2019

A DETECTION AND DATA ACQUISITION SYSTEM FOR PRECISION BETA DECAY SPECTROSCOPY

Aaron P. Jezghani
University of Kentucky, apjezghani@gmail.com
Digital Object Identifier: https://doi.org/10.13023/etd.2019.177

Right click to open a feedback form in a new tab to let us know how this document benefits you.

STUDENT AGREEMENT:

I represent that my thesis or dissertation and abstract are my original work. Proper attribution has been given to all outside sources. I understand that I am solely responsible for obtaining any needed copyright permissions. I have obtained needed written permission statement(s) from the owner(s) of each third-party copyrighted matter to be included in my work, allowing electronic distribution (if such use is not permitted by the fair use doctrine) which will be submitted to UKnowledge as Additional File.

I hereby grant to The University of Kentucky and its agents the irrevocable, non-exclusive, and royalty-free license to archive and make accessible my work in whole or in part in all forms of media, now or hereafter known. I agree that the document mentioned above may be made available immediately for worldwide access unless an embargo applies.

I retain all other ownership rights to the copyright of my work. I also retain the right to use in future works (such as articles or books) all or part of my work. I understand that I am free to register the copyright to my work.

REVIEW, APPROVAL AND ACCEPTANCE

The document mentioned above has been reviewed and accepted by the student's advisor, on behalf of the advisory committee, and by the Director of Graduate Studies (DGS), on behalf of the program; we verify that this is the final, approved version of the student's thesis including all changes required by the advisory committee. The undersigned agree to abide by the statements above.

<div align="right">

Aaron P. Jezghani, Student

Dr. Christopher B. Crawford, Major Professor

Dr. Christopher B. Crawford, Director of Graduate Studies

</div>

A DETECTION AND DATA ACQUISITION SYSTEM FOR PRECISION BETA
DECAY SPECTROSCOPY

DISSERTATION

A dissertation submitted in partial
fulfillment of the requirements for
the degree of Doctor of Philosophy
in the College of Arts and Sciences
at the University of Kentucky

By
Aaron P. Jezghani
Lexington, Kentucky

Director: Dr. Christopher B. Crawford, Professor of Physics & Astronomy
Lexington, Kentucky
2019

ABSTRACT OF DISSERTATION

A DETECTION AND DATA ACQUISITION SYSTEM FOR PRECISION BETA
DECAY SPECTROSCOPY

Free neutron and nuclear beta decay spectroscopy serves as a robust laboratory for investigations of the Standard Model of Particle Physics. Observables such as decay product angular correlations and energy spectra overconstrain the Standard Model and serve as a sensitive probe for Beyond the Standard Model physics. Improved measurement of these quantities is necessary to complement the TeV scale physics being conducted at the Large Hadron Collider. The UCNB, ^{45}Ca, and Nab experiments aim to improve upon existing measurements of free neutron decay angular correlations and set new limits in the search for exotic couplings in beta decay. To achieve these experimental goals, a highly-pixelated, thick silicon detector with a 100 nm entrance window has been developed for precision beta spectroscopy and the direct detection of 30 keV beta decay protons. The detector has been characterized for its performance in energy reconstruction and particle arrival time determination. A Monte Carlo simulation of signal formation in the silicon detector and propagation through the electronics chain has been written to develop optimal signal analysis algorithms for minimally biased energy and timing extraction. A tagged-electron timing test has been proposed and investigated as a means to assess the validity of these Monte Carlo efforts.

A universal platform for data acquisition (DAQ) has been designed and implemented in National Instrument's PXIe-5171R digitizer/FPGA hardware. The DAQ retains a ring buffer of the most recent 400 ms of data in all 256 channels, so that a waveform trace can be returned from any combination of pixels and resolution for complete energy reconstruction. Low-threshold triggers on individual channels were implemented in FPGA as a generic piecewise-polynomial filter for universal, real-time digital signal processing, which allows for arbitrary filter implementation on a pixel-by-pixel basis. This system is universal in the sense that it has complete flexible, complex, and debuggable triggering at both the pixel and global level without recompiling the firmware. The culmination of this work is a system capable of a 10 keV trigger threshold, 3 keV resolution, and maximum 300 ps arrival time systematic, even in the presence of large amplitude noise components.

Author's signature:_____Aaron P. Jezghani

Date:_____May 3, 2019

A DETECTION AND DATA ACQUISITION SYSTEM FOR PRECISION BETA
DECAY SPECTROSCOPY

By
Aaron P. Jezghani

Director of Dissertation: Christopher B. Crawford

Director of Graduate Studies: Christopher B. Crawford

Date: May 3, 2019

To my wife Maggie and our growing family - I have worked so that y'all can have
the life you deserve.

ACKNOWLEDGMENTS

earned her PhD and followed her dreams, she was the inspiration for me to follow. She has always provided the stable foundation that I was able to use to get to where I am today. With the coming arrival of our first daughter, Elizabeth, my work has so much more meaning. We have only just started our life together in Georgia, and I expect many wonderful years to come.

TABLE OF CONTENTS

LIST OF TABLES

Chapter 1: Introduction

Precision measurements of beta decay parameters, in particular of neutron decay, yield sensitive tests of the Standard Model and of new physics. The subject of this dissertation is the development of a silicon-based system for uniform detection of both beta decay electrons and protons in a magnetic spectrometer. This chapter motivates the underlying physics of these experiments, whose common detection elements are addressed through the collection of efforts in this work.

1.1 Neutron Beta Decay

Neutrons are electrically neutral particles of mass 939.565 MeV, and are comprised of two down and one up valence quarks. Together with protons, which have a slightly lower mass of 938.272 MeV, they form the nuclei of atoms. This mass difference energetically allows free neutron decay through the process:

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

The lifetime of the free neutron is 880.2 s [1]; bound neutron lifetimes vary by isotope, from extremely short-lived to completely stable.

While at low energy, beta decay appears as a contact interaction, the weak force is mediated by the exchange of a W⁻ boson, as shown in Figure 1.1. Mirroring the
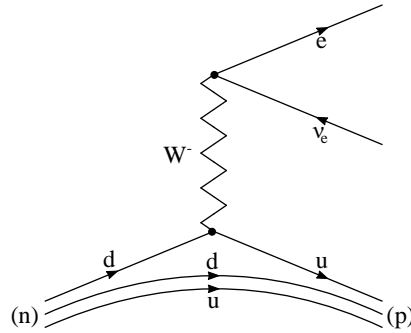
Figure 1.1: In beta decay, one of the bound down quarks of the neutron weakly transitions to an up quark, producing a proton. Meanwhile, the W⁻ boson decays into an electron and an electron antineutrino.

formalism of electromagnetic scattering, Enrico Fermi proposed a beta decay model from which the electron energy spectrum could be calculated as

$$\frac{\mathrm{d}\Gamma}{\mathrm{d}E_\mathrm{e}} = G^2 |M|^2 F\left(\pm Z, E_\mathrm{e}\right) p_\mathrm{e} E_\mathrm{e} \left(E_0 - E_\mathrm{e}\right)^2, \tag{1.1}$$

where G is an overall coupling constant for the process, M is the matrix element between initial and final states, E_e is the total electron energy, E_0 is the endpoint

energy, and $F\left(\pm Z, E_\mathrm{e}\right)$ is the Fermi function to account for Coulombic modification of the outgoing electron phase space. While this model was successful in describing the spectral shape, it was unable to account for the non-zero change in angular momentum observed in certain nuclear species such as thorium. George Gamow and Edward Teller proposed the linear superposition of another matrix element, such as an axial-vector interaction, to account for the observed change in angular momentum [2].

Two decades later, a more generalized interaction Hamiltonian for the decay of the neutron was proposed by Tsung-Dao Lee and Chen Ning Yang as an extension of the vector and axial-vector models proposed by Fermi, Gamow, and Teller [3]. Rather than restricting the theory of beta decay to vector and axial-vector components, they suggested including all types of exchanges:

$$H_\mathrm{int} = \sum_{j \in V,A,S,T,P} (\bar{\psi}_\mathrm{p}\Gamma_j\psi_\mathrm{n})(C_j\bar{\psi}_\nu\Gamma_j\psi_\mathrm{e} + C'_j\bar{\psi}_\nu\Gamma_j\gamma_5\psi_\mathrm{e}) + \text{Hermitian conjugate.} \quad (1.2)$$

The operators $\Gamma_\mathrm{V} = \gamma_\mu$, $\Gamma_\mathrm{A} = i\gamma_\mu\gamma_5$, $\Gamma_\mathrm{S} = 1$, $\Gamma_\mathrm{T} = \sigma_{\lambda\mu}$, and $\Gamma_\mathrm{P} = \gamma_5$ describe the vector, axial-vector, scalar, tensor, and pseudoscalar interactions, respectively, with C_j and C'_j the coupling constants for the parity conserving and non-conserving comoponents for each term. Two experiments, one by Chien-Shiung Wu and the low-temperature laboratory of the National Bureau of Standards, and the other by Richard Garwin, Leon Lederman, and Marcel Weinrich at Columbia's Nevis Cyclotron Laboratory, confirmed maximal parity violation in weak interactions [4,5].

The following year, John David Jackson, Sam Treiman, and Henry Wyld calculated the decay rate based on Equation 1.2, and determined a number of correlation measurements to ascertain each of the Lee and Yang couplings [6]. Focusing only on the terms addressed in the scope of this thesis, the differential decay rate for the neutron can be written:

$$\begin{aligned}
\frac{\mathrm{d}^3\Gamma}{\mathrm{d}E_\mathrm{e}\mathrm{d}\Omega_\mathrm{e}\mathrm{d}\Omega_\nu} \propto\ & p_\mathrm{e}E_\mathrm{e}\left(E_0 - E_\mathrm{e}\right)^2 \\
& \times\xi\left\{1 + a\frac{\vec{p_\mathrm{e}}\cdot\vec{p_\nu}}{E_\mathrm{e}E_\nu} + b\frac{m_\mathrm{e}}{E_\mathrm{e}} + \langle\vec{\sigma_\mathrm{n}}\rangle\cdot\left(A\frac{\vec{p_\mathrm{e}}}{E_\mathrm{e}} + B\frac{\vec{p_\nu}}{E_\nu} + \cdots\right) + \cdots\right\}.
\end{aligned} \quad (1.3)$$

In the above expression, each of the correlation coefficients (a, b, A, B) and ξ depend on the Fermi and Gamow-Teller matrix elements and the coupling coefficients C_j, C'_j.

Fully expanded in terms of these coefficients, then, the correlation terms are

$$
\begin{aligned}
a\xi &= |M_F|^2 \left(-|C_S|^2 + |C_V|^2 - |C_S'|^2 + |C_V'|^2\right) \\
&\quad + |M_{GT}|^2 \left(|C_T|^2 - |C_A|^2 + |C_T'|^2 - |C_A'|^2\right), \\
b\xi &= \pm 2\mathrm{Re}\left[|M_F|^2 \left(C_S C_V^* + C_S' C_V'^*\right) + |M_{GT}|^2 \left(C_T C_A^* + C_T' C_A'^*\right)\right], \\
A\xi &= 2\mathrm{Re}\left[\pm |M_{GT}|^2 \lambda_{J'J} \left(C_T C_T'^* - C_A C_A'^*\right)\right. \\
&\quad \left. + \delta_{J'J} |M_F||M_{GT}| \left(\frac{J}{J+1}\right)^{\frac{1}{2}} \left(C_S C_T'^* + C_S' C_T^* - C_V C_A'^* - C_V' C_A^*\right)\right],
\end{aligned}
$$

and

$$
\begin{aligned}
B\xi &= 2\mathrm{Re}\left\{|M_{GT}|^2 \lambda_{J'J} \left[\frac{m_e}{E_e}\left(C_T C_A'^* + C_T' C_A^*\right) \pm \left(C_T C_T'^* + C_A C_A'^*\right)\right]\right. \\
&\quad - \delta_{J'J} |M_F||M_{GT}| \left(\frac{J}{J+1}\right)^{\frac{1}{2}} \left[\left(C_S C_T'^* + C_S' C_T^* + C_V C_A'^* + C_V' C_A^*\right)\right. \\
&\quad \left.\left. \pm \frac{m_e}{E_e}\left(C_S C_A'^* + C_S' C_A^* + C_V C_T'^* + C_V' C_T^*\right)\right]\right\},
\end{aligned}
$$

where $\xi = |M_F|^2(|C_S|^2 + |C_V|^2 + |C_S'|^2 + |C_V'|^2) + |M_{GT}|^2(|C_T|^2 + |C_A|^2 + |C_T'|^2 + |C_A'|^2)$. The terms $\delta_{J'J}$ and $\lambda_{J'J}$ pertain to the selection rules for change in angular momentum. The former is the Kronecker delta symbol and the latter is

$$
\lambda_{J'J} = \begin{cases}
1, & J \to J' = J - 1 \\
1/(J+1), & J \to J' = J \\
-J/(J+1), & J \to J' = J + 1.
\end{cases}
$$

The electron-neutrino correlation coefficient, a, and the Fierz interference term, b, are measured using unpolarized neutrons, while the electron asymmetry, A, and the neutrino asymmetry, B, require polarized neutrons. This "alphabet" serves as the basis for a number of modern beta decay experiments to test the Standard Model and search for Beyond the Standard Model physics.

1.2 Tests of the Standard Model

According to the $V - A$ theory of the weak interaction in the Standard Model, only the vector and axial-vector coefficients are non-zero and only left-handed couplings are present. Assuming this to be true, the term ξ reduces to $g_V^2 + 3g_A^2$ for the free neutron, where g_V and g_A are the vector and axial-vector weak coupling constants, respectively. These two terms arise because the strong and weak eigenstates of the quarks are not the same. While the conserved-vector-current (CVC) hypothesis establishes that $g_V = +1$, no such constraint exists for g_A. As such, the coefficients in Equation 1.4 can be re-written in terms of the ratio of these coupling constants, $\lambda \equiv g_A/g_V$:

$$
a = \frac{1 - |\lambda|^2}{1 + 3|\lambda|^2}, \quad b = 0, \quad A = -2\frac{|\lambda|^2 + \mathrm{Re}(\lambda)}{1 + 3|\lambda|^2}, \quad B = 2\frac{|\lambda|^2 - \mathrm{Re}(\lambda)}{1 + 3|\lambda|^2}. \tag{1.4}
$$

Combined with the neutron lifetime, τ_n, the ratio λ can be used to determine the matrix element [7]

$$|V_{\mathrm{ud}}|^2 = \frac{4908.7(1.9)\mathrm{sec}}{\tau_{\mathrm{n}}(1 + 3\lambda^2)} \tag{1.5}$$

of the Cabibbo-Kobayashi-Maskawa (CKM) matrix,

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

which mixes the mass and weak eigenstates of the quarks [8, 9].

The Standard Model is complete in its description of the quark flavor changing due to the weak interaction. If it is indeed universal in that there are only three generations of quarks, then the sum of each squared element along any row or column must be exactly one:

$$\sum_k |V_{ik}|^2 = \sum_i |V_{ik}|^2 = 1. \tag{1.7}$$

At present, the accepted experimental values for the CKM matrix are [1]:

$$V_{\mathrm{CKM}} = \begin{pmatrix} 0.97420 \pm 0.00021 & 0.2243 \pm 0.0005 & 0.00394 \pm 0.00036 \\ 0.218 \pm 0.004 & 0.997 \pm 0.017 & 0.0422 \pm 0.0008 \\ 0.0081 \pm 0.0005 & 0.0394 \pm 0.0023 & 1.019 \pm 0.025 \end{pmatrix}. \tag{1.8}$$

Using these values, both the row and the column containing V_{ud} satisfy the condition of weak universality: the sum of the square of the elements is 1 within error bars. However, more precise results are needed to establish the extent to which this is true. As such, improved measurements of each of the CKM elements, including V_{ud}, are ongoing.

At present, the best values of V_{ud} are determined using superallowed $0^+ \rightarrow 0^+$ decays. In spite of the theoretical uncertainties associated with the calculation of the nuclear matrix elements involved in the decay, the systematic effects involved in such efforts are extremely small. A review of 20 isotopes from 222 measurements is reported by J. C. Hardy and I. S. Towner [10], and accounts for the experimental value shown in Equation 1.8. Nonetheless, V_{ud} extractions from neutron beta decay measurements are an invaluable, independent cross-check with different systematics.

While correlation tests and lifetime measurements from neutron beta decay can be used to determine V_{ud} with virtually no nuclear wavefunction contributions, experimental precision limits the competitiveness of such measurements in comparison to superallowed Fermi decays. Figure 1.2 shows the determination of V_{ud} from several methods, with a breakdown of the contributions to the uncertainty for each. The radiative corrections for each contribute roughly the same amount to the overall uncertainty, but the experimental systematics and nuclear corrections vary. Neutron beta decay correlation studies are promising in that they can exceed the precision of superallowed decays if the experimental contributions can be reduced to comparable levels.
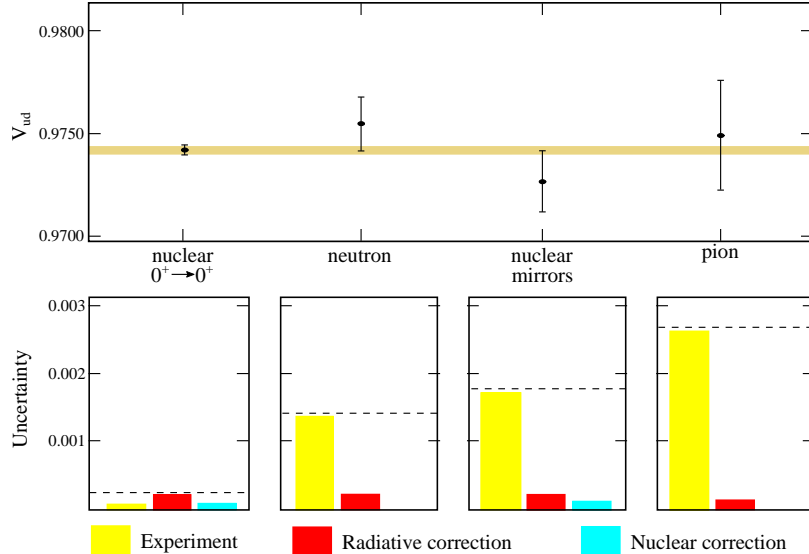
Figure 1.2: The precision of V_{ud} measurements from superallowed beta decays, neutron correlation measurements, nuclear mirror decays, and pion experiments. At present, the best value is determined from superallowed beta decay, but efforts are ongoing to improve the precision of beta decay correlation measurements. Reproduced from [11].

Calculation of V_{ud} from neutron beta decay requires two measurements, each with their own historical difficulties. Two methods used to determine the neutron lifetime, in-flight decay experiments and lifetime measurements in traps, are in contention with each, reducing the overall precision of τ_n [12]. In addition, there is the uncertainty in the extracted value of λ. The electron asymmetry, A, is the most sensitive correlation coefficient to λ, and is the easiest to measure. Several experiments such as PERKEO and UCNA have determined the electron asymmetry, A, yielding a reported average value of -0.1184 ± 0.0010, which includes an error scaling factor of 2.4 [1, 13, 14]. Omitted from the PDG average, however, is the most recent UCNA result, which reports a value of -0.12015 ± 0.00072 [15]; this will reduce the overall uncertainty in A, and hence λ.

The next most sensitive term to λ is the electron-neutrino correlation, a. Because of the difficulty in direct detection of the neutrino, and the marginally less difficult measurement of the proton, the uncertainty in measurement is worse than that of A. Efforts such as aSPECT and aCORN measured the proton spectrum via a retardation potential or indirectly detected the protons in a time-of-flight effort, respectively, to extract a [16, 17]; at present, the accepted value is -0.1059 ± 0.0028 [1]. In spite of the reduced sensitivity to λ and the greater experimental uncertainties obtained thus far, the electron-neutrino correlation term provides an additional cross-check in the test of the Standard Model.

Least sensitive to λ is the neutrino asymmetry, B. Like a, determining a value for B is challenging because it requires either detection of the neutrino or a coincident measurement of the electron and proton. A handful of experiments have measured

the neutrino asymmetry, but the best precision came from the measurements at the Institut Laue–Langevin (ILL) and the PERKEO spectrometer [18, 19]. The current PDG value is reported as 0.9807 ± 0.0030 [1].

Figure 1.3 summarizes the recent measurements from neutron beta decay correlation studies to determine λ. Due to the better precision in experiment, the value is almost exclusively comprised of measurements of the electron asymmetry. In the past year, a value of the electron-neutrino correlation has finally been determined with sufficiently small error to be included in the determination of λ, and further analysis of the data is expected to reduce the error bars for a more competitive result [20].

WEIGHTED AVERAGE
-1.2724±0.0023 (Error scaled by 2.2)

			χ^2
DARIUS	17	SPEC	
MENDENHALL	13	UCNA	1.1
MUND	13	SPEC	3.6
SCHUMANN	08	CNTR	
MOSTOVOI	01	CNTR	0.7
LIAUD	97	TPC	2.5
YEROZLIM	97	CNTR	11.7
BOPP	86	SPEC	4.3
			23.8

(Confidence Level = 0.0002)

-1.29   -1.28   -1.27   -1.26   -1.25   -1.24

$\lambda \equiv g_A/g_V$
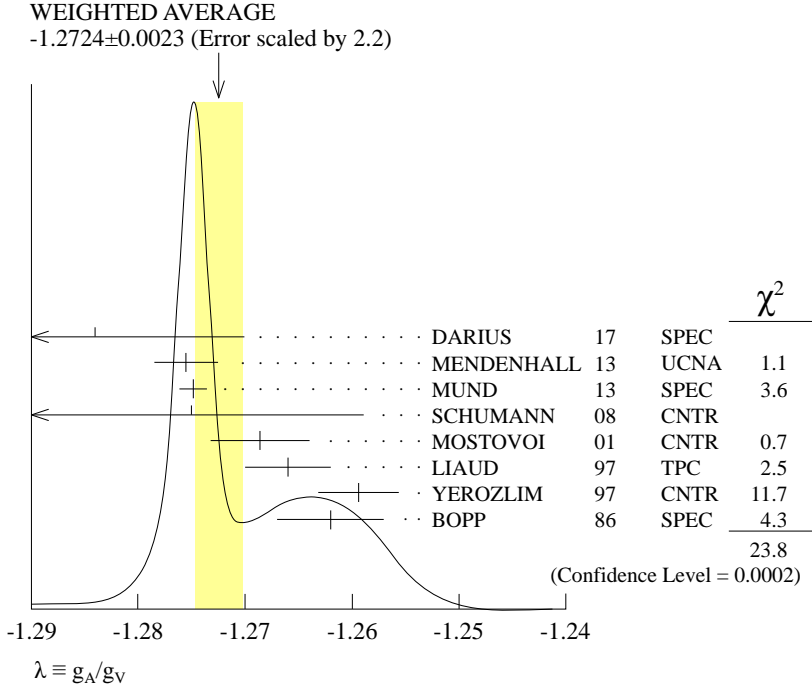
Figure 1.3: The current value of λ as reported in the 2018 PDG [1]. The average value comes from a combination of correlation measurements, and includes a 2.2 scale factor in the error.

Despite the double-peak in values of λ, the most recent experiments all seem to be converging on the left-most peak. As new experiments are developed that can determine λ competitively, old experiments with inconsistent data will become irrelevant.

1.3   Probing for Physics Beyond the Standard Model

The Fierz interference term, b, provides a sensitive observable to search for extensions to the Standard Model, as it is linearly proportional to proposed scalar and tensor couplings. Because the Standard Model excludes these interactions in beta decay at the tree level, any appreciably large value for the Fierz interference term ($\sim 10^{-3}$) would indicate new physics.

6

Unlike the other observables identified in Equation 1.3, b remains after integration over neutron polarization and the electron and neutrino phases spaces. As such, Fierz interference is manifest as a distortion to the observed beta spectrum, especially at lower energies due to its $1/E_e$ dependence. Figure 1.4 shows the variation of the free neutron beta spectrum for different values of b, exaggerated to emphasize the extent of this modification. In addition to the measurement of Fierz interference in
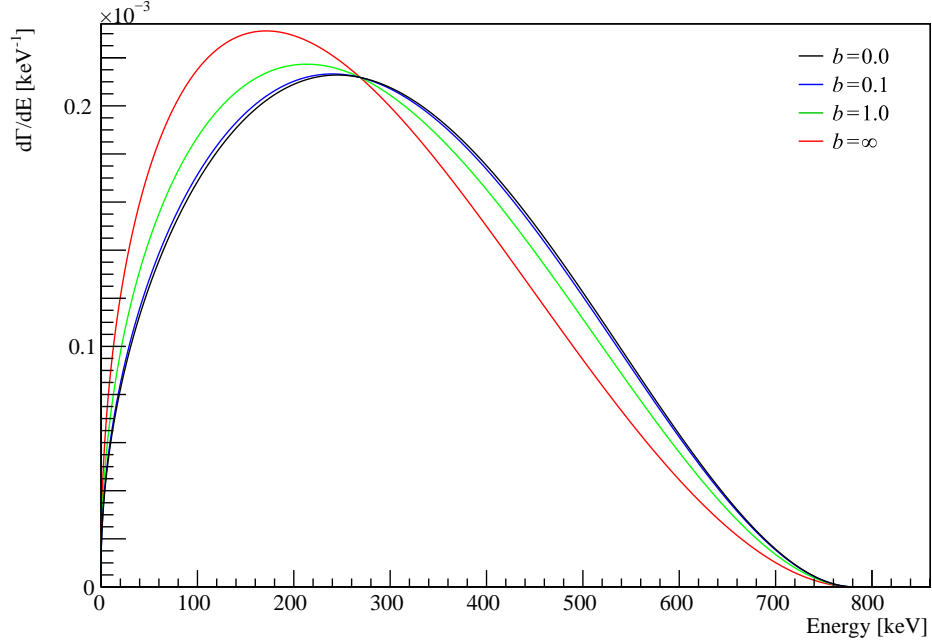
Figure 1.4: Distortion to the beta spectrum for non-zero b. Values of b=0.1,1, and inf are shown for emphasis, but if found to be non-zero, the Fierz interference term will likely be of the order 10^{-3}.

the mixed decay of the neutron, pure Fermi or Gamow-Teller decays can be studied to isolate the scalar or tensor contributions, respectively. As such, literature often distinguishes the different Fierz terms as b_n, which depends on both scalar and tensor couplings, b_F, which only includes scalar terms, or b_{GT}, which is sensitive to tensor currents.

Another measurement available in the search for exotic physics is that of the neutrino asymmetry, B. Apart from its ties to λ, B also depends linearly on any scalar or tensor couplings (see Equation 1.4). Because the term is multiplied by the same m_e/E_e as b_n, this additional term is denoted in literature as b_ν, and the neutrino asymmetry can be written as

$$B = B_0 + \frac{m_e}{E_e} b_\nu, \tag{1.9}$$

where B_0 is the Standard Model expression denoted in Equation 1.4 [21]. As the existence of scalar or tensor couplings would modify all of the correlation terms, it is insufficient to simply extract b_ν by fitting the $1/E_e$ dependence in B. Instead, the

residual $1/E_e$ dependence in the ratio of B/A would negate contributions from b_n, as well as systematics in polarimetry, and isolate b_ν for extraction.

A spectroscopic determination of b_n is exceptionally difficult: large statistical datasets must be accumulated, background events must be thoroughly characterized, and the detector response must be exceptionally understood. Other factors, such as residual nucleon polarization or experimental geometry contributions can further exacerbate this type of study. Consequently, other methods such as the analysis of β-decay $\mathcal{F}t$ values and extraction from the correlation terms a and A have been employed to constrain the presence of scalar and tensor couplings [10, 22, 23]. Limits on the new left-handed couplings from these efforts are

$$\frac{|C_S|}{|C_V|} \leq 0.065 \ \ (68\% \text{ confidence})$$

and

$$\left|\frac{C_T}{C_A}\right|^2 < 0.011 \ \ (95.5\% \text{ confidence}).$$

Recently, constraints on the value of b_n have been determined from the decay of the free neutron using data from the UCNA experiment, and a best-fit value of $b_n = 0.067 \pm 0.005_{\text{stat}} {}^{+0.090}_{-0.061}{}_{\text{sys}}$ was found [24]. The large systematic errors were attributed to the absolute energy reconstruction and spectrometer linearity, which further highlights the inherent difficulty in extraction of Fierz interference measurements via precision spectroscopy. Nonetheless, if these challenges can be overcome, as they should be, the improved sensitivity for determination of Fierz interference via spectroscopy will provide a higher precision search for new physics.

Chapter 2: Three Different Experiments, Three Similar Efforts

This chapter describes three experiments, that on the surface look very different, but are all unified by the their requirements for particle detection. Each experiment is designed to test a different facet of Equation 1.3: UCNB will measure the neutrino asymmetry, B, ^{45}Ca will measure the Fierz interference, b_{GT}, and Nab will measure the electron-neutrino correlation term, a and Fierz interference in neutron decay, b_n. In spite of the different physics tested by each experiment, they all share many of the same challenges.

The first challenge is the ability to detect protons from beta decay. Because of their extremely low kinetic energy (<751 eV) and large mass, protons are often detected via indirect methods. The novel detection system described in Chapter 3 allows for the direct detection of protons. Engineering challenges have required extensive iterations in development. As the UCNB experiment also required observation of protons from beta decay, one of its fundamental objectives was to develop and validate a system for such a purpose. The parallel data acquisition system development also meant that the verification of a low-threshold filter to independently trigger on proton signals could be determined as well. The underlying physics of the experiment provided a number of systematic checks to characterize proton detection in both systems.

The next issue is energy reconstruction. The energy-dependent responses of the detector, electronics, and data acquisition system all need to be thoroughly characterized for the success of Nab. Other effects, such as the multi-pixel deposition of electron energy due to charge sharing or backscattering, must also be considered. Furthermore, bremsstrahlung radiation and losses in the dead-layer cause additional complications for energy reconstruction. Calibration source measurements in UCNB and ^{45}Ca, as well as the measurement of b_{GT}, lead to the development of tools necessary for precision energy reconstruction.

