

# Novel technologies for future applications of nuclear reactor monitoring using antineutrinos

Robert J. Foster



Department of Physics & Astronomy

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

April 2024

---

## Abstract

Nuclear reactors produce copious number of antineutrinos as a result of the nuclear fission process. This antineutrino emission can be used to monitor the nuclear reactor, providing important information regarding its operating state, power level, and fissile content in the core. Efforts are underway to construct a large-scale water-based antineutrino detector that is able to detect and monitor a nuclear reactor at distance-scales of tens of kilometres. The baseline conceptual design is a kilotonne-scale water-Cherenkov detector doped with gadolinium that detects the “double flash” of a positron and a neutron resulting from the inverse beta decay process. However, novel detection technologies are being actively studied to improve performance or extract more information from antineutrino interactions.

One such technology is water-based liquid scintillator, a highly scalable detection medium with improved light yield compared to pure water. Detector simulations utilising water and water-based liquid scintillator detector fills suggests that such detectors would be sensitive to a nuclear reactor complex at a distance of around 150 km. Water-based liquid scintillator performs better in general, although despite the improved light yield from the addition of scintillator, gadolinium is still required to boost the neutron capture signal and suppress backgrounds.

Another key technology is the Large Area Picosecond PhotoDetector (LAPPD), a microchannel plate photosensor that offers millimetre-scale positional resolution and timing resolution of tens of picoseconds. The LAPPD performance is characterised and verified using a custom-designed laser test stand. In order to perform vertex reconstruction using LAPPDs, an algorithm to disambiguate photon hits on the LAPPD is developed allowing for the isolation of individual photon hits from hit clusters and the determination of the position and timing of these individual photons.

To achieve the goal of a large-scale demonstration of an antineutrino detector for nuclear security, smaller scale technology testbeds are being developed. One such testbed is the BUTTON detector at Boulby, which aims to deploy both water-based liquid scintillator and LAPPDs. The LAPPD characterisation and multiphoton disambiguation work detailed in this thesis will contribute directly to the deployment and operation of LAPPDs in BUTTON.

---

## Declaration

I, the author, confirm that this thesis is my own work, except as declared here or where references to other work have been made. This work has not previously been presented for an award at this or any other university.

The work performed in Chapter 4 is one small part of a much larger study taken on by the WATCHMAN collaboration as a whole, and performed by the Trade Study Team. I took ownership of the distant reactor study from Tom Brooks and performed the analysis, but the analysis used many tools developed by others in the collaboration, specifically the WATCHMAKERS software package developed by Marc Bergevin. Monte Carlo data was generated by Marc Bergevin and Morgan Askins. The parallel distant reactor analysis to which my results were compared was conducted by Gabriel Orebi Gann, Morgan Askins, Zara Bagdasarian, Stephane Zsoldos, and others.

Chapter 6 details characterisation of the LAPPD. This was performed by myself, with assistance and advice from Andrew Scarff. Stephen Wilson created the LabVIEW high voltage control system and assisted with setting up the gantry system. I directed the research and measurements, wrote the entirety of the analysis code, and wrote the majority of the testbench automation code.

Chapter 7 describes an algorithm for disambiguating photon hits on the LAPPD, which originated from a paper by Jocher et al. The detector response simulations, algorithm implementation and subsequent analysis and validation was performed entirely by myself. The integration of the LAPPDs into the BONSAI reconstruction was performed with the assistance of Liz Kneale.

A significant proportion of this work was carried out during the global coronavirus pandemic, which severely impacted access to University facilities, and delayed or cancelled several potential avenues of research. The entirety of the work in Chapter 4 was performed under work from home conditions.

*Robert J. Foster*

This research was conducted under the supervision of Dr. Matthew Malek and Dr. Patrick Stowell at the University of Sheffield, and supported by the Science and Technology Facilities Council.

---

## Acknowledgements

I never could have imagined when I began a PhD in particle physics that I would spend a whole year working in my bedroom and then the best part of three years working in a shed but, then again, I don't think anything in this PhD went quite as I expected. I've learned a hell of a lot in this time, mostly as a result of the fantastic people I have had the pleasure of working with.

Firstly, I would like to acknowledge my supervisor Matthew Malek for giving me the opportunity to begin this PhD and for his guidance. Thanks also to Patrick Stowell for stepping in to give guidance, advice, and feedback at the most crucial time in my PhD. I would like to express my gratitude to all the folks on the ANNIE experiment and the Fermilab community more widely for making my six-month stay so enjoyable (even if it was winter in the Midwest). I wish all the best for ANNIE, I know you can and will achieve great things. Additionally, I would like to thank all the folks at Incom for being so open, friendly, and willing to help with our many questions about the LAPPD and for their interest in the research we were performing.

Doing this PhD has only been bearable thanks to the amazing people I have worked with at Sheffield. The chances that this thesis would exist without the help of Sam Fargher and Andy Scarff are pretty much zero. Like so many others at Sheffield, my hardware skills are a direct result of Sam's crash course teaching style and FAFO attitude to research in general. Andy was a source of constant motivation and reassurance that my work was actually good enough for a PhD. It turns out you were right in the end. I think it's fair to say that you both went far above and beyond what was required; that is truly appreciated. Thank you to James, Steve, Jack, Jordan, and everyone else around the office and the lab for the constant laughs, the conversations I can never repeat, and for being a source of endless encouragement. It's a good job I only joined the chaos of the office a few months before the end, this thesis would never have been written otherwise.

Finally, I would like to say a massive thank you to all my friends and family. Without your advice and your belief in me, none of this would ever have been possible.





# Contents

<b>1</b>	<b>The neutrino</b>	<b>1</b>
1.1	The hidden particle . . . . .	2
1.2	First detection . . . . .	3
1.3	Neutrino flavours . . . . .	8
1.4	Neutrino oscillation . . . . .	12
1.5	Chapter summary . . . . .	15
1.6	Thesis outline . . . . .	15
<b>2</b>	<b>Reactor neutrinos</b>	<b>18</b>
2.1	Nuclear power . . . . .	18
2.2	Reactor safeguarding . . . . .	19
2.3	Reactor antineutrinos . . . . .	22
2.3.1	Antineutrino flux . . . . .	22
2.3.2	Antineutrino energy . . . . .	25
2.3.3	Antineutrino interactions . . . . .	26
2.4	Reactor antineutrino anomaly . . . . .	28
2.5	Demonstrations of antineutrino monitoring . . . . .	29
2.6	Technologies for antineutrino detection . . . . .	32
2.6.1	Water Cherenkov detectors . . . . .	32
2.6.2	Liquid scintillator detectors . . . . .	34
2.6.3	Gadolinium doping . . . . .	40
2.6.4	Water-based liquid scintillator . . . . .	41
2.7	Chapter summary . . . . .	45

<b>3</b>	<b>Reactor Antineutrino Detector Prototype</b>	<b>47</b>
3.1	The Advanced Instrumentation Testbed . . . . .	48
3.2	Experimental objectives . . . . .	49
3.3	The NEO detector . . . . .	50
3.3.1	Tank . . . . .	51
3.3.2	PMT support structure . . . . .	52
3.3.3	PMTs . . . . .	53
3.3.4	Calibration . . . . .	57
3.3.5	Data acquisition . . . . .	60
3.3.6	Advanced photodetectors . . . . .	61
3.3.7	Water treatment system . . . . .	62
3.4	Chapter summary . . . . .	64
<b>4</b>	<b>Distant reactor detection</b>	<b>66</b>
4.1	Heysham nuclear reactor complex . . . . .	67
4.2	Detector configurations . . . . .	68
4.3	Simulation framework . . . . .	69
4.3.1	IBD signal and backgrounds . . . . .	71
4.3.2	Radiological and cosmogenic backgrounds . . . . .	73
4.4	Analysis methodology . . . . .	75
4.4.1	Framework . . . . .	75
4.4.2	Event reconstruction . . . . .	75
4.4.3	Energy estimation . . . . .	77
4.4.4	Region of interest optimisation . . . . .	78
4.5	Parallel analysis . . . . .	81
4.6	Results . . . . .	81
4.7	Later work . . . . .	85
4.8	Conclusions . . . . .	86
4.8.1	BUTTON . . . . .	87
4.9	Chapter summary . . . . .	88
<b>5</b>	<b>Large Area Picosecond PhotoDetectors</b>	<b>90</b>
5.1	Photomultiplier tubes . . . . .	91

5.2	Microchannel plates . . . . .	93
5.3	Traditional fabrication process . . . . .	96
5.4	History of the LAPPD project . . . . .	98
5.5	The LAPPD package . . . . .	100
5.5.1	Ultem frame . . . . .	100
5.5.2	Glass window and photocathode . . . . .	101
5.5.3	MCP stack . . . . .	102
5.5.4	Anode . . . . .	105
5.6	LAPPD operation . . . . .	109
5.6.1	High voltage . . . . .	109
5.6.2	Fast digitisation . . . . .	110
5.7	LAPPD time resolution . . . . .	111
5.7.1	Variability in electron cascade timing . . . . .	112
5.7.2	Signal degradation . . . . .	114
5.8	Water Cherenkov LAPPD deployment . . . . .	115
5.8.1	Waterproof housing . . . . .	115
5.8.2	Signal digitisation . . . . .	116
5.8.3	Power . . . . .	120
5.8.4	Monitoring and slow control . . . . .	122
5.9	Chapter summary . . . . .	123
<b>6</b>	<b>LAPPD characterisation</b>	<b>124</b>
6.1	Test stand setup . . . . .	125
6.1.1	Dark box . . . . .	125
6.1.2	Laser gantry system . . . . .	126
6.1.3	Electronics . . . . .	129
6.2	Single pulse signal processing . . . . .	132
6.2.1	Waveform preprocessing . . . . .	133
6.2.2	Peak finding . . . . .	133
6.2.3	Timestamping . . . . .	134
6.2.4	Pair matching . . . . .	137
6.3	Longitudinal position reconstruction . . . . .	138
6.3.1	Method . . . . .	138

6.3.2	Stripline propagation velocity . . . . .	139
6.4	Transverse position reconstruction . . . . .	141
6.5	Characterisation measurements . . . . .	147
6.5.1	Pulse shape . . . . .	147
6.5.2	Dark rate . . . . .	147
6.5.3	Gain . . . . .	151
6.5.4	Transit time spread . . . . .	158
6.5.5	Spatial characterisation . . . . .	164
6.6	Chapter summary . . . . .	168
<b>7</b>	<b>Multiphoton disambiguation with LAPPDs</b>	<b>170</b>
7.1	Motivation . . . . .	170
7.2	LAPPD waveform simulator . . . . .	171
7.3	Deconvolutional method . . . . .	174
7.3.1	Wiener deconvolution . . . . .	174
7.3.2	Peak finding . . . . .	177
7.3.3	Pair matching . . . . .	177
7.4	Tests on simulated data . . . . .	181
7.5	Vertex reconstruction using LAPPDs . . . . .	185
7.6	Alternative approaches towards multiphoton disambiguation . . . . .	190
7.6.1	Formula fitting . . . . .	191
7.6.2	Machine learning . . . . .	191
7.7	Chapter summary . . . . .	192
<b>8</b>	<b>Conclusions and future work</b>	<b>194</b>
	<b>References</b>	<b>198</b>
<b>A</b>	<b>LAPPD 104</b>	<b>217</b>
<b>B</b>	<b>Single photoelectron operation</b>	<b>219</b>

# List of Figures

1.1	Original experimental setup of “Project Poltergeist” . . . . .	5
1.2	Savannah River experimental setup . . . . .	7
1.3	AGS experimental setup . . . . .	9
1.4	Hadron production cross-section at $Z^0$ boson resonance . . . . .	10
2.1	Fission fragment yields . . . . .	23
2.2	Emitted antineutrino spectrum . . . . .	24
2.3	Detection cross-section for antineutrinos . . . . .	25
2.4	Inverse beta decay Feynman diagram . . . . .	26
2.5	Re-evaluated antineutrino energy spectrum . . . . .	29
2.6	Antineutrino flux over reactor cycle at Rovno . . . . .	30
2.7	Antineutrino rate during reactor refuelling . . . . .	31
2.8	Cherenkov effect . . . . .	33
2.9	Singlet and triplet vibrational states of fluorescent molecule . . . . .	36
2.10	Emission and absorption wavelengths of scintillator components . . . . .	38
2.11	A micelle encapsulating organic scintillator . . . . .	41
2.12	Cherenkov-scintillation separation by wavelength . . . . .	43
2.13	Cherenkov and LAB/PPO emission spectra . . . . .	45
3.1	Conceptual NEO tank design . . . . .	51
3.2	Conceptual design PMT support structure top endcap . . . . .	53
3.3	$^{16}\text{N}$ decay scheme . . . . .	60
3.4	EGADS molecular bandpass filtration system . . . . .	63
3.5	Resin ion-exchange filtration system . . . . .	64

4.1	Map showing reactor locations . . . . .	67
4.2	Heysham signal and background spectrum . . . . .	72
4.3	NEO PMT hit time residuals . . . . .	76
4.4	Efficiency maps . . . . .	80
4.5	Heysham signal sensitivity as a function of fiducial volume . . . . .	82
4.6	Heysham signal sensitivity as a function of tank diameter . . . . .	83
4.7	Heysham signal sensitivity as a function of tank diameter . . . . .	84
4.8	Schematic of the BUTTON detector . . . . .	88
5.1	Schematic of a photomultiplier tube . . . . .	92
5.2	Schematic of a microchannel plate . . . . .	94
5.3	Exploded view of an LAPPD . . . . .	100
5.4	Fused silica window transparency . . . . .	101
5.5	Image of a glass capillary array . . . . .	102
5.6	Cutaway drawing of an LAPPD . . . . .	103
5.7	Diagram of MCP coated using ALD . . . . .	104
5.8	The LAPPD stripline readout . . . . .	107
5.9	Crosstalk on microstriplines . . . . .	108
5.10	Schematic of the mechanism for early, normal, and late pulses . . . . .	113
5.11	Impact of voltage jitter on time resolution . . . . .	115
5.12	Operation of DRS4 waveform digitiser . . . . .	117
5.13	PSEC4 digitising ASIC . . . . .	120
6.1	LAPPD characterisation test stand schematic . . . . .	126
6.2	Photograph of the LAPPD test stand . . . . .	127
6.3	Timing of the laser trigger and observed LAPPD pulse in acquisition window . . . . .	130
6.4	LAPPD high voltage configuration during characterisation . . . . .	132
6.5	Example of time slew . . . . .	135
6.6	Log-normal fit to an LAPPD waveform . . . . .	136
6.7	LAPPD differential timing versus laser position . . . . .	140
6.8	Reconstructed position as a function of laser position . . . . .	140
6.9	Strip line charge distribution . . . . .	142
6.10	Transverse position resolution . . . . .	143

6.11	Pulse height vs pulse integral linearity . . . . .	144
6.12	Fitted charge vs pulse integral . . . . .	145
6.13	Rise time distribution . . . . .	146
6.14	Rise time . . . . .	146
6.15	LAPPD dark rate as a function of photocathode voltage . . . . .	150
6.16	Strip-to-strip dark rate variation . . . . .	151
6.17	LAPPD gain distribution at varying operating voltages . . . . .	154
6.18	LAPPD pulse height distribution at varying operating voltages . . . . .	155
6.19	The LAPPD mean gain at (a) varying photocathode voltages and at (b) varying MCP voltages . . . . .	156
6.20	LAPPD gain as a function of laser repetition rate . . . . .	157
6.21	Transit time distributions at different photocathode voltages . . . . .	160
6.22	Transit time spread values as a function of photocathode voltage . . . . .	161
6.23	LAPPD afterpulse example . . . . .	163
6.24	LAPPD afterpulse ratio . . . . .	164
6.25	Relative arrival time vs amplitude of afterpulses . . . . .	165
6.26	Heat map of gain across photocathode . . . . .	166
6.27	Heat map of transit time spread across photocathode . . . . .	167
6.28	Transit time distributions across a stripline . . . . .	168
7.1	Simulated LAPPD response for single photon . . . . .	173
7.2	Simulated LAPPD response for twenty photons . . . . .	174
7.3	Point spread function . . . . .	175
7.4	Example of deconvolution on single stripline . . . . .	176
7.5	Time-of-arrival pulse-matching PDF . . . . .	179
7.6	Transverse position pulse-matching PDF . . . . .	180
7.7	Amplitude pulse-matching PDF . . . . .	181
7.8	Identified photon hits with 100 photons incident on the LAPPD . . . . .	184
7.9	Schematic of the BUTTON detector with LAPPDs . . . . .	186
7.10	LAPPD hit time residuals . . . . .	187
7.11	PMT hit time residuals . . . . .	188
7.12	Vertex reconstruction with PMTs or LAPPDs . . . . .	189
A.1	LAPPD 104 . . . . .	217



B.1 Probabilities for producing different numbers of photoelectrons . . . . 220

# List of Tables

4.1	Number of PMTs for each detector configuration . . . . .	69
4.2	Signal rates for Heysham reactor cores . . . . .	71
4.3	Simulated backgrounds in NEO detector . . . . .	73
5.1	Typical LAPPD operational voltages . . . . .	109
5.2	Comparison of LAPPD digitisation options . . . . .	121
6.1	A selection of performance values from Incom’s test and measurement report for LAPPD 104. . . . .	125
6.2	Comparison of transit time spread measurements using different timestamping methods at 850 V applied to the MCPs and 100 V applied to the photocathode. . . . .	162
6.3	A selection of performance values from the characterisation measurements of LAPPD 104. . . . .	169
7.1	Results of the deconvolutional method applied to simulated data . . .	183

# List of Acronyms

<b>ADC</b>	analogue to digital converter
<b>AIT</b>	Advanced Instrumentation Testbed
<b>AGR</b>	advanced gas-cooled reactor
<b>ALD</b>	atomic layer deposition
<b>AmBe</b>	Americium-Beryllium
<b>ANNIE</b>	Accelerator Neutrino Neutron Interaction Experiment
<b>ANL</b>	Argonne National Laboratory
<b>ASIC</b>	application-specific integrated circuit
<b>BUTTON</b>	Boulby Underground Technology Testbed Observing Neutrinos
<b>CE<math>\nu</math>NS</b>	coherent elastic neutrino nucleus scattering
<b>DAQ</b>	data acquisition
<b>DAC</b>	digital to analogue converter
<b>DRS4</b>	Domino Ring Sampler 4
<b>DT</b>	deuterium-tritium
<b>FPGA</b>	field-programmable gate array
<b>FRET</b>	Förster resonant energy transfer
<b>IAEA</b>	International Atomic Energy Agency
<b>IBD</b>	inverse beta decay
<b>IC</b>	integrated circuit
<b>LAPPD</b>	Large Area Picosecond PhotoDetector™
<b>MCP</b>	microchannel plate
<b>NEO</b>	Neutrino Experiment One
<b>PDF</b>	probability distribution function
<b>PLL</b>	phase-locked loop

<b>PMT</b>	photomultiplier tube
<b>POPOP</b>	1,4-bis(5-phenyloxazol-2-yl) benzene
<b>PPO</b>	2,5-diphenyl-oxazole
<b>PSF</b>	point spread function
<b>RAT-PAC</b>	Reactor Analysis Tool (Plus Additional Codes)
<b>RF</b>	radio frequency
<b>SHV</b>	safe high voltage
<b>WATCHMAN</b>	WATER CHerenkov Monitor for AntiNeutrinos
<b>WbLS</b>	water-based liquid scintillator

# Chapter 1

## The neutrino

“Before I came here I was confused about this subject. Having [written this thesis] I am still confused — but on a higher level.”

---

Enrico Fermi

The most concise description of the nature of the neutrino exists in its own name. It means “little neutral one” and, perhaps not surprisingly, the neutrino is a very light particle (thought for a long time to be massless) with no electrical charge. The neutrino is a lepton which, when combined with its lack of electrical charge, means that it only ever meaningfully interacts through the weak force. Despite being one of the most abundant particles in the universe, produced in copious numbers during astrophysical processes, the neutrino is incredibly difficult to detect due to its elusive, “ghostly” nature. Whilst a particle that does pretty much nothing but happily stream through the universe may seem boring, neutrinos are the key to understanding many of the most burning questions in physics including the nature of mass and the reason for the matter-antimatter asymmetry in the universe. It is for this reason that the international research effort into neutrino physics is massive. Bigger and more expensive experiments are continually being proposed and constructed to understand the fundamental

properties of the neutrino. As our understanding of the neutrino improves, so too does the potential of using neutrinos for practical applications. This thesis focuses on one such application — using neutrinos as a tool to detect and monitor nuclear reactors. This chapter provides a brief history of the many pioneering experiments that led to our current understanding of the neutrino.

## 1.1 The hidden particle

By 1930, tension had been building in the studies of radioactivity over the past two decades. Beta decay was thought to be a two body decay like alpha and gamma decay, in which one observes a very narrow band of outgoing particle energies resulting from the mass difference between the initial and final states. However, since 1911 measurements of the electron kinetic energy from beta decays had hinted at a continuous energy spectrum. This was later confirmed conclusively by Chadwick [1] and others [2–4] over the next couple of decades. Adding to the confusion, measurements of the nuclear spin of nitrogen-14 suggested that it, and other nuclei with even mass numbers, had integral spin, but the electron spin was  $1/2$ , implying that angular momentum was not conserved during beta decay. Niels Bohr at the time even hypothesised that maybe energy was only conserved in the statistical sense, rather than per decay [5].

In what he dubbed as a “desperate remedy” to this state of confusion in the field, Wolfgang Pauli postulated the existence of an as-yet-unseen particle in a 1930 letter to a meeting of physicists in Tübingen [6]. This hidden particle must have been contained within the nucleus, and would then be expelled in the process of beta decay. In order to explain the existing observations of beta decay, and the fact that it had not yet been observed, this particle should be electrically neutral. It should also have  $1/2$  spin and thus obey the exclusion principle, and be similar in mass to the electron. He had dubbed it the “neutron” and lamented that his proposed particle could never be observed because of the properties that must be required of a particle that had evaded detection for so long.

It took only three years after Pauli’s letter for a major breakthrough. Enrico Fermi, one of the most accomplished physicists of the 20<sup>th</sup> century, broke ground on

a mathematical description of the theory of beta decay [7]. He incorporated Pauli’s “neutron” into his theory but renamed it the neutrino to distinguish it from the particle Chadwick had observed [8] just one year earlier. Fermi’s theory described beta decay as a four fermion interaction, in which a neutron would become a proton and in doing so would produce an electron and a neutrino as so,

$$n \longrightarrow p + e^{-} + \bar{\nu}_e. \tag{1.1}$$

Fermi’s theory of beta decay proposed a major step away from conventional understanding at the time: that the electron and neutrino were not contained within the atom and expelled upon the decay but rather were created during the interaction. With this there was finally a proposal for the theory of weak interactions, although it was, now famously, rejected from publication in *Nature* for being too speculative. At a similar time, evidence was building that there was a concrete endpoint to the beta decay energy spectrum [9], solidifying the principle of energy conservation and the idea of the existence of the neutrino amongst the physics community. Efforts began immediately to search for experimental evidence of this new particle [10, 11], but it would take another two decades for the existence of the neutrino to be confirmed by observation.

## 1.2 First detection

While Fermi’s theory proved successful in explaining the mechanics of beta decay, it also made clear just how difficult a direct detection of the neutrino could be. After Fermi’s work was published it was shown that, if Fermi was correct in just *how* weak the weak interaction was, the neutrino should be able to easily pass through the entire Earth unobstructed and unnoticeable. A letter to *Nature* by Bethe and Peierls investigated the cross-section for neutrino interactions and concluded that “...there is no practically possible way of observing the neutrino” [12]. Despite the obvious challenge, efforts to make the first ever detection of the neutrino were undeterred.

Frederick Reines met Clyde Cowan while they were both working at Los Alamos National Laboratory in 1951. Up until then, evidence for the existence of

the neutrino had been indirect, through attempted recoil measurements using beta decay. A direct measurement of the free neutrino observed at some distance away from where it was produced was required to be conclusive. Cowan and Reines began “Project Poltergeist” [13], an attempt to perform a direct detection of the neutrino, partly because, as Reines later recalled, “...everyone said you couldn’t do it” [14]. Due to the tiny likelihood of a neutrino interacting, they knew they would need an intense source of neutrinos. Reines, who had been involved in the Manhattan Project during the Second World War, knew that the atomic bomb should produce massive numbers of neutrinos — specifically electron antineutrinos — from the beta decay of the many fission fragments that would be produced in the detonation. In addition, the neutrinos would be produced in a quick burst, and so the neutrino flux would dominate over any backgrounds as the bomb exploded.

Even if a reasonable neutrino flux could be achieved, a choice had to be made regarding how exactly they would detect the neutrino. Cowan and Reines knew from Fermi’s work that there should exist a reaction that was the inverse of beta decay, now aptly named inverse beta decay (IBD). IBD, covered in much more detail in subsection 2.3.3, goes as the following,

$$\bar{\nu}_e + p \longrightarrow n + e^+. \tag{1.2}$$

They chose IBD partly because of its simplicity and partly because liquid organic scintillators, a technology still very much in its infancy at the time [15–17], could be used to detect the charged positron. Cowan and Reines knew from the crude understanding of the expected antineutrino energy spectrum at the time that they would need a very large detector to observe even just a few neutrinos. The largest deployments of liquid organic scintillator at the time were only on the scale of a few litres but Cowan and Reines needed a detector volume roughly three orders of magnitude larger, a very large leap in the world of liquid scintillators.

The electron antineutrinos from the atomic bomb would be identified through observation of the positron produced in IBD. It was understood that, upon its production, the positron would almost immediately annihilate with an electron, producing two gammas that would Compton scatter and produce scintillation light






Figure 1.1: The originally proposed experimental setup for “Project Poltergeist”, featuring an atomic bomb and a detector that would be dropped into a pile of feathers and rubber. Image from [14].

in the detection medium. This light could then be observed by photomultiplier tubes organised around the edge of the detector volume. The detector would need to be located very close to the explosion, within 100 m. To protect the detector, it would be suspended inside a vertical vacuum tank, and, as the bomb was detonated, it would fall freely until after the shockwave had passed. During this time, the detector would gather data for later analysis. This proposed setup is shown in Figure 1.1. Cowan and Reines had predicted their detector would be sensitive to cross-sections of around  $10^{-40}$  cm<sup>2</sup>/proton, but could not be guaranteed below this due to backgrounds. This was a far cry from the predicted neutrino interaction cross-section of  $10^{-43}$  cm<sup>2</sup>/proton, but their experiment was nevertheless approved by senior lab management because it was three orders of magnitude more sensitive than the upper limit on the cross-section at the time [11].

It was later noticed that if one also observes the neutron resulting from IBD then backgrounds could be suppressed massively by requiring that the positron and neutron were detected in coincidence. The neutron would thermalise inside the detector and then capture on hydrogen with a characteristic timescale that no background could mimic at any significant rate. It was also known that materials like

## 1.2. First detection

---

cadmium had very high cross-sections for neutron capture and could be dissolved into a liquid to add neutron sensitivity. Cadmium produces much higher energy capture gammas than hydrogen ( $\sim 9$  MeV vs 2.2 MeV) as the excited cadmium nucleus drops down into a lower energy state through the following interaction,



The neutron capture also takes place on a much shorter timescale than hydrogen, reducing the potential for background contamination. If the backgrounds could be reduced in this way, then the huge neutrino flux from the nuclear bomb was no longer required. Instead, a nuclear reactor would be used to provide a less intense, but more stable (and safer), source of antineutrinos. With this realisation, Cowan and Reines designed and built the first generation of their detector and headed to the nuclear reactor at the Hanford Engineering Works in Washington.

The detector at Hanford was composed of 300 L of liquid scintillator doped with cadmium salt, instrumented with 90 photomultiplier tubes and surrounded by borax and paraffin wax to reduce neutron backgrounds from the reactor. Unfortunately, backgrounds due to cosmic rays and electrical interference prevented the team from reaching the sensitivity they had predicted. The results were promising but inconclusive; a tiny excess of IBD-like events were observed with the reactor switched on relative to when it was switched off, but certainly not enough to be proclaimed as a discovery.

Undeterred, a second detector was designed and constructed 12 m underground at the Savannah River nuclear reactor. The overburden of earth provided some much needed shielding from the cosmic ray backgrounds that had thwarted their efforts at Hanford. The monolithic target volume from the Hanford experiment was replaced with three tanks filled with 1,400 L of liquid scintillator interspersed with tanks filled with 200 L of water, as shown in Figure 1.2. The water was doped with cadmium chloride to add neutron sensitivity. Each tank of scintillator was instrumented with 55 photomultiplier tubes at each end of the tank. With this new detector configuration, when an antineutrino interacted via IBD in the water tank, the annihilation gammas resulting from the positron would penetrate into the




Figure 1.2: Experimental setup at the Savannah River Reactor. The tanks marked 1, 2, and 3 are filled with 1,400 L of liquid scintillator observed by 55 photomultiplier tubes at each end. A and B are filled with 200 L of water doped with cadmium chloride as a neutron-capture agent. Image from [14].

scintillator tanks and produce scintillation light. Meanwhile, the neutron would capture on the cadmium dopant and produce a delayed flash of light. The flash in two separate scintillator tanks followed by a flash in the water tank provided the unmistakable signature of a true antineutrino event, and was easily distinguishable from backgrounds. After taking data for 900 hours with the reactor on and 250 hours with the reactor off, followed by extensive analysis to fully characterise and understand the rate of background events, the results were published [18] and Cowan sent a telegram to Pauli in 1956 declaring that they had detected the neutrino. It was not until 1995, more than two decades after Cowan's death, that the Nobel Prize was awarded for their groundbreaking work [19], with Reines collecting the award on behalf of the pair. Even now, 70 years after their initial work, the delayed coincidence technique remains in use for antineutrino detection.

At the time that Cowan and Reines confirmed their observation of the antineutrino, it was still not clear whether this particle was distinct from the neutrino or if they were one and the same. It took only three years after the

discovery of the antineutrino for Raymond Davis to determine that they were, in fact, different particles. Davis was the first to utilise a radiochemical technique of detecting neutrinos, first proposed by Bruno Pontecorvo in 1948 [20], whereby a neutrino interacting with a chlorine atom would convert it to an argon atom through the neutrino capture process,



He constructed a detector, first at a reactor at Brookhaven National Laboratory [21], and then at the same Savannah River reactor used by Cowan and Reines, filled with 3,900 L of carbon tetrachloride [22]. From lepton number conservation, the interaction in Equation 1.4 must involve a neutrino, not an antineutrino. If the neutrino and antineutrino were distinct particles, then the following interaction induced by an antineutrino,



could not have occurred. Davis measured the relative cross-sections for each of these interactions, and failed to make any concrete observation of the interaction induced by antineutrinos. An upper limit was placed on the cross-section of the antineutrino interaction as being at least twenty times less than the theoretical expectation for the equivalent neutrino interaction, strongly suggesting that the neutrino and antineutrino were not the same.

## 1.3 Neutrino flavours

By 1962, it had been shown that not only were neutrinos and antineutrinos distinct particles, but that there was more than one type, or “flavour”, of neutrino. In the first ever neutrino accelerator experiment, shown in Figure 1.3, Lederman, Schwartz, and Steinberger used the Alternating Gradient Synchrotron at Brookhaven National Laboratory to accelerate protons up to 15 GeV into a beryllium target. Upon striking

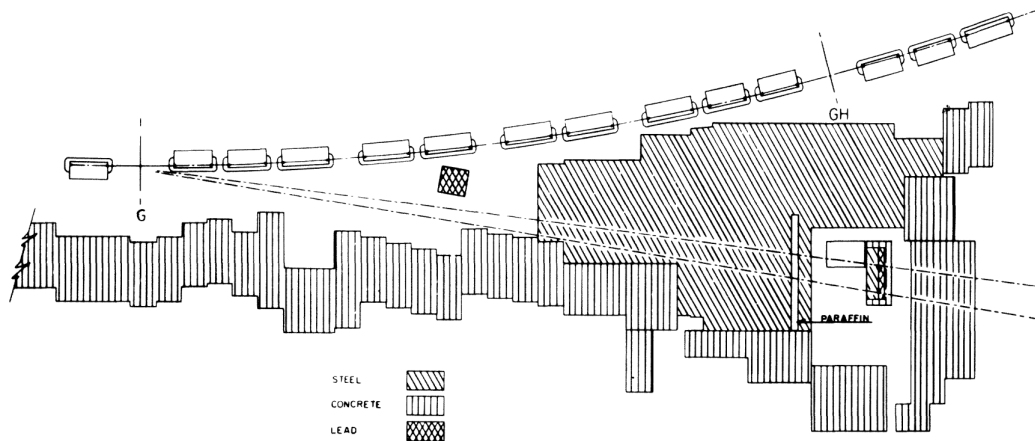


Figure 1.3: Experimental setup at the Alternating Gradient Synchrotron at Brookhaven National Laboratory. Muons and their associated neutrinos are produced from the decay of pions that are expelled from the proton beam dump. Image from [23].

the target, a beam of pions is produced which will then go on to decay in flight,

$$\pi^\pm \longrightarrow \mu^\pm + \nu/\bar{\nu}, \quad (1.6)$$

producing a beam containing muons and neutrinos. Steel shielding with a thickness of 13.5 m was used to stop most of the particles except for the neutrinos, resulting in a beam of muon-associated neutrinos with energies up to around 1 GeV. After the shielding, a 10 tonne aluminium spark chamber was placed to observe the sparks from muons produced when the neutrinos interact within the aluminium. A total of 34 muon tracks were identified, with 5 of them considered to be background produced from cosmic ray muons. The number of electron events observed in the spark chamber was also measured. If the neutrinos associated with muons were the same as those already observed associated with electrons, then it would be expected that the same number of electron events would be observed as muon tracks. However, only six electron event candidates were observed. This naturally led to the conclusion that the electron neutrino and the muon neutrino were indeed different particles, and that lepton number was conserved separately for electrons and muons [23]. Lederman, Schwartz, and Steinberger would later share the 1988 Nobel Prize for their discovery [24]. Their pioneering use of an accelerator to produce a beam




Figure 1.4: The hadron production cross-section near the  $Z^0$  boson resonance, showing the invisible  $Z^0$  decay width. The curves represent the Standard Model prediction for two, three, and four flavours of neutrino. Image from [25].

of neutrinos is a technique that has become widespread in its use today.

After the tau lepton was discovered by the SPEAR accelerator at SLAC in 1975 [26], the existence of its associated neutrino counterpart was generally expected amongst physicists. Evidence of the existence of the tau neutrino had gradually built up over the subsequent decades from indirect measurements of the tau lepton decay width and lifetime at the LEP collider [27], SPEAR [28, 29], and UA1 [30]. The  $W^+$ ,  $W^-$ , and  $Z^0$  bosons, whose existence was predicted by the Glashow, Salam, and Weinberg theory of electroweak interaction [31–33] as mediators of the force (along with the photon), were discovered by the UA1 and UA2 collaborations in 1983 [34]. Due to their large predicted mass, it took ten years from the first observation of weak neutral current interactions at Gargamelle [35] to construct colliders capable of observing the three gauge bosons directly. This discovery allowed

these experiments to begin work on using the  $Z^0$  boson to constrain the number of neutrino flavours. From observation of  $Z^0$  decays, and by subtracting all the observed decays from the total measured decay width of the  $Z^0$ , there remained an “invisible” width resulting from decays that were not observed. It was assumed that these decays were to neutrinos and thus provided a measurement of  $N_\nu$ , the number of light neutrino couplings to the  $Z^0$  boson. Analyses of the invisible width at ALEPH found  $N_\nu = 3.27 \pm 0.30$  [36], and a later combined analysis from the ALEPH, DELPHI, L3, and OPAL collaborations at LEP found  $N_\nu = 2.9840 \pm 0.0082$  [25], strongly supporting the three neutrino hypothesis. This measurement of the invisible width, as well as the Standard Model prediction for the two, three, and four neutrino hypotheses, is shown in Figure 1.4.

It took until 2000 to achieve the first direct detection of the tau neutrino [37]. Following the results from SPEAR, the DONUT (Direct Observation of the NU Tau) collaboration was formed with the aim of directly measuring tau neutrino interactions. Using the TeVatron at Fermilab, protons were accelerated up to 800 GeV into a tungsten beam dump. Proton interactions in the tungsten produced a variety of mesons including the strange-charmed meson,  $D_s$ . Neutrinos were produced via the decay of these mesons,

$$D_s \longrightarrow \tau + \bar{\nu}_\tau. \tag{1.7}$$

The tau lepton would then go on to produce further tau neutrinos when it decays. Some of these neutrinos would undergo charged current interactions in a nuclear emulsion detector 36 m downstream from the beam dump. Only around 5% of the neutrino interactions observed in the emulsion were expected to be associated with tau neutrinos, with the remainder being associated with the decay of other hadrons produced in the beam dump. The tau neutrino interactions would produce tau leptons, which would then decay to a single charged daughter with a branching ratio of 86%. This decay produced a “kinked” track in the emulsion which, due to the short lifetime of the tau lepton, was only a few millimetres in length and could be used to identify the track as a tau lepton. Four interactions with the characteristic kink were identified, and the number of confirmed species of neutrino rose to three.

With this observation, all the neutrinos predicted by the Standard Model had been observed experimentally. There exists three generations of lepton, each of which has its own associated neutrino and each conserves its own lepton number. Neutrinos are left-handed, antineutrinos are right-handed, and, for the moment at least, they are understood to be distinct particles. With the exception of the observation of the Higgs Boson, the Standard Model remains unchanged from this picture.

## 1.4 Neutrino oscillation

During the 1960s, neutrino physics was not only focussed on discovering new types of neutrinos, but also on using neutrinos as tools to investigate others areas of physics. One notable field of study was the mechanisms of energy generation in the Sun as described in the Standard Solar Model. Bahcall had made predictions of the solar neutrino flux [38], building on much earlier work by Bethe [39], and in 1964 he teamed up with Davis to create the Homestake experiment [40]. They used the same radiochemical method involving chlorine that Davis had pioneered a few years earlier in his studies on the neutrino and antineutrino. This radiochemical technique provided no sensitivity to direction or energy, and very little sensitivity to time, but the low threshold of 0.814 MeV made it sensitive to a reasonable proportion of the estimated solar neutrino flux resulting from the p-p chain of nuclear fusion.

The detector was constructed at the Homestake mine in South Dakota, composed of 390,000 L of tetrachloroethylene ( $C_2Cl_4$ , a commonly used cleaning fluid). The mine's overburden of 4,400 m.w.e. (metres water equivalent) provided significant shielding from backgrounds. It became obvious that there was a significant deficit in the number of observed electron neutrinos relative to the number predicted by Bahcall's model. Whether the issue lay with the experiment or with the theory was not certain. Homestake's final results showed that it had detected approximately one third of the expected neutrino flux [41], giving rise to the "Solar Neutrino Problem". Homestake's findings were supported by Kamiokande-II, an imaging water Cherenkov detector which observed much higher energy solar neutrinos using a completely different detection mechanism, which



reported a deficit of around 46% [42]. In addition, since Kamiokande-II used neutrino scattering interactions, and thus had sensitivity to the incoming neutrino direction, it could prove that the neutrinos it observed were really coming from the Sun, something that was not possible with the Homestake experiment. Other radiochemical experiments such as GALLEX [43] and SAGE [44], which could probe even lower-energy neutrinos than Homestake (down to 0.233 MeV), also reported a deficit in the observed number of electron neutrinos, although it was smaller (55-60% of the predicted neutrino flux) than had been reported at Homestake.

The Solar Neutrino Problem persisted for around three decades from when it was first uncovered. Finding the exact reason proved challenging. What was obvious was that all experiments observed some deficit in the expected number of electron neutrinos. Experiments that used the same technique, such as GALLEX and SAGE both using Gallium, largely agreed with one another, so the deficit was not just a result of detector effects. However, experiments that used different techniques, and thus observed different regions of the solar neutrino energy spectrum, saw different deficits. This suggested that there was some relation between the observed deficit and the neutrino energy. This energy relation must not be simply linear, however, because both the high and low energy experiments (Kamiokande-II and GALLEX/SAGE) saw smaller deficits than the moderate energy experiment (Homestake).

A similar problem was also discovered pertaining to atmospheric neutrinos, with both the IMB and Kamiokande experiments finding an unexplained deficit in the number of observed muon neutrinos relative to electron neutrinos [45, 46]. This became known as the atmospheric neutrino problem. Super-Kamiokande, the successor to the original Kamiokande experiment, confirmed results from its predecessor regarding the disappearance of muon neutrinos [47]. More strikingly, this disappearance seemed to be asymmetric with respect to the neutrino's zenith angle as the deficit only manifested for upwards-going muon neutrinos, i.e. those that had travelled into the detector through the Earth.

Ultimately, the resolution to these problems was neutrino oscillation, a

mechanism by which neutrinos can change flavour as they propagate. This had been proposed by Pontecorvo back in 1957 [48], but did not receive much attention at the time. When neutrinos are produced or interact, they are in a well-defined flavour state. However, during the time that they are propagating, this is not the case. The neutrino mass state is not well-defined when it is produced, instead it is in a superposition of three possible mass states. As the neutrino travels, the relative proportions of these mass states change as the different mass states propagate at different speeds. Since each flavour state is a different superposition of the three mass states, and the relative proportions of these mass states change as it propagates, it is now in a mixed flavour state when it comes to interact, so there is now some probability that it will interact as any one of the three flavour states. For the case of two flavour and two mass states, the probability that a neutrino produced as an electron neutrino will be observed as a muon neutrino is given by,

$$P(e \rightarrow \mu) = \sin^2(2\theta) \sin^2\left(1.267\Delta m^2 \frac{L}{E}\right), \quad (1.8)$$

where  $\theta$  is the mixing angle,  $\Delta m^2$  is the difference between the squares of the masses,  $L$  is the distance travelled in kilometres, and  $E$  is the energy in GeV. The value of  $L/E$  is determined by the experiment, so if the expected energy of the neutrino and propagation distance to the detector is known then the oscillation probability can be calculated. In the case of Super-Kamiokande, this explained the disappearance of upward-going muon neutrinos, since these neutrinos had travelled a much greater distance and had oscillated into other flavours.

The Sudbury Neutrino Observatory (SNO) experiment later categorically confirmed the existence of neutrino oscillation. SNO used a heavy water ( $D_2O$ ) detector that was capable of detecting not only charged current interactions, but also neutral current interactions ( $\nu_x + d \rightarrow p + n + \nu_x$ ), which are equally sensitive to all neutrino flavours. This allowed SNO to count the total number of all flavours of neutrino. Although the observed number of electron neutrinos showed a similar deficit to that observed in earlier experiments, when the total number of interactions from all neutrino flavours were counted, it matched the prediction of the Standard Solar Model [49]. This proved definitively that the deficit of neutrinos that had been observed in all other experiments was a direct

result of oscillation to different flavours. Takaaki Kajita and Arthur McDonald of the Super-Kamiokande and SNO collaborations respectively were awarded the 2015 Nobel Prize for this discovery [50].

## 1.5 Chapter summary

- The existence of the neutrino was first proposed as a measure to explain the observed energy spectrum of the outgoing electron in beta decay.
- The neutrino was first observed using a nuclear reactor as a source and coincident detection of the inverse beta decay interaction, a method which is still in use today.
- Three generations of neutrino corresponding to the three generations of lepton were subsequently observed. Experimental measurements have confirmed that these three are the only generations of light neutrino.
- Neutrino oscillation explains the observations of appearance and disappearance of different flavours of neutrinos. The probability of observing a given flavour depends on the distance the neutrino has travelled and its energy.
- There is a massive, ongoing research effort to determine the fundamental properties of the neutrino.

