Quieting the Chaos in the Cold:
Noise Cancellation Algorithms in CUORE's Search for $0\nu\beta\beta$ Decay

by

Kenneth J. Vetter

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Yury Kolomensky, Chair
Professor Gabriel Orebi Gann
Professor Emeritus Eric Norman

Summer 2024

Quieting the Chaos in the Cold:
Noise Cancellation Algorithms in CUORE's Search for $0\nu\beta\beta$ Decay

Abstract

Quieting the Chaos in the Cold:
Noise Cancellation Algorithms in CUORE's Search for $0\nu\beta\beta$ Decay

by

Kenneth J. Vetter

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Yury Kolomensky, Chair

Neutrinos are the most elusive particles known to humanity, and their nature has puzzled physicists for decades. While great strides in understanding neutrinos have been made in the past three decades, many open questions about them remain, including their absolute mass scale, the nature of their masses, and the extent to which they violate CP symmetry. The fact that neutrinos are indeed massive has shown that the Standard Model of particle physics, while wildly successful, is incomplete. This fact has motivated searches for Beyond Standard Model processes that may fundamentally change our understanding of nature. Among these potential processes is neutrinoless double beta decay ($0\nu\beta\beta$), a theorized lepton number violating process. Experimental evidence of $0\nu\beta\beta$ decay would conclusively show that neutrinos are Majorana fermions. This has profound implications for particle physics, as it would help to explain the matter-antimatter asymmetry in the universe while also providing insight into the origins of the masses of neutrinos.

The Cryogenic Underground Observatory for Rare Events (CUORE) experiment is an ongoing experiment searching for $0\nu\beta\beta$ decay in ^{130}Te. This dissertation outlines the current status of the ongoing global search for $0\nu\beta\beta$ and details how the CUORE experiment is contributing to this effort. The results presented here showcase the use of a novel noise cancellation algorithm to improve the CUORE data quality, including the RMS noise, energy resolution, and threshold. The application of this algorithm to CUORE's latest dataset with 2039 kg-years of TeO_2 exposure is then detailed. The results from CUORE's search for $0\nu\beta\beta$ decay using this dataset are also reported. No evidence for $0\nu\beta\beta$ decay is found, and an upper limit on the decay rate of $\Gamma_{0\nu} < 2.0 \times 10^{-26}$ yr^{-1} (90% C.I.) is placed, corresponding to a lower limit on the half-life of $T^{0\nu}_{1/2} > 3.5 \times 10^{25}$ yr. The 90% C.I. exclusion sensitivity to $0\nu\beta\beta$ decay is $T^{0\nu}_{1/2} > 4.4 \times 10^{25}$ yr, and the 3σ discovery sensitivity is $T^{0\nu}_{1/2} > 2.6 \times 10^{25}$ yr. Finally, future searches for $0\nu\beta\beta$ decay beyond CUORE and the prospects of noise decorrelation algorithms for these searches are discussed.

To my father

Be still my beating heart
It would be better to be cool
It's not time to be open just yet

– Sting

Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank the many people I have met on this journey, particularly those who have been my mentors along the way. First and foremost, thanks to my advisor, Professor Yury Kolomensky. I will forever appreciate our many brainstorming sessions, late-night conversations, and time together in the classroom. Thank you for helping to shape me into the physicist I am today.

I want to acknowledge the other members of my dissertation committee, Professors Gabriel Orebi Gann and Eric Norman, for their invaluable time and effort in shaping this dissertation. I would also like to acknowledge Professors Lawrence Hall and Kai Vetter for serving on my qualifying exam committee in 2021.

Thanks to Dr. Sergio Zimmermann of the LBL Engineering Division as well. Sergio, I have learned a great deal from you over the course of our years working together. To be able to work with you for so long has been an absolute joy. I wish you well in your retirement, and I am ecstatic to see that you are traveling the world. I know that we will keep in touch in the years to come.

I would like to thank Dr. Bradford Welliver next. Much of the work in this dissertation would not have been possible without his guidance. Thank you, Brad, for your endless dedication to helping me learn more. You have imparted so much knowledge to me, and I am excited to continue working with you in the future.

Many thanks also to Dr. Vivek Singh, who taught me an incredible amount during our time in the laboratory. You have played a pivotal role in my development as an experimental physicist.

Dr. Erin Hansen has been a fantastic mentor in every aspect of my work. Thank you, Erin, for helping me navigate life as a Ph.D. student and being there when I needed help. You have made me a better researcher, a better teacher, and a better mentor.

Thank you to Dr. Laura Marini for being a mentor from the beginning until the end. From writing my first ROOT code in 2019 to assembling and cooling down the CCVR2 this February, you have been teaching me new things every step of the way.

Thanks also to Dr. Irene Nutini for being an excellent leader and role model. You have taught me a great deal not only about how CUORE works and how to make DIANA do what I want it to do but also about being a leader in the collaboration. I have looked up to you since I started my Ph.D. and continue to do so today.

I would also like to thank Dr. Chiara Capelli for helping me with every aspect of my life during my first trip on-site. From working late hours underground to patiently helping me practice my Italian, your guidance and mentorship were incredible. Were it not for you, L'Aquila would never have felt like a second home as it does now.

Thanks to Dr. Mattia Beretta for being a great mentor and a friend with whom I can share jokes and stories about work, language, and music. I cherished our time together at Berkeley, and I am delighted that we will continue to be able to work together.

My deepest thanks to Professor Giovanni Benato as well. Gio, you have filled my graduate studies with knowledge, laughter, and fond memories. Here's to many more in the years to come.

To my fellow graduate students, present and past, whom I have met along the way – Roger Huang, Sachi Wagaarachchi, Tong Zhu, and Joe Camilleri – thank you all for being supportive, friendly, and inviting. You have all helped make 321 Physics South an office space I will remember fondly.

I sincerely thank all of my fellow members of the CUORE and CUPID collaborations. Collaborating with such an exceptional group of scientists and remarkable individuals has been a privilege and an honor. My graduate studies have created friendships that span the globe. I have made far too many friends in the past six years to list them all here. The memories of our time together will always be cherished. L'Aquila, Italy is now a place I hold dear, home to many of my cherished friends, mentors, and colleagues. To be a part of a collaboration with such a rich history is a truly humbling experience. As one of the fortunate individuals to be part of these endeavors, I want to express my gratitude and acknowledge everyone who has contributed to shaping the CUORE and CUPID collaborations into what they are today.

There have been many people within the UC Berkeley Physics community who have helped me through the years. In particular, I want to thank Austin Hedeman and Marjorie Shapiro, both of whom have been instrumental in my development as a teacher, and Joelle Miles, whose kindness knows no bounds and without whom I would never have made it this far.

Thank you to my good friend Roman Marcarelli for helping me with my coursework whenever I struggled and for being more than willing to discuss particle physics in detail to become a better teacher. I hope that our friendship will continue to thrive for years to come.

Thank you to my best friend in Berkeley, Max Smiley. After living together for six years, including being locked down together during the pandemic, I am proud to call you my friend. For much of my graduate studies, you were my rock. I am truly going to miss having you around.

Finally, to my loving parents, grandparents, brothers, aunts, uncles, and cousins, thank you all for your thoughts, kind words, and emotional support throughout this endeavor. None of this would be possible without you.

Many more people deserve to be acknowledged here. To those I have forgotten to mention, I sincerely apologize. Rest assured that everyone who helped me along the way made all of this work possible. Earning a Ph.D. is like raising a child, only in the sense that it takes a village to do so. Thankfully, earning a Ph.D. is done in about a third of the time.

Chapter 1

Introduction

Of the many unanswered questions in physics today, among the most intriguing is perhaps: *"Why are we here?"* A potentially satisfying answer may come from the generally accepted Big Bang theory [1], which proposes that the universe began as a highly concentrated region of space-time with an incomprehensibly large energy density which then expanded, cooled, and ultimately gave way to the vast network of galaxies, stars, and planets we observe today. Unfortunately, as many particle physicists and cosmologists are acutely aware, the current standard model of Big Bang cosmology, ΛCDM, [2] coupled with the Standard Model of particle physics[1] [3–6], predicts that our universe should consist of nearly equal parts matter and antimatter. As a universe which obeys this prediction expands and cools, one would expect the matter and antimatter to annihilate with each other, producing only photons. This is, of course, not the universe that we observe today. Therefore, our earlier question should slightly be rephrased as: *"Why is the observable universe composed of matter instead of antimatter?"*

It is important to note that we can only safely conclude that the *observable* universe is dominated by matter instead of antimatter. It is entirely possible that the entire observable universe exists within a volume in which there is an over-fluctuation of matter over antimatter, and that beyond this bubble, there are bubbles containing over-fluctuations of antimatter such that on average, the universe indeed has no preference for matter over antimatter. The Copernican principle [7] provides an epistemological ground on which to reject this idea, but the idea is certainly physically possible and has been the subject of many theoretical discussions [8–10].

The underlying reason for *why* a matter-antimatter asymmetry exists in the observable universe quickly enters the realms of philosophy and perhaps even theology. (As Richard Feynman once stated in a 1983 interview as a part of the 1983 BBC program *Fun to Imagine*, "How does a person answer why something happens?" [11]) Given the scope of this dissertation, it is ultimately more appropriate for us to frame the question in the following way:

---

[1] Throughout this dissertation, only the Standard Model of particle physics will be capitalized. Apologies to the astrophysics and cosmology communities.

"How can we explain an asymmetry between the amount of matter and antimatter in the observable universe in a way that is consistent with experimental results?"

Particular phrasing aside, this unanswered question is known as the Baryon Asymmetry Problem, and it is the subject of much research and debate in particle physics and cosmology. Here, I will explore one particular avenue out of the many available to the scientific community to address this question. I will focus on the ongoing effort to answer this question through the lens of particle and nuclear physics. Specifically, this dissertation will focus on the curious nature and still unknown properties of the neutrino, how the neutrino may play a role in resolving the matter-antimatter asymmetry problem, and how physicists are trying to unveil these properties using experimental nuclear physics techniques.

The first four chapters of this dissertation will provide relevant theoretical and experimental background information. Chapter 1 will provide a more technical description of the matter-antimatter asymmetry problem, its connection to our current understanding of particle physics, and why neutrinos may play a key role in its resolution. It will also detail the history of neutrino physics from both an experimental and theoretical point of view. Chapter 2 will discuss the phenomena of two-neutrino and neutrinoless double beta decay and why the existence of the latter would provide insight into the properties of neutrinos that may help resolve the matter-antimatter asymmetry problem. Chapter 3 will provide an overview of the Cryogenic Underground Observatory for Rare Events (CUORE) experiment, an ongoing search for neutrinoless double beta decay, and Chapter 4 will detail the data processing chain used to conduct this search.

The last five chapters of this dissertation will focus on my contributions to the CUORE experiment. Chapter 5 will detail the characteristics of the noise in the CUORE detectors and its sources; much of this work comes from studies I have performed over the past several years. Chapters 6 and 7 will describe the heart of my contribution to the CUORE experiment: a novel noise decorrelation algorithm which improves the CUORE data quality when added to the processing chain. Chapter 8 will outline the latest results from CUORE in its search for neutrinoless beta decay, an analysis in which I played a major role. Finally, chapter 9 will discuss the future of the search for neutrinoless double beta decay, the CUORE Upgrade with Particle Identification (CUPID) experiment, and the future technologies, including new noise decorrelation algorithms, that will be implemented to make CUPID a reality.

1.1 On the Origin of Everything

1.1.1 Baryogenesis

The Baryon Asymmetry Problem has been on the minds of physicists for decades. A major step forward in articulating and solidifying this problem in a particle physics context was a 1967 paper by Andrei Sakharov. This paper outlined a set of three conditions necessary for *Baryogenesis*, i.e. the production of an excess of baryons over anti-baryons during the evolution of the universe. Now known as the “Sakharov conditions,” these three conditions are [12]:

- Baryon number (B) violation
- Charge conjugation symmetry (C) violation and charge conjugation parity symmetry (CP) violation
- Interactions which take place out of thermal equilibrium

Baryon number is defined as

$$\mathsf{B} = \frac{1}{3}(n_q - n_{\overline{q}}) \tag{1.1}$$

where n_q is the number of quarks and $n_{\overline{q}}$ is the number of anti-quarks. A process that conserves baryon number must therefore produce quarks and anti-quarks in equal quantities or destroy them in equal quantities. The violation of baryon number is thus a self-evident condition for a matter-antimatter asymmetry to arise; you cannot have more baryons than anti-baryons in the observable universe without a process that produces more quarks than anti-quarks in the early universe. At low energies, the Standard Model provides no such interaction, however as discussed in more detail soon in section 1.1.2, there are interactions allowed by the Standard Model at high energies which can violate baryon number.

The fact that CP must be violated is effectively axiomatic because CP conjugation maps particles to anti-particles. Any baryon-generating mechanism that conserves CP would produce baryons and anti-baryons in equal numbers, which is obviously not what we need to explain what we see in our universe, thus Baryogenesis in our universe requires CP violation. CP violation is already described by the Standard Model and there is experimental evidence confirming its existence. In the quark sector, CP violation has been observed in many interactions that are consistent with the Standard Model. In the lepton sector, there is a solid theoretical framework which is analogous to that in the quark sector, but the experimental picture is less clear. There are many ongoing experiments still searching for CP violation in the lepton sector. In particular, the Tokai to Kamioka (T2K) collaboration has excluded CP conservation at 95% confidence in long-baseline neutrino and antineutrino oscillations. This confidence interval increases to 99% assuming the inverted neutrino mass ordering [13]. On the other hand, the NuMI Off-axis ν_e Appearance (NOνA) experiment has reported results that are consistent with both CP conservation and CP violation [14, 15].

The third Sakharov condition is perhaps the most subtle. It first assumes that CPT (the combination of CP and the time reversal operator T) is conserved. The CPT theorem [16] states that any quantum field theory that is unitary, causal, and Lorentz-invariant must be invariant under a CPT transformation. We generally expect that this will hold for any beyond Standard Model physics we may find. If CPT is conserved, then any process obeying the first two Sakharov conditions occurring in thermal equilibrium will still fail to create an excess of baryons. In thermal equilibrium, the rate R_b of the baryon-creating process will be equal to the rate of the time-reversed process $\mathsf{T}(R_b)$ that destroys baryons. The same is true for the creation rate ($R_{\bar{b}}$) and annihilation rate ($\mathsf{T}(R_{\bar{b}})$) of anti-baryons, so we must have a process occurring out of thermal equilibrium to have an excess of baryons in the universe. This condition is met during the expansion of the universe so long as the rate of excess baryon production is slower than the rate of expansion of the universe [17].

It may seem that there is no need for beyond Standard Model physics to explain the matter-antimatter asymmetry in our universe, since all three Sakharov conditions are met in the standard models of cosmology and particle physics. To analyze this more carefully, it is useful to define the baryon number asymmetry density:

$$\eta \equiv \frac{n_b - n_{\bar{b}}}{n_\gamma} \tag{1.2}$$

where $n_{b(\bar{b})}$ is the (anti-)baryon number density per unit volume, and n_γ is the photon number density per unit volume. Big Bang Nucleosynthesis (BBN) predicts abundances of light elements which are in good agreement with observational data. Using BBN based on the Standard Model with three neutrino flavors, and allowing η to be a free parameter in the model, these abundances are best explained by a value of $\eta_{BBN} = (6.04 \pm 0.12) \times 10^{-10}$ [18]. This is in good agreement with an analysis from the Planck collaboration using data from the cosmic microwave background (CMB) which corresponds to a value of $\eta_{CMB} = (6.12 \pm 0.04) \times 10^{-10}$ [19]. However, in the Standard Model, the amount of CP violation observed in the quark sector only allows for a value of $\eta \approx 10^{-26}$ [20]. This suggests that any baryogenesis due to electroweak interactions must have an additional source of CP violation that is so far unknown. Finally, since the universe was expanding extremely rapidly during the Big Bang, it is possible that even a universe obeying a matter-antimatter symmetry would have some primordial matter and antimatter that survives all the way until the conditions of the present-day, effectively "freezing out" as the universe cools, but this can only account for a value of $\eta \approx 10^{-20}$ [21]. The fact that this leaves us to explain another ten orders of magnitude of matter in the universe effectively necessitates the existence of physics beyond the standard models of particle physics and cosmology.

1.1.2 Leptogenesis

A common way to theoretically resolve the Baryon Asymmetry Problem is with grand unified theories (GUTs), which predict baryogenesis with the right order of magnitude in the early universe. Cosmic inflation complicates this picture, however, because the temperature of the universe was only above the GUT energy scale before inflation, and inflation as it is currently understood would significantly reduce the baryon density in the universe [22]. It is also extremely unlikely that anywhere in the universe reach temperatures above the GUT scale after inflation. In the 1980s, it was shown that non-perturbative electroweak processes which do not conserve baryon number could in fact exceed the expansion rate at temperatures above ∼100 GeV [23]. These processes, known as sphaleron transitions, violate the symmetry of $\mathsf{B}+\mathsf{L}$, but conserve $\mathsf{B}-\mathsf{L}$, where L is the lepton number operator. The baryon number asymmetry can be written exactly in terms of the $\mathsf{B}-\mathsf{L}$ asymmetry [24]:

$$n_{\mathsf{B}} = \left(\frac{8N_f + 4N_H}{22N_f + 13N_H}\right) n_{\mathsf{B-L}} \tag{1.3}$$

where N_f is the number of fermion families, and N_H is the number of Higgs doublets. This implies that if a lepton number asymmetry existed in the early universe, sphaleron transitions could have partially converted this asymmetry into a baryon asymmetry, thus providing a route to baryogenesis.

The Standard Model treats B and L as discrete symmetries. One way of extending the Standard Model is to introduce $U(1)_{\mathsf{B-L}}$ as a gauge symmetry which is spontaneously broken at a high energy scale on the order of $10^{10} - 10^{13}$ GeV. Following [25] and [26], we can write this extended Standard Model gauge group as:

$$\begin{aligned} SU(3)_C \times SU(2)_L \times U(1)_Y \times U(1)_{Y'} &\xrightarrow{\langle\chi\rangle = v'} SU(3)_C \times SU(2)_L \times U(1)_Y \\ &\xrightarrow{\langle\phi\rangle = v} SU(3)_C \times U(1)_{EM} \end{aligned} \tag{1.4}$$

where ϕ is the Standard Model Higgs doublet and χ is a new singlet Higgs field with vacuum expectation value v' which is needed to break the $U(1)_{Y'}$ symmetry. Omitting the $SU(3)_C$ group, the Lagrangian for this gauge group is:

$$\begin{aligned} \mathcal{L} \supset &- \frac{1}{4}\tilde{W}_{\mu\nu}\tilde{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}C_{\mu\nu}C^{\mu\nu} \\ &+ i\bar{\ell}\not{D}\ell + i\bar{\nu}_R\not{D}\nu_R + i\bar{e}_R\not{D}e_R + (D_\mu\phi)^\dagger(D^\mu\phi) + (D_\mu\chi)^\dagger(D^\mu\chi) - V(\chi,\phi) \\ &- (\bar{\ell}\phi g_\nu\nu_R + \bar{\ell}\bar{\phi} g_e e_R + \frac{1}{2}\chi\bar{\nu}_R^c h\nu_R + \text{h.c.}) \end{aligned} \tag{1.5}$$

where g_e, g_ν, and h are the Yukawa coupling matrices for the charged leptons, the left-handed neutrinos, and the right-handed neutrinos, respectively. The term which couples right-handed neutrino field to its charge conjugate field, $-\frac{1}{2}\chi\bar{\nu}_R^c h\nu_R$, implies the existence of a *Majorana mass*. In the above equation, the covariant derivative is the same as the Standard Model but with an additional contribution from the $U(1)_{Y'}$ group:

$$D_\mu = \partial_\mu - igW_\mu^\alpha T_\alpha - ig'B_\mu Y - ig'\sqrt{\frac{2}{3}}C_\mu Y' \tag{1.6}$$

where C_μ is the $U(1)_{Y'}$ gauge field. The hypercharge Y' is related to $\mathsf{B}-\mathsf{L}$ by:

$$Y' = Y - \frac{5}{4}(\mathsf{B}-\mathsf{L}) \tag{1.7}$$

Finally, the most general Higgs potential in this model can be written as:

$$V(\chi,\phi) = \mu_1\phi^\dagger\phi + \mu_2\chi^*\chi + \frac{1}{2}\lambda_1(\phi^\dagger\phi)^2 + \frac{1}{2}\lambda_2(\chi^*\chi)^2 + \lambda_3(\chi^*\chi)(\phi^\dagger\phi) \tag{1.8}$$

where $\lambda > -\sqrt{\lambda_1\lambda_2}$ to ensure the potential remains finite from below. When spontaneous symmetry breaking first occurs (see equation 1.4), a massive neutral Z' boson is created, and the right-handed neutrinos acquire Majorana masses $M = hv'$. Since v' is a new free parameter distinct from the Standard Model, it can potentially generate large Majorana masses. The existence of large Majorana masses is important because they may play a role in generating the small neutrino masses we observe today. This is discussed in more detail in section 1.2.3.

Assuming that the mass matrix h is real and diagonal, the weak eigenstates are equivalent to the Majorana mass eigenstates, i.e. the left-handed neutrinos we observe today are given by:

$$\nu = \nu_L + \nu_L^C \tag{1.9}$$

and the so-far unobserved right-handed neutrinos are given by:

$$N = \nu_R + \nu_R^C \tag{1.10}$$

The right-handed neutrinos, if sufficiently heavy, can decay into Standard Model leptons and Higgs bosons. The two decay channels are [22]:

$$\begin{aligned} N &\to \ell_L + \phi^\dagger \\ N &\to \bar{\ell}_L + \phi \end{aligned} \tag{1.11}$$

So long as CP is violated and these decays occur out of thermal equilibrium, there is a difference in the branching ratios of these two decay paths. Thus, the decay of heavy right-handed neutrinos could reasonably produce an excess of leptons in the early universe, paving the way for baryogenesis and the matter-dominated universe we observe today.

This model, which introduces a Higgs singlet, is called the singlet majoron model [27], so-named because the Goldstone boson associated with the breaking of the $U(1)_{Y'}$ symmetry is known as the majoron. This is of course not the only model that predicts leptogenesis in the early universe, but it is compelling for a few reasons: it is a minimal extension of the Standard Model which invokes only one new gauge field, it predicts the existence of heavy right-handed neutrinos with Majorana masses, and it provides a mechanism for leptogenesis and subsequent baryogenesis in the early universe.

Neutrinos are of great interest to the particle physics and cosmology communities; they are an excellent avenue to potentially find physics beyond the Standard Model of particle physics, and they may be the key to explaining leptogenesis and baryogenesis in the early universe. While right-handed neutrinos have no experimental confirmation, the existence of extremely light, left-handed neutrinos has been known for decades. There is a rich history of experimental efforts that have lead to this knowledge. We will now turn our attention away from theorized models and towards the established phenomenology of the left-handed neutrinos.

1.2 Neutrino Physics

Neutrinos were first hypothesized to exist by Wolfgang Pauli in 1930 to explain the observed spectra of beta decays. Later Enrico Fermi, who coined the term "neutrino," proposed an effective field theory describing the weak nuclear force which predicted the emission of antineutrinos during beta decay. In 1937, Ettore Majorana showed that unlike other fundamental fermions, neutrinos could possibly be their own antiparticle [28]. In 1956, Clyde Cowan and Frederick Reines provided the first experimental evidence for the existence of neutrinos by observing inverse beta decays induced by reactor antineutrinos [29]. In the late 1960s, Raymond Davis Jr. conducted an experiment utilizing a large tank filled with perchloroethylene located deep underground in the Homestake Gold Mine in South Dakota [30]. The experiment measured electron neutrino interactions with chlorine atoms in the tank mediated by the weak force:

$$^{37}\mathrm{Cl} + \nu_e \rightarrow ^{37}\mathrm{Ar} + e^- \tag{1.12}$$

It provided the first direct evidence of neutrinos from the sun. However, it also gave rise to a significant discrepancy between the number of neutrinos Davis expected to observe given theoretical models of the nuclear fusion processes in the sun and the number of neutrinos he actually detected [31]. This so-called "Solar Neutrino Problem" puzzled physicists for several decades, hinting at a gap in the understanding of either solar physics or neutrino properties. This was later resolved with the discovery of neutrino oscillation by the SNO and Super-K experiments, a phenomenon in which neutrinos change their flavor as they travel, which was not accounted for in the original predictions from solar models. This discovery showed that neutrinos have mass, providing the first evidence of physics beyond the Standard Model [32].

1.2.1 Neutrino Oscillations

With the development of the Standard Model came a more complete understanding of the lepton sector. The three generations of charged leptons e, μ, and τ each have an associated ν_e, ν_μ, and ν_τ, or an associated "flavor eigenstate" of the three neutrinos. The left-handed fields associated with these pair up to create three doublets:

$$\begin{pmatrix}\nu_e\\ e_L\end{pmatrix}, \begin{pmatrix}\nu_\mu\\ \mu_L\end{pmatrix}, \begin{pmatrix}\nu_\tau\\ \tau_L\end{pmatrix}$$

that transform under the electroweak $SU(2)_L$ gauge symmetry. After electroweak symmetry breaking, this results in the various interactions with the weak bosons such as $W^- \longrightarrow e^- + \overline{\nu}_e$ and $Z \longrightarrow \nu_e + \overline{\nu}_e$. The phenomenon of neutrino oscillations can be described by representing each of the three flavor eigenstates as a linear combination of "mass eigenstates," i.e.

$$|\nu_\ell\rangle = \sum_{i=1}^{i=3} U_{\ell i} |\nu_i\rangle \tag{1.13}$$

where ℓ runs over e, μ, τ and indexes the three flavors and i runs from 1 to 3 and indexes the three mass eigenstates. The matrix $U_{\alpha i}$ is a 3×3 mixing matrix known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [33, 34]. It is conventional to parameterize this matrix with three mixing angles θ_{12}, θ_{23} and θ_{13}, a CP-violating phase δ_{CP}, and two so-called Majorana phases α_1 and α_2 which would also give rise to CP violation. The full matrix can then be written as:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.14}$$

where $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$, with θ_{ij} being the mixing angles. It is instructive to see how the existence of separate mass and flavor eigenbases is directly responsible for neutrino oscillations. To outline this, we will temporarily use a two-flavor model in which only the electron and muon neutrinos exist as superpositions of two mass eigenstates, and there is a single mixing angle θ. We start with a neutrino created by a nuclear process which produces an electron neutrino $|\nu(t=0)\rangle = |\nu_e\rangle = U_{e1}|\nu_1\rangle + U_{e2}|\nu_2\rangle$. Using natural units, $(\hbar = c = 1)$, the neutrino will evolve in time according to[2]:

$$|\nu(t)\rangle = e^{-iEt} |\nu_e\rangle \tag{1.15}$$

where $E = \sqrt{p^2 + m^2}$. Assuming that the neutrino mass is small, we can expand this as $E \approx p + \frac{m^2}{2p}$. The mass-dependence of this term requires that we assign different phases to the different mass eigenstates, i.e.

$$|\nu(t)\rangle \approx U_{e1} e^{-it\left(p + \frac{m_1^2}{2p}\right)} |\nu_1\rangle + U_{e2} e^{-it\left(p + \frac{m_2^2}{2p}\right)} |\nu_2\rangle \tag{1.16}$$

The overall phase of e^{-ipt} does not cause any oscillation between the two mass eigenstates, so we can drop it. The probably that the particle is observed in the muon eigenstate $P(|\nu(t)\rangle = |\nu_\mu\rangle)$ is given by:

---

[2] This is admittedly a non-relativistic framework in which we assume the neutrino is produced in an energy eigenstate. It has been shown in [35] that such a framework still produces the same results as a relativistic framework regardless of the choice of "equal-momentum versus equal-energy."

$$P(\nu_e \to \nu_\mu) = |\langle \nu_\mu | \nu(t) \rangle|^2 = \left| U^*_{\mu 1} e^{-i\frac{m_1^2}{2p}t} U_{e1} + U^*_{\mu 2} e^{-i\frac{m_2^2}{2p}t} U_{e2} \right|^2 \tag{1.17}$$

Again using the assumption of a light neutrino (i.e. $E = p + O(m_i^2) \approx p$), the oscillation probability simplifies to:

$$P(\nu_e \to \nu_\mu) = \left| U^*_{\mu 1} U_{e1} e^{-i\frac{m_1^2}{2E}t} + U^*_{\mu 2} U_{e2} e^{-i\frac{m_2^2}{2E}t} \right|^2 \tag{1.18}$$

Expanding this and using the unitarity of the PMNS matrix (which in the two-flavor case reduces to simple trigonometric identities for the mixing angle θ), we have:

$$P(\nu_e \to \nu_\mu) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2}{4E} t\right) \tag{1.19}$$

where $\Delta m^2 = m_2^2 - m_1^2$ is the difference of the squares of the mass eigenvalues, and E is the energy of the neutrino. Finally, again assuming a light neutrino, we say $t \approx L$ to give us the often-used result:

$$P(\nu_e \to \nu_\mu) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{1.20}$$

In the three-flavor case, this becomes:

$$\begin{aligned} P(\nu_\ell \to \nu'_\ell) = \delta_{\ell\ell'} &- 4 \sum_{i>j} \mathrm{Re}(U^*_{\ell i} U_{\ell' i} U_{\ell j} U^*_{\ell' j}) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) \\ &+ 2 \sum_{i>j} \mathrm{Im}(U^*_{\ell i} U_{\ell' i} U_{\ell j} U^*_{\ell' j}) \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right) \end{aligned} \tag{1.21}$$

Finally, we can rewrite the final sum in terms of a lepton-sector analog of the Jarlskog invariant [36]:

$$\mathcal{J} = \sum_{i>j} \mathrm{Im}(U^*_{\ell i} U_{\ell' i} U_{\ell j} U^*_{\ell' j}) = J \sum_{\alpha,\beta} \epsilon_{\ell\ell'\alpha} \epsilon_{ij\beta} \tag{1.22}$$

Like its quark sector cousin, this is nonzero only if the PMNS matrix allows for CP-violation. In this context, it is proportional to $\sin(\delta_{CP})$. Right now, the best fits for δ_{CP} range from 1.08π to 1.58π [37], and the global best fit on δ_{CP} right now is $\delta_{CP} = 1.36^{+0.20}_{-0.16}$ [38]. CP-violation is thus slightly preferred given the current data but by no means confirmed. In the quark sector, the amount of CP-violation is well known: $\mathcal{J} = 3.08^{+0.15}_{-0.13} \times 10^5$ [38].

Neutrino oscillations are a natural consequence of a general unitary matrix U which mixes the flavor states with the mass states. The fact that the oscillation probability depends on the difference of the *squares* of the masses of each mass eigenstate is important to keep in mind; it implies that oscillation experiments are not directly sensitive to the absolute mass scale of the neutrino. (A technique that is directly sensitive to the absolute mass scale is described later in Chapter 2). However, Δm^2_{ij}, θ_{ij}, and Δ_{CP} are known with varying degrees of sensitivity today. It is worthwhile to discuss how each of these parameters has been (and is being) measured experimentally. The jargon often used to describe these parameters derives from experimental results, and this discussion will directly lead to an outline of the neutrino ordering problem.

1.2.2 Measuring the Neutrino Oscillation Parameters

When building a neutrino oscillation experiment, a common practice is to build neutrino beams and neutrino detectors at a known distance L from the creation point of the neutrinos. The neutrino energy E is a parameter that is either controlled (e.g. a beam energy) or predicted by theory (e.g. the expected energies of solar neutrinos). Equation 1.21 tells us that constructing such an experiment restricts the probability of a neutrino changing flavors to be functions of only the norms of the PMNS matrix values (i.e. the three mixing angles θ_{ij} and Δm^2_{ij}). Using reasonable units, we can use Equation 1.21 to write the oscillation probability of from one mass eigenstate to another:

$$P(\nu_i \longrightarrow \nu_j) = \sin^2(2\theta_{ij}) \sin^2\left(\frac{1.27\Delta m^2_{ij} L}{E} \frac{\text{GeV}}{\text{eV}^2\text{km}}\right) \tag{1.23}$$

Experiments are only sensitive to changes in expected numbers of neutrino interactions when the oscillation rate is large so that the statistics are sufficiently large. This means that the oscillation probability "angle" $\frac{\Delta m^2 L}{E}$ is of order 1. If this quantity becomes too small, then the oscillation probability goes to zero, and nothing can be observed along a reasonable length scale. If it becomes too large, then the oscillations occur much too rapidly for an experiment to meaningfully distinguish between the three mass states. Therefore, the ratio $\frac{L}{E}$ governs the range of Δm^2 to which an oscillation experiment is sensitive. In short baseline experiments, $L/E \approx 1\text{eV}^{-2}$. These experiments can broadly be broken down into reactor $\overline{\nu}_e$ disappearance experiments and accelerator ν_μ to ν_e oscillation experiments. Atmospheric neutrino measurements involve detecting neutrinos produced by cosmic rays interacting with the Earth's atmosphere. Atmospheric neutrino energies range from ∼100 MeV to 10 TeV, with the highest flux occurring at a few GeV [39]. Using the peak energy, this gives $L/E \approx 10^{-3}\text{eV}^{-2}$. These experiments are primarily sensitive to Δm_{32} and the mixing angle θ_{23} because they primarily govern the oscillations between ν_μ and ν_τ. For this reason, θ_{23} is known as the "atmospheric angle." Solar neutrino experiments detect neutrinos produced in fusion reactions within the Sun. Solar neutrinos also have a wide range of energies, ranging from tens of keV to above 10 MeV. The flux peaks at roughly 300 keV, for

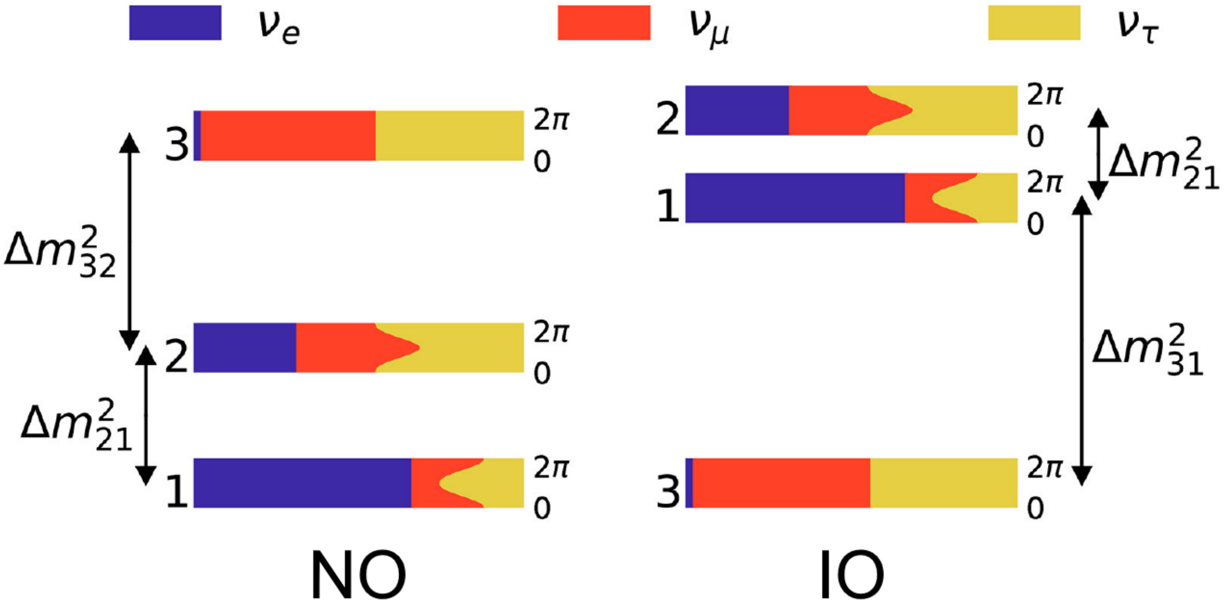

Figure 1.1: A diagram illustrating the two possible neutrino mass orderings. The normal ordering (NO) assumes $m_1 < m_2 < m_3$ while the inverted ordering (IO) assumes $m_3 < m_1 < m_2$. The colors represent the probability of finding a given mass eigenstate in a particular flavor eigenstate. The curved parts of the colored regions represent how these probabilities change as δ_{CP} varies from 0 to 2π. Reprinted from [41].

which $L/E \approx 10^{-13}\text{eV}^{-2}$. The key parameters for solar neutrino experiments are Δm^2_{21} and θ_{12}. The angle θ_{12} is known as the "solar angle" because it is primarily responsible for the oscillation of the ν_e emitted by the Sun.

We know from matter effects in the Sun that $\Delta m^2_{21} = m^2_2 - m^2_1$ is positive and much smaller than $|\Delta m^2_{31}|$, but we have yet to determine the sign of Δm^2_{31} (i.e., whether $m_3 \geq m_1, m_2$) [40]. The "neutrino mass hierarchy problem" or "neutrino mass ordering problem" refers to open question of whether the neutrino masses follow the normal ordering (NO) where $m_1 < m_2 < m_3$, or the inverted ordering (IO) where $m_3 < m_1 < m_2$.

From a global analysis of solar, atmospheric, reactor, and accelerator experiments, the best fit values for the squared-mass differences are $\Delta m^2_{21} = (7.41^{+.21}_{-.20}) \times 10^{-5}$ eV^2 for both NO and IO, $\Delta m^2_{31} = (2.51 \pm 0.03) \times 10^{-3}$ eV^2 for NO, and $\Delta m^2_{32} = (-2.50 \pm 0.03) \times 10^{-3}$ eV^2 for IO [37]. This analysis favors the normal ordering by 2.7σ. The masses of the three neutrinos for each ordering are shown in Fig. 1.1. Given these values, there must be at least two neutrinos with masses of at least 8 meV and at least one neutrino with a mass of at least 50 meV regardless of the ordering. This places a lower bound on the sum of the neutrino masses:

$$\begin{aligned} \sum m_\nu &> 60 \text{ meV (NO)} \\ \sum m_\nu &> 100 \text{ meV (IO)} \end{aligned} \tag{1.24}$$

While we are still unsure about the absolute mass scale of the neutrinos, an analysis of cosmological data from the Dark Energy Spectroscopic Instrument (DESI) collaboration has constrained the sum of the neutrino masses at 95% C.L. to be [42]:

$$\sum m_\nu < (73 - 195) \text{ meV} \tag{1.25}$$

Fig. 1.2.2 shows the posterior distribution on the sum of the neutrino masses using the ΛCDM+$\sum m_\nu$ model. This reported limit includes a range of masses depending on which external datasets are included in the DESI analysis and which extensions to the ΛCDM model are considered. Cosmological results cannot yet rule out the inverted mass ordering, but there is a 2σ preference for the normal ordering from these cosmological results. It should be stressed that these limits are highly dependent on the cosmological models used. Direct detection of the neutrino masses is more desirable in the particle physics community. If a direct measurement of the neutrino mass were to come into tension with cosmological limits, it could indicate new physics beyond ΛCDM, which would certainly be exciting.

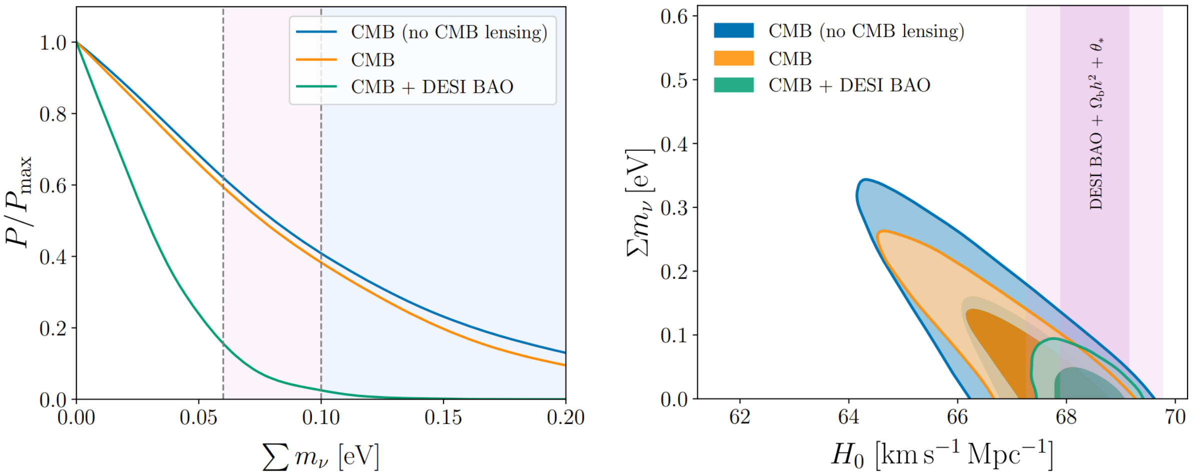

(a) (b)

Figure 1.2: Results on the sum of the neutrino masses from cosmological data. (a) The marginalized 1D posterior distribution and constraints on $\sum m_\nu$ using the single parameter extension flat ΛCDM+$\sum m_\nu$ model. (b) The joint marginalized 68% and 95% credibility intervals on $\sum m_\nu$ and H_0. Reprinted from [42].

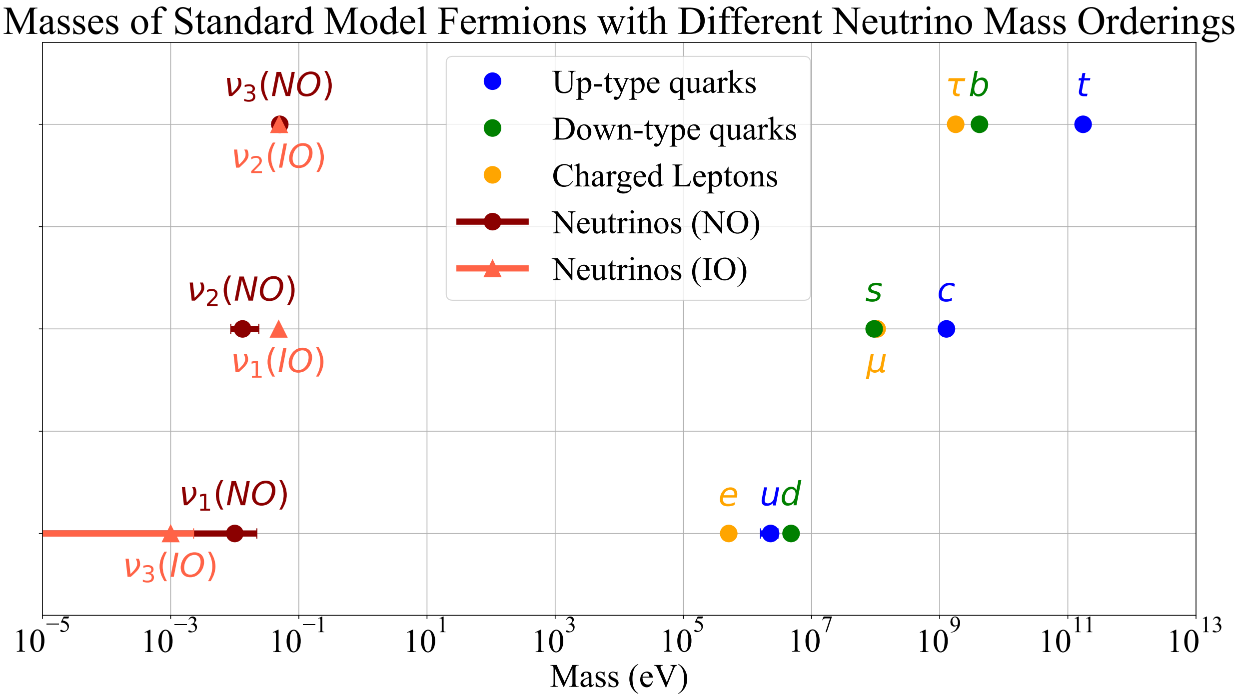

Figure 1.3: Masses of the leptons in the Standard Model and their uncertainties. The absolute mass scale of the neutrinos is still unknown. The ranges on the possible the neutrino masses come from the current values of the oscillation parameters and the limit on $\sum m_\nu$ from cosmological data.

1.2.3 Explaining the Nonzero Neutrino Masses

The masses of the charged leptons are spread across the six orders of magnitude between the electron mass (511 keV) and the top quark mass (173 GeV). The neutrino masses, on the other hand, lie at least six orders of magnitude away from the electron mass (see Fig. 1.3). In the particle physics community, it is widely accepted that it would be incongruous for the mechanism that generates the mass of the neutrinos to be the same as that which generates the masses of the other particles. This suggests that something else may be at play, including potentially new physics.

In the Standard Model, the charged lepton masses are described by a Yukawa term in the electroweak Lagrangian:

$$\mathcal{L}_{EW} \supset -y_\alpha \tilde{\phi}^\dagger \overline{L}_\alpha \ell_{\alpha R} + h.c. \tag{1.26}$$

where $\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$ is the Higgs doublet, $\tilde{\phi} = i\sigma_2\phi^*$, and $L_\alpha = \begin{pmatrix} \nu_{\alpha L} \\ \ell_{\alpha L} \end{pmatrix}$ is one of the six $SU(2)_L$ doublets:

$$\begin{pmatrix}\nu_e\\ e_L\end{pmatrix},\begin{pmatrix}\nu_\mu\\ \mu_L\end{pmatrix},\begin{pmatrix}\nu_\tau\\ \tau_L\end{pmatrix},\begin{pmatrix}u_L\\ d_L\end{pmatrix},\begin{pmatrix}c_L\\ s_L\end{pmatrix},\begin{pmatrix}t_L\\ b_L\end{pmatrix}$$

Each charged lepton has an associated right-handed $SU(2)_L$ singlet field:

$$e_R,\ \mu_R,\ \tau_R,\ u_R,\ c_R,\ t_R,\ d_R,\ s_R,\ b_R$$

Note the conspicuous absence of the right-handed neutrino fields ν_R. This is a reasonable model because of a property unique to neutrinos: *we have never observed a neutrino with right-handed chirality or an antineutrino with left-handed chirality*. This fact is crucial in neutrino physics because it suggests that neutrinos may not gain their mass via the mechanism that applies to the charged leptons. In the Higgs mechanism, all particles are massless at an energy above the electroweak scale, Λ_{EW}. Below this scale, the Higgs field gains a non-zero expectation value $\langle\phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0 & v\end{pmatrix}^T$, and the Yukawa term in the Lagrangian simplifies to:

$$\mathcal{L}_{EW} \supset -y_\alpha v \bar{\ell}_{\alpha L}\ell_{\alpha R} + h.c. \tag{1.27}$$

where $\frac{1}{\sqrt{2}}y_\alpha v = m_\alpha$ is the mass of each charged lepton. Here $v = (\sqrt{2}G_F)^{-1/2} \approx 246$ GeV. There are no right-handed neutrino fields to which the left-handed fields can couple, so the Yukawa couplings cannot give rise to the small neutrino masses we observe. One way to remedy this is to simply add a set of right-handed neutrino fields to the Standard Model and bound the Yukawa couplings. The new term in the Lagrangian is the *Dirac mass* term:

$$\mathcal{L}_N^D \supset -y_{i\alpha}\overline{N}_{iR}\tilde{\phi}^\dagger L_\alpha + h.c. \tag{1.28}$$

where N_{iR} are the right-handed neutrino fields. Using 49 meV $\leq m_{heaviest} \leq$ 120 meV, we get Yukawa couplings of order $10^{-13} - 10^{-12}$. Compare this to the charged lepton Yukawa couplings which range from $y_e \approx 2\times10^{-6}$ to $y_t \approx 0.7$. It could be that these coupling constants simply span across thirteen orders of magnitude. After all, if the Higgs mechanism describes the charged fermions, then they must already span across six orders of magnitude, so why not another seven? The community tends to agree that this fine-tuning of the Yukawa couplings would be rather unnatural. This is a major shortcoming of the Dirac-neutrino model. There are other models that do not suffer from the same problems. In fact, there is a Lagrangian we can write down that is more general than 1.28, provided that we allow for the violation of lepton number L. One can include a Majorana mass term which couples the right-handed neutrino field to its charge conjugate field instead of the left-handed field. (An example of model which does this was outlined earlier in section 1.1.2.) The Lagrangian then becomes:

$$\mathcal{L}_N^M \supset -y_{i\alpha}\overline{N}_{iR}\tilde{\phi}^\dagger L_\alpha - \frac{1}{2}\overline{N}_{iR}M_{ij}(N_{jR})^c + h.c. \tag{1.29}$$

where M_{ij} is a symmetric matrix of Majorana masses. This can be written as a matrix expression:

$$\mathcal{L}_N^M \supset -\frac{1}{2}\begin{pmatrix}\overline{\nu}^c_{\alpha L} & \overline{N}_{iR}\end{pmatrix}\begin{pmatrix}0 & (m^D_{\alpha j})^T\\ m^D_{i\beta} & M_{ij}\end{pmatrix}\begin{pmatrix}\nu_{\beta L}\\ (N_{jR})^c\end{pmatrix} \tag{1.30}$$

The mass matrix can then be diagonalized:

$$\mathcal{M} = \begin{pmatrix} 0 & (m^D_{\alpha j})^T \\ m^D_{i\beta} & M_{ij} \end{pmatrix} \longrightarrow \begin{pmatrix} D^\nu_{ab} & 0 \\ 0 & D^N_{kl} \end{pmatrix} \tag{1.31}$$

where:

$$D^\nu_{ab} = m_a \delta_{ab} \qquad D^N_{ij} = M'_i \delta_{ij} \tag{1.32}$$

We now impose the limit that the Majorana masses M'_i are much larger than the electroweak scale $\sim v$, known as the *seesaw limit*. Assuming that the Majorana mass matrix M_{ij} is itself diagonalized, we can approximate the total mass matrix as:

$$\mathcal{M} = \begin{pmatrix} m_{\alpha\beta} & 0 \\ 0 & M'_{ij} \end{pmatrix} \tag{1.33}$$

For low-energy phenomenology, we only need to consider the light neutrino mass matrix:

$$m_{\alpha\beta} = -(m^D_{\alpha i})^T {M'}^{-1}_{ij} m^D_{\beta j} \tag{1.34}$$

The eigenvalues m_1, m_2, m_3 are the masses of the light neutrino mass eigenstates. The light neutrino masses are thus suppressed by the ratio of the electroweak scale to the Majorana mass scale, m^D_α / M'_i. In the seesaw limit, this suppression factor is very small. If we use an O(1) Yukawa coupling for the neutrinos such that m_D is of order 100 GeV, then a 10 meV light neutrino corresponds to a Majorana mass of order 10^{15} GeV. Until now we have made no assumptions about how many right-handed neutrino fields are contained in N_{iR}. This mechanism for generating a set of very light particles using a very large mass scale is called the seesaw mechanism [43–45]. It turns out that two options present themselves here: a model with two right handed neutrinos (2RHN) and a model with three right handed neutrinos (3RHN), so M_{ij} can be either 2×2 or 3×3 [46, 47]. In the 3RHN, everything outlined above regarding neutrino oscillations still holds. The principle difference in the 2RHN is that because there is one fewer right-handed neutrino field, one of the two Majorana phases in the PMNS matrix (see Equation 1.14) can be set to zero. It is conventional to set α_1 to zero and let $\alpha = \alpha_2$ be the sole Majorana phase.

It should be stressed that this mechanism, known as a type I seesaw, is not the only mechanism that could explain a Majorana neutrino. Several other models exist including the type II seesaw[48, 49] in which a small left-handed Majorana mass is generated via a scalar Higgs scalar triplet, and the type III seesaw[50, 51] which adds three gauge-triplet fermions to the theory. By its nature, a Majorana neutrino mass must break the conservation of lepton number (L), also called *lepton number violation*, no matter the mechanism at play. This is exactly one of the conditions we need for leptogenesis.

Zooming out a bit, the Dirac-neutrino model anticipates new physics at the electroweak scale, while the Majorana-neutrino model anticipates new physics at the scale of Majorana masses, which we have good reason to believe are very large. The Majorana-neutrino model readily explains the small neutrino masses without invoking vastly different Yukawa couplings, instead invoking a higher energy scale at which we might expect new physics. If these masses are at sufficiently high enough scales, we could expect these Majorana masses to play a role in grand unified theories [52–54], provide a reasonable mechanism for leptogenesis, and explain why the light neutrinos are in fact so small. This makes studying the Majorana nature of neutrinos an inviting task for physicists today.

A question that still persists is how to build an experiment that can probe the seemingly natural Majorana mass value of 10^{15} GeV. The Large Hadron Collider at CERN has operated at energy scales up to 13 TeV [55], and the ICECUBE experiment has recorded neutrino events with energies $\gtrsim 1$ PeV [56] but even these energies leave us with six orders of magnitude still to surmount. It so happens that a good place to look for the effects of a Majorana mass is in the energy spectra of beta decays. This is the search for neutrinoless double beta decay ($0\nu\beta\beta$), and it is the overarching science goal of this dissertation. We must now turn our attention towards nuclear physics for a while and discuss $0\nu\beta\beta$, how it provides a direct way to verify the Majorana nature of the neutrinos, and how one can build an experiment to search for it.

Chapter 2

Beta Decay

Beta decay is a nuclear process described by the reaction:

$$ {}^{A}_{Z}X \longrightarrow {}^{A}_{Z+1} X' + e^- + \overline{\nu}_e \tag{2.1} $$

Beta (β) decay played a key role in the discovery of the neutrino. When the first measurements of β decay spectra were conducted, there was no knowledge of neutrinos. The continuous nature of the β decay spectrum (see Fig. 2.1a) was unexplained by the contemporary theories of β decay until Pauli introduced a neutrino to the theory. In Fermi's theory of β decay, the neutrino is a chargeless, massless spin-$\frac{1}{2}$ particle. It is instructive to start by analyzing β decay from first principles. To model the dynamics of β decay, we start in the rest frame of the parent nucleus. We assume that the three-momenta of the daughter nucleus ($\vec{p_d}$), electron ($\vec{p_e}$), and (anti)neutrino ($\vec{p_\nu}$) are all non-negligible. This means that the electron and anti-neutrino can each be emitted in any direction with one constraint:

$$ \vec{p_N} = \vec{p_e} + \vec{p_\nu} \tag{2.2} $$

We next assume that since $m_e \ll m_N$, the kinetic energy of the nucleus is negligible. Therefore, energy conservation implies:

$$ Q \equiv M_p - M_d - m_e = T_e + T_v \tag{2.3} $$

where T_i is the kinetic energy of particle i, M_p is the mass of the parent nucleus, M_d is the mass of the daughter nucleus. The Q-value is fixed for a given value of (A, Z). To find the β decay rate, we use Fermi's golden rule:

$$ \Gamma_{i\to f} = \frac{2\pi}{\hbar} F(Z, T_e) \left| \langle i|\hat{\mathcal{O}}|f\rangle \right|^2 \iint dN_\nu dN_e \delta(Q - T_e - T_\nu) \tag{2.4} $$

Here $F(Z, T_e)$ is the Fermi function, a correction term which models the Coulomb interaction between the emitted electron (or positron in the case of β^+ decay) and the daughter nucleus. For now, we will ignore it. The proceeding integral is the phase space factor for the decay rate. To solve it, we consider the density of states for a discrete lattice of momentum states in a box of volume V:

$$dN = \frac{4\pi V}{(2\pi\hbar)^3} p^2 dp$$

and we apply the appropriate relativistic transformations assuming a massless neutrino:

$$p_e = \sqrt{T_e^2 + 2m_e T_e} \tag{2.5}$$

$$dp_e = dT_e \frac{T_e + m_e}{\sqrt{T_e^2 + 2m_e T_e}} \tag{2.6}$$

$$dp_\nu = dT_\nu \tag{2.7}$$

The decay rate is therefore:

$$\Gamma_{i\to f} = \frac{2\pi}{\hbar} \left| \langle i|\hat{\mathcal{O}}|f\rangle \right|^2 \frac{16\pi^2 V^2}{(2\pi\hbar)^6} \int_0^Q dT_e \,(Q - T_e)^2 \,(T_e + m_e) \sqrt{T_e^2 + 2m_e T_e} \tag{2.8}$$

This integral scales as Q^5 to leading order, a fact which is known as Sargent's rule [57]. This can also be deduced solely from the kinematics of the decay. β decay is a "1-to-3" decay in which one particle decays into three particles. Each product particle has a unique momentum in each direction and energy, giving four equations per product particle. The on-shell condition requires that $E^2 = |\vec{p}|^2 + m^2$, eliminating one equation per product particle, while conservation of energy and momentum eliminates four more equations. In total, this gives $3N_P - 4$ unconstrained equations, where N_P is the number of product particles. For $N_P = 3$, this gives five unconstrained equations, so we expect the phase space integral to scale as Q^5. Sargent's rule is useful to keep in mind during isotope selection for an experiment, as higher values of Q give a faster decay rate and typically higher signal-to-background ratios. We are often interested in the electron kinetic energy spectrum, $\mathrm{d}\Gamma/\mathrm{d}T_e$, because we can measure it directly.

The nuclear matrix element $\mathcal{M} = \langle i|\hat{\mathcal{O}}|f\rangle$ is the subject of much study in β decay and double beta ($\beta\beta$) decay, and we will come back to it soon. For now, we will group it in with the constant of proportionality:

$$\frac{\mathrm{d}\Gamma}{\mathrm{d}T_e} \propto (Q - T_e)^2 \,(T_e + m_e) \sqrt{T_e^2 + 2m_e T_e} \tag{2.9}$$

When we account for a neutrino mass m_β, this becomes:[1]

$$\frac{d\Gamma}{dT_e} \propto (Q - T_e)\,(T_e + m_e) \sqrt{\left(T_e^2 + 2m_e T_e\right)\left[(Q - m_\beta - T_e)^2 + 2m_\beta\,(Q - m_\beta - T_e)\right]} \tag{2.10}$$

---

[1] Here we define Q such that $Q \equiv M_p - M_d - m_e$ still holds. This means $Q = T_e + T_\nu + m_\beta$. This function is 0 at $T_e = 0$ and $T_e = Q - m_\beta$. This is a helpful way to keep Q fixed as we vary the neutrino mass, but still enforce an endpoint of $Q - m_\beta$.

where m_β is calculated using the PMNS matrix:

$$m_\beta^2 = \sum_i |U_{ei}|^2 m_i^2 \tag{2.11}$$

This simplified model does not account for several effects. The Fermi function, which we have ignored, has a large effect on the overall shape of $\mathrm{d}\Gamma/\mathrm{d}T_e$, but it does not play a large role at the endpoint, $T_{max} = Q - m_\nu$ (see Fig. 2.1b). The non-zero neutrino mass does play a large role in the spectral shape at $T_e \approx Q$, where the shape is distorted and the endpoint is shifted to the left by m_β (see Fig. 2.1c). This presents an interesting opportunity to make a direct measurement of the neutrino masses by analyzing β decay events near the Q-value. This is done by measuring the kinetic energy of the electrons produced by β decay and fitting the energy spectrum near the endpoint to the spectral shape function while floating the endpoint. The KATRIN experiment [58] has conducted such an analysis using tritium decay:

$${}^3_1\mathrm{T} \longrightarrow {}^3_2\,\mathrm{He} + e^- + \overline{\nu}_e \tag{2.12}$$

to set a limit of $m_\beta < 0.45$ eV (90% C.L.) [59], and the experiment has a projected lifetime sensitivity of 0.2 eV.

Theoretically, because β decay produces a neutrino in a flavor eigenstate, ν_e, there is a non-zero probability of finding the emitted neutrino in each of the three neutrino mass eigenstates when reconstructing the electron energies. The probability of finding the neutrino in mass state i is equal to U_{ei}^2. As a result, the true spectrum from β decay is a sum of three spectra, each with different endpoints. This produces a kink in the spectrum (see Fig. 2.1d). In practice, this effect is small enough that the KATRIN experiment cannot resolve this kink, and this effect can be ignored. However, a similar effect due to the existence of sterile neutrinos could cause a more prominent kink in the spectrum away from the endpoint, at $T_e = Q - m_s$, where m_s is the mass of the sterile neutrino, and this effect could potentially be measured by KATRIN.

2.1 Nuclear Matrix Elements in β decay

Nuclear matrix elements (NMEs) are mathematical constructs that describe the probability amplitude for a nuclear transition. They are derived from the wavefunctions of the initial and final nuclear states and the relevant transition operator. The NME can be written as:

$$\mathcal{M} = \langle i|\hat{\mathcal{O}}|f\rangle \tag{2.13}$$

where $\langle f|$ and $|i\rangle$ are the final and initial nuclear states, respectively, and $\hat{O}$ is the transition operator. In the context of β decay, the transition operator includes contributions from both the vector and axial-vector currents, reflecting the weak interaction's nature.