Another requirement for the Nab experiment is a thorough understanding of the uncertainties in timing. While not as stringent in its requirements, the UCNB experiment also demanded tight timing to determine which detector was hit first in the event of multi-detector backscatters. Development of an understanding of the timing for this experiment was critical for ascertaining the challenges that would arise in improving the timing uncertainty by an order of magnitude for Nab.

The last challenge for the two systems is high event throughput. The Nab experiment expects roughly 50 kHz of singles triggers to find a coincidence rate of 500 Hz for electrons and protons. While the number of observed events in UCNB was substantially less than in Nab, the activity of the ^{45}Ca source was high enough that rate-dependent effects were readily apparent. As such, the ^{45}Ca effort served as a high-rate stress test for the detector and especially for the data acquisition system.

The UCNB and ^{45}Ca efforts are interesting in their own right. As with Nab, these experiments measure physics quantities complementary to efforts in both the low and high energy frontiers. The unique geometry of each experiment allowed independent measurements of Standard Model and BSM parameters, which places

even more stringent constraints on exotic interactions.

2.1 UCNB: The Neutrino Asymmetry with Ultracold Neutrons

The UCNB experiment was designed to measure the neutrino asymmetry, B, from the decay of ultracold neutrons (UCN) [25]. Conducted in the existing UCNA spectrometer (SCS) on Line B of the proton linear accelerator at the Los Alamos Neutron Science Center (LANSCE), at Los Alamos National Laboratory (LANL), the UCNB experiment directly measured same-side electron-proton coincidences from the free decay of UCN. Figure 2.1 shows the layout of Area B and highlights the pertinent components, as discussed below.
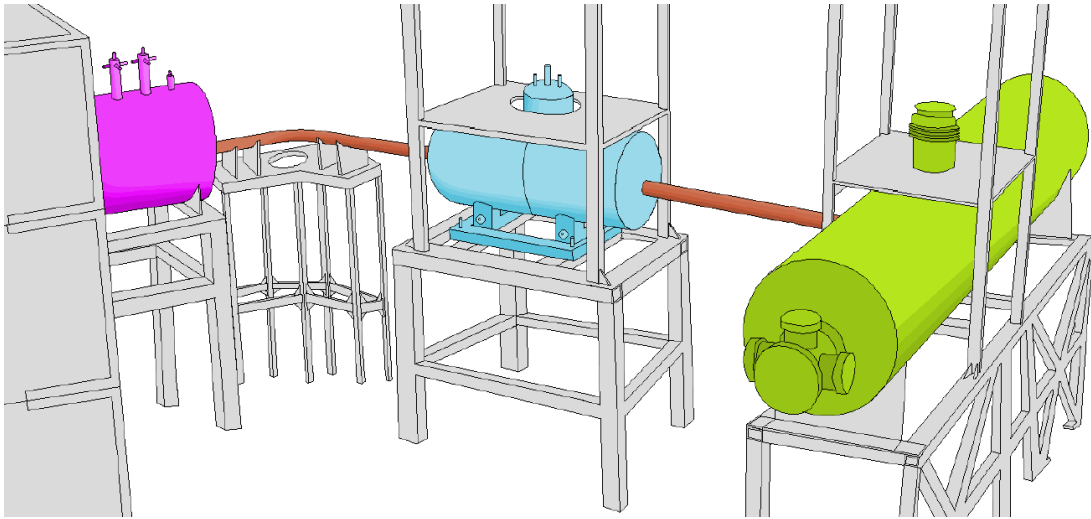
Figure 2.1: A diagram of the three magnets in Area B at LANSCE of the UCNB experiment. Once UCN are produced in the deuterium source, they propogate through a series of stainless steel, copper, and diamond-like carbon coated quartz tubes. The pre-polarizer magnet (magenta) uses a 6 T field to accelerate the UCN through a deuterium-confinement foil, after which they pass through the adiabatic fast-passage magnet (blue), which polarizes and flips the neutron spins to be parallel or antiparallel with the local field in the superconducting spectrometer (green), where they ultimately decay.

The ultracold neutron source at LANSCE is a solid deuterium lattice situated above a tungsten spallation target and encased in a cryogenically cooled polyethylene moderator [26]. As protons impact the tungsten target, neutrons are ejected and enter the hydrogen-rich moderator, where they scatter and achieve thermal equilibrium. These cold neutrons continue until they enter the deuterium volume, where those with sufficiently low energy can downscatter to the ultracold state via phonon exchange with the lattice. UCN are defined as neutrons having less than 300 neV of energy. A guide to the experimental area is situated above the deuterium lattice, and a flapper valve installed above the deuterium prevents UCN from upscattering in the deuterium, which would allow them to escape.

Once above the flapper valve, the neutrons propagate through a series of guides en route to the magnetic spectrometer. The UCN are nominally polarized after passing through the pre-polarizer magnet (PPM), which contains a thin foil to contain the deuterium of the UCN source volume, and are fully polarized once they pass through the 7 T barrier field of the adiabatic fast-passage magnet (AFP), which can apply an RF field to flip the spin of the neutrons relative to the field of the SCS. Successive runs with the neutron spin oriented parallel or anti-parallel to the SCS field are used to extract the neutrino asymmetry with respect to the neutron spin.

In the SCS, the neutrons bounce around until they decay or are lost through mechanisms such as absorption or upscatter. The charged decay products spiral along the magnetic field lines until they reach pixelated silicon detectors at each end. As the charged particles exit the decay region of SCS and approach the detectors, the magnetic field expands from 1 T to 0.6 T, so that the particles' momenta are longitudinalized, reducing the probability to backscatter. While the electrons carry sufficient energy to produce an easily detectable signal, the low-energy protons require an accelerating potential to penetrate the silicon dead-layer and generate sufficient ionization within the depletion region of the detectors for measurement. Electrodes just past the expansion region of the SCS impart the protons with 30 keV of kinetic energy for direct detection within the silicon. Figure 2.2 shows the SCS geometry and fields for the UCNB experiment.
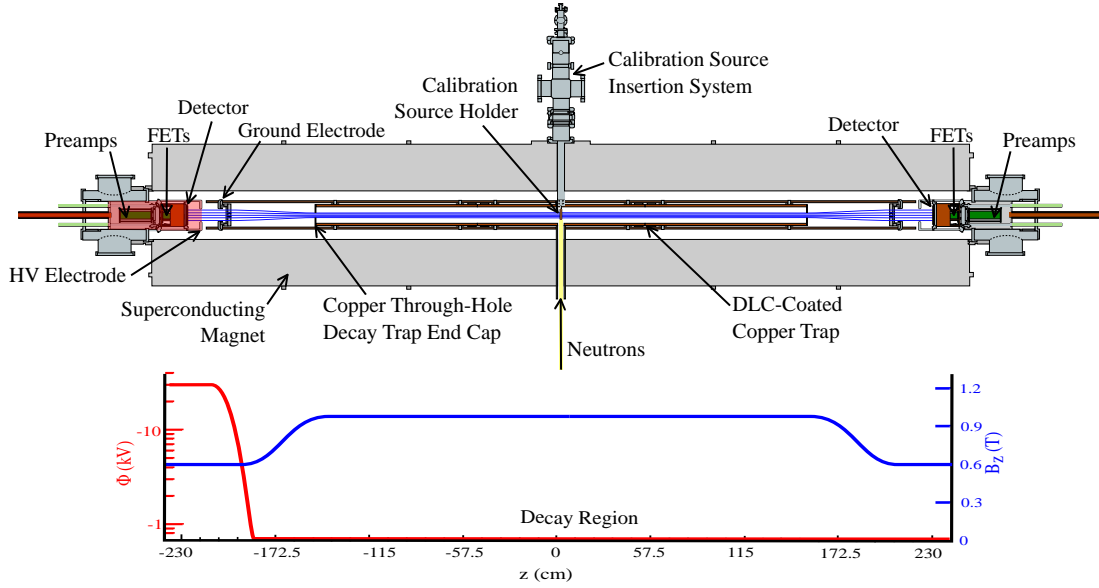
Figure 2.2: The SCS in the UCNB experiment.

As in the PERKEO experiment, the neutrino asymmetry is found via the detection of same-side electron-proton coincidences, where the experimental asymmetry [19]

$$B_{exp} = \frac{N^{--} - N^{++}}{N^{--} + N^{++}} \tag{2.1}$$

is related to the physics asymmetry via

$$B_{exp} = \frac{4P}{3} \begin{cases} \frac{A\beta(2r-3)+B(3-r^2)}{8-4r+a\beta(r^2-2)}, & r < 1 \\ \frac{-A\beta+2Br}{4r-a\beta}, & \geq 1 \end{cases}.$$ (2.2)

In the above expression, $r = \beta(E_e + m_e)/)E_{max} - E_e)$, with the usual definition $\beta = v/c$. In Equation 2.1, the signs indicate the whether the electron and proton emissions are in (+) or against (-) the neutron spin direction. Because the electron and proton emissions are also correlated with the neutron spin, Equation 2.2 relates the neutrino asymmetry to the experimental asymmetry. Unlike PERKEO, however, the protons will be directly detected in the UCNB experiment.

The goal of the UCNB experiment is to determine the neutrino asymmetry to a precision of $\delta B/B = 1 \times 10^{-3}$. This is sufficient to provide a measurement competitive with the current PDG value, but would provide little insight into any new physics. In order to determine a meaningful value of b_ν, statistics on the order of 10^9 decays are required, which is a significant challenge given the low event rate of UCN. The recent source upgrade marks a significant step towards achieving this goal, but additional efforts to improve the experiment will be needed [27].

2.2 ^{45}Ca: Probing for Tensor Couplings with Nuclear Beta Decay

The isotope ^{45}Ca β decays with an endpoint energy of 256 keV and a lifetime of just under 6 months. [28]. Figure 2.3 shows the decay scheme for ^{45}Ca. The decay is predominantly the $7/2^- \to 7/2^-$ transition between the two ground states, with a branching ratio of 99.9981(11)%, although a very small Fermi component could be present due to electromagnetic interactions [29]. The relatively low endpoint energy and nearly pure GT transition mean that ^{45}Ca is a sensitive probe of tensor coupling via the β spectrum shape.
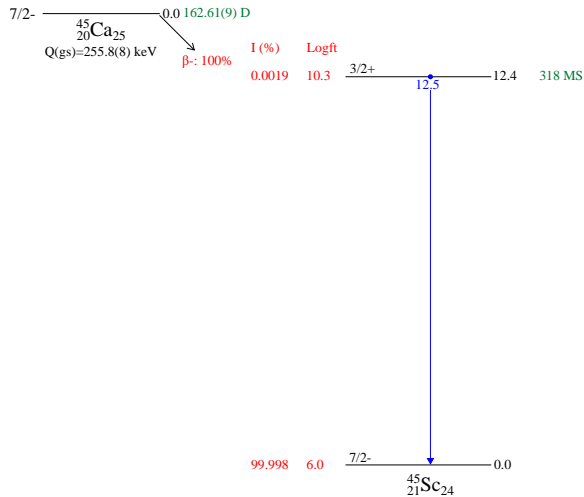
Figure 2.3: The decay scheme for ^{45}Ca.

As with the free neutron, a non-zero value for the Fierz interference term would present itself as a distortion to the beta spectrum. As a heavy nuclei, however, several

corrections to the expected shape must be included to fully describe the expected shape and possible contribution from tensor couplings. Hayen and others presented a detailed description for the necessary corrections to the spectral shape, followed by a generator to produce the appropriate spectra [30,31]. Figure 2.4 shows the calculated spectrum using this generator for ^{45}Ca with several large values of b_{GT} for emphasis.
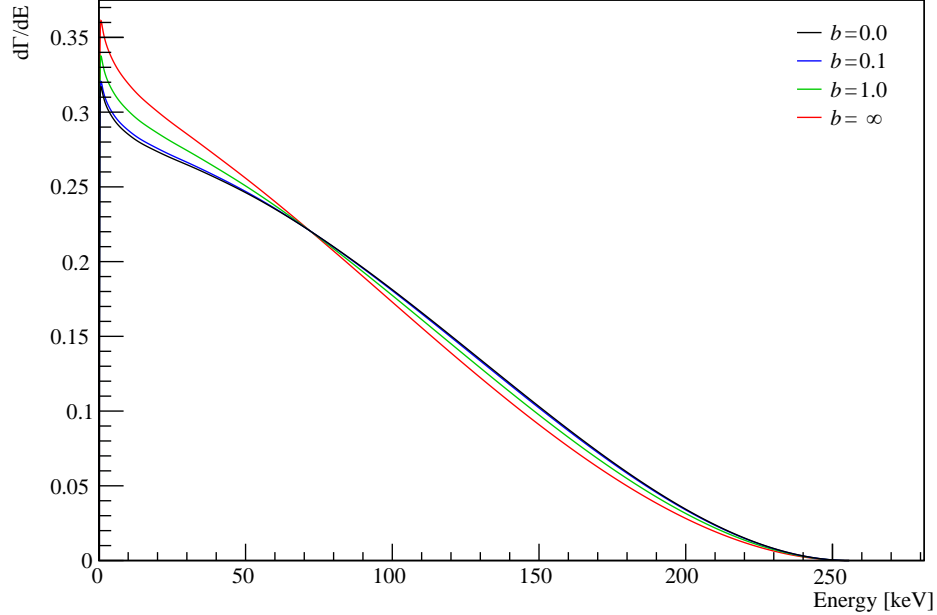
Figure 2.4: Fierz interference in the ^{45}Ca beta spectrum. Measurement of non-zero b_{GT} would be indicative of the presence of tensor currents.

The ^{45}Ca source preparation was a joint effort by groups at North Carolina State University and The Katholieke Universiteit Leuven. A 500 nm aluminized Mylar film mounted on a 2.54 cm aluminum ring was treated with an acid deposition of a calcium solution on its bare side and allowed to dry. The nominal activity of the source was about 1 kBq, allowing for observation of 10^8 decays in just over a day of continuous running. The ring was mounted in a repositionable plastic holder and inserted into the SCS using the guides for the decay trap components. The source was located approximately 71 cm towards the west detector, relative to the center of the decay trap. As with UCNB, calibration sources were inserted through the load-lock for guided electron calibration and were also mounted on the ends of the high-voltage caps for full detector illumination with photons and electrons. Figure 2.5 shows the spectrometer configuration for the ^{45}Ca experiment.

The goal of ^{45}Ca is to establish a limit on the Gamow-Teller Fierz interference term to a level of $\delta b_{GT} \leq 1 \times 10^{-2}$ via a precision spectroscopy measurement. To achieve this, several sets of 10^8 decays will need to be acquired for a full understanding of the spectrometer systematic effects. Measurement at this level would improve current limits on tensor interactions and contribute substantially to the search for new physics.
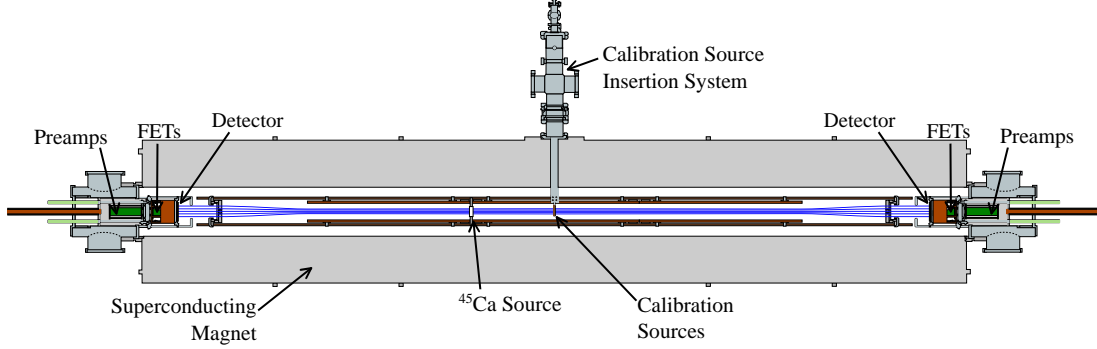
Figure 2.5: The SCS configured for the ^{45}Ca experiment. By centering the ^{45}Ca source in the center of the magnetic spectrometer, 4π acceptance of beta particles is achieved. Furthermore, any backscattered electron energy can be fully reconstructed by summing opposite pixels.

2.3 Nab: Precision a and b Determination

The Nab experiment, installed on the Fundamental Neutron Physics Beamline (FNPB) at the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory (ORNL), aims to make the most precise measurement of the electron-neutrino correlation term, a, and in doing so, improve the precision of the determination of V_{ud} from the neutron measurements. Additionally, a new limit on the Fierz interference term in neutron beta decay will be established from the neutron beta spectrum shape.

The Nab experiment will directly measure the electron-proton phase space in order to calculate the correlation in angle between the electron and antineutrino. Through conservation of momentum, $\cos\theta_{e\nu}$ depends entirely on the proton and electron momenta:

$$p_{\mathrm{p}}^2 = (\vec{p_e} + \vec{p_\nu})^2 = p_{\mathrm{e}}^2 + 2p_e p_\nu \cos\theta_{e\nu} + p_\nu^2. \tag{2.3}$$

This determines the allowed phase-space for the electron and proton, where the upper and lower bounds of the proton momentum squared correspond to the electron and neutrino traveling parallel or anti-parallel to each other, respectively. The distribution of allowed proton momenta for a given electron energy is trapezoidal in shape, where the slope along the top of the trapezoid is proportional to the electron-neutrino correlation term. The extraction of a from several of these trapezoids allows for a self-consistency check within the experiment. Figure 2.6 shows the allowed phase space for the electron and proton from free neutron beta decay, and several corresponding trapezoids for different electron energies.

To determine a, the Nab experiment will directly detect coincident electrons and protons from the neutron decay. Unpolarized cold neutrons pass through a magnetic spectrometer, where some portion will decay. As with UCNB, each coincidence event is comprised of a prompt electron signal followed by a proton within 100 μs. Unlike UCNB, however, the neutrons decay within a small volume, and thus the proton time-of-flight (TOF) can be used to ascertain its initial momentum.

Directly above the decay volume is a 4 T magnetic filter, which preferentially
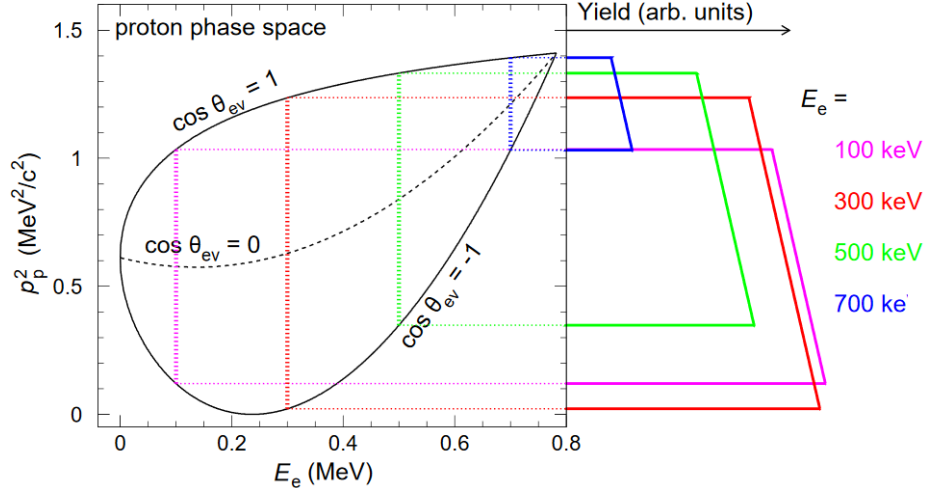
Figure 2.6: The allowed phase space for the coincident proton and electron in beta decay. The upper and lower limits to the proton energy occur when the neutrino is emitted parallel or anti-parallel, respectively, to the electron. The distribution of proton energy for a specific electron energy yields a trapezoidal distribution, where the slope at the top is proportional to the electron-neutrino correlation term, a.

accepts the upward particles with low transverse momentum, before adiabatically transitioning to the 4.5 m long TOF region. The momentum is further longitudinalized in this region via a field expansion to 0.1 T. A -30 kV electrostatic potential is applied after the TOF region to accelerate the protons for detection. A full experimental diagram is shown in Figure 2.7.

In addition to a, the Nab experiment is also designed to measure the Fierz interference term for the neutron, b_n. Much like ^{45}Ca, the electron spectrum will be fit to constrain scalar and tensor couplings. The spectrometer allows for full reconstruction of backscattered electron energies, and detection in coincidence with the proton will prevent background events from contaminating the beta spectrum. In the b configuration, the lower detector will be floating at -30 kV to accept protons in the closer of the two detectors, maximizing the overall $e - p$ coincidence rate.

The Nab experiment is designed to achieve a relative resolution of $\delta a/a = 0.1\%$ for the electron-neutrino correlation term. With a measurement of this precision, a value of λ competitive with A measurements can be achieved, which will reduce the experimental uncertainty for the extraction of V_{ud} from neutron beta decay. Additionally, the Nab experiment is targeting an absolute upper limit of $\delta b = 3 \times 10^{-3}$ for the Fierz interference term in free neutron decay.
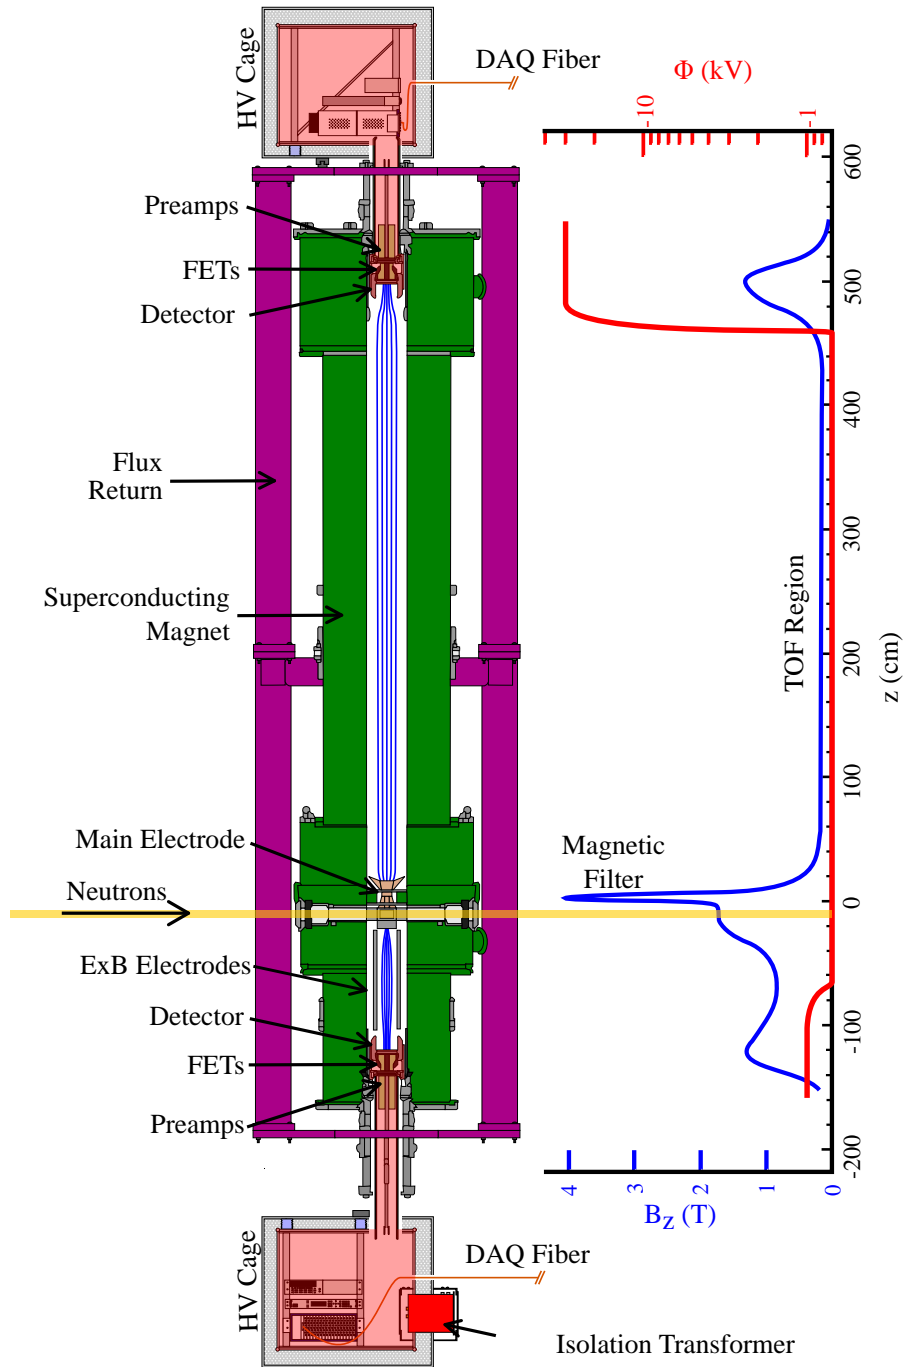
Figure 2.7: The Nab experiment on the Fundamental Neutron Physics Beamline at the Spallation Neutron Source in Oak Ridge, TN. Some neutrons entering the spectrometer decay, producing a proton, electron, and antineutrino. The charged decay products are guided to detectors above or below the beamline; protons are only measured in the upper detector, where an electrostatic field imparts sufficient energy for detection, while electrons are measured in either detector.

16

Chapter 3: Precision Particle Detection in Silicon

All three experiments carry the common requirement for low-threshold particle detection and extremely precise energy reconstruction over the range of $10 - 800$ keV. For UCNB, the low-threshold efficiency requirement is for direct detection of the co-incident protons from beta decay, while ^{45}Ca is most sensitive to the new physics at the lower energies of the beta spectrum; Nab is subject to both of these conditions. Furthermore, the relatively small distortion to the beta spectra shapes in Fierz interference searches in ^{45}Ca and Nab leads to the requirement that event-by-event energy reconstruction must have a linearity of 0.01%.

This chapter discusses the detector system hardware, including the high voltage assemblies used to bias the detector for proton detection. Additionally, simulation efforts were developed to understand the detector response and the subsequent analysis of the digitized waveform traces. Two methods, a Monte Carlo model to incorporate the physics of the detector response and a lightweight algorithm for the generation of large statistics, are presented.

3.1 Detector Hardware

The silicon detectors, manufactured by Micron Semiconductor, Ltd., are large area, thick, highly pixelated silicon PiN diodes. Each detector is a single 12.35 cm-diameter crystal oriented in the $\langle 100 \rangle$ plane, and implanted to form a PiN diode. The wafers were manufactured in thicknesses of 1.5 and 2.0 mm to stop all electrons across the neutron beta spectrum. The silicon bulk is ultra-high resistivity (nominally 25 k$\Omega \cdot$cm), phosphorus-doped silicon. To achieve such high resistivity silicon, float-zone crystal growth is used to constrain contaminants to the sub parts-per-trillion atom level.

The junction side features a shallow implant of boron to a depth of 100 nm for minimal dead-layer energy loss. The ohmic side is comprised of 127 hexagonal pixels of area 70 mm^2, developed from a deep phosphorus implant. Additionally, partial pixels comprise the remainder of the 11.7 cm active diameter. The pixel pitch is 100 µm, in which a p-spray is applied to provide high-resistance inter-pixel separation. The hexagonal pixels are metalized with aluminum for electrical contact with the readout electronics, and the junction side is covered with a 300 nm thick aluminum layer in a square grid geometry. This grid is comprised of 10 µm-thick lines spaced 4 mm apart, resulting in a 0.4% coverage of the active junction window.

For each detector, the silicon wafer is affixed to a ceramic package on the ohmic side. Holes in the ceramic provide a location through which wires are bonded from the hexagonal pixels to traces that route the signals to one of four 32-pin connectors for electronic readout. The partial pixels around the circumference of the active area of the silicon were originally intended to be single readout region, but the decision was made to ground them to minimize noise contributions. Figure 3.1 shows a schematic

of the detector for the UCNB and ^{45}Ca experiments; the geometry will be similar for Nab, although readout will be through pogo-pins instead of the multi-pin connectors.

(a) Ohmic side (bare silicon).

(b) Ohmic side with ceramic.

(c) Junction side.

