The University Of Sheffield.

Department of Physics & Astronomy

Coincidence-based reconstruction and analysis for remote reactor monitoring with antineutrinos.

Elisabeth Kneale

University of Sheffield & AWE

A thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

October 2021

Abstract

Antineutrinos from a nuclear reactor comprise an unshieldable signal which carries information about the core. Their interaction via inverse beta decay (IBD), in a gadolinium-doped water medium, produces a positron-neutron pair coincident in time and space. NEO (Neutrino Experiment One), a gadolinium-doped water-based Cherenkov detector, is planned to demonstrate remote reactor monitoring for nuclear non-proliferation. Such a monitor would, by its very nature, operate at the limit of its sensitivity, detecting low-energy antineutrinos, from a distance, in a complex nuclear landscape. In this thesis, I investigate ways to achieve optimal sensitivity to a reactor signal.

I implement a reconstruction which uses the combined light from both events in an IBD pair to accurately reconstruct the interaction vertex. I also develop a powerful analysis to reject coincident backgrounds and find the maximal sensitivity of a detector to a reactor signal. Finally, I evaluate the sensitivity of NEO detector designs to reactor signals at the planned site in Boulby Mine.

I find that the combined reconstruction improves vertex resolution for IBD positrons by up to a factor of two. IBD-neutron vertex resolution is found to improve by more than 20% at most energies. The coincidence analysis shows that a 22 m gadolinium-doped water-based liquid scintillator detector would be sensitive to the signal from the Heysham 2 nuclear reactor complex 149 km away within 22 months and could detect the dual-site signal from the reactors at the Heysham 2 and Torness complexes in just over 6 months.

A detector at Boulby Mine could detect signals from reactor complexes within the time constraints of the current schedule for reactor decommissioning in the UK. As a remote reactor monitoring prototype, NEO would be the first ever neutrino physics application.

Declaration

I, the author, confirm that the Thesis is my own work. I am aware of the University's Guidance on the Use of Unfair Means (www.sheffield.ac.uk/ssid/unfair-means). This work has not been previously been presented for an award at this, or any other, university.

Working within an international particle physics collaboration such as WATCHMAN, my work has inevitably been enmeshed with that of others. Chapters 5 to 7 present my original work but I would like to take this opportunity to clarify how the previous work of others influenced the development of the material presented there. Chapter 5 details the first use of a combined reconstruction for WATCHMAN. This constitutes an implementation and optimisation of the BONSAI combined reconstruction code developed by Michael Smy. Chapter 6 describes the coincidence analysis - Cobraa - that I have developed for WATCHMAN. This work began as part of the Watchmakers code. Watchmakers, to which I have also made a number of contributions, was originally written by Marc Bergevin. Cobraa eventually developed into a wholly separate analysis, founded on a partial modification of the simulation process and an alternative method of signal and background evaluation. Cobraa is now used by members of the WATCHMAN Collaboration as a separate analysis. Nevertheless, due to the nature of its development, Cobraa draws many of its structural features from Watchmakers. Chapter 7 presents the results of the new simulation-reconstruction-analysis toolchain.

Research carried out under the supervision of

Dr. Matthew Malek, University of Sheffield
Dr. Jon Burns & Mr Steve Quillin, AWE

and supported by EPSRC and AWE.

Acknowledgements

First of all, I would like to thank my supervisors - both official and unofficial.

My somewhat unusual route both to and through my PhD has, I believe, been a learning experience for both myself and my supervisor, Matthew Malek. From the start, Matthew has been quick to adapt to - and make allowance for - my particular (logistical) limitations. At the same time, he has encouraged me to work beyond the level at which I think I am capable of working and has valued my perseverance. He has also unerringly promoted my work within the WATCHMAN collaboration and beyond.

A reflection on my PhD journey would not be complete without reference to the Covid-19 pandemic and the exceptional challenges that it has brought with it. Covid-19, home schooling and PhD research do not mix well. Matthew's primary concern was my health and well-being - and his reassurance that all was not lost was a lifeline for me during that time.

It took me half a lifetime to 'find' physics and I am grateful to Matthew for giving me the opportunity to pursue it further, for valuing my work and for accommodating my circumstances with understanding and consideration.

I was lucky enough to also have two supervisors based at AWE over the course of this PhD - Jon Burns and Steve Quillin. When I met Jon Burns, I was impressed with his calm, straight-talking rationality. My opinion has not changed. He has been a reassuring and steadying presence over the last four years. I first worked closely with Steve Quillin on the acceptance testing for the initial order of PMTs for WATCHMAN. Then and since, I have at regular intervals valued his sharp mind and astute questioning - not least in my final stages of producing this thesis. That, and his appreciation of festivals and good beer.

My research work really started the year before my PhD began, when I worked on an early version of some of Jost Migenda's *sntools* supernova event-generation code. Jost is insightful and precise and has patience by the bucket load (which they needed at times!). I miss our chats over lunch in the Austin Room. On my first day working with them, Jost told me, "You will have a steep learning curve". You weren't wrong Jost - but I got there eventually, with your help.

The early days of my PhD were punctuated with frustration over new software and a sometimes overwhelming amount of knowledge to gain. Marc Bergevin was the person that I turned to for help and since then he has always been on hand to offer assistance, ideas, opportunities and enthusiasm. Marc, I've said this before and I'll say it again: your guidance has been vital to me throughout the past four years.

I would like to thank my parents (all of you!) for the practical and emotional help that you have given me. Your appreciation of what I do means a great deal. Rod, thank you for the stories. Lesley, thank you for the creativity. Mum, thanks for being an example and showing that a woman can come to a career later in life - and be successful. Dad, thank you for giving me the benefit of your experience from your brilliant academic career - if I can achieve a fraction of what you have achieved I will be happy.

Last and certainly not least, I want to thank my husband. Bob, you have never questioned my decision (or my sanity!) to go back and start my career again from the beginning. Thank you for being you, and for supporting me through it all - in many different ways. You can now rest assured - I am the eternal student no more.

To Bobby, Martha and Jude.

You are all amazing in your own, very different ways.

It was the best of times,

It was the worst of times,

It was the age of wisdom,

It was the age of foolishness,

It was the epoch of belief,

It was the epoch of incredulity,

It was the season of light,

It was the season of Darkness,

It was the spring of hope,

It was the winter of despair...

Charles Dickens, A Tale of a PhD

Contents

List of Figures

List of Tables

Neutrinos and Non-Proliferation

> *All things great are wound up with all things little.*
>
> – L. M. Montgomery *Anne of Green Gables*

There is a paradox that exists in the expansion of nuclear power globally.

Nuclear power can help meet the so-called *base-load* demand for electricity without the vilified carbon output of fossil fuels and the much-quoted instability of clean, renewable energies. As such, nuclear power is promoted worldwide as a clean source of energy[1]. The UK Government announced in its 2020 Energy White Paper [1] its intention to incorporate nuclear power in its array of *clean energy technologies* to replace fossil fuels over the next thirty years - and the UK government is not alone. The International Atomic Energy Agency (IAEA) projects that the global nuclear power capacity could increase by 80% to 715 gigawatt electrical (GWe) by 2050, up from 392 GWe in 2019 [2].

The paradox lies in the fact that while global energy security increases thanks to nuclear power, so too does the global risk posed by nuclear weapons.

Plutonium is a by-product of nuclear power generation due to the irradiation of ^{238}U by the abundant neutrons in a reactor core and it can be extracted by reprocessing, a chemical process by which the spent fuel is broken down into its constituent parts. Extracted plutonium may then be diverted for use in nuclear weapons. For this reason, nuclear reactors require careful monitoring. Enter the IAEA and the Non-Proliferation Treaty.

Non-Proliferation Treaties have a threefold focus - the peaceful use of nuclear energy, non-proliferation and disarmament [3] - and are supported by IAEA Comprehensive Safeguards Agreements. Where it is deemed necessary, Additional Protocols strengthen the safeguards agreements to allow careful monitoring of reactors and spent reactor fuel at a nuclear facility [4]. However, it can be difficult to balance resistance to the increased intrusion with confidence in adherence to the treaties.

In response to this perceived difficulty, the National Nuclear Security Administration (NNSA) of the U.S. Department of Energy published its Strategic Plan in 2011 [5], which called for the demonstration of remote reactor monitoring for nuclear non-proliferation. A remote monitor may be used as part of a remote monitoring toolkit alongside other monitoring techniques for verification of declared reactor activity (mid-field or wide-area monitoring) or for the detection of a hidden reactor (far-field monitoring).

WATCHMAN, a WATer CHerenkov Monitor for ANtineutrinos, is the scientific collaboration charged with demonstrating antineutrino technology with a water-based detector applied to remote monitoring for nuclear non-proliferation. Use of remote reactor antineutrino detection in the field would be the first ever neutrino physics application.

[1] Whether or not nuclear power can be classed as clean remains a topic of impassioned debate but that discussion will have to be saved for another time and place.

Overview of this Work

The challenge of remote reactor monitoring is that the low-energy antineutrinos from a reactor are at the very limit of the energy threshold in water Cherenkov detectors[2]. Backgrounds (e.g. from other reactors and natural radioactivity of detector components) further reduce sensitivity to the signal. This issue requires a multi-faceted approach to improving the sensitivity. With that in mind, this thesis focuses on improving event reconstruction and signal-to-background discrimination for optimal sensitivity to a remote reactor. This has been achieved with the implementation of a new reconstruction algorithm and a new post-reconstruction analysis. The simulation-reconstruction-analysis toolchain in this thesis presents the most comprehensive treatment to date of reactor antineutrino detection for WATCHMAN. The improved reconstruction and analysis have been used to evaluate the sensitivity of several detector configurations to a number of reactor signals [6] [7]. A selection of these results are presented in this thesis.

The following chapter details the history and development of neutrino physics. Chapter 2 covers reactor antineutrinos and the application of antineutrino detection to non-proliferation. Chapter 3 describes WATCHMAN's detector in its currently planned form and the signal and backgrounds expected in the detector are explained in Chapter 4. The new reconstruction is described in Chapter 5 and the sensitivity analysis is detailed in Chapter 6. In Chapter 7, the new reconstruction and analysis methods are applied to different discovery scenarios to evaluate and compare the sensitivity of proposed designs to a number of reactor signals. Chapter 8 summarises the findings and discusses possible extensions to the methods presented in this thesis.

[2] This technology will be explained in detail later.

1 Neutrinos - a History

"Begin at the beginning," the King said, gravely, "and go on till you come to an end; then stop."

– Lewis Carroll, *Alice in Wonderland*

1.1 A Desperate Remedy - Pauli's Postulated Particle and other theories.

The year was 1930. Sixteen years had passed since James Chadwick had shown that electrons from beta decay were, contrary to expectation, emitted in a continuous spectrum. In an attempt to solve the mystery of the 'missing' energy, Wolfgang Pauli, despite his better judgement, "hit upon a desperate remedy" [8] - a new particle, which he called the *neutron*. This new fermion would have to be electrically neutral with spin $\frac{1}{2}$ to conserve charge and angular momentum. It would also have very little (but not necessarily zero) mass and - crucially - be weakly interacting and therefore impossible to detect [3].

Following Chadwick's discovery of what we know today as the *neutron* in 1932, which was clearly too massive to be Pauli's neutron, Enrico Fermi adopted Edoardo Amaldi's new name for Pauli's postulated particle - the *neutrino*, or 'little neutral one'. Albeit after a shaky start[4], Fermi solidified the existence of the neutrino in the minds of physicists via its integral role in his theory of beta decay or Fermi theory [9], the first theory of weak interactions.

Two decades of confusion over the weak interaction ensued, punctuated by moments of experimental clarity (see section 1.2), until the proposal [10] and discovery [11] of parity violation in the weak interaction. In her test of parity conservation in the beta decay of ^{60}Co, Chien-Shiung Wu found that most of the electrons were emitted in the direction opposite to the nuclear spin vector. Since, if parity were conserved, the electron would have been produced in the same and opposite direction to the nuclear spin in equal proportions, parity conservation was overthrown. To the relief of many a physicist, anarchy was avoided when the complex of symmetries and asymmetries at work in weak interactions were distilled into the V - A theory [12] [13] [14] soon afterwards in 1958. The vector current of the Fermi theory was split into a combination of vector (V) currents, whose spatial components change sign under parity transformation, and axial-vector (A) currents, whose spatial components do not change sign under parity transformation. Parity violation arises from the interference between the two.

That same year, the neutrino was confirmed to be left-handed - its direction of spin is opposite to its direction of motion. By finding that the helicity of the photon in electron capture on Europium was negative, it was deduced that the neutrino also has negative helicity and by extension (since the neutrino has very little mass) is left-handed [15]. The combination of the V and A currents in the V - A theory leads to a charged weak interaction that only couples to left-handed chirality particles or right-handed anti-particles. Since neutrinos are only created by charged weak interactions, neutrinos must be left-handed and antineutrinos must be

[3] Or so, it is rumoured, Pauli thought.

[4] It is highly likely that *Nature* later regretted rejecting Fermi's *Tentative theory of β radiation* paper on the grounds that "it contained speculations too remote from reality to be of interest to the reader".

right-handed. So left-handed neutrinos and right-handed antineutrinos are a result of the interactions described by the V - A theory and vice versa. The development of the V - A theory into the description of neutrino interactions in the Standard Model of weak interactions in the Glashow-Weinberg-Salam Standard Model (see section 1.3) heralded a (momentary) end to the theoretical confusion.

So the neutrino was well-embodied in theory by this time and yet, due to its elusive nature, it was only a couple of years earlier that physicists had first managed to detect the neutrino. The remainder of this section will take us through the history of neutrino detection and the uncovering of some of what we now know about neutrinos.

1.2 Ghost Hunting

Pauli is rumoured to have said, "I have done a terrible thing, I have postulated a particle that cannot be detected." If he really did say that, however, he was mistaken.

To be fair, it is no easy matter to detect a particle which can pass through matter so freely as to earn it the title "ghost particle". Nevertheless, Frederick Reines and Clyde Cowan set out in 1953 to prove the existence of neutrinos.

Neutrino interactions observed via outgoing particles

The development of neutrino detection began with the observation of antineutrinos from a nuclear reactor. Electron antineutrinos are emitted during β-decay in a nuclear reactor or explosion. After first pursuing the idea of using a controlled nuclear explosion to produce a sufficiently high antineutrino flux as to be detectable, Reines and Cowan were thankfully able to design an alternative and rather safer experiment, which made it possible to detect the lower flux of antineutrinos from a fission reactor. They achieved their goal - detecting antineutrinos from the Savannah River nuclear reactor [16] [17] with two tanks of cadmium-doped water. Electron antineutrinos were detected by their inverse beta decay (IBD) interaction with protons in the water:

$$\bar{\nu}_e + p \rightarrow n + e^+.$$

This remains the principal reaction by which electron antineutrinos are observed today (see section 2.4.1 for further details). Subsequent positron-electron annihilation as well as neutron capture on cadmium resulted in a coincident[5] signal pair detected via scintillation light in three outer tanks of liquid scintillator instrumented with photomultiplier tubes. Reines scooped the Nobel Prize for the discovery in 1995.

Neutrinos and antineutrinos - distinct particles

The theory of lepton-number conservation was first proposed by Emil Konopinski and Hormoz Mahmoud [18] in 1953 and was subsequently incorporated into the V - A theory and Standard

[5] The term *coincident* will be used for the remainder of this thesis to refer to events which occur within a short distance and time of each other.

Model of Weak Interactions (see Section 1.3). Later that decade, Ray Davis attempted to observe the lepton-number-violating interaction of electron antineutrinos with ^{37}Cl:

$$\bar{\nu}_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^-.$$

His aim was to determine whether or not the neutrino and the antineutrino are distinct particles with opposite lepton numbers.

Davis, who was later to go on to conduct the famous Homestake experiment detecting neutrinos from the Sun, tried to observe antineutrinos from from the Brookhaven High Flux Research reactor [19] and Savannah River reactor [20] using a detector filled with 3800 litres of carbon tetrachloride. Instead, he was able to put an upper limit on the interaction cross section, which was a factor of 20 lower than the theoretical cross section for the equivalent neutrino interaction with ^{37}Cl, at 10^{-45}cm^2 [19].

This was one of the first hints that neutrinos and antineutrinos should be treated as distinct particles - and more variations were yet to be added to the neutrino family.

Three *flavours* of neutrino

The next triumph in neutrino detection came in 1962 with the demonstration of the existence of muon neutrinos (ν_μ) as particles distinct from electron neutrinos (ν_e). This was the first ever accelerator neutrino experiment, in which the Alternating Gradient Synchrotron at Brookhaven National Laboratory was used to produce a beam of protons which were fired at a beryllium target, producing large numbers of pions. This beam then passed through a 13.5 m-thick iron wall to filter out backgrounds. The remaining pions decayed into a muon plus a neutrino and the resulting neutrinos passed into a nearby 10-tonne spark chamber, where they would interact with protons in aluminium plates to produce either an electron or a muon. These could be identified from sparks produced as they ionised the neon fill between the plates, which were held under high voltage. If the ν_μ and ν_e were the same particle, the muons and electrons would have been produced in equal numbers. In fact, thirty four muons (5 of which, however, were considered to be cosmogenic in origin) and only six electron candidates (identified as showering events) were produced from neutrino interactions, leading to the conclusion that the ν_μ and ν_e are not the same particle [21]. In 1988 Jack Steinberger, Leon Lederman, and Mel Schwartz won the Nobel Prize for this discovery.

A third generation of neutrino, the tau neutrino (ν_τ), was shown to exist when it was detected in 2001 by the DONUT (Direct Observation of the NU Tau) collaboration [22]. The experiment was designed following the discovery of the τ lepton in 1975 by a group led by Martin Perl (who shared the Nobel Prize with Reines in 1995) at the Stanford Linear Accelerator Center (SLAC). After this, it was generally assumed that there must be a third generation of neutrino. The DONUT experiment was ultimately able to demonstrate this.

In the DONUT experiment, an 800 GeV beam of protons from the TeVatron at Fermilab was fired at a tungsten beam dump to produce charm D_s mesons (consisting of a charm quark and strange antiquark). These would then produce τ antineutrinos as they decayed into τ leptons and $\bar{\nu}_\tau$. The τ leptons decayed in turn into ν_τ (plus other particles depending on the specifics of the decay). The ν_τ then passed through 36m of shielding (to reduce backgrounds) into sheets of

emulsion targets which consisted of light-sensitive silver bromide crystals suspended in gelatin. Occasionally, the ν_τ would interact with a nucleon in the detector to produce charged particles, which left visible tracks in the emulsion. When the light-sensitive emulsion sheets were developed, the tracks of particles could be reconstructed by connecting the marks they left on adjacent sheets of emulsion. Neutrino interactions were identified by the appearance of several outgoing tracks without a preceding track. A tau neutrino was discriminated from other neutrinos by a kink, after a few mm, in one of the several tracks - indicating the decay of a τ lepton and thus the presence (and indeed existence) of a τ neutrino.

Around the same time, a limit was placed at three generations of light neutrino when the four Large Electron-Positron (LEP) collider experiments at CERN determined that there can be no more species of light neutrinos [23]. By measuring the invisible width (attributed to neutrino interactions) of the Z boson, it was deduced that there can be only three species of neutrinos with masses of less than half the Z-boson mass. This direct measurement was borne out also by experiments at Stanford and corroborated the number of neutrino species from cosmological calculations based on the composition of elements in the Universe which is affected by the number of neutrino species involved in the primordial nucleosynthesis [24].

And so the neutrino picture was complete for the moment: three generations, a left-handed neutrino for each of the electron, muon and tau leptons and their associated right-handed antiparticles, all distinct particles and each generation adhering to lepton number conservation. That is not to say this is the end of the road. Neutrinos continue to present puzzles for theorists but for now they have at least for the most part been enshrined in the Standard Model of Electroweak Interactions - an incongruous name for a model so extraordinary and all-encompassing (*almost!* - see section 1.4).

1.3 Neutrinos in the Standard Model

Neutrino interactions are described by the Standard Model of Electroweak Interactions, which has been incorporated into the Standard Model. The Standard Model of particle physics describes the strong, electromagnetic and weak interactions of elementary particles - fermions - mediated by gauge bosons. The model has been superlatively successful in explaining what has been observed experimentally.

Electromagnetic and weak interactions are mediated by the massive W^+, W^- and Z bosons and the massless γ (photon). Although neutrinos interact only through the weak force, the electromagnetic and weak interactions cannot be treated independently (unlike the strong interaction). These interactions are therefore unified into the Standard Model of Electroweak Interactions.

The precursor to the Standard Model of Electroweak Interactions was the Glashow Model proposed in 1961 by Sheldon Glashow [25]. This model incorporated the parity-violating V - A currents and predicted the existence of the Z vector boson and weak neutral currents.

The Glashow-Weinberg-Salam (GWS) Standard Model of Electroweak Interactions was adapted from the Glashow Model in 1967 to incorporate the Higgs mechanism [26] [27] [28] by Steven Weinberg [29] and Abdus Salam [30]. It was later combined with the theories of

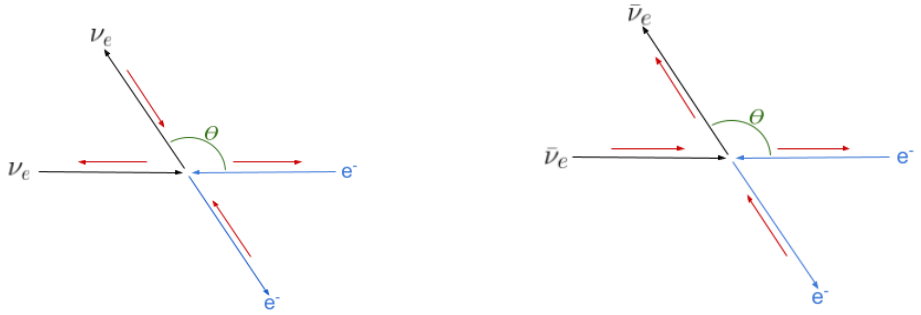

(a) Helicity of neutrino-electron scattering in the massless limit

(b) Helicity of antineutrino-electron scattering in the massless limit

Figure 1: Helicity in neutrino- and antineutrino-electron scattering, neglecting neutrino mass, in the centre-of-mass frame. Helicity is indicated by the shorter, red arrows. In neutrino-electron scattering (a), the left-handed chiral and helicity neutrino and electron have an initial total spin $J_z = 0$ along the interaction. If the scattering angle $\theta = \pi$ then the final total spin is still zero. By contrast, in antineutrino-electron scattering (b) a right-handed antineutrino and left-handed electron have an initial total spin $J_z = +1$. A scattering angle $\theta = \pi$ is therefore forbidden since the final total spin would be -1.

Quantum Electrodynamics and Quantum Chromodynamics to form the Standard Model as we know it today.

In the Standard Model, weak interactions violate parity symmetry, where parity transformation is the sign reversal of spatial co-ordinates. In the Standard Model of Electroweak Interactions, the weak V - A current is also the left-handed chiral projection operator. This means that only left-handed chiral leptons and right-handed chiral antileptons can take part in charged current weak interactions. Neutrinos are left-handed chiral (as we know from the Wu experiment - see Section 1.1) and because they are almost massless, almost all must have left-handed helicity (the frame-dependent version of the Lorentz-invariant chirality). Conversely, antineutrinos are right-handed with right-handed helicity. The chiral property of the weak interaction leads to differences in neutrino and antineutrino cross sections. For example, the cross section of the electron neutrino charged-current scattering on electrons is three times greater than the antineutrino equivalent because the conservation of total spin along the interaction axis places a limit on the scattering angles in the antineutrino case (see Figure 1).

Charge conjugation transformation reverses the charges associated with fundamental forces and charge-conjugation symmetry is also violated in the weak interaction. For particles other than the neutrino, charge conjugation is responsible for particle-antiparticle transformation. However, under charge conjugation, a left-handed neutrino becomes a left-handed antineutrino and since left-handed antineutrinos have never been observed, charge conjugation symmetry is violated. Instead, CP symmetry is responsible for the neutrino-antineutrino transformation. Taking a combination of the charge conjugation and parity transformations - CP transformation - a left-handed neutrino becomes a right-handed antineutrino.

Nevertheless, CP violation has already been observed in the quark sector and experiments such

as T2K (Tokai to Kamioka) in Japan are currently searching for evidence of CP violation in the oscillations of neutrinos and antineutrinos (see Section 1.4.2.

The Standard Model as a whole has been well-corroborated by experiment. The GWS Standard Model's prediction of the weak neutral current was confirmed in 1973 by the Gargamelle bubble-chamber experiment at CERN, with the discovery of the neutral-current muon neutrino and electron scattering interaction [31] and the equivalent interaction with a proton or neutron [32]. The model was further confirmed with the discovery of the W and Z bosons by the UA1 and UA2 experiments respectively at CERN's Super Proton Synchrotron (SPS) proton-antiproton collider in 1983 [33] [34] [35]. Then, in a moment of triumph for the Standard Model, the predicted Higgs boson was discovered at the ATLAS and CMS experiments at CERN's Large Hadron Collider in 2012 [36] [37].

The Standard Model of particle physics has been remarkably successful in describing the interactions we observe experimentally. There remain, however, some questions. The Standard Model does not explain the existence of the three generations of quarks and leptons, nor does it include gravity. That said, until the recent Muon g-2 combined experimental result, which evaluates the difference between experiment and Standard Model theory at 4.2 standard deviations [38], only neutrino oscillation experiments had definitively indicated interactions beyond the Standard Model.

Neutrino experiments have shown conclusively that neutrinos can change from one flavour to another (oscillate) as they travel. By extension, this proved that neutrinos have mass, while the Standard Model requires that neutrinos have no mass. To couple to the Higgs boson, a particle must have both a left-handed and right-handed field. Since the neutrino has no right-handed field, the Standard Model neutrino is massless.

The fact that neutrinos oscillate and therefore have mass raises the question of how neutrinos acquire their mass and whether neutrinos are *Dirac* or *Majorana* particles. Massive neutrinos are Dirac particles if the total lepton number L is conserved. In this case the neutrino and antineutrino of a particular flavour have the same mass and different lepton numbers (+1 and -1 respectively) - the neutrino and antiparticle are independent, distinct particles. Neutrinos are Majorana particles if there is no lepton number conservation - neutrinos are their own antiparticles. These are questions which are not answered by the Standard Model.

This disparity between theory and experiment has highlighted the fact that the Standard Model does not tell the whole story. At the same time, the proof that neutrinos oscillate resulted in a clear resolution of the Solar Neutrino Problem and Atmospheric Neutrino Anomaly. This will be the subject of the next section.

1.4 Neutrino Mass Matters

Neutrino physicist Professor David Wark once said that his entire career could be summed up in the word 'almost'. At the beginning of his career in neutrino physics, neutrinos were thought to have no mass, as predicted by the Standard Model. However, in 2015, Art McDonald and Takaaki Kajita were awarded the Nobel Prize for proving that neutrinos *do* have mass. *Almost* no mass - but mass all the same. In fact, they had simultaneously demonstrated that neutrinos

had mass, proved that neutrinos 'oscillate' and solved the solar/atmospheric neutrino problems, since these were all interlinked.

The idea of neutrino oscillations was first proposed by Bruno Pontecorvo in 1957. By the end of the 1990s it was becoming more and more widely acknowledged as a possible explanation for the solar-neutrino and atmospheric-neutrino deficits which had been observed (section 1.4.1). In his theory of neutrino oscillations, Pontecorvo had shown that if different flavours of neutrino exist then oscillation between the different flavours in a vacuum may be possible [39]. Vacuum oscillations lead to exactly the deficits observed in the atmospheric neutrino flux on Earth. Similarly, the theory of neutrino oscillations extended to matter would lead to predictions of precisely the observed solar flux.

1.4.1 The Mystery of the Disappearing Neutrinos

The Solar Neutrino Problem and Atmospheric Neutrino Anomaly were both symptoms of the same issue. Somehow neutrinos were disappearing on their way from their point of creation to the detectors on Earth - this much was established by a number of solar and atmospheric neutrino experiments. Ultimately both problems would be resolved by the theory of neutrino oscillations which was clearly demonstrated by experiment.

Solar Neutrino Problem

The Solar Neutrino Problem was first indicated by the results of the Homestake experiment in the 1960s. In a puzzling result, over a period of 20 years, the experiment saw far fewer neutrinos from the Sun than predicted by John Bahcall [40] and Davis [41] and the Solar Neutrino Problem was born.

The aim of the Homestake experiment was to observe the interaction

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

of solar electron neutrinos with chlorine. The Homestake detector, located around 1.5 km underground in the Homestake Mine in South Dakota, was a 390,000-litre tank filled with perchloroethylene (C_2Cl_4), a dry cleaning fluid and a good source of chlorine. From the solar neutrino flux calculated by Bahcall, Davis predicted that between 4 and 11 solar neutrinos per day would interact with the chlorine in the tank, which would in turn generate argon atoms an order of magnitude above the anticipated argon background [41]. These could then be counted using the method Davis had developed to extract the argon by purging the perchloroethylene with helium, adsorbing the argon in a carbon trap. The quantity of argon present was measured by counting the argon atom decays with a proportional counter, which enabled identification of argon decays by measuring the energy of incident radiation.

The result of around 20 years of data-taking was that the experiment in fact saw only around a third of the expected neutrinos from the sun: a capture rate of $2.56 \pm 0.16(\mathrm{stat}) \pm 0.16(\mathrm{syst})$ SNU[6] [42] compared to the prediction of the standard solar model of $7.6^{+1.3}_{-1.1}$SNU [43]. After ruling out theoretical and experimental error, the consensus was that something was

[6] Solar Neutrino Units - 10^{-36} captures per target atom per second.

happening to two-thirds of the neutrinos from the Sun. This discrepancy became known as the Solar Neutrino Problem.

Similar deficits were confirmed by other experiments. The Super-Kamiokande (Super-K) water Cherenkov detector looking at solar ^{8}B neutrinos confirmed the existence of a solar neutrino deficit [44], observing an interaction rate of ~0.45$\pm$0.02 SNU - just 45% of the predicted rate. Importantly, Super-K observed solar neutrinos via the elastic scattering reaction

$$\nu + e^- \rightarrow \nu + e^-$$

and was able to reconstruct the incoming neutrinos' directions from the directions of the outgoing electrons, which are largely forward-scattered. In this way, they were able to confirm that the neutrinos they observed did indeed come from the Sun.

The GALLEX, GNO, and SAGE gallium experiments had lower energy thresholds and were able to detect ^{8}B, ^{7}B and pp-chain solar neutrinos via the

$$\nu_e + {}^{71}\mathrm{Ga} \rightarrow {}^{71}\mathrm{Ga} + e^-$$

reaction [45] [46] [47]. They also confirmed the solar neutrino deficit.

Finally, the Sudbury Neutrino Observatory (SNO) published results in 2001 [48] which, combined with the existing results from Super-K, heralded an end to the Solar Neutrino Problem. The final piece of the puzzle came in 2002 when SNO published results of a solar neutrino flux measurement which was - crucially - independent of flavour.

The SNO detector was built deep underground and consisted of a 12-metre diameter acrylic sphere containing 1 kilotonne of heavy water (D_2O), suspended in a water-filled cavern for both buoyancy and shielding from backgrounds. Solar neutrinos were detected via three different interactions:

- Charged current interaction with a neutron in the deuterium (electron neutrino only)

$$\nu_e + D \rightarrow 2p + e^-,$$

- Elastic scattering of ν_e, ν_μ or ν_τ on electrons

$$\nu + e^- \rightarrow \nu + e^-,$$

- Neutral current interaction of ν_e, ν_μ or ν_τ breaking the deuterium into its constituent parts

$$\nu + D \rightarrow p + n + \nu.$$

Thanks to the overlap of these interactions, the SNO experiment was able to deduce the fluxes of each of the three neutrino flavours individually. Strikingly, the ν_e flux was found to be one third of the total flux [49]. Since the Sun emits only electron neutrinos, it was clear that two-thirds of the neutrinos had changed flavour by the time they reached Earth. The 'missing' neutrinos had been found and the observations of the other experiments were correct - there was indeed a deficit in the ν_e flux.

Atmospheric Neutrino Anomaly

The Atmospheric Neutrino Anomaly, discovered later - in the Kamiokande detector in Japan - was resolved in a similarly conclusive feat of experimentation. Atmospheric neutrinos are principally produced when cosmic rays interact with nuclei in our upper atmosphere to produce pions, which then rapidly decay in the following two ways:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad \text{then} \quad \mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$$

or

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \quad \text{then} \quad \mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu.$$

Around two-thirds of atmospheric neutrinos are muon neutrinos and muon antineutrinos, and the remainder are electron neutrinos and electron antineutrinos. However, there was found to be a deficit in the muon neutrino flux [50] and more specifically in the upward-going muon neutrino flux [51]. In 1998, Super-K published its *Evidence for Oscillation of Atmospheric Neutrinos* showing a deficit of muon neutrinos which was dependent on the direction with respect to the Earth's surface. This meant that it depended on the distance the muon neutrino had travelled. This was a confirmed deficit in the upward-going ν_μ flux which was not seen in the ν_e flux, nor in the downward-going ν_μ flux and it was consistent with two-flavour vacuum oscillations (see section 1.4) into tau neutrinos at the 90% confidence level [52].

The Solar Neutrino Problem and Atmospheric Neutrino Anomaly were ultimately resolved in a clean sweep by the theory of neutrino oscillations in a three-neutrino mixing model, the experimental proof of which, at Super-K and SNO, led to another Nobel Prize in neutrino physics.

1.4.2 Vacuum Oscillations (and the solution to the Atmospheric Neutrino Anomaly)

Five years after Pontecorvo's proposal of neutrino oscillations in a vacuum, and following the discovery at Brookhaven of the muon neutrino in 1962 [21], Ziro Maki, Masami Nakagawa and Shoichi Sakata proposed a model of two-neutrino mixing between neutrino flavour eigenstates ν_e and ν_μ and neutrino mass eigenstates ν_1 and ν_2 [53]. This was subsequently extended to a three-neutrino mixing model [54] to take into account the ν_τ flavour eigenstate and the ν_3 mass eigenstate. Resulting oscillations, which affect neutrinos as they propagate in a vacuum, account for the difference between the upward-going and downward-going muon neutrino fluxes observed in the Super-K atmospheric neutrino experiment.

Oscillations occur because of the mixing between the flavour and mass eigenstates. In the model, a neutrino is a state with a definite electron, muon or tau flavour or *flavour eigenstate* that is produced in a weak interaction. However, the flavour eigenstates are not exactly the same as the states of definite mass or *mass eigenstates*. A flavour eigenstate contains elements of the other two flavours but destructive interference means that we do not see the other flavours. However, if the masses are not all the same, then this results in a phase difference between different flavour states and so it is possible that the flavour eigenstate observed at a later time will be different to the original flavour.

Because we cannot measure the mass state in which the neutrino was created, we assume a coherent superposition of mass eigenstates and as such express the flavour eigenstates

($\nu_x, x = e, \mu, \tau$) as a linear combination of the mass eigenstates ($\nu_i, i = 1, 2, 3$):

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} = U_{PMNS} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \tag{1.1}$$

In this expression, U_{PMNS} is the PMNS (Pontecorvo-Maki-Nakagawa-Sakata) mixing matrix, which describes the mixing between the three mass eigenstates. It is composed of three rotation matrices with mixing angles θ_{12}, θ_{13} and θ_{23} between the orthogonal mass eigenstates denoted by the subscript and a complex phase δ_{cp}, which is responsible for charge-parity (CP) violation. An additional sub matrix includes contributions from the *Majorana phases* $e^{i\beta}$ and $e^{i\gamma}$ to describe a neutrino which is a Majorana particle (one which is its own antiparticle). The standard parameterisation of the PMNS matrix is:

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{cp}} \\ 0 & 1 & 0 \\ -s_{13}e^{\delta_{cp}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\beta} & 0 & 0 \\ 0 & e^{i\gamma} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.2}$$

where s_{ij} is $sin\theta_{ij}$ and c_{ij} is $cos\theta_{ij}$.

The on-axis Majorana phases do not enter into the neutrino oscillation probabilities. These contributions can therefore be ignored for the purposes of this discussion and the standard PMNS unitary mixing matrix for Dirac particles (particles with distinct antiparticles) becomes:

$$U_{PMNS} = \begin{pmatrix} c_{12}c_{13} & c_{13}s_{12} & s_{13}e^{-i\delta_{cp}} \\ -c_{23}s_{12} - c_{12}s_{13}s_{23}e^{i\delta_{cp}} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta_{cp}} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - c_{23}s_{12}s_{13}e^{i\delta_{cp}} & c_{13}c_{23} \end{pmatrix}. \tag{1.3}$$

The oscillation arises from the coherent superposition of mass states. The wavepackets of the mass states can be approximated as plane waves, which therefore propagate according to the Schrödinger equation:

$$i\frac{d}{dt}|\nu_j(t)\rangle = \hat{\mathcal{H}}|\nu_j(t)\rangle, \tag{1.4}$$

where $j = 1, 2, 3$.

For a neutrino of flavour $\alpha = e, \mu, \tau$, the initial state is described by

$$|\nu_\alpha(0,0)\rangle = U_{\alpha 1}|\nu_1\rangle + U_{\alpha 2}|\nu_2\rangle + U_{\alpha 3}|\nu_3\rangle \tag{1.5}$$

and after time t the state becomes

$$|\nu_\alpha(x,t)\rangle = U_{\alpha 1}|\nu_1\rangle e^{-i\phi_1 t} + U_{\alpha 2}|\nu_2\rangle e^{-i\phi_2 t} + U_{\alpha 3}|\nu_3\rangle e^{-i\phi_3 t}, \tag{1.6}$$

where the phase $\phi_j = E_j - p_j$ and E_j and p_j are the energy and momentum of the j^{th} mass state.

Expressing the initial mass states in terms of the flavour states:

$$|\nu_j(0,0)\rangle = \sum_\alpha U^*_{\alpha j} |\nu_\alpha(0,0)\rangle \tag{1.7}$$

we can see that the superposition of mass states becomes a superposition of flavour states at time $t > 0$:

$$|\nu_\alpha(x,t)\rangle = \sum_{\alpha,j} U_{\alpha j} U^*_{\alpha j} e^{-i\phi_j t} |\nu_\alpha\rangle \tag{1.8}$$

and we can express the probability that a neutrino of flavour α will oscillate into a neutrino of flavour β, the *transition probability*, as

$$P_{\nu_\alpha \to \nu_\beta}(t) = \left| \langle \nu_\beta | \nu_\alpha \rangle e^{-i(E_j - p_j)t} \right|^2 = \sum_{j,k} U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k} e^{-i(\phi_j - \phi_k)t}. \tag{1.9}$$

Since neutrinos are relativistic (mass $m_\nu \ll E_\nu$),

$$p_j = \sqrt{E_j^2 - m_j^2} \approx E_j - \frac{m_j^2}{2E_j} \tag{1.10}$$

so the phase difference between two states j and k, which is the mechanism responsible for the neutrino oscillations, is

$$\Delta\phi_{jk} \approx \frac{\Delta m^2_{jk} t}{2E_\nu}, \tag{1.11}$$

where we assume the mass eigenstates were created with the same energy.

Additionally, for ultrarelativistic particles, we can use the approximation $t = L$, where L is the distance travelled by the neutrino from source, and so the transition probability becomes

$$P_{\nu_\alpha \to \nu_\beta}(L, E_\nu) = \sum_{j,k} U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k} e^{-i \frac{\Delta m^2_{jk} L}{2E_\nu}}. \tag{1.12}$$

The probability that the neutrino will be in its initial flavour state at a later time is called the *survival probability*:

$$P_{\nu_\alpha \to \nu_\beta}(L, E_\nu) = 1 - 4 \sum_{j>k} |U_{\alpha j}|^2 |U_{\alpha k}|^2 \sin^2\left(\frac{\Delta m^2_{jk} L}{4E_\nu}\right). \tag{1.13}$$

As we saw in Section 1.4.1, the atmospheric neutrino results from Super-K can be explained in terms of vacuum oscillations principally of ν_μ to ν_τ transitions.

So we have a model describing oscillations in a vacuum arising from the phase differences between the mass eigenstates, which change the interference pattern. The maximum probability of oscillation is determined by the mixing angle $\sin^2(2\theta)$ and the oscillations depend on the mass-squared difference between the mass eigenstates Δm^2, the neutrino energy and

the distance from the point of creation/weak interaction. This model is well-demonstrated by experiment.

Vacuum oscillations have a consequence for reactor antineutrino detection. Reines and Cowan observed antineutrinos from a reactor using a detector situated next to a reactor core but, beyond 10 km from a nuclear reactor, vacuum oscillations of electron antineutrinos into muon antineutrinos cause a dramatic drop in the electron antineutrino flux. This reduces the observable antineutrino signal from a reactor using a remote detector and the vacuum survival probability for electron antineutrinos, assuming zero CP violation and normal mass ordering ($\nu_3 > \nu_2, \nu_1$), is given by:

$$\begin{aligned} \mathrm{P}(\bar{\nu}_\mathrm{e} \to \bar{\nu}_\mathrm{e}) = 1 &- \cos^4\theta_{13}\ \sin^2(2\theta_{12})\ \ \sin^2(1.27\Delta\mathrm{m}^2_{21}\mathrm{L}/\mathrm{E}_{\bar{\nu}_\mathrm{e}}) \\ &- \cos^2\theta_{12}\ \ \sin^2(2\theta_{13})\ \ \sin^2(1.27\Delta\mathrm{m}^2_{31}\mathrm{L}/\mathrm{E}_{\bar{\nu}_\mathrm{e}}) \\ &- \sin^2\theta_{12}\ \ \sin^2(2\theta_{13})\ \ \sin^2(1.27\Delta\mathrm{m}^2_{32}\mathrm{L}/\mathrm{E}_{\bar{\nu}_\mathrm{e}}), \end{aligned} \tag{1.14}$$

where L is the distance from the reactor in metres and $E_{\bar{\nu}_e}$ is the antineutrino energy in MeV.

On the other hand, it can change the detectable antineutrino spectrum, which is then dependent on the distance from a reactor (see Section 4 for more detail). This can make reactor ranging possible through comparing the detected spectrum with the expected spectrum at different distances.

1.4.3 Matter Oscillations (and the solution to the Solar Neutrino Problem)

Neutrino propagation in matter is affected by interactions with particles in the medium, which enhance the oscillation probabilities for the electron neutrino and antineutrino. Flux predictions for the Homestake and Super-K solar neutrino experiments, taking into account these matter effects in the Sun, were borne out in the electron neutrino fluxes observed.

