control and test pulses return to their nominal values. Compared to the perturbation with passive compensation, only a very small distortion of the energy scale remains.

4.5.4. Zero-field detector performance

As shown previously, compensating the DC magnetic field background has a sharpening effect on the superconducting transitions (see Fig. 4.17), and the active compensation mitigates the impact on magnetic transients. In combination, overall detector performance should improve with the compensation system running.

As the transition shape changes sometimes dramatically between field configurations, an independent detector optimization (bias current, operation point) is needed to ensure best performance in both configurations. In July 2019, after the commissioning of the compensation system, all detectors were re-optimized in a collaborative effort. This allows a comparison between two datasets within Run 35 using the same detectors: “coil-off” (acquired Mar. 04 - Jun. 10, 2019) and “coil-on” (acquired Aug. 02 - Sep. 04, 2019). Because of the limited reproducibility of the detector setup procedure, it is desirable to study many detectors to extract the general picture. This in turn makes it necessary to choose simple observables related to e.g. energy resolution achievable in a full dark matter analysis, without the complexity of a full-scale analysis effort.

Such a simple observable readily available for automatic analysis is provided by the test pulses that are continuously injected into CRESST detectors to ensure stability of the energy scale. For this analysis, one small (0.1 V injected amplitude) and one large (1 V injected amplitude) test pulse were selected. These amplitudes correspond to about 2 keV and 20 keV of injected energy.

In CRESST analyses, one typically finds that the energy resolution of a small testpulse (after careful energy reconstruction either via template fitting or the optimum filter) closely matches the energy resolution at zero energy, and is therefore indicative of the achievable energy threshold. Larger testpulses show a degraded energy resolution. One of the dominant factors is the beginning non-linearity (saturation) of the detector in combination with variations in operating temperature. This is addressed in the energy reconstruction techniques used in a full analysis. Skipping these steps by using only raw reconstructed amplitudes, one has a proxy for detector stability. Even without an absolute energy calibration, the relative resolution of a small testpulse is indicative of the achievable energy resolution, while the relative resolution of a large testpulse contains information on detector stability.

An example testpulse pattern comparison from detector F-Ph is shown in Fig. 4.23. Shown side-by-side are 700 h of data before and after commissioning of the field compen-




Figure 4.23.: Test pulse patterns for F-Ph, recorded with coils off (left, acquired Mar. 04 - Jun. 10, 2019) and coils on (right, acquired Aug. 02 - Sep. 04, 2019). The saturated pulse height increased, which can be an effect both of an increased bias current and an operation point deeper in the transition. The amplitudes of the 0.1 V and 1 V testpulse both increased (marked with arrows).

sation system. The amplitude of the largest testpulse is a function of bias current and operation point within the transition. Therefore, its increase does not in itself indicate an improvement in detector performance. F-Ph shows increased amplitudes in all testpulses. The rms scatter (in voltage units) is likewise increased, but the relative resolution (ratio of rms scatter and amplitude) is improved both for the 0.1 V (by about 30%) and 1 V testpulses (by over 50%). This indicates that F-Ph achieves a better energy resolution and improved detector stability with the compensation system running.

The complete list of testpulse resolutions for all operational detectors in Run 35 is given in Tables 4.2 (0.1 V testpulse) and 4.3 (1 V testpulse). The 1 V comparison contains one more detector, as the performance of G-L was insufficient to trigger the 0.1 V testpulse. In general, the testpulse amplitude is increased in the “coil-on” data (0.1 V: 14/17 detectors, 1 V: 12/18 detectors). This is also the case for the phonon detectors, whose transitions appear largely unaffected by magnetic field in Fig. 4.16. Likewise, the rms scatter is generally increased (0.1 V: the same 14/17 detectors, 1 V: 13/18 detectors). This picture alone would be consistent with the interpretation that the field compensation only enhances the transition sharpness, increasing the voltage gain, without changing the temperature signal and noise power-spectra. In that case, one would expect amplitude and rms of the testpulses to increase proportionally, their ratio remaining constant. Instead, the ratio changes in many detectors, sometimes significantly. An improvement greater than 10% is observed in 7/17 detectors for the 0.1 V testpulse (4/17 deteriorate by more than 10%). For the 1 V testpulse 7/18 detectors improve, 8/18 deteriorate.

Det.	coils off			coils on		
	amp. [mV]	rms [mV]	res.	amp. [mV]	rms [mV]	res.
A-Ph	227.5	7.44	3.27%	253.3	8.19	3.23%
B-Ph	134.6	3.76	2.79%	190.5	4.57	2.40%
D-Ph	32.85	0.986	3.00%	258.4	4.05	1.57%
F-Ph	164.0	5.91	3.60%	279.8	7.49	2.68%
J-Ph	411.4	8.12	1.97%	321.8	6.62	2.06%
A-L	40.60	1.25	3.08%	107.6	3.42	3.18%
B-L	39.01	2.31	5.92%	88.45	6.45	7.29%
E-L	52.93	2.38	4.50%	84.81	2.73	3.22%
F-L	44.10	2.14	4.85%	118.9	5.21	4.38%
I-L	16.61	0.806	4.85%	42.98	1.57	3.65%
J-L	3.23	0.49	15.2%	5.73	0.87	15.2%
L-L	53.45	1.40	2.62%	84.92	2.08	2.45%
C-Ph	86.76	3.92	4.52%	110.6	5.70	5.15%
C-L	124.0	5.80	4.68%	540.2	12.5	2.31%
G1-C	58.2	1.35	2.32%	13.10	0.44	3.36%
G1-R	97.6	1.36	1.39%	111.6	2.55	2.28%
G1-L	85.3	5.02	5.89%	55.55	2.49	4.48%

Table 4.2.: Comparison of coil off/coil on detector performance through the 0.1 V test-pulse. Amplitude, rms scatter and relative resolution are listed for detectors grouped by orientation. For the “coils on” data, color is used to denote improvement (green) or deterioration (red) larger than 10% . The first two columns are most strongly affected by the changed transition slope (temperature-voltage gain). The third column is the ratio of the first two and reveals the net effect on detector performance. See discussion in the text.

This indicates that other effects beyond the dominant increase in voltage gain play a role. For many detectors, the bias current was increased for the “coil-on” data, as critical current effects preventing this before were suppressed by the field compensation. Generally speaking, a higher bias current as well as increased transition slope will make the operating temperature harder to stabilize for the heater control loop. This effect was observed for example in detectors A-Ph, G1-C, G1-R, which suffered increased oscillations around the operation point. This has a clear effect on the 1 V testpulse rms, and can also affect the 0.1 V testpulse rms through changes in gain as a function of operation point.

Nevertheless, some detectors achieved dramatically improved performance, nearly reaching a factor of two in case of D-Ph and C-L. This shows the potential of the field compensation for improved sensitivity for low-mass dark matter search. As a lesson for the future, it is important during detector setup not to set the maximum possible bias current (but rather the lowest bias current within the window of optimum performance). Sharper transitions promise better energy resolution, but come with the responsibility of a careful detector setup.

Det.	coils off			coils on		
	amp. [mV]	rms [mV]	res.	amp. [mV]	rms [mV]	res.
A-Ph	1925	30.0	1.56%	1790	55.4	3.09%
B-Ph	1319	6.97	0.53%	1659	9.51	0.57%
D-Ph	361.7	4.37	1.21%	1291	9.78	0.76%
F-Ph	1480	12.4	0.84%	2470	9.60	0.39%
J-Ph	3094	26.7	0.86%	3921	7.81	0.20%
A-L	398.3	1.58	0.40%	889.4	14.4	1.62%
B-L	395.1	3.02	0.76%	776.5	18.5	2.38%
E-L	656.7	8.61	1.31%	620.5	10.6	1.71%
F-L	425.2	2.39	0.56%	762.4	18.3	2.40%
G-L	21.02	0.715	3.40%	68.1	2.31	3.39%
I-L	167.1	0.91	0.54%	420.6	1.87	0.44%
J-L	29.3	0.57	1.95%	53.2	1.06	1.99%
L-L	498.9	2.45	0.49%	543.7	4.44	0.82%
C-Ph	1110	7.33	0.66%	1234	5.83	0.47%
C-L	251.7	8.92	3.54%	1701	29.1	1.71%
G1-C	650.0	1.67	0.26%	139.9	0.751	0.54%
G1-R	1059	2.14	0.20%	1173	5.08	0.43%
G1-L	612.4	8.83	1.44%	588.5	3.22	0.55%

Table 4.3.: Comparison of coil off/coil on detector performance through the 1.0 V test-pulse. Amplitude, rms scatter and relative resolution are listed for detectors grouped by orientation. For the “coils on” data, color is used to denote improvement (green) or deterioration (red) larger than 10% .

4.6. Outlook

With the test results obtained at the end of Run 35, the CRESST collaboration decided to make operation with magnetic field compensation the standard procedure for future data taking. Before the start of Run 36, the coils were modified to allow easier opening and closing of the radon box. This step is necessary for insertion and removal of calibration sources (several times per run). The new coils feature plug connectors at all points where cables cross from one half of the radon box to the other. In this way, the cables remain mounted to the radon box also when opened, which increases long-term stability and reduces work in the clean room. In addition, the two x-coils (which span both halves of the radon box) were split along the radon box opening into two half-coils each. By selecting an optimal width of 16 cm for the separation of the half-coils across the split, the homogeneity of the expected field across the carousel is improved from 3.6% to 1.2%, while the expected strength is reduced by $\sim 20\%$. The improved homogeneity may help with the simultaneous field compensation at the different detector locations.

A mid-term improvement of the field compensation can be achieved by deploying a fluxgate array, for example at the positions of the AMR magnetometers symmetrically around the detector location. In this way, the measured high-resolution field values can be extrapolated to the carousel position, and linear gradients in the background field are correctly compensated. The ideal technology for magnetometry at CRESST is the cryogenic SQUID magnetometer, which reaches ultimate sensitivity and can measure in close proximity to the detectors. It can be implemented either with a superconducting flux transformer added to one of the SQUID sensors used for detector readout, or with dedicated SQUID magnetometers mounted in the carousel. Both solutions require new low-impedance wires between base temperature and 4 K, so their implementation must be coordinated with the wiring upgrade of CRESST-III Phase 2.

Alternatives to the radon box coils are under study. It appears feasible to generate the fields using coils mounted on the OVC instead. The OVC is the first vacuum-tight room temperature copper vessel surrounding the carousel, with a diameter of 56 cm. An advantage of this solution is that the much more compact coils can remain on the OVC vessel in one piece when opening the cryostat. Since the conductors are closer to the target region, the dissipated power for a given field is smaller. On the other hand, the wiring complexity needed for percent-level homogeneity across the carousel increases significantly. While a classic Barker coil design [259] or similar is sufficient for the vertical field, the coils for the horizontal field components have to be designed using numerical optimization techniques.

A less complex solution to magnetic field variations could be the use of a superconducting shield around the CRESST detectors. One of the “copper pots” surrounding the carousel can be outfitted or replaced with an aluminum layer. To avoid freezing in background field at the time of the superconducting transition, the field should be compensated by another method at that time. This can be done using coils, or by inserting a mu-metal shield around the cryostat. Mu-metal is undesirable inside the shielding during detector operation because of its trace radioactive contamination. Combined with the superconducting shield, the mu-metal can be inserted only during cooldowns and removed after reaching the transition temperature of Al. The radiopurity of an aluminum sample for a potential superconducting shield is under investigation using low-background gamma spectroscopy.

5. TES design for CRESST-III

5.1. Scaling of Transition Edge Sensors

In TES design in general, many factors can be studied, from material choice to film quality, material parameters and geometrical layout. Here the problem of the correct scale of the TES, matched to its target crystal, is investigated.

Finding the right scale of a TES for a given target crystal is a multi-layered task. As seen from Equation 2.9, the signal collection happens in a competition between the TES film and the crystal surface, which slowly thermalizes the athermal phonon population. Clearly, a high athermal phonon collection efficiency $\epsilon \approx 1$ is desirable for optimal signal height. On the other hand, scaling the TES up to cover the whole crystal surface is not the right approach. The increased heat capacity of the TES cancels out the gain in collection, such that one would expect a temperature rise independent of the TES area. Another subtlety arises because the TES film influences the duration of the athermal phonon signal (as in Equation 2.8). This effect can be understood from Equation 2.15:

$$A_n \propto \frac{\epsilon}{C_e} = \frac{\tau_c}{\tau_c + \tau_f} \cdot \frac{1}{C_e} = \frac{1}{1 + \tau_f/\tau_c} \cdot \frac{1}{C_e} \quad (5.1)$$

Making all dependences on TES area A explicit, we can write $\tau_f = \tau_{f0}/A$ and $C_e = C_{e0} \cdot A$:

$$A_n \propto \frac{1}{1 + \tau_{f0}/(A \cdot \tau_c)} \cdot \frac{1}{A \cdot C_{e0}} = \frac{1/C_{e0}}{A + \tau_{f0}/\tau_c} \quad (5.2)$$

Since the quantity τ_{f0}/τ_c has the dimension of an area, we can call it the critical area A_{crit} and find:

$$A_n \propto \frac{1}{A + A_{crit}}. \quad (5.3)$$

The critical area is the TES area at which $\tau_f = A \cdot \tau_{f0} = \tau_c$, i.e. half of the phonons are thermalized in the TES and $\epsilon = 0.5$. Equation 5.3 shows that the signal height drops with increasing TES area above $A \gtrsim A_{crit}$. For $A \ll A_{crit}$, the signal height becomes independent of A . Another way of discussing this result follows from a rearrangement:

$$\frac{\epsilon}{C_e} = \frac{\tau_n}{\tau_c \cdot A_{crit} \cdot C_{e0}}. \quad (5.4)$$

Here, the only quantity depending on the TES area is τ_n . When the thermometer dominates the phonon collection, τ_n becomes much smaller than τ_c and the signal height decreases. With only a small fraction of phonons collected in the TES, τ_n attains a value near its maximum τ_c and the signal height becomes independent of A .

In summary, the TES area should be several times smaller than A_{crit} , which is defined by the relative contribution of film and crystal surfaces to the athermal phonon collection. This derivation makes the assumption that the loss processes contributing to τ_c are

TES	l_{TES} [mm]	w_{TES} [mm]	l_{PC} [mm]	w_{PC} [mm]	l_{link} [mm]	A [mm ²]
L	1.2	3.4	2.0	8.0	0.17	36.08
M	0.85	2.4	1.45	5.7	0.61	18.57
S	0.6	1.7	0.95	4.0	1.0	8.62

Table 5.1.: Dimensions of TES layouts studied in this chapter. l_{TES} and w_{TES} are length and width of the tungsten rectangle forming the TES. l_{PC} and w_{PC} are length and width of each phonon collector. l_{link} is the length of the gold thermal link (the width is always 40 μm). A is the total area of the TES.

independent from the TES design and all energy transport within the TES happens much faster than τ_n . This ignores the potential impact of the phonon collectors which can, depending on TES design, contribute to losses and exhibit quasiparticle diffusion times slower than τ_n . This aspect is discussed further in section 5.1.3.

Finding the size of A_{crit} (which depends on crystal material, size and shape) is not straightforward, as τ_f and τ_c are not directly observable in a single detector: only their combination τ_n appears as the pulse rise time. The rise time can be determined from a noise-free template pulse created by averaging pulses from a small amplitude range in the linear regime of the detector. Fitting the pulse model from Equation 2.11 to the template pulse then gives τ_n as one of the fit parameters.




Figure 5.1.: TES geometries investigated for CRESST-III phonon detectors. Structures L (area 36.08 mm²) and M (area 18.57 mm²) were studied in [231] and deployed at LNGS in Run 34. Structure S (area 8.62 mm²) is a new development. The light detector TES (LD, area ~ 1 mm²) is shown for comparison.

To study the effects of TES size on athermal phonon lifetime τ_n , different TES geometries (shown in Figure 5.1) were tested. Their dimensions are listed in Table 5.1.

Study of prototype CRESST-III phonon detectors at MPP In [231], two 2 cm \times 2 cm \times 1 cm CaWO₄ detectors, one each with an L- and M-TES structure, were operated at the MPP cryogenic laboratory. Despite the high-rate conditions, good energy resolution could be achieved and phonon lifetimes of $\tau_n^L = 0.9681$ ms and $\tau_n^M = 1.297$ ms were mea-

interface	$\langle v_{\perp} \alpha \rangle$ [m/s]	$\bar{\eta}$	τ_f^L [ms]	τ_f^M [ms]	τ_f^S [ms]
CaWO ₄ → W	647.9	0.142	2.410	4.683	10.088
Si → W	1485.4	0.246	0.607	1.179	2.540
Al ₂ O ₃ → W	2547.5	0.290	0.300	0.583	1.256

Table 5.2.: Numerical input and results for calculation of athermal phonon collection time τ_f according to Equation 2.10, for different absorber materials and TES sizes. The values for $\langle v_{\perp} \alpha \rangle$ and $\bar{\eta}$ are taken from [260].

sured. For the interpretation in terms of τ_c and τ_f , two methods are introduced [231]¹:

Method 1: take τ_f in each crystal from the model calculation of Equation 2.10, and evaluate τ_c separately. Ideally, the same τ_c should result for the identical crystals.

Method 2: assume only that τ_f scales inversely with TES area, i.e. $\tau_f^L/\tau_f^M = A_M/A_L$, and additionally that τ_c is the same for both detectors.

The agreement between the two methods allows a statement on the validity of the assumptions (calculated τ_f in 1, identical τ_c in 2). In both cases, the phonon collection efficiency ε can be calculated and compared. For method 1, Equation 2.10 with the material constants listed in Table 5.2 yields values of $\tau_f^L = 2.410$ ms and $\tau_f^M = 4.683$ ms. This assumes that the whole area of TES and phonon collector follows the same assumptions for a CaWO₄-W interface. The resulting values for the crystal lifetime are $\tau_c^L = 1.6$ ms and $\tau_c^M = 1.7$ ms, supporting the validity of the calculation. Using method 2, a slightly longer $\tau_c = 1.8$ ms is found together with somewhat faster τ_f of 2.0 ms and 4.0 ms for L and M respectively. Overall, the methods show reasonable agreement with $\varepsilon^L = 40.0 - 47.7\%$ and $\varepsilon^M = 27.0 - 31.3\%$. As the heat capacity for M is half the value for L, but ε is reduced by less than a factor of two, a higher temperature signal (by $\sim 33\%$ [231]) is expected in the smaller TES structure. In the language of Equation 5.3, the numbers found with method 2 evaluate to a critical area:

$$A_{crit} = 31.1 \text{ mm}^2 \quad (5.5)$$

for the 24 g CaWO₄ CRESST-III phonon detector as measured with the two prototypes above-ground at MPP. This would imply that the L structure is too large, M should perform better and even smaller structures are potentially interesting. The next sections describe follow-up measurements undertaken since this study.

5.1.1. TES on CaWO₄ at LNGS

Of the ten standard modules of Run 34, three were equipped with an L-TES, seven with an M-TES. Pulse templates allowing a determination of τ_n could be found for 3×L (D, G, J) and 4×M (A, B, C, E). Detector G is the same crystal as the L-TES in [231]. At LNGS the pulse rise time was measured to be $\sim 9\%$ faster (hinting at the reproducibility of the procedure).

To interpret the data from more than two detectors, method 1 can be used in the same way. Method 2 can be generalized in two different ways.

¹Section 7.2.3: Thermalization of Non-Thermal Phonons




Figure 5.2.: Athermal phonon lifetimes measured in $2\text{ cm} \times 2\text{ cm} \times 1\text{ cm}$ CaWO_4 detectors, for TES structures M and L. Two orange data points are from prototype measurements at MPP described in [231]. Blue data points are from Run 34 detectors operated at LNGS. Lines are fits with equation 2.8. Two green data points are from test measurements for Run 35 at MPP. Dotted orange: considering only the first two prototypes. Continuous blue: all data points. Dashed blue: excluding the two detectors with longest τ_n .

Method 2a: assume identical τ_c in all detectors, and scaling inversely with area for the mean τ_f observed in the groups of different TES sizes.

Method 2b: assume the same mean τ_c in the TES size groups, and identical τ_f in each detector of a size group, with a scaling inversely with area between the groups.

To evaluate a “mean time constant”, it is practical to use the inverse average $\bar{\tau} = \langle \tau^{-1} \rangle^{-1}$ due to the form of equation 2.8. Method 2b has a better physical motivation, as it accounts for crystal-by-crystal variations in surface properties that can influence τ_c . Method 2a on the other hand allows for simple evaluation of quantities and displaying of a common scaling rule for τ_n with A , and is used in the following.

Figure 5.2 shows all measured pulse rise times τ_n for Run 34 (blue data points) and the values from [231] (orange data points) for comparison. Two additional M-TES (green data points) were measured at MPP in preparation for Run 35. They exhibit very similar rise times. The τ_n values in each TES-group agree within $\sim 20\%$, with the exception of two outliers for M (detectors B,C) with very long rise time. Interestingly, the data points from the prototype measurements fall within the scatter range, but represent the fastest M value and the slowest L-value². A fit following method 2 is shown as a dotted orange

²This is moderately unlikely ($p \sim 4\%$), but a plausible systematic effect that could influence M and L in opposite ways at MPP compared to LNGS was not identified.

line for the prototype data. The Run 34 data is interpreted using method 2a, leading to the continuous blue line, or the dashed blue line when ignoring the outliers B,C. Several conclusions can be drawn from the larger set of detectors observed underground:

- τ_c determined from extrapolating Equation 2.8 to $A = 0$ appears to be much larger than the ~ 2 ms found in [231]. A naive extrapolation yields 5.0 ms without B,C and 17.1 ms for all data points. The latter is indistinguishable from an infinite τ_c with the data at hand. Given the large uncertainty of the individual data points and extrapolation, the final conclusion is that τ_c is at least a few times larger than the observed τ_n .
- While the relative behaviour of the two prototype data points was consistent with the τ_f values from Equation 2.10, this is no longer the case for Run 34 data points. Method 1 yields $\overline{\tau_c^L} = 1.3$ ms and an inconsistent $\overline{\tau_c^M} = 2.4$ ms (2.0 ms without B,C). Conversely, method 2a yields $\tau_f = 0.9$ ms/1.7 ms for L and M (or 1.0 ms/2.0 ms without B,C), incompatible with the calculated 2.410 ms and 4.683 ms.
- From the complete dataset, the shorter τ_f and longer τ_c results in a smaller critical area $A_{crit} = 3.0$ mm² (8.7 mm² without B,C). This means that also the M thermometer is too large from a phonon lifetime perspective. From this dataset, it cannot be decided whether the S structure (never tested on CaWO₄) is small enough.

5.1.2. TES on Si and Al₂O₃ at MPP




Figure 5.3.: L-, M-, and S-TES on CRESST-III phonon detectors. Left: holder for test measurements of three CRESST-III phonon detectors at MPP. The holder is pictured without outer and top covers. The pictured detectors are 3×L-TES on CaWO₄ which were not selected for Run 34 (not measured in this configuration). Right: silicon detectors prepared for testing with 1×M-TES, 2×S-TES (eventually not selected for Run 35, two installed in Run 36).

To collect additional rise time data before the start of Run 35, the detector testing procedure at MPP was augmented to include full detector operation rather than just

a transition measurement. This requires weakly coupling the crystals in an enclosed detector holder. The setup for this is shown in Figure 5.3 (left): three phonon detectors are housed side-by-side, with a copper-kapton-copper partition preventing cross-talk via scintillation light. The detectors are resting on 3 Al_2O_3 spheres (1 mm diameter) and are clamped from the top (the mechanical contact is provided by an additional Al_2O_3 sphere glued on the clamp). For Run 35, one batch (3 crystals) each of CaWO_4 , Al_2O_3 and Si detectors were tested in early 2018. Of the CaWO_4 crystals, two were operated for a rise time measurement. For the third, only the transition temperature was measured. The two CaWO_4 data points are included in Figure 5.2 as green data points. The Al_2O_3 detectors were successfully operated, but their transition temperatures were deemed too low, and the TES were polished off and refabricated. Thus there are six data points from Al_2O_3 detectors.




Figure 5.4.: Athermal phonon lifetimes measured in $2\text{ cm} \times 2\text{ cm} \times 1\text{ cm}$ Si (left) and Al_2O_3 (right) detectors, for TES structures M and S. For Al_2O_3 , the two runs are shown in blue and green data points. All detectors were operated in high-rate conditions at MPP. Given the low number of detectors and high intrinsic scatter, multiple possible interpretations using Equation 2.8 are shown. The blue lines averages all data points. The dotted lines are the best fit assuming $\varepsilon = 1$, i.e. letting $1/\tau_c$ go to zero. The dashed lines are an attempt to find a plausible interpretation with smallest possible τ_c .

Figure 5.4 (left) shows rise time data for the three silicon detectors pictured in Figure 5.3 (right), measured at MPP. The two S-thermometers show rise times of 0.72 ms and 0.90 ms, for an inverse mean $\overline{\tau_n^S} = 0.80$ ms. The M-thermometer is significantly faster with $\tau_n^M = 0.45$ ms. With a naive extrapolation through the data points (shown in solid blue in Figure 5.4, left) one arrives at $\tau_c^{Si} = 2.4$ ms and $\tau_f^S = 1.2$ ms, $\tau_f^M = 0.55$ ms. Like in the case of CaWO_4 , the τ_f found in this way do not agree with the calculated values of 2.54/1.18 ms.

The agreement with the model for τ_f can be somewhat improved by choosing a shallower interpolation (dashed line in Figure 5.4, left), which goes through the faster S-rise

time and assumes a 13% larger $\overline{\tau_n^M}$ than the observed value. This interpolation, consistent with the three data points given the large scatter, yields $\tau_c^{Si} = 1.2$ ms and $\tau_f^S = 1.9$ ms, $\tau_f^M = 0.88$ ms. Still, tension with τ_f from Equation 2.10 remains. On the other hand, the data cannot exclude an infinite τ_c for silicon (i.e. athermal phonon absorption only in the TES). The dotted line following $\tau_n \propto A^{-1}$ fits the data equally well.

To summarize the Si measurement, with three detectors a τ_c of at least 1.2 ms could be demonstrated which translates to $A_{crit}^{Si} < 14$ mm². The data are consistent with all values for A_{crit}^{Si} down to zero. The M-TES are therefore clearly too large, while S could be correctly sized or still too large. A more precise statement is possible with a greater number of measured detectors. Faster-than-expected film collection times are also observed in Si.

For preparation of Run 35, two measurements with 3 Al₂O₃ detectors each were done, for a total of 4×M and 2×S. Data from the two runs are shown in Figure 5.4 (right) as blue and green data points. The rise times measured for the M-TES scatter by a factor of 2.4. Still, the rise times of both S-TES were larger than that of all the M-TES. As above, the data is interpreted using method 2a with three fits to Equation 2.8. Averaging all data points (solid line) yields $\tau_c = 8.1$ ms. An equally good fit is obtained setting $\tau_c \rightarrow \infty$ (dotted line). For the low- τ_c estimate, we ignore the faster two M-TES with a result of $\tau_c = 2.0$ ms. Like in the case of Si, this should be understood as a lower bound on τ_c , with a resulting upper bound on $A_{crit}^{Al_2O_3}$ of 14 mm².

For Al₂O₃, the disagreement with τ_f from Equation 2.10 (0.58/1.26 ms for M/S) is even more dramatic: the observed rise times are longer than the calculated τ_f . This observation cannot be reconciled even for $\tau_c \rightarrow \infty$. The inferred τ_f from method 2 are 0.7/1.6 ms for all data points (solid line), 1.6/3.37 ms for the low- τ_c estimate (dashed line).

5.1.3. Discussion

Conclusions With the available data, the standard interpretation (ignoring phonon collectors) leads to the conclusion that the M-TES with 18.57 mm² area, used on 2 cm × 2 cm × 1 cm crystals of CaWO₄, Al₂O₃ and Si, absorbs a fraction $\varepsilon > 0.5$ of athermal phonons and thus leads to a reduction of athermal phonon lifetimes in all three materials. For the smaller S-TES, the data are consistent with both sub-dominant collection $\varepsilon < 0.5$ or dominant collection $\varepsilon \approx 1$. This would imply that a TES structure with several times smaller area than S is needed to avoid reducing τ_n significantly below τ_c , and to reach the largest possible time-integrated energy input per unit TES area.

Dataset quality To extract more reliable statements, both quantity and quality of the dataset can be improved. As demonstrated here, the interpretation of results can change significantly going from 2 to 11 detectors (in the case of CaWO₄) or 3 to 6 detectors (for Al₂O₃). Beyond the characterisation of more detectors (especially with S- and smaller TES), longer recorded detector live time in each measurement can improve the fidelity of each rise time measurement. In the high-rate environment of MPP, it can be beneficial to collect data for at least 24 h, preferably several days. Compared to the few-hour measurements presented here, stronger selections can then be applied on pulses. This allows finding a better template pulse, created only from pulses following longer-

than-average recovery times after the previous particle hit. This could help reduce the observed scatter in rise times in identically-produced detectors. In addition, an independent measurement of ε can be done by determining the temperature rise induced by an event from a calibration source. These measures increase the complexity of the measurement beyond that of a simple transition measurement traditionally done to select detectors for CRESST.

Modeling and interpretation In addition to the dataset, also its interpretation can be improved. The discussion above ignores many details, maybe most importantly the phonon collectors, which can influence the measurement of τ_n in several ways.

First of all, the diffusion of quasiparticles into the TES happens at a finite speed, which can in principle distort or delay the rise of the signal. This could cause a pulse rise time slower than τ_n especially in large TES. A study of quasiparticle diffusion in CRESST detectors [238] has found delays of ~ 0.2 ms over a length of 4 mm (close to the extension of L-phonon collectors away from the TES).

Secondly, quasiparticle transport is not fully efficient, with recombination occurring during the diffusion. For large phonon collectors, this means that areas far from the TES may contribute only a small share of their collected athermal phonon energy to the signal. This “dead area” can be calculated from diffusion lengths as measured in [238]. Dead area contributes to athermal phonon losses without generating signal, and so effectively shortens τ_c . This would lead to more strongly reduced τ_n in larger TES, as observed here. This can be understood as a τ_c that’s actually smaller for larger TES sizes, as found by “method 1”-interpretations of the data. In a more refined analysis, the impact of phonon collector dead area can be included in the model.

Finally, quasiparticle transport is sensitive to the film quality of the evaporated aluminum. In [238], diffusion lengths of 1.2 – 2.5 mm were measured. Should this be due to intrinsic variation in produced film quality, a batch-by-batch scatter of dead area and time delay in phonon detectors is to be expected.

Outlook Despite the uncertainties, the results of this work motivate a thorough study of CRESST-III phonon detectors with S- and smaller TES. One has to point out that, in spite of the obvious difference in τ_n between the M and L detectors in Run 34, no clear trend in energy resolution is observed. The energy resolution finally achieved in a detector depends on many TES-, detector and operation parameters that make each detector highly individual.

Even if the phonon life times point toward smaller TES sizes, general considerations also show limits to scaling down the TES. Firstly, by collecting less of the signal energy, the fundamental noise from the finite number of excitations rises and becomes dominating at some point. The weaker thermal link necessary to provide constant relaxation time also becomes problematic, with increased thermal-fluctuation noise, higher self-heating limiting the operational bias current, and increased link heat capacity (discussed in the next section). All these factors have to be studied and balanced against each other to find the optimum TES area for a given absorber material and crystal size.

5.2. Thermal diffusion as a limit to weak link designs

The previous section describes a simple way of changing TES size for a CRESST detector, by isotropically scaling down the area of W and Al surfaces and increasing the length of the Au thermal link by the same factor. This section explores a limitation to this approach: scaling the TES area down too much, the thermal link stops acting as a lumped thermal conductance, and has to be modelled as a 1d-structure. First the basic physical phenomenon is introduced, then the relevant scales for the TES detectors in this thesis are evaluated. Finally a measurement to observe and quantify this effect is proposed, and consequences on future designs are given.

5.2.1. Diffusion in an extended thermal link

The differential equation describing the temperature evolution of a heat capacity C (representing the TES electron heat capacity) coupled to a thermal bath at temperature T_b with a point-like link of thermal conductance g (units $[g] = \text{W/K}$) is:

$$C \cdot \dot{T} = g \cdot (T_b - T). \quad (5.6)$$

The solution is exponential relaxation

$$T(t) = T_b + (T(0) - T_b) \cdot e^{-t/\tau} \quad (5.7)$$

with a time constant $\tau = C/g$.

An extended one-dimensional thermal link (without an attached lumped heat capacity) is described by a diffusion equation for the temperature distribution $T(x, t)$:

$$c_0 \cdot \dot{T}(x, t) = g_0 \cdot \partial_x^2 T(x, t). \quad (5.8)$$

with c_0 the specific heat capacity (units $[c_0] = \text{J/m}^3/\text{K}$) and g_0 the specific thermal conductance (units $[g_0] = \text{W/m/K}$). For boundary conditions of either specified temperatures (Dirichlet b.c.) or power inputs (Neumann b.c.), this can be solved analytically by separation of variables. For a finite energy packet initially localized at $x = 0$, the temperature evolution (up to arbitrary scaling and offset) is given by:

$$T(\hat{x}, \hat{t}) = \sum_{n=1}^{\infty} \cos \left[(2n-1) \cdot \frac{\pi}{2} \cdot \hat{x} \right] \cdot \exp \left[- \left((2n-1) \cdot \frac{\pi}{2} \right)^2 \hat{t} \right]. \quad (5.9)$$

Dimensionless space and time variables are introduced as $\hat{x} = x/L$, $\hat{t} = t/t_d$ with L the length of the link and $t_d = c_0/g_0 \cdot L^2$ the diffusion time along the link. The temperature evolution of the hot end of the link thus follows

$$T(0, \hat{t}) = \sum_{n=1}^{\infty} \exp \left[- \left((2n-1) \cdot \frac{\pi}{2} \right)^2 \hat{t} \right] \quad (5.10)$$

and consists of an infinite series of exponentials with increasingly fast decay times. Initially, the temperature of the hot end drops quickly, as energy spreads across the link length. After some time, only the first term in the series remains significant, and the hot end relaxes with a decay time $\left(\frac{2}{\pi}\right)^2 \cdot t_d$. This relaxation time depends on the microscopic properties c_0, g_0 and the link length, but not the cross-section and is thus




Figure 5.5.: Thermal relaxation of a heat capacity through an extended thermal link, by numerical solution of Eqn. 5.11 and 5.12. The cases shown have the same lumped decay time $\tau = C/(g_0 A/L)$ but varying link length, given in terms of the diffusion length $x_d = \sqrt{CL/(c_0 A)}$. The dashed black line shows the exponential relaxation of the lumped link. For the two longest links, the analytical solution for $C = 0$ is shown as a dotted line, giving good agreement for $L = 4 \cdot x_d$. A diffusive link first quickly absorbs energy into its distributed heat capacity, then releases it slowly leading to long tails.

independent from the lumped relaxation time, with depends on the lumped conductance $g = g_0 \cdot A/L$ with the link cross-section A .

The two relaxation regimes can be illustrated by the temperature profile along the link, which is linear for a link with negligible heat capacity, but parabolic ($\propto \cos(\pi/2 \cdot x/L)$) after internal equilibration of the extended link. The situation of an extended link with a lumped heat capacity attached interpolates between the two cases. It is described by a coupled system of equations:

$$C \cdot \dot{T}(0, t) = g_0 A \cdot \partial_x T(0, t) \quad (5.11)$$

for $x = 0$, and

$$c_0 \cdot \dot{T}(x, t) = g_0 \cdot \partial_x^2 T(x, t) \quad (5.12)$$

for $x > 0$. The first equation can be understood as a Robin boundary condition on $T(x, t)$ and leads to a Sturm-Liouville problem not generally solvable analytically. It is easily discretised and solved numerically.

Figure 5.5 shows numerical solutions for links with varying lengths but identical lumped decay time $\tau = C/g = C/(g_0 A/L)$. For link lengths several times shorter than the diffusion length $x_d = \sqrt{CL/(c_0 A)}$ one recovers the exponential relaxation of the lumped link. For L several times longer than x_d , the relaxation can be approximated




Figure 5.6.: Slowest decay time extracted from numerical solutions of Eqn. 5.11 and 5.12, as a function of link length L . Shown in black are the lumped decay time $t_{lump} = \tau$ (dashed) and the slowest decay time $t_{diff} = \left(\frac{2}{\pi}\right)^2 \cdot t_d$ for the “fully diffusive” case with $C = 0$. The numerical decay times closely follow an interpolation with the function $(t_{lump}^\xi + t_{diff}^\xi)^{1/\xi}$ with $\xi \approx 1.15$. Already for $L = x_d$, the slowest decay time is increased to about $1.3 \cdot \tau$.

by the analytical solution for the diffusive link without attached heat capacity (shown as dotted lines for two cases).

Of particular relevance for detectors with extended thermal links is the longest decay time in the system, which sets the timescale for the return to the operation point and thus the rate-tolerance of the detector. Figure 5.6 shows how the longest decay time, extracted from numerical solutions, interpolates between the lumped and the diffusive regime. Importantly, the longest decay time is always above the longer of the two limiting cases. Already at $L = x_d$, the decay time of the thermal link is increased by 30%.

5.2.2. Expected effect on different TES designs

A simple view of the multi-parameter problem emerges by writing the diffusion length as:

$$x_d = \sqrt{\frac{C}{C_{link}}} \cdot L = \sqrt{\frac{g_0}{c_0}} \cdot \tau. \quad (5.13)$$

From the first form, one can read that the link behaves like a lumped element as long as its heat capacity $C_{link} = c_0 \cdot A \cdot L$ is much smaller than the attached heat capacity C . This is a clear criterion for validating existing TES designs. The second form is more useful for finding new designs: with a targeted lumped relaxation time τ and the link material properties c_0, g_0 , the diffusion length x_d can be specified. The link length must then be several times shorter than x_d to stay fully in the lumped regime. To achieve the

TES	$V_W [10^{-15} \text{ m}^3]$	$V_{Au} [10^{-15} \text{ m}^3]$	C_{Au}/C_W
PD-L	816	0.34	$2.6 \cdot 10^{-4}$
PD-M	408	1.2	$1.8 \cdot 10^{-3}$
PD-S	204	2.0	$6.2 \cdot 10^{-3}$
LD, NUC-LD	8.4	3.0	0.22
NUC-M	2.1	3.0	0.44
NUC-S	1.0	1.5	0.42

Table 5.3.: Volumes of W (TES) and Au (thermal link) as well as relative heat capacity in the TES designs described in this thesis. PD: CRESST phonon detector, LD: CRESST light detector, NUC: NUCLEUS targets (see chapter 8).

necessary lumped conductance g implied by the targeted τ , very small link cross-section may then be required.

The behaviour of the thermal links in the TES designs described in this thesis can be evaluated from the low-temperature heat capacity of the tungsten and gold structures. At low temperature, the electron contribution to the heat capacity (linear in temperature) dominates. The electronic heat capacity is described by the Sommerfeld parameter ($\gamma_W = 1.008 \text{ mJ/mol/K}^2$ [261] for tungsten, $\gamma_{Au} = 0.67 \text{ mJ/mol/K}^2$ [198] for gold), leading to $c_W = 105 \text{ J/K/m}^3 \cdot \frac{T}{\text{K}}$ and $c_{Au} = 66 \text{ J/K/m}^3 \cdot \frac{T}{\text{K}}$. The volumes of tungsten and gold films in the various TES designs are listed in Table 5.3, along with the estimated relative heat capacity between gold and tungsten. Here, the normal-conducting electron heat capacity has been listed, in the superconducting transition, the TES heat capacity is expected to increase by a factor of up to 2.34.

As shown in the third column of Table 5.3, the link heat capacity is negligible compared to the TES heat capacity for the CRESST phonon detector TES. The CRESST light detector TES combines an order-of-magnitude smaller collection area (smaller tungsten film) with a proportionally stronger thermal link (larger gold volume). In consequence, the gold heat capacity is only a factor 4-5 below the tungsten heat capacity. The NUCLEUS target TES (see chapter 8 feature even smaller tungsten areas at similar thermal link strength, leading to an even larger gold heat capacity (a factor 2-3 below tungsten).

This shows that the thermal link in phonon detector TES should behave properly as a lumped element. In light detector TES and more so in the smaller NUCLEUS TES, the link comes close to the diffusive regime. This can have adverse effects on detector performance. In a lumped link that quickly relaxes to a linear temperature profile, half of the link heat capacity can be taken as a contribution to the effective TES heat capacity. As the TES heat capacity dominates in this limit, the link contribution can be neglected. For a diffusive link, the time-evolution of a signal changes. A part of the deposited energy quickly leaks into the link, reducing the pulse height of a slow calorimeter. The relaxation of the TES back to operation temperature is slower, leading to reduced rate tolerance.

5.2.3. Implications for future TES designs

A well-working thermal link must have a heat capacity negligible compared to the TES it cools. This implies a limit for simple scaling-down of TES layouts. A reduced-sized TES

film requires a correspondingly weaker thermal link to keep the nominal relaxation time constant fixed. The simple way of designing a weaker thermal link without changing feature sizes is making it longer. This increases the relative link-TES volume by the square of the scaling factor. Scaling below a critical size will bring the thermal link in the diffusive regime. The phonon detector TESs studied here are far from this limit. The light detector is close. The smaller NUCLEUS TES introduced in chapter 8 are fabricated with reduced feature sizes (width of the link): 20 μm for NUC-M, 10 μm for NUC-S (see Table 8.1) with a correspondingly shorter link length. Without this reduction in width, the gold heat capacity would be larger than the tungsten heat capacity (by a factor of 6-7 for NUC-S).

Should smaller, more weakly coupled designs become attractive in the future, fabrication techniques could be developed to avoid diffusive links. Beyond the wet-chemical photolithography used reliably in CRESST down to 10 μm features, smaller structures could be realized in positive lithography steps, e-beam lithography or reactive ion etching. Before investing in new technology, a thorough exploration of the feature size-limits of the standard techniques is warranted.

An alternative approach would be to identify alternative materials to fabricate weak thermal links. Several properties must be considered, such as oxidation, purity and magnetic properties. Purely in terms of thermal link performance, low specific heat c_0 and thermal conductance g_0 are desirable. The longest decay time achievable with a given material, while requiring the link capacity to be smaller than the TES capacity divided by a factor α , is

$$\tau_{max} = \frac{C_{TES}^2/\alpha}{c_0 \cdot g_0 \cdot A_{min}^2} \quad (5.14)$$

where A_{min} is the smallest link cross-section possible in fabrication. This relation allows studying trade-offs in material properties and fabrication methods.

5.2.4. Proposed experimental study

It was shown above that CRESST light detectors and NUCLEUS TES are close to the regime where the thermal link stops behaving as a lumped element. Nonetheless, observing an effect from thermal diffusion in standard detectors is not obvious. The standard pulse model containing two decay times may fit reasonably well also to a pulse that is additionally distorted by non-exponential thermal relaxation. A test structure with devices for comparison on a single chip may clarify the onset and impact of thermal diffusion.