(d) Side view.
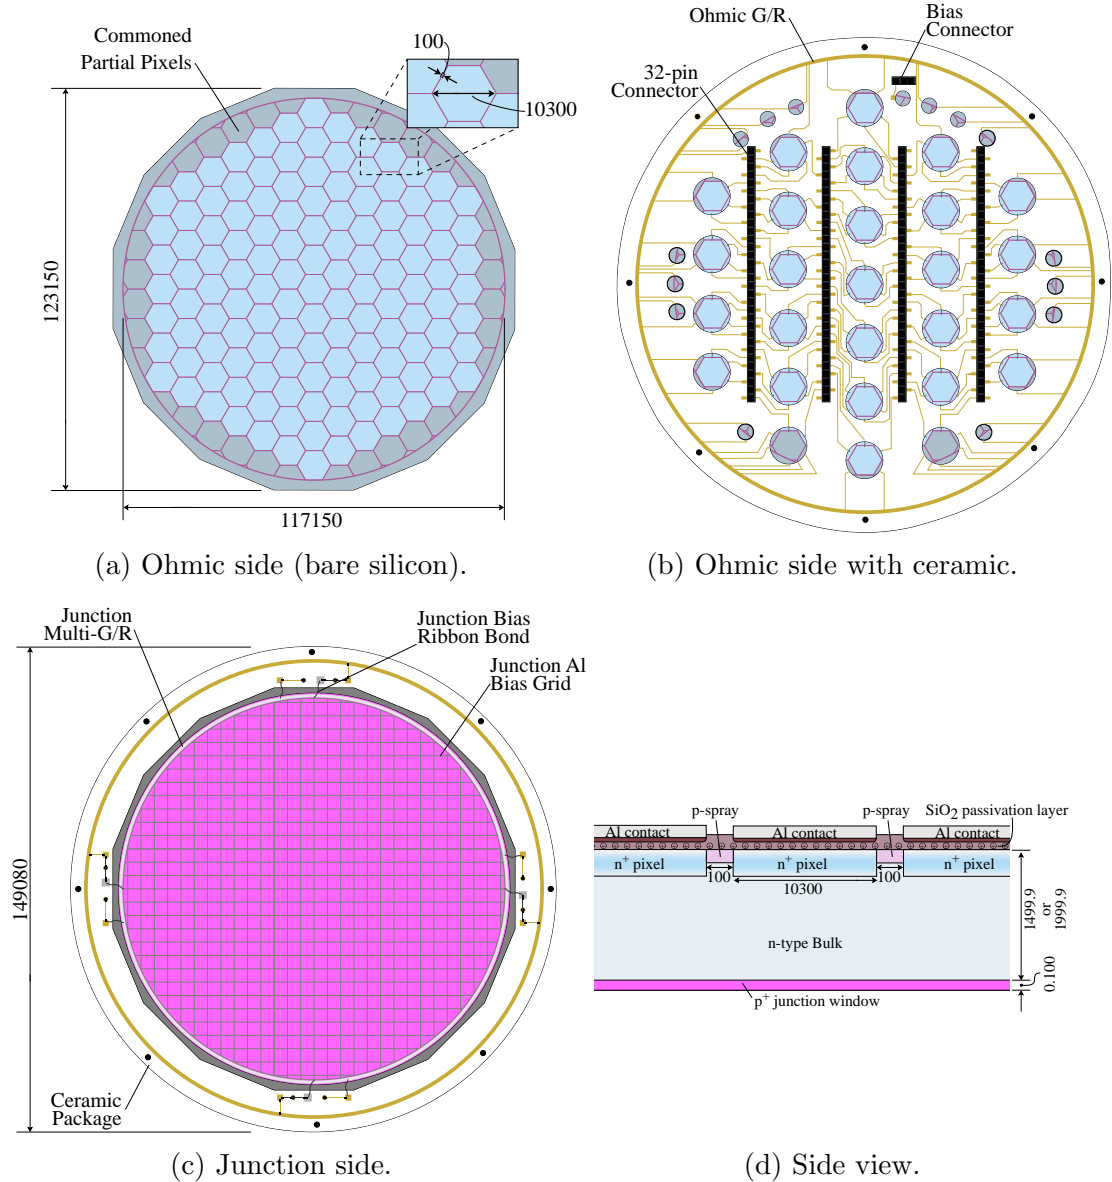
Figure 3.1: The silicon detector geometry, where all dimensions are in microns.

For the UCNB runs, the detector was biased with five A23 batteries for a nominal bias of -60 V to the junction side of the detector. The batteries were housed in an aluminum project box, which included a series-connected 100 kΩ, 2 nF low-pass filter for bias-line stability. A switch installed after the project box allowed for disconnection of the bias line, so that the batteries would not discharge as the detector warmed.

During the ^{45}Ca runs, the battery packs were upgraded to a Keithley 6487 DC voltage source and picoammeter. The unit combines a ±500 V DC power supply with a precision ammeter capable of reading the leakage current through the detector. By

monitoring this current, estimates of the detector temperature and its overall health can be determined remotely. Furthermore, a larger bias voltage can be applied across the silicon to improve the efficiency of charge collection for better precision in timing.

To assess the optimal bias potential a series of tests were conducted to study the effects of over-depletion in the 1.5 mm silicon. First, the baseline noise was examined as a function of bias potential. With increased voltage, the leakage current through the silicon can increase, thus causing additional noise concerns. Also, the structure of the noise can cause an apparent shift in the baseline, which can cause issues with energy resolution. Baseline traces of length 262 μs were recorded to analyze the correlation between the baseline noise and bias potential. As shown in Figure 3.2, the average DC offset was found to be independent of the depletion bias, while the RMS noise increased beyond −400 V; analysis of the baseline traces in the frequency domain confirmed this increase in noise.

(a) Baseline offset and RMS amplitude   (b) Baseline power spectrum
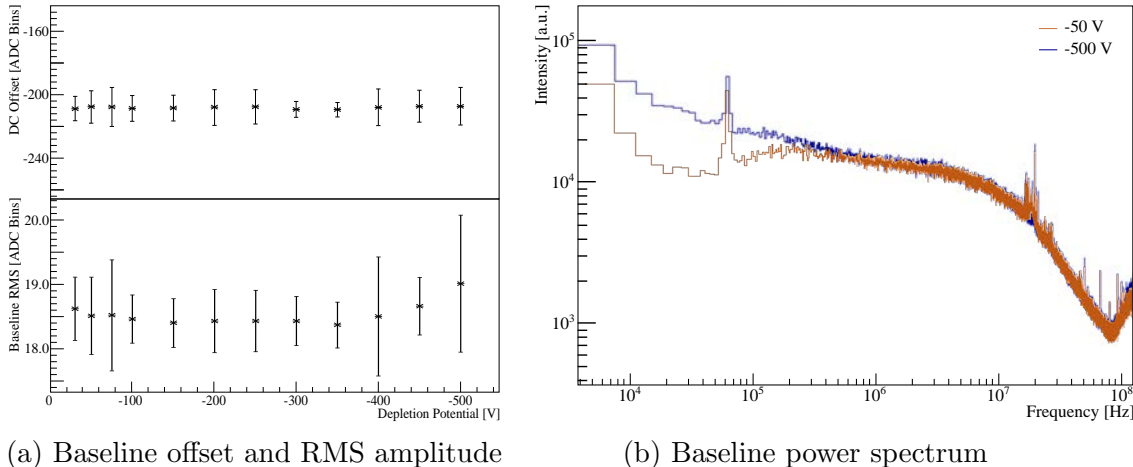
Figure 3.2: The baseline noise characteristics are largely independent of the detector bias potential in the range of -30 to −500 V. At larger potentials, the baseline fluctuations begin to increase in amplitude, likely due to an increase in leakage current across the silicon; the power spectrum shows the prominent frequencies for the additional noise.

Next, the peak resolution from the decay of ^{207}Bi was studied as a function of bias potential. Table 3.1 summarizes the prominent electron peaks from the decay, and their intensities. Additionally, a 569.698 keV gamma, which results in a Compton edge at 393 keV, is produced in 97.75% of decays. Figure 3.3 shows the measured spectra at three different bias potentials, and Figure 3.4 shows the effect of increased bias potential on the 481 and 975 keV peaks.

At low voltages, the silicon was only partially depleted, so the Compton edge is less pronounced and the high-energy electron peak is broadened due to incomplete charge collection. Furthermore, the low drift field resulted in lower overall carrier transport efficiency, so the peaks at 481, 553, and 565 keV are also widened. Increasing the voltage, the charge collection efficiency and depletion width both increase, so the resolution of all peaks improves, which allows for discrimination between the 553 and

565 keV electron lines. However, as the voltage continues increasing, the increased baseline noise results in minor degradation to the energy resolution, so the peaks begin to coalesce again.

Table 3.1: A selection of prominent electron peaks from the decay of [207]Bi and their intensities per decay, where the italicized digits indicate uncertainty in the respective value.

Energy (keV)	Intensity (%)
481.6935 *21*	1.537 *22*
553.8372 *21*	0.442 *6*
565.8473 *21*	0.111 *5*
975.651 *3*	7.08 *17*

(a) -30 kV   (b) -350 kV   (c) -500 kV
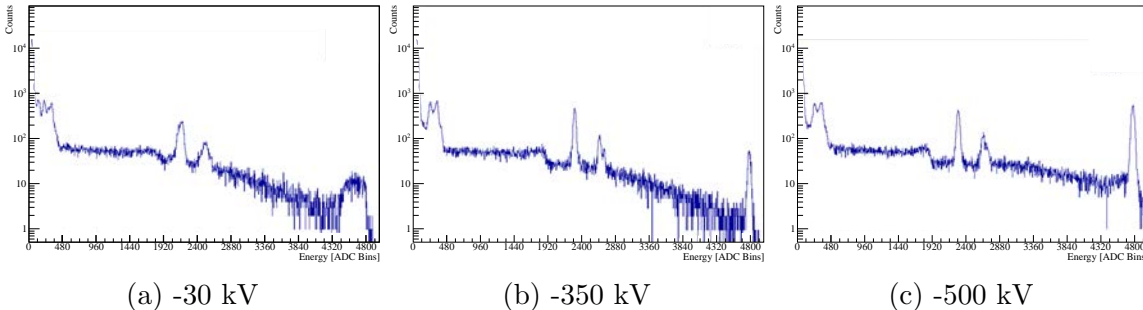
Figure 3.3: [207]Bi spectrum at three different bias potentials. At -30 kV, the peaks are broad and poorly defined; the 975 keV electron peak is especially broad, most likely due to the silicon only being partially depleted. At -350 kV, the peaks are much sharper, so much so that the 553 and 565 keV lines can be discerned as distinct peaks. At -500 V, the peaks maintain their structure, but the resolution is slightly degraded, as apparent in the melding of the peaks at 553 and 565 keV.

Lastly, the effect on charge carrier transport was considered via examination of the waveform rising edges. As discussed in Section 3.4, the mobility of the charge carriers increases with an increased drift field across the silicon, resulting in a decrease in the charge collection time. This effect can be observed by comparing the average 10-90% rise times for fixed amplitude pulses at different bias potentials. Figure 3.5 shows a two-dimensional histogram for 481 keV waveforms under a detector bias of -30 V and -500 V, as well as the average rise time versus applied potential. Beyond -200 V, no appreciable decrease in signal rise time can be observed.

As a result of these various measurements, it was determined that the original bias voltage of -60 V was inadequate for full depletion of the silicon. Partial depletion resulted in signals with longer rise times and broadened peaks. Similarly, the increase in noise beyond -400 V limits the energy resolution for the detector. Based on a balance of signal rise time and peak width, the optimal bias time is expected to be within -250 and -350 V.
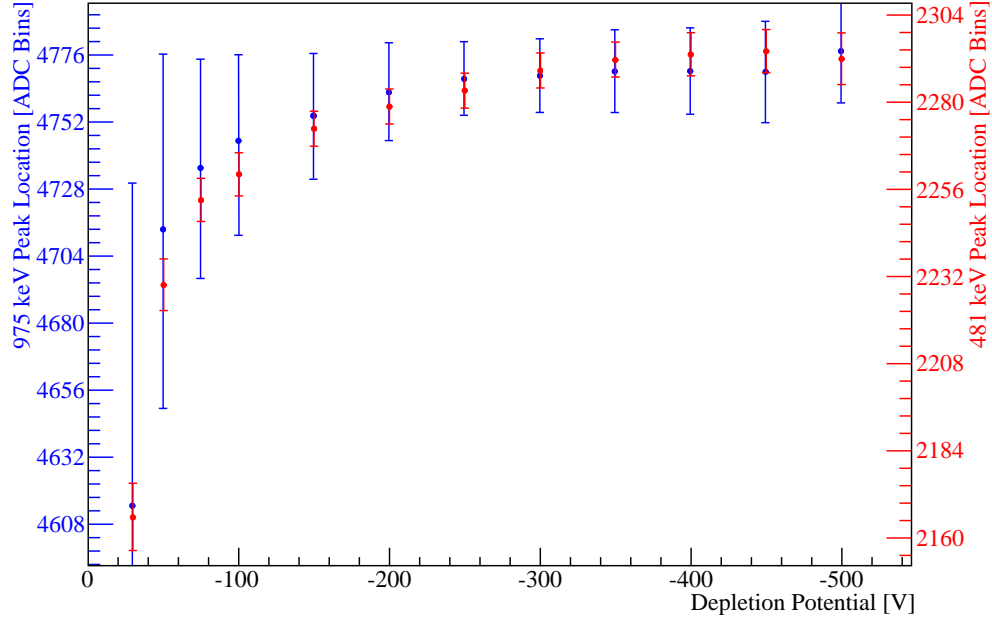
Figure 3.4: Shown are the peak locations (circular marker) and 1σ peak widths (vertical bars) for the 975 (blue) and 481 keV (red) electron peaks from the decay of ^{207}Bi as a function of depletion potential. For both lines, the peak location increases and the peak width decreases with increasing bias potential due to an increase in charge collection efficiency. The especially wide peaks for the 975 keV lines at lower bias potential are a result of partial detector depletion.

(a) 481 keV waveforms

(b) $10 - 90$ rise times
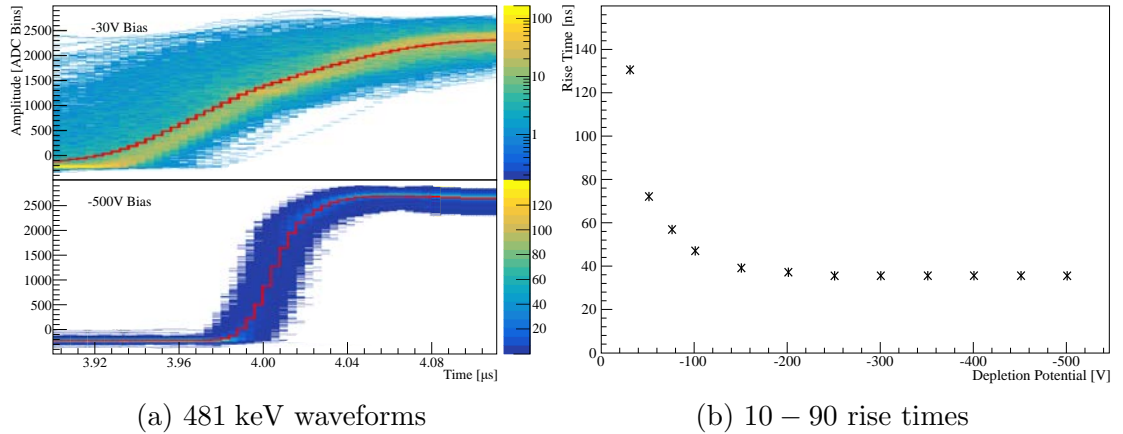
Figure 3.5: The charge carrier mobility increases with an increased depletion field, resulting in faster signal rise times. After roughly -200 V, the charge collection time does not noticeably improve.

3.2 Front-end Electronics

Signal readout is accomplished via a charge-sensitive amplifier followed by signal shaping electronics designed and developed by staff scientists at LANL [32]. The ohmic current for each pixel is integrated in a cryogenically-cooled BF862 low-noise

n-JFET with a 1 pF feedback capacitor and a 1 GΩ discharge resistor, which provides an effective integration time of 1 ms for full charge collection. The resultant voltage is then amplified by an AD8011 op-amp,followed by shaping by a CR-$(RC)^2$ filter. As the pre-amplifier and shaping stages are unable to operate at the low temperatures required by the detector and FET assembly, they are situated on the other side of a vacuum feedthrough. Figure 3.6 depicts the circuitry used for the UCNB and ^{45}Ca experiments.

(a) Prototype FET and amplifier boards

(b) Circuit diagram for the electronics

Figure 3.6: The readout electronics used in the UCNB and ^{45}Ca experimental runs. The electronics are divided into two boards - the cold FET electronics (blue) integrate the current induced on the ohmic side, and the warm preamps (red) shape the pulse before digitization.

The FET boards are designed as a single-sided layout, with 8 channels of electronics per board. The preamp cards are double-sided and contain 6 channels of electronics per board. The UCNB and ^{45}Ca setups used 3 FET and 4 preamp cards, for a total of 24 readout channels. Initially, the 19 central pixels and 4 channels of ganged pixels were instrumented for UCNB. It was discovered that a capacitive mismatch between the grouped pixels and the FETs produced overwhelming noise in the entire system, so they were short-circuited to ground in order to remove the noise.

The CR-$(RC)^2$ shaping parameters were selected to optimize performance with regards to several issues. During the UCNB runs, the pulse fall-time was chosen to be 1 μs to allow for discrimination of pile-up events while attenuating high-frequency

noise contributions for improved energy resolution, and the circuit rise time was ~20 ns to balance electronic stability and timing performance. However, the long shaping time relative to the signal fall-time, which could lead to energy resolution degradation via effects such as ballistic deficit, prompted a study into the optimal time constant to maximize the energy resolution of the electronics.

The sensitivity to noise can be decomposed into three contributions, as described by Bertuccio and Pullia [33]:

- Series white noise, which increases with detector capacitance and is inversely proportional to the electronic shaping time

- Parallel white noise, which results from FET shot noise and increases linearly with shaping time

- $1/f$ noise, which arises from issues such as random carrier fluctuations within the FET and is independent of shaping time

To determine the contribution from each, the energy resolution of a fixed-amplitude peak is examined as a function of the shaping time τ_S of a trapezoid filter. The data is fit to account for each contribution with

$$ENC^2 = \frac{h_1}{\tau_S} + h_2 + h_3\tau_S, \tag{3.1}$$

where the three coefficients $h_1, h_2,$ and h_3 relate the relative intensities of the series, $1/f$, and parallel noise respectively. The data and fits are shown in Figure 3.7, and the results are summarized in Table 3.2.

Table 3.2: The fit values for $h_1, h_2,$ and h_3 from Equation 3.1, as well as the calculated minimum ENC, where the arbitrary units are proportional to the unit charge, e, by a fixed constant.

Fall Time [μs]	h_1 [a.u.]	h_2 [a.u.]	h_3 [a.u.]	ENC_{min} [a.u.]
1.0	0.00502	0.0235	0.00478	0.182
4.0	0.00566	0.0241	0.00107	0.170
8.0	0.00669	0.0244	0.00067	0.169

Based on the results of these studies, the signal fall-time was extended to 5 μs to provide improved energy resolution via the mitigation of the parallel noise contributions. This change was implemented for the ^{45}Ca runs, and effectively reduced the noise contribution by about 7%. As can be seen in Figure 3.7 and Table3.2, however, the main contribution to the noise is the $1/f$ component, which indicates that work should be done to mitigate noise originating within the FETs.

The FET boards require +12 V and 15 mA per channel, while the warm amplifiers require ±6 V inputs with ~60 and 65 mA per channel. The electronics are powered using a Keysight N6700B low-profile modular power system mainframe, which is populated with one N6733B 20 V, 2.5 A and two N6732B 8 V, 6.25 A DC supply
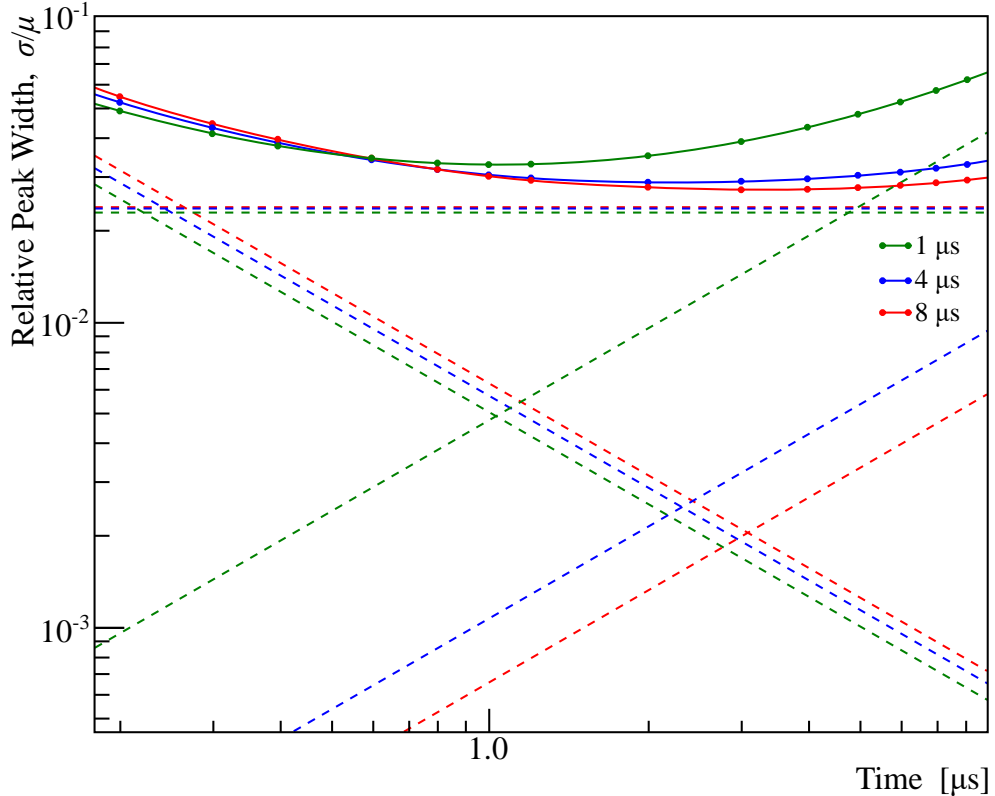
Figure 3.7: A shaping time scan was conducted on the prototype electronics with fall times of 1, 4, and 8 μs to determine the appropriate time constant. The solid lines represent a least-squares fit of the data to Equation 3.1, and the dashed lines represent the individual components of the noise (series, parallel, and $1/f$) for each fit. Table 3.2 contains the fit parameters, as well as the calculated minimum ENC for each electronic fall time.

modules. With the increased channel density for Nab, the N6732B supplies will be upgraded to N6773A 20 V, 15 A to accommodate the additional current draw.

In 2016, a high-frequency pulser circuit was added after the bias-line low-pass filter on the junction side of the silicon. This circuit allowed for the direct injection of a capacitively coupled pulse on the bias line, which produces a response on the entire detector array to provide feedback for timing and energy systematics, as well as a general assessment of the overall detector health. Figure 3.8 shows the complete detector bias line circuit with the addition of the pulser.

A Berkeley Nucleonics PB-5 precision tail pulse generator was used for injecting the signal onto the bias line. The pulse generator was selected for its specification of 35 ppm full-range integral non-linearity and fast rise time of 50 ns. Because incoming voltage pulses are capacitively coupled to the bias line, long tail pulses were favorable as they resulted in a unipolar output from the detector system. The pulse fall time was selected as 1 ms, and the polarity was chosen to be negative to match the real physics of electron current collection.
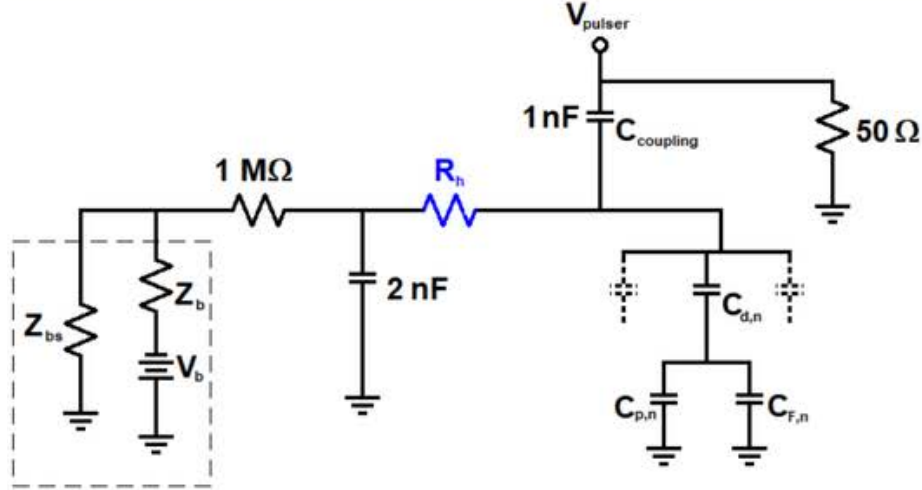
Figure 3.8: The full detector bias line circuit, including the pulser. Originally, the coupling capacitor was 1 nF, but after finding that the noise injected along the bias line degraded the detector resolution, it was modified to be 100 pF instead. Furthermore, R_h was accidentally omitted in the initial electronics, causing unexpected behavior; it was added back in for the Nab circuit to provide the correct response.

3.3   High Voltage Assembly

To float the silicon detector at -30 kV for proton detection, all of the associated electronics (amplifiers, DAQ, power supplies, temperature sensors, etc.) must also float at high voltage. An assembly, constructed as an inner and outer Faraday cage at high voltage and ground, respectively, is used to house the electronics for the experiment. The detector cabling is run through an insulated, conductive conduit from the inner Faraday cage to the ceramic and vacuum isolated detector mount. For UCNB and ^{45}Ca, the cryogenic lines were adapted with G-10 insulating breaks to cross the high voltage barrier at the detector mount; for Nab, these breaks will occur between the walls of the two Faraday cages.

Two prototype high voltage assemblies were prepared at LANL as part of the UCNB experiment. For the first, the inner "box" was an open frame constructed from copper tubing and the outer box was a solid aluminum cart, with the front open for access to the electronics. 1.25 cm polyethylene and polycarbonate slabs were inserted between the tube frame and the aluminum cart as the high-voltage standoff, and a detachable acrylic face was placed over the opening for safety. To dissipate the heat of the electronics, vent holes were cut in the bottom of the aluminum cart, and two ventilation fans were installed in the top. A 5 cm diameter copper tube surrounded by thick-walled PVC insulation connected the high voltage assembly to the detector mount, and carried the signal cables between the detector electronics and the hardware in the box.

Table 3.3 details the electronics that are contained within the high voltage assembly and their respective power requirements. This equipment list was used to determine the minimum size for the inner box, as well as the requirements for the

power consumption for each. In UCNB, only portions of the system were run, so components such as the data acquisition system and amplifier power supply drew less current. Additionally, only the essential elements were used for the experiment, so additional items for temperature monitoring and total remote communication were omitted.

Table 3.3: Electronics to be installed in the high voltage system.

Electronic Component	Max Power Consumption (W)
Data acquisition system	1133
Keysight N6700B (amplifier power supply)	1000
Keithley 6487 (detector bias and picoammeter)	50
Lakeshore 224 (temperature monitor, Nab only)	35
Compact NIM crate (Nab only)	227
Highland Tech. J730 (3x) (optical relay, Nab only)	2.4 (ea.)
RS232 to ethernet	41.4
5-Port ethernet hub (Nab only)	7.5
Ethernet to fiber converter	10
Total	**2511**

Power was delivered to the electronics within the high voltage assembly via a custom-made Stangenes SI-24769 isolation transformer, shown in Figure 3.9. For UCNB, the transformer was rated at 50 kV standoff potential and 1 kVA power draw, as only a portion of the above system was implemented. For Nab, the transformer was upgraded to a 3 kVA rating. To protect the coils of the transformer from corrosion, the whole assembly was cast in epoxy. The primary coil was wired through a fused disconnect switch for the safety of the user, while the secondary coil was connected to a power strip. In the Nab box, a low-pass filter with a 5 Hz cut-off frequency was added to the secondary coil for the prevention of high voltage breakdown through the electronics as shown in Figure 3.11.
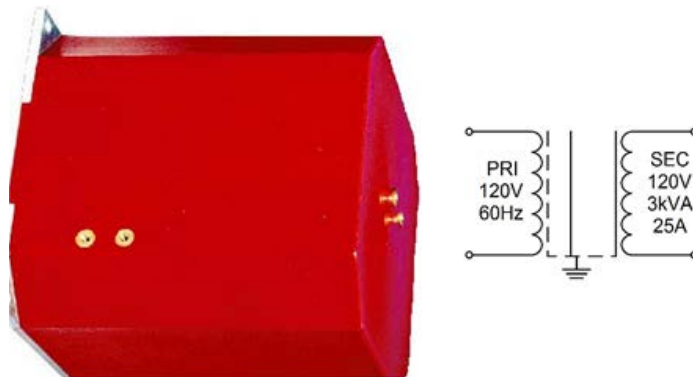
Figure 3.9: Stangenes SI-24769 isolation transformer.

For Nab, the potential of the inner box will be controlled by a Matsusada K12-30R-LC high voltage power supply, shown in Figure 3.10. The supply is rated to

30 kV DC, in both positive and negative polarities, and is controlled remotely via digital/analog I/O. The ground terminals of the supplies are connected to the SNS grounding plate. The supply is current limited to 0.4 mA, but an additional 1 GΩ resistor in series with the supply further limits the current to 0.03 mA. A 1.5 nF shunt capacitor bridges the connection between high voltage and earth-ground to reduce high frequency ripple in the system, to prevent instability and noise in the detection electronics.

The supply has an internal capacitance of approximately 1.5 nF, which stores about 0.6125 J when fully energized. Although the literature for the power supply does not characterize a discharge time nor an internal resistance, the polarity reversing time is referenced as 2 s [34]. Based on this information, the inherent discharge time for the power supply is assumed to be much less than 1 second.

Figure 3.10: Matsusada K12-30R-LC high voltage power supply.

In addition to the components already mentioned, a ground stick for safely discharging the de-energized high voltage assembly is attached to ground. The ground stick has a long metallic probe connected to ground through a resistor, to "softly" ground the inner high voltage box before removing the cover, and a hook directly connected to ground, to prevent the inner box from floating while open to users.
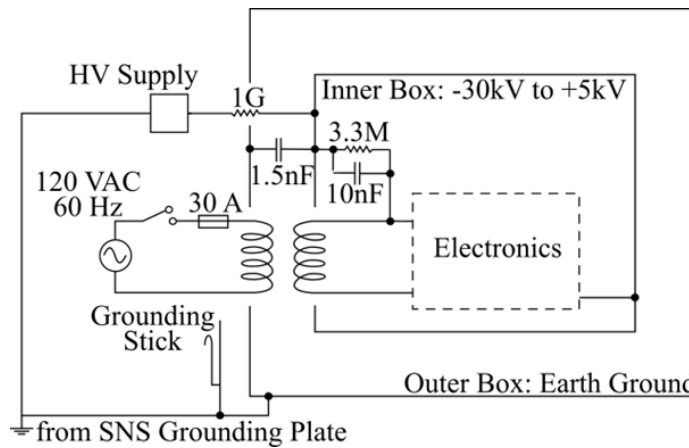
Figure 3.11: Electrical diagram for the high voltage assembly.

