Precision Measurement of the Neutron Lifetime

Thesis by
Eric Michael Fries

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended June 11, 2021

Eric Michael Fries
ORCID: 0000-0002-3788-8096

ACKNOWLEDGEMENTS

I'd like to thank my wife, Melissa, for supporting me throughout the difficulties of graduate school. Without her tremendous assistance, I would not have completed my doctorate.

I'd like to thank both my family and Melissa's family for all of the support and encouragement that they gave me. It was very difficult to move across the country and to be away from my family, but Melissa's family took me in and made me feel like one of their own. Although my family was nearly 3000 miles away, they were never more than a phone call away and assisted me with great advice whenever I needed it.

I'd like to thank my advisor, Brad Filippone, for all of the advice and mentoring that he gave me. Brad always made time to help me when I needed it.

I'd like to thank my labmates at Caltech, the UCNτ collaboration, and those at the Los Alamos Neutron Science Center. My camaraderie with my colleagues made my work very enjoyable, and I'm grateful to them for working with me and teaching me so much about ultra-cold neutrons.

I'd like to thank all of those who financially supported my work: Caltech; the David and Barbara Groce travel fund; the LANL LDRD program; the United States Department of Energy, Office of Science, Office of Nuclear Physics under contract numbers DE-FG02-ER41042, DE-AC52-06NA25396, DE-AC05-00OR2272, and 89233218CNA000001 under proposal LANLEEDM; NSF grants 1614545, 1914133, 1506459, 1553961, 1812340, 1714461, and 1913789; and the NIST precision measurements grant.

ABSTRACT

The neutron lifetime plays a critical role in Big Bang nucleosynthesis (BBN) calculations, and measurements of the neutron lifetime can also be used to probe the unitarity of the Cabibbo-Kobayashi-Maskawa quark weak mixing matrix. Most experiments that measure the neutron lifetime fall into two classes: "bottle" and "beam" experiments. A bottle experiment stores neutrons in a bottle and counts the number of neutrons that do not decay. A beam experiment counts the decay products of a beam of neutrons that passes through an electrostatic trap. An unresolved $\approx 4\sigma$ difference remains between the current global averages of measurements of the neutron lifetime using the bottle method and measurements using the beam method. This difference is the dominant uncertainty in BBN calculations of the helium mass fraction in the early universe. The UCNτ experiment is a bottle experiment which uses a magneto-gravitational trap to store ultra-cold neutrons (UCN) without any physical interactions between the UCN and the walls of the trap. The UCN that do not decay are counted with an *in situ* detector that is lowered into the trap. These two features stand in contrast to most past bottle experiments, which had to make significant corrections to the extracted lifetime to account for losses of neutrons due to material interactions with walls and losses while removing the neutrons from the bottle in order to be counted. This thesis will present an analysis of the 2017-2018 UCNτ data set that extracted a value for the neutron lifetime of 877.79 ± 0.27 (stat.) $^{+0.19}_{-0.12}$ (sys.) s. This measurement has an uncertainty of roughly half of the current global average for the neutron lifetime.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

1.1 Lifetime of a free neutron

A neutron is the bound state of one up quark (u) and two down quarks (d). Neutrons are usually found in the nuclei of atoms, and in many cases, atomic nuclei are stable and will not decay. However, free neutrons can undergo β-decay to generate a proton, an electron, and an electron antineutrino, as shown in Figure 1.1. This decay is energetically favorable and releases $Q = 0.782$ MeV of energy [1], which implies that free neutrons will decay with a finite lifetime.

The amplitude of the tree-level process shown in Figure 1.1 can be calculated using Feynman's rules for weak interactions. The vertex term for the electron-neutrino interaction is

$$\mathcal{J}^{\mu} = -\frac{ig_w}{\sqrt{8}}\gamma^{\mu}(1-\gamma^5), \tag{1.1}$$

where g_w is the weak coupling constant.

The vertex term for the quark-quark interaction is identical, except for two modifications: the first to account for the weak mixing of quarks of different generations, and the second to account for the presence of the two spectator quarks. The quark-quark vertex term is

$$\mathcal{K}^{\mu} = -\frac{ig_w}{\sqrt{8}}V_{ud}\gamma^{\mu}(1-\lambda\gamma^5), \tag{1.2}$$

where V_{ud} is the first element of the Cabibbo–Kobayashi–Maskawa quark weak mixing matrix (further discussed in Section 1.2) and λ is a measure of the relative scales of the vector and axial vector components of the quark-quark vertex term. Interactions between the quarks and gluons in the neutron modify the axial vector component of the quark-quark vertex term, but this modification is absorbed into λ.

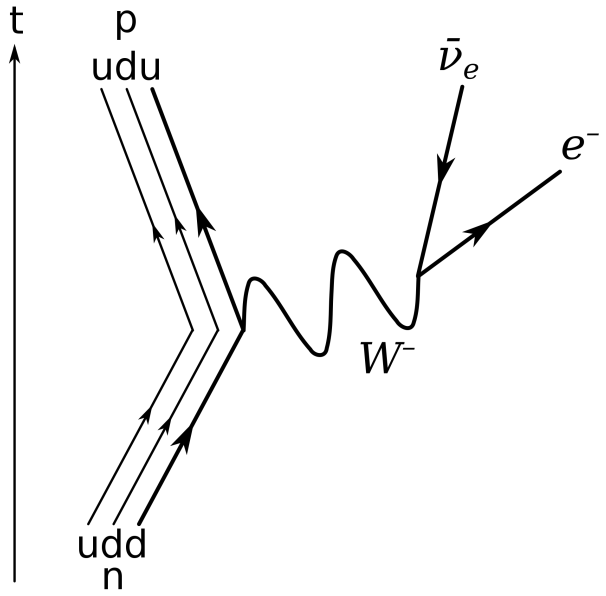

Figure 1.1: Feynman diagram for the tree-level approximation of the β-decay of a free neutron.

In the low-momentum limit ($Q \ll m_W c^2 = 80.379$ GeV [1]), the propagator term can be approximated as $\frac{ig_{\mu\nu}}{m_W^2}$. Then the amplitude of the tree-level process is

$$\begin{aligned}\mathcal{M} &= i\bar{\psi}_n \mathcal{K}_\mu \psi_p \frac{ig_{\mu\nu}}{m_W^2} \bar{\psi}_e \mathcal{J}^\nu \psi_{\nu_e} \\ &= \frac{g_w^2 V_{ud}}{8m_W^2} \bar{\psi}_n \gamma_\mu (1 - \lambda\gamma^5) \psi_p \bar{\psi}_e \gamma^\mu (1 - \gamma^5) \psi_{\nu_e},\end{aligned} \tag{1.3}$$

where ψ_{n,p,e,ν_e} are spinors for the neutron, proton, electron, and electron antineutrino.

Fermi's Golden Rule can be used to calculate a decay lifetime

$$\frac{1}{\tau} = \frac{m_e^5 c^2 g_W^4 \left|V_{ud}\right|^2}{64\pi^3 \hbar m_W^4} (1 + 3\lambda^2) f, \tag{1.4}$$

where $f = 1.6887$ is a phase-space statistical factor [2]. Equation 1.4 can be expressed in terms of the Fermi constant

$$G_F \equiv \frac{g_W^2 \hbar^3}{\sqrt{32} m_W^2 c}$$

as

$$\frac{1}{\tau} = \frac{m_e^5 c^4 \left|V_{ud}\right|^2 G_F^2}{2\pi^3 \hbar^7} (1 + 3\lambda^2) f. \tag{1.5}$$

Using values for $m_e, c, |V_{ud}|, G_F, \hbar$, and λ from [1] and the value for f from [2], this theoretical calculation results in a value of $\tau_{\text{calc}} = 915$ s. This differs from the current global average of measurements of the free neutron lifetime of $\tau_{\text{meas}} = 879.4 \pm 0.6$ s [1] by $\simeq 4\%$. This difference is mostly due to ignoring higher-order terms in the neutron β-decay process. Including these radiative corrections decreases τ by a factor of $1 + \text{RC} = 1.03947 \pm 0.00032$ [2], which leads to a calculated value for the lifetime of $\tau_{\text{calc,RC}} = 880.0 \pm 1.5$ s. This is within the statistical uncertainty of τ_{meas}. The purpose of the UCNτ experiment is to precisely measure the lifetime of the free neutron.

1.2 Unitarity of the Cabibbo–Kobayashi–Maskawa quark weak mixing matrix and the neutron lifetime

The Cabibbo–Kobayashi–Maskawa (CKM) quark weak mixing matrix quantifies the relationship between the mass and flavor eigenstates of the down-type quarks [3]. Specifically,

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}, \tag{1.6}$$

where the 3×3 matrix is the CKM matrix, d', s', and b' are the flavor eigenstates of the down-type quarks, and d, s, and b are the mass eigenstates of the down-type quarks. Using the CKM matrix to translate between the flavor and mass eigenstates should neither create nor annihilate quarks, so theoretically the CKM matrix should be unitary. However, new physics beyond the Standard Model (BSM) could modify the way that the elements of the CKM matrix are extracted from experimental measurements. An apparent non-unitary CKM matrix would imply BSM physics.

Equation 1.5 relates $\tau, |V_{ud}|$, and λ. After substituting in the current global averages for measurements for all other terms, Equation 1.5 reduces to

$$|V_{ud}|^2 = \frac{(4905.7 \pm 1.5) \text{ s}}{\tau(1 + 3\lambda^2)}. \tag{1.7}$$

In contrast, the current global averages for measurements of τ, λ, and $|V_{ud}|$ yield $\tau |V_{ud}|^2 \left(1 + 3\lambda^2\right) = 4902.4 \pm 9.1$ s. The uncertainty of τ contributes 3.3 s of the 9.1 s of uncertainty of $\tau |V_{ud}|^2 \left(1 + 3\lambda^2\right)$, and the uncertainty of λ contributes almost all of the rest of the uncertainty of $\tau |V_{ud}|^2 \left(1 + 3\lambda^2\right)$. Therefore, improved measurements of λ are necessary in order to probe the unitarity of the CKM matrix through measurements of τ.

1.3 Big Bang nucleosynthesis and the neutron lifetime

Big Bang nucleosynthesis (BBN) is the process by which nuclei with more than one nucleon were generated in the early universe. Initial temperatures and densities were sufficiently high that the proton-to-neutron ratio was held at roughly 1:1 via three nuclear reactions:

$$\begin{aligned} \mathrm{n} + \nu_e &\rightleftharpoons \mathrm{p} + \mathrm{e}^-, \\ \mathrm{n} + \mathrm{e}^+ &\rightleftharpoons \mathrm{p} + \nu_e, \text{ and} \\ \mathrm{n} &\longrightarrow \mathrm{p} + \mathrm{e}^- + \overline{\nu}_e. \end{aligned} \tag{1.8}$$

The first two of these nuclear reactions occurred in both directions in the very early universe, but as the temperature of the universe cooled, the nuclear reactions began to favor the generation of protons instead of neutrons. About 1 s after the Big Bang, the temperature of the universe had sufficiently decreased to "freeze out" the first two nuclear reactions. At this point the ratio of protons to neutrons was approximately 6:1. Over the next $\sim$ 200 s, the temperature of the universe continued to decrease and heavier nuclei began to form. During these $\sim$ 200 s, some of the neutrons underwent β-decay and generated more protons. Those decays shifted the ratio of protons to neutrons from 6:1 to 7:1. The exact ratio of protons to neutrons at this time drove the relative proportions of heavier nuclei that were generated after $\sim$ 200 s.

BBN theory predicts that $\sim$ 75% of the baryons remained as free protons, $\sim$ 25% were bound into ^{4}He nuclei, and the remaining few baryons formed trace amounts of ^{2}H, ^{3}He, and nuclei with $A > 4$. The helium mass fraction Y_p is defined as the fraction of all the baryons that were bound into ^{4}He nuclei. BBN calculations of the mass fractions of the various nuclei are all sensitive to the baryon density in the early universe (ω_b) and the neutron lifetime. Equation 1.9 shows the relationship between various mass fractions and ω_b and τ,

$$\begin{aligned} Y_p &\propto \left(\frac{\omega_b}{\omega_{b,0}}\right)^{+0.39} \left(\frac{\tau}{\tau_0}\right)^{0.72}, \\ ^2\mathrm{H/H} &\propto \left(\frac{\omega_b}{\omega_{b,0}}\right)^{-1.62} \left(\frac{\tau}{\tau_0}\right)^{0.41}, \\ ^3\mathrm{He/H} &\propto \left(\frac{\omega_b}{\omega_{b,0}}\right)^{-0.59} \left(\frac{\tau}{\tau_0}\right)^{0.15}, \text{ and} \\ ^7\mathrm{Li/H} &\propto \left(\frac{\omega_b}{\omega_{b,0}}\right)^{+2.12} \left(\frac{\tau}{\tau_0}\right)^{0.44}, \end{aligned} \tag{1.9}$$

where $\omega_{b,0}$ and τ_0 are fiducial values [4].

1.4 Past measurements of the neutron lifetime

1.4.1 Bottle lifetime experiments

There are two main experimental approaches that are used to measure the lifetime of a free neutron. The first approach is referred to as a "bottle" experiment. Bottle experiments store ultracold neutrons (UCN) in a bottle for varying lengths of time, and the UCN that remain in the bottle after the storage time are counted. The number of UCN that do not decay during the storage time is modeled as

$$\langle N(t) \rangle = N_0 e^{-t/\tau}, \tag{1.10}$$

where N_0 is the number of UCN in the bottle at time $t = 0$.

If two iterations of this experiment are done with different holding lengths (a short holding length t_S and a long holding length t_L) and the number of UCN that remain in the bottle at the end of the iterations are $N(t_S) = N_S$ and $N(t_L) = N_L$, then the lifetime can be extracted as

$$\tau = \frac{t_L - t_S}{\ln \frac{N_S/N_0}{N_L/N_0}} = \frac{t_L - t_S}{\ln \frac{N_S}{N_L}}. \tag{1.11}$$

The derivation of Equation 1.11 assumes that the number of UCN in the bottle at the beginning of the each iteration (N_0) is constant. In this case, the extracted lifetime is not dependent on the actual value for N_0. Furthermore, if the counting of surviving UCN is $r < 100\%$ efficient, the extracted lifetime would remain unbiased as long as r is the same for N_S and N_L.

In practice, different iterations usually have different values for N_0. In that case, the extracted lifetime is

$$\tau = \frac{t_L - t_S}{\ln \frac{N_S/N_{0,S}}{N_L/N_{0,L}}} = \frac{t_L - t_S}{\ln \frac{N_S N_{0,L}}{N_L N_{0,S}}}. \tag{1.12}$$

In order to extract a lifetime using Equation 1.12, there must be some method of estimating how many UCN are in the bottle at the beginning of the iteration. Unfortunately, the mechanism used to count the UCN in the bottle removes the UCN from the bottle. Therefore, the number of UCN in the bottle at the beginning of each iteration is estimated instead of directly measured. This estimation will be discussed in Sections 2.7.3 and 3.6. Figure 1.2 shows a simplified version of how a generic bottle experiment measures the neutron lifetime in the simplified case where $N_0 = N_{0,S} = N_{0,L}$.

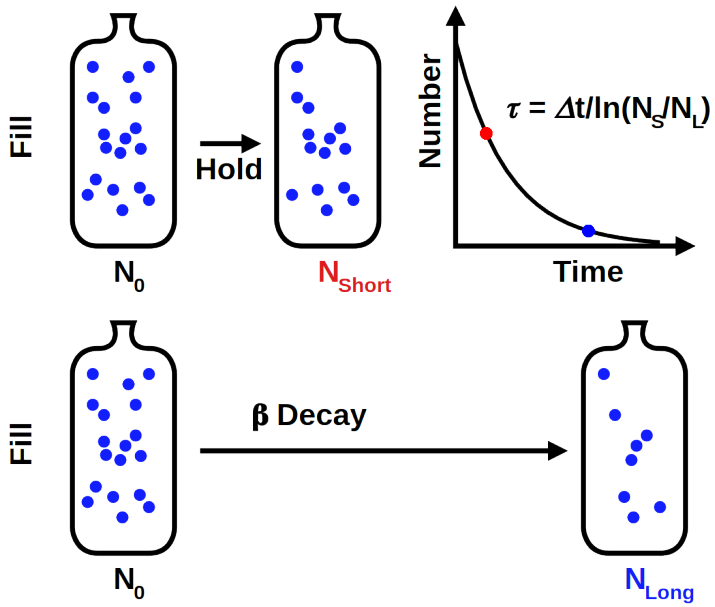

Figure 1.2: A simplified version of how a generic bottle experiment measures the neutron lifetime.

1.4.2 Beam lifetime experiments

The second approach that is used to measure the neutron lifetime is referred to as a "beam" experiment. In a beam experiment, a directed beam of cold neutrons are sent through an electrostatic trap. Most of the neutrons pass through the trap, but some undergo β-decay while in the trap. The protons generated by β-decay are trapped by the electric field, and later these protons are removed from the trap and counted. Figure 1.3 shows a schematic of the beam experiment at the National Institute of Standards and Technology's Center for Neutron Research (NCNR).

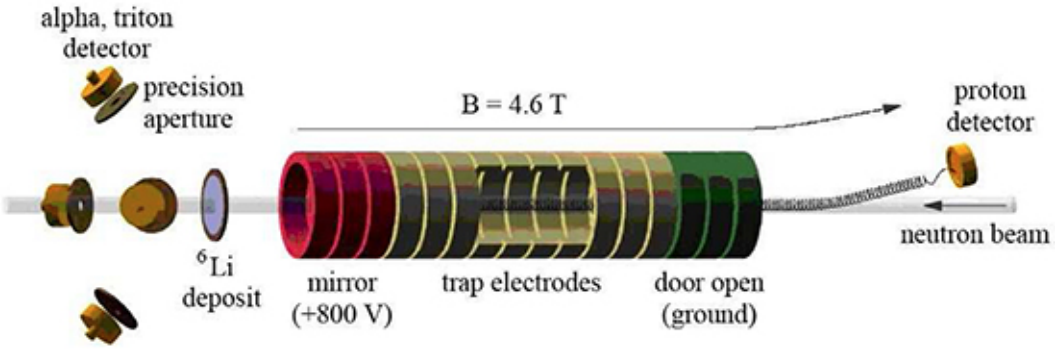

Figure 1.3: A schematic of the NCNR beam experiment used to measure the neutron lifetime [5]. This schematic shows the experiment in a state where trapped protons are being flushed out through the door to the proton detector.

An iteration of the NCNR experiment can be described as a three-step process:

1. The "door" and "mirror" electrodes are held at +800 V for 10 ms as a neutron beam passes through the electrostatic trap. This voltage is sufficiently high to trap all protons that would be generated by the β-decay of a neutron that is currently inside of the trap. Neutrons that do not decay pass through the trap and are captured by ^{6}Li via ^{6}Li(n, α)^{3}H. The α and the ^{3}H are counted by the detectors that observe the ^{6}Li deposit.

2. The door electrode is then "opened" by grounding the electrode. A slight graduated potential is applied to the trap electrodes to flush any trapped protons out through the door. These protons are counted by the proton detector over 76 μs.

3. The mirror and door electrodes are grounded for 33 μs (while maintaining the slight graduated potential on the trap electrodes) to force any charged particles to escape from the electrostatic trap.

The fluence rate of the neutron beam $I(v)$ is dependent on the velocity of neutrons in the beam. The time that a neutron spends within the trap is inversely proportional to the velocity of that neutron, so the expected number of neutrons that pass through the trap is

$$N_n = L \int I(v) v^{-1} dv. \tag{1.13}$$

Therefore, the rate at which the trap captures protons is

$$\dot{N}_p = \epsilon_p \times \tau^{-1} N_n = \frac{\epsilon_p L}{\tau} \int I(v) v^{-1} dv, \tag{1.14}$$

where ϵ_p is the efficiency of the trap at capturing protons generated by the β-decay of neutrons inside of the trap. The neutron capture cross section on the ^{6}Li deposit is inversely proportional to the velocity of the neutron. Thermal neutrons with a fixed velocity of $v_0 = 2200$ m/s are used to calibrate the total neutron detection efficiency ϵ_0 at $v = v_0$. The velocity-dependent efficiency of detecting neutrons is thus modeled as $\epsilon_{\mathrm{Li}}(v) = \epsilon_0 v_0 / v$ and the neutron detection rate for the ^{6}Li deposit is

$$\dot{N}_{\mathrm{Li}} = \int I(v) \times \frac{\epsilon_0 v_0}{v} dv = \epsilon_0 v_0 \int I(v) v^{-1} dv. \tag{1.15}$$

Equations 1.14 and 1.15 contain identical integrals. Setting these two integrals equal to each other results in a formula to extract the lifetime via

$$\tau = \frac{\dot{N}_{\mathrm{Li}}}{\dot{N}_p} \frac{\epsilon_p}{\epsilon_0} \frac{L}{v_0}. \tag{1.16}$$

A lifetime is measured by precisely measuring $\dot{N}_{\mathrm{Li}}, \dot{N}_p, \epsilon_p, \epsilon_0$, and L. The two different efficiencies in Equation 1.16 are efficiencies from two different detectors so their effects on the extracted lifetime are not suppressed. This stands in contrast to the efficiency in Equation 1.11, which is suppressed because a bottle experiment uses the same detector to count the UCN remaining in the trap after short and long runs.

1.4.3 Differences between bottle and beam lifetime experiments

Figure 1.4 shows the history of measurements of the neutron lifetime. The current global averages of lifetimes measured using the bottle and beam methods are $\tau_{\mathrm{bottle}} = 879.4 \pm 0.6$ s and $\tau_{\mathrm{beam}} = 888.0 \pm 2.0$ s [1]. The difference between these values is $> 4\sigma$ so it is extremely unlikely that the difference is due to statistical fluctuations alone. Measuring τ_{bottle} and τ_{beam} with increasing precision will likely have marginal-to-no impact on resolving the discrepancy between τ_{bottle} and τ_{beam}. A search for flaws in the two experimental methods is needed to better understand this discrepancy.

One plausible source for bias in bottle lifetime experiments is an underestimation of the rate at which UCN escape the bottle. This underestimation would cause τ_{bottle} to be less than it should be. An unaccounted for escape lifetime of $\sim 10^5$ s would be sufficient to explain the discrepancy between τ_{bottle} and τ_{beam}. A potential source for bias in beam lifetime experiments is mis-estimation of the efficiency of the detection of cold neutrons or the detection of protons generated by neutron decay. A 1% error in the estimation of either of these efficiencies would be sufficient to explain the discrepancy between τ_{bottle} and τ_{beam}.

Another potential explanation for the discrepancy between τ_{bottle} and τ_{beam} would be the existence of a significant decay channel other than the standard β-decay process of $n \rightarrow pe^-\overline{\nu}_e$. Bottle experiments measure the number of neutrons that **do not decay**, so bottle experiments are sensitive to any neutron decay channel. In contrast, beam experiments measure the number of free protons from neutrons that **do decay**, so beam experiments are not sensitive to neutron decay channels that do not generate free protons. A branching ratio of $\Gamma \simeq 10^{-2}$ to neutron decay channels that do not generate a free proton would be sufficient to explain the discrepancy between τ_{bottle} and τ_{beam}. One recent proposal hypothesized decay channels $n \rightarrow X\gamma$ and $n \rightarrow Xe^+e^-$, where X is a dark matter particle [28]. However, experimental observations ruled out the existence of these decay channels with a 1% branching ratio at levels of $> 4\sigma$ for $n \rightarrow X\gamma$ [29] and $\gg 5\sigma$ for $n \rightarrow Xe^+e^-$ [30].

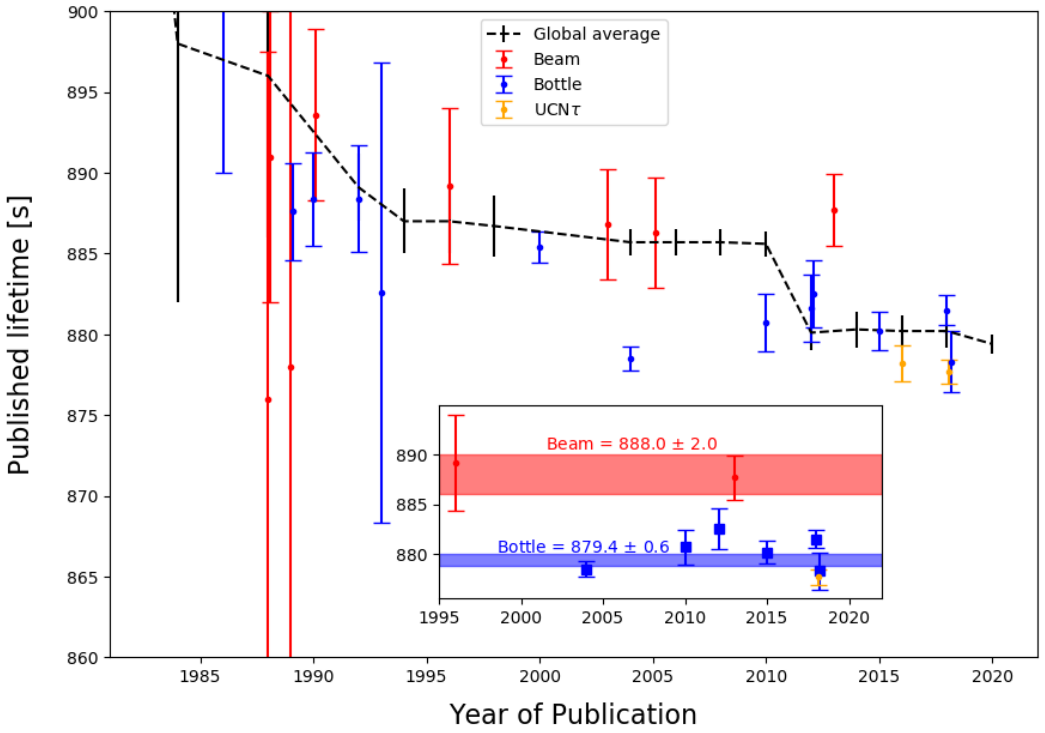

Figure 1.4: History of measurements of the neutron lifetime [1, 5–27]. Inset: modern measurements of the neutron lifetime, as well as the current global averages for lifetimes measured using the beam and bottle methods. Over time, the global average value has steadily decreased. In 2012, the global average shifted from combining multiple beam measurements to combining multiple bottle measurements. Since then, multiple measurements of the lifetime by bottle experiments, including those published by the UCNτ collaboration, have been in good agreement with the global average.

1.4.4 Implications of resolving the discrepancy between beam and bottle lifetimes

Section 1.2 discussed how a more precise measurement of λ would allow for a measurement of τ to be used a probe of CKM unitarity. The relationship between the current global averages for measurements of τ, λ, and certain components of the CKM matrix [1] is shown in Figure 1.5. If τ_{beam} is ignored, then the measured values for τ_{bottle} and λ are in good agreement with both of the values for $|V_{ud}|$ shown in Figure 1.5. However, the values for $|V_{ud}|$ directly measured by studying $0^+ \rightarrow 0^+$ decays and the value for $|V_{ud}|$ indirectly calculated by measuring $|V_{us}|$ and $|V_{ub}|$ and demanding unitarity of the CKM matrix differ by $\simeq 3.3\sigma$. Recent updates to radiative corrections [31, 32] are included in the values for $|V_{us}|$ and $|V_{ub}|$ used to calculate the value of $|V_{ud}|$ that satisfies CKM unitarity.

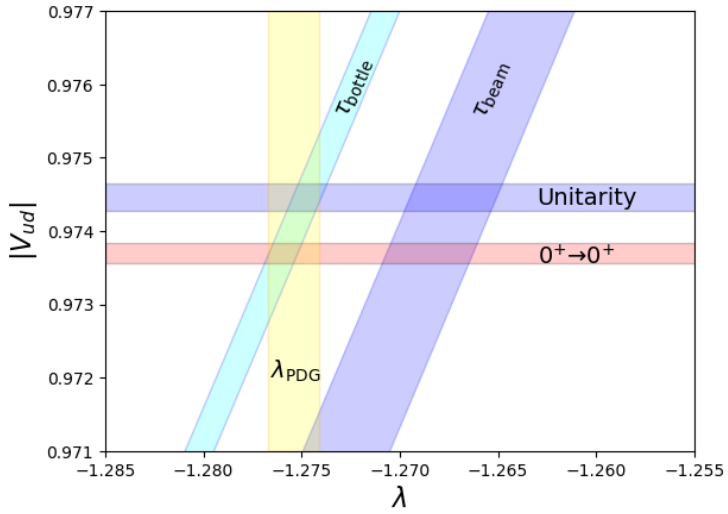

Figure 1.5: The relationship between the current global averages for measurements of τ, λ, and certain components of the CKM matrix [1]. Each band denotes 1σ uncertainties of the measured quantity. The "Unitarity" band is determined by using the current global averages for $|V_{us}|$ and $|V_{ub}|$ and demanding that $1 = |V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2$. The "$0^+ \to 0^+$" band is the current global average of measurements of $|V_{ud}|$ made by studying superallowed $0^+ \to 0^+$ nuclear β-decays. τ_{bottle} and τ_{beam} are calculated using Equation 1.7.

Section 1.3 discussed the role that the neutron lifetime plays in BBN. The current global average for the baryon density in the early universe is $\omega_b = 0.02230 \pm 0.00020$ [1]. Table 1.1 shows how the uncertainties of ω_b and τ contribute to uncertainties of various mass fractions calculated using Equation 1.9. The discrepancy between τ_{bottle} and τ_{beam} causes a much larger uncertainty of Y_p than the uncertainty of ω_b. However, rejecting τ_{beam} as incorrect would result in the uncertainty of τ being a negligible contribution to the uncertainty of theoretical calculations of Y_p. Therefore, an understanding of the difference between τ_{beam} and τ_{bottle} is necessary to decrease the uncertainty of BBN calculations of Y_p.

Shifts in physical constants	Y_p	^{2}H/H	^{3}He/H	^{7}Li/H
A shift from τ_{bottle} to τ_{beam}	0.7%	0.4%	0.1%	0.4%
1σ shift in ω_b	0.3%	1.4%	0.5%	1.9%
1σ shift in τ_{bottle}	0.05%	0.03%	0.01%	0.03%

Table 1.1: Shifts in the BBN calculations of the mass fractions of elements formed during BBN due to shifts in the values of ω_b and τ. The uncertainty of Y_p is dominated by the discrepancy between τ_{bottle} and τ_{beam}, but the uncertainties of all other mass fractions are dominated by the uncertainty of ω_b.

Chapter 2

UCNτ EXPERIMENTAL DESIGN

2.1 Properties of ultra-cold neutrons (UCN)

The defining property of UCN are that they are neutrons that can be stored in material bottles. The primary method of interaction between UCN and material walls is through strong interactions with the nuclei of the atoms that make up the walls. The interaction between a UCN and a single nucleus that is fixed in location is

$$V_F(\boldsymbol{r}) = \frac{2\pi\hbar^2}{m_n} a\delta^{(3)}(\boldsymbol{r}), \tag{2.1}$$

where m_n is the mass of a neutron and a is the neutron-nucleus scattering length. The scattering length is determined experimentally. Some materials have a positive scattering length and can serve as material walls, while other materials have a negative scattering length and accelerate incident UCN into the material. Following the derivation in Sections 2.3.2 and 2.3.3 of [33], the interaction between a UCN and a material wall made of nuclei with scattering length a can be considered as an interaction between a neutron and a potential step with height

$$V_F = \frac{2\pi\hbar^2}{m_n} Na, \tag{2.2}$$

where N is the number density of the nuclei in the material wall. This potential is often referred to as the Fermi potential of a material. UCN with $E < V_F$ are completely reflected by the material wall. ^{58}Ni is the material with the highest known Fermi potential of $V_F = 335$ neV, which establishes the energy threshold for neutrons to be considered UCN.

A neutron with $E = 335$ neV has a velocity of $\simeq 8$ m/s and can rise $\simeq 3.25$ m (about 1 cm/neV) through Earth's gravitational field before reaching its turning point, so experiments that store UCN have a size on the scale of a few meters. Typical UCN sources can only produce UCN densities of < 1000 UCN/cm^3, so the motion of UCN are best modeled using simple kinematics and ballistic trajectories instead of kinematic gas theory. However, if kinematic gas theory was used to model UCN the temperature of the UCN would be < 4 mK (hence, the name: "ultra-cold" neutrons).

The neutron is a spin one-half particle, so it can have a non-zero magnetic moment. The potential energy in the interaction between a neutron's magnetic moment $\boldsymbol{\mu}$ and an external magnetic field $\boldsymbol{B}$ is

$$V = -\boldsymbol{\mu} \cdot \boldsymbol{B}. \tag{2.3}$$

The maximum potential energy in this interaction is

$$V_{\max} = |\mu B| \,. \tag{2.4}$$

Given that the maximum energy of UCN is 335 neV, a > 5.6 T magnetic field can completely polarize UCN by their spin.

These properties of UCN make them useful for precision measurements of β-decay parameters. Section 2.2 will discuss how an experimental apparatus with a size on the scale of $\sim$ meters is used to measure the lifetime of the free neutron. Section 2.4 will discuss how the extremely low energy of UCN allows them to be transported from their source to the experimental apparatus by material guides. Section 2.5 will discuss how UCN are polarized by spin state. Section 2.6 will show why the polarization of UCN is essential for the UCNτ experiment, and will also discuss why it is important that UCN can only rise a short distance through Earth's gravitational field before they reach their turning point.

2.2 Overview

This section provides an overview of the various components of the UCNτ experiment and how they work together, and subsequent sections will provide detailed descriptions of the various components. Figure 2.1 shows a CAD drawing of the bulk of the UCNτ experimental apparatus. The source that produces UCN will be described in Section 2.3. UCN produced by the UCN source (not shown in Figure 2.1) diffuse through the guides to the UCN polarizer. The guides that transport UCN will be described in Section 2.4. The UCN polarizer selects UCN with the correct spin state to be stored in the UCN trap. The UCN polarizer will be described in Section 2.5 and the UCN trap will be described in Section 2.6. The various types of monitors that detect UCN will be discussed in Section 2.7. Section 2.8 will discuss how high-energy UCN are removed from the trap. Section 2.9 will discuss an upgrade made to the experiment between the 2017 and 2018 run cycles that smooths out fluctuations in the production of UCN.

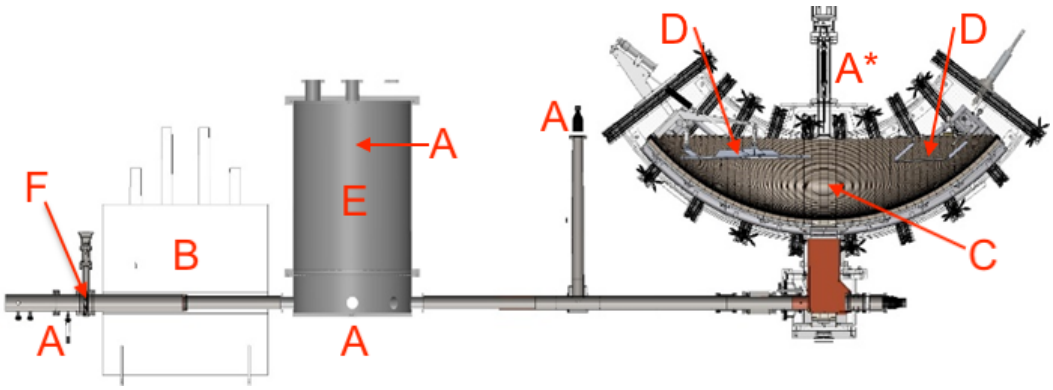

Figure 2.1: A CAD drawing of the UCNτ experimental apparatus. The UCN are produced in the source (left of figure) and diffuse rightward. A: various UCN detectors. The one detector marked A* is the primary detector and all other detectors are upstream monitors. B: UCN polarizer. C: UCN trap. D: trap cleaners. E: buffer volume. F: gate valve. The various components of the system are all connected with UCN guides (unlabeled). The primary detector is used to precisely count UCN in the trap. The upstream monitors are used to measure the local density of UCN in the guides. The UCN polarizer is used to select only one spin state of UCN, which is necessary because the UCN trap can only store one spin state of UCN. Part of UCN polarizing apparatus that is located between B and E is not shown in this figure. The two cleaners remove high-energy UCN from the trap. The buffer volume, which was installed between the 2017 and 2018 run cycles, smooths out fluctuations in the production of UCN. The gate valve can be closed to separate the UCN source from the rest of the UCNτ experiment.

All runs that measure the lifetime follow the same pattern:

1. Production of UCN begins with spallation neutrons that are produced from a 800 MeV proton beam incident on a W target. Simultaneously, the gate valve opens which allows UCN to diffuse from the source to the trap. UCN enter the trap through an opening in the bottom of the trap. The trap is filled with UCN while the cleaners and the primary detector are lowered to the cleaning height.

2. The gate valve and the opening in the bottom of the trap close, and production of UCN ceases. The primary detector is raised to the top of the trap, but the cleaners remain at the cleaning height for an extended period of time.

3. The cleaners are raised. UCN are held in the trap for various lengths of time, between 20 s and 5000 s.

4. The primary detector is lowered to count all of UCN that remain in the trap.

2.3 Production of UCN

The UCNτ experiment is located within the Los Alamos Neutron Science Center (LANSCE). LANSCE contains a linear accelerator which provides 800 MeV protons to multiple experiments. UCNτ uses these protons to produce neutrons via spallation from a W target. In typical running conditions, the proton beam delivers $\simeq$ 10 pulses of beam for a total of $\simeq 3 \times 10^{14}$ protons-on-target during bursts that last $\simeq$ 0.5 s. The W target is cooled for $\simeq$ 5 s before the next burst is delivered to the W target. The beginning of the first burst is synchronized with the beginning of each run of the experiment. UCNτ shares a branch of beamline with other non-UCN experiments, so UCN are typically unable to be produced between 6 AM and 8 PM on weekdays.

The neutrons produced by spallation are highly energetic ($E_n > 1$ MeV) and need to be cooled in order to become UCN. This cooling takes place in three steps:

1. Graphite and Be that surround the spallation target thermalize the spallation neutrons to room temperature and reflect them back toward the solid deuterium (sD_2) UCN source.

2. Cold ($\simeq$ 100 K) polyethylene beads that are between the Be and the sD_2 UCN source further cool the thermal neutrons to cold neutron energy levels.

3. Cold neutrons interact with the sD_2 UCN source and downscatter to UCN levels.

The sD_2 is cooled to $\sim$ 4 K with liquid helium. The W target is cooled by the exhaust helium gas from the sD_2 cooling loop. While the 800 MeV proton beam is being delivered to the W target the polyethylene beads warm up to $\lesssim$ 150 K [34].

The sD_2 is frozen in place inside of guides that are used to transport UCN from the source to the rest of the experiment. These guides will be described in Section 2.4. Cold neutrons are able to pass through the walls of the guides, but UCN are not. UCN that are produced inside of the sD_2 receive a 109 neV boost when they leave the sD_2 source. The UCN guides travel directly upward for $\simeq$ 1 m before they turn and exit the biological shield. This $\simeq$ 1 m of vertical travel opposite the direction of Earth's gravitational field partially cools the UCN produced by the sD_2 source.

Although the sD_2 UCN source is used to coherently downscatter cold neutrons to UCN energy levels, it can also incoherently upscatter UCN to cold neutron energy levels of $E > 335$ neV. In order to limit these losses, a "butterfly valve" is installed above the sD_2 UCN source. The control of the butterfly valve is synchronized to the timing of the beam bursts so that the butterfly valve opens slightly before the beginning of the beam burst and closes slightly after the end of the beam burst. The design of this source is shown in Figure 2.2. The entire source is housed under $\simeq 2$ m of steel surrounded by $\simeq 2$ m of concrete in order to enable experimenters to work in the experimental area while the beam is delivered to the target.

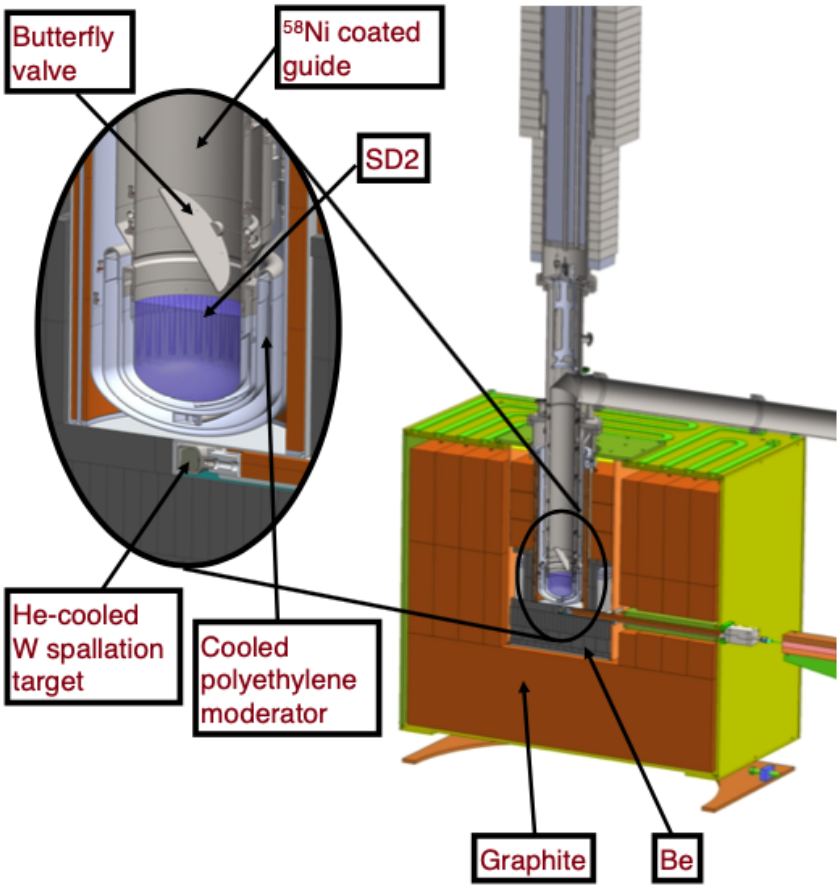

Figure 2.2: A CAD drawing of the LANSCE UCN source [34]. Neutrons are produced by spallation from the W target that is caused by protons that come into the source from the right. The graphite, Be, and cooled polyethylene moderator are used to cool the spallation neutrons to cold neutron energy levels. The sD_2 is used to downscatter cold neutrons to UCN energy levels. The butterfly valve (shown open) allows UCN to pass from the source to the guides, but when closed, it prevents UCN in the guides from upscattering on the sD_2 source and being lost from the guides. The ^{58}Ni coated guide transports UCN away from the source.

Under optimal conditions, the density of UCN that exit the biological shield is 184 ± 32 UCN/cm^3 [34]. However, many factors cause this rate to vary:

1. The current delivered by the proton beam varies due to various problems with operating the beamline. A decrease in the current delivered to the W target reduces the number of UCN produced.

2. The spin states of the two nuclei in the D_2 can either be aligned (orthodeuterium) or anti-aligned (paradeuterium). Orthodeuterium is the ground state of D_2, and the energy of paradeuterium is $\sim 10^7$ neV above the energy of orthodeuterium. Paradeuterium can down-convert to orthodeuterium by up-scattering a UCN to cold neutron energy levels. Prior to freezing the D_2 into place, it is run through a converter to convert as much as possible into orthodeuterium. Typically the paradeuterium contamination is on the order of $\sim 2\%$, but the level of contamination can vary.

3. The D_2 supplied by the manufacturer has some HD contamination, which has a large capture cross section for H (n, γ) D. Over $\gtrsim$ months there is a measured increase in the level of HD contamination. The cause of this increase is unknown.

4. While the source was operating, protons are delivered to the W target for $\simeq 0.5$ s, followed by $\simeq 5$ s without protons before the process repeats. This repetitive temperature cycling of the source causes layers of frost to form on the surface of the sD_2. This frost slowly accumulates and degrades the production of UCN [35] over timescales of days.

5. The quality of the bulk of the sD_2 changes over time. Ideally, the D_2 is slowly introduced to the source and freezes into one solid mass. However, this process is difficult to control and requires freezing in the D_2 over ~ 8 hours. Often the sD_2 is frozen in with some deformities instead of as a homogeneous mass. Furthermore, occasionally the sD_2 source spontaneously melts, partially vaporizes, and then refreezes during periods where protons are being delivered to the W target. If the bulk of the sD_2 source is in a very good state when a spontaneous excursion occurs, then the production of UCN is significantly degraded. If the bulk of the sD_2 source is in a poor state, then a spontaneous excursion has only a minor effect on the production of UCN.

The first item on the above list is completely out of the control of the UCNτ experiment. The second and third items on the list can be controlled, but the conversion from paradeuterium to orthodeuterium takes $\sim$ 1 week and the monetary budget for purchasing fresh D_2 is limited. The fourth and fifth items on the list can be fixed if the sD_2 source is reconditioned. A full reconditioning requires the sD_2 to be completely vaporized, removed from the guides, and slowly refrozen into the source. This process takes $\sim$ 8 hours to complete so it is not an efficient use of time to do this while the beam is available to produce UCN. However, this is done during the day on most weekdays while the beam is unavailable for use by UCNτ and UCN can not be produced. When the source needs reconditioning but the beam is available, the sD_2 source is quickly reconditioned by turning off the liquid helium cooling for $\sim$ 2 minutes, and then letting the source stabilize for $\sim$ 15 minutes. The goal of this quick reconditioning is to melt the layers of sD_2 frost on the surface of the source and to then freeze the melted D_2 back into place. This quick reconditioning does not improve the UCN production as much as a full reconditioning process but can be done without disrupting data gathering while the beam is available.

The sD_2 frost produced by the temperature cycling described in the fourth item in the above list not only decreases the overall number of UCN produced, but also hardens the energy spectrum of the UCN that are produced. This hardening has a small, but measurable, effect on the number of UCN loaded into the trap. This effect is significant to the experiment because the trap can only store the low-energy part of the spectrum of UCN that are produced by the source. The UCN source produces UCN with $E < 335$ neV, but the trap is only able to store UCN with $E \lesssim 60$ neV, measured relative to the height at which the guides exit the biological shield. This effect will be further discussed in Section 3.6. Both the quick and full reconditionings mostly reset the spectral hardening.

2.4 Transport of UCN

As discussed in Section 2.1, the defining characteristic of UCN are their ability to be stored in material bottles. The same materials that make good UCN bottles also make good UCN transport guides. Material guides connect the sD_2 UCN source to the UCNτ experiment, and the UCN diffuse through the guides.

Recall the overview of the Fermi potential V_F from Section 2.1. In practice, nuclei in a material wall have a non-zero cross section for upscatter or absorption of UCN. To quantify this, the potential V_F is updated as $U = V + iW$, where W quantifies the

loss probability for an interaction between a UCN and the material wall. Typical materials used for UCN material walls have $10^{-5} \lesssim |W/V| \lesssim 10^{-4}$, so the effect of the introduction of W on the reflection amplitude can be calculated with perturbation theory. These losses are small enough to not be a concern for the transport of UCN, but would significantly impact a lifetime measured using a trap with material walls.

2.5 Polarization of UCN

Section 2.6 will describe the trap that is used to store UCN. The trap uses magnetic walls to prevent material interactions between the stored UCN and the walls of the trap. However, not all of the UCN that are produced by the source can be stored in the trap. One basis that can be used to describe the spin state of UCN is high and low field seeking UCN. UCN with spin in the low-field-seeking (LFS) state are repelled by the magnetic walls and can be stored in the trap, but UCN with spin in the high-field-seeking (HFS) state are attracted to the magnetic walls and can not be stored in the trap.

A 6 T superconducting magnet that is coaxial with the UCN guide is used to select only UCN in the HFS state from the UCN that exit from the biological shield. UCN in the HFS state accelerate as they are attracted toward the point of highest magnetic field at the center of the magnet. Once the UCN pass the center of the magnet, they are traveling toward magnetic fields of lower magnitude, which cause the UCN to experience negative acceleration. The magnetic fields are symmetric, so the velocity gained as the UCN travels toward the center of the magnet is equal to the velocity lost as the UCN travels away from the center of the magnet. Any UCN that is in the LFS does not have enough energy to make it to the point of highest magnetic field and is repelled back toward the biological shield. Repelling UCN in the LFS state gives them the opportunity to switch their polarization to the HFS state, which increases the fraction of UCN produced by the source that can be stored in the trap.

UCN that pass through the 6 T superconducting magnet are polarized, but they are polarized in the HFS state. The trap that will be introduced in Section 2.6 can only store UCN in the LFS state. The polarized UCN in the HFS state are converted to polarized UCN in the LFS state using an adiabatic fast passage "spin flipper" coaxial with the UCN guide that is similar to the one used in the UCNA experiment [36]. The spin flipper applies an RF standing wave to the UCN at a frequency designed to match the precession frequency of the UCN that pass through the spin flipper. The frequency is tuned to 372 kHz to maximize the density of UCN in the trap.

2.6 UCN trap

2.6.1 Geometry

The geometry of the trap is comprised of two torus patches. The two tori share a common axis 1.5 m above the bottom of the trap. One torus has a minor axis $a = 1$ m and a major axis $R = 0.5$ m. The other torus has a minor axis $a = 0.5$ m and a major axis $R = 1$ m. Both tori have $a + R = 1.5$ m, which ensures that the two toroidal sections join smoothly. The asymmetry in the trap promotes rapid phase-space mixing of UCN within the trap. The need for rapid phase-space mixing will be discussed in Sections 3.8.2 and 4.7. The top of the trap is set at 0.5 m above the lowest point of the trap, and both torus patches do not continue above this height. UCN with sufficient energy to rise > 0.5 m above the bottom of the trap ($E_n/mg > 50$ cm) are not contained by the magnetic walls that will be discussed in Section 2.6.2, so there is no reason to extend the maximum height of the trap [37]. A cutout of the geometry of the trap is shown in Figure 2.3. The volume of the trap is $\simeq 600$ L, but for reasons that will be discussed in Section 2.8, only $\simeq 420$ L of the volume is used to store UCN. The volume of the trap is held inside of a vacuum jacket, and the pressure is held at $\sim 10^{-6}$ torr to suppress UCN from upscattering on gas molecules.

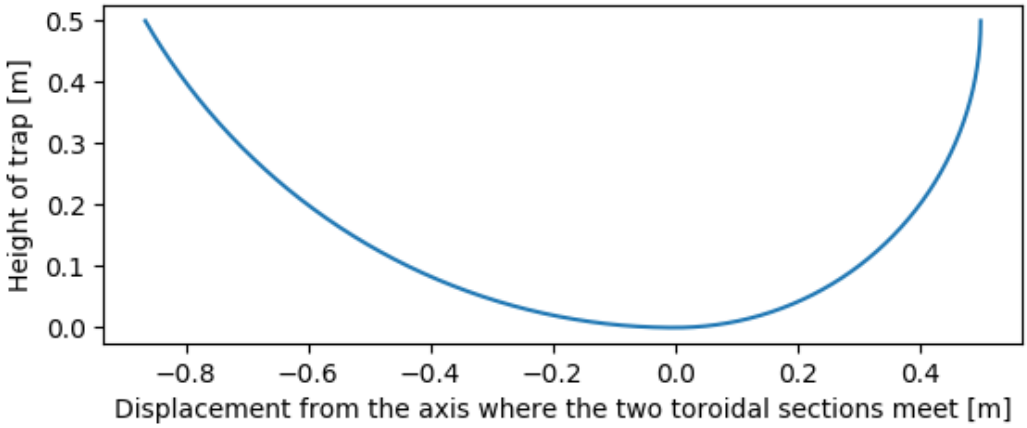

Figure 2.3: The surface of the trap at the point where the trap was the widest across the axis along which the two torus patches met. The two torus patches join at 0 m displacement, and both the surface and derivatives of all orders of the surface are continuous at the point where the two torus patches meet. The horizontal and vertical axes have the same scale, so the curvature shown in this figure is representative of the curvature of the actual trap.

2.6.2 Halbach array magnetic walls

Past experiments that stored UCN in order to measure the neutron lifetime used traps with material walls. The material of the walls were chosen to minimize losses due to interactions with the walls, but significant corrections were applied to the extracted lifetimes to account for the losses. The UCNτ experiment uses a magneto-gravitational trap to store UCN. Section 2.5 discussed how UCN are polarized into the low-field-seeking (LFS) state before they are loaded into the UCNτ trap. The bottom and sides of the trap consist of a magnetic field that repels the polarized UCN before they can physically interact with any material walls. There is no top to the trap, but the Earth's gravitational field prevents UCN with sufficiently low energy from escaping out of the top of the trap.

A magnetic field that is always normal to the surface of the trap would achieve the goal of repelling all UCN that were inside the trap. A Halbach array [38] is used to (approximately) generate this magnetic field. The Halbach array consists of a series of permanent magnets (PMs) that are arranged in an alternating fashion. An example of a one-dimensional Halbach array is shown in Figure 2.4.

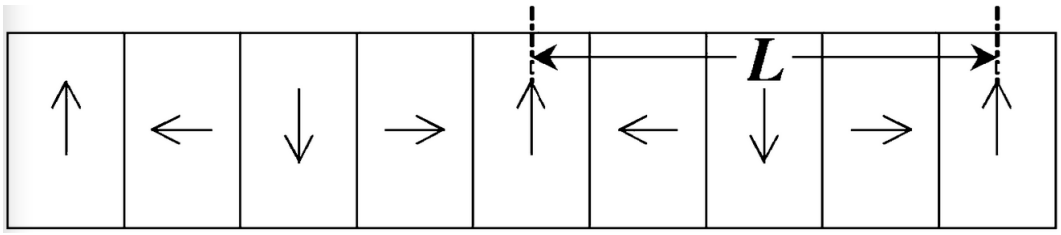

Figure 2.4: A subset of a one-dimensional Halbach array constructed of nine PMs [37]. The direction of the remnant field in each PM is shown with arrows. If "above" the Halbach array is defined as the direction that the left-most PM is pointing, then the magnetic field above the Halbach array decreases exponentially with the distance from the PMs with a characteristic length proportional to L. The magnetic field below the Halbach array has minimal magnitude. In the UCNτ experiment, UCN are stored above the Halbach array. Each PM used in the UCNτ experiment is 1×2 inches, with a depth of $\frac{1}{2}$ inches (not shown in this figure).

A two-dimensional Halbach array was built by combining many one-dimensional Halbach arrays, where the phase difference between each one-dimensional Halbach array and the previous one-dimensional Halbach array is one PM. The magnetic field generated by an infinite two-dimensional Halbach array was derived in [37] as

$$\boldsymbol{B}(x,y,z) = \frac{4B_{\text{rem}}}{\sqrt{2}\pi} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{4n-3} \left(1 - e^{k_n d}\right) e^{-k_n z} \left(\sin k_n y \hat{\boldsymbol{y}} + \cos k_n y \hat{\boldsymbol{z}}\right), \quad (2.5)$$

where B_{rem} is the remnant field generated by one of the PMs, $k_n = 2\pi(4n-3)/L$, d is the thickness of the Halbach array, x and y are orthogonal directions along the surface of the two-dimensional Halbach array, and z is the distance above the surface of the two-dimensional Halbach array. $\boldsymbol{B}$ is not symmetric in the plane of the two-dimensional Halbach array because the orientation of the PMs are not symmetric under a $\frac{\pi}{2}$ rotation. As shown in Figures 2.3 and 2.5, the trap is clearly not a flat two-dimensional plane. However, Equation 2.5 is a reasonable approximation of the magnetic field within a local region of the surface of the trap because the scale of the distances in Equation 2.5 are much less than the radius of curvature of the surface of the trap. When Equation 2.5 is used to approximate the magnetic field generated by the surface of the trap the coordinates are updated from the Cartesian frame (x, y, z) to the surface local frame (ξ, η, ζ), where ξ and η are orthogonal to each other and in the plane of the local surface of the the trap and ζ is the normal distance from the surface of the trap.

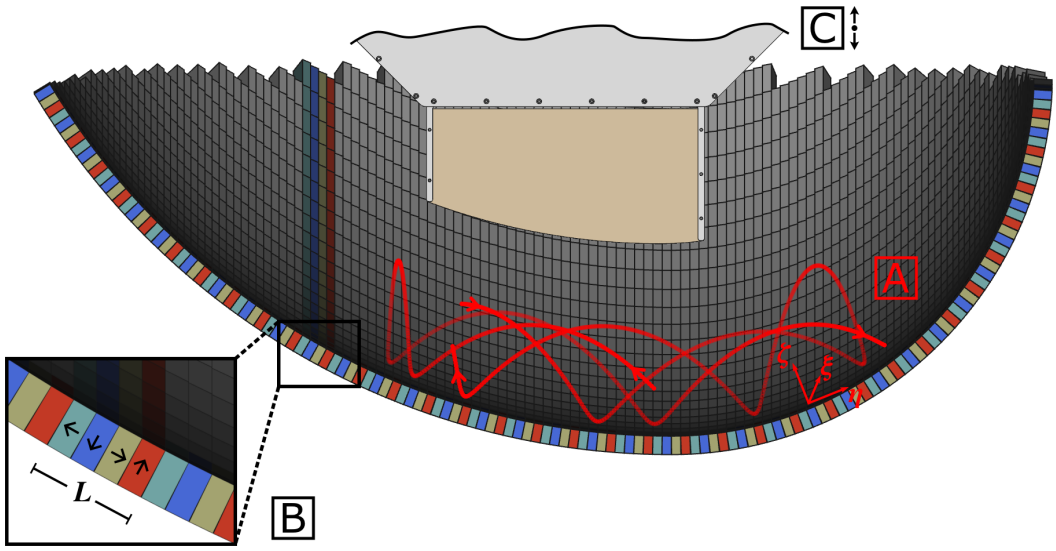

Figure 2.5: A CAD drawing of the Halbach array. A: a typical trajectory of a UCN confined within the trap. B: the alternating pattern of the PMs that make up the Halbach array. C: the primary detector that will be described in Section 2.7.2 that can be raised above the volume of the trap, and later lowered into the trap to count the UCN inside of trap. The tan section is the active surface of the detector. The silver section is the mounting that connects the active surface to an actuator that was located farther above the trap (not shown).

