A Search for Neutrinoless Double Beta Decay of ^{130}Te

by

Adam Douglas Bryant

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 Stuart Freedman
Professor Eric Norman

Spring 2010

A Search for Neutrinoless Double Beta Decay of ^{130}Te

Abstract

A Search for Neutrinoless Double Beta Decay of ^{130}Te
by
Adam Douglas Bryant
Doctor of Philosophy in Physics
University of California, Berkeley
Professor Yury Kolomensky, Chair

This dissertation describes an experimental search for neutrinoless double beta ($0\nu\beta\beta$) decay of ^{130}Te. An observation of $0\nu\beta\beta$ decay would establish that neutrinos are Majorana fermions and would constrain the neutrino mass scale. The data analyzed were collected by two bolometric experiments: CUORICINO and an R&D experiment for CUORE known as the Three Towers Test. Both experiments utilized arrays of TeO_2 crystals operated as bolometers at ∼10 mK in a dilution refrigerator. The bolometers measured the energy deposited by particle interactions in the crystals by recording the induced change in crystal temperature. Between the two experiments, there were 81 TeO_2 bolometers used in the analysis, each of which was an independent detector of nuclear decays as well as a source of ^{130}Te. The experiments were conducted underground at a depth of about 3300 meters water equivalent in Hall A of the Laboratori Nazionali del Gran Sasso in Assergi, Italy, in order to shield the detectors from cosmic rays. The data analyzed represent an exposure of 19.9 $\mathrm{kg} \cdot \mathrm{y}$ of ^{130}Te (18.6 $\mathrm{kg} \cdot \mathrm{y}$ from CUORICINO and 1.3 $\mathrm{kg} \cdot \mathrm{y}$ from the Three Towers Test). In addition to the combined analysis of the two experiments, an analysis of CUORICINO data alone is presented in order to compare with an independent analysis being carried out by collaborators at the Univerity of Milano-Bicocca.

No signal due to $0\nu\beta\beta$ decay is observed, and therefore a limit on the partial half-life for the decay is set. From a simultaneous fit to the 81 independent detectors, the rate of $0\nu\beta\beta$ decay of ^{130}Te is measured to be $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-0.6 \pm 1.4\ (\mathrm{stat.}) \pm 0.4\ (\mathrm{syst.})) \times 10^{-25}\ \mathrm{y}^{-1}$, which corresponds to a lower limit on the partial half-life for $0\nu\beta\beta$ decay of ^{130}Te of $T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 3.0 \times 10^{24}$ y (90% C.L.). Converting the half-life limit to an upper limit on the effective Majorana neutrino mass, $m_{\beta\beta}$, using a set of recent nuclear matrix element calculations results in $m_{\beta\beta} < 0.25$–0.68 eV (90% C.L.), where the range reflects the spread of calculated nuclear matrix element values. These results disagree by at least 1.2σ, depending on the nuclear matrix element calculation, with a claim of observation of $0\nu\beta\beta$ decay of ^{76}Ge, assuming that the dominant mechanism driving $0\nu\beta\beta$ decay is the exchange of light Majorana neutrinos.

To my family

Contents

List of Figures

List of Tables

Acknowledgments

I am grateful to many people, professionally and personally, for their help, support, and encouragement during the completion of this work. The success of the experiments on which this dissertation is based is due to dozens of individuals in the CUORICINO and CUORE Collaborations. Likewise, my work as a graduate student has been made possible by numerous people.

In Berkeley, I first would like to thank my advisor, Professor Yury Kolomensky, for bringing me on board the CUORE project and for being an excellent mentor. His advice and suggestions are always insightful and tremendously helpful. I would like to thank the members of my thesis committee, Professor Stuart Freedman and Professor Rick Norman, for their careful attention to my analysis and manuscript and for their suggestions for improvements. Other CUORE collaborators currently or formerly in Berkeley have my appreciation for their suggestions and input; they include Tom Banks, Thomas Bloxham, Patrick Decowski, Michelle Dolinski, Elena Guardincerri, Ke Han, Koichi Ichimura, Richard Kadel, and Laura Kogler. I wish to express my thanks to all the members of the Weak Interactions Group at Berkeley for creating an enjoyable environment in which to work. I would like to acknowledge my current and former KamLAND counterparts in the neutrino graduate student office, Tommy O'Donnell and Lindley Winslow, for being examples of excellent graduate students. I am grateful to Nu Xu for his help with my GSRA appointment at LBNL. The members of the support staff in the Department of Physics at Berkeley, especially Donna Sakima and Anne Takizawa, have my appreciation for all that they do for graduate students.

In Italy, I wish to thank the CUORE collaborators at Gran Sasso, Iulian Bandac, Carlo Bucci, Paolo Gorla, Maurizio Perego, and Stefano Pirro, for welcoming me into their laboratories, their hospitality, and their guidance. I would like to acknowledge Fabio Bellini for his attention to my analysis and his suggestions and Marco Vignati for developing the Diana software framework and for setting an excellent example of how to program. I feel lucky to have been involved in the close collaboration between the U.S. and Italian groups on the data analysis, and I would like to make a special acknowledgement of the other members of the teams that performed the first-level data analysis and data quality checks: for CUORICINO – Fabio Bellini, Marco Carrettoni, Laura Kogler, Maria Martinez, and Claudia Tomei; for the Three Towers Test – Silvia Capelli and Marco Vignati. I would like to thank our physics analysis coordinator, Maura Pavan, for her advice and suggestions and for keeping the analysis teams on track.

On a personal level, I wish to thank my parents for their many years of support and encouragement. Finally, I would like to thank my wife, Leslie, for her love and for understanding my long overseas trips to the lab.

Chapter 1

Introduction

In 1937 Ettore Majorana pointed out that the neutrino introduced by Wolfgang Pauli seven years earlier could be its own antiparticle, which would make the neutrino fundamentally different from the other elementary particles of matter [1, 2]. Being electrically neutral particles, neutrinos and antineutrinos can only be distinguished by their lepton number[1], which is conserved in the Standard Model due to an accidental symmetry of the theory. Fermions with distinct particle and antiparticle states, including the quarks and charged leptons, are Dirac fermions, while fermions whose particle and antiparticle states coincide are Majorana fermions. The distinction is closely connected to the form of the fermion mass term in the Lagrangian of the theory (Sect. 1.2). If the fermion is massless, so that no mass term appears, there is no physical distinction between the Dirac and Majorana cases [3]. Until the late 1990s, experiments sensitive to the neutrino masses found the masses consistent with zero, and neutrinos are massless in the Standard Model.

Recently, experiments have provided compelling evidence for neutrino oscillations from one flavor to another in neutrino fluxes originating in the sun [4, 5], the atmosphere [6], nuclear reactors [7], and particle accelerators [8, 9]. The fact that neutrinos oscillate implies that neutrinos have mass, and prompted by this discovery, the question of whether neutrinos are Dirac or Majorana particles has become one of the issues at the forefront of particle and nuclear physics. To answer this question, physicists are undertaking sensitive new experiments to search for neutrinoless double beta ($0\nu\beta\beta$) decay, in which a nucleus (A, Z) decays to $(A, Z+2) + 2e^-$ with no (anti)neutrinos in the final state. So far unobserved, neutrinoless double beta decay is a possible decay mode of some nuclides if and only if neutrinos are Majorana particles [10]. Attempting to detect $0\nu\beta\beta$ decay is currently the only practical experimental technique to determine whether neu-

---

[1] If neutrinos were massless, their helicity could also distinguish neutrinos (which are left-handed) and antineutrinos (which are right-handed).

trinos are Dirac or Majorana particles. Furthermore, if neutrinos are Majorana particles, a measurement of the rate of $0\nu\beta\beta$ decay of a nuclide would provide a constraint on the absolute mass scale of neutrinos.

This dissertation describes one of the most sensitive experimental searches for $0\nu\beta\beta$ decay to date. The data analyzed are from two experiments based on cryogenic bolometers, CUORICINO and the Three Towers Test, performed by the CUORE Collaboration at the Laboratori Nazionali del Gran Sasso (LNGS) in Italy. These two experiments served as prototypes for the upcoming CUORE experiment as well as sensitive detectors in their own right for the neutrinoless double beta decay of ^{130}Te. This chapter discusses the theoretical framework for neutrinos within the Standard Model, the distinction between Dirac and Majorana neutrinos, and recent experimental results on neutrino masses.

1.1 The Standard Model and neutrino masses

The Standard Model of particle physics is a renormalizable gauge theory characterized by the gauge symmetry $\mathrm{SU}(3)_{\mathrm{c}} \otimes \mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$. The field content of the Standard Model includes three left-handed neutrinos: ν_{eL}, $\nu_{\mu L}$, and $\nu_{\tau L}$. A left-handed neutrino $\nu_{\ell L}$ ($\ell \in \{e, \mu, \tau\}$) and the corresponding left-handed charged lepton ℓ_L transform as a doublet under the $\mathrm{SU}(2)_L$ gauge symmetry. The right-handed charged leptons, ℓ_R, are singlets under $\mathrm{SU}(2)_L$, and there are no right-handed neutrinos in the Standard Model. The lepton content of the Standard Model may be summarized as

$$L_\ell \equiv \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix}, \quad \ell_R \qquad \ell \in \{e, \mu, \tau\}, \tag{1.1}$$

plus antiparticles.

The absence of right-handed neutrinos implies that neutrinos are massless in the Standard Model. Bare mass terms are not allowed in the Lagrangian because they are not gauge invariant. The charged leptons and quarks acquire masses via the Higgs mechanism. The Yukawa interaction $\lambda \Phi^T \overline{L_\ell} \ell_R$, where $\Phi \equiv \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$ is the Higgs doublet, is a gauge invariant term in the Lagrangian. When the electroweak symmetry is spontaneously broken and the Higgs acquires a vacuum expectation value $\langle \Phi \rangle = \begin{pmatrix} 0 \\ v \end{pmatrix}$, the Yukawa interaction generates a mass term for the charged leptons of the form $\lambda v \overline{\ell_L} \ell_R$. A mass term of this form, involving left-handed and right-handed fields, is called a Dirac mass term. Every massive fermion in the Standard Model acquires its mass from such a Dirac mass term. Neutrinos are massless in the theory because without a right-handed neutrino, the Yukawa interactions that would give rise to neutrino masses do not exist.

Recent experiments have shown that neutrinos are not massless. Observations of neutrino oscillations, discussed in Sect. 1.3, demonstrate that at least two neutrinos have nonzero masses. How can neutrino masses be incorporated into the theory? The straightforward tack is to add right-handed neutrinos to the particle content of the Standard Model by analogy with the charged leptons or quarks. Then, neutrino masses would be generated by the usual Higgs mechanism. However, the introduction of right-handed neutrinos allows for new terms in the Lagrangian: Since a right-handed neutrino is uncharged under all the gauge symmetries, a mass term $m\overline{(\nu_{\ell R})^c}\nu_{\ell R}$ is gauge invariant. This term should be included in the Lagrangian according to the rule of constructing the most general gauge invariant and renormalizable Lagrangian. This type of mass term, involving fields of the same chirality, is called a Majorana mass term. Majorana mass terms are uniquely possible for neutrinos because for charged fermions the Majorana mass terms are not invariant under the U(1) gauge symmetry of electromagnetism.

A consequence of the existence of a Majorana mass term is that lepton number is not conserved. Conservation of lepton number is not associated with a gauge symmetry, like the conservation of electric charge is for example. Lepton number is an accidental symmetry of the Standard Model; its conservation is a result of the field content and the requirement of renormalizability [11]. As such, conservation of lepton number is not viewed as inviolable like the conservation of electric charge, and it is broken in many models extending the Standard Model. Lepton number is the only quantum number that distinguishes neutrinos and antineutrinos. If lepton number is not conserved, there is nothing to distinguish neutrinos and antineutrinos, and neutrinos would be identical to antineutrinos. A fermion that is its own antiparticle is called a Majorana fermion.

1.2 Dirac and Majorana neutrinos

A free spin-1/2 fermion, whether Dirac or Majorana, may be represented by a four-component spinor field ψ, which satisfies the Dirac equation

$$(i\gamma^\mu \partial_\mu - m)\psi = 0. \tag{1.2}$$

The field may be decomposed into left-handed and right-handed chiral projections: $\psi = \psi_L + \psi_R$, where $\psi_L = \left(\frac{1-\gamma^5}{2}\right)\psi$ and $\psi_R = \left(\frac{1+\gamma^5}{2}\right)\psi$. For Dirac fermions, the left-handed and right-handed components are independent.

The definition of a Majorana fermion is that its field satisfies the constraint

$$\psi = \psi^c, \tag{1.3}$$

where $\psi^c \equiv C\overline{\psi}^T$ and C is the charge-conjugation matrix. This constraint is called the Majorana condition and means the particle and antiparticle quanta of the field are the same. The Majorana condition implies $\psi_R = (\psi_L)^c$, so the left-handed and right-handed components of a Majorana field are not independent [2]. Therefore, even though a Majorana particle may be represented by a four-component spinor, the additional constraint of the Majorana condition implies that a Majorana field is actually a two-component object.[2]

The Majorana condition does not need to be imposed by hand. The fields of neutrinos with definite mass may automatically satisfy the Majorana condition, depending on the form of the neutrino mass terms in the Lagrangian. Introducing vectors of left-handed and right-handed neutrino fields,

$$\nu_L = \begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix}, \quad \nu_R = \begin{pmatrix} \nu_{eR} \\ \nu_{\mu R} \\ \nu_{\tau R} \end{pmatrix}, \tag{1.4}$$

the most general set of neutrino mass terms is

$$\mathcal{L}^{D+M} = -\frac{1}{2}\overline{(\nu_L)^c} M_L^M \nu_L - \frac{1}{2}\overline{\nu_R} M_R^M (\nu_R)^c - \overline{\nu_R} M^D \nu_L + \text{h.c.} \tag{1.5}$$

where M_L^M, M_R^M, and M^D are complex matrices [13]. The right-handed neutrinos, ν_R, are iso-singlets and do not enter into the interaction terms in the Lagrangian. The term with M^D in Eq. (1.5) is a Dirac mass term and is the type of mass term generated by the usual Higgs mechanism. The terms with M_L^M and M_R^M are Majorana mass terms.[3] The mass terms in Eq. (1.5) can be written as

$$\mathcal{L}^{D+M} = -\frac{1}{2}\overline{(n_L)^c} M n_L + \text{h.c.}, \tag{1.6}$$

where $n_L = \begin{pmatrix} \nu_L \\ (\nu_R)^c \end{pmatrix}$ and the mass matrix, M, is symmetric. After the symmetric mass matrix is diagonalized, the mass term takes the form

$$\mathcal{L}^{D+M} = \sum_i m_i \overline{\nu_i} \nu_i, \tag{1.7}$$

where the fields of definite mass satisfy $\nu_i = (\nu_i)^c$ if M_L^M or M_R^M is nonzero, implying that the mass-eigenstate neutrinos are Majorana particles [13]. If $M_L^M = M_R^M = 0$, the fields of definite mass do not satisfy the Majorana condition, in which case the fermions are of Dirac type.

---

[2] It is possible and perhaps more natural to describe the theory of Majorana particles using a two-component formalism from the outset [12]. The advantages of the four-component formalism are symmetry in form with the Dirac fermions and convenience in writing the charged-current interaction.

[3] In order for the term with M_L^M to appear, the Higgs sector of the theory must contain more than the minimal two-component Higgs field [12].

A Majorana mass term violates any global U(1) symmetry $\psi \to e^{i\phi}\psi$, such as lepton number. Correspondingly, the Majorana condition precludes a Majorana particle from having any nonzero conserved charge. Hence, Majorana neutrinos are truly neutral particles. Since nonzero M_L^M or M_R^M implies Majorana neutrinos, Majorana neutrinos are in this sense a more general possibility than Dirac neutrinos. If neutrinos are Dirac particles, it would strongly indicate that there is some unknown symmetry restricting the form of the neutrino mass terms. The nature of neutrinos, Dirac or Majorana, and whether lepton number is conserved in nature are open experimental questions. Experiments aiming to answer these questions attempt to observe a lepton-number-violating process, with neutrinoless double beta decay being the most promising candidate.

1.3 Neutrino mixing

The three neutrinos that participate in weak interactions are linear combinations of the neutrinos of definite mass defined by

$$\nu_\ell = \sum_{i=1}^{n} U_{\ell i}\nu_i, \tag{1.8}$$

where U is a unitary matrix called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) lepton-mixing matrix. In Eq. (1.8) the mixing is expressed in terms of fields. In terms of states,

$$|\nu_\ell\rangle = \sum_{i=1}^{n} U_{\ell i}^{*}|\nu_i\rangle. \tag{1.9}$$

The number of neutrinos of definite mass may be more than three, but from the oscillation experiments, we know that there are only three mass eigenstates light enough to be produced in the experiments. There may be more neutrinos with very heavy masses, such as the GUT-scale neutrinos of seesaw models. The 3×3 submatrix of U connecting the weak-eigenstate neutrinos with the light mass eigenstates is approximately unitary and is written in the standard parameterization as

$$\begin{aligned} U = &\begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \\ &\times \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \end{aligned} \tag{1.10}$$

In this notation, c_{ij} and s_{ij} stand for $\cos\theta_{ij}$ and $\sin\theta_{ij}$, respectively. The parameterization in Eq. (1.10) contains three mixing angles, θ_{12}, θ_{13}, and θ_{23}; the CP-violating phase δ; and two Majorana phases, α_1 and α_2. The phases α_1 and α_2 are called Majorana phases because they have observable consequences only if neutrinos are Majorana particles.

Because the neutrinos that are produced and absorbed in weak interactions are mixtures of states with different masses, a neutrino produced in a particular flavor can be detected as a neutrino of a different flavor. This phenomenon of flavor change is known as neutrino oscillations. The first experimental hints of neutrino oscillations came from experiments detecting neutrinos from the sun, starting with Ray Davis's chlorine experiment at the Homestake mine [14]. The Davis experiment, which ran between 1970 and 1994, utilized the reaction $\nu_e + {}^{37}\mathrm{Cl} \rightarrow e^- + {}^{37}\mathrm{Ar}$ to measure the flux of neutrinos produced in the sun. It was designed to test models of the chain of nuclear reactions that power the sun. Rather than confirming the solar models by finding the expected flux of neutrinos, the Davis experiment measured a flux of neutrinos that was only about one-third of what was predicted. This deficit of solar neutrinos was confirmed by the GALLEX [15] and SAGE [16] experiments, which utilized the reaction $\nu_e + {}^{71}\mathrm{Ga} \rightarrow e^- + {}^{71}\mathrm{Ge}$, although the measured flux in these gallium experiments was about one-half, rather than one-third, of the expected flux. We know now that the difference in electron neutrino fluxes between the chlorine and gallium experiments is due to the lower energy threshold of the gallium experiments and the energy dependence of neutrino oscillations. The solar neutrino deficit was also observed by the Super-Kamiokande experiment via neutrino-electron scattering in a water Cherenkov detector [17].

The deficit of ν_e from the sun hinted that the ν_e were oscillating into other flavors, but conclusive evidence that neutrino oscillations were the correct solution to the solar neutrino deficit came from the Sudbury Neutrino Observatory (SNO) experiment [4, 5]. The SNO experiment was the first experiment that could detect the appearance of ν_μ and ν_τ, not just the disappearance of ν_e. The SNO detector was a water Cherenkov detector that used heavy water, D_2O, instead of ordinary water. In addition to elastic scattering and charged-current interactions of neutrinos, thanks to the use of heavy water, SNO was also sensitive to neutral current interactions, $\nu_\ell + d \rightarrow \nu_\ell + p + n$, which can happen for any neutrino flavor with equal cross sections. The neutral-current interactions allowed SNO to measure the total flux of solar neutrinos of all flavors, and the total flux agrees with the value predicted by solar models for ν_e production (Fig. 1.1). Together with the reduced flux of ν_e measured in charged-current interactions, the SNO neutral-current rate demonstrates that electron neutrinos produced in the sun oscillate into different flavors on their way to the earth.

Neutrino-oscillation experiments have usually been analyzed in terms of the

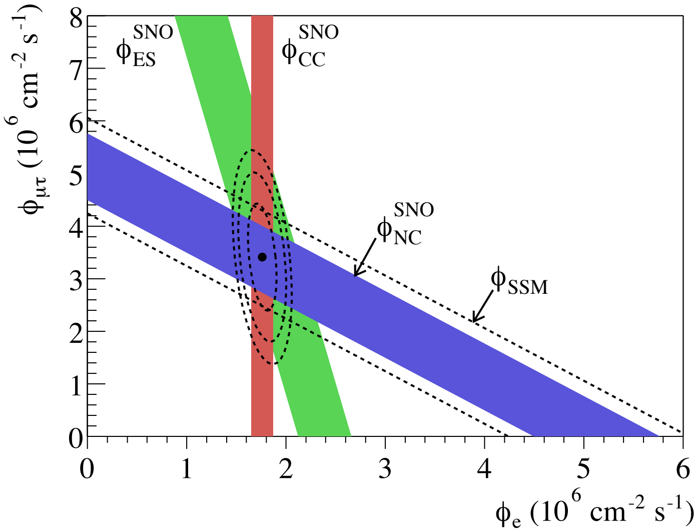

Figure 1.1: Neutrino fluxes measured by SNO in the elastic scattering (ES), charged current (CC), and neutral current (NC) channels. The horizontal axis represents the flux of electron neutrinos, and the vertical axis represents the flux of muon and tau neutrinos combined. The neutral current channel is equally sensitive to all neutrino flavors. Therefore, the flux measured in the NC channel, $\phi_{\rm NC}^{\rm SNO}$, could be due to fluxes of electron neutrinos and muon/tau neutrinos represented by any point in the purple band with slope -1. The charged current channel is sensitive only to electron neutrinos, so the allowed region from the $\phi_{\rm CC}^{\rm SNO}$ measurement is the red vertical band. The elastic scattering channel is sensitive to all neutrino flavors but has a larger cross section for electron neutrinos, resulting in an allowed region indicated by the green band for the $\phi_{\rm ES}^{\rm SNO}$ measurement. The intersection of the three bands determines the composition of the solar neutrino flux at the earth, proving that electron neutrinos produced in the sun oscillate into other neutrino flavors. The SNO measurement of the total flux agrees well with the neutrino flux predicted by the standard solar model indicated by the dashed lines. Figure from Ref. [4].

mixing of two neutrino flavors. In the two-neutrino-mixing case, two parameters describe the oscillations: a mixing angle θ and the difference in squared masses $\Delta m^2_{12} \equiv m^2_2 - m^2_1$. The survival probability of an electron neutrino is given by

$$P(\nu_e \to \nu_e) = 1 - \sin^2 2\theta \sin^2\left(1.27 \frac{\Delta m^2_{12}}{\mathrm{eV}^2} \frac{L}{\mathrm{km}} \frac{\mathrm{GeV}}{E}\right), \tag{1.11}$$

where L is the distance between the neutrino source and the detector and E is the neutrino energy [18]. Two-neutrino mixing is a good approximation for existing experiments because the value of θ_{13} in Eq. (1.10) turns out to be small and the two mass-squared differences measured in solar and atmospheric neutrino oscillations are very different. Therefore, the oscillation parameters are described in terms of the solar mixing angle, solar mass-squared difference, atmospheric mixing angle, and atmospheric mass-squared difference. These are identified with the parameters of the full three-neutrino oscillation picture (Eq. (1.10)) as $\theta_{\mathrm{solar}} = \theta_{12}$, $\Delta m^2_{\mathrm{solar}} = \Delta m^2_{12}$, $\theta_{\mathrm{atm}} = \theta_{23}$, and $\Delta m^2_{\mathrm{atm}} = \Delta m^2_{13}$.

The values of the solar mixing angle and mass-squared difference have been measured by combining SNO data with data from the KamLAND reactor antineutrino experiment. The KamLAND experiment has measured the flux of electron antineutrinos from nuclear reactors [7]. The KamLAND liquid scintillator detector in the Kamioka mine in central Japan is located at an average distance of ~180 km from many commercial power reactors. Electron antineutrinos produced in the beta decay of fission products in the reactors are detected when they induce inverse beta decay, $\overline{\nu}_e + p \to e^+ + n$, in the KamLAND detector. The delayed coincidence between the prompt signal from the positron and the delayed signal from the capture of the neutron is used to separate antineutrino events from backgrounds. Oscillations measured by KamLAND are characterized by the same mixing angle, θ_{12}, and mass-squared difference, Δm^2_{12}, as the solar neutrino experiments. The measured values of these parameters, obtained from a global analysis of all the experiments, are listed in Table 1.1. KamLAND has also produced the most direct evidence for neutrino oscillations by measuring the electron antineutrino survival probability as a function of L/E, shown in Fig. 1.2.

Neutrino oscillations have also been observed in atmospheric neutrinos, first by the Super-Kamiokande experiment. Atmospheric neutrinos are produced when cosmic rays bombard the air in the upper atmosphere. The collisions mainly produce pions, which produce neutrinos when they decay. The Super-Kamiokande detector is a water Cherenkov detector in the Kamioka mine that detects atmospheric neutrinos via charged-current interactions with nucleons, $\nu_\ell + N \to \ell + N'$. Because the earth is spherical, the flux of upward-going and downward-going neutrinos at the detector site would be symmetric if no oscillations occurred. Super-Kamiokande found a clear asymmetry in upward-going and downward-going neutrinos [6]. The oscillation effect was confirmed by the accelerator-based, long

Table 1.1: Measured values of neutrino mixing parameters from the global three-flavor analysis of Ref. [19] (September 2007 update).

Parameter	Best fit	2σ confidence interval
Δm^2_{12}	7.6×10^{-5} eV^2	$(7.3–8.1) \times 10^{-5}$ eV^2
$\lvert\Delta m^2_{13}\rvert$	2.4×10^{-3} eV^2	$(2.1–2.7) \times 10^{-3}$ eV^2
$\sin^2\theta_{12}$	0.32	0.28–0.37
$\sin^2\theta_{23}$	0.50	0.38–0.63
$\sin^2\theta_{13}$	0.007	≤ 0.033

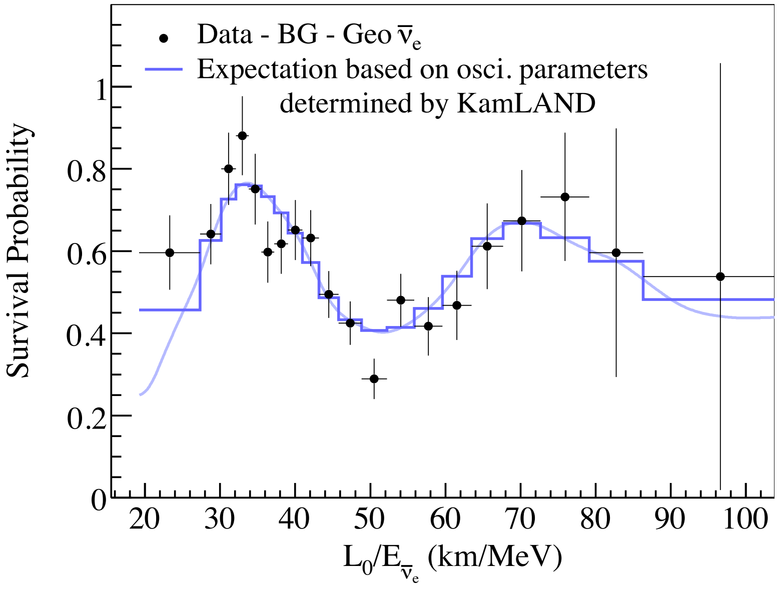

Figure 1.2: Neutrino oscillations observed by KamLAND. The survival probability of electron antineutrinos measured by KamLAND as a function of L/E agrees very well with the expectation based on neutrino oscillations. Figure from Ref. [7]

.

baseline experiments K2K [8] and MINOS [9], and the measured values for the oscillation parameters, θ_{23} and Δm^2_{13}, are listed in Table 1.1.

The limit on the third mixing angle, θ_{13}, in Table 1.1 comes primarily from the CHOOZ reactor antineutrino experiment in France [20]. The CHOOZ liquid scintillator detector was located about 1 km from two nuclear reactors. At this distance, the experiment was sensitive to oscillations at the frequency determined by Δm^2_{13} that are modulated by the size of $\sin^2(2\theta_{13})$. The data were consistent with no oscillations, and therefore an upper limit on θ_{13} was obtained.

Although the differences in neutrino mass-squared have been measured, oscillation experiments cannot measure the individual neutrino masses. Therefore, the absolute mass scale (the mass of the lightest neutrino) remains unknown. The absolute neutrino masses are constrained by the Mainz and Troitsk tritium-beta-decay experiments to be less than 2.5 eV (95% C.L.) [21, 22]. A stronger, but model-dependent, bound of $\sum m_\nu < 0.67$ eV (95% C.L.) has been obtained from cosmological evidence [23]. Moreover, the sign of Δm^2_{13} is unknown, leading to two possible hierarchies of neutrino masses, normal ($\Delta m^2_{13} > 0$) and inverted ($\Delta m^2_{13} < 0$). If the mass of the lightest neutrino is much greater than $\sqrt{|\Delta m^2_{13}|}$, the pattern of neutrino masses is said to be degenerate. The situation is illustrated in Fig. 1.3. If neutrinos are Majorana fermions, a measurement of the rate of neutrinoless double beta decay of a nuclide, combined with results from single beta decay experiments and cosmology, can determine the absolute neutrino mass scale and hierarchy (see Chapter 2).

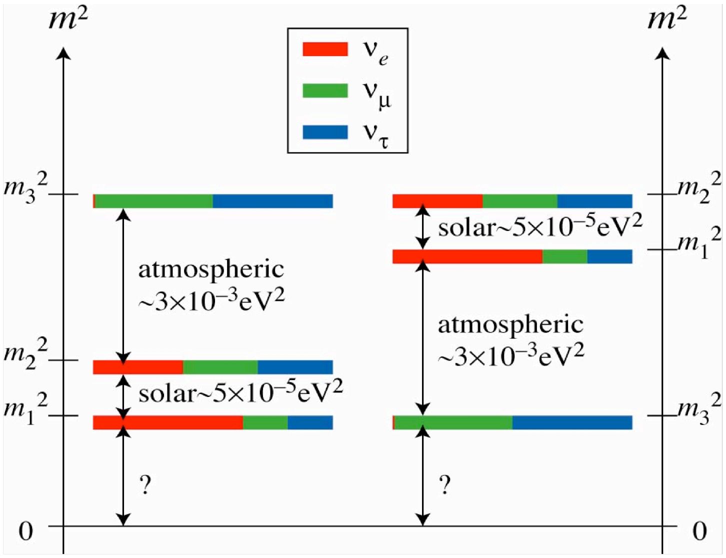

Figure 1.3: Two orderings of the neutrino masses are consistent with current data, depending on the sign of Δm^2_{13}. In the normal hierarchy (left side), the small mass splitting observed in solar neutrino oscillations is at the bottom, and in the inverted hierarchy (right side), the small mass splitting is at the top. The question marks indicate that the mass of the lightest neutrino is unknown.

Chapter 2

Double beta decay

Double beta decay is an extremely rare nuclear transition from a nucleus (A, Z) to its isobar $(A, Z+2)$ with the emission of two electrons. The transition may occur via a Standard Model allowed process in which two electron antineutrinos are emitted along with the electrons: $(A, Z) \rightarrow (A, Z+2) + 2e^- + 2\overline{\nu}_e$. This decay mode, known as two-neutrino double beta $(2\nu\beta\beta)$ decay, can be thought of as two simultaneous beta decays. The expected rate for $2\nu\beta\beta$ decay was first calculated by Goeppert-Mayer in 1935 [24], and the decay has now been observed in ten nuclides. Double beta decay half-lives for nuclei that undergo the process are very long, on the order of 10^{19}–10^{24} y, since the decay is second order in the weak interaction. The measured half-lives of double beta decaying isotopes are listed in Table 2.1 based on Ref. [25].

In principle, a nucleus (A, Z) can decay by double beta decay as long as the nucleus $(A, Z+2)$ is lighter. However, if the nucleus can also decay by single beta decay to $(A, Z+1)$, the branching fraction for double beta decay will be so small that it is practically impossible to observe the double beta decays in an experiment due to the overwhelming background rate from single beta decays. Therefore, candidate nuclei for experimental detection of double beta decay are even-even nuclei that, due to the nuclear pairing force, are lighter than the odd-odd $(A, Z+1)$ nucleus, making single beta decay kinematically forbidden. This situation is shown schematically in Fig. 2.1 and shown concretely for the $A = 130$ isobars in Fig. 2.2. The lightest, and therefore the only stable, $A = 130$ nuclide is ^{130}Xe; ^{130}Te is stable against single beta decay since ^{130}I is heavier, so it can only decay by double beta decay to ^{130}Xe. Two additional double beta decay candidates are ^{48}Ca and ^{96}Zr, for which single beta decay is kinematically possible but is greatly suppressed by a large difference in nuclear spin.

If neutrinos are Majorana fermions, there is an additional double beta decay mode, $(A, Z) \rightarrow (A, Z+2) + 2e^-$, with no (anti)neutrinos in the final state. This decay mode is known as neutrinoless double beta $(0\nu\beta\beta)$ decay and has never been

Table 2.1: Recommended half-life values for double beta decaying isotopes from Ref. [25].

Nuclide	$T_{1/2}^{2\nu\beta\beta}$ (y)
^{48}Ca	$4.4^{+0.6}_{-0.5} \times 10^{19}$
^{76}Ge	$(1.5 \pm 0.1) \times 10^{21}$
^{82}Se	$(9.2 \pm 0.7) \times 10^{19}$
^{96}Zr	$(2.3 \pm 0.2) \times 10^{19}$
^{100}Mo	$(7.1 \pm 0.4) \times 10^{18}$
^{116}Cd	$(2.8 \pm 0.2) \times 10^{19}$
^{128}Te	$(1.9 \pm 0.4) \times 10^{24}$
^{130}Te	$6.8^{+1.2}_{-1.1} \times 10^{20}$
^{150}Nd	$(8.2 \pm 0.9) \times 10^{18}$
^{238}U	$(2.0 \pm 0.6) \times 10^{21}$

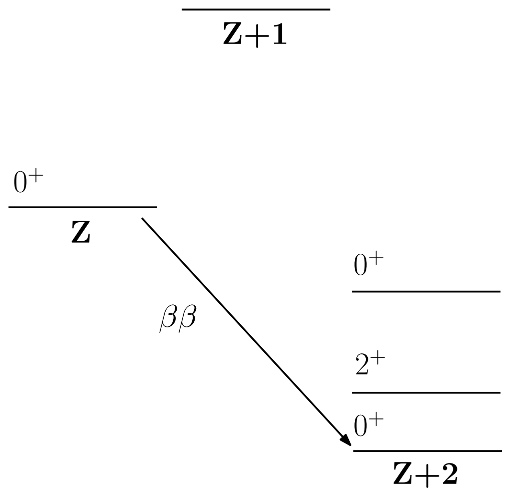

Figure 2.1: Schematic diagram of energy levels involved in double beta decay. Figure from Ref. [26].

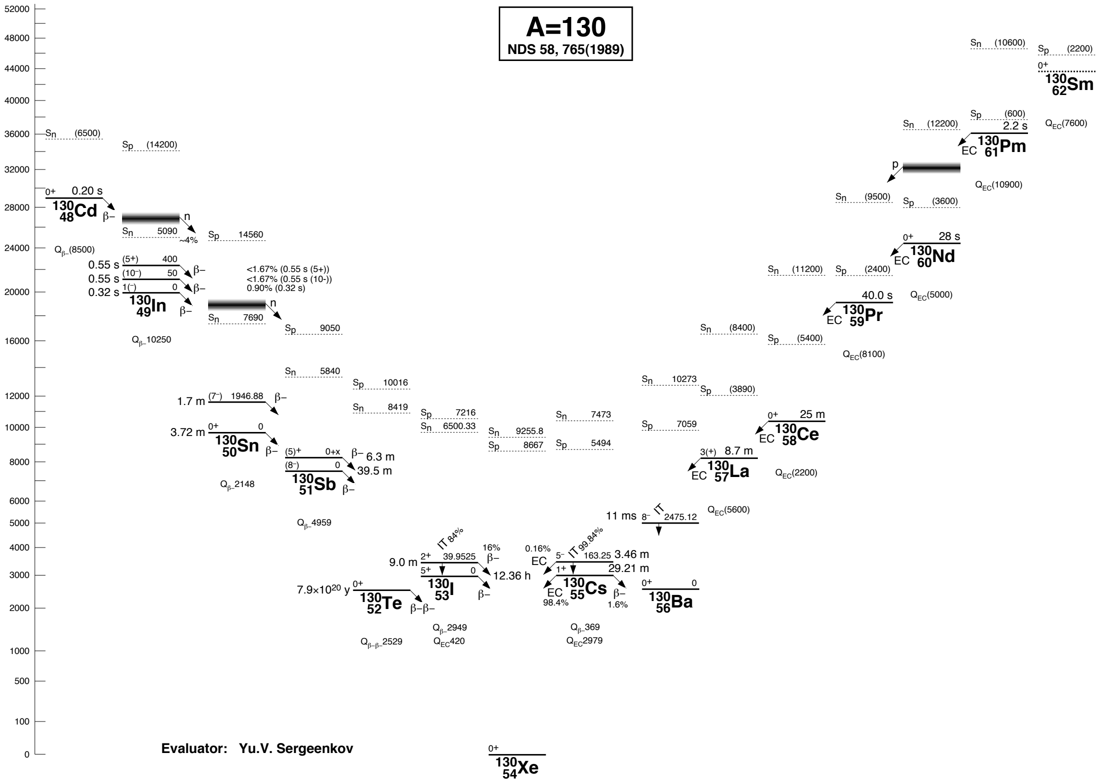

Figure 2.2: Masses of $A = 130$ isobars relative to ^{130}Xe. The values on the vertical axis are $(m(^{A}_{Z}\mathrm{X}) - m(^{130}\mathrm{Xe}))c^2$ in keV. Figure from the Table of Radioactive Isotopes [27].

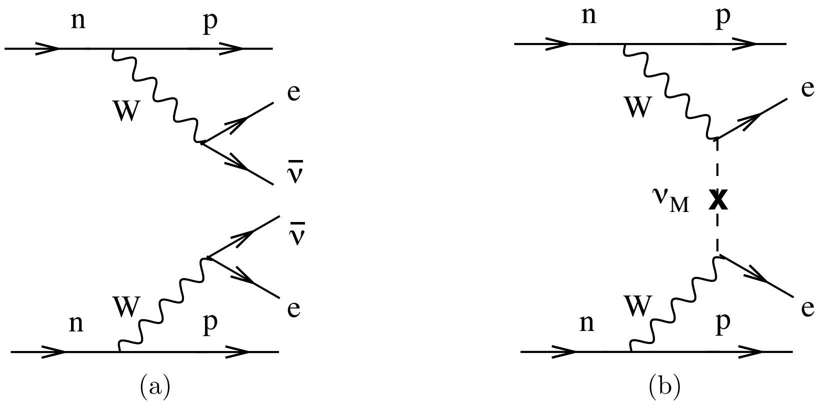

(a) (b)

Figure 2.3: Feynman diagrams for (a) two-neutrino double beta decay and (b) neutrinoless double beta decay mediated by light neutrino exchange. Figure from Ref. [26].

observed (except for one controversial claim discussed in Sect. 2.3.1). Neutrinoless double beta decay was first considered by Raccah in 1937 [28] soon after Majorana put forward his symmetric theory of particles and antiparticles. The first calculations of the rate for $0\nu\beta\beta$ decay, performed by Furry [29], yielded a much faster rate than for $2\nu\beta\beta$ decay, which prompted initial interest in experimental detection of $0\nu\beta\beta$ decay. However, at that time the chiral nature of the weak interaction was not yet known so a severe suppression of the $0\nu\beta\beta$ rate, discussed below, was not incorporated in the calculations.

Neutrinoless double beta decay is forbidden in the Standard Model since it manifestly breaks lepton number conservation (and $B - L$). Of course lepton number conservation is broken anyway if neutrinos are Majorana particles. Feynman diagrams for $2\nu\beta\beta$ and $0\nu\beta\beta$ decay are shown in Fig. 2.3. The $2\nu\beta\beta$ diagram contains only Standard Model interactions. The $0\nu\beta\beta$ diagram requires only the known $V - A$ interactions in addition to a massive Majorana neutrino. One can think of the virtual neutrino in the diagram as being produced as an antineutrino (equal to a neutrino since it is Majorana) at one vertex and absorbed as a neutrino at the other vertex. In addition to the Majorana equivalence of neutrino and antineutrino, a nonzero neutrino mass is required to flip the helicity since antineutrinos are right-handed and neutrinos are left-handed. The helicity flip and the smallness of the neutrino mass cause the rate of $0\nu\beta\beta$ decay, if it occurs at all, to be much lower than the rate of $2\nu\beta\beta$ decay.

The rate of $0\nu\beta\beta$ decay driven by the exchange of light Majorana neutrinos is given to a good approximation by [26]

$$\frac{1}{T_{1/2}^{0\nu\beta\beta}} = G_{0\nu}(Q,Z)|M^{0\nu}|^2\, m_{\beta\beta}^2, \tag{2.1}$$

where $G_{0\nu}(Q,Z)$ is a phase space factor proportional to Q^5, $M^{0\nu}$ is a nuclear matrix element, and $m_{\beta\beta}$ is the effective Majorana mass defined as

$$m_{\beta\beta} \equiv \left| \sum_{i=1}^{3} U_{ei}^2\, m_i \right| . \tag{2.2}$$

The particle physics information is contained in $m_{\beta\beta}$. The phase space factor, $G_{0\nu}(Q,Z)$, is calculable. Calculation of the nuclear matrix element, $M_{0\nu}$, is a challenging problem in nuclear theory (discussed further in Sect. 2.1). Experiments attempt to measure $T_{1/2}^{0\nu\beta\beta}$, and in the absence of a signal, they set a lower limit. Combining the measurement and the calculations, the value of $m_{\beta\beta}$ is deduced or an upper limit is set.

If neutrinos are Majorana particles, measuring or constraining the effective Majorana mass provides information on the neutrino mass scale and hierarchy. This is possible because there is a relationship between the effective Majorana mass and the mass of the lightest neutrino. The relationship depends on whether the hierarchy is normal or inverted because which neutrino mass eigenstate is the lightest depends on which hierarchy is realized in nature, as indicated in Fig. 1.3. For the normal hierarchy, m_1 is the lightest and therefore

$$\begin{aligned} m_{\beta\beta} &\equiv |U_{e1}^2 m_1 + U_{e2}^2 m_2 + U_{e3}^3 m_3| \\ &= \left| U_{e1}^2 m_1 + U_{e2}^2 \sqrt{\Delta m_{12}^2 + m_1^2} + U_{e3}^2 \sqrt{|\Delta m_{13}^2| + m_1^2} \right| \\ &= \Big| \cos^2\theta_{12}\cos^2\theta_{13} e^{i\alpha_1} m_1 + \sin^2\theta_{12}\cos^2\theta_{13} e^{i\alpha_2}\sqrt{\Delta m_{12}^2 + m_1^2} \\ &\quad + \sin^2\theta_{13} e^{-2i\delta}\sqrt{|\Delta m_{13}^2| + m_1^2} \Big| . \end{aligned} \tag{2.3}$$

A similar expression is easily derived for the inverted hierarchy in which m_3 is the lightest mass eigenvalue. Plugging in the measured values of the neutrino mixing angles and mass-squared differences from Table 1.1, a value for $m_{\beta\beta}$ is obtained for each value of the lightest neutrino mass, m_1 for the normal hierarchy or m_3 for the inverted hierarchy, and for a given set of values for the phases. Figure 2.4 shows the range of allowed values for $m_{\beta\beta}$ for each value of the lightest neutrino mass, obtained by allowing the unknown phases to vary over their possible values from 0 to 2π. There are distinct bands of allowed $m_{\beta\beta}$ depending on the hierarchy, though the bands overlap in the quasi-degenerate mass regime. Neutrinoless

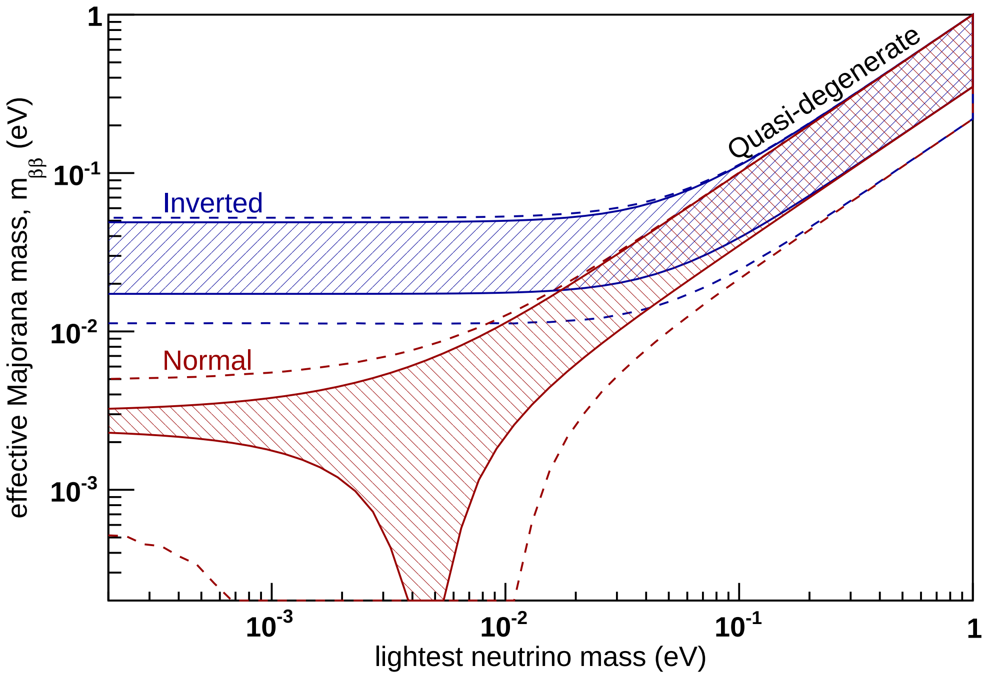

Figure 2.4: Allowed values of the effective Majorana mass as a function of the lightest neutrino mass based on the neutrino mixing parameters in Table 1.1. The width of the hatched regions is due to the unknown phases. The dashed lines show how the allowed regions expand if the neutrino mixing parameters are allowed to vary within their 2σ experimental errors.

double beta decay experiments set upper limits on $m_{\beta\beta}$ and therefore exclude a region from the top of Fig. 2.4. In this way, $0\nu\beta\beta$ decay experiments can rule out the quasi-degenerate mass regime under the assumption that neutrinos are Majorana particles. Future $0\nu\beta\beta$ decay experiments may have the sensitivity to rule out the inverted hierarchy.

Direct counting experiments search for double beta decay by measuring the sum of the electron energies and, in some experiments, the energy of the nuclear recoil. In a double beta decay, the total decay energy, or Q-value, is shared among all of the final state particles. The amount of the decay energy that goes into the kinetic energy of the daughter nucleus is negligible since the nucleus is so much heavier than the electrons (and the antineutrinos in the $2\nu\beta\beta$ case). In the spectrum of summed electron energies, $2\nu\beta\beta$ decay appears as a continuum from 0 up to the Q-value of the decay since the decay energy is shared among the electrons and the antineutrinos. In the $0\nu\beta\beta$ case, the electrons carry away all of the decay energy except for the negligible kinetic energy of the nuclear

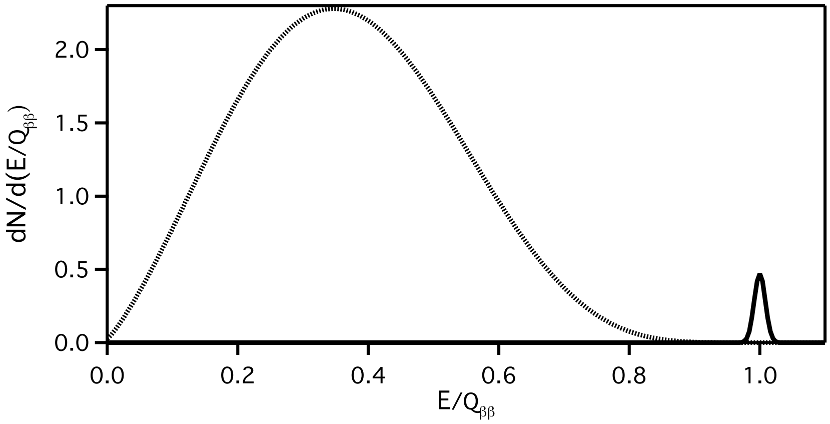

Figure 2.5: Summed electron energy spectral shapes for $2\nu\beta\beta$ decay (continuum) and $0\nu\beta\beta$ decay (peak). Figure from Ref. [26].

recoil. Therefore, $0\nu\beta\beta$ decay produces a mono-energetic peak at the Q-value in the summed electron energy spectrum (Fig. 2.5). The $0\nu\beta\beta$ peak will be spread only by the energy resolution of the detector. Experiments attempt to discover $0\nu\beta\beta$ decay by searching for this peak at the Q-value of the candidate nucleus. Bolometric experiments, such as CUORE, measure the nuclear recoil energy in addition to the electron energies; since the nuclear recoil energy is negligibly small, the spectral signatures of double beta decay in bolometric detectors are the same as for experiments that measure only the electron energies.

2.1 Nuclear matrix elements

The most significant source of uncertainty in drawing conclusions about neutrino masses from an experimental value of $T_{1/2}^{0\nu\beta\beta}$ for a particular isotope comes from theoretical uncertainty in the nuclear matrix element, $M^{0\nu}$ in Eq. (2.1). The calculation of nuclear matrix elements for double beta decay is a notoriously difficult problem in nuclear theory involving complicated operators acting on structurally complex nuclei. It requires a detailed description of the second-order weak transition from the parent to daughter nucleus through virtual states of an intermediate nucleus. The nuclear matrix element for $0\nu\beta\beta$ decay is a sum of Fermi and Gamow-Teller contributions:

$$M^{0\nu} = M_{\mathrm{GT}}^{0\nu} - \frac{M_{\mathrm{F}}^{0\nu}}{g_A^2}, \tag{2.4}$$

where g_A is the effective axial coupling in nuclear matter, not necessarily equal to the free nucleon value $g_A = 1.25$ [30]. The Fermi and Gamow-Teller matrix elements connect the 0^+ ground states in the initial and final nuclei as follows:

$$M_{\mathrm{F}}^{0\nu} = \langle 0_f^+ || \sum_{n,m} h(r) t_{n-} t_{m-} || 0_i^+ \rangle \tag{2.5}$$

$$M_{\mathrm{GT}}^{0\nu} = \langle 0_f^+ || \sum_{n,m} h(r) (\vec{\sigma}_n \cdot \vec{\sigma}_m) t_{n-} t_{m-} || 0_i^+ \rangle. \tag{2.6}$$

The sum in the matrix elements is over nucleons. The neutrino propagator arising from the presence of a virtual neutrino in the Feynman diagram leads to the neutrino potential, $h(r)$.

Approaches to evaluating $M^{0\nu}$ have generally fallen into two categories, the quasiparticle random phase approximation (QRPA) and the interacting nuclear shell model (NSM). The two kinds of models take complementary approaches in the approximations they make. The QRPA calculations [30] include a large number of single particle states but allow only a limited set of configurations of the states. On the other hand, NSM calculations [31] take into account all possible nucleon configurations but using a very restricted single particle basis. Because of the factorial growth in computational complexity of NSM calculations as the size of the nucleus under consideration increases and because double beta decay occurs in medium to heavy nuclei, most calculations of $0\nu\beta\beta$ matrix elements have employed the QRPA technique or a refinement of it.

There has been significant progress in the calculations of $M^{0\nu}$ in recent years. Calculations with QRPA performed by different groups have approximately converged. Rodin *et al.* have shown that if the particle-particle interaction strength, g_{pp}, is tuned so that measured $2\nu\beta\beta$ half-lives are correctly obtained, then almost all modern QRPA calculations of $M^{0\nu}$ agree to within about 30% [32, 33]. Figure 2.6 shows the spread of nuclear matrix element calculations for several $0\nu\beta\beta$ decay candidates. The two QRPA calculations shown in the figure agree well with each other and with an interacting boson model (IBM) calculation, while a recent NSM calculation is roughly a factor of two lower.

2.2 $0\nu\beta\beta$ decay implies Majorana neutrinos

The exchange of a light Majorana neutrino, the diagram in Fig. 2.3b, is the minimal mechanism for $0\nu\beta\beta$ decay in that it requires no new particles or interactions beyond neutrinos being Majorana particles. It is not the only possible mechanism for $0\nu\beta\beta$ decay, however. Grand unified theories of particle physics beyond the Standard Model present several mechanisms that could drive $0\nu\beta\beta$ decay (see Ref. [34] and references therein). Mechanisms involving lepton number violating quark-lepton interactions can drive $0\nu\beta\beta$ decay in both R-parity conserving and R-parity non-conserving supersymmetric extensions of the Standard

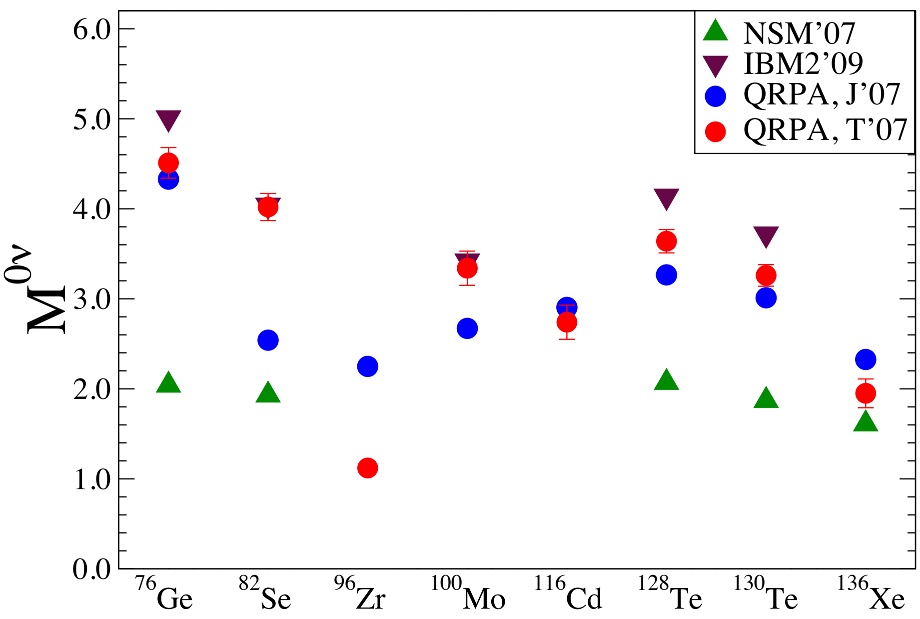

Figure 2.6: Nuclear matrix elements calculated by four different groups. Figure from Ref. [30].

Model. There is also a possible mechanism based on leptoquark exchange. If the symmetry responsible for lepton number conservation is spontaneously broken, there exists a Nambu-Goldstone boson associated with the spontaneous symmetry breaking, known as the Majoron. The Majoron, ϕ, couples to neutrinos, and its existence would lead to another possible mode of double beta decay: $(A, Z) \rightarrow (A, Z + 2) + 2e^- + \phi$. This neutrinoless double beta decay mode with a Majoron is easily distinguished from neutrinoless double beta decay without a Majoron by the spectrum of summed electron energies, which is continuous in the case that the decay energy is shared with the Majoron and mono-energetic in the standard case.

Experiments attempt to measure the partial half-life, $T_{1/2}^{0\nu\beta\beta}$, for $0\nu\beta\beta$ decay for some nuclide. Such a measurement is model-independent. Making a determination of the absolute neutrino mass scale and hierarchy, however, requires the assumption that the dominant mechanism driving $0\nu\beta\beta$ decay is the exchange of a light Majorana neutrino. Only under this assumption can Eq. (2.1) be used to turn a value of $T_{1/2}^{0\nu\beta\beta}$ into a value for $m_{\beta\beta}$. Nevertheless, Schechter and Valle [10] have shown that if $0\nu\beta\beta$ decay occurs in nature, neutrinos must be Majorana fermions, even if the $0\nu\beta\beta$ decay is dominantly driven by a mechanism other than light Majorana neutrino exchange. They constructed the diagram in Fig. 2.7, in

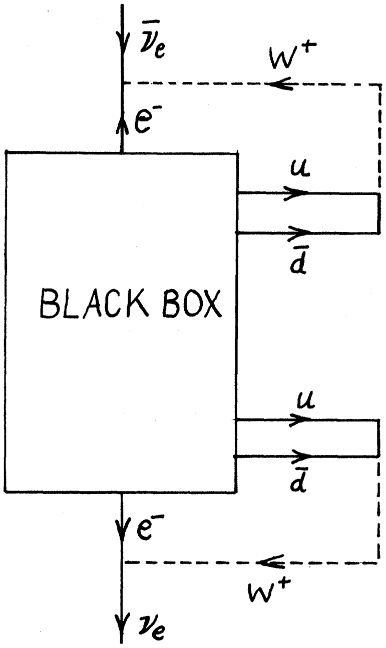

Figure 2.7: If $0\nu\beta\beta$ decay occurs, regardless of the dominant mechanism, this diagram (from Schechter and Valle [10]) modifies the neutrino propagator, inducing a Majorana mass term for neutrinos.

which whatever mechanism is responsible for $0\nu\beta\beta$ decay goes inside the "black box." By connecting the external lines appropriately, a contribution to the neutrino propagator is obtained that turns an anti-neutrino into a neutrino thereby inducing a Majorana mass term for the neutrino. Thus, an experimental observation of $0\nu\beta\beta$ decay would provide model-independent proof that neutrinos are Majorana fermions. To determine the mechanism driving $0\nu\beta\beta$ decay, measurements of $T_{1/2}^{0\nu\beta\beta}$ for multiple nuclides and well as accurate nuclear matrix elements will be necessary.

2.3 Status of experimental searches for $0\nu\beta\beta$ decay

The experimental signature for $0\nu\beta\beta$ decay is a mono-energetic peak, spread only by the detector resolution, in the spectrum of summed electron energies located at the Q-value of the nucleus undergoing the decay. The two principal requirements of an experiment designed to observe this peak are a very low background level and the capability to measure electron energies with high resolution.

A low background level is required to make a statistically significant discovery of a $0\nu\beta\beta$ decay signal, which is expected to be very small based on limits set by previous experiments. A high energy resolution prevents the tail of the $2\nu\beta\beta$ decay continuum spectrum from burying a small $0\nu\beta\beta$ decay peak, and an improvement in energy resolution directly reduces the background level by narrowing the region of interest in which the $0\nu\beta\beta$ peak is expected.

Several candidate isotopes are available for neutrinoless double beta decay searches. The choice of isotope used in an experiment is guided by the candidate isotope's Q-value, nuclear matrix element, and natural isotopic abundance, as well as the detector technology available to study the isotope. These properties for several isotopes of experimental interest are listed in Table 2.2. An ideal $0\nu\beta\beta$ decay candidate would feature a high Q-value, large nuclear matrix element, and large natural isotopic abundance. High Q-values are important so that the $0\nu\beta\beta$ decay peak occurs at a higher energy than natural gamma backgrounds, which extend up to 2615 keV. Higher Q-values also increase the expected number of $0\nu\beta\beta$ decay events because the phase space factor scales as Q^5. Higher nuclear matrix elements likewise increase the expected number of signal events. The natural isotopic abundance of a candidate isotope greatly affects the cost of an experiment. In order to field an experiment with a large source mass while keeping the total size of the apparatus within practical limits, the isotope of interest must usually comprise a significant fraction of the total mass, which requires isotopic enrichment for most candidates. The notable exception is ^{130}Te, the candidate isotope chosen for CUORICINO and CUORE. At 33.8% the natural isotopic abundance of ^{130}Te is far greater than that of the other $0\nu\beta\beta$ decay candidates with large Q-values. This feature of ^{130}Te allows CUORICINO and CUORE to achieve a high sensitivity per dollar by using natural, unenriched tellurium.

The most sensitive current limits on $0\nu\beta\beta$ decay have been set by the Heidelberg-Moscow and IGEX experiments using ^{76}Ge and by the CUORICINO experiment using ^{130}Te. These experiments utilized approaches in which "source = detector" to reduce backgrounds from natural radioactivity of detector components. "Source = detector" means that the double beta decaying isotope functions as an active part of the detector and makes up a large fraction of the total detector mass. The CUORICINO experiment, on which this dissertation is focused, used the bolometric technique to operate TeO_2 crystals as high resolution calorimeters (Chapter 3). The ^{76}Ge experiments built high-purity germanium (HPGe) detectors out of germanium enriched in ^{76}Ge.

The Heidelberg-Moscow Collaboration operated five HPGe detectors enriched to 86% in ^{76}Ge with a total source mass of 10.96 kg. Located in the Gran Sasso underground laboratory, the experiment ran from 1990 to 2003 and collected 71.7 $\text{kg} \cdot \text{y}$ exposure of ^{76}Ge [35]. It achieved a background level of 0.11 counts/$(\text{keV} \cdot \text{kg} \cdot \text{y})$ in the energy region around the ^{76}Ge Q-value of 2039 keV. In 1999 the Heidelberg-

Table 2.2: Properties of candidate isotopes for $0\nu\beta\beta$ decay. The Q-values and natural isotopic abundances, η, are from the Table of Radioactive Isotopes [27]. The last column is the expected partial half-life for $0\nu\beta\beta$ decay if $m_{\beta\beta} = 50$ meV, taken from Ref. [33]; it provides a measure of the nuclear matrix element combined with the phase space factor.

Nuclide	$Q_{\beta\beta}$ (keV)	η (%)	$T_{1/2}^{0\nu\beta\beta}(m_{\beta\beta} = 50\text{ meV})$ (y)
^{48}Ca	4276	0.2	
^{76}Ge	2039	7.4	8.6×10^{26}
^{82}Se	2992	8.7	2.4×10^{26}
^{96}Zr	3351	2.8	9.8×10^{26}
^{100}Mo	3034	9.6	2.4×10^{26}
^{116}Cd	2902	7.5	2.9×10^{26}
^{128}Te	867	31.7	4.5×10^{27}
^{130}Te	2527	33.8	2.2×10^{26}
^{136}Xe	2467	8.9	4.6×10^{26}
^{150}Nd	3368	5.6	2.2×10^{25}

Moscow Collaboration produced a limit of $T_{1/2}^{0\nu\beta\beta}(^{76}\text{Ge}) > 5.7 \times 10^{25}$ y (90% C.L.) [36]. Later a subset of the collaboration claimed discovery of $0\nu\beta\beta$ with the Heidelberg-Moscow data. This claim of discovery is discussed below in Sect. 2.3.1.