A second high voltage assembly, which functioned as a prototype for the Nab design, was constructed at LANL for the UCNB experiment. The inner box measured $35''$wide $\times$ $31''$deep $\times$ $20''$high and was surrounded by an air gap of 2.5" on all sides, except the bottom, where it was positioned on a piece of 1" polyethylene. Both the inner and outer box were constructed using 1" aluminum right angles, and were covered with 40% open, 14 gauge perforated aluminum mesh. The perforated mesh was chosen to allow for air flow to cool the electronics within the box.
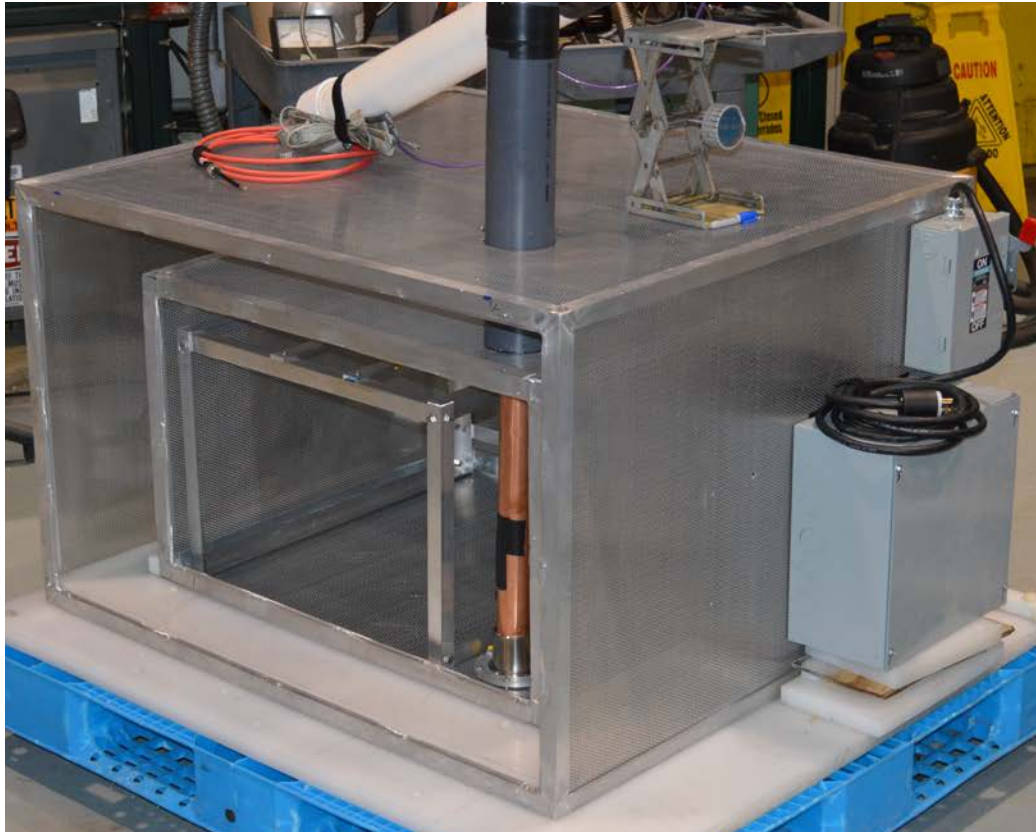
Figure 3.12: Prototype HV assembly at LANL.

A few design flaws were exposed with the LANL prototype. Most importantly, sharp edges on both the perforated aluminum and the corners of the angle resulted in a steady corona discharge. A file was used to round all exposed edges, which were then treated with corona dope to prevent further discharge, but the design was modified in for the Nab assemblies. Also, the boxes were too small for all of the electronics of the experiment. While all of the necessary equipment was installed for the run, connecting the cables and physically arranging the electronics proved difficult.

Based on the experience with the LANL prototype assembly, the Nab high voltage boxes were modified. The outer box retained the aluminum angle frame design used in the original assembly, but the inner box was framed with a rod-and-ball geometry. The edges of the box were hidden with 0.75" diameter aluminum rods which inserted into 1.125" balls at each of the corners. The rods were channeled to completely hide all sharp edges. Additionally, the air gap on all sides was extended to 4", and

the inner box was mounted on four polycarbonate standoffs. The standoffs were fabricated with a 0.5" ribbing along the sides to increase the minimal surface path for breakdown to 10", and each was fitted with four EZ-Lok press-inserts on both the top and bottom, to which the inner and outer boxes could be securely attached.

Before fabrication of the box began, the electric fields between the two plates were considered to assess the risk of arcing. The main concerns were the design of the insulating standoff and the abrupt transition at the corners of the inner box. Code was written using the finite difference method (FDM) in two dimensional cross-sections, with a uniform 800×800 grid with 0.01" spacing [35]. Dirichlet boundary conditions were implemented to fix the two boxes at the appropriate potential, and the open sides were assumed to follow a linear gradient according to an infinite parallel plate approximation. Figure 3.13 shows the results (electric potential gradient, equipotential lines, and electric field magnitude) for the box corners and insulating standoffs. The dielectric strength of air is typically reported as 3×10^6 V/m [36], or roughly 76 kV/inch, so the risk of breakdown was assumed to be negligible, as the maximum calculated field strength of 20 kV/inch.
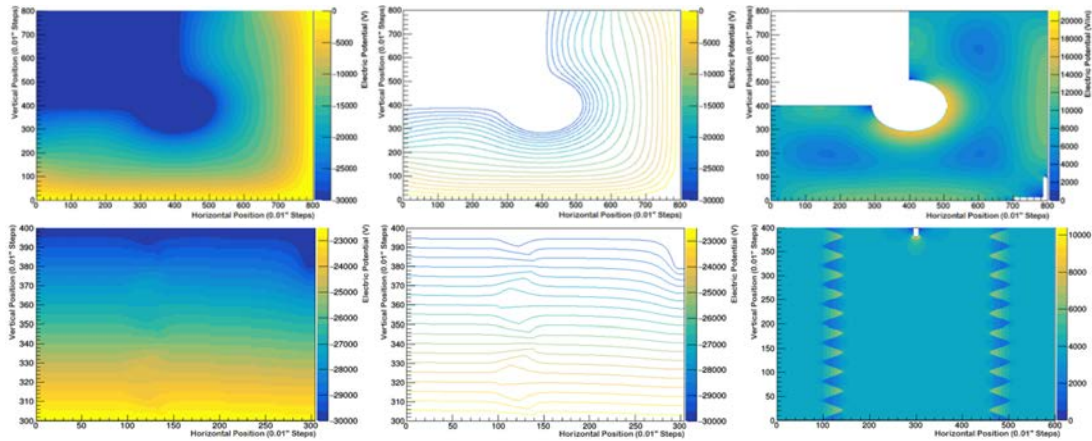
Figure 3.13: High voltage electrostatics calculations. The maximum calculated field of 20 kV/inch is well below the 76 kV/inch breakdown value from literature [36].

The size of the high voltage assembly was increased for the Nab experiment to accommodate the additional electronics necessary for the fully instrumented detector. The inner box was increased in size to 32"×31"×30", while the outer box's new dimensions were 42"×40"×40". Additionally, a 19" rack was integrated as part of the inner box design, so that electronics could use standard mounting hardware . Figure 3.14 shows the solid model for the Nab high voltage assemblies. Fabrication was a joint effort between the machine shops at the University of Kentucky and the University of Tennessee, and the high voltage assemblies are expected to be finished by mid 2019.
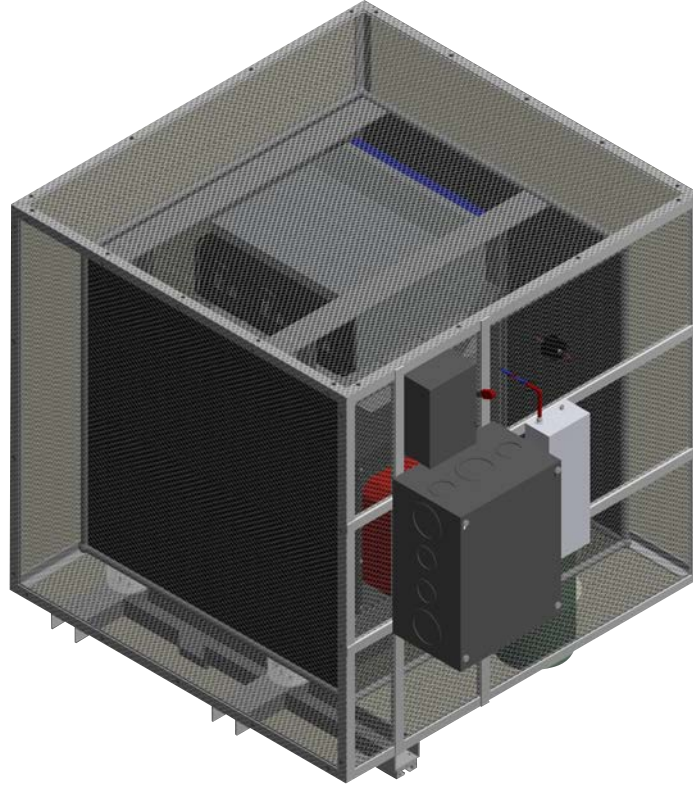
Figure 3.14: High voltage assembly for the Nab experiment. In this new design, the sharp edges of the perforated mesh are hidden in channels in the rod and ball geometry, and a 19" rack is added to the inner box for easy mounting of the electronics.

3.4 Signal Formation

Signal formation within the silicon detector depends on the type of radiation being deposited. Electrons, for example, primarily lose their energy by ionizing the detector atoms via Coulomb interactions with the orbital electrons and nuclei, although they also emit a small fraction of their energy as bremsstrahlung due to their deceleration through the detector. Because they are so light, electrons also change direction frequently as they traverse the detector, so the distribution of energy deposition is highly variable.

Protons, on the other hand, interact almost exclusively via Coulomb interactions; upon entry into the detector volume, the interact simultaneously with many electrons, creating compact clusters of charge carriers. Protons stop much quicker than electrons, and tend to maintain a uniform trajectory as they traverse the silicon.

As ionizing radiation transits the depleted region of a reverse-biased silicon detector, one electron-hole pair is created per 3.62-3.76 eV of energy deposited, depending on the temperature of the silicon [37]. Before they can recombine, as they would in an unbiased detector, the applied field sweeps the electrons and holes from the intrinsic region to the conductive contacts on either side of the detector. Along the way, the charge carriers experience thermal diffusion and Coulombic interactions with other

free charges, which both contribute to lateral migration in addition to the motion along the drift field. These secondary effects, however, are much smaller in magnitude than the motion induced by the applied voltage, so they contribute much less to the overall pulse shape.

The current output from the detector begins as soon as the charge carriers begin their migration, and continues until all motion has stopped. The induced current is described by the Shockley-Ramo theorem:

$$i = q\vec{v} \cdot \vec{E_0}. \tag{3.2}$$

In the above equation, q is the charge of the carrier, $\vec{v}$ is the drift velocity of the charge, and $\vec{E_0}$ is the weighting field. Unlike the true electric field, which solves the Poisson equation for the physical detector geometry, charges, and boundary conditions, the weighting field is the solution to the Laplace equation where all physical charges within the detector volume are ignored and artificial boundary conditions are applied. Specifically, all electrodes are fixed at ground with the exception of the desired readout pixel, which is set to unity.

The carrier velocity is a function of temperature, drift field, and silicon crystal orientation. As seen in Figure 3.15, the electron mobility is greater than the hole mobility for a given temperature and drift potential. As such, the incident particles for all experiments enter through the junction window to minimize the distance traveled by holes, and thus the total charge collection time.

(a) Electrons.

(b) Holes.
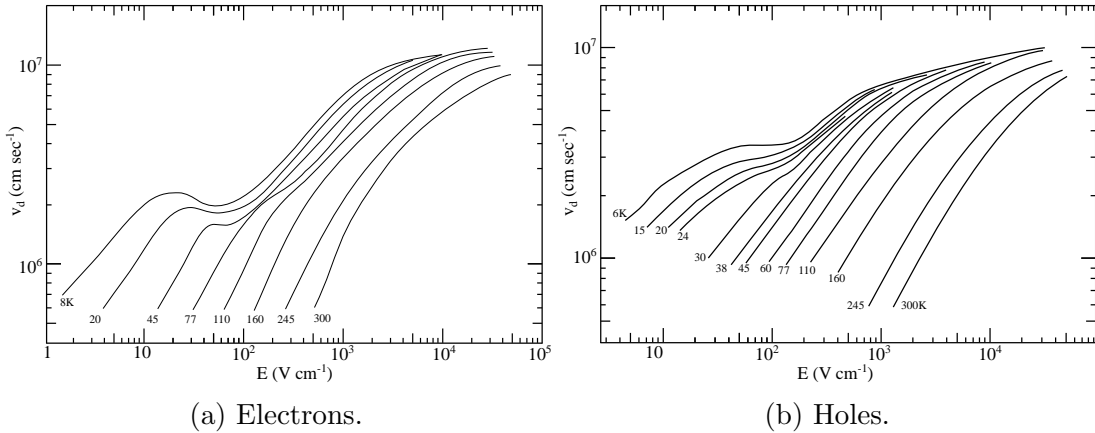
Figure 3.15: Charge carrier drift velocities as a function of temperature and drift field for silicon oriented in the $\langle 100 \rangle$ plane. The figure is a reproduction of Figures 2 and 5 from [38].

In the Nab experiment, the target silicon temperature is 110 K, and the applied drift field will likely be around 2000 V/cm. Based on these values, the expected charge carrier mobilities should be roughly 5.7×10^6 cm/s and 4.1×10^6 cm/s for the electrons and holes, respectively. With these numbers, then, full charge collection should take less than 30 ns.

3.5 Monte Carlo Signal Generation

The simulation efforts developed in Penelope and Geant4 for UCNB, ^{45}Ca, and Nab provide an abundance of information about systematic effects involving particle transport through the spectrometers and charge deposition within the volume of the silicon, but they fail to provide insight into the charge collection and signal formation processes, which are just as important for understanding the spectrometer response. To better understand these processes, a signal generation code was written to track charge carrier migration within the silicon and calculate the induced current for the purpose of investigating the impact of signal formation.

The signal generation code utilizes four main components: calculating the weighting and drift fields according to the detector geometry, simulating the electron path through the silicon, determining the charge carrier migration and resultant induced current, and using the current pulse to generate the observed waveform. The weighting and drift fields are determined separately using numerical methods on a sparse grid, and interpolation is used within the simulation to produce finer results. The electron paths are also determined ahead of the signal generation program, and are used to establish the charge carrier production within the silicon.

The weighting potential is calculated using FDM, as in Section 3.3. For these calculations, however, the algorithm extends to three dimensions, as the components of the weighting field, rather than the magnitude, are required to calculate the induced current. A uniform grid with spacing 15 μm is used to balance computation time and resolution in the output field. The grid is calculated approximately 2.5 mm beyond the pixel boundaries in both x and y so that the fringe fields are determined correctly. Figure 3.16 shows the results of the numerical calculations, with a two dimensional cross-section of the weighting potential and the potential as a function of z for fixed x and y.

At the center of the pixel, the weighting potential is fairly consistent with the potential function for a parallel plate capacitor. At the edges of the pixel, however, fringe fields begin to influence the electric potential profile. Induced current formation near the pixel boundaries is more strongly dependent on carrier motion nearest the collection plate, so the largest contribution to the current pulse is due to the electron motion within the silicon. Outside the pixel, an induced signal still appears due to stray capacitance; however, because the net change in electric potential transited by the charge carriers is zero, the signal is bipolar with an integral area of zero. This leads to differential crosstalk signals in the detector, which if used correctly, can provide sub-pixel spatial resolution for incident particles.

While the true drift field depends on the bias potential and the dopant concentration distributions, the drift field for these calculations was chosen to be a uniform field between the two sides of the silicon. The motivation for this decision was to focus on understanding the effects of charge migration on the observed pulse shape at the simplest level; fortunately, this approximation was appropriate for a few reasons: First, the distortions due to impurity presence or dopant gradients should be negligible as the silicon used in the detectors is ultra-high purity (Sec. 3.1). Second, the shallowness of the dopant implants on both the junction and ohmic sides of the detector

(a) 2D cross-section

(b) $\varphi_0(z)$ for fixed x and y
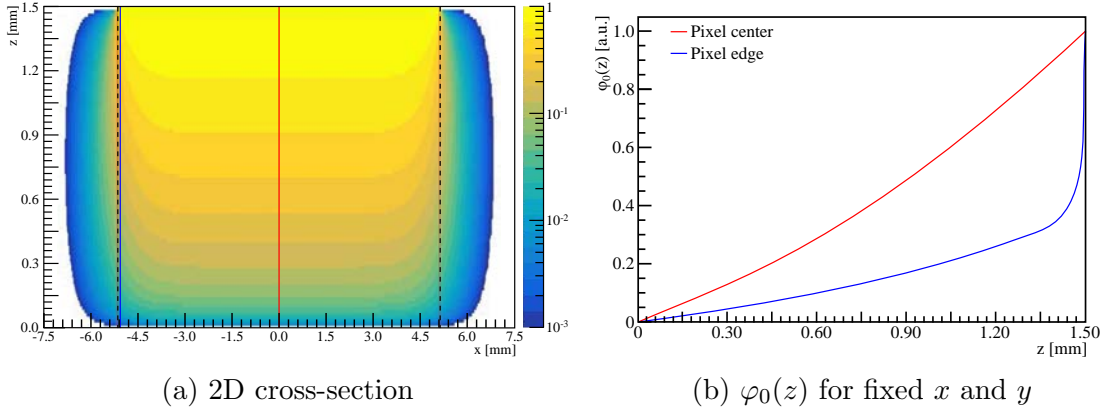
Figure 3.16: The numerical solution for the hexagonal pixel weighting potential, as determined by FDM. At the center of the pixel (red line), the potential is nearly linear in z, which is consistent with the parallel plate approximation. Near the pixel boundaries, however, the fringe fields leak into adjacent pixels, so the weighting potential favors carrier motion in close proximity to the collection plate for current induction.

resulted in a drift field that was completely dominated by the intrinsic region. Last, the detector was biased so far beyond depletion that the expected ratio of the slope of the electrostatic field to the constant component is very small. Thus, the field was approximated by a uniform field of magnitude 2 kV/cm for these calculations.

For rapid development of the simulation code, version 2.4.8.1 of the CASINO (monte CArlo SImulation of electroN trajectory in sOlids) software was used to model the electron trajectories within the silicon [39]. Originally written for scanning electron microscopy, CASINO provides a detailed profile of the electron's energy deposition as it traverses the silicon. These profiles are used to establish the starting position and carrier densities for the signal formation script. Figure 3.17 shows a sample simulation of the electron trajectories in a silicon target.

Running a simulation minimally requires definition of a target specimen and configuration of the electron microscope. The specimen consists of rectangular slabs of material, which can be defined according to pre-defined values or user input. The silicon detectors are modeled as 100 nm of silicon as the dead-layer on top of 1.4999 mm of silicon for the active region. The electron microscope is also user configurable with regards to energy, beam width, and incident angle. For these studies, normally incident, monoenergetic beams of radius 1 nm are used.

The output of the CASINO simulation is exported to ASCII data files, where each trajectory includes a header followed by an array of the data values until the electron tracking ends, as shown in Figure 3.18. The relevant data garnered from these files is limited to the position coordinates and energy of the electron throughout its path. For each step, the number of electron-hole pairs is determined by considering the energy loss between the successive steps divided by the mean energy to create the pair. The time step between each deposition is determined from the average energy of the electron between the two steps and the distance traveled.

33
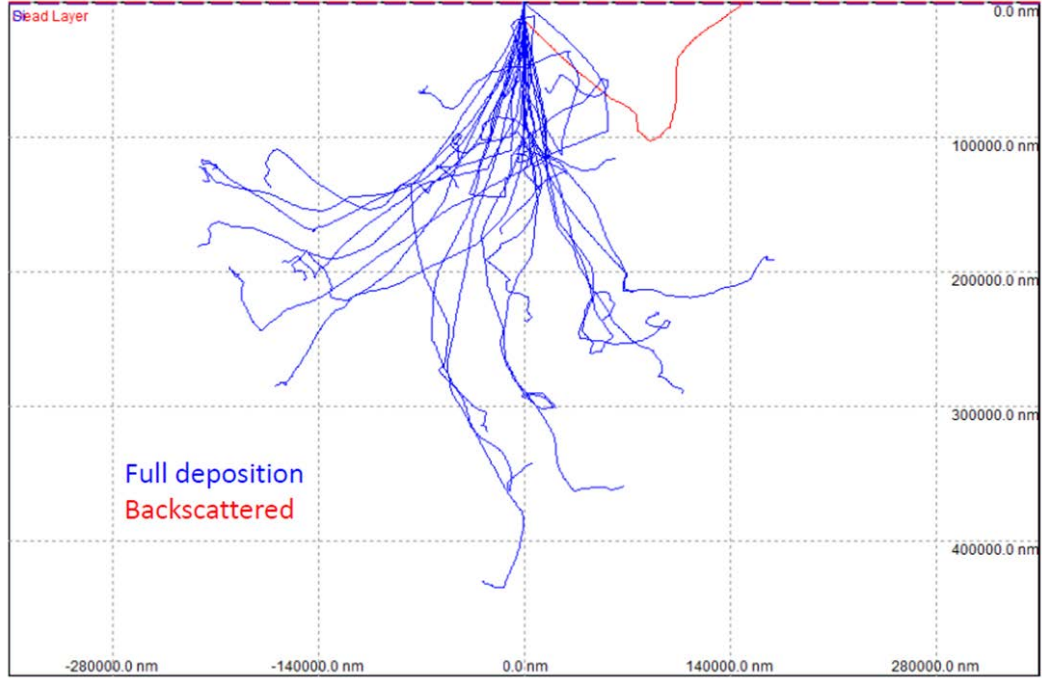
Figure 3.17: Electron trajectories through the silicon as calculated by the CASINO program. The blue paths are fully deposited, while the red paths scatter out of the detector.

```
-------------------------------------------
     Trajectory   0
BackScattered   Transmitted   Detected   Displayed   Nb Collisions   Max Depth   Mean Free Path   Mean dE/dS    Mean Phi    Mean Theta   Mean Position
no              no            no         yes         1779            62494.3     0                -0.00321689   3.14445     0.0626815    (22546.4, -17118.8, 46175.2)
X               Y             Z          Rx          Ry              Rz                  Energy       Region
0.381396        1.76402       0          0           0               1                   100          Dead
0.381396        1.76402       1.9821     0           -0              1                   99.9873      Dead
0.381396        1.76402       100        0           -0              1                   99.9251      Active
0.381396        1.76402       340.947    0           -0              1                   99.7445      Active
0.381396        1.76402       421.987    -0          -0              1                   99.6836      Active
0.381396        1.76402       490.664    0           -0              1                   99.6321      Active
0.381396        1.76402       539.428    -0          -0              1                   99.5955      Active
    :               :             :          :           :               :                   :            :
-39097.7        -2950.18      48362.2    0.225677    -0.447137       -0.948319   0.122035     Active
-39097.7        -2950.2       48362.2    0.223078    -0.770412       -0.677846   0.120846     Active
-39097.7        -2950.28      48362.1    0.483394    0.671796        0.614769    0.116481     Active
-39097.5        -2950.06      48362.3    0.46866     0.811665        0.811665    0.102557     Active
-39097.3        -2949.88      48362.7    -0.499092   0.578912        0.552798    0.0835192    Active
-39097.7        -2949.44      48363.1    -0.499092   0.578912        0.552798    0.05         Active
-------------------------------------------       -------------------------------------------
     Trajectory   1
BackScattered   Transmitted   Detected   Displayed   Nb Collisions   Max Depth   Mean Free Path   Mean dE/dS    Mean Phi    Mean Theta   Mean Position
no              no            no         yes         1884            28450.4     0                -0.0028651    3.11798     0.058557     (1665.08, 23436.8, 14946.4)
X               Y             Z          Rx          Ry              Rz                  Energy       Region
6.72026         -1.15695      0          0           0               1                   100          Dead
6.72026         -1.15695      100        0           0               1                   99.9251      Active
    :               :             :          :           :               :                   :            :
```

Figure 3.18: Example output data from a CASINO simulation.

The simulated charge collection was performed in 10 ps steps and followed directly from the generation of each electron-hole pair cluster. For each cluster, the position at each time step is determined via numerical integration. The weighting field at that position is then determined via cubic spline interpolation within the numerical grid, and the induced current for that time step is calculated according to Equation 3.2. An array indexed by time steps is used to track the cumulative induced current for all charge carrier clusters associated with that event, as shown in Figure 3.21(a).

The last step to determine the overall signal shape is to model the response of the amplification electronics to the input current pulse. The most robust method for

achieving this would entail explicit measurement of the system response function and convolving with the output from the charge collection simulation. However, this was difficult, as the equipment necessary to make such a measurement was unavailable.

The next best alternative was to use the Simulation Program with Integrated Circuit Emphasis, or SPICE, to determine the expected electronics response [40]. The circuit shown in Figure 3.6 was modeled in Analog Device's LTspice environment [41], where the detector was modeled as a series resistor and capacitor pair with a supplied voltage pulse. Figure 3.19 shows the model built to simulate the electronic response and the resultant signal shape.
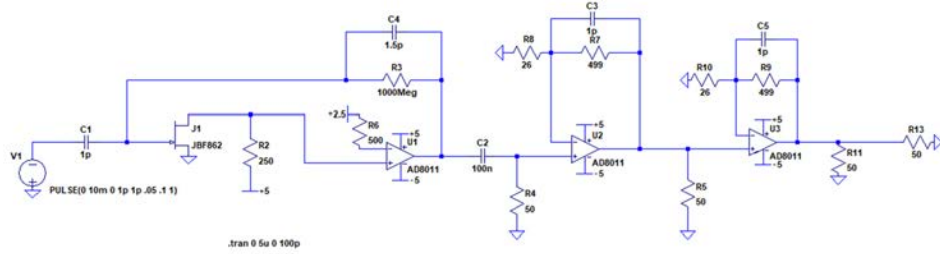
(a) LTspice circuit model
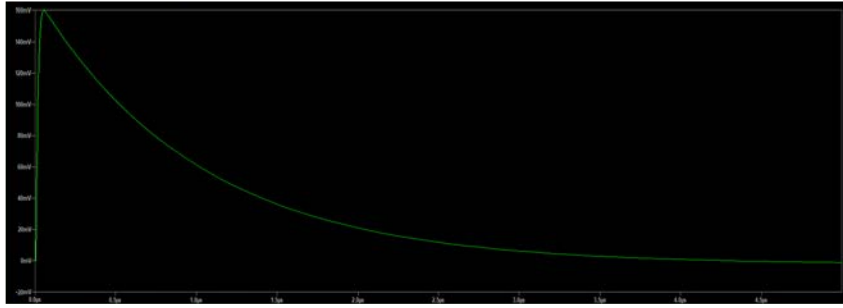
(b) Simulated signal shape

Figure 3.19: Using Analog Device's LTspice, the electronics were modeled to predict the output signal shape. The detector is modeled as a series resistor and capacitor combination with a an input voltage pulse.

To improve the speed of the code, a series of high- and low-pass filters implemented recursively is used to model the effect of the electronic shaping. In doing so, the filters can be integrated directly into the simulation code without additional libraries. The recursion relations for first-order high-pass and low-pass filters are

$$f_{\text{LP}}[t] = \alpha_{\text{LP}} \left(x[t] - f_{\text{LP}}[t-1] \right) + f_{\text{LP}}[t-1] \tag{3.3}$$

and

$$f_{\text{HP}}[t] = \alpha_{\text{HP}} \left(f_{\text{HP}}[t-1] + x[t] - x[t-1] \right), \tag{3.4}$$

where $\alpha_{\text{LP}} = 1/(\tau + 1)$ and $\alpha_{\text{HP}} = \tau/(\tau + 1)$. To model the response of the CR-$(\text{RC})^2$ shaping filter, the integrated charge pulse is sequentially processed by a single

35

high-pass stage with a 5 μs time constant, followed by two low-pass stages with time constants of 7 ns. Figure 3.20 shows a comparison between this light-weight algorithm and the more realistic LTspice output. The difference between the two simulations is within 5% of the full scale amplitude, confirming the validity of this lightweight solution. The amplitude of the resultant pulse is scaled according to empirical data to reflect the gain in the electronics.
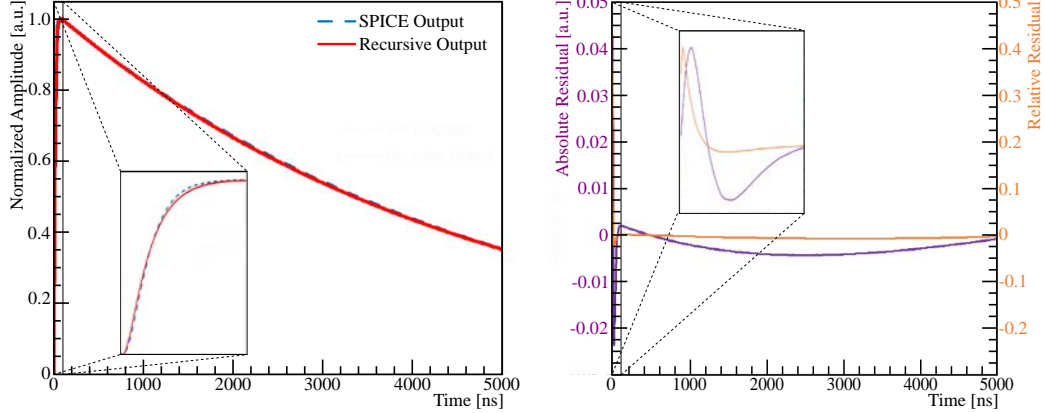
Figure 3.20: LTspice simulated output versus the recursive CR-$(RC)^2$ signal response (left), and the residual difference between the two (right). Since the difference between the two is within 5% of the full scale amplitude, the faster recursive approximation is utilized for the waveform generation.

