Extracting the time of core-bounce from core-collapse supernova neutrino signals in current and next-generation neutrino detectors

by
Remington W. Hill

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science (MSc) in Physics

The Office of Graduate Studies
Laurentian University
Sudbury, Ontario, Canada

THESIS DEFENCE COMMITTEE/COMITÉ DE SOUTENANCE DE THÈSE
Laurentian Université/Université Laurentienne
Office of Graduate Studies/Bureau des études supérieures

Title of Thesis Titre de la thèse	Extracting the time of core-bounce from core-collapse supernova neutrino signals in current and next-generation neutrino detectors	
Name of Candidate Nom du candidat	Hill, Remington	
Degree Diplôme	Master of Science	
Department/Program Département/Programme	Physics	Date of Defence Date de la soutenance September 25, 2023

APPROVED/APPROUVÉ

Thesis Examiners/Examinateurs de thèse:

Dr. Clarence Virtue
(Supervisor/Directeur(trice) de thèse)

Dr. Erica Caden
(Committee member/Membre du comité)

Dr. Caio Licciardi
(Committee member/Membre du comité)

Dr. Katherine Pachal
(External Examiner/Examinateur externe)

Approved for the Office of Graduate Studies
Approuvé pour le Bureau des études supérieures
Tammy Eger, PhD
Vice-President Research (Office of Graduate Studies)
Vice-rectrice à la recherche (Bureau des études supérieures)
Laurentian University / Université Laurentienne

ACCESSIBILITY CLAUSE AND PERMISSION TO USE

I, **Remington Hill**, hereby grant to Laurentian University and/or its agents the non-exclusive license to archive and make accessible my thesis, dissertation, or project report in whole or in part in all forms of media, now or for the duration of my copyright ownership. I retain all other ownership rights to the copyright of the thesis, dissertation or project report. I also reserve the right to use in future works (such as articles or books) all or part of this thesis, dissertation, or project report. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that this copy is being made available in this form by the authority of the copyright owner solely for the purpose of private study and research and may not be copied or reproduced except as permitted by the copyright laws without written authority from the copyright owner.

Abstract

Core-collapse supernovae (CCSNe) are amongst the most rare and energetic events in the galaxy. In the Milky Way, they are predicted to happen as infrequently as 1.64 ± 0.46 times per century. Over a duration of approximately ten seconds, a CCSN will convert $\approx 99\%$ of its iron core's gravitational binding energy into neutrinos. The initial wave of neutrinos is powered by the neutronization burst, which is generated by electron capture reactions on the collapsing core, which follows a critical time in the dynamics of a CCSN, core-bounce. It has been 36 years since a CCSN was observed via its neutrinos. With the observation of SN 1987A via its neutrino signal, a global effort has been undertaken to bring together all neutrino detectors under a common goal: providing an early alert to the astronomical community of an impending supernova and, if possible, triangulate to the CCSN using its neutrino signal. This effort is called the SuperNova Early Warning System (SNEWS). Triangulation simulations have recently seen tremendous success in determining where a supernova is positioned from its neutrinos, but these studies have made use of high statistics detectors such as HyperK, JUNO, and DUNE. This work implements six analytic techniques into the detectors HALO and HALO-1kT, with the intent of extracting a common reference time across all detectors to use for triangulation efforts. The common reference time chosen is the time of core-bounce (t_0), as it is followed by a rapid rise in ν_e events within ν_e sensitive neutrino detectors. Our analysis made use of the SNOwGLoBES event rate calculator, which quantifies event rates from supernova neutrino signals, which was then simulated through each detector's Monte Carlo simulation code. Various supernova models were taken into consideration to account for systematic uncertainties between different mass progenitors, equations of state, etc. Our analysis determined that for HALO and HALO-1kT, a constant fraction discriminator (CFD) technique was optimal in extracting the time of core-bounce from the neutrino signal at close distances ($<$ 3 kpc), while a negative log likelihood technique was optimal at further distances. At 1 kpc, HALO-1kT had a precision of 543 μs when using the CFD technique to extract t_0, which falls within the precision required to triangulate effectively ($<$ 1 ms, which HALO-1kT can obtain out to 3 kpc). With the intent of eventually implementing these techniques into all experiments involved in SNEWS 2.0, SNO+ was incorporated into our analysis in the later stages of this research. A preliminary exploration showed severely degraded performance in contrast to HALO-1kT, where the CFD technique could only obtain millisecond precision, not microsecond. Further analysis is encouraged.

Contents

List of Figures

List of Tables

Acknowledgements

I am indebted to Dr. Clarence Virtue for all the opportunities he has provided me over the past five years. In January 2018, Dr. Virtue took a chance on me by offering me a position as an undergraduate student researcher in the Summer of 2018. With plans to pursue a PhD at Queen's later this year, this too I owe to Dr. Virtue. It was through him that I was introduced to research opportunities available under Dr. Stephen Sekula.

All the support I received from Dr. Andrea Gallo Rosso and Dr. Tom Sonley, has not gone unappreciated. I look forward to working together again in the future. I also want to extend a thank you to Dr. Ubi Wichoski and Dr. Caio Licciardi, who I had the pleasure of working with for the last two years of my undergraduate degree.

A tremendous amount of appreciation goes to my parents, John and Donna, both of whom have provided a support network that has encompassed every aspect of my life. Without you, these past few years would have played out much differently.

I have been fortunate to have been surrounded by friends who were there for me every step of the way. Fellow graduate students Cass and Connor, thank you for all you have done assisting me along this journey. Awab, Willem, Jamie, and many others, thank you for being there for me as well.

Thank you to my wonderful fiancé, Emily. You truly are the best this world has to offer and your support throughout these past few years has been uncompromising and limitless. Your continued commitment to me and my pursuit of higher education is something I cherish deeply. Hope you are ready for four more years!

The term unprecedented has been thrown around with increasing frequency over the past three years, but I truly believe the events that transpired at Laurentian in April 2021 were just that. For a public post-secondary institution to prioritize profit over people, with such blatant disregard for the public good, including terminating our entire department, to me this foreshadows much darker things to come. It would appear that faculty, staff, and students bore the burden of Laurentian's financial woes rather than the corrupt administration whose actions directly resulted in the university going into creditor protection. An administration whose president has been rewarded with another opportunity at a major health centre, rather than punished for his failures. I came to the Laurentian physics department in September 2016, and have spent the past six years working alongside the most talented and hardworking individuals I have ever met. Wishing them all the best in whatever comes next.

CHAPTER 1

Introduction

Our modern understanding of astronomy and astrophysics has progressed to a point where our current picture of the universe is unimaginably different from that of astronomers centuries before. Research over the past century has allowed us to glean insight into the processes that govern our Sun and other stellar objects. Stellar evolution and nucleosynthesis have provided us clues to how heavy elements populate the cosmos. Deep in the heart of stars, hydrogen (H) and helium (He) are fused together to form progressively heavier elements. For stars that are sufficiently massive (> 8 $\mathrm{M}_{\odot}$), a phenomena known as Type II core-collapse supernovae is possible. For a star to terminate its life in such a manner, it must first proceed through millions of years of hydrogen burning. The length of this hydrogen burning phase is dependant on the mass and can be on the order of a billion years. After exhausting its hydrogen fuel, a shorter period of helium burning will follow. The star will proceed through successive stages of burning (carbon, neon, oxygen...) each with a shorter time scale than the last. The final stage, silicon burning, will yield iron group elements that dominate the central region of the core. The resulting star resembles that of an onion, with progressively heavier elements towards the core. This distribution of elements within the star will prove catastrophic. Iron cannot undergo nuclear fusion, as there is insufficient energy within the core to fuse these heavy elements together. Electron degeneracy can resist the immense gravitational force for a time, but will yield eventually and collapse of the core will begin. The inner core will collapse until it reaches nuclear densities, at which time it will rebound and generate an outward directed shock. The in-falling outer core, will super-sonically collide with the inner core, providing ample opportunity for this shock to propagate to the outer shells. Yet, simulations of core-collapse supernovae, conducted by numerous researchers in the latter half of the twentieth century, failed to yield an explosion, which is inconsistent with experimental observation. These simulations clearly indicated the shock wave will stall; a revival of the stalled shock can occur if the neutrinos created in the first hundreds of milliseconds post-bounce provide enough energy to the shock front, reinvigorating it and propel the outer shells into the cosmos.

These events are exceedingly rare in the Milky Way, so much so that they are only predicted to happen 1.63 ± 0.46 times per century [1] (in the Universe, it is several per second). It was 36 years ago that we were graced with a supernovae that was sufficiently close to be observed via its neutrinos. The progenitor of this SN had a mass of ≈ 20 $M_\odot$ and was identified as Sanduleak -69 202 at a distance of approximately 51 kpc. It would receive the name SN 1987A. The neutrino signal from SN 1987A was observed by three different neutrino detectors operating at the time, IMB [2], Baksan [3], and Kamiokande-II [4]. Only 25 neutrino events were observed across the three detectors. It was a groundbreaking achievement that confirmed many of the inner workings theorized about for core-collapse SN and provided world leading constraints on neutrino properties. In the years since, the SuperNova Early Warning System (SNEWS, a network of neutrino detectors that have come together to coordinate the response to the next CCSN

[5]) has come online and developed infrastructure aimed at providing an early alert to the astronomical community of an impending supernova. For the past two decades, SNEWS has been able to autonomously monitor incoming alarms from individual experiments and issue an alert to the community at large. Since coming online, no alerts have been made to the astronomical community.

Like all things, SNEWS is evolving to meet the challenges that face us and explore the new opportunities that have arisen with the advent of new detectors coming online in recent years and the years ahead (i.e. JUNO, DUNE, Super-Kamiokande). This thesis primarily contributes to the overall triangulation campaign. SNEWS has, since its inception, operated in a capacity that would be akin to someone telling you to duck when an errant puck flies over the boards without telling you where to look. Sure, it is useful to know something is coming, but directional information would be even more beneficial. In recent years, there has been a number of studies conducted to determine how effectively SNEWS can triangulate to a core-collapse supernova based off its neutrinos. Varying techniques have been explored, each of them coming to the same conclusion: with current and next-generation neutrino detectors, triangulation is not only possible, but partially useful out to distances of 10 kpc. Yet, missing from these studies was a generalized method that could be used by all detectors currently operating in SNEWS. With the primary focus on HALO-1kT (and by extension HALO), we must explore techniques that allow for detectors such as these to contribute to the SNEWS triangulation campaign. One example from a previous study [6] made use of the first event (that was followed by at least one more within a 15 ms window) in each detector to triangulate to a core-collapse supernovae. A problem with this approach is that the first event is intrinsically tied to the detector mass, interaction threshold, cross section, and supernova model. It is therefore of great interest to explore techniques that rely on a common reference time between each experiments. For this report, we elect to use the time of core-bounce (t_0), as it is the instant in time immediately prior to a rapid rise in ν_e production within the supernovae (primarily from electron capture reactions on in-falling material and within the core). An array of techniques will be implemented in HALO and HALO-1kT (with preliminary expansion into SNO+), with the intent of expansion into the wider SNEWS 2.0 framework in the future. What follows is a breakdown of each chapter.

Chapter 2 consists of a detailed discussion of historical supernovae, Type Ia supernovae, and core-collapse supernovae. This chapter establishes some core concepts such as the dynamics of core-bounce and neutrino production in core-collapse supernovae; and what is to be gained from early optical observations of a core-collapse supernovae. Chapter 3 establishes the set of core-collapse supernova models that are used in later analysis. Chapter 4 focuses on the SuperNova Early Warning System (SNEWS). Here, the focus is discussing the current infrastructure that SNEWS has made use of over the past two decades. With SNEWS evolving to meet the demands of the next-generation neutrino telescopes, a brief discussion is given to the updated infrastructure currently being implemented. However, the main focus of this chapter is to introduce the reader to the concept of triangulating with supernova neutrinos by providing an overview of historical and recent studies. From there, the motivation for this work (extracting t_0 from core-collapse neutrino signals) is discussed in earnest. Chapter 5 provides the reader with a comprehensive look at each of the neutrino detectors that are contained in the analysis chapter of this report, HALO, HALO-1kT, and SNO+. HALO is a dedicated supernova neutrino detector primarily sensitive to the ν_e flavour. It is composed of a 79 tonne lead volume and makes use of 128 ^{3}He proportional counters. HALO-1kT is its proposed successor.

These two experiments are the primary focus of this thesis; SNO+ was included towards the end of this research project, but only a preliminary exploration of the extracted t_0 was quantified. In Chapter 5, some historical commentary on how each detector came into existence, then the physics goals and neutrino observation shall be explored.

Chapter 6 illustrates how the interactions in the detector volume from a CCSN neutrino signal are calculated in each detector, and then how they are simulated (using Monte Carlo simulation methods). This chapter focuses on how we configure `SNOwGLoBES`, `sntools`, `SNEWPY`, and each detector's `Geant4` Monte Carlo simulation. Chapter 7 defines each analysis technique that is used in this thesis. Each technique has its own dedicated subchapter, where its general formulation is defined; necessary modifications are outlined; how we extract t_0; and how each technique accounts for systematic uncertainties. Chapter 8 compares each of the techniques from Chapter 7 and how well they extract the time of core-bounce in HALO-1kT. A recommendation is given on how best HALO-1kT can implement these techniques and contribute to the SNEWS triangulation campaign. Chapter 9 is the culminating analysis chapter that compares the performance of each technique in HALO-1kT, HALO, and SNO+. No recommendation for the technique to be used by SNO+ is made, as more analysis and further exploration of extraction techniques is required.

CHAPTER 2

Supernovae

2.1 Prelude

Organic life across Earth is marvelled for its diversity. Nature has a way of producing creatures that can survive some of the most inhospitable conditions present on Earth. Withstanding intense pressures and scarce resources, the bathynomus giganteus thrives in deep environments up to $\approx 2,200$ metres deep. Conversely, the frigid climate offered up by the northern most regions of our planet offers little reprieve for species that call these places home. Polar bears, snow owls, arctic hares, and many more species have populated these regions for thousands of years. In that time, humans have expanded to all corners of the Earth, reaching a point in our evolution that leaves our day to day life unfathomably detached from the experience of our ancestors. The farmers of ancient Rome, the feudal lords of thirteenth and fourteenth century Japan, and the factory worker in England throughout the industrial revolution lived vastly different lives. Their diets, beliefs, and education share little in common with one another. Their lives did share the most profound similarities in their beginning and their end. These similarities arise in all organic life, from bathynomus giganteus in the deep ocean to the Arctic hares roaming

Northern Canada. The fixed beginning and end points of life are not unique to life itself and present themselves within the grandest structures of the Universe, stars.

Nebulas, regions of space with high densities of hydrogen, helium, and cosmic dust, are the celestial wombs for the next-generation of stars, often referred to as stellar nurseries. One of the most famous examples is the Pillars of Creation, located approximately $2,100$ pc away from Earth (Figure 2.1). A star's beginning is humble, mainly hydrogen and helium collect together to form a volume of increasing mass and temperature. As the mass and temperature increase, sufficient energy becomes available for nuclei within the volume (mainly hydrogen and helium) to undergo nuclear fusion. This process occurs over the course of millions of years. Billions of years into the star's existence, it will terminate, just as our lives will. Should certain conditions be reached, the termination process can have a flair for the dramatic, leaving an imprint in the night sky for the whole galaxy to observe. These dramatic ends are referred to as supernovae, and come in a variety of forms, the most notable being core-collapse.

2.2 History & SN 1987A

The term supernova came into existence in 1929, courtesy of Walter Baade and Fritz Zwicky. Supernovae have had a profound legacy throughout human history, even as recently as towards the end of the Cold War; a time when many people alive today can offer first hand accounts.

Tuesday, February 24, 1987 was hardly different from any other day for large swaths

Figure 2.1: The Pillars of Creation nebula captured by the Hubble Space Telescope in the near-infrared view. Courtesy of [7].

of society; the same cannot be said for members of the astronomical community. This date would usher in a new era of multi-messenger astronomy, confirming the long held belief that supernovae would not only be observable optically, but also via a neutrino signal that would precede the optical signal. At 23:00 UTC, a team of researchers at Las Campanas Observatory in Chile observed a noticeably bright object emanating from the Large Magellanic Cloud [8]. This bright object would later become known as SN 1987A (Figure 2.2 shows the SN 1987A remnant as captured by the Hubble Space Telescope).

Although the optical observation of a core-collapse supernova is exciting, it is not without precedent. In 185 A.D., Chinese astronomers observed the emergence of singular bright star amongst an ocean of stars; today, it has been given the name RCW 86. Follow-up observations with modern tools at the disposal of astronomers would confirm that RCW 86 was a Type Ia supernova [9]. The mechanism that leads to a Type Ia event occurs when one star in a binary system sheds its material onto a white dwarf, inducing a runaway chain of nuclear reactions that will result in the explosion of the WD. A millennium later, various cultures across the globe bore witness to another supernova event in our night sky. The event was of the Type II variety, known as core-collapse. Type Ia and Type II SN shall be discussed in §2.4. This event left behind what is known today as the Crab Nebula, receiving its name from an early sketch done by the Irish astronomer William Parsons. The number of core-collapse supernovae observed throughout our history is ever growing, albeit at a rate considered tremendously slow on the time scales of a single human's life. Early estimates predicted a rate of 2-3 galactic core-collapse supernova per century [10]; more refined estimates in recent years now place it closer 1.63 ± 0.46 per century [1]. The estimates for CCSN rates make it clear that supernovae in the Milky Way are rare phenomena. Their rarity encourages us to be well prepared such that the events that occurred during the observation of SN 1987A do not happen again.

Figure 2.2: Observation of SN 1987A via the Hubble Space Telescope. Observation of SN 1987A via the Hubble Space Telescope, courtesy of [11].

The key feature of SN 1987A was the observation of the neutrino signal. In the present era, the odds of finding an underground particle physics lab without a detector capable of observing supernova neutrinos (≈ 10-20 MeV) is scant. SNOLAB in Canada hosts the Helium and Lead Observatory (HALO) and SNO+; each are capable of detecting supernova neutrinos. Laboratori Nazionali del Gran Sasso (LNGS) in Italy hosts LVD and Borexino; and Kamioka observatory hosts KamLAND and Super-Kamiokande (Super-K). This picture was not as rosy in 1987. Also absent was a global infrastructure to coordinate the response of these detectors to a core-collapse supernova. Even though the dynamics of a core-collapse supernova will result in the emission of all neutrino flavours prior to an intensely luminous optical signal that will linger for weeks (with a half-life of 50-80 days), SN 1987A first came to the attention of the astronomical community via its optical observation at Las Campanas Observatory. The neutrino community would discover 25 supernova neutrino events in three detectors across the globe, but only after an alert from the wider astronomical community.

The Irvine-Michigan-Brookhaven (IMB) detector was a water Cherenkov detector aimed at observing the postulated proton decay. The detector was located at Morton-Thiokel salt mine in the US state Ohio. IMB would observe a total of 8 neutrino events (presumed to be $\bar{\nu}_e$) over the span of ≈ 6 seconds. The full break down of event information can be found in Table 2.1.

Event Number	Time (UT)	Num. of PMTs	Energy (MeV)
33162	7:35:41.37	47	38
33164	7:35:41.79	61	37
33167	7:35:42.02	49	40
33168	7:35:42.52	60	35
33170	7:35:42.94	52	29
33173	7:35:44.06	61	37
33179	7:35:46.38	44	20
33184	7:35:46.96	45	24

Table 2.1: Event information for neutrinos observed in IMB from SN 1987A. Uncertainty on energies is quoted at ±25%, including both statistical and systematic uncertainties. Adapted from [2].

Kamiokande-II was a multi-purpose water Cherenkov detector that operated at the Kamioka Observatory in Japan from 1985 to 1990. Its primary physics goal was to address the solar neutrino problem, which had been established decades prior from solar neutrino measurements using ^{37}Cl as a target at the Homestake gold mine in South Dakota (see §5.1 for additional discussion). Owing to the larger detector volume in comparison to IMB, the Kamiokande-II detector would observe a dozen events across ≈ 12.5 seconds; these neutrinos were presumed to be primarily $\bar{\nu}_e$ [4]. The event information in Kamiokande-II can be found in Table 2.2.

Event Number	Time (s)	Num. of PMTs	Energy (MeV)
1	0	58	20.0 ± 2.9
2	0.107	36	13.5 ± 3.2
3	0.303	25	7.5 ± 2.0
4	0.324	26	9.2 ± 2.7
5	0.507	39	12.8 ± 2.9
6	0.686	16	6.3 ± 1.7
7	1.541	83	35.4 ± 8.0
8	1.728	54	21.0 ± 4.2
9	1.915	51	19.8 ± 3.2
10	9.219	21	8.6 ± 2.7
11	10.433	37	13.0 ± 2.6
12	12.439	24	8.9 ± 1.9

Table 2.2: Event information for neutrinos observed in Kamiokande-II from SN 1987A. Time indices of each event is with respect to the first event identified in the burst. Adapted from [4].

The third and final neutrino detector that observed SN 1987A was the Baksan Underground Scintillation Telescope (BUST). The detector was located 300 m underground, and consisted of 3156 standard detectors. Each standard detector was $\approx 70 \times 70 \times 30$ cm^3 of oil-based liquid scintillator that was viewed by a single PMT (photomultiplier tube) [3]. The event information in BUST from SN 1987A can be found in Table 2.3.

Event Number	Time (UT)	Num. of PMTs	Energy (MeV)
1	7:36:11.818	N/A	12 ± 2.4
2	7:36:12.253	N/A	18 ± 3.6
3	7:36:13:528	N/A	23.3 ± 4.7
4	7:36:19.505	N/A	17 ± 3.4
5	7:36:20.917	N/A	20.1 ± 4.0

Table 2.3: Event information for neutrinos observed in BUST from SN 1987A. Adapted from [3].

Intrinsic properties of core-collapse supernovae give rise to the features observed in the time delay between the neutrino and optical components of SN 1987A. The confirmation of the time delay between the optical and neutrino components would give rise to the SuperNova Early Warning System (SNEWS). SNEWS is a network of neutrino detectors around the globe with the singular purpose of providing an early alert to the astronomical community of an impending supernova. The reader is directed to §4 where SNEWS is discussed in more detail. For now, we must first consider the dynamics of collapse that give rise to the time delay between signals. The intrinsic properties of core-collapse supernova provide direction and assistance in achieving the mission statement of this thesis: improving current and next-generation detectors abilities to triangulate to the next core-collapse supernova by making use of intrinsic properties of the supernova to establish a common reference time and reduce systematic uncertainties that arise from detector types. What follows is a review of the different types of supernova and the dynamics of Type II supernova.

2.3 Type Ia supernovae

In the current understanding of the mechanisms that lead a star to terminate, supernovae are categorized into two overarching classes, Type I and Type II. Each of these have sub-classes, such as Type Ia, Ib, and Ic for Type I supernovae, and Type IIL, IIP, IIn, and IIb for Type II supernovae. The main optically observable characteristic that differentiates a Type Ia supernova from a Type II or core-collapse supernova, is the absence of hydrogen absorption lines in the Type Ia luminosity light curve observed by detectors/telescopes here on Earth. In most contexts, the luminosity light curve is traditionally referred to as the "light curve"; here there is added an additional qualifier as later analysis relies heavily on the term "light curve" when discussing the cumulative events observed in supernova neutrino detectors (see §7 for more discussion). Amongst the sub-classes of supernovae, Type Ia are the only type that are not brought about via core-collapse mechanisms.

Type Ia supernovae are thought to arise from binary star systems. Since the progenitor of such phenomena has been identified as a white dwarf (WD), the beacon of stability, external factors (i.e. a stellar companion) must drive the supernova. The more massive of the two stars would evolve, as our own Sun will, into a WD (assuming it is $< 1.4\ \mathrm{M}_\odot$), meaning that it would peter out, exhausting its fuel and eventually reducing its size to a radius comparable to Earth. The $1.4\ \mathrm{M}_\odot$ limit is known as the Chandrasekhar limit; it was formulated by Subrahmanyan Chandrasekhar in 1930 as an upper limit on stable WD masses [12]. A WD with a mass exceeding this would instead collapse into a black hole, as the electron degeneracy pressure within the core cannot provide an outward directed force with sufficient magnitude to push against the force of gravity within the star. A WD is inert, no longer a source of nuclear fusion reactions as the temperature within the core is insufficient to fuse heavier elements. It would be invisible to our optical and near-infrared instruments unless within close proximity to Earth due to their low luminosity. If this WD were in solitude, it would gradually cool, and all residual thermal energy would emanate into space, leaving behind a cold remnant with minimal activity within to produce light. The transition process from a white to black dwarf is on timescales larger than the current age of the Universe, as such, there is no indication that the Milky Way, or our Universe, is host to a black dwarf. Such an end is nothing more than a whimper. The same cannot be said with the introduction of a stellar companion.

2.3.1 Variation in binary star systems

The exact nature of the WD and its stellar companion is a centre point for debate within the astronomical community. Astronomers have considered a wide range of scenarios, a handful are listed below:

- Single-degenerate (SD) - The single-degenerate case is broken down into two classifications, hydrogen and helium-burning donors mixed with a WD [13]. For a hydrogen-burning donor, the donor will be a main-sequence star that fills or exceeds its Roche-lobe (a lop-sided figure eight region around a binary star system that maps the proximity that material is gravitationally bound to either star). The accretion of matter from the donor star will increase the central density, providing effective conditions for carbon to burn. The explosion mechanism outlined in §2.3.2 is now unstoppable. For a helium-burning donor, the same accretion process takes place, but the effective time-scale for thermonuclear detonation to occur is reduced.

- Double-degenerate (DD) - Double-degenerate Type Ia supernovae arise from a *pair* of carbon oxygen white dwarf (COWD) stars (in the mass range 0.7 - 1.0 $M_\odot$). The smaller of the two WD sheds its mass onto a disk surrounding the larger of the two WD, this disk then slowly accretes matter on the larger WD . If the rate of this accretion is sufficiently large [14], then the star will begin off-center burning of carbon towards the core, creating effective conditions for a core-collapse SN to occur. The accretion rate is the determining factor on whether the COWD terminates via Type Ia or core-collapse mechanisms.

- Double-detonation (DDET) - For a DDET to occur, a COWD that is below the Chandrasekhar mass must have a companion that accretes slowly enough such that off-center burning does not take place; instead, a layer of helium will form around the COWD [14]. If this helium layer could undergo detonation, the resulting shock wave would then disrupt the COWD and ignite the carbon towards the core. The companion falls under two classifications: helium-burning stars or helium rich degenerate stars [13]. A surface detonation can arise under the context of a violent double-degenerate scenario (a result of a direct collision between two COWDs [14]); this detonation would involve a binary system consisting of two COWDs, one with a thin layer of helium on its surface, or a COWD and helium donor. Thus, with the presence of helium in the surface of one of the stellar objects, a detonation could be induced, resulting in ignition of the carbon in the core.

2.3.2 Explosion mechanism of Type Ia supernovae

With the addition of a companion star to the system, the companion star will begin to shed its outer layers onto the WD. This was first hypothesized by Hoyle and Fowler [15]. A consequence of this accretion is a sharp rise in density within the WD. This increased density will provide sufficient energy to begin the fusion of the two elements most commonly found in WDs, carbon and oxygen. As matter accretes from the stellar companion onto the WD, thermonuclear flames (a propagation front that fuses elements together as it propagates) spread throughout the WD. When this flame reaches the core, it is ignited, causing a runaway chain of nuclear reactions. First carbon and oxygen are fused together, while subsequent fusions will form radioactive nickel; with a half-life on the order of months, the nickel will linger long after the supernovae as an "afterglow". The intensity of this afterglow has, historically, made supernovae distinguishable from stars in the night sky with the naked eye for weeks after the primary event. Two theories have been put forth on how this flame spreads, deflagration and detonation. In deflagration, the flame propagates through the WD at subsonic speeds. Detonation on the other hand is more violent and rapid. The flame would propagate outwards at supersonic speeds, heavily compressing the surrounding material as it radiates outward, providing effective conditions to fuse heavier elements than a deflagration wave front. Of the two, deflagration has been shown to be the most promising [16].

2.3.3 Neutrinos from Type Ia supernovae

For all the grandeur of Type Ia supernovae, their neutrino output is dwarfed in comparison to any of the Type II core-collapse supernova models that exist today. It is worth noting that the neutrino output from both core-collapse and Type Ia supernovae are highly model dependant. Early studies of the neutrino output from Type Ia supernovae indicted that the peak luminosity is achieved at $\approx 10^{50}$ ergs/s for the electron neutrino flavour (ν_e), emitted primarily from electron capture reactions on free protons [17]. When focusing on electron antineutrinos ($\bar{\nu_e}$), the outlook is more pessimistic, as the luminosity peaks at $\approx 10^{45}$ ergs/s. The resulting spectra of neutrinos emitted from the Type Ia supernova have a mean energy of ≈ 3 MeV. The exact model used by [17] in their calculations is a W7 model, a hydrodynamical model of a carbon deflagration supernovae. For a relatively "cool" core-collapse SN model [18] (which shall be discussed in more detail in §3.1) the neutrino luminosity peaks at $\approx 10^{52}$ ergs/s; additionally, the mean neutrino energy peaks at a higher value of $\approx$ 12-14 MeV, depending on the flavour of neutrino. These two key features of the CCSN spectra result in significantly fewer observable events for a typical Type Ia supernovae, over what is expected to be a conservative baseline for the output from a core-collapse supernovae. It is estimated that for a Type Ia supernova, even at a distance of 1 kpc, which by astronomical standards is in our neighbourhood, Super-Kamiokande (with an electron recoil threshold of 4 MeV and a detection efficiency of 35%) would only observe 0.28 events [17]. To contextualize this, at 10 kpc, estimates for the number of events in Super-K, from the core-collapse model in [18], range from 2110 to 2170 [6] (depending on the neutrino mass ordering).

The estimates from previous studies have been reinforced by recent studies of the neutrino emission from Type Ia supernovae in two different scenarios, the gravitationally confined detonation scenario (GCD) [19] and the deflagration-to-detonation scenario (delayed detonation transition, DDT) [20]. Both the GCD and DDT scenarios make use of only weak processes, electron (positron) capture on free protons (neutrons)

$$e^{-} + p \rightarrow n + \nu_e \tag{2.1}$$

$$e^{+} + n \rightarrow p + \bar{\nu}_e \tag{2.2}$$

where the above processes will also occur on bound neutrons and protons within nuclei.

Electron positron annihilation will also occur, resulting in the production of thermal neutrinos

$$e^{+} + e^{-} \rightarrow \nu_{e,\mu,\tau} + \bar{\nu}_{e,\mu,\tau} \,. \tag{2.3}$$

For both of these models, neutrino events were calculated in Super-Kamiokande, Hyper-Kamiokande, DUNE, JUNO, and IceCube. At a distance of 10 kpc, the expected number of events for the unoscillated case can be found in Table 2.4. It is worth noting that these results are averaged across eight lines of sight from the supernovae, as the neutrinos from Type Ia can have a preferential direction in which they are emitted. The variability between these lines of sight were negligible ($< 5\%$ change in total observed events, which arise for different neutrino mass orderings and models) [19, 20].

Detector	Events (DDT)	Events (GCD)
Super-K	0.154	0.0089
HyperK	1.725	0.0997
DUNE	0.138	0.0069
JUNO	0.063	0.0037
IceCube	1.320	0.0689

Table 2.4: Event rates for GCD and DDT Type Ia supernovae models in current and next-generation neutrino detectors at 10 kpc (unoscillated neutrino spectra). Courtesy of [19, 20].

The DDT model is most certainly a more optimistic result than [17]. Yet it would take a Type Ia supernova within close proximity of Earth (< 1 kpc) to see even a handful of events from the GCD model. With the overarching work of this thesis being the implementation of a new technique to assist in triangulating to supernovae (with primary interest in implementation into HALO and HALO-1kT), the sparse statistics from Type Ia supernovae make them less than ideal for any triangulation study. Instead, we turn our attention to core-collapse supernovae.

2.4 Type II supernovae - core-collapse supernovae (CC-SNe)

2.4.1 Stellar evolution

Put simply, stars initially sustain themselves over the period of millions to billions of years by fusing hydrogen into helium (the time scale is contingent on the mass). This is the dominant process in which energy is produced, not just within the Sun, but in all stars. The initial mass of the progenitor star dictates which of two possible hydrogen-burning modes is dominant. In stars that are more massive than 1.3 $M_\odot$, the CNO (carbon-nitrogen-oxygen) cycle is dominant. For smaller stars such as the Sun, the dominant process is the PP (proton-proton) chains, referred to as PPI, PPII, and PPIII. This work is primarily concerned with Type II (or core-collapse) supernovae, which are restricted to more massive stars (> 8 $M_\odot$ [21]), so only a cursory look shall be given at the PP chains.

2.4.1.1 PP chains

The abundance of hydrogen within the star will allow for a pair of protons to come together via

$$p + p \rightarrow d + e^+ + \nu_e \tag{2.4}$$

where d is the deuteron (an isotope of hydrogen that has been stripped of its electrons, often also denoted as ^{2}H or D). This reaction will start to occur around stellar temperatures of 2×10^7 K [22], and liberates $Q = 1.442$ MeV of energy [12] (once the positron annihilates with an electron, releasing a pair of photons in the process). The Q-value is the energy needed for or liberated from an interaction. As the density of deuterons within the star increases, ^{3}He can be produced following

$$d + p \rightarrow {}^3\text{He} + \gamma \tag{2.5}$$

where the Q-value is 5.493 MeV. Both Equations 2.4 and 2.5 are the first two steps for *all* PP chains. It is possible for ^{4}He to be produced from two deuterons ($d + d \rightarrow {}^4\text{He} + \gamma$), but this interaction is not favoured due to its low cross section and insufficient deuteron density [12].

Each of the three PP chains are shown in Figure 2.3, starting from the ^{3}He from Equation 2.5. The PPI chain is simply the formation of ^{4}He from two ^{3}He atoms. The resulting ^{4}He atom can be used, in combination with a ^{3}He atom, as a starting point for the PPII chain, which through three interactions (starting from the ^{4}He nuclei from the PPI chain) will eject a ν_e and yield two ^{4}He atoms. After the first step in the PPII chain, the PPIII chain can be initiated if a free proton captures on a ^{7}Be atom. The PPIII chain, like the PPII chain, will take three steps to eject a ν_e and yield two ^{4}He atoms. A consequence of both the PPII and PPIII chain being energetically allowed, is that the ν_e emitted in either chain will emerge with different energies, as one is a two body decay and the other is a three body decay. The out-going neutrino from $^7\text{Be}(e^-, \nu_e)^7\text{Li}$ carries either 0.863 MeV (90% of the time) or 0.385 MeV (10% of the time). The ν_e from β^+ decay of ^{8}B will carry with it up to 15 MeV in kinetic energy.

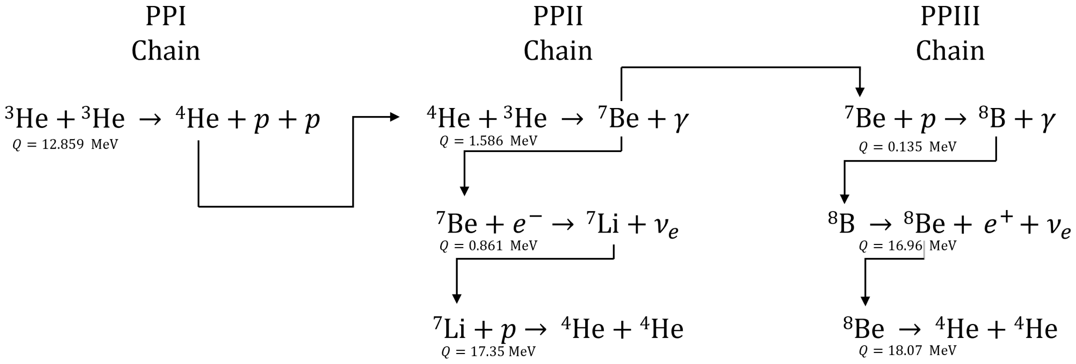

Figure 2.3: Overview of the PP chains which govern the mechanisms for fusing hydrogen into ^{4}He in stars [12, 22].

2.4.1.2 CNO cycle

The CNO cycle is much more complex than the PP chains and relies on, carbon, oxygen, and nitrogen as *catalysts*. The key difference between the PP chains and the CNO cycle, is that the CNO can only be initiated if an initial supply of carbon, or oxygen are present; the PP chains only require protons and appropriate thermal and matter density conditions to begin burning hydrogen. The specific isotope of carbon that initiates the cycle is ^{12}C. This cycle was first proposed by two independent researchers in 1938 by Hans Bethe [23] and Carl Friedrich von Weizsäcker [24, 25] (the original proposals didn't include oxygen as a part of the cycle, however the CN cycle would later be modified to include ^{16}O). Within the CNO cycle, the ^{12}C cycle is as follows [22, 12]

$$p + {}^{12}\mathrm{C} \rightarrow {}^{13}\mathrm{N} + \gamma \tag{2.6}$$

$${}^{13}\mathrm{N} \rightarrow {}^{13}\mathrm{C} + e^{+} + \nu_e \tag{2.7}$$

$$p + {}^{13}\mathrm{C} \rightarrow {}^{14}\mathrm{N} + \gamma \tag{2.8}$$

$$p + {}^{14}\mathrm{N} \rightarrow {}^{15}\mathrm{O} + \gamma \tag{2.9}$$

$${}^{15}\mathrm{O} \rightarrow {}^{15}\mathrm{N} + e^{+} + \nu_e \tag{2.10}$$

$$p + {}^{15}\mathrm{N} \rightarrow {}^{12}\mathrm{C} + {}^{4}\mathrm{He} \tag{2.11}$$

We see that through a series of β^{+} decays and proton captures, a ^{12}C nucleus can form a ^{4}He nucleus in a proton rich environment. This is a narrow view that does not include secondary contributions that arise from other sub-dominant chains. These other chains are well summarized in Figure 2.4, which depicts the other cyclic processes that branch off and return to the main ^{12}C cycle. These are the primary mechanisms that massive stars burn hydrogen. As a massive star depletes its hydrogen fuel, it turns to progressively heavier elements, starting with the product of the CNO cycle, ^{4}He.

2.4.2 Stages of burning

When the star depletes its hydrogen supply within the inner regions of the star, the delicate balance between gravitational pressure and thermal pressure is disrupted. With this disruption, the core will contract. In doing so, the density and temperature of the star will increase, opening up new pathways for the star to burn the ashes left over from hydrogen-burning. “Ashes” is the term used to refer to the nuclei that are the product

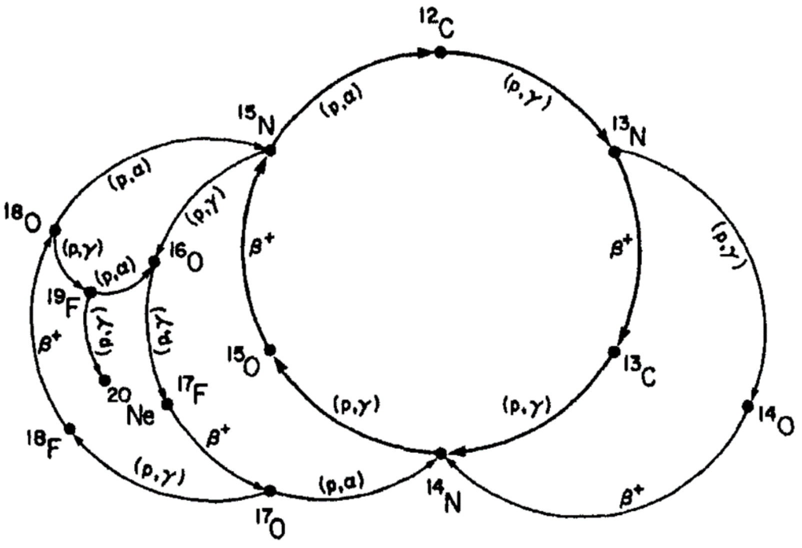

Figure 2.4: Depiction of the CNO cycle, courtesy of [22]. The main ^{12}C cycle described earlier, Equation 2.6, has different points where it can branch off and return at a later point. For example, before ^{13}N can β^+ decay to ^{13}C, the proton rich environment can allow it to capture a free proton and form ^{14}O; the ^{14}O will subsequently β^+ decay to ^{14}N, back into the main carbon cycle, releasing a ν_e in the process.

of the previous burning cycle, for the CNO cycle or PP chains, the ash would be ^{4}He. Once the star reaches a temperature of $T \approx 1.2 \times 10^8 - 2.3 \times 10^8$ K and a density of $\rho \approx 10^5 - 10^8$ kg/m^3 [22], the triple-α process can take place.

Within the core, ^{8}Be can be formed via

$$^4\text{He} + {}^4\text{He} \rightarrow {}^8\text{Be} \tag{2.12}$$

^{8}Be is unstable to α-decay, with a half-life of 8.9×10^{-17} seconds. For the star to synthesize heavier elements, it must fuse two ^{4}He nuclei at a rate that can sustain an equilibrium with the α-decay rate of ^{8}Be. At the aforementioned stellar temperature and densities, the ratio of ^{8}Be to other nuclei is as small as 1:10^9 [12]. As negligible as this might be, this is sufficient to allow the ^{8}Be to undergo another interaction with ^{4}He (hence the name the triple-α process), producing

$$^{4}\mathrm{He} + {}^{8}\mathrm{Be} \rightarrow {}^{12}\mathrm{C} + \gamma \tag{2.13}$$

^{12}C is the next heavy element synthesized in the star's core. It is possible with the formation of ^{12}C, that it too can burn the remaining ^{4}He ash in the core resulting in

$$^{4}\mathrm{He} + {}^{12}\mathrm{C} \rightarrow {}^{16}\mathrm{O} + \gamma \tag{2.14}$$

thus developing a mixed layer of ^{12}C and ^{16}O ash beneath the outer ^{4}He layer. After completing its hydrogen-burning phase by exhausting the PP chain and CNO cycle reactions, the star will then repeat the same process as before. It will contract, increase in temperature and density, and turn to the resulting ash as a new source of fuel. The next stage of burning is carbon-burning, using the ^{12}C ash from helium-burning as a source.

The carbon-burning processes will be dominated by two body interactions between two ^{12}C nuclei that yield pairs of the following: (^{24}Mg, γ), (^{23}Na, p), (^{20}Ne, α), (^{23}Mg, n), and (^{16}O, 2α) [22]. These interactions are ordered by decreasing Q-value, within the range [13.930, -0.114] MeV [12].

When the ^{12}C fuel is exhausted, the core will then briefly burn ^{20}Ne before quickly exhausting it and begin burning ^{16}O (aptly named neon and oxygen-burning respectively). The final stage that is expected to follow oxygen-burning is silicon-burning, which will begin to yield iron group elements as the ash of this burning phase. In this case iron group elements are those that are centered around the abundance peak of ^{56}Fe, which consist of Cr, Mn, Fe, Co, and Ni [26]. As the star progresses through these stages of burning, the timescale of each phase decreases one after the other. In Table 2.5, the time scale of burning phases is listed for a 15 $M_\odot$ star. In the initial hydrogen-burning phase, the star takes 11 My to exhaust the hydrogen fuel within the core, but by the time the star reaches the silicon phase, it is completed in as little as 18 days.

Stage	Time	Fuel	Ash
Hydrogen	11 My	H	He
Helium	2.0 My	He	C,O
Carbon	2000 y	C	Ne, Mg
Neon	0.7 y	Ne	O, Mg
Oxygen	2.6 y	O, Mg	Si, S, Ar, Ca
Silicon	18 d	Si, S, Ar, Ca	Fe, Ni, Cr, Ti...

Table 2.5: Time scale of each burning phase for a 15 $M_\odot$ star, alongside the fuels and ash for each respective cycle. Courtesy of [26] (additional information provided in regards to neutrino energy loss, core-collapse timing, density, and temperature, can be found in the original table [26]).

The observed decrease of burning length in each stage is a direct consequence of neutrino cooling. The star will be considered electrostatically neutral, meaning there is an approximate charge equilibrium between protons and electrons in the stars core. Through every stage of burning, alongside the ash produced are positrons (see interactions above); with the abundance of electrons in the core, these positrons can undergo annihilation processes that will yield a pair of photons or neutrinos ($e^+e^- \rightarrow \gamma\gamma$ or $e^+e^- \rightarrow \nu\bar{\nu}$). Photons have a short mean free path, so the initial thermal loses within the star will be dominated by the energy carried away by the neutrinos [22]. With the energy lost to neutrino cooling, the star's core begins to accelerate the burning of the nuclear ashes to sustain equilibrium between gravitational and thermal pressure.

After each successive stage of burning, the star develops layers of elements with increasing mass (towards the core). Such a structure is often compared to that of an onion. These layers are not entirely composed of any singular element, but rather each dominated by the ash of the previous burning cycle, and will remain active until shock breakout. The progressive stages of burning result in the formation of an iron group element core, which is unable to undergo nuclear fusion as a result of the nuclear binding energy per nucleon peaking in the range 50 < A < 60 [12]. Provided the star is close to the Chandrasekhar mass (M_{Ch}) [27]

$$M_{\mathrm{Ch}} \approx 1.44 \cdot \left(\frac{Y_e}{0.5}\right)^2 \left[1 + \left(\frac{s_e}{\pi Y_e}\right)^2\right] \cdot M_\odot \tag{2.15}$$

which for iron group nuclei cores is 1.3 $M_\odot \leq M_{\mathrm{Ch}} \leq 1.7$ $M_\odot$, then the star will undergo collapse. The variation is a consequence of the Chandrasekhar masses dependence on the progenitor star's core entropy (s_e) and electron fraction per baryon (Y_e). Core-collapse supernovae that possess non-iron group element cores are possible, but that discussion shall be deferred until §3.1, when that specific model is discussed. Here, we will briefly discuss the role photodisintegration plays prior to and directly in core-collapse before moving onto the climactic event.

2.4.3 Photodisintegration

When the star enters the heavy element burning phases, photodisintegration becomes an interaction mode that plays a critical role in the dynamics of a CCSN. Photodisintegration is the process of a nuclei absorbing a high energy γ, then emitting neutrons, protons, or α particles as it forms a lighter, more stable, nuclei. It plays a crucial role in the mixing of elements and pushing them towards stability during the later stages of a stellar evolution. Thermal conditions in the heavy element burning phases produce photons with sufficient energy to photodisintegrate surrounding nuclei. The first major contribution from photodisintegration arises during the oxygen-burning phase, with photodisintegration on ^{20}Ne

$$\gamma + {}^{20}\mathrm{Ne} \rightarrow {}^{16}\mathrm{O} + \alpha \tag{2.16}$$

yielding an ^{16}O daughter nuclei and an α particle [12]. The photodisintegration rate is proportional to the binding energy of the particle in the parent nucleus; thus photodisintegration will play a role in stripping from nuclei loosely bound protons, neutrons, or alphas. These stripped particles can then be captured on the more stable, tightly bound, nuclei. The process can be repetitive, as these ejected particles can be recaptured on loosely bound nuclei, only to undergo photodisintegration once more, repeating the process until the particle is captured on a more stable nuclei. This has profound impacts on stellar nucleosynthesis and the formation of iron group elements in the core.

2.4.4 Collapse

The stage has been set, and all the pieces are in place for the star to terminate. To summarize up to this point, a massive star will have formed layers within its core; each sequential layer is dominated by progressively heavier elements, with the central region consisting of predominantly iron group nuclei. This core, unable to fuse these iron group nuclei to sustain itself against the gravitational pressure, will begin to contract, just as it has done throughout each burning phase.

2.4.4.1 Proto-neutron star formation and core-bounce

The temperature within the core now exceeds $T \approx 4 \times 10^9$ K and the density after the silicon-burning phase reaches $\rho \approx 10^{10}$ km/m^3 [22]. The core will begin to contract inwards, where electron captures

$$e^- + p \to n + \nu_e \tag{2.17}$$

will drive the conversion of both protons in nuclei, and free protons, into neutrons, starting the formation of the proto-neutron star (PNS). The PNS is the first stage of the formation of the neutron star remnant; it is initially at extreme temperatures due to the energy liberated from the gravitational energy of the core. At the same time, photodisintegration once again plays a big role in tearing apart the iron group nuclei within the core. A photon's energy must exceed ≈ 7.6 MeV to expel an α from ^{56}Fe [22]; since photons produced at this stage can be above this threshold, the photodisintegration of iron group nuclei will begin to rob the core of stability and accelerate the collapse.

The inner core will continue to collapse until it reaches nuclear densities, $\rho \approx 2.7 \times 10^{14}$ g/cm^3, halting any further contraction [28]. This typically occurs at a radius of 20-30 km [29], equivalent to the radius of the PNS. The abrupt momentum change causes an outward directed shock. During this same period, the outer regions of the core, at this point collapsing at supersonic speeds, will come crashing into the inner core. Amongst supernova theorists, this moment in time has been labelled "core-bounce". It is often, and for the purposes of this thesis, designated as t_0, when the core reaches its maximum density, an over-compression that will result in reduced neutrino transparency. The outer core will absorb the shock wave launched by the rebounding inner core, and the outward directed shock wave will stall around a radius of 100-200 km [29, 27]. The formation of the PNS, starting from the time of core-collapse, will only take approximately one quarter of a second [27]. If the mass of the PNS is within the range 2.1 to 2.4 $M_\odot$, instead of following the below path towards a climactic end, it will collapse into a black hole and begin accreting the outer shells. Current estimates predict that massive stars below 20 $M_\odot$ have a higher probability of terminating as core-collapse supernovae, while stars above 30 $M_\odot$ will most likely collapse into a black hole [21].

In the time the PNS is formed, neutrinos are created at a tremendous rate (primarily through electron capture reactions, but so too through pair production processes which will yield all neutrino flavours). At this point in the CCSN, the neutrinos will emerge with upwards of 10s of MeV in kinetic energy. At these energies, their interaction cross section with matter is small, but still the dense core and surrounding mantle will play a role in attenuating them prior to them emerging from the star. To outside observers, the star has taken on the appearance of a "neutrino star" that has distinct regions within, where neutrinos of different flavours will emerge. These have been given the name "neutrinospheres" and are typically on the range of 30-60 km [21] from the center of the CCSN. With the shock stalled, resembling what is now an accretion shock (a shock traversing through in-falling matter but unable to push further towards the outer shells [27, 29]), some mechanism must now drive the supernova to explosion.

2.4.4.2 Stalled shock and shock breakout

Once collapse is initiated, and over the following seconds, the neutrinos created within the supernova will carry away with them ≈ 99% of the star's core gravitational binding energy. It was first postulated by S. Colgate and R. White, that provided sufficient momentum was transferred back into the shock wave, it could revitalize the shock and proceed towards ejecting the outer shells [30]. The dominant processes in which *neutrinos* could transfer momentum back into the shock are [21]

$$\nu_e + n \rightarrow e^- + p \tag{2.18}$$

$$\bar{\nu_e} + p \rightarrow e^+ + n \tag{2.19}$$

$$\nu_x + (A, Z) \rightarrow \nu_x + (A, Z) \tag{2.20}$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \tag{2.21}$$

where the first two are electron and anti-electron neutrino absorption respectively, and the latter two are inelastic scattering off nuclei and electrons. The inelastic scattering channel of transferring momentum is open to all neutrino flavour (where $\nu_x = \{\nu_e, \bar{\nu_e}, \nu_\mu, \bar{\nu_\mu}, \nu_\tau, \bar{\nu_\tau}\}$). These interactions primarily occur in the gain region, nested under the shock layer and just outside the neutrinospheres. Once the shock has been revived, the outer shells will be ejected with speeds up to 20,000 km/s [29]. The heavy elements that have been synthesized in the core and surrounding shells will populate the Universe; owing to intense thermal conditions, additional nucleosynthesis is predicted within the ejected material. This would include forming additional iron group nuclei through synthesis of the oxygen and silicon ash within the "inner ejecta" (the ash that is nested closest to the gain region)

[21]. The inner and outer mantles that surround the PNS are opaque to the optical signal emerging from the supernovae. Until the shock is able to reach the surface and eject the shells, the supernovae is primarily a neutrino emitter.

2.4.5 Neutrinos from core-collapse supernovae

The neutrinos that carry away the binding energy of the star and revive the shock, can initially carry away as high as ≈ 0.5 $M_\odot c^2$ per second [29] (this is equivalent to $\approx 10^{54}$ erg/sec, however no model used in this thesis reached luminosities of this magnitude, instead peaking at $\approx 10^{53}$ erg/sec). At first, the dominant species of neutrino emitted is the electron neutrino (ν_e) through electron capture on free protons and nuclei in the collapsing core. This is called the neutronization burst and occurs during the neutronization period. Electron capture on free protons and nuclei plays a crucial role in accelerating the collapsing core and formation of the PNS. The neutrinos produced during the neutronization burst carry away with them ≥ 10 MeV on average, sufficiently high such that detectors here on Earth can make an observation. Throughout the remaining phases, neutrinos of all flavours will continue to be produced, albeit at a rate that generally becomes undetectable here on Earth within ten or so seconds. §3 contains the supernova neutrino spectra for the various models utilized in this report.

What makes neutrinos from core-collapse supernovae interesting is the time delay between emission of the neutrinos and the optical signal. With the trapping of the optical signal underneath the dense outer layers of the star, that is until shock breakout, a time delay on the order of 30 minutes to 10 or more hours will exist. This is an observational fact, as we have discussed in the context of SN 1987A as observed by IMB, Baksan, and Kamiokande-II. From a particle physicist's perspective, there exists unique opportunities to provide world-leading constraints on neutrino properties, as was done in the aftermath of SN 1987A, owing to the large number of neutrino interactions that are expected to be observed in current and next-generation neutrino detectors. From an astrophysical point of view, the underlying dynamics of the supernova can be better understood when combining optical measurements alongside the spectral information from neutrino observation. Observing shock breakout, which would at this moment in time be only possible through luck, could become a trivial task if astronomers were provided early alerts (and directional information!) of an impending supernova. In the event of a massive star collapsing into a black hole, neutrino observation and the subsequent early alert becomes a life line for optical astronomers, as no will appear in the sky (akin to the nominal Type II events). These concepts and ideas are the main motivator for SNEWS, which brings together theorists and experimentalists to make the most of a once in a generation event.