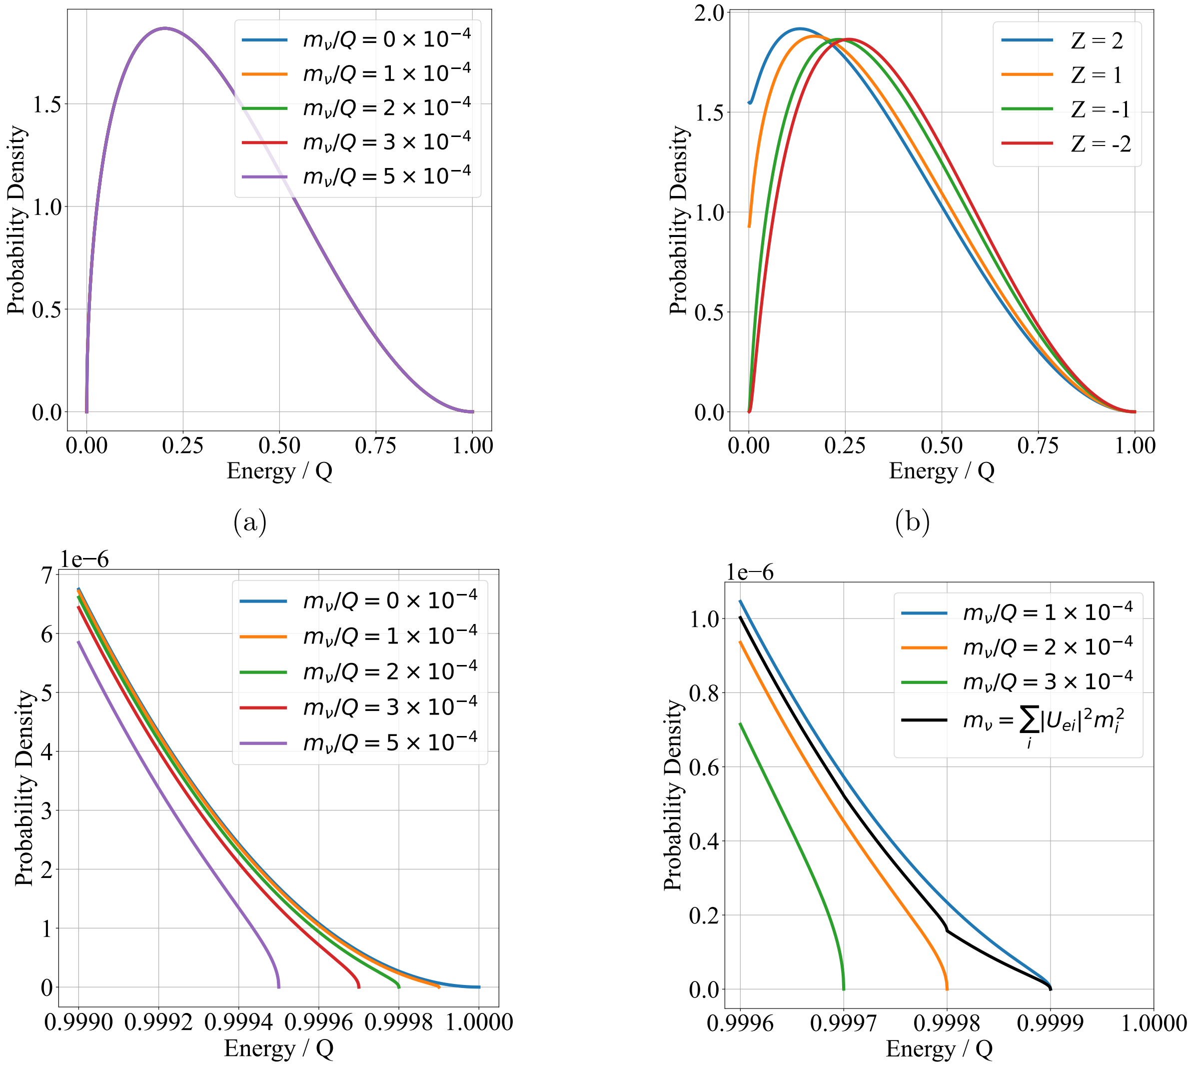

(a) (b) (c) (d)

Figure 2.1: Electron kinetic energy spectra from β decay with an endpoint $Q \ll m_e$. (a) Kinetic energy spectra using only the phase space factor. The vast majority of the spectrum is unaffected by varying a very small neutrino mass. (b) Effects of the Fermi function $F(Z, T_e)$ on the electron kinetic energy spectrum. Distributions with $Z < 0$ are the spectra of positrons emitted during β^+ decay. (c) Energy spectra at the β decay endpoint Q with different neutrino mass values assuming the electron neutrino is in a mass eigenstate. (d) Energy spectra including the spectrum with a neutrino flavor eigenstate which is a linear combination of three mass eigenstates. This spectrum, with kinks at m_2 and m_3, is most similar to what is expected for a true β decay that produces an electron neutrino.

Fermi matrix elements describe the transition involving the vector current. They assume no spin change between the initial and final states, and they are relevant for super-allowed Fermi transitions for which isospin changes but parity does not. Fermi matrix elements are typically written as:

$$\mathcal{M}_F = \langle f|\sum_k \tau_k^+|i\rangle \tag{2.14}$$

where τ_k^+ is the isospin raising operator and k indexes over each of the nucleons in the nucleus.

Gamow-Teller matrix elements, on the other hand, describe the transition involving the axial vector current. These operators account for changes in both spin and isospin between the initial and final states, and they are relevant for transitions where both spin and parity can change. Gamow-Teller matrix elements are typically written as:

$$\mathcal{M}_{GT} = \langle f|\sum_k \sigma_k\tau_k^+|i\rangle \tag{2.15}$$

where σ_k is the Pauli spin operator.

It is instructive to compute a simple NME for the case of β decay of tritium. This decay is of particular interest because tritium and helium-3 are mirror nuclei, i.e. the number of protons in one nucleus is equal to the number of neutrons in the other nucleus, and vice-versa. This allows one to simplify the calculation by assuming that the spatial wave-functions of the nuclei are identical and focusing on isospin. This reduces equation 2.14 a single term in the Fermi matrix element:

$$\mathcal{M}_F = \langle {}^3\mathrm{He}|\tau^+|{}^3\mathrm{H}\rangle = \sqrt{2} \tag{2.16}$$

This value comes from the Clebsch-Gordan coefficients for the transition of an isospin $I_3 = \frac{1}{2}$ state. The Gamow-Teller matrix element is also a single term:

$$\mathcal{M}_{GT} = \langle {}^3\mathrm{He}|\sigma\tau^+|{}^3\mathrm{H}\rangle = \sqrt{3} \tag{2.17}$$

which can be derived from the combination of spin and isospin transition probabilities. The total transition probability involves both $\mathcal{M}_F$ and $\mathcal{M}_{GT}$:

$$|\mathcal{M}|^2 = g_V^2|\mathcal{M}_F|^2 + g_A^2|\mathcal{M}_{GT}|^2 \tag{2.18}$$

where g_V and g_A are the vector and axial-vector coupling constants, respectively. For the weak interaction, $g_V \approx 1$ and $g_A \approx 1.27$. Plugging in these values, we have:

$$|\mathcal{M}|^2 = (\sqrt{2})^2 + 1.27^2 \cdot (\sqrt{3})^2 \approx 6.83 \tag{2.19}$$

This is a very simplified calculation, and more accurate calculations include corrections based on nuclear structure [60]. In practice, the Gamow-Teller NME is approximately $0.95\sqrt{3}$ [61]. Computing NMEs accurately is difficult even for small nuclei such as tritium. For large nuclei, this is an even greater challenge since the computational complexity of the problem grows substantially with the number of nucleons in the system. Therefore, one must make approximations and use different models to compute the NMEs while considering the tradeoff between accuracy and computational complexity. Several techniques are employed to compute NMEs for large nuclei, including:

- *Ab initio* Methods: These utilize first-principles calculations based on fundamental interactions between nucleons, aiming for highly accurate results without relying on phenomenological parameters. *Ab initio* methods are computationally intensive but offer a deep understanding of the underlying nuclear forces.

- Shell Model: In this model, the wavefunctions of the nucleons are well-described. The model effectively accounts for the discrete energy levels that nucleons occupy within a nucleus, providing a clear picture of the nuclear structure. However, the interactions between nucleons are described by a limited set of effective operators, known as residual interactions, which account for deviations from an independent particle model. This method is well-suited for describing the structure of nuclei with closed shells ($Z, N =$ 2, 8, 20, 28, 50, 82, or 126). but becomes increasingly complex for nuclei with partially filled shells.

- Quasiparticle Random Phase Approximation (QRPA): This extends the shell model by including collective excitations of the nucleus. The wavefunctions are described by considering quasiparticles that arise from the pairing correlations in the nucleus. The quasiparticle wavefunctions are constructed by solving the Hartree-Fock-Bogoliubov equations, which take into account both mean-field potentials and pairing interactions. QRPA then describes collective excitations as coherent superpositions of quasiparticle excitations, providing a more comprehensive description of the nuclear wavefunctions, especially for nuclei with open shells. QRPA can handle correlations between nucleons more effectively than the shell model alone.

- Nuclear Mean-Field Models: These models, which also include the Hartree-Fock-Bogoliubov method, consider mean-field potentials to describe nucleon interactions. Mean-field models simplify the problem by averaging the effects of all nucleons on a single nucleon, making the calculations more tractable. However, they usually need to be extended to include correlations and fluctuations beyond the mean field to improve accuracy.

Each of these techniques has its benefits and drawbacks, and, importantly, the results substantially differ between the different techniques. Fig. 2.2 shows the NMEs for various $0\nu\beta\beta$ candidate isotopes calculated using many different methods. While these techniques have not converged, recent developments in *ab initio* methods, which are typically seen as

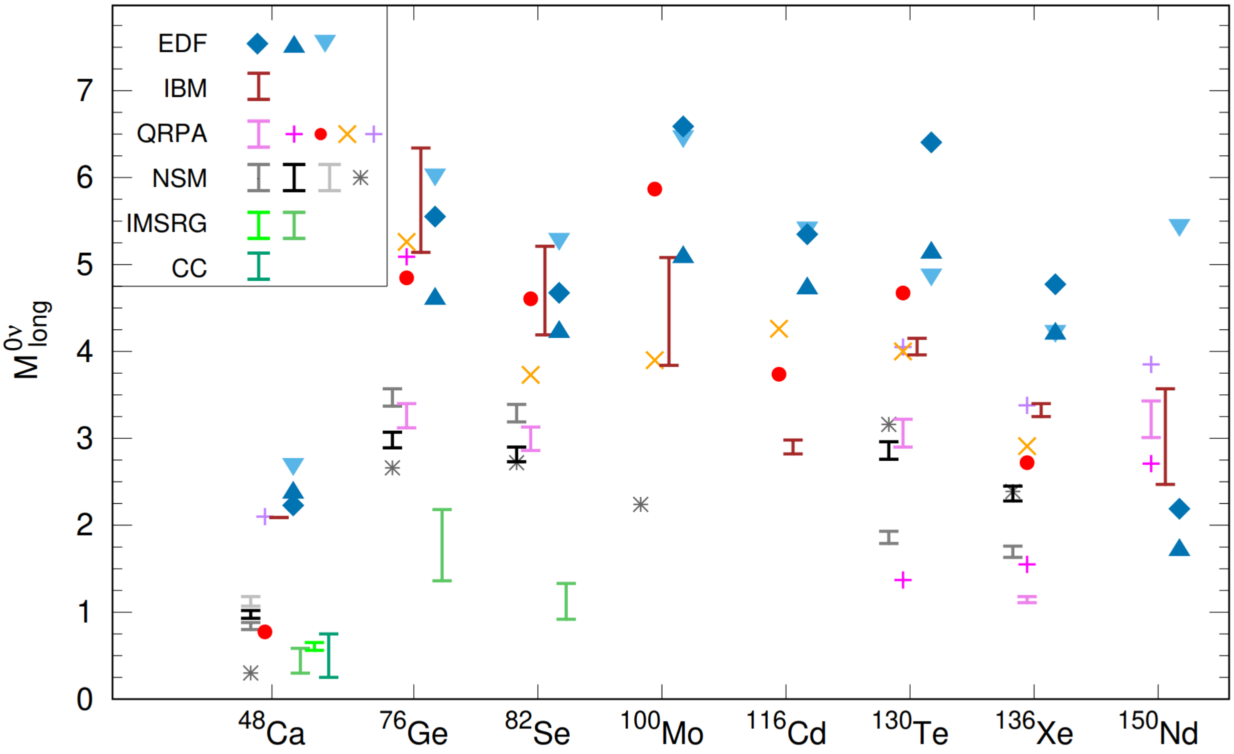

Figure 2.2: Nuclear matrix elements for light-neutrino exchange from different many-body methods. Image reprinted from [62].

the most accurate, have given insight into the shortcomings of different techniques and are challenging the many-body methods currently used in the field. As these insights develop and the theoretical work continues, we can expect to see the accuracy of NME calculations improve in the coming years.

2.2 Double Beta Decay

In 1935, Maria Goeppert-Mayer hypothesized the existence of two-neutrino double beta decay ($2\nu\beta\beta$) [63]:

$$ {}^{A}_{Z}X \longrightarrow {}^{A}_{Z+2}X' + 2e^{-} + 2\overline{\nu}_{e} \tag{2.20}$$

$2\nu\beta\beta$ is a second-order weakly-mediated Standard Model process that conserves lepton number. The Feynman diagram for the process is shown in Fig. 2.3a. While this process is allowed in the Standard Model, its decay rate is highly suppressed due to the small nature of the Fermi constant. Its half-life can be expressed as:

$$\frac{1}{T_{1/2}^{2\nu}} = G^{2\nu} g_A^4 \left|\langle \mathcal{M}^{2\nu} \rangle\right|^2 \tag{2.21}$$

where $G^{2\nu}$ is the phase space factor that scales as Q^{11}, g_A is the axial vector coupling, and $\mathcal{M}^{2\nu}$ is the nuclear matrix element. Note that the phase space integral for $2\nu\beta\beta$ is much more complicated than the one shown in equation 2.4. Unlike ordinary β decay, $2\nu\beta\beta$ is a 1-to-5 decay. Counting the kinematic degrees of freedom as we did for β decay, the exponent on the leading term of the phase space integral is $3N_P - 4 = 11$.

The axial vector coupling for a free nucleon is known with high precision to be ≈ 1.27 [64, 65]. In heavy nuclei, however, quenching effects may significantly change the effective value of g_A. Understanding these quenching factors is an ongoing pursuit of both experimentalists and theorists. For isotopes of interest for $\beta\beta$ decay ($20 \leq A \leq 250$), the quenching g_A^{eff}/g_A^{free} can range between between 0.3 and 1 [66]. This depends both on the analysis one chooses to endorse and on the isotope in question. The value of the quenching for each $\beta\beta$ decaying isotope is of great experimental importance because of the g_A^4 dependence in the decay rate.

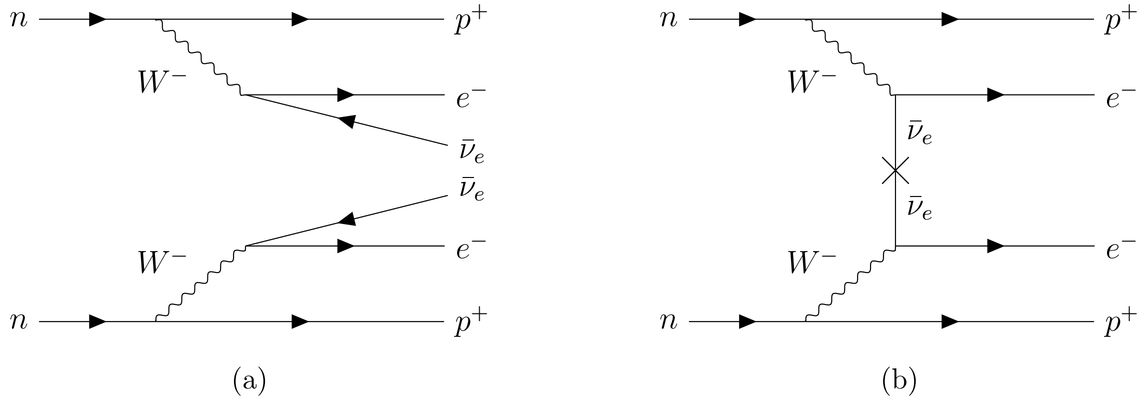

Figure 2.3: (a) A Feynman diagram for $2\nu\beta\beta$. (b) A Feynman diagram for $0\nu\beta\beta$.

The suppressed decay rate means that most radionuclei with fast β decay rates will completely wash out any $2\nu\beta\beta$ signal, making it very difficult to observe. There exist isotopes for which single β^- (or β^+) transitions are forbidden or highly suppressed, such that $2\nu\beta\beta$ is the dominant decay process. Table 2.1 lists the current values of the half-life of $2\nu\beta\beta$ in several of these isotopes. The half-lives of these decays range from $10^{18} - 10^{21}$ years.[2]. Notably, all the listed isotopes are "even-even," i.e. they have even numbers of protons and neutrons. This is because nuclei in which protons and neutrons can pair up amongst themselves tend to have higher binding energies and, thus, lower masses than their odd-odd counterparts. In the semi-empirical mass model, or liquid drop model [68, 69], this is modeled as a pairing

---

[2] To date, the slowest process ever measured is another second-order weakly-mediated process, the two-neutrino double electron capture (2νECEC) of ^{124}Xe, with a half-life of $(1.8 \pm 0.5) \times 10^{22}$ years [67]

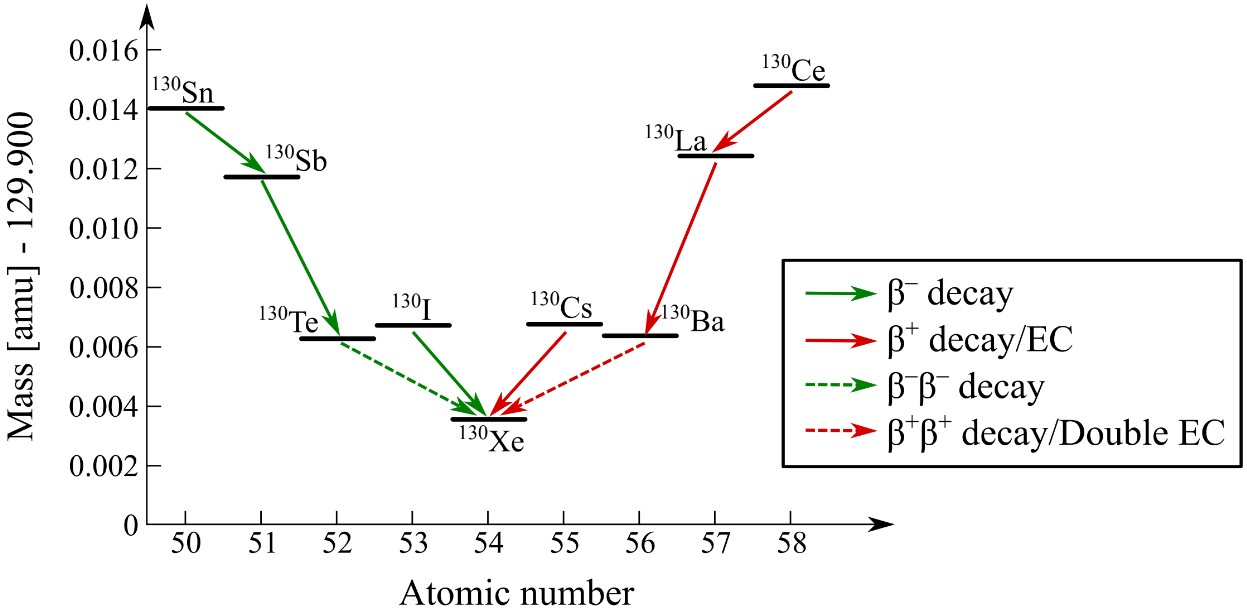

Figure 2.4: Masses of nuclei with $A = 130$ and their allowed decay modes. Because ^{130}I is heavier than ^{130}Te, the single β decay of ^{130}Te to ^{130}I is forbidden, and the dominant process is $2\nu\beta^-\beta^-$ to ^{130}Xe. Similarly, the decay of ^{130}Ba to ^{130}Cs via electron capture or β^+ decay is energetically forbidden, and $2\nu\beta^+\beta^+$ is favored. Image reprinted from [70].

term in the binding energy:

$$\Delta E_B = \begin{cases} \delta_0 A^\gamma & \text{if } Z, N \text{ both even} \\ 0 & \text{if } A \text{ odd} \\ -\delta_0 A^\gamma & \text{if } Z, N \text{ both odd} \end{cases} \tag{2.22}$$

where γ is typically chosen to be -0.5 or -0.75. The masses of several nuclei from the $A = 130$ isobar are shown in Fig. 2.4. This isobar is particularly interesting because two allowed $\beta\beta$ decay modes exist. ^{130}Te can undergo $2\nu\beta^-\beta^-$:

$$^{130}_{52}\text{Te} \longrightarrow ^{130}_{54}\text{Xe} + 2e^- + 2\overline{\nu}_e \tag{2.23}$$

and ^{130}Ba can undergo $2\nu\beta^+\beta^+$, double electron capture (2νECEC), or β^+/EC:

$$^{130}_{56}\text{Ba} \longrightarrow ^{130}_{54}\text{Xe} + 2e^+ + 2\nu_e \tag{2.24}$$

$$^{130}_{56}\text{Ba} + 2e^- \longrightarrow ^{130}_{54}\text{Xe} + 2\nu_e \tag{2.25}$$

$$^{130}_{56}\text{Ba} + e^- \longrightarrow ^{130}_{54}\text{Xe} + e^+ + 2\nu_e \tag{2.26}$$

Table 2.1: Averages of measured $2\nu\beta\beta$ half-lives in different isotopes from current-generation experiments. Values are taken from [71] using averaging methods from [72].

Isotope	$T^{2\nu}_{1/2}$ (year)	Experiments
^{48}Ca	$\left(5.3^{+1.2}_{-1.8}\right) \times 10^{19}$	Irvine TPC [73], TGV [74], NEMO-3 [75]
^{76}Ge	$\left(1.88^{+0.08}_{-0.07}\right) \times 10^{21}$	HEIDELBERG-MOSCOW [76], GERDA [77]
^{82}Se	$\left(0.87^{+0.02}_{-0.01}\right) \times 10^{20}$	NEMO-3 [78],CUPID-0 [79], Irvine TPC [80], NEMO-2 [81]
^{96}Zr	$(2.3 \pm 0.2) \times 10^{19}$	NEMO-2 [82], NEMO-3 [83]
^{100}Mo	$\left(7.06^{+0.15}_{-0.13}\right) \times 10^{18}$	NEMO-3 [84], CUPID-Mo [85], NEMO-2 [86], Irvine TPC [87], $ZnMoO_4$ bolometers [88]
^{116}Cd	$(2.69 \pm 0.09) \times 10^{19}$	NEMO-3 [89], Aurora [90], ELEGANT [91], Solotvina [92], NEMO-2 [93]
^{130}Te	$(7.91 \pm 0.21) \times 10^{20}$	CUORE-0 [94], CUORE [95], CUORICINO [96], NEMO-3 [97]
^{136}Xe	$(2.18 \pm 0.05) \times 10^{21}$	EXO-200 [98], KamLAND-Zen [99]
^{150}Nd	$(9.34 \pm 0.65) \times 10^{18}$	NEMO-3 [100]

2.3 Neutrinoless Double Beta Decay

In 1939, Furry [101] synthesized the idea of $\beta\beta$ decay with the idea of the Majorana neutrino to hypothesize the existence of neutrinoless double beta decay ($0\nu\beta\beta$):

$$^{A}_{Z}X \longrightarrow^{A}_{Z+2} X' + 2e^- \tag{2.27}$$

Fig. 2.3b shows the Feynman diagram for this process. In $0\nu\beta\beta$, the two neutrinos present in the $2\nu\beta\beta$ final state are instead virtually exchanged. This is particularly enticing to particle physics because it violates lepton number ($\Delta\mathsf{L} = 2$). As discussed in section 1.2.3, the presence of a Majorana mass guarantees a lepton-number violating process regardless of the mass generation mechanism. If massive Majorana neutrinos exist, then so should $0\nu\beta\beta$. We can still ask the converse, though: does the observation of $0\nu\beta\beta$ prove the existence of massive Majorana neutrinos? The Schechter-Valle theorem [102] says yes.[3] The theorem states that all of the beyond Standard Model physics at play in $0\nu\beta\beta$ can be put into a black box, then the $0\nu\beta\beta$ diagram can be rearranged into a diagram for the process $\overline{\nu} \longrightarrow \nu$ (see Fig. 2.5). The observation of $0\nu\beta\beta$ proves the existence of a Majorana mass term. Therefore, if Majorana neutrinos exist, then we can expect $0\nu\beta\beta$ to be observed except in a particular case of fine-tuning.

As discussed in section 1.2.1, the neutrino mass eigenstates are different from their flavor eigenstates. Using the light neutrino exchange model, we sum over the electron projections

---

[3] This theorem considers $0\nu\beta\beta$ only in the framework of a "natural" SU(2)×U(1) gauge theory. There are of course many other possibilities for a beyond Standard Model gauge group, but many of the common choices still contain SU(2)×U(1) as a subgroup.

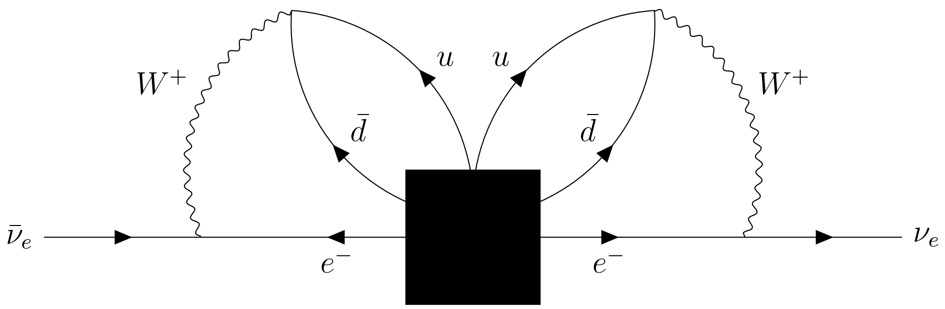

Figure 2.5: A Feynman diagram demonstrating the Schechter-Valle theorem.

of the mass eigenstates m_i to define the effective Majorana mass $m_{\beta\beta}$ as [103]:

$$m_{\beta\beta} \equiv \left| \sum_{i=1}^{i=3} U_{ei}^2 m_i \right| \tag{2.28}$$

Note the difference between this equation and Equation 2.11. m_β is a quadrature sum of the mass-weighted probability amplitudes, while $m_{\beta\beta}$ is a mass-weighted linear sum of three complex amplitudes. Expanding this sum we have:

$$m_{\beta\beta} = \left| c^2\theta_{12} c^2\theta_{13} e^{i\alpha_1} m_1 + s^2\theta_{12} c^2\theta_{13} e^{i\alpha_2} m_2 + s^2\theta_{13} e^{-2i\delta_{CP}} m_3 \right| \tag{2.29}$$

It is possible that these three terms sum to zero, but this requires that the parameters of the PMNS matrix be finely tuned. Our knowledge of θ_{ij} and Δm_{ij}^2 constrains the allowed range of $m_{\beta\beta}$ for a given lightest neutrino mass $m_{lightest}$ (see Fig. 2.6). The fine-tuning scenario in which $m_{\beta\beta} = 0$ is only allowed given the normal ordering and 1 meV $< m_{lightest} <$ 10 meV. Note that even if both of these conditions are met, $m_{\beta\beta}$ need not be 0. If no fine-tuning exists, then the nonzero values of $m_{\beta\beta} >$ 1 meV are overwhelmingly more likely assuming all Majorana phases are equally likely [104].

It is important to keep in mind that at the end of the day, experiments are only sensitive to the half-life (or the decay rate) of $0\nu\beta\beta$. Any mapping from the half-life to $m_{\beta\beta}$ will be model-dependent. As experimentalists, we often like to quote experimental sensitivities to $0\nu\beta\beta$ in terms of $m_{\beta\beta}$, as we do in Fig. 2.6. This is useful because it gives us a way to compare half-lives across different isotopes. However, doing so is highly model-dependent. The experimental community has generally adopted the practice of assuming light Majorana neutrino exchange as the dominant mechanism responsible for $0\nu\beta\beta$. Doing so allows us to express the half-life of $0\nu\beta\beta$ as:

$$\frac{1}{T_{1/2}^{0\nu}} = G^{0\nu} g_A^4 \left| \langle \mathcal{M}^{0\nu} \rangle \right|^2 \frac{m_{\beta\beta}^2}{m_e^2} \tag{2.30}$$

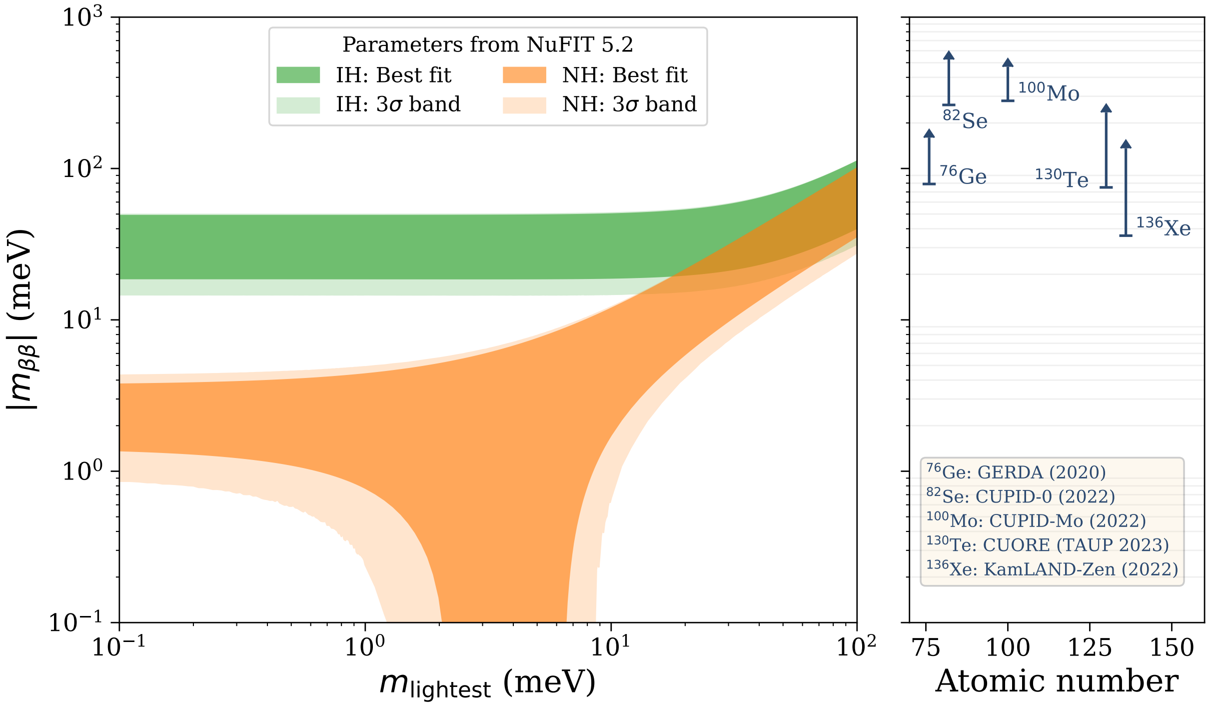

Figure 2.6: Regions of allowed values of $m_{\beta\beta}$ versus the lightest neutrino mass $m_{lightest}$ for both neutrino mass orderings. The right panel shows the current experimental limits on $m_{\beta\beta}$ for different isotopes. All data shown assume that $0\nu\beta\beta$ is mediated via light neutrino exchange. The range of each experimental limit is due to the uncertainty of the $0\nu\beta\beta$ nuclear matrix element. Image made with the LobsterPlot software package, available courtesy of J. Torres.

which is similar to 2.21, but with a different phase factor since $0\nu\beta\beta$ is a one-to-three decay, a different matrix element, and of course the Majorana mass dependence. Even after adopting this convention, the model-dependence does not end here. The nuclear matrix element $\mathcal{M}^{0\nu}$ is also not agreed upon by modern theories or computational models, thus there is still an uncertainty in $m_{\beta\beta}$ depending on the model one chooses. In practice, we tend to select the most modern from a running list of matrix element calculations using different techniques and report a range of possible values of $m_{\beta\beta}$ given a particular half-life.

2.3.1 Searching for Neutrinoless Double Beta Decay

Experimental Sensitivity and Scaling

The final state of $0\nu\beta\beta$ contains, unsurprisingly, no neutrinos. Therefore, the kinetic energy of the entire interaction must be carried away by the two outgoing electrons (up to the recoil energy of the nucleus, which is negligible in β decay). The experimental signature of this process is thus a sharp peak at $E = Q_{\beta\beta}$ in the spectrum of all $\beta\beta$ decay events. After selecting a candidate isotope to search for such a peak, the goal is to maximize the number of $\beta\beta$ events observed while minimizing the number of observed background events. The number of observed $0\nu\beta\beta$ events is:

$$n_{0\nu} = n_0 \left(1 - e^{-t \ln 2 / T^{0\nu}_{1/2}}\right) \approx \frac{n_0 t \ln 2}{T^{0\nu}_{1/2}} \tag{2.31}$$

where we have assumed that the $0\nu\beta\beta$ half-life is very long (an extremely reasonable thing to do in this context). Here n_0 is the initial number of candidate isotope nuclei being observed. We often rewrite this number as:

$$n_0 = \frac{m}{m_A} f \tag{2.32}$$

where m_A is the atomic mass of the candidate isotope, m is the mass of the detector, and f is the enrichment factor.[4] Assuming a signal detection efficiency of the detector ϵ (including all fiducialization and analysis efficiencies), the expected number of signal counts is:

$$\hat{n}_{0\nu} = \frac{f \epsilon_s m t}{m_A} \tag{2.33}$$

The numerator is the time-integrated amount of candidate isotope observed by the experiment. We call this quantity the exposure, and we will often use it when referring to the sizes of particular datasets. In an experimental setup, background events and background-like events can be induced by many sources including cosmogenic events, radioactive contamination, and noise. When searching for $0\nu\beta\beta$, we define a region of interest (ROI) of width ΔE centered around $Q_{\beta\beta}$. Typically ΔE is chosen to be the full-width at half maximum (FWHM) energy resolution of the experiment when quoting sensitivity values, but actual analyses use much wider windows when fitting for the $0\nu\beta\beta$ signal. The number of expected background events is:

$$\hat{n}_b = \mathrm{BI} \cdot \Delta E \cdot \epsilon_b \cdot mt \tag{2.34}$$

where BI is the background index, defined in units of counts/(Energy·Mass·Time) and ΔE is the energy resolution, typically quoted as the standard deviation σ, or the FWHM. Note that we have distinguished between the background event efficiency ϵ_b and the signal event efficiency ϵ_s. This reflects the fact that experiments can apply signal/background rejection techniques (e.g. pulse shape discrimination, light yield discrimination, etc.) to select for supposed signal events while rejecting background events. Increasing the ratio ϵ_s/ϵ_b is an ongoing effort across the experimental community. One definition of the sensitivity to $T^{0\nu}_{1/2}$ is the half-life required such that the number of signal events is equal to the uncertainty on the number of background events. Assuming the number of background events is Poisson distributed, we have:

$$\hat{n}_{0\nu} = \sqrt{\hat{n}_b} \tag{2.35}$$

---

[4] It is often of value to further rewrite the enrichment factor as $f' = f/a_{nat}$ where a_{nat} is the natural abundance of the particular isotope. This is useful because it scales well with the most fundamental unit, dollars.

$$\frac{f\epsilon_s mt\ln 2}{m_A T_{1/2}^{0\nu}} = \sqrt{\mathrm{BI}\cdot\Delta E\cdot\epsilon_b\cdot mt} \tag{2.36}$$

$$T_{1/2}^{0\nu} = \ln 2\frac{f\epsilon_s}{m_A\sqrt{\epsilon_b}}\sqrt{\frac{mt}{\mathrm{BI}\cdot\Delta E}} \tag{2.37}$$

In the background-free limit, the sensitivity scales proportionally to the number of expected signal events:

$$T^{0\nu}((\mathrm{BI})=0) = \ln 2\frac{f\epsilon_s}{m_A}mt \tag{2.38}$$

Current Experimental Framework

One may take many technological approaches when designing and building an experiment to search for $0\nu\beta\beta$. Liquid cryogenic detectors – such as those used in the EXO-200 experiment [105], its successor nEXO [106], and the next-generation NEXT experiment [107] – employ large volumes of liquid xenon cooled to cryogenic temperatures. These detectors leverage the scintillation light and ionization signals produced by particle interactions to

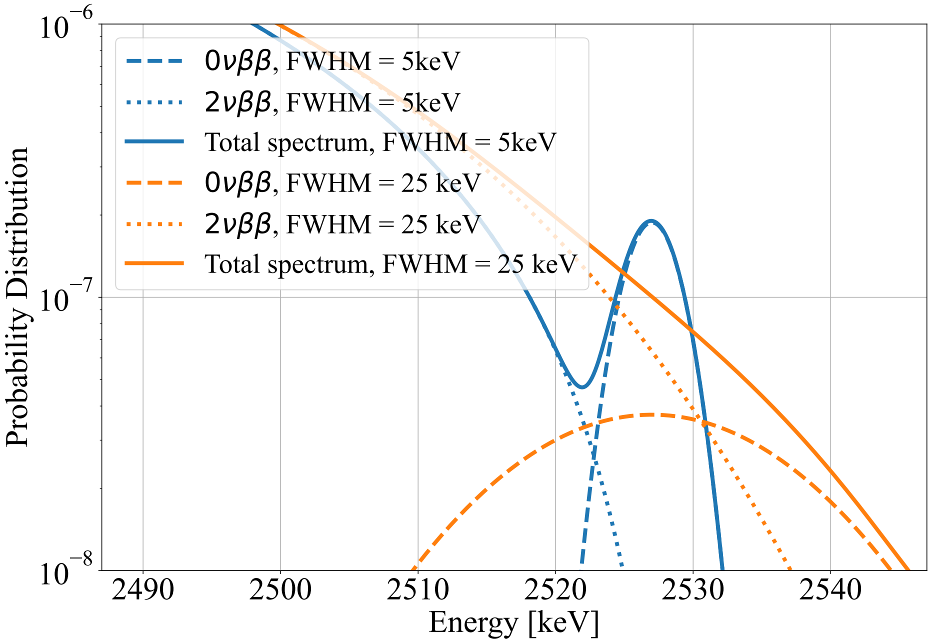

Figure 2.7: Predicted spectra of $2\nu\beta\beta$ and $0\nu\beta\beta$ near $Q_{\beta\beta} = 2527$ keV for one experiment with an FWHM energy resolution of 5 keV and one experiment with an FWHM energy resolution of 25 keV. No backgrounds are considered. As the resolution worsens, the number of $2\nu\beta\beta$ events in the ROI grows. The rate of $0\nu\beta\beta$ is assumed to be 10^{-6} times the $2\nu\beta\beta$ rate.

identify potential $0\nu\beta\beta$ events. Large liquid scintillator detectors, such as KamLAND-Zen [108] and SNO+ [109], use a liquid scintillator material to detect the light produced when a particle interacts with the liquid. This technology allows for the deployment of extremely large masses of the isotopes of interest, allowing for very large exposures within a reasonable data collection timeframe. Cryogenic solid-state calorimeters, such as those deployed in CUORE [110], CUPID-Mo [111], and AMoRE [112], utilize cryogenically cooled crystals of candidate isotopes. These detectors measure the tiny temperature rise caused by particle interactions by collecting phonons produced in the crystal lattice. Finally, solid-state semiconductor detectors, such as those used in the MAJORANA demonstrator [113], the GERDA experiment [114], and their successor LEGEND [115], measure the charge produced when ionizing radiation interacts with them. When a particle, such as an electron, interacts with the detector, it excites electrons from the valence band to the conduction band, creating electron-hole pairs. An electric field applied across the detector collects these charge carriers, resulting in a measurable electric current.

Each technology varies in its sensitivity and scalability. Liquid cryogenic detectors like EXO-200/nEXO offer energy resolutions around 1% at MeV scales and the capability to scale to large detector masses, enhancing sensitivity. However, the extensive cryogenic infrastructure required is complex and costly, posing significant engineering and operational challenges. Furthermore, as we enter the era of tonne-scale experiments, it will become challenging and costly to procure multiple tonnes of liquid xenon, especially after considering the enrichment costs. Large liquid scintillator detectors, such as SNO+, are easily scalable to large volumes, making them suitable for extensive, large-scale experiments. They generally offer moderate energy resolutions, around 5% at MeV scales. The large monolithic nature of these detectors enables them to be self-shielding, allowing them to have much lower intrinsic backgrounds than their counterparts. Ensuring uniform detector response across large volumes can also be challenging. Cryogenic calorimeters like those used in CUORE and AMoRE, sometimes referred to as bolometers, provide excellent energy resolution, about 0.1% at MeV scales, and their segmented nature allows for good background discrimination. However, scaling these detectors to larger volumes is challenging due to the complexity and cost of producing large numbers of high-purity crystals and maintaining cryogenic conditions. Additionally, the operation of cryogenic systems demands meticulous engineering and constant monitoring to ensure stable performance. Solid-state semiconductor detectors, typically made from high-purity Germanium, provide the best energy resolutions of the current technologies, and their extremely fast timing resolution allows for excellent background rejection using pulse shape discrimination techniques. One of the primary challenges associated with these detectors is the need for high-purity crystals made of enriched germanium, which are expensive and require specialized facilities for production and enrichment. The cryogenic systems needed to operate these detectors also introduce further complexity and cost.

Today, many experiments have reached the quasi-background-free regime in which the projected sensitivities to $0\nu\beta\beta$ are competitive, but the predicted number of background events in the ROI during the lifetime of the experiment is $\lesssim 1$. Even in the quasi-background free regime, $0\nu\beta\beta$ has an inescapable background associated with it: the $2\nu\beta\beta$ continuum near the endpoint[5] (see Fig. 2.7). This effect becomes worse as the detector resolution worsens, so the energy resolution becomes even more important than equation 2.37 may suggest.

The Past, Present, and Future of $0\nu\beta\beta$ Searches

The experimental search for $\beta\beta$ decay has developed a rich history over the past several decades. As mentioned earlier, Furry's proposal of $0\nu\beta\beta$ was only two years after Majorana's theory was published in 1937. Since then, the experimental pursuit has endured. For decades, the best limits on $0\nu\beta\beta$ were placed by geochemistry [116]. Excess quantities of products of radioisotopes in the Earth could be attributed to $\beta\beta$ decay, though this method is unable to distinguish between $2\nu\beta\beta$ and $0\nu\beta\beta$. The first evidence of $\beta\beta$ decay was found in ^{130}Te using the geochemical method, and the half-life was determined to be $1.4{\times}10^{21}$ years [117]. The first direct measurement of the half-life of $2\nu\beta\beta$ was made in 1987 [118]. The experiment employed a 1m^3 time projection chamber (TPC) with 14g of enriched ^{82}Se deposited on Mylar foil which composed the central electrode of the TPC. This experiment found evidence for $2\nu\beta\beta$ in ^{82}Se with a lifetime of $T^{2\nu}_{1/2} = 1.1^{+0.8}_{-0.3} \times 10^{20}$ yr. This is in good agreement with the current global best fit value of $0.87^{+0.02}_{-0.01} \times 10^{20}$. Today, the rates of $2\nu\beta\beta$ have been measured to at least 5σ precision in almost ten isotopes, as was shown in Table 2.1. $0\nu\beta\beta$ has proven more elusive than $2\nu\beta\beta$. In 1948, the first direct search for $0\nu\beta\beta$ in ^{124}Se set an upper bound of 3×10^{15} years on the half-life[119]. Now, three-quarters of a century after the first limits were placed, the best limits on $0\nu\beta\beta$ half-lives range from 10^{22} to 10^{28} years, and experimental searches continue to push these limits higher and higher. The current leading limits are listed in Table 2.2.

The best limit on $m_{\beta\beta}$ currently comes from the KamLAND-Zen experiment [108], which is located in Japan. The Yangyang Underground Laboratory in northeast South Korea is home to the AMoRE-I experiment, which has set the leading limit on the $0\nu\beta\beta$ half-life in ^{100}Mo. The EXO-200 experiment, which set the world-leading limit on the $0\nu\beta\beta$ half-life in ^{136}Xe at the time, and the MAJORANA demonstrator were both based in the United States. Europe is also home to two ongoing experiments, CUORE and LEGEND-200, which are located at the Laboratori Nazionali del Gran Sasso (LNGS) in Assergi, L'Aquila, Italy. LNGS was also home to the GERDA experiment which set the leading limit on $0\nu\beta\beta$ in ^{76}Ge [114]. This search is an active, international campaign. The rest of the 2020s are an optimistic time for the community. We hope to see something soon.

Contemporary $0\nu\beta\beta$ experiments are crossing the tonne-scale milestone, and the latest exclusion sensitivities are beginning to rule out sizable regions of the $m_{\beta\beta}$ parameter space, particularly when assuming the inverted ordering. The precise values of the neutrino oscillation parameters are solidifying, the upper limit on the sum of the neutrino masses is being pushed closer to the minimum possible value, and the neutrino ordering problem could easily

---

[5] In the limit of arbitrarily good resolution and no backgrounds, one could discriminate $2\nu\beta\beta$ events from $0\nu\beta\beta$ with absolute confidence. $2\nu\beta\beta$ decays have an endpoint that is shifted from the $0\nu\beta\beta$ endpoint by $2m_{\beta}$, allowing one to separate the two in principle. However, crafting a detector with the necessary resolution remains out of reach.

Table 2.2: Leading $0\nu\beta\beta$ half-life limits in different isotopes from current-generation experiments. Note that the listed leading limit in ^{130}Te excludes the latest result from CUORE, which is a major subject of the later chapters of this dissertation.

Isotope	$T^{0\nu}_{1/2}$ (years)	Experiment
^{48}Ca	$> 5.8 \times 10^{22}$	ELEGANT VI [120]
^{76}Ge	$> 1.8 \times 10^{26}$	GERDA [114]
^{82}Se	$> 4.6 \times 10^{24}$	CUPID-0 [121]
^{96}Zr	$> 9.2 \times 10^{21}$	NEMO-3 [83]
^{100}Mo	$> 3.0 \times 10^{24}$	AMoRE-I [112]
^{116}Cd	$> 2.2 \times 10^{23}$	Aurora [90]
^{128}Te	$> 3.6 \times 10^{24}$	CUORE [122]
^{130}Te	$> 2.2 \times 10^{25}$	CUORE [110]
^{136}Xe	$> 3.8 \times 10^{26}$	KamLAND-Zen [108]
^{150}Nd	$> 2.0 \times 10^{22}$	NEMO-3 [100]

be resolved in the coming decade. With that, the allowable parameter space for $0\nu\beta\beta$ decay given our typical assumptions comes into clearer view. Though $0\nu\beta\beta$ decay still evades us, the future of the search for $0\nu\beta\beta$ decay seems bright. In the United States, the Nuclear Science Advisory Committee Long Range Plan states [123]:

As the highest priority for new experiment construction, we recommend that the United States lead an international consortium that will undertake a neutrinoless double beta decay campaign, featuring the expeditious construction of tonne-scale experiments, using different isotopes and complementary techniques.

Chapter 3

The CUORE Experiment

The Cryogenic Underground Observatory for Rare Events (CUORE) [110] is an experiment primarily searching for $0\nu\beta\beta$ decay of ^{130}Te. It is located at the Laboratori Nazionali del Gran Sasso (LNGS) in Italy, the world's largest underground laboratory for neutrino and astroparticle physics. Home to many experiments including dark matter searches and rare event searches such as CUORE, LNGS is one of the few laboratories in the world in which many experiments can be built while shielded from cosmic rays. This shielding comes primarily from Gran Sasso d'Italia, a massif in the Apennine Mountain range which contains Corno Grande, the highest peak on the Italian peninsula. The massif provides 3800 meters water equivalent of shielding from cosmic rays, reducing the flux of cosmic ray muons from about 1 per cm^2 per minute at sea level to just 2×10^{-6} per cm^2 per minute at the LNGS underground labs [124]. While other underground laboratories provide greater shielding, LNGS has many practical advantages compared to its sister laboratories. The laboratory is accessible by vehicles, including cars and large trucks, allowing for the easy transport of researchers, equipment, and materials. This is particularly convenient for those who do not enjoy the experience of traveling down mine shaft elevators or dealing with their limited space.

As was discussed in chapter 2, there are several isotopes in which $2\nu\beta\beta$ has been measured. It turns out that not all of these isotopes are practical for use in an experiment given current technologies. For example, ^{48}Ca may seem like a fine choice of isotope. Of all the known double beta decaying isotopes, ^{48}Ca has the highest Q-value (4.27 MeV) and the smallest atomic mass. Its low mass number makes calculating the nuclear matrix elements more tractable than doing so for the other double beta decaying isotopes. Unfortunately, ^{48}Ca is an extremely scarce isotope of calcium, making up only 0.187% of natural calcium by molar fraction. This renders enrichment costs extremely high; a gram of ^{48}Ca costs hundreds of thousands of dollars, and global production is limited to roughly ten grams per year. This makes a large-scale $0\nu\beta\beta$ decay search using ^{48}Ca very impractical. ^{130}Te does not suffer from the same problems. ^{130}Te is the most abundant isotope of tellurium, with a natural

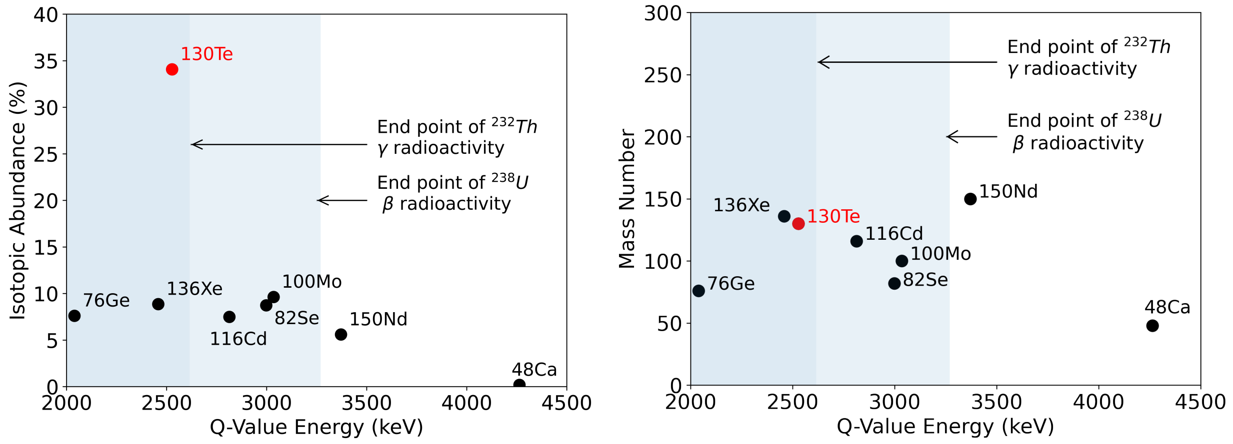

(a) (b)

Figure 3.1: A comparison of double beta decaying isotopes. (a) The natural isotopic abundance of different double beta decaying isotopes. (b) The Q-values of the double beta decay of different isotopes and the energies of some relevant backgrounds.

abundance of ∼34%[1][125]. This is by far the highest natural abundance of all known double beta decaying isotopes (See Fig. 3.1a). By selecting ^{130}Te as the $0\nu\beta\beta$ candidate isotope, CUORE can use natural tellurium, eliminating enrichment costs. The Q-value of double beta decay in ^{130}Te is 2.5275 MeV, which is higher than many sources of γ and β decay present in laboratories on Earth. A somewhat significant drawback of this isotope is that 2.5275 MeV is less than the 2.6145 MeV γ ray emitted by the de-excitation of ^{208}Pb after the β decay of ^{208}Tl, which is the last process in the 4n decay chain of ^{232}Th (See Fig. 3.1b). Luckily, the Q-value of ^{130}Te lies between the Compton edge of this γ ray and the peak energy value, so any backgrounds from ^{208}Tl that fall near the Q-value of ^{130}Te must come from multiple instances Compton scattering. Given the scattering cross-section in TeO_2, the interaction length of a 2.615 MeV γ ray turns out to be close to the length of a CUORE crystal, meaning that a multi-Compton event is less likely than a single scattering in a CUORE crystal. Multi-compton events that occur in multiple crystals can be tagged using coincidence analysis and then rejected when conducting a $0\nu\beta\beta$ search. Finally, a major benefit of using tellurium is that it can be manufactured into large, high-purity crystals. This keeps radioactive backgrounds to a minimum while also allowing one to instrument the crystals with detectors at the "human scale" as opposed to at the microscopic scale.

CUORE is the most recent of many experiments utilizing the bolometric technique, described in detail below, to search for $0\nu\beta\beta$ decay in ^{130}Te. Using the bolometric technique to search for $0\nu\beta\beta$ was first proposed by Ettore Fiorini in 1984. In 1994, Fiorini and oth-

---

[1] Interestingly, the second-most abundant isotope of tellurium, ^{128}Te, also double beta decays and has a natural abundance of ∼32%. However, the half-life of $2\nu\beta\beta$ in ^{128}Te is about 2000 times longer than that of ^{130}Te (2.2×10^{24} years versus 8.2×10^{20} years).

ers demonstrated the feasibility of this technique by operating a single TeO_2 crystal with a mass of 340 grams [126]. Further results on $0\nu\beta\beta$ of ^{130}Te were obtained with four [127] and eight [128] crystals with similar design in the late 1990s. In 1997, the Milano Double Beta Decay (MiDBD) experiment [129], a tower of 20 TeO_2 crystals, began collecting data. These experiments paved the way for CUORICINO [130], an array of 62 TeO_2 crystals arranged into 13 floors with a total mass of 40.8 kg and 11.8 kg of ^{130}Te. Cuoricino operated from 2003 to 2008 and set a limit on the half-life of $0\nu\beta\beta$ in ^{130}Te of 2.8×10^{24} years (90% C.L.) with 19.75 kg·yrs of ^{130}Te exposure.

The successor to Cuoricino was CUORE-0 [131], which operated from 2013 to 2015. CUORE-0 consisted of an array of 52 TeO_2 crystals, each $5 \times 5 \times 5\text{cm}^3$, arranged into a tower of 13 floors with four crystals per floor arranged in a two-by-two configuration. CUORE-0 was an important experiment designed to validate the background-reduction techniques intended for CUORE. These techniques included improving the detector production, developing better cleaning strategies to enhance the radio-purity of the detectors (particularly at the crystal surface), and procuring more radio-pure materials, including copper, for the detector support structure. These methods proved effective, as the background due to α contamination was six times less than that of Cuoricino. In addition, the improvements made to the detector assembly line improved the uniformity of the detectors, which showed more consistent resolutions than Cuoricino. After two years of data-taking, CUORE-0 set a limit on the half-life of $0\nu\beta\beta$ in ^{130}Te of 2.7×10^{24} years (90% C.L.) with only 9.8 kg · yrs of ^{130}Te exposure [132]. The fact that this limit is so similar to that of Cuoricino with less than half of the exposure can be explained by the improvement in detector energy resolution and the reduction of backgrounds in CUORE-0 compared to Cuoricino. When the data from Cuoricino and CUORE-0 were combined, the resulting limit on $0\nu\beta\beta$ in ^{130}Te was found to be 4.0×10^{24} years at 90% C.L. CUORE was constructed between 2015 and 2017 using effectively the same techniques that were realized for CUORE-0. The detectors are arranged into 19 towers, each consisting of 52 crystals with the same geometry as CUORE-0. Of these 988 detectors, 984 were successfully deployed, while four were declared dead before the start of data-taking.

Before the publication of this dissertation, the most stringent limits on the half-life of $0\nu\beta\beta$ in ^{130}Te come from the analysis of CUORE data with one (metric) tonne-year of TeO_2 exposure.[2] The most recent results from CUORE's search for $0\nu\beta\beta$ decay in ^{130}Te with two tonne-years of exposure will be discussed later in chapter 8. This chapter will focus on the experimental techniques used in the CUORE experiment and how they allow for the search of $0\nu\beta\beta$ decay.

---

[2] The word "tonne" will be used throughout this dissertation in an attempt to distinguish it from the American "ton."

3.1 Cryogenic Calorimeters

A calorimeter is any device used to measure energy. The earliest calorimeters were created in the eighteenth century; one of the earliest uses was to measure the latent heat of the melting of water. Calorimetry has developed into a rich field of research in the past quarter millennium. While most calorimeters used in high energy physics measure the energy loss of a particle interaction in the form of ionization or scintillation light, cryogenic calorimeters measure this energy in the form of ionization and heat [133]. A typical cryogenic calorimeter consists of an absorber, a thermometer, and a thermal link to a thermal bath with constant temperature T_0. When energy is deposited in the absorber, this causes a change in the thermometer – typically a change in its resistance, which leads to a measurable change in voltage when passing a current through it.

When a given amount of energy ΔE is deposited in a calorimeter, the temperature of the calorimeter increases by an amount

$$\Delta T = \frac{\Delta E}{C} \tag{3.1}$$

where C is the heat capacity of the calorimeter. When searching for $0\nu\beta\beta$ decay, we seek to measure energies of order 1 MeV with resolutions of at most 1% of the energy deposition. This is a challenging feat because we must build calorimeters in our macroscopic world. Operating at extremely low temperatures makes this possible thanks to Debye's Law, which states that the heat capacity of a dielectric crystal scales with temperature as

$$C(T) \propto \left(\frac{T}{\Theta_D}\right)^3 \tag{3.2}$$

where Θ_D is the Debye temperature of the crystal.[3] This means that at very low temperatures, a small amount of energy deposited in a dielectric crystal will result in a relatively large change in the crystal temperature. This suggests that operating calorimeters at cryogenic temperatures allows one to build a calorimeter sensitive to tiny energy deposits. For a $\mathrm{TeO_2}$ crystal, the Debye temperature is 232 K [134], and when operating at a base temperature of $T = 15$ mK, the corresponding heat capacity is roughly 1 MeV / 10^{-4} K.

No calorimeter is perfect; there will always be random fluctuations in the amount of energy measured by the system. It can safely be assumed that the amount of energy deposited in the absorber will be normally distributed. We can treat the calorimeter as a canonical ensemble to derive the energy resolution. The average energy $\langle E\rangle$ is given by:

$$-\frac{\partial \ln Z}{\partial \beta} = -\frac{1}{Z}\frac{\partial Z}{\partial \beta} = \frac{1}{Z}\sum_i E_i e^{-\beta E_i} = \langle E\rangle \tag{3.3}$$

---

[3] For conductors, the heat capacity is more accurately modeled as $C(T) \propto a\left(\frac{T}{\Theta_D}\right)^3 + bT$ where a and b are material-specific constants. The linear term comes from the electron contribution to the heat capacity, while the cubic term comes from the phonon contribution.

where Z is the partition function, the sum is over all microstates, and β is the thermodynamic beta $1/k_BT$, where k_B is the Boltzmann constant. Differentiating again, we have

$$\frac{\partial^2 \ln Z}{\partial \beta^2} = -\frac{\partial}{\partial \beta}\left[\frac{1}{Z}\sum_i E_i e^{-\beta E_i}\right] = \left[\frac{1}{Z}\sum_i E_i^2 e^{-\beta E_i}\right] - \left[\frac{1}{Z}\sum_i E_i e^{-\beta E_i}\right]^2 = \langle E^2 \rangle - \langle E \rangle^2 = \sigma_E^2 \tag{3.4}$$

Now, using the definition of heat capacity:

$$C = \frac{\partial \langle E \rangle}{\partial T} = -\frac{1}{k_BT^2}\frac{\partial \langle E \rangle}{\partial \beta} = \frac{1}{k_BT^2}\frac{\partial^2 \ln Z}{\partial \beta^2} = \frac{\sigma_E^2}{k_BT^2} \tag{3.5}$$

Rearranging to solve for the energy resolution, we have:

$$\sigma_E = \sqrt{k_BT^2C} \tag{3.6}$$

Note that this is not dependent on ΔE, so we can expect this contribution to the energy resolution to be constant regardless of the amount of energy deposited in the absorber. Of course, this is an intrinsic contribution to the energy resolution that does not consider outside sources of noise, which may further degrade it. In a TeO_2 crystal operating at $T = 15$ mK, this corresponds to a resolution of roughly 14 eV. In contrast, the energy resolution in a typical CUORE crystal is $1 - 10$ keV, and the observed energy resolution strongly depends on the energy deposited in the detector. Other contributions to the resolution therefore cannot be ignored. Studying the sources that further degrade the energy resolution is a focal point of this dissertation.

In CUORE, the cryogenic calorimeter technique has been realized through the careful design and fabrication of state-of-the-art detector components. Fig. 3.2 shows a schematic of one CUORE calorimeter. The absorber is the TeO_2 crystal, which contains the $0\nu\beta\beta$ decay candidate isotope, ^{130}Te. The fact that the $0\nu\beta\beta$ candidate isotope is itself a part of the detector is a crucial aspect that makes CUORE a competitive experiment. This “source = detector” technology gives rise to a high detection efficiency which would not be possible if the source and detector were separated into two different components.

The thermometers used in CUORE are neutron-transmutation doped (NTD) germanium thermistors [135, 136]. A thermistor is simply a resistor whose resistance depends on temperature. When a current is supplied to the thermistor, a change in temperature causes a change in the voltage drop across the thermistor, resulting in a measurable signal that can be read out. The NTD germanium is fabricated by irradiating samples of pure germanium with a controlled flux of thermal neutrons. Of the five stable isotopes of Germanium (A = 70, 72, 73, 74, 76), three transmute into dopants once they have captured a thermal neutron [137]:

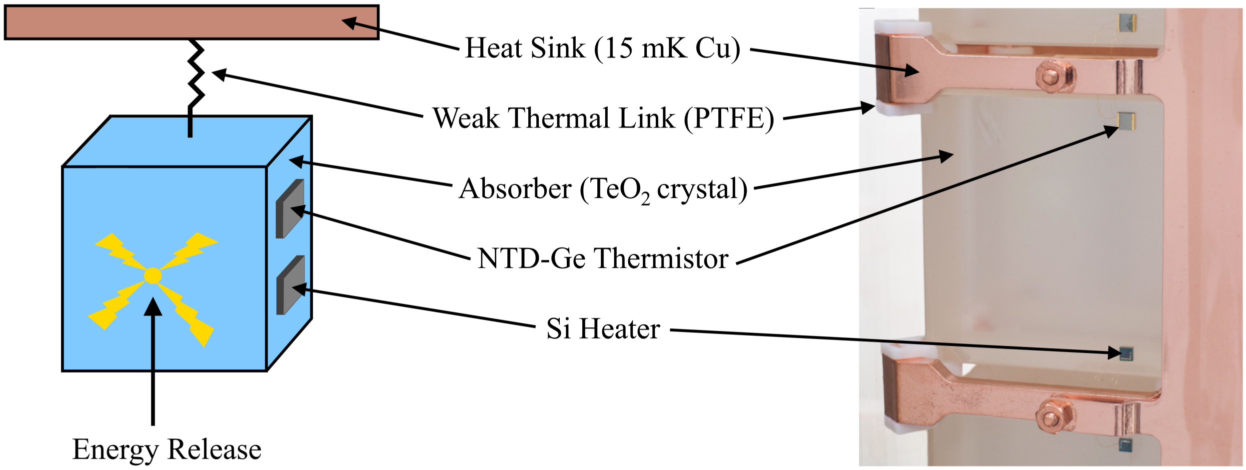

Figure 3.2: Schematic of a CUORE calorimeter. The absorber is the TeO$_2$ crystal which also contains the $0\nu\beta\beta$ decay candidate isotope, ^{130}Te. The thermometer is an NTD germanium thermistor glued to the absorber and read out via gold wire bonds. The crystal is connected to the copper frame of the detector by PFTE holders which serve as a weak thermal link. The copper frame is connected to the mixing chamber stage of the dilution refrigerator and serves as the heat sink.

$$^{70}_{32}\text{Ge} + n \rightarrow ^{71}_{32}\text{Ge} \xrightarrow[\text{T}_{1/2}=11.2\,\text{days}]{\text{EC}} {}^{71}_{31}\text{Ga} + \nu_e \tag{3.7}$$

$$^{74}_{32}\text{Ge} + n \rightarrow ^{75}_{32}\text{Ge} \xrightarrow[\text{T}_{1/2}=82.8\,\text{min}]{\beta^-} {}^{75}_{33}\text{As} + \bar{\nu}_e \tag{3.8}$$

$$^{76}_{32}\text{Ge} + n \rightarrow ^{77}_{32}\text{Ge} \xrightarrow[\text{T}_{1/2}=11.3\,\text{hrs}]{\beta^-} {}^{77}_{33}\text{As} + \bar{\nu}_e \xrightarrow[\text{T}_{1/2}=38.8\,\text{hrs}]{\beta^-} {}^{77}_{34}\text{Se} + \bar{\nu}_e \tag{3.9}$$

Compared to other techniques such as diffusion, this process results in a high uniformity of dopants throughout the volume of the germanium sample (so long as the thermal neutron source used is large compared to the sample). This is because in natural germanium, the stable isotopes are randomly distributed throughout a crystal and because the absorption cross sections for neutron capture are of order 10^{-24} cm^2 at the relevant neutron energies, effectively eliminating the self-shielding of the neutrons by the germanium [138].