Figure 5.7 shows such a structure with four LD-TES, identical up to the implementation of the thermal link. With widths of 10, 20, 40, 80 μm and correspondingly increasing lengths, the thermal links have the same lumped conductance but widely varying heat capacity. By observing simultaneously pulse shapes and amplitude responses to a calibration source, the impact of excessive thermal link heat capacity may become quantifiable.




Figure 5.7.: Proposed test structure to study experimentally whether and how CRESST light detector TES are affected by diffusion in the thermal link. The four LD-TES on the chip share identical lumped-element link strength, but length and width of the link each change by a factor 2 from one to the next (i.e. a factor 4 in gold heat capacity). The link widths are $10\ \mu\text{m}$ (bottom right) to $80\ \mu\text{m}$ (bottom left). The standard light detector is in the top left. Two heater structures are added for redundancy. The size of the square boundary is 10 mm.

6. The NUCLEUS experiment

NUCLEUS is a new CE ν NS experiment aiming to deploy ultra-low threshold cryogenic detectors at a nuclear power reactor. The project started with initial ideas at Max-Planck-Institut für Physik in Munich in 2016, a first detector prototype and publications [224, 262, 263] followed in 2017. The collaboration was officially founded in 2018. The basic premises of the NUCLEUS experiment were in part developed in the scope of this thesis.

6.1. Ultra-low threshold detectors for CE ν NS

The unique approach of NUCLEUS among CE ν NS experiments is using detector technology with recoil thresholds as low as 20 eV (demonstrated, aiming for below 10 eV long-term). This allows the detection not only of the high-energy tail of CE ν NS-induced recoils, but the majority of all interactions. The strongly increased rate per detector mass offered by this approach allows a corresponding reduction in detector size. Incidentally, the targeted ultra-low threshold requires a reduction in size of the individual detector element: a gram-scale cryogenic detector. For these two reasons, NUCLEUS is aiming to develop the smallest working neutrino detector of 10 g mass. The smallest demonstrated neutrino detector is the 14.6 kg CsI scintillator used by the COHERENT collaboration in the first observation of CE ν NS.

While the achievable energy threshold is a clear advantage for cryogenic detectors, some challenges arise from this approach. The technology providing the cryogenic environment, typically a $^3\text{He}/^4\text{He}$ dilution refrigerator, is a complex system in itself, requiring time-intensive maintenance and custom fine-tuning. The total target mass is limited by the size of individual detector elements and the number of electronic read-out channels, usually a main driver of cost and complexity for such an experiment. In addition, cryogenic detectors are inherently slow, with millisecond rise times. This removes the advantage of a pulsed neutrino source with a much faster time-structure. Also, it poses the problem of dead time when operating anti-coincidence vetos.

Overall, a cryogenic experiment such as NUCLEUS requires large neutrino fluxes, such as provided by nuclear power reactors. Fortunately, the small recoil energies induced by low-energy reactor neutrinos are less problematic than for other technologies. Innovation is needed on the side of background suppression, to allow a small target to attain high sensitivity in an above-ground location. This chapter goes into detail about all those aspects and how they come together to shape the strategy for NUCLEUS.

6.1.1. Neutrino flux model and recoil spectra

To study the physics potential of gram-scale cryogenic detectors at nuclear reactors, a neutrino flux model has to be specified. The study of neutrino emission from nuclear reactors is a rich research field in itself, so here only a summary and motivation of the model used in the sensitivity studies of this work will be given.

The dominant physical origin of the reactor antineutrinos is the successive beta decay of the fission products of the fissile isotopes ^{235}U , ^{239}Pu , ^{241}Pu and ^{238}U . Many hundreds of nuclides contribute, with a wide range of half-lives. Calculating the reactor antineutrino spectrum is thus a complex problem that depends on the initial fuel composition as well as the thermal power history of the reactor. Two approaches exist in the literature: an “ab initio” evolution of isotopes in the reactor using a full nuclear database, or a conversion approach taking as input the measured beta-spectra of the products of the fissile isotopes irradiated with thermal neutrons at ILL [264]. Ab initio calculations have to deal with incomplete or uncertain nuclear databases, while the conversion approach depends on assumptions such as the details of “virtual beta branches” fitted to the electron spectra. Modern approaches uniting both methods were developed to predict fluxes measured in IBD experiments [265, 266]. The precision of the flux prediction reaches 2-3 % over a wide energy range, but the conversion approach is limited to neutrino energies of above 2 MeV.

The neutrino flux below the IBD threshold has to be calculated “ab initio”. A large contribution to the neutrino flux in this energy range comes not from fission, but radiative neutron capture on ^{238}U . The neutrino flux prediction used for neutrino-electron scattering in TEXONO [179] uses the fission spectra from an older ab initio calculation [159] and adds a flux prediction for ^{238}U neutron capture. A more recent ab initio calculation [170] combines these contributions. Figure 6.1 shows a comparison of these neutrino spectra.

All flux calculations agree well at higher energies, falling within the 3%-error corridor given by [170] between 2-6 MeV. At lower energies, Vogel & Engel is lacking the contribution from ^{238}U neutron capture and the parametrisation from Mueller et al is outside its region of validity. There is reasonable agreement between Kopeikin (2012) and TEXONO (2006), although derived by different methods for different situations (generic pressure water reactor vs. particular power plant with given thermal power output history).

The sensitivity studies in this work use the generic flux model of [170] which agrees with the flux predictions widely used in IBD experiments at high energies and extends the energy range below the IBD threshold. A more detailed understanding of the neutrino flux at the specific experimental site will be required to interpret future data from NUCLEUS, and to correctly make the envisioned precision measurements. The sensitivity studies presented in chapter 7 are only weakly dependent on the details of the flux model.

Fig. 6.2 shows the relative importance of the different neutrino energies in $\text{CE}\nu\text{NS}$, by folding the neutrino flux with the $\text{CE}\nu\text{NS}$ cross-section, proportional to E_ν^2 . The solid green line shows that the most important contribution to $\text{CE}\nu\text{NS}$ at nuclear reactors stems from neutrino energies between 2 and 4 MeV, although a significant rate is induced by lower energies.

For a given target nucleus and experimental energy threshold E_{th} , the fraction of detectable recoils can be found from Eqn. 1.11 to be

$$f_{det} = (1 - E_{th}/T_{max})^2 \quad (6.1)$$

for maximum recoil energy $T_{max} = 2E_\nu^2/M$, as long as $T_{max} > E_{th}$. In this way the dashed, dash-dotted and dotted lines in Fig. 6.2 show the distribution of neutrino energies resulting in detectable recoils on a ^{184}W target at various thresholds. Neutrinos below




Figure 6.1.: Reactor antineutrino flux predicted at 40 m distance from a reactor of 4 GW thermal power. The plot compares different methods: Mueller et al (2011) [265], combining information from nuclear databases and reference electron spectra, as parametrised in Table VI valid above 2 MeV, Vogel, Engel (1989) [159], an ab initio approach including fission only, Texono (2006) [179], adding the neutrino flux from ^{238}U neutron capture, and Kopeikin (2012) [170], a more recent ab initio calculation. See main text for discussion.

1.8 MeV can be detected only at thresholds below 30 eV. Higher thresholds lead to a large suppression of the detectable $CE\nu NS$ rate.

Using the differential cross-section of Eqn. 1.12 with the neutrino flux of [170], the differential recoil rate can be calculated:

$$\frac{dR}{dT}(T) = \int_{E_{min}}^{\infty} dE_{\nu} \quad \Phi_{\nu}(E_{\nu}) \cdot \frac{d\sigma}{dT}(E_{\nu}) \quad (6.2)$$

with $E_{min} = \sqrt{\frac{1}{2}TMc^2}$ the minimal neutrino energy necessary to produce a recoil energy T . Figure 6.3 shows the expected $CE\nu NS$ recoil rates at 40 m distance to a 4 GW power plant, on targets of CaWO_4 , Al_2O_3 , Ge and Si.

6.1.2. Gram-scale cryogenic detectors

With the threshold model of section 2.3.4 and the recoil rates at a nuclear reactor from section 6.1.1, the choice of detector size can be studied in more detail. The rate in a single cube shows an interesting dependence on the edge length, as the rate initially increases with the mass of the cube, but then sharply drops as the increasing threshold makes the detector insensitive to the small recoil energies. Figure 6.4 shows the $CE\nu NS$ rate per cube (at 40 m from a 4 GW PWR) for CaWO_4 and Al_2O_3 detectors. The




Figure 6.2.: Visualisation of the distribution of neutrino energies relevant to a CE ν NS experiment. The antineutrino flux (blue) is weighted by the CE ν NS cross-section ($\propto E_\nu^2$, yellow), to yield the contribution to the CE ν NS rate as a function of neutrino energy (solid green line). There is a significant contribution from flux below the IBD-threshold (gray area), although kinematics and threshold effects further suppress this. In the additional green lines, the fraction of neutrinos producing a ^{184}W -recoil above threshold is taken into account, for $E_{th} = 10$ eV (dashed), $E_{th} = 30$ eV (dash-dotted) and $E_{th} = 100$ eV (dotted green).

shaded bands indicate 20% uncertainty on the threshold. The highest rate per cube is expected at 15-25 mm edge length in CaWO_4 , 60-70 mm in Al_2O_3 . These numbers should not be seen as the optimal cube size, but rather as the maximum size beyond which the mass gain is overcome by a loss in sensitivity to low-energy recoils. Total rate constraints will prevent a 60 mm Al_2O_3 cube and also 20 mm CaWO_4 cube from running with sub-keV thresholds at surface. A more valid argument can be made regarding the expected signal-to-background ratio. Assuming a mass-dependent background rate of 10 counts / kg day (dashed lines in Fig. 6.4) corresponding to 100 dru¹ over a ROI of 100 eV, one sees immediately that a) smaller cubes achieve better signal-to-background and b) signal-to-background in Al_2O_3 is far inferior to CaWO_4 . The optimal size of the detector cubes has to be chosen in a compromise between signal-to-noise (which keeps improving for smaller sizes) and absolute signal rate (which biases towards unrealistically large detectors). In addition, technical reasons (detector mounting, handling, fabrication of a scaled TES) place limits on the feasible detector size.

For the NUCLEUS prototypes and the initial physics phase, cubical targets of 5 mm edge length were chosen. This translates to target masses of 0.49 g for Al_2O_3 , 0.76 g for CaWO_4 , 0.67 g for Ge and 0.29 g for Si, justifying the term “gram-scale cryogenic detector”. The threshold scaling model predicts sub-10 eV thresholds when optimized, safely sufficient for CE ν NS observation. This size is a compromise between TES fabrica-

¹the “differential rate unit” dru corresponds to 1 count/keV/kg/day.




Figure 6.3.: Expected differential rates of $CE\nu NS$ on different target materials at 40 m distance from a reactor of 4 GW thermal power. The gray band indicates a range of background indices achieved at higher energies in different shallow low-background facilities [267, 268, 269, 182]. Below recoil energies of 100 eV the $CE\nu NS$ rates on targets containing heavy nuclei ($CaWO_4$, Ge) rises orders of magnitude above these background rates.

tion and detector handling (difficult for smaller sizes) and stable above-ground operation (potentially problematic for larger targets). The listed target materials are proven with the W-TES sensor technology within the CRESST experiment ($CaWO_4$, Al_2O_3 , Si) or used in other cryogenic rare-event searches and commercially available in high quality (Ge).

One should keep in mind from Figure 6.4 that for the application to $CE\nu NS$, 5 mm target size is on the small side. For the projected energy thresholds, a signal gain $\propto d^3$ is expected for target sizes at least up to 1 cm ($CaWO_4$) or 2 cm (Al_2O_3). Therefore, a careful investigation of detector stability for larger target sizes is warranted.

6.1.3. Science phases of NUCLEUS

Like many scientific projects, the NUCLEUS experiment will proceed in a staged approach. In the first science phase, all aspects of the experiment have to be proven at a small scale. At the same time, relevant results should be achievable with the small-scale setup. The success of this phase motivates an upgrade to larger target mass, enabling precision measurements and opening up new physics cases.

6.1.3.1. NUCLEUS-10g

The necessary size for a definite observation of $CE\nu NS$ using reactor neutrinos can be estimated from the neutrino flux model given in section 6.1.1. The recoil spectra for 40 m distance to a 4 GW power reactor, shown in Figure 6.3, integrate to (31.9, 5.2, 20.0,




Figure 6.4.: $\text{CE}\nu\text{NS}$ and background rates in a detector cube as a function of size. The solid lines show the $\text{CE}\nu\text{NS}$ rate above threshold in a single cube 40 m from a 4 GW PWR, the detector threshold is modeled as in section 2.3.4. The shaded bands indicate the effect of 20% uncertainty in the threshold. The dashed lines show the rate per cube caused by a background of 10 counts / kg day. The optimal cube size has to be found in a compromise between signal-to-background ratio and absolute signal rate.

6.2) counts per year per gram of target material above a threshold of 10 eV in the well-characterised cryogenic detector materials (CaWO_4 , Al_2O_3 , Ge, Si). This shows that few grams of the targets containing heavy elements (CaWO_4 , Ge) can be sufficient to make an observation with less than a year of measurement time. More detailed sensitivity studies for a first reactor- $\text{CE}\nu\text{NS}$ measurement under various conditions are presented in section 7.2. Naturally, larger signal statistics is always preferred from a scientific standpoint. The upper limit on the target mass for the first phase of NUCLEUS comes from a technical constraint: each target detector requires a dedicated SQUID amplifier for readout. This usually means the number of wires between the cryogenic setup and room temperature is proportional to the number of target elements. Classical readout systems can support several tens of channels in a cryostat without prohibitive wiring complexity or heat-load.

If the goal was simply to maximize neutrino counts, the detector array should consist of a few dozen CaWO_4 calorimeters. The sensitivity of NUCLEUS, as a rare-event search operating in the presence of backgrounds and at an unexplored energy regime, depends not only on the signal statistics. A marked difference in the $\text{CE}\nu\text{NS}$ rate between light and heavy targets is apparent in the count rates listed above, due to the N^2 -dependence of the coherent enhancement. Combining a light and a heavy target in the same setup therefore allows disentangling signal from (incoherent) background with different dependences on target material. This multi-target approach is easy to implement in NUCLEUS, as the cryogenic sensor technology can be combined with many different target materials. The strategy of multiple target materials can yield invaluable information for tuning background models. (A statistical study exploring the potential of this approach is described in section 7.2.2.) The payload for NUCLEUS will




Figure 6.5.: Conceptual drawings of planned NUCLEUS detectors. Left: NUCLEUS-10g, comprised of 3×3 arrays of CaWO_4 and Al_2O_3 gram-scale cryogenic calorimeters, enclosed in cryogenic veto detectors. Right: one component of the NUCLEUS-1kg target, a 15×15 array of Ge detectors that can be fabricated from a single 6-inch wafer. Several of these 150 g units can be stacked vertically. Figures taken from [262].

therefore consist of equal numbers of CaWO_4 and Al_2O_3 targets, chosen for their high hardness and the expertise gained with these materials within the CRESST experiment. Two 3×3 arrays, one for each of the target materials, add up to a total mass of 11.25 g, suggesting the name NUCLEUS-10g for the first science phase of the experiment. A conceptual drawing of the envisioned detector assembly is shown in Figure 6.5 (left). The 18 SQUID channels for the target arrays, plus a smaller number for cryogenic veto detectors described below, can be provided by a commercial SQUID system in a standard cryostat.

6.1.3.2. NUCLEUS-1kg

A future systematics-limited precision measurement of $CE\nu NS$ at a nuclear reactor needs to scale up to a larger target mass. For a statistical precision of 1%, $\mathcal{O}(10^4)$ events need to be detected on a timescale of few years. This requires kg-scale target masses, so this future phase will be tentatively called NUCLEUS-1kg. A realistic path towards kg-scale arrays of gram-scale cryogenic calorimeters combines two innovative technological approaches: parallel fabrication of sensor arrays, and multiplexed TES readout.

The photolithographic thin-film fabrication technology used for TES production naturally lends itself to manufacture of sensor arrays. Current CRESST target detectors ($2 \times 2 \times 1 \text{ cm}^3$) can be processed in groups of three in the sputtering and evaporation machines, but have undergone lithography individually. This is an obstacle to large-scale series production. The smallness of NUCLEUS target calorimeters allows to fabricate many sensors on one 5 mm thick wafer, which are subsequently diced into separate detectors. First experience with this process was gained with the prototype discussed in chapter 8. On six-inch wafers, a standard raw material for Si and Ge used in the semiconductor industry, 15×15 arrays (151 g for Ge) could be produced in a single fabrication cycle in this way. A drawing of such a unit is shown in Figure 6.5 (right). The commercial availability of high-quality large wafers motivates the choice of Ge and Si as heavy and light targets for NUCLEUS-1kg. This production methodology requires

larger fabrication machines and improved large-area deposition homogeneity, which have to be developed.

Such an increase in number of target elements similarly poses new challenges on the read-out side. Separately addressing an individually packaged SQUID amplifier per TES channel is no longer feasible in terms of wiring design and heat-load on the cryostat. SQUID multiplexing of TES detectors is a maturing technology, applied e.g. to microcalorimeters in X-ray spectroscopy and bolometers for cosmology applications. Compared to some of these detectors, the “particle-physics-TES” of NUCLEUS are slow ($\mathcal{O}(100)$ kHz bandwidth), which eases performance goals on the multiplexing scheme. On the other hand, the CE ν NS application demands highest energy resolution, so ideally the multiplexer should not contribute significantly to the total noise budget. Allowing $\mathcal{O}(10)$ multiplexer units in the setup, multiplexing factors of $\mathcal{O}(100)$ are needed for NUCLEUS-1kg.

A review of TES spectrometers introducing the techniques of time-division, code-division, (MHz) frequency-domain and (GHz) microwave SQUID multiplexing can be found in [218]. The largest multiplexing factors have been reached with microwave SQUID multiplexing, and applied to both microcalorimeter and bolometer TES arrays. A bolometric TES array for millimeter-wave astronomy with 64 multiplexed channels read out at 20 kHz has been used for on-sky observations [270]. In the laboratory, simultaneous readout at 62.5 kHz sampling rate of 128 TES microcalorimeters for gamma spectroscopy was demonstrated [271]. In both cases, the detector noise was not increased by the multiplexed readout. Upcoming experiments are developing microwave SQUID multiplexing of $\mathcal{O}(1000)$ channels for cosmology [272], neutrino mass measurements [273] and X-ray astronomy [274]. A system of several units with parameters similar to the currently demonstrated multiplexers will be sufficient for the readout of NUCLEUS-1kg. This shows the technical feasibility of reading out thousands of gram-scale cryogenic calorimeters. Nonetheless, the development of a multiplexing scheme for NUCLEUS-1kg, the necessary room-temperature electronics, and its integration with the experiment represent long-term R&D efforts for NUCLEUS.

6.2. A fiducial-volume cryogenic detector

Gram-scale cryogenic calorimeters open up ultra-low energy thresholds. To make a successful CE ν NS observation, several additional features are important for NUCLEUS. It is of central importance to achieve a low background rate at sub-keV energies. The use of an array of small cryogenic calorimeters as target allows to implement a new concept for background suppression, named the “fiducial-volume cryogenic detector”. Cryogenic detectors usually do not feature event localization, and are therefore particularly vulnerable to surface backgrounds. Operating such sensitive devices without the protection of the rock overburden provided at underground laboratories, penetrating radiation also presents new challenges. The fiducial-volume cryogenic detector addresses these challenges with separate cryogenic anti-coincidence veto detectors.

6.2.1. Target array

The target of the NUCLEUS experiment will consist of a large number of gram-scale cryogenic calorimeters, individually below a cm in size. The NUCLEUS-10g array will

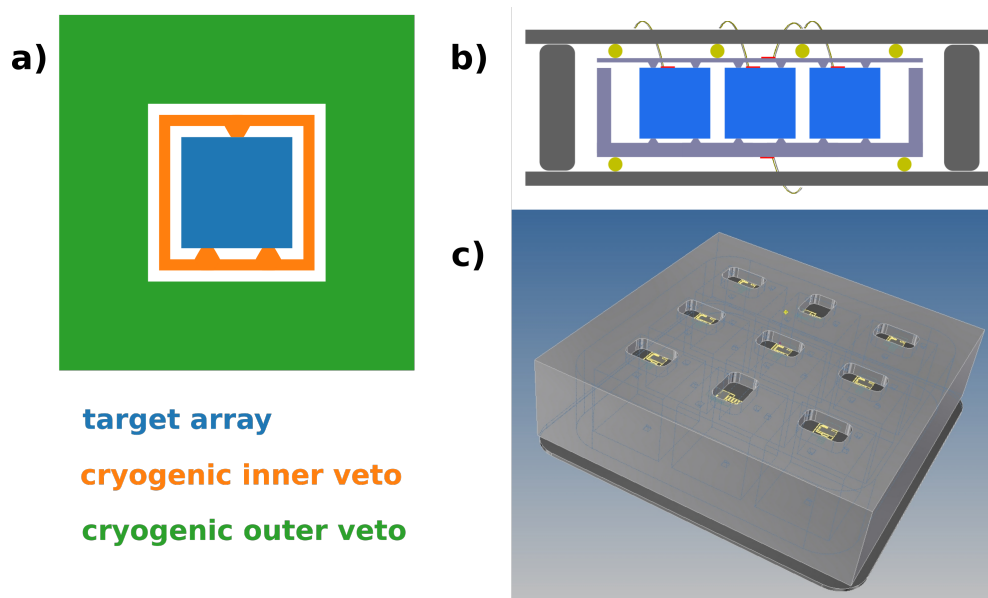


Figure 6.6.: a) Schematic drawing of the fiducial-volume cryogenic detector. The target array is held by the thin cryogenic inner veto, which blocks all line-of-sight from the target array to the surroundings. Both these elements are surrounded by the massive cryogenic outer veto. b) cross-section of a possible implementation for NUCLEUS-10g. Rigid Si support structures are shown in dark gray, inner veto components in light gray. The target elements are mechanically supported on point-like pyramid structures. Windows in the inner veto allow feeding through wire-bond connections from the support structures. c) Technical drawing of the planned NUCLEUS-10g inner veto, consisting of a rigid rectangular beaker and a flexible plate.

have a side length of only few cm. For NUCLEUS-1kg, this increases to $\mathcal{O}(10\text{ cm})$. The smallness of the target elements is important to achieve the necessary low thresholds, but has several other advantages. Firstly, a gram-scale calorimeter can run stably in a high-rate environment. Considering the (unshieldable) surface muon flux of $\mathcal{O}(1\text{ /cm}^2\text{/min})$, a larger detector would suffer multiple muon hits per minute, preventing stable operation at the TES transition temperature. Secondly, the small target array can be fully surrounded by other cryogenic detectors with less stringent requirements on energy threshold and stability.

6.2.2. Cryogenic inner veto

The cryogenic inner veto serves the dual purpose of detector holder and active surface veto. This addresses two challenges previously encountered with sensitive cryogenic target detectors in the CRESST experiment. Holding forces applied to thermal detectors can lead to stress-relaxation events in the crystal, in extreme cases observable as high rates of high-energy events associated with crystal cracking [275]. To exclude an analogous phenomenon at lower energies, an instrumented holder (itself a cryogenic detector with similarly low energy threshold) can act as an anticoincidence veto. A separate issue encountered in a low-background environment is related to non-instrumented surfaces




Figure 6.7.: Geant4-simulations of veto performance from [262]. The simulated geometry contains a 5 mm CaWO_4 target cube, surrounded by a 1 mm thick silicon shell (9 mm side length) as inner veto, and a 10 cm by 10 cm CaWO_4 cylinder as outer veto with a central 1 cm cubical cavity. Left: effect of the inner veto on a β -emitting contamination of ^{210}Pb on the inner surface of the inner veto. The differential rate of energy depositions in the target (black) can be lowered by about two orders of magnitude by an anticoincidence cut with the inner veto at a threshold of 30 eV (red). Right: effect of the outer cryogenic veto on an environmental γ -background modeled after the Dortmund Low Background Facility [269]. The rate of energy depositions in the target is shown without the outer veto (black), with a passive outer veto (blue), and an outer veto energy threshold of 1 keV, resulting in a suppression by more than three orders of magnitude.

facing the detector. Surface events from natural radioactivity, such as back-to-back alpha/nuclear recoil events, are likely to deposit only a part of the full energy, while the rest can be deposited in a non-instrumented surfaces. This causes e.g. “degraded alpha” events which range from zero energy up to the Q-value of the decay. This process, involving ^{206}Pb recoils from ^{210}Po decays, is visible as an important background in CRESST-II [240]. The event excess above the modeled background observed in [240] can be plausibly explained by an enhanced low-energy rate of such events via processes involving surface roughness of exposed copper surfaces facing the target crystal [241]. A subsequent design change in CRESST-detectors, called the “stick design” [276], avoids copper surfaces by implementing a fully scintillating housing comprised of reflecting foil and CaWO_4 sticks holding the crystal. The scintillation light, detected in a separate cryogenic detector, flags energy depositions in the inner surface of the detector holder. This design concept is taken one step further by making all surfaces facing the detector active (i.e. a cryogenic detector themselves). This ensures an efficient surface veto in the energy range relevant to $\text{CE}\nu\text{NS}$. Geant4 simulations demonstrate the potential of an active low-threshold surface veto. In an early study published in 2017 [262], an inner veto operating with an energy threshold of 30 eV was found to suppress background from surface contamination by about two orders of magnitude at low energies (see Figure 6.7 left).

From these top-level purposes of the inner veto, individual requirements can be derived. To be an efficient surface veto, the inner veto needs a low energy threshold. Ideally this should be as low as that of the target. For simulations, a threshold of 30 eV

was assumed. For the same reason, the inner veto should block all line-of-sight from the target array to the (non-instrumented) surroundings. This necessitates a segmented inner veto, to be able to mount and contact the target array. As a holder element in mechanical contact with the targets, the inner veto needs to provide a defined holding force. This is best provided by flexible elements. The holder of a cryogenic detector should ideally be thermally insulating, so as to not provide an escape path for the signal phonons to be detected in the thermometers. One way to approach this is to limit the contact area between target detectors and inner veto through point-like contacts.

In Figure 6.6 a), the inner veto is shown as a thin structure both enclosing and touching the target cube. Design drawings for NUCLEUS-10g, in Figure 6.6 b) and c), show an implementation for a 3×3 target array. The inner veto is segmented in two silicon cryogenic detectors. A thick-walled (1-2 mm) “rectangular beaker” is milled from a silicon slab and covers most of the solid angle around the target array. The cover is completed by a thin “lid”, a flexible 200 μm silicon wafer. Both beaker and wafer touch the target elements via pyramid structures fabricated in a wet-etching process. The pyramid tips expose a 100 μm patch of original wafer surface, ensuring equal height of the pyramids. A defined holding force is provided by the deformation of the thin wafer. The holding force can be tuned by adjusting the height of silicon posts (dark gray in Figure 6.6 b). The inner veto components are held via Al_2O_3 spheres glued to a thick silicon support structure (not instrumented). Laser-cut windows in one of the inner veto components allow feeding wire-bonds for thermal and electrical connections through to the target detectors, sacrificing some of the geometric coverage. Experience gained with a simplified version of this design is discussed in chapter 8.

6.2.3. Cryogenic outer veto

The setup of NUCLEUS must be placed in proximity to a nuclear power reactor, therefore necessarily in a surface location with limited control over the experimental environment. A main challenge for conducting a rare-event search in such a location is the rate of external particle backgrounds, potentially many orders of magnitude higher than in underground laboratories. Of particular concern are muons and muon-induced secondaries, cosmic ray-induced neutrons and environmental gamma radiation.

To help address these challenges, the suite of NUCLEUS cryogenic detectors will include a massive outer veto enclosing the “inner detector” composed of target array and inner veto. Its purpose is to identify penetrating ionizing radiation depositing energy both in the outer veto and a target detector. Operated in anticoincidence, such backgrounds can then be suppressed. Modern dry dilution refrigerators feature sample spaces of diameter > 30 cm. The goal of outer veto design is to fill this space with kilograms of cryogenic detectors with \sim keV energy thresholds and sufficiently fast time response to run with a particle rate of $\mathcal{O}(\text{Hz})$.

An important design choice is the material used for the outer veto detectors. Gamma radiation is most effectively stopped by high- Z materials, while neutrons transfer their energy most efficiently to light elements (ideally hydrogen). CaWO_4 , containing heavy as well as light nuclei, is therefore an interesting choice. A disadvantage is the lack of commercially available large-diameter crystals. The use of Li_2WO_4 in cryogenic detectors is investigated in the R&D project BASKET [277]. This material also combines light and heavy nuclei, with the additional advantage of containing the isotope ${}^6\text{Li}$. This

isotope has a high neutron absorption cross-section and could allow a simultaneous in-situ characterization of neutron backgrounds inside the NUCLEUS setup. The concept is presented in the context of the CRESST experiment using LiAlO_2 crystals in [93]. Germanium is a rather heavy target with the advantage of being a well-developed semiconductor material. Germanium detectors using charge readout enjoy widespread use in industry and particle physics. Charge readout is possible in a cryogenic environment and could offer fast timing (sub-microsecond rise times) and sufficient energy resolution.

Early simulations of the outer veto as a CaWO_4 cylinder (10 cm diameter, 10 cm height), shown in Figure 6.7 (right), promise a suppression of γ -background at low energies by $\sim 10^2$ with a passive outer veto and $\sim 10^3$ with a 1 keV threshold in the outer veto.

The NUCLEUS prototype discussed in chapter 8 features a silicon cylinder outer veto (5 cm diameter, 5 cm height). Silicon acts as a cheap and easily machinable stand-in for the more exotic materials, proving the basic module design without focus on the outer veto detector properties.

The outer veto design process, ongoing in the collaboration at the time of writing, is interlinked with the design of the passive shielding. The goal of both subsystems together is to suppress neutron and gamma backgrounds to tolerable levels. The passive shielding can be optimized for either of the two by choice of materials and layer ordering. The outer veto, on the other hand, cannot compete in neutron suppression due to its limited thickness. The best combined strategy then appears to be focusing the passive shielding on neutron suppression, and optimizing the outer veto against gamma radiation. The currently favored design for the NUCLEUS outer veto consists of $\mathcal{O}(10)$ individual germanium charge detectors with a thickness > 5 cm in all directions surrounding the inner detector.

6.3. The NUCLEUS experimental site

6.3.1. The Very-Near-Site at Chooz

A suitable experimental location near a nuclear reactor is a key ingredient for any reactor neutrino experiment. Building on the long-standing cooperation between the French power company Electricité de France (EDF) and CEA Saclay on the Double Chooz neutrino experiment (data-taking 2011-2016), a new experimental site at Chooz Nuclear Power Plant (CNPP) has been identified for the NUCLEUS experiment. CNPP operates two N4 pressurised water reactors with 4.25 GW thermal power output each (B-1 and B-2), which are typically running at full power with several weeks of refueling stops per year on an alternating schedule. Within Double Chooz, two identical 10 ton-scale liquid scintillator detectors were studying neutrino oscillation via IBD detection, at the so-called Near Site (~ 400 m baseline) and Far Site (~ 1 km baseline). For the planned NUCLEUS-10g cryogenic detector, the distances of both Double Chooz sites to the reactors are too high. A new and much smaller room closer to the reactor buildings was made available to the NUCLEUS collaboration, consequently named the “Very-Near-Site” (VNS) [278]. A map of the power plant area including VNS is shown in Figure 6.8. The experimental space consists of a 24 m^2 room in the basement of a five-story office building located between the two reactors. The distances to the cores are 72 m (B-1) and 102 m (B-2), resulting in a total neutrino flux at VNS of about $3 \cdot 10^{12} \bar{\nu}_e / (\text{s} \cdot \text{cm}^2)$.

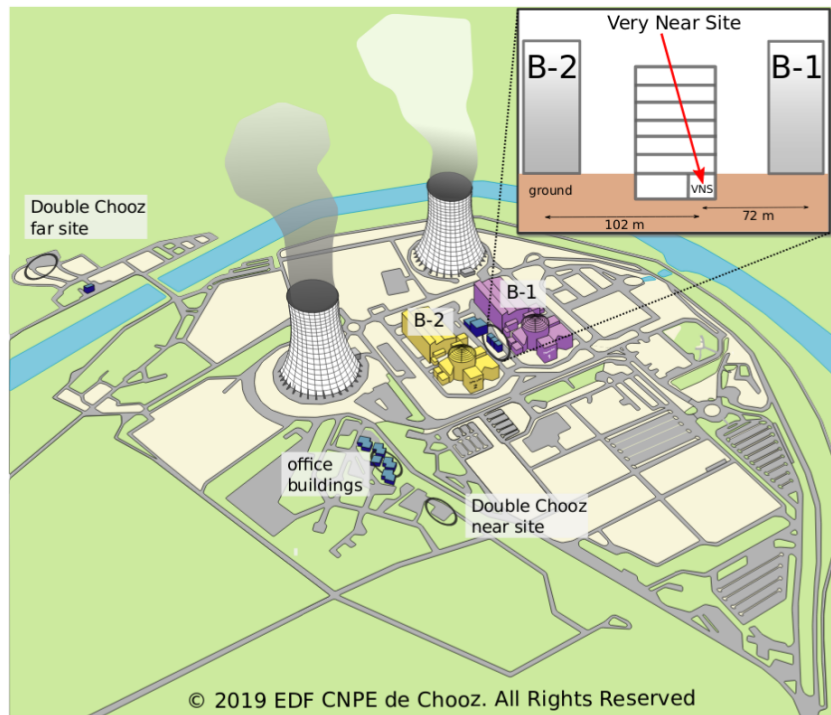


Figure 6.8.: Sketch of the Chooz nuclear power plant and the Very-Near-Site, from [278]. Since 2008, the power plant hosts the neutrino experiment Double Chooz with experimental locations called the Near Site (~ 400 m baseline) and Far Site (~ 1 km baseline). The Very-Near-Site is located on the connection line between reactor buildings B-1 and B-2, with distances to the cores of 72 m and 102 m. The inset shows the location of the experimental room in the basement of an office building.

The NUCLEUS setup is limited by available space to a size of a few m^3 and a weight of several tons, to avoid overloading the floor of the room. Background measurements of muon and neutron flux were performed at VNS (discussed below). Measurements of gamma radiation and vibration spectra are planned before the installation of detectors.

6.3.2. Radiogenic background

A detailed understanding of backgrounds in the NUCLEUS detectors will be crucial for a successful $\text{CE}\nu\text{NS}$ measurement. Here, the discussion of the topic in [278] is summarized. Reviews on backgrounds to low-rate experiments can be found in [279, 280].

A variety of radioactive processes can deposit energy in the target detectors and thus mimic a neutrino signal. Radiogenic backgrounds can be categorized as external, internal, reactor-correlated and cosmic-ray induced. Radiation from the reactors, in particular reactor-correlated gamma and neutron backgrounds, can be excluded thanks to the reactor containment and tens of meters of soil separating the experiment from the cores.

Natural radioactivity, in the materials of the building, surrounding environment and the detector materials themselves, comes in the form of α , β and γ radiation. Natural

γ radiation with energies up to 2.6 MeV is penetrating and interacts mainly via Compton scattering, with deposited energies ranging from zero to the Compton edge. γ -rays can be attenuated using high- Z materials. In addition, the tendency of γ -rays to induce multiple interactions can be exploited to suppress this background with the cryogenic outer veto.

β and α radiation has much shorter range (mm and tens of μm in solids, respectively). The full particle energies are far above the CE ν NS region of interest. Only the background originating in detector and surrounding surfaces is dangerous, as energy may only be partially deposited in the target detector by the electrons and α -particles. Internal contamination with radioactive isotopes can be controlled by cleanliness procedures in detector production and assembly. In addition, the inner veto is specifically designed to tag partial energy depositions in detector and surrounding surfaces. An important special case of surface backgrounds is the airborne radioactivity induced by ^{222}Rn . A gaseous intermediate decay product of the uranium decay chain, ^{222}Rn with its long half-life of 3.8 days can diffuse out of materials with natural uranium content. Subsequent deposition of radon and its decay daughters on detector surfaces is a main source for surface backgrounds in low-radioactivity experiments. Radon deposition has to be prevented by special cleaning procedures during detector assembly, and a radon-free environment has to be provided during operation of the experiment.

Cosmic-ray induced backgrounds are strongly suppressed in deep-underground laboratories and present a new challenge to rare event searches at shallow depth. Primary cosmic rays are predominantly protons, with a small admixture of alpha particles and heavier nuclei, hitting the Earth's atmosphere with high energies. The secondary particles reaching the Earth's surface are a mixture of muons, neutrons, electrons, neutrinos, protons and pions. A small overburden such as a building efficiently blocks electrons, protons and pions, so that muons and neutrons remain as relevant backgrounds. (Atmospheric neutrinos do not cause a measurable interaction rate.) Muons are the most penetrating charged particle species and active countermeasures in the form of e.g. a plastic scintillator muon veto are ubiquitous also in deep underground experiments. External neutrons can be slowed and stopped by elastic scattering, most efficiently with hydrogen-containing materials. Shielding materials doped with neutron-absorbing nuclei (such as boron or lithium) are useful to capture thermal neutrons.

In addition to the atmospheric neutrons, produced outside the experimental setup, an important background comes from muon-induced neutrons. These neutrons are produced in nuclear reactions around and within the shielding, particularly on heavy elements. As the muons cannot be shielded, the only handle on muon-induced neutron production is the material choice and layer sequence of the shielding. Particularly, high- Z layers close to the detectors are problematic due to muon-induced neutron production. As muon-induced neutrons are coincident with the primary muon, a high-efficiency muon veto helps suppressing this background.

Between October 2017 and May 2018, a first background characterization campaign was carried out at VNS [278]. The relative reduction of muon and fast neutron flux compared to outside the building was measured. For the muon measurement, a plastic-scintillator based muon telescope was used. A mean muon flux reduction by $29 \pm 1\%$ was found, corresponding to an average overburden of (2.9 ± 0.1) m.w.e. Fast neutrons were measured using liquid scintillator detectors (using EJ-301 by Eljen Technology) that allow pulse-shape discrimination between electron and proton recoils. At energies

above 100 keV_{ee}, a constant neutron flux reduction by a factor of 8.1 ± 0.4 was found.

In a first GEANT4 MC simulation using a preliminary setup configuration, the relative importance of the various backgrounds has been studied [278]. External γ -rays have been found negligible compared to other backgrounds after efficient suppression in the outer veto. Muon-induced neutrons are expected to be the highest background, while atmospheric neutrons are sub-dominant. Without the muon veto, 26 ± 1 muon-induced neutron events are expected below 100 eV per kg day exposure in CaWO₄, compared to 32 CE ν NS events. For Al₂O₃, the CE ν NS rate drops to 1.9 / (kg day), while the muon-induced neutron rate is similar to CaWO₄. This shows that the signal-to-background ratio in NUCLEUS decisively depends on the muon veto efficiency (with a target of 99%). Also, the multi-target approach may be useful to achieve an independent background-only measurement. The simulated energy spectrum for muon-induced neutrons is used as the expected dominant background for some of the sensitivity studies discussed in section 7.3 and shown in Figure 7.14.

6.3.3. Cryogenic infrastructure

The base temperature and holding times required for NUCLEUS make ³He/⁴He dilution refrigerators an essential technology. Operating in proximity to a nuclear power plant, without liquid helium infrastructure and under hazardous materials restrictions, a wet dilution refrigerator relying on a liquid helium bath is not a good solution for the experiment. Dry dilution refrigerators liquefy the circulating mixture using a Joule-Thomson valve thermally anchored to a mechanical pulse-tube cryocooler with a base temperature below 4 K. Without the liquid helium bath, these refrigerators offer larger experimental spaces, a simpler vacuum environment allowing faster cooldowns, and the potential for remote operation. A commercial dry dilution refrigerator (Bluefors LD400) is therefore selected as the cryostat for the NUCLEUS experiment. The operation of sensitive cryogenic detectors in close proximity to pulse-tube cryocoolers requires special care in vibration isolation [281]. The NUCLEUS cryostat includes a central port from the experimental space up to room temperature. A dedicated vibration isolation using a long spring pendulum (resonance around 1 Hz) is in the design phase. The cryostat will also host a calibration system using room-temperature LEDs coupled to the detectors using optic fibres. Via the known photon energy and Poisson statistics, this allows an absolute energy calibration of the cryogenic detectors, as e.g. described in [216].

6.3.4. Passive shielding

To operate a sensitive rare-event search in a surface building, a carefully optimized passive shielding is indispensable. Successful concepts optimized for germanium γ spectroscopy with keV-scale thresholds are described in [182, 269]. Typically, they feature alternating layers of high-Z (Pb) and low-Z (PE) materials. This is to suppress external γ radiation as well as the neutrons produced by muons and γ -rays in the high-Z materials. The innermost layer is typically copper, to alleviate the problem of neutron production close to the center. An active muon veto, sometimes with an inner layer, is an integral component of a successful shield. CONUS [181] uses an adaptation of such a shield specifically optimized for CE ν NS observation at lower energies than γ spectroscopy (few 100 eV). Notable features include an innermost layer of low-activity lead and a muon veto inside the outermost lead layer. This is to reduce the muon veto trigger rate due




Figure 6.9.: Sketch of the cryogenic facility planned at VNS for the NUCLEUS experiment. The centerpiece is a commercial dry dilution refrigerator (Bluefors LD400) with 29 cm diameter experimental space, sufficient to host the cryogenic detectors also for future NUCLEUS phases. Surrounding the cryogenic detectors is a m^3 -scale layered passive shield covered with an active muon veto (green). A section of the passive shielding (red) and muon veto will be installed between the cryogenic detectors and the cryostat mixing chamber. A suspension system isolating the detector from vibrations is under study.

to external gammas and allow a double-role of the muon veto to moderate neutrons produced in the first lead layer.

For NUCLEUS, the situation is slightly different due to the presence of the outer cryogenic veto, which ideally offers several orders of magnitude γ suppression and doubles as an inner muon veto. This allows (and forces) to optimize the passive shielding more towards neutron suppression, to reach a comparable final rate of both types of backgrounds. For this reason, high-Z materials will be fully avoided close to the detectors (as long as it is not itself instrumented, such as the outer veto). The shielding is planned with an outer lead layer and a thick inner polyethylene layer. Detailed optimization is ongoing.

6.3.5. Muon anticoincidence veto

With muon-induced neutrons causing the largest expected background to a $\text{CE}\nu\text{NS}$ measurement at VNS, a high-efficiency muon veto is of paramount importance to the success of the experiment. Efficiencies of around 99% have been achieved in the low-




Figure 6.10.: Detector dead time fraction versus muon-veto side length assuming different detector rise-times. The muon veto trigger rate is taken from MC simulations, the relation between rise time and time resolution is taken from [282]. With a targeted rise time of $100 \mu\text{s}$ (dashed blue line), the NUCLEUS target array can operate in a cubic-meter shielding with a dead time of few percent. The expected dead time for the first target prototype (discussed in chapter 8) with a $300 \mu\text{s}$ rise time is shown as a solid blue line. Slower detector technologies with millisecond rise times (black lines) cannot profit from a muon anticoincidence veto without prohibitive dead time.

background setups mentioned above. This is also the goal for NUCLEUS, to reach a signal-to-background ratio of 100.