Figure 2 shows the Feynman diagrams for the neutrino-electron elastic scattering interaction $\nu_\mathrm{x} + \mathrm{e}^- \to \nu_\mathrm{x} + \mathrm{e}^-$. All flavours of neutrino can interact via neutral current elastic scattering but electron neutrinos can additionally interact via charged current elastic scattering on electrons in matter. This leads to a difference in the elastic scattering potential V_e for ν_e compared to the potential V_α for the other flavours:

$$\Delta\mathrm{V} = 2\sqrt{2}\mathrm{E}_\nu\mathrm{G}_\mathrm{F}\mathrm{n}_\mathrm{e} \tag{1.15}$$

where G_F is the Fermi coupling constant for weak interactions and n_e is the electron number density in a medium. This leads to an additional term in the phase difference, which becomes

$$\Delta\phi_\mathrm{m} = \frac{\Delta\mathrm{m}^2_{jk}}{2\mathrm{E}_\nu}t + \sqrt{2}\mathrm{G}_\mathrm{F}\mathrm{n}_\mathrm{e}t = \frac{\Delta\mathrm{m}^2_{jk}L}{2\mathrm{E}_\nu} + \sqrt{2}\mathrm{G}_\mathrm{F}\mathrm{n}_\mathrm{e}L. \tag{1.16}$$

It is this difference in the phase due to the additional potential experienced by electron neutrinos which in practice leads to the additional effects in matter, such that the oscillations have a different amplitude and length compared to oscillations in a vacuum.

The *Mikheyev, Smirnov and Wolfenstein* (MSW) theory of oscillation in matter predicted that the solar electron neutrinos emitted from the Sun's core should be converted into equal parts ν_e, ν_μ

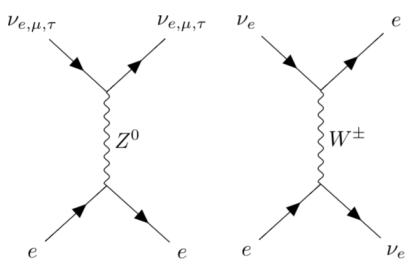

(a) Neutrino-electron neutral-current (left) and charged current (right) elastic scattering interactions.

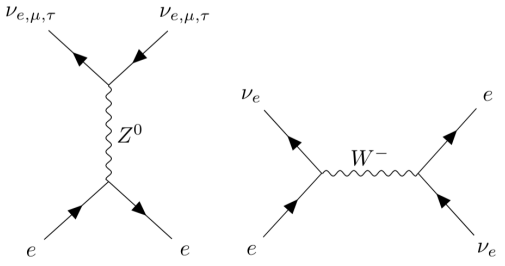

(b) Antineutrino-electron neutral-current (left) and charged current (right) elastic scattering interactions.

Figure 2: Neutrino and anti-neutrino elastic scattering interactions. Only electron neutrinos and antineutrinos can interact via the charged-current interaction, leading to an additional potential in their matter Hamiltonian which affects neutrino oscillations.

and ν_τ as they travel. This is the result of a combination of resonant enhancement of oscillations due to matter effects in the Sun[7] and vacuum oscillations. Electron neutrinos detected in the Homestake experiment should therefore be one third of the neutrino flux emitted by the Sun - which is exactly what was observed.

The Super-K solar neutrino results are similarly explained by matter effects. Super-K observed neutral-current elastic scattering of ν_μ and ν_τ off electrons, in addition to the neutral-current and charged-current elastic scatter of ν_e off electrons. The ν_μ and ν_τ interactions have a cross section of approximately a sixth of the ν_e elastic scattering cross section (0.157×10^{-43} (E_{ν_e}/10 MeV) cm^2 compared to 0.920×10^{-43} (E_{ν_e}/10 MeV) cm^2 [44]) and Super-K should be able to see 45% of the solar neutrino flux. This, again, is exactly what was observed.

Matter effects are small at reactor antineutrino energies and over the distances under consideration and as such are neglected for the purposes of this work.

1.5 Summary

- Reactor antineutrino detection has a long history going back to the first discovery of neutrinos.
- Three neutrino flavours exist: left-handed electron, muon and tau neutrinos and their right-handed antineutrino counterparts.
- Neutrino interactions, mediated by the weak force, are described by the Standard Model of Electroweak Interactions.
- Neutrinos have mass and are subject to oscillations between flavours as they travel.

[7] Resonance between the oscillation length and the refraction length in matter for a given electron number density can make total oscillation between flavours possible.

- Reactor $\bar{\nu}_e$ are subject principally to vacuum oscillations, which result in a reduced flux and altered energy spectrum at a distant detector.
- Vacuum oscillations are dependent on a number of parameters constrained by experiment (see Table 1).

Table 1: Current global best fits for neutrino oscillation parameters [55] for normal mass ordering (NO) and inverted mass ordering (IO) where applicable.

Oscillation parameter	1-sigma best fit
$\Delta m^2_{21}\left(10^{-5}\text{eV}^2\right)$	$7.50^{+0.22}_{-0.20}$
$\Delta m^2_{31}\left(10^{-3}\text{eV}^2\right)$	$2.55^{+0.02}_{-0.03}$ (NO) $2.45^{+0.03}_{-0.02}$ (IO)
$\sin^2\theta_{12}/10^{-1}$	3.18 ± 0.16
$\sin^2\theta_{23}/10^{-1}$	5.74 ± 0.14 (NO) $5.78^{+0.10}_{-0.17}$ (IO)
$\sin^2\theta_{13}/10^{-2}$	$2.200^{+0.069}_{-0.062}$ (NO) $2.225^{+0.064}_{-0.070}$ (IO)
δ_{cp}	$1.08^{+0.13}_{-0.12}$ (NO) $1.58^{+0.15}_{-0.16}$ (IO)

2 Reactor Antineutrinos for Monitoring

It's still magic even if you know how it's done.

– Terry Pratchett, *A Hat Full of Sky*

In all types of nuclear reactor, energy is produced by the fission of ^{235}U,^{238}U,^{239}Pu and ^{241}Pu into neutron-rich nuclei which then undergo a series of β decays to stability. This releases, on average, 6 antineutrinos per fission at energies up to 10 MeV. The antineutrino emission is dependent on the power of the reactor (how many fissions) and the composition of the reactor core (the fission fractions of the different isotopes).

The detectable flux decreases with distance from a reactor (the *standoff*) and is further altered by neutrino oscillations. Information about the core composition and location of a reactor is carried by the outgoing particles from antineutrino interactions in a detector. The precise composition of the core of a reactor and the time evolution of the core (*burnup*), including the refuelling frequency, depend on the reactor type.

All types of nuclear power reactor generate an isotropic antineutrino flux - $\mathcal{O}(10^{20}\text{s}^{-1})$ from a 1GW_{th} reactor [56] - so, although the interaction cross section of antineutrinos with matter is very small, the enormous number of antineutrinos released means that the antineutrino signal from a reactor can be seen in a variety of detectors. Indeed, the history of reactor antineutrino detection goes right back to the very first neutrino detector built at Savannah River by Reines and Cowan (see Section 1.2). Reactor antineutrinos can be detected for reactor monitoring through the outgoing particles from inverse beta decay and neutrino-electron elastic scattering using water Cherenkov and scintillation technologies.

This chapter discusses the background to reactor antineutrino monitoring. The physics and operation of nuclear reactors are discussed in Section 2.1. The reactor emission and information encoded in the antineutrino flux is detailed in Section 2.2. Interactions used for reactor antineutrino detection are covered in Section 2.3 and antineutrino detection technologies are discussed in Section 2.4. Section 2.5 explains how information about a reactor can be extracted from the detected signal.

2.1 Nuclear Reactors and Power Generation

The fission of atoms in a nuclear reactor produces energy which can be used to generate electricity. The fission reaction is initiated by firing neutrons at atoms, causing them to split (fission) into smaller atoms. This is accompanied by the emission of energy (gammas) and further neutrons. Moderators are used to slow down the neutrons released by fission so that they in turn can initiate fission of more atoms, which releases more gammas and neutrons. Control rods, e.g. made of boron, are inserted in the core to reduce the fission rate in the chain reaction. In this way, a sustained and controlled chain reaction is set in progress. The chain reaction produces the vast amounts of energy required to make the electricity. A coolant takes the energy produced in the fission reactions, in the form of heat from the core, for use in the next step of the power production process.

The Pressurised-Water Reactor (PWR) is the most common type of reactor. It is a type of Light Water Reactor (LWR), which runs on solid fuel and uses water as both the coolant and the

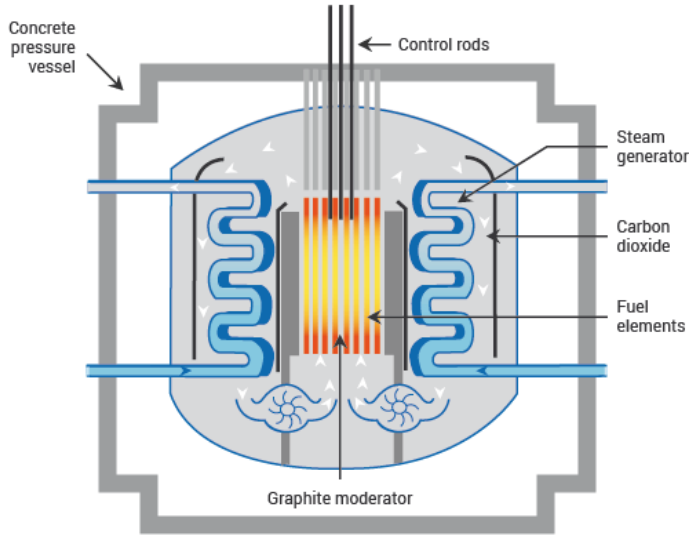

Figure 3: Advanced Gas-Cooled Reactor (AGR) design [57].

neutron moderator. The core in a PWR is cooled by normal ('light') water, which flows through it under very high pressure. Inside a LWR, the fuel rods are immersed in water. The water coolant is heated to ~300 °C. It is kept at high pressure (typically around 150 atm) to ensure it stays in liquid form as it is passed through a network of looped pipes in a heat exchanger. Here, a secondary water flow, at a much lower pressure, is converted to steam by the heat from the hot, pressurised water from the core.

The Advanced Gas-Cooled Reactor (AGR) is a form of the more general Gas-Cooled Reactor (GCR) type and it is the type of reactor most commonly run in the UK. GCRs use graphite as a neutron moderator and a gas as the coolant. In an AGR, the coolant is carbon dioxide. The fuel rods are contained in graphite channels which act as the moderator and form a network of channels through which the gas coolant can flow (see Section 4.2 for more detail). As the carbon dioxide gas coolant is circulated through the core, it is heated to ~650 °C. This creates steam via heat exchange in a similar way to the PWR. Figure 3 is a schematic showing the key components of the AGR design.

In both the PWR and AGR, the steam drives turbines and this kinetic energy is converted to electrical energy in a generator. Finally, the electricity is transformed to high voltage and sent to the national grid. The steam used to drive the turbines is cooled and condensed before being piped back in to be reused in the steam generator.

The nuclear fuel in PWRs and AGRs is ^{235}U. Natural uranium contains only 0.7% ^{235}U and so it is enriched prior to use in a nuclear power station to increase the ^{235}U concentration to between 3% and 5%. Enriched uranium oxide is formed into hollow, sintered[8] pellets which are strung together in sealed metal tubes called fuel pins or rods. Fuel rods are then grouped together to form a fuel assembly. The power output is dependent on the mass of uranium dioxide in the core. In the Sizewell B PWR core in the UK, 264 fuel pins make up each of 193 fuel assemblies

[8]Sintered means that it has been compacted and formed into a solid mass by heating or cold pressing, without liquefaction.

to give a thermal capacity of 3425 MW_{th} from 101 tonnes of uranium oxide [58].

The four main fissioning isotopes in a core are ^{235}U, ^{238}U, ^{239}Pu and ^{241}Pu. The principal fuel, ^{235}U contributes around 70% (55%) of the fissions in a reactor at the mid-point in the AGR (PWR) fuel cycle. ^{238}U is the most abundant isotope in natural uranium (>99%) - and still makes up most of the uranium in the enriched uranium fuel - but it has a fission fraction in a reactor of less than 10%. Mid fuel cycle, about a third of the power in a PWR and 20% of the power in an AGR is generated by the fission of ^{239}Pu, which is produced in a neutron-rich core by neutron capture on ^{238}U, followed by beta decay:

$$^{238}_{92}\mathrm{U} + \mathrm{n} \rightarrow {}^{239}_{92}\mathrm{U} \xrightarrow{\beta^-} {}^{239}_{93}\mathrm{Np} \xrightarrow{\beta^-} {}^{239}_{94}\mathrm{Pu}.$$

^{241}Pu is produced from ^{239}Pu in the core by the successive capture of two neutrons.

The AGR undergoes more frequent refuelling than the PWR: every ~160 days for an AGR, compared to every ~600 days for a PWR. The refuelling frequency balances safe operation with optimising fuel usage and minimising reactor downtime. In an AGR, neutron capture on the graphite moderator is minimal and so fuel pins can be spaced far enough apart to allow replacement of individual fuel channels during refuelling, making it a relatively quick process.

By contrast, in a PWR the pins and assemblies must be densely packed to reduce neutron capture on the water moderator. For this reason, refuelling necessitates removing all of the fuel at once in order to replace depleted assemblies. It is therefore a much more lengthy process and is carried out less frequently to reduce the impact of shutdowns, whereas the AGR refuelling schedules are geared towards optimal use of the fuel. [59]

The sub-optimal use of the fuel in PWRs is reflected in a marked reduction in the fission rate (and consequently the power output) with time as the core burnup (depletion of ^{235}U) progresses, as discussed in Section 2.2 and Figure 4. The change in rate due to burnup is less significant in the AGR due to the shorter fuel cycle.

2.2 Reactor Antineutrino Emission - Unshieldable Signal

To be most useful for application to non-proliferation, a reactor monitor must provide better or more information than the current monitoring methods. This is the strength of antineutrino monitoring. The weakly interacting nature of the neutrino makes the antineutrino flux unshieldable and continuous i.e. not dependent on external conditions, e.g. cloud cover. In addition, detailed information about the reactor core composition, power and even location is encoded in the antineutrino signal.

The reactor antineutrino flux and the information it is possible to glean from it is dependent on the reactor thermal power and core composition, on the nuclear physics of the fission of the isotopes in the core and on neutrino oscillations which alter the flux with distance from a reactor.

2.2.1 Emitted Reactor Antineutrino Flux

The antineutrino flux from a reactor core is primarily dependent on its power output however the fission fractions of each of the fissioning isotopes have a subdominant but important effect

on the reactor antineutrino spectrum. Determining the power of antineutrino detection for reactor monitoring, or indeed monitoring nuclear reactors for verification purposes, requires a precise understanding of the antineutrino spectrum emitted from a reactor. Accurate prediction of the $\bar{\nu}_e$ spectrum of a reactor requires detailed information about the fissile fraction and time evolution of the four fissioning isotopes, as well as an accurate understanding of the associated β spectra of the fission products of each of the isotopes.

The derivation of the antineutrino spectrum from a reactor is described in [60]. The antineutrino spectrum from a fissioning isotope in a reactor is related to the power output and composition of a reactor core by

$$\Phi_{\bar{\nu}_e,i}(E_{\bar{\nu}_e}) = P_{th}\frac{p_i\lambda_i(E_{\bar{\nu}_e})}{Q_i} \tag{2.1}$$

where P_{th} is the thermal power of the core, p_i is the fraction of the thermal power resulting from the fission of isotope i, Q_i is the average thermal energy emitted per fission and λ_i is the energy spectrum in antineutrinos per fission for fissioning isotope i given by

$$\lambda_i(E_{\bar{\nu}_e}) = \exp\left(\Sigma_{j=0}^{5} a_j E_{\bar{\nu}_e}^j\right), \tag{2.2}$$

where the coefficients a_j are fit parameters from the Huber-Mueller predictions [61] [62], which are derived from measurements of the β spectra from nuclear fissions.

A change in the antineutrino flux emitted by a reactor can be caused by a change in the core composition or reactor thermal power, or indeed by both. SONGS1 (see Section 2.4.3) demonstrated that the antineutrino flux from a reactor reflects the operating power and fuel evolution of the core [63]. The SONGS power station consisted of three Pressurised Water Reactors (PWRs). This type of reactor is switched off to undergo refueling every 18 to 24 months. Figure 4 shows the reported reactor power and antineutrino fluxes seen in the detector with time. The power outages reported by SONGS (Figure 4 (a)) are clearly reflected in the number of events seen in the detector (Figure 4 (b)).

The antineutrino flux is sensitive to the proportion of ^{235}U in a reactor core and the total flux varies according to the *fission fractions* - the relative weighting of the fissioning isotopes in the core. Neutrons from the fission of ^{235}U are captured on ^{238}U, which can decay to ^{239}Pu via two β decays, resulting in an accumulation of ^{239}Pu as the *burnup* of ^{235}U progresses. ^{239}Pu can undergo α decay to ^{235}U or can absorb neutrons and fission along with the ^{235}U. The antineutrino emission from ^{239}Pu is ~65% of the emission from a quantity of ^{235}U with the same thermal power output [64]. Thus the ^{235}U burnup is marked by a steady decrease in the antineutrino flux (Figure 4 (c)), despite a constant power output (Figure 4 (a)). Time-dependent fission fractions depend on the initial composition of a reactor, the power output, burnup and the specifics of the fuel and reactor design.

2.2.2 Reactor Anomaly

In 2011 the reactor antineutrino spectra from the four fissioning isotopes were modified by the widely used Huber-Mueller model [61] [62], which calculates antineutrino spectra by

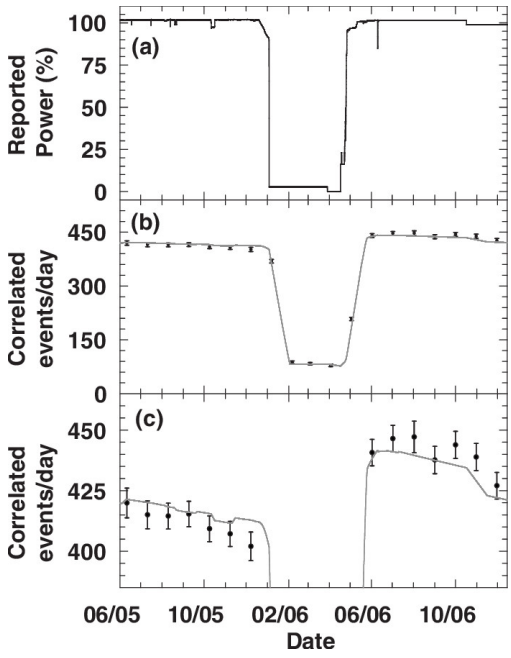

Figure 4: Power outages seen in the antineutrino flux in the SONGS1 detector showing (a) reported power, (b) detected events and (c) close-up of detected events showing a decrease in rate with fuel evolution. The markers show the rate measured and the lines show simulated rate based on the reactor power reported by the operator.

Re-printed from [63] Bowden et al. Observation of the isotopic evolution of pressurized water reactor fuel using an antineutrino detector, Journal of Applied Physics, 105(6) (2009), with the permission of AIP Publishing.

conversion from measurements of cumulative β spectra from fissions. However, it was found that this resulted in a $5.7 \pm 2.3\%$ deficit [65] in the detected antineutrino flux - the reactor antineutrino anomaly - as shown in Figure 5(a). The most recent measurements, from STEREO, confirmed this deficit, measuring the ^{235}U reactor antineutrino flux at the global average [66]. Additionally, results from experiments including Daya Bay, RENO and Double Chooz show an excess in the number of electron antineutrinos at ~5 MeV - the 5 MeV bump - as seen in Figure 5(b).

The Daya Bay and RENO collaborations [67] [68] found that the reactor anomaly could be largely resolved by re-evaluating the emission from ^{235}U. RENO has also seen a dependence of the ~5 MeV excess on reactor thermal power and a hint of a correlation between the shoulder and the fraction of ^{235}U in the core [68]. Experimental measurements were found to agree better up to ~5 MeV with fluxes calculated ab initio from fission and β yields. Re-evaluation with the Huber-Mueller model using updated flux predictions suggest that the deficit is due in full to ^{235}U [69].

The Huber-Mueller calculation of antineutrino spectra relies on measurements of β spectra using thin foils of uranium and plutonium irradiated by thermal neutrons. The measurements originally used for the calculations were carried out for ^{235}U, ^{239}Pu and ^{241}Pu in the 1980s. The measurement for ^{238}U was carried out in 2013 [70]. Updated measurements in 2021 for ^{235}U and ^{239}Pu [71] have led to a preliminary recalculation of the antineutrino spectra using the Huber-Mueller model with the newly-measured β spectra, which finds a 5% suppression in the spectra for both isotopes [72]. The newly-calculated spectra agree with the data, implying a resolution of the reactor antineutrino deficit if this result is confirmed.

Until there has been conclusive resolution of the reactor anomaly, it is important to take this anomaly into consideration when calculating the expected flux from a reactor. As yet, the shoulder in the data at 5 MeV is unresolved and should be factored in to the spectra used. The results presented in this thesis use spectra calculated prior to the recalculation in 2021. The

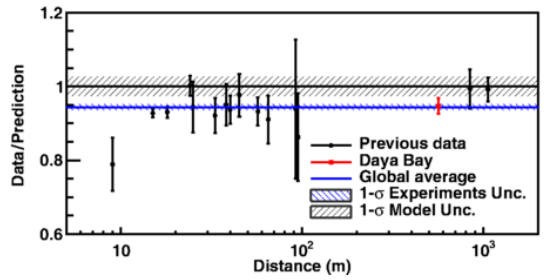

(a) **Reactor anomaly**: ~6% deficit in the global average reactor antineutrino flux measurement (blue) compared to the Huber-Mueller prediction at distances over 1 km from the reactor for the three-flavour neutrino model and normalised to the Huber-Mueller prediction at 1

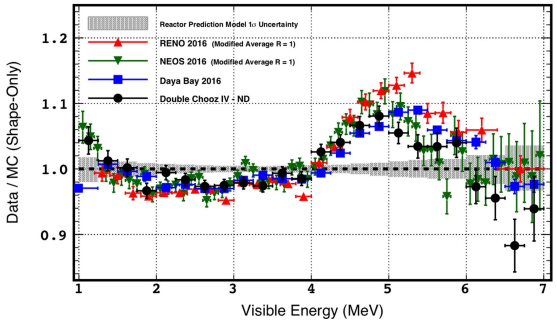

(b) **5 MeV bump**: excess at ~5 MeV in measured reactor antineutrino flux compared to Huber-Mueller model prediction. The normalised spectral flux from Double Chooz and Daya Baya data in particular agree well across the entire energy range

Figure 5: Reactor anomaly and 5 MeV bump from Daya Bay.

Reprinted figure (a) with permission from Prof. Kam-Biu Luk. F.P An et al (Daya Bay Collaboration), *Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay*, Phys. Rev. Lett. 116, 061801, 2016. Copyright 2016 by the American Physical Society.
Reprinted figure (b) by permission from Springer Nature Customer Service Centre GmbH: Nature, Nature Physics, De Kerret, H. et al., textitDouble Chooz theta(13) measurement via total neutron capture detection, Copyright Springer Nature 2020. doi:10.1038/s41567-020-0831-y

systematic uncertainty on the IBD rates quoted in Section 4.3.7 attempt to take into account the flux deficit and 5-MeV bump.

2.3 Reactor Antineutrino Interactions

Reactor antineutrinos can be detected for reactor monitoring through the outgoing particles from inverse beta decay and neutrino-electron elastic scattering using water Cherenkov and scintillation technologies. These interactions and technologies have differing characteristics which may be exploited for application to remote reactor monitoring and will be described in this and the next section.

As Reines and Cowan showed [16], the antineutrino emission from a reactor can be detected via the inverse β decay (IBD) weak interaction (Figure 6) of antineutrinos with protons in water or a hydrocarbon liquid:

$$\overline{\nu}_e + p \longrightarrow e^+ + n.$$

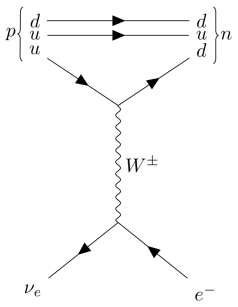

Figure 6: The inverse beta decay interaction. Time runs from left to right.

The IBD cross section is $\mathcal{O}(10^{-44})E_e p_e$ cm^2 [73]. Although small, it is relatively high compared to the cross sections of other antineutrino interactions in matter and it has been calculated to within 1% accuracy at low energies. To date, what is considered to be the most accurate cross section in the MeV to GeV range is given in [74]. The cross section outlined in [73] is also very accurate at the low energies of reactor antineutrinos, particularly $<$ 20MeV.

IBD is the dominant interaction of antineutrinos with energies of less than a few tens of MeV and has a low threshold energy E_{thr} which can be expressed in terms of the proton, neutron and

positron rest masses m_p, m_n and m_e as approximately:

$$E_{thr} = \frac{(m_n + m_e)^2 - m_p^2}{2m_p} = 1.8\text{MeV}. \tag{2.3}$$

in the laboratory frame.

The positron gains most of the kinetic energy of the antineutrino and with good energy resolution the positron can therefore provide information about the energy of the incoming antineutrino. The energy of the incoming antineutrino is related approximately to that of the positron by:

$$E_\nu = E_{e^+} + E_{thr} - m_e. \tag{2.4}$$

While the positrons are emitted almost isotropically, with a slight bias in the backwards direction, the neutron takes on most of the antineutrino's momentum and its initial direction is largely parallel to that of the incoming antineutrino. From the point of emission, the neutron then takes a random walk and thermalises in the detector medium through successive scatterings, which knock the neutron off its original path. Once thermalised, the neutron is captured e.g. on a hydrogen nucleus and a second signal arising from the de-excitation of the capture nucleus can be detected. This occurs within a short distance and time of the positron signal and results in a signal of coincident interactions which can be beneficial for background rejection. The time and distance between the positron and neutron events are dependent on the medium in which the interaction takes place.

An alternative detection mechanism is via the $\bar{\nu}_e e^-$ neutral-current and charged-current elastic scattering interactions. These have a cross section of $\mathcal{O}(10^{-45}) m_e E_{\bar{\nu}_e}$ [75] and theoretically no threshold energy, although in practice the detection of these interactions is limited by the detection threshold. Although the cross section is lower, the number of particles with which the antineutrino can interact in this way is higher. The outgoing electrons are predominantly forward-scattered, particularly at the higher end of the reactor antineutrino energy spectrum (Figure 7).

2.4 Reactor Antineutrino Detection Technologies

Water Cherenkov and scintillator detectors are the two principal types of antineutrino detection technology. Scintillator detectors are a proven technology for reactor antineutrino detection. A nascent water Cherenkov technology - gadolinium doping - widens the possiblities for reactor antineutrino detection. The combination of the two technologies into a water-based scintillator technology promises to exploit the best features of each method. This section will introduce water Cherenkov and scintillation technologies before going on to discuss the latest advances in antineutrino detection media.

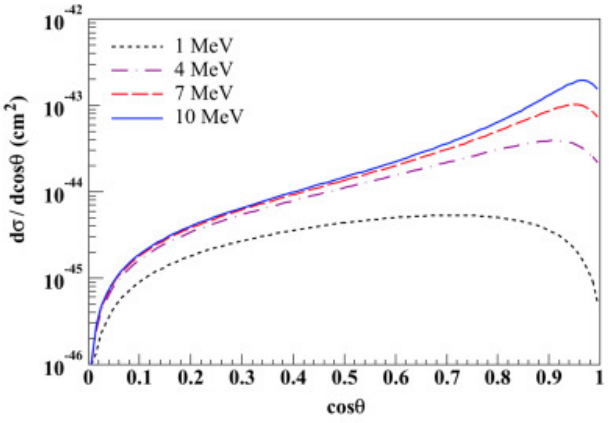

Figure 7: Angular distribution of scattered electrons with respect to the incoming antineutrino direction in antineutrino elastic scattering interactions.

Reprinted from [75] Hellfeld et al, Reconstructing the direction of reactor antineutrinos via electron scattering in Gd-doped water Cherenkov detectors, Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 841:130–138 Copyright 2017, with permission from Elsevier.

2.4.1 Water Cherenkov Detectors

Reactor antineutrinos can be detected via IBD and elastic scattering in a water Cherenkov detector. A water Cherenkov detector consists of a tank of ultra pure water instrumented with photomultiplier tubes (PMTs) close to the walls of the tank and pointing in towards a target volume. Although it must undergo intense purification to be used in a water Cherenkov detector, water is inexpensive and readily available and so very large detectors can be built.

Neutrino events are detected via Cherenkov radiation in the detector. Charged particles from neutrino interactions in water travel at a velocity v faster than the local speed of light in water:

$$v \geq \frac{c}{n}$$

where $n = 1.33$ is the refractive index of water. When a charged particle travels through water, electromagnetic radiation is emitted through the polarisation and subsequent de-excitation of molecules in the water. When a charged particle travels faster than the local speed of light, the disruption results in constructive interference of the polarisation field. This creates a characteristic blue *Cherenkov cone* of radiation, which can be detected via the PMTs on the walls of the detector as a Cherenkov ring or ellipse depending on the angle of approach.

The geometry of the cone (Figure 8) provides information about the velocity of the particle. The open angle of the cone is given by:

$$\cos\theta = \frac{1}{n\beta},$$

where $\beta = v/c$. The maximum Cherenkov angle in water is:

$$\theta_{max} = \cos^{-1}\left(\frac{1}{1.33}\right) = 41.2^{\circ}$$

and the Cherenkov threshold kinetic energy for electrons/positrons in water is:

$$E_{thr} = \frac{m_e c^2}{\sqrt{1 - \frac{1}{n^2}}} = 0.8\text{MeV}.$$

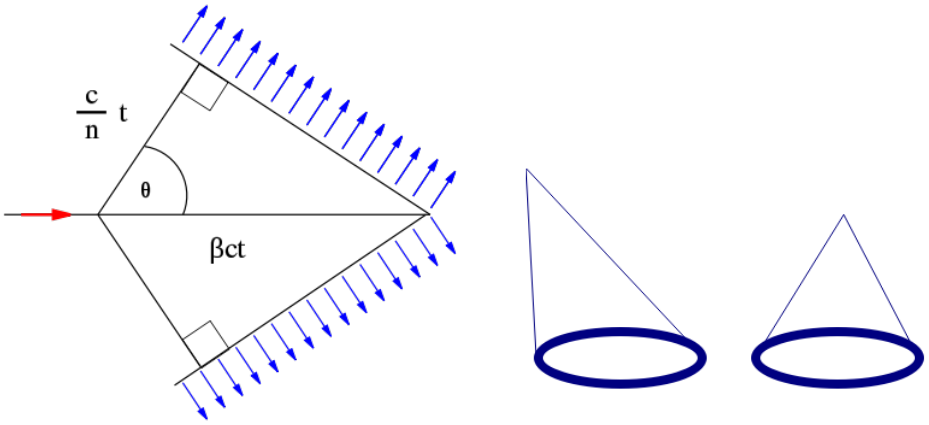

(a) The velocity of the charged particle can be deduced from the cone geometry.

A. Horvath, CC BY-SA 2.5
https://commons.wikimedia.org/wiki/File:Cherenkov.svg

(b) The interaction vertex and direction of the charged particle can be deduced from the pattern and time of arrival of light on the walls of a detector.

Figure 8: The geometry of the Cherenkov cone and the pattern of light can be used to deduce information about the velocity, vertex and direction of the charged particle.

The interaction vertex (position of the interaction in the detector) and direction of the charged particle can be reconstructed from the Cherenkov ring pattern and relative time of arrival of the light at the PMTs. For low-energy reactor antineutrinos, the energy of the particle can be calculated by taking into account the distance to the interaction vertex and the intensity of the light detected in terms of the number of hits on the PMTs. The vertex and energy resolution are limited by the relatively low light yield from Cherenkov light.

As yet, reactor antineutrinos have not been detected as a signal in a water Cherenkov detector. In pure water, the neutron from the inverse β decay reaction in the detector thermalises in the water and captures onto a hydrogen nucleus. The correlated signal consists of a prompt emission of Cherenkov radiation with a peak energy of ~2.5 MeV and a delayed neutron-capture emission of a single 2.2 MeV gamma ~200 μs later, as the deuterium nucleus de-excites. The hydrogen neutron-capture signal and much of the positron signal is below the currently achieved Super-K detection threshold ~3.5 MeV [76].

2.4.2 Gadolinium Water for Low-Energy Reactor Antineutrinos

The principle of using gadolinium (Gd) to delve into lower-energy neutrino detection was first introduced by [77] and developed by [78]. Gadolinium has a very high thermal neutron capture cross section (48, 800 barns (b) for natural Gd compared to ~0.3 b for hydrogen) and a relatively high-energy subsequent gamma cascade of ~8 MeV (mean total energy). This gives a more easily detectable correlated signal from the inverse beta decay reaction [78].

In gadolinium-doped water (Gd-H_2O) the neutron thermalises and captures after a mean time of ~30μs and within a distance of ~100 cm. Thanks to the higher neutron-capture cross section of Gd, it is captured preferentially onto the Gd nucleus. In Gd-H_2O, the delayed neutron-capture emission has a peak at ~4 MeV. This means that the positron-detection efficiency can be increased by looking for a positron-like signal (the *prompt* event) in

coincidence with the higher-energy (and thus easier to see) neutron-capture signal (the *delayed* event). Positron events at the lower end of the energy range would otherwise be lost amongst the background. In this way, Gd can lower the energy threshold of a water Cherenkov detector to increase sensitivity to the low-energy positrons for reactor antineutrino interactions via IBD. The characteristics of the neutron capture in Gd-H_2O compared to pure water are summarised in Table 2.

Table 2: Characteristics of neutron capture in pure water and Gd-H_2O

Characteristic	Pure Water	Gd-H_2O
Neutron-capture cross section (b)	0.3	49 000
Mean neutron capture time (μs)	~200	30
Energy released (MeV)	2.2	8

The optimal concentration of Gd in water has been found to be 0.1% Gd ions. At this concentration, over 80% and up to more than 90% [78] of the neutrons capture onto Gd. The remaining neutrons capture onto the hydrogen in the water. Beyond this concentration, the Gd begins to crystallise out of solution. This concentration can be achieved by adding 0.2% concentration of gadolinium sulphate ($Gd_2(SO_4)_3$).

The EGADS (Evaluating Gadolinium's Actions on Detector Systems) [79] gadolinium-doped water Cherenkov detector demonstrated that Gd-H_2O can be purified to ~100m attenuation lengths to over 90% of Super-K ultra-pure levels. This emerging Gd-H_2O technology has now been deployed in SK-Gd (Super-Kamiokande with Gadolinium) [80] and also in ANNIE (Accelerator Neutrino-Nucleus Interaction Experiment) [81] with a simplified water purification system [82]. Gd-H_2O would be the minimum fill required by WATCHMAN for reactor antineutrino monitoring.

2.4.3 Scintillator Detectors

Reactor antineutrinos have been detected with liquid and plastic scintillator detectors. Liquid scintillator detectors, like water Cherenkov detectors for reactor antineutrino detection, are enclosed tanks instrumented with PMTs but contain a liquid scintillator fill rather than water. A liquid scintillator is composed of an organic solvent containing a scintillating chemical in solution. Plastic scintillator antineutrino detectors may be of a layered or lattice design to combine the neutron-capture and scintillating elements. In a scintillating medium, the scintillator is hit by incoming particles, which impart energy to the scintillator. Excited scintillator particles then release this additional energy as light.

The Cowan-Reines experiment, detecting reactor antineutrinos, used organic liquid scintillator to detect the coincident signals from the back-to-back gammas from the annihilation of the IBD positron with an electron and the gammas from the capture of neutron on the cadmium in the cadmium chloride medium. The Kamioka Liquid Scintillator Antineutrino Detector (Kamland) [83] was able to make a measurement of the θ_{13} mixing angle using ultra-pure liquid scintillator (2,5-diphenyloxazole - PPO - in an organic solvent) contained within a balloon suspended within a larger spherical vessel to detect antineutrinos from a reactor.

Sudbury Neutrino Observatory Plus (SNO+) recently finished filling with a liquid scintillator fill. Designed primarily to search for neutrinoless double beta decay, which could point to whether or not neutrinos are Majorana particles, SNO+ will also be sensitive to reactor antineutrinos and will be able to make a measurement of Δm^2_{12} to a sensitivity of $0.2 \times 10^{-5}\text{eV}^2$ in about 7 years [84].

A number of small scintillator detectors designed to operate at very short range have been developed to detect reactor antineutrinos. SONGS1 was a feasibility study into antineutrino detection for nuclear non-proliferation and the first detector to be applied to reactor monitoring. It was a gadolinium-doped liquid-scintillator antineutrino detector located 25m from the 3.56-GW San Onofre Nuclear Generating Station (SONGS) and was able to demonstrate a correlation between reactor power and antineutrino detection rate as we saw in Section 2.2.

PROSPECT (Precision Oscillation and Spectrum Experiment) [85] is a segmented ^{6}Li – doped liquid scintillator detector and began data-taking in 2018. It is a very near-field reactor antineutrino detector designed to precisely measure the antineutrino flux and spectrum from nuclear reactors and also search for sterile neutrinos at metre-scale baselines.

Scintillators bring an increased light yield compared to water Cherenkov detectors, resulting in higher signal detection efficiencies, better position and energy resolution and lower energy thresholds. The scintillation light does not, however, carry the directional information that can be deduced from Cherenkov light and liquid scintillator fills are more expensive, less scalable to large detectors and less readily available than pure water.

2.4.4 Gadolinium-Doped Water-Based Liquid Scintillator

Gadolinium-doped water-based liquid scintillator (*Gd-WbLS*) has the potential to combine the benefits of liquid scintillator and gadiolinium-doped water Cherenkov detectors. The combination of Cherenkov and scintillation light ultimately brings increased IBD detection, with the added benefit of improved detection quality, particularly where it is possible to separate the Cherenkov and scintillation components.

Scintillation detectors bring a high light yield, low-energy detection and good energy and position resolution. However, they do not preserve directional information and are limited in size due to light absorption and the cost and availability of the medium. Water Cherenkov detectors bring directional information and are scalable to very large detectors. However, they have a low light yield and no sensitivity below the Cherenkov threshold. Combining the two media provides a solution which can be scaled to large sizes and results in a higher light yield, sensitivity down to lower energies, improved energy and position resolution and directional information from the Cherenkov light. The addition of gadolinium brings a further increase in light yield due to the neutron capture on gadolinium and enhanced background rejection due to the coincident signal pair which occur closer in space and time with gadolinium doping, as discussed in Section 2.4.2.

Water-based liquid scintillator (WbLS) is an emerging detector medium, which is still undergoing optimisation and improvement. A WbLS cocktail using the PPO (2,5-diphenyl-oxazole) wavelength-shifting scintillator in a linear alkylbenzene (LAB) solvent

has been developed [86] and gadolinium doping is under investigation. Pure liquid scintillator is a scintillating material in solution in an oily organic solvent. In WbLS, the scintillator is dissolved in an oily solvent in the same way. This solution is then further combined with pure water. The mixing between the oil and water in WbLS is achieved by the addition of a surfactant which creates micelles which have both hydrophilic (polar) and hydrophobic (non-polar) surfaces.

Cherenkov-scintillation separation disentangles the information about the particle direction and particle identification carried by the Cherenkov light, while preserving the excellent energy resolution and low energy threshold brought by the scintillation light. The Cherenkov and scintillation light differ in their angular, timing and wavelength distributions. The scintillation component is isotropic, while a Cherenkov cone is directional. The peak of the Cherenkov light from an event is earlier in time than that of the scintillation light from the same event. Cherenkov light spans a broad spectrum of wavelengths while scintillation light has a narrowband spectrum. These differing properties may be used to discriminate between the two components of the light.

The precise timing and wavelength distributions of the scintillation light depend on the decay-time constant and spectral response of the chosen scintillator and solvent mix. The LAB/PPO combination has a peak emission wavelength of 360 nm with a secondary peak at 380 nm [87]. The decay-time distribution has four exponential components - the rise time, fast decay and two successively slower decays. The fast decay generates the main component of the emitted light and, in a water-based LAB/PPO mixture, this component constitutes up to 90% of the scintillation light. The fast decay-time constant is particularly affected by the solvent used and the linear alkylbenzene (LAB) solvent results in a longer fast decay-time constant [88]. A longer fast decay-time constant increases the timing separation of Cherenkov and scintillation light, which is ideal for a WbLS detection medium. PPO concentration can in fact be tuned to optimise the Cherenkov-scintillation separation - lowering the PPO concentration can slow the scintillation decay time and reduce the light yield to increase the separation [87]. The Cherenkov light and the main component of the scintillation light in a 1% LAB/PPO WbLS mixture are separated on the order of nanoseconds.

Water-based liquid scintillator - particularly with the addition of Gd - is an emerging technology. The ANNIE [81] experiment is currently working towards a WbLS fill.

2.5 Detectable Reactor Antineutrino Spectra

The detection of antineutrinos above a known or unknown background can tell us about the existence of a reactor and can provide information about the power or reactor standoff through a straightforward counting experiment. With additional spectral information, there is the potential to achieve more precise ranging of a reactor and to extract information about the composition of a reactor core. With sensitivity to elastic scattering, it is also possible to obtain directional information and in theory to point back in the direction of a reactor.

The total detectable IBD or electron elastic scattering spectrum, prior to oscillations, is given by

$$R(E_{\bar{\nu}_e}) = \frac{N_p \sigma(E_{\bar{\nu}_e})}{4\pi d^2} \Phi_{\bar{\nu}_e}(E_{\bar{\nu}}) \tag{2.5}$$

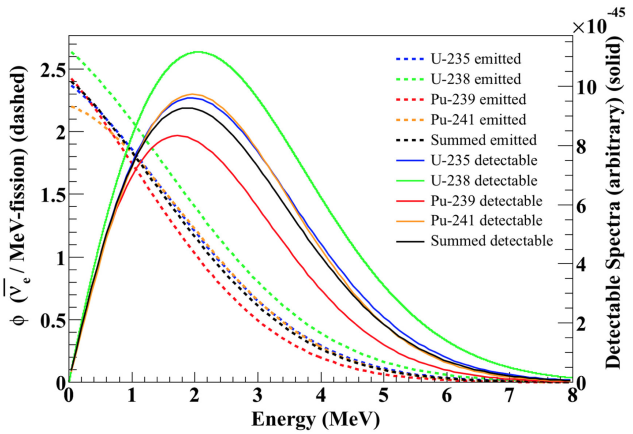

Figure 9: Emitted reactor antineutrino spectra and detectable elastic scattering spectra per fission for average mid-cycle fission fractions in a PWR.

Reprinted from [75] Hellfeld et al, Reconstructing the direction of reactor antineutrinos via electron scattering in Gd-doped water Cherenkov detectors, Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 841:130–138 Copyright 2017, with permission from Elsevier.

where N_p is the number of interacting particles (protons or electrons) in the detector target, d is the distance from the reactor and $\sigma(E_{\bar{\nu}_e})$ is the relevant cross section. Calculation of the total emitted reactor flux $\Phi_{\bar{\nu}_e}(E_{\bar{\nu}})$ summed over the four main isotopes in a reactor core as per Equation 2.1 is explained in Section 2.2.1. Figure 9 shows the emitted and detectable antineutrino spectra from elastic scattering from the emission spectra of each of the fissioning isotopes.

As equation 2.5 shows, there is some degeneracy between the reactor power, the core composition (fission fractions of the fissioning isotopes) and the distance of the detector from the reactor. Nevertheless, through the detection of antineutrinos, it is possible to put limits on these variables. By constraining one or two of the variables using external information, more detailed reactor information may be deduced.