The resistivity of NTD germanium has been found to obey the relation derived in [139]:

$$\rho = \rho_0 e^{\sqrt{T_0/T}} \tag{3.10}$$

where ρ_0 and T_0 are constants that depend on the doping. For cryogenic calorimeters, this has the important ramification that as $T \rightarrow 0$, the resistivity diverges. By applying a bias current to the thermistor at cryogenic temperatures, one can thus measure significant

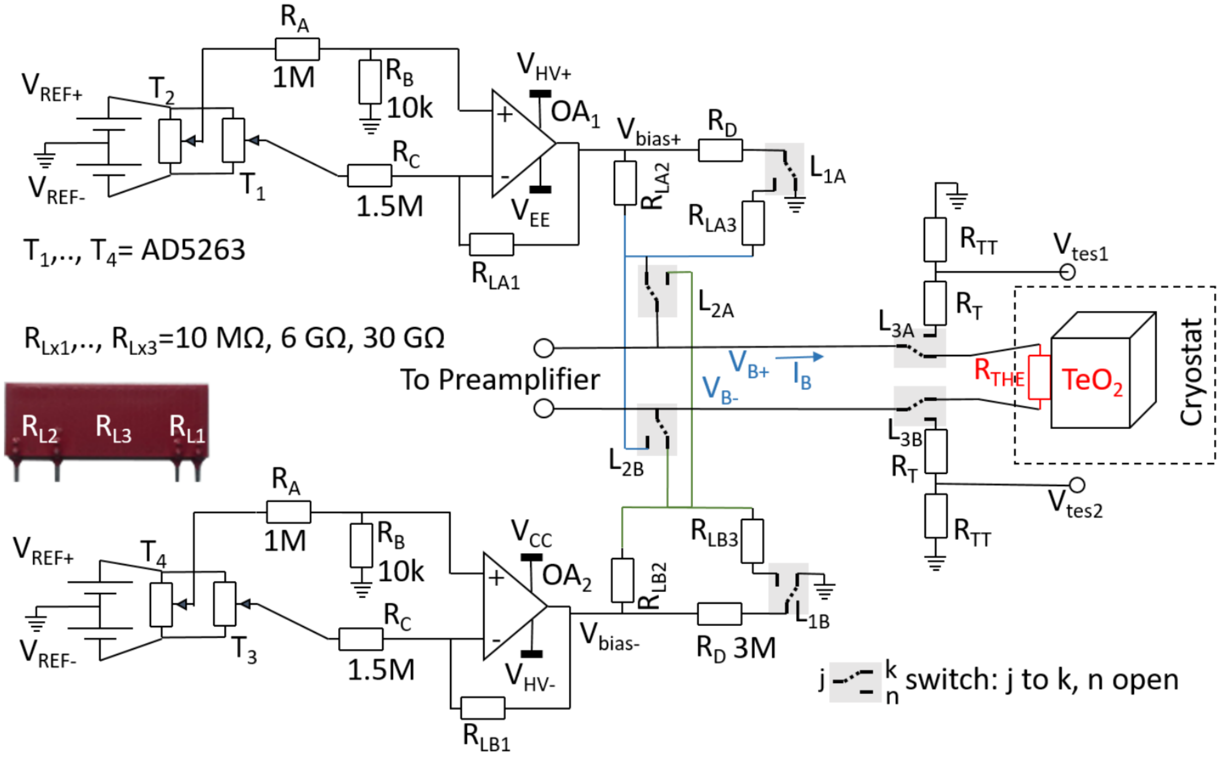

Figure 3.3: Schematic of the biasing circuit for the detector thermistor R_{THE}, glued on the TeO_2 crystal, not to scale. Reprinted from [140].

changes in temperature by reading the corresponding changes in the output voltage. A detailed schematic of the biasing circuit is shown in Fig. 3.3. Combined with the fact that the heat capacity of the absorber is extremely small at low temperatures, this high resistivity allows one to build a cryogenic calorimeter capable of measuring the energy deposition from single particle events.

In CUORE, each TeO_2 crystal is outfitted with one NTD-Ge thermistor and a silicon heater [141]. This heater allows one to inject a known amount of energy into the absorber in a controlled manner, allowing one to study the thermal response of the detector. For a particle event to be measured by the thermistor, the heat from the event must dissipate slowly enough that the thermistor and the absorber have time to be in thermal equilibrium with each other, but not with the thermal bath. This is ensured by attaching small pieces of Teflon to each corner of the crystal. The Teflon pieces are then directly connected to a copper support frame, which serves as the thermal bath at 15 mK. Along with the gold wires that connect the thermistor and heater to the readout cables, these Teflon pieces serve as weak thermal links connecting the calorimeter to the cryogenic structure.

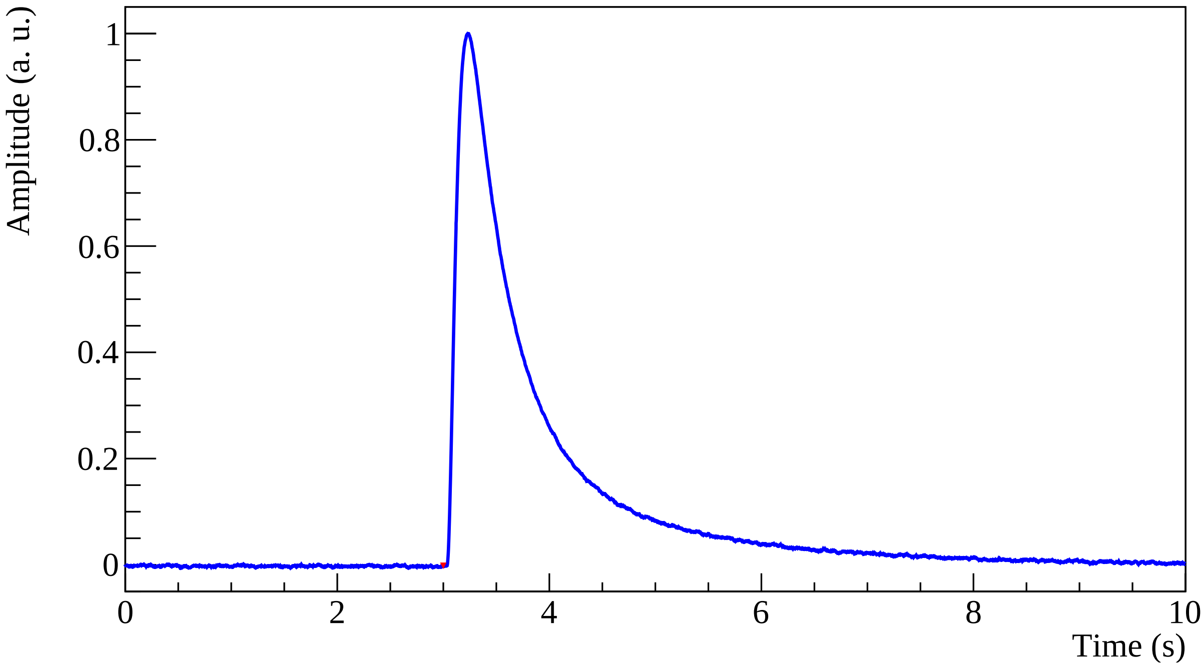

Figure 3.4: A typical CUORE pulse. The time axis is set such that the pulse is triggered at $t = 3$ seconds, and the pulse amplitude is normalized to 1. The amplitude of the pulse is proportional to the temperature change in the calorimeter, $\Delta T = \frac{E}{C}$ where E is the energy deposited in the detector by a particle interaction and C is the heat capacity of the detector. The decay time of the pulse is given by $\tau_D = C/G$ where G is the thermal conductance between the calorimeter and the thermal bath. A typical CUORE pulse has a decay time of order 1 second.

The characteristic output signal from a particle event is a thermal pulse with two time constants.[4] A typical CUORE pulse is shown in Fig. 3.4. The first time constant, τ_R, is proportional to the *rise time* of the pulse, and it is the amount of time it takes for the NTD-Ge thermistor to come into thermal equilibrium with the absorber. The second time constant, τ_D, is proportional to the *decay time* of the pulse. Analogous to the time constant of an RC circuit, this time constant is given by $\tau_D = C/G$, where C is the heat capacity of the detector and G is the thermal conductance between the calorimeter and the thermal bath. In CUORE, a typical pulse has a decay time of roughly 1 second.

---

[4] In practice, it is known that real pulses in CUORE often exhibit three time constants: a rise time and two decay times. One decay time is due to the heat capacity of the crystal and the thermal conductance of the gold wires, while the other is due to the heat capacity and thermal conductance of the Teflon. Here, we ignore the second decay time to simplify the explanation of the system.

The signal is sent to room-temperature electronics outside the cryogenic infrastructure to read out the voltage signals from the NTD-Ge thermistors. The signal first passes through a front-end electronics board which consists of a differential pre-amplifier and a second-stage programmable gain amplifier. It is then passed to an analog anti-aliasing Bessel-Thomson filter. This filter features a 120 dB/decade roll-off and a selectable -3 dB bandwidth ranging from 15 Hz to 120 Hz in four steps [142], with unity gain below the cutoff frequency. The Bessel-Thomson filter was chosen because of its maximally flat phase delay, which prevents distortion of the pulse shape after filtration. The filtered analog signal signal is then digitized at 1 kHz with a 24-bit $\Sigma\Delta$ analog-to-digital converter (ADC) and written to disk. The 24-bit ADC is chosen so that the digitization noise is negligible compared to other noise sources, and the 1 kHz sampling frequency ensures that the relevant signal band ($\lesssim 10Hz$) is well-covered. The timestreams from all 988 calorimeter channels are stored continuously, i.e. there is no data rejection upfront. This allows all digital signal processing, triggering, and pulse shape discrimination to be done offline, though triggering is also done online for the purposes of detector monitoring. More details on the front-end electronic system can be found in [140].

3.2 The Cryogenic Infrastructure

3.2.1 ^{3}He/^{4}He Dilution Refrigerators

The cryogenic calorimeters used in CUORE must be operated at temperatures well below 1 K to measure nuclear and particle interactions effectively. These temperatures are achieved with ^{3}He/^{4}He dilution refrigerators (DRs). ^{3}He/^{4}He DRs represent a pinnacle in cryogenic technology and low-temperature physics; they are uniquely capable of cooling large payloads to temperatures in the millikelvin range while keeping vibration and potential contamination to a minimum [143]. This makes them ideal for experiments that demand high thermal stability and low thermal noise. Today they are essential for experiments in quantum computing, condensed matter physics, experimental astrophysics, and nuclear physics.

Helium exists primarily as two stable isotopes: the relatively rare ^{3}He isotope with a single neutron and the more abundant ^{4}He isotope with two neutrons. Both isotopes become liquid at low temperatures but exhibit distinct quantum mechanical behaviors. The ^{4}He nucleus is composed of four spin-$\frac{1}{2}$ nucleons, thus it is a boson. When cooled below 2.17 K (at 1 atmosphere), the so-called "lambda point," ^{4}He undergoes Bose-Einstein condensation and becomes a superfluid. In contrast, ^{3}He is a fermion and only becomes a condensate once the ^{3}He atoms preferentially pair together in a process analogous to the pairing of electrons in Bardeen–Cooper–Schrieffer theory. The lambda point of ^{3}He is roughly 2.5 mK. This difference in lambda points is at the heart of the DR technology.

Dilution refrigeration leverages the unique thermodynamic properties of a mixture of ^{3}He and ^{4}He at low temperatures. The key phenomenon is a phase separation that occurs in a ^{3}He/^{4}He mixture below 0.87 K, depending on the concentration of ^{3}He in the mixture (See Fig. 3.5). At these temperatures, the mixture separates into two distinct phases: a concentrated phase rich in ^{3}He and a dilute phase containing a small proportion of ^{3}He dissolved in ^{4}He.

The refrigeration cycle of the DR is based on the continuous extraction of ^{3}He from the ^{3}He-rich phase to the dilute phase. This transfer process is endothermic; it absorbs heat from the surroundings, thereby cooling the system. The DR exploits this cooling effect to reach temperatures as low as a few millikelvin. For this refrigeration to be possible, one must first reach a temperature below 0.87 K so that the helium phase separation can occur. DRs can widely be classified into two main categories, each of which uses very different technologies to achieve this: so-called "wet" cryostats and "dry" cryostats.

In a wet DR, the cooling process begins with a "main bath" of liquid ^{4}He, which serves as a primary coolant. Liquid nitrogen is often used first to cool the bath from room temperature to 77 K, then the nitrogen is removed, and liquid ^{4}He is added. This dramatically reduces the amount of helium lost to evaporation, and liquid nitrogen is much cheaper and more readily available than liquid ^{4}He. The main bath is maintained at a temperature of around 4.2 K under normal pressure, but by reducing the pressure over the helium bath using a vacuum pump, the temperature can be lowered to around 1.5 K. This step is known as "pumping on the bath." Once this temperature is achieved, a Joule-Thomson (JT) expansion valve cools a gas mixture of ^{3}He and ^{4}He to lower temperatures. The cooled mixture then enters a series of heat exchangers and stages within the refrigerator, progressively cooling the helium mixture as it moves through each stage before finally arriving at the mixing chamber stage.

On the other hand, dry DRs, also known as cryogen-free DRs, eliminate the need for a liquid helium bath by utilizing a closed-cycle refrigeration system. This system eliminates the need for continuous liquid helium consumption, making dry DRs more sustainable and cost-effective for long-term operations. The primary cooling mechanism in dry DRs involves mechanical cryocoolers, such as pulse-tube coolers or Gifford-McMahon coolers, which compress and expand helium gas in a closed cycle, cooling the system to temperatures below 4 K. Once the system reaches these temperatures, a series of mechanical and thermal stages are employed to achieve further cooling. The pre-cooled ^{3}He/^{4}He mixture is circulated through a series of heat exchangers, where it progressively cools down through counter-flow heat exchange with the returning mixture from the mixing chamber. In these heat exchangers, the temperature of the ^{3}He decreases with each pass. The mixture is then expanded through a capillary or orifice, allowing the temperature to drop significantly as the gas undergoes JT expansion, as is the case in wet cryostats.

In both wet and dry DRs, once the mixture is sufficiently cooled via JT expansion and has reached the mixing chamber, the enthalpy of mixing drives the phase separation. A dilution unit is then needed to reach millikelvin temperatures. A typical ^{3}He/^{4}He DR consists of several key components which comprise the dilution unit:

- Mixing chamber: The coldest part of the DR, where the two phases of the helium mixture coexist and phase separation occurs. This serves as the site of the primary cooling effect.

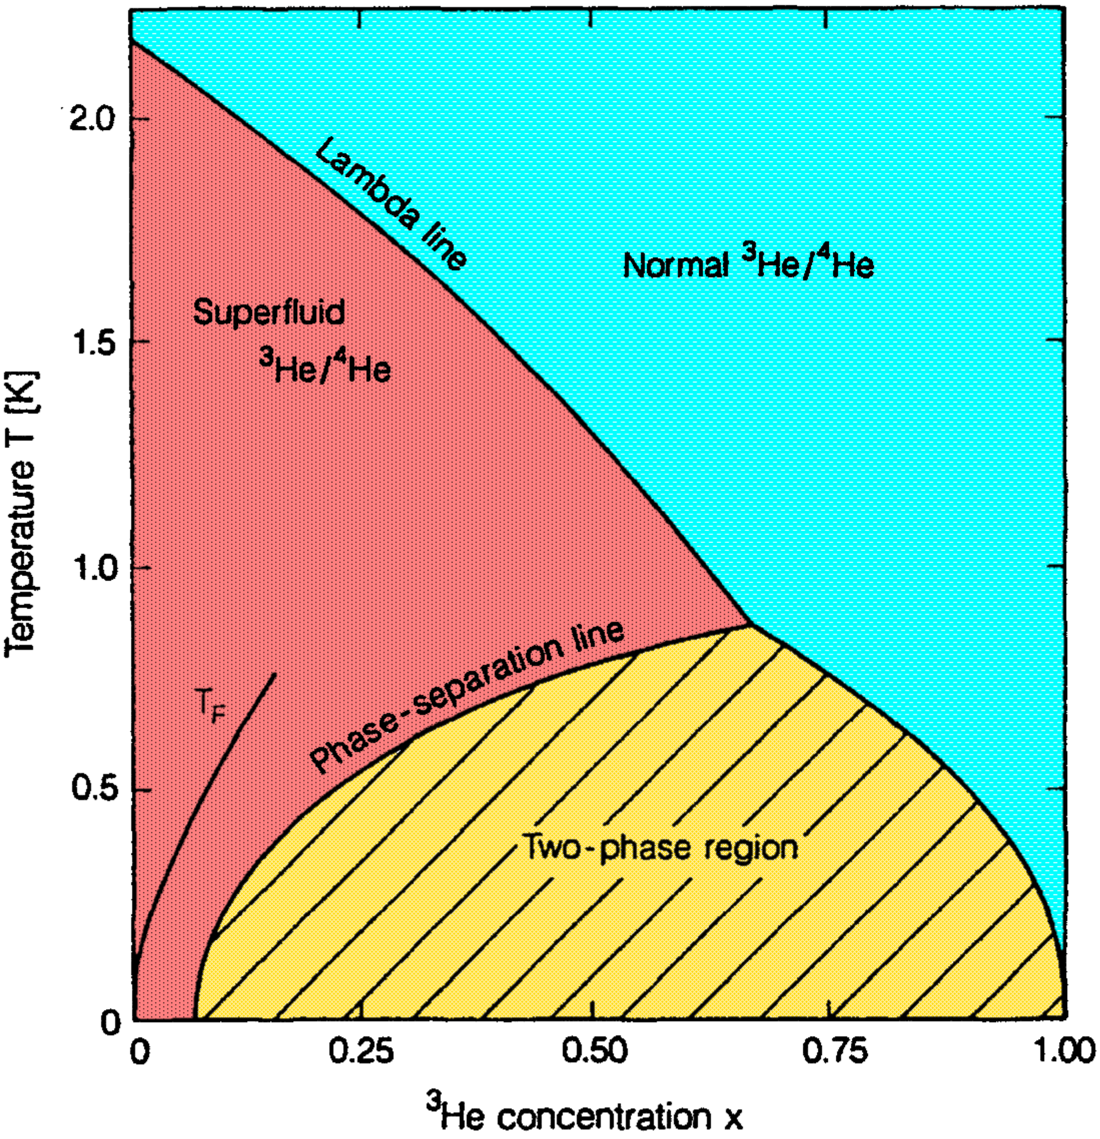

Figure 3.5: Phase diagram of liquid ^{3}He/^{4}He mixtures. At temperatures below 0.87 Kelvin, phase separation occurs, creating a ^{3}He-rich phase and a ^{3}He-dilute phase. The Lambda line is the line below which ^{4}He becomes a superfluid Bose-Einstein condensate, and T_F is the Fermi temperature. Reprinted from [143], then colorized.

- Still: A component that separates the evaporated ^{3}He from ^{4}He, with ^{3}He preferentially evaporating due to its higher vapor pressure, facilitating recycling. This is positioned above the mixing chamber and is slightly warmer.
- Heat exchangers: Crucial components that transfer heat from warmer parts of the system to colder regions. These enhance the DR's efficiency by pre-cooling the incoming ^{3}He using the outgoing flow from the mixing chamber.

A schematic of a dilution unit is shown in Fig. 3.6. Outside the dilution unit, ^{3}He is continuously circulated through the system. It is pumped through the refrigerator and then re-condensed in a condenser, then it re-enters the mixing chamber allowing the cycle to continue.

3.2.2 The CUORE cryostat

Much of the information reported in this section is taken from [144], and a more detailed description of the CUORE cryostat can be found there. The CUORE cryostat is a cryogen-free dilution refrigerator custom-made by Leiden Cryogenics. It is one of the most powerful cryostats in the world, providing more than 96 W of cooling power at 45 K and roughly 4 μW of cooling power at 10 mK (though this will be easily surpassed by the Colossus DR [145] which is currently being built at Fermilab). The CUORE cryostat has a fiducial volume of more than one cubic meter. When operating at 10-15 mK, this is the coldest cubic meter in the known universe [146].

The cryostat comprises six coaxial vessels at 300 K, 40 K, 4 K, 600 mK, 50 mK, and 10-15 mK (see Fig. 3.7). The 300 K and 4 K stages each form a vacuum chamber. The outer vacuum chamber (OVC) has a volume of 5.9 m^3 and a minimum pressure of approximately 10^{-6} mbar. The inner vacuum chamber (IVC) has a volume of 3.4m^3 and a minimum pressure of roughly 10^{-8} mbar. The five inner vessels are made using oxygen-free copper [147].

Two cold lead shields are placed inside the IVC to suppress background radiation from naturally occurring contaminants in the cryostat. The innermost shield is composed of Roman lead. This lead was found in a shipwreck off the coast of Sardinia from the days of the Roman Empire. It is therefore depleted of ^{210}Pb and other radioisotopes with half-lives shorter than several hundred years.[5] This shield is 6 cm thick and has a mass of roughly five metric tons. It is suspended from the 600 mK stage plate and is thermalized at 4 K. Another shield of copper and modern, natural lead is also placed inside the IVC, between the 10 mK plate and the Tower Support Plate. This shield is 30 cm thick and has a mass of roughly two metric tons. It is thermalized at 50 mK.

---

[5] Interestingly, the Roman lead used in CUORE has even less ^{238}U than would be expected in the lead from Ancient Roman times. Around the dawn of the Common Era, the Roman Empire was beginning to run out of silver. The Ancient Romans developed a method to extract silver from lead ore, and in the process, they extracted ^{238}U as well.

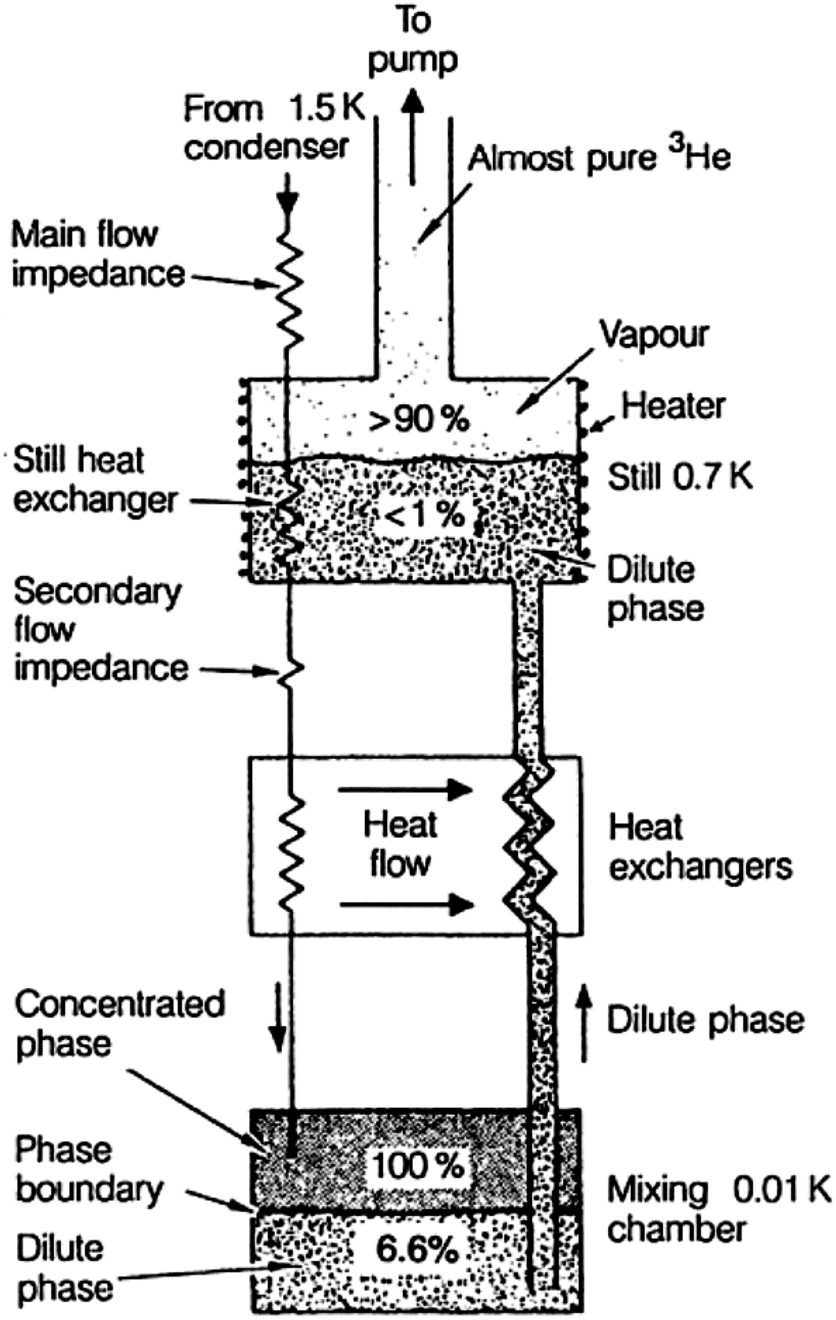

Figure 3.6: Schematic of the dilution unit of a ^{3}He/^{4}He dilution refrigerator. The percentages refer to the concentration of ^{3}He at each step. Reprinted from [143] and digitally edited for increased clarity.

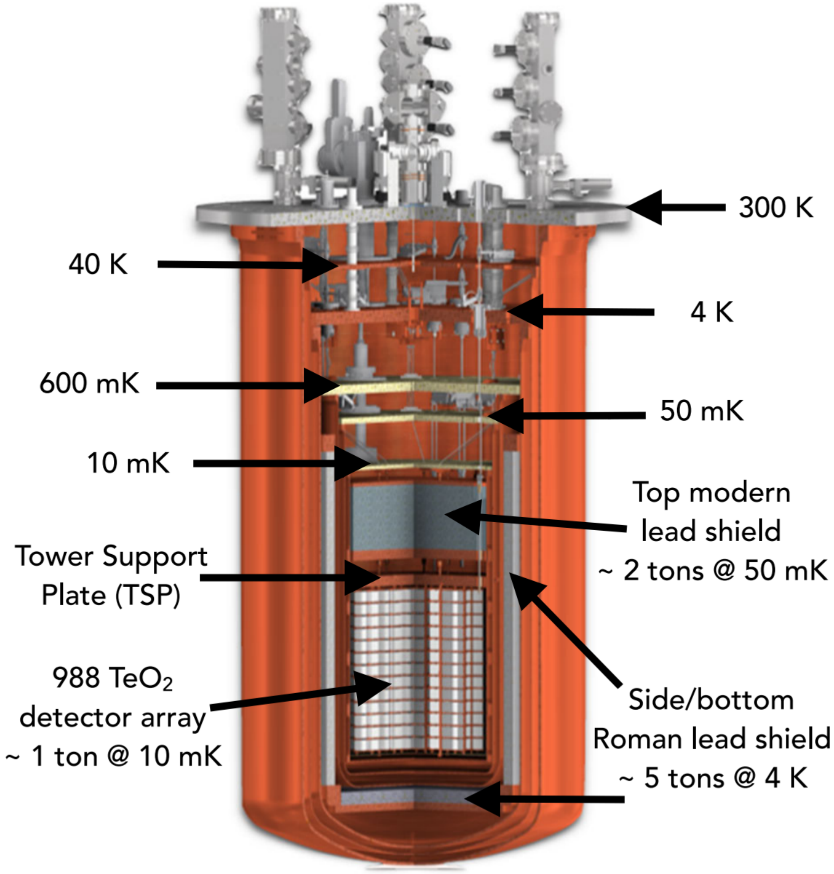

Figure 3.7: The CUORE cryostat structure, including all vessels and lead shieldings. Reprinted from [144].

Altogether, approximately 4.5 metric tons of material is cooled to 50 mK or below, and roughly 1.5 metric tons of material – including the mixing chamber plate, the Tower Support Plate, and the detector with 988 TeO_2 crystals – is cooled to the base temperature of 10-15 mK. To cool down this massive thermal load, the CUORE cryostat relies on three refrigeration systems, each capable of reaching a different temperature. The first stage is the "fast cooling system" (FCS) [148], which consists of heat exchangers that use 3 Gifford-McMahon refrigerators. Several sensors allow one to monitor and control the system during the cooldown procedure. The FCS cools the cryostat to ∼50 K and is switched off as the second stage begins. The second stage consists of five two-stage Cryomech PT415-RM pulse tubes coolers [149], hereafter referred to as simply "pulse tubes," each with a nominal cooling power of 1.2 W at 4.2 K and 32 W at 45 K. Many efforts are taken to dampen pulse tube vibrations and decouple them from the cryostat. These efforts will be further detailed in chapter 5. Finally, to achieve the base temperature of 10-15 mK, CUORE employs a DRS-CF3000 continuous-cycle $^3He/^4He$ DR, customized by Leiden Cryogenics. This third stage

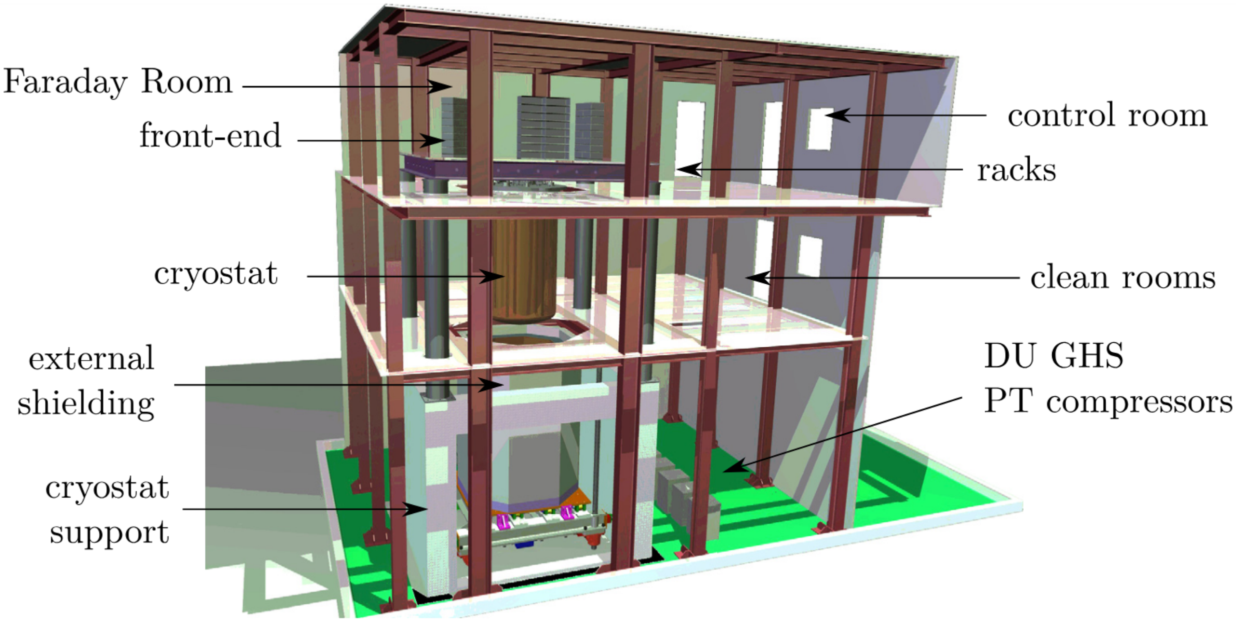

Figure 3.8: Diagram of the building housing the CUORE experiment in hall A of LNGS. Reprinted from [150].

operates simultaneously with the pulse tubes. With the full thermal load of CUORE, the dilution unit can reach a minimum temperature of roughly 7 mK.

3.3 The CUORE support and suspension systems

CUORE is located in a three-story building in Hall A of LNGS (see Fig. 3.8). The ground floor hosts the pulse tube compressors, the main cryostat support legs, and some external shieldings. The second floor (counting with American conventions) has a clean room to access the detector when the cryostat is open. The third floor contains the plumbing for the cryogenic infrastructure, the front-end electronics, and several parts of the suspension system. To electronically isolate the front-end electronics from the outside world, the parts of the experiment on the third floor of the building are housed inside a Faraday cage made of a high-permeability metallic glass alloy sandwiched by two aluminum sheets. This provides an attenuation of a factor of 15 at 50 Hz and a factor of 1000 from ∼1 kHz to ∼100 MHz [151].

The cryostat support structure was designed to mechanically isolate the detector from the building and the larger environment. LNGS is located in the heart of the Apennine mountains and is a quite active seismic area, so seismic safety requirements must also be met. The cryostat is mechanically connected to the ground via mechanical fuses. In the

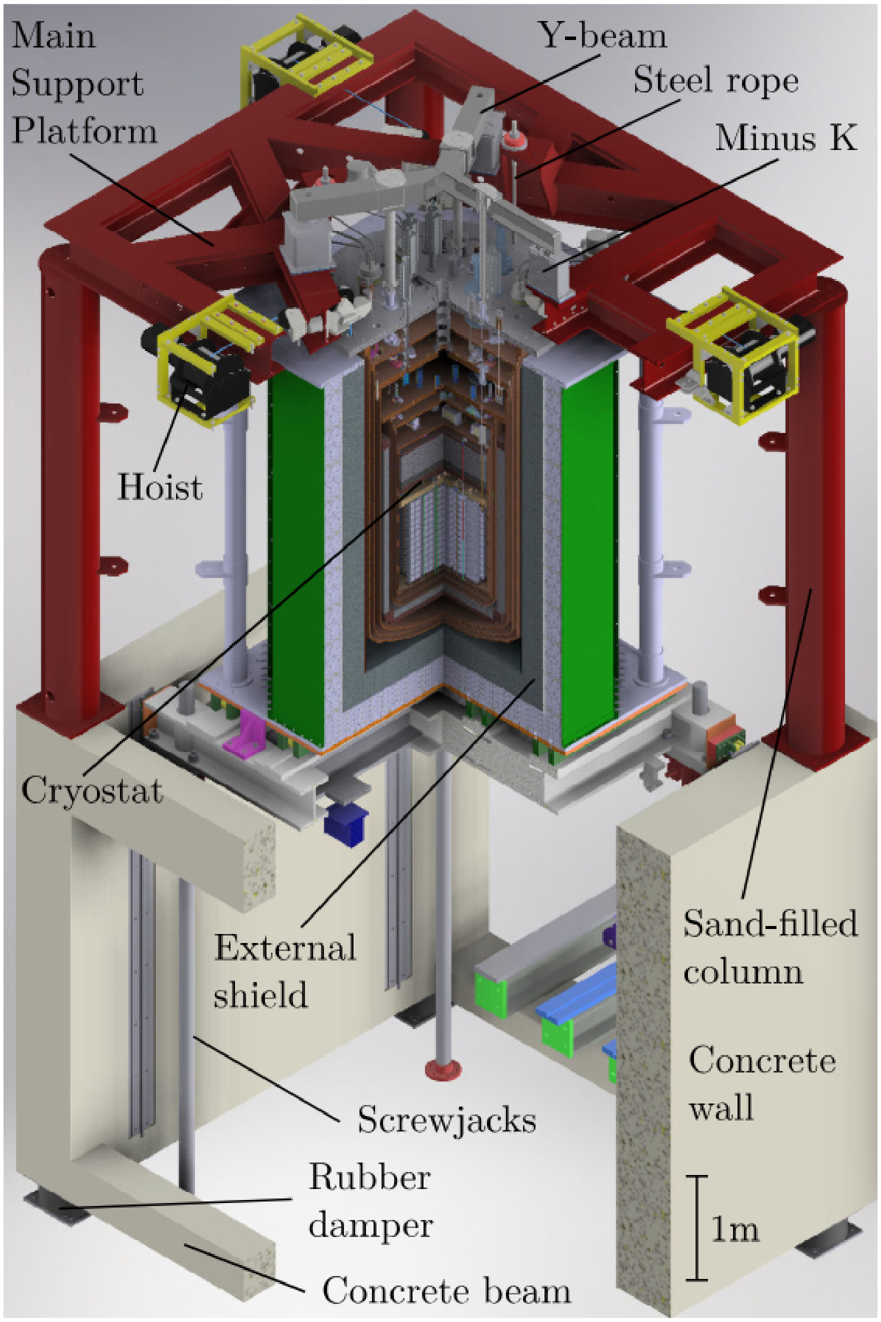

Figure 3.9: Diagram of the CUORE mechanical support structure. All labeled components are at room temperature. No cryogenic support structures are labeled here. Reprinted from [150].

event of a strong earthquake, the fuse breaks, mechanically isolating the cryostat and the mechanical support structure from the ground.

A diagram of the CUORE support structure is shown in Fig. 3.9. To minimize the impact of other external vibrations on the detector, the cryostat is suspended from the Main Support Platform, which is made of steel I-beams. The Main Support Platform supports the weight of the cryostat via three stainless steel ropes attached to the 300 K plate. A series of steel bars, Kevlar ropes, and copper bars are used to suspend each cryogenic stage plate from its outer neighbor. Notably, the Tower Support Plate is designed to be mechanically isolated from the other cryogenic stages; it is suspended from an external, room-temperature structure called the Y-beam. The Y-beam comprises three identical metal arms with square cross-sections joined symmetrically at 120-degree angles forming a Y shape. It is mechanically connected to the Main Support Platform via three Minus-K Technology mechanical insulators. The insulators act as low-pass filters which dampen vertical vibrations at frequencies above $\sim$0.5 Hz. Three segmented detector suspension bars connect the detector to the Y-beam. Each segment is thermalized to a different cryostat stage, down to the 600 mK stage. Finally, the Tower Support Plate is suspended from the Still stage.

Chapter 4

The CUORE Data Processing Chain

The raw CUORE data are stored as continuous digital waveforms for each channel. These channels include the 988 cryogenic calorimeters, several NTD-Ge thermistors used to monitor the temperature, and auxiliary devices discussed in greater detail in section 5.2. A significant amount of signal processing and analysis is needed to convert these waveforms to an energy spectrum that can be used to search for $0\nu\beta\beta$. Here, we describe this process in detail.

4.1 High-level Data Organization

The fundamental unit of CUORE data collection is a *run*. A run is a period of continuous data-taking that lasts roughly one day. The relevant types of runs in the context of this dissertation are calibration runs and physics runs. Calibration runs are those during which calibration sources are deployed, and the events from these sources dominate the resulting energy spectrum. Physics runs are those during which the experiment measures only background radiation and, ideally, $0\nu\beta\beta$ events from ^{130}Te. The data from these runs are used to construct the energy spectrum which is later analyzed to search for $0\nu\beta\beta$.

To analyze the data, several runs are grouped into datasets. A typical dataset comprises thirty to fifty physics runs sandwiched by three to six calibration runs on each end of the physics data-taking. Often, the “final calibration” runs for one dataset are also used as the “initial calibration” runs for the next dataset. The purpose of having both initial and final calibration runs is to roughly account for any changes to the cryogenic environment during the dataset. Precautions are taken to try to ensure that the cryogenic system is stable across an entire dataset, such as the use of a Proportional–Integral–Derivative (PID) controller that maintains a consistent temperature of the mixing chamber stage, but changes in the operating conditions of the cryostat still arise during a dataset. Datasets form the basic unit for most physics analyses, including the search for $0\nu\beta\beta$.

As physics data are collected, the collected TeO_2 exposure of the experiment grows at a rate of 743 kg years per year, and the ^{130}Te exposure grows at a rate of 253 kg years per year. Of course, the collected exposure will always be less than the analyzed exposure due to

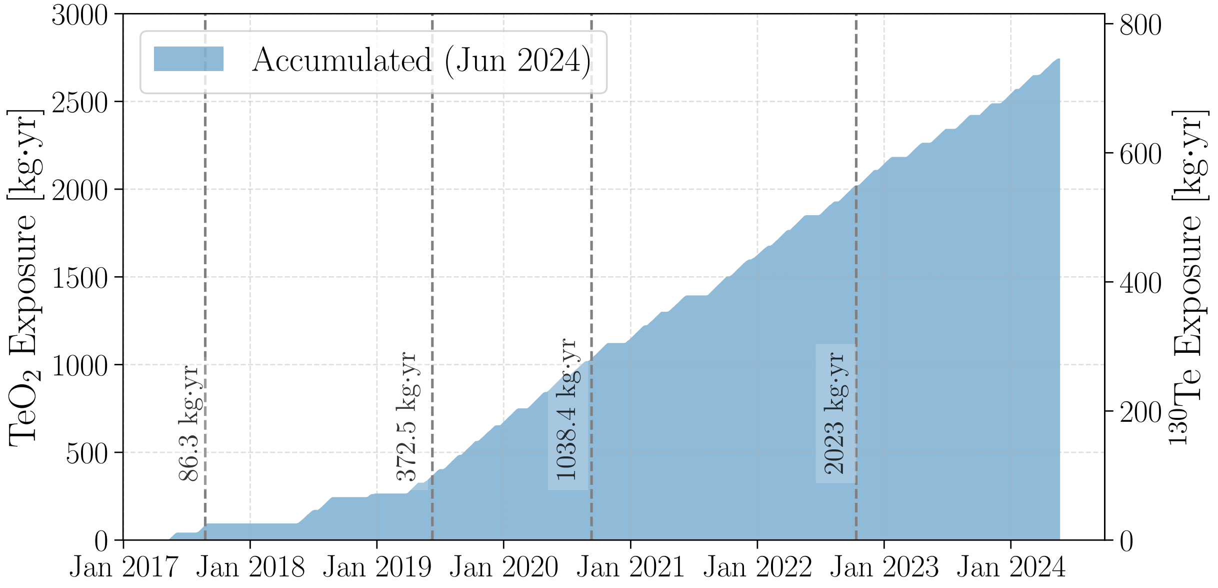

Figure 4.1: Collected exposure over time. Data-taking began in 2017, though there was significant down-time between 2017 and 2019. From January 2019 to May 2024, the detector demonstrated a 90% duty cycle with no cryogenic interventions necessary. During this time, the average exposure collected was roughly 500 kg years per year.

various efficiencies which are less than one. CUORE began collecting physics data in 2017. In that year, two datasets were collected, and then data-taking stopped for almost one year due to problems with the cryogenic infrastructure. Data-taking resumed in mid-2018, and two more datasets were acquired before another period of cryogenic interventions, and a fifth dataset was acquired toward the end of 2018. After a third period of cryogenic interventions lasting until mid-2019, data-taking resumed. Since the start of the sixth dataset, no cryogenic interventions were required until May 2024, during which time the average rate of TeO_2 exposure collection was roughly 500 kg years per year. During this time, the detector had a 90% duty cycle including the time spent calibrating the detector, which does not count toward the exposure. A plot showing the collected exposure over time is shown in Fig. 4.1.

4.2 Data Storage and Online Processing

CUORE saves all of the continuous data it takes during a run to disk for later analysis. With roughly 1000 channels sampled at 1 kHz with an 18-bit ADC, this corresponds to a data collection rate of roughly 200 GB per day. Given the relatively low sampling rate, no hardware trigger is needed, and a more flexible software trigger can be used. The `APOLLO` data acquisition (DAQ) software [152] performs both the DAQ and the online triggering algorithms on the data. While online triggering is not strictly necessary for analysis since the continuous data are saved, doing so offers the benefit of real-time data quality monitoring. The online triggering consists of two different triggering algorithms. The first, called the noise trigger, places a trigger at a pre-selected time. A ten-second long window is placed around this trigger, creating a "noise event." After some selection cuts to ensure they are not contaminated with a pulse, the noise events can be used to build a given calorimeter's average noise power spectrum during data-taking. The other type of trigger is called the derivative trigger (DT), and it is designed to fire when a pulse event occurs in a channel. The DT thresholds for each calorimeter must be set by hand because the noise levels and the energy response of each calorimeter channel can vary widely across the entire detector. When the DT fires, a ten-second window is placed around the trigger with a pre-trigger time of three seconds. This ensures that the baseline of the pulse, i.e. the first 2.25 seconds of the event, can be accurately measured in later processing steps. `APOLLO` then records the timestamp of the triggers in new files separate from the continuous data. In addition, every 570 seconds, identical square pulses are injected into the calorimeters via the silicon heaters described in section 3.1. These heater pulses (also called "pulser events" or simply "pulsers") are flagged explicitly using digital voltage signals that are synchronized in time with the pulser electronics. These digital signals are then used to flag all of the pulsers as their own type of event. The pulsers can then be used to perform a process called stabilization, which is described in more detail below, and to study the efficiency of the triggering algorithm.

A PostgreSQL database stores information about each channel, the configuration of the electronics, and the online triggering. It also stores higher-level bookkeeping information, such as each run's start and stop time, the run type, the location of the raw data, and the associated dataset number. As the DAQ process runs on the various underground computers, backups are automatically created and saved to an above-ground computer at LNGS, from which more backups are written to other machines in Italy and the United States, ensuring that the data are not lost.

4.3 Offline Processing

Once an entire dataset has been acquired and the data have been stored, the data can be processed with a more robust sequence of steps than the online triggering offers using the `DIANA` data analysis framework [153]. These steps include the construction of an optimum filter and retriggering of the data, stabilization, calibration, coincidence analysis, pulse shape discrimination, and energy reconstruction. This section outlines each of these techniques as they were employed for the latest CUORE result with two tonne-years of TeO_2 exposure.

4.3.1 Optimum Filter and Triggering

Optimum Filter

Over the course of an entire dataset, a typical channel will acquire several thousand noise events and anywhere from hundreds to thousands of valid physics events. Using this ensemble, we create the average pulse (AP) and the average noise power spectrum (ANPS). The AP $s[t]$ is simply a time-domain average of each pulse, i.e.:

$$s[t] = \frac{\sum_{i=1}^{N_{\text{events}}} s_i[t]}{\max_t \left[\sum_{i=1}^{N_{\text{events}}} s_i[t]\right]} \tag{4.1}$$

where $s_i[t]$ is the ith pulse from the dataset. The denominator ensures that the amplitude of the AP is 1. In practice, we use the events from calibration runs to construct the AP because the calibration runs have much higher statistics, and events triggered during calibration are likely “real” physics events. Thus, the AP should accurately represent the detector’s true response to a particle interaction. An example AP is shown in Fig. 4.2.

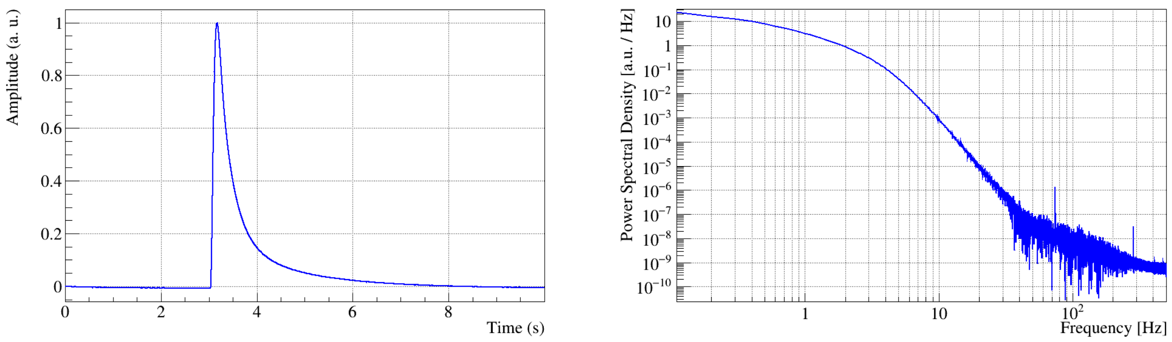

(a) (b)

Figure 4.2: An example of an average pulse (AP) from a typical CUORE channel. (a) The time-domain version of the AP. This is built by averaging events triggered by the derivative trigger with sufficiently large amplitude in the time domain. It is normalized such that the amplitude of the pulse is 1. (b) The power spectral density of the same average pulse.

The ANPS is more abstract and more difficult to compute than the AP, but it is a useful tool for analyzing the properties of the detector noise. The ANPS $J(f)$ of a signal $v(t)$ is defined as the Fourier transform of its autocorrelation function [154, 155]:

$$J(f) = \lim_{T\to\infty} \int_{-T/2}^{T/2} dt\, R(t) e^{-i2\pi f t} \tag{4.2}$$

where $R(t)$ is the autocorrelation function:

$$R(t) = \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2} dt'\, v(t')v(t+t') \tag{4.3}$$

Assuming that $v(t)$ has units of Volts, the ANPS has units of Volts^2 / Hz. This may seem unnatural as it is not a unit with any physical meaning. However, it is convenient because the variance of the noise in time is given by the integral of the ANPS over all frequencies:

$$\left\langle (v(t))^2 \right\rangle = \int_{-\infty}^{\infty} df\, [J(f)] \tag{4.4}$$

When dealing with signals of finite time T digitized at a sampling frequency f_s, the ANPS becomes:

$$J[f_n] = \sum_{i=0}^{N-1} \Delta t\, R[t_i] e^{-i2\pi f_n \Delta t} \tag{4.5}$$

where $N = f_s T$, $\Delta t = 1/f_s$, and $R[t_i]$ is the discrete autocorrelation function:

$$R[t_i] = \frac{1}{N} \sum_{j=0}^{N-1} v[t_i] v[t_{i+j}] \tag{4.6}$$

This can be re-written in another form [156]:

$$J[f_n] = T \left\langle |\tilde{v}_n|^2 \right\rangle \tag{4.7}$$

where $\tilde{v}_n$ is the nth element of the discrete Fourier transform of $v(t)$. This form is more in line with how the ANPS is typically calculated, which is known as Bartlett's method [157]. One collects a set of noise traces, calculates the mean norm-square of each Fourier amplitude, and then multiplies by T to ensure that the result has units of V^2/Hz. A typical ANPS for a typical CUORE channel is shown in Fig. 4.3. Unlike the AP, the ANPS is built using physics runs, since there are not enough noise events in calibration runs due to the higher event rate. When constructing the ANPS, several analysis cuts are placed on the noise events to ensure that the ANPS accurately represents the noise of the detector. For example, noise events with very high RMS and events containing derivative triggers are cut from the data.

The ANPS and AP are the necessary inputs to construct the matched filter, also called the optimum filter (OF) [158–161]. Applying the OF to the detector signal produces a new signal for which the signal-to-noise ratio is maximized [162]. Provided that the OF has unity gain, the maximum values of the filtered pulses are the optimal estimators of the pulse amplitudes. The least-squares error of a waveform $v[f]$ given the signal template $s[f]$ and the noise power spectrum J is:

$$\chi^2_{OF} = \sum_{n=0}^{N-1} \Delta f\, \frac{\left|\tilde{v}[f_n] - A\tilde{s}[f_n]\right|^2}{J[f_n]} \tag{4.8}$$

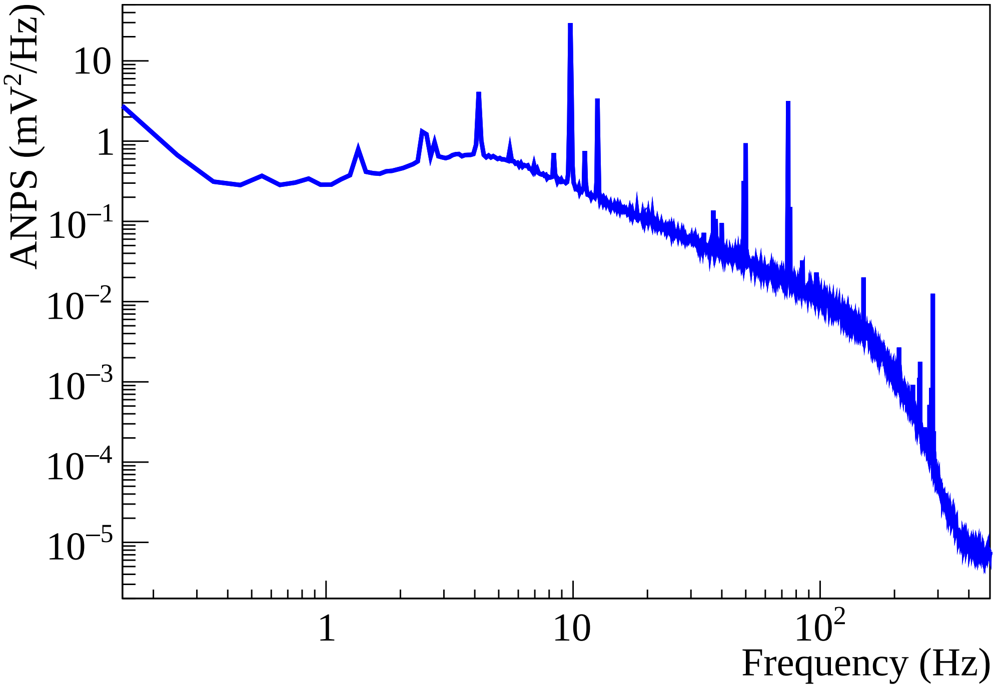

Figure 4.3: An example ANPS of a typical CUORE channel. The continuum from ∼1 Hz to ∼100 Hz is due to thermal fluctuation noise in the detector and Johnson noise in the read-out circuit, the roll-off above 100 Hz is due to the Bessel filtering of the data, the flattening out of the ANPS at close to 500 Hz is due to the digitization noise floor, and the peaks in the spectrum are due to a combination of excess vibrational and electrical noise.

Minimizing this quantity with respect to A gives the following estimator for the amplitude:

$$\left.\frac{\partial \chi^2_{OF}}{\partial A}\right|_{A=\hat{A}} = 0 \tag{4.9}$$

$$\Delta f \sum_{n=0}^{N-1} \frac{\left(\tilde{v}^*[f_n] - \hat{A}\tilde{s}^*[f_n]\right)(-\tilde{s}[f_n]) + \left(\tilde{v}[f_n] - \hat{A}\tilde{s}[f_n]\right)(-\tilde{s}^*[f_n])}{J[f_n]} = 0 \tag{4.10}$$

$$2\hat{A}\sum_{n=0}^{N-1} \frac{|\tilde{s}[f_n]|^2}{J[f_n]} = \sum_{n=0}^{N-1} \frac{\tilde{s}^*[f_n]\tilde{v}[f_n] + \tilde{s}[f_n]\tilde{v}^*[f_n]}{J[f_n]} \tag{4.11}$$

$$\hat{A} = \left[\sum_{n=0}^{N-1} \frac{|\tilde{s}[f_n]|^2}{J[f_n]}\right]^{-1} \sum_{n=0}^{N-1} \frac{\tilde{s}^*[f_n]\tilde{v}[f_n]}{J[f_n]} \tag{4.12}$$

where in the last line, we have used the fact that $\tilde{x}[f_n] = \tilde{x}^*[f_{N-n}]$ when $x[t]$ is real. The factor in front ensures that the time-domain maximum of the filtered signal template is 1. Thus, the amplitude estimator does not re-scale the signal. Treating this as a filtering of the original waveform $v[f]$, we can define the OF as:

$$\mathrm{OF}[f_n] = h\frac{\tilde{s}^*[f_n]}{J[f_n]} \tag{4.13}$$

where $h = \left[\sum_{n=0}^{N-1} \frac{|\tilde{s}[f_n]|^2}{J[f_n]}\right]^{-1}$ ensures that the filter is unity-gain. This can also be extended to account for a time delay in the signal by minimizing a new χ^2:

$$\chi^2_{OF}(t_0) = \sum_{n=0}^{N-1} \Delta f \, \frac{\left|\tilde{v}[f_n] - As[f_n]e^{-i2\pi f_n t_0}\right|^2}{J[f_n]} \tag{4.14}$$

A similar derivation as the one above gives:

$$\hat{A} = \left[\sum_{n=0}^{N-1} \frac{|\tilde{s}[f_n]|^2}{J[f_n]}\right]^{-1} \sum_{n=0}^{N-1} \frac{\tilde{s}^*[f_n]e^{i2\pi f_n \hat{t}_0}\tilde{v}[f_n]}{J[f_n]} \tag{4.15}$$

where $\hat{t}_0$ is the time delay of the signal, i.e. the time at which the filtered signal reaches its maximum amplitude.

From here, it is useful to define the amplitude resolution σ_A. The resolution is inversely proportional to the signal-to-noise ratio of the optimally filtered pulse shape, i.e.

$$\sigma_A = \left[T\sum_{n=0}^{N-1} \frac{|\tilde{s}[f_n]|^2}{J[f_n]}\right]^{-\frac{1}{2}} \tag{4.16}$$

where factor of T comes from the fact that $J[f_n]$ has units of V^2/Hz, and σ_A has units of Volts. This is a much more realistic description of the amplitude resolution than Eq. 3.6.

Optimum Triggering

Once the OF is constructed for each channel-dataset pair, the continuous data are filtered with the OF and stored as a new run. This must be done offline because the OF is an acausal filter. To reduce both the computation time of the filtering and the amount of data that must be stored, the continuous data and the OF are decimated by a factor of 8 to 125 Hz. This new sampling frequency is chosen because the signal power is much less than the noise power at frequencies above 62.5 Hz for virtually all channels. This decimation is only applied during the application of the OT algorithm; the later processing steps, including amplitude estimation, are done using the data sampled at 1 kHz.

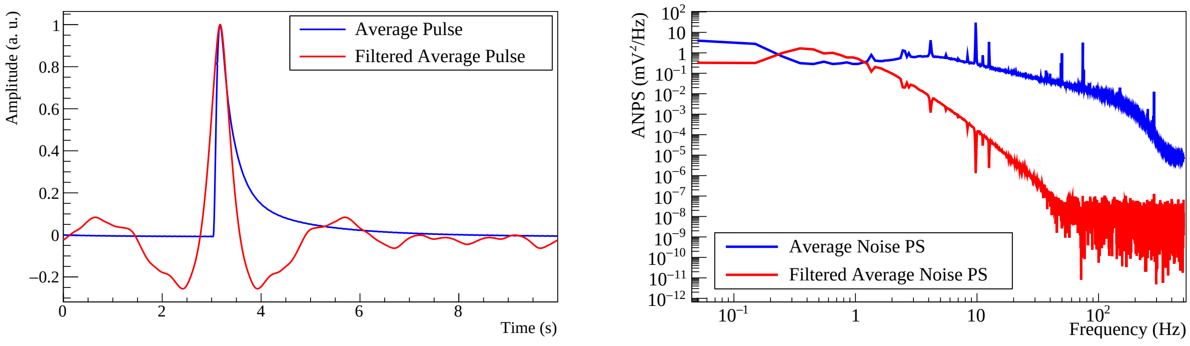

(a) (b)

Figure 4.4: Examples of the effect of Optimum Filter (OF) on the CUORE data. (a) The average pulse shown in 4.2 before and after applying the OF. (b) The average noise power spectrum shown in 4.3 before and after applying the OF.

The OF waveforms are then retriggered using the optimal trigger (OT). The OT is a simple amplitude threshold trigger; if the amplitude of the OF waveform is greater than four times the OF resolution, then a trigger is placed at the local maximum of the waveform. There is a dead time of a few seconds after an OT trigger is set to prevent "double triggering" on the side pulses that are created as a result of the filtering (see Fig. 4.4). After running the OT, the time of each event is saved, just like the events triggered by the DT. In addition to the timestamp, the OF amplitudes are also stored. As the OF amplitude is the best amplitude estimator, it becomes the best energy estimator once the data are stabilized and calibrated (again, using the data sampled at 1 kHz).

4.4 Thermal Gain Stabilization

A unique aspect of the cryogenic calorimetric technique is that changes in temperature in the system have a non-negligible effect on the detector response. As was discussed in section 3.1, the resistivity of the NTD-Ge thermistors decreases as operating temperature increases (see Eq. 3.10). This implies that changes in temperature over the course of a run can change the amplitude of pulses at a fixed energy. The pulser events can be used to account for this effect. The energy of a pulser event in a given channel is the same, with very high precision. Using the event baseline as a proxy for the operating temperature, one can fit a function of the OF amplitude of each pulser versus the baseline. This process is called *thermal gain stabilization*, or simply stabilization. The baseline of each pulser event is fit to a first-order polynomial, and the mean value of the fit is taken to be the baseline voltage. If the magnitude of the slope of this fit is sufficiently large, then the pulser event is excluded since it must have occurred either in coincidence with another pulse, i.e. as a pileup event, or when the detector was unstable, e.g. during an earthquake or when a truck traveling on the highway near LNGS hits a pothole. To perform the stabilization, the OF amplitudes of all pulser events are fit to a first-degree polynomial as a function of the baseline. The baseline dependence on the OF amplitude is rotated away, producing the "stabilized amplitude." This procedure is called heater-based thermal gain stabilization (heaterTGS), and it is used in most CUORE channels.