The International Germanium Experiment (IGEX) was similar to Heidelberg-Moscow in that it used 86% isotopically enriched HPGe detectors to search for $0\nu\beta\beta$ decay of ^{76}Ge. The collaboration originally operated detectors at three different underground laboratories: the Homestake gold mine in the United States, the Canfranc Tunnel in Spain, and the Baksan Neutrino Observatory in Russia. Then, three 2 kg detectors were operated at Canfranc [37]. The resulting limit on $0\nu\beta\beta$ decay was $T_{1/2}^{0\nu\beta\beta}(^{76}\text{Ge}) > 1.6 \times 10^{25}$ y (90% C.L.) [38].

Results from experimental searches for $0\nu\beta\beta$ decay in several other isotopes are listed in Table 2.3. The NEMO-3 experiment is notable because it has the capability to track charged particles and perform particle type identification. In the NEMO-3 detector, which is not a "source = detector" design, the double beta decaying source is contained in thin foils. Particles emitted from the foils are tracked inside a wire ionization chamber and their energies are measured by plastic scintillator calorimeters [45]. This type of detector provides a richer set of information than the "source = detector" experiments, including the number and type of particles in the decay, independent measurements of the energy of each final state particle, and angular distributions of decay products. This information

Table 2.3: Selection of current $0\nu\beta\beta$ decay half-life limits.

Nuclide	$T_{1/2}^{0\nu\beta\beta}$ (y)	Experiment
^{48}Ca	$> 1.4 \times 10^{22}$ (90% C.L.)	ELEGANT VI [39]
^{76}Ge	$> 5.7 \times 10^{25}$ (90% C.L.)	Heidelberg-Moscow [36]
^{82}Se	$> 1.0 \times 10^{23}$ (90% C.L.)	NEMO-3 [40]
^{96}Zr	$> 9.2 \times 10^{21}$ (90% C.L.)	NEMO-3 [41]
^{100}Mo	$> 4.6 \times 10^{23}$ (90% C.L.)	NEMO-3 [40]
^{116}Cd	$> 1.7 \times 10^{23}$ (90% C.L.)	Solotvina [42]
^{130}Te	$> 3.0 \times 10^{24}$ (90% C.L.)	CUORICINO + TTT [this work]
^{136}Xe	$> 1.2 \times 10^{24}$ (90% C.L.)	DAMA [43]
^{150}Nd	$> 1.8 \times 10^{22}$ (90% C.L.)	NEMO-3 [44]

greatly helps to disentangle signal events from backgrounds. It also allows for searches for $0\nu\beta\beta$ decay in many isotopes by changing the source foil. Thanks to its background suppression capabilities and module source design, NEMO-3 has made precision measurements of the $2\nu\beta\beta$ half-lives of several isotopes [40, 41, 44]. However, this approach is limited to relatively small source masses and has limited energy resolution.

2.3.1 Claim of discovery

In 2001 a subset of the Heidelberg-Moscow Collaboration, led by Klapdor-Kleingrothaus, announced 3.1σ evidence for $0\nu\beta\beta$ decay of ^{76}Ge in the Heidelberg-Moscow experiment [46]. Because this result was not endorsed by the full collaboration and because of several technical criticisms of the analysis [47, 48, 49], the claim of discovery has not been fully accepted by the community. Its proponents reasserted their claim with more data, more sophisticated analyses, and results with greater statistical significance [50, 51, 35]. In particular, they developed a pulse shape discrimination capability to distinguish between single-site events and multi-site events in their germanium detectors. Double beta decay events deposit the decay energy within a small, localized region and are therefore single-site events, whereas gamma backgrounds typically Compton scatter in the detector at least once before being fully absorbed and are therefore usually multi-site events. Pulse shape discrimination is in principle a powerful tool for separating a double beta decay signal from background in a germanium detector.

Application of pulse shape discrimination to the Heidelberg-Moscow data produced an extremely low background spectrum that contains a peak at the Q-value

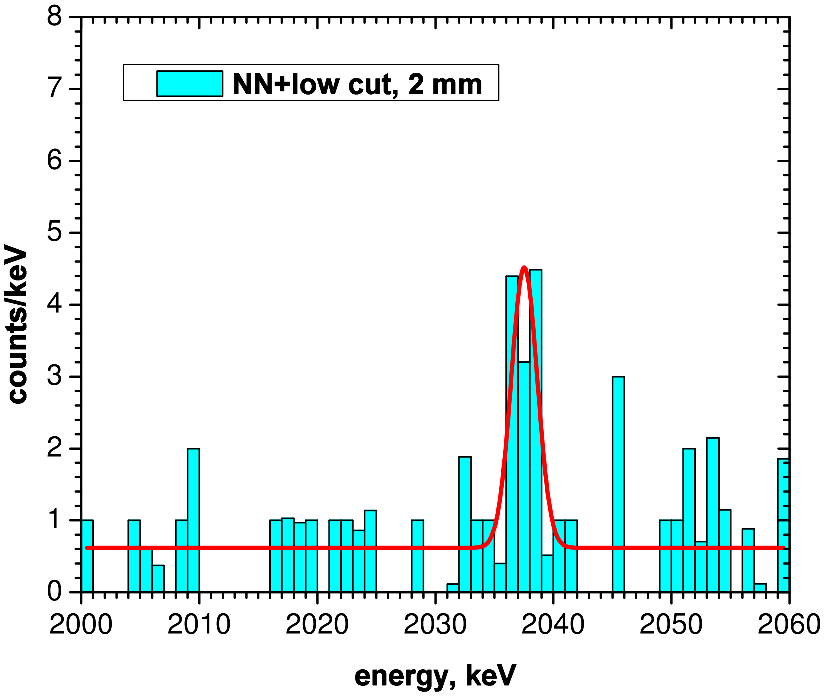

Figure 2.8: The Heidelberg-Moscow spectrum of single site events near the Q-value of ^{76}Ge after application of pulse shape discrimination. Figure from Ref. [35].

of ^{76}Ge with a statistical significance of greater than 6σ (Fig. 2.8) [35]. Klapdor-Kleingrothaus *et al.* obtain a measured value for the partial half-life for $0\nu\beta\beta$ decay of ^{76}Ge of $T_{1/2}^{0\nu\beta\beta}(^{76}\mathrm{Ge}) = (2.23^{+0.44}_{-0.31}) \times 10^{25}$ y, corresponding to an effective Majorana neutrino mass (assuming the light Majorana neutrino exchange mechanism dominates) of $m_{\beta\beta} = 0.32 \pm 0.03$ eV [35]. The community has still not found the analysis conclusive due in part to questions of validation of the pulse shape discrimination method. The method appears however to produce the correct shape for the $2\nu\beta\beta$ spectrum, as shown in Fig. 2.9. It is clear that the issue will only be settled by confirmation or refutation of the claim with data from future next generation experiments.

2.3.2 Next generation experiments

In the next few years, several next generation $0\nu\beta\beta$ decay experiments will start running. These experiments aim for the sensitivity to reach an effective Majorana mass of ~50 meV in order to cover the quasi-degenerate neutrino mass regime and perhaps to begin to probe the inverted hierarchy (see Fig. 2.4). To reach this goal within a few years of running will require hundreds of kilograms

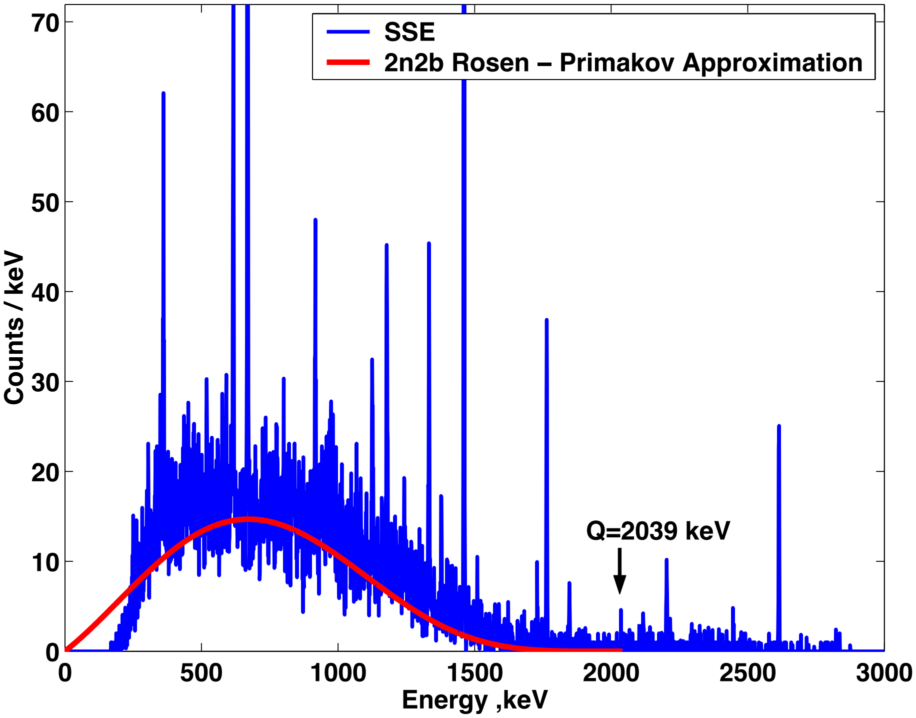

Figure 2.9: The Heidelberg-Moscow spectrum of single site events after application of pulse shape discrimination. The red curve represents shape of the $2\nu\beta\beta$ spectrum in the Primakov-Rosen approximation. Figure from Ref. [50].

of source mass. Table 2.2 indicates the half-life to which an experiment must be sensitive to reach $m_{\beta\beta} = 50$ meV.

The GERmanium Detector Array (GERDA) experiment will search for $0\nu\beta\beta$ decay of ^{76}Ge and is currently being commissioned at Gran Sasso [52, 53]. In its first phase, the enriched HPGe detectors from Heidelberg-Moscow and IGEX are being redeployed, providing ∼18 kg of ^{76}Ge. GERDA employs a novel scheme to cool the HPGe detectors: The bare semiconductor detectors are submerged in liquid argon. In addition to keeping the detectors cool, the liquid argon serves as a shield against external radioactivity. In Phase I the background level is expected to be as low as 10^{-2} counts/(keV·kg·y). An exposure of about 30 kg·y of ^{76}Ge will be collected in Phase I in order to test the Klapdor-Kleingrothaus *et al.* claim of observation. For the second phase, new segmented detectors are being developed that will bring the total source mass up to about 40 kg. Backgrounds causing multi-site energy depositions will be suppressed with pulse shape analysis. The possibility of reading out the liquid argon scintillation light is being investigated in order to turn the liquid argon bath into an active shield. The background level in Phase II is expected to be 10^{-3} counts/(keV·kg·y). The GERDA Collaboration plans to collect a total exposure of at least 100 kg · y.

The MAJORANA experiment also seeks to observe $0\nu\beta\beta$ decay of ^{76}Ge using HPGe detectors [54, 55]. It will be located in the Sanford Underground Laboratory, the first stage of the new Deep Underground Science and Engineering Laboratory (DUSEL) in the Homestake mine in South Dakota. The MAJORANA project is taking a phased approach. The first phase, termed the MAJORANA DEMONSTRATOR, will consist of a 60 kg module of high-purity Ge, of which 30 kg will be enriched to 86% in ^{76}Ge. The goals of the first phase are to demonstrate a background level at or below 1 count/(ton · y) in a 4 keV wide region of interest and to test definitively the claim of observation of $0\nu\beta\beta$ decay of ^{76}Ge. The MAJORANA Collaboration is engaged in R&D on detector design, especially p-type, point-contact Ge detectors, and on signal processing for identifying multi-site interactions in segmented Ge detectors. The collaboration eventually aims to build a ton-scale experiment and may join with the GERDA Collaboration to undertake the larger experiment.

The Enriched Xenon Observatory (EXO) experiment will use a Xe time projection chamber to search for $0\nu\beta\beta$ decay of ^{136}Xe [56]. The first phase, EXO-200, consists of a 200 kg detector currently being set up in laboratory space in the Waste Isolation Pilot Plant (WIPP), which is an underground facility in New Mexico for the long-term storage of nuclear waste from military applications. The EXO Collaboration eventually plans to construct a one-ton detector. An R&D effort is underway to develop a technique to tag double beta decay events by atomic spectroscopy of the daughter ^{136}Ba^{++} ions. The Ba tagging would provide the capability to reject all backgrounds except for $2\nu\beta\beta$ decay events.

The CUORE experiment, currently under construction at Gran Sasso, will search for $0\nu\beta\beta$ decay of ^{130}Te. It is a larger and more radiopure version of the CUORICINO detector containing 204 kg of ^{130}Te. Data taking is expected to begin in 2012. The CUORE detector is described in detail in Sect. 3.2.3.

Several other next generation $0\nu\beta\beta$ decay experiments have been proposed or are in the R&D phase. SuperNEMO is a proposed upgrade of NEMO-3 with at least 100 kg of source mass of ^{48}Ca, ^{82}Se, or ^{150}Nd [57]. The SNO+ Collaboration plans to fill the SNO acrylic vessel with neodymium-loaded liquid scintillator in order to search for $0\nu\beta\beta$ decay of ^{150}Nd [58]. In a similar fashion, the KamLAND Collaboration plans to turn the KamLAND detector into a $0\nu\beta\beta$ decay experiment by dissolving ^{136}Xe in the liquid scintillator [59]. The CANDLES project (CAlcium fluoride for the study of Neutrinos and Dark matters by Low Energy Spectrometer) uses undoped CaF_2 scintillators to attempt to detect $0\nu\beta\beta$ decay of ^{48}Ca [60].

Chapter 3

Bolometric detectors

A bolometer is a low temperature detector that measures the energy of an incident particle by the induced rise in temperature of the detector. Such small energy depositions result in measurable temperature changes because at low temperature the heat capacity, C, of a dielectric and diamagnetic crystal of mass m and molar mass M behaves according to the Debye law as

$$C = \frac{m}{M}\frac{12}{5}\pi^4 N_A k_B \left(\frac{T}{\Theta_D}\right)^3, \tag{3.1}$$

where Θ_D is the Debye temperature, which depends on the material [61]. The temperature change, given by $\Delta T = E/C$ for an energy deposition E, becomes large enough to be measured to high precision when C is made sufficiently small by cooling the detector to a very low temperature on the order of 10 mK.

Since most of the energy transferred from a particle to a detector is eventually converted into heat, a bolometer has a higher intrinsic resolution than other types of nuclear radiation detectors that measure the component of a particle's energy that goes into ionization or excitation of atomic electrons. Aside from the possibility of energy loss due to excitation of metastable states in the lattice, the theoretical resolution of a bolometer is limited only by fluctuations in the number of phonons exchanged with the heat sink that maintains the bolometer's base temperature. An estimate of this thermodynamic limit on the resolution of a bolometer may be obtained by considering that the elementary excitation, the energy required to create one phonon, is about $\epsilon = k_B T$ so that the number of phonons in a bolometer with energy $E = CT$ is $N = E/\epsilon = CT/k_B T$. Assuming this number of phonons fluctuates according Poisson statistics, the variation in the energy is

$$\Delta E = \sqrt{N} \cdot \epsilon = \sqrt{k_B C T^2}, \tag{3.2}$$

a quantity that is independent of energy. Plugging in typical values for CUORE bolometers ($C \approx 1$ MeV/0.1 mK and $T \approx 10$ mK) yields an energy fluctuation of $\Delta E \sim 10$ eV. While this value is impressive, roughly two orders of magnitude better than conventional detectors, it should be stressed that this is merely the thermodynamic limit on the detector technology. In a real detector, the energy resolution is degraded due to extrinsic sources of noise, and the contribution of Eq. (3.2) is negligible.

3.1 The CUORE bolometer module

CUORE and its predecessor experiments, including CUORICINO and the Three Towers Test, are arrays of independent bolometer modules. A bolometer module consists of three essential parts: an energy absorber, a temperature sensor, and a Joule heater for monitoring the gain. When the bolometer modules are assembled into arrays, Teflon spacers between the crystals and the copper structure of the array provide the low thermal conductivity connection between the energy absorber and the heat sink.

3.1.1 Energy absorber: TeO_2 crystal

The energy absorber is a TeO_2 crystal, which is also the source of double beta decaying ^{130}Te. For CUORE each crystal will be $5 \times 5 \times 5$ cm^3 with a mass of 750 g. The CUORICINO crystals were of two types: 44 were slightly larger with an average mass of about 790 g, and 18 were smaller, $3 \times 3 \times 6$ cm^3 and 330 g. A photograph of four of the $5 \times 5 \times 5$ cm^3 CUORICINO crystals is shown in Fig. 3.1. The Debye temperature of TeO_2 is (232 ± 7) K [62], leading to a heat capacity of about 1 MeV/0.1 mK for the larger crystals at an operating temperature of 10 mK. Crystals of tellurium dioxide are preferred over pure tellurium crystals because of their thermal and mechanical properties. Pure tellurium was tested as a bolometer, but the mechanical stress of thermal contraction caused excessive damage to the crystal [63]. Tellurium dioxide crystals, on the other hand, can undergo repeated thermal cycling with no observable damage or decrease in performance as bolometers. Moreover, TeO_2 crystals have a higher Debye temperature than pure tellurium crystals, yielding a lower heat capacity and therefore larger pulse amplitudes at the same working temperature. Tellurium dioxide crystals are also advantageous because they are readily available from experienced commercial manufacturers due to their use as acousto-optic materials in industrial applications. The TeO_2 crystals for CUORICINO and CUORE were grown by the Shanghai Institute of Ceramics, Chinese Academy of Sciences (SICCAS).

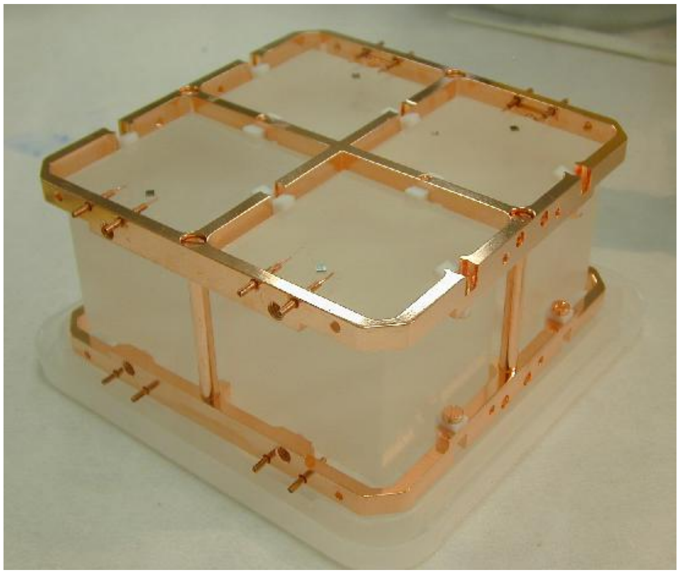

Figure 3.1: A CUORICINO floor with four $5 \times 5 \times 5$ cm^3 TeO$_2$ crystals.

3.1.2 Temperature sensor: NTD Ge thermistor

The temperature sensor that converts the thermal variation into an electrical signal is a neutron transmutation doped (NTD) germanium semiconductor thermistor. A semiconductor thermistor is a high resistance semiconductor with a doping density below the metal-insulator transition [64]. Conduction occurs when charges tunnel across the potential barrier between impurity sites. Different impurity sites have different energy levels, and the charges acquire or give up the necessary energy difference by absorbing or emitting a phonon. At low temperatures, with few high energy phonons available, charges may tunnel, or "hop," not only to nearest neighbor sites but over longer ranges in order to find a site with an energy difference matched to an available phonon. This conduction regime is known as variable range hopping. The resistivity of a semiconductor thermistor in the variable range hopping regime depends on its temperature as

$$\rho = \rho_0 e^{(T_0/T)^\gamma}, \tag{3.3}$$

with $\gamma = 1/2$. The parameters ρ_0 and T_0 must be determined experimentally.

The NTD technique produces germanium thermistors with highly homogeneous concentrations of dopants and good reproducibility [65]. To produce an NTD Ge thermistor, a wafer of natural germanium (Fig. 3.2) is exposed to the thermal neutron flux from a nuclear reactor. Neutron capture reactions on the stable germanium isotopes create unstable germanium isotopes that subsequently

Figure 3.2: Germanium wafers used to make NTD Ge thermistors.

decay to the desired dopant elements. The nuclear processes that take place are:

$$\begin{aligned} {}^{70}\text{Ge (21\%)} + n &\rightarrow {}^{71}\text{Ge } (\sigma_T = 3.43 \pm 0.17\text{ b}, \sigma_R = 1.5\text{ b}) \\ {}^{71}\text{Ge} &\rightarrow {}^{71}\text{Ga } (t_{1/2} = 11.4\text{ d}) \qquad \text{Acceptor} \end{aligned} \tag{3.4}$$

$$\begin{aligned} {}^{74}\text{Ge (36\%)} + n &\rightarrow {}^{75}\text{Ge } (\sigma_T = 0.51 \pm 0.08\text{ b}, \sigma_R = 1.0 \pm 0.2\text{ b}) \\ {}^{75}\text{Ge} &\rightarrow {}^{75}\text{As } (t_{1/2} = 83\text{ min}) \qquad \text{Donor} \end{aligned} \tag{3.5}$$

$$\begin{aligned} {}^{76}\text{Ge (7.4\%)} + n &\rightarrow {}^{77}\text{Ge } (\sigma_T = 0.16 \pm 0.01\text{ b}, \sigma_R = 2.0 \pm 0.35\text{ b}) \\ {}^{77}\text{Ge} &\rightarrow {}^{77}\text{As } (t_{1/2} = 11.3\text{ h}) \\ {}^{77}\text{As} &\rightarrow {}^{77}\text{Se } (t_{1/2} = 38.8\text{ h}) \qquad \text{Double donor} \end{aligned} \tag{3.6}$$

where σ_T and σ_R are the thermal and epithermal neutron capture cross sections, respectively. Since the neutron interaction probability is low, the entire volume of germanium is exposed to a nearly uniform neutron flux, leading to very uniform doping levels. The thermistors for CUORICINO were irradiated at the Missouri University Research Reactor (MURR), and the thermistors for CUORE were irradiated at the MIT Nuclear Reactor Laboratory (MIT-NRL).

After a waiting period of at least six months for the activity of the ^{71}Ge in the sample to decay to an acceptable level, the germanium wafers are cut into pieces to make individual thermistors. From Eq. (3.3) the resistance of a thermistor at temperature T is given by

$$R = R_0 e^{(T_0/T)^\gamma} \tag{3.7}$$

with $R_0 = \rho_0 \ell / A$ where ℓ and A are the thermistor's length and cross sectional area, respectively. The thermistors are characterized by measuring their resistance as a function of temperature by gluing the thermistor to a low temperature heat sink with a high thermal conductivity epoxy. The temperature of the heat sink is varied and monitored with a calibrated thermometer. From a fit to the $R(T)$ data points, the parameters R_0 and T_0 are obtained. Characteristic values for CUORE thermistors are $R_0 = 1.15\ \Omega$ and $T_0 = 3.35$ K, corresponding to a resistance of approximately 100 MΩ at 10 mK. A parameter quantifying the performance of a thermistor is the logarithmic sensitivity, η, defined by

$$\eta = \left| \frac{d \ln R(T)}{d \ln T} \right| . \tag{3.8}$$

Using Eq. (3.7), the logarithmic sensitivity of a CUORE thermistor is given by

$$\eta = \gamma \left(\frac{T_0}{T} \right)^{\gamma} , \tag{3.9}$$

which usually evaluates in the range 2–10.

The NTD Ge thermistor is glued to a $\mathrm{TeO_2}$ crystal with Araldit Rapid epoxy. The glue is applied in nine spots, arranged in a 3 × 3 grid, of 0.5 mm diameter. The spacing between the thermistor and the crystal is 50 μm. Compared with a single large glue spot, this arrangement of nine glue spots reduces the mechanical stress on the thermistor due to differences in thermal contraction between the thermistor and the crystal. Mechanical stresses can affect the resistance of the thermistor and degrade its performance as a thermometer.

3.1.3 Joule heater

The third element of a CUORE bolometer module is a silicon resistor, glued to the $\mathrm{TeO_2}$ crystal with the same Araldit Rapid epoxy used to bond the thermistor. The resistor is used as a Joule heater to inject a constant energy into the bolometer at regular intervals, usually once every 5 minutes. The resistor, typically 50–100 kΩ, is very stable with temperature and is pulsed with an ultra-stable pulser [66]. The heat is injected over a very short time compared with the response time of the detector, simulating the energy deposition of a particle interaction. The heater pulses are used to obtain frequent calibrations of the gain of the detector, which varies with temperature. Using the heater pulses, the gain is stabilized as part of the offline analysis (discussed in Sect. 4.7) [67].

3.1.4 Bolometer operation

To read out the signal from the NTD Ge thermistor, the thermistor is biased with the circuit shown in Fig. 3.3. The biasing circuit consists of a voltage source

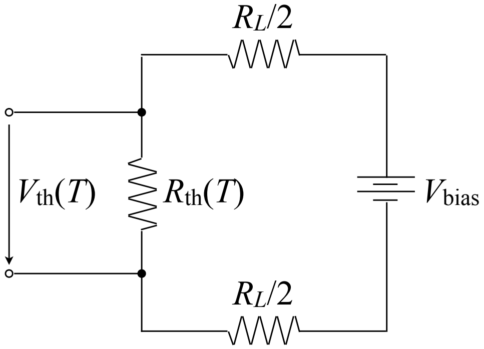

Figure 3.3: Biasing circuit for an NTD Ge thermistor.

and two load resistors in series with the thermistor. The total resistance of the load resistors is chosen to be much greater than the resistance of the thermistor at the working temperature so that an approximately constant current $I = V_{\mathrm{bias}}/(R_L + R_{\mathrm{th}}) \approx V_{\mathrm{bias}}/R_L$ flows through the thermistor. The voltage across the thermistor, $V_{\mathrm{th}} = IR_{\mathrm{th}} \approx V_{\mathrm{bias}}R_{\mathrm{th}}/R_L$, is proportional to the thermistor resistance. The thermal information contained in the thermistor resistance is read out by recording the voltage across the thermistor.

The optimal bias voltage is determined independently for each bolometer. As the bias voltage is increased from zero and current flows through the thermistor, power $P = IV_{\mathrm{th}}$ is dissipated as heat in the thermistor. The Joule heating of the thermistor increases its temperature and decreases its resistance, a phenomenon known as electrothermal feedback. The I-V relationship, or load curve (Fig. 3.4), for the thermistor begins approximately linearly at low bias voltages where electrothermal feedback is negligible. As the bias voltage increases, the slope of the I-V curve increases until reaching the inversion point where the thermistor voltage is maximal. At higher bias voltages, the curve reverses direction, and the thermistor voltage decreases while the current in the biasing circuit continues to increase. The working point of the thermistor is a particular point on the load curve set by the bias voltage. The optimal working point is the one where the ratio of signal amplitude to noise level is maximized. In practice the optimal working point is found approximately by scanning the bias voltage in steps and selecting the value for which the signal amplitude is maximized (Fig. 3.5).

Energy released in the crystal from particle interactions is dissipated as heat through a copper mounting structure, which is thermally linked to the mixing chamber of a dilution refrigerator. Since the crystals are large and the connection between the crystal and the copper frame has low thermal conductivity, the heat generated in a crystal by a particle interaction dissipates slowly, resulting in a

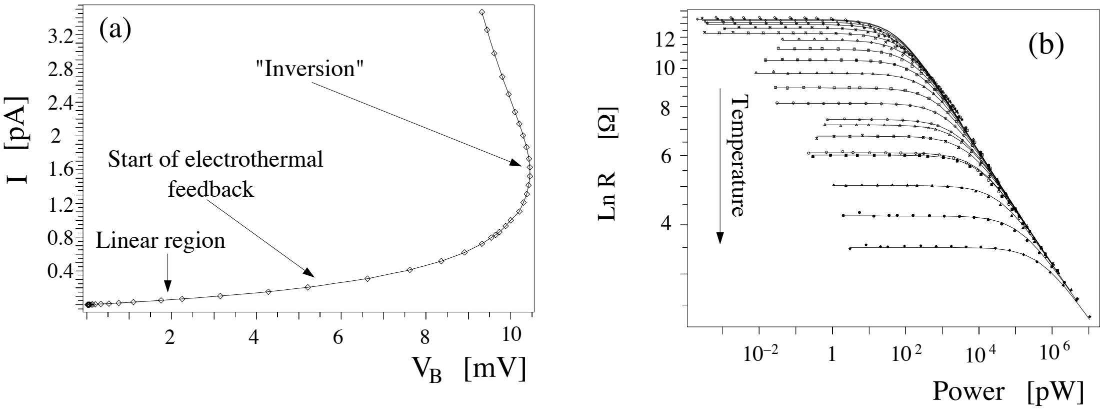

Figure 3.4: On the left, the load curve for an NTD Ge thermistor. On the right, measurements illustrating electrothermal feedback: As the power dissipated in the thermistor increases, the resistance of the thermistor decreases. The different curves were obtained at different base temperatures, i.e. the temperature at $P = 0$.

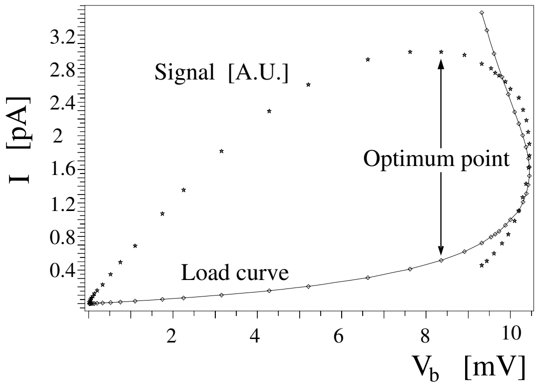

Figure 3.5: Superimposed on the load curve is the signal pulse height of a constant energy input from the heater as a function of the bias voltage for the NTD Ge thermistor. The bias voltage is optimized by choosing the value that maximizes the signal pulse height.

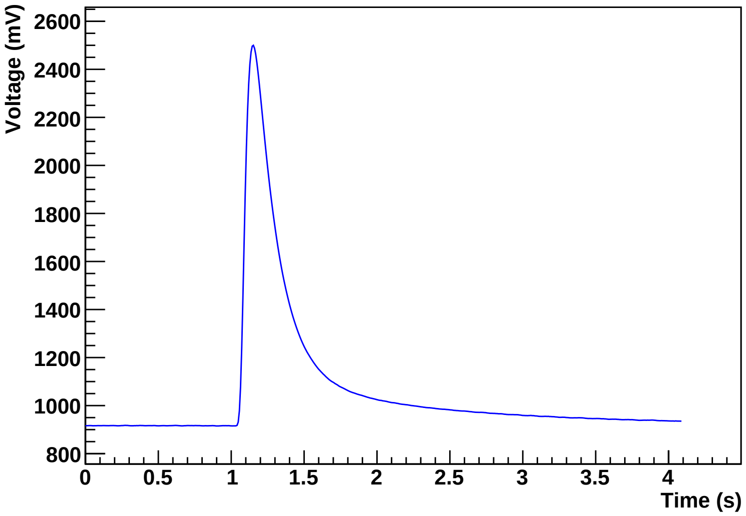

Figure 3.6: A 2614 keV pulse from CUORICINO channel 15, run 174 (event #1240).

temperature pulse with a characteristic decay time on the order of 1 second. A typical pulse is shown in Fig. 3.6.

3.2 Bolometric experiments for neutrinoless double beta decay searches

Since Fiorini and Niinikoski first proposed the use of bolometers for rare decay searches in 1984 [68], members of the CUORE Collaboration have operated a series of $0\nu\beta\beta$ decay experiments based on the bolometric technique. Starting with a single crystal bolometer [69], the detectors were increased in size to arrays of 4, 8, and 20 crystals [70, 71], leading up to the recently completed 62-crystal CUORICINO experiment [72]. Based on this experience, the CUORE Collaboration is currently engaged in constructing a much larger array, CUORE, which will consist of 988 crystals. In addition to the long-running experiments aimed at producing physics results, many short-time-frame R&D experiments have been conducted with the primary goal of reducing backgrounds from radioactive contaminations.

3.2.1 CUORICINO

The CUORICINO experiment was operated from 2003 to 2008 at the Laboratori Nazionali del Gran Sasso (LNGS), an Italian national laboratory in Assergi,

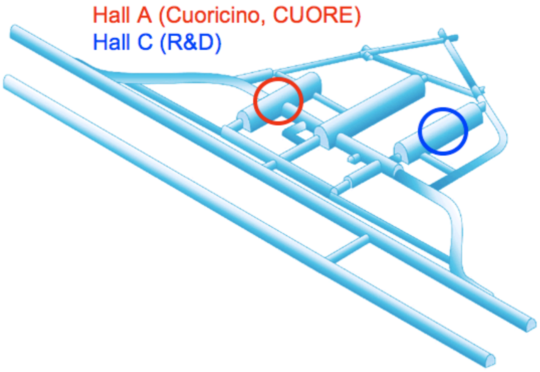

Figure 3.7: Layout of the underground laboratory at LNGS. CUORICINO and CUORE are located in Hall A, and a dedicated R&D facility for CUORE is located in Hall C.

Italy. The laboratory is located about 3500 meters water equivalent underground beneath the Gran Sasso mountain range. The rock overburden reduces the muon flux by a factor of $\sim 10^6$ in the underground laboratory compared with the surface. LNGS contains three large experimental halls hosting many large experiments for neutrino physics and dark matter physics; CUORICINO was located in Hall A (Fig. 3.7).

CUORICINO was comprised of 62 bolometer modules arranged in a tower structure with 13 floors. A picture of the tower is shown in Fig. 3.8. Eleven of the floors contained four $5 \times 5 \times 5$ cm^3, 790 g crystals. Two of the floors contained nine $3 \times 3 \times 6$ cm^3, 330 g crystals. The smaller crystals were reused from the previous 20-crystal array, and the larger crystals were produced specifically for CUORICINO. All of the 790 g crystals and 14 of the 330 g crystals were made from natural tellurium. Two of the smaller crystals were enriched to 75% in ^{130}Te, and two were enriched to 82.3% in ^{128}Te. The purpose of the enriched crystals was to facilitate a measurement of the $2\nu\beta\beta$ decay rate of ^{130}Te by using the crystals enriched in ^{128}Te as blanks for obtaining the background spectrum. The detector contains 40.7 kg of TeO$_2$, of which 11 kg is ^{130}Te.

The detector was housed in a cryostat and cooled to $\sim$8 mK by a dilution refrigerator. The tower was suspended from a 25 mm copper bar connected to a steel spring attached to the 50 mK plate of the cryostat. Hanging the tower from a spring mechanically decouples the detector from the cryostat, helping to prevent vibrations of the refrigerator from propagating to the detector and causing heating or inducing thermal noise. To shield against environmental radioactivity, lead shielding was placed outside and inside the cryostat, as shown in Fig. 3.9. To keep costs manageable, a layered approach was taken, using less costly but higher activity lead farther from the detector. Surrounding the cryostat were

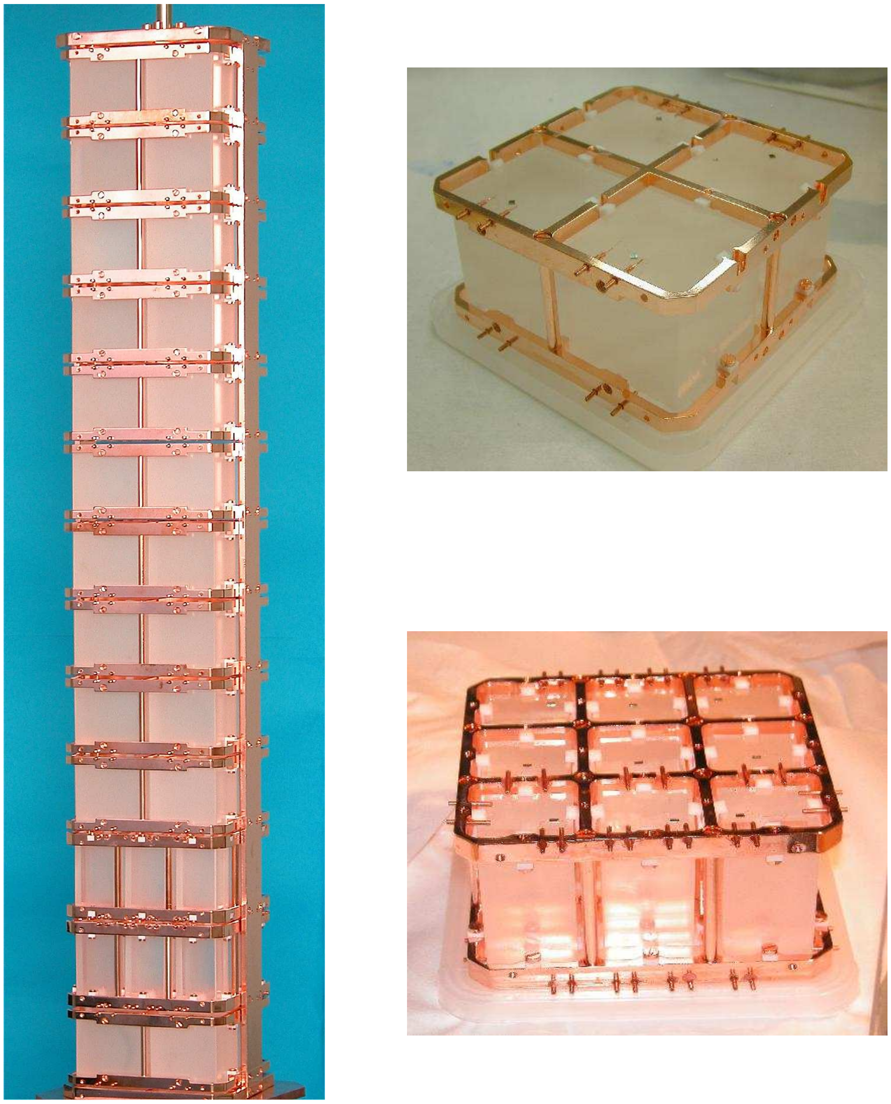

Figure 3.8: Photographs of the CUORICINO tower and two of the individual floors. On the top right is a floor of $5 \times 5 \times 5$ cm^3 crystals, and on the bottom right is a floor of $3 \times 3 \times 6$ cm^3 crystals.

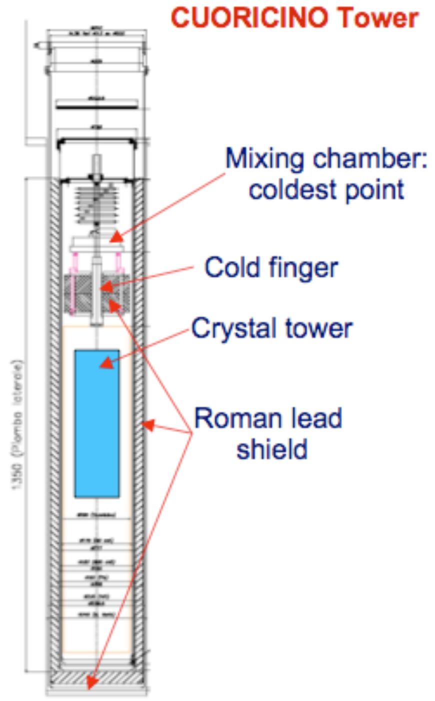

Figure 3.9: Diagram of internal lead shielding for CUORICINO.

two layers of commercial lead, each with 10 cm minimum thickness. The outer layer was common low radioactivity lead, and the inner layer was a special lead whose ^{210}Pb activity was measured to be (16 ± 4) Bq/kg. Inside of the cryostat, shielding the detector from the radioactive contaminations of the dilution unit, a 10 cm layer of "Roman" lead was located directly above the tower and below the dilution unit; the bar from which the tower was hanging passed through a gap in the center of this lead. The Roman lead was recovered from an ancient Roman shipwreck. The ^{210}Pb content in Roman lead has decayed, according to its 22.3 y half-life, to negligible levels (< 0.004 Bq/kg) over the millennia since the lead was separated from other elements. The thermal shields of the cryostat, made of electrolytic copper totaling at least 2 cm thickness, acted as both a shield against external radiation and a source of radioactive contaminations. A lateral 1.2 cm thick internal layer of Roman lead encircled the detector to shield it from activity from the thermal shields of the cryostat. The cryostat was enclosed in a Plexiglas box flushed with clean N_2 from a liquid nitrogen evaporator to expel radon from the vicinity of the detector. The setup was surrounded by a borated polyethylene neutron shield with a minimum thickness of 10 cm and enclosed inside a Faraday cage to suppress electromagnetic interference.

The crystals for CUORICINO were grown by SICCAS using low radioactivity materials previously assayed by the collaboration. They were shipped from China to Italy by sea to minimize the exposure to cosmic rays, which can activate the isotopes of tellurium and oxygen in the crystals. At LNGS the crystals were lapped with a low contamination polishing compound.

The CUORICINO detector was assembled in a clean room located directly above the cryostat where the experiment operated. The assembly was carried out in a glove box that was continuously flushed with nitrogen from liquid N_2 boil off in order to avoid contamination from radon. Crystals were mounted into floors, as shown in Fig. 3.8. The mounting structure was made of oxygen-free high conductivity (OFHC) copper. Copper was chosen as the material for the frame because OFHC copper has low levels of radioactive contaminations and because the thermal contraction of copper is well matched to that of TeO_2. The copper frame was thermalized with the mixing chamber of the dilution refrigerator by means of a metallic strip connecting the two components. Holding the crystals in the frame were Teflon spacers, which served as a weak thermal link between the crystals and the copper frame in addition to providing mechanical support.

Electrical connections to the thermistors were made by 50 μm gold wires bonded to metallic contacts on the thermistors. The 50 μm wires were crimped inside copper pins mounted near the crystal through holes in the copper frame with glue. On the other side of the pins, a twisted pair of wires carried the signal along the tower up to the mixing chamber. From there the signal was carried out of the cryostat by a pair of twisted coaxial cables, which plugged into the front-end electronics boards through a Fischer connector.

The front-end electronics include the biasing circuit, amplifiers, and an antialiasing Bessel filter [73]. The two load resistors for each bolometer's biasing circuit were 27 GΩ each. The bias voltage was adjustable in the range 0–10 V and was set independently for each bolometer as described in Sect. 3.1.4. The thermistor signal of approximately 200 μV/MeV was amplified by a differential voltage preamplifier with a gain of 220 V/V and by a second stage programmable amplifier with a gain between 1 and 45 V/V. The signal was then passed through an antialiasing 6-pole Bessel filter with an attenuation of 120 db/decade and a cutoff of 12 Hz. The filtered signal was acquired by a 16-bit analog-to-digital converter (ADC) sampling at 125 Hz (8 ms period). The dynamic range of the ADC was 0–10 V, and the gain of the amplifier was adjusted to match this range, with acquired pulses having amplitudes on the order of 1 V/MeV. A trigger with threshold typically in the range 20–100 keV, depending on the noise level of the bolometer, selected events to record to disk. For each triggered event, 512 data points (4.096 s) were recorded, including 1 s before the time of the trigger. The DC level of the detector before the triggering event provides a measure of the detector temperature at the time of the event, which is used for the offline gain stabilization. The front-end electronics for all of the small crystals and for 20 of the large crystals were at room temperature. For the other 24 large crystals, the preamplifier and load resistors were located near the detector at ∼100 K to reduce preamplifier noise and microphonic noise. One electronics channel was used for an electrothermal feedback system for stabilizing the tower temperature [74].

The CUORICINO experiment began collecting data in March 2003. In November 2003 the detector was removed from the cryostat to repair connections to 13 thermistors that were broken during the initial cooling down. It was redeployed in May 2004 with all but two thermistors and one heater functional and was operated without opening the cryostat until May 2008 when the experiment was completed. The two deployment phases of CUORICINO, before and after the repairs, are referred to as Run I and Run II.

During 2009–2010 the data from Run II were reprocessed using newly developed analysis software. Since the live time in Run I was small compared to Run II and since improvements were made to the data acquisition for Run II, we decided not to include Run I data in the reprocessing. The analysis procedures applied during the reprocessing campaign to produce calibrated energy spectra from the raw data are described in Chapter 4. The data collected in CUORICINO Run II constitute an exposure of 18.6 kg · y of ^{130}Te.

3.2.2 Three Towers Test

The Three Towers Test (TTT) was an R&D experiment for CUORE designed to compare the surface contamination levels of copper cleaned by three different methods. The detector consisted of three 12-crystal arrays (the three towers) separated from one another by copper shields (Fig. 3.10). The surfaces of the copper parts of the three towers were treated by three different procedures to remove radioactive contaminants. The detector was installed in the cryostat previously occupied by CUORICINO, and it collected data from September 2009 through January 2010. While the primary goal of the TTT was to measure the radioactive contamination levels of the copper parts facing the crystals, I have used the TTT data in combination with the CUORICINO data to increase the statistics available for the $0\nu\beta\beta$ decay search. The TTT augments the exposure of 18.6 kg · y of ^{130}Te from CUORICINO with an additional 1.3 kg · y with a slightly lower background level than CUORICINO.

Since this dissertation is focused on the $0\nu\beta\beta$ decay analysis with the TTT data, only a brief description of the copper cleaning techniques is given here. The copper for the top tower (Fig. 3.11a) was treated following the procedure that produced the best result in previous R&D tests. The copper pieces were cleaned with soap, treated with H_2O_2 + H_2O + citric acid, and wrapped with seven layers of polyethylene. The middle tower (Fig. 3.11b) was cleaned at LNGS by a new chemical process, starting with simple soap and water, followed by electroerosion

(a) (b)

Figure 3.10: Photographs of (a) the Three Towers Test detector suspended from the dilution refrigerator and (b) a closer view of the three towers. The polyethylene tower is on top, the LNGS tower is in the middle, and the Legnaro tower is on the bottom.

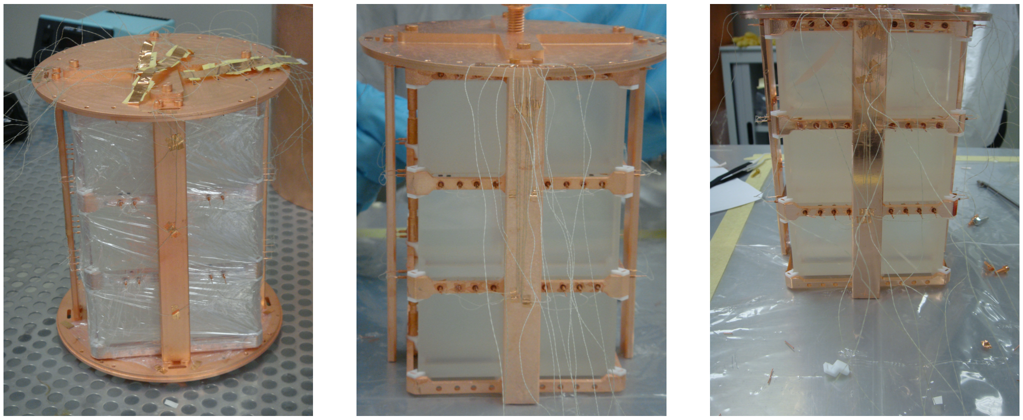
(a) (b) (c)

Figure 3.11: Photographs of the three 12-crystal arrays (the three towers) without their cylindrical shields: (a) the polyethylene tower, (b) the LNGS tower, and (c) the Legnaro tower.

with 85% phosphoric acid, 5% butanol, and 10% water, followed by chemical etching with nitric acid, and finally passivation with H_2O_2 + H_2O + citric acid. The copper for the bottom tower (Fig. 3.11c) was cleaned with a sophisticated multi-step procedure at the Laboratori Nazionali di Legnaro, abbreviated TECM for tumbling, electropolishing, chemical etching, and magnetron. The magnetron step is a plasma cleaning technique. The TECM procedure is the baseline for CUORE.

The setup of the TTT detectors was very similar to CUORICINO. The TTT crystals came from CUORICINO and other R&D tests. The crystals were lapped and polished following the CUORE procedure before being installed in the TTT. Many of the crystals were outfitted with two thermistors for redundancy in case some electrical connections were lost. Unfortunately, we did lose a large number of electrical connections during the cooling down. Circuits that were tested successfully at room temperature became broken as the temperature dropped to 10 mK. The cause of the broken channels is not completely understood but was traced to connectors on the towers that when subjected to thermal contraction must have severed the electrical contact. We attempted to repair the broken channels after the TTT had to be warmed up to room temperature because a strong earthquake struck near the laboratory in L'Aquila on April 6, 2009, forcing the underground laboratory to close for several days. Some of the channels that were broken at 10 mK even regained functionality as the temperature increased. Nevertheless, after the detector was cooled again, 10 of the 36 crystals had no operational thermistors, and 10 of the 26 active crystals did not have a working heater. The

Table 3.1: Active crystals in the Three Towers Test.

Tower	Number of crystals with at least one working thermistor	Number of crystals with a working thermistor and a working heater
Polyethylene	12	9
LNGS	7	3
Legnaro	7	4

working channels are illustrated on the channel map in Fig. 3.12, and a summary is provided in Table 3.1. Since there were enough active crystals in each tower to reach the sensitivity required to test the copper surface contaminations and since attempting another repair would have caused a delay of several weeks, we decided to proceed with the experiment and tolerate the loss of 10 crystals.

A significant difference in the TTT with respect to CUORICINO was the data acquisition system. Instead of the CUORICINO DAQ, the TTT used a new DAQ developed for CUORE and dubbed Apollo. Details of the Apollo system may be found in Ref. [75]. The Apollo DAQ features an expanded ADC range of -10.5 V to $+10.5$ V, compared to 0 V to 10 V for CUORICINO, and a longer acquisition window of 626 samples, compared to 512 samples for CUORICINO. As for CUORICINO, the Apollo ADCs sample at 125 Hz. Also, in addition to triggered data, Apollo stores the continuous readout from the bolometers, although this feature is not used in this analysis.

3.2.3 CUORE

The CUORE experiment is a next generation ^{130}Te $0\nu\beta\beta$ decay experiment based on the experience with CUORICINO [76]. It is funded by the INFN of Italy, the United States Department of Energy, and National Science Foundation of the United States. Infrastructure for CUORE is currently under construction in Hall A at LNGS next to the CUORICINO building, and components for CUORE are being manufactured around the world. The CUORE detector will be a tightly packed array of 988 TeO_2 bolometer modules, each $5 \times 5 \times 5$ cm^3 and 750 g, for a total mass of 741 kg of TeO_2. Since the tellurium is unenriched, 204 kg of the total mass is the isotope of interest, ^{130}Te.

The bolometer modules will be arranged in 19 towers of 13 floors each, with 4 crystals per floor (Fig. 3.13a). Each of the 19 towers is similar to CUORICINO in size and arrangement. The CUORE detector will be housed in a specially

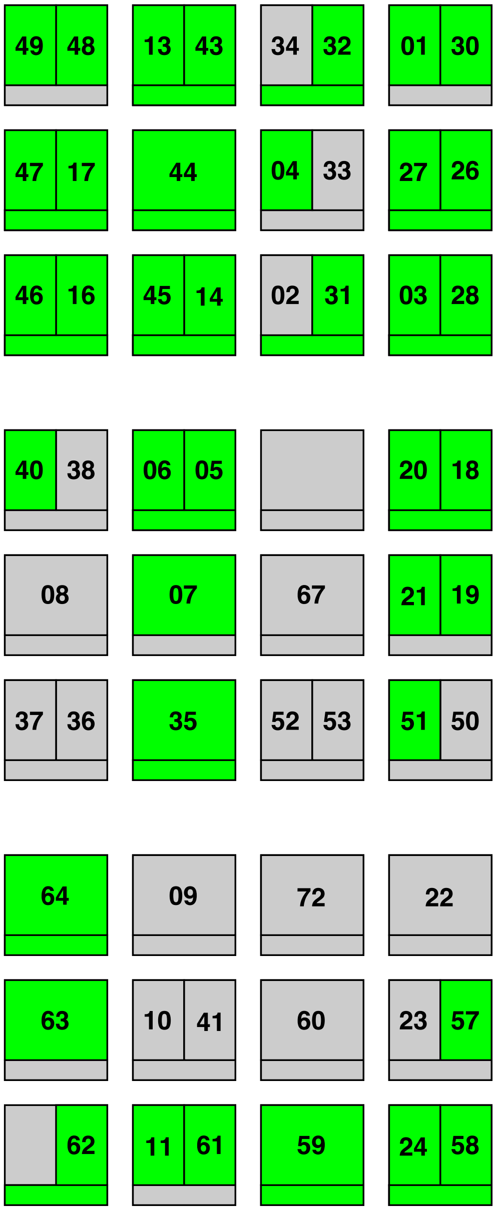

Figure 3.12: Channel map for the Three Towers Test. The towers are: top – Polyethylene, middle – LNGS, bottom – Legnaro. Crystals labeled with two channel numbers have two thermistors, called thermistor left and thermistor right. The horizontal band at the bottom of each crystal represents the heater. Green indicates working; gray indicates not working. Heaters on crystals without a working thermistor are shaded gray; we cannot tell whether they were working or not.

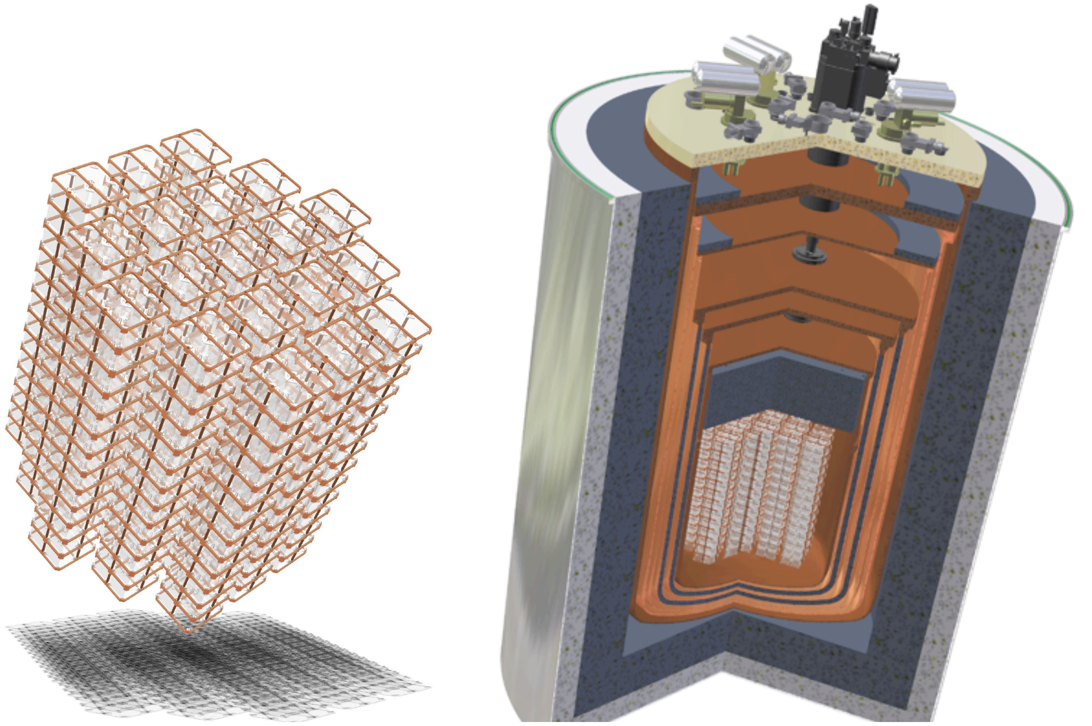

(a) (b)

Figure 3.13: CAD renderings of (a) the CUORE detector array and (b) the CUORE cryostat with the detector array inside.

built cryostat and cooled to ∼10 mK by a pulse-tube-assisted dilution refrigerator (Fig. 3.13b).

CUORE aims for a sensitivity on the order of 10^{26} y for the partial half-life for $0\nu\beta\beta$ decay of ^{130}Te in 5 years of running. To achieve this goal, the CUORE Collaboration has made an intense effort to understand and control the sources of backgrounds in CUORICINO and other cryogenic bolometer R&D experiments. Improvements in the cleanliness of materials and detector shielding lead to an expected background level for CUORE of 0.01 counts/$(\mathrm{keV} \cdot \mathrm{kg} \cdot \mathrm{y})$ in the region of interest around 2527.5 keV.

The main sources of backgrounds in CUORE are expected to be surface contaminations of the copper materials facing the crystals, bulk and surface contaminations of the TeO_2 crystals, and ^{232}Th contaminations in the copper radiation shields of the cryostat. External gamma, neutron, and muon induced backgrounds are estimated by simulations to contribute less than 0.01 counts/$(\mathrm{keV} \cdot \mathrm{kg} \cdot \mathrm{y})$ [77, 78]. Copper surface contaminations will be reduced by the TECM procedure described in Sect. 3.2.2 and tested in the Three Towers Test. The production of radiopure TeO_2 crystals has been the subject of extensive research and development by CUORE collaborators [79]. Dedicated production lines were set up at SICCAS for the growth and surface processing of CUORE crystals. Raw materials, reactants, consumables, ancillaries, and intermediate products used in the production of the crystals are screened for radioactive contaminations. The crystals are shipped from SICCAS to Italy by sea to minimize exposure to cosmic rays, and upon arrival at Gran Sasso, the crystals are immediately stored underground. As a final quality control procedure, four crystals are randomly selected from each production batch (of approximately 60 crystals) to undergo a bolometric test in the CUORE R&D cryostat in Hall C of LNGS. In the bolometric test, which lasts for a few weeks and is known as a CUORE Crystal Validation Run (CCVR), the crystals are operated as bolometers as they will be for CUORE and used to count their own radioactive contaminations. The bulk contamination levels of ^{238}U and ^{232}Th are measured (or upper limits set), and the bolometric performance is evaluated in terms of the energy resolution. Three CCVR tests have been completed so far, and the fourth is currently under way. All crystals tested have met the CUORE specifications.

The CUORE detectors will be shielded from environmental radioactivity and radioactive contaminations of the dilution refrigerator by several layers of lead shielding both inside and outside the cryostat, with Roman lead constituting the innermost layer [76]. The cryostat radiation shields will be built from high purity copper that has been stored underground at LNGS to prevent cosmogenic activation; the copper material will be brought to the surface just in time for the construction of the cryostat. The cryostat and lead shielding will be enclosed within a neutron shield made of 10 cm thick borated polyethylene, which will be continuously flushed with dry nitrogen to exclude radon from the space near the detectors.

The geometry of the CUORE bolometer array provides an intrinsic advantage over a single tower like CUORICINO for rejecting backgrounds that deposit energy in multiple crystals. Many backgrounds, such as alpha decays near the surface of a crystal and Compton-scattered gammas, cause interactions in multiple crystals that are effectively simultaneous. An anti-coincidence cut will be highly effective at suppressing these backgrounds in CUORE. Furthermore, the inner crystals are shielded by the outer crystals.

Other advances in the CUORE design relate to the mechanical structure and assembly. The tower design has been optimized to reduce the amount of copper near the crystals. The NTD thermistors feature a new design for the electrical contacts, which will make the assembly of the bolometer modules much easier and less error prone. The NTD thermistors for CUORICINO had electrical contacts on the sides, requiring the 50 μm signal wires to be bonded to a thermistor before the thermistor was glued onto a crystal. The fragile connections of the signal wires to the thermistors were prone to break during the procedure of gluing a thermistor to a crystal. For CUORE the electrical contacts wrap around from the sides onto the top of the thermistors, which allows a thermistor to be glued to the crystal first and the signal wires to be bonded to the top of the thermistor already attached to the crystal. The detector assembly will be carried out in a dedicated cleanroom in the underground laboratory. The bolometer modules and towers will be assembled in custom glove boxes flushed with nitrogen so that the crystals and other components of the towers do not come into contact with air. The final installation of the towers in the cryostat requires too large a space to be performed within a nitrogen environment and therefore will be performed in a section of the CUORE cleanroom containing nearly radon-free air supplied by an activated charcoal filtration system.

The CUORE goal for the energy resolution is 5 keV FWHM, which is a modest improvement over the 7 keV average for the CUORICINO $5 \times 5 \times 5$ cm^3 crystals. The improvement is due mainly to minimization of mechanical vibrations of the crystals. The CUORE crystals have more stringent tolerances on their dimensions than the CUORICINO crystals in order to ensure they fit snuggly in their holders, and the detector will be mechanically decoupled from the building structure, pumps, and cryocoolers by a carefully designed suspension system in order to suppress the propagation of vibrations to the crystals. In addition to improving the average energy resolution, these improvements in crystal uniformity and mechanical vibrations are expected to lead to much less variation in energy resolution between different detectors compared with CUORICINO.

Chapter 4

First-level data analysis

This chapter describes sequence of computational procedures applied to the raw data to reconstruct the energy of the events in the detectors. The main steps involved are pulse amplitude evaluation by a signal processing technique that mitigates the influence of noise, stabilization of pulse amplitudes to correct for gain variation using heater pulses, and determination of the calibration coefficients for converting stabilized pulse amplitudes to energies. The primary requirement on these first-level analysis procedures is to maintain the approximately 0.2% energy resolution of which the bolometric detectors are capable. A secondary requirement is that the computational procedures should be automated to the fullest extent possible in order to handle several years worth of data. In view of CUORE and its 988 channels, automation and robustness of low-level analysis procedures are of critical importance for the timely production of results.

4.1 Raw data

The data are organized into events defined by a trigger. The trigger is based on a threshold: If the thermistor voltage exceeds its baseline value by a programmable amount, the trigger fires and the corresponding voltage pulse is recorded. The value of the threshold is tuned for each channel independently and is set as low as possible without producing an unreasonably high trigger rate due to noise spikes firing the trigger. In CUORICINO and the TTT, the trigger thresholds are generally in the range of 30 to 50 keV. The raw data recorded for each event consists of pulse samples, the trigger time, the channel number, and a unique event number.