CHAPTER 3

Supernova models

With supernova theory behind us, this chapter will now focus on the set of models used throughout this report. Five models were chosen as listed below:

1) Hüdepohl et al. 8.8 $M_\odot$ model.
2) Sukhbold et al. SFHo z9.6 $M_\odot$ model.
3) Sukhbold et al. SFHo s27.0 $M_\odot$ model.
4) Sukhbold et al. LS220 z9.6 $M_\odot$ model.
5) Sukhbold et al. LS220 s27.0 $M_\odot$ model.

The first model can be found in [18], where the Garching group simulated an 8.8 $M_\odot$ model making use of the same equation of state (EOS) as Shen et al. in [31]. For years, this model has been incorporated in the `SNOwGLoBES` software package, which calculates neutrino events in the current and next-generation neutrino detectors. At the start of this project, it was the only time-dependant model contained within `SNOwGLoBES`. `SNOwGLoBES` technically can only handle fluences (time-integrated fluxes), so the neutrino spectral parameters are binned and their integral is used in event calculations. This preserves the temporal information of the CCSN signal, while putting it in a format that `SNOwGLoBES` accepts (see §6.1.1). For the studies contained in this thesis, a time-dependant model was a necessity so the choice of model was self-evident. The other four models are from [32] and make use of either the LS220 EOS from [33] or the SFHo EOS from [34]. The name of each equation of state is derived from its authors; Lattimer and Swesty for the LS220 EOS and Steiner, Hempel, and Fischer for the SFHo EOS. Further discussion of the Sukhbold et al. models and their implementation in neutrino studies can be found in [35, 36]. The Sukhbold et al. models are fully integrated into the SNEWS collaboration's new supernova neutrino tool `SNEWPY` (see §6.1.2). These five models cover an adequate range of progenitor masses and EOSs, making them ideal for quantifying the ability for neutrino detectors to extract t_0. Models that fail to explode, collapse into black holes, or even have multiple core-bounces are not considered in this work.

The following sections provide information on the emitted neutrino spectra and luminosity. Event rates for each model simulated through the three detectors utilized in this report can be found in §6.3. Other models incorporated in `SNEWPY` [37–42] were included in the early stages of this analysis; these models were left out of this report owing to constraints we elected to impose on our analysis. Further discussion on potential inclusion of additional models can be found in §10.

3.1 Hüdepohl et al. model

The Hüdepohl et al. 8.8 $M_\odot$ model for a core-collapse SN, within `SNOwGLoBES` referred to as the Garching model, is an electron-capture supernova with an O-Ne-Mg core. In our discussion of core-collapse supernovae, it was alluded to that non-iron group element cores are possible as a star proceeds through its burning phases. These supernovae typically have progenitor masses ranging from $9 - 9.25$ $M_{\odot}$ [28]; variations in the metallicity of the star can widen this range and other estimates place progenitor masses within $\approx 8 - 10$ $M_{\odot}$ [43]. The formation of a O-Ne-Mg core, and the resultant supernova, is a direct consequence of electron degeneracy being reached before temperatures sufficient to burn Ne are reached [28]. It is expected that O-Ne-Mg core-collapse supernovae will make up approximately 30% of all core-collapse supernovae [43]. The core masses of these supernovae are constrained, as iron group element cores are, to be around 1.4 $M_{\odot}$. Provided the core is able to reach this mass, then the star shall explode via a Type II supernova mechanism. In these supernovae, the collapse is primarily facilitated by electron captures, hence the name electron-capture supernova. What separates these from iron group element cores is the steep density profile decline from the inner core to the outer hydrogen and helium shells [43, 44], which are loosely bound to the star.

The Garching model, as well as other O-Ne-Mg core-collapse supernovae, serve as an excellent conservative baseline for any study utilizing the neutrinos emitted from core-collapse supernovae. The driving factor that makes electron-capture supernovae ideal candidates for neutrino studies is their low interaction yield, meaning they have fewer events in the detector volume compared to iron-core CCSN models. The time evolution of the neutrino luminosity and mean energy can be found in Figures 3.1 and 3.2 respectively. The neutrino luminosity peaks at 10^{50} ergs/sec, specifically for the electron neutrino flavour. This neutronization burst illustrates the dominant effect that electron capture reactions have on neutrino production in the early stages of collapse in O-Ne-Mg core-collapse supernovae. For ν_e sensitive detectors, the bulk of events observed in the

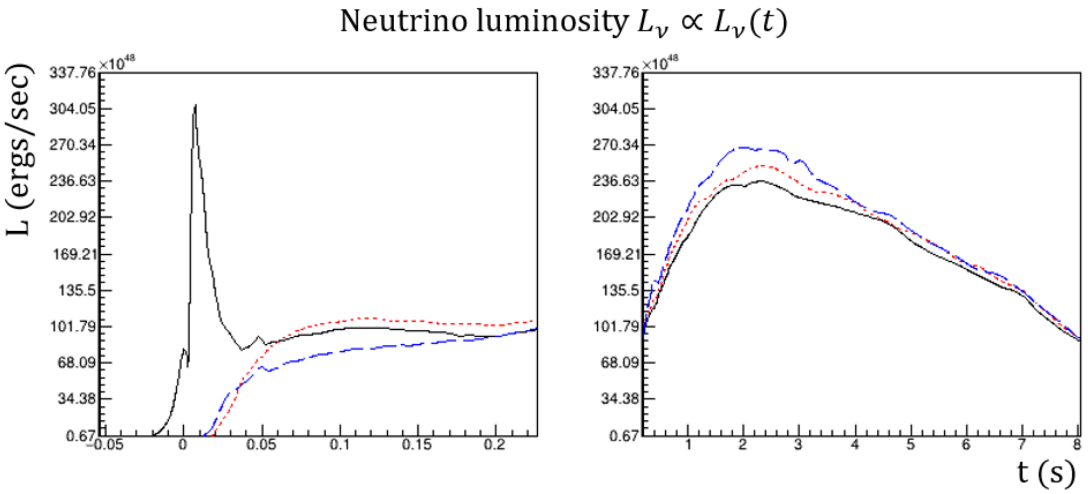

Figure 3.1: Time evolution of the neutrino luminosity for the Garching model, courtesy of [18]. Colour and line type correspond to different neutrino species. The solid black line is ν_e; the red, short dashed line is $\bar{\nu_e}$; the blue, long dashed line are the ν_μ, ν_τ, and their anti-neutrino counterparts, simply called ν_x.

neutrino burst will occur in the ≈ 50 ms following core-bounce. This is the unoscillated spectrum of neutrinos emitted and does not take into account oscillations between neutrino species (the effect neutrino oscillation has on the analysis undertaken in this thesis is discussed in §9.2). Accretion occurs in the immediate aftermath of core-bounce and the neutronization burst, until the neutrinos reverse the momentum of accreting material (≥ 0.2 seconds). At this time, the surface of the PNS begins to cool; while the PNS contracts and increases in temperature (Kelvin-Helmholtz cooling) [18]. The time scale of this cooling period is $O(10\text{s})$, but the neutrino output is only simulated out to ≈ 9 seconds for this model. For the purposes of this thesis, the cutoff serves a practical purpose (calculation and simulation duration) without compromising the regions that are rich in events.

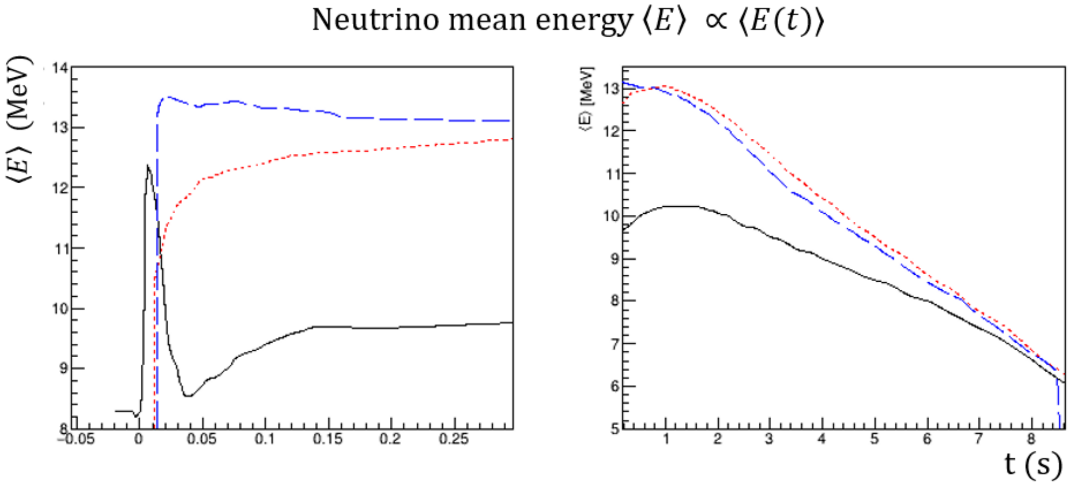

Figure 3.2: Time evolution of the mean neutrino energy for the Garching model, courtesy of [18]. Colour and line type correspond to different neutrino species. The solid black line is ν_e; the red, short dashed line is $\bar{\nu_e}$; the blue, long dashed line are the ν_μ, ν_τ, and their anti-neutrino counterparts, simply called ν_x.

3.2 Sukhbold et al. models

The four models chosen from [32] provide us with the means of stress testing our analytic techniques implemented in §7. They are 1D models that have been artificially exploded. For the two LS220 models, the 220 indicates the nuclear incompressibilty modulus for that particular EOS (in units of MeV) [36]. The LS220 and SFHo models each make use of unique equations of state in comparison to the Garching model. Including these in our analysis allows us to quantify systematic uncertainties that will arise in later analysis (see §7 for more details). In Figure 3.3 and 3.4 the reader will find the luminosity and mean neutrino energy for both Sukhbold et al. SFHo models. The Sukhbold et al. LS220 models spectra are found in Figure 3.5 and 3.6.

An important point of discussion between the Sukhbold et al. models and the Garching model, is their luminosities. It was noted that the flux of neutrinos leaving the Garching

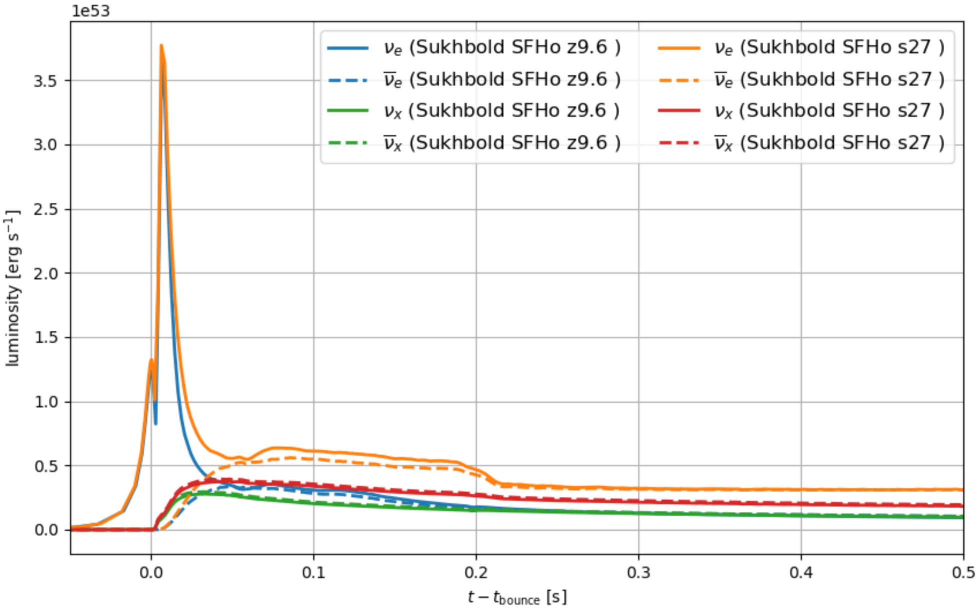

Figure 3.3: Luminosity as a function of time since core-bounce for the Sukhbold et al. SFHo models.

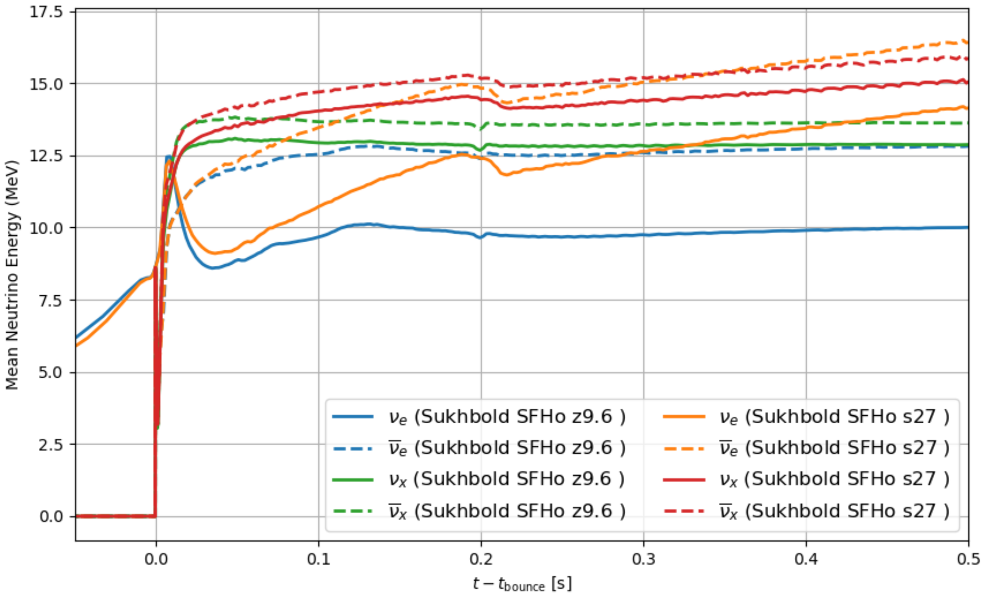

Figure 3.4: Mean neutrino energy as a function of time since core-bounce for the Sukhbold et al. SFHo models.

model for a core-collapse supernova was much higher compared to a Type Ia supernova. By CCSN standards, the Garching model is considered “cool”. For the Sukhbold et al. SFHo and LS220 models, their luminosities are on the order of 10^{53}, three orders of magnitude larger than the Garching model. As a consequence, the number of events expected for these models is much greater (Section 6 will illustrate how these events are calculated and simulated, §7 will have the number of events). A feature that is of significant importance is the shouldering observed in the luminosity curves in the 200 ms post bounce (this can be seen in Figures 3.3 and 3.5 for the 27 $M_\odot$ models). For the higher mass progenitors, instead of tapering off exponentially, their luminosities, and by consequence fluxes, are approximately double that for each flavour of neutrino than the lower mass counterpart. It should be expected that upon calculating the number of events from this signal, there will be a higher concentration of events in this shoulder region compared to the low mass models. This feature will require special consideration to reduce the systematic effects that arise in our fitting techniques when we “mix and match” different SN models. More on this in §7.

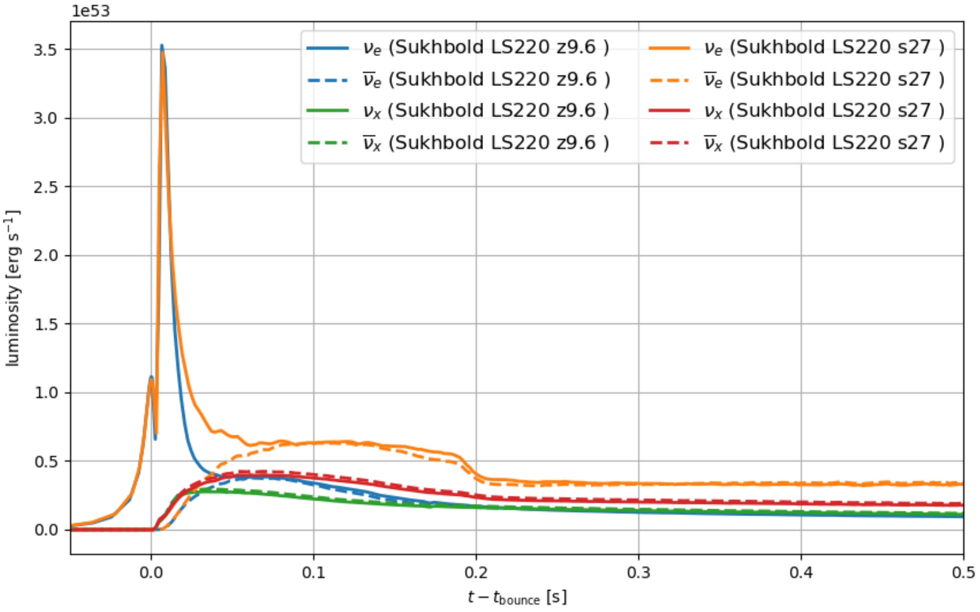

Figure 3.5: Luminosity as a function of time since core-bounce for the Sukhbold et al. LS220 models.

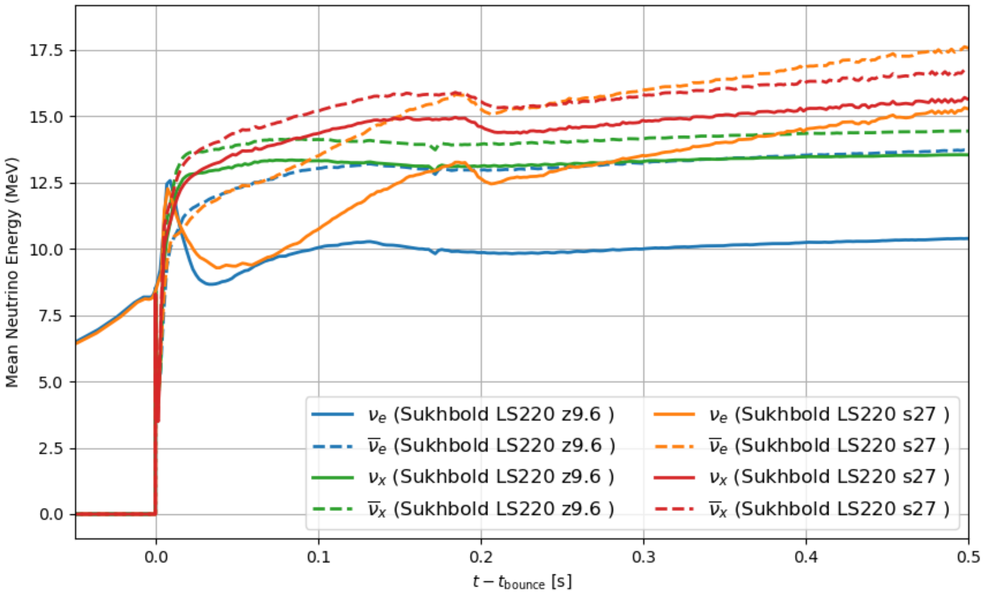

Figure 3.6: Mean neutrino energy as a function of time since core-bounce for the Sukhbold et al. LS220 models.

CHAPTER 4

SuperNova Early Warning System

The SuperNova Early Warning System (SNEWS) is a global collaborative effort with the aim of coordinating the response of neutrino detectors for the next galactic core-collapse supernova. This chapter will discuss the origins and history of SNEWS, its current infrastructure, and its evolution into SNEWS 2.0.

4.1 Origins and history

In §2.2, we touched upon perhaps the most consequential event in multi-messenger astronomy, the neutrino observation of SN 1987A; this event is considered the birth of neutrino astronomy. It was hypothesized that the neutrino wavefront from a CCSN would arrive prior to its optical counterpart. The relative time difference between these two signals was always in question. With the observation of the neutrino component of SN 1987A, not only was this hypothesis proven, but out of this observation came the birth of neutrino astronomy. Within 18 years, SNEWS would come online as an automated system primed to deliver alerts to members of the astronomical community if the neutrinos from an impending supernova were observed were observed in coincidence in two or more detectors.

4.2 Infrastructure

In its current configuration, the SNEWS communication network for neutrino experiments is run from a server at Brookhaven National Laboratory (BNL), with a backup server in Bologna, Italy. Each participating experiment can send an alarm to the SNEWS server with two fields, a packet type and a level flag. The packet field can be of the following types [45]:

- *PING* is for testing purposes. It will only fill the log stored on the local server with a printed message.
- *ALARM* contains information about a given detector alarm, which is then placed in a queue for comparison with other detector alarms.
- *RETRACTION* provides general information about previous packets that an experiment wishes to retract from the server's alarm queues.

The level field holds the values of:

- *TEST* informs the server that the packet is purely for testing purposes.
- *POSSIBLE* indicates that the packet sent occurred under suspicious conditions. These might include detector maintenance or calibration, where background events could trigger an individual detector's SN alarm. Experiments will generally have automated procedures to handle this.

- *GOOD* indicates that the information sent to the SNEWS server is believed to be a SN and not attributed to a non-normative detector state (i.e. maintenance or calibration).
- *RETRACTED* will retract the previous message, in combination with the packet field. Generally this will contain information about the *RETRACTION*.
- *OVERRIDE* will confirm to the server that the previous alert is good. If another detector has issued an alert that is *GOOD*, the *OVERRIDE* field will bypass the conditions listed below and issue a GOLD alert.

SNEWS will issue an ALERT, should two or more detectors issue alarm datagrams to the SNEWS server with timestamps that are within 10 seconds of one another and the following conditions on the server are met:

- That coincidence detectors are not co-located (an example of such would be HALO and SNO+, both located at SNOLAB). This condition ensures that any potential SN candidate is not from background sources present at a singular lab (such as geological activity in the region).
- Packets from each of the detectors coincident with one another must be considered "GOOD". What determines a "GOOD" alert is up to the experiment, as discussed before. For more details on this implementation in HALO see §5.3.
- For a given detector, their recent history of ALARMS issued to the SNEWS server is consistent with less than one burst per week on average. Such a condition allows for noisy detectors to be removed from a potential coincidence with other detectors. SNEWS aims to avoid being the boy who cried wolf, more details of how this is achieved can be found in §4.3.2.

Provided these conditions are met, an ALERT is then sent out to the astronomical community. The reader is encouraged to participate by signing up at: `https://snews.bnl.gov/alert.html`. SNEWS 2.0 alerts can also be found at: `https://snews2.org/alert-signup/`. Should any of these three conditions not be met, the alert is assigned a SILVER status, rather than GOLD. A flowchart illustrating the determination of GOLD and SILVER alerts can be found in Figure 4.1. GOLD alerts are sent to the astronomical community automatically; SILVER alerts will only be sent to experiments participating in SNEWS. SILVER alerts *must* be approved by each experiment prior to any announcement being made to the public [45].

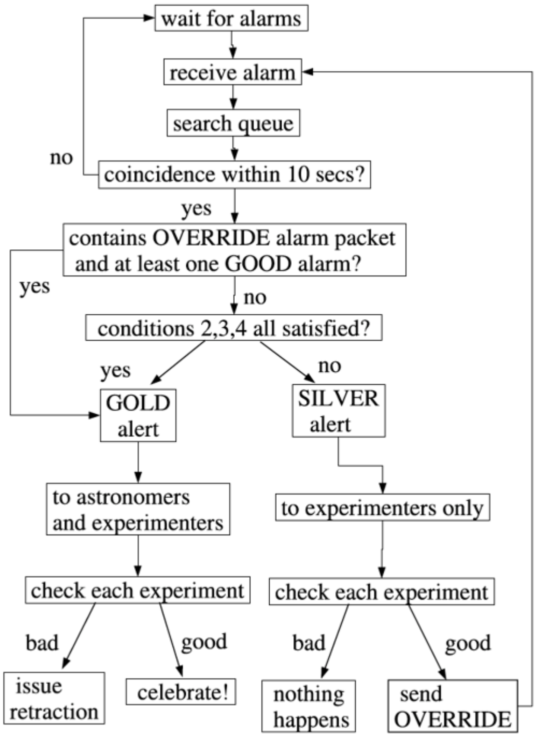

Figure 4.1: Flowchart illustrating the SNEWS decision process, courtesy of [5].

4.2.1 Participating experiments

When SNEWS was first brought online in 2004, there were four participating experiments: Super-K, SNO, IceCube, and LVD. Over the years, other experiments have joined and left. Borexino first joined in July 2009, followed by KamLAND in December 2013. Daya Bay would join in November 2014, with the most recent addition being HALO in October 2015.

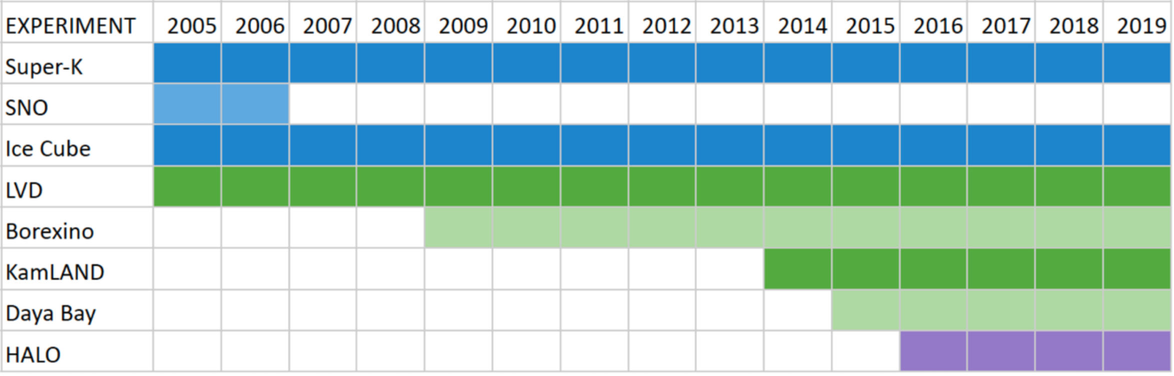

EXPERIMENT	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Super-K	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦
SNO	🟦	🟦													
Ice Cube	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦
LVD	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
Borexino					🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
KamLAND										🟩	🟩	🟩	🟩	🟩	🟩
Daya Bay											🟩	🟩	🟩	🟩	🟩
HALO												🟪	🟪	🟪	🟪

Figure 4.2: Progression of neutrino experiments joining (or leaving) SNEWS since it first came online (as of 2019). In the past four years, Borexino, KamLAND, Daya Bay, and LVD have ceased issuing SNEWS alerts, as they have ceased taking data or are in the process of decommissioning.

4.3 The three P's

The main pillars of the SNEWS early alert system are the "three P's": prompt, positive, and pointing. Each of these pillars informs all aspects of SNEWS and shall be discussed in detail below.

4.3.1 Prompt

The neutrino wavefront of SN 1987A arrived only two and a half hours prior to the optical signal. It is then necessary that the mechanisms in place to alert the astronomical community are efficient and automated. The goal of SNEWS is to issue an alert within five minutes of a coincidence alert from two or more detectors. Not only is the automation of alerts on the server side important, but so too are the ALARMS coming from the experiments themselves. Automated alerts and trigger conditions are currently being implemented for SNO+ and has been implemented in HALO since October 2015, when the experiment joined SNEWS. For HALO, SN ALARMS (as well as other alerts, spallation, coincidence, other, etc.) are sent to the collaboration via email. However only SN ALARMS containing the time of the burst are sent to SNEWS.

4.3.2 Positive

When the infrastructure of SNEWS was being developed, the astronomical community insisted that false alarms occur less than once per century, or "not in my lifetime". This is achieved by imposing a false alarm rate per detector of once per week. These rates are subjected to Poisson fluctuations. The accidental alert can be found for an n-fold coincidence between N different detectors [5, 45], with results shown in Figure 4.3; for ≥ 4 detectors, a minimum of a three fold coincidence must be used.

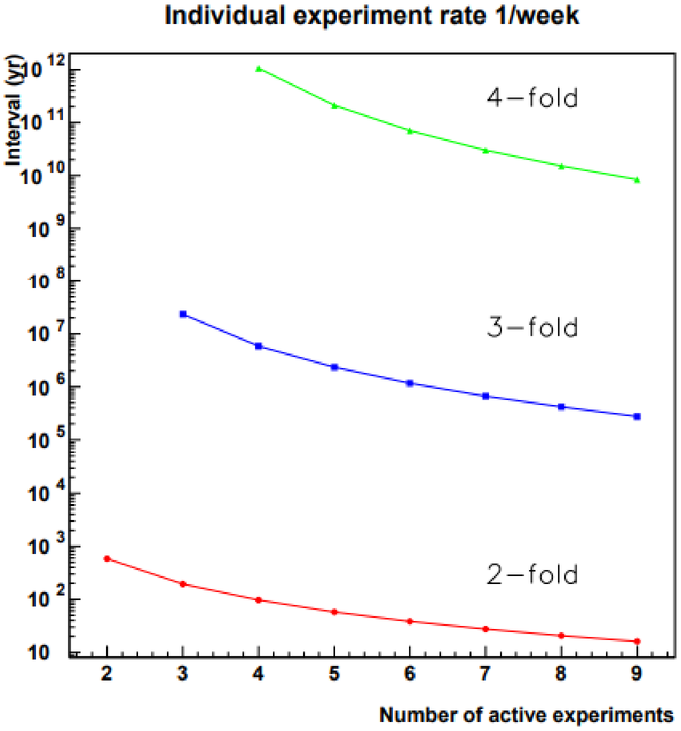

Figure 4.3: False alarm rate for SNEWS as a function number of active detectors, for two, three, and four-fold random coincidences. Courtesy of [45].

4.3.3 Pointing

One can make the argument that the most important feature of SNEWS is its ability to communicate the SN location to the astronomical community. There are two techniques one can use in an attempt to locate the supernova, pointing or triangulation. The former relies on event vertex reconstruction in combination with the light distribution in a given detector to point back towards the SN. In [46] by J. F. Beacom and P. Vogel, consideration was given to how well Super-K and SNO+ could use the directionality of $\nu e^- \rightarrow \nu e^-$ elastic scattering events to locate a core-collapse supernova. The scattering angle, α, between the electron and neutrino is given by

$$\cos\alpha = \frac{E_\nu + m_e}{E_\nu}\left(\frac{T}{T + 2m_e}\right)^{1/2} \tag{4.1}$$

where E_ν is the kinetic energy of the incident neutrino, m_e is the mass of the electron, and T is the experimental energy threshold for this interaction. The uncertainty of the angle from the best fit direction, θ, can be determined in each detector from

$$\delta\theta \approx \frac{\delta\alpha}{\sqrt{N_s}} \tag{4.2}$$

where N_s is the number of events in the detector and $\delta\alpha$ is the error on the scattering angle (or cone width). For a naive estimate of $\delta\alpha = 25°$ it was shown that a circle centered around the supernova with an angular width of 1.5° and 5.0°, for Super-K and SNO respectively, could be determined. This ignores the inclusion of isotropic backgrounds. Taking backgrounds into account, Beacom and Vogel were able to restrict the supernova direction to an angular width of 5° and 20° for Super-K and SNO respectively. In the years that followed, a second study by R. Thomas et al. [47] looked at non-isotropic interaction channels in Super-K, with background contributions from isotropic interaction channels $\bar{\nu_e} + p \rightarrow n + e^+$ and $\nu_e + {}^{16}\mathrm{O} \rightarrow X + e^-$, to determine how well it would be able to point back towards a CCSN. The absolute worst case was 8 degrees at 95% confidence level (CL); with an improved event tagging efficiency of 95%, this result would greatly improve to 3 degrees. A triumphant result that shall only improve with larger megaton-scale detectors.

Triangulation, on the other hand, uses the time difference of the observed neutrino signal between multiple detectors to locate the supernova. The time difference between two or more detectors must be on the order of milliseconds for the supernova to be located with adequate precision. When SNEWS was first established, current-generation detectors were not able to achieve this feat. One of the outcomes of early studies was that provided with sufficient statistics, commensurate with Super-K ($\approx 10^4$ events at 10 kpc) in each detector, triangulation via the method derived in [46] could accurately determine where a supernova was. Decades have passed since these techniques were first explored. In recent years, many triangulation techniques have been explored for the next-generation of neutrino telescopes, with promising results. These works are discussed in more detail in §4.4.2; how the scope of this thesis fits into the SNEWS 2.0 triangulation campaign is thoroughly discussed in §4.4.3.

4.4 SNEWS 2.0 and triangulation

SNEWS has entered a new era of neutrino astronomy. With next-generation detectors coming online in the years to come and, combined with the observation of gravitational waves over the past half a decade, improvements and expansions to the pre-existing SNEWS network are overdue. SNEWS submitted a proposal to the National Science Foundation (NSF) in 2018 to fund the next installment of the SNEWS program. SNEWS received funding the next year and the first SNEWS 2.0 workshop was held at Laurentian University in June 2019. In the years since, a tremendous amount of work has been done by members of the collaboration to update our event generation tools, triangulation routines, publishing tools, and more. The following subchapters will briefly discuss some of the ongoing work, as well as how the main objectives of this thesis fit neatly into the overall SNEWS 2.0 mission statement.

4.4.1 Upgrades for SNEWS 2.0

SNEWS still operates under the same client that was first deployed when it came online. This client, in the modern era, will not suffice if SNEWS wishes to provide additional information to optical astronomers. Extensive work has been undertaken to develop a new client that will be able to meet the challenges neutrino detectors face in expanding the network to gravitational wave experiments and handling triangulation calculations. This new client is referred to as snews_pt, known as "SNEWS Publishing Tools" [48]. Access to the client is restricted to SNEWS members. The snews_pt client makes use of the hop/scimma client [49], which is a publishing and subscription tool that has seen implementation in the multi-messenger astronomy (MMA) field. The open-access documentation for snews_pt can be found at [48]. As of May 2023, many of the experiments currently operating have been able to install the latest version of snews_pt, and have been stress testing the software suite through regular fire drills (the collaborations method of testing ALARM issuing and subscription, as well as identifying bugs and nuisances). This is not the sole improvement made to the SNEWS 2.0 infrastructure. With the advent of many high statistics detectors, dedicated software has been developed to handle the triangulation calculations, as well as to investigate different pointing and triangulation techniques across a wide array of experiments. This software suite is known as `SNEWPDAG` [50], which the work outlined in this thesis can be added too.

4.4.2 Summary of recent work

Triangulation will rely on contributions from multiple neutrino experiments. There have been numerous studies done over the past half a decade, here the focus is placed on two studies done by members of the SNEWS collaboration and their collaborators. The first study [6], aimed to use the time differences between the first event observed in current (Super-K) and next-generation detectors (i.e. HyperK, JUNO, and DUNE) to triangulate to a core-collapse supernova. To reduce errors induced by the presence of backgrounds, the "first" event is located by requiring that there are two events within 15 ms of one another. Using a similar technique that was derived in [51], the time difference between a pair of detectors is

$$t_{ij} = \frac{(\vec{\boldsymbol{r}}_{\boldsymbol{i}} - \vec{\boldsymbol{r}}_{\boldsymbol{j}}) \times \hat{\boldsymbol{n}}}{c} \tag{4.3}$$

where $\vec{\boldsymbol{r}}_{\boldsymbol{i}}$ and $\vec{\boldsymbol{r}}_{\boldsymbol{j}}$ are the location of detectors i and j respectively, and $\hat{\boldsymbol{n}}$ is the direction of the SN. A chi-squared formula is used to assess the probability that a given SN location gives rise to the time differences measured by the detectors at the aforementioned locations, with the following explicit definition

$$\chi^2(\alpha, \delta) = \sum_{i,j}^{i<j} \left(\frac{(t_{ij}(\alpha^{'}, \delta^{'}) + B_{ij}) - t_{ij}(\alpha, \delta)}{\sigma_{t,ij}} \right)^2 \tag{4.4}$$

where $\sigma_{t,ij}$ is the calculated time difference uncertainty; B_{ij} is the mean bias for a given pair of detectors; $t_{ij}(\alpha^{'}, \delta^{'})$ is the true time difference between a pair of detectors; and $t_{ij}(\alpha^{'}, \delta^{'}) + B_{ij}$ is the expected measured time difference. Here, α and δ are the right ascension and declination, respectively. The bias is intended to account for variations in the first event time that arise due to varying detector masses, flavour sensitivities, and detector thresholds. The explicit derivation will not be discussed here, for additional details the reader is directed to [6]. The supernova signals in each detector were generated using the event rate calculator `SNOwGLoBES` (see §6.1), making use of the time-dependant 8.8 $M_\odot$ model defined in §3.1. The CCSN is assumed to occur with right ascension $\alpha = -94.4°$

and declination $\delta = -28.92°$, at varying distances. For a CCSN at 10 kpc, observed by HyperK, DUNE, JUNO, and IceCube, with the assumption of normal mass ordering (NMO) and the Mikheyev-Smirnov-Wolfenstein (MSW) effect present, the skymap produced via Equation 4.4 can be found in Figure 4.4.

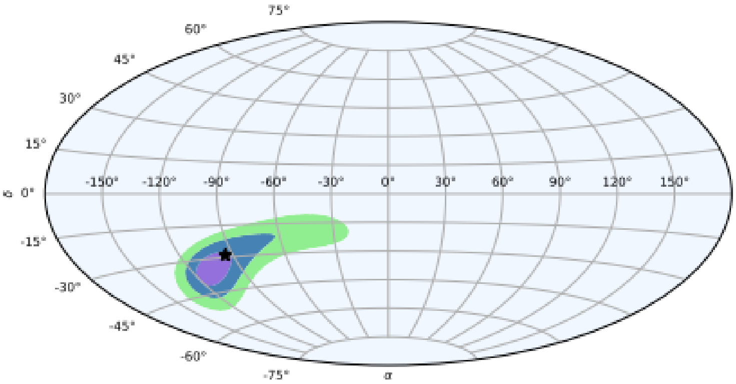

Figure 4.4: Sky area from combining first events times in IceCube, HyperK, DUNE, and JUNO at 10 kpc with NMO. Courtesy of [52].

Each coloured region (purple, blue, green), corresponds to confidence intervals (1σ, 2σ, 3σ), respectively. The results of this work illustrate that triangulating with current and next-generation neutrino detectors is possible and remarkably robust for such a simple technique. From this study, it is shown that accurate directional information can be provided for a CCSN, even if detectors that are able to point back to the supernova are offline.

The second triangulation investigation [52] makes use of the time differences between light curves observed in pairs of detectors. In this context, the "light curve" refers to the simulated/measured event rate in the detector, organized into 50 ms bins. In [52], two techniques were investigated to compare the observed light curves: a chi-squared technique and normalized cross-correlation method. For the chi-squared formulation, the explicit definition is

$$\chi^2(\tau) = \sum_{t_i=t_{min}}^{t_{max}} \frac{\left((n_{t_i-\tau} - m_{t_i}) - E(n_{t_i-\tau} - m_{t_i})\right)^2}{V(n_{t_i-\tau} - m_{t_i})} \tag{4.5}$$

where $n_{t_i-\tau}$ is the number of events observed in the first detector in time bin $t_i - \tau$, m_{t_i} is the number of events observed in detector two in time bin t_i, $E(n_{t_i-\tau} - m_{t_i})$ and $V(n_{t_i-\tau} - m_{t_i})$ are the expectation and variance, respectively, of the difference in events between the two detectors. For the normalized cross-correlation technique, it was implemented as

$$C(\tau) = (n * m) = \frac{1}{N} \sum_{t_i=t_{min}}^{t_{max}} n_{t_i} m_{t_i-\tau} \tag{4.6}$$

where τ is offset between the two detectors, n_{t_i} and $m_{t_i-\tau}$ is each detectors respective events in a given time bin. The results of both of these analysis techniques were comparable with one another. In comparing the light curves of IceCube, HyperK, JUNO, and KM3NeT, the CCSN at a distance of 10 kpc can be restricted to a celestial region of 70 ± 10 deg^2. The exploration of all these techniques, from both [6] and [52], gives the SNEWS toolkit more options that will assist it in its triangulation campaign.

4.4.3 Motivation for this work

Not all neutrino detectors are created equal. Owing to inherent flavour sensitivities, target masses, and interaction cross sections, each current and next-generation neutrino telescope will observe a varying degree of neutrino interactions for the next SN. Table 4.1 has the number of expected events in both current and next-generation neutrino detectors for a given model, adapted from [53].

Detector	Primary Flav. Sens.	Events	Status
HALO	ν_e, ν_x	30	Active
HALO-1kT	ν_e, ν_x	750	Future
SNO+	$\bar{\nu}_e$	300	Active
NOvA	$\bar{\nu}_e$	4000	Active
LZ	ν	100	Active
DUNE	ν_e	3000	Future
KM3NeT	$\bar{\nu}_e$	100,000	Future
LVD	$\bar{\nu}_e$	300	Offline
Borexino	$\bar{\nu}_e$	100	Offline
DarkSide-50	$\bar{\nu}_e$	100	Active
DarkSide-20k	ν	250	Future
XENON	ν	100	Active
Darwin	ν	800	Future
Super-K	$\bar{\nu}_e$	70,000	Active
KamLAND	$\bar{\nu}_e$	300	Active
Daya Bay	$\bar{\nu}_e$	100	Active
JUNO	$\bar{\nu}_e$	1,000	Active
PandaX	ν	100	Active
IceCube	$\bar{\nu}_e$	1,000,000	Active

Table 4.1: Tabulation of experiments currently or planning on participating in SNEWS 2.0 with their expected number of events, as calculated from a generic SN model at 10 kpc. Adapted from from [53].

What becomes abundantly clear is that both HALO and HALO-1kT will have lower overall statistics compared to other large scale detectors such as DUNE and HyperK (HALO and HALO-1kTs low statistics are a result of their interaction channels and discussed in more detail in §5.3.3). The studies carried out in [6, 52] and discussed in detail in §4.4.2 become less feasible for HALO and HALO-1kT, as the variation in first event time from a CCSN neutrino signal is both too large and model dependant. The variation in first event time for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc can be found in Figure 4.5 alongside the other models defined in §3.

The first event in a SN neutrino burst is highly dependant on the target mass, flavour sensitivity, and efficiency of the detector. Between SNO+ and HALO-1kT, making use

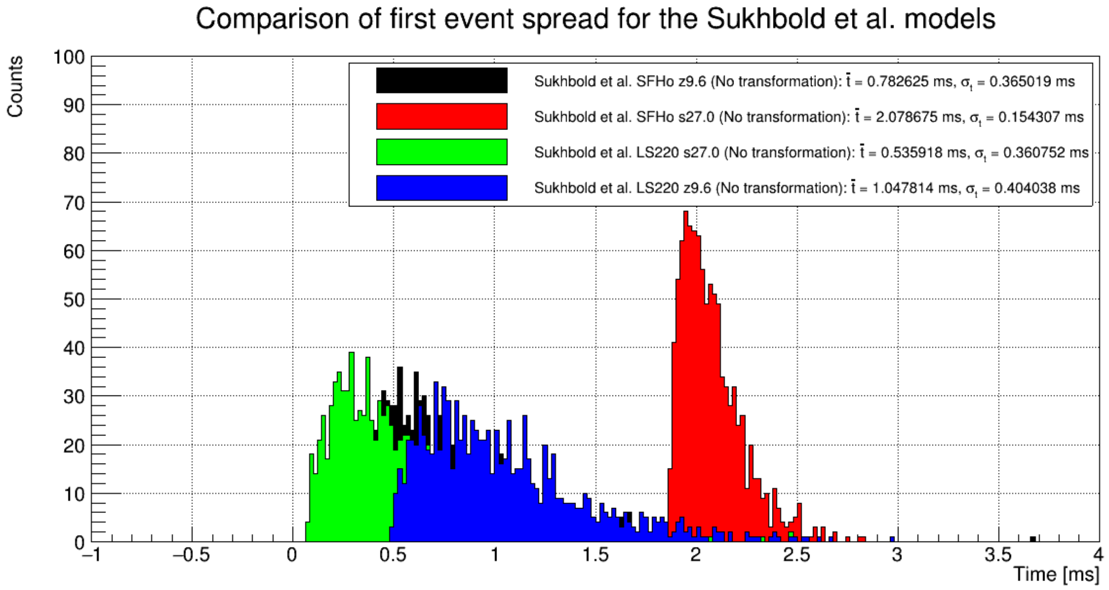

Figure 4.5: First event in 10^3 SN bursts simulated through HALO-1kT, for varying CCSN models at 1 kpc. Obtained using the event rate calculators and simulation tools outlined in §6.

of the same model at the same distance, an even greater variation between the two peaks would be observed since SNO+ will contain low energy proton-scattering events prior to t_0 in all models utilized in this thesis. To effectively triangulate, SNEWS 2.0 *must* establish a common reference time, as the first event is not an accurate representation of the neutrino wave fronts arrival at a given detector. Not only would this rid a triangulation calculation of any potential bias inherent to the wave fronts arrival time and first event, it would also open the door to allow other neutrino detectors to meaningfully contribute to the SNEWS 2.0 triangulation campaign.

Within the SNEWS 2.0 NSF proposal in 2018, the idea was put forth to use the time of core-bounce, as defined in §2.1 and generalized for each model contained in this report in §3. In supernova simulations, this is a well-defined time and traditionally set to occur at $t_0 = 0$ s. The benefit for t_0 is two-fold. Firstly, immediately following t_0 is the rapid rise in ν_e flux from electron capture reactions that produce the abundance of neutrons present in the PNS. In detectors that are primarily sensitive to ν_e, the t_0 in the observed signal shall be followed by a rapid rise in the number of events observed in a given time window. For HALO and HALO-1kT, the simulated models in this report saw up to 15% of all their events in the neutronization peak (within 20 ms of core-bounce). Such a large event density should make it possible for automated fitting routines to extract the time of core-bounce with excellent precision and communicate it to SNEWS. An additional benefit of using t_0, as opposed to some other reference time, is that for CCSN simulations that collapse into a black hole prior to detonation, t_0 is still well defined and the supernova is expected to emit a significant burst of neutrinos prior to black hole formation. Such a scenario is of tremendous importance as early follow-up observations will only be possible if an alert is issued from either neutrino or gravitational wave experiments.

This work will focus on three neutrino experiments, HALO, HALO-1kT, and SNO+ (all of which are outlined in §5). We will use both neutrino event rate and event calculators (Section 6.1) to develop a vector of events that will be simulated through our detector Monte Carlo simulation (here after, we elect to use the term Monte Carlo to refer to each detectors simulation code). Each vector provided to the Monte Carlo simulation will have its number of generated events Poisson fluctuated prior to simulation. A full detector Monte Carlo simulation treatment is utilized to ensure that detector responses can be accounted for. This will realistically account for detector capture efficiencies, propagation times, and additional interactions. A background will be added to the data post Monte Carlo simulation, and is a Poisson fluctuation of either observed or predicted background rates. From there, a wide range of analytical tools shall be utilized on the observed time series to quantify these detector's abilities to extract t_0.

As shall be outlined in the forthcoming chapters, these tools have seen full development in both HALO and HALO-1kT, with preliminary expansion into SNO+. The work is far from over. The intent has always been to expand these tools to other current and next-generation neutrino detectors. This work is currently underway, with these tools being incorporated into `SNEWPDAG`. The implementation of these tools into the SNEWS collaboration tool chains will allow the extraction of t_0 to be quantified in other detectors for which we do not have access to their Monte Carlo. Studies such as [6, 52] will also be performed to assess the feasibility of triangulating with a common reference time.

CHAPTER 5

Supernova neutrino experiments

5.1 Homestake and the Sudbury Neutrino Observatory

In the 1960s, Raymond Davis Jr., a researcher with the University of Pennsylvania, established a solar neutrino experiment at Homestake mine. Known formally as the Homestake experiment, the detector consisted of a 100,000 gallon ($\approx 370\,\mathrm{m}^3$) perchloroethylene target volume at 4,200 m.w.e depth (where m.w.e is meter water equivalent, which represents the amount of water that would be equivalent to the overburden above the detector with respect to cosmic ray attenuation). With the intention of measuring the solar neutrino flux, assumed to be entirely ν_e, the interaction mode for incident electron neutrinos was

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^- \tag{5.1}$$

where the neutrino must have ≥ 0.81 MeV in kinetic energy for this interaction to be available. ν_e arise from the PP chains described in §2.4.1.1, through the β^+ decay of ^{8}B or electron capture on ^{7}Be. Since the neutrino must have at least 0.81 MeV in kinetic energy, this interaction is primarily a probe of the PPIII chain (ν_e from ^{8}B) [54]. Measurements of these ν_e interactions in the perchloroethylene target is made possible through the ^{37}Ar instability. The half-life of ^{37}Ar is 35 days [55], where it will decay via electron capture

$$^{37}\text{Ar} + e^- \rightarrow {}^{37}\text{Cl} + \nu_e + \text{Auger } e^- \tag{5.2}$$

where the Q-value of this interaction is 0.813 MeV. Detection of the ^{37}Cl atom in Equation 5.2 is made possible through the emission of the Auger electron. The interactions between the solar ν_e and ^{37}Cl occur at a slow rate, happening every few days on average. Every few weeks, the ^{37}Ar was flushed out of the perchloroethylene target with a helium gas and concentrated into small proportional counters, where the Auger electrons were counted. In doing so, Davis and his team produced the solar ν_e flux measurement of $2.56 \pm 0.16 \pm 0.16$ solar neutrino unit (SNU) [56]. A SNU is equivalent to the neutrino flux that results in one capture per second for every 10^{36} target atoms. This measurement was a surprising result to researchers at the time, as the expected rate from solar models ranged anywhere from $\approx 2.5 - 3.5$ times larger, well outside the statistical or systematic errors reported by Davis and his team. This discrepancy became known as the "solar neutrino problem" and was a source of great theoretical and experimental interest for many decades.

The Sudbury Neutrino Observatory (SNO) was a neutrino experiment aimed at solving the solar neutrino problem. The detector consisted of an acrylic vessel (AV), 12 m in diameter, that housed 1 kilo-tonne of heavy water (D_2O). Surrounding the AV was an array of 9,600 PMTs, which were mounted on a PMT support structure. Within the D_2O volume, the following neutral current (NC) reaction, mediated by the Z^0 boson

$$\nu_x + d \rightarrow n + p + \nu_x \tag{5.3}$$

would occur. Where d is the deuterium atom and ν_x is any of the neutrino flavours. For future discussion, the charged current (CC, mediated by the $W^\pm$ boson) and elastic scattering (ES) processes in SNO were

$$\nu_e + d \rightarrow p + p + e^- \tag{5.4}$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad . \tag{5.5}$$

Equations 5.4 and 5.5 are directly observable via Cherenkov light from the charged particles propagating in the heavy water. SNO, to "solve" the solar neutrino problem, would have to show that the measured deficit of ν_e, was a consequence of ν_e transforming into ν_μ or ν_τ; this phenomena is known as neutrino oscillations. To do this, the NC current interaction noted above would need to be disentangled from the CC and ES interactions.

SNO operated from 1998 to 2006; the data taking runs occurred in three distinct phases. The three phases of the SNO experiment were as follows:

- Phase 1 - D_2O (heavy water).
- Phase 2 - D_2O enriched with salt (NaCl).
- Phase 3 - NCD array deployed within the D_2O volume.

The term NCD refers to "neutral current detector". Although they are ^{3}He proportional counters, the name stems from their use in detecting the out-going neutrons from the NC interaction. NCD is a legacy term that still sees use in HALO, where the proportional counters from SNO are now used (see §5.3). The measurement of the rate of the NC reaction (Equation 5.3) relies on the efficient detection of the out-going neutron. A different strategy was employed in each of the three phases. In the first phase, neutron capture on deuterium

$$n + d \rightarrow {}^{3}\mathrm{H} + \gamma \tag{5.6}$$

produces a 6.25 MeV γ. This gamma has a high probability of Compton scattering and producing an electron above the Cherenkov threshold (the minimum energy for a given charged particle to create Cherenkov radiation in a medium) within the heavy water volume. With the addition of salt in the second phase, there were two distinct benefits. The first is the addition of an element with a higher neutron capture cross section (chlorine) and the second is the out-going γs from this process total 8.6 MeV, an energy about one third higher than the single γ from neutron capture on deuterium. With the addition of salt in the heavy water volume, the neutron capture efficiency doubled and remained high at large radial positions (see Figure 5.1).

In the third phase, an array of NCDs were deployed in the heavy water volume. By doing so, the observation method of these out-going neutrons was drastically changed. The goal of the third phase was to verify systematic uncertainties in the measurement of the NC flux, reduce or eliminate correlations in the CC and NC signal separation, and improve the general statistical precision of the experiment [57]. By switching the method of observing the neutron from Cherenkov light (indirect) to neutron capture in a ^{3}He proportional counter, the NC and CC interactions can be disentangled in flux reconstruction

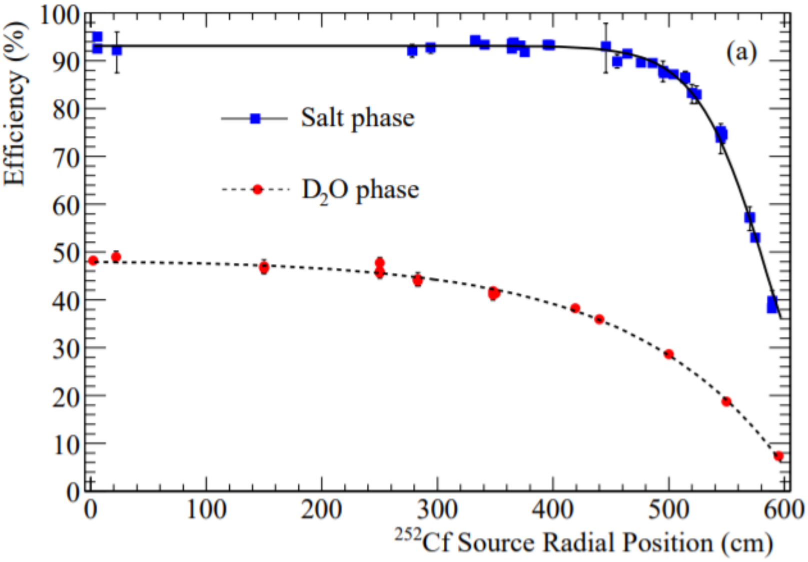

Figure 5.1: Courtesy of [57]. Depicted above is the neutron capture efficiency in the heavy water for pure D_2O and D_2O enriched with NaCl at 0.195% by weight.

since the CC interaction will not produce neutrons that undergo neutron capture on the ^{3}He in the counters.