A similar technique is employed using events from calibration data at known energies. During calibration, the highest-energy peak from γ rays in the data is the 2.615 MeV peak due to the decay of ^{208}Tl to ^{208}Pb. These events are sufficiently high in energy that fluctuations in the baseline do not smear out the peak, allowing one to select the events in the peak before stabilization. Assuming that the energies of these peaks are all at 2.615 MeV and, therefore, should have the same amplitude in a given channel, one can perform a routine similar to the heaterTGS to account for the changes in resistivity due to thermal fluctuations. This procedure is called calibration thermal gain stabilization (calibrationTGS).

For most CUORE channels, both heaterTGS and calibrationTGS are applied. In general, heaterTGS tends to perform better than calibrationTGS. The primary reason for this is that calibration data are only taken at the beginning and end of each dataset. Suppose a channel experiences a sudden jump in temperature during a physics run but returns to its original temperature before the final calibration begins. In that case, there will be no calibrationTGS data at these outlier baseline values, so the fit needs to be extrapolated to these values. In contrast, heater events occur throughout the data-taking, during physics and calibration runs, giving a distribution of events covering a broad range of baseline values and allowing for a more robust fit to the data. The heaterTGS and calibrationTGS methods are compared in Fig. 4.5. Of the 984 active channels deployed in CUORE, 29 have dysfunctional or broken silicon heaters, meaning only calibrationTGS can be used for thermal stabilization. For the other channels, the two stabilization techniques are compared after the detectors are calibrated, and the best technique is selected thereupon.

4.4.1 Calibration

For the first three datasets of CUORE, the Detector Calibration System (DCS) was used to deploy calibration sources into the cryostat and remove them without human intervention. These calibration sources included thoriated tungsten which contained ^{232}Th. For CUORE, this is a natural choice for calibration because the decay products of ^{232}Th include ^{208}Tl, which decays to ^{208}Pb and emits a 2.615 MeV γ ray in the process. This γ ray is close to the Q-value of ^{130}Te, so the detector response at energies close to $Q_{\beta\beta}$ can be studied well. Other γ ray peaks observed in the calibration spectrum include the 338, 911, and 969 keV γ rays from the de-excitation of ^{228}Ra, the 583 keV γ ray from the decay of ^{208}Tl, and the 239 keV γ ray from the decay of ^{212}Pb, all of which are decay products from the ^{232}Th chain. These energies were most commonly used to calibrate the detector for these datasets. The DCS was designed to irradiate the detector uniformly, and each detector had effectively the same rate of roughly 100 mHz. This rate provided each detector with sufficiently high statistics

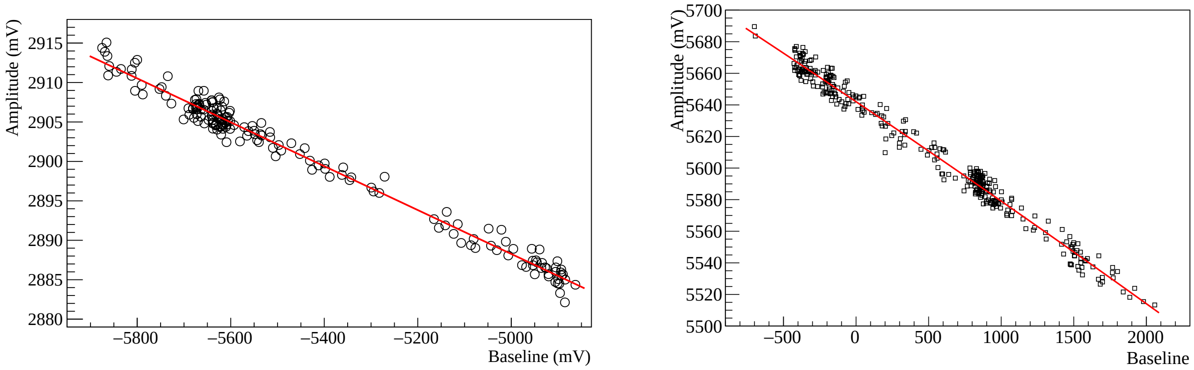

(a) (b)

Figure 4.5: Examples of the thermal gain stabilization (TGS) fit of amplitude versus baseline for two different channels. (a) An example of heaterTGS in which heater pulses are used to find the stabilization curve. (b) An example of calibrationTGS in which events from the 2.615 MeV peak are used to find the stabilization curve. In this case, the data cover the space of different baseline values well, but this is not guaranteed to be the case.

to calibrate the detector while avoiding the problem of too many pile-up events within a ten-second window.

After the first three datasets, the DCS was retired and replaced with the External Detector Calibration System (EDCS). The EDCS is a set of strings deployed outside the cryostat's lead shield. The EDCS initially used ^{232}Th sources like the DCS, but after one dataset, ^{60}Co were added to the calibration strings. In addition to those provided by the ^{232}Th decay products, the ^{60}Co sources also provide γ rays at 1.173 and 1.332 MeV. In practice, for the datasets utilizing the EDCS, the lines most commonly used to calibrate the detector are the two ^{60}Co peaks and the 583 keV and 2.615 MeV peaks from ^{208}Tl. Because the EDCS sources are all outside the lead shield, the channels toward the outer edge of the detector have a much higher calibration rate than the inner channels. For a typical dataset, the rate of events with energies close to 2.615 MeV is roughly 40 to 50 mHz for the outer channels but closer to 10 to 20 mHz for the inner channels. An example energy spectrum from a dataset calibrated with the EDCS including ^{60}Co and ^{232}Th is shown in Fig. 4.6.

The energies of the aforementioned γ peaks are well-known thanks to various experimental studies over the decades. These energies are used to fit the distribution OF amplitudes to a quadratic function of energy that crosses through the origin, i.e.

$$E(A) = c_1 A_S + c_2 A_S^2 \tag{4.17}$$

where E is the calibrated energy, A_S is the stabilized amplitude, and c_1 and c_2 are determined uniquely for each detector. The lack of a constant term asserts that a zero-amplitude event corresponds to zero energy, which should always be true. The quadratic term accounts for

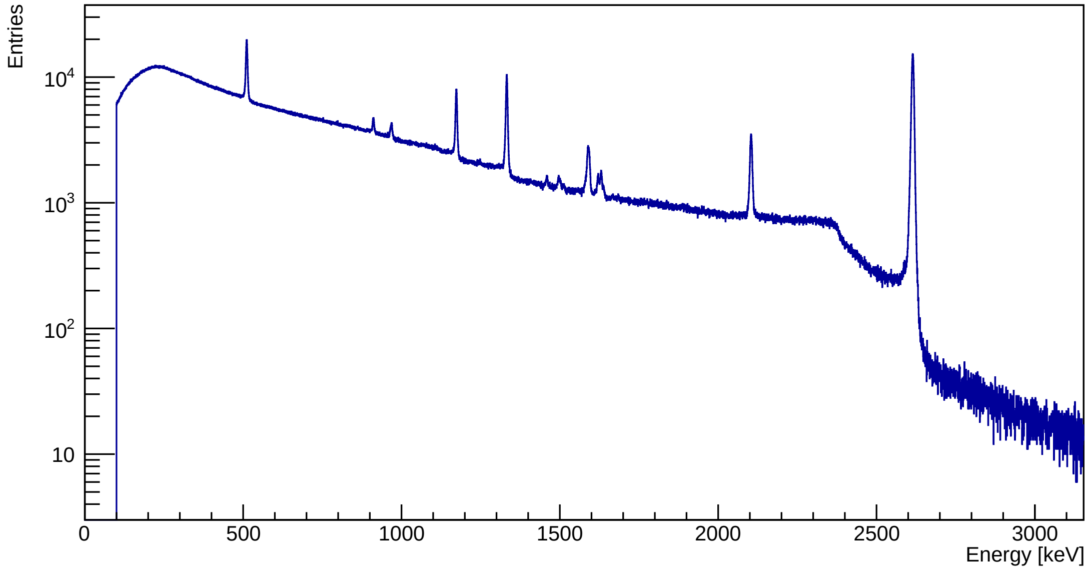

Figure 4.6: An energy spectrum built from the calibration events of one dataset. This dataset used the external detector calibration system (EDCS), and the spectrum is dominated by events coming from ^{60}Co and the decay products of ^{232}Th.

the fact that the detector response is not truly linear, but in practice, this term has a small effect on the calibration.

The calibration is run twice for channels with functional silicon heaters: once with the heaterTGS amplitudes and once with the calibrationTGS amplitudes. A separate analysis module called the Energy Selector determines which of the two thermal gain stabilization techniques is used to produce the final energy spectrum. The Energy Selector examines the calibration spectrum for each channel and selects the stabilization method that delivers the best energy resolution at 2.615 MeV while ensuring maximal analyzed exposure.

4.5 Coincidence Analysis

As described in Chapter 3, CUORE consists of multiple individual detectors arranged closely together into towers, and the towers are placed close to each other. The fact that CUORE is a segmented detector significantly enhances our ability to reduce radioactive backgrounds by tagging events. By utilizing the proximity of these channels, CUORE can effectively identify coincidental events occurring in adjacent detectors. When a particle interaction simultaneously results in signals in nearby channels, the event can be tagged as having a *multiplicity* greater than one. Of course, how one defines “simultaneous” and “nearby” is relative. For the $0\nu\beta\beta$ analysis, any events occurring within 5 ms of each other in detectors separated by at most 30 cm are considered part of the same multiplet. We also place an energy threshold of 40 keV to avoid false coincidences with noise. As a shorthand, we refer to multiplets containing N events as a "multiplicity-N," or simply "M-N," events.

A classic example of an event that the coincidence analysis can reject is the α decay of ^{210}Po to ^{206}Pb, which has a Q-value of 5.4 MeV. ^{210}Po is a decay product from the ^{238}U chain with a half-life of roughly 138 days. In the context of CUORE, it is a contaminant typically found on the surface of the TeO_2 crystal. When it decays, the α particle and the daughter nucleus travel in opposite directions, depositing energy in different detectors; the α particle carries away 5.3 MeV of kinetic energy, and the daughter nucleus recoils with 0.1 MeV of kinetic energy. The result is two events in nearby detectors above the 40 keV threshold, which can be tagged as an M2 event together. Other examples of high-multiplicity events include γ rays which Compton scatter from one detector into another and muons which pass through the detector, depositing energies in many detectors simultaneously.

From Monte Carlo simulations, it is expected that roughly 88% of the $0\nu\beta\beta$ decays in CUORE are expected to be M1 events [163]. Based on these simulations and measurements of the relevant backgrounds, it is expected that the experiment's sensitivity to $0\nu\beta\beta$ decay can be enhanced by analyzing only M1 events and cutting the higher multiplicity events with an anti-coincidence cut.[1]

4.6 Pulse Shape Discrimination

4.6.1 Principal Component Analysis

The last step in the analysis chain is pulse shape discrimination (PSD). In CUORE, this is done using principal component analysis (PCA). PCA is a statistical technique used to reduce the dimensionality of a high-dimensional dataset while retaining information about the underlying trends and variances in the data. In the context of CUORE, we run PCA separately for each channel, viewing each 10-second event (sampled at 1 kHz) as a vector $\vec{x}_n$ with dimension $D = 10^4$, where n indexes over the number of events in the dataset-channel. The goal of PCA is to rotate the data into a new D-dimensional orthonormal basis $\{u_i\}$, where i indexes over the D dimensions of the vector, then reduce the dimensionality by projecting the data onto the set of new basis vectors that contains the largest variance in the data. The variance of the ith is given by [164]:

$$\frac{1}{N}\sum_{n=1}^{N}(\vec{u}_i \cdot \vec{x}_n - \vec{u}_i \cdot \langle \vec{x} \rangle)^2 = \vec{u}_i \cdot S \cdot \vec{u}_i \tag{4.18}$$

---

[1] Though they are cut from the primary $0\nu\beta\beta$ analysis, M2 events in which one of the two produced electrons escapes to another detector, or both electrons escape into the same detector while the nuclear recoil occurs in the original detector of the decay, are expected to account for 6% of the $0\nu\beta\beta$ decays in CUORE and a search for $0\nu\beta\beta$ decay with M2 events is ongoing.

where S is the covariance matrix of the data:

$$\frac{1}{N}\sum(\vec{x}_n - \langle\vec{x}\rangle) \otimes (\vec{x}_n - \langle\vec{x}\rangle) \tag{4.19}$$

The basis vectors $\vec{u_i}$ are found by maximizing $\vec{u_i} \cdot S \cdot \vec{u_i}$ for $i = 1$, then $i = 2$, etc. This is done using D Lagrange multipliers λ_i while imposing orthonormality and maximizing the expression with respect to $\vec{u_i}$, i.e.:

$$\frac{\partial}{\partial\vec{u_i}}\left[\vec{u_i} \cdot S \cdot \vec{u_i} + \lambda_i(1 - \vec{u_i} \cdot \vec{u_i})\right] = 0 \tag{4.20}$$

Solving this yields:

$$S \cdot \vec{u_i} = \lambda_i \vec{u_i} \tag{4.21}$$

thus the vectors $\vec{u_i}$ are simply the eigenvectors of the covariance matrix of the data. Maximizing the variance of the projections of the data onto $\vec{u_i}$ is equivalent to projecting onto the eigenvector with the largest corresponding eigenvalue. We call these eigenvectors the *principal components* of the data. The problem of reducing the dimension of the data now amounts to choosing the first M principal components of the data.

When analyzing the CUORE data, one cannot simply take all events that survive through to the PSD step of the analysis and dump them into the PCA routine. PCA is highly sensitive to outliers in the data since the outliers will contribute heavily to the variance of the data. In the CUORE data, the outliers tend to be what we want to reject, and blindly using PCA without some event selection would not achieve this goal. To remedy this, we first select events likely to be "good" physical events, i.e. M1 events with energies within 40 keV of selected γ peaks from calibration data. These events serve as the training set from which we derive the principal components. In CUORE, we use only the first principal component of the data, which is extremely similar to the AP used to construct the OF. We use the AP as a proxy for the first principal component, which reduces the need to compute the entire eigenbasis for the dataset.

4.6.2 Event selection with PCA Reconstruction Error

While PCA is able to reduce the dimensionality of the signals in the dataset effectively, the problem of using it to cut spurious events remains. To build a quantity on which we can perform a cut, we define the reconstruction error (RE) of a given event as:

$$RE \equiv \sqrt{\sum_{i=1}^{D}(x_i - \vec{x} \cdot \vec{u}_1)u_i} \tag{4.22}$$

where $\vec{u}_1$ is the AP, our proxy for the first principal component. Events similar in shape to the AP will have a small RE, and those that differ from the AP will have a large RE. This quantity will naturally scale with the amplitude of the pulse, so we must also include an energy-dependent normalization of the RE. We define the normalized reconstruction error (NRE) as:

$$\text{NRE} = \frac{\text{RE} - f(E)}{\text{MAD}} \tag{4.23}$$

where E is the energy and MAD is the median average distance of the RE from $f(E)$. $f(E)$ is a monotonically increasing second-order polynomial found via an iterative fitting procedure described in [164]. The procedure ensures that the bulk population of the events is well modeled by $f(E)$, while outliers are free to be arbitrarily far away from $f(E)$. Allowing this normalization to be nonlinear with respect to energy also helps to account for any non-linearity in the detector response.

To determine the value of NRE on which to cut, we calculate the value of NRE for each event in a large window around the 2.615 MeV γ ray peak. We then use the quantity $\epsilon_S/\sqrt{\epsilon_B}$ as the figure of merit for the cut value, where ϵ_S and ϵ_B are the signal and background efficiencies of the cut, respectively.

When analyzing the data, one must choose the order of the cuts applied. One could easily apply the PSD cut to the individual pulses and then apply the anti-coincidence cut using only the waveforms that pass PCA, but in CUORE, the coincidence analysis is run first, and then the PSD cut is applied. The advantage of this choice is that we can select strictly M1 events when building the NRE function. A disadvantage is that the multiplets selected during the coincidence analysis may contain events rejected by PCA. A separate analysis module called the Multiplet Validator ensures that each multiple includes only events that pass the PSD cut. If an event in a multiplet is cut by PSD, then the entire multiplet is rejected from the analysis. This is a very conservative choice, as one could also decrease the multiplicity of the event by the number of cut events. This choice was made to maximize the effect of the anti-coincidence cut, which is appropriate in the context of the $0\nu\beta\beta$ decay search with M1 events, though it is not suitable for other analyses such as the M2 $0\nu\beta\beta$ decay search or analyses of delayed coincidences to measure specific radioactive backgrounds.

Chapter 5

Measuring and Characterizing Noise in CUORE

The CUORE detectors are subject to noise from a host of sources including thermal fluctuation noise in the detector, Johnson-Nyquist noise in the readout system, vibrational noise, and electrical noise. An ANPS of a typical CUORE detector is shown in Fig. 5.1. This chapter is dedicated to understanding the various sources of noise, the steps taken to mitigate the noise, and the techniques used to monitor the noise that persists.

5.1 Sources of CUORE Noise

Thermal noise

Johnson-Nyquist noise is an intrinsic electronic noise arising from the thermal motion of charge carriers inside a conductor. This noise is present in all conductors and is independent of any applied voltage. The RMS voltage of Johnson-Nyquist noise for a given bandwidth Δf is:

$$\sqrt{\langle V_J^2 \rangle} = \sqrt{4kTR\Delta f} \tag{5.1}$$

where k is Boltzmann's constant, T is the temperature, and R is the resistance. Johnson-Nyquist noise from the readout circuit contributes significantly to the continuum noise in CUORE; the largest sources are the bias resistor and the NTD-Ge thermistor. The maximum possible transfer of noise power within the band is:

$$P_{max} = \frac{\langle V_J^2 \rangle}{R} = 4kT\Delta f \tag{5.2}$$

The thermal system of a CUORE detector is much more complicated than a simple conductor. Instead, the power transfer is better described by a non-linear energy-dependent electro-thermal model [165]. This model includes a contribution to the noise power arising from the random fluctuations in the energy transfer from the electron system to the phonon

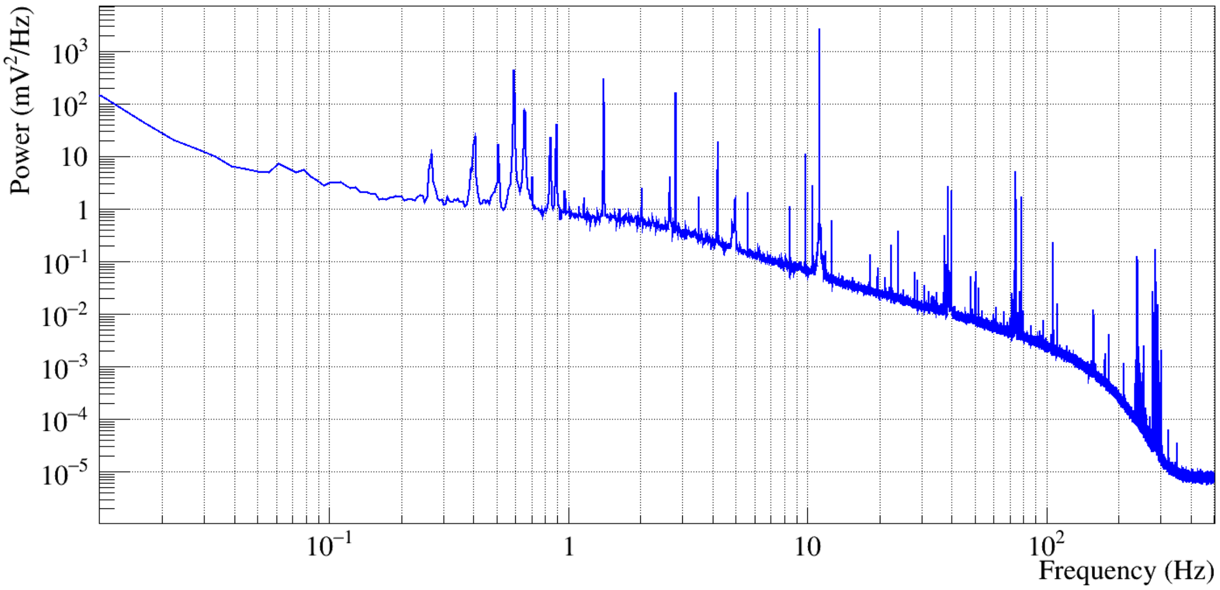

Figure 5.1: A typical ANPS of a CUORE detector using 180-second long noise windows. The shape of the continuum is predicted by an electro-thermal model described in [165]. The peaks at the harmonics of 1.4 Hz are due to the pulse tubes, while those at the harmonics of 50 Hz are due to parasitic electrical noise from the AC power lines. The sub-Hz peaks are likely due to the mechanical resonances of the support structure, including the Y-beam and the minus-K mechanical insulators.

system. Using the hot electron model [166], we assign this energy transfer an effective thermal conductivity $C = G_0 T^\beta$, so the power transfer is [167]:

$$P = \frac{G_0}{\beta + 1}\left(T_e^{\beta+1} - T_n^{\beta+1}\right) \tag{5.3}$$

where T_e is the electron temperature, T_n is the crystal lattice temperature. The constants G_0 and β can be fit to data measuring the device's resistance versus the applied bias power and scanning over base temperatures. For semiconductors including silicon and germanium, $\beta \sim 5$, rather than the expected value of 4 for metals. While still independent of frequency, this thermal fluctuation noise has a different temperature dependence than the Johnson noise.

Pulse Tube noise

The peaks in a typical CUORE ANPS arise from a plethora of sources. The most obvious of these is the pulse tubes, at least four of which are constantly pumping on the cryostat during data-taking. Mechanical vibrations from the pulse tube heads are transmitted to the 300 K plate, while the vibrations from the two cold pulse tube stages are transmitted to the 40 K and 4 K plates. While the detector suspension system is designed to mechanically isolate the detector from the rest of the cryostat, this isolation is imperfect, and vibrations often propagate from the warmer stages to the detector. Vibrations from the pulse tubes can also cause the signal wires to vibrate, inducing capacitive pickup which is subsequently read out. The pulse tubes operate at a fundamental frequency of 1.4 Hz, and their main effect on the CUORE detector ANPS is a series of peaks at the integer multiples of 1.4 Hz.

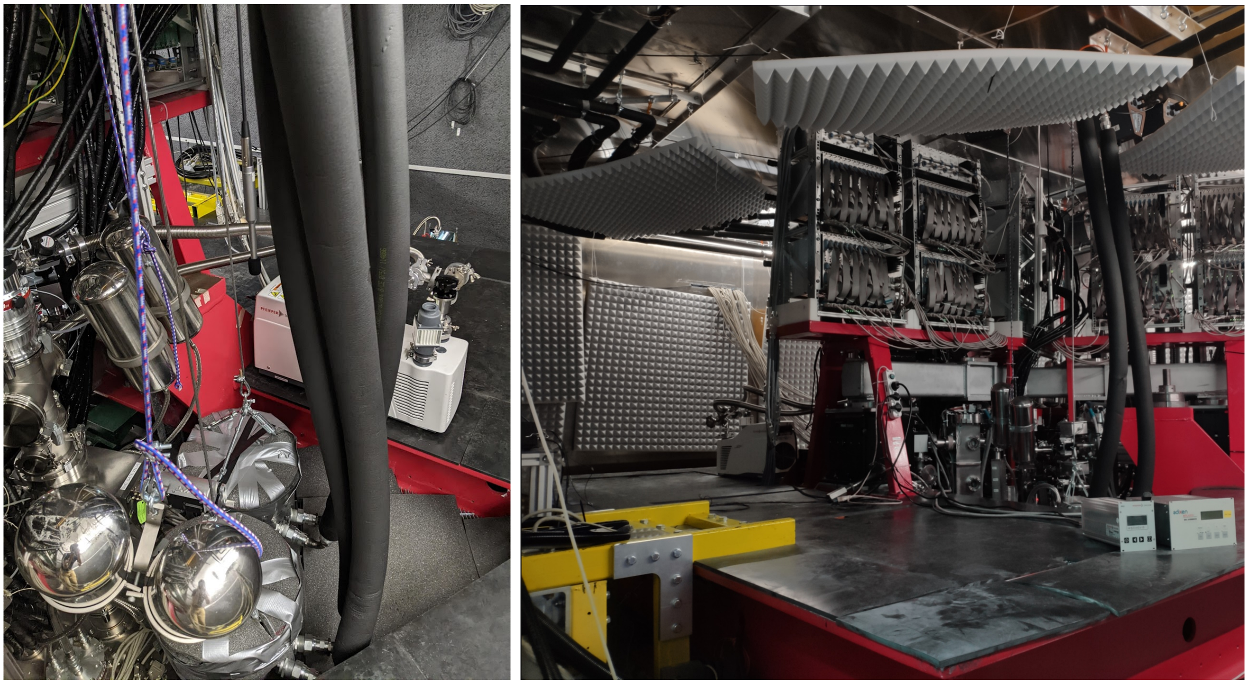
(a) (b)

Figure 5.2: Passive steps taken to reduce noise from pulse tubes in CUORE. (a) The pulse tube motor-heads are placed in sand-filled buckets, covered in foam, and suspended from the ceiling. The two buffer volumes are also suspended from the ceiling, and the compressor lines are covered in foam. (b) Acoustic foam is placed all around the Faraday room to reduce the ambient noise from the pulse tubes, compressor lines, and anthropogenic activity.

CUORE employs many passive techniques to reduce the vibrational noise from the pulse tubes. The pulse tube motor-heads are mechanically decoupled from the DR via flexible steel lines and are placed in sand-filled buckets covered in foam. Along with the compressor lines and buffer volumes, the buckets are suspended from the ceiling of the Faraday cage with elastic bands to further reduce the effect of vibrations on the DR (see Fig. 5.2a). In addition, acoustic foam has been placed around the walls and ceiling of the Faraday room to reduce the ambient acoustic noise due to the pulse tubes, compressor lines, and anthropogenic activities (see Fig. 5.2b). Finally, CUORE uses stepper motors with linear amplifiers to drive the pulse tubes. These "linear drives" smooth the rotation of the pulse tube rotary valves by dividing a full rotation into 25600 steps. The linear drives further reduce vibrations compared to the typical Cryomech drives usually employed for these pulse tubes.

While these passive techniques help mitigate vibrational noise, the effect of the pulse tubes on the detector is still significant. To further reduce the noise, CUORE employs an active noise cancellation technique that controls the relative phases of the pulse tubes [168]. The linear drive software allows one to identify relative phase configurations, move between phase configurations, and maintain a selected phase configuration via a feedback system. An optimal phase configuration is found by scanning through the space of possible phase configurations and measuring the noise at each configuration. Fig. 5.3 shows the median normalized noise as a function of phase configuration. The noise is defined as a weighted sum over the first ten harmonics of 1.4 Hz:

$$\mathcal{N}_c(\vec{\phi}) = \sum_{n=1}^{10} \mathrm{AP}_c(f_n)\, \mathcal{G}_c^{-1}\, \mathrm{PS}(f_n, \vec{\phi}) \tag{5.4}$$

where AP_c is the power spectrum of the average pulse of channel c, f_n is the nth harmonic of 1.4 Hz, $\mathcal{G}_c$ is the front-end gain of the electronics for channel c, PS is the power spectrum of the noise sample, and $\vec{\phi}$ is the particular phase configuration. The average pulse serves as the signal template for the channel, and it is included in the weighting so that noise in the signal band contributes more to the sum. The normalization is defined such that the noise over the entire detector is in the range $[0, 1]$ while also allowing for a comparison of the channels in terms of their relative noise regardless of their absolute noise levels.

In practice, any cryogenic maintenance that requires manual interventions in the Faraday room causes changes to the mechanical system, thus changing the optimal phase configuration. This means that conducting a pulse tube phase scan between datasets is usually necessary to find the new minimal-noise phase configuration. After a phase scan is complete, the best phase configuration is chosen, and the pulse tube software can ensure that this configuration is maintained to within a few degrees.

Low-frequency noise

Besides pulse tube noise, the other main contribution to the noise in the signal band is a set of low-frequency peaks below 1 Hz. These peaks are induced by microseismic activity, including small earthquakes and waves from the Tyrrhenian and Adriatic seas [169]. While microseismic activity from the sea waves tends to be broad-band noise ($0.02 \lesssim f \lesssim 1$ Hz), the peaks in the CUORE detector ANPS are likely due to the resonances of the support structure, particularly the Y-beam and minus-K mechanical insulators. This was verified by conducting vibrational tests and mechanically exciting the Y-beam. A schematic of the setup for these tests is shown in Fig. 5.4. For these tests, one accelerometer was placed on each arm of the Y-beam roughly equidistant from the center and oriented vertically. The Y-beam was then excited by pushing on it (gently so as not to damage the detectors) and

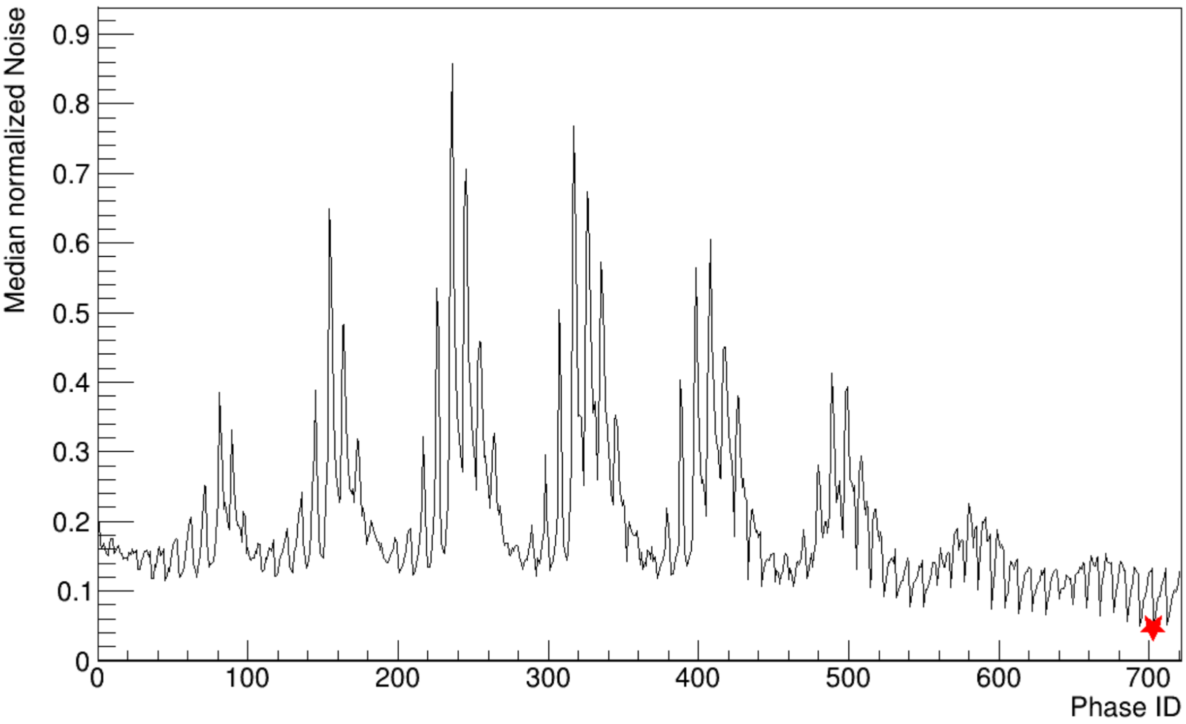

Figure 5.3: An example of how the CUORE noise changes during a pulse tube phase scan. The y-axis is the normalized noise in the median channel, and the phase ID is a unique identifier corresponding to a particular configuration of relative phases between pulse tubes. The red star identifies the optimal phase ID at which the normalized noise is minimized.

measuring the power spectrum of different accelerometer signal combinations. Adding the three accelerometer timestreams together approximates the Y-beam's vertical motion while canceling any differences between the different accelerometers due to rocking about any one arm. Subtracting one timestream from another approximates the rocking motion between the two arms while canceling out the vertical motion that both accelerometers experience simultaneously. Taking the power spectrum of the resulting signals reveals the resonant structure of the Y-beam (see Fig. 5.5). The vertical motion of the Y-beam exhibits a peak at roughly 0.7 Hz, while the rocking motion exhibits a peak at roughly 2.4 Hz. By adjusting the stiffnesses of the minus-K insulators, these peak frequencies can be changed by approximately 20%.

The low-frequency noise observed in the CUORE detectors is not due to only the Y-beam resonances. As was seen in Fig 5.1, a host of peaks at frequencies between 0.2 and 0.9 Hz are prominent in the ANPS of the detector. Some of these are likely due to beating between different mechanical resonances of the structural supports. Beating between pulse tube harmonics and the Y-beam can produce noise peaks at $f = (1.4n - 0.7)$ Hz and

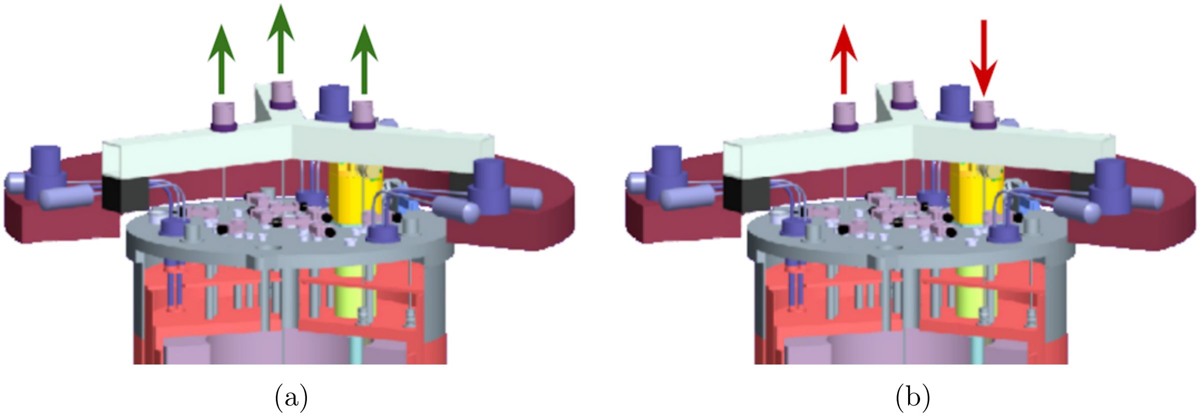

Figure 5.4: Schematic of the setup for the Y-beam vibrational tests. For these tests, one accelerometer was placed on each arm of the Y-beam and oriented vertically. (a) The time-domain sum of all three accelerometer signals approximates the vertical motion of the Y-beam. (b) The time-domain difference of each pair of accelerometers approximates the rocking motion of the Y-beam.

$f = |1.4n - 2.4|$ Hz for $n \in \mathbb{Z}^+$. For example, a peak at 0.4 Hz is featured prominently in the ANPS, and this may be due to a beating between the Y-beam rocking mode at 2.4 Hz and the second harmonic of the pulse tube noise at 2.8 Hz.

After conducting the resonance tests on the Y-beam, we also pushed on the Main Support Platform to excite its resonances and measured the resulting ANPS of the accelerometers. These ANPS are shown in Fig. 5.6. The Main Support Platform clearly has a sizeable mechanical resonance at ∼2 Hz and another less prominant resonance at ∼8.7 Hz. Excitations of this resonance can mix with the Y-beam rocking mode to produce beats at $f = |1.4n - 2|$ Hz, which may explain the prominent 0.6 Hz peak in the ANPS.

While these mechanical excitation tests on the Y-beam help us understand the sources of the noise peaks, the CUORE suspension and support structures are very complex, and we do not yet completely understand the mechanical system. Work is ongoing to simulate the mechanical resonances of the structure using 3D CAD models. Still, even with advanced computational tools, many simplifications to the structure need to be made for the simulations to be computationally feasible, and these simplifications could eliminate a structural component that gives rise to a particular noise peak.

High-frequency Peaks

Peaks above 10 Hz are also prominent in the detector ANPS. They are much less relevant than the lower-frequency peaks since the CUORE calorimeter signals are slow, and the bandwidth is only ∼50 Hz wide, but it is worth mentioning a few of these peaks and their sources. Electrical peaks at the harmonics of 50 Hz appear in some detectors more prominently than

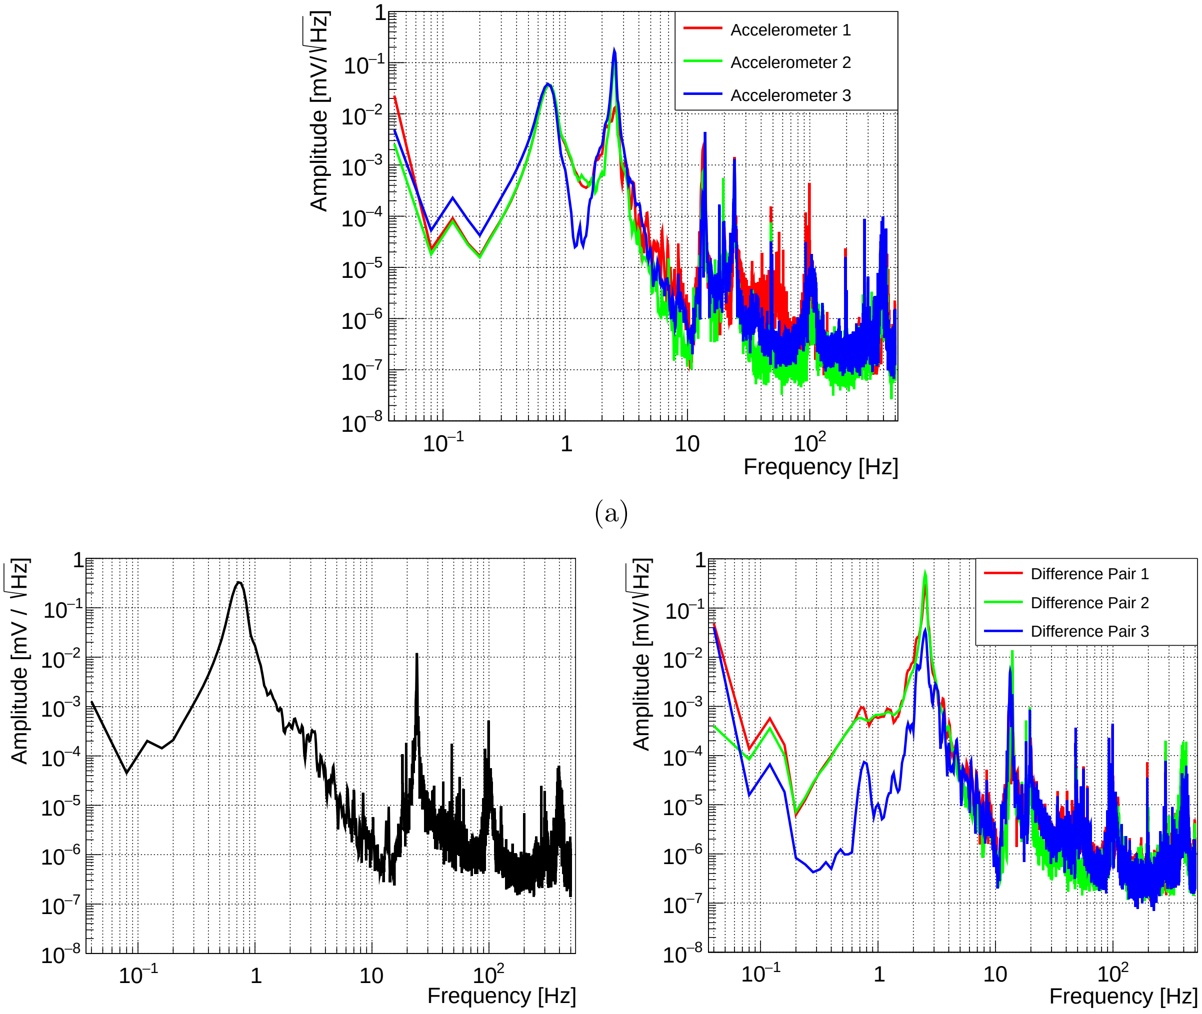

(a)

(b)

(c)

Figure 5.5: Power spectra of different accelerometer combinations during Y-beam tests. For these tests, one accelerometer was placed on each arm of the Y-beam roughly equidistant from the center. (a) The power spectrum of each accelerometer. Each accelerometer exhibits two large peaks at 0.7 Hz and 2.4 Hz. (b) The power spectrum of the time-domain sum of all three accelerometers exhibits only the 0.7 Hz peak, indicating that this is the resonance of the vertical motion of the entire Y-beam. (c) The power spectra of the time-domain difference of each pair of accelerometers. All three pairs exhibit the 2.4 Hz peak, while the 0.7 Hz peak is largely suppressed, indicating that the 2.4 Hz resonance is due to the rocking motion of the Y-beam.

in others, depending on the grounding of the detector. In addition, the pulse tube motors spin at a fixed frequency creating electrical peaks at harmonics of 35 Hz. Finally, the compressor

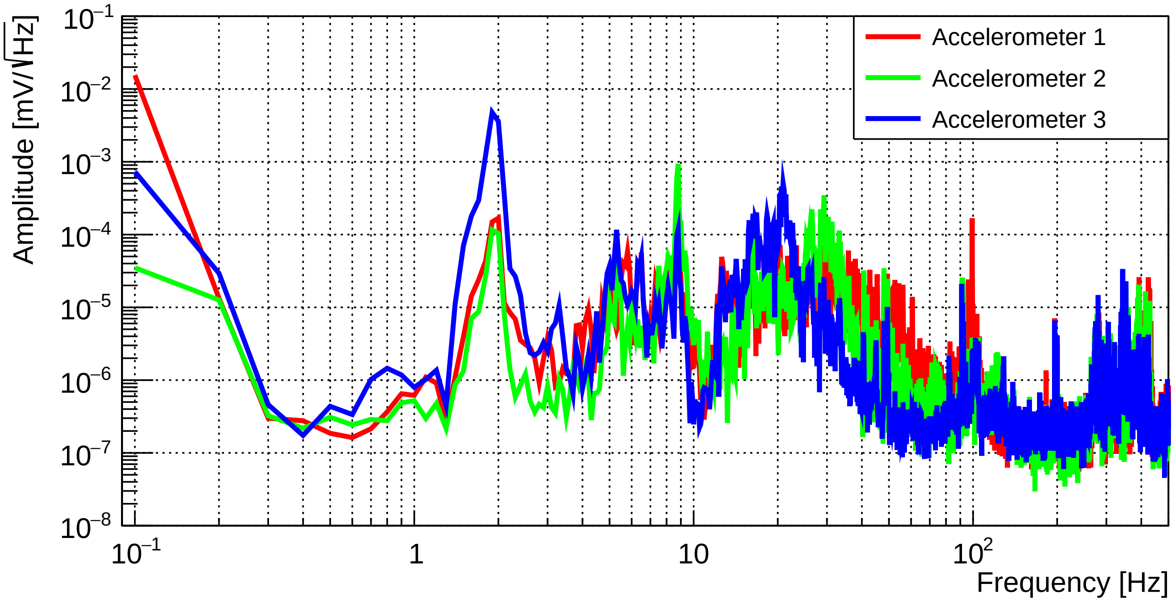

Figure 5.6: Power spectra of accelerometers during Main Support Platform tests. The accelerometers were still located on each arm of the Y-beam, but the Y-beam was not directly excited. Instead, we pushed on the Main Support Platform to excite its mechanical resonances. There is a clear resonance at 2 Hz and another less prominent resonance at ∼8.7 Hz.

lines that feed into each pulse tube also vibrate, giving rise to peaks at harmonics of ∼140 Hz.

5.2 Monitoring Noise with Auxiliary Devices

While many steps are taken to reduce the detector noise, virtually every channel is still subject to noise from extrinsic sources, including vibrational and electrical noise. Since this noise cannot be removed with the current noise suppression techniques, the next best thing to do is monitor it and try to remove it in post-processing. This is done using diagnostic devices, or *auxiliary devices*, which are specially designed to measure particular noise sources of interest.

Deploying auxiliary devices to monitor transient noise in an experiment is a common practice. For example, the the Laser Interferometer Gravitational-Wave Observatory (LIGO) [170] and Virgo [171] experiments are subject to anthropogenic microphonic noise and are sensitive to transient noise during earthquakes, so they use seismometers, accelerometers, magnetometers, and other devices to monitor non-Gaussian transient noise near their detectors [172, 173]. Analyses of these data are typically conducted over a range of frequencies above ∼20 Hz [173], though this cutoff could be reduced with the reduction of anthropogenic noise below 20 Hz [174]. In the case of charge-collecting detectors used in modern γ ray spectroscopy experiments, vibrational disturbances can change the capacitance of the system, leading to microphonic noise which degrades the detector performance, including its energy resolution [175].

Compared to the cost of typical detectors used in experiments today, auxiliary devices are typically affordable while providing sufficiently high fidelities to monitor the relevant noise sources. When the auxiliary data is correlated with the noise in the detectors of interest, noise decorrelation algorithms can be applied to the data. The result is a relatively low-cost way to reduce noise without major hardware interventions that can seriously impact an ongoing experiment. Noise cancellation algorithms using data from accelerometers have already been demonstrated [176] for possible use in the GRETINA [177] and GRETA [178] experiments. Noise cancellation techniques were also developed and applied to data from the CUORICINO [130] and CUORE-0 [131] experiments by decorrelating the noise from different detector channels [179, 180]. The use of noise cancellation algorithms in CUORE is described in detail in Chapters 6 and 7.

5.2.1 Auxiliary Devices in CUORE

The first auxiliary devices in CUORE were installed in 2017 when four Behringer ECM8000 microphones [181] were placed inside the Faraday room. Three microphones were placed close to the pulse tubes, while another was placed further away from the detector to measure the ambient noise in the Faraday room. Some examples of the microphone placement are shown in Fig. 5.7.

These microphones accurately measure the noise in their designed frequency range of 20 Hz to 20 kHz. The upper range is more than enough to measure any high-frequency noise in the CUORE detectors, but thanks to the slow response of the calorimeters, the relevant signal band is well below the 20 Hz cutoff. Fig. 5.8 shows the ANPS of one of the microphones placed near the pulse tubes overlaid on top of the ANPS of a sample CUORE channel. The roll-off of the microphone response below 20 Hz is clearly visible, and the 1.4 Hz and 2.8 Hz peaks cannot be seen in the microphone ANPS. At 280 Hz, however, the microphones can easily measure the noise peak from the compressor lines which feed into the pulse tubes.

In November 2020, a set of three PCB 393B31 accelerometers [182] were mounted triaxially to MSPs. These accelerometers are specifically designed to measure noise at frequencies down to ∼1 Hz. In August 2021, these accelerometers were moved to the Y-beam. An example ANPS from an accelerometer on the Y-beam is shown in Fig 5.9. The 1.4 Hz peak and the structural resonance at 2 Hz are clearly visible, as are the other low-frequency pulse tube harmonics.

In May 2022, four loop antennas were placed around the four active pulse tube motor heads to measure the electrical noise at harmonics of 35 Hz which emanate from the pulse tubes, and another two antennas were placed inside the Faraday room to measure noise due to ground loops at harmonics of 50 Hz. At the same time, an accelerometer set identical to the aforementioned one was installed on the Main Support Platform. This allowed us to

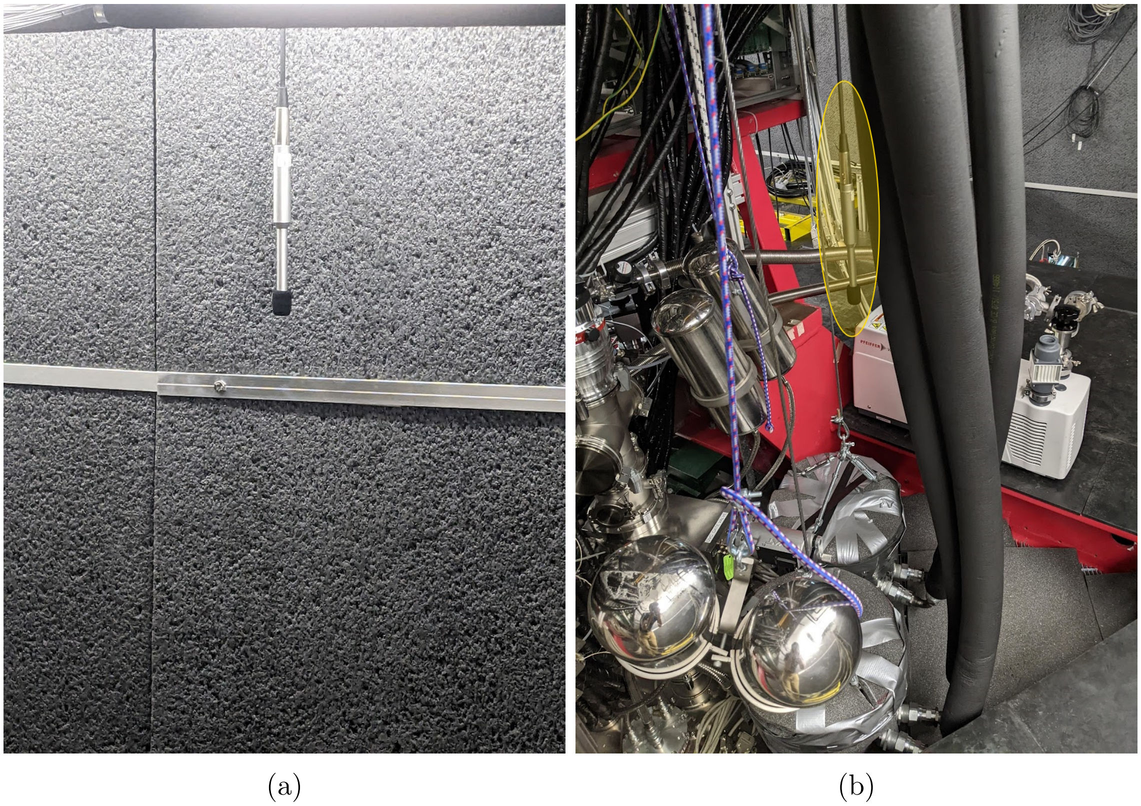
(a) (b)

Figure 5.7: Two of the four Behringer ECM8000 microphones in the CUORE Faraday room. (a) A microphone placed away from the detector, near the wall of the Faraday room. (b) A microphone (circled and highlighted in yellow for clarity) placed above two pulse tube motor-heads, each located in a bucket (below).

measure the vibrations on the Main Support Platform and the Y-beam with accelerometers at the same time.

Finally, in August 2023, a triaxial Sara Electronics SS-10 seismometer [183] was mounted to the Y-beam, allowing us to measure frequencies below 1 Hz with higher sensitivity than the accelerometers could. The seismometer is sensitive to microseismic activity below 1 Hz. This microseismic activity varies considerably over the course of a dataset, as shown in Fig. 5.10. The seismometer also observes some of the peaks due to the resonant structure at 0.6 Hz and 0.8 Hz, and the intensity of these peaks also varies in time with the microseismic activity. This indicates that the microseismic noise is one of the main drivers of the noise at these frequencies in the CUORE detectors.

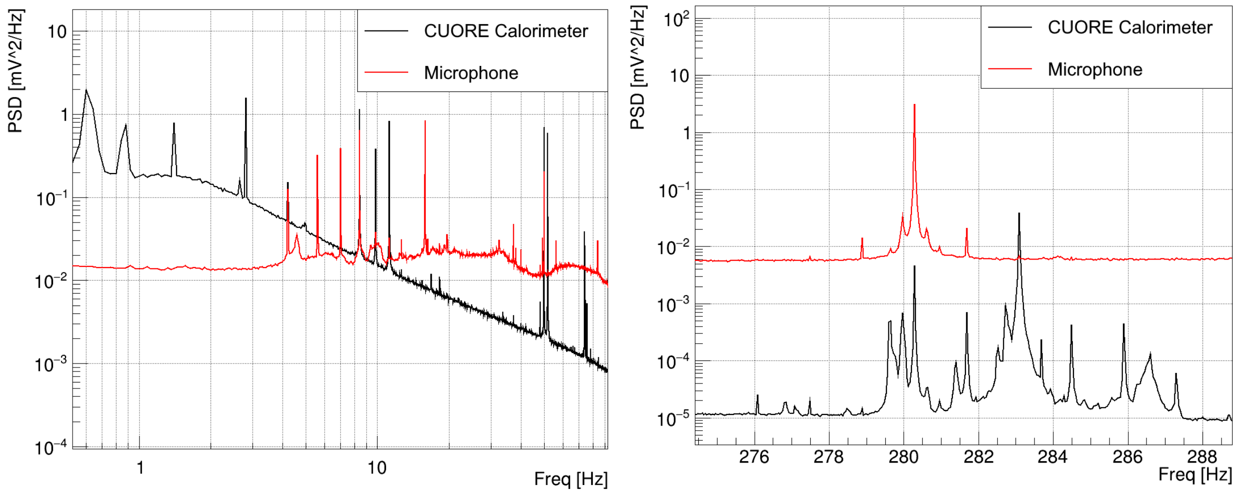

(a) (b)

Figure 5.8: Comparison of the ANPS of a CUORE detector and a microphone placed near the pulse tubes. (a) ANPS of frequencies below 90 Hz. The electrical noise at 50 Hz and the pulse tube noise at the harmonics of 1.4 Hz are clearly visible. (b) ANPS at frequencies near 280 Hz. The microphone exhibits one main peak and a few side peaks at these frequencies, but the CUORE detector exhibits a “forest” of peaks spanning several Hz.

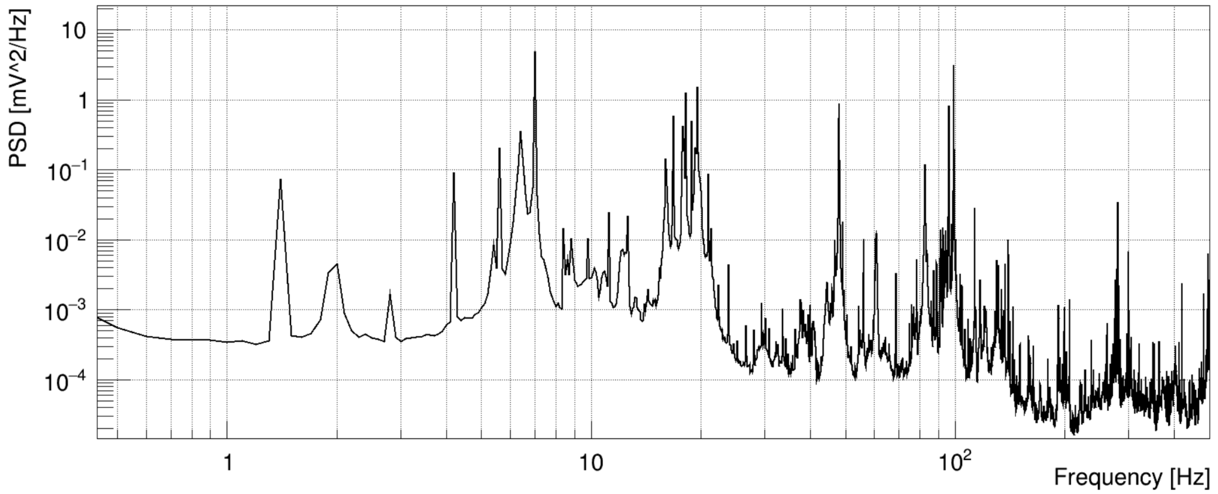

Figure 5.9: An example ANPS from an accelerometer on the Y-beam. The 1.4 Hz peak and the structural resonance at 2 Hz are clearly visible, as are the other low-frequency pulse tube harmonics.

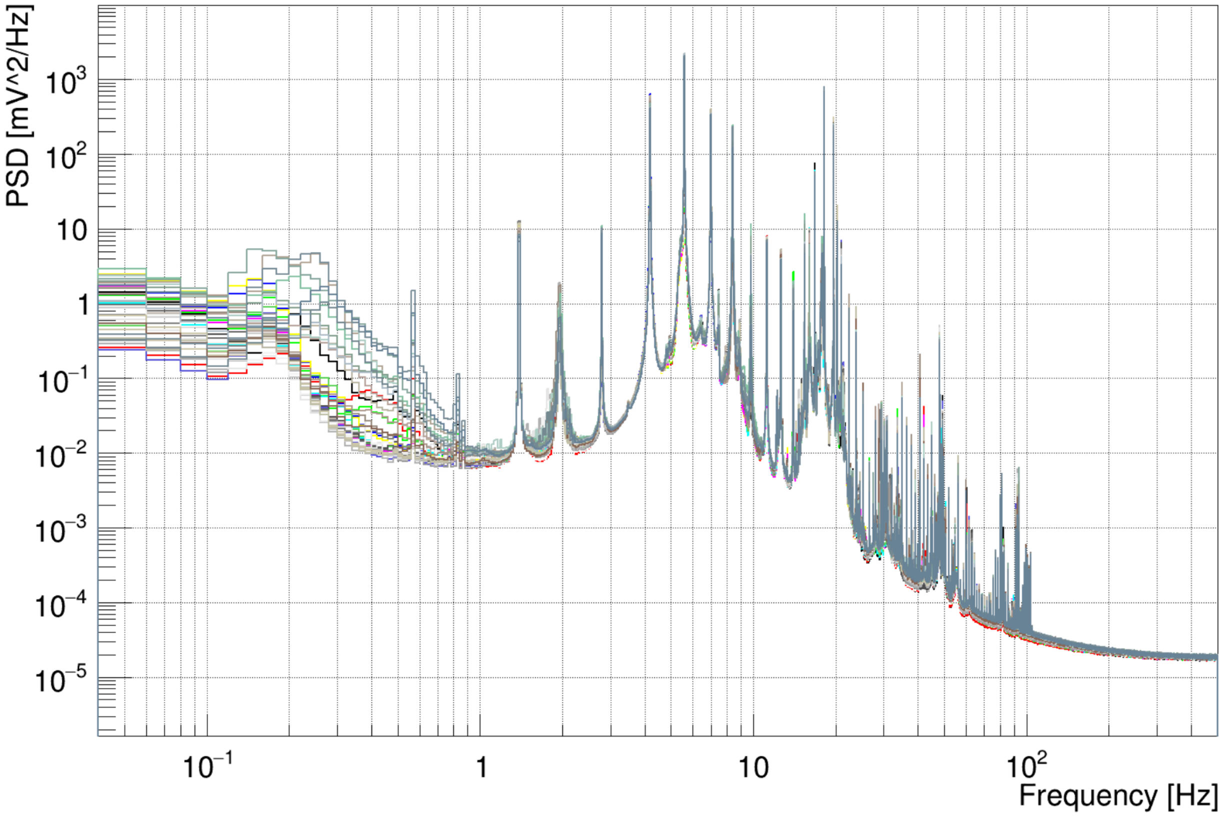

Figure 5.10: Several ANPS of one seismometer channel in CUORE during a dataset. Each line corresponds to a single run of physics data. Over the course of the dataset, the 1.4 Hz noise from the pulse tubes is stable, while the structural resonance at 2 Hz shows some variability across the dataset, and the microseismic activity below 1 Hz is even more variable across the dataset. The seismometer signal at 0.6 Hz and 0.8 Hz changes according to the microseismic activity.

5.3 Characteristics of CUORE Noise

5.3.1 Frequency Cross-correlation with Auxiliary Devices

To characterize the noise in CUORE, we first want to ensure that the auxiliary devices measure vibrations that are also seen in the CUORE detectors. To do this, we calculate the noise covariance matrix and the noise power spectra of the different devices. The noise covariance matrix is defined similarly to how the ANPS is defined in equation 4.7, but with two frequencies and two different devices, i.e.:

$$C_{m,n}[f_i, f_j] = T \langle \tilde{v}_m^*[f_i] \, \tilde{v}_n[f_j] \rangle \tag{5.5}$$

where $\tilde{v}_m[f_i]$ is the discrete Fourier transform of the signal from device m evaluated at frequency f_i, and the factor of T again ensures that the quantity has units V^2/Hz. The cross-correlation matrix of the frequency pair (f_i, f_j) is the magnitude of the noise covariance matrix divided by the square root of the product of the ANPS of each device:

$$\rho_{m,n}[f_i, f_j] = \frac{C_{m,n}[f_i, f_j]}{\sqrt{J_m[f_i]\, J_n[f_j]}} = \frac{|\langle \tilde{v}^*_m[f_i]\, \tilde{v}_n[f_j] \rangle|}{\sqrt{\langle \tilde{v}^*_m[f_i]\, \tilde{v}_m[f_i] \rangle \, \langle \tilde{v}^*_n[f_j]\, \tilde{v}_n[f_j] \rangle}} \tag{5.6}$$

This is a unique matrix for each pair of devices, and each entry must lie between 0 and 1 by the Cauchy-Schwarz inequality. An example of one of these matrices for a microphone and a CUORE detector is shown in Fig. 5.11. The expectation values are calculated using the noise events from a single physics run. The regions of nonzero correlation occur near the pulse tube frequencies starting at 4.2 Hz and near 280 Hz, the same regions of frequency space that are shown in Fig. 5.8. Interestingly, zooming in closely on the low-frequency region of the calorimeter noise reveals a small correlation (∼6%) in the noise at ($f_{\text{calorimeter}} = 0.6$ Hz, $f_{\text{microphone}} = 281.7$ Hz). The same correlation analysis is repeated with an accelerometer instead of the microphone, and a similar trend can be observed (see Fig. 5.12. The cross-correlation of the pulse tube harmonics is stronger in the accelerometer's case, and the non-zero correlation is observed at the 1.4 Hz and 2.8 Hz peaks, unlike in the microphone. This is expected based on the ANPS of the microphone and accelerometer shown above.

5.3.2 Bicoherence

Finally, the fact that there is a slight correlation between the calorimeters at 0.6 Hz and the microphones at ∼280 Hz, as well as the "forest" of peaks in the calorimeter ANPS, many of which are spaced apart by 0.6 Hz, motivates us to analyze the *bispectrum* of the calorimeter as well. The bispectrum measures the phase coupling at two different frequencies in a single signal. Note that this is a unary quantity; it is only a measurement of the calorimeter, not the calorimeter and an auxiliary device.