Any compact shielding is limited in size by the trigger rate of the muon veto surrounding it. An anti-coincidence window, in which detector pulses are ignored, has to be placed around all muon veto signals. Therefore the detector dead time increases proportionally with the muon veto trigger rate. For NUCLEUS, this is a particular challenge due to the slow cryodetector pulses in the target array. To be sure to veto target pulses caused by muon-associated events, the size of the anticoincidence window is chosen according to the cryodetector time resolution (which is much worse than the muon veto time resolution).

A quantitative version of this argument is shown in Figure 6.10. The trigger rate as a function of the edge length l of a cubical muon veto is found from MC simulations to be about $493 \text{ Hz/m}^2 \cdot l^2$. The relation between time resolution and detector rise-time is assumed to be linear, with the coefficient taken from a similar cryogenic detector operated at a neutron beam facility [282]. There, an anticoincidence window of $48 \mu\text{s}$ ($\pm 5\sigma$) was found sufficient for a detector rise time of $100 \mu\text{s}$. With these assumptions, the resulting dead time fraction can be plotted as a function of muon veto edge length for various rise times (Figure 6.10). For the first NUCLEUS target prototype discussed in chapter 8, the rise time was found to be $300 \mu\text{s}$. The blue line shows the corresponding

expected dead time. The goal for NUCLEUS is to speed up the target detectors to a rise time of $100 \mu\text{s}$, to reach a dead time of only 2% for a $(1 \text{ m})^3$ shielding.

7. Sensitivity studies for NUCLEUS

After having specified the basic parameters, plans and goals of the NUCLEUS experiment in chapter 6, this chapter describes the sensitivity calculations exploring the potential of gram-scale cryogenic calorimeters, which helped define and solidify the concept earlier introduced over the course of the thesis.

Discovery Scenarios The likelihood studies presented here have been instrumental in shaping the NUCLEUS experiment in the early stages, before the collaboration was formally established. The studies of various neutrino sources for a hypothetical cryogenic CE ν NS experiment (presented in section 7.1) showed that a nuclear power reactor is the most promising neutrino source for such a measurement.

Using the basic parameters of the prototype gram-scale cryogenic calorimeter (size, threshold), the discovery scenarios (detailed in section 7.2) provided the necessary target size (10g) and background index (100 dru) for a first observation of CE ν NS of reactor neutrinos.

The studies of neutrino flux at different reactor baselines in section 7.2.1 was an important input to the process of locating an experimental site for the experiment. With these calculations, the basic parameters of the NUCLEUS experiment were settled before the first observation of CE ν NS was announced by the COHERENT collaboration.

The advanced discovery scenarios (sections 7.2.2 and 7.2.3) explore more realistic experimental situations NUCLEUS might face. For the first CE ν NS observation by NUCLEUS-10g the calculations highlight the value of the multi-target approach.

Physics Scenarios The studies described in section 7.3 stake out the physics case for the NUCLEUS experiment, both in the first phase and for future upgrades with larger target mass.

The calculations of statistical precision achievable with one year with 10 g and 1 kg targets quantify the required systematics control. This is an important input for the collaboration, affecting aspects from detector design to neutrino flux modeling. Vice-versa, the calculations can be used to adjust target sizes to estimated systematic uncertainties, choosing the smallest target mass capable of achieving statistical uncertainties on the level of the estimated systematics budget. For example, with a 3% neutrino flux uncertainty, a 100-200 g target is sufficient to make a systematics-limited measurement of the CE ν NS cross-section. On the other hand, many other applications, such as the measurement of the neutrino magnetic dipole moment or light mediators, do not rely as strongly on signal normalization and keep profiting from larger targets.

The examples of BSM sensitivity to be achieved by NUCLEUS strongly motivate the construction of the experiment. More detailed versions of these studies can be used to quantify the impact of various detector effects (such as energy resolution, energy scale calibration and linearity) more realistically, and again provide quantitative feedback to the design process. They are also the starting point for studies of degeneracies between

different BSM scenarios and possibilities to break them, e.g. by employing various target materials.

7.1. Comparison of Neutrino Sources for NUCLEUS

7.1.1. Stopped-pion source

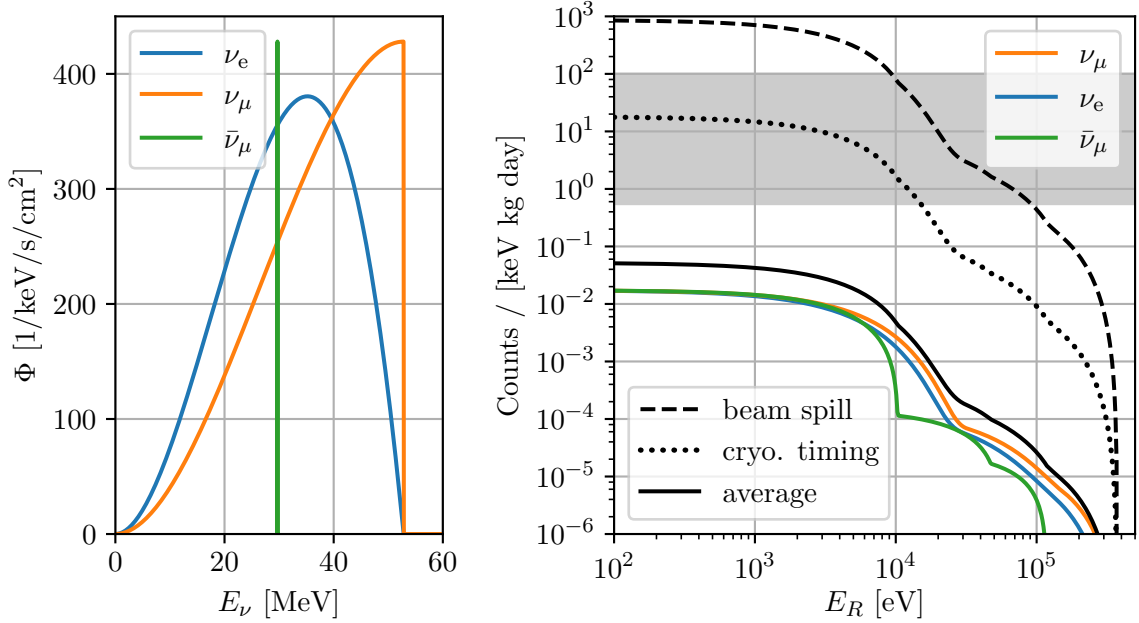


Figure 7.1.: Neutrino flux and recoil spectrum expected from a stopped-pion source. Left: neutrino spectrum for pion decay-at-rest, as parametrised in [168]. The normalization is the time-averaged flux at 19.3 m distance from SNS. Right: recoil spectrum on CaWO_4 at 19.3 m distance from SNS. The solid lines are the time-averaged rate and composition by neutrino species. The dashed line shows the instantaneous rate during a beam spill. The dotted line is the rate averaged over a 48 μs window around a beam spill, reflecting the limited timing performance of a NUCLEUS-detector. The gray band indicates the shallow lab-backgrounds from Figure 6.3.

When neutrino sources were introduced in section 1.2.4, stopped-pion sources were named as a favourable facility for $\text{CE}\nu\text{NS}$ detection due to neutrino energy and timing structure. To explore the potential of NUCLEUS at a spallation neutron source, we study a scenario in which a gram-scale cryogenic detector is placed at the COHERENT experimental site at SNS. The flux expected at the closest COHERENT location 19.3 m from the SNS target is shown in Figure 7.1 (left). With an expected 0.08 ν produced per flavor per proton-on-target (POT) at SNS and $1.76 \cdot 10^{23}$ POT delivered over 308.1 days of running [127], the average neutrino production of SNS evaluates to $5.3 \cdot 10^{14}$ ν / flavor / s, some five orders of magnitude smaller than that of a nuclear power reactor. The time-averaged recoil spectrum shown as solid lines in Figure 7.1 (right) is correspondingly reduced, far below the demonstrated background levels in shallow low-background

laboratories. The instantaneous recoil rate during a $1 \mu\text{s}$ beam spill (which happens 60 times per second) on the other hand reaches up to 10^3 counts/(keV·kg·day) (dashed line in Figure 7.1 right), allowing high signal-to-noise ratio for a detector with excellent timing performance. Cryogenic detectors are intrinsically slow devices, physically limited by phonon collection timescales (see also section 6.3.5). In a setup optimized for coincidence measurements, a cryogenic detector reached a time resolution of $4.8 \mu\text{s}$ (rms) [282]. Such a detector cannot resolve the beam spill in time, and a region-of-interest around it has to be specified, e.g. $\pm 5\sigma$. The average recoil rate over this $48 \mu\text{s}$ window is shown as the dotted line in Figure 7.1 (right). This recoil spectrum falls in the middle of the expected background range, promising at most mediocre signal-to-background ratio.

Another challenge for a cryogenic experiment is the comparatively small exposure that can be collected. In the COHERENT measurement using CsI, the signal expectation is given as 173 events over 308.1 days in a 14.6 kg detector [127], evaluating to a rate of 14 events / kg / year. The expectation calculated for CaWO_4 between 5 keV and 20 keV is 23.1 events / kg / year, slightly higher due to the larger W nucleus. This increases to 60.1 events / kg / year with a threshold of 1 keV and 85.7 events / kg / year with a threshold of 10 eV. Thus, target masses below ~ 100 g will not observe a large number of $\text{CE}\nu\text{NS}$ events per year, independent of threshold, and with a background rate that may be comparable to the signal rate.

In summary, for a $\text{CE}\nu\text{NS}$ measurement at a stopped-pion source, energy thresholds in the keV range are sufficient. Detectors with μs timing can subtract source-unrelated backgrounds and the low overall neutrino flux can be offset with large target masses. These requirements are a good match for conventional scintillation detectors, while ultralow-threshold gram-scale cryogenic calorimeters do not achieve the necessary timing performance and suffer from their smaller target mass. Their strongest advantage, an energy threshold in the 10 eV range, does not lead to a significantly higher signal rate.

7.1.2. Research reactor

Although a much weaker neutrino source compared to a power reactor, a research reactor is potentially attractive as such a site may allow easier access and better accommodation of experimental activities. To study the possibility of detecting $\text{CE}\nu\text{NS}$ with NUCLEUS at a research reactor, we define a new benchmark scenario in terms of signal and background. As a typical research reactor we take FRM-II, the neutron source operated by the Technical University of Munich in Garching. Compared to a typical power reactor, FRM-II uses fuel with a much higher content of ^{235}U , between 93% and 85% over the burning cycle [283]. This leads to a different normalization and energy dependence of the neutrino flux. A test using the flux parametrisation of the different fissile isotopes from [284] suggests the flux changes by $\mathcal{O}(20\%)$. This difference is ignored for the purpose of this study, which assumes the reactor neutrino spectrum discussed in section 6.1.1.

With 20 MW thermal power, the neutrino flux is 200 times lower than that of a typical power reactor core. To compensate, we assume an optimistic experimental location 3 m from the reactor core (right up to the concrete wall surrounding the reactor pool). This leads to a comparable flux (11% reduced) to the case of 40 m distance from a power reactor. At such a close distance, reactor-correlated backgrounds cannot be excluded




Figure 7.2.: Signal and background models for a $\text{CE}\nu\text{NS}$ detection scenario at a research reactor. The $\text{CE}\nu\text{NS}$ signals on CaWO_4 and Al_2O_3 are shown for 3 m distance from a 20 MW reactor (resulting in a similar flux compared to 40 m distance from a power reactor). The background shown is a two-exponential parametrisation of the simulated neutron background expected in the MINER experiment [187]. For illustration, the same background suppressed by a factor 10 is shown as a dashed line.

and the level and composition of backgrounds are difficult to estimate. To judge the feasibility of the experiment at FRM-II, we use the results of a detailed background measurement and simulation campaign performed by the MINER collaboration [187]. MINER plans to deploy CDMS-style cryogenic detectors in close proximity to the 1 MW TRIGA reactor of Texas A&M University. Gamma and thermal neutron backgrounds have been measured directly on-site and used for GEANT4 simulations of the planned shielding configuration. The shielding covers about 2.7 m between the reactor core and detectors and consists of layers of graphite, borated polyethylene, lead, high-density polyethylene and copper. We take the MINER simulated neutron spectrum (Figure 14 in [187]) as a possible realistic background scenario after shielding.

Figure 7.2 shows a comparison of expected background and signal rates with the neutrino flux at FRM-II. The background overwhelms the signal by nearly an order of magnitude at all energies. In this case, reactor-off periods are not expected to help, as the neutron background is presumably fully reactor-correlated. The multi-target approach may be of advantage by allowing a direct in-situ measurement of the neutron background.

In our sensitivity study, we assume identical backgrounds in Al_2O_3 and CaWO_4 with the known shape of the MINER neutron background. We vary the background normalization to study the additional suppression needed for a $\text{CE}\nu\text{NS}$ detection within one year of measurement time. The result of the study is shown in Figure 7.3. The assumed




Figure 7.3.: Significance of $\text{CE}\nu\text{NS}$ observation after one year of measurement at a research reactor, as a function of background normalization with respect to the simulated MINER neutron background. To make a significant detection of $\text{CE}\nu\text{NS}$ likely, the background would have to be suppressed by a factor 10 (using CaWO_4 alone) or 5 (for a combined measurement on Al_2O_3 and CaWO_4) with respect to the MINER simulation.

background has to be reduced by a factor of ten to allow a likely (95%) 5σ -observation of $\text{CE}\nu\text{NS}$ on CaWO_4 alone after one year. Combining CaWO_4 and Al_2O_3 relaxes this requirement to a factor of five of necessary suppression. Given that the simulated neutron background already assumes a shielding of several meters, achieving the required background level to observe $\text{CE}\nu\text{NS}$ at a research reactor seems challenging.

7.1.3. Radiogenic neutrino source

In section 1.2.4, radiogenic neutrino sources were introduced as an alternative with well-predicted neutrino flux allowing to operate the experiment underground, albeit with a low signal rate and small recoil energies. To further illustrate this point, we consider a scenario in which a source such as previously planned for the SOX experiment [172] is deployed in an underground facility at a distance of 1 m to a NUCLEUS-like detector. We assume a $^{144}\text{Ce}/^{144}\text{Pr}$ neutrino source with an initial activity of 150 kCi. The spectrum, taken from [285], is shown in Figure 7.4. The choice of this isotope is motivated by the strong anticorrelation of Q-value and decay time among beta-decaying isotopes. For a source experiment, one would like a high Q-value (to obtain energetic recoils) and a decay time of many months to ease the logistics constraints between manufacturing the source and operating the experiment. $^{144}\text{Ce}/^{144}\text{Pr}$ fulfills both of this: the initial decay of ^{144}Ce has a long half-life of 285 d and a correspondingly small Q-value of 318.7 keV, while the following decay of the daughter nucleus ^{144}Pr has a short half-life of 17 min




Figure 7.4.: Neutrino spectra and induced recoil spectra of a SOX-like radiogenic antineutrino source. Left: neutrino spectrum of the beta-decaying pair $^{144}\text{Ce}/^{144}\text{Pr}$ (solid blue line, from [285]). The dashed orange line shows the same spectrum weighted with E^2 to highlight the contribution to CE ν NS. Right: recoil spectra on various materials, calculated for 1 m distance from a 150 kCi source. The gray band indicates background levels measured (at higher energies) in underground labs using solid-state detectors [234, 286].

and a high Q-value of 2997.5 keV. While the ^{144}Ce decay does not yield a detectable neutrino flux, it sets the timescale for ^{144}Pr production and decay, which generates the signal. With an endpoint of around 3 MeV neutrino energy, all W recoils produced by the source in a CaWO_4 target fall below an energy of 105 eV.

In addition to the resulting stringent requirement on detector threshold, a source experiment faces the challenge of low signal rates. Even assuming the high activity of the SOX source at 1 m from the detector, the expected CE ν NS counts over a year of measurement are 0.154 per target cube of CaWO_4 and 0.019 per Al_2O_3 cube. The advantage of a source experiment is the choice of experimental location, meaning a well-shielded environment where low background levels can be achieved (provided source contaminations are strictly limited). The gray band in Figure 7.4 (right) shows background levels observed in underground laboratories, from 3.1 counts/(keV·kg·day) observed in CRESST (keV energy range) [234] to 0.1 counts/(keV·kg·day) measured above 40 keV in a high-purity germanium spectrometer at LNGS [286]. Assuming these background levels can be achieved at lower energies, the CE ν NS signal becomes observable below 100 eV for heavy nuclei (Ge, CaWO_4 targets) and below 300 eV for light nuclei (Si, Al_2O_3 targets) with a rate lower by 1-2 orders of magnitude.

Here we assume a detector threshold of 10 eV and targets of CaWO_4 and Al_2O_3 . In contrast to previous scenarios, the measurement time cannot be arbitrarily extended due to the finite life time of the source. Instead we consider a fixed measurement duration




Figure 7.5.: Significance of $\text{CE}\nu\text{NS}$ observation using a 150 kCi $^{144}\text{Ce}/^{144}\text{Pr}$ radiogenic neutrino source. The study assumes an energy threshold of 10 eV, a flat background of 1 counts/(keV·kg·day), and does not leverage the known time dependence of the neutrino flux (counting experiment). The significance reached after one year of measurement is shown as a function of target mass, given as the number of 5 mm cubes per target materials. For the given background model, Al_2O_3 does not contribute to the measurement. An observation becomes possible using 50 cubes of CaWO_4 (38 g) and likely with 100 cubes (75 g).

of one year, which includes 59% of the decays of the source. We vary the target size of the experiment, assumed to consist of an equal number of CaWO_4 and Al_2O_3 cubes of 5 mm side length. For the background we assume a flat rate of 1 counts/(keV·kg·day), a moderate improvement over the CRESST results [234] (extrapolated from keV down to the $\text{CE}\nu\text{NS}$ ROI). Under these conditions, the signal-to-background ratio in Al_2O_3 rises to 1 only at threshold. Given the low overall statistics, Al_2O_3 is not expected to contribute to the measurement. The likelihood study shown in Figure 7.5 confirms this. Given the benign background model, the additional background-only data is not necessary for the measurement. The experiment is limited by signal statistics and requires a CaWO_4 target mass roughly ten times that of NUCLEUS-10g to allow a likely significant observation of $\text{CE}\nu\text{NS}$.

Given the logistical challenge associated with production of the radiogenic neutrino source and the requirements of target mass, a source experiment does not appear favourable for a first observation of $\text{CE}\nu\text{NS}$ using cryogenic detectors. Should large-mass ultralow threshold detectors prove feasible in the future and $\text{CE}\nu\text{NS}$ anomalies arise in other experiments, a source experiment could be a clean follow-up taking advantage of the well-predicted neutrino energy spectrum.

7.2. CE ν NS Observation Potential at a Power Reactor

7.2.1. Benchmark case and variations

To explore the discovery potential of NUCLEUS-10g at a nuclear power reactor, Monte Carlo studies of different scenarios were performed using the unbinned likelihood scheme detailed in Appendix B. Important variables to study are the energy threshold of the detector as well as neutrino flux and background level at the experimental site.




Figure 7.6.: Possible locations for a neutrino experiment at a generic nuclear power reactor. The different options are explained in greater detail in the text and have different profiles in terms of neutrino flux, control of backgrounds, site availability and ease of access. Figure taken from [262].

Figure 7.6 shows possible locations of a neutrino experiment at a hypothetical nuclear power reactor. In the “near case” (15 m distance), the experiment is located inside the reactor containment, where the close proximity to the core ensures a high neutrino flux but reactor-correlated backgrounds cannot be excluded a priori. The “medium case” proposes to deploy the experiment in a building adjacent to the reactor (40 m distance), where the neutrino flux is lower, but access is easier and backgrounds from the reactor are excluded. In the “far case” (100 m distance), the experiment can be outside the power-plant area with unrestricted access but low neutrino flux. To compare the feasibility of a CE ν NS-observation at each of the sites, 300 random spectra were generated for each of the sites and a number of different exposure times. The neutrino flux is calculated for a single reactor core of 4.25 GW thermal power. The detector consists of two arrays of nine (5 mm)³ cubes each, one of CaWO₄, one of Al₂O₃. An energy threshold of 10 eV and a flat background of 100 dru were assumed. This energy threshold is better than currently achieved values (15-20 eV, see chapter 8) but not as low as values predicted by the scaling model in section 2.3.4. The background level is the upper range of backgrounds measured in the shallow laboratories mentioned in Figure 6.3. While such a rate appears realistically achievable in a shallow laboratory, no measurements exist in the energy range relevant for NUCLEUS.

The model fitted to these random spectra has two free parameters, one for the normalization of the flat background and one for the normalization of the CE ν NS signal. In




Figure 7.7.: Significance of $CE\nu NS$ observation as a function of time for different distances from a PWR neutrino source, calculated in an unbinned likelihood ratio test. The background is flat at 100 dnu, with the normalization fitted as a free parameter. The colored bands encompass 90% of the 300 toy-MC datasets generated for each exposure time. The median line (solid) is also drawn. A median significance of 5σ is reached after 1.5 days at 15 m, 19 days at 40 m and 314 days at 100 m. The dashed lines follow the expected significance calculated from the most likely dataset, using binned data and assuming a known background level. As expected, both methods converge for large exposure times.

the null model (background-only), the signal normalization is fixed to zero. Figure 7.7 shows the distribution of signal significance for the different sites as a function of measurement time t . The colored bands cover 90% of the toy-MC instances, the solid line follows the median significance. The dashed lines are the median significances obtained from the Asimov data set described in Appendix B. Compared to the toy-MC technique, this method assumes a known background level and uses binned spectra. The Asimov significance rises strictly with the square-root of measurement time, so that a single likelihood calculation for each case is sufficient. As expected, for large measurement times, when the background is well measured above the $CE\nu NS$ energy range, the two statistical methods give very similar results.

This first sensitivity study shows that a $5\text{-}\sigma$ discovery of $CE\nu NS$ at a nuclear power reactor is possible with the NUCLEUS-10g detector within few days in the near case, few weeks in the medium case, and after a year of measurement time in the far case. This makes the medium case the most favourable for the NUCLEUS experiment. While the near case would allow a precise day-to-day measurement of neutrino flux, the expected restrictions on cryogenic operations at such a site can be prohibitive. The low signal rate in the far case does not allow precision measurements over a realistic time scale. The

medium case offers both sufficient signal rate for a fast measurement and ease-of-access which allows constructing and operating the experiment.




Figure 7.8.: Significance of $\text{CE}\nu\text{NS}$ observation as a function of time for different energy thresholds of NUCLEUS-10g, at a distance of 40 m from the reactor core. The other experimental parameters are as in Figure 7.7. The median discovery times are 11 days for a hypothetical threshold-less detector and 27, 74, 274 days for thresholds of 20, 50 and 100 eV respectively.

Figure 7.8 studies the dependence on energy threshold for the medium case in a similar fashion. While the difference in time-to-discovery is less than a factor 3 between 20 eV and 0 eV threshold, an energy threshold of 100 eV increases measurement time by more than an order of magnitude compared to the benchmark 10 eV. This shows that an improvement of energy threshold below the 20 eV demonstrated so far is not crucial for NUCLEUS. On the other hand the measurement rapidly becomes impossible for thresholds significantly higher than a few tens of eV.

Figure 7.9 shows the variation of the sensitivity with the level of the flat background. An order-of-magnitude reduction (10 dru) with respect to the benchmark scenario speeds up the median discovery time from 19 to 10 days. In this regime, measurement time is dictated by signal statistics only, as shown by the background-free case in which the median time to discovery is 8 days. For an order-of-magnitude increase in background (1000 dru), the $\text{CE}\nu\text{NS}$ signal on CaWO_4 is covered in background down to 30 eV and time-to-discovery increases to 73 days.

7.2.2. Multi-target approach: background of unknown shape

Assuming a flat background in the NUCLEUS ROI, from which the $\text{CE}\nu\text{NS}$ signal clearly rises, may be an overly optimistic assumption. Flat backgrounds, such as expected from low-energy Compton scattering, have been observed at well-shielded germanium spec-




Figure 7.9.: Significance of $CE\nu NS$ observation as a function of time for different background levels in NUCLEUS-10g, at a distance of 40 m from the reactor core. The other experimental parameters are as in Figure 7.7. The median discovery times are 8 days for the background-free case and 10, 19, 73 days for the background levels of 10, 100 and 1000 dru.

trometers [268, 269, 182], but only at higher energies (tens of keV). The background in the NUCLEUS ROI has to be understood and thoroughly characterized through ancillary measurements at the experimental site. The multi-target nature of the NUCLEUS experiment can be an advantage towards understanding backgrounds. The distinct nuclear composition of the NUCLEUS targets implies different response to various backgrounds (through different cross-sections for e.g. neutron backgrounds and gamma backgrounds) and the neutrino signal.

To illustrate the power of the multi-target approach, we introduce a new likelihood description of the background, with a free shape parameter. The background in this model is composed of a flat and an exponential component, with three free parameters: $R(E) = C + A \cdot e^{-E/B}$. Like before, we assume an identical background response of the different target materials. The background+signal model again has a common signal normalization as its additional parameter.

The assumption of identical background response in both target materials is clearly simplistic, and serves as a placeholder for a detailed understanding (through simulation and ancillary measurements) of the relevant backgrounds and their relative impact in the target materials at the NUCLEUS site. Still, the simulations introduced in section 6.3.2 indicate that the muon-induced neutron background (expected to be dominant in NUCLEUS-10g) is very similar in $CaWO_4$ and Al_2O_3 .

The simulated background is chosen to represent a “worst-case scenario”, with a slope $B = 45$ eV similar to the $CE\nu NS$ signal in $CaWO_4$ and an amplitude $A = 3000$ dru covering the $CE\nu NS$ signal at all recoil energies. The flat background $C = 100$ dru is




Figure 7.10.: Modeled background spectra for the “worst-case scenario” of a background of unknown shape similar to the $\text{CE}\nu\text{NS}$ signal in CaWO_4 . The background begins to rise exponentially below ~ 300 eV and covers the $\text{CE}\nu\text{NS}$ signal at all recoil energies.

left unchanged. This background is shown in Figure 7.10 in comparison to the $\text{CE}\nu\text{NS}$ signal in the two target materials.

Figure 7.11 shows the result of a suite of likelihood-ratio tests using the binned likelihood and the Asimov data set approach (generating the expected value of the significance of the observation). On the horizontal axis the amplitude A of the exponential background is raised from zero to the value shown in Figure 7.10. The vertical axis shows $\text{CE}\nu\text{NS}$ significance after 1 year of observation at a distance of 40 m from a 4 GW_{th} power reactor.

Even for a background that is actually flat, no significant detection can be made from Al_2O_3 or CaWO_4 alone, as the exponential background model also fits the $\text{CE}\nu\text{NS}$ signal reasonably well. A small significance in favor of the presence of $\text{CE}\nu\text{NS}$ appears as the spectral shape of the signal is correctly modeled. The combination of both materials allows a significant detection even for large exponential backgrounds overwhelming the $\text{CE}\nu\text{NS}$ signal. In this simplistic model, the Al_2O_3 detectors (without appreciable $\text{CE}\nu\text{NS}$ signal) serve to constrain the background fit parameters with sufficient accuracy to discern the additional contribution of $\text{CE}\nu\text{NS}$ in CaWO_4 .

The power of the multi-target approach depends on the feasibility of transporting knowledge about backgrounds in the Al_2O_3 detectors to the CaWO_4 array. In a real experiment, the Al_2O_3 spectrum can be used to tune a Monte-Carlo background model, characterizing the external particle fluxes at the NUCLEUS location. This information then can be used to extract the $\text{CE}\nu\text{NS}$ signal from the CaWO_4 spectrum. The simplistic sensitivity study shown here demonstrates that, in principle, much can be gained employing this multi-target scheme.




Figure 7.11.: Significance of $CE\nu NS$ observation after 1 year of measurement time, allowing the background shape in the model to vary freely. The horizontal axis shows the amplitude of the exponential component of the simulated background. Green and orange lines show the significance that can be reached using data from Al_2O_3 and $CaWO_4$ detectors separately, the blue lines show the significance from a combined fit to both datasets.

7.2.3. Correlation with reactor power

Reactor neutrino experiments can profit from reactor-off periods, which allow studying the backgrounds at the experimental site in the absence of a neutrino signal [287]. The power of this notion for the NUCLEUS experiment can be studied by fitting simultaneously spectra with different neutrino signal strengths and a correspondingly reduced neutrino signal normalization parameter.

In a first artificial scenario, the general impact of reactor-off data is investigated. To this end, the “worst-case scenario” of section 7.2.2 is augmented by a second dataset containing only the unknown exponential background in both materials. The total measurement time of the experiment is fixed to one year, the relative lengths of “reactor-on” and “reactor-off” periods are varied against each other.

Figure 7.12 shows the result of this study, where “reactor duty cycle” denotes the relative duration of the “reactor-on” period. The significance of a measurement employing exclusively $CaWO_4$ detectors can be strongly enhanced by reactor-off data. Varying the relative duration of “reactor-on” and “reactor-off” datasets, an optimum significance after 1 year of measurement is achieved for a ratio of approximately 1/1. The median significance rises from about 1.5σ for “always on” to over 5σ for “half on, half off”. It has to be noted that the sharp decline in sensitivity with reduced reactor-off data is due to the freedom of shape in the background model, which can fit the (clearly visible) $CE\nu NS$ signal on $CaWO_4$ in the absence of constraining “background-only” data. This




Figure 7.12.: The impact of reactor-off data, considering one “reactor-on” dataset and one “reactor-off” dataset of 1 year total duration, using the “worst case” unknown background scenario of section 7.2.2. The median significance of $CE\nu NS$ observation after 1 year of measurement time is shown as a function of the relative length of the “reactor-on” dataset. Al_2O_3 alone remains insensitive to $CE\nu NS$. A measurement using $CaWO_4$ alone profits strongly from reactor-off data, with an optimal ratio of approximately 1/1 for this background scenario. A combined measurement on Al_2O_3 and $CaWO_4$ profits moderately from about 1/3 of “reactor-off” data. This is due to the additional availability of effectively background-only data from Al_2O_3 in this scenario.

matches well with the observation that a combined measurement on Al_2O_3 and $CaWO_4$ profits much less strongly from reactor-off data. The median significance of $CE\nu NS$ detection can be enhanced from 6σ (always on) to 6.5σ (3/4 on, 1/4 off). A simple explanation is that the Al_2O_3 measurement, in this scenario and model, provides the same kind of background-only data constraining the fit. A small amount of additional reactor-off data constrains the background better, before the loss of signal statistics becomes more relevant.

A second scenario is designed to model more closely the neutrino flux conditions expected for NUCLEUS at the VNS. From long-term experience from the Double Chooz Experiment, typically each reactor core of Chooz NPP is switched off for refueling for about one month per year. The refueling stops do not overlap, so that both cores are off simultaneously only in short unplanned intervals (few days per year). As the VNS is located at distances of 72 m and 102 m to cores B1 and B2 respectively, they contribute about 2/3 and 1/3 of the neutrino flux each. Figure 7.13 shows the median sensitivity to $CE\nu NS$ as a function of time in this scenario.

As before, the single-material measurement on $CaWO_4$ shows the clearest improve-




Figure 7.13.: Sensitivity to $CE\nu NS$ as a function of time, including neutrino flux modulation for a VNS-like year of operations. The gray shaded regions denote periods of reduced neutrino flux: the fourth month with a neutrino flux of $1/3$ (B1 off) and the tenth month with $2/3$ (B2 off). The last five days of the year are taken as neutrino-free (B1 and B2 off). The solid lines show the median significance in the given flux-modulation scenario, compared to the dashed square-root functions resulting from the “always-on” scenario. Relevant impact on $CE\nu NS$ significance is visible in the $CaWO_4$ -only measurement.

ment from periods of reduced neutrino flux. The median significance after 1 year is increased from 1.5σ (always-on) to 2.5σ (flux-modulated). On the other hand, the median significance of the combined measurement falls slightly below the “always-on” case shown by the dashed lines in Figure 7.13. This indicates that higher signal statistics on $CaWO_4$ is more valuable than additional data allowing a separation of signal and background. In this scenario, the Al_2O_3 target already allows constraining the background fit.

In the real experiment, a background below the “worst-case” scenario is expected (by definition). Simulations of backgrounds will provide constraints to the background model, which will have less freedom than the one studied in section 7.2.2. Depending on the success of ancillary measurements constraining the background, and the feasibility to fully include both target materials in the same signal+background fit, the flux modulation may contribute to the $CE\nu NS$ observation at VNS. In general, the available reactor-off periods are not expected to be crucially relevant.

7.3. Studying New Physics at VNS

After exploring the conditions under which a successful measurement of $\text{CE}\nu\text{NS}$ can be made with cryogenic detectors, the next section explores the physics potential of NUCLEUS assuming it becomes an effective $\text{CE}\nu\text{NS}$ experiment. Sensitivity calculations in different BSM physics cases are performed for a fixed scenario of NUCLEUS at VNS, for both the 10 g phase ($9\times\text{CaWO}_4$, $9\times\text{Al}_2\text{O}_3$ targets) and the 1 kg phase ($15\times 15\times 7$ Ge targets) introduced in section 6.1.3. The theoretical motivation of the physics cases and references to reviews and existing constraints are given in section 1.2.3.

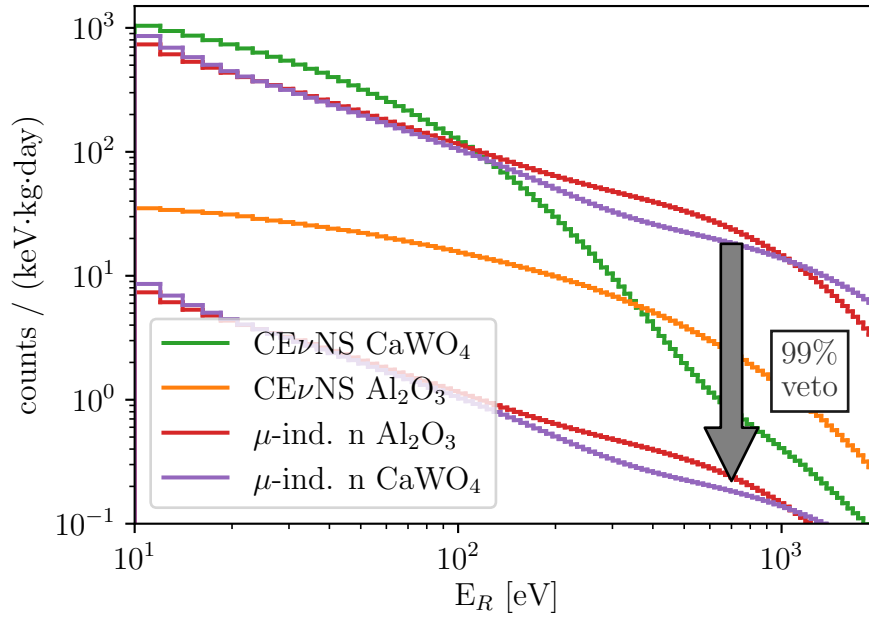


Figure 7.14.: Signal and background rates assumed for studies of physics reach for NUCLEUS. $\text{CE}\nu\text{NS}$ rates on CaWO_4 and Al_2O_3 are calculated from the SM expectation. The dominant background is assumed to be the muon-induced neutron background taken from simulation. It is similar for CaWO_4 and Al_2O_3 and taken to be suppressed by two orders of magnitude using a 99% efficient muon veto.

To showcase the possible future potential of NUCLEUS, we make the assumption that unknown backgrounds will be understood and suppressed in the near future, so that the measurement will be limited by the largest currently known background. By GEANT4 simulations done by the collaboration for an initial shielding configuration (introduced in section 6.3.2), the expected dominant background is caused by muon-induced neutrons. Figure 7.14 shows this background compared to the $\text{CE}\nu\text{NS}$ signal rates. The background swamps the signal for recoil energies above 100 eV and leads to a signal-to-noise ratio below 2 even at small recoil energies. Strong suppression of this background by a high-efficiency muon veto is mandatory. With a muon veto efficiency of 99%, the muon-induced neutron background is expected to be suppressed by 2 orders of magnitude, as indicated in Figure 7.14. Under these conditions, a $\text{CE}\nu\text{NS}$ measurement below 100 eV is possible with S/B above 100, and NUCLEUS can collect data limited

by signal statistics.

7.3.1. Achievable precision at VNS




Figure 7.15.: Precision ($1\text{-}\sigma$) of the $\text{CE}\nu\text{NS}$ -cross-section measurement achievable with NUCLEUS at VNS using a 10 g ($\text{Al}_2\text{O}_3 + \text{CaWO}_4$) target and a 1 kg Ge target. The signal normalization is the only free parameter in the fit to the expected spectrum. Statistical precision is shown as solid lines, for the dashed lines a 10% (1%) systematical uncertainty has been added in quadrature for NUCLEUS-10g (1kg). For comparison, a horizontal dashed line marks the 32% precision achieved in the first observation by the COHERENT collaboration.

Having fixed a model for signal and backgrounds for NUCLEUS at VNS, the achievable precision of the $\text{CE}\nu\text{NS}$ cross-section measurement can be extracted. To find the necessary exposure for a given precision, the Asimov method shown in Appendix B.5 is used. The SM $\text{CE}\nu\text{NS}$ rate is the expected signal s_0 , the alternative hypothesis s_1 is scaled by a constant factor. The exposure necessary to reject this scaling factor is found from the likelihood ratio test statistic. As expected for a signal normalization measurement with no further free parameters (and following from the form of Eqn. B.9), the precision improves with the square-root of the exposure time. For NUCLEUS-10g, a combined measurement on CaWO_4 and Al_2O_3 is calculated using one likelihood function for both binned spectra. For NUCLEUS-1kg, the single material germanium is taken. For lack of a dedicated Ge background simulation, the simulated muon-induced neutron spectrum on CaWO_4 (suppressed by the muon-veto) is taken instead.

In addition to the signal statistics, many uncertainties will affect a realistic experiment, such as neutrino flux uncertainty, detector energy resolution and energy scale. To account for this, a constant systematic uncertainty is added in quadrature to the statistical uncertainty derived above. A similar result is reached by adding the systematic uncertainty as an uncertain signal normalization using the method described in Appendix B.6. For NUCLEUS-10g, 10% systematical uncertainty is assumed, well above

	t [yr]	stat.	stat.+sys.
NUCLEUS-10g	1	9.4%	13.8%
NUCLEUS-10g	3	4.4%	11.4%
NUCLEUS-1kg	1	1.0%	1.4%
NUCLEUS-1kg	3	0.57%	1.15%

Table 7.1.: Statistical and total precision of a $\text{CE}\nu\text{NS}$ cross-section measurement achievable by the NUCLEUS experiment at VNS. The experimental scenario assumes a dominating muon-induced neutron background suppressed by two orders of magnitude using an efficient muon veto.

the flux uncertainty of 2-3% above 1.8 MeV (see section 6.1.1). For NUCLEUS-1kg, a large improvement to 1% is projected, corresponding to full understanding of detector and environment and improved neutrino flux calculations.

A more detailed error budget is necessary for the experiment and will evolve once a detector will be operated at the experimental site. Measured data corroborating the background expectations are needed to ensure that a detailed error budget will match the experiment.

Figure 7.15 shows the resulting precision as a function of measurement time. Both experimental stages reach a balance between statistical uncertainty and the respective assumed systematical uncertainty after one year of exposure. Data collection campaigns of a few years (accounting for reactor schedule, experimental efficiencies, calibration periods) therefore appear realistic given these systematic goals. Table 7.1 summarizes the results of the calculation for net data collection periods of 1 and 3 years.

7.3.2. Weak mixing angle

Ignoring radiative corrections, a measurement of the $\text{CE}\nu\text{NS}$ cross-section in one isotope (as described by equation 1.33) can directly be converted into a measurement of $\sin^2 \theta_W$. The relative precision depends on the neutron-to-proton ratio of the isotope according to:

$$\frac{\Delta \sin^2 \theta_W}{\sin^2 \theta_W} = \frac{Z \cdot (4 \sin^2 \theta_W - 1) + N}{8Z \sin^2 \theta_W} \cdot \frac{\Delta \sigma}{\sigma} \quad (7.1)$$

Interpreting the results of section 7.3.1 as an expected 11.4% measurement of the $\text{CE}\nu\text{NS}$ cross-section using ^{184}W for NUCLEUS-10g, and an expected 1.15% measurement on ^{74}Ge for NUCLEUS-1kg, the uncertainties of the measurements of $\sin^2 \theta_W$ are 8.6% and 0.76% respectively. The dependence on isotopic contribution of the target is small. The momentum scale μ of the measurement was estimated from the relation $\mu^2 = 2M \cdot E_R$ to target mass M and recoil energy E_R , assuming a dominant contribution of recoils from 10 eV to 100 eV.

Figure 7.16 shows these projections in comparison with existing precision and neutrino-based measurements. The left panel focuses on the lower momentum transfer range and zooms out on the vertical axis to show all neutrino-based measurements. The right panel shows the running of $\sin^2 \theta_W$ predicted in the SM, zooming in to highlight the experiments constraining the scale dependence. In the first stage, NUCLEUS is not competitive with the most precise measurements, but can add to the neutrino-based low-momentum-exchange measurements. The full sensitivity of NUCLEUS-1kg will allow a

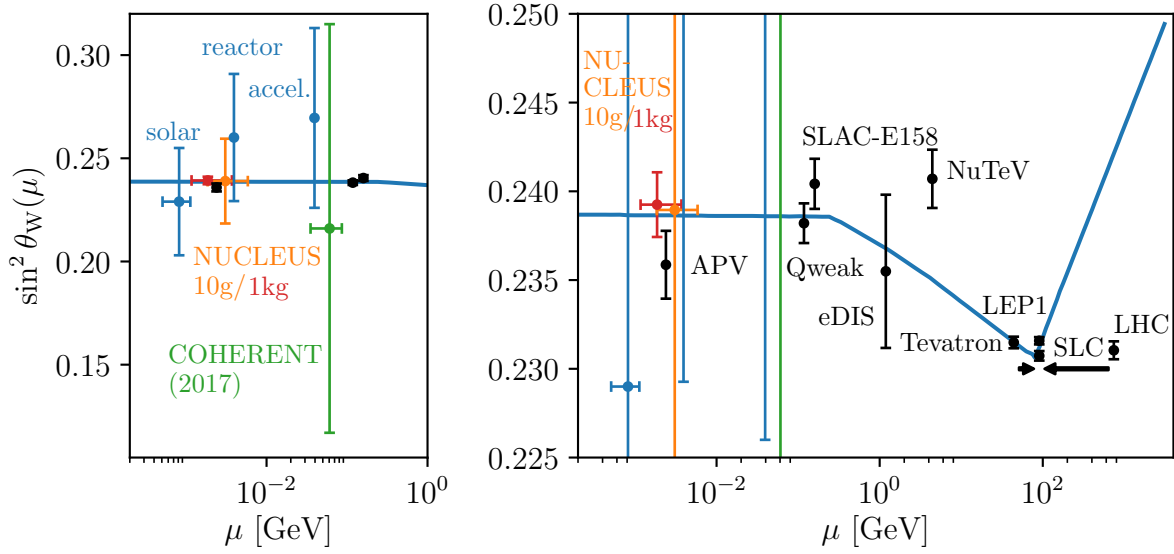


Figure 7.16.: Projected sensitivity of NUCLEUS to the weak mixing angle compared to existing measurements. The left figure is zoomed on the region of low momentum transfer, while the right plot is scaled to show the running of the weak mixing angle as predicted by the Standard Model. The continuous line shows the SM prediction in the $\overline{\text{MS}}$ renormalization scheme [137]. The black data points show the precision measurements listed in Fig. 10.2 of [129]. The blue data points show global fits to reactor- and accelerator-neutrino experiments [288] and an analysis of Borexino solar-neutrino data [289]. The green data point is the interpretation of the 2017 COHERENT result as a weak mixing angle-measurement by [290]. The orange (red) marker shows the expected precision with NUCLEUS-10g (NUCLEUS-1kg) after 3 years of live time.

neutrino-based constraint with similar precision, and in a comparable momentum range, to atomic parity violation measurements.