SNO proved to be a tremendous success. In combination with Super-Kamiokande, they confirmed that the solar deficit was a direct consequence of ν_e converting to ν_μ or ν_τ through the MSW effect. In combination with results from KamLAND, this would prove that the neutrinos were not massless particles. For this, Arthur B. MacDonald of the SNO experiment and Takaaki Kajita of the Super-Kamiokande experiment shared the 2015 Nobel Prize in Physics.

5.2 SNO+

5.2.1 Design and construction of SNO+

With the completion of SNO, it was thought that the experiment could be retooled to search for other, rare physical events. SNO+ was proposed as a neutrinoless double beta decay search ($0\nu\beta\beta$), making use of the primary infrastructure that had been used in the original SNO (i.e the AV, PMTs, PMT support structure). Instead of implementing D_2O within the AV, SNO+ will have three phases that make use of different configurations. They are

- Phase 1 - AV filled with ultra-pure water (UPW).
- Phase 2 - 780 tonnes of linear alkylbenzene (LAB).
- Phase 3 - 780 tonnes of LAB loaded with 3.9 tonnes of natural occurring tellurium (Te).

The main physics search will be undertaken in phase 3. To shield against external backgrounds from the surrounding rock, the AV is surrounded by 7 kt of UPW. The decision to use LAB over other liquid scintillators arose from the incompatibility between existing liquid scintillators and the AV. LAB was the natural choice as it would be compatible with the AV and had seen small scale deployment in other detectors, while SNO+ developed the infrastructure to pioneer the use of LAB on such a large scale.

5.2.2 Physics goals

Amongst the detectors utilized in this report, SNO+ is unique as it has a number of overarching physics goals. Comparatively, HALO and HALO-1kT only have one, super-

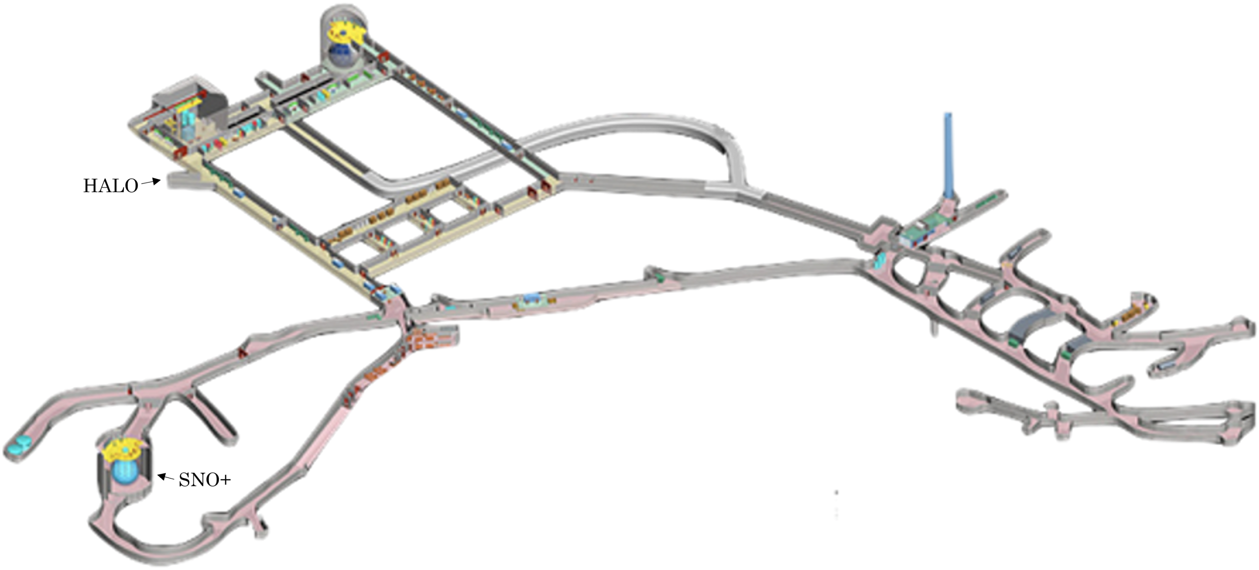

Figure 5.2: Overview of SNOLAB, located 2 km below surface in Creighton Mine, outside Sudbury, ON. The location of SNO+ and HALO are labelled. Courtesy of [58].

nova neutrinos. Since SNO+ is not the primary focus of this work, only a cursory look at neutrinoless double beta decay and supernova physics shall be given, which shall be discussed in §5.2.2.1 to §5.2.2.2. Other physics goals for SNO+ include reactor neutrinos, geoneutrinos, and like its predecessor, solar neutrinos. As exciting as these opportunities are, especially concerning reactor neutrinos which have been confirmed in the recent water phase of the detector [59] (opening up the possibility of low cost monitoring of nuclear sites), they do not contribute to the overall analysis contained within this thesis.

5.2.2.1 Neutrinoless Double Beta Decay

Beta decay is a well understood and thoroughly investigated phenomena. It was through this process that the neutrino would first be theorized. Studies of radioactive decay (whether that be alpha, beta, or gamma) had envisioned two products, the daughter nuclei and some emitted particle. It was for this reason that the measurement of the energy spectrum of out-going electrons from beta decay came as a shock to James Chadwick and his team. For a two body decay from one nucleus to another, where the ejected particle is either a positron or electron (in the case of β^+ or β^- decay respectively), the kinetic energy of the electron/positron would be equal to the Q-value of the interaction. In early measurements of the energy spectra of out-going electrons, the distribution was continuous, which would imply a third particle not known to physicists at the time. Wolfgang Pauli postulated that this unknown particle was a lightweight and neutral particle, as it did not interact within the detector volume. It would be given the name “neutrino”, meaning “little neutral one” [54]. It was discussed in §2.4.1.1 and 2.4.1.2 the role that β-decay plays in pushing nuclei to stability, while assisting a star to fuse heavier elements (either through the PP chains or CNO cycle). Whereas ordinary β-decay is allowed in a wide range of nuclei

$$(A, Z) \rightarrow (A, Z+1) + e^- + \bar{\nu}_e \tag{5.7}$$

the also allowed process of double β-decay ($2\nu\beta\beta$),

$$(A, Z) \rightarrow (A, Z+2) + 2e^- + 2\bar{\nu}_e \tag{5.8}$$

where two electrons and two electron antineutrinos emerge from the decay together, is very rare. Typical half-lives of β-decay processes range anywhere from fractions of a second, to years. In comparison, $2\nu\beta\beta$ half-lives are typically on the order 10^{18} to 10^{21} years [60]. If it is energetically allowed for a nucleus to undergo both β-decay and $2\nu\beta\beta$, only the the β-decay will be observed due to its orders of magnitude shorter half-life. There are few isotopes where β-decay is energetically forbidden and it is only energetically allowed for $2\nu\beta\beta$ to happen, making these isotopes ideal candidates for studying $2\nu\beta\beta$. In both β-decay and $2\nu\beta\beta$, lepton number conservation is preserved (where e^- and ν_e have $L = 1$,

e^+ and $\bar{\nu_e}$ have $L = -1$, where L is the lepton number of each corresponding lepton). Provided that the neutrino is a Majorana fermion (opposed to Dirac), the neutrinoless double beta ($0\nu\beta\beta$) mode of decay can become available to the nuclei

$$(A, Z) \rightarrow (A, Z+2) + 2e^- \tag{5.9}$$

where the neutrinos annihilate one another, as they are Majorana fermions, resulting in the electrons carrying away the full energy liberated from the decay. In practice, $0\nu\beta\beta$ is a tremendously difficult interaction to observe, if at all possible. For the traditional $2\nu\beta\beta$, the energy liberated in the decay will go into the kinetic energy of the out-going neutrinos and electrons. With the absence of the neutrinos in the decay, this energy is now shared entirely by the two electrons. For $0\nu\beta\beta$ searches, the signal is the measured energy of the out-going electrons, where the summed energy of these two electrons is the reaction Q-value. These searches have been underway for the better part of three decades, yet no observation of $0\nu\beta\beta$ has been made. Provided that the neutrino was a Majorana fermion, detecting $0\nu\beta\beta$ would still face many challenges. In $0\nu\beta\beta$, the electrons will carry a discrete energy for every decay; when measured, the energy is smeared by the detectors resolution. Since the decay of $2\nu\beta\beta$ is continuous up until just prior to the Q-value (as *some* of the energy liberated in the decay will be carried off and lost by the neutrinos, which doesn't contribute to the signal), detectors with poor energy resolution or high background are unable to observe this theoretical mode of decay. Nuclei that are capable of undergoing $2\nu\beta\beta$, with β-decay energetically unallowed, are uncommon [54].

SNO+ will use 1.3 tonnes of ^{130}Te, naturally occurring in the 3.9 tonnes Te that will be dissolved in the LAB (as mentioned earlier this will take up the third phase of the detector operations). ^{130}Te is chosen for a number of reasons [61], chiefly

- Of the isotopes known to undergo $2\nu\beta\beta$, ^{130}Te is amongst the nuclei with the longest $2\nu\beta\beta$ half-life. Currently, the best measurement of the half-life of $2\nu\beta\beta$ in ^{130}Te is $\{8.2 \pm 0.2\,(\text{stat}) \pm 0.6\,(\text{sys})\} \times 10^{20}$ yrs [62]. For ^{136}Xe, it has a measured $2\nu\beta\beta$ half-life of $\{2.165 \pm 0.016\,(\text{stat}) \pm 0.059\,(\text{sys})\} \times 10^{21}$ yrs from EXO [63]. The length of these half-lives makes both ^{130}Te and ^{136}Xe ideal candidates for $0\nu\beta\beta$ searches.

- The Q-value for the interaction is 2.527 MeV. Most γ's from the decay of ^{238}U and ^{232}Th, naturally occurring nuclei with long half-lives and decay chains, are well below this threshold. An exception is the 2.615 MeV γ that arises from the decay ^{208}Tl, in the ^{232}Th chain. The branching ratio of this decay is 0.3555 (or 35.58%). In the ^{238}U chain, the ^{214}Bi nuclei arises and can β^- decay with a Q-value of 3.269 MeV [44].

- An additional benefit in using ^{130}Te over other nuclei capable of decaying via $2\nu\beta\beta$, is that ^{130}Te has no absorption lines in the visible wavelengths, thus reducing potential loss in intensity of scintillation light produced from charged particles in the detector volume.

SNO+ plans for the third phase in 2024 and hopes to achieve a sensitivity lower limit on $0\nu\beta\beta$ of 2.1×10^{26} yrs [61]. With the completion of the main physics program, there are currently plans in place to expand the sensitivity of the detector by increasing the amount of ^{130}Te loaded into the detector, which could achieve sensitivities of $\geq 10^{27}$ yrs.

5.2.2.2 Supernova neutrinos

Of main interest to this work is the ability for SNO+ to observe supernova neutrino events. As it stands, members of the SNO+ collaboration are working on implementing a SNEWS 1.0 alert system within the data acquisition system (DAQ) of the detector. This will add yet another high-statistics neutrino detector to the SNEWS network, potentially opening the door for improved triangulation capabilities. Like all other detectors capable of observing supernova neutrinos, the experimental aim of SNO+ is to constrain the SN mechanism and neutrino properties. Neglected up to this point, has been a discussion of pre-supernova neutrinos. Super-K, SNO+, and KamLAND have all investigated the potential observation of pre-SN neutrinos as an additional early alert channel. These pre-SN neutrinos arise in the later stages of burning due to the accelerated rate of burning in the hours leading up to core-collapse. In the case of Super-K [64], when loaded with a concentration of 0.01% Gd (which assists in identifying the out-going neutrons from the IBD interaction channel), low energy pre-SN $\bar{\nu_e}$ should be produced in sufficient quantities to be observable out to 600 pc. As part of her Doctoral Thesis [65], Janet Rumleskie conducted similar studies for SNO+. SNO+ is now providing an early alert up to 100 hours prior to the onset of collapse, at a maximum distance of 640 pc (these results are dependant on supernova model, neutrino mass ordering and hierarchy). For detectors such as Super-K, SNO+, or KamLAND, an early alert of the impending supernova neutrinos provides other detectors that are down for any reason, potential time to get back online and ready to observe the burst. Having detectors that are capable of detecting a core-collapse supernovae via its neutrinos in the final stages of burning, SNEWS will have another layer of support to ensure that the neutrino signal is observed by as many detectors as possible and the most can be made of this once in a generation event.

5.2.3 Supernova neutrino observation

Although SNO+ will observe neutrino events from many sources other than supernovae, this work is solely concerned with neutrinos that trace their origin to a core-collapse supernova. As such, this subchapter will focus solely on interactions at supernova neutrino energies. SNO+ is primarily $\bar{\nu_e}$ sensitive, as its active detector volume consists of organic liquid scintillator. There exist two primary interaction modes that will make up the vast majority of interactions in the LAB. The first primary interaction mode is inverse beta decay (IBD), where $\bar{\nu_e}$ interact with free protons via

$$\bar{\nu_e} + p \rightarrow n + e^+ . \tag{5.10}$$

IBD events can be tagged with relative ease provided backgrounds are kept at a minimum. The signal consists of a prompt and delayed components, with the prompt component arising from the scintillation light produced from the positron as it deposits energy and annihilates. The $\bar{\nu_e}$ kinetic energy can be in-directly measured via [66]

$$E_{\bar{\nu_e}} = E_{e^+} + 1.3\,\mathrm{MeV} \tag{5.11}$$

The out-going neutron from the IBD interaction is the delayed component. The neutron will thermalize in the detector volume, with the possibility to undergo neutron capture on hydrogen via

$$n + \mathrm{H} \rightarrow d + \gamma \tag{5.12}$$

where the energy of the photon is 2.22 MeV. The time difference between the delayed and prompt signal is expected to be on the order of 200 μs [67].

The next primary interaction mode is neutrino-proton scattering via the NC interaction

$$\nu_x + p \rightarrow \nu_x + p \tag{5.13}$$

where the neutrino elastically scatters off the free proton in the detector volume, imparting a fraction of its kinetic energy to the free proton. This will be the dominant mode of neutrino interactions in SNO+; for the work done as a part of this thesis, an elevated energy threshold to suppress backgrounds from ^{14}C decay will minimize neutrino-proton scatterings contribution to the analysis found in §9.

For both Equation 5.10 and 5.13, the energy of the photon or prompt scintillation signal can be inferred from the number of PMTs hit for a given event ($\mathrm{N_{hit}}$). In the current detector configuration with pure LAB and no Te loading, the expected $\mathrm{N_{hit}}$ per MeV is approximately 500.

A sub-dominant interaction is neutrino electron scattering ($\nu - e$ scattering), which is through both the CC and NC channels, and is open to all neutrino flavours ν_x

$$\nu_x + e^- \rightarrow \nu_x + e^- \tag{5.14}$$

where the recoil momentum of the electron is strongly correlated to the incident neutrino's momentum. This interaction is a direct measure of the energy of the neutrino flux; since the rate of these interactions for SN neutrinos is small with respect to other channels, flux reconstruction using the observed events from this interaction mode in a single detector is a difficult process.

For interactions that occur on ^{12}C nuclei, the first is analogous to IBD, written as

$$\bar{\nu}_e + {}^{12}\mathrm{C} \rightarrow {}^{12}\mathrm{B} + e^+ \, . \tag{5.15}$$

Alongside this CC interaction, another interaction that can occur at supernova neutrino energies is

$$\nu_e + {}^{12}\mathrm{C} \rightarrow {}^{12}\mathrm{N} + e^- \, . \tag{5.16}$$

The sole NC interaction of interest is available to all neutrino flavours

$$\nu_x + {}^{12}\mathrm{C} \rightarrow {}^{12}C^* + \nu_x^{'} \tag{5.17}$$

where the ^{12}C nucleus will promptly de-excite and release a photon. The outlined interactions are not the only that can or will occur in the detector volume. As shall be discussed in §6.1.3, these are the only ones currently implemented in the neutrino event calculator sntools. Other interactions that occur on carbon are expected to make negligible contributions to the total neutrino yield, thus the simulations and analysis in later chapters shall not be hindered by the omission of these other interaction channels.

5.3 The Helium and Lead Observatory (HALO)

HALO is a dedicated supernova neutrino detector located at SNOLAB in Creighton Mine. A overview of the location of SNOLAB within Creighton Mine can be found back in Figure 5.2.

5.3.1 Design and construction of HALO

In the early 2000s, if one were to catalogue the neutrino telescopes capable of detecting the next core-collapse supernova, a particular pattern began to emerge. All of the neutrino detectors that observed SN 1987A, were primarily $\bar{\nu_e}$ sensitive, as they consisted of liquid scintillator (Baksan) and water Cherenkov detectors (Kamiokande-II and IMB). As SNEWS was formed, and more neutrino detectors came online, this discrepancy became more apparent. As of 2019, the bulk of the detectors in SNEWS are either liquid scintillator and water Cherenkov. Even with the addition of dark matter detectors for SNEWS 2.0, which can be sensitive to coherent elastic neutrino-nucleus scattering (CEvNS), there exists a need for a dedicated supernova neutrino detector that will be primarily sensitive to the ν_e flavour.

Explorations for a primarily ν_e sensitive detector composed of lead began in 1996 [68], when the Lead Astronomical Neutrino Detector (LAND) was proposed. The original Monte Carlo simulation design consisted of a tonne-scale lead detector instrumented with ^{10}BF$_3$ neutron counters previously used for detecting cosmic rays. The LAND detector existed solely in Monte Carlo simulations. Preliminary explorations with the LAND

Monte Carlo simulation showed that a tonne-scale lead detector would be competitive in its observation of the next core-collapse supernova. Additional exploration came in the form of OMNIS [69], the observatory for multiflavour interactions from supernovae. The OMNIS concept was an evolution of the Supernova Neutrino Burst Observatory (SNBO) which proposed using natural rock as a neutrino target, making use of ejected neutrons from nuclei as the signal. Neither of these experiments would exist outside of their Monte Carlo simulation, but elements of their design and overall physics aims were carried over to HALO.

SNO concluded its third phase in 2006. In this phase, the detector was instrumented with an array of ^{3}He proportional counters as a means of measuring the neutrons from the NC interaction in Equation 5.3; these neutrons were electron neutrinos from the decay of ^{8}B neutrinos above 2.2 MeV in the PPIII chain. The array of counters consisted of 36 proportional counters with a ^{3}He and CF_4 gas mixture, plus 4 additional counters with a gas mixture of ^{4}He and CF_4. ^{3}He was chosen for its large neutron capture cross section. As before, the ^{3}He proportional counters thus allowed for better disentanglement of the CC and NC flux. With the conclusion of the SNO experiment in 2006, much of the infrastructure would be reused in the years to come for the follow-up experiment SNO+. Prior to SNO being taken offline, investigations were done to quantify the ability for a future SNO+ to aid in the search for neutrinoless double beta decay ($0\nu\beta\beta$). SNO+ did not require the ^{3}He detectors, opening up the possibility for their use elsewhere. It was a perfect coincidence that around the period of SNO being decommissioned, a large volume of lead became available from a cosmic ray monitoring station in Deep River, Ontario. From the availability of these two components, HALO was born.

The Helium and Lead Observatory makes use of a 79 tonne lead volume instrumented with 128 ^{3}He proportional counters distributed uniformly in the lead volume. The lead is formed into an array of annuli, where each annulus has thick wings that allow for better stacking in the steel support structure. The term "block" is used to describe both the annulus and its wings. The isotopic composition of the lead is assumed to be given by the natural abundance, and can be found in Table 5.1.

Isotope	^{208}Pb	^{207}Pb	^{206}Pb	^{205}Pb
Abundance (%)	52.4	22.1	24.1	N/A

Table 5.1: Natural abundance of lead isotopes in HALO.

Each lead block is coated in a thin layer of paint to inhibit the formation of lead carbonate [70]. With all lead blocks placed within the steel support structure, there were 32 channels to place the counters in. Each channel was fit with a steel tube that houses the high density polyethylene (HDPE) moderator and ^{3}He proportional counters. This geometry is shown in Figure 5.3, where the front shielding of the detector has been removed and an internal view of the detector can be seen.

The proportional counters have a 5 cm diameter, with either a 2.5 m or 3 m length. With all 128 counters in place, they total 368 m in length, which are paired in 64 channels for readout. The nickel walls are 600 microns thick. The very high purity of these counters was required for the measurement of solar neutrinos in SNO, in HALO these purity requirements are not necessary but a nice benefit. The pressure from the gas inside the proportional counter is 2.5 atm, which is 85% ^{3}He and 15% CF_4 by pressure. The CF_4 is intended to act as a stopping agent, which reduces the track length of the daughter particles in the proportional counter. The shorter track length of the triton-

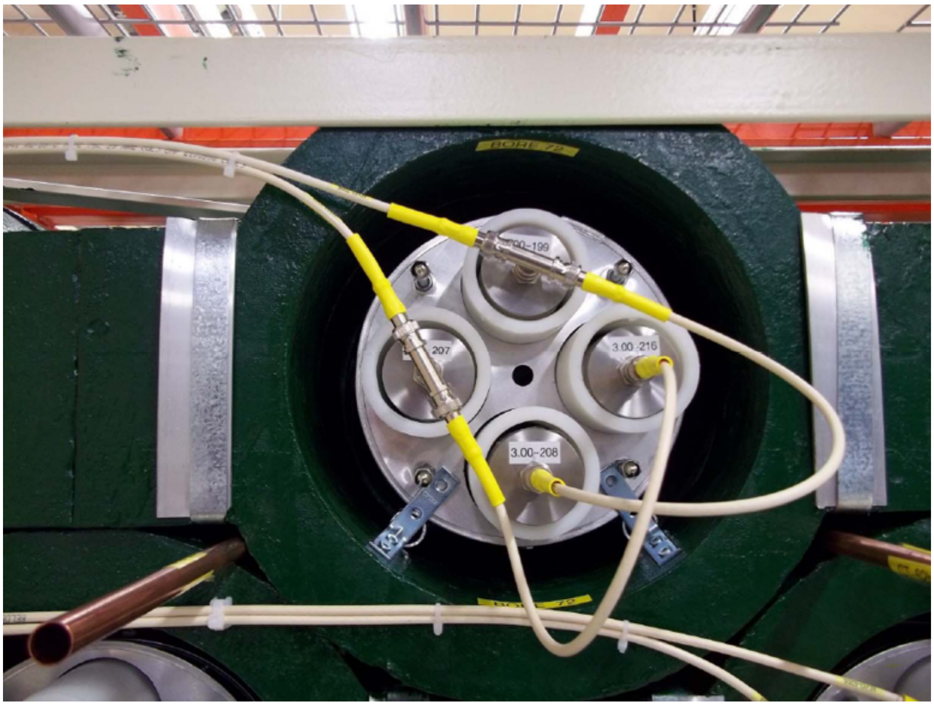

Figure 5.3: Internal view of the lead blocks alongside the HDPE moderators and ^{3}He proportional counters.

proton pair reduces the wall effect, as shall be discussed later. A schematic diagram of the ^{3}He proportional counters can be found in Figure 5.4.

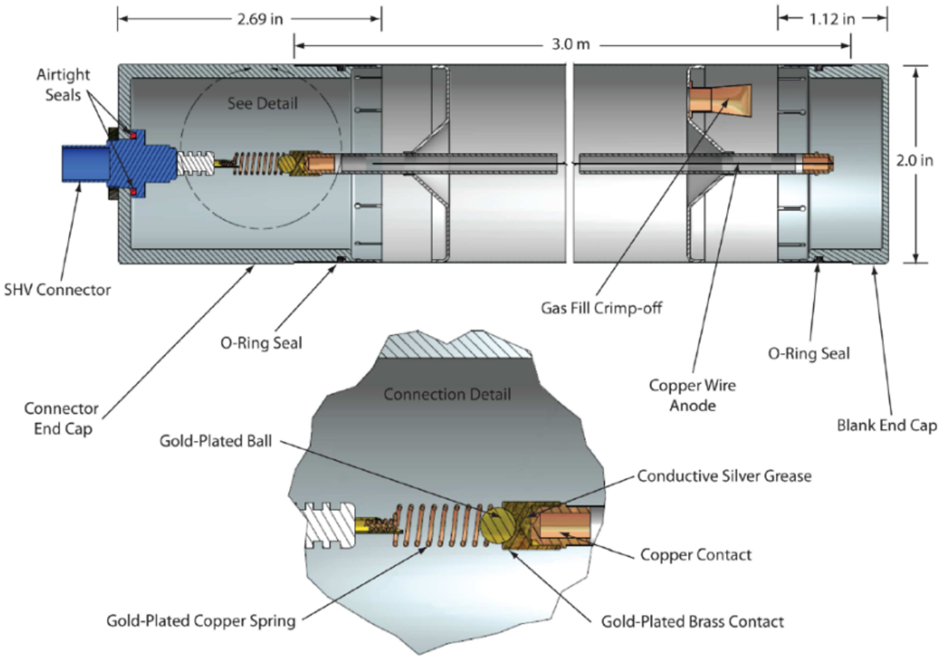

Figure 5.4: Schematic diagram of a ^{3}He proportional counter in HALO, courtesy of [71].

5.3.2 Physics goals

As the only physics goal of HALO is to observe core-collapse supernovae, the experiment bides its time, awaiting the next galactic supernova. There are two critical aims of HALO that have greatly informed nearly all aspects of its design. They are summarized as "long lifetime" and "high live-time". For any supernova neutrino detector, both of these features are of the utmost importance. Core-collapse supernovae are rare, it having been well over 30 years since the last was observed via its neutrino signal. For an individual experiment, it would be a tragedy if the detector was down for either maintenance or technical issues at the moment the neutrino signal from a CCSN reaches Earth. Alternatively, with many experiments operating under the pretense of having a set data collection period prior to being taken offline, it too would be problematic if CCSN neutrino detectors were decommissioned before an observation. This conundrum is what multi-purpose neutrino detectors face in the years to come. The design of HALO has tackled these problems head on.

To achieve a high live-time, HALO was designed with built-in redundancies for all of the major components. The experiment operates with two DAQs, only one of which is taking data at any given moment. They are both linked through an internal software routine called "sentry". Should one of the DAQs go offline, stop taking data, or encounter any other mission critical errors, sentry will automatically initiate a run (a 24 hr cycle of data taking) on the other machine. HALO uses the same software that the original SNO experiment (and SNO+) used, which is ORCA (Object-oriented Real-time Control and Acquisition). To power the pre-amps, which in turn readout and power the 64 channels of 128 proportional counters, there are two low voltage (LV) supplies. A pair of high voltage supplies exist to power the 128 proportional counters. For single points of failure, half the detector will continue operating and collecting data should the other half go offline. Design choices such as these contribute to HALO having a live-time of $\geq 99\%$ in recent years (see Figure 5.5).

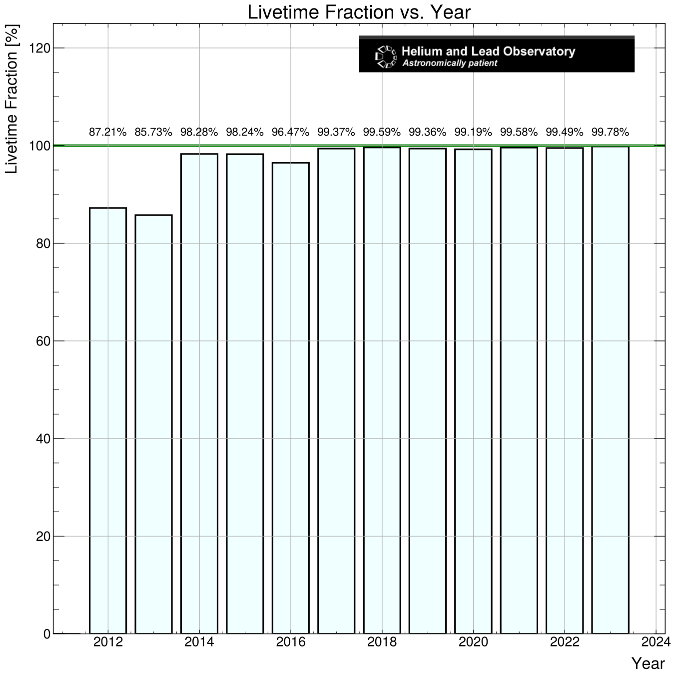

Figure 5.5: HALO live-time since coming online, courtesy of [72]. In recent years, HALO has been able to consistently achieve live-times in excesses of 99%.

5.3.3 Neutrino observation

HALO is primarily ν_e sensitive, making it an excellent flavour complement to existing supernova neutrino detectors. The modes of interaction are via the NC and CC interactions. The main contribution to observed interactions from a SN burst comes from the CC interaction on our lead target from electron neutrinos

$$\nu_e + (A, Z) \rightarrow (A, Z+1)^* + e^- \tag{5.18}$$

which has an electron neutrino kinetic energy threshold of 10.7 MeV for 1n reactions in ^{208}Pb (for 2n interactions, the threshold is 18.6 MeV). Since the primary target isotope is ^{208}Pb, the excited daughter is ^{208}Bi. The dexictation of the daughter nuclei is prompt, occurring via the following mode

$$(A, Z+1)^* \rightarrow (A-\eta n, Z+1) + \eta n + \gamma \tag{5.19}$$

where η is the number of neutrons emitted. In HALO, supernova neutrino energies allow for ν_e CC interactions to occur with $\eta = 1$ or 2. The ν_e CC interaction is expected to make up about $\approx 75\%$ of ν-Pb interactions from the SN burst. The ratio of 1n to 2n interactions is SN model dependant.

The next primary contributor of events in HALO is the NC interaction for neutrino scattering off heavy atoms

$$\nu_x + (A, Z) \rightarrow (A, Z)^* + \nu_x \tag{5.20}$$

where ν_x is any neutrino flavour and their anti-neutrino counterpart. The excited atom will de-excite via the following mode

$$(A, Z)^* \rightarrow (A-\eta n, Z) + \eta n + \gamma \tag{5.21}$$

The threshold for this ν-Pb NC interaction is 7.4 MeV for 1n (14.4 MeV for 2n) and is expected to make up the remaining $\approx 25\%$ (with the CC $\bar{\nu}_e$ interaction found in Equation 5.23) of the total signal. Since the only method of observation of the neutrino signal is via the proportional counters detection of the ejected neutrons, HALO is blind to both the γ

and e^- from both the ν_x-Pb NC interaction and ν_e-Pb CC interaction. The electron and gamma will be lost in the lead, as they will attenuate with insufficient energy to escape. In the statistically improbable case that either the electron or a Compton scattered electron from the photon reaches a counter, coincidence tagging with captured neutrons is not possible as the capture efficiency, the variation of counter that observed the event, and external backgrounds cannot be disentangled. To perform coincidence tagging, scintillator could be laced through the detector to observe any charged secondaries. Scintillators would compete with the primary signal for HALO as the out-going neutrons from the ν-Pb NC and ν_e-Pb CC interactions can capture on free protons via

$$n + p \rightarrow d + \gamma \tag{5.22}$$

where d is the deuterium atom, and the kinetic energy of the photon is 2.2 MeV. It is for this reason, amongst others, scintillator is not used to tag the different interaction modes. The only potential coincidence tagging that HALO can do is for 2n events.

The third and final interaction mode observed in ^{208}Pb is the $\bar{\nu_e}$-Pb CC interaction and occurs in the following way

$$\bar{\nu_e} + (A, Z) \rightarrow (A, Z-1) + e^+ \tag{5.23}$$

In liquid scintillators or water Cherenkov detectors, the analogous interaction is IBD. The kinetic energy threshold for IBD is 1.8 MeV in the aforementioned detectors, making it one of the best probes for neutrino supernova neutrinos. This channel is heavily suppressed in ^{208}Pb. ^{208}Pb is a double magic nuclei with a large neutron excess (^{208}Pb, Z = 82 & N = 126). Any process in lead that wishes to convert a proton to a neutron must contend with the Fermionic nature of these two particles and the few states available for the created neutron to occupy. This suppression effect is referred to as Pauli blocking, and in combination with the high Z number and the resulting Coulomb enhancements of Equation 5.18 (and Coulomb suppression of Equation 5.23), lead serves as an excellent probe of the ν_e component of the neutrino flux, as the $\bar{\nu_e}$ CC channel is negligible relative to the ν_e CC channel in terms of total events [73]. The neutrino-lead cross sections are discussed in §6.2, as these are not measured quantities and there has been a great deal of theoretical calculations done for varying lead isotopes, neutron multiplicity, and neutrino flavour.

The emitted neutrons from the outlined ν-Pb interactions will be ejected into the lead volume. As they scatter and interact with the surrounding matter, they will thermalize and quickly reach a temperature of $\approx$ 0.025 eV. ^{208}Pb has a small neutron absorption cross section (4.8×10^{-4} b) and a very large neutron scattering cross section (11.34 b) [74]. These two features reduce the probability of neutrons capturing in the lead instead of the counters and ensuring they are at thermal temperatures when they reach the ^{3}He gas. If the neutrons reach the gas, they can undergo neutron capture via

$$n + {}^{3}\mathrm{He} \rightarrow t + p + 764\,\mathrm{keV} \tag{5.24}$$

producing a triton-proton (t and p) pair. There is 764 keV of energy liberated in this interaction and carried away by the triton-proton pair. The proton carries 573 keV while the triton carries 191 keV. Deviations from this are possible if the neutron does not have adequate time to thermalize in the lead volume, resulting in more kinetic energy being available for the end products. For the out-going triton-proton pair, each particle will ionize the surrounding gas along their tracks (which are in opposite directions). The electrons will then drift towards the anode and the total deposited charge will by read-out via the DAQ. A consequence of the ionization occurring along the entire track is the possibility for partial energy depositions. If the neutron captures on ^{3}He sufficiently close to the wall, either particle can embed itself in the counters inner wall prior to a full deposit. Within HALO, this is known as the wall effect and an illustration can be found in Figure 5.6. Within the observed spectra from our counters lies two discrete shoulders that represent the complete loss of either the triton or proton. HALO has a measured average neutron capture efficiency of 28%, which came from calibration studies with a ^{252}Cf source in [75]; averaged because of the radial dependence on the neutron capture efficiency, peaking at the center of the detector.

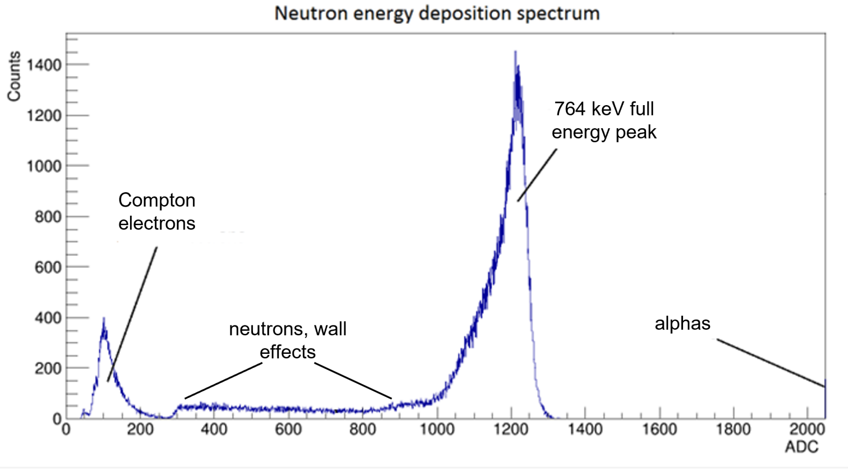

Figure 5.6: Observed energy spectrum in HALO from a californium neutron source. Courtesy of [75].

5.4 HALO-1kT

5.4.1 Design

Over the past decade, HALO has proven to be a tremendous success at its mission aims. HALO bides its time, patiently awaiting the next galactic supernova. For all its success, perhaps its greatest draw back is its relatively small scale. Convolving the relatively low detector mass with the small neutrino-lead cross sections, HALO will see few events at 10 kpc. This is reflected in the current SN burst trigger, which is only 4 events in a 2 second window. As such, to probe additional components of a SN from the observed neutrino signal, a larger detector is required. Even prior to HALO's construction, a proposal for a kilo-tonne scale lead detector designated HALO2 was put forth. With the quantity of proportional counters from SNO and total mass of lead from Deep River available, a kilo-tonne scale detector was out of the question.

The Oscillation Project with Emulsion-tRacking Apparatus (OPERA) experiment aimed to search for $\nu_\mu \leftrightarrow \nu_e$ oscillations in the same parameter region for atmospheric neutrino results in Super-Kamiokande, Soudan2, and MACRO [76]. The intent was to confirm atmospheric neutrino results using a fixed beam and measure $|\Delta m^2_{23}|$ and θ_{23} [77]. The detector was located in Hall C at Gran Sasso National Laboratories (LNGS) and made use of the CERN to Gran Sasso neutrino beam (CNGS) as its neutrino source, where the neutrinos travelled 732 km underground prior to reaching their target. The primary target volume in OPERA was bricks consisting of 56 layers of lead and 57 emulsion layers, where each layer of lead was 1 mm thick. The resulting bricks had a surface area of 10.2×12.7 cm^2, with a depth of 7.5 cm. A total of 62 target walls were arranged into 2 super-modules, each super-module consisting of 31 target walls which are made up of lead-emulsion bricks. The total detector mass was 1.25 kt. The CNGS beam primarily consisted of ν_μ neutrinos at a mean energy of 17 GeV, with minor contamination from the other neutrino species. OPERA concluded data taking in December of 2012, freeing up the lead to be used for another detector.

The next-generation of lead supernova detectors is the proposed HALO-1kT. HALO-1kT is currently in active R&D, which has been funded by the Natural Sciences and Engineering and Research Council (NSERC). The core-principles of HALO have not changed; long lifetime and high live-time. The detector would consist of 1 kt of lead, a factor of 12.66 larger than that in HALO. This effectively increases the SN distance sensitivity and statistics, which would allow deeper probes of the SN from the neutrino signal. In Figure 5.7, the current Monte Carlo simulation model of HALO-1kT is depicted next to the HALO schematic layout.

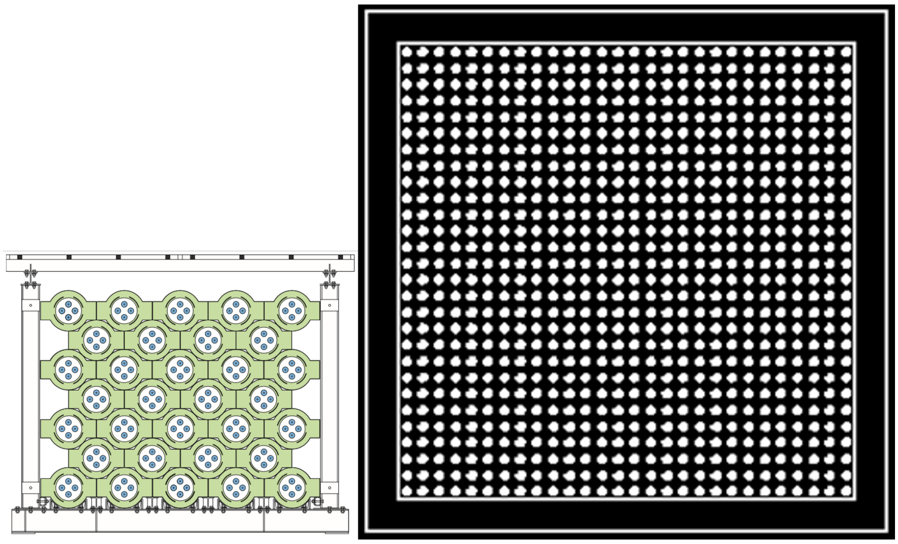

Figure 5.7: Comparison of scale between HALO and HALO-1kT. The image of HALO-1kT is from the detectors Monte Carlo.

In the current Monte Carlo, HALO-1kT consists of a 1 kilo-tonne of lead volume that is $4 \times 4 \times 5.5$ m^3. The lead volume is surrounded by water shielding that extends 30 cm outward from the lead volume. A 28×28 array of polystyrene moderators are interlaced throughout the detector in fixed increments. These moderators are 5.5 m in length and envelope the proportional counters, such that their inner radii are in contact with the counters and their outer radii are in contact with the lead. Within these moderators, stainless steel proportional counters are placed, each with a diameter of 2 in; a thickness of 0.02 in; and a length of 5.5 m. The gas composition is derived from the counters used in HALO, which is 15% ^{3}He and 85% CF_4 by pressure with a total pressure of 2.5 atm. The anode within the counters is modelled along the full length of the counter and is 50 microns in diameter.

Detector Name	Mass (t)	Capture Eff. (%)
HALO	79	28 (measured)
HALO-1kT	1000	53 (simulated)

Table 5.2: Comparison of basic properties for both the HALO and HALO-1kT detectors.

5.4.2 Outlook

HALO-1kT would use the same observation channels as HALO for ν-Pb interactions outlined in §5.3.3. The overarching physics goals would also be unchanged. With the improvements to the design and the magnitude increase in target mass, HALO-1kT would serve as a more effective probe of supernova neutrinos, peering further into the cosmos than its predecessor. The detector efficiency rises from 28% in HALO, to 53% in HALO-1kT (the important distinction is that the HALO-1kT efficiency is simulated, not measured). In a study by Andrea Gallo Rosso [78], it was shown that by combining the neutrino signal in HALO-1kT, Super-K, and JUNO, one could constrain some of the emission parameters of the neutrino flux (mainly the mean neutrino energy and total emitted energy). HALO-1kT would be unable to constrain these parameters on its own, owing to its lack of statistics and inherent flavour sensitivity, but by combining it with $\bar{\nu}_e$ sensitive detectors, one obtains a greatly improved constraint on the aforementioned parameters. Since detectors like JUNO and Super-K, amongst many other neutrino telescopes, have supernova neutrino detection as a secondary physics search, many of these detectors could be decommissioned and taken offline in the decades that follow. The intent is that HALO-1kT could remain, as HALO has, and the neutrino signal from a core-collapse supernova can be constrained even if other detectors have gone offline.

The experiment that was of primary focus for the analysis in the following subchapters was HALO-1kT. This work is but a small contribution amongst a number of studies by graduate students in recent years. These studies have focused on developing HALO-1kT and exploring its physics impact. Divyaben Patel, as part of her MSc thesis work [79], explored the implementation of a graphite layer on the outer edges of the detector. Graphite is an excellent neutron reflector/moderator, thus keeping external neutrons from the surrounding rock and lab out of the detector, and keeping those that are created from neutrino interactions in the lead confined to the detector volume. Her study concluded that the graphite should be of nuclear grade and ≈ 15 cm thick. Esther Weima, as part of her MSc thesis explored the backgrounds induced in proportional counters of varying outer shell composition. HALO was fortunate to have been lent the ^{3}He proportional counters from SNO, after it completed its third and final phase. SNO's ^{3}He proportional counters, with ultra-pure CVD (carbon vapour deposition) nickel bodies, are prohibitively expensive and the cost of filling a tonne scale detector with only these counters is not economically feasible (for HALO, as stated previously, only 368 m of ^{3}He counters were needed; in HALO-1kT, the total length would be 4.3 km). Esther spearheaded the effort looking into the solutions to this problem, which have been under investigation for the better part of a decade now. Four sample proportional counters were provided by commercial manufacturers and taken underground to SNOLAB for "counting". As provided, these detectors do not meet the strict 1 Hz requirement [80], as such, further investigation is required. What follows is my contribution to the excellent work carried out by my fellow students and collaborators.

CHAPTER 6

Geant4 Monte Carlo simulations

For this thesis, the most up-to-date versions of each experiment's Monte Carlo simulation was utilized. HALO, HALO-1kT, and SNO+ all utilize `Geant4` [81–83] as the main simulation software. Our simulation and analysis codes also make use of the software framework `ROOT` [84], and its companion software `RooFit` [85]. As is typical for SN neutrino simulations, an external neutrino event generator was used to determine the number of events for each interaction channel. For HALO and HALO-1kT, this work made use of `SNOwGLoBES` [86] as it has been used in numerous SN neutrino studies and has access to time-dependent fluxes (as discussed in §3). For SNO+, the detector material (LAB) has never been implemented into `SNOwGLoBES`. The only scintillator material built into `SNOwGLoBES` has a different composition of hydrogen, carbon, and oxygen. In early 2022, Sammy Valder, of the SNO+ collaboration, was looking into adopting a more modern and well maintained supernova neutrino event generator for the SNO+ Monte Carlo simulation. `sntools` [87] was chosen as the suitable candidate, which allowed for convenient access to supported neutrino flavour transformations and time-dependant fluxes provided by the `SNEWPY` [88] software package. Integration of SNO+ into `sntools` has been validated (i.e. number of events for a given channel are similar to previous or current neutrino event generators. Small differences in the yield arise due to the variation in scintillator compositions, but these are expected.

The following subchapters discuss the two SN neutrino event generators used throughout this thesis and how their outputs are loaded into each detector's Monte Carlo simulation. We then discuss how these simulation outputs are staged for analysis.

6.1 Neutrino event rate calculations

6.1.1 SuperNova Observatories with General Long Baseline Experiment Simulator (`SNOwGLoBES`)

`SNOwGLoBES` is an event rate generator that folds neutrino fluxes with detector cross sections and is used for a wide range of neutrino physics (i.e. supernova, solar, etc.). Public access is available at [86]. In the latest release of `SNOwGLoBES`, there are three supernova models present, Livermore [89], GVKM [90], and the Garching model (as discussed in §3). The Livermore and GVKM models are not used in this thesis as they are implemented as fluences (meaning time integrated flux) and do not provide the necessary time-dependence to extract t_0. The Garching model's neutrino flux is time-dependent by

$$\phi(t, E_\nu) = \mathcal{N}\left(\frac{E_\nu}{\langle E_\nu(t)\rangle}\right)^{\alpha(t)} \exp\left[-(\alpha(t)+1)\frac{E_\nu}{\langle E_\nu(t)\rangle}\right] \tag{6.1}$$

where E_ν is the neutrino energy, $\langle E_\nu(t)\rangle$ is the mean neutrino energy, $\alpha(t)$ is the Fermi-Dirac pinching parameter, and $\mathcal{N}$ is the normalization constant. `SNOwGLoBES` uses four neutrino "flavours" (the Garching parameterization), ν_e, $\bar{\nu_e}$, ν_x, and $\bar{\nu_x}$ (where ν_x represents the ν_μ and ν_τ flavours, and $\bar{\nu_x}$ represents the $\bar{\nu_\mu}$ and $\bar{\nu_\tau}$ flavours). The flux's time-dependence arises from the Fermi-Dirac pinching parameter and the mean neutrino energy. Since `SNOwGLoBES` was originally intended to work with fluences as input rather than fluxes, the reading in of fluxes requires its separation into discrete time bins. The Garching model is separated into 300 time bins by default, each bin has an associated time integrated flux, or fluence, associated with it. The format for the flux files (one for each time bin, in this case 300) is 7 columns by 501 rows. The 7 columns represent the discrete neutrino energies and the flux for each of the six neutrino flavours, which have been integrated across the time bin. It is worth noting here, that four of the columns share the same flux value since under the Garching parameterization, the non-ν_e or $\bar{\nu_e}$ flavours have the same flux value. The 501 rows are for neutrino energies ranging from 0.1 to 100 MeV in 0.2 MeV increments. The standard `SNOwGLoBES` protocols are run for each of the 300 time bin "fluxes", producing output files with total number of events per channel.

For the additional set of Sukhbold et al. models outlined previously (§3), there was a conscious choice to increase the number of time bins used when parsing the flux. Each of the four models had their integrated time bins increased to 500, up from the original 300 used by default within `SNOwGLoBES`. This was done as these models are defined over longer time periods, where as the Garching model is cut-off at ≈ 9 seconds after t_0. The increased number of time bins also allows time bins near the neutronization peak to be sufficiently narrow and thus preserving the shape of the flux. The time bins are not uniform in size but increase the further along in the burst. The technique used for the Garching model in `SNOwGLoBES` was implemented by the developers as

$$t_i = t_s \cdot 10^{i t_w} \tag{6.2}$$

$$dt_i = t_s \times \left[10.0^{(i+0.5)\cdot t_w} - 10.0^{(i-0.5)\cdot t_w} \right] \tag{6.3}$$

where t_i is the starting time of the i-th bin, dt_i is the bin width, t_s is the model starting time, and t_w is the step size defined as

$$t_w = \frac{\log_{10}\left(\frac{t_e}{t_s}\right)}{N+1} \tag{6.4}$$

where t_e is the end time of the model and N is number of bins. The results are logarithmically increasing bin widths over the entire SN burst.

This work utilizes the "halo" and "halo2" configurations built into `SNOwGLoBES`, which are simplified 79 tonne and 1 kilo-tonne lead targets respectively. CC and NC neutrino-lead interactions that yield 1 neutron (1n) or 2 neutron (2n) events are both utilized, while inelastic neutrino-electron scattering is ignored. The interactions in HALO and HALO-1kT are listed in Table 6.1.

Interaction	Channel
$^{\mathrm{A}}\mathrm{Pb}(\nu_e,\, n + e^-)^{\mathrm{A\text{-}1}}\mathrm{Bi}$	CC
$^{\mathrm{A}}\mathrm{Pb}(\nu_e,\, 2n + e^-)^{\mathrm{A\text{-}2}}\mathrm{Bi}$	CC
$^{\mathrm{A}}\mathrm{Pb}(\nu_x,\, n)^{\mathrm{A\text{-}1}}\mathrm{Bi}$	NC
$^{\mathrm{A}}\mathrm{Pb}(\nu_x,\, 2n)^{\mathrm{A\text{-}2}}\mathrm{Bi}$	NC
$^{\mathrm{A}}\mathrm{Pb}(\bar{\nu}_x,\, n)^{\mathrm{A\text{-}1}}\mathrm{Bi}$	NC
$^{\mathrm{A}}\mathrm{Pb}(\bar{\nu}_x,\, 2n)^{\mathrm{A\text{-}2}}\mathrm{Bi}$	NC

Table 6.1: Interaction modes present in `SNOwGLoBES` for HALO and HALO-1kT. In the CC interactions, an electron is ejected for both 1n or 2n reactions, but go unobserved in HALO and HALO-1kT.

Upon running `SNOwGLoBES` for the 300 time bin files, there is an output file per interaction channel and multiplicity pair (i.e. `pinched_0_nc_nue_Pb208_1n_halo2_events.dat`). Each output file is accompanied by a smeared output based on "smearing matrices" that incorporate the detector response and interaction product spectra. For this analysis, we opt to use the raw weighted files as we account for detector response in our use of the detector's `Geant4` Monte Carlo simulations and the interaction product spectra via theoretical models (both of these are discussed in §6.2). The format of the output is a 2 column by 200 row text file, where the first column is the neutrino energy (from 1 to 100 MeV in 0.5 MeV increments) and the second is the event rate integrated across the time bin. The theoretical model we utilize for the kinetic energy distribution of ν-induced neutrons does not provide discrete distributions for each of the neutrino energies listed in the output file (more on this in §6.2). Instead, it is parameterized by a mean neutrino energy and Fermi-Dirac pinching parameter. In lieu of this, we sum the expected events per incident neutrino energy in each output file. To account for statistical deviations in the number of ν-induced neutrons, we Poisson fluctuate the sum events per file, where the mean (λ) is equal to the sum. We opt to fluctuate each files sum events rather the the total events expected over the entire burst to preserve the temporal information provided by `SNOwGLoBES`. This discussion will be picked up in §6.2, where the output from `SNOwGLoBES` is loaded into HALO and HALO-1kT.

6.1.2 Supernova Neutrino Early Warning Models for Python (`SNEWPY`)

`SNEWPY` made its first release to the SNEWS collaboration in 2021, opening up the possibility for additional models to be easily incorporated into our pipeline. `SNEWPY` does not calculate neutrino events and is rather a wrapper package for `SNOwGLoBES` in its current configuration. With permission from the authors, `SNEWPY` has integrated dozens of SN models that can be used to perform studies such as this. However, many of these models are tabulated in a format the respective groups deem desirable. As such, `SNEWPY` exists to facilitate the conversion from one format to another. For specific details on how the conversion process is handled on a case-by-case basis, the reader is directed to the User Manual on the git page [88].