We define the bispectrum of a signal $v[t]$:

$$\mathcal{B}[f_i, f_j] = \langle \tilde{v}^*[f_i] \tilde{v}^*[f_j] \tilde{v}[f_i + f_j] \rangle \tag{5.7}$$

and we define the bicoherence, b using the following normalization convention:

$$b[f_i, f_j] = \frac{|\mathcal{B}|}{\sqrt{\left\langle |\tilde{v}[f_i]\tilde{v}[f_j]|^2 \right\rangle \left\langle |\tilde{v}[f_i + f_j]|^2 \right\rangle}} \tag{5.8}$$

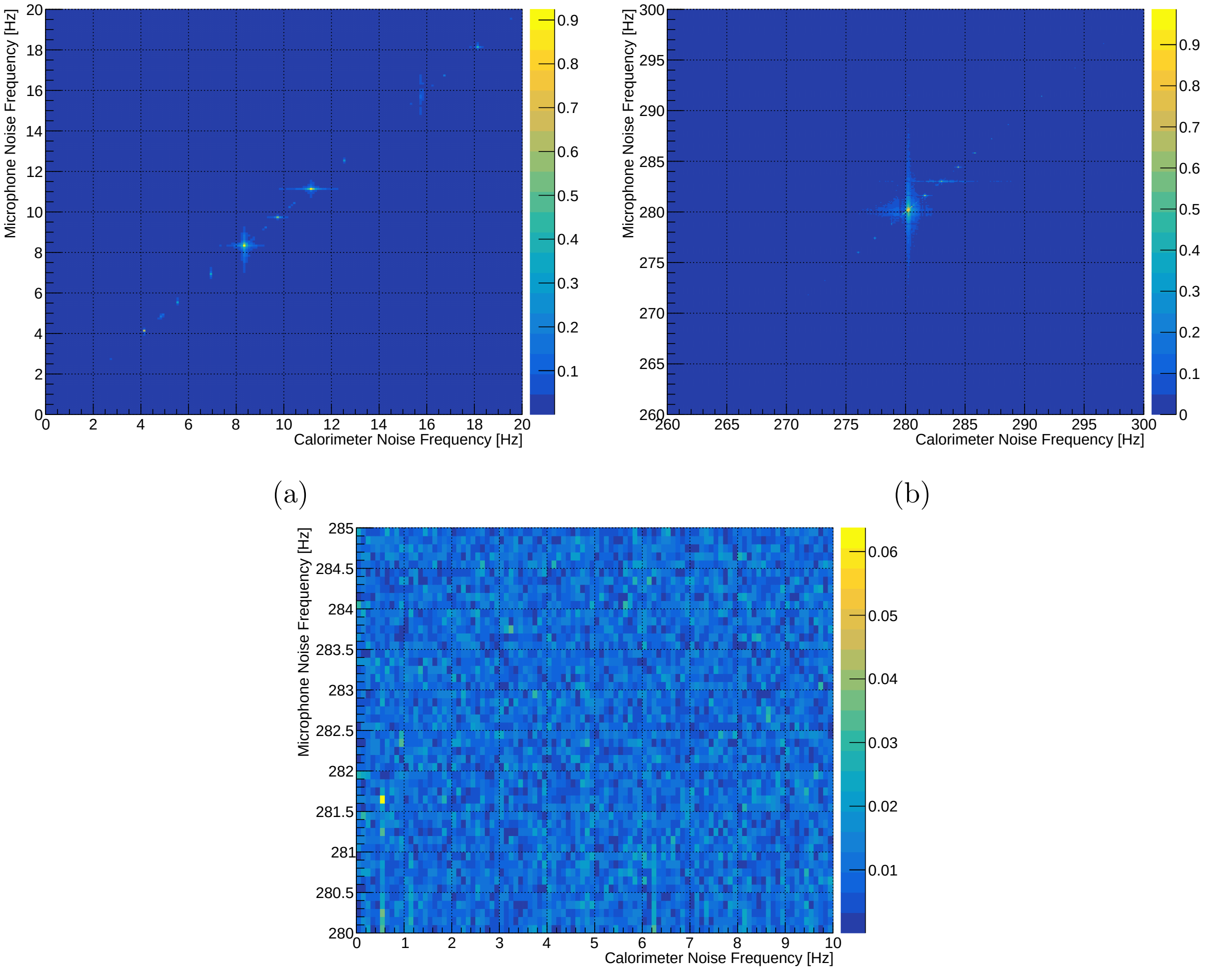

(a)

(b)

(c)

Figure 5.11: Cross-correlation of the noise between a microphone and a CUORE calorimeter taken from a physics run. (a) Correlated noise in the low-frequency region is mostly on-diagonal at the harmonics of 1.4 Hz due to the pulse tubes. (b) The strongest correlation in the noise is at 280 Hz, which is the frequency of the pulse tube compressor line. (c) A small correlation of $\sim 6\%$ is observed at ($f_{\text{calorimeter}} = 0.6$ Hz, $f_{\text{microphone}} = 281.7$ Hz), indicating an interesting relationship between the noise at these frequencies.

The bicoherence is restricted to be between 0 and 1, and it measures the proportion of the signal power at any frequency pair that is phase-coupled to its sum frequency.[1] This is sometimes interpreted to be a measurement of the nonlinearity present in a system, however

---

[1] We could just as easily use the difference frequency $f_1 - f_2$, and the results will still measure the phase coupling of the system. The sum frequency is what is conventionally used when defining the bicoherence.

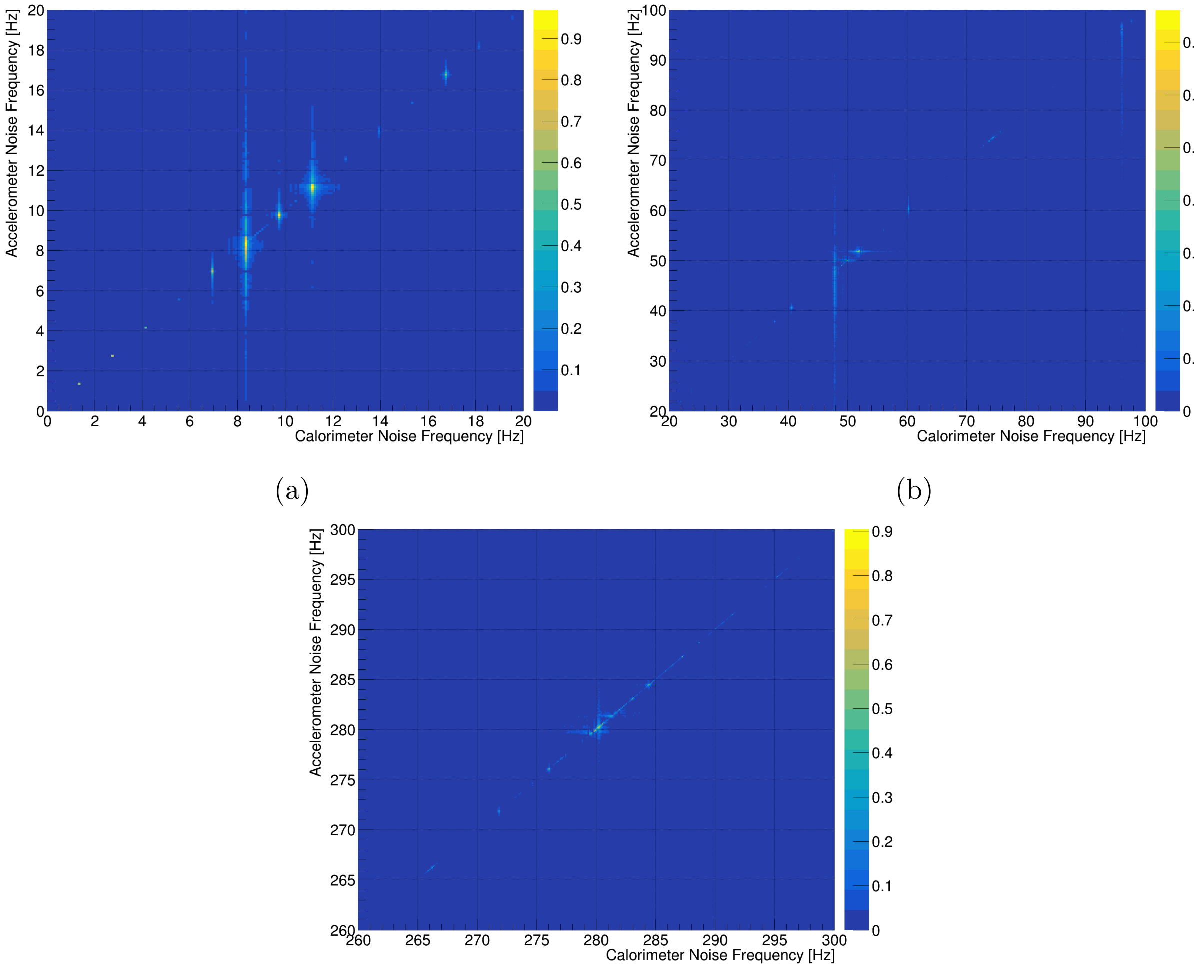

(a)

(b)

(c)

Figure 5.12: Cross-correlation of the noise between an accelerometer and a CUORE calorimeter taken from a physics run. (a) Correlated noise in the low-frequency region is mostly on-diagonal at the harmonics of 1.4 Hz due to the pulse tubes. The accelerometer cross-correlation is non-zero even at the fundamental frequency. (b) In the accelerometers, unlike the microphones, a non-zero cross-correlation is clearly noticeable at frequencies beyond 20 Hz and up to 100 Hz (c) As seen with the microphones, there is a strong correlation between the accelerometers and the calorimeters at 280 Hz.

there are subtleties related to this. Consider a signal of the form:

$$x_L(t) = \sum_i A_i \cos(2\pi f_i t + \phi_i(t)) \tag{5.9}$$

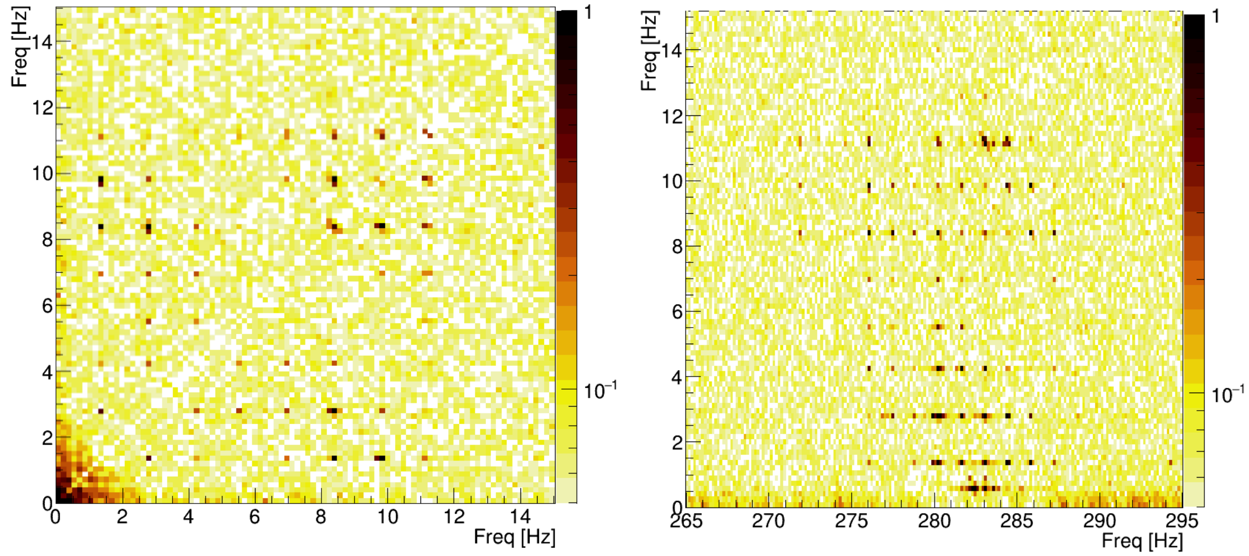

(a) (b)

Figure 5.13: Bicoherence of a CUORE channel. (a) Region of non-zero bicoherence in CUORE data at harmonics of 1.4 Hz induced by the pulse tube phase coupling. (b) Region of non-zero bicoherence in CUORE data indicating mixing of frequencies up to ~12 Hz with frequencies near 280 Hz.

where A_i are all real coefficients. There is no nonlinear mixing in this system. Now consider three frequencies f_i, f_j, f_k where $f_i + f_j = f_k$. If there is no phase coupling between f_i, f_j, and f_k, then $b^2(f_i, f_j)$ approaches 0 after averaging over a sufficiently long period of time. If the phases are coupled such that over a long time $\langle \phi_k - (\phi_i + \phi_j) \rangle \neq 0$, then $b^2(f_i, f_j) > 0$. This phase coupling is expected in systems driven by an external periodic noise source with multiple harmonics. The pulse tubes in CUORE are an example of such an external source, so we expect non-zero bicoherence at harmonics of the fundamental pulse tube frequency (1.4 Hz) even without the presence of a nonlinear mixing term.

Now, we analyze the bicoherence of a system with a quadratic mixing term. Consider a simple signal of containing two frequencies with a nonlinear mixing term. Using $\omega_i = 2\pi f_i$, we have:

$$\begin{aligned} x_{NL}(t) &= A_1 \cos(\omega_1 t + \phi_1) + A_2 \cos(\omega_2 t + \phi_2) + 2A_{12} \cos(\omega_1 t + \phi_1) \cos(\omega_2 t + \phi_2) \\ &= A_1 \cos(2\pi\omega_1 t + \phi_1) + A_2 \cos(\omega_2 t + \phi_2) \\ &\quad + A_{12} \cos((\omega_1 + \omega_2)t + \phi_1 + \phi_2) + A_{12} \cos((\omega_1 - \omega_2)t + \phi_1 - \phi_2) \end{aligned} \tag{5.10}$$

The bicoherence of this signal at (ω_1, ω_2) is:

$$b(f_1, f_2) = \frac{\left\langle A_1 e^{i\phi_1} A_2 e^{i\phi_2} A_{12} e^{-i(\phi_1+\phi_2)} \right\rangle}{A_1 A_2 A_{12}} = 1 \tag{5.11}$$

In this system, all of the power at $f = f_1 + f_2$ is due to the mixing of the signals at f_1 and f_2, so the bicoherence attains the maximum possible value.

Analysis of these two signals shows that while a non-zero bicoherence is a natural consequence of a quadratic nonlinearity in a system, it does not guarantee the existence of such a nonlinearity. Therefore, the bicoherence is helpful in examining the phase-coupled components of a signal, but one should be cautious before attributing such phase coupling to an underlying nonlinearity. We discuss why this quadratic mixing model may still be appropriate for the CUORE system in Section 6.5.1.

The bicoherence of a single CUORE channel is plotted in Fig. 5.13 for regions of frequency space where the bicoherence is non-negligible. We see a large amount of bicoherence between the different pulse tube frequencies, but we expect the harmonics of the pulse tube phases to be coupled already. An interesting region exists around $f_1 \approx 280$ Hz, $f_2 < 10$ Hz, the same region where the forest of peaks is visible in the calorimeter ANPS. There, we see several points with a significant bicoherence term. The distances between these points are integer multiples of 0.6 and 1.4 Hz, both appearing independently in the calorimeter ANPS. Here, no obvious harmonic effect could contribute to the bicoherence, so a nonlinear mixing term may be at play.

Chapter 6

A Denoising Algorithm for Cryogenic Calorimeters

As we have seen, experiments require low threshold detectors with low noise to meet increasingly demanding physical sensitivities. In the case of low-temperature calorimetric detectors, non-negligible amounts of noise can arise from the vibrations induced by the cryogenic facility which dissipate power, thereby adding microphonic and thermal noise to the system [184]. In experiments employing such detectors (e.g. [110, 111, 185–187]), the relevant signal bandwidths range from tens to hundreds of Hz [188]. In CUORE, despite the excellent energy resolutions and low thresholds of the detectors, the intrinsic limit of the energy resolution is far out of reach as extrinsic noise degrades the signal quality, as was shown in the previous chapter. We now seek to build an algorithm that uses the information from the auxiliary devices described in Section 5.2 to remove noise from the CUORE detector and improve their performance, namely their threshold, energy resolution, and timing resolution. We do this by implementing a noise decorrelation algorithm based on a multi-input, single-output model. In Section 6.2, we demonstrate the algorithm on simulated data. In Section 6.4, we apply the algorithm to a cryogenic calorimeter operated as a light detector instrumented with neutron transmutation doped (NTD) germanium [135, 136] as our detector of interest. We use an antenna, three accelerometers, and a separate NTD vibration sensor as auxiliary devices and demonstrate that our algorithm improves both the signal-to-noise ratio (SNR) of the detector and its energy resolution. Finally, in Section 6.5 we expand this algorithm to include nonlinear terms and apply this nonlinear version of the algorithm to a CUORE channel. Using accelerometers and seismometers as auxiliary devices, we again show an improvement in the timing resolution, energy resolution, and threshold of the detector.

6.1 Denoising Algorithm

The starting point of this algorithm draws heavily from the multiple-input, single-output model outlined in [189]. We begin with n timestreams $x_i[t]$ digitized at a sampling frequency of f_s. Here $i = 0, .., n-1$ indexes the different inputs, e.g. the auxiliary devices described in Section 5.2. We also require that the timestream of the detector we seek to denoise, $y[t]$, be sampled at f_s.[1] The goal is to find the array of transfer functions $H_{x_i y}[t]$ which describes the multiple-input, single-output system. In practice, $y[t]$ will contain noise and signals of interest for the experiment. (Henceforth, we will refer to this signal of interest as a *pulse*, while "signal" will be reserved to refer to digital signals such as x_i and y.) Thus, one must subdivide the timestreams $x_i[t]$ and $y[t]$ into events of length $N = Tf_s$, where T, the length of the filter, is a tunable parameter of the algorithm. The subdivision is done by triggering the raw data stream, after which we can divide the data into two subsets. In the ideal case, one subset contains only events which are free of pulses, $\{y^{(n)}[t]\}$, which we will call noise events. The other subset, $\{y^{(p)}[t]\}$, contains all of the pulses present in the data. We take the set of noise events in the detector of interest as well as the corresponding set of events $\{x_i^{(n)}[t]\}$, which occur at the same times as the elements of $\{y^{(n)}[t]\}$. We then take the discrete Fourier transform of all noise events, giving us the frequency domain sets $\{\tilde{y}^{(n)}[f]\}$ and $\{\tilde{x}_i^{(n)}[f]\}$. With these two sets, we construct the following objects:

$$G_{yy}[f] = \frac{2}{T}\left\langle \tilde{y}^{*(n)}[f]\tilde{y}^{(n)}[f]\right\rangle \tag{6.1}$$

$$G_{x_i y}[f] = \frac{2}{T}\left\langle \tilde{x}_i^{*(n)}[f]\tilde{y}^{(n)}[f]\right\rangle \tag{6.2}$$

$$G_{x_i x_j}[f] = \frac{2}{T}\left\langle \tilde{x}_i^{*(n)}[f]\tilde{x}_j^{(n)}[f]\right\rangle \tag{6.3}$$

The expectation values are taken over the sets of noise events from each device that occur at the same time. For all these objects, f is defined within the range $[0, f_s)$ and discretized into N bins of width $\Delta f = \frac{1}{T}$. For convenience, we will drop $[f]$ hereafter, and it will be implied that we are operating in the frequency domain unless otherwise noted. It is worthwhile to describe these objects in terms of familiar quantities. G_{yy} is the ANPS of the output signal, which is a real non-negative $1 \times N$ vector. $G_{x_i y}$ is a complex $n \times N$ matrix containing the cross-spectral densities of each input with the output. $G_{x_i x_j}$ is a complex $n \times n \times N$ 3-dimensional array. For each frequency, it contains a Hermitian matrix of the cross-spectral densities between the input signals. The on-diagonal terms $G_{x_i x_i}$ comprise the ANPS of the input signals. After averaging over all noise events, the transfer functions from the inputs to the detector $H_{x_i y}$ can be calculated:

$$H_{x_i y} = G_{x_i x_j}^{-1} G_{x_j y}. \tag{6.4}$$

$H_{x_i y}$ is the same size as $G_{x_i y}$ and also complex. The denoised signal, $\tilde{y}'$, is the difference between the original signal and the sum of the auxiliary signals filtered with the transfer function:

$$\tilde{y}' = \tilde{y} - H_{x_i y}^T \tilde{x}_i. \tag{6.5}$$

---

[1] Technically, the common sampling frequency requirement is a simplification. Because the transfer functions are calculated in the frequency domain, if the sampling rates differ, they can still be calculated by truncating the signal with the higher sampling rate in the frequency domain or by downsampling it in the time domain.

This process can then be applied to pulse events to predict the noise that is correlated with the input devices during the event. Assuming that the noise in the detector of interest has components which are correlated with the auxiliaries, the result will be a denoised version of the original set of pulse events. This assumption is usually true given the careful choice of auxiliary devices used in the experiment. However, in cases where the matrix $G_{x_i x_j}$ is dominated by uncorrelated noise in the auxiliary devices, this denoising technique will not work.

One can verify the success of the algorithm by comparing the noise power spectra of the output signal and the denoised signal, G_{yy} and $G_{y'y'}$. One can then proceed to analyze the denoised data, leading to more precise measurements. It is important to note that this algorithm is not a replacement for other filtering techniques such as optimum filtration. Rather, this algorithm can be used first to improve the SNR, resulting in new coefficients for the optimum filter which can then be applied to the data. This procedure and applying many common filters, including the optimum filter, are linear operations, so the two can be combined into a single step [180].

6.2 Simulation of Calorimeter Data

To measure the efficacy of the multivariate algorithm under known conditions, we simulate a multi-input, single-output system to generate mock data of a cryogenic calorimeter. This mock signal of interest contains both noise and pulses. The simulated auxiliary timestreams represent three accelerometers measuring vibrations and one antenna measuring electromagnetic interference. We first create two signals representing the original sources of electrical and vibrational noise. The electrical source signal is a sum of many sinusoids at harmonics of 60 Hz, each with a randomized phase.[2] The vibrational noise is modeled as the sum of three independent sources. The first source signal is a series of sinusoids at harmonics of 1.4 Hz representing the noise from pulse tubes. The second source signal is another sinusoid at 100 Hz, simulating a pump or fan with a particular operating frequency. The third source signal is a randomly distributed set of infrequent impulses passed through band-pass filters at three low-frequency values. The impulses represent sudden disturbances in the system, e.g. earthquakes, and the narrow-band responses represent mechanical resonances in the system that are excited by these disturbances. These impulses cause the noise to be non-stationary, which better represents the dynamic nature of noise in actual experimental environments. This provides a more authentic test of the algorithm's performance, as real-world conditions often involve fluctuating and unpredictable noise sources that must be accurately identified and mitigated.

To create our mock antenna signal, $x_0[t]$, we add $1/f$ noise and Gaussian white noise to the electrical source signal. The addition of white noise simulates thermal noise in the antenna, and the $1/f$ noise represents the flicker noise often present in electronic amplifiers.

---

[2] Here we are simulating electrical noise in data taken in North America, where the standard AC frequency is 60 Hz. Later, in Section 6.5, we will use data taken in Europe, where the standard AC frequency is 50 Hz.

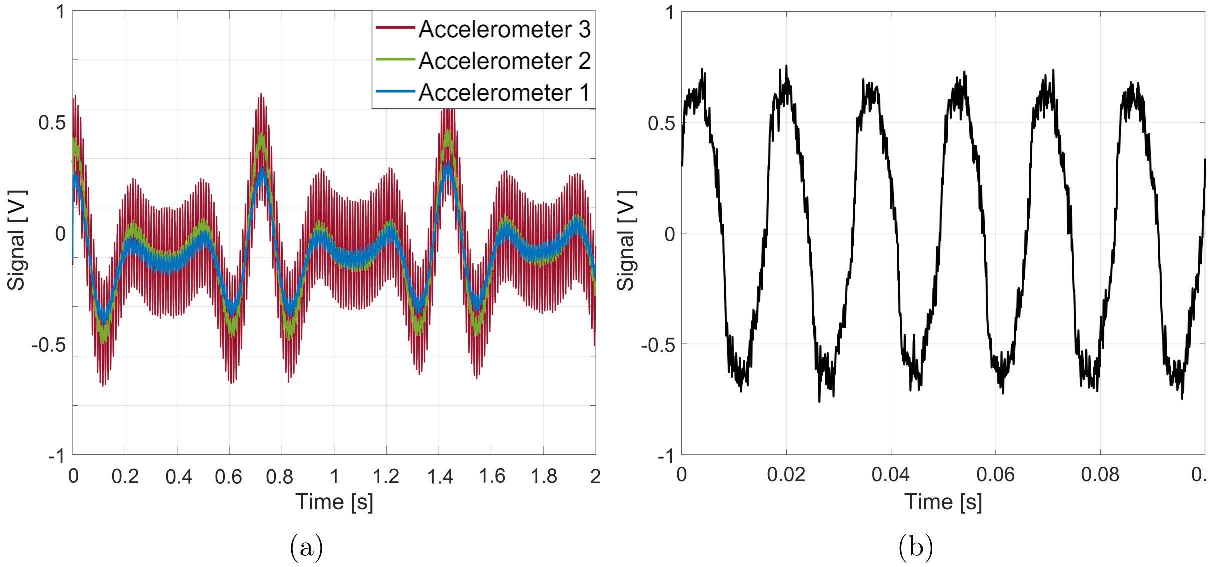

(a) (b)

Figure 6.1: [Input auxiliary signals for the denoising simulation. (a) Simulated accelerometer signals. These signals simulate vibrations in the laboratory from sources such as pulse tubes, fans, and transient events such as earthquakes. (b) Simulated antenna signal. The signal simulates electromagnetic pickup in the laboratory with a large 60 Hz component.

We generate three accelerometer timestreams by individually passing the vibration source signal through unique band-pass filters and adding $1/f$ noise and white noise to each output. The band-pass filters represent the response function of each accelerometer, which we slightly vary for each signal. The resulting signals $x_1[t], x_2[t], x_3[t]$, correspond to accelerometer signals in the $\hat{x}$, $\hat{y}$, and $\hat{z}$ directions. An example of each input signal and their power spectra are shown in figures 6.1 and 6.2.

To create the mock detector signal, we again generate random white noise and $1/f$ noise. These terms are similar to some of the noise terms used in [188]. We then take the vibrational source signal and pass it through a unique band-pass filter. This models the detector's response to vibrations, which should be different from that of the auxiliaries. We add the result and the electrical source signal to the detector signal. Finally, we take a template pulse shape and generate a train of pulses which are randomly distributed in time with the caveat that the minimum time between pulses is longer than the length of the template to prevent pileup events. The pulse amplitude is randomly chosen to be 1 or 2 units with equal probability. The pulse amplitudes represent physical events with energies E_0 and $E_1 = 2E_0$ assuming linearity of the detector response. After generating the set of valid noise events and pulse events, we apply the noise decorrelation algorithm to the detector signal. Using the simulated antenna and accelerometer signals as auxiliary inputs, we construct the transfer functions $H_{x_i y}$ following equations 6.2, 6.3, and 6.4, where the expectation values

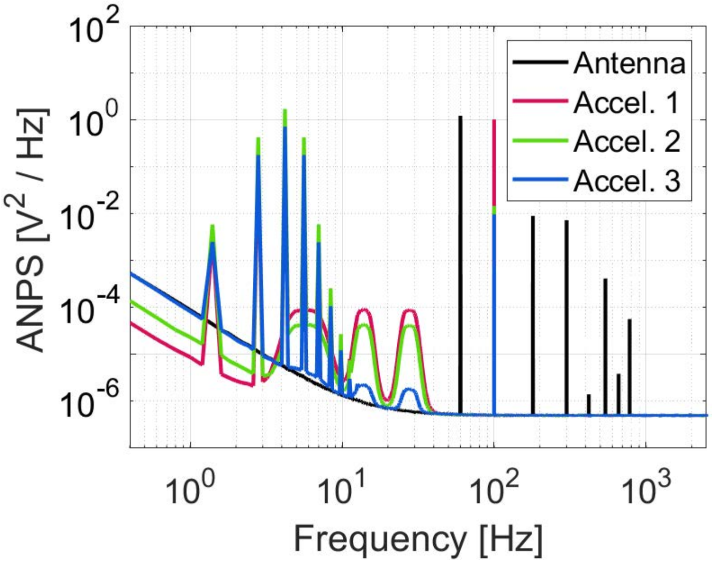

Figure 6.2: Average noise power spectra of simulated auxiliary input signals. The simulated antenna signal (black) measures electrical noise and contains harmonics of 60 Hz. The simulated accelerometer signals (red, green, and blue) measure vibrational noise. The vibrational noise is composed of harmonics of 1.4 Hz which represents pulse tube noise, and broader peaks below 100 Hz which represent the excitation of mechanical resonances in the system due to sudden disturbances.

are taken over 5-second long noise windows. The results of the noise decorrelation on the output signal are shown in Fig. 6.3.

6.3 Optimal filtration of Denoised Data

We now seek to show that the amplitude reconstruction and resolution of the simulated data are improved with the addition of the denoising algorithm. We can empirically measure the OF noise resolution by choosing an amplitude value from the same random time point within each noise window. The standard deviation of these values is a good approximation for the OF noise resolution of the detector.

To analyze the results of the noise decorrelation algorithm, we run two analyses in parallel. First, we construct the OF using the original simulated signal and use it to reconstruct the pulse amplitudes and OF noise resolution. Separately, we construct a new OF using the denoised data and apply it to the denoised data. The results before and after denoising are shown in Fig. 6.4 and Table 6.1. The amplitude resolution is improved at the noise peak and both pulse amplitudes. In this particular iteration of the simulation, the amplitude resolution consistently improves by 22–23%. Both before and after denoising, the amplitude resolutions

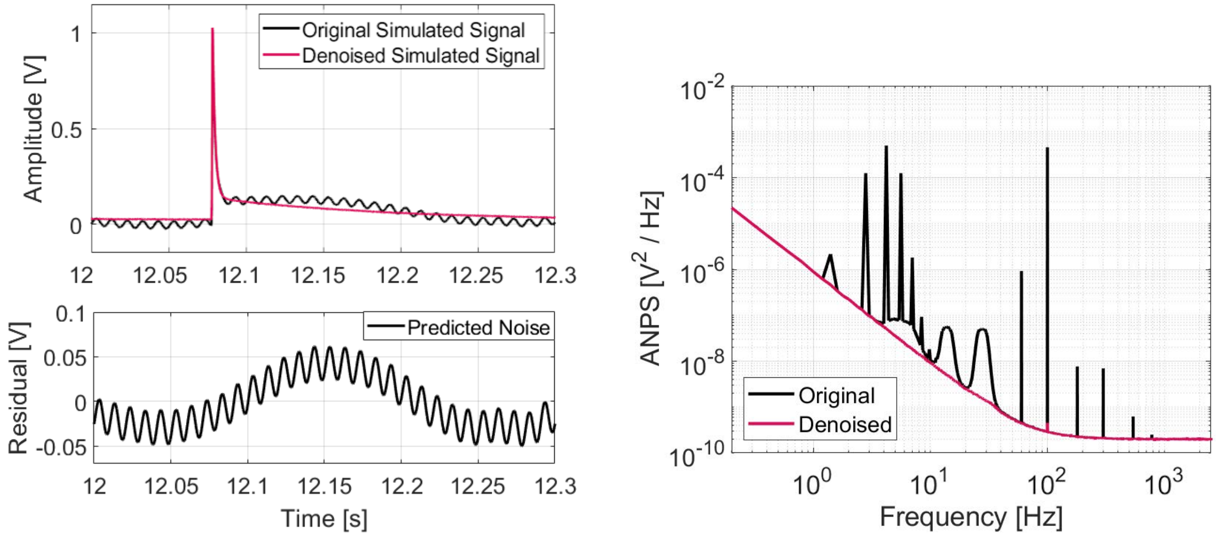

(a)

(b)

Figure 6.3: Effects of the denoising algorithm on the simulated data. (a) Simulated pulse event in the noisy detector signal (black) and denoised signal (red), and the noise predicted by the algorithm (below). (b) ANPS of the simulated detector signal before denoising (black) and after denoising (red). The noise that is correlated with the auxiliary devices is removed.

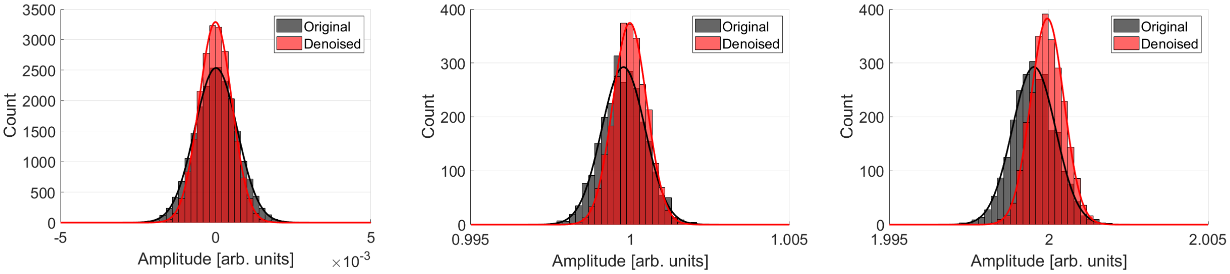

(a)

(b)

(c)

Figure 6.4: Optimally filtered amplitudes of simulated pulses without denoising (black) and with denoising (red) applied to the data. Each peak shows an improvement in resolution when the denoising algorithm is applied. Note that the denoised data better reconstruct the original signal amplitudes of 1 and 2 units, reducing the bias in the system. (a) OF amplitude distribution of noise events. (b) OF amplitude distribution of pulse events with amplitude 1. (c) OF amplitude distribution of pulse events with amplitude 2.

across all three values are compatible with each other. This is also expected because no energy-dependent resolution scaling is built into the simulation.

Table 6.1: Amplitude resolution and reconstruction error of simulated noise events and pulse events after denoising. The reconstruction error is the difference between the measured pulse amplitude and the simulated pulse amplitude. The denoising algorithm reduces the magnitude of the reconstruction error and improves the resolution at all amplitudes. The resolutions at all amplitudes are compatible before and after denoising. This is expected because no energy-dependent resolution scaling is built into the simulation.

Simulated Pulse Amplitude [V]	0 (Noise)	1	2
Original Reconstruction Error [μV]	9 ± 5	-210 ± 10	-460 ± 10
Denoised Reconstruction Error [μV]	-7 ± 4	-8 ± 10	-40 ± 10
Original Resolution [μV]	684 ± 3	680 ± 10	680 ± 10
Denoised Resolution [μV]	527 ± 3	530 ± 10	520 ± 10

One important ramification of this algorithm is that it can change the mean value of the OF amplitudes. This behavior is not typically expected of linear filters, and it is worth exploring in some detail. When the denoising is applied to the data, the noise power of the detector at a particular frequency is reduced *only* if the auxiliary devices are correlated at that frequency. In addition, the assumptions typically made about the noise conditions of the system when applying the OF are not actually valid in this simulation. These assumptions include a linear, time-invariant system and a signal which is uncorrelated with the noise. In this analysis, we construct new average pulses after denoising, resulting in cleaner signal templates. The fact that the signal template is not the same when constructing the different OFs after denoising demonstrates that the noise in the system is in fact correlated with the AP templates built during the simulation. A simple way to view this mathematically is as follows:

A waveform composed of a signal of amplitude A and noise can be written as

$$y[f] = As[f] + n[f]$$

where s is the signal template with amplitude 1 and n is the noise. Let $J[f] = T\langle n^*[f]n[f]\rangle$ be the noise power spectrum of the system. The OF is defined as:

$$OF[f] \equiv \left[\sum_f \frac{|s[f]|^2}{J[f]}\right]^{-1} \frac{s^*[f]}{J[f]} \tag{6.6}$$

and optimally filtered waveform is:

$$\left[\sum_f \frac{|s[f]|^2}{J[f]}\right]^{-1} \frac{s^*[f]\,(As[f] + n[f])}{J[f]} \tag{6.7}$$

The amplitude estimator is the sum of all filtered frequency components:

$$\hat{a} = \left[\sum_f \frac{|s[f]|^2}{J[f]}\right]^{-1} \sum_f \left[\frac{A|s[f]|^2}{J[f]} + \frac{s^*[f]n[f]}{J[f]}\right] \tag{6.8}$$

Suppose the signal and noise are uncorrelated and the noise is stationary. In that case, the numerator in the last term, $\sum_f [s^*[f]n[f]/J[f]]$, has an expectation value of zero and $\hat{a} = A$, i.e. the estimator is unbiased. However, if the signal template s is correlated with the noise, i.e. $\langle s^*[f]n[f]\rangle \neq 0$ for some frequencies, then the last term in the above is not necessarily 0 for all waveforms, and the estimator can become biased. In practice, since the signal template has less noise after denoising.

In addition, $J[f]$ is not a good representation of the noise when non-stationary noise arises in the detector. Again, this causes the final term in the above equation 6.8 to acquire a non-zero expectation value, thus biasing the detector. In this simulation, the denoising algorithm reduces the non-stationary noise; thus the OF amplitudes of the pulses change. In practice, this serves to unbias the data, and the reconstruction error of the amplitudes is significantly reduced after the denoising is applied.

6.4 Denoising Experimental Data

6.4.1 Experimental Setup

We performed the measurements to demonstrate the denoising algorithm in an Oxford Instruments Triton 400 DR cooled with a cryogen-free pulse tube cooler at an operating temperature of ∼10 mK. Inside the refrigerator, we mounted a low-temperature light detector which serves as our detector of interest in this demonstration. The light detector is composed of a 45×45×0.5 mm^3 silicon wafer that absorbs the light increasing its temperature. This temperature change is measured by a 3×3×1 mm^3 NTD germanium thermistor which is epoxied to the wafer. The silicon wafer was rigidly connected to a square copper frame that was connected to the experimental stage of the DR. The front-end electronics used in this setup are identical to those used in CUORE [190]. We ran a fiber-optic cable to near this detector with the axis of the cable approximately normal to the substrate surface. We connected the fiber to a commercially available LED with a wavelength of (600±12) nanometers to generate a train of LED light pulses with durations and separation times that can be easily controlled by a signal generator with pulse width modulation (PWM). This allowed us to create periodic light signals which we could use to characterize our detectors.

Alongside the light detector, we also mounted another NTD-Ge thermistor which was adhered to a small shard of a silicon wafer. This “NTD vibration sensor” was not rigidly connected to the copper frame, but instead left free to vibrate during the data-taking. Because the NTD vibration sensor was adhered to a relatively small wafer, we expected that the rate of pulses it observed would be negligible compared to that of the light detector. We could therefore use it as another auxiliary device while making cuts to the data in the rare case of a pulse event on the auxiliary device. This was a useful addition to the setup because we expected that the vibrations it measured would be closely correlated to those measured in the light detector due to their close proximity. We also expected that the electrical noise on each of the NTD signals will be similar since they are being read out in identical ways.

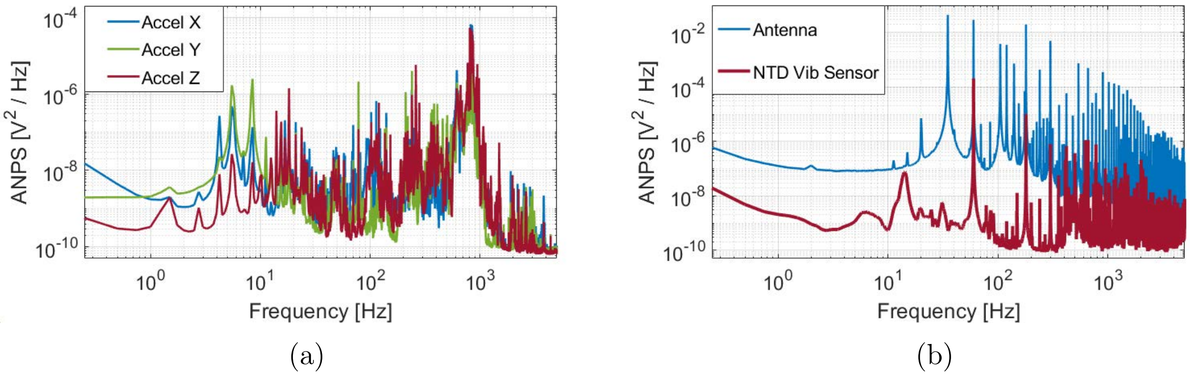

Figure 6.5: ANPS of auxiliary devices used to denoise data from the Oxford DR. (a) ANPS of accelerometers. Pulse tube noise at harmonics of 1.4 Hz and many other high-frequency peaks are visible. The noise power from ∼10 Hz to 100 Hz is reduced in the $\hat{z}$-oriented accelerometer because the experimental setup is equipped with dampeners that reduce vertical vibrations. (b) ANPS of antenna and NTD vibration sensor. The antenna measured a signal dominated by components at 35 Hz, 60 Hz, and their harmonics. The NTD vibration sensor signal is dominated by the odd harmonics of 60 Hz, though there is also a source of broadband noise between 10 and 20 Hz.

To measure the electrical noise present in our system, we attached an antenna to our DAQ and placed it inside the Faraday cage which houses the experiment. The antenna we used was a hand-made Helmholtz coil with $\sim$ 50 turns and a radius of $\sim$ 30 cm, and we amplified the antenna signal by a factor of 5,000. To measure the vibrational noise, we used three PCB 393B31 accelerometers [182] mounted triaxially to the 300 K plate of the DR. The accelerometer signals were amplified with a PCB 482C15 signal conditioning pre-amplifier [191] and read out to the same DAQ system as the antenna, NTD vibration sensor, and light detector. The ANPS of the auxiliary devices are shown in Fig. 6.5.

As a starting ground for testing the algorithm, we collected data while a ground loop persisted in our readout chain. This gave rise to noise peaks at 60 Hz and subsequent harmonics from the AC power line with sufficient power to dominate the detector signal. In addition to the electrical noise, we observe noise peaks at 1.4 Hz and subsequent harmonics due to vibrations from the pulse tube which vibrate the internal cryogenic structure, including the detector and the cold electronics. To construct an average noise template, we first took a long run of data without injecting any pulses to ensure a sufficient number of noise events

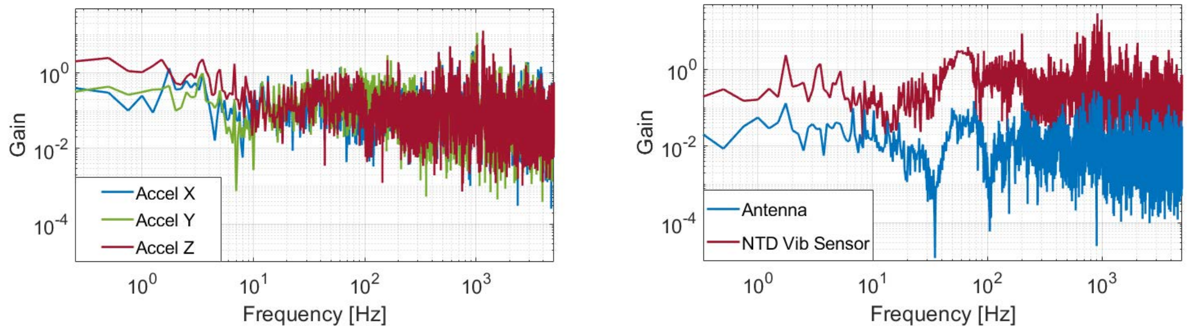

(a) (b)

Figure 6.6: Transfer functions from auxiliary devices to light detector in the Oxford DR. (a) Accelerometer transfer functions. The peaks at the harmonics of 1.4 Hz, particularly in the $\hat{z}$-oriented accelerometer, indicate a correlation between the pulse tube noise observed in the accelerometers and the microphonic noise in the light detector. (b) Antenna and NTD vibration sensor transfer functions. Note the wide-band region of high gain near 60 Hz for both of these devices. The large troughs at 35 Hz and 105 Hz indicate that the antenna signal is not correlated with the light detector at these frequencies.

free of pulses. The data from these runs were used to construct the transfer functions, $H_{x_i y}$, which are shown in Fig. 6.6. After, we acquired two runs of data, each with a set of LED pulses at a fixed frequency (1.7 Hz). The first set of LED pulses had a duration of 100 μs, and the second set had a duration of 200 μs and thus twice the energy of the first set. We split the PWM signal and acquired it with our DAQ to efficiently trigger on the light pulses with few backgrounds. We denoised these LED runs by applying the previously calculated transfer functions to the auxiliary data from the LED runs and subtracting the result from the light detector signal.

6.4.2 Analysis Techniques

Since we could trigger the LED pulses, we could still analyze our signal despite the low SNR in the system. We used the same analysis procedure described in Section 6.2, but this time we conducted three separate analyses. We again analyze the original data and denoised data, but we also analyze the data after applying a notch filter centered at 60 Hz to the signal. Since the signal is dominated by 60 Hz noise, one might expect that notch filtering the data will lead to an improvement comparable to that of the noise decorrelation algorithm. We selected 0.5-second windows between LED pulses and making RMS-based cuts to veto windows with spurious events (e.g. caused by cosmic rays) either on the detector signal or NTD vibration sensor. We then construct the average noise power spectrum for each of the three signals. We use the LED triggers to build an average pulse which is also 0.5 seconds long. Together, the noise events and average pulse template allow us to denoise the data and build the OF. We separately construct the OF using the original data, the band-pass filtered data, and the denoised data.

To align the gains of the three filters, we perform a pseudo-calibration assuming that the second set of LED pulses has twice the energy of the first one. We align the LED pulse amplitudes to arbitrary energy values of E_0 and $2E_0$. We construct histograms of the amplitudes obtained by each analysis of each LED event, then fit the amplitudes to a normal distribution using an unbinned fit. We then use a second-order polynomial fixed at the origin to align the measured amplitude values to a common scale. Following the work of [192], one can model the energy resolution of the detector as $\sigma(E) = \sqrt{\sigma_0^2 + f(E)\sigma_E^2}$, where σ_0 is the baseline resolution of the detector, σ_E is a scale factor that arises due to the stochastic nature of the LED signals, and $f(E)$ is a monotonically increasing function. We expect this behavior of the detector response because the magnitude of the Poisson fluctuations in the number of photons captured by the light detector will increase as the amplitude of the LED pulse increases. With this model, we expect that the denoising algorithm does not affect $f(E)$ or σ_E, but does lower the value of σ_0. We thus expect to see the greatest relative improvement in energy resolution at the noise peak where the energy-dependent effects are small. We define the ratio of resolutions $R(E) = \hat{\sigma}_{denoised}(E)/\hat{\sigma}_{original}(E)$ and calculate it at each energy for the three signal processing techniques.

6.4.3 Results of Denoising

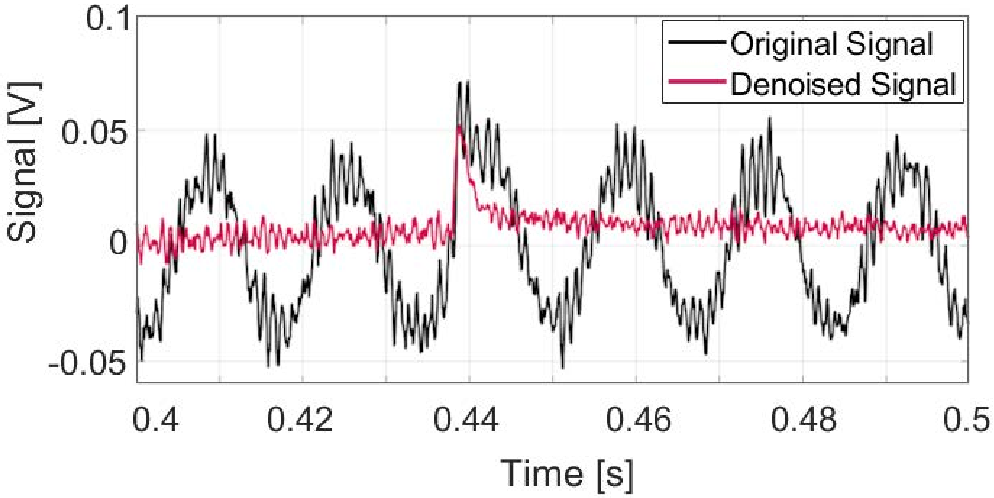

Figure 6.7: Light detector signal containing an LED pulse before and after denoising. Before denoising, the pulse is almost completely washed out by the noise which is dominated by a 60 Hz component. After denoising, the pulse is visible by eye.

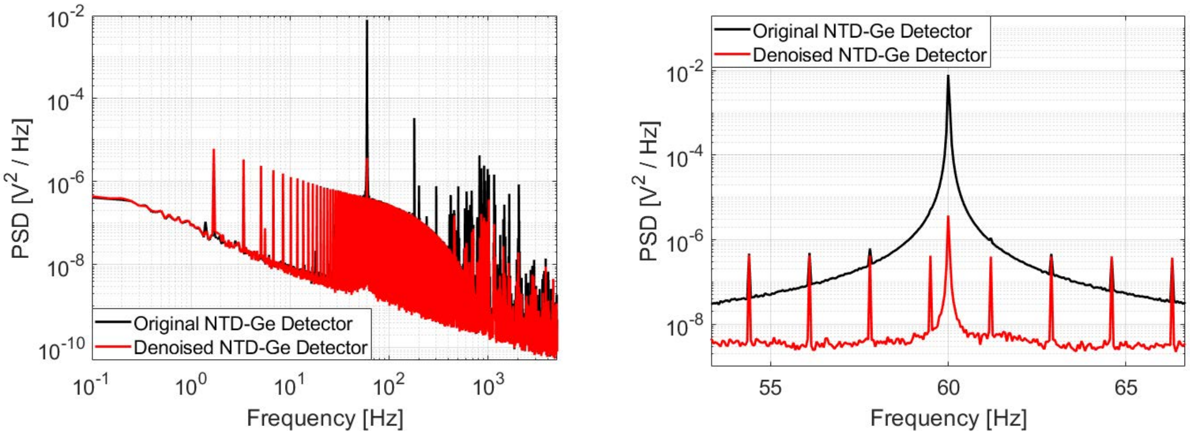

(a) (b)

Figure 6.8: (a) Periodogram of the light detector signal before and after denoising. The events used to build these power spectra include the data from the train of LED pulses which produces the evenly spaced peaks separated by 1.7 Hz. Noise peaks across the frequency spectrum, including the 1.4 Hz vibrational peak and the 60 Hz electrical peak, are eliminated after denoising. (b) The same periodogram zoomed to the 60 Hz noise peak. The peak is reduced by ∼33 dB, and the FWHM of the peak is reduced by approximately 10%. The evenly spaced peaks due to the LED pulse train become more apparent after denoising. This structure is largely unaffected by the denoising, indicating an improved signal-to-noise ratio near 60 Hz.

After denoising, there was a visible improvement in the signal-to-noise ratio of the detector as seen in Fig. 6.7. We also observed a reduction in both vibrational and electrical noise in the light detector periodogram (see Fig. 6.8). The vibrational peak at 1.4 Hz was eliminated, and many of the electrical peaks ranging from 60 Hz to 5 kHz were significantly reduced. The 60 Hz noise peak was reduced by 33 dB, and the full width at half-max (FWHM) was reduced by approximately 10%. This peak was extremely wide due to small modulations of the ground loop frequency, but these modulations were measured by the antenna and the NTD vibration sensor so we could successfully remove them from the data. This is reflected in the wide-band region of high gain near 60 Hz in the transfer functions for both of these devices as was shown in Fig. 6.6. Outside of this region, the main effect of the denoising is a reduction in noise peaks. The algorithm does not have an apparent effect on the broadband detector noise. Before denoising, high frequency noise with a large 60 Hz component contaminates the average pulse. Applying the notch filter induces a ringing in all pulses which coherently add together when constructing the average pulse. When the denoising algorithm is applied, the noise in the average pulse is reduced without inducing any ringing (see Fig. 6.9).

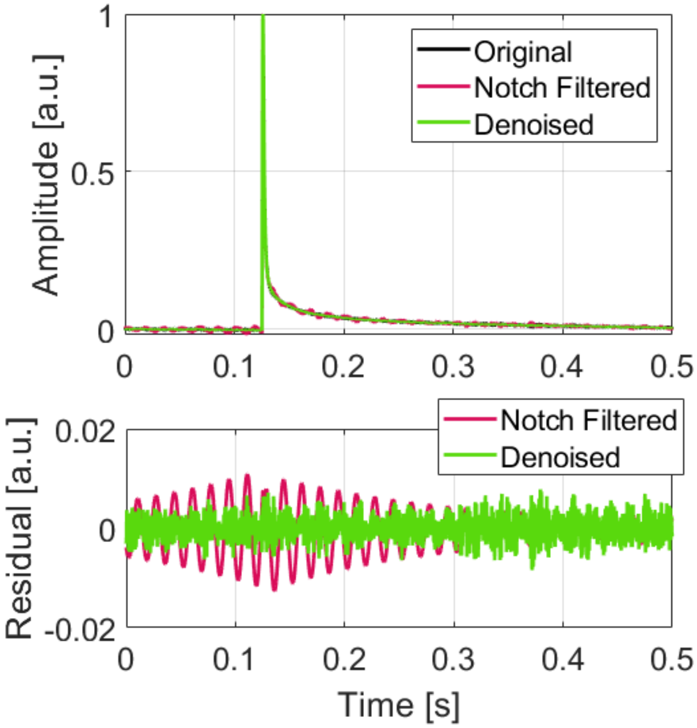

Figure 6.9: Average pulses from the light detector constructed with different signal processing techniques. The notch filter produces a ringing effect which can be seen in the plot of the residuals. The denoising procedure removes 60 Hz noise as well as other high-frequency noise from the average pulse, producing a smoother and more accurate pulse template for later use in the OF.

Next, we examine the effect of the denoising algorithm on the resolution of noise and LED pulse events. As expected, the resolution worsens monotonically with energy. At all amplitudes, the denoising improves the measured detector resolution (see Table 6.2). We see the greatest relative improvement at lower pulse amplitudes, which was also expected. It should be noted that the improvement in amplitude resolution at the first LED peak ($E = E_0$) is compatible with the improvement at the noise peak, though it has a greater uncertainty.

We also observe that the denoising algorithm performs as well as or better than the notch filtering approach (see Fig. 6.10). This effect is especially visible at the noise peak, where the improvement in the detector resolution from the denoising (4.1%) is 1.7 times that of from the notch filtering (2.4%). At the highest-amplitude peak, we do not expect the denoising to have as much of an effect on the resolution, but we still observe an improvement. At this amplitude, we see that the denoising and notch filtering have similar effects on the data These results are summarized in Table 6.3.

Table 6.2: Best-fit estimators of energy resolution with and without denoising applied. Here $\hat{\sigma}_0$ is the best estimator of the energy resolution of the detector noise events. All energies are scaled to that of the first LED pulse amplitude. All energy resolutions are calculated from the optimally filtered data.

Signal Type	$\hat{\sigma}_0(E=0)$	$\hat{\sigma}(E=E_0)$	$\hat{\sigma}(E=2E_0)$
Original	$(2.89 \pm .01) \times 10^{-2}$	$(3.15 \pm .04) \times 10^{-2}$	$(3.16 \pm .04) \times 10^{-2}$
Notch Filtered	$(2.82 \pm .01) \times 10^{-2}$	$(3.13 \pm .03) \times 10^{-2}$	$(3.10 \pm .03) \times 10^{-2}$
Denoised	$(2.77 \pm .01) \times 10^{-2}$	$(3.00 \pm .03) \times 10^{-2}$	$(3.10 \pm .03) \times 10^{-2}$

Table 6.3: Ratio of denoised and notch filtered energy resolutions to original energy resolution. The ratio of resolutions is defined as $R(E) = \hat{\sigma}_{denoised}(E)/\hat{\sigma}_{original}(E)$. The denoising algorithm outperforms the notch filter given its capability to remove noise across a large range of the frequency domain. Note that the main effect of the denoising is in the reduction in baseline resolution R_0. This is expected given that the resolution is dependent on the LED amplitude due to the Poisson fluctuation of detected photons.

Signal Type	$\hat{R}_0$	$\hat{R}(E=E_0)$	$\hat{R}(E=2E_0)$
Notch Filtered	0.976 ± 0.003	0.99 ± 0.01	0.98 ± 0.01
Denoised	0.959 ± 0.003	0.95 ± 0.01	0.98 ± 0.01

It is important to note that the improvement in resolution is $\leq 5\%$ in all cases. This is also expected since the OF is by definition the filter that results in the highest SNR given a particular choice of noise and pulse template. The improvements from the denoising are thus limited by the amount of noise that can be removed from the noise and signal templates, which can in turn be limited by the "SNR" of the auxiliary devices, i.e. the ratio of the amount of noise correlated with the detector of interest to the amount of uncorrelated noise. Having verified the efficacy of the denoising technique in a controlled experiment, we now turn our attention to decorrelating the noise from CUORE data.

6.5 Denoising CUORE data

To validate the efficacy of the denoising algorithm on CUORE data, we selected six of the auxiliary devices outlined in section 5.2. We use these devices as the six input signals when we apply the noise decorrelation algorithm. Here, we will also use a nonlinear version of the algorithm for thermal detectors. Using the data from one CUORE detector channel during a physics run, we demonstrate both the linear and nonlinear versions of the denoising algo-

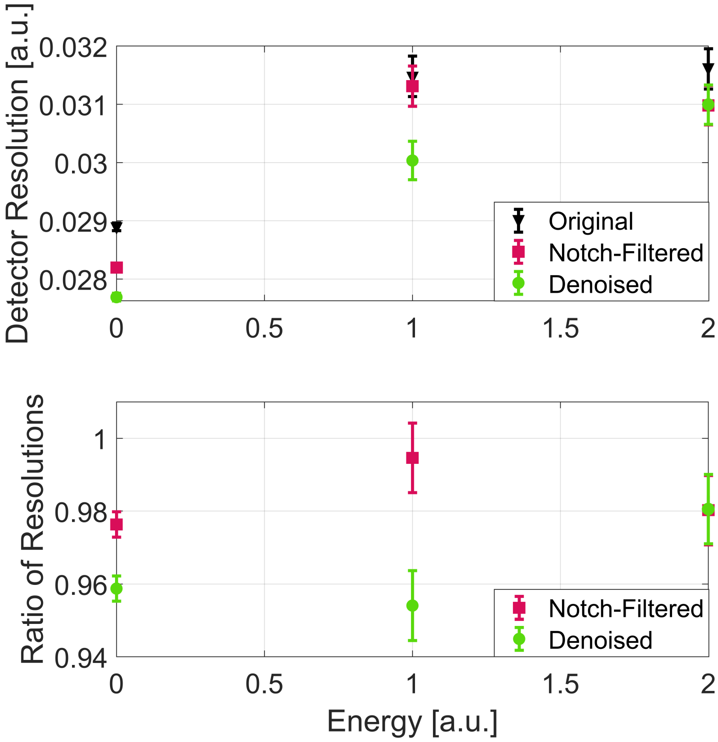

Figure 6.10: Energy resolution of noise events and LED events before and after denoising or notch filtering is applied. The fact that the energy resolution increases with energy is expected due to the Poisson fluctuation of detected photons. At all energies, the denoising algorithm improves the resolution of the detector. At all but the highest energies, the algorithm outperforms the 60 Hz notch filtering of the data.

rithms. We then analyze the noise events and the heater events in the detector to evaluate the algorithm's performance. The particular channel we choose has two pulser amplitudes during this run, the greater of which is used for thermal gain stabilization. The two amplitudes are comparable to physical pulses with energy depositions of 960 keV and 3500 keV. These energy values were determined during the online CUORE data analysis. We will calibrate the detector by using the two pulser event energies and the noise peak, then use the resulting calibration curve to estimate the improvement in energy resolution in units of energy.

6.5.1 The Nonlinear Algorithm

It is expected that quadratic mixing of acceleration signals should occur in the CUORE thermal detectors. This is motivated both by the non-zero bicoherence observed in CUORE (see section 5.3.2) and also by first principles. A simplified explanation for this is that any power incident on the detector due to vibrations heats a thermal detector regardless of the directionality of the vibration. This power is non-trivially related to the square of the waveform describing this vibration. Consider an object of mass m vibrating in 1-dimension whose acceleration is described by a discrete signal $a[t]$. The power supplied by this oscillation is equal to $p[t] = m\, a[t]\, v[t] = m\, a[t] \sum_{t'<t} a[t']dt'$. (We will drop the m so that p has units of $t(a[t])^2$). Taking the Fourier transform of the power, we have:

$$\tilde{p}[f] = \tilde{a}[f] \circledast \frac{\tilde{a}[f]}{f} = i \sum_{f' \neq f} \frac{\tilde{a}[f']\tilde{a}[\Delta f']}{\Delta f'} \tag{6.9}$$

where $\Delta f' = f - f'$ and $\circledast$ represents the convolution. One can replace $(\Delta f')$ with $(2\pi f_s + \Delta f')$ if $f' > f$ when working on the non-negative frequency domain. Now consider the square of the acceleration signal $(a[t])^2$. Its Fourier transform is:

$$\tilde{a}^2[f] = \sum_{f' \neq f} \tilde{a}[f']\tilde{a}[\Delta f'] \tag{6.10}$$

The transfer function from this signal to the power signal is:

$$H_{a^2p} = \frac{\left\langle \left(\tilde{a}^2[f]\right)^* \tilde{p}[f] \right\rangle}{\left\langle \left(\tilde{a}^2[f]\right)^* \left(\tilde{a}^2[f]\right) \right\rangle} \tag{6.11}$$

Plugging in, we see the transfer function is averaged over all mixing frequencies, the $\tilde{p}[f]$ term of which is averaged in a frequency-weighted way:

$$H_{a^2p} = \frac{i \left\langle \left(\sum_{f'} \tilde{a}^*[f']\tilde{a}^*[\Delta f'] \right) \left(\sum_{f'} \frac{\tilde{a}[f']\tilde{a}[\Delta f']}{\Delta f'} \right) \right\rangle}{\left\langle \left| \sum_{f'} \tilde{a}[f']\tilde{a}[\Delta f'] \right|^2 \right\rangle} \tag{6.12}$$

The bispectrum $\mathcal{B}^2[f', \Delta f']$ of the original acceleration signal is now a useful quantity, since it tells us which frequency pairs will contribute to the expectation value of the above sum. If the matrix is sparse, then the number of frequency pairs over which one must sum becomes significantly smaller, and evaluating the sum for each f becomes significantly less computationally expensive. Nonetheless, the transfer function H_{a^2p} still provides a "frequency-smeared" approximation of the expected value of the power signal given the square of the accelerometer signal. In the limit that a single frequency pair $(f_0, f - f_0)$ dominates the sums in equations 6.9 and 6.10, we can approximate the signals as:

$$\begin{aligned} \tilde{p}[f] &= i\tilde{a}[f_0]\tilde{a}[f - f_0] \left(\frac{1}{f_0} + \frac{1}{f - f_0} \right) \\ &= \frac{if\tilde{a}[f_0]\tilde{a}[f - f_0]}{f_0(f - f_0)} \end{aligned} \tag{6.13}$$

$$\tilde{a}^2[f] = 2\tilde{a}[f_0]\tilde{a}[f - f_0] \tag{6.14}$$

Plugging this into 6.12, the factors of $\left\langle |\tilde{a}[f_0]\tilde{a}[f-f_0])|^2 \right\rangle$ cancel assuming there is no uncorrelated noise in the acceleration signal at f. The transfer function is:

$$H_{a^2p} = \frac{if}{2f_0(f-f_0)} \tag{6.15}$$

In this case, the transfer function accurately describes the frequency mixing behavior while accounting for the difference in units between $p[t]$ and $(a[t])^2$. Of course, this approximation of a single perfectly bicoherent frequency pair contributing to the sum in 6.12 is not valid in the CUORE system since there are multiple frequency pairs exhibiting bicoherence that sum to the same frequency (see Fig. 5.13). Still, using the frequency-smeared coefficients from the squares of the auxiliary signals in the denoising algorithm results in an improvement over the linear version.