7.3.3. Non-standard Neutrino Interactions

CE ν NS phenomenology To describe the CE ν NS cross-section including these NSI contributions, the following substitution in Eqn. 1.12 is necessary:

$$\begin{aligned} [g_V^p \cdot Z + g_V^n \cdot N]^2 \longrightarrow & [(g_V^p + 2\varepsilon_{ee}^{uV} + \varepsilon_{ee}^{dV}) \cdot Z + (g_V^n + \varepsilon_{ee}^{uV} + 2\varepsilon_{ee}^{dV}) \cdot N]^2 + \\ & + [(2\varepsilon_{\tau e}^{uV} + \varepsilon_{\tau e}^{dV}) \cdot Z + (\varepsilon_{\tau e}^{uV} + 2\varepsilon_{\tau e}^{dV}) \cdot N]^2, \end{aligned} \quad (7.2)$$

where the parameters $\varepsilon_{\alpha\beta}^{qV}$ describe the additional amplitude for the scattering $q + \nu_\alpha \rightarrow q + \nu_\beta$. Note that the flavor-changing contribution has to be separately squared, as the process does not interfere with the SM process due to the distinct final state.

The NSI parameter space is high-dimensional and contains many degeneracies. We consider pairs of operators $(\varepsilon_{ee}^{uV}, \varepsilon_{ee}^{dV})$, $(\varepsilon_{\tau e}^{uV}, \varepsilon_{\tau e}^{dV})$, $(\varepsilon_{ee}^{dV}, \varepsilon_{\tau e}^{dV})$ which we allow free while setting all others to zero. The dependence of the CE ν NS cross-section of ^{184}W on these




Figure 7.17.: CE ν NS rate on the typical tungsten isotope ^{184}W relative to the SM rate, as a function of combinations of NSI parameters. Red signifies enhancement, blue suppression with respect to the SM. The black (green) lines indicate a rate equal to the SM prediction for ^{184}W (^{27}Al). Varying slopes appear for different isotopes due to the distinct relative number of protons and neutrons. Left: non-universal ν_e coupling on u,d-quarks. Interference with the SM occurs. Right: flavor-changing $\nu_e - \nu_\tau$ interaction on u,d-quarks. No interference with the SM, cross-section increases $\propto \varepsilon^2$.

parameters is shown in Figure 7.17. A single typical isotope is selected as the different proton-to-neutron number ratio leads to slightly different pictures for the isotopes of an element. Typically, the few isotopes in an element with appreciable natural abundance are similar enough that no measurable difference occurs by treating an element as mono-isotopic.

For $(\varepsilon_{ee}^{uV}, \varepsilon_{ee}^{dV})$ and $(\varepsilon_{\tau e}^{uV}, \varepsilon_{\tau e}^{dV})$, the CE ν NS rate on a given isotope is insensitive to NSI changes in the direction $\Delta\varepsilon^{uV}/\Delta\varepsilon^{dV} = -(A+N)/(A+Z)$. This is due to a cancellation of NSI contributions on neutrons and protons.

For $(\varepsilon_{\tau e}^{uV}, \varepsilon_{\tau e}^{dV})$ (top right panel in Fig. 7.17), one band of parameter values reproduces the SM cross-section (i.e. when NSI contributions internally cancel). Otherwise, this type of NSI can only enhance the CE ν NS cross-section. The enhancement is proportional to the square of the NSI parameters. For $(\varepsilon_{ee}^{uV}, \varepsilon_{ee}^{dV})$ (left panel in Fig. 7.17), interference with the SM CE ν NS reaction occurs, leading to two bands in parameter space where the SM cross-section is matched. The first comes from internal cancellation of NSI contributions, for the second solution the NSI amplitude has twice the size of the SM amplitude but opposite sign. Between the two bands the NSI amplitude can suppress and even completely cancel SM CE ν NS. Due to this interference, the CE ν NS cross-section is modified to first order in the NSI parameters, which makes CE ν NS particularly sensitive to this type of NSI.

The degeneracy in direction $k = -(A+N)/(A+Z)$ can be lifted by combining measurements on targets with different neutron-to-proton ratio. This is illustrated in Fig. 7.17 by comparing the black line (SM-degenerate parameter points for CE ν NS on ^{184}W , $k = -1.140$) and the green line (SM-degenerate parameter points for CE ν NS




Figure 7.18.: CE ν NS rate on ^{184}W relative to the SM rate, as a function of combinations of NSI parameters. Analogous to Figure 7.17, but showing non-universal ν_e versus flavor-changing $\nu_e - \nu_\tau$ interaction on d-quarks.

on ^{27}Al , $k = -1.025$). For $(\varepsilon_{ee}^{uV}, \varepsilon_{ee}^{dV})$, one SM-degenerate parameter point at $(-1 + \frac{8}{3} \sin^2 \theta_W, 1 - \frac{4}{3} \sin^2 \theta_W)$ remains, where both the neutron and proton contributions to the CE ν NS amplitude separately flip signs with respect to the SM. This ambiguity is specific to CE ν NS and the “second island” is already excluded by the LHC monojet constraint.

For the combination $(\varepsilon_{ee}^{dV}, \varepsilon_{\tau e}^{dV})$ (Fig. 7.18), where one component interferes with the SM and the other does not, lines of equal cross-sections are circles in the parameter space. The radius of the SM-degenerate circle again depends on neutron-to-proton ratio, so that a multi-target measurement can strongly improve constraints.

NUCLEUS constraints The NSI constraints projected for NUCLEUS at VNS, given the observation of the SM expected CE ν NS signal, are shown in Figures 7.19 and 7.20. Exposures of one year with Al_2O_3 and CaWO_4 are assumed for NUCLEUS-10g, and also one kg-year of exposure using germanium for NUCLEUS-1kg. In addition, the Figure shows the potential of a kg-year exposure using silicon as a target (assuming background as in Al_2O_3 and systematics as in germanium), to underline the potential of a combined precision measurement using targets of widely different neutron-to-proton ratio.

The calculation uses counting statistics only, treating systematics as an unknown signal normalization as described in Appendix B. CE ν NS on all isotopes present in the targets (using the isotopic abundances from [292]) is taken into account. This does not appreciably alter the picture discussed so far for single isotopes, as all targets have either similar nuclei in terms of k (Al_2O_3 , Ge, Si) or a single element dominating the cross-section (CaWO_4).

Table 7.2 summarizes the constraints expected on individual NSI parameters (i.e. setting all parameters but one to zero). The center of the degenerate second solution for $\varepsilon_{ee}^{u,d}$ is indicated separately, where applicable. The combination of Al_2O_3 and CaWO_4 in NUCLEUS-10g is sufficient to exclude the second solution for ε_{ee}^{uV} at 90% C.L. assuming $\varepsilon_{ee}^{dV} = 0$, see Figure 7.19. The second solution for ε_{ee}^{dV} assuming $\varepsilon_{ee}^{uV} = 0$ is only dis-

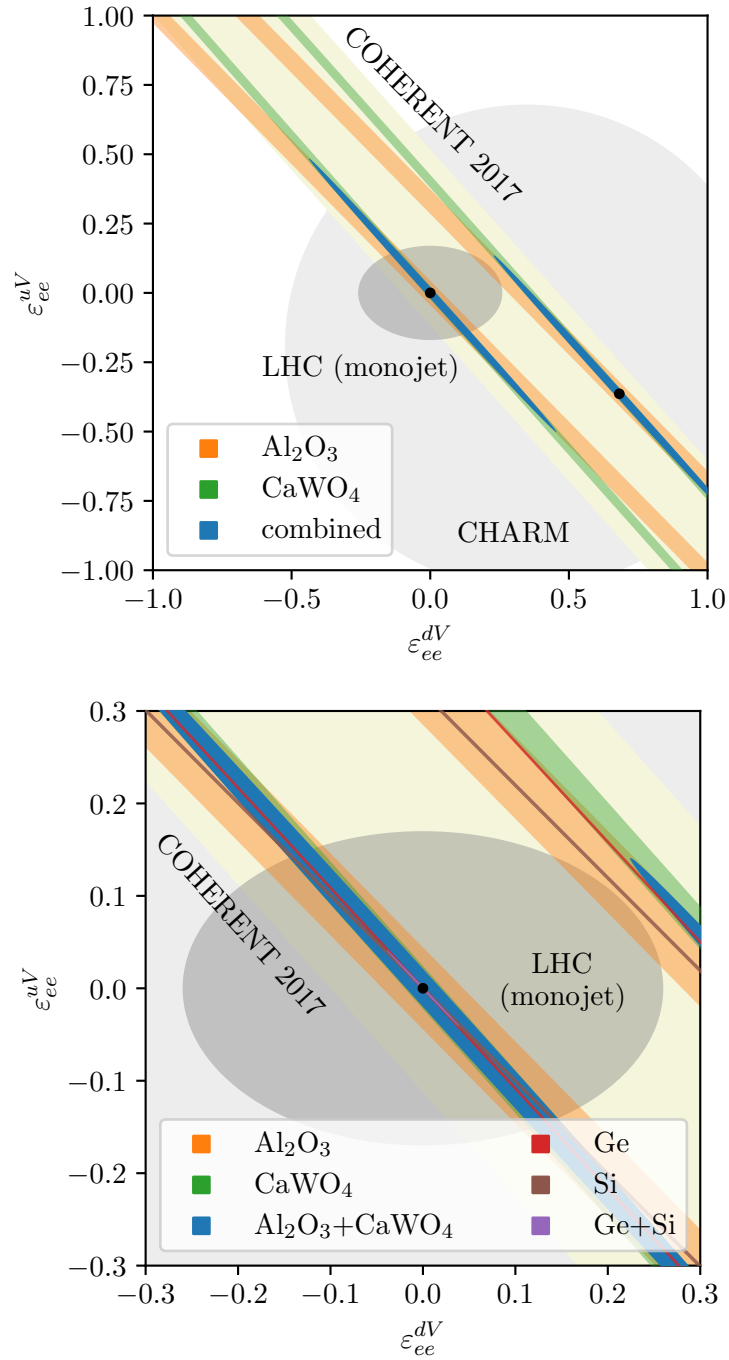


Figure 7.19.: Projected constraints (90% C.L.) on combinations of NSI parameters ϵ_{ee}^{dV} and ϵ_{ee}^{uV} from NUCLEUS compared to existing measurements. Top panel: NUCLEUS-10g constraints from Al_2O_3 and CaWO_4 . Bottom panel: zoom showing additionally the constraint from NUCLEUS-1kg using germanium and silicon. Existing constraints come from the CHARM [144] experiment [140], LHC monojet studies [141] and the COHERENT observation [291].

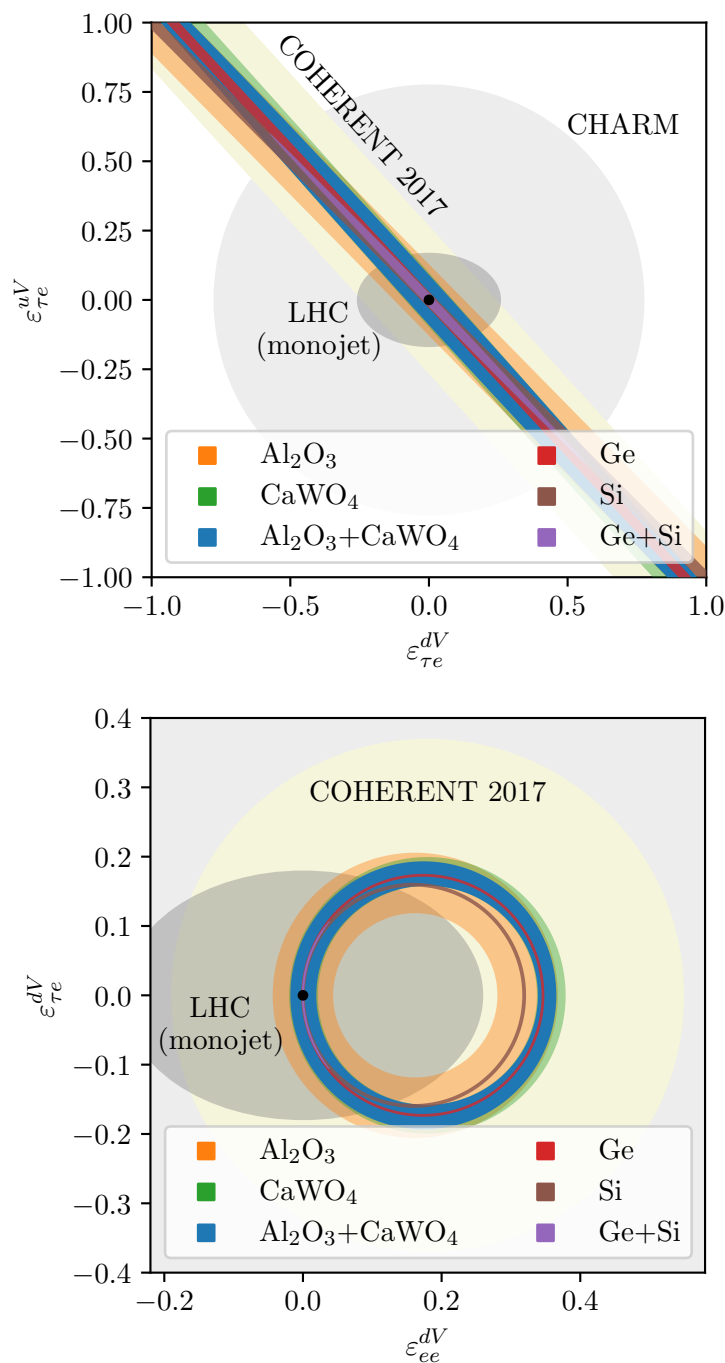


Figure 7.20.: Projected constraints (90% C.L.) from NUCLEUS (both phases) on flavor-changing NSI parameters. Top panel: flavor-changing $\nu_\tau \rightarrow \nu_e$ interactions on u- and d-quarks. Bottom panel: flavor changing versus non-universal coupling to d-quarks. For $\text{Al}_2\text{O}_3 + \text{CaWO}_4$, 10 g · y total exposure and 10% signal normalization uncertainty was assumed, while limits for Ge and Si are calculated with 1 kg · y exposure each and 1% signal normalization uncertainty. Existing constraints as in Figure 7.19.

material	$ \varepsilon_{ee}^{uV} $	deg.	$ \varepsilon_{ee}^{dV} $	deg.	$ \varepsilon_{\tau e}^{uV} $	$ \varepsilon_{\tau e}^{dV} $
Al ₂ O ₃	< 0.044	0.329	< 0.044	0.324	< 0.129	< 0.127
CaWO ₄	< 0.024	0.403	< 0.022	0.358	< 0.098	< 0.087
Al ₂ O ₃ +CaWO ₄	< 0.021	–	< 0.019	0.352	< 0.090	< 0.082
Si	< 0.0027	0.320	< 0.0027	0.319	< 0.0296	< 0.0295
Ge	< 0.0022	0.375	< 0.0020	0.346	< 0.0286	< 0.0284
Si + Ge	< 0.0017	–	< 0.0016	–	< 0.0245	< 0.0233

Table 7.2.: NSI constraints (90% C.L.) on individual parameters achievable with NUCLEUS (both phases). For each limit, all other NSI parameters are fixed to zero. The larger (Al₂O₃, CaWO₄) intervals for $|\varepsilon_{ee}^{u,dV}|$ are slightly asymmetric and have been enlarged to allow the shorthand notation.

favoured at 75% C.L. From a kg-year CE ν NS measurement using both Ge and Si, a large improvement in NSI limits is expected, and no degeneracies remain at a single-operator level.

In summary, NUCLEUS-10g will allow unprecedented NSI constraints for ν_e -quark interactions beyond the current limits from the first CE ν NS observation by COHERENT and LHC monojet studies. NUCLEUS-1kg can probe below 0.2% in non-universal ν_e -quark couplings and below 3% for flavor-changing ν_e - ν_τ interactions. These numbers are comparable with the sensitivities expected from a future Neutrino Factory [140], where NSI would be probed in deep inelastic scattering. The coherence property of CE ν NS leads to a unique dependence of the cross-section on individual NSI couplings, enhancing the sensitivity strongly for those parameters where the NSI contribution interferes with the SM process.

7.3.4. Neutrino electromagnetic form factors

Neutrino charge radius in NUCLEUS Treating the charge radius as a NSI contribution with $\varepsilon_{ee}^{uV} = 0.0176 \cdot \langle r_{\nu_e}^2 \rangle / 10^{-32} \text{ cm}^2$, $\varepsilon_{ee}^{dV} = -0.0088 \cdot \langle r_{\nu_e}^2 \rangle / 10^{-32} \text{ cm}^2$, we can use the calculation of section 7.3.3 to derive the sensitivity of NUCLEUS. The 90% C.L. interval for NUCLEUS-10g has an extent of $\pm 3.2(5.1) \cdot 10^{-32} \text{ cm}^2$ around the expectation for CaWO₄ (Al₂O₃). Thus no observation of the SM charge radius $\langle r_{\nu_e}^2 \rangle = -0.83 \cdot 10^{-32} \text{ cm}^2$ is expected in this phase. For NUCLEUS-1kg, the interval sizes are $\pm 0.27(0.31) \cdot 10^{-32} \text{ cm}^2$ for Ge (Si). This should allow a 5.0σ measurement of $\langle r_{\nu_e}^2 \rangle$ in the standard scenario of NUCLEUS-1kg.

While the SM charge radius discussed so far is associated with one neutrino flavor, also flavor-changing transition charge radii are possible (although not allowed in the SM) [156]. The only limits in the literature are derived from COHERENT data: $(|\langle r_{\nu_e \mu}^2 \rangle|, |\langle r_{\nu_e \tau}^2 \rangle|, |\langle r_{\nu_\mu \tau}^2 \rangle|) < (22, 38, 27) \cdot 10^{-32} \text{ cm}^2$ [156]. As with flavor-changing NSI, also in this case the cross-section adds incoherently to the SM cross-section due to the distinct final state. As above for the flavor-diagonal charge radii, the cross-section can be translated to an NSI one (leaving the neutron coupling undisturbed) with $\varepsilon_{\tau e}^{uV} = 0.0176 \cdot \langle r_{\nu_{el}}^2 \rangle / 10^{-32} \text{ cm}^2$, $\varepsilon_{\tau e}^{dV} = -0.0088 \cdot \langle r_{\nu_{el}}^2 \rangle / 10^{-32} \text{ cm}^2$ for flavors $l = \mu, \tau$. The NSI sensitivity derived for NUCLEUS above in section 7.3.3 can thus be translated into limits on flavor-changing transition charge radii $|\langle r_{\nu_{el}}^2 \rangle| < (14.9, 12.7, 3.55, 3.27) \cdot 10^{-32} \text{ cm}^2$ for $l = \mu, \tau$ and target materials (Al₂O₃, CaWO₄, Ge, Si). Note that silicon is expected




Figure 7.21.: Coherent neutrino-nucleus scattering induced by electromagnetic neutrino form factors compared to the standard weak process. Recoil rates are shown for CaWO_4 (blue) and Al_2O_3 (orange). The standard weak $\text{CE}\nu\text{NS}$ rate is shown as solid lines. The colored bands indicate the modification of the $\text{CE}\nu\text{NS}$ rate via a non-standard neutrino charge radius in the range allowed by laboratory limits ($-4.2 < \langle r_{\nu_e}^2 \rangle / 10^{-32} \text{ cm}^2 < +6.6$). The dashed lines show the contribution from a neutrino magnetic dipole moment at the current experimental sensitivity. The dotted lines depict scattering induced by a neutrino charge at the level allowed by direct measurements.

to yield the best limit in NUCLEUS-1kg, as the precision in our scenario is mostly systematics-limited and the higher proton-to-neutron ratio increases the sensitivity of a light target in this case. This is apparent in Figure 7.20 (top), as the Si allowed region is of comparable width but shallower than that of Ge, while the NSI-direction along which the $\langle r_{\nu_{el}}^2 \rangle$ act is steeper, giving Si a smaller overlap.

Neutrino magnetic dipole moment in NUCLEUS The differential cross-section for neutrino magnetic dipole scattering on spin-zero nuclei is [159]:

$$\frac{d\sigma}{dT} = \frac{\pi\alpha^2}{m_e^2} \cdot Z^2 \cdot \left[\frac{1}{T} - \frac{1}{E_\nu} + \frac{T}{4E_\nu^2} \right] \cdot \left(\frac{\mu_\nu}{\mu_B} \right)^2 \quad (7.3)$$

For nuclei with spin, the third term in the square brackets is absent. In any case, it is negligible for $\text{CE}\nu\text{NS}$ as it is suppressed by at least $(E_\nu/M)^2$ or six orders of magnitude with respect to the first term. In addition to the scattering of the magnetic dipole in the Coulomb field, there is also dipole-dipole scattering for nuclei with spin, but this interaction is likewise suppressed. For a $\text{CE}\nu\text{NS}$ experiment, it is therefore justified to consider only the first term and neglect nuclear details.

Adding this cross-section to the one of $\text{CE}\nu\text{NS}$ and varying μ_ν , a likelihood-ratio test as described in Appendix B.6 produces the 90% C.L. upper limit on μ_ν in the




Figure 7.22.: Expected limits on the effective electron antineutrino magnetic dipole moment from the NUCLEUS experiment. Limits from NUCLEUS-10g (CaWO_4 , Al_2O_3) and -1kg (Ge, Si) are calculated as a function of detector threshold. The dashed horizontal line shows the best current laboratory limit $\mu_\nu < 2.9 \cdot 10^{-11} \mu_B$ from neutrino-electron scattering [157]. NUCLEUS-10g lacks the statistics necessary to improve on the existing constraint. For thresholds below 10 eV (25 eV), the NUCLEUS-1kg Ge and Si detectors have the potential to set a new best limit.

given experimental scenario. In Figure 7.22, the resulting magnetic dipole limits for NUCLEUS-10g and -1kg are shown as a function of the experimental threshold. As to be expected from Figure 7.21, the limits in NUCLEUS-10g fall in the $10^{-10}\mu_B$ -range. For CaWO_4 , the $\text{CE}\nu\text{NS}$ rate is more than one order of magnitude above the maximum allowed one from dipole scattering down to 10 eV. For Al_2O_3 , the 4 g · y exposure is not sufficient to detect a potential contribution from dipole scattering at lowest energies. The increased statistics of NUCLEUS-1kg changes the picture. For thresholds below 10 eV in germanium (25 eV in silicon), NUCLEUS-1kg can achieve leading sensitivity on the effective reactor $\bar{\nu}_e$ magnetic dipole moment. Unlike the precision of the $\text{CE}\nu\text{NS}$ cross-section, the dipole limit will still be limited by signal statistics after 1 kg year exposure, as only a small energy range close to threshold contributes to the dipole signal. As the scattering rate is proportional to μ_ν^2 , improvement of statistical sensitivity by a factor 2 requires 16 times the measurement time.

Uncertain signal normalization does not strongly impact the dipole limit, as the concurrent $\text{CE}\nu\text{NS}$ observation will provide a neutrino flux measurement, and the dipole moment sensitivity is driven by spectral-shape information rather than rate information. By the same token, threshold effects and energy resolution have a potential impact on a real measurement. In this simplistic study, they are ignored apart from choosing a bin size on the order of the expected energy resolution ($E_t/5$).

In summary, a kg-year exposure of NUCLEUS detectors with thresholds in the 10 eV range has the potential to achieve leading sensitivity on the effective neutrino magnetic dipole moment via neutrino-nucleus scattering. One should keep in mind that such a measurement intensifies all the challenges associated with a $\text{CE}\nu\text{NS}$ measurement: low absolute rates, small recoil energies and discrimination from backgrounds near detector thresholds.

7.3.5. New Light Mediators

Light Scalar Mediators For the light scalar mediator Φ , we assume an interaction with SM fermions via the Lagrangian [162]:

$$\mathcal{L}_\Phi = \Phi \left[g_\nu \bar{\nu}_R \nu_L + g_\nu^* \bar{\nu}_L \nu_R + \sum_q g_q \bar{q} q \right]. \quad (7.4)$$

We consider couplings to neutrinos and quarks, necessary to create a $\text{CE}\nu\text{NS}$ signal. For simplicity we take the couplings of the scalar to all quarks to be identical. Couplings to charged leptons are not considered here, as they are not relevant to $\text{CE}\nu\text{NS}$. Assumptions about the scalar coupling to charged leptons are decisive when comparing $\text{CE}\nu\text{NS}$ bounds to other laboratory constraints.

Neutrino interactions with a scalar mediator are chirality-changing and do not interfere with SM Z-exchange. In addition to the lepton-number conserving Lagrangian involving right-handed neutrinos shown here, there is a second (lepton-number violating) possibility discussed in detail in [163]. It leads to the same scattering cross-section but comes with different constraints from supernova explosions and neutrinoless double-beta decay.

The additional neutrino-nucleus scattering cross-section from scalar exchange is given by [293, 163]:

$$\frac{d\sigma_\Phi}{dT} = \frac{g_\nu^2 Q_\Phi^2}{4\pi} \cdot \left(\frac{Mc^2}{E_\nu}\right)^2 \cdot \frac{(\hbar c)^2 \cdot T}{(2Mc^2 \cdot T + (m_\Phi c^2)^2)^2} \cdot F^2(q^2) \quad (7.5)$$

with the exchanged momentum $q^2 = 2Mc^2 \cdot T$ and the nuclear charge under Φ exchange Q_Φ . Q_Φ can be derived from nucleon scalar form factors f_q^N which are calculated in chiral perturbation theory for first-generation quarks [294] and from lattice QCD for strange quarks [295]. Following the values collected in [163], but setting the Φ -couplings to all quarks equal, we find

$$Q_\Phi \approx (17.1 \cdot Z + 17.1 \cdot N) \cdot g_q \quad (7.6)$$




Figure 7.23.: CE ν NS spectra enhanced by the presence of a light scalar mediator. Three different models with varying scalar mass m_Φ and effective coupling $g_\Phi \equiv \sqrt{g_\nu \cdot g_q}$ are compared to the standard CE ν NS recoil spectrum on CaWO₄.

Equation 7.5 describes an enhancement of the CE ν NS recoil spectrum as a function of the two parameters m_Φ and $g_\Phi \equiv \sqrt{g_\nu \cdot g_q}$. The shape of the scalar-induced recoil spectrum is distinct from the CE ν NS one, and as long as $m_\Phi^2 \lesssim q^2$ the spectral shape carries information about m_Φ . Figure 7.23 shows some examples for combined CE ν NS+scalar recoil spectra for different parameters.

With the same likelihood framework used throughout section 7.3, projected constraints can be calculated for the parameter g_Φ as a function of m_Φ . In Figure 7.24 the median sensitivity is shown for the target materials of NUCLEUS-10g and -1kg. For comparison, we show the constraint derived from the CE ν NS observation of COHERENT (as calculated in [164], but rescaled to the coupling assumptions and Q_Φ used here). Additional constraints are placed by the reactor experiment CONNIE [186] without observing a CE ν NS signal (drawn in purple). The dashed line and shaded region correspond to different assumptions regarding the quenching factor in silicon at low energies (for details see [186]). While the published limit ends at $m_\Phi = 1 \text{ MeV}/c^2$, the sensitivity extends




Figure 7.24.: Expected sensitivity of NUCLEUS to light scalar mediators. Median exclusion curves (90% C.L.) are shown for one year of exposure for NUCLEUS-10g (CaWO_4 , Al_2O_3) and -1kg (Ge, Si). The grey region shows the parameter space excluded by the first $\text{CE}\nu\text{NS}$ observation by COHERENT [164]. The purple region and dashed line show the limit from CONNIE [186] discussed in the text.

almost unchanged to lower masses, as the cross-section becomes independent of m_Φ for $m_\Phi^2 \ll q^2$.

For large scalar masses, as the mediator becomes kinematically inaccessible, the limit rises linearly with m_Φ . This results from the simpler mass-dependence of the cross-section formula, which becomes proportional to $(g_\Phi/m_\Phi)^4$ for $m_\Phi^2 \gg q^2$. The “turning point”, above which the scalar interaction can be described as a four-point contact interaction, is therefore given by the maximum momentum exchange or twice the neutrino energy used in the experiment. This explains why the COHERENT limit follows this dependence only above ~ 100 MeV, compared to ~ 10 MeV for reactor experiments.

For small scalar masses $m_\Phi^2 \ll q^2$, the cross-section becomes independent of the scalar mass and proportional to $1/(TE_\nu^2)$. This shows the advantage of reactor neutrinos compared to accelerator neutrinos for this physics observable, where ~ 10 times smaller neutrino energies are used and ~ 100 times smaller recoil energies are observed, leading to about an order of magnitude better sensitivity to g_Φ for equal measurement precision. At large scalar masses, the cross-section becomes proportional to T/E_ν^2 , rising with recoil energy (see the red line in Figure 7.23) and independent of E_ν at the kinematic endpoint $T_{max} = 2E_\nu^2/M$. Therefore, observing lower energy neutrinos and smaller recoil energies has no intrinsic advantage for the case of a heavy scalar mediator and the achieved measurement precision decides over the sensitivity.

The fact that a $\text{CE}\nu\text{NS}$ experiment can only obtain information about the scalar mass in the intermediate range, where q^2 is on the same order as m_Φ^2 , underlines the importance of precise $\text{CE}\nu\text{NS}$ measurements both at reactors and accelerator facilities, where different q^2 -regions are probed.

Light Vector Mediators For vector mediators Z' , the interaction is indistinguishable from standard $\text{CE}\nu\text{NS}$. Therefore the amplitudes of the two processes have to be added and interference between Z - and Z' -exchange occurs. We consider a Lagrangian

$$\mathcal{L}_{Z'} = Z'_\mu \cdot \left[\sum_{q=u,d} g_q \cdot \bar{q} \gamma^\mu q + g_\nu \cdot \bar{\nu} \gamma^\mu \nu \right] \quad (7.7)$$

for the interaction, focusing on vector (γ^μ) rather than axial ($\gamma_5 \gamma^\mu$) currents. As for $\text{CE}\nu\text{NS}$, the vector nuclear charge dominates, while spin-dependent axial effects are absent or smaller by $\sim A^2$. By setting the couplings to first-generation quarks equal $g_u = g_d \equiv g_q$, the nuclear charge under Z' -exchange becomes:

$$Q_{Z'} = [(2g_u + g_d) \cdot Z + (g_u + 2g_d) \cdot N] \cdot g_\nu = 3A \cdot g_q g_\nu \quad (7.8)$$

As the Z' - and Z -contributions interfere, care has to be taken with the relative sign. The couplings $g_V^{p,n}$ introduced in section 6.1.1 are for neutrinos, there is a sign-flip for reactor antineutrinos:

$$Q_Z = [\bar{g}_V^p \cdot Z + \bar{g}_V^n \cdot N] = [-g_V^p \cdot Z - g_V^n \cdot N]. \quad (7.9)$$

This has been irrelevant in the discussion up to this point, as the nuclear charge gets squared for standard $\text{CE}\nu\text{NS}$. The modified recoil spectrum can be recovered from the $\text{CE}\nu\text{NS}$ -one with the following substitution, taking into account the $(m_{Z'}^2 + q^2)$ momentum-dependence of the light vector interaction [165, 164]:

$$\left(\frac{G_F}{(\hbar c)^2} \cdot [\bar{g}_V^p \cdot Z + \bar{g}_V^n \cdot N] \right)^2 \rightarrow \left(\frac{G_F}{(\hbar c)^2} \cdot [\bar{g}_V^p \cdot Z + \bar{g}_V^n \cdot N] - \frac{\hbar c}{\sqrt{2}} \cdot \frac{3A \cdot g_q g_\nu}{2Mc^2 \cdot T + m_{Z'}^2 c^4} \right)^2. \quad (7.10)$$

Like in the case of a scalar mediator, the Z' -model in the context of $\text{CE}\nu\text{NS}$ contains two parameters, $m_{Z'}$ and $g_{Z'} \equiv \sqrt{g_\nu \cdot g_q}$. The phenomenology is richer than in the scalar case, as interference can result in suppression, enhancement and distortion of the observable recoil spectrum. Example cases for three interesting parameter combinations are shown in Figure 7.25.

The projected limits for NUCLEUS-10g and -1kg are shown in Figure 7.26, again compared to the limit derived from the COHERENT measurement in [164] and the constraints from CONNIE [186]. The same points apply regarding quenching factors and low-mass sensitivity as in the vector case.

For heavy mediators $m_{Z'} \ll q^2$, the scattering again turns into a contact interaction where only the spectral normalization is influenced. This case is exactly degenerate with flavor-conserving NSI (with a particular choice regarding the relative coupling strengths to u- and d-quarks). In analogy to the SM-degenerate band in flavor-conserving NSI, the Al_2O_3 -limit in Figure 7.26 has a thin gap opening up above ~ 10 MeV: a heavy scalar mediator can reproduce the SM $\text{CE}\nu\text{NS}$ cross-section by flipping the overall sign of the amplitude. In the band below the slice, the destructive interference is weaker, leading to a detectable suppression of the rate. Below 10 MeV, the slice disappears, as any detectable Z' -coupling causes significant spectral distortions.

Slices should also appear in the limits from other target materials, but they are not captured by the grid resolution used in the likelihood study. They are slightly displaced




Figure 7.25.: $\text{CE}\nu\text{NS}$ spectra modified by the presence of a light vector mediator interfering with the SM process. Three example cases with different values for mediator mass $m_{Z'}$ and coupling to fermions $g_{Z'}$ are selected to visualise the possible distortion, enhancement and suppression of the SM $\text{CE}\nu\text{NS}$ rate on CaWO_4 .




Figure 7.26.: Expected sensitivity of NUCLEUS to light vector mediators, analogous to Figure 7.24. Median exclusion curves (90% C.L.) are shown for one year of exposure for NUCLEUS-10g (CaWO_4 , Al_2O_3) and -1kg (Ge, Si). The grey region shows the parameter space excluded by the first $\text{CE}\nu\text{NS}$ observation by COHERENT [164]. The purple region and dashed line show the limit from CONNIE [186].

with respect to one another due to the varying neutron content in the target affecting the $CE\nu NS$ amplitude. They are also much thinner due to the higher $CE\nu NS$ statistics in the other target materials. In principle, multi-element or even just multi-isotope targets allow breaking such degeneracies (as with NSI), but only at high measurement precision.

Overall, *NUCLEUS*-10g will allow an improved constraint on heavy vector mediators by around a factor of two with respect to *COHERENT*. For sub-MeV vector mediators, this increases to 1.5 orders of magnitude improvement. With *NUCLEUS*-1kg, constraints at all mediator masses can again be improved by a factor of 5 with respect to *NUCLEUS*-10g.

At high masses, these constraints rely on information about the signal normalization (e.g. a precise neutrino flux prediction for *NUCLEUS*). At low masses, spectral distortions have to be controlled using precise knowledge of calibration and detector linearity.

8. The NUCLEUS-1g prototype

The testing of cryogenic detectors for NUCLEUS began in 2017 with the first individual target cubes. After good performance was demonstrated, development of a “1g prototype” was started. The design goals were a detector holder scalable to many target cubes, including the cryogenic veto detectors necessary for the NUCLEUS experiment. The NUCLEUS-1g prototype consists of a single target cube combined with inner and outer vetos. This prototype allows to study the performance and interplay of the cryogenic detectors of the NUCLEUS experiment. The mechanical, electrical and thermal design was validated with much of the complexity necessary for the first science phase of NUCLEUS. Anticoincidence vetoing and background suppression using the inner veto was shown in a proof-of-principle. The solid angle around the target detector was not completely covered. This chapter introduces the NUCLEUS-1g prototype, detailing its components, the measurement environment and finally the experimental runs with their results. For the first prototype run, analysis methods are presented in great detail. For the other three, emphasis is put on explaining the results obtained.

The milestones achieved with the prototype runs are:

- **Prototype Run 1:** First operation of a NUCLEUS target detector. Energy threshold of 19.7 eV, new constraints on low-mass dark matter
- **Prototype Run 2:** Demonstration of the flexible silicon holder. Energy resolution improved over Run 1. Low-energy background measurement with and without calibration source
- **Prototype Run 3:** Assembly of the full prototype. Outer vetos operated for the first time.
- **Prototype Run 4:** Simultaneous operation of target and inner vetos. Demonstration of anticoincidence surface veto.

8.1. Prototype components

The prototype aims to demonstrate the function and interplay of the cryogenic detector components for NUCLEUS. To this end, the cryogenic vetos are implemented around a single target detector. With outer dimensions of ~ 8 cm diameter and ~ 12 cm height and a total of 5 SQUID channels (target + $2\times$ inner veto + $2\times$ outer veto) needed for readout, the prototype can be operated in standard wet cryostats used for detector characterisation above ground. Figure 8.1 shows an overview of the assembled prototype and its components.




Figure 8.1.: Components of the NUCLEUS-1g prototype. Center: cut-away of CAD drawing. Left: photograph of the assembled detector module showing the outer veto detectors. Right: photograph of the central inner detector. From top to bottom: thick silicon support wafer (a) with electrical and thermal contacts (b), sapphire balls (c) as spacers, top inner veto (d) with etched pyramids (e) and window for target contacts, target detector (f), bottom inner veto (g). The tension of the inner vetos is adjusted via the precise machining of the vertical silicon standoffs (h) setting the overall height of the assembly. Figure taken from [296].

8.1.1. Target detectors

Target detectors for the NUCLEUS prototype were produced three at a time from $(20 \times 10 \times 5)$ mm³ substrates, shown in Figure 8.2 (left). The fabrication procedure consists of alternating deposition and lithographic patterning of tungsten, aluminum and gold layers, followed by dicing of the substrate into individual detectors and polishing. This procedure can be scaled up to larger substrates without increase in complexity, therefore many more target cubes can in principle be produced in a single fabrication cycle.

Three TES structures were designed for testing on prototype targets: “light detector” (LD), “medium” (M) and “small” (S), shown in Figure 8.3. The feature sizes are listed in Table 8.1. The LD structure is functionally identical to the CRESST-III light detector TES in use on $(20 \times 20 \times 0.4)$ mm³ Al₂O₃ plates. The layout is rearranged to allow contacting through the narrow window of an inner veto. With its large surface area and fast relaxation time, this TES structure is not optimized for NUCLEUS targets. Its motivation is the proven reliability of manufacture, it was intended as a starting point for NUCLEUS TES optimization. The S structure is correctly scaled to cover the same surface area fraction as CRESST phonon TES, aiming for a similar phonon collection fraction. To account for the very small TES area and heat capacity, an extremely weak thermal link is needed. This is realized by a circular gold stripe of only 10 μm width. The M structure was introduced to hedge against potential fabrication issues with “S”. It covers the same total surface area, but uses twice the tungsten area (higher heat capacity) which allows to use a 20 μm thermal link. All three structures have been fabricated with good reliability and operated cryogenically.

While target detectors from Al₂O₃, CaWO₄ and Si have been produced, all prototype measurements so far were using Al₂O₃ targets, which is the easiest material to handle. Due to its high hardness, Debye temperature and sound speeds, a moderately




Figure 8.2.: Target cubes for the NUCLEUS prototype. Left: substrates under processing, with tungsten layers for three targets each. Right: finished Al_2O_3 targets with three different TES designs.

better detector performance is expected for Al_2O_3 e.g. compared to CaWO_4 , so that a characterisation of all target materials is needed in the future.

8.1.2. Inner Veto

The purpose of the inner veto is to fully surround the target detectors with instrumented surfaces, at the same time providing a defined force for holding the target. In the prototype, this is solved by pressing the target detector between two elastic silicon wafers of $200\ \mu\text{m}$ thickness, shown in Figure 8.4. The inner surfaces of the veto wafers are equipped with pyramidal structures by wet etching. The pyramid tips leave small patches of original wafer surface providing a point-like mechanical contact to the target. This design aims at a reproducible suppression of thermal contact between target and inner veto. While the prototype inner veto does not fully cover the target detector, it allows testing the operation of cryogenic detectors in mechanical contact, as necessary for this vetoing strategy. For the NUCLEUS experiment, one of the flexible wafers will be replaced with a milled silicon beaker, surrounding the targets and covered by the flexible wafer.




Figure 8.3.: Overview of TES structures used in CRESST and NUCLEUS. From left to right: CRESST-III phonon L, CRESST-III light detector, NUCLEUS LD, NUCLEUS medium, NUCLEUS small. Tungsten layers appear blue, aluminum gray and gold yellow.

Feature size [μm]	S	M	LD
W length	50	72	140
W width	100	144	300
Al length	210	210	500
Al width	320	320	1000
Au length	3000	3000	1500
Au width	10	20	40

Table 8.1.: Feature sizes of NUCLEUS TES structures “S” (small), “M” (medium), “LD” (adapted CRESST-III light detector TES). Film thicknesses are 200 nm for W, 50 nm for Au and 1 μm for Al.

8.1.3. Holding Structure

A holding structure for the “inner detector” (target + inner vetos) is constructed of 1 mm silicon wafers (see Figure 8.5). The structure allows the inner detector to be supported independently from the surrounding outer veto, important to avoid mechanical contact. It carries electrical and thermal connections formed by sputtered aluminum and gold layers. The backside of the upper holding wafer is sputtered with a thick gold layer to improve thermal equilibration. Target and upper inner veto are contacted via wire-bonds through a central window in the upper support wafer. A row of copper posts feeds the connections to the bottom support wafer for contacting the bottom inner veto.