## 1.6 Thesis outline

This thesis focuses on the practical application of detecting reactor antineutrinos for the purposes of nuclear nonproliferation. More specifically, two new technologies that could be utilised in far-field reactor antineutrino detection are presented, water-based liquid scintillator and Large Area Picosecond PhotoDetectors. Both of these nascent technologies require much study to determine their potential impact.

- Chapter 2 is an overview of nuclear safeguarding and how reactor antineutrino detection could be introduced as an additional tool for this purpose. The

theory of reactor antineutrino detection is presented in detail, along with the technologies and techniques that are currently used, and what information these techniques can reveal about the nuclear reactor being observed.

- Chapter 3 introduces the WATer CHerenkov Monitor for AntiNeutrinos (WATCHMAN) collaboration and its technical and scientific objectives. In this chapter, the most probable conceptual design of the WATCHMAN detector is presented, as of the period of time when this work took place, spanning from 2019 to early 2022. The chapter details the design, construction, operation, potential capabilities, and future plans of this conceptual detector design.
- Chapter 4 details a simulation study undertaken in late 2020 to assist in finalisation of the detector conceptual design based on a number of key performance criteria. The study presented in this chapter analyses the sensitivity of many possible detector designs to a nuclear reactor complex situated at a distance of approximately 150 km from the detector. In particular, the study investigates the impact that a water-based liquid scintillator detection medium could have on the detector's capabilities.
- Chapter 5 introduces the Large Area Picosecond PhotoDetector, a newly developed ultra-fast photosensor with positional sensitivity that is under consideration for use by the WATCHMAN collaboration as well as many other future neutrino experiments. The theory, operation, and capabilities of this detector are discussed, as are the techniques and challenges that are present in fast timing scenarios.
- Chapter 6 presents the results from the characterisation campaign of a Large Area Picosecond PhotoDetector using a laser test stand. Many properties of the detector, including the gain, dark rate, and timing resolution are determined.
- Chapter 7 introduces a method for reconstructing photon hits on the Large Area Picosecond PhotoDetector in scenarios where the photodetector is being struck by many photons in a short period of time. The performance of the algorithm is verified using simulations of the photodetector geometry, and

a detailed simulation of the detector response. In addition, the integration of this algorithm and the photodetector itself into a larger water Cherenkov antineutrino detector is presented.

- Chapter 8 presents the conclusions of this thesis, and the avenues for future work.


# Chapter 2

## Reactor neutrinos

### 2.1 Nuclear power

The general principle underpinning nuclear power is that unstable, neutron rich isotopes can undergo induced fission when they absorb a neutron, causing the isotope to split into smaller fission fragments, accompanied by a release of energy. This energy produces heat, which can be used to boil water into vapour to spin a turbine resulting in the generation of electricity. The fission process releases a number of free neutrons that will then proceed to induce more fission reactions if there is additional fissionable material present. In this way, a chain reaction of fission can occur which can produce vast amounts of electricity. In order to be absorbed by the atom and cause fission, a neutron must be slowed down to thermal energies using a substance known as a moderator. The fission neutrons elastically scatter on the moderator nuclei until enough energy is lost that they can be captured by a fissionable isotope. Typical moderators used in nuclear reactors include water and graphite. Control rods, made from neutron absorbing material such as cadmium or boron, are used in a nuclear reactor to restrain the fission rate so that the chain reaction can continue sustainably. The more neutron absorbing material that is present, the fewer neutrons can go on to induce more fission, slowing down the fission rate and reducing the power output of the reactor. The inverse is true if the neutron absorbing control rods are removed from the reactor.

There are four dominant fissionable isotopes that account for the vast majority of energy released in a nuclear reactor. These are  $^{235}\text{U}$ ,  $^{238}\text{U}$ ,  $^{239}\text{Pu}$ , and  $^{241}\text{Pu}$ . Reactor fuel is often enriched to increase the proportion of fissile  $^{235}\text{U}$  as it only makes up around 0.7% of natural uranium.  $^{238}\text{U}$  (which makes up >99% of natural uranium) by contrast is fissionable, but it is not fissile, and as such is not able to sustain a fission chain reaction. The initial fuel loading of a new reactor might be around 3.5%  $^{235}\text{U}$ , with the rest made up of  $^{238}\text{U}$  [51].



$^{239}\text{Pu}$  and  $^{241}\text{Pu}$  will start being produced as soon as the reactor is switched on through neutron capture interactions on  $^{238}\text{U}$ , as shown in Equation 2.1. These two plutonium isotopes will also contribute to the overall generation of energy as they are both fissionable ( $^{239}\text{Pu}$  is also fissile). Despite this, there is still more plutonium produced throughout the reactor cycle than is removed through fission reactions. A standard pressurised water reactor has a net positive production of around 200 kg of plutonium per year [51]. From the perspective of nuclear nonproliferation, this is the area where safeguarding is so crucial. It is vital to be sure that the operator of a nuclear reactor is not diverting fissile material that is produced throughout the reactor cycle during the refuelling process, as this material could be used in the creation of a weapon.

## 2.2 Reactor safeguarding

The International Atomic Energy Agency (IAEA) is the regulatory body responsible for ensuring that nuclear reactors around the world are only being used for peaceful purposes. This is achieved through the process of “safeguarding”, an array of technical measures to verify that the activities of a state’s nuclear facilities are not being used for the development of nuclear weapons. Safeguards are generally applied through legally binding agreements with nations, which authorise the IAEA to confirm compliance through the implementation of these technical measures at nuclear facilities. The IAEA verifies the activities of over 1,300 nuclear facilities around the world, and monitors around 200,000 Significant

Quantities of nuclear material (a Significant Quantity is the approximate amount of nuclear material from which the possibility of producing a weapon can not be excluded, for  $^{239}\text{Pu}$  this is 8 kg [52]).

The IAEA has a number of tools available for the purposes of safeguarding [53]:

- The most basic tool available is nuclear material accountancy, whereby the declared amount of nuclear material possessed by a state is verified by inspectors. This may include fuel assemblies, rods, and masses of powdered uranium or plutonium, with the aim being to discover if any bulk collections of nuclear material are unaccounted for.
- Non-destructive analysis, such as neutron counting or gamma ray spectroscopy, are used to determine if there is any difference between the declared and measured amount of an item in the inventory. In addition, destructive analysis techniques that provide better sensitivity may be used when required.
- Containment and surveillance techniques ensure that nuclear material only follows declared routes. This includes methods such as CCTV monitoring in crucial locations and the use of tamper-resistant seals that allow inspectors to verify that the containment of a declared material has not been breached.
- Environmental sampling using sensitive analytical techniques can provide assurance that no nuclear material has been present in the location the sample was taken. Environmental samples are therefore often taken in areas where there is no legitimate reason for the presence of nuclear material.

While the tools currently at the disposal of the IAEA have largely been effective in safeguarding nuclear reactors, there are limitations that have been identified. The IAEA have noted that the current safeguarding regime has focussed largely on declared nuclear material, but was ineffective for clandestine, undeclared nuclear activities [54]. The IAEA can only inspect facilities that are made known through declarations, and there are currently very few options to discover undeclared facilities except for through standard intelligence gathering methods. Nations operating clandestine nuclear facilities may also be reticent to allow IAEA inspectors into declared facilities as these inspections can be perceived



to be intrusive. In addition, many of the techniques in use do not perform direct measurement of the isotopic content of the reactor core, instead relying on regular inspections of seals on fuel assemblies and CCTV monitoring of spent fuel. When direct measurements are made, they are performed at the beginning or end of reactor fuelling cycles, and as such no real-time verification is possible. The current safeguarding regime also requires significant manpower. With the number of nuclear facilities whose regulatory compliance falls under the remit of the IAEA rising every year [54], this has resulted in increased strain on the IAEA's limited personnel and budget.

Reactor antineutrinos offer a promising opportunity to resolve some of the issues outlined with traditional methods of reactor safeguarding. All nuclear reactors emit huge numbers of antineutrinos. Due to the nature of the neutrino, this signal is impossible to shield or alter, either by the operators of the nuclear reactor, or by external conditions (unlike, for example, how satellite imaging can be adversely affected by cloud cover). Furthermore, the antineutrino emission is inextricably linked to the fission occurring inside the reactor core, resulting in a direct, real-time measurement of the fissile content, something that is not offered by traditional safeguarding methods. The antineutrino emission is continuous, and so a detector could be set up to autonomously monitor the antineutrino emission over time, saving the cost and manpower requirements of regular inspections after the detector has been constructed.

There are two foreseeable options as to how antineutrino detection could be incorporated into the existing safeguarding framework. The first is in a cooperative monitoring scenario, whereby a nation would agree to host an antineutrino detector at the nuclear facility under scrutiny. The detector would verify that the nation is operating the reactor according to its obligations as part of the cooperative monitoring agreement. The detector could be placed just outside the boundary of the nuclear facility to minimise site intrusion and to help build trust with the nation in question. The alternative is in a non-cooperative scenario, whereby a much larger detector that is sensitive to the antineutrino emission at much greater distances would be constructed outside the borders of a nation under scrutiny. The detector would then be able to identify if reactors are

operating that have not been declared, or if there are large changes in the operation of the known reactors.

## 2.3 Reactor antineutrinos

The emission of antineutrinos from a reactor is a direct result of the fission process. Each fission releases about 200 MeV of energy and produces two (or, rarely, three) unstable, neutron-rich, fission fragments that subsequently preferentially undergo beta decay (as shown in Equation 2.2) three times each on average to reach nuclei with a stable ratio of protons to neutrons. This process results in a typical production of 6 electron antineutrinos per fission [55, 56] with energies of up to around 10 MeV, or around 5% of the total energy released in the fission.



This translates to an antineutrino flux in the region of  $10^{20}$  antineutrinos per second per gigawatt of thermal reactor power [57]. The flux is isotropic, and so the inverse-square law determines the flux at a given distance from the reactor. This flux will also be modulated by an additional contribution resulting from neutrino oscillation as some fraction of the electron antineutrinos will have changed flavour depending on the distance from the reactor at which they are detected.

### 2.3.1 Antineutrino flux

Measurement of the reactor antineutrino flux by an antineutrino detector can provide an indication of the reactor's operating state, power output, and core content. Determining the operating state of a reactor using antineutrinos is straightforward if the detector is close to the reactor; many more antineutrinos are produced whilst the reactor is producing energy, and the drop in the antineutrino rate when the reactor powers down is significant. Aside from the operating state, the reactor antineutrino flux is primarily determined by the thermal power of the reactor. The antineutrino flux will increase as the thermal power increases, as more fissions are occurring per second, and decreases when the thermal power is




Figure 2.1: The fission fragment yields for each of the four major fissioning isotopes. Image from [61].

reduced. However, there is a non-negligible impact on the overall flux from the core composition. Each of the four dominant fissioning isotopes produces a distinct distribution of beta-decaying fission fragments, as shown in Figure 2.1. Each of these fragments have varying decay rates and outgoing electron energy spectra. This results in a differing average number of antineutrinos produced for each isotope, and as such the overall antineutrino emission is the combination of the individual contributions from each fissioning isotope. The change in the fission fractions over the reactor cycle is known as the burnup effect. Simply by measuring the antineutrino flux over time, a measurement of the relative change in the core composition can be performed. In general, the antineutrino flux reduces as burnup continues because the fission fraction of  $^{235}\text{U}$  decreases as the fission fraction of  $^{239}\text{Pu}$  increases, and  $^{239}\text{Pu}$  produces significantly fewer antineutrinos per fission for the same energy output [58, 59]. However, this pattern can depend on the reactor type and the initial core composition [60].

The exact antineutrino flux and energy spectrum produced by a nuclear reactor is complex and depends on many parameters. The spectral estimate of the

### 2.3. Reactor antineutrinos

---

antineutrino emission per fission of a single fissioning isotope is given by,

$$\lambda(E_{\bar{\nu}}) = \exp\left(\sum_{j=0}^5 a_j E_{\bar{\nu}}^j\right), \quad (2.3)$$

where  $a_j$  are predictions of the fit parameters by Huber and Mueller [62, 63]. Figure 2.2 shows the spectral estimate for each of the main fissioning isotopes. The average energy of the emitted antineutrinos for each spectrum can be used to calculate the average thermal energy per fission,  $Q_i$ . The contribution to the overall reactor antineutrino energy spectrum from a given reactor for the  $i$ th fissioning isotope is given by,

$$\phi_{\bar{\nu},i}(E_{\bar{\nu}}) = P_{th} \frac{p_i \lambda_i(E_{\bar{\nu}})}{Q_i}, \quad (2.4)$$

where  $P_{th}$  is the thermal power output of the reactor core,  $p_i$  is the fraction of the thermal output that is produced by the  $i$ th isotope, and  $Q_i$  is the average energy produced per fission.




Figure 2.2: The energy spectrum of emitted reactor antineutrinos for the four dominant isotopes. Image from [64].

To relate this to the detectable antineutrino emission at a detector, the antineutrino interaction cross-section must be taken into account. For elastic scattering or inverse beta decay interactions (see subsection 2.3.3), the total




Figure 2.3: The cross-section for detecting antineutrinos from each of the dominant fissioning isotopes, given by the product of the inverse beta decay cross-section and the neutrino flux at that energy. Image from [61].

detectable antineutrino spectrum at a distance  $d$  from the reactor for each fissioning isotope, before taking into account neutrino oscillation, can be described by,

$$R_i(E_{\bar{\nu},i}) = \frac{N\sigma(E_{\bar{\nu}})}{4\pi d^2} \phi_{\bar{\nu}}(E_{\bar{\nu},i}), \quad (2.5)$$

where  $N$  is the number of target particles, and  $\sigma(E_{\bar{\nu}})$  is the antineutrino interaction cross-section for the given interaction. The summation of this over each fissioning isotope gives the total detectable antineutrino spectrum.

### 2.3.2 Antineutrino energy

It is also possible to determine information about the reactor core content through measuring the energy spectrum of the emitted antineutrinos. Since, as mentioned in subsection 2.3.1, each fissioning isotope has a distinct distribution of beta-decaying fission fragments, the energy distribution of the antineutrinos resulting from each of the fissioning isotopes will also be distinct. Figure 2.3 shows that the fissions of  $^{235}\text{U}$  result in a harder antineutrino energy spectrum than fission from  $^{239}\text{Pu}$ . Since the likelihood of detecting an antineutrino increases with the antineutrino energy, and there is typically some lower energy threshold (see subsection 2.3.3) below which




Figure 2.4: The Feynman diagram for inverse beta decay. Time runs from left to right.

no detection can occur, this further shows that antineutrinos from  $^{235}\text{U}$  will be preferentially detected over  $^{239}\text{Pu}$ . The observed antineutrino energy spectrum will be the combination of the individual contributions from each fissioning isotope, weighted by their associated fission fraction. The antineutrino energy spectrum is important to understand since, as mentioned in section 1.4, the neutrino survival probability depends on  $L/E$ , and so predicting the detectable antineutrino flux at any given distance from a reactor requires accurate predictions of the expected energy spectrum at the reactor.

### 2.3.3 Antineutrino interactions

Antineutrinos can be observed through IBD or antineutrino scattering interactions. Each interaction has its own signature and provides different information about the incoming antineutrino.

IBD, shown in Figure 2.4, was used by Cowan and Reines in the first observation of the antineutrino, and remains the primary method by which to detect reactor antineutrinos. It has a cross-section on the order of  $10^{-44} E_e p_e \text{ cm}^2$  [65, 66]. The

threshold for this interaction is 1.8 MeV (a straightforward estimate<sup>1</sup> of this is given by the mass difference between the initial and final states,  $E_{\text{thresh}} \sim m_n + m_e - m_p$ ) Only about 30% of reactor antineutrinos have energies above this threshold, significantly reducing the observed antineutrino flux. The visible energy of the positron is given by  $E_{\text{vis}} = E_\nu - 1.8 \text{ MeV} + 2 \times 0.511 \text{ MeV}$ , and as such there is a direct correlation between the energy of the initial antineutrino and the observed energy deposition by the positron. This relationship allows for reconstruction of the antineutrino energy using the detector response. The positron resulting from the IBD interaction is produced with a slight backwards bias relative to the direction of the incident antineutrino at reactor antineutrino energies [65]. The neutron retains the majority of the antineutrino's momentum and will always be produced in the forward hemisphere. It will then quickly scatter as it begins its random walk through the detector medium before finally thermalising. As such, the neutron retains some memory of the initial antineutrino direction [67, 68] but the random walk through the series of elastic scatters makes it difficult to use the neutron to deduce direction information pertaining to the antineutrino. Upon thermalisation, the neutron will then be captured in the detection medium, usually on hydrogen unless the medium is doped with neutron capture agents. The resulting de-excitation produces a second signal. The requirement that the positron and neutron signals are observed in coincidence, i.e. that the signals are closely related spatially and temporally within the detector, can be used as a powerful tool for suppressing detector backgrounds.

Antineutrino scattering interactions are an alternative to IBD for observing reactor antineutrinos. Elastic scattering of an antineutrino on an electron,

$$\bar{\nu}_e + e^- \longrightarrow \bar{\nu}_e + e^-, \quad (2.6)$$

has a smaller cross-section at around  $10^{-45} m_e E_\nu \text{ cm}^2$  [69, 70]. There is no threshold for this interaction in theory, however in practice the detector will have an effective threshold below which these interactions will be lost amongst the backgrounds. Since only the electron in the final state is observable, elastic scattering interactions do not benefit from the coincident detection technique meaning that discrimination

---

<sup>1</sup>More precisely, in the laboratory frame with the proton at rest:  $E_{\text{thresh}} = \frac{(m_n + m_e)^2 - m_p^2}{2m_p} = 1.806 \text{ MeV}$

of these events can be difficult, especially at lower energies in the case of reactor antineutrinos. One potential benefit of elastic scattering interactions in the context of reactor monitoring is the directionality provided by the final state electron is strongly tied to the initial direction of the antineutrino. If an antineutrino detector could reliably reconstruct electron directions from elastic scattering interactions, this could provide the ability for detectors to determine the direction of previously unknown nuclear reactors [70], although this capability has never been demonstrated in practice.

Another scattering interaction that could be used for reactor antineutrino detection is coherent elastic neutrino nucleus scattering (CE $\nu$ NS) [71], given by,

$$\bar{\nu}_e + A \longrightarrow \bar{\nu}_e + A. \quad (2.7)$$

The cross-section of CE $\nu$ NS is the largest of all low-energy neutrino couplings and scales with the square of the neutron number of isotope  $A$ . Despite this, CE $\nu$ NS has proven difficult to observe experimentally because it requires detecting the nuclear recoil signature. Since it has only recently been observed, CE $\nu$ NS is not likely to be useful for reactor monitoring applications for the foreseeable future.

## 2.4 Reactor antineutrino anomaly

The antineutrino fluxes from the four dominant fissioning isotopes have been calculated multiple times since the Cowan and Reines experiment. In 2011 the antineutrino spectrum was recalculated using a combination of the summation method [63], which uses the fission fractions and a database of beta decay information to calculate the overall contribution from every possible decay branch, and an improved electron to antineutrino conversion [62] using measurements of the beta spectrum from thermal neutron fission made in the 1980s [72–74]. This recalculation resulted in an increase in the expected antineutrino flux of around 6% compared to previous calculations. All previous observations made by short baseline antineutrino experiments, which were formerly in agreement with expectations, now showed a deficit in the expected antineutrino rate. This discrepancy is known as the reactor antineutrino anomaly [75].






Figure 2.5: The ratio of the re-evaluated summation model (SM) to the Huber-Mueller model (HM) for each fissionable isotope. Image from [82].

In addition, the RENO experiment [76, 77], and later the Daya Bay [78] and Double Chooz [79] experiments, revealed an excess in the antineutrino flux at around 5 MeV — known as the “5 MeV bump”. Lots of potential explanations for this bump have been explored, including incorrect modelling of certain fissionable isotopes [80] and non-standard neutrino interactions [81], but the source of the bump is yet to be confirmed.

The reactor antineutrino anomaly has largely been resolved by re-evaluation of the antineutrino energy spectra using a new summation method and updated information from nuclear databases [82]. This re-evaluation had the largest impact on the modelled antineutrino energy spectra from  $^{235}\text{U}$  fission, where the updated spectrum is normalised around 5% lower than the Huber-Mueller model, as shown in Figure 2.5. This re-evaluation reduced the expected detectable antineutrino flux, largely bringing the theory into agreement with experimental findings.

## 2.5 Demonstrations of antineutrino monitoring

Scientists in the Soviet Union were the first to recognise the potential application of antineutrino detection for monitoring nuclear reactors in 1978 [84]. Pioneering antineutrino experiments were conducted at the Rovno Atomic Energy Station in Ukraine throughout the 1980s, with some of the measurements being directly relevant to reactor safeguarding [83]. The detector used at Rovno was perhaps the first demonstration of applied antineutrino physics. It consisted of 1,050 L of




Figure 2.6: The measured antineutrino rate over the course of the reactor cycle at the Rovno detector. The smooth solid line shows the prediction based on simulation of the reactor core. Image from [83].

gadolinium-doped organic liquid scintillator, separated by light-reflective material into a 510 L inner target volume and a 540 L outer volume used as a shield against backgrounds, and was instrumented with 84 photomultiplier tubes. The detector was constructed directly below a 440 MWe pressurised water reactor with 18 m of overburden, providing an extremely high antineutrino flux and a high degree of shielding from cosmogenic backgrounds. With the reactor on, the detector observed approximately 900 antineutrino events per day, compared to around 150 with the reactor off. The Rovno experiment clearly showed, through measurements of the antineutrino flux, that the detector was sensitive to both the operating status of the reactor and the evolution of the core content. The observed 6% change in the antineutrino rate, shown in Figure 2.6, was well-matched with predictions based on simulated core evolution. The thermal power was predicted using only measurements of the antineutrino flux with a precision of around 2%. From measurement of the antineutrino energy spectrum, the group also predicted the mass of fissile material consumed over the cycle, which closely agreed with predictions based on the records of the reactor thermal power.

Later in the United States, an independent effort to demonstrate the practical application of antineutrinos took place at the San Onofre Nuclear Generating Station (SONGS). In the early 2000s, the SONGS1 detector was constructed, consisting of 600 kg of gadolinium-doped liquid scintillator in four cells, each instrumented with two photomultiplier tubes. The detector was located approximately 20 m from the




Figure 2.7: The antineutrino rate observed by the SONGS1 detector throughout a single refuelling outage of a pressurised water reactor. Image from [85].

reactor. The detector operated largely unattended for five years, observing around 500 antineutrino events per day. A muon veto was used to reduce cosmogenic backgrounds. SONGS1 made very similar findings to the earlier Rovno experiments. It was demonstrated that the detector was sensitive to the reactor state, the reactor power, and the evolution of the core fuel content over the course of the reactor cycle [85–87]. Figure 2.7 shows the obvious reduction in antineutrino rate as the reactor underwent refuelling, followed by an obvious increase in antineutrino rate when the reactor was switched back on.

All demonstrations for the purpose of reactor monitoring have thus far been “near-field” detectors i.e. located very close (within a few metres) to the nuclear reactor. Mid-field to far-field detectors (at distances on the kilometre scale) have not yet demonstrated applications for reactor monitoring. Experiments such as KamLAND [88], which uses a kilotonne-scale liquid scintillator detector, have observed the combined antineutrino emission from many reactors at distances of over 100 km. These detectors, at distances beyond the near-field deployments, have been used for the purpose of precision neutrino oscillation measurements, but their existence lays the groundwork for detectors at similar scales to be applied for

nonproliferation purposes.

## 2.6 Technologies for antineutrino detection

Since the primary interaction of interest for most antineutrino experiments is IBD, the detection medium must be a dense source of protons in order to achieve a reasonable interaction rate. For this reason, there are two dominant technologies for antineutrino detection: water Cherenkov detectors and scintillation detectors.

### 2.6.1 Water Cherenkov detectors

Water Cherenkov detectors are large volumes of water instrumented with photodetectors that utilise the Cherenkov effect [90] to detect charged particles. Cherenkov radiation is electromagnetic radiation produced by a charged particle that is travelling faster than the phase velocity of light within a medium, i.e.,

$$\beta > \frac{1}{n}, \quad (2.8)$$

where  $\beta = v/c$  and  $n$  is the refractive index of the material the particle is travelling through (in the case of pure water this value is roughly 1.33). Cherenkov radiation is often compared to the sonic boom that is produced when an aircraft travels faster than the speed of sound in air. The electromagnetic radiation forms spherical wavefronts of constructive interference produced along the particle's path of travel, as shown in Figure 2.8, resulting in a characteristic Cherenkov cone of blue light. The open angle of the cone,  $\theta$ , is given by,

$$\cos(\theta) = \frac{1}{n\beta}. \quad (2.9)$$

As such, for a particle travelling at velocity  $v \approx c$ , the open angle in water will be,

$$\arccos\left(\frac{1}{1.33}\right) = 41.2^\circ. \quad (2.10)$$




Figure 2.8: Diagram showing the production of wavefronts of electromagnetic radiation (blue arrows) as a result of the particle (red arrow) travelling with velocity  $v > c/n$ . Image from [89].

For electrons or positrons travelling in water, the kinetic energy required to pass the Cherenkov threshold is,

$$E = \frac{m_e c^2}{\sqrt{1 - \frac{1}{n^2}}} = 0.8 \text{ MeV}, \quad (2.11)$$

where  $m_e$  is the electron mass.

Cherenkov emission covers a broad spectrum of wavelengths but is peaked in the blue and near-ultraviolet. The total energy emitted as a particle travels through the medium is governed by the Frank-Tamm equation [91], which states that the energy  $dE$  emitted per unit of the particle path length  $dx$  per unit frequency  $d\omega$  is given by,

$$\frac{\partial^2 E}{\partial x \partial \omega} = \frac{q^2}{4\pi} \mu(\omega) \omega \left( 1 - \frac{c^2}{v^2 n^2(\omega)} \right), \quad (2.12)$$

which is valid as long as,

$$\beta = \frac{v}{c} > \frac{1}{n(\omega)},$$

where  $\mu(\omega)$  and  $n(\omega)$  are the permeability and refractive index of the medium, respectively, both of which are dependent on the frequency.

Since the number of photons emitted increases with the energy of the particle, the energy of the particle can be estimated by counting the number of photons

observed either through counting the number of triggered photodetectors or integrating the total observed charge in the photodetectors, depending on the regime the detector is operating in. The position of the interaction vertex can be determined by analysis of the spatial and temporal distribution of the signals observed by photodetectors. In general, the ability to reconstruct position and energy improves as the number of photons increases. As a result, higher energy particles tend to be reconstructed more accurately than those with lower energies. As will be discussed in more detail in subsection 2.6.2, Cherenkov detectors have a relatively low light yield — the number of photons produced per unit of deposited energy — compared to alternative detection technologies such as liquid scintillator. This limits the vertex and position resolution that can be achieved in water Cherenkov detectors. Since Cherenkov light is highly directional, it is also possible to reconstruct the direction of the particle by analysing the geometry of the triggered photodetectors and predicted interaction vertex. In a large enough detector with enough photodetectors, such as Super-Kamiokande, the Cherenkov ring can be observed simply from the pattern of triggered photosensors.

Water-based detectors have the advantage that the volume of the detection medium can be easily scaled up without complications, allowing for incredibly large detectors at the scale of tens of kilotonnes. The directional nature of Cherenkov light also provides insights into the particle directions. This may be especially useful in the case of nuclear nonproliferation, where the statistical distribution of the positron directions resulting from antineutrino elastic scattering interactions (see subsection 2.3.3) could provide a hint as to the direction of the reactor relative to the detector. However, the Cherenkov threshold of  $\beta > 1/n$  enforces a hard limit on the detector threshold of about 0.8 MeV for electrons, although in reality the actual analysis threshold will be much higher than this (typically  $> 3$  MeV [92]) to remove detector backgrounds arising from impurities in the detector itself. The limited light yield is also a limiting factor on the possible energy and position resolution.

### 2.6.2 Liquid scintillator detectors

Reactor antineutrinos have also been observed using liquid scintillator detectors, as in the case of the first neutrino detection by Cowan and Reines. Liquid scintillator

detectors are fundamentally very similar to water Cherenkov detectors in that they are large tanks of liquid instrumented with photodetectors to observe the flashes of light from neutrino interactions. The major difference is that they use an organic material that exhibits scintillation when a charged particle ionises the medium, rather than the Cherenkov effect. Plastic scintillator has also been used to observe reactor antineutrinos as in the case of the MiniCHANDLER experiment [93]. The process of detection is similar to that of liquid scintillators, but will not be explored further in this thesis.

### Scintillation mechanism

Organic scintillators are made up of aromatic hydrocarbons as a solvent. These contain phenyl groups that are rich in  $\pi$  electrons — delocalised electrons that allow the molecule to absorb energy from ionising radiation. Charged particles deposit energy in an organic scintillator mainly as a result of Coulomb interactions with bound electrons. Gamma rays can also deposit energy despite their lack of electric charge, predominantly through three processes: the photoelectric effect, Compton scattering on orbital electrons, and pair production of an electron and positron. These three processes produce energetic charged particles that can then interact with bound electrons through Coulomb interactions as previously mentioned. The ionising radiation excites the  $\pi$  electrons into higher energy states, as shown in Figure 2.9. Electrons that are excited into vibrational states quickly transition to excited states non-radiatively on the order of picoseconds. The transition from an excited state back to the ground state can occur through two different mechanisms depending on whether they were excited into a singlet or triplet state. For the singlet state, the electron relaxes through emission of a photon. Solving the Schrödinger equation for these  $\pi$  orbitals reveals energy states in the range of 3 to 6 eV [94], leading to the emission of photons in the blue or ultraviolet wavelengths. This process occurs quickly, typically on the scale of a few nanoseconds depending on the molecule, and is known as *fluorescence*. Alternatively, the electron can reach a triplet state through inter-system crossing from the singlet state. The triplet state is very stable and thus long-lived; the electron can only relax back to the ground state by emitting a photon when in the presence of another excited molecule. This process

takes much longer than singlet transitions, typically milliseconds or longer, and is known as *phosphorescence*. The energy of the emitted photon is often less than the absorbed energy, mainly as a result of the relaxation from vibrational energy states. This difference between the absorbed and emitted energy is known as the Stokes shift.




Figure 2.9: The singlet,  $S$ , and triplet,  $T$ , vibrational states for a fluorescent molecule. Solid lines represent energy transitions resulting in fluorescence or phosphorescence whereas dashed lines represent non-radiative transitions. Image from [95].

### Fluors

Although the medium is able to produce light from ionising radiation, it is not especially practical for detecting radiation as the solvent will self-absorb the majority of light that it emits. Notably, this is not the case for noble gas liquid scintillators as the excitation process differs. To reduce the self-absorption of an organic scintillator a fluorescent molecule, known as a fluor, is added to shift the wavelength of the emitted photons by absorbing and then re-emitting at longer wavelengths. The



most efficient method by which energy is transferred from the solvent to the fluor is through non-radiative dipole-dipole coupling known as Förster resonant energy transfer (FRET) [96]. The energy transfer is dependent on the distance between solvent and fluor molecules; if the fluor concentration is too low then very few photons will be shifted to wavelength regions where the solvent is transparent. In contrast, high concentrations can cause the fluor to self-absorb in much the same way as the solvent can before adding fluors.

The addition of a secondary fluor allows for increased concentrations of the primary fluor by shifting the wavelength again and further reducing the self-absorption. At this stage, the energy transfer is mostly radiative since the concentrations of the primary and secondary fluor are typically too low for FRET to have a major effect, usually  $\mathcal{O}(10^{-3})$  g/g and  $\mathcal{O}(10^{-6})$  g/g, respectively. The secondary fluor must have an absorption spectrum that significantly overlaps the emission spectrum of the primary fluor and an emission spectrum that is significantly separated from the absorption spectrum of the primary fluor to be efficient in reducing self-absorption (i.e. a significant Stokes shift). A common combination of wavelength shifters uses 2,5-diphenyl-oxazole (PPO) as a primary fluor and 1,4-bis(5-phenyloxazol-2-yl) benzene (POPOP) as a secondary fluor, due to their favourable absorption and emission regions. An example of how the wavelength is shifted as the energy is absorbed and re-emitted is shown in Figure 2.10.

### Scintillator properties

For neutrino experiments the most pertinent properties of the scintillator are often the light yield and time profile of the light emission, and how these differ with varying fluor concentrations. In general, increasing the fluor concentration also increases the light yield as energy that would otherwise be eventually lost to heat in the solvent is transferred to the fluor resulting in photon emission. Increasing the fluor concentration also makes the scintillator “quicker”, i.e. the peak of the emission time profile is earlier. This is because a higher fluor concentration reduces the distance between molecules and so FRET occurs more efficiently. By adjusting the fluor concentration, an experiment can tune the scintillator to meet their requirements




Figure 2.10: An example of the emission and absorption wavelengths of the different scintillation components. Image from [95].

for the time profile and light yield.

### Quenching

In the ideal case, the scintillator light output is directly proportional to the energy of the particle. For this to be the case, the ionisation of the scintillating medium must take place over the course of a path length whereby excited molecules are separated by several molecular lengths, ensuring that interactions between excited molecules can be neglected. However, this ideal case does not always occur, and the light output does not linearly increase with the particle energy. This phenomenon is known as “quenching”.

The most prominent form of quenching is “ionisation quenching”. This often occurs when heavy charged particles, such as alpha particles, deposit much of their energy in a very small path length, resulting in regions of localised high excitation density and a reduction in the efficiency of the scintillation mechanism. Birks’ formula was developed in 1951 [97] to describe a scintillating medium’s response to a charged particle. It describes the light output per unit of distance travelled in

medium by the particle, and includes the quenching probability.

$$\frac{dL}{dx} = \frac{L_0 \frac{dE}{dx}}{1 + kB \frac{dE}{dx}}, \quad (2.13)$$

where  $\frac{dL}{dx}$  is the light output per unit path length,  $L_0$  is the light yield constant,  $\frac{dE}{dx}$  is the energy deposition per unit path length, and  $kB$  is Birks' constant, or the quenching parameter, which depends on the scintillating material. The effect of quenching is strongest when a particle has high specific energy loss,  $dE/dx$ . An alpha particle can have a detector response greater than ten times less than an electron of equivalent energy due to ionisation quenching [95].

An unrelated form of quenching is the result of impurities in the scintillator, often due to oxygen dissolving in the material. These impurities do not exhibit fluorescence and so any excitation energy that is transferred to them during the scintillation process will not be emitted as light and is ultimately lost. Impurity quenching can be reduced by purging the scintillator of impurities using nitrogen. Furthermore, shielding the detector fluid from ambient air using a blanket of an inert gas such as nitrogen can reduce the rate of impurities dissolving into the scintillator.

### Comparison to water Cherenkov

Liquid scintillators produce orders of magnitude more photons per unit of energy deposited than the Cherenkov effect, depending on the exact scintillator. In the case of a neutrino detector, this greatly helps with position and energy reconstruction, resulting in a better resolution simply due to the greater number of photons. There is also effectively no threshold for a charged particle to produce scintillation light and the overall detector threshold can be lower than in an equivalent water Cherenkov detector. Furthermore, liquid scintillators can be purified to very high radiopurity levels [98] that have not been matched by water Cherenkov detectors [99]. However, scintillator detectors do also have drawbacks. Since scintillation light is produced isotropically, it is difficult to retain any directional information if the path length of the particle in the detector is short. Additionally, there is the practical issue that scaling up liquid scintillator is notoriously difficult as most organic scintillators are some combination of toxic, combustible, or expensive. Some scintillators, such as the

now commonly used linear alkylbenzene, are more suited than others to deployment in large quantities due to its limited toxicity, low cost, and high flashpoint [100].

### 2.6.3 Gadolinium doping

In the context of nuclear nonproliferation, where IBD is the primary interaction of interest, the neutron capture is of great interest for the coincident detection technique detailed in section 1.2. However, the neutron capture on the hydrogen nucleus only releases a single 2.2 MeV gamma, which will deposit energy in the detector through Compton scattering on electrons in the detection medium. This amount of energy deposition is well below the analysis threshold in most water Cherenkov detectors [101]. One solution to this problem is doping the water with gadolinium, which was first proposed in [102]. Gadolinium doping has previously been used in multiple liquid scintillator neutrino detectors [103, 104].

Gadolinium has an incredibly large cross-section for neutron capture that is approximately five orders of magnitude greater than hydrogen (48,800 b vs 0.3 b [105]). Additionally, when the gadolinium nucleus de-excites, it does so by emitting a gamma cascade with a total energy of approximately 8 MeV, easily observable above background in a liquid scintillator or water Cherenkov detector. By doping a water Cherenkov detector with just 0.2% concentration of gadolinium sulphate by mass, this results in up to 90% of neutron captures occurring on a gadolinium nucleus [106]. The mean time required for the neutron to be captured is also reduced from around 120  $\mu$ s in the case of hydrogen, to 30  $\mu$ s on gadolinium. This reduces the probability that the coincidence window will be contaminated by an uncorrelated background event, resulting in a false pair.

The EGADS project was created in 2009 [107] to demonstrate that gadolinium-loaded water could be a practical detection medium for large-scale neutrino experiments like Super-Kamiokande. The research and development program at EGADS successfully showed that the addition of gadolinium sulphate would not have any negative effect on the detector components of Super-Kamiokande. The proposed bandpass filtration system proved capable of adding and removing gadolinium from the detection medium whilst maintaining




Figure 2.11: A micelle encapsulating a droplet of organic scintillator inside a volume of water. Image from [110].

water quality with attenuation lengths in the region of 100 m. In addition, EGADS investigated the new neutron backgrounds that arise from gadolinium-doping. The low number of free neutrons in the water will now become visible after their capture on gadolinium, and the gadolinium sulphate itself contain impurities which will contribute to the overall radiopurity of the detector. After the success of EGADS, Super-Kamiokande has recently entered a new phase of operation which includes gadolinium-doping, and thus much improved neutron tagging capabilities.

#### 2.6.4 Water-based liquid scintillator

Water-based liquid scintillator (WbLS) is a novel detection media under development at Brookhaven National Laboratory [108, 109] that attempts to combine the advantageous properties of both water and liquid scintillator. In practice, WbLS would allow for a large improvement in light yield and thus energy resolution, due to the presence of the scintillator, whilst greatly reducing the issues regarding safety and cost when scaling liquid scintillators to large volumes.

Molecules of organic scintillator do not naturally mix with water. Instead, a surfactant is required to bring the water and oily molecules into a stable mixture. The surfactant has hydrophilic (lipophobic) and lipophilic (hydrophobic) regions on either end of the molecule, reducing the surface tension between the organic and water molecules. This can emulsify the organic molecules in the water molecules

by forming a hydrophilic shell around a droplet of organic molecules, known as a micelle (shown in Figure 2.11). These micelles have diameters on the order of 10 nm and can retain stability on the timescale of years. Higher concentrations of surfactant allows higher concentrations of organic scintillator to be emulsified in water, but this can be at the cost of the optical properties of the medium. Higher surfactant concentrations can cause the opacity of the medium to increase, harming the attenuation length which is vital for large-scale neutrino detectors. Despite this, the idea of a tunable scintillator concentration is very appealing. With an active filtration system, a detector's target medium could even be tuned during operation. The tuning would change not only the light yield, but also the time profile of light emission.

In a more advanced scenario, it is possible separate the Cherenkov and scintillation signals from one another [111], allowing different analyses to be performed on the different components of the total light output. The scintillator component would provide the large number of photons necessary for improved performance in vertex and energy reconstruction. Meanwhile, the Cherenkov component can provide directional information using the topology of the Cherenkov ring and also timing information from the event since the Cherenkov light is produced immediately compared to the slower scintillation light. It may also be possible to perform particle identification in a large-scale WbLS detector, because particles of different masses will produce different ratios of Cherenkov and scintillation photons. Furthermore, there is likely to be a difference in scintillator quenching effects between proton recoil interactions and electron-like particles that result in differing scintillation time profiles. It has been shown that it is possible to differentiate between neutrons and gamma rays in water-based scintillators using pulse shape discrimination since neutrons generally produce scintillation light with longer tails in the timing distribution [112].

Cherenkov and scintillation light can be separated in two ways: by time [113, 114] or by wavelength [115]. Separation by timing uses the fact that Cherenkov light is emitted instantaneously as a particle travels through the medium, whereas scintillation light is delayed due to the time required for the medium to be excited and then emit photons by de-excitation. Depending on the concentration and the

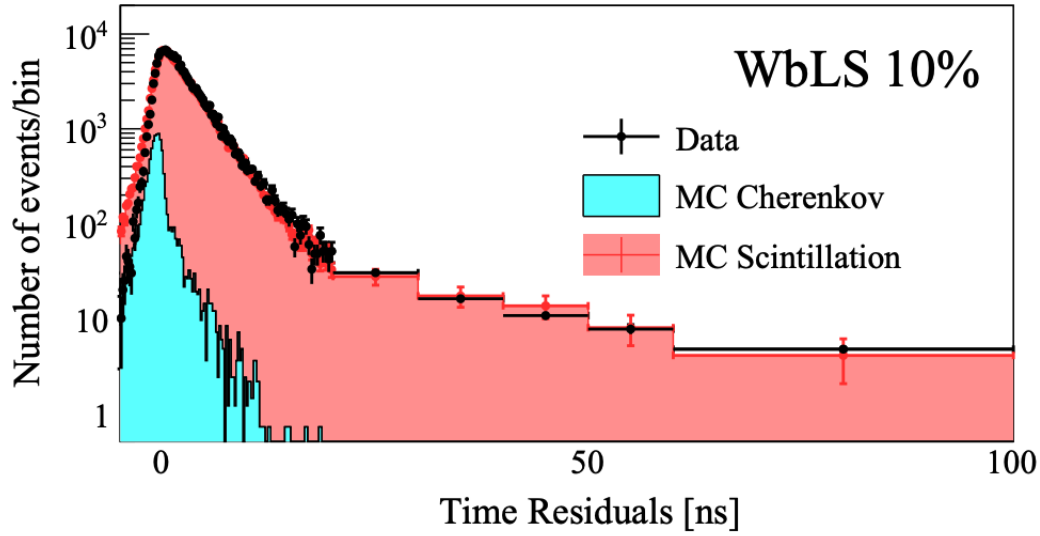


Figure 2.12: Plot showing the different emission time profiles for the Cherenkov and scintillation components of WbLS, allowing for separation of the light by wavelength. The red and blue distributions show the results of Monte Carlo (MC) simulation and the black dots show the real data. Image from [113].

exact scintillator used, the delay between the peak of the timing distributions of the Cherenkov and scintillation light might be on the order of nanoseconds or less. For this reason, fast photosensors are required such as Large Area Picosecond PhotoDetectors (detailed from chapter 5 onwards). With a transit time spread of less than 100 ps, one should see a distinct doubled-peaked timing distribution, with the earlier peak resulting from Cherenkov light and the delayed peak from scintillation light. Figure 2.12 shows how Cherenkov and scintillation light can be separated by timing.

Separation by wavelength exploits the difference in the emission spectra between Cherenkov and scintillation light, shown in Figure 2.13. Scintillators tend to emit in a relatively narrow band of wavelengths whilst Cherenkov light peaks in the ultraviolet and blue wavelengths but is ultimately broadband in its emission. If one only selects photons with longer wavelength, by choosing a photodetector with appropriate quantum efficiency in the wavelength regions of interest or through the use of bandpass filters, then it is possible to get a relatively pure sample of Cherenkov photons. Likewise, it is possible to get a sample of largely scintillation

photons by only selecting photons with wavelengths inside the scintillator's emission region. Due to the broadband nature of Cherenkov, there is likely to be some contamination of Cherenkov photons. If using higher concentrations of scintillator, where the number of produced scintillation photons largely outnumber the Cherenkov photons, this can still provide a relatively pure sample of scintillation photons. However, it is worth noting that in this scenario, the Cherenkov sample may observe a large amount of contamination from scintillation photons simply due to leakage through the bandgap filter. The advantage of using this technique is that one can use slower, and therefore cheaper, photomultiplier tubes (PMTs), which will result in a greater area coverage of photodetectors per unit cost. There is a further possibility for separating the light components using topological information since Cherenkov light will be emitted in the direction of the particle and scintillation light will be isotropic. By observing the angular dependence of PMT hits in a detector, it may be possible to observe a higher density of PMT hits in the direction of the particle. In a large-scale detector employing WbLS, it is possible that one or all of these techniques may be used together to most effectively separate Cherenkov and scintillation light.