Here, we use the square of each auxiliary signal $x_i^2[t]$ as a proxy for the thermal response to vibrational noise. Approaches similar to this, which analyze nonlinear systems by including higher-order powers of the input signals, are well described in [193, 194]. We still keep the original auxiliary signals $x_i[t]$ since they may be separately correlated with the detector noise, for example due to capacitive pickup in the detector wiring. We proceed with the denoising algorithm using the new set of input signals, which is twice as large as the original set. In practice, the Fourier amplitudes $\tilde{x}_i[f]$ and $\tilde{x}_i^2[f]$ can be strongly correlated, so one must ensure that the matrices are well-conditioned before proceeding. While we do not expect this model to capture the full nonlinear behavior of the thermal response to the input signals, we do expect to see an improvement in the noise decorrelation over the linear model.

6.5.2 Results of Denoising

The CUORE data have a much higher SNR than the data used to demonstrate the algorithm in Section 6.4. We are therefore able to trigger on pulse events directly and create a set of noise events. The only possible contamination of this set is from very low-amplitude pulses not detected by the CUORE triggering algorithm. This effect is expected to be small because the pulse rate is small (∼3 mHz) and the small amount of signal power due to these pulses is negligible. We first run the algorithm with 50-second long noise windows using the accelerometer and seismometer signals as inputs, then we repeat the algorithm on the original data using the same signals and their squares as inputs. We conduct an independent analysis after each of the different denoising methods is applied to the raw data. We show that the denoising produces a cleaner pulse template as well as a reduced noise power spectrum. Comparing the noise power spectra, we see that the denoising performs excellently on the pulse tube noise, but does not substantially impact the noise below 1 Hz and does little to change the broadband continuum noise below the various peaks (see Fig. 6.11). The nonlinear method is more effective than the linear one, though this is limited to specific bands in frequency space. The most noticeable differences between the linear and nonlinear algorithms are seen at the harmonics of 1.4 Hz and in a region around 280 Hz with many noise peaks (see Fig. 6.12). These regions of frequency space are exactly the ones where a

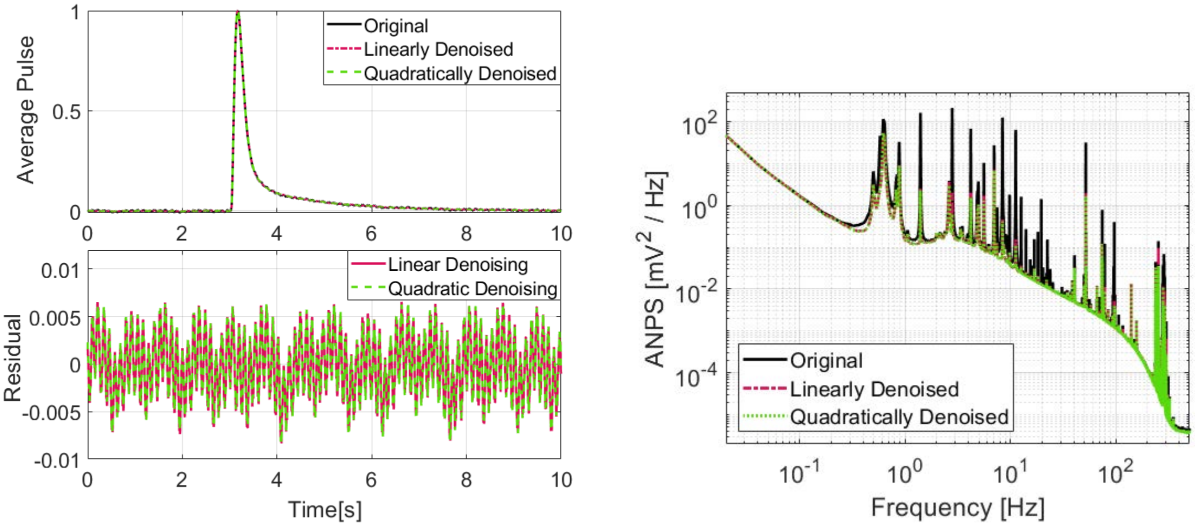

(a) (b)

Figure 6.11: Effects of linear and nonlinear denoising techniques before constructing the OF. (a) Average pulse created using the original, linearly denoised, and nonlinearly denoised data. The denoising procedure removes noise containing many harmonics of 1.4 Hz from the average pulse. (b) ANPS of a CUORE channel before denoising and after the two versions of the denoising algorithm. The linear version removes noise across the signal band, from less than 1 Hz to several hundred Hz. At several frequencies including 2.8 Hz, the nonlinear version of the algorithm further reduces the noise.

significant bicoherence is observed. Together, these observations suggest the presence of a nonlinearity in the mechanical system which is described by the transfer function from the auxiliary devices to the CUORE channel.

To evaluate the performance of the different denoising methods, we first examine the distributions of noise events which are shown in Fig. 6.13. Before the OF is applied, the denoising algorithm reduces the RMS of a typical noise event by more than a factor of 2. We build the OF using a 10-second long signal template. To make the calculation of the OF coefficients as straightforward as possible, we also use 10-second long noise windows to construct the OF as opposed to the 50-second long noise windows we use to perform the denoising. The OF naturally further reduces the RMS noise, but combining the OF with the denoising produces a smaller mean noise event RMS. This value is further reduced when the OF is combined with the nonlinear algorithm. We confirm this effect by building the ANPS of the detector after each signal processing technique is applied (see Fig. 6.14). Taking the square root of the integral of each of the ANPS gives the expected RMS detector noise, the values of which are tabulated in Table 6.4. The denoising improves the noise RMS of the OF signal by 1.0 mV (∼8%). The quadratic version of the algorithm again out-performs the linear one, improving the noise RMS by 1.2 mV (∼10%).

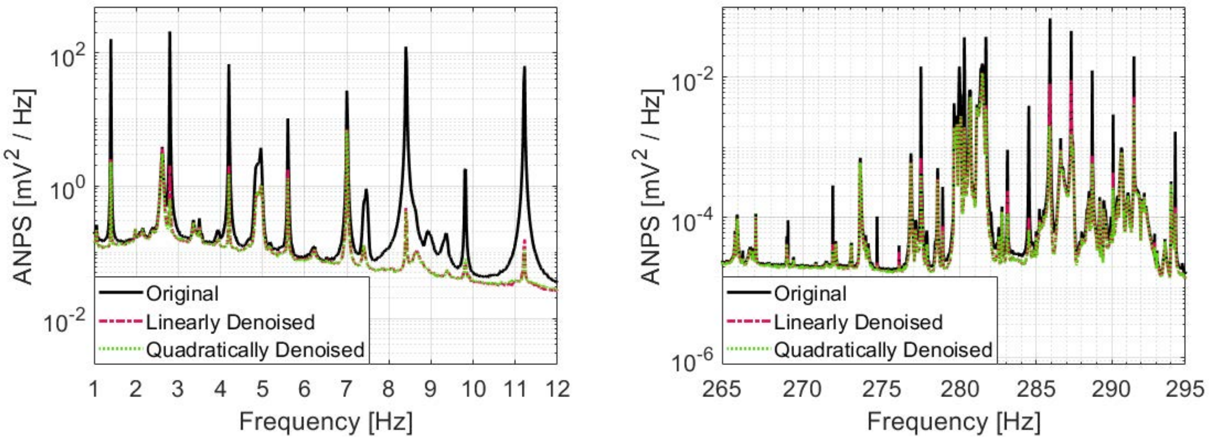

(a) (b)

Figure 6.12: Regions of the CUORE channel ANPS before and after different denoising techniques are applied. (a) ANPS in the region dominated by pulse tube noise at the harmonics of 1.4 Hz. The nonlinear version of the algorithm reduces the noise peaks only slightly more than the linear version. This difference is most noticeable at 2.8, 4.2, and 5.6 Hz. (b) The same ANPS near 280 Hz. Here, there are many peaks where the noise is reduced noticeably more by the nonlinear algorithm than by the linear one, indicating that a nonlinearity in the system is responsible for these peaks.

We now analyze the noise events and pulser events together using the same analysis techniques as described in Section 6.4.2. Fitting the OF amplitude spectrum and calibrating using the known pulser energies gives us our final energy spectrum. The obtained resolutions are reported in Table 6.5, while the ratios of resolutions are reported in Table 6.6. Looking at the noise peak, where the statistical uncertainty on the resolution is smallest, the relative improvement in resolution is 9% when the linear denoising is applied and 10% when the nonlinear version is applied. This again shows the better performance of the nonlinear method. This 10% improvement in energy resolution is similar to the improvement in the noise RMS, verifying that the calibration method does not qualitatively change the effects of the denoising. The pulser resolutions are comparable to the noise resolution as expected, though the resolution of the 3500 keV pulser is slightly worse than expected. Even so, the relative improvement in energy resolution is compatible at all energies (see Fig. 6.15). Together, these results support the hypothesis that the resolution of the pulsers is approximately the same as that of the noise. An energy-dependent contribution to the pulser resolution should not be removed by the denoising algorithm because the auxiliary device signals are completely uncorrelated to the pulse energy. If a non-negligible energy-dependent contribution were present, then the relative change in the resolution of the pulsers after denoising would be different as a function of energy. We observe that the improvements in pulser resolution are

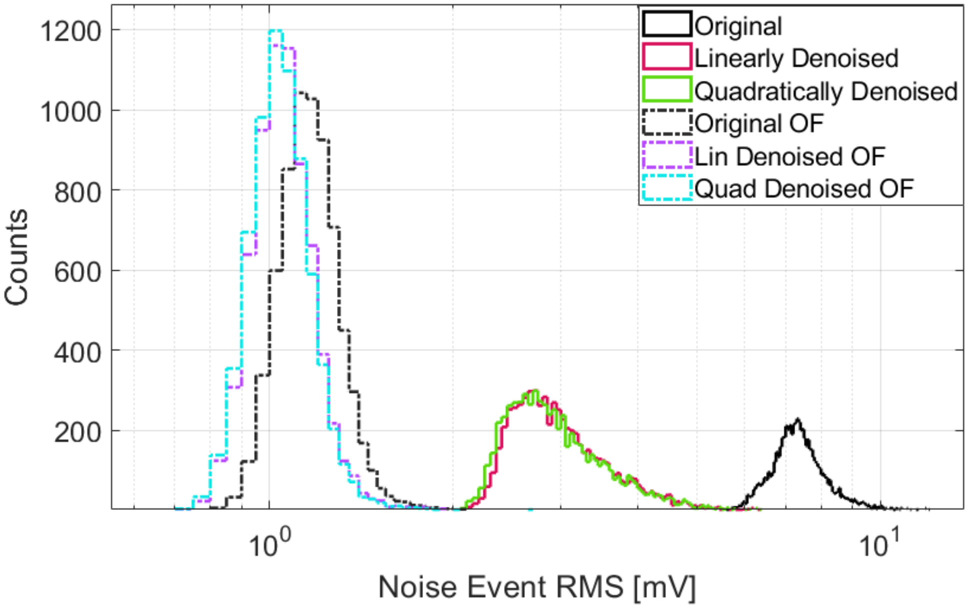

Figure 6.13: Distribution of RMS of noise events before any signal processing, after denoising, and after applying the OF. The lowest noise RMS is achieved when the data are denoised with the nonlinear version of the denoising algorithm and then optimally filtered.

Table 6.4: Average noise RMS of the CUORE channel after applying different combinations of processing techniques. As expected, the OF alone outperforms the noise decorrelation, but the combination of the OF with the denoising gives the best performance. The noise RMS is minimized when the quadratic denoising algorithm is used along with the OF.

Signal Type	Noise RMS [mV]
Raw Signal	7.52
Linearly Denoised Signal	3.16
Quadratically Denoised Signal	3.12
Optimal Filtered (OF) Signal	1.19
OF Linearly Denoised Signal	1.09
OF Quadratically Denoised Signal	1.07

compatible at all measured energies, which suggests that the energy-dependent contribution to the pulser energy resolution is negligible. In summary, the denoising algorithm consistently improves the energy resolution of noise events and pulser events by ∼10%. The nonlinear version of the algorithm adds another ∼1% improvement, though this improvement is within the uncertainty of the measurement.

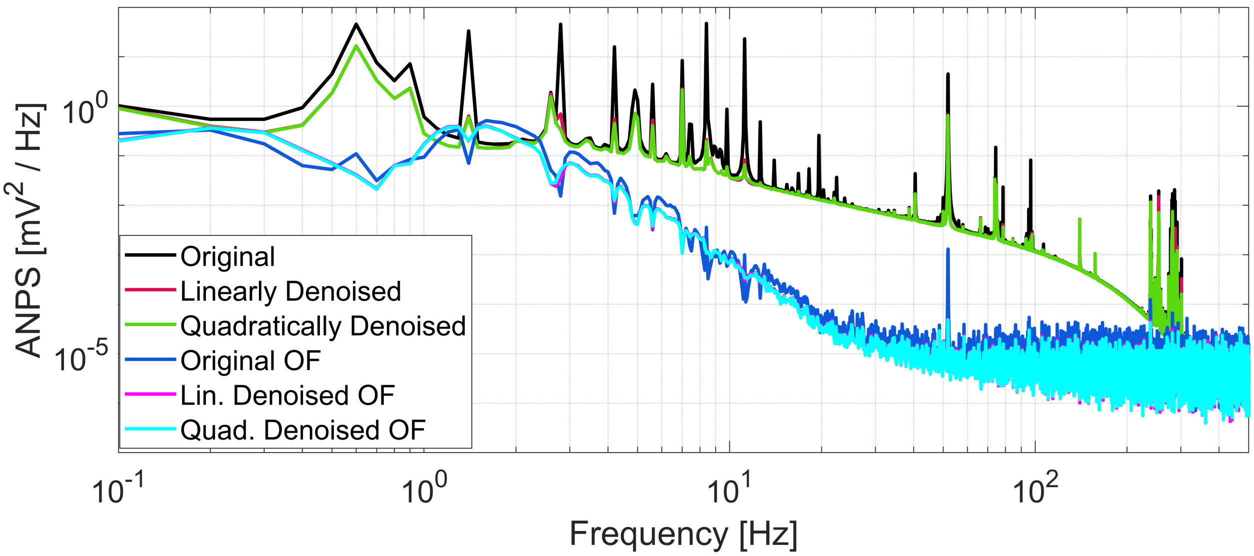

Figure 6.14: Effects of the denoising and the OF on the noise in the CUORE channel. The original and denoised curves (before the OF) are the same as those in Fig. 6.11, but with a coarser binning. This is because the denoising algorithm uses 50-second noise windows, while the OF uses 10-second signal and noise windows. Note the reduction of notches at 1.4 Hz and subsequent harmonics when the data are denoised. Since the OF is normalized to ensure unity gain, reducing the notches has the effect of lowering the continuum of the noise, resulting in lower total noise power in the system.

We can also estimate the timing resolution of the detector before and after the denoising techniques are applied. For a signal of amplitude A, the timing resolution of the digital signal is given by [156]:

$$\sigma_{t_0} = \frac{1}{A}\left[T\sum_{n=-N/2}^{N/2-1}(2\pi f_n)^2\frac{|s(f_n)|^2}{N(f_n)}\right]^{-\frac{1}{2}} \tag{6.16}$$

where T is the length of the signal in seconds, $s(f_n)$ is the Fourier transform of the average pulse, and $N(f_n)$ is the ANPS. The timing resolution is inversely proportional to the amplitude of the signal, and it is a single number for a given average pulse and ANPS. We therefore use this proportionality constant $A\sigma_{t_0}$ as our figure of merit for the different signal processing techniques. We calculate it using the average pulse with amplitude 1 mV and multiply by the value of the calibration function at $A = 1$ mV to report the value in keV·s (see Table 6.7). We find that the linear denoising improves the timing resolution by 1.4% while the quadratic version improves the timing resolution by 1.9%. We note that the improvement in timing resolution is significantly less than that of the energy resolution. This

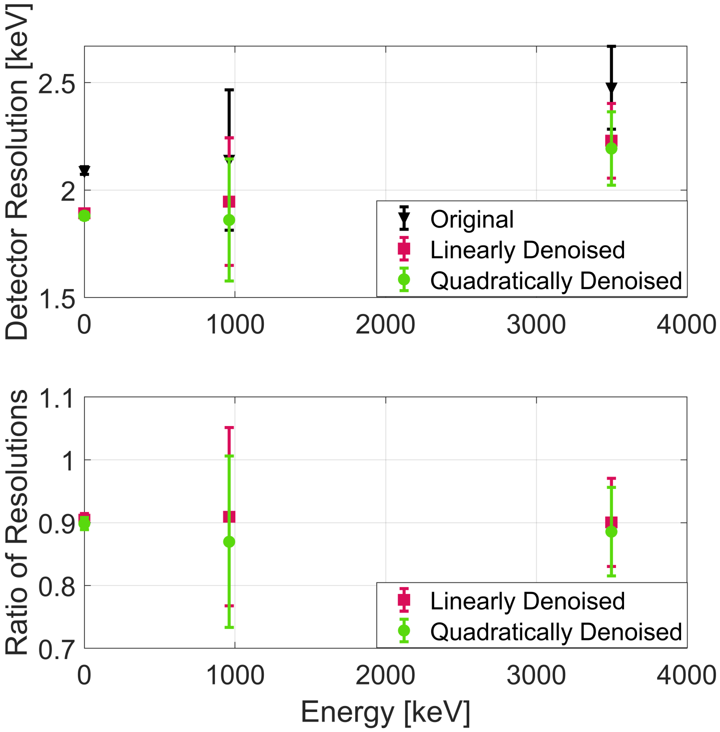

Figure 6.15: Energy resolution of noise events and pulser events before and after the different denoising techniques are applied. The energy resolutions of the pulsers are similar to that of the noise. At all energies, the denoising algorithm improves the resolution of the detector and the nonlinear version further improves the resolution.

suggests that in this channel, vibrational noise does not contribute significantly to the timing resolution.

The main limitations of the algorithm are likely related to the auxiliary devices. First, it is crucial that the auxiliary devices be sensitive to noise in relevant signal bands. The magnitude of the OF applied to the raw timestream peaks near 1.4 Hz, where the pulse tube noise is very well measured. At frequencies below ∼1 Hz, the accelerometer responses roll off significantly. The accelerometers are therefore not sensitive to the sub-Hz noise observed in the CUORE data. The seismometers used in CUORE are designed to measure sub-Hz noise with better sensitivity, so the algorithm can reduce the sub-Hz noise peaks in the CUORE data. These peaks are reduced by 4 to 5 dB and are far from eliminated.

The performance of the denoising is also directly tied to the coherence of the auxiliary devices with the CUORE channel, i.e. the fraction of the power in the auxiliary signals that is correlated with the detector at a particular frequency. It is possible that the coherence of sub-Hz noise between the seismometer signal and detector noise can be increased by moving the seismometers elsewhere on the cryostat, as their locations have not yet been optimized.

Table 6.5: Energy resolution of noise events and pulser events after linear and nonlinear denoising techniques are each applied. The denoising improves the resolution by ∼10%. The quadratic version of the algorithm adds another ∼1% improvement, though this improvement is not statistically significant. Pulser energies were determined during online CUORE data analysis. All energy resolutions are calculated from the optimally filtered and calibrated data.

Signal Type	$\hat{\sigma}_0$ [keV]	$\hat{\sigma}(E = 960\text{ keV})$ [keV]	$\hat{\sigma}(E = 3500\text{ keV})$ [keV]
OF Original	2.09 ± 0.02	2.1 ± 0.4	2.5 ± 0.2
OF Linearly Denoised	1.89 ± 0.02	1.9 ± 0.3	2.2 ± 0.2
OF Quadratically Denoised	1.88 ± 0.02	1.9 ± 0.3	2.2 ± 0.2

Table 6.6: Ratio of denoised energy resolution to original energy resolution using linear and nonlinear versions of the algorithm. The denoising algorithm reduces the noise resolution and the pulser energy resolution by compatible amounts. This is expected because the main contribution to the pulser resolution is the noise. The quadratic version of the algorithm further reduces the resolution, though this reduction is not statistically significant. The resulting resolutions are again compatible across all energies.

Signal Type	$\hat{R}_0$	$\hat{R}(E = 960\text{ keV})$	$\hat{R}(E = 3500\text{ keV})$
OF Linearly Denoised	0.91 ± 0.01	0.91 ± 0.14	0.90 ± 0.07
OF Quadratically Denoised	0.90 ± 0.01	0.87 ± 0.13	0.89 ± 0.07

Table 6.7: Estimated timing resolution of the CUORE channel after applying different combinations of processing techniques. The estimated timing resolution is minimized when the quadratic denoising algorithm is used along with the OF. The improvement is of order 1% after the denoising.

Signal Type	Timing Resolution [keV·ms]
Optimal Filtered (OF) Signal	26.91 ± 0.11
OF Linearly Denoised Signal	26.54 ± 0.11
OF Quadratically Denoised Signal	26.41 ± 0.11

It is also likely that the algorithm can be improved by introducing auxiliary devices in close proximity to the CUORE detector. This suggests that introducing cryogenic vibration sensors into future cryogenic bolometric experiments is to be investigated further. A basic version of this technique using an NTD vibration sensor was demonstrated in Section 6.4.1.

This algorithm is also limited to removing noise that is correlated on timescales of ∼50 seconds. When calculating the transfer functions using 50-second time windows, noise which is only correlated on short timescales is effectively washed out and may not add coherently in the sum. The cross-correlation spectrum of two signals which are correlated on short timescales must have a broad frequency spectrum. This would explain why the peaks in the spectrum are well denoised, but the continuum is not.

Finally, it is instructive to compare the effects of this denoising technique with the noise decorrelation technique used in CUORICINO [179]. In CUORICINO, neighboring calorimeter channels were used as auxiliary devices to remove noise that was correlated between them. The CUORICINO technique effectively removed sharp peaks from the noise power spectrum much like the technique presented here. However, the CUORICINO technique was effective at removing continuum noise below 1 Hz, which our technique fails to do. This suggests that the noise measured in the auxiliary devices used in CUORE is not correlated with the continuum noise in the CUORE detector. In [179], it was found that most CUORICINO channels do not exhibit more than a 5% improvement in OF noise resolution, though the largest improvements were more than 50%. It is unclear whether the CUORE channel showed here is one of the most-improved or least-improved CUORE channels using our noise correlation algorithm, but the 10% improvement in OF resolution we report here out-performs the vast majority of channels in the CUORICINO analysis. In chapter 7, we demonstrate the effects of denoising on the entire CUORE detector, and the results are similar to those seen in CUORICINO.

It should also be noted that the CUORICINO noise decorrelation was done using 11 calorimeter channels as auxiliary channels. In practice, this creates a problem due to "side-pulses," i.e. pulses due to physical events which are present on neighboring calorimeter channels in coincidence with a pulse on the channel of interest. The presence of side-pulses creates difficulties as one must find a way to remove them, either by fitting the pulses on the auxiliary channels and subtracting them or by removing the channel(s) containing side-pulses and re-computing the transfer functions between the remaining channels. The probability of encountering side-pulses grows as the number of channels used in the decorrelation increases. Furthermore, neighboring calorimeters are the most likely to have correlated noise, but they are also the most likely to have coincident pulses due to multi-site physics events. The denoising technique presented here does not incur such problems since there are no side-pulses on the auxiliary devices used in CUORE. Furthermore, this technique denoises calorimeter channels against auxiliary devices, which has not been attempted before in bolometric experiments, and it introduces nonlinear terms motivated by detector response.

Chapter 7

Effects of Denoising on CUORE Data

Having demonstrated an improvement in the quality of CUORE data, we implemented the non-linear multivariate denoising algorithm in the CUORE processing chain described in chapter 4. Because the denoising affects the continuous waveforms, the denoising happens at the very beginning of the processing chain. As a first demonstration of denoising at this scale, we analyze the data from one dataset. Running two parallel analyses – one with the "raw" waveforms that are not denoised and one with denoised waveforms – allows us to isolate the effects of the denoising. We run the processing chain for both analyses only up to the calibration and energy reconstruction step; no coincidence analysis or PSD is applied. After validating the efficacy of the denoising with a single dataset, we turn our attention to analyzing the CUORE data totaling two tonne-years of TeO_2 exposure.

7.1 Denoising a Single Dataset

For this analysis, we restrict the set of auxiliary devices to three accelerometers placed on the Y-beam.[1] As always, we use the accelerometer timestreams and their squares as the input signals. The algorithm is run in parallel, treating each channel as a separate output. The transfer functions from the accelerometers and their squares to a calorimeter channel for a single physics run are shown in Fig. 7.1.

7.1.1 Preliminary Considerations

Due to the very different natures of the physics and calibration data, there are some considerations one must make before proceeding with denoising an entire dataset. For each physics run, the denoising is applied using 50-second windows. The set of noise events is built using the DT triggers and selecting all possible windows in which a trigger is not in the window.

---

[1] During this dataset, the microphones were not functioning correctly, and no other auxiliary devices had been installed yet. This dataset was chosen because I completed the online analysis shift for it, so it was the one with which I had the greatest familiarity.

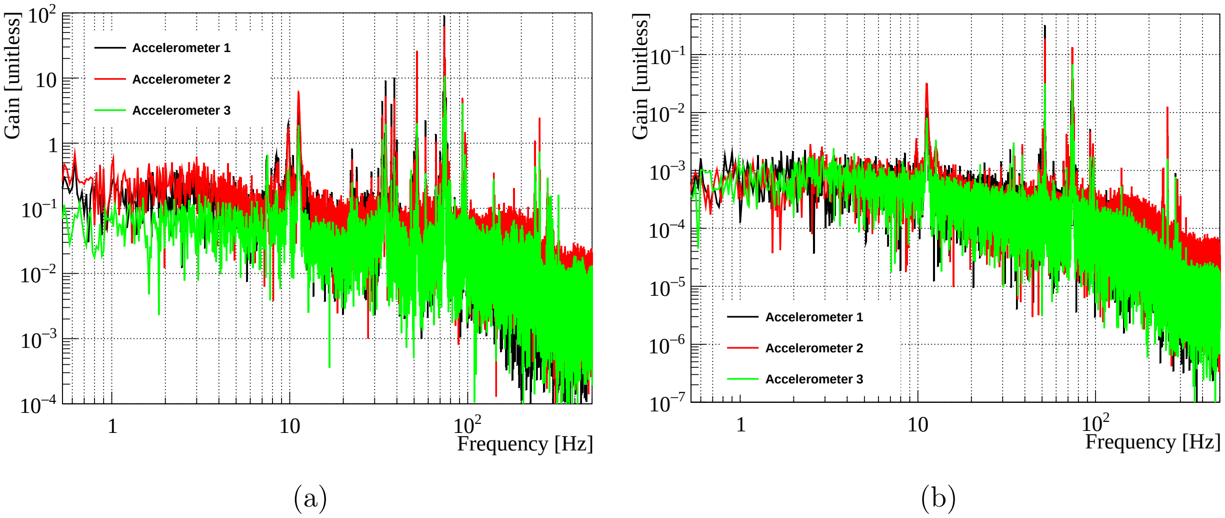

(a)

(b)

Figure 7.1: Magnitudes of the transfer functions from accelerometers and their squares to a CUORE calorimeter using data from one physics run. These transfer functions are built with 50-second-long noise windows and are then used to denoise the calorimeter data. Note the different y-axis values used in each graph. (a) Transfer functions from the three accelerometers to the calorimeter. The third accelerometer points in the vertical direction, and vertical vibrations are dampened by the minus-K system. Hence, the magnitude of the third transfer function is smaller than that of the other two. (b) Transfer functions from the squares of the three accelerometers to the calorimeter. The gain of these transfer functions is much lower than that of the linear transfer functions because the amplitudes of the square accelerometer signals are much larger than the linear signals.

We also place a 20-second dead time after each trigger to ensure that the falling edges of any pulses triggered by the DT are not included in the set of noise windows. The first step of the denoising algorithm is to estimate the expectation values of the power spectral densities and cross-spectral densities of the inputs and output. If the expectation values are taken over too few events, the resulting transfer functions likely do not accurately model the actual transfer function from the auxiliaries to the calorimeter of interest. We therefore require that at least 50 noise events be selected during the physics run to apply the denoising. The DT algorithm is sometimes not tuned adequately during online data-taking, resulting in a "hot trigger" that fires very often, typically on the rising edge of some periodic noise. In the most extreme cases, this results in the channel only having a few 50-second noise windows in a physics run, and the channel cannot be denoised. Unfortunately, this means that sometimes the noisiest channels in CUORE cannot be denoised because of the hot trigger. 983 CUORE channels were active during this dataset for five initial calibration runs, 48 physics runs, and five final calibration runs. Of the 47184 channel-run pairs of physics data, only 349 (or 0.7% of the data) could not be denoised because of a hot trigger. The choice of a 50 noise event cutoff was arbitrary, and a more rigorous study can be done to find the optimal balance between the accuracy of the transfer functions and the amount of data that is not denoised.

For the calibration data, the denoising is slightly more complicated. The higher event rate in calibration data results in very few noise events being collected, so we cannot denoise the calibration data as we do the physics data. However, because the AP is built using signal events from calibration runs, while the ANPS is built using noise events from physics runs (see Section 4.3.1), it is imperative that the calibration data be denoised in some way.[2] To do this, we first denoise the physics data. We then average the transfer functions from the first several physics runs and apply them to the initial calibration data. Likewise, we average the transfer functions from the last several physics runs and apply them to the final calibration data. The number of physics runs averaged together can vary from dataset to dataset depending on the time between physics and calibration data-taking. In this analysis, we averaged the transfer functions from the six nearest physics runs for each set of calibration data. If a channel is not denoised for n of the first six physics runs, then the remaining $6-n$ transfer functions are averaged together and applied to the initial calibration data. If $n=6$, then the initial calibration data is not denoised. (The same logic applies to the last six physics runs and the final calibration.) In this dataset, there was only one channel for which the initial calibration data were not denoised, and all final calibration data were denoised.

Once the denoising was complete, the denoised continuous waveforms were saved to disk for later reprocessing. The DT was re-applied to these waveforms using the same triggering parameters from the online processing. The effects of this step are discussed below.

7.1.2 Results of Denoising

RMS noise

The most natural thing to check immediately after applying the denoising is the ANPS of each channel. Some examples are shown in Fig. 7.2. For virtually all channels, the accelerometers do an excellent job of removing the noise induced by the pulse tubes, as seen in the reduction of the peaks at the harmonics of 1.4 Hz. The low-frequency peaks are not reduced like in the single-channel test discussed in chapter 6, which also uses the seismometers as auxiliary devices. Instead, in some rare cases, the uncorrelated noise in the accelerometers increases the noise below 1 Hz. Taking the square root of the integral of each ANPS to find the RMS noise, we find that 978 of the 983 active channels show an improvement in the RMS noise. In rare cases in which the denoising increases the RMS noise, the denoised data is still used in the analysis. Recall that the RMS noise is defined as the square root of the integral of the ANPS, i.e. $\sqrt{\sum_f J[f]\Delta f}$ where J is the ANPS. Fig. 7.3 shows the RMS noise of each channel before and after denoising and the fraction of RMS noise removed from each channel. Of all

---

[2] A separate study in which only the physics data were denoised showed that the energy resolution of calibration data was worsened using this method. This is expected since the optimum filter built from the denoised noise events is no longer optimal when applied to calibration data that has not been denoised.

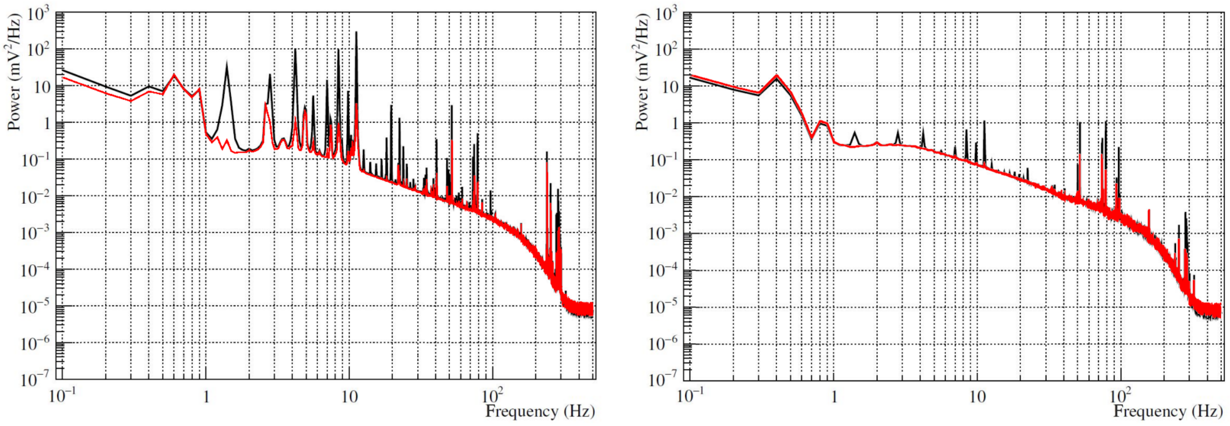

(a) (b)

Figure 7.2: Two sample ANPS before denoising (black) and after denoising (red). In both channels, the main improvement is in the peaks at harmonics of 1.4 Hz due to the pulse tube noise. (a) A channel that is significantly improved by the denoising. The RMS noise after denoising is reduced by 61%. (b) A channel which is slightly worsened by the denoising. The RMS noise after denoising is increased by 1.6%. This occurs because this channel does not experience much pulse tube noise, and the low-frequency noise increases slightly because the accelerometer signals are uncorrelated with the calorimeter signals at frequencies below 1 Hz. In rare cases such as these in which the denoising increases the RMS noise, the denoised data is still used in the analysis.

983 active channels, the best-performing channel exhibits a 74% reduction in RMS noise, while the worst-performing channel exhibits a 7% increase in RMS noise.

We can quantify the noise reduction across the detector in several ways. The first is the total fraction of RMS noise reduced in the entire detector, i.e.:

$$\frac{\sum_{c=1}^{N_{\text{channels}}} \sqrt{\sum_f J'_c[f]\Delta f}}{\sum_{c=1}^{N_{\text{channels}}} \sqrt{\sum_f J_c[f]\Delta f}} = \frac{2.03 \text{ mV}}{3.36 \text{ mV}} = 0.604 \tag{7.1}$$

where c indexes over all active channels, J is the original ANPS of the channel before denoising, and J' is the denoised ANPS. For this dataset, the total noise in the detector is reduced by almost 40%. We can also calculate the average reduction in noise by channel, i.e.:

$$\frac{1}{N_{\text{channels}}} \sum_{c=1}^{N_{\text{channels}}} \sqrt{\frac{\sum_f J'_c[f]\Delta f}{\sum_f J_c[f]\Delta f}} = 0.713 \tag{7.2}$$

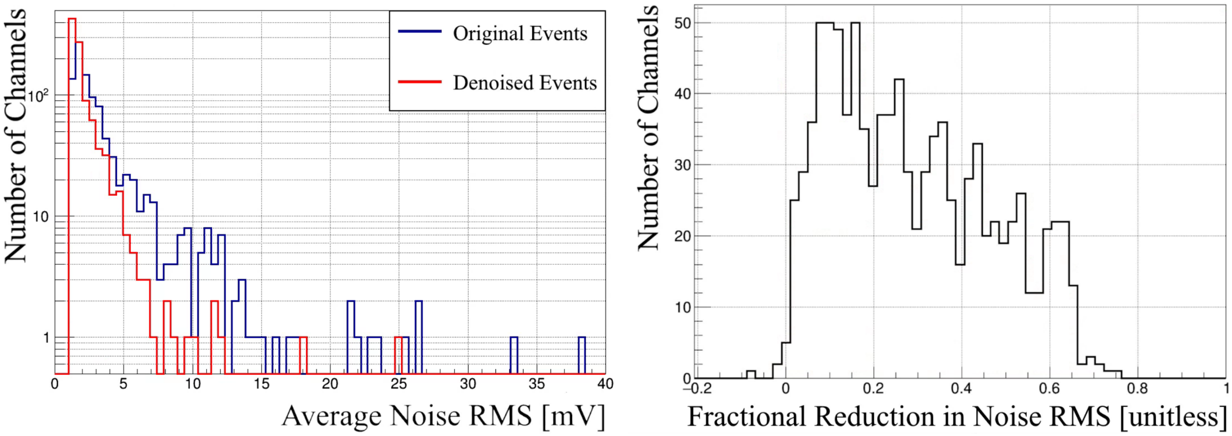

(a) (b)

Figure 7.3: Effect of denoising on the RMS noise of each channel. (a) Overlaid histograms of the RMS noise of each channel before denoising (dark blue) and after denoising (red). The total RMS noise is reduced by 40%, from 3.36 mV to 2.03 mV. (b) A single histogram showing the fraction of RMS noise removed by the denoising. The best-performing channel exhibits a 74% reduction in RMS noise, while the worst-performing channel exhibits a 7% increase in RMS noise.

For this dataset, the average channel has a roughly 29% reduction in RMS noise. Lastly, we can take the median value of the noise reduction by channel, i.e.:

$$\text{med}\left[\sqrt{\frac{\sum_f J'_c[f]\Delta f}{\sum_f J_c[f]\Delta f}}\right] = 0.742 \tag{7.3}$$

For this dataset, the median channel has a roughly 26% reduction in RMS noise.

The fact that these three quantities are not the same suggests a variation in the amount of noise removed from each channel. Fig. 7.4 illustrates this variation. Channels at the top of the detector tend to have more noise removed than those at the bottom. Channels in columns 27 and 28 are also more affected by the denoising than their neighbors. Channels on floor 11 are denoised more effectively than any other floor, as the pulse tube noise dominates their ANPS the most.[3] Of the five channels with a worse RMS noise after the denoising, three of them are located on the bottom floor of the detector.

---

[3] It is well-known that the channels on floors 12 and 13 have the highest amount of noise, both from pulse tubes and from other sources. The channels on floor 11, however, have less noise from sources besides the pulse tubes, thus the greatest fraction of their noise comes from the pulse tubes, making their relative change in noise the greatest.

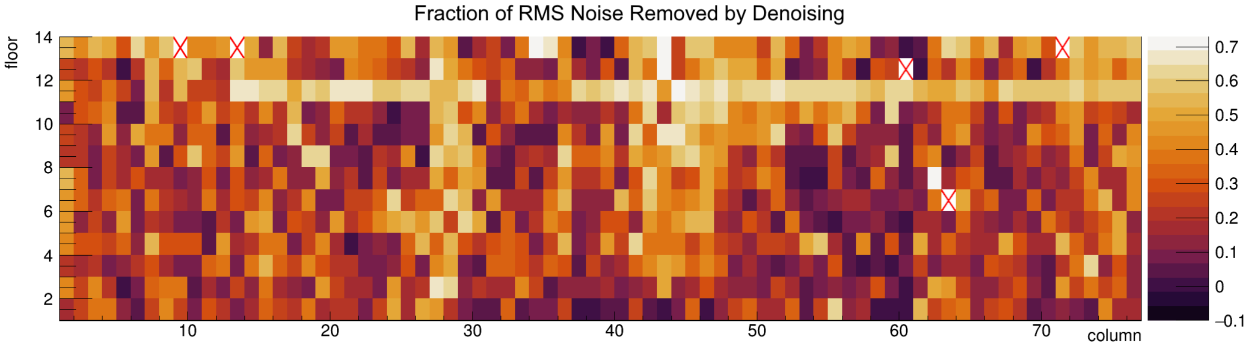

Figure 7.4: Fraction of noise reduction in each CUORE detector achieved by applying the denoising algorithm. Five of the 988 detectors are inactive and marked with a red X in the plot. 978 out of the 983 active detectors show a reduction in RMS noise after applying the denoising algorithm. Channels on floor 11 are denoised more effectively than any other floor, as the pulse tube noise dominates their ANPS the most.

Triggering and Pulse Reconstruction

There are other more subtle effects of the denoising that are also worth mentioning. Because the DT algorithm was re-applied to the denoised data, there was a significant reduction in the number of false triggers, reducing the number of coincident triggers for many channels. This actually aids in retaining channels that can otherwise be lost due to high noise. In the analysis using the raw waveforms, there was one channel on the top floor for which the AP could not be constructed because there were no events containing a single trigger in a 10-second window. The OF could therefore not be built, and the channel was cut from the analysis. After applying the denoising, however, an AP was constructed from 5 events in the calibration data. This channel happened to lie in a “sweet spot” where the hot trigger effect prevented the AP from being built but did not prevent the denoising from occurring, and the channel was “recovered” by the denoising algorithm, leading to a slight increase in exposure for this dataset. This is admittedly a rare effect, but it is a positive effect that the denoising has on the analysis.

Another repercussion of mitigating the hot trigger effect is shown in Fig. 7.5. For a different channel on the top floor, the AP was originally constructed using only seven events. This channel is known to have a very long decay time compared to a typical channel, and the AP is usually cut off prematurely. The high noise in this channel also significantly affected the trigger timing, causing a misalignment of the pulses used to build the AP. The denoising ameliorates this by again reducing the number of false coincident triggers. After denoising, 257 events were used to build the AP. As expected, the denoised AP still has a longer-than-typical decay time, but the premature truncating is not nearly as bad as it was before denoising.

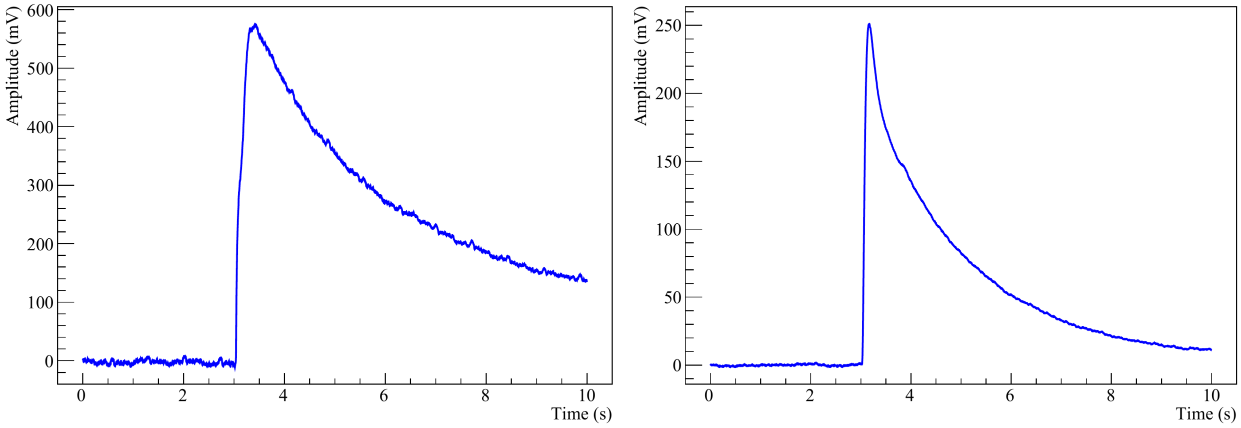

(a) (b)

Figure 7.5: Effect of denoising on the AP of a poorly behaved channel. (a) AP before applying the denoising. This AP is built from seven events which are clearly misaligned based on the bend in the rising edge. The AP is truncated, as the falling edge does not fit within the 10-second window. (b) AP after applying the denoising. This AP is built from 257 events. The alignment is still not great based on the hiccup in the falling edge at $t \approx 4$ seconds, but the truncation is improved, and the AP better represents a true signal in this channel.

Energy Resolution

While the denoising positively impacts the DT, the construction of the AP, and the RMS noise, the most crucial aspect of the denoising is its effect on the detector's energy resolution. For this analysis, we only considered the channels with active silicon heaters, and we only applied the heaterTGS version of the thermal gain stabilization. After completing the processing chain through calibration and energy reconstruction, 940 channels survived all analysis cuts in both rounds of processing. We can therefore directly compare the energy resolutions of these channels. The impact of the denoising expected to be most significant at the baseline energy resolution, i.e. the resolution at the noise peak ($E = 0$ keV), because energy-dependent factors contribute to the energy resolution in a meaningful way. This is discussed in more detail in section 8.1.2.

Fig. 7.6 shows the spatial distribution of the improvement in energy resolution after applying the denoising algorithm. 825 channels (88%) show an improvement in baseline energy resolution. Of the remaining channels, only two show a significant worsening of the energy resolution. Both of these channels show a substantial improvement in baseline RMS voltage, and their APs are virtually unaffected by the denoising, so in theory, they should also exhibit an improved energy resolution. This suggests that a later step in the processing chain, likely stabilization, caused a problem with the energy reconstruction. Even when including these outliers, the average channel has a 4.5% improvement in baseline energy resolution after denoising. This increases to 4.6% after excluding these outliers. The best improvement

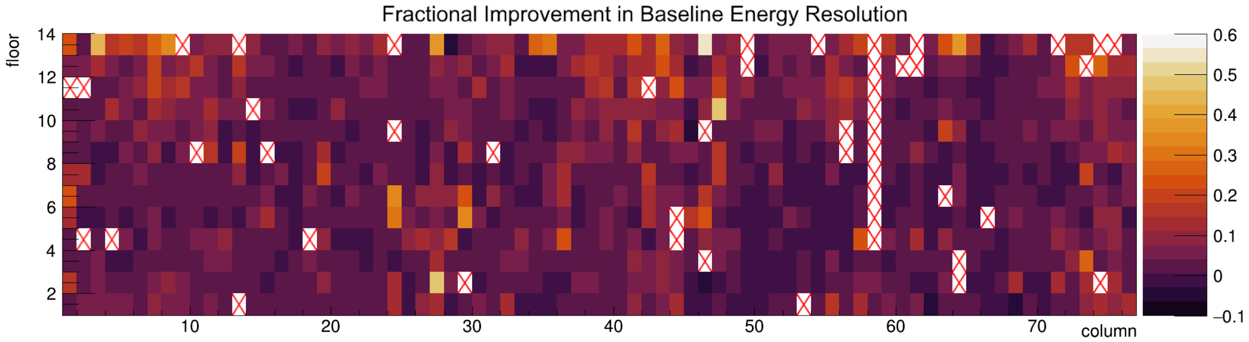

Figure 7.6: Improvement in baseline energy resolution in each CUORE detector achieved by applying the denoising algorithm. 48 of the 988 detectors are not used in the analysis and are marked with a red X in the plot. 825 of the remaining 940 detectors show an improvement in baseline energy resolution after applying the denoising algorithm.

in any one channel is 54%, from 30.4 keV FWHM to 14.1 keV FWHM. The baseline energy resolution (FWHM) of each channel before and after denoising and the percent improvement are shown in Fig. 7.7.

The results from this analysis are comparable to the improvements in energy resolution seen in CUORICINO using nearest neighbor calorimeters to remove correlated noise [179]; the best channels exhibit a decrease of more than 50% in OF noise resolution, while the majority of channels are either unaffected or show a slight improvement of a few percent. It should be stressed that this method, using auxiliary devices, does not suffer from the same problems as the method used in CUORICINO, as discussed earlier in section 6.5.2. This method is not plagued by the issue of side-pulses, and it is highly scalable, as the noise from the auxiliary devices is decorrelated from each calorimeter channel independently.

7.2 Denoising the Two Tonne-Year Dataset

Having demonstrated an improvement in the energy resolution across the entire detector, we now focus on analyzing the CUORE data totaling two tonne-years of TeO_2 exposure. This “2TY” dataset encompasses 28 datasets accrued from 2017 to 2023. This section is intended only to explain the denoising of this dataset. The analysis of the 2TY dataset to search for $0\nu\beta\beta$ decay is the focus of chapter 8.

23 of the 28 datasets comprising the 2TY dataset were denoised. Two of the five datasets that were not denoised were acquired before any auxiliary devices were installed. During the remaining three datasets, the available auxiliary devices suffered from data quality issues, rendering the denoising ineffective. For each of the denoised datasets, a subset of the available auxiliary devices was selected based on the data quality in each device, as some auxiliary devices suffered from high levels of electronic noise during particular datasets. Table 7.1 lists

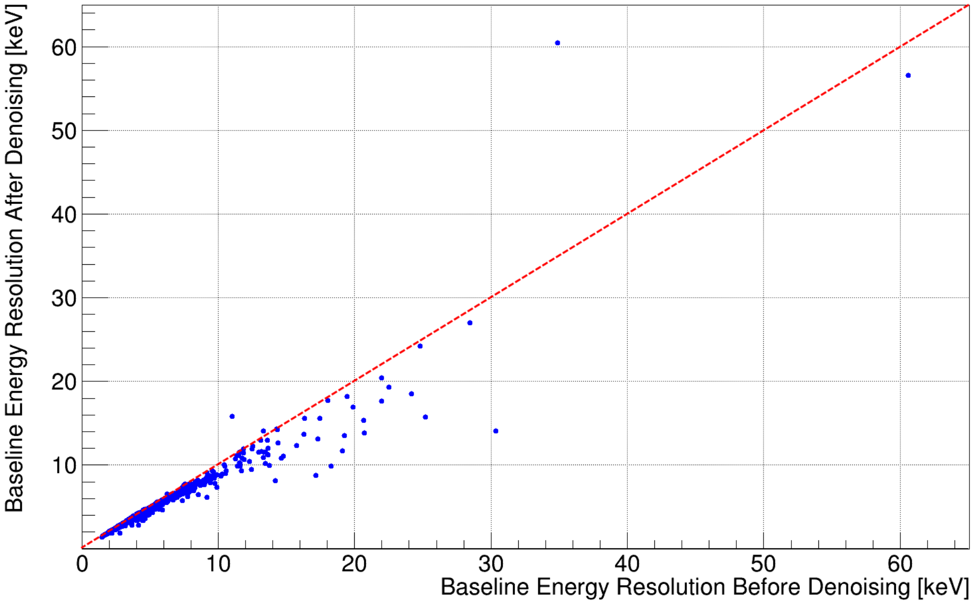

(a)

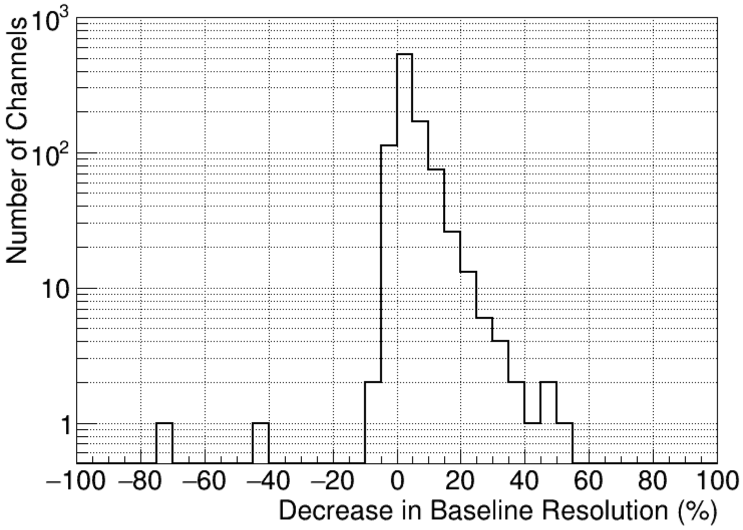

(b)

Figure 7.7: Effect of denoising on the baseline energy resolution of each channel. (a) Scatter plot of the energy resolution (FWHM) of each CUORE channel before the denoising versus after the denoising. The red dashed line represents no change in the energy resolution. (b) Distribution of baseline energy resolution improvements for each CUORE channel when the denoising algorithm is applied. The average improvement is 4.5%. The two channels that exhibit a significantly worse baseline resolution after applying the denoising likely reconstructed poorly due to a problem with the stabilization procedure.

the auxiliary devices used for each dataset. In all cases, both the auxiliary device signals and their squares are used to denoise the data. The analysis described in the previous section was performed using dataset 20. As was mentioned, this dataset was denoised only with accelerometers. The improvements in RMS noise and energy resolution from this analysis are expected to be better than those achieved using microphones but not as good as those obtained using seismometers.

As was mentioned in Section 4.1, calibration data is often shared across datasets. In these instances, we denoise the data once using the transfer functions from the physics runs closest in time to the calibration. For example, the final calibration runs for dataset 3 are also the initial calibration runs for dataset 4. The physics data from dataset 3 concluded six days before the calibration began, while the physics data from dataset 4 started 33 hours after the calibration concluded. It was therefore decided to denoise these runs using the first six physics runs from dataset 4. There are some exceptions to this rule of thumb. If a pulse tube phase scan occurs between physics data-taking and calibration data-taking, we denoise the calibration data using the physics data with the same pulse tube configuration. Also, if calibration data is shared by datasets that are denoised with different auxiliary devices, we select the physics data based on the quality of the auxiliary device data during the calibration data-taking.

A table summarizing the denoising of the 2TY dataset is shown in Table 7.2. Of the datasets that were denoised, 98.6% of the physics data were successfully denoised, and 99.5% of the calibration data were successfully denoised. Dataset 4 performed the worst, with 93.8% of physics data and 98.3% of calibration data denoised, while dataset 8 performed the best, with 99.7% of physics data and 100% of calibration data denoised.

The data comprising the 2TY dataset were denoised before any further processing. All datasets were then processed following the steps outlined in chapter 4.

Table 7.1: Auxiliary devices used to denoise the 2TY CUORE dataset. In all cases, both the auxiliary device signals and their squares are used to denoise the data.

Dataset	**Auxiliary Devices Used**
1	None
2	None
3	3 Microphones
4	4 Microphones
5	4 Microphones
6	None
7	None
8	4 Microphones
9	4 Microphones
10	4 Microphones
11	4 Microphones
12	4 Microphones
13	4 Microphones
14	4 Microphones
15	4 Microphones
16	3 Accelerometers on Main Support Platform
17	3 Accelerometers on Main Support Platform
18	3 Accelerometers on Main Support Platform
19	None
20	3 Accelerometers on Y-beam
21	3 Accelerometers on Y-beam
22	2 Microphones and Accelerometers on Y-beam
23	2 Microphones and Accelerometers on Y-beam
24	2 Microphones and Accelerometers on Main Support Platform
25	1 Seismometer, 2 Microphones, and 1 Accelerometer on Main Support Platform
26	3 Seismometers and 1 Accelerometer on Main Support Platform
27	3 Seismometers, 1 Accelerometer on Main Support Platform, and 2 Accelerometers on Y-beam
28	3 Seismometers, 2 Accelerometers on Main Support Platform, and 2 Accelerometers on Y-beam

Table 7.2: Summary of the denoising of the 2TY dataset.

	Physics Data			**Calibration Data**		
Dataset	**# channel-run pairs denoised**	**# channel-run pairs not denoised**	**% channel-run pairs denoised**	**# channel-run pairs denoised**	**# channel-run pairs not denoised**	**% channel-run pairs denoised**
1	Not denoised					
2	Not denoised					
3	37837	1523	96.13	15638	106	99.33
4	35994	2382	93.79	8709	147	98.34
5	11573	235	98.01	7848	24	99.70
6	Not denoised					
7	Not denoised					
8	41221	107	99.74	8856	0	100.00
9	40202	142	99.65	9840	0	100.00
10	45612	636	98.62	9825	15	99.85
11	43970	310	99.30	7852	20	99.75
12	44119	161	99.64	8852	4	99.95
13	45987	261	99.44	7856	16	99.80
14	36840	552	98.52	7828	44	99.44
15	47703	513	98.94	7836	36	99.54
16	47740	476	99.01	8856	0	100.00
17	35886	522	98.57	8856	0	100.00
18	44280	984	97.83	13544	232	98.32
19	Not denoised					
20	46835	349	99.26	9825	5	99.95
21	37198	156	99.58	9820	10	99.90
22	43979	256	99.42	8829	18	99.80
23	38858	462	98.83	6825	56	99.19
24	37071	283	99.24	7860	4	99.95
25	43194	1041	97.65	7816	48	99.39
26	39681	622	98.46	7792	72	99.08
27	34969	419	98.82	7812	52	99.34
28	36674	680	98.18	6788	93	98.65
Total	917423	13072	98.60	205563	1002	99.51

Chapter 8

The Search for $0\nu\beta\beta$ Decay in CUORE with Two Tonne-Years of Exposure

We now turn our attention to the culmination of everything described so far in this dissertation: the search for $0\nu\beta\beta$ decay in CUORE with two tonne-years (2TY) of TeO_2 exposure. This chapter details the high-level analysis, including the fitting procedure for the $0\nu\beta\beta$ decay, and the results thereof. The 2TY dataset has a total TeO_2 exposure of 2039.0 kg years, or 567.0 kg years of ^{130}Te.

8.1 Analysis Procedure

8.1.1 Blinding

Before beginning the $0\nu\beta\beta$ analysis, we first blind the data to ensure that any choices made during the analysis procedure are not biased in a way that may produce a false positive discovery of $0\nu\beta\beta$ decay (or a more competitive limit on the $0\nu\beta\beta$ half-life in the event that no discovery is made.) The blinding is done by "salting" the data with fake $0\nu\beta\beta$ events to create a peak at 2.528 MeV. A random fraction of events within ±50 keV of the 2.615 MeV γ ray peak is selected, and the energies are artificially reduced by 87 keV. We also perform the reverse process, selecting a random fraction of events within ±50 keV of $Q_{\beta\beta}$ and artificially increasing their energies by 87 keV. If a significant number of real $0\nu\beta\beta$ events exist in the data, the salting procedure thoroughly washes them out or moves them to the 2.615 MeV peak, rendering the analyzers blind to the actual $0\nu\beta\beta$ decay rate. A comparison of the blinded and unblinded data in the region of interest near $Q_{\beta\beta}$ is shown in Fig. 8.1. The analysis is first carried out on the blinded data, then all decisions are finalized, and then the analysis is repeated on the unblinded data.

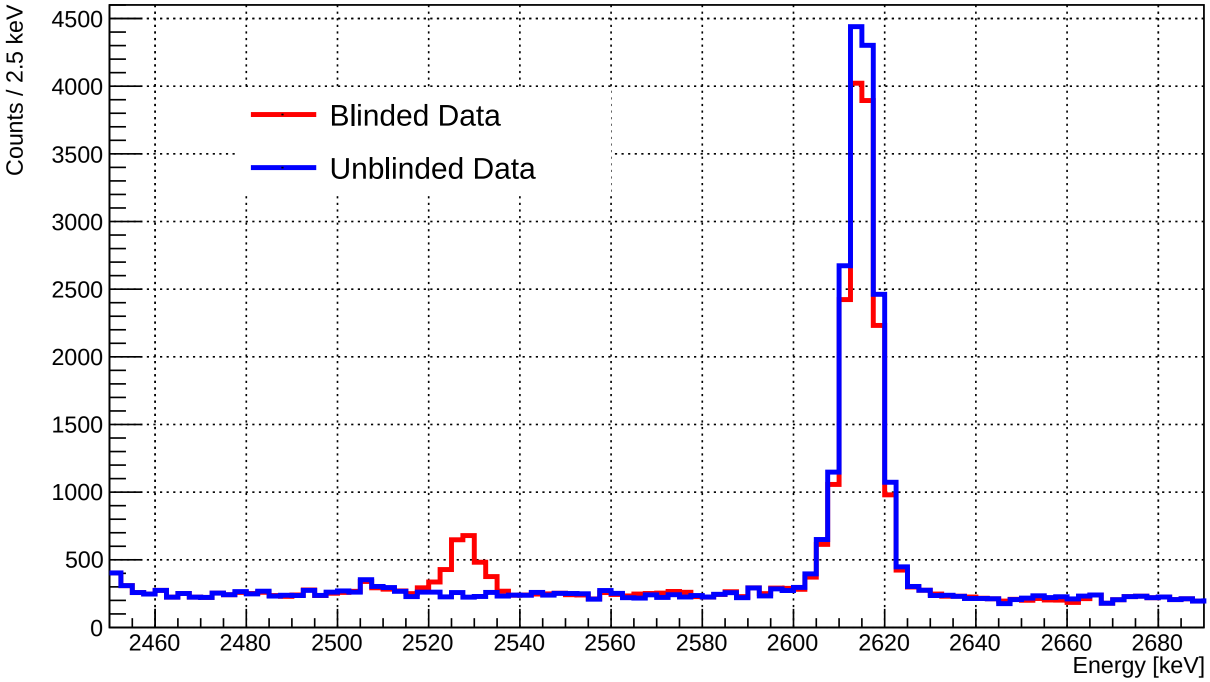

Figure 8.1: Effect of blinding on the energy spectrum of physics events in the region of interest near $Q_{\beta\beta}$. The blinding is done by moving events within ±50 keV of the 2.615 MeV γ ray peak to the region within ±50 keV of $Q_{\beta\beta}$ and vice versa.