For assembly, the upper and lower holding wafer are screwed together at four corners, with precisely machined silicon posts ensuring a reproducible spacing. Each holding wafer contacts its inner veto plate via three Al_2O_3 balls glued to its inside (see Figure 8.1, right). The height of the spacer posts determines the mechanical load on the inner veto wafers, which are under tension to act as an elastic holder for the target detector. Tests have shown that the inner vetos can be safely deformed up to about 50 μm . Within that range, friction solidly holds the inner detector together and prevents slipping of the target cube. Initial low-temperature tests using brass screws sometimes led to cracking holding wafers, blamed on differential thermal contraction and mechanical imperfection




Figure 8.4.: Left: Prototype inner veto silicon wafers ($11 \text{ mm} \times 11 \text{ mm} \times 200 \mu\text{m}$). The top veto has a laser-cut $2.5 \text{ mm} \times 2 \text{ mm}$ window for routing contacts to the target. The inner surfaces are patterned with pyramid structures in a wet-etching process. Right: SEM image of a pyramid (taken at MPG-HLL), showing a $\sim 100 \mu\text{m} \times 100 \mu\text{m}$ patch of original wafer surface at the tip. The white scale bar is $100 \mu\text{m}$.

on the brass surfaces facing the silicon. This could be overcome by switching to silicon screws and using a controlled torque of 10 mNm for tightening the screws.

8.1.4. Outer Veto

The outer veto is designed to function as a massive barrier surrounding target and inner veto. Its purpose is to allow anti-coincidence vetoing double-scattering backgrounds such as gamma rays and, to a lesser extent, neutrons. Due to its large size, it has to operate at a high particle count rate (many per second). Its energy resolution will not be comparable to the inner detectors, instead, its time resolution is an important parameter. For later stages of NUCLEUS, a heavy-element target is preferred. This increases the interaction cross-section with gamma rays. For the prototype, silicon was selected as the outer veto material due to its easier availability. The outer veto consists of two cylinders (diameter 5 cm , height 5 cm) that can be stacked with a gap of 0.5 mm using three Al_2O_3 balls fitting in grooves (see Figure 8.6). Both halves have a milled cavity that houses the inner detector without mechanical contact.

The TES design for the outer veto (see Figure 8.6, right) is adapted to the characteristics of the outer veto. For technical reasons, the TES is fabricated on a small carrier chip and then glued to the large outer veto crystal. The TES itself shares most features with CRESST-II phonon TES, with the goal of speeding up the response to particle energy depositions. The thermal contact is not provided via a weak link (sputtered gold stripe) but through a bond wire directly on the tungsten film. This makes the TES relaxation time faster than the athermal phonon burst following a particle energy deposition. This regime is called “bolometric operation” [109] (see section 2.3.1); all other TES discussed here are designed for “calorimetric operation”. The relaxation time of a particle pulse is therefore given by the athermal phonon lifetime in the crystal. Phonon collectors are omitted, as they are expected to release energy into the TES by quasiparticle diffusion on a slower timescale, and therefore do not contribute to the pulse amplitude.




Figure 8.5.: Holding structure for the NUCLEUS-1g prototype. Silicon wafers (a and b, 1 mm thickness) provide mechanical support, while only the inner veto plates (c and d) touch the target. a) top support wafer, with wings allowing to suspend the inner detector inside the outer veto cavity without touch. Electrical and thermal connections are patterned from aluminum and gold layers. b) bottom support wafer. A row of copper posts (not pictured) feeds the connections for the bottom veto from the top to the bottom support wafer. The inner veto plates are held by the support wafers via 1 mm Al_2O_3 spheres (not pictured). c) top inner veto with cut-out for target bond-wires. d) bottom inner veto plate.

8.2. Experimental setup

8.2.1. Cryogenic infrastructure

Measurements with the NUCLEUS prototypes were performed in the cryogenic detector laboratory at the Max-Planck-Institut für Physik, Munich. Two wet dilution refrigerators are in operation: “cryostat 1”, an Oxford Instruments Kelvinox100, and “cryostat 2”, a Leiden Cryogenics refrigerator with smaller experimental volume. Cryostat 1 is housed inside a Faraday cage and equipped with six SQUID channels (APS-581, Supracon, Magnicon) while cryostat 2 operates without an EMI shield and houses 4 APS-581 SQUID channels.

8.2.2. Data Acquisition

Both setups can be used with two independent data acquisition systems: a classical “hardware-triggered” DAQ saving pulse traces using an Incaa Computers VD80 transient recorder, and a “software-triggered” data acquisition that saves the continuous data stream from a detector to disc (using a NI USB-6218 BNC with 200 kHz sampling rate) for analysis in software. The hardware-triggered DAQ allows online amplitude evaluation and rough analysis, which is used for detector stabilization via feedback, and valuable for detector optimization. The streaming DAQ on the other hand allows using an optimized trigger scheme that can be tuned after completion of the measurement, enabling full use of the detector performance at low energies. For these reasons, both DAQ systems were




Figure 8.6.: NUCLEUS prototype outer veto, consisting of two $5\text{ cm} \times 5\text{ cm}$ silicon cylinders. Left: lower outer veto showing the milled cavity surrounding the inner detector. Right: glued outer veto TES on a silicon carrier chip featuring a large W film, no phonon collectors and a direct thermal contact (gold bond wire).

used in the prototype measurements described below.

8.3. Measurements

Over several measurement campaigns, the various parts of the prototype have been operated in different combinations. While the full five detector channels have never been run simultaneously, all core features of the prototype have been demonstrated.

- Prototype Run 1: first operation of a $5\text{ mm Al}_2\text{O}_3$ cube as a target for NUCLEUS. Its main result is the demonstration of an ultra-low energy threshold below 20 eV .
- Prototype Run 2: a different NUCLEUS target operated inside the silicon holding structure, without instrumenting the inner veto. The measurement features a different TES structure, the best energy resolution so far and calibration and background datasets.
- Prototype Run 3: the first attempt at operating all five detectors together. The outer vetos were operated successfully, while a bad thermal contact prevented operation of the inner detector. After this partial success, the outer veto was again removed to focus on the inner detector.
- Prototype Run 4: first simultaneous operation of the target and two inner veto detectors.

8.3.1. Prototype Run 1: first target detector

This measurement performed in February 2017 represents the first operation of a NUCLEUS-type gram-scale cryogenic detector. The results obtained in this measurement lead to




Figure 8.7.: The first NUCLEUS prototype target cube assembled for cryogenic measurement. Left: sketch of the present prototype components. Green signifies an operational detector, grey a present but passive component, and white a component not present in the measurement. Right: photograph of the setup. a) cubical 5 mm Al_2O_3 calorimeter. b) clamp holding the detector via a 1 mm Al_2O_3 sphere, with glued Cu-kapton-Cu bondpad for electrical and thermal connections. c) contacts for heater and bias lines. d) 3 Al_2O_3 spheres glued on a copper plate to support the detector from below.

a publication on the detector performance [224], one paper on a new experimental approach to measure $\text{CE}\nu\text{NS}$ at a nuclear reactor [262], and new results on sub-GeV dark matter [263].

The prototype setup for this run consists of a single Al_2O_3 cube clamped on a copper plate via 4 Al_2O_3 spheres (diameter 1 mm, 3 below, one above) as shown in Figure 8.7. The electrical and thermal connections were made through wire-bonds routed on Cu-Kapton-Cu pads glued on the top clamp. The TES on the target was produced from the “small” layout. The detector was faced with an ^{55}Fe source with an activity of 0.6 Bq.

In the cryogenic test inside cryostat 1 at MPP, the TES showed a transition at a temperature of 22 mK. The mixing chamber of the cryostat was stabilized at 11 mK. The detector working point was optimized, resulting in a choice of $1 \mu\text{A}$ bias current and an appropriate detector heater setting to keep the TES at a sensitive yet stable point in its transition. A dataset of several hours was acquired with the hardware-triggered DAQ in parallel with stream data for software triggering. The hardware-triggered dataset is discussed below for a detailed determination of energy resolution and threshold, the driving parameter for the NUCLEUS target. The energy threshold thus determined can then be used to re-trigger the stream dataset. Exploiting the unprecedentedly low energy threshold, new limits on light dark matter models were derived from the software-triggered dataset, as discussed in the following section.




Figure 8.8.: Detector response to particle events. Left: Template pulse generated by averaging detector signals selected from a narrow amplitude range. The average of 40 hand-selected events (black) is overlaid with a parametric fit (blue) composed of an athermal (green dashed) and a thermal component (orange dotted). The inset zooms on the time around the event onset. Right: example pulses of different energies (0.02 keV, 0.16 keV, 0.54 keV, 1.35 keV, 4.88 keV, 8.60 keV) showing the linear range and saturation region of the detector.

8.3.1.1. Detector response and template pulse

The simplest method to determine pulse amplitude is to take the maximum of the trace and subtract an average of the pre-trigger window. This is good for automated quick analysis and can be displayed online while the measurement is ongoing. This method is susceptible to noise fluctuations and does not take into account pulse shape information. A better strategy is to fit an event trace with a scaled ideal detector response. The best-fit scaling factor then gives the amplitude.

The ideal detector response to a particle energy deposition in time-domain is given by the “standard event” or “template pulse”. It is obtained by averaging the traces of several dozen particle events with a similar amplitude and identical pulse onset. The template obtained in this way for the NUCLEUS target is shown in Figure 8.8 (left). The averaging serves to remove the noise, an identical pulse onset is important to avoid distorting the average pulse shape. To be sure that no artifacts enter the template pulse average, the selected events are vetted by manual inspection. As it can be challenging to collect enough events in a small amplitude range to completely suppress the noise, it is practical to fit the template pulse with a theoretical model and use this parametric fit instead of the actual template. The theoretical pulse shape is derived in [109] and discussed in section 2.3.1. The best-fit pulse model (see Figure 8.8, left) consists of two exponentials, with a common rise-time of $\tau_N = 0.30$ ms, a fast decay time of $\tau_f = 3.46$ ms and a slow decay time of $\tau_t = 28.17$ ms.

This “template fit method” works well as long as the pulse shape does not change as a function of deposited energy, i.e. in the linear range of the detector. The dynamic range

of a TES is limited, since the resistance change is proportional to temperature change only in a small range around the working point. As soon as the TES is heated above its transition, its resistance becomes constant and the detector saturates. The effect of detector nonlinearity and saturation is shown in Figure 8.8 (right). With the chosen working point and SQUID amplifier settings, saturation occurs at a pulse height around 1.3 V.

8.3.1.2. Calibration and truncated template fit




Figure 8.9.: Illustration of the truncated template fit method. Left: A 100 eV pulse and a 5.9 keV ^{55}Fe iron event overlaid with the best-fit pulse template. The fit is restricted to the linear range of the detector response. For clarity, the 100 eV event is displaced by 5 ms with respect to its trigger time. Right: fit to the Mn X-ray fluorescence lines of the ^{55}Fe source at 5.895 keV (K_α) and 6.490 keV (K_β). The fitted amplitude ratio between K_β and K_α is 1.097 ± 0.003 , in good agreement with the expectation of 1.101. This demonstrates the linearity of the truncated template fit method up to several keV in energy.

With the X-rays of the ^{55}Fe source fully saturating the detector, the template fit method cannot be used to find the calibration factor relating (linear range) pulse amplitudes to deposited energies. To extend the linear range of the detector to include the calibration lines, an extension called the “truncated template fit” is used. For this, the linear range of the detector has to be identified, the template fit then discards any data points beyond this amplitude and minimizes the rms between the template and the remaining data points. This is illustrated in Figure 8.9, where the left panel shows how the “extrapolated amplitude” of a saturated pulse can be found from fitting only the tail of the pulse. The truncation limit must be set in the linear range of the detector. The main frame of Figure 8.12 shows a plot of template fit rms residuals as a function of fitted amplitude. In the linear range of the detector the residuals are independent of the amplitude. As soon as nonlinearity sets in, the pulse-shape is distorted and the template fit residuals start rising. The truncation limit is set accordingly, in this case to

0.4 V.

With this method, the amplitudes of pulses well into saturation can be determined. This allows fitting the X-ray lines of the ^{55}Fe source, shown in Figure 8.9 (right). The detector resolves the K_α line at 5.895 keV and K_β at 6.490 keV. The resolution of the truncated template fit at this energy is found to be (135.3 ± 3.7) eV, the calibration factor allowing to convert pulse height and detector resolution between voltage and energy units is 1490.9 ± 1.1 eV/V. The linearity of the truncated fit up to this energy is confirmed by the amplitude ratio between K_β and K_α : the fitted value is 1.097 ± 0.003 , which compares well with the literature expectation of 1.101.

8.3.1.3. Optimum filter method and theoretical energy resolution




Figure 8.10.: Comparison of signal and noise power spectra from the first NUCLEUS target detector. The noise power spectrum is estimated from 283 empty baseline samples obtained from random triggers. The normalized signal power spectrum is given both as a discrete Fourier transform of the pulse template (scaled to an amplitude of 1 V) and as the continuous Fourier transform of the analytical pulse model. The analytical model can be decomposed into thermal and athermal components. Dotted vertical lines show the “poles” of the signal power spectrum, where the roll-off changes. Dashed vertical lines mark the frequency range contributing 90% weight to the detector resolution.

The template fit method is simple and useful as it can be extended to deal with detector non-linearity, which is necessary for calibration. Still, it does not take into account information on the noise power spectrum. There is a way to improve energy

resolution using the observed noise power spectrum, known as the optimum filter method (or sometimes, the optimum matched filter). The filter maximising signal-to-noise ratio in amplitude estimation given the ideal pulse shape and noise power spectrum has been textbook-knowledge for many decades [297, 298], but of limited practical use as it cannot be implemented in hardware (for a clear introduction see Appendix B in [299]). The optimum filter is acausal and needs information of the future signal to calculate the filter output at each point in time. This can be overcome approximately by introducing a large time-delay [300], or in modern implementations by filtering in software during post-processing. The filter can naturally be applied to hardware-triggered data, but this opens up an ambiguity as the trigger necessarily must use a different amplitude estimation scheme (with worse resolution) than the subsequent energy reconstruction. This is a problem for low-threshold applications, which are forced to either trigger large amounts of noise or waste some precious low-energy sensitivity. To avoid this issue, the complete data stream can be saved and triggering performed on optimally filtered data in post-processing, as is now the standard procedure in the CRESST experiment [243].

Here, the optimum filter is used on the hardware-triggered dataset to find the theoretically possible energy resolution (this section) as well as the achieved energy resolution and possible trigger threshold (next section), in preparation for the analysis of the stream dataset.

For an acquired trace (N samples spaced by a time interval Δt), the optimum filter output at time t_0 is given by:

$$A(t_0) = \frac{\sum_j e^{i\omega_j t_0} \cdot s_j^* \cdot v_j / |n_j|^2}{\sum_j |s_j|^2 / |n_j|^2} \quad (8.1)$$

where $j = -N/2 \dots N/2 - 1$ runs over the discrete Fourier components, $\omega_j = 2\pi f_j = 2\pi j/T$ is the associated frequency (using the trace length $T = N\Delta t$), and s_j , n_j , v_j are discrete Fourier transforms (DFTs) of the normalized ideal signal, the detector noise and the trace in question, respectively. The $*$ denotes complex conjugation. Discrete Fourier transforms g_j are calculated from the time series $\hat{g}(t)$ as¹:

$$g_j = 1/N \cdot \sum_k \hat{g}(t_k) \cdot e^{-i\omega_j t_k} \quad (8.2)$$

with $t_k = k\Delta t$. The denominator in Eqn. 8.1 normalizes the filter, as can be seen by setting $v_j = s_j$ and $t_0 = 0$. Intuitively, the optimum filter weights each Fourier component with its signal-to-noise ratio calculated from the expected pulse shape. For white noise $n_j = \text{const.}$, this reduces to a frequency-space formulation of the template fit method which correlates the observed trace with the expected signal.

The filter output (as well as all DFTs) are complex-valued (containing amplitude and phase information), but in practice the expectation value of the imaginary part of the filter output is zero and only the real part is used.

¹A word on units: in the definition above, the DFT has dimensions of amplitude per frequency channel. To arrive at the expected $\text{V}/\sqrt{\text{Hz}}$, one has to divide by $\sqrt{\Delta f} = 1/\sqrt{T}$. This is not necessary in the filter algorithm, as such rescaling cancels out. Doing so is useful for displaying the amplitude spectra independent from the sampling rate. As an additional source of confusion, $\hat{v}(t)$ and $\hat{n}(t)$ carry voltage units, while the normalized signal $\hat{s}(t)$ is dimensionless to give $A(t_0)$ the right units. Again, for displaying on a common axis, one can imagine normalizing $\hat{s}(t)$ to 1 V rather than 1, while for filtering a dimensionless $\hat{s}(t)$ makes more sense.




Figure 8.11.: The optimum filter applied to large and small signals. Left: a large signal is shaped into a time-symmetric localized peak by the action of the filter. Right: close to threshold, the optimum filter uses pulse shape information to distinguish small pulses from noise fluctuations. In the given example, the unfiltered signal crosses the threshold (dashed line at 13 mV) several times, while the filter output suppresses the noise and triggers only on the simulated 19.7 eV pulse.

The signal and noise amplitude spectra for the first NUCLEUS target are shown in Figure 8.10. s_j can be directly found from the discrete Fourier transform of the template pulse. To obtain n_j , the DFTs of many “empty baselines” collected from random triggers are averaged. Correctly selecting the empty baselines is a crucial point: it is necessary to remove random coincidences with particle events, as otherwise the filter would also suppress signal frequencies. On the other hand, a strict selection can also remove true high noise fluctuations, thus biasing the calculation of expected energy resolution and noise trigger rate. In general, for construction of filters any baseline selection is valid (not necessarily useful). For evaluation of filter and trigger performance, the only selection cuts allowed are those also made on the dataset independent from the trigger (called “quality cuts”). The empty baselines used here have been selected using a cut on the difference between last and first sample (“right-minus-left baseline cut”). Evidence of signal contamination is discussed below.

The expected variance of the optimum filter amplitude estimator is [299]:

$$\sigma_{\text{OF}}^2 = \left[\sum_j \frac{|s_j|^2}{|n_j|^2} \right]^{-1} \quad (8.3)$$

For the signal and noise amplitude spectra in Figure 8.10, a theoretical resolution of $\sigma_{\text{OF,theo.}} = 2.40 \pm 0.01$ mV is found. The uncertainty comes from the averaging of a finite number of empty baselines and was derived using case resampling (bootstrapping) of the individual baselines. Converting this to energy units, one obtains $\hat{\sigma}_{\text{OF,theo.}} = 3.58 \pm 0.02$ eV. This is the best possible energy resolution with the given signal pulse shape and noise power spectrum. The next section explores how close the real detector

approached this value. The formula for theoretical energy resolution allows to gauge the importance of different frequency bands to the detector performance. It was found that the central 90% weight of the sum in Eqn. 8.3 come from the frequency range $f_5 = 15.3$ Hz - $f_{95} = 903.3$ Hz. These frequencies are marked by vertical dashed lines in Figure 8.10.

The optimum filter formalism also provides an easy way to calculate the theoretically possible time resolution [299]:

$$\sigma_t^2 = \left[A^2 \cdot \sum_j \omega_j^2 \frac{|s_j|^2}{|n_j|^2} \right]^{-1} \quad (8.4)$$

Thus, the time resolution is inversely proportional to the signal amplitude A and high frequency signal components contribute more strongly. For the power spectra at hand, this expression evaluates to $\sigma_t = 1.08 \text{ ms} \cdot (A/\text{eV})^{-1}$. This is an important input for anticoincidence vetoing schemes, as discussed in section 6.3.5. The time resolution of the NUCLEUS target will be important for the successful suppression of muon-induced events [278].

Figure 8.11 shows the optimum filter applied to large and small example signals. Large pulses are shaped into time-symmetric localized peaks by the filter. Close to threshold, the filter reduces noise fluctuations and thus allows setting a lower threshold for a given accepted noise trigger rate.

8.3.1.4. Combination of energy reconstruction methods

Figure 8.12 compares the different amplitude reconstruction methods used in this analysis. A first pass with the template fit (orange data points) allows finding the truncation limit appropriate for the truncated template fit (TTF, blue data points). The TTF allows extracting the calibration factor (amplitude to energy conversion) which can be used by the optimum filter (OF) method at low energies. The inset in Figure 8.12 compares the energies reconstructed with these two methods event-by-event. Below the truncation limit, the reconstructed energies agree within uncertainty. At higher energies, the optimum filter underestimates pulse amplitudes due to saturation effects. The optimum filter thus provides the best energy resolution, but is limited to the linear range. The truncated template fit is valid also at higher energies, but has a worse energy resolution. For truncated pulses the resolution of the TTF degrades further, as fewer samples contribute to the fit and only partial pulse-shape information can be used.

8.3.1.5. Energy resolution from simulated events

Simulated events are used for the study of the trigger efficiency as described in the next section. Quite naturally, they also yield a measurement of energy resolution at low energy. A simulated event is shown in Figure 8.13 (left). To generate an artificial event with a known “true energy”, an appropriately scaled template pulse is added on an empty baseline sample. This is done for a set of empty baselines to obtain a representative sample of noise realisations. The artificial events can then be reconstructed with various techniques, and the noise-induced scatter around the known “injected amplitude” can be deduced. Figure 8.13 (right) shows histograms of 283 artificial events with an injected amplitude of 25 mV, reconstructed with the template fit and optimum filter techniques.

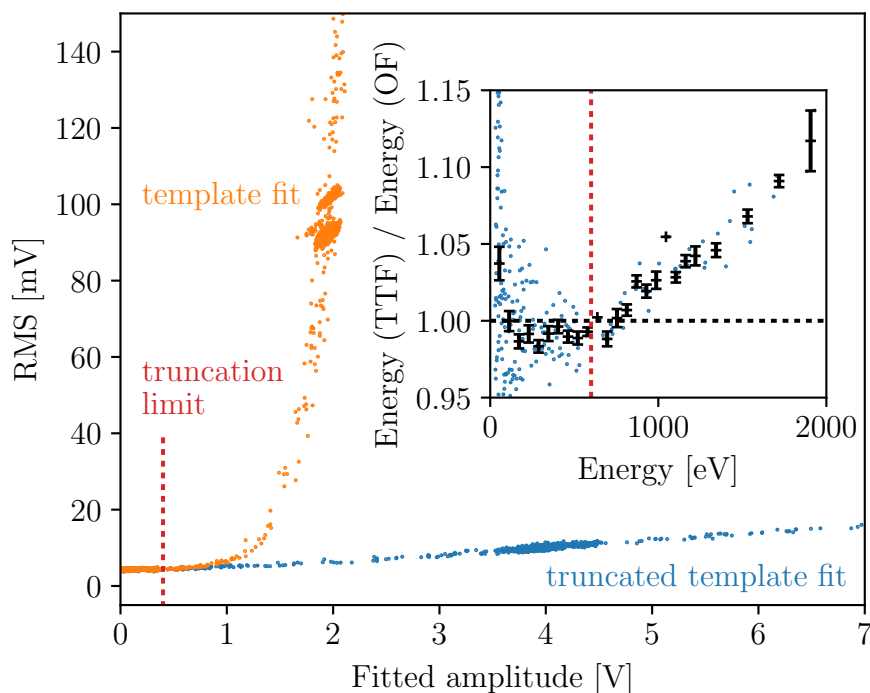


Figure 8.12.: Comparison of energy reconstruction methods. Main frame: fitted amplitudes and residual rms values are compared for the template fit and truncated template fit methods. The sharp rise in template fit rms above 1 V indicates saturation of the detector. The truncated template provides a better fit up to much higher amplitudes with only a moderate rise in rms. Inset: comparison of energy reconstruction by truncated template fit (TTF) and optimum filter (OF) methods. Below the truncation limit (indicated by the dashed red line), the methods agree on the percent level. Above, the onset of non-linearity causes the OF to underestimate pulse amplitudes.

The resolutions in energy units are 6.71 ± 0.28 eV (template fit) and 3.84 ± 0.16 eV (optimum filter). The performance of the optimum filter method measured with simulated events is within 10% (or 2σ) of the theoretical value $\sigma_{\text{th}} = 3.58$ eV calculated from pulse shape and noise power spectrum, and improved by a factor 1.8 with respect to the template fit. The small deviation from the theoretical filter energy resolution can be caused by low statistics of the simulated events, or violated assumptions like noise stationarity.

8.3.1.6. Threshold from trigger efficiency simulation

The low value of the energy resolution observed in Run 1 allows setting a correspondingly low trigger threshold without suffering from a high noise trigger rate. The threshold of the NUCLEUS target prototype was set to 13 mV after on-line study of the behaviour of the trigger rate. This choice is justified below. To characterise detector response at low energies, the “trigger-efficiency curve” is calculated. It gives the probability of an

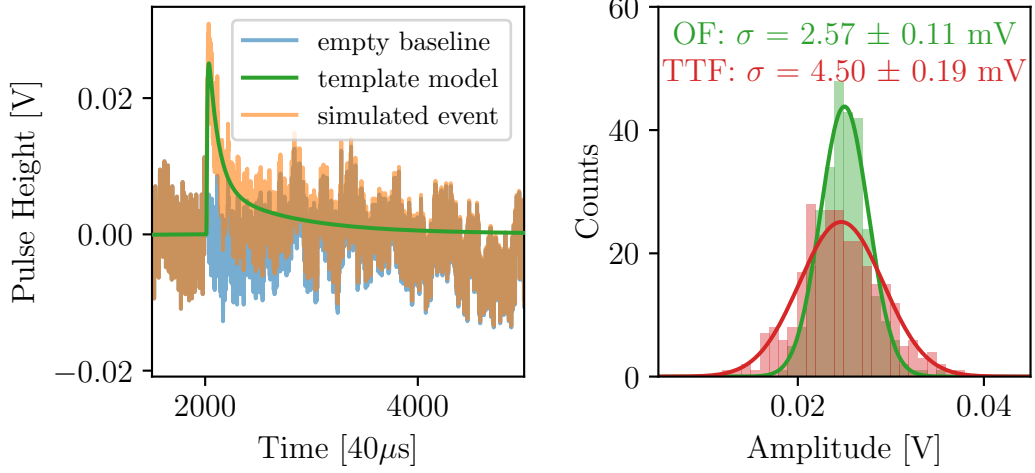


Figure 8.13.: Simulated events for determining energy resolution. Left: an artificial event with known “true energy” is composed of an empty baseline and a template pulse scaled to the appropriate amplitude. Right: an ensemble of simulated events from a large number of empty baselines can be used to determine the resolution of different amplitude reconstruction methods. A Gaussian fit shows the resolution of the truncated template fit (TTF) and optimum filter (OF) methods.

event triggering the detector as a function of its true energy. The trigger efficiency as a function of reconstructed energy is a step function. The finite energy resolution of the detector smears the trigger efficiency curve into an error function (the convolution of step function and Gaussian resolution function). As there is a non-zero probability of triggering noise (i.e. an event with zero true energy), the error function has a positive offset, and can be described by three parameters (noise trigger probability p_0 , threshold μ , energy resolution σ):

$$P_{\text{trig}}(E) = 1 - (1 - p_0) \cdot \frac{1 - \Phi(E, \mu, \sigma)}{1 - \Phi(0, \mu, \sigma)} \quad (8.5)$$

or

$$P_{\text{trig}}(E) = \frac{p_0 + (1 - p_0) \cdot \Phi(E, \mu, \sigma) - \Phi(0, \mu, \sigma)}{1 - \Phi(0, \mu, \sigma)} \quad (8.6)$$

with the Gaussian cumulative function:

$$\Phi(x, \mu, \sigma) = \frac{1}{2} \cdot \left[1 + \text{Erf} \left(\frac{(x - \mu)}{\sqrt{2}\sigma} \right) \right]. \quad (8.7)$$

The trigger efficiency curve essentially is an appropriately shifted Gaussian cumulative. The offsets and scaling factors ensure $P_{\text{trig}}(0) = p_0$ and $\lim_{E \rightarrow \infty} P_{\text{trig}}(E) = 1$.

In practice, the trigger efficiency curve is generated by counting the fraction of simulated events with a reconstructed amplitude above threshold, as a function of injected energy. Figure 8.14 shows a trigger efficiency curve generated by successively superimposing the same 283 empty baseline samples with differently scaled standard events. Overlaid is the best fit with the model of Eqn. 8.5. The energy resolution derived from




Figure 8.14.: Trigger efficiency curve generated from simulated events at many different energies, reconstructed using the optimum filter method. A trigger setting of 13 mV leads to an energy threshold of 19.54 ± 0.07 eV. The width of the best-fit error function $\sigma_{th} = 3.63 \pm 0.08$ eV agrees with the energy resolution from simulated events. A small noise trigger probability p_0 can be seen as a pedestal in the efficiency curve.

the trigger efficiency curve is 3.63 ± 0.08 eV, in agreement with the value found directly from the simulated events. The fitted threshold (mid-point of the error function) is 19.54 ± 0.07 eV, close to the expected value of 19.38 ± 0.01 eV. With a noise trigger probability of 1.05 %, one would estimate a noise trigger rate of $p_0 / (8192 \cdot 40 \mu\text{s}) = 3.2 \cdot 10^{-2} \text{ s}^{-1}$ or 115 noise triggers per hour. A closer look is warranted, since the uncertainty is large and does not take into account that the empty baseline ensembles to which the efficiency curve is fitted are not statistically independent.

8.3.1.7. Estimate of noise trigger spectrum

Noise trigger rates can be calculated in a simple statistical model, as described in [301]. The samples in a trace are understood to be randomly distributed according to a normal distribution with a standard deviation given by the energy resolution. The probability of a random fluctuation above threshold can be calculated with the methods of order statistics. One needs to find the probability distribution of the highest sample, known as the “distribution of the extreme” in statistics literature [302]. It can be quickly derived by considering its cumulative distribution function. For the largest sample to be smaller than a given value, all N individual samples have to be smaller than that value:

$$CDF_{\max}(x) = CDF(x)^N. \quad (8.8)$$

The probability distribution function is found by differentiating once:

$$P_{\max}(x) = N \cdot P(x) \cdot CDF(x)^{N-1}. \quad (8.9)$$

This pdf is plotted for a Gaussian $P(x)$ and various values of N in Figure 8.15. The samples in a detector trace are not statistically independent, as the noise power spectrum




Figure 8.15.: Probability distributions for the highest of a set of N samples from a standard normal distribution. The values of N are chosen such that the expectation value for the highest sample is an integer.

is typically not white. This can be captured by introducing the “effective number of independent samples” $N_{\text{eff}} < N$, which must be determined concurrently. The strategy in [301] is to take a distribution of maxima from a large sample of empty baselines, and simultaneously fit for σ and N_{eff} . This has been successfully applied to CRESST-data [228].

The maxima of 283 filtered empty baseline samples are shown as the blue histogram in Figure 8.16, along with the best-fit model of Eqn. 8.9. The fitted energy resolution is much larger than the values from trigger efficiency or simulated events. Repeating the analysis with all baselines multiplied by -1 (orange histogram and fit), the picture changes significantly. The histogram is much narrower and centered at lower values. The energy resolution is now slightly underestimated. This behaviour shows that the baseline samples are distributed asymmetrically. An explanation could be the presence of undetected (sub-threshold) signal pulses. While noise is expected to appear symmetrically, detector pulses have a fixed polarity, and are converted to upward spikes by the filter. Sub-threshold pulses therefore bias the maximum-distribution upwards. After inverting, the maximum search is only affected if the would-be maximum on the trace is aligned with the pulse. The pulse then removes the actual maximum and leads to the second-highest sample being incorrectly identified as the maximum. Therefore the maximum distribution from the inverted baselines is expected to be biased downwards, but less severely than the untreated baselines. This qualitatively matches the observed behaviour of the empty baseline sample in Figure 8.16.

Lacking a cleaner empty baseline sample, one can use these two cases to find upper and lower bounds on the noise trigger rate. The noise trigger rate in the statistical model is given by

$$NTR(E_{\text{th}}) = (1 - CDF_{\text{max}}(E_{\text{th}})) / (m \cdot T) \quad (8.10)$$

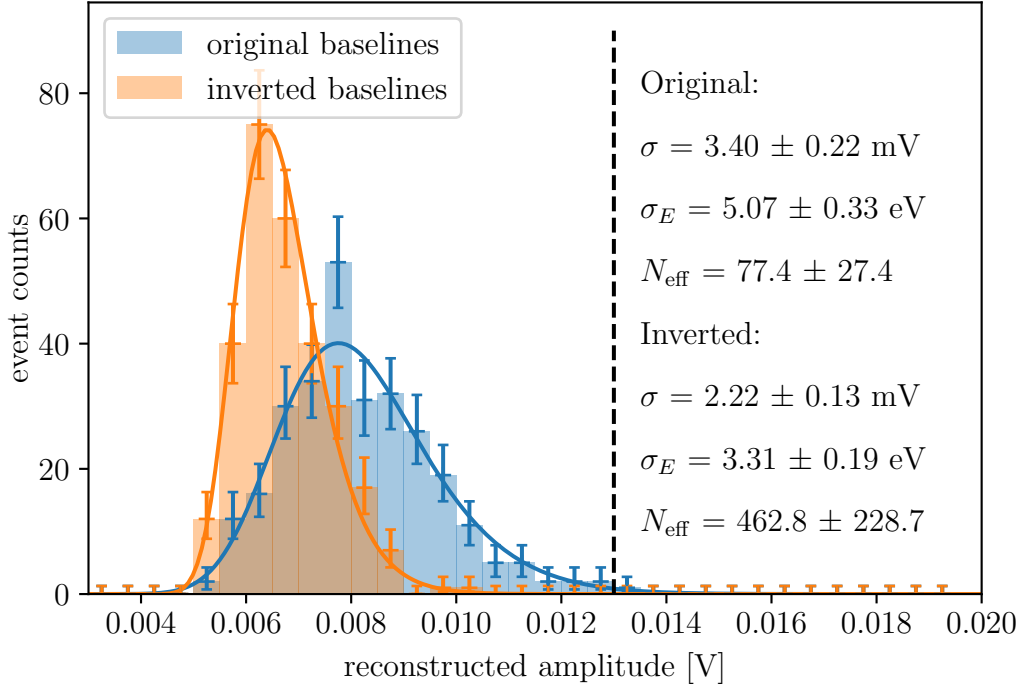


Figure 8.16.: Distributions of maxima of 283 filtered empty baseline samples. Blue: histogram of maxima with best fit according to the model of Eqn. 8.9. Orange: same 283 baselines multiplied by -1 before filtering. The best fit energy resolution matches the value found from simulated events better for the inverted baselines, suggesting the presence of sub-threshold events in the “empty” baselines.

with window length T and detector mass m normalizing the expected noise triggers (units counts / kg day) to exposure. This function, with the best-fit parameter values for original and inverted baselines is plotted in Figure 8.17 (left). The uncertainty range on the noise trigger rate is large, since one is estimating the tail area of a distribution far away from the peak. For the original baselines, one finds a noise trigger rate of $2.7 \cdot 10^6$ (kg · day) $^{-1}$, with the error band ranging between $0.9 \cdot 10^6$ (kg · day) $^{-1}$ and $6.9 \cdot 10^6$ (kg · day) $^{-1}$. For a 5.3 h measurement, this translates to between 98 and 742 noise triggers. For the inverted baselines, one finds a best-fit value of $5.9 \cdot 10^2$ (kg · day) $^{-1}$ (range from $0.7 \cdot 10^2$ (kg · day) $^{-1}$ to $38.3 \cdot 10^2$ (kg · day) $^{-1}$), or between 0.007 and 0.41 noise triggers in 5.3 h.

This huge range shows that the “empty baseline maximum distribution” cannot precisely constrain the expected noise trigger rate in the given conditions. This is due to the statistics and quality of empty baseline samples, and therefore does not affect the applicability of the method to larger datasets acquired in underground conditions.

In addition to the noise trigger rate, also the noise trigger energy distribution can be predicted, it is given by Eqn. 8.9. Figure 8.17 (right) shows the expected noise trigger energy distribution (cut at threshold) for the same best-fit parameters. (The data from the hardware-triggered DAQ stops at a higher threshold around 30 eV.) As the differential noise trigger rate contains a factor of $P_{\text{max}}(x)$, it is expected to fall as a

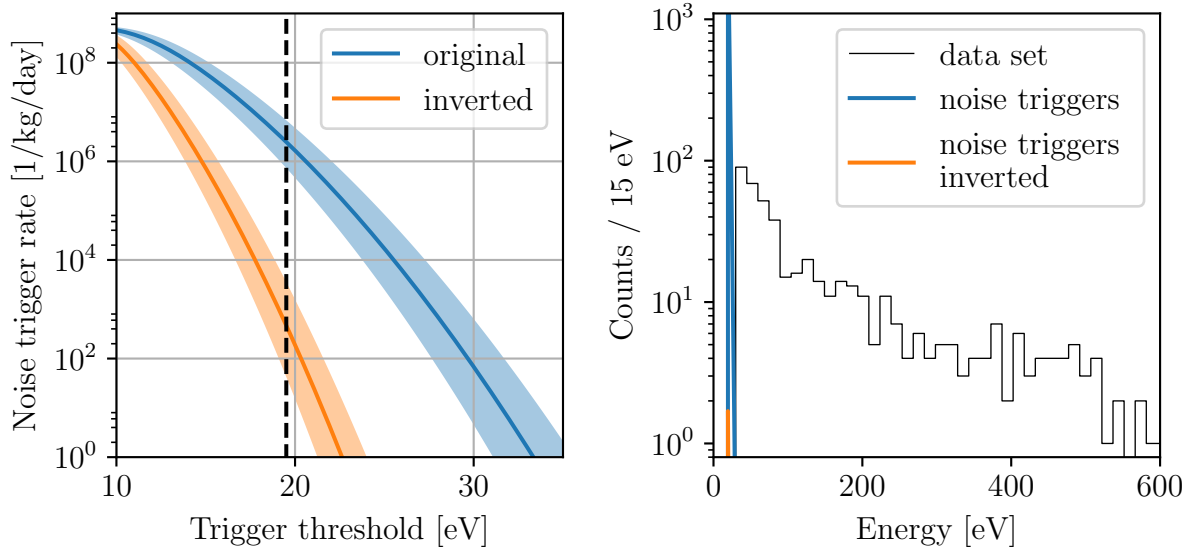


Figure 8.17.: Noise trigger distributions extrapolated from the empty baseline sample. Left: noise trigger rate as a function of threshold, for the original (blue) and inverted (orange) baseline sample. Right: expected energy distribution of noise triggers, compared to the observed spectrum using the higher-threshold hardware trigger.

normal distribution having a width given by the energy resolution. Therefore, although a significant number of noise triggers may be present, their rate above the lowest energy bin is expected to be negligible.

8.3.1.8. Stream dataset: Dark Matter Analysis

The low threshold of the first prototype NUCLEUS target implies an unprecedented sensitivity to low-energy nuclear recoils. For this reason, the dataset from Prototype Run 1 was used to set new limits on dark matter-nucleus interactions (published in [263]). With the short duration and high background, only very basic data selection was applied. Figure 8.18 shows the stability cut: unstable periods were marked by hand and removed from the dataset. This reduced the live-time of the detector from 5.31 h to 3.26 h. Two cuts on simple pulse parameters (shown in Figure 8.19) are used to remove artifacts: requiring a decay time faster than 400 samples removes test pulses and misreconstructed saturated pulses. Selecting for a difference of at least -1.8 V between the first and last sample in a trace (right-left baseline) removes SQUID resets whose amplitude is not correctly reconstructed. As the population of signal events does not extend into the cut-regions, a survival probability of 1 is assumed. Finally, the region-of-interest (ROI) is restricted to the linear energy range of 19.7 eV-600 eV. All events outside the ROI are counted as detector dead-time, reducing the live time from 3.26 h to 2.27 h. The final exposure of the measurement is thus 0.046 kg day. Figure 8.20 shows the energy spectrum of the selected events. The inset zooms on the 511 events in the ROI with a finer binning.

The ROI events, with the exposure of 0.046 kg day and the signal survival probability




Figure 8.18.: Stability cut for Prototype Run 1, showing Optimal Filter energy as a function of time. The grey regions were manually selected and removed from consideration. The dashed black line shows the truncation limit (upper end of the detector’s linear range).

of Figure 8.14 is converted into a limit on dark matter-nucleus interactions. The low-energy background is conservatively assumed a potential dark matter signal, so that only larger-than-observed count rates are excluded. The limit is derived using the Yellin Optimum Interval Method [95]. Figure 8.21 shows the limit curve, calculated for dark matter masses between $140 \text{ MeV}/c^2$ and $10 \text{ GeV}/c^2$. In Figure 1.9, the same limit appears as the dotted red line (CRESST surface 2017). At the time of publication in [263], unexplored parameter space below $0.5 \text{ GeV}/c^2$ could be constrained for the first time. As the detector was operated with negligible overburden, these constraints are valid also for dark matter interacting strongly enough to be significantly attenuated in the Earth’s crust, and thus invisible to underground searches. This interpretation is explored in [303], where limits are calculated for Strongly Interacting Massive Particle (SIMP) dark matter.




Figure 8.19.: Data quality cuts for Prototype Run 1. Left: decay time cut, removing test pulses and misreconstructed saturated pulses. Right: right-left baseline cut, removing SQUID resets and pile-up events.




Figure 8.20.: Final energy spectrum reconstructed in Prototype Run 1. Main frame: complete energy range up to the ^{55}Fe calibration lines, reconstructed with the truncated template fit. Inset: zoom on the region-of-interest (19.7-600 eV), reconstructed with the optimum filter.




Figure 8.21.: Limit on dark matter-nucleus scattering derived from the 511 events observed in the ROI using the Yellin Optimum Interval Method [95]. The sensitivity provided by the lighter oxygen nuclei dominates below ~ 1 GeV/ c^2 . At higher masses, the contribution from the larger aluminum nuclei is more important. In Figure 1.9, this limit is shown in the context of results from other searches.

8.3.2. Prototype Run 2: background measurement in silicon holder




Figure 8.22.: Detector setup for the first test of the silicon detector holder. Left: sketch of the present prototype components. Right: photograph of the setup. The NUCLEUS target is equipped with an LD-TES and sandwiched between inner veto plates (not instrumented in Run 2). On the left, the detachable iron source is visible.

The second NUCLEUS prototype target detector was operated from January 18-29, 2018 in cryostat 2 at MPP Munich. It features the first LD-TES (see Figure 8.3) on a Al_2O_3 target and was operated in the silicon inner detector holder including inner veto plates. These were not equipped with TES in this first run. In addition, a removable ^{55}Fe source was mounted to the detector holder. In addition to testing the holding scheme mechanically and thermally, this allows to assess whether the very-low-energy background is source-related.

Figure 8.22 shows a photograph of the detector setup. To the left of the silicon structure carrying the detector, an arrangement of copper wires holds the source (which is horizontally facing the detector). A nylon string secures the source via a latch, and runs up to a bellows on the room-temperature side of the cryostat vacuum chamber. Extending the bellows pulls the string, releases the latch and causes the source to drop away from the detector (into the mixing chamber shield).