6.1.3 Supernova tools (`sntools`)

`sntools` is a Monte Carlo event generator for supernova neutrinos in current and next-generation neutrino detectors. It was developed by Jost Migenda for their PhD thesis and recently made part of the SNEWS 2.0 software package. It can be found at (https://github.com/JostMigenda/sntools). Analogous to `SNOwGLoBES`, `sntools` breaks the neutrino flux into time bins to evaluate the expected number of events throughout the burst. `sntools` by default uses 1 ms bins for every point in the burst. For any time bin, the number of expected events is given by

$$N(t) = \int\int \frac{\mathrm{d}\phi(t, E_\nu)}{\mathrm{d}E_\nu} \frac{\mathrm{d}\sigma(E_\nu, E_e)}{\mathrm{d}E_e} \mathrm{d}E_e \mathrm{d}E_\nu \tag{6.5}$$

where $\phi(t, E_\nu)$ is the neutrino flux and $\sigma(E_\nu, E_e)$ is the detector material cross section. `SNEWPY` is fully integrated into `sntools`, resulting in the full library of models being accessible for neutrino event calculation. For a full treatment on how each model's format is integrated into `sntools`, please consult [87]. Event rate calculations assume four neutrino species, ν_e, $\bar{\nu_e}$, ν_x, and $\bar{\nu_x}$. As was the case in `SNOwGLoBES`, ν_x represents ν_μ and ν_τ, while

$\bar{\nu_x}$ represents $\bar{\nu_\mu}$ and $\bar{\nu_\tau}$. For each time bin, `sntools` provides a Poisson fluctuation of the expected number of events, using $N(t)$ as the mean value of the distribution. As shall be discussed in §6.2, this was integrated into `SNOwGLoBES` for the scope of this work.

The interaction modes available for LAB in `sntools` are listed in Table 6.2. These are direct implementations of the interactions outlined in §5.2.3. The bulk of the work required to get SNO+ implemented into `sntools` was implementing low energy proton scattering, as well as the correct geometry for SNO+. All other interactions listed in Table 6.2 were already present for other liquid scintillator detectors previously implemented within `sntools`.

Interaction	Channel
^{12}C(ν_e, e^-)^{12}N*	CC
^{12}C($\bar{\nu_e}$, e^+)^{12}B*	CC
^{12}C(ν,ν')^{12}C*	NC
$\bar{\nu_e}$ $(p,n)e^-$	CC
$\nu(p,p)\nu$	NC
$\nu(e^-,e^-)\nu$	CC/NC

Table 6.2: Supernova neutrino interactions in LAB implemented in `sntools`.

6.2 Input into the Monte Carlo simulations

6.2.1 HALO and HALO-1kT

With the output from `SNOwGLoBES`, it almost stands ready to be simulated through the HALO and HALO-1kT Monte Carlo simulations. Since `SNOwGLoBES` only provides temporal and multiplicity information (1n or 2n), additional assumptions about positioning and initial momentum vectors are required to simulate the neutrons in `Geant4`. After these assumptions are made, the last information required to create our input vector for event simulation in HALO-1kT and HALO is the initial neutron energy. The neutrino interactions in the lead are assumed to be uniformly distributed throughout the volume of lead. This assumption is justified as the ν-Pb cross section is sufficiently small such that neutrino wave front does not experience any meaningful attenuation as it propagates through the lead. `Geant4` has integrated tools in place to handle event vertex simulation via volume confinement. The tools in place use an accept-reject method, where a point is chosen uniformly within a box with the same half-lengths in (x,y,z) as the lead physical volume. Initial position vectors $(\vec{r})$ are randomly selected by scaling three random numbers, one for each component of the position vector, by the lead volumes dimensions along each axis. If $\vec{r}$ is within the lead volume, the particle is simulated, otherwise a new position is sampled until this criterion is satisfied. 2n interactions make use of the same initial vertex position. For both HALO and HALO-1kT, 2D scatter plots for the initial event locations are shown in Figures 6.1 and 6.2 respectively (the former is for HALO and only contains the XY plane). The distribution of initial positions along the XY plane is as expected, discrete gaps are present where the annuli are located. The same behaviour in regards to the annuli is observed along the YZ plane;‘ since the proportional counters run length-wise along the z-axis the gaps are instead replaced by low initial event population lines.

For both 1n and 2n interactions, the momentum cosines (u_x, u_y, u_z) are assumed to be isotropic and are determined via sampling

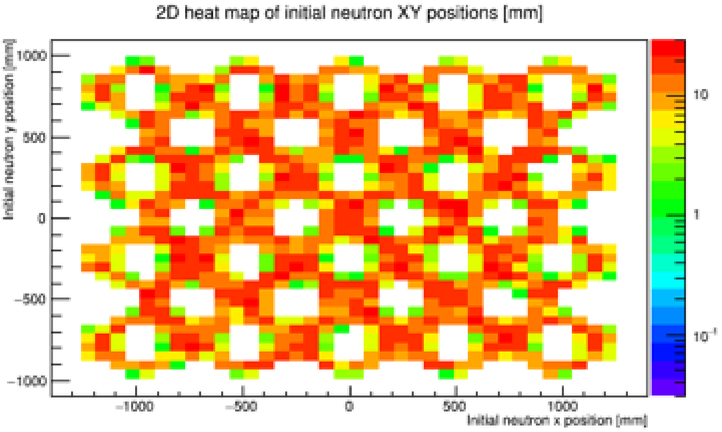

Figure 6.1: 2D scatter plot of the initial vertex positions in HALO for 10^4 neutrons simulated in the lead.

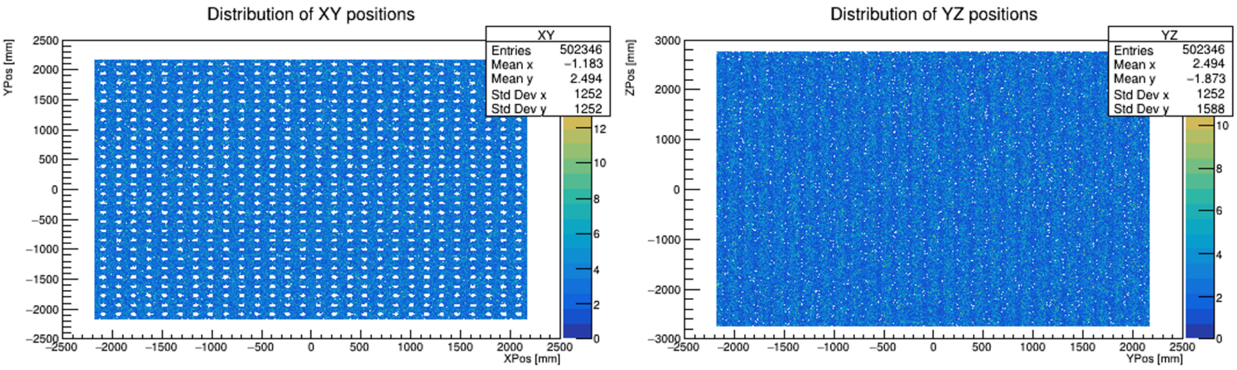

Figure 6.2: 2D scatter plots of the initial vertex positions in HALO-1kT for 10^4 SN signals at 10 kpc (Garching model).

$$\theta = 2\pi \cdot R_\theta \tag{6.6}$$

$$\phi = \arccos\left(1 - 2R_\phi\right) \tag{6.7}$$

where R_θ and R_ϕ are random numbers between [0,1]. In Figure 6.3 the random polar (θ) and azimuthal (ϕ) angle are plotted for the SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT.

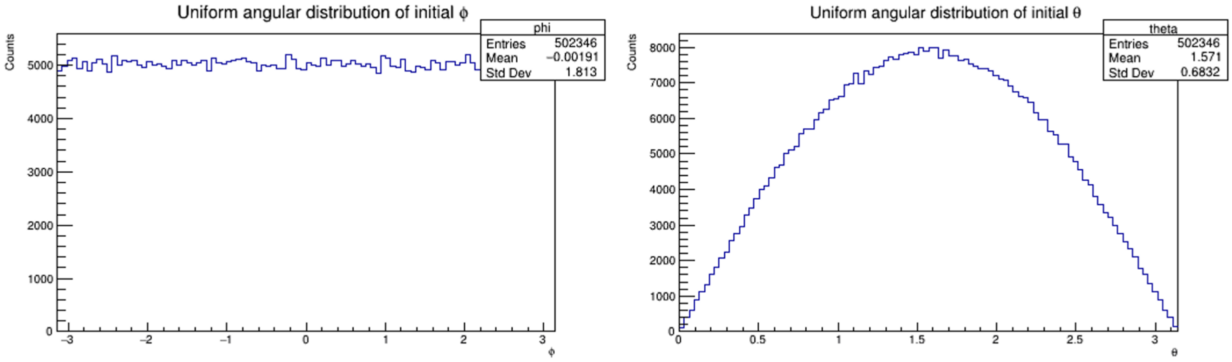

Figure 6.3: Uniform azimuthal (left) and polar (right) angular distributions from generated neutrons in HALO-1kT. This particular set is 10^4 SN signals at 10 kpc (Garching model).

As of this work, there has been no direct measure of the neutron energy spectrum from neutrino-lead interactions, nor a measurement of neutrino-lead cross sections at supernova relevant energies. The latter is the primary physics goal of the miniHALO experiment that has plans to be located at Oak Ridge and make use of the Spallation Neutron Source (SNS), which provides an intense pulsed neutrino source from stopped π^+ decays in the neutron production target. Several theoretical calculations have been done for the neutrino-lead cross sections [91–95]. Even with the theoretical calculations for the neutrino-lead cross section, there persists gaps in our knowledge of the neutrino-lead cross section for different isotopes of lead, multiplicity, and neutrino species. `SNOwGLoBES` elected to use [91] in its calculations for neutrino events in lead-based supernova detectors. Additionally, we make use of Kolbe et al. [95] for the out-going neutron energy spectra for varying SN fluxes. These fluxes are parameterized by the Fermi-Dirac spectrum

$$n(E_\nu) = \frac{1}{F_2(\alpha)T^3} \frac{E_\nu^2}{\exp[(E_\nu/T) - \alpha] + 1} \tag{6.8}$$

where T is the neutrino temperature, α is the pinching parameter, E_ν is the neutrino energy and F_2 is a normalization factor. The cross section for varying out-going neutron energies from a typical supernova neutrino burst (making use of the spectrum in Equation 6.8) can be found in Figure 6.4 for both the CC and NC interactions. These two figures have been normalized to facilitate a shape comparison, each normalized such that each model's area is equal to one. Of course, there is a strong positive correlation between increasing neutrino temperature and increasing out-going neutron temperature.

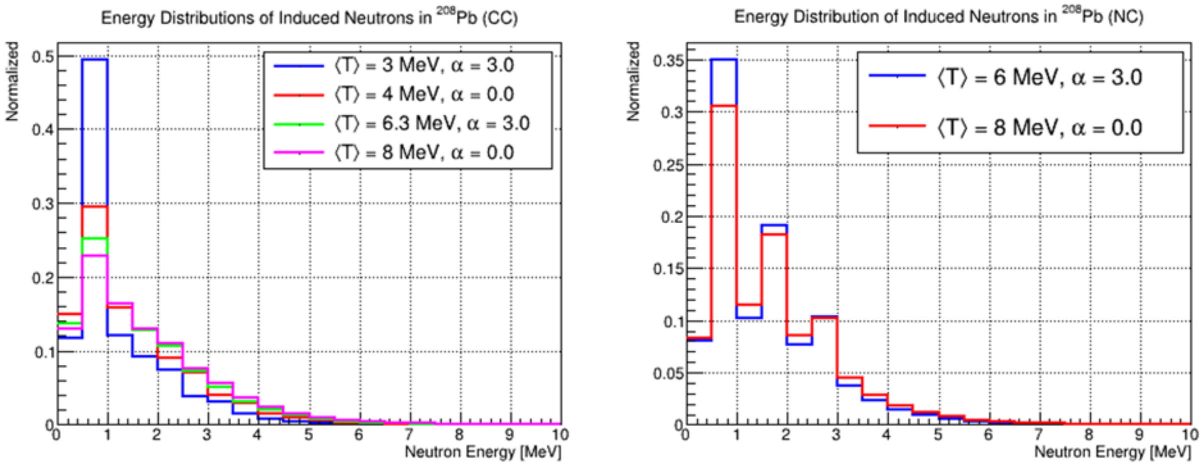

Figure 6.4: Neutron energy distribution (courtesy of [95]) for both the CC and NC interactions in lead (for varying pinching parameters and neutrino temperature configurations).

For the purposes of this report, the NC and CC distributions with the highest average out-going neutron energy were chosen. This correlates to the energy distributions with

$\langle T \rangle = 8\,\mathrm{MeV}$ and $\alpha = 0.0$. This could affect the capture efficiency of neutrons in the lead volume. Provided $|\vec{r}|$ is close to the half-length of the width, length, or height of the detector, higher energy neutrons that need to undergo additional interactions in the lead before thermalizing, can have a lower capture efficiency as a result of their energy and proximity to the edge. To ensure that unwanted errors were not introduced into our analysis (i.e. a reduction in the signal yield due to the higher neutron energy), we simulated 10^6 neutrons in HALO-1kT for each distribution found in Figure 6.4, randomly assigning momentum vectors and positions in the manner described above. The neutron capture efficiency is determined simply by

$$\epsilon = \frac{n_c}{N_s} \tag{6.9}$$

where n_c is the number of captured neutrons and N_s is the total number simulated. The findings of these simulations can be found in Table 6.3. The largest variation in capture efficiency between two spectra is $\sim 0.26\%$. HALO-1kT is not sensitive to the variation in shape between these energy distributions for a typical SN burst, especially when accounting for the Poisson fluctuations implemented earlier.

Channel	α	$\langle T \rangle$ [MeV]	ϵ
CC	3	3	0.4939
CC	0	4	0.4928
CC	3	6.3	0.4922
CC	0	8	0.4913
NC	3	3	0.4910
NC	0	8	0.4909

Table 6.3: Neutron capture efficiencies for the six energy distributions present in Figure 6.4. Poisson errors are not depicted as they are of order 10^{-4} and do not affect the findings.

The additional collisions required to thermalize the neutron prior to capture will increase the capture times. The capture times can be found in Figure 6.5. On average, the increase in capture time is $\approx 10\,\mu$s. Such a slight increase is negligible and will have no impact on the analysis in later chapters.

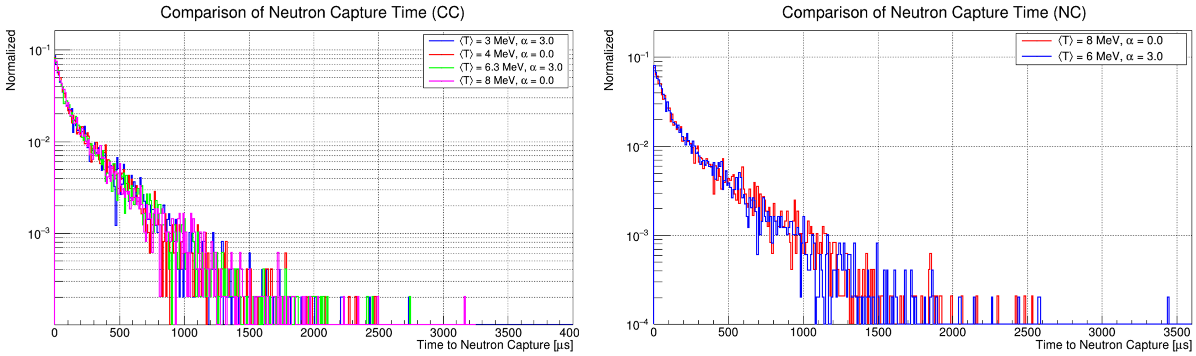

Figure 6.5: Neutron capture time in HALO-1kT for the CC (left) and NC (right) neutron energy distributions.

6.2.2 SNO+

The output format of `sntools` can be either NUANCE or RATPAC. Since the SNO+ Monte Carlo simulation is configured to use RATPAC format files, the latter option was chosen. For each interaction implemented in `sntools`, there are different methods applied to obtain the parameters (i.e. daughter momentum vector, position vector, kinetic energy, etc.) necessary to pipe it into the simulation. This work was already built into `sntools` and for the purposes of this report, no changes or additional implementations were neccesarry to carryout the analysis in the following chapters. If the reader is interested as to how this was obtained for each interaction channel, they may consult the `sntools` documentation [87].

6.3 Output from the Monte Carlo simulations

To perform the extraction of t_0 from each detector's Monte Carlo simulation, a little information is required. For HALO and HALO-1kT, only the event capture times are necessary since our analysis later on requires only the time series of events as simulated in the detector. These are the Monte Carlo truth times as recorded by `Geant4` upon the creation of the triton-proton pair. Any electronic readout effects are ignored for this analysis. On time scales of the SN burst (seconds) and each neutron's thermalization and capture time ($\approx 200\ \mu s$), the electronic readout will happen on time scales of 10's of μs.

6.3.1 Addition of Poisson background

6.3.1.1 HALO and HALO-1kT

In keeping with accounting for a realistic detector response, a Poisson background was added into the each detector's simulated signal. For HALO, the background rate within the 764 keV peak and tail region is 15 mHz [96], summed across all background channels. The background events present in region of interest (ROI, the region that corresponds to the 764 keV peak and the tail) come from spontaneous fission, internal alphas, fast neutrons from the surrounding rock, cosmic muons, etc. An internal study by Andrea Gallo Rosso [97] was aimed at identifying and quantifying the external backgrounds for HALO-1kT at LNGS. Currently, the design goal for HALO-1kT is to have a background rate of ≤ 1 Hz.

Throughout the burst, background events were added to the HALO simulation data by incrementing through the burst in 1 second intervals and sampling a Poisson distribution with $\lambda = 0.015$. For each of the number of events returned (n_p), we select a time stamp randomly between $[t_j,t_{j+1})$, where $j = 1, 2, 3...10$ and is used as the time window identifier. The same procedure was carried out for HALO-1kT, except with a background rate of 1 Hz. These background events added to HALO and HALO-1kT are not simulated, rather time stamps are added into the Monte Carlo output. Any potential proportional counter wall effects outlined in §5.3 are ignored. The Poisson probability distributions are shown in Figure 6.6 for HALO-1kT. These background events will have little impact on the analysis. This discussion will be picked up in §7.1, where the impact of the background on picking out the first event in the SN burst will be explored.

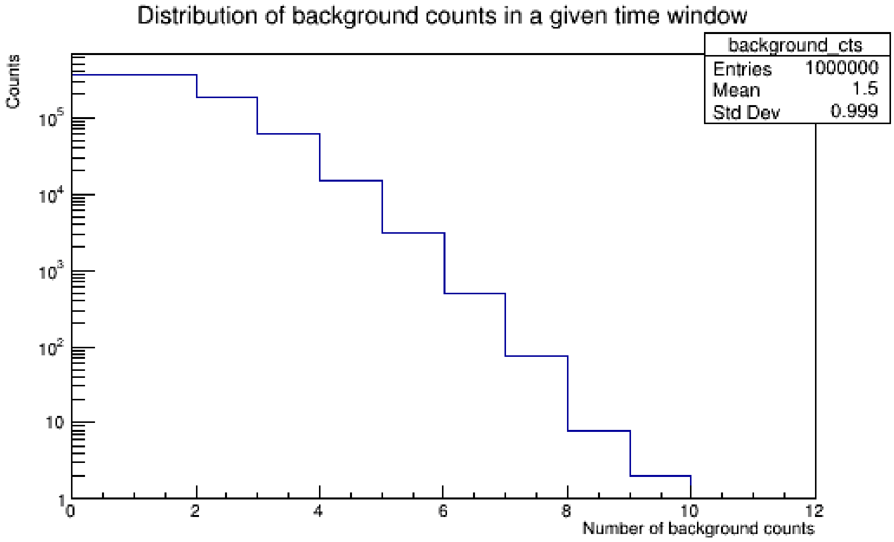

Figure 6.6: Background counts in a one second window within HALO-1kT for 10^6 sampled time windows (if normalized, this would be equivalent to the Poisson p.d.f).

6.3.1.2 SNO+

For SNO+, there is no pre-determined background rate expected during the neutrino signal [98]. The simplified model applied to HALO and HALO-1kT cannot be used here.

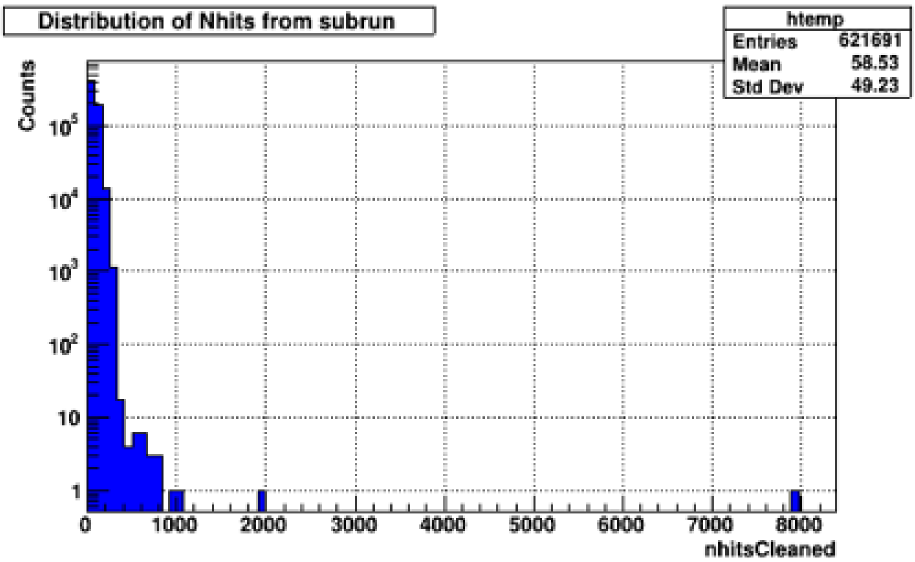

Figure 6.7: Cleaned N_{hit} spectrum for a subrun fetched from Nearline.

However, since SNO+ is currently collecting data in the scintillator phase, we can make use of the data stored on Nearline (the computing cluster located at SNOLAB) taken during a physics run. Of interest is the cleaned N_{hit} and time spectra. To add in the background to each SN burst simulated in the SNO+ Monte Carlo, the following procedure was applied. First we select a random data file from the list of data taking runs. We then fetch the cleaned N_{hit} spectra and time series of events. Next we randomly select a first event from the UTC time branch in the `ROOT` file (the cleaned N_{hit} spectrum is the N_{hit} spectrum that has passed through filters to remove spurious bursts that arise from electronic noise). We then increment forward through the time series and fetch all event times that fall within our model's simulation duration. Each of the events we sample will have a corresponding cleaned N_{hit}. Each cleaned N_{hit} is Poisson fluctuated, where the mean of each Poisson PDF used in the fluctuation corresponds to the cleaned N_{hit} retrieved from the file. This background event is then added into the Monte Carlo output. Figure 6.7 shows an example of the cleaned N_{hit} spectra as observed in SNO+. Figure 6.8 shows the corresponding GPS clock time for the events.

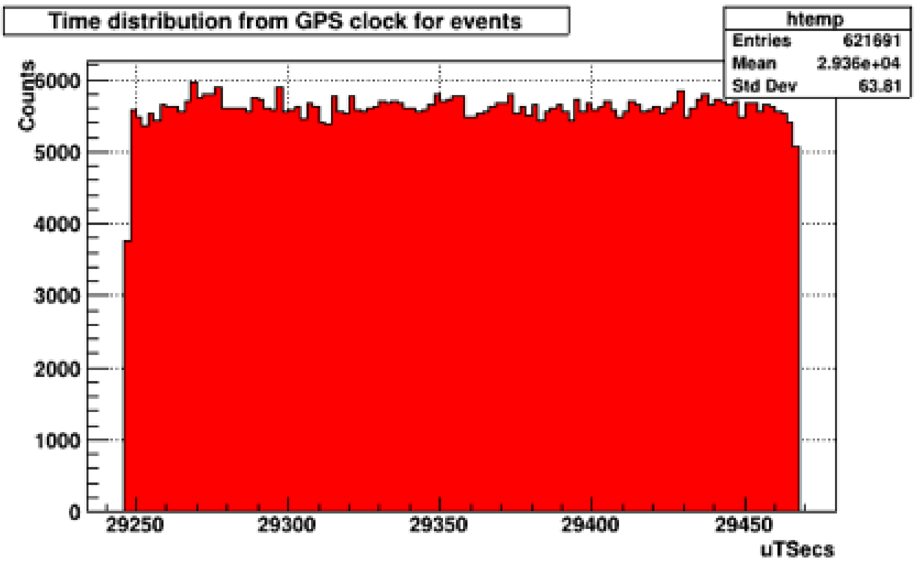

Figure 6.8: GPS clock time for events in a subrun fetched from Nearline.

CHAPTER 7

Extracting the time of core-bounce from SN neutrino signals

With the Monte Carlo simulation output, we now look to perform the mission statement of this work, extracting t_0, the time of core-bounce. This chapter will not cover all techniques used to extract t_0 from a SN burst explored throughout the duration of this thesis; instead, we shall focus on five of the most promising techniques. §7.1 will cover the techniques used to determine the first event of the SN burst in each of the three detectors; §7.2 to 7.6 will go through each technique individually, defining their general formulations, potential corrections to their definitions, and how to extract t_0. For an in-depth comparison between the techniques, the reader is directed to §8. Contained in this chapter is only a brief discussion on the relative performance of one technique against another.

For this chapter, the results shall be presented with the following assumptions:

- HALO-1kT will be the primary focus when presenting the results for extracted t_0. The results for HALO and SNO+ can be found in §9, where a more rigorous comparison between the three experiments will take place. The intent is to present the results of this thesis in the order they were developed. HALO-1kT was the first experiment to have its extraction capabilities quantified in this thesis; the scope of work was expanded to HALO; then SNO+.

- No flavour transformations are applied to the neutrino flux between emission and observation. This includes oscillations in matter or vacuum. For each model described in §3, the neutrino spectrum from the SN simulations was used as presented.

- Results will be shown for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc unless otherwise stated. The motivation to present this model ahead of others is attributed to its excellent integration into the `SNEWPY` pipeline. Although the Hüdepohl et al. 8.8 $M_\odot$ model was the primary focus for the first stages of this thesis, its integration into `sntools` has yet to be completed, where all neutrino event calculations are done for SNO+ in this thesis. Therefore, for a first pass comparison between detectors, the Sukhbold et al. SFHo z9.6 $M_\odot$ model makes an ideal candidate.

Some useful terms will be defined here to ensure clarity in this chapter and beyond:

- For the likelihood, Anderson-Darling, Kolmogorov-Smirnov, and cross-correlation techniques we use the term *model* to describe the time series of data that has been simulated through the detector Monte Carlo simulation. When referencing the probability distribution function (PDF), this is either the mean cumulative event count or the mean event rate that is formed from a large number of the simulated bursts for a specific model and distance. First discussed in §7.2.6, expressions such as "the Sukhbold et al. SFHo z9.6 model was fit to the Sukhbold et al. LS220 s27.0 PDF" refers to fitting Sukhbold et al. SFHo z9.6 time series data to either the Sukhbold et al. LS220 s27.0 mean cumulative count or mean event rate (depending on the technique). These techniques are employed to estimate the systematic uncertainties in the extraction of $\hat{t_0}$ when fitting an observed burst to a different PDF. More information can be found in the aforementioned subchapter.

- *Light curve* refers to the cumulative event count formed from the time series of data that is observed in an experiment. For the i-th bin, the lower bin edge corresponds to the event time and the width is the time difference between the i and $i+1$ event. When we form the mean cumulative event count, this is the mean light curve.
- t_0 is the Monte Carlo truth time of core-bounce for each of our SN models as defined in §3.
- $\hat{t_0}$ is the estimator of t_0, and is extracted from the various techniques outlined in this work. For each of the techniques explored in this chapter there is the potential for the $\hat{t_0}$ to be systematically offset from the Monte Carlo truth value of $t_0 = 0$ seconds. For further details the reader is encouraged to consult the discussion in the subchapters that follow and §8.

7.1 Determination of the first event

With the addition of the Poisson background described in §6.3.1, careful attention must be paid to ensure that background events in the leading edge of the SN neutrino signal are ignored in the the algorithms used to extract t_0. The following two subchapters will briefly touch on the methods used in this thesis to determine the appropriate start point in the SN burst for this analysis, in each of the respective detectors.

7.1.1 HALO and HALO-1kT

For HALO, the background rate within the 764 keV peak (and down to the end of the 191 keV tail) is quoted at 15 mHz. Over the duration of a SN burst ($\approx$ 10s of seconds), this exceedingly low background rate will prevent any background channel from significantly contributing to the observed time series. For this analysis we utilize the standard SN trigger in HALO, with no modifications. This current trigger configuration is 4 events in a 2 second window. This window is independent of the time window that shall be applied for individual fits and is solely intended to identify the start of the burst.

The trigger condition for HALO-1kT requires a more nuanced approach. The design goal background rate is 1 Hz for HALO-1kT, two orders of magnitude larger than in HALO. The primary contribution to this rise in background rates comes from a shift in the material used in the neutron proportional counters within in the detector and external backgrounds at the proposed location. This 1 Hz Poisson background has a significant contribution on the observed time series at distances around 10 kpc, where the expected number of neutrons observed in HALO-1kT for the Sukhbold et al. SFHo z9.6 model is on the order of ≈ 43. The Poisson background will be 8-28% of the total signal at 10 kpc, depending on the models tested in this thesis. The two order of magnitude increase in the background rate yields a trigger configuration similar to HALO's unusable. With the benefit of foresight, we are aware that the ability to extract t_0 at 5 kpc, with sufficient precision and accuracy to contribute to the SNEWS 2.0 triangulation program, is modest at best. The approach to HALO-1kT's SNEWS trigger is two-fold: we wish to keep the background events in the time window sufficiently low to ensure that SNe at 5 kpc can be picked up, but also keep the rolling time window's length small enough that multiple events from the 1 Hz background do not present themselves. A time window of 50 ms is sufficient to cut out the contribution from $>$ 1 background events. The distribution of total events in the first 50 ms after the first event can be found in Figure 7.1. From this we see that at 5 kpc, the average number of events from the Monte Carlo simulation is 15 within a 50 ms window (across all models simulated in HALO-1kT). For an aggressive threshold we place the minimum number of events at 5. This minimum is sufficient to trigger off all models simulated in HALO-1kT. As for the Hüdepohl et al. 8.8 $M_\odot$ model, it is the only model to reach low enough statistics at 5 kpc that for ≤ 1 % of the simulated bursts, our method will fail to pick out the first SN related event.

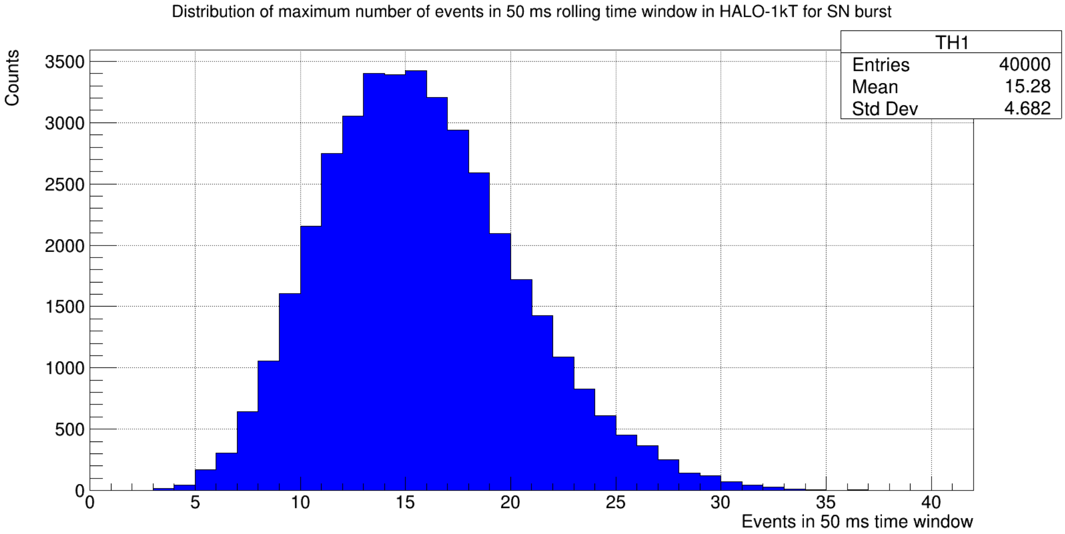

Figure 7.1: Distribution of events in a 50 ms time window for HALO-1kT. This is for all four models simulated in the detector Monte Carlo.

7.1.2 SNO+

The sparse events in the early part of the SN signals simulated in SNO+ pose difficulties in all fitting techniques that shall be explored in this chapter. It is in the best interest of this work to reduce the prevalence of these events and to introduce an aggressive trigger condition that helps remove the leading edge. As such, we introduce the following criteria to identify the start of the burst; a total of 5 events within the LAB volume; each with an N_{hit} higher than 500; all occurring in a rolling time window of 20 ms. This effectively truncates the events that occur prior to core-bounce, but improves the performance in extracting t_0, which for SNO+, shall be discussed in §9. The N_{hit} spectrum from backgrounds in the detector for a given subrun was shown in Figure 6.7; by introducing an aggressive trigger condition outlined previously, the backgrounds are effectively removed from our time series.

7.2 Negative Log Likelihood (NLL)

7.2.1 Motivation

HALO-1kT is unique amongst the next-generation of supernova neutrino detectors, in large part due to its flavour sensitivity and relatively low statistics compared to other neutrino detectors. Owing to HALO-1kT's low statistics, the aim is to avoid methods of extraction that require binning the neutrino signal. At 10 kpc, the expected events from the Hüdepohl et al. 8.8 $M_\odot$ model in HALO-1kT is ≈ 25 events, excluding the contribution from the Poisson background. If an attempt were made to extract t_0 at this distance, bin widths on the order of milliseconds would be too sparsely populated and the alternative of increasing bin size would sacrifice the temporal resolution needed to effectively triangulate. The variance and asymmetry of the first event times also rule out using techniques implemented by [6, 52] (as discussed in §4.4.2). As a result, we turn our attention to the unbinned maximum likelihood technique or Negative Log Likelihood

(NLL).

7.2.2 General formulation

In our formulation of the likelihood function, the only parameter of interest that shall be estimated is t_0. As such, the formal definition used here is as follows [99]

$$\mathcal{L}(t_0) = \prod_{i}^{N} P[n(t_i); \bar{n}(t_i - t_0)] \tag{7.1}$$

where N is the number of events in the burst, P is the probability density function (PDF), $n(t_i)$ is the number of events observed since the start of the SN burst (each event has a time stamp t_i) and $\bar{n}(t_i - t_0)$ is the mean cumulative event count evaluated at the SN burst times, with an offset relative to t_0. The choice of PDF should describe the variations in burst size we encounter with the introduction of the Poisson fluctuations. A time dependant Poisson PDF for our likelihood function more accurately reflects the distribution of events in the early burst times versus a Gaussian PDF. Our Poisson PDF can simply be written as

$$P[n(t_i); \bar{n}(t_i - t_0)] = \frac{\bar{n}(t_i - t_0)^{n(t_i)} \cdot e^{-\bar{n}(t_i - t_0)}}{n(t_i)!} \tag{7.2}$$

To ease the computation of the likelihood, we can transform it by taking the negative natural log of both sides to obtain

$$\ell(t_0) = 2 \cdot \sum_{i=1}^{N} \bar{n}(t_i - t_0) + \ln\big(n(t_i)!\big) - n(t_i) \ln(\bar{n}(t_i - t_0)) \tag{7.3}$$

In our calculations, the second term is removed as it is constant at each offset time we evaluate the function. Additionally, for bursts that occur at sufficiently close distances, this term rapidly explodes to infinity, resulting in a failed fit. More impactful changes to this formalization are covered in §7.2.4. How to extract t_0 from the likelihood function, with the outlined modifications, will be discussed in §7.2.5.

7.2.3 Forming the mean cumulative event count

To form the mean cumulative event count, we implement an average of the set of simulated SN signals (for a given model and distance). The motivation behind using the simulated signals as opposed to the calculated values from the event rate calculators, is the inclusion of realistic detector effects (i.e propagation times and capture efficiencies). In the case of HALO and HALO-1kT, this has a temporal effect on time series of the order a couple hundred microseconds.

Two methods of forming $\bar{n}(t)$ were explored, only one of which proved effective. Each burst in the detector can be represented as a time series

$$n_k(t) = \{t_1, t_2, t_3 ... t_i ... t_N\} \tag{7.4}$$

where t_i is the event time of the i-th event; N is the total number of events in the burst that pass the cuts applied by each detector; and $1 \leq k \leq N_k$ (where N_k is the set of simulated bursts for a given model and distance). For HALO-1kT and HALO, a set of $N_k = 10^4$ bursts were simulated for each of the four models at 5 distances. For SNO+, we reduced the number of bursts simulated within a set to 500. The individual bursts can be concatenated and sorted (with respect to their event times) into a singular data series of the form

$$\bar{n}(t) = \frac{1}{C_{\bar{n}}} \cdot \text{sorted}\{n_1(t), n_2(t), n_3(t)...n_{N_K}(t)\} \tag{7.5}$$

A normalization coefficient ($C_{\bar{n}}$) is introduced into the sorted time series. This coefficient is determined by averaging the total number of events for each burst. The sorted time series will serve as the mean cumulative event curve for each detector, where an example in HALO-1kT in shown in Figure 7.2.

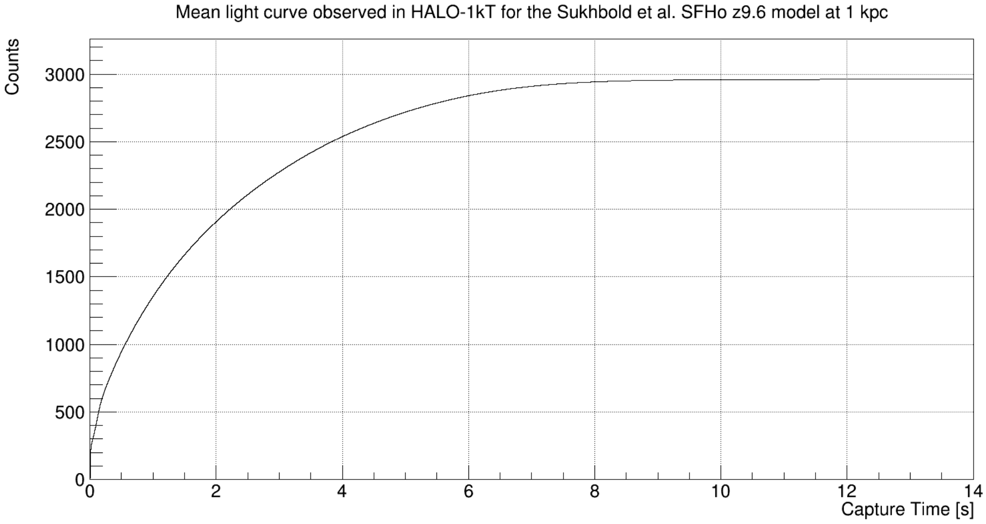

Figure 7.2: Mean light curve for Sukhbold et al. model observed in HALO-1kT.

7.2.4 Modifications

In its current configuration, the likelihood fit is inconsistent. Minor modifications are required to ensure that the fit does not reach out of bound regions and can deal with the Poisson fluctuations we implemented into the simulation data. These are discussed below.

7.2.4.1 Normalization of mean cumulative count

We defined the NLL as

$$\ell(t_0) = 2 \cdot \sum_{i=1}^{N} \bar{n}(t_i - t_0) - n(t_i)\ln(\bar{n}(t_i - t_0)) \tag{7.6}$$

With the NLL technique, there is no consideration that the size of the SN burst varies for each of our simulated bursts in the set. This is a direct consequence of the Poisson fluctuations introduced previously. When mixing models and PDFs (Section 7.2.6), the effects of varying burst size is more dramatic. To account for this variation we introduce a normalization coefficient into the NLL technique as

$$\ell(t_0) = 2 \cdot \sum_{i=1}^{N} \eta \cdot \bar{n}(t_i - t_0) - n(t_i)\ln(\eta \cdot \bar{n}(t_i - t_0)) \tag{7.7}$$

where η scales the mean cumulative event count, $\bar{n}(t_i - t_0)$, down to the number of events in the simulated time series under consideration. The normalization will be calculated on an individual basis and is defined as

$$\eta = \frac{n(t_N)}{\bar{n}(t_N)} \tag{7.8}$$

where $n(t_N)$ is the event number of the final event in the burst and $\bar{n}(t_N)$ is the final value of the mean cumulative event count. The result is fixing the mean cumulative neutron count to the same stationary end point as the burst under consideration. The resolution of our extracted t_0 is greatly improved (Figure 7.3), with the standard deviation reducing greatly when the normalization is turned on.

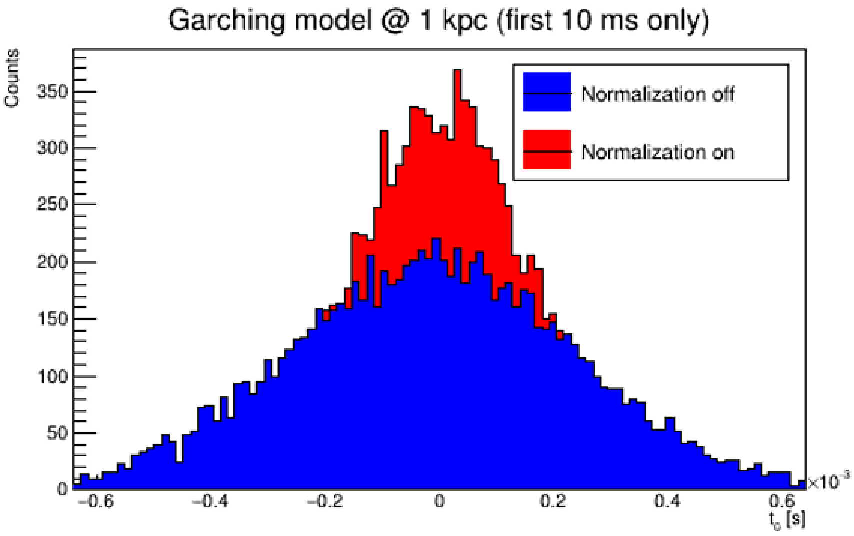

Figure 7.3: Impact of a normalization coefficient on the NLL fit in HALO-1kT for the Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc, for the first 10 ms.

7.2.4.2 Extension of the leading edge

Many of the modifications that are included in these techniques are reactionary, rather than preemptive. Although not used when comparing the experiments to one another in §9, the Hüdepohl et al. 8.8 $M_\odot$ model was used in the analysis for HALO-1kT. The Hüdepohl et al. 8.8 $M_\odot$ model for a CCSN posed many issues when attempting to extract t_0 using the NLL technique at far distances (> 5 kpc). At these distances, a "cool" SN model such as this will have fewer events and we expect its precision to extract t_0 to be greatly reduced. The likelihood function must then be defined (no discontinuities, singularities, etc.) within a few milliseconds around the true t_0 value, otherwise the function can terminate prior to reaching its minimum. In the definition of our likelihood, the function shall be undefined should the following be satisfied

$$\bar{n}(t_i - t_0) \leq 0 \tag{7.9}$$

Of the five models simulated for HALO-1kT, the Hüdepohl et al. 8.8 $M_\odot$ model was unique as it was the only model that had a significant observed neutrino signal prior to t_0. These early events have Monte Carlo truth time indices that extend out to -15 milliseconds. If the likelihood function does not reach its minimum within a few milliseconds of t_0, Equation 7.9 can be satisfied and the function will terminate. To alleviate this, we extend the value of the mean light curve 50 ms prior to core-bounce, treating it as constant. An example of this is given in Figure 7.4. Other techniques were considered, such as discarding the event and reevaluating the likelihood. However, discarding the leading events can become recursive and in doing so, systematic uncertainties will be introduced as a consequence. An alternative would be to discard the event and continue the likelihood evaluation, opposed to restarting it; this too possesses problems as the function can reach multiple minima and systematically offset $\hat{t_0}$. The sensible solution is to simply extend the mean light curve's leading edge.

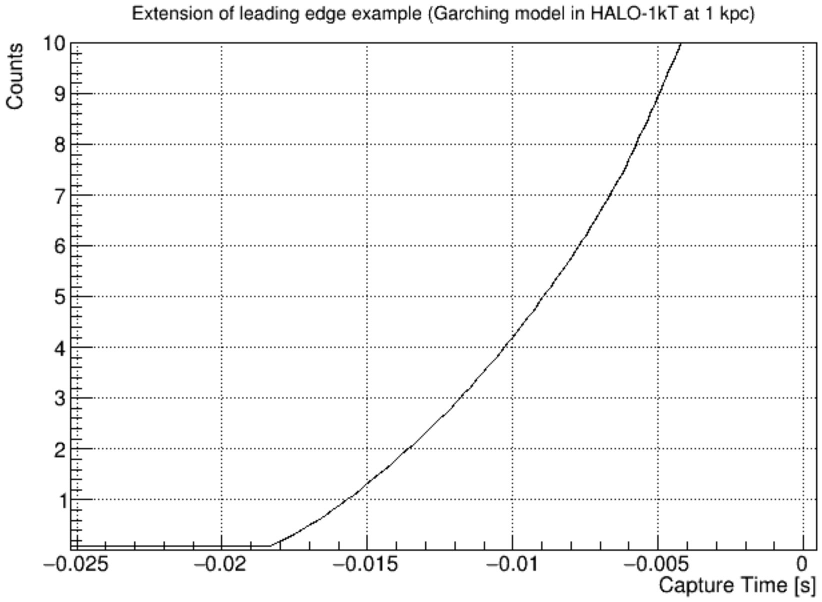

Figure 7.4: Extension of leading edge (from ≈ -15 ms to -65 ms) for the Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc in HALO-1kT.

7.2.5 Extraction of t_0

With the creation of the likelihood function, the extraction of t_0 is a trivial process. The package `RooFit` is used to make use of the MINUIT [100] minimization routines. For the Sukhbold et al. SFHo z9.6 $M_\odot$ model, we provide a simulation time series to fit to its own PDF, with the resulting likelihood function shown in Figure 7.5. The minimum t_0 extracted is our estimator ($\hat{t_0}$) of the true t_0. After the minimum is located, the errors on $\hat{t_0}$ can be found by scanning around the minimum and finding the points on each side of the minimum where the function increased by one; this is a 1σ error, where σ_- is the lower error and σ_+ is the upper error. Although the example found in Figure 7.5 is symmetric (within error), there is a tendency for these errors to be asymmetric. Asymmetries primarily arise throughout the mixing of models and PDFs in §7.2.6.

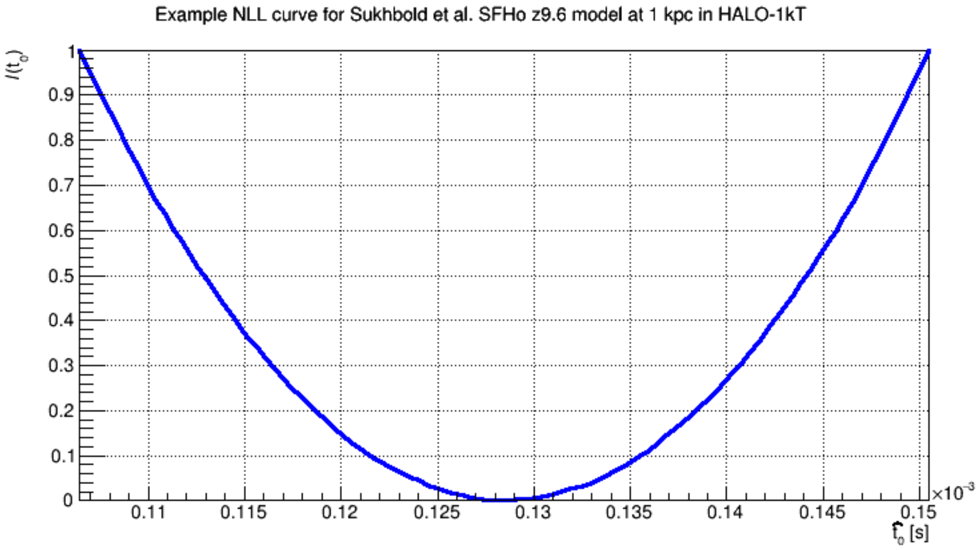

Figure 7.5: NLL curve for Sukhbold et al. SFHo z9.6 burst at 1 kpc in HALO-1kT.

Instead of computing the likelihood for each of the 10^4 bursts, which at 1 kpc is quite computationally intensive, we shall first turn our attention elsewhere. Thus far, we have neglected any discussion on two key components of this fitting technique. Firstly, expecting the observed time series of data to conform perfectly with any of our simulated models is naive. Secondly, for all models simulated in HALO-1kT, after the first 500 ms post core-bounce, the burst has already reached 32-45% of the events expected across the entire burst (this number comes from the values the mean light curves take at 500 ms post bounce). Is the tail end of the burst necessary and can the accuracy to which we extract $\hat{t_0}$ be improved? Both of these shall be discussed in detail in the subchapters that follow.

7.2.6 Mixing models and PDFs

The intent of this thesis is for each experiment to have a set of PDFs simulated and stored locally. Any observed burst would then be fit to these PDFs and a $\hat{t_0}$ could be extracted (by averaging the $\hat{t_0}$ providing from each model and PDF combination). Yet, we do not know which PDF best corresponds to a real SN signal and it is unrealistic to expect any one of the supernova simulations used to perfectly describe an observed time series. There is no escaping the systematic uncertainties that will present themselves when fitting our observed time series to these PDFs. To quantify these uncertainties, we mix all combinations of model and PDF, performing the same analysis as before.

First we consider fitting the Sukhbold et al. SFHo z9.6 $M_\odot$ model to the Sukhbold et al. LS220 z9.6 PDF. Although of different EOS, their progenitors share the same starting mass and their mean light curves do not diverge to the same extent as the 27.0 $M_\odot$ models. A systematic offset on the order of milliseconds was found to be introduced into the extracted $\hat{t_0}$ distribution when making use of the *entire* signal. This was anticipated as each model has slight variations in the length and size of the observed neutrino burst. On top of this, their spectra vary greatly. Additionally, if one were to compare the low mass models to the high mass models, the higher progenitor mass models maintain a higher event following the neutronization burst. There is a method of reducing both the systematic offset and tightening the standard deviation when mixing models and PDFs; to do this, we introduce a fixed time window that starts at the first event in the burst.

7.2.7 Selection of a time window

The implementation of a time window arose out of necessity. Provided HALO-1kT has sufficient statistics ($> 10^3$ events, ≈ 1 kpc), the inherent shape difference between the model's time series and the PDFs will be too pronounced and the fit will fail with remarkable consistency.

To get around this problem, we look to when the mean cumulative event curves are most similar. Figure 7.6 shows each model's respective light curve out to the models specific end time. Here, we focus on the first 500 ms. Each model posseses similarities in the leading edge, as the primary contribution in this regime is the ν_e (HALO and HALO-1kT's main sensitivity) from the neutronization burst. Even though the magnitude of the neutronization burst varies, the overall shape is similar. It is only around 100 ms after t_0 that the larger solar mass models begin to deviate substantially. To reduce systematic offsets, we then introduce a 100 ms time window on our time series. The first event is selected as per the requirements outlined in §7.1 for each experiment and indicates the start of the time window. Within the time series, we discard any events that have time indices that are larger than the first event time plus 100 ms. The results for the Sukhbold et al. SFHo z9.6 solar mass model fit using all Sukhbold et al. models is shown in Figure 7.7.

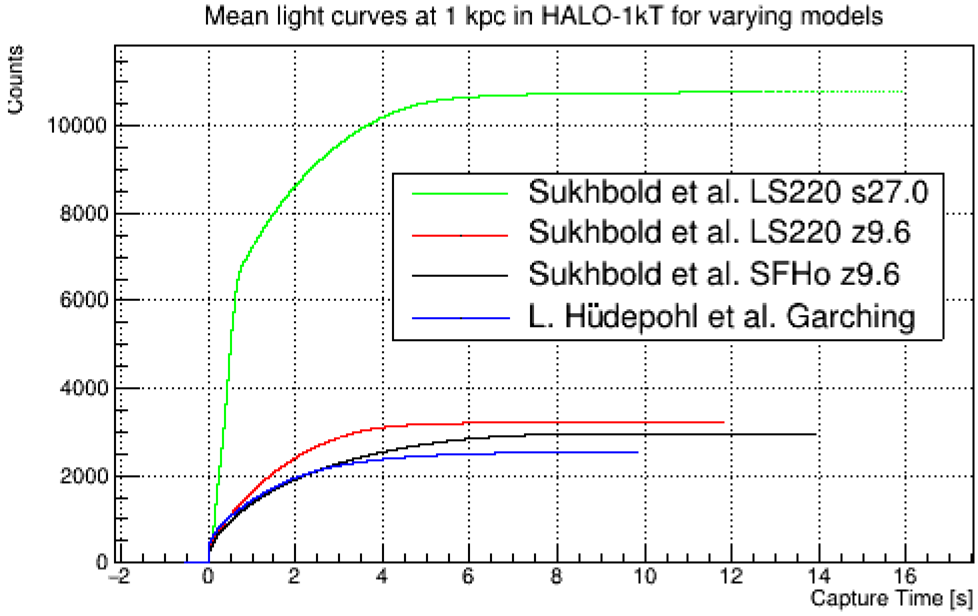

Figure 7.6: Mean light curves for various simulated models in HALO-1kT at 1 kpc.