8.1.2 Resolution and Bias Scaling

The next step in the analysis is to estimate the expected energy resolution of each detector at $Q_{\beta\beta}$. The energy resolution exhibits a strong dependence on the energy, but because we are blind to any peak that exists at $Q_{\beta\beta}$, we are unable to measure the resolution at $Q_{\beta\beta}$ directly.[1] Instead, we model the detector response using data from the calibration runs then apply a process called resolution scaling to estimate the expected resolution in physics data at $Q_{\beta\beta}$.

The Detector Response Function

To model the response function of each detector, we use only the events near 2.615 MeV in the calibration data. These events mostly come from the γ ray emitted during the decay of ^{208}Tl, so they model the detector response to a mono-energetic set of events. Using the calibration

---

[1] In practice, even if we were not blind to a peak at $Q_{\beta\beta}$, we know that the rate of $0\nu\beta\beta$ decay is so slow that only a few $0\nu\beta\beta$ events would ever be discovered across the entire detector. Therefore, it would still be impossible to measure the energy resolution of each of the 984 active detectors at $Q_{\beta\beta}$ directly for each dataset even without blinding.

data ensures an ample number of events in the peak in each detector, which is not typically the case in the physics data, allowing the resolution of the detector to be measured with some precision. The response function of an ideal detector is a Gaussian distribution $\mathcal{G}(E;\mu,\sigma)$, where μ is the mean of the distribution and σ is its standard deviation. In CUORE, we use a sum of three Gaussians with some constraints:

$$\begin{split}&\mathcal{G}_3(E;E_\gamma,\sigma,\delta,\vec{a}) = \\ &\frac{1}{1+A_L+A_R}\left[\mathcal{G}(E;(E_\gamma+\delta),\sigma)+A_L\mathcal{G}(E;a_L(E_\gamma+\delta),\sigma)+A_R\mathcal{G}(E;a_R(E_\gamma+\delta),\sigma)\right]\end{split} \tag{8.1}$$

Here, E_γ is the expected energy of the γ ray, and δ is the bias of the detector at E_γ. We also impose the constraints $0.996 \leq a_L < 1$, $1.004 \geq a_R > 1$, and $0 \leq A_L, A_R \leq 0.4$ to ensure that there is one main Gaussian centered at $E_\gamma + \delta$ and a smaller side-Gaussian on each side of the main one. We define the vector of parameters $\{a_L, a_R, A_L, A_R\}$ as $\vec{a}$ for later conciseness. While we still do not fully understand the basis for this shape, we know that this "3-Gauss" response function more accurately models the response of the CUORE detectors than a single Gaussian distribution. This may be due to uncertainties from the stabilization trend line and the calibration fitting procedure propagating to the final energy spectrum and causing events at 2.615 MeV to deviate from the expected Gaussian behavior [164]. The response function is one of several components used to fit the calibration data near E_γ. The other components include:

- A multi-Compton shelf. This accounts for 2.615 MeV γ ray events which Compton scatter multiple times, depositing almost but not all of the γ ray energy in the detector. This is modeled as a complementary error function centered at the mean of the main peak of the response function: $\mathrm{cerf}\left[(E-(E_\gamma+\delta))/(\sqrt{2}\sigma)\right]$.
- An X-ray escape peak. This accounts for 2.615 MeV γ ray events that produce a 30 keV X-ray in the tellurium which subsequently escapes the detector. This is modeled as $\mathcal{G}_3(E;E_\gamma - 30\text{ keV},\sigma,\delta,\vec{a})$.
- An X-ray coincidence peak. This accounts for 2.615 MeV γ ray events that occur in coincidence with the aforementioned X-ray escaping from a nearby detector. This is modeled as $\mathcal{G}_3(E;E_\gamma + 30\text{ keV},\sigma,\delta,\vec{a})$.
- A γ escape + coincidence peak.[2] This accounts for 2.615 MeV γ ray events that pair produce creating two 511 keV photons, one of which then escapes the detector, in coincidence with the 583 keV γ ray that is produced in ~85% of ^{208}Tl decays. The result is another peak at 2. This is modeled as $\mathcal{G}_3(E;E_\gamma + 30\text{ keV},\sigma,\delta,\vec{a})$.
- A flat background. This models any spurious events that might survive the PSD cut and degraded α events. This is modeled simply as a constant.

---

[2] This component is only included in the datasets that used the EDCS calibration system.

For each dataset, the calibration data from all detectors in a tower fit this model simultaneously. The response function parameters σ, δ, and $\vec{a}$ are determined separately for each detector, while the amplitudes of the other fit components are shared for all detectors in the tower due to insufficient statistics in each detector. Fig. 8.2 shows a simultaneous fit of one tower-dataset pair.

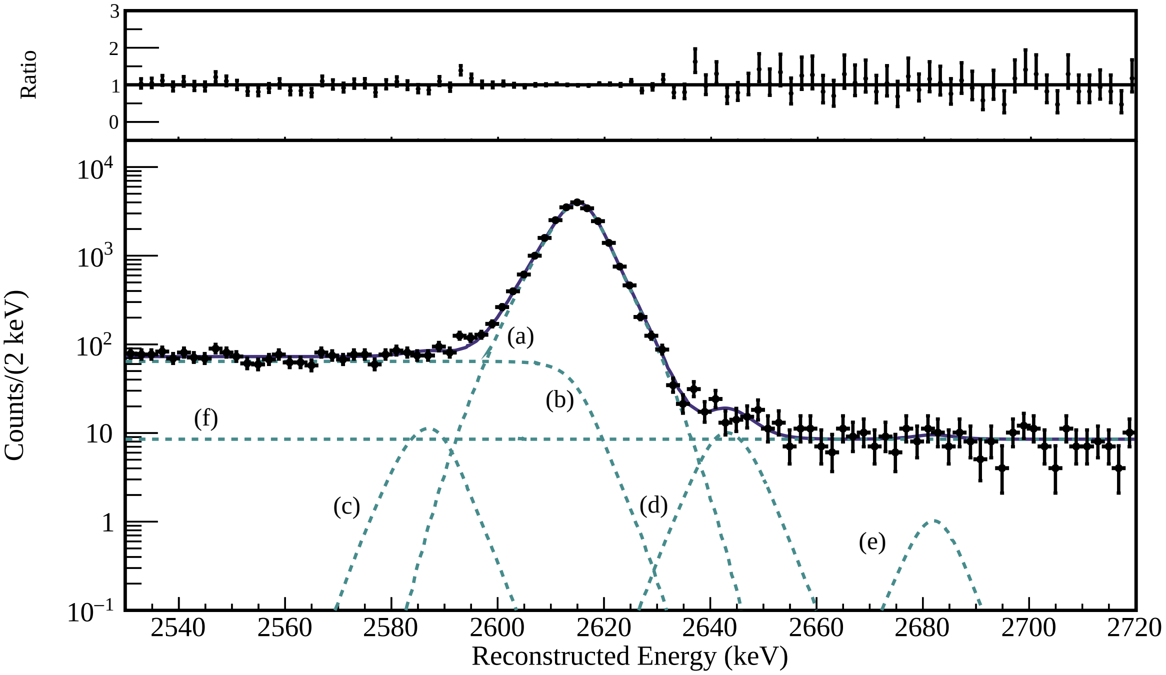

Figure 8.2: Fit of the detector response function at 2.615 MeV from calibration data. The data from all detectors in a tower are fit simultaneously for each dataset. The fit contains the following components: a) A “3-Gauss” response function models the 2.615 MeV γ ray. b) A multi-Compton shelf. c) An X-ray escape peak. d) An X-ray coincidence peak. e) A γ ray escape + coincidence peak. All of these components are described in detail in the text.

The Resolution Scaling Model

Having evaluated the response function of each detector at 2.615 MeV, we now seek to estimate the expected resolution of each detector in the physics data at $Q_{\beta\beta}$. We do this by measuring the resolutions of various peaks in the physics data across a wide range of energies. Unlike in previous CUORE analyses, the 2TY analysis also uses information from the baseline resolutions of each detector. The energy-dependent contributions to the resolution include variations in the event size and the cross-sections of pair-production, Bremsstrahlung, and Compton scattering. These contributions should be independent of the baseline resolution which should be determined by the detector noise (See equation 4.16). Therefore, we assume that these two contributions add to the resolution in quadrature separately in the physics data and the calibration data. Because the physics data do not have ample statistics to measure the energy resolution of each channel individually at 2.615 MeV, we must measure the resolution of the physics data globally, incorporating all events from the entire detector, and then scale the resolution from the 2.615 MeV peak in calibration data to the resolutions of the peaks in the physics data. This gives the following model for the energy resolution of a given detector:

$$\sigma_{\text{physics}}(E) = \sqrt{\sigma^2_{0,\text{physics}} + S(E)\left[\sigma^2_{\text{calibration}}(2.615\text{ MeV}) - \sigma^2_{0,\text{calibration}}\right]} \tag{8.2}$$

where σ_0 is the baseline resolution of the detector, $\sigma(2.615\text{ MeV})$ is the value of σ output by the fit to the calibration data described in the previous section, and $S(E)$ is the resolution scaling function. To find a reasonable functional form of $S(E)$, we analyzed the resolutions of different peaks in the physics data.[3] This analysis revealed a turn-on effect at around 1 MeV. At energies below the turn-on effect, the resolutions of each peak are consistent with the baseline resolution, while at energies above it, the resolutions increase monotonically with energy, so we use the following function to describe the resolution scaling phenomenologically:

$$S(E) = \frac{1}{2}\left[\text{erf}\left(\frac{E - E_0}{\Sigma}\right) + 1\right]\left(c_1 E + c_2 E^2\right) \tag{8.3}$$

Here E_0 is the central value of the turn-on effect, which is constrained to lie between 1.099 MeV and 1.106 MeV. Σ parameterizes the sharpness of the turn-on effect and is constrained to lie between 415 and 597 keV. These constraints come from global fits of the resolution scaling function using all datasets simultaneously. The parameters c_1 and c_2 are constrained to be non-negative to ensure that the resolution scaling increases monotonically with energy. For each dataset, this function is fit to peaks from the physics data, which are pre-selected based on their goodness of fit to Gaussian distributions, using the Bayesian analysis toolkit (BAT) software package [195]. Using BAT allows us to quantify the credibility intervals (C.I.) of the resolution scaling function at any energy which we then use to create a probability distribution for the resolution scaling function at $Q_{\beta\beta}$. The probability distribution of the scaling is used as a dataset-dependent informative prior in the final fit. Fig. 8.3 shows an example of one of these fits for a single dataset and the prior probability distribution on the resolution scaling at $Q_{\beta\beta}$ which is output by the fit.

---

[3] While the resolutions of different peaks from γ rays are used to estimate the detector resolution scaling, it is expected that a $0\nu\beta\beta$ signal will behave quite differently than a γ ray interaction. It is plausible that a $0\nu\beta\beta$ event will more closely resemble a double-escape event in the CUORE detector. If this is the case, it would be more appropriate to use only double-escape peaks when fitting for the energy dependence of the resolution scaling. Studies are ongoing at UC Berkeley to investigate this possibility.

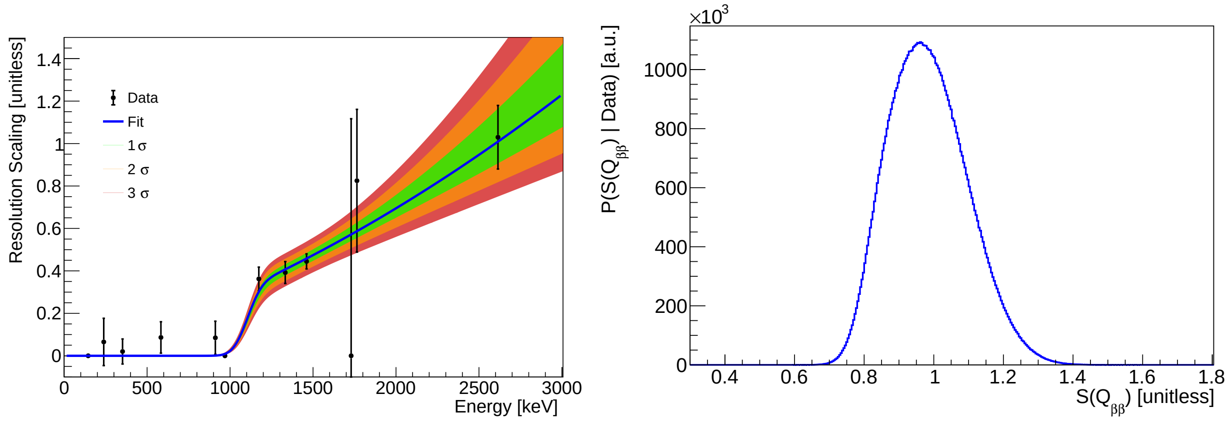

(a) (b)

Figure 8.3: (a) Fit of the energy resolution scaling function for one dataset. The data points come from pre-selected peaks from the physics data which are well-modeled by Gaussian distributions. The regions in color show the 1σ, 2σ, and 3σ C.I. of the scaling. (b) The probability distribution of the resolution scaling at $Q_{\beta\beta}$.

Some pathological cases must be carefully handled when using this technique. For some calorimeter-dataset pairs, $\sigma^2_{0,\text{calibration}} > \sigma^2_{\text{calibration}}(2.615 \text{ MeV})$.[4] In these cases, equation 8.2 can become imaginary, which is obviously problematic when estimating an energy resolution. To overcome this, we assert that $\sigma_{\text{physics}}(E) = \sigma_{0,\text{physics}} \forall\, E$ for these calorimeter-dataset pairs. This is equivalent to setting the bracketed term in equation 8.2 to zero instead of letting it be negative, so it is a conservative estimation of the resolution scaling.

The bias scaling is carried out similarly to the resolution scaling, but the function is much more straightforward. The same global studies of the peaks in physics data used to constrain the resolution scaling function parameters also revealed that the bias function is well described by a quadratic function, $\delta(E) = \delta_0 + \delta_1 E + \delta_2 E^2$. Using the same pre-selected peaks from the physics data that are used in the resolution scaling fit, we fit $\delta(E)$ for each dataset using BAT. The output gives a prior probability distribution on the bias, which is also used as an informative prior in the final fit. Fig. 8.4 shows the fit of the bias for a single dataset and its prior probability distribution.

Combining the fit results from all calorimeter-dataset pairs, we use the exposure-weighted harmonic mean of the resolutions as a figure of merit for the resolution of the full detector. The harmonic mean is used instead of the arithmetic mean because the sensitivity to the half-life of $0\nu\beta\beta$ scales as $1/\sqrt{\Delta E}$, where ΔE is the resolution of a detector, and the resolutions

---

[4] This behavior is admittedly strange, as we do not expect the resolution of a detector to improve with increasing energy. In detectors where the baseline resolution dominates the energy-dependent contributions to resolution at 2.615 MeV, statistical fluctuations can reasonably cause this pathological behavior. However, this behavior is also seen in channels that do not have particularly large baseline resolutions, and further studies are underway to explain these cases.

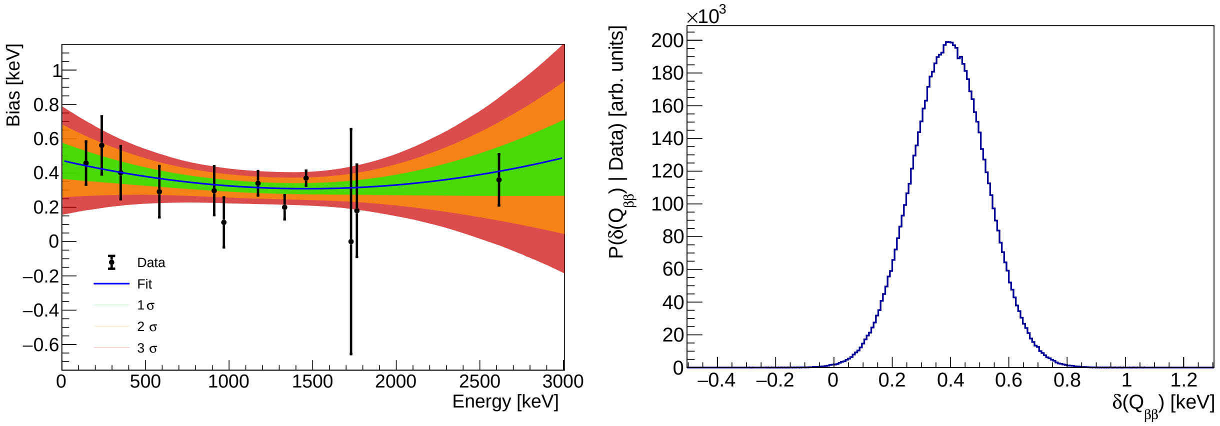

(a)

(b)

Figure 8.4: (a) Fit of the bias function for one dataset. Like in the resolution scaling, the data points come from pre-selected peaks from the physics data which are well-modeled by Gaussian distributions. The regions in color show the 1σ, 2σ and 3σ C.I. of the bias. (b) The probability distribution of the bias at $Q_{\beta\beta}$.

of each detector add in quadrature. The distribution of the FWHM resolution of all detectors for each dataset included in the analysis is shown in Fig. 8.5. Before computing the total resolution, we cut the calorimeter-dataset pairs with FWHM resolutions greater than 23 keV as events from background radioactive peaks can overlap with the $0\nu\beta\beta$ signal in these detectors. This cut value is similar to the one used in CUORE's 1TY $0\nu\beta\beta$ decay analysis (19 keV), which was chosen such that the sensitivity to $0\nu\beta\beta$ decay is only reduced by 0.5% [164]. Pairs with a bias of more than one standard deviation, i.e. $\delta > \sigma$, are cut as well. After applying these cuts, we find that the FWHM resolution of the full detector is (7.540 ± 0.024) keV at 2.615 MeV in calibration data and (7.320 ± 0.024) keV at $Q_{\beta\beta}$ in physics data.

8.1.3 Efficiencies

Because several cuts are applied to the data throughout the processing chain, there is a possibility that a $0\nu\beta\beta$ decay event is removed from the data, thereby reducing the signal efficiency of the experiment. The efficiency of each of these steps must be accounted for to accurately report the rate of an observed $0\nu\beta\beta$ signal (or the limit on the rate if no such signal is observed). We account for several efficiencies in the experiment using the heater pulses from the physics data which are used for stabilization. The energy of these pulses is $\sim$3 MeV, so we can expect that the efficiencies measured using the heater pulses will be close to those at $Q_{\beta\beta}$. The efficiencies we measure with the heater events are:

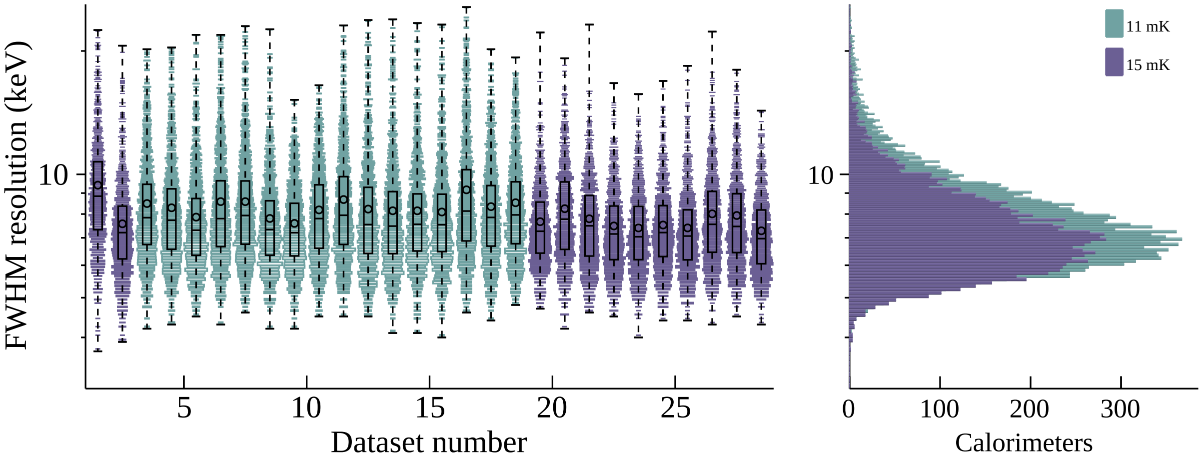

Figure 8.5: Energy resolution of all calorimeters for each dataset. The colors indicate the base temperature of the detector during each dataset. The right panel shows the projection of all calorimeter-dataset pairs onto the energy axis for each base temperature. Calorimeter-dataset pairs with FWHM resolutions greater than 23 keV, which are shown here, are cut from the final $0\nu\beta\beta$ decay analysis and are not shown here.

- Trigger efficiency: the probability of triggering a heater event. This is calculated simply by dividing the number of triggered heater events by the number of heater events injected into the detector.
- Energy reconstruction efficiency: the probability that a triggered heater event reconstructs within 3σ of the mean heater event energy divided by the probability that it reconstructs within 10σ of the mean energy. Here σ is the standard deviation of a Gaussian fit of the reconstructed heater event energies.
- Pileup efficiency: the probability that a triggered heater event does *not* occur in coincidence with another triggered pulse, i.e. the probability of no pileup.

We also calculate two other efficiencies that cannot be estimated using the heater events:

- Heater pile-up efficiency: the probability that a triggered physics event occurs in coincidence with a heater event. This is simply $(1 - t/T)$, where t is the 10-second event window and T is the time between heater pulses in a given detector (typically 570 seconds).

- Heater cross-talk efficiency: the probability that a triggered physics event happens in coincidence with a pulse due to a heater pulse in a *different* calorimeter. This cross-talk is not uniform across the detector; only 10 – 20 detectors experience a notable amount of heater cross-talk in a single dataset, but those that do can have a significant reduction in exposure, ranging from 10% to 80%. Overall, the cross-talk efficiency for the entire detector is ∼99%.

For each calorimeter, these five efficiencies are multiplied together to give the "reconstruction efficiency" (not to be confused with the energy reconstruction efficiency described above). The exposure-weighted reconstruction efficiencies of each calorimeter are then put into a histogram which serves as a probability distribution of the reconstruction efficiency for the entire dataset.

The anticoincidence cut described in chapter 4 also comes with a dataset-dependent signal efficiency. This is the probability that a physical M1 event is moved to M2 (or higher) due to an accidental coincidence. To do this, we select the events from the physics data at a peak centered at 1.46 MeV. This peak is due to the decay of ^{40}K to ^{40}Ar via electron capture, which comprises about 11% of ^{40}K decays. It is expected that the vast majority of these events should be M1 events. The only potential exception is a surface event in which a ^{40}K nucleus on the surface of a calorimeter decays, the nuclear recoil or an atomic X-ray deposits energy in the calorimeter, and the γ ray is deposited in a nearby detector. This effect is negligible, however, as ^{40}K is a bulk contaminant in the CUORE calorimeters. To estimate the anticoincidence cut efficiency, we count the total number of M1 events in the peak at 1.46 MeV and the total number of M>1 events in the same peak, accounting for the background. This is done via a Bayesian counting analysis in BAT, and the result is a probability distribution for the anticoincidence efficiency.

The next efficiency comes from the PSD cut, which was also described in chapter 4. This is estimated using the γ ray events from the physics data.[5] We again conduct a Bayesian counting analysis, this time estimating the total number of signal events in each peak that pass the PSD cut and the number of events that fail the PSD cut. The posterior distributions of these numbers are combined to produce a probability distribution for the PSD efficiency at every γ peak. The efficiency values are then plotted versus energy and fit with a straight line. Similar to the resolution and bias scaling, the PSD efficiency at $Q_{\beta\beta}$ is determined by slicing this fit at $E = Q_{\beta\beta}$ to get a probability distribution.

Another option for computing the PSD efficiency is to analyze M2 events. In this case, the events that pass the cut are multiplets in which both events pass PSD and the energies sum to one of the energies of the ^{60}Co, ^{40}K, or ^{208}Tl γ ray peaks. Events that fail the cut are multiplets in which one of the two events fails PCA and the energies also sum to one of the aforementioned γ ray peaks. Using M2 events produces a slightly higher PSD efficiency than using M1 events, so we use the M1 approach to be conservative.

---

[5] Of course, we exclude the 511 keV "γ ray" peak due to pair production since events in this peak often occur in coincidence with single- and double-escape events.

Table 8.1: Exposure-weighted efficiencies across the entire detector from the 2TY analysis

Total analysis efficiency (from data)	93.4(18)%
Reconstruction efficiency	95.624(16)%
Anti-coincidence efficiency	99.80(5)%
PSD efficiency	97.9(18)%
Containment efficiency (from Monte Carlo)	88.35(9)% [163]

To account for the systematic uncertainty on all of the efficiencies mentioned thus far, we take the probability distributions of the reconstruction, anticoincidence, and PSD efficiencies and multiply them together. The resulting probability distribution serves as an informative prior distribution on the efficiency. This is done for each dataset independently.

The last efficiency we need to consider is the containment efficiency, or the probability that a $0\nu\beta\beta$ decay event will actually manifest as an M1 event in the CUORE detector. As was mentioned in section 4.5, this is calculated using Monte Carlo simulations of the detector. We account for the systematic uncertainty of the containment efficiency by assigning it a Gaussian prior with a standard deviation given by the Monte Carlo simulation. Table 8.1 shows the values of the different efficiencies for the whole 2TY analysis and their uncertainties.

8.1.4 Fitting the $0\nu\beta\beta$ signal

With all inputs to the $0\nu\beta\beta$ fit prepared, we now use a multi-parameter model to fit the energy spectrum in the region of interest (ROI) of [2.465,2.575] MeV. The ROI is chosen to completely cover the predicted $0\nu\beta\beta$ signal given the resolution of the detector. It is not centered at $Q_{\beta\beta}$ because a background peak at 2.506 MeV exists due to the coincidence of the two γ rays emitted by the decay of ^{60}Co, and the ROI must sufficiently cover this peak as well. To measure the rate of $0\nu\beta\beta$ decay observed in the data, we use a fit function that includes the following parameters:

- $\Gamma_{0\nu}$, the decay rate of $0\nu\beta\beta$ decay in inverse years. This is a global parameter that is shared across all datasets.
- $A_{^{60}\mathrm{Co}}$, the activity of ^{60}Co in the experiment. This is a global parameter, but the activity in each dataset is given by $A_{\mathrm{Co}}(t_{\mathrm{ds}}) = A_{^{60}\mathrm{Co}} \exp\{-t_i/\tau\}$, where t_i is the time between dataset 1 and dataset i, and τ is 7.61 years, the mean lifetime of ^{60}Co.[6]

---

[6] Cosmogenic activation of the copper frame of the detector could in principle add a constant term to the ^{60}Co activity, but it was found to be negligible in CUORE.

- BI, the background index in counts / keV / kg / year (ckky). We model this as a flat background that accounts for degraded α particles and multi-Compton scattering of 2.615 MeV γ rays. This is a dataset-dependent parameter.[7]

We fit this function to the energy spectrum in the ROI using BAT, taking the Bayesian approach. These parameters are all assigned uninformative, uniform priors with high cutoff values so as not to bias the result. In addition, for each dataset, we include the informative priors for the analysis efficiency, the resolution scaling, and the energy bias. In total, this gives $4N_{\mathrm{ds}} + 2$ parameters in the fit, where N_{ds} is the number of datasets.

We perform an unbinned fit of all energies $\{E_i\}$ in the ROI to find the best-fit rate and posterior distribution of $\Gamma_{0\nu}$. We seek to maximize the extended likelihood function considering all calorimeter-dataset pairs, i.e.:

$$\begin{aligned}\mathcal{L}(\{E_i\}) = \prod_{ds=1}^{N_{\mathrm{ds}}} \prod_{c=1}^{N_{\mathrm{cal}}} \frac{e^{-\lambda}\lambda^n}{n!} \prod_i \frac{1}{\lambda} \Big[& S\,\mathcal{G}_3(E_i; Q_{\beta\beta}, \sigma_{c,ds}(Q_{\beta\beta}), \delta_{c,ds}, \vec{a}_{c,ds}) \\ & + C\,\mathcal{G}_3(E_i; 2.506\text{ MeV}, \sigma_{c,ds}(2.506\text{ MeV}), \delta_{c,ds}, \vec{a}_{c,ds}) + B\,\frac{1}{W}\Big]\end{aligned} \tag{8.4}$$

where S is the number of signal events, C is the number of ^{60}Co background events, B is the number of flat background events, $\mathcal{G}_3$ is the 3-Gauss detector response function, and W is the width of the ROI. $\lambda = S + C + B$ is the number of events in the ROI for a given calorimeter-dataset pair. The expected number of each type of event is calculated separately:

$$\begin{aligned} S &= \epsilon_a \epsilon_c \Gamma_{0\nu} \rho \frac{\mathcal{E}}{m_A} \\ C &= \epsilon_a A_{\mathrm{Co}}(t_{\mathrm{ds}})\,\mathcal{E} \\ B &= (\mathrm{BI})\,W\mathcal{E} \end{aligned} \tag{8.5}$$

where ϵ_a is the analysis efficiency, ϵ_c is the containment efficiency, $A_{\mathrm{Co}}(t_{\mathrm{ds}})$ is the expected ^{60}Co activity for the given dataset, ρ is the fraction of the detector mass that is ^{130}Te, and $\mathcal{E}$ is the exposure of the dataset after analysis cuts (which is less than mt, the actual mass of the detector times the live-time of the dataset).

8.1.5 Running the fit

When running the fit routine in BAT on the entire 2TY dataset, we experienced significant convergence issues. After over a day of running the fit on a node of Perlmutter, the supercomputer hosted and maintained by the National Energy Research Scientific Computing Center (NERSC), the fit procedure still had not been completed. We decided to divide the data

---

[7] If a global background index is used, then changes in the background index over time, either due to a change in the activity of the backgrounds or Poisson fluctuations from dataset to dataset, could cause the posterior distribution on the analysis efficiency to change in kind. We want to avoid this possibility.

into two equal partitions each comprising 14 datasets. The datasets were randomly chosen while requiring that each partition contain at least seven of the 15 datasets from the first tonne-year and six of the 13 datasets from the second tonne-year. We then ran the fit on each partition independently, treating each partition as a separate pseudo-experiment. Once the fit is complete, we marginalize over all parameters except $\Gamma_{0\nu}$ to produce a posterior distribution on the $0\nu\beta\beta$ decay rate. We then multiply the posteriors from each partition, giving the final posterior distribution of $\Gamma_{0\nu}$.

8.2 $0\nu\beta\beta$ Decay Fit Results

8.2.1 $\Gamma_{0\nu}$ and $m_{\beta\beta}$

The main results from the fit using the Bayesian approach are shown in Fig. 8.6. We observe no statistically significant evidence of $0\nu\beta\beta$ in ^{130}Te. The best-fit decay rate is $\hat{\Gamma}_{0\nu} = 5.5^{+7.3}_{-5.5} \times 10^{-27}$ yr^{-1}, and the 90% C.I. limit on the decay rate is $\Gamma_{0\nu} < 2.0 \times 10^{-26}$ yr^{-1}, corresponding to a 90% C.I. limit on the half-life of $T^{0\nu}_{1/2} > 3.5 \times 10^{25}$ yr.

We also use the frequentist approach outlined in [196] to cross-check the Bayesian result. In the frequentist approach, we calculate the maximum likelihood estimate $\hat{\Gamma}^{(F)}_{0\nu}$ using the likelihood from equation 8.4 and include Lagrange multipliers for the efficiencies, resolution scaling, and energy bias parameters (the superscript (F) serves to distinguish this from the best-fit value from the Bayesian analysis, which can be different). To find the 90% C.L., we calculate the likelihood ratio for different values of $\Gamma_{0\nu}$ in a neighborhood around $\hat{\Gamma}_{0\nu}$. Assuming Wilks' theorem [197] holds in this case, the $\Delta\chi^2$ distribution is given by:

$$\Delta\chi^2(\Gamma_{0\nu}) = -2\ln\left[\mathcal{L}(\Gamma_{0\nu})/\mathcal{L}\left(\hat{\Gamma}^{(F)}_{0\nu}\right)\right] \tag{8.6}$$

This function is shown in Fig. 8.7. The value of $\Gamma_{0\nu}$ at 90% confidence level (C.L.) is the positive value of the rate at which $\Delta\chi^2 = 2.706$, identified by the dotted red line. For this analysis, the 90% C.L. limit on the decay rate is $\Gamma_{0\nu} < 2.1 \times 10^{-26}$ yr^{-1}, or a 90% C.L. limit on the half-life of $T^{0\nu}_{1/2} > 3.4 \times 10^{25}$ yr, which is consistent with the Bayesian result.

Assuming that light neutrino exchange is the dominant mechanism responsible for $0\nu\beta\beta$ decay, we can convert the limit on the half-life to a limit on the Majorana mass of the neutrino, $m_{\beta\beta}$, using equation 2.30. Using the different matrix elements for $0\nu\beta\beta$ in ^{130}Te calculated in [198–206] and assuming $g_A = 1.27$, we place a limit of $m_{\beta\beta} < 70 - 250$ meV. This is the most stringent limit on $m_{\beta\beta}$ obtained using ^{130}Te.

8.2.2 Exclusion and Discovery Sensitivity

Given that the 2TY dataset does not show significant evidence for $0\nu\beta\beta$ decay, a natural question is: what is the largest value of $\Gamma_{0\nu}$ in ^{130}Te that we can expect to exclude 50% of the time given our experimental parameters? We call this value the *median exclusion sensitivity.*

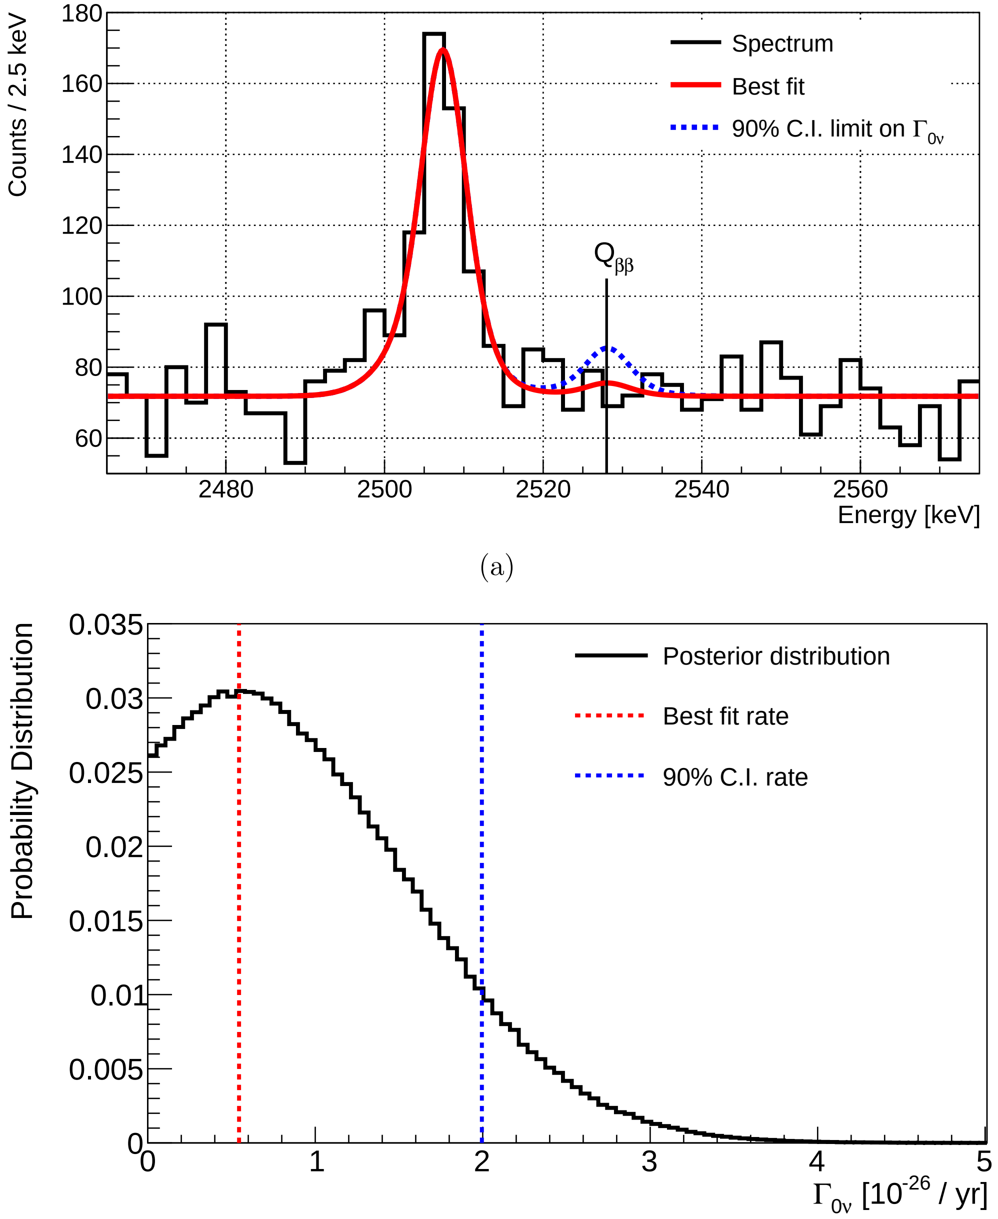

(a)

(b)

Figure 8.6: Results of the $0\nu\beta\beta$ decay fit. (a) The energy spectrum in the ROI. The solid red line is the best fit of the $0\nu\beta\beta$ signal + background model, which is described in detail in the text. The dashed blue line is the sum of the best fit of the background and the 90% C.I. limit on the $0\nu\beta\beta$ rate. (b) The posterior distribution on the $0\nu\beta\beta$ decay rate. The best-fit rate is $\hat{\Gamma}_{0\nu} = 5.5^{+7.3}_{-5.5} \times 10^{-27}$ yr^{-1}, and the 90% C.I. limit on the rate is $\Gamma_{0\nu} < 2.0 \times 10^{-26}$ yr^{-1}, corresponding to a 90% C.I. limit on the half-life of $T^{0\nu}_{1/2} > 3.5 \times 10^{25}$ yr.

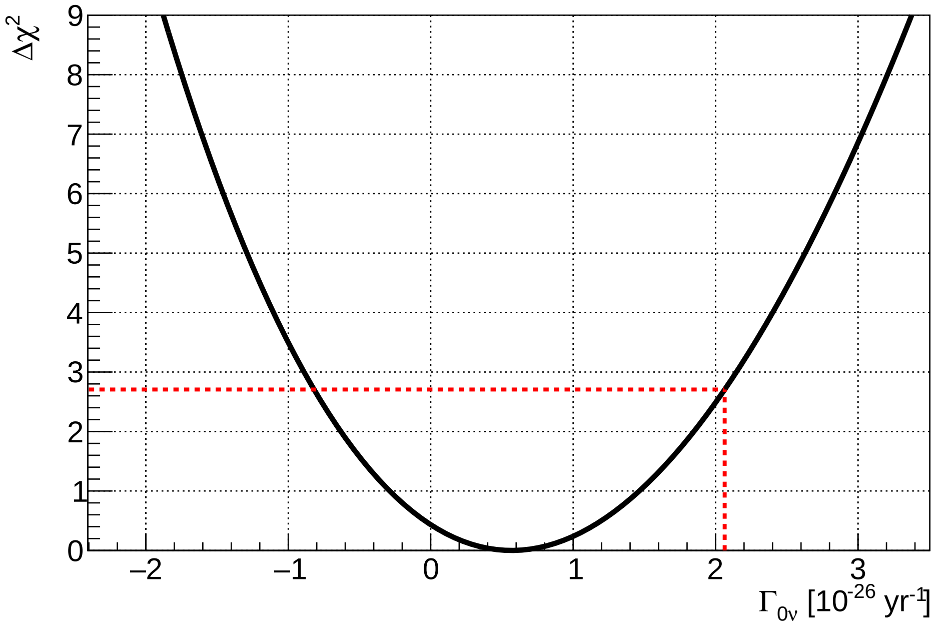

Figure 8.7: Log-likelihood function for the rate of $0\nu\beta\beta$ using the frequentist approach. The 90% C.L. limit on the $0\nu\beta\beta$ decay rate is $\Gamma_{0\nu} < 2.1 \times 10^{-26}$ yr^{-1}, which corresponds to a 90% C.L. limit on the $0\nu\beta\beta$ half-life of $T^{0\nu}_{1/2} > 3.4 \times 10^{25}$ yr.

To answer this question, we generate toy Monte Carlo experiments that use the background index, resolution and bias parameters, and ^{60}Co activities obtained from the $0\nu\beta\beta$ fit.[8] We then generate 10^4 energy spectra with these experimental parameters and fit the $0\nu\beta\beta$ decay signal to each. A histogram of the 90% C.I. $T^{0\nu}_{1/2}$ rates from the toy experiments is shown in Fig. 8.8. The median exclusion sensitivity is $T^{0\nu}_{1/2} > 4.4 \times 10^{25}$ yr. Using the frequentist approach, we find a median exclusion sensitivity of $T^{0\nu}_{1/2} > 4.7 \times 10^{25}$ yr. Since the best-fit value of $\Gamma_{0\nu}$ in the original experiment is positive, we expect the chance of obtaining a stronger limit than the one reported above to be greater than 50%. Using the Bayesian approach, we find that the probability of obtaining a 90% C.I. limit greater than our actual limit is 74%. Using the frequentist approach, this probability is 73%.

We can also ask a different but related question: what is the smallest value of $\Gamma_{0\nu}$ in ^{130}Te for which would we expect to measure the $0\nu\beta\beta$ signal with 3σ significance 50% of the time given our experimental parameters? We call this the *discovery sensitivity*. To answer this, we generate toy Monte Carlo experiments similar to those used to calculate the exclusion sensitivity. For this analysis, we scan across different values of $\Gamma_{0\nu}$ and generate 100 toy

---

[8] We actually re-fit the data assuming no $0\nu\beta\beta$ signal to obtain the background index for the toy experiments. We do this because the best-fit $0\nu\beta\beta$ rate is positive, so some of the background is counted as $0\nu\beta\beta$ signal in the original fit, which slightly lowers the background index.

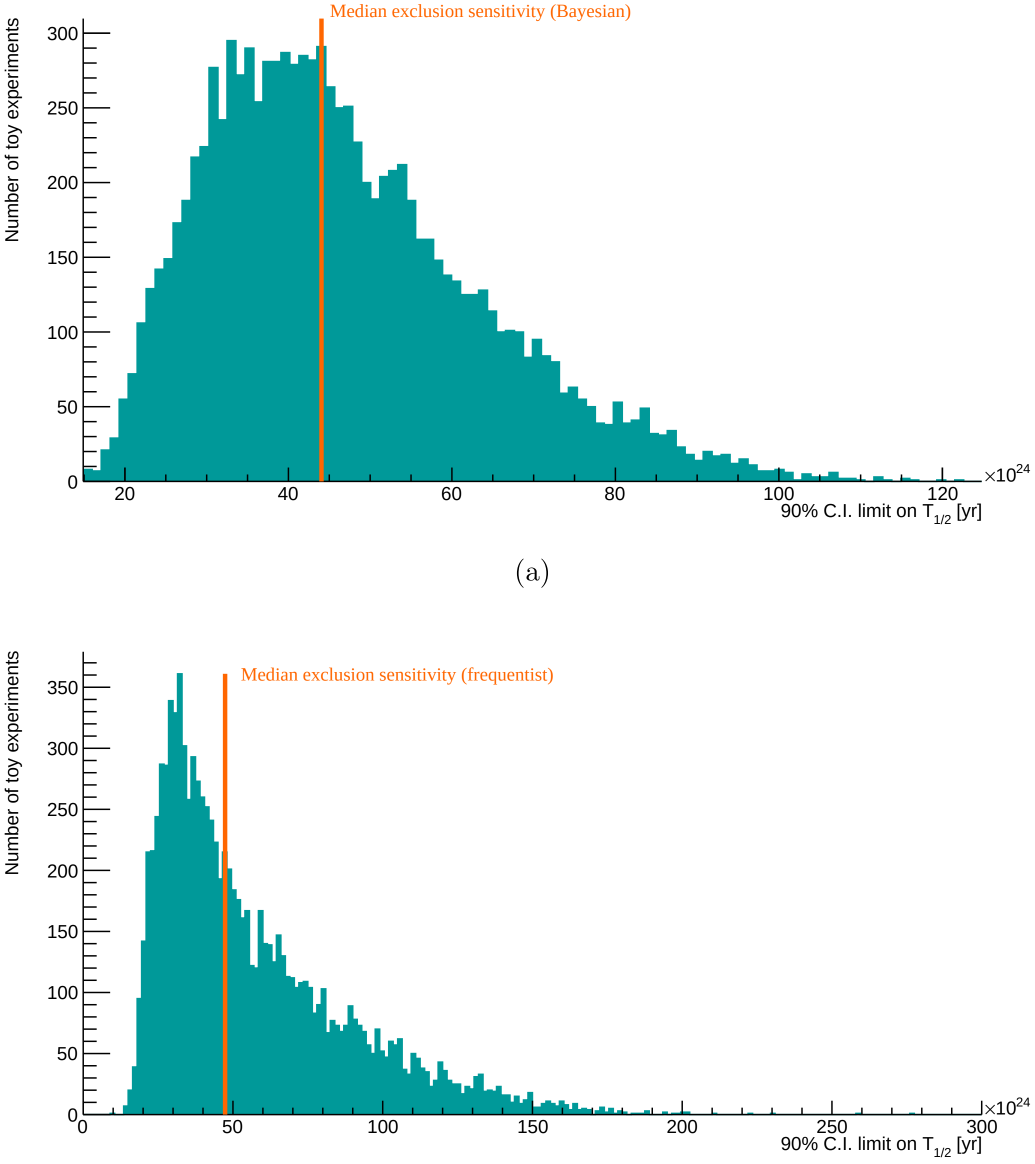

(a)

(b)

Figure 8.8: Exclusion sensitivity from toy Monte Carlo experiments. 10^4 toy Monte Carlo experiments are generated using the background index, resolution and bias parameters, and ^{60}Co activities obtained from the $0\nu\beta\beta$ fit assuming a $0\nu\beta\beta$ decay rate of 0. A best-fit $0\nu\beta\beta$ decay rate is obtained for each toy experiment, and the median best-fit value defines the exclusion sensitivity. (a) The median exclusion sensitivity using the Bayesian approach is $T^{0\nu}_{1/2} > 4.4 \times 10^{25}$ yr. (b) The median exclusion sensitivity using the frequentist approach is $T^{0\nu}_{1/2} > 4.7 \times 10^{25}$ yr.

experiments for each one. For each set of toys, we calculate the following quantity:

$$\Delta\chi^2_{(0)}(\hat{\Gamma}_{0\nu}) = -2\ln\left[\mathcal{L}(\hat{\Gamma}_{0\nu})/\mathcal{L}(0)\right] \tag{8.7}$$

The discovery sensitivity is the smallest value of $\hat{\Gamma}_{0\nu}$ for which half of the experiments have $\Delta\chi^2_{(0)} > 9$, i.e. evidence for $0\nu\beta\beta$ decay with 3σ (frequentist) significance. Fig. 8.9 shows the results from the toy experiments. Because we only fit 100 toy experiments per $0\nu\beta\beta$ decay rate, there is some noise in the data. Thus, we fit the fraction of toy experiments with $\Delta\chi^2_{(0)} > 9$ to an error function and quote the discovery sensitivity as the point where the fit crosses 0.5. We find that the discovery sensitivity is $\Gamma_{0\nu} < 2.6\times10^{-26}$ yr^{-1}, or $T^{0\nu}_{1/2} > 2.6\times10^{25}$ yr, which is 1.8 times weaker than the frequentist exclusion sensitivity.

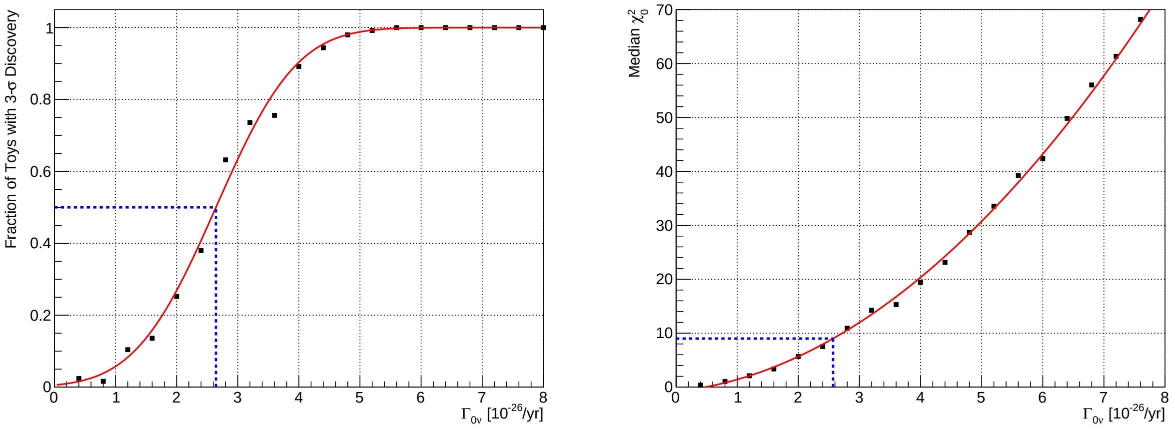

(a)

(b)

Figure 8.9: Discovery sensitivity using toy Monte Carlo experiments. (a) The fraction of toy experiments with $\Delta\chi^2_{(0)} > 9$ as a function of the injected $0\nu\beta\beta$ decay rate (where $\chi^2_{(0)}$ is defined in equation 8.7). There is some noise in the data because we only fit 100 toy experiments per $0\nu\beta\beta$ decay rate, so the data are fit to an error function. The fit crosses 0.5 when the injected rate is 2.6×10^{-26} yr^{-1} (b) The median value of $\Delta\chi^2_{(0)}$ as a function of the injected $0\nu\beta\beta$ decay rate. These data fit a second-degree polynomial. The fit crosses 9 when the injected rate is 2.6×10^{-26} yr^{-1}, in agreement with the result from (a).

8.2.3 Background Index

As discussed above, the background index is a dataset-dependent parameter in the $0\nu\beta\beta$ fit. Fig. 8.10 shows the background index as a function of time. The background index for the entire 2TY dataset is $(1.53\pm0.03)\times10^{-2}$ ckky, and it shows a statistically significant decrease over time from ~0.018 ckky to ~0.011 ckky over more than four years.[9] A natural hypothesis is that a radioactive contaminant with a known half-life is present in the data, and over time, the activity of this contaminant decreases. Thus, we fit the background index to an exponential function of the form $Ae^{-\ln(2)t/T_{1/2}}$. Using all datasets, the best-fit value of $T_{1/2}$ is 14.4 ± 3.5 years. In addition, we observe that the background index increases to its maximum measured value after the first three datasets, indicating possible contamination of the cryogenic environment during the interventions conducted between 2017 and 2018. Excluding the first three datasets, the best-fit value of $T_{1/2}$ is 10.2 ± 2.2 years. Neither of these measured half-lives is consistent with a radioisotope with a significant presence in the CUORE data. The most likely explanation at this time is that a combination of ^{228}Ra, with a half-life of 5.75 years, and ^{210}Pb, with a half-life of 22.3 years, is responsible for this decrease in background index. Given that the background in the ROI is dominated by degraded α particles, it is reasonable to suspect that the background should also include a constant term with respect to time. When fitting the background indices to an exponential decay plus a positive constant, the constant term is consistent with zero, and the reduced χ^2 of the fit is worse than that of the fit without a constant term. In addition, fitting two exponential functions simultaneously yields no information as the fit parameters become highly correlated and have such large errors that one or both half-lives are consistent with zero.

---

[9] Though the background index is observed to decrease with time, this decrease does not exhibit a statistically significant energy dependence. Fitting the background of each dataset to a first-degree polynomial gives background slopes that are consistent with zero.

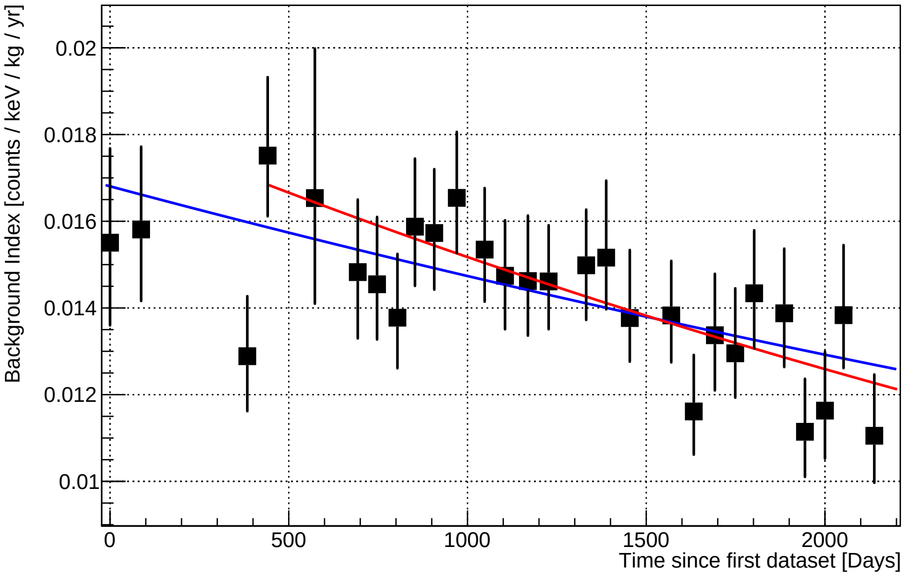

Figure 8.10: Measured background index near $Q_{\beta\beta}$ over time. The blue line is a fit to a single exponential using all datasets, and the best-fit value of $T_{1/2}$ is 14.4 ± 3.5 years. The red line is a fit excluding the first three datasets, and the best-fit value of $T_{1/2}$ is 10.2 ± 2.2 years. Neither of these measured half-lives is consistent with that of a radioisotope which is present in the CUORE data. The observed drop in background index is likely due to the decay of two or more isotopes.

Chapter 9

Looking Forward

As the previous chapters of this dissertation have hopefully shown, CUORE has been an incredibly successful experiment so far. Not only has CUORE achieved the most stringent limit on $0\nu\beta\beta$ in ^{130}Te in the world, but it has also proven that a cryogenic experiment close to the tonne-scale can be operated stably for years of data-taking. CUORE therefore represents a significant step forward in the search for $0\nu\beta\beta$ and in cryogenic technology, which has major implications in many areas of physics and engineering. The non-linear denoising algorithm described in chapters 6 and 7 can be applied to many experiments, not just cryogenic calorimetry experiments, and we hope to see similar techniques used in searches for dark matter interactions and other rare events. In this final chapter, we turn our attention to the future of both $0\nu\beta\beta$ searches using cryogenic calorimeters and denoising algorithms to improve their data quality.

9.1 The Future of Denoising Algorithms for Cryogenic Calorimeter Experiments

The current denoising algorithm used in CUORE builds a transfer function from an array of auxiliary devices to each calorimeter channel. This transfer function is unique to a single physics run, and it is built using only the information from physics runs. One potential shortcoming of this technique is that the transfer functions describing the mechanical system of auxiliary devices and calorimeters may change over the course of a physics run. A static transfer function applied over an entire physics run can fail to account for these changes. An intriguing option is to implement a multivariate nonlinear adaptive filter based on the single-input linear adaptive filter demonstrated in [175]. An adaptive filter could be seeded with the transfer functions calculated using the current technique, then allowed to evolve over the course of a physics run. Of course, this becomes much more computationally expensive, as one must continuously measure the error between the noise in the calorimeter and that which is predicted by the algorithm. This may be feasible to do for 1000 calorimeter channels with modern computing resources such as dedicated GPU farms. Active filters such as these may then have the potential to be integrated into the online data-taking of future experiments.

Another potential shortcoming of the current denoising technique is that vibrations from pulse tubes usually dominate the noise measured by the auxiliary devices during a physics run. Thus, the algorithm can accurately reconstruct the transfer functions at the pulse tube frequencies, but the uncertainty on the transfer functions is potentially much greater at other frequencies where there is very little noise power measured by the auxiliary devices. One possible improvement to this method is to measure the transfer functions while subjecting the detector to a vibrational impulse, such as a sudden hit from a hammer or the dropping of a large object close to the detector. This allows us to measure the transfer function across all frequencies with a high signal-to-noise ratio in the calorimeters and the auxiliary devices, which is not achievable in the physics data. These transfer functions could then be applied to physics and calibration data, leading to better performance of the denoising algorithm.

This approach is loosely analogous to dithering in digital signal processing, where a known noise signal is intentionally added to the input to ensure the sampling process more accurately represents the signal across the entire frequency spectrum [207]. Dithering techniques have previously been applied to evaluate the systematics in Møller scattering experiments [208, 209]. Just as dithering smooths out quantization errors and enhances the accuracy of digital recordings, introducing controlled vibrations to measure transfer functions would reduce the error of the transfer functions arising from very little signal power in the physics data.

We conducted a preliminary test of this dithering denoising technique by jumping on the ground floor of the CUORE hut once every ten seconds. Each jump serves as a vibrational impulse. The resulting transfer functions, shown in Fig. 9.1, are noticeably different from the typical transfer functions found using only physics runs (see Fig. 7.1). While the transfer functions from physics runs are typically flat with prominent peaks at specific frequencies, the transfer functions from the dithering exhibit a significant $1/f$ component. This suggests that the excess $1/f$ noise, which is not well-described by the dynamical thermal model described in chapter 5, may be due to external vibrations. For these tests, we build the transfer functions from 10-second-long noise windows instead of the typical 50-second-long windows to increase the number of noise windows used to construct the transfer functions and to match the rate of the jumping.

Once the transfer functions were acquired, we denoised the data from the dithering tests using the accelerometers located on the Main Support Platform. Fig. 9.2 shows the results of this process for a single channel. Finally, integrating the transfer function over the entire frequency domain gives a single number which we use as a proxy for the overall susceptibility of a calorimeter to vibrational noise. Fig. 9.3 shows this value for every channel and a scatter plot of this value versus the energy threshold for each calorimeter. We observe that channels on the top floors of CUORE, which are known to be more sensitive to vibrations, have the greatest values of the integrated transfer function. The energy threshold is directly proportional to the baseline energy resolution, σ_0. The threshold and the integrated transfer function are $\sim$30% correlated, indicating that a higher susceptibility to vibrations is associated with a worse baseline energy resolution.

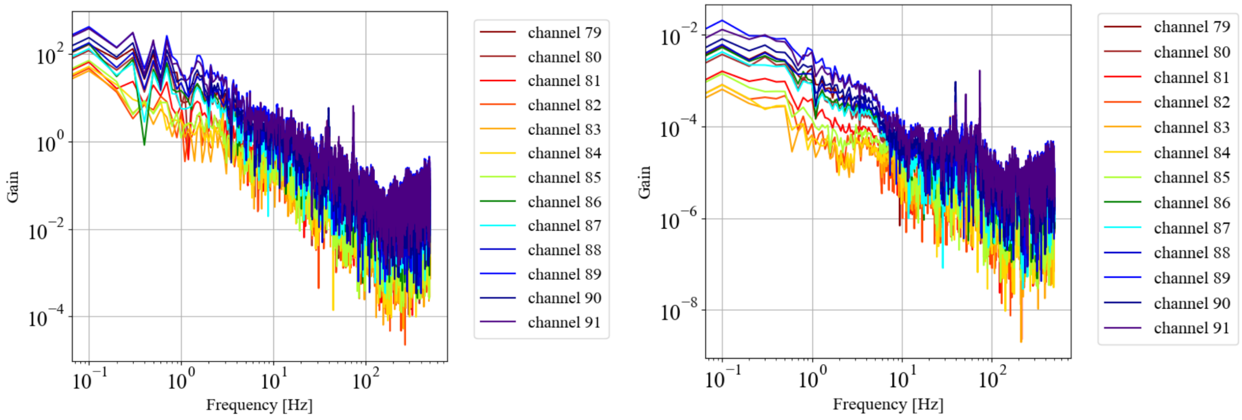

(a) (b)

Figure 9.1: Transfer functions from dithering tests. During the dithering tests, we jumped on the ground floor of the CUORE hut at regular intervals to produce several impulse responses, then use the measured impulse responses of the auxiliary devices and the calorimeter channels. (a) Transfer functions from a vertically-oriented accelerometer on the Main Support Platform to each calorimeter in one column of the CUORE detector. (b) Transfer functions from the square of the accelerometer signal to the same calorimeters. While the transfer functions from physics runs are typically flat with prominent peaks at the pulse tube frequencies, these transfer functions exhibit a significant $1/f$ component.

The results of denoising CUORE data from chapter 7 show that pulse tube noise, while a substantial contribution to the noise itself, is not the main contribution to the energy resolution, even at baseline. Therefore, there must exist a noise source in the signal band that dominates the energy resolution which is not captured by the denoising algorithm. These dithering tests potentially give some insight into this noise source. The transfer functions produced in these tests are well-described by a power law. For $f < 10$ Hz, the transfer functions from the accelerometer to the calorimeters roughly scale as A/f with $A \sim$ 1-100. The transfer functions from the accelerometer squared to the calometers are also scale as A/f, with $A \sim 10^{-5} - 10^{-3}$.

The fact that $1/f$ noise is a significant contribution to the noise, and therefore to the resolution, is not surprising. The fact that it may arise from vibrational noise is more surprising, and it is something that is clearly missing from our current model of vibrations in CUORE. Even if this low-frequency continuum noise is the most significant contribution to the energy resolution, it is not guaranteed that the significant contribution to that noise is due to vibrations. It is well known that electronics can give rise to some $1/f$ noise, but the results from the dithering tests suggest that there is a vibrational component to the $1/f$ noise.