For both CUORICINO and the TTT, the thermistor voltage is digitized with an 8 ms sampling period (125 Hz). In CUORICINO 512 pulse samples are stored for each event, and in the TTT 626 pulse samples are stored. The first approximately 1 s of the acquisition window is used to store the pre-pulse baseline level of

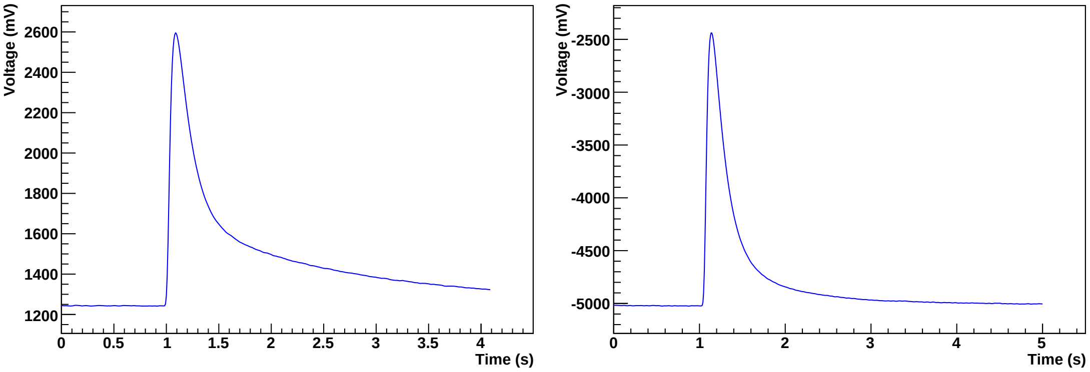

Figure 4.1: Left: A 2615 keV pulse in CUORICINO (channel 6, run 1203, event #396). Right: A 2615 keV pulse in the TTT (channel 11, run 100045, event #48751). On some CUORICINO channels, such as the one displayed here, the pulse does not fully recover to the baseline level within the acquisition window. Therefore, a longer acquisition window is used for the TTT.

the detector in order to have a measure of the detector temperature at the time of the pulse. The remaining ∼3 s of a CUORICINO acquisition window is not long enough for the pulses on some channels to recover to the baseline level within the acquisition window; an example of such a channel is shown in Fig. 4.1. For this reason a longer acquisition window is used in the TTT.

The data collection is organized by runs, with a run lasting approximately one day. The daily runs, with a lowercase "r," should not be confused with the two deployments of CUORICINO, Run I and Run II with a capital "R." Runs are grouped into data sets defined by calibration runs. A data set generally contains about one month of data and ideally starts and ends with calibration runs. Some data sets have only an initial calibration or only a final calibration due to problems such as power outages in the laboratory forcing a warm up of the cryostat before the planned final calibration.

4.2 Analysis software framework

The number of channels in CUORE will be over an order of magnitude greater than in CUORICINO. To meet the challenges of more channels and higher data rates posed by CUORE, the collaboration has developed in recent years a new analysis software framework, named Diana, which is employed for the analysis in this thesis. The Diana framework defines the event class and standard data storage format and provides a mechanism for the work flow of reading events from a file, calling user-created code for each event, and writing the (possibly modified or appended to) events back to file. All of the code is written in C++. The framework supports a flexible approach to data processing through a design in which user code is placed in modules, which are C++ classes inheriting from a common base class. Modules can be selected at run-time via plain text configuration files. A collection of DIANA configuration files drives the analysis procedures.

4.3 Analysis database

The collaboration maintains a database for storing information about the detectors and about each run, including electronics and DAQ parameters, run type, start and stop times, bad time intervals, and other information relevant for the analysis. It is a relational database implemented with the open-source PostgreSQL Database Management System [80]. The database server is currently hosted on `cuoredb.lngs.infn.it`. Analysis code connects remotely to the database server to retrieve the parameters needed for the computation being carried out. The DIANA framework supplies a set of standard database interface functions to allow access to information in the database using only C++ code.

4.4 Pulse amplitude evaluation

To a good approximation, the shape of a bolometer pulse is independent of the pulse amplitude. The acquired voltage signal as a function of time can be expressed as

$$V(t) = b + a \cdot s(t - t_0) + n(t), \tag{4.1}$$

where b is the baseline DC level of the detector, a is the pulse amplitude, t_0 is the start time of the pulse, $s(t)$ is the detector response function (the pulse shape without noise, normalized to unit height), and $n(t)$ is the stochastic noise. The baseline, b, is estimated by averaging over the $\sim$1 s pre-pulse interval. The simplest estimate of the pulse amplitude is the pulse maximum minus the baseline value. An estimate with better signal-to-noise ratio can be obtained since the whole pulse contains information that can be exploited in estimating the amplitude of the pulse. A commonly used technique to calculate the amplitude of a signal superimposed with stochastic noise is the optimal filter. The optimal filter technique weights different Fourier components of the pulse differently depending on the expected relative signal-to-noise ratio of each frequency component. Frequencies that are highly influenced by noise are down-weighted, thus mitigating the influence of the noise on the pulse amplitudes. The optimal filter is implemented as a digital filter with transfer function

$$H(\omega) = \frac{S^*(\omega)}{N(\omega)} e^{-i\omega t_M}, \tag{4.2}$$

where $S(\omega)$ is the Fourier transform of the ideal pulse without noise, $N(\omega)$ is the average noise power spectrum, and t_M is the time of the pulse maximum. An approximate representation of the ideal pulse for each channel is obtained by point-wise averaging in the time domain over many high energy pulses. In computing the point-wise average, the pulses are shifted so that their maxima are aligned by the following procedure. Since bolometer pulses have very flat tails, the derivative of the signal is assumed to go to zero at the beginning and at the end of the acquisition window. The derivative of each pulse contributing to the average is cyclicly shifted so that the position of the maximum of the raw signal falls at a chosen reference point, and the derivative is integrated to recover the pulse shape. The average noise power spectrum for each channel is obtained by averaging the power spectrum of many noise samples that are acquired by a random trigger and determined not to contain a pulse. In CUORICINO all data sets use the same set of average pulses and average noise power spectra, but in the TTT separate collections of average pulses and average noise power spectra are produced for each data set. Figure 4.2 displays the average pulse and average noise power spectrum for CUORICINO channel 1. The pulse amplitude can be evaluated in the time domain by applying Eq. (4.2), transforming back to the time domain, and measuring the amplitude based on the maximum of the filtered pulse. The amplitude can also be evaluated in the frequency domain by integrating over frequency components after applying Eq. (4.2). The time domain amplitude is the standard one used in CUORICINO and the TTT.

4.5 Identification of re-triggered pulses

It can happen that the same pulse is contained in two or more acquisition windows if it is preceded within ~3 seconds by another pulse or a noise spike that also causes the trigger to fire. An example is shown in Fig. 4.3. This may lead to the amplitude of the same pulse being attributed to more than one triggered event because the optimal filter algorithm measures the amplitude of the largest pulse in the acquisition window. If the situation is actually a pileup of physical pulses, the pileup is rejected by a pulse shape cut. In the case where the preceding trigger corresponds to a noise spike rather than a physical pulse, this multiply acquired pulse is called a re-triggered pulse, and it is not rejected by the pulse shape cut since its shape is good. To avoid double counting the same event, an algorithm is run after the optimal filter that identifies re-triggered pulses and selects the best

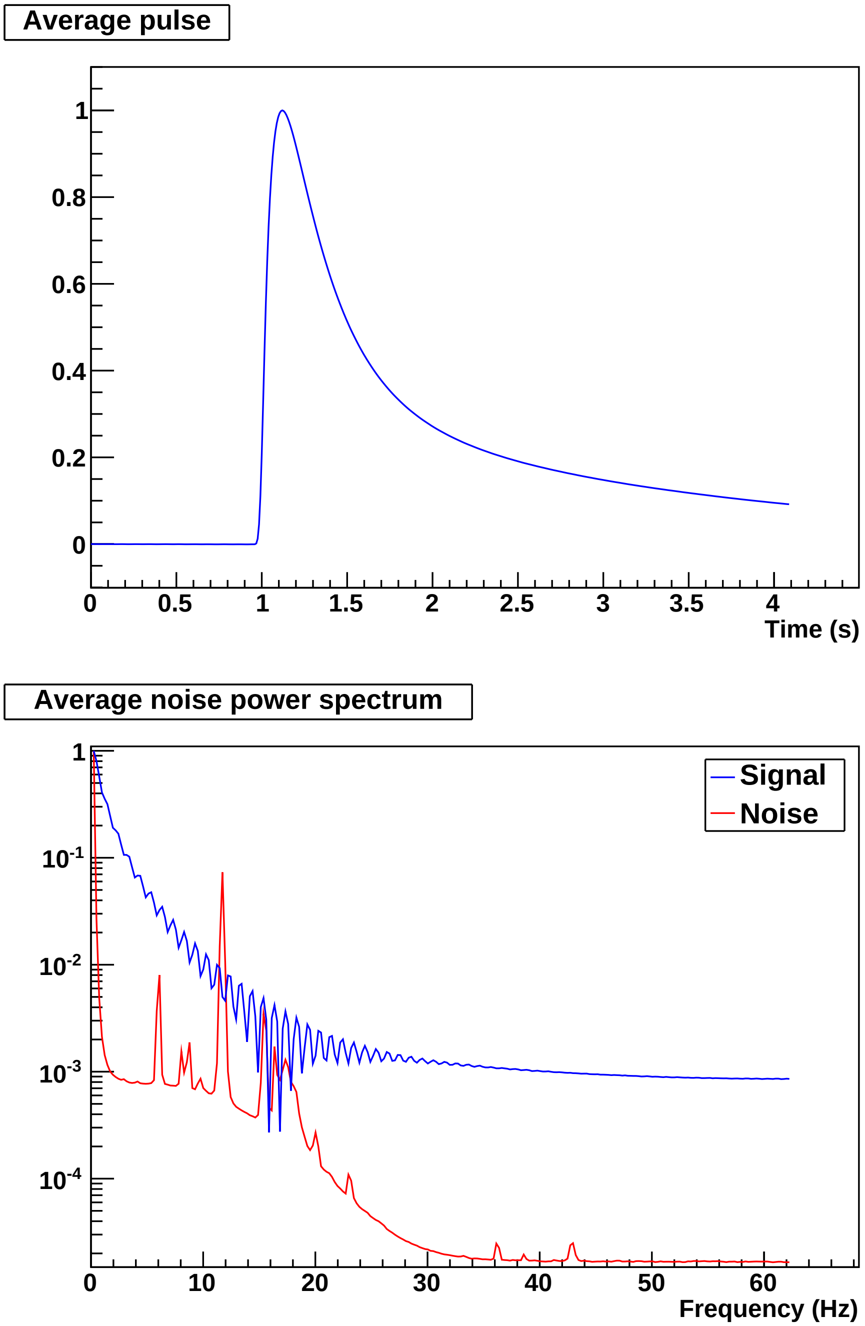

Figure 4.2: Top: The average pulse for CUORICINO channel 1. Bottom: The average noise power spectrum (red) and the modulus of the discrete Fourier transform of the average pulse (blue) for CUORICINO channel 1 (scaled to be equal at the lowest frequency).

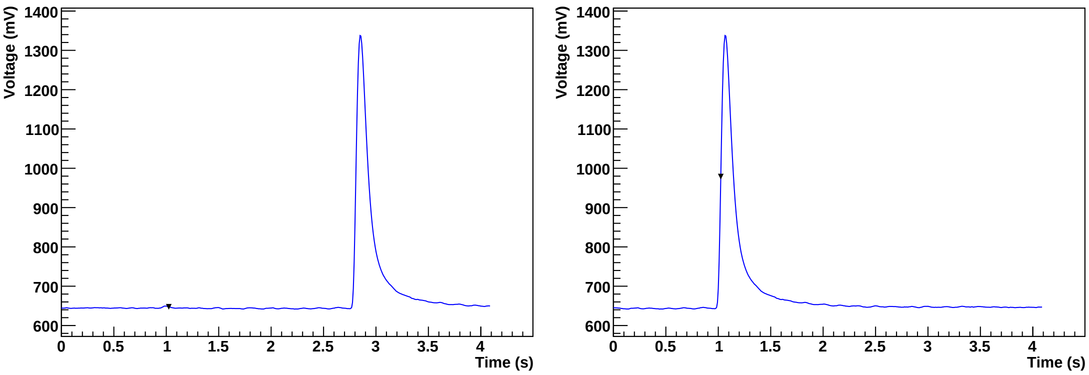

Figure 4.3: Example of a re-triggered pulse. On the left the trigger was fired by a noise spike. The trigger time is roughly indicated by the black triangle. The large pulse ~1.8 seconds after the noise trigger also has its own triggered event, which is on the right. The optimal filter evaluates the amplitude of both triggered events as the amplitude of the same pulse; the amplitude evaluates to 694.3 mV for the event on the left and 694.4 mV for the event on the right. Events are from channel 21, run 540 (event #16573 and #16575).

triggered event for each one, rejecting the other triggered events containing the same pulse. The algorithm identifies re-triggered events based on the absolute time of the pulse maximum for the pulse chosen by the optimal filter. It chooses the best acquisition of the same pulse by picking the one with the pulse closest to the expected trigger time within the acquisition window but not earlier.

4.6 Offline heater flagging

Heater pulses, bolometer events induced by heat injected via the Si resistor, are flagged by the CUORICINO DAQ after run 293. Before run 293 the heater pulses are not flagged in the raw data and must be identified and flagged offline in order to remove them from the spectrum and to select the heater events for use in gain stabilization. Furthermore, after run 293 the hardware flagging of heater pulses by the DAQ is not 100% efficient, leaving some heater pulses flagged as signal events. The Apollo DAQ used for the TTT identifies all heater pulses. An offline heater flagging algorithm was developed to identify heater pulses in CUORICINO with virtually 100% efficiency and negligible false positive rate.

The algorithm is based on the regular timing of the heater pulses and the grouping of heaters into heater channels. The heaters on the same heater channel are connected in parallel so those heaters always fire in coincidence. Table 4.1 lists

Table 4.1: CUORICINO heater channels. Heaters on the same channel are connected in parallel and therefore fire in coincidence.

Heater channel	Bolometer channels
E1-0	7, 16, 31, 39, 46, 55, 63, 67
E1-1	9, 18, 19, 23, 36, 37, 48, 53, 61
E1-2	8, 10, 15, 17, 21, 24, 32, 34, 38, 47, 51, 62
E1-3	11, 12, 20, 22, 33, 35, 49, 50, 52
E2-1	1, 2, 13, 42, 43, 59, 60, 66
E2-2	3, 4, 14, 41, 44, 57, 58, 65
E2-3	5, 6, 40, 45, 54, 56, 64, 68

Table 4.2: Time intervals between stabilization heater pulses during CUORICINO Run II.

Runs	Interval between stabilization heater pulses
174–293	300.224 s
294–1218	305.536 s
1219–1226	300.200 s

the heater channels and corresponding bolometer channels. The pattern and the spacing of heater pulses varied during the experiment. In addition to stabilization heater pulses (medium energy pulses intended to be used for gain stabilization), there were sometimes high energy heater pulses and low energy heater pulses. Table 4.2 summarizes the time intervals between stabilization heater pulses. The most common heater pattern, active during runs 409–1218, consisted of eight stabilization heater pulses separated by 305.536 s, then a gap of 267.224 s followed by a high energy heater pulse, then a gap of 305.536 s followed by a low energy heater pulse, and finally a gap of 343.848 s before the pattern repeats.

The offline heater pulse identification algorithm operates on each run as follows:

1. Do this step for every event in the run. Let t denote the time of the event under consideration, and let T denote the interval between stabilization heaters from Table 4.2. Consider the set of times separated from t by a

multiple of T:

$$\mathcal{T} = \{t + nT \mid n = \dots, -2, -1, 0, 1, 2, \dots\}. \tag{4.3}$$

Let N denote the number of times in $\mathcal{T}$ that are within ±1 s of an event on the same channel with an amplitude greater than 1000 mV. If the event under consideration is a stabilization heater pulse, N will be large, roughly equal to the number of stabilization heater events in the run. If the event under consideration is not a heater pulse, N will be close to 0.

2. Repeat the previous step but adjust the time of each event by −267.244 s because the high energy heater is separated by an interval of 267.244 s from the preceding stabilization heater on some channels. That is, consider

$$\mathcal{T}' = \{t - 267.244\ \mathrm{s} + nT \mid n = \dots, -2, -1, 0, 1, 2, \dots\}. \tag{4.4}$$

Let N' denote the number of times in $\mathcal{T}'$ that are within ±1 s of an event on the same channel with an amplitude greater than 1000 mV.

3. Define a goodness parameter $G \equiv \max\{N, N'\}$, which is large for heater events and small for non-heater events.

4. Set a cut on G based on a maximum false positive rate of 0.1% under the hypothesis that the events with amplitudes greater than 1000 mV are randomly distributed. Under the random distribution hypothesis, the probability of an event occurring within ±1 s of any specified time is given by $p = 1 - e^{-2r(1\ \mathrm{s})}$, where r is the average event rate (here the rate of events with amplitude greater than 1000 mV). The probability of k of the times in $\mathcal{T}$ being within ±1 s of an event with amplitude greater than 1000 mV is given by the binomial distribution, $P(k) = \mathrm{Binomial}(k; p, n)$, where the number of trials, n, is the number of times in $\mathcal{T}$ that are within the start and stop times of the run. The probability that $N > k$ under the random distribution hypothesis is given by

$$\sum_{k'=k+1}^{n} P(k'), \tag{4.5}$$

which is also the probability that $N' > k$. If the cut on G is set at $G > k$, then to fail the cut requires both $N \leq k$ and $N' \leq k$. Therefore, the probability of failing the cut under the random distribution hypothesis is given by

$$\left(\sum_{k'=0}^{k} P(k') \right)^2 . \tag{4.6}$$

The cut on G is set at the smallest goodness parameter value, k, such that the false positive rate (the probability of passing the cut) under the random distribution hypothesis is less than 0.1%:

$$1 - \left(\sum_{k'=0}^{k} P(k') \right)^2 < 0.1\%. \tag{4.7}$$

5. In addition to requiring $G > k$ to flag an event as a heater event, there is a coincidence requirement. If $G > k$ and there are at least two other events on the same heater channel in coincidence (within ±50 ms) with the event under consideration, flag the event as a heater event.

6. As an independent method for identifying heater events, flag an event as a heater event if it is highly likely to be a heater event based on coincidence events alone: If there are at least 5 other events on the same heater channel in coincidence (within ±50 ms) with the event under consideration and no more than 2 other coincident events not on the same heater channel, flag the event as a heater event.

The effect of the offline heater identification on the spectrum of one data set is illustrated in Fig. 4.4. The figure shows that many heater events that were not identified by the hardware flag at the time of data acquisition are successfully identified offline. The hardware flag fails to identify heater events when the heater events occur at slightly the wrong time, which happens when the PC that instructs the pulser to fire the heaters becomes busy and delays the firing instruction past the scheduled time. Usually, the delay is less than 1 s so the heater events pass the cut on the timing goodness parameter, G. If the delay is greater than 1 s, the heater events are identified by the independent method based on coincident events on the same heater channel.

An example of a situation the offline heater identification algorithm cannot handle is the sharp peak in Fig. 4.4 just below 6000 keV, which is due to heater events on channel 66 and is not identified by the offline algorithm. The offline algorithm does not flag those events because they were not triggered by heater pulses. The heater pulses are contained in the acquisition window of a previous trigger, as shown in Fig. 4.5. The situation is similar to a re-triggered pulse (Sect. 4.5), but these heater pulses are not re-triggered pulses because there is not an acquisition window triggered by the heater pulse itself. The acquisition window in Fig. 4.5 is the only acquisition window containing that heater pulse. The DAQ does not trigger on the heater pulse itself because the voltage level is rising between the previous trigger and the pulse, and the trigger algorithm requires the voltage level to reach a local maximum before it will trigger twice within 4 s.

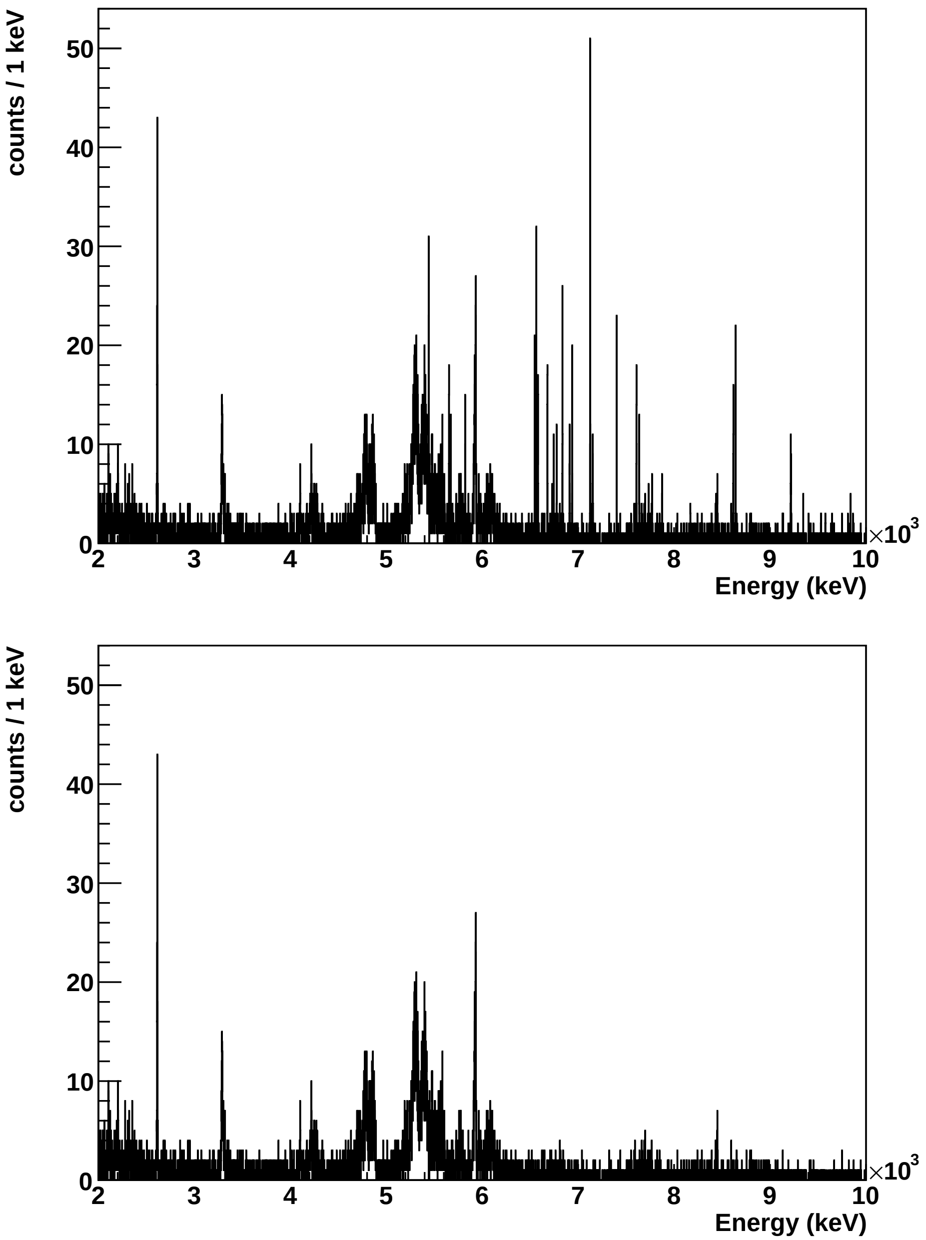

Figure 4.4: The spectrum of data set 32 before (top) and after (bottom) application of the offline heater identification algorithm. The sharp peaks above 5000 keV in the top figure are due to heater events that were not flagged by the data acquisition. They are effectively removed by the offline heater identification algorithm, except for a heater peak just below 6000 keV, due to channel 66, which is not eliminated by the offline algorithm because the events were contained in acquisition windows that triggered before the start of the heater pulses, as shown in Fig. 4.5.

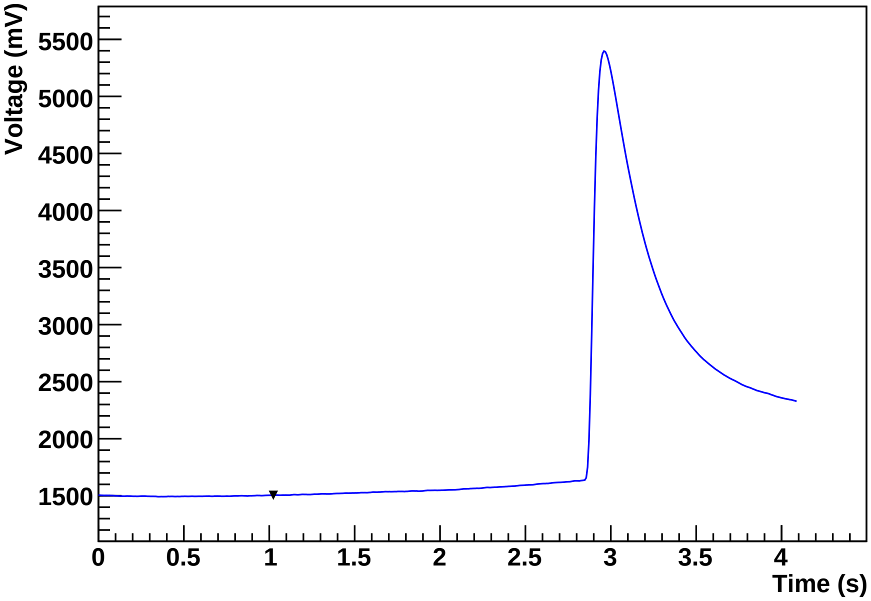

Figure 4.5: A pulse from the unidentified heater peak just below 6000 keV in Fig. 4.4. The pulse, from channel 66, was contained in the acquisition window of a previous trigger before the pulse itself could trigger the DAQ. The pulse itself does not cause the trigger to fire because the voltage level continually rises between the trigger position (indicated by the triangle) and the start of the pulse; the trigger will only fire again after the voltage level reaches a local maximum or at least 4 s elapse since the previous trigger.

4.7 Gain stabilization

The gain of a bolometer depends on its temperature, primarily because of the highly nonlinear dependence of the thermistor resistance on temperature and also because of the temperature dependence of the heat capacities and thermal conductances of the detector components. Since the temperature of the detector drifts over time, to maintain the optimal energy resolution, it is necessary to stabilize the gain against variations induced by a varying temperature [67]. The stabilization heater pulses are used to monitor the gain and establish its dependence on temperature. The pre-pulse baseline level, b, of the detector provides a measure of the detector temperature at the time of the pulse. The gain stabilization procedure consists of plotting the heater pulse amplitude vs. baseline (Fig. 4.6) and fitting a line to the points. For CUORICINO it is observed empirically that a straight line provides a good fit even when the baseline changes by a large amount. The best fit line establishes the heater amplitude, $a_h(b)$, as a function of baseline. The heater amplitude is related to the (constant) energy of the heater event, E_h, by the baseline-dependent gain, $G(b)$:

$$a_h(b) = G(b)E_h. \tag{4.8}$$

The assumption underlying the stabilization method is that the gain, $G(b)$, is independent of the amount of energy deposited in the crystal. Empirically, this assumption is at least approximately true, and any small deviations can be corrected by the calibration function. Once the heater amplitude trend is established, it is used to factor out the gain from the amplitude of a particle pulse. Considering a pulse of energy E, amplitude a, and baseline b,

$$a = G(b)E = a_h(b)\frac{E}{E_h}. \tag{4.9}$$

Therefore, the ratio $a/a_h(b)$ is a constant function of baseline. It is proportional to the energy of the underlying event, up to small nonlinear corrections to be established by calibration. The stabilized amplitude is defined as

$$a_s \equiv a/a_h(b) \times 5000. \tag{4.10}$$

The scaling by 5000 is done for convenience so that one unit of stabilized amplitude is $\mathcal{O}(1\text{ keV})$.

It sometimes happens that the amplitude vs. baseline trend of the heater events abruptly changes during a run. The reason this happens is not completely understood but may be due to a shift in the offset of the front-end electronics or due to a higher order effect of the temperature of part of the tower suddenly changing. An algorithm was developed to identify multiple trends in amplitude

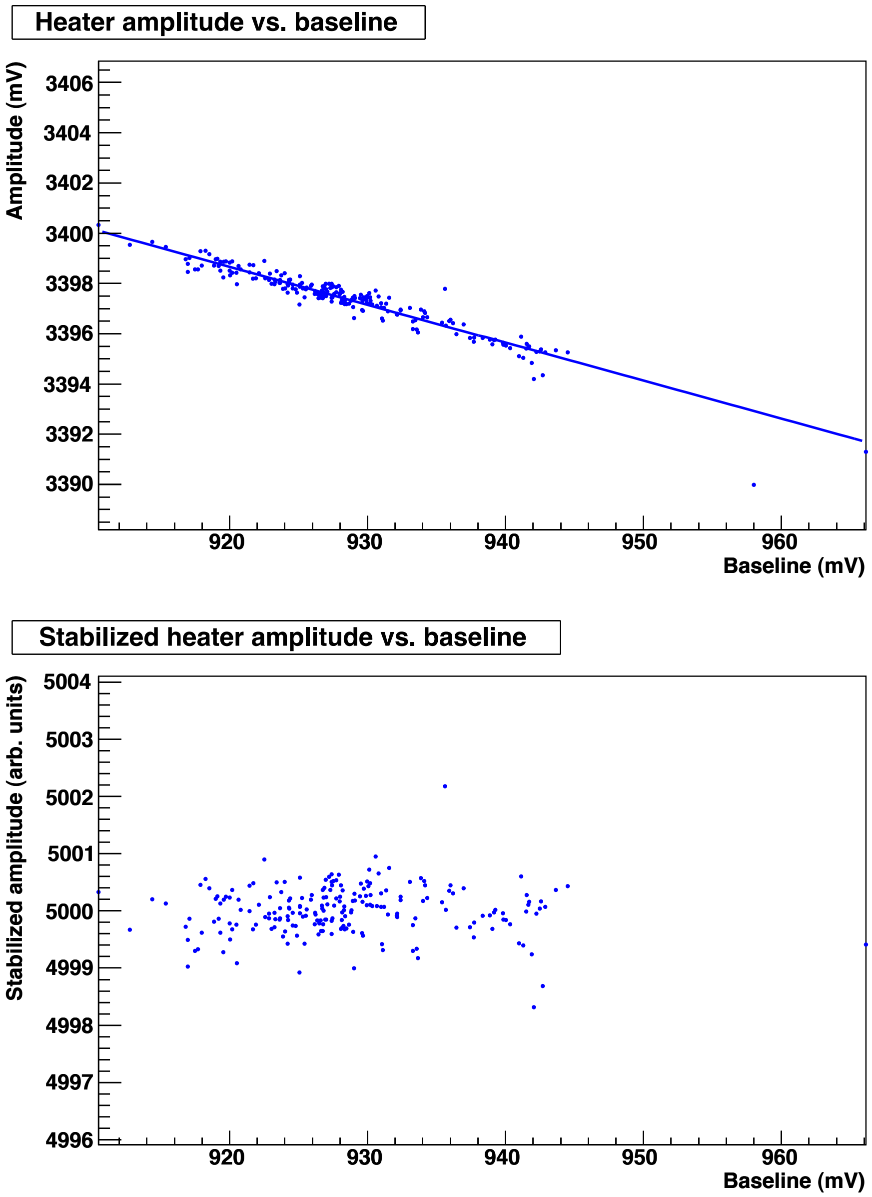

Figure 4.6: Example of a stabilization fit (channel 8, run 1214).

vs. baseline within the same run. The algorithm operates on one channel at a time and takes a brute-force approach. It first performs a single linear fit to the amplitude vs. baseline graph of heater points and records the sum of squared errors from the best fit line. Let $SSE(0)$ denote the sum of squared errors from the fit without splitting the run. It then tries splitting the run into two contiguous segments. It finds the best time to split the run by attempting to divide it at every possible time (that is, the time of every heater pulse); the best time to divide the run is the time such that the total sum of squared errors from the two best fit lines in the two segments is minimized. Let $SSE(1)$ denote the total sum of squared errors for the best single split. Similarly, the algorithm is allowed to try two and three splits, finding the best times to divide the run into three or four contiguous segments and computing $SSE(2)$ and $SSE(3)$. To decide how many splits the data actually prefer, for each number of splits an error on the points is estimated by setting the χ^2 per degree of freedom equal to 1 and solving for σ:

$$\begin{aligned} \frac{\chi_i^2}{N-2(i+1)-i} &= \frac{SSE(i)}{\sigma_i^2}\frac{1}{N-2(i+1)-i} = 1 \\ \Rightarrow \sigma_i^2 &= \frac{SSE(i)}{N-2(i+1)-i}, \end{aligned} \tag{4.11}$$

where N is the number of points. The smallest of the σ_i values, $\sigma_{\min}$, is taken as the estimate for the error on the points. Then, χ_i^2 for $i = 0, 1, 2, 3$ is recomputed using $\sigma_{\min}$ as the error: $\chi_i^2 = SSE(i)/\sigma_{\min}^2$. Also, the p-value associated with each of these χ^2 values with $N - 2(i + 1) - i$ degrees of freedom is computed; that is, the probability of obtaining a larger value of χ^2 by chance. The number of times to split the run is chosen to be the minimum number with a p-value greater than 1%. Figure 4.7 shows an example of a run that was split twice for the stabilization fit.

In the Three Towers Test, 10 of the 26 active bolometers did not have functioning heaters. Therefore, an alternative mechanism was employed to reduce or eliminate the gain variation. This mechanism is described in Sect. 4.10.

4.8 Energy calibration

Calibration sources were inserted near the detector approximately once per month in order to obtain the calibration function that maps the stabilized amplitude of a pulse to the energy deposited in the detector. The sources consisted of two thoriated tungsten wires that were inserted on opposite sides of the detector between the external lead shields and the cryostat. Each calibration typically lasted about two days. Figure 4.8a shows the spectrum of stabilized amplitudes recorded by a single CUORICINO channel with the calibration sources in place.

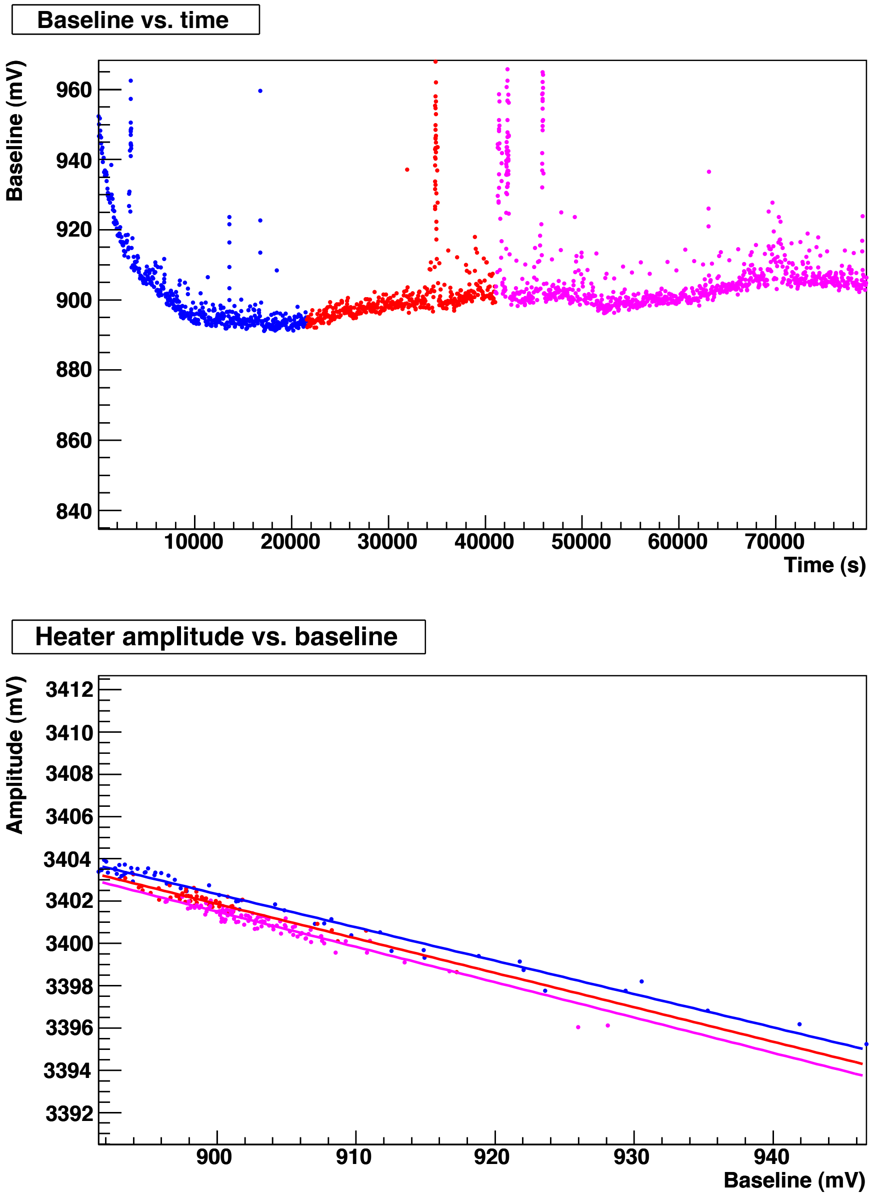

Figure 4.7: Example of a run split twice for stabilization with best fit lines (channel 8, run 1209). The colors in the plot of baseline vs. time indicate the intervals into which the run is split and correspond to the colored points and lines in the plot of heater amplitude vs. baseline.

The peaks in Fig. 4.8a are identified as gamma lines from the decay of nuclei in the ^{232}Th decay chain. The most intense gamma lines, listed in Table 4.3, are used to establish a set of points, $\{(a_s, E)\}$, which the calibration function should pass through; these points are represented by the red dots in Fig. 4.8b.

For CUORICINO we parameterize the calibration function in two ways. The standard parameterization used in this analysis is a third-order polynomial,

$$E = c_0 + c_1\, a_s + c_2\, a_s^2 + c_3\, a_s^3. \tag{4.12}$$

The third-order polynomial is viewed as the leading terms of the Taylor series expansion of the true calibration function. Since the calibration function is close to linear in the energy region up to the highest calibration peak at 2614.5 keV, including up to third-order terms in the parameterization provides a good approximation to the true calibration function up to 2614.5 keV. However, the third-order polynomial function does not extrapolate well up to the energies of alpha particles, which are greater than 4000 keV. The other parameterization that we compute for the calibration function is a second-order log polynomial defined by

$$\ln(E/(1\text{ keV})) = c_0 + c_1 \ln a_s + c_2\, (\ln a_s)^2, \tag{4.13}$$

which has been found phenomenologically to produce good energy estimates when extrapolated up to the energies of alpha particles. For the Three Towers Test, since we apply the thermal response transformation (discussed in Sect. 4.10), which removes the nonlinearities due to the thermistor, we parameterize the calibration function as a second-order polynomial.

For each channel and data set, the calibration coefficients, c_i, are calculated by an automatic code developed by collaborator Laura Kogler that determines the locations of the calibration peaks in the stabilized amplitude spectrum and fits each parameterization of the calibration function to the collection of (a_s, E) points. The algorithm divides the calibration peaks into two classes, primary and secondary, as indicated in Table 4.3. The primary peaks are the strongest and are clearly present in virtually every calibration spectrum. The secondary peaks are weaker and may or may not be discernible in a given calibration spectrum, depending on the length of exposure, relative positions of the sources and the crystal, and the energy resolution of the channel. The code first locates the primary peaks either by an automated peak search or by inverting an existing calibration function supplied by the user. After the primary peaks are found, the locations of the secondary peaks are estimated by interpolation. Then, an unbinned maximum likelihood fit is performed to each peak with the line shape listed in Table 4.3 plus a linear background. The line shapes are Gaussian except for two cases of nearby peaks separated by approximately the detector resolution, which are fit jointly with a double Gaussian line shape. Examples of the line shape

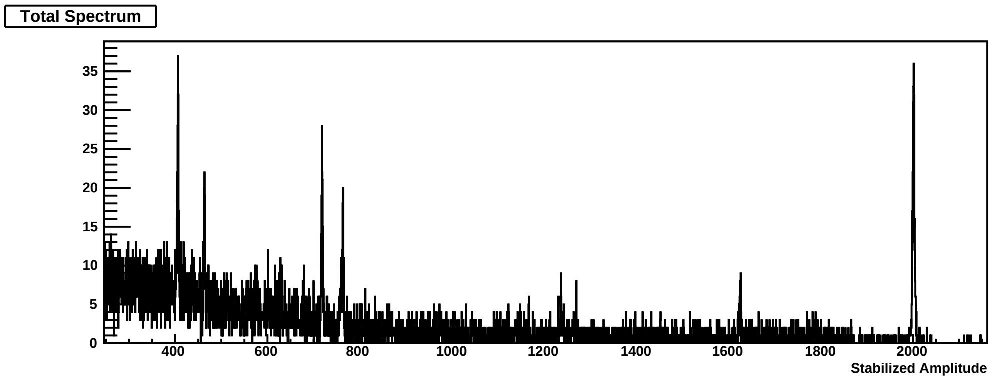

(a) Stabilized amplitude spectrum of calibration sources recorded by one channel.

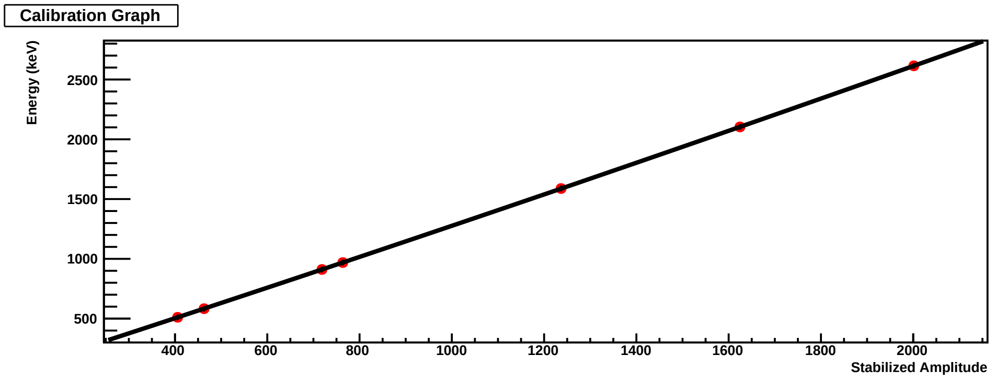

(b) Calibration function fit to graph.

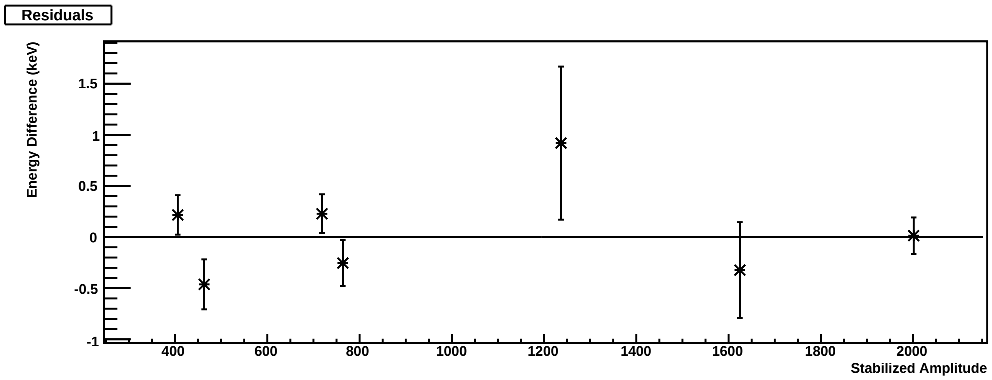

(c) Difference between the calibration function evaluated at the stabilized amplitude of each calibration peak and the known energy of the peak.

Figure 4.8: Calibration spectrum, calibration fit, and fit residuals. Figure from Laura Kogler.

Table 4.3: Gamma lines in the ^{232}Th decay chain used for calibration.

Source	Energy (keV)	Line shape	Classification
^{208}Tl	2614.5	Gaussian	Primary
2614.5 keV single escape	2103.5	Gaussian	Secondary
2614.5 keV double escape ^{228}Ac	1592.5 1588.2	Double Gaussian	Secondary
^{228}Ac ^{228}Ac	969.0 964.8	Double Gaussian	Primary
^{228}Ac	911.2	Gaussian	Primary
^{208}Tl	583.2	Gaussian	Secondary
e^+e^- annihilation ^{208}Tl	511.0	Gaussian	Primary

fits to the calibration peaks are shown in Fig. 4.9. After the peak fits have been performed, a graph is produced of the known energy of each peak versus the mean stabilized amplitude returned by the line shape fit; Figure 4.8b is an example of one such graph. Calibration peaks that fail statistical significance criteria are omitted from the graph. The points in the graph include horizontal error bars, which represent the error on the mean returned by the fitter. The calibration coefficients, c_i, are computed by fitting the parameterization of the calibration function to the graph.

As mentioned above, this analysis uses energies calibrated with the third-order polynomial parameterization of the calibration function for CUORICINO. However, the third-order polynomial calibration was unsuccessful for 17 (channel, data set) pairs, which constitute approximately 1% of the data, where the second-order log polynomial calibration succeeded. For these few cases, the energies from the second-order log polynomial calibration are used instead of the third-order polynomial calibration.

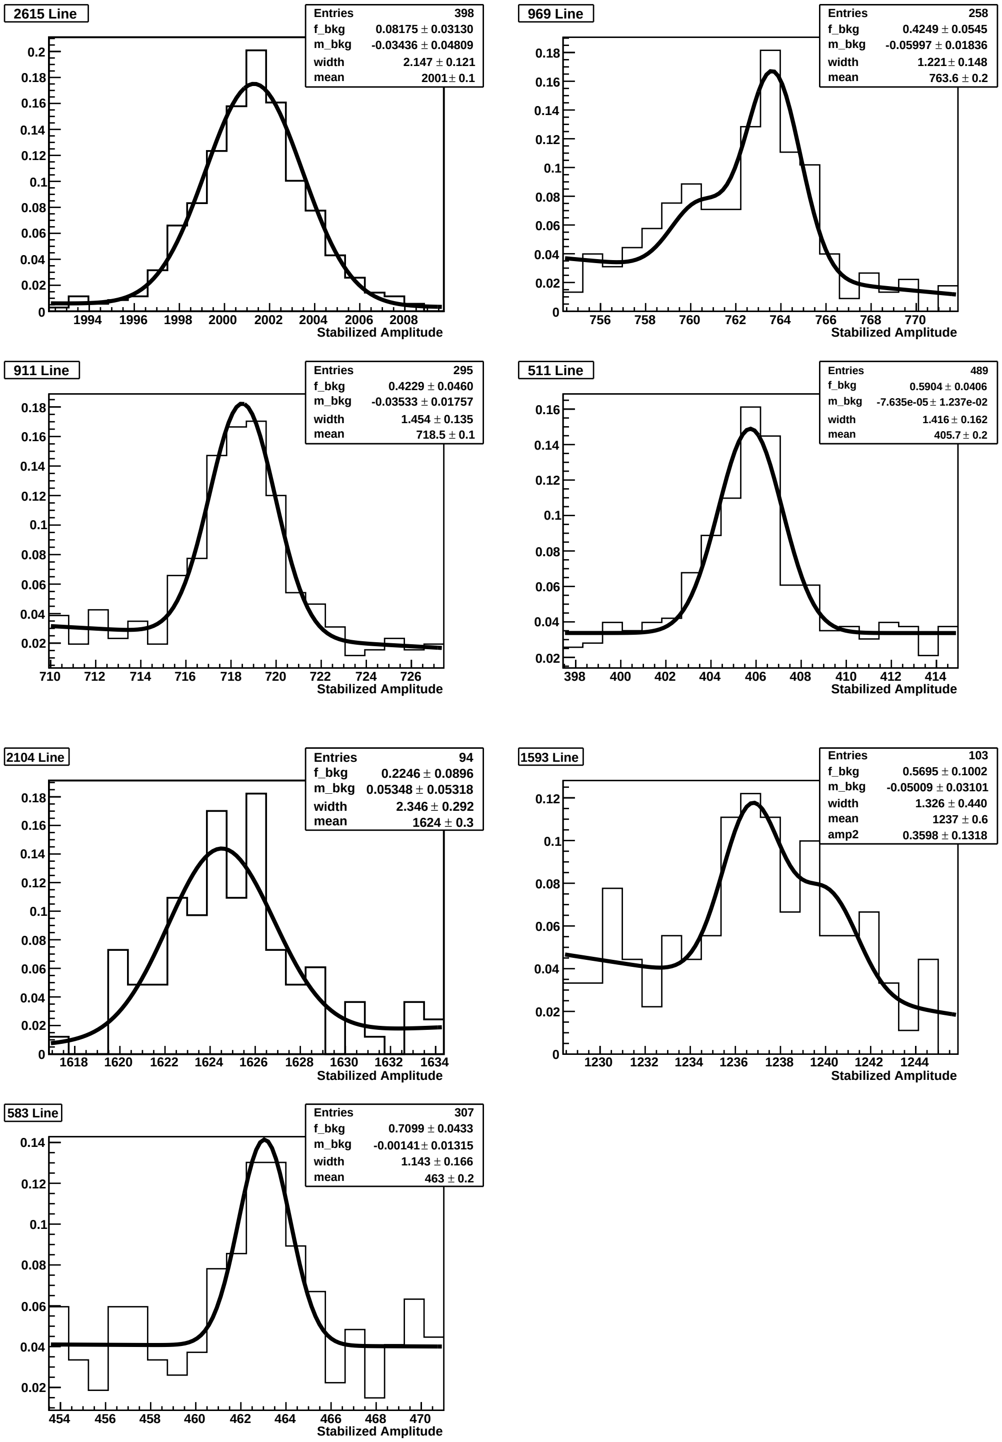

Figure 4.9: Calibration peaks with line shape fits. The plot titles indicate the energies of the peaks. Figure from Laura Kogler.

4.9 Pulse shape discrimination

Some of the triggered events are spurious signals, which include radioactive decays inside the thermistor, vibration-induced rapid temperature variations, electronics crosstalk, and pileup of multiple pulses. Examples of spurious signals are shown in Fig. 4.10. The energies of these spurious signals are typically evaluated to be small (less than $\sim$100 keV), but in rare cases their energies may be evaluated in the region of interest for $0\nu\beta\beta$ decay. In a low background experiment, the contribution of spurious events to the total background level may not be negligible. For CUORICINO, we discriminate between real physical events and spurious signals based on pulse shape parameters using a technique developed by collaborator Marco Carrettoni.

The pulse shape parameters that are considered are the rise time (t_r), decay time (t_d), and the ratio $R_{\mathrm{OF}} \equiv a_{\mathrm{TD}}/a_{\mathrm{FD}}$, where a_{TD} and a_{FD} are the pulse amplitude calculated by the optimal filter in the time domain and in the frequency domain, respectively. All physical pulses from a single detector have approximately the same rise time and decay time, except for a slight energy dependence that becomes most evident at high energies, as shown in Fig. 4.11. The energy-dependent trends of the rise time and decay time are fitted with second order polynomials, $f_{t_r}(E)$ and $f_{t_d}(E)$, respectively. The fit employs a robust regression method that is less sensitive to outliers than a least squares fit. The rise and decay times are then transformed to linearized, energy-independent quantities, $t_r^{\mathrm{lin.}} \equiv t_r/f_{t_r}(E)$ and $t_d^{\mathrm{lin.}} \equiv t_d/f_{t_d}(E)$. Physical pulses have $t_r^{\mathrm{lin.}} \approx t_d^{\mathrm{lin.}} \approx 1$.

For physical pulses, the ratio R_{OF} is close to 1 since the shape of physical pulses is close to the shape of the average pulse used in the optimal filter calculation. This ratio deviates from 1 for spurious signals since their shape does not match the shape of the average pulse. However, this ratio has a complicated energy dependence due to a subtle variation of the pulse shape with energy. In a similar manner as for the rise and decay times, R_{OF} is linearized by mapping out the main trend, $f_{R_{\mathrm{OF}}}(E)$, and transforming to $R_{\mathrm{OF}}^{\mathrm{lin.}} \equiv R_{\mathrm{OF}}/f_{R_{\mathrm{OF}}}(E)$; three examples are shown in Fig. 4.12. The algorithm implemented by Mr. Carrettoni to establish $f_{R_{\mathrm{OF}}}(E)$ in an automated fashion for a wide variety of possible trends and in the presence of many outliers is one of the keys to the success of this method.

In order to set a cut on $R_{\mathrm{OF}}^{\mathrm{lin.}}$ to select physical pulses, the mean and standard deviation of $R_{\mathrm{OF}}^{\mathrm{lin.}}$ for physical pulses are estimated as a function of energy. First, the events are binned in equal steps in $\ln(E/1\ \mathrm{keV})$. Then, for each bin, a histogram of the values of $R_{\mathrm{OF}}^{\mathrm{lin.}}$ is produced; the histogram is expected to contain a peak centered at 1 due to the physical pulses and a continuum background due to spurious signals. Each histogram is fit with a Gaussian to estimate the mean and standard deviation of $R_{\mathrm{OF}}^{\mathrm{lin.}}$ for physical pulses in the corresponding energy bin. The parameter used to make a cut for selecting physical pulses and rejecting spu-

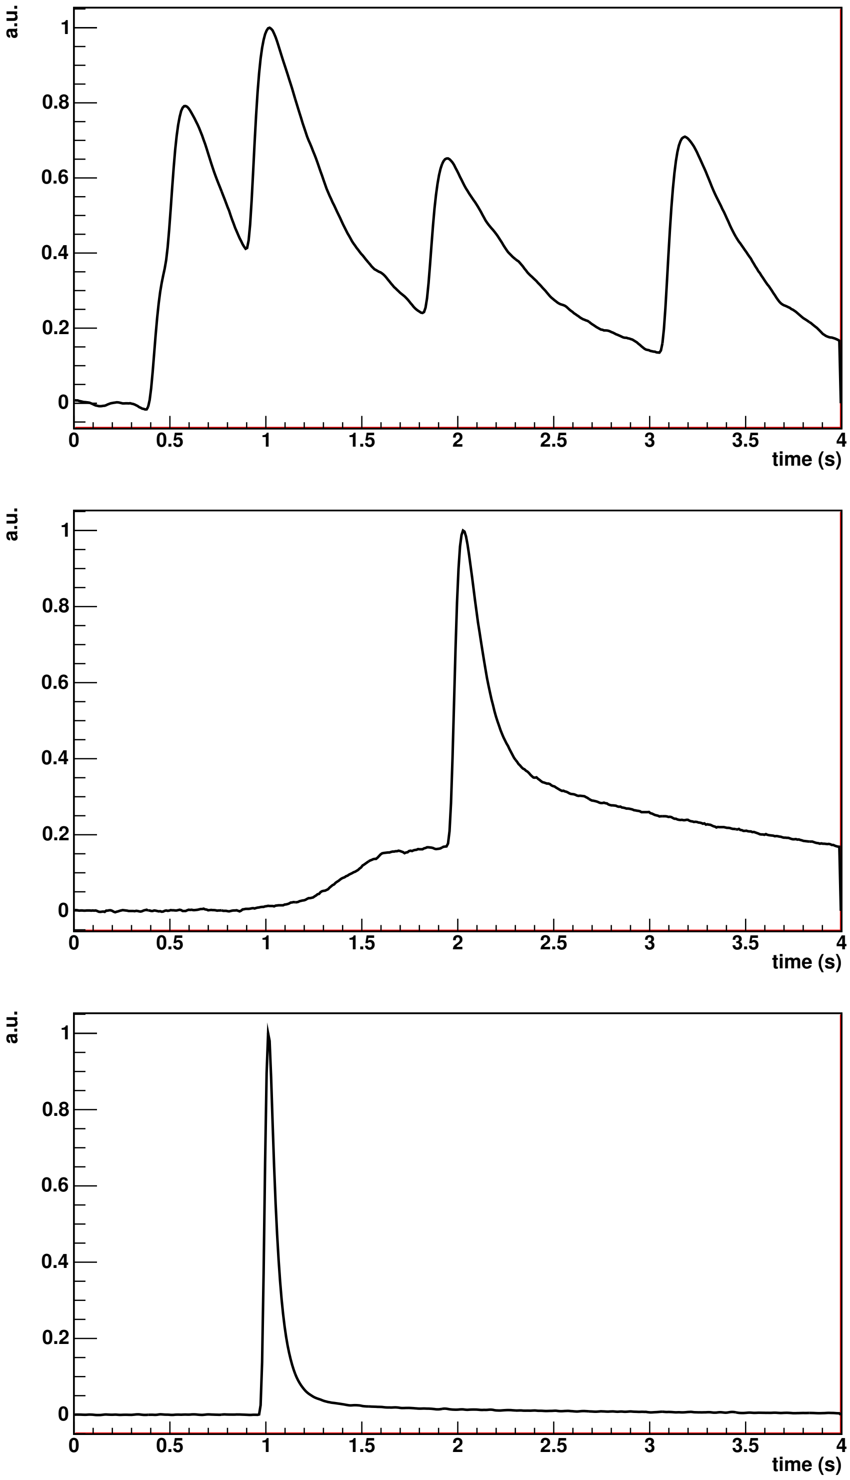

Figure 4.10: Examples of spurious signals: pile-up (top), pileup with a non-particle event (middle), and a spike possibly due to a radioactive decay in the thermistor (bottom). Figure from Marco Carrettoni.

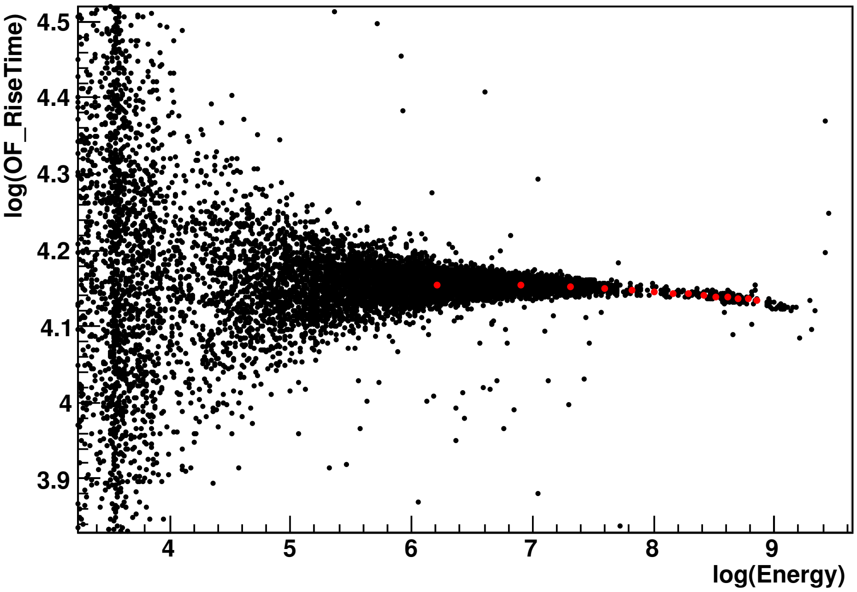

Figure 4.11: Scatter plot of the logarithm of the rise time versus the logarithm of the energy. The energy range is 0–10 MeV, and the red dots identify intervals of 500 keV. The nonlinearity becomes most pronounced at high energies. Figure from Marco Carrettoni.

rious ones is the number of sigma deviation from the mean, $n_\sigma^{\mathrm{PS}} \equiv (R_{\mathrm{OF}}^{\mathrm{lin.}} - \mu)/\sigma$, where μ and σ are the mean and sigma of the Gaussian fit, respectively. An example of the result of this procedure is shown in Fig. 4.13.

The pulse shape cut used in this analysis is a logical combination of cuts on several variables. The first requirement is a simple cut on the time of the pulse maximum for the pulse whose amplitude is evaluated by the optimal filter; the pulse maximum is required to be less than 1.4 s from the beginning of the acquisition window to ensure that enough of the pulse shape is contained in the acquisition window for an accurate estimation of the amplitude. A very loose anti-coincidence requirement is included in the pulse shape cut to eliminate a class of spurious heater events caused by noise pickup in the heater circuit that resulted in almost every detector receiving simultaneous impulses from the heaters: The number of coincident events with energies greater than 80 keV within a 10 ms time window is required to be less than or equal to 7. The linearized rise and decay times are required to be within 20% of their nominal values: $|t_r^{\mathrm{lin.}} - 1| < 0.2$ and $|t_d^{\mathrm{lin.}} - 1| < 0.2$. Finally, we require $n_\sigma^{\mathrm{PS}} < 2$, which is the most powerful component of the pulse shape cut. The efficiency of this combination of cuts for $0\nu\beta\beta$ decay signal events is evaluated in Sect. 6.2.2.

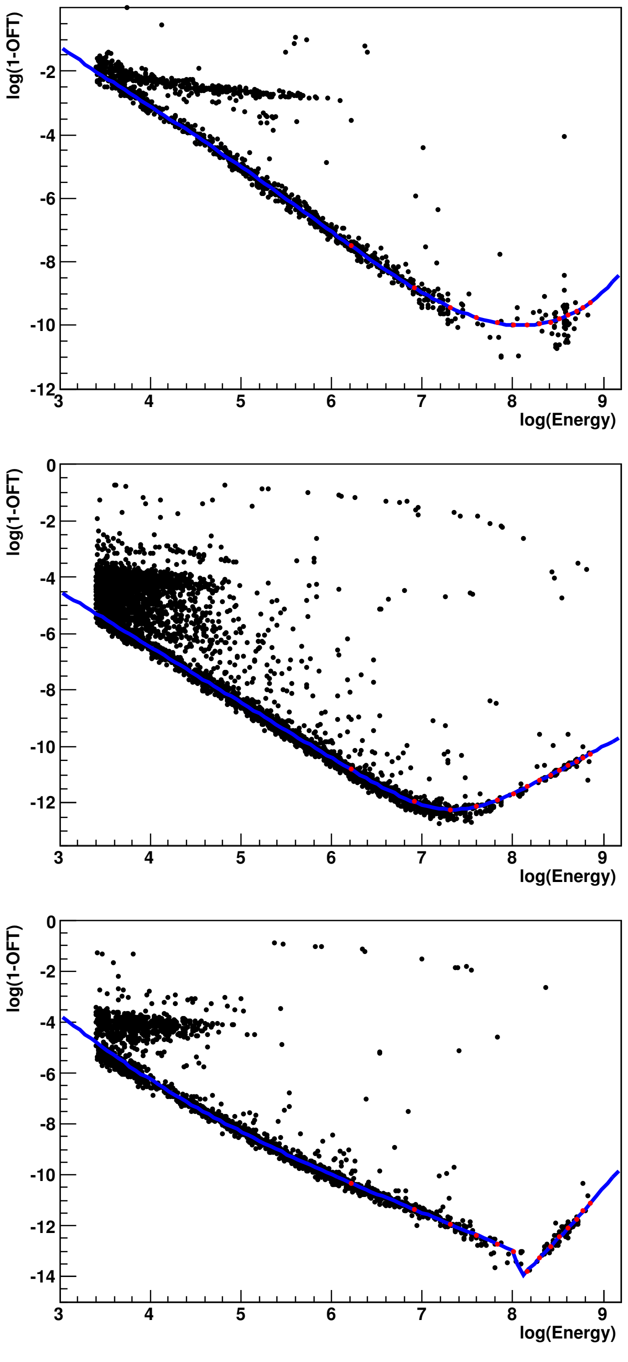

Figure 4.12: Examples of the R_{OF} parameter (denoted OFT in the plots) energy dependence and the functions $f_{R_{\mathrm{OF}}}(E)$ (blue curves) used in the linearization. The energy range is 30–7000 keV, and the red dots identify intervals of 500 keV. Outliers are most likely pileups and vibration-induced rapid temperature variations, while the high density regions far from the main trends are spikes or low energy noise events. Figure from Marco Carrettoni.