Active development is ongoing into doping WbLS with neutron-capture agents such as gadolinium to improve the neutron sensitivity of the medium. Current challenges in the development of WbLS include the optimisation of optical parameters such as the attenuation length, as well as the design of a suitable filtration system. Filtration presents a particular challenge for the utilisation of WbLS. The water and organic components must be separated as they need to be treated separately by the filtration system. The water then goes through a standard water purification system as used in previous water Cherenkov detectors. The organic component passes through a nanofiltration system that must remove impurities without damaging the integrity of the micelle structures. The two components must then be recombined at the end of the filtration system before returning to the detector fluid.



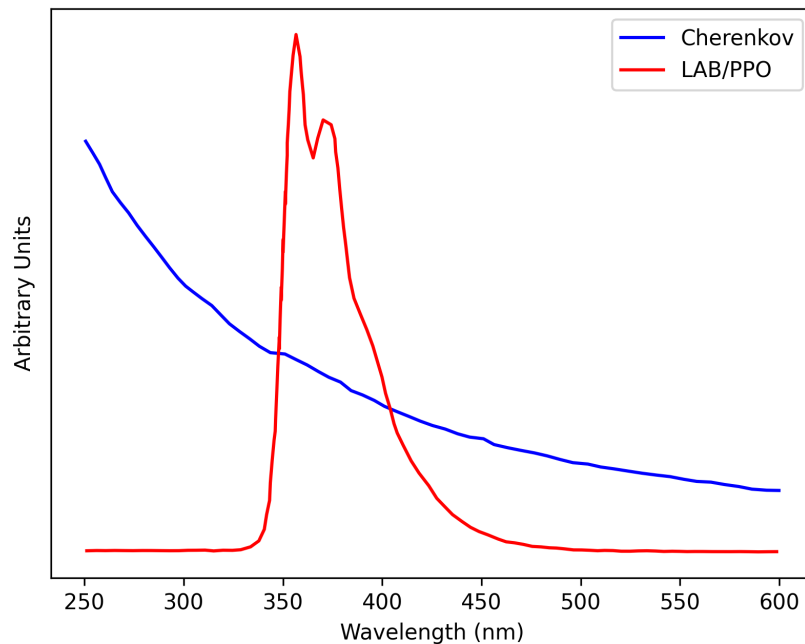


Figure 2.13: The emission spectrum for Cherenkov light and the common scintillation cocktail LAB/PPO. Above 450 nm one can select a relatively pure sample of Cherenkov photons.

## 2.7 Chapter summary

- Nuclear reactors are the most powerful artificial source of antineutrinos, producing around  $10^{20}$  per second per gigawatt of thermal energy.
- Observing the antineutrino emissions has been proposed as a new technique for detecting and monitoring nuclear reactors.
- The antineutrino emission from a nuclear reactor can provide information as to the reactor's operating state, power output, and fuel composition.
- Near-field detection of nuclear reactors has been demonstrated multiple times using liquid scintillator detectors with target volumes of a few hundred tonnes.
- Water Cherenkov detectors provide the potential to scale up a detector to several thousands of tonnes in order to observe nuclear reactors at much greater

distances.

- Liquid scintillator detectors have much higher light yields than water Cherenkov and as such are capable of more precise measurements of the antineutrino interactions but are more difficult to scale up to large volumes.
- Gadolinium doping offers much better sensitivity to the neutron in inverse beta decay interactions allowing for suppression of backgrounds using the coincident detection technique.
- Water-based liquid scintillator is a novel detection medium that aims to combine the benefits of both water and liquid scintillator.

## Chapter 3

# Reactor Antineutrino Detector Prototype

The WATER CHerenkov Monitor for AntiNeutrinos (WATCHMAN) scientific collaboration was created with the intention of demonstrating the ability of a kilotonne-scale water-based antineutrino detector to remotely monitor a nuclear reactor using its antineutrino emission. WATCHMAN is a joint effort between US and UK academic institutions and national laboratories with funding from the Office of Defence Nuclear Non-Proliferation in the US as well as from the Science and Technology Facilities Council and the Ministry of Defence through the Atomic Weapons Establishment. The collaboration consists of over 120 collaborators across 21 universities and 5 national laboratories.

The project is split into two branches: the Advanced Instrumentation Testbed (AIT) and Neutrino Experiment One (NEO). These make up the facility and the initial antineutrino detector respectively. This chapter outlines the goals of the collaboration, proposed designs for the facility and detector, as well as a breakdown of the proposed detector subsystems. These were outlined in the 2019 conceptual design report [116]. As the collaboration remains at the conceptual design stage, many details are likely to change in the future.

## 3.1 The Advanced Instrumentation Testbed

The AIT is intended to be a world-class deep-underground facility for the research and development of new technologies enabling antineutrino monitoring of nuclear reactors. It will be the host facility for the initial detector proposed by the WATCHMAN collaboration, NEO. The AIT will be constructed at the STFC Boulby Underground Laboratory in North Yorkshire, England. The underground lab is located within Boulby mine, an active salt, potash and polyhalite mine, owned and operated by Israel Chemicals Ltd UK. It is the second deepest mine in Europe, with tunnels as deep as 1.4 km underground. The lab, located 1.1 km under the surface, has hosted numerous deep underground experiments dating back to the 1990s. In the past these were predominantly dark matter experiments, but currently there is a broad multidisciplinary array of active experiments ranging from dark matter to studies of microbial life. Additionally, Boulby Underground Laboratory houses the Boulby Underground Germanium Suite (BUGS), a set of germanium semiconductor gamma spectrometers for ultra-low background material screening studies.

Due to the lab's overburden of 2,800 m.w.e., the cosmic ray muon flux is attenuated by  $\mathcal{O}(10^6)$  relative to the surface [117]. Furthermore, the lab has the lowest level of ambient radon of any deep underground lab with an average radon concentration of less than  $3 \text{ Bq m}^{-3}$  [118]. These aspects of the Boulby Underground Lab make it an attractive environment for the WATCHMAN collaboration since both cosmic ray muons and ambient radon produce backgrounds that are troublesome for neutrino detectors.

Boulby Underground Laboratory is located 26 km from the EDF Hartlepool Nuclear Power Station, which will be the primary source of reactor antineutrinos for the NEO detector. Further afield, at distances of 148 km and 187 km respectively are the Heysham and Torness reactor complexes, also operated by EDF. This collection of nearby nuclear reactors at a range of standoffs adds to the opportunities for demonstrating practical reactor monitoring at the Boulby site.

As part of AIT, a new cavern would be constructed in the Boulby mine some distance away from the existing underground lab, likely in an area of tunnels

branching out underneath the North Sea. The specification of the facility is not yet completed, but it is expected that the cavern will be large enough to host the initial detector itself, as well as associated clean rooms, a room to host data acquisition (DAQ) electronics, and staff rooms.

## 3.2 Experimental objectives

The primary nonproliferation objectives of AIT-NEO relate to the Hartlepool Nuclear Power Station. There are a number of demonstrative scenarios that are potentially of interest to the nuclear nonproliferation community.

The first is the detection of an unknown nuclear reactor. In this scenario, the two cores of Hartlepool will be treated as an unknown nuclear reactor and every other reactor in the world is treated as part of the known reactor background. The NEO detector will then measure the antineutrino flux at Boulby until the Hartlepool reactor signal is observed at a  $3\sigma$  significance over background. This is the most basic measurement, and is considered to be the baseline scenario. It emulates a situation where a reactor is surreptitiously constructed by a state that may intend to use it for non-peaceful purposes.

The second is a similar scenario except one of the two Hartlepool reactor cores is treated as part of the known world reactor background and the other is treated as the unknown signal. This is a much more challenging scenario, largely due to the size of the background contribution from the known Hartlepool core that is very large. This scenario emulates a situation where a clandestine reactor is obscured by constructing it in close proximity to a larger, declared, and legally operating reactor.

There are also more ambitious nonproliferation objectives of the experiment. One of these is the detection of nuclear reactor complexes at distances even farther than that of Hartlepool. At a distance of 148 km, the Heysham nuclear complex on the west coast of the UK allows for the demonstration of a potential non-cooperative monitoring scenario. In this scenario, an antineutrino detector might be positioned many hundreds of kilometres from reactors of interest, most likely outside that country's borders.

NEO will also have an extensive scientific program to contribute to neutrino physics. Firstly, NEO will be able to make measurements of neutrino oscillation parameters, in much the same vein as experiments such as Daya Bay, Double Chooz, and KamLAND. Whilst the primary baseline will be the 26 km distance to Hartlepool, there is also the potential to use the Heysham and Torness nuclear complexes further afield. This measurement would require energy resolution that will most likely not be possible with a standard gadolinium-doped water fill but may be possible with advanced detection media such as WbLS that provide a higher light yield.

Secondly, in the event of a galactic supernova, NEO will be able to observe the antineutrinos emitted during the event. The supernova neutrinos could be detected through the IBD interaction like reactor antineutrinos. NEO will have a supernova alarm which constantly watches for bursts of IBD events that could be indicative of a supernova.

The WATCHMAN collaboration also aims to achieve a number of auxiliary objectives. The first of these is to increase collaboration between the nonproliferation and particle physics communities in the UK and the US. Related to this is the desire to encourage and train early-career scientists both in the UK and US to participate in nonproliferation science and policy. Finally, WATCHMAN will aim to develop new technologies and techniques enabling reactor monitoring with antineutrinos.

### **3.3 The NEO detector**

The NEO detector will be the first experiment constructed at the AIT. Although the design has not yet been finalised at the time of writing, the NEO conceptual design is a kilotonne-scale gadolinium-doped water Cherenkov detector. The tank size is in the region of 20 m in diameter and 20 m in height; the optimal size and configuration of the detector was subject to a trade study, the results of which are detailed in chapter 4. The active volume will be instrumented with 1,600 to 6,400 Hamamatsu R7081 10 inch PMTs, resulting in a photocoverage of between 10% and 40%. Surrounding the active volume is a veto region instrumented with outward-facing PMTs to enable tagging of background events induced by cosmic




Figure 3.1: Cutaway CAD model of a NEO detector conceptual design by Jan Boissevain. The diagram shows the steel truss structures at the top and bottom, the surface lab, the access spaces between the tank and cavern walls, and the inner detector volume with PMTs.

rays.

#### 3.3.1 Tank

The tank refers to the main vessel that contains the detection medium and structures required to support the photosensors. The tank will be constructed from welded sections of 316L stainless steel to ensure compatibility with the gadolinium dopant and a potential WbLS fill in the future. One of the conceptual designs for the tank is shown in Figure 3.1. It will be an upright cylinder, with the exact height and diameter still to be determined. It is expected that the resulting fiducial volume will be in the region of 1 kilotonne. Other tank design geometries, including a spherical tank, a rectangular tank, and acrylic inner volumes, were disfavoured on either cost

or logistical grounds.

Below the tank is a concrete base built into the cavern, which supports the floor of the tank on a set of steel girders, that can be shimmed to ensure that the structure is level. The floor of the tank is slightly sloped (around  $2^\circ$ ) concentrically towards the centre, to allow for a pipe to be connected to drain the tank when necessary.

At the top of the tank is an arrangement of steel support trusses that are welded to the inside of the tank. These provide support for the tank lid, which doubles up as the floor for the lab above the tank, and the internal PMT support structure. The tank lid will also provide a gas-tight seal, allowing for an inert gas to be pumped across the surface of the contained fluid to prevent degradation of the purity of the detection medium. A series of holes will be cut into the tank lid to create ports for calibration access, water circulation, personnel access, as well as for routing signal, high voltage, and slow control cables. All of these ports are accessible from the top-deck cleanroom, which sits on top of the tank lid. In addition to the ports, the cleanroom will also contain the electronics racks to operate the photosensors and monitoring systems.

The tank will be wrapped in thermally insulating material to ensure temperature stability of the detection medium. Magnetic compensation coils will be attached to the inside wall of the tank to reduce the detrimental affect of the Earth's magnetic field on the collection efficiency of the PMTs.

#### **3.3.2 PMT support structure**

The PMT support structure is the internal cylindrical tank structure to which all photosensors are affixed. The structure comprises two almost identical endcaps at the top and bottom and a vertical cylinder. The two endcaps provide a series of concentric circles of mounting points for PMTs facing both inwards (towards the fiducial volume) and outwards (for veto PMTs). A conceptual design of the top endcap can be seen in Figure 3.2. The main difference between the two endcaps is that the top endcap must accommodate a central plug of 3 m in diameter. The PMTs mounted directly below the plug will not be on the main truss structure but fixed to the underside of the plug. The vertical structure will be supported from






Figure 3.2: A conceptual design of the top endcap of the PMT support structure. Image from [116].

the top truss structure and attached to the floor of the tank. Like the endcaps, the vertical rods of the vertical structure will provide mounting points for both inwards- and outwards-facing PMTs. Between the inward- and outward-facing PMTs, layers of black opaque sheeting will be placed to optically isolate the inner target volume from the outer veto volume.

Signal and high voltage cables will be run along the rods making up the endcaps and vertical structure which run from each PMT through a feedthrough port in the tank lid and then into the electronics racks.

#### 3.3.3 PMTs

More detail on PMTs and their operation can be found in section 5.1. NEO will utilise thousands of Hamamatsu R7081-100 [119] high quantum efficiency 10 inch PMTs with low radioactivity glass. Each PMT is provided with a waterproof potting assembly, making it suitable for deployment in a water Cherenkov detector. The window is made from borosilicate glass and features a bialkali photocathode with a spectral response from 350 nm to 600 nm, peaking at 420 nm.

The quantum efficiency is around 35% at 390 nm. Borosilicate glass naturally contains radioactive material, predominantly  $^{238}\text{U}$  chain isotopes,  $^{232}\text{Th}$  chain isotopes and  $^{40}\text{K}$ . The low radioactivity version of this PMT model uses low radioactivity glass which reduces the concentration of these isotopes by up to a factor of ten.

The R7081-100 contains ten dynode stages and offers a typical gain of  $10^7$  at a nominal anode-to-cathode supply voltage of 1,500 V, although the supply voltage can be increased to 2,000 V if higher gain is required at the cost of higher dark rate. At nominal voltage, the typical dark rate is in the region of 8 kHz when measured at 25°C. It includes a built-in voltage divider in the PMT base to supply each of the ten dynode stages. The photocathode is grounded, and positive voltage is supplied at the anode. As such, the signal output is superimposed onto the DC high voltage supply line meaning the PMT requires only a single cable. A splitter circuit, made up of a high-pass RC circuit, is required to isolate the signal and pass it on to any other front-end electronics. The resulting signal has a nominal rise time of around 3.8 ns, whilst the overall PMT timing resolution is in the region 3.4 ns (FWHM).

The PMTs will undergo a characterisation campaign that will determine and compare the following parameters of an initial batch of 100 Hamamatsu R7081-100 PMTs with the vendor datasheet:

- Dark count rate
- Supply voltage that achieves a gain of  $10^7$
- Peak-to-valley ratio
- Afterpulse rate

The dark count characterisation measures the thermionic emission from the photocathode. The measurement is made by counting the number of single photoelectron pulses produced by the PMT in a total amount of live time, where the live time is given by the total time for which the digitising electronics are sampling the PMT waveforms. It is important to know the expected dark rate for a collection of PMT, since this will affect how triggering and reconstruction algorithms are implemented. For example, the triggering algorithm will typically

count the number of PMT hits in a rolling time window. As such, it is important to understand how many dark pulses are expected to occur in this time window to ensure that a detector-wide trigger is not issued because of a few coincident dark hits. All the 100 tested PMTs showed dark rates below the acceptance limit of 10 kHz, with a mean of approximately 1.8 kHz [120].

The gain calibration tested the operating voltage required to reach a gain of  $10^7$  by increasing the supply voltage and finding the single photoelectron peak in the gain distribution. If the required gain could not be achieved before exceeding the nominal maximum supply voltage of 2,000 V, the PMT would be rejected. The measurements showed that all 100 PMTs met this requirement. Whilst the characterisation campaign has thus far only been performed on the initial batch of 100 PMTs, ultimately all PMTs will be fully characterised before deployment.

Light enhancement options are being evaluated as a means to achieve greater effective photocoverage without incurring the cost of purchasing additional photosensors. One such option is Winston cone light concentrators, such as those used in the SNO experiment [121]. A light concentrator is a cone with a parabolic profile that fits around the glass window of the PMT. The inside of the concentrator is coated in a specular reflective surface such that any light that enters the cone will be reflected until it reaches the PMT surface, effectively widening the acceptance area of the PMT. Winston cones are one such implementation of the light concentrator, which uses the tangent-ray principle to ensure that any photon entering the cone with an incident angle below some critical angle (defined by the height of the cone and the diameter of the upper aperture) will always eventually reach the PMT surface. The advantage of light concentrators is that they can be produced relatively cheaply compared to the cost of an additional PMT — likely around 10-20% of the PMT cost. However, the introduction of reflections will produce additional uncertainties in PMT timing which has a direct impact on the performance of event reconstruction algorithms. Furthermore, ageing effects on the reflective surface of the concentrator must be considered, as these will reduce the light collection efficiency over time.

The possibility of encapsulating each PMT in an acrylic enclosure is currently

under consideration. This would remove any concerns regarding the PMT's compatibility with the detection medium or ability to withstand the pressure of submersion, as well as protecting neighbouring PMTs in the case of an implosion event. It also isolates much of a PMT's natural radioactivity to within its own encapsulation, reducing the chances of escaping into the water and producing backgrounds in the detector. The drawbacks of the encapsulation approach include the additional cost of manufacturing the acrylic enclosures, as well as introducing more opportunities for photons to reflect or refract in the acrylic. The refractive index must be matched as closely as possible and the boundary between the PMT window and the acrylic must have good optical coupling to retain light collection efficiency.

The PMTs, or the encapsulations if they are chosen, will be fully tested for compatibility with the liquid medium, especially with the gadolinium dopant. The dominant compatibility uncertainty is between the low-radioactivity glass and gadolinium, since this has not been tested before. The procedure for this testing involves submerging two PMTs in a solution containing 1% gadolinium sulphate for multiple years. This gadolinium concentration of this solution is five times higher than the proposed concentration to be used in the detector in order to simulate the effects of many years of exposure to gadolinium on a shorter timescale. The absorbency of the solution is measured using UV-VIS spectroscopy; increased absorption at UV and optical wavelengths indicates deterioration in the PMT material. A further ten PMTs are submerged in a 0.2% gadolinium-sulphate solution - the same concentration as is expected in the detector.

In addition to compatibility testing, the PMTs will also be pressure tested to ensure their mechanical strength is sufficient to withstand the pressure of the liquid medium. Like with the compatibility tests, the PMTs will be tested at a higher pressure than expected inside the detector to simulate many years of submersion in the tank. The PMTs that underwent gadolinium-soaking for the compatibility tests will also be pressure tested and compared with non-soaked PMTs to study if the gadolinium had any lasting impact on the mechanical strength.

The PMTs will also undergo radiological screening to ensure that the low

radioactivity glass matches the supplier specifications. Radiological screening involves gamma spectrometry and neutron activation analysis to determine the levels of contamination. Preliminary studies have yielded results congruent with previous testing campaigns at the Double Chooz and Borexino experiments. Ensuring the radiopurity of the PMT glass is vital since the PMTs are a major contributor to the overall radiological contamination of the detector.

#### 3.3.4 Calibration

The NEO tank features a number of access ports intended for calibration systems. An array of these ports are placed at multiple locations in the tank lid from which calibration sources can be deployed into the detector medium from gloveboxes. There will be an increased density of ports near the edges of the detector, as this is where spatial variations in the detector response is expected to be largest.

There are two broad types of calibration systems that will be deployed in the NEO detector: photon sources and radiological sources.

##### Photon calibration sources

A multitude of LEDs and laser fibres will be mounted throughout the detector, both in the inner detector volume and in the veto volume. This will ensure that every PMT can be illuminated by more than one photon source as these will be used to run continuous gain calibration of all PMTs throughout the detector's operation. This continuous gain calibration is required as a PMT's gain is expected to drift throughout its lifetime and its operational voltage may need to change to bring it back into accordance with all other PMTs. Laser fibres will allow for continuous monitoring of the quality of the liquid medium through attenuation and scattering measurements.

In addition to the permanently fixed photon sources inside the detector, a deployable light injection system is also under consideration. This system will consist of a laser fibre deployed from one of the lid access ports coupled to a diffuser ball that will be lowered into the detector. The diffuser ball will scatter the laser light, resulting in an almost isotropic and instantaneous light source. The

isotropic light will illuminate several PMTs, depending on the intensity of the laser light, which can then be used to calibrate the relative PMT timing and charge characteristics.

#### **Radiological calibration sources**

These calibration systems consist of gamma or neutron radiation sources that are deployed from the ports on the lid of the tank. The proposed radiological calibration sources are:

- $^{232}\text{Th}$ , a 2.6 MeV gamma source<sup>1</sup>
- Americium-Beryllium (AmBe), a 4.4 MeV gamma source with a reverse-tagged<sup>2</sup> neutron
- $^{16}\text{N}$ , a 6.1 MeV gamma source with a tagged electron

The energies of these sources offers wide coverage of the reactor antineutrino energy spectrum (and thus the IBD positron energy spectrum) for calibrating the detector's energy response. In addition, the ability to deploy these sources from several positions on the lid allows for studies of the spatial variations in the detector response at different energies, which is important for defining the detector's fiducial volume. The tagged sources are particularly powerful, as the tag provides an unambiguous signal that separates the calibration source from backgrounds.

An AmBe source consists of  $^{241}\text{Am}$  and  $^9\text{Be}$ .  $^{241}\text{Am}$  is an alpha source which, when paired with  $^9\text{Be}$  produces excited carbon and a neutron through the following reaction,



The excited carbon atom subsequently de-excites producing a 4.4 MeV gamma. The neutron will enter the detector medium before capturing, most likely on gadolinium,

---

<sup>1</sup>The gamma emission results from the decay of  $^{208}\text{Tl}$ , a member of the  $^{232}\text{Th}$  decay chain.

<sup>2</sup>Tagging refers to the coincident detection of the calibration particle in conjunction with another easily observable particle which acts as a trigger. Reverse-tagging is when the trigger occurs after the production of the calibration particle, so the data acquisition system must look backwards to find the calibration event.

and emit an isotropic gamma cascade summing to around 8 MeV. The source will be coupled with a gamma detector consisting of a scintillating crystal and PMT that will provide the reverse tag to after a neutron capture has been observed. This source will be used to calculate the detection efficiency of neutron captures, a crucial parameter to understand for a detector that will primarily observe IBD interactions. In addition to this, the AmBe system will enable studies of the relative isotropy of single gamma and neutron capture events.

$^{16}\text{N}$  has been used as a calibration source in other neutrino experiments such as Super-Kamiokande [122] and SNO [123]. As shown in Figure 3.3, it commonly ( $\sim 66\%$ ) emits a 6.1 MeV gamma following its beta decay (with an electron endpoint energy of 4.3 MeV) to an excited state of  $^{16}\text{O}$ , making it a useful high energy calibration point as it is largely monoenergetic.  $^{16}\text{N}$  is produced through the  $(n, p)$  interaction on  $^{16}\text{O}$ ,



A deuterium-tritium (DT) generator produces 14.1 MeV neutrons through the fusion of deuterium and tritium,



and has proven to be a suitable neutron source to produce  $^{16}\text{N}$  in appreciable quantities. At the energy of a DT neutron, the cross-section for the interaction shown in Equation 3.2 is approximately 42 mb. In this calibration system, gaseous carbon dioxide is used as the source of oxygen atoms. A chamber is constructed around the fusion region of the DT generator and carbon dioxide is pumped in. The DT generator bombards the chamber with neutrons, producing  $^{16}\text{N}$  atoms in the chamber. The DT generator cannot be located close to the detector because of the risk to personnel from neutron radiation and to prevent stray neutrons from reaching the active volume of the detector. This constraint, combined with the nuclide's short half life of 7.1 seconds, means that the gaseous mixture of carbon dioxide and  $^{16}\text{N}$  must be pumped at high pressure from the DT generator to the lid of the tank and into the detector volume through one of the access ports. Inside the detector, a decay volume is lined with a thin layer of scintillator and a




Figure 3.3: The decay scheme of  $^{16}\text{N}$ . Beta and gamma decays are represented by solid and dashed arrows respectively.

PMT. This system allows the gamma to escape into the detector volume, but the electron will be stopped in the scintillator and act as the tag for coincident detection of both particles. The  $^{16}\text{N}$  source will primarily be used for calibrating the detector’s absolute energy scale (including measuring spatial non-uniformities) as well as providing a cross-check on energy and vertex reconstruction algorithms.

### 3.3.5 Data acquisition

The DAQ system is responsible for extracting timing and charge information from the tank PMTs by conditioning and digitising the raw PMT waveforms. It will also implement the triggering logic to initiate detector readout on an event of interest as well as packaging the data produced by every photosensor into an “event” that can be saved for future offline analysis. In implementing the detector trigger logic, the DAQ system will utilise fast algorithms to reconstruct the vertex and energy of detector events, the output of which will determine if the event is likely to be signal and should trigger a full detector readout.

Only DAQ solutions that allow for full waveform digitisation are under consideration. This is because it is assumed that digital signal processing techniques will be used to counteract the expected large effects of electromagnetic



noise on the detector due to the surrounding mine environment. Additionally, having access to the full waveforms is expected to improve the ability to reconstruct muon events, which produce many photons in the detector, in order to improve the overall detector deadtime.

The DAQ system has been designed to meet the following requirements for the initial gadolinium water detector fill, assuming a nominal PMT gain of  $10^7$ :

- Single PMT trigger threshold of 0.25 photoelectrons
- Time resolution of 1.0 ns
- analogue to digital converter (ADC) dynamic range of 12 bits
- Detector trigger threshold of 6 hits in a 100 ns time window
- Able to manage a maximum average PMT dark rate of 10 kHz

The requirements will likely be more stringent in the case of a gadolinium-loaded WbLS detector fill as a result of the higher light yield. Downward-going cosmogenic muons occur at a rate of around 0.1 Hz and will produce in the region of 20,000 photons, taking into account detection efficiency. The system will need to be able to handle these events without saturation in order to track the muon as it travels through the detector. This track will be used to establish a cylindrical veto region along the muon's trajectory in an effort to eliminate backgrounds resulting from the decay of muon spallation-induced radionuclides.

#### 3.3.6 Advanced photodetectors

NEO will also be a testbed for future neutrino detection technologies that may offer benefits for nuclear nonproliferation. The research and development effort into deploying advanced photodetectors in NEO is very much partnered with the development of alternative detection media such as the aforementioned WbLS. As discussed in subsection 2.6.4, traditional photosensors will not have the capabilities to separate the Cherenkov and scintillation components of the light and thus the additional particle information that WbLS offers will not be exploitable.

One of the advanced photosensor options is the Large Area Picosecond

PhotoDetector™ (LAPPD), due to its potential timing resolution in the tens of picoseconds and its millimetre scale positional resolution. The LAPPD is discussed in much more detail from chapter 5 onwards.

Another advanced photosensor option under consideration is the dichroicon [115]. The dichroicon is a Winston cone light concentrator made from dichroic filters coupled to a pair of PMTs. Since dichroic filters reflect a set of wavelengths above or below some bandpass value and transmit the rest, they can be used to separate out the Cherenkov and scintillation light by their wavelength and guide each component of the light to a separate photosensor. This negates the need for expensive, fast timing photosensors like the LAPPD.

#### 3.3.7 Water treatment system

The detection medium must be continuously purified in order to maintain reasonable attenuation and scattering lengths. It is envisaged that the water treatment system will be based on a scaled-up version of the EGADS system, shown in Figure 3.4, that has proven to be effective for filtration of gadolinium-doped water. The required circulation rate scales as a function of the surface area of the detector divided by the detector volume, since most impurities are leached from the tank walls. Therefore, it is expected that the water treatment system will require a flow rate in the region of 35 tonne/h.

The water treatment system will use a molecular bandpass filtration system that is made up of several stages. The first stage includes an ultraviolet lamp and microfilters to kill and remove bacteria in the fluid. The next stage is a chiller to lower the water temperature. This has two purposes. Firstly, gadolinium dissolves more easily at lower temperatures but later stages in the filtration system such as pumps and filters will increase the water temperature, so this helps to avoid precipitation of the gadolinium. Secondly, bacterial growth will be slowed down in the detector and filtration system due to the lower temperature. After the chiller is an ultrafilter which is used to remove large impurities.

### 3.3. The NEO detector

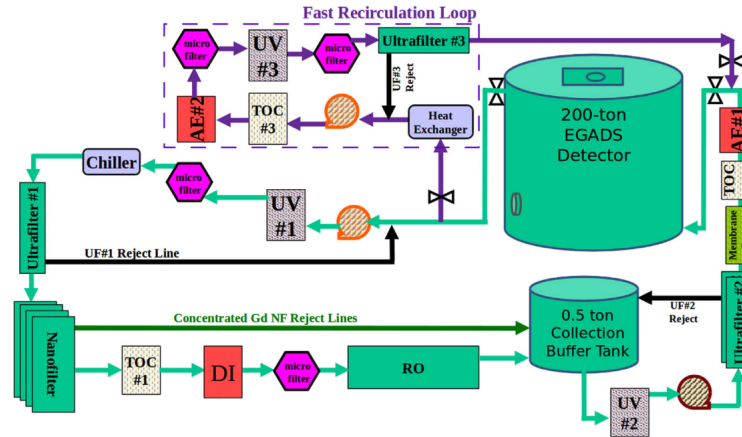


Figure 3.4: Schematic of the EGADS molecular bandpass filtration system. Image from [106].

Following on from this is a series of nanofilters. The reject line from the nanofilters contains gadolinium-rich water and the other is gadolinium-less to less than 1 ppm. The reject line goes into a buffer tank, whilst the gadolinium-less fluid goes through a standard water treatment system of deioniser, microfiltration, and then reverse osmosis to remove any other impurities in the water. The gadolinium and filtered water are then recombined in the buffer tank before undergoing further sterilisation by ultraviolet lamp, ultrafiltration and then re-entering the detector. A fast recirculation system also runs in parallel with the loop described previously. This system consists of a heat exchanger, microfilters, and ultraviolet treatment and does not remove any gadolinium.

An alternative system that could be used for the NEO detector is the resin ion-exchange system [124] that has been used in both Accelerator Neutrino Neutron Interaction Experiment (ANNIE) and Super-Kamiokande. This system uses ion-exchange resins in conjunction with ultraviolet lamps and microfilters to perform the filtration. A schematic of such a system is shown in Figure 3.5. By using a hydroxide resin, the filtration system will preferentially remove nitrates over sulphates, and as such impurities in the water and gadolinium-sulphate can be removed without a major impact on the gadolinium sulphate. Whilst the molecular bandpass system performs better in separating the gadolinium from the water with almost no loss of gadolinium, the resin ion-exchange system is much

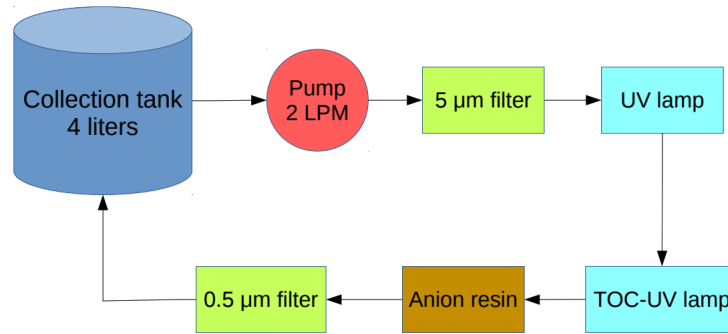


Figure 3.5: Schematic of the test setup for the resin ion-exchange filtration system at UC Davis. Image from [124].

cheaper and is much more suitable to being scaled up to the size required by NEO. The resin ion-exchange system also requires less power and uses less floor space. Overall, the resin ion-exchange system could be the favoured option as long as some degree of gadolinium loss is tolerable.

Since the performance of NEO will be sensitive to the water quality, the attenuation length will be monitored in real time. For this purpose, a system similar to the UDEAL device [125] currently in use at EGADS will be utilised. The UDEAL system is made up of a light source with seven possible wavelengths and a measurement arm containing detector fluid. An integrating sphere is located at both ends of the arm. The light is split into two paths, the first goes directly to the first integrating sphere which acts as an intensity monitor, and the second travels along the entire length of the arm into the second integrating sphere which measures the attenuation relative to the first. In this way, the wavelength-dependent attenuation length of the water can be continuously monitored to ensure optimal detector performance.

### 3.4 Chapter summary

- NEO is a proposed detector that would be located at the Boulby Underground Laboratory and observe the antineutrino emission from the Hartlepool Nuclear Power Station.

- Nonproliferation goals include verification of the reactor cycle at Hartlepool and different reactor discover scenarios.
- Possible scientific goals include measurements of neutrino oscillation parameters and detection of supernova neutrinos.
- Many designs are under consideration with a baseline conceptual design of a 20 m diameter and height right-cylinder of gadolinium-doped water, instrumented with roughly 4,000 photomultiplier tubes.
- A characterisation campaign for the Hamamatsu R7081-100 photomultiplier tube is underway.
- Calibration ports have been designed into the 316L stainless steel tank to support a wide array of calibration methods, based on both light injection and radioactive sources.
- Novel detection media, such as water-based liquid scintillator, and novel photosensors, such as LAPPDs and dichroicons, are under consideration.
- Water filtration systems based on those utilised by Super-Kamiokande and ANNIE are under consideration for purifying the gadolinium-doped water detection medium.

# Chapter 4

## Distant reactor detection

In late 2020, the WATCHMAN collaboration undertook a trade study to evaluate the performance of the conceptual NEO detector (outlined in chapter 3) in a wide variety of scenarios and use cases, both nonproliferation and science-focused. The study intended to create a path towards a final detector design, and aimed to optimise an array of design parameters including the tank size, photocoverage, veto region thickness and alternative light collection technologies. Included in this study was the first evaluation of the capabilities of the novel WbLS detection media for application in a large antineutrino detector.

Although a number of use cases were evaluated, this chapter will detail the distant reactor detection use case, which used the Heysham reactor complex as the target. This was considered a stretch goal for the detector, since the Hartlepool reactor complex was the primary target. Nevertheless, the Heysham complex was chosen to be studied as it may provide a viable contingency plan in the event that the Hartlepool reactor halted operations before the detector could collect enough data to make a measurement. This study assumes a scenario whereby the Heysham reactor complex is fully operational whilst the Hartlepool complex has been shut down.

The figure of merit for this study was the detector sensitivity to the reactor (defined as  $\sigma = S/\sqrt{S+B}$ , where  $S$  is the total signal counts and  $B$  is the total

## 4.1. Heysham nuclear reactor complex

---

background counts) within a time period of 365 days. It is assumed that all four reactor cores at Heysham are operating at their nominal power output (see section 4.1) throughout this time period, since no reliable information of the reactor power cycle was available at this time. The study was performed as a preliminary investigation to determine if the Heysham reactor complex could be considered as being in-scope for a detector design such as those considered for NEO. Follow-on studies have been performed which have sought to address many of the limitations and assumptions made in this study.

## 4.1 Heysham nuclear reactor complex

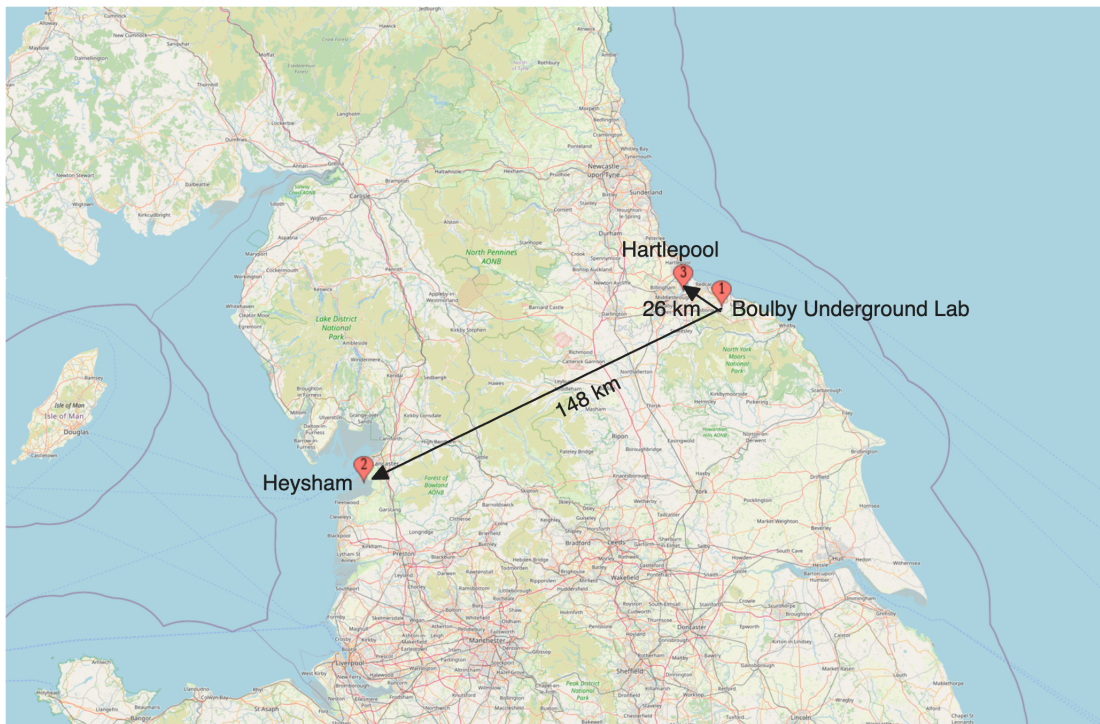


Figure 4.1: The locations of the Hartlepool and Heysham nuclear complexes with their respective distances from the detector at the Boulby Underground Laboratory.

The Heysham nuclear reactor complex (“Heysham”) is situated in Lancashire, 148km from the NEO deployment site at Boulby Underground Laboratory as can be seen in Figure 4.1. The Heysham site is made up of two separate power station

units, Heysham 1 and Heysham 2, both of which are advanced gas-cooled reactors (AGRs). Each power station contains two nuclear reactor cores, making a total of four reactor cores at Heysham. Each unit has a total thermal capacity of approximately 1,500 MWth, although they are rarely operated at this output.

## 4.2 Detector configurations

An array of parameters, including the tank diameter, the PMT support structure radius, the photocoverage and the detection medium were varied to produce a wide range of different detector configurations to be studied. This was to better understand the dependence of the overall ability of the detector to achieve different nonproliferation and physics goals on the detector design. The detector parameters that were altered were as follows:

- Tank diameters of 14 m, 16 m, 18 m, and 20 m were studied. The tank is a major cost driver and a larger tank allows for a larger buffer region between the tank wall and the PMT support structure.
- The PMT photocoverage was varied between 10%, 20%, and 40%. As PMTs are typically the single largest expenditure for many neutrino detectors, it is important to understand the lower limit of the photocoverage that can still achieve the goals of the experiment.
- PMT support structure radii of 5.7 m, 6.7 m, 7.7 m, 8.7 m were studied. This essentially defines the size of the possible fiducial volume. The number of PMTs required for each combination of photocoverage and support structure radius is shown in Table 4.1.
- Several detection media were compared,  $\text{H}_2\text{O}$ , 1%-loaded WbLS, and 3%-loaded WbLS.
- For each detection media, options with and without gadolinium were considered.

Almost all permutations of these parameters were considered, resulting in Monte Carlo generation for a total of 75 different detector configurations. With this matrix



Photocoverage	PMT support structure radius			
	5.7 m	6.7 m	7.7 m	8.7 m
10%	1,154	1,603	2,178	2,799
20%	2,379	3,302	4,338	5,585
40%	4,753	6,636	8,674	11,148

Table 4.1: The number of PMTs required for the given photocoverage and PMT support structure radius of each detector configuration.

of detector configurations, a decision could be made on the most suitable NEO detector design given budget constraints and the evaluated performance in a number of key metrics.

### 4.3 Simulation framework

The study was performed using the Reactor Analysis Tool (Plus Additional Codes) (RAT-PAC) [126] software package, which has been extensively modified and built upon by the WATCHMAN collaboration [127]. RAT-PAC is a tool that performs Monte Carlo simulation of the detector and integrates this with an event-based analysis loop. It is built on top of a number of other software packages. The Geant4 physics simulation framework [128–130] provides particle tracking and simulation of particle interactions. The CLHEP library [131] provides useful utilities such as random number generation, geometry functions, and linear algebra. PMT response and scintillation light production is modelled by the GLG4Sim (Generic Liquid Antineutrino Detector or GenericLAND) software package, a Geant4 extension designed specifically for neutrino physics derived from the simulation efforts of the KamLAND experiment. Output data is written to files that are analysed using the ROOT analysis framework [132].

All signal and background events are simulated in RAT-PAC which models the interactions, particle propagation and overall detector response. Events are initiated using event generators from GLG4Sim. These event generators define which particle should be produced, as well as the initial position, direction, and energy of the particle. These parameters can be sampled from probability distribution functions depending on the interaction being modelled. RAT-PAC

also allows for the implementation and simulation of the DAQ system, such that realistic models of triggering, electronics saturation, etc are included. An optical model of the PMTs is implemented using an adapted version of `GLG4PMTOpticalModel` from `GLG4Sim`. This parameterised model takes in a charge and timing distribution and produces a simulated PMT response drawing from these distributions as well as the wavelength dependent quantum efficiency value. These PMT hit times and charges are saved for each event and are used as the input to the event reconstruction described in subsection 4.4.2.

For each detector design under consideration, the detector geometry is implemented in RAT-PAC. A 1 m thickness region of rock surrounds the cavern; throughout this region is the starting point for the simulation of the fast neutron background. Inside the rock region is a volume of concrete that covers all sides except for the top of the tank. Inside the concrete is the tank and many complex structures such as the steel trusses and the PMT support structure. These structures are simplified for ease of design and simulation, but the bulk masses of these structures closely match the expected masses from the conceptual designs of the full detector.

The WbLS optical parameters represent the best understanding of the medium at the time of the study. As it is a detection medium under active development, these parameters are likely to change. The WbLS light yield is taken from measurements made at the CHESS experiment [113]. The emission spectra and time profiles were calculated using x-ray luminescence measurements made by Bourret et al. [133]. The scattering length is taken from preliminary measurements made at Brookhaven National Laboratory. Two different measurements are used, a nominal scattering length, and a conservative scattering length based on expected improvements and current measurements respectively. Background activities are assumed to be the same as for gadolinium-doped water due to the low scintillator concentration, although scintillators typically have a higher cleanliness ceiling than water.

Reactor core	Assumed power (MWth)	IBD rate (events/year/kton)
Heysham A-1	1362.0	18.7
Heysham A-2	1040.2	14.3
Heysham B-1	1087.0	14.9
Heysham B-2	1460.1	19.8
Total	4949.3	67.7

Table 4.2: Signal rates for each of the Heysham reactor cores at the proposed detector site. The assumed power is the mean power reported by the IAEA throughout the year 2018.