The Halbach array that is the bottom and walls of the trap consists of 5310 NdFeB PMs that are each $1 \times 2 \times \frac{1}{2}$ in^3. Each individual PM has a remnant magnetic field of $B_{\text{rem}} = 1.2$ T. These values can be used with Equation 2.5 to find that the magnitude of the magnetic field 2 mm above the surface of the trap is $B_{\text{surface}} = 0.88 \pm 0.07$ T, where the uncertainty denotes ripples in the field. The minimum magnitude of the magnetic field 2 mm above the surface of the trap is 0.81 T. Setting the magnetic potential energy of a UCN in this field ($E = 48.9$ neV) equal to the gravitational potential energy of a UCN in the LFS state at height h results in $h = 47.6$ cm, which is slightly less than the 50 cm that was chosen for the maximum height of the trap. Therefore, the magnitude of the magnetic field generated by the Halbach array is sufficient to trap nearly all UCN that have insufficient energy to escape out of the top of the trap.

2.6.3 Holding field

The magnitude of the magnetic field produced by Equation 2.5 decreases exponentially as the distance from the surface of the trap increases. Therefore, the magnetic fields from the PMs are fairly localized to the surface of the trap and do not contribute significantly to the magnetic fields in the bulk of the volume of the trap. The UCN that are loaded into the trap are polarized into the LFS state using the methods from Section 2.5. If the UCN inside of the trap pass through a region of zero magnetic field then they could spin-flip from the LFS state to the HFS state. A UCN in the HFS state would not be repelled by the Halbach array and would be quickly lost from the trap.

To ensure that no field zeros exist, the entire volume of the trap is placed inside of a magnetic holding field. The holding field is generated by 10 water-cooled copper coils that surround the vacuum jacket. A current of 300 A is run through the coils to generate a $\sim 6 - 13$ mT holding field throughout the trap. Equation 2.5 shows that the magnetic field produced by the Halbach array has components in the direction normal to the surface of the trap, as well as in one of the two directions along the surface of the trap. To prevent the holding field from weakening the magnetic fields produced by the Halbach array and possibly producing field zeros, the holding field is oriented orthogonally to the magnetic fields produced by the Halbach array. Figure 2.6 shows how the coils that generate the magnetic holding field surround the trap.

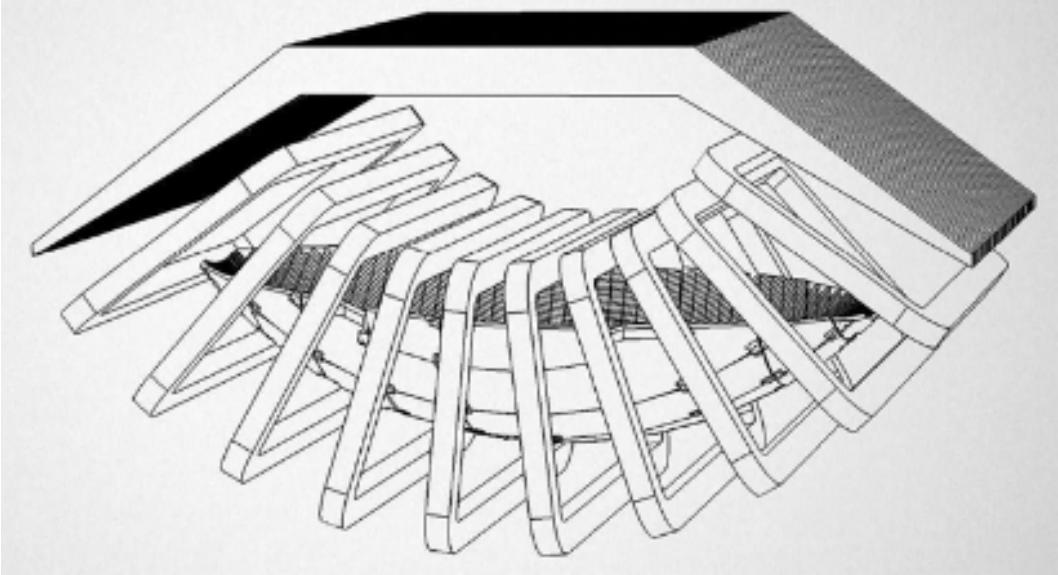

Figure 2.6: A CAD drawing of the coils that generate the magnetic holding field that surrounds the trap [37]. A magnetic flux return is shown above the trap, but the flux return was never built.

2.6.4 Putting UCN into the trap

UCN in the trap are confined below by the magnetic fields generated by the Halbach array, but there is no surface confining UCN above the trap. Instead, Earth's gravitational field prevents low-energy UCN from escaping out of the top of the trap. However, this presents a problem for getting UCN into the trap. Any UCN that is *freely* dropped into the trap has sufficient energy to escape from the trap. Therefore, UCN enter from the bottom of the trap. (Section 5.4 will discuss a proposed next-generation mechanism that would fill the trap with UCN from the top by slowly lowering the UCN into the trap.)

In order to accommodate this, a 6×6 in^2 section of the Halbach array located in the middle of the bottom of the trap was designed to be lowered to create an opening through which UCN can enter. This section is referred to as the "trapdoor." After the trapdoor is lowered, a copper plate ($V_{\text{Fermi}} = 168$ neV [33], which is sufficiently high to reflect all trappable UCN) is positioned at a $\frac{\pi}{4}$ angle and redirects UCN that are travelling along the guides upwards and through the trapdoor. After the trap is filled with UCN, the copper plate is shifted out of the path of the trapdoor, and then the trapdoor is raised back into place to seal the bottom of the trap.

2.7 UCN detectors

2.7.1 ^{10}B-coated ZnS:Ag

Neutrons are neutral particles so it is not efficient to directly detect them with scintillators. Instead, UCN are captured using reactions that produce charged particles, and those charged particles are detected with scintillators. The detection of UCN takes place in three steps, as shown in Figure 2.7:

1. A UCN is captured by ^{10}B, which has a $\sim$ 50 Mb capture cross section.

2. After ^{10}B captures a UCN two charged particles are produced via ^{10}B(n, α)^{7}Li. This reaction releases 2.79 MeV of energy. There is a 96% branching ratio to a decay channel in which a 0.48 MeV photon is released, and in the other 4% of the decays no photon is released. The combined kinetic energy of the α and the ^{7}Li are $\geq$ 2.31 MeV. The α and the ^{7}Li are nearly back-to-back and travel in nearly opposite directions, or in the case where a photon is not emitted the α and the ^{7}Li are exactly back-to-back. Either the α or the ^{7}Li travels through ZnS doped with Ag (ZnS:Ag) that is located adjacent to the ^{10}B.

3. When one of the charged particles passes through the ZnS:Ag it generates scintillation light. This light is detected by a photomultiplier tube (PMT) that generates a voltage signal proportional to the amount of light detected.

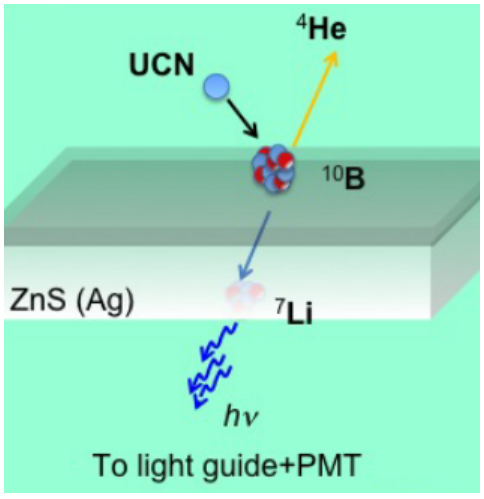

Figure 2.7: A schematic for the process by which UCN are detected [39]. UCN are captured by the ^{10}B layer. Back-to-back charged particles are produced via ^{10}B(n, α)^{7}Li. One of the charged particles passes through the ZnS:Ag and produces scintillation light. The scintillation light travels through a light guide to a PMT.

Clear polyester sheets coated with 3.25 ± 0.25 mg/cm^2 of ZnS:Ag were purchased from a commercial vendor [40]. The thickness of the ZnS:Ag layer is $\sim$ 10 μm. Vacuum evaporation is used to deposit $\sim$ 20 nm of ^{10}B on top of the ZnS:Ag. The thickness of the ^{10}B layer was chosen while considering that a thicker layer increases the UCN capture efficiency per interaction, but also decreases the scintillation light yield in the ZnS:Ag. The clear polyester sheets are backed with $\sim$ few mm of poly(methyl methacrylate) (PMMA, or sometimes referred to as acrylic) which provides a rigid structure to the UCN detector. The probabilities that a UCN incident on the detector is captured or reflected can be calculated by modelling the detector as a series of layers with different Fermi potentials and considering a quantum mechanical wave incident on the layers, as discussed in Section 2.1 of [41]. All UCN detectors in the UCNτ experiment use this detection method. The general process by which the ^{10}B and Zn:S were layered is common among all UCN detectors, but the way in which the scintillation light is collected and detected by PMTs differs among different types of detectors. Sections 2.7.2 and 2.7.3 will discuss how PMTs collect the scintillation light from different types of detectors.

2.7.2 Primary detector

The primary detector counts the number of UCN in the trap. The number of UCN in the trap at the beginning and the end of each storage time are direct inputs to extracting a lifetime, so it is critical that the primary detector accurately and precisely counts the number of UCN in the trap. Therefore, the primary detector was built to a higher standard than the upstream monitors, which will be described in Section 2.7.3.

The key feature of the primary detector is the coincident detection of scintillation light in two PMTs. This coincident detection significantly suppresses the backgrounds caused by the dark rates of the PMTs. A set of 1×1 mm^2 grooves with 2 mm spacing are machined out of a PMMA plate that is 3 mm thick. Wavelength shifting fibers (WLSF) are glued into each groove using optical epoxy. The WLSF are coupled to two PMTs in an alternating fashion so that any two neighboring WLSF are coupled to different PMTs. A second 3 mm thick PMMA plate is glued onto the first to completely surround the WLSF with PMMA. ^{10}B-coated ZnS:Ag screens described in Section 2.7.1 are glued to both sides of the 6 mm of PMMA such that the ^{10}B-coated surface faces outward [42]. Figure 2.8 shows a schematic of the layers of the primary detector. Figure 2.9 shows the primary detector in various stages of construction.

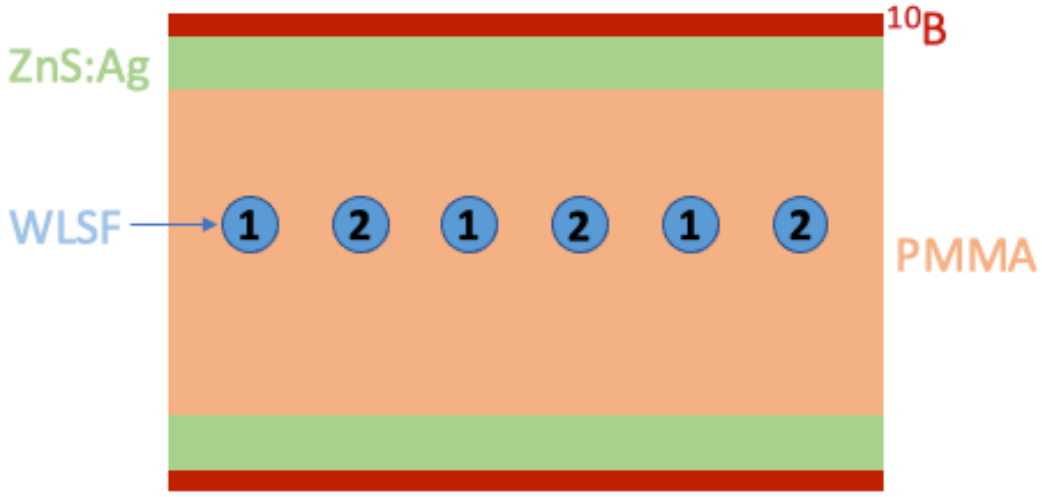

Figure 2.8: A schematic of the layers of the primary detector. The WLSF alternate which PMT they are connected to, as denoted by the numbering. The relative locations and size of the WLSF and the PMMA are as shown, but the layers of ZnS:Ag and ^{10}B are actually much thinner than shown.

The primary detector is connected to an actuator that raises and lowers it through the volume of the trap. The primary detector is lowered in multiple steps to sample the energy spectrum of UCN in the trap. The primary detector is positioned in the middle of the midplane of the trap and covers $\sim 20\%$ of the area of the midplane. The curved bottom edge of the primary detector was chosen to match the curve of the bottom of the trap. The primary detector is raised above the entire usable volume of the trap while UCN are stored in the trap. When a UCN is captured by the primary detector scintillation light is produced. Then, some of that scintillation light enters the PMMA, is transferred to the WLSF, and is directed to the PMTs. The voltage signals from the two PMTs are amplified by a $10\times$ gain and are discriminated at thresholds of $\frac{1}{6}$ and $\frac{1}{3}$ photoelectrons. The two thresholds allow systematic effects due to gain drifts and background to be studied. The higher threshold data are more sensitive to gain drifts, so only the lower threshold data are used in the analyses presented in Chapters 3 and 4.

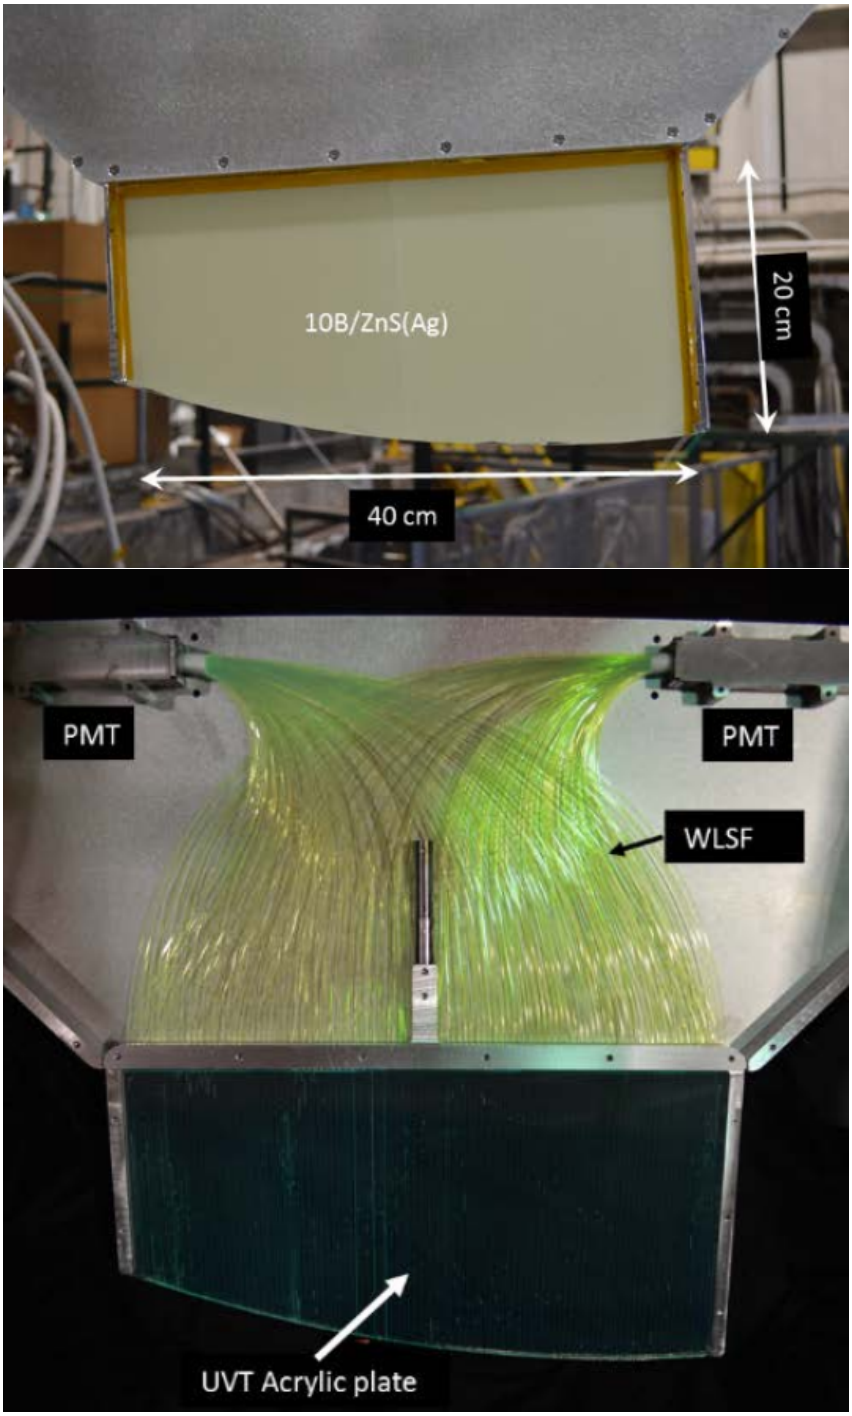

Figure 2.9: The primary detector. Top: a fully-constructed version of the primary detector. Bottom: the internal workings of the primary detector, observed prior to completing the construction [42]. ^{10}B-coated ZnS:Ag sheets are mounted on both sides of the acrylic (PMMA) plate. WLSF are interleaved through the acrylic plate which directs the scintillation light to the two PMTs. The WLSF are interleaved in an alternating manner so that any source of significant scintillation light is detected in both PMTs.

2.7.3 Upstream monitors

The primary detector described in Section 2.7.2 counts the number of UCN in the trap, but in doing so it removes the UCN from the trap. The number of UCN in the trap at the beginning and the end of each storage time are used to extract a lifetime from the data. The number of UCN in the trap are not directly measured at the beginning of each storage time, but instead the number is estimated using measurements of the UCN density at various positions in the guides between the UCN source and the trap via the upstream monitors.

The UCN density in the guides are measured using ^{10}B-coated ZnS:Ag detectors. It is much less critical to make high-precision measurements of the local UCN densities than it is to make a high-precision measurement of the number of UCN in the trap, so the upstream monitors have a simpler design than the primary detector. The ^{10}B-coated ZnS:Ag screens of each upstream monitor are optically observed by a PMT, but the PMT is not physically coupled to the scintillator with WLSF. The voltage signal from each PMT is integrated by a shaping amplifier and then discriminated. The parameters of the shaping amplifier are tuned with the goal that one captured UCN should produce one discriminated count.

Some of upstream monitors sample, and therefore remove, UCN that can be stored in the trap. In order to limit the decrease in the number of UCN loaded into the trap, upstream monitors that sample UCN with energies that can be stored in the trap are coupled to the UCN guides with only very small openings in the guides. One monitor of this type is the "gate valve" monitor, which is located near the gate valve and is shown in Figure 2.1.

The sD_2 source produces, and the guides contain, many UCN with too much energy to be stored in the trap. These high-energy UCN can not be used to directly measure the lifetime, but they can be used to measure the density of UCN in the guides. The "standpipe" monitor is installed inside of a vertical offshoot of the guide system at a height just above the lowered height of the cleaners. The standpipe monitor is also shown in Figure 2.1. At this height no trappable UCN can be sampled by the standpipe monitor, so there is no need to limit the coupling of the standpipe monitor to the guide system. Therefore, the standpipe monitor is designed with a large active area and detects UCN at a rate $\sim 7\times$ the rate at which UCN are detected in the gate valve monitor.

Between the 2017 and 2018 run cycles, a large buffer volume (also shown in Figure 2.1) was installed between the UCN source and the trap. Upstream monitors are installed at the bottom of the buffer volume and near the top of the buffer volume. Following the same logic in the above two paragraphs, the upstream monitor at the bottom of the buffer volume (the "buffer volume monitor") is coupled to the interior of the buffer volume with a very small opening. In contrast, the upstream monitor at the top of the buffer volume (the "buffer volume pre-cleaner") has ~ 1000 cm^2 of active surface.

2.8 Removing above-threshold UCN from the trap

The UCN trap discussed in Section 2.6 is 50 cm tall and has magnetic walls sufficient to contain all UCN with $E/m_n g \leq 47.6$ cm. This height of 47.6 cm is determined by considering a UCN that falls through Earth's gravitational field directly onto the weakest point of the magnetic wall. However, UCN in the trap follow trajectories that are not exactly vertically. Therefore, some UCN with $E/m_n g > 50$ cm are reflected by the magnetic walls. UCN with certain classes of trajectories and with $E/m_n g > 50$ cm remain in the trap for $\sim$ hundreds of seconds before being lost from the trap. The loss of high-energy UCN out of the trap would bias the extracted lifetime, so it is necessary to remove these high-energy UCN from the trap after the filling process is complete.

Two "cleaners" are used to quickly remove above-threshold UCN from the trap. The "large cleaner", which is shown in Figure 2.10, has a $\simeq 0.86$ m^2 surface of polyethylene, which has a high cross section for upscattering and absorbing UCN. The "small cleaner" has a $\simeq 0.23$ m^2 active surface composed of the ^{10}B-coated ZnS:Ag that is described in Section 2.7.1. The active surface of the small cleaner is observed by four PMTs to detect scintillation light that is generated when UCN are captured by the ^{10}B surface. The PMTs observe the scintillation light from a distance instead of being physically coupled to the scintillator. The voltage signals from the PMTs that observed the small cleaner provide some amount of information about the energy spectrum of the UCN that are loaded into the trap, but that information is not used in the analyses presented in Chapters 3 and 4. The shapes of the cleaners are designed to minimize the spacing between the cleaners and the sides of the trap while the cleaners are in their lowered positions. This reduces the number trajectories of above-threshold UCN that can avoid the cleaners.

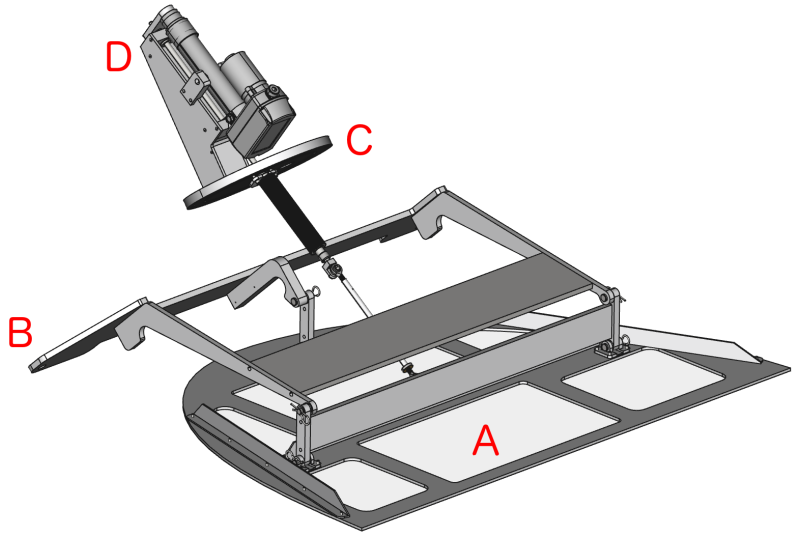

Figure 2.10: A CAD drawing of the large cleaner assembly and actuator. The small cleaner has a similar assembly. A: the cleaner that is used to remove high-energy UCN from the trap. B: the mounting bracket that attaches to the vacuum jacket and supports the cleaner assembly. C: vacuum port with a feedthrough that connects the cleaner to a stepper motor. D: The stepper motor that is used to raise and lower the large cleaner.

In order to remove above-threshold UCN from the trap, the two cleaners are lowered to 38 cm above the bottom of the trap while the trap is being filled with UCN. The surfaces of the two cleaners span $\simeq$ 65% of the horizontal cross-sectional area of the trap while they are 38 cm above the bottom of the trap. The two cleaners remain 38 cm above the bottom of the trap for 50 s after the filling process ends, and then they are raised to 43 cm above the bottom of the trap. The two cleaners remain 43 cm above the bottom of the trap while UCN are stored in the trap. The height of the cleaners when lowered and the duration for how long the cleaners remain lowered are chosen through a combination of experimentation and Monte Carlo simulations. As mentioned in Section 2.7.2, the primary detector is lowered in multiple steps to sample the energy spectrum of UCN in the trap. The primary detector can be lowered in nine steps (as opposed to the typical three steps) to finely sample the energy spectrum of UCN in the trap. This energy spectrum measured with nine steps of the primary detector is used as an input to to the Monte Carlo simulations. The energy spectra of real and simulated UCN match when measured with three steps, which shows that the simulated evolution of UCN trajectories within the trap matches reality.

If the lowered height of the cleaners were to be decreased, or the duration for how long the cleaners are lowered were to be increased, then both the number of above-threshold UCN in the trap and the total number of UCN in the trap would be reduced. Monte Carlo simulations find that lowering the cleaners to 38 cm for 50 s is 99.99% efficient at removing the population of UCN with $E/m_n g > 43$ cm [43]. Some of the UCN with $E/m_n g > 43$ cm that remain in the trap can reach the raised cleaners during the holding process and be lost from the trap. These losses can bias the extracted lifetime, but the Monte Carlo simulations find that for this level of losses of UCN the bias is < 0.05 s. If the height of the cleaners is lowered, or the cleaners are left at their lowered height for a longer duration of time, then the bias in the extracted lifetime would decrease. However, either of these options would also decrease the statistical reach of each run of the experiment.

2.9 Upstream buffer volume

Between the 2017 and 2018 run cycles, a cylindrical buffer volume was installed between the UCN polarization assembly and the standpipe monitor. The purpose of the buffer volume is to smooth out fluctuations in the production of UCN, which improves the ability to estimate the number of UCN that are loaded into the trap. The buffer volume is made of Al and the interior is coated with NiP, which has a Fermi potential of $V_F = 213$ neV. The buffer volume has a volume of $\simeq 770$ L. The buffer volume pre-cleaner, discussed in Section 2.7.3, is positioned at a height near the middle of the buffer volume. This height is chosen to match the height of the two cleaners from Section 2.8 that are inside of the trap. Therefore, only $\simeq 380$ L of the buffer volume is usable for smoothing out fluctuations in the production of UCN. This volume is similar to the $\simeq 420$ L of the trap that are used to store UCN.

UCN in the buffer volume have already passed through the UCN polarization assembly and are in the LFS state. To prevent UCN in the LFS state from spin-flipping, the entire buffer volume is placed inside of a holding field. Approximately 10 A of current are run through ~ 125 turns of copper wire that are wound around the buffer volume, which generates a holding field of ~ 2 mT throughout the buffer volume.

The ports through which UCN enter and exit the buffer volume are located directly opposite each other. To prevent UCN from quickly passing through the buffer volume, which would defeat the purpose of the buffer volume, large bent copper strips are installed along the bottom of the buffer volume. These copper strips randomize the trajectories of UCN within the buffer volume.

Section 3.7.9 will compare the data gathered during the 2017 run cycle (without the buffer volume) to the data gathered during the 2018 run cycle (with the buffer volume). The buffer volume accomplishes its goal of smoothing out fluctuations in the production of UCN (this will be quantified in Section 3.7.9). However, the buffer volume also reduces the number of UCN loaded into the trap by $\sim$ 50%. The contributions to the statistical uncertainty of an extracted lifetime from both the finite number of UCN counted in the trap and the uncertainty of the number of UCN loaded into the trap will be quantified in Section 3.11.

2.10 Data acquisition

The discriminated counts from the detectors from Section 2.7 and the small cleaner from Section 2.8 are recorded by two event-time digitizers that record the time of each event with 800 ps precision [44]. The internal clocks of the two digitizers are synchronized at the beginning of each run, and the beginning of data acquisition is synchronized with the beginning of production of UCN. The internal clocks have frequency stability of 3×10^{-8} cycles per cycle, which corresponds to a stability in the extracted lifetime of $\simeq 26$ μs. Only 10 input channels are available for use, so in order to reduce bandwidth the data from the four PMTs of the small cleaner are combined into one data stream prior to being sent to the digitizers.

The photoelectrons from the two PMTs of the primary detector, as well as the integrated currents from the normalization monitors and the small cleaner, are recorded and saved to be analyzed later. Each recorded event has three values: the time at which it occurred, the detector in which it occurred, and tags which record the state of the experiment when the event occurred. The tags record the following information:

1. If the proton beam is being delivered to the target;
2. If the cleaners are raised or lowered;
3. If the gate valve is opened or closed;
4. If the trapdoor is in motion; and
5. If the primary detector is in motion.

These tags are generated by either the motion of the components or by the voltage signals that control the components. The rate at which events of any type are recorded is $\gtrsim$ 300 Hz so the tags are used to reconstruct the history of the states of the components with high precision. The mean ratio between the reconstructed holding lengths and the desired holding lengths for runs in the 2017 run cycle is $1 + (38.1 \pm 0.7) \times 10^{-6}$. The same ratio is $1 + (19.5 \pm 0.6) \times 10^{-6}$ for runs the 2018 run cycle. It is unknown what caused the ratio to change between run cycles, but the differences of both ratios from 1 are caused by the imperfections of the CPU clock used to time each run. For some runs the tags were not correctly saved to the data files, so the mean ratios listed in this paragraph are used to estimate the reconstructed holding lengths of those runs.

2.11 Experimental control system

The various components of the UCNτ experiment change state during each run of the experiment. The production of UCN (Section 2.3) starts at the beginning of each run and ends after the trap is filled with UCN. The bottom of the trap (Section 2.6.4) is opened while the trap is being filled with UCN and is then closed. The primary detector (Section 2.7.2) changes height throughout each run. The cleaners (Section 2.8) are lowered into the trap before the holding process, and are then raised above the trap for the rest of the run. The data acquisition system (DAQ, Section 2.10) receives a signal at the beginning of each run in order to synchronize the beginning of the run with the beginning of data acquisition. A control system is needed to manage the timing of the actuation of all of these components of the UCNτ experiment.

The runs from the 2017 and 2018 run cycles were controlled by a system built with a custom C++ program that used QT Creator [45] to create a graphical user interface (GUI). The timing of the experiment was controlled with a timer that was part of QT Creator. This timer ran on the CPU of the computer that displayed the GUI. At the desired times, the C++ program sent commands to a LabJack U3 device [46], which generated voltage signals used to control the various components of the experiment.

The precision of the timing of the experiment was dependent on the precision of the CPU clock. For a typical CPU clock, temperature-dependent frequency drift is $\simeq -4.3 \times 10^{-7}$ cycles per cycle per degree [47]. The experimental area does not have climate control. Over the course of a year the temperature in the experimental area varies by $\sim$ 30 degrees. This drift can cause a $\simeq$ 0.01 s difference between lifetimes

extracted when the experimental area is at its warmest and coldest. A stable, but imperfect, CPU clock that drifts one second per day would bias the extracted lifetime by 0.01 s. Although these shifts are much less than the statistical uncertainty of the extracted lifetime that will be reported in Chapter 3, they could be significant in a next-generation measurement of the lifetime with an uncertainty $\lesssim 0.10$ s.

Section 2.10 reported that the reconstructed and desired holding lengths differ by $\sim 20 - 40$ ppm. The stability of the clocks in the DAQ are orders of magnitude greater than the stability of a typical CPU clock. Therefore, the clocks in the DAQ are assumed to be exactly correct. If the desired holding lengths are used instead of the holding lengths reconstructed from the tags, then the extracted lifetime would be biased by ~ 0.04 s.

Following the 2018 run cycle the control system was upgraded to use a dedicated off-board clock instead of the CPU clock. The off-board clock is provided by a Artix-7 35T Arty FPGA Evaluation Kit [48]. The FPGA is programmed to perform all of the timing and signalling necessary during a run. Prior to the beginning of each run the GUI sends the run sequence to the FPGA. The FPGA controls the generation of low-current 3.3 V signals that are used to initiate and synchronize the start of the run with the start of data acquisition and the start of production of UCN. A printed-circuit board (PCB) with integrated circuits (ICs) is used to convert the low-current 3.3 V signals to 100 mA, 5 V signals. This level is sufficient to serve as a logic signal, but some signals need to be further converted to 24 V to actuate solenoids that control some components of the experiment and to be able to open and close gate valves. An additional ~ 1500 lines of code were developed at Caltech to integrate the FPGA with the existing GUI.

The difference between the clock in the FGPA and the clocks in the DAQ was measured by comparing the reconstructed holding lengths described in Section 2.10 to the desired holding lengths. The mean difference between the reconstructed and desired holding lengths is $\simeq 6$ ppm, which is almost an order of magnitude smaller than the difference between the CPU clock and the DAQ clock. If unaccounted for, this difference would correspond to a shift in the extracted lifetime of $\simeq 5$ ms, which is far less than the statistical uncertainty that will be presented in Chapter 3 or the statistical uncertainty of any next-generation experiment to measure the free neutron lifetime.

Figure 2.11 shows the hardware that allows the evaluation kit to interface with the rest of the experiment. Figure 2.12 shows a schematic of the two types of circuits included in the PCB. A ribbon cable was used to connect between the FPGA and the PCB, and a separate ribbon cable was used to connect between the PCB and the rest of the UCNτ experiment. The updated control system was installed between the 2018 and 2019 run cycles.

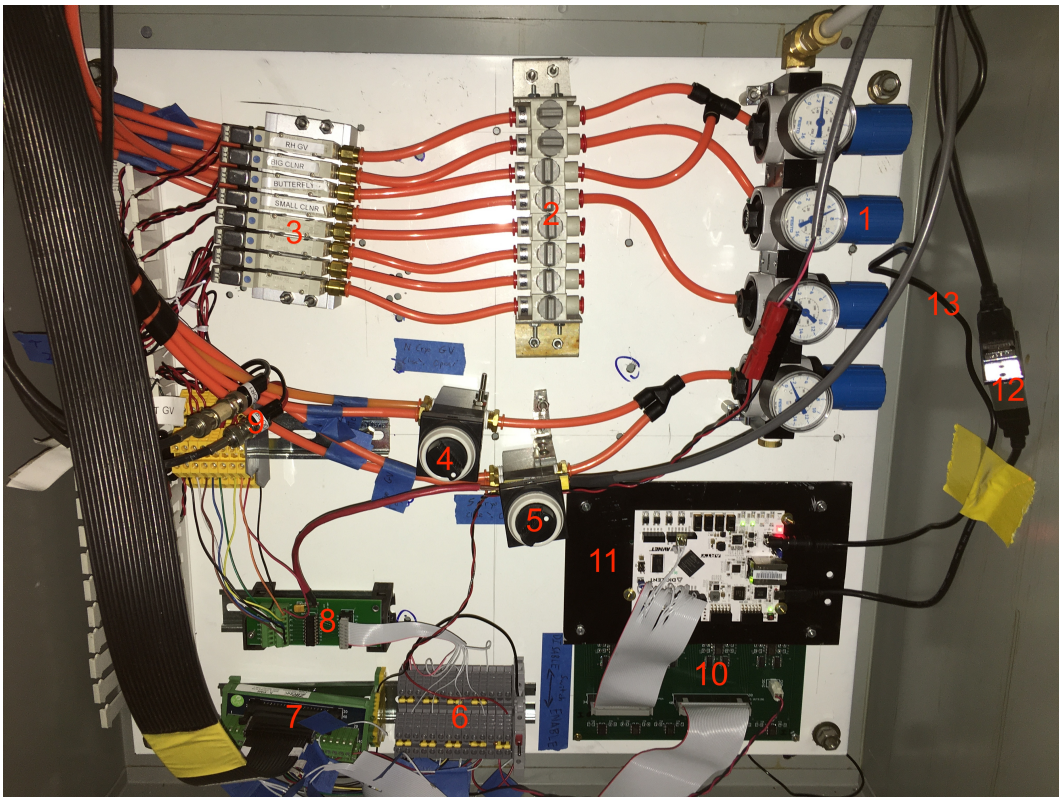

Figure 2.11: The hardware that the evaluation kit uses to interface with the rest of the experiment. 1: four regulators that limit the pressurized air supply to four different levels. 2: manual switches used to control the supply of air to solenoids so that the air lines from the solenoid can be de-energized. 3: solenoids used to control various components of the experiment. 4 and 5: manual switches used to open and close gate valves that are used as part of the vacuum system. 6 and 7: breakout panels where most of the tags are separated from the control signals. 8: converter that transforms 5 V signals to 24 V signals needed to actuate solenoids and gate valves. 9: two BNC cables that carry 24 V signals to control the trapdoor. 10: a printed circuit board that optically isolates the FPGA from the rest of the experiment, and increases the voltage and current of the signals generated by the FPGA. See Figure 2.12 for more details. 11: the FPGA mounted on a support board. 12: a USB connection that communicates between the FPGA and the GUI. 13: a 5 VDC power supply.

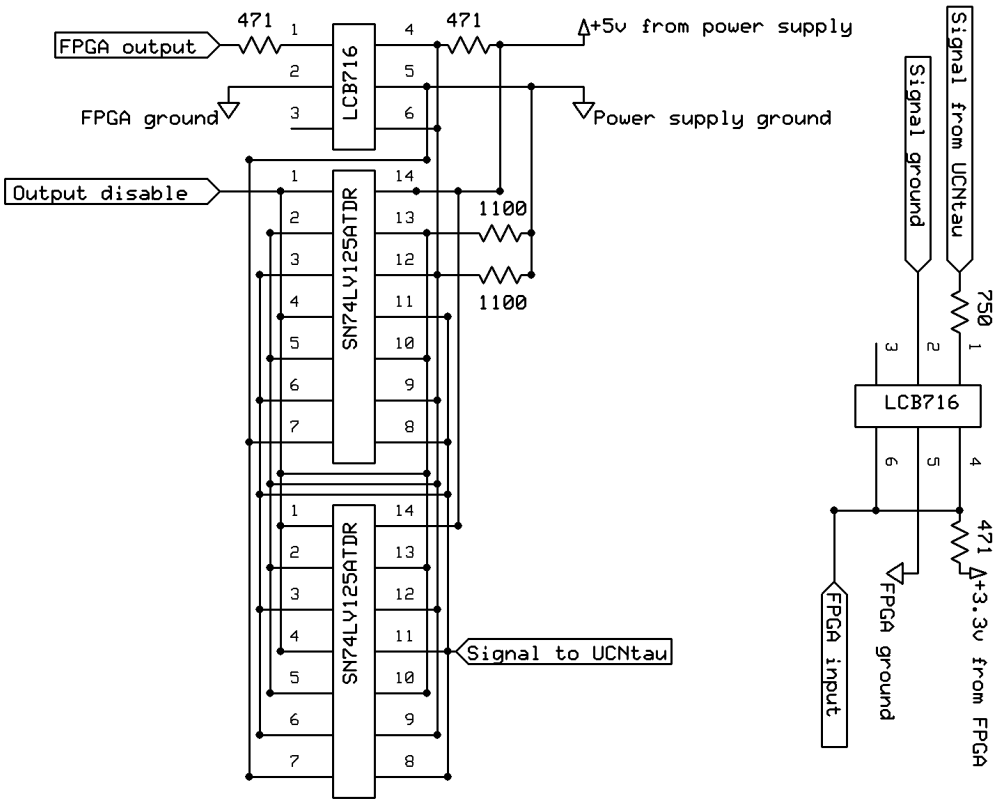

Figure 2.12: Schematics of the two types of circuits used in the PCB. Left: a schematic of the circuit that is part of the PCB that converts an output signal from the FPGA to a 5 V, 100 mA signal that can control a component of the UCNτ experiment. Right: a schematic of the circuit that takes one signal from the UCNτ experiment and creates an input signal to the FPGA. The LCB716 ICs [49] are opto-isolators that protect the FPGA from the outside world. The SN74LV125ATDR ICs [50] are buffer amplifiers that amplify the signal from the FPGA. Each IC has four channels, and each channel can produce a 5 V, 16 mA signal. Connecting eight of these channels in parallel produces a 5 V, 128 mA signal, which is greater than the 100 mA needed to drive the various components of the experiment. Some of these 5 V signals drive switches that control 24 V signals (not shown in this Figure). The PCB contains 24 copies of the "output" circuit (left) and 6 copies of the "input" circuit (right), and was developed at Caltech with assistance from an electrical engineer.

Chapter 3

EXTRACTION OF THE NEUTRON LIFETIME

3.1 Introduction

In general, one run of UCNτ works as follows:

1. Some unknown number of UCN are put into the trap, and the number of UCN in the trap was estimated as N.
2. The UCN are held in the trap for some length of time t, and during this time some of the UCN underwent β-decay and are lost from the trap.
3. The surviving UCN, as well as some background, generate U counts in the primary detector.
4. The number of background counts in the primary detector is estimated as B.

The expected proportion of UCN that remain in the trap is $e^{-t/\tau}$. The *yield* $Y \equiv \frac{U-B}{N}$ is the measured proportion of UCN that remain in the trap. Setting the two equal to each other results in a formula to extract the lifetime via

$$\tau = -\frac{t}{\ln Y}.$$

Section 3.2 will discuss the structure of the data that were used to extract a lifetime and Section 3.3 will justify why data were either included in or excluded from the analysis. Section 3.4 will develop a method to identify when a UCN was captured by the primary detector and Section 3.5 will estimate the rate of background events that were in the primary detector. Fluctuations in the production of UCN resulted in different runs having different numbers of UCN loaded into the trap at the beginning of each run. Section 3.6 will develop a model for estimating the number of UCN that were loaded into the trap at the beginning of each run. Section 3.7 will develop a likelihood model to quantify the statistical reach of the data and will use this likelihood model to find the optimal parameters of the normalization model from Section 3.6. Section 3.8 will use Monte Carlo simulations to study what times should be used to mark the beginning and the end of the length of time that UCN were stored in the trap. Section 3.9 will develop a method of extracting a lifetime that combines

the results from Sections 3.4 through 3.8 with the assumption that all of the UCN in all runs had the same β-decay lifetime. Section 3.10 will develop a different method of extracting a lifetime that also combines the results from Sections 3.4 through 3.8, but pairs together short and long runs that were approximately adjacent in time in order to extract many nearly-independent values for the lifetime.

3.2 Data structure

During the 2017 and 2018 run cycles, the two main types of runs that were performed were production and background runs. During production runs UCN were loaded into the trap, held in the trap for various lengths of time, and then the surviving UCN were counted. During background runs no UCN were loaded into the trap, so the only counts in the primary detector were background counts. Most of the runs were either production or background runs, but some other special types of runs were performed to study systematic effects that may impact the extracted lifetime. The use of these systematic runs will be discussed in Sections 4.2, 4.6, 4.7, and 4.8.

3.2.1 Production runs

Production runs contained the majority of the data that went into extracting a lifetime. Production runs consisted of a series of four processes:

1. Filling: UCN were produced from the source and diffused through the guides. UCN entered the trap through the open trapdoor. In 2017 the trap was filled for 150 s and roughly 20,000 to 40,000 UCN were loaded into the trap. In 2018 the was filled for 300 s and roughly 10,000 to 20,000 UCN were loaded into the trap. The increase in filling length was required because the buffer volume increased the length of time that it took for the UCN density in the trap to reach equilibrium.

2. Cleaning: The trapdoor was closed. UCN with energy sufficient to reach the cleaning height (38 cm above the bottom of the trap) were removed from the trap by the cleaners during the 50 s of cleaning.

3. Holding: The cleaners were raised $\geq$ 5 cm above the cleaning height. UCN were held in the trap for varying holding lengths:

    - Short runs: 20 s, 50 s, 100 s, and 200 s
    - Long runs: 1280 s, 1350 s, and 1550 s
    - Extra-long runs: 3000 s, 4000 s, and 5000 s

Most of the data were gathered in a "short, long, long, short" pattern to reduce possible systematic effects due to differences in UCN source recovery time. The short run holding length cycled through all four options to ensure an even distribution. Multiple short holding lengths were used to provide sensitivity to any loss mechanisms with a lifetime less than the neutron lifetime. Following a hardware upgrade in early 2017 that enabled holding lengths > 1280 s, most long runs had a holding length of 1550 s. As shown in Appendix A, this value was chosen to maximize the statistical reach of the data gathered per hour.

The 800 MeV proton linear accelerator required to begin the production of UCN was usually unavailable between 6 AM and 8 PM on weekdays. It was known in advance that UCN would not be able to be produced during these times, so an extra-long run usually began immediately prior to loss of production of UCN. This allowed for an additional ~ hour of production running before having to wait for the production of UCN to resume. Additionally, the extra-long runs provided some sensitivity to loss mechanisms with a lifetime much greater than the neutron lifetime.

4. Unloading: The primary detector lowered into the center of the trap. The lowering occurred in three steps:

    a) Peak 1: The primary detector lowered from the top of the trap to the cleaning height (38 cm above the bottom of the trap). At this height the primary detector was sensitive to above-threshold UCN that were not properly removed during the cleaning process, as well as to UCN whose energy may have increased during the holding process. Peak 1 had a length of 40 s. Roughly 0.005% of the UCN in the trap were counted during Peak 1. This section will explain why it was useful to have some Peak 1 counts, so it was important to count for some amount of time at Peak 1. However, the primary detector sampled only a very small fraction of the UCN in the trap while at the Peak 1 height. Further increasing the length of Peak 1 would have increased the number of UCN that decayed before being captured by the primary detector, which would have reduced the statistical reach of each run.

    b) Peak 2: The primary detector lowered from the cleaning height to 25 cm above the bottom of the trap. At this height the primary detector was sensitive to the higher-energy trappable UCN, but not to the lower-energy trappable UCN. This provided some information about the energy

spectrum of UCN inside the trap at the end of the run. Peak 2 had a length of 20 s. Roughly 5% of the UCN in trap were counted during Peak 2, which was a sufficient number of UCN to measure the energy spectrum of UCN in the trap. However, the primary detector could only sample a fraction of the UCN in the trap while at the Peak 2 height. Increasing the length of Peak 2 would have increased the number of UCN that decayed before being captured by the primary detector, which would have reduced the statistical reach of each run.

c) Peak 3: The primary detector lowered from 25 cm above the bottom of the trap to 1 cm above the bottom of the trap. During the next 100 s $\gtrsim$ 99.99% of the remaining UCN in the trap were captured by the primary detector. However, the reconstruction algorithm that will be introduced in Section 3.4 was only $\sim$ 90% efficient, so not all of the captured UCN were identified as UCN events. Peak 3 continued for an additional 50 s, for a total length of 150 s.

This number of UCN counted in Peak 1 was too small to significantly affect the statistical precision of the extracted lifetime. The number of UCN counted in Peak 1 was roughly the same between short and long runs, so including UCN counted in Peak 1 in an extraction of the lifetime had a negligible effect on the extracted lifetime. All extracted lifetimes that will be presented do not include Peak 1 counts unless otherwise noted. Instead, Peak 1 counts will be used in Sections 4.7 and 4.8 to constrain systematic uncertainties.

3.2.2 Background runs

Production of UCN was usually unavailable between 6 AM and 8 PM on weekdays. When production of UCN was unavailable for an extended period of time, background runs were performed in order to measure the background rates in the primary detector. During these "beam-off" background runs all parts of the experiment remained stationary, except for the primary detector which moved among four heights. The primary detector spent 250 s at each of the four heights before the beam-off background run ended.

A small number of background runs were also performed while the proton beam was available to produce UCN. During these "beam-on" background runs all parts of the experiment operated as they would have in a production run, except the primary gate valve that allowed UCN to pass from the source to the polarizing magnet remained

closed. This prevented all UCN from entering the trap, so only background counts were recorded in the primary detector. These beam-on background runs more accurately mimicked the structure of a production run. The background models that will be developed in Section 3.5 were compared to the backgrounds observed during the beam-on background runs and no statistically significant different was observed. However, very few beam-on background runs were performed, so the statistical uncertainty of the background rates that were measured during the beam-on background runs was too large to be able to be used to reliably estimate the background rates in this analysis.

3.2.3 Blinding the data

Previous measurements [26] of the neutron lifetime have been published by the UCNτ collaboration. An analyzer who compared the lifetime extracted using a version of their analysis to previously published values for the lifetime could have been biased to make analysis choices that extracted similar lifetime values. This analysis was developed using a blinded version of the data to prevent analyzer bias.

The holding lengths were blinded by squeezing the holding length by a ratio b. The value of $b = 0.98869$ was unknown to the analyzers until after their analysis methods were finalized. A run that was called a "20 s run" really had a holding length of $20b$ s. The timestamps on events in the blinded data files were adjusted so that a "20 s run" did appear to have a holding length of exactly 20 s. This blinding process shifted the measured value of the lifetime by approximately 10 s. After an analysis of the blinded data was finalized, the exact same analysis was performed on the unblinded data to confirm that the process by which the data was blinded was correctly understood.

3.3 Run selection

3.3.1 Introduction

Not every production run from the 2017 and 2018 run cycles could have been used in this analysis. A non-negligible fraction of the production runs were aborted before they finished due to problems with the production of UCN, problems with the data acquisition system, or problems with the experimental apparatus. Some of the runs that did finish were also excluded from this analysis. Section 3.3.2 will discuss which completed production runs were rejected from the analysis due to problems with the production of UCN. Section 3.3.3 will mention a mistake in the program that controlled the experimental apparatus that caused about 10% of production runs to be performed incorrectly. Section 3.3.4 will discuss how changes in the behavior of the primary detector necessitated separating the data that survived the cuts from Sections 3.3.2 and 3.3.3 into different subsets.

3.3.2 Fill quality

A repeatable filling procedure was critical to making an estimate of the number of UCN loaded into the trap. Abnormalities during the production of UCN made it more difficult to estimate the number of UCN loaded into the trap. There were two factors that caused most of the abnormal production of UCN:

- The production of UCN began with an 800 MeV proton linear accelerator. Sometimes this accelerator suffered "glitches" which caused a temporary loss of the proton beam and disrupted the production of UCN.
- The sD_2 source was sometimes accidentally exposed to the UCN in the guides between the source and the trap for much longer than it should be. Some of the UCN upscattered on the sD_2 and were lost from the system.

During the 2017 run cycle, almost all runs with even minor abnormalities in the production of UCN were restarted. As described in Section 2.9, a buffer volume was installed between the UCN source and the trap before the 2018 run cycle. The buffer volume slowed the rate at which the UCN density in the trap increased and lowered the overall maximum UCN density in the trap, but the buffer volume also decreased the effect of abnormalities in the production of UCN on the number of UCN loaded in the trap. During the 2018 run cycle, only runs with extreme abnormalities in the production of UCN were restarted. Figures 3.1 and 3.2 show the effect of beam glitches on the UCN rates in various upstream monitors.

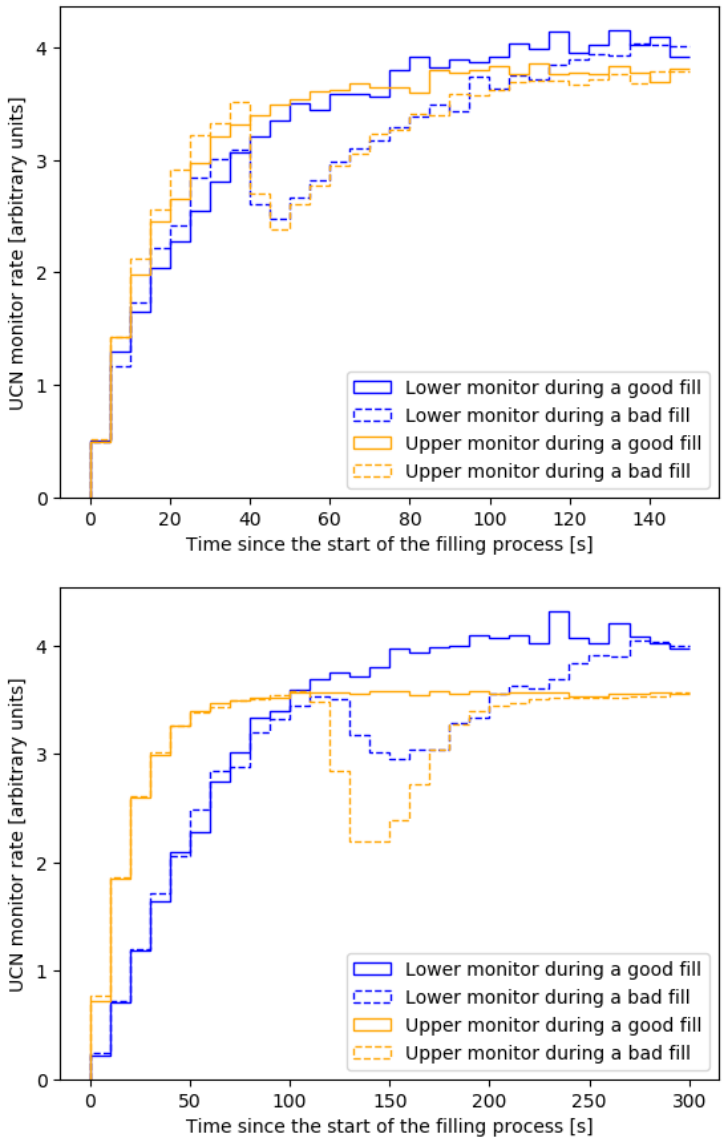

Figure 3.1: Comparison of UCN monitor rates with and without abnormal production of UCN. Top: good and bad fills before the buffer volume was installed. Bottom: good and bad fills after the buffer volume was installed. In the top plot the UCN rates in the upper and lower monitors were similarly affected by the abnormal production of UCN. In the bottom plot both the lower and upper monitors were located inside of the buffer volume. In contrast to the top plot, the bottom plot shows that abnormal production of UCN affected the UCN rate in the lower monitor less than the UCN rate in the upper monitor because the buffer volume smoothed out the effect. This smoothing reduced the effect of abnormal production of UCN on the number of UCN loaded into the trap and the UCN rate in the lower monitor. However, the effect on the UCN rate in the upper monitor was not reduced because high-energy UCN were quickly lost from the guides and the trap, so the density of high-energy UCN was strongly coupled to the very recent production of UCN.

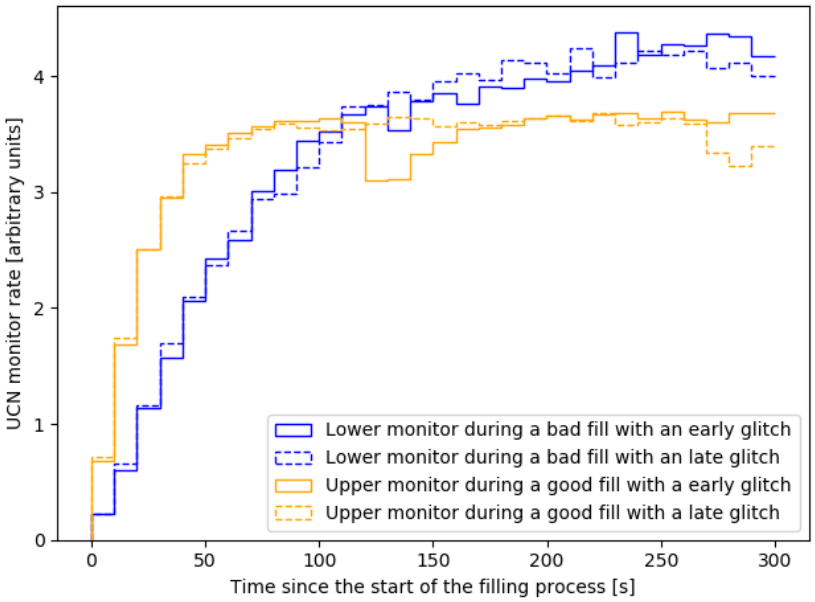

Figure 3.2: UCN monitor rates for runs with glitches that caused abnormal production of UCN at times that were early and late in the filling process. If a beam glitch occurred sufficiently early during the filling process then the system had time to recover and the effect on the number of UCN loaded into the trap was minimal. However, there was not sufficient time for the system to recover from a beam glitch that occurred near the end of the filling process, which made it more difficult to accurately estimate the number of UCN loaded into the trap. This figure only includes data from runs in which the buffer volume was installed. Prior to the installation of the buffer volume, almost all runs with even minor abnormalities in the production of UCN were aborted and restarted.

During a good fill, the rate in the gate valve monitor could be well-modeled as

$$\hat{r}(t) = A\left(1 - e^{-t/\kappa}\right). \tag{3.1}$$

The measured UCN rate $r(t)$ of each run was fit with Equation 3.1. The sum of the squares of the differences between the measured UCN rate and the fit provided a quantitative measure of the quality of the fill. Abnormalities in the production of UCN that occurred in the later part of the filling process had a significantly greater affect on the number of UCN loaded into the trap than abnormalities that occurred in the early part of the filling process. Therefore, a weighted sum of the squares of the residuals was a better measure of the quality of the fill. The weighting function

$$w(t) = \left[1 + \exp\left(-\frac{t - T_{\text{fill}}/2}{T_{\text{fill}}/10}\right)\right]^{-1} \tag{3.2}$$

assigned more significance to abnormal events that occurred later during the fill.