Sensitivity to spectral information - given sufficient statistics, background rejection and energy resolution - can provide information about the composition of a core, since the emitted antineutrino spectrum is dependent on the precise fission fractions in the core. Achieving sensitivity to reactor composition would rely on excellent energy resolution and would most likely also require some degree of directional sensitivity. A spectral analysis could also make use of neutrino oscillations, which significantly alter the detectable reactor antineutrino spectrum over distances greater than around 10 km. Although this reduces the detectable spectrum, it also offers the potential to extract information about the distance to a reactor by comparing the detected spectrum with the expected spectrum at different reactor standoffs.

Directionality - the ability to point in the direction of a reactor - would be a standout non-proliferation goal. Other reactors are likely to be a significant background in the field, particularly where a small reactor has been 'hidden' near a larger reactor. The ability to resolve the signal from two reactors would be a significant non-proliferation goal. In theory, it would be possible to extract directional information from either the IBD or the elastic scattering interaction. However, the practicalities of both methods for remote monitoring are challenging.

Reactor antineutrino detection focuses on detecting the outgoing particles from the IBD interaction due to the higher interaction cross section of IBD compared to all other interactions, including elastic scattering. The positrons from IBD interactions are slightly back-scattered, while the neutron is forward-scattered. It may be possible to reconstruct the direction of the incoming neutrino to sufficient accuracy with a combination of neutron and

positron direction. This would rely on obtaining the directional information of the neutron within just a few scattering recoils and may be possible in a segmented detector. However, making a segmented detector large enough to operate remotely presents its own difficulties in terms of the cost and availability of materials.

As we saw in Section 2.3, elastic scattering gives us the most information about the direction of the incoming antineutrino. Electrons from neutrino-electron elastic scattering are largely forward-scattered, giving the potential to reconstruct the direction of the incoming neutrino with a good degree of accuracy. An investigation into using elastic scattering for directionality in the application to non-proliferation in a Gd-H_2O detector found that a very large detector and extremely stringent levels of cleanliness, which have not yet been achieved in any detector, would be required in order to disentangle the elastic scattering signal from backgrounds due to the radioactive decay of radon in the water [75]. While a gadolinium-doped water-based medium offers the potential to scale to a much larger detector, creating a large enough detector with sufficiently low backgrounds would present an obstacle.

The work presented in this thesis focuses on a counting experiment to detect the existence of a reactor, however the updated reconstruction algorithm, through lowering the energy threshold and improving energy resolution, could also increase the accuracy of reactor ranging (see Section 8.1).

2.6 Summary

- Nuclear reactors emit antineutrinos at a rate of $\mathcal{O}(10^{20}\ \mathrm{s}^{-1}\ \mathrm{GW}_{\mathrm{th}}^{-1})$.
- Reactor antineutrinos carry information about the power, composition and location of a core.
- The reactor anomaly is a deficit in the antineutrino flux from a reactor according to current calculations at the time of writing.
- Reactor antineutrinos can be detected via inverse beta decay (IBD), which emits a positron and a neutron, in water-based Cherenkov detectors.
- Gadolinium can enhance the sensitivity to reactor antineutrinos through tagging of the IBD neutron.
- Scintillators give an increased light yield but lose the valuable directional information carried by the Cherenkov light.
- Water-based liquid scintillator with gadolinium doping retains the best of Gd-H_2O and scintillator technologies.
- Reactor antineutrino detection can provide information about the existence of a core via a counting analysis.
- Spectral analysis can provide information about the reactor *standoff* and core composition.

- Obtaining directional information about a reactor would require significant background reduction or a very large detector to provide sensitivity to reactor antineutrino-electron elastic scattering.
- The work presented in this thesis focuses on using WATCHMAN's detector as a counting experiment to detect reactor antineutrinos via inverse beta decay interactions.

Reactor antineutrinos present a real possibility for remote reactor detection for non-proliferation and a prototype is now needed to understand the capabilities of antineutrino detection for remote monitoring. The planned prototype will be described in the next chapter.

3 Remote Reactor Monitor Prototype

I have been impressed with the urgency of doing. Knowing is not enough; we must apply.
– Leonardo da Vinci

The WATCHMAN (WATer CHerenkov Monitor for Anti-Neutrinos) scientific collaboration was formed to demonstrate the application of antineutrino detection to non-proliferation. The aim is to demonstrate the potential of antineutrino detection to provide information about a reactor's power, composition and location. The first experiment will demonstrate the ability both to discover an unknown reactor and to verify data reported about a known reactor. The WATCHMAN scientific collaboration is a U.S.-UK collaboration, with primary funding via the USA's Office of Defense Nuclear Non-Proliferation (DNN) under the remit of the NNSA's 2011 Strategic Plan [5] and in the UK from the Science and Technologies Facilities Council (STFC), with additional funding from the Ministry of Defence via the Atomic Weapons Establishment (AWE) in the UK.

The project is divided into the Advanced Instrumentation Testbed (AIT) civil construction and, initially, the Neutrino Experiment One (NEO) detector. The primary goal of AIT is to create a facility for the development and testing of advanced technologies for deep underground science, with initial non-proliferation applications mid- to far-field (tens to hundreds of km) remote monitoring of nuclear reactors. Fundamental physics research goals will also be pursued. NEO, a water-based Cherenkov detector, is the planned first experiment at AIT and it is currently in its design stage.

Details of the AIT-NEO project are set out in this chapter. Sections 3.1 and 3.2 introduce the AIT site and non-proliferation goals. While the non-proliferation goals are the focus of this thesis, physics goals are an additional focus in the UK and the principal physics potential is outlined in 3.3. The remainder of the chapter describes the NEO detector design to date.

My principal contribution to the WATCHMAN collaboration as described in this chapter has been the evaluation and optimisation of detector configurations and instrumentation in the progression towards a final design. Some of these configurations are discussed in Section 3.4. Other significant contributions include my work on the initial PMT acceptance testing and in particular the development of the analysis used for gain calibration (Section 3.7.1).

3.1 The AIT (Advanced Instrumentation Testbed) Site

The planned location of the AIT site is in the ICL Boulby Mine on the coast near Whitby in North Yorkshire. Boulby Mine is an ultra-low background environment and hosts the Boulby Underground Laboratory, which has been home to deep underground physics experiments since the 1990s and is now operated by the STFC. Boulby Underground Laboratory currently houses facilities including part of the CYGNUS directional dark matter detection programme and the Boulby UnderGround Screening (BUGS) facility for ultra-low background radioactive assay. As such, the proposed site benefits from the lab's pre-existing working relationship with the mine company, with established infrastructure and operational protocol, and from its experience hosting underground science experiments both inside and outside the lab.

The low-background environment in Boulby Mine can be attributed to low radioactivity rates in Boulby Mine's rock salt layers, a low cosmic muon rate and the world's lowest ambient air radon concentration in an underground lab (~3 Bq m^{-3}). Boulby Underground Laboratory and the proposed AIT site are at a depth of 1.1 km underground and an effective depth of 2.8 km.w.e. (km water equivalent) with a flat overburden. This results in a significant $\mathcal{O}(10^7)$ reduction in cosmic muons, which produce dangerous backgrounds for rare-event detectors. For comparison, the cosmic muon rate in Boulby Mine is a factor of ten lower than the rate in Kamioka Mine (home of Super-K) but a factor of 100 higher than in SNOLab (Sudbury Neutrino Observatory Laboratory) which, with a flat overburden and an effective depth of almost 6 km m.w.e., has the lowest cosmic muon rate in the world $\mathcal{O}(10^{-13})$. Radioactivity and cosmic muon rates are discussed in more detail in Chapter 4.

The proposed cavern location is close to the existing Boulby lab at a 26 km standoff from the EDF Hartlepool Nuclear Power Station, which is intended to be the primary signal for the non-proliferation goal of the experiment, subject to its continuing operation. The Heysham and Torness nuclear power stations offer the potential for detecting more distant signals, at a distance of 149 km and 187 km from the cavern respectively (see Chapter 4 for further discussion of the reactor signals). The precise details of the facility are yet to be confirmed but the designs currently under consideration have been analysed with respect to the potential signals at Hartlepool, Heysham and Torness [7] and the results are presented in Chapter 7.

In addition to the main NEO detector, a 20 to 100-tonne science and technology demonstrator has been proposed. The demonstrator would serve to test out innovative technology that may be used in later experiments and would be installed within the existing lab at Boulby.

3.2 AIT-NEO Non-Proliferation Goals

The principal non-proliferation goal of the AIT-NEO Project is to demonstrate reactor antineutrino detection as a non-proliferation tool for monitoring a reactor. To address the potential non-proliferation use cases to which the technology may be applied, the WATCHMAN collaboration has identified verification and discovery scenarios:

- Verification scenario - verify the operational cycles of reactors. This assumes full and accurate knowledge of reactors and could vary in difficulty from simple verification of the power cycle of reactors to more complex verification of the composition of reactors.
- Discovery scenarios - discovering an unknown reactor, with varying degrees of information about backgrounds and existing reactors.

The focus of the analysis presented in this thesis addresses discovery scenarios since these are the most challenging goals for remote detection. See Chapter 7 for a more detailed interpretation of the discovery scenarios.

Additional goals of the project, as defined in 2020 by the project's UK/U.S. Joint Governance Board, are to:

- *"Innovate, advance and demonstrate technologies and associated instrumentation for neutrino detection and analysis."*

- *"Increase engagement between nonproliferation and fundamental physics communities in the U.S. and UK."*
- *"Attract and train early career scientists and engineers to increase scientific and nonproliferation understanding and collaboration throughout and between the U.S. and UK."*

3.3 Physics Goals of AIT-NEO

3.3.1 Supernova Neutrinos

Detection of neutrinos from a galactic, core-collapse supernova is the standout physics goal for AIT and will be possible in the initial phase with the NEO detector. The supernova signal would be detected primarily through IBD interactions of supernova electron antineutrinos. Neutron tagging in a Gd-doped detector lends power to supernova detection through the detection of coincident IBD events and through the detection of neutron events where the positron energy is below the detectable threshold.

A large number of IBD interactions in the detector in a short space of time would flag a supernova explosion. Since the neutrino signal from a supernova event would arrive before the visible light, the astrophysics community could be alerted to the imminent arrival of the light. Extracting data from a supernova explosion on an event-by-event basis would demonstrate how the supernova neutrino emission evolves over time and may provide insight into the particular mechanisms at work in a supernova explosion.

Prior to a supernova explosion, there is an emission of very low-energy (~2 MeV) antineutrinos from silicon burning in the star in its final stage of fusion. The silicon burning in a large, nearby star would generate a steep increase in the number of IBD interactions in the hours before a supernova explosion. While the positron from these interactions would mostly be too low in energy to detect, an increase in the number of neutron captures could provide advance warning of a supernova explosion.

Using neutron tagging for IBD events would also allow the neutrino-electron elastic scattering signal to be distiguished from the more dominant IBD interactions. Elastic scattering interactions would provide directional information (as discussed in Section 2.3) to point back towards the origin of the supernova explosion.

NEO would therefore be able to detect and communicate the advent of a core-collapse supernova in a timely manner and disentangle the IBD and elastic scattering interactions to provide vital information on the physics at work in a supernova explosion[9].

3.3.2 Neutrino Oscillation Measurements

Reactor antineutrino experiments can be used to measure neutrino oscillation parameters. At a 26 km standoff from the Hartlepool reactor, the AIT site sits on the falling edge of the electron antineutrino survival probability (ratio of neutrinos observable in the detector to neutrinos at the source) as a function of distance from the reactor (Figure 10). With sufficient energy resolution, which may be achieved with a Gd-WbLS medium (see Section 2.4.4), NEO could

[9] So all we need now is a supernova explosion some time towards the end of this decade.

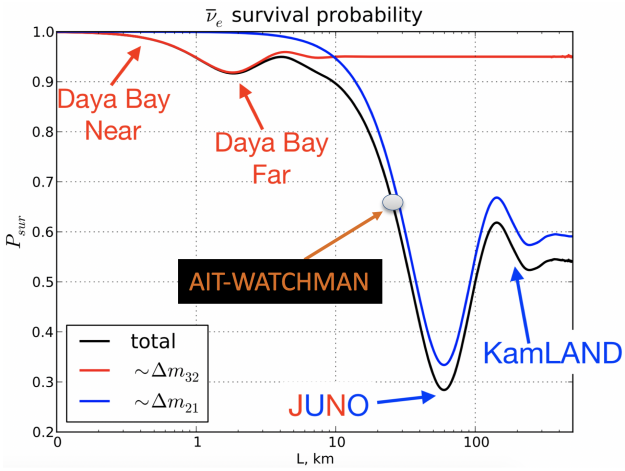

Figure 10: Electron antineutrino survival probability P_{sur} as a function of distance L from a reactor. NEO and subsequent detectors at the AIT site could provide a vital data point between existing Kamland/Daya Bay experiments and the upcoming JUNO experiment. [91]

have sensitivity to the oscillation parameter Δm^2_{12} through a measurement of the survival probability. AIT will be positioned midway between the baselines (distance from reactor or antineutrino source) of the Daya Bay and KamLAND experiments and, as such, could provide a vital data point to bridge the gap between the Daya Bay [89] and KamLAND [90] results.

3.4 The NEO (Neutrino Experiment One) Detector

NEO is the planned first experiment at the AIT site and its principal aim is to demonstrate the ability to remotely detect a reactor. The NEO detector is to be a kilotonne-scale, gadolinium-doped, water-based Cherenkov detector. The detector design has not yet been finalised and studies of the kind presented in Chapter 7 of this thesis will help to determine the optimal design to balance cost and sensitivity to remote reactor detection.

A number of design proposals have been considered over the last year as part of a down-selection process. The two major design styles considered were an upright cylinder and a horizontal rectangular prism shape and these were investigated in a number of configurations of outer volume size, inner volume size, PMT coverage and detector fill. Up to five different detector fill media were considered. Many of the configurations have now been rejected and only two detector sizes remain under consideration - a *baseline* design, which is the design towards which the collaboration is now working, and a *maximal* design, evaluation of which will demonstrate the maximum achievable results within the physical constraints of the AIT site, other constraints not withstanding. This thesis presents results for the two remaining designs, each with two fill media, for a number of reactor signals.

The preferred fill is Gd-WbLS, possibly with a transitory Gd-H_2O first phase. A water-based medium is the most scalable in terms of cost and availability of a material. Gadolinium doping has been shown by simulation to be required for the detection of more distant reactors. NEO is likely to be the first experiment to use Gd-WbLS but as yet the medium is still under development. LAB/PPO concentrations of 1%, 3% and 5% in water have been investigated for the NEO detector. A 1% WbLS solution with relative LAB/PPO concentrations of ~1% LAB and 0.002% PPO and doped with 0.1% gadolinium has been selected for NEO. A 1% LAB/0.002% PPO WbLS mixture gives a good separation of Cherenkov and scintillation light, which could

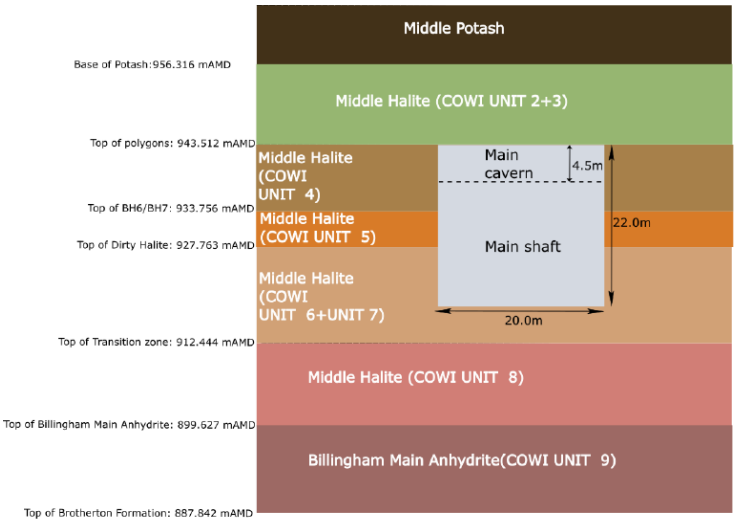

Figure 11: Geology at the AIT site in Boulby Mine. The cavern will be excavated in the more stable halite layers. Reduction in stability due to the span of the cavern necessitates the engineering of hanging galleries for the larger detector/cavern configurations.

Source: COWI UK Ltd.

be detected with fast photosensors in NEO.

Although the final tank dimensions and detector specifications have yet to be decided, the design concept for the detector system (the tank, instrumentation and subsidiary systems and structures) has been well developed and there are many features which would be common to any final configuration. Common features are described in the following subsections, based on the detail given in [91] and adapted to the current baseline design, with additional details from [92]. These features have also been adopted for evaluation of the maximal design for the purposes of this thesis. Since the project is still in its design stage, it should be noted that any of the details presented may be subject to change before the design is finalised.

3.4.1 Rejected Designs

A 20 m x 20 m upright cylinder was designed and costed. Cavern dimensions are constrained by the mine geology as shown in Figure 11. Ensuring the stability of a cavern large enough to house a cylinder of this size would necessitate extensive specialist engineering of hanging galleries built into the halite layers above the cavern to support the detector from above, coming at considerable cost in terms of both time and money.

In an effort to reduce costs, the horizontal rectangular prism design was proposed as a way to reduce the width and height of the detector, thus reducing the span and height of the cavern. This would use the conventional mining techniques employed for excavating roadways in the mine and dispense with the need for the hanging galleries. Two horizontal lengths of this type of rectangular, prism-shaped detector were considered - 50 m and 80 m, both of which were 8 m high x 8 m wide. The rectangular prism design was demonstrated by simulation to be untenable due to the dominance of backgrounds from fast neutrons. These scale with the surface area of the detector, which was high in proportion to the inner volume for this shape, and are usually reduced through shielding provided by the outer detector, which was absent in this design.

Smaller cylinders of 10 m and 12 m were also studied. Like the rectangular prism designs, these would also reduce the span and height of the cavern, avoiding the additional engineering

works to the cavern. Simulation showed that these smaller tanks were unable to achieve the project goals within a reasonable time frame due to their small size and lack of outer detector, as in the horizontal rectangular prism designs.

A matrix of detector configurations was considered in the down-selection process. Each of the above tank designs and sizes investigated and ultimately rejected was studied with a range of photocoverages, varying from 10% to 40%. A earlier version of the analysis presented in Section 6 was used to evaluate and compare the sensitivity of many of these configurations [6].

3.4.2 Current Baseline Design

The baseline design is intended to balance the costs and engineering works to the cavern with the non-proliferation and science potential of the project. The detector would be a 16 m x 16 m upright cylinder in an 18 m wide and 20 m high cavern. An inner PMT radius of 5.7 m has been proposed and this would give a 2.3 m-wide volume between the tank wall and the inner volume, which would serve to reduce the backgrounds from cosmogenic muons and natural radioactivity from the tank steel and rock/concrete in the cavern (see Chapter 4). In the baseline design this is a passive buffer. This could be instrumented with outward-pointing photomultiplier tubes to form an active muon veto (see Section 4.3.2). The inner volume of the detector would be instrumented with ~1600 (~2500) inward-pointing low radioactivity 10" Hamamatsu R7081 high-quantum efficiency photomultiplier tubes (*the 10" PMTs*) for a 10% (15%) photocoverage. A 15% photocoverage has been used in this thesis.

3.4.3 Current Maximal Design

The maximal design has been conceived to evaluate the full potential of the AIT site, given reduced constraints on costs and engineering. The maximum possible cavern size at the AIT site would be a cavern 27 m high by 27 m in diameter, even with the construction of hanging galleries in the halite layer above the cavern. To allow for access on top of the detector tank, the maximum tank dimensions in a cavern of this size would be a 22 m right cylinder. The proposed inner PMT radius for this sized detector is 9 m, leaving a 2 m passive buffer between the tank wall and inner volume. A 15% photocoverage would require ~4600 of the 10" PMTs to instrument the inner volume.

3.5 NEO Tank System

The tank will be an enclosed cylinder constructed from 316L stainless steel bolted panels with a seal weld on the inside. This material has been selected for compatibility with Gd-H_2O and Gd-WbLS, both of which are chemically aggressive. Figure 12 shows the conceptual design for the tank, the inner PMT structure and other support structures, the stairwell between the tank and lined cavern walls (for access around the detector) and the cleanroom area above the tank.

The bottom of the cylinder will be supported on a concrete base and installed on an arrangement of carbon steel girders to form a shallow conical base with a 2° slope down to a 150 mm-diameter central drainage port.

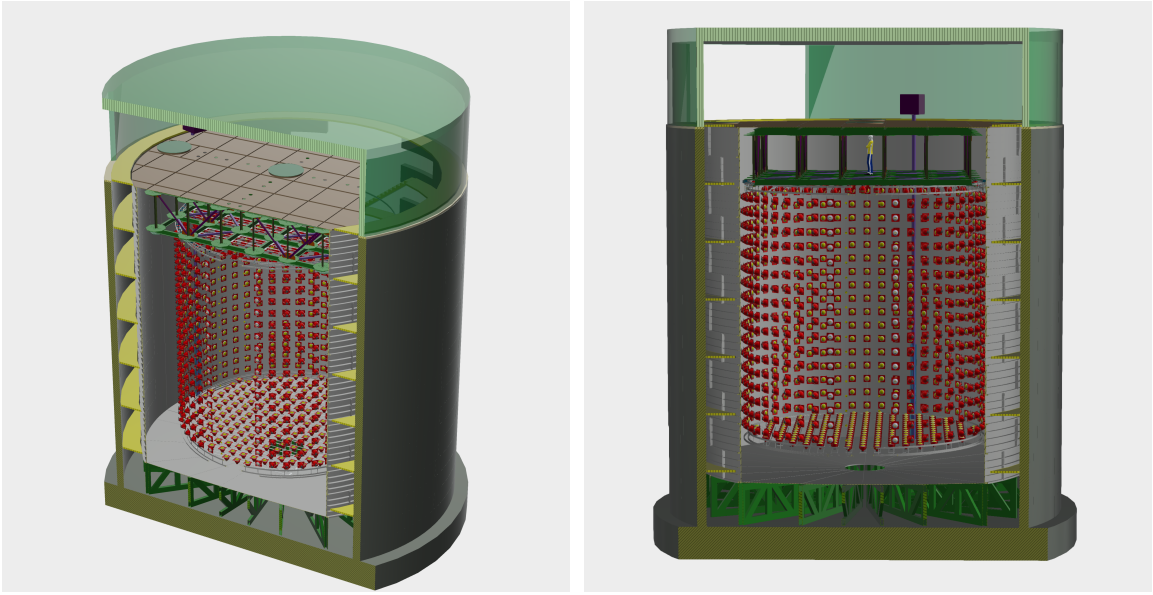

Figure 12: Schematics of the NEO baseline conceptual design by Jan Boissevain (University of Pennsylvania), showing the roof deck cleanroom with calibration glovebox and racetrack area between the tank and cavern walls, tank with top steel truss structure, top and bottom PMT support structure endcap and inner PMTs, tank base with water drainage connection point, stairwell and steel staircase outside the tank and concrete base and steel support girders under the tank.

A stainless steel deck at the top of the cylinder will serve as both a lid for the tank and a floor for the space above the tank. It will be fixed to the tank walls with a gas-tight seal to preserve the integrity and stability of a gas layer at the top of the tank. Above the tank, the lid will constitute a robust floor for an area of, at minimum, 'office-like' cleanroom standard. A set of approximately 2 m-high steel support trusses will be welded to the top of the cylinder (under the lid) to support both the floor deck and the PMT support structure inside the tank (see Section 3.6). An additional *racetrack* area will extend the top deck beyond the tank walls towards the cavern walls.

The top-deck cleanroom will protect electronics racks and ports for calibration, water circulation, high voltage cables and electronics cables and slow control cables for data acquisition (see below for details of these subsystems).

There will be around ten calibration ports across the top deck for the insertion of calibration sources into the detector. The calibration ports are likely to be ~30 cm in diameter with 3 m-long 316 stainless steel or PTFE tubes attached, which will guide sources past the truss system and into the inner detector. Four light injection ports on the top of the tank will allow the cables for the light injector calibration system to pass through to the inner volume. Two ports may also be included for the water attenuation monitor (see Section 3.13) should it be installed in the outer detector.

A total of four ~20 cm-diameter ports will be provided for PTFE water input and output pipes

for the inner and outer detectors. Water will be input from the base of the tank and output from the top to create the desired water-flow direction and all pipes will be split into four to create a more even flow in the tank. Water will be sampled for the purposes of water monitoring at different depths in the detector via 3 small ports in the tank top.

Other ports will include four optical-detector cable feed-through ports and five installation hatches on the tank roof, plus a 2 m x 2 m square access door in the wall of the tank near the base.

There will be an ullage gas layer at the top of the tank above the detector liquid which will allow for the movement and expansion of of the liquid. This will be controlled and purified by an ullage gas system, which will extract and purify air from the lab above the tank, before entering into the tank via ports in the tank roof. Gas will leave the tank via a pressure relief valve.

Outside the tank, thermal insulation to help stabilise the temperature of the detector is provided for through a 10 cm-thick layer on the sides of the tank and plastic flooring at the top of the tank. The outside of the tank will additionally house a stainless steel staircase for access around the tank and, potentially, a vertical water transparency monitor (see Section 3.13). On the inside wall of the tank there will be a support system for the installation of magnetic field compensation cables to avoid the detrimental effect that a perpendicular field has on the collection efficiency of PMTs.

3.6 NEO PMT Support Structure

The PMT support structure will comprise a vertical, tubular structure made of cylindrical steel rods and a top and bottom modular endcap. The structure will be supported from above by the previously mentioned ~2 m-high truss structure and attached at the base to a set of ~1 m-high floor trusses. The modularity of the top cap allows the PMTs to be removed separately to allow access via the central access hatch on the tank roof.

Each pair of vertical rods in the upright portion of the support structure and each steel in the endcaps will support the inward-facing PMTs, which will be mounted in plastic moulded housings. Outward-facing veto PMTs would be supported in a similar way, should an active muon veto be selected.

High voltage (HV) and signal cables will be attached to the vertical and endcap supports, secured in cable trays positioned throughout the detector volume and passed through flanges in the lid and deck into electrical racks on the racetrack area of the top deck.

The PMT housings (see Figure 13) have been designed to provide protection from a chain reaction in the event of implosion of one of the PMTs. PMT implosion due to the failure of the glass envelope can set off a chain reaction of multiple implosions in a detector. The housings, designed for what was originally to be the water Cherenkov far detector for LBNE (Long Baseline Neutrino Experiment), will be purpose-cast and designed to be fitted interchangeably to either a pair of rods in the vertical PMT supports or to one of the steels in the endcap modules. The design of the housings could also incorporate transparent acrylic enclosures with an acrylic-gel fill, which would be fitted over the front faces of the PMTs for additional protection from implosion. This may also help to reduce the backgrounds in the detector arising

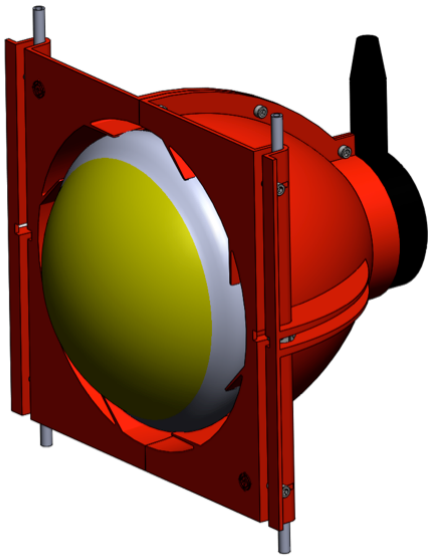

Figure 13: Schematic drawing of PMT housing design by Sam Wolcott (University of Wisconsin-Madison). The housings will be equipped with attachments for both the vertical and endcap supports and will provide protection from implosion.

from radioactive decays in the PMT glass. Pressure and materials compatibility testing of the PMTs prior to installation will also help to mitigate the risk of implosion (see Section 3.7.5).

A light barrier in the form of opaque black plastic sheeting will form the division between the inner and outer detector. This will reduce reflection within the inner volume, which can have a detrimental effect on vertex and energy reconstruction. An additional reflective Tyvek outer layer can improve the efficiency of an active muon veto by increasing the light captured by veto PMTs, if installed.

3.7 Photomultiplier Tubes (PMTs)

A PMT is a type of photosensor and consists of an evacuated envelope made of borosilicate glass (or other material e.g. fused silica) and containing a photocathode and a current multiplier. A photon is converted to a photoelectron at the photocathode and an electron multiplier, made up of a number of stages, multiplies the current through the emission of secondary electrons at each stage. The signal is read out at the anode. The multiplier *gain* is the number of electrons generated per originating photoelectron. With a gain of the order of 10^6 or more, it is possible to count single photoelectrons. Multiple-photon pulses may be evaluated by the charge read out at the anode.

The photocathode is deposited as a thin layer on the inside front of the envelope. Photocathodes vary in their sensitivity to different wavelengths and in their quantum efficiency - the probability that a photon incident on the photocathode is absorbed and converted to a photoelectron. Bialkali photocathodes for water Cherenkov applications have a peak quantum efficiency at blue wavelengths. High quantum efficiency is considered to be up to 35% for this type of photocathode. Commonly used bialkali materials include antimony-rubidium-caesium (Sb-Rb-Cs) and sodium-potassium-antimony (Na-K-Sb).

As previously stated, the detector will be instrumented with low radioactivity 10" Hamamatsu

R7081 high-quantum efficiency photomultiplier tubes (*the 10" PMTs*). The 10" PMTs multiply the signal from the photocathode down a chain of ten dynodes. The main body of the envelope is approximately hemispherical but the PMTs are being re-designed for the WATCHMAN collaboration with a more spherical shape to ensure they exceed the pressure rating of 2 bars required for the NEO detector (see Section 3.7.5). They have a bialkali photocathode coating with a spectral response between 300 and 650 nm, peaking at 420 nm, and a quantum efficiency of 25% to 35% at 390 nm [93]. The diameter of the PMT is ~253 mm and the diameter of the minimum effective photocathode area is ~220 mm.

The PMT envelopes are made from borosilicate glass manufactured using low-radioactivity materials and with a process which minimises radioactivity in the final product. The radioactivity rates for ^{238}U, ^{232}Th and ^{40}K are up to ten times lower in low-radioactivity glass and so this substantially lowers backgrounds in the detector.

Other PMT characteristics which can have significance for the experiment are the timing and charge response and dark noise background. The timing and charge response are discussed in more detail below. Dark noise is due principally to the thermionic emission of electrons from the photocathode.

Preliminary acceptance tests were carried out on a sample of one hundred 10" PMTs [94]. The testing procedure is now being refined in preparation for acceptance testing of all PMTs prior to installation.

The sample of 10" PMTs was tested at a rate of up to eight per day. A blue LED (470 nm wavelenth) was triggered by an LED driver at 10 kHz to give 1-2 ns light pulses. A neutral density filter was used to attenuate the light to ensure that the majority of the signals were from single-photoelectrons, with a mean number of photoelectrons between 0.1 and 0.3 per trigger. The signal was passed from the PMT anode through an amplifier for 10 times amplification and the waveforms were then recorded using a CAEN V1730B digitiser, which was also triggered at 10 kHz by a gated signal from the LED driver. The PMTs were tested in dark tents to minimise pollution from external light and were placed in the dark tents for a *cool-down period* of at minimum 18 hours at a nominal voltage before testing. This allowed the dark noise rate to reduce and stabilise following exposure to light during installation.

Integrated waveforms from the digitiser were combined to produce charge distributions for each of the tests. A fixed integration window of 50 ns starting 15 ns before the mean signal peak time (measured from the time of the LED signal) was used for most of the tests. The integration window was increased to 70 ns for the gain calibration to take into account the wider charge distribution at higher voltages.

The aim of the acceptance tests is to verify that the following PMT characteristics are within the required ranges for NEO:

- High voltage for a gain of 10^7
- Dark count rate
- Peak-to-valley ratio
- Rate of after pulses

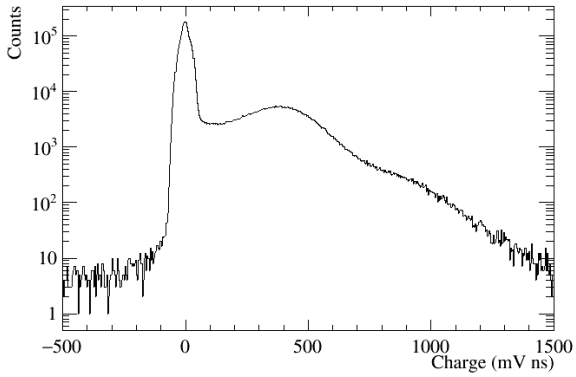

Figure 14: Typical charge spectrum for one of the Hamamatsu R7081 10" PMTs measured during the preliminary acceptance tests. Triggers where no pulses were recorded populate the *pedestal* around zero. The peak at 450 mV ns is due to single-photoelectron pulses and the double-photoelectron peak is at around 900 mV ns.

3.7.1 Gain Calibration

PMT gain is the ratio of the charge output at the anode to the number of photoelectrons produced by conversion at the photocathode. Gain calibration ensures the correct interpretation of the charge output at the PMT anode and is essential in setting operating voltages to achieve a uniform charge response across all PMTs in the detector.

For the purpose of the PMT acceptance testing, the gain tests verify that the required charge gain can be achieved within the limits of the range of voltages at which the PMT can be operated and set the operating voltage at which the remaining tests should be performed. The required gain for NEO is 10^7. The Hamamatsu R7801 PMT can be operated at voltages up to 2000 V. A PMT might be rejected if the required 10^7 gain cannot be achieved below the operating limit of 2000 V, or some specified value below this limit to allow for PMT degradation.

The *operating voltage* required to give 10^7 gain was calculated from the fit to charge distributions (Figure 14) over a range of high voltages for each PMT as described in Appendix A.

The results obtained for the first one hundred PMTs were largely correlated with the Hamamatsu shipping data. All PMTs were found to achieve 10^7 gain at less than 2000 V. For those PMTs requiring close to the 2000 V limit for 10^7 gain, long-term investigations into the extent and timescale of gain deterioration are underway.

Data-taking to measure the dark count rate, peak-to-valley ratio and delayed pulsing rate was carried out using the results of the gain calibration to set the high voltage to the operating voltage value, such that the single-photoelectron peak gives the expected charge output.

3.7.2 PMT Timing Distribution

The shape of the PMT anode output signal reflects a PMT's timing and resolution [95]. The transit time distribution for a single photoelectron (SPE) response is used to characterise this

signal. Transit time is the time from the creation of a photoelectron at the photocathode to the output of the resulting signal at the anode at the back of the PMT. The SPE transit time distribution has a number of features as discussed in [96]. Most notably, there is a prompt Gaussian signal peak (where a photon converted to a photoelectron at the photocathode constitutes a signal), delayed pulsing peaks and a continuous but random distribution of dark noise events (see below).

Signal-induced backgrounds to the SPE signal in a PMT hit cause delayed pulses in the PMT. These can be categorised according to the originating processes and the expected timing in the 10" PMT as follows:

- Double pulsing: undersized pulses due to inelastic scattering off the first dynode. Accompanies an undersized prompt signal pulse and occurs at around 50 ns after the trigger.
- Late pulsing: due to elastic scattering off the first dynode. Occurs without a prior prompt signal at around 70 ns after the trigger.
- After pulsing: undersized pulses due to ionisation of gases in an imperfect vacuum in a PMT. Occurs with a prior prompt signal beyond ~200 ns after the trigger.

The transit time spread of the PMT is the statistical fluctuation of the prompt signal peak, often expressed in terms of standard deviation. The dispersion is due in part to a combination of the launch angle and emission energy of the photoelectron at the photocathode and at each dynode, as well as the way in which the charge is collected by the anode. A large contribution to the dispersion also arises due to the varying distance between the photocathode and the first dynode, depending on the photon's point of arrival on the face of the PMT. The size of the spread has an impact on the ability to reconstruct the vertex. [95]

For the purposes of the acceptance tests, double and after pulses were counted to be pulses over a fixed threshold and delayed pulsing was calculated as a percentage of the signal pulsing rate. Delayed pulsing for the 10" PMTs was found to be consistent with previous measurements.

Table 3 shows the timing characteristics of the 10" PMTs as determined by Hamamatsu. The transit time spread in terms of the standard deviation of the prompt signal is 1.44 ns [93]. This value affects the accuracy of any vertex reconstruction based on PMT timing. Full timing characterisation of a selection of the PMTs will be performed prior to installation in the detector.

Table 3: Timing characteristics of the Hamamatsu R7081 HQE PMT with operating voltage set at 1500 V to give gain = 10^7 [97].

Transit time spread σ_{prompt} (ns)	1.44
Transit time spread FWHM (ns)	3.4
Dark noise (Hz)	8000.
Late pulsing (%)	4.0

3.7.3 Dark Count Rate

Dark count in the PMT (as distinct from dark current, which is a function of gain) is dominated in air by thermionic emission from the photocathode. Where multiple PMTs are submerged together in water, it is expected that there is also some contribution from radioactive decay in the glass of other PMTs and other backgrounds.

Dark noise affects triggering and reconstruction, particularly at low energies. A reduction in the dark rate simplifies triggering and improves vertex resolution. The dark rate can be calculated as the number of pulses in a pulse-height distribution above a suitable threshold, where, for best results, the threshold is set from the value of the single photoelectron charge Q_{SPE} calculated by a fit (Equation A.5 in the Appendix) to the SPE integrated charge distributions of the PMTs. Out of the initial one hundred PMTs, all had acceptable measured dark rates below 10 kHz and a mean of ~2.5 kHz.

Since dark noise is largely due to thermionic emission from the photocathode, it is strongly dependent on temperature [95]. As such, monitoring and control of PMT temperature will be implemented for the full acceptance testing and the detector will be kept at below ~15° C during stable operation.

3.7.4 Peak-to-Valley

The peak-to-valley ratio is a measure of the signal-to-noise of a PMT and is the ratio of the height of the single-photoelectron peak to the minimum of the charge spectrum (Figure 14) in the valley region between the *pedestal*, where no pulses were recorded after the trigger, and the SPE peak. The ratio is calculated by fitting to the valley and SPE peaks to extract the minimum and maximum of the valley and the SPE peak respectively. A nominal acceptable peak-to-valley ratio is 2 and care must be taken to ensure that there is little to no external noise which may skew the results.

3.7.5 Mechanical Performance and Chemical Compatibility Testing

Pressure testing is intended to ensure that:

1. the 10" PMTs, manufactured with the weaker, low-radioactivity glass, conform to the required pressure rating for NEO and that the glass is not degraded by long-term immersion in Gd-H_2O or Gd-WbLS and
2. the implosion of a single PMT does not lead to a chain reaction resulting in multiple failures.

Pressure tests are being performed to verify pressure ratings and evaluate the protection from shock-induced multiple PMT failure afforded by the moulded PMT housings and, additionally, acrylic enclosures or CUPPA (Contained Underwater Photosensor and Pulse-Extraction Apparatus) modules (see Section 3.8). The pressure vessel has been designed to hold any enclosure or CUPPA module and at the time of writing a bare PMT and a number of enclosures have been tested to destruction.

The low-radioactivity glass of the 10" PMT has not previously been tested for compatibility with ultra-pure Gd-H_2O. As such, a compatibility testing regime has been developed as follows:

1. Two of the 10" PMT and waterproof potting assemblies have been submerged in a 1% gadolinium sulphate solution since 2018. The soak solution is tested regularly for signs of material degradation in the form of absorbance in the optical range of 200-900 nm.

2. Ten PMTs have been installed in a 0.2% concentration gadolinium sulphate solution in the ANNIE detector and will be run under operating conditions for two years.

3. Soaked PMTs will subsequently be pressure tested and their pressure rating compared to that of unsoaked PMTs.

A WbLS-compatible potting has been developed by Hamamatsu and the full PMT and WbLS-compatible potting assembly would need to undergo similar soak tests in WbLS followed by mechanical strength testing prior to installation in a WbLS detector medium without encapsulation.

3.8 Light Collection Enhancements

Optical detection is an area in which innovation could be particularly beneficial. A focus on optical detector upgrades and light collection has resulted in the investigation of a number of options.

A minimum of 20% photocoverage (and up to the 40% photocoverage adopted by Super-K) was originally under consideration for NEO. A 20% photocoverage was adopted as the baseline for the costing and review of the 20m tank design [91]. The photodetectors are one of the most expensive detector components and the need to drive down costs led to investigation of the detector sensitivity with a lower percentage photocoverage. It was found that the sensitivity to reactor-detection goals dropped off steeply below a 15% photocoverage, particularly in the smaller tanks [6] [98]. With any photocoverage below 20%, the reduced light collection and associated reduction in IBD detection degrades the detector performance and enhanced light collection that achieves improved detector sensitivity is always desirable. There are a number of light collection enhancements which may be incorporated into NEO and subsequent experiments.

Wavelength-shifting (WLS) plates and light concentrators are two relatively inexpensive and established forms of light collection enhancements. Increasing the effective light collection area with light collection enhancements such as these may increase detector sensitivity or indeed reduce the cost of instrumenting the detector. However, while they increase the number of photons collected, they can have an unwanted impact on timing and reconstruction and therefore the detector sensitivity does not necessarily scale with the number of photons collected.

A WLS plate [99] is an acrylic plate doped with a wavelength shifting chemical and fitted closely around a PMT, in contact with the glass envelope at the outer edge of the photocathode surface (the photocathode does not necessarily cover the full diameter of the PMT envelope). Plates designed for a Cherenkov detector absorb UVA and UVB light with wavelengths between 280 nm and 400 nm and re-emit the photons at a longer wavelength in the region of the peak quantum efficiency of the PMT photocathode at around 420 nm. Wavelength-shifted

photons are emitted isotropically after a time defined by the decay time constant of the fluorescent, wavelength-shifting material. Photons which are trapped within the plate by total internal reflection at the interface between the acrylic and the detector medium may then travel towards the PMT at the centre, with an additional associated delay. By absorbing and re-emitting light which would otherwise miss the PMT, WLS plates can significantly increase light collection. Reflective coating at the outer edges of the plates re-directs photons towards the centre, thus increasing the number of photons reaching the PMT. The shape of the plate can also be significant for the number of photons reaching the PMT. Square or petal shapes, rather than circular plates, reduce the number of photons that are permanently trapped in the plate and increase the chance that they are directed towards the PMT at the centre.

The timing delays in the re-emission and propagation of the wavelength-shifted photons can be problematic for reconstruction. As such, WLS plates are often used to increase the effective photocoverage in an outer detector where the number of photons collected is key and timing and vertex reconstruction are less important.

Light concentrators are of conical geometry with a specular reflective coating, which combine to collect light from around the PMTs and reflect it towards the PMT photocathode at the centre of the cone's narrower aperture. A common design for the conical shape is an off-axis parabola or Winston Cone [100] and the geometry can be tuned for the optimal acceptance angle and length depending on the detector design. Like WLS plates, light concentrators affect the PMT signal timing and thus can create issues with the vertex reconstruction, which can be particularly difficult to calibrate in a non-spherical detector.

Less established technology for increasing light collection is the retro reflector. This is a backing plate which reflects light from around the PMT in the same direction as the incoming light but offset by a small distance, such that it would be detected by a PMT on the opposite side of the detector. The effect in a water-based Cherenkov detector would be a primary, prompt Cherenkov ring followed by a secondary, delayed Cherenkov ring on the far side. This has been found to improve vertex and angular resolution in a water Cherenkov detector [101] and may be a method of improving light collection without detriment to the PMT signal timing and vertex reconstruction.