### 4.3.1 IBD signal and backgrounds

The primary signal to be observed in this study is the reactor antineutrino signal from the Heysham complex. All four Heysham cores were summed to form the signal. The IBD rates used in this study from each Heysham core are shown in Table 4.2, and the reactor antineutrino spectrum is shown in Figure 4.2. At the time of the study, no IBD event generator had been implemented in RAT-PAC. Therefore, the IBD signal is simulated as “singles”, with the positron and neutron simulated separately. These must then be correlated as part of the analysis process, which will be described in subsection 4.4.4.

However, there are other antineutrino sources that contribute to the background. The dominant source of IBD background is the overall contribution from all other reactors in the world, known as the “world reactor” background. Reactor antineutrino rates (both the Heysham signal and world reactor background) are provided by [geoneutrinos.org](http://geoneutrinos.org) [64], which uses Equation 2.5 summed over each fissioning isotope, with assumed fission fractions based on the type of the reactor. The distance from every reactor to the detector site at Boulby is also used to apply the electron antineutrino survival probability to account for neutrino oscillation which modulates the visible antineutrino spectrum. The positron energy spectrum is then produced by shifting the antineutrino energy spectrum to account for the threshold energy of the IBD interaction. The reactor power values used to calculate the IBD rates is the average power for each reactor core provided by the IAEA for the year 2018, which implicitly accounts for the nominal reactor fuel cycles, but this did result in the assumption that there were

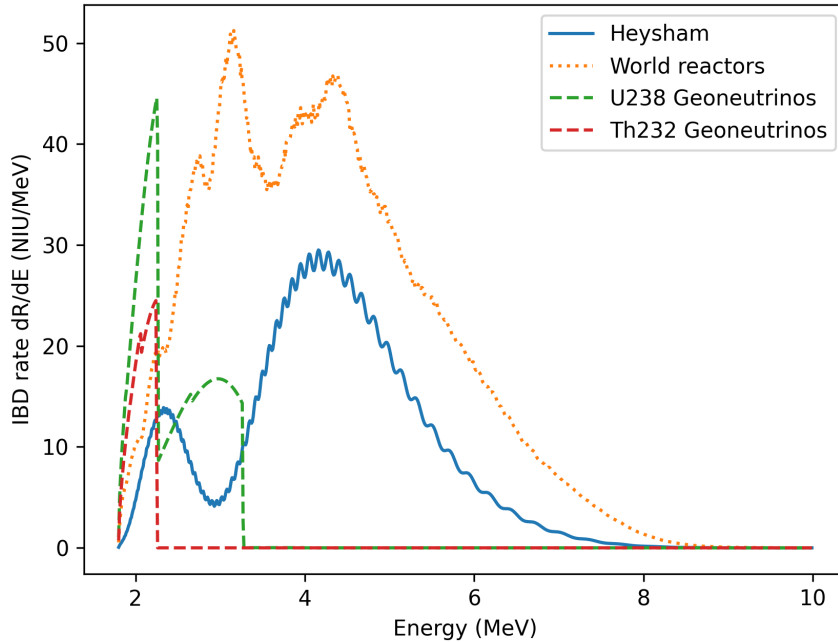


Figure 4.2: Detectable IBD spectrum at Boulby for the combined Heysham reactor complex, all other world reactors (excluding Hartlepool), and geoneutrinos. The energy spectrum is given in Neutrino Interaction Units (NIU) which corresponds to 1 interaction per  $10^{32}$  free protons per year. Spectral information generated from [64] with the effects of neutrino oscillation included.

no unforeseen reactor shutdowns for any of the reactors that were modelled. Furthermore, the Hartlepool reactors were assumed to have been turned off permanently, and so their antineutrino production does not contribute to the world reactor background.

Another notable source of IBD background is geological electron antineutrinos, known as geoneutrinos. These are produced from the beta decay of  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chain elements as well as  $^{40}\text{K}$  present in the crust and mantle of the Earth. Estimates of the geoneutrino energy spectrum at Boulby are also taken from [geoneutrinos.org](http://geoneutrinos.org), which are calculated by summing across all the decays in the relevant decay chain taking into account the branching ratio. The energy spectrum for all modelled IBD signal and backgrounds can be seen in Figure 4.2.

### 4.3. Simulation framework

---

Background component		Activity
PMT	$^{238}\text{U}$	0.043 ppm
	$^{232}\text{Th}$	0.134 ppm
	$^{40}\text{K}$	0.004 ppm
Liquid	$^{238}\text{U}$	$1 \times 10^{-6} \text{ Bq kg}^{-1}$
	$^{232}\text{Th}$	$1 \times 10^{-7} \text{ Bq kg}^{-1}$
Cosmogenic radionuclide	$^9\text{Li}$	$1.258 \times 10^{-6} \text{ Hz/kton}$
	$^{17}\text{N}$	$1.265 \times 10^{-6} \text{ Hz/kton}$
Antineutrino	Geoneutrinos	$3.51 \times 10^1 \text{ events/year/kton}$
	World reactors	$1.52 \times 10^2 \text{ events/year/kton}$

Table 4.3: Rates of each background component modelled in the study. “Liquid” refers to both water and WbLS and are assumed to have the same inherent radiopurity.

#### 4.3.2 Radiological and cosmogenic backgrounds

Table 4.3 shows the backgrounds that were considered in this study as well as their associated rates. These baseline rates were used to scale the absolute rates for each detector configuration depending on the number of PMTs or tank size.

Background contributions were only considered from the detection medium and PMT detector regions, as well as cosmogenic-induced backgrounds. Contributions from radioactive contamination of the steel structures (mostly  $^{137}\text{Cs}$  and  $^{60}\text{Co}$ ) were not included as previous studies had shown that, at the rates expected from spectroscopy measurements of 316L stainless steel, very few events will ever reconstruct inside the fiducial volume [134], and simulation of these events was computationally wasteful.

For radiological backgrounds, all beta-emitters from the  $^{238}\text{U}$ ,  $^{232}\text{Th}$ , and  $^{222}\text{Rn}$  decay chains were simulated, excluding those with very small branching ratios ( $< 0.1\%$ ) or those with endpoint energies below 0.5 MeV, which are below any realistic analysis threshold that NEO may achieve. These are simulated throughout the detector medium and PMT volumes. Alpha emitters were excluded. In water, alpha particle will not cross the Cherenkov threshold, and in WbLS alpha particles will cause strong ionisation quenching of the scintillator due to the high specific energy loss (dE/dx) of alpha particles relative to betas and gammas. It is assumed

that quenching will push these alpha interactions below the analysis threshold.

Two radionuclides produced from cosmogenic muon spallation on oxygen in the detector medium are also included in simulations,  ${}^9\text{Li}$  and  ${}^{17}\text{N}$ . These are included as they undergo beta-neutron decay with a beta endpoint energy that can mimic a true IBD event, and their relatively long half-lives (hundreds of milliseconds to seconds) make vetoing these decays challenging. It is assumed that the bright Cherenkov signals produced from cosmogenic muons travelling through the detector medium can be vetoed, but these radioisotopes are long-lived enough to not be cut out by a standard veto around the path of the muon. Other isotopes that are typically produced in spallation interactions on oxygen are omitted. These isotopes either have such a low yield that their impact on the total background is negligible (as in the case of  ${}^8\text{He}$  [135]), or have a short enough lifetime that they will be excluded using any reasonable veto scheme. A muon rate of 0.15 Hz is used based on the muon flux at the depth of the lab. This value is based on the nominal 20 m tank and was scaled for alternate tank sizes.

Accurate simulations of the cosmogenic-induced fast neutron background were not immediately available, and so were omitted for the initial study. The fast neutron background was later calculated by sampling the energy of fast neutrons produced in the rock around the detector using the Mei and Hime model of cosmogenic muon fluxes [136], which has been validated against measured cosmogenic muon fluxes at a number of underground laboratories, including Boulby. Interactions from these fast neutrons in the detector volume itself were simulated using FLUKA [137, 138] due to its more advanced neutron interaction and transport libraries. The particle information, including energies and positions, were recorded. These were then used as inputs into the full detector Geant4 simulation to record the detector response to these events. This simulation and analysis technique for fast neutron backgrounds was validated using the WATCHBOY detector [139]. For gadolinium-doped water configurations, only the di-neutron capture rate is relevant, as this can mimic the IBD double flash signature. For WbLS, cosmogenic fast neutrons can scatter on a proton and cause a recoil which can mimic the prompt positron signal in a scintillator, but this effect will be invisible in a water medium as the proton recoil is below the Cherenkov

threshold.

No systematic uncertainties on the background were included in this study, and so only statistical uncertainties are reported. Later studies have built on these results with a full systematic treatment of all backgrounds, as will be discussed in section 4.7. Instead, two parallel analysis paths were used to provide a cross-check on one another. The maximum difference between the two analysis paths, around 10%, is used as a conservative uncertainty on all the quoted sensitivity values to account for systematics on the methods. See section 4.4 for more detail on the two parallel analyses.

## 4.4 Analysis methodology

### 4.4.1 Framework

The WATCHMAKERS framework is used to manage the generation of simulation macros, the submission of batch jobs to computing clusters, and to run analysis on the resulting files. Using this framework, there is a consistent series of steps to perform the sensitivity analysis consisting of simulation, event reconstruction, application of cuts to define a region of interest, and sensitivity evaluation. These steps are detailed in the following subsections.

### 4.4.2 Event reconstruction

Vertex reconstruction in the detector is performed using BONSAI [140], a low energy vertex reconstruction tool designed for water Cherenkov detectors. It was originally used in Super-Kamiokande but has been modified for use by the WATCHMAN collaboration [141].

BONSAI is a time-of-flight fitter, meaning that it uses the PMT hit times to attempt to determine the interaction vertex. BONSAI calculates the hit timing residuals — the difference between a PMT’s hit time and the time at which the photon was emitted — taking into account the photon time-of-flight. For a given

test vertex, the hit time residual at the  $i$ th PMT is given by,

$$\delta t_i(x) = t_i - t_{\text{tof}} - t_0, \quad (4.1)$$

where  $x$  is the test vertex,  $t_i$  is the hit time of the  $i$ th PMT,  $t_0$  is the emission time, and  $t_{\text{tof}} = |x_i - x|/(c/n_{\text{water}})$  is the time-of-flight from the test vertex to the  $i$ th PMT. For a given detector configuration, a probability distribution function of the hit timing residuals is generated using simulation or, once a detector is constructed, through the calibration methods described in subsection 3.3.4. Figure 4.3 shows an example distribution of the hit timing residuals. These distributions are the result of many different detector effects, predominantly the PMT timing effects such as late pulsing and afterpulsing, but also optical effects in the detection medium such as scattering and, in the case of scintillating media, absorption and re-emission.




Figure 4.3: An example of the hit time residuals for a WbLS-configured NEO detector. Hits extend out over hundreds of nanoseconds as a result of the delayed production of light by the scintillator.

At least four PMT hits are required for BONSAI to reconstruct an event. If more than four hits are present in an event, the hits are grouped into collections of four hits and an initial test vertex is generated for each group. This helps to prevent the iterative search from getting “stuck” in areas of local maxima by spreading out the initial test vertices throughout the detector.



BONSAI then iterates through the series of test vertices and performs the likelihood maximisation, generating a dodecahedron-shaped grid of additional test vertices around each original test vertex. This shape of grid allows the algorithm to efficiently scan large regions of the detector for the most optimal interaction vertex. Once BONSAI has concluded its grid search, it returns the highest likelihood vertex along with a “goodness-of-fit” metric that defines how well the vertex matches the expected hit timing residuals distribution for that event. This goodness-of-fit parameter can be used as a data quality cut to ensure that only well-reconstructed events are analysed, and events that are known to be poorly reconstructed, likely due to being near the edge of the fiducial volume, are discarded.

### 4.4.3 Energy estimation

An energy reconstruction algorithm was not implemented in this work. Instead, the number of hit PMTs within a given time period is used as an analogue to the energy. At the energy of reactor antineutrinos in large water Cherenkov detectors like NEO, it is reasonable to assume that each PMT will observe a maximum of a single hit per event. Since the number of produced Cherenkov photons is approximately proportional to the positron energy, and the positron energy is linearly proportional to the incident antineutrino energy, the assumption that the number of hit PMTs represents the antineutrino energy will hold true for most antineutrino events in the detector.

For most configurations of the NEO detector, the number of PMT hits with a hit time residual of less than 9 ns,  $n_9$ , is used as the energy analogue. The 9 ns time window spans from 3 ns before the calculated interaction timestamp to 6 ns afterwards. This narrow time window is used to ensure that most PMT hits are from unscattered light, and to restrict the number of dark pulses that contaminate the hit selection. Given the size of the NEO detector and the water detection medium where most light is produced instantaneously, 9 ns is an optimal time window. However, for scintillating media, where much of the light production is delayed relative to the initial interaction, the window may need to be extended to exploit the improvements in position and energy resolution that result from

collecting more light. For the WbLS cocktails considered in this study, a time window of 100 ns, *n100*, made up of 10 ns pre- and 90 ns post-interaction, was used as the energy analogue. Previous studies had determined that *n100* provided the optimal energy resolution as it included the majority of the unscattered light resulting from the delayed emission from the scintillator whilst excluding most scattered light that has a negative impact on vertex resolution [142].

#### 4.4.4 Region of interest optimisation

The sensitivity to the reactor signal in the presence of all modelled backgrounds is evaluated using a counting analysis. A region of interest in the detector is created by applying a number of cuts on the data selection intended to remove backgrounds. These cuts are:

- The distance from the event to the nearest PMT i.e. the fiducial cut
- The energy analogue threshold (*n9* or *n100*) for the prompt positron
- The energy analogue threshold (*n9* or *n100*) for the delayed neutron
- The interaction vertex position goodness
- The maximum time between pairs of events
- The maximum distance between pairs of events

The maximum time and distance cuts are fixed due to the method of simulating all events as singles. The fiducial cut and prompt and delayed cuts are iterated through for every component, producing two-dimensional efficiency maps like those seen in Figure 4.4 showing the detection efficiency as a function of the cut values. An underlying assumption of this analysis method is that the observed reactor antineutrino signal rate ( $R_{\text{coincident}}$ ), consisting of IBD interactions, can be estimated as,

$$R_{\text{coincident}} = R_{\text{IBD}} \times \epsilon_{e^+} \times \epsilon_n \times \epsilon_{\Delta t} \times \epsilon_d, \quad (4.2)$$

where  $R_{\text{IBD}}$  is the true IBD rate,  $\epsilon_{e^+}$  ( $\epsilon_n$ ) is the fraction of positron (neutron) events that pass the fiducial and prompt (delayed) energy threshold cut respectively,  $\epsilon_{\Delta t}$  and  $\epsilon_d$  are the fraction of events that pass the cuts on the maximum time and

distance between the positron and neutron components of the IBD interaction and are not optimised. Furthermore, the contamination of accidental coincidences is determined based on the singles rates of all modelled backgrounds. The accidental coincidence rate ( $R_{acc}$ ) is estimated by,


$$R_{acc} = R_p \times R_d \times \Delta t \times \epsilon_d, \quad (4.3)$$

where  $R_p$  and  $R_d$  are the rate of background events that pass the fiducial cut and the prompt and delayed energy threshold cuts respectively,  $\Delta t$  is the medium-dependent coincidence time window and  $\epsilon_d$  is the fraction of the total accidental rate remaining after the distance cut of 2 m,  $d$ . The fraction of accidental contamination that passes these cuts is recalculated and applied in the analysis for each detector configuration.


The prompt and delayed energy threshold, as well as the fiducial cut, are optimised through iterative evaluation of the signal to background ratio. The signal and background counts are evaluated at the optimal data selection cuts using Equation 4.2 and Equation 4.3 in conjunction with the efficiency maps in Figure 4.4, and the metric  $\sigma = S/\sqrt{S+B}$  is calculated for 365 days of observation time.

#### 4.4. Analysis methodology


---




(a)  $^{210}\text{Tl}$  background originating in the PMTs with water detection medium.




(b)  $^{210}\text{Tl}$  background originating in the PMTs with WbLS detection medium.




(c)  $^{210}\text{Tl}$  background originating in the fluid with water detection medium.



(d)  $^{210}\text{Tl}$  background originating in the fluid with WbLS detection medium.



(e) Heysham antineutrino signal with water detection medium.



(f) Heysham antineutrino signal with WbLS detection medium.

Figure 4.4: Example efficiency maps showing the effect of ROI optimisation on detection efficiency for each component. (a-d) show non-correlated backgrounds resulting from radiological contamination in detector subsystems, (e-f) show the antineutrino signal from Heysham.

## 4.5 Parallel analysis

In order to investigate methodological effects on detector sensitivity, a second analysis was also conducted in parallel. This was used primarily to study alternative detection media, since the baseline analysis was developed for gadolinium-doped water. The parallel analysis used the CHESSE/THEIA version of RAT-PAC to perform event simulation. The Path Fitter hit-time based reconstruction algorithm, originally used in SNO, was used for event reconstruction. Unlike the baseline analysis, the parallel analysis combined all backgrounds into a merged dataset which intersperses events in time, allowing for realistic evaluation of the accidental contamination. This allowed for optimisation of the region of interest using dynamic cuts on the coincidence time and distance, as well as the cuts used in the baseline analysis on the prompt and delayed energy thresholds, and the distance from the PMTs. Results between the two analysis were found to agree to within 10-20%, with the parallel analysis consistently outperforming the baseline analysis, likely as a result of the improved methodology for optimising the region of interest. Sensitivity values for gadolinium-doped water and 3% WbLS are from the baseline analysis, whereas results from 1% WbLS, 1% Gd-WbLS, and 3% Gd-WbLS are from the parallel analysis but conservatively normalised to the baseline analysis performance to enable easier comparison of detection media.

## 4.6 Results

The results of the distant reactor study as well as the studies of other use cases were detailed in the WATCHMAN Tank Size Trade Study Report [143]. The signal significance is reported for each detector configuration that was considered. It was found that many of the potential detector designs could achieve the goal of  $3\sigma$  signal significance within 365 days of reactor-on time. As is shown in Figure 4.5, sensitivity is improved as the photocoverage or tank diameter is increased. Perhaps most crucially, the results highlight the importance of gadolinium doping over increasing the light yield of the detection medium for low-statistics analyses such as this study. The IBD detection efficiency is improved due to the increased




Figure 4.5: Detector sensitivity to Heysham reactor discovery depending on the PMT coverage, and detection medium, as a function of PMT radius, an analogue for the fiducial volume radius. The dashed horizontal line shows the  $3\sigma$  sensitivity within 365 days target.

light yield, but this is counteracted by a larger relative increase in the background rate as a result of the increased capture time on hydrogen relative to gadolinium leading to a larger contamination of accidentals. Furthermore, the shape of the signal and background spectra may work against WbLS in this case, since the antineutrino signal from Heysham remains appreciable above the 1.8 MeV IBD threshold, whilst many backgrounds are concentrated at lower energies, nullifying much of the possible gains from the removal of the Cherenkov threshold for WbLS fills. The sensitivity as a function of the detection medium is shown in Figure 4.6.

Figure 4.7 shows the sensitivity with the inclusion of fast neutron backgrounds as a function of the tank diameter (where the tank diameter is given by the PMT radius plus the size of the buffer region) for gadolinium-doped water detection medium. Larger buffer regions are expected to reduce the impact of the fast




Figure 4.6: Detector sensitivity as a function of the PMT radius depending on the detection medium at a photocoverage of 20%. The dashed horizontal line shows the  $3\sigma$  sensitivity within 365 days target.

neutron background on the sensitivity, however this comes at the expense of a smaller fiducial volume. Larger tank diameters allow for greater potential to reduce fast neutron backgrounds, but this is obviously more expensive. The results show that the reduction in the fast neutron background from larger buffer volumes is not greater than the increased signal resulting from retaining the largest possible fiducial volume. The overall reduction in sensitivity from the inclusion of fast neutron backgrounds is around 30%, showing that much more in-depth modelling of cosmogenic-induced backgrounds, as well as more advanced region of interest optimisation such as asymmetric fiducial volume definitions, should be considered in the future. With a WbLS detector medium, the fast neutron background could potentially be further reduced by applying particle identification techniques that use the measured ratio of Cherenkov to scintillation light or the time profile of the observed scintillation light to discriminate these events from IBD interactions.




Figure 4.7: Detector sensitivity to Heysham reactor discovery with a gadolinium-doped water fill depending on the PMT coverage, and buffer radius, as a function of tank diameter. The fast neutron component of the background is included. The dashed horizontal line shows the  $3\sigma$  sensitivity within 365 days target.

$3\sigma$  sensitivity within 365 days of reactor-on time is achievable for tank diameters of 18 m or larger, and at 20% photocoverage or larger for all considered detection media with the exception of non-gadolinium-doped 1% WbLS. If the photocoverage were to be limited to 10%, then a WbLS detector fill is required to achieve the sensitivity target. If tank sizes smaller than 18 m were to be chosen, then gadolinium-loaded WbLS mixtures with higher scintillator concentrations, or photocoverages greater than 40% would be required otherwise  $3\sigma$  sensitivity is not possible. These recommendations produce numerous options for the final detector design depending on the priorities of the different physics and nonproliferation goals, as well as the relative costs or savings of up/downgrading various detector systems (e.g. the cost of increasing the tank size compared to increasing the photocoverage). Ultimately these decisions will decide the goals that the initial



NEO detector will be able to achieve.

## 4.7 Later work

This preliminary study has limitations resulting from assumptions or simplifications that were made to meet time constraints. Firstly, since the fast neutron backgrounds were not immediately available, they were added into the signal significance calculation without a full reoptimisation of the region of interest. It is likely that performing the optimisation with the fast neutron backgrounds included from the start will have resulted in different fiducial volume cuts, which may have provided improved discrimination of fast neutron event and hence better signal-to-background values. Secondly, no systematic uncertainties on the backgrounds were included. This could have a large effect, since the systematic uncertainties on some background components, especially the cosmogenic-induced backgrounds, can be large due to limitations of the muon flux models and their application to the underground lab at Boulby. Thirdly, the lack of an IBD event generator required that all events were simulated as singles and then correlated later in the analysis using fixed distance and timing cuts as well as assumptions on the efficiency of these cuts.

These limitations were all addressed in a follow-on study [144] into the sensitivity to distant reactors that expanded the scope to also include the next closest reactor complex, Torness. This study developed two improved methodologies, one based on machine learning to reduce fast neutron backgrounds and a likelihood analysis to reduce uncorrelated radiogenic backgrounds, and another based on coincidence reconstruction of the IBD event. This study included fast neutron simulations from the beginning, and included systematic uncertainties on all backgrounds. In addition, an IBD event generator was used, which simulated the positron and neutron in the same event allowing for realistic coincidence cuts to be made on the distance and time between the particles. Different detector designs were considered in the follow-on study, and so the results are not directly comparable to the results presented here, but agree reasonably well with each other. In the follow-on study, a 22m tank with

gadolinium-doped water fill, with a 9 m PMT radius and 15% photocoverage achieved  $3\sigma$  sensitivity within 241 days. This can be compared to the 295 days required to achieve  $3\sigma$  sensitivity in this study for a 20 m tank with gadolinium-doped water fill, with a 7.6 m PMT radius and 20% photocoverage.

Overall, it is clear that a large water-based detector has the potential to be used to detect undeclared reactors at large distances. Further studies have investigated more specific use-cases that may be of interest to the nonproliferation community, such as the size required for a detector to exclude the existence of reactors within a given distance, as well as the ability to detect an undeclared reactor operating in the presence of a declared reactor [145].

## 4.8 Conclusions

This chapter detailed a first look at distant reactor detection using traditional gadolinium-doped water technology as well as the novel WbLS detection medium. It was seen that detection of the Heysham reactor could be considered as being in-scope for the NEO detector for many of the detector configurations under consideration, although care should be taken with the limitations of the study that have been detailed and more study is required to confirm this.

Since 2022, development of a large-scale reactor antineutrino demonstrator was suspended due to uncertainties related to the shutdown schedule for several key nuclear reactors including the Hartlepool and Heysham sites. Efforts to continue the project and build a large-scale demonstrator using a nuclear reactor and underground mine in the USA are ongoing. However, smaller-scale research into new technologies that can be applied to antineutrino reactor monitoring is continuing. One of these successor projects stemming from the WATCHMAN collaboration is the Boulby Underground Technology Testbed Observing Neutrinos (BUTTON) experiment. These projects aim to de-risk new antineutrino technologies and build towards a large-scale deployment of a detector for demonstrating an application of antineutrinos for nonproliferation in the future.

### 4.8.1 BUTTON

BUTTON is intended to be a testbed for antineutrino detection technologies. It will be located at the Boulby Underground Laboratory. At the time of writing, the design is still being finalised, but it is expected that the BUTTON detector will be a cylinder of approximately 1.7 m in radius and around 3.1 m in height, resulting in a total volume of around 30 tonnes. This volume will initially be filled with gadolinium-doped water, although it is expected that this will be replaced by WbLS once the technology is mature enough for a larger deployment. The detector will be instrumented with 96 PMTs placed as close to the outside of the cylinder as possible to maximise the target volume. An image of the current design of the BUTTON detector is shown in Figure 4.8. Due to its size, it will be constructed in the main experimental hall at Boulby and as such will not require any major engineering works as NEO would have done.

A number of ideas have been proposed for BUTTON's physics programme. Firstly, BUTTON could make measurements of the cosmogenic muon-induced neutron backgrounds at the lab, which have historically been one of the largest systematic uncertainties when evaluating the sensitivity of the large-scale water Cherenkov detectors for the WATCHMAN collaboration. In addition, BUTTON expects to become a testbed for characterising advanced photosensors which could be used in larger detectors in the future. Amongst these photosensors include silicon photomultipliers coupled with wavelength-shifting plates, which could allow for cost-effective large-area instrumentation. One of the most promising advanced photosensor options is the LAPPD, mentioned briefly in subsection 2.6.4 and subsection 3.3.6, due to its large photosensitive area, and impressive time and position resolution. Together with a deployment of WbLS, LAPPDs in BUTTON could demonstrate Cherenkov-scintillation separation, a major milestone in the development of hybrid detection media which are promising options for future neutrino detectors. The characterisation and application of LAPPDs for BUTTON-like neutrino experiments are the subject of the subsequent chapters of this thesis.

There are already proposals to build a larger BUTTON-like detector, perhaps




Figure 4.8: A schematic of the simulated geometry of the BUTTON detector shown in Geant4. The support structure is arranged in segments each holding four PMTs.

100 to 400 tonnes of liquid volume, to demonstrate the scalability of the technology that will be developed at BUTTON. These detectors, along with other WATCHMAN successor experiments focussed on advancing novel neutrino detection technologies, hope to guide a path towards a deployment of a large, kilotonne-scale antineutrino detector for nonproliferation purposes in the future.

## 4.9 Chapter summary

- A Trade Study was undertaken by the WATCHMAN collaboration to evaluate the performance of numerous potential designs for Neutrino Experiment One.

- Many nonproliferation and science objectives were studied, including an investigation of NEO's capabilities for discovering nuclear reactors at much larger distances such as the Heysham Nuclear Power Station.
- A matrix of detector designs with varying tank sizes, PMT coverages, and detection media was created.
- The IBD signal, as well as many predicted backgrounds, were simulated using RAT-PAC.
- Vertex and energy reconstruction was performed, and a region of interest was defined based on energy cuts on the prompt and delayed events, as well as an IBD coincidence cut.
- Results showed that a  $3\sigma$  discovery of the Heysham reactors was possible within 365 days for designs with gadolinium-doped detection media and higher photocoverages. Many smaller tank designs or lower photocoverage designs could not achieve this goal.
- Follow-on studies have taken place and addressed many of the shortcomings of this initial study.
- A large-scale antineutrino detector demonstration will not take place at Boulby Underground Laboratory in the near future. However, a smaller detector, BUTTON, will be constructed and used as a testbed for new antineutrino detection technologies such as WbLS and advanced photosensors.

## Chapter 5

# Large Area Picosecond PhotoDetectors

The LAPPD is a novel imaging photosensor technology that is under strong consideration for use in neutrino experiments. Coupled with new hybrid liquid detection media, such as WbLS that has been discussed in chapter 4, these photosensors could provide greatly improved vertex resolution during event reconstruction and the possibility to perform Cherenkov-scintillation separation as a result of their impressive timing capabilities. There is a great amount of interest in LAPPDs across many areas of high energy physics, and significant effort is being invested to overcome the challenges associated with deploying them into neutrino experiments.

The subsequent chapters will provide an overview of the LAPPD, and detail the work that has been performed to characterise and evaluate the capabilities of LAPPDs in the context of water-based Cherenkov detectors.

This chapter first gives an overview of the most basic photomultiplier, the photomultiplier tube, before summarising microchannel plates, the photomultiplication technology on which the LAPPD is built. Subsequently, it provides an overview of the development of the LAPPD, its construction, its operation, and the challenges associated with utilising LAPPDs in water

Cherenkov neutrino experiments. In order to develop further understanding of LAPPDs and their operation, a visit to Incom was made to see first-hand how LAPPDs were constructed and tested by the manufacturer. Additionally, the author joined the on-site team for the ANNIE experiment at Fermilab for six months to get hands-on experience with LAPPD characterisation and deployment.

## 5.1 Photomultiplier tubes

Before discussing LAPPDs, it is important to understand the basic photodetection technologies used in neutrino experiments. Conventional photomultiplier tubes (PMTs) are the workhorse photon detector of many particle physics experiments. The role of any photomultiplier is to convert small numbers of photons into observable electronic signals. A PMT, shown in Figure 5.1, consists of a vacuum tube and a glass window coated with a photocathode - a material that has a low work function. Upon being struck by a photon, the photocathode emits an electron, known as a photoelectron, through the photoelectric effect. The probability of producing a photoelectron in response to an incident photon is known as the quantum efficiency. The quantum efficiency is dependent on the photon wavelength and the composition of the photocathode, and typically peaks at around 15-30% for typical photocathode materials sensitive to visible light [146].

In order to generate an observable electrical current, the single photoelectron must be multiplied by several orders of magnitude. This signal amplification is achieved using a dynode chain, a series of electrodes made from secondary-emitting material (material that, upon being struck by an electron, will eject several more electrons). Beginning at the photocathode, each dynode is held at a gradually increasing positive voltage to accelerate the electrons along the vacuum tube towards an anode. The initial photoelectron is focussed onto the first dynode, then the electrons produced at the first dynode are focussed onto the second dynode, and so on. Each stage produces more and more electrons until they are eventually collected at the anode. A typical photomultiplier might have between five and twenty dynodes in the chain, and each dynode will produce a few electrons for each incident electron depending on the secondary electron yield of the material and the kinetic energy

## 5.1. Photomultiplier tubes

---




Figure 5.1: Schematic demonstrating the structure and basic operating principle of a photomultiplier tube. Image from [147].

of the incident electron. In this way, a single photoelectron can be multiplied such that there can be as many as  $10^8$  electrons collected at the anode, producing an observable electrical signal even for single photons. This value, the ratio of the final number of electrons compared to the initial number of photoelectrons, is known as the gain, and is a key parameter for any photomultiplier. The amplification process also has an efficiency factor, known as the collection efficiency, which is the ratio of pulses that are observed at the anode to the number of photoelectrons produced at the photocathode. This is mostly dependent on the fraction of photoelectrons that are successfully focussed onto the first dynode. This collection efficiency, combined with the quantum efficiency, produce the overall likelihood that a photomultiplier will produce an electrical signal when struck with a photon.

PMTs are incredibly mature instruments, having first been invented in the 1930s and steadily developed and improved upon ever since [148]. Nowadays, PMTs are available in many different shapes and sizes, with varying gains, timing capabilities, photocathode materials (and as such sensitivity to different regions of the electromagnetic spectrum), tolerance to magnetic fields, and so on.

The PMT is an example of a discrete dynode photodetector. These have a single dynode path that is followed by every produced photoelectron. The PMT acts as essentially one single large pixel, because there is no difference in the



observed electrical signal depending on where on the photocathode the photon hit (although there may be efficiency effects, for example an angular dependence on the collection efficiency). Therefore, these PMTs do not have any positional resolution. Multianode PMTs do exist [149] which, as the name suggests, feature multiple discrete dynode paths and multiple anodes to collect the electrons on each path. By looking at which anode produced the signal, it is possible to determine where the photon struck on the surface, typically with only a coarse level of resolution.

## 5.2 Microchannel plates

In contrast to the discrete dynode design of most conventional PMTs, a microchannel plate (MCP) is an example of a continuous dynode photomultiplier. This is the photodetection technology utilised by the LAPPD.

MCPs were originally developed primarily for military purposes as the component responsible for electron multiplication in image intensifier tubes — the operational element in night-vision goggles [151, 152]. An MCP, shown in Figure 5.2, is made up of a substrate, typically fabricated from glass, that contains an array of many tiny capillaries arranged parallel to one another. These capillaries, known as microchannels, are generally 15 to 50  $\mu\text{m}$  in diameter [153] and typically have a length to diameter ratio of between 40 and 100 [154]. The microchannels are made from, or coated with, a secondary emissive material such that each microchannel acts as an independent electron multiplier, resulting in a typical gain of around  $10^3$  [152]. The microchannels are not arranged perpendicular to the face of the plate, but rather are tilted off-axis by a small angle of around  $8^\circ$  [154]. This ensures that the electron actually collides with the microchannel rather than passing straight through. In addition, the electron is more like to collide at an angle relative to the microchannel wall whereby it reaches the optimal depth into the microchannel wall such that the produced secondary electrons are able to escape the material. This optimum angle is in the region of 8 to  $15^\circ$  [155]. Tilting the MCP also has the effect of reducing ion feedback, explained in more detail in section 5.3. Since a single MCP only

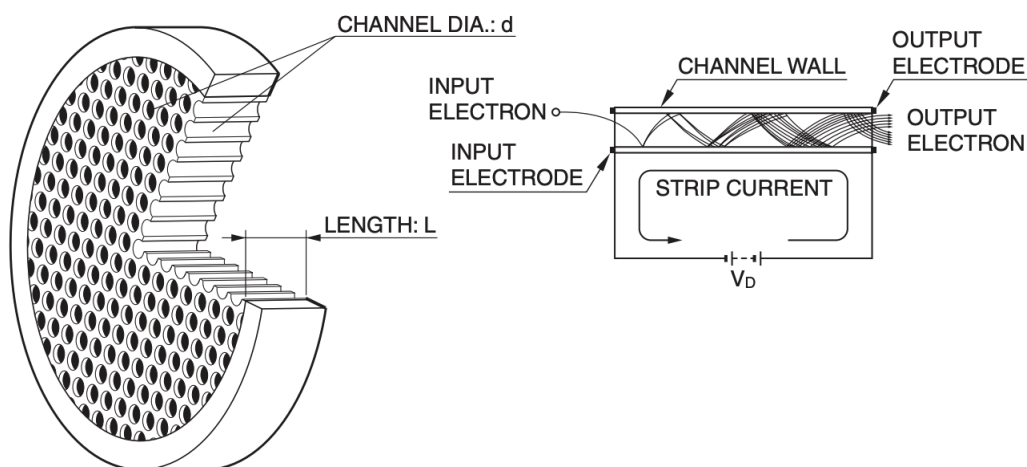


Figure 5.2: Left: A cutaway schematic of the structure of a microchannel plate. Right: A schematic of a single microchannel demonstrating how electrons are multiplied. Image from [150].

produces a gain in the region of  $10^3$ , MCPs are commonly stacked together to increase the gain further to approximately  $10^6$ . When MCPs are coupled together, they are often arranged in a “chevron pair” where the tilt of the microchannels from the surface normal is inverted for each MCP to ensure optimal multiplication of the cloud of electrons that exits the first MCP. The electrons are collected at the bottom of the MCP stack at an anode, the design of which is flexible and can be selected based on the needs of the experiment.

The compact nature of MCPs and the high electric field results in excellent timing properties suitable for ultrafast timing measurements. The overall transit time is often on the order of a few nanoseconds or less and fluctuations in the transit time can be as low as tens of picoseconds for single photons [152]. The rise and fall time of pulses is also rapid which is beneficial when operating in high rate environments.

The spatial resolution of an MCP-based detector is the product of multiple factors. Since the MCP confines the initial cloud of electrons within a single microchannel, it is possible for MCP-based detectors to preserve to a large degree the spatial information of the incident photon. The electron cloud produced within

a single microchannel will spread out when it reaches the exit of the microchannel. If the cloud then enters a second MCP, electrons will enter into multiple microchannels, each of which will produce its own electron cloud that will then disperse once it leaves the microchannel. Therefore, the potential spatial resolution of a given stack of MCPs will predominantly depend on the pore diameter, the pore spacing, distance between each stage of the MCP stack, and the voltage applied between each MCP in the stack. This feature of MCP-based photosensors means that the position of the incoming photon can be inferred from the location of the charge cloud at the anode. The spatial resolution of the full detector will inevitably depend strongly on the design chosen for the anode, as well as the readout electronics. It is often in the interest of the user to allow the electron cloud to spread rather than placing the anode immediately below an MCP since analysis techniques such as centroiding allow for spatial resolution better than the smallest readout element on the anode. However, this needs to be balanced against the fact that as the charge cloud spreads wider, less charge will be deposited onto each readout element and reducing the signal-to-noise ratio.

MCP-PMTs typically have improved tolerance to magnetic fields [156, 157] over conventional PMTs, which can show reduced collection efficiency and gain due to the impact of the magnetic field on electron trajectories at the first dynode and dynode chain respectively [158]. Many larger PMTs require mu-metal shielding even to fields as weak as the Earth's magnetic field, and so operating such a PMT in any artificial magnetic field will have an obvious impact on performance. Due to the compact form factor of MCP-PMTs, the electron trajectories are far shorter and thus less impacted by magnetic fields. They have short distance at the crucial points of the electron trajectories: the photocathode to dynode gap, and the dynode to anode gap. Furthermore, a very high voltage is typically applied (up to around 3,000 V), providing the electrons with increased energy relative to most PMTs and reducing the sensitivity of the electrons to the magnetic field.

There are considerations when using MCP-PMTs that do not have to be taken into account when using standard PMTs. The charge lost from the walls of a microchannel during the formation of an electron cascade must be replenished before the microchannel can perform electron multiplication again. When high voltage is

applied across the plate, a strip current flows through the bulk material which replenishes the charge that was lost. This strip current is present even with no incident light due to the finite conductivity of the material, and might typically be around a few microamps. The result of this non-trivial recovery time is essentially a “dead time” for that specific microchannel. If this microchannel is struck again whilst it is still recovering, it will not be able to produce the same charge in response to an incoming electron and a lower gain will be observed. Following the treatment in [154] of the microchannels within an MCP behaving as a parallel plate capacitor with a recovery time constant of  $\tau = R_C C_C$ , with typical values of the channel resistance,  $R_C$ , of  $3 \times 10^{14} \Omega$  and channel capacitance,  $C_C$  of  $7 \times 10^{-17} \text{ F}$ , a dead time of around 20 ms should be expected. The reduction in gain will become larger as the frequency of the microchannel being struck increases. It is also worth noting that the recharge time is not consistent along the length of a single microchannel. Since electron multiplication is an exponential process, more charge is dislocated from the end of the microchannel than from the start and so will take longer to recover. This concept also applies to entire MCPs if they are in a stack. The uppermost MCP will be able to recharge much faster than lower down MCPs since less charge is dislocated during multiplication. The exact value of the recharge time will depend on the fabrication of the MCP. The time required to replenish the lost charge can be reduced by lowering the bulk resistance of the MCP material and thus increasing the strip current. This will increase the power requirements for operating the MCP but may be required for high rate applications. In the context of neutrino experiments, it is unlikely that the photon occupancy across the LAPPD photocathode area will be high enough for this effect to be an issue.

## 5.3 Traditional fabrication process

Prior to the development of the LAPPD, the fabrication process MCPs has remained fairly consistent for decades [159]. The precursor to the individual microchannels is formed by concentrically drawing lead glass tubes, known as the “cladding” glass, that are filled with glass rods, known as the “core” glass. The core glass is soluble in a chemical etchant, whereas the cladding glass is not. Each drawn fibre is then stacked in parallel into a hexagonal array pattern, and then drawn for a second time

to form hexagonal multifibre arrays. Many multiples of these multifibre arrays can then be stacked in parallel and fused together to form a boule. The boule is sliced at angle almost perpendicular to the fibre direction, typically at an angle of 8-15°, to form a thin wafer, often with a thickness on the order of a millimetre. A chemical etchant is then used to dissolve away the central core glass of each fibre, resulting in an array of many pores of around 10-20  $\mu\text{m}$  in diameter. The wafer is then hydrogen fired, which chemically reduces the lead oxide surface of the microchannel walls to conductive metallic lead which is secondary emissive. The level of hydrogen reduction has a strong effect on the final properties of the secondary emissive surface of the microchannel walls [160]. Electrodes are vapour deposited onto the top and bottom sides of the MCP to allow for the application of high voltage across the plate. Finally, the MCP is placed into the detector assembly along with the glass window containing the photocathode and any other associated components before the whole assembly is outgassed to produce a near-vacuum and then sealed.

There are certain drawbacks to the use of lead glass in MCPs. First of all, it is very mechanically weak which results in a constraint on the length to diameter ratio of a lead glass microchannel plate before it fails under its own weight. Secondly, lead glass contains many radiological impurities, the most dominant of which is  $^{40}\text{K}$ , which will produce secondary electrons as a result of the radioactive decay. This is thought to be the dominant cause of the contribution to the background count rate from the MCPs themselves [161].

A common area of concern in MCP fabrication is the lifetime of a single plate. During the fabrication process, large amounts of gas can be adsorbed onto the microchannel walls which is not completely eliminated during the vacuum sealing process. The formation of an electron cascade within a microchannel can ionise these gas molecules in the channel walls, which will then be accelerated in the electric field towards the photocathode. This ion bombardment can severely degrade the photocathode sensitivity and reduce the useful lifetime of the plate. Electron scrubbing, the practice of firing high energy electrons at the MCP in a vacuum environment, ionises and removes many of these impurities and is often used as a final pre-conditioning step before the detector is fully assembled. Whilst this process improves the effective lifetime of the plate, electron scrubbing can also produce

defects on the surface of the microchannels, which has a negative effect on the charge that can be extracted from a single microchannel, and hence the overall gain [162].

Atomic layer deposition (ALD) has begun to be utilised [163] as an additional post-processing step after the traditional fabrication process detailed earlier. By depositing a secondary electron emissive material with a higher secondary electron yield, the gain can be improved. In addition, the deposited material seals away the contaminants and gasses that become absorbed into the nanoscale holes of the glass, preventing them from becoming ionised as the electron cloud is formed and bombarding the photocathode. This can improve the lifetime of the MCP-PMT. ALD is discussed in more detail in subsection 5.5.3.

## 5.4 History of the LAPPD project

The desire for a photosensor with picosecond timing capabilities arose from a proposal for a time-of-flight detector to measure particle velocities at the Fermilab TeVatron in 2004 [164]. Due to the intrinsic relationship between the physical size of a photosensor and its timing resolution, determined by the variation in the time-of-arrival of the fastest electron to traverse the amplification stage, the detector must be small enough to achieve timing capabilities on the picosecond scale. However, a large-area detector is also desirable because cost per unit area generally reduces as the photosensitive area increases, resulting in large cost savings in experiments where photosensors typically dominate the budget.