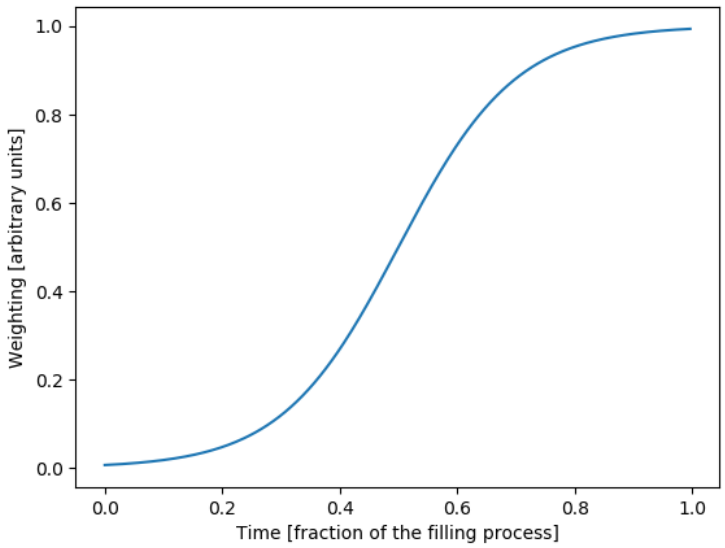

Figure 3.3: Fill quality weighting of different parts of the filling process, as defined in Equation 3.2. Later parts of the filling process were more highly correlated with the number of UCN loaded into the trap, so the later parts of the filling process were more heavily weighted when fill quality was calculated.

The fill quality was then defined as

$$Q = \ln\left[\frac{1}{\sum_t w(t)}\sum_t w(t)\left(\frac{r(t)-\hat{r}(t)}{\sqrt{\hat{r}(t)}}\right)^2\right]. \tag{3.3}$$

Runs with high-quality fills had smaller values for Q than runs with poor-quality fills. Figure 3.5 shows distributions for Q for the 2017 and 2018 data sets. In order to assess the effect of including runs of different fill qualities on the extracted lifetime, the runs were sorted by fill quality and lifetimes were extracted as runs of marginally worse fill quality were included in the analysis. The results of this are shown in Figure 3.4.

The buffer volume was not installed during 2017 run cycle. Abnormal production of UCN had a significant effect on the number of UCN loaded into the trap so great care was taken to abort and discard any runs that had an abnormal fill. In 2017, the lifetime extracted using all of the runs did not significantly vary from the lifetime extracted using only the 70% of runs with the best fill qualities. However, including the 2% of runs with the worst fill qualities increased the statistical uncertainty of the extracted lifetime, as seen in the tapering of the uncertainty bands in the top plot in Figure 3.4. Therefore, the 98% of the runs from 2017 with the best fill qualities were used in this analysis and the other 2% are rejected.

The buffer volume was installed before the 2018 run cycle. After its installation the effect of abnormal production of UCN on the number of UCN loaded into the trap was suppressed, so many more runs with filling abnormalities were run to completion instead of being manually aborted. In 2018, the lifetime extracted using all of the runs did significantly differ from the lifetime extracted using only the 70% of runs with the best fill qualities. The 89% of runs with the best fill qualities could be used before the extracted lifetime significantly deviated from the lifetime extracted using only the 70% of runs with the best fill qualities, as seen in the bottom plot of Figure 3.4. Therefore, the 89% of the runs from 2018 with the best fill qualities were used in this analysis and the other 11% are rejected.

3.3.3 Blinding bug

Subsection 3.2.3 details how and why the data were blinded. Of the $\sim$ 5000 runs that were blinded, 345 were improperly blinded. Some of these 345 runs would have been unusable, even if properly blinded, due to poor fill quality or other issues. The final analysis used 3853 runs to extract a lifetime. The unblinded data were not affected by this bug. The cause of the improper blinding was determined and fixed between the 2018 and 2019 run cycles.

The cause of the improper blinding was well understood and the improperly blinded runs could have been recovered and used in an unblinded analysis. It was estimated that this would have decreased the statistical uncertainty of the extracted lifetime by $\simeq$ 4%. Out of an overabundance of caution, these improperly blinded runs were excluded from this analysis.

3.3.4 Changes in the condition of the primary detector

Occasionally, actions were taken that caused step changes in the behavior of the primary detector. Some examples of these actions included:

- Accidentally damaging the primary detector;
- Adjusting the system used to cool the photomultiplier tubes (PMTs) of the primary detector;
- Removing the primary detector from the trap and exposing it to light and air for a significant amount of time; and
- Adjusting what other detectors were mounted inside of the trap.

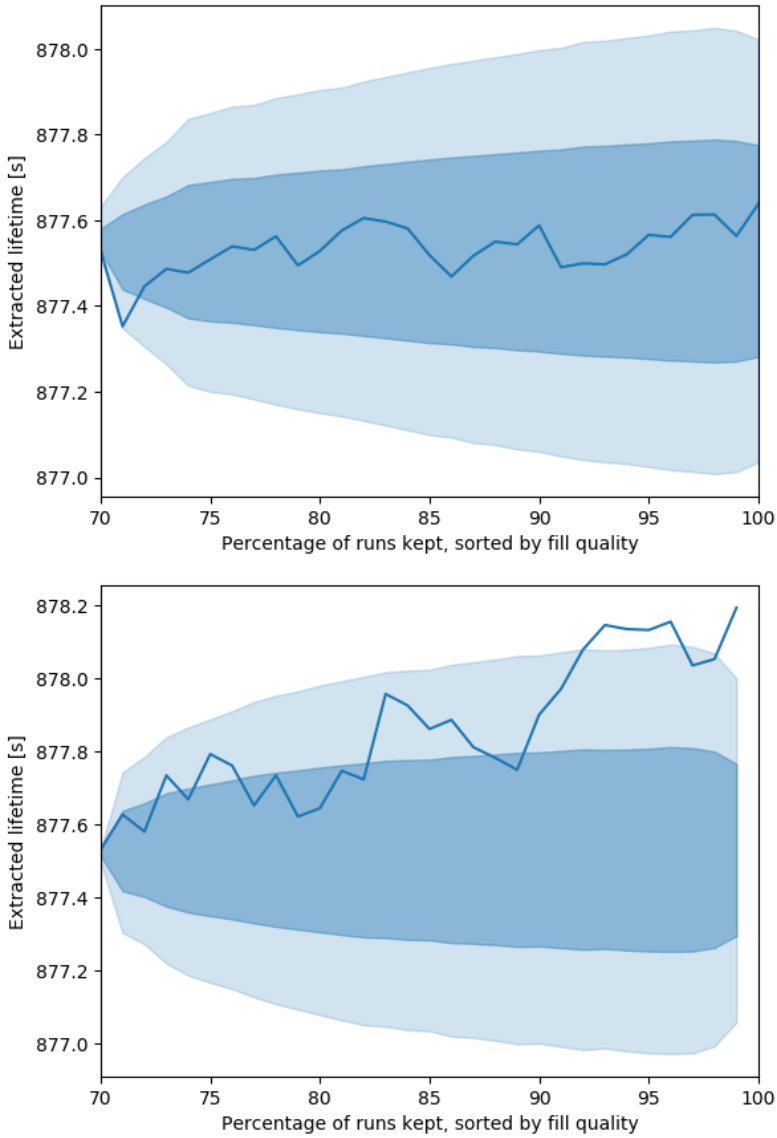

Figure 3.4: Extracted lifetimes as runs of marginally worse fill quality were included in the analyses. The data were separated between 2017 (top) and 2018 (bottom) because of the installation of the buffer volume and the change in running procedures. As additional runs were included in the analyses the extracted lifetimes could vary due to statistical fluctuations alone, but because each additional run made up a small part of each data set the extracted lifetimes should have only slightly varied. The shaded areas denote the 1σ and 2σ statistical uncertainties for how much the extracted lifetime should have varied from a lifetime extracted using just the 70% of runs with the best fill quality, assuming that the additional runs added to the extraction were of good fill quality. If the fill quality of the additional runs was sufficiently bad that it biased the extracted lifetime then the lifetime could have deviated beyond statistical fluctuations.

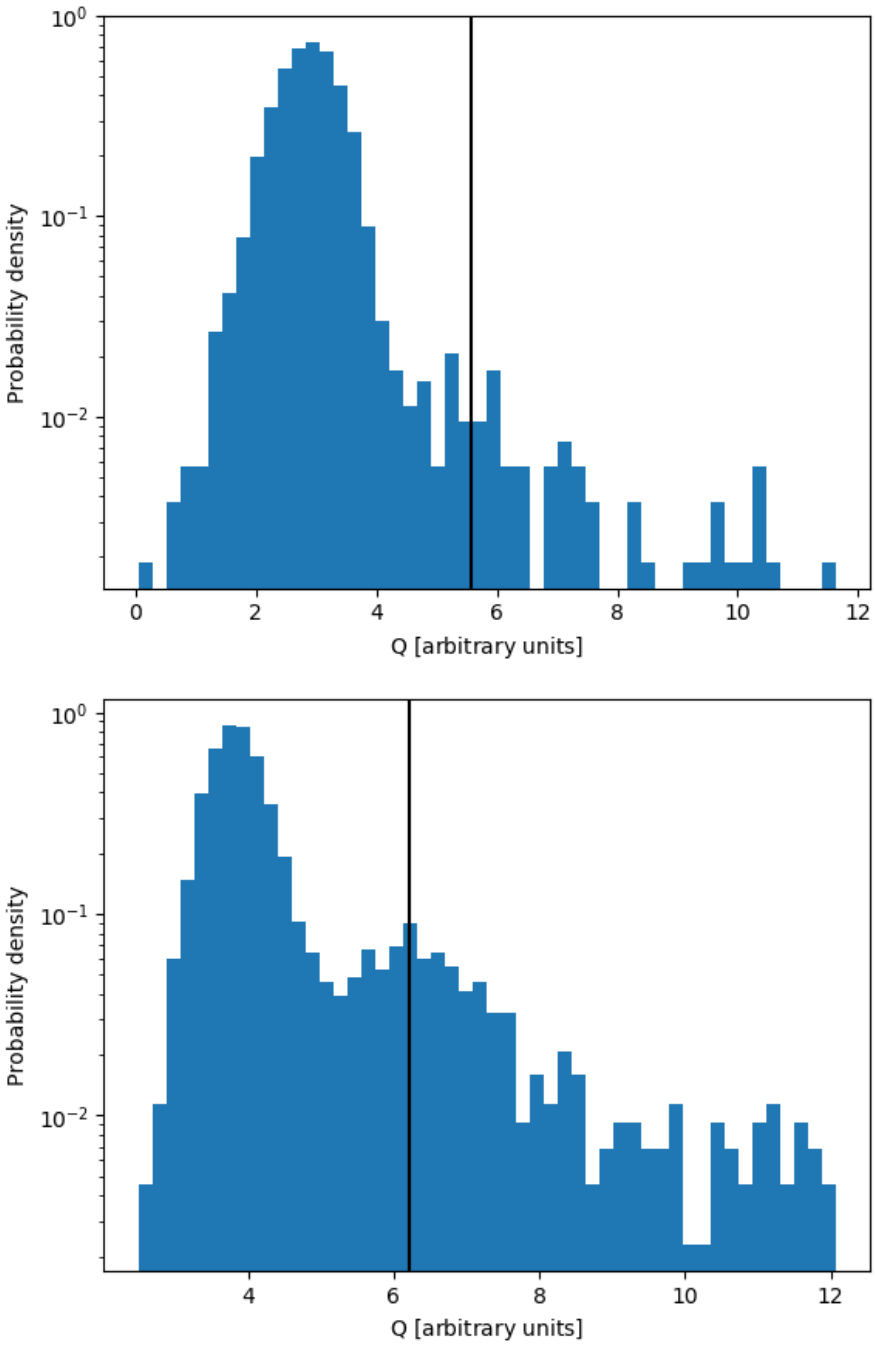

Figure 3.5: Fill quality of each run. Top: fill quality from each run of the 2017 data set. Bottom: fill quality from each run of the 2018 data set. The black lines represent the threshold of fill quality at which runs were rejected from the analysis. Values for Q from the different data sets can not be directly compared to each other because they were calculated by summing over different amounts of data. Changes in running procedures between 2017 (top) and 2018 (bottom) changed the fraction of the runs that were rejected due to poor fill quality.

Across the 2017/2018 data set, nine epochs of primary detector behavior were observed. Table 3.1 lists the nine different epochs, as well as how many usable production and background runs were in each epoch. Estimates of the background rate (Section 3.5) and detector efficiency (Section 3.7.4) of the primary detector were made independently for each epoch of primary detector behavior.

Run numbers	Usable production runs	Usable background runs
4200 - 5440	420	68
5441 - 6730	610	69
6731 - 7611	460	34
7612 - 8592	398	60
8593 - 10606	492	100
10607 - 11310	237	84
11311 - 12424	409	107
12425 - 13171	376	110
13172 - 14731	451	224

Table 3.1: Run numbers of the nine different epochs of primary detector behavior, as well as how many usable runs were in each region.

3.3.5 Conclusion

396 runs were removed from the analysis due to the cuts on fill quality developed in Section 3.3.2. 345 runs were removed from the analysis due to the error in the blinding mechanism discussed in Section 3.3.3, although some of these 345 runs would have been removed for other reasons even if there was no bug in the blinding mechanism. After these two cuts, and another minor cut that will be developed in Section 3.6.4, a total of 3853 runs remained in the analysis. The two cuts developed in this section removed about 16% of the runs from the analysis. If these runs were not problematic and could have been included in the analysis, then a rough scaling suggests that the statistical uncertainty of an extracted lifetime would have decreased by about 8%. The potential marginal decrease in the statistical uncertainty gained by including some of these runs in the analysis did not justify possibly contaminating the analysis with runs of poor quality.

3.4 Identification of UCN events in the primary detector

3.4.1 Introduction

The most important process that took place during a production run was the counting of the UCN in the trap during the unloading process. The primary detector discussed in Section 2.7.2 was used to capture the UCN in the trap. Each UCN that was captured by the primary detector generated a burst of scintillation light that was transformed to voltage signals by two photomultiplier tubes (PMTs). The voltage signal from each PMT was digitized by a discriminator. Throughout this analysis, one digital count in a discriminator is referred to as a photoelectron (PE). Some of the digital counts in the discriminators were caused by the dark rates of the PMTs. Although the digital counts caused by the dark rates were not caused by photons from the scintillator that generated photoelectrons in the cathodes of the PMTs, they will still be referred to as PEs. Generically, all non-UCN induced PEs are referred to as background PEs.

A measure of the number of UCN that were captured by the primary detector during the unloading process of each run could have been constructed by summing the number of PEs recorded during each unloading process and subtracting an estimate for the number of those PEs that were background PEs. The ratio of the number of UCN-induced PEs to the number of background PEs that were recorded during the unloading process of a typical short run was $\sim 13/1$. The same ratio was $\sim 2.3/1$ for a long run. If the number of background PEs could have been accurately estimated then an accurate value for the lifetime could have been extracted, and the large PE backgrounds would have negligibly contributed to the statistical uncertainty of the extracted lifetime. However, a systematic 0.1% bias in the estimates of the number of background PEs would have biased the extracted lifetime by $\simeq 0.15$ s. 0.15 s is about half of the total uncertainty of the lifetime that will be extracted in this analysis, so this level of bias was unacceptable. Insufficient data were gathered during the 2017 and 2018 run cycles to ensure that an estimate of the number of background PEs recorded during each unloading process was not systematically biased at a level that would have induced unacceptably large biases in the extracted lifetime.

Most of the background PEs were caused by uncorrelated dark rates in the two PMTs. In contrast, a captured UCN generated a burst of PEs that were coincident in the two PMTs. This section will discuss how PEs were grouped together to reconstruct UCN events. These events will be referred to as coincidence chains (CCs). Figure 3.6 shows an example of the PEs generated when one UCN was captured by the primary

detector. The goal of the reconstruction of UCN events was to form one CC from patterns of PEs that were similar to those shown in Figure 3.6. The reconstruction process began by searching for coincident PEs in the two PMTs, and then continued to add some number of subsequent PEs to the CC. This reconstruction suppressed the background rate by a factor of $\sim$ 500, and also increased the ratio of UCN-induced events to background events during a typical unloading process by a factor of $\sim$ 10. The backgrounds of CCs will be discussed in Section 3.5.

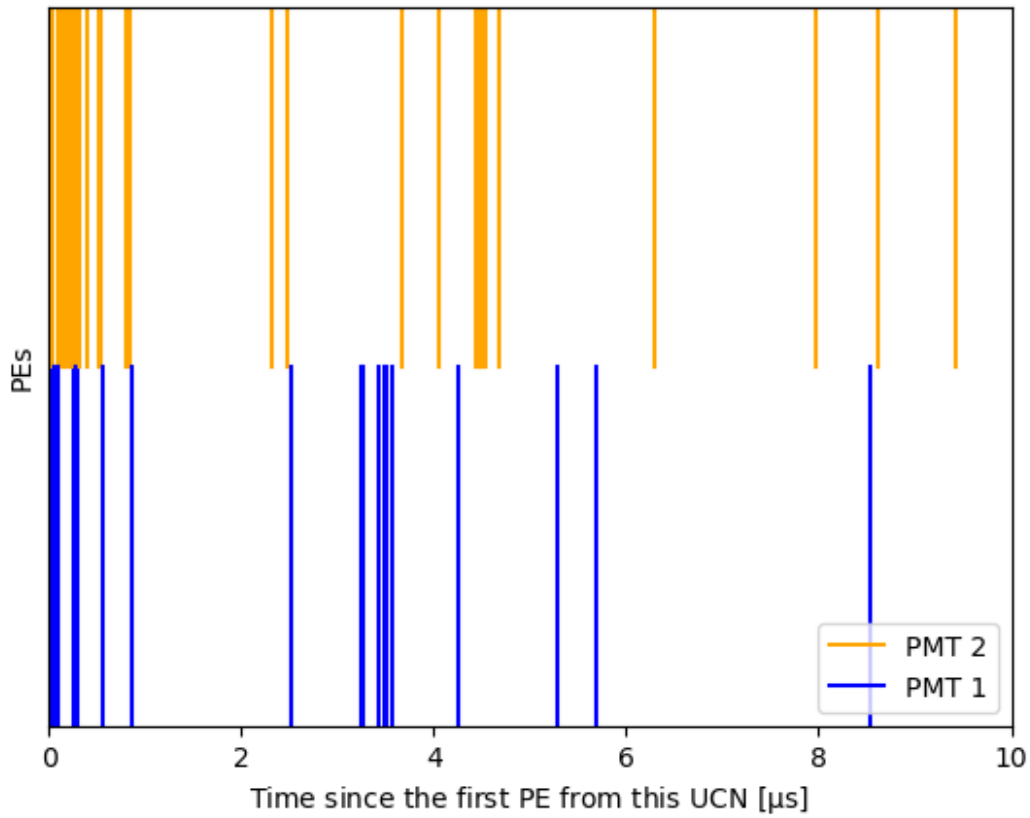

Figure 3.6: PEs generated when one UCN was captured by the primary detector. There was an initial burst of PEs that were coincident in both PMTs, followed by a slowly decaying tail of PEs.

The features of the reconstruction algorithm used to group PEs into CCs should match the distribution of PEs generated when one UCN was captured by the primary detector. However, in order to tune the reconstruction algorithm to match the features of the PEs generated by one UCN there must first be a way to identify when UCN were captured by the primary detector. Section 3.4.2 will use an initial reconstruction algorithm with lenient criteria to identify when UCN were captured by the primary detector, and will then use the identified PEs to construct models of the PEs that

were generated when one UCN was captured by the primary detector. Section 3.4.3 will discuss how these models were used to simulate the capture of multiple UCN on the primary detector in rapid succession. Section 3.4.4 will use the simulated data to develop an algorithm to form CCs. Sections 3.4.5 through 3.4.7 will discuss corrections that needed to be made to the CCs due to interactions between UCN that were captured by the primary detector in rapid succession. Section 3.4.8 will optimize the parameters of the reconstruction algorithm.

3.4.2 Modeling

Data from UCN captures on the primary detector were used to build models of the PEs generated when a UCN was captured by the primary detector. However, in order to select these data, there first must have been some way to identify when a UCN was captured by the primary detector. To prevent biasing the models by selecting only a subset of CCs that did not represent all UCN captures, very lenient standards were used to identify the CCs used to build the models.

This section will show that a UCN captured by the primary detector generated a quick burst of PEs, followed by a long tail of PEs. Algorithm 1 from Appendix B was used to identify CCs that were used to study the characteristics of UCN events. This algorithm required a PE to be found in both PMTs within a coincidence window (T_C = 100 ns). This coincidence criterion significantly suppressed the background PE rate. After finding a coincidence, the algorithm continued to add PEs to the candidate CC as long as the inter-arrival time between consecutive PEs was less than the tail window (T_T = 2000 ns). Once the inter-arrival time exceeded T_T the candidate CC ended. If the number of PEs in the candidate CC was at least the threshold (N_{PE} = 6) then a CC had been found. The final version of the reconstruction algorithm that will be chosen in Section 3.4.8 will be much more stringent than this algorithm. However, for the reasons discussed in the above paragraph, a lenient algorithm was used to identify CCs that were used to build models that described the PEs generated by UCN captured by the primary detector.

In order to model the PEs observed when one UCN was captured by the primary detector, three empirical distributions were built by sampling CCs from the data:

1. The times when UCN were captured on the primary detector, relative to the beginning of Peak 2 of the unloading process.

2. The times when PEs from a given UCN were recorded in either of the PMTs of the primary detector, relative to the time of the first PE of that UCN. Four distributions of this kind were generated by crossing two choices: in which PMT this PE was observed, and in which PMT the first PE of the CC was observed.

3. A three-dimensional distribution for the number of PEs from UCN that were recorded in each of the PMTs of the primary detector, as well as which PMT recorded the first PE in the CC.

Figure 3.7 shows the time distribution of UCN captured by the primary detector during the unloading process. The highest rates of UCN capture occurred between 20 s and 60 s after the start of Peak 2. (See Section 3.2.1 for the definition of Peaks 1-3.) As the rate of UCN capture increased, so does the likelihood that a CC would contain PEs from multiple UCN. In order to avoid contaminating the models of UCN events, CCs between 20 s and 60 s after the start of Peak 2 were excluded. Figure 3.8 shows that the rate of PEs from one captured UCN relaxed to the background PE rate after $\sim$ 10 ms. Therefore, to further reduce contamination of the models of UCN events, CCs that occurred within 16 ms of another CC were excluded to reduce cross-contamination. The 16 ms of data following each remaining CC were collected to form the distributions shown in Figures 3.8 and 3.9. Figures 3.7 through 3.9 were constructed using data from the first epoch of primary detector behavior listed in Table 3.1, but the general behavior was the same for all different epochs.

A few of the key characteristics of UCN events were:

- Half of the UCN counted during the unloading process were counted during a 7.2 s period that contained the time of maximum rate of UCN capture.

- Half of the PEs from a UCN event were detected by the PMTs within 3.4 μs after the first PE was detected.

- The number of PEs in a UCN event could vary widely, from as low as 6 (the chosen threshold) to over 100, with an average of 46.6 PEs per UCN event.

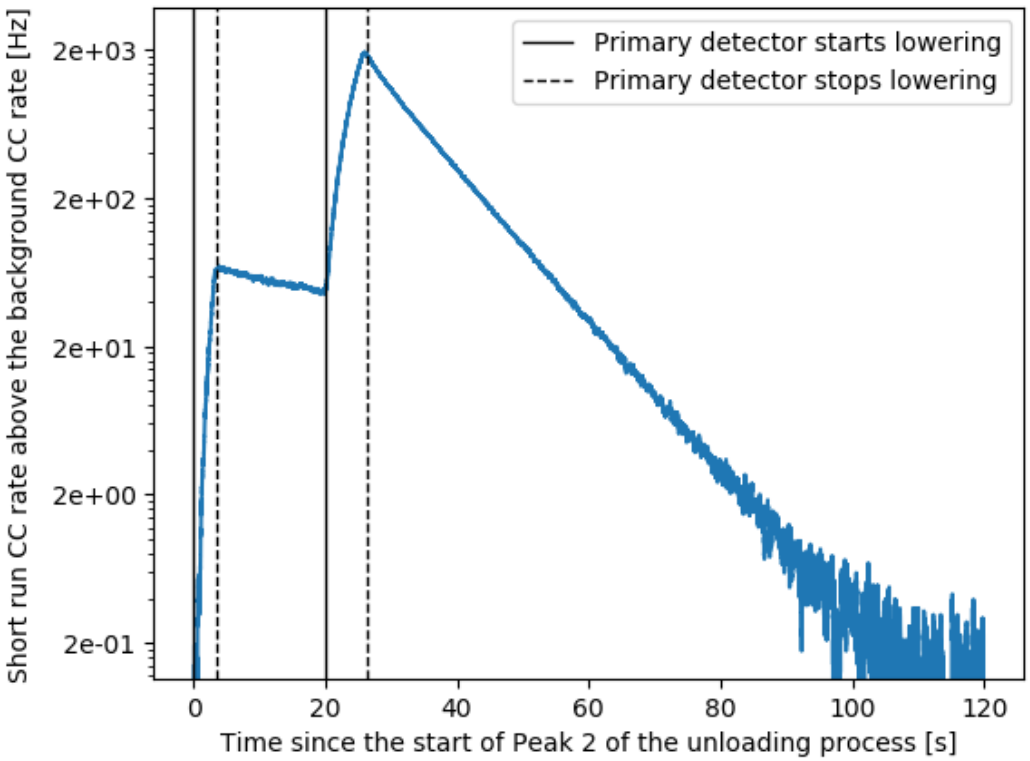

Figure 3.7: Time distribution for capture of UCN on the primary detector. The distribution was made by sampling CCs from the real data that were identified with a very lenient algorithm that was designed to identify UCN that generated even a small number of PEs. The distribution has been smoothed with a rolling average. The sharp increases in CC rate around 0 s and 20 s were caused by the primary detector lowering deeper into the trap, which allowed the primary detector to access many more of the UCN in the trap. At 0 s the primary detector lowered from the Peak 1 height of 38 cm above the bottom of the trap to the Peak 2 height of 25 cm above the bottom of the trap and some of the UCN with sufficiently high energy to reach to primary detector were counted. At 20 s the primary detector lowered from 25 cm above the bottom of the trap to the Peak 3 height 1 cm above the bottom of the trap and all remaining UCN were counted.

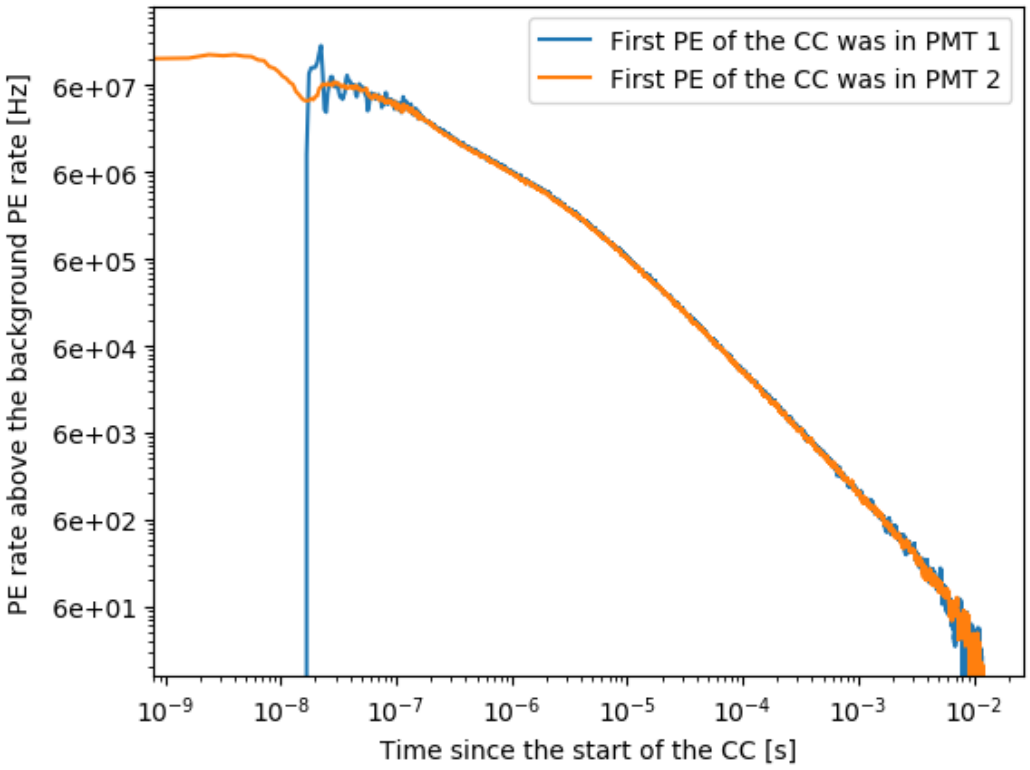

Figure 3.8: Time distribution for detection of PEs in PMT 1 from one UCN. The time distributions for detection of PEs in PMT 2 from one UCN (not shown) were nearly identical to the distributions shown in this figure. The distributions were made by sampling CCs from the real data that were identified with a very lenient algorithm that was designed to identify UCN that generate even a small number of PEs. If the first PE of the CC was recorded in PMT 1 then no more PEs could have been recorded in PMT 1 for 16 ns due to hardware dead time (blue curve), but if the first PE of the CC was recorded in PMT 2 then PEs could have been recorded in PMT 1 during the initial 16 ns (orange curve). The tail of the distributions have been smoothed with rolling averages. After initial differences due to hardware dead time, the two probability density curves became very similar. After $\sim 10^{-2}$ s the PE rate from one UCN event had relaxed to the background PE rate.

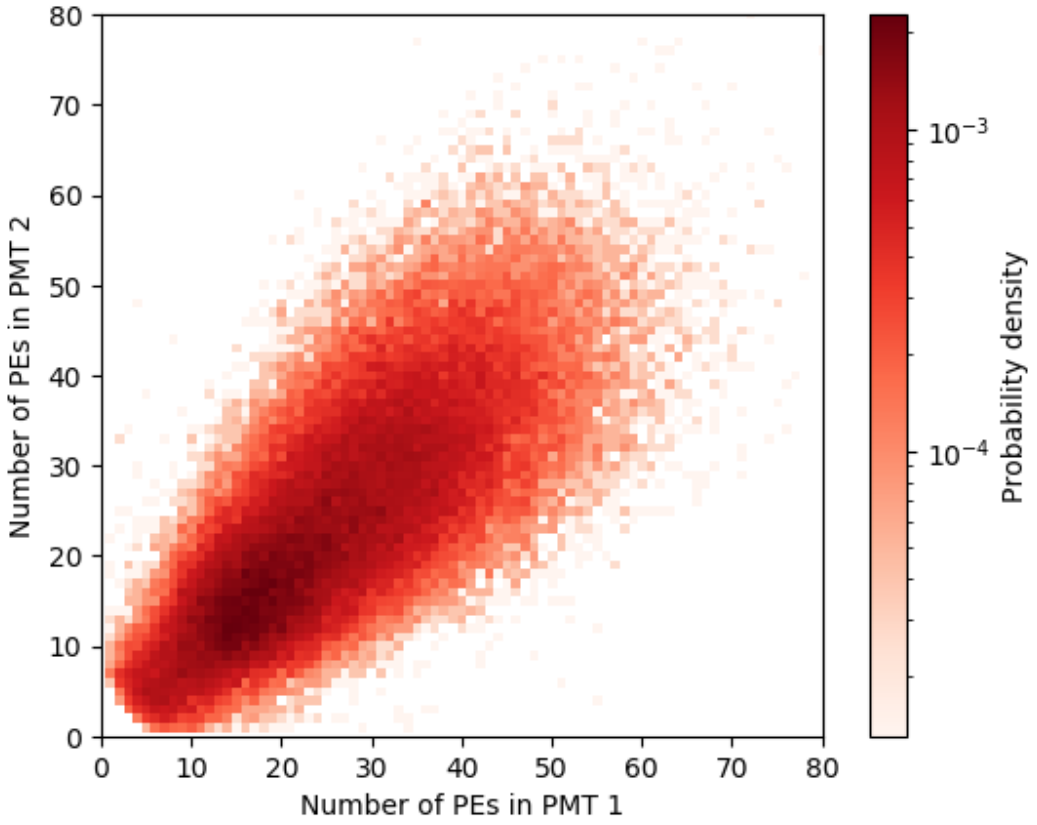

Figure 3.9: Distribution for the number of PEs from one UCN event that were recorded in each of the two primary detector PMTs. The distribution was made by sampling CCs from the real data that were identified with a very lenient algorithm that was designed to identify UCN that generated even a small number of PEs. The region of zero probability density nearest the origin was due to the choice of threshold of $N_{\mathrm{PE}} = 6$ PEs to qualify as a CC.

3.4.3 Simulating UCN events using Monte Carlo simulations

Section 3.4.2 used real data from the primary detector to build three types of empirical distributions that described the PEs generated by one UCN that was captured by the primary detector. This section will use those models to simulate the PEs generated when multiple UCN were captured by the primary detector. Of particular interest are UCN that were captured in rapid succession, where the PEs from multiple UCN were intermingled. This intermingling of PEs made it more difficult to properly reconstruct events. Sections 3.4.5 through 3.4.7 will develop corrections for errors in reconstructions caused by these difficulties. The probability that multiple UCN were captured by the primary detector in rapid succession increased with the overall number of UCN in the trap, so short and long runs were affected differently.

The following process was used to sample the empirical distributions from Section 3.4.2 to generate Monte Carlo data:

1. N UCN captures on the primary detector were drawn from the distribution in Figure 3.7.

2. For each of the N simulated UCN captures:

    a) One entry from the distribution in Figure 3.9 was drawn for the number of PEs in PMT 1 (n_1), the number of PEs in PMT 2 (n_2), and which PMT recorded the first PE of the UCN event.

    b) n_1 time of detection of PEs in PMT 1 generated by one UCN event were drawn from a distribution similar to those in Figure 3.8, depending on which PMT recorded the first simulated PE of this event. If any of the independently drawn PEs were within 16 ns (the hardware dead time) of each other, then the latter PE that did not respect the hardware dead time was removed and a replacement was drawn. This process continued until a full set of n_1 PEs that respected the hardware dead time were drawn.

    c) The previous step was repeated to draw n_2 values for the time of detection of PEs in PMT 2.

3. A sorted list of all PEs in PMT 1 was created. All PEs that did not respect the hardware dead time were removed, but replacements were **not** drawn.

4. The previous step was repeated for all PEs in PMT 2.

Steps 2b and 2c did draw replacements for PEs removed by the 16 ns hardware dead time, but Steps 3 and 4 did not. Steps 2b and 2c simulated the PEs from one UCN captured by the primary detector. These simulations were based on empirical distributions built from the real PE data observed when one UCN was captured by the primary detector. As such, the data that built these empirical distributions naturally included the effect of the 16 ns hardware dead time. If Steps 2b and 2c did not draw replacements for PEs removed by the 16 ns hardware dead time, then the effect would have been erroneously applied twice. In contrast, Steps 3 and 4 simulated how PEs from different UCN intermingled. In order to simulate the effect where higher rates of UCN capture led to more intermingling of PEs and more hardware-dead-time interactions, PEs from one UCN that were removed by the 16 ns hardware dead time were not replaced.

Values for N_{short} and N_{long} were chosen that approximately matched the mean unload counts per run and approximately matched τ. In contrast to a physical experiment, each simulated run had exactly N_{short} (or N_{long}) simulated UCN unloaded. This eliminated variations due to statistical fluctuations in the filling and holding processes and allowed the Monte Carlo simulation to focus on variations due to reconstruction alone.

3.4.4 Reconstructing UCN events from PEs

In analyses of past UCNτ data sets, the "standard" UCN reconstruction algorithm was Algorithm 1 from Appendix B, with $T_C = 100$ ns, $T_T = 1000$ ns, and $N_{\text{PE}} = 8$. (See Section 3.4.2 for a brief explanation of how the algorithm worked.) This algorithm was used to reconstruct simulated UCN events to study the properties of the reconstruction.

1000 pairs of runs were generated based on the models built from runs 4200-5440. The short runs had $N_{\text{short}} = 17382$ UCN (modeled after a $t_{\text{short}} = 20$ s run) and the long runs had $N_{\text{long}} = 3055$ UCN (modeled after a $t_{\text{long}} = 1550$ s run). These four constants corresponded to a value of the true Monte Carlo lifetime of

$$\tau_{\text{MC}} \equiv \frac{t_{\text{long}} - t_{\text{short}}}{\ln \frac{N_{\text{short}}}{N_{\text{long}}}} \simeq 879.990 \text{ s}.$$

The values for N_{short} and N_{long} were chosen as the mean of the number of UCN that were detected at the end of the 20 s and 1550 s runs with run numbers between 4200 and 5440. Each simulated PE was tagged with the corresponding simulated UCN so that it was possible to analyze PEs from each simulated UCN on their own. This significantly reduced the UCN-to-UCN interactions, but did not eliminate them entirely due to the simulated 16 ns hardware dead time.

Even in the complete absence of UCN-to-UCN interactions, a reconstruction algorithm could still have flaws. An ideal reconstruction algorithm would result in exactly one CC per UCN. An imperfect algorithm sometimes fails to reconstruct a UCN event. Also, an imperfect algorithm sometimes *retriggers* and finds multiple CCs when there was only one UCN. Analysis of a single UCN that resulted in ≤ 1 reconstructions had 0 retriggers, but analysis of a single UCN that resulted in $k \geq 2$ reconstructions had $k - 1$ retriggers.

Table 3.2 presents three key metrics that quantify the performance of the reconstruction algorithm. A lifetime was extracted by scaling N_{short} and N_{long} by the appropriate values for the number of CCs per UCN. The values reported in Table 3.2 induced a -0.196 s bias in the extracted lifetime. The difference in CCs per UCN between short and long runs was due to the difference in retriggers per UCN, and not due to a difference in the fraction of UCN events that were not reconstructed.

	Short runs	Long runs
Fraction of UCN not reconstructed	0.08100 ± 0.00006	0.08100 ± 0.00016
Retriggers per UCN	0.03575 ± 0.00004	0.03538 ± 0.00011
CCs per UCN	0.95475 ± 0.00008	0.95438 ± 0.00019

Table 3.2: Performance metrics of Algorithm 1 from Appendix B when used with $T_C = 100$ ns, $T_T = 1000$ ns, and $N_{\text{PE}} = 8$ to reconstruct UCN events. The sum of the fraction of UCN events that were not reconstructed and the CCs per UCN, less the retriggers per UCN, was identically equal to 1. These metrics were calculated from simulated data that were generated using the procedure from Section 3.4.3.

Recall that these performance metrics were extracted after significantly reducing UCN-to-UCN interactions in Monte Carlo simulations. The presence of these UCN-to-UCN interactions could have exacerbated the difference between short and long runs, which would have further increased the bias in the extracted lifetime.

In order to suppress retriggering, Algorithm 1 from Appendix B was modified to Algorithm 2 from Appendix B. The major change was the introduction of the prompt window. The reconstruction algorithm with a prompt window required at least a threshold number of PEs within a fixed amount of time at the beginning of a CC. In contrast, the reconstruction algorithm without a prompt window only required the total number of PEs in a CC to be at least as large as the threshold.

Algorithm 1, without a prompt window

1. Require each of the two primary detector PMTs to observe at least one PE within T_C (the coincidence window) of each other.
2. As long as the inter-arrival time between PEs is less than T_T (the tail window), then continue adding PEs to the **candidate** event.
3. If the number of PEs in the candidate event is at least N_{PE} (the threshold) then a **reconstructed** UCN event has been found.

Algorithm 2, with a prompt window

a. Require each of the two primary detector PMTs to observe at least one PE within T_C (the coincidence window) of each other.

b. Require at least N_{PE} (the threshold) PEs to be observed within T_P (the prompt window) from the first observed PE in this candidate event, without restriction as to which PMT observes the PEs.

c. As long as the inter-arrival time between PEs is less than T_T (the tail window), then continue adding PEs to the **reconstructed** UCN event.

The Monte Carlo data previously analyzed in this section was reanalyzed with the new reconstruction algorithm. The same key metrics are presented in Table 3.3.

	Short runs	Long runs
Fraction of UCN not reconstructed	0.14375 ± 0.00008	0.14350 ± 0.00020
Retriggers per UCN	0.00183 ± 0.00001	0.00190 ± 0.00002
CCs per UCN	0.85809 ± 0.00008	0.85837 ± 0.00020

Table 3.3: Performance metrics of Algorithm 2 from Appendix B when used with $T_C = 100$ ns, $T_P = 1000$ ns, $T_T = 1000$ ns, and $N_{\mathrm{PE}} = 8$ to reconstruct UCN events. Note that the sum of the fraction of UCN events that were not reconstructed and the CCs per UCN, less the retriggers per UCN, was identically equal to 1. These metrics were calculated from simulated data that were generated using the procedure from Section 3.4.3.

There was no evidence of a difference in the fraction of UCN events that were not reconstructed between short and long runs. However, there was evidence (p-value = 0.016) that the retriggers per UCN differed between short and long runs. (This p-value is the probability of observing a difference in retriggers per UCN between short and long runs that was larger than the one that was actually observed in Monte Carlo simulations. This probability was calculated under the assumption that there was no difference in the retriggers per UCN between short and long runs. Future use of p-values will not explicitly restate how the p-values were calculated.) If all of the data from short and long runs were combined, the fraction of UCN events that were not reconstructed was 0.14371 ± 0.00008. Combining this value with the measured values for retriggers per UCN for short and long runs led to a bias of 0.038 s in the extracted lifetime. This bias was significantly reduced when compared to the bias in the extracted lifetime of -0.196 s from earlier in this section.

The key update to the reconstruction algorithm was the introduction of the prompt window. The prompt window required a CC to have a quick burst of PEs. This matched the shape of the distributions shown in Figure 3.8. The likelihood of finding a quick burst of PEs in the tail of the distributions for PEs within a UCN event was much lower than finding a similar burst at the beginning of a UCN event. In contrast, a reconstruction algorithm without a prompt window could have reconstructed an event that consisted of one PE in each of the PMTs of the primary detector, followed by a slow, but steady, trickle of PEs. This was still less likely to occur in the tail of a UCN event than at the beginning, but this reconstruction criterion was much less strict than the prompt window so it occurred more often and increased the number of retriggers per UCN. The addition of the prompt window decreased the fraction of UCN events that were reconstructed, but Section 3.11 will show that this was a small contribution to the overall statistical uncertainty of the measurement of τ.

3.4.5 Software dead time

In addition to the 16 ns hardware dead time of the discriminators that digitized the voltage signals from the primary detector PMTs, the reconstruction algorithm had software dead time. If a second UCN was captured by the primary detector prior to the end of the reconstruction of a first UCN, then the two UCN events would have been reconstructed as one CC. This effect was more significant at higher rates of UCN capture. Short runs had more UCN in the trap during the unloading process than long runs, so runs of different lengths were affected differently. If not corrected for, this would have induced a bias in the extracted lifetime.

Three simplifying assumptions were initially made in order to estimate the number of events that were missed due to either software or hardware dead time. Flaws in these assumptions will be explored in the next few paragraphs.

1. The dead time of the detector was *constant*: each event had the same amount of dead time.

2. The dead time of the detector was *non-paralyzable*: if the detector was "dead" and another event occurred in the detector, then the dead time did not reset. Instead, the dead time ended when it would have had the second event not occurred.

3. $\delta \ll r^{-1}$, where δ is the amount of dead time per event and r is the rate of events.

$\delta \ll r^{-1}$ implies that there was approximately no overlap among events in the detector, so the fraction of time that the detector was 'live" was $1 - r\delta$. Therefore, the corrected rate of events in the detector was

$$r' = \frac{r}{1 - r\delta}. \tag{3.4}$$

As long as the three assumptions held true, Equation 3.4 could be used to correct the rate of PEs for the hardware dead time of the discriminators and to correct the rate of CCs for the software dead time of the reconstruction algorithm.

Individual PEs did follow all three of the assumptions laid out in the above paragraph. However, the structure of the CCs did not match the first two assumptions. First, the length of each CC depended on the inter-arrival time between PEs. A typical CC had a length of a few μs. Second, the reconstruction algorithm was paralyzable. If a constant stream of UCN were captured by the primary detector at a rate of 1 MHz, and the length of each CC was >1 μs, then the reconstruction algorithm would have collected all of the UCN into one CC, and would have continued to do so until the UCN rate decreased. Although the structure of the CCs did not match the first two assumptions, an adapted version of Equation 3.4 was used to apply a correction to the CCs to account for the software dead time. Section 3.4.7 will use Monte Carlo simulations to test the validity of this correction.

Equation 3.4 assumes that the rate of CCs was constant throughout the entire run. Figure 3.7 shows that the rate of CCs varied by almost four orders of magnitude during the 210 s unloading process of a typical short run. The following algorithm was used to construct instantaneous estimates of the rate of CCs, and then used those instantaneous rates to apply a software-dead-time correction to each CC in the unloading process of each run:

1. A list of the CCs was constructed, where the start time t_i and length δ_i were known for each CC.

2. A histogram C was constructed of the t_is with bin width w. This histogram summed the number of CCs that occurred in each bin.

3. A histogram Δ was constructed of the t_is with bin width w, where each entry was weighted by δ_i. This histogram summed the total software dead time in each bin.

4. The fraction of time in the j^{th} bin that was "live" time was $\frac{w-\Delta_j}{w}$, so the number of dead-time-corrected CCs in each bin was $C'_j = \frac{wC_j}{w-\Delta_j}$

Section 3.4.7 will show that $w = 0.1$ s was a reasonable choice for the bin width. This value balanced two considerations:

- A large bin width ensured that each bin had a sufficient number of CCs to suppress statistical fluctuations between the measured rate of CCs and the expected rate of CCs. The maximum rate at which UCN were captured by the primary detector was ~ 4 kHz, so the measured rate of CCs in a histogram with a bin width $w \lesssim \frac{1}{4\text{ kHz}} = 0.25$ ms would have been strongly dependent on statistical fluctuations of the time of capture of UCN by the primary detector.

- A small bin width ensured that changes in the expected rate of CCs within one bin as the unloading process progressed were sufficiently small as to not bias the software-dead-time correction applied to each CC in that bin. The rate at which UCN were captured by the primary detector changed much faster than 10 s (see Figure 3.7), so the measured rate of CCs in a histogram with a bin width $w = 10$ s would not have accurately reflected the probability that an individual CC would have been missed due to software dead time.

The start time t_i of each CC was the time of the first PE in that CC, measured relative to the beginning of the unloading process. The length of each CC δ_i was the length of time between the first and last PEs in the CC, plus the tail window T_T. The tail window was included in the length because the reconstruction algorithm continued to search for T_T after the last PE in the CC, and if any additional PEs were found then the reconstruction would have continued. Therefore, it was not possible to begin a new CC during that time so it was part of the software dead time. Section 3.4.7 will show that $w = 0.1$ s was a reasonable choice, and will also explore the effect of choice of w on on the extracted lifetime.

Figure 3.10 shows the scale of the correction to the CCs to account for software dead time as a function of the total number of CCs found during the unload of each run. For runs with a large number of UCN in the trap, up to 1% of CCs were missed due to software dead time. Short runs had significantly more events missed due to software dead time because there were significantly more UCN in the trap during the unloading process. The installation of the buffer volume between the 2017 and 2018 run cycles (around run number 9000) led to significantly fewer UCN entering the trap. This decreased the fraction of CCs missed due to software dead time because the maximum UCN counting rate was lower in 2018 than in 2017.

Each CC resulted in software dead time. There were many opportunities for CCs to be missed due to software dead time, but the probability that any one CC was missed due to software dead time was small. Poisson uncertainties were assigned to the estimated number of CCs missed due to software dead time.

3.4.6 UCN-event-tail interactions

Figure 3.8 shows that it took $\sim$ 10 ms for the PEs from a UCN event to relax to the background PE rates of the PMTs in the primary detector. If each CC lasted 10 ms, then the software dead time would have been unacceptably large. Half of the PEs in a UCN event occurred within the first $\sim$ 3.4 μs of the UCN event, so a relatively short CC could still capture the majority of the PEs from a UCN event.

If two UCN events occurred within $\sim$ 10 ms of each other, some of the PEs from the first UCN event that were not captured in the first CC could have intermingled with the second UCN event. This intermingling could have increased the probability that the second UCN event was reconstructed, and could have also increased the probability that the reconstruction algorithm retriggered on the second UCN event. Either of these changes could have induced a bias in the extracted lifetime.

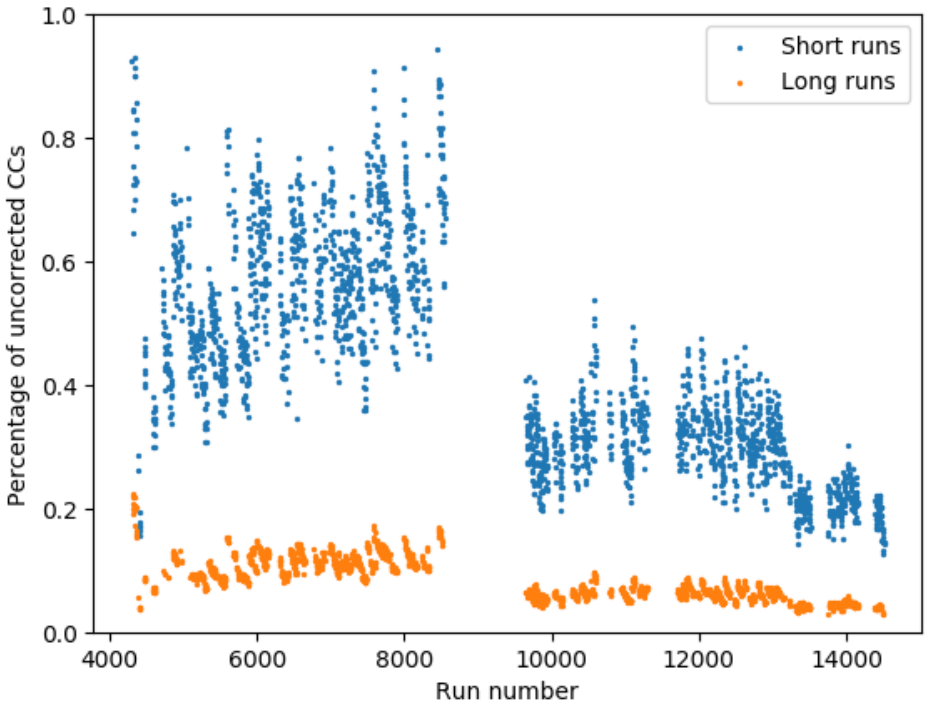

Figure 3.10: Estimated number of CCs missed due to software dead time, expressed as a percentage of the number of uncorrected CCs. These estimates were calculated from the real PE data from each run. The step change in behavior around run number 9000 occurred between the 2017 and 2018 run cycles. Fewer UCN were loaded into the trap during 2018, which also decreased the effects of the software dead time.

Algorithm 3 from Appendix B was used to correct for these UCN-event-tail interactions. This algorithm estimated the number of PEs in the prompt window of the current CC that came from previous UCN events. Then, the algorithm compared this estimate to the total number of PEs recorded in the prompt window of the current event, as well as the threshold number of PEs needed to be considered a CC. The current CC was de-weighted by the probability that the number of the PEs recorded in the prompt window of this CC that came from another UCN event were sufficient to push the current event below threshold.

The correction for UCN-event-tail interactions made two assumptions:

1. The PEs in a given UCN event occurred independently of one another (aside from the 16 ns of hardware dead time).

2. The distribution for PEs in a UCN event followed the distributions shown in Figure 3.8, regardless of the number of PEs in the UCN event or the location of the UCN event in the primary detector.

The first assumption could have been false if the discriminator that digitized the voltage signals from the primary detector PMTs was prone to retriggering. The second assumption could have been false if different areas of the primary detector had different UCN capture efficiency or different scintillation properties. Section 4.4 will discuss how the uncertainty of the validity of these assumptions translated to a systematic uncertainty of the measured lifetime, and Section 3.4.8 will discuss how this systematic uncertainty of the measured lifetime was minimized.

If the two assumptions in the previous paragraph held true, then the UCN-event-tail correction algorithm was theoretically sound in the case that the neighboring CCs corresponded to different UCN events. However, the correction fails when the reconstruction algorithm retriggered and the two CCs corresponded to the same UCN event. If the number of PEs in the tail of a UCN event was $\lesssim$ the expected number, then the reconstruction algorithm likely did not retrigger and no UCN-event-tail correction was applied. If the number of PEs in the tail of a UCN event was $\gg$ the expected number, then the reconstruction algorithm could have retriggered. The UCN-event-tail correction was proportional to the expected number of PEs in the tail, but retriggering only occurred when the actual number of PEs in the tail was much greater than expectations, so the UCN-event-tail correction underestimated the correction needed to eliminate the problem of retriggering. Nonetheless, the UCN-event-tail correction did reduce the problem of retriggering. Any bias in the extracted lifetime due the failures of the UCN-event-tail correction will be explored in Section 3.4.7.

Figure 3.11 shows the scale of the UCN-event-tail correction as a function of the total number of CCs found during the unload of each run. The fraction of extra CCs due to UCN-event-tail interactions was greater in short runs with higher UCN counting rates, but the difference between short and long runs was much smaller than the difference in CCs missed due to software dead time between short and long runs. Most of the UCN-event-tail corrections were due to retriggering, and not due to UCN-to-UCN interactions, so the UCN counting rate did not significantly affect the fraction of extra CCs due to UCN-event-tail interactions. There was a decrease in the fraction of extra CCs due to UCN-event-tail interactions after run number 13171, when the primary detector was damaged and the number of PEs generated per captured UCN decreased.

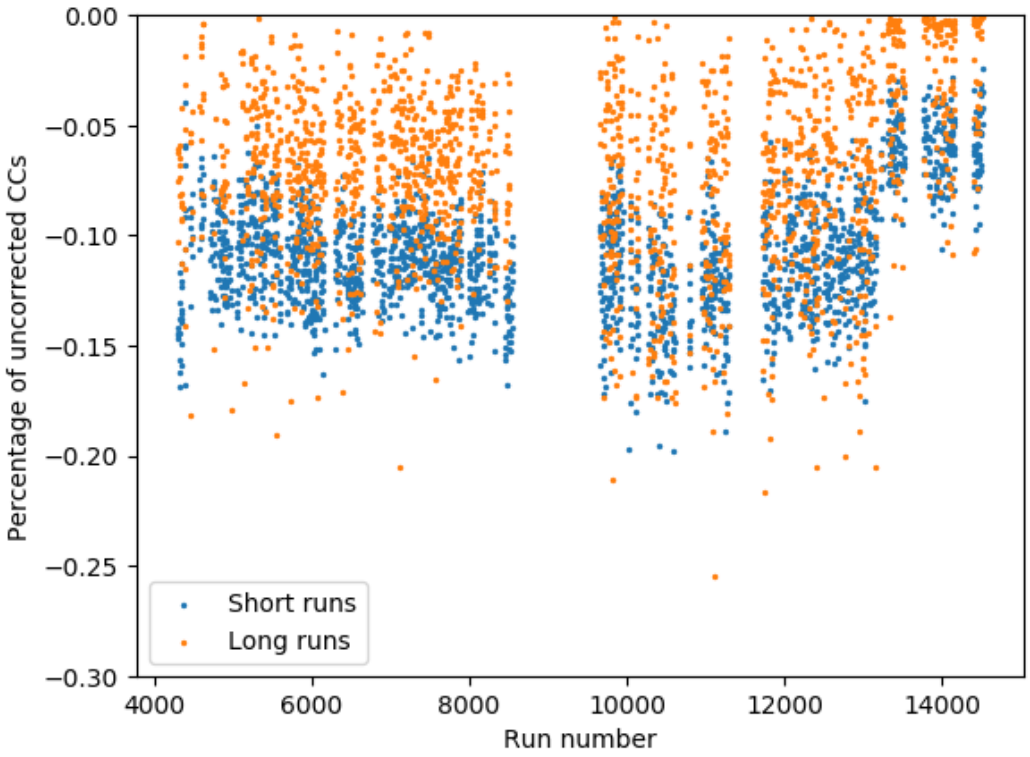

Figure 3.11: The UCN-event-tail correction to the number of CCs, expressed as a percentage of the number of uncorrected CCs. These estimates were calculated from the real PE data from each run. The UCN-event-tail correction was negative because it decreased the number of CCs.

Each CC had a tail of PEs that were not captured by the reconstruction algorithm. There were many opportunities for the reconstruction algorithm to have retriggered on these tails, but the probability of any one tail being retriggered on was $\simeq 2\times10^{-3}$. Poisson uncertainties were assigned to the estimated number of CCs missed due to UCN-event-tail interactions.

3.4.7 Testing the software-dead-time and UCN-event-tail corrections on simulated UCN events

Algorithm 2 from Appendix B was used with $T_C = 100$ ns, $T_P = 1000$ ns, $T_T = 1000$ ns, and $N_{\mathrm{PE}} = 8$ to create CCs from the simulated UCN event data generated in Section 3.4. These CCs were corrected for software dead time and UCN-event-tail interactions using the methods detailed in Sections 3.4.5 and 3.4.6. Section 3.4.5 describes the software-dead-time correction algorithm for a general bin width w. Different bin widths were tested to determine the optimum value for w.