3.9 Alternative Photodetection

To take full advantage of the benefits of WbLS requires discrimination of the dual contributions from Cherenkov and scintillation light. This allows access to the directional information of the Cherenkov light as well as potential improvement in vertex, directional and energy resolution using the respective timing of the Cherenkov and scintillation light. Cherenkov-scintillation discrimination might be achieved with fast photosensors for timing discrimination or *dichroicons* for discrimination by wavelength.

The Large Area Picosecond PhotoDetector (LAPPD) is a fast, planar-geometry photosensor with millimetre spatial resolution and picosecond timing resolution [102]. The LAPPD has a 350-cm^2 detection area coated with a bialkali Na_2KSb photocathode. The signal produced at the photocathode is amplified, with a 10^7 gain, by a pair of micro-channel plates stacked behind the front plate. The amplified signal is then read out at the anode. The LAPPD has a quantum

efficiency of between 20% and 30%, similar to that of the 10" PMTs, and was found to have mm-scale spatial resolution and up to ~50 picosecond single-photoelectron timing resolution using a 6 GHz readout [103].

Timing resolution on the picosecond scale could enable discrimination of the Cherenkov and scintillation light, which are separated in time on the order of 10s of nanoseconds. LAPPDs have been installed with a readout system incorporating the 10 GSPS (giga samples per second) PSEC4 waveform recording chip [104] in the ANNIE detector, alongside conventional PMTs with the ADC (analogue-to-digital converter) readout system referred to in Section 3.11. Characterisation of the LAPPD by WATCHMAN is underway, and will include the development of analysis to deconvolve the signals, which is particularly important where there are multiple photon hits. There is a limited supply of LAPPDs and they are substantially more expensive per unit ($50k compared to $3k per unit), with a factor of ~100 smaller active area than the 10" PMT. As such, it is likely that the LAPPD, if deemed suitable for NEO, would be installed in tandem with other photodetectors, as in the ANNIE detector.

An alternative fast photosensor might be the multiple photomultiplier tube (mPMT) which contains a number of photomultiplier tubes inside a single envelope. This maximises the photocathode area as a percentage of the total photosensor area compared to using a single PMT. Discrimination of multi-photon hits through directional sensitivity and the separation of the photocathode area across individual tubes brings the benefit of increased charge resolution, which in turn provides better vertex and energy reconstruction. One such example is the mPMT optical module developed for KM3NeT, with thirty-one 3" PMTs inside an optical module with high-voltage supply and digitisation [105]. Although mPMTs have tens of picosecond timing rather than picosecond timing, the benefit of mPMTs is their lower cost per unit active area compared to LAPPDs and improved timing and reconstruction compared to single PMTs. Monte-Carlo simulations performed for KM3NeT show three times fewer mPMT optical modules were required to achieve the same performance as optical modules containing 10" PMTs.

A dichroicon [106] is a dichroic Winston cone - a light concentrator with a dichroic filter which can perform photon sorting by wavelength. An example configuration might use a dichroic short-pass filter to selectively filter and reflect longer wavelengths (450 nm to 900 nm) from Cherenkov light through a long-pass filter to a red-sensitive PMT at the inner aperture, while allowing shorter wavelengths (350 nm to 450 nm) to pass through to a blue-sensitive PMT behind the dichroicon (Figure 15). A second, standard light concentrator may be used to direct light to a blue-sensitive PMT or the dichroicon may be used without the second concentrator in combination with tiled photodetectors such as the LAPPD. Up to a 90% Cherenkov purity has been achieved in the lab and a large-scale detector simulation of PMTs with dichroicons has demonstrated clear Cherenkov-scintillation separation.

3.10 High Voltage

High voltage (HV) will be supplied via commercial, multi-channel units, with at least 1 mA per channel and which can provide bias in the range of the PMT operating voltage between 1000V and 2000V. Commercial HV modules have a standard ripple of less than 10 mVpp (millivolts peak to peak), which is sufficiently low for single-photoelectron sensitivity. The model will be

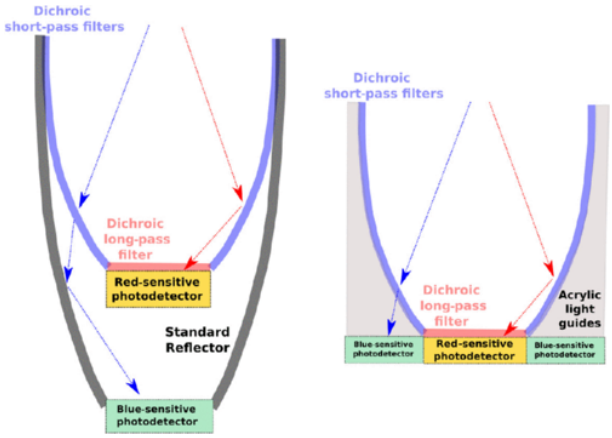

Figure 15: Suggested uses of dichroicon for Cherenkov-scintillation separation.

Reprinted figure with permission from T. Kaptanoglu. [106] Kaptanoglu et al, Spectral photon sorting for large-scale cherenkov and scintillationdetectors, Phys. Rev. D, 101:072002 Copyright 2020 by the American Physical Society.

selected subject to measurements of the single photoelectron response and must also be compatible with HV control and monitoring software, which can be incorporated into the data acquisition and readout software.

The HV system will be able to supply the specific operating voltages required for each individual PMT to give a gain of 10^7. The signal will be separated from the HV and fed into the data acquisition (DAQ) system on a channel-by-channel basis with a custom-designed splitter circuit.

3.11 Data Acquisition System

The DAQ system as a whole encompasses the triggering, extraction, pre-processing, monitoring and storage of data. The DAQ system design is largely driven by the PMT dark noise rate, which affects the triggering requirements, and the rate of downward-going cosmogenic muons in the fiducial volume, which affects saturation requirements.

Preliminary measurements indicate a PMT dark noise rate in water of around 3 kHz (subject to stringent testing). NEO will detect between 9 and 10 photoelectrons per MeV in Gd-H_2O and current reconstruction requires a minimum of 10 PMT hits[10]. The DAQ requirements for Gd-H_2O (subject to change for Gd-WbLS) have been defined nominally as follows:

Single Photoelectron Threshold	0.25 p.e.
Time Resolution	1.0 ns
Dynamic Range	12 bits
Trigger Threshold	6 hits in 100 ns
Maximum average PMT noise rate	10 kHz

[10]This threshold is also currently used for the new CoRe reconstruction (see Section 5) but requirements with CoRe could be lowered on the basis of detailed study.

The cosmogenic muon rate (see Section 4.3) in the detector is expected to be $\mathcal{O}(0.1)$ Hz and each muon event will result in $\mathcal{O}(10^4)$ detected photoelectrons in the PMTs. The DAQ system must have sufficient capacity to support the volume of data from such an event and to minimise saturation from the high number of photoelectrons expected, for the purposes of accurately reconstructing muon events.

3.11.1 Triggering and Digitisation

The DAQ itself must incorporate waveform digitisation for optimal flexibility in triggering and analysis of the data. There are three options being considered:

- TARGET (TeV Array Readout Electronics with GSPS sampling and Event Trigger) chip [107], with custom ASICs (Application-Specific Integrated Circuits) for triggering and for sampling and digitisation.
- ANNIE/KOTO DAQ [81] scalable, modular system of 500 MSPS (mega samples per second) commercial analog-to-digital converters (ADCs).
- CAEN commercial off-the-shelf option: 500 MSPS ADCs

Custom triggering electronics will be installed in a single, central downstream location.

The full waveform digitisation as well as the intrinsic modularity of the above designs mean that the DAQ system would also be compatible with higher-resolution photosensors such as LAPPDs.

3.11.2 Readout, Event-Building and Monitoring System

Triggered data from the DAQ will be fed into the readout, event-building and monitoring (REBAM) system. This system can be divided into two parts: data readout and monitoring. Data readout will build events from the data and perform fast reconstruction for data quality evaluation prior to physics monitoring and storage. *Slow controls* will monitor the DAQ system, environment, high voltage and detector functioning and will be used to control data and calibration runs and feed calibration constants back into the DAQ system.

The REBAM infrastructure will use commercial, off-the-shelf solutions and will be housed underground, with sufficient local storage for 1 week of data in the event of communications failure.

Data Readout and Event Building

The main function of the data readout is to implement DAQ hardware control and to extract and pre-process data from the DAQ hardware for storage and offline analysis. The pre-processing infrastructure will include event building, data formatting and fast event reconstruction.

Real-time, continuous pulse extraction of digitised waveforms from single-photoelectron signals will reduce both the data flow from the digitisers and the downstream computational load to ensure the system can cope with the expected trigger rates. A pulse-extraction

algorithm will identify single-photoelectron signals and extract their time of arrival and integrated charge. More complex waveforms will be passed downstream of the data readout for more sophisticated processing.

The readout and event-building software will be the interface between the DAQ hardware and the physics-monitoring infrastructure. As such, it will perform a fast physics reconstruction and will collect physics-monitoring metrics including event rate, position distribution and PMT occupancy for use in downstream physics monitoring.

Slow Control and Physics Monitoring

The detector monitoring or slow controls will interface with all detector subsystems and collect metrics which will be fed back to operators for validation of detector operation and data quality. Warning systems and fast-response procedures will be invoked in the event of sub-standard operational conditions in the detector.

Control of the calibration system with external trigger inputs will support detector calibration. Triggered calibration devices in the detector may require temporary changes to the trigger threshold and readout time windows. Untriggered calibration sources will require calibration-event tagging.

The data readout and pre-processing will feed into the physics-monitoring system. This will quantify the physics data in terms of the monitoring metrics.

For the purposes of detecting supernovae, the WATCHMAN collaboration will implement a supernova burst trigger system.

3.12 Contained Underwater Photosensor and Pulse-Extraction Apparatus Option

An alternative to the baseline PMT housing is the CUPPA (Contained Underwater Photosensor and Pulse-Extraction Apparatus) [108] fully encapsulated module. In the baseline design, the PMTs will be mounted in the plastic moulded mounting with the possibility of additional acrylic enclosures, both of which are described in more detail in Section 3.6. In this case, the signal would be transmitted via cables from the back of the PMT to the data acquisition system in the electronics racks on top of the tank.

Inside the CUPPA, the PMT would have its own 1-2 kV low-ripple HV supply and electronics. The electronics would perform triggering, pulse shaping and digitisation, as well as calibrations and self tests. Optical gel would be used to fill any gaps in transition regions e.g. between acrylic and glass. Digitisation would be performed by either one or two 250-MSPS commercial ADCs. The trigger, dataflow, on-board calibrations and self tests would be performed by custom firmware on a commercial FPGA (Field-Programmable Gate Array). This would be able to extract timing and amplitudes from single-photoelectron signals, which could also seed algorithms for multi-photoelectron signals on a microcontroller unit (MCU). Timing and amplitude could also be extracted by the MCU from more complex waveforms due to throughgoing cosmogenic muons. Power and data could be transmitted by a Category 6 twisted-pair Ethernet cable connected either above the detector or underwater, close to the CUPPA module, and then to PoE (Power over Ethernet) switches on top of the detector.

Full encapsulation in CUPPA modules would simplify detector assembly, avoid materials compatibility issues, reduce radon transmission into the inner volume and remove the need for HV cables in the detector volume. A CUPPA module could also incorporate any of the light collection enhancements and developing technologies which are being considered. Some modules could additionally incorporate remotely controlled calibration light sources or other devices.

3.13 Water Purification and Monitoring

Water purification systems exist independently for Gd-H_2O and WbLS and a purification system for Gd-WbLS is under development for NEO.

The water purification system in the NEO tank will be based on the EGADS [79] system. This system incorporates a *molecular band-pass filter*, which selectively filters out different sized particles, allowing the gadolinium sulphate to be extracted from the water and then returned to the tank once the water has undergone further filtering to remove impurities. With very little maintenance required, the EGADS experiment:

- demonstrated the ability to purify Gd-H_2O to ultra-pure water levels,
- established the process for large-scale loading of gadolinium,
- saw no loss of gadolinium sulphate over 650 passes of the entire water system and
- verified that the gadolinium did not degrade the detector materials with years of loading with 0.2% gadolinium sulphate.

The EGADS system was designed to be modular and is scalable to larger detectors with the addition of more filtration elements. This has been demonstrated by the doping of the Super-K detector with 10% of its final planned loading of gadolinium. Because impurities tend to enter the detector medium from the walls and PMT structure, the required re-circulation rate scales with the ratio of the detector surface area to the volume.

Gd-H_2O passing through the full filtration system will undergo a series of processes. The first of these are sediment filtration (5 μm) and UV sterilisation and microfiltration (0.2 μm) to kill and remove bacteria. At this stage, prior to passing into the main, band-pass filtration system, the Gd-H_2O passes into a chiller. Gadolinium sulphate is more soluble at lower temperatures but some of the filtering elements increase the water temperature. Reducing the temperature at this point ensures that the gadolinium sulphate remains in solution throughout the filtering process. Reducing the temperature has the added benefit of arresting/slowing bacterial growth in the filtration system and detector. The band-pass filtration system incorporates the following stages:

1. Ultrafilter ($\mathcal{O}(0.01)$ μm) to remove suspended solids.
2. Series of nanofilters ($\mathcal{O}(0.001)$ μm) to remove Gd (in the form of concentrated Gd-H_2O).
3. Total organic compounds (TOC) photooxidation lamp to destroy organic chemicals.

4. De-ionisation, microfiltration and reverse osmosis to remove all remaining impurities.

Recombination of ultra-pure filtered water with Gd-H_2O in a buffer tank is followed by further UV sterilisation, ultrafiltration, membrane degasification (to remove dissolved air), TOC photooxidation and deionisation (e.g. anion exchange resin to remove uranium) before the ultra-pure Gd-H_2O is returned to the tank.

An additional, *fast re-circulation system* consisting of elements which do not remove gadolinium (ultrafiltration, UV and TOC lamps, deionisation), equipped with a heat exchanger to keep temperatures low, provides additional cleaning power and runs in parallel with the full band-pass filtration.

To purify Gd-WbLS, the oil solvent containing the WbLS would be separated out from the Gd-H_2O at the nanofiltration stage. The Gd-H_2O would undergo the purification process described above and the oil would be directed to an additional filtration module for scintillator purification, which would include the following components of the kind described in [109]:

- Exchange/assay columns with silica-gel (metal scavenger) adsorption and/or solvent-to-solvent water extraction to remove uranium, thorium, radium, potassium and lead (and daughters where applicable) and suspended, ultrafine particles.
- Gas (air/nitrogen) stripping column to remove volatile impurities such as radon, argon, krypton and oxygen.

Purified scintillator would be re-combined with the Gd-H_2O and would pass through the final purification stages before being returned to the detector.

The water purification system and associated infrastructure includes pre-treatment, tanks for mixing the Gd-H_2O or Gd-WbLS and additional cylindrical tanks for the storage of concentrated Gd-H_2O (and liquid scintillator where applicable) for the purpose of detector re-fills. The whole facility including the water purification system, mixing tank and storage tanks will be located together close to the main detector.

Water transparency as a measure of water purity is likely to be monitored in real time using a system similar to the Underground Device Evaluating Attenuation Length (UDEAL) [79], which was used to monitor water quality in EGADS. It would be positioned either vertically in or just outside the outer detector or horizontally in a nearby drift. The UDEAL system consists of:

- a light injector at the top for laser beams with varying wavelengths,
- a vertical pipe and
- integrating spheres at the top and bottom of the pipe.

The light injector is configured to span a range of wavelengths in the Cherenkov spectrum and photomultiplier tube sensitivity range. The integrating spheres are diffuse reflectors which remove directional information, to minimise the effects of the beam jitter, while preserving the intensity. UV-enhanced silicon photodiodes in the walls of the spheres detect the diffuse light from which the measurements are made.

Upon injection, the laser beam hits a beam splitter and is directed towards the integrating sphere at the top, which monitors fluctuations in the beam intensity, and down the pipe towards the integrating sphere at the bottom. The integrating sphere at the bottom monitors the intensity of the beam after attenuation in the water in the pipe.

Water can be taken from different positions in the detector to monitor the variation in water quality through the detector. Measurements are taken for each wavelength at several heights and the absolute attenuation can be calculated for each wavelength from these measurements. A weighted average attenuation across all wavelengths measured can then be calculated as a metric for the water quality.

3.14 Calibration

Deployable calibration sources will be suspended into the tank in various locations inside and outside the fiducial volume via gloveboxes (one fixed, one mobile) through calibration ports located in the top deck. Gloveboxes, purged with nitrogen prior to the opening of the port in the case of the mobile glovebox, will ensure that there is no ingress of mine air into the tank during deployment. Calibrations will be carried out in multiple locations to perform calibration of the detector response as a function of position, particularly near the edge of the detector where variations are expected to be most significant.

3.14.1 Radioactive Sources

Gamma-ray and neutron sources will be deployed to determine the energy scale and neutron detection efficiency.

Three gamma-ray sources are to be deployed in NEO:

- ^{232}Th: 2.6 MeV gamma
- Americium-beryllium (AmBe): 4.4 MeV gamma
- Tagged ^{16}N: 6.1 MeV gamma

These three sources span the IBD positron spectrum and thus provide a method to determine the linearity of the detector response over the range. They can also be used to calibrate event position in the detector, which is used to define the fiducial volume.

AmBe and ^{16}N sources are particularly useful as calibration sources as they both have a coincident event that enables discrimination of the calibration gamma ray from a background source.

An AmBe neutron-gamma source contains ^{241}Am and ^{9}Be. The decay of ^{241}Am emits alpha particles which induce the following nuclear reaction:

$$^{9}\mathrm{Be} + \alpha \rightarrow \mathrm{n} + ^{12}\mathrm{C}^{*}$$

The subsequent decay of the excited ^{12}C, via the emission of a 4.4 MeV gamma ray, is coincident with the capture of the neutron on Gd and emission of the characteristic 8 MeV gamma cascade.

The tagged AmBe source will be used to calibrate the neutron detection efficiency. Detection of the gamma ray e.g. via a cerium-doped YAP (yttrium aluminium perovskite $YAlO_3$) scintillator crystal coupled to a PMT [110] provides a tag for the neutron capture, allowing the neutron detection efficiency to be measured directly. The relative isotropy of the neutron capture to that of the single gamma-ray event may also be measured using the AmBe source.

A ^{16}N source uses a deuterium-tritium (DT) generator to produce neutrons with energies of 14.2 MeV. These produce ^{16}N via the reaction

$$^{16}\mathrm{O} + \mathrm{n} \rightarrow ^{16}\mathrm{N} + \mathrm{p}.$$

^{16}N has a half life of 7.1 s and undergoes beta decay with the emission of a 4.3 MeV electron and a 6.1 MeV gamma ray. A currently proposed design, based on the design used in SNO [111], is to create ^{16}N through the irradiation of CO_2 with the neutrons from the DT generator and transfer it to a small, scintillator-lined decay chamber inside the detector volume. This is designed to allow the 6.1 MeV gamma rays to exit but to detect the coincident β-particles via scintillation light incident on a PMT within the decay chamber. The ^{16}N tagging is particularly strong, enabling each gamma ray to be identified. This makes it a particularly powerful calibration source for the energy scale and energy scale position dependence and can also be used to verify the detector energy resolution as well as event position in a detector and reconstruction algorithms (see Section 8.1).

3.14.2 Light sources

Permanently fixed LEDs and lasers will also be used for PMT gain and timing and attenuation/scattering calibration. LED sources can be permanently situated at various points in the detector to provide degeneracy in PMT illumination for PMT gain calibrations. Automated calibration of PMT gain may be achieved by regular illumination of certain LEDs over the course of detector operation. Laser fibres, for PMT timing and water attenuation and scattering measurements can also be mounted in fixed positions throughout the detector.

A light injection system able to produce diffuse, isotropic light as well as a narrow beam of light is in development. Deployed in the centre of the detector via a glovebox to emit light isotropically, such a device could be used to perform PMT timing and charge calibrations.

3.15 Summary

- The planned NEO detector at the AIT site in the low-background environment of Boulby Mine will aim to detect antineutrinos from the reactor cores at Hartlepool Power Station and further afield.
- The WATCHMAN collaboration will address reactor verification and discovery scenarios for non-proliferation.
- The project science goals include supernova neutrino detection and neutrino oscillation measurements.

- Two cylindrical detector designs are being considered - a 16 m high by 16 m diameter detector with an inner PMT radius of 5.7 m and a 22 m high by 22 m diameter detector with an inner PMT radius of 9 m.
- Two detector fills are under consideration - gadolinium-doped water and the more nascent gadolinium-doped water-based liquid scintillator.
- The detector will be a 316L stainless steel insulated tank with multiple service and calibration ports and a 316L stainless steel inner PMT structure.
- The detector will be instrumented with Hamamatsu 10" R7081 low-radioactivity, high quantum efficiency photomultiplier tubes (PMTs), which will undergo a series of quality assurance tests and calibrations prior to installation. A sample of one hundred PMTs has already undergone acceptance testing.
- A number of light collection enhancements and alternative photodetection options are being considered, including an encapsulated and fully-contained PMT and pulse-extraction system.
- Water or water-based liquid scintillator purification will be carried out via an EGADS-style molecular bandpass filtration system and calibration systems will include both radioactive and light sources.

The design for the detector has not yet been finalised. Assessing the capabilities of the NEO detector and optimising the design for the application of antineutrino detection to remote monitoring requires a thorough understanding of the signal and background in the detector. The signal and backgrounds as they are currently understood are described in the following chapter.

4 Signal and Backgrounds in NEO

What is repugnant to every human being is to be reckoned always as a member of a class and not as an individual person.

– Dorothy L Sayers, *Are Women Human?*

The *simulation* → *reconstruction* → *analysis* process begins with a thorough understanding of the signal and backgrounds in the detector. This chapter describes the signal and backgrounds as simulated for the purposes of this thesis. Section 4.1 describes the general simulation method used and the remainder of this chapter describes the signal and background simulations in more detail. For this thesis, pre-existing custom event generators were used to simulate events. The simulation of radioactive decays was modified as discussed in Section 4.3.6. Expected signal and background rates were calculated for the two detector sizes according to the methods set out in the relevant sections. Some of the material in this Chapter now forms the basis of a technical note for the WATCHMAN collaboration [112].

4.1 Simulations of Signal and Background in NEO configurations

Full detector simulations were carried out with RAT-PAC (Reactor Analysis Tool - Plus Additional Codes) [113], which has been adapted by the WATCHMAN collaboration [114] and which is based on the physics simulation framework GEANT4 [115] [116], the CLHEP physics library [117], the GLG4sim (Generic Liquid-scintillator Anti-Neutrino Detector or *GenericLAND*) Geant4 simulation for neutrino physics [118] and the data analysis framework ROOT [119].

RAT-PAC models the event-by-event detector response to the signal and background. Events are produced for NEO with custom GLG4sim Monte Carlo event generators and particles are propagated in the detector medium with GEANT4. Light emission and PMT response is managed by GLG4sim. RAT-PAC also handles the triggering and data acquisition before the data are output in ROOT format.

For each of the detector configurations simulated, the large-scale, complex structures (*e.g.* I-beams, trusses, PMT support structure and tank) were simplified to approximate position and volume. Each detector was set within a cavern surrounded by a layer of rock with 2 m thickness. This depth of rock is considered sufficient, in combination with a detector buffer region of 2 m or more, given that the total neutron flux attenuation is in the region of two orders of magnitude beyond a distance of 3.5 m in rock or water [120] and very few backgrounds due to radioactivity penetrate into the detector fiducial volume from the rock. Consistent with previous studies, backgrounds from the rock were generated in the inner 10 cm of this layer. The outer detector is assumed to be a passive buffer and, as such, PMTs have been simulated only in the inner detector.

Event generators have been configured by WATCHMAN for NEO to produce the initial signal and background particles with spectral and angular distributions as described in the remainder of this chapter.

4.2 Signal

The primary signal for the NEO detector may be the twin core reactor at the Hartlepool Power Station (26 km from Boulby), which is now operated by EDF Energy. The complex consists of two cores rated with a maximum power output of 1.575 GW_{th} each. The power station met ~2.9% of the UK's 284.4 TWh (terawatt hours) electricity demand in 2020 [121] [122].

The Hartlepool Power Station is currently expected to run until 2024. The lifetime of a reactor can be extended if it is considered safe to do so but this is not guaranteed. Beyond the closure of Hartlepool, the WATCHMAN collaboration can focus on the detection of antineutrinos from the more distant reactors at Heysham (149 km) and Torness (187 km). The Heysham Nuclear Power Station, also operated by EDF Energy, consists of two reactor complexes (Heysham 1 and Heysham 2), which each have two cores rated at ~1.5 GW_{th}. Heysham 1 is the same type of reactor as the Hartlepool reactor, both being the first type of AGR to be built in the UK (AGR-1). Like Hartlepool, Heysham 1 is currently scheduled to shut down in 2024. Heysham 2 is of the same, later AGR design (AGR-2) as the Torness Nuclear Power Station. Heysham 2 and the Torness reactor complex are currently due to shut down in 2030.

As an AGR (see Section 2.1), the Hartlepool reactor has a CO_2 coolant. It is fuelled with 110 tonnes of pellets of ~3%-enriched uranium dioxide. The fuel elements are enclosed in graphite blocks or *channels* (Figure 16(a)), which compose the core and provide a network of cooling and fuel channels in the core (Figure 16(b)). Each channel has a fuel assembly of eight 1-metre high elements containing 36 stainless steel fuel pins which contain the fuel pellets [123] [124]. Each core at Hartlepool has 332 graphite channels to give a thermal capacity of 1.575 GW_{th} from 110 tonnes of uranium oxide.

The power stations are usually run continuously, such that one core (or more) is kept running while another is shut down for refuelling. The refuelling process is undertaken twice per year and takes around two weeks per core. EDF are able to provide the WATCHMAN collaboration with information about the assembly types, hourly reactor power variations, daily mean channel power and daily mean channel irradiation. This information is key in building an accurate model of the reactor antineutrino flux, as are the fission fractions of the radioactive isotopes in the cores, particularly for more complex spectral analyses.

The WATCHMAN collaboration currently uses an emitted antineutrino spectral flux calculated as per Equation 2.1 using monthly-averaged load factors and fission fractions estimated at the midpoint of the fuel cycle and summed over the main fissile isotopes in the core [60]. Detailed studies are underway to calculate a precise spectral flux using reactor power over a fuel cycle and more detailed fission fraction data provided by EDF Hartlepool. Due to the relatively short fuel cycle, the variation in the antineutrino flux through the cycle is likely to be $\mathcal{O}(1\%)$.

Some of the emitted antineutrinos are converted to μ and τ flavour as they travel, which reduces the detectable antineutrino spectrum. The emitted spectrum is convolved with the electron antineutrino vacuum survival probability. Matter effects, which are expected to be small at reactor antineutrino energies and the distances under consideration, are neglected and zero CP violation is assumed (see Equation 1.14).

To obtain the detectable spectrum, the oscillated antineutrino emitted spectrum is corrected for the distance from the reactor and convolved with the IBD cross section. The tree-level IBD

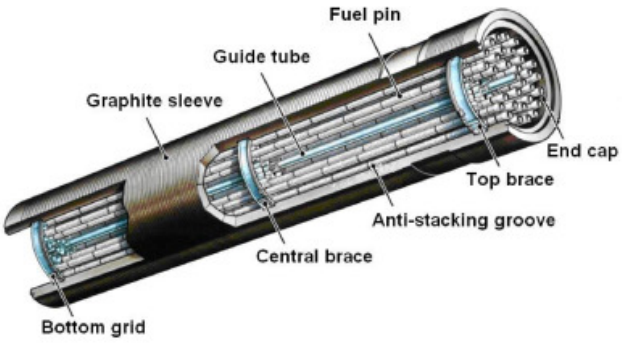

(a) AGR fuel element consisting of 36 fuel pins which contain the fuel pellets and are enclosed in a graphite sleeve. Each fuel pin is around 1 m long. The hollow fuel pellets are fed into the pins, which are filled with helium before the end caps are welded on.

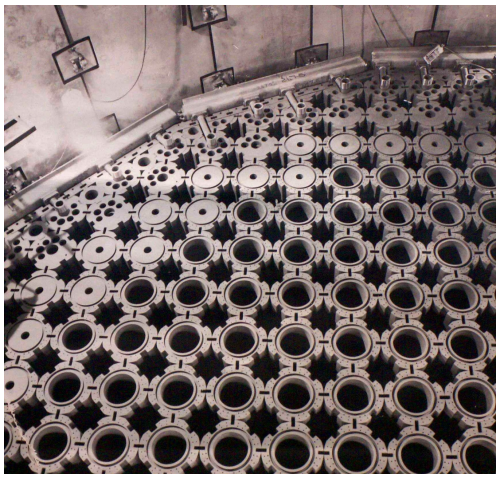

(b) Partially complete AGR graphite core composed of graphite blocks of the type used in Heysham 2 and Torness. These hold the fuel elements, act as a moderator and provide channels for the gas coolant.

Figure 16: Structure of typical AGR core.

Figure (a) reprinted from Progress in Nuclear Energy, Volume 109. James B.Goode, David Harbottle, Bruce C.Hanson, *Vacuum drying of advanced gas reactor fuel*, Pages 145-158, Copyright (2018), with permission from Elsevier. Figure (b) courtesy of EDF Energy.

cross section is given by:

$$\frac{d\sigma_{IBD}}{dE_e} = 2m_p \frac{G_F^2 \cos^2\theta_C}{2\pi(s - m_p^2)^2} |\mathcal{M}^2| \tag{4.1}$$

with

$$|\mathcal{M}^2| = A(t) - (s-u)B(t) + (s-u)^2 C(t), \tag{4.2}$$

where E_e is the positron energy, G_F is the Fermi coupling constant, θ_C is the Cabibbo quark mixing angle, m_p is the proton mass and $\mathcal{M}$ is the matrix element such that $|\mathcal{M}^2|$ gives the probability of the transition from initial to final state. Additionally, $s = (p_{\bar{\nu}_e} - p_p)^2$, $t = (p_{\bar{\nu}_e} - p_e)^2$ and $u = (p_{\bar{\nu}_e} - p_n)^2$ are the Mandelstam variables with momenta $p_{\bar{\nu}_e}$, p_p, p_e and p_n of the electron antineutrino, proton, positron and neutron respectively. A, B and C are functions of the particle masses and Mandelstam variable t, and are detailed in [74].

The cross section can be expressed more simply as:

$$\sigma_{IBD}(E_e) = 10.0 \times 10^{-44}\ p_e\ E_e\ E_{\bar{\nu}_e}^{\alpha}\ cm^2, \tag{4.3}$$

where $E_{\bar{\nu}_e}$ is the antineutrino energy and $\alpha = -0.07056 + 0.02018 \ln E_{\bar{\nu}_e} - 0.001953 \ln^3 E_{\bar{\nu}_e}$ [74].

Figure 17 shows the total theoretically detectable antineutrino spectra at Boulby from [60] for the three reactor sites. The oscillated spectrum of emitted antineutrinos, weighted by average AGR fission fractions and using average thermal power reported by the IAEA for the year 2019,

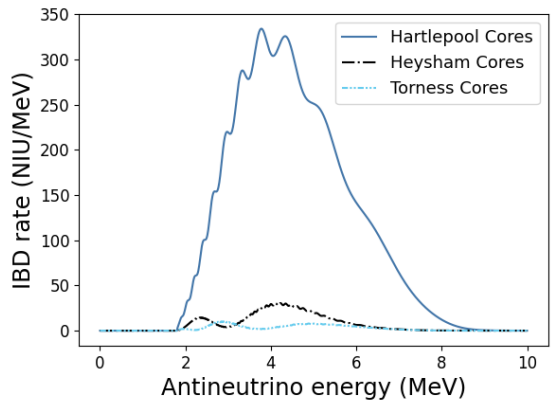

(a) Detectable IBD spectra.

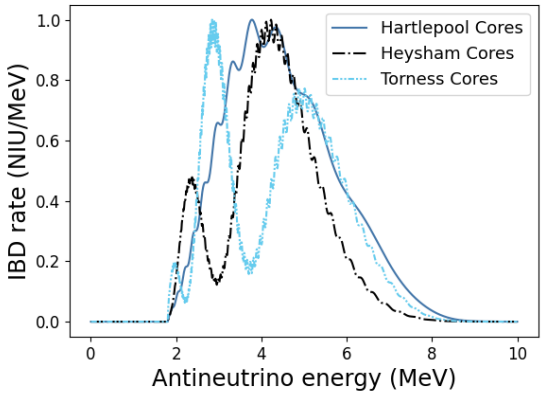

(b) Normalised spectra showing effect of oscillations.

Figure 17: Detectable IBD signal spectra (left) and normalised spectra (right) at Boulby for the three reactor sites at Hartlepool, Heysham and Torness. The emission spectrum is convolved with the neutrino oscillation probability and the IBD cross section to produce the detectable antineutrino spectrum at Boulby. NIU (Neutrino Interaction Units) is the interaction rate per 10^{32} protons per year. 1 ktonne H_2O contains 6.686×10^{31} free protons.

has been convolved with the IBD cross section given in [74] and adjusted for neutrino oscillations.

In general, the sensitivity of a detector to an antineutrino interaction increases with energy. Neutrino oscillations change the spectrum of detectable antineutrinos from a reactor so that, for example, a greater or lower fraction of the antineutrinos interacting in the detector are above the detection threshold. In Figure 17, the effect of oscillations is to increase the fraction of antineutrinos above the threshold from Heysham and increase the fraction at or below the threshold from Torness.

Oscillations notwithstanding, much of the associated positron spectrum (shown in Figure 18 for the Hartlepool reactor) is below previously achieved energy thresholds in water Cherenkov detectors (~3.5 MeV in Super-Kamiokande [76]). Doping with gadolinium will reduce the detector threshold due to improved background rejection and complementary improvements in the technology, triggering, reconstruction and analysis are, at least in part, aimed at attaining very low-energy sensitivity, ideally below the signal peak at ~2.5 MeV. In particular, improvements in the reconstruction and analysis, demonstrated through simulation and discussed in Chapters 5 and 6, indicate that the energy threshold could be lowered with a process adapted to search for IBD in a gadolinium-doped medium.

The IBD generator developed for WATCHMAN draws antineutrino energies from the relevant signal spectra input from [60] as described above. The positrons are slightly back-scattered overall and the angles between the positron and antineutrino directions are weighted by the $\cos\theta$-dependent cross section. The positron and neutron energies and momenta are calculated to the first order correction for finite nucleon mass as set out in [73].

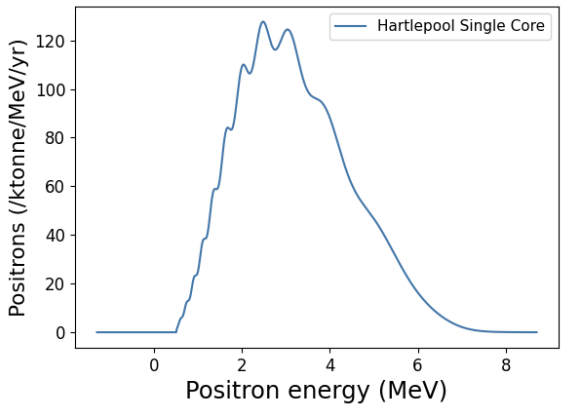

Figure 18: Detectable positron signal spectrum at Boulby due to reactor antineutrino emission from a single core at Hartlepool. Positrons above the Cherenkov threshold for positrons of ~0.8 MeV are in theory detectable in a water Cherenkov detector. In practice, much of the detectable positron spectrum is likely to be below the detector threshold. Advances in technology (e.g. Gd loading and WbLS), reconstruction and analysis aim to open up more of the spectrum to detection.

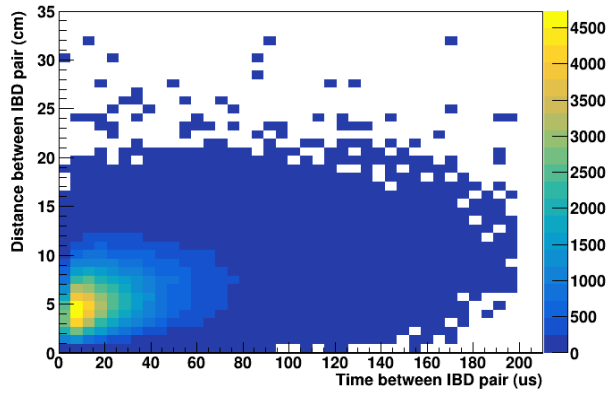

Figure 19: IBD positron-neutron pair time (dT) and spatial (dR) coincidence in Gd-H_2O from simulation with RAT-PAC in the baseline 16 m cylinder. The mean time between the pair is 28μs and the mean distance between the pair is 6 cm. For most of the IBD events, dT < 100 μs and dR < 10cm.

Figure 19 shows the distributions of time and distance between the positron and neutron in IBD pairs from the RAT-PAC simulation. The mean time between the positron and neutron in Gd-H_2O is ~28 μs and the mean distance is ~6 cm.

The integrated IBD rate R_{NIU} in interactions per 10^{32} targets per year (NIU) is calculated from the antineutrino spectra [60]. These are then converted to a total annual detectable rate as follows:

$$R_{IBD} = R_{NIU} \div 10^{32} \times N_{ktonne} \times N_p \tag{4.4}$$

where N_{ktonne} is the mass of water in the detector and $N_p = 6.686 \times 10^{31}$ is the number of free protons in 1 ktonne of water. The IBD signal rates at Boulby are given in NIU in Appendix B.

4.3 Backgrounds

The overall background rate is significantly reduced by requiring that a pair of prompt and delayed events must occur within the time and distance characteristic of the correlated IBD positron and subsequent Gd neutron capture signal. However, there remain correlated and uncorrelated background events which can mimic the correlated signal.

Correlated backgrounds are cosmogenic muon-induced radionuclide and fast neutron backgrounds and IBD backgrounds from other reactors and geoneutrinos. Uncorrelated backgrounds are radioactive decays of isotopes which occur naturally in the detector and its environment. These can happen to occur in coincidence with each other and these coincident, uncorrelated events are referred to as *accidental coincidences*. Illustrative raw event rates for the 16 m baseline detector for each type of background are shown in Table 4. β decays from ambient radioactivity are by far the most numerous background event before data reduction. A breakdown of the rates of these backgrounds is given along with the signal rates in Appendix B.

Table 4: Illustrative rates of each type of background in the 16 m detector.

Type of background	Rate (Hz)
Cosmogenic radionulide β-n decays	2.0×10^{-5}
Cosmogenic fast neutrons	1.9×10^{-2}
Reactor IBD (beyond Hartlepool)	2.2×10^{-6}
Geoneutrino IBD	3.6×10^{-7}
Ambient radioactive isotope β decays	3.4×10^{5}

Spontaneous fission of ^{238}U and ^{232}Th in the detector medium creates a correlated background of a prompt gamma ray flash and delayed events due to the capture of multiple emitted neutrons. By rejection of events where there are more than two interactions, the number of events is expected to be subdominant ($\mathcal{O}(1)$ per year) [125] and this type of background has been neglected for the purposes of this thesis.

4.3.1 Correlated Cosmogenic muon-induced backgrounds

Cosmogenic muons are created principally by the interaction of primary cosmic rays with nuclei in the atmosphere and subsequent decay of the pions generated:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$
$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$$

At high energies, kaons also contribute to the production of cosmogenic muons.

Throughgoing muons can induce backgrounds by interacting both in the material outside the detector (e.g. rock and any shielding such as lead) and inside the detector.

Muons interacting outside the detector - in the rock in the case of NEO - induce fast neutrons by muon spallation and also neutron evaporation along a muon track. It is common for these neutrons to be produced in multiplicity, which can generate neutron pairs in the detector volume. Fast neutron pairs in the detector volume thermalise and capture, producing a detector response which mimics the reactor anti-neutrino signal.

Muons interacting inside the detector with the water or water-based target can produce unstable beta-neutron emitters with half-lives of a few seconds in the detector volume. The decay of these long-lived radionuclides, such as ^{9}Li, generates correlated signals which have a time interval and energies similar to the antineutrino signal.

Locating a detector underground reduces the cosmogenic muon flux. At Boulby the flux is reduced significantly - by a factor of almost 10^8 [120]. Thanks to its flat overburden, this is an order of magnitude greater reduction than in the Kamioka mine where Super-K is located, which is at a similar depth but has a mountainous overburden.

Muogenic backgrounds can be mitigated by the inclusion of a veto volume between the tank wall and the PMT support structure in the detector. A passive buffer is effective in reducing the fast neutron background due to their rapid attenuation in water. Outward-facing PMTs to give an active veto region can be used to identify muons entering the tank and thus reduce both types of cosmogenic muon backgrounds.

WATCHBOY, a tonne-scale Gd-H_2O detector funded under the 2011 Strategic Plan, was designed to test sensitivity to radionuclide backgrounds [126] for the WATCHMAN collaboration. The MARS Multiplicity and Recoil Spectrometer was a high-energy neutron spectrometer fielded in the Kimballton Underground Research Facility (KURF) - under the auspices of the NNSA Strategic Plan - to measure the fast neutron rate at different depths. Similar fast neutron studies at Boulby e.g. [127] can be used to improve our simulation and analysis of the fast neutron background for the Boulby location and to reduce systematic uncertainties on the rates.

4.3.2 Fast Neutron Backgrounds in NEO

The theoretical spectral flux for fast neutrons between 0 MeV and 500 MeV in Boulby Mine is shown in Figure 20(a). The fast neutron spectrum is defined in [128] as:

$$\frac{dN}{dE_n} = A\left(\frac{e^{-0.7E_n}}{E_n} + B(E_\mu)e^{-2E_n}\right), \tag{4.5}$$

where A is a normalisation factor and $B(E_\mu) = 0.52 - 0.58e^{-0.0099E_\mu}$ is a function of the muon energy E_μ. An alternative definition is given in [120]. Fast neutrons from the rock below 10 MeV are not considered to be of concern for NEO as they are insufficiently penetrating. However, a significant number of fast neutrons have energies greater than 10 MeV, with an average neutron energy of 88 MeV and an average muon energy at Boulby of 264 GeV [120]. The fast neutron flux in the Boulby Mine at the cavern wall interface, calculated through a combination of simulation and experimental data for fast neutrons above 10 MeV [120], is:

$$\Phi_n(> 10\ \text{MeV}) = 1.11 \times 10^{-9}\ \text{cm}^{-2}\text{s}^{-1} \times \text{cavern surface area}.$$

Note that there is a very high uncertainty of 27 per cent on this flux due to uncertainties on the muon flux and neutron production rates.

Cosmogenic muons are strongly biased in the downward direction, particularly at the high energies required to penetrate underground to the depth of the cavern at Boulby. The muon angular distribution with respect to the Earth's surface at the detector location is described by the distribution given in [129]. Attenuation of muons in the Earth varies as a function of this muon angle with the distance travelled through the Earth and is also described in [129].