Figure 3.21 shows the successive stages of waveform simulation. After the final pulse shape is produced, random noise is superimposed according to a measured power spectrum from the silicon detectors, such as in Figure 3.2(b). To produce arbitrary noise, the phase for each frequency bin is randomly assigned before using the inverse Fourier transform to convert to the time domain. Figure 3.22 shows a comparison between a simulated and a measured signal, where a slight time offset is inserted between the two pulses for clarity. Although the synthetic noise differs from the measured features, as modifications to the electronics changed the spectrum, the rising and falling edges of the pulse agree nicely between the two waveforms.

10000 electrons were simulated at 17 and 300 keV in order to motivate further discussion in the signal formation process and how it pertained to Nab. Much like protons, the 17 keV electrons deposited all of their energy in a very small portion of the silicon, and thus produced a consistent pulse shape. The 300 keV electrons, on the other hand, exhibited a wide array of paths and penetration depths, which resulted in a collection of pulse shapes. Figure 3.23 shows the distribution of induced current pulses for the two electron energies.

Using the simulated current pulses, the full waveform shape was produced according to the procedure described above. The output of these simulations is shown in Figure 3.24. The noise was omitted for inspection of the resultant pulse shape, but a full-blown analysis will require a simulation with all pertinent features.

While these simulations approximated several important features of the underlying physics, they highlighted a significant issue for the Nab experiment. With the

(a)

(b)

(c)
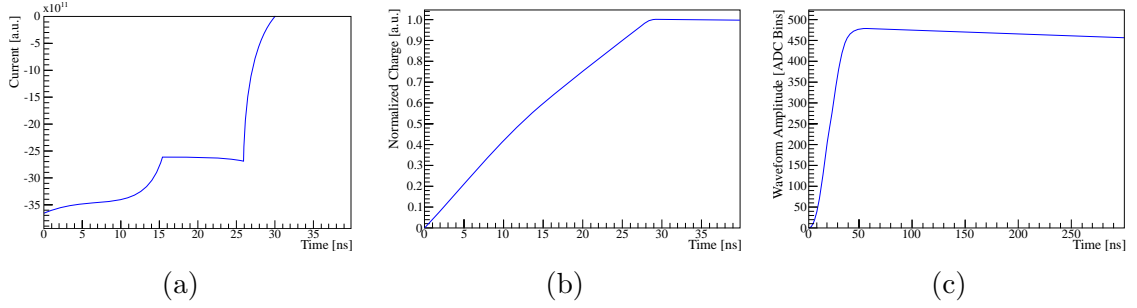
Figure 3.21: To simulate waveform generation, (a) the Shockley-Ramo theorem is applied to determine the induced current pulse, (b) which is then integrated and normalized and (c) run through a CR-$(RC)^2$ shaping filter to produce the resultant pulse shape.
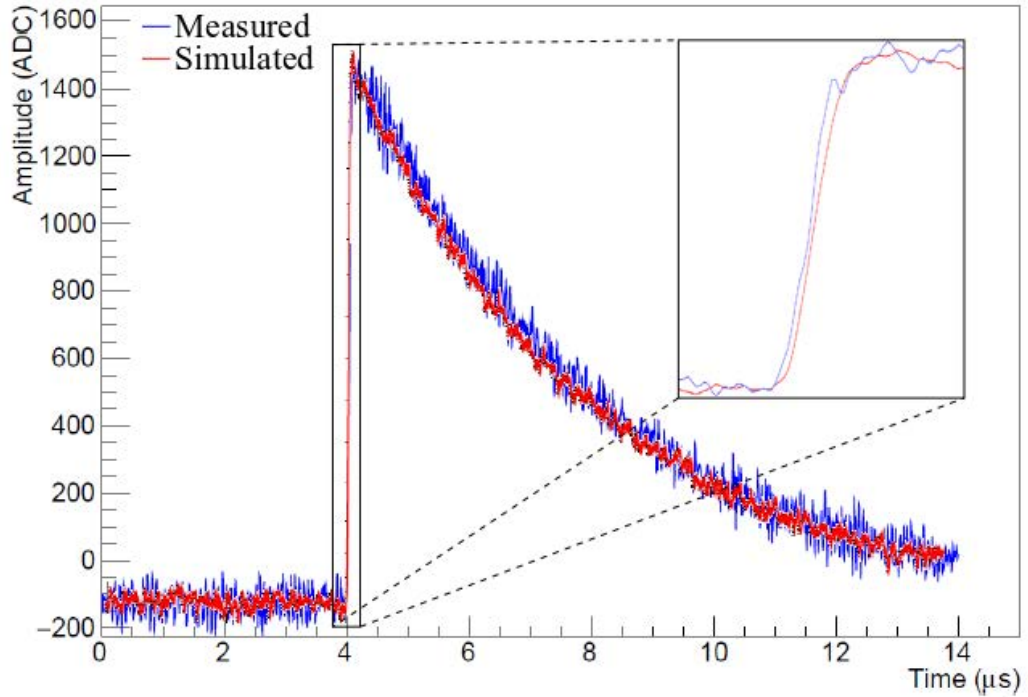
Figure 3.22: Measured and simulated signals for the silicon detector. Modifications to the electronics produce a different noise spectrum between simulation and measurement, but the calculated rising and falling edges of the pulse agree well.

requirement of a 300 ps systematic bias in particle arrival time, signal formation as a function of energy and location need to be understood better. Inspection of the rising edges of the waveforms for the two different energies, as shown in Figures 3.25 and 3.26, illuminates that the charge collection process could easily produce an uncertainty on order of 1 ns if not taken into account.

37

(a) 17 keV current.
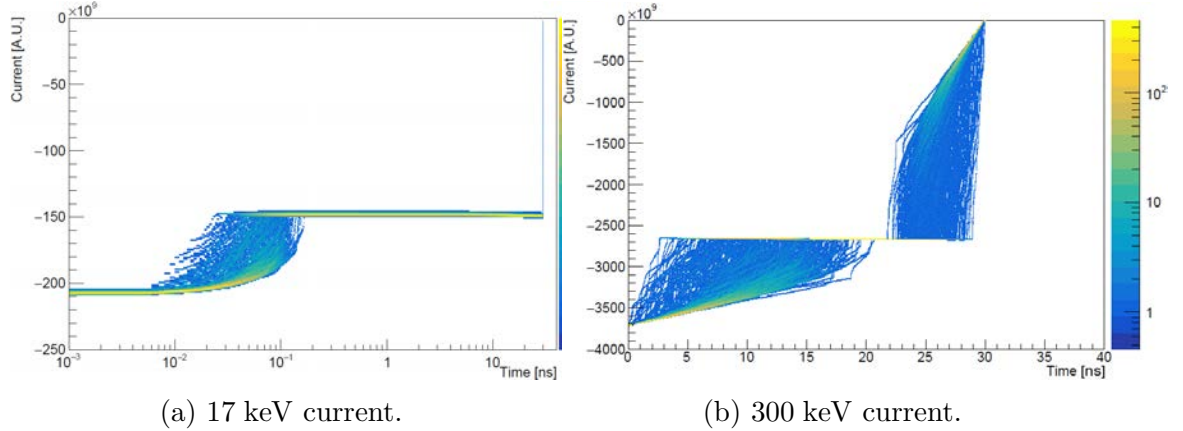
(b) 300 keV current.

Figure 3.23: Induced current on ohmic readout at 17 and 300 keV. The low-energy signal time axis is shown on a logarithmic scale to emphasize the very minor variation in current pulse shape, while the higher-energy signals vary substantially.

(a) 17 keV waveform.
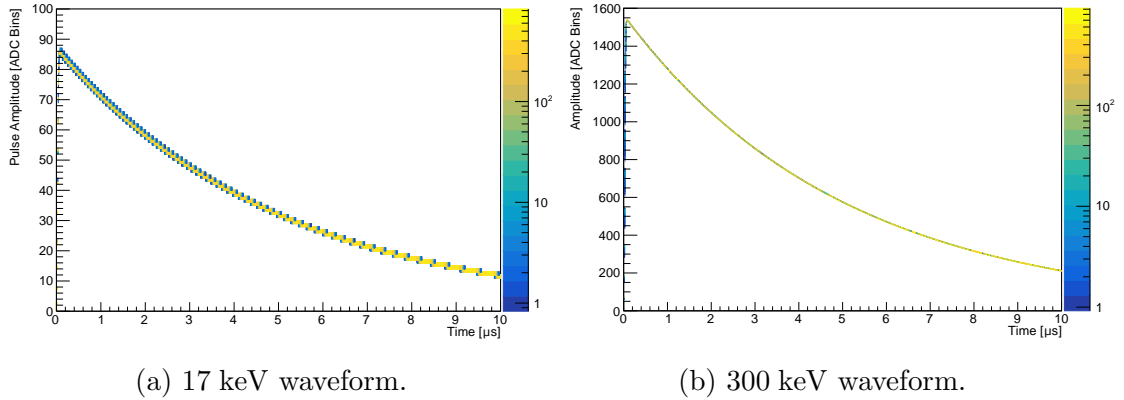
(b) 300 keV waveform.

Figure 3.24: Simulated waveforms for 17 and 300 keV electrons. The output pulse shape is coerced to integer bins and 4 ns ticks to reflect the digitized records in the actual experiment.

3.6 Lightweight, Linear Convolutions for Signal Approximation

Simulations such as those discussed in Section 3.5 produced high fidelity data for systematic studies, but significant resources and time were required to generate appreciable statistics. In order to efficiently conduct exploratory data analysis, such as for filter design, dead-time characterization, or pile-up identification logic purposes, synthetic signals were needed on a much quicker time scale. Previously used methods to approximate the shaped detector response included a difference of exponentials [42], piecewise Gaussian functions [43], and a Gaussian-convolved exponential pulse [44].

While these methods modeled the resultant pulse-shape nicely, they did not allow for an easy study of specific features such as the charge collection time. The rise time from more commonly used methods was designed to model the total system rise time, which includes both the charge collection and electronic shaping times. Be-

(a) 17 keV rising edge.

(b) 300 keV rising edge.
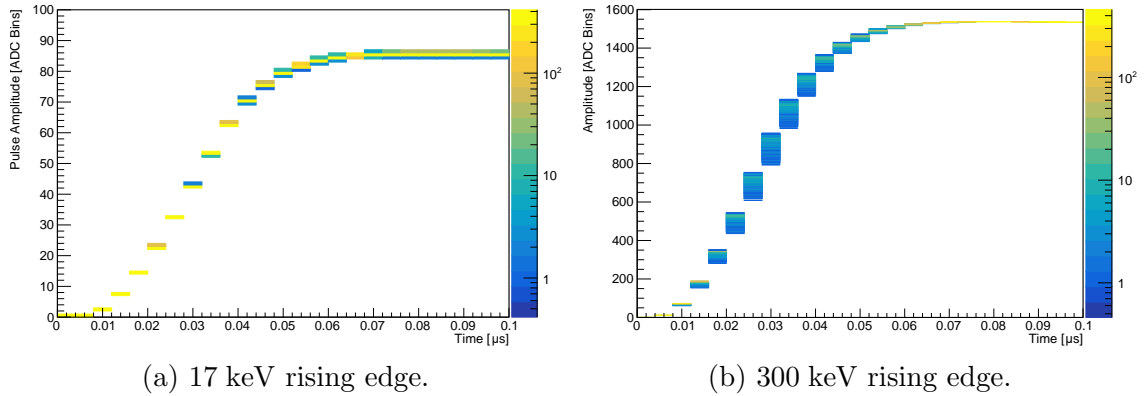
Figure 3.25: Inspection of rising edges. The consistency of the 17 keV current pulses results in a homogeneous pulse shape, while the variability in induced current for the 300 keV electrons is apparent in the overall pulse shape
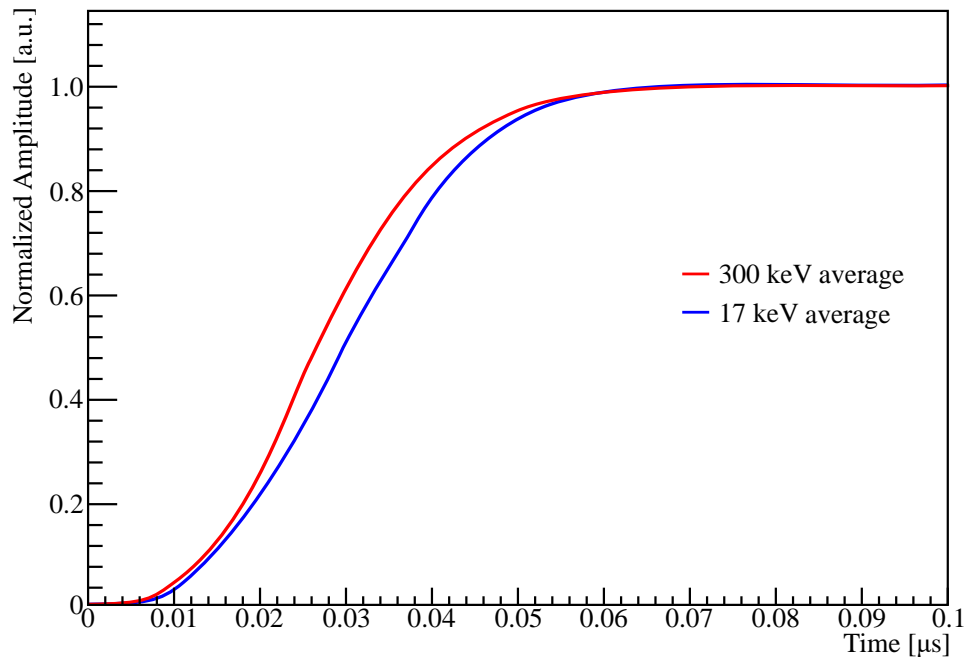
Figure 3.26: Average rising edges for 17 and 300 keV pulses at the center of the detector. Due to the energy deposition profile of the incident electrons and the process of current induction, the two shapes are distinct enough to merit consideration to avoid a systematic bias in timing.

cause of the intertwined nature present in these pulse models, accurately reflecting the contributions of just the charge collection dependency was difficult. Furthermore, modeling effects such as pileup required calculating each pulse separately and adding the output from each pulse. For high-fidelity pileup analysis, a large number of waveforms with a variety of pulse spacing and rise times must be calculated, and ideally, this should be done with a randomized stream of these events. The aforementioned methods, while useful for fixed length traces, proved cumbersome for these types of

simulations.

A simple work-around was established by using a series of linear convolutions to emulate the detector current pulse, followed by the integration and linear convolutions used in the Monte Carlo applications of the previous section. The pulse amplitudes could be randomly generated as a series of steps, which, when convolved with current pulse and shaping response functions, would produce a reasonable approximation for real detector signals. Furthermore, in the event of successive pulses, they would be properly superimposed with no additional cost for simulation. Figure 3.27 shows the progression from step function to shaped waveform.

(a) Step function

(b) Convolved current pulse

(c) Integrated charge

(d) Shaped waveform
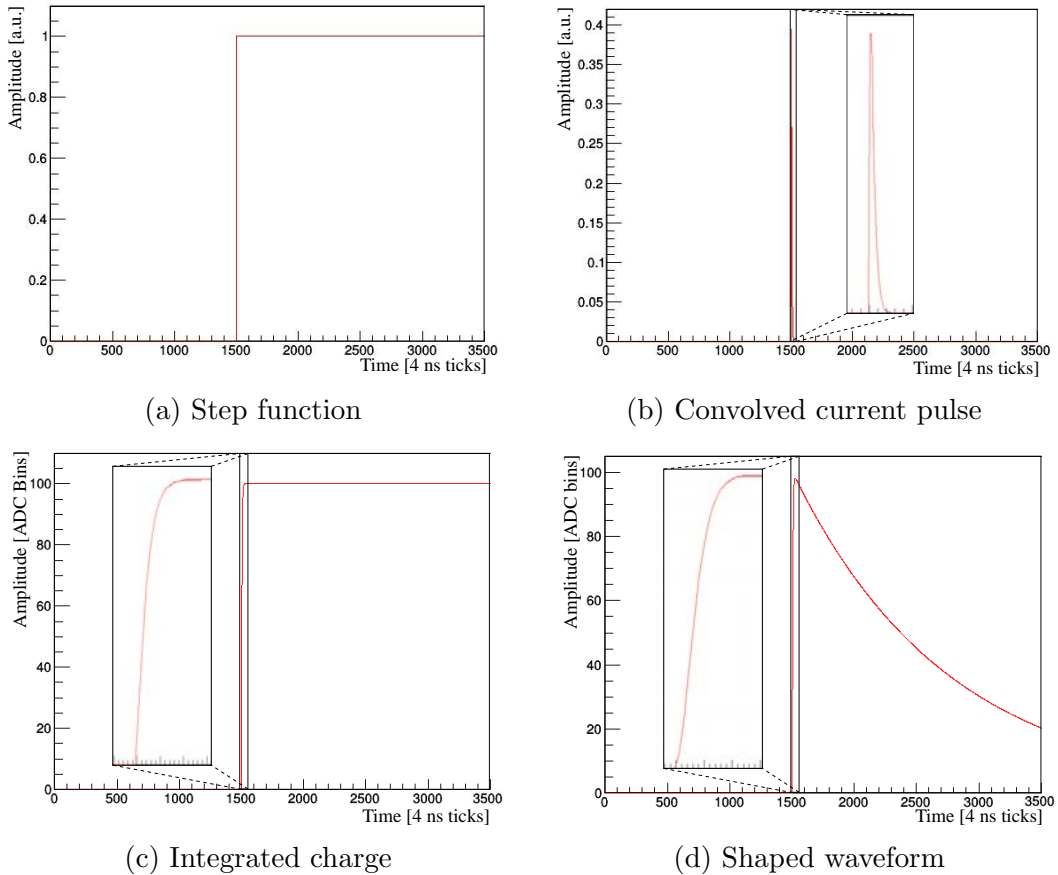
Figure 3.27: Synthetic pulse generation using linear convolutions was comprised of four steps: (a) the initial step generated at the start of each waveform, (b) convolution with a semi-Gaussian to mimic the current pulse shape, (c) integration to determine the total current and scaling to ADC bins, and (d) fast filter convolution to produce the shaped waveform.

These synthetic waveforms, while not useful for predicting the detector physics, were helpful for waveform analysis and digital filter design studies. The full Monte Carlo simulation of waveforms required significantly more time for generation, so any tests that required large numbers of waveforms or a modification of the run parameters would be hindered. However, these synthetic waveforms mimicked real pulses

well enough to test code for high statistics pile-up analysis, multi-pixel energy reconstruction, and trigger efficiency studies. Furthermore, the efficiency of the algorithm allows it to be implemented in FPGA, as discussed in greater detail in Section 5.6.

Chapter 4: Timing Resolution Studies

In order to meet the target precision of $\delta a/a \leq 1 \times 10^{-3}$ in the Nab experiment, a 300 ps limit on the systematic bias to particle arrival time must be maintained. The average reconstructed time of arrival of the electron and proton must be within 300 ps of the true arrival time. The challenge in this task lies in the fact that the rising edge of detector signals is affected not only by the detector effects described in section 3.5, but also the particle momentum and type. Low energy electrons provide a fairly homogeneous charge deposition profile, with the bulk of the induced current resulting from the motion of the negative charge carriers in the silicon. As the incident electron energy increases, however, the deposition profile depth and lateral distribution increase, resulting in a greater contribution to the induced current from hole motion.

In this chapter, a method for detector model validation is demonstrated using tagged electron sources. Temperature, dopant concentration distributions, depletion gradients, bias voltage, and particle deposition profiles all heavily influence the time of arrival T_0 inferred from rising edge.

4.1   Using Tagged Electrons to Measure Transit Time

In order to benchmark the detector response model, tagged electrons were used to assess energy and pixel dependence in the rising edge shapes due to charge collection effects. Comparison of the variation in effects as a function of electron energy, location within a single pixel, and location across the detector can be used to ascertain the effective charge carrier mobilities. The tag for the electron is a coincident particle measured in a nearby detector with higher precision timing capabilities. In our case, the tagged electrons were produced by electron capture sources, with the simultaneous photon detected in a fast scintillator.

Isotopes that decay via electron capture start with an initial capture of an electron, typically from the K or L shells, followed by a cascade of nuclear relaxations until they reach their ground state configuration. Each nuclear de-excitation can occur via the emission of a gamma or a conversion electron, which carries energy equal to that of the gamma minus the electron's binding energy. Each transition level has varying probabilities for conversion gamma or electron emission; additionally, the relative probabilities for the shell from which the electron is ejected varies by line.

In conjunction with the conversion lines for the decay, each vacancy produced from the capture or emission of an electron is capable of producing additional particles. Typically, a characteristic X-ray is produced as an outer shell electron relaxes to fill the vacancy. When the initial vacancy is within the K shell, the electron that relaxes usually originates in the L shell, resulting in the Kα X-ray lines, or the M shell, which produces the Kβ lines. Similarly, a hole that is originally within the L shell is filled by an electron from the M shell or beyond, emitting a much lower energy L X-ray.

Vacancies beyond the L shell do not produce X rays of sufficient energy for detection in our apparatus.

Another possibility is Auger production, where the relaxing electron transfers its energy to a secondary emission electron. When the initial vacancy is within the K shell, the hole is usually filled by a relaxing L shell electron, which can either eject another L shell electron (resulting in two additional vacancies) or an M or higher shell electron (only one additional detectable vacancy); the prior Auger process is referred to as a KLL Auger (K shell vacancy filled by L shell electron that results in emission of an Auger from the L shell), while the latter process is addressed as a KLX Auger. Figure 4.1 summarizes the four mechanisms for radiation production in an electron capture source.
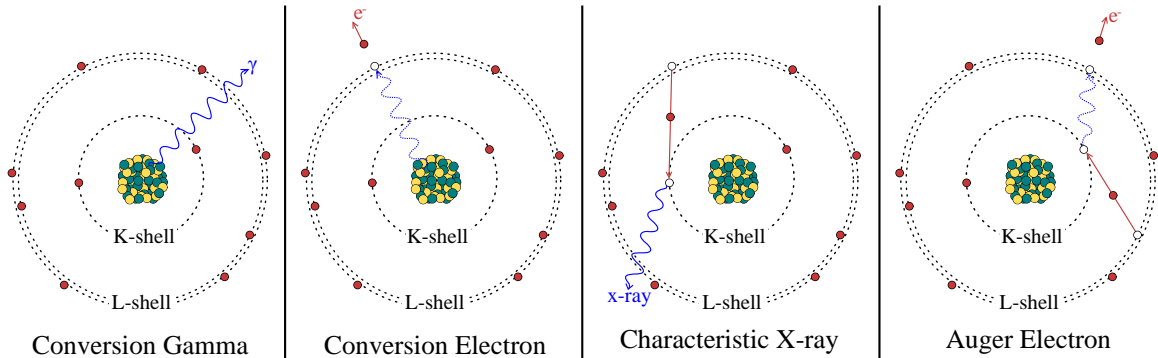
Figure 4.1: The four main mechanisms of radiation production in an electron capture source. After the initial electron capture, the nucleus is in an excited state. As it de-excites to its ground state, it can emit a photon or electron corresponding to the energy level differences. Meanwhile, inner shell vacancies are filled by outer shell electrons, producing characteristic x-rays and augers. Depending on the complexity of the decay chain, various combinations of electrons and photons can be emitted for each decay event.

A single decay event is capable of producing multiple decay products in coincidence, which can be used for precision timing measurements. The lifetimes for each of the nuclear transition states are well-documented, while the filling of vacancies is nearly instantaneous. Thus, the relative timing distribution for electrons and photons measured in coincidence can be predicted with better precision than the required 300 ps systematic for Nab.

4.2 Electron Capture Source Event Generators

A Monte Carlo generator for several commonly used electron capture sources has been written [45]. Each decay is simulated as a cascade, where the probabilities for the processes described in the previous section are produced according to empirical and tabulated data from nuclear databases [46]. The two sources initially considered are ^{139}Ce and ^{133}Ba. Their decay schemes are shown in Figure 4.2, which highlights the initial transition state probabilities after each electron capture, the available con-

version lines for the cascade to ground, and the average lifetimes of each state. To minimize auto-correlation within the simulation, a pseudo-random number generator was implemented using the `mt19937_64` engine, based on the 64-bit Mersenne twister, from the C++11 random library [47].

(a) ^{139}Ce   (b) ^{133}Ba

Figure 4.2: The decay schemes for ^{139}Ce and ^{133}Ba. Each isotope decays via electron capture, followed by a cascade of conversion gammas and electrons, as well as x-rays and auger electrons. Given the well-understood lifetimes of each state, these sources can be used in tagged electron timing studies.

At the start of the program, the number of loop iterations is determined according to the desired simulation rate, the number of events to be simulated, and the time step for each iteration, which is 100 ps by default. In each iteration, a random number is thrown and compared to each decay probability to determine whether to generate that event. Each decay process follows a strict pattern: determine the initial capture, fill the vacancy, and continue the process of relaxation to the ground state. If a vacancy is in the M-shell or higher, no further emission is recorded, as the energy of these decay products would be too small to observe.

For each emission, three uniform random numbers between 0 and 1 are thrown. The first determines the specific emission line, while the next two determine the direction of the outgoing particle. To populate 4π space evenly, random θ and ϕ are selected according to

$$\theta = \mathrm{acos}\left(1 - 2u\right), \quad \phi = 2\pi u, \tag{4.1}$$

where u represents each random number. For conversion lines, an additional random number is generated to determine the time of the nuclear relaxation. Once again, a uniform random number between 0 and 1 is used to generate numbers according to the Poisson distribution,

$$t_{emission} = t_0 - \tau \ln\left(1 - u\right). \tag{4.2}$$

In the above expression, τ represents the nuclear transition state average lifetime.

10^6 decay events each were simulated for ^{139}Ce and ^{133}Ba sources using the generators. Each output file was analyzed to determine the ideal timing difference distribution between the emission of an electron and any coincident photons. Due to

44

the possibility of multiple emissions for a single decay, the script considered 5 (13) events before or after each electron for the ^{139}Ce (^{133}Ba) data. The coincident energy spectra were also calculated, but without a simulation of detector response, the plots were not particularly illuminating. Figures 4.3 and 4.12 show the simulated timing difference distributions for the prominent electron lines of the two isotopes. Additionally, Tables 4.1 and 4.2 list the mean and standard deviation for each of the distributions, as these properties validate the timing algorithm used to process the silicon waveforms.

Figure 4.3: Ideal $\Delta T_{\gamma-e}$ distributions for the prominent electron lines from ^{139}Ce decay. Due to detector resolution effects, electron lines within a few keV are grouped together as a single line and L X-rays are omitted.