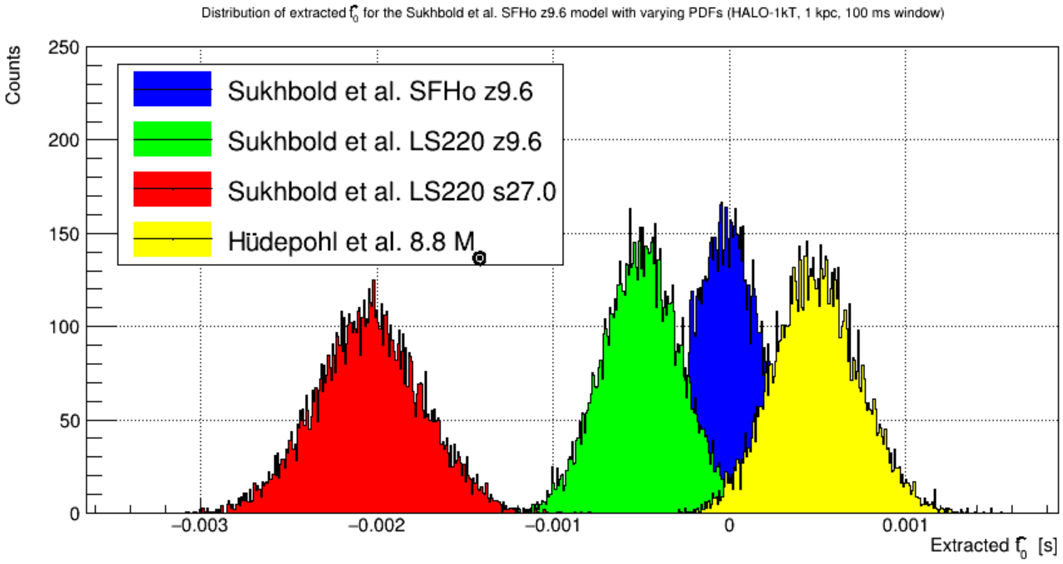

Figure 7.7: Distribution of extracted $\hat{t}_0$ for varying PDFs, using Sukhbold et al. SFHo z9.6 time series data at 1 kpc.

The systematic uncertainties introduced from fitting the lower mass Sukhbold et al. SFHo z9.6 model to the Sukhbold et al. LS220 z9.6 PDF have been reduced. This is not entirely unexpected, as even though there are fewer events after the neutronization burst and the accretion phases, their contribution to the likelihood function is negligible. The offset from this particular model and PDF combination has been reduced to -500 μs. The standard deviation of the distribution is also much smaller (230 μs), resulting in a more precise extracted $\hat{t_0}$. The improvements observed in the intermixing of models and PDFs also extends to fitting a model to its own PDF. With the introduction of our time window, the offset is only 14 μs and the standard deviation is reduced by a factor to 220 μs, as seen in Figure 7.8.

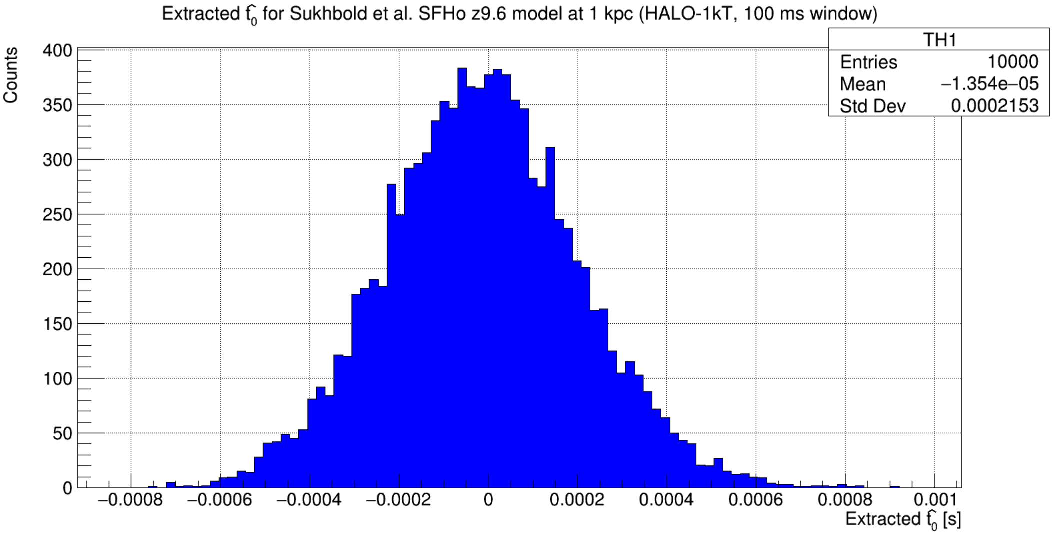

Figure 7.8: Distribution of extracted $\hat{t_0}$ from the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT with a 100 ms time window.

The initial time window chosen to experiment with in this work was arbitrary. There are any number of points where the model's mean light curves could have been cut prior to their divergence from one another. Since it was observed that the reduction in time window reduced the overall systematic offset and standard deviation, we elected to determine the optimal window size. To do so, we carry out the same fitting routines as listed above, for time window sizes between 10 and 100 ms in 10 ms increments. The optimal time window is selected such that the offset and standard deviation is minimized. For the Sukhbold et al. SFHo z9.6 $M_\odot$ model fit to itself, the results are shown in Figure 7.9. As expected, an increase in the time window size will cause the distribution of $\hat{t_0}$ to drift further from t_0 (and with an increased standard deviation). Although optimal performance is achieved with a 10 ms time window for the Sukhbold et al. SFHo z9.6 $M_\odot$ data, this is not the case for the Hüdepohl et al. 8.8 $M_\odot$ model, as seen in Figure 7.10. The Hüdepohl et al. 8.8 $M_\odot$ model in HALO-1kT has a relatively sparse event distribution in the initial 15 ms prior to t_0; this sparsity will result in a low population in the selected window and as a consequence, an extracted $\hat{t_0}$ that is inaccurate. By making use of a 30 ms window, this will begin to include the neutronization burst and a sufficient number of events will be within the window to extract $\hat{t_0}$. For this technique, we then settle upon an optimal window size of 30 ms for HALO and HALO-1kT. For the preliminary analysis of this technique in SNO+, a $\approx$ 50 ms time window was used (approximate in the sense that the simulation was cut 50 ms after core-bounce).

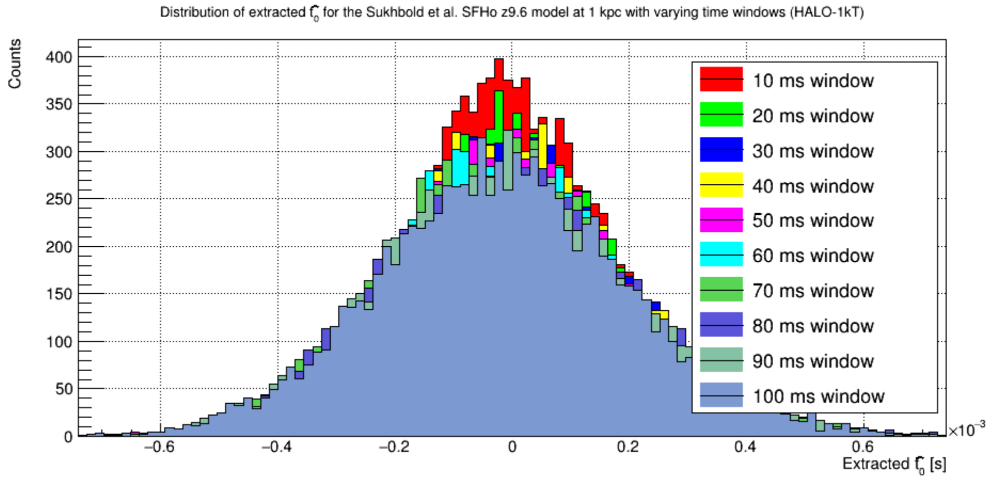

Figure 7.9: Impact of time window length on extracted $\hat{t_0}$ precision for Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT

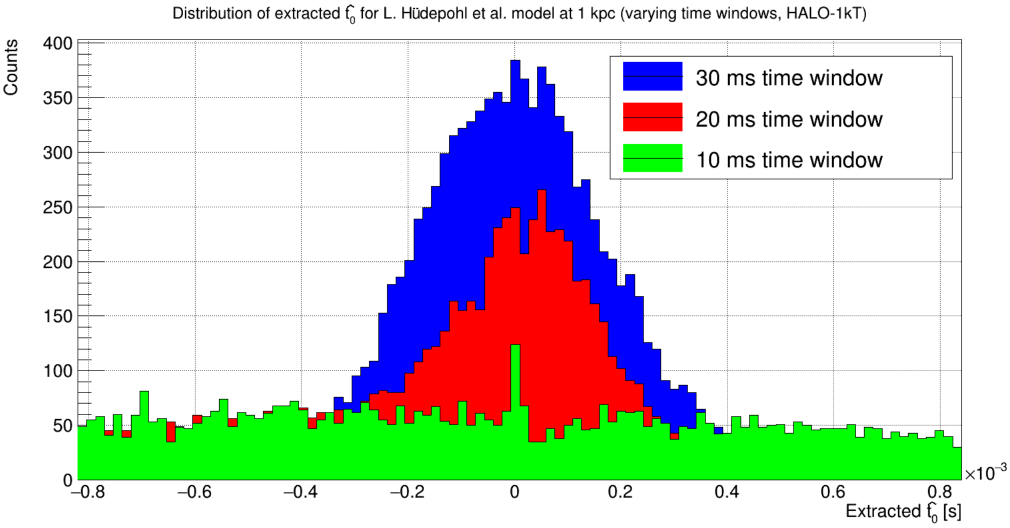

Figure 7.10: Impact of time window length on extracted $\hat{t_0}$ precision for Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc in HALO-1kT. The Hüdepohl et al. 8.8 $M_\odot$ models performance collapses below a time window size of 30 ms, unlike any other models simulated in HALO-1kT.

7.2.8 Estimating systematic uncertainties

To finalize our discussion on the extraction of $\hat{t_0}$ using the NLL technique, we seek a distribution of the extracted $\hat{t_0}$ that accounts for all systematic uncertainties. To achieve this goal, we sample all the histograms of $\hat{t_0}$ obtained by mixing models and PDFs, treating them all of equal weight and fill a 1D histogram. The result can be found in Figure 7.11 for the extracted $\hat{t_0}$ at 1 kpc in HALO-1kT. With the systematic uncertainties accounted for, the distribution is offset by -20 μs from the true t_0 and has a standard deviation of 430 μs. This uncertainty on the extracted $\hat{t_0}$ falls in line with the precision needed to triangulate (as discussed in §4.4). The next question we seek to answer is how well this precision can be maintained at further distances.

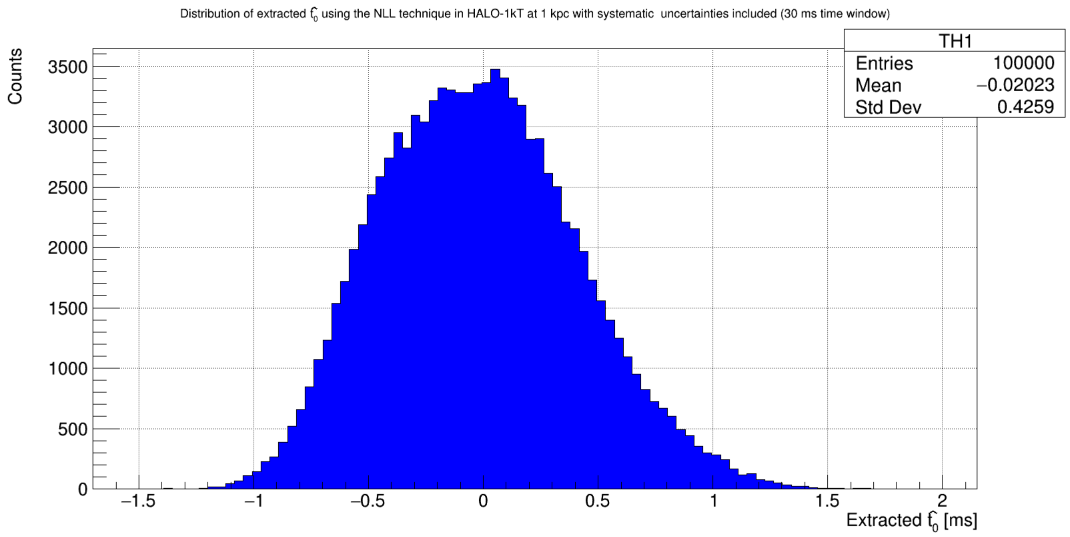

Figure 7.11: Distribution of extracted $\hat{t_0}$ using the NLL technique at 1 kpc in HALO-1kT with systematic uncertainties included.

7.2.9 Performance as a function of distance

The same procedure, applying the NLL fit to Monte Carlo simulation data, was carried out for SN bursts at distances of 2-5 kpc. The results for this can be found in Figure 7.12, which shows the extracted $\hat{t_0}$ in HALO-1kT at distances 1-5 kpc with the systematic uncertainties included. The resolution of the extracted $\hat{t_0}$ is no surprise, however, the deviation from the Monte Carlo truth, t_0, with increasing SN distance is. The distribution of extracted $\hat{t_0}$'s at 1 kpc is not perfectly Gaussian, with a shoulder on the positive $\hat{t_0}$ side of the distribution that tapers off at a slower rate. This feature is not present at further distances and the distribution approaches a state of symmetry as it shifts towards negative $\hat{t_0}$'s. In Figure 7.13 the precision of the extracted $\hat{t_0}$'s is shown with the integrated 1σ errors contained within the region of the graph. This discussion shall be picked up in §8, where the techniques illustrated in this chapter are compared to one another and a formal recommendation on the technique to be implemented for HALO-1kT is made.

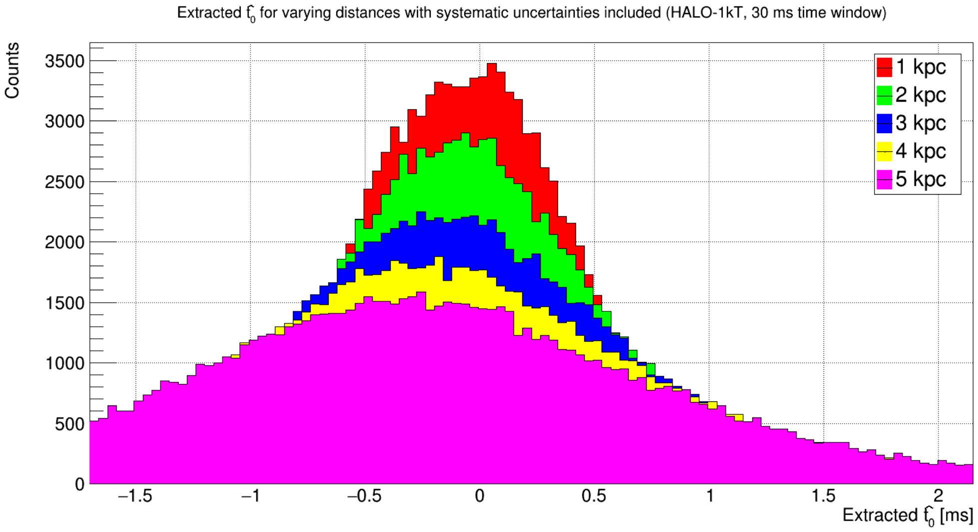

Figure 7.12: Distribution of extracted $\hat{t_0}$ for varying distances from the NLL fit in HALO-1kT, with systematic uncertainties included.

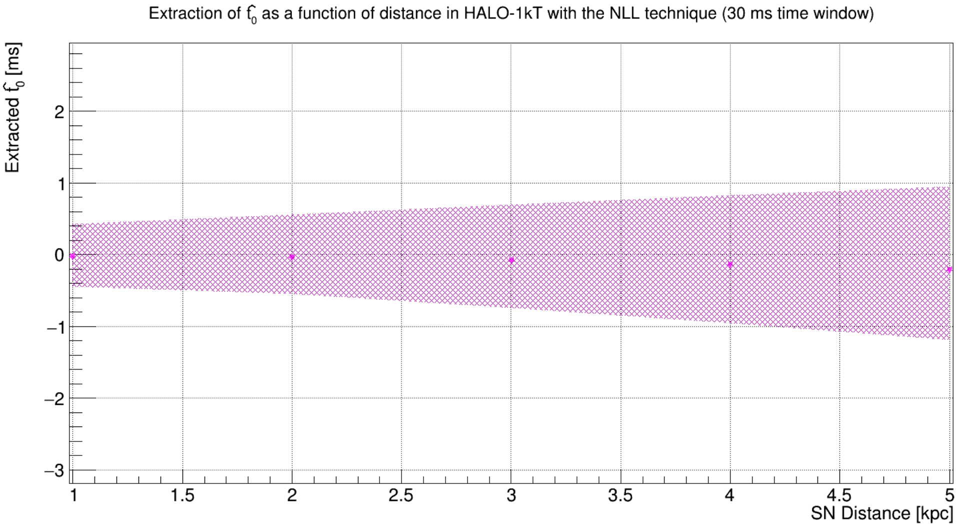

Figure 7.13: Performance curve for extracted $\hat{t_0}$ from the NLL technique in HALO-1kT.

7.3 Constant Fraction Discriminator (CFD)

7.3.1 Motivation

When making use of the NLL technique to extract $\hat{t_0}$, the distribution of extracted $\hat{t_0}$ will shift from the true t_0 as the simulation distance is increased. As seen in Figure 7.12, when accounting for systematic uncertainties and using a 30 ms time window, the difference in central value from 1 to 5 kpc is ≈ 163 μs. We call this the “walk”: a drifting of the trigger time as a function of signal amplitude. This feature of the NLL technique is present even if systematic uncertainties are not included. It is also present if the normalization is excluded and the time window is removed. Simply put, this is inherent to the technique used and not a result of the modifications included to optimize its configuration. To correct this, we sought to implement a digital constant fraction discriminator (CFD). Traditionally, these are used in signal processing as an alternative to threshold triggering and expressly to remove walk from trigger times. Although this technique was initially intended to be used as a correction to the NLL technique, its performance at estimating t_0 on its own surpassed expectations.

7.3.2 General formulation

For the CFD technique, there is no explicit mathematical formulation or minimization needed. Instead, this technique relies solely on the observed time series in the detector. The method of extracting $\hat{t_0}$ can be broken down as follows

- Place the observed time series into a histogram with non-uniform bins as described in the preamble of §7.

- Clone the observed time series.

- Invert the cumulative signal and scale it *down* by an arbitrary coefficient. We label this coefficient the constant fraction, c_f.

- Shift the inverted signal ahead of the Monte Carlo truth signal by some time t_s, which we call the temporal offset.

- Sum the inverted and Monte Carlo truth signal together. Since the bins are of non-uniform width and vary in range, the summed histogram is created by incrementing 1 ms forward from the first "event" in the inverted signal. At each increment, a bin is created in the summed histogram and given a value that corresponds to the sum of the two histograms at this time.

- Associate $\hat{t_0}$ with the time of zero-crossing of the composite signal.

Using the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc, we observe the effect of processing a SN burst through our digital CFD in Figure 7.14. On the left is the original light curve as observed in HALO-1kT, where each lower bin edge corresponds to the time stamp of a new event. The middle histogram is the inverted and delayed signal, which has been scaled and offset by factors c_f and t_s respectively. The culmination of summing these together yields the histogram depicted in the right frame. For discussion on the selection of two parameters of this fit, see §7.3.5. Of interest is not the point where the function is minimized, but rather the point where the sum of two histograms is zero.

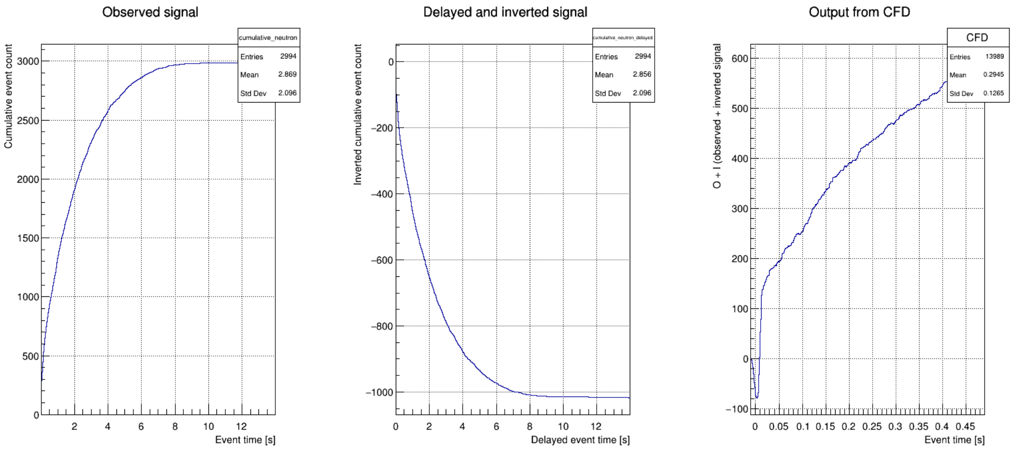

Figure 7.14: Example of CFD output using the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT.

7.3.3 Extraction of t_0

To extract $\hat{t_0}$, the summed histogram is parsed until the point it crossed the x-axis is located. We label this point the intersection point. Referring back to Figure 7.14, we note that the for this particular burst the intersection point is located at $\approx t = 10$ ms in the right plot. We repeat the same procedure for all 10^4 simulated bursts for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT to obtain the distribution of extracted intersection point in Figure 7.15. Any discussion on the resolution of this distribution shall be relegated to §7.3.5, where the selection of c_f and t_s will drastically impact the precision to which $\hat{t_0}$ is extracted. Instead, here we shall briefly discuss the distributions positioning which is 8.48 ms offset from the Monte Carlo truth t_0. An excepted result since this technique is not parameterized as a function of t_0. This offset is of little concern since we can use its relative location with respect to t_0, provided it is consistent amongst our library of models, as a direct estimator of $\hat{t_0}$. This discussion will be picked up in the next subchapter where the model dependence of this technique is explored.

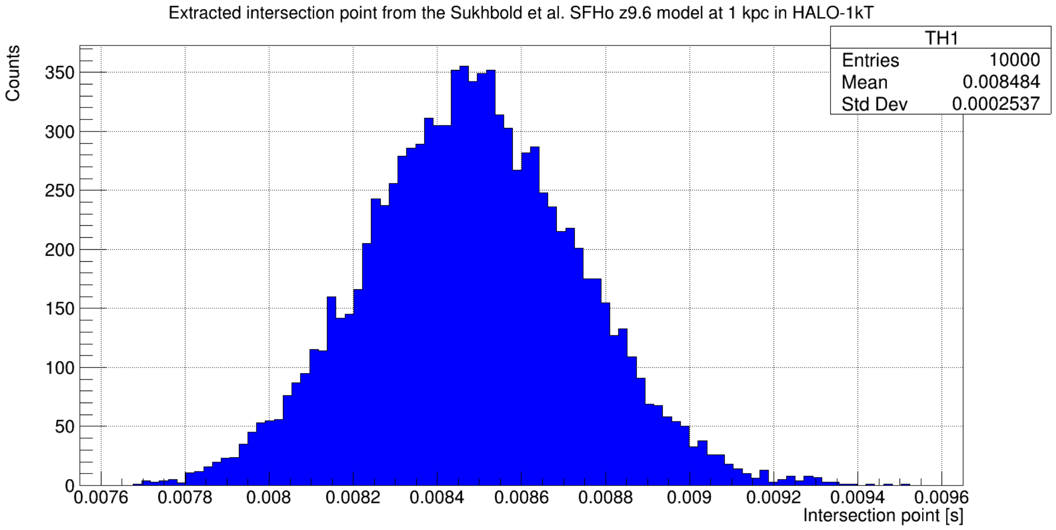

Figure 7.15: Distribution of extracted intersection points for the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT. The above does not have an optimized configuration for the two parameters.

7.3.4 Supernova model dependence

The positioning of the intersection point is highly model dependent. This dependence arises from the neutronization burst, which will vary in length and event rate for the library of models chosen for this thesis. To assess the potential impact of this model dependence we process each of the 10^4 bursts for the remaining models through our digital CFD. Depicted in Figure 7.16 shows the distribution of intersection points for each model at 1 kpc in HALO-1kT. Between the models depicted in the aforementioned figure, there is an $\approx 800\ \mu s$ separation between the two outer most distributions. The goal of using this technique is to ensure that this separation is minimized; in its current configuration this technique would severely underperform if the systematics of the model dependence were accounted for and a direct comparison was undertaken. To alleviate this we look to optimize our parameters such that the model dependence can be minimized.

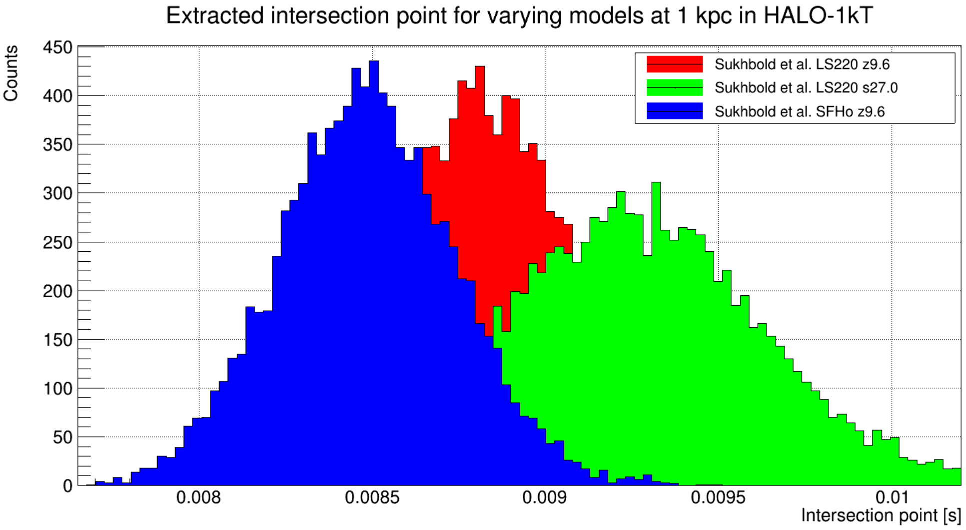

Figure 7.16: Distribution of extracted intersection points for varying models at 1 kpc in HALO-1kT. The above distributions do not have an optimized configuration for the two parameters of our CFD.

7.3.5 Optimal configuration

Previous results shown for the CFD technique had not made use of an optimal configuration for the two parameters (c_f and t_s). These can be finely tuned to significantly impact the performance of the CFD technique, for better or for worse. Since the neutronization burst is heavily pronounced in HALO and HALO-1kT, owing to its flavour sensitivity, the constant fraction plays an integral role in determining when the summed histogram ramps back up across the x-axis. If a value is chosen for the constant fraction that is sufficiently small, the fit can fail to pass the x-axis consistently. This is especially true for models with sparse event populations prior to t_0 (i.e. Hüdepohl et al. 8.8 $M_\odot$). The temporal offset also plays a significant role in the location of the intersection point and its consistency. If the temporal offset is too large, the inverted signal will have the dominant contribution to the summed histogram in the early times, once again owing to the neutronization burst. This would not be an issue if it did not introduce a large degree of variance in the zero-crossing point of the summed histograms. An example of the results from a CFD configured with an unoptimized temporal offset (30 ms) and constant fraction (40%) can be found in Figure 7.17, for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT.

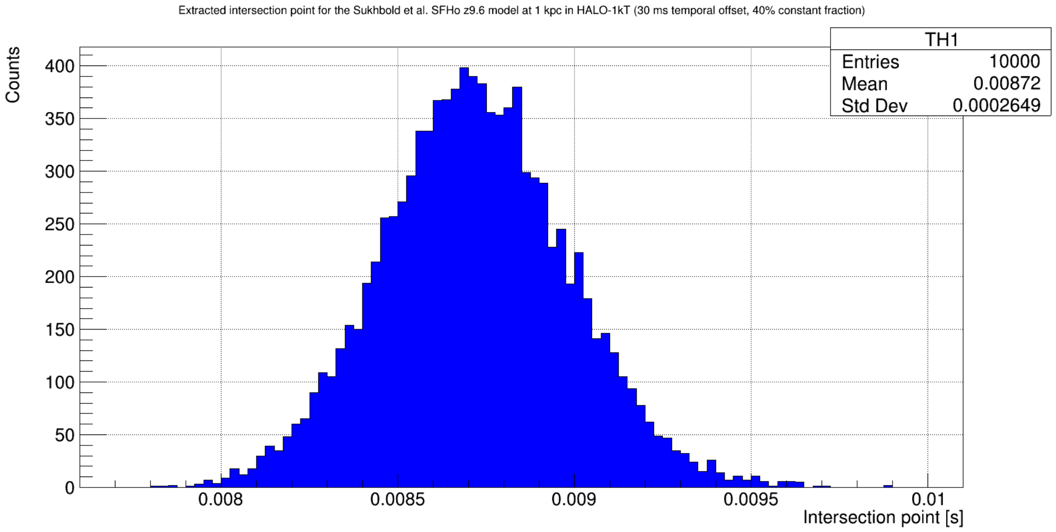

Figure 7.17: Distribution of extracted intersection points for Sukhbold et al. SFHo z9.6 model with an unoptimized configuration.

To best optimize our CFD parameters, we perform the CFD fit for a range of parameter values. For the constant fraction, we elect to use values between 0.1 and 0.9 in 0.1 increments. The temporal offset is chosen between 10 and 100 milliseconds, in 10 millisecond increments. Hence, there are 81 different combinations for the two parameters. To determine the optimal configuration, we fetch the extracted intersection point distribution for each model in a given parameter set. We then record the maximum distance between the central values of the two histograms on the outer edges of the four histograms. Once this procedure is done for the 81 combinations of parameters, we make use of the parameters that resulted in the narrowest spacing between the four histograms. The results of this preliminary analysis yielded an optimize constant fraction of 0.3 and a temporal offset of 10 ms. To further refine this we perform the same optimization on a narrower range of parameter options. For the constant fraction, a range of 0.2 to 0.4 is chosen with a 0.01 increment. The temporal offset is similarly refined to a range of 10 ms to 20 ms with a 1 ms increment. This optimization yielded a new optimized configuration of 13 ms for the temporal offset and 0.34 for the constant fraction. The distribution of intersection points for the various models can be found in Figure 7.18, showcasing excellent consistency between our models. As such, we can begin to account for the systematics to showcase the precision of the CFD technique for extracting $\hat{t_0}$.

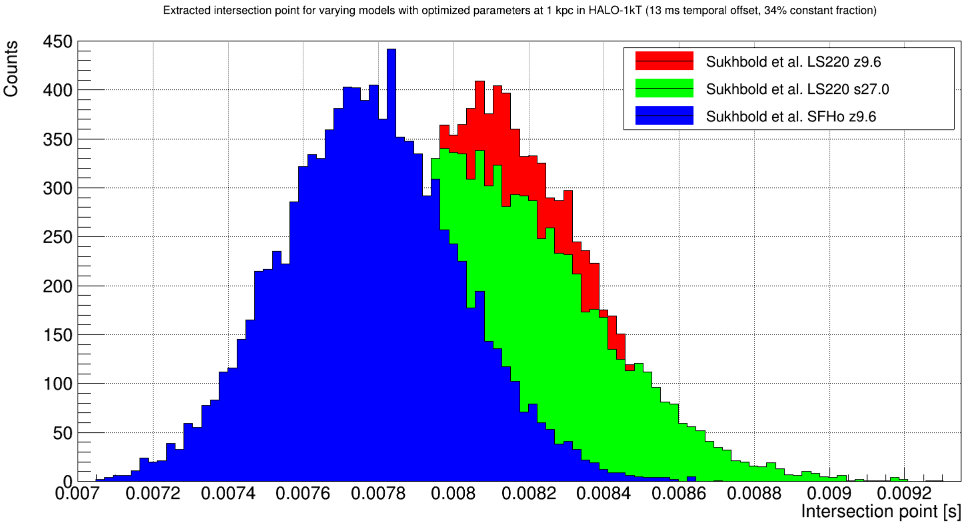

Figure 7.18: Distribution of extracted intersection points for varying models with optimized parameters (results are for HALO-1kT at 1 kpc).

7.3.6 Estimating systematic uncertainties

Once again, systematic uncertainties that introduce themselves via applying this technique to varying SN models need to be taken into account. We treat each output from the CFD technique for the 10^4 bursts per model with equal weight, sampling them 10^4 times and filling a 1D histogram for the intersection point with systematic uncertainties included. The result of this can be found in Figure 7.19. Recall that, with systematics included, the NLL fit extracted a $\hat{t_0}$ with an offset of -20 μs and a standard deviation of 430 μs. The CFD technique improves upon this result by a factor of 1.56, narrowing the standard deviation to 270 μs. Further discussion on the comparison between techniques shall be taken up in §8.

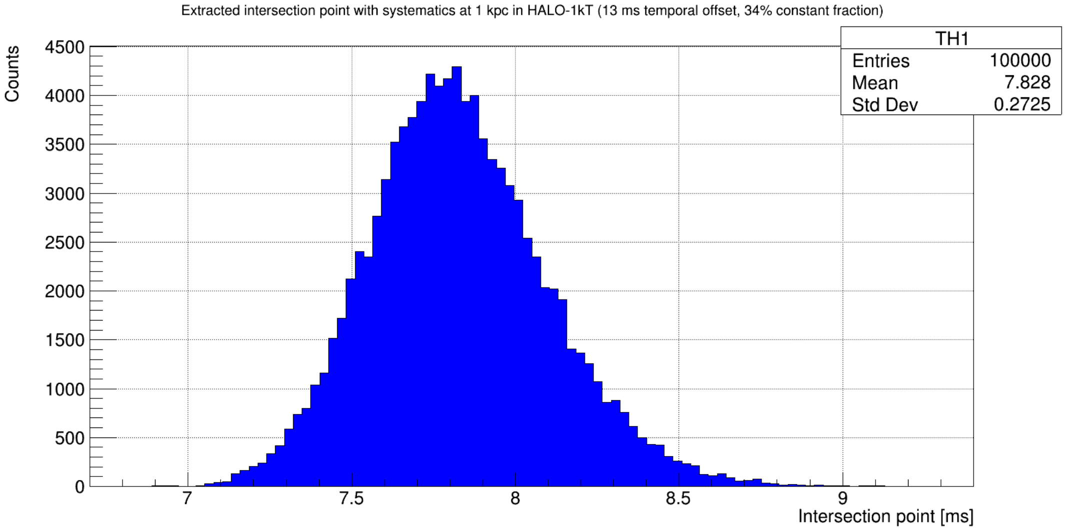

Figure 7.19: Distribution of extracted intersection points at 1 kpc in HALO-1kT with systematics included.

Here, we introduce the concept of an "offset" correction for the CFD technique. When this term is used it is meant to signify the expected output from the CFD at 1 kpc (for HALO-1kT, for HALO this shall be 500 pc, more on this in §9), with the systematic uncertainties included. This offset correction is used to shift all values extracted from the CFD technique by the expected average. As was demonstrated prior, the CFD in its optimum configuration should expect an offset of 7.8 ms from the Monte Carlo truth t_0. This offset correction is applied to all simulated SN distances. Recall that the intent for introducing this analysis tool was to potentially mitigate the walk observed in the NLL technique. By applying the offset correction to distances greater than 1 kpc, the hope is that each intersection point distribution shall be centered at zero.

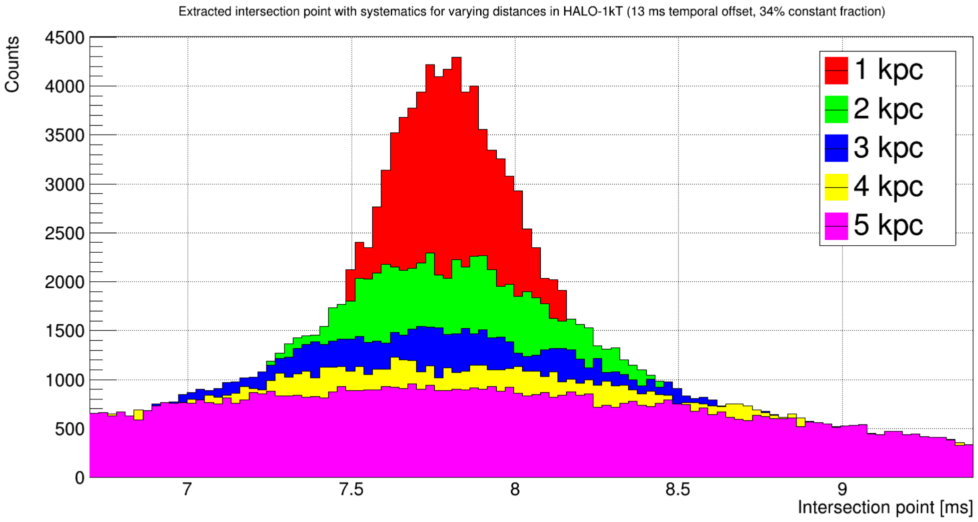

Figure 7.20: Distribution of extracted intersection points at varying distances in HALO-1kT with systematics included.

7.3.7 Performance as a function of distance

We now conduct the same fitting technique as described previously to all of the simulated distances. The results can be found in Figure 7.20, which is the extracted intersection point distribution (or the intersection point distribution without the offset correction applied) for each of the five distances. What becomes clear from this comparison is that the walk, as identified in the NLL technique, is still present in the CFD technique. After applying the offset correction, the 5 kpc distribution is off center by 250 μs. By integrating the CFD technique into our analysis, our principle objective of reducing the walk has not been rectified. With that said, we have added a competitive technique to our toolkit, provided sufficiently high statistics of events are observed.

7.4 PDF independent linear fit

7.4.1 Motivation

Perhaps the greatest drawback of the NLL technique is the requirement to fit the observed time series of data to a library of PDFs. A tremendous amount of simulation time is required to obtain adequate statistics to create these PDFs. Not only is the simulation time of concern, but so too are the systematic uncertainties introduced when mixing models and PDFs, especially if we expand this analysis to dozens of SN models. As such, there is a desire to explore solutions that give us a free-form PDF that the observed time series could be fit to. Preliminary discussion centered around making use of a ≈ 5 degree polynomial as a PDF. The parameters of this polynomial could be used to provide an estimate of t_0, provided this was an appropriate description of the observed burst. Five parameters was deemed sufficient as the aim of this fit is not to perfectly fit the cumulative time series; its singular purpose is to provide an accurate estimate of t_0. As the polynomial fit was being explored, a substitute was put forth that would be more efficient. In HALO and HALO-1kT, the mean light curves have two regions where the event rate is approximately linear. In Figure 7.21, these regions correspond to the neutronization burst and accretion phase. Although not perfectly linear and subject to curvature as the electron neutrino flux ramps up or begins to cool, if an accurate t_0 could be extracted, we would be satisfied. The following subchapters will go into detail on the formulation of our linear fit; its application; and a discussion of its performance as a function of distance and model.

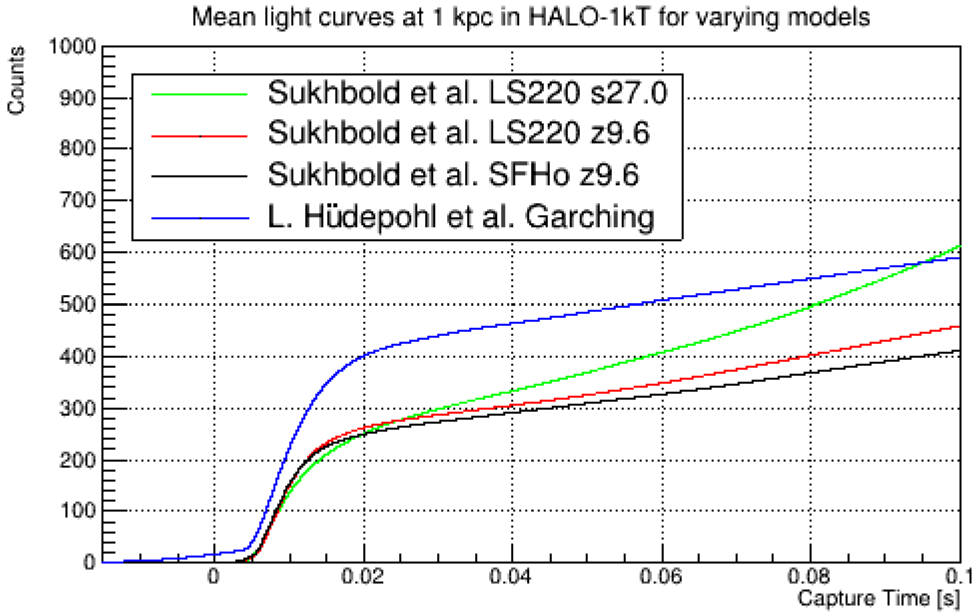

Figure 7.21: Mean light curves observed in HALO-1kT for each supernova model at 1 kpc, cutoff at 100 ms past t_0.

7.4.2 General formulation

In an attempt to fit the observed light curve with our linear approximation, a simple linear fit was implemented as

$$f(t) = p_0 + p_1 \cdot (t - p_2) \cdot (t > p_2) \tag{7.10}$$

where p_0, p_1, and p_2 are parameters of the fit. The additional (t $> p_2$) term constrains the fit to times greater than p_2. A more thorough discussion of the parameters occurs in §7.4.2.1 and §7.4.2.2. The `ROOT` data analysis framework is used to handle the fit, where Equation 7.10 is implemented as a `TF1` (1D function). For each observed light curve, the first event was used as the lower bound on the fit range. For the end point, preliminary explorations used the derivative of the cumulative time series (i.e. event rate) to determine the maximum slope (this would occur shortly after t_0 during the neutronization peak). Once located, the fit was constrained to 5 ms past this point. Unfortunately, the further the observed SN burst was simulated, locating the maximum event rate became more problematic owing to fewer events in the neutronization burst. To circumvent this, we implemented two additional terms into our fit as

$$f(t) = p_0 + \Big(p_1 \cdot (t - p_2) \cdot (t > p_2) \cdot \big(t < (p_4 + p_2)\big)\Big) + \Big(p_1 \cdot p_4 + p_3 \cdot \big(t - (p_4 + p_2)\big)\Big) \cdot \Big(t > (p_4 + p_2)\Big) \tag{7.11}$$

where p_3 and p_4 are additional parameters implemented to assist in locating the appropriate end point of the neutronization burst. As shall be seen in the coming subchapters, $f(t)$ in its current form proved remarkably efficient and accurate at providing an estimate on $\hat{t_0}$.

7.4.2.1 Physics interpretation of parameters

Equation 7.11 is intended to fit the neutronization *and* a portion of the accretion phase of the neutrino signal. Although this fitting technique is rather simplistic, it is worthwhile discussing each parameter.

- p_0 is the intercept point of both lines. If a line is drawn along the neutronization burst, it will intersect with the y-axis at some arbitrary point. This was done for each of the four models and the average was used for the initial condition.
- p_1 is the event rate of the SN burst during the neutronization burst. It is obtained by averaging the event rate in the neutronization burst for each model. The bounds were chosen to accommodate models with large event rates.
- p_2 is the offset of the fitted line to the neutronization burst with respect to the true t_0. Expected value for p_2 comes from the Monte Carlo truth value for t_0. 10 ms

bounds on each side were chosen as other techniques explored (NLL, CFD, etc.) had all extracted $\hat{t_0}$ within these bounds.

- p_3 is the event rate of the SN burst in the accretion phase. It is obtained by averaging the event rate throughout the accretion phase.

- p_4 is the width of neutronization burst. Summed with p_2, this parameter will estimate the transition from the neutronization burst to accretion phase. By averaging the width of the neutronization peak for each model, p_4 was obtained. Symmetric 5 ms bounds were chosen to accommodate fluctuations

7.4.2.2 Initial parameters

When `ROOT` fits a `TF1` to a `TH1` (1D histogram), the success of the fit is highly dependant on the initial conditions provided by the user. Listed below are the initial conditions used for each each of the five parameters in $f(t)$.

Parameter	Initial Value	Limits
p_0	0.0025	[0, 0.005]
p_1	13	[-5, 25]
p_2	0	[-0.01, 0.01]
p_3	1	[-1, 3]
p_4	0.01	[0.005, 0.015]

Table 7.1: Initial parameters for the linear fit. The bounds of each parameter are generally symmetric about the mean value. Slight variances from this symmetry arise from the fit encountering rare SN bursts that fall outside of the original bounds. It is worth noting that the parameters encountered here are specific to HALO and HALO-1kT.

With these initial conditions, we can apply our linear fit to the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT. In Figure 7.22, the observed light curve is fit with $f(t)$. As expected, the fit does not perfectly describes the nuances of the SN signal. This of course is of little concern, as we are only interested in obtaining an accurate and precise estimate of t_0.

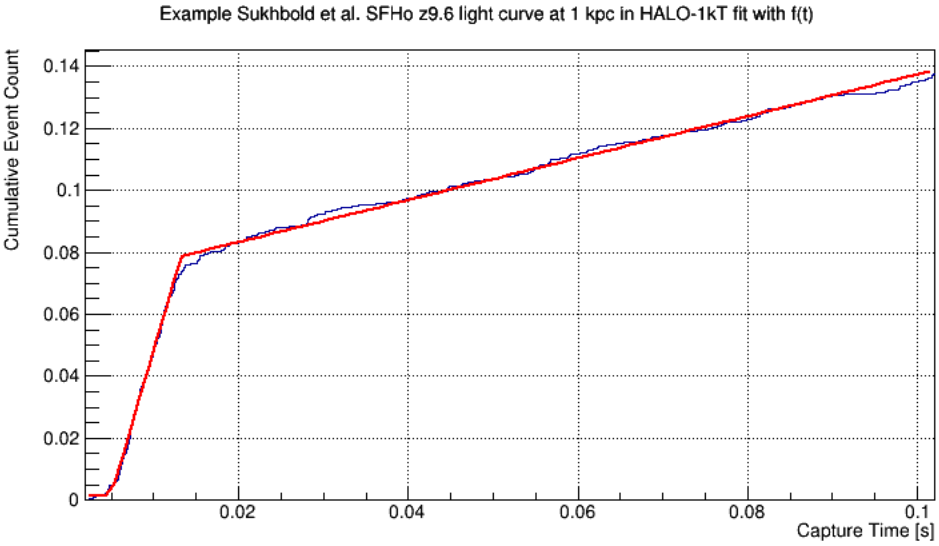

Figure 7.22: Observed Sukhbold et al. SFHo z9.6 light curve in HALO-1kT with linear fit

7.4.3 Extraction of t_0

Similar to the CFD technique, when the linear fit is applied to the SN burst, the returned p_2 value will deviate from the Monte Carlo truth t_0. This is of little concern as this can once again be corrected if the deviation from t_0 is known a priori. This will be discussed in further detail in §7.4.5. As such, the returned value of p_2 in our fit from the previous subchapter can be used as an estimate of t_0 should its displacement be consistent. We repeat the fitting procedure for 10^4 Sukhbold et al. SFHo z9.6 $M_\odot$ bursts in HALO-1kT at 1 kpc and fill a histogram with each of their respective p_2's. The results can be found in Figure 7.23. Each of the extracted p_2's are in excellent agreement with one another and are well distributed about the central value of 4.9 ms.

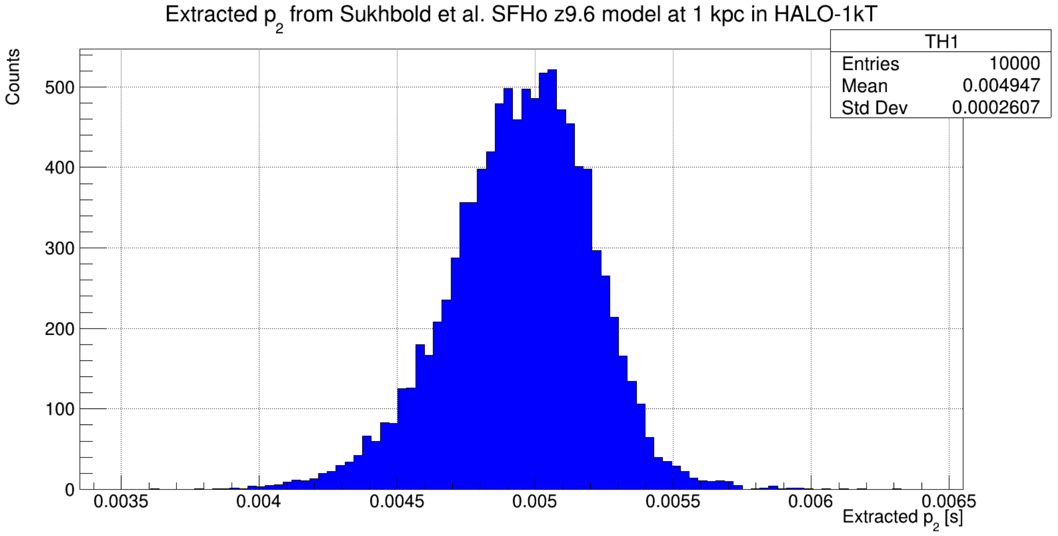

Figure 7.23: Distribution of extracted p_2 for the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT

Asymmetry arises in the p_2 histogram as the fit has difficulties reconciling the two linear regions of the curve. This feature is observed in a 2D scatter plot of p_2 and p_4, in Figure 7.24. The other four parameters can be found in Figure 7.25 for this particular fit.

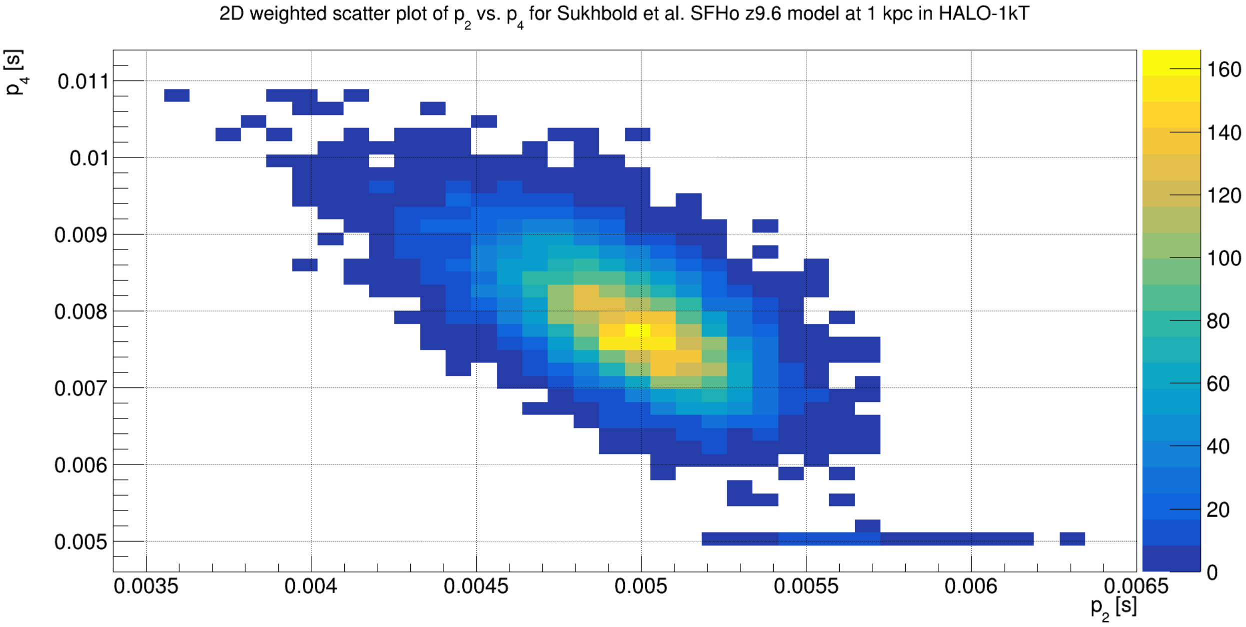

Figure 7.24: 2D weighted scatter plot of the anti-correlation of p_2 and p_4 for the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT

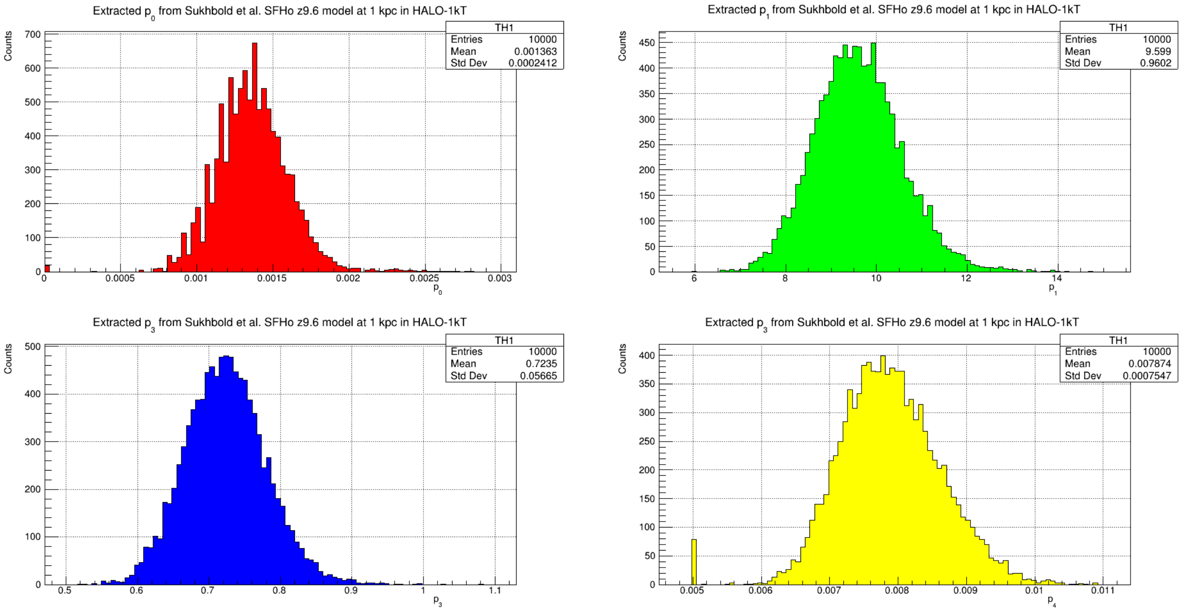

Figure 7.25: Distributions of p_0, p_1, p_3, and p_4 for the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT.