Detector optimization resulted in a bias current of $9.2 \mu\text{A}$, much larger than the $1.0 \mu\text{A}$ used in Run 1. The observed time constants are $\tau_n = 71 \mu\text{s}$, $\tau_{in} = 1.09 \text{ ms}$ and $\tau_t = 6.87 \text{ ms}$. τ_n is much faster compared to Run 1, as expected due to the increased TES area. τ_{in} and τ_t can be reduced by electrothermal feedback due to the much higher bias heating.

Figure 8.23 shows test and particle pulse amplitudes recorded as a function of time, and gives an overview of the measurement. The first 15 h are the calibration dataset




Figure 8.23.: Pulse heights as a function of time for the calibration and background datasets of Prototype Run 2. After ~ 15 h, the calibration source was released using the pull-string mechanism, leading to a significant disturbance of the cryostat temperature. Data taking was restarted once the detector returned to its operation point. The calibration lines around 2 V amplitude are absent in the second (background) dataset, while the continuously injected test pulses ensure a fixed energy scale between the two datasets.

recorded with the final detector settings. Around 2 V amplitude, the K_α and K_β lines from ^{55}Fe are visible. Test pulses start at low energies and extend beyond the energy of the calibration lines. The detector fully saturates around 1 V amplitude, and the truncation limit was set to 0.4 V.

After recording the calibration dataset, the ^{55}Fe source was removed by pulling the string. The mixing chamber warmed up to ~ 40 mK and took one hour to return to base temperature. After the detector returned to its operation point, recording of the background dataset (right in Figure 8.23) was started. The iron lines are absent as expected, but the continuing testpulses ensure a consistent energy scale between the two measurements.

The energy scale is found as in Run 1, determining the calibration factor from the amplitude of the Mn K_α line at 5.895 keV. The fit to the iron lines, shown in Figure 8.24 (right), gives a value of 3.135 ± 0.001 keV/V.

The optimum filter resolution is then evaluated on simulated events created from empty baseline samples. The distribution of optimum filter energies reconstructed for the simulated events is shown in Figure 8.24 (left). The optimum filter resolution is:

$$\sigma_{OF} = 0.93 \pm 0.03 \text{ mV} = 2.92 \pm 0.11 \text{ eV}.$$

This is the best baseline energy resolution obtained with a massive cryogenic calorime-




Figure 8.24.: Determination of the energy resolution for Run 2. Left: empty baseline samples give an optimum filter resolution of 0.93 ± 0.03 mV. Right: the Mn K_α line, reconstructed at 1800.5 ± 0.6 mV yields a calibration factor of 3.135 ± 0.001 keV/V. The baseline energy resolution is then 2.92 ± 0.11 eV. The energy resolution at 6 keV is 58.3 ± 1.5 eV.

ter to date and allows triggering at a threshold of 14.5 eV with a similar expected noise trigger rate as for the higher threshold in Run 1.

Figure 8.25 gives an overview of the event populations observed in Run 2. Optimum filter amplitude is plotted against decay time, defined as the number of samples after the pulse maximum at which the amplitude falls below $1/e$ of the maximum. Particle events in the target appear as the band from a)-f) (for example pulses see Figures 8.26 and 8.27). Test pulses are the populations with discrete energies and slower decay times b). The populations around d) and e) are artefacts. “rectangle events” d) appear periodically and are thought to originate in the TES operation point hopping between two stable points near a micro-feature in the transition curve. Pulse pile-ups e) lead to incorrect reconstruction of amplitude and decay time.

A new population c) is observed in Run 2, with a significantly slower decay time than particle or test pulses. An example event is shown in Figure 8.27. This population is attributed to particle energy depositions in the veto plates, which are touching the target. Thermal cross-talk, or phonon transmission, causes the target TES to respond also to these events. The strength of the cross-talk, i.e. the relative amount of energy transmitted from the veto into the target, must be assessed with the veto instrumented, and will be studied in section 8.3.4. Even without the instrumented veto, the two populations a) and c) can be distinguished due to their different pulse shape.

To compare the calibration and background datasets and assess the low energy spectrum, identical pulse-shape cuts are applied. Figure 8.28 shows the energy from 40 eV to 8 keV, reconstructed with the truncated template fit. Target energy depositions are




Figure 8.25.: Event populations in the amplitude-decay time plane, for the background dataset. The decay time parameter allows to distinguish particle events in the target (a, f) from other types of events (b, c, d, e).

selected using a cut on the template fit RMS. The rate in the calibration dataset is higher than in the background dataset, not only at the energies of the X-ray fluorescence lines, but also at sub-keV energies and in a region around 5 keV. This shows that some of the low-energy background in Run 2 is source-related, for example from secondary electrons not stopped efficiently in the mylar foil covering the ^{55}Fe source.

Figure 8.29 shows the comparison of the Run 2 low-energy spectra (triggered with the optimum filter at 14.5 eV, selected for target events by the decay-time cut shown in Figure 8.25) with the spectrum obtained in Run 1. The Run 2 background spectrum at low energies is comparable to the Run 1 spectrum in the presence of a source. This demonstrates that, while ^{55}Fe sources can cause sub-keV backgrounds, the main contribution to the low-energy rate is not source-related. In addition, the low-energy rate was found comparable between the two cryostats at MPP Munich.




Figure 8.26.: Low-energy pulses from populations marked in Figure 8.25. a) target particle event. b) test pulse. c) slow event attributed to thermal cross-talk after an energy deposition in the silicon holder plates.




Figure 8.27.: Artefacts and high-energy pulses from populations marked in Figure 8.25. d) “rectangle event” attributed to a local feature in the transition. e) pile-up of saturated events. f) saturated target particle event.




Figure 8.28.: Full energy spectra (40 eV - 8 keV) reconstructed with the truncated template fit, for calibration and background datasets. Identical cuts on the fit residuals were used to select events in the target.




Figure 8.29.: Low-energy spectra reconstructed with the optimum filter, from the region of interest outlined in dashed black in Figure 8.25. The linear energy range of the detector extends up to 1.25 keV. For comparison, the Run 1 spectrum up to 600 eV is drawn as well.

8.3.3. Prototype Run 3: first outer veto operation

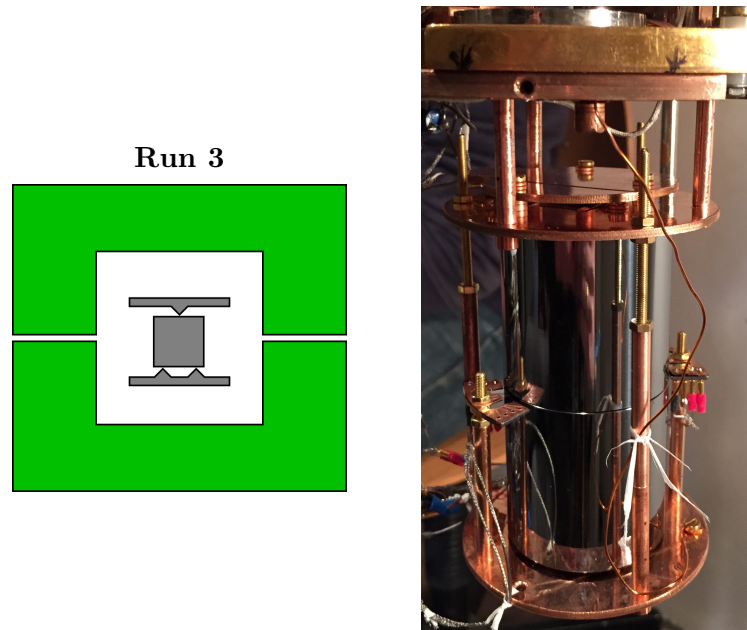


Figure 8.30.: The full NUCLEUS-1g prototype mounted in cryostat 1 at MPP in March 2018. Left: sketch of the present prototype components. Right: photograph of the setup. Between the outer veto detectors, the support wafer holding the inner detectors is visible.

In March 2018, the full NUCLEUS-1g prototype was assembled for the first time and cooled down in cryostat 1 at MPP. Figure 8.30 shows the assembled detectors mounted below the mixing chamber plate of the cryostat. Besides the silicon outer veto cylinders, the support wafer independently holding the inner detectors can be seen. During the measurement, the outer veto detectors could be operated, but no transitions were observed on any of the three “inner detector” channels. Post-run it was found that the brass screws holding the inner detector had been set too loosely and likely mechanical vibrations prevented cooldown of the inner detector. All electrical connections survived the thermal cycling without problems.

The outer veto detectors were operated in coincidence for several hours. The 5 cm silicon cylinders show a high rate of particle events (tens of events per second) in the unshielded setup. Figure 8.31 shows an example trace of the two outer vetos including a large energy deposit saturating both TES and many smaller events preventing relaxation to a defined operation point.

Despite the high rate and unstable conditions, enough clean pulses could be selected to create meaningful template pulses. Figure 8.32 shows the templates and corresponding fits with the bolometric pulse-shape model (see Equation 2.11 and Figure 2.6). The fit parameters summarized in Table 8.2 show well-separated time scales ($\tau_{in} \ll \tau_n \ll \tau_t$) and a dominant athermal component, indicating a TES design well matched to the absorber crystals.

An energy calibration was not attempted in the high-rate environment at MPP. For its later use as anticoincidence veto in NUCLEUS, the time resolution of the outer veto is




Figure 8.31.: Example trace from the outer veto detectors operated without any shielding. Particle events at a rate above 10 Hz prevent stable operation.

outer veto	top	bottom
A_n/A	0.8649 ± 0.0005	0.7773 ± 0.0006
τ_f	1.316 ± 0.005 ms	0.957 ± 0.005 ms
τ_n	13.70 ± 0.02 ms	9.34 ± 0.04 ms
τ_t	71.2 ± 0.3 ms	60.3 ± 0.2 ms

Table 8.2.: Parameters of the thermal model fit to the pulse templates shown in Figure 8.32.

an important parameter. As a first study of this property, a search for coincident events in the two outer veto detectors was performed. Figure 8.33 shows a histogram of time differences between events triggered in the top and bottom outer vetos. The distribution is modeled as a flat component (representing uncorrelated events) and a Gaussian centered at zero (for events triggering both detectors with a finite time resolution). The best-fit time resolution determined in this way is 1.26 ± 0.10 ms. Given the \sim ms rise time of the outer veto response, this number is expected to improve with a more advanced analysis.

With the operation of the outer veto detectors demonstrated, they were again removed from the test assembly. In this way, more rapid iterations were possible to bring the inner detector to full operation.

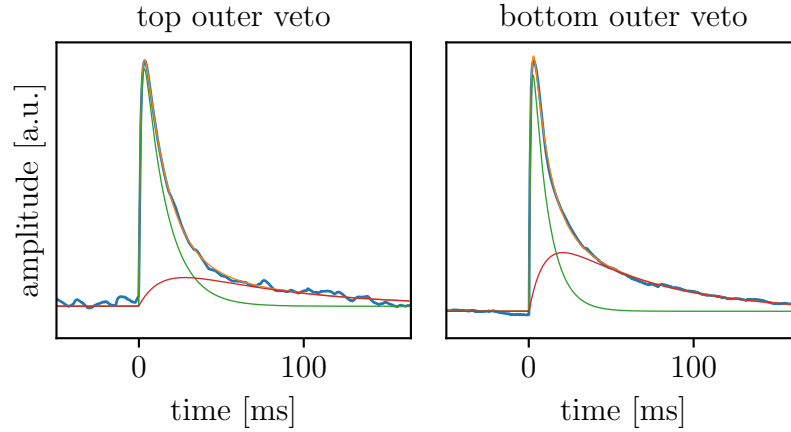


Figure 8.32.: Template pulses (blue) derived for top (left) and bottom (right) outer veto. The fit with the thermal model of equation 2.11 is shown in orange, with athermal (green) and thermal (red) components. The templates show the remnants of numerous pile-up events. Still, the time scales of the pulse template can be determined with confidence.




Figure 8.33.: Coincidence search between top and bottom outer veto. The histogram shows the distribution between time differences of all triggered events. The fitted model consists of a constant part and a Gaussian at zero time difference. The width of the Gaussian component gives the time resolution.

8.3.4. Prototype Run 4: target and active inner veto




Figure 8.34.: Detector configuration of Prototype Run 4. Left: sketch of the present prototype components. Right: photograph of the setup. Compared to Run 2, new silicon screws and electrical connections also to the top inner veto are visible. On the bottom side of the holder similar connections to the bottom inner veto are placed.

In March 2019 the inner detector of the NUCLEUS prototype (target and two inner veto detectors) was fully operated for the first time (in cryostat 1 at MPP). The measurement was performed in compensated magnetic field conditions, as described in section 4.2.3. The detector holder is pictured in Figure 8.34. It looks similar to the setup of Run 2, except for the silicon screws now used on three of the posts securing the veto stack, and the electrical contacts to the veto that are visible through the cut-out in the holding wafer. Similar connections to the bottom veto are on the bottom side of the holder. The Al_2O_3 target cube was instrumented with an M-TES structure. The detector assembly is again equipped with the removable source, but its mechanism was not used in this run.

The transition temperatures were measured to be (16, 20, 18) mK for target, top and bottom inner veto respectively. The three detectors are in mechanical contact (as required for the instrumented holder concept). In consequence, each detector responds also to a heating applied to the others. The optimization can therefore not be done independently and upper limits exist for heater settings (to avoid overheating other TES channels). Several configurations of bias currents and heater settings were studied before a final setting was found: (2.0, 1.0, 3.3) μA on target, top and bottom inner vetos. As all detectors were illuminated by the ^{55}Fe source, an independent energy calibrations can be performed. The baseline resolutions found for the final configuration are (9.6, 14.7, 8.3) eV. The inner vetos thus easily reach their targeted 100 eV thresholds. The performance of the target was not comparable with previous runs. This is possibly related to the optimization constraints imposed by operation of detectors in mechanical contact. A proposal to mitigate the effects of cross-heating is discussed in section 8.3.5.

In the same way as the steady heating applied to the detectors, also the phonons created by particle events are partially transmitted between the detectors. Due to the different pulse shapes for direct hits versus cross-talk events it is not straightforward to quantify the amounts of phonon transmission. The phenomenon can be understood by looking at truncated fit energy spectra reconstructed for the different detectors. Fig-




Figure 8.35.: Simultaneous response of the three detector channels to a ^{55}Fe event depositing 6 keV in the target. The target trace is multiplied by 10 as it uses a SQUID electronics with weaker amplification. The target pulse is strongly saturated (saturation limit ~ 2 keV) while the veto responses induced by phonon transmission are small.

Figure 8.36 shows events observed in a 9.1 h calibration dataset, by the top and bottom inner veto detector channels. Three populations immediately stand out, corresponding to direct hits in the three detectors. The iron calibration lines can be identified as direct hits in each inner veto. Also the iron direct hits in the target can be identified using only the cross-talk amplitudes induced in the vetos. The appearance of the target iron events near (1 keV, 0.5 keV) in Figure 8.36 does not imply a leakage of 1/6 and 1/12 of the deposited energy: the veto pulse shapes are significantly different for target and veto hits, and the energy scale derived for a direct veto hit cannot be transported to target hits. No significant cross-talk is observed between the two veto plates up to 8 keV deposited energy, consistent with the fact that there is no direct mechanical contact between them allowing phonon transmission.

Figure 8.37 shows the same events reconstructed by the target and top veto detectors. The direct hit populations are significantly tilted, as there is direct contact between the two plotted detectors. When reconstructed in this way, the cross-talk appears asymmetric: for an iron deposition in the target, the top veto reconstructs a higher amplitude than vice-versa. This can be explained from the different volumes of target and veto. For a given transmitted phonon fraction, different phonon densities and thus temperature rises result. The bottom veto direct iron hits appear around 100 eV in the target and cannot be clearly separated from target and top veto amplitude alone. Events strongly saturating the bottom veto can be identified due to their cross-talk amplitudes in target and top veto.

To turn the inner veto detectors into an anticoincidence veto, a cut on the inner veto amplitude is developed. To accurately judge the effect of the cut also at low energies,




Figure 8.36.: Event populations from a 9.1 h calibration dataset, as observed by the two inner veto detectors. Three populations corresponding to particle events in the three detectors can be distinguished.

simulated events following the observed target hit population are created. Figure 8.38 shows the model used to create the simulated target events. The veto yield, defined as the relative amplitude in veto and target, is plotted against target energy for both vetos. The target events follow a curved band which is well-described by a fourth-order polynomial. At low energies, the veto yields are 0.132 and 0.044 for top and bottom inner veto respectively. Artificial events of various energies are created by superimposing template pulses scaled to the amplitudes required by the polynomial model to empty baseline samples of the detectors. The artificial events thus scatter around the polynomials in Figure 8.38 by construction. Figure 8.39 shows the cut efficiency derived by passing the artificial events through cuts on top veto yield (acceptance window $-0.1 \dots 0.2$) and bottom veto yield (acceptance window $0.0 \dots 0.2$). The cut efficiency determined with 544 artificial events reaches 100% above 1 keV and falls below 50% at ~ 150 eV.

Figure 8.40 shows a comparison of target energy spectra before and after applying the veto cut. Both spectra are first cleaned from artifacts by a right-left baseline cut. The post-cut rate shown in orange is corrected dividing by the energy-dependent survival probability shown in Figure 8.39. Although the veto cut removes events from the spectrum at energies as high as the iron lines, the effect becomes most dramatic below 2 keV in the target. Where the before-cut spectrum rises sharply due to numerous cross-talking veto events, the post-cut spectrum stays flat down to ~ 300 eV. Figure 8.41 shows a zoom of the same spectra from 60-600 eV. Features at ~ 100 eV and ~ 200 eV due to ^{55}Fe energy depositions in the vetos are removed by the cut, leaving a sharp rise in the rate from 100 eV towards threshold.




Figure 8.37.: Event populations from a 9.1 h calibration dataset, as observed by the target and top inner veto detectors.




Figure 8.38.: “Veto yield” for top and bottom inner veto, as a function of energy in the target. The veto yield is defined as the truncated fit amplitude ratio between veto and target detectors. The red data points represent the median of the event yields (blue) for equidistant energy bins. The orange line is a fit by a polynomial of 4th degree. This line is used as a model to generate artificial events with a realistic cross-talk strength, to determine the survival probability of the veto cut for target events.




Figure 8.39.: Survival probability for the veto cut defined in the text, derived using simulated events following the veto yield model shown in Figure 8.38.




Figure 8.40.: High-energy spectra of Run 4 target events, before and after applying the inner veto cut. The post-cut spectrum is corrected for the cut survival probability shown in Figure 8.39. Below ~ 2 keV, the veto cut removes a significant number of cross-talk events.




Figure 8.41.: Low-energy spectra of Run 4 target events, before and after applying the inner veto cut. In the energy range from 60 eV to 600 eV, the veto cut removes the direct hits of the ^{55}Fe source in the vetos from the target spectrum. The remaining spectrum is flat above ~ 300 eV. Below ~ 100 eV, a sharp rise of surviving target events remains.

8.3.5. Conclusions from Prototype Runs




Figure 8.42.: New TES structures designed for NUCLEUS target and inner veto detectors after the Prototype Runs. Both TES have larger bond pads compared to previous designs, which simplifies repeated contacting at the cost of some dead area. The inner veto TES features an ohmic heater directly on the TES film, electrically insulated by a silicon oxide layer (drawn in orange). This design is adapted from CRESST iStick TES and allows heating the TES more directly with a reduced impact on the crystal temperature.

In the measurements presented here, all components of the NUCLEUS-1g prototype were operated successfully. Run 1 allowed placing new leading limits on interactions of low-mass dark matter. In Run 2, the silicon holder was successfully introduced and a new best energy resolution of 2.9 eV was demonstrated. The cryogenic, mechanical and electrical design of the prototype was verified in Run 3 with the assembly of all components. The outer vetos were successfully operated, while the inner detector suffered a mechanical problem and failed to show transitions on its TES. Efforts were subsequently focused on the inner detector, which was fully operational in Run 4.

In Run 4, operation parameters of target and inner veto detectors proved constrained by the cross-heating of the detectors in mechanical contact. As a mitigation measure, new TES designs were developed for future prototypes (shown in Figure 8.42). The inner veto TES features a heater structure overlapping the TES tungsten film (electrically isolated using a silicon oxide layer). The same concept was used in the CRESST-III iStick detectors. The overlapping heater allows much smaller powers to control the TES temperature, thus reducing the problem of cross-heating. This will hopefully allow the target (and also several targets) to be operated with ultimate energy resolution while in contact with active inner veto detectors.

In addition to operating all components of the NUCLEUS-1g prototype, initial studies on the anticoincidence veto performance were done. The outer vetos showed a time resolution of 1.26 ms using an optimum filter-based trigger method. The inner vetos in Run 4 allowed a lower background rate between 1-2 keV compared to Run 1 and a lower energy ~ 300 eV down to which the background was observed to be flat. This indicates that covering the surfaces of the target detector with instrumented material removes contributions to low-energy backgrounds. Both outer and inner veto capabilities can be extended much beyond the first demonstration shown in this work. More advanced

onset determination techniques and optimized readout electronics can improve the time resolution of the outer veto. The inner veto cut can be significantly refined by taking into account also the (veto and target) pulse shape that carries information about the origin of an event.

The backgrounds observed with the 1g-prototype are far from the levels that need to be achieved for a successful $CE\nu NS$ measurement. This is especially true for low energies (the $CE\nu NS$ region of interest < 100 eV) where an exponential rise is observed in all detectors. At least a factor of 10^6 suppression is needed before $CE\nu NS$ comes in reach. Important lessons on properties and nature of this background will be learned from operating all prototype components simultaneously, and under increasingly shielded conditions. These next steps are planned to occur on the way to the NUCLEUS blank assembly in the underground lab of TUM-UGL.

9. Conclusion and Outlook

Several avenues to improve low-threshold cryogenic detectors for rare event searches have been presented in this work. These approaches include magnetic field compensation and stabilization, the exploration of TES designs, the demonstration of lower thresholds on smaller target sizes and the characterisation of cryogenic veto detectors against surface and external backgrounds. These efforts have been applied to two experiments searching for low-energy nuclear recoils: CRESST, in direct dark matter search, and NUCLEUS, aiming for a detection of coherent nuclear scattering of reactor antineutrinos. In addition, detailed sensitivity studies have been developed to show the feasibility and physics potential of the NUCLEUS experiment.

A systematic study of the influence of magnetic fields on TES sensors was performed, motivated by the observation of external disturbances in CRESST detectors. At MPP Munich, it was shown that a magnetic field shifts the transition temperature of a large, unstructured tungsten film by 1 mK for every 20 μT of residual transverse field. For structured TES including superconducting phonon collectors, the temperature shift was larger and the transition was observed to get wider in the presence of magnetic fields. Most severe distortions occurred at high bias currents and for small tungsten areas, showing the interrelation of currents and fields in superconductors. The dependence on TES geometry points to flux expulsion by the aluminum phonon collectors. These studies demonstrate that the Earth magnetic field and fields from large equipment (cars, cranes, machinery) close to the experiment impact the performance and stability of CRESST detectors, which rely on a tightly controlled operating resistance in the sharp superconducting transition. In response, a field compensation system was designed and constructed at the CRESST facility. It consists of three coil pairs surrounding the cryostat and a fluxgate magnetometer installed below, allowing to constantly update the currents in response to magnetic field transients. The field compensation was characterised and the detectors were operated in suppressed field conditions for several weeks. Improved stability under external magnetic perturbations was demonstrated. TES gains enhanced by up to a factor 2 were observed, due to the sharper transitions. These improved detector performances, however, made temperature stabilization more challenging for some detectors. Unlocking the natural sharpness of TES transitions with magnetic field compensation allows for higher sensitivity, but comes with the burden of more careful detector setup and stabilization.

The CRESST-III experiment extended its reach below 100 eV nuclear recoil energy by scaling down absorber sizes from 300 g to 24 g. Together with the absorber crystal volume, also the TES area should be scaled down for maximum collection of athermal phonons per TES area. Previously [231], TES designs with 36 mm² (L) and 18 mm² (M) were tested on CaWO₄ crystals. From the observed pulse rise times, it was inferred that the design L thermalizes nearly half the athermal phonons. This motivated the study of smaller TES designs, as the heat capacity is expected to shrink more strongly with TES area than the athermal phonon collection in this regime. In this work, a new TES design

(S, 9 mm²) was introduced and characterised on Al₂O₃ and silicon crystals for future CRESST runs. The additional rise time data from seven CRESST detectors operated at LNGS with L- and M-TES was studied. In spite of significant scatter between the individual detectors, a strong increase of athermal phonon life-time for smaller TES areas emerges for all three studied crystal materials. Taking the data at face value, the M-structure is shown to be too large for optimal signal collection (on CaWO₄, Al₂O₃ and Si). This is in tension with the initial study of [231] and with model calculations for athermal phonon absorption by the TES. To address this tension, new measurements and analysis methods are needed. A first step for future studies should be the characterisation of an even smaller TES, such as a modified light detector TES. Longer datasets should be acquired and combined with a precise transition measurement for each detector. These additional measurements are challenging above ground, but could give important further input to the interpretation of the rise time data. On the modeling side, including the absorption of the phonon collectors and their partial energy delivery to the TES could alleviate the tensions between the data and existing calculations.

CE ν NS is unique among neutrino interactions with its high cross-section and low-energy signature, a nuclear recoil analogous to the signature of dark matter interactions. The goal of the NUCLEUS experiment, making reactor neutrinos measurable with cryogenic detectors, motivated an attempt to obtain more sensitive, even smaller detectors. A combination of sensitivity and technical considerations led to gram-scale cryogenic detectors, with cubical crystals of 5 mm edge lengths. Based on the success of the 24 g CRESST-III detectors (best energy threshold 30 eV), such detectors are extrapolated to reach energy thresholds around 10 eV, which makes them well-suited to search for CE ν NS of low-energy reactor neutrinos.

Statistical studies of various discovery scenarios identified a nuclear power reactor as the preferred neutrino source, and showed that a 10 g target at a distance of 40 m can make a significant CE ν NS detection within a few weeks. A flat background of 10² counts/(keV · kg · day) in the region of interest (from threshold up to ~200 eV) can be tolerated. This defines the basic parameters for the first phase of the NUCLEUS experiment. More detailed statistical studies showed that a background of unknown shape can be handled by deploying targets with light nuclei (Al₂O₃) in addition to the heavy CE ν NS targets (CaWO₄), as long as a common background model can be fitted to the spectra of both materials. Reactor-off periods, as available at the Chooz nuclear power plant planned as the site for NUCLEUS, do not strongly enhance the sensitivity of the experiment. The precision of the CE ν NS cross-section measurement with a 10 g detector can reach 10% with one year of measurement time. A future larger detector array with 1 kg target mass can achieve 1% precision, subject to systematic uncertainties.

Detailed model calculations exploring the sensitivity of the NUCLEUS phases to physics scenarios beyond the Standard Model were performed in this thesis. Assuming the background goal of 10² counts/(keV · kg · day) can be met, NUCLEUS can make a new measurement of the weak mixing angle at low momentum transfer with 10 g target mass, constrain several parameters describing non-standard interactions of electron neutrinos with quarks and search for new light mediators that could modify the weak interactions of neutrinos at low energies. A future 1 kg target allows precision measurements in the former cases and additionally opens sensitivity to neutrino electromagnetic form factors.

The NUCLEUS experiment attempts to open a previously unexplored energy range

for rare-event searches. The background levels necessary for CE ν NS detection were not previously demonstrated with cryogenic detectors operating above ground. An innovative combination of cryogenic anticoincidence detectors was developed to address this challenge. It comprises an instrumented holder for the target detectors, which doubles as a 4π surface veto. A more massive outer cryogenic veto suppresses multiple-scattering backgrounds from penetrating α and γ radiation. Prototype measurements with the target, silicon holder and instrumented inner and outer vetos were performed in this work. The first NUCLEUS target detector showed an energy resolution of 3.7 eV at low energies. A dataset obtained with a threshold of 19.7 eV allowed placing new limits on light dark matter with masses down to 140 MeV/ c^2 . This demonstrates the required sensitivity for a CE ν NS observation. The second prototype, operated in the silicon holder, had an energy resolution of 2.9 eV and could be stably operated. The background rate in the unshielded setup was high ($\mathcal{O}(10^5)$ counts/keV/kg/day) and a sharply rising background below ~ 200 eV was shown not to be related to the ^{55}Fe calibration source. The massive outer veto, consisting of 5 cm silicon cylinders, was commissioned separately. Despite a high particle rate (tens of counts per second), the detectors could be operated successfully. Individual pulses were identified and showed an optimum-filter time resolution of 1.3 ms. Further studies of its operation are planned in a dedicated shielded environment. The final run performed in the framework of this thesis operated a target detector simultaneously with the two inner veto holding plates. All three detectors achieved baseline energy resolutions around 10 eV, which fulfills the design goal of sub-100 eV inner veto thresholds. The instrumented vetos are shown to efficiently tag energy depositions in the vetos. This leads to a flat background down to ~ 150 eV in the target. Cross-talk mediated by phonon transmission through the point-like contacts between target and vetos was observed for particle energy depositions. Similarly, the electric heating applied to each detector for temperature stabilization affected also the other detectors, which restricted the operating point optimization for the target. An updated design of the electric heaters should mitigate the latter issue in the future and enable a target energy resolution comparable to the first runs. A detailed study of the veto efficiency has to follow with a high-performing target detector.

All prototype measurements show a rising background in the CE ν NS region of interest, superficially similar to the low-energy background observed at much lower rates in CRESST and other low-threshold experiments. The detailed study, identification and reduction of this background by around 6 orders of magnitude is mandatory before a successful CE ν NS measurement is possible. This work has assembled a set of tools that allow addressing one of the most important challenges of rare-event searches at present time: the understanding and suppression of low-energy backgrounds that currently limit the sensitivity to new physics.

A. Appendix to active magnetic field compensation for CRESST

A.1. Fields of rectangular coils

The field of an ideal rectangular coil can be calculated analytically. A straight conductor from $(-x_0, y_0, z_0)$ to $(+x_0, y_0, z_0)$, carrying current nI produces the field

$$\vec{B}(x, y, z) = \frac{\mu_0 n I}{4\pi \rho^2} \cdot \begin{pmatrix} 0 \\ -(z-z_0) \\ y-y_0 \end{pmatrix} \cdot \left(\frac{x+x_0}{(x+x_0)^2 + \rho^2} - \frac{x-x_0}{(x-x_0)^2 + \rho^2} \right) \quad (\text{A.1})$$

with $\rho^2 = (y-y_0)^2 + (z-z_0)^2$. A rectangular coil pair can be modeled as a superposition of such terms.

The on-axis field of a single square coil of side length a works out to

$$\vec{B}(z) = \frac{4\mu_0 n I}{\pi a} \cdot \vec{e}_z \cdot (1 + (2z/a)^2)^{-1} \cdot (2 + (2z/a)^2)^{-1/2} \quad (\text{A.2})$$

The Helmholtz condition for vanishing first three derivatives in a coil pair of spacing d reduces by symmetry to $\partial_z^2 |\vec{B}(z)|_{z=d/2} = 0$, which has the analytic solution

$$\frac{d}{a} = \left(\left(\frac{1}{3} - \frac{2}{6^3} \sqrt{610} \right)^{\frac{1}{3}} + \left(\frac{1}{3} + \frac{2}{6^3} \sqrt{610} \right)^{\frac{1}{3}} - 1 \right)^{\frac{1}{2}} \approx 0.544506 \quad (\text{A.3})$$

A sketch of a cubical 3-axis Helmholtz compensation system is shown in Figure A.1 (left). The right panel shows the absolute B-field generated by one coil pair of side length 1, as a function of position along the symmetry axis. The deviation from the central value is below 1% for $|x| < 0.34$.

A.2. Determining TES orientations

Given that the resistance of a thin-film TES is a function of the normal magnetic field component, but not of the parallel ones, the relative orientation of the TES and the applied coil fields can be determined. In this way, TES misalignments can be probed (by looking for outliers among the TES) or distortions of the applied magnetic field studied (which would appear as a trend in computed orientation over the experimental space).

The method is illustrated in Figure A.2: the data from a two-component magnetic field sweep is projected along a direction, which is then varied to match the TES orientation. In this case, the clear image of the TES response to normal fields appears. For a deviating angle, the projection is misaligned and the resulting plot shows more scatter and less clear structure. The applied normal field is given by

$$B_n = B_x \cos \alpha + B_y \sin \alpha \quad (\text{A.4})$$

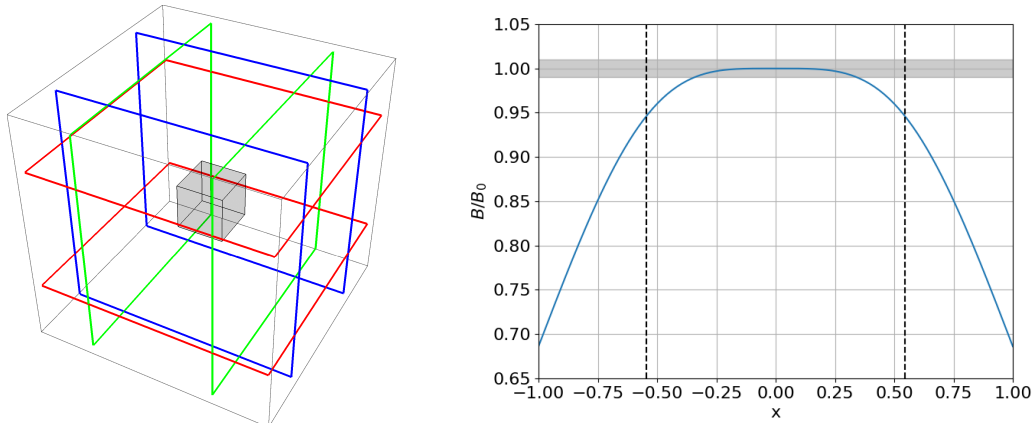


Figure A.1.: Left: ideal 3-axis cubical Helmholtz compensation coils. Right: field strength on the axis of a square Helmholtz coil pair with side length 1. The field strength is normalized to the central value, the positions of the coils are indicated by vertical dashed lines. The shaded region shows $\pm 1\%$ deviation from the central value.

for a given rotation angle α . A numerical estimator is used to quantify the amount of structure in the resulting dataset of step size versus applied normal field. A simple “sharpness criterion” works by binning the data in B_n and adding up the standard deviations in step size of the individual bins. Minimizing the sharpness criterion as a function of α gives the TES orientation angle. Changing the number of bins used and cross-checking the result by eye for a few TES, the precision of the method is estimated to be about 1.5° .

Assuming the coils produce mutually perpendicular fields of equal strength at the TES location, the “normal field vector” directly corresponds to the geometric vector perpendicular to the TES film. Since two-dimensional magnetic sweeps are used, the resulting angle is between the coil field directions and the projection of the TES normal vector on the plane of the magnetic field sweep.

The data used to determine the orientations of the group representatives is shown in Figures 4.15 (xy-sweep) and A.3 (zy-sweep). For most TES, the data contains multiple repetitions of scans across the low-normal-field region, offset in the slowly sweeping field direction. This allows relatively precise determination of the angles. The vertical TES are an exception: being almost orthogonal to the slowly sweeping z-coil pair, only a slow single scan across the response pattern is recorded. Thus, a misalignment of the projection angle does not lead to scatter in the response plot, but only to a distortion of the slowly scanned response. Without using the shape of the response itself as an observable, the orientation angle can only be constrained to less than the opening angle of the zig-zag line of the two-component sweep (9.8° in the yz-sweep). A larger misalignment would lead to non-monotonous change of the normal field component, which would be visible in the response plot. This shortcoming could be avoided by complementing the zy-sweep with a yz-sweep covering the same magnetic field configurations, but sweeping the z-coil fast and the y-coil slow.

Table A.1 shows the results of the angle measurements. As expected, light and phonon

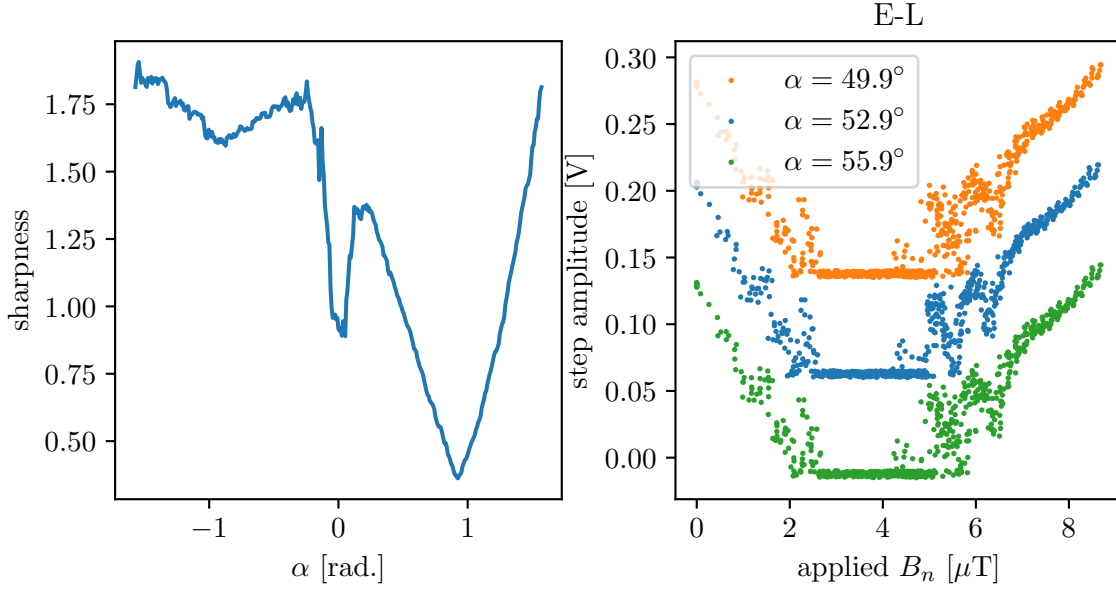


Figure A.2.: Illustration of TES orientation finding. The left panel shows the xy-sweep data of detector E-L also displayed in Figure 4.15 (right). The step amplitude is plotted against applied normal field, $B_n = B_x \cos \alpha + B_y \sin \alpha$, which is calculated for three different assumed rotation angles α of the TES. The clearest self-overlap occurs at $\alpha = 52.9^\circ$. Projections mis-rotated by $\pm 3^\circ$ are shown (with an arbitrary vertical offset) to illustrate the accuracy of the method. The right panel shows the sharpness criterion, a bin-wise standard deviation of the data, as a function of rotation angle. Its minimum gives an automated determination of the rotation angle.

detectors are separated by about 90° in the xy-plane, while the deviations within each group are below the precision of the angle determination. Phonon and light detector TES are also expected to be vertical, i.e. $\theta_{zy} = 90^\circ$. While the light detector are found to be vertical within a degree, the phonon detectors deviate much more strongly within a range of almost 10° . It should be noted that the angle θ_{zy} , due to the projection onto the yz-plane, is larger than the angle δ between the normal vector and the xy-plane. They are related through

$$\tan \delta = \frac{\sin \theta_{xy}}{\tan \theta_{zy}}. \quad (\text{A.5})$$

As a local coil field distortion should affect light and phonon TES alike, this deviation from vertical in the orientation of the phonon detectors can be attributed to the holding system using the crystal sticks. Apparently the sticks allow for enough freedom in crystal orientation that a few-degree deviation from the nominal position is possible. This corresponds to a displacement of the crystal edges by a 1-2 millimeters. This is important to keep track of, but the holder design should provide enough safety margin to avoid touches (e.g. 4.8 mm spacing between phonon and light detectors).

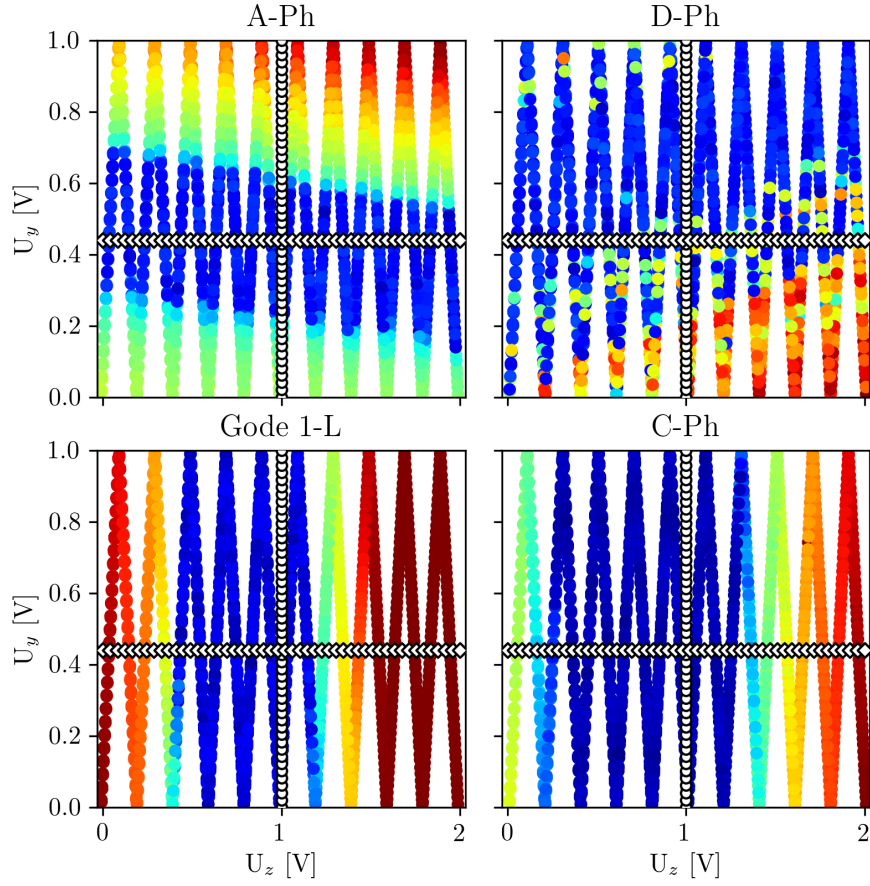


Figure A.3.: Two-component sweep analogous to Figure 4.15, using the z- and y-coil pairs. Phonon and light TES show various small tilts (top row) while vertical TES (bottom row) respond almost exclusively to the z-coil pair.