Table 4.1: ^{139}Ce $\Delta T_{\gamma-e}$ distribution characteristics

E_e [keV]	Mean $\Delta T_{\gamma-e}$ [ns]	Std. Dev. [ns]
27	0.822 ± 0.006	1.719
127	-0.615 ± 0.002	1.217
160	-1.269 ± 0.009	1.452
All	-0.332 ± 0.002	1.502

Each of the two sources has its advantages and disadvantages. ^{139}Ce, for example, offers a much cleaner spectrum for analysis, but the higher energy electrons are all in such close proximity that there are really only two unique data points to consider for
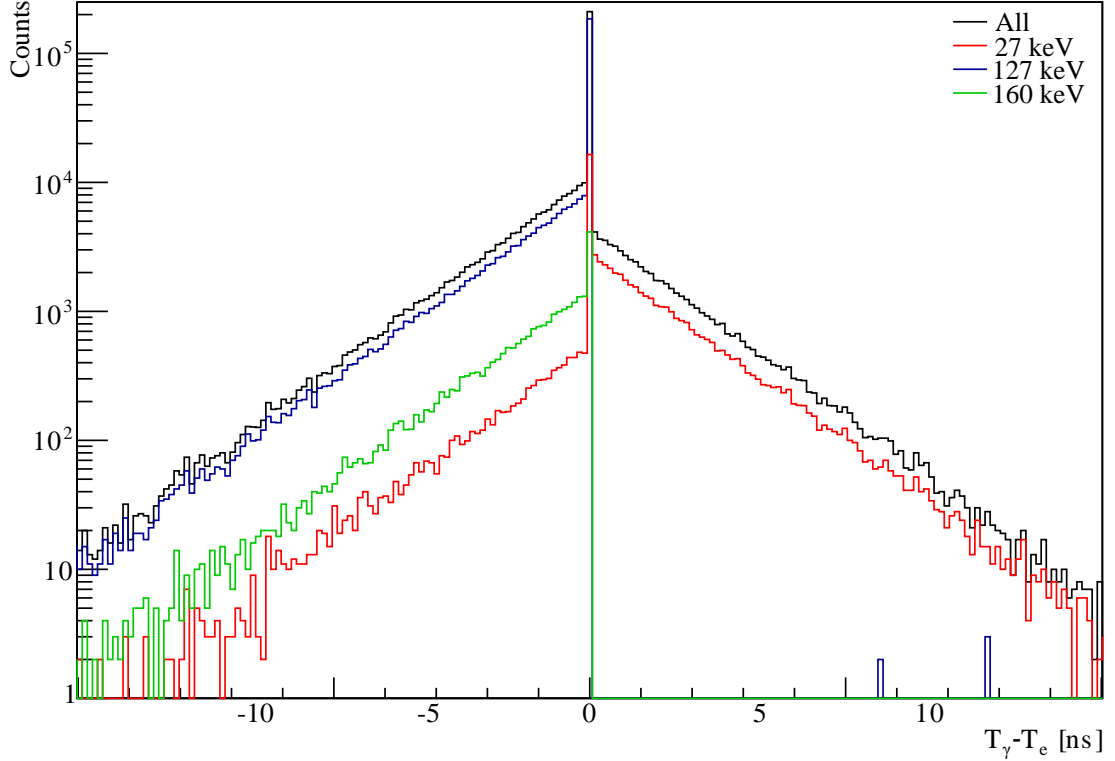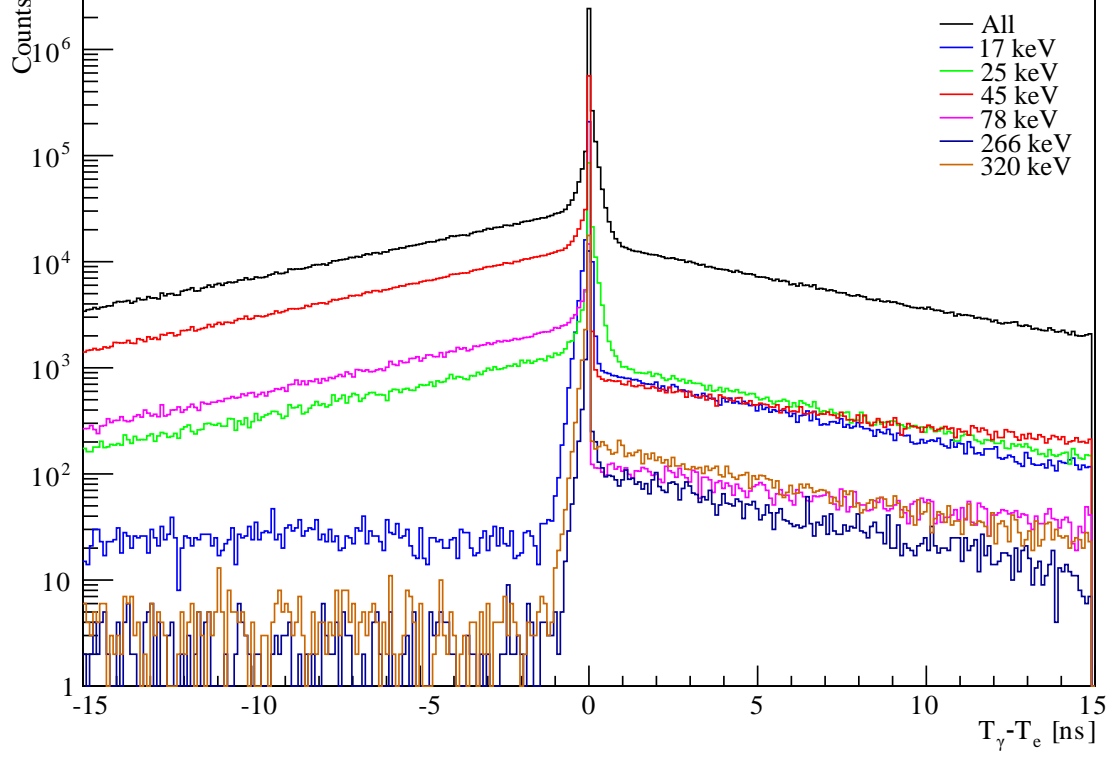
Figure 4.4: Ideal $\Delta T_{\gamma-\mathrm{e}}$ distributions for the prominent electron lines from ^{133}Ba decay. Due to detector resolution effects, electron lines within a few keV are grouped together as a single line and L X-rays are omitted.

Table 4.2: ^{133}Ba $\Delta T_{\gamma-\mathrm{e}}$ distribution characteristics

$E_{\mathbf{e}}$ [keV]	Mean $\Delta T_{\gamma-\mathbf{e}}$ [ns]	Std. Dev. [ns]
17	0.763 ± 0.005	2.704
25	-0.297 ± 0.010	4.570
45	-2.457 ± 0.003	4.141
78	-3.764 ± 0.011	4.536
266	1.245 ± 0.022	3.230
320	1.482 ± 0.019	3.493
All	-1.722 ± 0.003	4.299

charge collection characterization. Furthermore, the transition lifetime is 1.499 ns, much longer than the 300 ps allowed for the Nab experiment.

On the other hand, ^{133}Ba emits many different energy electrons between 17 and 320 keV, most of which come from states with lifetime below 200 ps. Unfortunately, almost all decays have to proceed through the 81 keV excited state of ^{133}Cs, which has a lifetime of 6.3 ns. Exclusion of this transition state is difficult, since most decays include this transition, and the main emissions, a 45 keV electron and an 81 keV gamma, overlap with other lines of interest. Fortunately, however, this only produces a long tail around the prompt peaks, so with a sufficient coincident rate, the fast

components should still dominate.

4.3 Experimental Setup

The prompt photon detector used for the tagged electron timing studies was a CeBr$_3$ scintillator, chosen for many reasons. After excitation, the decay to its ground state is fast, with an average lifetime of 19 ns. The crystal density is high (5.1 g/cm^3), so the detection efficiency is good, even for a small crystal. The light yield is high, with an average production of 60 photons per keV of deposited energy. Lastly, the energy resolution is good, with reported peak widths of approximately 4% for 662 keV gamma lines. Table 4.3 details the interesting properties of CeBr$_3$ scintillators.

Table 4.3: General properties of CeBr$_3$ scintillators.

Property	Value (units)
Decay Time	19 ns
Light Yield	60 (photons/keV)
Density	5.1 g/cm^3
Wavelength (Max. Emission)	380 nm
Wavelength Emission Range	$325 - 475$ nm
Energy Resolution @662 keV	3.8 keV
Background Radioactivity	0.004 Bq/cm^3

Due to its use of high Z materials, photoelectric absorption is the most probable photon interaction mechanism for CeBr$_3$ scintillators for photons up to about 200 keV (see Figure 4.5). Above these energies, Compton scattering is the dominant interaction mechanism, until pair production takes over after 7 MeV. The energy resolution as a function of incident energy is approximated by the function [48]

$$R(\%) \approx \frac{108}{\sqrt{E}}. \tag{4.3}$$

The energy dependence of the resolution is caused by a combination of non-proportionality of the response of the scintillator and statistical effects in the light collection process.

The scintillator, purchased from Advatech UK Ltd., was a cylindrical crystal, with a radius of 9 mm and a height of 12 mm. The scintillator was housed in an aluminum casing with a hermetically-sealed quartz optical window. The aluminum was 1.0 mm thick on the sides, while the entrance window was 0.5 mm. A layer of Teflon between the CeBr$_3$ crystal and aluminum housing served as a reflective surface to maximize the photon count exiting the housing through the quartz window.

Light readout for the scintillator was performed using a SensL ArrayJ-60035-4P SiPM. This board included four 6.07x6.07 mm^2 J-series SiPM dies, each with 22,292 microcells of area 35x35 μm^2, constituting an active area of 75%. The anode-cathode rise-time specification was 300 ps, with a microcell recharge time of 48 ns. As an evaluation board, the SiPM was already mounted to an 8-pin connector for signal readout.
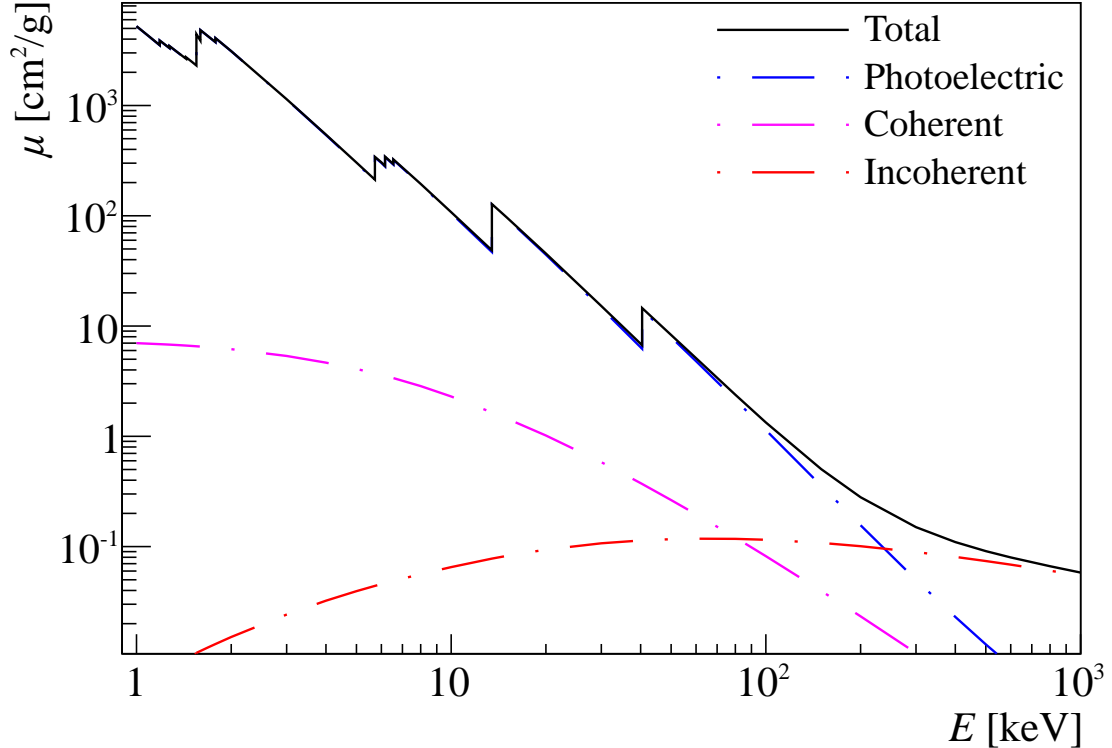
Figure 4.5: Photon cross sections as a function of energy in CeBr3. The crystal's density and use of high Z materials allow for high detection efficiency and a mostly photoelectric spectrum up to about 200 keV.

The readout electronics for the SiPM were prepared using conductive copper tape traces on a G10 substrate as a temporary printed circuit board, using a combination of through-hole and surface mount components. Per SensL literature, the circuit was designed with a shunt capacitor to match the SiPM pixel capacitance on the anode side and a resistor to ground on the cathode side, as shown in Figure 4.6 [49]; an additional resistor was placed in the bias line to filter transient signals. The signal was read out across a $50\,\Omega$ resistor, chosen to match the ADC input impedance. To simplify the electronics, the four SiPM pixels were junctioned, so the value of the shunt capacitor was chosen as 40 nF to match the capacitance of the parallel pixels. Lastly, the low-pass filter resistor was chosen to be 12 kΩ, resulting in a cutoff frequency of about 340 Hz. Table 4.4 lists the parts used and their purpose, and Figure 4.7 shows the circuit along with the scintillator.

A quasi-sealed ^{133}Ba source was used for prototyping the tagged electron timing apparatus. The source was sandwiched between two 500 nm Mylar foils with a 10 nm aluminum coating, and supported by a 1″ aluminum ring. A small vent-hole in the aluminum ring prevented accidental rupture of the films due to pressure changes. The source was mounted in the opening of the high voltage electrode of the west detector in the UCNB spectrometer, about 5 cm in front of the detector. The scintillator was affixed in front of the opening to maximize the source exposure, and thus, the coincident event rate.

Table 4.4: Components used to populate the SiPM readout board.

Component	Description
ED3108-ND	8-pin connector to interface with SiPM chip
478-10354-1-ND	0.039 µF±5% ceramic shunt capacitor
311-12KJRCT-ND	12 kΩ bias line filter resistor
541-2477-1-ND	50 Ω readout resistor
WM5544-ND	SMA jacks for bias and readout connection
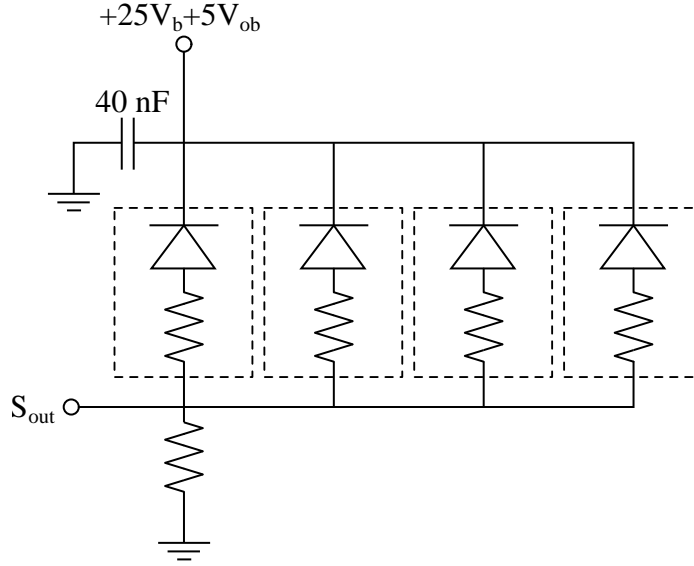
Figure 4.6: Basic readout circuitry for the SiPM as indicated by SensL [49].
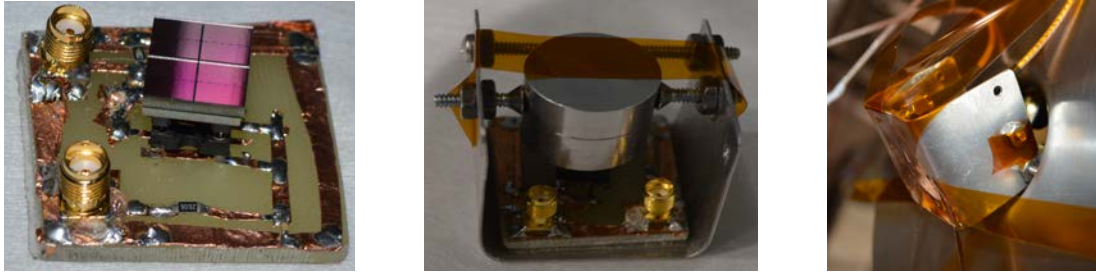
Figure 4.7: The prototype readout electronics for the SiPM were constructed with copper tape traces and surface mount components on a G10 substrate. The CeBr$_3$ scintillator was held against the SiPM using two set screws, and the whole assembly was taped against the opening of the high voltage endcap.

The bias pulser was attached to the detector for stability monitoring over the course of the experiment. The rate was chosen to be 2 Hz, so that the coincident event rate would not be affected by excess pulser triggers. Six pixels from the silicon detector were read to benchmark the timing source, and to begin an understanding of channel-by-channel variations within the detector and front-end electronics. The data acquisition system was configured to trigger on either the SiPM channel or a

TTL generated by the bias pulser; in the event of either trigger, the six silicon pixels would also be read at that moment. Approximately 30 hours of data were acquired in this first test.

4.4   Tagged Electron Proof of Concept Results

From the data taken with the ^{133}Ba source, just over 1.2×10^7 events were recorded in the SiPM, corresponding to a rate of approximately 110 Hz. From these events, about 20000 events were recorded in each pixel, corresponding to a per-pixel coincidence rate of approximately 0.2 Hz. Of these events, roughly 4% were discarded due to a data integrity issue discussed in Section 5.3.

The SiPM signals ranged in amplitude from 0-400 ADC bins, which suggested that each ADC bin might correspond to roughly 1 keV. The baseline RMS noise was found to be 1.7 ADC LSB, which was consistent with the digitization noise of the DAQ. The SiPM pulses were found to have an average rise time of 18 ns and an average fall time of 40 ns.

In order to understand observed measured energy spectrum in the scintillator, a Monte Carlo simulation was written in Geant4. The scintillator and housing were modeled according to the dimensions provided by Advatech and placed in a world comprised of the G4_Galactic material. The ^{133}Ba source was treated as a floating point source, and was situated at a displacement of 1.0 cm radially and 1.5 cm axially from the center of the scintillator.

For each event, particles were generated using the per-decay scheme in a modification of the EC Event Generation library used to build the time difference distributions. At the beginning of each event, a single decay was generated to determine the decay products emitted in that decay. The particles directions were independently generated according to a uniform, spherically symmetric distribution. A cumulative sum of energy deposited in the CeBr$_3$ scintillator was tracked and recorded in a histogram at the end of the event.

Given the true energy deposited within the scintillator, the next step was to modify the output spectrum to reflect the energy-dependent resolution mentioned above. For every bin i in energy, each event j was modified via the addition of a normally distributed random number with standard deviation $108\%/\sqrt{E}$:

$$E_{i,j\mathrm{mod}} = E_i + \mathcal{N}\left(0, (1.08\sqrt{E_i})^2\right)$$

A calibration for the measured spectrum was determined by iteratively fitting the three peaks at 35, 80 and 350 keV to a Gaussian distribution. Following the initial guess, each peak was fit to $\mu \pm \sigma$ until the values for μ and σ stabilized. Figure 4.8 shows a linear fit of the ADC value as a function of simulated energy. The error bars for each point are the fit errors of μ.

Based on this fit, then, the conversion from amplitude in ADC bins to peak amplitude in keV is calculated via

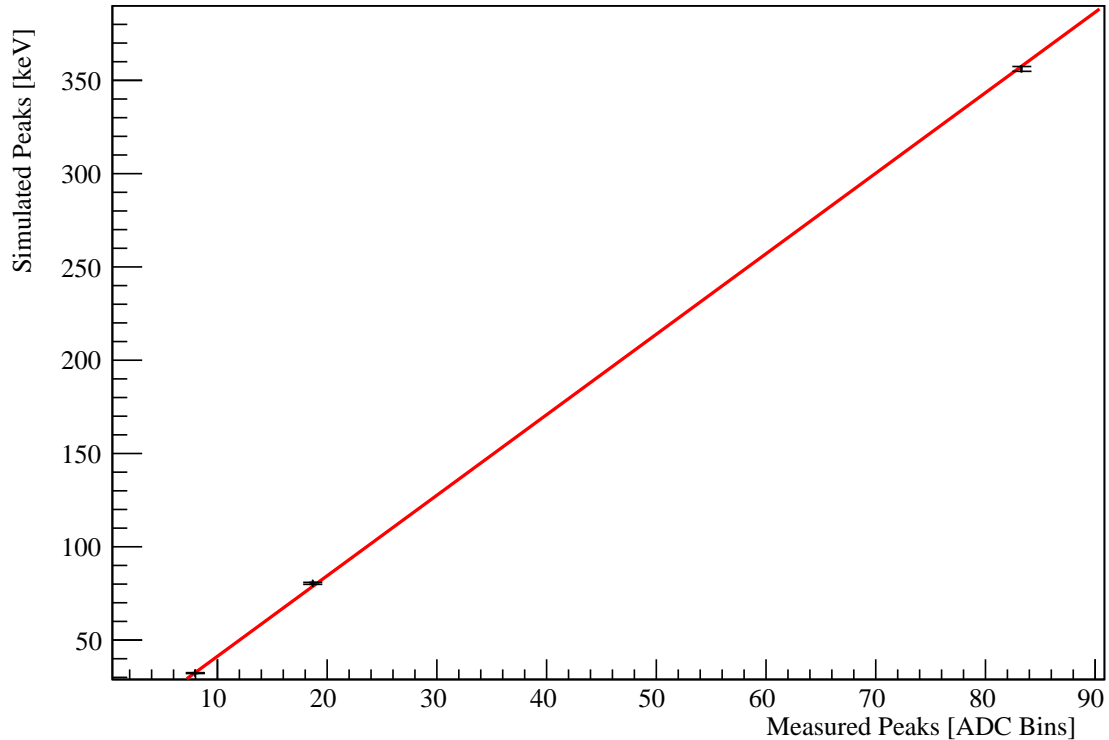$$\text{Energy (keV)} = 4.31551 * \text{Amplitude} - 1.85459. \tag{4.4}$$

Figure 4.8: A linear fit calibration relating measured ADC amplitudes to simulated energies. Each point was determined via relaxation of an iterative Gaussian fit for repeatability.

A comparison of the simulated spectrum and the measured spectrum is shown in Figure 4.9. For both spectra, the ~80 keV peaks were fit to a Gaussian distribution; the amplitude of the centroid of each fit was used as the scaling factor for that spectrum. While there was excellent agreement in the photoelectric peak amplitudes, the region from $100 - 250$ keV showed a discrepancy between the measured and simulated events. The two most likely explanations of this difference are the presence of some unknown contaminant in the source that contributed to the counts in that region, which has been seen in previous source experiments, and incompleteness of secondary tracks in the simulation.

Figure 4.10 shows the spectrum measured in one of the silicon pixels, as determined from offline analysis of the waveform data. A trapezoid filter with a rise time of 5 μs and a flat top of 1 μs was used to extract the pulse amplitude in each waveform; more details about this filter are available in Chapter 6. Three types of events were observed in the waveform traces: null triggers, coincident particles, and bias pulser triggers.

As a first pass analysis, the arrival times of particles in the scintillator/SiPM and in the silicon were estimated using an RC-$(CR)^2$ filter, which is discussed in Chapter 6. The trigger time was determined via linear interpolation to find the filter zero-crossing, as shown in Figure 4.11. While this method is susceptible to variations in the rising edge and poor performance due to high frequency noise, with both being

51
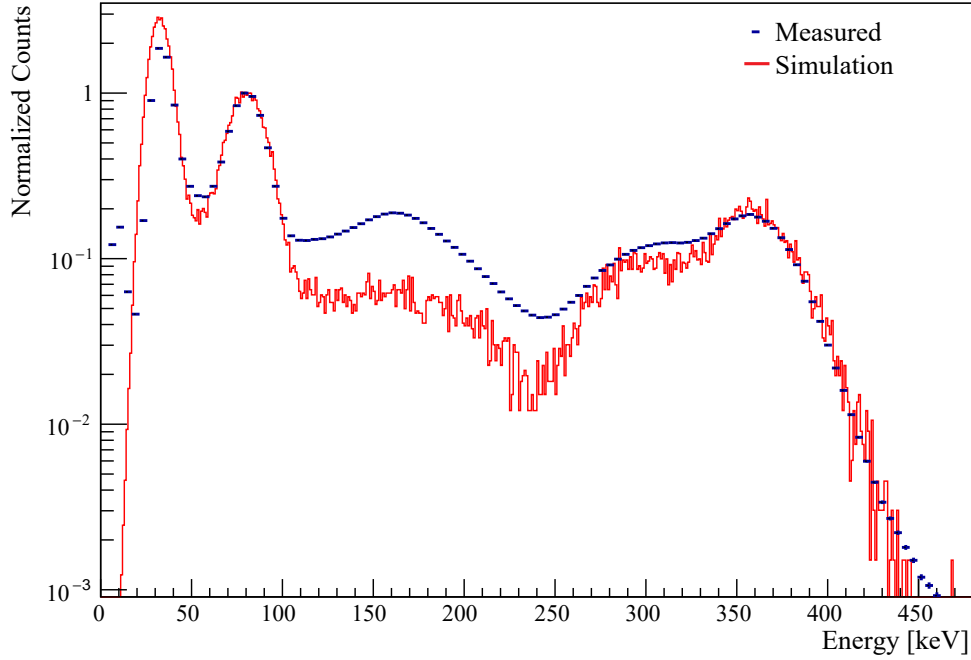
Figure 4.9: Measured and simulated ^{133}Ba spectra in the CeBr$_3$ scintillator, with both distributions normalized to the centroid of Gaussian fits to the $\sim$80 keV peak. While the photopeak intensities were consistent, the discrepancy from $100 - 250$ keV is likely due to the incomplete geometry of the simulation.

expected for the silicon detector signals, it was reasonable as a qualitative check of the system functionality.

The mean of the $\Delta T_{\gamma - e}$ distribution for the 45 keV events was found to be -178.05 ± 0.48 ns, with a standard deviation of 18.19 ns. For the distribution from the 78 keV events, the mean was -179.00 ± 0.91 ns and the standard deviation was 14.10 ns. Ideally, the offset between these distributions was expected to be around 1.3 s, which was consistent with the measured distributions. Ultimately, however, these data were statistically limited, so the main observation is that at the surface, this type of experiment may be useful for silicon timing characterization.

Figure 4.10: Summed coincidence spectra from the six silicon pixels

(a) CeBr$_3$/SiPM response.

(b) Silicon response

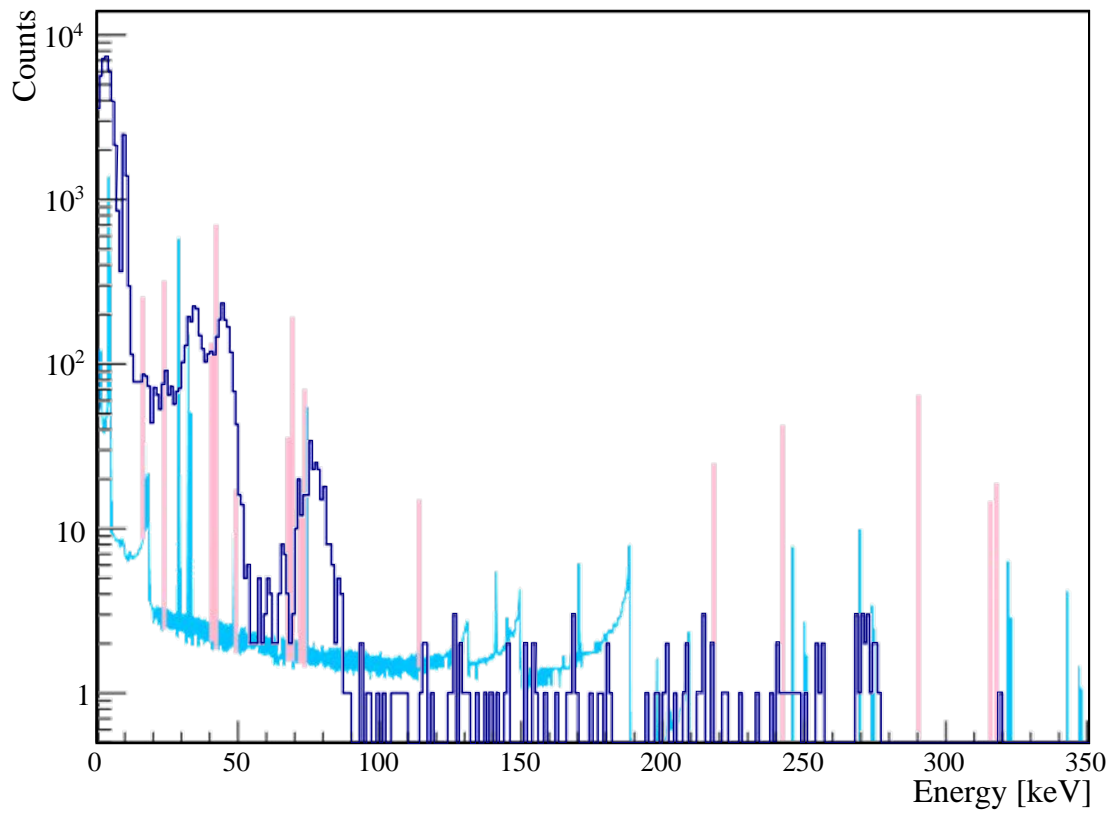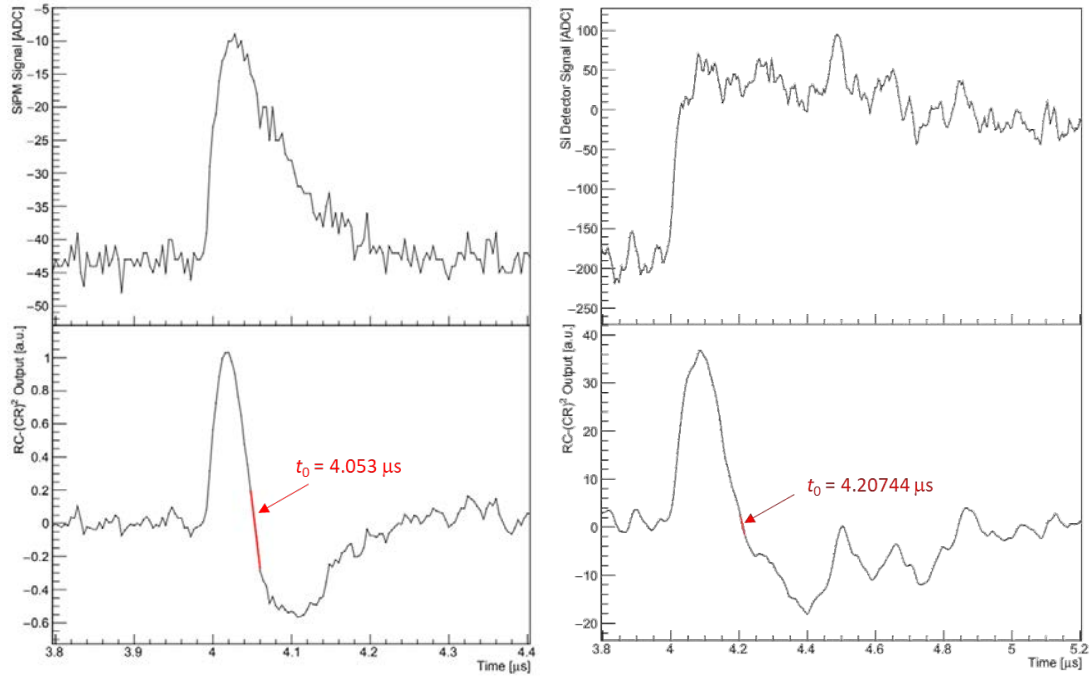Figure 4.11: Linear interpolation to find the zero-crossing of an RC-(CR)^2 filter was used to determine the particle arrival time for both the SiPM and silicon responses.
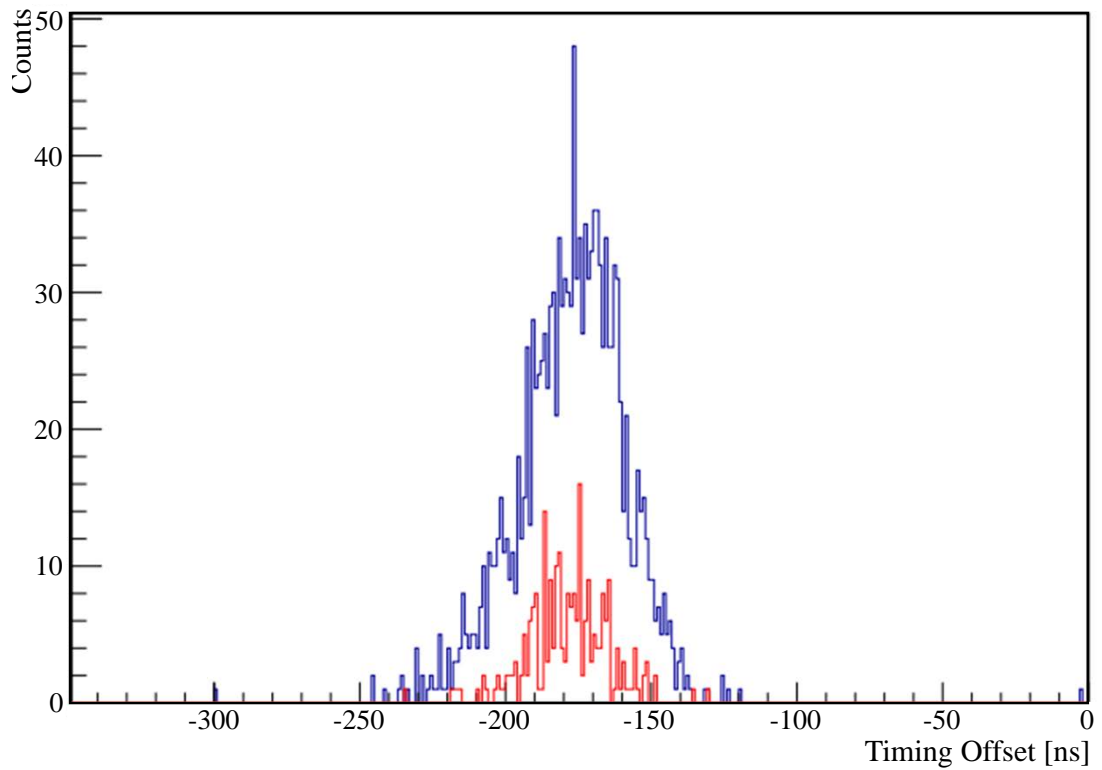

Figure 4.12: [^{133}Ba $\Delta T_{\gamma-e}$ distribution.