7.4.4 Supernova model dependence

Since the linear fit is not parameterized directly as a function of t_0, there is concern that the p_2 extracted by the fit, for the varying models, may not be in agreement with one another. Since the extracted p_2 is dependent on the slope of the neutronization burst and the accretion phase, each model's unique profile will determine its relative position of the extracted p_2. For each of the simulated models, we fit all 10^4 bursts with the linear function and extract their p_2. The results can be found in Figure 7.26.

There is excellent agreement between our library of models, where the distance between the peaks of the distributions is ≤ 1 ms. This variance is not unexpected, nor can it be removed, simply reduced (a potential fix to this is addressed in §7.4.7). The shift in the distributions are much in line with the systematic uncertainties observed in the NLL technique when mixing models and PDFs.

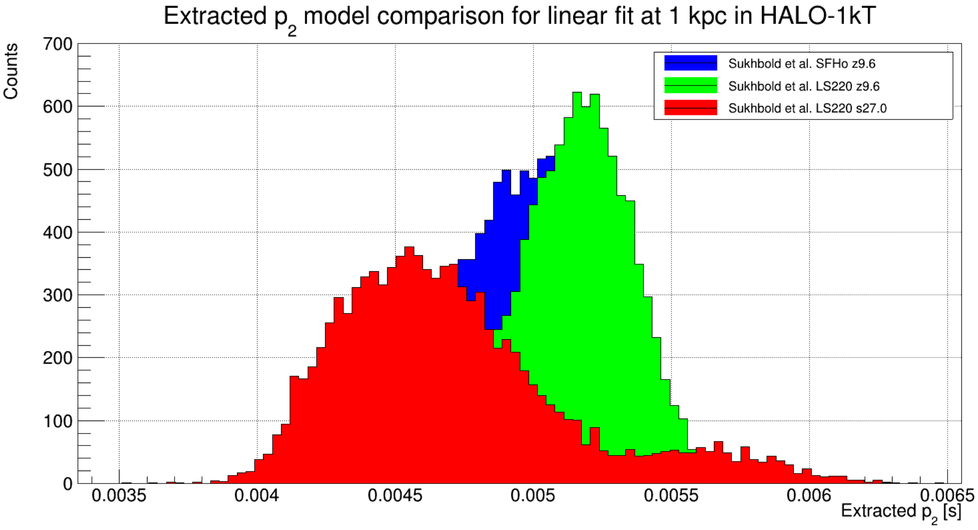

Figure 7.26: Extracted p_2 for varying simulated models in HALO-1kT at 1 kpc.

7.4.5 Estimating systematic uncertainties

We now seek to account for the systematic uncertainties that present themselves outlined in the previous subchapter. Since there is no intermixing of models and PDFs, we simply treat the four distributions found in Figure 7.26 with equal weight and populate a 1D histogram with 10^4 samples. The results can be found in Figure 7.27. At 1 kpc, the performance of this technique is comparable to the other techniques thus far, yielding excellent results for such a simple model. The mean of the distribution, what represents the offset correction for the linear fit, is 4.9 ms with a standard deviation of 350 μs.

7.4.6 Performance as a function of distance

As the linear fit is applied to further distances, the distributions of extracted p_2 for each respective distance smear into one another. Their central values will tend to coalesce around some mean value such that each mean extracted p_2 is equivalent across the SN models. This trend can be observed in Figure 7.28, where the extracted p_2 for each

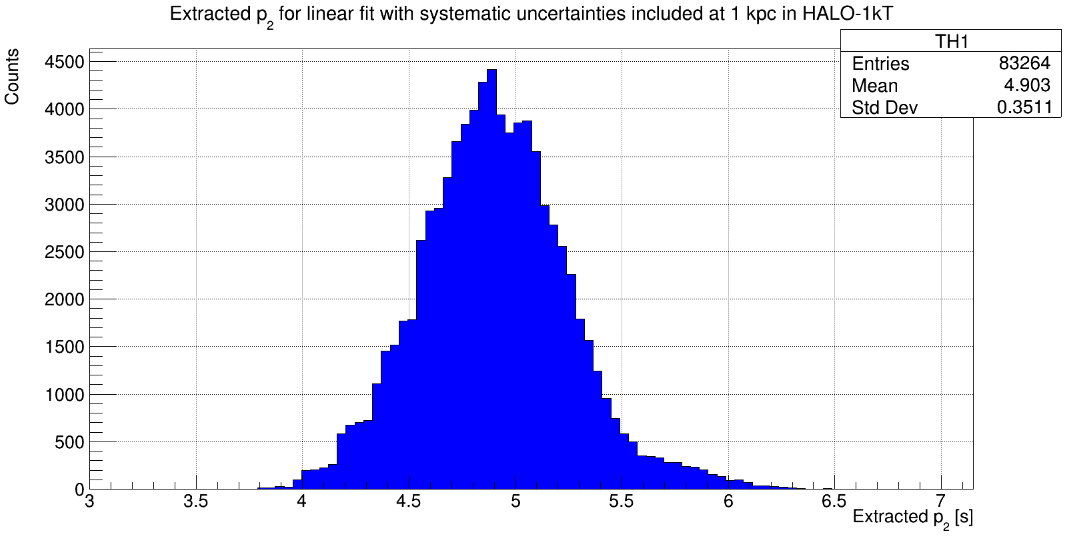

Figure 7.27: Distribution of extracted p_2 with systematic uncertainties included at 1 kpc in HALO-1kT

distance in HALO-1kT is shown, with the systematic uncertainties included. The walk observed in each technique thus far is also present in the linear fit, resulting in the peak of the distribution shifting further from the offset observed at 1 kpc. We once again utilize the concept of offset correction to shift our extracted p_2 distributions back to the Monte Carlo truth t_0. For the linear technique, the offset correction, as observed in Figure 7.27, is 4.90 ms. When comparing results with the other techniques in §8, all results for the linear technique shall have the offset correction applied.

7.4.7 Additional modifications

To facilitate a direct comparison between the results that come from SNO+ and HALO-1kT, the same models were simulated to ensure that any comparison was fair. As such, even though the initial analysis carried out for HALO-1kT primarily made use of the Hüdepohl et al. 8.8 $M_\odot$ model, it was dropped later on since it was not integrated into `sntools` and its integration was outside the time-scale of this thesis. However, there

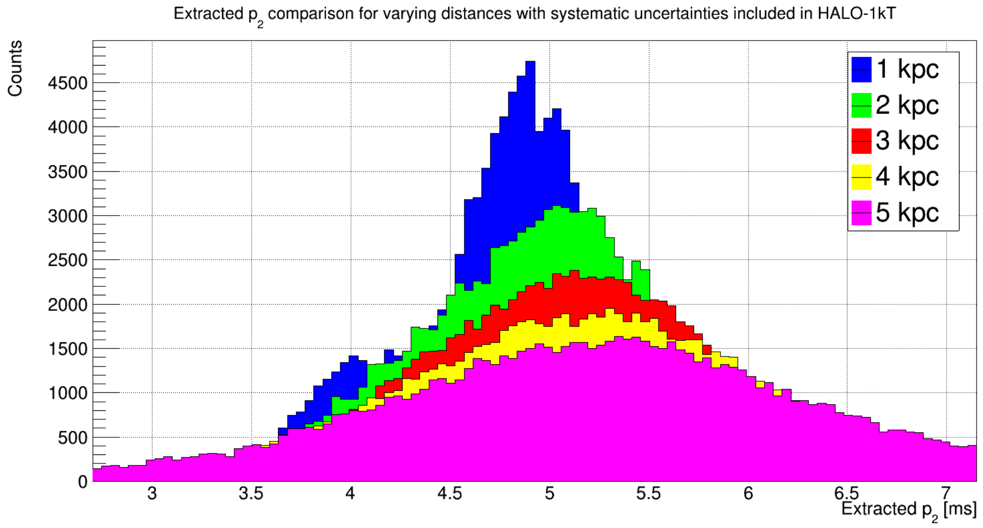

Figure 7.28: Distributions of extracted p_2 for varying distances, each with systematic uncertainties included, in HALO-1kT.

are still important lessons that can be learned from this model and addressing potential issues when applying the linear fit to Hüdepohl et al. 8.8 $M_{\odot}$ data will aid in any later potential integration of other models. As such, we carry out the same fitting procedure from the subchapters noted prior and plot the distribution of extracted parameters p_2 for the Hüdepohl et al. 8.8 $M_{\odot}$ model in Figure 7.29. What we immediately conclude is that the Hüdepohl et al. 8.8 $M_{\odot}$ models p_2 is systematically offset from the other models tested. This will pose potential issues, as when the procedure to address systematic errors is carried out, we yield the distribution of p_2 found in Figure 7.30. The distribution has two peaks, where the left peak is attributed to the Hüdepohl et al. 8.8 $M_{\odot}$ model and the tails of the four Sukhbold et al. models. When the offset correction is applied to our linear fit, we note that each peak shall be located on their respective side of the Monte Carlo truth t_0.

As we apply the linear fit to distances beyond 1 kpc, the observed trend of dual peaks in the extracted p_2 distribution will disappear. In Figure 7.31 the extracted p_2 with the

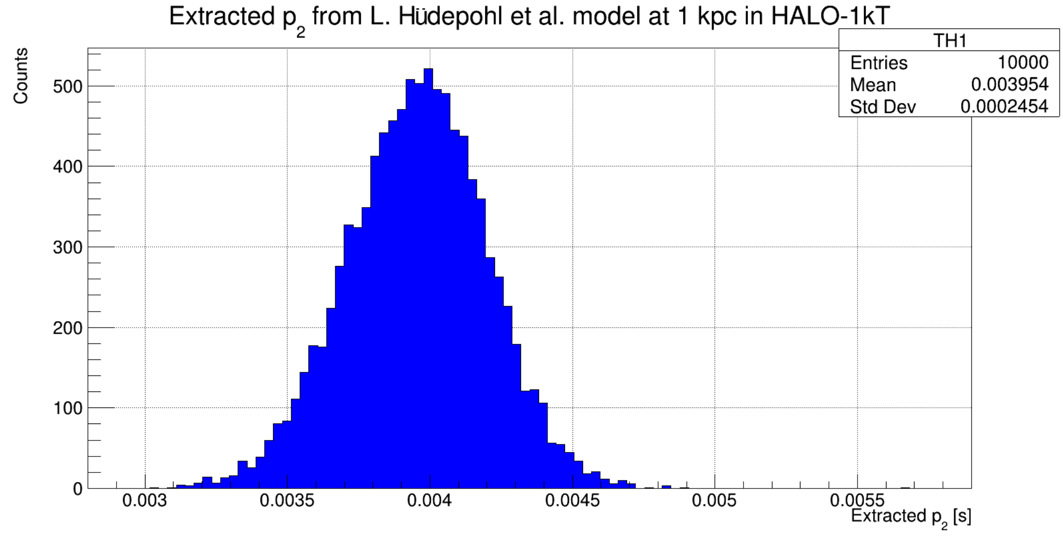

Figure 7.29: Distribution of extracted p_2 for the Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc in HALO-1kT.

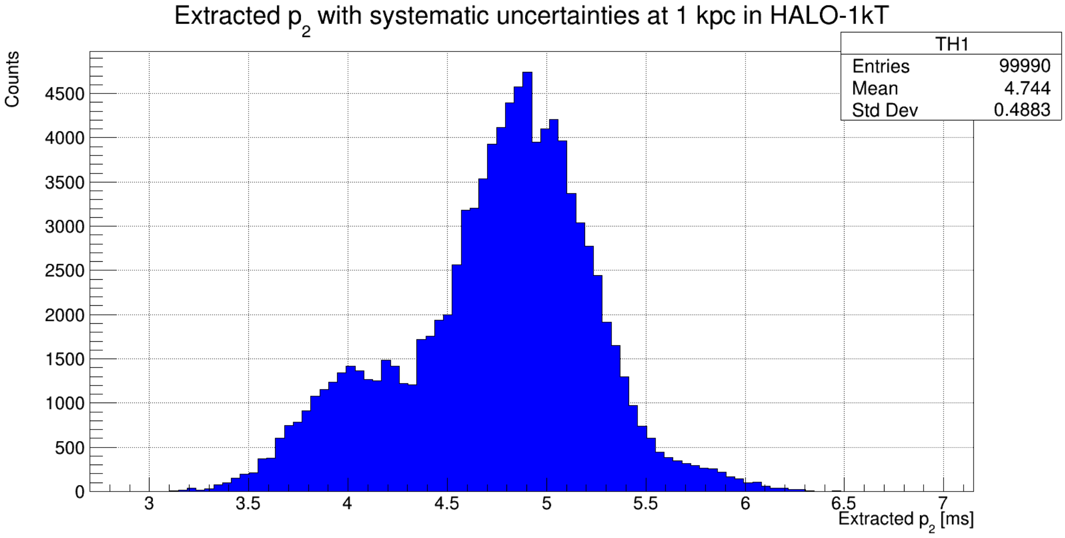

Figure 7.30: Distribution of extracted p_2 with both the Hüdepohl et al. 8.8 $M_\odot$ model and systematic uncertainties included at 1 kpc in HALO-1kT.

systematic uncertainties included for all five models is presented without the offset correction. The distribution is well described by a normal distribution and lacks the two distinct peaks observed at 1 kpc. This isn't solely attributed to the distributions of extracted p_2 smearing into one another. As the simulated SN signal distance is increased, the Hüdepohl et al. 8.8 $M_\odot$ model shifts back towards the central value of the other respective distributions. There is clearly a feature that is present in the Hüdepohl et al. 8.8 $M_\odot$ bursts at 1 kpc, but disappears later on. This feature is an artifact that arises based on how different authors choose to supply their data to `SNEWPY`.

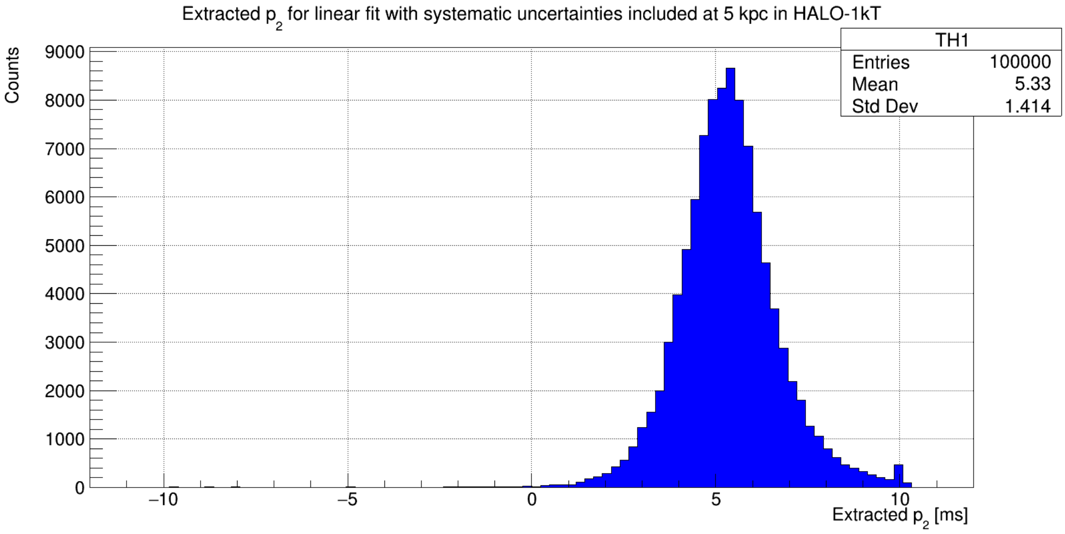

Figure 7.31: Depicted above is the extracted p_2 distribution at 5 kpc in HALO-1kT with both the Hüdepohl et al. 8.8 $M_\odot$ model and systematic uncertainties included.

Since this artifact is only present in the Hüdepohl et al. 8.8 $M_\odot$ data, we seek an additional cut that will bring the Hüdepohl et al. 8.8 $M_\odot$ model closer to the others, *without* impacting the Sukhbold et al. models. Recall that in the definition for the linear fit, we discussed how the shape of the mean light curve informed the initial parameter selection. In Figure 7.32, the Hüdepohl et al. 8.8 $M_\odot$ model has been included amongst the Sukhbold et al. models and their observed mean light curves. The overall shape of the leading edge is similar, however it differs in magnitude. To shift the overall mean light curve down, an early event cut can be applied. We it label as such, since the intent is to remove events from the early times prior to t_0 that cause the Hüdepohl et al. 8.8 $M_\odot$ model to differ by a constant factor early on in the burst. The best approach to apply this cut is to scan the observed time series and find where the maximum event density in a rolling 20 ms window is located. For the simulated distances in this report, this will always occur in the neutronization burst. The scan is only applied from the first event time plus 30 ms. The 30 ms upper bound is put in place to prevent the conditions of this early event cut to be satisfied further into the observed burst for the high mass models; as the high mass models maintain a high event rate out to 700 ms. Once the maximum event density is found along the burst, any events prior to 15 ms of the central value of the maximum event density window are cut. In Figure 7.33, the distribution of extracted p_2 from the Hüdepohl et al. 8.8 $M_\odot$ model an be found with the inclusion of the early event cut. The implementation of the early event cut had the intended effect as the central value shifts from 3.95 ms (Figure 7.29) to 4.32 ms (Figure 7.33), which brings it into better agreement with the other distributions. The precision is also increased from 490 μs to 230 μs. For the Hüdepohl et al. 8.8 $M_\odot$ model, the introduction of the early event cut has no trade-off and only serves to improve its location with respect to the other models. At 5 kpc, the effects of the early event cut go unobserved as the number of events in the leading edge is reduced.

The initial application of this early event cut was too aggressive and resulted in large quantities of the signal being cut from all models in our library. To ensure that in its current configuration the only model impacted was the Hüdepohl et al. 8.8 $M_\odot$ model, we shall look at the performance of extracted p_2 for the other models both prior to and after

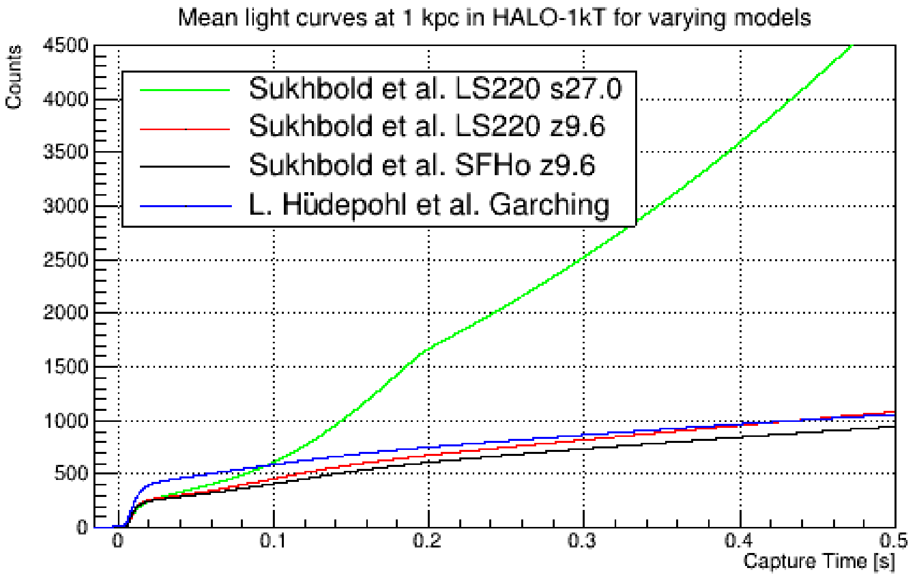

Figure 7.32: Mean light curves of all models at 1 kpc up to the first 500 ms post-bounce.

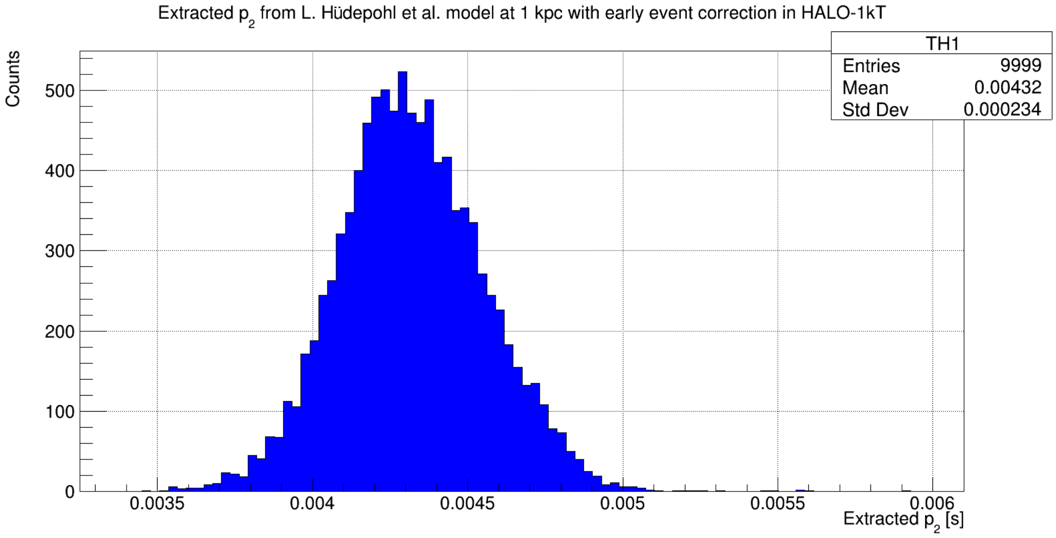

Figure 7.33: Distribution of extracted p_2 for the Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc with the early event cut applied to each observed time series.

the early event cut was applied. An example is shown in Figure 7.34 for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc.

Each of the models tested are unique in the effects that the early event cut imposes on the extracted p_2 distribution. The objective is to maintain or improve upon the degree of precision observed in all models prior to the early event cuts application. At closer distances, there is no statistically significant observed effect on the Sukhbold et al. models p_2 distributions. With the inclusion of the early event cut, the negative effect observed is the population of extracted p_2 that are in the left most tail of the distribution is increased (an increase in asymmetry). On the scale of the number of fits that were done, this effect is negligible and we justify ignoring these minor deficiencies due to the relative increase in performance the Hüdepohl et al. 8.8 $M_\odot$ model sees via the introduction of the early event cut. The Sukhbold et al. LS220 s27.0 $M_\odot$ model saw general improvements with the introduction of the early event cut, further justifying keeping its implementation in place.

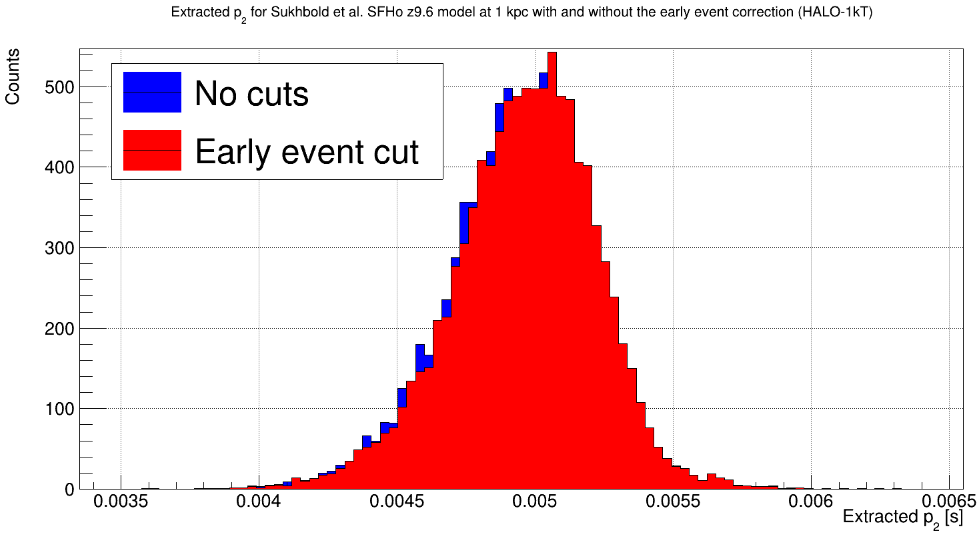

Figure 7.34: Comparison of the effect of the early event cut on the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT.

7.5 Kolmogorov-Smirnov and Anderson-Darling tests

7.5.1 Kolmogorov-Smirnov test

7.5.1.1 Motivation

The Kolmogorov-Smirnov (KS) test is a statistical method used to test if a specific sample belongs to a parent distribution. The test can be configured to determine the maximum deviation between the cumulative time series and the parent CDF, or the probability that the sample belongs to the parent distribution. Only the maximum deviation is of interest here. Early on in the implementation of the NLL technique, there were numerous issues with systematic offsets at distances approaching 10 kpc. These offsets were present prior to mixing models and PDFs. It was believed that the KS test could be implemented to observe if any correlation existed between the offset of $\hat{t_0}$, and the value returned from the KS test. An analysis of the two quantities yielded no correlation between the maximum deviation obtained from the KS test and the NLL estimate of t_0. It was clear that the maximum deviation played no role in the extraction of $\hat{t_0}$ from the NLL technique. The KS test was shelved and any additional implementation was paused. It wasn't until our group thought about reviving the technique as a dedicated extraction technique, that the KS test saw further exploration. This technique, like the others, makes use of unbinned data, once again preserving the temporal resolution needed to effectively communicate t_0 to SNEWS.

7.5.1.2 General formulation

The KS test is defined as [101]

$$D(t_0) = \max_{1 \leq i \leq N} \left(\mathcal{N} \cdot \bar{n}(t_0 - t_i) - \frac{i-1}{N}, \frac{i}{N} - \mathcal{N} \cdot \bar{n}(t_0 - t_i) \right) \tag{7.12}$$

where, as defined in the NLL technique, $\bar{n}(t_0 - t_i)$ is the mean cumulative events, N is the number of events in the burst, $D(t_0)$ is the maximum deviation, and $\mathcal{N}$ is the normalization to unity as defined by

$$\mathcal{N} = \frac{1}{\bar{n}(t_e)} \tag{7.13}$$

where t_e is the SN model's end time. The KS test is fully integrated into the `ROOT::TMath` namespace; we used the KS test in a custom `RooFit` function configured to handle the offset. It is expected that the the minimum of the maximum deviation function will correspond to our true t_0. As such, MINUIT is once again utilized to locate and extract $\hat{t_0}$.

7.5.1.3 Issues arising with the fit

If we take a Sukhbold et al. SFHo z9.6 $M_\odot$ time series, simulated in HALO-1kT at 1 kpc, and form the maximum deviation function using the Sukhbold et al. SFHo z9.6 $M_\odot$ mean light curve ($\bar{n}$), with the result shown in Figure 7.35. No time window has been applied to this fit and the entire burst is used. Around the true t_0 value, the behaviour of the function is disjointed. This feature arises from using a discrete time series and a "continuous" PDF. Labelling the PDF as continuous is a slight misnomer. It is in fact made up of 10^4 discrete simulated time series (as discussed in §7.2). Any interpolation required by `RooFit` is so infinitesimal (especially around the neutronization burst) that discrete features do not present themselves in the fit. By introducing a normalization to the mean light curve, both the observed time series and mean light curve will be fixed to the same stationary initial and end points. Without the normalization, the maximum deviation will occur towards the end of the burst, as the Poisson fluctuations cause a ballooning effect away from the mean light curve at later times. With the normalization applied, the maximum deviation will instead occur where there is greater variability in the expected number of events, which occurs during the neutronization burst. This is especially true for high statistics observations, which occur at distances around 1 kpc. At this point, around the neutronization burst, the maximum deviation between the mean light curve and time series has a linear response as the mean light curve is shifted laterally. This is attributed to the bin width of the observed time series, which corresponds to the time difference between two events. At the minimum maximum deviation, the time scale of the time difference between two events in the neutronization burst is sufficiently small that its distance to the mean light curve is linearly related to its horizontal translation.

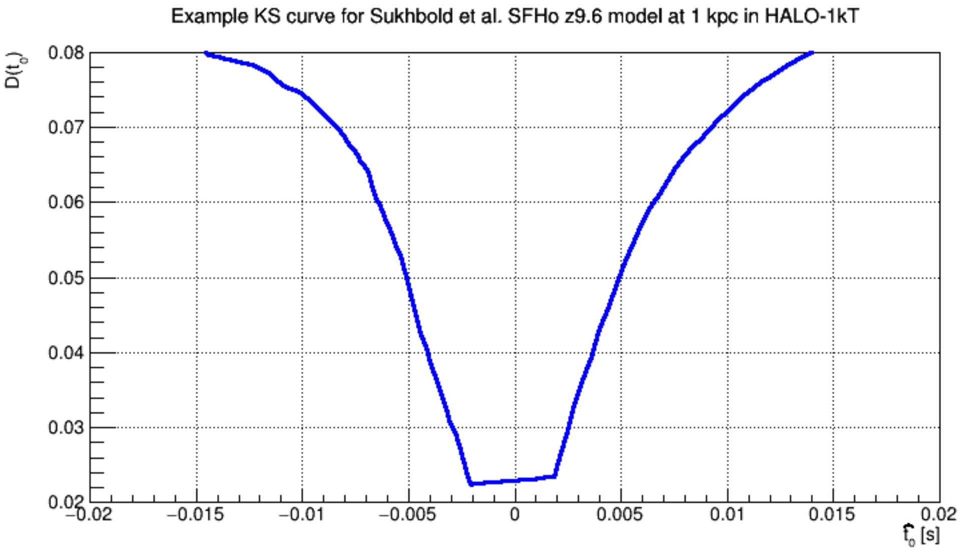

Figure 7.35: The KS test parameterized as a function of the offset at 1 kpc in HALO-1kT for the Sukhbold et al. SFHo z9.6 $M_\odot$ model.

Owing to the shape of this curve around the true t_0, MINUIT has a difficult time locating the minimum. As such we turn to a custom method of finding the minimum and by consequence, t_0. If we plot the distribution of t_0 by simply parsing until the minimum of the function is found, we obtain the results in Figure 7.36. The two peaks are attributed to the slope of the linear region around t_0. A positive slope corresponds to the peak on the left and a negative slope to the peak on the right.

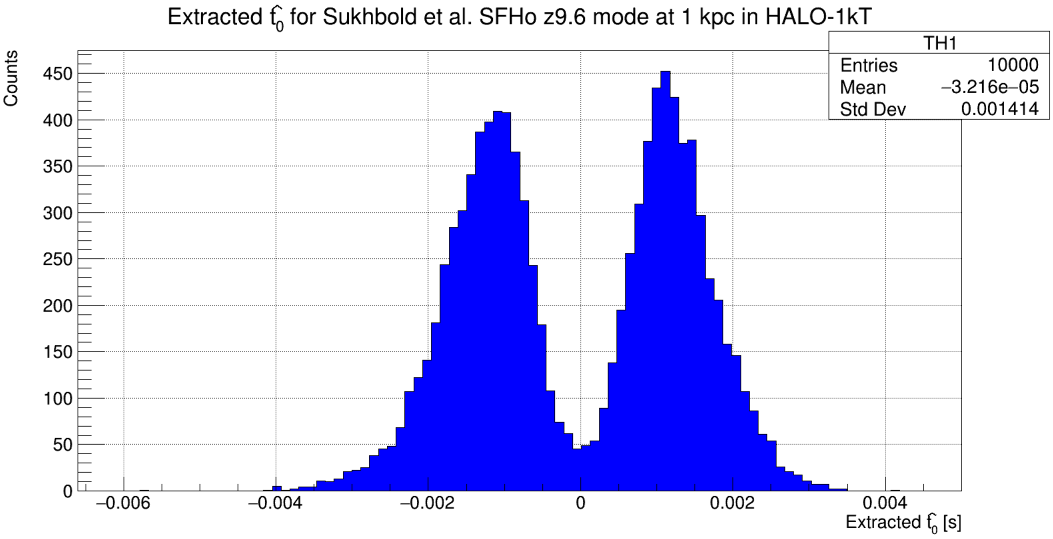

Figure 7.36: The extracted $\hat{t_0}$ from the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT. These $\hat{t_0}$ were obtained via parsing the maximum deviation function for the minimum, with no correction applied.

We exploit the fact that this curve is parabolic on each side of the linear region by finding the x values that correspond to $1.5 \cdot D(t_{\text{min}})$, where t_{min} is the parsed minimum of the function. We label these the upper and lower symmetry points, indicating which side of the minimum they are on. From here, the corrected $\hat{t_0}$ can be obtained by averaging the two time values that correspond to $1.5 \cdot D(t_{\text{min}})$. The distribution of the corrected $\hat{t_0}$ can be found in Figure 7.37 for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT. This is offset from the true t_0 value by 550 μs and has a standard deviation of 360 μs. In contrast to the NLL method, this technique fails to extract t_0 with the same precision.

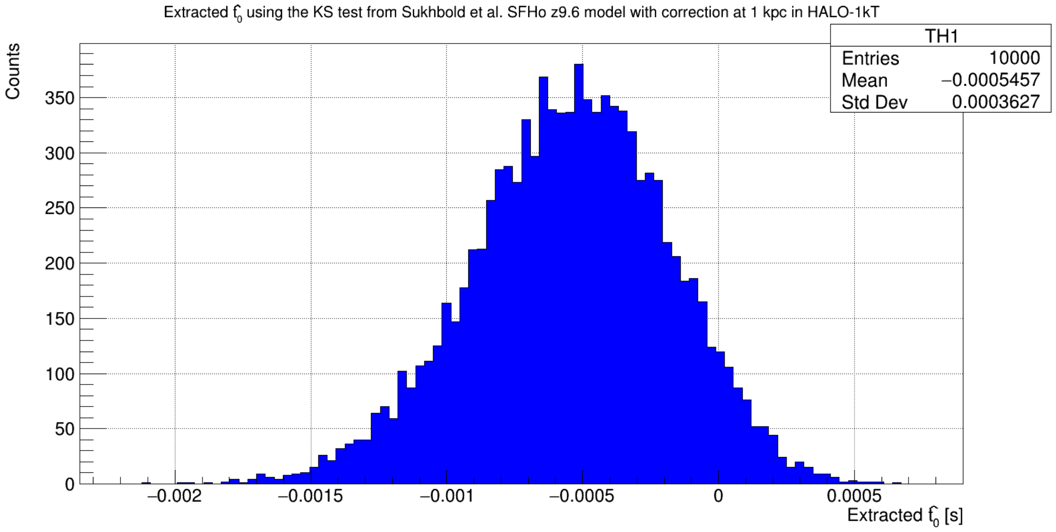

Figure 7.37: The extracted $\hat{t_0}$ from the Sukhbold et al. SFHo z9.6 model at 1 kpc for HALO-1kT. These $\hat{t_0}$ were obtained via parsing the maximum deviation function for the minimum, with corrections applied as described in the text.

7.5.1.4 Motivation to abandon

The results presented for the KS test to this point are not useless. It is entirely possible that as this analysis is expanded to encompass additional next-generation neutrino experiments, the KS test could be uniquely suited for a specific detector. With how different the detectors in SNEWS can be, the KS test could fare better in another experiment. Only by carrying out similar studies to this thesis can we determine if the KS test is a good fit for a given neutrino telescope. As the KS test was rolled out across all simulated models at larger distances, the maximum deviation function began to lose critical features, as the uncertainties increased, that allowed for a precise estimate of t_0. An example is shown in Figure 7.38. With the lack of statistics, there is too much variability in the maximum deviation as a function of offset. Attempts were made to salvage this technique by applying a time window to the observed time series, akin to the NLL technique. This time window would also need to be applied to the PDF, as the time series for this fit should be of the same temporal length. There is tremendous variability in the first event times even at distances of 1 kpc, as seen in Figure 4.5, so there is no method to consistently predict where the first event lies on the PDF, potentially inducing systematic errors in our fitting procedure. Unfortunately, no matter how the time window was applied to the PDF, the KS test could not extract $\hat{t_0}$ at these distances. As such it was shelved once more.

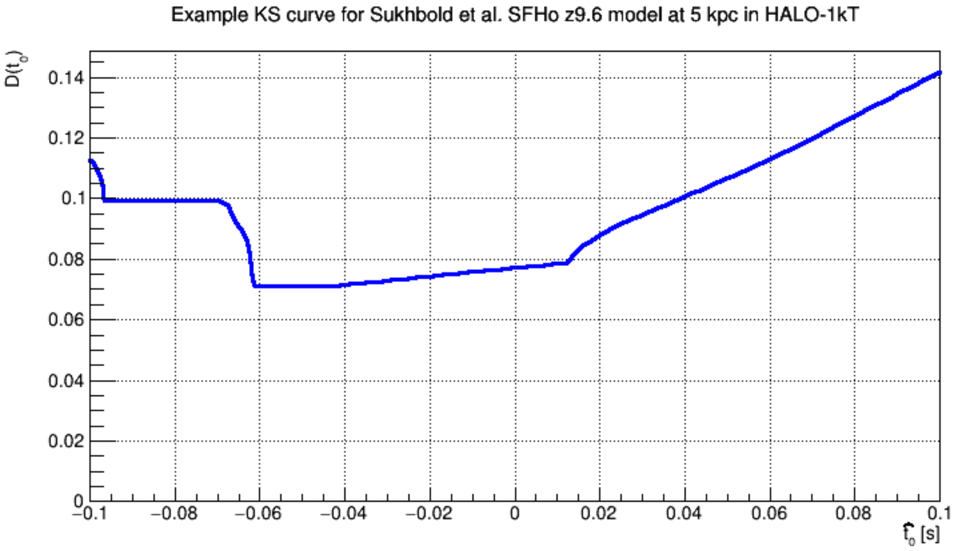

Figure 7.38: An example of the maximum deviation function at 5 kpc for the Sukhbold et al. SFHo z9.6 model in HALO-1kT.

7.5.2 Anderson-Darling test

7.5.2.1 Motivation

With the KS test able to extract $\hat{t_0}$ with ease at 1 kpc but not at greater distances, the intent was to find an alternative test that was competitive beyond 1 kpc. The Anderson-Darling (AD) test was chosen as it is an extension of the KS test, giving more weight to the tails. Like the KS test, its intent is to assess the probability that a sample belongs to some parent distribution. The AD test also avoids binning the signal, preserving the time series temporal information. As will be demonstrated in the subchapters that follow, transitioning to this technique added an additional competitive technique to our library of analytic techniques to extract $\hat{t_0}$.

7.5.2.2 General formulation

The AD test can be written as [102]

$$A^2(t_0) = -N - S(t_0) \tag{7.14}$$

$$S(t_0) = \sum_{i=1}^{N} \frac{(2i-1)}{N} \Big[\ln \bar{n}(t_i - t_0) + \ln \big(1 - \bar{n}(t_{N+1-i} - t_0)\big) \Big] \tag{7.15}$$

where A is the AD test statistic, N is the number of events in the SN burst, and $\bar{n}$ is the mean light curve. Since the AD test makes use of the normalized cumulative event curve, we introduce $\mathcal{N}$, the same normalization defined in Equation 7.8.

$$S(t_0) = \sum_{i=1}^{N} \frac{(2i-1)}{N} \Big[\ln(\mathcal{N} \cdot \bar{n}(t_i - t_0)) + \ln \big(1 - \mathcal{N} \cdot \bar{n}(t_{N+1-i} - t_0)\big) \Big] \tag{7.16}$$

We do not make use of the test statistic for hypothesis testing per se. Instead we parameterize it as a function of t_0 such that $A^2(t_0)$ can be minimized and t_0 can be estimated.

7.5.2.3 Modifications

The implementation of the this technique will require modifications in the same vein as the NLL fit. Since the fit is expressed in logarithmic terms, singularities will arise if either condition

$$\bar{n}(t_i - t_0) \leq 0 \tag{7.17}$$

$$\bar{n}(t_{N+1-i} - t_0) \geq 1 \tag{7.18}$$

is satisfied. The correction applied in §7.2.4.2 is applied here, avoiding the situations where the first condition might arise. The second condition arises under the following two situations

- **Mixing of models** - Each SN model integrated into `SNEWPY` is of varying length with respect to the neutrino flux. As outlined previously, the variation between two models can reach up to ≈ 6 seconds. This temporal variance will result in a failed AD test. Suppose a model with a final event at time t_{N_1} is fit to a PDF with a final event time t_{N_2}, where $t_{N_1} \geq t_{N_2}$. In this scenario, the CDF is not defined in these regions. These issues can be rectified if a time window is applied, however, new issues will present themselves.

- **Application of time window** - When a time window is applied to the SN burst and only the first 10s of milliseconds are used, the fit will fail unless the same time window is applied to the mean light curve. Similar to the KS test, the AD test expects that the two samples are of equal temporal length. To apply a time window to the mean light curve, we normalize the point along the curve, at the time window length, to unity. Since we chose the Monte Carlo truth time along the mean light curve to apply our cut, if the first event is after t_0, the second condition in Equation 7.17 will always occur. One might suggest that the time window applied to the Monte Carlo truth times of the mean light curve should be extended on the order of ms to prevent singularities from arising. By increasing the time window applied to the mean light curve we introduce systematic offsets of our t_0 distribution. The same behaviour is observed in the NLL technique. The magnitude of the introduced offset is highly model and PDF dependant. If the time window is applied without the normalization condition the fit will fail to reach a satisfactory minimum. This is not surprising, as without the normalization condition the mean light curve will not equate to unity at the end of the burst. We then encounter the issue of temporal length once more.

If either of the aforementioned singularities occur when processing the SN burst, the function is immediately incremented by 10^4 to avoid an arithmetic fault. The impact of this increment shall be discussed in §7.5.2.8. As a summary, a time window will once again be applied to the observed time series. A time window of the same length will also be applied to the mean light curve using the Monte Carlo simulation truth time.

7.5.2.4 Extraction of t_0

In keeping with the format outlined in the previous subchapters, we apply the Anderson-Darling test to the Hüdepohl et al. 8.8 $M_\odot$ model at 1 kpc in HALO-1kT. For a single simulated burst fit using Equation 7.14, using the Hüdepohl et al. 8.8 $M_\odot$ model mean light curve, the resulting $A^2(t_0)$ can be found in Figure 7.39. The linear region of the fit encountered in the KS test is not present in the AD test, thus no custom minimum solution is required and MINUIT can once again be utilized to locate the function minimum.

We apply the same fitting process to 10^4 of the Sukhbold et al. SFHo z9.6 $M_\odot$ bursts in HALO-1kT and obtain the distribution of extracted $\hat{t_0}$ in Figure 7.40. The results

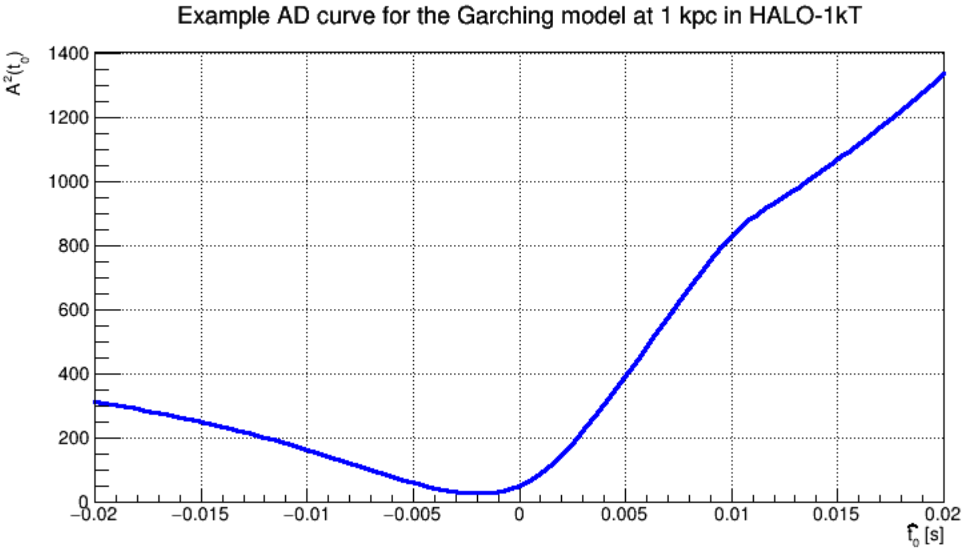

Figure 7.39: Example of the AD test statistic curve at 1 kpc in HALO-1kT for the Hüdepohl et al. 8.8 $M_\odot$ model.

here are excellent, with a mean extracted $\hat{t_0}$ of 19 μs. Moreover, the results are dramatically improved over the KS test. The systematic offset observed in Figure 7.37 is largely eliminated and the standard deviation has been narrowed down to 190 μs. This fitting method does possess some numerical issues, mainly the bias towards the left side of the distribution, which can barely be identified in Figure 7.40. These numerical issues are model dependant and discussion surrounding them shall be tabled until §7.5.2.8.

7.5.2.5 Mixing models and PDFs

To estimate the systematic uncertainties present when fitting an observed time series to our library of models, we once again must investigate the model dependence. We elect to perform the AD test for Sukhbold et al. SFHo z9.6 $M_\odot$ time series data with both the Sukhbold et al. LS220 z9.6 $M_\odot$ PDF and the Sukhbold et al. LS220 s27.0 $M_\odot$ PDF. The former distribution of extracted $\hat{t_0}$ can be found in Figure 7.41 and the latter distribution of extracted $\hat{t_0}$ is in Figure 7.42. Both of these distributions are unique amongst

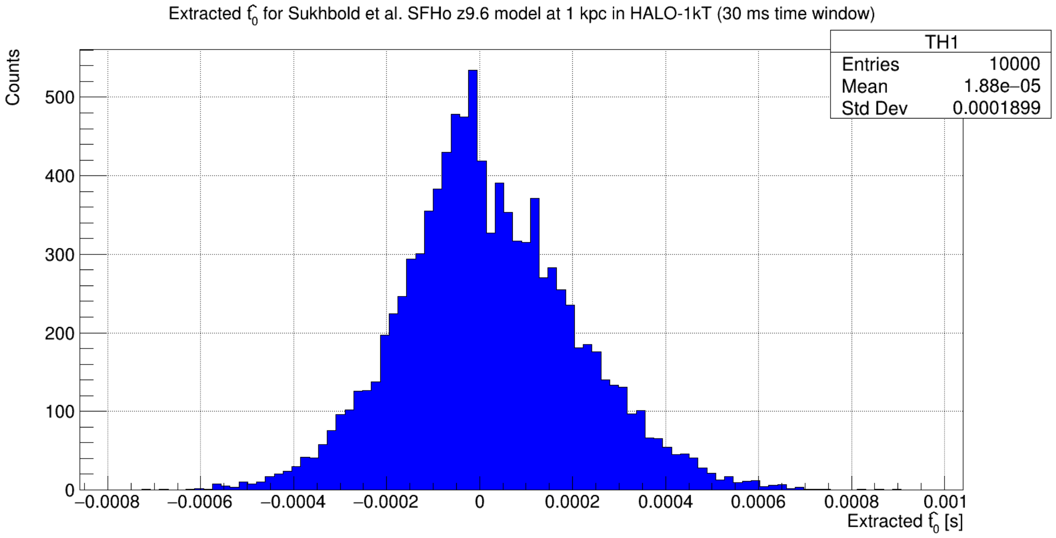

Figure 7.40: Distribution of extracted $\hat{t_0}$ from the Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT using the AD test.

the techniques tested thus far as they are asymmetric and non-Gaussian. In Figure 7.41, the right shoulder of the distribution is primarily attributed to $A^2(t_0)$ curves that successfully reach their minimum prior to an invalid t_0 being reached as per the conditions outlined in Equation 7.17. The extracted $\hat{t_0}$'s to the left of this shoulder are attributed to both $A^2(t_0)$ that successfully reach their minimum prior to satisfying the outlined conditions and those $A^2(t_0)$ that don't. The former is the source for the underlying Gaussian characteristics in the distribution while the latter gives rise to the linear decline towards negatively offset $\hat{t_0}$. In the case of Figure 7.42, the right shoulder is heavily diluted with $A^2(t_0)$ that fail to reach their minimum, but the underlying Gaussian distribution that is expected, with a systematic offset, is also present. Its relative location is offset around -700 μs from t_0. Even though the function fails to reach a minimum in these cases, we can still use it to estimate the Monte Carlo truth t_0, as such, we do not discard the extracted $\hat{t_0}$ if the minimum is not reached prior to the singularities. Techniques were explored to potentially correct the issues present in the AD test and are explored in §7.5.2.8.

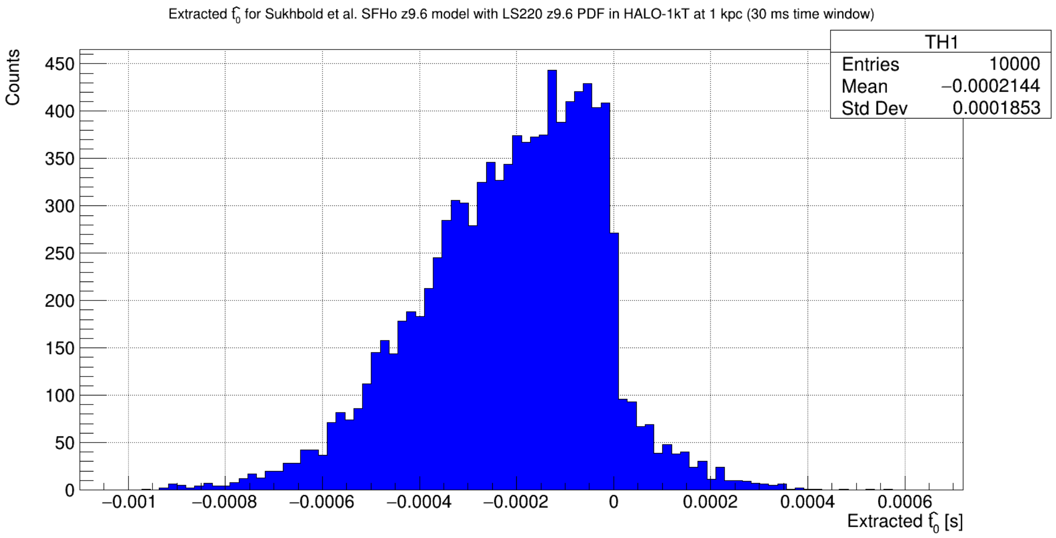

Figure 7.41: Distribution of extracted $\hat{t_0}$ for the Sukhbold et al. SFHo z9.6 model tested against the Sukhbold et al. LS220 z9.6 PDF (HALO-1kT, 1 kpc).

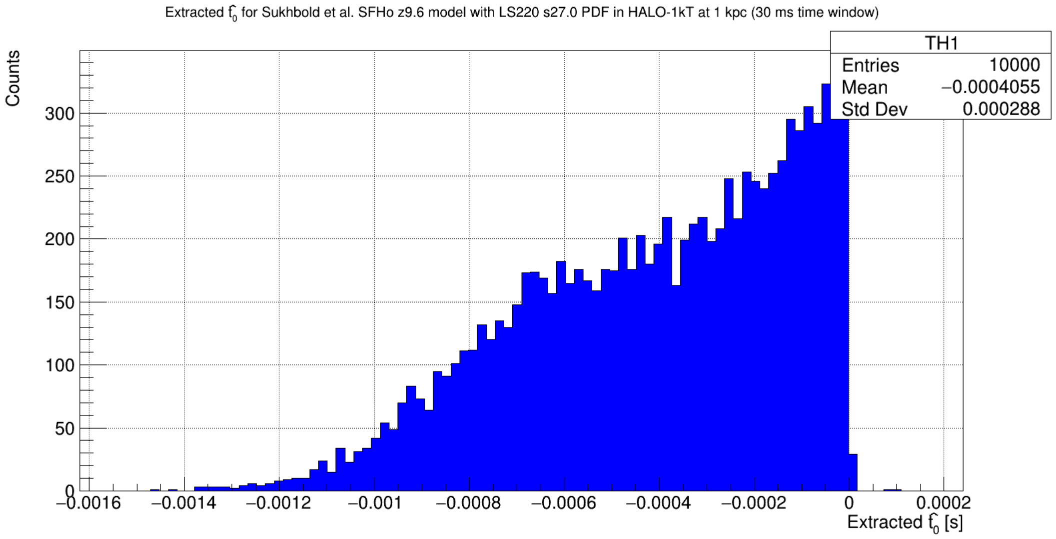

Figure 7.42: Distribution of extracted $\hat{t_0}$ for the Sukhbold et al. SFHo z9.6 model tested against the Sukhbold et al. LS220 s27.0 PDF (HALO-1kT, 1 kpc).

7.5.2.6 Estimating systematic uncertainties

We once again follow the same procedure to account for the systematic uncertainties as was done for the NLL fit, since both techniques share the dependence on a choice in PDF. We treat each of the extracted $\hat{t_0}$ histograms obtained from mixing models and PDFs with equal weight and sample them 10^4 times. The results can be found in Figure 7.43. We observe the same behaviour noted previously with respect to the shoulder in the fits that terminate before reaching a minimum. The pronounced tails on the distribution are attributed to the issues with finding a minimum prior to termination; this issue is not unique to a single model or PDF combination, and is intrinsic to how we have implemented this technique. Overall performance is competitive with NLL technique but is edged out by the CFD method. The reader is directed to §8 for additional discussion of the comparison between techniques.