The Large Area Picosecond PhotoDetector<sup>TM</sup> (LAPPD) collaboration, led by Argonne National Laboratory (ANL), began in 2009 to develop such a photosensor. MCP technology was the natural choice given the requirements as they were already known for their fast timing capabilities along with sufficient gain and low variability in performance across large areas. However, there were a number of challenges that required solutions. Mature, commercially available MCP-based detectors were small, with most only being a few centimetres in diameter at most. They were also typically expensive for the photosensitive area, with a difficult and extended fabrication process and, once they were fabricated and ready for use, the lifetime

of the device was limited as a result of ion feedback damage to the photocathode. Finally, there was also the challenge that, not only did this new detector have to be designed and fabricated, but the electronics required to read out the signals — with sufficient timing resolution to match the detector itself — would need to be developed as well.

The collaboration brought together expertise in a wide variety of subjects from US national laboratories, universities, and industry. Three high-priority areas of research and development were identified [165]:

1. The development of high quantum efficiency photocathodes
2. Utilising ALD to apply the emissive and resistive layers to the walls of MCP pores to provide them with their secondary-emissive properties
3. The development of high-speed waveform sampling integrated circuits (ICs) to read out the high bandwidth MCP-PMT signals

Furthermore, research occurring in parallel to this effort focussed on producing an alternative to the typical ceramic hermetically-sealed package utilising an all-glass module. This provided the potential for a cheaper package due to fewer components and simpler assembly.

The US company Incom Inc [166], based in Sturbridge, Massachusetts, specialise in the production of glass optical devices such as fibre optics and lenses. Incom led the development of the glass substrates containing pores with diameters of 20 microns. These would then be sent to ANL to become an MCP-PMT via the recently developed ALD technique before then being tested. By 2013, the technology was mature enough to begin the commercialisation process.

The expertise and techniques that were developed throughout the course of the project has now been transferred to Incom, who are now devoting significant effort into mass manufacturing of the LAPPD. As of the time of writing, approximately two hundred LAPPDs have been produced (including successes and failures), across two generations of design. For a full documentation of the history of the LAPPD project, the reader is guided to [165].




Figure 5.3: Exploded view of the Generation I LAPPD. This image shows an older design which uses X-spacers rather than the newer bar spacers. Image from Incom Inc.

## 5.5 The LAPPD package

The LAPPD, shown in Figure 5.3, as it is now produced by Incom is a hermetically-sealed glass package of approximately  $30 \times 30$  cm of which  $20 \times 20$  cm is the area of the photocathode. The components that comprise the LAPPD are as follows:

### 5.5.1 Ultem frame

The LAPPD frame serves much of the same purposes as housings on traditional PMTs. It provides mechanical stability to the entire package and protects the fragile glass package to ensure the integrity of the vacuum seal. Additionally, the frame houses the five safe high voltage (SHV) connectors that distribute the high voltage to the different stages of the LAPPD through electrodes that penetrate the vacuum seal. Finally, the frame also provides mounting holes to allow the LAPPD to be fixed onto another surface without the risk of damage. The frame is 3D printed using Ultem® thermoplastic.






Figure 5.4: An example of the quantum efficiency for a bialkali photocathode with either a fused silica or borosilicate glass window, demonstrating the extended transparency of fused silica glass deep into the ultraviolet region of the spectrum. Image from [146].

### 5.5.2 Glass window and photocathode

The glass window is made from an ultraviolet-grade fused silica glass which has good transparency across the visible wavelength ranges and into the ultraviolet as far as 165 nm [167, 168], as can be seen in Figure 5.4. On the underside of this window is the photocathode material which is applied onto the glass using evaporative deposition. The photocathode material is potassium-sodium-antimony ( $K_2NaSb$ ). These low-noise, blue-sensitive bialkali photocathodes have very high temperature tolerances, with the additional benefit being that it has very low dark current at room temperature, an ideal characteristic for applications at the single photon level [169]. The spectral response of the photocathode peaks at around 365 nm with a quantum efficiency that can be in excess of 30% [168].

A considerable amount of effort has been invested into the photocathode deposition procedure, as there is a significant challenge in uniformly coating such a large area of glass. Due to the inconsistencies that exist in the deposition process, there is likely to be some variation in the quantum efficiency, and indeed the overall performance of the detector, over the entire surface area of the LAPPD.




Figure 5.5: Image of a  $20 \times 20$  cm glass capillary array fabricated by Incom Inc. The glass capillary array appears translucent because the surface contains many  $20 \mu\text{m}$  pores and the open area is greater than 60%. This will later be converted into a microchannel plate through atomic layer deposition of secondary emissive materials. Image from Incom Inc.

### 5.5.3 MCP stack

The glass window is sealed to the sidewalls using a molten indium alloy to form a hermetically-sealed glass package which contains the MCP stack in a near-vacuum environment. The side walls and the bottom plate of the package are made from borosilicate float glass. The high voltage connectors pass the supply voltages for the LAPPD onto pins that penetrate through the glass package onto the top and bottom of each of the two MCPs that are arranged in a chevron-pair inside the vacuum. This arrangement is shown in Figure 5.6. The top and bottom surfaces of each MCP in the stack is coated with a film of nichrome to form an electrode. These electrodes provide a path for the delivery of high voltage to the MCP stack to accelerate the electrons and restore charge in the microchannels. This layer is about  $200 \text{ nm}$  in thickness and is applied from a set angle relative to the surface normal of the MCP to ensure that deposition of the nichrome inside of the microchannels is as uniform as possible. In general, the microchannel walls are coated in nichrome at the top and bottom up to a height of about one to two times the microchannel diameter during the process — this is known as “end spoiling”. This reduces secondary emission in the “spoiled” region near the output of the microchannel, resulting in




Figure 5.6: A cutaway image demonstrating the internal design of the LAPPD. The full stack consisting of the glass window, photocathode, chevron pair of MCPs, (pores not drawn to scale) and the anode can be seen. The blue hatched area shows the electron cascade as it travels through the MCP. The high voltage connections to the LAPPD are shown on the left-hand side.

greater collimation of the electrons leaving the microchannel, at the expense of gain [170–172].

Above and below each of the MCPs in the stack is a spacer. In earlier LAPPD designs such as is shown in Figure 5.3, the spacers had an X-shape, but later designs use a series of smaller linear spacers. The purpose of the spacer is to prevent deflection of the glass window, and the entire LAPPD package, under atmospheric pressure which can occur due to the large cross-sectional area of the LAPPD. However, these spacers introduce dead regions into the photocathode area of the LAPPD as they obscure the microchannels and so within these regions no photons can be detected. The positioning of the spacers sets the gaps between the photocathode, each of the MCPs, and the anode, which can result in significant differences in the performance characteristics of the LAPPD.

Each plate typically has a pore diameter of  $20\ \mu\text{m}$ . LAPPDs utilising a  $10\ \mu\text{m}$  pore diameter have also been developed, which will provide improvements in time and position resolution, although this comes with the drawback of being more difficult to fabricate. The MCPs have an open area ratio of up to approximately 74%, which gives the pre-treated plate a translucent appearance, as shown in Figure 5.5. The length to diameter ratio is 60:1.

The LAPPD is a result of two major innovations in the MCP fabrication process:




Figure 5.7: Borosilicate glass capillary array functionalised by the application of thin films of resistive and emissive layers using ALD. The off-axis tilt of the microchannels is not shown.

the use of borosilicate glass capillary arrays, and ALD [173]. Unlike traditional MCPs that are fabricated with lead glass, the MCPs that make up the LAPPD are fabricated by drawing borosilicate glass into hollow tubes. Borosilicate glass is much stronger than lead glass, effectively eliminating the mechanical limitations on the size of the MCP, and also can be made much radiologically cleaner. The background count rate is around five times lower than lead glass MCPs due to the lower concentration of  $^{40}\text{K}$  [174].

Since borosilicate glass is not naturally secondary emissive, it is functionalised through the use of ALD [175], a type of chemical vapour deposition that does not require line of sight and is well-suited for coating the microchannels of an MCP. ALD is used to apply a thin film made up of a resistive and emissive layer to the microchannels. These layers are shown in Figure 5.7. The resistive layer is added first; this layer consists of a conductive material such as nanocomposite tungsten-aluminium oxide ( $\text{W}:\text{Al}_2\text{O}_3$ ) with a thickness of around 5-10 times that of the emissive layer. This allows the strip current to flow through the bulk material of the plate and replenish charge to the emissive layer after an electron cascade is formed in the microchannel. The resistivity of the bulk material can be tuned by adjusting the cycle percentage of tungsten when applying the  $\text{W}:\text{Al}_2\text{O}_3$  layer; a

change in the tungsten cycle percentage between 10 and 30% can change the bulk resistivity by around four orders of magnitude [176]. MCPs with high bulk resistances may be desirable as it reduces the power requirements for operation but at the cost of reducing the strip current and thus reducing the rate at which charge can be resupplied to the microchannels. For high rate applications, a low bulk resistance is required to ensure the recharge rate is as high as possible. The emissive layer is then applied onto the resistive layer. A material with high secondary electron yield is used, such as MgO or Al<sub>2</sub>O<sub>3</sub>. The total thickness of the deposited films is around 100 nm.

This fabrication method has a number of advantages over the traditional method. Firstly, the use of ALD to apply the secondary emissive material effectively decouples the properties of the substrate and the active material, allowing for each to be optimised according to their purpose. ALD is also much more tunable than the hydrogen reduction process used in lead glass MCPs, allowing for optimisation of the film properties by selecting different materials and thicknesses. It is possible to achieve much higher gains by using ALD-coated MCPs. The lead sub-oxide emissive layers typically have a secondary electron yield of around 2, whereas Al<sub>2</sub>O<sub>3</sub> and MgO have secondary electron yields of 3 and 7 [177], respectively, and can be tuned to some degree by the thickness of the emissive layer.

### 5.5.4 Anode

Once the electron multiplication has occurred, an observable signal must be produced. This is most commonly achieved by collecting the electrons at an anode that makes up the final layer of the stack within the vacuum volume, although it is worth noting that it is also possible to simply measure the current through the MCP stack (after decoupling the high voltage with a suitable splitter circuit) to observe the discharge and recharge of the plates which can act like a photon counter. The anode will then form a signal that can be propagated on to the front-end digitising electronics.

The choice of anode design is application specific. If position information is

not required, a single monolithic anode can be used to record timing and bulk charge. Otherwise, the anode is generally segmented in some way to provide a degree of position resolution, which is naturally dependent on the segmentation. The Generation I LAPPD uses a series of 28 silver microstriplines printed onto the inner side of the borosilicate glass that makes up the bottom window of the LAPPD package. The strips pass through the sidewall of the hermetically sealed package, which is glass frit bonded onto the lower window, so that signals can be transmitted outside the vacuum chamber. The signals are passed from the strips into traces on a printed circuit board backplane with a characteristic  $50\ \Omega$  impedance which are routed to SMA connectors. Each strip is 5.2 mm in width with 1.7 mm spacing in between, for a strip pitch of 6.9 mm. In addition to the 28 striplines carrying signals from the LAPPD there is an additional grounding strip at the top and bottom.

Using a stripline anode, pulses are formed as the strips collect the electron cascade from the MCPs and travel in both directions towards each end of the strip. As discussed in more detail in section 6.3, the position of the incident photon can be inferred from the difference in a pulse's time of arrival at each end of a given stripline, and from the relative pulse height measured across multiple neighbouring striplines. An example of this is shown in Figure 5.8. The stripline approach moves much of the complexity of determining the photon position to software-based signal processing algorithms that are performed after data is collected. The benefits of the stripline approach is that it offers very good position and timing resolution with the minimum number of readout channels. It may even be possible to reduce the number of readout channels further by only reading out a single end of the strip and calculating the position using the reflection of the pulse off the unterminated end of the stripline. However, this is likely to introduce difficulties in distinguishing signals if multiple photons strike the same stripline in quick succession.

Any anode with neighbouring conductors will suffer from crosstalk to some degree. Crosstalk is the effect where a signal travelling in a circuit induces a signal, often undesirably, in other regions of the circuit. This is distinct from the "charge-sharing" effect of MCP-PMTs, whereby the electron cloud spreads radially as it travels down the multiplier stack. Crosstalk in the stripline anode is a result of the mutual inductance and capacitance between neighbouring striplines. The




Figure 5.8: The stripline anode readout on the Generation I LAPPD. The photon location along the stripline can be inferred from the difference in the time of arrival at each end. The location across striplines can be inferred from the centroid of the relative pulse heights as shown in the image.

mutual inductance induces current in the opposite direction of the driving current as a result of Lenz's law. Conversely, mutual capacitance passes current from the driving stripline to its neighbour that travels in both directions on the neighbouring stripline. Crosstalk can be observed on the near-end and far-end of the neighbouring stripline, where the near-end refers to the end of the stripline that is being driven on the driver line. At the near-end, the inductive and capacitive currents sum and the resulting waveform on the neighbouring stripline is always positive. However, at the far-end the inductive current is subtracted from the capacitive current resulting in a waveform on the neighbouring stripline that is usually negative. An example of near-end and far-end crosstalk can be seen in Figure 5.9. In the context of an LAPPD with stripline readout, the situation becomes more complicated than the idealised example because the driven line is not driven from one end, but rather current is induced on the driven line at some




Figure 5.9: An example of crosstalk occurring on neighbouring striplines. Capacitive crosstalk is shown on the left and inductive crosstalk is shown on the right. Image from [179].

point on the stripline as the electron cloud travels past. However, it is still important to understand the crosstalk process, because any waveforms produced on striplines as a result of the electron multiplication process will be modulated by mutual crosstalk between striplines. Minimising crosstalk is one of the ways in which the stripline anode design has been optimised, along with maximising the bandwidth. It has been found that ceramic anodes display far more desirable electromagnetic properties than glass anodes. The higher dielectric constant results in higher bandwidth and less crosstalk between striplines. A full characterisation of the crosstalk on the stripline anode is beyond the scope of this thesis, however preliminary studies of crosstalk on different anode designs can be found in [178].

Incom has also developed the Generation II LAPPD, which differs from the Generation I primarily in the design of the anode. In comparison to the direct coupling on the Generation I, the Generation II uses a resistive metallic anode, which can then be capacitively-coupled with conductors on an external readout board. Therefore, the signal pickup occurs outside the hermetically sealed housing, and has the effect of decoupling the detector itself from the signal pickup geometry. The result is that the implementation of the readout board can freely be modified at any time even after the LAPPD has been fabricated and sealed. Incom's default pickup board design for the Generation II utilises a pixelated readout board with 64



pixels organised in an  $8 \times 8$  grid for a total of 64 channels, but this board can easily be redesigned for a specific application and swapped onto a Generation II LAPPD. A number of different anode patterns have been designed for the readout board [180]. If desired, one could even choose a readout board consisting of microstriplines similar to the Generation I. Capacitively-coupled readout offers significant benefits from a manufacturing perspective [181]. It removes all the signal connectors that penetrate through the vacuum seal, reducing the risk of a vacuum failure. Additionally, since the readout is now decoupled from the detector itself, the design of the vacuum enclosure can be identical for every LAPPD produced, simplifying the manufacturing process and reducing costs.

## 5.6 LAPPD operation

### 5.6.1 High voltage

The LAPPD takes five high voltage inputs. From the bottom of the package to the top these voltages are applied to: the exit of the exit MCP, the entrance of the exit MCP, the exit of the entrance MCP, the entrance of the entrance MCP and the photocathode. These connections are shown on the LAPPD system diagram in Figure 5.6. The supply voltages are typically specified by the potential difference between the stages. Typical supply voltages for LAPPD 104 are shown in Table 5.1

There are multiple methods for powering the LAPPD, each with advantages

Connector	Typical gap voltage (V)	Typical supply voltage (V)	Typical current ( $\mu\text{A}$ )
Exit of Exit	200	-200	192
Entry of Exit	850	-1050	210
Exit of Entry	200	-1250	1222
Entry of Entry	850	-2100	87
Photocathode	200	-2300	0

Table 5.1: The high voltage connections to the LAPPD and the typical gap and absolute supply voltages for LAPPD 104. The optimal values will vary from LAPPD to LAPPD. The locations of each high voltage connector in the LAPPD stack are shown in Figure 5.6.

and disadvantages. The first is to simply use five independent voltage supplies. It is important to ensure that the current limit of the power supply is high enough to provide enough current to the LAPPD, since certain channels require in excess of 1 mA. However, for deploying an LAPPD in a neutrino detector, it may be unwieldy to have five high voltage connections for every LAPPD.

Since the most optimal voltage gaps across each stage of the MCP stack are typically consistent over a reasonable period of time for a given LAPPD, it is possible to produce a voltage divider that takes the photocathode voltage as an input, and that voltage is then divided down with a series of resistors to provide the desired voltages to the remaining four inputs to the LAPPD. There is a subtlety in powering the LAPPD using a voltage divider in that the MCPs themselves have multi-megaohm resistance and so will contribute to the overall resistance in the divider. Furthermore, the two MCPs in each LAPPD are not guaranteed to have comparable resistance — in fact they can differ by tens of megaohms. This introduces a further challenge in that it is desirable to have the same voltage across each MCP but if they have significantly different resistances then the voltage will not divide evenly across each plate. To rectify this, another resistor can be added in parallel with the higher resistance MCP in order to reduce the overall resistance of that branch of the divider to match the resistance of the other MCP. This will divide the voltage in the most desirable fashion.

### 5.6.2 Fast digitisation

In order to exploit the timing capabilities of the LAPPD, the inherent time resolution of the LAPPD must be matched by the digitising electronics. Firstly, in order to adequately sample a signal, the Nyquist-Shannon limit [182] must be met. That is, for a signal with bandwidth  $B$ , the signal can be completely determined if it is sampled at a rate of  $2B$  samples per second. The bandwidth of a signal is the frequency of the highest significant sine wave frequency component that is present in the signal. This can be estimated using the rule of thumb that  $B = 0.35/\tau_r$  [183] where  $B$  is the analogue bandwidth and  $\tau_r$  is the 10%-90% rise time of the signal. If we approximate the rise time of an average LAPPD pulse to be 0.6 ns, then the analogue bandwidth of the LAPPD signal is  $0.35/0.6 \times 10^{-9} \approx 600$  MHz,

and so  $1.2\text{GS/s}^1$  is the minimum sampling frequency required to meet the Nyquist-Shannon limit. It is still desirable to sample the signal at a rate greater than the Nyquist-Shannon limit if possible as this adds redundant points which improve the signal-to-noise ratio. Since timing is usually determined by analysis of the rising edge of a pulse, increasing the sample size in this region reduces the impact of uncorrelated noise from the digitising electronics, such as voltage and timing jitter (see subsection 5.7.2 for discussion of this topic).

It is clear from the requirements on bandwidth and channel density that digitising signals from an LAPPD is a challenging problem. Most commercial ADCs do not have sampling rates into the GS/s range and if they do, they are often incredibly expensive or have drawbacks in other areas such as the channel count or resolution. High channel density and high speed digitisation is likely to be the next most expensive component after the photodetector itself and as such, significant effort must be exerted to ensure that a suitable digitisation solution is chosen. The exact system will be specific to each experiment, since the digitisation system must be able to meet the experiment's requirements in terms of data acquisition and interfacing with other systems.

Cost-effective, high sample rate digitising circuits have been developed using switched capacitor array technology. Two such options based on this technology are explored in the context of neutrino experiments in subsection 5.8.2.

## 5.7 LAPPD time resolution

LAPPDs have the potential to achieve very impressive timing resolutions on the order of tens of picoseconds. When using high speed detectors such as the LAPPD, there are many effects that can limit the time resolution which must be kept under consideration, both within the photosensor itself and in the digitising electronics. The two most limiting factors on the timing resolution in the case of the LAPPD are [184]:

- Fluctuations in the time difference between photoelectron production and the

---

<sup>1</sup>Sampling frequencies are referred to in units of samples per second to avoid confusion with analogue bandwidth, measured in Hz.  $1.2\text{GS/s}$  is 1.2 giga samples per second.

collection of the electron cascade at the anode.

- Signal degradation within the transmission and digitisation stages of the frontend electronics.

These two factors are explored in the subsequent subsections.

### 5.7.1 Variability in electron cascade timing

The first reason for timing fluctuations in the electron cascade is variability in the time required for the photoelectron to initiate the electron cascade in the first microchannel after it is emitted from the photocathode. Much of this variability is due to the multiple possible paths that the photoelectron can take to reach the first microchannel, as shown in Figure 5.10.

During the course of producing a “normal” pulse, the photoelectron strikes at an intermediate depth into the microchannel and begins the electron cascade partway down the pore. Pre-pulsing can occur when the incident photon is transmitted through the photocathode and strikes the microchannel directly, resulting an undersized pulse as fewer secondary electrons are produced.

Another possible path for the photoelectron to take results in it hitting the surface of the plate between the pores rather than entering into a microchannel at the first attempt. The photoelectron bounces back upwards before being refocussed by the electric field into a microchannel. This incurs a delay in the time for the electron cascade to reach the anode and is known as late-pulsing. Late pulses typically arrive  $\mathcal{O}(1)$  ns later than the average pulse.

Afterpulsing is a similar but unrelated effect caused by ion feedback from residual gaseous molecules that remained after the enclosure was vacuum sealed. The ion drifts upwards in the electric field and strikes the photocathode resulting in the release of multiple electrons that go on to produce their own electron cascades. This effect typically causes very large pulses and, unlike pre and late-pulsing, is always correlated with a normal initial pulse but occurs  $\mathcal{O}(1)$   $\mu$ s later.

The fluctuation in the electron cascade reduces statistically with increasing numbers of incident photons (and hence produced photoelectrons), and so single




Figure 5.10: Schematic showing the mechanisms for how early, normal, and late pulses occur within the LAPPD. The photoelectrons in both the normal and late pulses reach the top of the microchannel plate at the same time, but the late pulse initiates the electron cascade much later due a longer photoelectron path length. Prepulses occur when the photocathode transmits the incident photon which then goes on to strike the MCP directly and produce secondary electrons.

photoelectrons will have the worst timing jitter contribution from the MCPs. It is worth noting that this relationship for multiple photon hits is not exact, and can be affected by factors such as whether the photoelectrons enter the same or multiple microchannels and whether any of those microchannels saturate as a result.

The single photoelectron timing variability is improved by increasing the number of secondary electrons produced in the first collision. This has the effect of reducing fluctuations resulting from the randomised secondary electron energies and directions. The number of secondary electrons produced can be increased by increasing the kinetic energy of the photoelectron at the point it collides with the microchannel wall by increasing the voltage across the photocathode to first MCP

gap or by increasing the secondary emission yield of the material coating the microchannel walls.

### 5.7.2 Signal degradation

Even if the LAPPD behaved in an ideal manner and generated pulses with zero timing variability, there would still be limitations on the possible timing resolution. The presence of noise in the electronics produces an uncertainty in a pulse's threshold crossing time that can be reduced, but never avoided completely, by better signal conditioning and processing.

Figure 5.11 shows how voltage jitter in the digitising electronics can have an effect on the time resolution. Timing measurements are typically performed by analysing the rising edge of the pulse and extracting the time at which some threshold is passed. In order to improve the time resolution beyond the sampling interval, the waveform is typically interpolated (usually linear or cubic depending on the application) as the threshold crossing point is likely between samples. Jitter on the voltage measurement will induce jitter in the threshold crossing time as the gradient, and thus the interpolated samples, between two adjacent samples will change. Noise on the voltage that increases the measured voltage for the sample of interest will make the calculated crossing time earlier than if there were no noise and likewise noise that decreases the measured voltage at the sample will make the crossing time later. The impact of voltage jitter is reduced if the rise time is quicker, although the rise time is determined by the detector and is generally not (easily) adjustable. The alternative way to reduce the impact of the voltage jitter on timing is to sample more points along the rising edge. If the voltage noise is stochastic, which should be the case for voltage jitter in an ADC, then the timing jitter will reduce with  $\sqrt{N}$  where  $N$  is the number of samples taken along the rising edge [185]. It is therefore clear that the intuition that sampling the waveform at a higher frequency will improve the timing resolution is correct, at least to first order. In reality, the time resolution will also have a dependence on the analogue bandwidth of the digitising electronics and the detector.




Figure 5.11: The impact that voltage jitter in the digitising electronics can have on the resulting the timing resolution. Image from [184].

## 5.8 Water Cherenkov LAPPD deployment

Deploying an LAPPD into a water Cherenkov detector presents a series of challenges. The BUTTON experiment, a UK-based successor to the WATCHMAN collaboration is planning to deploy at least two LAPPDs. The BUTTON experiment is planned to be a roughly 30 tonne detector, initially filled with gadolinium-doped water, but will later be filled with WbLS. It will act as a testbed for new technologies for neutrino detection, and aims to perform measurements including LAPPDs. The ANNIE collaboration [186] was the first to successfully deploy LAPPDs into a water Cherenkov experiment. The lessons learned by ANNIE during this process, as well as the experiences of the author assisting with the deployment of LAPPDs in the ANNIE detector, will be used to guide the development required to deploy LAPPDs in BUTTON. A number of considerations are detailed in the following subsections, along with the current best estimate of how these will be solved.

### 5.8.1 Waterproof housing

A waterproof housing is required to ensure the LAPPD and electronics are not exposed to the outside environment. The housing itself can be constructed from any gadolinium-compatible plastic that can easily thermoformed or injection-moulded, such as polyvinyl chloride. The housing will have a front window made from UV-transparent acrylic to allow transmission of photons to the

LAPPD photocathode. UV-transparent acrylic transmits photon wavelengths down to the near-UV which, whilst not as deep into the UV as the fused silica window of the LAPPD itself, is adequate for photon wavelengths produced by Cherenkov radiation and typical scintillators. The backplane of the housing will be made from gadolinium-compatible stainless steel (such as 316L stainless) which will allow some degree of heat dissipation from the inside of the enclosure into the detector medium. The window and backplane will be fixed to the housing using screws, with each side containing a groove and o-ring to ensure a water-tight seal. Cable feedthroughs will be provided using high-density waterproof connectors and waterproof cables will be run from outside of the tank into the housing to provide power, signal, and communications lines.

### 5.8.2 Signal digitisation

With 56 readout channels on the Generation I LAPPD, it is not practical to run 56 signal cables from a submerged LAPPD to typical rack-mounted digitisers outside of the tank. Instead, digitisation will be performed inside the waterproof enclosure, as close to the LAPPD as possible. As of the time of writing, there are no commercially available digitiser boards that can provide full waveform digitisation at the required sampling rates and be deployed inside the limited space of the waterproof enclosure. At least two prototype digitisation options are under consideration for use in BUTTON. These are detailed in the subsequent sections.

#### DRS4

The Domino Ring Sampler 4 (DRS4) is a very mature digitising application-specific integrated circuit (ASIC) based on switched capacitor array technology [187]. A switched capacitor array is a chain of capacitors that acts as a sample and hold circuit. A series of inverters generates a pulse which connects a switch to each sample and hold circuit in sequence, charging the capacitor and thus sampling the waveform for a given period of time before the switch disconnects and the next capacitor begins sampling. After the capacitor is disconnected, its content is then passed on to an external, commercially available ADC via a shift register for digitisation at the desired resolution. A phase-locked loop (PLL) is



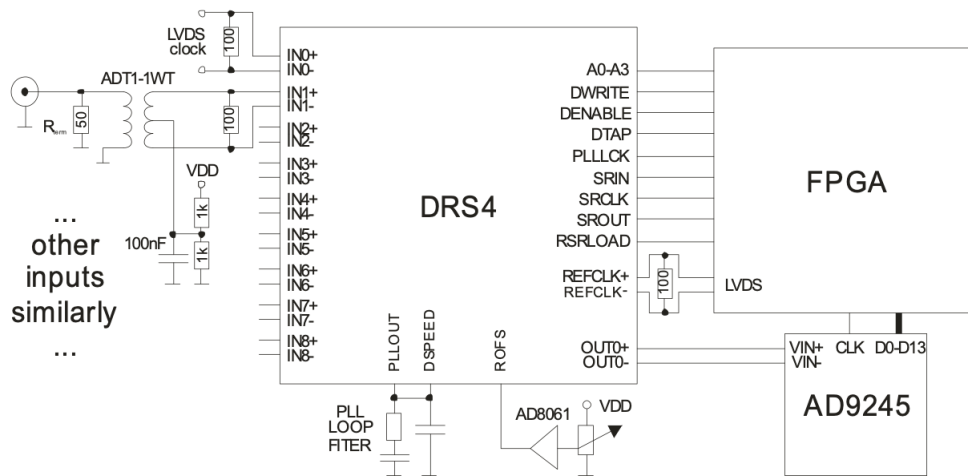


Figure 5.12: The typical operation of the DRS4 chip. The choice of ADC and FPGA will depend on the specific application. Image from [188].

used to precisely control the delay between each switch opening to set the sample rate of the system and to stabilise the system from the affect of external factors, such as variations in supply voltage or temperature, on the inverter chain. Like many modern, high speed digitising circuits, the DRS4 uses differential signalling for its analogue inputs as this provides the benefit of doubling the dynamic range and rejecting common-mode noise and interference. This topology of digitisation circuit using switched capacitors allows for very high sample rates, up to around  $\sim 15$  GS/s as of the time of writing. A field-programmable gate array (FPGA) is used to control the DRS4, and to retrieve and process the digital output from the ADC. Rudimentary threshold-based triggering logic can be created by feeding the OUT0+ signal through a hardware comparator with the desired reference voltage, the output of which is then passed into the FPGA. The FPGA then processes the comparator output and can issue a trigger signal to the DRS4 if necessary.

The DRS4, along with other digitisers based on the switched capacitor array technology, follow a mantra of “sample fast, digitise slow” — the switched capacitor array can sample the waveform at very high speeds before passing those samples on to a much slower commercial ADC for digitisation. The result of this is that all switched capacitor array-based digitisers share a disadvantage that it takes significant time to discharge and read out the capacitors into the ADC which

results in dead time when data acquisition cannot occur. Readout can take up to around 30 ns per capacitor cell in the case of the DRS4, so in order to digitise 56 channels with a typical sampling depth of 1,024 capacitor cells per channel, the dead time is approximately 1 ms, resulting in a maximum acquisition frequency of around 1 kHz. This reason for this dead time on the DRS4 is because the multiplexed differential output line of the shift register (labelled `OUT0+` and `OUT0-` in Figure 5.12) is clocked at 33 MHz which is the input to an external ADC operated at a typical sample rate of around 50 MS/s. This limits switched capacitor array digitisers to low trigger rate situations where good time resolution is required. Luckily, neutrino experiments tend to fit into exactly this mould. In general, the time resolution of switched capacitor array digitisers lies in the range of 5 ps to 25 ps. The lower end of this range, and even values below that, are achievable only if the proper calibration is conducted to account for voltage and timing errors present in the system. For a well-calibrated DRS4 system, the time resolution is typically around 5 ps and so although the timing uncertainty from the electronics noise is not trivial, it is likely dominated by the timing uncertainty produced by the physical effects inside the MCPs themselves, as explained in subsection 5.7.1.

A prototype DRS4-based readout board designed specifically for the LAPPD has been designed by the University of Hawaii and Ultralytics LLC [189]. This was accompanied by the development of open-source data acquisition firmware and software. The board uses 8 DRS4 chips, which each have 8 input channels. Of these input channels, 7 are used for LAPPD signals and the final channel is used to inject a 100 MHz high-precision sinusoidal oscillator signal as a timing calibration source. Signal digitisation is performed by two 32-channel ADS52J90 ADCs. A central FPGA provides control signals for all DRS4 and ADC chips, and provides triggering capabilities, a soft-core processor, and gigabit Ethernet transceiver. Unfortunately, development has paused and any experiment wishing to use this option would need to restart the firmware and software development and perform extensive validation testing before this could be used in an experiment.

### PSEC4

The second option is the PSEC4 [190] digitising ASIC used by the ANNIE collaboration. The PSEC4, designed at the University of Chicago, is a six-channel switched capacitor array digitiser that is capable of sampling at rates up to 15 GS/s with a buffer depth of 256 samples. The PSEC4 offers a very favourable analogue bandwidth of around 1.5 GHz, which is more than suitable for digitising LAPPD pulses. A block diagram of the PSEC4 is shown in Figure 5.13.

ANNIE uses the “ACDC” board which contains five PSEC4 chips and so two ACDC boards are required to read out an entire LAPPD. An FPGA coordinates the timing and control signals for the PSEC4 chips and correlates their digital outputs. The ACDC boards can be attached to the LAPPD via a custom designed “pickup” board that takes the signal outputs from the LAPPD striplines and feeds them into the PSEC4 inputs through a high density SAMTEC connector. Whilst the PSEC4 chip has a discriminator for internal triggering, ANNIE uses a custom-designed trigger board as a more robust solution. The trigger board replaces one of the standard Incom passive sideboards. It serves the same purpose as the passive sideboard; it routes the signals from the striplines into the backplane, but it adds additional comparator circuitry that monitors every stripline against some programmable reference voltage and issues a trigger pulse if any stripline voltage exceeds the reference. This provides essentially a basic threshold trigger. This trigger signal can then be fed back into all the PSEC4 chips to begin the data acquisition simultaneously on every chip.

The PSEC4 chips have been extensively used by ANNIE and have proven capable of digitising the fast signals from an LAPPD. However, the shallow buffer depth and the usage of components that are now end-of-life presents challenges for new experiments hoping to leverage the PSEC4. These have been addressed by the introduction of the PSEC4a, which has an equivalent buffer depth to the DRS4, but the availability of these chips is not clear.

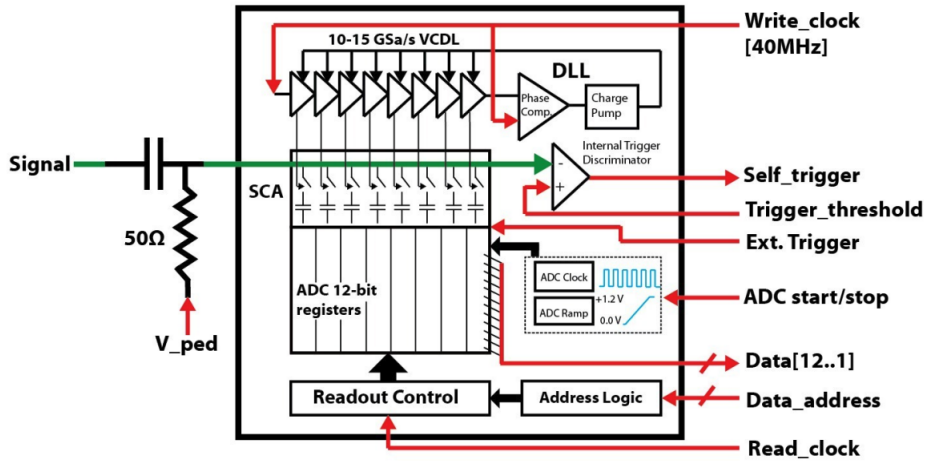


Figure 5.13: Block diagram of the PSEC4 digitising ASIC. Image from [190].

## Comparison

Table 5.2 presents a summary of the two digitiser options that are under consideration for an LAPPD deployment.

### 5.8.3 Power

The LAPPD module will require high voltage for the amplification stages of the LAPPD itself, as well as low voltage for all the associated frontend electronics.

#### High voltage

The high voltage connections can be provided in two ways. The first is to pass the five high voltage connections down a cable from commercial high voltage power supplies out of the tank into the LAPPD module. This carries the disadvantage of involving more cables to manage, and the possibility for introducing cross-talk into the LAPPD signal path if the two paths are too close together. The other option, which was chosen by the ANNIE collaboration, is to produce the high voltages inside the module using a custom low voltage - high voltage board that produces all the necessary supply voltages on a single board. The required high voltages can be produced in-situ using a high precision, low ripple, remotely programmable boost

## 5.8. Water Cherenkov LAPPD deployment

---

Characteristic	DRS4 [191]	PSEC4 [190]
Manufacturer	PSI	U. Chicago
Channel count	8 + 1 trigger	6
Sampling rate	0.7 - 5 GS/s	4 - 15 GS/s
Resolution	Determined by external ADC	12 bits
Buffer depth /ch	1,024 samples	256 samples (1,024 for PSEC4a)
Bandwidth	950 MHz*	1.5 GHz
Dynamic range	1 V <sub>pp</sub>	0.8 V <sub>pp</sub>
Noise	350 $\mu$ V RMS	700 $\mu$ V RMS
Dead time	30 ns per cell	12.5 ns per cell
Max power dissipation	$\sim$ 350 mW	$\sim$ 330 mW

Table 5.2: Comparison of the potential digitising ASICs to be used in an LAPPD deployment. \* Bandwidth of the DRS4 only, the bandwidth may be limited by the external ADC.

converter. These boost converters have a control voltage input, which sets the output voltage of the converter. A programmable digital to analogue converter (DAC) can be used as the control voltage input which can be programmed to ramp the voltage at a rate that is safe for the LAPPD. The boost converter reads this control voltage and sets the output voltage appropriately. By using a voltage divider, as explained in detail in subsection 5.6.1, this output voltage can be split to provide the suitable voltage to each of the five stages of the LAPPD.

When designing a custom high voltage supply, one must ensure there is adequate overvoltage protection to prevent damage to the LAPPD. Zener diodes can be placed at each stage of the voltage divider to ensure that the voltage across the stage does not exceed some predetermined safe value. Due to the differences in high voltage performance from one LAPPD to the next, these safe voltage limits should be determined independently for each LAPPD.

For BUTTON, both options of high voltage configuration remain open. Since BUTTON will act as a testbed and there is a high likelihood that the high voltage configuration will be adjusted on short timescales, using commercial high voltage power supplies and feeding five high voltage cables per LAPPD will likely be the best option. If or when longer data taking runs are conducted by BUTTON, a low voltage - high voltage converter similar to that used by ANNIE can be developed.

### **Low voltage**

The low voltage system powers everything else inside the LAPPD module. The digitising electronics will use ASICs or FPGAs that will have supply voltages of 5 V or less. If the high voltage is being produced on the same board, it is practical to have a higher supply voltage than this to power the boost converter such as 12 V. Linear regulators or buck converters can then be used to step down this input voltage to the desired level. It is important to keep the heat dissipation under consideration. Primarily, it is important to ensure that no voltage regulation is performed in proximity to the LAPPD itself. This is because the glass window is vacuum sealed to the body of the LAPPD using an indium alloy, which has a relatively low melting point of approximately 60 °C. Since linear regulators, and to a lesser extent buck converters, convert much of the input voltage to thermal energy, they can easily heat up the environment to a temperature that could be dangerous for the integrity of the vacuum seal. For this reason, regulation of the low voltage inputs will be performed away from the digitisation electronics (which will likely be close the LAPPD), most likely on the same board used to generate the high voltage.

### **5.8.4 Monitoring and slow control**

The LAPPD module will be continuously monitored whilst submerged in the detector to ensure that the operating conditions of the system are recorded for future reference and any catastrophic failures (such as water ingress into the waterproof housing, voltage spikes, or dangerously high temperatures) are noticed promptly. The values that will be reported in the monitoring system include the output of all voltage regulators for the low voltage system and the DAC output voltage that controls the output of the high voltage boost converter. In addition, a humidity sensor and thermistor will monitor the ambient humidity and temperature. If either of these values crosses some preprogrammed safe value the power to the module will be cut. Finally, a leak sensor will be positioned at the lowest edge of the module which will trigger in the case that the waterproof seal fails and cut the power to the module to prevent as much damage as possible.

## 5.9 Chapter summary

- MCP-PMTs are continuous dynode photosensors that are capable of detecting single photons with timing resolutions in the tens of picoseconds. Position sensitivity is also possible due to the “pixellated” nature of the microchannel plate.
- The LAPPD is a novel MCP-based photosensor technology manufactured by Incom, Inc that uses borosilicate glass functionalised by atomic layer deposition. Many years of development were required to develop the techniques and experience required to commercialise such a device.
- LAPPDs have many advantages over traditional lead glass MCP-PMTs including higher gain, improved mechanical robustness, and lower background count rates.
- There are two generations of LAPPD, the first of which uses directly coupled stripline anode readout. The second uses a flexible capacitively coupled readout, allowing the user to design their own readout board to best suit the application.
- High-speed digitising electronics are required in order to make use of the ultrafast detection capabilities of the LAPPD. There are many challenges present in utilising such systems.
- There exists many challenges to deploying LAPPDs in neutrino experiments. Solutions must be found for waterproofing, thermal management, compact frontend digitisation, and provision of high voltage.

## Chapter 6

# LAPPD characterisation

It is important to understand the behaviour of any photosensor before it is deployed in an active experiment. The dark rate, gain, and timing of a photosensor can have a large impact on any subsequent event building and analysis, and so understanding these properties is a vital step in the characterisation of the device. It is doubly important for novel photosensors still in development such as LAPPDs, where the performance can vary drastically from one unit to the next, or even across a single device. These variations will ultimately result in systematic biases in any subsequent event building and reconstruction that must be accounted for by thorough characterisation of each device. The key parameters required to understand the behaviour of an LAPPD include the dark rate, single photoelectron gain, position resolution, and timing resolution. It is also required to have an understanding of how these parameters are affected by the operating voltages, such that an optimal configuration can be reached depending on the objectives that an experiment wishes to achieve. Many of these parameters have been measured in-house by Incom before shipping the device to confirm that the device passes basic performance requirements [192]. The measurements made in this chapter will also act as a verification of these factory measurements.

The characterisation measurements were performed on LAPPD 104, one of two devices purchased by the University of Sheffield. It is a Generation I LAPPD with stripline readout. Some basic performance characteristics for LAPPD 104, as



## 6.1. Test stand setup

---

Parameter	Value	Notes
Max MCP voltage	900 V	Recommended
Max PC voltage	300 V	Recommended
Typical operating voltages	200 V anode to MCP 825 V across MCP 200 V between MCPs	-
Mean QE	24.3 %	At 365 nm
Gain	$5.44 \times 10^6$	At 850 V MCP, 30 V PC
Dark rate	414 Hz cm <sup>-2</sup>	At 850 V MCP, 30 V PC
Transit time spread	79.2 ps	Optimal measurement
Longitudinal resolution	2.5 mm	Typical, not measured for this unit
Transverse resolution	0.8 mm	Typical, not measured for this unit

Table 6.1: A selection of performance values from Incom’s test and measurement report for LAPPD 104.

measured by Incom, is shown in Table 6.1. Additional information on LAPPD 104 is provided in Appendix A.

## 6.1 Test stand setup

An optical test stand was constructed to perform the characterisation measurements on the LAPPD. A schematic diagram of the full test stand setup is shown in Figure 6.1.

### 6.1.1 Dark box

The characterisation measurements were performed on an optical breadboard in a large dark box measuring 1.25 m × 0.65 m × 0.58 m. A cable feedthrough was constructed to pass SHV and SMA cables into the dark box for the LAPPD, as well as any other cables necessary for electrical equipment such as the laser gantry system. A Hamamatsu R9880U PMT was placed inside the dark box as a monitor for light leaks; the rate of this PMT was measured on an oscilloscope every time before any LAPPD characterisation measurements were taken. The LAPPD was affixed onto a vertically mounted 30 cm × 30 cm optical breadboard using the mounting holes on the LAPPD frame. A photograph of the setup is shown in Figure 6.2.