Chapter 5: Nab Data Acquisition System

The detection system performance metrics drive the data acquisition system (DAQ) specifications; additionally, the physics resulting from the spectrometer geometry further extends these requirements. As such, the DAQ capabilities are pivotal to the success of the experiment. From Chapter 3, the minimum rise time for the shaped signals from the detector is approximately 40 ns, which means that the signal should be digitized at least every 10 ns to acquire sufficient data on the rising edge for precise determination of particle arrival time. The experiment relies on high efficiency detection of protons, accelerated to 30 keV, just above the dead layer and noise floor, so the DAQ must have a robust low-threshold trigger, with high efficiency and energy resolution even in the case of electronic oscillations due to cooling.

In summary, the DAQ is required to meet the following capabilities in order to achieve the aforementioned experimental goals:

- At least 100 MS/s digitization rate, for determining the first hit of a backscattered electron

- At least 10-bit resolution for electron spectroscopy with 3 keV binning

- An efficient low-threshold (10 keV) trigger for proton event detection

- Global synchronization stable to 100 ps for TOF spectroscopy

- Multipixel logic for detectors separated by high voltage

5.1 Universal Architecture for Spectroscopy

The newest generation of DAQs feature much more powerful Field Programmable Gate Arrays (FPGAs), huge local memory banks, and high bandwidth data buses like PCIe over optical fiber, versus lower bandwidth parallel buses like VME. This opens up the possibility of new architectures to utilize these resources efficiently. Whereas old filter and trigger logic needed to be as lightweight as possible, with only a few tuneable parameters, we now have the opportunity for completely generalized programmable filters and trigger logic. While previous systems recorded brief fixed records in segmented memory buffers with physical trigger lines in between modules and low level trigger logic in firmware to validate events before readout, deep memory buffers and low-latency buses enable much more general and powerful trigger logic. Finally, the data throughput can be orders of magnitude higher than in VME systems, which are limited to 320 MB/s.

This new improved hardware: resource-rich polylithic FPGAs, GB banks of DRAM, and GB/s PXIe back-planes offer exciting new possibilities. Our goal was to design a novel approach to DAQ architecture exploiting the expanded capability of this new technology. While the limitations of FPGA logic still restrict real-time processing, resource usage can be less efficient to allow for more generalized filter design. To make

global logic less restrictive, the incoming data stream can be continuously written to a single deep ring buffer memory to provide time for CPU software to implement complex trigger logic and request arbitrary waveform segments, even in untriggered modules or noncoincident times.

The resulting new DAQ architecture and its hardware and software implementation developed in collaboration with National Instruments (NI) is described in this chapter. Rather than describing the system following the progression of input signal to output data, we break it down by hardware components. Because each component is completely modular, any modification of the system could be accommodated by considering the necessary I/O for the component.

5.2 Nab DAQ Hardware and Software Platform

The data acquisition system was built around the NI PXIe-5171 Reconfigurable Oscilloscope module. Each card provides eight channels of 250 MS/s, 14 bit ADCs backed by a Xilinx Kintex-7 410T FPGA. Sixteen of these digitizer cards are installed in an eighteen-slot NI PXIe-1085 chassis, with the extra slots used for the PCIe-PXIe 8381 fiber optic controller and the PXIe-6674T (PXIe-6672) timing and synchronization module in the primary (secondary) chassis. A host computer controls the hardware in both chassis, executes the global trigger logic for record acquisition, and disseminates the data to a waveform analysis and storage server. Figure 5.2 shows the fully instrumented primary DAQ chassis, with the controller module, sixteen oscilloscope modules, and the primary timing module. Figure 5.3 depicts a schematic of the full DAQ system, including both chassis, the host CPU, and the analysis and storage server.

The motivation for using the NI system was to have total control over the firmware and explore new ideas for the DAQ architecture. The learning curve for LabVIEW and LabVIEW FPGA was much smaller than that of a traditional hardware design language (HDL), so we were able to produce a basic DAQ system in a very short period of time. The abundance of libraries and examples provided from NI, as well as the information available on the users forum both helped develop our application. Based on these features, we were able to successfully implement and optimize this design of a universal DAQ for nuclear spectroscopy applications.

LabVIEW and LabVIEW FPGA

LabVIEW is a system-design platform and development environment for NI's graphical programming language, "G" [50]. Any LabVIEW program, known as a *virtual instrument*, or VI, consists of two components: the front panel, which serves as a graphical user interface, and the block diagram, which describes the data flow and manipulation from input to output. The front panel displays user input and output terminals, called *controls* and *indicators*, which can be any combination of buttons, text boxes, dials and graphs. The block diagram contains code, which consists of these terminals connected to other components (sub-VIs) by wires; as the program
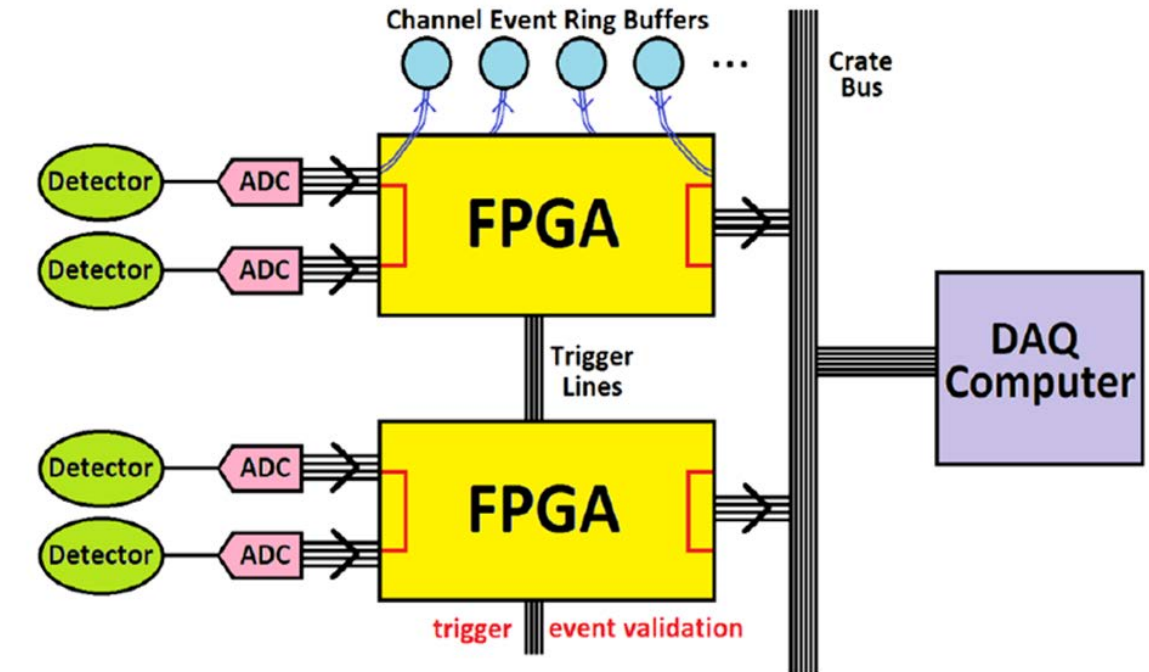
(a) Typical architecture
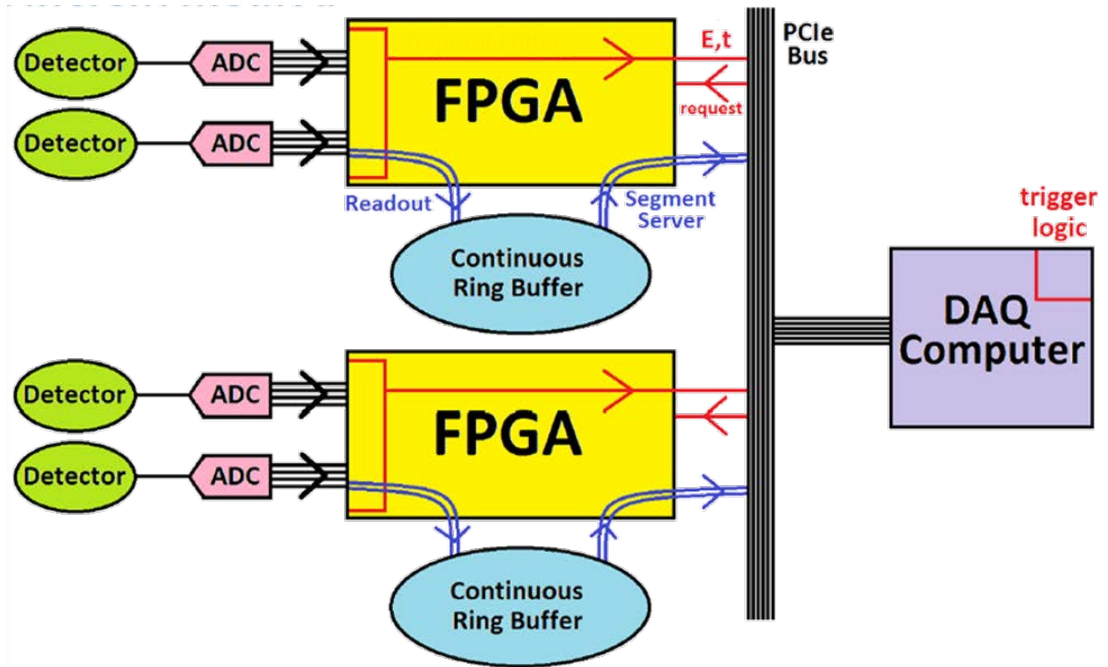
(b) Universal architecture

Figure 5.1: A comparison of traditional DAQ architectures, where global logic was achieved with physical trigger lines and only small records of specific channels were recorded in memory before readout, and our universal DAQ architecture, which continuously stores the last 400 ms of all channels to allow migration of all global logic to software on the host CPU.

Figure 5.2: A photograph of the instrumented primary DAQ chassis. The secondary chassis contains identical hardware, with the exception of the timing module installed in slot 10.
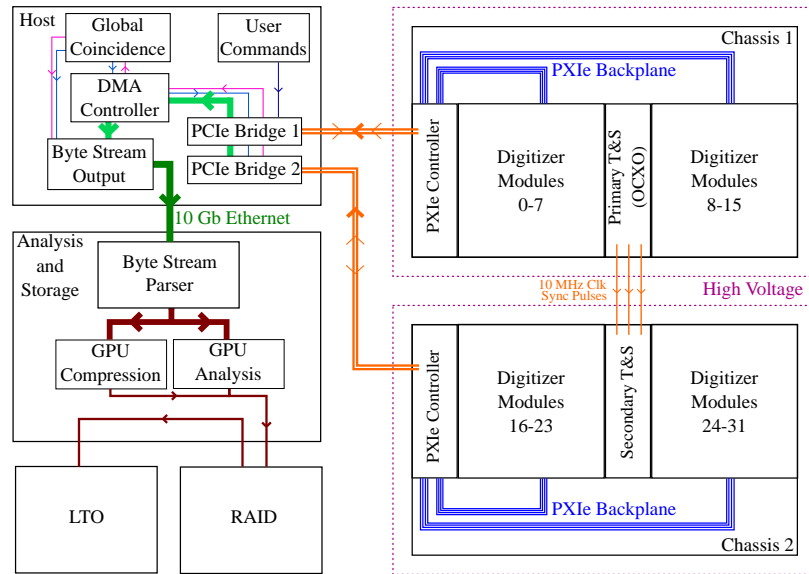
Figure 5.3: Diagram of the hardware connectivity for the DAQ.

is executed, data flows between nodes, and is modified according to functional operations of each component along the wire path.

Unlike code written in many other languages, LabVIEW is inherently parallel: independent wires are executed simultaneously, albeit asynchronously unless otherwise specified, until they reach a stopping point, such as an indicator or multi-wire functional input. In addition to dedicated user-interface and timing threads, the LabVIEW environment generates up to five execution systems, each comprised of four threads, per physical CPU [51]. Within a specific data flow branch, sub-VIs operate within the parent's execution system for efficiency purposes. However, the resource allocation is dynamic; if the code is written sequentially and does not lend itself to parallelization, only a single execution system will be established.

LabVIEW FPGA is an add-on to LabVIEW, and gives the capability of high-level FPGA synthesis. The data-flow and parallelization inherent to LabVIEW readily translate to an understanding of the implementation of code within the FPGA circuit. The module is used with NI FPGA devices, and can be considered at the levels of

the FPGA and host. The FPGA level instantiation of the code, which describes the logic to be performed on the FPGA, is converted to hardware design language and compiled to the binary run instructions. The host side version of the code details the interface to the FPGA hardware and establishes local resources necessary for the flow of data between CPU and FPGA.

The FPGA logic operates around single-cycle timed loops (SCTL), which are while loops operating at a specific clock frequency. In order to maintain absolute synchronization, all logic within each loop must complete within a single clock tick. Available libraries such as high-throughput implementations of the standard functions are available to optimize code implementation [52], but sometimes the physical signal propagation does not settle within the timing requirements. In these cases, there are two paths forward: the firmware can be recompiled, in the hopes that a shorter path may be generated, or signal pipe-lining can be introduced into the data path. Pipelining breaks up segments of code into multiple execution cycles with shift-register delays, which are registers to hold the signal between successive stages of the calculation during each clock cycle. In this way, the timing requirement is applied only to segments in between each shift-register rather than along the whole path.

When data needs to cross between single-cycle clock domains, or longer delays within a single clock domain are required, a first-in, first-out (FIFO) structure is employed. FIFOs can be implemented in shift-registers or Block Random Access Memory (BRAM), but if they cross clock-domains, they are required to be instantiated in BRAM. As the name implies, a FIFO is simply an ordered queue, in which the elements are read out one at a time after a fixed interval in the order that they were stored. In LabVIEW FPGA, when declaring a FIFO, the data type, requested number of elements, and implementation preferences are configured according to the specific needs of the application. A FIFO that operates within a single FPGA device is referred to as "locally scoped."

The host refers to the CPU that controls the FPGA device, and is responsible for bit-file selection and instructing the device to start running. In addition to being able to monitor the status of the FPGA, the host is able to exchange data with the device. The two mechanisms for passing data between the HOST and FPGA are controls/indicators and Direct Memory Access (DMA) FIFOs. Controls and indicators are just like the input and output terminals used in vanilla LabVIEW, but rather than a memory-based implementation, they are defined in FPGA as shift-registers. Accessing the values stored in controls and indicators involves non-deterministic timing, as there is an uncertainty associated with the time that it takes the CPU to posit the request and an unknown delay to latch the value in FPGA due to behind the scenes handshaking to validate the data integrity. As such, the use of controls and indicators are best utilized for applications where explicit handshaking has been defined by the user, or where the data exchange process need not be real-time in nature.

For applications that necessitate large transfers of data, or an exchange according to specific timing parameters, the use of a DMA FIFO is preferred. These FIFOs are implemented in BRAM onboard the FPGA and in DRAM on the host; a low-latency operation sends data as it is available from the FPGA BRAM to the host DRAM, or

vice-versa. The host software, meanwhile, reads and writes in large segments, so that data throughput is balanced with memory access latency. The advantage of such an architecture is that the abundance of memory on the host system allows for retention of data until the CPU is free to access it for use in software; likewise, data being sent to the FPGA can be held in memory until the FPGA is ready to use it according to a deterministic clock. Just like the FPGA-scoped FIFOs, the depth of the FIFO onboard the host can be set as needed.

The FPGA firmware was written in the NI LabVIEW 2015 FPGA module and compiled to a deployable bit-file with LabVIEW 2015 FPGA Module Xilinx Tools - Vivado 2014.4. While new versions of these software packages are available, it was discovered that some of the libraries used in the DAQ firmware were deprecated after the 2015 version, resulting in errors that prevented modification of the firmware. The Xilinx compile toolkit interprets the LabVIEW block diagram to generate intermediate VHDL code, which is then used for synthesis of the deployable bitfile. As the binary is generated, the toolkit validates that the code will execute within the timing requirements of each domain. Given the complexity and density of logic, the compiler was set to optimize for performance. This resulted in slightly longer build times, averaging about 2.5 h per synthesis, but allowed for greater probability of success per attempted build.

PXIe-5171 Modules

The workhorse of the DAQ system is the reconfigurable oscilloscope module. Each card contains eight channels of ADC wired directly to a single FPGA for real-time signal processing; additionally, a 12 Gbit DRAM bank wired directly to the FPGA serves for temporary waveform storage. The bank, comprised of two DRAM chips, each with 2^{24} 384-bit words, is used to store the latest 400 ms of waveform data from each channel in a giant ring-buffer while waiting to serve data requests from the trigger logic on the host.

The input coupling and voltage range are software configurable; for our applications, we chose DC-coupled input in the 5 Vpp (-2.5 V to $+2.5$ V) setting to read the full range of the detector electronics. The default calibration offset and gain settings were used, and the 120 MHz anti-aliasing filter was enabled.

Timing and Synchronization

The PXIe-6674T timing and synchronization module serves two functions: generation of an extremely stable global reference clock, to which all digitizer modules in both chassis are synchronized, and routing the back-plane trigger lines to the other chassis. The timing source onboard this device is an oven-controlled crystal oscillator (OCXO): the crystal used to generate the reference clock resides in an oven controlled by a temperature feedback circuit, which stabilizes the global reference clock to 50 ppb/year. This signal was routed out and split into two lines: one back to its own clock-in, and the other to the clock-in of the 10 MHz clock on the secondary PXIe-6672 timing module in the other chassis, which does not need its own OCXO. The

PXIe trigger lines 0 and 1 of the chassis backplane were routed out from the primary timing module into the secondary timing module and distributed to the secondary chassis backplane.

Three sets of Highland Technologies fiber-optic converters were used to carry the signals across the 30 kV barrier between the two chassis. Originally, three J720 passive converters were used to produce the fiber signals, which were read out via J730 receivers. It was found, however, that by forking the 10 MHz clock, the voltage was insufficient to drive the J720 diode, and so the converter was upgraded to the active J724 driver. Each of these devices was specified with a typical jitter below 12 ps. Due to the delays in these diodes and the lengths of the cables, the maximum offset between any two channels was 300 ns.
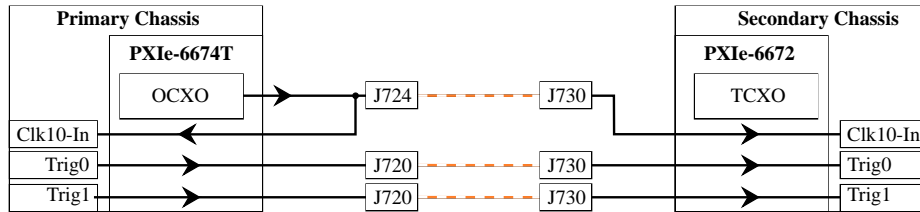

Figure 5.4: The hardware used for DAQ timing and synchronization. The PXIe-6674T generated a stable reference clock, which was distributed to the back-planes of the primary and secondary chassis for distribution to the 32 digitizer modules. Additionally, the device routed the PXI trigger lines for synchronization purposes. All signals routed between the two chassis used fiber converters to cross the high-voltage barrier.

Synchronization of the thirty two digitizer modules was achieved using the NI `Synch` API [53]. When the hardware is instructed to start a run, each module locks its ADC clock phase to the 10 MHz global reference clock using a Phase Lock Loop (PLL). At this point, the relative offset between any two modules is fixed. Next, each FPGA module derived a 25% duty cycle, 5 MHz TClk from its ADC clock. After initiation by the host CPU, one board designated as the synchronization master sends a pulse to the back-plane, instructing all boards to perform a synchronization measurement. Each board measures the time difference from the next falling edge of its TClk to the next rising edge of the global PXI Clk10, and reports the value to the host. This process is repeated many times, and the average value is used to adjust the TDCs to match the difference reported by the synchronization master. Figure 5.5 depicts the steps of this synchronization process.

NI provided a 300 ps specification in jitter of the relative offset after use of the `Synch` API. To verify this, we retrieved the measurements of the time difference between each of the falling edges of the TClk and the next rising edge of PXI Clk10, and repeated the measurement process several times to see how it varied. Figure 5.6(a) shows the average values for the relative offsets for multiple synchronization process iterations. After 2 synch processes, the spread was less than 30 ps. Figure 5.6(b) shows the distribution of offsets for one of the secondary boards after two synch processes, confirming that the jitter could indeed be restricted to 30 ps. The likely
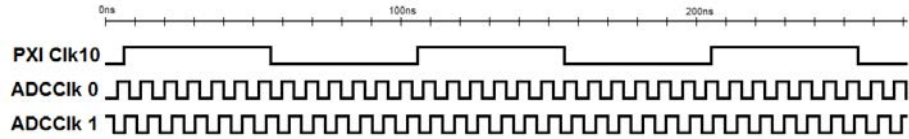
(a) Whenever hardware was initialized, all clock phases were random relative to each other. Step one was phase-locking the clocks to keep them fixed relative to each other.
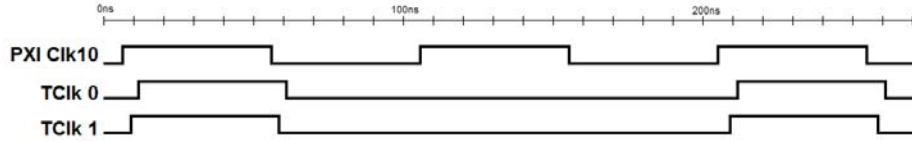
(b) Each board derived a 25% duty-cycle, 5 MHz clock called TClk from its own ADC Clock for precision synchronization.
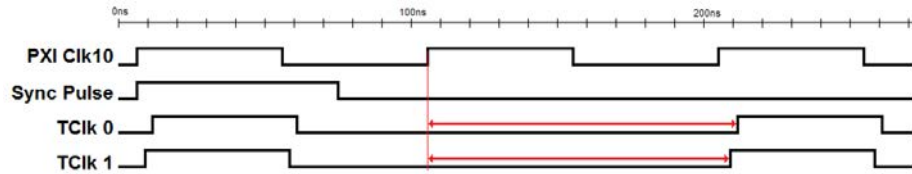
(c) After each Sync Pulse, every board measured the difference from the next rising edge of PXI Clk10 to its own TClk.
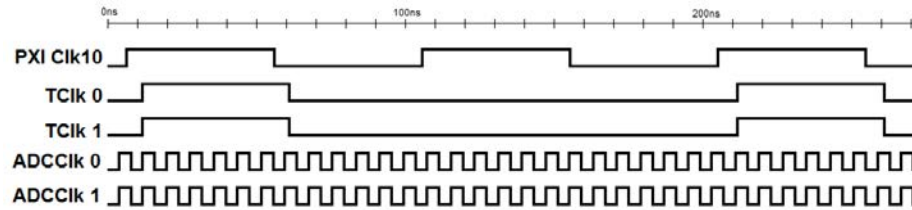
(d) The average time difference was used to fine tune the TDCs to phase align all boards with respect to PXI Clk10.
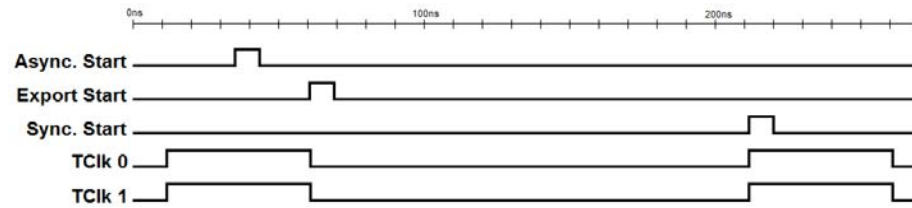
(e) At the start of data acquisition, all boards waited for the Export Start signal, and started on the next rising edge of their TClk.

Figure 5.5: NI's synchronization process used to phase lock and align all thirty two FPGA modules, as well as provide a low-jitter start of acquisition.

reason for this improved performance is that initially, the measurement logic may be split between two successive edges of PXI Clk10, causing an artificially large jitter. The correction from this first measurement shifts the TClk relative to PXI Clk10, so

that the measurements from the second synch process consistently hit the same edge. Ultimately, the DAQ hardware contribution to the timing uncertainty was at most 30 ps.

(a) Relative timing offsets (b) Relative offset distribution
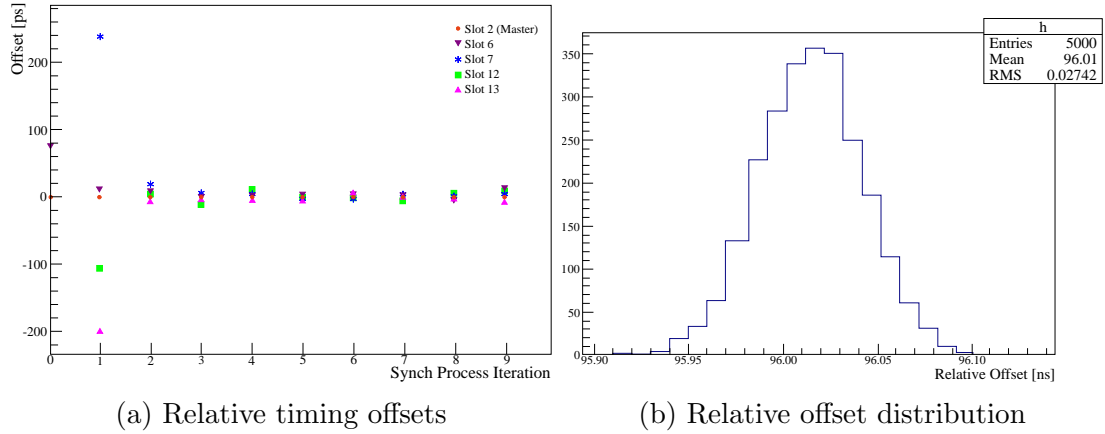
Figure 5.6: The result of multiple synchronization processes.

5.3 FPGA Firmware

The FPGA firmware was divided into three main components: hardware initialization, signal processing and triggering, and waveform serving. A 64-bit counter within the signal processing and triggering loop was used to give an absolute time stamp relative to the start of the each run. Figure 5.7 depicts the layout of logic within the FPGA.
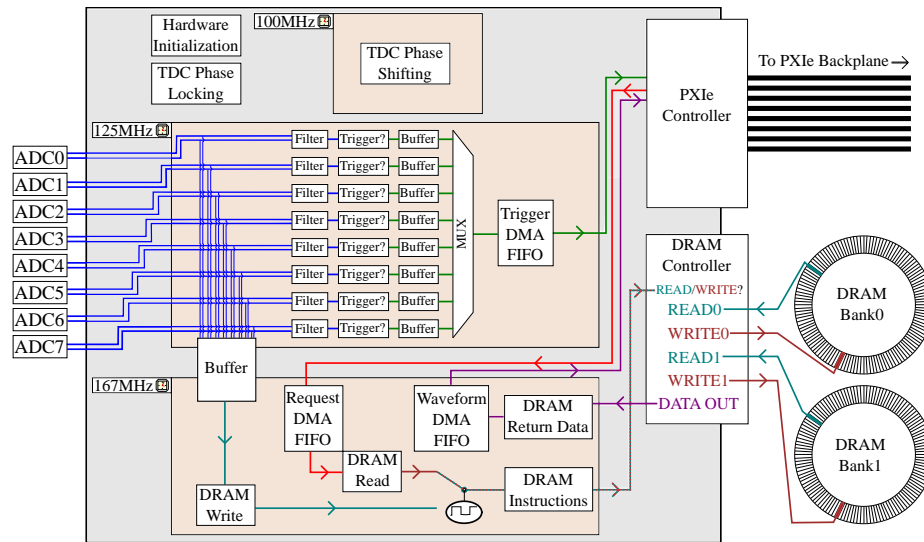
Figure 5.7: Logic diagram for the FPGA.