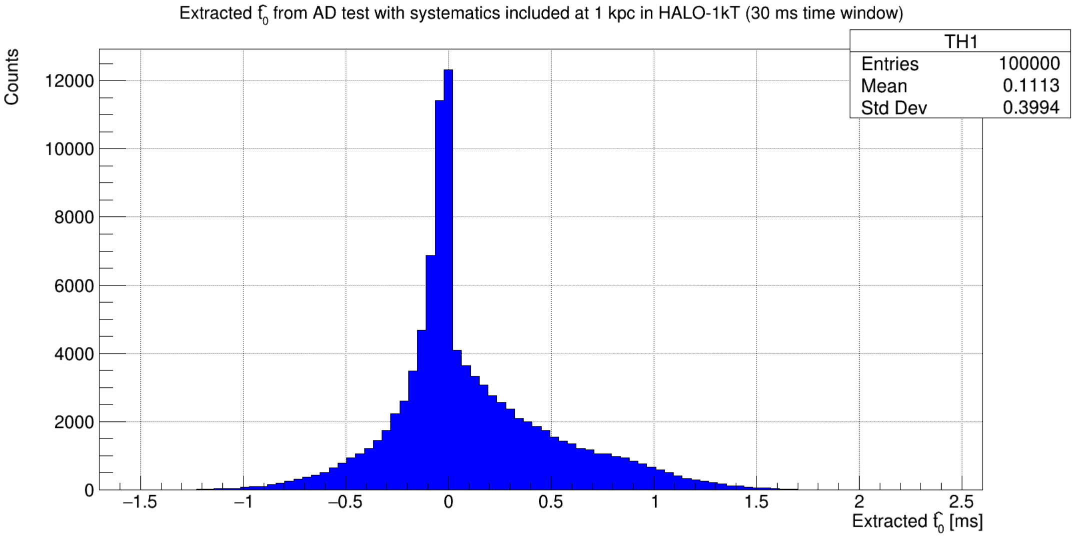

Figure 7.43: Extracted $\hat{t_0}$ for the AD test with systematic uncertainties included at 1 kpc in HALO-1kT.

7.5.2.7 Performance as a function of distance

Before we delve into a brief discussion on the performance of each technique as a function of distance with systematic uncertainties included, we will first turn our attention to issues that presented themselves as a result of the conditions outlined in Equation 7.17. As the supernova simulation distance is increased, the number of observed events falls off as the inverse square law, and issues encountered when mixing models and PDFs disappear. In HALO-1kT the distortion on the extracted $\hat{t_0}$ distribution is entirely gone by the time a simulated distance of 5 kpc is reached. Depicted in Figure 7.44 is the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 5 kpc fit to the Sukhbold et al. LS220 s27.0 $M_\odot$ PDF with a 30 ms time window. The 5 kpc distribution of the extracted $\hat{t_0}$ is much more symmetric, however there is the anticipated increase in uncertainty. The source of the termination issues, and their subsequent disappearance, is the event density in the early signal and neutronization burst. Additional details can be found in the following subchapter. The primary motivation of briefly addressing the return to symmetry here is to assure the reader this is an expected result and not an issue in technique.

We now turn our attention back to the discussion of the performance as a function of distance. Depicted in Figure 7.45 is the performance of the AD test in extracting $\hat{t_0}$ for distances from 1 to 5 kpc. Since the issue of $A^2(t_0)$ terminating (reaching an undefined state) prior to reaching the true minimum (i.e. the minimum if the function did not terminate prior to reaching this $\hat{t_0}$ value) at larger distances is reduced, its effect on the overall extracted $\hat{t_0}$ distribution with systematic uncertainties included is heavily reduced. This technique on its own maintains a sub-millisecond precision out to 3 kpc, breaching this threshold at 4 kpc with a offset of 120 μs and standard deviation of 1.09 ms. We

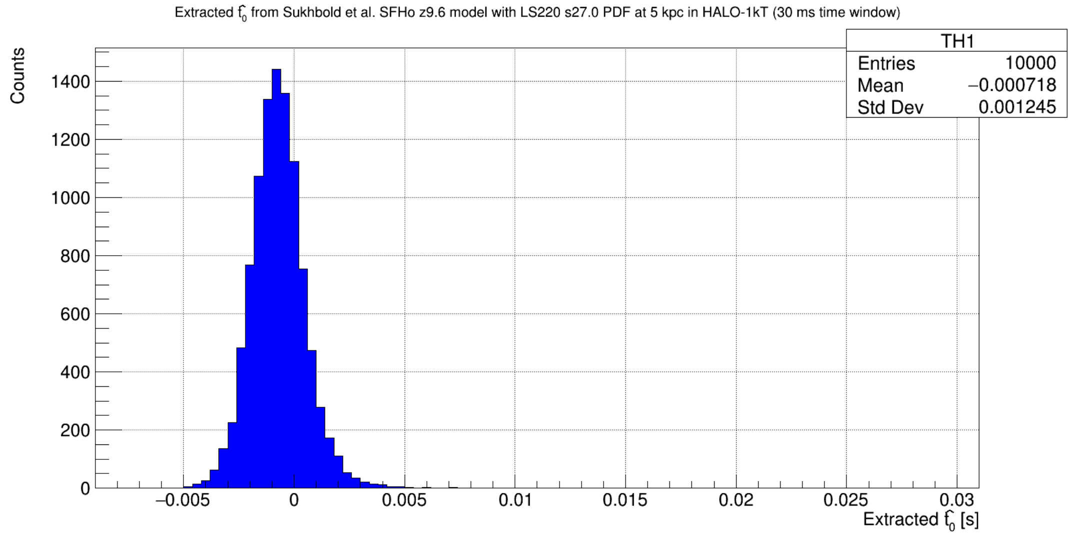

Figure 7.44: Extracted $\hat{t_0}$ from the AD test using Sukhbold et al. SFHo z9.6 time series data tested against the Sukhbold et al. LS220 s27.0 PDF at 5 kpc in HALO-1kT.

note that the walk present in all techniques thus far is evident once more. The difference in the offset between the initial distribution and final is $\approx 50\ \mu s$ (from Figure 7.45).

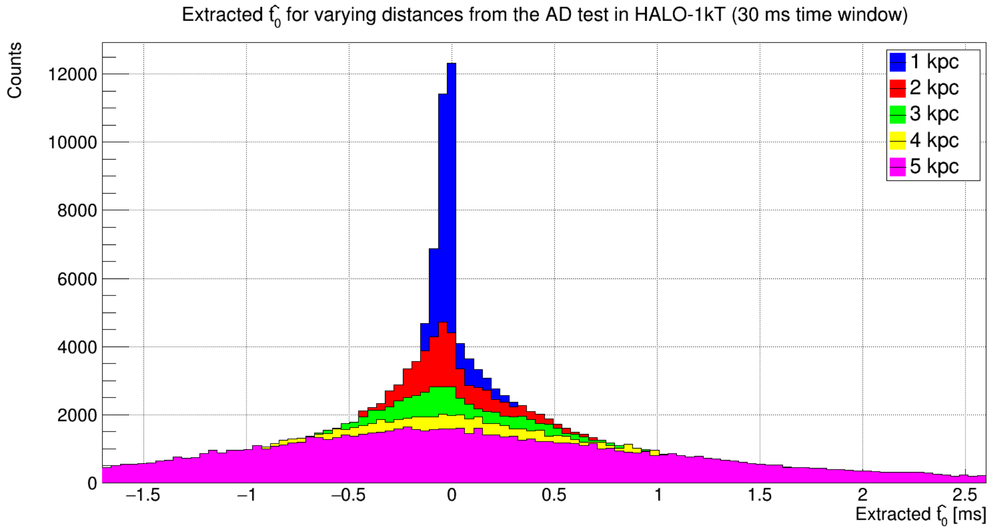

Figure 7.45: Distribution of extracted $\hat{t_0}$ using the AD test with systematic uncertainties included at varying distances.

7.5.2.8 Brief aside on Anderson-Darling minimum issues

As alluded to in previous subchapters, the AD test experiences issues with events early on in the signal, analogous to the linear fit, as well as with events in the neutronization burst. Once again this issue is evident with greater severity when making use of Hüdepohl et al. 8.8 $M_{\odot}$ time series data. To demonstrate this, we apply the Anderson-Darling test to the Hüdepohl et al. 8.8 $M_{\odot}$ data with a 30 millisecond window, using its own mean light curve at 1 kpc. The results can be found in Figure 7.46. Similar features that were present when intermixing model and PDF are present here, however, this is fitting Hüdepohl et al. 8.8 $M_{\odot}$ data to itself. Other models do encounter the issues similar to those observed here, when fitting them to themselves, but the severity is drastically reduced.

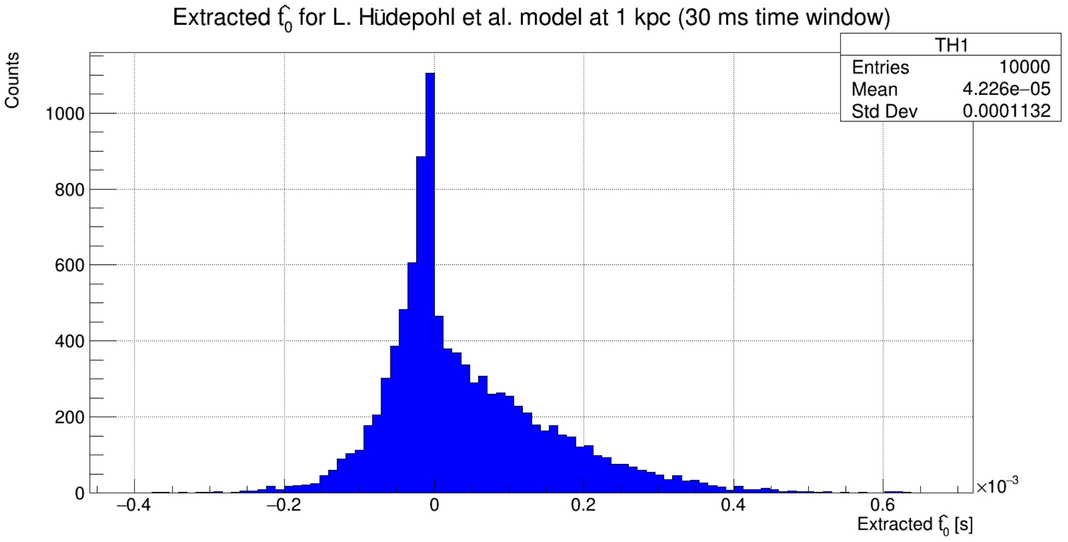

Figure 7.46: Distribution of extracted $\hat{t_0}$ for the Hüdepohl et al. 8.8 $M_{\odot}$ model, making use of the AD test and a 30 ms time window. The results above are for 1 kpc in HALO-1kT.

Corrections were explored to address the termination of $A^2(t_0)$ at certain t_0 values. The first method to correct the failure to reach a minimum prior to termination was to supply an "early event cut" in the same vein as the linear fit. Making use of the same configuration for early event cut, we scan in a rolling 20 ms window for the maximum event density. Once located, we cut any data from the time series that occurs 15 ms prior to the central value of the window with the peak event density. The results can be found in Figure 7.47, which is the extracted $\hat{t_0}$ from the Hüdepohl et al. 8.8 $M_\odot$ model fit to itself at 1 kpc in HALO-1kT. The result of this cut yields a distribution of $\hat{t_0}$ that is much more symmetric, however it comes at a steep price. The distribution is systematically offset from the Monte Carlo t_0 by 245.3 μs. This is of little surprise, as the mean light curve for the Hüdepohl et al. 8.8 $M_\odot$ model retains the early events that have been cut from the observed time series. In an attempt to rectify this, the exact same cut can be supplied to the PDF, which we label the mean light curve correction. It is once again determined by scanning the mean light curve and finding the maximum event rate in a rolling 20 ms windows. We then fix the point 15 ms prior to this to $y = 0$, effectively shifting the mean light curve down by some arbitrary value. This correction proved detrimental, as the AD test failed to extract $\hat{t_0}$ with an adequate degree of consistency. Regardless of how the PDF is "corrected", the introduction of systematic uncertainties and increased uncertainty due to the early event cut resulted in its use being dismissed for the AD test.

7.6 Cross-correlation

7.6.1 Motivation

At the outset of this research project, aspects of this work had been deemed essential (SN bursts should be passed through a full detector Monte Carlo, the SN signal should be Poisson fluctuated to account for statistical variability in the observed signal, etc.).

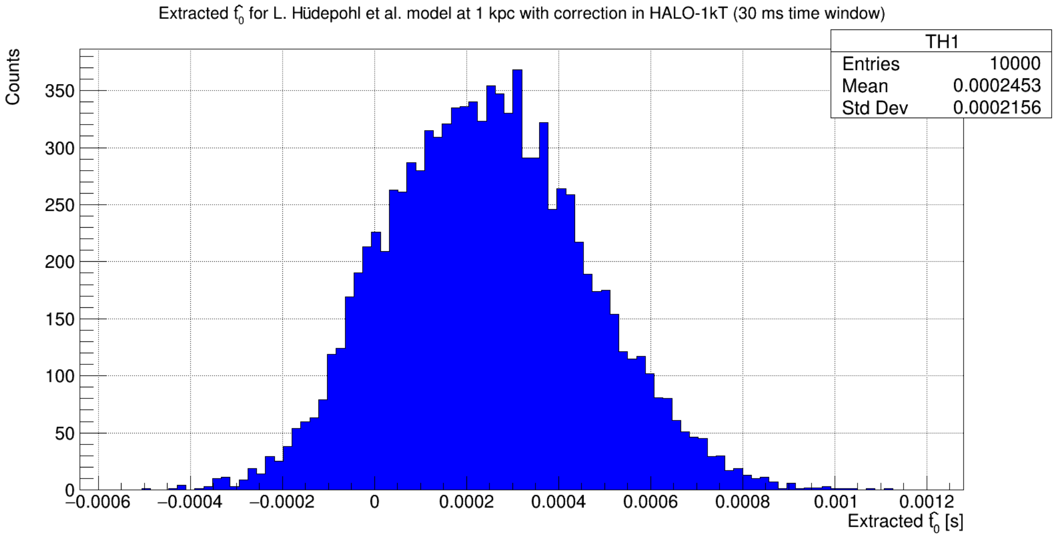

Figure 7.47: Distribution of extracted $\hat{t_0}$ for the Hüdepohl et al. 8.8 $M_\odot$ model, making use of the AD test, a 30 ms time window, and the early event cut. The results above are for 1 kpc in HALO-1kT.

The exact technique to be used to extract t_0 was largely undecided. As discussed in the §7 preamble, the criterion we deemed necessary was to avoid binning at all cost. It was believed that by including binning into our techniques, we limit the precision to which we can communicate $\hat{t_0}$ to SNEWS. As the the cross-correlation technique was implemented, results for the Hüdepohl et al. 8.8 $M_\odot$ model at 10 kpc in HALO-1kT were remarkably inconsistent and introduced systematic uncertainties that were difficult to reconcile. The likelihood technique by comparison, proved to be much more consistent at providing an accurate estimate of t_0. Yet the extracted $\hat{t_0}$ proved inadequate if HALO-1kT wished to contribute to the triangulation effort, as the uncertainty on $\hat{t_0}$ exceeded the millisecond precision needed to effectively triangulate. The hopes of attaining an estimate of t_0 at 10 kpc were dashed, as it became clear that results obtained in [6, 52] would not be possible with the event statistics in HALO-1kT for a SN at 10 kpc. We deemed it of more interest to shift our focus to distances ranging from 1 through 5 kpc. At the same time, the cross-correlation technique was shelved. As this project has come to its conclusion, and the intended expansion of this work to include many of the next-generation neutrino experiments became of interest, the need for a binned technique proved necessary. IceCube will not be able to communicate a time series of data to SNEWS [103]. Instead, it shall only provide a histogram of event times in discrete bins. For this reasoning I revisited the cross-correlation technique.

7.6.2 General formulation

Cross-correlation (CC) is the measure of similarity between two data series and is defined analytically as

$$(H_n * H_{\bar{n}})(t - t_0) = -\sum_{t=-\infty}^{+\infty} H_n(t) H_{\bar{n}}(t - t_0) \tag{7.19}$$

where H_n is the observed time series and $H_{\bar{n}}(t - t_0)$ is the mean event rate evaluated at some offset t_0, both of which are stored in 1D histograms. The event rate is formed by populating a histogram (from -1 to 20 seconds, with 100 μs wide bins) with the time series of events from all 10^4 SN bursts simulated in the detector for a given model. The event rate histogram is then normalized such that its integral is equal to one. The mean event rate for the Sukhbold et al. SFHo z9.6 $M_{\odot}$ model at 1 kpc in HALO-1kT can be found in Figure 7.48. Since the Monte Carlo truth of t_0 is zero seconds, when the simulated data is cross-correlated with the mean event rate, Equation 7.19 should reach its minimum at $t_0 = 0$.

When storing the time series in a histogram, the data is placed into 1 ms bins from 100 ms prior to the first event and 20.9 seconds after. Although no simulated model in this work extends out to 20 seconds, this was purely proactive in the event that additional

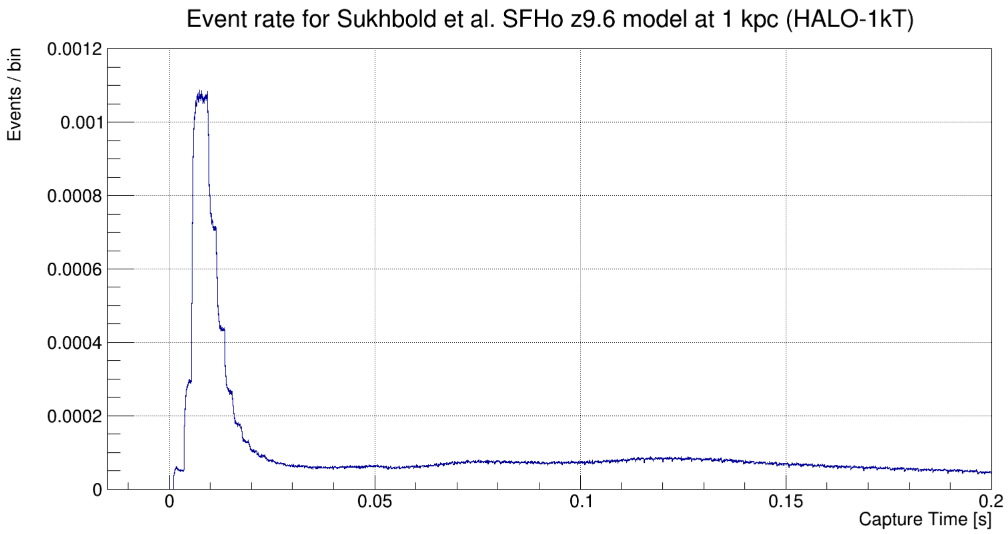

Figure 7.48: Observed mean event rate for the Sukhbold et al. SFHo z9.6 model in HALO-1kT at 1 kpc.

models are incorporated later and have bursts of this length. The choice of initial bin location is to remove artifacting present in the $\hat{t_0}$ distribution that will be described in §7.6.4.

7.6.3 Modifications

7.6.3.1 Average between two minima

The CC technique can be directly applied to an observed SN neutrino burst with few modifications necessary. For a SN neutrino burst from the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc in HALO-1kT, application of the CC technique, with a Sukhbold et al. SFHo z9.6 $M_\odot$ PDF, yields the $(H_n * H_{\bar{n}})$ curve found in Figure 7.49. Since the time series and mean event rate histograms take discrete values in finite bins, the step function behaviour observed in Figure 7.49 is anticipated.

The absence of a smooth behaviour means that if MINUIT is used, a minimum will not be found with a satisfactory consistency. There are a few techniques that can be

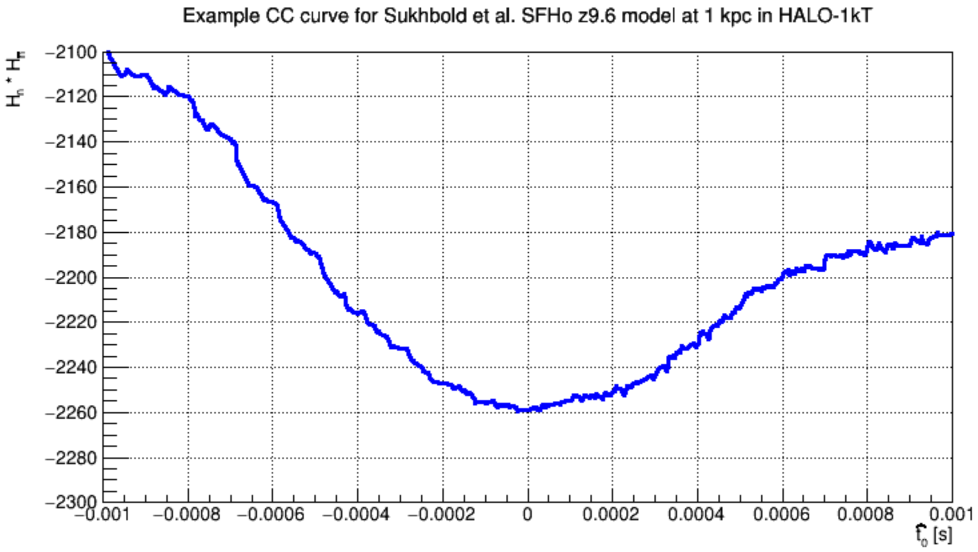

Figure 7.49: Example of the cross-correlation parameterized as a function of the offset (t_0) for the Sukhbold et al. SFHo z9.6 model in HALO-1kT at 1 kpc.

utilized to find the "true" $\hat{t_0}$ from this curve. In principle one could use the lower edge, or upper edge, of the bin that corresponds to the minimum step of the function. A N$^{\text{th}}$ degree polynomial could be fit to account for potential asymmetries on either side of the minimum. We elect to simply use the central value of the minimum bin, avoiding any potential systematic offsets that could present themselves from using either edge of the minimum or potential biases in the polynomial fit. As such the minimum is found via

$$\frac{t_u + t_l}{2} \tag{7.20}$$

where t_u and t_l correspond to the upper and lower edge of the step at the minimum. To find the lower bin edge, $(H_n * H_{\bar{n}})$ is simply scanned for increasing t_0, storing the current value of the function in a buffer variable if and only if it is less than the current value. From there, it is trivial to locate the upper bin edge as this is known a priori (100 μs). To assess any systematic offsets we shall cross-correlate all simulated SN bursts from the Sukhbold et al. SFHo z9.6 M$_\odot$ model.

7.6.4 Extraction of t_0

The CC technique was utilized for all 10^4 bursts of the Sukhbold et al. SFHo z9.6 $M_\odot$ model simulated in HALO-1kT. The distribution of extracted $\hat{t_0}$ is found in Figure 7.50. The results of this technique are quite surprising. Preliminary analysis had ruled out binned techniques as these were observed to fail at a higher rate than that of unbinned techniques such as the NLL fit. Comparing these two techniques, cross-correlating Sukhbold et al. SFHo z9.6 $M_\odot$ data to its own event rate, the NLL produced an extracted $\hat{t_0}$ distribution centered at -12 μs with a standard deviation of 190 μs, where as the CC technique produced a $\hat{t_0}$ distribution of 39 μs with a standard deviation of 210 μs. The CC technique maintains a competitive offset with respect to the extracted $\hat{t_0}$, but does deviate from the Monte Carlo truth t_0 by over triple the mean $\hat{t_0}$ provided by the NLL technique. We will hold off on making any further remarks with respect to the performance of these two techniques. Once systematic uncertainties are accounted for, a clearer picture of the CC's performance will arise. This can be found in §8. For now we turn our attention to the mixing of models and event rates.

7.6.5 Mixing models and event rates

Since the CC technique requires an event rate of the SN burst to cross-correlate, we will once again encounter systematic uncertainties when cross-correlating one model to another. In Figure 7.51, the Sukhbold et al. SFHo z9.6 $M_\odot$ model is cross-correlated with the Sukhbold et al. LS220 z9.6 $M_\odot$ event rate. Observed effects fall within expectations. Since each model's event rate does not differ largely in their neutronization and accretion phases the introduction of a minor systematic offset of -83 μs is anticipated (which is a

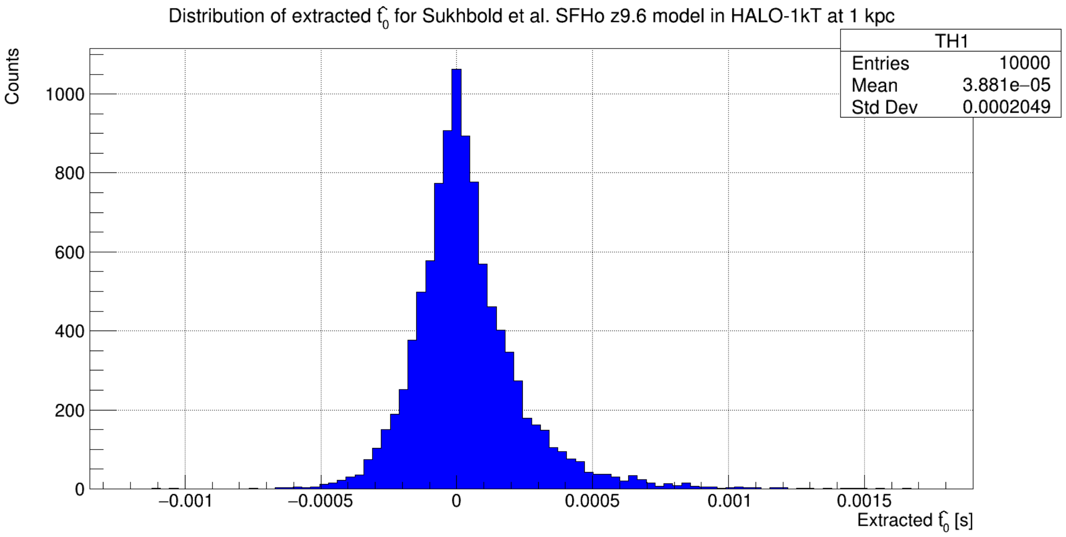

Figure 7.50: Distribution of extracted $\hat{t_0}$ for Sukhbold et al. SFHo z9.6 model at 1 kpc in HALO-1kT, using the CC technique.

factor of two larger than when the Sukhbold et al. SFHo z9.6 $M_\odot$ model is cross-correlated with itself)

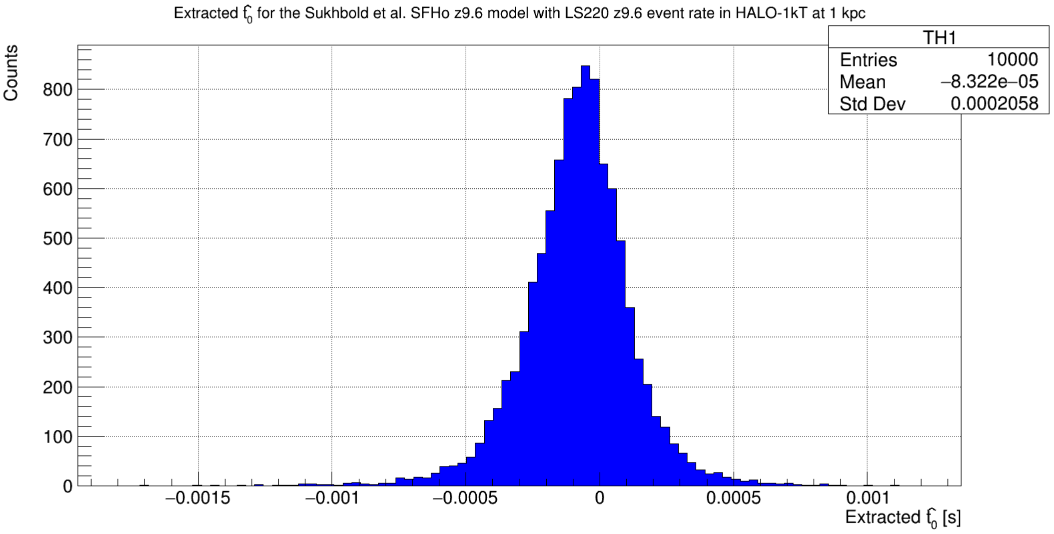

Figure 7.51: Distribution of extracted $\hat{t_0}$ obtained by cross-correlating Sukhbold et al. SFHo z9.6 time series data with the Sukhbold et al. LS220 z9.6 event rate. Results are in HALO-1kT at 1 kpc.

Systematic effects of greater magnitude are introduced when cross-correlating the Sukhbold et al. SFHo z9.6 $M_\odot$ time series data to either the Sukhbold et al. LS220 s27.0 $M_\odot$ or Sukhbold et al. SFHo s27.0 $M_\odot$ event rates. The results for the former can be found in Figure 7.52. We observe the narrowing and shift in the central peak from the extracted $\hat{t_0}$ distribution. Both of these features are attributed to the event rate of the 27.0 $M_\odot$ models greatly exceeding that of the lower mass models, especially in the 200 ms post-bounce (noted as the shouldering in the luminosity in §3.2). As a consequence, when populating the 1D histogram with the events as observed in HALO-1kT for the 9.6 $M_\odot$ models and cross-correlating it with the 27.0 $M_\odot$ models event rate, there is a degree of uncertainty with respect to where the peak event density in the burst correlates to the PDF. The inverse can be observed if we cross-correlate Sukhbold et al. LS220 s27.0 $M_\odot$ data to the Sukhbold et al. SFHo z9.6 $M_\odot$ event rate, the results are found in Figure 7.53.

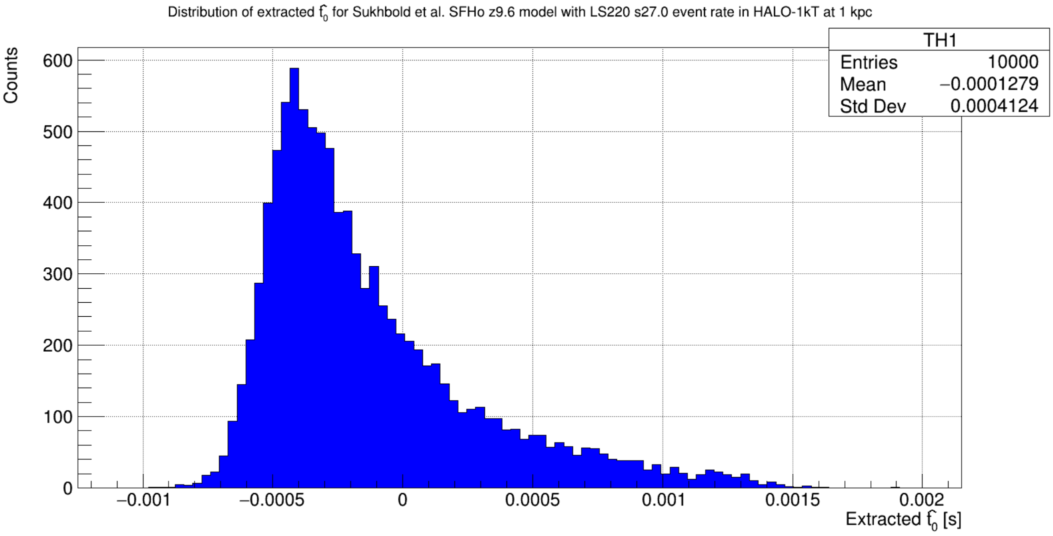

Figure 7.52: Distribution of extracted $\hat{t_0}$ obtained by cross-correlating Sukhbold et al. SFHo z9.6 time series data with the Sukhbold et al. LS220 s27.0 event rate. Results are in HALO-1kT at 1 kpc.

7.6.6 Estimating systematic uncertainties

To account for the systematic uncertainties present when mixing models, we once again follow the procedure in the same vein as the NLL fit and AD test. For the SN models simulated through the HALO-1kT Monte Carlo, each of the 10^4 simulated bursts from

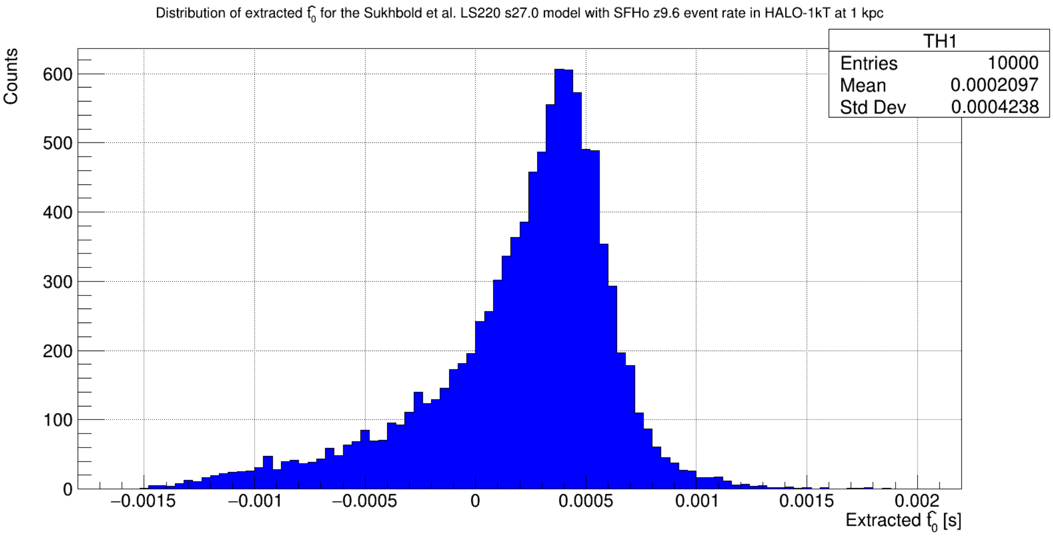

Figure 7.53: Distribution of extracted $\hat{t_0}$ obtained by cross-correlating Sukhbold et al. LS220 s27.0 time series data with the SFHo z9.6 event rate. Results are in HALO-1kT at 1 kpc.

each model are cross-correlated (Equation 7.19) with each PDF's event rate. This will yield sixteen extracted $\hat{t_0}$ distributions at each of the simulated distances. We then treat each distribution with equal weight and populate a new 1D histogram of the extracted $\hat{t_0}$ distribution by sampling the array of histograms randomly. The resulting distribution of $\hat{t_0}$ will have systematics included, as found in Figure 7.54. The yield of this distribution is a 35 μs offset from the Monte Carlo truth t_0 and a standard deviation of 370 μs. An unexpected result considering the performance of the preliminary implementation early on in this work. With this result the CC technique is only edged out by the CFD at a distance of 1 kpc. Whether the performance is maintained at distances beyond 1 kpc remains to be seen.

7.6.7 Performance as a function of distance

To quantify the potential walk in our extracted $\hat{t_0}$ and the performance at distances beyond 1 kpc, we carry out the same procedure noted above for distances 2 to 5 kpc.

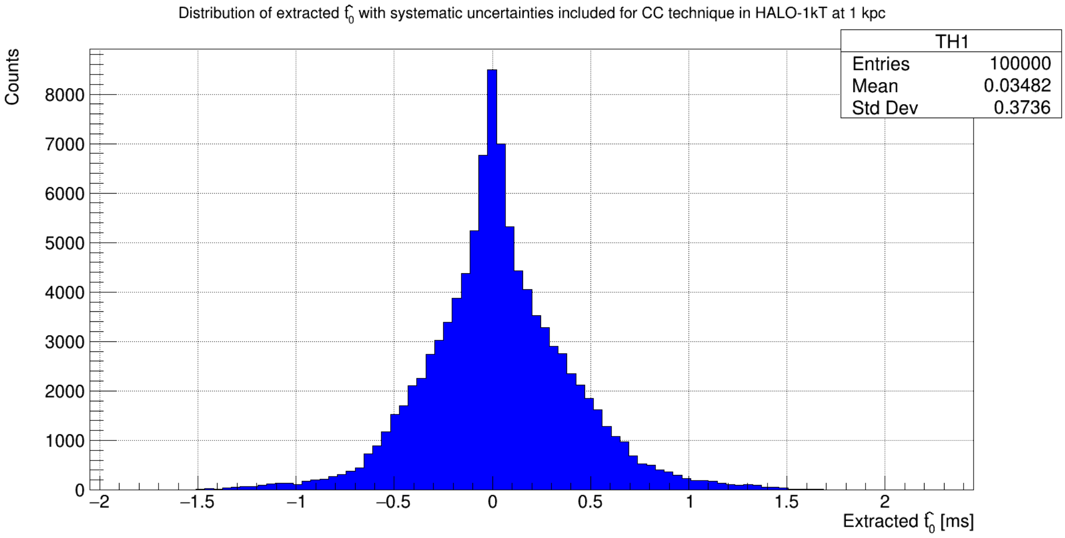

Figure 7.54: Extracted $\hat{t_0}$ from the CC technique at 1 kpc in HALO-1kT with systematic uncertainties.

The distributions of extracted $\hat{t_0}$ can be found in Figure 7.55 for all distances, while the performance curve can be found in Figure 7.56. In the other techniques explored, we noted the walk of the peak of the distribution away from the Monte Carlo truth t_0 (as a function of distance). In the CC technique, the mean of the distribution skews away from t_0, however the peak remains centered about the Monte Carlo truth t_0. By the time a distance of 5 kpc is reached, the right shoulder of the distribution is skewed to larger offsets from t_0. Its performance, compared to the NLL technique at this distance, is the inverse behaviour observed at 1 kpc, where the CC technique's precision is a factor of 1.5 times larger. Yet, this is an expected result as its reliance on a binned time series will cause the performance to degrade with a relatively low signal yield at a much faster rate. With all the techniques defined, we now seek to discuss the pros and cons of each, concluding by making recommendations for HALO and HALO-1kT.

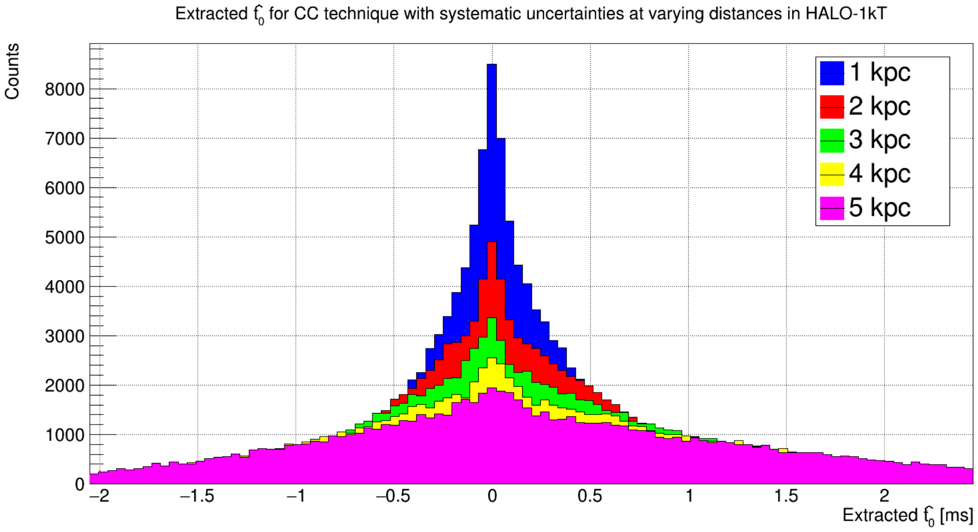

Figure 7.55: Extracted $\hat{t_0}$ from the CC technique at varying distances in HALO-1kT with systematic uncertainties.

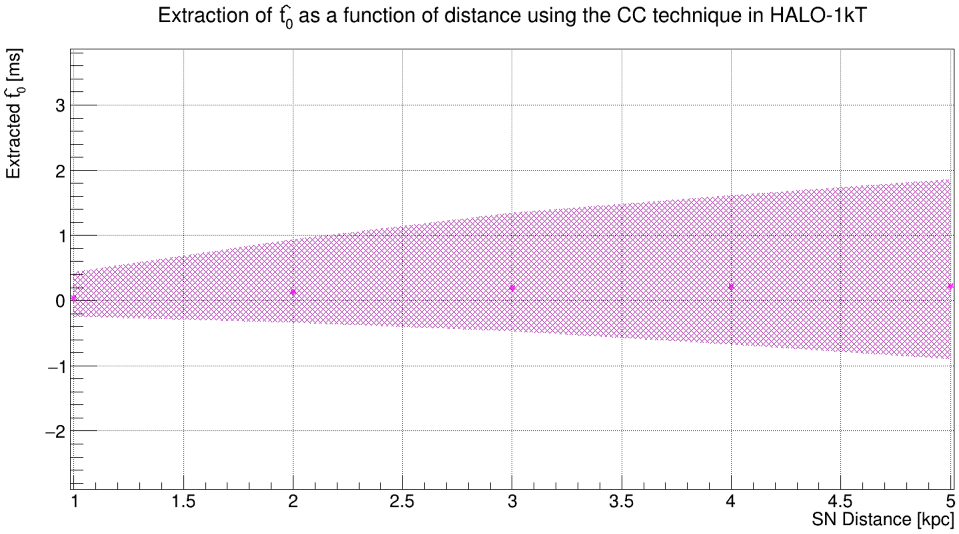

Figure 7.56: Performance curve for the CC technique in HALO-1kT, with the systematic uncertainties from model mixing included.

CHAPTER 8

Comparison of analytic techniques

In this section, we compare each of the techniques defined and modified in §7. Only HALO-1kT will be discussed here. For a discussion involving HALO and SNO+, please see §9. §8.1 tabulates and plots each of the techniques used; §8.2 explores how each technique can fail and the impact that this has on our analysis; and §8.3 provides a recommended approach for HALO-1kT to extract $\hat{t_0}$ and communicate it to SNEWS 2.0.

8.1 Precision of the extracted $\hat{t_0}$

In Figure 8.1, a graph of the extracted $\hat{t_0}$ as a function of distance and the accompanying projection for each distance is shown; Figure 8.1 includes systematic errors. The errors are estimated to be 1σ; we say estimated as the extracted $\hat{t_0}$ distributions can be non-Gaussian (mainly for close distances in the linear fit and AD test), so our integrated errors are an approximation at best. The exact values (i.e. $\hat{t_0}$ and integrated errors, at 68% confidence interval) are shown in Table 8.1. These are the final results, which quantify the ability of HALO-1kT to effectively extract t_0 from a CCSN signal.

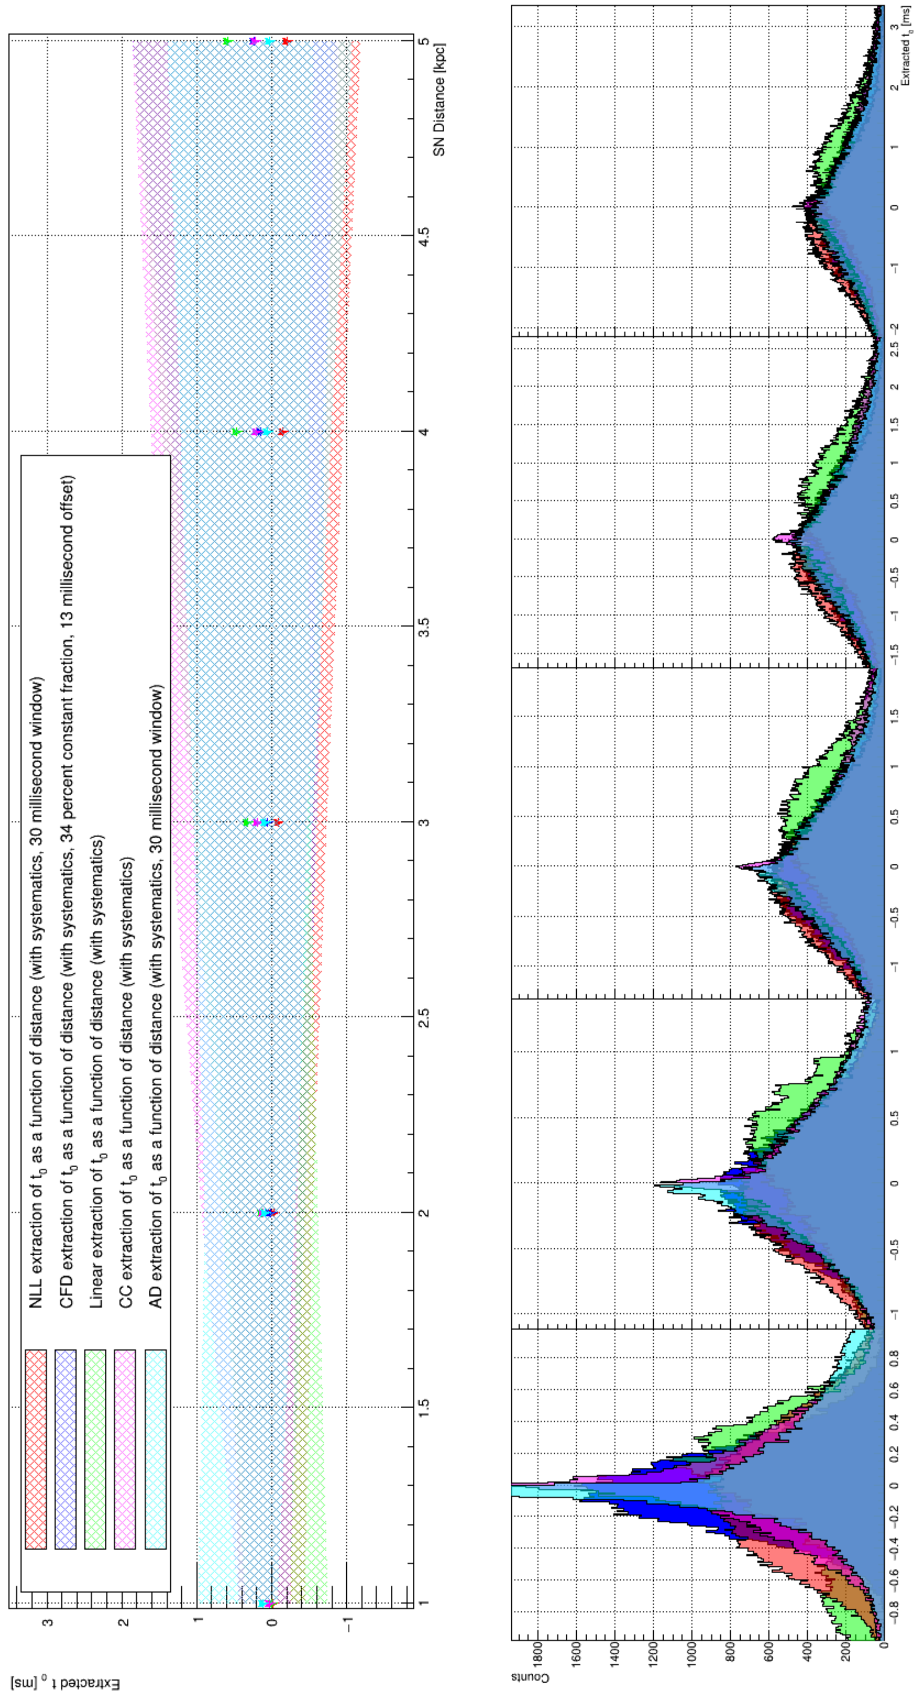

Figure 8.1: Extracted $\hat{t_0}$ as a function of distance for each of our techniques in their optimized configurations (top). Accompanying the graph is the projected performance for each of the five simulated distances (bottom). We note that the extracted $\hat{t_0}$ scale (x-axis) increases with SN distance.

Technique	1 kpc			2 kpc			3 kpc			4 kpc			5 kpc		
	$\hat{t_0}$ [μs]	σ_- [μs]	σ_+ [μs]	$\hat{t_0}$ [μs]	σ_- [μs]	σ_+ [μs]	$\hat{t_0}$ [μs]	σ_- [μs]	σ_+ [μs]	$\hat{t_0}$ [μs]	σ_- [μs]	σ_+ [μs]	$\hat{t_0}$ [μs]	σ_- [μs]	σ_+ [μs]
NLL	-20	420	460	-37	510	600	-81	660	780	-140	820	970	-210	980	1200
CFD	0	240	300	21	460	550	55	660	870	140	860	1200	230	1100	1600
Linear	0	750	360	140	750	520	330	840	720	480	970	990	590	1100	1300
CC	35	280	410	120	460	820	190	660	1200	200	880	1400	220	1100	1600
AD	110	190	880	90	400	810	83	620	970	63	860	1200	24	1100	1400

Table 8.1: Extracted $\hat{t_0}$ for all the techniques and the associated upper/lower integrated errors (σ_- and σ_+ respectively) obtained from Figure 8.1.

8.2 Points of failure and failure rates

With any fitting technique or test, there are scenarios in which they can “fail”. In dubbing some tests and fits as failures, it is intentionally amorphous; in this context a failed fit can still provide an estimate of t_0. Each fitting technique or test will have different scenarios in which they fail. We have already discussed at length how each technique can encounter issues; and how we have elected to address potential issues that arise with sufficient frequency. It is worth noting that rates of failure *do not* contribute to our overall recommendation for the technique that will be used in HALO-1kT. Firstly, when each of the techniques begins to fail, this is usually at a distance (or low event yield) that prevents us from extracting $\hat{t_0}$ with adequate precision to begin with. A simple reality is that beyond a certain distance, in this case $\approx 4 - 5$ kpc, HALO-1kT will be unable to achieve the required precision that is useful to SNEWS. For all techniques the increase in failure rate lags well behind the drop in precision of extracted $\hat{t_0}$. Secondly, the rates of these failures, are small, at distances of 5 kpc, occurring in $\leq 1\%$ of SN bursts for an analysis technique like cross-correlation.

8.2.1 NLL

If the estimator fails when using the NLL technique, `RooFit` will return the initial $\hat{t_0}$ value provided to the NLL function (in this study, the initial $\hat{t_0}$ value was 0 s; when implemented within the HALO and HALO-1kT DAQs, the first event time will be used). The NLL technique will fail only when a burst drops below an adequate number of events, as identified earlier in our discussion of time windows. Our analysis settled on a time window of 30 ms, which was wide enough for a sufficient number of events to be within the window to achieve an accurate and precise estimate of t_0. A feature that begins to arise at 5 kpc, and is only more pronounced at further distances, is an insufficient number of events within the time window. From Figure 7.1, it was shown that the average number of events across the simulated models was 15 (within the first 50 ms). There were simulated bursts that had as few as 3 events within the first 50 ms; our optimal time window of 30 ms would thus have a distribution of events in the time window shifted to fewer events. This is where the effects of a sparse event population in the leading edge arise more frequently. To address this, a non-static time window can easily be applied to ensure that a preset minimum number of events is included within the time window. For example, if at 6 kpc a burst had fewer than 5 events within the first 30 ms, the time window could be lengthened until 5 events were tallied in the window. The effect this would have on systematic uncertainties, particularly at distances of 5 kpc and beyond, is not known at this time. A non-static time window would lead to the need for a cost benefit analysis on how improved the analysis is by including additional events, but also degraded when accounting for systematic uncertainties (akin to §7.2.8). Such an analysis was not conducted as a part of this work.

8.2.2 CFD

The CFD technique, in its current configuration and tested from distances 1 to 5 kpc, does not fail. The largest complications that it faces are multiple intersection points. This was briefly addressed in §7.3.5, under the pretext of ill advised initial choices for both the constant fraction and temporal offset. With fewer events in the SN burst, the CFD output can oscillate around zero if the impact of the neutronization peak is reduced (this reduction being a product of the constant fraction, temporal offset, and low event yield). The solution to address these oscillations is to simply take the final intersection point, which is the most accurate and consistent representation of t_0. At further distances, the final intersection point becomes more varied in its location with respect to the offset correction established at 1 kpc. This is intrinsic to the technique and there are no further changes that can be introduced to alleviate this problem.

8.2.3 Linear fit

The linear fit faces few hurtles when attempting to extract an estimate of t_0. The primary complication, arising due to inherent differences in the spectra from varying supernova models, had led us to impose an early event cut on the data that would allow the models to better equate to one another. Other techniques will require automated minimization routines to obtain an estimate of t_0; the linear fit will not run into any issues in this context as minimization routines are absent from the linear fit in its entirety. A reduction in its performance is entirely attributed to a decrease in event statistics. Nonetheless, there are still some features that arise in our fit as a consequence of the decisions we have made thus far. With a reduction in total event yield, the observed events that are correlated with the neutronization burst will resemble a linear rise in events less and less. This is particularly true at distances of 10 kpc. As a direct consequence, the fit will have a more difficult time pinning down where t_0 is; this is an inevitable problem. Owing to the constraints that were placed on each parameter, chiefly their limits (see Table 7.1), the distributions of all parameters will be confined within these limits, regardless of whether the fitting routines within `ROOT` think this is the best fit for a linear line around the neutronization burst. This would be analogous to the overflow bins one would observe on an energy distribution through ADC channels. This problem is hardly a failure, as no other artifacts are observed in the p_2 distribution. The limits could be broadened, but in doing so, the distribution of p_2 will be widened as the fit is given more flexibility across all parameters. For this technique, it is important to remember; consistency is key, much more so than the accuracy of p_2 with respect to t_0, as this can always be corrected with what we had termed the offset correction.

8.2.4 Anderson-Darling test and Cross-correlation technique

The Anderson-Darling test and cross-correlation technique warrant little discussion with regards to the points of failure. For the AD test, this is in large part owing to its similarities to the NLL technique. The primary points of failure in the AD test arise due to a lack of event statistics in the time window (which can be rectified by making a variable time window based on event yield) and the minimization routine exiting without locating a minimum (which will, once implemented in real-world hardware, use the first event time). With regards to the cross-correlation technique, with dwindling statistics the performance of the technique will rapidly deteriorate (the rate of which is much faster than that of the unbinned techniques). This deterioration is a consequence of the fixed bin width on the histogram containing the simulated time series; as fewer of these bins will contain non-zero entries at distances beyond 5 kpc. The result, is the potential for binary bins; as discussed previously, these binary bins greatly hinder the performance of the CC technique. An attempt to increase the bin width would sacrifice the precision of the extracted $\hat{t_0}$.