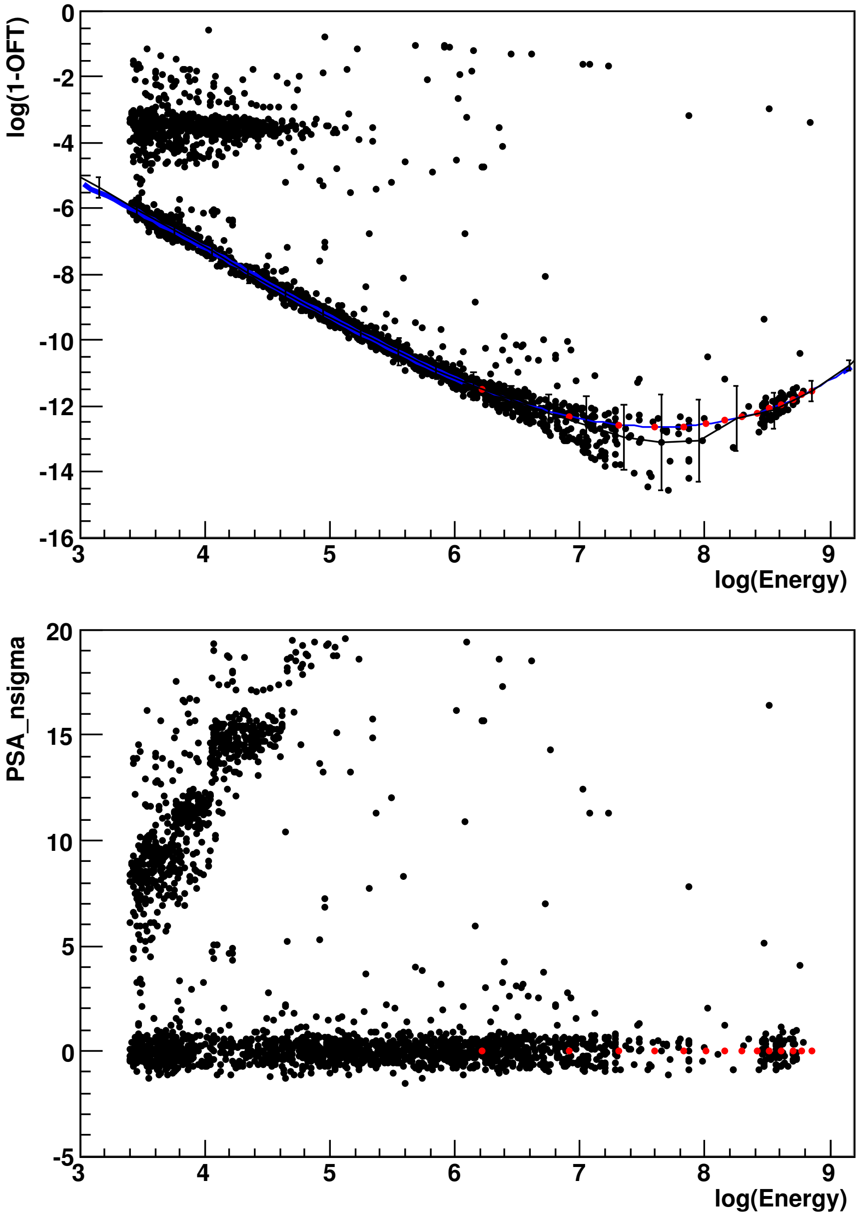

Figure 4.13: Example of the process of linearization and normalization of the R_{OF} parameter (denoted OFT in the plot). The energy range is 30–7000 keV, and the red dots identify intervals of 500 keV. Top: The scatter plot is analyzed in energy bins of equal size in $\ln(E/1\ \mathrm{keV})$; the points with vertical error bars indicate mean and standard deviation computed by the method described in the text. Bottom: The distribution of n_σ^{PS} as a function of energy. Figure from Marco Carrettoni.

4.10 Thermal response transformation

For the Three Towers Test, we applied an additional processing step at the very beginning of the first-level processing sequence. Because 10 of the 26 active crystals did not have a working heater, we needed an alternative method to correct for the gain instabilities due to temperature fluctuations of the cryostat. For the first TTT data set, ^{40}K sources were placed in the positions of the usual calibration sources so that the 1461 keV gamma line of ^{40}K could be used in place of the heaters for the stabilization fits. Unfortunately, with the ^{40}K sources being relatively far from the crystals, outside of the lead shield, the continuum of Compton scattered gammas below 1461 keV was large compared to the photopeak, and relatively strong sources were required to produce enough events in the photopeak. Because the high event rates produced by the ^{40}K sources could obscure some of the features of surface contaminations of the copper parts, including low energy gammas from the ^{238}U and ^{232}Th decay chains, we decided to remove the ^{40}K sources after the first TTT data set and to use a newly developed analysis technique to mitigate the gain instabilities induced by varying temperature.

The technique, developed by collaborator Marco Vignati, transforms the voltage samples recorded by the ADC into “thermal response” (TR) samples proportional to the thermistor temperature variation [81]. The thermal response transformation deconvolves the effect of the thermistor and the electrical properties of its biasing circuit. It is shown in Section 3.2.8 of Ref. [81] that the dominant cause of the pulse amplitude variation with detector temperature is the nonlinear response of the thermistor, despite the dependence of intrinsic thermal properties of the detector, such as the heat capacity, on temperature. For the TTT analysis, we replaced the thermistor voltage samples with TR samples as the first step in the data processing. An example of a TR transformed pulse is shown in Fig. 4.14. The details of the transformation, which is rather complicated, are described in Chapters 4 and 5 of Ref. [81]. We applied the optimal filter as usual to the TR samples but did not do any gain instabilities correction (except for the first data set, which we stabilized with the ^{40}K source).

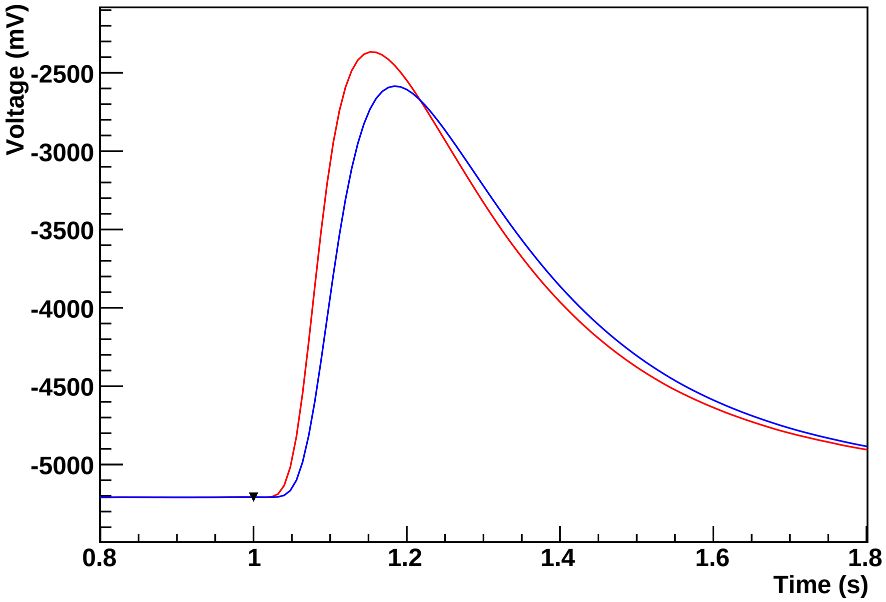

Figure 4.14: Comparison between a pulse before (blue) and after (red) the thermal response transformation (channel 7, run 100294, event #693).

Chapter 5

Data quality checks and data selection criteria

This chapter describes the data selection criteria and the checks that were performed to ensure that the detector was operating normally while the data used in the analysis were acquired. Of particular concern is the stability of the detector resolution. Following the steps described here, we identified periods when the detector was experiencing a problem that degraded the resolution or caused an artificial dead time. Data acquired during problematic periods were excluded from the analysis.

The first level of data quality checks were performed using the CUORE Online/Offline Run Check (CORC) system, which is a tool that we developed to monitor the detector operation and data quality. For every run we checked the CORC output described in Sect. 5.1. If a problem was found, then in the analysis database, we either marked the whole run as bad or marked a channel as bad for a certain time interval within the run. The types of bad runs and bad intervals are described in Sect. 5.2. The procedure we followed for the CORC checks is described in Sect. 5.3.1. After the CORC checks were completed, we reprocessed all of the data so that certain parts of the analysis code could ignore data in bad intervals if desired. For example, if a channel was saturated (pulses exceeding the maximum range of the ADC) for half of a run, the stabilization fit to establish the heater amplitude vs. baseline trend is performed on only the good half of the run. During the reprocessing, the events that occured in bad intervals were flagged so that they may be removed from the final analysis with a cut. After the reprocessing, we checked the consistency of the detector resolutions for every run using the heater events. The procedure for the resolution consistency check is described in Sect. 5.3.2. Finally, we checked the consistency of the calibrations taken at the beginning and end of each data set. The calibration consistency requirements are described in Sect. 5.3.3.

5.1 CUORE Online/Offline Run Check (CORC)

The purpose of CORC is to provide a web-based collection of standard plots and automated warnings for identifying and diagnosing problems with the detector. It is used by the on-site shifters during data taking and by collaborators working on analyses of the data. The current location of the CORC webpages is `http://cuoreweb.lngs.infn.it/~cuore/corc/`.

For each run several CORC webpages are generated. At the top of each page is a header that lists the data set(s) to which the run belongs, the number of days since the last calibration, the run type (background, calibration, or test), the `stopstatus` (which indicates whether the whole run has been flagged as bad), and a link to the database webpage for the run.

5.1.1 CORC pages

The pages currently available on CORC are:

History The History page displays summary data about each channel for the selected run and the previous nine runs in a tabular format. An example is shown in Fig. 5.1. The variables available to be displayed on the History page are listed in Table 5.1. The layout of the History page facilitates comparison of current values with recent historical values to allow for recognition of abnormal values and determination of when a problem began. An optional reference value with a minimum and maximum range may be specified. If a reference exists, it is displayed in the left-most column and the cells in the table are color-coded depending on whether the value is within (green), below (blue), or above (red) the reference range.

History Graph The History Graph page displays the same data as the History page in a graphical format. There is a page for each channel that displays a scatter plot for each of the summary variables in Table 5.1. The plot includes the previous twenty-four runs and gives a visual presentation of how the variable has been changing over time. Some examples of the plots are shown in Fig. 5.2.

Summary The Summary page facilitates identifying problematic channels in a particular run. It consists of bar graphs of the summary variables in Table 5.1. Each bar graph provides a quick look at the values for all the channels of one run and allows easy comparison among channels and comparison with reference ranges. The reference ranges are indicated with a gray band in the background. Examples of plots from the Summary plot page are shown in Fig. 5.3.

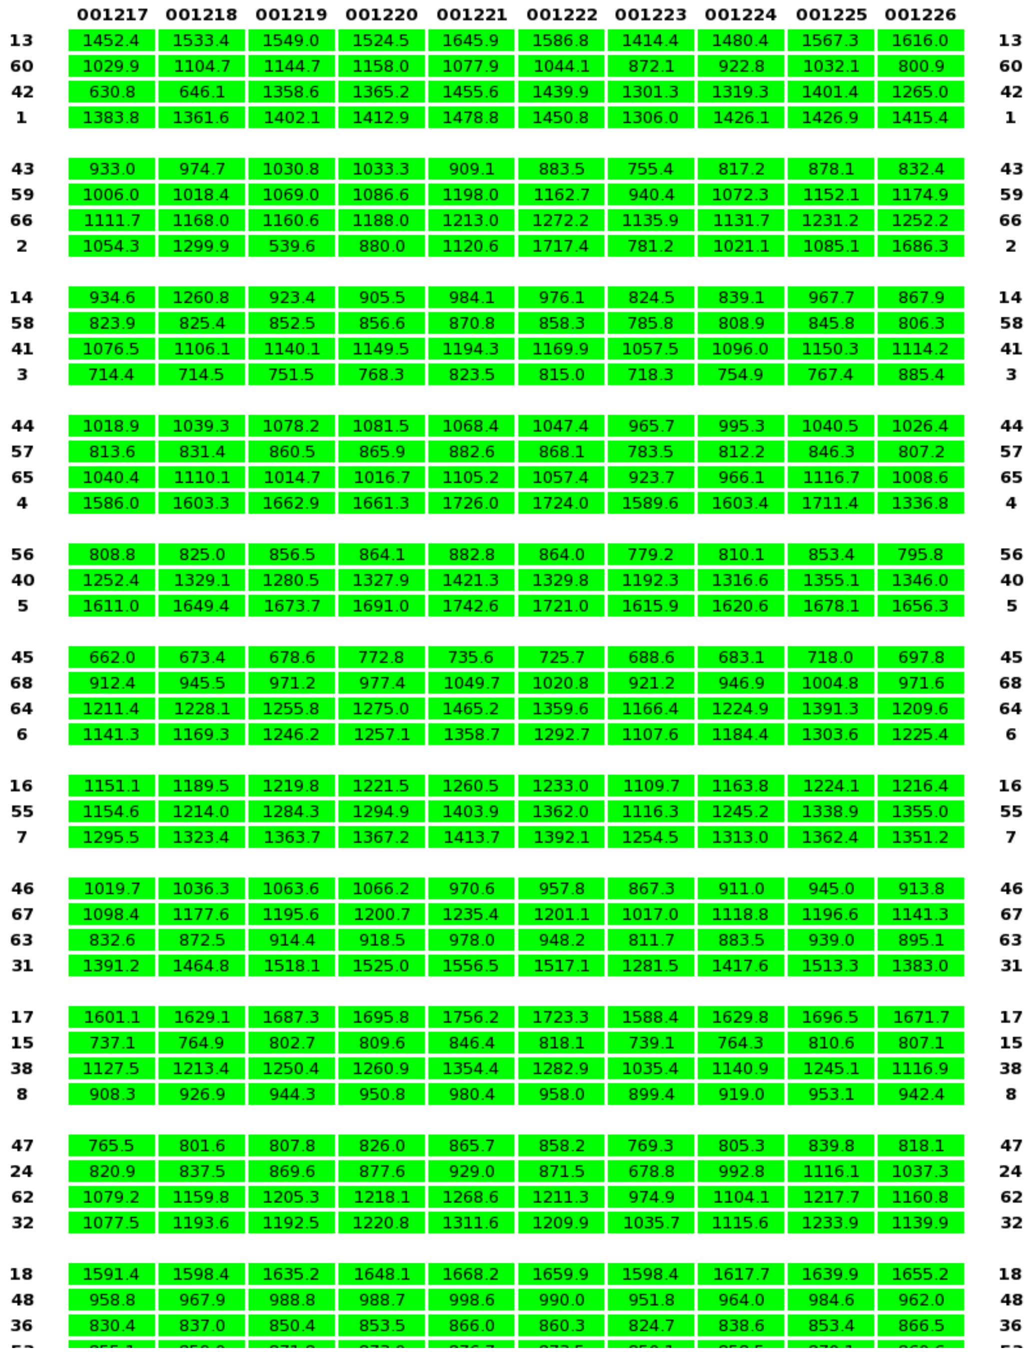

	001217	001218	001219	001220	001221	001222	001223	001224	001225	001226	
13	1452.4	1533.4	1549.0	1524.5	1645.9	1586.8	1414.4	1480.4	1567.3	1616.0	13
60	1029.9	1104.7	1144.7	1158.0	1077.9	1044.1	872.1	922.8	1032.1	800.9	60
42	630.8	646.1	1358.6	1365.2	1455.6	1439.9	1301.3	1319.3	1401.4	1265.0	42
1	1383.8	1361.6	1402.1	1412.9	1478.8	1450.8	1306.0	1426.1	1426.9	1415.4	1
43	933.0	974.7	1030.8	1033.3	909.1	883.5	755.4	817.2	878.1	832.4	43
59	1006.0	1018.4	1069.0	1086.6	1198.0	1162.7	940.4	1072.3	1152.1	1174.9	59
66	1111.7	1168.0	1160.6	1188.0	1213.0	1272.2	1135.9	1131.7	1231.2	1252.2	66
2	1054.3	1299.9	539.6	880.0	1120.6	1717.4	781.2	1021.1	1085.1	1686.3	2
14	934.6	1260.8	923.4	905.5	984.1	976.1	824.5	839.1	967.7	867.9	14
58	823.9	825.4	852.5	856.6	870.8	858.3	785.8	808.9	845.8	806.3	58
41	1076.5	1106.1	1140.1	1149.5	1194.3	1169.9	1057.5	1096.0	1150.3	1114.2	41
3	714.4	714.5	751.5	768.3	823.5	815.0	718.3	754.9	767.4	885.4	3
44	1018.9	1039.3	1078.2	1081.5	1068.4	1047.4	965.7	995.3	1040.5	1026.4	44
57	813.6	831.4	860.5	865.9	882.6	868.1	783.5	812.2	846.3	807.2	57
65	1040.4	1110.1	1014.7	1016.7	1105.2	1057.4	923.7	966.1	1116.7	1008.6	65
4	1586.0	1603.3	1662.9	1661.3	1726.0	1724.0	1589.6	1603.4	1711.4	1336.8	4
56	808.8	825.0	856.5	864.1	882.8	864.0	779.2	810.1	853.4	795.8	56
40	1252.4	1329.1	1280.5	1327.9	1421.3	1329.8	1192.3	1316.6	1355.1	1346.0	40
5	1611.0	1649.4	1673.7	1691.0	1742.6	1721.0	1615.9	1620.6	1678.1	1656.3	5
45	662.0	673.4	678.6	772.8	735.6	725.7	688.6	683.1	718.0	697.8	45
68	912.4	945.5	971.2	977.4	1049.7	1020.8	921.2	946.9	1004.8	971.6	68
64	1211.4	1228.1	1255.8	1275.0	1465.2	1359.6	1166.4	1224.9	1391.3	1209.6	64
6	1141.3	1169.3	1246.2	1257.1	1358.7	1292.7	1107.6	1184.4	1303.6	1225.4	6
16	1151.1	1189.5	1219.8	1221.5	1260.5	1233.0	1109.7	1163.8	1224.1	1216.4	16
55	1154.6	1214.0	1284.3	1294.9	1403.9	1362.0	1116.3	1245.2	1338.9	1355.0	55
7	1295.5	1323.4	1363.7	1367.2	1413.7	1392.1	1254.5	1313.0	1362.4	1351.2	7
46	1019.7	1036.3	1063.6	1066.2	970.6	957.8	867.3	911.0	945.0	913.8	46
67	1098.4	1177.6	1195.6	1200.7	1235.4	1201.1	1017.0	1118.8	1196.6	1141.3	67
63	832.6	872.5	914.4	918.5	978.0	948.2	811.7	883.5	939.0	895.1	63
31	1391.2	1464.8	1518.1	1525.0	1556.5	1517.1	1281.5	1417.6	1513.3	1383.0	31
17	1601.1	1629.1	1687.3	1695.8	1756.2	1723.3	1588.4	1629.8	1696.5	1671.7	17
15	737.1	764.9	802.7	809.6	846.4	818.1	739.1	764.3	810.6	807.1	15
38	1127.5	1213.4	1250.4	1260.9	1354.4	1282.9	1035.4	1140.9	1245.1	1116.9	38
8	908.3	926.9	944.3	950.8	980.4	958.0	899.4	919.0	953.1	942.4	8
47	765.5	801.6	807.8	826.0	865.7	858.2	769.3	805.3	839.8	818.1	47
24	820.9	837.5	869.6	877.6	929.0	871.5	678.8	992.8	1116.1	1037.3	24
62	1079.2	1159.8	1205.3	1218.1	1268.6	1211.3	974.9	1104.1	1217.7	1160.8	62
32	1077.5	1193.6	1192.5	1220.8	1311.6	1209.9	1035.7	1115.6	1233.9	1139.9	32
18	1591.4	1598.4	1635.2	1648.1	1668.2	1659.9	1598.4	1617.7	1639.9	1655.2	18
48	958.8	967.9	988.8	988.7	998.6	990.0	951.8	964.0	984.6	962.0	48
36	830.4	837.0	850.4	853.5	866.0	860.3	824.7	838.6	853.4	866.5	36

Figure 5.1: Example of a CORC History display of average baselines at the end of CUORICINO, cropped to fit the page. The channels are arranged by floor in the tower.

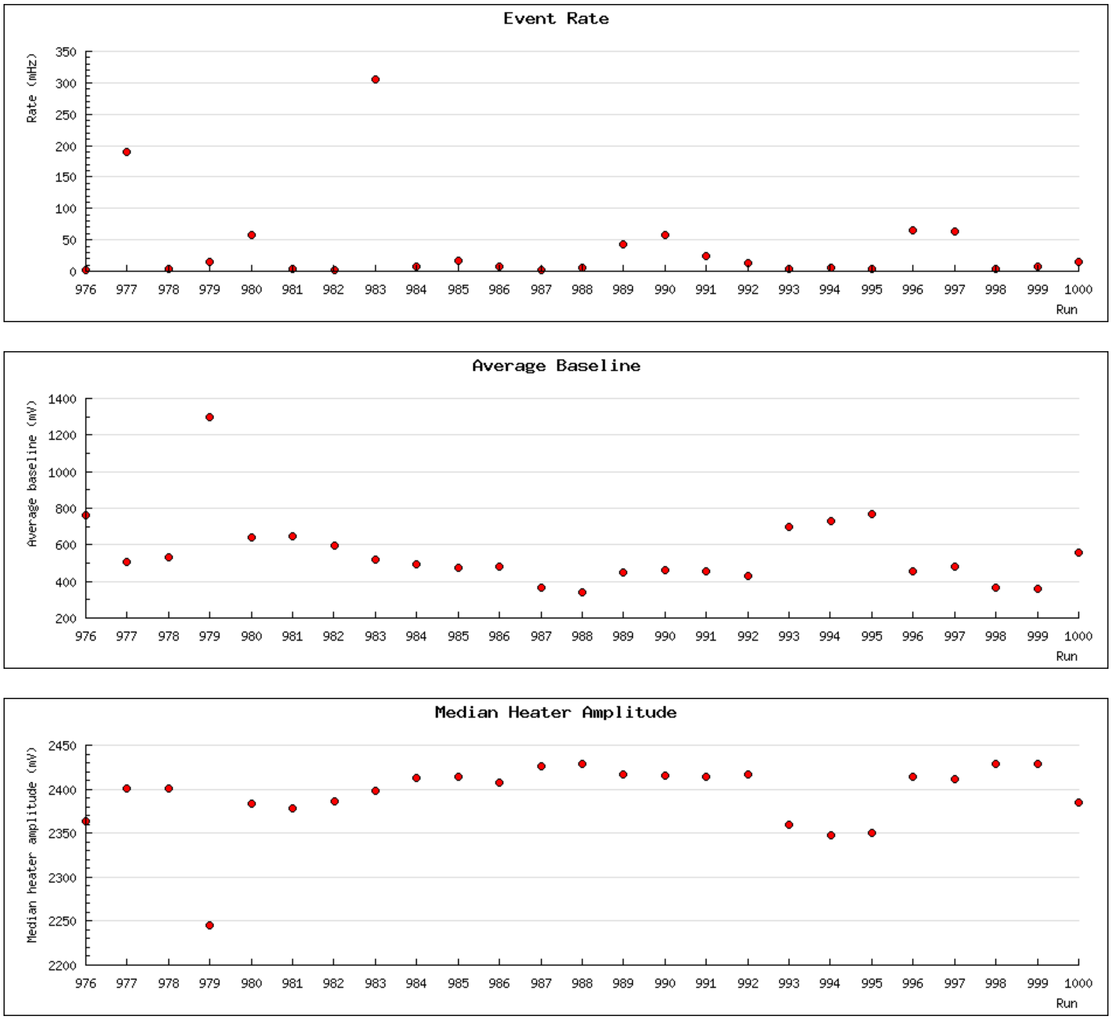

Figure 5.2: Examples of CORC History Graph plots for channel 1 and run 1000.

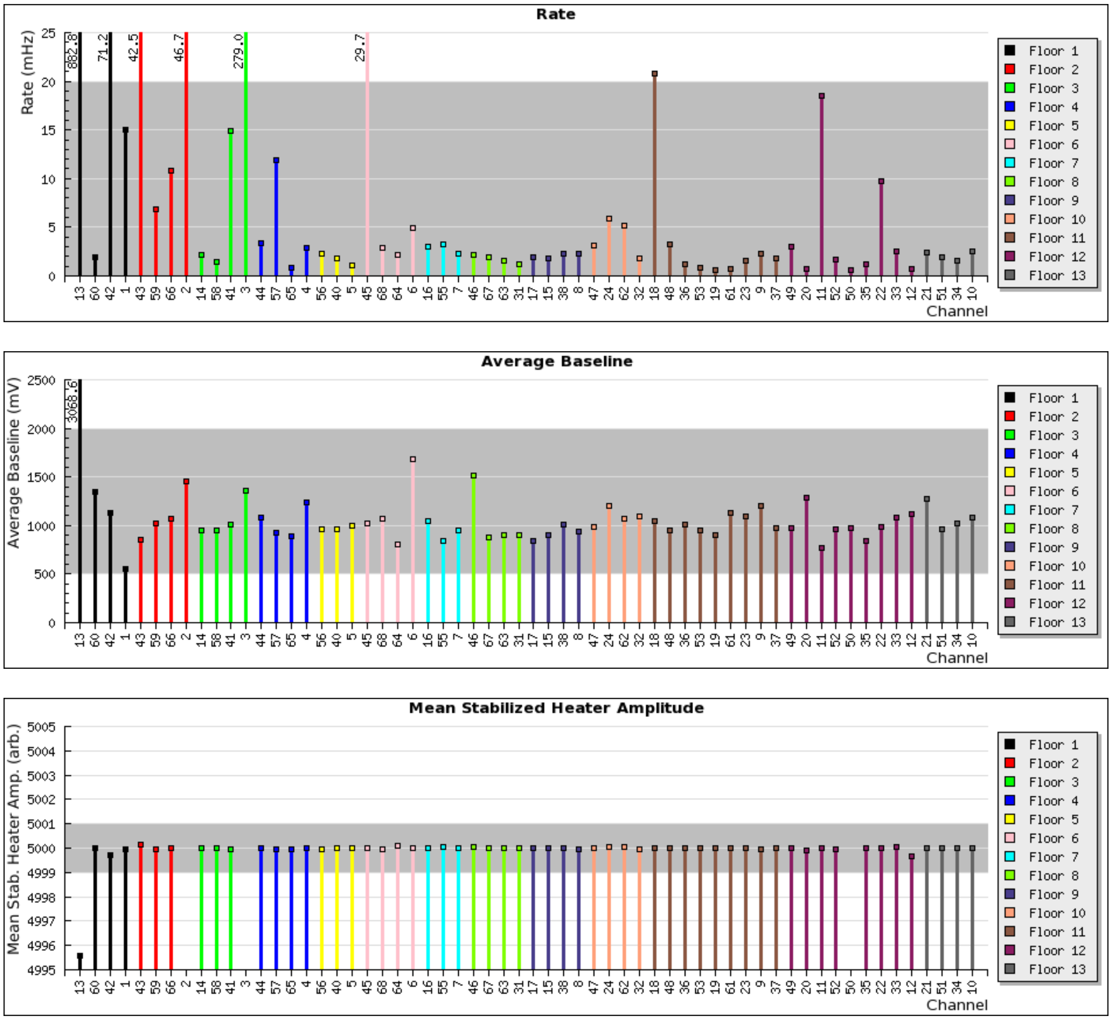

Figure 5.3: Examples of CORC Summary plots for run 1000.

Table 5.1: Summary variables available on CORC History, History Graph, and Summary pages.

Summary variable	Description
Rate	Trigger rate (mHz)
Average Baseline	Average pre-trigger level of all acquired pulses (mV)
Baseline Fluctuation	RMS spread of baselines of all acquired pulses (mV)
Average Baseline RMS	Average baseline RMS of all acquired pulses (mV)
Median Heater Amplitude	Median amplitude of stabilization heater pulses (mV)
RMS Stabilized Heater Amplitude	Gaussian σ of fit to histogram of stabilized heater amplitudes (arbitrary units)

Report The CORC Report page is where automatically generated warnings are presented to the user. Warnings are generated when a bolometer is saturated (i.e. the amplified thermistor voltage exceeds the range of the analog-to-digital converter), when there is a time hole in signal events greater than two hours, when there are coincident triggers on more than eight channels, and when the yield of a Gaussian fit to the stabilized heater amplitude peak deviates from the expected number of heater events by more than 5%. The automated warnings alert the on-site shifters to potential problems with the detector and help with the quality checks of data collected before the CORC system was in place.

Channel While the previously described CORC pages display summary information about each channel for each run, the Channel page presents a detailed picture of the behavior of each channel during a run. A CORC Channel page currently displays eight plots: baseline vs. time, baseline RMS vs. time, amplitude vs. time, heater amplitude vs. baseline, energy vs. time, stabilized heater amplitude vs. baseline, stabilized heater amplitude vs. time, and a histogram of stabilized heater amplitude. The Channel page can be used, for example, to observe noisy periods within a run, to check the normal operation of the heater, and to verify the success of the stabilization. Part of a Channel plot may have a gray background, which indicates an interval of time that was flagged as bad. The plot of heater amplitude vs. baseline also shows the stabilization fits. If a run was broken into multiple time intervals for the purpose of stabilization, all the Channel plots and the stabilization fits are color-coded by time interval. Clicking on one of the plots brings up a high resolution pdf version. Examples of plots from a Channel page are shown in Fig. 5.4.

Floor The CORC Floor pages display the same plots as the Channel pages but organized by the floor of the crystal in the tower rather than by channel number. Looking at the same kind of plots for all the channels on a floor can help identify and diagnose some problems, such as elevated noise levels, that tend to be correlated with position in the tower. The Floor page also provides a more efficient way to step through all the channel-by-channel plots of a certain type for the whole detector.

5.1.2 CORC technology and operation

To produce its output, CORC utilizes a collection of technologies. The first step is a set of DIANA modules that extract the relevant summary quantities from the data and produce the plots for the Channel page using ROOT [82, 83]. The CORC modules can be run in online mode or offline mode. The online mode is used while a run is in progress and the modules are running within the APOLLO DAQ process. The online mode is more restrictive in order to lessen the CPU demands on the DAQ computer, and the online mode cannot include quantities that require post-processing of the whole set of data collected during a run. Therefore, the online mode computes amplitudes with a simple maximum-minus-baseline algorithm and does not compute quantities requiring the stabilization step. After a run is finished, CORC is executed in offline mode for that run as part of the prompt data reconstruction code, which includes a preliminary application of the full data processing chain. The CORC summary quantities are stored in the CUORE analysis database, which is implemented in PostgreSQL and hosted on the computer `cuoredb.lngs.infn.it`. In offline mode, the CORC DIANA modules communicate directly with the database to write the summary quantities. In online mode, in order not to make the DAQ process dependent on the network connection to the PostgreSQL server, the summary quantities are written to a plain text file. The summary files and the channel-by-channel plots are updated and written out periodically, typically once per hour. A script on the computer `cuorehome.lngs.infn.it` periodically checks for updated CORC output on the DAQ machine, and when it finds new output, it copies the files to the web-accessible area, parses the text file with summary quantities, and fills the database with the values from the file.

The CORC webpages are written in PHP, a server-side scripting language.

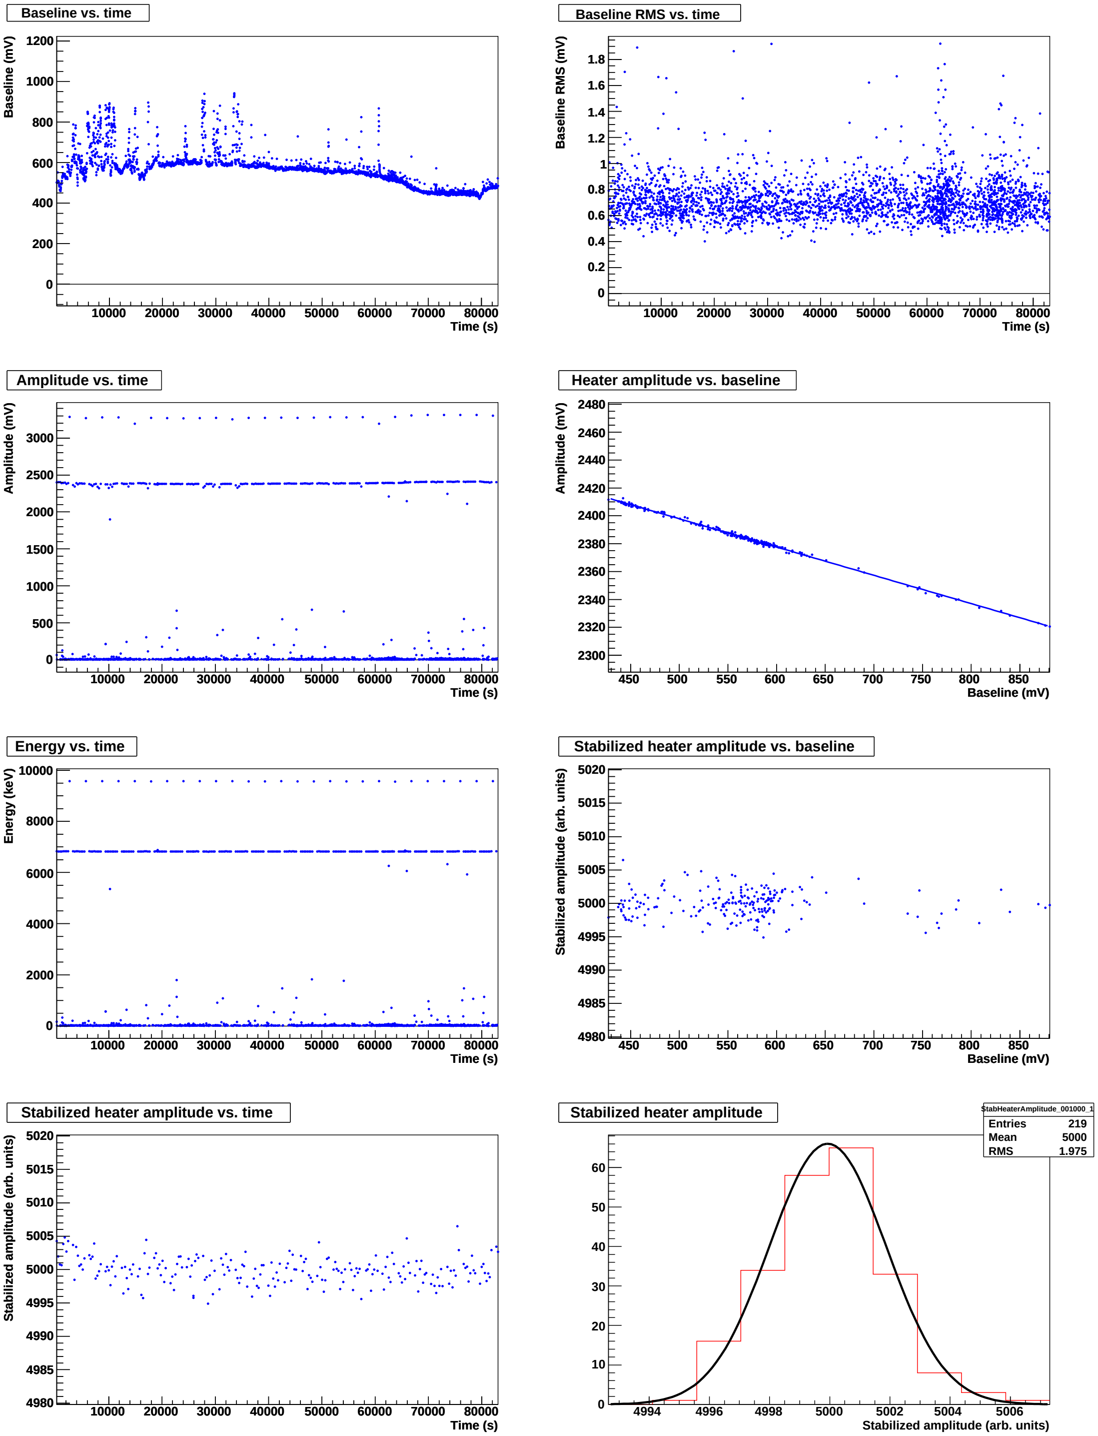

Figure 5.4: Examples CORC Channel plots for channel 1 and run 1000.

When a user accesses a CORC webpage, the PHP code communicates with the database to lay out the page with the appropriate selection of channels, depending on the detector, which is deduced from the first digit of the run number. The plots of summary quantities on the History, History Graph, and Summary pages are generated dynamically when the webpage is requested based on the data in the database. The dynamic generation of plots with PHP is implemented with the JpGraph package [84]. The database-driven nature of the CORC webpages makes it easy to update the output by updating the database; separate updating of the webpages is not required. Similarly, implementing new detector setups, such as Hall C tests, does not necessitate any modification to the CORC webpages because the layout is inferred from the detector information in the database, which is already required for the analysis software.

5.2 Bad runs and bad intervals

The infrastructure for keeping track of bad runs and bad time intervals is implemented in the analysis database. The `runs` table, where the basic information about each run is stored, contains a column `stopstatus` that serves as the flag that indicates whether the run is good or bad. The `stopstatus` is initially set by the on-site shifters. When the shifters terminate a run, the DAQ program asks if the run ended normally, and if it did not end normally, the program asks for the reason for the abnormal end. The shifters' response became the basis for the `stopstatus` until we verified it by our data quality checks. Good runs have `stopstatus` = OK or `stopstatus` = CRASH, and bad runs have a `stopstatus` that indicates why they are bad. The CRASH designator means that there was no stop file (one of the standard files written by the DAQ when the run ends) for this run; without a stop file, the `stopstatus` could not be determined because the stop file contains the shifters' answer to whether the run ended normally. A missing stop file could indicate that the DAQ program crashed upon terminating the run, but the data may still be good. The DAQ configuration for early CUORICINO runs did not produce stop files so all runs before 295 have `stopstatus` = CRASH or one of the bad run indicators. The values of `stopstatus` that we used are listed in Table 5.2.

Bad intervals within a run are stored in the `bad_channels` table. A row in the `bad_channels` table specifies the run, channel, type of bad interval, start time, and stop time of the bad interval. The current possible types of bad intervals are listed in Table 5.3. The current entries in the `bad_channels` table can be viewed on a webpage at:
`http://cuoreweb.lngs.infn.it/~cuore/corc/hallA/bad_channels.php`

In the stabilization and anti-coincidence steps of the first-level analysis pro-

Table 5.2: Types of bad runs indicated by the `stopstatus` column of the `runs` table in the analysis database.

`stopstatus`	Description
OK	Good run
Bad for analysis	Run not suitable for analysis but not for one of the following problems
CRASH	No stop file but the run is good
Cryogenics	Bad run due to problems with cryogenics
DAQ	Bad run because of data acquisition problems
Electronics	Bad run due to problems with electronics

Table 5.3: Types of bad intervals in the `bad_channels` table in the analysis database.

Bad interval type	Description
bad calibration resolution	calibration peak FWHM too large
bad heater resolution	heater resolution too large, possibly due to bad stabilization
corc: saturated high	baseline too high, as determined by CORC
corc: saturated low	baseline too low, as determined by CORC
incompatible calibrations	initial and final calibration peaks shifted
inverted polarity	thermistor bias voltage was reversed, resulting in upside down pulses
large calibration uncertainty	cannot identify 2615 keV peak or residual at 2615 keV is too large
no triggered events	there were no signal events recorded, possibly due to incorrect DAQ or front-end settings
noisy	too much noise
saturated high	baseline too high
saturated low	baseline too low
unable to calibrate	calibration spectrum has no discernible peaks and cannot be calibrated

cessing, it is desirable not to use data contained in certain types of bad intervals. For the stabilization fits, all types of bad intervals are rejected if at least 3000 s of live time remain for the channel under consideration after rejection of all bad intervals. This amount of time is enough for at least eight stabilization heater events with which to perform the linear fit to the heater amplitude vs. baseline plot. If rejecting all bad intervals would result in less than 3000 s of live time, bad intervals of type "bad calibration resolution," "bad heater resolution," "incompatible calibrations," "large calibration uncertainty," "noisy," and "unable to calibrate" are not rejected when creating the graph of heater amplitude vs. baseline for the stabilization fit. For anti-coincidence, bad intervals are rejected except those of type "bad calibration resolution," "bad heater resolution," "incompatible calibrations," "large calibration uncertainty," and "noisy." That is, if there are coincident events on channels C and C' but the event on channel C' is in a bad interval other than one of the listed types, the event on channel C will not be rejected by the anti-coincidence cut.

5.3 Data selection

5.3.1 CORC check

We began the CORC check for each run by looking over the CORC Report page. For any channels referenced in the automatic reports of saturated intervals, time intervals of longer than two hours without triggered events, or the number of observed heater events being inconsistent with the expected number of heater events, we investigated the problem by checking the CORC Channel page. Saturated intervals can be an indication of a very noisy channel. A time hole without triggered events can signal rare problems with the DAQ setup or front-end electronics settings. The comparison of the observed number of heater events with the expected number is a powerful global check that the DAQ was operating normally, the heater was firing normally, there was not excessive noise degrading the pulse amplitude evaluation, and the stabilization fits were successful.

We also looked over the CORC Summary page for each run. For any channels outside of the normal range for any variable, we checked the channels on the CORC Channel page to determine if there was a real problem with the detector. We used the CORC History and History Graph pages to compare abnormal values with historical values for the same channel.

If there was a clear problem with the detector, such as an extremely high noise level or an absence of triggered events, we inserted into the analysis database the corresponding type of bad interval from Table 5.3 for the affected channel for a time period covering the duration of the problem. If the entire run was unusable for all channels, we changed the `stopstatus` to one of the bad types in Table 5.2.

We only marked data as bad if it was obviously unusable and would amount to an artificial dead time or severe degradation of the energy resolution. By marking the data as bad, we can explicitly treat bad intervals as dead time when computing the live time of a run and channel. Marginal cases in which a channel appeared worse than normal but not clearly unusable were not marked as bad. These marginal channels are handled by the signal efficiency calculations in Sect. 6.2.

Detailed logs of the CORC checks are kept on a web page accessible to the CUORE Collaboration at:
`http://wiki.hep.wisc.edu/cuore/AnalysisTopics/FirstLevelAnalysisNotes`
Additional logs about problems with CUORICINO runs are available at:
`http://wiki.hep.wisc.edu/cuore/AnalysisTopics/RunNotes`

5.3.2 Resolution consistency requirements

In the $0\nu\beta\beta$ decay analysis presented in Chapter 6, we assume the resolutions measured during calibration runs are representative of the resolutions during the whole data set. If a run is very noisy, the resolution may be degraded such that this assumption is violated. To check the consistency of the resolutions of the background runs, we monitored the resolutions of the heater peaks. For every run and channel, CORC fits a Gaussian to the heater peak in the stabilized amplitude spectrum and records the σ of the fit in the analysis database. An example of one of the fits is the plot at the lower right in Fig. 5.4. For each data set and channel, we plot the Gaussian σ of the heater peak as a function of run number, and we compute the median of the Gaussian σ values and the median absolute deviation (m.a.d.) from the median.

For a set of points $\{x_i\}$, the m.a.d. is defined as

$$\text{m.a.d.} \equiv \text{median}\{|x_i - \text{median}\{x_i\}|\}. \tag{5.1}$$

The m.a.d. provides a robust estimator of the width of a distribution; unlike the r.m.s. of the distribution, the m.a.d. is insensitive to outliers. For a Gaussian distribution, $1.4826 \times \text{m.a.d.}$ is equivalent to one sigma.

For each data set and channel, we set a cut on the Gaussian σ of the fit to the stabilized heater amplitude peak at $\sigma < \text{median} + 5 \times 1.4826 \times \text{m.a.d.}$ Figure 5.5 shows an example of one of the plots of resolution consistency with the cut value. If a run in the data set failed the cut, the channel and run were marked "bad heater resolution" in the database. The cut is rather loose from a statistical point of view; it is possible to make a loose cut on the heater resolution because the particle resolution at 2615 keV (or 2527 keV) is much larger than the heater resolution. The reason the particle resolution is worse is not fully understood, but there is evidently an intrinsic resolution for particle interactions, while the heater resolution is primarily determined by the noise. At 2527 MeV the intrinsic

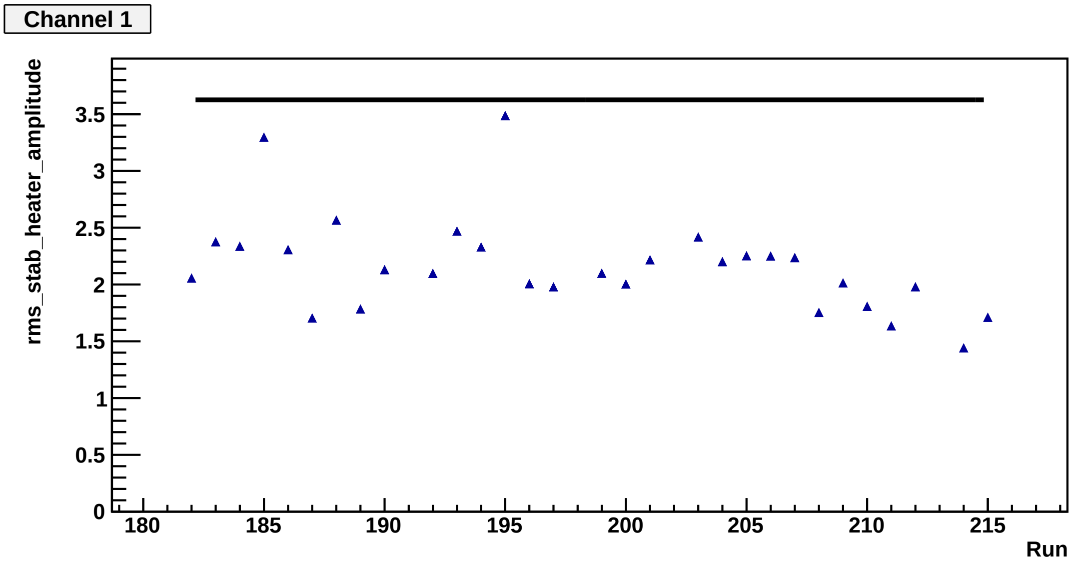

Figure 5.5: Example of a resolution consistency check for data set 2, channel 1. The points represent the heater resolution, retrieved from the `rms_stab_heater_amplitude` column of the `corc_summary` table of the analysis database. The horizontal line represents the cut at median + 5 × 1.4826 × m.a.d.

resolution dominates over the noise resolution, so the noise resolution measured on the heater peak may deviate a large amount from its average value without significantly affecting the particle resolution. Therefore, with this check we were mainly looking for instances of severe degradation of the noise resolution.

5.3.3 Calibration quality and consistency requirements

As described in Sect. 4.8, the calibration function that maps the stabilized amplitude of a pulse to the energy deposited in the detector was determined independently for each channel and data set. The results were checked according to the criteria described in this section, and the `bad_channels` database table was updated with the channels and data sets that failed the acceptance criteria.

For the cases where the automatic peak finding algorithm did not succeed, the calibration peaks were located by hand with the help of a GUI tool, leaving only a small number of channels and data sets that could not be calibrated because too few peaks were distinguishable from the background due to poor statistics. These channels and data sets were flagged as “unable to calibrate” in the analysis database.

A check was performed on the residual of the calibrated energy the 2615 keV peak, defined as the difference between the calibration function evaluated at the fitted mean of the peak in the stabilized amplitude spectrum and the known energy. Channels and data sets with an absolute value of the residual greater than 1 keV were flagged as "large calibration uncertainty." Additionally, channels and data sets for which the 2615 keV peak could not be located in the calibration spectrum were flagged as "large calibration uncertainty."

Instabilities in the energy injected by the heater, possibly due to ground loops or noise pickup by the heater circuit, can cause a shift in the calibration function. We checked for a shift in the calibration function by comparing the initial and final calibrations for each data set that has both. If the calibration peaks are shifted by an amount that is smaller than the detector resolution, when the initial and final spectra are added together to obtain the combined calibration spectrum, the small shift results in an enlargement of the combined peak width, as shown in Fig. 5.6a. If the distance between the center of the initial and final calibration peaks is large compared to their widths, the peak fitting algorithm is able to resolve both peaks, and the result of the fit is not representative of the whole data set (Fig. 5.6b). A noticeable incompatibility between the initial and final calibrations is undesirable because it is clear evidence of heater instabilities during the background runs. The parameter used to quantify the compatibility is $|x_i - x_f|/\overline{\mathrm{FWHM}}$, where x_i and x_f represent the fitted mean of the 2615 keV peak in the initial and final calibrations, respectively, and $\overline{\mathrm{FWHM}}$ is the mean FWHM resolution of this peak for the corresponding channel over all the data sets. Figure 5.7 shows the distribution of this parameter. Channels and data sets with $|x_i - x_f|/\overline{\mathrm{FWHM}} > 1$ were flagged as "incompatible calibrations" in the analysis database. For a channel and data set with $|x_i - x_f|/\overline{\mathrm{FWHM}} < 1$, the combined data from the initial and final calibrations form a single peak; the width of this peak, possibly broadened by a small shift between initial and final calibrations, is used in the $0\nu\beta\beta$ decay analysis as the energy resolution for the corresponding channel and data set.

An additional flag in the `bad_channels` database table was set to mark channels and data sets with poor resolutions. The flag is called "bad calibration resolution" and was set for channels and data sets with FWHM $>$ 11 keV for the big crystals or FWHM $>$ 20 keV for the small crystals, where the FWHM was measured on the 2615 keV calibration peak. However, the $0\nu\beta\beta$ decay analysis in this dissertation does not consider this flag because there is no loss of sensitivity from including data with poor resolutions. The "bad calibration resolution" flag may be useful for other analyses that sum spectra from different channels and data sets.

Table 5.4 summarizes the calibration acceptance criteria and the corresponding dead time introduced by excluding the data that failed the criteria.

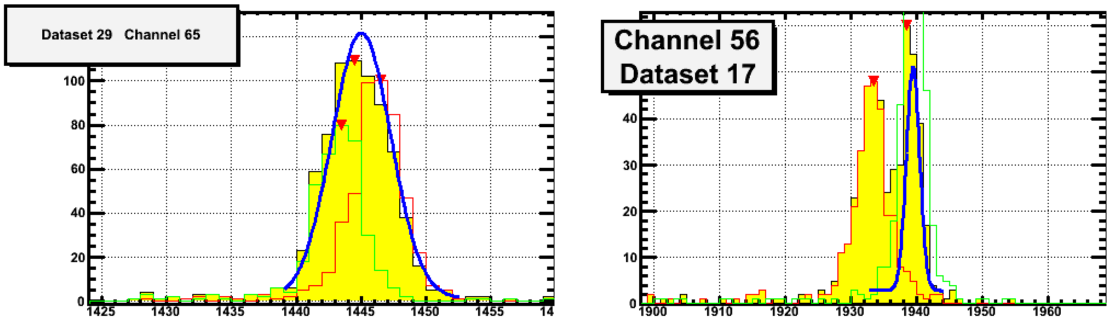

(a)

(b)

Figure 5.6: Two examples of channels and data sets for which the initial and final calibrations (red and green lines, respectively) are shifted. The horizontal axis is the stabilized amplitude, and the peak is the 2615 keV gamma line of ^{208}Tl. The blue line is the result of the calibration peak fitting algorithm. In example (a), the shift is small, and the initial and final calibration peaks merge into a single broader peak. In example (b), the shift is large, both peaks are resolved, and the calibration peak fit returns an incorrect estimate of the average position of the 2615 keV events and of the overall resolution. Figure from Maria Martinez.

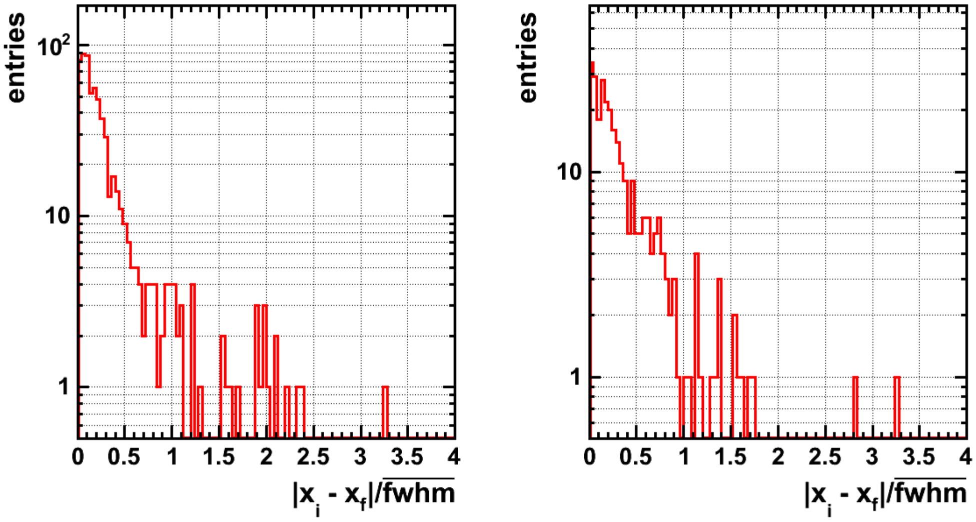

(a) Big crystals

(b) Small crystals

Figure 5.7: Distribution of the calibration compatibility parameter. Figure from Maria Martinez.

Table 5.4: Types of bad intervals in the `bad_channels` table in the analysis database related to the calibration acceptance criteria and the corresponding dead time. The $0\nu\beta\beta$ decay analysis in this dissertation does not reject data flagged as “bad calibration resolution,” but other CUORICINO analyses may choose to.

Bad interval type	Description	Dead time
unable to calibrate	calibration spectrum has no discernible peaks and cannot be calibrated	0.85%
large calibration uncertainty	cannot identify 2615 keV peak or residual at 2615 keV is > 1 keV	1.0%
incompatible calibrations	$\lvert x_i - x_f \rvert > \overline{\mathrm{FWHM}}$ @ 2615 keV	3.6%
bad calibration resolution	$\mathrm{FWHM}_{2615\ \mathrm{keV}} > 11$ keV (Big crystals) $\mathrm{FWHM}_{2615\ \mathrm{keV}} > 20$ keV (Small crystals)	4.6%

5.3.4 Selection criteria for analysis

Only runs with `stopstatus` = OK or `stopstatus` = CRASH are used in the analysis. Bad intervals in the `bad_channels` table are excluded from the analysis except for those of type “bad calibration resolution” because this analysis can assign a poor resolution to a channel for an individual data set without worsening the overall sensitivity. Other types of analyses may prefer to exclude all bad intervals, including those of type “bad calibration resolution.” CUORICINO channels 2 and 3 are not used because they were too noisy and problematic. CUORICINO channel 50 is not used because it did not have a working heater. The two crystals enriched in ^{128}Te, CUORICINO channels 35 and 52, are not used because they contribute a negligible amount of ^{130}Te. The live time for each run and channel is computed by subtracting the length of the bad intervals that are excluded (avoiding double counting overlapping bad intervals) from the run duration.

Chapter 6

CUORICINO $0\nu\beta\beta$ decay fit and limit

In addition to a combined analysis of data from CUORICINO and the Three Towers Test, an analysis of only the CUORICINO data has been carried out so that it can be compared with an independent analysis of the CUORICINO data by collaborators at the University of Milano-Bicocca. This chapter presents the results of the measurement of the ^{130}Te $0\nu\beta\beta$ decay rate with the CUORICINO data and the statistical approach used to obtain them. The general strategy is to treat each crystal as an independent detector and to perform a simultaneous fit to all the detectors. In the fit the response function of each detector is matched to the resolution of the respective detector measured during calibration runs. The data are further subdivided by time according to data set because the energy resolution varies over time. The combined analysis of CUORICINO and the Three Towers Test is presented in the next chapter.

The chapter is organized as follows: First, in Sect. 6.1 the cuts used in the analysis are described. Then, in Sect. 6.2 the calculation of the efficiency for detecting $0\nu\beta\beta$ decays is discussed. In Sect. 6.3 the approach for obtaining the resolutions used in the fit is presented. The fitting technique and the probability density function that defines the simultaneous fit are presented in Sect. 6.4. Section 6.5 contains the results of the fit for the $0\nu\beta\beta$ decay rate of ^{130}Te including the statistical errors. Systematic errors on the decay rate are considered in Sect. 6.6. The technique employed to set an upper limit, incorporating systematic errors, on the $0\nu\beta\beta$ decay rate (or lower limit on the partial half-life) of ^{130}Te is explained in Sect. 6.7 together with the results. Finally, toy Monte Carlo studies to validate the fitting procedure and to compute the expected sensitivity of the experiment are presented in Sect. 6.8.

6.1 Cuts

In addition to the data selection criteria in Sect. 5.3.4, the following cuts are applied. The re-triggered pulses algorithm introduced in Sect. 4.5 is used to identify pulses that are present in more than one acquisition window and to select only one instance of each multiply acquired pulse. Events that are consistent with being heater events according to the offline heater flagging algorithm described in Sect. 4.6 are removed from the spectrum. The pulse shape cut discussed in Sect. 4.9 is applied with $n_\sigma^{\mathrm{PS}} < 2$ to reject spurious events and pileup. An anti-coincidence cut is applied to reduce backgrounds from alpha decays on the crystal surfaces that deposit energy in two neighboring crystals, from gammas that Compton scatter in one crystal before interacting in another crystal, and from gammas that are part of a cascade of gammas emitted in a radioactive decay (e.g. the decay of ^{208}Tl). The anti-coincidence cut rejects events that occurred within ±100 ms of a non-heater event on another bolometer that passed the pulse shape cut; the time difference is based on the trigger times of the events. A fixed dead time around each heater event is imposed so that the efficiency loss due to pileup with heater events can be calculated exactly. The fixed dead time begins 3.1 s before a heater event and lasts until 4.0 s after the heater event.

6.2 Signal efficiency

The signal efficiency is the probability that a $0\nu\beta\beta$ decay event is detected, its energy is reconstructed accurately, and it passes the data selection cuts. The average signal efficiency is evaluated to be $(81.7 \pm 3.3)\%$ for the CUORICINO big crystals and $(80.0 \pm 3.3)\%$ for the small crystals. Mechanisms for loss of signal efficiency are escape of a β from the crystal, the pulse shape cut, the anti-coincidence cut, signal degradation due to noise, and pileup. These mechanisms are studied in the following sections, and their contributions to the signal efficiency are summarized in Table 6.1.

6.2.1 Escape of a β

If a $0\nu\beta\beta$ decay event occurs near the surface of a crystal (within ~2 mm), one or both of the β-particles may escape from the crystal before depositing their full kinetic energy in the crystal. Since only a fraction of the Q-value would be registered in the bolometer, the event would fall outside of the region in the spectrum where a $0\nu\beta\beta$ decay peak is expected and therefore would not be counted as a signal event. The efficiency loss due to decays near the surface is evaluated using a GEANT4-based Monte Carlo simulation. In the simulation $0\nu\beta\beta$ decays are generated uniformly throughout the volume of a crystal with the correct energy

Table 6.1: Contributions to the CUORICINO $0\nu\beta\beta$ decay signal efficiency.

Source	Signal efficiency (%)
Escape of a β	86.3 ± 3.5 (big crystals) 84.5 ± 3.5 (small crystals)
Pulse shape cut	98.5 ± 0.3
Anti-coincidence cut	99.3 ± 0.1
Noise	99.1 ± 0.1 (average of all channels and data sets)
Pileup with heater pulses	97.7 ± 0.0
Total	81.7 ± 3.3 (average of big crystals) 80.0 ± 3.3 (average of small crystals)

spectrum and angular distribution for the individual electrons. The efficiency is the number of decays for which both β-particles stop in the crystal divided by the total number of simulated decays. Since the probability of escape depends on the crystal geometry, in particular the ratio of surface area to volume, the efficiency is slightly different for the big ($5 \times 5 \times 5$ cm^3) crystals and the small ($3 \times 3 \times 6$ cm^3) crystals. From the Monte Carlo simulation, the efficiency due to escape is evaluated to be 86.3% for the big crystals and 84.5% for the small crystals. Based on variations in the efficiency result depending on threshold cuts for secondaries production in GEANT4 and depending on the version of the Monte Carlo code, I assign a conservative error of 3.5% on the efficiency estimates.

6.2.2 Pulse shape cut

A pulse shape cut is applied to remove events that are not consistent with a single physical pulse in the acquisition window. Some events rejected by the pulse shape cut are spurious fluctuations of the thermistor voltage due to noise or temperature instabilities that are large enough to trigger the data acquisition. The pulse shape cut also removes pileup involving multiple events in the same acquisition time interval, including most cases of pileup on the rise of the pulse, except when two events in the same crystal occur so close in time as to be indistinguishable from a single event with the sum of their energies. Efficiency loss occurs if the pulse shape cut rejects any good events. The pulse shape cut is designed to be conservative, in the sense that good events should be retained with close to 100% efficiency.