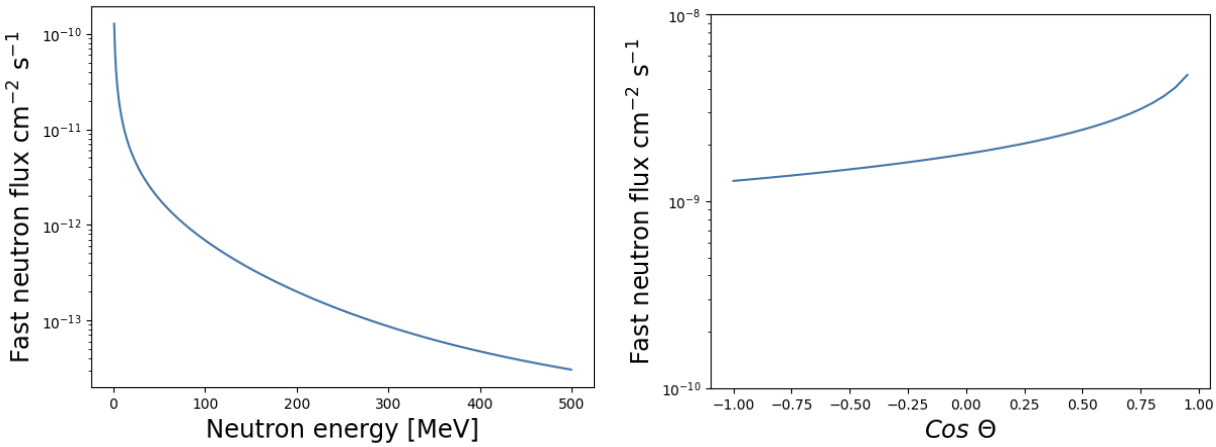

(a) Fast neutron spectrum at Boulby in the range from 0 to 500 MeV. Fast neutrons below 10 MeV are unlikely to penetrate into the detector from outside.

(b) Neutron angular distribution with respect to the muon direction for an average muon energy of 264 GeV. At neutron energies below 100 MeV, the distribution is more isotropic.

Figure 20: Fast neutron energy and angular distributions at Boulby from [128] and [120] respectively.

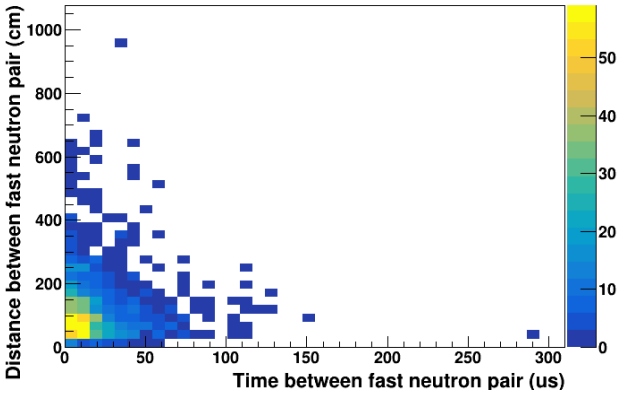

Figure 21: Fast neutron pair time (dT) and spatial (dR) coincidence in Gd-H_2O from simulation with RAT-PAC in the baseline 16 m cylinder. The mean time between the pair is 20 μs and the mean distance between the pair is 124 cm. Most of the fast neutron coincidences occur within a similar range to the IBD pair coincidences.

Fast neutrons from muon spallation are peaked in the direction parallel to the muon direction. Secondary neutrons from neutron evaporation (nuclear de-excitation with neutron emission) along the path of the muon are emitted isotropically [128]. The resulting neutron angular distribution with respect to the muon direction (Figure 20b) is peaked in the direction of the muon, with a flattening effect due to the secondary neutrons. The highest-energy neutrons tend to be emitted in the same direction as the muon but lower-energy neutrons are more isotropically distributed. The highest-energy, downward-going neutrons are most easily rejected. The more isotropic, lower-energy and secondary neutrons are likely to be more of a concern for NEO.

The mean time and distance between fast neutron events are 20 μs and 124 cm respectively. The time coincidence of fast neutron pairs is similar to that of the IBD event pairs (see Figure 21) and many of the fast neutron events occur within a distance similar to the spatial range of the IBD pair.

The neutron yield per muon is dependent on the muon residual energy and thus also on the depth, since higher-energy muons are more likely to penetrate deeper underground. It is additionally dependent on the material in which the muon is interacting. A hundred or more coincident neutrons may be produced by a muon in the rock at Boulby, although the majority of muons result in less than twenty coincident fast neutrons and the average expected multiplicity at Boulby is 6.03 per muon, given an assumed rock density of 2.7 g cm^{-3} [120].

A number of measures, both in detector design and at the analysis stage, can be taken to reject these backgrounds. Since attenuation in the veto region and identification of incoming muons in an active veto can reduce fast neutron backgrounds, optimisation of the buffer width and/or active veto are important processes at the design stage. Applying a time veto after a very high-energy event in the detector at the triggering or analysis stage can help to reject fast neutron events. Rejecting multiply coincident neutrons in the detector to ensure that there are no more than two coincident events offers further reduction. The isotropy and energy of the prompt event in a fast neutron pair can also be leveraged to reject fast neutron backgrounds. A more detailed examination of mitigation methods is presented and applied analytically in Chapter 6.

4.3.3 β-Neutron Radionuclide Backgrounds in NEO

Cosmogenic muons in the detector produce abundant daughter particles that, through successive spallation processes, generate electromagnetic and hadronic showers. Hadronic showers are the principal mechanism by which unstable but relatively long-lived radioisotopes are generated. These form the cosmogenic muon-induced radionuclide backgrounds at the low energies of reactor antineutrino detection, which have the same time and spatial coincidence as the IBD signal. It was found that, of the spallation-induced radioisotopes which are considered to be backgrounds in Super-K, only 11% are produced by muon spallation directly and of these only 4% are from throughgoing muons [130].

The radionuclide backgrounds in a water-based Cherenkov detector are due to spallation on oxygen in water. Table 5 shows the spallation-induced radionuclides which are potential backgrounds in NEO. The radionuclides of particular concern are ^{17}N and ^{9}Li due to their relatively high yields combined with β-n branching ratios as shown.

The total muon flux in Boulby mine is $\Phi_\mu = 4.09 \pm 0.15 \times 10^{-8}$ cm^{-2} s^{-1} [134]. The β-neutron decay rates of the radionuclides of concern are calculated according to Equation 4.6.

$$R_{iso}(s^{-1}) = R_\mu \times L_\mu(cm) \times Y_{iso}(\mu^{-1}g^{-1}cm^2) \times br \times \rho(g\ cm^{-3}) \times \left(\frac{E_{\mu,Boulby}}{E_{\mu,Super-Kamiokande}}\right)^{\alpha} \tag{4.6}$$

where $R_\mu(s^{-1}) = \Phi_\mu \times$ tank surface area, L_μ is the muon path length, Y_{iso} is the isotope yield, br is the branching ratio, and $\rho = 1$ g cm^{-3} for water. The muon path length is taken to be the vertical height of the detector in the calculation of the yields in [130], which makes the simplifying assumption that the muons are all downward going. Rates for NEO are calculated for the full detector volume since simulations are performed in the full volume and the muon path length is taken to be the full height of the detector. There is a depth-related correction for

Table 5: β-neutron backgrounds. Isotope yields in water calculated with FLUKA [130] for Super-K. Principal beta endpoint energies with approximate branching ratios taking into account β-n branching ratio taken from [131] [132] [133].

Isotope	Half life (s)	Total yield ($10^{-7}\mu^{-1}g^{-1}cm^2$)	Production process	β endpoints (MeV)
^{17}N	4.173	0.59	^{18}O(n,n+p)	3.3 (50%), 4.1 (38%)
^{16}C	0.747	0.02	^{18}O(π^-,n+p)	4.7(84%), 3.7(16%)
^{11}Li (85% br)	0.0085	0.01	^{16}O(π^+, 5p + π^+ + π^0)	16.6(22%)*, 12.5(16%)*
^{9}Li (50.8% br)	0.178	1.9	^{16}O(π^-,α+2p+n)	11.2(29%), 10.8(12%)
^{8}He (16% br)	0.119	0.23	^{16}O(π^-,3H+4p+n)	5.3(8%), 7.5(8%)

* Total branching ratio

the average muon energy of E_μ^α where $\alpha = 0.73 \pm 0.10$ [120] since, on average, higher-energy muons will survive to greater depths.

The measured yield at Super-K for ^{9}Li was found to be around a factor of four less than expected at $0.50 \times 10^{-7}\mu^{-1}g^{-1}cm^2$ (without neutron tagging) and $0.51 \times 10^{-7}\mu^{-1}g^{-1}cm^2$ (with neutron tagging), with a systematic uncertainty of 18.1%. The measured yield of ^{8}He was consistent with zero and the 90% upper confidence limit of $0.9 \times 10^{-7}\mu^{-1}g^{-1}cm^2$ was adopted [135]. It was proposed that the large disparity between the theoretical calculation and the experimental measurement for ^{9}Li is due to the spallation ejecting many nucleons of the ^{16}O, since a large disparity was also seen in the case of ^{8}Li/^{8}B, whereas the difference between theory and experiment was within 2σ otherwise. For NEO, the experimental value for the ^{9}Li yield must be corrected for the depth-dependent average muon energy from 220 GeV in Super-K to 264 GeV in Boulby Mine.

To use the rates calculated using the experimental yields, the difference between the experimental and theoretical values should be incorporated into the systematic uncertainties. The uncertainty on the experimental rate for ^{9}Li is thus:

$$\frac{\sigma_{R_{iso}}}{R_{iso}} = \sqrt{\frac{\sigma^2_{Y_{iso}}}{Y_{iso}} + \frac{\sigma^2_{E^\alpha_\mu}}{E^\alpha_\mu}} \tag{4.7}$$

where the uncertainty on the yield is $\sigma_{Y_{iso}} = \left(\frac{|Y_{theor.} - Y_{exp.}|}{Y_{theor.}}\right)$ and the uncertainty on the muon flux is not included since no assumption of the muon flux was made for the experimental measurement. The 18.1% systematic uncertainty on the Super-K measurement would be a minimum given the discrepancy between theory and experiment.

As a cautious approach, the theoretical values have been used for the results presented in Section 7. The uncertainty on the theoretical rates is:

$$\frac{\sigma_{R_{iso}}}{R_{iso}} = \sqrt{\frac{\sigma^2_{\Phi_\mu}}{\Phi_\mu} + \frac{\sigma^2_{E^\alpha_\mu}}{E^\alpha_\mu}} \tag{4.8}$$

where the uncertainty on the muon flux $\sigma_{\Phi_\mu} = 0.15 \times 10^{-8}$. The uncertainty on the energy

dependence is :

$$\sigma_{E_\mu^\alpha} = \sqrt{(E_\mu^\alpha \ln(E_\mu)\sigma_\alpha)^2 + (\alpha E_\mu \sigma_{E_\mu})^2} \tag{4.9}$$

where $\alpha = 0.73$ and the uncertainty on the energy dependence $\sigma_\alpha = 0.10$. Uncertainty on the average muon energy σ_{E_μ} is due to uncertainties in the atmospheric muon energy spectrum and muon energy loss in the rock but is not highly correlated with the rock composition or precise depth [134] and so is assumed to be negligible for this work. It should be noted that although the calculated systematic uncertainty on the rate from the theoretical yield is very low (see Table 6 at the end of this chapter), adopting the theoretical yield remains a more cautious approach than using the experimental rates even with their higher associated systematic uncertainties.

The rates do not account for the differences in muon energy loss between the Super-K fiducial volume height and the NEO full detector volume height. This is assumed to be negligible since the energy loss is small compared to the muon energy and thus the interaction cross sections do not change significantly over the path length in the detector [130].

Correlations between radionuclide-like events and a muon track can be used to reduce the radionuclide backgrounds. Backgrounds from long-lived radionuclides with half lives on the order of seconds would survive a time veto after a muon event. With effective muon tracking, using the transverse distance from a muon track can help to reduce these backgrounds without introducing excessive dead time in the detector. Using a neutron cloud to locate a hadronic shower may similarly help. More detail is given in Section 6.2.4.

4.3.4 Reactor IBD Backgrounds

Neutrinos from reactor cores other than the signal core(s) produce a true IBD background which will not be distinguishable from the signal in NEO. The reactor background at Boulby will depend on the particular signal being observed. Figure 22 shows the reactor background at Boulby for a single-core (a) and twin-core (b) signal at Hartlepool. The background to the single-core signal includes the second core. The reactor backgrounds exceed the signal in the single-core case.

The spectra and rates are calculated in a similar way to the signal and summed over all other reactors.

4.3.5 Geoneutrino IBD Backgrounds

Geological electron antineutrinos (geoneutrinos) emanate principally from beta decays of the ^{238}U and ^{232}Th decay chains and ^{40}K in the Earth's crust and mantle. Where they have an energy over the IBD threshold, they can interact in a water-based detector to form a subdominant IBD background to reactor antineutrinos. Potassium may also capture an electron to form argon, which emits mono-energetic geoneutrinos [60] but these are not of concern, even as an uncorrelated background for NEO, due to the low β endpoint energy of 1.5 MeV and the relative suppression of the decay (11% branching ratio).

Like the reactor antineutrino spectra, geoneutrino spectra (also shown in the plots in Figure 22) are also taken for NEO from [60]. The geoneutrino spectrum for each isotope is estimated by

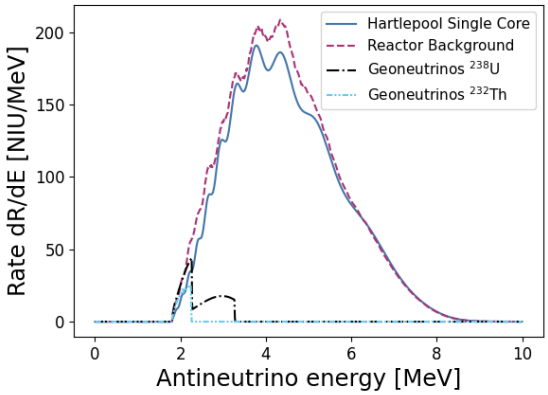

(a) Detectable IBD spectra for the single-core Hartlepool signal.

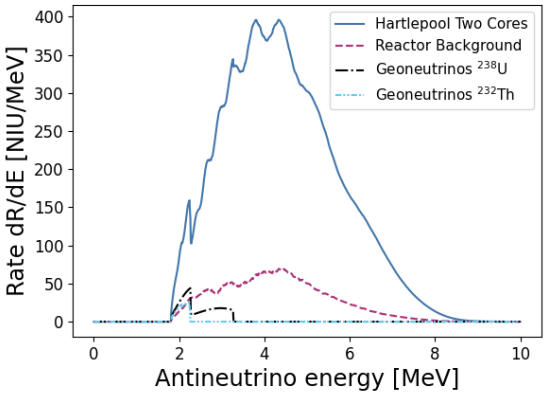

(b) Detectable IBD spectra for the twin-core Hartlepool signal.

Figure 22: Detectable background IBD spectra at Boulby for a single Hartlepool core (a) and both Hartlepool cores (b). Reactor backgrounds include IBD interactions from antineutrinos emitted by all nearby and world reactors excluding the signal reactor(s). Geoneutrinos are a sub-dominant IBD background for all signals investigated.

summing the spectra from each beta decay in the chain, weighted by its branching ratio. The geoneutrino IBD rates for NEO are calculated by integrating the product of the differential geoneutrino flux, IBD cross section, geoneutrino energy spectrum and neutrino survival probability over the spectrum. Pre-calculated geoneutrino surface fluxes are adjusted for variations in the density of the mantle and natural abundances of each isotope.

4.3.6 Uncorrelated backgrounds

Accidental coincidence of two physically independent interactions due to natural levels of radioactivity in the detector components and environment can mimic the time-correlated signal of IBD and neutron capture.

The rate of accidental coincidences depends on the composition of the detector (PMTs, detector medium, steel) and cavern in which the detector is situated (concrete, surrounding rock and air), the energy threshold and vertex position, vertex resolution, and the distance of the fiducial volume from the detector and cavern materials. Isotopes which decay with a non-negligible half-life and energy around or above the Cherenkov threshold are considered a background in NEO. Radioactive isotopes naturally present in the detector components and environs are shown in Table 14 in Appendix B, as is the nominal radioactivity currently adopted by WATCHMAN.

The principal contribution to accidental backgrounds comes from ^{238}U and ^{232}Th in the PMT glass (see Figure 23) and efforts are planned to mitigate this background, including using PMTs made with low-radioactivity glass. In addition, many of these events can be rejected by applying a fiducial cut, which rejects all events within a given distance from the PMTs.

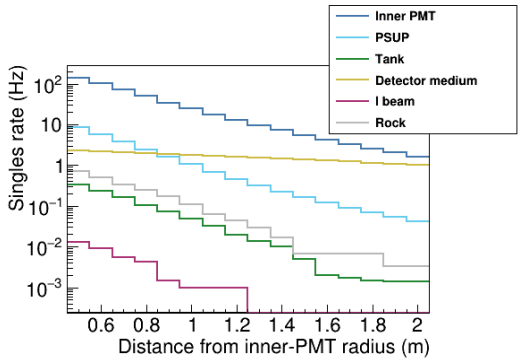

Figure 23: Single events due to natural radioactivity of components in and around the detector which can result in accidental coincidences and mimic the signal. Shows results for the 16 m tank with a Gd-WbLS fill. Relative rates in Gd-H_2O are similar.

Supported radon in the water, from the decay of ^{238}U in the PMTs and absorption from the mine air, plus other radioactivity in the water, tends to be the next most significant background, once a fiducial cut has been applied, which eliminates many of the contributions from the steel of the PMT support structure. The radon levels in Boulby Mine are comparatively low at 3 Bq m^{-3}. Nevertheless, the rate of radon diffusion from glass into water is not well known and these backgrounds occur throughout the detector volume, making them difficult to reject.

Only beta decays are considered. Alpha decay products are not relativistic and are therefore not visible in a Gd-H_2O Cherenkov detector. Similarly, they are not expected to contribute to backgrounds in Gd-WbLS due to their rapid attenuation.

While the beta decays of ^{40}K, ^{60}Co and ^{137}Cs are single decays, the other isotopes represent chains which include a number of beta decays (Table 15 in Appendix B). Where the branching fraction is below 0.1% and/or the endpoint energy is below 0.5 MeV, and as such well below the reasonably attainable detection threshold, the decay is not of concern for NEO.

Some of the beta decays shown in Tables 14 and 15 rarely pass low-level trigger or analysis cuts. As such, these have been neglected for the purposes of this thesis. Only uncorrelated backgrounds with a non-zero trigger efficiency and a singles rate $> 10^{-3}$ Hz are included in the final analysis. The trigger efficiency is the fraction of events which pass the DAQ trigger of 6 hits within 100 ns. Total rates due to isotopes in the detector components are shown in Table 16 in Appendix B.

The uncorrelated backgrounds were combined into a single simulation, interleaved at their respective rates, for the new simulation-reconstruction-analysis tool presented in this thesis.

Cleanliness requirements for the detector, in terms of limits on the radioactive content of the detector components, must be specified in the procurement of materials and equipment. A tool to asses the radiopurity requirements for NEO has been developed to help plan for this [136].

4.3.7 Systematic Uncertainties on Backgrounds

Systematic uncertainties on the reactor, geoneutrino, fast neutron and radionuclide backgrounds are given in Table 6. Uncertainties on the radionuclide backgrounds are also discussed in Section 4.3.3.

Table 6: Systematic uncertainties on raw background rates, not including detector effects.

Background	Uncertainty
Reactor	6% [60]
Geoneutrino	25% [60]
Fast neutron	27% [120]
Radionuclides	$< 1\%$

4.4 Summary

- Signal and backgrounds in the NEO detector will be coincident event pairs.
- The signal will consist of a coincident positron-neutron pair from the IBD interaction of reactor antineutrinos with protons in the detector medium.
- The primary signal used may be the Hartlepool Power Station but signals from the Heysham and Torness Power Stations could be targeted in the event of a Hartlepool shutdown.
- The detectable IBD positron spectrum peaks at around 2.5 MeV and is calculated by adjusting the reactor antineutrino spectrum for neutrino oscillations before convolving it with the IBD cross section.
- There are two types of correlated background. These are IBD events and cosmogenic muon-induced backgrounds.
- IBD backgrounds come from other reactors and geoneutrinos
- Cosmogenic muons will induce fast-neutron pairs and β-neutron decays of radionuclides in the detector.
- Uncorrelated backgrounds are due to accidental coincidences of independent, single radioactive decays of radioactive contamination in the detector components and environment.
- In some cases, there are large uncertainties on the background rates.
- Actual coincidences of signal and backgrounds are simulated in a Geant4-based simulation at their respective expected rates calculated from theory, experiment, or a combination.

The coincident signal and backgrounds have motivated the development of a method of reconstruction and a new analysis, which together form the focus of this thesis. These attempt to take full advantage of the coincidence of the signal to maximise sensitivity and achieve powerful background rejection. The new coincident-event reconstruction will be discussed in Chapter 5 and the newly developed coincidence analysis will be presented in Chapter 6.

5 Reconstructing the IBD heartbeat

In the space between chaos and shape there was another chance.
– Jeanette Winterson, *The World and Other Places*

This chapter describes the *combined* vertex reconstruction implemented specifically to capitalise on the coincidence of events in a gadolinium-doped detector medium, by reconstructing pairs of events together.

The signal and background events described in Chapter 4 are reconstructed from the timing of PMT-hit signals combined with the pattern they make on the walls of the inner detector in simulated or observed data. Although WATCHMAN currently relies on simulated data, the implementation of the vertex reconstruction can be readily adapted for use, ultimately, on observed data. Observed data will undergo pre-processing in the DAQ system's data-readout and event-building software (see Section 3.11) prior to reconstruction but will contain an inherent contribution from dark noise. This is currently added at the reconstruction stage for simulated data at a nominal rate of 3 kHz.

The accuracy of the vertex reconstruction is important for reducing systematic error on our definition of the fiducial volume, for background rejection and for reconstruction of the positron and therefore antineutrino energy. Improvements at lower energies in particular can improve overall sensitivity and help to lower the energy threshold of the detector.

In the reactor antineutrino energy range, the light yield from the IBD positron event is often very low, particularly with low photocoverage in the detector, and this can make vertex reconstruction unreliable. The light yield from the IBD neutron capture is generally higher and the distance of the neutron-capture vertex from the IBD interaction and positron vertex (<10 cm) is less than the current vertex resolution. The additional light from the neutron event can therefore be used to improve the reconstruction of the positron event.

Reconstruction in the standard WATCHMAN tools is achieved using BONSAI (Branch Optimisation Navigating Successive Annealing Iterations). BONSAI has been used for many years to perform low-energy reconstruction up to ~100 MeV for Super-Kamiokande [137]. It has since been implemented for WATCHMAN [138]. The combined vertex reconstruction constitutes an adaptation of BONSAI, which I have implemented and optimised for the detector configurations under consideration. Section 5.1 describes the BONSAI single-event reconstruction which forms the basis of the combined vertex reconstruction. Section 5.2 explains how the single-event reconstruction was adapted to perform a combined reconstruction. The remainder of the chapter details the *CoRe* [139] implementation of the combined reconstruction for WATCHMAN and its improvements to the vertex resolution and event selection/rejection.

5.1 BONSAI Vertex Reconstruction

BONSAI was originally written to reconstruct low-energy events from Cherenkov light in water Cherenkov detectors. It is a maximum likelihood fitter to the PMT hit timing. The likelihood is based on the hit time residuals of the Cherenkov signal in Gd-H_2O (or Cherenkov +

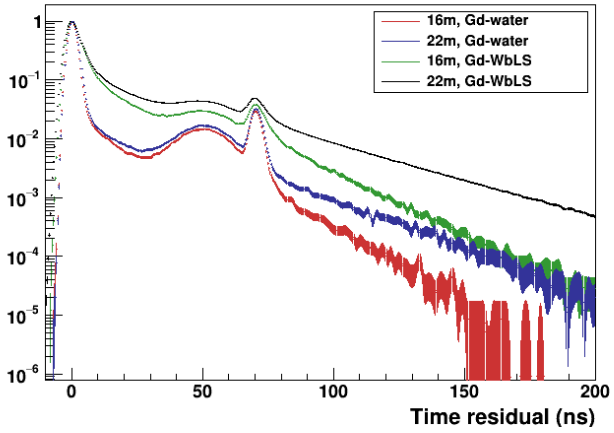

Figure 24: True hit-time residuals (calculated here from the MC vertex using simulation) are largely influenced by the PMT timing properties, which give the peaks at 0 ns, ~50 ns and ~70 ns, but include effects due to photocoverage, detector medium and scattering. Larger tank sizes result in a longer tail due to increased scattering. WbLS gives less defined prompt and second peaks due to the absorption and re-emission by the scintillator.

scintillation signal in Gd-WbLS) and dark noise background. It is calculated for a selection of test vertices and is given by:

$$\ln\mathcal{L}(\mathbf{x}, t_0) = \ln(\prod_{i=1}^{N} P(\Delta t_i(\mathbf{x}))). \tag{5.1}$$

The hit time residual $\Delta t_i(\mathbf{x})$ is:

$$\Delta t_i(\mathbf{x}) = t_i - tof_i(\mathbf{x}) - t_0 \tag{5.2}$$

where $\mathbf{x}$ is the test vertex, t_i is the hit time at the i^{th} PMT, t_0 is the emission time and $tof_i = |\mathbf{x_i} - \mathbf{x}|/c_{water}$ is the time of flight from the reconstructed vertex to the PMT vertex for hit i.

$P(\Delta t_i)$ is a probability density function (pdf) populated with the hit time residuals from true vertices. The timing residual pdfs fold in the effects of PMT timing features, photocoverage and scattering in the detector medium. In Figure 24, showing the timing residuals for the detector configurations under analysis, the shape is dominated by the PMT timing features discussed in 3.7.2 with the prompt peak at zero, the double-pulsing peaking at ~50 ns and the late-pulsing peak at ~70 ns. Scattering in increasing tank sizes results in increasing tails out to longer times. The addition of liquid scintillator in Gd-WbLS results in a wider prompt peak due to absorption and re-emission by the scintillator and consequently less well-defined prompt and double-pulsing peaks. The pdfs of time residuals are currently derived from simulation but will ultimately be calculated from calibration data.

For a given triggered event, hits are first passed through a hit selection criterion, which creates a list of hits that can be used to generate a sample of vertices which form the starting point for the likelihood maximisation. This is done by requiring that for any one pair of PMT hits separated by time Δt, the distance that could be travelled by direct light in the time between the hits ($\frac{c}{n}\Delta t$) is less than the distance between the two hit PMTs. This ensures that in principle the light is unscattered and could have come from the same interaction.

A minimum of four hits is required to reconstruct a vertex in 3D space. Sets of four hits are selected from the list of direct hits and used to define a test vertex by requiring that all four hit-time residuals are zero. In this way, each quadruple of hits defines a point in the detector

and a list of potential initial test vertices for the maximum likelihood vertex search is generated. Having more than one starting point helps to avoid mis-reconstruction due to local maxima.

Since the number of quadruples is proportional to the number of hits N^4, some limits are applied to increase speed. The number of quadruples is reduced by giving preference to clusters of hits in a variable time window. There is an optimal number of hits, which gives the optimum balance between the number of combinations and the desired number of quadruples. A time window containing the optimal number of hits is selected and maximisation is performed over all combinations in that window. When test vertices for all combinations in that window have been evaluated, the number of starting points is further reduced by averaging over nearby points in steps of 60 cm and 150 cm.

From the final, reduced list of starting points, likelihood maximisation, with free parameters for emission time and dark noise rate, is carried out for successive iterations of searches of test vertices in a dodecahedron grid surrounding the original test vertex for that branch or sub-branch. The dodecahedron grid shape was selected to give the optimal coverage of the space to balance accuracy and speed. The vertex corresponding to the highest likelihood is chosen as the reconstructed vertex.

5.2 Combined Fit in BONSAI

The combined fit for event pairs in BONSAI uses the same building blocks as the standard BONSAI single-event fit. It works on the assumption that the vertices of the positron and neutron event are sufficiently close as to be considered the same and outputs a single combined reconstructed vertex for an event pair, using the light from both events.

The hits from each event go through the hit selection process independently. Once the selected direct (unscattered) hits have been through the four-hit selection, a list of test vertices is output for each event. At this point, the lists of test vertices from the two events are combined into a single, larger list of test vertices which are used as starting points for the vertex search.

The hit information for each event is retained and the likelihood for each event is calculated simultaneously for each test vertex. The combined fit is achieved via the maximisation of the sum of log likelihoods for the prompt (positron or positron-like) and delayed (neutron or neutron-like) events with free parameters for prompt and delayed emission times and the dark noise rate:

$$\ln\mathcal{L}(\mathbf{x}, t_0) = \ln\mathcal{L}_p(\mathbf{x}, t_0) + \ln\mathcal{L}_d(\mathbf{x}, t_0) \tag{5.3}$$

where the log likelihoods for the prompt and delayed events are as given in Equation 5.1.

Combining the starting solutions for each individual event into a larger list of starting solutions improves rejection of local maxima in the likelihood maximisation. The addition of data points (PMT hits) is particularly helpful where light yields from one or both individual events are low, which is often the case for the positron event.

5.3 CoRe Implementation of the BONSAI Combined Fit

The CoRe (Combined Reconstruction) implementation was first developed to reconstruct pairs of events using Cherenkov light in Gd-H_2O and optimised for the best possible vertex resolution for IBD positron-neutron pairs in NEO.

BONSAI constrains the vertex search using a quadratic constraining term with a *constraining angle* of $\theta_c = 44.75°$, close to the maximum Cherenkov open angle for positrons in water. The Cherenkov open angle is the angle between the direction of the particle and the direction of the Cherenkov photons. The constraint is achieved by correcting the log likelihood for the difference $\Delta\theta(\mathbf{x}) = \theta_c - \theta_{fit}$ between the constraining angle and the fitted open angle θ_{fit} for a given reconstructed vertex, such that the log likelihood becomes:

$$\ln\mathcal{L}'(\mathbf{x}, \lambda) = \ln\mathcal{L}(\mathbf{x}, t_0) - \lambda(\Delta\theta(\mathbf{x}))^2 \tag{5.4}$$

where $\lambda = \sqrt{\frac{\alpha}{2}}$ is a constant. The value of α here depends on whether the open angle is less than the constraining angle (a good fit to Cherenkov light) or greater than the constraining angle (a poor fit to Cherenkov light). This results in an asymmetric Gaussian likelihood which gives preference to the vertices which, when combined with the PMT hit pattern, reflect a Cherenkov light distribution.

The light from an IBD positron source is a single Cherenkov cone. In water, the maximum Cherenkov open angle for positrons is 41.2°, as shown in Section 2.4.1. The Cherenkov light from IBD neutron events is generally more isotropic than the light from the positron, since it is made up of a number of Cherenkov cones emitted by multiple particles resulting from the neutron-capture gamma rays (total energy ~8 MeV) interacting with the fluid. For the CoRe implementation, the constraint was relaxed and it was found that a wider angle constraint resulted in better vertex resolution for neutrons and - additionally - for positrons. The optimal constraint was found to increase with the size of the detector. A 90°-angle constraint gave the optimal vertex resolution for both positrons and neutrons in Gd-H_2O in the 22 m detector. The optimal vertex resolution in the 16 m detector was found to be 80°.

In order to extend the CoRe implementation to reconstruct pairs of events in Gd-WbLS, the constraint on the angle was turned off completely to allow for the fully isotropic scintillation light, which does not arrive with sufficient separation from the Cherenkov light to require a separate treatment. If it were possible to separate out the Cherenkov and scintillation light, for example by methods discussed in Section 3.9, then separate treatment of the Cherenkov and scintillation light may improve results for WbLS fills.

In both the Gd-H_2O and Gd-WbLS simulations, a trigger is considered to be any event which passes the DAQ trigger of 6 hits in 100 ns on PMTs instrumenting the inner and/or outer detector. In both the WATCHMAN BONSAI implementation and in CoRe, there is a second-level trigger threshold requiring a minimum total light yield from an event of 10 hits on the PMTs instrumenting the inner volume. Ten hits is considered to be the minimum light yield required for a reliable reconstruction in BONSAI and, as such, this secondary trigger rejects events which would reconstruct poorly. This trigger was retained in CoRe for the purpose of direct comparison and under the assumption that the reconstruction with fewer than 10 hits would be similarly poor. In CoRe, a cut-off at a maximum number of 150 inner-PMT hits in

Gd-H_2O and 300 inner-PMT hits in Gd-WbLS is applied to reject fast neutrons and associated events at energies well above the energy range of the signal. A time veto after a very high-energy event is also applied to help reject fast neutron backgrounds.

CoRe iterates over all triggers and attempts to reconstruct all pairs occurring within a specified time of each other - 200 μs in Gd-H_2O and 500 μs in Gd-WbLS. For successfully reconstructed pairs, the data output include the time between events, as well as the total charge and total number of PMT hits for each event. Additional information for each event includes a measure of the fit quality, an energy analogue and a newly-implemented variable which provides information about the relative isotropy/anisotropy of the events in the pair.

5.3.1 BONSAI Timing Goodness - Fit Quality

Where events are poorly reconstructed, the coincidence of the hit time residuals as calculated from the reconstructed vertex is also poor. BONSAI outputs the coincidence of the time residuals as a measure of the vertex fit quality - the timing goodness. The time residuals are calculated using the reconstructed time of emission, which is extracted from a fit to the peak of the time-of-flight-subtracted PMT hit times at the reconstructed vertex **x**.

More specifically, the timing goodness is given by a Gaussian distribution for the Cherenkov timing resolution, weighted by a second, wider Gaussian:

$$g(\mathbf{x}) = \frac{\sum_{\text{hits}} w_i e^{-0.5\left(\frac{\Delta t_i(\mathbf{x})}{\sigma}\right)^2}}{\sum_{\text{hits}} w_i} \tag{5.5}$$

Here, σ is the timing resolution expected for Cherenkov events and w_i are weights based on the hit time residuals using a wider effective time resolution. The hit weights are given by a Gaussian of width ω:

$$w_i = e^{-0.5\left(\frac{\Delta t_i(\mathbf{x})}{\omega}\right)^2} \tag{5.6}$$

Only nearly-in-time residuals are included in the weighted Gaussian to reduce contributions from background and scattered light and the fitted dark noise rate is subtracted from both. An ideal reconstruction would result in a timing goodness $g(\mathbf{x}) = 1$.

This measure of fit quality is less well-adapted to Gd-WbLS because of the wider prompt peak in the pdf of the time residuals and the convolution of the Cherenkov and scintillation light but, for the purposes of this thesis, it was found to be sufficient as a relative, rather than absolute, measure. It should be adapted to provide an accurate measure of the fit quality for reconstruction in Gd-WbLS in the future. The results presented in this thesis for both media have timing goodness values calculated using Gaussian distributions with widths of $\sigma = 4$ns and $\omega = 50$ns.

5.3.2 Energy Analogue - Unscattered Light

The number of PMT hits due to unscattered light from a single event (e.g. either a positron or a neutron-capture in a signal-event pair) can be used as an energy analogue. Indeed, by

correcting the number of unscattered hits for various detector effects such as PMT occupancy and light attenuation, and for the distance to the interaction vertex from which the light originated, a simple relation can be found between the unscattered hits and the energy of the originating particle. This is the method used for low-energy reconstruction in Super-K [140] and is particularly useful in the case of the positron since the positron energy is directly related to the antineutrino energy as discussed in Section 2.4.1.

Only PMT hits in a narrow time window are selected in order to filter out hits from dark noise and scattered light. To maximise the number of hits from unscattered light and minimise the dark hits and scattered light, a time window of 9 ns is used by WATCHMAN for Gd-H_2O - 3 ns before the peak of the light from a single event and 6 ns after. A time window of 100 ns (10 ns before and 90 ns after) is used for Gd-WbLS in line with previous studies. The *n9* and *n100* energy analogues are therefore the number of PMT hits with a residual time within a 9 ns or 100 ns time window around the peak of the light. For the 22 m detector with a Gd-H_2O fill, it was found that the *n9* estimator was not optimal. For this reason *n100* has been used for the results presented in Section 7 for the 22 m Gd-H_2O detector. Optimisation of the values for the larger detector and the Gd-WbLS fill should be performed in the future to take into account expected noise effects in real data (see Section 8.1).

5.3.3 Isotropy Variables

In an IBD event, the light from the neutron-capture gamma cascade is isotropic, whereas the light from the positron comes from a single Cherenkov cone. A new output variable, added to the CoRe implementation of the BONSAI combined fit, provides a measure of the isotropy of each of the events in a reconstructed pair. The aim was to provide an additional way of discriminating between true IBD pairs and other pairs of coincident events such as fast neutron pairs, in which the light from both events would be isotropic.

The variables are derived from Legendre polynomials and give a measure of the spread of PMT hits in an event. A similar method was used in the SNO salt phase to discriminate between events with different relative isotropies [141].

The Legendre polynomials of order n are defined as:

$$P_n(\cos\theta_{ij}(\mathbf{x})) = \frac{1}{2^n n!}\frac{d^n}{(d\cos\theta_{ij}(\mathbf{x}))^n}(\cos^2\theta_{ij}(\mathbf{x}) - 1)^n \tag{5.7}$$

where $\theta_{ij}(\mathbf{x})$ is the angle subtended between the i^{th} and j^{th} hits from the reconstructed vertex $\boldsymbol{x}$.

The variables are then approximated as:

$$\beta_n = \frac{2}{N(N-1)}\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} P_n(\cos\theta_{ij}(\mathbf{x})) \tag{5.8}$$

where N is the number of PMT hits in an event. The full derivation can be found in Appendix C. In the isotropic case, the Legendre polynomial expansion contains only the zeroth order term and the β coefficient tends towards zero. The higher the absolute value of the variables, the more anisotropic the light.

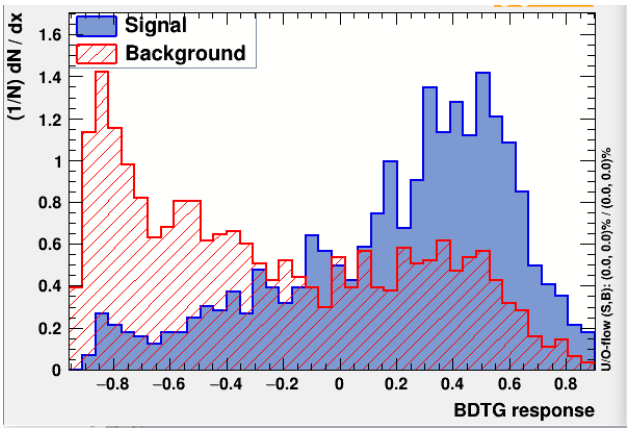

Figure 25: Discrimination of IBD positrons and neutrons using β coefficients of Legendre polynomials as a measure of the isotropy of light from the events. Discrimination was performed with the ROOT TMVA boosted decision tree with gradient boosting (BDTG). In this decision tree, positrons are classed as 'signal' and neutrons as 'background'. The BDTG response represents a degree of classification from -1 (most neutron-like) to 1 (most positron-like).

The implementation in CoRe selects only those hits which occur in a 9 ns window around the event in order to select the unscattered light from the event and filter out dark noise and scattered light, as per the energy analogue above. This time window is also used in the implementation for the Gd-WbLS detection medium in an attempt to select a maximal proportion of hits from the Cherenkov light.

Positrons and neutrons are well-discriminated by these variables. Figure 25 shows the separation that can be achieved between positrons and neutrons using a boosted decision tree with gradient boosting (BDTG) to classify the events. The variables have already been used to reject fast neutrons using machine learning [7] and can be readily adopted into a cut-based analysis (see Section 8.1).

5.4 Vertex Resolution with CoRe

The combined fit improved the IBD vertex resolution compared to the BONSAI single-event fit for all detector configurations studied. Figures 26 and 27 show the results for Gd-H_2O and Gd-WbLS respectively for both the 16 m and 22 m detectors. The improvement achieved using the combined fit is particularly beneficial for positrons at lower energies in Gd-H_2O, where the low light yield from such events makes reconstruction without the additional light from the neutrons difficult.

The results for 2.5 MeV and 5 MeV IBD positrons using BONSAI and CoRe are summarised in Table 7. The vertex resolution generally improves with the addition of WbLS thanks to the additional, scintillation light but worsens with increasing detector size. The peak of the positron signal is at 2.5 MeV. At this energy, the resolution is improved by a factor of almost 2 in Gd-H_2O - from 74.6 cm to 39.1 cm in the 16 m detector and from 83.9 cm to 44.4 cm in the 22 m detector. In Gd-WbLS the vertex resolution at the same energy improved by more than 25% from 48.6 cm to 36.2 cm in the 16 m detector. In the 22 m detector, the resolution is improved by around 20% from 53.3 cm to 42.7 cm in Gd-WbLS. The two methods tend to converge at higher energies in the reactor IBD positron spectrum for all detector configurations. The vertex resolution for IBD neutrons improved with CoRe by greater than 20% at most energies.

The improvement in the positron vertex resolution towards 1 MeV, which is observed particularly in the single-event BONSAI reconstruction, is due to a combination of factors

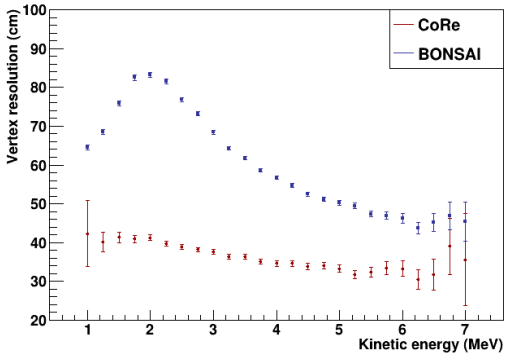

(a) Positrons in the Gd-H_2O 16 m detector.

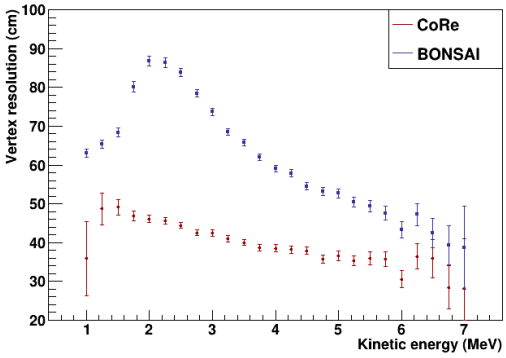

(b) Positrons in the Gd-H_2O 22 m detector.

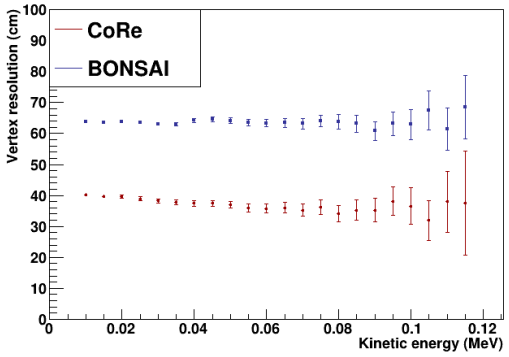

(c) Neutrons in the Gd-H_2O 16 m detector.