Figure 3.12 shows the impact that the software-dead-time and UCN-event-tail-interaction correction algorithms had on the extracted lifetime. An additional 4000 pairs of simulated runs were generated, for a total of 5000 short-long pairs. After the software-dead-time correction with a bin with of $w = 0.1$ s and a UCN-event-tail correction were applied to the reconstructed events, the extracted lifetime was 0.002 ± 0.058 s greater than the true Monte Carlo lifetime. Section 3.4.5 explained why very large and very small values of w induced a bias in the extracted lifetime. A bin width of $w = 0.1$ s was used in all further software-dead-time corrections of CCs.

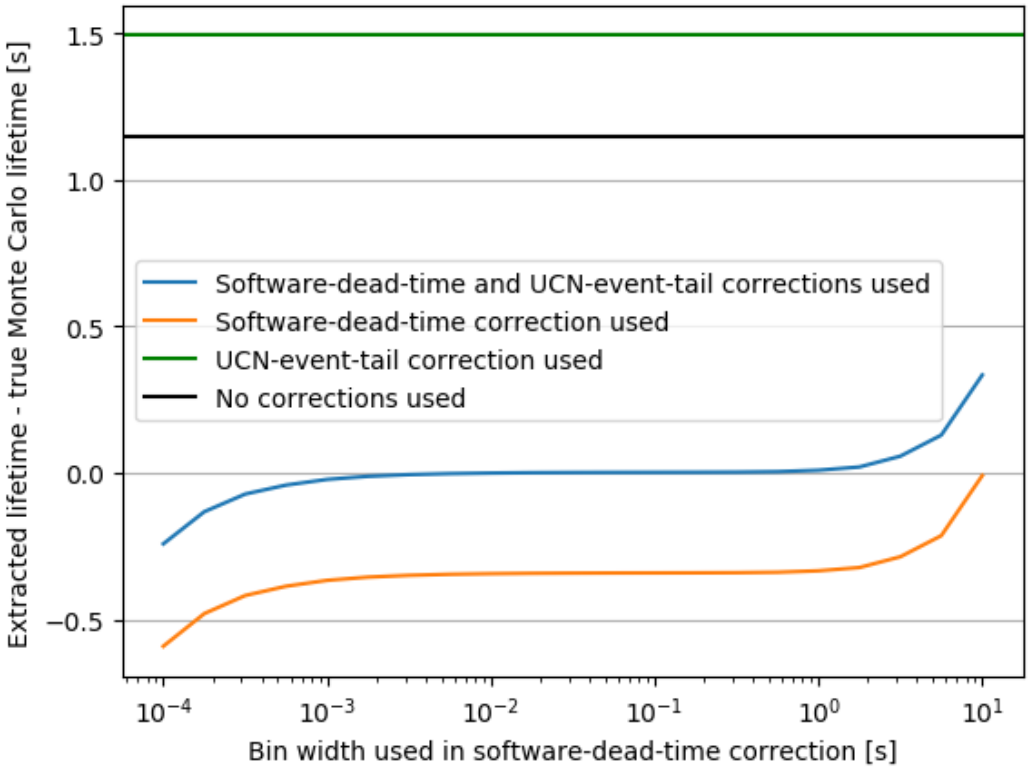

Figure 3.12: Extracted lifetimes from 5000 simulated short-long pairs of runs, with and without the use of various CC correction algorithms. All data shared a common 0.058 s statistical uncertainty that is due to generating a finite amount of simulated data. These lifetimes were extracted from simulated data that were generated using the procedure from Section 3.4.3.

3.4.8 Choice of the parameters of the reconstruction algorithm

Section 3.4.7 shows that when the software-dead-time and UCN-event-tail corrections were applied to the CCs, there was no statistical evidence of a difference between the lifetime extracted from simulated UCN events and the true Monte Carlo lifetime. Those reconstructions were done using Algorithm 2 from Appendix B, with $T_C = 100$ ns, $T_P = 1000$ ns, $T_T = 1000$ ns, and $N_{\mathrm{PE}} = 8$ PEs. To study the effect of the choice of parameters on the extracted lifetime, 125 different combinations of parameters of the reconstruction algorithm were used to reconstruct events from the simulated data generated in Section 3.4.3. The coincidence window was fixed at $T_C = 100$ ns, but the other three parameters were swept over a range of reasonable values:

- The prompt window T_P was swept over 600, 800, 1000, 1200, and 1400 ns;
- The tail window T_T was swept over 600, 800, 1000, 1200, and 1400 ns; and
- The threshold N_{PE} was swept over 6, 7, 8, 9, and 10 PEs.

Figure 3.13 shows the variation in extracted lifetime that arose from the choice of parameters of the reconstruction algorithm. There was a class of parameters where the software-dead-time and UCN-event-tail corrections failed to appropriately correct the CCs and led to an error in extracting the lifetime. Eliminating choices of parameters that led to $> 10^{-2}$ retriggers per UCN removed most of the tail of the distribution for differences between extracted lifetimes and the true Monte Carlo lifetime.

Figure 3.14 provides additional insight as to why some extracted lifetimes remained in the left tail of the distribution shown in Figure 3.13. For any choice of parameters T_P and N_{PE} of the reconstruction algorithm, a decrease in T_T increased the number of retriggers per UCN. This is most clearly seen in Figure 3.14 when looking at the set of five data points with an extracted lifetime that was ~ 0.15 s less than the true Monte Carlo lifetime, and had between 10^{-7} and 10^{-4} retriggers per UCN. These five data points shared common values for T_P and N_{PE} and swept over the five explored values for T_T. In contrast to this type of behavior, the extracted lifetime for reconstruction algorithms that resulted in high number of retriggers per UCN was affected by changes in T_T because the high number of retriggers per UCN induced more bias in the extracted lifetime.

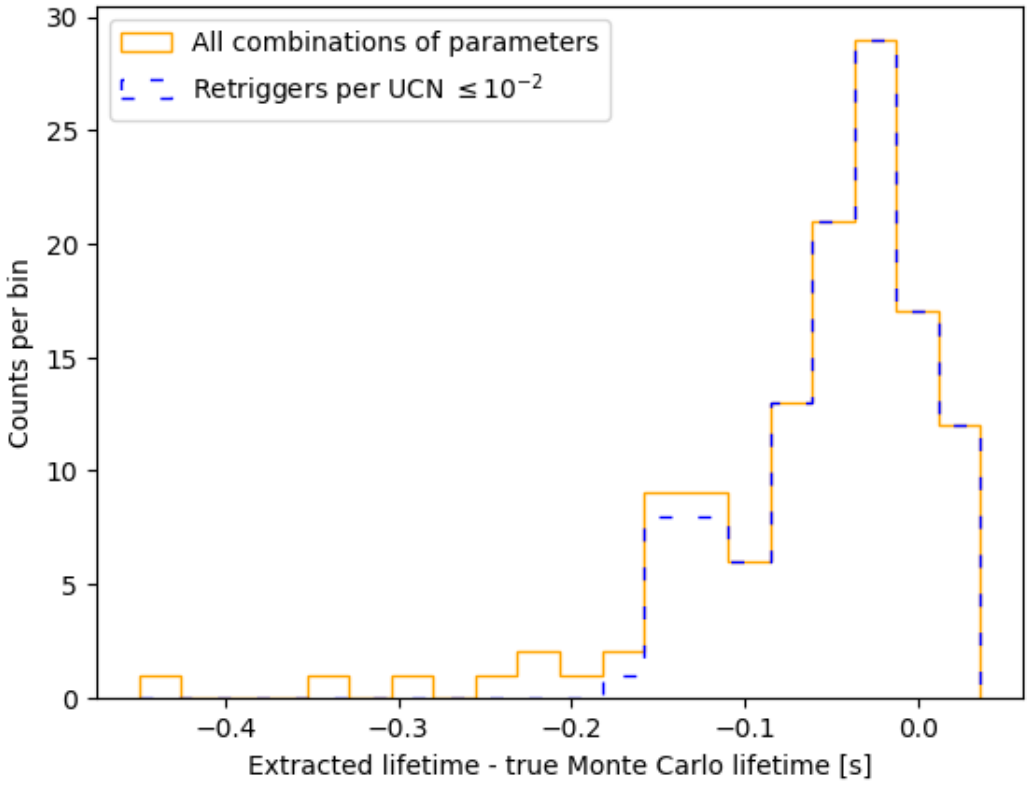

Figure 3.13: Distribution for extracted lifetimes for various choices of the parameters of the reconstruction algorithm. After removing choices that resulted in $> 10^{-2}$ retriggers per UCN, the remaining extracted lifetimes had a mean difference of -0.04 s from the true Monte Carlo lifetime. Each extracted lifetime had a statistical uncertainty of $\simeq 0.06$ s. The lifetimes were extracted from the same underlying Monte Carlo data, so the statistical uncertainties were very highly correlated. Therefore, the mean difference of -0.04 s was not significantly different from 0.

In contrast to the values presented in Tables 3.2 and 3.3, the estimates for the number of retriggers per UCN and the number of CCs per UCN used to generate Figure 3.13 did include full UCN-to-UCN interactions, and also included the software-dead-time and UCN-event-tail corrections that corrected for these interactions. Some retriggering was tolerable, but $> 10^{-2}$ retriggers per UCN led to a bias in the extracted lifetime.

One set of parameters for the coincidence algorithm must have been chosen to use in this analysis. The set $T_P = 600$ ns, $T_T = 1400$ ns, and $N_{\mathrm{PE}} = 6$ PEs was chosen because, of all the combinations tested, it minimized corrections to account for UCN-event-tail interactions. The two assumptions made in Section 3.4.6 used to calculate the corrections for UCN-event-tail interactions were reasonable, but

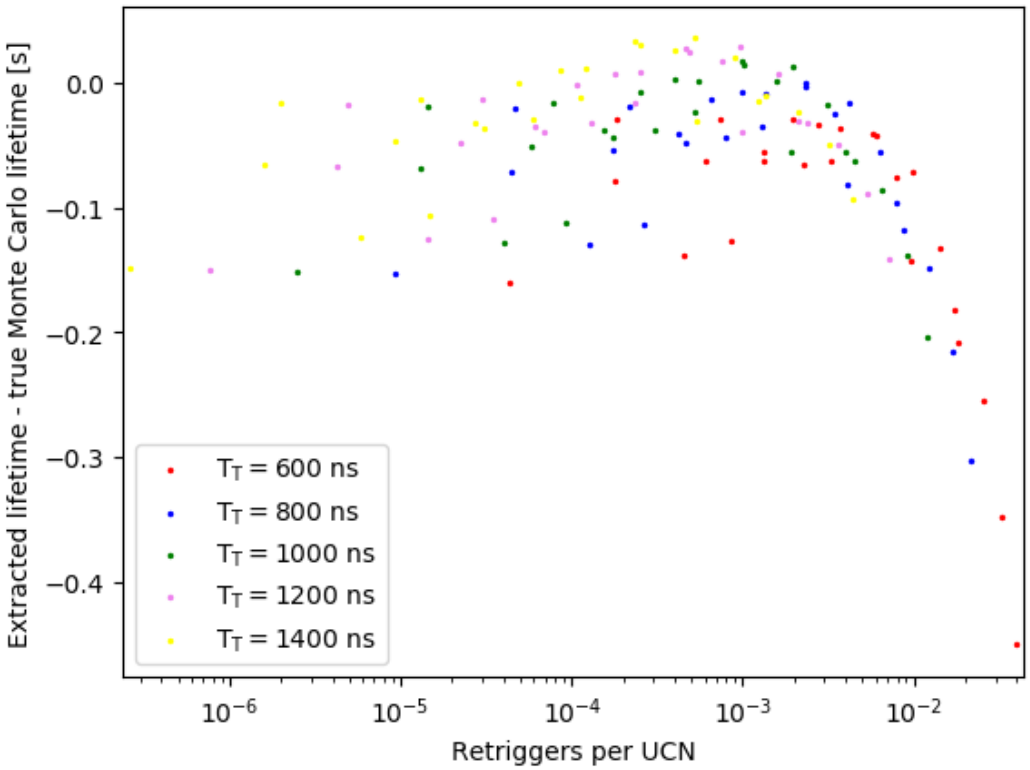

Figure 3.14: Extracted lifetimes for various choices of parameters of the reconstruction algorithm. The data had been separated by choice of tail window T_T, and plotted against the number of retriggers per UCN. For choices of T_P and N_{PE} that resulted in a relatively low number of retriggers per UCN, changes in T_T had no significant effect on the extracted lifetime. For choices of T_P and N_{PE} that resulted in a relatively high number of retriggers per UCN, changes in T_T had a significant effect on the extracted lifetime due to uncorrected retriggering.

difficult to verify. If either of the assumptions were not true, then there could have been some unknown bias. Minimizing UCN tail interactions minimized any potential unknown bias in the extracted lifetime. As seen in Figure 3.13, most reconstruction algorithms resulted in extracted lifetimes that did not significantly differ from the true Monte Carlo lifetime. Section 4.4 will discuss how the systematic uncertainty of the extracted lifetime due to the fact that choosing different seemingly valid reconstruction algorithms resulted in slightly different extracted lifetimes was quantified.

3.4.9 Conclusion

Section 3.4.2 developed models for PEs that were generated by one UCN captured by the primary detector and Section 3.4.3 used these models to simulate the capture of multiple UCN. UCN-to-UCN interactions caused difficulty in reconstructing the events in a way that would have biased an extracted lifetime, but Sections 3.4.4 through 3.4.7 developed corrections for these issues. These corrected CCs represented UCN that were counted in the trap, so the uncertainty of the values were assigned using counting statistics. These representations of UCN formed the basis for the rest of this analysis.

After forming the corrected CCs, the corrected CCs and the counts in all other detectors were collected into 0.1 s bins to reduce the computational scope of this analysis. This effectively reduced the precision with which each event was recorded from 800 ps to 0.1 s. This precision was much less than the precision of the lifetime that will be extracted in this chapter. Additionally, the binning of the data could have only affected the extracted lifetime if events were shifted into or out of the period of time during which the remaining UCN in the trap were counted. The delineations for the beginning and end of this period were $\sim$ 50 s before and after significant numbers of UCN were counted, so a potential shift of 0.1 s did not cause any UCN to be missed. Therefore, this binning did not significantly affect the extracted lifetime.

3.5 Backgrounds

3.5.1 Introduction

Some of the coincidence chains (CCs) that were constructed using the methods described in Section 3.4 were not caused by UCN, but were caused by various sources of backgrounds. After the holding process, long runs had far fewer UCN remaining in the trap than short runs. Therefore, the fraction of total CCs that were background CCs was greater in long runs than short runs. The extracted lifetime would have been biased by $\sim$ 20 s if an estimate of the number of background CCs were not subtracted from the measured number of CCs. Section 3.5.2 will investigate the nature of the background CCs and develop a method of rejection that will remove CCs that were characteristically very different from the models developed in Section 3.4.2. Section 3.5.3 will develop a model to estimate the background CC rates in the primary detector when it was at different locations within the trap. Section 3.5.4 will look for changes in the background rates of PEs and CCs that occurred within one run while the primary detector remained at a constant location. Section 3.5.5 will provide evidence that the estimates of the background rates of CCs were correct.

3.5.2 Classification of background events

The two photomultiplier tubes (PMTs) in the primary detector each had a background photoelectron (PE) rate of $\sim$ 100 Hz. Requiring coincident PEs in both of the PMTs to form a CC suppressed the background rate to $\simeq$ 0.4 Hz. If the background PE rates in the two PMTs in the primary detector were uncorrelated Poisson noise, and the only criterion to form a CC was coincident PEs within 100 ns, then the expected background CC rate would have been $(\sim 100\ \text{Hz})^2 \times 100\ \text{ns} \sim 1\text{mHz}$. In reality, the background PE rates in the two PMTs of the primary detector were not uncorrelated Poisson noise, and the criteria to form a CC were much more stringent than just finding coincident PEs within 100 ns, but this estimation alone shows that the background of CCs cannot be caused by just uncorrelated noise in the two PMTs of the primary detector. Instead, the majority of the background of CCs in the primary detector must have been caused by physical events in the detector. Figures 3.15 and 3.16 show distributions for CCs observed during and after the unloading process. The distributions show the number of PEs in each CC and the length of time between the first and last PEs in that CC. There were two classes of CCs. Class I made up a very small fraction of each of the distributions and was characterized by CCs that had an extremely high number of PEs, relative to the length of the CC. Specifically, CCs where the length of time between the first PE and the last PE was less than 20 ns per PE in the CC are Class I CCs. Class II consisted of all other CCs. Class I CCs were clearly very different from the UCN events modeled in Section 3.4.2. Class I CCs consisted of an extremely quick burst of PEs with no long tail of PEs. The distribution of PEs in Class I CCs did not favor either PMT. In contrast to the quickness of Class I CCs, UCN events had an extremely long tail of PEs.

Figure 3.17 shows the rates at which Class I CCs were detected during and after the unloading process. There was a baseline rate of Class I CCs that occurred even when there were no UCN remaining in the trap. A similar baseline rate was also present during the beam-off background runs, so the background of Class I CCs could not be explained by byproducts (either radioactive gasses or cold neutrons) of the production of UCN. Most Class I CCs had roughly the same number of PEs in each of the two PMTs, which ruled out the possibility that Class I CCs were caused by just one of the PMTs undergoing a high-voltage discharge. The physical cause of the baseline rate of Class I CCs was unknown.

Additionally, Figure 3.17 shows that there was an excess of Class I CCs around the time that the bulk of the UCN in the trap were counted by the primary detector. Monte Carlo simulations that used the models developed in Section 3.4.2 generated Class I CCs at a rate comparable to the increase in the Class I CC rate seen in Figure 3.17. The Class I CCs that were generated by simulated UCN events contained $\lesssim 15$ PEs. Figure 3.15 suggests that part of the distribution of CCs that behaved like the UCN events modeled in Section 3.4.2 were classified as Class I CCs. Therefore, it was likely that some Class I CCs were reconstructed from the PEs generated by a UCN that was captured by the primary detector.

The definition of Class I CCs captured a class of CCs that were clearly unlike the UCN events modeled in Section 3.4.2, but also captured some CCs which were almost certainly caused by UCN events in the primary detector. The fraction of CCs that behaved like UCN events, but were classified as Class I CCs, was $\lesssim 0.3\%$, so the exclusion of these CCs from an analysis had a negligible effect on the statistical uncertainty of an extracted lifetime. Section 3.5.5 will develop a method of verifying that the estimate of the background CC rate at Peak 3 was correct. When the Class I CCs were not excluded from this analysis, the method from Section 3.5.5 detected a $\simeq 0.1$ s bias in the extracted lifetime. The background rate of Class I CCs differed slightly between production runs and beam-off background runs. If the Class I CCs were not rejected from the analysis, then the estimates of the background CC rates would have been slightly biased, which in turn would have biased the extracted lifetime. The necessity to suppress this bias far outweighed the negligible decrease in the statistical uncertainty of an extracted lifetime due to rejecting 0.3% of the CCs generated by UCN captured by the primary detector. Therefore, all Class I CCs were rejected and were not considered in this analysis. This rejection was also applied to all models and simulations built in Section 3.4.

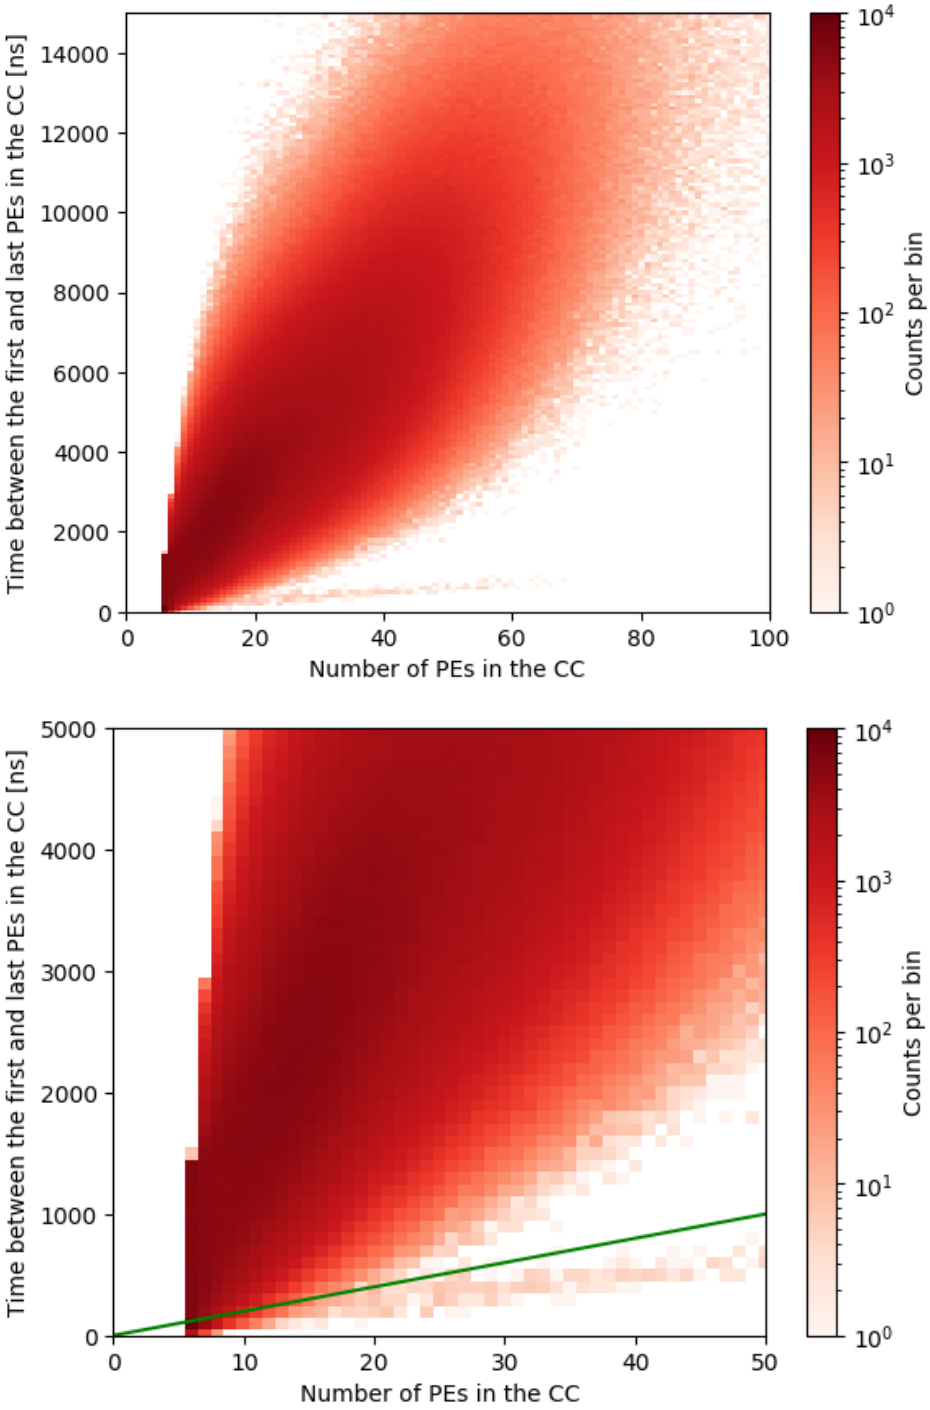

Figure 3.15: CCs in the primary detector during the unloading process. The green line in the lower plot represents a CC that consisted of PEs that arrived at a rate of 1 PE per 20 ns. CCs that fell below the green line were Class I CCs. The bottom and top plots share the same data and color scales, but the bottom plot shows only a subset of the data. *n.b.*: CCs in this figure were constructed using Algorithm 1 from Appendix B with T_c = 100 ns, T_T = 1500 ns, and N_{PE} = 6 PEs.

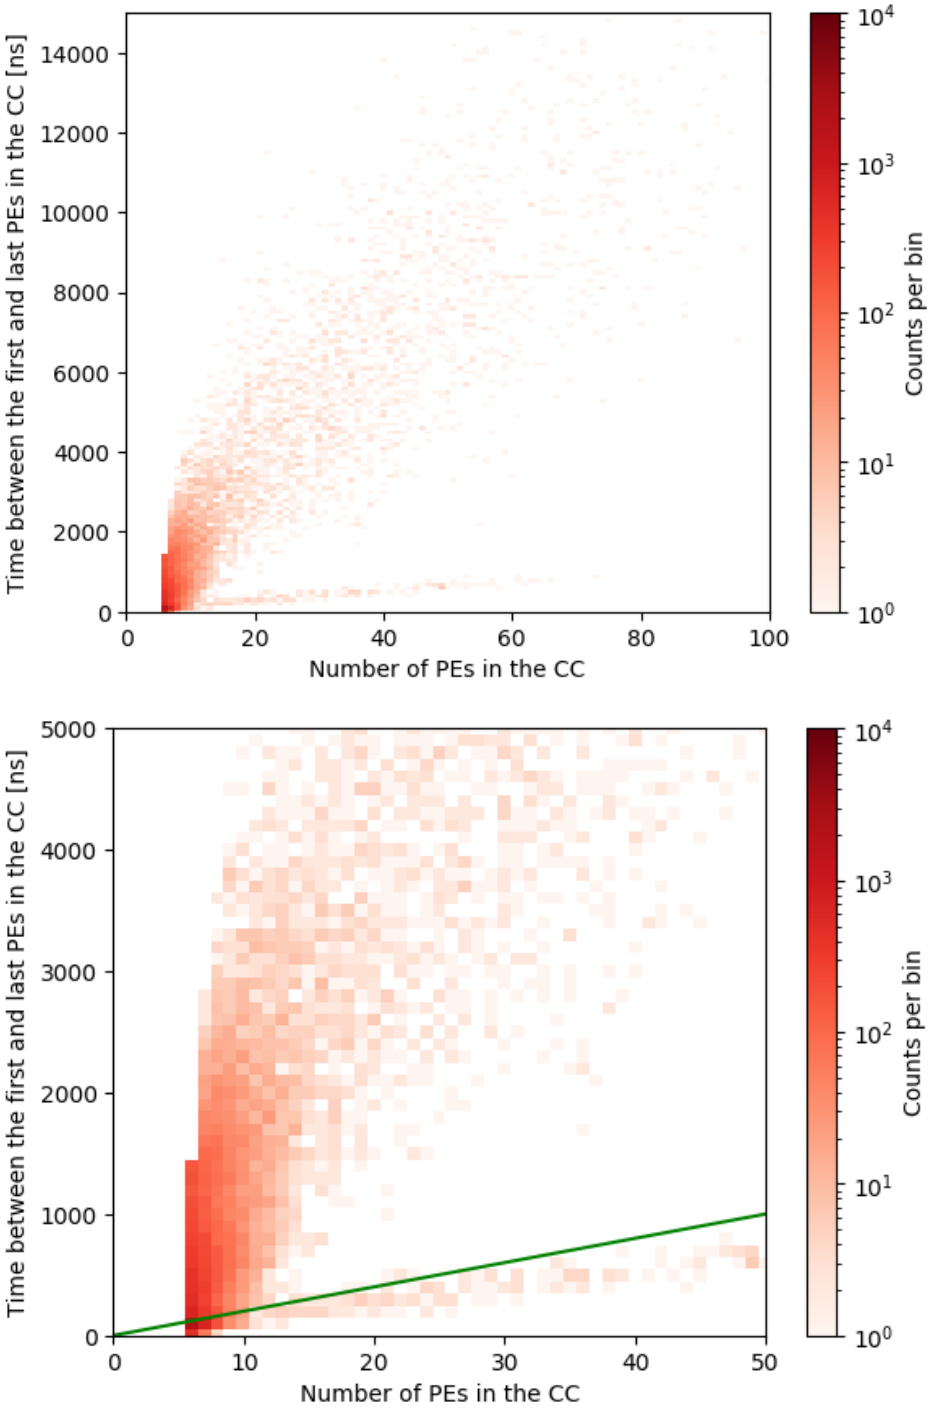

Figure 3.16: Background CCs in the primary detector after the unloading process. The green line in the lower plot represents a CC that consisted of PEs that arrived at a rate of 1 PE per 20 ns. CCs that fell below the green line were Class I CCs. The bottom and top plots share the same data and color scales, but the bottom plot shows only a subset of the data. *n.b.*: CCs in this figure were constructed using Algorithm 1 from Appendix B with $T_c = 100$ ns, $T_T = 1500$ ns, and $N_{\text{PE}} = 6$ PEs.

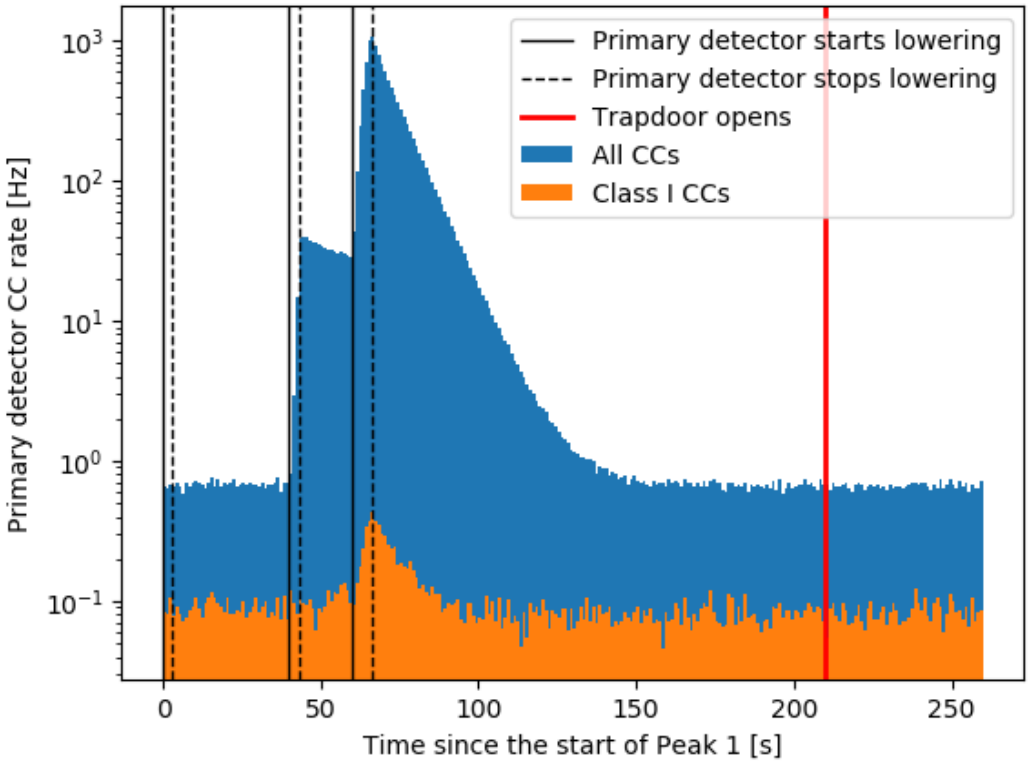

Figure 3.17: Time distributions for CCs recorded during and after the unloading process. When the primary detector lowered at 40 s and 60 s, the more UCN in the trap were able to reach the primary detector, which increased the rate at which UCN were captured by the primary detector. The CC rates then decreased as UCN were captured and the number of UCN remaining in the trap decreased. The lowering of the primary detector and the opening of the trapdoor did not cause the background CC rates to change. The baseline rate of Class I CCs was $\simeq$ 0.07 Hz. The rate of Class I CCs increased in tandem with an increase in the rate of all CCs. The maximum rate of all CCs was 1018 Hz above the baseline rate of all CCs. The maximum rate of Class I CCs was 0.33 Hz above the baseline rate of Class I CCs. The ratio of these increases was 3085 to 1. This ratio implied that the rate of Class I CCs should have increased by only 0.01 Hz when the primary detector is lowered at 40 s, so it was not surprising that there was not a clearly visible increase in the rate of Class I CCs at 40 s. *n.b.*: CCs in this figure were constructed using Algorithm 1 from Appendix B with T_c = 100 ns, T_T = 1500 ns, and N_{PE} = 6 PEs.

3.5.3 Height-dependent background rates

During the unloading process, the primary detector lowered from the top of the trap to the bottom of the trap. The background CC rate during the unloading process could not be measured because UCN were being captured by the primary detector. It was not possible to determine if a specific CC recorded during the unloading process was a CC caused by a UCN that was captured by the primary detector or a background CC. Some method of estimating the background CC rate was needed in order to estimate the total number of background CCs that were recorded during the unloading process. Two methods were considered:

1. If the background CC rate remained constant over $\sim$ weeks, then the ideal way to estimate the background CC rate during each unloading process would have been to sum together many background runs to make a high-statistics measurement of the background CC rate.

2. If the background CC rate remained constant only over $\sim$ minutes, then the ideal way to estimate the background CC rate during each unloading process would have been to measure the background CC rate at the end of each run.

Significantly more beam-off background runs were gathered during the 2017 and 2018 run cycles than beam-on background runs. (See Section 3.2 for a description of beam-off and beam-on background runs.) Therefore, the beam-off background runs were used to study the background CC rates. The mean background CC rate in the primary detector was slightly different at different heights, as shown in Table 3.4. Figure 3.18 shows the background CC rates measured in the primary detector during 68 beam-off background runs.

Height above the bottom of the trap	Background CC rate [Hz]
"Top" of trap: 45 cm	0.427 ± 0.005
Peak 1: 38 cm	0.443 ± 0.005
Peak 2: 25 cm	0.452 ± 0.005
Peak 3: 1 cm	0.449 ± 0.005

Table 3.4: Measured background CC rate at each of the four heights at which the primary detector spends a significant amount of time, averaged over 68 beam-off background runs between run numbers 4200 and 5440.

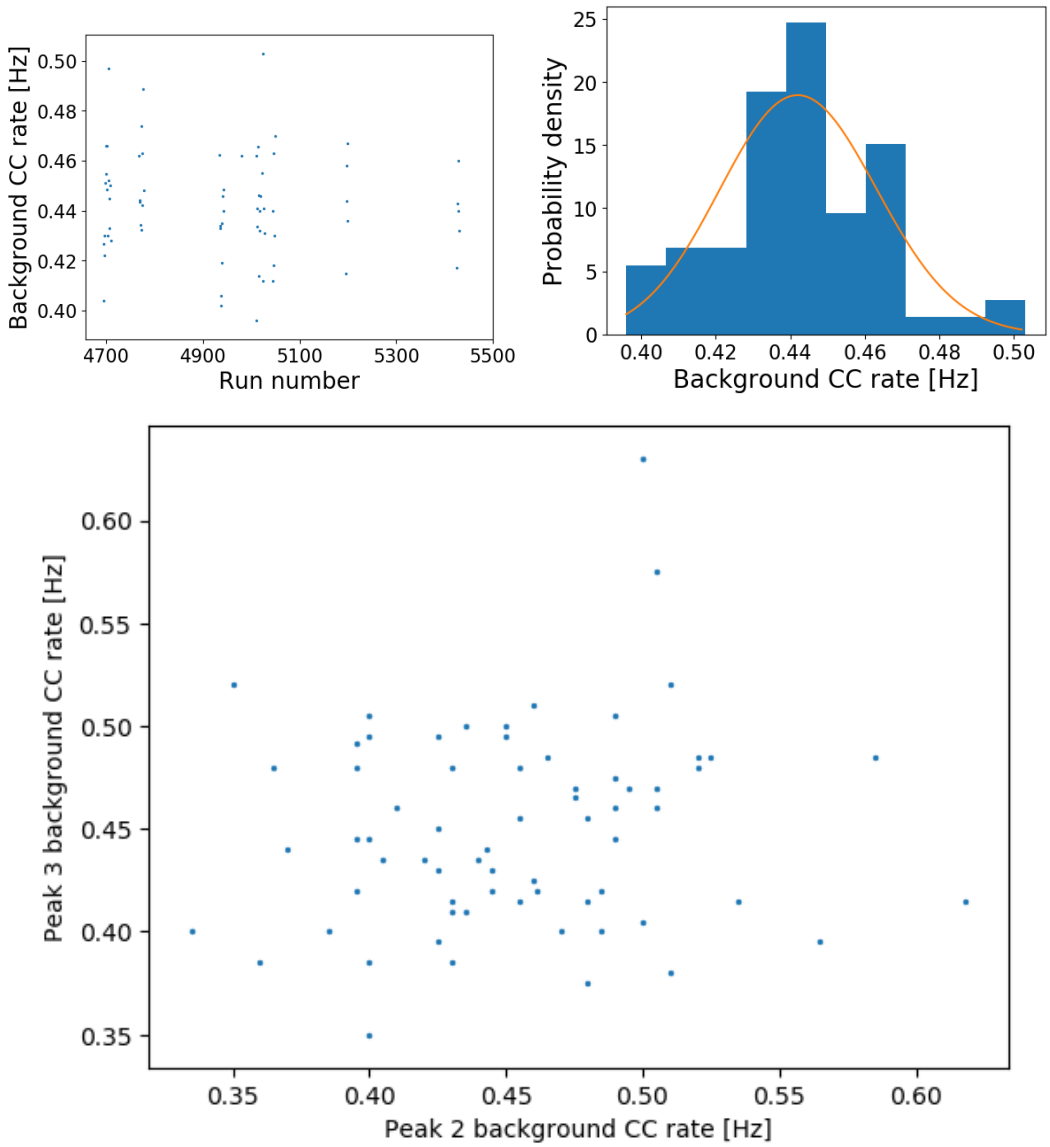

Figure 3.18: Background CC rates measured in the primary detector during the 68 beam-off background runs between run numbers 4200 and 5440. Top left: background CC rate averaged over all heights versus run number. Top right: the blue histogram shows the background CC rate averaged over all heights measured during each beam-off background run. The orange curve shows a Poisson distribution with a mean calculated from the mean number of background CCs recorded during one run. Bottom: Peak 3 background CC rate versus Peak 2 background CC rate. Each data point represents the background CC rates measured during one beam-off background run and has uncertainties along both axes of $\simeq 0.07$ Hz. A Kolmogorov–Smirnov (KS) test found no difference between the distribution of measured background CC rates and a Poisson distribution with mean background CC rate (top right, p-value = 0.80). No evidence was found of a linear correlation between the background CC rates measured at Peak 2 and Peak 3 (bottom, p-value = 0.31).

Figure 3.18 compares the background CC rates measured during beam-off background runs between run numbers 4200 to 5440 to the distribution of rates that would be expected if the number of background CCs recorded during each beam-off background run as a Poisson random variable (Poisson-distributed rates). No evidence of a difference was found between the measured background CC rates and Poisson-distributed rates. This comparison was repeated for each of the nine epochs of primary detector behavior. A meta-analysis of these nine comparisons found no evidence that the background CC rates differed from Poisson-distributed rates (p-value = 0.13). A similar meta-analysis was performed to search for a correlation between the background CC rates measured at Peak 2 and Peak 3, and no evidence that a correlation existed was found (p-value = 0.18).

If a linear relationship did exist between the background CC rate measured at Peak 2 (r_2) and the background CC rate measured at Peak 3 (r_3), then there would have existed some constants $m > 0, b$ such that $r_3 = mr_2 + b + \epsilon$, where the ϵs were residuals. This relationship could have been transformed by standardizing r_2 and r_3. In this context, standardization involved centering the mean of a data series to 0, and then scaling the standard deviation of a data series to 1. After $r_{2,3}$ were standardized to $r'_{2,3}$, the relationship between the two variables would have been $r'_3 = r'_2 + \epsilon'$. On the other hand, if no linear relationship existed between r_2 and r_3, then r'_3 should have been distributed around 0 without any relationship to r'_2.

Figure 3.18 also compares the background CC rates measured at Peaks 2 and 3. The likelihoods of observing these data were calculated under the assumptions that there was a linear relationship between the two background CC rates, and also that there was no relationship between the two background CC rates. The likelihood calculated under the assumption of no relationship was $> 10^{11}\times$ greater than the likelihood calculated under the assumption of a linear relationship. Similar differences in the likelihoods were found for all epochs of primary detector behavior.

These studies strongly suggested that background CC rate in the primary detector at a given height remained constant over long periods of time. Therefore, the background CC rates were estimated by summing together many beam-off background runs to make a high-statistics measurement of the background CC rates. The beam-off background runs were averaged during each epoch of primary detector behavior, and the background CC rate was estimated for each of the four heights at which the primary detector spent a significant amount of time. No increases in the background CC rates were observed while the primary detector was moving between positions.

3.5.4 Time-dependent background rates

The production of UCN began by using an 800 MeV proton beam to produce spallation neutrons from a W target. The W target was contained under a shield of steel concrete. Despite best attempts, some amount of radioactive gasses escaped from the shielding and spread through the experimental area. Evidence of this was seen in the background PEs in both PMTs of the primary detector. Data from the primary detector gathered during the holding processes of long runs was averaged to study this effect. During the holding processes of long runs, the primary detector was 7 cm above the cleaning height so it did not encounter any UCN. Figure 3.19 shows how the background PE rate measured in the PMTs of the primary detector changed over time. If not corrected for, this time-dependent PE background would have biased a lifetime extracted using raw PE counts in lieu of CCs because the background PE rate would have been different for short and long runs.

A similar average was constructed for the background of CCs and no time-dependent behavior was found. The time-dependent exponential behavior of the background PE rates had an amplitude of $\simeq 5.5$ Hz. As mentioned in Section 3.4.2, the average number of PEs generated when one UCN was captured by the primary detector was 46.6 PEs. If the time-dependent PE backgrounds were due to captured UCN, then there should have been a time-dependent background of CCs with an amplitude of 0.12 Hz over the baseline rate of background CCs of 0.39 Hz. This amplitude would have been apparent if it existed, but there was no evidence that any time-dependent background of CCs existed. Therefore, no estimate of a time-dependent background of CCs was used while extracting a lifetime.

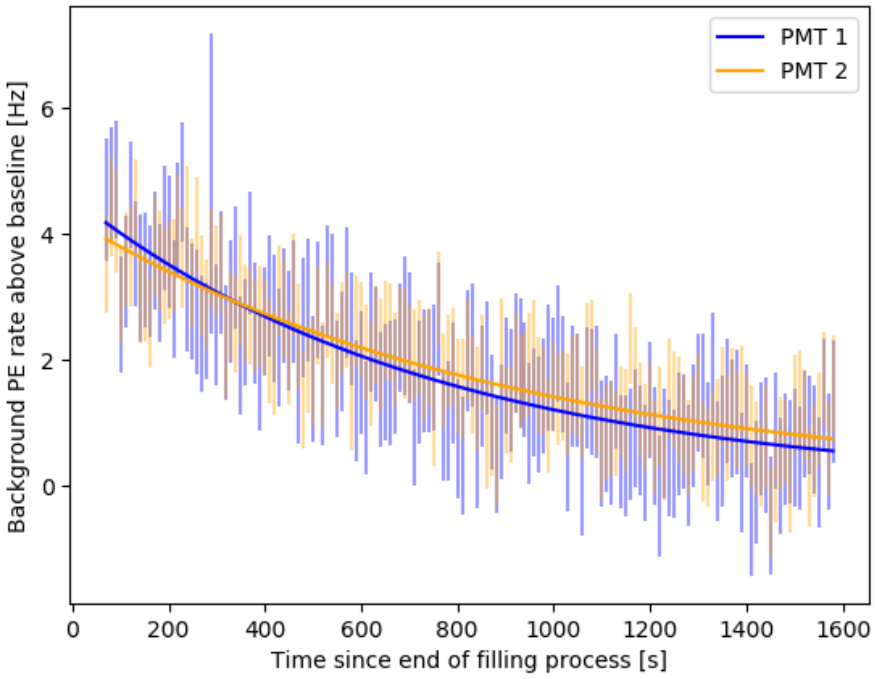

Figure 3.19: Mean background PE rates in both PMTs of the primary detector over the baseline background PE rates. These data were gathered during the holding process of 190 long runs between run numbers 4200 and 5441. The background PE rates above baseline were fit with single exponentials. The fit of PMT 1 had an amplitude of 5.8 ± 0.2 Hz and a time constant of 590 ± 70 s. The fit of PMT 2 had an amplitude of 5.2 ± 0.2 Hz and a time constant of 750 ± 100 s. The baseline background PE rates were ≃ 150 Hz for PMT 1 and ≃ 90 Hz for PMT 2.

3.5.5 Check of the estimation of the background CC rates

Figure 3.7 shows the distribution for time of detection of UCN on the primary detector during the unloading process. The unloading process continued for 170 s after the start of Peak 2, with 20 s spent at the Peak 2 height and 150 s spent at the Peak 3 height. Figure 3.7 shows that the UCN counting rate decreased to background levels around 100 s after the beginning of Peak 3. On average, < 1 UCN were counted during the last 50 s of Peak 3 of the unloading process. In contrast, ~ 30 background CCs were expected to be counted during the last 50 s of Peak 3.

If the estimation of background CC rates was incorrect, then the extracted lifetime should have drifted as the length of Peak 3 increased from 100 s to 150 s. Figure 3.20 shows that the extracted lifetime did not significantly change as the length of Peak 3 was extended from 100 s to 150 s. Even a small bias in the estimation of background CC rates can have caused a significant change to the extracted lifetime. A mis-estimate of 3% would have changed the extracted lifetime by 0.25 s.

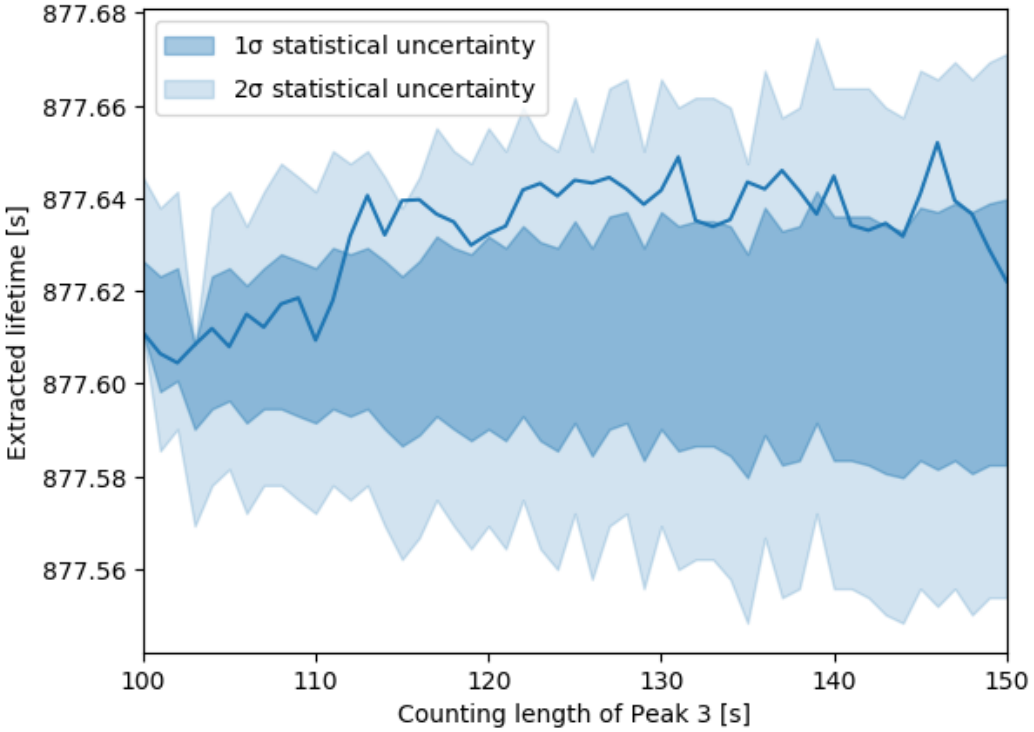

Figure 3.20: Extracted lifetime as a function of the counting length of Peak 3. In the last 50 s of Peak 3, almost all counts were background counts. If the background CC rates were mis-estimated, then the extracted lifetime would have linearly drifted as the counting time increased. The 1σ and 2σ uncertainty bands show the statistical variation in the extracted lifetime due to marginally increasing the number of UCN counted in Peak 3 as the length of Peak 3 increased. Almost all of the UCN counted in Peak 3 were counted within the first 100 s, so a small number of additional UCN counted in the last 50 s of Peak 3 could have only slightly changed the extracted lifetime.

3.5.6 Conclusion

The combination of the results from Sections 3.4 and 3.5 provided an estimate of the number of UCN that remained in the trap at the end of each run. Once the definition of a CC was finalized, it was straightforward to determine the number of CCs that occurred during an unloading process, but it was not straightforward to estimate the number of background CCs that occurred during the same unloading process. Section 3.5.2 presented a type of background event that was characteristically very different from the models of UCN events presented in Section 3.4.2 and rejected these events from the analysis. Section 3.5.3 showed that the background CC rate in the primary detector varied as the height of the primary detector changed. Section 3.5.4 showed that an excess of background PEs in both PMTs was present at the beginning of holding processes but exponentially decayed. Building CCs from the PEs completely suppressed this time-varying excess background PE rate. Section 3.5.5 demonstrated that continuing to count CCs after almost all of the UCN had already been removed from the trap did not significantly affect the extracted lifetime because the estimates of the background CC rates were sufficiently precise. The number of UCN that remained in the trap at the end of each run had been well measured. To complete the extraction of a lifetime, an estimate for the number of UCN that were in the trap at the beginning of the run was needed. This estimate will be developed in Sections 3.6 and 3.7.

3.6 Normalization

3.6.1 Introduction

The primary detector described in Section 2.7.2 was used to count the number of UCN in the trap, but in doing so, the UCN were removed from the trap. This counting was performed at the end of each run to determine the number of UCN that did not decay during the holding process, but obviously could not have been used to count the number of UCN in the trap before the holding process. Instead, the number of UCN in the trap before the holding process was estimated using other data sources. Section 3.6.2 will discuss how the normalization monitors from Section 2.7.3 were used to gain some information about the UCN that were produced by the source. Section 3.6.3 will develop a model that uses this information to predict the number of UCN that were loaded into the trap. Section 3.6.4 will list reasons why the same set of parameters of the model could not have been used for the entire 2017/2018 data set and how many different sets of parameters were needed.

3.6.2 Normalization monitors

For each run, UCN rates from two normalization monitors were used to estimate the number of UCN loaded into the trap. The monitors used in the normalization estimate changed as hardware and monitors were added or removed from the experiment, but in general there was one monitor that was at the same height as the UCN guide (the *lower* monitor) and one monitor that was at an elevated height (the *upper* monitor). Both monitors provided a good measurement of the overall number of UCN produced by the source. The cleaning process tried to remove all UCN with energy $\geq$ 38 neV from the trap, but the source produced and the guides transported UCN with energy up to $\sim$ 213 neV, as discussed in Section 2.4. As discussed in Section 2.3, when the UCN source aged the energy spectrum of UCN that it produced hardened. The relative number of UCN measured in the upper and lower monitors contained information about what fraction of the UCN produced by the source could have been be stored in the trap.

During the 2017 run cycle, the lower monitor was the gate valve monitor and the upper monitor was the standpipe monitor. The number of counts recorded in the gate valve monitor during the filling process was $\sim 3.3\times$ the number of UCN loaded into the trap. The number of counts recorded in the standpipe monitor was $\sim 21\times$ the number of UCN loaded into the trap. (Section 2.7.3 explained why the upper monitors counted significantly more UCN than the lower monitors.) During most of the 2018 run cycle, the lower monitor was the buffer volume density monitor and the upper monitor was the buffer volume pre-cleaner. The number of counts recorded in the buffer volume density monitor during the filling process was $\sim 6.8\times$ the number of UCN loaded into the trap. The number of counts recorded in the buffer volume pre-cleaner during the filling process was $\sim 59\times$ the number of UCN loaded into the trap.

As discussed in Section 2.7.3, the normalization monitors sometimes recorded more than one count per captured UCN. Even if the the monitors recorded exactly one count per captured UCN, an unbiased uncertainty could not be directly assigned to the number of UCN captured by the monitor. Under Poisson statistics, if a monitor that was expected to observe M counts actually observed M' counts, then the uncertainty of the measured value M' is $\sqrt{M}$ and not $\sqrt{M'}$. Therefore, a model of the expected number of UCN counts recorded in each normalization monitor would be needed in order to assign an unbiased uncertainty of the measured counts in each monitor. Developing such a model would have required precise simulations of UCN

diffusion from the source through the guides. Instead, Section 3.6.3 will circumvent this problem and develop a method of assigning uncertainties to the estimates of the number of UCN loaded into the trap without considering the uncertainties of the number of counts recorded in the normalization monitors.

Following the same logic as in Section 3.3.2, monitor counts later in the filling process were given more weight than monitor counts early in the filling process. Let $l(t)$ and $u(t)$ represent the background-subtracted rates in the upper and lower monitors during the filling process. Then define

$$M_L = \sum_t w(t) l(t)$$

and

$$M_U = \sum_t w(t) u(t),$$

where $w(t)$ is Equation 3.2 and t is summed over the entire filling process. M_L and M_U were weighted integrals of the UCN rates recorded in the lower and upper monitors during the filling process. The effect of choice of the weighting function on the extracted lifetime will be explored in Section 4.5.

The background rate in the monitors was measured at the end of each long run. The end of short runs were not used to measure the monitor background rates because there were still a detectable number of UCN in the guides. The estimate of the monitor background rate at the beginning of each run was made by using the nearest measurement of the monitor background rate. Sometimes the estimate of the monitor background rate at the beginning of a long run was different than the measured monitor background rate at the end of that long run. Consider the example where two long runs were done consecutively. The estimate of the monitor background rate in the second run was the measured background rate at the end of the first run because the time of measurement and the time at which the rate was estimated only differ by $\simeq$ 1 minute. If the measured monitor rate from the end of each long run was used to estimate the monitor at the beginning of the same long run, then there would have been $\sim$ 30 minutes of time between the measurement of the monitor rate and when the monitor rate was estimated.

Figure 3.21 shows how the background rate in the buffer volume pre-cleaner changed over time. These changes in background rate were driven by changes in the temperature of the photomultiplier tube (PMT) that observed the buffer volume pre-cleaner. Additionally, the background rates in various normalization monitors occasionally changed due to changes in the detector settings.

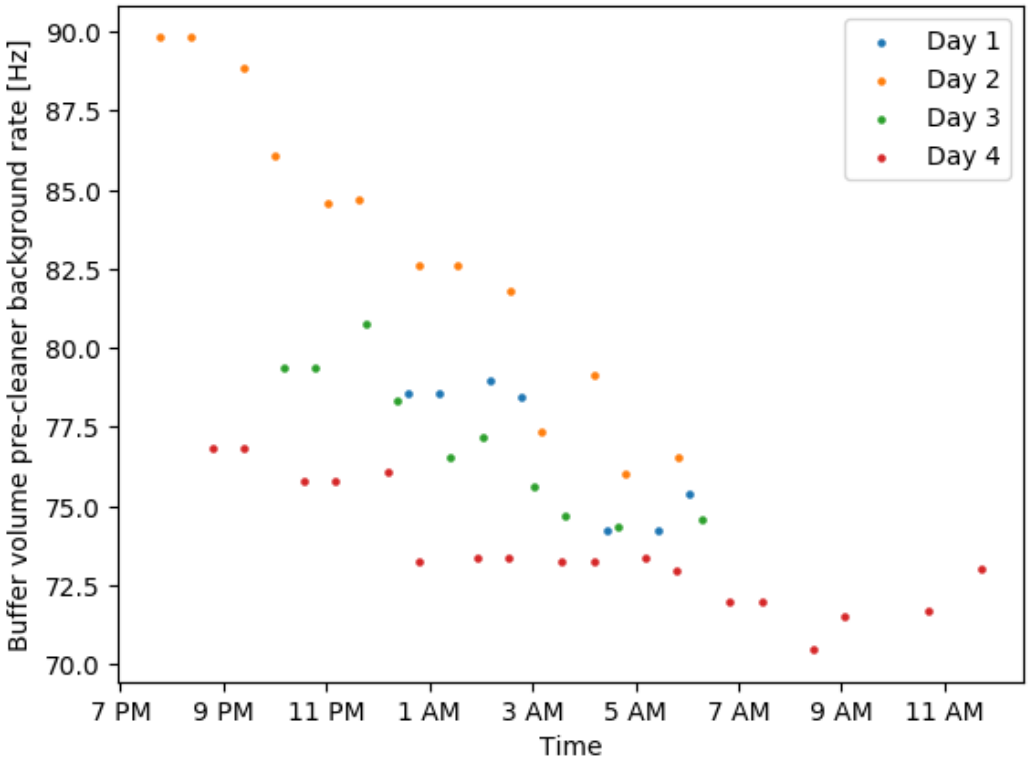

Figure 3.21: Changes of the background rate in the buffer volume pre-cleaner as a function of time of day. The experimental area did not have sufficient climate control. As the temperature of the experimental area decreased overnight, the temperature of the PMT that observed the buffer volume pre-cleaner also decreased, which reduced the background rate in the buffer volume pre-cleaner. During the filling process of these runs, the foreground rate in the pre-cleaner was around 2500 Hz. An unaccounted for 20 Hz shift in the background rate would have led to a 0.8% bias in the estimate of the number of foreground counts.