To evaluate the efficiency of the pulse shape cut, the background photopeak at 2615 keV due to ^{208}Tl is used as a proxy for the $0\nu\beta\beta$ decay peak expected at 2527.5 keV. The 2615 keV peak is chosen because the pulse shape cut is energy-dependent, and the 2615 keV peak is the closest to 2527.5 keV of all of the large background peaks. The peak provides a large number of good events with which to compute the efficiency of the cut. The region of the spectrum containing the 2615 keV background peak is plotted with the pulse shape cut applied and also with the complementary cut applied (that is, the events that are rejected by the pulse shape cut). Figure 6.1 shows the two spectra with the accepted events on the top and the rejected events on the bottom. The efficiency, $\varepsilon_{\mathrm{PS}}$, of the pulse shape cut is determined from a simultaneous fit to both spectra.

The simultaneous fit is parameterized in terms of the total number of signal events, N_{signal}; the signal efficiency, $\varepsilon_{\mathrm{PS}}$; the total number of background events, N_{bkg}; and the background efficiency, $\varepsilon_{\mathrm{PS}}^{\mathrm{bkg}}$. The signal yield for the accepted events is $\varepsilon_{\mathrm{PS}} N_{\mathrm{signal}}$, and the signal yield for the rejected events is $(1-\varepsilon_{\mathrm{PS}})N_{\mathrm{signal}}$. Similarly, the background yield for the accepted events is $\varepsilon_{\mathrm{PS}}^{\mathrm{bkg}} N_{\mathrm{bkg}}$, and the background yield for the rejected events is $(1-\varepsilon_{\mathrm{PS}}^{\mathrm{bkg}})N_{\mathrm{bkg}}$. By including $\varepsilon_{\mathrm{PS}}$ directly in the parameterization of the fit, correlations among the fit parameters are automatically taken into account by the fitter when it determines the error on $\varepsilon_{\mathrm{PS}}$. The advantage of fitting accepted and rejected events rather than accepted and total events is that the numbers of accepted and rejected events are uncorrelated since the number of events is a random variable following a Poisson distribution; the number of accepted events is correlated with the total number of events, which would lead to an incorrect error on the efficiency if the correlation is not taken into account.

The same signal and background shapes are used to fit both spectra. The signal shape is taken to be a sum of two Crystal Ball functions. A Crystal Ball function is a Gaussian with a power law tail to account for energy loss defined by [85]

$$f(x;\alpha,n,\bar{x},\sigma) = \begin{cases} \exp\left(-\frac{(x-\bar{x})^2}{2\sigma^2}\right), & \text{for } \frac{x-\bar{x}}{\sigma} > -\alpha \\ A\cdot\left(B-\frac{x-\bar{x}}{\sigma}\right)^{-n}, & \text{for } \frac{x-\bar{x}}{\sigma} \leq -\alpha \end{cases} \tag{6.1}$$

where

$$\begin{aligned} A &= \left(\frac{n}{|\alpha|}\right)^{n}\cdot\exp\left(-\frac{|\alpha|^2}{2}\right), \\ B &= \frac{n}{|\alpha|}-|\alpha|. \end{aligned} \tag{6.2}$$

The two Crystal Ball signal components have the same Gaussian mean but different Gaussian sigmas and different values of the Crystal Ball parameters α and n. The sum of two Crystal Ball functions with different widths better represents the signal shape than a single Crystal Ball function since the signal is obtained by

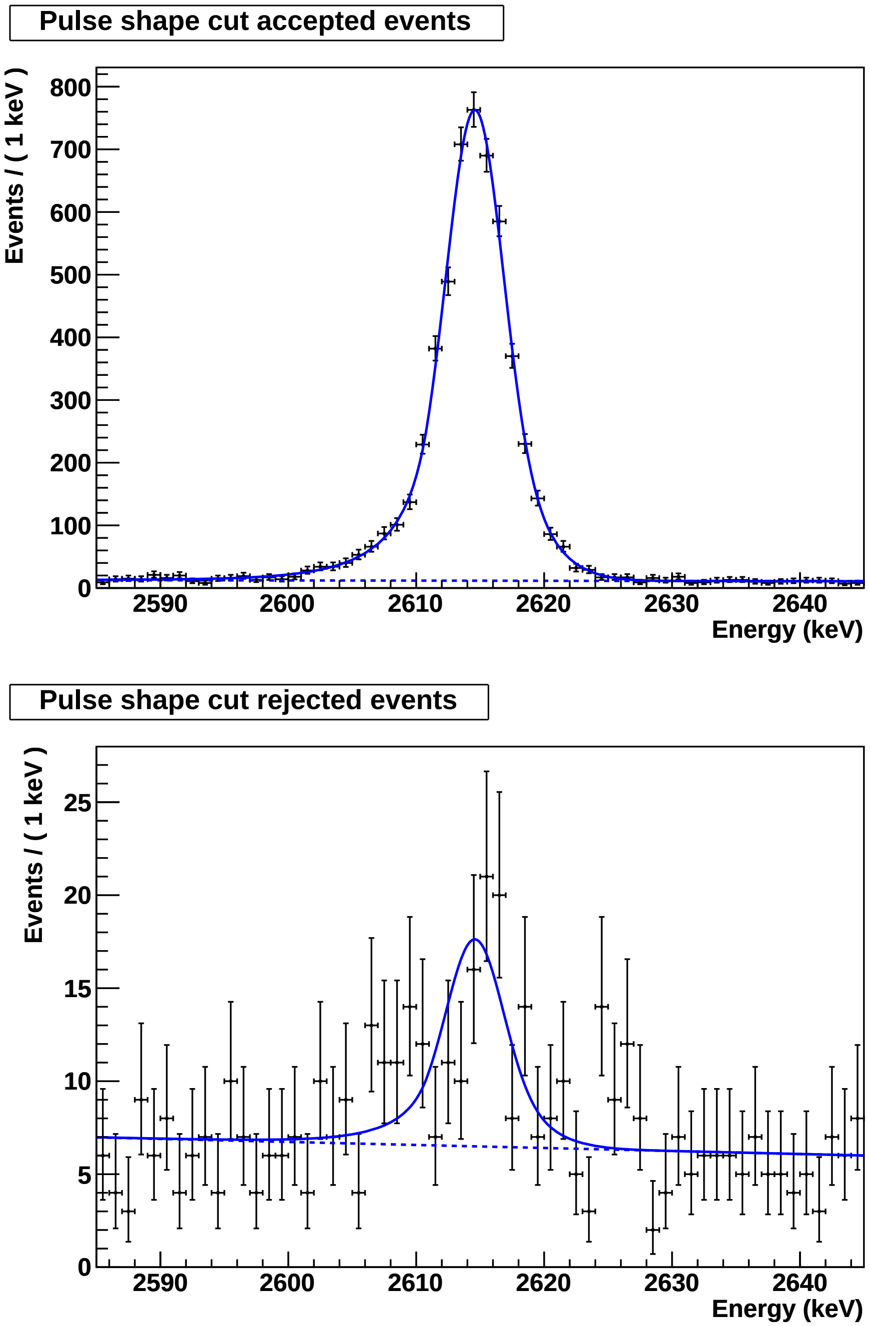

Figure 6.1: Fit to the ^{208}Tl spectrum to determine the pulse shape cut efficiency. On the top are the events that are accepted by the pulse shape cut, and on the bottom are the events that are rejected. The efficiency of the cut is determined from a simultaneous fit to the two spectra.

summing many detectors with different resolutions. All signal shape parameters, including the relative fraction of the two Crystal Ball components, are floating in the fit. The background shape is taken to be linear. The result from the fit is $\varepsilon_{\mathrm{PS}} = (98.5 \pm 0.3)\%$ for CUORICINO. The efficiency for background is $\varepsilon_{\mathrm{PS}}^{\mathrm{bkg}} = (64 \pm 2)\%$ for CUORICINO. From these results the pulse shape cut is seen to be very powerful, rejecting approximately 36% of the events in the continuum background while retaining 98.5% of the signal events in the peak.

6.2.3 Anti-coincidence cut

An anti-coincidence cut is applied to suppress backgrounds from alpha decays on the crystal surfaces and gamma interactions in multiple crystals. The anti-coincidence cut rejects any events that occur within 100 ms of an event in another crystal. Since $0\nu\beta\beta$ decay events that would lead to a bolometer signal at the Q-value are inherently single-crystal events (unlike the events with a β-particle escaping described in Sect. 6.2.1), the only efficiency loss from the anti-coincidence cut is due to the possibility of a random coincidence with another event. The probability of an event randomly occurring within $\pm\Delta t$ of a particular event is given by

$$\mathrm{Prob(random\ coincidence)} = 1 - e^{-2r\Delta t}, \tag{6.3}$$

where r is the total event rate on all other bolometers. The signal efficiency, being the probability that a random coincidence does not occur, is $e^{-2r\Delta t} \approx 1 - 2r\Delta t$. For typical values of $r \approx 1$ mHz/crystal $\times\, 56$ crystals $= 56$ mHz and $\Delta t = 100$ ms, the expected signal efficiency evaluates to 98.9%.

For this analysis the efficiency due to the anti-coincidence cut is evaluated directly on the largest background peak, the photopeak at 1461 keV due to ^{40}K, following the same procedure used for the pulse shape cut efficiency in Sect. 6.2.2. The spectra of events accepted and rejected by the anti-coincidence cut are fit simultaneously using the sum of two Crystal Ball functions and a linear background. As in Sect. 6.2.2, the anti-coincidence cut efficiency, $\varepsilon_{\mathrm{AC}}$, is directly included in the parameterization of the fit. The fit is shown in Fig. 6.2. For CUORICINO the fit result for the efficiency of the anti-coincidence cut is $\varepsilon_{\mathrm{AC}} = (99.3 \pm 0.1)\%$.

6.2.4 Signal degradation due to noise

The noise superimposed on the signal pulses distorts the pulse shape and introduces an error on the reconstructed energy of the event that grows with the size of the noise. For this reason very noisy runs are rejected from the analysis as described in Chapter 5. Within generally good runs used in the analysis, some channels occasionally experience momentary periods of elevated noise levels due, for example, to mechanical vibrations from the dilution unit. Figure 6.3 shows

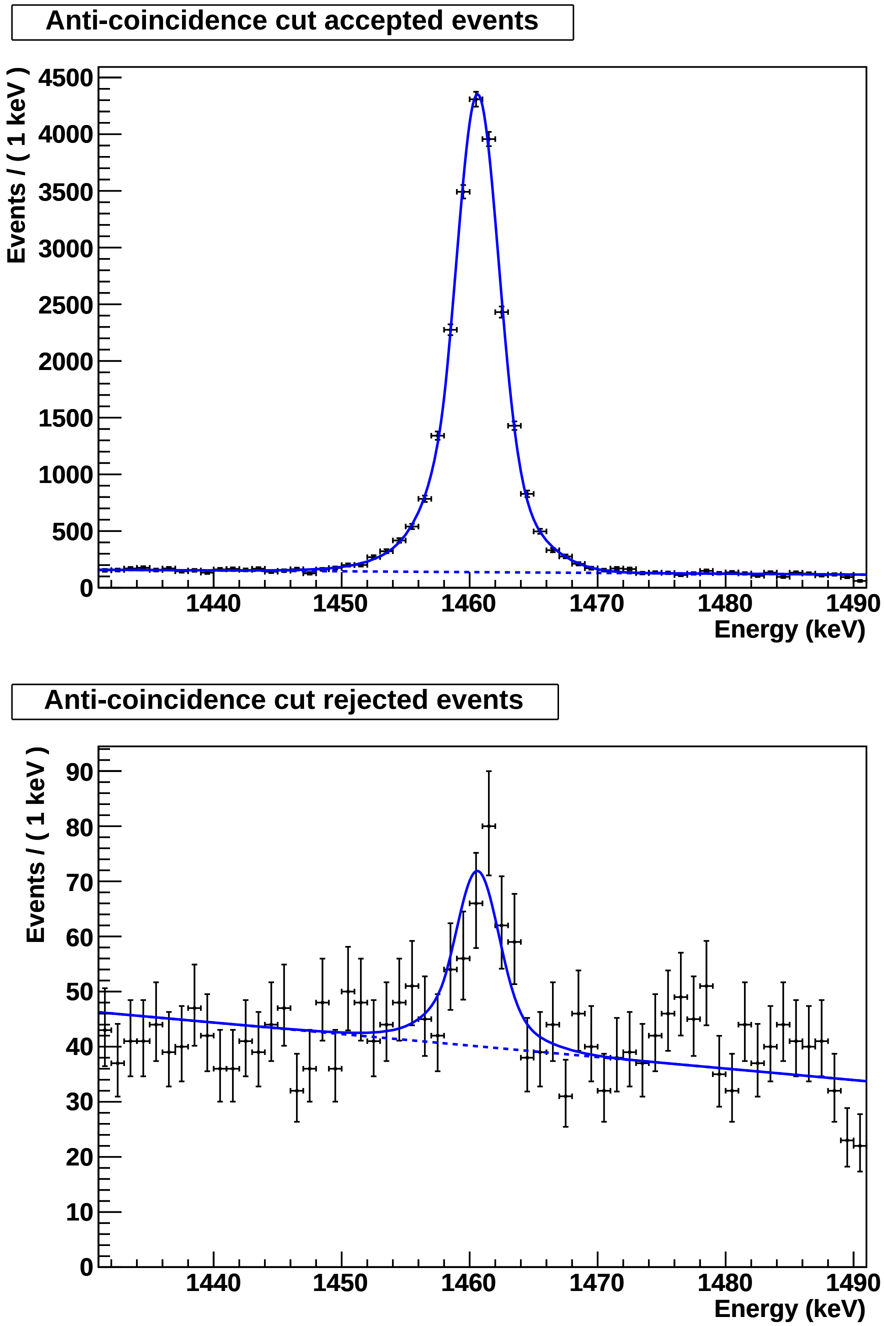

Figure 6.2: Fit to the ^{40}K spectrum to determine the anti-coincidence cut efficiency. On the top are the events that are accepted by the anti-coincidence cut, and on the bottom are the events that are rejected. The efficiency of the cut is determined from a simultaneous fit to the two spectra.

an example of the onset of a period of increased noise and the corresponding degradation of the energy resolution for heater events. A representative pulse from the noisy period is shown in Fig. 6.4; distortion of the pulse shape leads to a large error in the pulse amplitude evaluation as well as a large error in the baseline estimation.

To estimate the efficiency loss due to noise degrading the signal and spoiling the energy reconstruction, the number of events in the heater peak in the spectrum is compared to the expected number of heater events: $\varepsilon_{\text{noise}} = N_{\text{observed}}/N_{\text{expected}}$. Since the prevalence of elevated noise levels depends on the channel and the data set, a separate efficiency due to this effect is computed for each channel and data set; that is, one for each subset of the $0\nu\beta\beta$ decay simultaneous fit. The observed number of heater events is obtained from a fit to the heater peak in the spectrum of stabilized amplitudes. The fit function is a sum of two Gaussians with their means fixed to 5000 (where the heater peak is expected in the stabilized amplitude spectrum) but different widths and a flat background. The Gaussian widths are required to be less than the width of the signal response function used for the $0\nu\beta\beta$ decay fit. The expected number of heater events is carefully reconstructed based on the expected time of each heater event and whether that time is contained in a bad time interval that is rejected. The error on the efficiency for each channel and data set is calculated from

$$\sigma_{\varepsilon_{\text{noise}}} = \sqrt{\frac{\varepsilon_{\text{noise}}(1-\varepsilon_{\text{noise}})}{N_{\text{expected}}}}, \tag{6.4}$$

which is valid for a binomial process with a fixed number of trials. Figure 6.5 shows the distribution of efficiencies due to noise for all channels and data sets. The average efficiency due to noise is $\langle\varepsilon_{\text{noise}}\rangle = 99.1\%$, and the average error on the efficiency is $\langle\sigma_{\varepsilon_{\text{noise}}}\rangle = 0.1\%$. Most channels and data sets have efficiencies greater than 99%, and almost all are greater than 95%. Figure 6.6 depicts the channels and data sets with an efficiency less than 95%.

6.2.5 Pileup

Pileup is the occurrence of two or more events in the same crystal within a time period comparable to the length of a pulse. The optimal filter cannot accurately estimate the amplitude of events involved in pileup, and pileup events are rejected by the pulse shape cut. The probability of an event randomly occurring on a particular channel within a time interval of length Δt is given by $1 - e^{-r\Delta t}$, where r is the event rate on the same channel. For pileup considerations, the relevant time interval extends from about 3 s before an event to about 3 s after the event for CUORICINO, resulting in $\Delta t \approx 6$ s. For a typical single-crystal event rate of

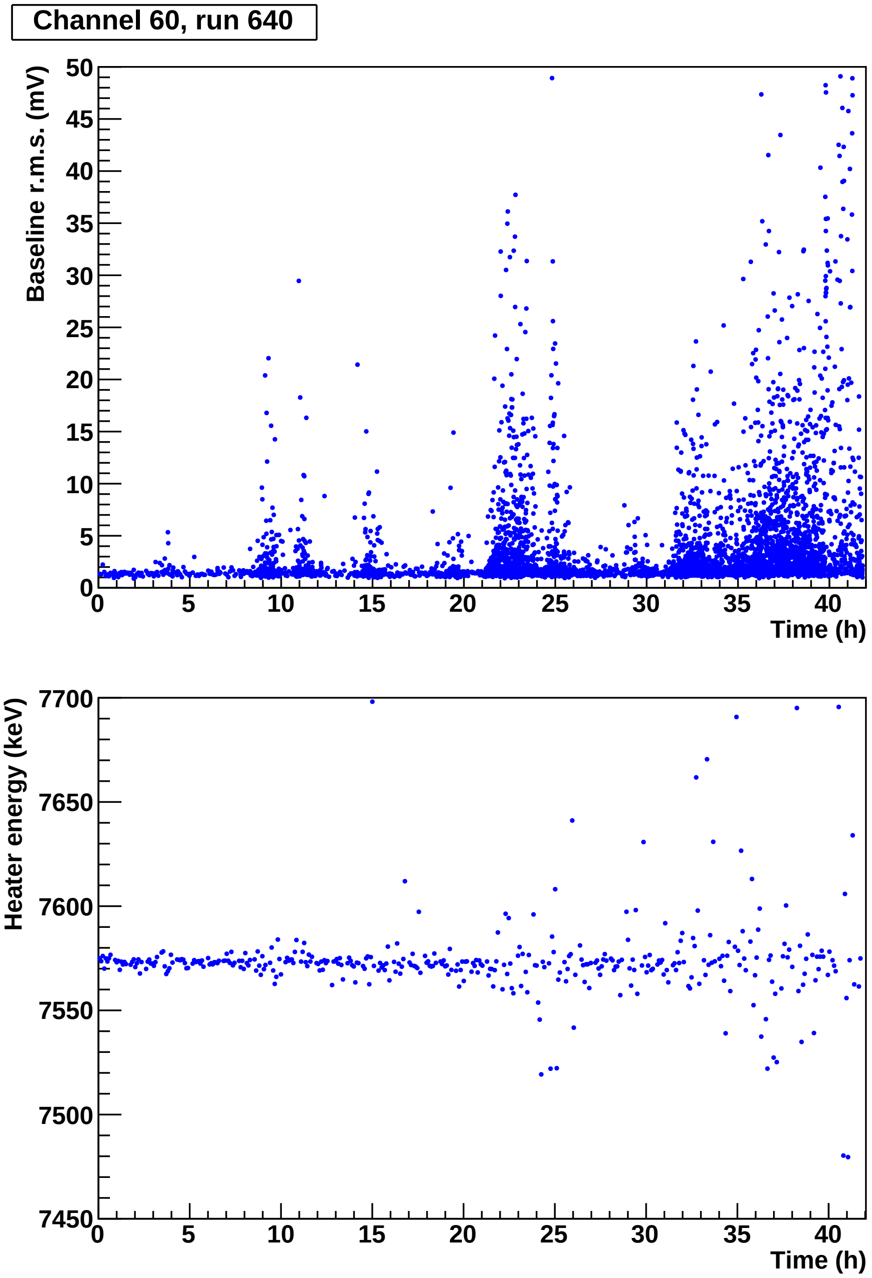

Figure 6.3: Example of noisy time intervals on channel 60 in run 640. The baseline r.m.s. measures the noise level. After ~9 h into the run, the noise level spikes several times. The scatter in measured heater energies demonstrates the degradation of the energy reconstruction caused by the noise.

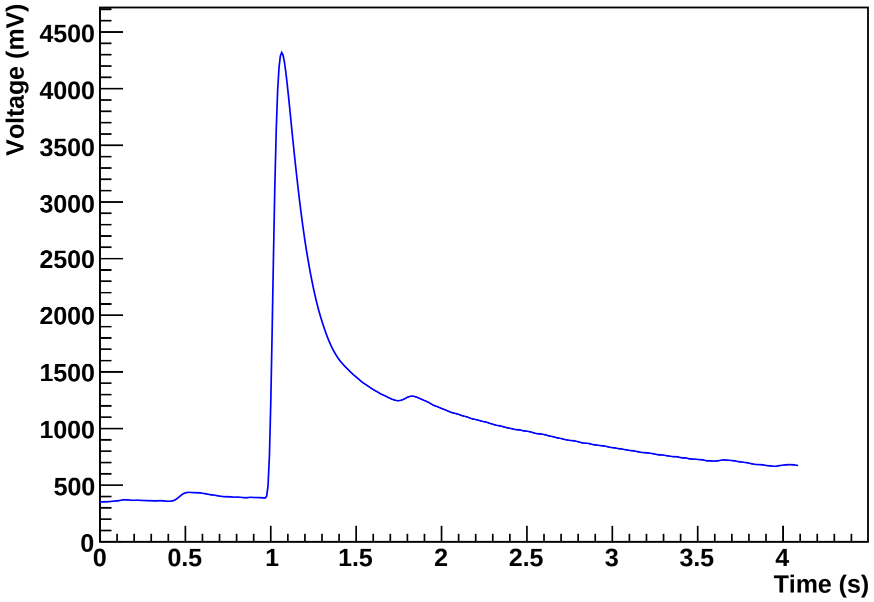

Figure 6.4: Example of a heater pulse from a noisy time interval in Fig. 6.3. Distortions superimposed on the normal pulse shape impede accurate measurements of the pulse amplitude and the pre-pulse baseline level. Such deterioration of pulses lowers the signal efficiency since it can cause the reconstructed energy of $0\nu\beta\beta$ decay events to fall outside of the expected peak location.

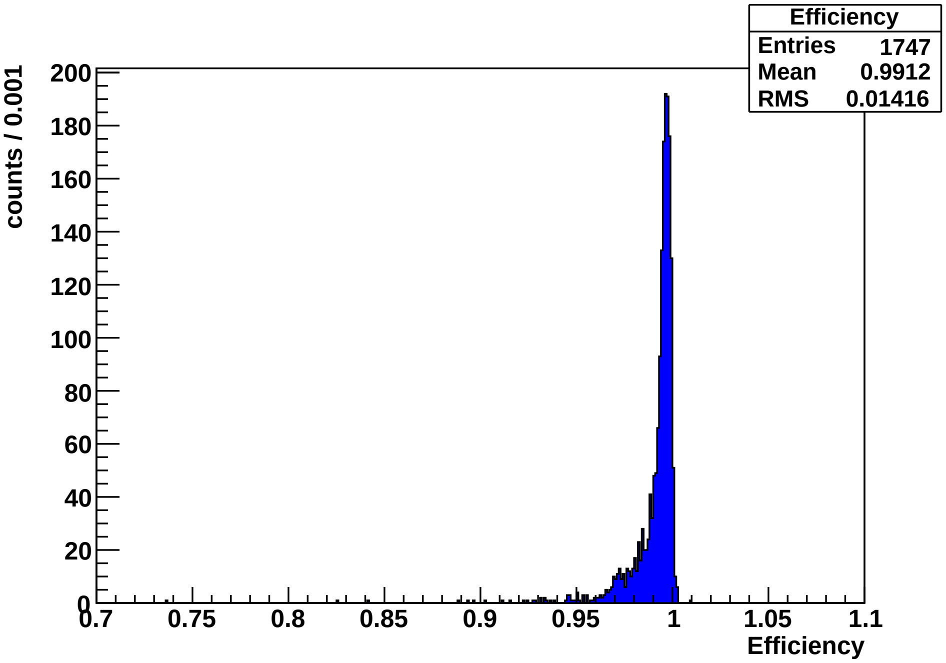

Figure 6.5: Distribution of the ratio of the observed number of heater events to the expected number of heater events, used to estimate the efficiency loss due to signal degradation by noise. Each entry in the histogram corresponds to one channel and data set.

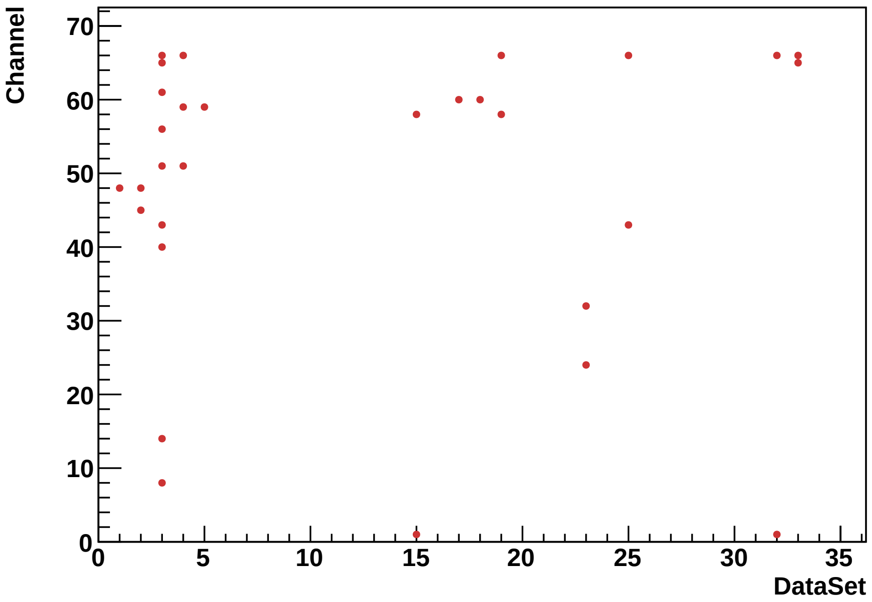

Figure 6.6: Channels and data sets with efficiency due to noise less than 95%.

$r \approx 1$ mHz, the probability of a given event having another event within a window of length $\Delta t \approx 6$ s around it is about 0.6%. Therefore, the contribution to the signal efficiency from pileup is expected to be about 99.4%. This contribution to the signal efficiency is already incorporated into the efficiencies evaluated for the signal degradation due to noise (Sect. 6.2.4) and due to the pulse shape cut (Sect. 6.2.2). If a pileup occurs, the underlying events may be lost from the signal if either there is a large error in the energy evaluation or there is not a large error in the energy but the event is rejected by the pulse shape cut. The former possibility is accounted for in the fit of the heater events used to estimate the efficiency loss due to noise, and the latter possibility is accounted for in the fit to the background peaks used to evaluate the efficiency loss due to the pulse shape cut. Therefore, there is no independent contribution to the efficiency budget from pileup with non-heater pulses.

A special kind of pileup is pileup with a heater pulse, which is not a random process. Heater events occur every approximately 305 s, and any signal event within about 3 s before or about 4 s after a heater event will be lost due to pileup. A fixed dead time of 3.1 s before and 4 s after every heater event is imposed by a cut, resulting in an efficiency of $(305 - 7.1)/305 \approx 97.7\%$.

6.3 Energy resolutions

The detector response function used in the $0\nu\beta\beta$ decay fit requires the energy resolution of every bolometer for every data set. The energy resolutions are measured using calibration data collected while the detectors were exposed to a ^{232}Th source. The spectrum of the calibration source includes a strong photopeak at 2615 keV due to the decay of ^{208}Tl, which is part of the thorium decay chain. Since 2615 keV is close to the ^{130}Te Q-value of 2527.5 keV, the measured resolution of the 2615 keV calibration peak is used as the detector resolution for the $0\nu\beta\beta$ decay fit. An unbinned maximum likelihood fit with a Gaussian signal shape and a linear background measures the full width at half maximum (FWHM) of the 2615 keV peak. Examples of the fits are shown in Fig. 6.7. The FWHM resolutions for each channel and data set are listed in Tables B.2, B.3, and B.4 in Appendix B, and histograms of the resolutions for the big crystals and small crystals are plotted in Fig. 6.8.

6.4 Probability density function

The $0\nu\beta\beta$ decay rate is measured using a simultaneous, unbinned maximum-likelihood fit. In a simultaneous fit, the data are divided into subsets according to one or more discrete observables, and one or more fit parameters are shared among the subsets. For the $0\nu\beta\beta$ decay fit, the data are split into subsets according to channel and data set, and the $0\nu\beta\beta$ decay rate parameter is shared among all subsets. A subset contains the data from one channel and one data set. Many subsets are empty since the background rates are extremely low, but this is not a problem for the fit. The probability density function (p.d.f.), $f(E, C, D)$, is a function of one continuous observable, energy (E), and two discrete observables, channel (C) and data set (D):

$$\begin{aligned} f(E,C,D) = {} & \frac{N_{\text{bkg}}(C,D)}{N_{\text{total}}}\text{polynomial}(E) \\ & + \frac{N_{^{60}\text{Co}}(C,D)}{N_{\text{total}}}\text{Gaussian}(E;\mu_{^{60}\text{Co}},\sigma(C,D)) \\ & + \frac{N_{0\nu\beta\beta}(C,D)}{N_{\text{total}}}\text{Gaussian}(E;Q,\sigma(C,D)). \end{aligned} \tag{6.5}$$

The first term in Eq. (6.5) represents the continuum background, modeled by a polynomial function for each subset. The polynomial function of degree N, normalized to unit area over the energy range used in the fit, is defined as

$$\text{polynomial}(E) = \frac{\sum_{n=0}^{N} c_n E^n}{\int_{E_{\min}}^{E_{\max}} \sum_{n=0}^{N} c_n E'^n \, dE'}, \tag{6.6}$$

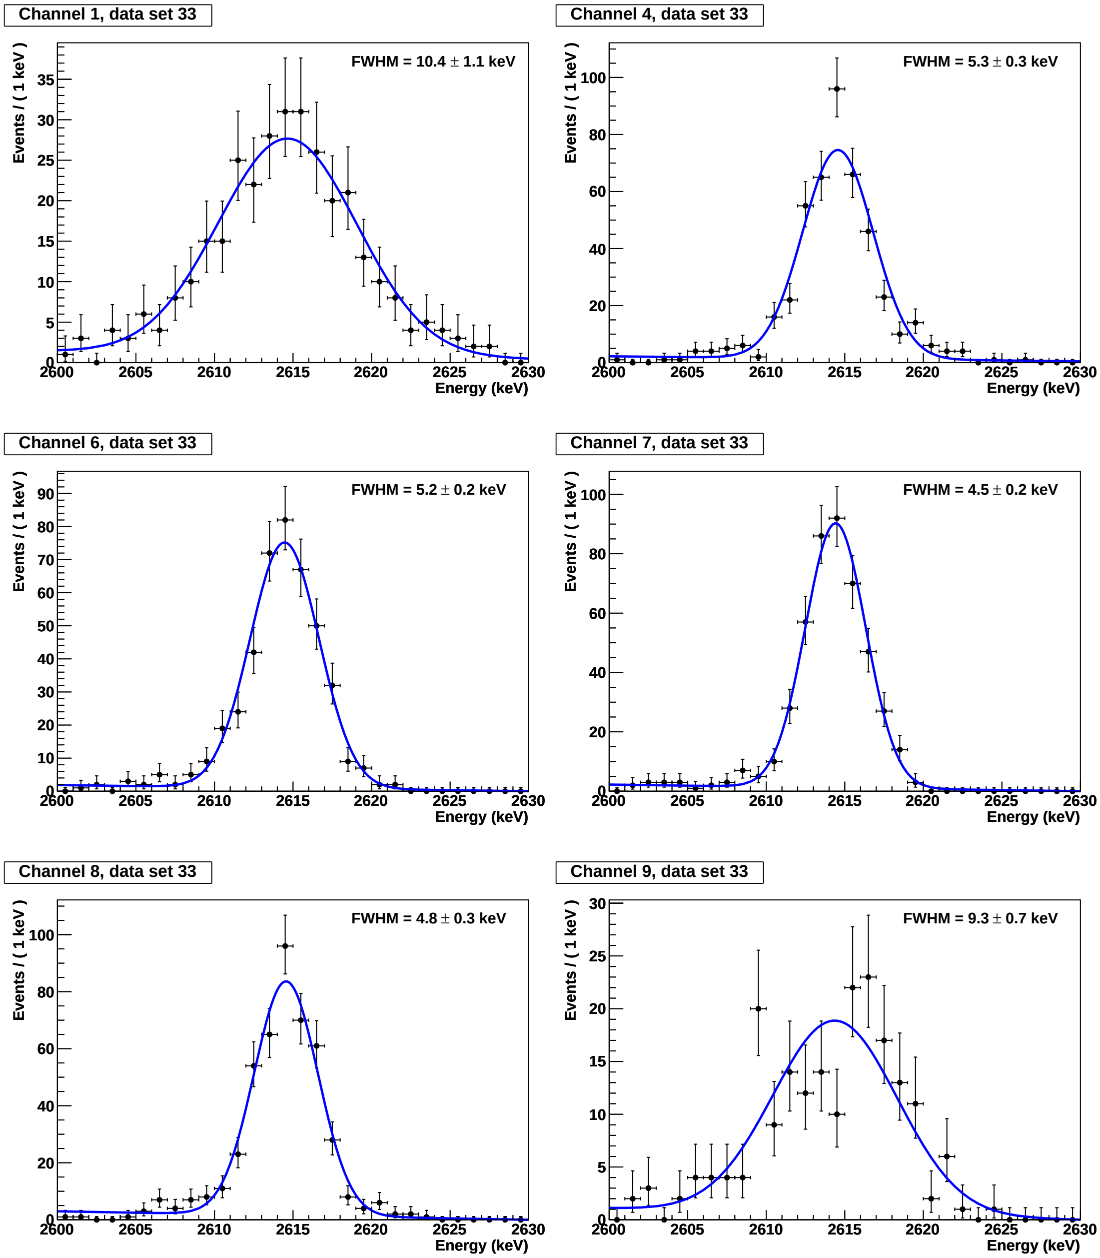

Figure 6.7: Examples of fits to the 2615 keV calibration peak for determining the energy resolution.

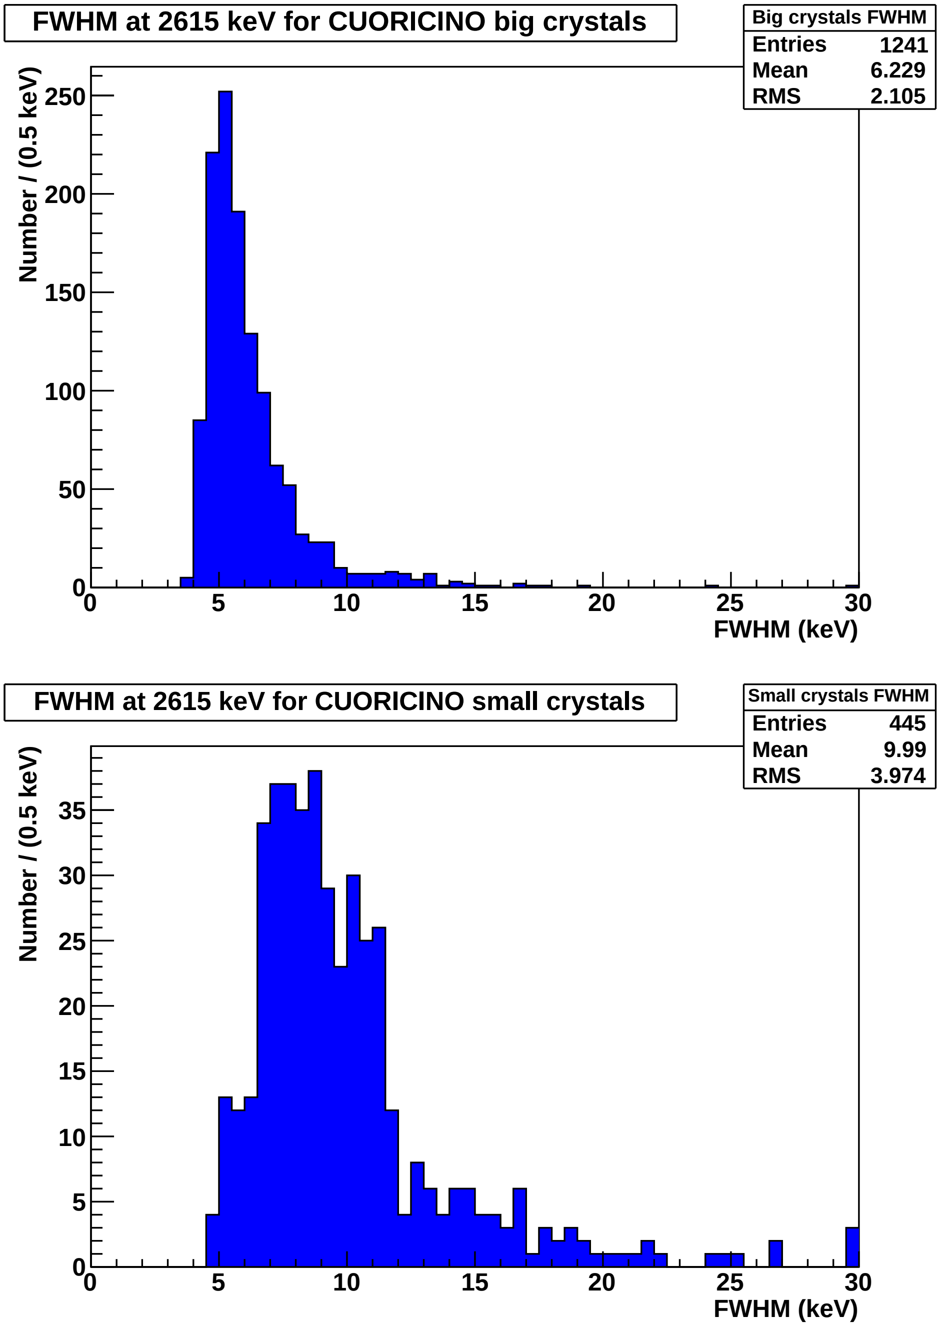

Figure 6.8: Distributions of energy resolutions for the CUORICINO big crystals (top) and the CUORICINO small crystals (bottom). There is one entry for each channel and data set used in the analysis.

where E_{min} and E_{max} are the minimum and maximum energies in the fit range, respectively. The standard version of the fit models the continuum background as a flat, 0-degree polynomial. Higher degree polynomial background models are considered in Sect. 6.6 on systematic errors. The expected number of background events in the subset corresponding to channel C and data set D is

$$N_{\text{bkg}}(C, D) = r_{\text{bkg}}(C) \times \Delta E \times m(C) \times T(C, D), \tag{6.7}$$

where $r_{\text{bkg}}(C)$ is the background rate in counts/(keV · kg · y) for channel C, ΔE is the width of the energy window used in the fit (nominally 2465–2595 keV, $\Delta E = 130$ keV), $m(C)$ is the mass of crystal C, and $T(C, D)$ is the live time for channel C and data set D. The background rate, $r_{\text{bkg}}(C)$, is floating in the fit and is constrained to be the same for all crystals of the same type. The types of crystals are big, small (natural tellurium), and enriched. Therefore, there are three background rates floating in the fit. The parameters ΔE, $m(C)$, and $T(C, D)$ are fixed in the fit.

The second term in Eq. (6.5) represents the sum peak of the 1173.2 keV and 1332.5 keV gammas that are emitted in a cascade in the β decay of ^{60}Co. The ^{60}Co was produced by cosmogenic activation of the copper parts of the detector mounting structure and cryostat thermal shields while the copper was above ground. Since the ^{60}Co nuclei are contained within the bulk of the copper, the electron emitted in the β decay, which has a relatively low endpoint energy of 318 keV, is stopped within the copper. Therefore, the detector may observe one or both of the 1173.2 keV and 1332.5 keV gammas but does not detect the electron. The individual gamma peaks are prominent in the background spectrum (Fig. A.1) and do not show any evidence of a high energy tail that would be present if the electron entered the detector. The ^{60}Co sum peak is due to the simultaneous energy deposition of both gammas in the same crystal. It is modeled by a Gaussian for each subset with the Gaussian sigma fixed to the value obtained from a Gaussian fit to the 2615 keV calibration peak for the same subset. The Gaussian function, normalized to unit area, is defined as

$$\text{Gaussian}(E; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(E-\mu)^2/2\sigma^2}. \tag{6.8}$$

The peak is expected to be centered at 2505.7 keV, but the mean of the Gaussian function, $\mu_{^{60}\text{Co}}$, is allowed to float in the fit to accommodate a small degree of uncertainty in the calibration. The expected number of ^{60}Co events in the subset corresponding to channel C and data set D is

$$N_{^{60}\text{Co}}(C, D) = r_{^{60}\text{Co}}(C) \times m(C) \times T(C, D) \times (1/2)^{t(D)/t_{1/2}}, \tag{6.9}$$

where $r_{^{60}\text{Co}}(C)$ is the rate of ^{60}Co events in counts/(kg · y) at time $t = 0$ (the beginning of the experiment). Like the continuum background rate, $r_{^{60}\text{Co}}(C)$ is floating in the fit and is constrained to be the same for all big crystals and the same for all small crystals. No ^{60}Co peak is included for the enriched crystals (the rate is fixed to 0) because the statistics from the two enriched crystals are too limited to resolve the ^{60}Co peak. In the last factor in Eq. (6.9), which accounts for the exponential decay of ^{60}Co with a half-life of $t_{1/2} = 5.3$ y, $t(D)$ is the time at the start of data set D since the beginning of the experiment. To be exact, the exponential function representing the decay of ^{60}Co should be averaged over the data set, omitting bad intervals that are not used, but since the ^{60}Co half-life of 5.3 y is much longer than the length of any data set, the rate of ^{60}Co events is nearly constant within a single data set, so it is a good approximation simply to use $r_{^{60}\mathrm{Co}}(C) \times (1/2)^{t(D)/t_{1/2}}$ as the ^{60}Co rate for data set D.

The third term in Eq. (6.5) represents $0\nu\beta\beta$ decay of ^{130}Te modeled by a Gaussian centered at the double beta decay Q-value of ^{130}Te. The Gaussian function is normalized to unit area and is defined according to Eq. (6.8). Two precise measurements of the ^{130}Te double beta decay Q-value have recently been made: 2527.518 ± 0.013 keV [86] and 2527.01 ± 0.32 keV [87]. For this analysis, the Q-value has been taken to be 2527.518 keV. The uncertainty on the Q-value measurement has a negligible effect on this analysis since it is much smaller than the energy resolution of the CUORICINO detectors. For each subset, the Gaussian sigma is fixed to the value obtained from a Gaussian fit to the 2615 keV calibration peak for the same subset. The expected number of $0\nu\beta\beta$ decay events in the subset corresponding to channel C and data set D is related to the $0\nu\beta\beta$ decay rate, Γ, as

$$N_{0\nu\beta\beta}(C, D) = \Gamma \times N_{^{130}\mathrm{Te}}(C) \times T(C, D) \times \varepsilon(C, D), \tag{6.10}$$

where $N_{^{130}\mathrm{Te}}(C)$ is the number of ^{130}Te nuclei in crystal C and $\varepsilon(C, D)$ is the signal efficiency for channel C and data set D. The number of ^{130}Te nuclei in each crystal and the signal efficiencies are fixed in the fit.

The total expected number of events, N_{total}, is the sum of the expected number of events in each subset:

$$N_{\mathrm{total}} = \sum_{C,D} N_{\mathrm{bkg}}(C, D) + N_{^{60}\mathrm{Co}}(C, D) + N_{0\nu\beta\beta}(C, D). \tag{6.11}$$

The p.d.f., $f(E, C, D)$, represents a joint probability distribution in the three observables, energy (E), channel (C), and data set (D). It is normalized to unit probability, as any p.d.f. must be:

$$\sum_{C,D} \int_{E_{\mathrm{min}}}^{E_{\mathrm{max}}} f(E, C, D) dE = 1, \tag{6.12}$$

where $E_{\min}$ and $E_{\max}$ are the minimum and maximum energies in the fit range, respectively, and the C and D indices run over all channels and data sets, respectively. The p.d.f. enables a simultaneous fit because of the presence of the discrete observables, channel (C) and data set (D).

With the p.d.f. constructed in this way, the simultaneous fit is carried out following the standard procedure for an unbinned, extended maximum likelihood fit [88]. The likelihood function is constructed from the p.d.f. and the data as

$$\mathcal{L}(\Gamma, \vec{p}) = \text{Poisson}(N_{\text{obs}}; N_{\text{total}}) \times \prod_{i \in \{\text{data}\}} f(E_i, C_i, D_i), \tag{6.13}$$

where Γ, the decay rate, is the parameter of interest and $\vec{p}$ is the vector of all other floating parameters, and the index in the product runs over all data points with energies in the fit range. The fit contains seven floating parameters: Γ, three continuum background rates (for big, small, and enriched crystals), two ^{60}Co rates (for big and small crystals), and the ^{60}Co Gaussian mean. In an extended maximum likelihood fit, the number of events observed in the experiment, N_{obs}, is treated as a Poisson distributed observable, which contributes the factor of

$$\text{Poisson}(N_{\text{obs}}; N_{\text{total}}) = \frac{(N_{\text{total}})^{N_{\text{obs}}} e^{-N_{\text{total}}}}{N_{\text{obs}}!} \tag{6.14}$$

in the likelihood function. The best fit values of the floating parameters are those that maximize the likelihood function or, equivalently, minimize the negative log likelihood. The statistical uncertainties on the parameters are obtained from the variation of the negative log likelihood function around the minimum. The fit is implemented in RooFit [89], a toolkit for statistical data analysis built on ROOT [82, 83].

6.5 Fit results

Table 6.2 lists the results of the simultaneous fit applied to the CUORICINO data with a total exposure of 18.6 kg · y of ^{130}Te. No signal attributable to $0\nu\beta\beta$ decay is observed, and Γ is consistent with 0.

Figure 6.9 shows the fit projected onto the spectrum summed over all channels and data sets. Figure 6.10 shows the profile negative log likelihood function for the $0\nu\beta\beta$ decay rate, Γ. The profile negative log likelihood function is defined as the negative log likelihood function minimized over all other floating parameters with Γ held fixed at each point; similarly, the profile likelihood function is the likelihood function maximized over all parameters besides Γ. The profile likelihood function is used in Sect. 6.7 for setting an upper limit on Γ.

Table 6.2: Results of the CUORICINO $0\nu\beta\beta$ decay fit. Errors are statistical only.

Parameter	Fit result
Γ	$(-0.2 \pm 1.5) \times 10^{-25}$ y^{-1}
$r_{\rm bkg}$(big)	0.161 ± 0.005 counts/(keV · kg · y)
$r_{\rm bkg}$(small)	0.163 ± 0.014 counts/(keV · kg · y)
$r_{\rm bkg}$(enriched)	0.47 ± 0.05 counts/(keV · kg · y)
$r_{^{60}\rm Co}$(big)	3.1 ± 0.36 counts/(kg · y)
$r_{^{60}\rm Co}$(small)	2.9 ± 1.1 counts/(kg · y)
$r_{^{60}\rm Co}$(enriched)	0.0 counts/(kg · y) [fixed]
$\mu_{^{60}\rm Co}$	2506.5 ± 0.32 keV

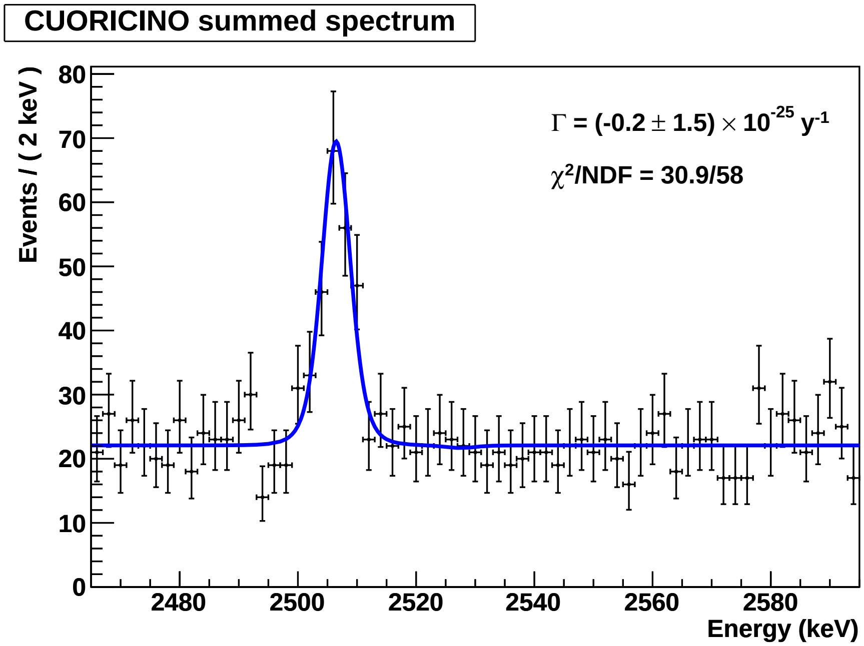

Figure 6.9: CUORICINO $0\nu\beta\beta$ decay fit projected onto the summed spectrum of all channels and data sets. The χ^2 and number of degrees of freedom (NDF) for this graph are displayed to provide a measure of the goodness-of-fit, although this χ^2 is not the quantity that is minimized to determine the fit parameters.

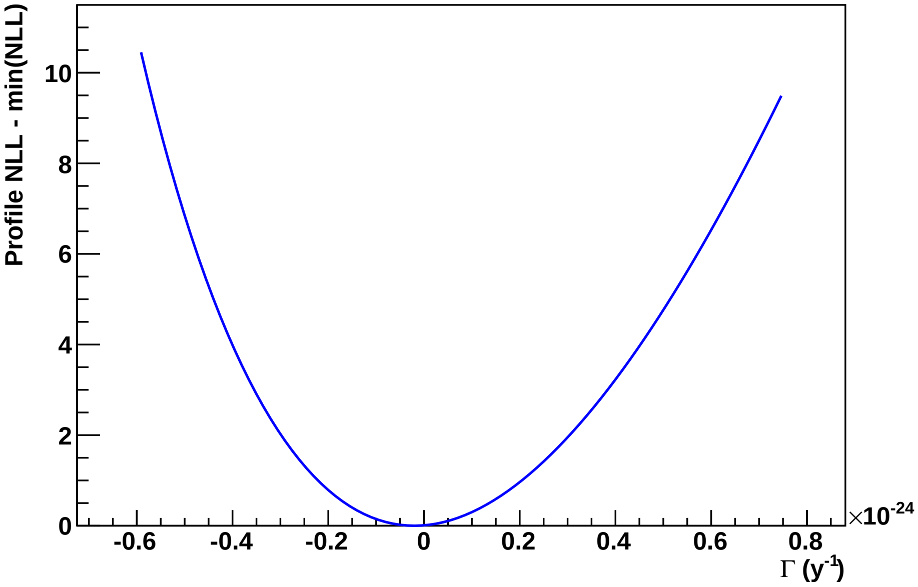

Figure 6.10: Profile negative log likelihood function. For each value of the decay rate, the value of the function is the negative log likelihood minimized over all other fit parameters.

Table 6.3: Contributions to the CUORICINO systematic error on the $0\nu\beta\beta$ decay rate, Γ. For comparison the statistical error is 1.5×10^{-25} y^{-1}.

Source	Systematic error on Γ
Energy scale uncertainty	0.21×10^{-25} y^{-1}
Background parameterization	0.27×10^{-25} y^{-1}
Efficiency uncertainty	3.3%
Total	0.34×10^{-25} y^{-1}

6.6 Systematic errors

The sources of systematic errors on the $0\nu\beta\beta$ decay rate result are uncertainty on the energy scale, uncertainty in the background shape, and uncertainty on the signal efficiency. The contributions from each of these sources to the systematic error are described in this section, and the procedure employed to incorporate systematic errors in the upper limit on the $0\nu\beta\beta$ decay rate is detailed in Sect. 6.7. Table 6.3 lists the systematic error due to each source of uncertainty. The total systematic error is taken as the quadrature sum of the individual contributions because the errors due to different sources are expected to be uncorrelated.

The uncertainty on the energy scale arises from lack of knowledge of the true calibration function that maps stabilized pulse amplitudes to particle energies, as well as statistical uncertainty in the calibration coefficients for the chosen calibration function due to fitting peaks in the calibration spectrum. At the Q-value for ^{130}Te, the calibration uncertainty for CUORICINO has been estimated to be about $\delta E = 0.4$ keV, but I use a more conservative value of $\delta E = 0.8$ keV taken from the discrepancy between the best fit value of the Gaussian mean of the ^{60}Co peak and the expected position of the peak. Because of the energy scale uncertainty, a peak due to $0\nu\beta\beta$ decay could appear in the spectrum anywhere within δE of the Q-value. To study the dependence of the fit result on the expected location of the $0\nu\beta\beta$ decay peak, the center of the $0\nu\beta\beta$ Gaussian is varied between $Q - \delta E$ and $Q + \delta E$ in steps of $\delta E/10$, and the fit is performed for each step. The energy scale contribution to the systematic error is taken as the sample standard deviation of the resulting fit values for Γ, which is 0.21×10^{-25} y^{-1}.

The uncertainty in the background shape stems from the limited number of events in the $0\nu\beta\beta$ decay region of interest. There are not enough data points to determine the shape of the continuum background with certainty. In the fit the continuum background is parameterized as a uniform function of energy. The sensitivity of the fit result for Γ to the background parameterization is estimated by carrying out the fit using higher order polynomials to parameterize the background. In addition to the zeroth order polynomial background function, a first order and a second order polynomial background are considered (Fig. 6.11). Furthermore, the range of the fit is chosen to be 2465–2595 keV, but there is a degree of arbitrariness in the choice of fit range. The fit must only avoid background peaks above and below this range. Changing the fit range affects the background level slightly. The sensitivity of the fit result for Γ to the fit range is estimated by varying the range used in the fit by $+15$ and $+30$ keV on the lower end and by -15 and -30 keV on the upper end of the range. The systematic uncertainty on Γ due to the background parameterization is taken as the sample standard deviation of the fit results for Γ when the degree of the background polynomial and the fit range are jointly varied. The resulting systematic error is 0.27×10^{-25} y^{-1}.

The uncertainty on the signal efficiency is reported in Sect. 6.2 to be 3.3%.

Summing the systematic error contributions in quadrature, the total systematic error on Γ is $\sigma_\Gamma(\text{syst.}) = 0.34 \times 10^{-25}$ y^{-1}, which is small compared to the statistical error of $\sigma_\Gamma(\text{stat.}) = 1.5 \times 10^{-25}$ y^{-1}. The result of the CUORICINO measurement of the $0\nu\beta\beta$ decay rate of ^{130}Te is

$$\boxed{\Gamma^{0\nu\beta\beta}(^{130}\text{Te}) = (-0.2 \pm 1.5 \text{ (stat.)} \pm 0.3 \text{ (syst.)}) \times 10^{-25} \text{ y}^{-1}.} \tag{6.15}$$

6.7 Limit technique and results

Since no signal for $0\nu\beta\beta$ decay is observed, an upper limit on the $0\nu\beta\beta$ decay rate is set. The calculation of the limit is complicated by two factors: First, the best fit value is near a physical boundary since the decay rate is intrinsically non-negative, and in principle the best fit value is equally likely to be positive or negative if $\Gamma = 0$ in reality. Second, there are six nuisance parameters (the floating fit parameters besides Γ). To deal with these two difficulties and to incorporate systematic errors, this analysis uses a Bayesian technique that is based on the profile likelihood function. Denoting by $\hat{\Gamma}$ and $\hat{\vec{p}}$ the values of the fit parameters that minimize the negative log likelihood function, it can be shown that the quantity,

$$\chi^2_{\text{stat.}}(\Gamma) \equiv 2\left(\min_{\vec{p}}\{-\ln\mathcal{L}(\Gamma,\vec{p})\} - (-\ln\mathcal{L}(\hat{\Gamma},\hat{\vec{p}}))\right) \tag{6.16}$$

follows a χ^2 distribution with one degree of freedom (for the true value of Γ) [90]. To include systematic errors, the following quantity is constructed:

$$\chi^2_{\text{syst.}}(\Gamma) = \frac{(\Gamma - \hat{\Gamma})^2}{\sigma^2_{\text{syst}}}. \tag{6.17}$$

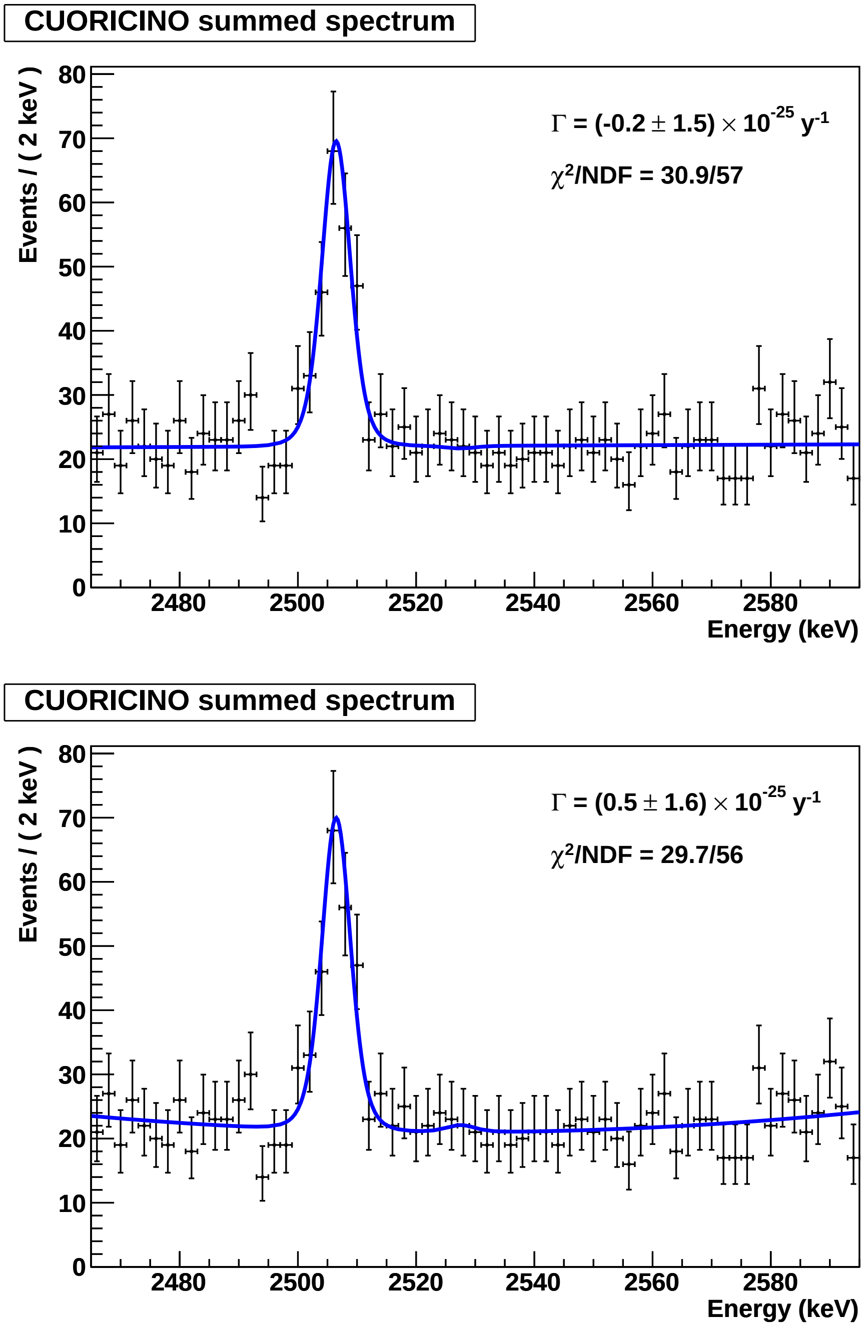

Figure 6.11: Fit projected onto the summed spectrum with a first order polynomial background (top) and a second order polynomial background (bottom). The χ^2 and number of degrees of freedom (NDF) for these graphs are displayed to provide a measure of the goodness-of-fit, although this χ^2 is not the quantity that is minimized to determine the fit parameters.