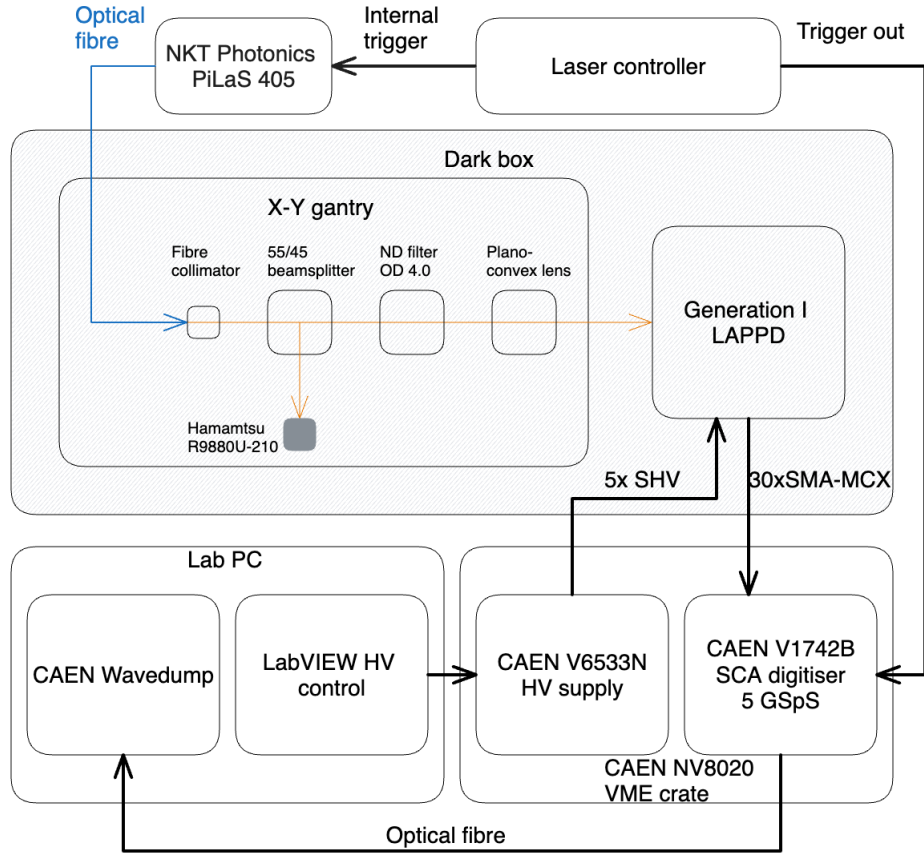


Figure 6.1: Schematic diagram of the dark box and electronics setup used for the LAPPD characterisation test stand.

### 6.1.2 Laser gantry system

The primary light source used in the characterisation measurements was an NKT Photonics PiLaS 405 [193], a 405 nm pulsed diode fibre-coupled laser with a pulse width of around 50 ps, depending on the exact settings used. The laser was chosen due to its short pulse width, which allows for probing of the LAPPD transit time spread, and the emission wavelength which is well-matched to both the peak quantum efficiency of the LAPPD bialkali photocathode and the emission spectrum of Cherenkov radiation. A 2 m single-mode optical fibre was coupled to the laser head and passed into the dark box via the feedthrough. The laser pulse passes down the optical fibre and into a fixed-focus fibre collimator with an alignment wavelength of 405 nm. The collimated light then reaches a 55/45




Figure 6.2: Photograph of the LAPPD test stand setup. The laser light enters the optical apparatus through the yellow optical fibre. The gantry system can be seen on the left and the LAPPD is on the right. It can be seen that the lower half of the LAPPD channels is connected to the digitiser whilst the upper half is terminated with  $50\ \Omega$  SMA terminators.

reflection/transmission pellicle beamsplitter. The transmitted portion was directed towards a Hamamatsu R9880U PMT, which was used as a cross-check on the laser trigger and a secondary timing reference for the arrival time of the laser pulse. The reflected portion of the pulse passes through a single neutral density filter with an optical density of 4.0, corresponding to a reduction in the intensity of the light by a factor of 10,000. In initial testing, a combination of neutral density filters with a total optical density of 4.0 was used, but this was replaced with a single filter to reduce the probability of optical non-idealities resulting from smudges or imperfect coupling between filters. The attenuated pulse then passes through a plano-convex

lens with a focal length of  $f = 15$  cm to focus the light and illuminate a spot of roughly 2 mm diameter on the LAPPD window.

The optical setup consisting of fibre collimator, pellicle beamsplitter, neutral density filter, and plano-convex lens was mounted onto the gantry system, approximately 15 cm from the LAPPD. The gantry system was composed of two linear rails such that the optical setup can be moved in two dimensions in the same plane as the LAPPD face. The working area of the gantry system was slightly larger than the LAPPD window, approximately  $25 \text{ cm} \times 25 \text{ cm}$ , but due to constraints on the positioning of the gantry rails the centre of the gantry working space and LAPPD window were not aligned. Despite this, full coverage of the LAPPD photocathode was possible. Two stepper motors control the movement of the optical setup. A NEMA17 stepper motor was used on the horizontal axis and a NEMA23 stepper motor was used on the vertical axis, which has a higher torque that is able to hold and lift the weight of the rail and optical setup against gravity. The stepper motors were each controlled by a D1 stepper motor controller produced by Igus GmbH. Movement commands were sent to the D1 controller via a Python SDK, allowing for programmatic control of the position of the optical setup, and thus the position of the laser pulse on the LAPPD. Limit switches were placed on one end of each rail which allowed for the optical assembly to be homed prior to each data taking run.

The laser is provided with a controller unit that is able to generate self triggers at a programmed frequency, or alternatively it can be triggered externally using TTL logic signals. In most measurements, the laser was set up to self trigger at a programmed repetition rate of 100 Hz, which was chosen as a good balance between the acquisition rate and ensuring that the microchannels had sufficient time to fully recharge before the next pulse, as well as ensuring that the data rate did not overwhelm the digitising electronics. The laser has a “tune” setting, which is essentially a proxy for the power output of the laser. In most cases, the tune was set to 50%.

The single photoelectron operating regime was identified using statistical arguments as in [194, 195]. At this level of light intensity, controlled by the laser

tune setting and the optical density of the neutral density filter, 92 to 94% of laser pulses resulted in no response from the LAPPD. Assuming that the number of photoelectrons produced follows a Poisson distribution, any LAPPD pulses that do occur are overwhelmingly single photoelectron ( $> 95\%$ ) and the probability of producing more than a single photoelectron is strongly suppressed. More detail is provided in Appendix B.

### 6.1.3 Electronics

The Generation I LAPPD has 56 readout channels across 28 striplines. A CAEN V1742B 32+2 channel VME switched capacitor array digitiser [196] was used as the primary data acquisition system. This digitiser is not expected to be used in the final deployment of LAPPDs in a BUTTON, but is based on the same DRS4 digitising ASIC that is under consideration. Due to channel constraints, this meant that only around half of the LAPPD channels were digitised during any given measurement. In the majority of cases, both ends of all strips in either the top or bottom half of the LAPPD were connected to the digitiser, such that there was a large continuous area of the LAPPD that was instrumented at any given time. All unused strips were terminated with  $50\ \Omega$  SMA terminators to reduce crosstalk into the active strips from reflected pulses on the unused strips. The CAEN V1742B uses multiple DRS4 switched capacitor array chips (discussed in subsection 5.6.2), partnered with Analog Devices 9222 65 MS/s octal 12-bit ADCs. The digitiser samples the input signals at a rate of 5 GS/s with a 1,024 sample buffer depth for a total acquisition window of 204.8 ns.

The laser controller issued a trigger out TTL pulse upon signalling the laser to fire with a very small timing jitter of  $< 3$  ps RMS according to the manufacturer specifications. This was used as an input to the digitiser's fast trigger input, which can trigger an acquisition for a single input group comprising 16 channels. Due to time delays resulting from cable lengths and delays in the triggering circuitry on the digitiser, as well as the inherent difference in the time at which the laser light is produced and the time at which the trigger out pulse is produced, the laser light is observed in the data approximately 40 ns after the digitising electronics initiates the trigger. This is shown in Figure 6.3. Additionally, since only primitive threshold




Figure 6.3: The time of the laser trigger and the maximum signal amplitude on the LAPPD in the 200 ns acquisition window of the DRS4. The blue peak shows the time at which the laser pulse strikes the photocathode. The standard deviation of this peak is 1.4 ns, which represents the total jitter on the triggering circuitry of both the laser and the digitiser.

triggering schemes are available on the digitiser, the position on the rising edge of the TTL pulse at which the trigger initiates shows jitter on the order of a few nanoseconds. For this reason, the time at which the trigger is initiated by the digitiser was not used for any timestamping purposes. Instead, the entire trigger signal is also digitised and a robust algorithm is used in the post-acquisition analysis to determine the trigger time from which all other timestamps are measured relative to (see subsection 6.2.3). The PMT signal at the beamsplitter could also have been digitised and used as a trigger and a timestamp for the beginning of the event. However, the PMT unfortunately did not have a sufficiently small transit time spread to provide a better timing reference than the laser controller trigger itself. A fast photodiode could potentially replace the PMT in future characterisation tests to provide a more precise secondary timestamp.

The four DRS4 chips on the CAEN V1742B were calibrated and synchronised by CAEN before shipment. Adjustments to these calibrations were not made. The exact procedure followed by CAEN to calibrate the DRS4 is not known, but there are a few known effects that occur in the DRS4 chip which must be accounted for [185]:

- Each sampling capacitor is read out by an individual buffer which has a typical DC voltage offset of 10 mV to 20 mV. These offsets must be measured and a correction applied to every waveform.
- There is a time-dependent small voltage offset (around 2 mV) caused by the changing power requirements of the chip when it swaps from sampling to readout mode.
- The gain can vary from cell to cell because of small differences in the capacitors resulting from the manufacturing process.
- The sampling interval is not the same from cell to cell. These sampling intervals are constant over time, however.

It is not known if the factory calibration provided by CAEN is the most optimal for achieving the best possible timing resolution. Future characterisation tests will seek to perform an independent calibration of the DRS4 chips to determine if any improvements can be made over the factory calibration.

High voltage was supplied to the LAPPD using a CAEN V6533N [197] VME high voltage power supply with six channels and a current limit of 3 mA. This allowed the entire LAPPD to be powered from a single supply. The high voltage was controlled and monitored using a LabVIEW interface. During this characterisation phase, the high voltage configuration was chosen to allow for maximum control over all components of the LAPPD. The current paths for each of the MCPs were isolated from one another by providing separate ground paths for each MCP's strip current. By not grounding the first MCP's strip current through the second MCP, any anomalies in either of the MCPs can be noticed quickly since the strip current of each MCP can be monitored independently. The first option to achieve this is to use a high voltage power supply with a floating ground to power the entry MCP,




Figure 6.4: The LAPPD high voltage configuration used during the characterisation measurements. Yellow lines denote connections to the LAPPD.

and use a standard ground-referenced high voltage power supply for the other. The second option is to provide a path to ground through a resistive chain from the high voltage connection at the exit of each MCP. For the characterisation tests detailed in chapter 6, the latter configuration was used. This is shown in Figure 6.4. One additional benefit of this high voltage configuration is that the photocathode voltage can be adjusted without changing any of the MCP voltages, and thus the MCP gain. This allows for investigation of the impact of photocathode voltage on the gain of the LAPPD, and also for investigation of dark pulses arising purely from the MCPs by “turning off” the photocathode by providing a positive bias relative to the first MCP so that no electrons can escape the surface. This high voltage configuration is useful during the testing and characterisation process, but would likely be replaced by a voltage divider in the eventual deployment.

## 6.2 Single pulse signal processing

A single pulse in this case refers to a waveform from a single readout channel. Signal pulses from the Generation I LAPPD are negative going but, for the sake of clarity, terminology will be used as though the pulses are positive going (e.g. *maximum* refers to the most negative point on an LAPPD pulse).



### 6.2.1 Waveform preprocessing

LAPPD waveforms underwent a series of preprocessing steps before entering into the analysis chain. Firstly, the last ten samples were removed because of an effect in the DRS4 chip whereby the final few capacitor cells show a consistent positive gradient in ADC counts. These samples were simply removed and the acquisition time of these samples was removed from any live-time calculations.

Next, each waveform must be baseline subtracted since the ADC channels have arbitrary DC offsets. Two different methods were considered and implemented. The first method calculated the mean of the waveform after removing all samples that are more than one standard deviation away from the mean, with the intention being that outliers will be ignored from the calculation and so spikes in the ADC samples, either from noise or LAPPD pulses, will not compromise the baseline calculation. This value was then removed from all samples in the waveform to apply the baseline correction. A second, simpler method was also implemented which subtracted the median value of the full waveform from every sample in the waveform. This was found to be computationally quicker, with a negligible difference on the value of the baseline.

The final step of waveform post-processing was the conversion of waveform units from ADC counts to millivolts. Since the ADCs have 12-bit resolution and a full-scale range of 1 Vpp, each ADC count below baseline was equivalent to  $\frac{1}{4096} = 0.24$  mV.

The CAEN WaveDump software outputted the waveforms for each channel in a binary data format to maximise the write speed. This binary data was decoded and preprocessed using a custom C++ module before being loaded into NumPy arrays for subsequent analysis in Python.

### 6.2.2 Peak finding

A peak can be found in the waveform by simply finding the maximum amplitude present in the waveform. However, the acquisition window of the digitiser is large in comparison to the average width of an LAPPD pulse and so it is certainly possible for multiple pulses to occur in the window. A more complex peak finding algorithm

was used which found all points above some threshold (typically 4 mV in this work) and sorted them in descending amplitude. These points were iterated over and were marked as a peak only if the time above threshold was consistent with an LAPPD pulse, and there were no other marked peaks within an LAPPD pulse width of the peak. This process ranks peaks firstly on their amplitude, but also on their distance from other peaks to ensure that electronics noise overlaid on the rising or falling edge of a peak was not itself counted as a separate peak.

### 6.2.3 Timestamping

Timing is of the utmost importance when working with the LAPPD, so it is important to ensure that there is a robust method for timestamping the arrival time of pulses. There are multiple methods that can be used for determining the timestamp of a pulse. The methods detailed in the following subsections largely follow those explored in [198], whilst LAPPD-specific applications of these methods are detailed in [184]. The optimal method is obviously that which provides the best overall time resolution, although there are other factors that will always need to be under consideration, such as the computing time required to perform the algorithm, and robustness against pulses that are misshapen by noise.

#### Constant fraction discrimination

A constant fraction discriminator measures the time at which the pulse meets some fraction of its total height. In this work, the 25% pulse height crossing time was typically used as the constant fraction threshold. This was calculated by finding the peak sample as mentioned in subsection 6.2.2, and counting the samples backwards until the sample amplitude dips below 25% of the total pulse height. The value for the crossing time was then linearly interpolated between the two samples either side of the crossing point. This method guaranteed that the timestamping was agnostic to the height of the pulse, since, due to the effect of time walk, higher amplitude pulses will reach their peak faster than smaller pulses [199]. The time walk effect is shown in Figure 6.5. Constant fraction discrimination is not the only possible method that can be used for timestamping and in fact the method described here is not a true implementation of the




Figure 6.5: The effect of pulse height on timing, known as “time walk”. Two Gaussian-shaped pulses are shown arriving simultaneously. The dashed lines indicate the crossing time for each pulse at a constant threshold (left) and a 20% height constant fraction threshold (right). The constant fraction threshold is much more robust against varying pulse height than a simple constant threshold.

algorithm which is typically implemented using analogue electronics in a timing discriminator unit. A true constant fraction discrimination algorithm splits the input signal into two, where the first part is attenuated to the desired fraction,  $f$ , of the total pulse height. The second part is inverted and delayed in order to make the point on the leading edge where the pulse is at height  $f$  align with the peak amplitude of the attenuated pulse. These two parts are subsequently added together to form a bipolar pulse where the zero-crossing point, which is independent of the initial pulse height, represents the time at which the pulse crossed the constant fraction threshold. This approach, designed for analogue circuits, adds complexity and computation time to an otherwise straightforward software-based algorithm, and has not been implemented in this work.

### Rising edge fitting

LAPPD and, more generally, photomultiplier pulses are not Gaussian, and so a naive Gaussian fit will often produce large residuals in crucial portions of the waveform




Figure 6.6: A typical LAPPD response to a single photon fitted with a log-normal function.

(i.e. in the fast rise), resulting in a deterioration of the time resolution relative to more simple methods. However, if the fit range is applied on to the rising edge of the pulse up to a point just past the peak, then this “half-Gaussian” can provide a good approximation to the pulse’s rising edge. Firstly, a Gaussian is fitted to the full range, which provides the peak time of the pulse. This peak time can then be used to define the end of the fit range that will be used to fit the rising edge. A second Gaussian is then fit only between the baseline and the peak of the pulse.

### Full waveform fitting

In contrast to only fitting the rising edge of the waveform, one can also attempt to apply a functional fit to the entire waveform. Due to the asymmetry between the rise and fall times of a typical photomultiplier pulse, and as such an LAPPD, a Gaussian fit is not well suited. Instead, a log-normal function was used. One such

parameterisation of the log-normal distribution is given as,

$$f(t) = B - \left( \frac{Q}{(t + t_0) \cdot \sqrt{2\pi\sigma^2}} \right) \cdot \exp \left( -\frac{\ln^2 \left( \frac{t+t_0}{m} \right)}{2\sigma^2} \right), \quad (6.1)$$

where  $t$  is the time sample,  $B$  is the baseline,  $Q$  is the scale and represents the charge contained in the waveform,  $m$  is the geometric mean,  $t_0$  is an arbitrary time offset that allows the waveform to shift across the  $t$  window, and  $\sigma$  represents the width of the pulse.

The log-normal function was fitted to the waveform using least-squares minimisation. An example waveform with its associated fit is shown in Figure 6.6. The function could then be evaluated at any point in time to perform timestamping. For example, this approach could also be combined with the constant fraction discrimination method by iteratively evaluating the function in small timesteps backwards from the peak until the difference between the constant fraction of the pulse height and the evaluation of the fit function is minimised.

### Comparison of methods

A comparison of these timestamping methods is detailed in section 6.5.4 when each method was used to calculate the transit time spread of the LAPPD.

### 6.2.4 Pair matching

When a pulse is produced on a stripline, it travels in both directions towards each end of the strip and is digitised by two separate readout channels. Each half of the pulse pair must then be matched up in analysis. In the case where there is only one pulse pair in the acquisition window, this process is trivial. However, the presence of multiple pulse pairs, dark noise, and electronic noise can complicate this process and so an algorithmic method of pulse matching is used. Each half of the pulse pair will share similar characteristics if they originated from the same photon. Assuming the signal is not significantly degraded when transmitted along the stripline (which is not always the case), pulse pairs that share a causal origin will have a similar pulse height and will arrive at their respective end of the strip no

greater than  $L/v$  seconds from one another, where  $L$  is the length of the stripline (20 cm for Generation I LAPPDs) and  $v$  is the signal propagation speed along the stripline. For Generation I LAPPDs, this time is typically in the region of 1 ns, assuming a typical propagation velocity of  $v \approx 0.6c$ .

These constraints on the relative pulse heights and arrival times can then be used to pair up pulses with a reasonable degree of certainty when operating in environments of low photon occupancy. For the characterisation measurements, pulses were paired together if their arrival times were within a calibrated  $L/v$  value and if the relative pulse heights were within 25%. In the rare case where there were multiple potential pairs within the  $L/v$  window, pulses were paired based on the closest matching pulse height.

## 6.3 Longitudinal position reconstruction

### 6.3.1 Method

The longitudinal position (the position along the length of a stripline) of a photon strike on the photocathode was determined using the relative arrival times of the pulses on each end of a strip, along with prior knowledge of the signal propagation velocity along the striplines. In the process of matching the pulses together in a pair, the arrival time of each pulse was calculated using the constant fraction discrimination method as described in subsection 6.2.3. The difference in these arrival times allows for calculation of the longitudinal position of the incident photon using the following equations:

$$d_L = \frac{L}{2} - v \left( \frac{t_L - t_R}{2} \right), \quad (6.2)$$

$$d_R = L - d_L, \quad (6.3)$$

where  $d_L$  and  $d_R$  are the distances of the photon from the left and right ends of the stripline respectively and  $t_L$  and  $t_R$  are the pulse arrival times at the left and right end of the stripline respectively.

### 6.3.2 Stripline propagation velocity

In order to calculate the longitudinal position of the photon in the direction of the striplines, the propagation velocity of pulses along the striplines must be known. The propagation velocity can be characterised using laser measurements by scanning the laser along the direction of the stripline and digitising several thousand waveforms at each position. For each dataset, pulses were identified on both ends of the stripline and the difference in the arrival times of those pulses was calculated using the constant fraction discrimination method. This produced a Gaussian-distributed set of measurements around the mean difference in arrival time. Repeating this over a number of different points along the striplines and plotting the mean arrival time difference against the known laser position produced a set of points that could be fitted with a linear fit. The gradient of the fit could then be used to calculate the stripline propagation velocity, as is shown in Figure 6.7. In theory, one can also calculate the velocity by only taking data at the extrema of striplines, since the time delta would then represent the time required to travel along the entire length of the stripline, but taking data along the entire length of the stripline reduces uncertainty in the results.

Using the fit in Figure 6.7, the stripline propagation velocity was calculated as  $0.59 \pm 0.04 c = 17.7 \pm 1.2 \text{ cm/ns}$ . This was in good agreement with previous characterisation of the microstripline propagation velocity of a prototype LAPPD anode of  $0.57 \pm 0.07 c$  [178]. Averaged across all laser positions for a single stripline, the longitudinal position resolution was  $4.3 \pm 0.5 \text{ mm}$ . The results of the longitudinal position reconstruction for each laser position, which uses the previously calculated stripline signal propagation velocity and the differential arrival time of the pulse at each end of the stripline, are shown in Figure 6.8. It is noteworthy that the reconstructed position distributions are less Gaussian when the photon is incident close to the ends of the stripline, likely due to edge effects where the pulse arrival time at one end of the stripline is incredibly short.

It should also be noted that there can be subtle effects on the calculation of the longitudinal position. Baseline shift can have a significant effect on the constant fraction timing technique, especially if one end of the stripline is

### 6.3. Longitudinal position reconstruction

---




Figure 6.7: Mean differential arrival time between left and right side of stripline as a function of the laser gantry position.




Figure 6.8: Mean reconstructed photon position using differential timing with  $v = 0.59c$  as a function of the laser gantry position.



baseline-shifted relative to the other. Since the constant fraction is calculated relative to the zero-level baseline, this results in each waveform being processed with effectively a different fractional threshold and since the maximum time delta is in the region of 1 ns, if the constant fraction timestamping is thrown off by just a few tens of picoseconds this can have a large effect on the calculated position.

## 6.4 Transverse position reconstruction

The position of the photon in the dimension transverse to the stripline direction was reconstructed using the charge sharing effect. The electron cloud produced during the multiplication process spreads out radially from the initial microchannel depending on the size of the gaps between the first and second MCP, as well as the distance from the second MCP to the anode. As a result, the stripline closest to the photon position on the photocathode will observe the largest induced charge, and neighbouring striplines will see a reduced charge dependent on the distance between the stripline and the photon position. By comparing the distribution of induced charges on the striplines, one can infer the transverse position of the photon. In this work, the centroid of the stripline charges was used. If the  $i$ th stripline is at a position of  $y_i$  and observes an induced charge,  $q_i$ , then the inferred transverse position,  $y$ , is given by the centre-of-charge,

$$y = \frac{\sum_i y_i q_i}{\sum_i q_i}. \quad (6.4)$$

This equation is essentially analogous to the determination of a centre-of-mass of a system of particles as the charge-weighted distances relative to the centre-of-charge sum to zero, i.e.,  $\sum_i q_i (y_i - y) = 0$ . An alternative is to fit a Gaussian to the stripline charge distribution and use the mean of the Gaussian as the reconstructed transverse position.

Figure 6.9 shows the charge distribution across the striplines for a single laser event and the Gaussian fit and centroiding methods that were used to reconstruct the transverse position of the photon. The overall resolution for a dataset of laser events is shown in Figure 6.10, demonstrating that a transverse position resolution

of approximately 1 mm was possible for single photon events.




Figure 6.9: The distribution of charge across seven striplines for a single laser event. The shaded grey areas show the positions of the striplines. The transverse position is reconstructed using both a Gaussian fit (blue,  $y = 30.67$  mm) and the centroid method (red,  $y = 30.60$  mm) shown in Equation 6.4, which agree very closely.

The transverse position reconstruction becomes more accurate if more striplines observe an induced charge above background. It is therefore beneficial if:

- The density of signal pickups is increased.
- The charge cloud is allowed to spread more radially (by increasing the microchannel pore size or the gap distances).

However, increasing the number of signal pickups means increasing the number of readout channels required and thus increasing cost. Increasing the spread of the charge cloud results in lower induced charge per pickup, assuming that the gain is unchanged, resulting in a lower signal-to-noise on each stripline and reduced timing performance. The LAPPD is designed to operate in the middle ground




Figure 6.10: The resolution of the transverse position reconstruction for laser events at a single position.

between these trade-offs. In general, the charge cloud typically induces an above-background charge across a maximum of four to five striplines. In the characterisation measurements detailed in this chapter, eight striplines were used when performing the centroiding. However, this must be adjusted depending on the circumstances the LAPPD is operating in; in a situation where it is likely that the LAPPD will observe multiple photon hits in an event, it is important to constrain the centroiding process to as few strips as possible to reduce the possibility of pulses overlapping and contaminating the centroid calculation with charges induced by other photons.

So far, the charge of the pulse on each stripline had been used in the transverse position calculations. However, if the pulse height scales linearly with the pulse integral, then computing time can be saved by simply measuring the height of the pulse rather than calculating its integral. The exact same centroiding method can then be used but utilising the pulse heights in place of the integrals.




Figure 6.11: Comparison between measured pulse height and calculated pulse integral. Deviation from linearity is seen at larger pulse heights.

Figure 6.11 shows the relationship between the calculated pulse integral and the measured pulse height. The pulse integral was calculated using trapezoidal integration, and the pulse height was simply the maximum sample value in the waveform. A linear fit was applied using a linear least-squares regression method. The resulting fit predicted the pulse integral as  $2.07h + 7.22$ , where  $h$  is the pulse height, with a Pearson correlation coefficient of 0.82. A cut was applied such that only pulses with a measured pulse height of at least 4 mV were included in the dataset. At low pulse heights, the relationship between the amplitude and integral is strongly linear. However, at larger amplitudes, there are large deviations from linearity as the pulse integral becomes larger than expected for a given pulse height. Figure 6.12 shows the residuals distribution when using the linear fit from Figure 6.11 as a proxy for the true charge as measured by the pulse integral. The presence of a tail on the negative side of the residuals distribution demonstrates the deviation from linearity at large pulse integrals.

It is worth noting that, in general, one must be careful in using a constant




Figure 6.12: The distribution of the difference between the fitted charge from Figure 6.11 and the pulse integral. An RMS residual of approximately 5 mV ns is observed.

amplitude threshold when analysing events across multiple striplines. Since the charge, and thus pulse height, is shared across multiple strips in a distribution according to the vertical position of the incident photon on the LAPPD window, it is possible to set a threshold where only photons striking directly above a stripline will trigger the readout. Photons striking in between striplines will produce (roughly) the same overall charge, but the charge is spread more evenly across adjacent striplines in which no single stripline produces a pulse with sufficient height as to cause a trigger. This will create an effect where fewer events will be seen in the midpoint between adjacent striplines.




Figure 6.13: The 10%-90% rise time distributions for different photocathode voltages at an MCP voltage of 850 V.




Figure 6.14: The 10%-90% rise time as a function of the photocathode voltage at different MCP voltages.

## 6.5 Characterisation measurements

### 6.5.1 Pulse shape

To begin with, the single photoelectron pulse shape of the LAPPD was studied using the full waveform fitting technique as described in section 6.2.3. The laser was directed at a single point on the LAPPD photocathode and single photoelectron events were recorded. For each event in which any sample crossed the basic threshold value of 4 mV, the waveform is fitted using the log-normal equation in a time window of 1 ns either side of the threshold crossing sample. The time of the threshold crossing sample was used to set the initial estimate of  $m$  from Equation 6.1.

The 10%-90% rise time,  $\tau_R$ , was measured using the calculated fits to the waveform. It is expected that as the photocathode or MCP voltage increases, the pulse should exhibit a sharper rising edge. Figure 6.14 shows the rise time as a function of the photocathode voltage for different MCP voltages, and Figure 6.13 shows the rise time distributions for different photocathode voltages. Whilst the rise time did reduce with higher operating voltages, this increase was not significant. Similarly, the shape parameters,  $m$  and  $\sigma$ , likewise did not show significant variations with changing operating voltages.

If the rise time is taken as 700 ps based on the midpoint of the data points in Figure 6.14 and using the general rule of thumb that the bandwidth can be estimated as  $0.35/\tau_R$ , the highest frequency components present in the LAPPD signal pulses are in the region of 500 MHz. This is not surprising as the CAEN V1742B digitiser has a quoted bandwidth of 500 MHz. It is therefore likely that LAPPD pulses are at least somewhat limited by the digitiser bandwidth and that some of the highest frequency components were being clipped. However, digitisation of LAPPD pulses using a Tektronix 6 series mixed signal oscilloscope [200] with 1 GHz input bandwidth suggested that this does not have a significant effect on the pulse shape.

### 6.5.2 Dark rate

Dark noise refers to pulses that are produced by a photodetector in the absence of any photons. The number of dark pulses produced per second is the dark rate.

The most prominent cause of dark noise is from the thermionic emission of electrons within the photocathode material. The electrons have a statistical distribution of thermal energies given by the Maxwell-Boltzmann distribution, some of which will be greater than the work function of the material. It is important to characterise the dark rate so that the expected background count rate is well known and can be accounted for in any subsequent analyses.

### **Data acquisition**

The LAPPD window was covered using blackout fabric material for the duration of the dark rate measurements. This was done to ensure that the only photoelectrons produced by the photocathode are as a result of thermionic emission and that any stray photons entering the dark box due to an undetected light leak would not interfere with the measurements. The LAPPD was kept sealed in the dark box for at least 12 hours to allow the photocathode to cool down from previous exposure to ambient light. After this, the MCP stages of the LAPPD were powered up for 30 minutes to allow the LAPPD to electrically stabilise. The photocathode voltage was then incrementally stepped up, allowing 15 minutes for the LAPPD to stabilise at each voltage before beginning data acquisition.

Data was acquired from both ends of the striplines, with 15 striplines connected to the digitiser at a time. The digitiser was externally triggered using a 1 kHz TTL logic signal produced by a signal generator. Each run was 15 minutes in duration, resulting in a total of 900,000 waveforms per channel being recorded. Since the acquisition window is approximately 200 ns, a single dataset represents around 180 ms of real time. The main limitation on the run time was data storage, since a single run can produce over 200 GB of data. This methodology effectively provides a random sampling of signals produced by the LAPPD, and as such the dark rate can then be calculated by counting the number of dark pulses present in the dataset and accounting for the total live time.

### **Analysis**

The analysis procedure for the dark rate measurements did not ensure that pulses are correlated between each end of a single stripline. Each 200 ns acquisition was



scanned to check if any sample in the full record is above a minimum threshold of 4 mV. The threshold was selected by observing dark pulses from a single stripline on an oscilloscope and adjusting the trigger level until no spurious triggers were occurring due to electronic noise. If a sample passed above this threshold, a second maximum threshold of 70 mV was applied to remove any records that contain large-amplitude radio frequency (RF) noise which can be picked up on the digitiser from the environment. After all the peaks were identified using the method described in subsection 6.2.2, the width of the pulse associated with each peak was calculated and is rejected if it does not lie within 0.6 ns and 1.8 ns. The number of pulses that passed all the cuts were counted, and this process was repeated for every record acquired in the dataset.

The dark rate for a given dataset was calculated using the following equation,

$$R_{\text{dark}} = \frac{N_{\text{dark}}}{t_{\text{record}} \times (N_{\text{records}} - N_{\text{rejected}})}, \quad (6.5)$$

where  $R_{\text{dark}}$  is the dark rate in Hz,  $N_{\text{dark}}$  is the number of identified dark pulses in the data,  $t_{\text{record}}$  is the record length in seconds, and  $N_{\text{records}}$  and  $N_{\text{rejected}}$  are the total number of records and the number of rejected records respectively. The denominator represents the total acquisition time of the dataset.  $N_{\text{rejected}}$  is composed of records that were rejected at the maximum threshold cut and are removed from contributing to the denominator, as it is possible that a dark pulse was contained within the record but could not be observed due to the noise.

## Results

The measured dark rates are shown in Figure 6.15. As expected, the dark rate increases with increasing photocathode and MCP voltages. Large increases in the dark rate are observed above 150 V on the photocathode with 850 V on the MCP. These measurements show generally lower dark rates than those conducted in-house by Incom, which were measured using the trigger rate on an oscilloscope. Interestingly, the Incom measurements do not show the large jump in dark rate except at much higher MCP voltages. From the results, it seems clear that 150 V on the photocathode with 850 V on the MCPs should be considered an upper limit




Figure 6.15: LAPPD 104 dark rate per unit area averaged across 32 channels (16 strips) as a function of the photocathode voltage at varying MCP voltages.

on the operating voltage when considering only the dark rates. Below this limit, the dark rates are reasonably manageable at a maximum of around 40 kHz averaged over the entire photocathode area. One area that might be concerning is the channel-to-channel variation in the dark rate, shown in Figure 6.16. It is not clear at this point if the non-uniformity is due to the photocathode itself, the application of the electric field across the photocathode gap, or methodological issues.

Future tests will also investigate the temperature dependence of the LAPPD dark rate. Previous testing of Hamamatsu R7081 PMTs has shown strong dependence of the measured dark rate on ambient temperature [201]. Given the surface area-normalised dark rates and the overall large area of the LAPPD, it is important to understand how high the dark rates may become in a real experimental environment. This will become especially pertinent if the LAPPD is deployed with on-board readout electronics in an encapsulation inside a liquid volume, where the ability to perform thorough thermal management may be very




Figure 6.16: The strip-to-strip variation in the dark rate for the upper half of the LAPPD 104. Although variation is present, it tends to remain consistent across operational voltages.

limited.

### 6.5.3 Gain

The LAPPD gain is defined by the ratio of the number of electrons observed at the anode to the number of initial photoelectrons produced. The gain was measured by illuminating the LAPPD with the laser, ensuring that the light level is attenuated down to the single photon level using the neutral density filter. The charge at the anode can be determined using the definition of the ampere as one Coulomb per second, and Ohm's law, like so,

$$I = \frac{V}{R_T}, \quad (6.6)$$

$$Q = \int I dt = \frac{1}{R_T} \int_{t_{\text{start}}}^{t_{\text{end}}} V dt, \quad (6.7)$$

where  $I$  is the output current,  $V$  is the measured voltage at time  $dt$ ,  $Q$  is the charge, and  $R_T$  is the termination resistance, in this case  $R_T = 50 \Omega$ . In practice, this measurement is performed by integrating the waveform produced by the LAPPD in response to the single photon laser light. Since the charge of the single photoelectron is given by the elementary charge,  $e$ , calculating  $Q/e$  gives the ratio of the number of electrons at the anode and photocathode. The gain values quoted in this section refer to this unitless value.

### Analysis

A dataset of 30,000 events was acquired in this setup across a series of photocathode and MCP voltages to study the effect of these voltages on the gain. For each event, a 20 ns window around the expected arrival time of the LAPPD pulse was created. The position of this window within the full acquisition window was determined by calculating the time sample of maximum amplitude for each waveform, as shown in Figure 6.3. The time window denoting the region of interest was centred around the peak of the histogram. Within this window, peaks were detected by a simple threshold cut at 6 mV and the peak height was recorded. If a peak was found, a 6 ns window consisting of 3 ns either side of the peak was constructed and used as the integration window. Trapezoidal integration was used to calculate the total charge contained within the waveform using Equation 6.7 (this method was also cross-checked by applying the full waveform fit, detailed in section 6.2.3, and extracting  $Q$ ; no appreciable difference was found between the methods aside from the improved computation speed of the numerical integration). If a peak was not found, the centre of the integration window is set by the maximum amplitude within the wider window, and these events would then contribute to the pedestal of the gain distribution. For each set voltage, a histogram representing the gain distribution was created. The pedestal, single photoelectron, and double photoelectron peaks were fitted with a sum of three Gaussian functions using RooFit [202]. The means of the single and double photoelectron Gaussians were constrained such that the mean of the double photoelectron peak must be double that of the single photoelectron peak. The mean of the single photoelectron Gaussian fit provided the gain value at the given

voltage.


### Results

The gain distributions of LAPPD 104 at varying MCP and photocathode voltages can be seen in Figure 6.17, with the associated pulse height distributions shown in Figure 6.18. Both sets of distributions show the expected characteristic series of peaks. The first peak is the pedestal, which characterises the response of the LAPPD when no light is incident on the photocathode; the width of this peak can be used as a measure of the noise present in the system. The second peak is the single photoelectron peak; the mean of this peak is shown in Figure 6.19. As expected, the gain increases with an increase in photocathode voltage. It can be seen that the gain has a stronger dependence on the MCP voltage than the photocathode voltage. This is largely expected, since the secondary emission occurs within the microchannels and the voltage at the photocathode only has an effect on the initial photoelectron, whereas the MCP voltage affects all electrons produced as a result of secondary emission inside the microchannels.


To maximise the signal-to-noise ratio of the LAPPD waveforms, the voltage across the microchannel plates should be raised as much as possible. However, this also needs to be balanced against the increase in dark rate as observed in subsection 6.5.2. The dark rate most strongly depends on the photocathode voltage, whereas the gain most strongly depends on the microchannel plate voltage, and so it is possible to control each with some small degree of independence from one another.

### Gain as a function of laser repetition rate

As explained in section 5.2, the microchannels of the MCP take some time to replenish the charge that is delocalised during the electron cascade. If the microchannel is struck again before it has fully recharged, it is expected that there will be a reduction in the number of secondary electrons emitted, thus resulting in a lower gain. Since the laser used in this characterisation was highly collimated, there was a significant chance that the laser would repeatedly strike the same collection of microchannels on the LAPPD. By increasing the repetition rate of the laser, the effect of the microchannel recharge time on the gain of the LAPPD




(a)



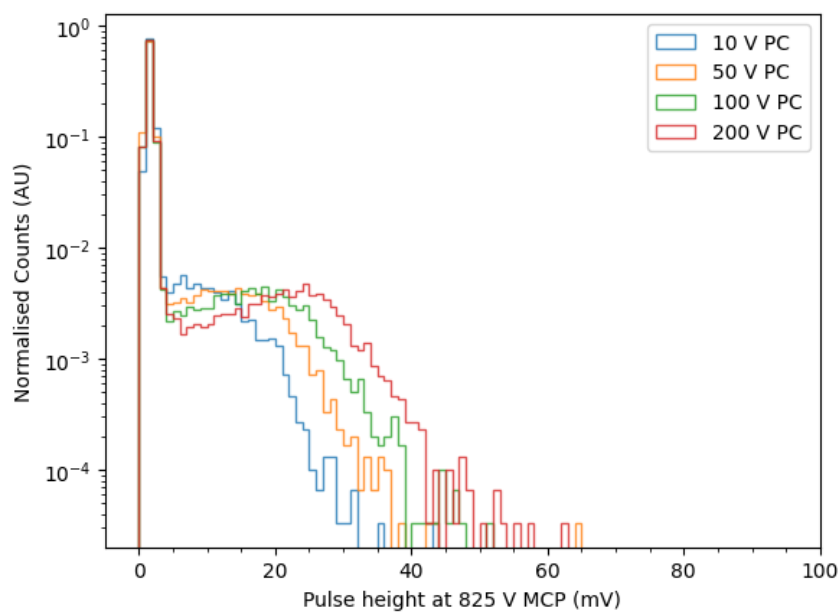
(b)

Figure 6.17: The LAPPD gain distribution at (a) 200 V on the photocathode at varying MCP voltages and at (b) 825 V on the MCPs with varying photocathode voltages.

## 6.5. Characterisation measurements




(a)




(b)

Figure 6.18: The LAPPD pulse height distribution at (a) 200 V on the photocathode at varying MCP voltages and at (b) 825 V on the MCPs with varying photocathode voltages.



(a)



(b)

Figure 6.19: The LAPPD mean gain at (a) varying photocathode voltages and at (b) varying MCP voltages



response can be investigated. With the laser illuminating a constant position on the photocathode, the repetition rate was gradually increased from 100 Hz to 1 MHz. The gain of each event was calculated using the same method as previously described.




Figure 6.20: LAPPD gain as a function of the laser repetition rate. The laser position is kept constant such that a single point on the photocathode is repeatedly illuminated.

The LAPPD gain distribution at each value of the laser repetition is shown in Figure 6.20. The distributions at a laser repetition rate of 100 Hz and 1 kHz are very similar, with differences resulting only from the statistics present in each dataset. This is expected as the microchannels are known to take around 1 ms to recharge, and so up to a repetition rate of 1 kHz no change in the gain is expected. At repetition rates above 1 kHz, the gain rapidly begins to reduce. At a repetition rate of 1 MHz, the mean single photoelectron gain is reduced by approximately a factor of three relative to the nominal value.

Overall, this feature of the LAPPD does not pose a major concern for its application in neutrino experiments, since it is unlikely that the photon occupancy

will ever be high enough in such a concentrated area on the photocathode for the microchannel recharge rate to have any influence. However, it may become a concern for detectors in high luminosity collider experiments where very high rates may be expected. In this case, care should be taken to select LAPPDs with low resistance MCPs such that the microchannel recharge time can be reduced as much as possible.

### 6.5.4 Transit time spread

The transit time is the time between a photon striking the photocathode and the production of an electrical signal at the anode. It measures the time required for the electron multiplication to take place. In general, the larger the photomultiplier or the greater the number of multiplication stages, the higher this value will be. The transit time spread measures the fluctuation in the transit time, and is the primary metric for describing the time resolution of a photodetector.

The data acquisition process was largely identical to the process used for measuring the gain. However, in this case the trigger waveform from the laser was also digitised for post-acquisition timing analysis.

#### Analysis

Pulses were identified using the same pulse finding methodology as is described in subsection 6.5.3. Regions of interest were defined for both the LAPPD and trigger waveforms depending on their characteristic location within the acquisition window. Within these regions of interest, the LAPPD and trigger waveforms were cubic-spline interpolated to upsample each by a factor of ten. The waveforms were then timestamped using the constant fraction discrimination method with the threshold set at 50% of the total pulse height. The time difference between these timestamps was then calculated; these form a Gaussian-shaped distribution where the mean is the average time difference between the LAPPD pulse and the laser trigger. With the various delays resulting from cable lengths and trigger circuitry, this value is mostly arbitrary. If these delays were to be thoroughly characterised and subtracted, which is not performed here, this value would represent the transit

time of the LAPPD itself. The standard deviation of the distribution represents the transit time spread and provides the best measurement of the timing resolution of the LAPPD.

However, this distribution is inclusive of uncertainties in the timing from the laser itself, as the laser has a non-trivial pulse width ( $<60$  ps) that is comparable to the estimated value of the transit time spread. Furthermore, the trigger out signal from the laser driver unit has an associated jitter. This value is claimed by the manufacturer to be  $<3$  ps RMS. The measurements of the transit time spread can be “corrected” for the laser pulse width and laser trigger jitter by subtracting their contributions to the uncertainty in quadrature. We assume that the measured time resolution is given by,

$$\sigma_{\text{measured}}^2 = \sigma_{\text{LAPPD}}^2 + \sigma_{\text{laser width}}^2 + \sigma_{\text{laser trigger}}^2, \quad (6.8)$$

and so by subtracting out these contributions in quadrature we can estimate the true transit time spread of the LAPPD,  $\sigma_{\text{LAPPD}}$ . In reality, there will be numerous other contributions that broaden the measured transit time spread. The most significant of which, not including the laser pulse width, is the variation in the DRS4 cell timing widths which can result in an overall timing resolution in the region of 20 ps unless a thorough timing calibration of each cell of every channel is performed, which is not done here. A Gaussian distribution is fitted to the data to determine the standard deviation and thus the transit time spread. Examples of the timing distributions at various photocathode voltages can be seen in Figure 6.21.