8.3 Concluding remarks and recommendation for HALO-1kT

With the performance and considerations for failure out of the way, we are now poised to make our recommendation for HALO-1kT. By exploring multiple techniques, we have given ourselves a high degree of flexibility in how best to assist the SNEWS triangulation campaign. In Table 8.1, it becomes clear that at closer distances, the CFD technique provides the best performance at extracting $\hat{t_0}$. This extracted $\hat{t_0}$ distribution has been corrected to shift back to 0 seconds, but the majority of the data (68%) is located in a 543 μs range, providing an ideal estimate of t_0. The cross-correlation technique is also competitive at sufficiently close distances; however both the CC and the CFD technique drop out of favour the further away the SN is. When we reach the end of our simulated distances (5 kpc), both of these are surpassed by the NLL technique; meaning we should adopt a hybrid approach in how we can best communicate $\hat{t_0}$ to SNEWS 2.0. For high statistics observations (SN signals with $\geq 10^3$ events), we recommend that the CFD be used as the estimator of t_0. If we observe SN signals with hundreds of events or fewer, rather than thousands, typically at distances of 3 kpc and beyond, we recommend that the NLL technique replace the CFD, and be used for these low yield observations.

CHAPTER 9

Comparison of performance in HALO, HALO-1kT and SNO+

9.1 HALO vs. HALO-1kT

With HALO-1kT being the successor and bigger brother to HALO, it naturally follows that they would be similar detectors. Both detectors utilize the same detection mechanism; the same target medium; have the same neutrino flavour sensitivity, and so on and so forth. They are distinctly different in other aspects, chiefly the target mass and detector efficiency. Yet, from a signal point of view, both of these experiments are just counting neutrons that are induced from neutrino interactions in the lead. It is therefore expected that should HALO and HALO-1kT observe a supernova, the total neutron yield as a function of distance will differ by some arbitrary scaling factor provided backgrounds are minimized. This factor would be correlated to the change in detection efficiency (28% to 53%) and detector mass (79 t to 1 kt). This is the exact feature observed when comparing the normalized light curves for an arbitrary SN burst observed in HALO and HALO-1kT. In Figure 9.1, the two light curves only differ as a result of the neutron background, which causes minor augmentations to each curve.

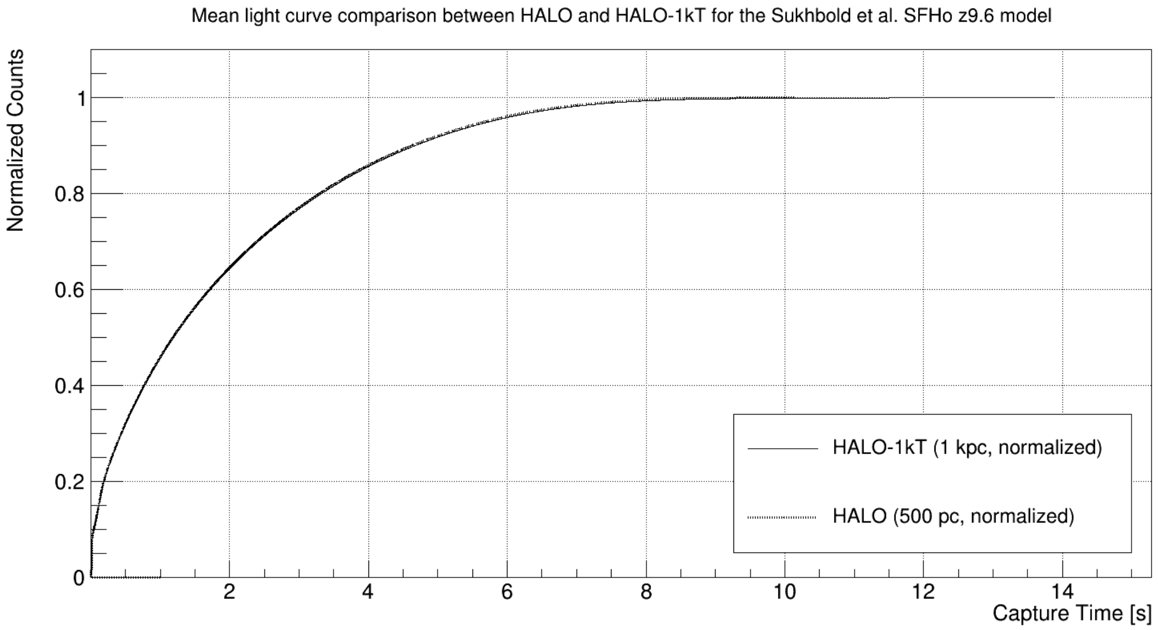

Figure 9.1: Mean light curves in HALO and HALO-1kT from the Sukhbold et al. SFHo z9.6 $M_\odot$ model.

Recall that for HALO-1kT, the target background neutron rate in the full-energy peak is 1 Hz; for HALO, it has a measured background rate of approximately 15 mHz. For HALO-1kT, in a typical core-collapse supernova burst that lasts $\approx$ 10 seconds, an average of 10 background events is expected based off current estimates. With HALO, 15 mHz is sufficiently small that on the time scales of 10s of seconds, no background events should occur during a typical CCSN burst. For low neutron yields in HALO-1kT, for example the Garching model at 10 kpc, the light curve will have contributions from background that make up roughly 1/3 of the total yield (25 neutrons observed from ν-Pb interactions, 10 background neutrons). When constructing the mean light curve at 10 kpc, the result is a relatively background free neutronization burst in the first 500 ms (where the ν_e flux peaks), but is then suppressed by the large contribution from the 1 Hz background. This can be seen in Figure 9.2, where the mean light curve is plotted for the Garching model simulated through the HALO-1kT Monte Carlo stimulation code at 1 kpc and 10 kpc, with the background added in post-simulation.

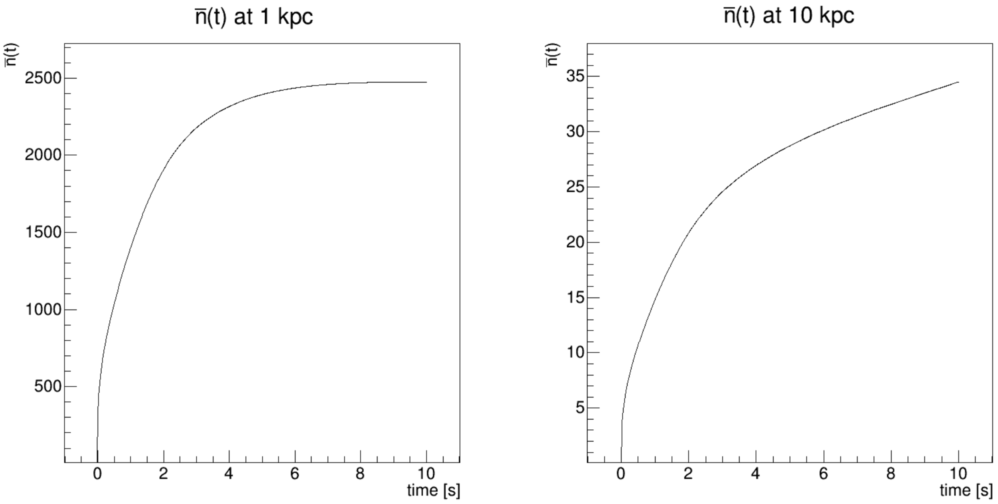

Figure 9.2: Mean light curve simulated in HALO-1kT from the Garching model at 1 and 10 kpc (left/right respectively).

For HALO, owing to its reduced backgrounds, if a SN burst results in a low neutron yield, the background will not dominate the latter half of the burst as observed in HALO-1kT, where the shape of the mean light curve deviates little from a high statistics observation (where the background is negligible). It is with this is mind that we can then explain the behaviour observed in Figure 9.2.

What does this mean for the analysis that was carried out in §7? Well, the results observed in HALO-1kT, with regards to the extraction of $\hat{t_0}$, will only differ from the HALO by some secondary scaling factor. Since much of our analysis has dealt with high statistic bursts and made use of an optimized time window equal to 30 ms, we can treat our signal as background "free". It would be of little use to compare HALO and HALO-1kT at the same distance, as the latter expects well over an order of magnitude increase in the total event yield. Instead, we consider a more fruitful comparison; in Figure 9.3 the distribution of extracted $\hat{t_0}$ is plotted for HALO at 500 pc and HALO-1kT at 3 kpc. Both plots are formed making use of the NLL technique with a 30 ms time window and accounting for systematic uncertainties. The performance in HALO and HALO-1kT is much more comparable (a direct consequence of equating distances that yield similar total event yields). Slight variations do occur between the two distributions, but these differences are a direct consequence of the average events expected at these two distances not being perfectly equivalent to one another. The culminating point of this subchapter is that HALO's performance curve is HALO-1kT's performance curve shifted to closer distances, due to the lower neutron yield. As a direct consequence, the recommendations for HALO-1kT that are made in §8 also apply here.

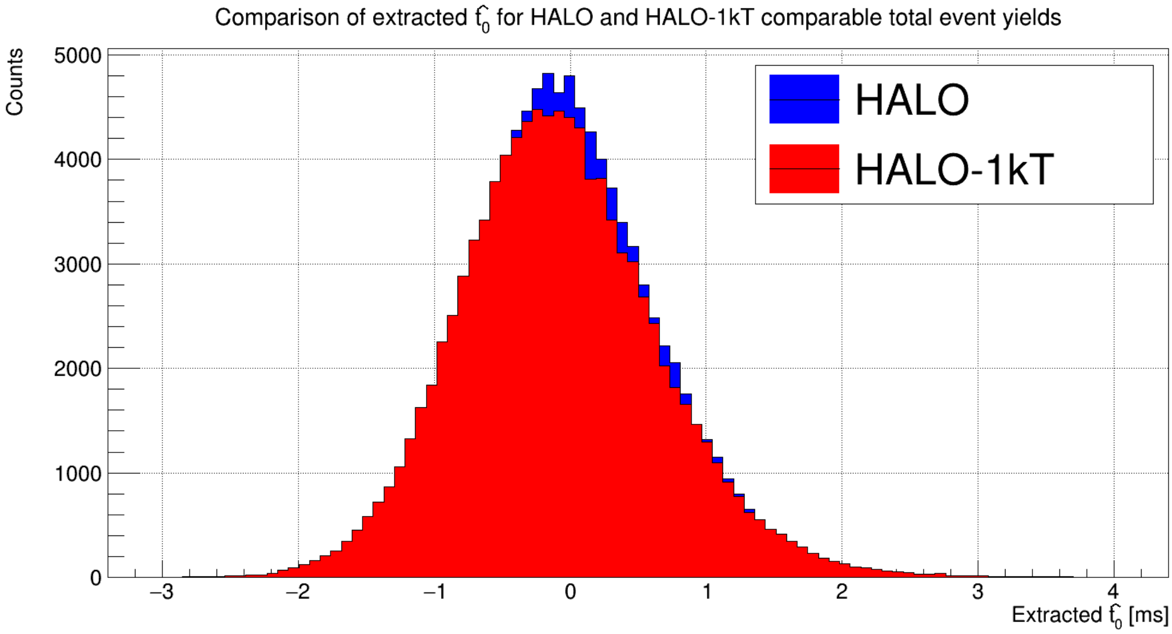

Figure 9.3: Comparison of the NLL technique performances in HALO and HALO-1kT at 500 pc and 3 kpc respectively (systematic uncertainties included, using a 30 ms time window).

9.2 SNO+ vs. HALO-1kT

From the outset, the intent was always to expand our implementation of the extraction of t_0 to other detectors involved in the SNEWS framework. Since I was a member of the HALO/HALO-1kT collaboration, the natural starting point was there. With my supervisor also being a member of the SNO+ collaboration, it was simple to gain access to this detector's Monte Carlo simulation, so the same analysis techniques can be easily implemented. The expansion of this work is preliminary; as shall be discussed in the coming subchapters, more work and discussion is needed to improve the results that follow. It was believed, by both myself and other members of the SNO+ collaboration, that whatever performance could be achieved in HALO-1kT, SNO+ would be able to achieve results on par with HALO-1kT or better. This assumption arose since SNO+ will observe a larger number of events for an arbitrary CCSN model than HALO-1kT. As the reader shall quickly see, this was not the case.

9.2.1 NLL

To facilitate a comparison between HALO-1kT and SNO+, we will once again make use of the optimal configuration derived in §7.2 (namely a 30 ms time window). For SNO+, the intent was to raise the trigger condition for the analysis such that the background can be reduced (or entirely mitigated). As seen in Figure 6.8, the distribution of background events should be uniform throughout the burst. These background events tend to be low N_{hit}. It is then trivial to raise the threshold to 800 N_{hit}, and remove the background entirely. We identify the first event as outlined in §7.1.2. The downside of such a high threshold is that low energy (N_{hit}) supernova-related events will be removed along with the background events. As a consequence, the low energy p-scattering events are completely removed from the burst, thereby reducing the total yield from the SN burst.

It is worth noting that before the aggressive cuts were placed on the SNO+ burst, the total number of neutrino interactions, and therefore yield, is much larger than HALO-1kT. This remains unchanged. In Figure 9.4, the mean light curve (with the aforementioned cuts applied where appropriate) for both HALO-1kT and SNO+ is shown for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc. In the first 20 ms past t_0, HALO-1kT observes on average 2.5 times more events. This is a consequence of the neutronization burst being primarily made up of ν_e. As the flux of this species reduces, and the other flavours increase in quantity, SNO+ will begin to dominate with respect to the number of events seen. The mean light curves between the two detectors differ as a direct result of differences in detector mass, cross sections, flavour sensitivity, and thresholds. It is worth mentioning that only up to the first 50 ms post core-bounce of the SN burst is calculated and simulated in SNO+, hence the abrupt cutoff in the data. The sole reason is to ensure an adequate number of simulations can be simulated without clogging the computing cluster at SNOLAB for months; for HALO-1kT, each model of supernova was simulated 10^4 times, at 5 distances; the same goes for HALO. Had this been done for SNO+, the event rate calculations and simulation time would take months, based on preliminary estimates. Applying a time cut of 50 ms post core-bounce was the reasonable solution.

For a direct comparison, the extracted $\hat{t_0}$ distribution for both SNO+ and HALO-1kT, from the NLL fit at 1 kpc for the Sukhbold et al. SFHo z9.6 model, can be found in Figure 9.5. Contrary to expectations, the SNO+ performance is severely degraded. The precision to which SNO+ can extract $\hat{t_0}$ is reduced by a factor of 3, even though the total event yield in the analysis is almost a factor of 2 larger. Traditionally, and as observed in HALO-1kT in §8, regardless of technique, an increase in statistics should allow for increased precision of the $\hat{t_0}$. This is not the observed feature in SNO+ (in comparison to

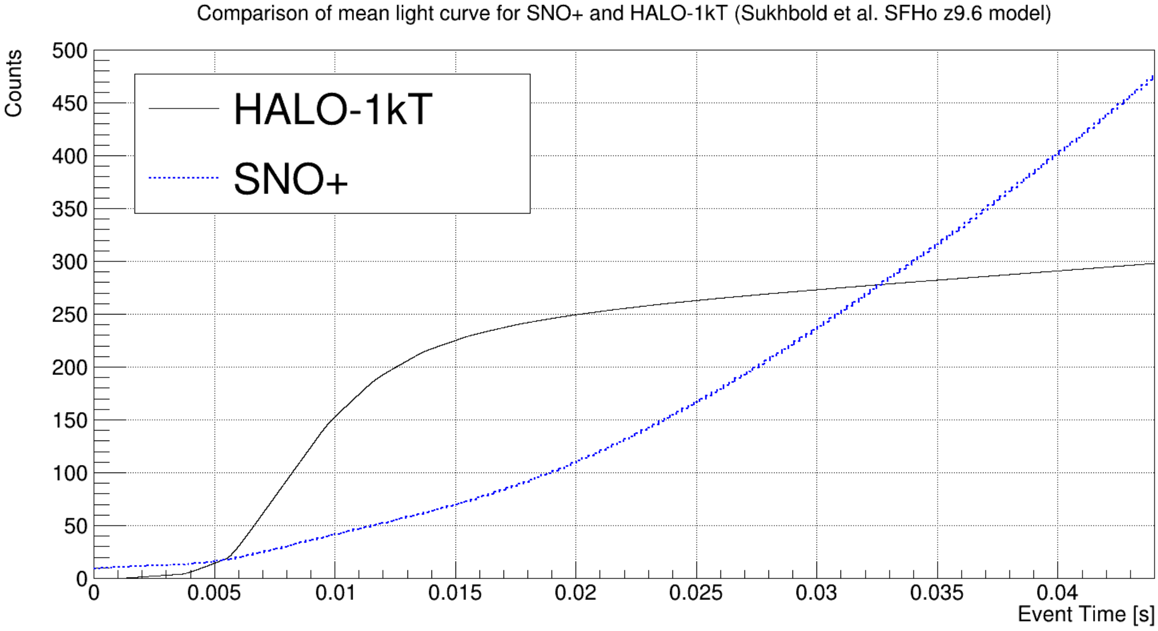

Figure 9.4: Mean light curves for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc simulated in both HALO-1kT and SNO+.

HALO-1kT).

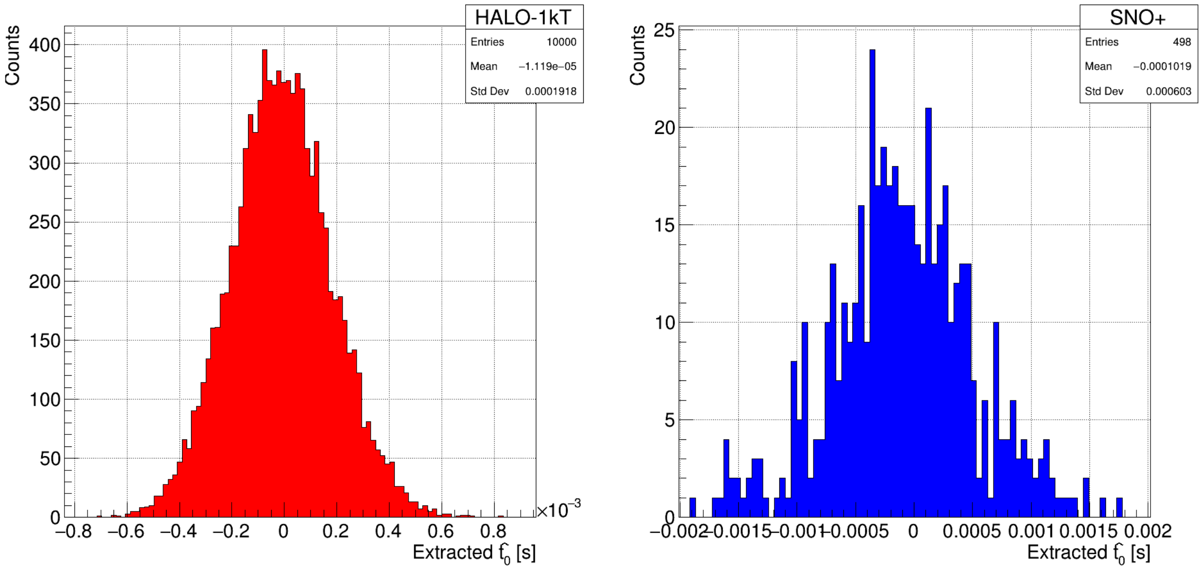

Figure 9.5: Comparison between extracted $\hat{t_0}$ distribution for HALO-1kT and SNO+ using the NLL technique for the Sukhbold et al. SFHo z9.6 $M_\odot$ model at 1 kpc.

When systematic uncertainties are accounted for, the outlook is not improved, as can be seen in Figure 9.6. My naive assumption was false, that based purely on statistics, the performance in SNO+ would be greatly improved over HALO-1kT. One might suggest that the cuts were too aggressive, or that cutting the data at 50 ms post core-bounce is to strict, or that an insufficient amount of the burst was used. These amongst other causes were investigated, with solutions implemented, to no avail. If the cuts on the background are reduced, additional contributions to the SN signal from p-scattering events arise, but so too does the background, degrading performance. If the number of events above a certain $\mathrm{N_{hit}}$ to trigger the analysis is reduced, as per the condition to identify the first event and trigger the analysis, then sporadic events in the tens of milliseconds prior to t_0 can trigger the analysis condition early, resulting in a failed fit. If we opt to increase the simulation duration, and simulate beyond the first 50 ms past t_0, we then run the risk of increasing systematic uncertainties (akin to HALO-1kT) as the light curves will differ greatly based on the model under consideration.

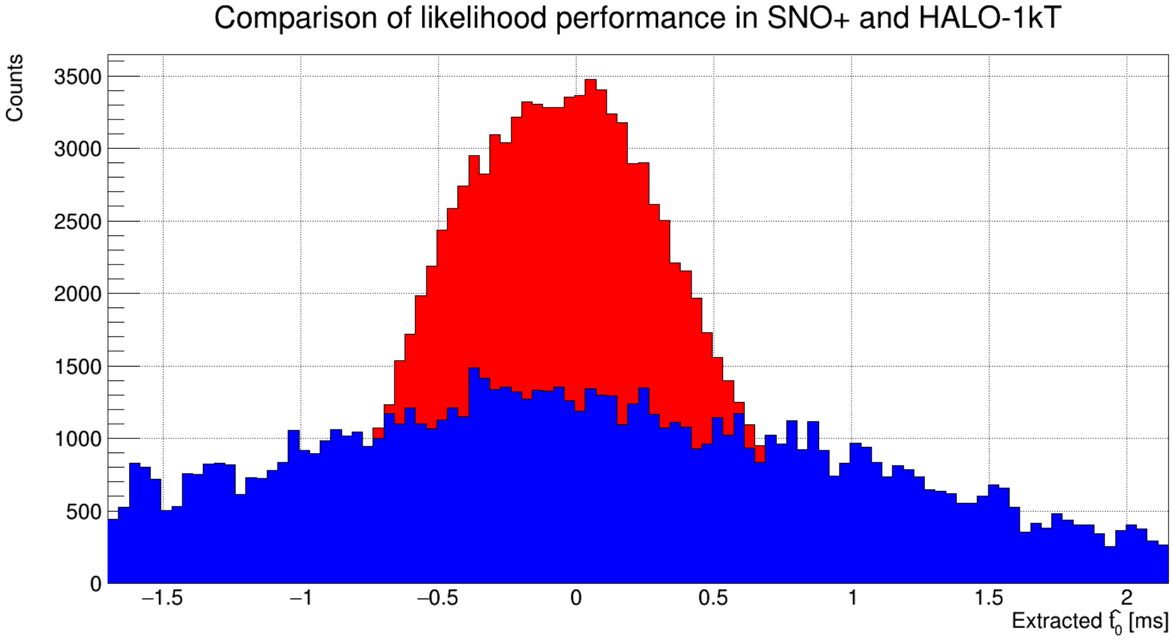

Figure 9.6: Comparison between extracted $\hat{t_0}$ distribution for HALO-1kT and SNO+ using the NLL technique with systematic uncertainties included at 1 kpc.

9.2.2 CFD

The same trigger conditions that were used in the analysis of the NLL were used here, chiefly that a minimum of 5 events above 500 N_{hit} fall within a rolling 20 ms window. We made use of the same optimal parameters as HALO-1kT, a constant fraction of 0.34 and a temporal offset of 13 ms. The expectation isn't that SNO+ would end up using the exact same configuration, as this would be dependant on light curve; as a first comparison between the two detectors it was useful to get an approximate idea of performance. The results of performing this technique on the Sukhbold et al. SFHo z9.6 $M_\odot$ at 1 kpc data set can be found in Figure 9.7 (for SNO+). As was the case for the NLL fit, the CFD performance is degraded in SNO+ in comparison to HALO-1kT. Even with an optimal configuration for the CFD, SNO+ results cannot achieve comparable results to HALO-1kT.

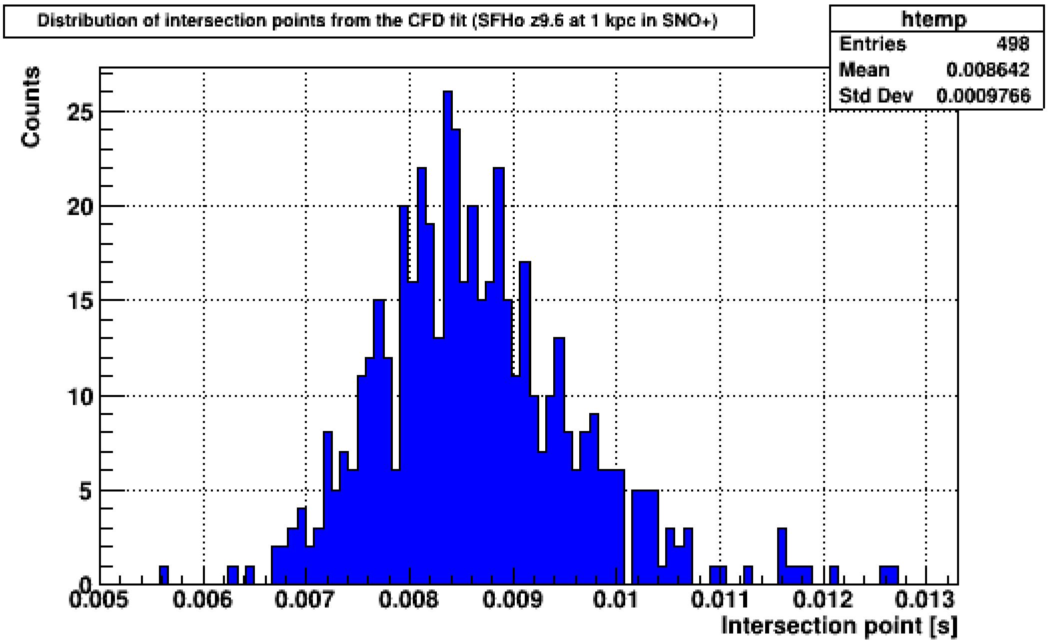

Figure 9.7: Distribution of extracted intersection point for the Sukhbold et al. SFHo z9.6 model at 1 kpc in SNO+ using the CFD technique (prior to the offset correction).

9.2.3 SNO+ considerations

Both the NLL and CFD technique suffer from poor performance in SNO+, compared to its relatively low statistics partner HALO-1kT. If it is not a problem that can be addressed by the total yield, some other underlying condition must drive this discrepancy. Owing to the preliminary nature of this analysis in SNO+, there has been no rigorous investigation into the source or solution to this poor performance. What I would like to do here is briefly discuss what I believe the source to be, and how to potentially mitigate it.

It has been described in detail and stated plainly many times throughout this report, HALO-1kT and SNO+ are two detectors that could not differ more than they already do. Their only shared trait is one of SNO+ physics goals. When we looked at the mean light curve in HALO-1kT and SNO+, there was a tremendous difference in the leading edge. The event rate in HALO-1kT ramped up aggressively around the neutronization burst, while SNO+ maintained a quasi-static event rate in the tens of milliseconds after core-bounce. In the case of the Sukhbold et al. SFHo z9.6 model at 1 kpc, HALO-1kT has 100 events within 80 ms of core-bounce; it takes SNO+ upwards of 200 ms past t_0 to reach a similar number of events (as seen in the aforementioned figure). The result is a low "event density" in the leading edge of the burst. This is not unique to the Sukhbold et al. SFHo z9.6 $M_\odot$ model, as seen in Figure 9.8, all other Sukhbold et al. models exhibit this behaviour. This feature, as discussed earlier, is expected and a direct consequence of the flavour sensitivity difference between the two detectors. The low yield during the neutronization burst must drive the variability in CFD and NLL performance.

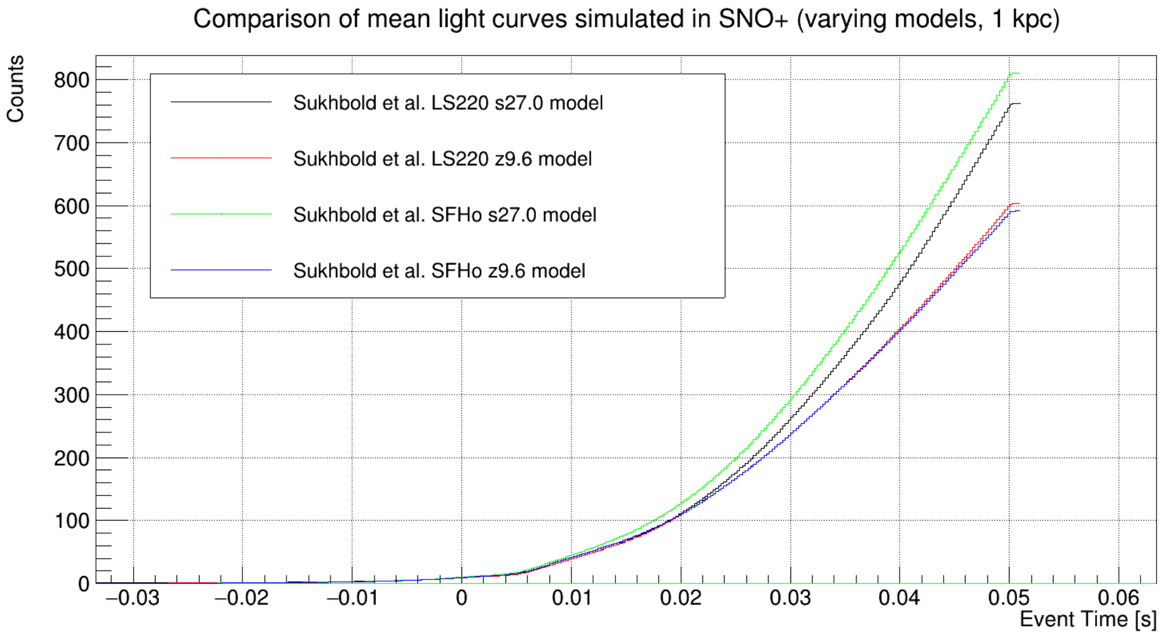

Figure 9.8: Mean light curves for the Sukhbold et al. models simulated in SNO+ at 1 kpc.

CHAPTER 10

Final thoughts and what is to come

Throughout this thesis, the ability to extract $\hat{t_0}$ has been quantified for six different techniques in HALO and HALO-1kT; a preliminary exploration of these techniques has been done in SNO+. We determined, that for high statistics bursts in HALO and HALO-1kT, the CFD technique provides the best estimate of t_0 over all other techniques; while at larger distances, with fewer statistics, the NLL fit provides the optimal performance. There are points of discussion that are worth mentioning, and are discussed in the sections that follow.

10.1 Neutrino flavour transformations

There has been little discussion on the impact of neutrino flavour transformations on our ability to extract $\hat{t_0}$. The results found in Table 8.1 are a best case scenario for HALO-1kT. `SNEWPY` has within it, a good variety of oscillation descriptions that can be used to transform the neutrino fluxes. A handful are; no transformation, complete exchange, adiabatic MSW for the normal mass ordering (NMO), and adiabatic MSW for the inverted mass ordering (IMO). For a thorough descriptions of these transformation scenarios, please consult the `sntools` documentation [87]. For HALO-1kT, since the neutronization burst is dominated by ν_e, any flavour transformation description used in the neutrino event calculations will reduce the total number of ν_e, while simultaneously increasing the number of ν_μ or ν_τ. Since HALO and HALO-1kT are primarily ν_e sensitive, any oscillation scenario that reduces the total number of ν_e in favour of ν_μ or ν_τ will also reduce the total number of events observed in that SN burst. We know that from our results in HALO-1kT and SNO+, any reduction in statistics will bring with it an increase in the uncertainty of $\hat{t_0}$. For HALO-1kT, we used no flavour transformations. For SNO+, at 1 kpc we simulated the Sukhbold et al. SFHo z9.6 $M_\odot$ model with each of the four aforementioned flavour transformations. Each of these simulation sets had their mean light curves formed (as described in Section 7.2.3) and the NLL fit was then performed on each simulation in the set. The extracted $\hat{t_0}$ can be found in Figure 10.1 for each of the flavour transformations. The effects of oscillations on the SNO+ data are marginal, and no sizeable performance increase or decrease is observed in its extraction of $\hat{t_0}$. The original intent was to observe, under a more realistic scenario, if the mean light curve is augmented enough that the degraded performance in SNO+ that we tied to the shape of the mean light curve could be improved; this was not the case. A similar analysis will need to be carried out for HALO and HALO-1kT, as the inclusion of oscillations will most certainly modify every aspect of our analysis (i.e. the amount each technique walks, the offset correction for the CFD and linear tests, perhaps which is the ideal technique at close and far distances, etc.).

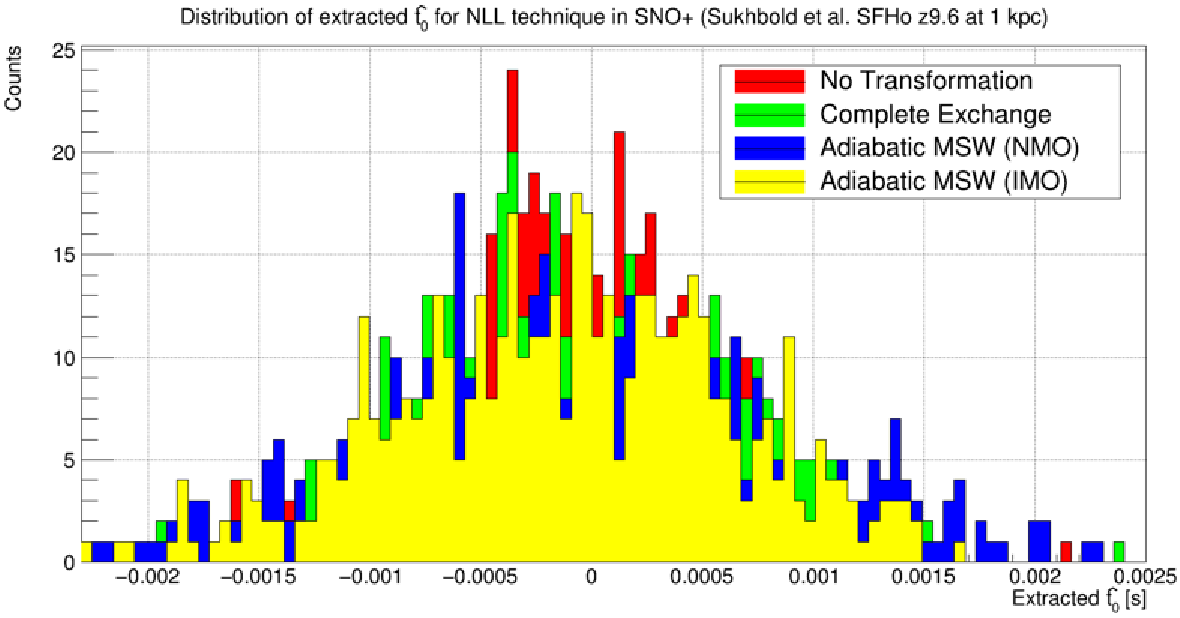

Figure 10.1: Distribution of extracted $\hat{t_0}$ for the Sukhbold et al. SFHo z9.6 model at 1 kpc (varying flavour transformations). Each is fit to its own respective PDF, in other words, there is no mixing of models and PDFs in the above data.

10.2 Inclusion of other SN models

The analysis contained in §7 to §9, was limited to 4-5 models (the total number being dependent on the detector under consideration). `SNEWPY` has dozens of models, meaning that these additional models could be incorporated into our analysis with ease, and additional probing can be done to determine if other issues will arise in our analysis that are intrinsic to these untested models. We limited ourselves to 4-5 models in the analysis (we did simulate many more, as stated previously), solely to prevent the time scales of this thesis and analysis to keep expanding. Of particular interest are the more massive models that collapse into black holes, as these cover a subset of SN simulations not covered in the analysis pertaining to extraction precision.

10.3 Outlook

HALO-1kT is able to achieve the recommended precision to triangulate out to distances of $\approx$ 3 kpc. The studies carried out in [6, 52] indicated that next-generation detectors would be able to triangulate out to distances of 10 kpc and beyond. Unfortunately this will not be possible in HALO or even its larger brother HALO-1kT. This should come as no surprise as the total yield of events from a SN burst can be orders of magnitude less than other detectors (mainly JUNO, Hyper-K, DUNE, etc.), greatly reducing the precision to which we can estimate $\hat{t_0}$. That being said, HALO and HALO-1kT are not without benefits. Galactic supernovae are exceedingly rare, meaning that for a detector to contribute to the triangulation campaign in any meaningful capacity, it must be collecting data for extended periods of time. With almost half of the detectors that originally comprised SNEWS having gone offline in recent years, the reality is that these high statistics experiments will collect their data, meet their primary physics goals, and then subsequently go offline. Who could be left to provide astronomers with an estimate of the supernova's location? HALO will continue operating as it has for the past 11 years; if HALO-1kT is built, it would do the same. Low cost, low maintenance, and high live-time detectors such as these could play an integral role in communicating $\hat{t_0}$ to SNEWS should other high statistics detectors come offline.

With all these techniques fully deployed in HALO and HALO-1kT, we now look to expand this work to other detectors. The hope is that all of the techniques explored in this thesis will be integrated into `SNEWPDAG`, the SNEWS 2.0 collaboration's tool for triangulation calculations. In late 2022, I began work on integrating the CFD and NLL techniques into `SNEWPDAG`, with the former producing the same estimate for t_0 as my own code package. From here, the ability to extract $\hat{t_0}$ can be quantified for current and next-generation detectors, and similar studies such as [6, 52] can be undertaken.

Bibliography

[1] K. Rozwadowska et al. On the rate of core collapse supernovae in the milky way. *New Astronomy*, 83:101498, 2021.

[2] R. M. Bionta et al. Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud. *Phys. Rev. Lett.*, 58:1494–1496, Apr 1987.

[3] E.N. Alexeyev et al. Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope. *Physics Letters B*, 205(2):209–214, 1988.

[4] K. Hirata et al. Observation of a neutrino burst from the supernova SN1987A. *Phys. Rev. Lett.*, 58:1490–1493, Apr 1987.

[5] K. Scholberg. SNEWS: The supernova early warning system. In *AIP Conference Proceedings*. AIP, 2000.

[6] N. B. Linzer et al. Triangulation pointing to core-collapse supernovae with next-generation neutrino detectors. *Physical Review D*, 100(10), nov 2019.

[7] NASA. Pillars of Creation (Hubble and Webb Images Side by Side). https://webbtelescope.org/contents/media/images/2022/052/01GF44EV0PPW2BHJS9HMA1AGEK, 2022. Accessed: 11/06/2023.

[8] W. Kunkel et al. Supernova 1987A in the Large Magellanic Cloud. IAU Circ., 4316:1, February 1987.

[9] F. Zhao et al. The Guest Star of AD185 must have been a Supernova. *Chinese Journal of Astronomy and Astrophysics*, 6(5):635, oct 2006.

[10] G. A. Tammann et al. The Galactic Supernova Rate. *ApJS*, 92:487, June 1994.

[11] NASA. Hubble Captures Wide View of Supernova 1987A. https://hubblesite.org/contents/media/images/2017/08/3987-Image.html, 2017. Accessed: 11/06/2023.

[12] D. D. Clayton. *Principles of Stellar Evolution and Nucleosynthesis*, chapter 5. University of Chicago Press, 1st edition, 1983.

[13] W. Hillebrandt et al. Towards an understanding of Type Ia supernovae from a synthesis of theory and observations. *Frontiers of Physics*, 8:116–143, april 2013.

[14] D. Maoz et al. Observational Clues to the Progenitors of Type Ia Supernovae. *Annual Review of Astronomy and Astrophysics*, 52(1):107–170, aug 2014.

[15] F. Hoyle et al. Nucleosynthesis in Supernovae. *ApJ*, 132:565, November 1960.

[16] W. Hillebrant et al. How to Blow Up a Star. *Scientific American*, 295:43–49, 2006.

[17] T. Kunugise et al. Neutrino Emission from Type Ia Supernovae. *Publications of the Astronomical Society of Japan*, 59(6):L57–L61, 12 2007.

[18] L. Hüdepohl et al. Neutrino Signal of Electron-Capture Supernovae from Core Collapse to Cooling. *Phys. Rev. Lett.*, 104:251101, Jun 2010.

[19] W. P. Wright et al. Neutrinos from type Ia supernovae: The gravitationally confined detonation scenario. *Phys. Rev. D*, 95:043006, Feb 2017.

[20] W. P. Wright et al. Neutrinos from type Ia supernovae: The deflagration-to-detonation transition scenario. *Physical Review D*, 94(2), jul 2016.

[21] A. Burrows et al. Core-collapse supernova explosion theory. *Nature*, 589(7840):29–39, jan 2021.

[22] S. S. M. Wong. *Introductory Nuclear Physics*, chapter 10. University of Chicago Press, 2nd edition, 1998.

[23] H. A. Bethe. Energy Production in Stars. *Phys. Rev.*, 55:434–456, Mar 1939.

[24] C. F. von Weizsäcker. Über Elementumwandlungen in Innern der Sterne I. *Physikalische Zeitschrift*, pages 176–191, 1937.

[25] C. F. von Weizsäcker. Über Elementumwandlungen in Innern der Sterne II. *Physikalische Zeitschrift*, pages 633–646, 1938.

[26] S. Woosley et al. The physics of core-collapse supernovae. *Nature Physics*, 1(3):147–154, dec 2005.

[27] S. M. Couch. The mechanism(s) of core-collapse supernovae. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2105):20160271, 2017.

[28] H. T. Janka. Explosion Mechanisms of Core-Collapse Supernovae. *Annual Review of Nuclear and Particle Science*, 62(1):407–451, nov 2012.

[29] A. Burrows. Supernova explosions in the Universe. *Nature*, 403(6771):727–733, feb 2000.

[30] S. A. Colgate et al. The Hydrodynamic Behavior of Supernovae Explosions. *ApJ*, 143:626, March 1966.

[31] H. Shen et al. Relativistic equation of state of nuclear matter for supernova and neutron star. *Nuclear Physics A*, 637(3):435–450, jul 1998.

[32] T. Sukhbold et al. CORE-COLLAPSE SUPERNOVAE FROM 9 TO 120 SOLAR MASSES BASED ON NEUTRINO-POWERED EXPLOSIONS. *The Astrophysical Journal*, 821(1):38, apr 2016.

[33] J. M. Lattimer et al. A generalized equation of state for hot, dense matter. *Nuclear Physics A*, 535(2):331–376, 1991.

[34] A. W. Steiner et al. CORE-COLLAPSE SUPERNOVA EQUATIONS OF STATE BASED ON NEUTRON STAR OBSERVATIONS. *The Astrophysical Journal*, 774(1):17, aug 2013.

[35] D. Kresse et al. Stellar Collapse Diversity and the Diffuse Supernova Neutrino Background. *The Astrophysical Journal*, 909(2):169, mar 2021.

[36] A. Mirizzi et al. Supernova neutrinos: Production, oscillations and detection. *La Rivista del Nuovo Cimento*, 39(102):1–112, Feb 2016.

[37] M. L. Warren et al. Constraining Properties of the Next Nearby Core-collapse Supernova with Multimessenger Signals. *The Astrophysical Journal*, 898(2):139, jul 2020.

[38] I. Tamborra et al. Neutrino emission characteristics and detection opportunities based on three-dimensional supernova simulations. *Physical Review D*, 90(4), aug 2014.

[39] T. Kuroda. Impact of a Magnetic Field on Neutrino–Matter Interactions in Core-collapse Supernovae. *The Astrophysical Journal*, 906(2):128, jan 2021.

[40] K. Nakazato et al. SUPERNOVA NEUTRINO LIGHT CURVES AND SPECTRA FOR VARIOUS PROGENITOR STARS: FROM CORE COLLAPSE TO PROTO-NEUTRON STAR COOLING. *The Astrophysical Journal Supplement Series*, 205(1):2, feb 2013.

[41] S. Zha et al. Progenitor Dependence of Hadron-quark Phase Transition in Failing Core-collapse Supernovae. *The Astrophysical Journal*, 911(2):74, apr 2021.

[42] L. Walk et al. Neutrino emission characteristics of black hole formation in three-dimensional simulations of stellar collapse. *Physical Review D*, 101(12), jun 2020.

[43] S. Wanajo et al. Nucleosynthesis in Electron Capture Supernovae of Asymptotic Giant Branch Stars. *The Astrophysical Journal*, 695:208, 03 2009.

[44] S. Zhu et al. Nuclear Data Sheets for A=214. *Nuclear Data Sheets*, 175:1–149, 2021.

[45] P. Antonioli et al. SNEWS: the SuperNova Early Warning System. *New Journal of Physics*, 6:114–114, sep 2004.

[46] J. F. Beacom et al. Can a supernova be located by its neutrinos? *Physical Review D*, 60(3), jul 1999.

[47] R. Tomà et al. Supernova pointing with low- and high-energy neutrino detectors. *Physical Review D*, 68(9), nov 2003.

[48] SNEWS Collaboration. SuperNova Early Warning System (SNEWS) Publishing Tools, 2023.

[49] SCiMMA. Hop Client, 2023.

[50] SNEWS Collaboration. SuperNova Early Warning System (SNEWS) 2.0 Alert and Pointing Calculations, SNEWPDAG, 2023.

[51] V. Brdar et al. Neutrino astronomy with supernova neutrinos. *Journal of Cosmology and Astroparticle Physics*, 2018(04):025–025, apr 2018.

[52] A. Coleiro et al. Combining neutrino experimental light-curves for pointing to the next galactic core-collapse supernova. *The European Physical Journal C*, 80(9), sep 2020.

[53] SNEWS Collaboration. SNEWS 2.0. Funding application for National Science Foundation Windows on the Universe: The Era of Multi-Messenger Astrophysics (NSF WoU-MMA). Unpublished material., 2018.

[54] B. Martin. *Nuclear and Particle Physics: An Introduction.* John Wiley & Sons Ltd, 2019.

[55] R. Davis. A review of the homestake solar neutrino experiment. *Progress in Particle and Nuclear Physics*, 32:13–32, 1994.

[56] B. Cleveland et al. Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector. *The Astrophysical Journal*, 496(1):505, mar 1998.

[57] M.C. Chen. The Sudbury Neutrino Observatory: Comparison of Phases. *Nuclear Physics B - Proceedings Supplements*, 145:5–10, 2005. NOW 2004.

[58] F. Duncan. SNOLAB Technical Drawing. Internal SNOLAB Document.

[59] A. Allega et al. Evidence of Antineutrinos from Distant Reactors Using Pure Water at SNO+. *Phys. Rev. Lett.*, 130:091801, Mar 2023.

[60] R. Saakyan. Two-Neutrino Double-Beta Decay. *Annual Review of Nuclear and Particle Science*, 63(1):503–529, 2013.

[61] V. Albanese et al. The SNO+ experiment. *Journal of Instrumentation*, 16(08):P08059, aug 2021.

[62] C. Alduino et al. Measurement of the two-neutrino double-beta decay half-life of ^{130}Te with the CUORE-0 experiment. *The European Physical Journal C*, 77(1), jan 2017.

[63] J. B. Albert et al. An improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector. *Physical Review C*, 89(1), jan 2014.

[64] L. N. Machado et al. Pre-supernova Alert System for Super-Kamiokande. *The Astrophysical Journal*, 935(1):40, aug 2022.

[65] J. Rumleskie. *SNO+ Sensitivities to Pre-supernova and Supernova Neutrinos*. PhD thesis, Laurentian University, 2021. https://omni.laurentian.ca/permalink/01OCUL_LU/14p4p5m/alma991010222759805165.

[66] J. Rumleskie et al. Supernovae and SNO+. *Journal of Physics: Conference Series*, 1342(1):012135, jan 2020.

[67] M. Stringer. *Sensitivity of SNO+ to supernova neutrinos*. PhD thesis, University of Sussex, 2019. http://sro.sussex.ac.uk/id/eprint/81233.

[68] C.K. Hargrove et al. A lead astronomical neutrino detector: LAND. *Astroparticle Physics*, 5(2):183–196, 1996.

[69] P.F. Smith. OMNIS—an improved low-cost detector to measure mass and mixing of mu/tau neutrinos from a Galactic supernova. *Astroparticle Physics*, 8(1):27–42, 1997.

[70] T. C. Shantz. Design and construction of the helium and lead observatory for supernova neutrinos, 2010.

[71] M. A. Schumaker et al. Data acquisition for the Helium and Lead Observatory. In *2010 IEEE Nuclear Science Symposium, Medical Imaging Conference, and 17th Room Temperature Semiconductor Detectors Workshop*, pages 1860–1865, 2010.

[72] S. Sekula. Personal communication, 2023.

[73] K. Zuber. HALO, a supernova neutrino observatory. *Nuclear and Particle Physics Proceedings*, 265-266:233–235, 2015. Proceedings of the Neutrino Oscillation Workshop.

[74] V. F. Sears. Neutron scattering lengths and cross sections. *Neutron News*, 3(3):26–37, 1992.

[75] C. Bruulsema. Calibration and commissioning of the Helium and Lead Observatory, 2017. `https://omni.laurentian.ca/permalink/01OCUL_LU/14p4p5m/alma991010221423705165`.

[76] G. Giacomelli et al. The Opera Experiment, 2006.

[77] N. Agafonova et al. Final Results of the OPERA Experiment on ν_τ Appearance in the CNGS Neutrino Beam. *Phys. Rev. Lett.*, 120:211801, May 2018.

[78] A. G. Rosso. Supernova neutrino fluxes in HALO-1kT, Super-Kamiokande, and JUNO. *Journal of Cosmology and Astroparticle Physics*, 2021(06):046, jun 2021.

[79] D. A. Patel. Study of Neutron Reflector for the HALO-1kT Supernovae Neutrino Detector. Master's thesis, University of Regina, 2020. `http://hdl.handle.net/10294/9348`.

[80] E. Weima. HALO-1kT prototype He-3 counters: background studies. Master's thesis, Laurentian University, 2023. `https://zone.biblio.laurentian.ca/handle/10219/4000`.

[81] J. Allison et al. Recent developments in Geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.

[82] J. Allison et al. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.

[83] S. Agostinelli et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[84] R. Brun et al. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth. A*, 389:81–86, 1997.

[85] W. Verkerke et al. The RooFit toolkit for data modeling. *eConf*, C0303241:MOLT007, 2003.

[86] J. Albert et al. SNOwGLoBES: SuperNova Observatories with GLoBES, 2018.

[87] J. Migenta. Documentation for sntools, 2022.

[88] A. L. Baxter et al. SNEWPY: A Data Pipeline from Supernova Simulations to Neutrino Signals. *The Astrophysical Journal*, 925(2):107, jan 2022.

[89] T. Totani et al. Future Detection of Supernova Neutrino Burst and Explosion Mechanism. *The Astrophysical Journal*, 496(1):216–225, mar 1998.

[90] J. Gava et al. Dynamical Collective Calculation of Supernova Neutrino Signals. *Phys. Rev. Lett.*, 103:071101, Aug 2009.

[91] J. Engel et al. What can be learned with a lead-based supernova-neutrino detector? 67:013005, Jan 2003.

[92] R. Lazauskas et al. Neutrino beams as a probe of the nuclear isospin and spin–isospin excitations. *Nuclear Physics A*, 792(3-4):219–228, aug 2007.

[93] W. Almosly et al. Theoretical estimates of supernova-neutrino cross sections for the stable even-even lead isotopes: Charged-current reactions. *Phys. Rev. C*, 94:044614, Oct 2016.

[94] W. Almosly et al. Neutral-current supernova-neutrino cross sections for ^{204,206,208}Pb calculated by Skyrme quasiparticle random-phase approximation. *Phys. Rev. C*, 99:055801, May 2019.

[95] E. Kolbe et al. Role of $\nu-$induced reactions on lead and iron in neutrino detectors. *Phys. Rev. C*, 63:025802, Jan 2001.

[96] HALO Collaboration. Internal Technical Report, 2012.

[97] A. G. Rosso. Background Sources at Gran Sasso Laboratory. Technical report, Laurentian University, 2019.

[98] J. Rumleskie. Personal communication, 2022.

[99] L. Lista. *Statistical Methods for Data Analysis in Particle Physics*, volume 909, chapter 5. Springer, 2016.

[100] F. James. MINUIT Function Minimization and Error Analysis: Reference Manual Version 94.1. 1994.

[101] NIST. Kolmogorov-Smirnov Goodness-of-Fit Test. Unpublished material. `https://www.itl.nist.gov/div898/handbook/eda/section3/eda35g.htm`.

[102] NIST. Anderson-Darling Test. Unpublished material. `https://www.itl.nist.gov/div898/handbook/eda/section3/eda35e.htm`.

[103] J. Tseng. Personal communication, 2022.

APPENDIX A

`SNEWS EveGeneSiS` code package

To support the work included in this thesis, I developed a package with the working name `SNEWS EveGeneSiS` (SNEWS [Eve]nt [Gene]ration and [Si]mulation of [S]upernova signals). I would be the first to admit that the name is mediocre at best (you are more than welcome to email me suggestions). This code, in short, with the input from a single macro file, run `SNOwGLoBES` for a given detector, distance, flavour transformation, supernova model, upper and lower times, and total bins. It was created before SNEWPY was released to the collaboration, so any attempt to use SNEWPY with `SNEWS EveGeneSiS` (as I have done) must be done manually. The package will then take the `SNOwGLoBES` output, Poisson fluctuate the number of events and assign energies, positions, and momentum vectors as described in Section 6 for HALO and HALO-1kT. These will then be put in `ROOT` files that can be simulated through HALO and HALO-1kT, which the package will do for you if the appropriate flag is set. For this thesis, all post-simulation analysis (i.e. extraction of t_0) was done in separate analysis files from the main `SNEWS EveGeneSiS` package. These analysis routines have now been fully incorporated into the pipeline and can be used freely.

Appendix B was intended to walk the reader (and other interested parties) through the process of running the analysis code in this report. Since this thesis has already ballooned in size well beyond my expectations, I have taken the liberty to upload a User Manual to my personal git page, where this code is stored. You can find my personal git page here: `https://github.com/RemiHill/SNEWS_EveGeneSiS`. You are more than welcome to make use of this code in any manner you see fit and if you have any questions, or want to contribute some changes to the code, feel free to let me know!