A.3. Simulation of compensation algorithm

To visualize the action of the algorithm, it helps to look at an example generated in a simulation of the magnetic system and its response to an external disturbance. In this example, we consider only the wall/door magnetometers, resulting in a 6×3 response matrix. The response matrices are idealized versions of the ones derived for the CRESST setup:

$$M_c = \begin{bmatrix} 3.3 & 0 & 0 \\ 0 & 3.3 & 0 \\ 0 & 0 & 5.7 \end{bmatrix} \mu\text{T/A}, \quad M = \begin{bmatrix} 1.6 & 0 & 0 \\ 0 & 4.4 & 0 \\ 0 & 0 & 2.5 \\ 1.6 & 0 & 0 \\ 0 & 4.4 & 0 \\ 0 & 0 & 2.5 \end{bmatrix} \mu\text{T/A}. \quad (\text{A.6})$$

The reduction operation for the wall/door magnetometers is a simple average between each component. Due to their symmetric location on opposite sides of the carousel, this correctly takes into account linear gradients in all field components. The operation can

Det.	A-Ph	D-Ph	E-L	J-L	C-Ph	Gode 1-L	Gode 1-C
θ_{xy}	-35.5°	-36.5°	+52.9°	+52.2°	–	–	–
θ_{zy}	+87.3°	+97.1°	+91.1°	+89.0°	< 9.8°	< 9.8°	< 9.8°
δ	+1.6°	-4.2°	+0.8°	-0.9°	–	–	–

Table A.1.: TES orientations measured from two-component magnetic field sweeps. θ_{xy} is the angle between the x-axis and the projection of the TES normal vector onto the xy-plane. Similarly, θ_{zy} is the angle between the z-axis and the projection of the TES normal vector onto the yz-plane. The vertical TES C-Ph, Gode 1-L and Gode 1-C show no visible response to the x- and y-coils, and their deviation from the z-axis can only be constrained weakly, as the z-coil was sweeping slowly in the zy-scan.

be expressed as a matrix multiplication:

$$\left(\vec{B} - M\vec{I}\right)_{red} = R \cdot \left(\vec{B} - M\vec{I}\right), \quad R = \frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}. \quad (\text{A.7})$$

Since the virtual magnetometer response matrix is diagonal (because the coil axes are nominally aligned with the magnetometer axes), the compensation matrix is simply:

$$C = -\left(M_c^T M_c\right)^{-1} M_c^T = -M_c^{-1} = - \begin{bmatrix} 1/3.3 & 0 & 0 \\ 0 & 1/3.3 & 0 \\ 0 & 0 & 1/5.7 \end{bmatrix} \text{ A}/\mu\text{T}. \quad (\text{A.8})$$

At each time step, the measured fields at the magnetometers are calculated from equation 4.4 plus a random noise contribution. The current on the coils for the next step is calculated from equation 4.8. The disturbance is taken to be zero in the beginning, then slowly ramp up to a constant value and later down again to zero. An important assumption is spatial homogeneity (i.e. the disturbance is the same at the measurement and control locations). For simplicity, the time dependent external field is taken to be equal in all field components. The coils start out with zero current. Figure A.4 shows the results of this simulated example, displaying the time-varying background field, the measured field at one of the magnetometers, the extrapolated field at the control location, and the coil currents.

As expected with perfect knowledge of M and M_c , the central fields scatter about zero when the background field is stable. The noise on the central fields comes from the magnetometer noise, which is reduced by averaging and multiplied by the compensation matrix (since we apply a noisy estimate of the optimal current) and the central response matrix. (Since the compensation matrix equals M_c^{-1} , the central field noise reduces to the measured field noise divided by \sqrt{n} with the number of magnetometers n .)

The plot shows how the coil currents slowly ramp to a steady negative value (needed to cancel the positive disturbance), and ramp back following the end of the disturbance. The imperfect canceling during the ramping of the external field is due to the delay (of one step) with which the algorithm reacts to the new situation. To cancel the external field (equal in each component), a smaller current is needed on z than on x and y, since the Helmholtz z-pair is more effective at the center than the too-widely-spaced x- and y-pairs.

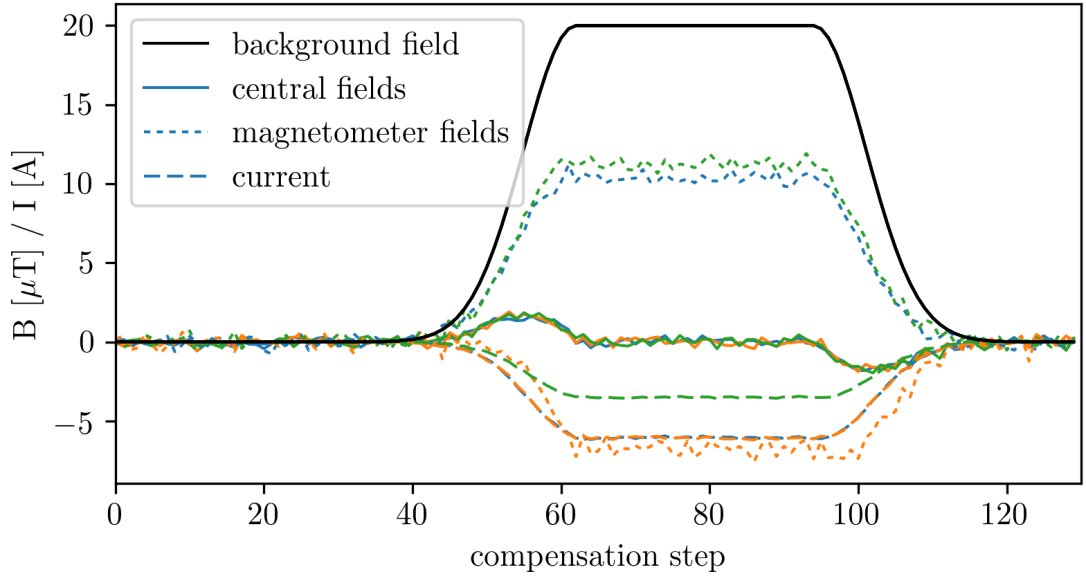


Figure A.4.: Simulated performance of the compensation algorithm canceling a time-dependent homogeneous background field. The calculation uses idealized matrices approximating the real situation at CRESST. Plotted are the external field ($x=y=z$, black), as well as the components of the central field, measured magnetometer field and the applied current (x : blue, y : orange, z : green).

It can be seen that the measured magnetometer fields do not compensate the external disturbance. Instead, for the x - and z -components about half of the external disturbance is canceled out at the magnetometer location, while the y -component is overcompensated. This is exactly as necessary to compensate at the control location, since in x and z the control location is more affected by the coils than the magnetometer location, while in y (due to the close proximity of the coils), the magnetometer is more affected. In the real setup, this leads to an inconvenient situation: the actual controlled field (in the cryostat) cannot be measured, so the effectiveness of the compensation can only be inferred from the TES performance.

If for example the elements of M_c are wrongly estimated by a factor α , the central field gets a systematic offset proportional to the external disturbance. The central field will scatter around $B_0 \cdot (1 - \frac{1}{\alpha})$ instead of zero, i.e. only a fraction $1/\alpha$ will be compensated. With a 10%-level agreement between the geometric model and the measured M , it seems reasonable to assume a similar accuracy for M_c . This means that external disturbances will be reduced by an order of magnitude by the active compensation. The DC-part of the magnetic field is not affected by this statement, as it is taken directly from TES observations and goes into the B_{tar} term.

A.4. Configuration of field compensation at CRESST

Setting up the compensation system, there are some choices to be made and further details to explain. Initially, the plan was to have the precise fluxgate as an independent




Figure A.5.: Responses of the magnetometer channels wall-x (left) and fluxgate-x (right) to a disturbance caused by moving the experimental hall crane on June 19, 2019. The scale of the magnetic field change is about $0.2 \mu\text{T}$, which is enough to cause a distortion of the detector energy scales. This is well resolved by the fluxgate, but at the edge of the resolution limit for the arduino-based magnetometers.

cross-check measurement not used in the compensation algorithm. Instead, five simpler arduino-based magnetometers were planned to be deployed symmetrically around and below the experiment, to get a field map capable of interpolating the magnetic field around CRESST. Initially only two of the five magnetometers, wall and door, were mounted and incorporated into the data taking system. Their precision of $0.2 \mu\text{T}$ was deemed sufficient to get a reasonable measurement of the earth magnetic field, and therefore to compensate it. A first test of the installed system with an external disturbance showed this to be insufficient. Figure A.5 shows the responses of magnetometer wall-x and fluxgate-x to a changing magnetic field induced by the hall crane being parked on top of the experiment, and moved back again half an hour later (similar to the observation of Figure 4.8). The resolution of the arduino-based magnetometers is not good enough to precisely detect and correct for the influence of the crane, which has previously been shown to be relevant to detector stability. The simple magnetometers may be enough to measure and compensate the earth magnetic field, but the requirements for stable detector operation are more stringent. The resolution of the fluxgate magnetometer is necessary.

Another way to gauge the needed magnetometer precision comes through the noise on the controlled current. As discussed in the previous section, the magnetometer noise propagates to the coil current and thus to the field at the carousel. One bit ($20 \text{ V} / 2^{16}$) of control voltage corresponds to a current step of 0.6 mA at the power supply, which is on the order of the specified ripple of 1.5 mA at 300 kHz bandwidth. The smallest possible current step thus corresponds to a field of $2\text{-}3 \text{ nT}$ (in x , y and z). The fluxgate magnetometer approaches this precision, which makes it possible to have the current analog control voltage stable to the last bit (in fixed field conditions). Under these circumstances, the field stability is limited by the power supply current ripple. The less precise magnetometers instead cause a significantly higher current fluctuation (e.g. $1/3.3 \text{ A}/\mu\text{T} \cdot 0.2\mu\text{T} / \sqrt{2} = 43 \text{ mA}$).

For these reasons, the active field compensation was set up using only the fluxgate magnetometer as input. The algorithm described above was implemented as a new process called `fluxcompensator` (or `flux` for short) in the CRESST control software `ccs`. The response matrix was set to the first three rows of Table 4.1 (left). The reduction operation, instead of averaging over multiple magnetometers, just takes the fluxgate measurement as the virtual magnetometer reading at the detector carousel. The effect of this is that any field gradient between the fluxgate location and detector carousel is lost to the compensation. It is always assumed that a disturbance affects the detectors with the same direction and magnitude as it does the fluxgate, and needs to be compensated accordingly. The compensation matrix C in “natural units” (coil control voltage / magnetometer output voltage) is

$$C = -(M_c^T M_c)^{-1} M_c^T = -M_c^{-1} = - \begin{bmatrix} +1.502 & 0 & 0 \\ 0 & -1.502 & 0 \\ 0 & 0 & -0.871 \end{bmatrix}. \quad (\text{A.9})$$

The two minus-signs are needed because the y-coil is oriented opposite the y-magnetometer channel, and the z-magnetometer channel (of the fluxgate, and thus also the virtual central magnetometer) is oriented opposite the z-coordinate axis.

```

CCS>flux
***** Global settings *****
Wait time between two dvm reading -w 0,1s
Number of dvm readings to average -n 50
*****

***** Axis settings *****
***** X-Axis *****
DAC Channel -c 24
DAC Maximum -m 2V
Virt. Setpoint -s -0,1851V
***** Y-Axis *****
DAC Channel -c 25
DAC Maximum -m 2V
Virt. Setpoint -s 0,2543V
***** Z-Axis *****
DAC Channel -c 26
DAC Maximum -m 2V
Virt. Setpoint -s 0,656V
*****

***** Magnetometer settings *****
Magnetometer 1 is the virtual magnetometer.
***** X-DVM Channel (-xdvm) [integer] *
1- 8| -1 11
***** Y-DVM Channel (-ydvm) [integer] *
1- 8| -1 12
***** Z-DVM Channel (-zdvm) [integer] *
1- 8| -1 13
***** X-DVM averaged value [V] *****
1- 8| -0,4835 -0,3575
***** Y-DVM averaged value [V] *****
1- 8| 0,5106 0,4077
***** Z-DVM averaged value [V] *****
1- 8| 1,8216 1,0762
*****

```

Figure A.6.: Help screen of the process `flux` controlling the field compensation system within the CRESST control software `ccs`, showing the configuration parameters.

Figure A.6 shows the help-text of the `flux` program, including all configurable parameters. With a command like `flux 1 -c24 -m2 -s-0.1851` one can set the DAC output channel for the x-coil control voltage (to channel number 24), limit the output voltage to 2 V and set the virtual magnetometer x-setpoint to -0.1851 V. The command `flux 2 -xdvm 11 -ydvm12 -zdvm13` assigns the correct dvm channels of the fluxgate to magnetometer 2 (magnetometer 1 is reserved for the virtual central magnetometer, and has no dvm channels).

Three operationally important parameters beyond the basic configuration are waiting time, number of dvm readings and damping. Instead of reading the magnetometer once and then updating the currents, many readings can be collected, then averaged and used in a single current update. This reduces the impact of magnetometer noise at the




Figure A.7.: Comparison of two settings for the active compensation. In both settings, 50 magnetometer readings taken in an interval of 5 s are averaged before the calculation step. To show the approach of the compensation to a new current configuration, the target magnetic field was changed to a different value and back again. In the left panel, the currents are set fully according to the calculation. In the right panel, the current is updated to 70% of the old value and 30% of the new value. This results in a slower approach to the new setpoint with stronger suppression of fluctuations.

cost of update speed. The damping parameter is a number between 0 and 1 reducing the magnitude of the current update. Instead of forgetting the old output currents I_0 and fully applying the newly calculated I_1 , the linear combination $(1 - d) \cdot I_0 + d \cdot I_1$ is applied. This further suppresses magnetometer noise and turns a steplike response to a field change into an exponential approach. Figure A.7 illustrates the effect of damping by showing the algorithm response to sudden changes of the magnetic field setpoint. In the left panel, no damping is used, and the current control voltage follows the setpoint change in one step. Oscillations by one bit, caused by magnetometer noise, can be seen in the flat sections of the control output. In the right panel, the parameter d is set to 30%. This suppresses the oscillations at the cost of a slower approach to the new setpoint. As the convergence happens over less than one minute, this is viewed as a good compromise and chosen for the running configuration of the active compensation.

B. Statistical methods for sensitivity studies

B.1. Likelihood-Ratio-Test

Making a statistically significant observation of $\text{CE}\nu\text{NS}$ boils down to the task of rejecting the null hypothesis “absence of $\text{CE}\nu\text{NS}$ ” for a given background model and dataset (containing contributions from background and $\text{CE}\nu\text{NS}$). This is a classic model selection problem with two nested models, since the background only “null model” can be formulated as a special case of the “free model” including $\text{CE}\nu\text{NS}$. The free model contains a description of the background with some parameters and one parameter for $\text{CE}\nu\text{NS}$ signal normalization. The null model is contained in the free model, by setting the signal strength parameter to zero. A classic frequentist test developed for this situation is the Likelihood-Ratio-Test.

The likelihood function $\mathcal{L}(D, \Theta)$ describes the probability of observing the dataset D given values of the model parameters Θ , and is considered a function of Θ . Separate likelihood functions $\mathcal{L}_{null}(D, \Theta_{null})$ and $\mathcal{L}_{free}(D, \Theta_{free})$ are constructed for the null and free models, where the model parameters $\Theta_{free} = (\Theta_{bg}, \Theta_{sig})$ contain background parameters and a signal strength parameter. For the null model, the signal strength parameter is set to zero: $\Theta_{null} = (\Theta_{bg}, \Theta_{sig} = 0)$.

A test statistic W is constructed from the ratio of separately maximized likelihood functions:

$$W = 2 \log \frac{\hat{\mathcal{L}}_{free}}{\hat{\mathcal{L}}_{null}} = 2 \log \frac{\mathcal{L}_{free}(D, \hat{\Theta}_{free})}{\mathcal{L}_{null}(D, \hat{\Theta}_{null})} \quad (\text{B.1})$$

where $\hat{\Theta}$ is the parameter value maximizing the respective likelihood function.

By Wilks’ theorem, this test statistic for nested models follows a χ^2 -distribution with the number of extra degrees of freedom in the free model. For one extra degree of freedom (signal strength), one has a χ^2 -distribution with one degree of freedom, so that \sqrt{W} follows a standard normal distribution. For a given dataset, the significance of $\text{CE}\nu\text{NS}$ observation in σ can therefore be directly read off from \sqrt{W} .

To make such calculations, we have to specify how to calculate the likelihood function, dependent on dataset and model, in our context. In our case, the model specifies the expected total event rate as a function of some parameters. The dataset is an observed “spectrum”, i.e. a number of particle events with an energy. Several choices are possible in specifying the likelihood function for a given dataset and model.

B.2. Unbinned likelihood

The most generic, but computationally intensive, likelihood function is an unbinned likelihood directly computed from the events without first creating a histogram. The probability of observing a given spectrum is the product of the probability to observe the given total number of events, and the probabilities of observing the individual events

at their respective energies:

$$\mathcal{L}(D, \Theta) = \frac{\mu^n}{n!} e^{-\mu} \cdot \prod_{i=1..n} p(E_i) \quad (\text{B.2})$$

for n observed events with energies E_i , $i = 1..n$. $p(E)$ is the event rate as a function of energy $R(E)$ (as specified by the model) normalized to unity. μ is the model expectation value for the number of observed events, i.e. the integral over $R(E)$ (and the normalization factor between $p(E)$ and $R(E)$). As the only relevant quantity is the ratio of two likelihood functions evaluated for different models, but the same dataset, any normalization independent of the model parameters (such as the factor $n!$) can be dropped. For computational reasons, the negative natural logarithm of the likelihood function is used:

$$-\log \mathcal{L}(D, \Theta) = \mu - \sum_{i=1..n} \log R(E_i) \quad (\text{B.3})$$

which can be safely numerically minimized over the model parameters. The term $n \cdot \log \mu$ from the Poissonian likelihood cancels against the normalization of the event rate.

If the dataset consists of several spectra, i.e. observed in different detectors or under different signal or background conditions, the negative-log-likelihood functions for the different spectra can be added (as the separate observations are statistically independent).

This unbinned approach allows a direct calculation of likelihood values from the model rate and observed spectra. It assumes perfect energy resolution of the detector, and does not introduce any loss of information from binning.

B.3. Binned likelihood

The effort of evaluating the unbinned likelihood rises linearly with the size of the dataset. This approach becomes unpractical for very large datasets. Therefore it is more common in particle physics to take a histogram of the observed spectrum to be the dataset on which statistical analysis is performed. The number and size of bins has to be chosen correctly to avoid losing relevant information from the spectrum. In this approach, the number of bins sets the complexity of the likelihood function, so that larger datasets can be handled.

The probability of observing a given number of events n_i in each energy bin i is a product of Poissonian terms:

$$\mathcal{L}(D, \Theta) = \prod_{i=1..n} \frac{\mu_i^{n_i}}{n_i!} e^{-\mu_i} \quad (\text{B.4})$$

where the bin expectation value μ_i is the integral of the event rate over each bin. As before for the unbinned case, some simplifications are possible:

$$-\log \mathcal{L}(D, \Theta) = \mu - \sum_{i=1..n} n_i \cdot \log \mu_i \quad (\text{B.5})$$

Like before the evaluation of the likelihood functions “costs” an integration of the event rate over the energy range of interest, this time subdivided into the energy bins to find

the μ_i . The description of the dataset as a list of n_i is much more compact and reduces the number of terms in the likelihood function.

Again, the negative-log-likelihoods from multiple spectra can be added. Here, the finite energy resolution of the detector introduces correlations between neighboring bins, which is neglected in this approach.

The choice of the bins is a compromise between preserving spectral information and keeping the total number of bins low for speed. In the case of CE ν NS, the most interesting part of the spectrum is close to threshold, while higher-energy regions are practically signal-free and can be useful for background information. To keep the total number of bins manageable while still capturing the energy dependence of the CE ν NS recoil spectrum, a binning with exponentially increasing bin size was chosen. The binning is specified by the following criteria: 100 bins span the range from threshold to a maximum energy of 2 keV. The first bin width is one fifth of the energy threshold. Each following bin increases in width by a factor chosen to reach the maximum energy after the given total number of bins.

B.4. Toy Monte Carlo

The last missing ingredient for a sensitivity study (apart from specifying the model, i.e. dependence of $R(E)$ on Θ) is a way to generate random datasets for each model scenario. For the binned likelihood, the dataset is a list of event numbers in the energy bins. It can be generated by integrating the total rate over each bin and drawing a Poisson-distributed random number with this expectation value. For the unbinned likelihood, the computational problem is to draw random numbers directly from the probability distribution $p(E)$ given by normalizing the event rate. This can be done through “inverse transform sampling”, i.e. evaluating the inverse cumulative distribution of the desired pdf with uniformly distributed random numbers. Figure B.1 shows a visualisation of the method. In this way, samples from the model’s background and signal spectra can be generated. The number of samples to be drawn is determined from a Poissonian random variable given the integral over the rate as expectation value.

B.5. Asimov data set

The normal procedure of a toy Monte Carlo study requires generating many random instances of the observed spectrum. The model is then fitted to each instance, allowing to extract the distributions of all parameters of interest. For large studies, the computation cost can become prohibitive, so that simplifications are necessary. An approach developed in high-energy physics uses the so-called “Asimov data set” [94]. By evaluating the test statistic with the “most-likely dataset” (i.e. setting all bin contents to their non-integer expectation values), one obtains the median of the test statistic W . Since the significance is a monotonic function of W , the median significance can be derived.

Explicitly, for a counting experiment, the likelihood ratio contains the Poissonian probabilities for the bin contents:

$$\frac{\mathcal{L}_{free}}{\mathcal{L}_{null}} = \frac{\prod_i \mu_i^{n_i} e^{-\mu_i}}{\prod_i \nu_i^{n_i} e^{-\nu_i}} \quad (\text{B.6})$$




Figure B.1.: Visualisation of Inverse Transform Sampling. The desired pdf of the samples (green) is a normal distribution with mean 5 and variance 1 in this example. Its cumulative distribution is shown in orange. Samples from a uniform distribution (blue) are drawn. Intercepting the orange graph horizontally (i.e. evaluating the inverse cumulative distribution) converts the uniform samples into random samples drawn from the desired pdf (green).

with expected bin counts $\mu_i = s_i + b_i$ and $\nu_i = b_i$ composed of signal and background. To find the expected (median) significance of the rejection of the background-only hypothesis in the presence of signal, we set $n_i = s_i + b_i$:

$$W_A = 2 \cdot \log \frac{\mathcal{L}_{free}}{\mathcal{L}_{null}} = 2 \cdot \sum_i (\mu_i \log \mu_i - \mu_i - \mu_i \log \nu_i + n u_i) \quad (\text{B.7})$$

$$= 2 \cdot \sum_i (-s_i + (s_i + b_i) \cdot \log(1 + s_i/b_i)) \quad (\text{B.8})$$

With this formula, given the signal and background bin contents s_i and b_i , one can quickly evaluate the expected significance of signal discovery $\sqrt{W_A}$. Note that this simple calculation assumes known values of the background, with no free parameters to be fitted.

To study the power to exclude possible deviations from the expected signal, one can replace $b \rightarrow b + s_0$, $s \rightarrow s_1 - s_0$ and set $n \rightarrow b + s_0$. Now the truth is assumed to be the expected signal s_0 with background b , and one studies the sensitivity to distinguish the alternative signal s_1 from s_0 . The corresponding likelihood ratio test statistic is

$$W_{hyp} = 2 \cdot \log \frac{\mathcal{L}_1}{\mathcal{L}_0} = 2 \cdot \sum_i \left(s_{0,i} - s_{1,i} + (s_{0,i} + b_i) \cdot \log \left(\frac{s_{1,i} + b_i}{s_{0,i} + b_i} \right) \right). \quad (\text{B.9})$$

A given deviation $s_1 - s_0$ can be considered excluded at 90% confidence level if $\sqrt{W_{hyp}} > 1.645$.

B.6. Nuisance parameters

A realistic model comparison will contain many uncertainties. These can be expressed as nuisance parameters in the background model, which are additionally varied in the likelihood maximisation. An important example that can be incorporated at this level is the uncertainty of the signal normalization (e.g. when searching for a deviation of the signal shape). Assuming one wants to distinguish the signals s_1 and s_0 while the signal expectation is known with a relative uncertainty σ , one can add a parameter η to the likelihood, replacing $s_{0,1} \rightarrow \eta \cdot s_{0,1}$ and including a pull-term $e^{-(\eta-1)^2/2\sigma^2} / \sqrt{2\pi\sigma^2}$ to represent the uncertainty of η . The nuisance parameter η is separately maximized in the null and free likelihoods. Since the observed spectrum in the Asimov dataset follows $s_0 + b$, the null likelihood

$$\log \mathcal{L}_0 = \left(\sum_i -(\eta \cdot s_{0,i} + b_i) + (s_{0,i} + b_i) \cdot \log(\eta s_{0,i} + b_i) \right) - \frac{(\eta - 1)^2}{2\sigma^2} \quad (\text{B.10})$$

is maximal for $\eta = 1$. In the likelihood for the alternative model, some of the signal difference can be hidden by a deviating normalization:

$$\log \mathcal{L}_1 = \left(\sum_i -(\eta \cdot s_{1,i} + b_i) + (s_{0,i} + b_i) \cdot \log(\eta s_{1,i} + b_i) \right) - \frac{(\eta - 1)^2}{2\sigma^2}. \quad (\text{B.11})$$

The value η_{max} maximizing \mathcal{L}_1 can be found by finding the root of the first derivative:

$$\partial_\eta \log \mathcal{L}_1 = \left(\sum_i s_{1,i} \cdot \frac{s_{0,i} - \eta s_{1,i}}{b_i + \eta s_{1,i}} \right) - \frac{(\eta - 1)}{\sigma^2} \stackrel{!}{=} 0. \quad (\text{B.12})$$

The likelihood ratio test statistic including the uncertain signal normalization is

$$W_{norm} = 2 \cdot \log \frac{\mathcal{L}_1(\eta=\eta_{max})}{\mathcal{L}_0(\eta=1)} \quad (\text{B.13})$$

$$= 2 \cdot \sum_i \left(s_{0,i} - \eta_{max} s_{1,i} + (s_{0,i} + b_i) \cdot \log \left(\frac{\eta_{max} s_{1,i} + b_i}{s_{0,i} + b_i} \right) \right) - \frac{(\eta_{max} - 1)^2}{2\sigma^2}. \quad (\text{B.14})$$

Bibliography

- [1] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.
- [2] William Thomson Baron Kelvin. *Baltimore lectures on molecular dynamics and the wave theory of light*. CJ Clay and Sons, 1904.
- [3] Henri Poincaré. The milky way and the theory of gases. *Popular Astronomy*, 14:475–488, 1906.
- [4] Ernst Öpik. Selective absorption of light in space, and the dynamics of the Universe. *Bull. de la Soc. Astr. de Russie*, 21(150):5, 1915.
- [5] J. H. Jeans. The Motions of Stars in a Kapteyn-Universe. *Monthly Notices of the Royal Astronomical Society*, 82(3):122–132, 01 1922.
- [6] J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bulletin of the Astronomical Institutes of the Netherlands*, 6:249, August 1932.
- [7] Gianfranco Bertone and Dan Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, October 2018.
- [8] Marc Schumann. Direct detection of WIMP dark matter: concepts and status. *Journal of Physics G: Nuclear and Particle Physics*, 46(10):103003, August 2019.
- [9] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419–421, July 1965.
- [10] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. *Astrophysical Journal*, 142:414–419, July 1965.
- [11] Planck Collaboration, N. Aghanim, Y. Akrami, et al. Planck 2018 results. VI. Cosmological parameters, 2018.
- [12] Daniel J. Eisenstein, Idit Zehavi, David W. Hogg, et al. Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies. *The Astrophysical Journal*, 633(2):560–574, November 2005.
- [13] Chandra Photo Album. <http://chandra.harvard.edu/photo/2006/1e0657/more.html>, 2016. [Online; accessed July 10 2020].
- [14] Anthony H. Gonzalez, Suresh Sivanandam, Ann I. Zabludoff, and Dennis Zaritsky. Galaxy cluster baryon fractions revisited. *The Astrophysical Journal*, 778(1):14, October 2013.

- [15] Richard Massey, Thomas Kitching, and Johan Richard. The dark matter of gravitational lensing. *Reports on Progress in Physics*, 73(8):086901, July 2010.
- [16] Douglas Clowe, Maruša Bradač, Anthony H. Gonzalez, et al. A Direct Empirical Proof of the Existence of Dark Matter. *The Astrophysical Journal Letters*, 648(2):L109, 2006.
- [17] E. Corbelli and P. Salucci. The extended rotation curve and the dark matter halo of M33. *Monthly Notices of the Royal Astronomical Society*, 311:441–447, January 2000.
- [18] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 $/R = 4\text{kpc}/$ to UGC 2885 $/R = 122\text{kpc}/$. *Astrophysical Journal*, 238:471–487, June 1980.
- [19] A. Bosma. *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types*. Dissertation, Groningen Univ., 1978.
- [20] J. P. Ostriker, P. J. E. Peebles, and A. Yahil. The size and mass of galaxies, and the mass of the universe. *Astrophysical Journal Letters*, 193:L1–L4, October 1974.
- [21] J. Einasto, A. Kaasik, and E. Saar. Dynamic evidence on massive coronas of galaxies. *Nature*, 250:309–310, July 1974.
- [22] J. C. Kapteyn. First Attempt at a Theory of the Arrangement and Motion of the Sidereal System. *The Astrophysical Journal*, 55:302, May 1922.
- [23] J I Read. The local dark matter density. *Journal of Physics G: Nuclear and Particle Physics*, 41(6):063101, May 2014.
- [24] Floor Van Leeuwen. *Hipparcos, the new reduction of the raw data*, volume 350. Springer Science & Business Media, 2007.
- [25] Martin C. Smith, S. Hannah Whiteoak, and N. W. Evans. Slicing and dicing the Milky Way disk in the Sloan Digital Sky Survey. *The Astrophysical Journal*, 746(2):181, February 2012.
- [26] AGA Brown, A Vallenari, T Prusti, et al. Gaia Data Release 2 – Summary of the contents and survey properties. *Astronomy & astrophysics*, 616:A1, 2018.
- [27] Pablo F. de Salas. Dark matter local density determination based on recent observations. *Journal of Physics: Conference Series*, 1468:012020, February 2020.
- [28] Tisserand, P., Le Guillou, L., Afonso, C., et al. Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *A&A*, 469(2):387–404, 2007.
- [29] C. Alcock, R. A. Allsman, D. R. Alves, et al. MACHO Project Limits on Black Hole Dark Matter in the 1–30 M_{\odot} Range. *The Astrophysical Journal*, 550(2):L169–L172, April 2001.

- [30] Timothy D. Brandt. Constraints on MACHO dark matter from compact stellar systems in ultra-faint dwarf galaxies. *The Astrophysical Journal*, 824(2):L31, June 2016.
- [31] M. Milgrom. A modification of the Newtonian dynamics - Implications for galaxies. *Astrophysical Journal*, 270:371–389, July 1983.
- [32] Jacob D. Bekenstein. Relativistic gravitation theory for the modified Newtonian dynamics paradigm. *Phys. Rev. D*, 70:083509, October 2004.
- [33] Benoît Famaey and Stacy S McGaugh. Modified Newtonian dynamics (MOND): observational phenomenology and relativistic extensions. *Living reviews in relativity*, 15(1):10, 2012.
- [34] Constantinos Skordis. The tensor-vector-scalar theory and its cosmology. *Classical and Quantum Gravity*, 26(14):143001, June 2009.
- [35] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119:161101, October 2017.
- [36] S. Boran, S. Desai, E. O. Kahya, and R. P. Woodard. GW170817 falsifies dark matter emulators. *Phys. Rev. D*, 97:041501, February 2018.
- [37] M. Taoso, G. Bertone, and A. Masiero. Dark matter candidates: a ten-point test. *Journal of Cosmology and Astroparticle Physics*, 3:022, March 2008.
- [38] Keith A Olive, Gary Steigman, and Terry P Walker. Primordial nucleosynthesis: theory and observations. *Physics Reports*, 333-334:389 – 407, 2000.
- [39] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astroparticle Physics*, 6:87–112, December 1996.
- [40] M. C. Smith, G. R. Ruchti, A. Helmi, et al. The RAVE survey: constraining the local Galactic escape speed. *Monthly Notices of the Royal Astronomical Society*, 379:755–772, August 2007.
- [41] Walter Dehnen and Justin I Read. N-body simulations of gravitational dynamics. *The European Physical Journal Plus*, 126(5):55, 2011.
- [42] Michael Kuhlen, Mark Vogelsberger, and Raul Angulo. Numerical simulations of the dark universe: State of the art and the next decade. *Physics of the Dark Universe*, 1(1):50 – 93, 2012. Next Decade in Dark Matter and Dark Energy.
- [43] Ricardo A. Flores and Joel R. Primack. Observational and Theoretical Constraints on Singular Dark Matter Halos. *Astrophysical Journal Letters*, 427:L1, May 1994.
- [44] Ben Moore. Evidence against dissipation-less dark matter from observations of galaxy haloes. *Nature*, 370(6491):629–631, 1994.

- [45] Katherine Freese, Paolo Gondolo, and Heidi Jo Newberg. Detectability of weakly interacting massive particles in the Sagittarius dwarf tidal stream. *Phys. Rev. D*, 71:043516, February 2005.
- [46] Volker Springel, Carlos S Frenk, and Simon DM White. The large-scale structure of the Universe. *Nature*, 440(7088):1137–1144, 2006.
- [47] Vid Iršič, Matteo Viel, Martin G Haehnelt, et al. New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data. *Physical Review D*, 96(2):023522, 2017.
- [48] Graciela B Gelmini. Light weakly interacting massive particles. *Reports on Progress in Physics*, 80(8):082201, June 2017.
- [49] Gary Steigman, Basudeb Dasgupta, and John F. Beacom. Precise relic WIMP abundance and its impact on searches for dark matter annihilation. *Phys. Rev. D*, 86:023506, July 2012.
- [50] Howard Baer, Ki-Young Choi, Jihn E Kim, and Leszek Roszkowski. Dark matter production in the early Universe: beyond the thermal WIMP paradigm. *Physics Reports*, 555:1–60, 2015.
- [51] Leszek Roszkowski, Enrico Maria Sessolo, and Sebastian Trojanowski. WIMP dark matter candidates and searches—current status and future prospects. *Reports on Progress in Physics*, 81(6):066201, May 2018.
- [52] Kathryn M. Zurek. Asymmetric dark matter: Theories, signatures, and constraints. *Physics Reports*, 537(3):91 – 121, 2014.
- [53] Lawrence J Hall, Karsten Jedamzik, John March-Russell, and Stephen M West. Freeze-in production of FIMP dark matter. *Journal of High Energy Physics*, 2010(3):80, 2010.
- [54] Yonit Hochberg, Eric Kuflik, Tomer Volansky, and Jay G. Wacker. Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles. *Phys. Rev. Lett.*, 113:171301, October 2014.
- [55] Anatoly Klypin, Andrey V. Kravtsov, Octavio Valenzuela, and Francisco Prada. Where Are the Missing Galactic Satellites? *The Astrophysical Journal*, 522(1):82–92, September 1999.
- [56] Ben Moore, Sebastiano Ghigna, Fabio Governato, et al. Dark Matter Substructure within Galactic Halos. *The Astrophysical Journal*, 524(1):L19–L22, October 1999.
- [57] WJG De Blok. The core-cusp problem. *Advances in Astronomy*, 2010.
- [58] David N. Spergel and Paul J. Steinhardt. Observational Evidence for Self-Interacting Cold Dark Matter. *Phys. Rev. Lett.*, 84:3760–3763, April 2000.
- [59] Michael Boylan-Kolchin, James S. Bullock, and Manoj Kaplinghat. Too big to fail? The puzzling darkness of massive Milky Way subhaloes. *Monthly Notices of the Royal Astronomical Society: Letters*, 415(1):L40–L44, July 2011.

- [60] R. Adhikari, M. Agostini, N. Anh Ky, T. Araki, et al. A White Paper on keV sterile neutrino Dark Matter. *Journal of Cosmology and Astroparticle Physics*, 2017(01):025–025, January 2017.
- [61] Scott Tremaine and James E. Gunn. Dynamical Role of Light Neutral Leptons in Cosmology. *Phys. Rev. Lett.*, 42:407–410, February 1979.
- [62] John Preskill, Mark B Wise, and Frank Wilczek. Cosmology of the invisible axion. *Physics Letters B*, 120(1-3):127–132, 1983.
- [63] Laurence F Abbott and P Sikivie. A cosmological bound on the invisible axion. *Physics Letters B*, 120(1-3):133–136, 1983.
- [64] Michael Dine and Willy Fischler. The not-so-harmless axion. *Physics Letters B*, 120(1-3):137–141, 1983.
- [65] Roberto D Peccei and Helen R Quinn. CP Conservation in the Presence of Instantons. *Phys. Rev. Lett*, 38(1440-1443):328, 1977.
- [66] Steven Weinberg. A new light boson? *Physical Review Letters*, 40(4):223, 1978.
- [67] Frank Wilczek. Problem of Strong p and t Invariance in the Presence of Instantons. *Physical Review Letters*, 40(5):279, 1978.
- [68] Andreas Ringwald. Exploring the role of axions and other WISPs in the dark universe. *Physics of the Dark Universe*, 1(1):116 – 135, 2012. Next Decade in Dark Matter and Dark Energy.
- [69] Ann E Nelson and Jakub Scholtz. Dark light, dark matter, and the misalignment mechanism. *Physical Review D*, 84(10):103501, 2011.
- [70] Paola Arias, Davide Cadamuro, Mark Goodsell, et al. WISPy cold dark matter. *Journal of Cosmology and Astroparticle Physics*, 2012(06):013–013, June 2012.
- [71] M. Aaboud, G. Aad, B. Abbott, et al. Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp-collisions at $\sqrt{s}=13$ TeV using the ATLAS detector. *Physical Review D*, 94(3), August 2016.
- [72] Felix Kahlhoefer. Review of LHC dark matter searches. *International Journal of Modern Physics A*, 32(13):1730006, 2017.
- [73] Maria Beltrán, Dan Hooper, Edward W. Kolb, Zosia A. C. Krusberg, and Tim M. P. Tait. Maverick dark matter at colliders. *Journal of High Energy Physics*, 2010(9), September 2010.
- [74] Jessica Goodman, Masahiro Ibe, Arvind Rajaraman, et al. Constraints on dark matter from colliders. *Physical Review D*, 82(11), December 2010.
- [75] Yang Bai, Patrick J. Fox, and Roni Harnik. The Tevatron at the frontier of dark matter direct detection. *Journal of High Energy Physics*, 2010(12), December 2010.

- [76] O. Buchmueller, Matthew J. Dolan, and Christopher McCabe. Beyond effective field theory for dark matter searches at the LHC. *Journal of High Energy Physics*, 2014(1), January 2014.
- [77] Jalal Abdallah, Adi Ashkenazi, Antonio Boveia, et al. Simplified Models for Dark Matter and Missing Energy Searches at the LHC, 2014.
- [78] Jalal Abdallah, Henrique Araujo, Alexandre Arbey, et al. Simplified models for dark matter searches at the LHC. *Physics of the Dark Universe*, 9-10:8–23, September 2015.
- [79] Daniel Abercrombie, Nural Akchurin, Ece Akilli, et al. Dark Matter benchmark models for early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum. *Physics of the Dark Universe*, 27:100371, January 2020.
- [80] CMS Collaboration et al. Search for dark matter produced with an energetic jet, or a hadronically decaying W or Z boson using 12.9 fb⁻¹ data at $\sqrt{s}=13$ TeV. Technical report, Technical Report, CERN, Geneva, 2016. CMS-PAS-EXO-16-037.
- [81] Jennifer M. Gaskins. A review of indirect searches for particle dark matter. *Contemporary Physics*, 57(4):496–525, 2016.
- [82] Daniel J. H. Chung, Edward W. Kolb, and Antonio Riotto. Superheavy dark matter. *Phys. Rev. D*, 59:023501, November 1998.
- [83] WB Atwood, Aous A Abdo, Markus Ackermann, et al. The large area telescope on the Fermi gamma-ray space telescope mission. *The Astrophysical Journal*, 697(2):1071, 2009.
- [84] BS Acharya, M Actis, T Aghajani, et al. Introducing the CTA concept. *Astroparticle Physics*, 43:3–18, 2013.
- [85] Tansu Daylan, Douglas P Finkbeiner, Dan Hooper, et al. The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter. *Physics of the Dark Universe*, 12:1–23, 2016.
- [86] Oscar Adriani, GC Barbarino, GA Bazilevskaya, et al. An anomalous positron abundance in cosmic rays with energies 1.5–100 GeV. *Nature*, 458(7238):607–609, 2009.
- [87] Stefano Profumo. Dissecting cosmic-ray electron-positron data with Occam’s Razor: the role of known Pulsars. *Open Physics*, 10(1):1–31, 2012.
- [88] Mark W. Goodman and Edward Witten. Detectability of Certain Dark Matter Candidates. *Phys. Rev.*, D31:3059, 1985. [325(1984)].
- [89] S. P. Ahlen, F. T. Avignone, R. L. Brodzinski, et al. Limits on Cold Dark Matter Candidates from an Ultralow Background Germanium Spectrometer. *Phys. Lett.*, B195:603–608, 1987.
- [90] F. Donato, N. Fornengo, and S. Scopel. Effects of galactic dark halo rotation on WIMP direct detection. *Astroparticle Physics*, 9:247–260, October 1998.

-
- [91] A. Liam Fitzpatrick, Wick Haxton, Emanuel Katz, Nicholas Lubbers, and Yiming Xu. The effective field theory of dark matter direct detection. *Journal of Cosmology and Astroparticle Physics*, 2013(02):004–004, February 2013.
- [92] A.H. Abdelhameed et al. First results from the CRESST-III low-mass dark matter program. *Phys. Rev. D*, 100(10):102002, 2019.
- [93] A. H. Abdelhameed, G. Angloher, P. Bauer, et al. Cryogenic characterization of a LiAlO₂ crystal and new results on spin-dependent dark matter interactions with ordinary matter. *The European Physical Journal C*, 80(9), September 2020.
- [94] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2):1554, February 2011.
- [95] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D*, 66:032005, August 2002.
- [96] J. Billard, L. Strigari, and E. Figueroa-Feliciano. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Phys. Rev. D*, 89(2):023524, 2014.
- [97] A. Gütlein et al. Impact of coherent neutrino nucleus scattering on direct dark matter searches based on CaWO₄ crystals. *Astropart. Phys.*, 69:44–49, 2015.
- [98] Francis Froberg and Alan R. Duffy. Annual Modulation in Direct Dark Matter Searches. *J. Phys. G*, 47(9):094002, 2020.
- [99] Dominik Jeschke. *Modulations of the Cosmic Muon Flux & Identification of Atmospheric Neutrino Interactions in Borexino*. Dissertation, Technische Universität München, München, 2018.
- [100] G. Bellini et al. Final results of Borexino Phase-I on low energy solar neutrino spectroscopy. *Phys. Rev. D*, 89(11):112007, 2014.
- [101] R. Bernabei et al. First model independent results from DAMA/LIBRA-phase2. *Nucl. Phys. Atom. Energy*, 19(4):307–325, 2018.
- [102] F. Mayet et al. A review of the discovery reach of directional Dark Matter detection. *Phys. Rept.*, 627:1–49, 2016.
- [103] H. Jiang et al. Limits on Light Weakly Interacting Massive Particles from the First 102.8 kg × day Data of the CDEX-10 Experiment. *Phys. Rev. Lett.*, 120(24):241301, 2018.
- [104] W. Zhao, Q. Yue, K. J. Kang, et al. Search of low-mass WIMPs with a *p*-type point contact germanium detector in the CDEX-1 experiment. *Phys. Rev. D*, 93:092003, May 2016.
- [105] A. Aguilar-Arevalo et al. Results on low-mass weakly interacting massive particles from a 11 kg-day target exposure of DAMIC at SNOLAB. 7 2020.