The systematic error, $\sigma_{\text{syst.}}$, consists of absolute and fractional errors, which sum in quadrature as $\sigma^2_{\text{syst.}} = \sigma^2_{\text{syst.,abs.}} + \sigma^2_{\text{syst.,frac.}}\Gamma^2$. The absolute systematic error is equal to the quadrature sum of the errors due to the energy scale and the background parameterization, while the fractional systematic error is equal to the uncertainty on the efficiency. Motivated by the goal of summing the statistical and systematic errors in quadrature, the χ^2-like quantities due to statistical and systematic errors are combined as

$$\frac{1}{\chi^2_{\text{total}}} = \frac{1}{\chi^2_{\text{stat.}}} + \frac{1}{\chi^2_{\text{syst.}}}. \tag{6.18}$$

The function $\chi^2_{\text{total}}/2$ is similar to the profile negative log likelihood (cf. Eq. (6.16)) but is broadened to account for the systematic error (Fig. 6.12). Therefore, $\chi^2_{\text{total}}/2$ is used as the profile negative log likelihood function would be in a Bayesian analysis. It is exponentiated to obtain the profile likelihood function, $e^{-\chi^2_{\text{total}}(\Gamma)/2}$, which is treated as a p.d.f. for Γ with a Bayesian prior, $\pi(\Gamma)$, equal to zero in the unphysical ($\Gamma < 0$) region and flat in the physical ($\Gamma \geq 0$) region:

$$\pi(\Gamma) = \begin{cases} 0, & \Gamma < 0 \\ 1, & \Gamma \geq 0 \end{cases} \tag{6.19}$$

Figure 6.13 shows the profile likelihood function before and after the inclusion of systematic errors. The upper limit on the decay rate at a given confidence level (C.L.) is the value of Γ^{limit} such that

$$\frac{\int_0^{\Gamma^{\text{limit}}} e^{-\chi^2_{\text{total}}(\Gamma)/2} d\Gamma}{\int_0^{\infty} e^{-\chi^2_{\text{total}}(\Gamma)/2} d\Gamma} = \text{C.L.} \tag{6.20}$$

Following this procedure, the upper limit on the $0\nu\beta\beta$ decay rate from CUORICINO is determined to be

$$\boxed{\Gamma^{0\nu\beta\beta}(^{130}\text{Te}) < 2.7 \times 10^{-25}\ \text{y}^{-1}\ (90\%\ \text{C.L.}),} \tag{6.21}$$

and the corresponding lower limit on the partial half-life for $0\nu\beta\beta$ decay, calculated as $T^{\text{limit}}_{1/2} = \ln 2/\Gamma^{\text{limit}}$, is

$$\boxed{T^{0\nu\beta\beta}_{1/2}(^{130}\text{Te}) > 2.6 \times 10^{24}\ \text{y}\ (90\%\ \text{C.L.}).} \tag{6.22}$$

Without systematic errors, the limits would be

$$\Gamma^{0\nu\beta\beta}(^{130}\text{Te}) < 2.6 \times 10^{-25}\ \text{y}^{-1}\ (90\%\ \text{C.L.}) \tag{6.23}$$

and

$$T^{0\nu\beta\beta}_{1/2}(^{130}\text{Te}) > 2.7 \times 10^{24}\ \text{y}\ (90\%\ \text{C.L.}). \tag{6.24}$$

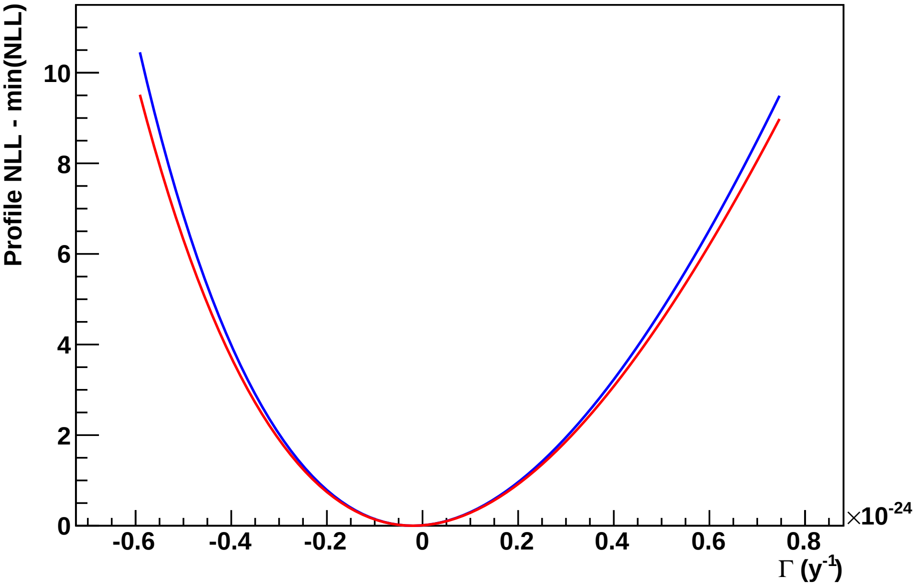

Figure 6.12: Profile negative log likelihood function without (blue) and with (red) systematic errors included. The blue curve is equal to $\chi^2_{\text{stat.}}(\Gamma)/2$, and the red curve is equal to $\chi^2_{\text{total}}(\Gamma)/2$.

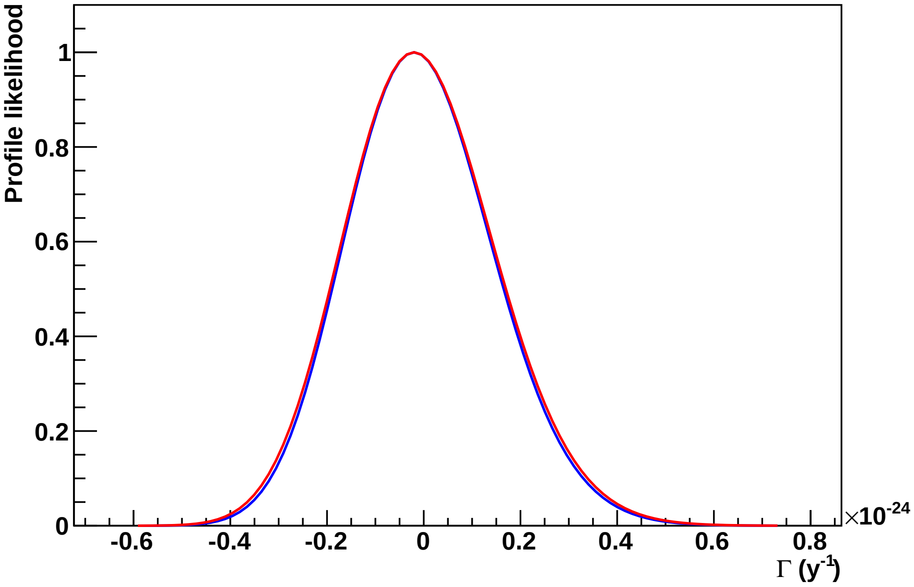

Figure 6.13: Profile likelihood function without (blue) and with (red) systematic errors included. The blue curve is equal to $e^{-\chi^2_{\text{stat.}}(\Gamma)/2}$, and the red curve is equal to $e^{-\chi^2_{\text{total}}(\Gamma)/2}$.

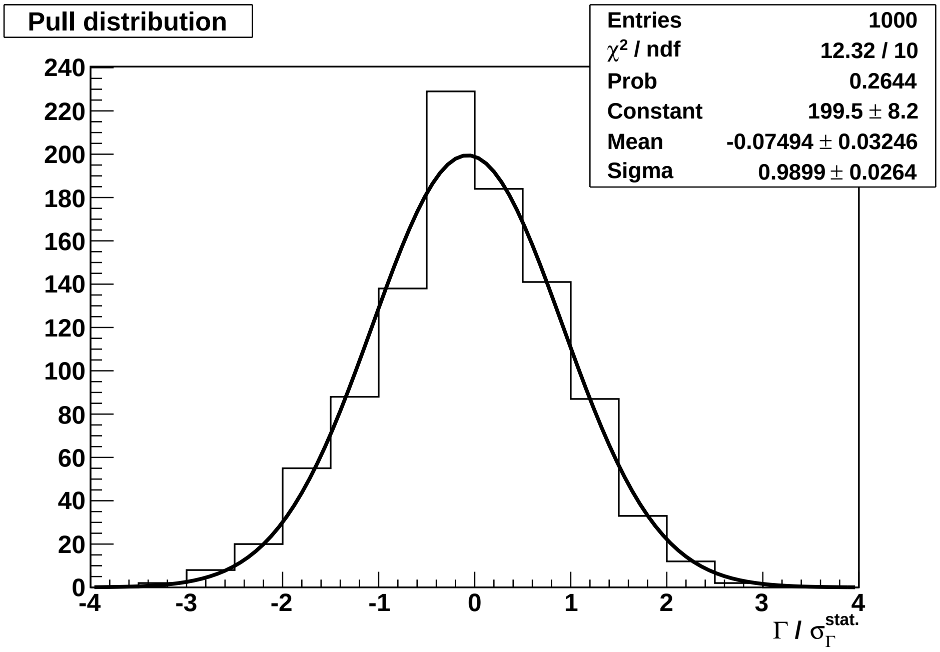

Figure 6.14: Pull distribution from toy Monte Carlo simulations. The pull distribution for an unbiased fit procedure with correct error estimation should be Gaussian with mean equal to 0 and sigma equal to 1.

6.8 Fit validation with toy Monte Carlo simulations

For small data samples, a maximum likelihood fit is not guaranteed to produce unbiased parameter estimates and to yield correct errors [90]. To check the fitting procedure for bias and the correctness of the errors returned by the fitter, 1000 toy Monte Carlo data sets are generated from the p.d.f. described in Sect. 6.4 with $\Gamma = 0$ and the background parameters set to the values in Table 6.2. The fit is then performed on each of the simulated data sets. The pull distribution is plotted in Fig. 6.14. For each toy data sample, the pull is equal to $(\hat{\Gamma} - \Gamma_{\text{true}})/\sigma_\Gamma^{\text{stat.}}$, where $\hat{\Gamma}$ is the best fit value, Γ_{true} is the true value used to generate the Monte Carlo ($\Gamma_{\text{true}} = 0$ in this case), and $\sigma_\Gamma^{\text{stat.}}$ is the statistical error on Γ calculated by the fitter. If the fit is unbiased and the errors are correctly estimated, the distribution of the pulls should be Gaussian with mean equal to 0 and sigma equal to 1. Based on the Gaussian fit in Fig. 6.14, the pull distribution is consistent with these criteria except that the mean is only marginally consistent with 0 (the mean is 2.3σ less than 0), which may indicate a negligibly small negative bias.

The toy Monte Carlo data are also used to quantify the expected sensitivity of the experiment by calculating the distribution of possible results from many identical experiments. Figure 6.15 shows the distributions of best fit values and errors for the decay rate for the 1000 Monte Carlo data sets with the true decay rate equal to 0. Figure 6.16 plots the distributions of upper limits on $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te})$ and lower limits on $T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te})$ obtained from the simulated data sets. Given the distribution of possible results, the results obtained for CUORICINO are reasonable in the sense that they fall at a point of high probability within the distribution. The distributions of decay rate limits and half-life limits are rather broad, with r.m.s. values greater than 1/3 of the mean. This shows that the limit obtained from an experiment like CUORICINO is highly dependent on statistical fluctuations of the background and can easily deviate by up to $\sim$50% from the expected sensitivity of the experiment, where the expected sensitivity is defined as the expectation value of the limit. On the other hand, as a measure of the sensitivity of the experiment, the statistical error on the decay rate is robust, varying by only $\sim$10% from its average value. The statistical fluctuation of the best fit value for Γ leads to the fluctuation in the half-life limit, as demonstrated in Fig. 6.17.

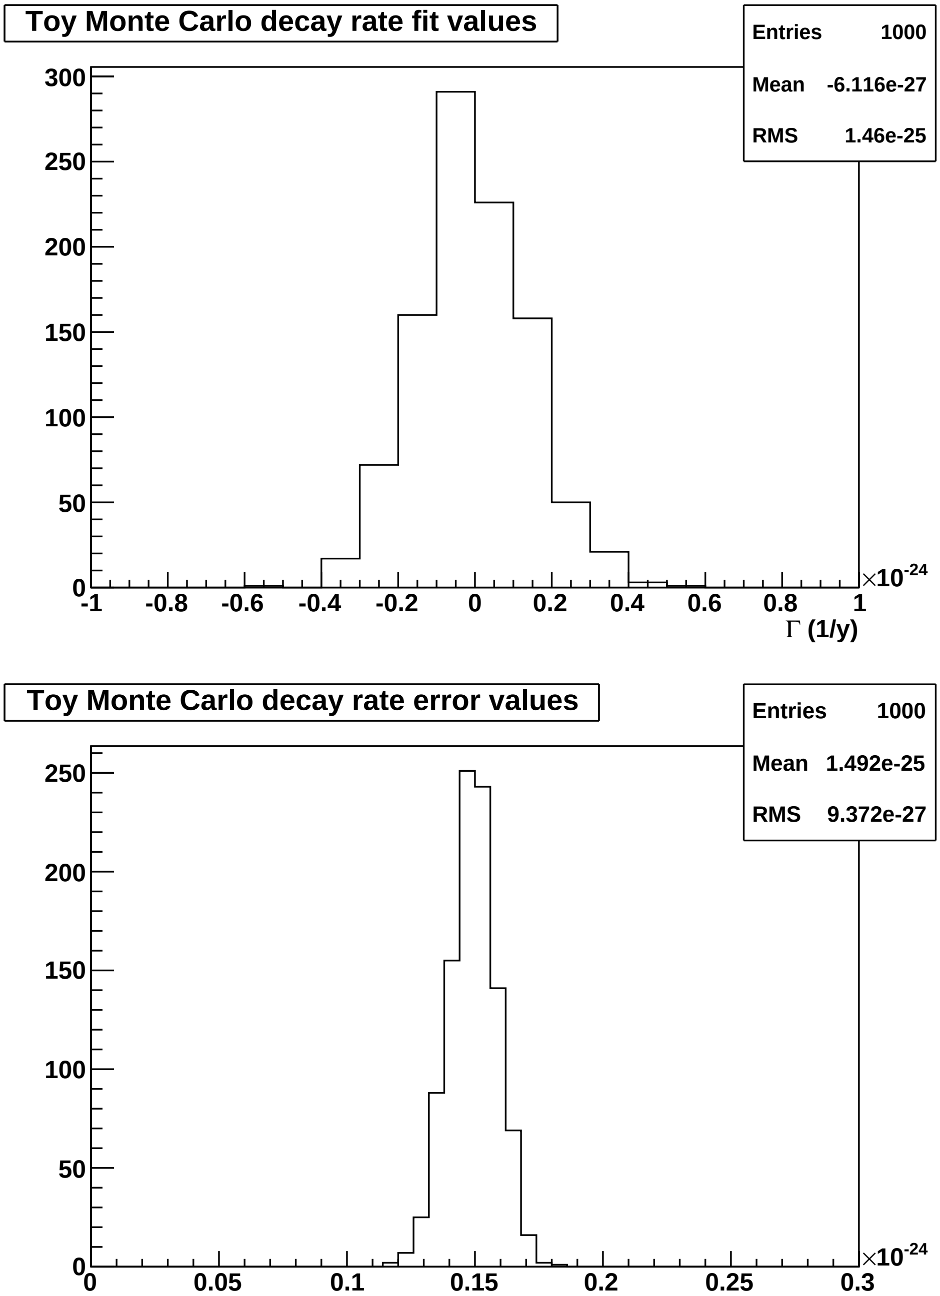

Figure 6.15: Distributions of decay rate best fit values (top) and errors (bottom) from toy Monte Carlo simulations.

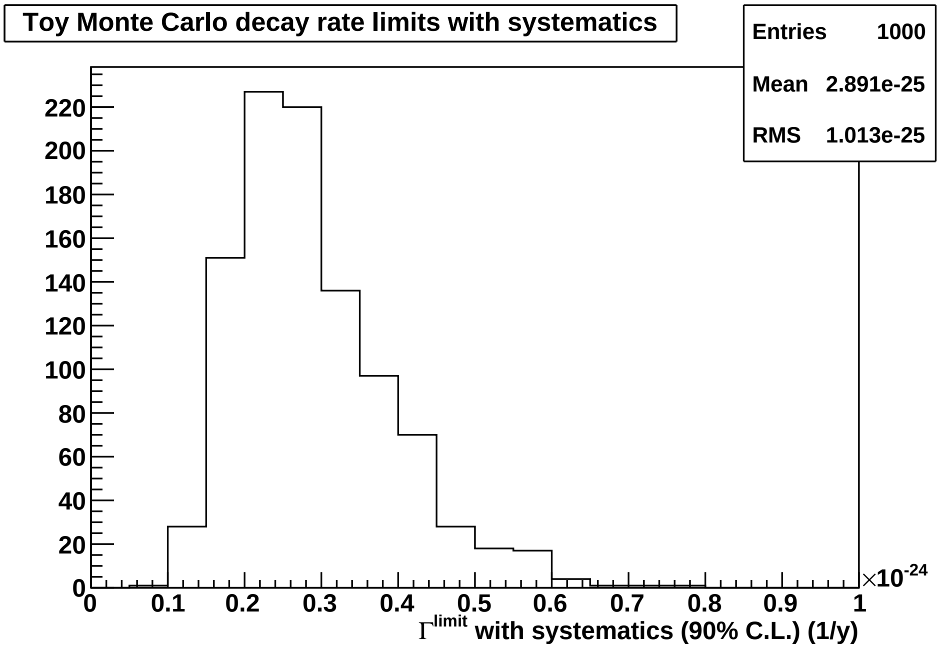

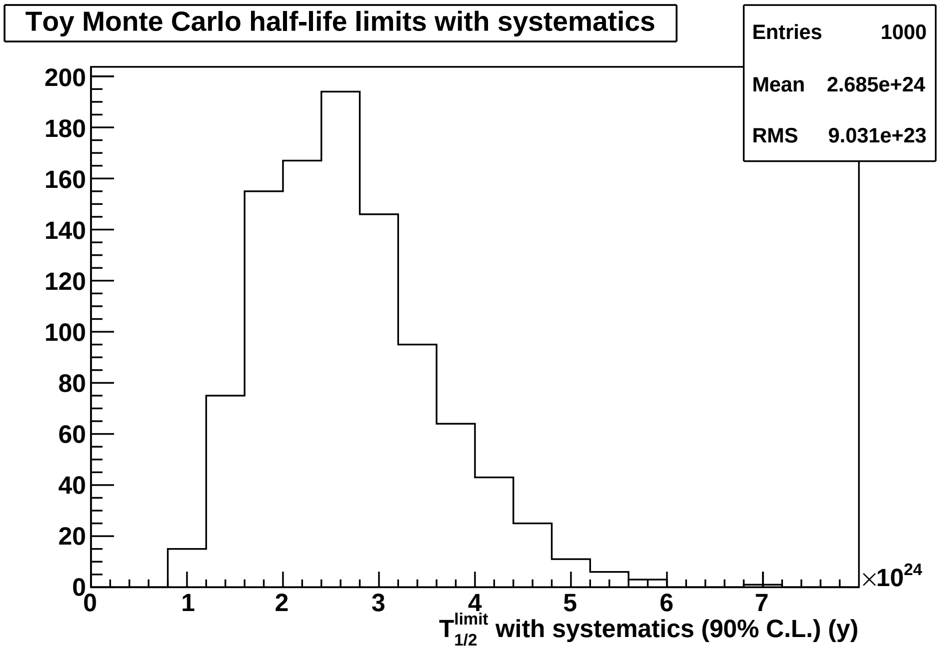

Figure 6.16: Distributions of decay rate limits (top) and half-life limits (bottom), both with systematic errors included, from toy Monte Carlo simulations.

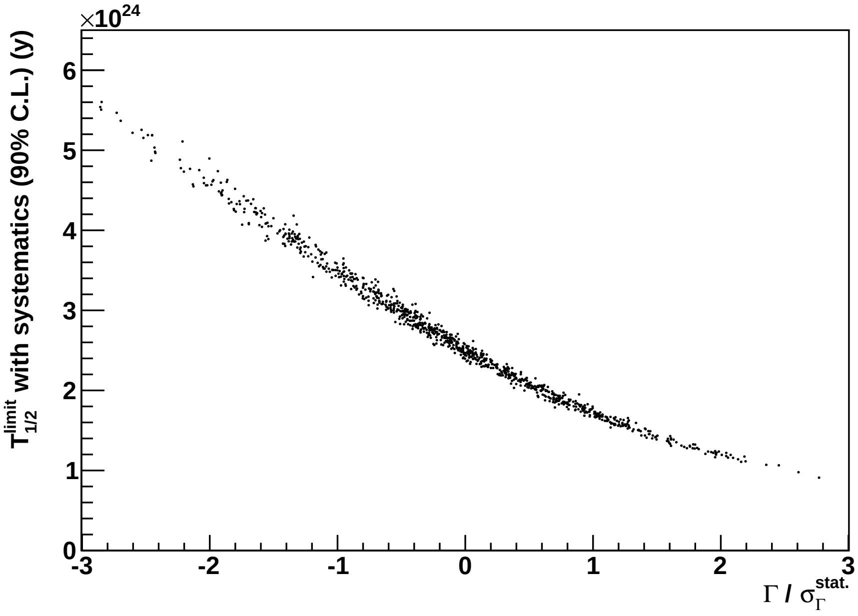

Figure 6.17: Dependence of half-life limit on best fit value from toy Monte Carlo simulations.

Chapter 7

Three Towers Test and CUORICINO combined analysis

In this chapter the $0\nu\beta\beta$ decay analysis performed for CUORICINO in Chapter 6 is extended to a combined analysis of CUORICINO and the Three Towers Test. The analysis technique remains the same: A simultaneous fit is performed to each channel and data set. The 26 channels and four data sets from the TTT add $26 \times 4 = 104$ new subsets to the simultaneous fit. As indicated in Table 7.1, the TTT data contribute an additional 7% of ^{130}Te exposure with a slightly lower background level than CUORICINO.

In Sect. 7.1 the cuts applied to the TTT data are described, and in Sect. 7.2 their efficiencies are calculated. Section 7.3 contains information on the energy resolutions of the TTT detectors. The results from the combined fit and systematic errors are given in Sect. 7.4. Finally, Sect. 7.5 contains the limit on $T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te})$ obtained from the combined analysis.

7.1 TTT cuts

In the combined CUORICINO and TTT analysis, the cuts in Sect. 6.1 are applied to the CUORICINO data. In this section the cuts applied to the Three Towers Test data are described. For the TTT detectors with two thermistors, the thermistor with better resolution was chosen to be used in the analysis. The selected channels for crystals with two thermistors are (in the order shown on the channel map in Fig. 3.12): 49, 13, 1, 17, 27, 16, 14, 3, 5, 18, 19, 61, and 58. Intervals of data that do not satisfy the data selection criteria in Sect. 5.3 are rejected. Double counting of re-triggered pulses is avoided as described in Sect. 4.5. The offline heater flagging code is not applied to the TTT data because the Apollo DAQ flags heater events with 100% efficiency. The anti-coincidence

Table 7.1: Exposures and background levels for CUORICINO and the Three Towers Test. The TTT data augment the CUORICINO data with about 7% more exposure with a lower background.

Crystal type	Exposure ($\mathrm{kg} \cdot \mathrm{y}$ ^{130}Te)	Background (counts/($\mathrm{keV} \cdot \mathrm{kg} \cdot \mathrm{y}$))
CUORICINO big	15.78	0.161 ± 0.005
CUORICINO small, natural Te	2.02	0.163 ± 0.014
CUORICINO ^{130}Te enriched	0.77	0.47 ± 0.05
CUORICINO total	18.57	
Three Towers Test	1.32	0.129 ± 0.015

cut applied to the TTT data rejects events that occurred within ±50 ms of a non-heater event on another detector that had an energy greater than 50 keV; the time difference is based on the trigger times of the events. A fixed dead time is imposed around heater events, starting 4 s before a heater event and lasting until 4 s after the heater event. The amount of dead time before each heater event is longer than for CUORICINO (it was 3.1 s for CUORICINO) because the acquisition window for the TTT was longer than for CUORICINO.

The pulse shape cut used for CUORICINO (Sect. 4.9) has not been computed for the Three Towers Test. Since this analysis is concerned only with events that have energies close to the ^{130}Te Q-value, a greatly simplified form of the pulse shape cut is effective. The cut is based on the ratio of the pulse amplitude evaluated in the time domain to the pulse amplitude evaluated in the frequency domain, $a_{\mathrm{TD}}/a_{\mathrm{FD}}$, which should be close to 1 for single, real pulses. This ratio is plotted versus energy in Fig. 7.1. From the figure, it can be seen that the main cluster of points, corresponding to good pulses, is above 0.999. Therefore, as a pulse shape cut for the TTT data, $a_{\mathrm{TD}}/a_{\mathrm{FD}} > 0.999$ is required.

7.2 TTT signal efficiency

The efficiency for detection of $0\nu\beta\beta$ decay events in the Three Towers Test is evaluated by the same procedure used for CUORICINO and described in Sect. 6.2. The contributions to the loss of signal efficiency are summarized in Table 7.2. As for CUORICINO, the largest source of inefficiency is the possibility that one of

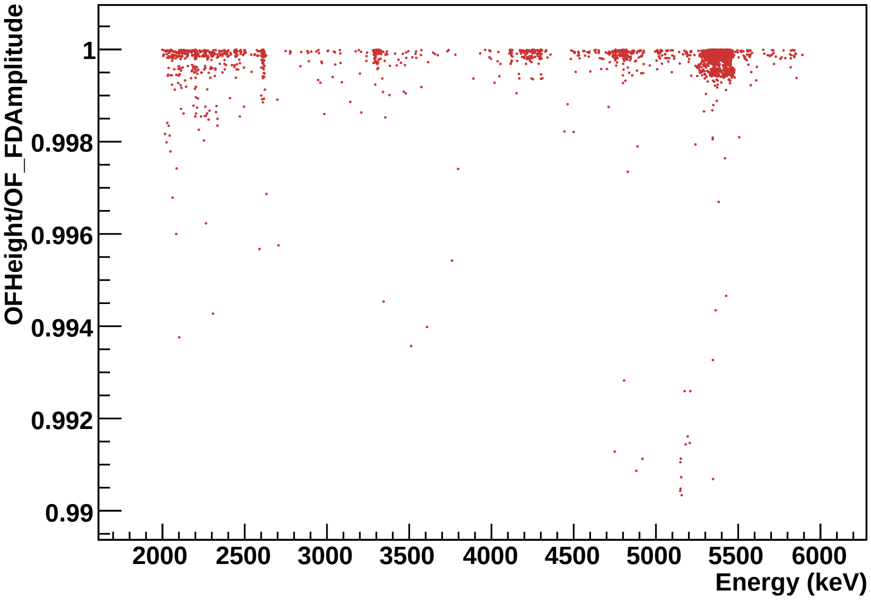

Figure 7.1: The ratio of the pulse amplitude evaluated in the time domain to the pulse amplitude evaluated in the frequency domain, used as a pulse shape parameter for the Three Towers Test data, plotted versus energy. The data plotted are from all channels and all background runs in data set 1004. The pulse shape cut applied for the TTT requires the ratio to be greater than 0.999.

Table 7.2: Contributions to the TTT $0\nu\beta\beta$ decay signal efficiency. The efficiency of the anti-coincidence cut depends on whether the detector was exposed to the ^{40}K sources. There is an efficiency loss due to pileup with heater pulses only for channels with a working heater and only for data sets 1001–1003 because the heaters were disconnected during data set 1004.

Source	Signal efficiency (%)
Escape of a β	86.3 ± 3.5
Pulse shape cut	97.4 ± 1.1
Anti-coincidence cut	97.48 ± 0.07 (data set 1001) 99.7 ± 0.3 (data sets 1002–1004)
Noise	99.1 ± 0.1
Pileup with heater pulses	97.3 (with heater) 100 (without heater)

the electrons emitted in a $0\nu\beta\beta$ decay escapes from the crystal. Since the TTT crystals are the same size as the CUORICINO $5 \times 5 \times 5$ cm^3 crystals, the efficiency due to escape of a β is $(86.3 \pm 3.5)\%$ from Sect. 6.2.1.

The efficiencies of the pulse shape cut and the anti-coincidence cut are estimated by the same type of simultaneous fits described in Section 6.2.2 and Sect. 6.2.3. The fit to determine the pulse shape cut efficiency for the TTT is shown in Fig. 7.2, and the fit results are $\varepsilon_{\mathrm{PS}} = (97.4 \pm 1.1)\%$ and $\varepsilon_{\mathrm{PS}}^{\mathrm{bkg}} = (65 \pm 1)\%$. The efficiency of the anti-coincidence cut depends on whether the ^{40}K sources were inserted because the event rate was much higher with the ^{40}K sources present. The sources were present during data set 1001 but not during data sets 1002–1004. The fits to determine the anti-coincidence cut efficiencies for data set 1001 and data sets 1002–1004 are shown in Fig. 7.3 and Fig. 7.4, respectively. The fit results are $\varepsilon_{\mathrm{AC}} = (97.48 \pm 0.07)\%$ for data set 1001 and $\varepsilon_{\mathrm{AC}} = (99.7 \pm 0.3)\%$ for data sets 1002–1004.

For the efficiency loss due to noise, the average value obtained for CUORICINO, $(99.1 \pm 0.1)\%$, is taken as an estimate for the TTT. Since the heaters were inoperative for 10 of the 26 crystals in the Three Towers Test and the heaters were disconnected during data set 1004, it is not possible to make a comprehensive evaluation of the efficiency loss due to noise using heater events as was described in Sect. 6.2.4. Furthermore, because efficiency loss due to noise was

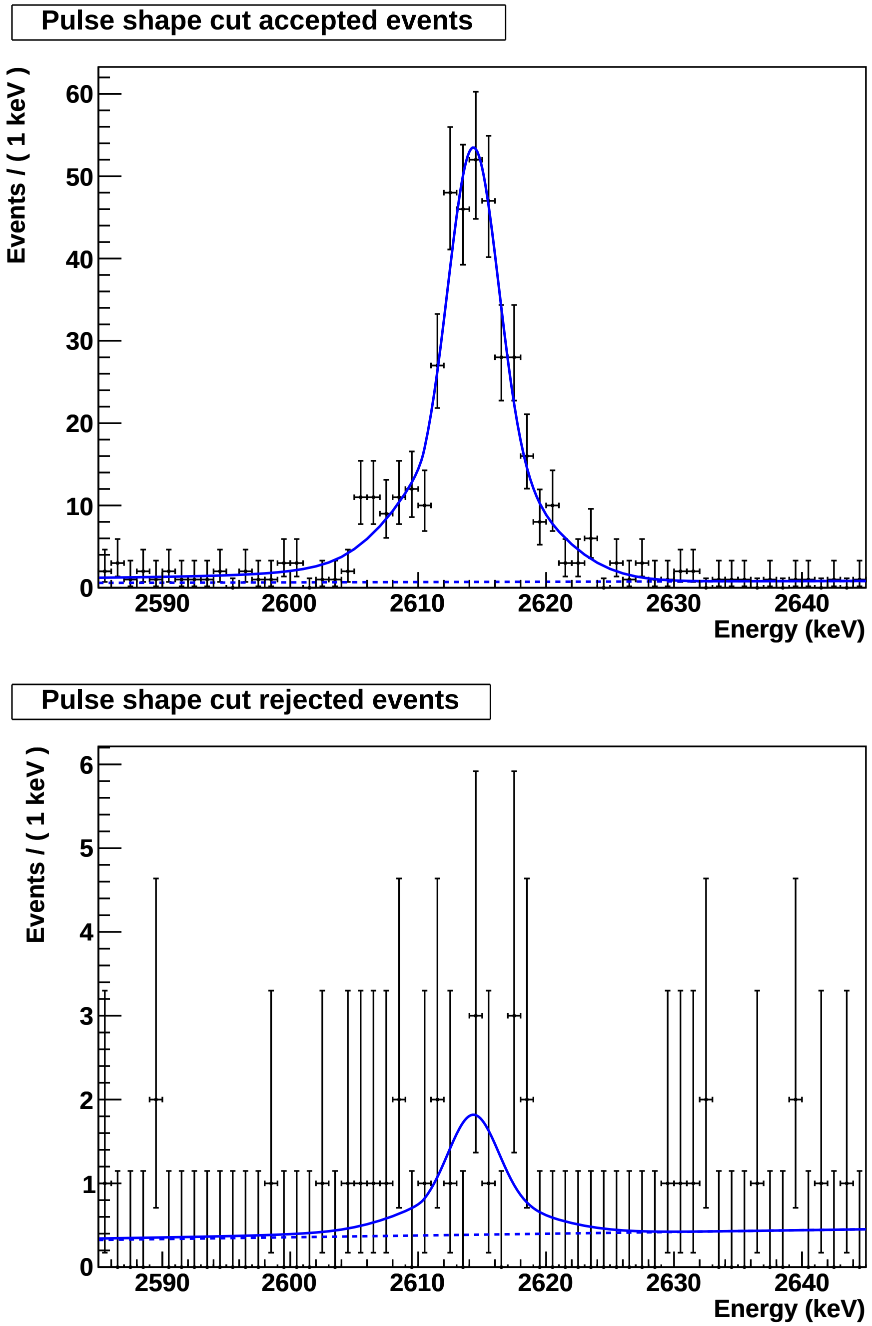

Figure 7.2: Fit to determine the pulse shape cut efficiency for the Three Towers Test. On the top are the events that are accepted by the pulse shape cut, and on the bottom are the events that are rejected. The efficiency of the cut is determined from a simultaneous fit to the two spectra.

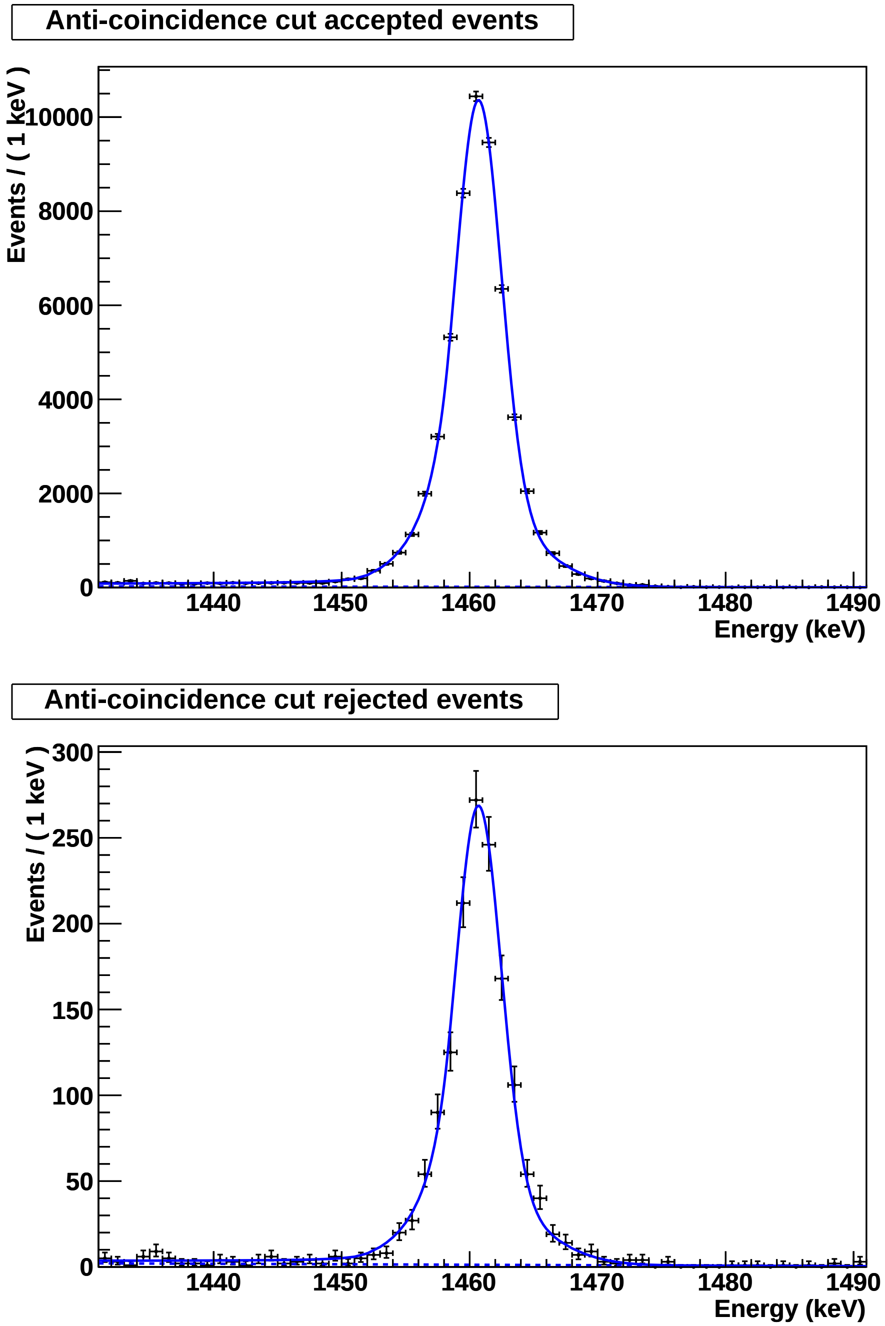

Figure 7.3: Fit to determine the anti-coincidence cut efficiency for data set 1001 of the Three Towers Test, during which the detector was exposed to ^{40}K sources. On the top are the events that are accepted by the anti-coincidence cut, and on the bottom are the events that are rejected. The efficiency of the cut is determined from a simultaneous fit to the two spectra.

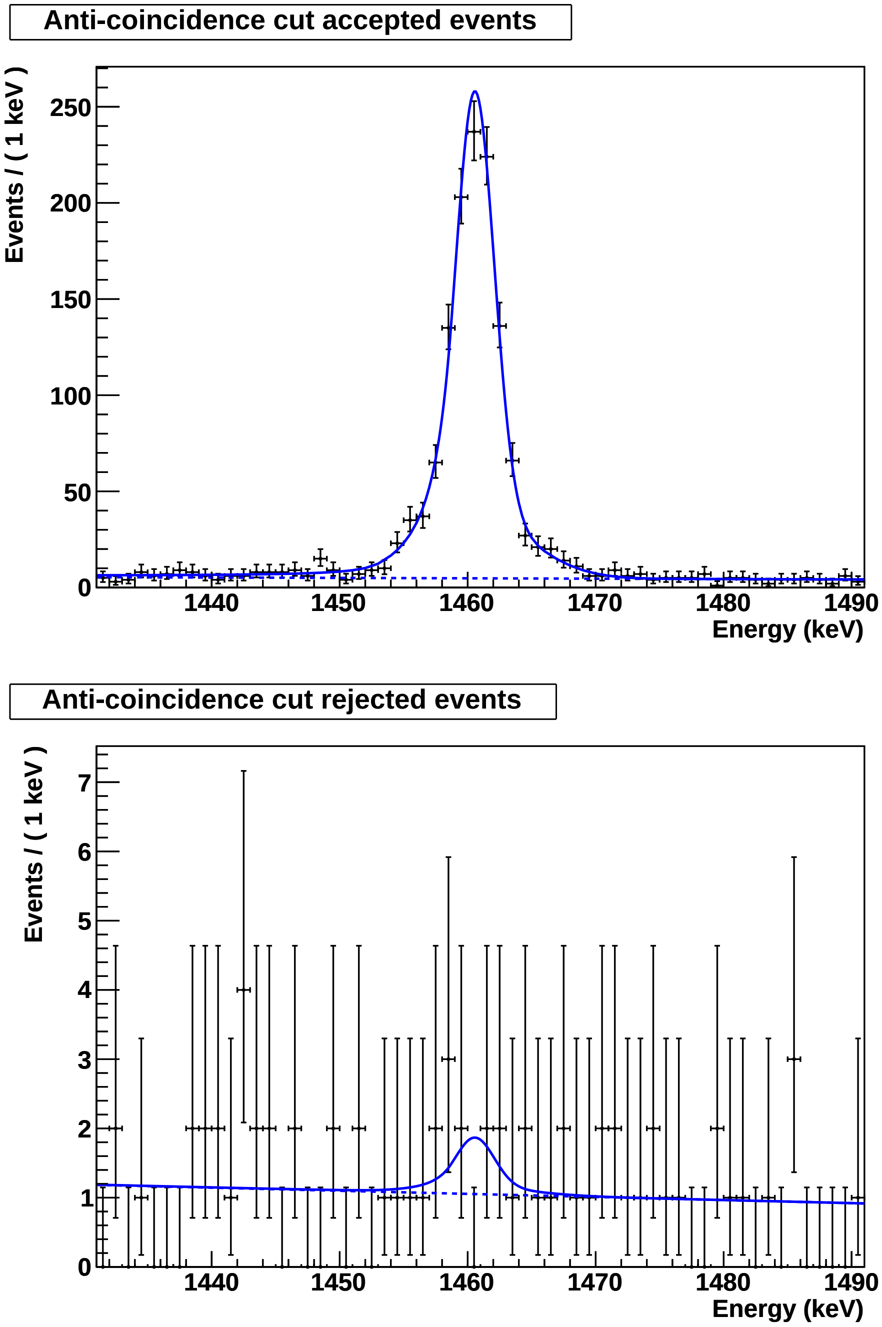

Figure 7.4: Fit to determine the anti-coincidence cut efficiency for data sets 1002–1004 of the Three Towers Test, during which the ^{40}K sources were not present. On the top are the events that are accepted by the anti-coincidence cut, and on the bottom are the events that are rejected. The efficiency of the cut is determined from a simultaneous fit to the two spectra.

demonstrated not to be a significant problem for CUORICINO and because the TTT data were closely monitored for noise during data taking, the average value of the noise efficiency from CUORICINO is a good estimate for the Three Towers Test.

The fixed dead time of ±4 s imposed around each heater pulse contributes an efficiency of $(300 - 8)/300 \approx 97.3\%$ for the 16 channels with a working heater in data sets 1001–1003. The corresponding efficiency is 100% for the 10 channels without a working heater and for all channels in data set 1004, during which the heaters were disconnected.

7.3 TTT energy resolutions

As for CUORICINO, the energy resolutions of the TTT detectors are measured using calibration data collected while the detectors were exposed to a ^{232}Th source. For each detector the resolution of the 2615 keV calibration peak, measured with an unbinned maximum likelihood fit with Gaussian signal shape and a linear background, is used as the detector resolution for the $0\nu\beta\beta$ decay fit. The distribution of FWHM resolutions obtained for each channel and data set is shown in Fig. 7.5. The distribution is similar to that of the CUORICINO big crystals (Fig. 6.8 (top)).

7.4 Combined fit results and systematic errors

The same p.d.f. defined in Eq. (6.5) is used to fit the combination of TTT and CUORICINO data. The TTT crystals are treated as an additional crystal type; that is, the TTT crystals share a background rate and a ^{60}Co rate. There are nine floating parameters in the fit: four continuum background rates, three ^{60}Co rates (no ^{60}Co peak is included for the CUORICINO enriched crystals due to limited statistics), the ^{60}Co Gaussian mean, and the ^{130}Te $0\nu\beta\beta$ decay rate.

The results of the combined fit are shown in Table 7.3, and the fit projected on the summed spectrum is shown in Fig. 7.6. The best fit value for the decay rate is $\Gamma = (-0.6 \pm 1.4 \text{ (stat.)}) \times 10^{-25}$ y^{-1}. The statistical error is improved modestly compared with the CUORICINO-only fit, as expected from adding 7% more exposure.

Systematic uncertainties are investigated following the same approach described in Sect. 6.6. The contributions of each of the three sources of systematic errors are listed in Table 7.4. Adding the contributions in quadrature yields a total systematic error of $\sigma_\Gamma(\text{syst.}) = 0.35 \times 10^{-25}$ y^{-1}, which is small compared to the statistical error of $\sigma_\Gamma(\text{stat.}) = 1.4 \times 10^{-25}$ y^{-1}. The result of the combined CUORICINO and Three Towers Test measurement of the $0\nu\beta\beta$ decay rate of

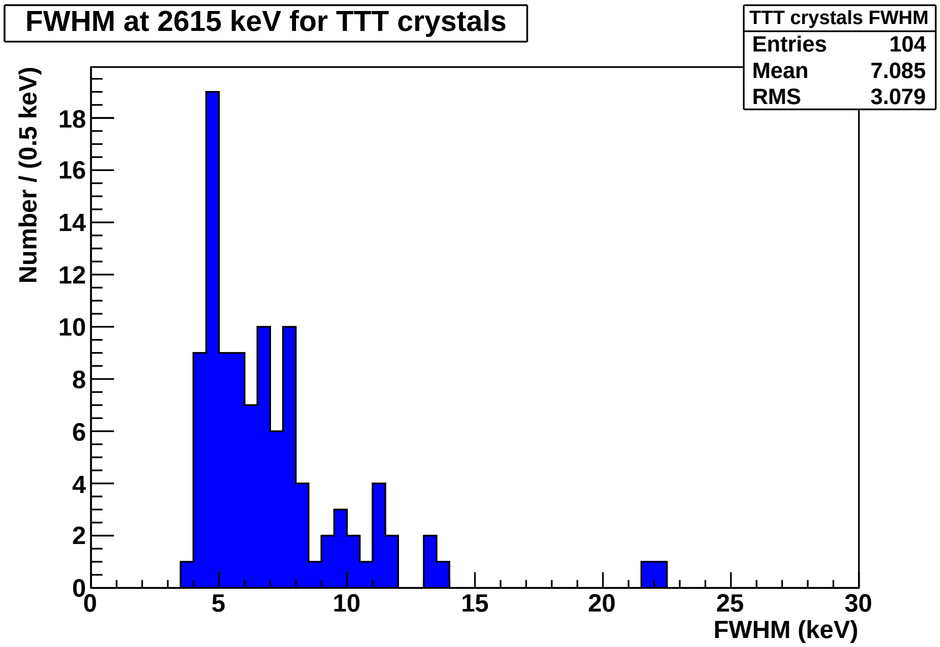

Figure 7.5: Distribution of energy resolutions for the TTT detectors. There is one entry for each channel and data set.

Table 7.3: Results of the combined CUORICINO and TTT $0\nu\beta\beta$ decay fit. Errors are statistical only.

Parameter	Fit result
Γ	$(-0.6 \pm 1.4) \times 10^{-25}$ y^{-1}
r_{bkg}(big)	0.161 ± 0.005 counts/(keV $\cdot$ kg $\cdot$ y)
r_{bkg}(small)	0.163 ± 0.014 counts/(keV $\cdot$ kg $\cdot$ y)
r_{bkg}(enriched)	0.47 ± 0.05 counts/(keV $\cdot$ kg $\cdot$ y)
r_{bkg}(TTT)	0.129 ± 0.015 counts/(keV $\cdot$ kg $\cdot$ y)
$r_{^{60}\mathrm{Co}}$(big)	3.1 ± 0.36 counts/(kg $\cdot$ y)
$r_{^{60}\mathrm{Co}}$(small)	2.9 ± 1.1 counts/(kg $\cdot$ y)
$r_{^{60}\mathrm{Co}}$(enriched)	0.0 ± 0.0 counts/(kg $\cdot$ y) [fixed]
$r_{^{60}\mathrm{Co}}$(TTT)	0.3 ± 0.6 counts/(kg $\cdot$ y)
$\mu_{^{60}\mathrm{Co}}$	2506.5 ± 0.32 keV

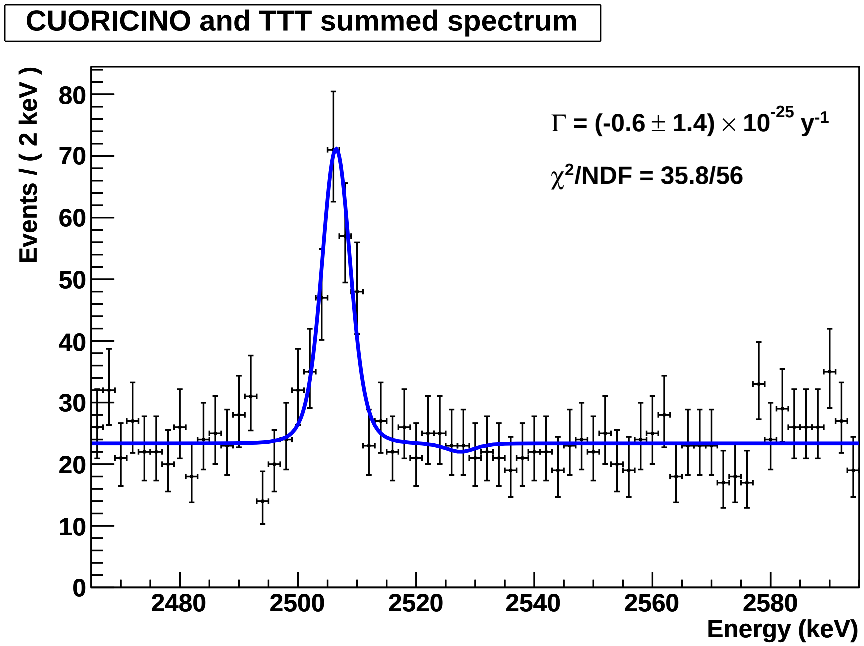

Figure 7.6: Combined CUORICINO and TTT $0\nu\beta\beta$ decay fit projected onto the summed spectrum of all channels and data sets. The χ^2 and number of degrees of freedom (NDF) for this graph are displayed to provide a measure of the goodness-of-fit, although this χ^2 is not the quantity that is minimized to determine the fit parameters.

Table 7.4: Contributions to the systematic error on the combined CUORICINO and TTT measurement of the $0\nu\beta\beta$ decay rate, Γ. For comparison the statistical error is 1.4×10^{-25} y^{-1}.

Source	Systematic error on Γ
Energy scale uncertainty	0.15×10^{-25} y^{-1}
Background parameterization	0.32×10^{-25} y^{-1}
Efficiency uncertainty	3.3%
Total	0.35×10^{-25} y^{-1}

^{130}Te is

$$\boxed{\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-0.6 \pm 1.4\ (\mathrm{stat.}) \pm 0.4\ (\mathrm{syst.})) \times 10^{-25}\ \mathrm{y}^{-1}.} \tag{7.1}$$

7.5 Limit results

Since the result of the decay rate measurement in Eq. (7.1) is consistent with zero, an upper limit on the decay rate is set, following the technique discussed in Sect. 6.7. Figure 7.7 shows the profile negative log likelihood function without systematic errors (blue) and with systematic errors included (red). The log likelihood function is exponentiated to obtain the likelihood function, shown in Fig. 7.8. The Bayesian posterior p.d.f. for Γ is the product of the likelihood function and a prior p.d.f. which is taken to be zero in the unphysical ($\Gamma < 0$) region and flat in the physical ($\Gamma \geq 0$) region. The 90% C.L. upper limit, Γ^{limit}, is the value such that 90% of the area under the posterior p.d.f. in the physical region is contained between $\Gamma = 0$ and $\Gamma = \Gamma^{\mathrm{limit}}$. The result for the upper limit on the decay rate for $0\nu\beta\beta$ decay of ^{130}Te including systematic errors is

$$\boxed{\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) < 2.3 \times 10^{-25}\ \mathrm{y}^{-1}\ (90\%\ \mathrm{C.L.}).} \tag{7.2}$$

The corresponding limit on the partial half-life for $0\nu\beta\beta$ decay of ^{130}Te, calculated as $T_{1/2}^{\mathrm{limit}} = \ln 2/\Gamma^{\mathrm{limit}}$, is

$$\boxed{T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 3.0 \times 10^{24}\ \mathrm{y}\ (90\%\ \mathrm{C.L.}).} \tag{7.3}$$

Without systematic errors the limits would be

$$\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) < 2.2 \times 10^{-25}\ \mathrm{y}^{-1}\ (90\%\ \mathrm{C.L.}) \tag{7.4}$$

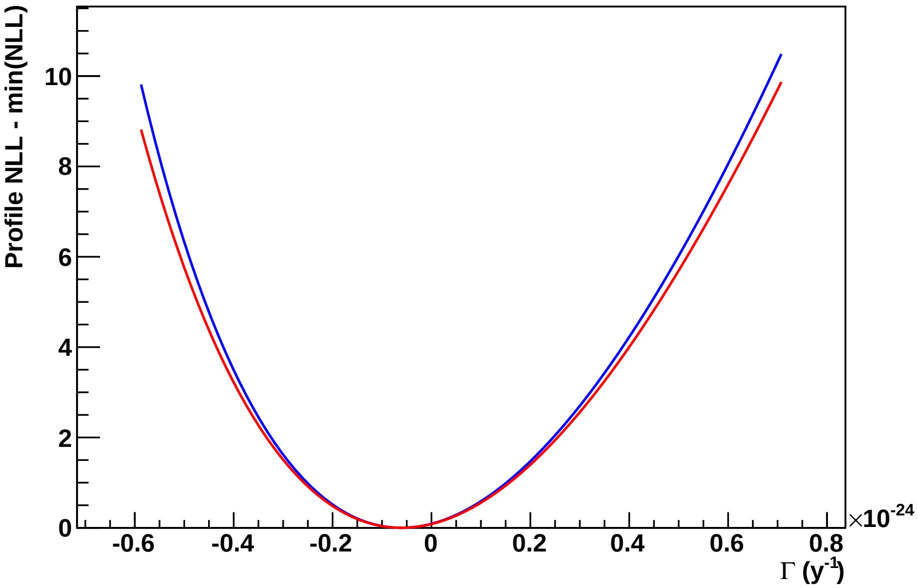

Figure 7.7: Profile negative log likelihood function without (blue) and with (red) systematic errors included. The blue curve is equal to $\chi^2_{\text{stat.}}(\Gamma)/2$, and the red curve is equal to $\chi^2_{\text{total}}(\Gamma)/2$.

and

$$T_{1/2}^{0\nu\beta\beta}(^{130}\text{Te}) > 3.1 \times 10^{24} \text{ y (90\% C.L.)}. \tag{7.5}$$

In comparison with the CUORICINO-only limit (Eq. (6.22)), the limit obtained from the combined analysis benefits from a slight downward fluctuation of the background. It was shown in Fig. 6.17 of Sect. 6.8 that the limit is very sensitive to the best fit value for Γ. Thus, for the limit, the additional 7% of exposure has a disproportionately large impact.

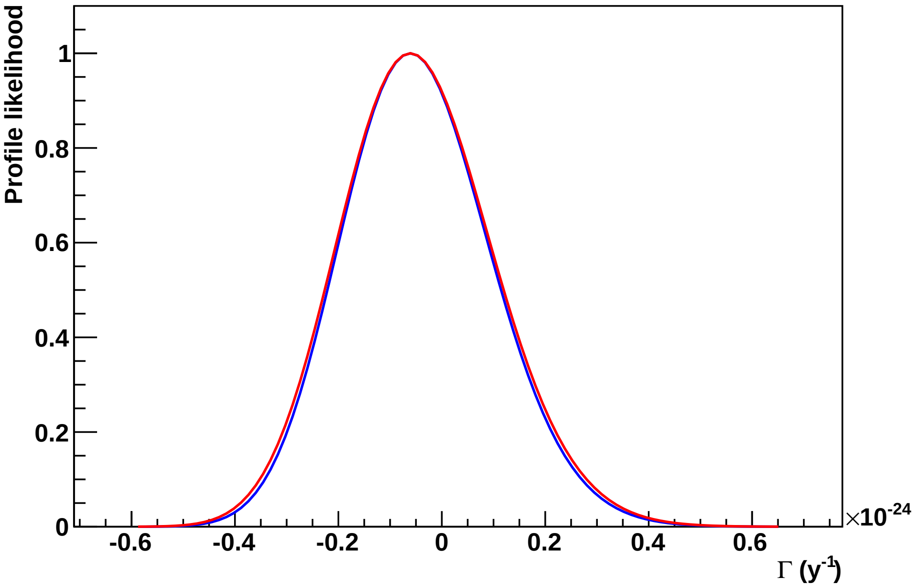

Figure 7.8: Profile likelihood function without (blue) and with (red) systematic errors included. The blue curve is equal to $e^{-\chi^2_{\text{stat.}}(\Gamma)/2}$, and the red curve is equal to $e^{-\chi^2_{\text{total}}(\Gamma)/2}$.

Chapter 8

Conclusion

This dissertation has presented a measurement (consistent with zero) of the decay rate for neutrinoless double beta decay of ^{130}Te made from combining the data from CUORICINO and the Three Towers Test. Since no signal was observed, an upper limit on the decay rate and a lower limit on the partial half-life for neutrinoless double beta decay of ^{130}Te have been set. The principal results are summarized in Table 8.1.

8.1 Comparison with previous CUORICINO results

The most recent published CUORICINO result on $0\nu\beta\beta$ decay was obtained from an analysis of 11.83 $\mathrm{kg} \cdot \mathrm{y}$ (^{130}Te) of data [72]. This analysis of all of CUORICINO Run II together with the Three Towers Test includes 68% more data. In order to compare the previous result directly with the one presented here, it is useful to convert the numbers in Ref. [72] for the best fit value for the number of signal counts and its error and for the exposure to a best fit value and error for Γ. After the conversion, the previous result was $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-2.9 \pm 1.8\ (\mathrm{stat.})) \times 10^{-25}\ \mathrm{y}^{-1}$. The result of this analysis, $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-0.6 \pm 1.4\ (\mathrm{stat.}) \pm 0.4\ (\mathrm{syst.})) \times 10^{-25}\ \mathrm{y}^{-1}$, has a statistical error that is reduced as expected from the increase in exposure. When comparing the results, the different values used for the signal efficiency have to be taken into account. The previous analysis used 86.3% and 84.5% as the signal efficiency for the big crystals and the small crystals, respectively, whereas this analysis used 81.7% and 80.0% (averaged over all CUORICINO channels and data sets). Systematic errors are explicitly included for the first time in this analysis. The calculation of systematic errors presented here verifies the previous expectation that they are small compared with the statistical error. Furthermore, the systematic errors have a

Table 8.1: Principal results of this work.

CUORICINO + Three Towers Test
$\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-0.6 \pm 1.4\ (\mathrm{stat.}) \pm 0.4\ (\mathrm{syst.})) \times 10^{-25}\ \mathrm{y}^{-1}$
$\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) < 2.3 \times 10^{-25}\ \mathrm{y}^{-1}$ (90% C.L.)
$T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 3.0 \times 10^{24}$ y (90% C.L.)

CUORICINO
$\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) = (-0.2 \pm 1.5\ (\mathrm{stat.}) \pm 0.3\ (\mathrm{syst.})) \times 10^{-25}\ \mathrm{y}^{-1}$
$\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te}) < 2.7 \times 10^{-25}\ \mathrm{y}^{-1}$ (90% C.L.)
$T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 2.6 \times 10^{24}$ y (90% C.L.)

statistical component and will therefore remain smaller than the statistical error for CUORE.

The previous result represents a 1.6σ downward fluctuation of the background in the signal region. In the new result, the best fit value is again negative but only slightly, representing a 0.4σ downward fluctuation. The change in best fit value comes from the additional data, an update of the Q-value from 2530.3 keV to 2527.5 keV, and the reshuffling of event energies within their errors in the new DIANA processing compared with the previous processing. The change in best fit value has a large impact on the results for the limit on $T_{1/2}^{0\nu\beta\beta}$. The limit obtained in the previous analysis was $T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 3.0 \times 10^{24}$ y (90% C.L.). The limit obtained in this work from the combination of the full CUORICINO data set and the Three Towers Test has the same value despite the increase in exposure and decrease in statistical error on Γ because the limit is strongly dependent on the best fit value for Γ, as demonstrated in Fig. 6.17.

8.2 Limit on the effective neutrino mass

An upper limit on the effective Majorana neutrino mass, $m_{\beta\beta}$, is obtained from the upper limit on $\Gamma^{0\nu\beta\beta} = \ln 2/T_{1/2}^{0\nu\beta\beta}$ and the relation (reproduced from Eq. (2.1))

$$\frac{1}{T_{1/2}^{0\nu\beta\beta}} = G_{0\nu}(Q,Z)|M^{0\nu}|^2\, m_{\beta\beta}^2. \tag{8.1}$$

Although making a measurement of $m_{\beta\beta}$ requires the assumption that the dominant mechanism driving $0\nu\beta\beta$ decay is the exchange of light Majorana neutrinos, setting a limit on $m_{\beta\beta}$ is possible regardless of the dominant mechanism. The conversion from $1/T_{1/2}^{0\nu\beta\beta}$ to $m_{\beta\beta}$ requires knowledge of the nuclear matrix element, $M^{0\nu}$. Since there is currently a significant spread in the values of $M^{0\nu}$ calculated with different models, a standard practice is to tabulate the limits on $m_{\beta\beta}$ obtained from a set of contemporary calculations of $M^{0\nu}$. Table 8.2 contains such a tabulation using recent calculations of the nuclear matrix element by four groups. The list includes calculations based on the quasiparticle random phase approximation (QRPA), renormalized QRPA (RQRPA), the interacting shell model (ISM), and the interacting boson model (IBM). For some of these models, different treatments of short range correlations are considered, including the coupled-cluster method (CCM), Miller-Spencer Jastrow, and the unitary correlation operator method (UCOM). For each calculation in the table, the values used for the axial vector coupling constant, g_A, and the nuclear radius parameter, r_0 ($R = r_0 A^{1/3}$), are listed because the nuclear matrix element and the phase space factor, $G_{0\nu}$, depend on both of these parameters. Since the effect of changing g_A on the $0\nu\beta\beta$ decay half-life must be ascertained by considering both $M^{0\nu}$ and $G_{0\nu}$, several authors of articles on nuclear matrix elements have found it convenient to use one value of $G_{0\nu}$ evaluated for $g_A = 1.25$ and to absorb the g_A^4 dependence of $G_{0\nu}$ in to a modified matrix element defined as

$$M'^{0\nu} = \left(\frac{g_A}{1.25}\right)^2 M^{0\nu}. \tag{8.2}$$

Defined in this way, it is possible to consider only $M'^{0\nu}$ when comparing different nuclear matrix element calculations. In Table 8.2, the matrix elements are expressed as $M'^{0\nu}$. In two cases, the value of $M^{0\nu}$ in the referenced article has been converted to $M'^{0\nu}$; this conversion is explicitly indicated in the table with a factor of $(1/1.25)^2$. For the other entries in Table 8.2, the original authors present their results in terms of $M'^{0\nu}$ or they use $g_A = 1.25$ in which case $M'^{0\nu} = M^{0\nu}$. The range of limits from the calculations in Table 8.2 is

$$m_{\beta\beta} < 0.25\text{–}0.68 \text{ eV (90\% C.L.)}. \tag{8.3}$$

8.3 Comparison with the claim of discovery

As discussed in Sect. 2.3.1, a subset of the Heidelberg-Moscow Collaboration has claimed to have observed $0\nu\beta\beta$ decay of ^{76}Ge. In this section, the statistical significance of the discrepancy between the non-observation of $0\nu\beta\beta$ decay in this work and the claimed observation is evaluated. In order to compare the results

Table 8.2: Upper limits on the effective Majorana neutrino mass obtained from $T_{1/2}^{0\nu\beta\beta}(^{130}\mathrm{Te}) > 3.0 \times 10^{24}$ y (90% C.L.) and recent calculations of the nuclear matrix element. The value of $G_{0\nu}$ used with the models with $r_0 = 1.1$ fm is from Table II of Ref. [91], and the value used with the models with $r_0 = 1.2$ fm is from Table 6 of Ref. [92].

Model	$M'^{0\nu} =$ $(g_A/1.25)^2 M^{0\nu}$	g_A	r_0 (fm)	$G_{0\nu}(g_A = 1.25)$ $(\mathrm{y}^{-1} \cdot \mathrm{eV}^{-2})$	$m_{\beta\beta}^{\mathrm{limit}}$ (eV) (90% C.L.)
(R)QRPA [93] CCM SRC	2.92–5.04	1.00 & 1.25	1.1	2.12×10^{-25}	0.25–0.43
QRPA [94] Jastrow	−2.993	1.25	1.2	1.59×10^{-25}	0.48
QRPA [94] Jastrow	−4.061 $\times(1/1.25)^2$	1.00	1.2	1.59×10^{-25}	0.56
QRPA [94] UCOM	−4.221	1.25	1.2	1.59×10^{-25}	0.34
QRPA [94] UCOM	−5.442 $\times(1/1.25)^2$	1.00	1.2	1.59×10^{-25}	0.42
ISM [95] Jastrow	2.12	1.25	1.2	1.59×10^{-25}	0.68
ISM [95] UCOM	2.65	1.25	1.2	1.59×10^{-25}	0.55
IBM [96]	4.059	1.25	1.2	1.59×10^{-25}	0.36

obtained with two different isotopes, the results for $\Gamma^{0\nu\beta\beta}(^{76}\mathrm{Ge})$ from Klapdor-Kleingrothaus *et al.* and $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te})$ from this work are converted into measurements of $m^2_{\beta\beta}$ using Eq. (8.1) because $m^2_{\beta\beta}$ has Gaussian errors since it is proportional to the number of observed signal events. The conversion is performed separately for each nuclear matrix element calculation because the calculated nuclear matrix elements for different isotopes are highly correlated [97].