3.6.3 Normalization model

As discussed in Section 2.3, the quality of the sD_2 UCN source deteriorated as the integrated number of protons on the W target increased. This deterioration continued until the sD_2 was reconditioned. The deterioration of the sD_2 UCN source caused it to become less efficient at producing UCN and hardened the energy spectrum of the UCN that were produced. The relationship between the overall number of UCN produced and the hardness of the energy spectrum of those UCN was non-linear. Additionally, the relationship between the energy spectrum of the UCN that were produced and the number of UCN captured by the normalization monitors was non-trivial to model and would have required Monte Carlo simulations to understand.

As shown in Figure 3.22, the effect of the change in the energy spectrum of the UCN produced by the source on the number of UCN captured by the normalization monitors was small and the relationship between M_U and M_L was very close to linear. M_U and M_L captured mostly the same information (within a normalization subset the average correlation between the two was $r = 0.96$). Within each normalization subset, Principal Component Analysis (PCA) [51] was used to produce two new data series P_1 and P_2. PCA changed the basis to (P_1, P_2), which contained the same information as (M_U, M_L) but P_1 and P_2 were uncorrelated. PCA was used on each normalization subset as follows:

1. M_U and M_L were standardized
2. PCA was used to transform from (M_U, M_L) to (P_1, P_2)
3. P_1 and P_2 were standardized

The standardization process in Steps 1 and 3 used a linear transformation to shift the mean of each data series to 0 and the standard deviation of each data series to 1. The standardization in Step 1 was necessary to ensure that PCA captured only the variation in M_U and M_L and not the non-zero means. The standardization in Step 3 resulted in P_1 and P_2 having the same scale, which was useful because it meant that coefficients in a soon-to-be-introduced model will represent the scale of the effects captured by the two principal components. P_1 captured about 98% of the variance in (M_U, M_L) because M_U and M_L were so heavily correlated, and the rest of the variance was captured in P_2. PCA did not provide an explanation for what P_1 and P_2 physically represent, but informally P_1 can be thought of as a measure of the overall number of UCN produced by the source, and P_2 can be thought of as a measure of the energy spectrum of those UCN. Including a measure of the energy spectrum of the UCN in the normalization model corrected the normalization estimate for the changes in the UCN energy spectrum.

The normalization model was defined as

$$N = \alpha + \beta_1 P_1 + \beta_2 P_2, \tag{3.5}$$

where N is the estimate of the number of UCN loaded into the trap. A method to determine α, β_1, and β_2 will be developed in Section 3.7. It was important to include some measure of the UCN energy spectrum in the normalization estimate because the trap can only store UCN from the lower part of the energy spectrum.

Consider two runs where the same number of UCN were produced but the UCN energy spectrum was harder in the second run. Fewer of the produced UCN were of the appropriate energy to be stored in the trap, but if a measurement of the UCN energy spectrum was not included in the normalization estimate, then the two runs would have the same estimate for the number of UCN loaded into the trap.

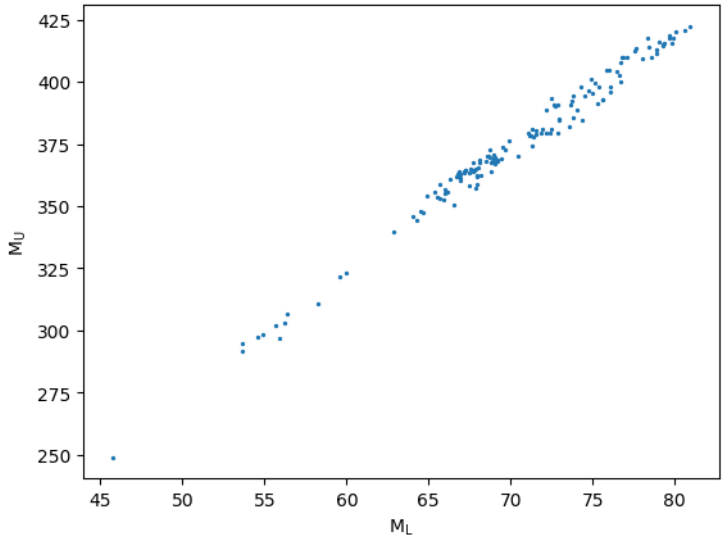

Figure 3.22: The relationship between the weighted sum of upper (M_U) and lower (M_L) monitor counts during the filling process. The number of counts in both detectors was primarily determined by the number of UCN produced by the source. The wide range of the data demonstrates the wide range of the rate of production of UCN as the source aged.

The effect of not including a measure of the UCN energy spectrum in the normalization estimate can be seen by comparing a lifetime extraction using Equation 3.5 to a lifetime extraction using

$$N' = \alpha' + \beta' P_1. \tag{3.6}$$

This simpler model still uses most of the information contained in (M_U, M_L), but does not correct for changes in the UCN energy spectrum. Comparison of lifetimes (and uncertainties) extracted from the same data set using the normalization models defined in Equations 3.5 and 3.6 provided a measure of how changes in the UCN spectrum affected the normalization estimate. Using the normalization model from Equation 3.6 instead of the model from Equation 3.5 increased the extracted lifetime by 0.01 s, but significantly increased the statistical uncertainty of the extracted lifetime from 0.27 s to 0.29 s. Equation 3.5 will be used for all normalization estimates, unless otherwise noted.

Equations 3.5 and 3.6 estimated the number of UCN loaded into the trap, but did not provide uncertainties of these estimates. Section 3.6.3 discussed the challenges in estimating the uncertainty of the number of counts recorded in the normalization monitors. If known, these uncertainties could have been used to calculate the uncertainty of the number of UCN which were loaded into the trap, which is the uncertainty that is actually used while extracting a lifetime. In ideal running conditions, the actual number of UCN which were loaded into the trap would have been a random variable $n \sim \text{Poisson}(N)$. For sufficiently large N, $\text{Poisson}(N) \simeq \text{Gaussian}(N, \sqrt{N})$. Instability in the production of UCN during the filling process induced additional uncertainty of the number of UCN loaded into the trap. To avoid issues with assigning uncertainties to the number of counts recorded in the normalization monitors, the uncertainty of the estimates of the number of UCN loaded into the trap was modeled as

$$\sigma_N^2 = fN, \tag{3.7}$$

and the actual number of UCN loaded into the trap was modeled as a random variable

$$n \sim \text{Integer}\left[\text{Gaussian}(N, \sqrt{fN})\right]. \tag{3.8}$$

f was the scale of the uncertainty of the normalization estimate σ_N relative to the ideal case of $\sigma_N = \sqrt{N}$. In the ideal case of a normalization estimate with Poisson uncertainties $f = 1$, but the uncertainties of the normalization estimates in excess of Poisson uncertainties corresponded to a value of $f > 1$. Different values for f were allowed for the data from 2017 and the data from 2018 due to the installation of the buffer volume and the associated change in running procedure, but to prevent overfitting only one value for f was allowed for each year's data.

Section 3.7.7 will construct an equation for a χ^2 that compares the actual and expected numbers of coincidence chains (CCs) detected during the unloading process of each run. The number of expected CCs was a function of the estimated number of UCN loaded into the trap. f was a free parameter in the equation for this χ^2, and f was determined by demanding that the χ^2 was equal to the number of degrees of freedom of the fit used to determine the normalization parameters. Section 3.7.9 will present the exact detail of how N was estimated for each run, will justify why it was appropriate to allow different values for f for the 2017 data and the 2018 data, and will develop a method to calculate f.

Equation 3.7 models the variance of the normalization estimate as a linear scaling of the normalization estimate. If there were non-linear terms in the model for the uncertainty of the normalization estimate, then this would have been evident by comparing the performance of the model on runs with a low normalization estimate to runs with a high normalization estimate. For example, if the uncertainty of the normalization estimate was truly modeled as $\sigma_N^2 = fN + gN^2$ (with $f, g > 0$ and $fN + gN^2 \geq N$), but the model from Equation 3.7 was used in the analysis, then the uncertainty of the normalization estimates of runs with a low normalization estimate would have been underestimated, and *vice versa*. This would have been evident when the distributions for residuals between the observed data and the estimated normalization and lifetime were compared. Figure 3.23 shows the distributions for the residuals for different levels of normalization estimate. The lack of a difference suggests that the model from Equation 3.7 is correct.

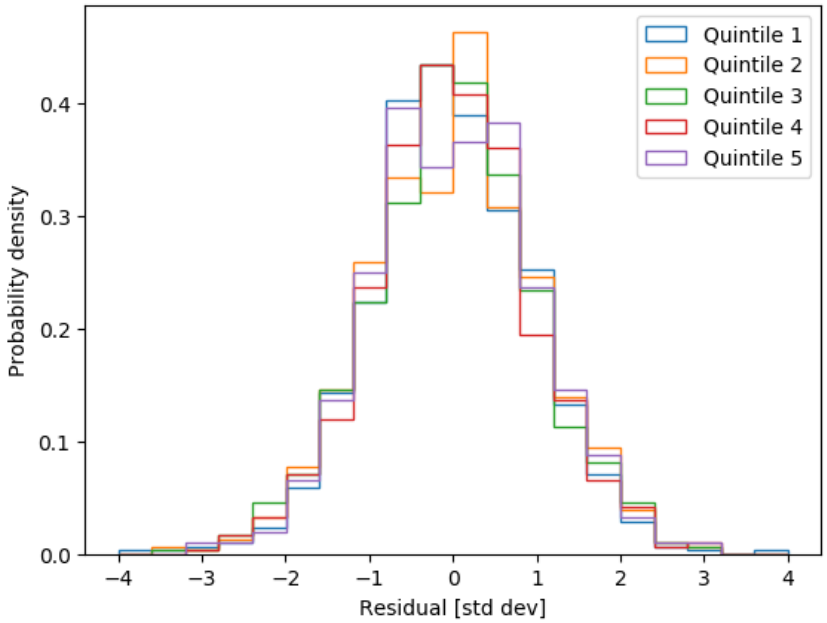

Figure 3.23: Distributions for the residuals between the measured and expected number of CCs from the unloading process at the end of each run. The runs were divided into five quintiles based on the normalization estimate. Quintile 1 was the 20% of the runs with the lowest normalization estimates and Quintile 5 was the 20% of the runs with the highest normalization estimates. A meta-analysis of Kolmogorov–Smirnov tests found no significant difference between the distribution for the residuals for runs with different levels of normalization estimates (p-value = 0.62). A difference between the distributions for different quintiles would have suggested that Equation 3.7 did not accurately model the relationship between the normalization estimate and the uncertainty of that estimate.

3.6.4 Changes in running conditions

Consecutive runs usually had very similar conditions and should have followed the same normalization model. However, there were many cases where there were significant changes between runs and two runs that followed each other had different normalization models. Some examples of these changes included:

1. The sD_2 source was reconditioned;
2. Long breaks were taken between runs;
3. Hardware or monitors were added or removed from the experiment, or detector settings were changed; and
4. Beam conditions changed.

185 significant changes were identified in the 2017/2018 data set, resulting in 186 distinct normalization subsets. Sometimes significant changes occurred rapidly and few runs were performed in a normalization subset. Any subset with less than eight short runs was excluded from the analysis. After making this exclusion, and considering that there were many subsets where no attempt was made to gather production data, there were a total of 101 different normalization subsets that were analyzed to extract a lifetime.

3.6.5 Conclusion

This section developed a way to estimate the number of UCN that were loaded into the trap at the beginning of each run. Sections 3.6.2 and 3.6.3 developed a model that used the counts in the normalization monitors to estimate the number of UCN that were loaded into the trap. Section 3.6.3 also developed a model for the uncertainty of the estimate of the number of UCN that were loaded into the trap. Section 3.6.4 explained why multiple sets of parameters of the model were needed for different sections of data that had different filling characteristics. This section did not discuss how to find the optimal values for the parameters of these models. A way of quantifying the likelihood of the observed data was needed in order to find the optimal values. Section 3.7 will develop a likelihood model to meet this need, and will use that likelihood model to develop a process to find the optimal values for the parameters of the normalization model and to calculate the uncertainty of the estimate of the number of UCN that were loaded into the trap.

3.7 Likelihood model

3.7.1 Introduction

Different runs from across the 2017/2018 data set had significantly different numbers of UCN that were loaded into the trap, different background CC rates, and variation in other quantities of interest. This section will develop a likelihood model in order to quantify how each run should contribute to the extraction of a lifetime. Each run will be modeled as a five-step process. Section 3.7.2 will model the uncertainty of the estimate of the number of UCN that were loaded into the trap. Section 3.7.3 will develop a model for the stochastic behavior of the β-decay of UCN in the trap during the holding process. Section 3.7.4 will extend the model from Section 3.7.3 to include statistical fluctuations due to detection inefficiency of the primary detector. Section 3.7.5 will further extend this model to account for how the software-dead-time and UCN-event-tail-interaction corrections from Sections 3.4.5 and 3.4.6 affected the number of coincidence chains (CCs). Section 3.7.6 will extend this model even further to account for variations in the actual number of background CCs, compared to the estimate for the number of background CCs. Section 3.7.7 will combine the modelled behavior of the runs to construct a likelihood equation for the entire data set, and Section 3.7.8 will use Monte Carlo simulations to test the validity of some of the assumptions used to build these models. Section 3.7.9 will combine the values extracted in Sections 3.4 through 3.6 and the models developed in this section to find the optimal values for the parameters of the normalization model and will assign an uncertainty to the estimate of the number of UCN that were loaded into the trap. Sections 3.7.2 through 3.7.7 will use the shorthand $p \equiv e^{-t/\tau}$.

3.7.2 Step 1: Loading UCN into the trap

Section 3.6.3 developed a model for estimating the number of UCN that were loaded into the trap at the beginning of each run. Equation 3.8 models the process of loading UCN into the trap as

$$n \sim \text{Integer}\left[\text{Gaussian}(N, \sqrt{fN})\right],$$

where $f \geq 1$. This model was first introduced in Section 3.6.3.

3.7.3 Step 2: Holding UCN in the trap

If there were exactly n UCN loaded into the trap, and the holding length was t, then the number of UCN that did not decay during the holding process was a random variable

$$d \sim \text{Binomial}\,(n, p)\,. \tag{3.9}$$

However, n was not exactly known, but instead was estimated as N with variance fN. In the case where $f = 1$ and $n \sim \text{Poisson}(N)$ (instead of making a Gaussian approximation), it can be analytically shown that $d \sim \text{Poisson}(Np)$, as shown in Appendix C.

A similar analytical proof could be constructed when the approximation in Equation 3.8 was used. Instead, Monte Carlo simulations were used in Section 3.7.8 to show that for sufficiently large Np and $N(1-p)$, d is approximately distributed as

$$d \sim \text{Integer}\left[\text{Gaussian}\left(Np, \sqrt{fNp^2 + Np(1-p)}\right)\right]. \tag{3.10}$$

3.7.4 Step 3: Detection efficiency of UCN in the trap

The primary detector did not have 100% efficiency for counting UCN. The simulated UCN events from Section 3.4.3 were used to estimate the efficiency of reconstruction of events in the primary detector. For reasonable choices of parameters of the reconstruction algorithm, the primary detector was $s \sim 90\%$ efficient for detecting UCN (when not damaged). Given a known, exact number of simulated UCN events, the variation in the number of CCs was close to Binomial. Most UCN events resulted in zero or one CCs, but some resulted in more than one due to retriggering. Most UCN events had sufficient time separation from other UCN events that it was reasonable to treat them as independent, but there were some UCN-to-UCN interactions. The variation in the number of CCs per UCN was estimated using the simulated UCN events from Section 3.4.3.

After assuming that Np was sufficiently large, the number of CCs due to UCN capture on the primary detector was approximately distributed as

$$c \sim \text{Integer}\left[\text{Gaussian}\left(Nps, \sqrt{fNp^2s^2 + Np(1-p)s^2 + Np\sigma_s^2}\right)\right]. \tag{3.11}$$

3.7.5 Step 4: Software-dead-time and UCN-event-tail interactions

The first three steps in this process have so far ignored the fact that the number of CCs needed to be corrected using the adjustments presented in Sections 3.4.5 and 3.4.6. Both the software-dead-time and UCN-event-tail-interaction processes followed Poisson statistics. These corrections did not change the expectation value of the corrected number of CCs, but did increase the variance on said quantity. If the software-dead-time correction increased the number of CCs by $\delta_d > 0$, and the UCN-event-tail correction decreased the number of CCs by $\delta_t > 0$, then the total effect on the number of CCs was a shift of $C = \delta_d - \delta_t$, with a standard deviation of $\sigma_C = \sqrt{\delta_d + \delta_t}$. Therefore, the corrected number of CCs due to UCN capture on the primary detector was approximately distributed as

$$c' \sim \text{Gaussian}\left(Nps, \sqrt{fNp^2s^2 + Np(1-p)s^2 + Np\sigma_s^2 + \sigma_C^2}\right). \tag{3.12}$$

3.7.6 Step 5: Background

The measured unload counts, U, were the sum of the surviving UCN that were counted c and the background counts b. The background counts were assumed to be distributed $b \sim \text{Poisson}(B)$. Then

$$U \sim \text{Gaussian}\left(Nps + B, \sqrt{fNp^2s^2 + Np(1-p)s^2 + Np\sigma_s^2 + \sigma_C^2 + B}\right). \tag{3.13}$$

3.7.7 Constructing a likelihood equation

A likelihood equation could now be constructed:

$$\mathcal{L} = -\ln L = -\sum_i \ln \text{Gaussian}\left(U_i | \mu_i, \sigma_i\right), \tag{3.14}$$

where i ran over all runs in the analysis, and

$$\mu_i = N_i p_i s_i + B_i$$

and

$$\sigma_i = \sqrt{fN_i p_i^2 s_i^2 + N_i p_i (1-p_i) s_i^2 + Np_i \sigma_{s_i}^2 + \sigma_{C_i}^2 + B_i}.$$

Equation 3.14 was adapted to construct

$$\chi^2 = \sum_i \frac{(U_i - B_i - N_i p_i s_i)^2}{fN_i p_i^2 s_i^2 + N_i p_i (1-p_i) s_i^2 + N_i p_i \sigma_{s_i}^2 + \sigma_{C_i}^2 + B_i}. \tag{3.15}$$

Section 3.7.9 will present the details of how Equation 3.15 was used to calculate f.

3.7.8 Checking the assumptions using Monte Carlo simulations

Equation 3.13 approximated the entirety of the five-step process that was modeled in Sections 3.7.2 through 3.7.6 as one Gaussian random variable. Some, but not all, of the individual steps were individually modeled as Gaussian random variables (up to rounding the random variable to the nearest integer). The model for the stochastic behavior of the β-decay of UCN in the trap during the holding process was Binomial instead of Gaussian, but Equation 3.13 approximated it as a Gaussian random process. If the difference between the exact and approximate distributions was sufficiently large, then it could have caused a bias in the extracted lifetime. Figure 3.24 compares the Gaussian distribution in Equation 3.13 to Monte Carlo data generated by sequentially drawing random values following the processes laid out in Sections 3.7.2 through 3.7.6. There was a small, but visible, difference between the model and the Monte Carlo data. The effect of this bias will be discussed in Section 3.9.3.

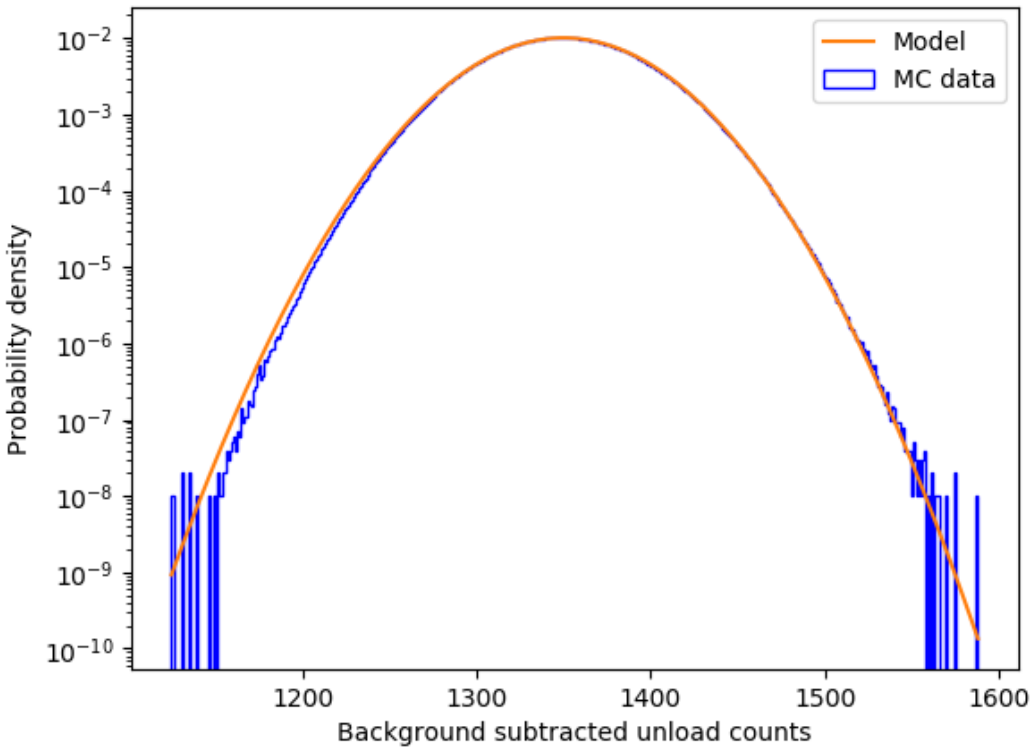

Figure 3.24: Monte Carlo simulation of the processes from Sections 3.7.2 through 3.7.6 with $N = 10,000$, $f = 2$, $e^{-t/\tau} = 0.15$, $s = 0.9$, $\sigma_s = \sqrt{s(1-s)} = 0.3$, $\sigma_C = 0$, and $B = 30$. The histogram has 100 million entries. Section 3.9.3 will show that the difference between the model and the simulated Monte Carlo data did not bias the extracted lifetime.

3.7.9 Normalization optimization

The number of UCN that were loaded into the trap was not directly measured, but the number of UCN that were unloaded from short runs were directly measured and had high sensitivity to the number of UCN loaded into the trap. The number of UCN unloaded from long and extra-long runs had low sensitivity to the number of UCN loaded into the trap, and were highly correlated with the lifetime, so only short runs were used in the normalization estimation process. There were four short holding lengths (20 s, 50 s, 100 s, and 200 s) so a value for the lifetime was needed in order to combine data from runs of different holding lengths. Due to the high correlation between the number of UCN unloaded from long and extra-long runs and the lifetime, including the long and extra-long runs in the normalization estimation process would have increased the statistical uncertainty of the extracted lifetime from 0.27 s to $\simeq$ 0.30 s.

101 different normalization subsets were used in this analysis. Each subset had three free normalization parameters. After including the lifetime, a full simultaneous optimization of the free parameters over all the data required searching over a 304 free parameters. This was numerically expensive, especially because the normalization parameters from different subsets were correlated because they shared a common lifetime. If the lifetime was fixed, then the 101 normalization subsets decoupled from one another, and an optimization over 304 parameters reduced to 101 optimizations, each of which was a search over three parameters.

Recall that Equation 3.7 had a parameter f which was the scale of the uncertainty of the normalization estimate σ_N relative to the ideal case of $\sigma_N = \sqrt{N}$. Even in the ideal case where $f = 1$, some normalization subsets would have more (or less) variation than Poisson statistics due to random fluctuations. If f was determined independently for each normalization subset, then the uncertainty of the normalization estimate would have always exactly matched the observed variation in the data, which would have clearly been overfitting. Instead, f was determined over larger subsets of the data. The major change in the normalization process was the installation of the buffer volume between the 2017 and 2018 run cycles, so one value for f was allowed for each of the two run cycles.

A value for the lifetime was needed to find the maximum likelihood estimate (MLE) of the normalization parameters. The normalization parameters were needed to make the normalization estimates, which were needed to extract the lifetime. These two optimizations were done iteratively until both converged. The methods used to extract the lifetime will be explained in Sections 3.9 and 3.10. Given a fixed trial value for the lifetime τ^*, the following approach was used to solve for the MLE of the normalization parameters for all normalization subsets:

1. Initial values for α, β_1, β_2 (or α', β' if not using the spectral correction) were picked for all normalization subsets.

2. Using those initial values, two values of f (one for 2017 before the buffer volume was installed, and one for 2018 after the buffer volume was installed) were solved for such that Equation 3.15 was equal to the number of degrees of freedom of the normalization estimate. The number of degrees of freedom was the number of short runs from each year, less three (or two, if not applying a spectral correction to the normalization estimate) per normalization subset. If $f < 1$, then set $f = 1$. The optimal values for f are reported in 3.5.

3. The trial values for the normalization parameters, τ^*, and f were used to calculate the likelihood from Equation 3.14 while summing over only the short runs.

4. The normalization parameters from Step 1 were varied, and Steps 1 through 3 were repeated until normalization parameters were found that maximized the negative log likelihood from Equation 3.14 computed in Step 3.

Buffer volume used	f when a spectral correction was used	f when a spectral correction was not used
No	2.55	4.63
Yes	1.92	2.35

Table 3.5: The scale of the uncertainty of the normalization estimate σ_N relative to the ideal case of $\sigma_N = \sqrt{N}$. The buffer volume smoothed out fluctuations in the filling process and reduced the variations in the filling process. The spectral correction also reduced variations in the filling process by correcting for changes in the UCN energy distribution.

Within each normalization subset,

- α was a measure of the average number of UCN loaded into the trap;
- β_1 was a measure of the range of how many UCN were loaded into the trap; and
- β_2 was a measure of how changes in the UCN energy spectrum affected the number of UCN loaded into the trap.

Over the various normalization subsets from the 2017-2018 data set, the average number of UCN loaded into the trap varied between 9,400 and 38,600 UCN. The number of UCN loaded into the trap varied by up to 11,200 from the average of that subset, but the mean displacement from the average was only 1,100. Changes in the UCN energy spectrum caused the number of UCN loaded into the trap to vary from the average of that subset by up to 1,800, but the mean displacement from the average was only 110.

Monte Carlo simulations were used to demonstrate that values of $f > 1$ were needed for the likelihood model to match the characteristics of the data observed during the 2017 and 2018 run cycles. Two sets of Monte Carlo data were simulated by following the five-step process from Sections 3.6.3 through 3.7.6. The first set was generated using the following process to match the characteristics of the data observed during the 2017 and 2018 run cycles:

1. Each run from the analysis had an estimate N for the number of UCN loaded into the trap, and an uncertainty of that estimate $\sigma_N = \sqrt{fN}$. A Monte Carlo value for the number of UCN loaded into the trap was drawn as $n \sim \text{Integer}[\text{Gaussian}(N, \sqrt{fN})]$.
2. Each run from the analysis had a holding length t. A Monte Carlo value for the number of UCN that remained in the trap after the holding process was drawn as $d \sim \text{Binomial}(n, e^{-t/\tau})$. τ was chosen to match the value extracted in Section 3.9.
3. Each run had an estimate for the efficiency of the primary detector s, and associated uncertainty σ_s. A Monte Carlo value for the number of CCs recorded in the primary detector was drawn as $c \sim \text{Integer}[\text{Gaussian}(sd, \sigma_s\sqrt{d})]$.

4. Each run had estimates for the number of CCs that were missed due to software dead time ($\delta_d > 0$) and the number of additional CCs that were erroneously counted due to UCN-event-tail interactions ($\delta_t > 0$). δ_d was the difference between the number of CCs when both the software-dead-time and UCN-event-tail corrections were used and the number of CCs when only the UCN-event-tail correction was used. δ_t was the difference between the number of CCs when only the UCN-event-tail correction was used and the number of CCs when no corrections were used. A Monte Carlo value for the corrected number of CCs recorded in the primary detector was drawn as $c' \sim c + [\delta_d - \text{Poisson}(\delta_d)] - [\delta_t - \text{Poisson}(\delta_t)]$.

5. Each run had an estimate for the number of background CCs that are recorded during the unloading process. A Monte Carlo value for the background-subtracted corrected number of CCs recorded in the primary detector was drawn as $u \sim c' + \text{Poisson}(B) - B$.

6. $Y = u/N$ was calculated for each run.

7. Steps 1 through 7 were repeated 100,000 times for each run from the analysis to produce a high-statistics Monte Carlo data set that matched the characteristics of the real data set.

The second set was generated in the same way as the first, except with $f = 1$. Figure 3.25 shows the values for Y for these simulations and compared them to the values for Y for the runs in the 2017/2018 data set.

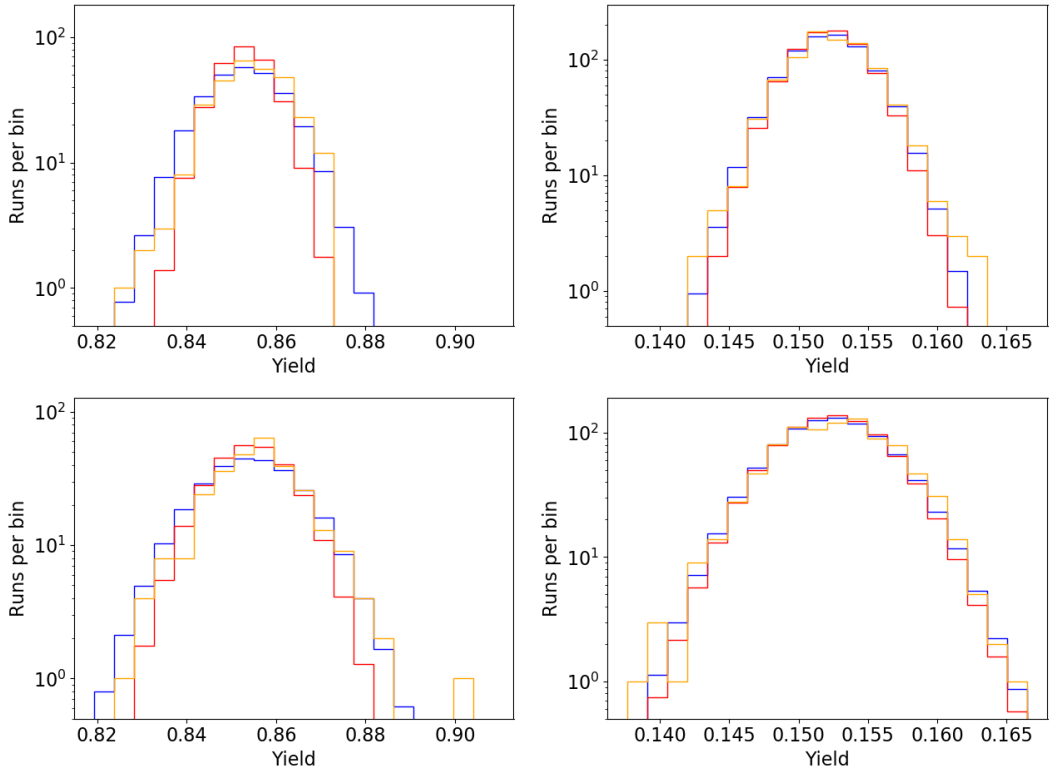

Figure 3.25: Observed and simulated yields. Top left: 20 s runs without the buffer volume installed. Top right: 1550 s runs without the buffer volume installed. Bottom left: 20 s runs with the buffer volume installed. Bottom right: 1550 s runs with the buffer volume installed. In each plot the orange histogram is the actual data from either the 2017 or 2018 run cycle, the blue histogram is Monte Carlo data simulated with values for f that match the values from Table 3.5, and the red histogram is Monte Carlo data simulated with with $f = 1$. In the two plots of 20 s runs (left), it is clear that the simulated data with values for f that match the values from Table 3.5 more closely matched the data, and that the simulated data with $f = 1$ did not match the data. In 20 s runs, the variation in the yields was composed of roughly equal contributions from the variations in the number of UCN loaded into the trap and variations in the number of UCN that undergo β-decay before being captured by the primary detector during the unloading process. In the two plots of 1550 s runs (right), the two different sets of simulated data are not clearly different from one another, nor are they different from the actual data. In 1550 s runs, the variation in the yields was almost entirely composed of the variation in the number of UCN that undergo β-decay before being captured by the primary detector during the unloading process.

3.7.10 Conclusion

Sections 3.7.2 through 3.7.8 built a likelihood model in order to quantify how each run should contribute to the extraction of a lifetime. This model broke each run into a five-step process that began with UCN being loaded into the trap and ended with the surviving UCN being counted from the trap. Section 3.7.9 used this likelihood model to extract the optimal values of the parameters of the normalization models from Chapter 3.6. These parameters were used to estimate the number of UCN in the trap at the beginning of each run, as well the uncertainty of those estimates. Sections 3.9 and 3.10 will use the estimates for the number of UCN in the trap at the beginning and end of runs (and associated uncertainties) to extract lifetimes.

3.8 Holding length

3.8.1 Fixed offsets in holding length

Recall the overview of a production run from from Section 3.2.1. For the sake of discussion in this section, assume that the cleaning process was perfect. That is, the cleaning removed all UCN with sufficient energy to have reached the cleaning height, and that the UCN with insufficient energy to have reached the cleaners were not affected by the cleaning process. From the perspective of a UCN with insufficient energy to have reached the cleaners, time spent in the trap during the cleaning process was no different than time spent in the trap during the holding process.

Starting from Equation 3.14 and making the simplifying assumptions that $s = 1$ and $B = \sigma_s = \sigma_C = 0$,

$$\begin{aligned}\mathcal{L}\left(\tau, N_i | u_i, t_i\right) &= -\sum_i \ln \text{Gaussian}\left(u_i \middle| N_i e^{-t_i/\tau}, \sqrt{f N_i e^{-2t_i/\tau} + N_i e^{-t_i/\tau}\left(1 - e^{-t_i/\tau}\right)}\right) \\ &= -\sum_i \ln \text{Gaussian}\left(u_i \middle| N_i e^{-t_i/\tau}, \sqrt{N_i e^{-t_i/\tau}\left[\left(f-1\right) e^{-t_i/\tau} + 1\right]}\right),\end{aligned} \tag{3.16}$$

where i sums over all runs in the analysis. In the best-case scenario of $f = 1$, σ_i^2 reduced to $N_i e^{-i_i/\tau}$. This was equal to the mean of the Gaussian likelihood function, which must be equal to the expected number of UCN counted during the unloading process. A fixed offset in the holding length (for example, including or excluding the 50 s from the cleaning process in the holding length) would have changed $e^{-t_i/\tau}$. However, $N_i e^{-t_i/\tau}$ must have remained constant because it was the expectation value of a physical quantity that was invariant to definitional changes in t_i.

When $f > 1$, the total variance was dependent on fixed offsets in holding length. $N_i e^{-t_i/\tau}$ remained invariant to fixed offsets in the holding length, but the variance included an additional term and was proportional to $(f-1)e^{-t_i/\tau}+1$. A Monte Carlo simulation was done to determine if including the 50 s of cleaning length in the holding length affected the extracted lifetime as follows:

1. 3000 simulated weighted normalization monitor counts were assigned as $M_i = 14000+2i$ for $i \in \{1, 2, \ldots, 3000\}$. Only one normalization monitor was used in this simulation, so Equation 3.6 was used as the normalization model.

2. For each M_i, a Monte Carlo value was drawn as
$$n_i \sim \text{Integer}\left[\text{Gaussian}(M_i, \sqrt{fM_i})\right].$$

3. Each run was assigned a holding length t_i by repeating the standard octet pattern: 20 s, 1550 s, 1550 s, 50 s, 100 s, 1550 s, 1550 s, and 200 s.

4. A known value for τ_{MC} was used to draw a Monte Carlo value for $u_i \sim \text{Binomial}(n_i, e^{-(t_i+50)/\tau_{\text{MC}}})$, where the additional 50 s simulated the length of the cleaning process.

5. The methods described in Sections 3.9 and 3.10 were used to extract a lifetime while either including or excluding the cleaning length in the holding length.

6. Steps 1 through 5 were repeated a total of 250 times to make a high-statistics measurement of the bias in the lifetime that was induced by not including the cleaning length in the holding length.

7. Steps 1 through 6 were repeated for different trial values of f.

Figure 3.26 shows the results of the Monte Carlo simulation for different values of f. As f increased, the bias in the global lifetime calculated when the cleaning length was not included in the holding length also increased. For large values of f, this bias was comparable to the statistical uncertainty of the extracted lifetime. This bias was due to a systematic mis-estimation of the normalization estimate. When the cleaning length was appropriately included in the values for the holding lengths, then there was no bias in the global lifetime. The paired lifetime method also systematically mis-estimated the normalization when the cleaning length was not included in the holding length, but the paired lifetime method significantly suppressed this bias through means that will be discussed in Section 3.10.

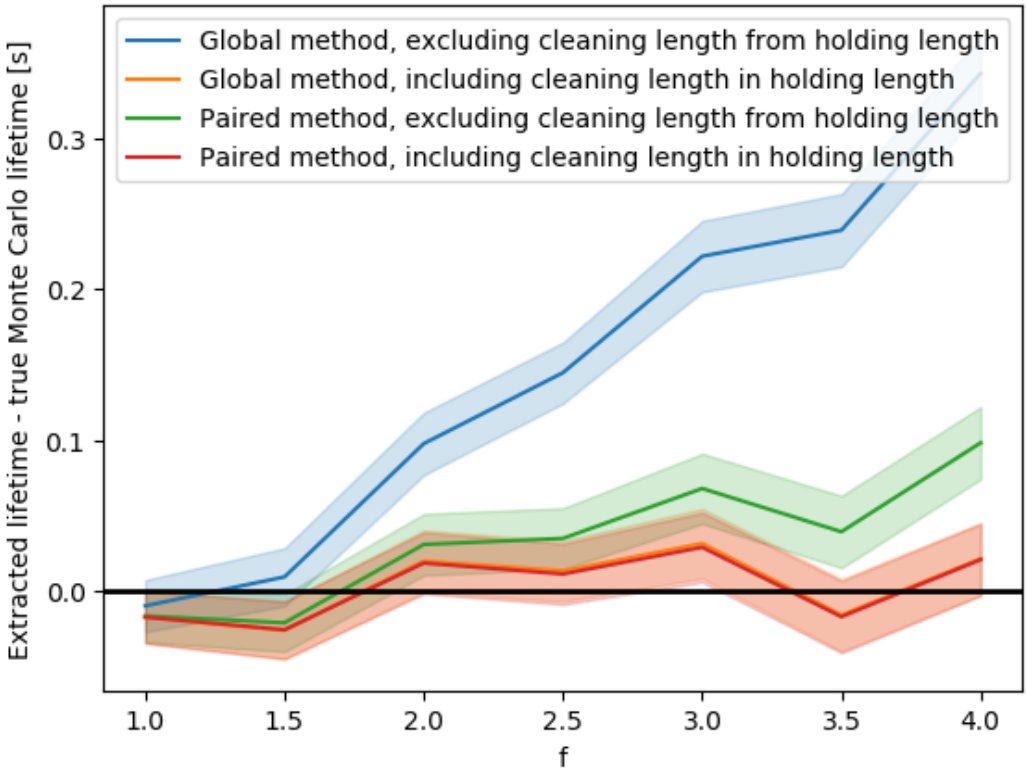

Figure 3.26: Bias in the extracted lifetime as a function of the scale of the uncertainty of the normalization estimate σ_N relative to the ideal case of $\sigma_N = \sqrt{N}$. When the cleaning length was included in the holding length, the global and paired methods extracted almost exactly identical lifetimes. As such, the orange curve lies almost directly beneath the red curve.

This Monte Carlo simulation demonstrated that the cleaning length should have been included in the holding length used to extract a lifetime. For the same reason, some amount of time should have been added to the holding length to account for the fact that there was some delay between the end of the holding process and when the UCN were captured by the primary detector. Section 3.8.2 will show that the average length of time from the end of the holding process to the captured of UCN was roughly 70 s. This length of time was about 1.4 s shorter in 2018 than 2017 because the UCN capture efficiency per interaction with the primary detector was increased between the two years. Section 3.8.2 will explore how run-to-run variations in the average amount of time between the end of the holding length and the detection of UCN affected the extracted lifetime.

3.8.2 Time evolution of UCN trajectories within the trap

The time evolution of the trajectories of UCN in the trap has been studied with Monte Carlo simulations [43]. This time evolution caused the distributions for arrival times of UCN to differ slightly between short and long runs, as shown in Figure 3.27 from Section 3.4.2. This difference in distributions for arrival times caused the fraction of UCN that underwent β-decay during the unloading process (but before being captured by the primary detector) to have differed between short and long runs, which could have biased the extracted lifetime.

Figure 3.7 demonstrated that UCN were counted at different times during the unloading process. UCN that were counted at later times had a greater chance of having undergone β-decay before being counted. Therefore, the different UCN in the trap had different probabilities of surviving long enough to have been counted by the primary detector during the unloading process. The probability that each individual UCN did not undergo β-decay before it was counted at the end of the run was modeled as a Binomial process with a single trial. However, the different UCN had different probabilities of not decaying, so the total number of UCN that did not undergo β-decay before being counted was not a Binomial random variable.

To better understand the consequences of each UCN having a different probability of undergoing β-decay before being counted, consider the example in which there were M UCN in the trap at the end of the run. If not for β-decay, each of those M UCN would have been captured by the primary detector at time t_i, where $t = 0$ was the start of the unloading process. The expected number of UCN that would have been captured by the primary detector could have been found by summing the probabilities that each individual UCN would have not undergone β-decay before being captured by the primary detector. This expected value was

$$\begin{aligned}
\langle M' \rangle &= \sum_{i=1}^{M} e^{-t_i/\tau} \\
&= \sum_{i=1}^{M} \sum_{j=0}^{\infty} \frac{(-t_i/\tau)^j}{j!} \\
&= M \sum_{j=0}^{\infty} \frac{\left\langle (-t_i/\tau)^j \right\rangle}{j!} \\
&\simeq M \sum_{j=0}^{\infty} \frac{\langle -t_i/\tau \rangle^j}{j!} \\
&= M e^{-\langle t_i \rangle/\tau}.
\end{aligned}$$

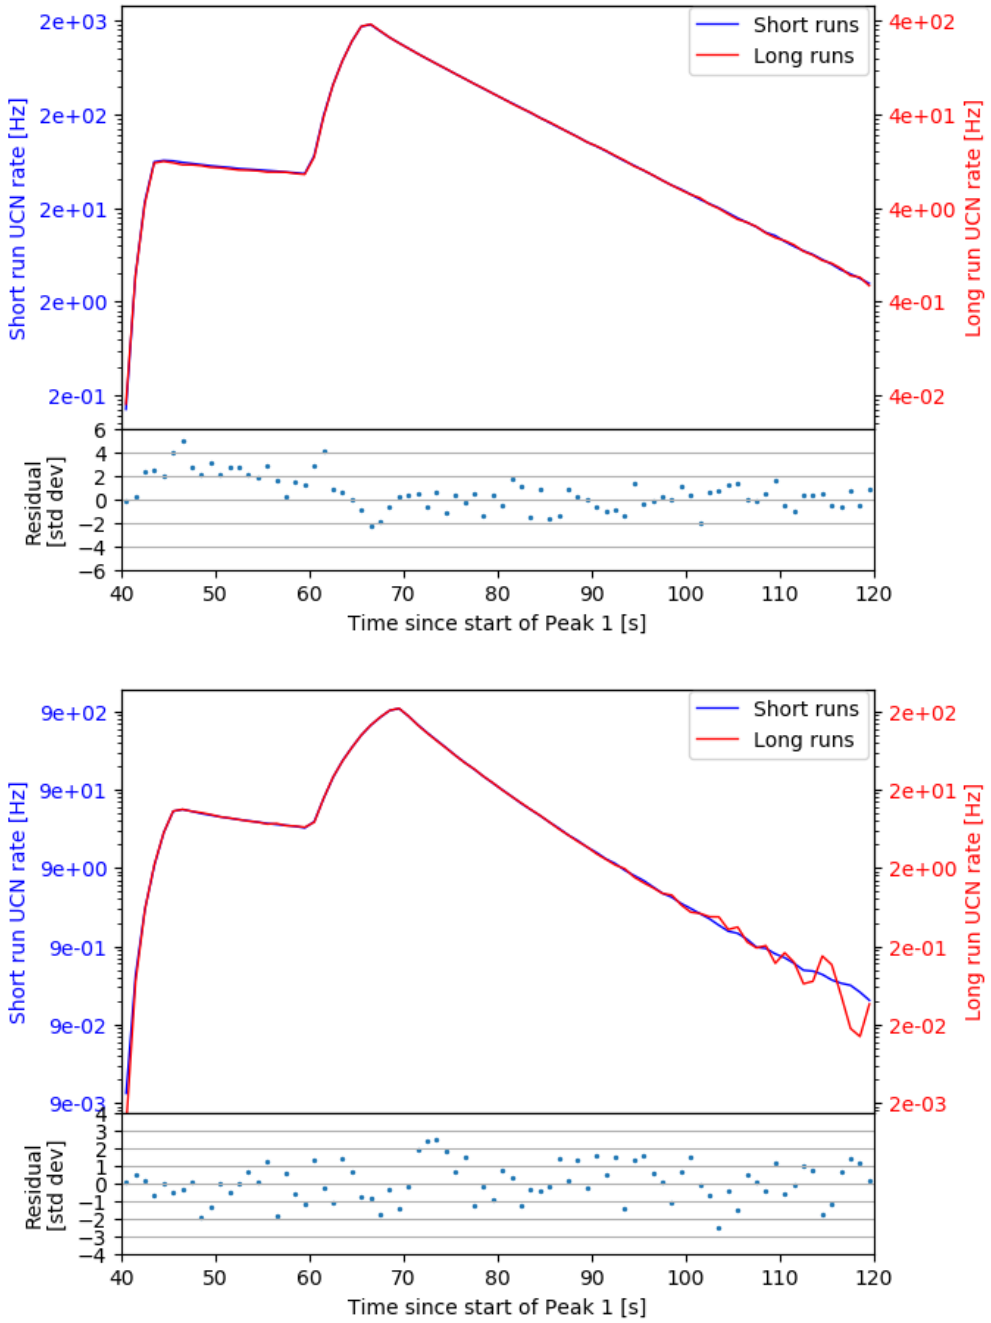

Figure 3.27: Distributions for arrival times for short and long runs. The top figure shows the distributions from 2017 and the bottom figure shows the distributions from 2018. As will be discussed in Section 4.2, there was an aluminum block accidentally inserted into the trap for part of the 2017 data. Figure 4.1 shows that this block affected the UCN counted during Peak 2 much more than the UCN counted during Peak 3, which matches the effect seen in the differences between the short and long arrival time distributions during 2017.

Monte Carlo simulations that used the distributions from Figure 3.7 found that

$$\frac{e^{-\langle t_i \rangle/\tau} - \left\langle e^{-t_i/\tau} \right\rangle}{\left\langle e^{-t_i/\tau} \right\rangle} \simeq -5.5 \times 10^{-5}.$$

This section will soon show that the change in the extracted lifetime due to considering the time evolution of UCN trajectories within the trap was small, so this difference of -5.5×10^{-5} was negligible.

To attempt to correct for any potential biases in the extracted lifetime caused by the time evolution of UCN trajectories within the trap, the mean arrival time $\bar{t}$ of all UCN captured by the primary detector during the unloading process was added to the holding length. The coincidence chains (CCs) were collected into 0.1 s bins, and then the estimated background CC rates were subtracted from each bin. For an unloading process where each bin had background-subtracted counts c_i and bin center b_i, the mean arrival time was defined as

$$\bar{t} \equiv \frac{\sum_i c_i b_i}{\sum_i c_i}. \tag{3.17}$$

Figure 3.28 shows the mean arrival times for each run in the global analysis. Most of the UCN counted during the unload process were counted during Peak 3, where there was no evidence of a difference between the distributions for short and long arrival times. The relatively large number of UCN in Peak 3 outweighed the difference in the Peak 2 component of the distributions for arrival times. Therefore, the difference in distributions was not apparent in the mean arrival times, as shown in Table 3.6. However, Figure 3.27 clearly shows that there was a difference between the distributions for arrival times between short and long runs.

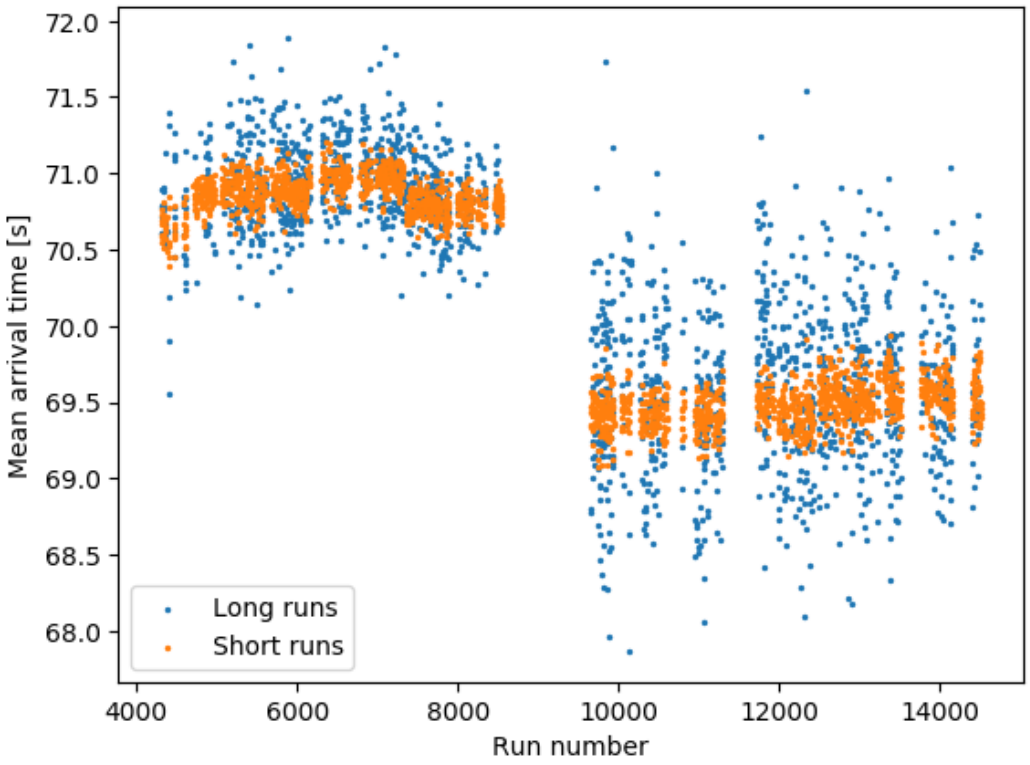

Figure 3.28: Mean arrival time for each run in the global analysis, measured relative to the beginning of the unloading process. The clear change in behavior around run 9000 was due to changing the primary detector between the 2017 and 2018 run cycles. The primary detector used during 2018 had a higher UCN capture efficiency per interaction than the primary detector used during 2017. The capture efficiency per interaction in 2018 was $\lesssim 2\times$ the efficiency from 2017. The exact ratio of efficiencies depended on the perpendicular energy of each UCN. The improved primary detector in 2018 counted the UCN in the trap more quickly. The faster counting rate can be seen by comparing the top and bottom plots in Figure 3.27. Short runs had less variation in mean arrival time than long runs because the short runs had a higher number of UCN in the trap, so statistical fluctuations were smaller and the signal-to-noise ratio was larger.

Year	Short runs [s]	Long runs [s]
2017	70.91 ± 0.04	70.94 ± 0.06
2018	69.50 ± 0.04	69.50 ± 0.07

Table 3.6: Mean arrival time of UCN on the primary detector, split between short and long runs and split by years. The mean arrival times from 2018 were less than those from 2017 because the primary detector was upgraded between 2017 and 2018, which increased the capture efficiency per interaction between a UCN and the primary detector.

The differences in the distributions for arrival times were more significant in 2017 than in 2018, so the distributions from 2017 were used in a Monte Carlo simulation to estimate bias on the extracted lifetime. Monte Carlo data were generated as follows:

1. Exactly 100 million arrival times t were drawn from both the distributions for short and long arrival times. For the purposes of this Monte Carlo simulation, the distributions were truncated to only include times between 40 s and 120 s after the beginning of Peak 1.

2. 50 s of cleaning length were added to each arrival time. 20 s were added to the arrival times of short runs and 1550 s were added to the arrival times of long runs.

3. For each arrival time, a random value was drawn as $r \sim \text{Uniform}(0, 1)$, and that random value was compared to $p = e^{-t/\tau_{\text{MC}}}$. If $r < p$, then the simulated UCN did not decay before being counted. If $r \geq p$, then the simulated UCN decayed and was not counted.

4. For both the short and long runs, the mean arrival time of surviving simulated UCN was calculated and the total number of surviving UCN were counted.

5. Equations E.1 and E.2 were used to extract lifetimes while either including or excluding the measured mean arrival times in each pair's Δt.

6. 2000 iterations of Steps 1 through 5 were performed.

7. The weighted averages of the 2000 iterations were used to extract two high-statistics lifetimes by either including or excluding the measured mean arrival times in each pair's Δt.

When the mean arrival times of each run were not included in the calculation of the Δts, the extracted lifetime was 0.018 ± 0.003 s less than the true Monte Carlo lifetime. When the mean arrival times of each run are included in the calculation of the Δts, the extracted lifetime was 0.0006 ± 0.003 s less than the true Monte Carlo lifetime. Excluding the mean arrival times from the calculation of the Δts decreased the lifetime extracted from the Monte Carlo data by 0.017255 ± 0.00003 s. The results of this Monte Carlo simulation showed that including each run's mean arrival time in the holding length corrected for shifts in the extracted lifetime due to differences in the distributions for arrival times between short and long runs.

Section 3.8.1 demonstrated that excluding fixed lengths of time from the holding length induced a bias in the extracted lifetimes. In order to quantify the effect of just the variation of the mean arrival times on the extracted lifetime, each run's mean arrival time was decomposed into two parts:

$$\bar{t}_i = \langle \bar{t}_i \rangle + (\bar{t}_i - \langle \bar{t}_i \rangle) \, .$$

The $\langle \bar{t}_i \rangle$s were calculated separately for 2017 and 2018 because of the significant change in mean arrival times between the two years. Two lifetimes were extracted: the first included $\langle \bar{t}_i \rangle$ in each holding length, and the second included $\bar{t}_i$ in each holding length. The second extracted lifetime was 0.04 ± 0.01 s greater than the first extracted lifetime, which roughly agreed with the results of the Monte Carlo simulation developed in this section.

3.9 Global analysis

3.9.1 Introduction

A global analysis began with the assumption that there was one value for the neutron lifetime that accurately represented all of the data in the analysis. There were problems with this assumption (see Sections 4.2 and 4.3), but corrections were made for these problems. The maximum likelihood estimate (MLE) of τ and the normalization parameters over the entire 2017/2018 data set were reported as the global lifetime. Section 3.9.2 will develop a method to find the MLE of τ and will calculate the uncertainty of the MLE of τ. Section 3.9.3 will use Monte Carlo simulations to search for any potential bias in the method developed in Section 3.9.2. Section 3.9.4 will present the results from the global analysis.

3.9.2 Extracting a global lifetime

Maximizing Equation 3.14 required searching over a 304-parameter space where the different parameters were correlated. (The runs from the 2017/2018 data set are split into 101 different normalization subsets. Three parameters were used to make the normalization estimates for each normalization subset. The lifetime was the one additional parameter.) Normalization parameters from different normalization subsets were correlated because they shared one common value for the lifetime. It was difficult to ensure that a numerical maximization algorithm found the global maximum while searching over this many dimensions. If the value for the lifetime was fixed, then the normalization parameters from different normalization subsets would have decoupled and one 304-parameter optimization would have reduced to 101 independent 3-parameter optimizations.

The process from Section 3.7.9 that found the optimal parameters of the normalization model required a value for the lifetime as an input. Extracting a global lifetime required the normalization estimates as inputs. The MLE of τ and the optimal parameters of the normalization model were simultaneously solved for. Golden-section search [52] was used to search for the MLE of τ between 800 and 1000 s. Golden-section search is a convex optimization algorithm that can optimize a convex function to arbitrary precision without calculating any derivatives. In this analysis, the MLE of τ was found to within a precision of 10^{-4} s. For a fixed value of τ, the MLE of the normalization parameters were found using the method described in Section 3.7.9. For each trial value of τ, the log-likelihood from Equation 3.14 was summed over only the long and extra-long runs to reduce correlation between τ and the normalization parameters. After the MLE of τ and the normalization parameters were found and $\mathcal{L}_{\text{min}}$ was calculated, the uncertainty of τ was estimated using the following method:

1. τ was displaced above the MLE of τ;
2. Using the displaced value of τ, the normalization parameters were optimized using the method described in Section 3.7.9;
3. $\mathcal{L}'$ was calculated by summing over only the long and extra-long runs, using the displaced value of τ and the normalization parameters found in Step 2;
4. The displacement of τ was adjusted and Steps 2 and 3 were repeated until $\mathcal{L}' = \mathcal{L}_{\text{min}} + \frac{1}{2}$;
5. Steps 1 through 4 were repeated with an initial displacement of τ below the MLE of τ.