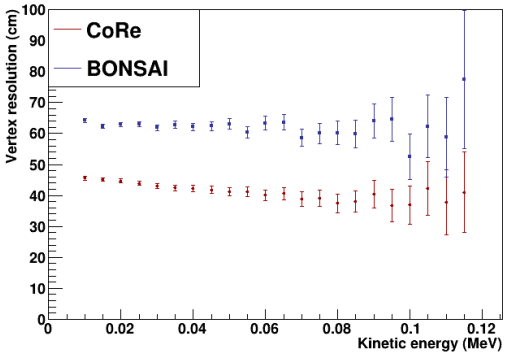

(d) Neutrons in the Gd-H_2O 22 m detector.

Figure 26: Comparison of mean vertex resolution for IBD positrons and neutrons in Gd-H_2O in the 16 m [top and bottom left] and 22 m [top and bottom right] detectors as a function of kinetic energy, using the standard BONSAI reconstruction (blue) and CoRe (red).

Table 7: Vertex resolution in cm with statistical error at selected energies for IBD positrons in the 16 m detector (top) and 22 m detector (bottom). A threshold timing goodness of 0.6 has been applied.

16 m Gd-H_2O	2.5 MeV	5 MeV	16 m Gd-WbLS	2.5 MeV	5 MeV
BONSAI	74.6 $\pm$0.50	50.0 $\pm$0.61	BONSAI	48.6 $\pm$0.50	39.1 $\pm$0.50
CoRe	39.1 $\pm$0.40	33.3 $\pm$0.65	CoRe	36.2 $\pm$1.63	32.7 $\pm$1.63
22 m Gd-H_2O	2.5 MeV	5 MeV	22 m Gd-WbLS	2.5 MeV	5 MeV
BONSAI	83.9 $\pm$1.05	52.7 $\pm$1.15	BONSAI	53.3 $\pm$0.52	44.7 $\pm$0.83
CoRe	44.4 $\pm$0.79	36.6 $\pm$1.19	CoRe	42.7 $\pm$1.87	39.0 $\pm$3.38

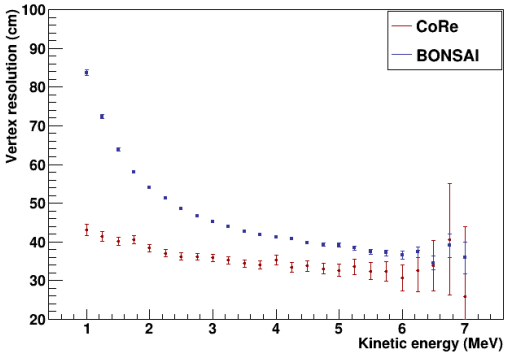

(a) Positrons in the Gd-WbLS 16 m detector

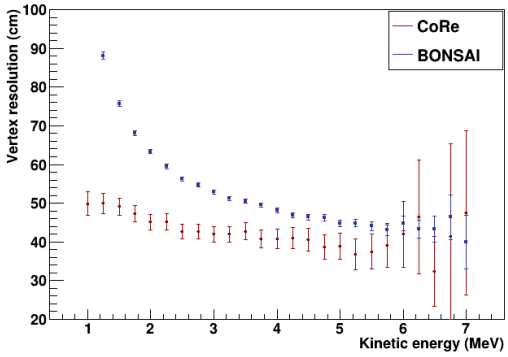

(b) Positrons in the Gd-WbLS 22 m detector

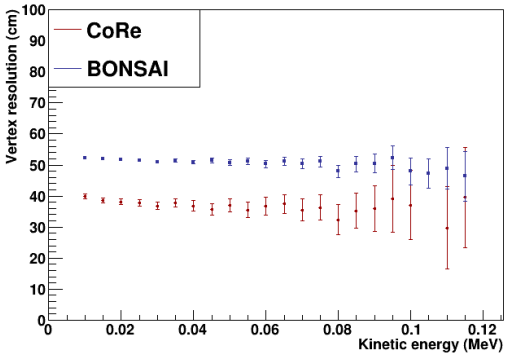

(c) Neutrons in the Gd-WbLS 16 m detector

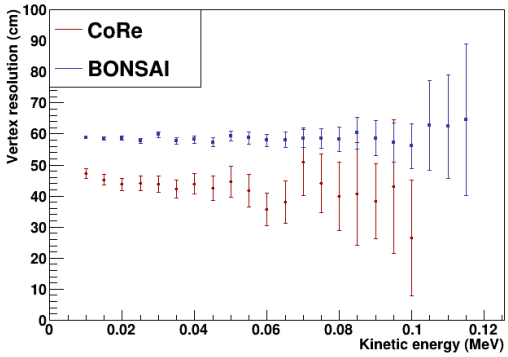

(d) Neutrons in the Gd-WbLS 22 m detector

Figure 27: Comparison of mean vertex resolution for IBD positrons and neutrons in Gd-WbLS in the 16m [top and bottom left] and 22m [top and bottom right] detectors as a function of kinetic energy, using the standard BONSAI reconstruction (blue) and CoRe (red).

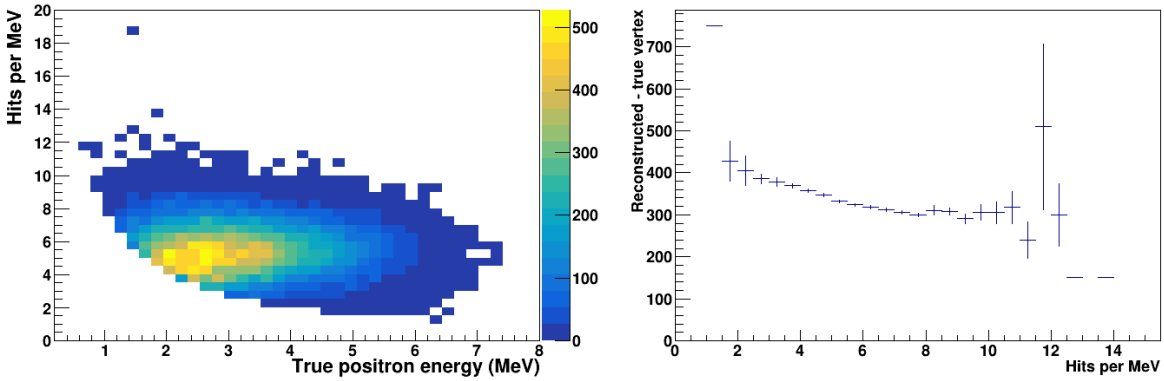

(a) Hits per MeV as a function of true energy

(b) Vertex resolution as a function of hits per MeV

Figure 28: Extrinsic effects on the vertex resolution in the 16 m Gd-H_2O detector. Second-level threshold of 10 PMT hits results in a higher mean number of hits per MeV below 2 to 2.5 MeV (left). Vertex resolution improves with the mean number of hits per MeV (right). The resulting feature is a dip in the vertex resolution below 2 MeV in Gd-H_2O.

shown in Figure 28. The second-level trigger threshold of 10 total inner PMT hits leads to a higher mean number of hits per MeV at 1 MeV, where the resolution is better, compared to at 2 MeV, where the resolution dips. In addition, the mean vertex resolution increases with the number of hits per MeV. The effect is diluted with CoRe due to the addition of the additional hits from the neutron capture and in Gd-WbLS due to the additional light yield from scintillation.

With the addition of WbLS, the BONSAI results are significantly improved since more of the lower-energy positrons have sufficient light to achieve a reasonable fit. Although the benefit of CoRe is less marked in Gd-WbLS for this reason, the improvement in vertex resolution achieved with the combined reconstruction remains significant. It is most significant in the smaller tank, which is likely due to the lower light level making the BONSAI fit less robust despite the addition of the scintillation light.

The vertex reconstruction worsens with increasing tank size with both BONSAI and CoRe for both detector fills. This is due in particular to the difficulty reconstructing vertices as the distance from the PMTs increases, making reconstruction towards the centre of the detector more difficult as the tank size increases. This effect is much more marked with the BONSAI reconstruction and in Gd-H_2O. The CoRe reconstruction offers the most improvement overall in both tank sizes with the Gd-H_2O medium.

The fit is most difficult to constrain in the Cherenkov direction. This results in the *pull* of the reconstructed vertex forwards or backwards along the Cherenkov direction with respect to the true vertex, which is shown in Figure 29 as a function of distance from the centre of the detector for positrons. In Gd-H_2O, this is particularly noticeable at the extremes of the detector - close to the centre and close to the PMTs, with the BONSAI reconstruction close to the centre of the detector and is significantly worse in the larger detector. Conversely, the pull forward is significant close to the PMTs. The additional light from the neutron helps the reconstruction to

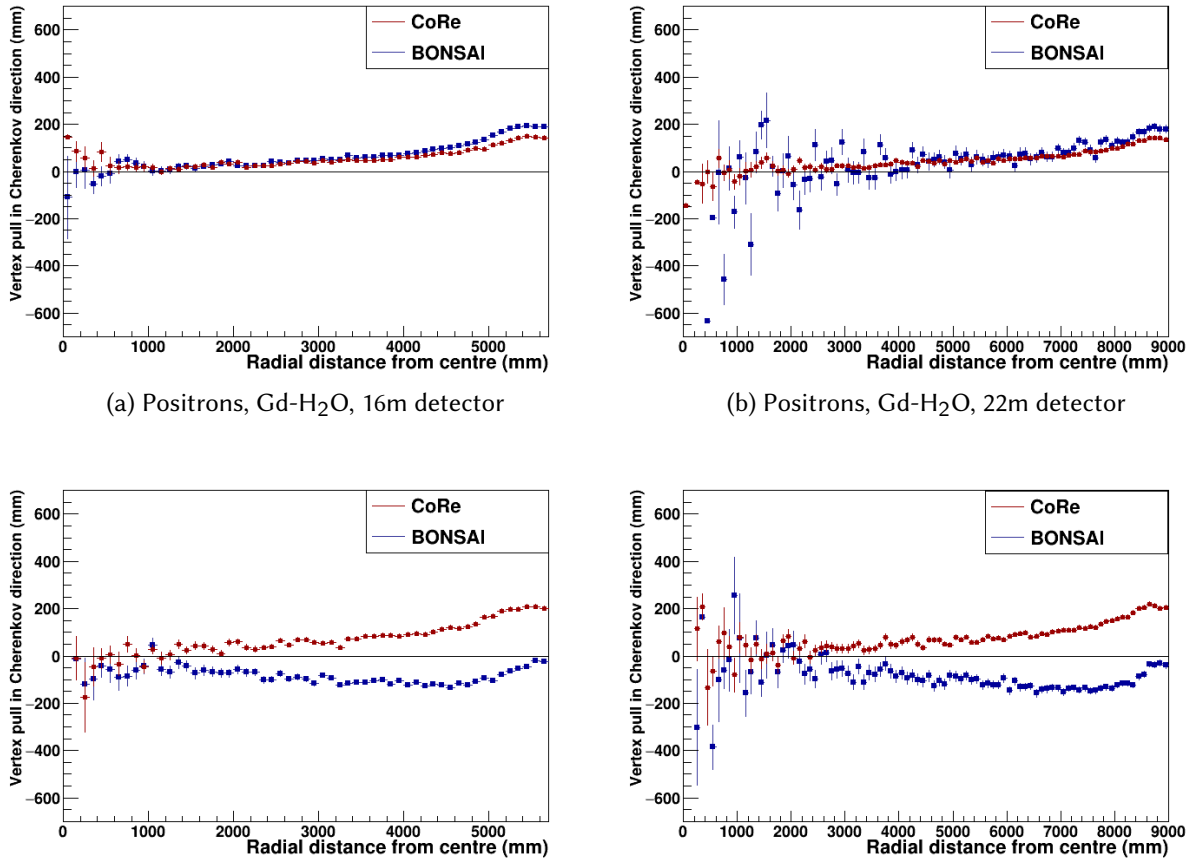

(a) Positrons, Gd-H_2O, 16m detector

(b) Positrons, Gd-H_2O, 22m detector

(c) Positrons, Gd-WbLS, 16m detector

(d) Positrons, Gd-WbLS, 22m detector

Figure 29: Comparison of the pull in the Cherenkov direction for IBD positrons in Gd-H_2O in the 16 m [top left] and 22 m [top right] detectors and in Gd-WbLS in the 16 m [bottom left] and 22 m [bottom right] detectors as a function of true radial distance from centre, using the standard BONSAI reconstruction (blue) and CoRe (red).

converge on a point in the Cherenkov direction and flattens the pull to a degree in the CoRe reconstruction in Gd-H_2O, although there remains a clear forward bias close to the edge of the detector. In Gd-WbLS, there is improvement close to the centre of the detector using CoRe but there is an overall pull in the forward direction, despite the improvement in vertex resolution with CoRe, indicating that some additional optimisation and adaptation may further improve the reconstruction in Gd-WbLS.

5.5 Background-Rejection Power of CoRe

The combined fit brings the additional benefit of powerful background rejection capability. False pairs i.e. accidental coincidences of uncorrelated events tend to result in a poorer vertex reconstruction. Consequently, the timing goodness can help to reject false pairs which pass the coincidence cuts. Figure 30 shows that a threshold value of the timing goodness as low as 0.2

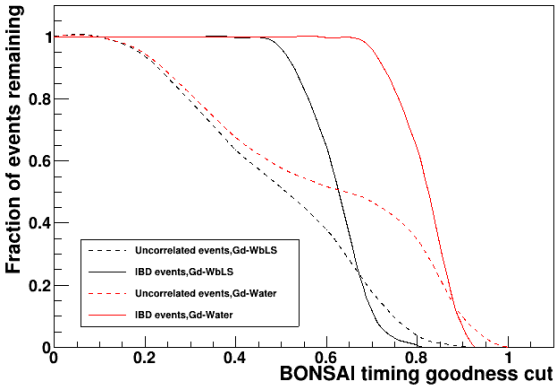

Figure 30: Fraction of correlated IBD events (solid) and uncorrelated accidental coincidences (dashed) remaining as a function of a fit quality threshold as measured by the BONSAI timing goodness, for Gd-H_2O (red) and Gd-WbLS (black).

has the power to reject uncorrelated background events, while a threshold of up to 0.6 will retain almost all of the signal events in both Gd-H_2O and Gd-WbLS. A timing goodness threshold has a more limited effect on the single-event reconstruction in BONSAI.

Conversely, this means that the timing goodness can be helpful in selecting true pairs from data which is a combination of all types of events (see Section 8.1).

5.6 Summary

- A method of combined reconstruction for coincident events has been incorporated into the BONSAI low-energy fitter.
- The BONSAI fitter is a maximum-likelihood fitter to the PMT hit times from an event. The combined reconstruction maximises the sum of the likelihoods for the prompt and delayed events.
- I have developed the CoRe Combined Reconstruction, which is an implementation of the combined reconstruction for WATCHMAN and has been optimised for the current NEO detector designs.
- A new set of variables which are a measure of isotropy was developed to simultaneously improve signal selection efficiency and the rejection of fast neutron backgrounds.
- CoRe improves the vertex resolution for IBD positrons and neutrons in both Gd-H_2O and Gd-WbLS and in both detector sizes under consideration.
- Further development and optimisation for Gd-WbLS could improve the fit further in that detector medium.
- CoRe offers the potential for powerful background rejection in combination with the improvement in vertex resolution, which brings associated improvements in our understanding of the reactor antineutrino emission.

The next stage of the data processing, after the reconstruction, is the final analysis. Data output from CoRe is fed into the coincidence evaluation and analysis presented in the next

chapter. This optimises the signal over background in a number of dimensions before application to the detector scenarios discussed in Chapter 7.

6 Coincidence Analysis

Any fact becomes important when it's connected to another.

– Umberto Eco, *Foucault's Pendulum*

This chapter describes a new simulation and analysis tool Cobraa (Coincident-Background Reactor Antineutrino Analysis) [142], which was developed as part of the software chain described in this thesis to perform a full evaluation of coincident signal and background events in NEO. Cobraa was developed from the WATCHMAKERS simulation and analysis tool [143] and as such draws its main structural features from WATCHMAKERS. Cobraa differs in its methods of simulation and coincidence evaluation in that it is based solely on the evaluation of true coincidences. The original WATCHMAKERS code, to which I have contributed and from which Cobraa was developed, performed analytical evaluation of IBD and accidental coincidences using single events for the correlated IBD signal and backgrounds and the uncorrelated events due to natural radioactivity. Using combined simulations, Cobraa performs a full evaluation of accidental coincidences of uncorrelated backgrounds and of true coincidences of both signal and correlated background events, additionally including fast neutron and radionuclide backgrounds. Cobraa can be used with either the single-event WATCHMAN BONSAI reconstruction or CoRe. For the purposes of this thesis, the Cobraa analysis was used in combination with CoRe.

The Cobraa-CoRe analysis chain takes the user from simulation in RAT-PAC of coincident correlated and uncorrelated signal and backgrounds as described in Section 4, through the CoRe combined vertex fitting as described in Section 5, to the optimisation of signal over all backgrounds in a total of six dimensions with the Cobraa analysis described in this chapter.

The Cobraa-CoRe process results in optimal signal efficiency and background reduction, in particular with improved rejection of uncorrelated backgrounds compared to previous analyses. This is achieved through the combination of a full coincidence simulation, more accurate vertex reconstruction and very fast optimisation in six dimensions with full evaluation of coincidences. Uncorrelated and fast neutron backgrounds are now at a minimum.

Section 6.1 describes the simulation process in Cobraa and Section 6.2 describes the Cobraa analysis and signal over background optimisation. Application of this optimisation to discovery scenarios are discussed in Chapter 7.

6.1 Full coincidence simulation with Cobraa

Cobraa creates all of the RAT-PAC macros required to run a full-detector simulation for reactor antineutrino analysis. For the development and optimisation of the analysis for each detector configuration, the signal and each type of correlated background are simulated individually at the relevant expected rates in the detector. Events due to natural radioactivity are combined into a single simulation where the decays are interleaved at their respective rates, as discussed in Chapter 4 and shown in Appendix B.

There is an option to select either the 16 m or the 22 m detector and to switch between the Gd-H_2O and Gd-WbLS detector media. Additional options are detailed in [142].

6.2 The Cobraa Coincident-Background Reactor Antineutrino Analysis

Reconstructed events from BONSAI or CoRe, as described in Chapter 5, are passed to the analysis. The analysis differs slightly depending on the particular reconstruction used. For the purposes of this thesis, the discussion and results presented will relate to the CoRe vertex fitting and associated analysis.

6.2.1 Optimisation of Dwell Time

The *dwell time* is the observation time required to achieve a given non-proliferation aim (see Chapter 7 for a discussion of signal significance and sensitivity metrics). The Cobraa-CoRe analysis optimises the dwell time via data-reduction cuts in six dimensions:

1. Fiducial cut - rejects all events below a minimum distance from the inner PMT radius.
2. Prompt-event energy threshold cut $E_{p,min}$ - rejects pairs with a lower prompt-event energy.
3. Delayed-event energy threshold cut $E_{d,min}$ - rejects pairs with a lower delayed-event energy.
4. Cut on maximum time between events dT_{max} - rejects pairs where the time between the two events $dT>dT_{max}$.
5. Minimum timing goodness g_{min} - rejects poorly reconstructed events.
6. Maximum prompt-event energy $E_{p,max}$ - rejects pairs with a higher prompt-event energy.

These cuts are targeted at rejecting the backgrounds in NEO. The first three optimisations were previously developed to reject uncorrelated backgrounds. These were extended in Cobraa to incorporate optimisation over the time between events and the timing goodness to improve uncorrelated-background rejection. The maximum prompt-energy cut was developed for Cobraa as a way to reject cosmogenic muon backgrounds following a study which showed that the fast neutron background, in particular, could be detrimental to detector sensitivity [6]. The range over which the optimisation is performed can be specified for each variable, as can the granularity within the range. The fill-appropriate *n9* or *n100* energy analogue as discussed in Section 5.3.2 is used for the cuts on energy. Using a reconstructed energy or correcting the energy analogue for detector effects and the location of the event vertex offers the potential for improvement in the optimisation. The fiducial cut is applied as a minimum distance from the PMT and creates a virtual inner target volume within the inner detector.

The first stage of the analysis is to evaluate the rate of coincidences for each combination of cuts. An additional multiplicity cut is applied to reduce the fast neutron background (see Section 6.2.3). Where the number of coincidences is evaluated to zero, the 95% Poisson[11] upper confidence limit is used. The coincidence rates are output in two-dimensional slices as a function of prompt-event energy threshold and fiducial cut for each combination of the remaining cuts. Figure 31 shows one such slice for the IBD signal.

[11] Strictly, the binomial upper confidence level is more appropriate here but the Poisson upper confidence limit gives a more conservative estimate of the uncertainty.

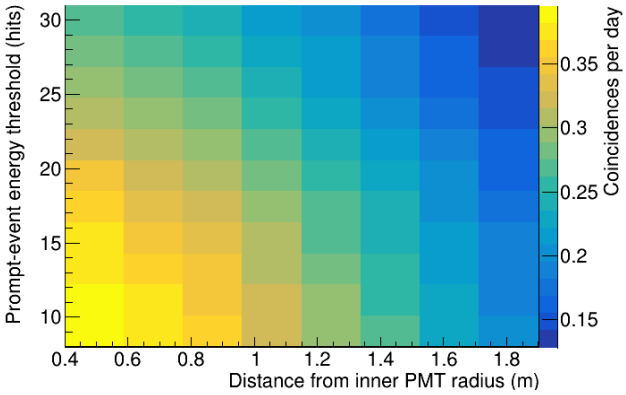

Figure 31: Two-dimensional profile output by Cobraa showing the coincident-signal rates per day in Gd-WbLS in the 16 m detector. The fiducial cut and prompt energy threshold are allowed to vary while the other cuts are held constant.

The second stage of the analysis performs the final sensitivity optimisation. It reads in the rates from the histograms in the previous step, adjusts the radionuclide rates for an analytical time-after-muon veto (see Section 6.2.4) and evaluates the signal significance and associated uncertainties at each combination of cuts. Plots of the signal significance and final signal and background coincidence rates are output in two-dimensional slices, as in the previous stage, for reference. Details of the optimal signal significance and associated optimal cuts, rates and uncertainties are also output.

Optimising on the dwell time has the effect of optimising on either the signal over background significance or the uncertainty on the background, depending on whether statistics or backgrounds and systematic uncertainties are the limiting factor.

6.2.2 Rejection of Accidental Coincidences

The first five of the cuts over which Cobraa optimises, and which are described above, are the most effective for the rejection of accidental coincidences. Figure 23 shows that a fiducial cut to remove events close to the PMTs can be highly effective in reducing many of the uncorrelated radioactive-decay events.

The prompt and delayed energy thresholds can remove more than 99% of the accidental coincidences which pass the DAQ and reconstruction trigger thresholds. Separate energy thresholds for the prompt and delayed event was first developed within the Watchmakers analysis. This increases the rejection of accidental coincidences in Gd-H_2O by allowing the energy threshold of the delayed, neutron capture-like event to be higher, while preserving as much of the positron signal as possible. In WbLS without gadolinium, the optimal prompt-event energy threshold can be higher than the optimal delayed-event energy threshold. Allowing the energy threshold of the delayed event to be either higher or lower than the prompt event can be beneficial in the Gd-WbLS detector medium at the 1% WbLS concentration under consideration.

Many of the remaining events are removed by the cuts on the time between events and the fit quality. Figure 32 shows the typical time and distance between consecutive uncorrelated events. A cut to remove events separated by more than ~100 μs in Gd-H_2O removes more than 50% of the uncorrelated events which pass the trigger threshold. The optimisation is carried

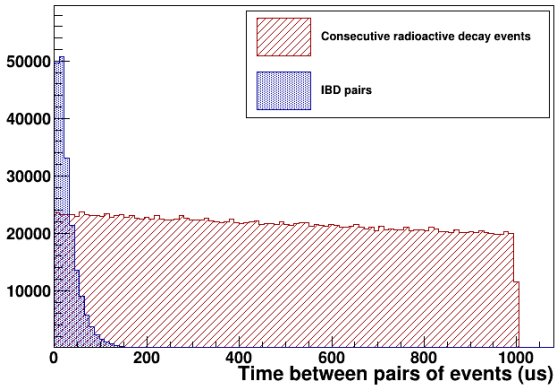

Figure 32: Typical time between consecutive uncorrelated events due to radioactivity compared to time between IBD positron-neutron pairs. A cut-off at 1000 μs has been applied to the time between consecutive uncorrelated events but in fact the mean time between events 6000 μs and the distribution is an exponential which extends to tens of 1000s of μs. Many of the uncorrelated events can be removed by the dT and g_{min} cuts on the time between events and reconstruction quality.

out over a range around a value close to the expected time between the positron and neutron events in an IBD event pair. Optimisation of the cut on the minimum timing goodness exploits the false-pair rejection power of CoRe shown in Figure 30. A pair of uncorrelated events which are far apart in distance are likely to have a poor measure of quality for a fit which outputs a single vertex for both events. For this reason, a cut on the fit quality can be viewed, in part, as an analogue for a cut on the distance between events.

Accidental coincidences have been of primary concern to WATCHMAN in the past. The combination of the CoRe reconstruction and subsequent Cobraa analysis tends to remove all of the accidental coincidences in Gd-H_2O. In this case, an upper confidence limit is given. Even in Gd-WbLS, which has a lower energy threshold and 'sees' more of the uncorrelated events, accidental coincidences are reduced to a subdominant background.

6.2.3 Fast Neutron Rejection

In a fast neutron pair, the first (prompt) event is likely to be on average of much higher energy than a positron in an IBD event. A maximum prompt-event energy cut, which is optimised in a range between the peak positron energy and the maximum positron energy, is an effective tool in rejecting fast neutron events where the prompt-event energy is not typical of a positron.

Additional fixed cuts were developed to deal with the fast neutron background:

- Re-trigger threshold cut
- Multiplicity cut

A threshold of 1 μs between events helps to reject re-triggers. These are particularly problematic at the very high energies of fast-neutron events.

Fast neutron events can consist of up to (and sometimes in excess of) eight coincident events in the NEO detector. Rejecting multiply-coincident events can therefore help to reject fast neutron events. The Cobraa analysis accepts only pairs of events which have no other coincident event before or after the pair. This can reduce the fast neutron rate in the fiducial volume by greater than 90%.

The effectiveness of the multiplicity and maximum prompt-event energy cut-off is evident in Figure 33, which gives an indication of the relative numbers of events remaining, as a function

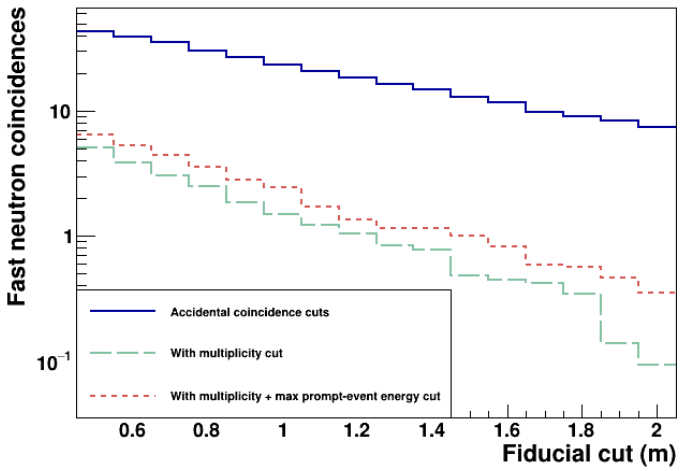

Figure 33: Fast neutron coincidences in Gd-H_2O in the 16 m detector as a function of fiducial cut with the 'basic' cuts for rejection of accidental coincidences (blue, solid), additional cut on multiplicity (red, short-dashed) and additional cut on multiplicity plus a fixed $E_{p,max}$ of 40 *n9* hits (green, long-dashed).

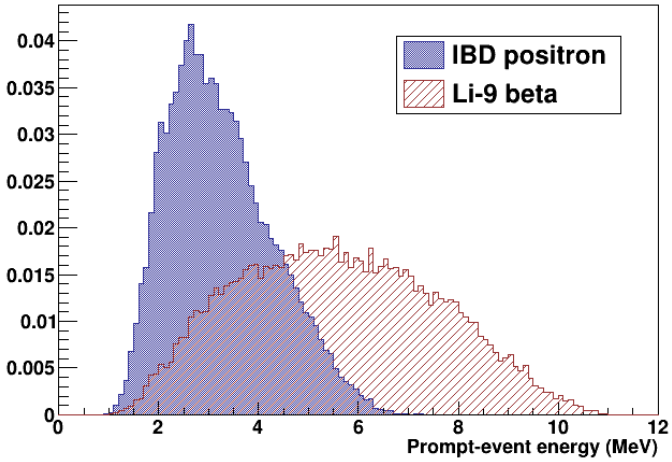

Figure 34: Energies of the prompt events in the IBD (positron) and ^{9}Li β-neutron decay (β particle). The mean prompt energy in the ^{9}Li decay tends is higher than that of the positron and a maximum prompt-energy requirement is effective in cutting out many ^{9}Li backgrounds.

of the fiducial cut expressed as the distance from the inner PMT radius, after the multiplicity and maximum prompt-event energy cuts have been applied.

6.2.4 Radionuclide rejection

The $E_{p,max}$ cut used to reject fast neutron pairs with a high-energy prompt event is also effective in reducing the radionuclide backgrounds. It is particularly useful in the case of ^{9}Li, where the endpoint energy of the β can be higher than the expected energy of the IBD positron as Figure 34 shows.

We can expect to reject some of the β-neutron radionuclide backgrounds using a vetofraction after a muon in the detector at the DAQ stage. Since the muons are not currently simulated, this has been applied analytically in Cobraa. Adapted from [144], the quantity of radionuclide activity remaining after a veto time t_{veto} with a fractional muon detection efficiency ϵ is given by:

$$R_{i,t_{veto}} = \left(1 - \epsilon + \epsilon \frac{\int_{t_{veto}}^{\infty} e^{-\ln(2)t/t_{1/2}} dt}{\int_{0}^{\infty} e^{-\ln(2)t/t_{1/2}} dt}\right) R_{i,tot}. \quad (6.1)$$

With an active veto, the muon-detection efficiency might reasonably be expected to be as high as 99.9%, allowing for electronics issues. With a passive buffer rather than an active veto, the muon-detection efficiency in the fiducial volume might be assumed to be the same. However, making some correction for muons which pass through the veto but induce spallation in the fiducial volume, a conservative muon-detection efficiency of 95% for a passive veto is assumed in the Cobraa analysis and the fractional reductions in radionuclide rates with a 1 s veto are:

$$R_{^9\mathrm{Li,1s\ veto}} = 0.069\ R_{^9\mathrm{Li,tot}}$$
$$R_{^{17}\mathrm{N,1s\ veto}} = 0.85\ R_{^{17}\mathrm{N,tot}}$$

Making the reasonable assumption that it will be possible in NEO to track the muons passing through the fiducial volume, the 1s time veto is assumed to be imposed only over a limited transverse distance from the muon track and thus incurs negligible detector dead time. The detector deadtime can be calculated as follows:

$$\text{Detector dead time}(t_{veto}) = t_{veto}(s) \times \frac{V_{veto}}{V_{fiducial}} \times R_{\mu} \times \epsilon$$

where t_{veto} is the duration of the veto after a muon and V_{veto} is the veto volume e.g. the volume within a transverse distance of 1 m along the length of the muon track.

The veto is applied analytically in Cobraa after the coincidence rates have been evaluated. NEO's muon-tracking capabilities with and without an active (instrumented) veto should be optimised and evaluated quantitatively with muon simulations once the final detector design has been decided. A further potential improvement to the muon veto is discussed in Section 8.1.

6.3 Summary

In combination with CoRe, Cobraa has the potential to increase signal efficiency and background rejection through:

- Full coincidence simulation,
- More accurate vertex reconstruction,
- Fast optimisation with full coincidence evaluation in six dimensions,
- Tailored rejection for backgrounds due to radioactivity, fast-neutrons and β-neutron emitting radionuclides,
- Analytical muon veto.

The optimised signal-over-background significance can then be applied to reactor discovery scenarios as discussed in the following chapter.

7 Sensitivity for Reactor Monitoring

> *If you are under the impression you have already perfected yourself, you will never rise to the heights you are no doubt capable of.*
>
> – Kazuo Ishiguro, *Remains of the Day*

The final step in the Cobraa-CoRe simulation-reconstruction-analysis toolchain is to apply the optimisation for application to remote reactor monitoring for non-proliferation.

The sensitivity of the experiment is expressed as the expected required observation time to achieve a given goal, or *discovery scenario*. The required observation time is referred to as the experiment *dwell time*. The dwell time is the sensitivity metric used for comparison of the detector configurations under consideration. The derivation of the sensitivity metric depends on the discovery scenario.

For the purposes of this thesis, the discovery scenarios have been identified as follows:

1. **Discovery of unknown reactor with known backgrounds (partially blind).**
   Discovery of one or more unknown reactor to reject the background-only or zero-reactor hypothesis with known backgrounds. Discovering a second, unknown reactor in the presence of a known, nearby reactor assumes that we have accurate knowledge of the power output of the known reactor.

2. **Discovery of unknown reactor with no reactor or background information (totally blind).**
   Discovery of one or more unknown reactor to reject the no-reactor hypothesis with unknown backgrounds. This relies on periods of refuelling to make a measurement of the background rate when there is no reactor in operation. Discovering a second, unknown reactor, in the presence of another reactor but without knowledge of either reactor or backgrounds would require a measurement of the background rate when there is only one reactor in operation.

The cores at Hartlepool are the primary signal for NEO if the reactors' operation is extended beyond 2024. A number of combinations of more distant reactor signals and reactor backgrounds have additionally been evaluated for each detector configuration for all applicable discovery scenarios. Listed with provisional decommissioning dates, these combinations are:

1. Hartlepool 1 & 2 — 2024, 2024
2. Hartlepool 1
3. Hartlepool 2
4. Heysham 1 & 2 — 2024, 2030
5. Heysham 2 & Torness — 2030, 2030
6. Heysham 2

Table 8 shows the reactors included in the signal and background for each of the reactor signals listed above. 'World' reactor backgrounds include all reactors more distant than Torness[12]. Rates are calculated from the average monthly thermal power for 2019, as reported by the IAEA. These scenarios take into account the expected shutdown of the Hartlepool cores and Heysham 1 in 2024.

Table 8: Combinations of reactor signals and backgrounds evaluated. Reactors marked black are included in the signal and reactors marked grey are included in the background. Reactors marked with a slash are assumed to be decommissioned.

Signal combination	Hartlepool 1	Hartlepool 2	Heysham 1	Heyhsam 2	Torness	World
Hartlepool 1 & 2	■	■	□	□	□	□
Hartlepool 1	■	□	□	□	□	□
Hartlepool 2	□	■	□	□	□	□
Heysham 1 & 2	⧸	⧸	■	■	□	□
Heysham 2 + Torness	⧸	⧸	⧸	■	■	□
Heysham 2	⧸	⧸	⧸	■	□	□

Cobraa treats NEO as a Poisson counting experiment. The scenarios detailed above call for different sensitivity metrics, which define our ability to perform reactor monitoring and are derived in full in Appendix D. Section 7.1 describes the sensitivity metrics which have been applied to reactor monitoring with known backgrounds in the context of a number of reactor signal/background combinations. Section 7.2 describes the sensitivity metric which has been applied to the discovery of an unknown reactor with unknown backgrounds.

7.1 Anomaly Detection with Known Backgrounds

The first, partially blind, Discovery Scenario 1 has been considered in the context of the signals in Table 8, with the exception of the Hartlepool 1 signal, which has a similar outcome to the Hartlepool 2 case in this method.

The simplest way to approach this scenario is to reject the *background-only* hypothesis where there is one or more unknown reactor core. This is referred to as *anomaly detection*. Events detected over the known background rate would serve to reject the background-only hypothesis. With a second-reactor signal, events detected over the expected rate from the known reactor and backgrounds would serve to reject the background-only, single-reactor hypothesis. A rejection of the null hypothesis to 3σ (or equivalent) significance is considered to be sufficient in the application to non-proliferation for this type of scenario. Two sensitivity metrics have been derived to test the background-only hypothesis, which are selected according to the strength of the signal relative to the background, as explained below.

[12]Data are projected to 2028, when the Hunterston and Hinkley Point B reactors will have been decommissioned and the Hinkley Point C reactor will be on.

The sensitivity metric for an anomaly measurement of one or more reactor core in the presence of known backgrounds is most simply derived from the Gaussian significance:

$$N_\sigma = \frac{st}{\sqrt{\Sigma b_i t + (\Sigma \sigma_{b_i} t)^2}}, \tag{7.1}$$

where N_σ is the number of Gaussian standard deviations from the expected background counts, st is the expected total signal counts at rate s in time t, b_i is the expected rate of the i^{th} background, $\Sigma b_i t$ is the expected total background counts and $\Sigma \sigma_{b_i} t$ is the total systematic uncertainty on all backgrounds. This gives an equation for the dwell time as follows:

$$t_{dwell} = \frac{N_\sigma^2 \Sigma b_i}{s^2 - N_\sigma^2 \Sigma \sigma_{b_i,sys}^2} \tag{7.2}$$

This is derived in Appendix D.1 and is calculated for $N_\sigma = 3$.

This approximation is valid where $s \ll b$ and where the number of counts is greater than 20 to 30. Where these criteria are not met, the Gaussian approximation is replaced with a Poisson distribution and the total systematic uncertainty on the background is incorporated as a Gaussian-distributed nuisance parameter. The resulting expression for the Poisson significance Z is more complex than Equation 7.1:

$$Z = \left[-2(st + \Sigma b_i t) \ln \frac{\hat{B}_0}{st + \Sigma b_i t} - \frac{(\Sigma b_i t - \hat{B}_0)^2}{2\Sigma(\sigma t)^2} - \hat{B}_0 + st + \Sigma b_i t \right]^{\frac{1}{2}} \tag{7.3}$$

where the background rate $\hat{B}_0$ for $st = 0$ is:

$$\hat{B}_0 = \frac{1}{2}\left(\Sigma b_i t - (\Sigma \sigma_{b_i} t)^2 + \sqrt{(\Sigma b_i t)^2 - 2\Sigma b_i t (\Sigma \sigma_{b_i} t)^2 + 4N(\Sigma \sigma_{b_i} t)^2 + (\Sigma \sigma_{b_i} t)^4} \right). \tag{7.4}$$

This is derived in Appendix D.2.

Equation 7.3 for the significance can be expressed as a function of a free parameter t and in this way the value for t_{dwell} can be extracted for the required significance Z=3, where Z is the Poisson equivalent of the Gaussian N_σ.

7.2 Reactor discovery with unknown backgrounds

The second discovery scenario is the sort that might be encountered in the field in a non-collaborative situation. In this case, we assume that we have no knowledge of the reactors or the backgrounds. In reality, it will be possible to constrain backgrounds to a degree but the assumption is made here that no information on backgrounds will be available. For this reason, this can be considered a worst-case scenario. This scenario is much more challenging than the case with known backgrounds and, as such, it has been applied only in the context of the three Hartlepool signals as set out at the start of this chapter. It should be noted that the Hartlepool 2 core is currently operating at a lower power than the Hartlepool 1 core (446 NIU compared to

595 NIU). This makes for an interesting scenario - looking at the Hartlepool 2 core against the background of the other core gives insight into the ability to discover a less powerful reactor operating close to a more powerful declared reactor, which would be a significant non-proliferation goal.

Confirming the presence of a reactor where both the signal and backgrounds are initially unknown requires longer observation times. Throughout the duration of the observation, periods when a reactor is suspected to be on can be separated out from periods when the reactor is suspected to be powered down for refuelling. There are therefore two observation periods: T_{on} when the reactor is turned on (or suspected to be on) and T_{off} when the reactor is turned off, during which the background rate can be measured. A period when the reactor is suspected to be on might be marked by a higher mean total rate in this worst-case scenario where there is no information about the backgrounds. Information about the reactor schedule is not assumed. It will be necessary to observe for a sufficient length of time to establish a background-only rate.

For this type of measurement, the excess rate under reactor operation must be sufficiently high to first minimise the probability of a false positive on measurement of the background-only rate and then minimise the probability of a false negative on the signal plus background rate. A 3σ positive detection of a signal at the 95% confidence level, giving a false positive rate of 5% and a false negative rate of 0.27%, has been proposed by WATCHMAN for this type of measurement. A less stringent requirement may be acceptable in practice, particularly in combination with complementary monitoring, and so a 2σ positive detection at the 90% confidence level has also been investigated.

In this scenario, the significance is dependent on the statistical error in the measurements made both during T_{on} and during T_{off}. The expression for the dwell time is derived from the total number of counts required for a 3σ detection at the 95% confidence level. The total number of counts is calculated according to the method presented in [145] and the dwell time is given by:

$$\mathrm{t_{dwell}} = (2 + \mathrm{R_{onOff}}) \left(\frac{1.645\sqrt{2\Sigma \mathrm{b_i}(\mathrm{R_{onOff}} + 1)} + 3\sqrt{(\mathrm{sR_{onoff}} + 2\Sigma \mathrm{b_i}(\mathrm{R_{onOff}} + 1)}}{\mathrm{sR}} \right)^2 \tag{7.5}$$

where s is the expected signal rate, b is the expected total background rate and $R_{onOff} = \frac{T_{on}}{T_{off}}$. For the purposes of this thesis, the ratio R_{onOff} is taken to be the expected ratio given the typical refuelling schedule of an AGR, where each core is shut down for two weeks every six months. The full derivation is presented in Appendix D.3.

7.3 Reactor Sensitivity with Current Detector Designs

The CoRe reconstruction and Cobraa analysis were applied to the two discovery scenarios described in Sections 7.1 and 7.2. The 16 m and 22 m detectors were investigated, each with both the Gd-H_2O and the Gd-WbLS fills. As might be expected, the sensitivity increases with the addition of WbLS and with the larger detector size in all scenarios and the dwell time increases with the stringency of the dwell time metric.

Increasing the detector size improves sensitivity for a given fill by effectively increasing the size of the fiducial volume, which increases the signal rate. Although some background rates also

scale with size, the higher signal rate makes it possible to cut more aggressively on the fiducial volume and energy cuts, resulting in proportionately better rejection of background rates. This results in a net increase in sensitivity and a decrease in dwell time with increasing detector size.

Upgrading the detector fill from Gd-H_2O to Gd-WbLS improves sensitivity to the signal thanks to the increase in light, particularly at the lower end of the reactor spectrum. This also increases sensitivity to backgrounds - again, particularly at lower energies - but as in the case of the increase in tank size, the higher signal rate enables more aggressive background rejection and so overall sensitivity is increased. This results in a decrease in dwell time with the addition of WbLS.

7.3.1 Dwell Time to Reactor Discovery

I begin with an examination of the most challenging *reactor discovery* scenario - the positive detection of a reactor in the presence of unknown backgrounds. This uses the metric introduced in Section 7.2 above.

The sensitivity of the four detector configurations to a 3σ reactor discovery at the 95% confidence level in the presence of unknown backgrounds is shown in Table 9. The dwell time comes down as the sensitivity increases. If the lifetime of the Hartlepool Power Station were to be extended beyond the planned 2024 permanent shutdown date, a reactor discovery measurement of Hartlepool 1 & 2 with the 22 m detector and Gd-WbLS fill could be made in around four months and a reactor discovery measurement of the individual Hartlepool 1 core could be made in around 19 months. The 22 m Gd-H_2O detector could make a reactor discovery measurement of the dual-core signal in under 6 months and of Hartlepool 1 just over two years. The 16 m Gd-WbLS configuration has sensitivity only to the dual-core Hartlepool 1 & 2 signal within two years. Results have been highlighted where a reactor discovery could be achieved within one year, which might be considered a reasonable timeframe for reactor discovery in a non-collaborative situation. Only the 22 m configurations could achieve a reactor discovery of the dual-core signal within one year.