## Results

The measured transit time spread as a function of the photocathode voltage can be seen in Figure 6.22, which shows the expected behaviour that increasing the photocathode voltage results in a reduction in the transit time spread. The measured transit time spread improves from a value of 186 ps at 10 V on the photocathode to 61.5 ps at 200 V on the photocathode, after correcting for a conservatively small laser pulse width of 30 ps. The photocathode voltage is expected to be the most important parameter in relation to the transit time spread




Figure 6.21: The LAPPD transit time distributions at various photocathode voltages. The standard deviation of the distribution, as well as the mean time delta between the LAPPD pulse and the laser trigger, reduces with increased photocathode voltage.

since the measured transit time spread is dominated by the time taken for the initial photoelectron to trigger the electron cascade within the first microchannel. An extended tail is observed, especially at higher photocathode voltages, likely as a result of late pulsing. It is expected that some proportion of photoelectrons will collide with the area between the pores on the first microchannel plate, resulting in a secondary timing distribution centred  $\mathcal{O}$  1 ns after the primary timing distribution. It is observed that this secondary distribution becomes more prominent at higher photocathode voltages, indicating that, whilst a higher photocathode voltage reduces the time required to initiate the electron cascade in most cases, it also increases the chance that the photoelectron will miss a microchannel in the first microchannel plate.

The uncertainty arising from the trigger out from the laser driver unit could be removed by the addition of another photosensor to observe the laser and act as a reference timestamp. The beamsplitter would divert part of the laser pulse to the




Figure 6.22: The transit time spread values as a function of photocathode voltage at an MCP voltage of 825 V. The values taken are the standard deviation of the timing distribution at each voltage. Also shown is the corrected transit time spread value with a laser pulse width of 30 ps subtracted out in quadrature.

reference photodetector whilst the other light continues through the neutral density filters and onto the LAPPD photocathode. The reference photodetector will see a large pulse from the unattenuated laser light and a constant fraction discriminator can be used to generate a timestamp from the resultant waveform. Since the light pulse is large, this photodetector will see a statistical reduction in the transit time spread relative to the single photon transit time spread, allowing for lower uncertainty on the reference timestamp. The trigger from the laser driver unit is still used to trigger the digitiser, but the timestamp to find the time difference comes from the reference photodetector instead of the trigger itself. This cross-check was not performed here, but should be implemented in future characterisation measurements of the timing response, and could potentially result in even better timing resolution than was achieved in this work.

Method	Measured transit time spread (ps)
Gaussian fit	$81.0 \pm 1.8$
Rising edge fit	$77.0 \pm 1.6$
Constant fraction discriminator	$70.9 \pm 1.3$
Log-normal fit	$68.8 \pm 1.3$

Table 6.2: Comparison of transit time spread measurements using different timestamping methods at 850 V applied to the MCPs and 100 V applied to the photocathode.

### Comparison of timestamping methods

Table 6.2 shows a comparison of measured transit time spreads using the different timestamping methods detailed in subsection 6.2.3. The best timing resolution was achieved when using the full log-normal template fitting method as would be expected. Constant fraction discrimination was ultimately used for timestamping in most cases as the timestamping resolution was only slightly worse than the full template fitting but was significantly faster to perform computationally.

### Afterpulsing

Afterpulsing is the production of a delayed pulse (distinct from, and more delayed, than late pulsing) resulting from ionised gas molecules striking the photocathode, and is another important component of the LAPPD characterisation. Afterpulses typically have amplitudes of up to a few photoelectrons, and appear on characteristic timescales after the production of the initial nominal pulse. This timescale is dependent on the mass and charge of the gas molecule that is ionised, much like in a mass spectrometer. An example of an afterpulse is shown in Figure 6.23.

Understanding the afterpulsing characteristics of the LAPPD is important for building robust event reconstruction algorithms. For example, timing-based vertex fitters for water Cherenkov neutrino experiments require prior knowledge of the typical hit time distribution in order to perform vertex reconstruction (as was discussed in subsection 4.4.2). The time distribution of afterpulsing should be included in the hit time distribution when performing the likelihood maximisation.




Figure 6.23: An example showing a nominal single photoelectron pulse at  $t \approx 90$  ns with an afterpulse at  $t \approx 150$  ns. The waveforms from both the left and right side of the stripline are shown.

The afterpulse ratio, the ratio of the number of identified afterpulses to the number of single photoelectron pulses, was characterised using the same laser dataset. Afterpulses were identified by searching a time window from the end of an observed single photoelectron pulse to the end of the acquisition window for any pulse that was of at least the same amplitude as the single photoelectron pulse. The absolute amplitude, the amplitude relative to the single photoelectron pulse amplitude, and the timestamp were then calculated. The dependence of the afterpulse ratio on the photocathode voltage is shown in Figure 6.25. As is expected, the afterpulse rate increases with increasing photocathode voltage, although not with a clear trend, as shown in Figure 6.24. An afterpulse rate of around 4% at operational voltages is in good agreement with tests performed by the manufacturer and MCP-PMTs more generally.

The timing distribution of afterpulses is shown in Figure 6.25. The dominant




Figure 6.24: The percentage of total pulses that result in an afterpulse as a function of the photocathode voltage for different MCP voltages.

population of afterpulses is observed at a time of approximately 60 ns relative to the initial pulse, with a subdominant population at around 110 ns. These distinct groupings of afterpulses are likely a result of the release of different ions, however identifying these ions based on the afterpulse timing was not explored as part of this work.

### 6.5.5 Spatial characterisation

In addition to the characterisation measurement of the aforementioned parameters at a single position of the LAPPD, it is also of interest to know how the value of these parameters varies across the entire surface of the LAPPD. The laser was scanned across the surface of the LAPPD in increments of 10 mm in the  $x$  direction and 5 mm in the  $y$  direction. Data was acquired with the laser at a repetition rate of 1 kHz for a duration of 15 s, resulting in a dataset of 15,000 events at each position. As the digitiser can only trigger half (16) of its channels using the fast trigger, the group of channels being triggered was swapped when the scan reaches half of the






Figure 6.25: The relative arrival time and amplitude of afterpulses identified after a laser-induced single photoelectron pulse. Note that there are likely to be additional populations of afterpulses at larger relative arrival times, but these were not observed as they would be beyond the acquisition window of the digitising electronics.

total height. The scan row at  $y = 55$  mm was nearest to the dividing line between channel groups and so is repeated in both channel groups such that the neighbouring strips above and below that position are both included in the total dataset. The two halves of the scan were then stitched together. The scan data was taken with 850 V applied to the MCPs and 100 V applied to the photocathode.

### Gain

At each laser position, the gain distribution was calculated for all striplines in the channel group. The strip closest to the incident photon position was determined from observation of the measured gain distributions. If the incident light was close to a single strip, demonstrated by a much extended gain distribution relative to other strips, the gain values from that strip and the two neighbouring strips were




Figure 6.26: The variation in the median gain across the lower half of the LAPPD photocathode.

summed. If two strips show similar distributions, then those two strips along with their neighbouring strips above and below were summed. Bicubic interpolation was performed across all the points to produce the heatmap shown in Figure 6.26.

The median gain is observed to be uniform across the measured area of the photocathode surface to within around 30%.

### Transit time spread

At each laser position, the distribution of the time delta between the laser trigger and the waveform timestamp was calculated for every strip. Similarly to the gain mapping, the closest strip to the incident laser light and its neighbouring strips were selected, and the average transit time spread across these strips was calculated and plotted in the heatmap shown in Figure 6.27. Figure 6.28 shows the time delta distribution at several positions along the same strip, demonstrating that there is effectively no variation in performance within a single stripline. The horizontal




Figure 6.27: The variation in the transit time spread across the lower half of the LAPPD photocathode. The dark region in near the centre corresponds to the position of the linear spacer, and as a result very few viable waveforms are produced that can be analysed.

dark bar feature present in Figure 6.27 seemingly corresponds to the position of the linear spacer. Data taken at this  $y$  position showed very few viable pulses, and those that could be analysed showed very poor time deltas. Ignoring the data at this  $y$  position, overall variation in the transit time spread across the measured area of the photocathode was also around 30%.




Figure 6.28: Transit time spread distributions for several points along the same stripline, demonstrating that there is very little variation in the distributions.

## 6.6 Chapter summary

The key results of the characterisation measurements are summarised in Table 6.3. Overall, we see a good level of agreement between the factory measurements and these characterisation tests when comparing like-for-like operating conditions for most parameters. The largest discrepancy occurs in the measured dark rate of the LAPPD. This is most likely a result of methodological differences. It is thought that Incom measures the dark rate through observation of the trigger rate after connecting a selection of striplines to an oscilloscope. In contrast, the methodology detailed in these measurements is much more robust, using random sampling of the LAPPD waveforms and pulse selection after the data acquisition. Additionally, it is also possible that the protocol followed in these measurements kept the LAPPD in a light-tight environment for a longer period of time before data acquisition. This can have a large impact on the dark rate since photocathodes have a non-trivial “cool-down” duration after being exposed to ambient light.

- An optical test stand was set up consisting of dark box, a pulsed picosecond

## 6.6. Chapter summary

---

Parameter	Incom value	Test stand value	Notes
Mean QE	24.3 %	-	At 365 nm
Dark rate	414 Hz cm <sup>-2</sup>	44.9 Hz cm <sup>-2</sup>	At 850 V MCP, 30 V PC
Gain	5.44 × 10 <sup>6</sup>	5.11 × 10 <sup>6</sup>	At 850 V MCP, 30 V PC
Transit time spread	79.2 ps	68.4 ps	Optimal measurement, uncorrected
Longitudinal resolution	2.5 mm*	4.3 mm	*Typical, not measured for this unit
Transverse resolution	0.8 mm*	1.0 mm	*Typical, not measured for this unit
Afterpulsing rate	-	1.9%	At 850 V MCP, 30 V PC

Table 6.3: A selection of performance values from the characterisation measurements of LAPPD 104.

laser on an X-Y gantry system, high voltage power supply, and 5 GS/s switched capacitor array digitiser.

- A procedure for pulse processing was implemented to detect pulses on the striplines of the LAPPD and correlate the signals from each side of the stripline.
- Methods were developed for reconstructing the position of the photon in the directions transverse and longitudinal to the stripline position.
- Multiple methods were considered for producing an accurate timestamp for ultrafast photosensors such as the LAPPD, with constant fraction discrimination determined to be the most practical algorithm whilst providing good timing resolution.
- Fundamental parameters of the LAPPD including the gain, dark rate, and transit time spread were characterised and found to be in good agreement with measurements made in-house by Incom.
- Laser scans over large portions of the surface area of the LAPPD found good uniformity of gain and transit time spread.
- Single-ended readout was considered as an option. The ability to correlate pulses deteriorated as was expected, but this option could be practical in the scenario where instrumenting all the readout channels of the LAPPD proves too costly.

# Chapter 7

## Multiphoton disambiguation with LAPPDs

### 7.1 Motivation

There are significant challenges to operating LAPPDs in environments of high photon occupancy. Unfortunately, many applications where LAPPDs may yield additional insights, such as neutrino beam experiments like ANNIE or experiments investigating novel scintillating media such as BUTTON, are likely to result in high photon occupancies across the LAPPD. The difficulty arises because, as more photons are present, interpreting the signals from the LAPPD becomes more challenging as signals will begin to overlap, especially if the photons are incident on the photocathode within a narrow time window. This overlapping occurs in both longitudinal and transverse directions as charge resulting from distinct initial photons can be collected onto the same series of striplines. In order to retain the temporal and spatial resolution offered by the LAPPD in high photon occupancy environments, a robust algorithm is required to disambiguate individual photon hits from the overall detector response. The algorithm should determine how many photons are present, and subsequently where and when the photons struck the photocathode.

There are multiple potential methods to approach this problem, which are all

based on the principle of isolating a known signal or series of signals that combine to form the full detector response that is observed. This chapter reports on a preliminary study investigating the effectiveness of a deconvolutional method of multiphoton disambiguation on the LAPPD using a simulated detector response developed using the detector characterisation detailed in the previous chapter. This method is evaluated using simulated photon hits resulting from liquid scintillator. A few possible alternative approaches are summarised and contrasted with the deconvolutional method.

## 7.2 LAPPD waveform simulator

To assist in developing this multiphoton disambiguation method, a waveform simulator was created to emulate the response of the LAPPD to a given input. The waveform simulator can take a series of photon hits on the LAPPD window, each made up of a longitudinal position (*along* the stripline),  $x$ , a transverse position (*across* the stripline),  $y$ , and a hit time,  $t$ , and generate an approximation of the complete LAPPD detector response across all striplines. The simulator was designed such that it can take arbitrary user input, or the output of photon hits from the RAT-PAC Monte Carlo simulation package. Waveforms are generated using the same parameterisation of the log-normal function as is described in section 6.2.3, using outputs from the simulation as inputs into the waveform generation.

The basic process for the simulating the LAPPD response from a single photon is as follows:

1. The photon hit time is smeared with a Gaussian distribution representing the transit time spread of the LAPPD as measured in the characterisation process. This has been implemented using an adaptation of the PMT timing response simulation in RAT-PAC.
2. A charge value is drawn from a probability distribution function (PDF) representing the single photon charge response of the LAPPD. This has been implemented using an adaptation of the charge response simulation for

PMTs in RAT-PAC.

3. The distance to the left and right ends of the LAPPD striplines is calculated from the photon's  $x$  position on the photocathode. The stripline signal propagation velocity is used to calculate the time delay for the pulse to reach each end of the stripline. This is then used as the timestamp for each generated pulse.
4. The distance from the photon's  $y$  position to every stripline is calculated. The charge cloud produced from the exit of the bottom MCP is assumed to have a Gaussian profile, with a two sigma width defined by the distance between the bottom MCP and the anode, in this case 4 mm. This distribution is used to determine the ratio of the total produced charge that is induced onto each stripline. It is assumed that the entirety of the charge cloud is deposited onto the striplines.
5. A log-normal function is used to emulate the shape of the LAPPD single photon response. The median of the log-normal is given by the timestamp and the integral is given by the charge induced on the stripline, after conversion from picoCoulombs to millivolt-nanoseconds. This is repeated for every stripline.
6. The log-normal function is sampled at the 5 GS/s sampling rate of the digitising electronics, and any associated jitter in the digitising electronics can be added here if desired.
7. An arbitrary time offset can be added at this point to move the pulse within the acquisition window to represent the pre- or post-trigger window of the digitising electronics.
8. If multiple photons are present in the event, this process is repeated for every photon. The total LAPPD response is given by the element-wise stacking of the entire  $(n \times m)$  matrix, where  $n$  is the number of striplines and  $m$  is the number of samples in the acquisition window.
9. Gaussian noise of a given RMS value is generated on top of the entire LAPPD response to represent the electronic noise, if desired.



The simulator does not account for the dead regions due to the linear spacers, or more complex detector effects, such as crosstalk or overshoot which will be present in a real system. These detector effects could be added at a later date to produce a much more thorough simulation of the detector response, although any electrical effects may require more advanced modelling through the use of a finite-element electrostatics simulation tool, such as Garfield++ [203]. This has not been performed here. Afterpulsing is also not included at this time, but could be incorporated in a relatively straightforward manner by using the afterpulse ratio to determine whether an afterpulse should be generated. The time after the initial pulse and the relative amplitude could be drawn from probability distribution functions created using the afterpulse characterisation data in section 6.5.4. A benefit of this approach to waveform simulation is that each simulated event is completely independent, and as such multiprocessing techniques can be used to vastly improve the simulation speed by processing multiple events in parallel.

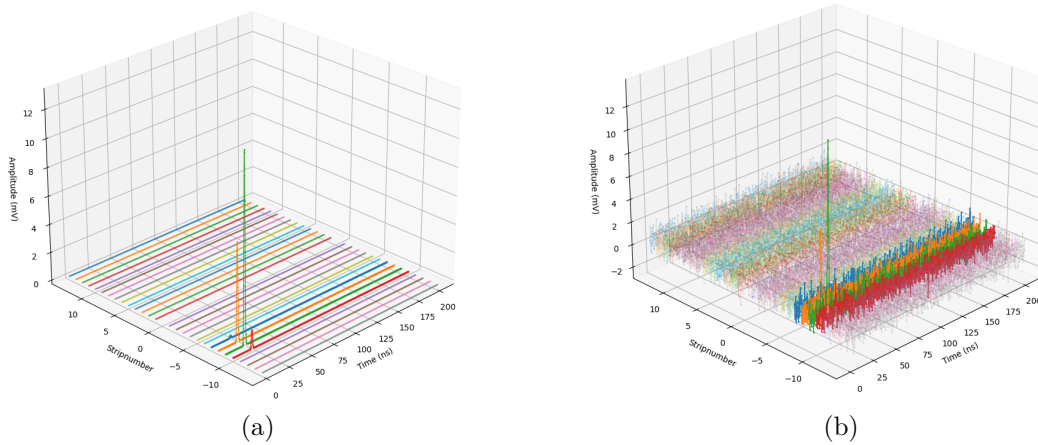


Figure 7.1: The full simulated LAPPD response for a single photon incident on the LAPPD with (a) no noise added and (b) a realistic 0.7 mV RMS Gaussian noise simulated and overlaid. All but the four striplines with appreciable deposited charge are made transparent to aid visualisation.

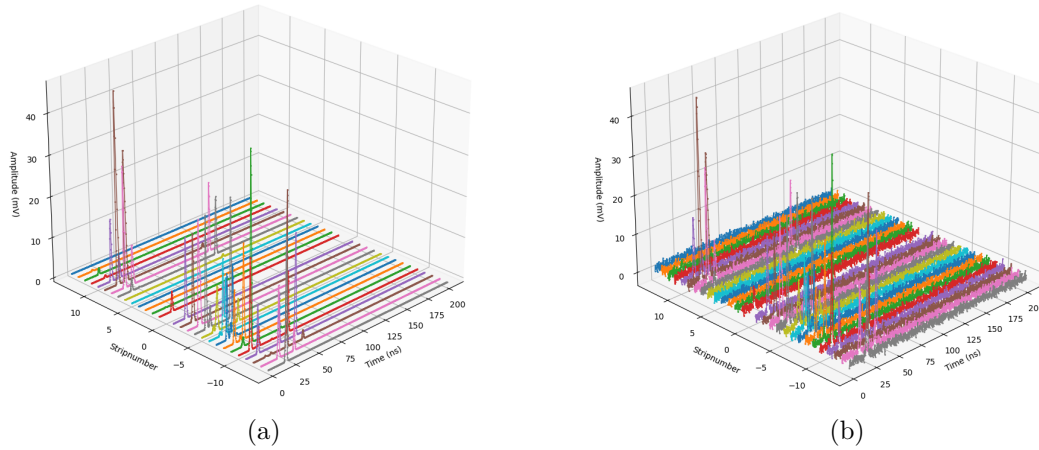


Figure 7.2: The full simulated LAPPD response for twenty photons incident on the LAPPD with (a) no noise added and (b) a realistic 0.7 mV RMS Gaussian noise simulated and overlaid.

## 7.3 Deconvolutional method

This deconvolutional method was first proposed to be used for disambiguating photon hits on the LAPPD in [204]. In this method, the waveforms from each of the 28 striplines are stacked as if to form an image with a height of 28 pixels, and a width given by the acquisition window of the digitiser, in this case 1,024 samples. The voltage amplitude of the waveform at each sample on a stripline gives the value of that “pixel”.

### 7.3.1 Wiener deconvolution

The task now is to search for features in the image that bear resemblance to a pulse from an LAPPD. Wiener filtering is a very common image processing technique used for the purpose of removing noise from an image. Wiener deconvolution applies the Wiener filter to the deconvolution process, which can frequently leave noisy artefacts. The process utilises a template image, known as a point spread function (PSF), with which the image is deconvolved. In this case, the template image represents the single photon response of the LAPPD. The PSF can be created either using an average of many thousands of confirmed LAPPD single photon pulses, or by a model of the LAPPD single photon response. An example of a modelled PSF for




Figure 7.3: 3D image of the point spread function used to model the LAPPD single photon response.

the LAPPD single photon response is shown in Figure 7.3.

Wiener deconvolution functions by taking the fast Fourier transform of both the image and the PSF to transform both into the frequency domain. Wiener filtering is performed in the frequency domain by applying the following equation,

$$G(f) = \frac{|H(f)|^*}{|H(f)|^2 + K}, \quad (7.1)$$

where  $G(f)$  is the transfer function of the Wiener filter,  $H(f)$  is the Fourier transform of the PSF and thus  $|H(f)|^2$  is the power spectrum of the frequency-domain representation of the PSF, and  $K$  is the regularisation parameter that represents the noise-to-signal power ratio and controls the trade-off between suppressing noise and retaining the original image. This is then deconvolved by dividing the filtered frequency domain representation of the

### 7.3. Deconvolutional method

---




Figure 7.4: Example of the deconvolution algorithm applied to a simulated waveform on a single LAPPD stripline. Pulses that were previously overlapping and difficult to analyse become distinct, and the baseline noise is removed.

original image with the PSF,

$$Y(f) = \frac{G(f)X(f)}{H(f)}, \quad (7.2)$$

where  $Y(f)$  is the deconvolved image in the frequency domain and  $X(f)$  is the original image. The deconvolved image is then converted back into the time domain with an inverse fast Fourier transform. The process suppresses regions of the image which do not resemble the template, whilst preserving regions which do. The `wiener` function from the SciKit-image Python library is used to implement this methodology in this work which has the advantage of utilising efficient NumPy and SciPy code for improved computational performance. An example of the deconvolution algorithm applied to the waveform from a single stripline is shown in Figure 7.4.

There is lots of room for optimisation in this step, from the value of the regularisation parameter to the exact form of the template which used as the PSF. That optimisation has not been performed to the fullest possible extent in this

work and there is still much room for improvement. Furthermore, the optimal parameters and PSF used in the deconvolution process are likely to change somewhat depending on the operating voltages of the LAPPD. It is worth noting that although the deconvolution process has the ability to isolate pulses in the presence of noise or other effects, it will not be able to pick out pulses that have been misshapen due to bandwidth limitations at the digitisation stage. This is because if the signal is bandwidth-limited then those frequency components will not be present in the resulting waveforms. Therefore, it is still important to emphasise the selection of the most appropriate digitising electronics as has been discussed in subsection 5.6.2.

After deconvolution, the entire image is cubic interpolated across the 2D grid of pixels to upsample in both the transverse and longitudinal direction by a factor of 10. Like in chapter 6, this is performed so that, when timestamping is performed later in the process, the achievable time resolution is better than the sampling period of the digitiser.

#### **7.3.2 Peak finding**

As the deconvolution processes should ideally suppress most of the noise, any pulses are very obvious. A two-dimensional maximum filter is used which scans over the entirety of the image with a rolling window of a similar two-dimensional size to the PSF and finds areas of local maxima. A baseline threshold is set to ensure that any remaining small spikes in the noise or artefacts resulting from the deconvolution process are not found as peaks by the maximum filter. The typical level for this threshold can be as low as 2 mV due to the noise suppression resulting from the deconvolution. The result of this process is a list of identified peaks on both the left and right sides of the LAPPD.

#### **7.3.3 Pair matching**

Unlike in subsection 6.2.4, events analysed using this method have the potential to contain many pulses, and so the assumption that a left-right pulse pair is correlated as long as the peaks fit within the time constraint of the stripline propagation

velocity breaks down. Instead, a maximum likelihood method is used to pair left and right peaks. The output of the maximum filter is used to create a matrix of  $N$  left-sided peaks  $\times$   $M$  right-sided peaks. Each of these peaks encodes three important values: the time sample in which the peak occurred, the interpolated stripline on which the peak exists, and the amplitude of the peak in volts. PDFs are generated for the difference between the left and right values for each of these variables which are then used to perform a maximum likelihood calculation. These PDFs are produced using empirical data from the LAPPD characterisation process.

The PDF representing the difference in arrival times between left and right peaks is produced using the calculated stripline propagation velocity. If the difference in the time of arrival between a left and right peak is more than the time required for a pulse to propagate the entire length of the stripline, then there is zero probability that the peaks originate from the same causal event. Thus, the PDF looks like a step function. The reason the PDF is not a perfect vertical step is because there is timing uncertainty arising from the digitising electronics and the LAPPD itself. The functional form of the PDF is,

$$\Delta_t = 1 - \frac{1}{2} \left( 1 + \operatorname{erf} \left( \frac{|t_L - t_R| - D/v}{\sqrt{2}\sigma_t} \right) \right), \quad (7.3)$$

where  $t_L$  and  $t_R$  are the left and right arrival times respectively,  $D$  is the length of the stripline,  $v$  is the stripline propagation velocity, and  $\sigma_t$  is the timing uncertainty, assumed in this work to be 10 ps. The PDF is shown in Figure 7.5.

The PDF representing the difference in interpolated stripline position between a left and right peak is a Gaussian distribution. A left and right peak are more likely to have originated from the same photon if there is a smaller difference there is in their interpolated stripline position. The width of the Gaussian is derived from the uncertainty in the calculation of a photon's  $y$  position from the single photon analysis during the characterisation studies. The functional form of the PDF is,

$$\Delta_y = \frac{1}{\sqrt{2\pi}\sigma_y} \exp \left( -\frac{(y_L - y_R)^2}{2\sigma_y^2} \right), \quad (7.4)$$

where  $y_L$  and  $y_R$  are the position of the left and right peak in the direction




Figure 7.5: The PDF used for matching pulses based on their relative time of arrival.

perpendicular to the stripline respectively.  $\sigma_y$  is the uncertainty of the  $y$  position of the photon, which can be tuned to allow matching between a wider or narrow selection of striplines. The PDF is shown in Figure 7.6.

The PDF representing the difference in measured amplitude between a left and right peak is also a Gaussian distribution. A left and right peak are more likely to have originated from the same photon if their amplitudes are very similar. The presence of electronic noise and potential attenuation of the pulse as it travels along the stripline can both have an impact on the observed amplitude of the waveform, and the width of the Gaussian distribution takes this into account. The functional form of the PDF is,

$$\Delta_a = \frac{1}{\sqrt{2\pi}\sigma_a} \exp\left(-\frac{(a_L - a_R)^2}{2\sigma_a^2}\right), \quad (7.5)$$

where  $a_L$  and  $a_R$  are the amplitudes of the pulses on the left and right side of the stripline, respectively.  $\sigma_a$  is the uncertainty of the amplitude of the pulse, which can be tuned to allow matching between a larger or smaller difference in pulse amplitudes. The PDF is shown in Figure 7.7.

A likelihood is formed from the combination of these three PDFs for every possible combination of left and right peaks. These likelihood values are used to




Figure 7.6: The PDF used to match pulses based on their relative inferred location perpendicular to the stripline direction.

create an  $N \times M$  matrix, as shown in Equation 7.6, where  $N$  is the number of left-sided peaks and  $M$  is the number of right-sided peaks. The matrix is iterated over and on each iteration the highest likelihood value is selected. The left and right peak corresponding to that likelihood value are paired together, and that row and column of the matrix are removed from further consideration. This iteration occurs until all left-sided and right-sided peaks are matched, or until the likelihood value falls below some minimum threshold below which peaks are considered to not be correlated. In this case, unmatched peaks are left unresolved.

$$\Lambda = \begin{bmatrix} \Lambda_{11} & \dots & \Lambda_{1M} \\ \vdots & \ddots & \vdots \\ \Lambda_{N1} & \dots & \Lambda_{NM} \end{bmatrix} \quad (7.6)$$

Once the pair matching is complete, the pair information is used to compute the position and timestamp of each identified photon hit. For each half of the pair, a constant fraction discriminator is applied to find the time of arrival of the pulse. The average time of arrival between the left and right pulses provides the photon arrival timestamp. The difference between the times of arrival on the left and right side of






Figure 7.7: The PDF used for matching pulses based on their relative amplitude.

the stripline provides the  $x$  position, once it has been converted to a distance based on the characteristic stripline propagation velocity. The  $y$  position is calculated using the centre-of-charge technique, where 51 interpolated striplines are used in the calculation, comprising the stripline where the photon was first identified, and the 25 interpolated striplines above and below this point.

## 7.4 Tests on simulated data

To validate the performance of the multiphoton reconstruction, a RAT-PAC simulation was used to generate a realistic model of photon hits on the LAPPD. The  $20 \times 20 \text{ cm}^2$  active area of the photocathode is simulated, along with a small volume of linear alkylbenzene at a distance of 15 cm. 5 MeV electrons are generated in the direction of the scintillating volume and produce scintillation light. The linear alkylbenzene is fully modelled by RAT-PAC with realistic light yield, emission spectrum and time spectrum. The study was conducted using a liquid scintillator similar to what is under consideration to be used in cocktails of WbLS that may be deployed in BUTTON in the future. Using liquid scintillator provides a high light yield which allows for the study of the disambiguation algorithm up to very high photon occupancies. Additionally, since the liquid scintillator is

accurately modelled in terms of its emission time profile, a realistic temporal distribution of photon hits on the LAPPD can be achieved, resulting in a more representative idea of what could be achieved by the algorithm in a real detector.

The LAPPD is modelled using an adaptation of the same `GLG4PMTOpticalModel` parameterised simulation model as was used to model the PMT response in chapter 4. The differences of note from the original code are a re-implementation of the geometry that defines the “envelope” region of the model to account for a square photodetector geometry, as well as switching off any magnetic effects to simulate the improved magnetic tolerance of the LAPPD relative to the PMT. A `G4SensitiveDetector` is assigned to the LAPPD geometry such that all photons hits on LAPPD are stored, allowing for their position and timing information to be used to generate the detector response post-simulation.

The study investigated the impact of the number of incident photons on the ability to reconstruct photon hits on the LAPPD surface. Each simulated event produces several hundred scintillation photons. Of the total number of photons that hit the LAPPD, a random subset of  $N$  photons are sampled to be used in the generation of the LAPPD detector response, where  $N$  ranged from 1 to 100. A total of 1,000 events were generated for each value of  $N$ , with a different subset of photons chosen each time. The detector response is generated, and the reconstruction algorithm is applied. The first metric is the “efficiency” — the ratio of the number of reconstructed hits and  $N$ . In addition, the position resolutions in the longitudinal ( $x$ ) and transverse ( $y$ ) dimensions are calculated, as well as the timestamp resolution which is defined as the difference from when the photon struck the photocathode in simulation and the timestamp calculated by the constant fraction discrimination algorithm. The timestamping resolution will have a natural floor at the transit time spread used in the simulation.

All information connecting a photon generated in Monte Carlo to a reconstructed hit is lost during the deconvolution process. At low photon occupancies the connection can be obvious, but at high occupancies a reconstructed hit may be made up of the contributions from several photons in the

#### 7.4. Tests on simulated data

---

N photons	$x$ resolution (mm)	$y$ resolution (mm)	Timestamp resolution (ps)	Efficiency (%)
1	6.67	0.381	80.4	99.3
5	7.01	0.408	81.5	87.4
25	8.27	1.56	89.5	61.9
100	13.3	3.38	143	29.7

Table 7.1: The results of the deconvolutional multiphoton disambiguation algorithm for a series of different photon occupancies. The  $x$  (longitudinal),  $y$  (transverse), and timestamping resolutions degrade with increasing photon occupancies as would be expected.

Monte Carlo. Parameter resolutions are determined by the difference between the reconstructed parameter and the value of the parameter for the true photon, where the true photon is chosen by finding the Monte Carlo hit with the closest three-dimensional proximity in  $x$ ,  $y$ , and  $t$ . There is no requirement that the pairing between reconstructed hits and true photon hits must be unique, however it is enforced that reconstructed left-right pairs must be unique in events where more than a single hit is reconstructed. In this study, the number of reconstructed hits is almost always less than the number of true photons; every reconstructed hit has an associated true photon, but not every true photon will have an associated reconstructed hit.

The single photon performance shows position and timing resolutions that largely agree with the LAPPD characterisation measurements using real data. Sub-millimetre position resolution is achievable transverse to the stripline direction using the centroiding method. The longitudinal position resolution is around 7 mm for single photons. The single photon timestamping resolution is largely limited by the simulation of the transit time spread, but at higher photon occupancies becomes limited by hit clustering which can affect the constant fraction discriminator. Table 7.1 shows the average efficiency and position and time resolutions for each subset of  $N$  photons.

As would be expected, the performance of the reconstruction with respect to each parameter degrades as more photons are added. The results suggest that on average a maximum of around 30 photoelectrons are able to be resolved for this configuration, although it is worth noting that this will be different for other scintillators with different time profiles and thus different patterns of hit clustering.




Figure 7.8: An example of an event with 100 photons incident on the LAPPD. The true photon locations are shown as blue dots with the reconstructed hits shown as red crosses. In this example, of a total 100 incident photons, around 40 are reconstructed.

Figure 7.8 shows a comparison of the locations of true photon hits on the LAPPD compared to where photons are reconstructed. As can be seen, where photons are well separated they are easily reconstructed. However, when photons are clustered together this results in only a few hits being reconstructed, or a hit is reconstructed at the rough centroid of a cluster as the deconvolution was not able to adequately disambiguate the individual pulses. These severely overlapping clusters of pulses are what cause the overall timing and position resolutions to degrade as the photo occupancy increases. In addition, mismatches in the matching algorithm between left and right pairs is likely to be another major contributor to the degradation in resolution. Overall, the position and timing resolutions are likely to be acceptable for any foreseeable usage of the LAPPD in neutrino experiments, and still offer vast improvements over typical PMTs in every metric. The only potential worry

is the efficiency value at high photon occupancies, where pixelated readouts may perform better and do not require advanced algorithms to determine hits, albeit at the expense of position resolution depending on the pixel size.

Overall, the stripline readout offers a virtually unlimited position resolution that is limited by the electronics (specifically, the sample rate) and the development of advanced hit disambiguation algorithms as presented here. Any experiment wishing to utilise stripline readout with fast photosensors such as LAPPDs will need to consider whether it wishes to exert a great amount of effort into developing these algorithms to reliably disambiguate many hits or spend the significant money required to instrument a pixelated detector to achieve the same resolution as is possible using the stripline readout. It is not yet clear which direction is likely to be followed by the LAPPD community. The Generation II LAPPD comes as standard with a pixelated readout board but the capacitively-coupled nature of the board allows for different readout boards to be attached — it is feasible to produce a capacitively-coupled stripline readout board if desired. If an experiment requires fast-timing in high photon occupancy environments and does not have position resolution requirements more stringent than a few centimetres, pixelated readout is the obvious choice. For precision neutrino experiments, especially where LAPPDs are being used to enhance vertex reconstruction in water Cherenkov or liquid scintillator detectors, it is possible that it is worth the effort to extract as much resolution out of the stripline readout as possible. For water Cherenkov detectors where the light yield is much lower, it is certainly feasible that most photon hits will be fully reconstructable with negligible losses due to multiphoton effects.

## 7.5 Vertex reconstruction using LAPPDs

In a water Cherenkov experiment such as BUTTON, the deployed LAPPDs must be integrated into the wider detector systems. Part of this integration is including the LAPPD hits into the vertex reconstruction algorithm. This should provide improved vertex resolution due to the positional resolution provided by the LAPPD as well as the much improved timing resolution over typical PMTs. In this study,




Figure 7.9: Schematic of the simulated geometry of the BUTTON detector shown in Geant4. All PMTs have been replaced with LAPPDs. Only the LAPPD geometry itself is simulated, any auxiliary equipment such as waterproof housings are omitted.

the BUTTON detector geometry (explained in detail in subsection 4.8.1) is used to evaluate the impact that LAPPDs could have on vertex reconstruction in a water Cherenkov detector. The BUTTON detector geometry with the PMTs replaced by LAPPDs is shown in Figure 7.9. Only a water fill is considered in this study; future studies could evaluate the reconstruction performance in a WbLS or liquid scintillator detector fill.

The BONSAI time-of-flight vertex fitter was used as the basis for this work as it had already been implemented for the BUTTON detector instrumented with PMTs only. As a first step, the detector geometry was edited to replace all the PMTs with




Figure 7.10: The hit time residuals calculated from the time-of-flight between the true electron vertex and the reconstructed hit position on the LAPPD.

LAPPDs. As discussed in subsection 4.4.2, BONSAI requires a photodetector hit timing likelihood to perform the log-likelihood minimisation. In order to create this likelihood, the distribution of hit time residuals must be calculated. 5 MeV electrons were simulated throughout the water volume of the detector with randomised initial directions using RAT-PAC. The truth information for photon hits on the LAPPD was tracked throughout the simulation and stored inside a ROOT TTree. This data was then used to simulate the detector response of all LAPPDs with photon hits in the event using the methodology detailed in section 7.2. Photon hits on each LAPPD were disambiguated and reconstructed using the methodology detailed in section 7.3. Since the photon disambiguation and reconstruction produces  $x$  and  $y$  positions in the local coordinate system of the LAPPD, a coordinate transform is applied — taking into account the rotation and translation of the photodetector within the global coordinate system — to transform these into the global detector coordinate system to be used as an input into BONSAI.

The distribution of time residuals was then created by calculating the time-of-flight between the true electron vertex and the global position of the detected hit on the LAPPD. This distribution is shown in Figure 7.10. The distribution for the BUTTON detector instrumented with only PMTs is shown in




Figure 7.11: The distribution of hit time residuals for a BUTTON detector instrumented with PMTs only.

Figure 7.11. As is expected, there is a much narrower timing distribution for the primary peak due to the much better timing resolution of the LAPPD relative to the PMT used in BUTTON. Note that the PMT timing distribution also includes simulation of afterpulsing, shown as the bump in the distribution at  $t = 60$  ns, but this is not included in the LAPPD response simulation.

The vertex resolution is then studied by performing simulations of monoenergetic electrons throughout the detector volume in randomised directions over several electron energies. The vertex resolutions in the  $x$ ,  $y$ , and  $z$  dimensions for both the all-PMT and all-LAPPD detector are shown in Figure 7.12. The results show that, although the resolution is very similar for very low energy events around 1 MeV (which will most likely be below the analysis threshold anyway), replacing PMTs with LAPPDs results in a consistent improvement of around a factor of three from 4 MeV upwards. The vertex resolution in the  $z$  dimension has a slightly different shape at lower energies compared to the  $x$  and  $y$  dimensions. This could potentially be a result of the asymmetrical cylindrical design of the tank as the  $z$  axis runs along the height of the cylinder. This would then cause a slight variation in the average photon transit time across the tank in the  $z$  direction resulting in a small difference in the shape of the curve at lower energies



## 7.5. Vertex reconstruction using LAPPDs

---




Figure 7.12: Resolution on the reconstructed vertex in  $x$ ,  $y$ , and  $z$  dimensions for a BUTTON detector instrumented with only PMTs and only LAPPDs.

and an improved vertex resolution overall.

Improvements in the ability to reconstruct the event vertex will help to reduce backgrounds in water Cherenkov detectors. Better vertex resolution allows for better fiducialisation of the detector volume, expanding the fiducial volume closer to the PMTs, resulting in an improved signal-to-noise ratio. In the context of reactor neutrino experiments, where a cut is placed on the maximum distance between the measured positron and neutron vertices, better vertex resolution will result in improved IBD pair matching and improved rejection of accidental backgrounds. All of this will result in improved sensitivity for detecting nuclear reactors, although a full study must be conducted in order to quantify the scale of this improvement.

A more realistic detector would include a handful of LAPPDs while the rest of the detector is instrumented with PMTs. In order to take this into account it would require combining the timing likelihoods for both photodetector types and selecting the correct distribution to draw from based on which photodetector was hit. This was not implemented in this work, but is a high priority for future work since the first deployment of LAPPDs in BUTTON will occur in conjunction with the roughly 100 PMTs that will have already been deployed as part of detector commissioning.

The work detailed here has demonstrated a toolchain for full end-to-end simulation of a water Cherenkov detector containing LAPPDs: from simulation of Cherenkov light in the detector resulting in photon hits on the LAPPD to simulation of the photodetector response as a result of those hits to performing hit reconstruction on the simulated photodetector response and finally to performing event vertex reconstruction using the reconstructed hits on the LAPPD. This toolchain will be valuable for water Cherenkov experiments such as BUTTON and ANNIE moving forward since, due to the complexities involved in analysing LAPPD signals, it is not enough to simply count hits as one would with a detector instrumented with PMTs. The deconvolution-based photon disambiguation and waveform simulator codes have been open-sourced to allow for these collaborations expecting to utilise LAPPDs in water Cherenkov detectors in the future to further develop and improve the codes as they see fit.

## 7.6 Alternative approaches towards multiphoton disambiguation

As of the time of writing, there is much interest amongst the users of Generation I stripline LAPPDs in performing multiphoton disambiguation. The deconvolution approach detailed in this chapter appears to be the most mature solution, but it is not necessarily the only or best approach to solving the problem. There are at least two alternative approaches that should be investigated as potential alternatives.

### 7.6.1 Formula fitting

A similar but alternate approach is a formula or template fitting approach. One such method is the non-negative least squares minimisation, detailed in [205]. In this, a noiseless template is fitted by a weighted sum of individual components. The equation to be solved is  $\text{argmin}|\mathbf{A}\vec{x} - \vec{b}|$ , where  $\mathbf{A}$  is a matrix that encodes a number of ideal component waveforms with scale factors given by the elements of  $\vec{x}$ .  $\vec{b}$  represents the samples of the original waveform to be fitted. Using this method, it is possible to make a determination of the total number of template pulses which best fits the waveform, providing the total number of photon hits on the LAPPD. Once this has been calculated, the time,  $x$  position, and  $y$  position can be calculated as normal since each individual component is known. A full study of photon disambiguation performance using this method has not yet been conducted. It is possible that the deconvolution and formula fitting approaches could complement each other if it becomes obvious that they are best suited to different types of events (for example, low or high photon occupancy events may be better disambiguated by one method over another).

### 7.6.2 Machine learning

One approach that could potentially be very powerful is the application of machine learning. In a similar way to the deconvolutional method whereby multiphoton disambiguation is treated more like an image analysis problem, one could conceivably use image recognition algorithms utilising convolutional neural networks.

The convolutional neural network would be trained on simulated data, using the waveform simulator detailed in section 7.2. The training data would be labelled with the number of photons present in the event, and the time sample at which the photon struck the photocathode, and the perpendicular stripline position of the photon. The network would be provided with the full simulated LAPPD response, all 28 striplines and 1,024 time samples, and would use the training data to learn to associate the location and presence of peaks in the data with the number and locations of the photons. A typical convolutional neural network for this image recognition task could be made up of a series 2D convolutional layers, separated by maximum pooling

layers to downsample the image. The input would then be flattened before being passed into a series of dense layers, finishing with a dense layer with an output dimensionality of 3 representing the three parameters of interest: the number of photons, and their positions in space and time. Obviously, the hyperparameters representing the number and types of each layer will need to be tuned to achieve optimal performance.

A simple mean squared error loss function that represents the difference between the true and estimated photon positions and times could be used to perform the training. It is likely that the training would need to be performed over a very wide range of events. First of all, many training events with a single photon would be desirable, so that the network learns the properties of the single photon LAPPD response. Following this, the number of photons present in the event would be increased up to the maximum expected in the experiment, to train the network in how to differentiate between pulses and to recognise the presence of multiple photons in what may appear to be single peaks.

It seems likely that a machine learning approach would have the highest performance ceiling in terms of the number of photons that can be disambiguated, and the overall position and temporal resolution that can be achieved. There is a chance that there will be some similarity to the deconvolution approach where, at very high photon occupancies, there simply aren't enough features present in the very large peaks to be able to determine with any level of precision the number of photons present. However, with a large and varied enough training dataset, which would be enabled through further development of the waveform simulator detailed previously, there is a good chance this limitation could be overcome.