3.9.3 Statistical bias

Many simplifying assumptions were made in Section 3.7 while constructing Equation 3.14. Some of these assumptions could have introduced a bias into the extraction of a global lifetime. A Monte Carlo data set was produced using the first five steps of the Monte Carlo simulation detailed in Section 3.7.9. The inputs to the simulation (the estimate of the number of UCN loaded into the trap and associated uncertainty, the holding length of each run, the estimate of the efficiency of the primary detector and associated uncertainty, the estimate of the software-dead-time and UCN-event-tail corrections, and the estimate of the background CCs) were chosen to match the

values from the analysis of the real data. However, this process did not simulate values for the number of counts recorded in the normalization monitors during the filling process. The simulated data were combined with the real values for the holding lengths and the counts in the normalization monitors to produce a combined data set with characteristics that were similar to the real data set.

The global analysis from Section 3.9.2 was used to extract a lifetime from the combined data set. This procedure was repeated 1000 times to measure any statistical bias to high precision. Each extracted lifetime had a statistical uncertainty that was comparable to the statistical uncertainty from the analysis of the real data. The mean of the lifetimes extracted from the 1000 iterations of the Monte Carlo simulation differed from the true Monte Carlo lifetime by 7 ± 10 ms (p-value = 0.48), so no evidence of a bias was found.

3.9.4 Results and conclusion

The global extracted lifetime was $\tau_g = 877.62 \pm 0.27$ s. When a spectral correction was not used in the normalization estimate (see Equation 3.6), the global lifetime was $\tau_g' = 877.63 \pm 0.29$ s. Figure 3.29 shows the results of a global analysis performed on 3853 runs.

These lifetimes were extracted using the method developed in Section 3.9.2, and Section 3.9.3 used Monte Carlo simulations to confirm that this method did not induce a bias in the extracted lifetimes. The Cramér-Rao bound confirmed that using the method of maximum likelihood estimation was the most efficient way to use the available data to extract the lifetime and normalization estimates [53]. The analysis developed in this section was close to, but was not exactly, a full maximum likelihood estimation over the entire data set. All runs contained some information about both the normalization and the lifetime, but the short runs were used to only estimate the normalization and the long runs were used to only estimate the lifetime. Section 3.10.7 will demonstrate that the information about the lifetime contained in the short runs was insignificant to this analysis.

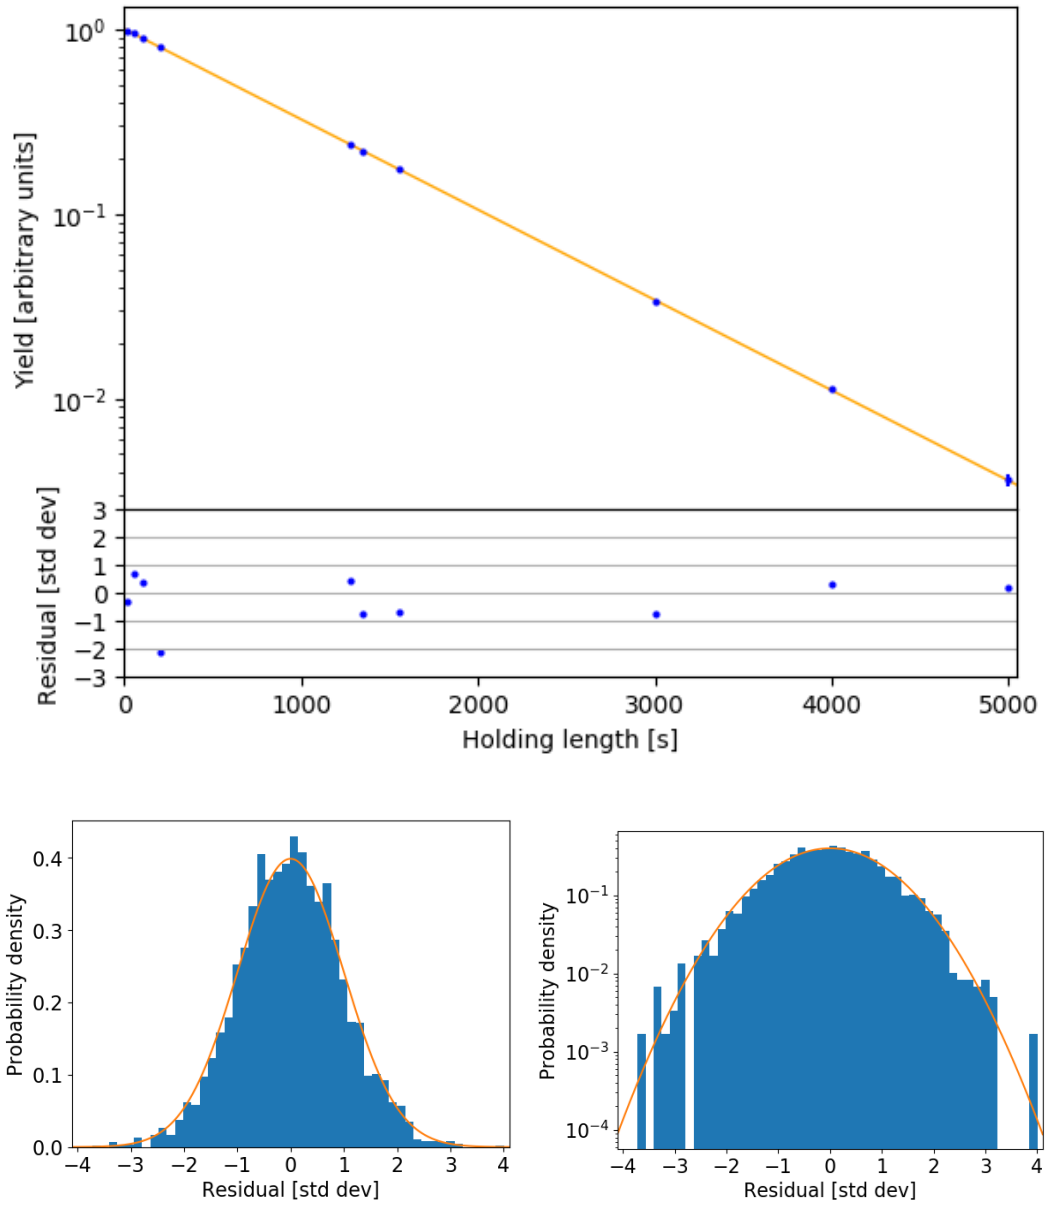

Figure 3.29: The results of a global lifetime fit of all runs. Top: average yields for each holding length, as well as residuals for each average yield. Bottom: differences between the measured and expected number of CCs recorded during the unloaded process of each run, divided by the uncertainty of that difference. In the top plot, all average yields are represented with central values and statistical uncertainty bars, but the bars are too small to see in all but the 5000 s runs. The bottom two plots display the same data, but the right plot uses a logarithmic vertical scale to better demonstrate the behavior of the tails of the distribution. The orange curve is a Gaussian distribution with a mean of 0 and a standard deviation of 1. Both plots in the bottom section share the same value of $\chi^2/\nu = 3657.2/3549$ (p-value = 0.10).

3.10 Paired analysis

3.10.1 Introduction

In contrast to a global analysis, a paired analysis extracted many values for the lifetime and then averaged them together. Each value for the lifetime was extracted using one short run and one long run. Most pairs were constructed from consecutive runs, but sometimes the two runs of a pair had a small number of non-production runs between the two runs of the pair. For a run with measured unload U, background estimate B, and normalization estimate N, a yield was defined as $Y = \frac{U-B}{N}$. For a pair of runs, one short and one long, a lifetime was extracted as

$$\tau = \frac{\Delta t}{\ln r}, \tag{3.18}$$

where $\Delta t = t_L - t_S$, $r = Y_S/Y_L$, and the subscripts S and L denote short and long runs.

There were two obvious benefits to extracting a lifetime using Equation 3.18. First, if there was a constant, but unknown, factor k by which the yields were mis-estimated, taking the ratio of the yields would have cancelled this mis-estimation: $(kY_S)/(kY_L) = Y_S/Y_L = r$. Even if k was not constant but instead only slowly changed, then consecutive runs should have had similar values for k. Small, slow, changes to the UCN source could have led to bias in the normalization, but taking the ratio of the yields of runs from adjacent runs could have reduced this bias. Second, if there was a constant, but unknown, offset δ by which the holding lengths were mis-estimated, then taking the difference of the holding lengths would have canceled this mis-estimation: $(t_L + \delta) - (t_S + \delta) = t_L - t_S$.

Section 3.10.2 will discuss how extracting a lifetime using this method introduces multiple biases to the lifetime and will develop corrections to suppress these biases. Section 3.10.3 will provide an overview of the algorithm used to match short and long runs to form pairs. Section 3.10.4 will develop an algorithm that extracts a lifetime for each pair of runs and averages these lifetimes together to extract one high-statistics measurement of the lifetime. Section 3.10.5 will make use of the Monte Carlo simulation developed in Section 3.9.3 to search for any potential bias that was not suppressed by the corrections in Section 3.10.2. Section 3.10.6 will present the results from the paired analysis. Section 3.10.7 will compare the statistical reach of the global and paired analyses.

3.10.2 Bias corrections

Section 3.7.2 discussed how, in an ideal version of this experiment, the actual number of UCN loaded into the trap at the beginning of each run would have been a Poisson random value. Section 3.7.3 modeled the stochastic behavior of β-decay of UCN in the trap during the holding process as a Binomial random process. As shown in Appendix C, if the number of UCN loaded into the trap was a Poisson random value $n \sim \text{Poisson}(N)$ and the number of UCN that did not decay during the holding process was a random value $U \sim \text{Binomial}(n, e^{-t/\tau})$, then the number of UCN that were in the trap at the end of the holding process was a random value $U \sim \text{Poisson}(Ne^{-t/\tau})$. If two iterations of this idealized experiment were done (one with a short holding length t_S and one with a long holding length t_L), then two values for the number of UCN unloaded from the trap were recorded: $S \sim \text{Poisson}(Ne^{-t_S/\tau})$ and $L \sim \text{Poisson}(Ne^{-t_L/\tau})$.

If S was exactly equal to $Ne^{-t_S/\tau}$ and L was exactly equal to $Ne^{-t_L/\tau}$, then the lifetime could have been exactly recovered as

$$\tau = \left.\frac{\Delta t}{\ln \frac{S}{L}}\right|_{\substack{S=Ne^{-t_S/\tau} \\ L=Ne^{-t_L/\tau}}} = \frac{\Delta t}{\ln \frac{Ne^{-t_S/\tau}}{Ne^{-t_L/\tau}}} = \frac{\Delta t}{\ln e^{\Delta t/\tau}} = \frac{\Delta t}{\frac{\Delta t}{\tau}} = \tau, \tag{3.19}$$

where $\Delta t \equiv t_L - t_S$. However, S and L were not exactly $Ne^{-t_S/\tau}$ and $Ne^{-t_L/\tau}$, but were Poisson random values drawn from distributions with those means. Statistical fluctuations caused the extracted lifetime $\Delta t / \ln \frac{S}{L}$ to differ from the true lifetime τ. One strategy to reduce the effect of the random variations was to repeat this idealized experiment many times and average the results together. $\langle S \rangle$ was the best estimate $Ne^{-t_S/\tau}$, and similarly $\langle L \rangle$ was the best estimate $Ne^{-t_L/\tau}$. Therefore, repeated iterations of this idealized experiment could have been combined to extract the lifetime as

$$\tau = \frac{\Delta t}{\ln \frac{\langle S \rangle}{\langle L \rangle}}. \tag{3.20}$$

The properties of the actual UCNτ experiment did not match the properties of the idealized experiment described in the above two paragraphs. First, Section 3.2.1 discussed why various short and long holding lengths were used instead of using just one short holding length and one long holding length. These different holding lengths caused different pairs of runs to have different values for Δt. Second, the introduction to Section 3.10 discussed why a paired analysis calculated a ratio for each pair of runs. This does not match what was described in the previous paragraph,

where all short and long runs would have been separately averaged before calculating a ratio. A paired lifetime that allowed for different pairs to have different values for Δt and that calculated a ratio on a pair-to-pair basis was extracted as

$$\tau = \left\langle \frac{\Delta t_i}{\ln \frac{S_i}{L_i}} \right\rangle . \tag{3.21}$$

This structure captured the two benefits of a paired analysis described in the introduction to Section 3.10, but it was not equal to Equation 3.20. Equation 3.20 was the proper and unbiased way to extract a paired lifetime, so Equation 3.21 needed to be corrected to be used to extract an unbiased paired lifetime.

Correction terms were derived using Taylor expansions so that a lifetime could be extracted using a similar formula as Equation 3.21. The average of the ratio of yields can be corrected to be approximately equal to the ratio of the average of yields via Eq. The effect of this correction is shown in Figure 3.30.

$$\frac{\langle Y_S \rangle}{\langle Y_L \rangle} \simeq \left\langle \frac{Y_S}{Y_L} \right\rangle \left[1 - \left(\frac{\sigma_{Y_L}}{\langle Y_L \rangle} \right)^2 \right] . \tag{3.22}$$

The average of the pairs of lifetimes can be corrected to be approximately equal to the lifetime extracted from the average ratio of yields via Equation 3.23. Figure 3.31 shows the effect of this correction when used with and without the correction from Equation 3.22.

$$\frac{\Delta t}{\ln \langle r \rangle} \simeq \left\langle \frac{\Delta t}{\ln r} \right\rangle \left[1 - \left(1 + \frac{1}{2} \ln \langle r \rangle \right) \left(\frac{\sigma_r}{\langle r \rangle \ln \langle r \rangle} \right)^2 \right] . \tag{3.23}$$

The combination of both of Equations 3.22 and 3.23 resulted in a method to extract lifetimes for each short-long pair that were averaged to extract a non-biased lifetime. This allowed for a paired lifetime extraction to make use of the two benefits listed at the beginning of this section, while not inducing a bias in the extracted lifetime. The derivation of the corrections in Equations 3.22 and 3.23 are shown in Appendix D.

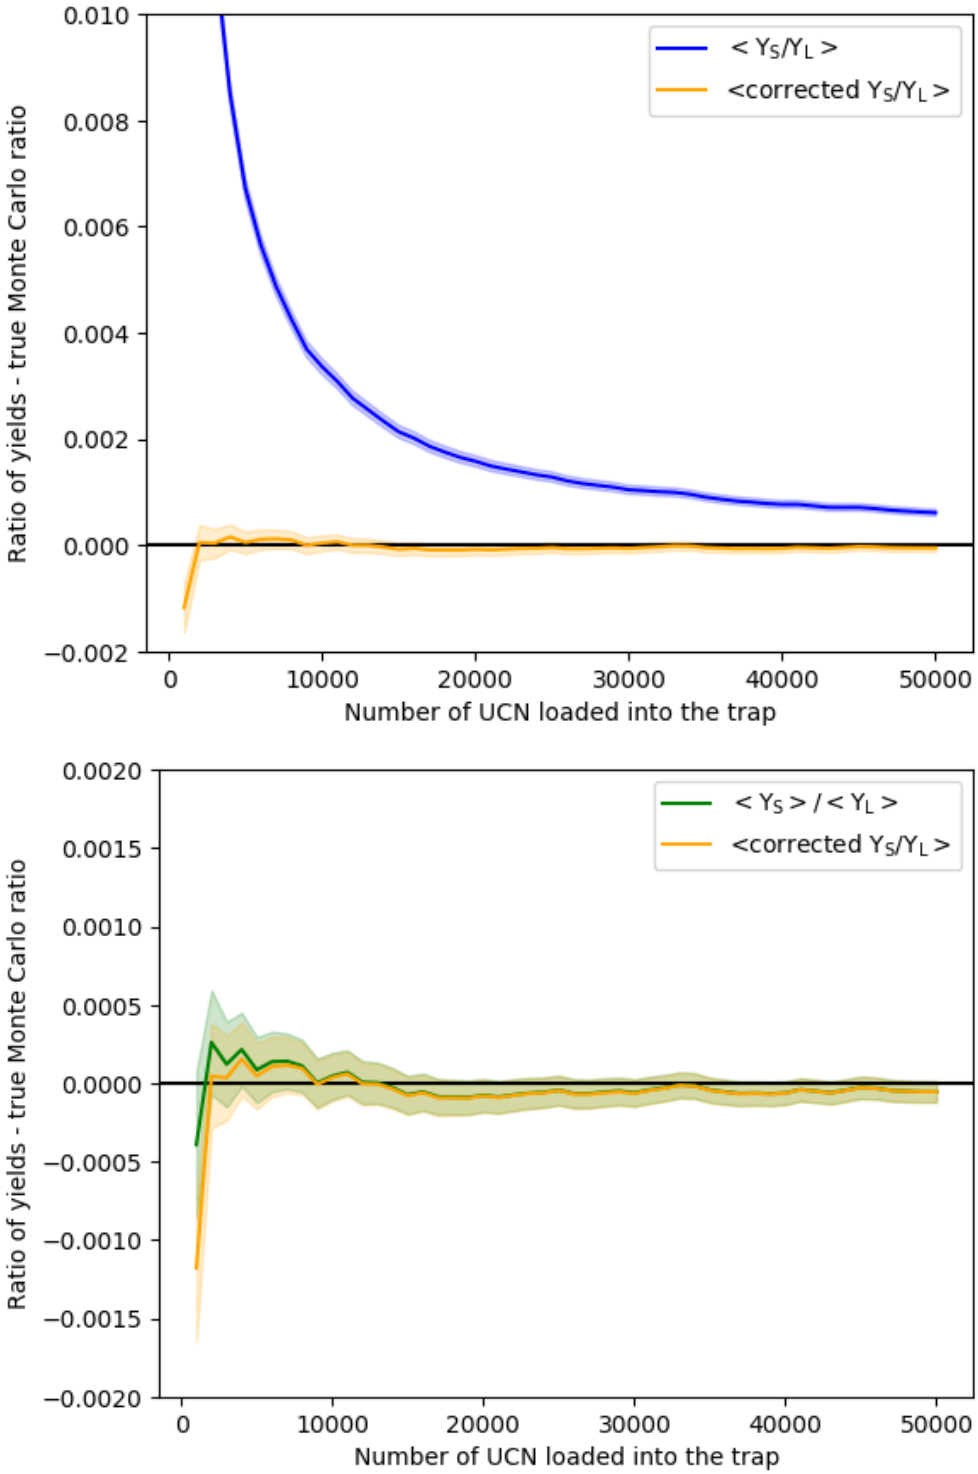

Figure 3.30: Correction to the average of the ratio of yields in a paired lifetime extraction. Top: the effect of the bias correction from Equation 3.22. Bottom: the corrected ratio was approximately equal to the unbiased ratio. The data shown were generated with one million simulated short-long pairs of runs.

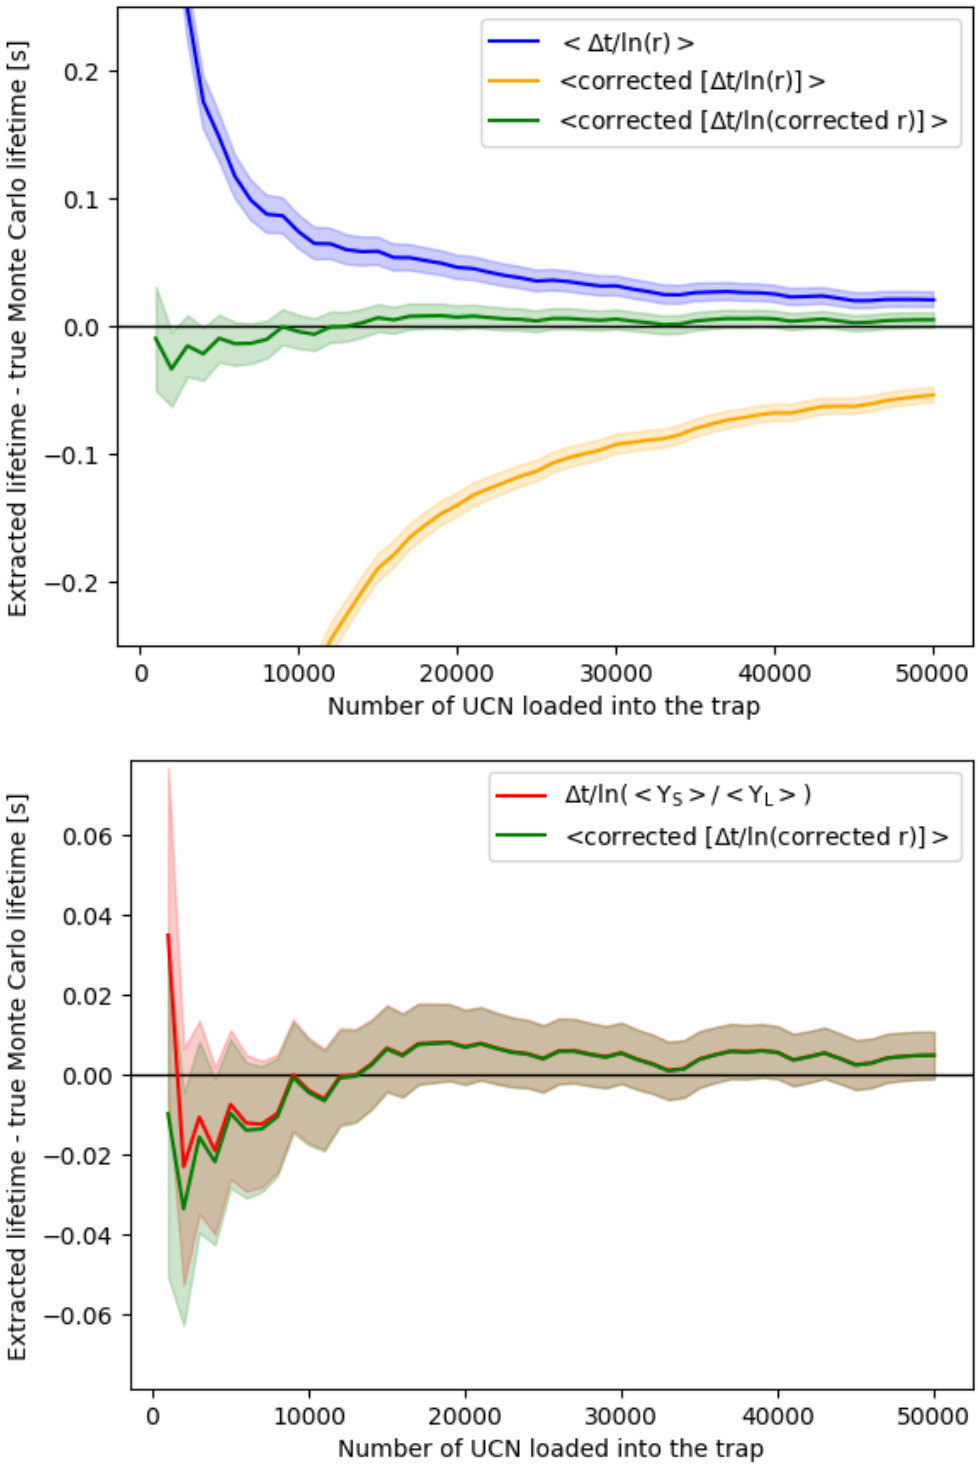

Figure 3.31: Corrections to the average of the paired lifetimes in a paired lifetime extraction. Top: the effect of the bias corrections from Equations 3.22 and 3.23. Bottom: the corrected lifetime was approximately equal to the unbiased lifetime. The data shown were generated with one million simulated short-long pairs of runs.

3.10.3 Forming run pairs

Short runs were paired with a long run from the same normalization subset. All extra-long runs were excluded from a paired analysis because the signal-to-noise ratio during the unloading process was low, so the uncertainty of a lifetime from a short-long pair that includes an extra-long run would have been extremely high. Pairs were chosen so that there was no overlap in time between pairs. That is, both runs in a pair must have occurred either before or after both runs in any other pair. The demands that pairs did not overlap in time, as well as that each pair must have been composed of one short run and one long run, meant that a paired analysis could make use of all of the runs used in a global analysis. The pairing algorithm constructed 1665 pairs from 3853 runs.

3.10.4 Extracting a paired lifetime

The process from Section 3.7.9 that found the optimal parameters of the normalization model required a value for the lifetime as an input. Extracting a paired lifetime required the normalization estimates as inputs. The iterative approach detailed in Appendix E was used to simultaneously solve for τ and the maximum likelihood estimate of the normalization parameters. This approach used 1665 short-long pairs of runs, extracted a lifetime for each pair of runs, and then combined the lifetimes using a weighted average. Past UCNτ analyses avoided using weighted averages to combine lifetimes extracted from short-long pairs of runs because of their tendency to bias the average, even though unweighted averages made less efficient use of the available data. The weighted average used in Step 8 of the algorithm in Appendix E did not induce a significant bias (this will be confirmed in Section 3.10.5) because the uncertainties σ_{Y_i}, σ_{r_j}, and σ_{τ_j} were properly calculated. The key to this process, which was overlooked in the past analyses, was to use the expected values (as opposed to the measured values) while estimating various uncertainties in the algorithm in Appendix E: $\langle U_i - B_i \rangle$ in Step 5, $\langle Y_S \rangle$ and $\langle Y_L \rangle$ in Step 6, and $\left\langle r_j \right\rangle$ in Step 7. These expected values were only able to be calculated given a value for τ, which further necessitated the use of an iterative approach to extract a paired lifetime.

3.10.5 Statistical bias

The Monte Carlo data generated in Section 3.9.3 were also analyzed using the paired lifetime method. The mean of the lifetimes extracted from 1000 iterations of the Monte Carlo simulation was 25 ± 10 ms (p-value = 0.01) less than the true Monte Carlo lifetime. This difference was < 10% of the statistical uncertainty of the extracted lifetime so it was sufficiently small to ignore.

3.10.6 Results

Figure 3.32 shows the lifetimes extracted for 1665 short-long pairs. Each short-long pair had one value for the lifetime with a statistical uncertainty of $\sim 10 - 20$ s, depending on levels of normalization. A weighted average of the lifetimes resulted in a paired lifetime of $\tau_p = 877.70 \pm 0.32$ s. When a spectral correction was not used in the normalization estimate (see Equation 3.6), the paired lifetime was $\tau_p' = 877.68 \pm 0.36$ s.

The Monte Carlo simulation from 3.9.3 was also used to explain the difference between the global lifetime extracted in Section 3.9 and the paired lifetime extracted in this section. The exact same simulated data were analyzed with both the global and paired methods. The pair-to-pair difference was found to be $\langle \tau_g - \tau_p \rangle = 18.1 \pm 2.6$ ms. This standard error was calculated from 1000 iterations of the Monte Carlo simulation, so the standard deviation of the distribution was 0.08 s. The difference between the global and paired lifetimes extracted from the real data was $\tau_g - \tau_p = -0.08$ s. The magnitude of this difference was not statistically significant (p-value = 0.23) when compared to the difference of 18.1 ms measured in Monte Carlo simulations.

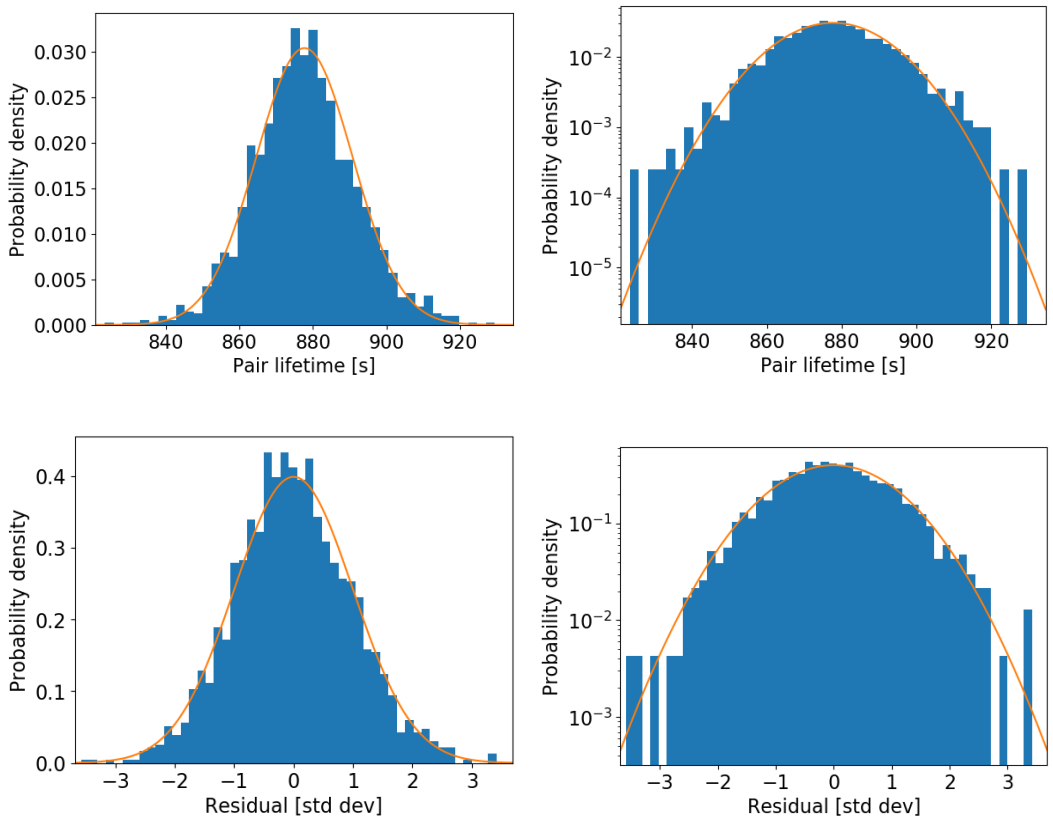

Figure 3.32: The results of a paired lifetime extraction from all pairs of runs. Top: lifetimes extracted from each short-long pair of runs. Bottom: differences between the lifetime extracted from each short-long pair of runs compared to the mean of all the paired lifetimes, divided by the uncertainty of that difference. In both the top and bottom sections, the two plots display the same data, but the right plots use logarithmic vertical scales to better demonstrate the behavior of the tails of the distribution. Each lifetime in the top plots has its own statistical uncertainty that was determined by the characteristics of the two runs in that pair. The statistical uncertainty of the lifetime of one pair was ~10-20 s. Most of the variation of statistical uncertainty was due to variation of the number of UCN loaded into the trap at the beginning of each run. The orange curve in the top plots is a Gaussian distribution with a mean of τ_p and a standard deviation of $\sigma_{\tau_p} \times \sqrt{1663}$. The subset of lifetimes with the largest statistical uncertainties ($\sim$ 20 s) had a distribution that was wider than the orange Gaussian curve, but this was not the correct comparison to make. A better comparison is shown in the residuals in the bottom plots, which all had an expected value of 0 and a standard deviation of 1. The orange curve in each of the bottom plots is a Gaussian distribution with a mean of 0 and a standard deviation of 1. After appropriately considering the statistical uncertainty of each individual lifetime, the distribution of residuals did match a standard Gaussian distribution. All four plots share the same value of $\chi^2/\nu = 1605.8/1663$ (p-value= 0.84).

3.10.7 Inefficiencies of the paired analysis

The statistical uncertainty of a paired lifetime was larger than the statistical uncertainty of a global lifetime for two reasons. First, the paired lifetime extraction used less of the available data because not every run could have been reasonably fit into a short-long pair. The global lifetime extraction made use of 1799 long and extra-long runs, and during the unloading process of those runs, a total of 5.7 million coincidence chains (CCs) were counted above background. The paired lifetime extraction made use of 1665 long runs, and during the unloading process of those runs, a total of 4.5 million CCs were counted above background. Second, the paired lifetime extraction made less efficient use of the data that it did use. Section 3.9.4 discussed how the global lifetime extraction was nearly as efficient as a maximum likelihood analysis of the entire data set. Monte Carlo simulations have found that when the exact same underlying data were analyzed with the global and paired lifetime extraction methods, the statistical uncertainty of the paired lifetime was $\simeq 8\%$ greater than the statistical uncertainty of the global lifetime. When the 0.27 s of statistical uncertainty from the global lifetime was inflated by $\sqrt{5.7 \div 4.5} \times 1.08$, it increased to 0.33 s. This was sufficiently close to the 0.32 s of statistical uncertainty from the paired lifetime to claim that the difference between statistical uncertainties from the global and paired lifetime extractions was understood.

3.10.8 Conclusion

This section developed an alternative method of extracting the lifetime. The paired method was designed to suppress some potential errors in the normalization. Section 3.10.2 discussed how the paired method was a biased way of extracting the lifetime and developed methods to suppress these biases. Section 3.10.4 explicitly stated how these corrections were used to extract a lifetime from each pair of runs and how these lifetimes were averaged to extract a high-statistics lifetime. Section 3.10.5 used Monte Carlo simulations to confirm that the corrections developed in Section 3.10.2 suppress any biases to a level that was negligible in this analysis. Section 3.10.6 presented the results of extracting a paired lifetime and Section 3.10.7 explained why the paired method was less efficient at utilizing the available data than the global method.

Given the results shown in this section, the global lifetime from Section 3.9 was reported as *the* lifetime of a free neutron extracted from the 2017/2018 UCNτ data set, and the paired lifetime was considered partial validation of the extracted global lifetime.

3.11 Conclusion

Analysis of the 2017/2018 UCNτ data set resulted in a global lifetime of $\tau_g = 877.62 \pm 0.27$ s and a paired lifetime of $\tau_p = 877.70 \pm 0.32$ s. A total of 3.3×10^7 background-subtracted reconstructed events caused by UCN that were captured by the primary detector were used to extract the global lifetime. The difference between the global and paired lifetimes of $\tau_g - \tau_p = -0.08$ s was not statistically significant, and the differences in the statistical uncertainties were explained in Section 3.10.7.

The Monte Carlo simulation from Section 3.9.3 was repurposed to measure contributions to the statistical uncertainty. Recalling the likelihood formulation developed in Section 3.7, an ideal version of UCNτ would have had the following features:

1. The normalization estimate would have had Poisson uncertainties ($f = 1$);
2. The primary detector would have had perfect efficiency ($s = 1$ and $\sigma_s = 0$);
3. The primary detector would not have had software dead time, UCN-event-tail interactions, or other rate-dependent effects ($\sigma_C = 0$); and
4. There would have been no backgrounds ($B = 0$).

All of these features were implemented in Monte Carlo simulations to study their contributions to the statistical uncertainty. The results from this study are listed in Table 3.7, along with the 0.01 s statistical uncertainty due to fluctuations in the mean arrival time that was measured in Section 3.8.2. Given the finite number of UCN captured by the primary detector, the imperfections in the experiment contributed only $\sim 15\%$ of the total statistical uncertainty of the extracted lifetime. Chapter 5 will discuss some of the ways that a future iteration of the UCNτ experiment can decrease the statistical uncertainty of the extracted lifetime.

Source of statistical uncertainty	Amount of statistical uncertainty [s]
Finite number of UCN in the trap	0.23
Additional normalization uncertainty	0.10
Primary detector inefficiency	0.09
Background estimation	0.04
Primary detector rate-dependent effects	0.01
Fluctuations in mean arrival times	0.01
Total [uncorrelated sum]	0.27

Table 3.7: Sources of statistical uncertainty of the extracted global lifetime.

Chapter 4

SYSTEMATIC EFFECTS

4.1 Introduction

The lifetimes extracted in Chapter 3 were the lifetimes of UCN lost from the trap by any mechanism. The dominant loss mechanism was β-decay of UCN. The existence of other loss mechanisms would have caused the extracted lifetime to be less than the neutron β-decay lifetime (τ_β). The extracted lifetime was related to τ_β by

$$\frac{1}{\tau_{\text{extracted}}} = \frac{1}{\tau_\beta} + \frac{1}{\tau_{\text{other}}}.$$

For $\tau_{\text{extracted}} = 877.62$ s, $\tau_{\text{other}} = 10^7$ s would have caused $\tau_\beta - \tau_{\text{extracted}} = 0.08$ s. Even very small imperfections in the trap could have caused non-negligible shifts in the extracted lifetime.

Most other experiments that stored UCN to measure the lifetime of the free neutron used traps with material walls [23, 25]. The material walls of the traps in these other experiments were chosen to be very good reflectors of UCN, but there were still losses due to material interactions and significant corrections (~ 13 s and $\lesssim 4$ s, in the same order as the citations in the previous sentence) needed to be applied to the extracted neutron lifetimes. In contrast, the UCNτ experiment uses a magnetic field as the walls of the trap. UCN in the UCNτ trap did not interact with any material walls while they were stored in the trap. However, there were two cases where UCN in the UCNτ trap did interact with other materials. Section 4.2 will discuss interactions between UCN and an aluminum block that was accidentally inserted into the UCNτ trap. Section 4.3 will discuss interactions between UCN and residual gasses that remained inside of the UCNτ trap. The corrections that will be developed in Sections 4.2 and 4.3 will be applied on a run-to-run basis before extracting a lifetime.

Sections 4.4 and 4.5 will discuss systematic uncertainties due to choices in how to perform an analysis of these data. Sections 4.6 through 4.8 will discuss three non-β-decay mechanisms through which UCN could have been lost from the trap, and how limits were placed on the rates of these losses. The potential losses from Sections 4.6 through 4.8 accumulated as the holding length increased, so long runs would have been more susceptible to these losses than short runs. Therefore, corrections

to account for these non-β-decay losses of UCN could have only increased the extracted lifetime. In contrast to the corrections that will be developed in Sections 4.2 and 4.3, the corrections that will be developed in Sections 4.4 through 4.8 were applied to the lifetime after it was extracted from the data.

4.2 Aluminum block interactions

During a routine maintenance period in 2017, the vacuum jacket around the trap was opened to adjust some of the hardware inside of the trap. At this time it was discovered that there was an aluminum block resting on the Halbach array. Upon further examination, it was determined that this block was part of the assembly that actuated the small cleaner. The location of the block was well-documented and the block was removed from the trap.

Some trapped UCN interacted with the aluminum block. Although aluminum is a very good reflector of UCN, it is not perfect, so some UCN were lost from the trap due to interactions with the aluminum block. UCN in long runs had more opportunities to interact with the aluminum block so there were more losses per UCN in long runs than in short runs. If not corrected for, this difference would have biased the extracted lifetime to be less than it should have been.

4.2.1 Identifying when the aluminum block fell into the trap

The aluminum block was removed from the trap before run 7327. It was impossible to know exactly when the aluminum block fell from the small cleaner into the trap, but that time was well-estimated. The trajectories of UCN in the trap evolved during the holding process, as directly measured in Section 3.8.2. Figure 4.1 shows the effect that the aluminum block had on the trajectories of UCN inside of the trap. The magnitude of the increase in the fraction of UCN counted in Peak 2 that occurred when the aluminum block was removed before run 7327 was similar to the magnitude of the decrease in the fraction of UCN counted in Peak 2 that occurred before run 4620. These changes were also seen, albeit less clearly, in changes to the mean arrival time as shown in Figure 3.28. It was unknown what caused the smaller change around run 6200. One possible explanation could be that the aluminum block shifted position along the surface of the Halbach array. For the purposes of studying the effect of the aluminum block on the extracted lifetime, it was assumed that the block fell into the trap just before run 4620 and did not move again until it was removed just before run 7327. Runs 4620-7326 contained 32% of the production runs that were used to extract lifetimes in Chapter 3.

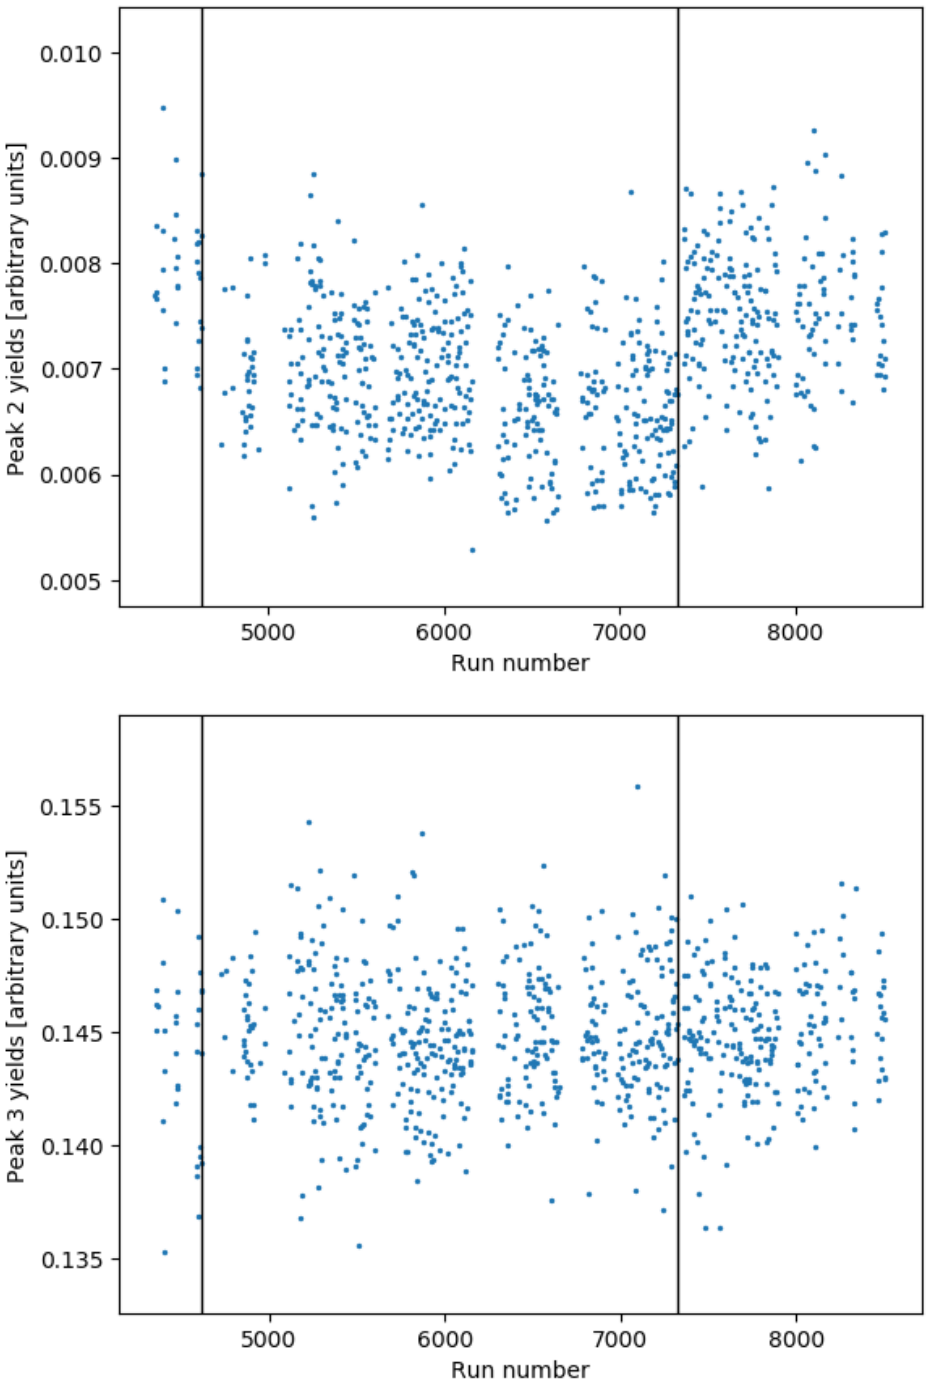

Figure 4.1: Fraction of total background-subtracted CCs found during Peaks 2 (top) and 3 (bottom) of the unloading process for all of the production runs gathered during the 2017 run cycle. The black lines at run 7327 show where the aluminum block was removed from the trap. The black lines at run 4620 are the estimate of when the aluminum block fell into the trap. It was unknown what caused the smaller change in Peak 2 behavior around run 6200. The aluminum block affected high-energy UCN from Peak 2 much more than low-energy UCN from Peak 3.

4.2.2 Estimating the effect of the aluminum block on trapped UCN

As discussed in Section 2.4, aluminum is a very good reflector of UCN and has a loss per bounce of $\mu_{\mathrm{Al}} < 10^{-4}$. The aluminum block's effect on the extracted lifetime was much smaller than the statistical uncertainty of the lifetime extracted using only the data gathered when the aluminum block was believed to be in the trap. In order to measure the effect of the aluminum block on UCN in the trap, a polyethylene block of the same shape and size as the aluminum block was installed into the same location in the trap. The aluminum block was located 6.32 cm above the bottom of the trap, so some UCN in the trap had insufficient energy to interact with the aluminum block. Polyethylene has a loss per bounce of $\mu_{\mathrm{poly}} \sim 1$, so the effect on the extracted lifetime was significantly amplified. 27 runs were performed in 2017 with the polyethylene block installed in the trap, and the effect of interactions between UCN in the trap and the polyethylene block were measured.

The probability that UCN were not lost from the trap during the holding process was updated as shown in Equation 4.1, where κ' is the loss lifetime of UCN that could have interacted with the polyethylene block, A is the fraction of UCN in the trap that could have interacted with the polyethylene block, and τ is the extracted global lifetime. Note that κ' is the loss lifetime by any mechanism, which included β-decay and losses from interactions with the aluminum block.

$$e^{-t/\tau} \rightarrow Ae^{-t/\kappa'} + (1 - A)e^{-t/\tau} \tag{4.1}$$

Due to the small number of runs with the polyethylene block in the trap, the value for f was held fixed at the value from the rest of the 2017 data instead of calculating it based on the data gathered with the polyethylene block installed in the trap. There were eight free parameters used to fit the data: A, κ', and six normalization parameters (three each for two normalization subsets). This relatively low number of free parameters allowed for finding the maximum likelihood estimate of the free parameters by directly minimizing a version of Equation 3.14 that had been updated using Equation 4.1. During the minimization, the value for τ was held fixed at the the value for τ_g reported in Section 3.9. Figure 4.2 shows the result of this minimization, which resulted in $A = 0.74 \pm 0.02$ and $\kappa'_{\mathrm{poly}} = 424 \pm 11$ s. This value for κ'_{poly} represented the lifetime for losses of any mechanism for UCN that have enough energy to interact with the polyethylene block. The lifetime for losses only due to interactions with the polyethylene block was

$$\kappa_{\mathrm{poly}} = \left[\left(\kappa'_{\mathrm{poly}}\right)^{-1} - \tau^{-1}\right]^{-1} = 812 \pm 41 \text{ s}.$$

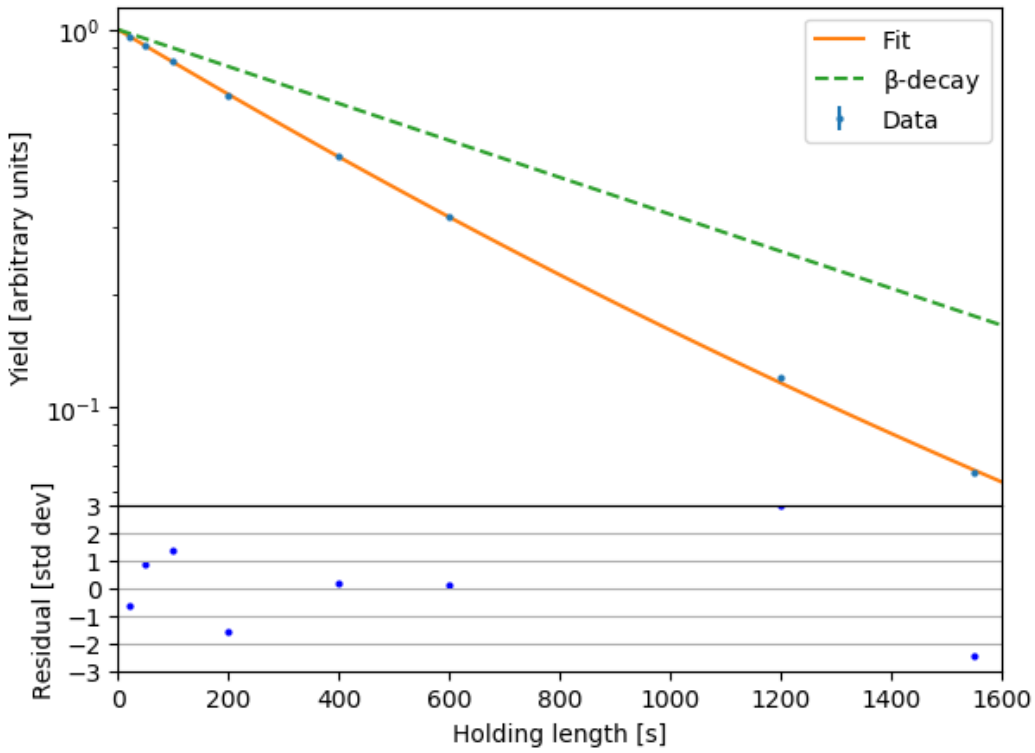

Figure 4.2: The results of fitting the data with a polyethylene block installed in the trap. The data points are the average yields for each holding length, as well as the associated residuals. The residuals of the fit have $\chi^2/\nu = 35.7/19$ (p-value = 0.011). The scale of the uncertainty of the normalization estimate was held constant at the value from the entirety of the 2017 production data, which could have contributed to the high value for χ^2/ν.

The loss per bounce of UCN on aluminum is

$$\mu_{\mathrm{Al}}(E,\theta) = 2f_{\mathrm{Al}}\sqrt{\frac{E\cos^2\theta}{V_{\mathrm{Al}} - E\cos^2\theta}}, \tag{4.2}$$

where $f_{\mathrm{Al}} = 2.25\times10^{-5}$ is the UCN loss factor, E is the UCN energy, θ is the angle of incidence relative to normal, and $V_{\mathrm{Al}} = 54$ neV is the aluminum wall potential [33].

Polyethylene has a negative Fermi potential of $V_{\text{poly}} = -8.7$ neV [33]. As discussed in Section 2.4.1 of [33], the loss per bounce of UCN on any material with a negative Fermi potential can be calculated by considering a three dimensional quantum mechanical wave incident on a one-dimensional negative step potential. Therefore, the loss per bounce of UCN on polyethylene was

$$\mu_{\text{poly}}(E,\theta) = 1 - \left[\frac{\sqrt{E\cos^2\theta} - \sqrt{E\cos^2\theta - V_{\text{poly}}}}{\sqrt{E\cos^2\theta} + \sqrt{E\cos^2\theta - V_{\text{poly}}}}\right]^2. \tag{4.3}$$

The aluminum block was found, and the polyethylene block was positioned, 6.32 cm above the bottom of the trap. During the cleaning process, the cleaners removed almost all UCN from the trap with sufficient energy to reach 38 cm of height. This height difference of 31.68 cm corresponded to a maximum UCN energy of 32.5 neV. The energy distribution for UCN within the trap was not well-understood, so extreme cases were used to place limits on the effect of the aluminum block on the extracted lifetime. A lower-bound of the UCN energy distribution was a uniform energy distribution between 0 and 32.5 neV. An upper-bound of the UCN energy distribution assumed that all UCN have 32.5 neV of energy.

Using the shorthand $f(\theta) = \sin 2\theta \sin\theta$, the lower-bound constraint on the UCN energy distribution resulted in

$$\left\langle \frac{\mu_{\text{poly}}}{\mu_{\text{Al}}} \right\rangle = \frac{1}{\int_0^{32.5\text{ neV}} dE \int_{-\pi/2}^{\pi/2} f(\theta)d\theta} \int_0^{32.5\text{ neV}} dE \int_{-\pi/2}^{\pi/2} f(\theta)d\theta \frac{\mu_{\text{poly}}(E,\theta)}{\mu_{\text{Al}}(E,\theta)} = 7114.$$

The upper-bound constraint on the UCN energy distribution resulted in

$$\left\langle \frac{\mu_{\text{poly}}}{\mu_{\text{Al}}} \right\rangle = \frac{1}{\int_{-\pi/2}^{\pi/2} f(\theta)d\theta} \int_{-\pi/2}^{\pi/2} f(\theta)d\theta \frac{\mu_{\text{poly}}(32.5\text{ neV},\theta)}{\mu_{\text{Al}}(32.5\text{ neV},\theta)} = 4315.$$

These values were further increased and decreased by 25% to account for potential losses due to rough surfaces, which resulted in

$$3452 \leq \left\langle \frac{\mu_{\text{poly}}}{\mu_{\text{Al}}} \right\rangle \leq 8893.$$

If UCN in the trap with the aluminum block followed the same trajectories as UCN in the trap with the polyethylene block, then the rate of bounces per second on the blocks would be a conserved quantity, up to losses on the blocks.

$$\frac{\text{bounces}}{\text{second}} = \frac{\text{losses}}{\text{second}} \div \frac{\text{losses}}{\text{bounce}} = \kappa^{-1} \div \mu = \frac{1}{\kappa\mu},$$

so $\kappa\mu$ was also a conserved quantity. Therefore, $\kappa_{\text{Al}} = \frac{\mu_{\text{poly}}}{\mu_{\text{Al}}}\kappa_{\text{poly}}$, so

$$2.8 \times 10^6 \text{ s} \leq \kappa_{\text{Al}} \leq 7.2 \times 10^6 \text{ s}.$$

Recall that κ_{Al} was the loss rate only due to UCN interactions. The loss rate by any mechanism of UCN that had sufficient energy to interact with the aluminum block was

$$\kappa'_{\text{Al}} = \left(\frac{1}{\kappa_{\text{Al}}} + \frac{1}{\tau}\right)^{-1},$$

so 877.35 s $\leq \kappa'_{\text{Al}} \leq$ 877.51 s. These values for κ'_{Al} corresponded to a reduction in the lifetime in the range of 0.11 to 0.28 s.

4.2.3 Correcting the extracted lifetime for losses from interactions with the aluminum block

Section 3.7.4 discussed the detector efficiency of the primary detector. The number of coincidence chains (CCs) per UCN (s) and the variance in the number of CCs per UCN (σ_s^2) were updated to correct for losses due to the aluminum block. When the aluminum block was not in the trap, the fraction of UCN in the trap that were expected to still be in the trap after the holding process was $p = e^{-t/\tau}$. When the aluminum block was in the trap, the fraction of UCN in the trap that were expected to still be in the trap after the holding process was $p' = Ae^{-t/\kappa'_{\text{Al}}} + (1 - A)e^{-t/\tau}$. The values for s and σ_s^2 for runs 4620 through 7326 were updated using

$$\begin{aligned} s &\rightarrow s - s\left(1 - \frac{p'}{p}\right) = s\left(\frac{p'}{p}\right) \\ \sigma_s^2 &\rightarrow \sigma_s^2 + s\left(1 - \frac{p'}{p}\right)\left[1 - \left(1 - \frac{p'}{p}\right)\right] = \sigma_s^2 + s\left(\frac{p'}{p}\right)\left(1 - \frac{p'}{p}\right). \end{aligned} \tag{4.4}$$

Section 4.2.2 estimated that the aluminum block decreased the lifetime of UCN in the trap that could have interacted with it by between 0.11 s and 0.28 s. The lifetime of UCN in the trap that did not interact with the aluminum block was not affected. Of all of the UCN in the trap, a fraction $A = 0.74 \pm 0.02$ interacted with the aluminum block and a fraction $1 - A$ did not interact with the aluminum block. A simple scaling of the shift in lifetime by the factor A led to the estimate that the aluminum block decreased the lifetime of UCN in the trap by between $A \times 0.11$ s = 0.08 s and $A \times 0.28$ s = 0.21 s.

The corrections from Equation 4.4 were applied to the runs where it was believed that the aluminum block was in the trap. This correction shifted the lifetime extracted from only the data where the aluminum block was believed to be in the trap by between 0.08 s and 0.21 s, which exactly matched the earlier estimation of $A(\tau - \kappa'_{\mathrm{Al}})$. The correction shifted the lifetime extracted from all of the data by between 0.03 s and 0.09 s. Therefore, the correction to the extracted lifetime was reported as a systematic correction of 0.06 ± 0.03 s.

4.3 Residual gas interactions

When the aluminum block from Section 4.2 was not in the trap, trapped UCN did not interact with any material walls. However, trapped UCN still interacted with residual gas molecules that were inside the vacuum jacket. These interactions could have upscattered or captured trapped UCN and caused them to be lost from the trap. Longer holding lengths presented more opportunities for a residual gas molecule to upscatter or capture a UCN, so long runs had more losses than short runs. Therefore, the systematic effect of losses of UCN due to interactions with residual gasses can have only decreased the extracted lifetime.

Typical gas pressures within the vacuum jacket were $\sim 2 \times 10^{-6}$ torr. However, the trap pressure was sometimes brought up to atmospheric levels in order to open the vacuum jacket and adjust the hardware installed inside of it. N_2 was flushed through the trap while this was done in order to limit the amount of other gasses from the experimental area that entered the trap, but some gasses still made it in. After resealing the vacuum jacket, the volume was pumped out until pressures reached $\sim 10^{-5}$ torr. This pressure was roughly 5× higher than ideal running conditions, but it was sufficiently low to safely re-energize the photomultiplier tubes (PMTs) on various detectors and to resume running the experiment. Over the next few days, the pressure inside the vacuum jacket slowly decreased to levels $\sim 2 \times 10^{-6}$ torr.

A residual gas analyzer (RGA) and multiple cold-cathode pressure gauges were used to measure the variations of the composition and pressure of residual gasses inside of the vacuum jacket. Pressure data were gathered continuously. However, actively using the RGA to measure the composition of the residual gasses increased the background PE rates in the PMTs of multiple detectors, so RGA measurements were only sparsely gathered. The sparseness of RGA measurements was not a problem because once the vacuum jacket was sealed and pumped down, significant quantities of new gasses could not have entered the trap.