Table 9: Cobraa results summary for unknown *reactor discovery* with no knowledge of backgrounds - dwell time in days for rejection of the no-reactor hypothesis to 3σ significance at the 95% confidence level. Results are highlighted in bold where a reactor discovery could be achieved **within one year**.

Detector	Hartlepool 1 & 2	Hartlepool 1	Hartlepool 2
16 m Gd-H_2O	892	4133	8043
16 m Gd-WbLS	549	2540	5133
22 m Gd-H_2O	**168**	799	1573
22 m Gd-WbLS	**121**	576	1124

Table 10 shows how the dwell times would change if the requirement were relaxed to a 2σ detection at the 90% confidence level. The dwell times, in this case, are approximately halved

and 22 m Gd-H_2O detector would be sensitive to the dual-core signal within 3 months and the single-core Hartlepool 1 signal in just over a year. At this level of confidence, the 22 m Gd-WbLS detector would be sensitive to discovery of the Hartlepool 2 core, which is operating at a lower power than the Hartlepool 1 core, within two years (in 19 months).

Table 10: Cobraa results summary for unknown *reactor discovery* with no knowledge of backgrounds - dwell time in days for rejection of the no-reactor hypothesis to 2σ significance at the 90% confidence level. Results are highlighted where a discovery could be achieved **within one year**.

Detector	Hartlepool 1 & 2	Hartlepool 1	Hartlepool 2
16 m Gd-H_2O	435	2042	3986
16 m Gd-WbLS	**266**	1254	2543
22 m Gd-H_2O	**81**	395	780
22 m Gd-WbLS	**59**	**285**	557

This is an attempt to examine what might be achieved in the field in a more challenging detection scenario. Assuming no knowledge of backgrounds at all is the worst-case scenario. In practice, some knowledge of background rates from theoretical calculations and sideband measurements outside the signal region would be used to constrain the backgrounds. In particular, the rate of radioactive decays is very high and the background due to accidental coincidences could be well-constrained. In addition, antineutrino detection technology is likely to be supported by more conventional monitoring methods such as visual and thermal imaging. This will reduce the stringency on the false positives requirement in particular. Reducing the false positive requirement and constraining backgrounds would bring down dwell times significantly and could open up the possibility of making a discovery of a more distant reactor with the 22 m Gd-WbLS detector configuration.

7.3.2 Dwell Time to Anomaly Detection

I now move on to the anomaly detection method using the metric described in Section 7.1. Table 11 summarises the results obtained for the *anomaly detection* scenario. The entries highlighted in bold in Table 11, which shows the results for anomaly detection, are the options considered to be possible within the practical constraints outlined.

- The AGR-1 reactors Hartlepool 1 & 2 will be decommissioned in 2024,
- the AGR-1 reactors at Heysham 1 will also be decommissioned in 2024,
- NEO data taking begins in 2028 and
- the AGR-2 reactors at Heysham 2 and the Torness complex will be decommissioned two years after NEO is due to be commissioned, in 2030.

Table 11: Cobraa results summary for anomaly detection - dwell time in days for rejection of the background-only hypothesis to 3σ significance. The results highlighted in bold are those which could be achieved between the commissioning of NEO and the planned shutdowns of the AGR-2 reactors.

Detector	Hartlepool 1 & 2	Hartlepool 1	Heysham 1 & 2	Heysham 2 + Torness	Heysham 2
16 m Gd-H_2O	12	61	2327	3488	8739
16 m Gd-WbLS	7	35	738	1022	3589
22 m Gd-H_2O	3	11	241	**232**	985
22 m Gd-WbLS	2	8	152	**192**	**647**

Although the 16 m detector with a Gd-H_2O fill is sensitive to the reactors at Hartlepool, it is not sufficiently sensitive to detect the signal from more distant reactors within the two-year constraint. With a Gd-WbLS fill, the sensitivity would increase but this configuration could still not make an anomaly detection of the Heysham 2 or Heysham 2 & Torness signals within two years. No detection will be viable in either of the 16 m configurations if the planned shutdowns in 2024 and 2030 go ahead.

The only option for NEO, with sensitivity to a single-site signal, out of the configurations investigated here would be the 22 m detector with a Gd-WbLS fill. This configuration would have sensitivity to the Heysham 2 & Torness signal in just over 6 months and to the Heysham 2 signal within 22 months.

The 22 m Gd-H_2O option would serve as risk mitigation if the Gd-WbLS technology was not ready for large-scale deployment before the commissioning of NEO. With Gd-H_2O, the 22 m detector should be able to make an anomaly detection of the Heysham 2 & Torness signal in around 8 months but detection of the Heysham 2 signal would take over two years - around 2 years and 8 months in this result, which extends beyond the planned permanent shutdown of the complex.

Figure 35 shows the Heysham 2 signal significance after 6, 12, 18 and 24 months of observation time for the 22 m detector with Gd-H_2O and Gd-WbLS fills. Signal, background and systematic error rates are taken from the result of the optimisation on dwell time. Table 12 shows the precise values of the significance at these observation times for both configurations. With a Gd-WbLS fill, a significance of 3.1 would be reached within the two-year period. A significance of 2.6 would be reached within the same time frame with the Gd-H_2O fill and it should be noted that this may be considered sufficient by the non-proliferation community.

The optimised rates for each of the signals are given for the 22 m detector configurations in Table 13. These rates were used to produce Figure 35 and Table 12 above. For *anomaly detection*, the 16 m detector with a Gd-H_2O fill is statistics limited. Table 13 clearly shows that the 22 m Gd-H_2O detector becomes statistics limited for the single-site Heysham 2 signal. The detectors with the Gd-WbLS fill tend to be more systematics limited, particularly the 22 m configuration, and reducing the systematics on the backgrounds could bring the dwell time down. Reduction of systematic uncertainties on the backgrounds could be achieved through sideband analysis or through measurements with the demonstrator planned for the existing Boulby Underground

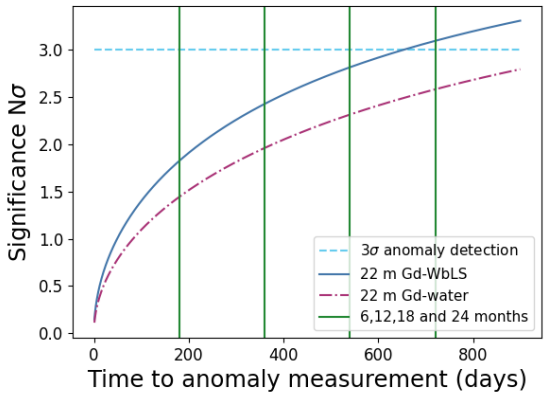

Figure 35: The sensitivity of the 22 m detector, with both the Gd-H_2O and the Gd-WbLS fill, to anomaly detection of the Heysham 2 complex. Shown as a function of observation time with vertical lines marking 6, 12, 18 and 24 months from start of observation.

Table 12: Significance for anomaly detection of Heysham 2 at 6, 12, 18 and 24 months of observation time with the 22 m detector. The Gd-WbLS configuration achieves 3σ significance in under two years.

Fill medium	6 months	12 months	18 months	24 months
Gd-H_2O	1.50	2.03	2.40	2.68
Gd-WbLS	1.83	2.43	2.82	**3.10**

Laboratory (see Section 3.1). Further improvements to the analysis, for example those outlined in Section 8.1, may also lessen the dwell time.

Table 13: Representative rates per day from optimisation in the 22 m detector with Gd-H_2O and Gd-WbLS fill media. Accidental coincidences and geoneutrino IBDs are omitted as they are sub-dominant backgrounds in the Cobraa results.

Gd-H_2O	s	Σb_i	^{9}Li & ^{17}N	Fast n	$IBD_{reactor}$	$\Sigma\sigma_{bi}$
Hartlepool 1	1.4	2.0	1.2×10^{-1}	3.3×10^{-1}	1.6	1.3×10^{-1}
Heysham	0.10	0.26	6.9×10^{-2}	1.1×10^{-2}	1.8×10^{-1}	1.1×10^{-2}
Heysham & Torness	0.10	0.24	7.4×10^{-2}		1.5×10^{-1}	9.6×10^{-3}
Heysham 2	0.059	0.26	6.9×10^{-2}		1.8×10^{-1}	1.1×10^{-2}
Gd-WbLS	s	Σb_i	^{9}Li & ^{17}N	Fast n	$IBD_{reactor}$	$\Sigma\sigma_{bi}$
Hartlepool 1	2.2	3.2	2.7×10^{-1}	5.4×10^{-1}	2.39	2.0×10^{-1}
Heysham	0.18	0.48	1.3×10^{-1}	4.4×10^{-2}	3.0×10^{-1}	2.1×10^{-2}
Heysham & Torness	0.13	0.33	8.5×10^{-2}	3.8×10^{-2}	2.1×10^{-1}	1.6×10^{-2}
Heysham 2	0.091	0.39	9.7×10^{-2}	3.4×10^{-2}	2.6×10^{-1}	1.8×10^{-2}

To date, Cobraa has been demonstrated to achieve excellent sensitivity through its refined simulation and analysis process, particularly in combination with the CoRe reconstruction. Its built-in flexibility ensures it can be readily adapted to new detector designs and signals, for fast evaluation of sensitivity with a variety of signal and detector combinations.

7.4 Summary

- The Cobraa analysis has the flexibility to optimise on signal significance or dwell time. Optimising on dwell time is key when systematic uncertainties become the limiting factor.
- Cobraa is able to evaluate sensitivity using different metrics which apply to different discovery scenarios and signal strengths.
- The sensitivity of the 16 m and 22 m detector configurations has been evaluated both for anomaly detection (rejection of the background-only hypothesis with known backgrounds) and for reactor discovery (rejection of the no-reactor hypothesis with unknown backgrounds).
- Reactor discovery with no knowledge of backgrounds is a challenging scenario but could be achieved with NEO in the 22 m detector if the lifetime of the Hartlepool reactor complex is extended.
- The 22 m Gd-H_2O configuration is able to perform a dual-site anomaly detection within the two-year time constraint and could serve as risk mitigation in the event that Gd-WbLS is not ready for full-detector deployment in NEO.
- The 22 m Gd-WbLS option could achieve a single-site anomaly detection of the Heysham 2 reactor complex within the two years from detector commissioning to reactor decommissioning under the currently-planned schedule.

8 Conclusions and Future Work

What she was now was what she herself had decided to become.

– Michael Ondaatje, *The English Patient*

In the development of the reactor-sensitivity toolchain presented in this thesis, I have addressed a key question for remote reactor monitoring for nuclear non-proliferation:

What is the optimal sensitivity of a given detector to a given reactor signal for application to real-world discovery scenarios?

I will now review the work with reference to this question and the development of the NEO detector. The key results will be discussed in the context of the NEO detector and the research area as a whole, including its implications for the wider remote monitoring research area. Finally, I will consider the limitations of the work and will look at how the work can be further developed in the future - both for the application to non-proliferation and for wider, fundamental physics goals.

WATCHMAN aims to build the NEO detector as the first experiment at AIT to demonstrate the application of reactor antineutrino detection to remote monitoring for non-proliferation. The detector technology selected as the most applicable to remote monitoring, for its balance of sensitivity and scalability, is a gadolinium-doped water-based liquid scintillator detector. A fill with gadolinium-doped water is also being considered to mitigate the risk associated with the development and large-scale deployment of the nascent Gd-WbLS technology. Two sizes of detector are currently under consideration by WATCHMAN for the AIT site in Boulby Mine in the UK: a cylinder 16 m high by 16 m in diameter with an inner PMT radius of 5.7 m and a cylinder 22 m high by 22 m in diameter with an inner PMT radius of approximately 9 m.

The aim of this work was to develop a simulation-reconstruction-analysis toolchain which gives the best possible sensitivity to a remote reactor signal in a gadolinium-doped, water-based detector. This was intended to be used to investigate the capability of reactor antineutrino detection for remote monitoring in a counting experiment and to inform detector design decisions made by the WATCHMAN collaboration.

In the pursuit of optimal sensitivity to a reactor signal, I developed a method to fully exploit the fact that, in a gadolinium-doped water-based detector, the signal consists of a pair of events which are coincident in time and distance.

A modified simulation method was created to reproduce true coincidences of all signal and background events. This constitutes an accurate reproduction of real data and paved the way for a full coincidence reconstruction and analysis.

A new position reconstruction was implemented using the BONSAI combined maximum-likelihood fitter to take the light from two coincident events and reconstruct a combined vertex. This is the first implementation of combined reconstruction for NEO and the first time, to my knowledge, that the combined reconstruction has been fully implemented and applied in an analysis. The new reconstruction improved the vertex resolution by a factor of almost 2 in the Gd-H_2O detectors for 2.5 MeV positrons, at the peak of the reactor signal. The reconstruction improved by 25% and 20% in the 16 m and 22 m Gd-WbLS configurations at the

same energy and the vertex resolution increased by more than 20% at most energies in all detectors.

Finally, I developed an analysis to optimise the sensitivity to a reactor by looking for coincident pairs which have the expected properties of the reactor antineutrino signal. The sensitivity can be optimised as a function of six parameters and includes a new treatment to reject fast-neutron and radionuclide backgrounds.

Sensitivity metrics - in terms of the *dwell time* required to reject a given hypothesis - were developed to assess detector sensitivity to (a) reject the background-only hypothesis and (b) confirm the presence of a reactor signal. I derived dwell time calculations based on Poisson and Gaussian significances to reject the background-only hypothesis. Systematic uncertainties on the background rates were incorporated for the first time and the option to optimise the dwell time rather than the signal significance was introduced to help mitigate the effect of large systematic uncertainties. Dwell time calculations to minimise false positives and false negatives were also derived. Having a variety of sensitivity metrics lends flexibility to assess the application to real-world scenarios.

The reactor sensitivity toolchain was used to evaluate four detector configurations under consideration for deployment in Boulby Mine, and the results are presented in this thesis. The results were considered in the context of planned reactor operations in the UK up until 2030. The three reactor complexes closest to Boulby Mine were considered as signals: Hartlepool, Heysham and Torness. Under the current schedule, the AGR-1 reactors at the Hartlepool and Heysham 1 complexes will be decommissioned in 2024 and the remaining reactors at Heysham and Torness will be decommissioned in 2030. The results showed that the most sensitive configuration - the 22 m tank with a Gd-WbLS fill - could assure a dual-site detection of the reactors at Heysham 2 and Torness, and a single-site detection of the reactor at Heysham 2, in the two years before Heysham 2 and Torness are due to be shutdown. The 22 m Gd-H_2O configuration could achieve a dual-site Heysham 2 & Torness detection within the two years in the event of a delay in the deployment of the Gd-WbLS technology.

The simulation, reconstruction and analysis software has been used to investigate a number of configurations, including those discussed in this thesis. This demonstrates the flexibility of the software, which can be readily adapted to a variety of configurations of water-based Cherenkov detectors.

When the software was first developed, the focus was on detecting the single- and dual-core signals at Hartlepool. Since learning of the likely decommissioning of the reactors at Hartlepool in 2024, I have shifted the focus to more distant reactors. The rapid pivoting of the analysis in this way is a testament to its robustness and flexibility. Additionally, the development of the work in response to these more challenging scenarios has led, in the end, to a better analysis, which has more application to real-world scenarios. In the field, the AIT-NEO technology will be most useful to detect signals at greater distances, against higher backgrounds from other reactors. These are the types of scenario that have been investigated by application of the analysis to the more distant reactors at Heysham and Torness.

8.1 The Future

8.1.1 Cobraa Analysis

Cobraa treats NEO as a counting experiment and would benefit from making use of additional information, for example spectral data. Incorporating a spectral analysis into the process could enhance Cobraa's power to discriminate signal from background. This could be particularly applicable when looking at more distant reactors, since neutrino oscillations change the energy spectrum of the antineutrinos with distance from the reactor source. Indeed, future work evaluating sensitivity to even more distant reactors including Hinkley Point C (which is due to come online in 2025), following the decommissioning of Heysham 2 and Torness, could drive this change.

In implementing the CoRe reconstruction, I incorporated a measure of the isotropy of the light from an event with respect to the reconstructed vertex (see Section 5.3.3) based on the method using Legendre polynomials used by SNO during the salt phase. I have demonstrated this to differentiate well between the positron and the neutron event in the IBD pair, since the light from the neutron event is isotropic, while the light from a positron event is limited to a single Cherenkov cone. The principal motivation behind developing these was to reduce fast neutron backgrounds, where both events in the pair are isotropic. Indeed, the variables have been used successfully in machine learning algorithms [146] [147] to help reject fast neutron backgrounds. These backgrounds are of particular concern in NEO due to the high uncertainty on the rates. Optimisation using the isotropy variables as an additional parameter should now be incorporated into Cobraa to improve fast-neutron pair rejection beyond what has already been achieved.

The reduction of radionuclide backgrounds via a time-after-muon veto has been applied analytically in Cobraa. This is subject to assumptions about the muon-detection efficiency and our ability to track muons and - as a conservative measure - does not draw on the very latest work on muon vetos for radionuclide rejection.

Most radionuclides in a water-based medium are produced by low- to medium-energy hadronic secondaries in muon-induced hadronic showers [130] and the fraction of the muons that shower hadronically is $< 1\%$ of all showers [148] assuming an average muon energy of 270 GeV, similar to the average muon energy in Boulby Mine. According to [148], radioisotope production is higher in high-energy showers. Total radiative muon energy loss along a muon track can be found by subtracting the minimum average muon energy loss and, in the case of multiple showers, the highest-energy shower can be identified using the light profile along a muon track. Muons with radiative energy losses of over 30 GeV, indicative of high-energy showers, make up only 2% of all muons and it was found that they are responsible for 60% of the radioisotopes produced through muon-induced spallation in Super-K. A 20-second transverse cut of 3m along the tracks of only those muons with high energy losses was proposed to cut 58% of isotopes in Super-K with only 4% deadtime.

A more intelligent muon veto for NEO, geometrically confined to the location of a hadronic shower along a muon track, could reduce the ^{9}Li background to negligible and remove most of the ^{17}N. A full muon simulation would be required in order to more accurately assess the effect of such a muon veto on the radionuclide background and detector deadtime.

An important step in the evaluation of the results from Cobraa will be to incorporate detector systematics and to assess the degree of uncertainty in the sensitivity metrics output by Cobraa. In addition, the creation of 'mock data' runs, in the absence of real data, for a blinded analysis of the detector sensitivity and background rejection power will serve to validate the reconstruction and analysis results.

8.1.2 Coincidence Reconstruction

A happy by-product of improved vertex resolution is an improvement in energy resolution and the potential to see down to lower energies. For consistency with the single-event reconstruction, a hit threshold of 10 inner-PMT hits has been used in the CoRe implementation in this thesis. However, given the improvement in the reconstruction of lower-energy positrons with CoRe, it is likely that a lower threshold is possible.

CoRe and the subsequent Cobraa analysis combine to eliminate all or most of the backgrounds due to radioactivity. This, along with the ability to reconstruct well at lower energies, could bring the analysis energy threshold down closer to the DAQ threshold. This should now be studied in detail. If it proved possible to reduce the analysis threshold, the effect would be to increase the sensitivity of a detector and consequently bring down the dwell times output by the Cobraa analysis.

The benefit of improved energy resolution and a lower energy threshold could also be extended to other work. This might include a reactor ranging study that has been performed for WATCHMAN, and which currently uses a maximum likelihood fit to the expected energy spectra at varying distances to constrain the distance to a reactor. The ability to constrain the distance to a reactor would be an important achievement in terms of the non-proliferation goals of the project.

The CoRe reconstruction should now be validated on data. Using a DT generator to produce ^{16}N by irradiating CO_2 (see Section 3.14) would generate a tagged coincident β and neutron pair which could be used in the first instance to test the reconstruction.

A number of variables are built into the BONSAI reconstruction used by CoRe. These include, for example, the correction to the open angle discussed in Section 5.3. This was originally optimised for pure H_2O and may need to be adjusted for Gd-H_2O and, in particular, Gd-WbLS. Developing a method by which the optimisation can be performed automatically for a new detector or fill would add flexibility to the reconstruction software.

Indeed, the best results for CoRe were achieved for Gd-H_2O and there is a need to fully adapt it for Gd-WbLS. This is likely to include incorporating charge data into the fit and could make use of a machine learning algorithm for pair reconstruction, which is already under investigation.

8.1.3 Beyond NEO

Software such as Cobraa could be adapted to be used as a tool for non-proliferation to evaluate the sensitivity of any given detector to any given signal, with some assumptions about background rates and uncertainties on backgrounds. This could be useful in the future, particularly if antineutrino detectors are adopted for remote reactor monitoring beyond NEO and subsequent AIT phases.

Implementation of the combined reconstruction for inverse beta decay in a gadolinium-doped detector could have wider application outside non-proliferation as the emerging Gd-H_2O and Gd-WbLS technologies are adopted. Super-K has already deployed gadolinium in its detector to a 0.01% concentration - a tenth of the planned final concentration. In Super-K, searches such as the hunt for supernova relic neutrinos rely on inverse beta decay and the addition of gadolinium is seen as vital in this search. Expected event rates are between 2 and 6 per year above 10 MeV [78] and improving the vertex reconstruction and, perhaps more significantly, background rejection with an implementation of the combined reconstruction could help to increase the number of supernova relic neutrinos detected.

8.2 Summary

- The NEO detector is the planned first experiment at the AIT facility and is intended to demonstrate the ability to use antineutrino detection to monitor nuclear reactors for non-proliferation.
- A new simulation-reconstruction-analysis chain was developed for the evaluation of sensitivity to the coincident signal from reactor antineutrinos for remote monitoring.
- The implementation and optimisation of a vertex reconstruction, using the light from two events in a coincident IBD event pair, significantly improved the vertex resolution for positrons and neutrons in the IBD signal.
- A new analysis to assess detector sensitivity to remote reactor antineutrino signals has been developed. This flexible software has been used to evaluate the sensitivity of a number of detector configurations and some of these results are presented in this thesis.
- Sensitivity metrics were developed to address the background-only and no-reactor hypotheses with known and unknown backgrounds for application to discovery scenarios that might be encountered in the field.
- Sensitivity to near and distant reactors was investigated to assist design and optimisation of the NEO detector.
- The software toolchain represents the most comprehensive analysis package for remote reactor detection for WATCHMAN to date and provides a strong basis for wider applications beyond NEO and the WATCHMAN collaboration.

Appendices

A Gain calibration

The charge response signal has two contributions: the photoelectric conversion at the photocathode and the amplification at the dynodes.

The photoelectric conversion is the convolution of a Poisson process for the number of photons hitting the PMT and a binary process for the photoelectric conversion. The resulting distribution is Poissonian

$$P(n, \mu) = \frac{\mu^n e^{-\mu}}{n!} \tag{A.1}$$

This is the probability that n photoelectrons will be observed when the mean number of photoelectrons collected at the first dynode is μ.

The response of the dynode system is a Poisson distribution but if the amplification is at least 4 per stage (and preferably > 10), then the n-*photoelectron peaks* can be approximated with Gaussian functions:

$$G_n(x) = \sum_{n=0}^{\infty} \frac{1}{\sigma_{SPE}\sqrt{2\pi n}} \exp\left\{\left(-\frac{(x - nQ_{SPE})^2}{2n\sigma_{SPE}^2}\right)\right\} \tag{A.2}$$

where σ_{SPE} is the width of the single-photoelectron (SPE) distribution, x is the charge variable and Q_{SPE} is the SPE charge output.

The ideal PMT response is then a convolution of the capture and photoelectric conversion with the amplification:

$$S_{ideal}(x) = \sum_n P(n, \mu) G_n(x) \tag{A.3}$$

There are two types of background. The low-charge processes which occur when no photoelectron is emitted from the photocathode give rise to the *pedestal* at around zero charge, which has a Gaussian distribution. The processes which can accompany the signal e.g. spontaneous emission from the photocathode and dynodes give rise to the *valley* between the pedestal and single-electron peak. These can be approximated by an exponential. The backgrounds are the sum of these processes:

$$B(x) = P(0, \mu)\frac{1}{\sigma_0\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma_0^2}\right) + P_{exp}\theta(x)\alpha \exp(-\alpha x) \tag{A.4}$$

where $P(0, \mu)$ is the Poisson probability that 0 photoelectrons are produced, σ_0 is the width of the pedestal, P_{exp} is the probability that the second type of background is present and α is the coefficient of the exponential decrease of the valley.

$\theta = \begin{cases} 0 & x \leq 0 \\ 1 & x > 0 \end{cases}$ ensures that there is an exponential component only if $x > 0$.

The total charge distribution can therefore be modelled as follows:

$$Q(x) = P(0, \mu)G(x)_{ped} + P_{exp}Exp_{val} + \sum_{n=0}^{\infty} P_n G(x)_n \tag{A.5}$$

where G_{ped} and G_n are the Gaussian fits to the pedestal and n-photoelectron peaks, the Poisson distributions give the relative fractions of each peak and Exp_{val} is the exponential fit to the valley.

The single photoelectron (SPE) charge output Q_{SPE} was extracted from this fit to the total charge distribution of integrated waveforms (Figure 36a) and the gain was then calculated as follows:

$$\text{Gain} = \frac{Q_{SPE} \times \text{splitting} \times 10^{-12}}{e \times \text{amplification} \times \text{impedance}} \tag{A.6}$$

where signal amplification prior to readout amplification $= 10$, cable impedance impedance $= 50\,\Omega$ and splitting $= 2$, which takes into account the splitting of signal and high voltage from the PMT's single SHV cable.

Each PMT was tested at five voltage values in steps of 100 V in a range around the nominal voltage quoted in the Hamamatsu shipping data to give 10^7 gain. The gain data were then fit according to equation (A.7)

$$\text{Gain} = \left(\frac{V}{V_{opt}}\right)^{\beta} \tag{A.7}$$

where V_{opt} is the voltage for 10^7 gain and both V_{opt} and β are free parameters in the fit [149]. The *operating voltage* required to give 10^7 gain was then calculated from the fit for each PMT (Figure 36b).

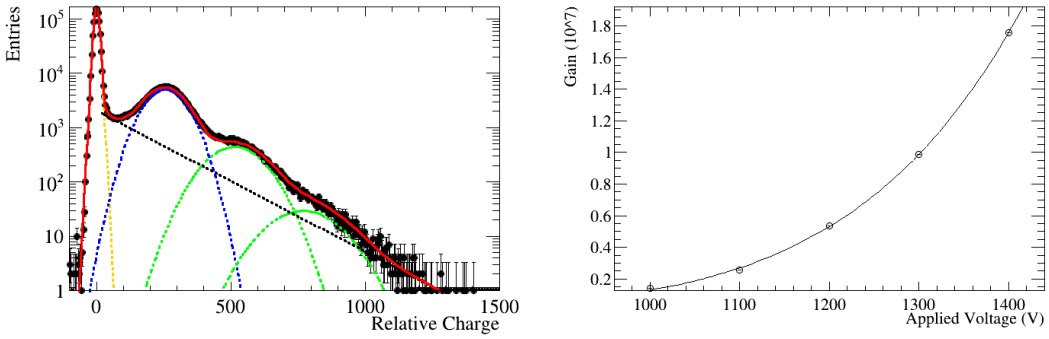

(a) Fit to charge spectrum.

(b) Fit to gain curve.

Figure 36: Results based on full fit to the charge distribution with double-Gaussian fit to the pedestal (yellow), Gaussian approximations to single- (blue), double- and triple-photoelectron (green) peaks and exponential approximation (black) to the valley. Error bars on gain-curve fit are within the circular markers.

B Signal and Background Rates

Table 14: Radioactive isotopes present in detector components and cavern. These rates are taken either from WATCHMAN or other testing data or manufacturer specifications, as referenced.

Component	Isotope	Activity (ppm)
PMTs [150]	^{238}U	0.064
	^{232}Th	0.172
	^{40}K	0.01
Gd-H_2O	^{238}U	8.10×10^{-8}
	^{232}Th	2.46×10^{-8}
	^{235}U	8.20×10^{-8}
	^{40}K	1.29×10^{-4}
PSUP stainless steel 316	^{238}U	9.49×10^{-3}
	^{235}U	8.38×10^{-5}
	^{232}Th	4.19×10^{-3}
	^{40}K	1.75
	^{60}Co	1.79×10^{-12}
	^{137}Cs	2.47×10^{-11}
Tank stainless steel 316	^{238}U	9.49×10^{-3}
	^{232}Th	4.19×10^{-3}
	^{40}K	1.75
	^{60}Co	1.79×10^{-12}
	^{137}Cs	2.47×10^{-11}
Concrete	^{238}U	4.94
	^{232}Th	7.38
	^{40}K	1.59×10^{4}
Rock	^{238}U	10×10^{-3}
	^{232}Th	220×10^{-3}
	^{40}K	750×10^{-3}

Table 15: β decays due to radioactive isotopes. Decays with a branching fraction below 0.1% or an endpoint energy below 0.5 MeV, and as such well below the reasonably attainable detection threshold, are not of concern for NEO.

Isotope	Decay	Total endpoint energy
^{238}U	^{234}Th	0.273 MeV (omitted)
	^{234}Pa	2.19 MeV
	^{214}Pb	1.02 MeV
	^{214}Bi	3.27 MeV
	^{210}Bi	1.16 MeV
	^{210}Tl	5.48 MeV (0.02% branching fraction)
^{232}Th	^{228}Ra	0.046 MeV (omitted)
	^{228}Ac	2.12 MeV
	^{212}Pb	0.57 MeV
	^{212}Bi	2.25 MeV
	^{208}Tl	5.00 MeV
^{235}U	^{231}Th	0.3916 MeV (omitted)
	^{227}Ac	0.048 MeV (omitted)
	^{223}Fr	1.15 MeV
	^{211}Pb	1.37 MeV
	^{211}Bi	0.57 MeV
	^{207}Tl	1.42 MeV
^{40}K	^{40}K	1.31 MeV
^{60}Co	^{60}Co	2.82 MeV
^{137}Cs	^{137}Cs	1.18 MeV

Table 16: Singles rates for uncorrelated backgrounds due to radioactive isotopes in the detector and rock or radiogenic neutrons in the rock for the 16 m Gd-H_2O detector with 15% photocoverage. A trigger is any event which passes the DAQ trigger of 6 PMT hits in 200 ns. Singles rates include all events which pass basic threshold energy and fiducial cuts.

Component	Origin	Principal contributor	Trigger rate (Hz)	Singles rate (Hz)
Gd-H_2O	^{232}Th	^{208}Tl	1.04×10^{-1}	2.78×10^{-2}
Gd-H_2O	^{235}U	^{207}Tl	8.37×10^{-3}	7.43×10^{-5}
Gd-H_2O	^{238}U	^{214}Bi	1.33	1.54×10^{-1}
IBeam	^{232}Th	^{208}Tl	1.17×10^{-1}	1.08×10^{-2}
IBeam	^{238}U	^{214}Bi	1.38×1^{-1}	3.78×10^{-3}
IBeam	40K	40K	1.47×10^{-3}	0
PMT	^{232}Th	^{208}Tl, ^{212}Bi	760	26.0
PMT	^{238}U	^{214}Bi	1120	7.20
PMT	^{40}K	^{40}K	902	4.03×10^{-1}
PSUP	^{232}Th	^{208}Tl	10.1	9.78×10^{-1}
PSUP	^{238}U	^{214}Bi	29.3	9.22×10^{-1}
PSUP	^{40}K	^{40}K	1.04	6.87×10^{-3}
Rock	^{232}Th	^{208}Tl	5.50	4.20×10^{-1}
Rock	Radiogenic	Neutrons	1.85×10^{-2}	5.43×10^{-3}
Tank	^{232}Th	^{208}Tl	7.65×10^{-1}	6.36×10^{-2}
Tank	^{238}U	^{214}Bi	1.02	4.04×10^{-2}
Tank	^{40}K	^{40}K	1.49×10^{-2}	1.92×10^{-4}

Table 17: IBD signal and background rates for the signals under consideration. Rates in NIU neutrino interactions per 10^{31} target particles [60]. Mean power as reported by IAEA used for calculation of the rates. Reactor background for all but Hartlepool signals assumes reactors at Hunterston and Hinkley Point B off after 2022 and Hinkley Point C on after 2025.

Signal	Signal rate	Reactor background	Geoneutrino background
Hartlepool 1 & 2	1041	206.6	36
Hartlepool 1	595	652.6	36
Hartlepool 2	446	801.6	36
Heysham 1 & 2	69.6	146.7	36
Heysham 2	39.6	146.7	36
Heysham & Torness	64.6	121.7	36
Torness	25	161.3	36

C Isotropy variables

This derivation has been adapted from [151].

Legendre polynomials $P_n(x)$, $n = 0, 1, 2, \ldots$ are a complete set of orthogonal functions defined, on the interval $-1 \leq x \leq 1$, for $x = \cos\theta$ as:

$$p_n(\cos\theta) = \frac{1}{2^n n!}\frac{d^n}{d(\cos\theta)^n}(\cos^2\theta - 1)^n \tag{C.1}$$

where:

$$p_1 = \cos\theta \tag{C.2}$$

$$p_2 = \frac{1}{2}(3\cos^2\theta - 1) \tag{C.3}$$

$$p_2 = \frac{1}{2}(5\cos^3\theta - 3\cos\theta) \tag{C.4}$$

and so on.

A function $f(x)$ defined on the interval from -1 to 1 can be expanded as a weighted sum of Legendre polynomials:

$$f(\cos\theta) = \sum_{n=0}^{\infty} A_n P_n(\cos\theta). \tag{C.5}$$

Legendre polynomials obey the orthogonality relation (using the normalisation $P_n(1) = 1$):

$$\int_{-1}^{1} P_n(x)P_m(x)dx = \frac{2}{2n+1}\delta_{mn}, \tag{C.6}$$

where δ_{mn} is the Kronecker delta. Multiplying both sides of Equation C.5 by $P_m(x)$ and integrating:

$$\int_{-1}^{1} P_m(x)f(x)dx = \sum_{n=0}^{\infty} A_n \int_{-1}^{1} P_n(x)P_m(x)dx \tag{C.7}$$

then combining this with Equation C.6:

$$\int_{-1}^{1} P_m(x)f(x)dx = \sum_{n=0}^{\infty} A_n \frac{2}{2m+1}\delta_{mn} \tag{C.8}$$

$$= \frac{2}{2m+1}A_m \tag{C.9}$$

and the definition for the coefficients A_n is obtained:

$$A_n = \frac{2n+1}{2}\int_{-1}^{1} f(\cos\theta)P_n(\cos\theta)dx. \tag{C.10}$$

To measure the isotropy of the light from an event, PMT hits are taken to be random events in 3D space, associated with points on a sphere with its centre at the interaction vertex

(reconstructed or otherwise). The angle θ_{ij} is the angular separation between pairs of hits, calculated from the vertex, and $x = \cos\theta_{ij}$ is in the interval -1 to 1.

The function $f(x)$ is a pdf for the angular separation, so the probability that x takes some value must be 1. Therefore:

$$\int_{-1}^{1} f(x)dx = 2A_0 = 1 \tag{C.11}$$

$$\text{i.e. } A_0 = 1/2 \tag{C.12}$$

In the isotropic case, the Legendre polynomial expansion of the pdf for $x = \cos\theta$ contains only the term $A_0 = \frac{1}{2}$. In the anisotropic case, the coefficients A_n provide information about the degree and nature of anisotropy. If we assume that we have a large number of independent samples S of $\cos\theta$, and that $\cos\theta$ is random, then the integral of $f(x)P(\cos\theta)$ is the mean value of $P(\cos\theta)$ so that:

$$A_n = \frac{2n+1}{2}\frac{1}{S}\sum_{S=1}^{S} P_n(\cos\theta_{ij}). \tag{C.13}$$

To then calculate the coefficient over the angles between all hits in an event we have the expression:

$$\beta_n = \frac{2}{N(N-1)}\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} P_n(\cos\theta_{ij}) \qquad i \neq j \tag{C.14}$$

where N is the number of PMT hits and the number of distinct samples S for $i \neq j$ is $\frac{N(N-1)}{2}$ since $\theta_{ij} = \theta_{ji}$. We remove the factors $\frac{2n+1}{2}$ because they have no effect on the final result.

D Sensitivity metrics

To find our sensitivity in terms of the significance of the signal over the background, we make use of the likelihood ratio

$$\lambda(\mathrm{st}) = \frac{\mathcal{L}(\mathrm{st})}{\mathcal{L}(\hat{\mathrm{st}})} \tag{D.1}$$

where t is the observation time, st is the number of signal counts observed and $\hat{s}t$ is the expected (most likely) number of signal counts. Here, $\mathcal{L}(st)$ and $\mathcal{L}(\hat{s}t)$ are the likelihoods for the signal based, for example, on the Gaussian or Poisson probabilities.

With sufficiently large counts $st + bt$, we can use Wilk's theorem[13] and define the significance:

$$\mathrm{Z}(\mathrm{st}) = \sqrt{-2\ln\lambda(\mathrm{st})}. \tag{D.2}$$

D.1 Gaussian significance

Cobraa models the NEO experiment as a Poisson counting experiment. For a large number of observed events $n > 20$ to 30, the Poisson distribution approximates a Gaussian distribution. We can then use the likelihood ratio method to calculate a sensitivity metric in the form of a Gaussian significance. function f The Gaussian likelihood for the signal $\mathcal{L}_s$, with total signal counts st and total background counts $\Sigma b_i t$ summed over all backgrounds, is defined as:

$$\mathcal{L}_{\mathrm{s}} = \frac{1}{\sqrt{2\pi(\mathrm{st} + \Sigma\mathrm{b_i t})}} \exp\left\{-\frac{(\mathrm{N} - (\mathrm{st} + \Sigma\mathrm{b_i t}))^2}{2(\mathrm{st} + \Sigma\mathrm{b_i t})}\right\}, \tag{D.3}$$

where N is the total number of counts observed in time t. The likelihood ratio for the signal to determine the experimental sensitivity to reject the background-only hypothesis where $st=0$ is then:

$$\lambda_0 = \frac{\mathcal{L}_{\mathrm{s}}(\mathrm{s} = 0)}{\mathcal{L}_{\mathrm{s}}(\hat{\mathrm{st}})} \tag{D.4}$$

and for the Asimov data set, where the observed values are equal to the expected values, we replace N with $st+bt$, this becomes:

$$\lambda_0 = \sqrt{\frac{\mathrm{st} + \Sigma\mathrm{b_i t}}{\Sigma\mathrm{b_i t}}} \exp\left\{-\frac{(\mathrm{st})^2}{2\Sigma\mathrm{b_i t}}\right\} \tag{D.5}$$

so that the signal significance is:

[13] Wilk's theorem says that as the sample size $\mathrm{N} \to \infty$, $-2\ln\lambda$ asymptotically approaches the χ^2 distribution under the null hypothesis.

$$Z = \sqrt{-2\ln\lambda_0} = \sqrt{\frac{(st)^2}{\Sigma b_i t} - \ln\left(\frac{st + \Sigma b_i t}{\Sigma b_i t}\right)}. \tag{D.6}$$

Where $st \ll bt$, this simplifies to:

$$Z = \frac{st}{\sqrt{\Sigma b_i t}}. \tag{D.7}$$

Accounting for the systematic error on the background rates, where the total background systematic by replacing the statistical error $\sqrt{\Sigma b_i t}$ with the statistical and systematic errors added in quadrature, we have our final Gaussian significance to reject the background-only hypothesis:

$$Z = \frac{st}{\sqrt{\Sigma_{b_i} t + (\Sigma\sigma_{b_i} t)^2}}, \tag{D.8}$$

where σ_{b_i} is the systematic uncertainty on the i^{th} background and $\Sigma\sigma_{b_i}$ is the total systematic uncertainty on the background rates. This can then be solved for t to give the required dwell time for a 3σ rejection of the null hypothesis:

$$t_{dwell} = \frac{N_\sigma^2 \Sigma b_i}{s^2 - N_\sigma^2 \Sigma\sigma^2_{b_i,sys}} \tag{D.9}$$

D.2 Poisson significance

The Gaussian approximation in Equation D.8 is valid only for high counts where $st \ll bt$, and tends to overestimate the significance where these criteria are not met, as demonstrated by Figure 37.

Where the counts are low or the signal rate is not significantly less than the background rate, it is necessary to revert to the Poisson significance for the sensitivity metric. The Poisson likelihood is expressed as:

$$\mathcal{L}_s = (st + \Sigma b_i t)^N \frac{e^{-(st+\Sigma b_i t)}}{N!} \tag{D.10}$$

To reject the background-only hypothesis where *st=0*, the likelihood ratio for the Asimov data set is thus:

$$\lambda_0 = \left(\frac{\Sigma b_i t}{st + \Sigma b_i t}\right)^{(st+\Sigma b_t)} e^{st} \tag{D.11}$$

and the approximate Poisson significance is given by:

$$Z = \sqrt{2\left((st + \Sigma b_i t)\ln\left(\frac{\Sigma b_i t}{st + \Sigma b_i t}\right) - st\right)} \tag{D.12}$$

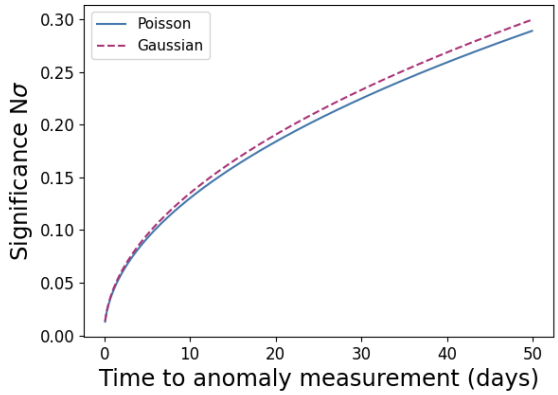

(a) Low signal s«b

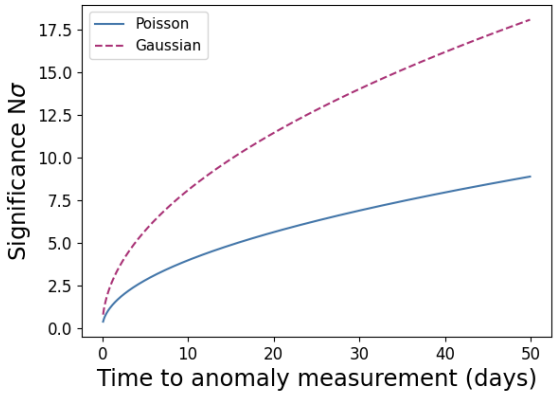

(b) High signal s>b

Figure 37: Comparison of Gaussian and Poisson significance for $s \ll b$ (left) and $s > b$ (right). The Gaussian approximation is only valid where the signal is low and the counts are high.