The number of observed signal events in the Klapdor-Kleingrothaus *et al.* analysis of the Heidelberg-Moscow data is 11 ± 1.8, which leads to a partial half-life of $T^{0\nu\beta\beta}_{1/2}(^{76}\mathrm{Ge}) = 2.23^{+0.44}_{-0.31} \times 10^{25}$ y [35]. The central value for $T^{0\nu\beta\beta}_{1/2}(^{76}\mathrm{Ge})$ is equivalent to a decay rate of $\Gamma^{0\nu\beta\beta}(^{76}\mathrm{Ge}) = 3.11 \times 10^{-26}\ \mathrm{y}^{-1}$. From the error on the number of counts, the error on $\Gamma^{0\nu\beta\beta}(^{76}\mathrm{Ge})$ is calculated to be $\sigma_{\Gamma^{0\nu\beta\beta}}(^{76}\mathrm{Ge}) = (1.8/11) \times 3.11 \times 10^{-26}\ \mathrm{y}^{-1} = 0.51 \times 10^{-26}\ \mathrm{y}^{-1}$. Using this $0\nu\beta\beta$ decay rate value and error for ^{76}Ge and the value obtained in this work from the combination of CUORICINO and the Three Towers Test for the $0\nu\beta\beta$ decay rate of ^{130}Te, the values for $m^2_{\beta\beta}$ are calculated from a selection of recent nuclear matrix element calculations, and the results are listed in Table 8.3. The comparison is presented graphically in Fig. 8.1. The last column in Table 8.3 expresses the difference between the values of $m^2_{\beta\beta}$ obtained from the two measurements in number of sigma; it is calculated as

$$\frac{\Delta m^2_{\beta\beta}}{\sigma_{\Delta m^2_{\beta\beta}}} = \frac{m^2_{\beta\beta}(^{76}\mathrm{Ge}) - m^2_{\beta\beta}(^{130}\mathrm{Te})}{\sqrt{\sigma^2_{m^2_{\beta\beta}(^{76}\mathrm{Ge})} + \sigma^2_{m^2_{\beta\beta}(^{130}\mathrm{Te})}}} \tag{8.4}$$

and ranges from 1.2 to 1.7 for the set of nuclear matrix element calculations in the table. Taking the minimum, 1.2σ, as a measure of the discrepancy between the two results, the probability of a discrepancy this large or larger happening by chance is

$$\mathrm{Prob}\left(\frac{|\Delta m^2_{\beta\beta}|}{\sigma_{\Delta m^2_{\beta\beta}}} > 1.2\right) = 1 - \int_{-1.2}^{+1.2} \frac{1}{\sqrt{2\pi}} e^{-x^2/2}\, dx = 23\%. \tag{8.5}$$

Thus, the inconsistency between the result of this work and the result of Klapdor-Kleingrothaus *et al.* is not statistically significant enough to exclude the possibility that the signal observed in the Heidelberg-Moscow experiment is due to neutrinoless double beta decay of ^{76}Ge driven by the exchange of light Majorana neutrinos.

Using the same approach, we can calculate how long it will take CUORE to reach a result that is inconsistent with the claimed observation at a given level of statistical significance, assuming that CUORE does not observe a signal and assuming that the nuclear matrix element models are reliable at least for calculating the ratio of the matrix elements for ^{76}Ge and ^{130}Te. In the absence of a signal, the expected statistical error on $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te})$ for CUORE is determined by the Poisson fluctuations of the background. Considering a very simple counting analysis within an energy interval of width ΔE around the Q-value, the error on Γ from Poisson statistics is

$$\sigma_\Gamma = \frac{\sqrt{b(\Delta E)m}}{N_0 \varepsilon \sqrt{t}}, \tag{8.6}$$

where b is the background level in counts/(keV $\cdot$ kg $\cdot$ y), m is the total mass of $\mathrm{TeO_2}$, N_0 is the number of ^{130}Te nuclei, ε is the signal efficiency, and t is the live time. If CUORE achieves its background goal of $b = 0.01$ counts/(keV$\cdot$kg$\cdot$y) and energy resolution goal of $\Delta E = 5$ keV, plugging in the values for CUORE results in an expected uncertainty of

$$\sigma_\Gamma = \frac{7.9 \times 10^{-27}\ \mathrm{y}^{-1/2}}{\sqrt{t}}, \tag{8.7}$$

assuming $\varepsilon = 0.82$, similar to the CUORICINO big crystals. With this precision, after just two months of taking data, the CUORE measurement of $\Gamma^{0\nu\beta\beta}(^{130}\mathrm{Te})$ will be inconsistent with the claim of observation of $0\nu\beta\beta$ decay at greater than 3σ significance for all of the nuclear matrix element calculations considered here, even if CUORE observes a 1σ upward fluctuation of the background.

Table 8.3: Comparison between the result of this work and the claim of discovery.

	Model	g_A	Nuclide	$M'^{0\nu}$	$m^2_{\beta\beta}$ (eV^2)	$\frac{\Delta m^2_{\beta\beta}}{\sigma_{\Delta m^2_{\beta\beta}}}$
(a)	RQRPA [93] CCM CD-Bonn	1.25	^{76}Ge	5.440	0.050 ± 0.008	1.4
			^{130}Te	4.400	-0.021 ± 0.049	
(b)	RQRPA [93] CCM CD-Bonn	1.00	^{76}Ge	4.620	0.069 ± 0.011	1.2
			^{130}Te	3.380	-0.036 ± 0.083	
(c)	QRPA [93] CCM CD-Bonn	1.25	^{76}Ge	6.320	0.037 ± 0.006	1.4
			^{130}Te	4.920	-0.017 ± 0.039	
(d)	QRPA [93] CCM CD-Bonn	1.00	^{76}Ge	5.160	0.055 ± 0.009	1.2
			^{130}Te	3.770	-0.029 ± 0.067	
(e)	RQRPA [93] CCM Argonne	1.25	^{76}Ge	4.970	0.060 ± 0.010	1.4
			^{130}Te	3.910	-0.027 ± 0.062	
(f)	RQRPA [93] CCM Argonne	1.00	^{76}Ge	4.210	0.083 ± 0.013	1.2
			^{130}Te	3.020	-0.045 ± 0.104	
(g)	QRPA [93] CCM Argonne	1.25	^{76}Ge	5.810	0.044 ± 0.007	1.3
			^{130}Te	4.370	-0.021 ± 0.050	
(h)	QRPA [93] CCM Argonne	1.00	^{76}Ge	4.770	0.065 ± 0.010	1.2
			^{130}Te	3.380	-0.036 ± 0.083	
(i)	QRPA [94] Jastrow	1.25	^{76}Ge	-4.029	0.114 ± 0.018	1.2
			^{130}Te	-2.993	-0.061 ± 0.142	
(j)	QRPA [94] Jastrow	1.00	^{76}Ge	-3.229	0.178 ± 0.029	1.3
			^{130}Te	-2.582	-0.082 ± 0.191	
(k)	QRPA [94] UCOM	1.25	^{76}Ge	-5.355	0.065 ± 0.010	1.3
			^{130}Te	-4.221	-0.031 ± 0.072	
(l)	QRPA [94] UCOM	1.00	^{76}Ge	-4.168	0.107 ± 0.017	1.4
			^{130}Te	-3.461	-0.046 ± 0.106	
(m)	ISM [95] Jastrow	1.25	^{76}Ge	2.300	0.351 ± 0.057	1.6
			^{130}Te	2.120	-0.121 ± 0.283	
(n)	ISM [95] UCOM	1.25	^{76}Ge	2.810	0.235 ± 0.038	1.7
			^{130}Te	2.650	-0.078 ± 0.181	
(o)	IBM [96]	1.25	^{76}Ge	5.465	0.062 ± 0.010	1.2
			^{130}Te	4.059	-0.033 ± 0.077	

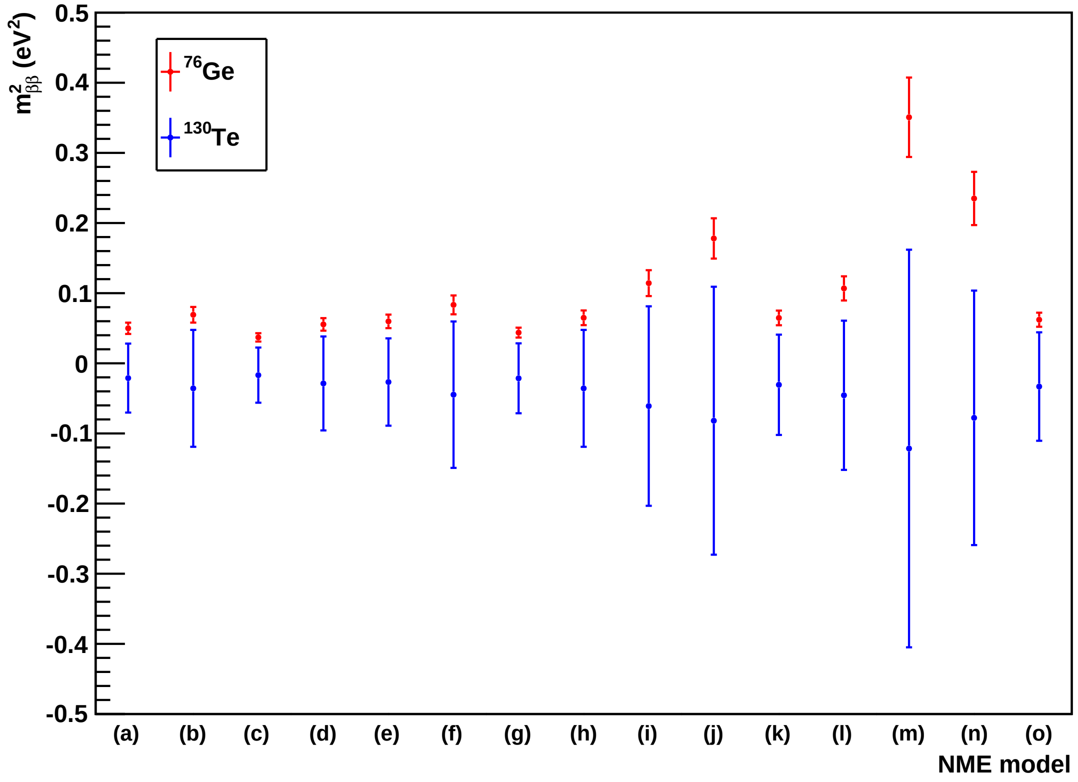

Figure 8.1: Comparison between the result of this work and the claim of discovery based on the $m^2_{\beta\beta}$ values and one-sigma errors in Table 8.3. The labels for the nuclear matrix element (NME) models correspond to the first column of Table 8.3.

Bibliography

[1] E. Majorana, Nuovo Cimento **14**, 171 (1937).

[2] S. M. Bilenky, (2006), arXiv:hep-ph/0605172.

[3] B. Kayser, Phys. Rev. D **26**, 1662 (1982).

[4] SNO Collaboration, Q. R. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 (2002), arXiv:nucl-ex/0204008.

[5] SNO Collaboration, S. N. Ahmed *et al.*, Phys. Rev. Lett. **92**, 181301 (2004), arXiv:nucl-ex/0309004.

[6] Super-Kamiokande Collaboration, Y. Ashie *et al.*, Phys. Rev. **D71**, 112005 (2005), arXiv:hep-ex/0501064.

[7] KamLAND Collaboration, S. Abe *et al.*, Phys. Rev. Lett. **100**, 221803 (2008), arXiv:0801.4589.

[8] K2K Collaboration, M. H. Ahn *et al.*, Phys. Rev. **D74**, 072003 (2006), arXiv:hep-ex/0606032.

[9] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **101**, 131802 (2008), arXiv:0806.2237.

[10] J. Schechter and J. W. F. Valle, Phys. Rev. **D25**, 2951 (1982).

[11] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Westview, 1995).

[12] J. Schechter and J. W. F. Valle, Phys. Rev. D **22**, 2227 (1980).

[13] S. M. Bilenky and S. T. Petcov, Rev. Mod. Phys. **59**, 671 (1987).

[14] B. T. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).

[15] GALLEX Collaboration, W. Hampel *et al.*, Phys. Lett. **B447**, 127 (1999).

[16] SAGE Collaboration, J. N. Abdurashitov *et al.*, Phys. Rev. **C60**, 055801 (1999), arXiv:astro-ph/9907113.

[17] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1158 (1998), arXiv:hep-ex/9805021.

[18] A. Strumia and F. Vissani, (2006), arXiv:hep-ph/0606054.

[19] M. Maltoni, T. Schwetz, M. A. Tortola, and J. W. F. Valle, New J. Phys. **6**, 122 (2004), arXiv:hep-ph/0405172v6.

[20] CHOOZ Collaboration, M. Apollonio *et al.*, Eur. Phys. J. **C27**, 331 (2003), arXiv:hep-ex/0301017.

[21] C. Weinheimer *et al.*, Phys. Lett. **B460**, 219 (1999).

[22] V. M. Lobashev *et al.*, Nucl. Phys. Proc. Suppl. **91**, 280 (2001).

[23] WMAP Collaboration, E. Komatsu *et al.*, Astrophys. J. Suppl. **180**, 330 (2009), arXiv:0803.0547.

[24] M. Goeppert-Mayer, Phys. Rev. **48**, 512 (1935).

[25] A. S. Barabash, Phys. Rev. **C81**, 035501 (2010), arXiv:1003.1005.

[26] F. T. Avignone, III, S. R. Elliott, and J. Engel, Rev. Mod. Phys. **80**, 481 (2008), arXiv:0708.1033.

[27] L. P. Ekström and R. B. Firestone, WWW Table of Radioactive Isotopes, database version 2/28/99 from URL http://ie.lbl.gov/toi.

[28] G. Racah, Nuovo Cimento **14**, 322 (1937).

[29] W. H. Furry, Phys. Rev. **56**, 1184 (1939).

[30] V. Rodin, (2009), arXiv:0910.5866.

[31] E. Caurier, G. Martinez-Pinedo, F. Nowacki, A. Poves, and A. P. Zuker, Rev. Mod. Phys. **77**, 427 (2005), arXiv:nucl-th/0402046.

[32] V. A. Rodin, A. Faessler, F. Simkovic, and P. Vogel, Phys. Rev. **C68**, 044302 (2003), arXiv:nucl-th/0305005.

[33] V. A. Rodin, A. Faessler, F. Simkovic, and P. Vogel, Nucl. Phys. **A766**, 107 (2006), arXiv:0706.4304.

[34] A. Faessler and F. Simkovic, J. Phys. **G24**, 2139 (1998), arXiv:hep-ph/9901215.

[35] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina, Mod. Phys. Lett. **A21**, 1547 (2006).

[36] L. Baudis *et al.*, Phys. Rev. Lett. **83**, 41 (1999), arXiv:hep-ex/9902014.

[37] IGEX Collaboration, C. E. Aalseth *et al.*, Phys. Rev. **C59**, 2108 (1999).

[38] IGEX Collaboration, C. E. Aalseth *et al.*, Phys. Rev. **D65**, 092007 (2002), arXiv:hep-ex/0202026.

[39] I. Ogawa *et al.*, Nucl. Phys. **A730**, 215 (2004).

[40] NEMO-3 Collaboration, R. Arnold *et al.*, Phys. Rev. Lett. **95**, 182302 (2005), arXiv:hep-ex/0507083.

[41] NEMO-3 Collaboration, J. Argyriades *et al.*, (2009), arXiv:0906.2694.

[42] F. A. Danevich *et al.*, Phys. Rev. **C68**, 035501 (2003).

[43] R. Bernabei *et al.*, Phys. Lett. **B546**, 23 (2002).

[44] NEMO-3 Collaboration, J. Argyriades *et al.*, Phys. Rev. **C80**, 032501 (2009), arXiv:0810.0248.

[45] R. Arnold *et al.*, Nucl. Instrum. Meth. **A536**, 79 (2005), arXiv:physics/0402115.

[46] H. V. Klapdor-Kleingrothaus, A. Dietz, H. L. Harney, and I. V. Krivosheina, Mod. Phys. Lett. **A16**, 2409 (2001), arXiv:hep-ph/0201231.

[47] C. E. Aalseth *et al.*, Mod. Phys. Lett. **A17**, 1475 (2002), arXiv:hep-ex/0202018.

[48] F. Feruglio, A. Strumia, and F. Vissani, Nucl. Phys. **B637**, 345 (2002), arXiv:hep-ph/0201291.

[49] Y. G. Zdesenko, F. A. Danevich, and V. I. Tretyak, Phys. Lett. **B546**, 206 (2002).

[50] H. V. Klapdor-Kleingrothaus, I. V. Krivosheina, A. Dietz, and O. Chkvorets, Phys. Lett. **B586**, 198 (2004), arXiv:hep-ph/0404088.

[51] H. V. Klapdor-Kleingrothaus, A. Dietz, I. V. Krivosheina, and O. Chkvorets, Nucl. Instrum. Meth. **A522**, 371 (2004), arXiv:hep-ph/0403018.

[52] GERDA Collaboration, J. Janicsko-Csathy, Nucl. Phys. Proc. Suppl. **188**, 68 (2009).

[53] GERDA Collaboration, A. A. Smolnikov, (2008), arXiv:0812.4194.

[54] MAJORANA Collaboration, C. E. Aalseth *et al.*, J. Phys. Conf. Ser. **203**, 012057 (2010), arXiv:0910.4598.

[55] MAJORANA Collaboration, C. E. Aalseth *et al.*, AIP Conf. Proc. **1182**, 88 (2009), arXiv:0907.1581.

[56] M. Danilov *et al.*, Phys. Lett. **B480**, 12 (2000), arXiv:hep-ex/0002003.

[57] SuperNEMO Collaboration, E. Chauveau, AIP Conf. Proc. **1180**, 26 (2009).

[58] SNO+ Collaboration, K. Zuber, AIP Conf. Proc. **1180**, 150 (2009).

[59] KamLAND Collaboration, A. Terashima *et al.*, J. Phys. Conf. Ser. **120**, 052029 (2008).

[60] CANDLES Collaboration, Y. Hirano *et al.*, J. Phys. Conf. Ser. **120**, 052053 (2008).

[61] C. Kittel, *Introduction to Solid State Physics*, 7th ed. (Wiley, 1996).

[62] M. Barucci *et al.*, J. Low Temp. Phys. **123**, 303 (2001).

[63] A. Alessandrello *et al.*, Phys. Lett. **B247**, 442 (1990).

[64] D. McCammon, Appl. Phys. **99**, 35 (2005), arXiv:physics/0503086.

[65] E. E. Haller, J. Appl. Phys. **77**, 2857 (1995).

[66] C. Arnaboldi, G. Pessina, and E. Previtali, IEEE Trans. Nucl. Sci. **50**, 979 (2003).

[67] A. Alessandrello *et al.*, Nucl. Instrum. Meth. **A412**, 454 (1998).

[68] E. Fiorini and T. O. Niinikoski, Nucl. Instr. Meth. **A224**, 83 (1984).

[69] A. Alessandrello *et al.*, Phys. Lett. **B335**, 519 (1994).

[70] A. Alessandrello *et al.*, Phys. Lett. **B433**, 156 (1998).

[71] A. Alessandrello *et al.*, Phys. Lett. **B486**, 13 (2000).

[72] CUORICINO Collaboration, C. Arnaboldi *et al.*, Phys. Rev. **C78**, 035502 (2008), arXiv:0802.3439.

[73] C. Arnaboldi *et al.*, Nucl. Instrum. Meth. **A520**, 578 (2004).

[74] C. Arnaboldi *et al.*, IEEE Trans. Nucl. Sci. **52**, 1630 (2005).

[75] S. Di Domizio, *Search for double beta decay to excited states with CUORICINO and data acquisition system for CUORE*, PhD thesis, Università degli Studi di Genova, 2009.

[76] R. Ardito *et al.*, (2005), arXiv:hep-ex/0501010.

[77] F. Bellini *et al.*, Astropart. Phys. **33**, 169 (2010), arXiv:0912.0452.

[78] E. Andreotti *et al.*, Submitted to Astropart. Phys. (2010), arXiv:0912.3779.

[79] C. Arnaboldi *et al.*, Submitted to J. Crystal Growth (2010).

[80] PostgreSQL Database Management System, http://www.postgresql.org.

[81] M. Vignati, *Model of the Response Function of CUORE Bolometers*, PhD thesis, Sapienza – Università di Roma, 2009.

[82] R. Brun and F. Rademakers, Nucl. Instrum. Meth. **A389**, 81 (1997).

[83] ROOT | A Data Analysis Framework, http://root.cern.ch.

[84] JpGraph, http://www.aditus.nu/jpgraph/.

[85] Crystal Ball function – Wikipedia, the free encyclopedia, http://en.wikipedia.org/wiki/Crystal_Ball_function.

[86] M. Redshaw, B. J. Mount, E. G. Myers, and F. T. Avignone, III, Phys. Rev. Lett. **102**, 212502 (2009), arXiv:0902.2139.

[87] N. D. Scielzo *et al.*, Phys. Rev. **C80**, 025501 (2009), arXiv:0902.2376.

[88] G. Cowan, *Statistical Data Analysis* (Oxford University Press, 1998).

[89] W. Verkerke and D. Kirkby, (2003), arXiv:physics/0306116.

[90] F. James, *Statistical Methods in Experimental Physics*, 2nd ed. (World Scientific, 2006).

[91] F. Simkovic, G. Pantis, J. D. Vergados, and A. Faessler, Phys. Rev. **C60**, 055502 (1999), arXiv:hep-ph/9905509.

[92] J. Suhonen and O. Civitarese, Phys. Rept. **300**, 123 (1998).

[93] F. Šimkovic, A. Faessler, H. Muther, V. Rodin, and M. Stauf, Phys. Rev. **C79**, 055501 (2009), arXiv:0902.0331.

[94] O. Civitarese and J. Suhonen, J. Phys. Conf. Ser. **173**, 012012 (2009).

[95] J. Menéndez, A. Poves, E. Caurier, and F. Nowacki, Nucl. Phys. **A818**, 139 (2009), arXiv:0801.3760.

[96] J. Barea and F. Iachello, Phys. Rev. **C79**, 044301 (2009).

[97] A. Faessler *et al.*, Phys. Rev. **D79**, 053001 (2009), arXiv:0810.5733.

[98] M. E. Wieser and M. Berglund, Pure Appl. Chem. **81**, 2131 (2009).

[99] G. Audi, A. H. Wapstra, and C. Thibault, Nucl. Phys. **A729**, 337 (2002).

Appendix A

Energy spectra from CUORICINO and the Three Towers Test

In this appendix, the full energy spectra from CUORICINO and the Three Towers Test are displayed. Figures A.1, A.2, and A.3 contain the CUORICINO spectrum. Figure A.4 contains the Three Towers Test spectrum below 2000 keV obtained while the ^{40}K sources were in place. Figure A.5 contains the Three Towers Test spectrum below 2000 keV without the ^{40}K sources. Figures A.6 and A.7 contain the Three Towers Test spectrum above 2000 keV, combining the data collected with and without the ^{40}K sources. For CUORICINO, the third-order polynomial calibration function is used to produce the spectrum below 3000 keV, and the second-order log polynomial calibration function is used to produce the spectrum above 3000 keV. For the Three Towers Test, the second-order polynomial calibration function has been used for the entire spectrum. The cuts used to produce the plots in this appendix are the same as those used for the $0\nu\beta\beta$ decay analysis, except that no anti-coincidence cut is applied, and for the TTT no pulse shape cut is applied.

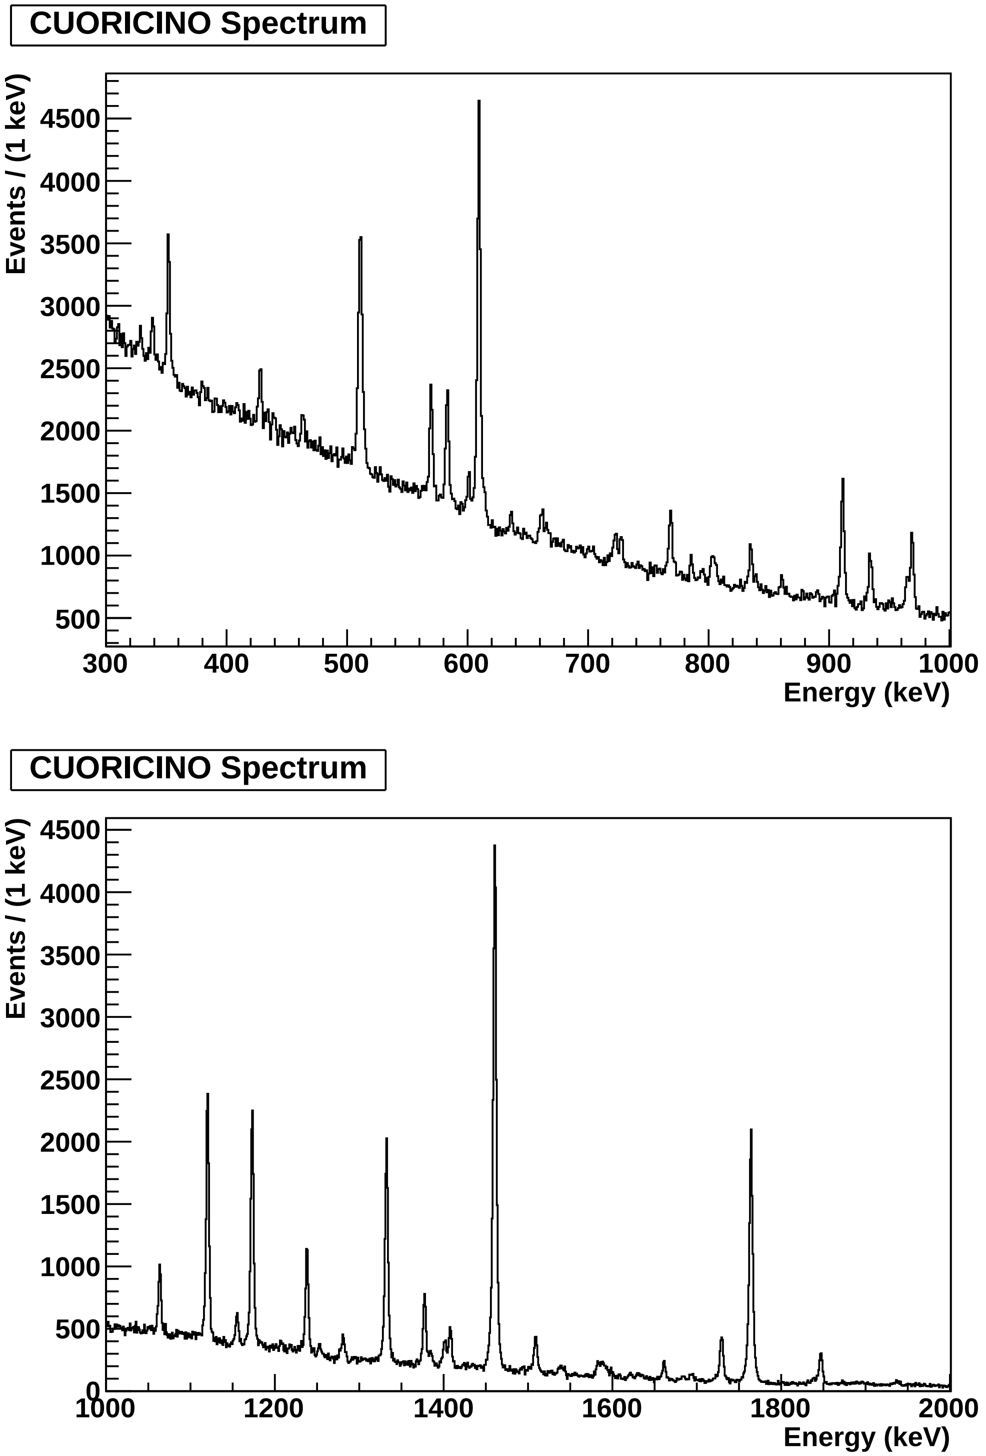

Figure A.1: CUORICINO spectrum calibrated with a third-order polynomial function.

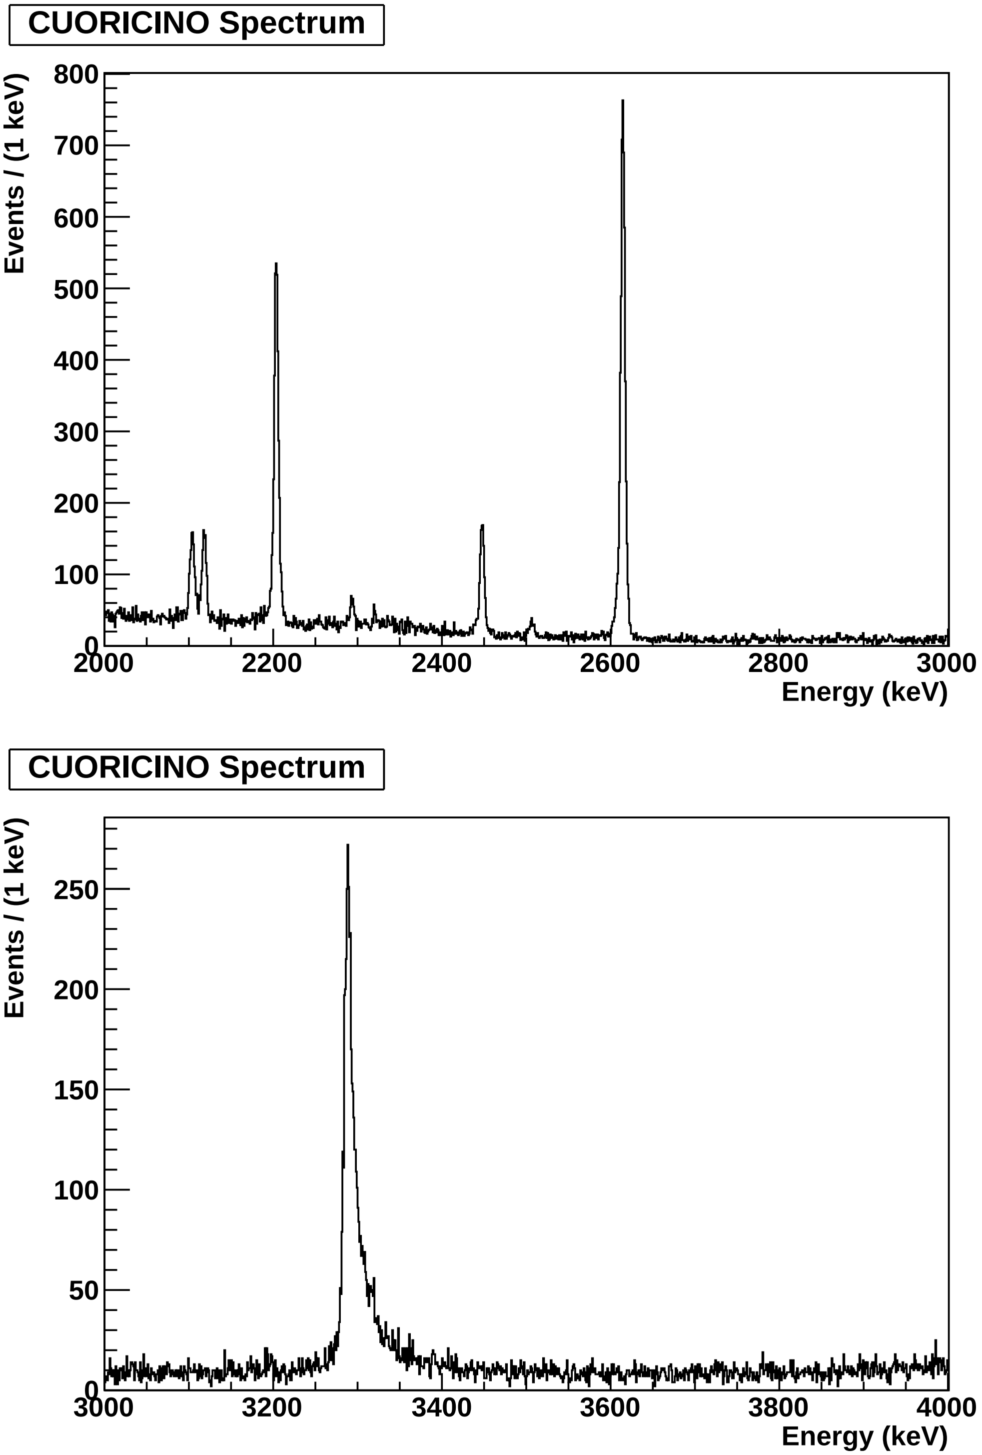

Figure A.2: CUORICINO spectrum calibrated with a third-order polynomial function below 3000 keV and calibrated with a second-order log polynomial function above 3000 keV.

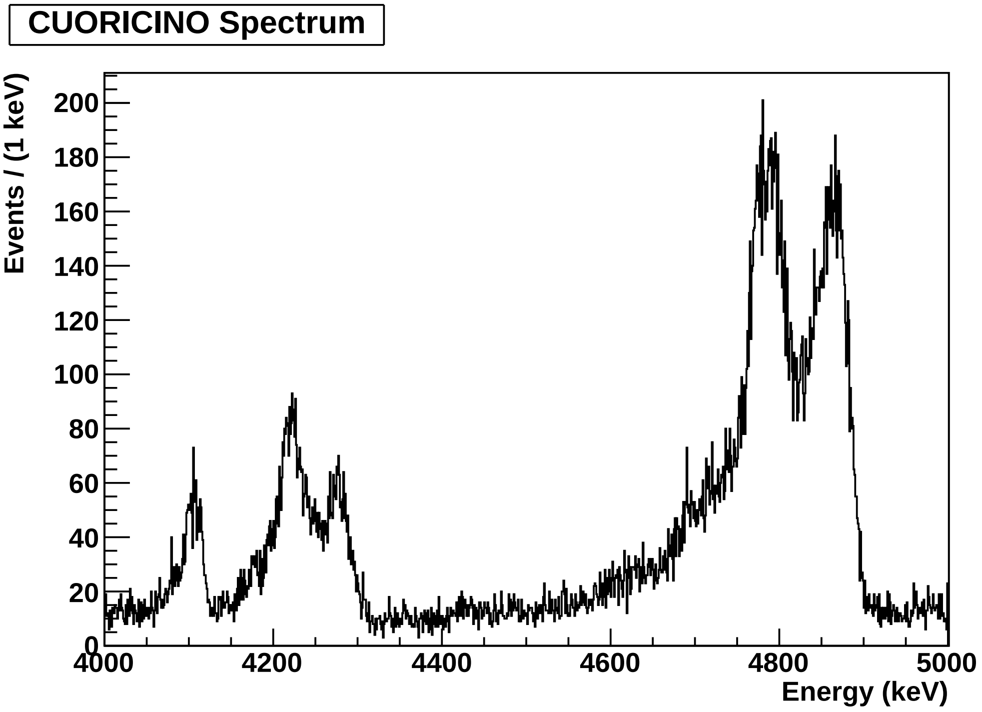

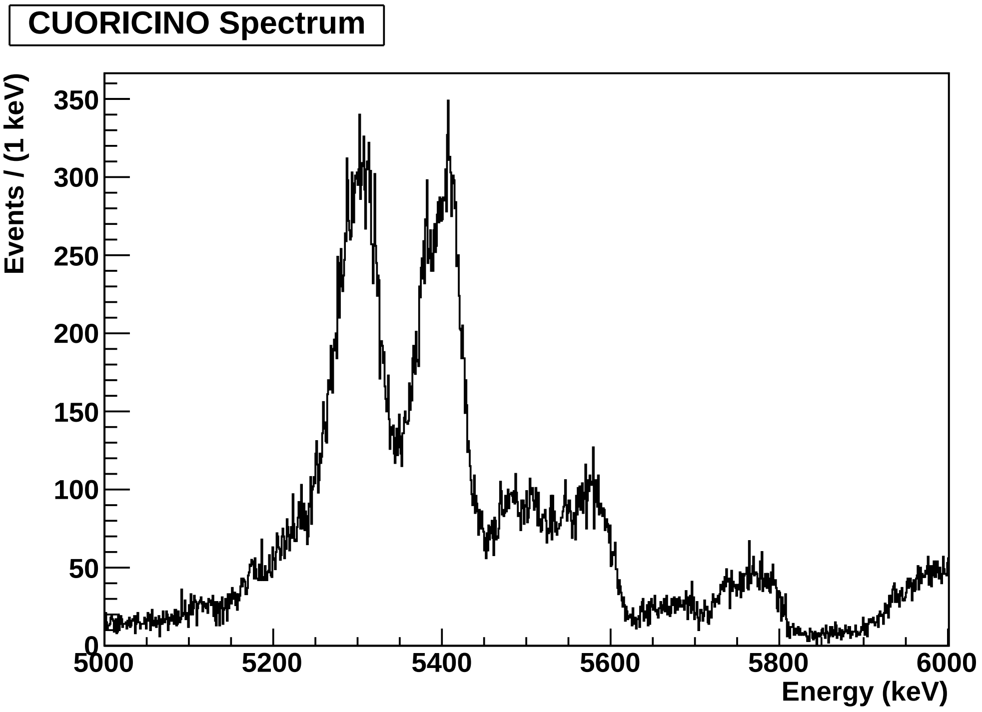

Figure A.3: CUORICINO spectrum calibrated with a second-order log polynomial function.

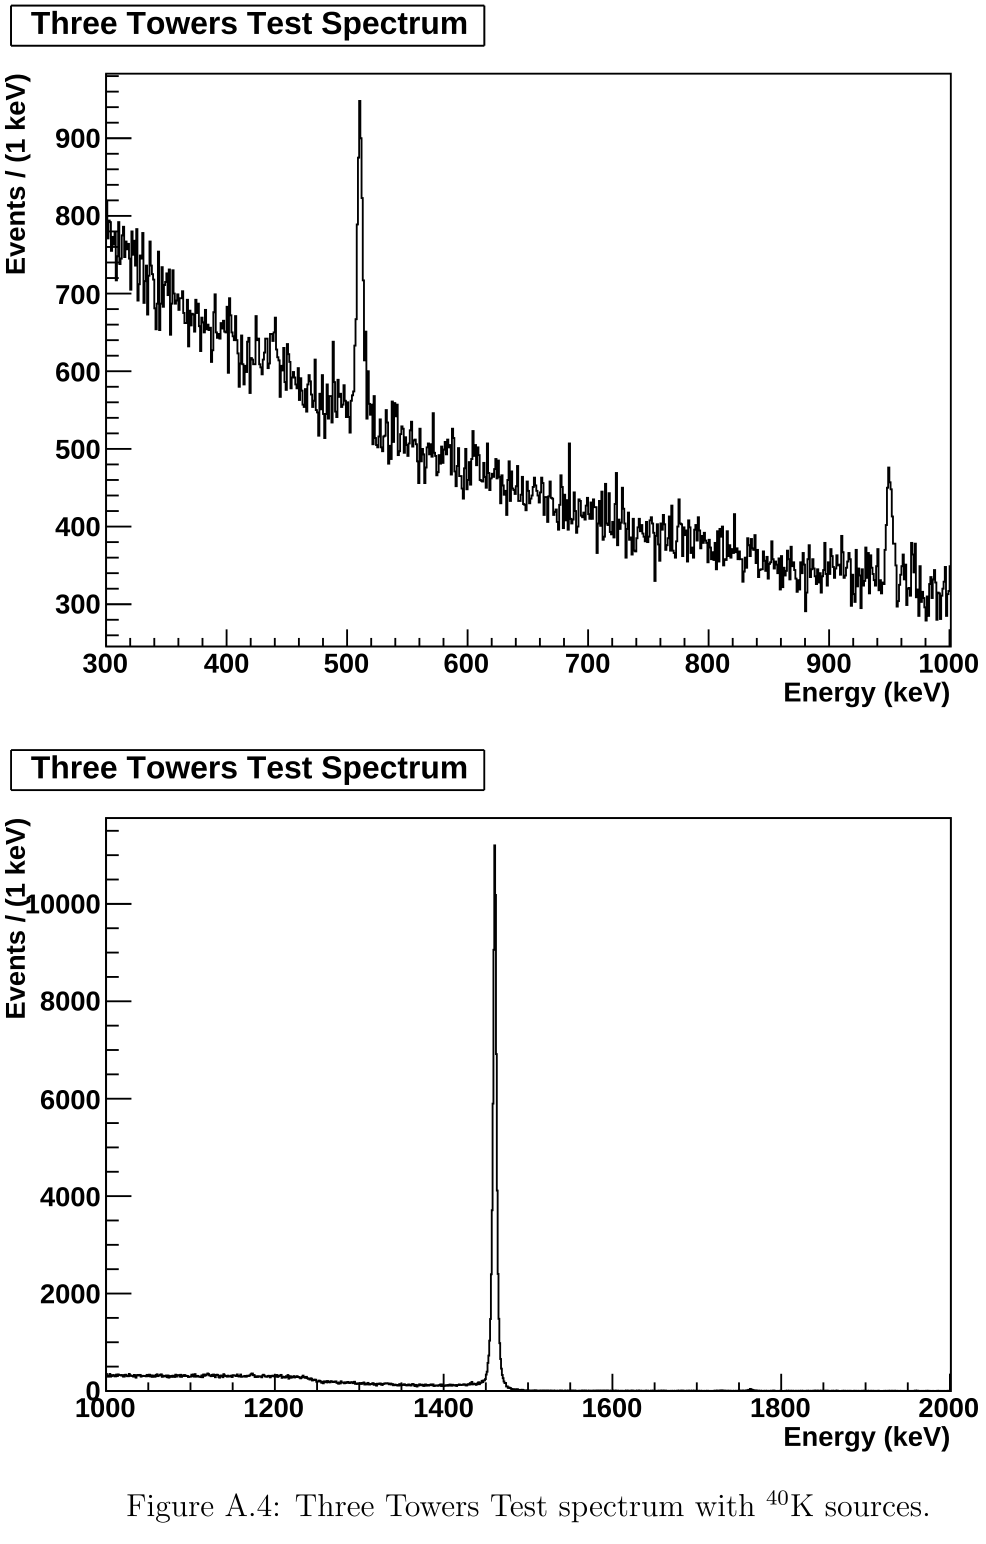

Figure A.4: Three Towers Test spectrum with ^{40}K sources.

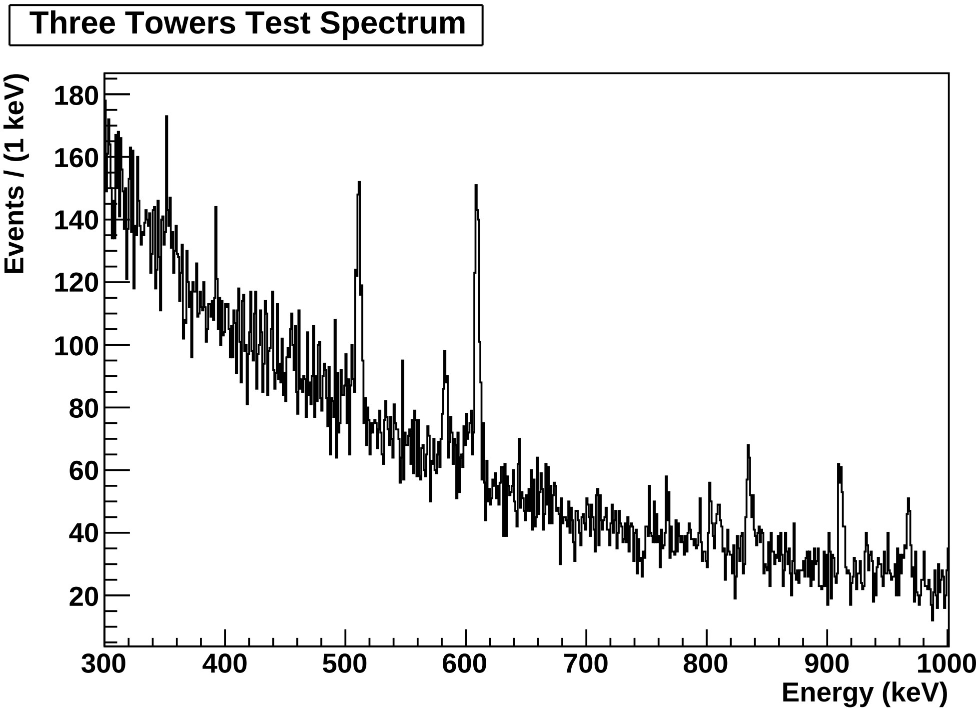

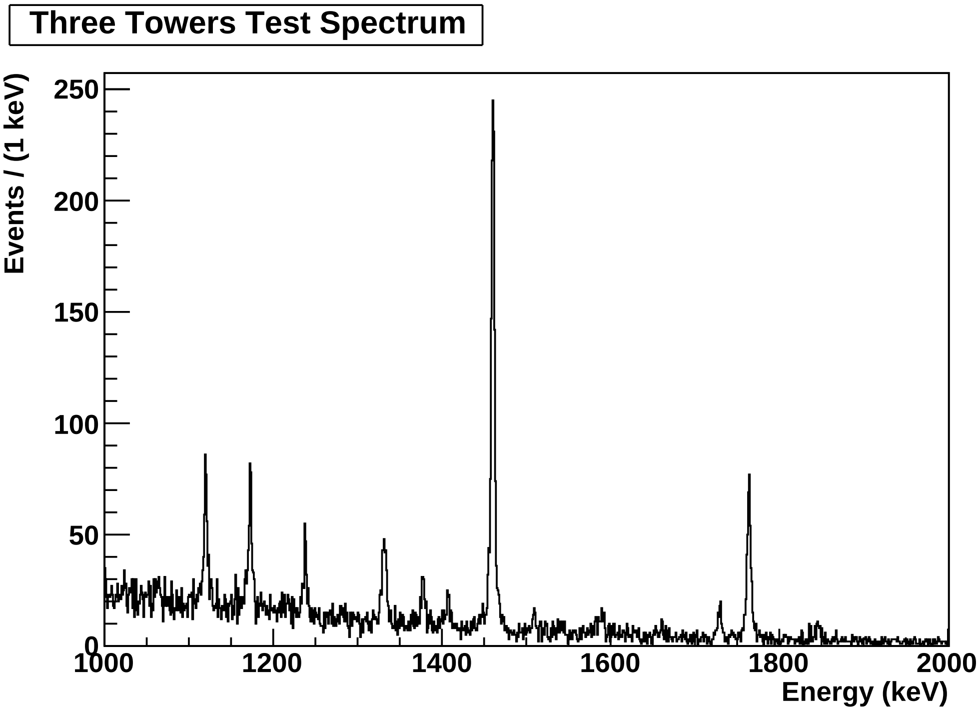

Figure A.5: Three Towers Test spectrum without ^{40}K sources.

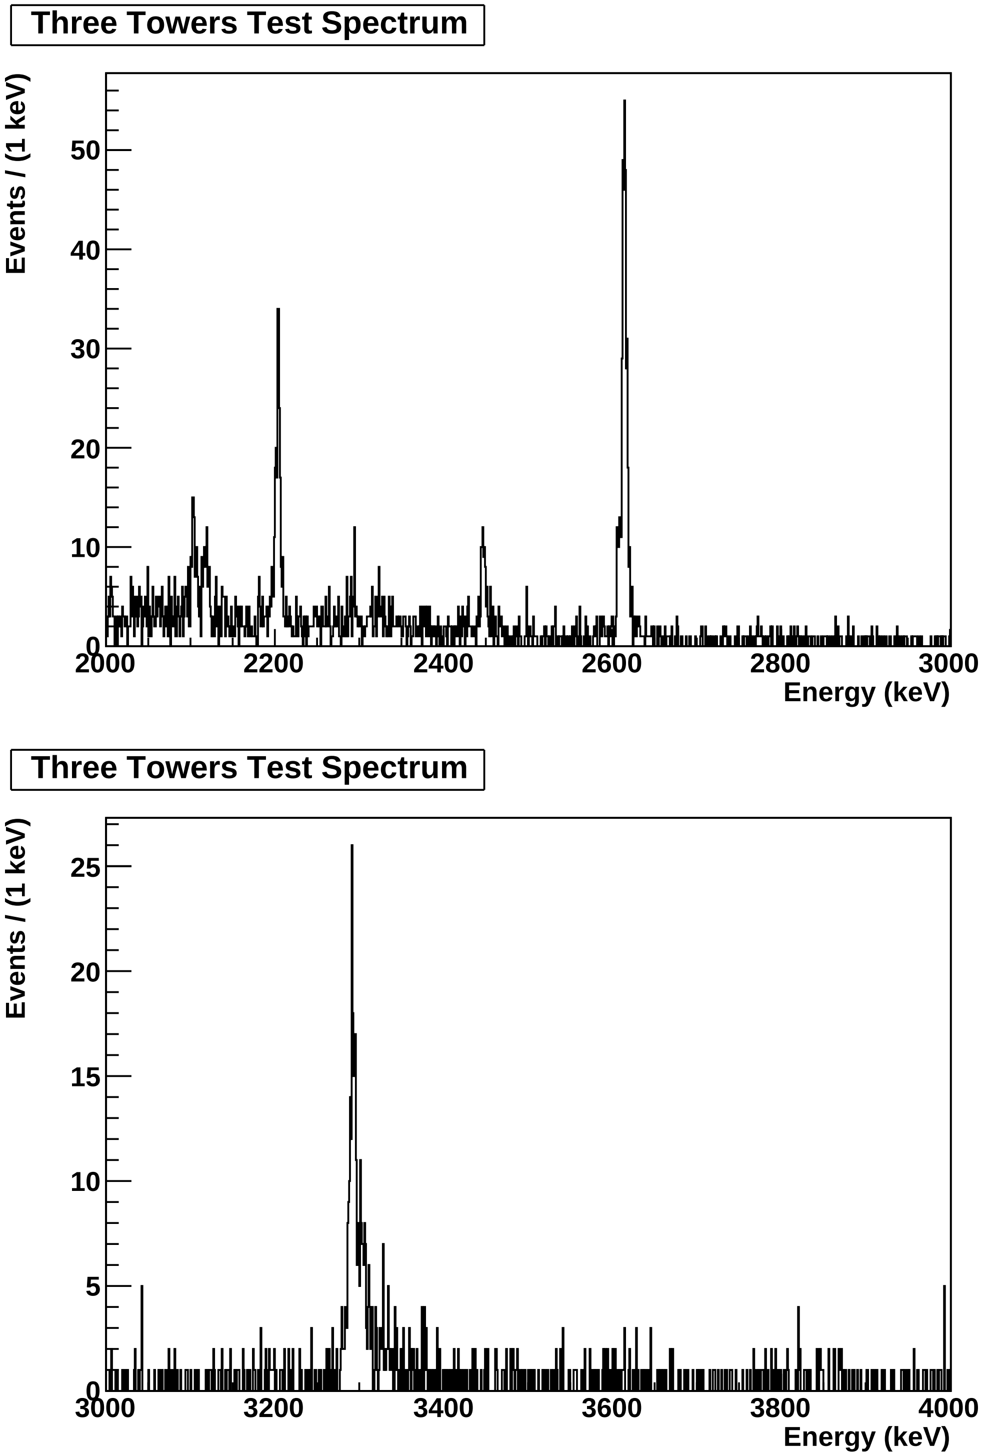

Figure A.6: Three Towers Test spectrum.

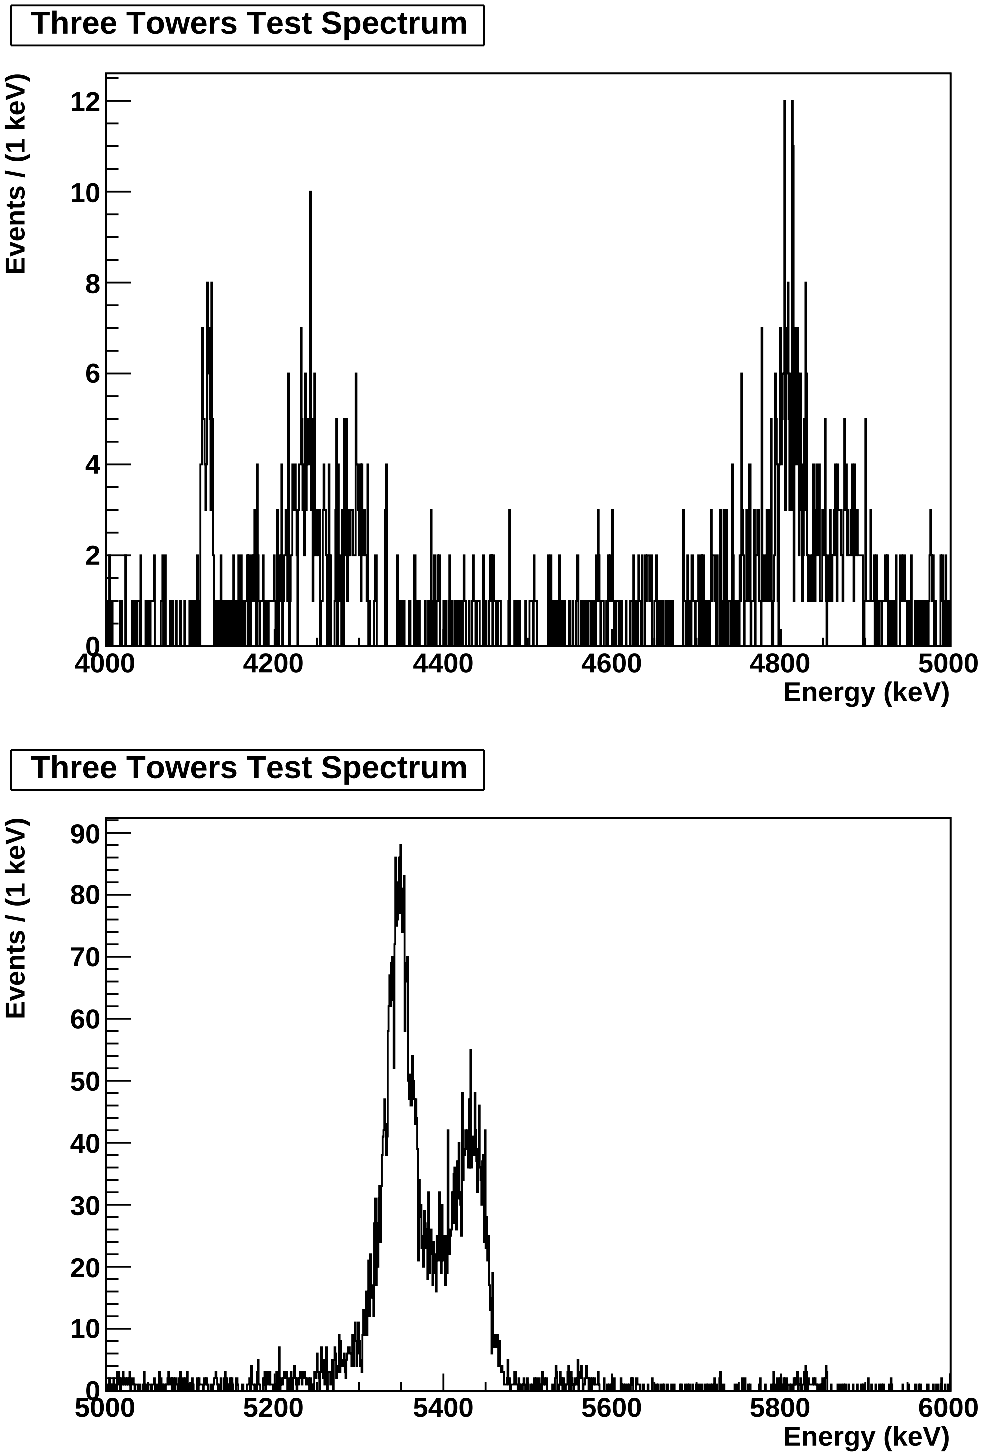

Figure A.7: Three Towers Test spectrum.

Appendix B

Tables of physical parameters, resolutions, and efficiencies

Values of physical parameters used in the analysis are listed in Table B.1. In Tables B.2, B.3, B.4, and B.5, the full width at half maximum (FWHM) energy resolutions used in the analysis are listed. The resolutions are obtained from an unbinned maximum likelihood fit with a Gaussian signal shape and a linear background to the 2615 keV calibration peak for each channel and data set. Examples of the fits are shown in Fig. 6.7. In Tables B.6, B.7, and B.8, the noise efficiencies for each CUORICINO channel and data set are listed. The calculation of the noise efficiencies is described in Sect. 6.2.4. Blank entries in the tables indicate channels and data sets that are excluded from the analysis according to the data selection criteria.

Table B.1: Values of physical parameters used in the analysis. Values without references are deduced from the other numbers.

Parameter	Value	Reference
Isotopic abundance of ^{130}Te	33.80%	[27]
Standard atomic weight of tellurium	127.60	[98]
Standard atomic weight of oxygen	16.00	[98]
Molar mass of TeO_2	159.60 g/mol	
Atomic mass of ^{130}Te	129.91	[99]
Fraction by mass of ^{130}Te in TeO_2	27.5%	
$\beta\beta$ decay Q-value of ^{130}Te	2527.518 ± 0.013 keV	[86]

Table B.2: CUORICINO energy resolutions (FWHM in keV) at 2615 keV.

	Data set											
Channel	1	2	3	4	5	6	7	8	9	10	11	12
1	7.0	6.9	9.1	8.1	6.8	7.8	6.9	9.8	6.9	6.4	7.4	6.2
4	5.3	4.9	6.5	6.2	6.2	5.9	6.0	5.6	4.7	4.3	5.0	5.4
5	5.1	5.7	8.1	4.9	4.8		4.6	5.2	5.0	4.4	5.3	5.1
6	4.8	4.6	4.8	4.8	4.6	5.1	4.5	4.4	4.2	4.8	5.0	5.3
7	5.1	4.5	4.9	4.7	4.4		4.9	4.5	4.9	6.0	5.0	5.7
8	6.8	6.4	4.9	4.5	4.9	5.7	5.0	4.8	5.3	5.6	5.7	4.9
9		8.6	8.4	8.7	8.5	7.4	9.5	9.6	9.3	9.9	8.4	
10	6.1	5.9	5.5	5.6	5.2	6.7	5.1	5.8	5.6	5.0	5.8	6.1
11	7.5	6.6	10.1		8.4	11.0	9.9	10.2	12.1	9.3	8.9	11.1
12	9.8	21.5	16.4	12.4	13.9	13.3	14.3	17.6	16.8	12.7	16.8	19.1
13	8.2	7.3	7.0	6.5	6.9		6.6	7.2	7.5	6.3	6.8	7.4
14		8.9	4.3	4.7	4.5	6.2	4.9	4.7	8.4		8.2	
15	5.0	5.3	4.7	4.4	4.8	4.9	4.7	4.6	4.9	5.9	6.0	5.1
16	4.5	4.8	4.6	4.4	4.4	4.7	4.6	4.5	4.5	4.8	4.7	4.6
17	4.8	6.2	4.0	4.6	4.2		6.6	4.4	4.3	6.1	5.9	5.1
18		8.8	9.6	10.9	8.4	10.7	12.0	10.9	9.9	8.2	8.1	10.5
19	10.2	8.8	9.9	11.2	9.2	11.4	11.1	16.0	9.1	9.9	11.2	12.0
20	30.0	22.0	10.1	18.0	14.9	12.8	14.7	12.5	15.7	24.9	18.2	15.8
21	7.8	9.1	5.0	4.7	4.4	4.6	4.7	4.8	4.6	4.6	5.3	4.7
22	10.0	9.2	8.1	8.0	8.3	8.9		9.9	8.8		10.9	
23	11.3	7.9	9.6	8.7	7.3	8.2	11.3	11.9	7.2	9.4	8.9	13.8
24	9.8	9.5	7.8	7.0	7.3	13.2	8.0	8.2	8.6	8.8	8.4	8.8
31	5.5	5.3	5.5	5.5	5.1	6.5	7.0	5.6	6.2	6.2	5.4	5.8
32	5.0	5.9	6.0	6.0	6.2	6.1	6.1	5.7	5.7	5.7	5.8	5.9
33	20.9	11.0	12.4	10.2	10.9	10.8	12.6	11.4	8.5	11.3	9.4	16.4
34	5.4	5.9	4.5	4.9	4.4		5.0	5.0	4.9	5.2	5.4	5.3
36	7.7	7.1	8.5	7.6	7.9	7.2	8.3	7.2	7.6	7.6	7.0	8.4
37	5.5		5.6	5.8	5.2	6.2	7.7	5.5	5.3	7.0	5.9	7.5
38	4.8	5.3	4.7	5.2	4.3	6.7	5.1	4.5	5.6	5.3	4.9	5.0
40	5.3	5.0	5.9	7.0	6.0	5.8	4.9	5.4	5.1	5.3	5.3	5.0
41	7.5	7.1	6.8	7.0	6.0		5.4	6.2	5.9	5.6	5.6	5.9
42	5.4	5.2	6.6	5.9	6.3	7.8	6.1	5.7	6.7	5.7	6.5	5.9
43	24.4	15.5	11.5	19.5	15.7	17.0	13.3	11.8	11.9	11.7	12.3	11.8
44	5.5	5.7	4.3	4.8	4.6	6.6	4.3	4.1	4.3	6.2	7.2	6.2
45	6.4	6.7	8.4	7.1	9.1	8.5	6.7	6.8	7.4	7.8	7.3	7.0
46	6.1	6.0	5.2	5.7	4.8	6.0	4.8	4.7	4.6	5.1	5.3	5.9
47	4.4	4.6	4.6	4.6	4.8	5.4	5.0	4.8	4.8	4.9	5.6	5.2
48	6.5	6.7	7.6	6.8	6.3	7.0	6.6	6.6	7.0	6.9	6.2	7.0
49	6.6	6.7	7.8	6.6	7.1	6.6	7.3	5.6	7.3	6.7	6.7	7.6
51	6.5	6.0	7.2	9.0	6.7	6.4	5.8	5.7	6.2	6.0	6.2	5.6
53	7.9	8.0	5.5	6.5	9.6	8.0	9.3	6.9	6.8	8.3	7.5	10.8
55	5.0	4.9	4.9	4.5	4.8	4.6	4.7	4.7	4.8	5.1	4.6	5.3
56	5.1	5.4	5.2	4.9	5.0	5.2	4.4	4.4	4.6	5.1	5.3	4.8
57	4.3	4.8	4.2	4.6	4.0		4.8	4.4	4.6	5.6	5.8	5.0
58	6.6	6.7	6.7	6.3	6.5	7.2	6.8	6.5	6.9	8.6	8.8	7.3
59	6.1	6.1	7.2	6.7	10.5		4.9	4.9	6.1	5.7	6.5	6.3
60	8.6	8.0	7.7	7.7	5.7	5.9	5.2	5.0	5.1	5.1	5.7	5.3
61	26.6		8.7		10.5	9.0			9.2	9.4	8.5	10.0
62	8.6	8.8	9.4	8.6	9.1	13.7	9.1	9.5	8.6		8.7	8.3
63	5.3	5.6	4.7	5.2	4.5	9.2	4.9	10.8	13.2	9.4	9.3	9.5
64	6.8	5.5	5.3	6.2	6.2		5.5	5.1	6.0	5.9	5.5	8.1
65	6.9	6.9	10.8	9.0	9.5	8.5	7.4	7.4	7.1	6.4	6.5	7.1
66	4.9	5.1	8.8	7.5	5.5	6.5	5.0	5.0	5.4	5.0	4.6	5.2
67	5.1	4.7	4.4	5.0	4.2	4.7	4.7	4.4	4.9	5.3	4.8	4.8
68	6.5	6.3	6.8	5.2	5.3		5.7	5.5	5.5	5.8	6.3	6.0

Table B.3: CUORICINO energy resolutions (FWHM in keV) at 2615 keV.