Previous studies done by the UCNτ collaboration used a combination of physical measurements and Monte Carlo simulations to calculate the total (upscatter and capture) cross sections between UCN and various gas molecules [54, 55]. The major components of the residual gasses in vacuum jacket were H_2O, N_2, and O_2, but trace amounts of heavy hydrocarbons were also present. The rate of loss of UCN due to interactions with residual gasses was

$$\frac{1}{\tau_{\text{gas}}} = \bar{v} \sum_i \rho_i \sigma_i, \tag{4.5}$$

where $\bar{v} \simeq 1.5$ m/s is the mean velocity of UCN in the trap, ρ_i is the number density of a type of gas molecule, σ_i is the total (upscatter and capture) cross section between UCN and a type of gas molecule, and i sums over all different types of residual gasses in the trap. The data from RGA measurements were used to decompose the average pressure in the trap during each run into partial pressures for various gasses, which were then used to calculate the number density for each type of gas in the trap. τ_{gas} was calculated for each run, and these values are shown in Figure 4.3.

A correction similar to the one developed in Section 4.2.3 was used to correct for losses due to interactions between UCN and residual gasses. In the absence of residual gasses, the fraction of UCN in the trap that were expected to still be in the trap after the holding process was $p = e^{-t/\tau}$. When residual gasses remained in the trap, the fraction of UCN in the trap that were expected to still be in the trap after the holding process was

$$p' = \exp\left[-t\left(\tau^{-1} + \tau_{\text{gas}}^{-1}\right)\right].$$

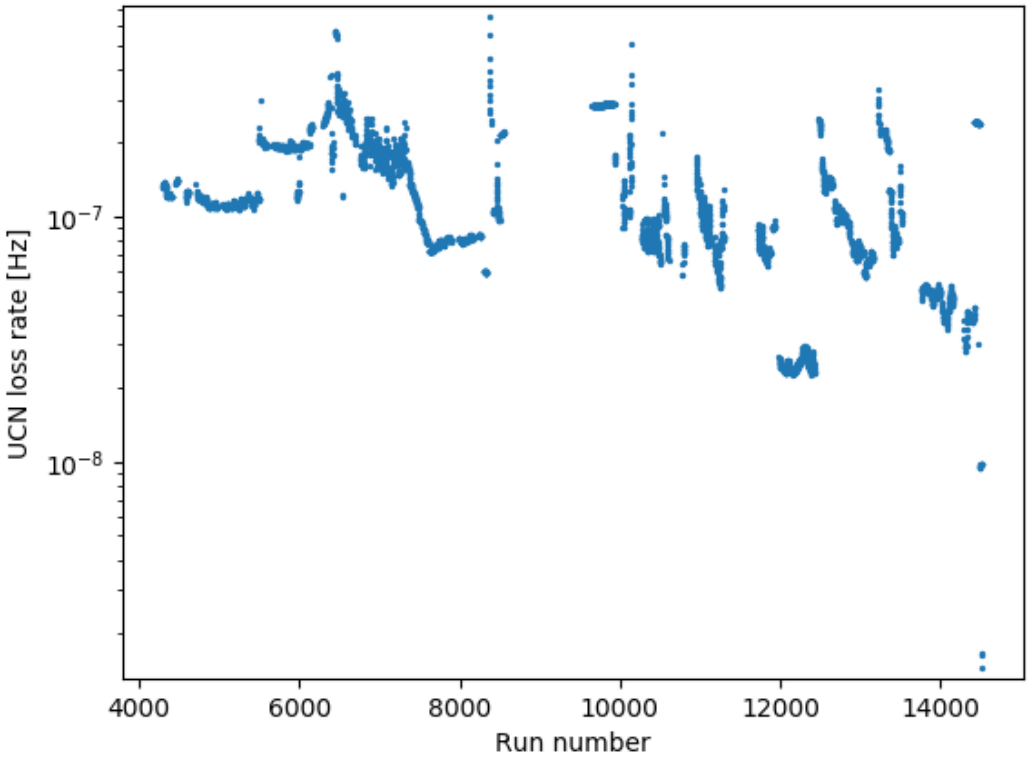

Figure 4.3: Rate of loss of UCN due to upscattering on residual gasses that remained inside of the trap. The rates follow a pattern of decreasing over time as the pressure inside of the trap continued to decrease as the vacuum system continued to pump on it, followed by sudden spikes when the trap was opened. Pressure data was not properly recorded for a small subset of data around run number 9500, so a constant estimated loss rate was used for all of those runs.

$p'/p = e^{-t/\tau_{\text{gas}}}$ was calculated using each run's value for τ_{gas}. Each run's value for p'/p was used to update each run's values for s and σ_s^2 using Equation 4.4. These corrections increased the extracted lifetime by 0.11 ± 0.06 s, and this shift was reported as a systematic correction to the extracted lifetime due to interactions between UCN and residual gasses. The uncertainty of 0.06 s was extracted by varying all runs' value for τ_{gas} by $\sigma_{\tau_{\text{gas}}}$ in the same direction. This assumed maximal correlation in the uncertainties between runs, which was a reasonable assumption because the uncertainties were all due to the systematic uncertainties of measuring the partial pressures in the trap.

4.4 Definition of a reconstructed UCN event in the primary detector

Section 3.4 discussed the characterization of UCN events in the primary detector and how these UCN events were reconstructed. Section 3.4.8 considered 125 different sets of parameters for the reconstruction algorithm. A Monte Carlo simulation was used to demonstrate that certain classes of parameters induced a bias in the extracted lifetime, but most sets of parameters resulted in extracted lifetimes that were not significantly different from the true Monte Carlo lifetime. In the end, one set of parameters was chosen with the goals of maximizing reconstruction efficiency and limiting retriggering.

Although one set of parameters was chosen and used in all analyses presented in this document, there was no guarantee that this set of parameters did not induce a bias that would have resulted in an extracted lifetime that was different from the true neutron lifetime. In order to estimate this uncertainty, the entire analysis presented in Chapter 3 was redone with the 125 different sets of parameters for the reconstruction algorithm. The distribution for the lifetimes extracted using these different sets of parameters is show in Figure 4.4. The same distribution, but for lifetimes extracted from Monte Carlo simulations instead of from the real data, is shown in Figure 3.13.

10 of the 125 sets of parameters of the reconstruction algorithm had $\geq 10^{-2}$ retriggers per UCN. Section 3.4.4 discussed why high retriggering rates could have caused the extracted lifetime to be biased. In order to suppress this potential bias, the 10 sets of parameters that caused $\geq 10^{-2}$ retriggers per UCN were rejected. The remaining 115 sets of parameters were used to extract 115 different values for the lifetime, all of which were reasonably correct.

In Monte Carlo simulations, the 115 values for the extracted lifetime had a mean of 0.04 s below the true Monte Carlo lifetime and a standard deviation of 0.05 s. When the real data were analyzed using the global method from Section 3.9, the 115 values for the extracted lifetime had a mean of 877.60 s (0.02 s below the reported global lifetime) and a standard deviation of 0.08 s. When the paired method from Section 3.10 was used, the distribution had a mean of 877.70 s (equal to the reported paired lifetime) and a standard deviation of 0.09 s. Thus, 0.08 s was reported as a systematic uncertainty of the extracted lifetimes due to uncertainty of the definition of a reconstructed UCN event in the primary detector.

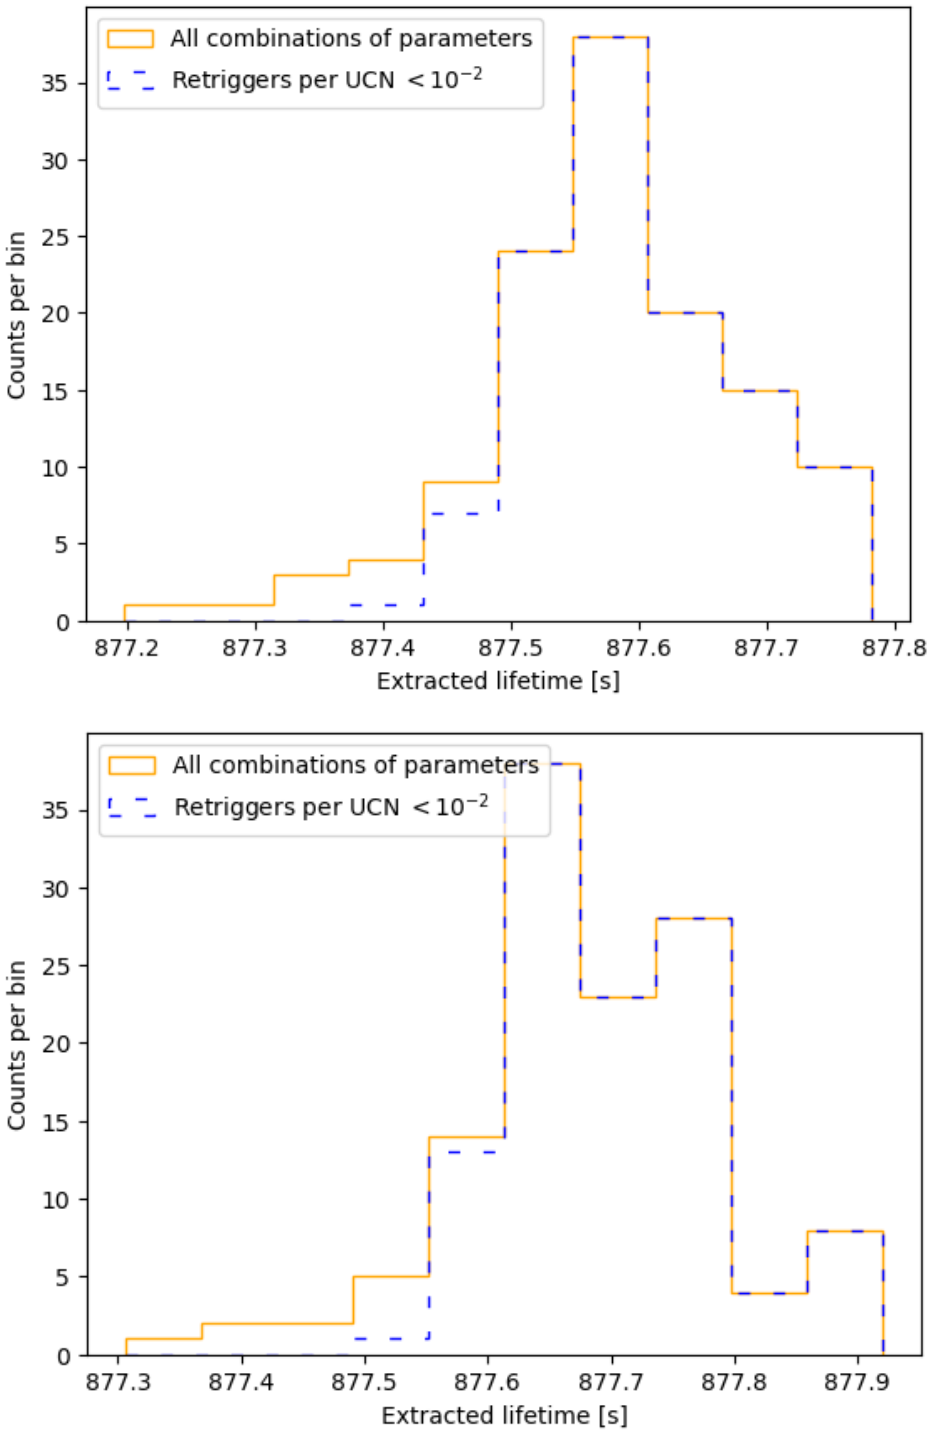

Figure 4.4: Distributions for the lifetimes extracted from the real data by applying the reconstruction algorithm with the various choices of parameters listed in Section 3.4.8. Top: the lifetimes extracted using the global analysis from Section 3.9. Bottom: the lifetime extracted using the paired analysis from Section 3.10.

4.5 Choice of normalization weighting function

Section 3.6 detailed how the number of UCN in the trap at the end of each filling process was estimated. This estimate began by taking weighted integrals of the counts recorded in the normalization monitors during the filling process. The weighting function that was used in the integral was Equation 3.2. The two constants in that function, $T_{\text{fill}}/2$ and $T_{\text{fill}}/10$, were chosen arbitrarily. $T_{\text{fill}}/2$ was the "turning location," which controlled where the weighting function would transition from low to high weighting. $T_{\text{fill}}/10$ was the "turning rate," which controlled how quickly the weighting function would transition from low to high weighting.

In order to quantify the effect of this choice on the extracted lifetime, the analysis from Section 3.9 was rerun using various values for the turning rate and location constants. Figure 4.5 shows the extracted lifetimes, and the uncertainties of those lifetimes, as the turning location and turning rate were varied. The extracted lifetimes shown in Figure 4.5 had a standard deviation of 0.06 s, which is reported as a systematic uncertainty due to choice of normalization model. This study was repeated while analyzing just the data from just the 2017 run cycle (without the buffer volume) and the data from just the 2018 run cycle (with the buffer volume). The systematic uncertainties due to choice of normalization model were 0.09 s in 2017 and 0.05 s in 2018, which provided further evidence that the buffer volume improved the ability to estimate the number of UCN loaded into the trap.

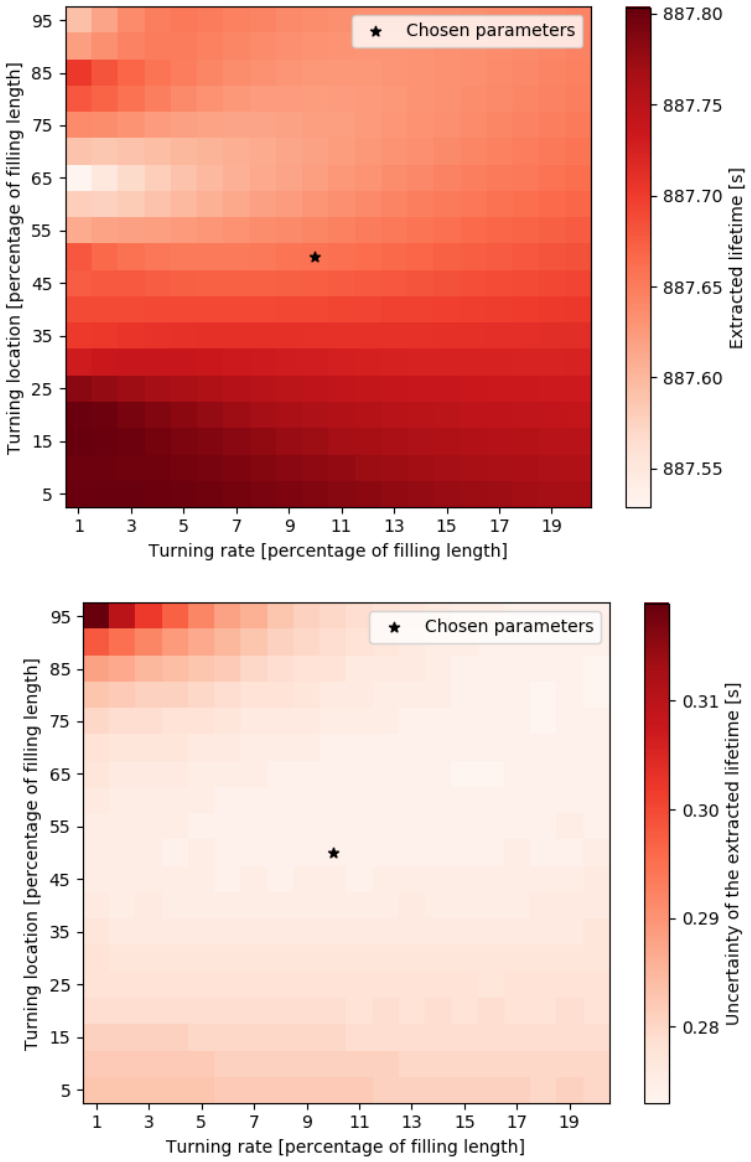

Figure 4.5: Effect of choice of normalization weighting function on the extracted lifetime. Top: extracted lifetimes. Bottom: uncertainty of those extracted lifetimes. In both plots, the black star represents the turning location and turning rate that were used in the analyses from Chapter 3. As the turning point was decreased, the extracted lifetime increased. Earlier turning points gave more weight to the counts recorded in the normalization monitors earlier in the filling process, even though these counts are mostly uncorrelated to the number of UCN in the trap at the end of the filling process. The uncertainty of the extracted lifetime increased as the turning rate decreased and the turning location increased. This is because most of the counts recorded in the normalization monitors during the filling process were de-weighted so the relative scale of statistical fluctuations of the number of counts recorded in the normalization monitors increased.

4.6 Depolarization of trapped UCN

The polarizing magnet and the adiabatic spin flipper from Section 2.5 were used to select only UCN with spin in the low-field-seeking (LFS) state so that UCN that entered the trap were repelled by the Halbach array. If any of the UCN in the trap depolarized, then they would have been attracted to and collided with the Halbach array, which would have resulted in a non-β-decay loss of UCN. UCN in long runs had a longer length of time, and therefore a higher probability, to depolarize than UCN in short runs. Therefore, a correction to the extracted lifetime due to the loss of depolarized UCN could have only increased the extracted lifetime.

If a UCN encountered a region with no magnetic field, then the UCN would have been free to depolarize from the LFS state. To prevent depolarization, the entire UCNτ trap was wrapped with coils to generate a magnetic holding field with magnitude $B_0 = 6.4$ mT. Theoretical calculations suggested that for magnetic fields ≥ 5 mT, the rate of depolarization of UCN in the trap was proportional to B^{-2} [56], where B was the magnitude of the magnetic holding field. This relationship was not exact, but was a useful estimate. When the B^{-2} relationship held true, the lifetime of UCN in the trap was

$$\frac{1}{\tau} = \frac{1}{\tau_n} + \left(\frac{B_0}{B}\right)^2 \frac{1}{\tau_{\text{depol}}}, \tag{4.6}$$

where τ_{depol} is the loss lifetime due to depolarization when $B = B_0$.

In order to estimate τ_{depol}, multiple low-statistics measurements of the lifetime due to all loss mechanisms were made in 2015 while varying the magnitude of the magnetic holding field between 6.4 mT and 6.4 μT. The statistical uncertainties of these lifetimes varied between 1 and 15 s. A previous publication from the UCNτ collaboration used these measurements to estimate a depolarization lifetime of $\tau_{\text{depol}} = \left(1.1^{+4.4}_{-0.5}\right) \times 10^7$ s [26]. Figure 4.6 shows the fit used to calculate this value. Although Equation 4.6 was designed to be used with $B \geq 5$ mT, all low-statistics measurements of the lifetime with $B > 0.5$ mT were used in the fit. The estimated depolarization lifetime would have caused the extracted lifetime to be 0.07 ± 0.06 s below the true neutron lifetime.

The estimated shift in the lifetime differed from 0 by 1.25σ (p-value = 0.11). No attempt was made during the 2017 and 2018 run cycles to directly observe depolarization of UCN in the trap, and there was only moderate statistical evidence that depolarization occurred. Furthermore, the fit of the data in Figure 4.6 extended well below the 5 mT cutoff developed in [56]. Therefore, no systematic correction

to the lifetime due to depolarization of UCN was reported. Instead, the 0.07 s was reported as a systematic uncertainty. As explained at the beginning of this section, any correction to the extracted lifetime due to the loss of depolarized UCN must be strictly non-negative. Therefore, a systematic uncertainty due to the loss of depolarized UCN of +0.07 s was reported.

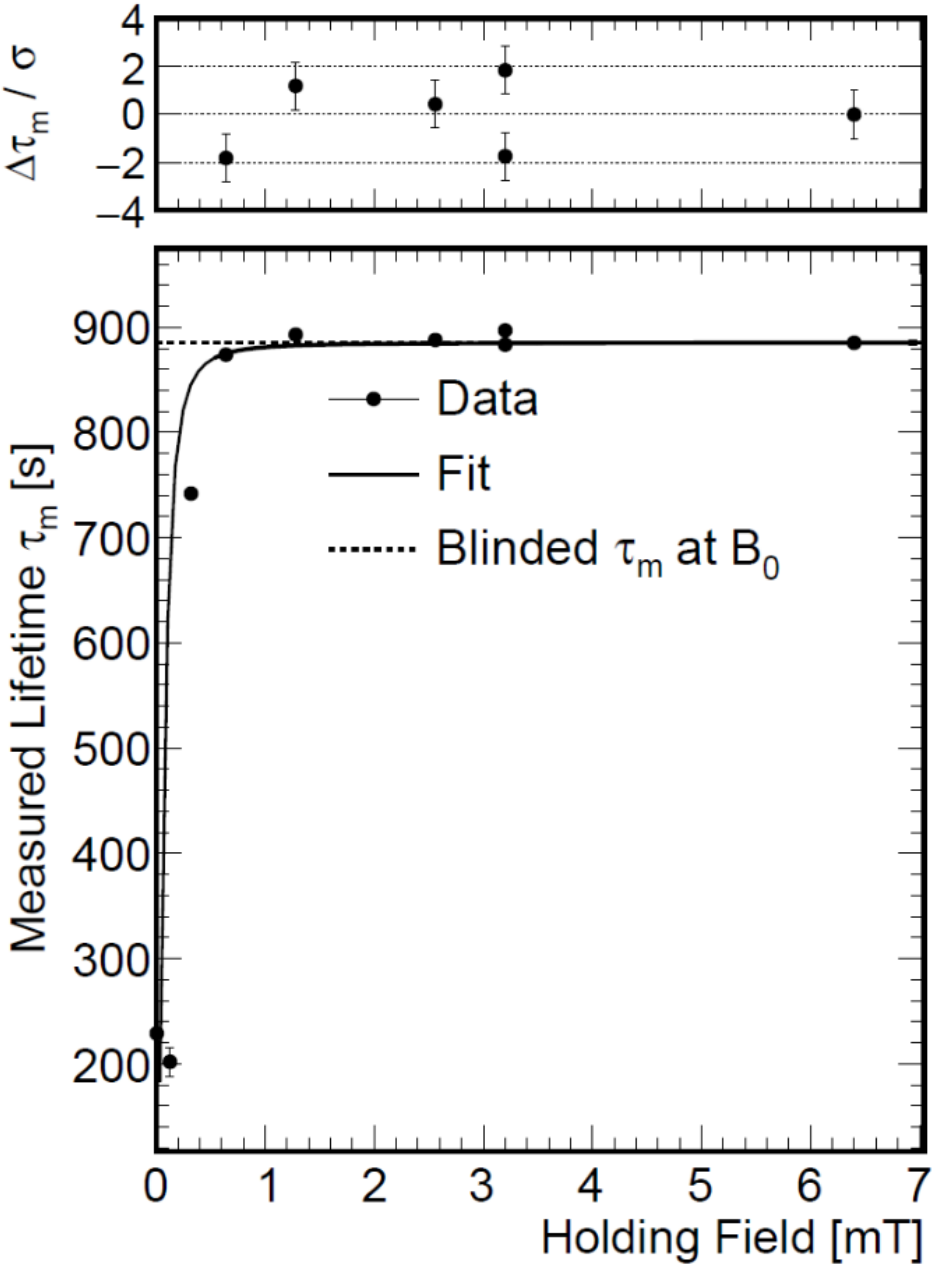

Figure 4.6: Measured lifetimes of UCN in the trap as a function of the magnetic holding field [26]. The fit shown used the model from Equation 4.6, but the parameters of the fit were extracted using only lifetimes measured with $B > 0.5$ mT.

4.7 Insufficient cleaning of quasi-bound UCN

Section 2.8 described how quasi-bound UCN were removed from the trap before the holding process begins. Monte Carlo simulations have shown that this process is not perfect [43]. The probability that a quasi-bound UCN was lost out of the trap increased as the length of the holding process increased. Therefore, a correction to the extracted lifetime due to insufficiently cleaned quasi-bound UCN that escaped from the trap could have only increased the extracted lifetime.

In a typical production run, the cleaners were lowered to 38 cm above the bottom of the trap during the filling and cleaning processes. Additionally, the primary detector was lowered to the same height as the cleaners during the filling process, but was raised to 43 cm above the bottom of the trap during the cleaning process. In order to study the systematic effect of insufficient cleaning of quasi-bound UCN, 30 runs were performed in 2017 and an additional 30 runs were performed in 2018 while the cleaners were held at their raised height of 43 cm above the bottom of the trap during the filling and cleaning processes. These runs were referred to as "no-cleaning" runs. The no-cleaning running configuration significantly increased the number of UCN with $E/m_n g > 38$ cm that were in the trap after the cleaning process.

A method similar to the global method of extracting a lifetime presented in Section 3.9 was used to separately analyze the data from each year's no-cleaning runs. Due to the small number of no-cleaning runs, the values for f were held fixed at the value from the rest of each year's data instead of calculating it based on the no-cleaning runs. In 2017, there were 10 free parameters used to fit the data: a lifetime $\kappa_{\text{no clean, 2017}}$, and nine normalization parameters (three each for three normalization subsets). In 2018, there were seven free parameters used to fit the data: a lifetime $\kappa_{\text{no clean, 2018}}$, and six normalization parameters (three each for two normalization subsets). In both years, the relatively low number of free parameters allowed for finding the maximum likelihood estimate of the free parameters by directly minimizing a version of Equation 3.14 that had been updated to use $\kappa_{\text{no clean}}$ instead of τ. These fits resulted in values of $\kappa_{\text{no clean, 2017}} = 848.8 \pm 2.8$ s and $\kappa_{\text{no clean, 2018}} = 862.3 \pm 3.7$ s. Figure 4.7 shows the results of these fits. The values for $\kappa_{\text{no clean}}$ differed significantly between 2017 and 2018 due to the instillation of the buffer volume pre-cleaner. In the 2018 no-cleaning runs, the two cleaners inside of the trap remained raised during the filling and cleaning processes, but the buffer volume pre-cleaner remained at its standard lowered height.

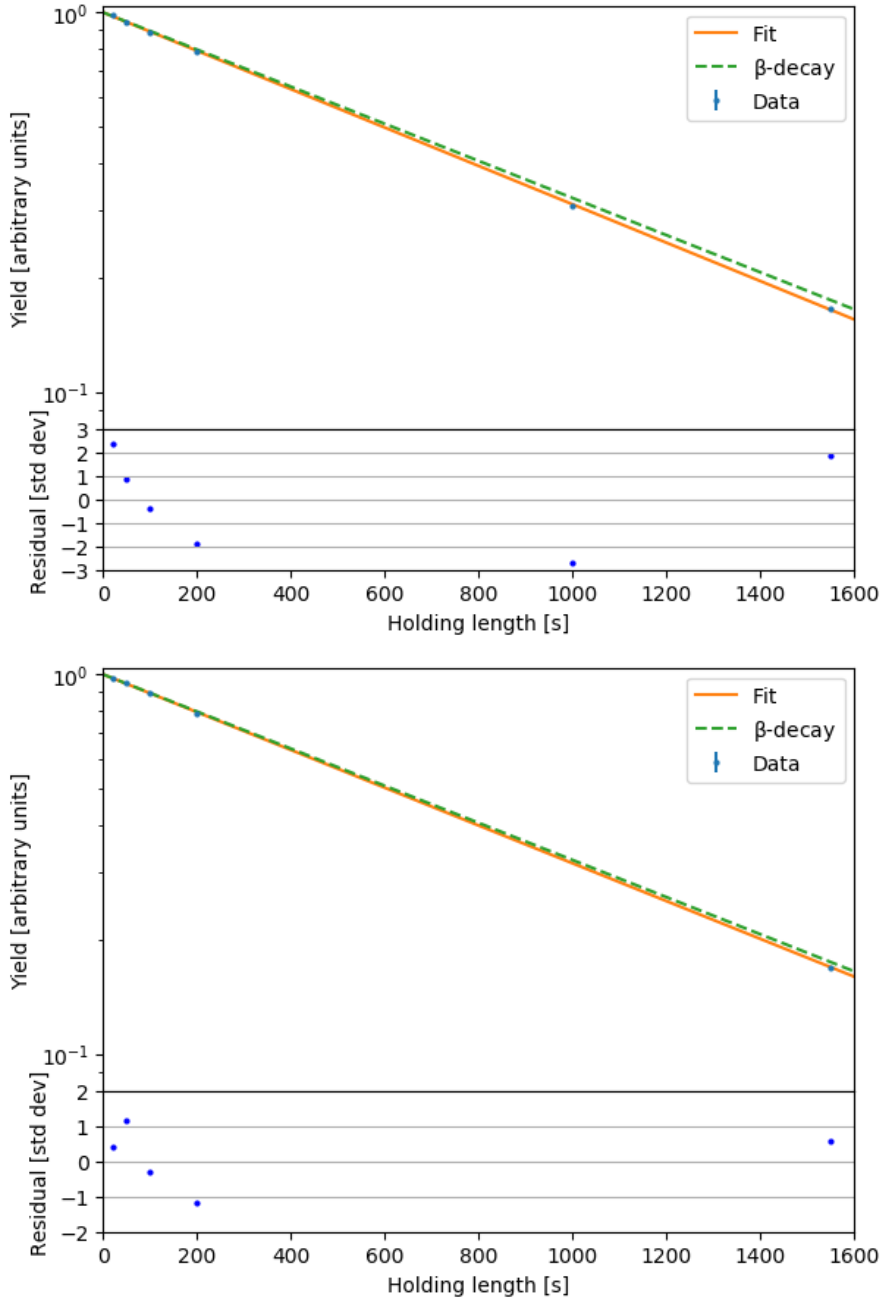

Figure 4.7: The results of fitting the no-cleaning data. Top: the 2017 no-cleaning data. Bottom: the 2018 no-cleaning data. In 2017, the residuals of the fit had $\chi^2/\nu = 42.1/20$ (p-value = 0.003). The scale of the uncertainty of the normalization estimate was held constant at the value calculated from the entirety of the 2017 production data, and therefore may not haven been a good representation of the uncertainty of the normalization estimate for these 30 runs. This could have contributed to the high value for for χ^2/ν. In 2018, the residuals of the fit had $\chi^2/\nu = 26.8/23$ (p-value = 0.26).

The values for the $\kappa_{\text{no clean}}$s differed from the neutron lifetime τ because some of the insufficiently cleaned quasi-bound UCN that were counted in the unloading process of short runs escaped from the trap during the holding process of long runs. The difference in the extracted lifetime between production runs and no-cleaning runs was roughly half as large in 2018 than in 2017 because the cleaning from the pre-cleaner in the buffer volume significantly reduced the number of UCN with $E/m_n g > 38$ cm that ever make it to the trap. The effect on the extracted lifetime from insufficiently cleaned quasi-bound UCN that escaped from trap can be analytically calculated by considering a short-long pair of runs with short holding length t_S, long holding length t_L, and $\Delta t \equiv t_L - t_S$. If:

1. exactly N trappable UCN were loaded into the trap,
2. exactly N_{IC} insufficiently cleaned quasi-bound UCN were loaded into the trap,
3. none of the insufficiently cleaned quasi-bound UCN escaped from the trap before the unloading process of the short run, and
4. all of the insufficiently cleaned quasi-bound UCN escaped from the trap before the unloading process of the long run,

then the extracted lifetime from that short-long pair would have been

$$\kappa = \frac{\Delta t}{\ln \frac{(N+N_{\text{IC}})e^{-t_S/\tau}}{Ne^{-t_L/\tau}}} = \frac{\Delta t}{\ln\left[\left(1+\frac{N_{\text{IC}}}{N}\right)e^{\Delta t/\tau}\right]}. \tag{4.7}$$

Equation 4.7 was rearranged as

$$r = \exp\left[\left(\kappa^{-1} - \tau^{-1}\right)\Delta t\right] - 1, \tag{4.8}$$

where $r \equiv \frac{N_{\text{IC}}}{N}$.

It was assumed that the number of insufficiently cleaned quasi-bound UCN that escaped from the trap between the short and long unloading processes was proportional to the Peak 1 yields for runs with a holding length of 20 s. Therefore,

$$r_{\text{clean}} = \frac{\left\langle Y^1_{\text{clean}}\right\rangle}{\left\langle Y^1_{\text{no clean}}\right\rangle} r_{\text{no clean}}, \tag{4.9}$$

where Y^1 was the Peak 1 yield.

Equations 4.7 through 4.9 were combined to derive

$$\kappa_{\text{clean}} = \frac{\Delta t}{\ln\left\{\left[1 + \frac{\langle Y^1_{\text{clean}}\rangle}{\langle Y^1_{\text{no clean}}\rangle}\left(\exp\left[\left(\kappa^{-1}_{\text{no clean}} - \tau^{-1}\right)\Delta t\right] - 1\right)\right] e^{\Delta t/\tau}\right\}}. \tag{4.10}$$

The difference between κ_{clean} and τ was the estimate of the shift in the extracted lifetime due the insufficiently cleaned quasi-bound UCN that escaped from the trap. Choice of Δt affected the calculated value for r, but the effect of choice of Δt on κ_{clean} was insignificant.

Table 4.1 shows the Peak 1 yields separated by year and by run type. The Peak 1 yields for no-cleaning runs were roughly 100× the Peak 1 yields of production runs. The only difference between production runs and no-cleaning runs was the height of the cleaners. Quasi-bound UCN that would have avoided the cleaners in production runs would still have avoided the cleaners in no-cleaning runs. UCN with non-stable trajectories and $E/m_n g > 38$ cm that would have been removed from the trap by the cleaners during a production run were able to remain in the trap until the unloading process of short no-cleaning runs. These UCN with non-stable trajectories were not a concern during production runs because they were rapidly removed by the cleaners, but the potential for the trajectories of insufficiently cleaned quasi-bound UCN to have evolved into non-stable trajectories and to have escaped from the trap was a concern.

Year	$\langle Y^1_{\text{clean}}\rangle$	$\langle Y^1_{\text{no clean}}\rangle$	$\langle Y^1_{\text{no clean}}\rangle / \langle Y^1_{\text{clean}}\rangle$
2017	$(5.8 \pm 1.6) \times 10^{-5}$	$(3.4 \pm 0.1) \times 10^{-3}$	57 ± 11
2018	$(-0.4 \pm 2.9) \times 10^{-5}$	$(15.2 \pm 0.6) \times 10^{-3}$	$284 \pm 153^{(*)}$

Table 4.1: Peak 1 yields of both production and no-cleaning short runs, separated by year. $^{(*)}$: The Peak 1 yield of production runs in 2018 was not significantly different than zero, so the value of $(-0.4 + 2 \times 2.9) \times 10^{-5}$ was used when calculating the ratio of the Peak 1 yields from 2018. $Y^1 = 10^{-4}$ corresponds to having counted $\sim 1 - 4$ UCN in Peak 1.

The average Peak 1 yield from production runs in 2018 was smaller than that of 2017. These yields were a measure of the number of quasi-bound UCN in the trap, and the pre-cleaning in the buffer volume significantly reduced the number of potential quasi-bound UCN that ever made it to the trap. In contrast, the average Peak 1 yield from no-cleaning runs was greater in 2018 than in 2017. As explained in the above paragraph, these yields were mostly a measure of the number of UCN with non-stable trajectories and $E/m_n g > 38$ cm in the trap. One factor that contributed to

this increase was the aluminum block discussed in Section 4.2. As shown in Figure 4.1, the aluminum block caused high-energy UCN to be lost from the trap, but did not significantly affect low-energy UCN. The no-cleaning runs from 2017 were done (unknowingly) while the aluminum block was in the trap, which decreased the Peak 1 yields from the no-cleaning runs from 2017. Upgrades to the primary detector also caused the Peak 1 yields from the no-cleaning runs from 2018 to be greater than in 2017. The layer of ^{10}B on the primary detector was twice as thick in 2018 as compared to 2017, which roughly doubles the efficiency per interaction of capturing low-velocity UCN. However, there was significant uncertainty of the thickness of ^{10}B on the primary detector from 2017, so this factor could have been even greater than two. In each version of the primary detector, the efficiency of capturing low-velocity UCN is thought to be $\leq 20\%$. UCN that are not captured are reflected and counted later in the unloading process, likely after the end of Peak 1. Additionally, the primary detector in 2017 was accidentally dropped 45 cm onto the Halbach array, which damaged the detector and decreased the efficiency of UCN capture on the lower part of the primary detector. Imperfections and impurities in the ^{10}B could have caused even more differences in the efficiency of UCN capture, but no data exist to support this hypothesis.

Given all of these differences (and potential differences) between the primary detectors of 2017 and 2018, a factor of 5 difference in the calculated value of $\langle Y^1_{\text{no clean}} \rangle / \langle Y^1_{\text{clean}} \rangle$ was not unreasonably large. What mattered for estimating the uncertainty of the extracted lifetime due to insufficient cleaning of quasi-bound UCN was not the number of UCN that were counted at Peak 1, but the actual number of UCN that had quasi-bound trajectories. The addition of the pre-cleaner in the buffer volume significantly decreased shifts in the lifetime due to not cleaning the population of UCN in the trap. If all else was held equal between the two years, then this should have resulted in a smaller value of $\langle Y^1_{\text{no clean}} \rangle / \langle Y^1_{\text{clean}} \rangle$ in 2018 than in 2017, but the opposite was observed. Several reasons have been presented to explain this discrepancy. The measurement of this ratio in 2018 was made using upgraded versions of multiple pieces of hardware, and theoretically the ratio should have been greater in 2017 than in 2018, so the measured ratio for 2018 was used to estimate the uncertainty of the extracted lifetime due to insufficient cleaning of quasi-bound UCN in both 2017 and 2018.

Using $\left\langle Y^1_{\text{no clean}} \right\rangle / \left\langle Y^1_{\text{clean}} \right\rangle = 284 \pm 153$ resulted in extracted lifetimes of $\kappa_{\text{clean, 2017}} = 877.50 \pm 0.34$ s and $\kappa_{\text{clean, 2018}} = 877.57 \pm 0.46$ s. These extracted lifetimes corresponded to decreases in the extracted lifetime of $\tau_{2017} - \kappa_{\text{clean,2017}} = 0.12 \pm 0.03$ s and $\tau_{2018} - \kappa_{\text{clean,2018}} = 0.06 \pm 0.03$ s. A lifetime extracted from just the 2017 production data had a statistical uncertainty of 0.34 s. A lifetime extracted from just the 2018 production data had a statistical uncertainty of 0.46 s. An average of the decreases in extracted lifetimes that was weighted by the statistical variances of the lifetimes extracted from the production data resulted in a combined decrease of $\tau - \kappa_{\text{clean}} = 0.10 \pm 0.03$ s.

The analysis presented in this section estimated that insufficiently cleaned quasi-bound UCN that escaped from the trap caused the extracted lifetime to be 0.10 ± 0.03 s less than it otherwise would be. If this 0.10 ± 0.03 s decrease was correct, then the extracted lifetime should have been increased by 0.10 ± 0.03 s to correct for it. This correction was estimated using multiple assumptions that were difficult to verify. There was too much uncertainty to justify applying a 0.10 s correction to the extracted lifetime, so 0.10 s was reported as a systematic uncertainty instead of a systematic correction. As explained at the beginning of this section, any correction to the extracted lifetime due to insufficiently cleaned quasi-bound UCN that escaped from the trap must have been strictly non-negative. Therefore, the systematic uncertainty of the extracted lifetime due to insufficient cleaning of quasi-bound UCN was reported as +0.10 s.

4.8 Vibrational heating of UCN

The energy of UCN in the trap could have increased due to mechanical vibrations of the trap, which could have resulted in vibrationally heated UCN that escaped from the trap. Each interaction of a UCN with the vibrating trap could have either increased or decreased the energy of the UCN. Decreases in energy could have shifted the mean arrival time of UCN during the unloading process. This effect was discussed in Section 3.8.2 and was not a significant concern. However, increases in energy could have led to vibrationally heated UCN that escaped from the trap. UCN in long runs had more opportunities than UCN in short runs to accumulate energy through interacting with a vibrating trap, so the probability that a vibrationally heated UCN escaped from the trap increased as the length of the holding process increased. Therefore, the correction to the extracted lifetime due vibrationally heated UCN that escaped from the trap could have only increased the extracted lifetime.

An accelerometer was used to measure the vibrations of the trap. The acceleration data were built into a Monte Carlo simulation that simulated trajectories of UCN in the trap [43]. A small number of vibrationally heated simulated UCN acquired sufficient energy to reach the cleaners at their raised height of 43 cm above the bottom of the trap (which is 5 cm above the lowered height of the cleaners) and were erroneously cleaned out of the trap. This resulted in an extracted lifetime that was 0.03 s less than the true Monte Carlo lifetime. When the amplitudes of the simulated vibrations were increased to 40× the measured amplitudes, the extracted lifetime was 0.15 s less than the true Monte Carlo lifetime.

This Monte Carlo simulation had three flaws:

1. The simulation modeled the magnetic fields from the Halbach array using the analytical expression from Equation 2.5, which approximated the Halbach array as an infinite two-dimensional plane of permanent magnets;

2. The permanent magnets in the trap had known defects which caused the actual magnetic field of the trap to have further differed from the model, which will be further discussed in Section 5.6; and

3. UCN with $E/m_n g \simeq 38$ cm bounced in the trap with frequencies $\gtrsim 2$ Hz, but the accelerometer data were gathered at higher frequencies.

The first two flaws both led to a magnetic field model that had fewer imperfections than the true magnetic field in the trap. These imperfections disrupted the periodic trajectories of quasi-bound UCN and made them more quickly leave the trap, which decreased the number of quasi-bound UCN that remained in the trap after the cleaning process. The Monte Carlo simulation did not have these features, so the fraction of simulated UCN that populated quasi-bound trajectories was greater in the simulation than in the actual experiment. The third flaw could have been problematic if the amplitude of vibrations with frequencies near 2 Hz was significant. In that case, those vibrational modes could have contributed more to vibrational heating than the vibrations measured with the accelerometer.

A data-driven calculation of the systematic effect of vibrational heating was developed to compare to the Monte Carlo estimate of 0.03 s. Peak 1 counts were again used to constrain this systematic effect. A count in Peak 1 of a long run could have come from one of three sources:

1. A UCN with 38 cm $< E/m_n g <$ 43 cm that should have been, but was not, removed from the trap by the cleaners during the cleaning process. $E/m_n g >$ 38 cm was a sufficient amount of energy that the UCN could have bounced high enough to reach the primary detector in Peak 1. $E/m_n g <$ 43 cm was insufficient energy for the UCN to have escaped from the trap during the holding process.

2. A UCN that started the run with $E/m_n g <$ 38 cm that then gained sufficient energy during the long holding process due to vibrational heating that it could have reached the primary detector in Peak 1.

3. A background event.

Peak 1 yields were calculated after subtracting the estimated background from the measured number of counts. Therefore, on average, only UCN contributed to the Peak 1 yields. Some of these UCN began the long holding process with enough energy to be part of Peak 1, and others accumulated enough additional energy from vibrational heating during the holding process. Two assumptions were made to estimate the systematic effect of vibrational heating of UCN on the extracted lifetime:

1. All UCN counted in Peak 1 of long runs were UCN that were vibrationally heated from Peak 2 to Peak 1; and

2. Vibrational heating did not cause any UCN to escape from the trap during the holding process.

The first assumption was a worst-case assumption about the counts recorded in Peak 1 of long runs, but the second assumption was a best-case assumption about the rate at which vibrationally heated UCN escape from the trap. After these assumptions were made, a lifetime that included the effect of vibrational heating of UCN was calculated using the methodology in Section 3.9, but where values for the long runs were calculated using data from Peaks 1 through 3 instead of just Peaks 2 and 3. This lifetime was $\tau_{\text{heated}} = 877.67 \pm 0.27$ s, which corresponded to an increase in the extracted lifetime of 0.05 ± 0.02 s.

This increase was calculated using one worst-case assumption and one best-case assumption. Monte Carlo simulations had found that the best-case assumption was not true, and that the vibrationally heated UCN that did escape from the trap increased the extracted lifetime by 0.03 s. A conservative 0.08 ± 0.02 s estimate of the correction to the extracted lifetime due to vibrational heating of UCN was established by summing the two estimates. This was likely an overestimate of the correction because it was made using a worst-case assumption, so 0.08 s was reported as a systematic uncertainty instead of as a systematic correction. As explained at the beginning of this section, any correction to the extracted lifetime due to vibrationally heated UCN that escape from the trap must have been strictly non-negative. Therefore, the systematic uncertainty of the extracted lifetime due to vibrational heating of UCN was reported as +0.08 s.

4.9 Summary of systematic effects

The systematic corrections from Sections 4.2 to 4.8 were collected in Table 4.2. The non-zero systematic corrections due to interactions between UCN and the aluminum block and residual gasses were calculated by applying corrections on a run-by-run basis before extracting a lifetime. The additional five systematic uncertainties were calculated separately from the process of extracting a lifetime, and then the uncertainties were added in quadrature to the other systematic uncertainties. The systematic corrections to the extracted lifetime were less than the statistical uncertainty of 0.27, and the systematic uncertainties did not significantly contribute to the total uncertainty. After applying all systematic corrections, the extracted lifetime increased to 877.79 ± 0.27 (stat.) $^{+0.19}_{-0.12}$ (sys.) s. Summing the statistical and systematic uncertainties of quadrature resulted in a total uncertainty of $^{+0.33}_{-0.30}$ s.

Source of systematic effect	Correction [s]	Uncertainty [s]
Aluminum block interactions	+0.06	±0.03
Residual gas interactions	+0.11	±0.06
Defining a reconstructed UCN event	0	±0.08
Choice of normalization weighting function	0	±0.06
Depolarization of trapped UCN	0	+0.07
Insufficient cleaning of quasi-bound UCN	0	+0.10
Vibrational heating of UCN	0	+0.08
Total [uncorrelated sum]	0.17	$^{+0.19}_{-0.12}$

Table 4.2: Systematic corrections to the extracted lifetime.

4.10 Implications of this measurement

This thesis presented 877.79 ± 0.27 (stat.) $^{+0.19}_{-0.12}$ (sys.) s as a measurement of the neutron lifetime. This value is consistent with values for the neutron lifetimes of that were previously published by the UCNτ collaboration [24, 26], but is 1.61 s less than the current global average for the neutron lifetime [1]. When the uncertainty of this measurement is combined in quadrature with the uncertainty of the global average, the difference is 2.35σ (p-value = 0.019).

Figures 1.4 and 1.5 were updated to include the measurement presented in this thesis, and these updates are presented in Figures 4.8 and 4.9, respectively. Although the difference between the value for the lifetime presented in this thesis and the current global value for the lifetime has not reached the typical 5σ significance level needed to declare a discovery in particle physics, it does hint that the current global average for the neutron lifetime is above the true value for the neutron lifetime. Given the constraints of Equation 1.5, the lifetime presented in this thesis provides moderate evidence that the value for $|V_{ud}|^2$ calculated as $1 - |V_{us}|^2 - |V_{ub}|^2$ is more consistent with the Standard Model than the value for $|V_{ud}|^2$ determined by studying superallowed $0^+ \rightarrow 0^+$ nuclear β-decays. However, more precise measurements of λ are needed in order to more clearly discern which value for $|V_{ud}|^2$ is consistent with the Standard Model. This improved measurement of the neutron lifetime using a “bottle” technique is vastly different from past measurements that used the ”beam” technique, so this measurement does not help resolve the discrepancy discussed in Section 1.4.3. Therefore, it also did not help clarify what value for the neutron lifetime should be used in Big Bang nucleosynthesis calculations.

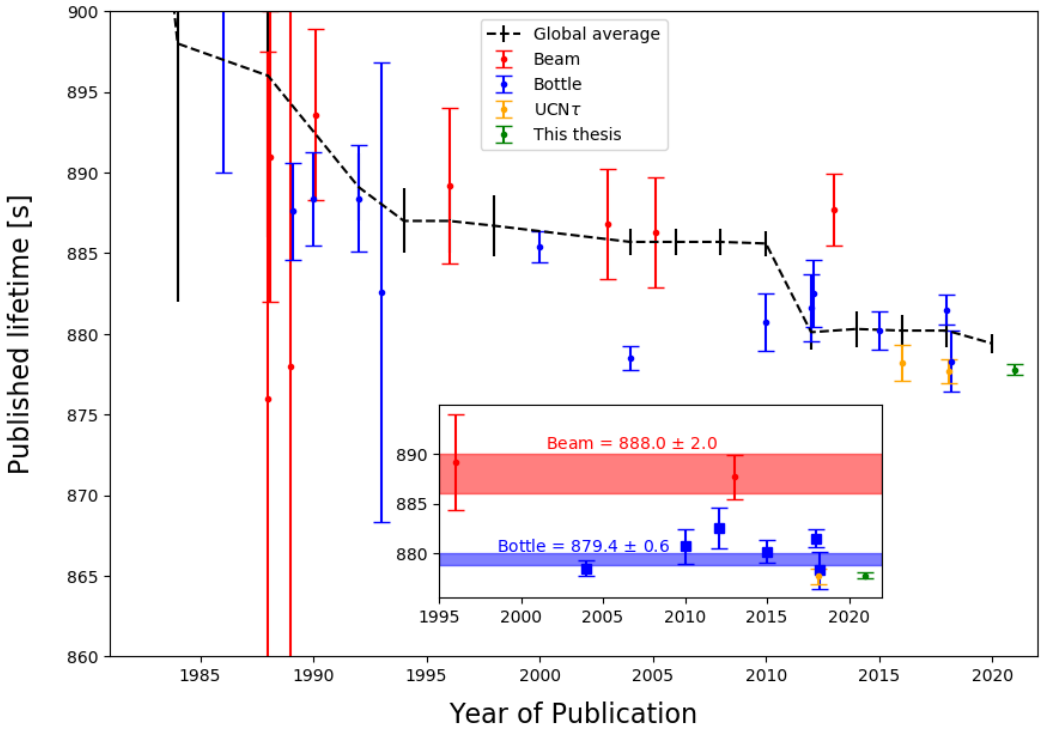

Figure 4.8: History of measurements of the neutron lifetime [1, 5–27]. Inset: modern measurements of the neutron lifetime, as well as the current global averages for lifetimes measured using the beam and bottle methods. Over time, the global average value has steadily decreased. In 2012, the global average shifted from combining multiple beam measurements to combining multiple bottle measurements. Since then, multiple measurements of the lifetime by bottle experiments, including those published by the UCNτ collaboration, have been in good agreement with the global average. However, this analysis of the 2017-2018 UCNτ data set resulted in an uncertainty that was small enough to lead to a 2.35σ difference between the current global average and the lifetime presented in this analysis.

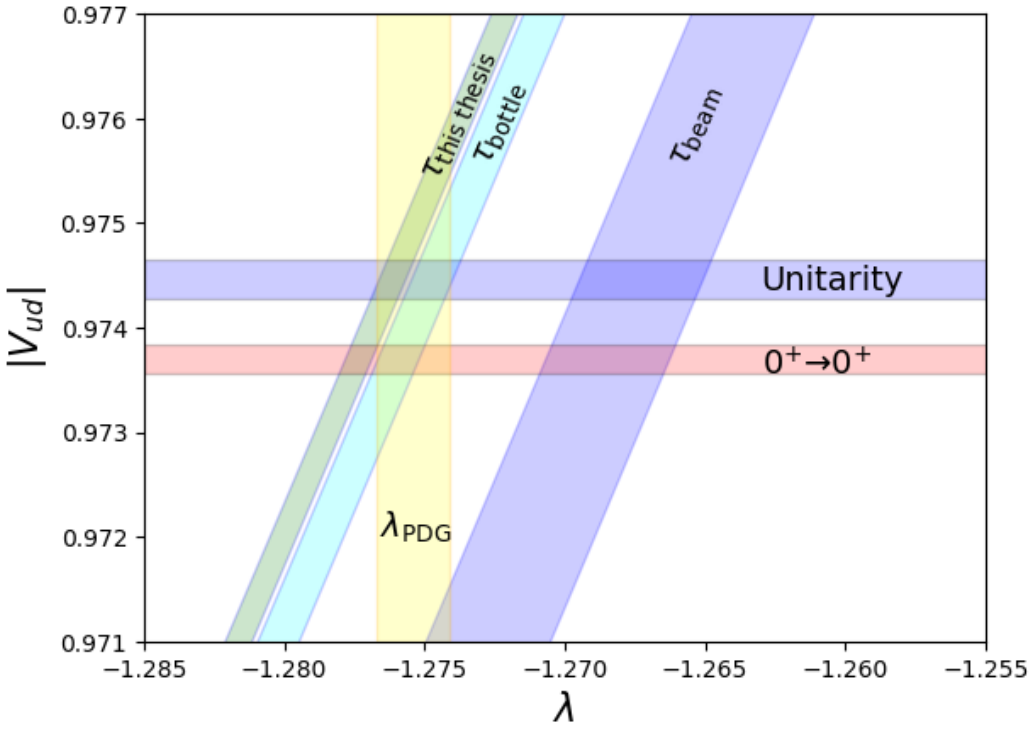

Figure 4.9: The relationship between the current global averages for measurements of τ, λ, and certain components of the CKM matrix [1], as well as the measurement of τ presented in this analysis. Each band denotes 1σ uncertainties of the measured quantity. The "Unitarity" band was determined by using the current global averages for $|V_{us}|$ and $|V_{ub}|$ and demanding that $1 = |V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2$. The "$0^+ \to 0^+$" band is the current global average of measurements of $|V_{ud}|$ made by studying superallowed $0^+ \to 0^+$ nuclear β-decays. τ_{bottle}, τ_{beam}, and $\tau_{\text{this thesis}}$ were calculated using Equation 1.7.

Chapter 5

THE FUTURE OF UCNτ

5.1 Introduction

A previous publication from the UCNτ collaboration reported an extracted lifetime of 877.7 ± 0.7 (stat.) $^{+0.4}_{-0.2}$ (sys.) s. Summing these two uncertainties of quadrature resulted in a total uncertainty of $^{+0.81}_{-0.73}$ s. The weighted average of that lifetime and the lifetime extracted in this analysis was $877.78^{+0.31}_{-0.28}$ s. The 2019 and 2020 run cycles did not gather a significant amount of production data, but did gather additional data to try to improve the understanding of systematic effects. The plan for the 2021 and 2022 run cycles is to gather more production data to reduce the statistical uncertainty, and to test and develop new hardware. After the 2022 run cycle, it is expected that the UCNτ will cease running and efforts will be redirected to other UCN experiments and to building the next-generation version of UCNτ.

It is expected that after all data from the UCNτ experiment is analyzed, a combined lifetime with a total uncertainty of between 0.20 and 0.25 s will be extracted. The next-generation version of the UCNτ experiment must have significant upgrades in order to gather sufficient data to measure the lifetime with an uncertainty of $\leq$ 0.1 s in a reasonable amount of time. Section 5.2 will provide an overview of a proposed next-generation UCN source. Section 5.3 will discuss a new method of using UCN that do not make it all the way from the UCN source to the trap to improve the normalization estimates. Section 5.4 will discuss a proposal for how to increase the efficiency of the transport of UCN from the UCN source into the trap. Section 5.5 will discuss potential upgrades to the primary detector to accommodate increased UCN counting rates. Sections 5.6 and 5.7 will provide an overview of how some of the systematic uncertainties can be better constrained.

5.2 Increasing the production of UCN

Table 3.7 decomposed the 0.27 s of statistical uncertainty of the extracted lifetime into six components. Five of these six components would decrease if the overall number of UCN that were loaded into the trap was increased. The sixth component, primary detector rate-dependent effects, would increase with an increase in the number of UCN loaded into the trap. Section 5.5 will discuss how a next-generation version of the primary detector will suppress rate-dependent effects. If the statistical

uncertainty due to primary detector rate-dependent effects is suppressed, then a 7× increase in the size of the data set would be sufficient to decrease the statistical uncertainty of the extracted lifetime to 0.10 s. If the systematic uncertainty of a future extraction of the lifetime was 0.05 s, then a 10× increase in the size of the data set would be required in order to achieve a total uncertainty of 0.10 s on the extracted lifetime. One way to quickly increase the size of the data set by 10× would be to increase the number of UCN produced by the UCN source.

A model for a proposed next-generation UCN source was developed [57]. In Monte Carlo N-Particle Transport Code (MCNP) [58] simulations, this source increased UCN production by 1000× compared to the current LANSCE UCN source. In contrast to the source discussed in Section 2.3, the proposed UCN source would use superfluid liquid Helium (He-II) to downscatter cold neutrons to UCN energy levels. Figure 5.1 shows a model of the proposed source that was tested in MCNP simulations. If this proposed next-generation UCN source can be built and the 1000× increase in UCN production can be achieved, then the measurement of the neutron lifetime to within a total uncertainty of 0.10 s would certainly be achieved.