D.2.1 Poisson uncertainty on backgrounds

Where the backgrounds are taken from a measurement with the reactor off, the systematic uncertainty on the background can be treated as a Poisson-distributed nuisance parameter as discussed in [152] such that the likelihood becomes:

$$\mathcal{L}_s(st, \Sigma b_i t) = (st + \Sigma b_i t)^N \frac{e^{-(st+\Sigma b_i t)}}{N!} (\tau \Sigma b_i t)^M \frac{e^{-\tau \Sigma b_i t}}{M!} \tag{D.13}$$

where $M \sim \tau \Sigma b_i t$ and τ is the ratio between the observation times t_M/t_N in the measurement of the backgrounds with no signal present e.g. during reactor refuelling. In this case, the likelihood ratio becomes:

$$\lambda(s, b) = \frac{\mathcal{L}(s, \hat{B}(s))}{L(\hat{s}, \hat{b})}, \tag{D.14}$$

where $\hat{B}$ is the maximum likelihood value of background rate b given a signal rate s and $\hat{s}$ and $\hat{b}$ are the expected values of s and b.

Using $Z = -2ln\lambda_0$ (where λ_0 is defined above), the significance Z is:

$$Z = -2\left[n \ln\left(\frac{\hat{B}(0)}{\hat{s}t + \Sigma \hat{b}_i t}\right) - \hat{B}(0) + \hat{s}t + \Sigma \hat{b}_i t + m \ln\left(\frac{\hat{B}}{\Sigma \hat{b}_i t}\right)\right]^{1/2}. \tag{D.15}$$

Using $\hat{s} = n - m/\tau$, $\hat{b} = m/\tau$ and $\hat{B}(0) = \frac{n+m}{1+\tau}$, this becomes:

$$Z = -2\left[n \ln\left(\frac{n+m}{n(n+m)}\right) + m \ln\left(\frac{\tau(n+m)}{m(1+\tau)}\right)\right]^{1/2} \tag{D.16}$$

Expressing tau in terms of the uncertainty on the background σ_{b_i}, $\tau = \frac{\Sigma b_i t}{(\sigma_{b_i} t)^2}$, the approximate Poisson significance with Poisson uncertainty on the background can be expressed as:

$$Z = \left[2 \left((st + \Sigma b_i t) \ln \left[\frac{(st + \Sigma b_i t)(\Sigma b_i t + (\Sigma \sigma_{b_i} t)^2)}{(\Sigma b_i t)^2 + (st + \Sigma b_i t)(\Sigma \sigma_{b_i} t)^2} \right] - \frac{(\Sigma b_i t)^2}{(\Sigma \sigma_{b_i} t)^2} \ln \left[1 + \frac{(\Sigma \sigma_{b_i} t)^2 st}{\Sigma b_i t (\Sigma b_i t + (\Sigma \sigma_{b_i} t)^2)} \right] \right) \right]^{1/2} \quad \text{(D.17)}$$

D.2.2 Gaussian uncertainty on backgrounds

Uncertainty on the backgrounds may be Gaussian, particularly where the background rates come from a model or other process. This is the simplifying assumption made for the uncertainties on background rates currently included in Cobraa. In the case of a Gaussian uncertainty on a background rate, given a rate for the i^{th} background of $b_i \pm \sigma_{b_i}$, the uncertainty is introduced as a Gaussian-distributed nuisance parameter, as discussed in [153]. This results in a likelihood which is the product of a Poisson and a Gaussian distribution:

$$\mathcal{L}_s = (st + \Sigma B_i t)^N \frac{e^{-(st + \Sigma B_i t)}}{N!} \frac{1}{\sqrt{2\pi \Sigma (\sigma_{b_i} t)^2}} \exp \left\{ - \frac{(\Sigma b_i t - \Sigma B_i t)^2}{2\Sigma (\sigma_{b_i} t)^2} \right\} \quad \text{(D.18)}$$

where $B_i t$ is the number of background counts observed.

Using the likelihood ratio λ_0 for the Asimov data set we generate an expression for the Poisson significance with a Gaussian-distributed background:

$$Z = \sqrt{-2 \ln \lambda_0} = \left[-2 \left((st + \Sigma b_i t) \ln \frac{\hat{B}_0}{st + \Sigma b_i t} - \frac{(\Sigma b_i t - \hat{B}_0)^2}{2\Sigma (\sigma t)^2} - \hat{B}_0 + st + \Sigma b_i t \right) \right]^{\frac{1}{2}} \quad \text{(D.19)}$$

where the background rate $\hat{B}_0$ for $st = 0$ is the positive solution for the maximum of $\mathcal{L}_0$:

$$\hat{B}_0 = \frac{1}{2} \left(\Sigma b_i t - (\Sigma \sigma_{b_i} t)^2 + \sqrt{(\Sigma b_i t)^2 - 2\Sigma b_i t (\Sigma \sigma_{b_i} t)^2 + 4(st + \Sigma b_i t)(\Sigma \sigma_{b_i} t)^2 + (\Sigma \sigma_{b_i} t)^4} \right). \quad \text{(D.20)}$$

The two equations D.19 and D.17 for the significance can be expressed as a function of a free parameter t and in this way the value for t_{dwell} can be extracted for a given significance.

D.3 Knoll Method

A dwell time calculation to give a 3σ positive detection significance with a confidence level CL = 95% *in the absence of any knowledge of the reactor or backgrounds* can be derived using the Knoll derivation in [145].

In this scenario, systematic errors on assumed background rates are replaced with large statistical errors on unknown background rates. Systematic errors due to detector effects and reconstruction will remain but for now we assume statistical errors only.

In this derivation, we consider two observation periods T_{on} and T_{off}. The total counts observed for the time T_{on} that the reactor is on (or suspected to be on) would equate to:

$$N_{on} = (s + b)T_{on} \tag{D.21}$$

for signal rate *s* and background rate *b*. The total counts observed during the time T_{off} that the reactor is off or suspected to be off is:

$$N_{off} = bT_{off} \tag{D.22}$$

The fraction of the observation time with the reactor on or suspected to be on is

$$R_{onOff} = T_{on}/T_{off}. \tag{D.23}$$

If the number of counts in the total number of events N_{on} and the total number of background events N_{off} observed is sufficient, the Poisson distributions can be modelled with Gaussian distributions. Since this type of measurement will require measurement over longer periods, it is reasonable to make this assumption in the calculation of the significance.

Knoll defines two values for the signal count, both given by $N_s = N_{tot} - N_{b,tot}$. These values are the number of 'signal' events observed with no reactor present and the number of signal events observed in the presence of a reactor. With no reactor present the number of signal events detected is

$$N_{s,noreactor} = N_{tot,noreactor} - N_{b,tot}$$

which gives a Gaussian distribution around zero with statistical error:

$$\sigma_{N_{s,noreactor}} = \sqrt{\sigma^2_{N_{tot,noreactor}} + \sigma^2_{N_b}} \tag{D.24}$$

$$= \sqrt{2\sigma^2_{N_b}} \tag{D.25}$$

$$= \sqrt{2bT_{off}(1/R_{onOff} + 1)} \tag{D.26}$$

where we replace T_{on} with $1/R_{onOff}$ to simplify the later calculation. With a reactor present the number of signal events detected is

$$N_{s,reactor} = N_{tot,reactor} - N_{b,tot}$$

which gives a Gaussian distribution around the mean 'signal' count. The statistical error on this count is:

$$\sigma_{N_{s,reactor}} = \sqrt{\sigma^2_{N_{tot,reactor}} + \sigma^2_{N_b}} \tag{D.27}$$

$$= \sqrt{(s + b)T_{off}/R_{onOff} + 2bT_{off}(1/R_{onOff} + 1)} \tag{D.28}$$

For a 95% confidence level of a positive detection, we wish to minimise false positives and so require a maximum 5% probability of a false positive (we are not concerned at this point with the false negatives since these will just be classed as no detection). There is a 90% probability that a random sample from a Gaussian distribution will lie within 1.645σ of the mean. Thus

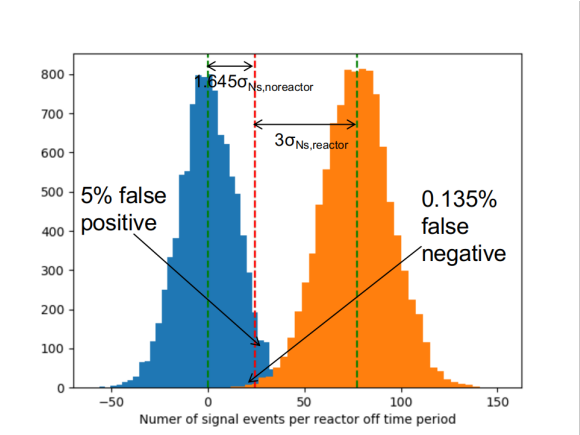

Figure 38: Signal events required for 3σ positive detection of a reactor at the 95% confidence level, using arbitrary but representative signal and background rates.

there is a 5% probability of a false positive and thus for 95% confidence we have $Z_{CL} = 1.645$. We then take $1.645\sigma_{N_{st,noreactor}}$ to be the number of events we require to reach 95% confidence, as shown in Figure 38. The number of signal events corresponding to this confidence level then acts as the trigger threshold for the 3σ detection in the presence (suspected or otherwise) of a reactor, since below this level we have 95% confidence that there is no reactor.

We then wish to minimise false negatives. By selecting a 3σ positive detection, we reduce the chance of a false negative to $\leq 0.135\%$ so that the number of events above the 95% confidence-level threshold is $3\sigma_{N_{st,reactor}}$ as shown in Figure 38. In this way, 99.87% of the signal events counted in the presence of a reactor are above the threshold and are taken to be a positive detection of a reactor.

This gives us the number of signal events N_s required for 3 sigma positive detection at 95% confidence:

$$N_s = 1.645\sigma_{N_{s,noreactor}} + 3\sigma_{N_{st,reactor}} \tag{D.29}$$

Using the definitions given in Equations D.26 and D.28 above, this becomes:

$$N_s = 1.645\sqrt{2bT_{off}(1/R_{onOff} + 1)} + 3\sqrt{(s+b)T_{off}/R_{onOff} + 2bT_{off}(1/R_{onOff} + 1)}. \tag{D.30}$$

Rearranging this in terms of T_{off} and obtaining the equivalent expression in terms of T_{on}, we have an expression for the dwell time:

$$t_{dwell} = \left(\frac{1.645\sqrt{2b(R_{onOff}+1)} + 3\sqrt{sR_{onOff} + 2b(R_{onOff}+1)}}{sR_{onOff}}\right)^2 + \left(\frac{1.645\sqrt{2b(1+1/R_{onOff})} + 3\sqrt{s/R_{onOff} + 2b(1+1/R_{onOff})}}{s}\right)^2 \tag{D.31}$$

References

[1] Department for Business Energy and Industrial Strategy. *The energy white paper: Powering our net zero future* (2020). (Page 10)

[2] International Atomic Energy Agency. *Energy, Electricity and Nuclear Power Estimates for the Period up to 2050* (2020). (Page 10)

[3] United Nations. *Treaty on the Non-Proliferation of Nuclear Weapons* (1970). (Page 10)

[4] International Atomic Energy Agency. *Model Protocol Additional to the Agreements(s) between the State(s) and the International Atomic Energy Agency for the Application of Safeguards* (1997). (Page 10)

[5] NNSA. *Strategic Plan - Making the world a safer place* (2011). (Page 10, 41)

[6] WATCHMAN Collaboration. *The AIT-NEO Project: Results from the NEO Performance Trade Study* (2021). Unpublished. (Page 11, 46, 54, 93)

[7] WATCHMAN Collaboration UK. *Alternative reactor sensitivity study: risk mitigation for the NEO project* (2021). Unpublished. (Page 11, 42, 85)

[8] L. M. Brown. The idea of the neutrino, *Physics today*, **31**(9):23–28 (1978); doi:10.1063/1.2995181 . (Page 12)

[9] E. Fermi. Tentativo di una teoria dei raggi β, *Il Nuovo Cimento*, 11:1–19 (1934); doi:10.1007/BF02959820 . (Page 12)

[10] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions, *Phys. Rev.*, 104:254–258 (1956); doi:10.1103/PhysRev.104.254 . (Page 12)

[11] C. S. Wu, E. Ambler, R. W. Hayward, et al. Experimental test of parity conservation in beta decay, *Phys. Rev.*, 105:1413–1415 (1957); doi:10.1103/PhysRev.105.1413 . (Page 12)

[12] E. C. G. Sudarshan and R. E. Marshak. Chirality invariance and the universal Fermi interaction, *Phys. Rev.*, 109:1860–1862 (1958); doi:10.1103/PhysRev.109.1860.2 . (Page 12)

[13] J. J. Sakurai. Mass reversal and weak interactions, *Nuovo Cim.*, 7:649–660 (1958); 10.1007/BF02781569 . (Page 12)

[14] R. P. Feynman and M. Gell-Mann. Theory of the Fermi interaction, *Phys. Rev.*, 109:193–198 (1958); doi:10.1103/PhysRev.109.193 . (Page 12)

[15] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos, *Phys. Rev.*, 109:1015–1017 (1958); doi:10.1103/PhysRev.109.1015 . (Page 12)

[16] C. L. Cowan, F. Reines, F. B. Harrison, et al. Detection of the free neutrino: a confirmation, *Science*, **124**(3212):103–104 (1956); doi:10.1126/science.124.3212.103 . (Page 13, 31)

[17] C. L. Cowan and F. Reines. The neutrino, *Nature* (1956); doi:10.1038/178446a0 . (Page 13)

[18] E. J. Konopinski and H. M. Mahmoud. The universal Fermi interaction, *Phys. Rev.*, 92:1045–1049 (1953); doi:10.1103/PhysRev.92.1045 . (Page 13)

[19] R. Davis. Attempt to detect the antineutrinos from a nuclear reactor by the $Cl^{37}(\overline{\nu}, e^-)Ar^{37}$ reaction, *Physics Review*, 97:766–769 (1955); doi:10.1103/PhysRev.97.766 . (Page 14)

[20] R. Davis, Jr. and D. S. Harmer. Attempt to observe the $Cl^{37}(\bar{\nu}e^-)Ar^{37}$ reaction induced by reactor antineutrinos, *Bull. Am. Phys. Soc.*, 4:217 (1959); . (Page 14)

[21] G. Danby, J.-M. Gaillard, K. Goulianos, et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos, *Phys. Rev. Lett.*, 9:36–44 (1962); doi:10.1103/PhysRevLett.9.36 . (Page 14, 20)

[22] K. Kodama, N. Ushida, C. Andreopoulos, et al. Observation of tau neutrino interactions, *Physics Letters B*, **504**(3):218–224 (2001); doi:10.1016/s0370-2693(01)00307-0 . (Page 14)

[23] M. Acciarri, O. Adriani, M. Aguilar-Benitez, et al. Determination of the number of light neutrino species from single photon production at LEP, *Physics Letters B*, **431**(1):199–208 (1998); doi:10.1016/S0370-2693(98)00519-X . (Page 15)

[24] G. Steigman, D. N. Schramm, and J. E. Gunn. Cosmological limits to the number of massive leptons, *Physics Letters B*, **66**(2):202–204 (1977); doi:10.1016/0370-2693(77)90176-9 . (Page 15)

[25] S. L. Glashow. Partial-symmetries of weak interactions, *Nuclear Physics*, **22**(4):579–588 (1961); doi:10.1016/0029-5582(61)90469-2 . (Page 15)

[26] P. W. Higgs. Broken symmetries and the masses of gauge bosons, *Phys. Rev. Lett.*, 13:508–509 (1964); doi:10.1103/PhysRevLett.13.508 . (Page 15)

[27] F. Englert and R. Brout. Broken symmetry and the mass of gauge vector mesons, *Phys. Rev. Lett.*, 13:321–323 (1964); doi:10.1103/PhysRevLett.13.321 . (Page 15)

[28] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global conservation laws and massless particles, *Phys. Rev. Lett.*, 13:585–587 (1964); doi:10.1103/PhysRevLett.13.585 . (Page 15)

[29] S. Weinberg. A model of leptons, *Phys. Rev. Lett.*, 19:1264–1266 (1967); doi:10.1103/PhysRevLett.19.1264 . (Page 15)

[30] A. Salam. Weak and electromagnetic interactions, *Conf. Proc. C*, 680519:367–377 (1968); doi:10.1142/9789812795915_0034 . (Page 15)

[31] F. Hasert, H. Faissner, W. Krenz, et al. Search for elastic muon-neutrino electron scattering, *Physics letters. B*, **46**(1):121–124 (1973); doi:10.1016/0370-2693(73)90494-2 . (Page 17)

[32] F. Hasert, S. Kabe, W. Krenz, et al. Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment, *Physics letters. B*, **46**(1):138–140 (1973); doi:10.1016/0370-2693(73)90499-1 . (Page 17)

[33] G. Arnison, A. Astbury, C. Bacci, et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at s=540 GeV, *Physics letters. B*, **122**(1):103–116 (1983); doi:10.1016/0370-2693(83)91177-2 . (Page 17)

[34] M. Banner, R. Battiston, P. Bloch, et al. Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider, *Physics letters. B*, **122**(5-6):476–485 (1983); doi:10.1016/0370-2693(83)91605-2 . (Page 17)

[35] P. Bagnaia, M. Banner, R. Battiston, et al. Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ collider, *Physics Letters B*, **129**(1):130–140 (1983); doi:10.1016/0370-2693(83)90744-X . (Page 17)

[36] G. Aad, T. Abajyan, B. Abbott, et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, *Physics Letters B*, **716**(1):1–29 (2012); doi:10.1016/j.physletb.2012.08.020 . (Page 17)

[37] S. Chatrchyan, V. Khachatryan, A. Sirunyan, et al. Observation of a new boson at a mass of 125 gev with the CMS experiment at the LHC, *Physics Letters B*, **716**(1):30–61 (2012); doi:10.1016/j.physletb.2012.08.021 . (Page 17)

[38] B. Abi, T. Albahri, S. Al-Kilani, et al. Measurement of the positive muon anomalous magnetic moment to 0.46 ppm, *Phys. Rev. Lett.*, 126:141801 (2021); DO:10.1103/PhysRevLett.126.141801 . (Page 17)

[39] B. Pontecorvo. Neutrino experiments and the problem of conservation of leptonic charge, *Zh. Eksp. Teor. Fiz.*, 53:1717–1725 (1967); . (Page 18)

[40] J. N. Bahcall. Solar neutrinos. I. Theoretical, *Phys. Rev. Lett.*, 12:300–302 (1964); doi:10.1103/PhysRevLett.12.300 . (Page 18)

[41] R. Davis. Solar neutrinos. II. Experimental, *Phys. Rev. Lett.*, 12:303–305 (1964); doi:10.1103/PhysRevLett.12.303 . (Page 18)

[42] B. T. Cleveland, T. Daily, J. Raymond Davis, et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector, *The Astrophysical Journal*, **496**(1):505–526 (1998); doi:10.1086/305343 . (Page 18)

[43] J. N. Bahcall, M. H. Pinsonneault, and S. Basu. Solar models: current epoch and time dependences, neutrinos, and helioseismological properties, *The Astrophysical Journal*, **555**(2):990–1012 (2001); doi:10.1086/321493 . (Page 18)

[44] Y. Suzuki. Solar neutrino results from Super-Kamiokande, *Nuclear Physics B - Proceedings Supplements*, **91**(1):29–35 (2001); doi:10.1016/S0920-5632(00)00919-1 . (Page 19, 24)

[45] W. Hampel et al. GALLEX solar neutrino observations: results for GALLEX IV, *Phys. Lett. B*, 447:127–133 (1999); doi:10.1016/S0370-2693(98)01579-2 . (Page 19)

[46] M. Altmann, M. Balata, P. Belli, et al. GNO solar neutrino observations: results for GNO I, *Physics Letters B*, **490**(1):16–26 (2000); doi:10.1016/S0370-2693(00)00915-1 . (Page 19)

[47] V. Gavrin, J. Abdurashitovi, T. Bowles, et al. Solar neutrino results from SAGE, *Nuclear Physics B - Proceedings Supplements*, **91**(1):36–43 (2001); doi:10.1016/S0920-5632(00)00920-8 Neutrino 2000. (Page 19)

[48] Q. R. Ahmad, R. C. Allen, T. C. Andersen, et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory, *Phys. Rev. Lett.*, 87:071301 (2001); doi:10.1103/PhysRevLett.87.071301 . (Page 19)

[49] Q. R. Ahmad, R. C. Allen, T. C. Andersen, et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory, *Physical Review Letters*, **89**(1) (2002); doi:10.1103/physrevlett.89.011301 . (Page 19)

[50] K. Hirata, T. Kajita, M. Koshiba, et al. Experimental study of the atmospheric neutrino flux, *Physics letters. B*, **205**(2):416–420 (1988); doi:10.1016/0370-2693(88)91690-5 . (Page 20)

[51] Y. Fukuda et al. Atmospheric muon-neutrino / electron-neutrino ratio in the multiGeV energy range, *Physics Letters B*, 335:237–245 (1994); doi:10.1016/0370-2693(94)91420-6 . (Page 20)

[52] Y. Fukuda, T. Hayakawa, E. Ichihara, et al. Evidence for oscillation of atmospheric neutrinos, *Physics Review Letters*, 81:1562–1567 (1998); doi:10.1103/PhysRevLett.81.1562 . (Page 20)

[53] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles, *Progress of theoretical physics*, **28**(5):870–880 (1962); doi:10.1143/PTP.28.870 . (Page 20)

[54] S. Bilenky and B. Pontecorvo. Lepton mixing and neutrino oscillations, *Physics Reports*, **41**(4):225–261 (1978); doi:10.1016/0370-1573(78)90095-9 . (Page 20)

[55] P. deSalas, D. Forero, S. Gariazzo, et al. 2020 global reassessment of the neutrino oscillation picture, *J. High Energ. Phys. 2021*, 71 (2021); doi:10.1007/JHEP02(2021)071 . (Page 25)

[56] A. Porta, V. M. Bui, M. Cribier, et al. Reactor neutrino detection for non-proliferation with the NUCIFER experiment, *IEEE Transactions on Nuclear Science*, **57**(5 PART 2):2732–2739 (2010); doi:10.1109/TNS.2009.2035119 . (Page 26)

[57] W. N. Association. *Nuclear Power Reactors* (2016). Accessed 22 February 2022. (Page 27)

[58] G. Meyer and E. Stokke. *Description of Sizewell B Nuclear Power Plant* (1997). Institutt for energiteknikk (IFE). OECD Halden Reactor Project, Halden, Norway. (Page 28)

[59] R Mills (2022). Private communication. (Page 28)

[60] S. Dye and A. Barna. Global Antineutrino Modeling for a Web Application (2021). arXiv:1510.05633v3. (Page 29, 66, 67, 68, 69, 75, 78, 117)

[61] P. Huber. Determination of antineutrino spectra from nuclear reactors, *Physical Review C - Nuclear Physics*, **84**(2):1–16 (2011); doi:10.1103/PhysRevC.84.024617 . (Page 29)

[62] T. A. Mueller, D. Lhuillier, M. Fallot, et al. Improved predictions of reactor antineutrino spectra, *Physical Review C - Nuclear Physics*, **83**(5):1–17 (2011); doi:10.1103/PhysRevC.83.054615 . (Page 29)

[63] N. S. Bowden, A. Bernstein, S. Dazeley, et al. Observation of the isotopic evolution of pressurized water reactor fuel using an antineutrino detector, *Journal of Applied Physics*, **105**(6) (2009); doi:10.1063/1.3080251 . (Page 29, 30)

[64] E. Christensen, P. Huber, and P. Jaffke. *Antineutrino reactor safeguards - a case study* (2014). arXiv:1312.1959. (Page 29)

[65] G. Mention, M. Fechner, T. Lasserre, et al. Reactor antineutrino anomaly, *Phys. Rev. D*, 83:73006 (2011); doi:10.1103/PhysRevD.83.073006 . (Page 30)

[66] H. Almazán, L. Bernard, A. Blanchet, et al. Improved sterile neutrino constraints from the stereo experiment with 179 days of reactor-on data, *Phys. Rev. D*, 102:52002 (2020); doi:10.1103/PhysRevD.102.052002 . (Page 30)

[67] F. P. An, A. B. Balantekin, H. R. Band, et al. Evolution of the reactor antineutrino flux and spectrum at Daya Bay, *Physical Review Letters*, **118**(25):1–8 (2017); doi:10.1103/PhysRevLett.118.251801 . (Page 30)

[68] The RENO Collaboration. Fuel-composition dependent reactor antineutrino yield at RENO, *Physics Review Letters*, 122:8–13 (2019); doi:10.1103/PhysRevLett.122.232501 . (Page 30)

[69] J. M. Berryman and P. Huber. Re-evaluating reactor antineutrino anomalies with updated flux predictions, *Phys. Rev. D*, 101:15008 (2020); doi:10.1103/PhysRevD.101.015008 . (Page 30)

[70] N. Haag, A. Gütlein, M. Hofmann, et al. Experimental determination of the antineutrino spectrum of the fission products of ^{238}U, *Physical Review Letters*, **112**(12) (2014); doi:10.1103/physrevlett.112.122501 . (Page 30)

[71] V. I. Kopeikin, Y. N. Panin, and A. A. Sabelnikov. Measurement of the ratio of cumulative spectra of beta particles from ^{235}U and ^{239}Pu fission products for solving problems of reactor-antineutrino physics, *Phys. Atom. Nucl.*, **84**(1):1–10 (2021); doi:10.1134/S1063778821010129 . (Page 30)

[72] V. Kopeikin, M. Skorokhvatov, and O. Titov. *Re-evaluating reactor antineutrino spectra with new measurements of the ratio between ^{235}U and ^{239}Pu β spectra* (2021b). arXiv:2103.01684v2. (Page 30)

[73] P. Vogel and J. F. Beacom. Angular distribution of neutron inverse beta decay, *Physical Review D*, **60**(5):10 (1999); doi:10.1103/PhysRevD.60.053003 . (Page 31, 68)

[74] A. Strumia and F. Vissani. Precise quasielastic neutrino/nucleon cross-section, *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* (2003); doi:10.1016/S0370-2693(03)00616-6 . (Page 31, 67, 68)

[75] D. Hellfeld, A. Bernstein, S. Dazeley, and C. Marianno. Reconstructing the direction of reactor antineutrinos via electron scattering in Gd-doped water Cherenkov detectors, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 841:130–138 (2017); doi:10.1016/j.nima.2016.10.027 . (Page 32, 33, 38, 39)

[76] Y. Suzuki. The Super-Kamiokande experiment, *Eur. Phys. J. C*, 78 (2019); doi:10.1140/epjc/s10052-019-6796-2 . (Page 34, 68)

[77] A. Bernstein, T. West, and V. Gupta. An assessment of antineutrino detection as a tool for monitoring nuclear explosions, *Science & Global Security*, **9**(3):235–255 (2001); doi:10.1080/08929880108426496 . (Page 34)

[78] J. F. Beacom and M. R. Vagins. GADZOOKS! Antineutrino spectroscopy with large water Cerenkov detectors, *Physics Review Letters*, 93:1–4 (2003); doi:10.1103/PhysRevLett.93.171101 . (Page 34, 35, 111)

[79] L. Marti, M. Ikeda, Y. Kato, et al. Evaluation of Gadolinium's Action on water Cherenkov Detector Systems with EGADS, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 959:163549 (2020); doi:10.1016/j.nima.2020.163549 . (Page 35, 60, 61)

[80] K. Abe, C. Bronner, Y. Hayato, et al. First gadolinium loading to super-kamiokande, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, page 166248 (2021); doi:10.1016/j.nima.2021.166248 . (Page 35)

[81] A. R. Back, J. F. Beacom, M. Bergevin, et al. *Accelerator Neutrino Neutron Interaction Experiment (ANNIE): preliminary results and physics phase proposal* (2017). arXiv:1707.08222. (Page 35, 37, 58)

[82] V. Fischer, J. He, M. Irving, and R. Svoboda. Development of an ion exchange resin for gadolinium-loaded water, *Journal of instrumentation*, **15**(7):7004 (2020); doi:10.1088/1748-0221/15/07/P07004 . (Page 35)

[83] A. Gando, Y. Gando, K. Ichimura, et al. Constraints on θ_{13} from a three-flavor oscillation analysis of reactor antineutrinos at KamLAND, *Physical review. D, Particles, fields, gravitation, and cosmology*, **83**(5) (2011); doi:10.1103/PhysRevD.83.052002 . (Page 35)

[84] S. Andringa, E. Arushanova, S. Asahi, et al. Current status and future prospects of the sno+ experiment, *Advances in High Energy Physics*, 2016:1–21 (2016); doi:10.1155/2016/6194250 . (Page 36)

[85] J. Ashenfelter, A. Balantekin, C. Baldenegro, et al. The PROSPECT reactor antineutrino experiment, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 922:287–309 (2019); doi:10.1016/j.nima.2018.12.079 . (Page 36)

[86] M. Yeh, S. Hans, W. Beriguete, et al. A new water-based liquid scintillator and potential applications, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **660**(1):51–56 (2011); doi:10.1016/j.nima.2011.08.040 . (Page 37)

[87] D. R. Onken, F. Moretti, J. Caravaca, et al. Time response of water-based liquid scintillator from x-ray excitation, *Materials Advances*, **1**(1):71–76 (2020); doi:10.1039/d0ma00055h . (Page 37)

[88] T. Marrodán Undagoitia, F. von Feilitzsch, L. Oberauer, et al. Fluorescence decay-time constants in organic liquid scintillators, *Review of Scientific Instruments*, **80**(4):43301 (2009); doi:10.1063/1.3112609 . (Page 37)

[89] A. Balantekin, H. Band, G. Cao, et al. Measurement of electron antineutrino oscillation based on 1230 days of operation of the Daya Bay experiment, *Physical review. D*, **95**(7) (2017); doi:10.1103/PhysRevD.95.072006 . (Page 44)

[90] S. Abe, T. Ebihara, S. Enomoto, et al. Precision measurement of neutrino oscillation parameters with KamLAND, *Physical review letters*, **100**(22):221803 (2008); doi:10.1103/PhysRevLett.100.221803 . (Page 44)

[91] WATCHMAN Collaboration. *Conceptual Design Overview of the Advanced Instrumentation Testbed (AIT) and the WATer CHerenkov Monitor of ANtineutrinos (WATCHMAN)* (2019). Unpublished. (Page 44, 45, 54)

[92] A. Scarff. *Tank - design*. Technical report, University of Sheffield (2021). Unpublished. (Page 45)

[93] HAMAMATSU Photonics. Large photocathode area photomultiplier tubes. https://www.hamamatsu.com/resources/pdf/etd/LARGE_AREA_PMT_TPMH1376E.pdf. Accessed 6 September 2021. (Page 50, 52)

[94] A. Ezeribe, L. Kneale, M. Needham, et al. *Tests of first 101 AIT-WATCHMAN PMTs at Boulby*. Technical report, WATCHMAN Collaboration (2020). Unpublished. (Page 50)

[95] A. G. Wright. *The Photomultiplier Handbook*. Oxford University Press, 2017. (Page 51, 52, 53)

[96] J. Brack, B. Delgado, J. Dhooghe, et al. Characterization of the Hamamatsu R11780 12 inch photomultiplier tube, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 712:162 – 173 (2013); doi:10.1016/j.nima.2013.02.022 . (Page 52)

[97] Hamamatsu Photonics. Photomultiplier tubes, *Experimental Methods in the Physical Sciences* (2013); doi:10.1016/B978-0-12-387695-9.00003-2 . (Page 52)

[98] WATCHMAN Collaboration. *The AIT-NEO Project: Results from the Second Tank Size Trade Study* (2020). Unpublished. (Page 54)

[99] R. Claus, S. Seidel, L. Sulak, et al. A waveshifter light collector for a water Cherenkov detector, *Nuclear instruments & methods in physics research. Section A, Accelerators, spectrometers, detectors and associated equipment*, **261**(3):540–542 (1987); doi:10.1016/0168-9002(87)90366-4 . (Page 54)

[100] R. Levi-Setti, D. Park, and R. Winston. The corneal cones of limulus as optimised light concentrators, *Nature*, page 115–116 (1975); doi:10.1038/253115a0 . (Page 55)

[101] L. Berns. *A novel water-Cherenkov detector design with retro-reflectors to produce antipodal rings* (2018). arXiv:1808.09623v1. (Page 55)

[102] B. W. Adams, K. Attenkofer, M. Bogdan, et al. *A brief technical history of the Large-Area Picosecond Photodetector (LAPPD) Collaboration* (2016). arXiv:1603.01843. (Page 55)

[103] A. Lyashenko, B. Adams, M. Aviles, et al. Performance of Large Area Picosecond Photo-Detectors (LAPPD™), *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 958:162834 (2020); doi:10.1016/j.nima.2019.162834 . (Page 56)

[104] M. Bogdan, E. Oberla, H. J. Frisch, and M. Wetstein. *A modular data acquisition system using the 10 GSa/s PSEC4 waveform recording chip* : 2016 IEEE-NPSS Real Time Conference (RT), pages 1–3, 2016; doi:10.1109/RTC.2016.7543167. (Page 56)

[105] H. Löhner, Q. Dorosti-Hasankiadeh, E. Heine, et al. The multi-PMT optical module for KM3NeT, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 718:513–515 (2013); doi:10.1016/j.nima.2012.11.049 . (Page 56)

[106] T. Kaptanoglu, M. Luo, B. Land, et al. Spectral photon sorting for large-scale cherenkov and scintillation detectors, *Phys. Rev. D*, 101:072002 (2020); doi:10.1103/PhysRevD.101.072002 . (Page 56, 57)

[107] S. Funk, D. Jankowsky, H. Katagiri, et al. *TARGET: A digitizing and trigger ASIC for the Cherenkov telescope array* : AIP conference proceedings, volume 1792, 2017; doi:10.1063/1.4969033. (Page 58)

[108] O. A. Akindele, T. Anderson, M. Askins, et al. *A Snowmass 2021 Letter of Interest for Encapsulation of Photosensors in kton–Mton Scale Neutrino Detectors* (2021). (Page 59)

[109] R. J. Ford. *A scintillator purification plant and fluid handling system for SNO+* : AIP Conference Proceedings, volume 1672, United States, 2015; doi:10.1063/1.4927998. (Page 61)

[110] J. Scherzinger, J. Annand, G. Davatz, et al. Tagging fast neutrons from an 241am/9be source, *Applied Radiation and Isotopes*, 98:74–79 (2015); doi:10.1016/j.apradiso.2015.01.003 . (Page 63)

[111] M. Dragowsky, A. Hamer, Y. Chan, et al. The ^{16}N calibration source for the Sudbury Neutrino Observatory, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **481**(1):284–296 (2002); doi:10.1016/S0168-9002(01)02062-9 . (Page 63)

[112] E. Kneale. *Radionuclide Rates.* Technical report, University of Sheffield & AWE (2021). Unpublished. (Page 65)

[113] S. Seibert. *RAT-PAC (Reactor Analysis Tool, Plus Additional Codes)* (2014). https://rat.readthedocs.io/en/latest/index.html. Accessed 1 August 2021. (Page 65)

[114] RAT-PAC v0.1.0 (Reactor Analysis Tool Plus Additional Codes) for AIT-WATCHMAN. https://github.com/AIT-WATCHMAN/rat-pac. Last accessed October 2021. Unpublished. (Page 65)

[115] J. Allison, K. Amako, J. Apostolakis, et al. Recent developments in Geant4, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225 (2016); doi:10.1016/j.nima.2016.06.125 . (Page 65)

[116] S. Agostinelli, J. Allison, K. Amako, et al. GEANT4 — a simulation toolkit, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **506**(3):250–303 (2003); doi:10.1016/S0168-9002(03)01368-8 . (Page 65)

[117] L. Lönnblad. CLHEP—a project for designing a C++ class library for high energy physics, *Computer Physics Communications*, **84**(1):307–316 (1994); doi:10.1016/0010-4655(94)90217-8 . (Page 65)

[118] G. Horton-Smith. *GLG4sim* (2005). www.phys.ksu.edu/personal/gahs/GLG4sim/. Accessed 1 August 2021. (Page 65)

[119] R. Brun and F. Rademakers. ROOT — an object oriented data analysis framework v6.18/02, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment.*, **389**(1):81–86 (1997); doi:10.5281/zenodo.3895860 . (Page 65)

[120] D.-M. Mei and A. Hime. Muon-induced background study for underground laboratories, *Phys. Rev. D*, 73:053004 (2006); doi:10.1103/PhysRevD.73.053004 . (Page 65, 71, 72, 73, 74, 78)

[121] EDF Energy. Great year for Hartlepool. https://www.edfenergy.com/media-centre/news-releases/great-year-hartlepool. Accessed 8 July 2021. (Page 66)

[122] Department for Business, Energy and Industrial Strategy. *Energy Trends: UK, October to December 2020 and 2020* (2021). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/976000/Energy_Trends_March_2021.pdf. Accessed 8 July 2021. (Page 66)

[123] R Mills. *Anti-neutrino source term from the EDF Hartlepool reactors* (2021). Unpublished. (Page 66)

[124] International Atomic Energy Agency. General design and principles of the advanced gas-cooled reactor (AGR). https://nucleus-new.iaea.org/sites/graphiteknowledgebase/wiki/Guide_to_Graphite/Home.aspx. Accessed 8 July 2021. (Page 66)

[125] O. Akindele. *Exclusion and verification of remote reactors* (2020). Unpublished. (Page 70)

[126] S. Dazeley, M. Askins, M. Bergevin, et al. A search for cosmogenic production of β-neutron emitting radionuclides in water, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 821:151–159 (2016); doi:10.1016/j.nima.2016.03.014 . (Page 71)

[127] E. Tziaferi, M. Carson, V. Kudryavtsev, et al. First measurement of low intensity fast neutron background from rock at the Boulby Underground Laboratory, *Astroparticle Physics*, **27**(5):326–338 (2007); doi:10.1016/j.astropartphys.2006.12.005 . (Page 71)

[128] Y.-F. Wang, V. Balic, G. Gratta, et al. Predicting neutron production from cosmic-ray muons, *Physical review. D, Particles and fields*, **64**(1) (2001); doi:10.1103/PhysRevD.64.013012 . (Page 71, 72)

[129] A. Tang, G. Horton-Smith, V. A. Kudryavtsev, and A. Tonazzo. Muon simulations for Super-Kamiokande, KamLAND, and CHOOZ, *Phys. Rev. D*, 74:053007 (2006); doi:10.1103/PhysRevD.74.053007 . (Page 71)

[130] S. W. Li and J. F. Beacom. First calculation of cosmic-ray muon spallation backgrounds for MeV astrophysical neutrino signals in Super-Kamiokande, *Phys. Rev. C*, 89:045801 (2014); doi:10.1103/PhysRevC.89.045801 . (Page 73, 74, 75, 109)

[131] S. Chu, L. Ekström, and R. Firestone. The Lund/LBNL nuclear data search: Table of radioactive isotopes. http://nucleardata.nuclear.lu.se/toi/. Accessed August 2021. (Page 74)

[132] TUNL Nuclear Data Evaluation Project. Energy level diagrams. https://nucldata.tunl.duke.edu/nucldata. Accessed August 2021. (Page 74)

[133] C. Jollet and A. Meregaglia. ^{9}Li and ^{8}He decays in GEANT4, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 949:162904 (2020); doi:10.1016/j.nima.2019.162904 . (Page 74)

[134] M. Robinson, V. Kudryavtsev, R. Lüscher, et al. Measurements of muon flux at 1070m vertical depth in the Boulby underground laboratory, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **511**(3):347–353 (2003); doi:10.1016/S0168-9002(03)01973-9 . (Page 73, 75)

[135] Y. Zhang, K. Abe, Y. Haga, et al. First measurement of radioactive isotope production through cosmic-ray muon spallation in Super-Kamiokande IV, *Phys. Rev. D*, 93:012004 (2016); doi:10.1103/PhysRevD.93.012004 . (Page 74)

[136] E. Kneale, T. Brooks, and A. Healey. *CLEANWATCH: radioactivity budgeting software for AIT-WATCHMAN detector design*. Technical report, University of Sheffield (2020). Unpublished. (Page 77)

[137] M. Smy. *Low energy event reconstruction and selection in Super-Kamiokande-III* : 30th International Cosmic Ray Conference, (2007). (Page 79)

[138] Fred (Functions to Reconstruct Events in the Detector). https://github.com/AIT-WATCHMAN/fred. Last accessed Sept 2021. Unpublished. (Page 79)

[139] L. Kneale. Core v0.1.0 (Combined Reconstruction for WATCHMAN). https://github.com/ekneale/core. Last accessed 18 Oct 2021. (Page 79)

[140] K. Abe, Y. Hayato, T. Iida, et al. Solar neutrino results in Super-Kamiokande-III, *Physical Review D*, **83**(5) (2011); doi:10.1103/physrevd.83.052010 . (Page 84)

[141] A. Bellerive, J. Klein, A. McDonald, et al. The Sudbury Neutrino Observatory, *Nuclear Physics B*, 908:30–51 (2016); doi:10.1016/j.nuclphysb.2016.04.035 Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015. (Page 84)

[142] L. Kneale. Cobraa v0.1.0 (Coincident-Background Reactor Antineutrino Analysis). https://github.com/ekneale/cobraa. Last accessed 18 Oct 2021. (Page 92)

[143] M. Bergevin. Watchmakers v0.1.0 (Water Cherenkov Monte carlo and Analysis Kit for Reactor Sensitivity). https://github.com/AIT-WATCHMAN/watchmakers. Last accessed January 2021. Unpublished. (Page 92)

[144] F. Sutanto, O. Akindele, M. Bergevin, et al. *Progress report: Study on the effect of varying veto thickness on the sensitivity of* Gd – H_2O *filled tank*. Technical report, Lawrence Livermore National Laboratory (2021). Unpublished. (Page 96)

[145] G. F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, Inc., 2000. (Page 101, 123)

[146] N. Holland. Using a boosted decision tree to improve the sensitivity of the watchman detector for remote reactor monitoring. Master's thesis, University of Sheffield, 2021. (Page 109)

[147] S. Farrell. *Machine learning approaches to particle classification in WATCHMAN* (2021). (Page 109)

[148] S. W. Li and J. F. Beacom. Spallation backgrounds in super-kamiokande are made in muon-induced showers, *Physical review. D, Particles, fields, gravitation, and cosmology*, **91**(10) (2015); . (Page 109)

[149] C. Bauer, E. Borger, R. Hofacker, et al. Qualification tests of 474 photomultiplier tubes for the inner detector of the Double Chooz experiment, *Journal of Instrumentation*, **6**(06):P06008–P06008 (2011); doi:10.1088/1748-0221/6/06/p06008 . (Page 113)

[150] D. Milind. *PMT Working Group 2020 Q2Report*. Technical report, Brookhaven National Laboratory (2020). Unpublished. (Page 115)

[151] J. Gribble. *Legendre polynomials and anisotropy of events* (2019). Internal communication. (Page 118)

[152] G. Cowan. *Discovery sensitivity for a counting experiment with background uncertainty*. Technical report, Royal Holloway, University of London (2012). (Page 122)

[153] G. Vianello. The significance of an excess in a counting experiment: Assessing the impact of systematic uncertainties and the case with a gaussian background, *The Astrophysical Journal Supplement Series*, **236**(1):17 (2018); doi:10.3847/1538-4365/aab780 . (Page 123)