	Data set											
Channel	13	14	15	16	17	18	19	20	21	22	23	24
1	8.5	8.0	12.5	6.1		6.8	5.9	6.0	6.3	7.3	6.7	6.9
4	5.8	4.8	4.6	4.5	7.9	4.7	4.5	4.5	4.4	5.1	5.1	5.1
5	6.3	5.2	4.7	4.5		12.8	11.2	14.5	14.4	6.1	5.4	5.8
6	5.0	4.5	4.3	5.0	6.1	5.1	4.9	4.6	4.6	5.1	5.2	4.7
7	6.8	5.1	4.9	5.3	9.5	5.2	5.2	5.2	5.3	5.7	5.2	5.2
8	6.2	5.4	5.1	5.6		5.6	5.1	5.2	5.0	6.0	5.3	5.2
9	10.6	10.7	8.6	6.8		7.5	7.2	8.4	8.6	8.9	7.6	8.7
10	7.8	6.6	5.7	6.1	16.9	6.4	5.9	5.5	5.5	7.2	5.4	5.0
11	10.2	9.2	7.9	12.5	8.3	9.2	10.4	8.8	7.9	9.3		8.1
12	20.3	13.0	16.7	8.6	30.0	21.2	17.9	17.0	16.9		15.1	15.3
13	6.9	6.8	6.9	7.4	7.8	6.5	4.2	6.4	6.3	8.1	7.5	6.9
14	7.4	6.3	5.7	5.6	8.0	5.8	5.8	5.7	5.8	5.8	5.5	4.7
15	6.5	5.2	4.6	5.2		6.1	5.6	5.9	5.9	5.1	4.6	4.8
16	5.1	4.9	4.8	6.0	5.8	4.6	4.6	4.6	4.6	5.1	4.6	4.6
17	5.5	4.6	4.6	4.4		5.3	5.1	5.0	4.8	5.6	4.2	4.4
18	9.6	8.7	10.7	10.0		12.7	13.3	11.5	11.5	11.2	10.3	10.1
19	9.5	9.7	10.8	14.4		9.5	10.4	9.5	8.8	11.1	10.2	9.4
20	18.8	16.6	14.7	17.3	18.7	19.8	17.9	14.6	15.0	19.3	18.8	12.8
21	4.9	4.5	4.4	4.3		5.5	4.9	5.2	4.7	5.2	4.6	4.5
22	10.6	10.3	8.5	8.4		7.6	8.1	7.3	7.0		7.5	7.6
23	11.3	8.4	10.4	7.7		9.1	10.5	8.7	8.7	13.5	8.5	9.2
24	8.2	8.7	8.3	9.7	8.0	7.4	7.2	7.9	7.7	9.0	9.1	9.6
31	5.8	5.4	5.5	7.2	9.3	5.8	5.0	5.6	5.6	6.1	6.1	5.8
32	6.0	5.0	5.5	5.4	9.2	6.5	5.8	6.3	6.3	6.3	6.8	6.9
33	13.7	11.7	11.2	10.9	10.4	11.1	12.6	11.5	11.3	11.6	10.3	10.2
34	6.1	5.3	4.7	4.7		5.7	5.2	5.4	5.0	5.9	4.7	4.9
36	8.6	7.5	8.1	7.4		8.0	7.4	8.0	8.2	9.1	7.4	7.2
37	6.4	6.1	7.2	6.9		5.9	5.6	5.4	5.3	5.9	5.3	5.3
38	5.4	6.2	4.9	5.0	7.3	5.8	4.6	5.5	5.4	5.6	5.8	5.4
40	5.4	5.3	5.2	5.0	7.1	5.2	5.0	5.3	5.3	5.5	5.2	5.0
41	7.9	6.4	5.7	6.3	9.4	5.9	5.2	5.7	5.7	6.1	6.2	7.2
42	6.2	5.5	6.0	7.2	6.2	5.5	4.8	7.4	7.4	6.4	5.6	5.5
43	10.8	10.8	12.3	17.2	12.1	12.3	13.0	12.3	12.2	13.4	13.1	12.5
44	6.9	6.8	5.3	4.0	6.6	5.0	4.7	4.9	5.0	4.4	4.5	4.8
45	8.0	9.7	6.2	11.0	10.4	10.6	8.2	9.9	10.2	6.6	6.1	6.6
46	6.1	6.3	5.1	7.1	5.7	5.2	5.1	5.3	5.3	5.7	5.2	5.3
47	6.1	5.9	4.6	5.3		5.7	5.2	5.5	5.4	6.3	6.1	5.9
48	6.6	6.9	7.2	6.3	7.2	7.0	6.8	6.9	6.8	6.3	6.5	6.6
49	7.3	7.5	6.9	7.2	8.6	6.3	7.4	6.6	6.6	8.3	6.7	5.8
51	6.2	6.3	6.6	5.0	9.0	7.3	5.7	6.4	6.1	6.4	6.2	6.4
53	7.6	6.4			11.1	7.8	7.0	8.0	7.9	7.6	7.4	7.1
55	5.2	4.9	5.0	6.1	5.7	4.9	5.1	5.2	5.1	5.2	4.7	4.8
56	7.4	5.1	5.2	5.0		4.7	5.0	5.0	5.0	5.8	4.8	4.9
57	6.2	4.6	3.9	3.7		4.5	4.1	4.3	4.2	5.4	5.0	5.2
58	7.7	7.4	11.7	7.4	11.5	8.4	7.5	7.9	7.9	6.5	6.3	6.9
59	6.9	4.6	4.8	5.4	7.6	5.7	5.6	5.3	5.3	5.8	5.9	6.2
60	5.1	5.8	5.8	5.1	7.7	6.0	5.3	5.0	5.0	5.4	5.3	5.7
61	8.0	8.9	10.9	7.4		8.6	8.9	10.1	10.4		8.5	7.4
62	7.8	7.7	9.0	8.7	9.0	8.1	7.5	8.1	8.1	8.3	7.4	7.4
63	11.0	14.9	17.8	6.2	6.5	5.3	5.1	5.5	5.5	5.9	5.4	5.8
64	6.2	6.1	5.4	7.8	9.5	7.8	7.9	6.3	6.3	5.9	6.0	7.3
65	7.0	6.9	7.2	6.8	8.6	7.6	7.9	7.3	7.3	8.5	7.9	7.6
66	6.8	6.0	8.8	5.1		6.7	8.3	5.4	5.4	6.2	5.6	5.8
67	5.0	5.1	4.9	5.9		5.0	4.6	5.3	5.1	5.5	5.3	5.5
68	6.7	6.0	5.8	6.1	6.9	6.1	30.0	5.7	6.1	6.3	6.0	6.5

Table B.4: CUORICINO energy resolutions (FWHM in keV) at 2615 keV.

	Data set							
Channel	25	26	27	28	29	31	32	33
1	7.7	7.3	7.0	7.6	7.9	7.5	10.1	10.4
4	5.5	5.0	4.7	5.5		4.9	5.3	5.3
5	6.7	6.1	5.6	5.1	4.7	5.4		
6	4.4	4.7	5.1	4.6	4.4	4.6	5.2	5.2
7	5.8	5.5	4.5		6.9	6.8	4.6	4.5
8	5.0	5.1	5.5	6.7	6.3	5.0	4.5	4.8
9	10.1	10.0	7.6	11.3	10.8	9.4	9.3	9.3
10	6.2	5.9	5.0	6.4		5.9	5.0	4.8
11	6.1	14.1	14.3		9.2	9.6	8.9	8.9
12	10.6	11.5	7.1	25.2	26.9	14.1	10.8	10.9
13	6.7	6.7		7.2	7.3	6.8	7.2	7.0
14	4.6	5.8	5.1	5.7		5.0	4.2	4.1
15	5.9	5.3	4.6	5.7	5.2	5.0	4.2	4.2
16	5.2	5.1	4.4		5.0	5.8	4.6	4.4
17	5.1	5.0	4.6	5.1		5.6	4.5	4.5
18	10.5	11.4	11.5	14.5	15.0	12.4	10.7	10.7
19	10.7	12.0	10.3		11.4	8.3	11.6	11.5
20	24.5	29.9	15.6	21.6	14.4	11.4	9.6	10.5
21	5.2	4.8	4.3	5.3	4.7	4.7	4.2	4.2
22	7.3	11.7	7.8	10.2	8.1	8.8	7.8	7.8
23	8.8	9.8	9.6	15.2	10.1	8.3	8.3	8.4
24	6.6	6.4	6.6	11.1	9.9	7.6	7.1	7.2
31	6.6	5.8	4.7	7.9	5.7	7.1	5.6	5.5
32	5.8	5.9	6.3	5.9	5.3	4.6	5.6	5.6
33	10.4	15.9	10.6	13.5	11.2	10.1	10.4	10.3
34	5.4	5.1	4.6	5.6		5.2	5.0	5.1
36	8.1	8.8	7.7	11.2	10.6	7.8	9.0	9.0
37	5.5	6.7	5.4	9.7	8.7	5.5	5.0	5.0
38	5.0	4.9	4.8	5.0	5.0	4.4	4.2	4.3
40	5.1	5.2	5.6	4.9	5.4	5.4	5.5	5.6
41	7.7	6.9	6.0	9.3	10.7	6.2	5.9	6.2
42	7.4	6.6	6.2	10.1	9.4	8.2	6.7	6.7
43	12.8	11.9	10.2	14.1	12.7	11.9	13.4	14.8
44	4.9	4.7	4.5	5.5	7.3	4.3	3.6	3.8
45	8.0	7.0	6.1	8.0		8.5	5.5	5.5
46	5.6	5.7	5.5	7.8	5.5	6.9	4.2	4.3
47	6.1	5.7	5.4	5.7	4.9	5.1	4.8	4.8
48	7.2	8.3	6.9			7.1	4.9	4.9
49	6.5		4.6		6.4	5.9	5.2	4.9
51	5.4	5.8	6.5	6.7	6.1	6.3	6.3	6.2
53	5.7	7.3		9.6	8.8		9.1	9.0
55	5.7	5.5	4.9	7.3	4.7	5.7	4.1	4.1
56	5.2	5.4	5.9	5.1	5.8	5.0	4.8	4.8
57	4.1	4.0	4.1	4.8	6.9	4.1	4.2	4.2
58	6.8	6.8	6.7	7.8	9.0	6.4	6.4	6.4
59	5.2	5.2	5.1	5.4		5.2	5.3	4.6
60	5.2	5.3	5.7	6.2	6.7	4.8	5.0	5.0
61	7.8	8.8		13.7	13.2	8.4	8.5	8.4
62	7.6	7.9	9.9	8.0		8.0	7.0	6.9
63	6.0	5.8	5.0	7.0	5.8	6.3	5.1	5.4
64	5.9	5.9	6.0	5.9	7.2	5.5	5.8	5.9
65	7.7	7.6	7.4	7.5	10.2	7.3	11.3	11.3
66	6.3	5.7	5.8	6.9	7.0	5.1	7.7	7.5
67	5.2	5.1	4.5	7.3	7.0	5.8	5.0	5.0
68	6.5	6.4	6.4	5.9	5.6	5.8	5.3	5.1

Table B.5: Three Towers Test energy resolutions (FWHM in keV) at 2615 keV.

	Data set			
Channel	1001	1002	1003	1004
1	4.8	4.9	5.1	4.9
3	4.8	4.6	4.3	5.9
4	5.5	7.6	5.0	5.1
5	4.6	4.3	4.4	4.8
7	4.6	4.3	4.1	4.7
13	4.7	4.5	4.6	4.7
14	5.5	4.5	4.7	4.4
16	6.6	5.6	5.8	7.1
17	7.8	8.5	7.9	6.7
18	8.1	11.2	10.5	11.4
19	6.4	5.6	5.0	6.7
27	4.4	4.5	4.7	4.4
31	7.8	6.2	3.9	13.4
32	21.5	22.1	13.8	11.8
35	5.4	4.5	4.8	8.0
40	6.5	5.9	5.4	7.4
44	11.5	8.4	8.0	11.1
49	6.4	5.9	7.7	6.2
51	11.2	7.9	7.4	13.5
57	8.4	9.5	8.0	9.6
58	9.5	6.8	7.6	6.8
59	7.3	6.9	7.1	7.0
61	6.8	6.3	6.1	6.0
62	5.4	5.3	5.2	5.0
63	9.7	10.8	10.0	9.4
64	6.5	6.6	7.0	7.9

Table B.6: CUORICINO noise efficiencies (%) for data sets 1–12.

	Data set											
Channel	1	2	3	4	5	6	7	8	9	10	11	12
1	99.5	96.5	94.4	96.3	99.7	99.0	99.0	99.4	99.2	99.1	99.5	96.0
4	99.5	97.4	95.3	96.8	99.8	99.7	99.5	99.7	99.3	99.8	99.8	99.6
5	99.1	97.8	94.4	95.4	99.6		99.0	99.4	99.5	99.8	99.8	99.7
6	98.4	97.6	95.3	96.1	99.5	99.7	99.6	99.6	99.6	99.7	99.8	99.8
7	99.6	97.7	93.8	94.9	99.4		99.3	99.5	99.6	99.8	99.6	99.7
8	99.4	98.0	93.2	94.3	99.4	99.2	99.1	99.4	99.5	99.6	99.7	99.6
9		98.0	97.4	97.4	100.0	100.0	99.6	99.8	99.6	99.6	99.6	
10	98.9	97.7	96.4	97.8	99.6	99.2	99.4	99.2	99.6	99.7	99.6	99.5
11	99.5	98.4	96.7		99.7	99.6	99.5	99.6	99.9	99.7	99.6	99.4
12	99.5	98.1	97.3	98.4	99.4	99.6	100.0	99.6	99.8	99.9	100.0	99.8
13	99.3	97.5	96.2	97.9	99.0		99.6	99.5	99.8	99.8	99.9	100.0
14		99.1	94.3	95.2	99.4	99.1	98.9	99.5	99.2		99.8	
15	99.5	98.1	94.6	95.3	99.7	99.6	99.4	99.6	99.6	99.8	99.5	99.7
16	99.4	97.9	94.9	96.0	99.8	99.7	99.5	99.6	99.7	100.1	99.6	99.9
17	99.2	98.1	94.0	94.9	99.5		99.4	99.4	99.5	99.4	99.6	99.5
18		98.6	97.4	98.4	100.0	100.0	99.5	99.9	100.0	100.0	100.0	100.0
19	99.7	98.3	96.7	98.0	100.0	99.8	99.6	99.7	100.0	99.9	99.9	99.9
20	99.6	98.5	98.3	98.8	100.0	98.9	99.5	99.5	100.1	99.8	99.2	99.8
21	99.2	98.3	96.7	97.4	99.7	99.7	99.7	99.6	99.5	99.7	99.8	99.7
22	99.2	98.3	96.2	97.3	100.2	99.5		99.7	99.8		99.8	
23	99.3	100.2	97.2	97.7	99.8	99.9	99.7	99.6	99.6	99.9	99.9	99.8
24	98.7	96.7	96.8	98.1	99.4	99.7	99.6	99.6	99.9	99.8	99.7	99.8
31	99.1	98.1	96.6	97.8	99.7	99.4	99.4	99.4	99.5	99.6	99.5	99.5
32	99.3	98.3	95.5	96.6	99.5	99.3	98.8	99.2	98.7	99.7	99.5	99.3
33	99.3	98.4	98.1	98.6	99.9	99.7	99.7	99.1	99.9	99.8	99.9	100.0
34	99.3	98.0	94.9	96.5	99.6		99.3	99.3	99.4	99.7	99.5	99.7
36	99.5	98.2	95.7	95.7	99.5	99.7	99.4	99.5	99.9	99.7	100.0	99.4
37	99.2		92.8	93.9	99.5	99.3	99.4	99.6	99.5	99.5	99.7	99.8
38	99.1	96.6	92.8	94.2	99.5	99.5	99.2	99.2	99.5	99.2	99.7	99.4
40	98.2	96.2	90.5	95.3	99.5	99.5	98.8	99.1	99.0	99.4	99.3	99.4
41	99.3	97.9	93.3	94.5	99.4		99.4	99.4	99.2	99.5	99.6	99.7
42	99.1	98.2	95.3	95.3	99.7	98.7	97.6	99.1	99.0	99.7	99.7	99.8
43	99.2	98.1	87.9	92.3	98.2	95.8	98.7	99.4	99.0	96.2	99.2	99.7
44	99.5	98.0	95.3	96.2	97.5	99.8	99.2	99.7	99.3	99.6	99.8	99.7
45	97.9	91.5	93.6	93.3	94.8	98.4	98.0	99.5	99.2	99.6	99.8	99.4
46	99.4	97.7	96.1	97.1	99.7	99.4	99.1	99.3	99.4	99.8	99.7	99.8
47	98.5	97.0	91.2	93.3	99.6	99.5	99.2	99.0	99.3	99.5	99.4	99.5
48	91.5	91.1	96.4	96.2	99.7	99.6	99.0	99.6	99.6	99.8	99.7	99.6
49	99.1	96.9	96.7	97.3	99.6	99.2	99.2	99.4	99.6	99.8	99.8	99.7
51	99.1	97.3	91.2	91.7	99.6	99.3	98.7	99.4	99.1	99.6	99.6	99.5
53	99.3	97.4	92.5	93.9	99.6	99.3	98.7	99.5	99.8	99.7	99.7	99.8
55	98.8	97.8	92.6	93.4	99.4	99.4	99.4	99.4	99.1	99.3	99.6	99.5
56	99.0	97.7	94.2	95.1	99.6	99.6	99.5	99.5	99.6	99.4	99.7	99.6
57	98.8	97.3	93.6	95.6	99.5		98.1	99.3	99.5	99.6	99.7	99.7
58	98.3	96.7	95.6	96.4	99.7	99.4	99.6	99.5	99.5	97.5	96.3	97.1
59	99.0	95.7	89.2	77.4	82.8		98.1	99.5	99.3	96.9	99.8	96.7
60	99.4	97.5	92.6	93.9	97.4	99.6	99.3	99.5	99.4	99.6	99.8	99.7
61	99.5		91.5		98.3	98.1			99.2	99.6	99.2	99.7
62	99.0	98.4	97.1	98.5	99.9	99.5	99.5	99.4	99.6		99.8	99.6
63	98.8	98.2	94.8	95.7	99.5	65.7	65.8	85.7	92.7	99.0	98.1	99.5
64	99.2	97.9	96.3	97.7	99.5		99.7	99.5	99.6	99.7	99.9	99.6
65	98.9	97.5	94.6	95.1	91.9	97.5	96.7	99.5	99.5	99.6	99.6	99.7
66	97.4	96.6	89.6	84.7	80.6	97.5	92.6	98.7	99.2	99.0	99.5	99.1
67	98.9	98.0	95.4	97.1	99.4	99.4	99.2	99.4	99.4	99.4	99.5	99.5
68	98.8	97.0	93.4	94.0	99.5		99.5	99.4	99.4	99.5	99.2	99.6

Table B.7: CUORICINO noise efficiencies (%) for data sets 13–24.

	Data set											
Channel	13	14	15	16	17	18	19	20	21	22	23	24
1	98.5	99.8	60.4	100.0		98.8	99.7	99.0	99.7	97.2	97.2	98.5
4	100.1	99.5	100.6	99.7	98.7	97.5	99.7	99.9	99.8	99.3	98.8	99.4
5	99.9	99.7	99.4	99.3		94.8	92.5	95.7	97.2	99.7	99.6	100.0
6	100.1	99.7	99.7	97.0	98.4	99.5	99.9	99.5	99.9	99.5	99.2	99.8
7	99.8	99.9	99.7	99.5	98.5	99.6	99.4	99.7	99.7	99.7	99.4	99.5
8	100.1	99.8	99.7	99.6		98.9	99.6	99.6	99.4	99.3	98.7	98.7
9	100.1	99.9	99.7	100.0		99.7	99.9	100.0	99.9	99.1	99.7	100.2
10	99.6	99.9	99.7	99.7	99.3	99.3	99.5	99.6	99.7	99.5	99.0	99.0
11	99.5	99.3	99.5	99.5	99.7	99.9	99.5	99.6	100.0	99.6		99.6
12	99.7	100.1	99.7	89.6	99.9	99.9	99.9	99.9	99.9		99.7	99.2
13	100.1	99.6	98.3	99.3	96.2	98.7	98.7	99.7	99.8	99.3	98.9	99.8
14	100.2	98.4	98.6	99.6	98.6	99.7	99.6	99.9	99.5	99.5	99.0	99.9
15	99.9	99.6	99.5	99.2		99.2	99.5	99.6	99.8	99.3	98.8	99.4
16	100.0	99.7	99.8	100.1	99.6	100.0	99.7	99.7	99.4	99.7	99.6	99.8
17	99.9	99.9	99.7	99.5		99.3	99.5	99.2	99.7	99.2	99.2	99.6
18	100.2	100.0	99.8	99.8		99.7	99.8	99.8	99.9	99.9	99.7	99.7
19	100.0	100.0	100.0	99.9		100.0	99.7	100.0	99.9	99.9	99.9	99.9
20	100.2	100.0	98.3	99.6	100.0	98.4	96.9	99.6	98.2	99.9	100.0	100.0
21	99.9	100.0	99.8	99.6		99.8	99.8	99.9	99.5	99.7	99.6	99.8
22	100.1	99.9	99.8	99.9		99.8	99.8	99.8	99.4		99.7	99.3
23	99.9	99.8	100.0	100.0		99.8	99.4	99.9	99.9	99.8	99.8	100.0
24	100.1	99.9	99.9	99.8	98.6	99.7	99.7	99.4	99.9	94.6	93.8	95.1
31	100.0	99.4	98.1	99.5	99.3	99.4	99.7	99.2	99.5	99.5	99.1	99.5
32	100.0	99.8	99.5	99.9	97.8	98.6	99.7	99.8	99.8	96.2	93.5	99.3
33	99.9	99.6	97.6	99.7	98.9	100.0	99.2	99.2	99.6	99.8	99.7	99.7
34	99.9	99.8	99.4	99.6		99.8	99.4	99.3	99.7	99.5	99.2	98.8
36	100.1	99.7	99.7	99.9		99.7	99.7	99.5	99.9	99.6	99.6	99.4
37	100.1	99.6	100.0	99.5		99.7	99.7	100.0	99.7	99.6	99.1	99.5
38	100.0	99.8	99.6	98.7	97.6	98.9	99.4	98.7	99.4	99.1	98.2	98.9
40	99.5	99.6	99.5	99.3	97.3	99.5	99.0	98.9	99.2	99.5	98.8	99.4
41	99.5	99.7	99.7	99.7	97.3	99.0	99.5	99.3	99.8	99.3	99.1	99.5
42	99.9	99.7	99.7	96.4	98.3	99.4	99.5	99.7	99.9	99.6	98.7	99.8
43	99.8	99.1	99.5	99.9	98.7	98.4	97.7	97.1	99.3	98.2	97.1	99.4
44	99.8	99.0	98.2	99.2	98.6	99.6	99.6	99.9	99.7	99.6	99.5	99.8
45	100.0	99.6	99.6	98.3	97.0	96.1	93.8	98.0	99.1	99.4	99.4	99.6
46	99.9	99.5	99.5	99.9	99.1	99.3	99.3	99.3	99.7	99.3	99.1	99.8
47	99.5	99.4	99.4	99.8		99.3	99.2	99.6	99.0	98.4	98.4	99.2
48	100.1	99.8	98.7	99.9	99.6	99.7	99.8	99.9	99.7	99.5	99.5	99.8
49	100.0	99.7	99.7	99.9	99.6	99.8	99.7	99.5	99.6	99.7	99.6	99.7
51	99.7	99.7	99.4	98.8	98.4	99.5	99.5	99.6	99.8	99.1	99.3	99.1
53	99.9	99.8			99.5	99.5	99.8	99.3	99.8	99.6	99.7	99.7
55	99.7	99.4	99.4	99.9	99.2	99.6	99.5	98.9	99.6	99.4	99.2	99.9
56	100.0	99.5	99.5	99.7		99.0	99.6	99.9	99.6	99.5	99.2	99.5
57	99.8	99.5	99.5	99.9		99.1	99.4	99.4	99.6	98.7	97.8	98.5
58	99.3	97.9	73.7	95.3	98.8	95.6	89.2	97.2	98.3	99.6	99.4	99.6
59	99.3	92.1	99.6	100.0	99.1	99.3	99.8	99.7	99.8	99.3	98.5	97.5
60	99.7	99.5	99.7	99.7	81.1	84.4	98.9	99.5	99.9	99.1	99.1	99.3
61	99.9	99.8	99.1	99.9		99.9	99.7	100.0	99.8		99.4	99.7
62	99.8	99.8	99.8	99.4	98.3	98.9	99.6	99.7	99.9	98.9	98.5	99.6
63	99.8	99.5	96.1	81.2	98.9	99.6	99.5	99.1	99.3	99.5	99.5	99.3
64	100.1	99.9	99.6	96.4	98.1	99.5	99.9	99.9	100.1	99.2	98.6	98.6
65	99.8	99.8	99.8	99.6	98.6	99.2	99.4	99.5	98.7	99.1	98.9	99.7
66	99.7	99.4	93.7	99.4		96.8	72.5	94.9	98.8	98.6	98.1	98.9
67	99.9	99.8	99.5	99.2		99.2	99.4	99.6	99.6	99.3	99.1	99.6
68	99.7	99.3	99.4	99.1	97.0	97.8	99.8	99.7	99.6	99.1	98.8	99.3

Table B.8: CUORICINO noise efficiencies (%) for data sets 25–33.

	Data set							
Channel	25	26	27	28	29	31	32	33
1	96.5	98.9	98.7	98.9	99.1	98.1	93.3	92.6
4	99.4	99.6	99.9	99.7		99.8	96.6	97.6
5	99.4	99.7	99.8	99.7	99.9	99.7		
6	99.4	99.8	99.8	99.8	99.9	99.8	98.4	98.2
7	99.2	99.5	99.6		99.7	99.7	98.6	98.4
8	99.0	99.0	98.8	99.3	99.6	99.6	98.4	97.8
9	99.7	99.8	99.8	99.2	99.9	99.9	99.1	98.9
10	99.3	99.6	99.6	99.7		99.8	98.7	98.4
11	98.8	99.6	100.0		99.6	99.8	98.8	99.0
12	99.7	99.9	95.1	99.9	100.1	99.8	99.1	99.3
13	99.5	94.6		99.6	99.8	99.8	99.0	99.1
14	99.3	99.6	99.6	99.3		99.8	97.0	97.4
15	99.4	99.5	99.6	99.8	100.0	99.8	98.9	98.3
16	99.5	99.8	100.0		100.0	99.8	99.0	98.5
17	97.9	99.5	99.6	99.7		98.6	97.5	97.8
18	99.7	100.0	99.9	99.9	100.1	99.9	99.2	99.6
19	99.5	99.9	99.9		100.1	99.9	99.0	99.1
20	99.8	99.7	100.1	97.5	98.8	99.0	98.7	99.4
21	99.6	99.8	99.8	99.9	100.0	99.7	98.9	98.8
22	99.4	99.7	99.7	96.8	99.9	99.5	97.7	98.3
23	99.7	99.8	100.0	99.8	100.0	99.8	98.9	99.1
24	95.7	98.7	99.6	99.2	97.6	98.1	98.4	99.3
31	98.7	99.6	99.9	99.4	99.7	99.5	98.4	98.3
32	99.3	99.5	99.1	99.4	99.5	99.7	98.6	98.6
33	99.6	99.8	99.9	99.8	100.1	99.8	98.9	98.6
34	99.2	99.6	99.9	99.6		99.2	98.7	98.4
36	99.1	99.8	99.5	99.9	99.9	99.8	99.0	97.3
37	99.5	99.5	99.8	99.8	99.9	99.8	98.9	95.0
38	99.2	99.4	99.5	99.3	99.8	99.4	98.6	93.0
40	99.3	99.3	99.4	99.7	99.6	99.3	97.9	96.1
41	96.0	98.9	92.8	96.4	95.4	97.3	97.7	96.1
42	97.4	98.8	97.3	98.1	95.6	98.0	98.9	94.6
43	86.9	94.6	92.0	97.0	98.9	97.7	98.6	99.1
44	98.4	99.7	99.5	99.8	99.8	99.7	98.7	98.8
45	96.9	97.9	94.8	93.0		97.6	97.5	96.2
46	99.0	99.5	99.6	99.7	99.9	99.8	98.4	98.8
47	98.9	99.3	99.7	99.2	99.6	99.5	98.7	97.1
48	99.4	99.7	100.0			99.8	98.8	98.3
49	99.5		99.9		99.9	99.6	98.9	98.7
51	99.3	99.4	99.6	99.4	99.8	99.6	98.6	97.7
53	99.5	99.7		99.8	99.9		98.8	98.5
55	99.1	99.5	100.0	99.6	99.6	99.7	98.7	97.2
56	99.2	99.6	99.8	99.8	99.8	99.7	98.7	97.9
57	99.2	99.5	99.8	99.6	99.6	99.5	98.8	98.2
58	99.2	99.7	99.7	99.7	100.0	99.7	96.0	98.7
59	99.2	99.5	99.7	99.5		99.6	97.7	97.6
60	94.8	98.3	98.6	99.2	99.0	98.4	98.5	98.3
61	99.4	99.4		99.4	99.8	99.5	98.8	97.9
62	99.5	99.4	99.8	99.6		99.7	98.9	98.9
63	99.1	99.6	100.0	99.6	99.7	99.8	98.6	97.9
64	99.3	99.7	99.8	99.7	99.9	99.7	98.5	98.1
65	99.0	99.5	99.9	99.5	99.6	99.7	96.7	94.6
66	89.1	99.4	99.7	98.9	97.4	97.6	86.7	88.4
67	98.9	99.5	99.7	99.5	99.7	99.8	98.4	98.3
68	99.4	99.4	99.9	99.0	99.6	99.5	98.4	97.8

Appendix C

Generalized pulse amplitude measurement algorithm

One of the main steps in the first-level data analysis of CUORICINO or CUORE bolometers is to measure the pulse amplitudes in a way that maximizes the signal-to-noise ratio in order to obtain the best possible resolution. Different approaches to solving this problem can be taken, including, for example, filtering or pulse fitting. The standard method used in CUORICINO analyses is the optimal filter, discussed in Sect. 4.4. This appendix describes an alternative algorithm that was developed to measure the pulse amplitudes with minimum variance due to noise. The method has similarities to a frequency-domain digital filter but is not a filter per se. It may be considered to be more general than the optimal filter because it considers the full covariance matrix of the Fourier components of the noise rather than only the noise power spectrum.

C.1 Description of the problem

A CUORICINO pulse consists of 512 digitized values recorded by a 16-bit analog-to-digital converter (ADC). An example CUORICINO pulse is shown in Fig. C.1. The time interval spanned by the 512 values is 4.096 seconds, and the range of 16-bit values (0 to 65535) corresponds to 0 V to 10 V. For CUORE the sampling rate may be higher and the acquired time interval longer, but the principle of the amplitude algorithm is the same.

The resolution goal for CUORE is 5 keV FWHM at 2527.5 keV. Translating FWHM to Gaussian sigma, the allowed fractional error on the amplitude is

$$\frac{\sigma_E}{E} = \frac{5\text{ keV}/(2\sqrt{2\ln 2})}{2527.5\text{ keV}} = 8.4 \times 10^{-4}. \tag{C.1}$$

Therefore, pulse amplitude measurements must be accurate to better than 0.08%.

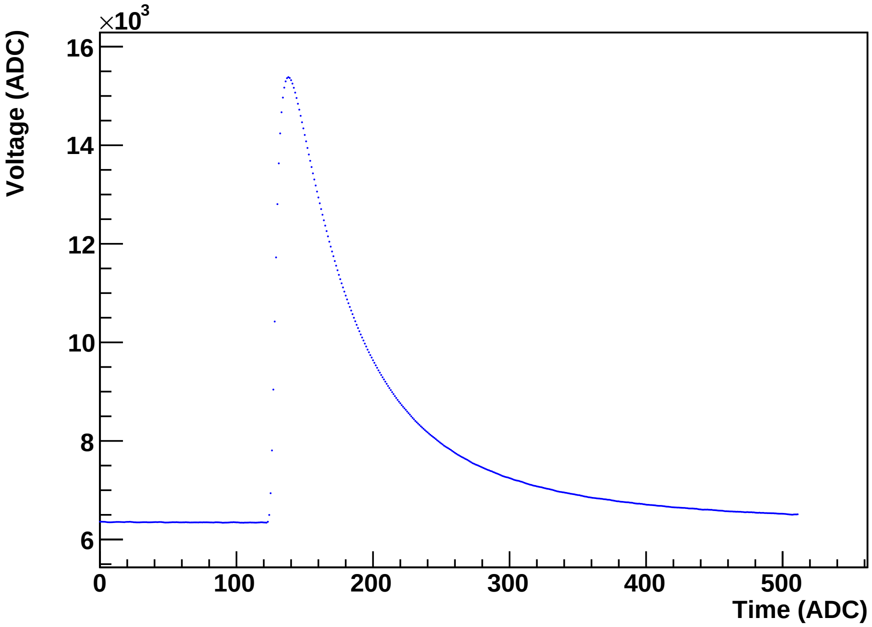

Figure C.1: A CUORICINO pulse on channel 7 of run 737. This pulse was caused by a gamma from the 2615 keV line of ^{208}Tl. The units on the vertical axis are ADC values, and the units on the horizontal axis are digital samples.

C.2 Variation in pulse onset time, t_0

The recording of pulses is controlled by a threshold trigger. For CUORICINO the first 125 points (1 s) of the acquired time window are before the trigger fired, allowing the pre-pulse baseline level to be measured. Since the triggering is based on the pulse exceeding the baseline level by a specified amount, pulses with different amplitudes cause the trigger to fire at different points in their evolution: Larger amplitude pulses exceed the threshold earlier in their evolution, i.e. closer to the actual start time, t_0, of the pulse. Therefore, a larger amplitude pulse is shifted to the right within the acquired window compared to a smaller amplitude pulse (Fig. C.2). This effect is known as slewing.

The first step in the amplitude calculation is to correct for this variation in pulse onset time. The correction is necessary because the amplitude algorithm is based on the shape of the pulse in the frequency domain (as will be described in Sect. C.3). It might appear that a time translation of the pulse merely affects the phase of the Fourier components, and this is true if and only if the pulse returns to the baseline level within the acquired window. The pulse shape is different for every channel, and for some channels the pulse has a long decay time and does not return to the baseline within the window (Fig. C.3). For these channels, the Fourier component magnitudes, in addition to their phases, depend on the onset time of the pulse within the window. Therefore, it is important to perform an alignment of the pulses. The alignment is performed by determining the t_0 offset with respect to the average pulse and truncating an appropriate number of points at the beginning and end of the pulse as described in the following two subsections.

C.2.1 Determining the t_0 offset

To align the pulses on a particular channel, it is enough to know the offset between the t_0 of each pulse and that of the average pulse for the channel, t_0^{avg}:

$$t_{\text{offset}} \equiv t_0 - t_0^{\text{avg}}. \tag{C.2}$$

Determining the offset is much easier than determining the actual t_0. The offset is determined by fitting the derivative of the average pulse to the derivative of the pulse allowing the average pulse derivative to shift left or right while normalizing the height of the average pulse to the height of the pulse, where for this purpose the height is taken as the maximum minus the baseline (Fig. C.4). One could instead allow the height to float in the fit; this was tested, and no significant difference was found, so the height is fixed in order to improve the speed of the fit. In order to obtain a continuous function representing the average pulse derivative, a cubic-spline interpolation is used. The function is set to 0 outside the bounds of the

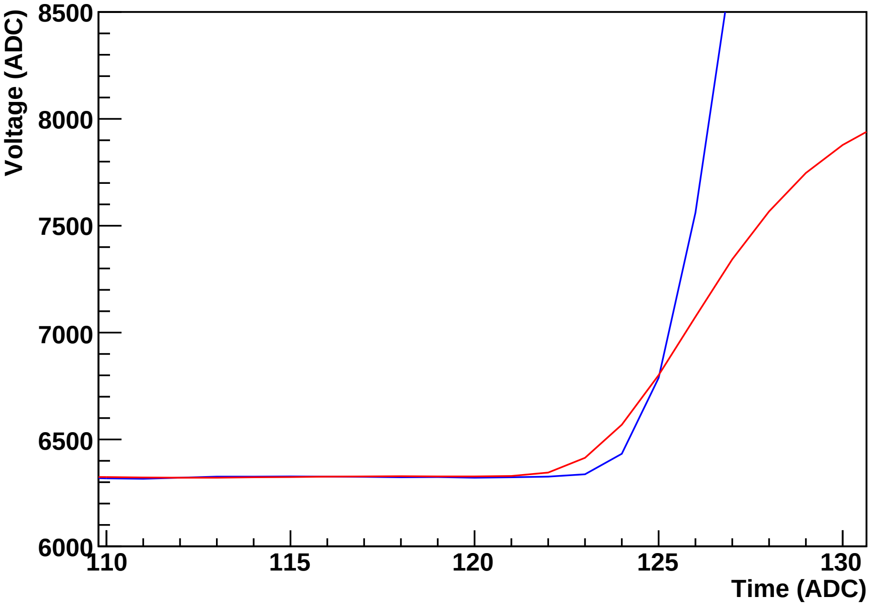

Figure C.2: The pulse onset time, t_0, measured with respect to the start of the acquired window varies with pulse amplitude. The larger amplitude pulse (blue) is shifted to the right with respect to the smaller amplitude pulse (red); this is easy to see from the fact that for a time the smaller pulse exceeds the larger one. The energy of the blue pulse is 2615 keV, and the energy of the red pulse is 511 keV. Point 125, where the two pulses exceed their baseline levels by the same amount, corresponds to the time that the threshold was crossed.

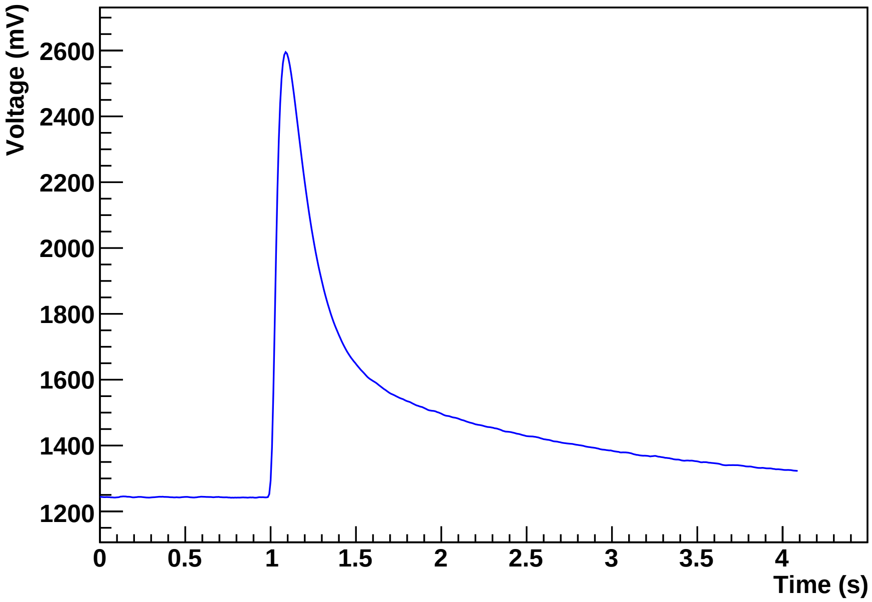

Figure C.3: This 2615 keV CUORICINO pulse on channel 6 of run 1203 has a long decay time and does not return to the baseline within the acquired window.

original acquired window, and the fit is performed excluding the first 16 points and last 16 points in the window so that the average pulse derivative function should not have to be evaluated outside the bounds of the original acquired window. The derivatives are computed as two-sided derivatives by averaging the derivative to the left and the derivative to the right.

The error on the offset may be estimated by adding noise samples to the average pulse and running the noise-added pulses through the algorithm. Of course, the error is greater for smaller amplitude pulses so the average pulse should be scaled appropriately. As an example, the error on the offset for a pulse on channel 1 with an energy close to 2.5 MeV is estimated to be 0.017, i.e. 1.7% of the spacing between digitized samples (Fig. C.5).

In Fig. C.6, the pulse amplitude vs. offset is plotted for channel 1 of the calibration run 737. The figure shows how lower amplitude pulses are shifted left, toward more negative offsets. The figure also shows the inherent jitter in the trigger, such that for a given amplitude the offsets vary within a range of about 1 unit.

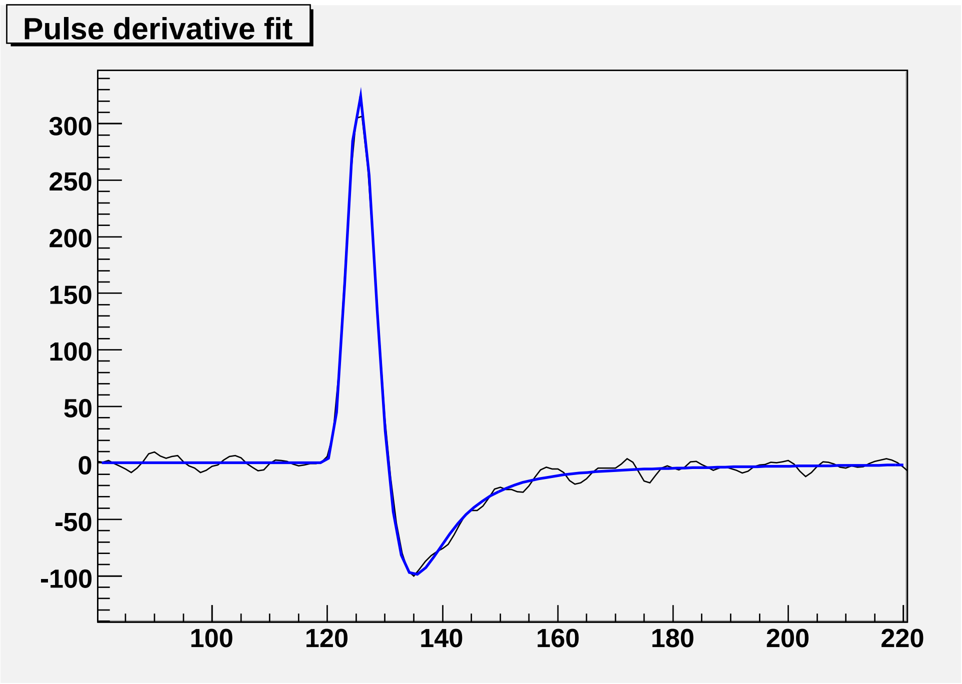

Figure C.4: Fit of the derivative of a pulse with a cubic-spline interpolation of the average pulse derivative. This pulse is from channel 31 of run 737, and the fit determined the offset to be $t_{\text{offset}} = -2.58$ in units of digital samples.

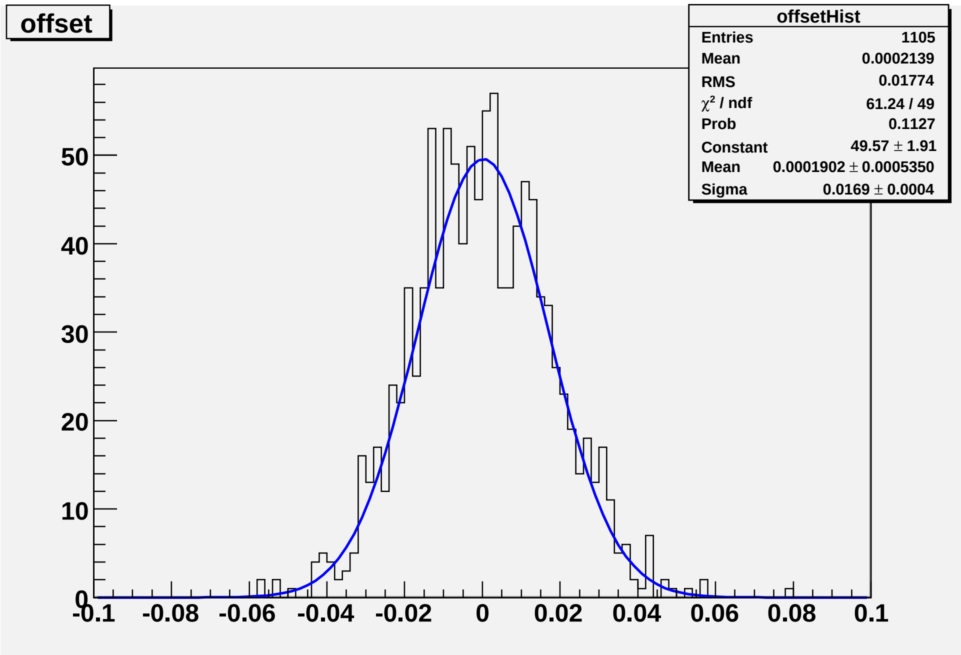

Figure C.5: Histogram of offsets (in units of digital samples) obtained by adding noise to the average pulse scaled to roughly the amplitude of 2.5 MeV events. The data are from channel 1 of run 738. From the sigma of the Gaussian fit, the error on the offset is estimated to be 1.7% of the spacing between digitized values.

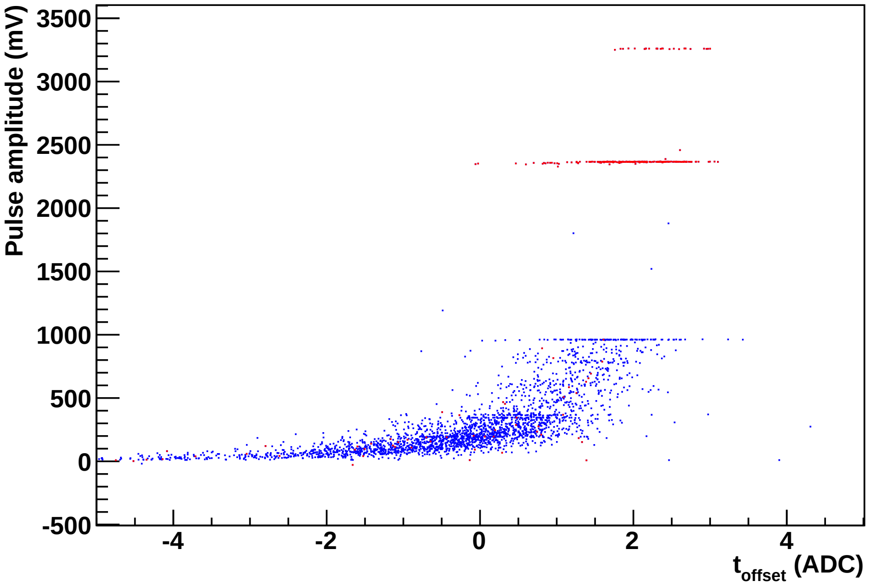

Figure C.6: Pulse amplitude vs. offset for channel 1 of run 737. The blue points are normal triggered events, and the red points are heater events.

C.2.2 Truncating the pulse

To align the pulses based on their offsets, it is necessary to shift the pulses left or right until they are aligned with the average pulse. It is not possible to retain all 512 points after such a shift: If the pulse is shifted left, there are missing points at the right side, and vice versa. Therefore, we truncate the average pulse by cutting off the sides; we remove the first 16 points on the left and the last 16 points on the right. Then, each pulse can be shifted until aligned with the average pulse by truncating the appropriate number of points from the left and from the right. For example, if a pulse has an offset of -2, then 14 points would be removed from the left and 18 from the right, leaving the pulse aligned with the average pulse. For the purpose of truncation, the offset determined from the fit to the derivative is rounded to the nearest integer.

C.2.3 Residual offset

After truncation there remains a residual offset between the truncated pulse and the truncated average pulse equal to the difference between the offset and the rounded offset. To correct for the residual offset, for each pulse the average pulse is resampled according to the residual offset as follows: First, the average pulse is interpolated by a cubic-spline. Since the average pulse should be smooth and noise-free, the cubic-spline interpolation should provide a good continuous representation of the average pulse, at least for the low frequencies we will use to construct the amplitude. Then, the spline-interpolation of the average pulse is resampled to obtain 480 discrete points ($= 512 - 16 - 16$) matching the offset of each pulse. The resampled average pulse for each pulse is used as input to the amplitude algorithm described in the next section.

C.3 Derivation of the amplitude algorithm

We assume that the function, $p(t)$, representing a pulse as a function of time within the acquired window may be represented to a good approximation as

$$p(t) = b + a \cdot s(t - t_0) + n(t), \tag{C.3}$$

where b is a constant baseline level, a is the pulse amplitude, $s(t)$ is the ideal pulse shape, t_0 is the onset time of the pulse – $s(t - t_0) = 0$ for $t < t_0$, and $n(t)$ is the noise. Technically, $s(t)$ should be normalized to unit amplitude. Written in discrete time,

$$p_m = b + a \cdot s_m + n_m, \tag{C.4}$$

where $m = 0, \ldots, N - 1$ with $N = 480$ for a truncated CUORICINO pulse. After the alignment described in the previous section, there is no need to mention t_0 explicitly anymore; the ideal pulse shape, s_m, is taken as the resampled average pulse, which includes the correct t_0 for each pulse.

The amplitude algorithm operates in the frequency domain where the pulse is represented as

$$\tilde{p}_k = \sum_{m=0}^{N-1} p_m \, e^{-2\pi i k m/N}, \tag{C.5}$$

where $k = 0, \ldots, N - 1$. Only the $k = 0$ frequency component, $\tilde{p}_0$, depends on b under the assumption that b is constant for the duration of the pulse. For $k > 0$,

$$\tilde{p}_k = a \cdot \tilde{s}_k + \tilde{n}_k. \tag{C.6}$$

To isolate a, we divide by $\tilde{s}_k$:

$$a = \frac{\tilde{p}_k}{\tilde{s}_k} - \frac{\tilde{n}_k}{\tilde{s}_k}. \tag{C.7}$$

In practice, we do not know $\tilde{n}_k$ and must work with only the first term on the right-hand-side of Eq. (C.7). Averaging over many pulse samples, $\langle \tilde{n}_k / \tilde{s}_k \rangle =$

$\langle \tilde{n}_k \rangle / \tilde{s}_k = 0$ since the phase of $\tilde{n}_k$ is randomly distributed according to a uniform distribution between 0 and 2π. Therefore,

$$\left\langle \frac{\tilde{p}_k}{\tilde{s}_k} \right\rangle = a, \tag{C.8}$$

so $\tilde{p}_k/\tilde{s}_k$ is an unbiased estimator of a for each $k > 0$. Although $\tilde{p}_k/\tilde{s}_k$ averages to a real number, for an individual pulse it will have a small imaginary part (due to the noise). Therefore, we work with the real part, which is clearly also an unbiased estimator of a and which we will denote by c_k to mean the contribution from the k^{th} frequency:

$$c_k \equiv \mathcal{R}e \left(\frac{\tilde{p}_k}{\tilde{s}_k} \right). \tag{C.9}$$

Each c_k is an estimator of the amplitude, a, and we want to form a weighted average of them in such a way that minimizes the variance due to noise. We write the estimator, $\hat{a}$, of the amplitude as a linear combination of the c_k with weights w_k and seek to choose weights that minimize $\mathrm{Var}(\hat{a})$:

$$\hat{a} = \sum_{k=1}^{M} w_k \, c_k, \tag{C.10}$$

where M is some maximum frequency, which could be as high as the Nyquist frequency but is typically chosen smaller to limit the size of the covariance matrix which must be inverted (to be described below). We require that the weights sum to 1 so that $\langle \hat{a} \rangle = a$, where the average is taken over an ensemble of pulses with true amplitude a in the presence of noise. The variance of $\hat{a}$ is exactly given by linear propagation of errors to be, in matrix notation,

$$\mathrm{Var}(\hat{a}) = w^T V w, \tag{C.11}$$

where V is the covariance matrix for the c_k. How to determine V from the data will be discussed in Sect. C.4. Thus, the problem of choosing the weights becomes the problem of minimizing the quadratic form in Eq. (C.11) subject to the constraint

$$\sum_{k=1}^{M} w_k = 1. \tag{C.12}$$

The constraint may be expressed in matrix notation by introducing a column matrix, u, whose elements are all 1: $u_k = 1$ for each k. Then, Eq. (C.12) may be written as

$$u^T w = 1. \tag{C.13}$$

We solve the minimization problem by the method of Lagrange multipliers. Introducing a Lagrange multiplier λ, we minimize

$$w^T V w + \lambda(u^T w - 1). \tag{C.14}$$

The derivative with respect to λ recovers the constraint equation, and differentiating with respect to the M components of w and writing the result in a column matrix, we have

$$\begin{aligned} \partial_w(w^T V w + \lambda u^T w) &= 2Vw + \lambda u && \text{(C.15)} \\ &= 0 && \text{(C.16)} \\ \Rightarrow w &= -\frac{\lambda}{2} V^{-1} u. && \text{(C.17)} \end{aligned}$$

After solving for λ from the constraint equation, we have the result for the weights:

$$w = \frac{V^{-1}u}{u^T V^{-1} u}. \tag{C.18}$$

By substituting this result into Eq. (C.11), we obtain the minimum variance and theoretical amplitude resolution to be:

$$\begin{aligned} \mathrm{Var}(\hat{a}) &= w^T V w && \text{(C.19)} \\ &= \frac{(V^{-1}u)^T V (V^{-1}u)}{(u^T V^{-1} u)^2} && \text{(C.20)} \\ &= \frac{1}{u^T V^{-1} u}, && \text{(C.21)} \\ \sigma_{\hat{a}} &= \frac{1}{\sqrt{u^T V^{-1} u}}. && \text{(C.22)} \end{aligned}$$

C.4 Determining the covariance matrix

The formula for the optimal weights, Eq. (C.18), requires the inverse of the covariance matrix of the c_k. Here we show how to obtain the covariance matrix from the noise samples. Starting with the definition of covariance,

$$\begin{aligned} V_{ij} &= \langle c_i c_j \rangle - \langle c_i \rangle \langle c_j \rangle && \text{(C.23)} \\ &= \left\langle \left(a + \mathcal{R}e\left(\frac{\tilde{n}_i}{\tilde{s}_i}\right)\right) \left(a + \mathcal{R}e\left(\frac{\tilde{n}_j}{\tilde{s}_j}\right)\right) \right\rangle - a^2 && \text{(C.24)} \\ &= \left\langle \mathcal{R}e\left(\frac{\tilde{n}_i}{\tilde{s}_i}\right) \mathcal{R}e\left(\frac{\tilde{n}_j}{\tilde{s}_j}\right) \right\rangle && \text{(C.25)} \\ &= \langle c_i^{\text{noise}} c_j^{\text{noise}} \rangle, && \text{(C.26)} \end{aligned}$$

where the last line contains frequency contributions to the amplitude from pure noise samples. Since the offset is irrelevant for pure noise, it is set to 0 in computing the frequency contributions from the noise samples.

C.5 Comparison with the optimal filter

During the processing of the CUORICINO data, the pulse amplitudes were measured with both the optimal filter and the algorithm described here. The gain stabilization and energy calibration were performed independently for the two sets of pulse amplitude measurements. The energy spectra produced from the optimal filter pulse amplitudes and from the pulse amplitudes computed by this algorithm are shown in Fig. C.7. Comparison of spectra shows no apparent difference in performance between the two techniques. In particular, the energy resolutions are virtually identical. Therefore, the standard optimal filter is used in this analysis. Nevertheless, having two independent applications of the first-level data analysis procedures, starting from the two different pulse amplitude measurements, provides a useful cross-check on the data processing. For example, if small spectral features are consistent between the two energy spectra, it builds confidence that the features are real physical effects and not artifacts such as miscalibration of one channel for one data set.

C.6 Implementation in Diana modules

The algorithm described in this appendix is implemented in the DIANA module MBOBFilter, which employs two classes: QOffset for computing the offset by the method of Sect. C.2 and QBAmplitude for computing the contributions from each frequency by the method of Sect. C.3. The optimal weights are computed by the module MBComputeWeights, which forms the covariance matrix from Eq. (C.26) and computes the weights according to Eq. (C.18).

C.7 Use in rejection of spurious pulses

The quantities computed in the course of generating the amplitudes may also be useful in the rejection of spurious triggered events. The contribution from the k^{th} frequency, c_k, defined in Eq. (C.9) probes the shape of a pulse at the frequency k. For real pulses, c_k has an expectation value (namely, the amplitude of the pulse) and an error (found from Eq. (C.26)). Cuts can be set on each c_k at a chosen number of sigma from their weighted average, $\hat{a}$, or a cut can be set on the χ^2 statistic that measures the compatibility of the c_k with $\hat{a}$.

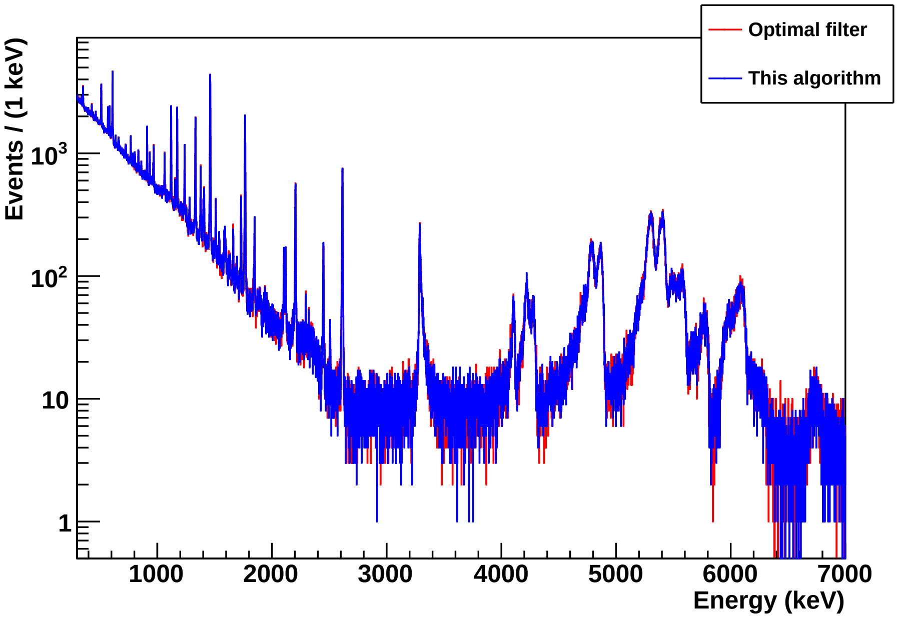

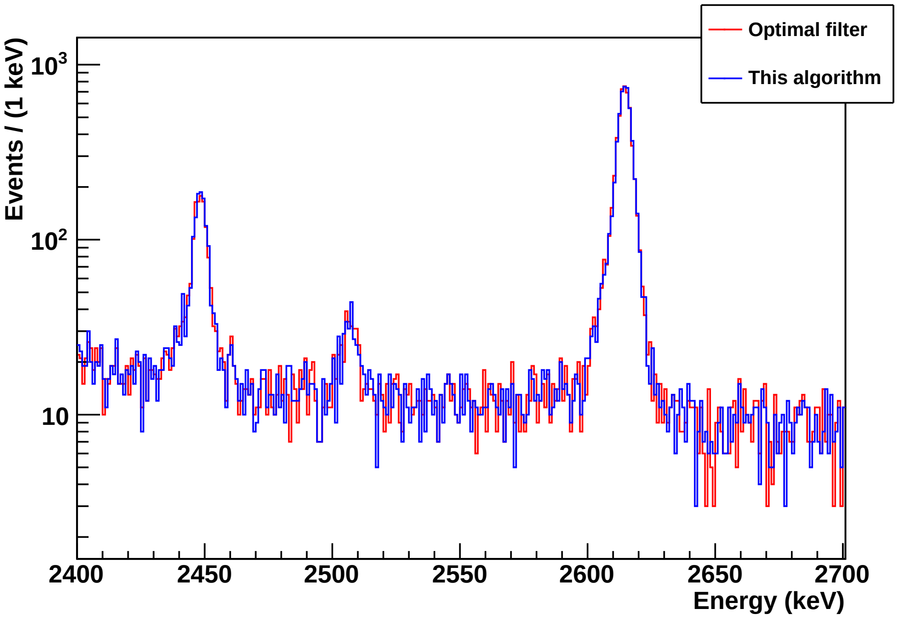

Figure C.7: CUORICINO spectrum obtained from pulse amplitudes measured with the optimal filter (red) and the algorithm described here (blue). The bottom plot is a zoom of the top plot around the $0\nu\beta\beta$ decay region of interest. The calibration function is parameterized as a second-order log polynomial in both cases, although the calibration coefficients are computed separately.