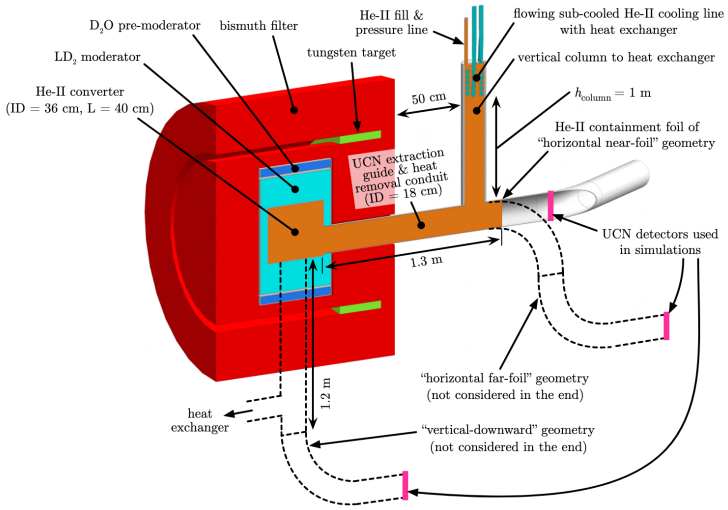

Figure 5.1: A model of the proposed next-generation UCN source [57]. Three extraction geometries were studied and the "horizontal near-foil" geometry was chosen in order to maximize the production of usable UCN while minimizing engineering constraints. In contrast to the current LANSCE UCN source that uses a solid cylinder of W, the proposed next-generation UCN source has a hollow cylinder of W that encircles the He-II source. A beam raster system moves the proton beam around this cylinder which reduces the local energy deposition. A location for this proposed next-generation UCN source has not been chosen, but this source would be compatible with the LANSCE 800 MeV proton beam.

5.3 Using the buffer volume as a normalization monitor

Section 3.6.2 discussed how the normalization monitors were used to estimate the number of UCN loaded into the trap. The rates in the normalization monitors were integrated during the filling process, and those integrals were used as inputs to a normalization model. The counts in the normalization monitors in the first $\sim$ half of the filling process had very low correlation with the number of UCN loaded into the trap, but nonetheless those counts were used as part of the normalization estimate.

Any measure of a UCN density that was proportional to the density of UCN in the trap could be used as a normalization monitor. The buffer volume and the trap were connected to each other via UCN guides during the filling process. After a sufficiently long filling process such that the UCN densities in the buffer volume and the trap had reached equilibrium, the densities in the two vessels should have been proportional to each other. The buffer volume and the trap stored UCN of similar energies so the constant of proportionality between the two densities should have remained fairly constant as the condition of the UCN source changed. Therefore, the number of UCN in the buffer volume at the end of a filling process should have been proportional to the number of UCN in the trap at the end of a filling process.

During a typical run from the 2020 run cycle, there were $\sim 10^6$ UCN in the buffer volume at the end of the filling process. In contrast, there were only $\sim 10^4$ UCN in the trap at the end of a filling process. The UCN in the buffer volume did not end up in the trap so they could not be directly used to measure the lifetime. The volumes of the buffer volume and the trap were comparable to each other, so the fact that the buffer volume had $\sim 100\times$ as many UCN as the trap suggested that transport of UCN from the buffer volume to the trap was very inefficient. Section 5.4 will discuss how to improve the transport of UCN from the buffer volume to the trap.

The number of UCN in the buffer volume after the filling process was $\sim 2-5\times$ greater than the number of counts recorded in the upper monitors during the filling process, and was $15-30\times$ greater than the number of counts recorded in the lower monitors during the filling process. An increase in the number of counts of $\sim 2-5\times$ may seem marginal, but the two numbers could not be directly compared to each other. First, the counts in the upper monitors were only from UCN that had too much energy to be stored in the trap, so they did not directly measure the population of UCN that are of interest. Second, many of the counts in both the upper and lower monitors were recorded before the UCN densities in the guides reached equilibrium, so many of the counts were not highly correlated with the number of UCN that were

loaded into the trap. The use of the UCN remaining in the buffer volume at the end of the filling process as a normalization monitor should significantly reduce the uncertainty of the normalization estimates.

Neither of the two normalization monitors installed in the buffer volume during the 2018 run cycle could have counted all of the UCN in the buffer volume. The pre-cleaner could have only counted high-energy UCN. The density monitor had a very small active surface and counted UCN much more slowly than the rate of loss of UCN from interactions with the walls of the buffer volume. In order to quickly count the remaining UCN in the buffer volume, a "dump" monitor with a large active surface was installed on a port at the bottom of the buffer volume. Gate valves were installed on three ports of the buffer volume:

1. The port where UCN entered the buffer volume from the UCN source;
2. The port where UCN exited the buffer volume as they traveled towards the trap; and
3. The port where the dump monitor was mounted.

During the filling process, the first two gate valves were open, but the third was closed. Immediately after the filling process ended, all gate valves were closed. Shortly after this, the gate valve to the dump monitor was opened, and all UCN that remained in the buffer volume were counted. The purpose of the delay between the end of the filling process and the opening of the gate valve to the dump monitor was to allow for most of the high-energy UCN in the buffer volume to be lost or captured by the pre-cleaner so that they did not contaminate the counts in the dump monitor. Figure 5.2 shows the counting rates for the three monitors in the buffer volume.

Equation 3.7 modeled the uncertainty of an estimate of the normalization N as $\sigma_N = \sqrt{fN}$, where $f \geq 1$ is a constant that quantified the uncertainty of the estimates of the number of UCN loaded into the trap. Therefore, the fractional uncertainty of the normalization estimate was $\sigma_N/N = \sqrt{f/N}$. The ideal way to decrease σ_N/N would be to increase the number of UCN loaded into the trap while f remained constant. However, N was limited by the ability of the source to produce UCN. If f was decreased while N was held constant, then σ_N/N would have also decreased.

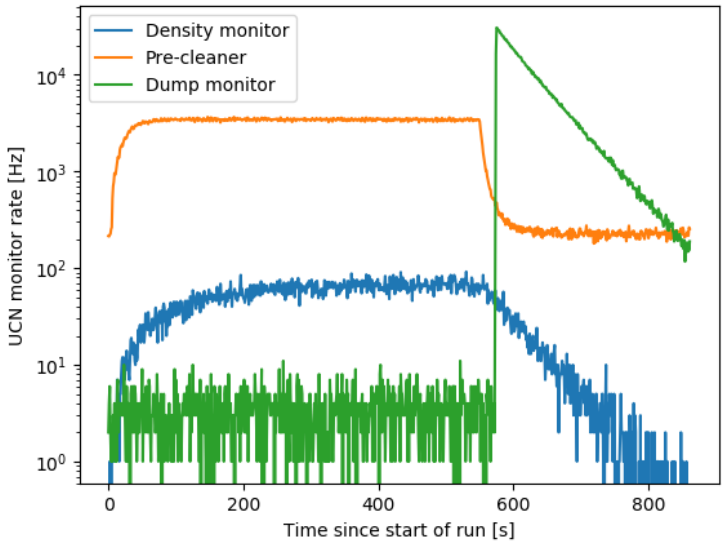

Figure 5.2: UCN counting rates in the three monitors of the buffer volume. For this test, the filling length was extended from the standard 300 s to $\simeq$ 550 s to ensure that the buffer volume was completely saturated with UCN.

This section explained why using the UCN that remained in the buffer volume after the filling process as a normalization monitor would be a significant improvement over previous normalization monitors. This improvement in the relationship between the counts in the normalization monitor and the number of UCN loaded into the trap should decrease f and σ_N/N. The test shown in Figure 5.2 that was done during the 2020 run cycle showed significant promise, and this method of normalization should significantly reduce the values for f during the 2021 and 2022 run cycles.

5.4 Improving transport of UCN from the buffer volume to the trap

Section 5.3 stated that the number of UCN that remain in the buffer volume after the filling process was $\sim 100\times$ greater than the number of UCN loaded into the trap. A new method has been realized for using these UCN to decrease the uncertainty of the estimates of the number of UCN loaded into the trap. Table 3.7 decomposed the 0.27 s of statistical uncertainty of the lifetime extracted from the 2017/2018 data set into various components. The component of the statistical uncertainty due to additional normalization uncertainty over $\sigma_N = \sqrt{N}$ contributed 0.10 s of statistical uncertainty to the total statistical uncertainty. Even if the improved normalization method presented in Section 5.3 completely suppressed this 0.10 s of statistical uncertainty, the total statistical uncertainty of the lifetime extracted from the 2017/2018 data set would only decrease from 0.27 s to $\sqrt{(0.27\text{ s})^2 - (0.10\text{ s})^2} \simeq 0.25$ s.

A better way to use these additional $\sim 100\times$ UCN would be to load them into the trap and use them to directly measure the lifetime. As discussed in Section 5.2, a $\sim 10\times$ increase in the number of UCN loaded into the trap could enable a measurement of the neutron lifetime with a total uncertainty of $\lesssim 0.10$ s. Therefore, a method of transport of UCN from the buffer volume to the trap that is even $\sim 10\%$ efficient could enable a 0.10 s measurement of the lifetime.

A new method of transporting UCN from the buffer volume to the trap has been proposed. In this proposed method, the buffer volume is redesigned so that it is able to be detached from all guides and lowered into the trap. Once lowered into the trap, the bottom of the buffer volume would open and the UCN stored inside of it would be quickly released into the trap. The buffer volume would then retract from the trap and the cleaning process would begin. Initial Monte Carlo simulations have shown that lowering the buffer volume at reasonable speeds does not significantly increase the velocity of UCN in the buffer volume, and can even decrease the velocity of some of the UCN due to interactions with receding walls. This ability to decrease the velocity of some of the UCN in the buffer volume would allow for a greater fraction of the UCN produced by the UCN source to be stored in the trap. Figure 5.3 shows a CAD drawing of this proposed mechanism.

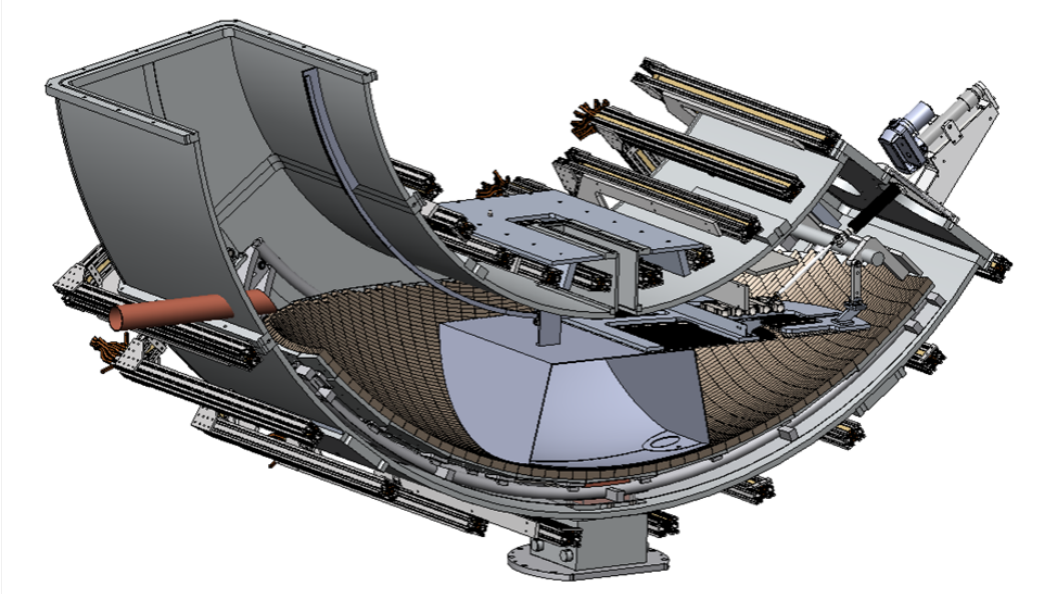

Figure 5.3: A CAD drawing of a proposed buffer volume mechanism to lower UCN into the trap. The red UCN guide (left) connects the redesigned buffer volume (grey, in the middle of the trap) to the UCN source when the buffer volume is in its raised height. After filling with UCN, the buffer volume lowers into the trap and releases all of the UCN that are inside of it.

5.5 Next-generation primary detector

Section 5.2 discussed how the production of UCN can be increased and Section 5.4 discussed how to more efficiently transport UCN from the buffer volume to the trap. A 10× increase in the number of UCN loaded into the trap could enable gathering sufficient data over two run cycles in order to measure the neutron lifetime with a statistical uncertainty $\lesssim 0.10$ s. An increase in the number of UCN stored in the trap helps with almost every aspect of this experiment. However, a 10× increase in the number of UCN loaded into the trap would also increase the maximum rate at which UCN are captured by the primary detector and would significantly increase the difficulty of reconstructing UCN events (see Section 3.4) because the number of UCN-to-UCN interactions would increase roughly 10×. The systematic uncertainty due to how to define a coincidence chain (CC) in the primary detector (see Section 4.4) would contribute significantly to the overall uncertainty of an extracted lifetime.

One strategy to circumvent this issue would be to extract a lifetime using raw photoelectron (PE) counts instead of reconstructed UCN events. For the 2017/2018 data set, the inability to accurately estimate the background PE rates led to shifts in the extracted lifetime of $\simeq 0.5$ s. If these background PE rates remained constant, but the number of foreground PEs increased by 10×, then these shifts would decrease to $\simeq 0.05$ s, which is well within the targeted 0.10 s total uncertainty.

CCs could still be used to extract a lifetime with a total uncertainty of 0.10 s if the number of UCN-to-UCN interactions could be suppressed. Figure 3.8 shows the distribution for times that PEs are observed during a UCN event in the primary detector. The long tail of this distribution was the cause of the significant number of UCN-to-UCN interactions that need to be corrected for using the methods developed in Sections 3.4.5 and 3.4.6. The primary detectors used in 2017 and 2018 used ZnS:Ag as the scintillator that eventually generated PEs. A next-generation primary detector could use $Lu_{2(1-x)}Y_{2x}SiO_5$ doped with Ce (LYSO:Ce) as a scintillator to generate PEs [59]. Figure 5.4 shows that the tail of the distribution for PEs is significantly suppressed when LYSO:Ce is used as the scintillator in the primary detector. In the prototype LYSO:Ce primary detector, the number of PEs generated per UCN was significantly less than in the ZnS:Ag primary detectors used in the past. Development work is ongoing to increase the number of PEs generated per UCN captured by a LYSO:Ce primary detector.

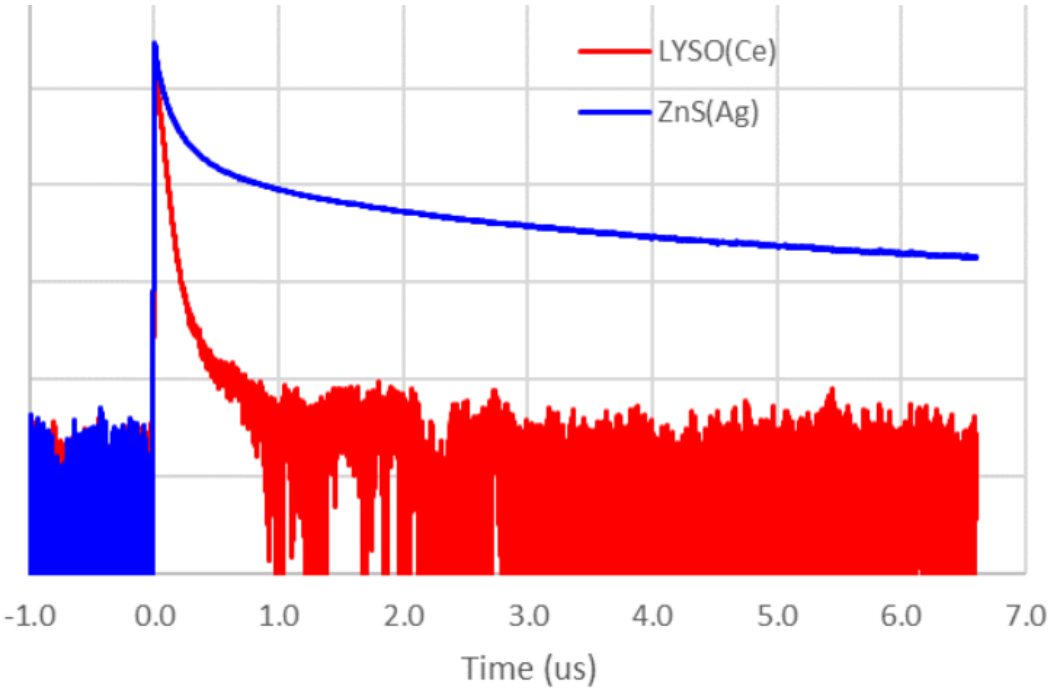

Figure 5.4: Distributions for times that PEs are observed during a UCN event for primary detectors built with two different scintillating materials [59]. The vertical axis is measured in arbitrary units that are proportional to the PE rate. The time distribution of PEs observed in a prototype of a primary detector that uses LYSO:Ce as the scintillator had a much shorter tail, but also detected far fewer PEs. The PE rate in the primary detector with LYSO:Ce relaxed to PE rates near the background PE rate within 1 μs, which would significantly reduce the number of UCN-to-UCN interactions in the primary detector, even if the maximum rate at which UCN are counted was significantly increased.

Another feature that could be implemented in a next-generation version of the primary detector is pixelation of its active surface. Each pixel would contain interleaved wavelength-shifting fibers (WLSF) in a manner similar to the layout of the current primary detector. Separating the active surface of the primary detector into 10 pixels would (roughly) keep the number of UCN-to-UCN interactions constant, even if the number of UCN in the trap was increased 10× and ZnS:Ag was still used as the scintillator in the primary detector. One problem with this proposal is that the number of photomultiplier tubs (PMTs) used to detect the PEs increases with the number of pixels, and it would be a difficult engineering task to mount all of these PMTs on the primary detector actuator. No work has been done on developing this feature.

5.6 High-precision measurements of magnetic fields in the trap

Monte Carlo simulations of UCN trajectories in the trap [43] are dependent on models for the magnetic fields inside of the trap. Two flaws in the existing models were mentioned in Section 4.8. Measurements of the actual magnetic fields in the trap are ongoing. The purposes of these measurements are to improve the models of the magnetic fields inside of the trap, and to search for regions of zero magnetic field where UCN could depolarize and be lost from the trap [60]. After removing the primary detector from the trap, a mapping arm with a Hall probe on the end of the arm was inserted as shown in Figure 5.5. The mapping arm is automatically moved around the trap to make measurements of the magnetic field throughout the volume of the trap.

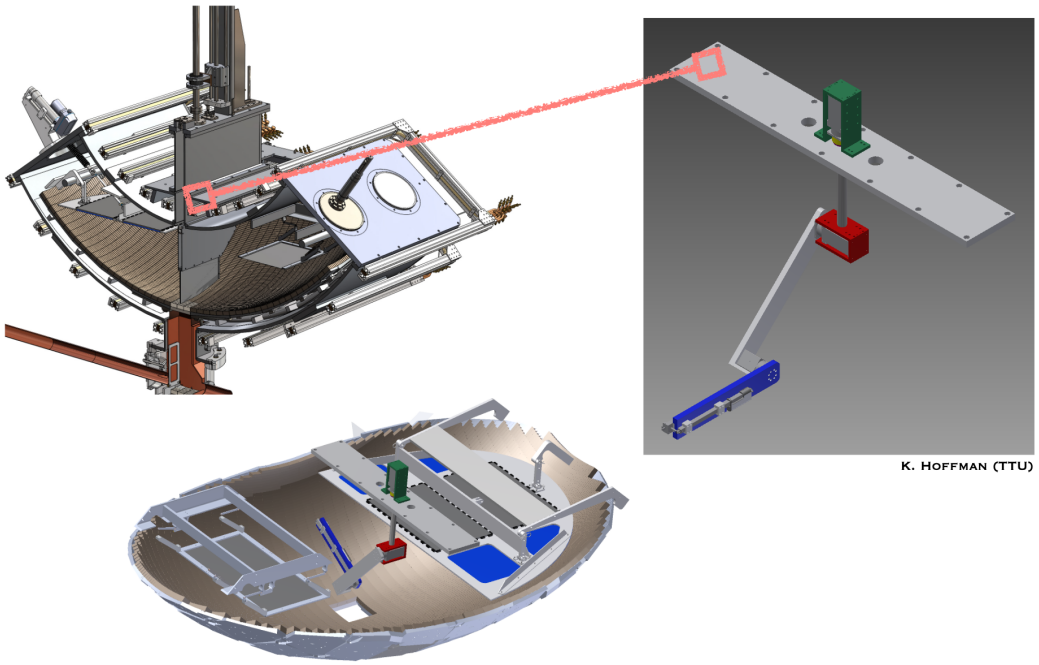

Figure 5.5: A CAD drawing of the mapping arm and field probe [60]. Top left: the UCNτ trap with the primary detector installed. Top right: the mapping arm and field probe, as well as a rectangular mount. Bottom: the UCNτ trap with the mapping arm installed. The mapping assembly can only be installed after removing the primary detector assembly from the apparatus because they are mounted on the same location of the vacuum jacket (orange line).

Magnetic field models that more closely resemble the actual magnetic fields in the trap would improve the accuracy of simulations of UCN trajectories in the trap. These improvements could further improve the understanding of UCN trajectories in the trap, and could also lead to an improvement of the Monte Carlo estimates for the rates at which insufficiently cleaned quasi-bound UCN and vibrationally heated UCN escape from the trap. Additionally, these magnetic field models could improve simulations that study the possibility that UCN could depolarize in the trap. However, all of these simulations would have to implement an interpolation of the magnetic field measurements in a way that does not violate Maxwell's equations, which is a non-trivial numerical problem.

5.7 Direct measurement of depolarized UCN

Section 4.6 discussed how depolarization of trapped UCN would have biased the extracted lifetime to be lower than it otherwise would have been, and also discussed the measurements that were made to estimate the rate of depolarization of UCN in the trap. The estimate of the rate of depolarization was made by extrapolating a theoretical calculation that the rate of depolarization is proportional to B^{-2}, where B is the strength of the magnetic holding field [56]. The validity of this extrapolation is suspect. Even if the extrapolation is valid, the amount of time that would need to be dedicated to this systematic study in order to reduce the systematic uncertainty due to depolarization of UCN would take a significant fraction of the time available during the next two run cycles.

If a UCN does depolarize in the trap, it would be attracted to, instead of repelled by, the Halbach array. The depolarized UCN would accelerate into the Halbach array and would be lost from the trap. The UCN density in the trap is greatest near the bottom of the trap and the probability that a UCN could depolarize is roughly uniform over the volume of the trap, so the rate of UCN depolarization is greatest near the bottom of the trap. These depolarized UCN could be detected using the same materials that all of the other active detectors are composed of: a thin ^{10}B surface to capture UCN on top of a ZnS:Ag scintillator to generate light. During the 2020 run cycle, this material was mounted on the surface of the trapdoor. This location was chosen because it is within the region where the rate of UCN depolarization is thought to be the greatest, and it was relatively easy to access. The ^{10}B-coated ZnS:Ag material was sufficiently thin that trappable UCN with the correct polarization were repelled by the Halbach array before reaching the detector, but depolarized UCN could still be detected.

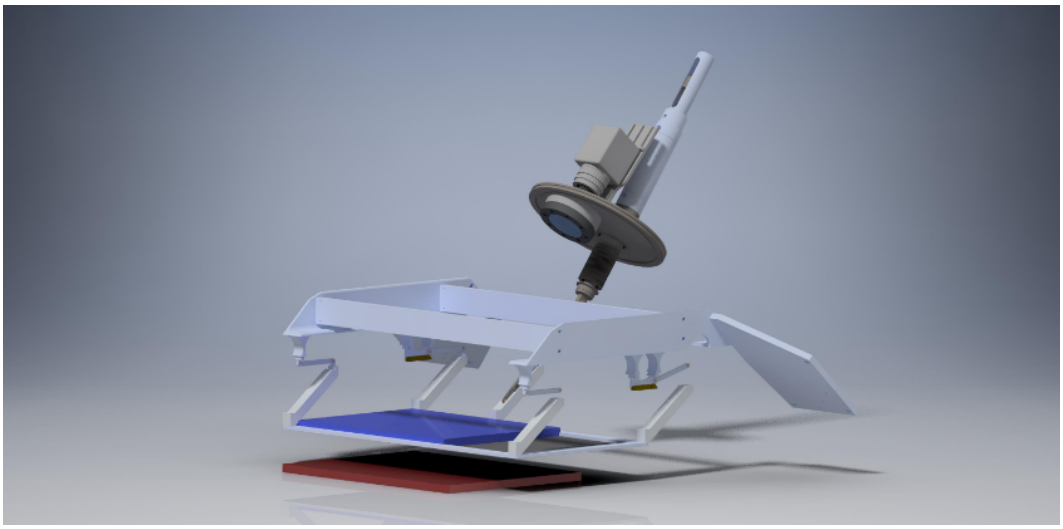

Figure 5.6: A CAD drawing of how an optical camera could be mounted outside of the vacuum jacket [61]. The camera would look for light from a scintillator mounted on the bottom of the trap (not shown) that would generate light when depolarized UCN are captured by the ^{10}B layer on top of the ZnS:Ag scintillator.

It was not feasible to use WLSFs to collect the photoelectrons from the scintillator mounted on the trapdoor without interfering with the trajectories of UCN stored in the trap. Instead, an optical camera was mounted above the trap to observe the scintillator at the bottom of the trap and look for flashes of light. Figure 5.6 shows a CAD drawing of how the optical camera was mounted outside of the vacuum jacket. No direct evidence of depolarization of UCN was found during the 2020 run cycle. However, the quality of the measurement was poor. There is ongoing work to improve the system's sensitivity to depolarized UCN. Direct measurement of depolarized UCN could lead to a reduction in the systematic uncertainty due to depolarization of trapped UCN, and could even enable the calculation of a systematic correction (if necessary).

5.8 Conclusion

This section discussed upgrades to the UCNτ experiment that could enable the measurement of the free neutron lifetime to a total uncertainty of less than 0.10 s. Section 5.2 discussed how increasing the number of UCN produced by the source would decrease the statistical uncertainty of an extracted lifetime, and Sections 5.3 and 5.4 discussed how to make better use of however many UCN may be available. Section 5.5 discussed the need to upgrade the primary detector to accommodate an increase in the number of UCN loaded into the trap. Sections 5.6 and 5.7 discussed how to decrease some of the largest systematic uncertainties of this experiment. If most of the potential upgrades presented in this section can be realized, then a future iteration of the UCNτ experiment will be able to achieve a measurement of the free neutron lifetime to a total uncertainty of less than 0.10 s.

Appendix A

OPTIMIZING THE HOLDING LENGTH

Consider a simplified version of a UCNτ with the following assumptions:

1. The number of UCN loaded into the trap is a random variable

$$n \sim \text{Integer}\left[\text{Gaussian}(N, fN)\right],$$

   where $f \geq 1$ is fixed and known and N is fixed but unknown

2. The number of background events during the unloading process is $b \sim \text{Poisson}(B)$, where B is fixed and known

3. Other than background, there are no systematic effects

4. All UCN are counted exactly 70 s after the beginning of the unloading process, and the counting of UCN is perfectly efficient

For a holding length t, the proportion of UCN that are expected not to decay between the end of the filling process and when the UCN are counted is $p = e^{-(120+t)/\tau}$. The 120 s in the exponent comes from two sources: 50 s of cleaning, as well as 70 s of delay between the beginning of the unloading process and when the UCN are counted.

As explained in Section 3.2.1, there is benefit to having multiple short holding lengths. An *octet* is a set of eight runs: four short runs with holding lengths of 20 s, 50 s, 100 s, and 200s; and four long runs with a holding length of t_{long}. Starting from Equation 3.14 and making the simplifying assumption that $s = 1$ and $\sigma_S = \sigma_C = 0$, the likelihood function simplifies to

$$\mathcal{L}\left(\tau, N | u_i, t_i\right) = -\sum_i \ln \text{Gaussian}\left(u_i \middle| Np_i + B, fNp_i^2 + Np_i(1-p_i) + B\right), \quad \text{(A.1)}$$

where i sums over all runs in the analysis.

The maximum likelihood of N and τ, as well as the uncertainties of those estimates (σ_N, σ_τ), was found using the methodology described in Section 3.9. σ_τ is inversely proportional to the square of the amount of data gathered, so maximizing the data gathered per unit of time (Δt) is equivalent to maximizing $\sigma_\tau^{-2}/\Delta t$.

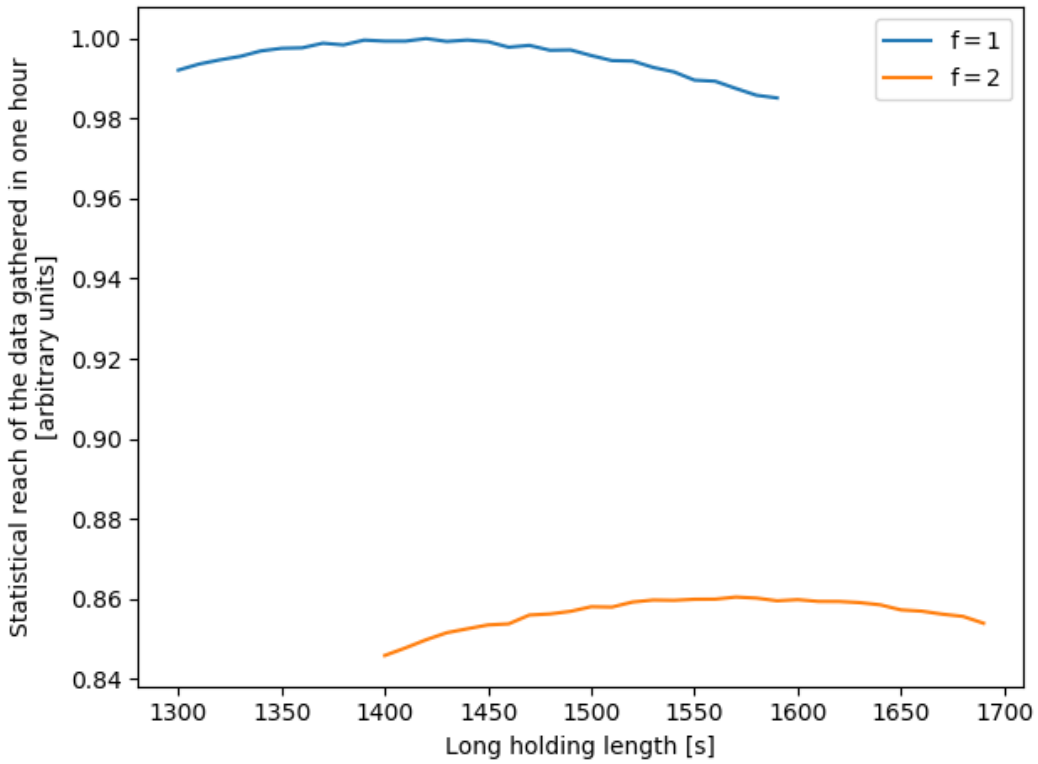

Figure A.1: Relative amounts of data gathered per hour as a function of t_L.

Each run takes a total time $t + \delta$ to complete, where t is the holding length of this run and δ is some constant amount of overhead time per run. An octet takes a total time of 370 s + $4t_L$ + 8δ to complete, where t_L is the holding length of the long runs and 370 s = (20 + 50 + 100 + 200) s. Monte Carlo simulations were modelled after the simplified version of the UCNτ experiment described at the beginning of this appendix. These simulation were repeated with different choices for the long holding length t_L. Figure A.1 shows the rate of data gathered per hour as a function of long holding length for reasonable choices of $N = 10,000$, $B = 50$, and $\delta = 630$ s, and for two options for f. Larger values for f (larger scales of the uncertainty of the normalization estimate σ_N relative to the ideal case of $\sigma_N = \sqrt{N}$) led to a larger value of the optimum holding length and decreased the statistical precision of a measurement made from a fixed number of runs.

Appendix B

UCN EVENT RECONSTRUCTION ALGORITHMS

In Algorithm 1:

- t_i is the time that the i^{th} photoelectron was observed
- $\text{ch}_i \in \{1, 2\}$ is which photomultiplier tube of the primary detector the i^{th} photoelectron was observed in
- T_C is the coincidence window
- T_T is the tail window
- N_{PE} is the photoelectron threshold

Algorithm 1 UCN reconstruction without a prompt window

INPUTS: $\{(t_i, \mathrm{ch}_i)\}_{i=1}^{L}$ where $t_i \leq t_{i+1}$
PARAMETERS: $T_C, T_T, N_{\mathrm{PE}}$

```
i ← 1
while i ≤ L do
    j, n, state ← i + 1, 1, 0
    while j ≤ L do
        if state = 0 then
            if t_j − t_i > T_C then
                break loop over j
            else
                n ← n + 1
                if ch_i ≠ ch_j then
                    state ← 1
                end if
            end if
        else
            if t_j − t_{j−1} > T_T then
                break loop over j
            else
                n ← n + 1
            end if
        end if
        j ← j + 1
    end while
    if state = 1 and n ≥ N_PE then
        push back reconstructed UCN event with start time t_i, end time t_{j−1} + T_T,
        and n photoelectron
        i ← j
    else
        i ← i + 1
    end if
end while
return reconstructed UCN events
```

In Algorithm 2:

- t_i is the time that the i^{th} photoelectron was observed
- $\text{ch}_i \in \{1, 2\}$ is which photomultiplier tube of the primary detector the i^{th} photoelectron was observed in
- T_C is the coincidence window
- T_P is the prompt window
- T_T is the tail window
- N_{PE} is the photoelectron threshold

Algorithm 2 UCN reconstruction with a prompt window

INPUTS: $\{(t_i, \mathrm{ch}_i)\}_{i=1}^{L}$ where $t_i \leq t_{i+1}$
PARAMETERS: $T_C, T_P, T_T, N_{\mathrm{PE}}$

```
i ← 1
while i ≤ L do
  n_P, n_T, state, j ← 1, 1, 0, i + 1
  while j ≤ L do
    if state = 0 then
      if t_j − t_i > T_C then
        break loop over j
      else
        n_P, n_T ← n_P + 1, n_T + 1
        if ch_i ≠ ch_j then
          state ← 1
        end if
      end if
    else if state = 1 then
      if t_j − t_i > T_P then
        if n_P ≥ N_PE then
          state ← 2
        else
          break loop over j
        end if
      else
        n_P, n_T ← n_P + 1, n_T + 1
      end if
    end if
    if state = 2 then
      if t_j − t_{j−1} > T_T then
        break loop over j
      else
        n_T ← n_T + 1
      end if
    end if
    j ← j + 1
  end while
  if state = 2 then
    push back reconstructed UCN event with start time t_i, end time t_{j−1} + T_T,
    n_P prompt photoelectrons, and n_T total photoelectrons
    i ← j
  else
    i ← i + 1
  end if
end while
return reconstructed UCN events
```

In Algorithm 3:

- t_i is the start time of the reconstructed UCN event
- δ_i is the length of each reconstructed UCN event
- N_i^P is the number of photoelectrons found in the prompt window of each reconstructed UCN event
- N_i^T is the total number of photoelectron found in each reconstructed UCN event
- T_P is the prompt window from the reconstruction algorithm
- T_T is the tail window from the reconstruction algorithm
- N_{PE} is the photoelectron threshold from the reconstruction algorithm
- cdf is a cumulative distribution function of the measured time of detection of photoelectron in a reconstructed UCN event, relative to the time of the first photoelectron in a reconstructed UCN event

Algorithm 3 UCN-event-tail correction

INPUTS: $\left\{\left(t_i, \delta_i, N_i^P, N_i^T\right)\right\}_{i=1}^{L}$ where $t_i \leq t_{i+1}$, $\delta_i \geq T_T$, and $N_{\mathrm{PE}} \leq N_i^P \leq N_i^T$
PARAMETERS: $T_P, T_T, N_{\mathrm{PE}}$, cdf

```
{F_i}_{i=1}^L , {W_i}_{i=1}^L ← 0, 1
for i ← 1 to L do
    A_i ← N_i^T / cdf(δ_i)
    for j ← i + 1 to L do
        F_j ← F_j + A_i [cdf(t_j − t_i + T_P) − cdf(t_j − t_i)]
        N_j^T ← N_j^T − A_i [cdf(t_j − t_i + δ_j) − cdf(t_j − t_i)]
    end for
    for j ← N_i^P − N_PE + 1 to N_i^P do
        W_i ← W_i − Poisson(j|F_i)
    end for
end for
return {W_i}_{i=1}^L
```

Appendix C

THE THINNED POISSON DISTRIBUTION

Denote the probability of observing a random value x drawn from a Poisson distribution with mean μ as $P_P(x|\mu)$. Denote the probability of observing a random value y from a Binomial distribution with N trials, each of which has a probability of success p, as $P_B(y|N, p)$.

Let $n \sim \text{Poisson}(N)$, and then let $x \sim \text{Binomial}(n, p)$. Then

$$\begin{aligned}
P(x|N,p) &= \sum_Y P_B(x|Y,p) \cdot P_P(Y|N) \\
&= \sum_{Y=x}^{\infty} \frac{Y!}{x!(Y-x)!} p^x (1-p)^{Y-x} \cdot \frac{N^Y e^{-N}}{Y!} \\
&= \frac{p^x e^{-N}}{x!} \sum_{Y=x}^{\infty} \frac{N^Y (1-p)^{Y-x}}{(Y-x)!} \\
&= \frac{(Np)^x e^{-N}}{x!} \sum_{Y=x}^{\infty} \frac{[N(1-p)]^{Y-x}}{(Y-x)!} \\
&= \frac{(Np)^x e^{-Np}}{x!} \sum_{Y=x}^{\infty} \frac{[N(1-p)]^{Y-x} e^{-N(1-p)}}{(Y-x)!} \\
&= \frac{(Np)^x e^{-Np}}{x!} \sum_{Y=0}^{\infty} \frac{[N(1-p)]^{Y} e^{-N(1-p)}}{Y!} \\
&= \frac{(Np)^x e^{-Np}}{x!} \cdot 1 \\
&= \frac{(Np)^x e^{-Np}}{x!} \\
&= P_P(x|Np).
\end{aligned}$$

Therefore $x \sim \text{Poisson}(Np)$.

Appendix D

CORRECTIONS TO BIASES INDUCED WHILE CALCULATING A PAIRED LIFETIME

Where

$$r \equiv \frac{Y_S}{Y_L},$$

$$\begin{aligned}
\frac{Y_S}{Y_L} \simeq & \left.\frac{Y_S}{Y_L}\right|_{\substack{Y_S=\langle Y_S\rangle \\ Y_L=\langle Y_L\rangle}} + \left.\frac{\partial r}{\partial Y_S}\right|_{\substack{Y_S=\langle Y_S\rangle \\ Y_L=\langle Y_L\rangle}} (Y_S - \langle Y_S\rangle) + \left.\frac{\partial r}{\partial Y_L}\right|_{\substack{Y_S=\langle Y_S\rangle \\ Y_L=\langle Y_L\rangle}} (Y_L - \langle Y_L\rangle) \\
& + \frac{1}{2}\left.\frac{\partial^2 r}{\partial Y_S^2}\right|_{\substack{Y_S=\langle Y_S\rangle \\ Y_L=\langle Y_L\rangle}} (Y_S - \langle Y_S\rangle)^2 + \frac{1}{2}\left.\frac{\partial^2 r}{\partial Y_L^2}\right|_{\substack{Y_S=\langle Y_S\rangle \\ Y_L=\langle Y_L\rangle}} (Y_L - \langle Y_L\rangle)^2 \\
= & \frac{\langle Y_S\rangle}{\langle Y_L\rangle} + \frac{1}{\langle Y_L\rangle}(Y_S - \langle Y_S\rangle) - \frac{\langle Y_S\rangle}{\langle Y_L\rangle^2}(Y_L - \langle Y_L\rangle) + \frac{\langle Y_S\rangle}{\langle Y_L\rangle^3}(Y_L - \langle Y_L\rangle)^2 \\
\left\langle \frac{Y_S}{Y_L} \right\rangle \simeq & \left\langle \frac{\langle Y_S\rangle}{\langle Y_L\rangle} + \frac{1}{\langle Y_L\rangle}(Y_S - \langle Y_S\rangle) - \frac{\langle Y_S\rangle}{\langle Y_L\rangle^2}(Y_L - \langle Y_L\rangle) + \frac{\langle Y_S\rangle}{\langle Y_L\rangle^3}(Y_L - \langle Y_L\rangle)^2 \right\rangle \\
= & \frac{\langle Y_S\rangle}{\langle Y_L\rangle} + \frac{1}{\langle Y_L\rangle}\langle Y_S - \langle Y_S\rangle\rangle - \frac{\langle Y_S\rangle}{\langle Y_L\rangle^2}\langle Y_L - \langle Y_L\rangle\rangle + \frac{\langle Y_S\rangle}{\langle Y_L\rangle^3}\left\langle (Y_L - \langle Y_L\rangle)^2 \right\rangle \\
= & \frac{\langle Y_S\rangle}{\langle Y_L\rangle} + \frac{\langle Y_S\rangle}{\langle Y_L\rangle^3}\sigma_{Y_L}^2 \\
= & \frac{\langle Y_S\rangle}{\langle Y_L\rangle}\left[1 + \left(\frac{\sigma_{Y_L}}{\langle Y_L\rangle}\right)^2\right] \Rightarrow \\
\frac{\langle Y_S\rangle}{\langle Y_L\rangle} \simeq & \left\langle \frac{Y_S}{Y_L} \right\rangle \left[1 + \left(\frac{\sigma_{Y_L}}{\langle Y_L\rangle}\right)^2\right]^{-1} \\
\simeq & \left\langle \frac{Y_S}{Y_L} \right\rangle \left[1 - \left(\frac{\sigma_{Y_L}}{\langle Y_L\rangle}\right)^2\right].
\end{aligned}$$

Where

$$\tau \equiv \frac{\Delta t}{\ln r},$$

$$\begin{aligned}
\frac{\Delta t}{\ln r} &\simeq \left.\frac{\Delta t}{\ln r}\right|_{r=\langle r\rangle} + \left.\frac{\partial \tau}{\partial r}\right|_{r=\langle r\rangle} (r - \langle r\rangle) + \frac{1}{2}\left.\frac{\partial^2 \tau}{\partial r^2}\right|_{r=\langle r\rangle} (r - \langle r\rangle)^2 \\
&= \frac{\Delta t}{\ln \langle r\rangle} - \frac{\Delta t}{\langle r\rangle \ln^2 \langle r\rangle}(r - \langle r\rangle) + \frac{1}{2}\frac{\Delta t(2 + \ln \langle r\rangle)}{\langle r\rangle^2 \ln^3 \langle r\rangle}(r - \langle r\rangle)^2 \\
\left\langle \frac{\Delta t}{\ln r} \right\rangle &\simeq \left\langle \frac{\Delta t}{\ln \langle r\rangle} - \frac{\Delta t}{\langle r\rangle \ln^2 \langle r\rangle}(r - \langle r\rangle) + \frac{1}{2}\frac{\Delta t(2 + \ln \langle r\rangle)}{\langle r\rangle^2 \ln^3 \langle r\rangle}(r - \langle r\rangle)^2 \right\rangle \\
&= \frac{\Delta t}{\ln \langle r\rangle} - \frac{\Delta t}{\langle r\rangle \ln^2 \langle r\rangle} \langle r - \langle r\rangle\rangle + \frac{\Delta t(2 + \ln \langle r\rangle)}{2 \langle r\rangle^2 \ln^3 \langle r\rangle} \left\langle (r - \langle r\rangle)^2 \right\rangle \\
&= \frac{\Delta t}{\ln \langle r\rangle} + \frac{\Delta t(2 + \ln \langle r\rangle)}{2 \langle r\rangle^2 \ln^3 \langle r\rangle}\sigma_r^2 \\
&= \frac{\Delta t}{\ln \langle r\rangle}\left[1 + \left(1 + \frac{1}{2}\ln \langle r\rangle\right)\left(\frac{\sigma_r}{\langle r\rangle \ln \langle r\rangle}\right)^2\right] \Rightarrow \\
\frac{\Delta t}{\ln \langle r\rangle} &\simeq \left\langle \frac{\Delta t}{\ln r} \right\rangle \left[1 + \left(1 + \frac{1}{2}\ln \langle r\rangle\right)\left(\frac{\sigma_r}{\langle r\rangle \ln \langle r\rangle}\right)^2\right]^{-1} \\
&\simeq \left\langle \frac{\Delta t}{\ln r} \right\rangle \left[1 - \left(1 + \frac{1}{2}\ln \langle r\rangle\right)\left(\frac{\sigma_r}{\langle r\rangle \ln \langle r\rangle}\right)^2\right].
\end{aligned}$$

Appendix E

ITERATIVELY CALCULATING A PAIRED LIFETIME

1. Initialize τ as a random value drawn from Uniform(800, 1000).

2. Hold τ fixed and optimize the normalization parameters using the method described in Section 3.7.9.

3. Calculate $N_i = \alpha^{(j)} + \beta_1^{(j)} P_{1,i} + \beta_2^{(j)} P_{2,i}$ and $\sigma_{N_i} = \sqrt{f N_i}$, where the i^{th} run is part of the j^{th} normalization subset.

4. Calculate $\sigma^2_{U_i - B_i} = N_i e^{-t_i/\tau}(1 - e^{-t_i/\tau}) s^2 + N_i e^{-t_i/\tau} \sigma_s^2 + B_i$.

5. Calculate
$$Y_i = \frac{U_i - B_i}{s N_i}$$
and
$$\sigma_{Y_i} = \langle Y_i \rangle \sqrt{\left(\frac{\sigma_{U_i - B_i}}{\langle U_i - B_i \rangle}\right)^2 + \left(\frac{\sigma_{N_i}}{N_i}\right)^2} = \frac{1}{s N_i} \sqrt{\sigma^2_{U_i - B_i} + f N_i s^2 e^{-2t_i/\tau}}.$$

6. For each short-long pair of runs, calculate
$$\Delta t_j = t_L - t_S,$$
$$r_j = \frac{Y_S}{Y_L}\left[1 - \left(\frac{\sigma_{Y_L}}{Y_L}\right)^2\right]$$
and
$$\begin{aligned}\sigma_{r_j} &= \sqrt{\left(\frac{\sigma_{Y_S}}{\langle Y_L \rangle}\left[1 - \left(\frac{\sigma_{Y_L}}{\langle Y_L \rangle}\right)^2\right]\right)^2 + \left(\frac{\langle Y_S \rangle \sigma_{Y_L}}{\langle Y_L \rangle^2}\left[1 - 3\left(\frac{\sigma_{Y_L}}{\langle Y_L \rangle}\right)^2\right]\right)^2} \\ &= \sqrt{\left(\sigma_{Y_S} e^{t_L/\tau}\left[1 - \left(\sigma_{Y_L} e^{t_L/\tau}\right)^2\right]\right)^2 + \left(\sigma_{Y_L} e^{(2t_L - t_S)/\tau}\left[1 - 3\left(\sigma_{Y_L} e^{t_L/\tau}\right)^2\right]\right)^2}.\end{aligned}$$

7. For each short-long pair of runs, calculate

$$\tau_j = \frac{\Delta t_j}{\ln r_j}\left[1 - \left(1 + \frac{1}{2}\ln r_j\right)\left(\frac{\sigma_{r_j}}{r_j \ln r_j}\right)^2\right] \tag{E.1}$$

and

$$\begin{aligned}\sigma_{\tau_j} &= \frac{\Delta t}{\left\langle r_j\right\rangle^3 \ln^4\left\langle r_j\right\rangle}\left[\left(\left\langle r_j\right\rangle^2 - \sigma_{r_j}^2\right)\ln^2\left\langle r_j\right\rangle - 3\sigma_{r_j}^2\left(1 + \ln\left\langle r_j\right\rangle\right)\right]\sigma_{r_j} \\ &= \frac{\tau^4}{e^{3\Delta t/\tau}\Delta t^3}\left[\left(e^{2\Delta t/\tau} - \sigma_{r_j}^2\right)\left(\frac{\Delta t}{\tau}\right)^2 - 3\sigma_{r_j}^2\left(1 + \frac{\Delta t}{\tau}\right)\right]\sigma_{r_j}.\end{aligned} \tag{E.2}$$

8. Save the weighted average of the $(\tau_j, \sigma_{\tau_j})$s as $\tau' \pm \sigma_{\tau'}$, using the $\sigma_{\tau_j}^2$s as the weights in the average.

9. If $|\tau' - \tau| < 0.0001$, s then report $\tau' \pm \sigma_{\tau'}$ as the paired lifetime. Otherwise, $\tau \leftarrow \tau'$, and return to step 2.

BIBLIOGRAPHY

[1]P. A. Zyla et al., "Review of particle physics", Progress of Theoretical and Experimental Physics **2020**, 083C01 (2020).

[2]A. Czarnecki, W. J. Marciano, and A. Sirlin, "Radiative corrections to neutron and nuclear beta decays revisited", Physical Review D **100** (2019).

[3]M. Kobayashi and T. Maskawa, "CP-violation in the renormalizable theory of weak interaction", Progress of Theoretical Physics **49**, 652–657 (1973).

[4]F. Iocco et al., "Primordial nucleosynthesis: From precision cosmology to fundamental physics", Physics Reports **472**, 1–76 (2009).

[5]J. S. Nico et al., "Measurement of the neutron lifetime by counting trapped protons in a cold neutron beam", Physical Review C **71**, 055502 (2005).

[6]Y. Y. Kosvintsev, V. I. Morozov, and G. I. Terekhov, "Measurement of neutron lifetime through storage of ultracold neutrons", Journal of Experimental and Theoretical Physics **44**, 571–574 (1986).

[7]P. E. Spivak, "Neutron lifetime from atomic-energy-institute experiment", Journal of Experimental and Theoretical Physics **67**, 1735–1740 (1988).

[8]J. Last et al., "Pulsed-beam neutron-lifetime measurement", Physical Review Letters **60**, 995–998 (1988).

[9]R. Kossakowski et al., "Neutron lifetime measurement with a helium-filled time projection chamber", Nuclear Physics A **503**, 473–500 (1989).

[10]W. Mampe et al., "Neutron lifetime measured with stored ultracold neutrons", Physical Review Letters **63**, 593–596 (1989).

[11]W. Paul et al., "Measurement of the neutron lifetime in a magnetic storage ring", Zeitschrift für Physik C Particles and Fields **45**, 25 (1989).

[12]V. P. Alfimenkov et al., "Results of neutron lifetime measurements with gravitational UCN trap", Journal of Experimental and Theoretical Physics Letters **52**, 373–378 (1990).

[13]J. Byrne et al., "Measurement of the neutron lifetime by counting trapped protons", Physical Review Letters **65**, 289–292 (1990).

[14]V. Nesvizhevsky et al., "Measurement of the neutron lifetime in a gravitational trap and analysis of experimental errors", Journal of Experimental and Theoretical Physics **75** (1992).

[15]W. Mampe et al., "Measuring neutron lifetime by storing ultracold neutrons and detecting inelastically scattered neutrons", Journal of Experimental and Theoretical Physics Letters **57**, 82–87 (1993).

[16]J. Byrne et al., "A revised value for the neutron lifetime measured using a penning trap", Europhysics Letters **33**, 187–192 (1996).

[17]S. Arzumanov et al., "Neutron lifetime measured by monitored storing of ultra-cold neutrons", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment **440**, 511–516 (2000).

[18]M. S. Dewey et al., "Measurement of the neutron lifetime using a proton trap", Physical Review Letters **91**, 152302 (2003).

[19]A. Serebrov et al., "Measurement of the neutron lifetime using a gravitational trap and a low-temperature fomblin coating", Physics Letters B **605**, 72–78 (2005).

[20]A. Pichlmaier et al., "Neutron lifetime measurement with the UCN trap-in-trap MAMBO II", Physical Letters B **693**, 221–226 (2010).

[21]A. Steyerl et al., "Quasielastic scattering in the interaction of ultracold neutrons with a liquid wall and application in a reanalysis of the Mambo I neutron-lifetime experiment", Physical Review C **85**, 065503 (2012).

[22]S. Arzumanov et al., "Analysis and correction of the measurement of the neutron lifetime", Journal of Experimental and Theoretical Physics Letters **95** (2012).

[23]S. Arzumanov et al., "A measurement of the neutron lifetime using the method of storage of ultracold neutrons and detection of inelastically up-scattered neutrons", Physics Letters B **745**, 79–89 (2015).

[24]C. L. Morris et al., "A new method for measuring the neutron lifetime using an in situ neutron detector", Review of Scientific Instruments **88**, 053508 (2017).

[25]A. P. Serebrov et al., "Neutron lifetime measurements with a large gravitational trap for ultracold neutrons", Physical Review C **97**, 055503 (2018).

[26]R. W. Pattie et al., "Measurement of the neutron lifetime using a magneto-gravitational trap and in situ detection", Science **360**, 627–632 (2018).

[27]V. F. Ezhov et al., "Measurement of the neutron lifetime with ultra-cold neutrons stored in a magneto-gravitational trap", Journal of Experimental and Theoretical Physics Letters **107**, 671–675 (2018).

[28]B. Fornal and B. Grinstein, "Dark matter interpretation of the neutron decay anomaly", Physical Review Letters **120** (2018).

[29]Z. Tang et al., "Search for the neutron decay $n \to X\gamma$, where X is a dark matter particle", Physical Review Letters **121** (2018).

[30]X. Sun et al., "Search for dark matter decay of the free neutron from the UCNA experiment: $n \to Xe^+e^-$", Physical Review C **97** (2018).

[31]S. Aoki et al., *Review of Lattice Results Concerning Low-Energy Particle Physics*, 2016.

[32]J. A. Bailey et al., "$|V_{ub}|$ from $B \to \pi \ell \nu$ decays and (2+1)-flavor lattice QCD", Physical Review D **92** (2015).

[33]R. Golub, *Ultra-Cold Neutrons* (Adam Hilger, Bristol Philadelphia, 1991).

[34]T. M. Ito et al., "Performance of the upgraded ultracold neutron source at Los Alamos National Laboratory and its implication for a possible neutron electric dipole moment experiment", Physical Review C **97** (2018).

[35]A. Anghel et al., "Solid deuterium surface degradation at ultracold neutron sources", The European Physical Journal A **54** (2018).

[36]A. T. Holley et al., "A high-field adiabatic fast passage ultracold neutron spin flipper for the UCNA experiment", Review of Scientific Instruments **83**, 073505 (2012).

[37]P. Walstrom et al., "A magneto-gravitational trap for absolute measurement of the ultra-cold neutron lifetime", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment **599**, 82–92 (2009).

[38]K. Halbach, "Design of permanent multipole magnets with oriented rare earth cobalt material", Nuclear Instruments and Methods **169**, 1–10 (1980).

[39]Z. Wang et al., "A multilayer surface detector for ultracold neutrons", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment **798**, 30–35 (2015).

[40]*Alpha Detection: EJ-440, EJ-442*, May 2018.

[41]B. Clement et al., "Boron-10 conversion layer for ultra-cold neutron detection", Journal of Instrumentation **14**, P09003–P09003 (2019).

[42]C. L. Morris et al. (UCNτ Collaboration), "A new method for measuring the neutron lifetime using an in situ neutron detector", Review of Scientific Instruments **88**, 053508 (2017).

[43]N. Callahan et al. (UCNτ Collaboration), "Monte Carlo simulations of trapped ultracold neutrons in the UCNτ experiment", Physical Review C **100**, 015501 (2019).

[44]*Model MCS6A, 64 Bit 5/(6) Input 100 ps Multistop TDC, Multiscaler, Time-Of-Flight*, Last accessed March 1, 2021.

[45]*Design and Development Framework for Cross-platform Applications*, Last accessed on May 2, 2021.

[46]*LabJack U3-HV Multifunction DAQ - USB*, Last accessed on May 2, 2021.

[47]H. Marouani and M. R. Dagenais, "Internal clock drift estimation in computer clusters", The Journal of Network and Systems Management **2008** (2008).

[48]*Artix-7 35T Arty FPGA Evaluation Kit*, Last accessed March 1, 2021.

[49] *LCB716 Single-Pole, Normally Closed OptoMOS Relay*, June 2018.

[50] *SN74LV125AT Quadruple Bus Buffer Gate with 3-state Outputs*, Aug. 2005.

[51] K. Pearson, "On lines and planes of closest fit to systems of points in space", The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **2**, 559–572 (1901).

[52] J. Kiefer, "Sequential minimax search for a maximum", Proceedings of the American Mathematical Society **4**, 502–506 (1953).

[53] H. Cramér, *Mathematical Methods of Statistics* (Princeton University Press, Sept. 1946).

[54] S. J. Seestrom et al. (UCNτ Collaboration), "Total cross sections for ultracold neutrons scattered from gases", Physical Review C **95**, 015501 (2017).

[55] S. J. Seestrom et al. (UCNτ Collaboration), "Upscattering of ultracold neutrons from gases", Physical Review C **92**, 065501 (2015).

[56] A. Steyerl et al., "Spin flip loss in magnetic confinement of ultracold neutrons for neutron lifetime experiments", Physical Review C **95**, 035502 (2017).

[57] K. K. H. Leung et al., "A next-generation inverse-geometry spallation-driven ultracold neutron source", Journal of Applied Physics **126**, 224901 (2019).

[58] C. J. Werner et al., *MCNP Version 6.2 Release Notes*, Feb. 2018.

[59] C. Morris (UCNτ Collaboration), ^{10}B *Coated LYSO Screens for Detecting Ultra-Cold Neutrons*, American Physical Society Division of Nuclear Physics Meeting Abstracts, 2020.

[60] A. Holley (UCNτ Collaboration), *Magnetic Field Characterization of the UCNτ Magneto-gravitational Ultracold Neutron Trap*, American Physical Society Division of Nuclear Physics Meeting Abstracts, 2019.

[61] A. Holley (UCNτ Collaboration), *Spin Dynamics in the UCNτ Magneto-gravitational Ultracold Neutron Trap*, American Physical Society Division of Nuclear Physics Meeting Abstracts, 2020.