- [106] Liron Barak et al. SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper-CCD. 4 2020.
- [107] Ning Wang, F. C. Wellstood, B. Sadoulet, E. E. Haller, and J. Beeman. Electrical and thermal properties of neutron-transmutation-doped Ge at 20 mK. *Phys. Rev. B*, 41:3761–3768, February 1990.
- [108] K.D. Irwin, S.W. Nam, B. Cabrera, B. Chugg, and B.A. Young. A quasiparticle-trap-assisted transition-edge sensor for phonon-mediated particle detection. *Rev. Sci. Instrum.*, 66:5322, 1995.
- [109] F. Pröbst, M. Frank, S. Cooper, et al. Model for cryogenic particle detectors with superconducting phase transition thermometers. *Journal of Low Temperature Physics*, 100:69–104, July 1995.
- [110] E. Armengaud et al. Searches for electron interactions induced by new physics in the EDELWEISS-III Germanium bolometers. *Phys. Rev. D*, 98(8):082004, 2018.
- [111] W. Rau. SuperCDMS SNOLAB - Status and Plans. *J. Phys. Conf. Ser.*, 1342(1):012077, 2020.
- [112] E. Aprile et al. Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys. Rev. Lett.*, 121(11):111302, 2018.
- [113] E. Aprile et al. Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment. 7 2020.
- [114] D.S. Akerib et al. Projected WIMP sensitivity of the LUX-ZEPLIN dark matter experiment. *Phys. Rev. D*, 101(5):052002, 2020.
- [115] P. Agnes, I. F. M. Albuquerque, T. Alexander, et al. Darkside-50 532-day dark matter search with low-radioactivity argon. *Phys. Rev. D*, 98:102006, November 2018.
- [116] C.E. Aalseth et al. DarkSide-20k: A 20 tonne two-phase LAr TPC for direct dark matter detection at LNGS. *Eur. Phys. J. Plus*, 133:131, 2018.
- [117] C. Amole et al. Dark Matter Search Results from the PICO-60 C₃F₈ Bubble Chamber. *Phys. Rev. Lett.*, 118(25):251301, 2017.
- [118] J.B.R. Battat et al. Low Threshold Results and Limits from the DRIFT Directional Dark Matter Detector. *Astropart. Phys.*, 91:65–74, 2017.
- [119] Bill Pope, Anne Druyan, Steven Soter, et al. Deeper, Deeper, Deeper Still. *National Geographic Channel, RGB Media, Cosmos: A Spacetime Odyssey*, Episode 6, April 2014.
- [120] D. Griffiths. *Introduction to Elementary Particles*. Physics textbook. Wiley, 2008.
- [121] M. G. Aartsen et al. Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption. *Nature*, 551:596–600, 2017.

- [122] D Z Freedman, D N Schramm, and D L Tubbs. The Weak Neutral Current and its Effects in Stellar Collapse. *Annual Review of Nuclear Science*, 27(1):167–207, 1977.
- [123] Madappa Prakash, James M. Lattimer, Raymond F. Sawyer, and Raymond R. Volkas. Neutrino propagation in dense astrophysical systems. *Annual Review of Nuclear and Particle Science*, 51(1):295–344, 2001.
- [124] Daniel Z. Freedman. Coherent effects of a weak neutral current. *Phys. Rev. D*, 9:1389–1392, March 1974.
- [125] F.J. Hasert, S. Kabe, W. Krenz, et al. Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment. *Physics Letters B*, 46(1):138 – 140, 1973.
- [126] F.J. Hasert, H. Faissner, W. Krenz, et al. Search for elastic muon-neutrino electron scattering. *Physics Letters B*, 46(1):121 – 124, 1973.
- [127] D. Akimov et al. Observation of Coherent Elastic Neutrino-Nucleus Scattering. *Science*, 357(6356):1123–1126, 2017.
- [128] Juan Barranco, Omar G Miranda, and Timur I Rashba. Probing new physics with coherent neutrino scattering off nuclei. *Journal of High Energy Physics*, 2005(12):021–021, December 2005.
- [129] M. Tanabashi, K. Hagiwara, K. Hikasa, et al. Review of Particle Physics. *Phys. Rev. D*, 98:030001, August 2018.
- [130] Richard H. Helm. Inelastic and Elastic Scattering of 187-MeV Electrons from Selected Even-Even Nuclei. *Phys. Rev.*, 104:1466–1475, December 1956.
- [131] P. Vilain, G. Wilquet, R. Beyer, et al. Precision measurement of electroweak parameters from the scattering of muon-neutrinos on electrons. *Physics Letters B*, 335(2):246 – 252, 1994.
- [132] J.V. Allaby et al. A Precise Determination of the Electroweak Mixing Angle from Semileptonic Neutrino Scattering. *Z. Phys. C*, 36:611, 1987.
- [133] G. P. Zeller, K. S. McFarland, T. Adams, et al. Precise Determination of Electroweak Parameters in Neutrino-Nucleon Scattering. *Phys. Rev. Lett.*, 88:091802, February 2002.
- [134] P.L. Anthony et al. Precision measurement of the weak mixing angle in Møller scattering. *Phys. Rev. Lett.*, 95:081601, 2005.
- [135] Michael T. Gericke. The Q_{weak}^p experiment: A test for physics beyond the standard model via a precision measurement of the proton weak charge. *AIP Conf. Proc.*, 1149(1):237–240, 2009.
- [136] CS Wood, SC Bennett, Donghyun Cho, et al. Measurement of parity nonconservation and an anapole moment in cesium. *Science*, 275(5307):1759–1763, 1997.

- [137] Jens Erler and Rodolfo Ferro-Hernández. Weak mixing angle in the Thomson limit. *Journal of High Energy Physics*, 2018(3):196, March 2018.
- [138] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, May 1978.
- [139] O. G. Miranda, M. A. Tortola, and J. W. F. Valle. Are solar neutrino oscillations robust? *JHEP*, 10:008, 2006.
- [140] Sacha Davidson, Carlos Peña-Garay, Nuria Rius, and Arcadi Santamaria. Present and future bounds on non-standard neutrino interactions. *Journal of High Energy Physics*, 2003(03):011–011, March 2003.
- [141] Alexander Friedland, Michael L. Graesser, Ian M. Shoemaker, and Luca Vecchi. Probing nonstandard standard model backgrounds with LHC monojets. *Physics Letters B*, 714(2):267 – 275, 2012.
- [142] Tommy Ohlsson. Status of non-standard neutrino interactions. *Reports on Progress in Physics*, 76(4):044201, March 2013.
- [143] Yasaman Farzan and Mariam Tórtola. Neutrino oscillations and non-standard interactions. *Frontiers in Physics*, 6:10, 2018.
- [144] J. Dorenbosch, J.V. Allaby, U. Amaldi, et al. Experimental verification of the universality of electron-neutrino and muon-neutrino coupling to the neutral weak current. *Physics Letters B*, 180(3):303 – 307, 1986.
- [145] M. Cadeddu, C. Giunti, K. A. Kouzakov, et al. Neutrino Charge Radii from COHERENT Elastic Neutrino-Nucleus Scattering. *Phys. Rev.*, D98(11):113010, 2018.
- [146] Carlo Giunti and Alexander Studenikin. Neutrino electromagnetic interactions: A window to new physics. *Rev. Mod. Phys.*, 87:531–591, June 2015.
- [147] Nicole F. Bell, Mikhail Gorchtein, Michael J. Ramsey-Musolf, Petr Vogel, and Peng Wang. Model independent bounds on magnetic moments of Majorana neutrinos. *Physics Letters B*, 642(4):377 – 383, 2006.
- [148] G. Bressi, G. Carugno, F. Della Valle, et al. Testing the neutrality of matter by acoustic means in a spherical resonator. *Phys. Rev. A*, 83:052101, May 2011.
- [149] J. Bernabéu, J. Papavassiliou, and J. Vidal. The neutrino charge radius as a physical observable. *Nuclear Physics B*, 680(1):450 – 478, 2004.
- [150] M. Cadeddu, F. Dordei, C. Giunti, Y. F. Li, and Y. Y. Zhang. Neutrino, Electroweak and Nuclear Physics from COHERENT Elastic Neutrino-Nucleus Scattering with a New Quenching Factor. 2019.
- [151] Georg G Raffelt. Limits on neutrino electromagnetic properties — an update. *Physics Reports*, 320(1):319 – 327, 1999.

- [152] M. Deniz, S. T. Lin, V. Singh, et al. Measurement of $\bar{\nu}_e$ -electron scattering cross section with a CsI(Tl) scintillating crystal array at the Kuo-Sheng nuclear power reactor. *Phys. Rev. D*, 81:072001, April 2010.
- [153] M. Cadeddu, C. Giunti, K. A. Kouzakov, et al. Neutrino charge radii from COHERENT elastic neutrino-nucleus scattering. *Phys. Rev. D*, 98:113010, December 2018.
- [154] Dimitrios K. Papoulias. COHERENT constraints after the Chicago-3 quenching factor measurement, 2019.
- [155] Amir N. Khan and Werner Rodejohann. New physics from COHERENT data with an improved quenching factor. *Phys. Rev. D*, 100:113003, December 2019.
- [156] Konstantin A. Kouzakov and Alexander I. Studenikin. Electromagnetic properties of massive neutrinos in low-energy elastic neutrino-electron scattering. *Phys. Rev. D*, 95:055013, March 2017.
- [157] A. G. Beda, V. B. Brudanin, V. G. Egorov, et al. The results of search for the neutrino magnetic moment in GEMMA experiment. *Adv. High Energy Phys.*, 2012:350150, 2012.
- [158] A. G. Beda, V. B. Brudanin, V. G. Egorov, et al. Gemma experiment: The results of neutrino magnetic moment search. *Phys. Part. Nucl. Lett.*, 10:139–143, 2013.
- [159] P. Vogel and J. Engel. Neutrino electromagnetic form factors. *Phys. Rev. D*, 39:3378–3383, June 1989.
- [160] M. Deniz, S. Bilmiş, İ. O. Yildırım, et al. Constraints on nonstandard neutrino interactions and unparticle physics with $\bar{\nu}_e - e^-$ scattering at the kuo-sheng nuclear power reactor. *Phys. Rev. D*, 82:033004, August 2010.
- [161] Jalal Abdallah et al. Simplified Models for Dark Matter Searches at the LHC. *Phys. Dark Univ.*, 9-10:8–23, 2015.
- [162] David G. Cerdeño, Malcolm Fairbairn, Thomas Jubb, et al. Physics from solar neutrinos in dark matter direct detection experiments. *JHEP*, 05:118, 2016. [Erratum: *JHEP*09,048(2016)].
- [163] Yasaman Farzan, Manfred Lindner, Werner Rodejohann, and Xun-Jie Xu. Probing neutrino coupling to a light scalar with coherent neutrino scattering. *JHEP*, 05:066, 2018.
- [164] D. Aristizabal Sierra, Bhaskar Dutta, Shu Liao, and Louis E. Strigari. Coherent elastic neutrino-nucleus scattering in multi-ton scale dark matter experiments: Classification of vector and scalar interactions new physics signals. *JHEP*, 12:124, 2019.
- [165] Julien Billard, Joseph Johnston, and Bradley J. Kavanagh. Prospects for exploring New Physics in Coherent Elastic Neutrino-Nucleus Scattering. *JCAP*, 1811(11):016, 2018.

- [166] L Michel. Interaction between Four Half-Spin Particles and the Decay of the μ -Meson. *Proceedings of the Physical Society. Section A*, 63(5):514–531, May 1950.
- [167] Kate Scholberg. Prospects for measuring coherent neutrino-nucleus elastic scattering at a stopped-pion neutrino source. *Phys. Rev. D*, 73:033005, February 2006.
- [168] D. Baxter et al. Coherent Elastic Neutrino-Nucleus Scattering at the European Spallation Source. *JHEP*, 02:123, 2020.
- [169] V. Kopeikin, L. Mikaelyan, and V. Sinev. Reactor as a source of antineutrinos: Thermal fission energy. *Phys. Atom. Nucl.*, 67:1892–1899, 2004. [*Yad. Fiz.*67,1916(2004)].
- [170] V. I. Kopeikin. Flux and spectrum of reactor antineutrinos. *Physics of Atomic Nuclei*, 75(2):143–152, February 2012.
- [171] W Hampel, G Heusser, J Kiko, et al. Final results of the ^{51}Cr neutrino source experiments in GALLEX. *Physics Letters B*, 420(1):114 – 126, 1998.
- [172] G. Bellini, D. Bick, G. Bonfini, et al. SOX: Short distance neutrino Oscillations with BoreXino. *Journal of High Energy Physics*, 2013(8):38, August 2013.
- [173] C. Bellenghi, D. Chiesa, L. Di Noto, et al. Coherent elastic nuclear scattering of ^{51}Cr neutrinos. *Eur. Phys. J.*, C79(9):727, 2019.
- [174] A. Drukier and Leo Stodolsky. Principles and Applications of a Neutral Current Detector for Neutrino Physics and Astronomy. *Phys. Rev.*, D30:2295, 1984. [,395(1984)].
- [175] D. Akimov, J. B. Albert, P. An, et al. First constraint on coherent elastic neutrino-nucleus scattering in argon. *Phys. Rev. D*, 100:115020, December 2019.
- [176] Jacob Zettlemoyer. New Results from a CEvNS Search with the CENNS-10 Liquid Argon Detector (Fermilab Joint Experimental-Theoretical Physics Seminar). <https://vms.fnal.gov/asset/detail?recid=1962782>, 2020. [Online; accessed March 17, 2020].
- [177] D. Akimov et al. First Detection of Coherent Elastic Neutrino-Nucleus Scattering on Argon. arXiv:2003.10630, March 2020.
- [178] Henry Tsz-King Wong. Taiwan EXperiment On Neutrino — History and Prospects. *The Universe*, 3(4):22–37, 2015.
- [179] H. T. Wong, H. B. Li, S. T. Lin, et al. Search of neutrino magnetic moments with a high-purity germanium detector at the Kuo-Sheng nuclear power station. *Phys. Rev. D*, 75:012001, January 2007.
- [180] Henry Wong. Neutrino-nucleus Coherent Scattering with Reactor and Solar Neutrinos. 10.5281/zenodo.1287671, June 2018.
- [181] C. Buck et al. A novel experiment for coherent elastic neutrino nucleus scattering: CONUS. *J. Phys. Conf. Ser.*, 1342(1):012094, 2020.

- [182] G. Heusser, M. Weber, J. Hakenmüller, et al. GIOVE: a new detector setup for high sensitivity germanium spectroscopy at shallow depth. *The European Physical Journal C*, 75(11):531, November 2015.
- [183] J. Hakenmüller et al. Neutron-induced background in the CONUS experiment. *Eur. Phys. J.*, C79(8):699, 2019.
- [184] Werner Maneschg. The Status of CONUS. 10.5281/zenodo.1286927, June 2018.
- [185] Alexis Aguilar-Arevalo et al. Exploring low-energy neutrino physics with the Coherent Neutrino Nucleus Interaction Experiment. *Phys. Rev.*, D100(9):092005, 2019.
- [186] Alexis Aguilar-Arevalo et al. Search for light mediators in the low-energy data of the CONNIE reactor neutrino experiment. *JHEP*, 04:054, 2020.
- [187] MINER Collaboration, G. Agnolet, W. Baker, et al. Background Studies for the MINER Coherent Neutrino Scattering Reactor Experiment, 2016.
- [188] Bhaskar Dutta, Yu Gao, Rupak Mahapatra, et al. Sensitivity to oscillation with a sterile fourth generation neutrino from ultra-low threshold neutrino-nucleus coherent scattering. *Phys. Rev.*, D94(9):093002, 2016.
- [189] I. Alekseev et al. Neutrino Physics at Kalinin Nuclear Power Plant: 2002 – 2017. *J. Phys. Conf. Ser.*, 934(1):012006, 2017.
- [190] V. Belov et al. The ν GeN experiment at the Kalinin Nuclear Power Plant. *JINST*, 10(12):P12011, 2015.
- [191] D. Yu. Akimov et al. Status of the RED-100 experiment. *JINST*, 12(06):C06018, 2017.
- [192] D. Yu. Akimov et al. Prospects for observation of neutrino-nuclear neutral current coherent scattering with two-phase Xenon emission detector. *JINST*, 8:P10023, 2013.
- [193] D. Yu Akimov et al. First ground-level laboratory test of the two-phase xenon emission detector RED-100. *JINST*, 15(02):P02020, 2020.
- [194] Joseph A. Formaggio, E. Figueroa-Feliciano, and A. J. Anderson. Sterile Neutrinos, Coherent Scattering and Oscillometry Measurements with Low-temperature Bolometers. *Phys. Rev.*, D85:013009, 2012.
- [195] A. J. Anderson, J. Conrad, E. Figueroa-Feliciano, et al. Coherent Neutrino Scattering with Cryogenic Semiconductor Detectors. In *2012 Electroweak Interactions and Unified Theories: Proceedings of the 47th Rencontres de Moriond on Electroweak Interactions and Unified Theories, La Thuile, March 3-10, 2012*, pages 425–429, 2012.
- [196] J. Billard et al. Coherent Neutrino Scattering with Low Temperature Bolometers at Chooz Reactor Complex. *J. Phys.*, G44(10):105101, 2017.

- [197] A. Juillard et al. Low-noise HEMTs for Coherent Elastic Neutrino Scattering and Low-Mass Dark Matter Cryogenic Semiconductor Detectors. In *18th International Workshop on Low Temperature Detectors (LTD 18) Milano, Italy, July 22-26, 2019*, 2019.
- [198] Neil W Ashcroft, N David Mermin, et al. *Solid state physics*. New York: Holt, Rinehart and Winston,, 1976.
- [199] R Viswanathan. Heat capacity of sapphire between 2 and 10 K by ac technique. *Journal of Applied Physics*, 46(9):4086–4087, 1975.
- [200] CUORE – A search for neutrinoless double beta decay: Images of CUORE. <https://cuore.lngs.infn.it/en/images>. [Online; accessed September 28, 2020].
- [201] Marco Vignati. *Model of the Response Function of the CUORE Bolometers*. Dissertation, Rome U., April 2011.
- [202] Allen Miller and Elihu Abrahams. Impurity conduction at low concentrations. *Physical Review*, 120(3):745, 1960.
- [203] Kohei M Itoh, EE Haller, WL Hansen, et al. Neutron transmutation doping of isotopically engineered Ge. *Applied physics letters*, 64(16):2121–2123, 1994.
- [204] C. Alduino et al. First results from the CUORE experiment. *J. Phys. Conf. Ser.*, 1342(1):012002, 2020.
- [205] Michelle J. Dolinski, Alan W.P. Poon, and Werner Rodejohann. Neutrinoless Double-Beta Decay: Status and Prospects. *Ann. Rev. Nucl. Part. Sci.*, 69:219–251, 2019.
- [206] S Kempf, A Fleischmann, L Gastaldo, and C Enss. Physics and applications of metallic magnetic calorimeters. *Journal of Low Temperature Physics*, 193(3-4):365–379, 2018.
- [207] A Fleischmann, L Gastaldo, S Kempf, et al. Metallic magnetic calorimeters. In *AIP Conference Proceedings*, volume 1185, pages 571–578. American Institute of Physics, 2009.
- [208] Loredana Gastaldo, Klaus Blaum, K Chrysalidis, et al. The electron capture in ^{163}Ho experiment–ECHO. *The European Physical Journal Special Topics*, 226(8):1623–1694, 2017.
- [209] HS Jo, S Choi, FA Danevich, et al. Status of the AMoRE experiment searching for neutrinoless double beta decay using low-temperature detectors. *Journal of Low Temperature Physics*, 193(5-6):1182–1189, 2018.
- [210] L. Cardani, I. Colantoni, A. Cruciani, et al. Energy resolution and efficiency of phonon-mediated kinetic inductance detectors for light detection. *Applied Physics Letters*, 107(9):093508, 2015.
- [211] L. Cardani, N. Casali, I. Colantoni, et al. High sensitivity phonon-mediated kinetic inductance detector with combined amplitude and phase read-out. *Applied Physics Letters*, 110(3):033504, 2017.

- [212] Philip Daniel Mauskopf, S Doyle, Peter Barry, et al. Photon-noise limited performance in aluminum LEKIDs. *Journal of Low Temperature Physics*, 176(3-4):545–552, 2014.
- [213] BA Mazin, PK Day, J Zmuidzinas, and HG LeDuc. Multiplexable kinetic inductance detectors. In *AIP Conference Proceedings*, volume 605, pages 309–312. American Institute of Physics, 2002.
- [214] Simon Doyle, Philip Mauskopf, J Naylon, Adrian Porch, and C Duncombe. Lumped element kinetic inductance detectors. *Journal of Low Temperature Physics*, 151(1-2):530–536, 2008.
- [215] Seth R Meeker, Benjamin A Mazin, Alex B Walter, et al. DARKNESS: a microwave kinetic inductance detector integral field spectrograph for high-contrast astronomy. *Publications of the Astronomical Society of the Pacific*, 130(988):065001, 2018.
- [216] L Cardani, N Casali, A Cruciani, et al. Al/Ti/Al phonon-mediated KIDs for UV–vis light detection over large areas. *Superconductor Science and Technology*, 31(7):075002, May 2018.
- [217] Kent D Irwin and Gene C Hilton. Transition-edge sensors. In *Cryogenic particle detection*, pages 63–150. Springer, 2005.
- [218] Joel N Ullom and Douglas A Bennett. Review of superconducting transition-edge sensors for x-ray and gamma-ray spectroscopy. *Superconductor Science and Technology*, 28(8):084003, July 2015.
- [219] DJ Cousins, Shaun N Fisher, AM Guénault, et al. An advanced dilution refrigerator designed for the new Lancaster microkelvin facility. *Journal of low temperature physics*, 114(5-6):547–570, 1999.
- [220] A.H. Abdelhameed et al. Deposition of Tungsten Thin Films by Magnetron Sputtering for Large-Scale Production of Tungsten-Based Transition-Edge Sensors. *J. Low Temp. Phys.*, 199(1-2):401–407, 2020.
- [221] Johannes Rothe. *Achieving Low Thresholds: Cryogenic Detectors for Low-Mass Dark Matter Searches*. Master thesis, Ludwig-Maximilians-Universität München, München, 2016.
- [222] S Henry, N Bazin, H Kraus, et al. The 66-channel SQUID readout for CRESST II. *Journal of Instrumentation*, 2(11):P11003–P11003, November 2007.
- [223] Michael Tinkham. *Introduction to Superconductivity (2nd Edition)*. Dover Publications, 1996.
- [224] R. Strauss, J. Rothe, G. Angloher, et al. Gram-scale cryogenic calorimeters for rare-event searches. *Physical Review D*, 96(2):022009, 2017.
- [225] O. Meier et al. Active thermal feedback for massive cryogenic detectors. *Nucl. Instrum. Meth.*, A444:350–352, 2000.

- [226] G. Angloher et al. Results on low mass WIMPs using an upgraded CRESST-II detector. *Eur. Phys. J.*, C74(12):3184, 2014.
- [227] G. Angloher et al. Results on light dark matter particles with a low-threshold CRESST-II detector. *Eur. Phys. J.*, C76(1):25, 2016.
- [228] M. Mancuso et al. A Low Nuclear Recoil Energy Threshold for Dark Matter Search with CRESST-III Detectors. *J. Low. Temp. Phys.*, 193(3-4):441–448, 2018.
- [229] G. Angloher, M. Bruckmayer, C. Bucci, et al. Limits on WIMP dark matter using sapphire cryogenic detectors. *Astroparticle Physics*, 18:43–55, August 2002.
- [230] D.-M. Mei and A. Hime. Muon-induced background study for underground laboratories. *Phys. Rev. D*, 73:053004, March 2006.
- [231] Anja Tanzke. *Low-Threshold Detectors for Low-Mass Direct Dark Matter Search with CRESST-III*. Dissertation, Technische Universität München, München, 2017.
- [232] G. Bellini, J. Benziger, D. Bick, et al. Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 5:015, May 2012.
- [233] R. Strauss, G. Angloher, A. Bento, et al. Beta/gamma and alpha backgrounds in CRESST-II Phase 2. *Journal of Cosmology and Astroparticle Physics*, 2015(06):030, 2015.
- [234] A. H. Abdelhameed et al. Geant4-based electromagnetic background model for the CRESST dark matter experiment. *Eur. Phys. J.*, C79(10):881, 2019.
- [235] Sabine Roth. *The Potential of Neganov-Luke Amplified Cryogenic Light Detectors and the Scintillation-Light Quenching Mechanism in CaWO_4 Single Crystals in the Context of the Dark Matter Search Experiment CRESST-II*. Dissertation, Technische Universität München, München, 2013.
- [236] Münster, Andrea Ruth. *High-Purity CaWO_4 Single Crystals for Direct Dark Matter Search with the CRESST Experiment*. Dissertation, Technische Universität München, München, 2017.
- [237] J. Rothe, G. Angloher, P. Bauer, et al. TES-Based Light Detectors for the CRESST Direct Dark Matter Search. *Journal of Low Temperature Physics*, 193(5-6):1160–1166, 2018.
- [238] Marc Wüstrich. *Improving Particle Discrimination and Achieving a 4-Veto Detector Concept for the CRESST Experiment*. Dissertation, Technische Universität München, München, 2020.
- [239] G. Angloher et al. Performance of a CRESST-II Detector Module with True 4π -veto. arXiv:1708.01581, August 2017.
- [240] G. Angloher, M. Bauer, I. Bavykina, et al. Results from 730 kg days of the CRESST-II Dark Matter search. *European Physical Journal C*, 72:1971, April 2012.

- [241] M. Kuźniak, M.G. Boulay, and T. Pollmann. Surface roughness interpretation of 730 kg days CRESST-II results. *Astropart. Phys.*, 36:77–82, 2012.
- [242] Andreas Josef Zöller. *Artificial Neural Network Based Pulse-Shape Analysis for Cryogenic Detectors Operated in CRESST-II*. Dissertation, Technische Universität München, München, 2016.
- [243] Nahuel Ferreiro Iachellini. *Increasing the sensitivity to low mass dark matter in CRESST-III with a new DAQ and signal processing*. Dissertation, Ludwig-Maximilians-Universität München, München, February 2019.
- [244] Philipp Martin Michael Bauer. *Data Analysis for the CRESST Experiment: New Methods, improved Alpha Analysis, and Results on Light Dark Matter and Backgrounds*. Dissertation, Technische Universität München, München, 2020.
- [245] Martin Stahlberg. *Probing Low-Mass Dark Matter with CRESST-III - Data Analysis and First Results*. Dissertation, Technische Universität Wien, Vienna, 2020.
- [246] H. Kamerlingh Onnes. *Further experiments with Liquid Helium. G. On the Electrical Resistance of Pure Metals, etc. VI. On the Sudden Change in the Rate at which the Resistance of Mercury Disappears.*, pages 267–272. Springer Netherlands, Dordrecht, 1991.
- [247] W. Meissner and R. Ochsenfeld. Ein neuer Effekt bei Eintritt der Supraleitfähigkeit. *Naturwissenschaften*, 21(44):787–788, November 1933.
- [248] J. W. P. Hsu and A. Kapitulnik. Superconducting transition, fluctuation, and vortex motion in a two-dimensional single-crystal Nb film. *Phys. Rev. B*, 45:4819–4835, March 1992.
- [249] K.D. Irwin and G.C. Hilton. *Transition-Edge Sensors*, pages 63–150. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- [250] J.N. Ullom, W.B. Doriese, G.C. Hilton, et al. Suppression of excess noise in transition-edge sensors using magnetic field and geometry. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 520(1):333 – 335, 2004. Proceedings of the 10th International Workshop on Low Temperature Detectors.
- [251] Antoine R. Miniussi, Joseph S. Adams, Simon R. Bandler, et al. Design of Magnetic Shielding and Field Coils for a TES X-ray Microcalorimeter Test Platform. *Journal of Low Temperature Physics*, 194(5):433–442, March 2019.
- [252] Arnaud Chulliat, William Brown, Patrick Alken, et al. Out-of-Cycle Update of the US/UK World Magnetic Model for 2015-2020. Technical Report <https://doi.org/10.25921/xhr3-0t19>, 2019. This item has been internally reviewed, but not externally peer-reviewed.
- [253] NCEI Geomagnetic Calculators. <https://www.ngdc.noaa.gov/geomag/calculators/magcalc.shtml>. [Online; accessed July 31, 2020].

- [254] Efthimios Kaxiras. *Atomic and electronic structure of solids*. Cambridge Univ. Press, 2003.
- [255] R. C. Dynes and T. A. Fulton. Supercurrent Density Distribution in Josephson Junctions. *Phys. Rev. B*, 3:3015–3023, May 1971.
- [256] Sean Hart, Hechen Ren, Timo Wagner, et al. Induced superconductivity in the quantum spin Hall edge. *Nature Physics*, 10:1745–2481, September 2014.
- [257] A. D’Addabbo et al. The CUORE cryostat. *J. Low Temp. Phys.*, 193(5-6):867–875, 2018.
- [258] OpenStreetMap. Node CRESST: <https://www.openstreetmap.org/node/2739238931>, Node CUORE: <https://www.openstreetmap.org/node/2739238932>. [Online; accessed August 29, 2019].
- [259] J R Barker. New coil systems for the production of uniform magnetic fields. *Journal of Scientific Instruments*, 26(8):273–275, August 1949.
- [260] Emilija Pantic. *Performance of Cryogenic Light Detectors in the CRESST-II Dark Matter Search*. Dissertation, Technische Universität München, München, 2008.
- [261] B. B. Triplett, N. E. Phillips, T. L. Thorp, D. A. Shirley, and W. D. Brewer. Critical field for superconductivity and low-temperature normal-state heat capacity of tungsten. *Journal of Low Temperature Physics*, 12(5):499–518, 1973.
- [262] R. Strauss, J. Rothe, G. Angloher, et al. The ν -cleus experiment: a gram-scale fiducial-volume cryogenic detector for the first detection of coherent neutrino–nucleus scattering. *The European Physical Journal C*, 77(8):506, 2017.
- [263] G. Angloher, P. Bauer, A. Bento, et al. Results on MeV-scale dark matter from a gram-scale cryogenic calorimeter operated above ground. *The European Physical Journal C*, 77(9):637, 2017.
- [264] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch. Determination of the antineutrino spectrum from ^{235}U thermal neutron fission products up to 9.5 MeV. *Physics Letters B*, 160(4):325 – 330, 1985.
- [265] Th. A. Mueller, D. Lhuillier, M. Fallot, et al. Improved predictions of reactor antineutrino spectra. *Phys. Rev. C*, 83:054615, May 2011.
- [266] Patrick Huber. On the determination of anti-neutrino spectra from nuclear reactors. *Phys. Rev.*, C84:024617, 2011. [Erratum: *Phys. Rev.*C85,029901(2012)].
- [267] D. S. Akerib, M. J. Attisha, L. Baudis, et al. Low-threshold analysis of CDMS shallow-site data. *Phys. Rev. D*, 82:122004, December 2010.
- [268] M Schwaiger, F Steger, T Schroettner, and C Schmitzer. A ultra low level laboratory for nuclear test ban measurements. *Applied Radiation and Isotopes*, 56(1):375 – 378, 2002. Proceedings of the Conference on Radionuclide Metrology and its Applications, ICRM’01.

- [269] Holger Gastrich, Claus Gößling, Reiner Klingenberg, et al. The Dortmund Low Background Facility — Low-background gamma ray spectrometry with an artificial overburden. *Applied Radiation and Isotopes*, 112:165 – 176, 2016.
- [270] S. M. Stanchfield, P. A. R. Ade, J. Aguirre, et al. Development of a Microwave SQUID-Multiplexed TES Array for MUSTANG-2. *Journal of Low Temperature Physics*, 184(1):460–465, July 2016.
- [271] J. A. B. Mates, D. T. Becker, D. A. Bennett, et al. Simultaneous readout of 128 X-ray and gamma-ray transition-edge microcalorimeters using microwave SQUID multiplexing. *Applied Physics Letters*, 111(6):062601, 2017.
- [272] Mayuri Sathyanarayana Rao et al. Simons Observatory Microwave SQUID Multiplexing Readout – Cryogenic RF Amplifier and Coaxial Chain Design. *J. Low Temp. Phys.*, 1:10, 2020.
- [273] M. Wegner, N. Karcher, O. Krömer, et al. Microwave SQUID Multiplexing of Metallic Magnetic Calorimeters: Status of Multiplexer Performance and Room-Temperature Readout Electronics Development. *Journal of Low Temperature Physics*, 193(3):462–475, November 2018.
- [274] Douglas A. Bennett, John A. B. Mates, Simon R. Bandler, et al. Microwave SQUID multiplexing for the Lynx x-ray microcalorimeter. *Journal of Astronomical Telescopes, Instruments, and Systems*, 5(2):1 – 10, 2019.
- [275] J. Åström et al. Fracture Processes Observed with A Cryogenic Detector. *Phys. Lett. A*, 356:262–266, 2006.
- [276] R. Strauss, G. Angloher, A. Bento, et al. A detector module with highly efficient surface-alpha event rejection operated in CRESST-II Phase 2. *The European Physical Journal C*, 75(8):1–8, 2015.
- [277] A. Aliane et al. First test of a $\text{Li}_2\text{WO}_4(\text{Mo})$ bolometric detector for the measurement of coherent neutrino-nucleus scattering. *Nucl. Instrum. Meth. A*, 949:162784, 2020.
- [278] G. Angloher, F. Ardellier-Desages, A. Bento, et al. Exploring CEvNS with NUCLEUS at the Chooz nuclear power plant. *The European Physical Journal C*, 79(12):1018, December 2019.
- [279] G. Heusser. Low-radioactivity background techniques. *Ann. Rev. Nucl. Part. Sci.*, 45:543–590, 1995.
- [280] J.A. Formaggio and C.J. Martoff. Backgrounds to sensitive experiments underground. *Ann. Rev. Nucl. Part. Sci.*, 54:361–412, 2004.
- [281] R. Maisonobe, J. Billard, M. De Jesus, et al. Vibration decoupling system for massive bolometers in dry cryostats. *JINST*, 13(08):T08009, 2018.
- [282] R. Strauss et al. Energy-dependent light quenching in CaWO_4 crystals at mK temperatures. *Eur. Phys. J. C*, 74(7):2957, 2014.

- [283] TUM FRMII – FAQs. <https://www.frm2.tum.de/en/fuel-development/faqs/>. [Online; accessed November 25, 2019].
- [284] Achim Gütlein. *Feasibility Study for a First Observation of Coherent Neutrino Nucleus Scattering Using Low-Temperature Detectors*. Dissertation, Technische Universität München, München, 2013.
- [285] J. Gaffiot, T. Lasserre, G. Mention, et al. Experimental parameters for a Cerium 144 based intense electron antineutrino generator experiment at very short baselines. *Phys. Rev. D*, 91:072005, April 2015.
- [286] M Laubenstein, M Hult, J Gasparro, et al. Underground measurements of radioactivity. *Applied Radiation and Isotopes*, 61(2):167 – 172, 2004. Low Level Radionuclide Measurement Techniques - ICRM.
- [287] Y. Abe, C. Aberle, J. C. dos Anjos, et al. Direct measurement of backgrounds using reactor-off data in Double Chooz. *Phys. Rev. D*, 87:011102, January 2013.
- [288] B.C. Cañas, E.A. Garcés, O.G. Miranda, M. Tórtola, and J.W.F. Valle. The weak mixing angle from low energy neutrino measurements: A global update. *Physics Letters B*, 761:450 – 455, 2016.
- [289] S.K. Agarwalla et al. Constraints on flavor-diagonal non-standard neutrino interactions from Borexino Phase-II. *JHEP*, 02:038, 2020.
- [290] D. K. Papoulias and T. S. Kosmas. COHERENT constraints to conventional and exotic neutrino physics. *Phys. Rev. D*, 97:033003, February 2018.
- [291] C. Giunti. General COHERENT Constraints on Neutrino Non-Standard Interactions. 2019.
- [292] John R De Laeter, John Karl Böhlke, Paul De Bièvre, et al. Atomic weights of the elements. Review 2000 (IUPAC Technical Report). *Pure and applied chemistry*, 75(6):683–800, 2003.
- [293] Enrico Bertuzzo, Frank F. Deppisch, Suchita Kulkarni, Yuber F. Perez Gonzalez, and Renata Zukanovich Funchal. Dark Matter and Exotic Neutrino Interactions in Direct Detection Searches. *JHEP*, 04:073, 2017.
- [294] Martin Hoferichter, J. Ruiz de Elvira, Bastian Kubis, and Ulf-G. Meißner. High-Precision Determination of the Pion-Nucleon Term from Roy-Steiner Equations. *Phys. Rev. Lett.*, 115:092301, 2015.
- [295] Parikshit Junnarkar and Andre Walker-Loud. Scalar strange content of the nucleon from lattice QCD. *Phys. Rev.*, D87:114510, 2013.
- [296] J. Rothe, G. Angloher, F. Ardellier-Desages, et al. NUCLEUS: Exploring Coherent Neutrino-Nucleus Scattering with Cryogenic Detectors. *Journal of Low Temperature Physics*, pages 1–8, 2019.
- [297] V. Radeka and N. Karlovac. Least-square-error amplitude measurement of pulse signals in presence of noise. *Nuclear Instruments and Methods*, 52(1):86 – 92, 1967.

- [298] Ralph J Schwarz and Bernard Friedland. *Linear systems*. McGraw-Hill, 1965.
- [299] Sunil Golwala. *Exclusion Limits on WIMP-Nucleon Elastic Scattering Cross-Section from the Cryogenic Dark Matter Search*. Ph.D. Thesis, Department of Physics, University of California, Berkeley, Berkeley, 2000.
- [300] E. Gatti and P.F. Manfredi. Processing the Signals From Solid State Detectors in Elementary Particle Physics. *Riv. Nuovo Cim.*, 9N1:1–146, 1986.
- [301] M. Mancuso, A. Bento, N. Ferreiro Iachellini, et al. A method to define the energy threshold depending on noise level for rare event searches. *Nucl. Instrum. Meth. A*, 940:492–496, 2019.
- [302] H.A. David and H.N. Nagaraja. *Order Statistics*. Wiley Series in Probability and Statistics. Wiley, 2004.
- [303] Jonathan H. Davis. Probing Sub-GeV Mass Strongly Interacting Dark Matter with a Low-Threshold Surface Experiment. *Phys. Rev. Lett.*, 119(21):211302, 2017.

Acknowledgements

To Wolfgang, who taught me much
in too short a time.

It has been almost six years since I first visited the MPP cryodetector group to talk about a possible Master thesis project. After this long time, I have much to be grateful for. I thank my thesis advisor Prof. Stefan Schönert for giving me the opportunity to pursue my PhD at the MPP cryodetectors group.

To Federica: it was great to have a supervisor that remembers this well the position of the student. Thank you for your patience with my extensive writing style, for your confidence that I'm able to finish in time, and for choosing the right moment to set some deadlines and make it happen.

To Karo, who gave the spark that lit the fire. Without that brief conversation at LNGS I may never have found to cryodetectors and MPP. It's great to have you back!

To my supervisor of a few months, before well-deserved and very active retirement: Franz, it is hard to put in words how fundamental you have been to my work. Your dedication is the basis for so many aspects of the experiment. Your contagious fascination with all the details are what got me into all this in the first place. I was lucky to have you just a brief email away, for detector physics discussions, analysis software feature requests and DAQ bugfixes on the weekend. Your commitment goes way beyond what can be reasonably expected. It will continue to be a great example.

Raimund, without you, this thesis would be half of what it became, and the future would be less clear. Thank you for your vision, your approach to work and life, your organization and the adventures. I'm looking forward to the next few years!

Special thanks go to my proofreaders Federica, Lucia, Antonio, Raimund and Franz. When the writing is done, the author is dependent on the efforts of others to close the feedback loop. I never had to wait long for your comments, no matter how many pages I buried you with. Thank you for that! I have great respect for the focus and dedication that was necessary from my proofreaders, especially towards the end.

My work owes much to Dieter, who made possible all the technical developments in this thesis. Your experience with crystals, evaporation machines, photolithography and chemical processing was invaluable. Thank you for your advice, patience and creativity. Without your constant efforts, we could not even be thinking about where we are today.

Thanks to my predecessors Anja and Marc, who taught me most of what I know about software and hardware in our experiment in my early time. To them and my colleagues Michele, Nahuel, Elia, Ahmed: thank you for creating an atmosphere of growth and learning, fueled by endless questions. It has been a privilege.

To Gode and Ahmed, our creative renegades! Thank you for showing how much there is we don't understand, and how many ways there are to fix that. Convincing yourselves is only the first step though!

To Lucia, the person I have known longest in the field. So nice to have my internship supervisor turn into a collaborator and finally colleague and office-mate. We shared an office much too briefly! You are excused for not proofreading the later chapters of my thesis. You really had much better things to take care of!

To the Gran Sasso crew: Paolo, Nello, Miriam, Stefano and Luca S., thanks for doing so much to keep the experiment running in good shape, now more than ever! Vanessa, thanks for sharing your office when I was on shift - it was a pleasure.

Many thanks to Michele, Luca P., Flo, Ahmed, Paolo, Nello, Miriam, Stefano, Elia, who all helped tremendously to demonstrate the importance of magnetic effects and implement countermeasures. Special thanks to Si at the workshop, for power supplies, magnetometers, almost 100 m of CERN legacy cables, and rapid prototyping of current sources.

To my old and new colleagues, at TU Munich, in Paris, Vienna and in Rome. It has been a pleasure and I'm looking forward to more work with you!

To Stefan Petersen, for creating an environment for physics discussion, exchange and enrichment. Thank you for making the German Physics Olympiad a place where students can keep in touch and contribute! It's been a pleasure to be a part of this all these years. And thanks to my supervisors for always considering this business trips!

An meine Familie: ihr habt mich seit ich denken kann auf so viele Arten unterstützt. Ohne all das, was ihr mir mitgegeben habt, wäre das sicher nichts geworden mit der Dunklen Materie! Ich denke gern an die zahlreichen Experimente mit meinem Großvater, die mir schon früh den Spaß am Experimentieren vermittelt haben. Vielen Dank an meine Schwester! Für den guten Draht und die zahlreichen Besuche, die du trotz des langen Weges unternimmst, wenn die Umstände es erlauben. Es macht Spaß mit euch!

Marie, du burde ikke bare stå i takksigelsen til doktoravhandlingene mine, men i den til hele mitt liv. Jeg er så glad at du deler ditt med meg!

Johannes Rothe, October 2020