## 7.7 Chapter summary

- Disambiguation of signals when multiple photons strike the LAPPD in close succession will be a challenge for any neutrino experiment wishing to utilise LAPPDs, especially in high light yield media such as liquid scintillator.
- An LAPPD waveform simulator has been developed to assist in developing

multiphoton disambiguation algorithms, as well as assisting with the integration of LAPPDs into a wider detector DAQ system through accurate simulation of the LAPPD detector response.

- A deconvolution-based approach to multiphoton disambiguation has been shown to be effective, although there is still much room for optimisation of the parameters controlling the deconvolution process.
- Integration of the reconstructed photon hits on the LAPPD into the BONSAI event reconstruction algorithm has been demonstrated, resulting in the first end-to-end simulation of LAPPDs in a water Cherenkov detector.
- LAPPDs have been successfully deployed into the ANNIE detector and will be deployed into the BUTTON detector in the near future. The methods detailed in this chapter will be crucial for analysing the data acquired by the LAPPDs.

# Chapter 8

## Conclusions and future work

“Don’t cry because it happened,  
smile because it’s over.”

---

In this thesis I have evaluated some of the potential benefits offered by two new technologies which are expected to be operationalised by neutrino detectors in the near future. The scope for future work is incredibly wide.

I performed a preliminary investigation into the effectiveness of the novel WbLS detection medium for the detection of nuclear reactors at great distances. I found that WbLS offered great promise but that gadolinium-doping was crucial to realise performance gains relative to a typical water Cherenkov detector. A gadolinium-doped water-based detector of the scale proposed by WATCHMAN would be capable of confirming the presence of a Heysham-like reactor at a distance of 150 km, demonstrating the applicability of antineutrino detection for monitoring nuclear reactors over large distances. Although the WATCHMAN collaboration will not ultimately construct a detector at the Boulby Underground Laboratory, there is a continued effort into research and development of new technologies for antineutrino detection. Three experiments have begun which are direct successors of the WATCHMAN collaboration based at Boulby Underground Laboratory, Brookhaven National Laboratory, and Lawrence Berkeley National Laboratory.

At Boulby, the BUTTON experiment is currently in the initial phases of constructing the proposed  $\sim 30$  tonne detector as of the time of writing. BUTTON will act as a technology testbed for antineutrino detection technologies, including alternative formulations of WbLS, LAPPDs, novel wavelength shifting plates utilising silicon photomultiplier readout, and more.

At Brookhaven, a 30 kilotonne deployment of WbLS is being planned to allow long-term precision studies of the stability and optical properties of the medium, as well as investigating compatibility with potential detector components. This will make it the largest ever deployment of WbLS and will mark a major milestone in the plan to eventually reach kilotonne-scale deployments of WbLS for future neutrino experiments.

Berkeley are spearheading the EOS programme, which is intended to evaluate the use of hybrid detection technologies, such as WbLS, to inform the design of the next generation of neutrino detectors [206]. EOS will utilise a detector with four tonnes of fiducial mass and will focus on performing Cherenkov-scintillation separation in bulk detection media using advanced photosensors such as LAPPDs or dichroicons. The volume will also be large enough to enable the use of time-of-flight information, allowing for full vertex reconstruction to be performed.

At the time of writing, it does not seem likely that a kilotonne-scale antineutrino detector for nuclear security will be constructed in the immediate future. Instead, the technology developments will be applied towards the next generation of neutrino experiments and near-field reactor antineutrino detection. The AIT infrastructure project at Boulby Underground Laboratory still has the potential to go ahead, albeit with a focus on other large-scale physics experiments, most likely a next generation liquid xenon dark matter detector. This would see historic levels of investment in physics both in the United Kingdom and especially in the North East of England.

The LAPPD characterisation campaign that I performed has demonstrated that LAPPDs are currently mature enough to be deployed in neutrino experiments. The LAPPD has shown to have sufficient gain to make it practical for single photon detection in the context of water Cherenkov neutrino detectors. I have demonstrated that the timing performance is close to the predicted

performance and is well below 100 ps transit time spread at reasonable operating conditions. In addition, I have shown that millimetre-scale position resolution in both dimensions of the LAPPD are very readily achievable with careful signal processing techniques. Overall, the characterisation measurements have shown that the single photon timing and positional performance of the LAPPD is unlikely to be matched by any existing commercial photosensor with a comparable photosensitive area.

The photon disambiguation capabilities of the Generation I stripline LAPPD in high photon occupancy environments still remain largely unknown. My simulation efforts have shown that there is good disambiguation ability up to about 30 photons incident on the LAPPD at the same time, with the caveat that these photons are arriving in very quick succession due to the time profile of the scintillator used. The disambiguation performance in other scenarios with different detector media could vary significantly. In addition, this technique has yet to be studied in real world tests. With the ANNIE experiment currently deploying a WbLS subvolume inside the main water Cherenkov detector, this could provide the first data to study different photon disambiguation algorithms in a realistic environment.

LAPPDs are seeing their first application in particle physics experiments. The ANNIE experiment currently has three deployed, with another two set to follow soon after. It is expected that LAPPDs will become utilised much more as the first hybrid detection media begin to be deployed in large volumes. The idea of Cherenkov-scintillation separation is present in the plans of many current and proposed experiments, and LAPPDs will likely play a large role in this. ANNIE's WbLS subvolume is likely to provide the first demonstration of Cherenkov-scintillation separation at large scale, as opposed to the current demonstrations on benchtop. The LAPPDs will be deployed downstream from the neutrino beamline, ensuring that a large proportion of the Cherenkov light will be observed by the LAPPDs and utilising the ultrafast timing to provide a high resolution timestamp. The delayed scintillation light will then be observed by the PMTs, which has a less stringent requirement on the timing resolution. By comparing the temporal and spatial distributions of hits on the PMTs and LAPPDs, it is expected that Cherenkov-scintillation separation with reasonable



purities will be achievable.

Despite the setbacks that have occurred in regard to the WATCHMAN collaboration, the future looks bright for the development of new technologies for neutrino detection. The first deployments of LAPPDs and WbLS in BUTTON look set to happen within the next few years. Measurements from ANNIE using the WbLS subvolume and LAPPDs should come even sooner than that.

# References

- [1] J. Chadwick, 'Intensitätsverteilung im magnetischen Spectrum der  $\beta$ -Strahlen von radium B+ C', Verhandl. Dtsc. Phys. Ges. **16**, 383 (1914).
- [2] C. D. Ellis and W. A. Wooster, 'The average energy of disintegration of radium E', Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **117**, 109 (1927).
- [3] L. Meitner and W. Orthmann, 'Über eine absolute Bestimmung der Energie der primären  $\beta$ -Strahlen von Radium E', Zeitschrift für Physik **60**, 143 (1930).
- [4] C. D. Ellis and N. F. Mott, 'Energy relations in the  $\beta$ -ray type of radioactive disintegration', Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **141**, 502 (1933).
- [5] N. Bohr, 'Faraday lecture. Chemistry and the quantum theory of atomic constitution', Journal of the Chemical Society (Resumed), 349 (1932).
- [6] W. Pauli, 'Dear radioactive ladies and gentlemen', Phys. Today **31N9**, 27 (1978).
- [7] E. Fermi, 'Tentativo di una Teoria Dei Raggi  $\beta$ ', Il Nuovo Cimento **11**, 1 (1934).
- [8] J. Chadwick, 'The existence of a neutron', Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **136**, 692 (1932).

- [9] W. J. Henderson, ‘The upper limits of the continuous  $\beta$ -ray spectra of thorium C and C’’’, Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences **147**, 572 (1934).
- [10] M. E. Nahmias, ‘An Attempt to Detect the Neutrino’, Mathematical Proceedings of the Cambridge Philosophical Society **31**, 99 (1935).
- [11] H. R. Crane and J. Halpern, ‘New Experimental Evidence for the Existence of a Neutrino’, Physical Review **53**, 789 (1938).
- [12] H. Bethe and R. Peierls, ‘The Neutrino’, Nature **133**, 532 (1934).
- [13] N. G. Cooper, ‘The Reines-Cowan experiments: Detecting the Poltergeist’, Los Alamos Sci. **25**, 4 (1997).
- [14] F. Reines, ‘The neutrino: from poltergeist to particle’, Reviews of Modern Physics **68**, 317 (1996).
- [15] I. Brüser and H. Kallmann, ‘Über die Anregung von Leuchtstoffen durch schnelle Korpuskularteilchen I (Eine neue Methode zur Registrierung und Energiemessung schwerer geladener Teilchen)’, Zeitschrift für Naturforschung A **2**, 441 (1947).
- [16] F.-H. Marshall, J. W. Coltman and A. I. Bennett, ‘The Photo-Multiplier Radiation Detector’, Review of Scientific Instruments **19**, 744 (1948).
- [17] H. Kallmann, ‘Scintillation Counting with Solutions’, Physical Review **78**, 621 (1950).
- [18] C. L. Cowan et al., ‘Detection of the Free Neutrino: a Confirmation’, Science **124**, 103 (1956).
- [19] Nobel Prize Outreach, *The Nobel Prize in Physics 1995*, <https://www.nobelprize.org/prizes/physics/1995/summary/>.
- [20] B. Pontecorvo, ‘Inverse beta process’, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **1**, 25 (1991).
- [21] R. Davis, ‘Attempt to detect the antineutrinos from a nuclear reactor by the  $\text{Cl}^{37}(\bar{\nu}e^-)\text{Ar}^{37}$  reaction’, Physical Review **97**, 766 (1955).
- [22] R. Davis and D. S. Harmer, ‘Attempt to observe the  $\text{Cl}^{37}(\nu, e^-)\text{Ar}^{37}$  reaction induced by reactor antineutrinos’, Bull. Am. Phys. Soc **4**, 217 (1959).

- [23] G. Danby et al., ‘Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos’, *Physical Review Letters* **9**, 36 (1962).
- [24] Nobel Prize Outreach, *The Nobel Prize in Physics 1988*, <https://www.nobelprize.org/prizes/physics/1988/summary/>.
- [25] The ALEPH, DELPHI, L3, OPAL, SLD Collaborations, the LEP Electroweak Working Group, the SLD Electroweak and Heavy Flavour Groups, ‘Precision electroweak measurements on the Z resonance’, *Physics Reports* **427**, 257 (2006).
- [26] M. L. Perl et al., ‘Evidence for Anomalous Lepton Production in  $e^+ - e^-$  Annihilation’, *Physical Review Letters* **35**, 1489 (1975).
- [27] P. Aarnio et al., ‘Measurement of the mass and width of the  $Z^0$ -particle from multihadronic final states produced in  $e^+ e^-$  annihilations’, *Physics Letters B* **231**, 539 (1989).
- [28] J. Kirkby, ‘Review of  $e^+e^-$  Reactions in the Energy Range 3-GeV to 9-GeV’, in 9th International Symposium on Lepton and Photon Interactions at High Energy (October 1979).
- [29] C. Blocker et al., ‘A study of the decay  $\tau \rightarrow \pi^- \nu_\tau$ ’, *Physics Letters B* **109**, 119 (1982).
- [30] V. Vuillemin, ‘ $W^\pm$  and  $Z^0$  Production in the UA1 Experiment at the CERN Proton–Antiproton Collider’, *Annals of the New York Academy of Sciences* **461**, 99 (1986).
- [31] S. L. Glashow, ‘Partial-symmetries of weak interactions’, *Nuclear Physics* **22**, 579 (1961).
- [32] S. Weinberg, ‘A Model of Leptons’, *Phys. Rev. Lett.* **19**, 1264 (1967).
- [33] A. Salam, ‘Weak and Electromagnetic Interactions’, *Conf. Proc. C* **680519**, 367 (1968).
- [34] L. D. Lella and C. Rubbia, in *Advanced Series on Directions in High Energy Physics* (WORLD SCIENTIFIC, July 2015), pp. 137–163.

- [35] F. Hasert et al., ‘Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment’, *Physics Letters B* **46**, 138 (1973).
- [36] D. DeCamp et al., ‘Determination of the number of light neutrino species’, *Physics Letters B* **231**, 519 (1989).
- [37] K. Kodama et al., ‘Observation of tau neutrino interactions’, *Physics Letters B* **504**, 218 (2001).
- [38] J. N. Bahcall, ‘Solar Neutrinos’, *Physical Review Letters* **17**, 398 (1966).
- [39] H. A. Bethe, ‘Energy Production in Stars’, *Physical Review* **55**, 434 (1939).
- [40] R. Davis, ‘A review of the homestake solar neutrino experiment’, *Progress in Particle and Nuclear Physics* **32**, 13 (1994).
- [41] B. T. Cleveland et al., ‘Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector’, *The Astrophysical Journal* **496**, 505 (1998).
- [42] K. S. Hirata et al., ‘“Observation of B-8 Solar Neutrinos in the Kamiokande-II Detector’, *Physical Review Letters* **63**, 16 (1989).
- [43] M. Cribier et al., ‘Results of the whole GALLEX experiment’, *Nuclear Physics B - Proceedings Supplements* **70**, 284 (1999).
- [44] J. Abdurashitov et al., ‘Results from SAGE (The Russian-American gallium solar neutrino experiment)’, *Physics Letters B* **328**, 234 (1994).
- [45] D. Casper et al., ‘Measurement of atmospheric neutrino composition with the IMB-3 detector’, *Physical Review Letters* **66**, 2561 (1991).
- [46] K. Hirata et al., ‘Observation of a small atmospheric  $\nu_\mu/\nu_e$  ratio in Kamiokande’, *Physics Letters B* **280**, 146 (1992).
- [47] Y. Suzuki, ‘Kamiokande solar neutrino results’, *Nuclear Physics B - Proceedings Supplements* **38**, 54 (1995).
- [48] B. Pontecorvo, ‘Mesonium and Antimesonium’, *Zhur. Eksptl'. i Teoret. Fiz.* **33** (1957).
- [49] A. McDonald et al., ‘First neutrino observations from the Sudbury Neutrino Observatory’, *Nuclear Physics B - Proceedings Supplements* **91**, 21 (2001).

- [50] Nobel Prize Outreach, *The Nobel Prize in Physics 2015*, <https://www.nobelprize.org/prizes/physics/2015/summary/>.
- [51] M. Cribier, ‘Reactor Monitoring with Neutrinos’, XXII International Conference On Neutrino Physics And Astrophysics (Neutrino 2006), Santa Fe, New Mexico : États-Unis (2006), 10.1016/j.nuclphysbps.2011.03.094 (2007).
- [52] International Atomic Energy Agency, *IAEA Safeguards Glossary, 2001 (International Nuclear Verification Series, No. 3)* (International Atomic Energy Agency, 2001), p. 218.
- [53] International Atomic Energy Agency, *Safeguards techniques and equipment* (International Atomic Energy Agency, 2011), p. 146.
- [54] International Atomic Energy Agency, *IAEA Safeguards Serving Nuclear Non-Proliferation* (2015).
- [55] P. Vogel et al., ‘Reactor antineutrino spectra and their application to antineutrino-induced reactions. II’, *Physical Review C* **24**, 1543 (1981).
- [56] S. V. Silaeva and V. V. Sinev, *The reactor antineutrino spectrum calculation*, arXiv:2012.09917, 17th December 2020.
- [57] V. I. Kopeikin and V. V. Sinev, ‘Energy spectrum of reactor antineutrinos and searches for new physics. (Recent developments)’, *Part. Nucl. Lett.* **108**, 41 (2001).
- [58] V. I. Kopeikin, L. A. Mikaelyan and V. V. Sinev, ‘Components of antineutrino emission in nuclear reactor’, *Physics of Atomic Nuclei* **67**, 1963 (2004).
- [59] E. Christensen, P. Huber and P. Jaffke, ‘Antineutrino reactor safeguards - a case study’, (2013).
- [60] S. Bogetic et al., *Anti-Neutrino Flux from the EdF Hartlepool Nuclear Power Plant*, 2023.
- [61] A. Bernstein et al., ‘Neutrino Detectors as Tools for Nuclear Security’, *Rev. Mod. Phys.* 92 (2020) no.1, 011003, 10.1103/RevModPhys.92.011003 (2019).
- [62] P. Huber, ‘Determination of antineutrino spectra from nuclear reactors’, *Physical Review C* **84**, 024617 (2011).

- [63] T. A. Mueller et al., ‘Improved predictions of reactor antineutrino spectra’, *Physical Review C* **83**, 054615 (2011).
- [64] S. Dye and A. Barna, ‘Global Antineutrino Modeling for a Web Application’, (2015).
- [65] P. Vogel and J. F. Beacom, ‘Angular distribution of neutron inverse beta decay,  $\bar{\nu}_e + p \rightarrow e^+ + n$ ’, *Physical Review D* **60**, 053003 (1999).
- [66] A. Strumia and F. Vissani, ‘Precise quasielastic neutrino/nucleon cross-section’, *Physics Letters B* **564**, 42 (2003).
- [67] K. A. Hochmuth, M. Lindner and G. G. Raffelt, ‘Exploiting the directional sensitivity of the double Chooz near detector’, *Physical Review D* **76**, 073001 (2007).
- [68] Y. V. Nikitenko, ‘Measuring Antineutrino Direction Using Inverse Beta Decay’, *Physics of Atomic Nuclei* **82**, 1635 (2019).
- [69] F. Reines, H. S. Gurr and H. W. Sobel, ‘Detection of  $\bar{\nu}_e - e$  Scattering’, *Physical Review Letters* **37**, 315 (1976).
- [70] D. 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 (2017).
- [71] A. Drukier and L. Stodolsky, ‘Principles and applications of a neutral-current detector for neutrino physics and astronomy’, *Physical Review D* **30**, 2295 (1984).
- [72] F. von Feilitzsch, A. Hahn and K. Schreckenbach, ‘Experimental beta-spectra from  $^{239}\text{Pu}$  and  $^{235}\text{U}$  thermal neutron fission products and their correlated antineutrino spectra’, *Physics Letters B* **118**, 162 (1982).
- [73] K. Schreckenbach et al., ‘Determination of the antineutrino spectrum from  $^{235}\text{U}$  thermal neutron fission products up to 9.5 MeV’, *Physics Letters B* **160**, 325 (1985).
- [74] A. Hahn et al., ‘Antineutrino spectra from  $^{241}\text{Pu}$  and  $^{239}\text{Pu}$  thermal neutron fission products’, *Physics Letters B* **218**, 365 (1989).

- [75] G. Mention et al., ‘Reactor antineutrino anomaly’, *Physical Review D* **83**, 073006 (2011).
- [76] S.-H. Seo, ‘New Results from RENO and The 5 MeV Excess’, 10.48550/ARXIV.1410.7987 (2014).
- [77] RENO Collaboration, ‘Observation of Energy and Baseline Dependent Reactor Antineutrino Disappearance in the RENO Experiment’, 10.48550/ARXIV.1511.05849 (2015).
- [78] Daya Bay Collaboration et al., ‘Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay’, 10.48550/ARXIV.1508.04233 (2015).
- [79] Y. Abe et al., ‘Measurement of  $\theta_{13}$  in Double Chooz using neutron captures on hydrogen with novel background rejection techniques’, 10.48550/ARXIV.1510.08937 (2015).
- [80] P. Huber, ‘NEOS Data and the Origin of the 5 MeV Bump in the Reactor Antineutrino Spectrum’, *Physical Review Letters* **118**, 042502 (2017).
- [81] J. M. Berryman, V. Brdar and P. Huber, ‘Particle physics origin of the 5 MeV bump in the reactor antineutrino spectrum?’, *Physical Review D* **99**, 055045 (2019).
- [82] M. Estienne et al., ‘Updated Summation Model: An Improved Agreement with the Daya Bay Antineutrino Fluxes’, *Physical Review Letters* **123**, 022502 (2019).
- [83] Y. V. Klimov et al., ‘Neutrino method remote measurement of reactor power and power output’, *Atomic Energy* **76**, 123 (1994).
- [84] A. A. Borovoi and L. A. Mikaelyan, ‘Possibilities of the practical use of neutrinos’, *Soviet Atomic Energy* **44**, 589 (1978).
- [85] N. S. Bowden, ‘Reactor monitoring and safeguards using antineutrino detectors’, *Journal of Physics: Conference Series* **136**, 022008 (2008).
- [86] N. S. Bowden et al., ‘Observation of the isotopic evolution of pressurized water reactor fuel using an antineutrino detector’, *Journal of Applied Physics* **105**, 064902 (2009).



- [87] A. Bernstein et al., ‘Monitoring the thermal power of nuclear reactors with a prototype cubic meter antineutrino detector’, *Journal of Applied Physics* **103**, 074905 (2008).
- [88] K. Eguchi et al., ‘First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance’, *Physical Review Letters* **90**, 021802 (2003).
- [89] A. Horvath, *Cherenkov.svg*, <https://commons.wikimedia.org/wiki/File:Cherenkov.svg>. This work is licensed under the Creative Commons Attribution-ShareAlike 2.5 Generic license. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/2.5/>, 2006.
- [90] P. A. Čerenkov, ‘Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light’, *Physical Review* **52**, 378 (1937).
- [91] I. M. Frank and I. E. Tamm, ‘Coherent visible radiation of fast electrons passing through matter’, *Compt. Rend. Acad. Sci. URSS* **14**, 109 (1937).
- [92] Y. Suzuki, ‘The Super-Kamiokande experiment’, *The European Physical Journal C* **79**, 10.1140/epjc/s10052-019-6796-2 (2019).
- [93] A. Haghighat et al., ‘Observation of Reactor Antineutrinos with a Rapidly Deployable Surface-Level Detector’, *Physical Review Applied* **13**, 034028 (2020).
- [94] J. P. Doering, ‘Electronic energy levels of benzene below 7 eV’, *The Journal of Chemical Physics* **67**, 4065 (1977).
- [95] C. Buck and M. Yeh, ‘Metal-loaded organic scintillators for neutrino physics’, *J. Phys. G: Nucl. Part. Phys.* **43** (2016) 093001, 10.1088/0954-3899/43/9/093001 (2016).
- [96] T. Förster, ‘Zwischenmolekulare Energiewanderung und Fluoreszenz’, *Annalen der Physik* **437**, 55 (1948).
- [97] J. B. Birks, ‘Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations’, *Proceedings of the Physical Society. Section A* **64**, 874 (1951).

- [98] J. Benziger et al., ‘A scintillator purification system for a large scale solar neutrino experiment’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **417**, 278 (1998).
- [99] Y. Nakano et al., ‘Measurement of the radon concentration in purified water in the Super-Kamiokande IV detector’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **977**, 164297 (2020).
- [100] H. Yang et al., ‘Light attenuation length of high quality linear alkyl benzene as liquid scintillator solvent for the JUNO experiment’, *Journal of Instrumentation* **12**, T11004 (2017).
- [101] K. Abe et al., ‘Neutron tagging following atmospheric neutrino events in a water Cherenkov detector’, *Journal of Instrumentation* **17**, P10029 (2022).
- [102] J. F. Beacom and M. R. Vagins, ‘Antineutrino Spectroscopy with Large Water Čerenkov Detectors’, *Physical Review Letters* **93**, 171101 (2004).
- [103] M. G. T. Lasserre, ‘Double Chooz, A Search for the Neutrino Mixing Angle  $\theta_{13}$ ’, (2006).
- [104] F. P. An et al., ‘Observation of Electron-Antineutrino Disappearance at Daya Bay’, *Physical Review Letters* **108**, 171803 (2012).
- [105] G. Leinweber et al., ‘Neutron capture and total cross-section measurements and resonance parameters of gadolinium’, *Nuclear Science and Engineering* **154**, 261 (2006).
- [106] L. Marti 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).
- [107] C. Xu, ‘Current status of SK-Gd project and EGADS’, *Journal of Physics: Conference Series* **718**, 062070 (2016).

- [108] M. Yeh 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**, 51 (2011).
- [109] L. Bignell et al., ‘Characterization and modeling of a Water-based Liquid Scintillator’, Journal of Instrumentation **10**, P12009 (2015).
- [110] E. Boutet, *Micelle\_scheme-en.svg*, [https://commons.wikimedia.org/wiki/File:Micelle\\_scheme-en.svg](https://commons.wikimedia.org/wiki/File:Micelle_scheme-en.svg). This work is licensed under the Creative Commons Attribution-ShareAlike 2.5 Generic license. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/2.5/>, 12th October 2007.
- [111] J. Caravaca et al., ‘Probing Cherenkov and Scintillation Light Separation for Next-Generation Neutrino Detectors’, Journal of Physics: Conference Series **888**, 012056 (2017).
- [112] M. J. Ford et al., ‘Pulse-shape discrimination in water-based scintillators’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1036**, 166854 (2022).
- [113] J. Caravaca et al., ‘Characterization of water-based liquid scintillator for Cherenkov and scintillation separation’, Eur. Phys. J. C **80**, 867 (2020), [10.1140/epjc/s10052-020-8418-4](https://doi.org/10.1140/epjc/s10052-020-8418-4) (2020).
- [114] T. Kaptanoglu et al., ‘Cherenkov and scintillation separation in water-based liquid scintillator using an LAPPDTM’, The European Physical Journal C **82**, [10.1140/epjc/s10052-022-10087-5](https://doi.org/10.1140/epjc/s10052-022-10087-5) (2022).
- [115] T. Kaptanoglu et al., ‘Spectral photon sorting for large-scale Cherenkov and scintillation detectors’, Physical Review D **101**, 072002 (2020).
- [116] A. Bernstein, *Conceptual Design Overview of the Advanced Instrumentation Testbed (AIT) and the WATER CHerenkov Monitor of ANtineutrinos (WATCHMAN)*, tech. rep. (March 2019).

- [117] M. Robinson et al., ‘Measurements of muon flux at vertical depth in the Boulby underground laboratory’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **511**, 347 (2003).
- [118] P. Scovell et al., ‘Low-background gamma spectroscopy at the Boulby Underground Laboratory’, *Astroparticle Physics* **97**, 160 (2018).
- [119] Hamamatsu Photonics K.K., *Large Photocathode Area Photomultiplier Tubes*, (2019) [https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99\\_SALES\\_LIBRARY/etd/LARGE\\_AREA\\_PMT\\_TPMH1376E.pdf](https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/LARGE_AREA_PMT_TPMH1376E.pdf) (visited on 10th July 2023).
- [120] O. A. Akindele et al., *Acceptance tests of Hamamatsu R7081 photomultiplier tubes*, arXiv:2306.09926, 2023.
- [121] G. Doucas et al., ‘Light concentrators for the Sudbury Neutrino Observatory’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **370**, 579 (1996).
- [122] E. Blaufuss et al., ‘ $^{16}\text{N}$  as a calibration source for Super-Kamiokande’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **458**, 638 (2001).
- [123] M. Dragowsky et al., ‘The  $^{16}\text{N}$  calibration source for the Sudbury Neutrino Observatory’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **481**, 284 (2002).
- [124] V. Fischer et al., ‘Development of an ion exchange resin for gadolinium-loaded water’, *Journal of Instrumentation* **15**, P07004 (2020).
- [125] P. Fernández, ‘GADZOOKS! (SuperK-Gd): status and physics potential’, in *Proceedings of The 34th International Cosmic Ray Conference — PoS(ICRC2015)* (August 2016).
- [126] S. Seibert, *RAT-PAC - Reactor Analysis Tool (Plus Additional Codes)*, (2014) <https://rat.readthedocs.io/en/latest/> (visited on 29th May 2023).

- [127] The WATCHMAN Collaboration, *RAT-PAC - Reactor Analysis Tool (Plus Additional Codes) v0.1.0*, (2015) <https://github.com/AIT-WATCHMAN/rat-pac> (visited on 29th May 2023).
- [128] S. Agostinelli et al., ‘Geant4—a simulation toolkit’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**, 250 (2003).
- [129] J. Allison et al., ‘Geant4 developments and applications’, *IEEE Transactions on Nuclear Science* **53**, 270 (2006).
- [130] J. Allison et al., ‘Recent developments in Geant4’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **835**, 186 (2016).
- [131] L. Lönnblad, ‘CLHEP - a project for designing a C++ class library for high energy physics’, *Computer Physics Communications* **84**, 307 (1994).
- [132] R. Brun et al., *root-project/root: v6.18/02*, 2019.
- [133] D. R. Onken et al., ‘Time response of water-based liquid scintillator from X-ray excitation’, *Materials Advances* **1**, 71 (2020).
- [134] L. Kneale, ‘Coincidence-based reconstruction and analysis for remote reactor monitoring with antineutrinos’, PhD thesis (University of Sheffield, October 2021).
- [135] Y. Zhang et al., ‘First measurement of radioactive isotope production through cosmic-ray muon spallation in Super-Kamiokande IV’, *Physical Review D* **93**, 012004 (2016).
- [136] D.-M. Mei and A. Hime, ‘Muon-induced background study for underground laboratories’, *Physical Review D* **73**, 053004 (2006).
- [137] A. Ferrari et al., *FLUKA: A Multi-Particle Transport Code*, tech. rep. (December 2005).
- [138] T. Böhlen et al., ‘The FLUKA Code: Developments and Challenges for High Energy and Medical Applications’, *Nuclear Data Sheets* **120**, 211 (2014).

- [139] F. Sutanto et al., ‘Measurement of muon-induced high-energy neutrons from rock in an underground Gd-doped water detector’, *Physical Review C* **102**, 034616 (2020).
- [140] M. Smy, ‘Low Energy Event Reconstruction and Selection in Super-Kamiokande-III’, in *International Cosmic Ray Conference*, Vol. 5, *International Cosmic Ray Conference* (January 2008), pp. 1279–1282.
- [141] The WATCHMAN Collaboration, *FRED (Functions to Reconstruct Events in the Detector)*, <https://github.com/AIT-WATCHMAN/fred> (visited on 29th May 2023).
- [142] S. Wilson, ‘Energy reconstruction, WATCHMAN UK Simulation, Analysis & Software, Internal presentation’, 7th December 2020.
- [143] WATCHMAN Collaboration, ‘The AIT-NEO Project: Results from the Second Tank Size Trade Study’, Unpublished, November 2020.
- [144] L. Kneale et al., *Sensitivity of an antineutrino monitor for remote nuclear reactor discovery*, 2022.
- [145] O. A. Akindele et al., *Exclusion and Verification of Remote Nuclear Reactors with a 1-Kiloton Gd-Doped Water Detector*, 2022.
- [146] A. G. Wright, *The Photomultiplier Handbook* (Oxford University Press), pp. 50–51.
- [147] Hamamatsu Photonics K. K., *Photomultiplier Tubes, Basics and Applications*, Fourth Edition (2017), p. 14.
- [148] B. Lubsandorzhev, ‘On the history of photomultiplier tube invention’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **567**, 236 (2006).
- [149] M. Salomon and S. Williams, ‘A multi-anode photomultiplier with position sensitivity’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **241**, 210 (1985).

- [150] Hamamatsu Photonics K. K., *Photomultiplier Tubes, Basics and Applications*, Fourth Edition (2017), p. 212.
- [151] I. P. Csorba, ‘Image tubes’, (1985).
- [152] T. Gys, ‘Micro-channel plates and vacuum detectors’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **787**, 254 (2015).
- [153] B. Leskovar, ‘Microchannel plates’, *Physics Today* **30**, 42 (1977).
- [154] J. L. Wiza, ‘Microchannel plate detectors’, *Nuclear Instruments and Methods* **162**, 587 (1979).
- [155] Z. Insepov, V. Ivanov and H. Frisch, ‘Comparison of candidate secondary electron emission materials’, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **268**, 3315 (2010).
- [156] E. Morenzoni et al., ‘Performance of microchannel plates in high magnetic fields’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **263**, 397 (1988).
- [157] M. Hattawy et al., ‘Characteristics of fast timing MCP-PMTs in magnetic fields’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **929**, 84 (2019).
- [158] V. Sulkosky et al., ‘Studies of relative gain and timing response of fine-mesh photomultiplier tubes in high magnetic fields’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **827**, 137 (2016).
- [159] W. Morrow Jr., J. Rennie and W. Markey, *Development and Manufacture of the Microchannel Plate (MCP)*, tech. rep. (Center for Night Vision and Electro-Optics, Fort Belvoir, Va, 1st February 1988).

- [160] Y. Zhang et al., ‘Effect of Hydrogen Reduction on Properties of Lead Silicate Glass for Microchannel Plates’, IOP Conference Series: Materials Science and Engineering **423**, 012167 (2018).
- [161] O. Siegmund, J. Vallergera and B. Wargelin, ‘Background events in microchannel plates’, IEEE Transactions on Nuclear Science **35**, 524 (1988).
- [162] J. Lian et al., ‘Effect of electron scrubbing on gain and dynamic range of microchannel plate’, in Optical Sensing and Imaging Technologies and Applications, edited by D. Liu et al. (December 2018).
- [163] W. Cao et al., ‘High-Sensitivity and Long-Life Microchannel Plate Processed by Atomic Layer Deposition’, Nanoscale Research Letters **14**, 10.1186/s11671-019-2983-1 (2019).
- [164] T. Credo et al., ‘Picosecond time-of-flight measurement for colliders using Cherenkov light’, in IEEE Symposium Conference Record Nuclear Science 2004. ().
- [165] B. W. Adams et al., ‘A Brief Technical History of the Large-Area Picosecond Photodetector (LAPPD) Collaboration’, (2016).
- [166] Incom, Inc, *Incom USA*, <https://incomusa.com/> (visited on 23rd June 2023).
- [167] Hamamatsu Photonics K. K., *Photomultiplier Tubes, Basics and Applications*, Fourth Edition (2017), p. 39.
- [168] S. Shin et al., *Advances in the Large Area Picosecond Photo-Detector (LAPPD): 8" × 8" MCP-PMT with Capacitively Coupled Readout*, 2022.
- [169] K. Nakamura et al., ‘Latest bialkali photocathode with ultra high sensitivity’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **623**, 276 (2010).
- [170] D. J. Ruggieri, ‘Microchannel Plate Imaging Detectors’, IEEE Transactions on Nuclear Science **19**, 74 (1972).



- [171] S. W. Thomas and G. D. Power, 'Unique microchannel plate process doubles MCPI resolution', in SPIE Proceedings, edited by U. Kim, J.-S. Chang and S.-H. Park (May 1995).
- [172] A. Stankovic et al., 'Effects of Endspoiling on Microchannel Plate Performance', in 2019 IEEE 31st International Conference on Microelectronics (MIEL) (September 2019).
- [173] M. A. Popecki et al., 'Microchannel plate fabrication using glass capillary arrays with Atomic Layer Deposition films for resistance and gain', *Journal of Geophysical Research: Space Physics* **121**, 7449 (2016).
- [174] O. H. W. Siegmund et al., 'Performance characteristics of atomic layer functionalized microchannel plates', in SPIE Proceedings, edited by O. H. Siegmund (September 2013).
- [175] R. W. Johnson, A. Hultqvist and S. F. Bent, 'A brief review of atomic layer deposition: from fundamentals to applications', *Materials Today* **17**, 236 (2014).
- [176] A. U. Mane and J. W. Elam, 'Atomic Layer Deposition of W:Al<sub>2</sub>O<sub>3</sub> Nanocomposite Films with Tunable Resistivity', *Chemical Vapor Deposition* **19**, 186 (2013).
- [177] S. J. Jokela et al., 'Secondary Electron Yield of Emissive Materials for Large-Area Micro-Channel Plate Detectors: Surface Composition and Film Thickness Dependencies', *Physics Procedia* **37**, 740 (2012).
- [178] H. Grabas et al., 'RF strip-line anodes for Psec large-area MCP-based photodetectors', *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **711**, 124 (2013).
- [179] Sierra Circuits, *Handle Crosstalk in High-Speed PCB Design*, <https://www.protoexpress.com/blog/crosstalk-high-speed-pcb-design/> (visited on 25th September 2023).

- [180] A. Kiselev et al., ‘Capacitively Coupled LAPPDs with 2D Pixelated Readout Planes for Time of Flight and Ring Imaging Cherenkov Applications’, in 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) (2021), pp. 1–6.
- [181] B. W. Adams et al., ‘Capacitive signal coupling through the vacuum wall in LAPPD™ with a conductive metal anode’, in 2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD) (October 2016).
- [182] C. Shannon, ‘Communication in the Presence of Noise’, Proceedings of the IRE **37**, 10 (1949).
- [183] E. Bogatin, *Signal and power integrity–simplified*, Second Edition (Prentice Hall, 2010), p. 62.
- [184] B. Adams et al., ‘Timing characteristics of Large Area Picosecond Photodetectors’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **795**, 1 (2015).
- [185] D. Stricker-Shaver, S. Ritt and B. J. Pichler, ‘Novel Calibration Method for Switched Capacitor Arrays Enables Time Measurements with Sub-Picosecond Resolution’, 10.1109/TNS.2014.2366071 (2014).
- [186] A. R. Back et al., *Accelerator Neutrino Neutron Interaction Experiment (ANNIE): Preliminary Results and Physics Phase Proposal*, 2017.
- [187] S. Ritt, ‘Design and performance of the 6 GHz waveform digitizing chip DRS4’, in 2008 IEEE Nuclear Science Symposium Conference Record (October 2008).
- [188] S. Ritt, R. Dinapoli and U. Hartmann, ‘Application of the DRS chip for fast waveform digitizing’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **623**, 486 (2010).
- [189] K. Croker et al., ‘Design and first performance results of waveform sampling readout electronics for Large Area Picosecond Photodetector’, Journal of Instrumentation **15**, C11016 (2020).

- [190] E. Oberla et al., ‘A 15GSa/s, 1.5GHz bandwidth waveform digitizing ASIC’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **735**, 452 (2014).
- [191] S. Ritt, *DRS4: 9 Channel, 5 GSPS Switched Capacitor Array*, tech. rep. (Paul Scherrer Institute, August 2020).
- [192] Incom, Inc., *Measurement and Test Report LAPPD 104*, tech. rep. (18th June 2021), Private Communication.
- [193] NKT Photonics, *PILAS DX - Picosecond pulsed diode lasers*, <https://www.nktphotonics.com/products/pulsed-diode-lasers/pilas/> (visited on 3rd August 2023).
- [194] B. Adams et al., ‘Invited Article: A test-facility for large-area microchannel plate detector assemblies using a pulsed sub-picosecond laser’, Review of Scientific Instruments **84**, 10.1063/1.4810018 (2013).
- [195] J. Wang et al., ‘Development and testing of cost-effective, 6 cm×6 cm MCP-based photodetectors for fast timing applications’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **804**, 84 (2015).
- [196] CAEN S.p.A, *CAEN V1742 32+2 Channel 12bit 5 GS/s Switched Capacitor Digitizer*, <https://www.caen.it/products/v1742/> (visited on 17th August 2022).
- [197] CAEN S.p.A, *6 Channel 4 kV/3 mA VME HV Power Supply Module*, <https://www.caen.it/products/v6533/> (visited on 19th May 2023).
- [198] J.-F. Genat et al., ‘Signal processing for picosecond resolution timing measurements’, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **607**, 387 (2009).
- [199] A. G. Wright, *The Photomultiplier Handbook* (Oxford University Press), p. 356.

- [200] Tektronix, Inc., *6 Series B MSO Mixed Signal Oscilloscope*, <https://www.tek.com/en/products/oscilloscopes/6-series-mso> (visited on 26th June 2023).
- [201] S. T. Wilson et al., ‘Characterisation of the Temperature-dependent Dark Rate of Hamamatsu R7081-100 10” Photomultiplier Tubes’, 10.48550/ARXIV.2306.10751 (2023).
- [202] W. Verkerke and D. Kirkby, *The RooFit toolkit for data modeling*, 2003.
- [203] Garfield++ Contributors, *Garfield++*, <https://gitlab.cern.ch/garfield/garfieldpp> (visited on 29th May 2023).
- [204] G. R. Jocher et al., ‘Multiple-photon disambiguation on stripline-anode Micro-Channel Plates’, NIM A, Volume 822, 2016, 10.1016/j.nima.2016.03.079 (2018).
- [205] E. Angelico, ‘Development of Large-Area MCP-PMT photo-detectors for a Precision Time-of-Flight System at the Fermilab Test Beam Facility’, PhD thesis (University of Chicago, 2020).
- [206] T. Anderson et al., ‘Eos: conceptual design for a demonstrator of hybrid optical detector technology’, Journal of Instrumentation **18**, P02009 (2023).

# Appendix A

## LAPPD 104




Figure A.1: Picture of LAPPD 104 immediately after delivery.

LAPPD 104, shown in Figure A.1, is a Generation I stripline photodetector designed and manufactured by Incom, Inc. It was vacuum sealed on June 2021 and was delivered to the University of Sheffield in October 2021.

Since the resistances of the MCPs have a large impact on the overall performance of the device with respect to power requirements and recharge rates, these were measured at varying points over the testing period. The resistances of the two MCPs used in LAPPD 104 were measured using a digital multimeter to be  $8.9\text{ M}\Omega$  for the entry MCP and  $3.8\text{ M}\Omega$  for the exit MCP, although during and immediately after operation these were observed to fall to  $5.1\text{ M}\Omega$  and  $1.9\text{ M}\Omega$  respectively. Although

Incom makes an attempt to match MCP resistances for MCPs within the same LAPPD, there is always likely to be some variation as is the case here. The resistance mismatch likely necessitates the use of additional parallel resistors if the device is to be powered using a voltage divider and equal voltage over each MCP is desired.

The resistances of each stripline were measured by probing across the SMA connector at each end of the stripline. A small number of striplines showed significant deviation from the nominal negligible resistance, likely because of damage to the external silver striplines caused during shipping. This damage caused some impedance mismatch at one or both ends of the affected strips, resulting in a deterioration of the signal integrity because of reflections at the boundary between the internal and external striplines. Despite this, the majority of striplines were unaffected and the rest of the system behaved nominally.

# Appendix B

## Single photoelectron operation

For laser characterisation of the LAPPD, the laser light intensity must be reduced down to the single photon level so that the vast majority of pulses generated by the LAPPD are the result of the production of a single photoelectron. To avoid a complex calibration of the laser intensity, one can try many light levels by adjusting the optical density of the neutral density filters which attenuate the laser light and use statistical arguments to identify the regime that the system is in. The most straightforward parameter to measure is the fraction of laser pulses that do not produce any response in the LAPPD, so-called “zero photoelectron” events. This method determines the fraction of all laser pulses that produce a response in the LAPPD that are single photoelectron events, as a function of the fraction of zero photoelectron events.

Following the treatment in [195], we assume that photoelectron production follows a Poisson distribution,

$$P(N_{pe}) = \frac{\overline{N}_{pe}^{N_{pe}} \times e^{-\overline{N}_{pe}}}{N_{pe}!}, \quad (\text{B.1})$$

where  $N_{pe}$  is the number of photoelectrons produced, and  $\overline{N}_{pe}$  is the average number of photoelectrons produced. Substituting in  $N_{pe} = 0$ , the average number of photoelectrons produced can be related to the probability of producing zero

photoelectrons as so,

$$\bar{N}_{pe} = -\ln(P(0)), \quad (\text{B.2})$$

where  $P(0)$  is the probability that no photoelectron is produced. From here, one can estimate the single photoelectron fraction,  $\eta$ , which is the ratio of the number of laser pulses in which exactly one photoelectron was produced compared to the number of laser pulses in which at least one photoelectron was produced,

$$\eta = \frac{P(1)}{1 - P(0)}, \quad (\text{B.3})$$

where  $P(1)$  is the probability that one photoelectron is produced. This is then related to the zero photoelectron probability by substituting  $N_{pe} = 1$  into Equation B.1, which gives,

$$\eta = \frac{-P(0) \ln(P(0))}{1 - P(0)}. \quad (\text{B.4})$$

This can be seen in Figure B.1. If the fraction of zero photoelectron events is at least 80%, then the single photoelectron fraction is greater than 90%, which we define to be the single photoelectron regime.




Figure B.1: The probability of producing a given number of photoelectrons as a function of the fraction of zero photoelectron events. The blue shaded region shows the area at which the single photoelectron fraction,  $\eta$ , is greater than 90%.