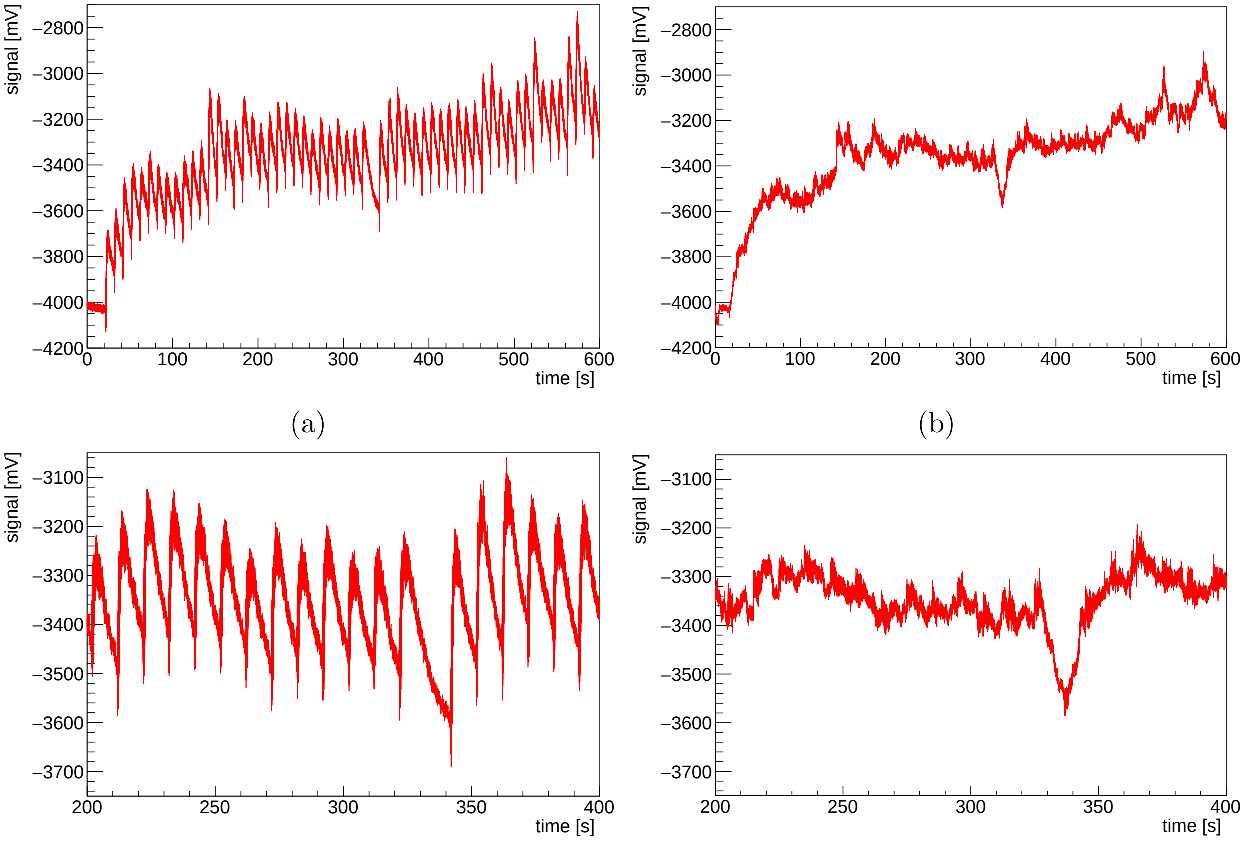

(a) (b)

(c) (d)

Figure 9.2: Results of denoising the dithering data. (a) Timestream of a single calorimeter during the dithering tests. (b) The same timestream after denoising. The denoising technique fails to capture the slow rise in voltage over the 600-second period but manages to remove many of the impulse responses. (c) A zoomed image of 200 seconds during which 19 impulses were created. At $t \approx 340$ seconds, one impulse was left out from the regular series of impulses (d) The same 200 seconds after denoising. At $t \approx 340$ seconds, there is a noticeable deviation in the voltage stream which is rather flat otherwise.

In the future, we hope to use this dithering technique with more auxiliary devices, particularly seismometers, to denoise physics data taken after the first two tonne-years of exposure. Should this technique prove fruitful, it may open the door to even more significant improvements in CUORE's energy resolution.

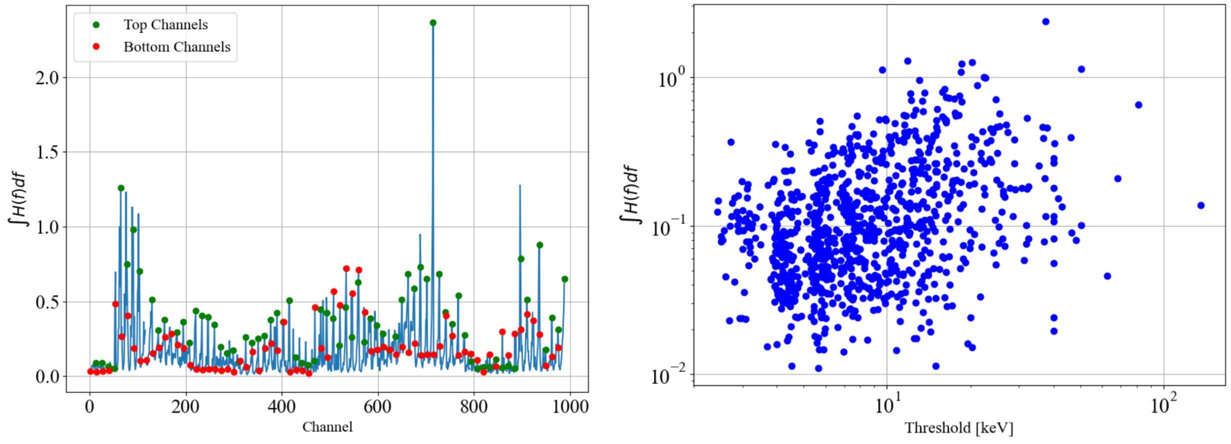

(a) (b)

Figure 9.3: Integrated transfer functions from dithering tests. (a) The integral of the transfer function is a proxy for the susceptibility to vibrational noise. The calorimeters with the highest vibrational susceptibility tend to be located at the top of the detector. (b) Scatter plot of the energy threshold of each calorimeter versus the vibrational susceptibility proxy. These quantities are ∼30% correlated, indicating that a higher susceptibility to vibrations is associated with a worse threshold and a worse baseline energy resolution.

9.2 The Future of CUORE

CUORE's ultimate goal is to reach an exposure of roughly three tonne-years of TeO_2, corresponding to over 800 kg years of ^{130}Te exposure. While CUORE has undoubtedly been a successful experiment so far, there are still many goals that CUORE has not yet reached.

In CUORE-0, the FWHM resolution at $Q_{\beta\beta}$ was estimated to be 5.1 ± 0.3 keV using resolution scaling techniques somewhat similar to those outlined in section 8.1.2 [163]. CUORE's resolution is worse by a factor of ∼44%. One supposed reason for this was that CUORE-0 was operated in a wet cryostat, while CUORE is operated in a dry cryostat with pulse tubes which add noise to the data. The fact that the denoising algorithm successfully removes pulse tube noise from the data, and yet the resolution at $Q_{\beta\beta}$ is mainly unaffected, calls into question this hypothesis. Another possible explanation is that the resistances of the NTD-Ge thermistors used in CUORE are very different from those used in CUORE-0. The CUORE thermistors tend to have larger T_0 values than the CUORE-0 thermistors (see equation 3.10), leading to a stronger temperature dependence of their resistivities. This may make them more susceptible to vibrational noise and cause them to have worse energy resolutions.

As discussed in section 8.2.3, the background index in CUORE at $Q_{\beta\beta}$ shows a statistically significant decrease over time. Should this trend continue, we expect that CUORE will meet its goal of a background index of 0.01 ckky during the third tonne-year of exposure. Couple this with the fact that the array of auxiliary devices used to denoise the data is now larger than ever before, with six fully functional accelerometers and three seismometers taking high-quality data, and the sensitivity gain from a third tonne-year of exposure will likely be more than factor of $\sqrt{1.5}$ that one might expect from a factor of 1.5 increase in exposure.

Ultimately, no matter the improvements made to CUORE in its final years of data-taking, it will always be a background-limited experiment, as α particles and multi-Compton scattering of 2.615 MeV γ rays will inevitably produce background events in the ROI. However, with some new technologies and a different isotope selection, these problems can be addressed in a future experiment.

9.3 CUPID

The CUPID (**C**UORE **U**pgrade with **P**article **ID**entification) experiment [210] is a proposed next-generation experiment that will search for $0\nu\beta\beta$ decay using ^{100}Mo. CUPID is the next in the long line of cryogenic calorimeter experiments based in Italy that was outlined at the beginning of chapter 3. CUPID is currently planned to be housed inside the CUORE cryostat once CUORE is decommissioned. Like CUORE, CUPID will be a segmented detector comprised of O(1000) cryogenic calorimeters, and many of the analysis techniques outlined in this dissertation will readily apply to CUPID. Thus, we will be able to build on our experiences from the years of CUORE data-taking and analysis.

One of the most impactful upgrades from CUORE to CUPID will be the introduction of scintillating calorimeters made from Li_2MoO_4 crystals. These crystals produce scintillation light which can be measured with light detectors placed in close proximity to the calorimeter crystals. When an α particle deposits energy in a calorimeter crystal, the ratio of scintillation light produced to the heat produced is much smaller than that of a β or γ particle. By combining the information from the light detectors and thermal detectors, one can apply a light-yield cut to tag α particles. The α particles can be rejected when conducting a $0\nu\beta\beta$ decay search, significantly reducing the backgrounds. This is the “Particle ID” from which CUPID derives its name. Light detector development for CUPID is an active area of research, with possible technologies including NTD-Ge thermistors amplified via the Neganov-Trofimov-Luke effect [211–214] as well as transition edge sensors [192]. The scintillating calorimeter technology has already been realized in two demonstrator experiments, namely CUPID-0 and CUPID-Mo,[1] which placed world-leading limits on the half-lives of $0\nu\beta\beta$ in ^{82}Se and ^{100}Mo, respectively, at the time of their measurements [186, 215].

In addition, the Q-value of double beta decay in ^{100}Mo is 3.034 MeV, so multi-Compton scattering of 2.615 MeV γ rays from ^{208}Tl will not contribute to the CUPID background in the ROI. The isotopic abundance of ^{100}Mo is only 9.7% [216], but the plan for CUPID is to enrich the molybdenum to $> 95\%$ ^{100}Mo before turning into Li_2MoO_4 (LMO) crystals. The higher Q-value of ^{100}Mo is also a benefit, as the $0\nu\beta\beta$ rate scales as Q^5. The dominant

---

[1] The near overloading of demonstrator experiment names has admittedly caused some confusion in the community.

background in CUPID will be the pile-up of $2\nu\beta\beta$ events, i.e. the occurrence of two $2\nu\beta\beta$ events in a single crystal at the same time. It is quite possible that two events in the $2\nu\beta\beta$ continuum have energies whose sum is close to the Q-value of ^{100}Mo. If these events happen simultaneously in a single detector, then this becomes a background in the CUPID ROI that cannot be cut via particle ID. The half-life of $2\nu\beta\beta$ in ^{100}Mo is roughly 100 times faster than that of ^{130}Te [217]. In CUORE, this is a negligible effect since the background is dominated by α particles and multi-Compton γ ray scattering. In CUPID, tagging pile-up events will be vital, and studying these events in cryogenic bolometers is a key research focus within the collaboration [218].

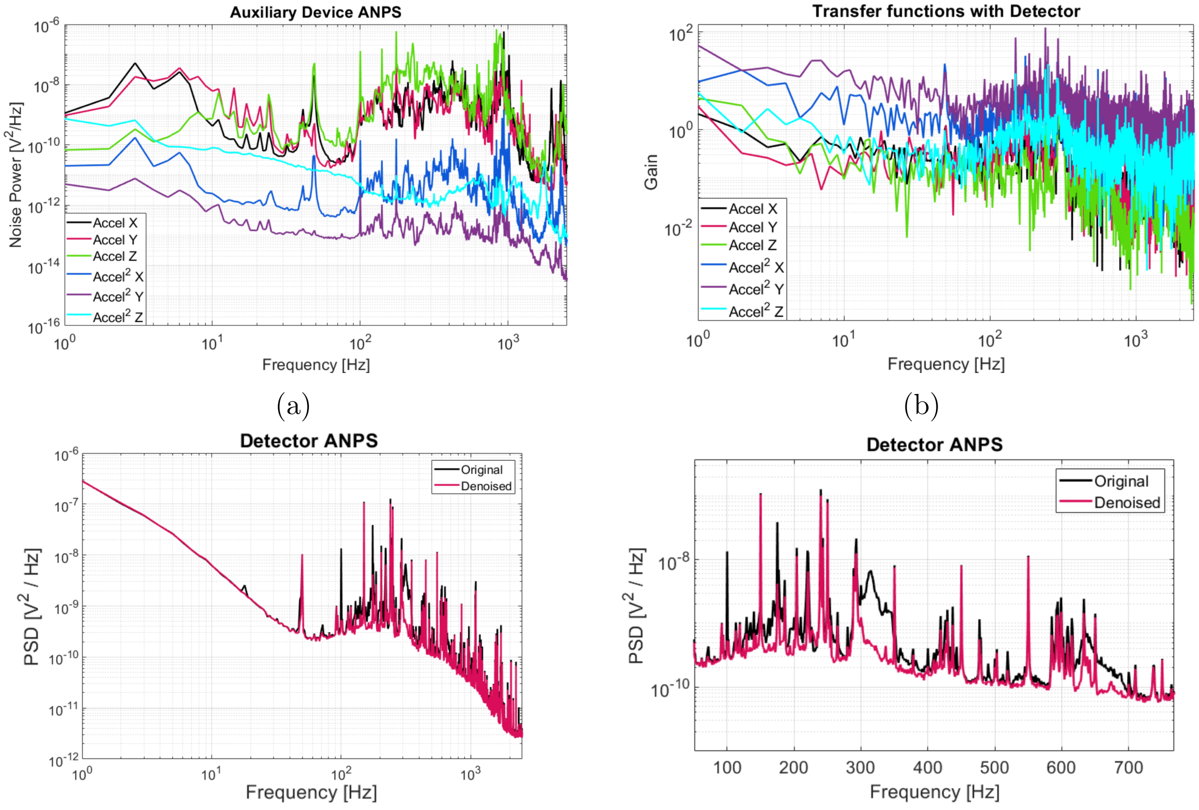

(a) (b)

(c) (d)

Figure 9.4: Results from denoising a thermal detector made with an LMO crystal in an above-ground DR. (a) The ANPS of the accelerometers and their squares when attached to the DR. (b) The transfer functions from each of the auxiliary device signals to the LMO thermal detector. (c) The ANPS of the LMO detector before denoising (black) and after denoising (red). (d) The same LMO detector ANPS before and after denoising, zoomed to a region where the continuum noise due to mechanical resonances at $\sim 300 - 350$ Hz and $\sim 600 - 700$ Hz is substantially suppressed by the denoising.

In the context of noise decorrelation for CUPID, we hope to use data from the accelerometers used in CUORE to predict the vibrational noise in CUPID. In R&D setups with "CUPID-esque" LMO thermal detectors, we have measured vibrations with accelerometers identical to those utilized in CUORE. With these accelerometers, we have denoised the LMO detectors which were cooled to ~16 mK in an above-ground DR. Some results from this exercise are shown in Fig. 9.4. In the future, we hope to apply this technique to the light detectors and conduct dithering tests similar to those performed in CUORE to improve the signal-to-noise ratio of the accelerometers. By taking the transfer functions from the accelerometers to the CUPID-esque detectors and applying them to the accelerometer noise measured in CUORE, we plan to estimate the expected vibrational noise in the CUPID detectors before CUPID even begins to take data.

9.4 Conclusions

The future of the search for $0\nu\beta\beta$ decay is a bright one. While current experiments, including CUORE, report ever-higher limits on the $0\nu\beta\beta$ half-lives in various isotopes, the next generation of experiments is currently being designed and crafted. As a full quantitative description of neutrino oscillations begins to take shape, the nature of the neutrino masses still eludes us. The search for evidence of Majorana neutrinos is a major focus in current physics research, and the global scientific community is fully committed to this effort, as emphasized by the Nuclear Science Advisory Committee Long Range Plan. As was stated at the end of chapter 2, the rest of the 2020s are an optimistic time for the community.

It is often said that we stand on the shoulders of giants, and I sincerely hope that this dissertation reinforces this saying. From Pauli to Fermi, from Cowan and Reines to Fiorini and Avignone, from the giants of old to those who walk among us now, countless individuals have dedicated their lives to furthering our understanding of the universe. It is humbling to be a small part of that story now.

In this dissertation, I have tried to convey what I believe are the most important and intriguing aspects of the search for neutrinoless double beta decay. While theoretical and experimental physics often end up in their own separate corners, I believe that neutrinoless double beta decay is a particularly beautiful area of physics in large part due to the interplay between theory and experiment, which is alive and well. Answering the most fundamental questions about the nature of the universe requires input from many of us. Trying to do so alone is not only impractical but also not nearly as fun as doing so with others. Like so many in the community, I am optimistic that we will answer these questions together and make many friends and fond memories along the way. Perhaps this is the answer to the question: *"Why are we here?"*

Bibliography

[1] Abbé G. Lemaître and A. S. Eddington. “The Expanding Universe”. In: *Monthly Notices of the Royal Astronomical Society* 91.5 (1931), pp. 490–501. ISSN: 0035-8711. DOI: `10.1093/mnras/91.5.490`. eprint: `https://academic.oup.com/mnras/article-pdf/91/5/490/3079989/mnras91-0490.pdf`. URL: `https://doi.org/10.1093/mnras/91.5.490`.

[2] P. J. E. Peebles and Bharat Ratra. “The Cosmological Constant and Dark Energy”. In: *Reviews of Modern Physics* 75 (2003), pp. 559–606. DOI: `10.1103/RevModPhys.75.559`.

[3] S. Weinberg. “A Model of Leptons”. In: *Physical Review Letters* 19.21 (1967), pp. 1264–1266. DOI: `10.1103/PhysRevLett.19.1264`.

[4] A. Salam. *Weak and Electromagnetic Interactions*. Ed. by N. Svartholm. Almqvist & Wiksell, 1968, pp. 367–377.

[5] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. In: *Physical Review D* 2.7 (1970), pp. 1285–1292. DOI: `10.1103/PhysRevD.2.1285`.

[6] Mary K. Gaillard, Paul D. Grannis, and Frank J. Sciulli. “The standard model of particle physics”. In: *Rev. Mod. Phys.* 71 (2 1999), S96–S111. DOI: `10.1103/RevModPhys.71.S96`. URL: `https://link.aps.org/doi/10.1103/RevModPhys.71.S96`.

[7] Hermann Bondi and Ian W Roxburgh. *Cosmology*. Courier Corporation, 2010.

[8] Ya. B. Zeldovich, I. Yu. Kobzarev, and L. B. Okun. “Cosmological Consequences of the Spontaneous Breakdown of Discrete Symmetry”. In: *Zh. Eksp. Teor. Fiz.* 67 (1974), pp. 3–11.

[9] A.D. Dolgov, S. Godunov, and A. Rudenko. “Domain Walls and Matter-Antimatter Domains in the Early Universe”. In: *EPJ Web of Conferences* 158 (2017), p. 05001. DOI: `10.1051/epjconf/201715805001`.

[10] A. Dolgov. “Matter-antimatter domains in the universe”. In: *Nuclear Physics B - Proceedings Supplements* 95.1 (2001). Proceedings of the 6th International Conference on Frontiers in Particle Astrophysics and Cosmology, pp. 42–46. ISSN: 0920-5632. DOI: `https://doi.org/10.1016/S0920-5632(01)01051-9`. URL: `https://www.sciencedirect.com/science/article/pii/S0920563201010519`.

[11] R. Feynman. *Fun to Imagine*. BBC Television. 1983.

[12] A. D. Sakharov. "Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe". In: *ZhETF Pisma Redaktsiiu* 5 (1967), p. 32.

[13] "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations". In: *Nature* 580.7803 (2020), pp. 339–344.

[14] Tomasz Wachala. "Recent Results on the CP Violation Search in the Accelerator Neutrino Oscillations". In: *Acta Phys. Polon. B* 48 (2017), p. 1969. DOI: `10.5506/APhysPolB.48.1969`.

[15] Son Cao. "Long Baseline Neutrino Oscillation Results (T2K and NOvA)". In: 2023. arXiv: `2310.09855 [hep-ex]`.

[16] Gerhart Lüders. "Proof of the TCP theorem". In: *Annals of Physics* 2.1 (1957), pp. 1–15.

[17] Antonio Riotto and Mark Trodden. "RECENT PROGRESS IN BARYOGENESIS". In: *Annual Review of Nuclear and Particle Science* 49.1 (1999), pp. 35–75. DOI: `10.1146/annurev.nucl.49.1.35`. URL: `https://doi.org/10.1146/annurev.nucl.49.1.35`.

[18] Tsung-Han Yeh et al. "Probing physics beyond the standard model: limits from BBN and the CMB independently and combined". In: *Journal of Cosmology and Astroparticle Physics* 2022.10 (2022), p. 046.

[19] Planck Collaboration et al. "Planck 2018 results - VI. Cosmological parameters". In: *A&A* 641 (2020), A6. DOI: `10.1051/0004-6361/201833910`. URL: `https://doi.org/10.1051/0004-6361/201833910`.

[20] Patrick Huet. "Electroweak baryogenesis and the standard model". In: *1st International Conference on Phenomenology of Unification: from Present to Future*. 1994, pp. 77–91. arXiv: `hep-ph/9406301`.

[21] Edward W. Kolb and Michael S. Turner. *The Early Universe*. Vol. 69. 1990. ISBN: 978-0-201-62674-2. DOI: `10.1201/9780429492860`.

[22] Masataka Fukugita and Tsutomu Yanagida. "Baryogenesis without grand unification". In: *Physics Letters B* 174.1 (1986), pp. 45–47.

[23] V.A. Kuzmin, V.A. Rubakov, and M.E. Shaposhnikov. "On anomalous electroweak baryon-number non-conservation in the early universe". In: *Physics Letters B* 155.1 (1985), pp. 36–42. ISSN: 0370-2693. DOI: `https://doi.org/10.1016/0370-2693(85)91028-7`. URL: `https://www.sciencedirect.com/science/article/pii/0370269385910287`.

[24] S.Yu. Khlebnikov and M.E. Shaposhnikov. "The statistical theory of anomalous fermion number non-conservation". In: *Nuclear Physics B* 308.4 (1988), pp. 885–912. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(88)90133-2`. URL: `https://www.sciencedirect.com/science/article/pii/0550321388901332`.

[25] Wilfried Buchmüller, Christoph Greub, and Peter Minkowski. "Neutrino masses, neutral vector bosons and the scale of B- L breaking". In: *Physics Letters B* 267.3 (1991), pp. 395–399.

[26] Wilfried Buchmüller, Roberto D Peccei, and Tsutomu Yanagida. "Baryogenesis and lepton number violation". In: *Annual Review of Nuclear and Particle Science* 55 (2005), pp. 311–355.

[27] Y. Chikashige, R.N. Mohapatra, and R.D. Peccei. "Are there real goldstone bosons associated with broken lepton number?" In: *Physics Letters B* 98.4 (1981), pp. 265–268. ISSN: 0370-2693. DOI: `https://doi.org/10.1016/0370-2693(81)90011-3`. URL: `https://www.sciencedirect.com/science/article/pii/0370269381900113`.

[28] Ettore Majorana. "Teoria simmetrica dell'elettrone e del positrone". In: *Il Nuovo Cimento (1924-1942)* 14.4 (1937), pp. 171–184.

[29] Clyde L Cowan Jr et al. "Detection of the free neutrino: a confirmation". In: *Science* 124.3212 (1956), pp. 103–104.

[30] Raymond Jr. Davis. "Solar Neutrinos. II. Experimental". In: *Physical Review Letters* 12 (1964), pp. 303–305.

[31] John N. Bahcall et al. "Solar Neutrinos: The First Thirty Years". In: *Reviews of Modern Physics* 70.4 (1998).

[32] Arthur B. McDonald. "The Nobel Lecture: The Sudbury Neutrino Observatory: Observation of Flavor Change for Solar Neutrinos". In: *Reviews of Modern Physics* 75.3 (2003).

[33] B. Pontecorvo. "Mesonium and anti-mesonium". In: *Sov. Phys. JETP* 6 (1957), p. 429.

[34] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the unified model of elementary particles". In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: `10.1143/PTP.28.870`.

[35] Andrew G Cohen, Sheldon L Glashow, and Zoltan Ligeti. "Disentangling neutrino oscillations". In: *Physics Letters B* 678.2 (2009), pp. 191–196.

[36] C. Jarlskog. "Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation". In: *Phys. Rev. Lett.* 55 (10 1985), pp. 1039–1042. DOI: `10.1103/PhysRevLett.55.1039`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.55.1039`.

[37] Ivan Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2020.9 (2020), pp. 1–22.

[38] R. L. Workman et al. "Review of Particle Physics". In: *PTEP* 2022 (2022), p. 083C01. DOI: `10.1093/ptep/ptac097`.

[39] AA Kochanov et al. "High-energy spectra of the atmospheric neutrinos: Predictions and measurements". In: *arXiv preprint arXiv:2109.13000* (2021).

[40] A Yu Smirnov. “The MSW effect and matter effects in neutrino oscillations”. In: *Physica Scripta* 2005.T121 (2005), p. 57.

[41] Pablo F De Salas et al. “Neutrino mass ordering from oscillations and beyond: 2018 status and future prospects”. In: *Frontiers in Astronomy and Space Sciences* 5 (2018), p. 36.

[42] DESI Collaboration et al. *DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations.* 2024. arXiv: `2404.03002 [astro-ph.CO]`. URL: `https://arxiv.org/abs/2404.03002`.

[43] Peter Minkowski. “$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?” In: *Phys. Lett. B* 67 (1977), pp. 421–428. DOI: `10.1016/0370-2693(77)90435-X`.

[44] Murray Gell-Mann, Pierre Ramond, and Richard Slansky. “Complex Spinors and Unified Theories”. In: *Conf. Proc. C* 790927 (1979), pp. 315–321. arXiv: `1306.4669 [hep-th]`.

[45] Tsutomu Yanagida. “Horizontal symmetry and masses of neutrinos”. In: *Progress of Theoretical Physics* 64.3 (1980), pp. 1103–1105.

[46] Vedran Brdar et al. “Type I seesaw mechanism as the common origin of neutrino mass, baryon asymmetry, and the electroweak scale”. In: *Phys. Rev. D* 100 (7 2019), p. 075029. DOI: `10.1103/PhysRevD.100.075029`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.100.075029`.

[47] Carlo Giunti and Marco Laveder. “Neutrino mixing”. In: (2003). arXiv: `hep-ph/0310238`.

[48] Rabindra N. Mohapatra and Goran Senjanovic. “Neutrino Masses and Mixings in Gauge Models with Spontaneous Parity Violation”. In: *Phys. Rev. D* 23 (1981), p. 165. DOI: `10.1103/PhysRevD.23.165`.

[49] Rabindra N Mohapatra and Alexei Y Smirnov. “Neutrino mass and new physics”. In: *Annu. Rev. Nucl. Part. Sci.* 56 (2006), pp. 569–628.

[50] Eugeni Akhmedov et al. “Dynamical left-right symmetry breaking”. In: *arXiv preprint hep-ph/9509255* (1995).

[51] SM Barr. “A different see-saw formula for neutrino masses”. In: *arXiv preprint hep-ph/0309152* (2003).

[52] Zackaria Chacko et al. “Predictive Dirac and Majorana Neutrino Mass Textures from $SU(6)$ Grand Unified Theories”. In: *Phys. Rev. D* 102.3 (2020), p. 035020. DOI: `10.1103/PhysRevD.102.035020`. arXiv: `2005.05413 [hep-ph]`.

[53] Rabindra N. Mohapatra and Goran Senjanovic. “Neutrino Masses and Mixings in Gauge Models with Spontaneous Parity Violation”. In: *Phys. Rev. D* 23 (1981), p. 165. DOI: `10.1103/PhysRevD.23.165`.

[54] Frank Wilczek and A. Zee. “Operator Analysis of Nucleon Decay”. In: *Phys. Rev. Lett.* 43 (1979), pp. 1571–1573. DOI: `10.1103/PhysRevLett.43.1571`.

[55] G. Arduini et al. "High Luminosity LHC: challenges and plans". In: *Journal of Instrumentation* 11.12 (2016), p. C12081. DOI: `10.1088/1748-0221/11/12/C12081`. URL: `https://dx.doi.org/10.1088/1748-0221/11/12/C12081`.

[56] R. Abbasi et al. "IceCube high-energy starting event sample: Description and flux characterization with 7.5 years of data". In: *Phys. Rev. D* 104 (2 2021), p. 022002. DOI: `10.1103/PhysRevD.104.022002`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.104.022002`.

[57] BW Sargent. "The maximum energy of the β-rays from uranium X and other bodies". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 139.839 (1933), pp. 659–673.

[58] The KATRIN collaboration et al. "The design, construction, and commissioning of the KATRIN experiment". In: *Journal of Instrumentation* 16.08 (2021), T08015. DOI: `10.1088/1748-0221/16/08/T08015`. URL: `https://dx.doi.org/10.1088/1748-0221/16/08/T08015`.

[59] M Aker et al. "Direct neutrino-mass measurement based on 259 days of KATRIN data". In: *arXiv preprint arXiv:2406.13516* (2024).

[60] V.E. Lewis. "Beta decay of tritium". In: *Nuclear Physics A* 151.1 (1970), pp. 120–128. ISSN: 0375-9474. DOI: `https://doi.org/10.1016/0375-9474(70)90972-3`. URL: `https://www.sciencedirect.com/science/article/pii/0375947470909723`.

[61] A Baroni et al. "Tritium β decay in chiral effective field theory". In: *Physical Review C* 94.2 (2016), p. 024003.

[62] Matteo Agostini et al. "Toward the discovery of matter creation with neutrinoless $\beta\beta$ decay". In: *Rev. Mod. Phys.* 95 (2 2023), p. 025002. DOI: `10.1103/RevModPhys.95.025002`. URL: `https://link.aps.org/doi/10.1103/RevModPhys.95.025002`.

[63] M. Goeppert-Mayer. "Double Beta-Disintegration". In: *Phys. Rev.* 48 (6 1935), pp. 512–516. DOI: `10.1103/PhysRev.48.512`. URL: `https://link.aps.org/doi/10.1103/PhysRev.48.512`.

[64] B. Märkisch et al. "Measurement of the Weak Axial-Vector Coupling Constant in the Decay of Free Neutrons Using a Pulsed Cold Neutron Beam". In: *Phys. Rev. Lett.* 122.24 (2019), p. 242501. DOI: `10.1103/PhysRevLett.122.242501`. arXiv: `1812.04666 [nucl-ex]`.

[65] C. Patrignani. "Review of Particle Physics". In: *Chinese Physics C* 40.10 (2016), p. 100001. DOI: `10.1088/1674-1137/40/10/100001`. URL: `https://dx.doi.org/10.1088/1674-1137/40/10/100001`.

[66] Jouni T Suhonen. "Value of the axial-vector coupling strength in β and $\beta\beta$ decays: A review". In: *Frontiers in Physics* 5 (2017), p. 55.

[67] "Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T". In: *Nature* 568.7753 (2019), pp. 532–535.

[68] George Gamow. “Mass defect curve and nuclear constitution”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 126.803 (1930), pp. 632–644.

[69] C. F. V. Weizsacker. “Zur Theorie der Kernmassen”. In: *Z. Phys.* 96 (1935), pp. 431–458. DOI: `10.1007/BF01337700`.

[70] Jeremy Stein Cushman. “A search for neutrinoless double-beta decay in tellurium-130 with CUORE”. PhD thesis. Yale University, 2018.

[71] Juan José Gómez-Cadenas et al. “The search for neutrinoless double-beta decay”. In: *La Rivista del Nuovo Cimento* (2024), pp. 1–74.

[72] Alexander Barabash. “Precise half-life values for two-neutrino double-β decay: 2020 review”. In: *Universe* 6.10 (2020), p. 159.

[73] A. Balysh et al. “Double Beta Decay of ^{48}Ca”. In: *Phys. Rev. Lett.* 77 (26 1996), pp. 5186–5189. DOI: `10.1103/PhysRevLett.77.5186`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.77.5186`.

[74] VB Brudanin et al. “Search for double beta decay of ^{48}Ca in the TGV experiment”. In: *Physics Letters B* 495.1-2 (2000), pp. 63–68.

[75] R Arnold et al. “Measurement of the double-beta decay half-life and search for the neutrinoless double-beta decay of Ca 48 with the NEMO-3 detector”. In: *Physical Review D* 93.11 (2016), p. 112008.

[76] Ch Dörr and HV Klapdor-Kleingrothaus. “New Monte-Carlo simulation of the HEIDELBERG-MOSCOW double beta decay experiment”. In: *Nucl. Instrum. Meth. A* 513.3 (2003), pp. 596–621.

[77] GERDA Collaboration et al. “Results on $\beta\beta$ decay with emission of two neutrinos or Majorons in ^{76}Ge from GERDA Phase I: GERDA Collaboration”. In: *The European Physical Journal C* 75 (2015), pp. 1–12.

[78] R Arnold et al. “Final results on ^{82}Se double beta decay to the ground state of ^{82}Kr from the NEMO-3 experiment”. In: *The European Physical Journal C* 78.10 (2018), pp. 1–15.

[79] O Azzolini et al. “Evidence of single state dominance in the two-neutrino double-β decay of Se-82 with CUPID-0”. In: *Physical Review Letters* 123.26 (2019), p. 262501.

[80] S. R. Elliott et al. “Double beta decay of ^{82}Se”. In: *Phys. Rev. C* 46 (4 1992), pp. 1535–1537. DOI: `10.1103/PhysRevC.46.1535`. URL: `https://link.aps.org/doi/10.1103/PhysRevC.46.1535`.

[81] R Arnold et al. “Double-β decay of ^{82}Se”. In: *Nuclear Physics A* 636.2 (1998), pp. 209–223.

[82] R Arnold et al. “Double beta decay of ^{96}Zr”. In: *Nuclear Physics A* 658.4 (1999), pp. 299–312.

[83] J Argyriades et al. “Measurement of the two neutrino double beta decay half-life of Zr-96 with the NEMO-3 detector”. In: *Nuclear Physics A* 847.3-4 (2010), pp. 168–179.

[84] R Arnold et al. “Detailed studies of ^{100}Mo two-neutrino double beta decay in NEMO-3”. In: *The European Physical Journal C* 79.5 (2019), pp. 1–11.

[85] E Armengaud et al. “Precise measurement of $2\nu\beta\beta$ decay of ^{100}Mo with the CUPID-Mo detection technology”. In: *The European Physical Journal C* 80.7 (2020), p. 674.

[86] D Dassie et al. “Two-neutrino double-β decay measurement of Mo-100”. In: *Physical Review D* 51.5 (1995), p. 2090.

[87] A De Silva et al. “Double β decays of ^{100}Mo and ^{150}Nd”. In: *Physical Review C* 56.5 (1997), p. 2451.

[88] Laura Cardani et al. “First bolometric measurement of the two neutrino double beta decay of ^{100}Mo with a $ZnMoO_4$ crystals array”. In: *Journal of Physics G: Nuclear and Particle Physics* 41.7 (2014), p. 075204.

[89] R Arnold et al. “Measurement of the $2\nu\beta\beta$ decay half-life and search for the $0\nu\beta\beta$ decay of ^{116}Cd with the NEMO-3 detector”. In: *Physical Review D* 95.1 (2017), p. 012007.

[90] AS Barabash et al. “Final results of the Aurora experiment to study 2β decay of ^{116}Cd with enriched $^{116}CdWO_4$ crystal scintillators”. In: *Physical Review D* 98.9 (2018), p. 092007.

[91] Hiroyasu Ejiri et al. “Double beta decays of ^{116}Cd”. In: *Journal of the Physical Society of Japan* 64.2 (1995), pp. 339–343.

[92] Fedor A Danevich et al. “Search for 2β decay of cadmium and tungsten isotopes: Final results of the Solotvina experiment”. In: *Physical Review C* 68.3 (2003), p. 035501.

[93] NEMO collaboration et al. “Double-β decay of ^{116}Cd”. In: *Zeitschrift für Physik C: Particles and Fields* 72 (1996), pp. 239–247.

[94] C Alduino et al. “Measurement of the two-neutrino double-beta decay half-life of ^{130}Te with the CUORE-0 experiment”. In: *The European Physical Journal C* 77 (2017), pp. 1–18.

[95] Irene Nutini et al. “The CUORE detector and results”. In: *Journal of Low Temperature Physics* 199.1 (2020), pp. 519–528.

[96] C. Arnaboldi et al. “A calorimetric search on double beta decay of ^{130}Te”. In: *Physics Letters B* 557.3 (2003), pp. 167–175. ISSN: 0370-2693. DOI: `https://doi.org/10.1016/S0370-2693(03)00212-0`. URL: `https://www.sciencedirect.com/science/article/pii/S0370269303002120`.

[97] Roger Arnold et al. “Measurement of the $\beta\beta$ decay half-life of ^{130}Te with the NEMO-3 detector”. In: *Physical Review Letters* 107.6 (2011), p. 062504.

[98] JB Albert et al. “Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector”. In: *Physical Review C* 89.1 (2014), p. 015502.

[99] Azusa Gando et al. “Search for Majorana neutrinos near the inverted mass hierarchy region with KamLAND-Zen”. In: *Physical review letters* 117.8 (2016), p. 082503.

[100] R Arnold et al. “Measurement of the $2\nu\beta\beta$ decay half-life of ^{150}Nd and a search for $0\nu\beta\beta$ decay processes with the full exposure from the NEMO-3 detector”. In: *Physical Review D* 94.7 (2016), p. 072003.

[101] W. H. Furry. “On Transition Probabilities in Double Beta-Disintegration”. In: *Phys. Rev.* 56 (12 1939), pp. 1184–1193. DOI: `10.1103/PhysRev.56.1184`. URL: `https://link.aps.org/doi/10.1103/PhysRev.56.1184`.

[102] Joseph Schechter and José WF Valle. “Neutrinoless double-β decay in SU(2)$\times$ U(1) theories”. In: *Physical Review D* 25.11 (1982), p. 2951.

[103] C. Giunti and E. M. Zavanin. “Predictions for neutrinoless double-beta decay in the 3+1 sterile neutrino scenario”. In: *Journal of High Energy Physics* 2015, 171 (2015), p. 171. DOI: `10.1007/JHEP07(2015)171`. arXiv: `1505.00978 [hep-ph]`.

[104] Matteo Agostini, Giovanni Benato, and Jason A. Detwiler. “Discovery probability of next-generation neutrinoless double-β decay experiments”. In: *Phys. Rev. D* 96 (5 2017), p. 053001. DOI: `10.1103/PhysRevD.96.053001`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.96.053001`.

[105] G. Anton et al. “Search for Neutrinoless Double-β Decay with the Complete EXO-200 Dataset”. In: *Phys. Rev. Lett.* 123 (16 2019), p. 161802. DOI: `10.1103/PhysRevLett.123.161802`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.123.161802`.

[106] G Adhikari et al. “nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity”. In: *Journal of Physics G: Nuclear and Particle Physics* 49.1 (2021), p. 015104. DOI: `10.1088/1361-6471/ac3631`. URL: `https://dx.doi.org/10.1088/1361-6471/ac3631`.

[107] Miryam Martínez-Vara. “NEXT, a neutrinoless double beta decay experiment”. In: *PoS* EPS-HEP2023 (2024), p. 169. DOI: `10.22323/1.449.0169`.

[108] S. Abe et al. “Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset”. In: (2024). arXiv: `2406.11438 [hep-ex]`.

[109] The SNO+ collaboration et al. “The SNO+ experiment”. In: *Journal of Instrumentation* 16.08 (2021), P08059. DOI: `10.1088/1748-0221/16/08/P08059`. URL: `https://dx.doi.org/10.1088/1748-0221/16/08/P08059`.

[110] D. Q. Adams et al. “Search for Majorana neutrinos exploiting millikelvin cryogenics with CUORE”. In: *Nature* 604.7904 (2022), pp. 53–58.

[111] E Armengaud et al. “The CUPID-Mo experiment for neutrinoless double-beta decay: performance and prospects”. In: *The European Physical Journal C* 80.1 (2020), pp. 1–15.

[112] A Agrawal et al. “Improved limit on neutrinoless double beta decay of ^{100}Mo from AMoRE-I”. In: *arXiv preprint arXiv:2407.05618* (2024).

[113] I. J. Arnquist et al. “Final Result of the Majorana Demonstrator’s Search for Neutrinoless Double-β Decay in ^{76}Ge”. In: *Phys. Rev. Lett.* 130 (6 2023), p. 062501. DOI: `10.1103/PhysRevLett.130.062501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.130.062501`.

[114] M. Agostini et al. “Final Results of GERDA on the Search for Neutrinoless Double-β Decay”. In: *Phys. Rev. Lett.* 125 (25 2020), p. 252502. DOI: `10.1103/PhysRevLett.125.252502`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.125.252502`.

[115] Anna Julia Zsigmond and for the LEGEND Collaboration. “LEGEND: The future of neutrinoless double-beta decay search with germanium detectors”. In: *Journal of Physics: Conference Series* 1468.1 (2020), p. 012111. DOI: `10.1088/1742-6596/1468/1/012111`. URL: `https://dx.doi.org/10.1088/1742-6596/1468/1/012111`.

[116] T Tomoda. “Double beta decay”. In: *Reports on Progress in Physics* 54.1 (1991), p. 53.

[117] Mark G. Inghram and John H. Reynolds. “Double Beta-Decay of ^{130}Te”. In: *Phys. Rev.* 78 (6 1950), pp. 822–823. DOI: `10.1103/PhysRev.78.822.2`. URL: `https://link.aps.org/doi/10.1103/PhysRev.78.822.2`.

[118] SR Elliott, AA Hahn, and MK Moe. “Direct evidence for two-neutrino double-beta decay in Se-82”. In: *Physical Review Letters* 59.18 (1987), p. 2020.

[119] AS Barabash. “Experiment double beta decay: Historical review of 75 years of research”. In: *Physics of Atomic Nuclei* 74 (2011), pp. 603–613.

[120] S Umehara et al. “Neutrino-less double-β decay of ^{48}Ca studied by CaF_2 (Eu) scintillators”. In: *Physical Review C—Nuclear Physics* 78.5 (2008), p. 058501.

[121] O. Azzolini et al. “Final Result on the Neutrinoless Double Beta Decay of ^{82}Se with CUPID-0”. In: *Phys. Rev. Lett.* 129 (11 2022), p. 111801. DOI: `10.1103/PhysRevLett.129.111801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.129.111801`.

[122] D. Q. Adams et al. “New Direct Limit on Neutrinoless Double Beta Decay Half-Life of ^{128}Te with CUORE”. In: *Phys. Rev. Lett.* 129 (22 2022), p. 222501. DOI: `10.1103/PhysRevLett.129.222501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.129.222501`.

[123] C Aidala et al. *A New Era Of Discovery: The 2023 Long-Range Plan For Nuclear Science.* Tech. rep. Lawrence Livermore National Laboratory (LLNL), Livermore, CA (United States), 2023.

[124] G. Bellini et al. "Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth". In: *Journal of Cosmology and Astroparticle Physics* 2012.05 (2012), p. 015. DOI: `10.1088/1475-7516/2012/05/015`. URL: `https://dx.doi.org/10.1088/1475-7516/2012/05/015`.

[125] FG Kondev et al. "The NUBASE2020 evaluation of nuclear physics properties". In: *Chinese Physics C* 45.3 (2021), p. 030001.

[126] A Alessandrello et al. "A new search for neutrinoless $\beta\beta$ decay with a thermal detector". In: *Physics Letters B* 335.3-4 (1994), pp. 519–525.

[127] A Alessandrello et al. "Improvements in ^{130}Te double beta decay search with cryogenic TeO$_2$ array detectors". In: *Nuclear Physics B-Proceedings Supplements* 48.1-3 (1996), pp. 238–240.

[128] A Alessandrello et al. "Preliminary results on double beta decay of ^{130}Te with an array of twenty cryogenic detectors". In: *Physics Letters B* 433.1-2 (1998), pp. 156–162.

[129] A. Alessandrello et al. "New experimental results on double beta decay of ^{130}Te". In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 486 (2000), pp. 13–21.

[130] C. Arnaboldi et al. "New Limit on the Neutrinoless $\beta\beta$ Decay of ^{130}Te". In: *Phys. Rev. Lett.* 95 (14 2005), p. 142501. DOI: `10.1103/PhysRevLett.95.142501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.95.142501`.

[131] L Canonica et al. "Results from the CUORE-0 experiment". In: *Journal of Physics: Conference Series* 718.6 (2016), p. 062007. DOI: `10.1088/1742-6596/718/6/062007`. URL: `https://dx.doi.org/10.1088/1742-6596/718/6/062007`.

[132] K. Alfonso et al. "Search for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE-0". In: *Phys. Rev. Lett.* 115 (10 2015), p. 102502. DOI: `10.1103/PhysRevLett.115.102502`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.115.102502`.

[133] Klaus Pretzl. "Cryogenic calorimeters in astro and particle physics". In: *Nucl. Instrum. Meth. A* 454.1 (2000), pp. 114–127.

[134] Marco Barucci et al. "Measurement of low temperature specific heat of crystalline TeO$_2$ for the optimization of bolometric detectors". In: *Journal of Low Temperature Physics* 123 (2001), pp. 303–314.

[135] EE Haller et al. "NTD germanium: a novel material for low temperature bolometers". In: *Neutron Transmutation Doping of Semiconductor Materials*. Springer, 1984, pp. 21–36.

[136] Ning Wang et al. "Electrical and thermal properties of neutron-transmutation-doped Ge at 20 mK". In: *Physical Review B* 41.6 (1990), p. 3761.

[137] VS Shirley. "Table of nuclides". In: (2008).

[138] EE Haller, KM Itoh, and JW Beeman. "Neutron transmutation doped (NTD) germanium thermistors for sub-mm bolometer applications". In: (1996).

[139] Boris Isaakovich Shklovskii and Alex L Efros. *Electronic properties of doped semiconductors*. Vol. 45. Springer Science & Business Media, 2013.

[140] C. Arnaboldi et al. "A front-end electronic system for large arrays of bolometers". In: *Journal of Instrumentation* 13.02 (2018), P02026. DOI: `10.1088/1748-0221/13/02/P02026`. URL: `https://dx.doi.org/10.1088/1748-0221/13/02/P02026`.

[141] Erica Andreotti et al. "Production, characterization, and selection of the heating elements for the response stabilization of the CUORE bolometers". In: *Nucl. Instrum. Meth. A* 664.1 (2012), pp. 161–170. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2011.10.065`.

[142] C Arnaboldi et al. "A programmable multichannel antialiasing filter for the CUORE experiment". In: *Nucl. Instrum. Meth. A* 617.1-3 (2010), pp. 327–328.

[143] Frank Pobell. *Matter and methods at low temperatures*. Vol. 2. Springer, 2007.

[144] A D'Addabbo et al. "The CUORE cryostat". In: *Journal of Low Temperature Physics* 193 (2018), pp. 867–875.

[145] Matthew I Hollister, Ram C Dhuley, and Grzegorz L Tatkowski. "A large millikelvin platform at Fermilab for quantum computing applications". In: *IOP Conference Series: Materials Science and Engineering*. Vol. 1241. 1. IOP Publishing. 2022, p. 012045.

[146] Jonathan Ouellet. "The coldest cubic meter in the known universe". In: *arXiv preprint arXiv:1410.1560* (2014).

[147] DR Artusa et al. "Searching for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE". In: *Advances in High Energy Physics* 2015.1 (2015), p. 879871.

[148] Carmine Pagliarone et al. "The CUORE Fast Cooling System". In: *PoS* EPS-HEP2017 (2018). Ed. by Paolo Checchia et al., p. 634. DOI: `10.22323/1.314.0634`.

[149] Cryomech. `https://www.cryomech.com/products/pt415/`.

[150] D.Q. Adams et al. "CUORE opens the door to tonne-scale cryogenics experiments". In: *Progress in Particle and Nuclear Physics* 122 (2022), p. 103902. ISSN: 0146-6410. DOI: `https://doi.org/10.1016/j.ppnp.2021.103902`. URL: `https://www.sciencedirect.com/science/article/pii/S0146641021000612`.

[151] Carlo Bucci et al. "The Faraday room of the CUORE experiment". In: *Journal of Instrumentation* 12 (2017). DOI: `10.1088/1748-0221/12/12/P12013`.

[152] S Di Domizio et al. "A data acquisition and control system for large mass bolometer arrays". In: *Journal of Instrumentation* 13.12 (2018), P12003.

[153] Giorgio Del Castello. "Development of energy calibration and data analysis systems for the NUCLEUS experiment". In: *arXiv preprint arXiv:2302.02843* (2023).

[154] Norbert Wiener. "Generalized harmonic analysis". In: *Acta mathematica* 55.1 (1930), pp. 117–258.

[155] Alexander Khintchine. “Korrelationstheorie der stationären stochastischen Prozesse”. In: *Mathematische Annalen* 109.1 (1934), pp. 604–615.

[156] Sunil Ramanlal Golwala. “Exclusion limits on the WIMP-nucleon elastic-scattering cross-section from the Cryogenic Dark Matter Search”. PhD thesis. University of California, Berkeley, 2000.

[157] Maurice S Bartlett. “Smoothing periodograms from time-series with continuous spectra”. In: *Nature* 161.4096 (1948), pp. 686–687.

[158] G. Turin. “An introduction to matched filters”. In: *IRE Transactions on Information Theory* 6.3 (1960), pp. 311–329. DOI: `10.1109/TIT.1960.1057571`.

[159] D.O. North. “An Analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems”. In: *Proceedings of the IEEE* 51.7 (1963), pp. 1016–1027. DOI: `10.1109/PROC.1963.2383`.

[160] E. Gatti and P. F. Manfredi. “Processing the Signals From Solid State Detectors in Elementary Particle Physics”. In: *Riv. Nuovo Cim.* 9N1 (1986), pp. 1–146. DOI: `10.1007/BF02822156`.

[161] A Branca et al. “Performance of the low threshold Optimum Trigger on CUORE data”. In: *Journal of Physics: Conference Series*. Vol. 1468. 1. IOP Publishing. 2020, p. 012118.

[162] S Di Domizio et al. “Lowering the energy threshold of large-mass bolometric detectors”. In: *Journal of Instrumentation* 6.02 (2011), P02007. DOI: `10.1088/1748-0221/6/02/P02007`. URL: `https://dx.doi.org/10.1088/1748-0221/6/02/P02007`.

[163] C. Alduino et al. “Analysis techniques for the evaluation of the neutrinoless double-β decay lifetime in ^{130}Te with the CUORE-0 detector”. In: *Phys. Rev. C* 93 (4 2016), p. 045503. DOI: `10.1103/PhysRevC.93.045503`. URL: `https://link.aps.org/doi/10.1103/PhysRevC.93.045503`.

[164] Roger Guo Huang. “Searching for $0\nu\beta\beta$ Decay with CUORE and CUPID”. PhD thesis. University of California, Berkeley, 2021.

[165] DQ Adams et al. “An energy-dependent electro-thermal response model of CUORE cryogenic calorimeter”. In: *Journal of Instrumentation* 17.11 (2022), P11023.

[166] J Zhang et al. “Non-Ohmic effects in hopping conduction in doped silicon and germanium between 0.05 and 1 K”. In: *Physical Review B* 57.8 (1998), p. 4472.

[167] Dan McCammon. “Semiconductor thermistors”. In: *Cryogenic Particle Detection* (2005), pp. 35–62.

[168] A. D’Addabbo et al. “An active noise cancellation technique for the CUORE Pulse Tube cryocoolers”. In: *Cryogenics* 93 (2018), pp. 56–65. ISSN: 0011-2275. DOI: `https://doi.org/10.1016/j.cryogenics.2018.05.001`. URL: `https://www.sciencedirect.com/science/article/pii/S0011227517304460`.

[169] L Aragão et al. “The environmental low-frequency background for macro-calorimeters at the millikelvin scale”. In: *arXiv preprint arXiv:2404.13602* (2024).

[170] Junaid Aasi et al. “Advanced LIGO”. In: *Classical and quantum gravity* 32.7 (2015), p. 074001.

[171] Maddalena Mantovani. “Virgo and the gravitational interferometry”. In: *EPJ Web Conf.* 280 (2023), p. 08005. DOI: `10.1051/epjconf/202328008005`.

[172] Fausto Acernese et al. “The Virgo O3 run and the impact of the environment”. In: *Classical and quantum gravity* 39.23 (2022), p. 235009.

[173] Benjamin P Abbott et al. “A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals”. In: *Classical and Quantum Gravity* 37.5 (2020), p. 055002.

[174] Beverly Berger and LIGO Scientific Collaboration Team. “Saga of mysterious 15 Hz ground motion at LIGO Livingston Observatory”. In: *APS April Meeting Abstracts*. Vol. 2022. APS Meeting Abstracts. 2022, S16.006.

[175] Sergio Zimmermann. “Active microphonic noise cancellation in radiation detectors”. In: *Nucl. Instrum. Meth. A* 729 (2013), pp. 404–409.

[176] Sergio Zimmermann et al. “Development of next-generation nuclear physics integrated readout electronics for GRETINA”. In: *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*. IEEE. 2013, pp. 1–7.

[177] IY Lee et al. “GRETINA: A gamma ray energy tracking array”. In: *Nuclear Physics A* 746 (2004), pp. 255–259.

[178] M.A. Deleplanque et al. “GRETA: utilizing new concepts in γ-ray detection”. In: *Nucl. Instrum. Meth. A* 430.2 (1999), pp. 292–310. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/S0168-9002(99)00187-4`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900299001874`.

[179] C Mancini-Terracciano and M Vignati. “Noise correlation and decorrelation in arrays of bolometric detectors”. In: *Journal of Instrumentation* 7.06 (2012), P06013. DOI: `10.1088/1748-0221/7/06/P06013`. URL: `https://dx.doi.org/10.1088/1748-0221/7/06/P06013`.

[180] Jonathan Ouellet. “The Search for Neutrinoless Double Beta Decay in ^{130}Te With CUORE-0”. PhD thesis. University of California, Berkeley, 2015.

[181] Behringer. `https://www.behringer.com/behringer/product?modelCode=0506-AAA`.

[182] PCB Piezotronics. `https://www.pcb.com/products?m=393b31`.

[183] Sara Electronic Instruments. `https://www.sara.pg.it/prodotti.php?id=60`.

[184] S Pirro et al. “Vibrational and thermal noise reduction for cryogenic detectors”. In: *Nucl. Instrum. Meth. A* 444.1 (2000), pp. 331–335. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/S0168-9002(99)01376-5`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900299013765`.

[185] C Alduino et al. “CUORE-0 detector: design, construction and operation”. In: *Journal of Instrumentation* 11.07 (2016), P07009.

[186] O. Azzolini et al. “Final Result on the Neutrinoless Double Beta Decay of ^{82}Se with CUPID-0”. In: *Phys. Rev. Lett.* 129 (11 2022), p. 111801. DOI: `10.1103/PhysRevLett.129.111801`.

[187] Han Beom Kim et al. “Status and performance of the AMoRE-I experiment on neutrinoless double beta decay”. In: *Journal of Low Temperature Physics* 209.5-6 (2022), pp. 962–970.

[188] M Carrettoni and M Vignati. “Signal and noise simulation of CUORE bolometric detectors”. In: *Journal of Instrumentation* 6.08 (2011), P08007. DOI: `10.1088/1748-0221/6/08/P08007`. URL: `https://dx.doi.org/10.1088/1748-0221/6/08/P08007`.

[189] Julius S Bendat and Allan G Piersol. *Random data: analysis and measurement procedures.* John Wiley & Sons, 2011.

[190] C. Arnaboldi et al. “A front-end electronic system for large arrays of bolometers”. In: *Journal of Instrumentation* 13.02 (2018), P02026. DOI: `10.1088/1748-0221/13/02/P02026`. URL: `https://dx.doi.org/10.1088/1748-0221/13/02/P02026`.

[191] PCB Piezotronics. `https://www.pcb.com/products?m=482c15`.

[192] V. Singh et al. “Large-area photon calorimeter with Ir-Pt bilayer transition-edge sensor for the CUPID experiment”. In: *Phys. Rev. Appl.* 20 (6 2023), p. 064017. DOI: `10.1103/PhysRevApplied.20.064017`. URL: `https://link.aps.org/doi/10.1103/PhysRevApplied.20.064017`.

[193] HJ Rice and JA Fitzpatrick. “A generalised technique for spectral analysis of nonlinear systems”. In: *Mechanical Systems and Signal Processing* 2.2 (1988), pp. 195–207.

[194] Julius S Bendat and Allan G Piersol. “Spectral analysis of non-linear systems involving square-law operations”. In: *Journal of Sound and Vibration* 81.2 (1982), pp. 199–214.

[195] Allen Caldwell, Daniel Kollár, and Kevin Kröninger. “BAT–The Bayesian analysis toolkit”. In: *Computer Physics Communications* 180.11 (2009), pp. 2197–2209.

[196] Wolfgang A Rolke and Angel M Lopez. “Confidence intervals and upper bounds for small signals in the presence of background noise”. In: *Nucl. Instrum. Meth. A* 458.3 (2001), pp. 745–758.

[197] Samuel S Wilks. “The large-sample distribution of the likelihood ratio for testing composite hypotheses”. In: *The annals of mathematical statistics* 9.1 (1938), pp. 60–62.

[198] Nuria Lopez Vaquero, Tomás R. Rodríguez, and J. Luis Egido. "Shape and pairing fluctuations effects on neutrinoless double beta decay nuclear matrix elements". In: *Phys. Rev. Lett.* 111.14 (2013), p. 142501.

[199] Juhani Hyvärinen and Jouni Suhonen. "Nuclear matrix elements for $0\nu\beta\beta$ decays with light or heavy Majorana-neutrino exchange". In: *Phys. Rev. C* 91.2 (2015), p. 024613.

[200] Mihai Horoi and Andrei Neacsu. "Shell model predictions for ^{124}Sn double-β decay". In: *Phys. Rev. C* 93 (2016), p. 024308.

[201] L. S. Song et al. "Nuclear matrix element of neutrinoless double-β decay: Relativity and short-range correlations". In: *Phys. Rev. C* 95 (2017), p. 024305.

[202] J. Menéndez. "Neutrinoless $\beta\beta$ decay mediated by the exchange of light and heavy neutrinos: the role of nuclear structure correlations". In: *Journal of Physics G: Nuclear and Particle Physics* 45.1 (2017), p. 014003.

[203] Dong-Liang Fang, Amand Faessler, and Fedor Šimkovic. "$0\nu\beta\beta$-decay nuclear matrix element for light and heavy neutrino mass mechanisms from deformed quasiparticle random-phase approximation calculations for ^{76}Ge, ^{82}Se, ^{130}Te, ^{136}Xe, and ^{150}Nd with isospin restoration". In: *Phys. Rev. C* 97 (2018), p. 045503.

[204] Fedor Šimkovic, Adam Smetana, and Petr Vogel. "$0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements evaluated in closure approximation, neutrino potentials and SU(4) symmetry". In: *Phys. Rev. C* 98 (2018), p. 064325.

[205] L. Coraggio et al. "Calculation of the neutrinoless double-β decay matrix element within the realistic shell model". In: *Phys. Rev. C* 101 (2020), p. 044315.

[206] Frank F. Deppisch et al. "Analysis of light neutrino exchange and short-range mechanisms in $0\nu\beta\beta$ decay". In: *Phys. Rev. D* 102 (2020), p. 095016.

[207] Lawrence Roberts. "Picture coding using pseudo-random noise". In: *IRE Transactions on Information Theory* 8.2 (1962), pp. 145–154.

[208] Paul A Souder et al. "Measurement of parity violation in the elastic scattering of polarized electrons from C 12". In: *Physical review letters* 65.6 (1990), p. 694.

[209] W. Emam. "A search for new physics at the TeV scale via a precise measurement of the weak mixing angle in Møller scattering". PhD thesis. Syracuse University, 2004.

[210] K Alfonso et al. "CUPID: The Next-Generation Neutrinoless Double Beta Decay Experiment". In: *Journal of Low Temperature Physics* (2022), pp. 1–9.

[211] L. Bergé et al. "Complete event-by-event $\alpha/\gamma(\beta)$ separation in a full-size TeO_2 CUORE bolometer by Neganov-Luke-magnified light detection". In: *Phys. Rev. C* 97 (3 2018), p. 032501. DOI: `10.1103/PhysRevC.97.032501`. URL: `https://link.aps.org/doi/10.1103/PhysRevC.97.032501`.

[212] L. Gironi et al. "Cerenkov light identification with Si low-temperature detectors with sensitivity enhanced by the Neganov-Luke effect". In: *Phys. Rev. C* 94 (5 2016), p. 054608. DOI: `10.1103/PhysRevC.94.054608`. URL: `https://link.aps.org/doi/10.1103/PhysRevC.94.054608`.

[213] K. Alfonso et al. "A first test of CUPID prototypal light detectors with NTD-Ge sensors in a pulse-tube cryostat". In: *Journal of Instrumentation* 18.06 (2023), P06033. DOI: `10.1088/1748-0221/18/06/P06033`. URL: `https://dx.doi.org/10.1088/1748-0221/18/06/P06033`.

[214] V. Novati et al. "Charge-to-heat transducers exploiting the Neganov-Trofimov-Luke effect for light detection in rare-event searches". In: *Nucl. Instrum. Meth. A* 940 (2019), pp. 320–327. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2019.06.044`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900219308897`.

[215] C Augier et al. "Final results on the $0\nu\beta\beta$ decay half-life limit of ^{100}Mo from the CUPID-Mo experiment". In: *The European Physical Journal C* 82.11 (2022), pp. 1–20.

[216] John R. de Laeter et al. "Atomic weights of the elements. Review 2000 (IUPAC Technical Report)". In: *Pure and Applied Chemistry* 75.6 (2003), pp. 683–800. DOI: `doi:10.1351/pac200375060683`. URL: `https://doi.org/10.1351/pac200375060683`.

[217] C. Augier et al. "Measurement of the $2\nu\beta\beta$ Decay Rate and Spectral Shape of ^{100}Mo from the CUPID-Mo Experiment". In: *Phys. Rev. Lett.* 131 (16 2023), p. 162501. DOI: `10.1103/PhysRevLett.131.162501`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.131.162501`.

[218] A. Armatol et al. "Novel technique for the study of pileup events in cryogenic bolometers". In: *Phys. Rev. C* 104 (1 2021), p. 015501. DOI: `10.1103/PhysRevC.104.015501`. URL: `https://link.aps.org/doi/10.1103/PhysRevC.104.015501`.