The hardware initialization loop operated on the FPGA base clock of 100 MHz. This loop was designed to establish the registers and clocks necessary for basic FPGA

functioning, the registers and FIFOS necessary for interfacing with the host CPU, and phase-locking and phase-aligning the ADC clocks on each FPGA device. This code came entirely from the NI Synch and NI Oscilloscope libraries, and consists almost exclusively of LabVIEW primitives, which are proprietary code blocks designed by NI.

The signal processing and triggering loop was directly tied to the ADCs, and as such operated on a 125 MHz clock dubbed the data clock. Starting with the arrival of the synchronized start signal, two ADC samples were read from each of the eight channels every clock tick, with flags indicating that a value was properly read and that it was within the dynamic range of the ADCs. Provided that both of these flags were valid, which was always the case, the data stream was then forked to the filter for L1 processing and to a small buffer instantiated with shift registers before being passed to the ring buffer. If, however, one of the flags registered as invalid, an error would be reported to the host CPU for diagnostic purposes. For the sake of efficiency, the incoming data were formatted as signed 14-bit fixed point numbers, with an integer word size of 1-bit; this translated to a dynamic range of -1 to $(1 - 1/2^{13})$ in steps of $1/2^{13}$.

Rather than attempting to perform filter logic on the full 250 MS/s data stream, the two samples from each clock cycle were averaged and processed in the data clock domain. The filter and trigger logic are detailed in Chapter 6. Any implementation must result in an energy and timestamp for the triggered channel. Triggers for the eight channels on each board were passed to the host via a single DMA FIFO, structured to accept 64 bit unsigned integers. The data structure of each trigger was comprised of two 64 bit integers: the first a bitwise concatenation of the 8 bit channel number and 16 bit energy, plus 40 empty bits, and the second a 64 bit timestamp. The remaining 40 bits were left empty, although their use for something like a trigger counter for systematic studies has been considered.

The DRAM ring buffer occupies both entire banks of DRAM adjacent to the FPGA. To efficiently use the 384 bit word length, we write six samples per channel to each address, with the samples 0, 2, and 4 in DRAM Bank 0 and samples 1, 3, and 5 in DRAM Bank 1. The forty eight samples are flattened to two 336 bit binary sequences, each concatenated with half of a 96 bit hash for data integrity, and written to memory as a single 384 bit word. The address is kept in a counter that increases by 1 after each write operation and wraps around after 2^{24} operations to cycle through all 2^{24} addresses in each bank.

Originally, the DRAM write operation was part of the 125 MHz loop, which simplified interface logic to the DRAM controller, but it was discovered during the ^{45}Ca run that this led to data corruption, discussed in Section 5.3. The solution was to buffer and pass the data to a separate waveform server loop on a separate clock domain. Two buffers, each capable of storing 1000 384-bit words, bridge the data from the acquisition loop to the waveform server loop. This buffer length accommodates the 600-sample round robin cycling between read and write access to the DRAM banks, for maximum efficiency during deterministic read/write cycles.
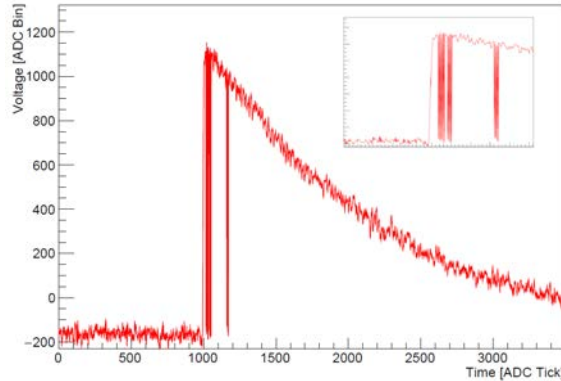
Figure 5.8: corruption

Restructuring the DRAM Buffer

A significant issue of occasional waveform corruption arose during testing of the DAQ under the high rate conditions of the ^{45}Ca experiment. These corrupt waveforms looked like a tail pulse, but with apparently random samples that dropped to a different baseline, such as in Figure 5.8. The corruption appeared to be rate dependent, as it was most prominent on traces retrieved from boards with high rates, but identification of the origin proved challenging. It was affected by many confusing factors, such as whether or not the mouse was moved around during data acquisition.

The first challenge in resolving the problem was identification of the corrupt waveforms. With the dramatic discontinuities present in the waveforms, an obvious first solution was numerical differentiation; any waveform which exhibited several significant steps after the initial rise would likely be due to waveform corruption. Unfortunately, this solution had two problems: First, numerical differentiation was only applicable for large amplitude waveforms, so only a portion of the corrupt events could be identified. Second, certain noise features such as transient ringing would create a false corruption trigger, which further complicated the situation. Another method that was employed was a comparison of the integrated area of the waveform versus the height of the leading edge; unfortunately this method was susceptible to the baseline oscillations, as they caused too large of deviations in the integral area to serve as a useful metric.

Ultimately, a robust corruption identification algorithm was employed, looking for excess amplitude in the high-end of the frequency spectrum. Using a Fast Fourier Transform (FFT), it was found that all corrupt waveforms included a peak in frequencies above 100 MHz. The cutoff frequency for the fast electronics was about 100 MHz, and the anti-aliasing filter at the ADC had a cutoff frequency of 120 MHz, so any significant features near the Nyquist frequency were likely due to data corruption. Figure 5.9 shows the FFT of waveforms with (red) and without (blue) corruption for a tail pulse and for baseline noise.

Further inspection of the corrupt waveforms indicated that the dropped bins occurred in groups of six samples, with either all six samples, the three even samples,
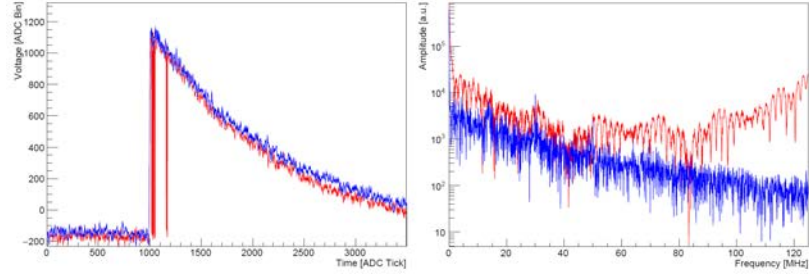
(a) Corruption in a tail pulse.
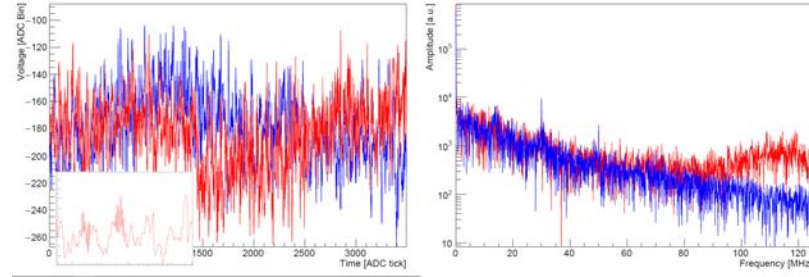
(b) Corruption in baseline noise.

Figure 5.9: FFT analysis of waveforms with (red) and without (blue) corruption. Even in baseline noise, the corruption creates a distinguishable feature in the region of $100 - 125$ MHz in the frequency space of the waveform.

or the three odd samples corrupted. Because the DRAM ring buffers onboard the FPGA were grouped in six elements per address, with the three even and the three odd samples written to separate banks, the connection to the DRAM access became clear. To further assess the issue, a second timestamp indicating the current time in the acquisition loop at the start of the waveform request was added to the waveform header.

The timestamps of events flagged as corrupt by the FFT analysis were compared against the reported DRAM read/write access records within the event headers, and it was found that whenever a corrupt event occurred, the waveform server was attempting to read from DRAM at the same time waveforms were being written, as shown in Figure 5.11. Investigation into the available literature on DRAM access within the FPGA indicated that the DRAM controller utilized round-robin arbitration in the event of simultaneous read/write requests [54].

While the DRAM write method utilized a *read request* method to query for data when the controller was available to perform the operation, the DRAM write method defaulted to a failed write if the controller was unable to process the data as it arrived. Furthermore, because each address contained six samples for each of the eight channels divided between the two DRAM banks, multiple samples could be lost. These facts were consistent with the observed nature of the corruption: even and odd sample corruption indicated a failed write to DRAM Bank 0 or 1, respectively.

To address the problem of data corruption, the data write process was pipe-lined and limited to periods when the DRAM read functionality was idle. Pipe-lining of the

66
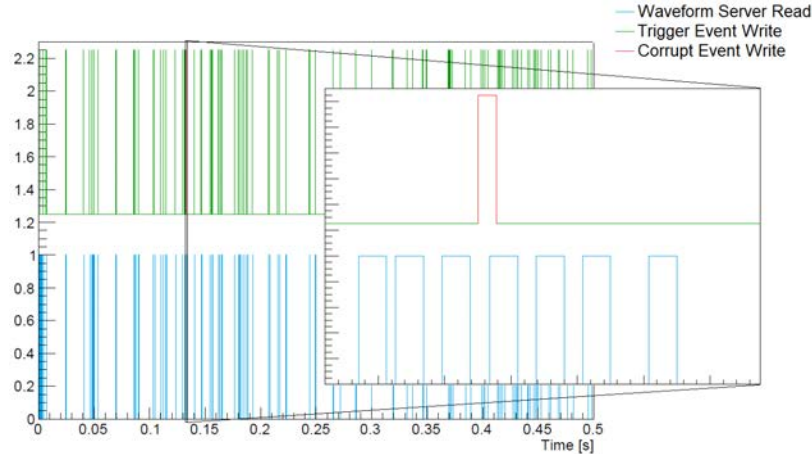
Figure 5.10: Using the FFT corruption flag,

incoming data stream was accomplished via instantiation of an intermediate buffer between the signal processing and triggering loop and the waveform server loop. The buffer length was a static allocation of 3000 events (3 events per 384 bit word) for each bank, so that adequate resources were available for accumulation while the read process occurred. Additionally, explicit read/write access was implemented via a flag set according to the status of the read utility. Write access was granted anytime the read utility was in an idle state, and a 600-sample choke was added to the read utility to manually arbitrate read/write access. Since the addition of these changes to the FPGA firmware, no corruption has been observed.

5.4 Host Software

Much like the FPGA firmware, the host software was partitioned so that LabVIEW's inherent parallelization could be utilized. After initialization of the hardware, three parallel processes are instantiated to handle user commands, trigger processing, and waveform readout. Initialization included designation of the hardware to be used, software parameter configuration, and instantiation of all buffers between the FPGAs and other computers. The synchronization process, which averaged 10 s per board, occurred in this portion of the code. Due to the substantial downtime accrued during initialization, the DAQ streams data continuously and run segmentation is performed separately on the data server.

Of the three parallel processes, the user command loop was the simplest. At programmed intervals, nominally once per second, the loop processed any change of state in the controls responsible for pausing the acquisition, manually triggering a user-defined collection of channels, or stopping the acquisition. If no change of state was detected by the end of the designated interval, the system returned the status of each of oscilloscope module, with details such as the acquisition status, the current state of the waveform server, and the status of each of the FIFOs. This diagnostic feature mainly served to debugging the code, but was kept in production mode, given the relatively low load and frequency of operation.

The second loop was responsible for processing the singles triggers, constructing different classes of events, and generating waveform requests. Optimal throughput was achieved by performing a bulk read of all triggers available from the Trigger DMA FIFO of each module. Because each trigger was comprised of two 64 bit integers, the FIFO with n 64 bit words contained $n/2$ or $(n-1)/2$ complete triggers, for n even or odd, respectively. Once read, each trigger was decomposed into its channel, energy, and timestamp. The thirty two FIFOs for all modules were processed, corrected for relative timing offsets, and concatenated into a single list, which was then time-ordered using the built-in LabVIEW sort function. This list was then appended to the last N triggers from the previous loop iteration to allow for continuity in global coincidence logic.

The trigger processing logic controlled the relationship between the singles triggers and the requested waveforms. The simplest version was a one-to-one relationship, where each trigger generated a single request. The entire UCNB data set was acquired using this method, as the DAQ software was still being developed at the time. A slightly more sophisticated version, in which complementary pairs of pixels were requested for each trigger, was implemented for the ^{45}Ca runs to aide in the analysis of systematic effects. Retrieving the opposite pixel for every trigger allowed for offline analysis of both waveforms for complete energy reconstruction, even if the pixel never triggered. Additionally, a decimated global trigger was encoded, where every 1000^{th} trigger generated a request for waveforms from all channels.

A robust trigger processor was developed to implement the global electron-proton coincidence logic required for the Nab experiment. The concatenated array of triggers is time-ordered using the `Sort 1D Array` function within LabVIEW, and then scanned for proton-like events. For development purposes, a nominal proton-like event was defined as an single-pixel event with energy in the range of 50-120 LSB. For each proton-like event, the previous 40 events were scanned for a coincident electron. If the timestamp of an event fell within the nominal 40 μs window preceeding the proton arrival time, then additional coincidence criteria were imposed. Otherwise, the proton-like event was discarded and the search continued for the next candidate.

The additional coincidence criterion of matched pixels was implemented as a truth table of pixel combinations loaded from a easily modifiable configuration file. For better intuition, the truth table was constructed using pixel names, rather than the board and channel numbers referenced by the DAQ software. Upon software initialization, a sub-VI maps the names to the appropriate device number according to the `multipixel_logic.config` file. While no energy requirements were dictated for the electron, they could readily be established as with the proton. If all criterion passed, a coincidence event was constructed, and the appropriate waveforms were requested for further processing offline. The multipixel requests were determined by a second truth table, read from the `multipixel_request.config` file, which used the same pixel names to describe the channels to be read. The waveforms from each channel were all read simultaneously in the current implementation, because the maximum offset of 300 ns between channels is much less than the 4 μs pre-trigger and 10 μs post-trigger.

If the timing window was not exceeded, or fewer than 40 previous triggers had been

processed, the algorithm would continue searching for additional possible matches. Due to the high event rates expected in Nab, some fraction of events are expected to meet basic $e - p$ coincidence criterion, creating false positives. The current logic was designed to allow for the recording of all false positives, as it was anticipated that applying a statistical correction to a complete set of data would be more effective than estimating the portion of false and true triggers being acquired.
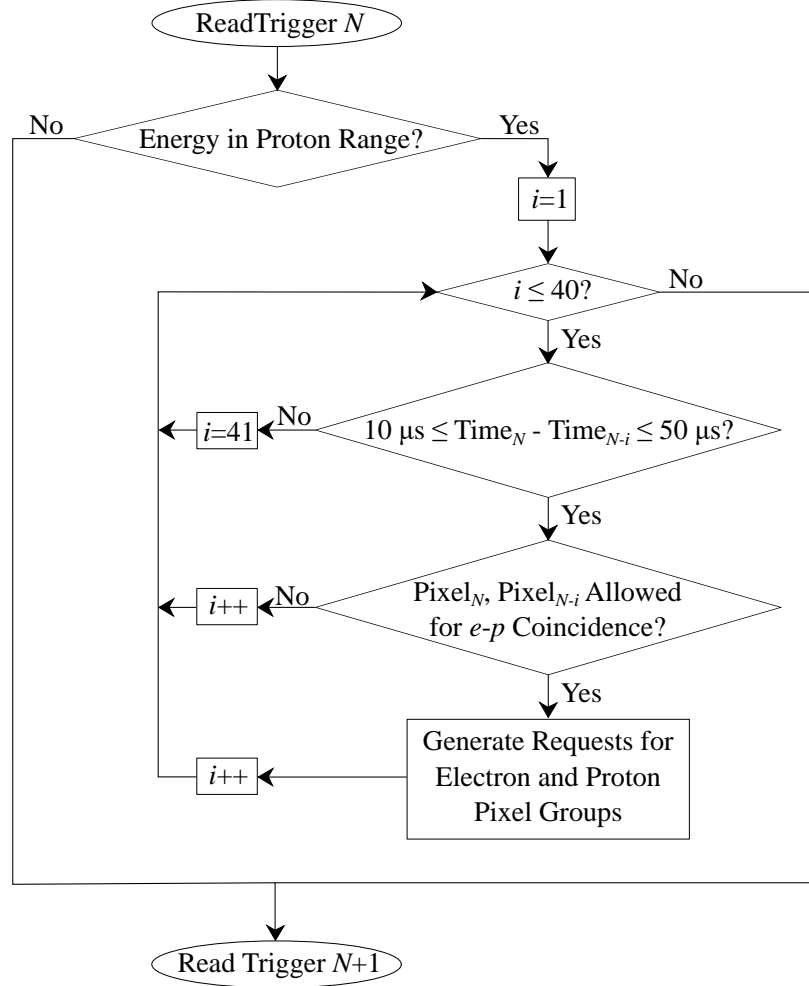
Figure 5.11: Nab electron-proton online global coincidence logic. The time-sorted list of triggers is parsed event-by-event for proton-like events, in which case the past 40 events are searched for possible electron matches. Because the logic is instantiated in software, it can be easily modified to meet experimental needs.

Testing of the coincidence logic was performed using a utility to mimic the distribution of $e - p$ coincidences in the Nab experiment. Random electron triggers were generated with the energy sampled according to the neutron beta spectrum and the pixel selected from a uniform distribution on both detector arrays. The time stamps were randomly assigned to mimic the average 50 kHz event rate. On average, 4% of the electrons were chosen to have a coincident proton, which was assigned a random pixel according a nearest neighbor geometry, with an energy normally distributed

about 18 keV. The time offset between the electron and proton was randomly assigned between $10 - 40$ μs. This list was then divided into the 32 boards, as in the actual DAQ program.

The test tracked the time to process the triggers and generate coincident waveform requests, as well as the success rate for the coincidence identification logic. Runs of 10^4 and 10^6 electron triggers were conducted to benchmark the processing capability of the software; it was found that an average of 400 μs and 40 ms for the two sample sizes, respectively, were required to sort, conduct coincidence logic, and generate waveform requests. Given that the current trigger processing loop requires less than 1 ms for completion at present, the anticipated increase in time for the full coincidence logic is expected to be negligible. From these same studies, it was found that false triggers occurred at an average rate of 1.5-2.0%, which is consistent with expectations based on the rate of low-energy electrons. In summary, the trigger logic is capable of meeting the Nab requirements, after tuning the coincidence definition and readout configuration files.

The third loop in the host software was responsible for readout of waveform data from the FPGA servers. This loop was the most time critical, as the volume of waveform data could easily exceed allowable memory allocations. The buffer size limitations were balanced against the latency associated with multiple DRAM accesses. For high rate data acquisition, waveform readout in batches of 500-1000 at a time provided the best overall throughput, without data loss. For lower event rates, however, waveforms should be read on a per waveform basis to ensure total readout.

When the DAQ software is stopped, all FPGA interfaces are closed to prevent resource locking between DAQ cycles. In this way, unprocessed waveform requests would not be read out after rerunning the DAQ.

The host utility can be operated directly using the graphical user interface or remotely through commands sent via a TCP/IP socket. The GUI implementation is more resource intensive, as it displays real-time updates for the DAQ, and requires redundant copies of all controls and indicators to be accessed on the front panel. However, for debugging purposes and local running, the GUI implementation is the preferred method for interfacing with the DAQ. The remote operation is the intended method of running for the Nab experiment, as all hardware will be controlled by the slow DAQ computer. In this mode, all data is transferred to the storage and GPU analysis servers via 10Gb ethernet, reserving processing power on the host for trigger and readout functionality.

5.5   DAQ Throughput Characterization

The expected data rate for the Nab experiment will depend heavily on undetermined background noise, but estimates based on Monte Carlo simulations indicate the range of 28-150 MB/s. The throughput in each stage of the DAQ was analyzed against the hardware specification to determine the efficiency of the firmware/software, and to ensure that the system can operate at these sustained rates without issue. The bottleneck of the system was the limitation of the host CPU to process waveforms from each FPGA FIFO, and the ability of a single FPGA to retrieve waveforms

from the DRAM buffer in excess of 150 MB/s. Other potential choke points, such as the trigger processing loop or the L1 filter and trigger logic on the FPGA were insignificant.

Considerations for the Future

While the DAQ exceeds all performance specifications for the Nab, and has been refined through the development process, there is still room for improvement, as the theoretical limitations of the system platform have yet to be reached. The single board throughput could be increased by restructuring the ring buffer to store 24 samples from a single channel per address, restructuring the waveform server DMA FIFO with 64 bit numbers, compressing the data onboard the FPGA before readout, optimizing the readout firmware to speed up the clock frequency.

The current ring buffer scheme requires one DRAM access to each bank per six samples, incurring a latency cost in addition to the time needed to parse the data and pass it to the host. From our studies of the waveform corruption, it was found that the latency per DRAM access is on the order of 60 ns; repackaging the ring buffer with one channel per address would reduce the DRAM access latency cost by a factor of 8, which, including the DRAM access latency, translates to a 55% reduction in per-waveform service for a 14 μs trace.

Modification of the DMA FIFO to a 64 bit architecture would allow four samples to be passed per clock tick, rather than the current implementation which allows for one sample per tick. The reduction in time necessary for passing a waveform scales linearly, so the portion of time spent writing to the DMA FIFO would be reduced by 75%. Unfortunately, the DRAM latency still dominates the overall process, so modification of the DMA FIFO word size would only result in a 28% net gain in throughput.

Another way to reduce the time spent writing to the DMA FIFO is to reduce the average size of each sample via compression. After the data is read from the ring buffer and before it is passed to the host via the DMA FIFO, a simple encoding scheme could be added. GPU tests of delta and Golomb coding on recorded data have achieve up to 70% data compression. Both of these algorithms have low complexity and could be implemented in FPGA as part of the waveform server loop. Although the DRAM latency is still the dominant contribution to the waveform service time, and compression would only improve throughput by about 26% in a single FPGA, the reduced volume of data would have a major effect on global readout over the PCIe bus, increasing throughput by 3.4 times. These savings would reduce the load on the entire system, including reduced data processing, storage on tape, and faster reading of data into GPUs for waveform analysis.

Finally, optimization of the code within the waveform server loop could allow for operation on a faster clock. The current 167 MHz domain is limited by synthesis of the firmware on FPGA. However, before the waveform server loop developed to its current state, it successfully compiled and ran at 200 MHz. A combination of code optimization and pipe-lining to segment the code within the loop could potentially speed up the clock domain, giving a minor boost to performance. Table 5.1

summarizes these four modifications, and the expected improvement to single-board throughput for each. In theory, implementation of all four changes could push the single board throughput to well over 1 GB/s, at which point the hardware limitations of the PXIe controller would becomes the bottleneck even for a limited number of modules.

Table 5.1: An overview of single-board throughput optimizations for future development, based on readout of a 14 µs trace.

Optimization Description	Service Time (µs)	Speed-up (%)
Current Implementation	56	-
DRAM ring buffer with 1 channel/address	25	55
DMA FIFO changed to 64-bit word size	40	28
Delta/Golomb compression onboard FPGA	41	26
Clock domain speed up to 200 MHz	52	6
All four changes	6	90

A significant improvement in host-side performance, could be achieved by migrating to LabVIEW's in-place structure. One of the pitfalls of coding in LabVIEW as opposed to C is the lack of low-level transparency of memory management. In C, a variable can be modified in place by a function using pass-by-reference. Similarly, in LabVIEW, one can use the in-place construct to directly operate on variables in memory, rather than redundant copies. In particular, the use of in-place operations on DMA FIFOs requires the use of the "Acquire Read Region" method for FIFO access. By passing a reference to the FIFO values rather than copying the data in the trigger processing and waveform retrieval loops, the global throughput of the system can likely be improved, although it is uncertain by what amount. It is worth mentioning that the current global throughput uses a considerable fraction of the 10 Gb ethernet bandwidth, so that improvement of readout performance would require additional hardware.

5.6 Random Tail Pulse Generator

Using NI's FlexRIO FPGA platform with a digital to analog (DAC) adapter module, a reconfigurable detector emulator was developed from the ideas for synthetic signal generation described in section 3.6. Unlike the PXIe-5171 oscilloscopes used to develop the DAQ, NI's FlexRIO architecture consists of dedicated FPGA modules with flexible I/O adapter modules according to application need. We purchased a PXIe-7962 FPGA board with a Virtex-5 SX50T FPGA and 512 MB of adjacent DRAM for early prototyping of the Nab DAQ. After moving to the integrated PXI 7151 platform, this FGPA board was available for use with an Active Technologies AT-1212 DAC adapter, which provides two channels of 1.2 GS/s DAC with an analog bandwidth of 400 MHz. The firmware was written and compiled using the same software version used for the DAQ.
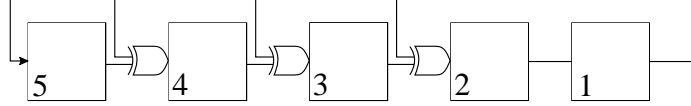
Figure 5.12: A 5-bit Galois LFSR with 4 taps, as determined by the polynomial $x^5 + x^4 + x^3 + x^2 + 1$.

The detector emulator logic consists of two main components: tail pulse generation and pseudo-random number production. The efficiency of the recursive implementation of linear convolutions used for synthetic waveforms in CPU readily translates to FPGA, so we adopted a similar scheme here. The current pulse width and simulated electronic shaping time are set by the user at run-time, as are the desired simulation rate and output pulse amplitude.

A pseudo-random number generator (PRNG) was developed utilizing an algorithm to generate a sequence of numbers whose properties are approximately the same as a true random number sequence, but any sequence can be exactly reproduced from run to run. PRNGs find utility in a variety of applications, including Monte Carlo simulations for statistical analysis, procedural generation within video games, and cryptographic encoding of information. PRNGs are limited to a finite period, exhibit correlations between successive values, and can fail to produce a uniformly distributed sequence of numbers. Various algorithms exist to generate pseudo-random numbers, and while many can overcome the aforementioned shortcomings, the cost of improved performance is an increase in calculation complexity. As such, the choice in PRNG is very much application and architecture specific.

A Galois linear feedback shift register (LFSR), named for the French mathematician Évariste Galois, is a typical FPGA-friendly PRNG. Each time the LFSR is clocked, all bits are shifted one position to the right, with the exception of the right-most bit, which is rotated to the left-most position. Special bits known as taps, whose locations are determined according to the feedback polynomial, are XORed with the right-most bit prior to being shifted, as shown in Figure 5.12. The maximal-length LFSR is achieved if and only if the feedback polynomial is primitive, which requires an even number of setwise coprime taps. The maximal-length feedback polynomials are not unique, but tables of optimal tap locations are readily available [55].

While a Galois LFSR is capable of generating a pseudo-random sequence of bits, a single LFSR does not necessarily produce a good pseudo-random sequence of numbers, especially not one for use in a Monte Carlo generator. However, the quality of an LFSR PRNG can be improved via multiple shifts per clock cycle and by XORing the output of two or more different sized LFSRs together [56].

This project can be extended to develop a fully functional detector emulator. Additional features that would be necessary for this include the addition of a correlated random trigger on the second output channel and the superposition of baseline noise on the output. The foundation for these efforts already exist, so it should be a small step to transition to this vision for the project.

Chapter 6: Developing Robust Digital Signal Processing Techniques for Physics

A vital part of data acquisition and online data analysis is the ability to determine signals of interest in real-time. When the signal-to-noise ratio (SNR) is large, this can be accomplished simply by discriminating the crossing of an amplitude threshold. However, as experiments progress in complexity and requirements in precision become more strenuous, greater sophistication in local trigger logic is necessary to observe low amplitude events. Digital filters can greatly improve the SNR and thus energy and timing precison, allowing for lower threshold, higher efficiency triggers.

This chapter discusses digital filters commonly used in nuclear physics for energy and time extraction, followed by some specialized techniques for precision measurements in the presence of noise. These methods provided the foundation of our efforts to develop more robust filters for the low-threshold trigger logic. We implemented these algorithms as universal filters in the digitizer FPGA logic that can be adapted to a wide range of algorithms without firmware recompilation. Changing the filter impulse response only requires providing new parameters at run-time, so multiple filters can be compared for performance with minimal loss of time. More importantly, as the filter is comprised of finite impulse response (FIR) segments, it is guaranteed to be stable as long as the coefficients are properly designed. Because any FIR filter can be approximated by our technique, we refer to it as a universal filter.

6.1 Fast Filters for Energy and Time Determination

Numerous FPGA-friendly finite impulse response (FIR) filters have been developed for the determination of energy and T_0 for particle detectors, and are well-documented in literature [57, 58]. Each of these filters is designed with considerations for real-time performance with high sample-rate digitizers and limited resource usage for implementation on an FPGA. Unfortunately, most filters are optimized either for time or for energy, but not both. Timing filters are designed to maintain the integrity of fast signals, so they do not attenuate high-frequency noise as well as energy filters; similarly, the low-pass characteristics of energy filters cause a temporal distortion that obfuscates relevant information with regards to particle arrival. Consequently, most commercially available DAQ systems integrate two parallel filters to calculate time and energy.

For semiconductor detectors in particular, flat-top trapezoid or cusp-like filters are commonly used for event energy determination, as they optimally improve signal-to-noise for step-like signals mixed with white noise, correct for the typical exponential baseline in peak integration, and provide a steady-state integrated amplitude for ballistic-deficit immunity. When programming these filters, the free parameters are the trapezoid risetime and the flat-top duration, plus an additional decay lifetime constant parameter characteristic of the detector front-end electronics. The risetime determines the amount of baseline and pulse averaging, thus noise attenuation of the

74

filter; a brief scan of peak resolution as a function of risetime can be used to determine the optimal value. The appropriate length of the flat-top duration is determined by the total rise time of the signal; the flat-top should be long enough to allow full-integration of charge, but not so long that it rejects the high amplitude segment of the tail pulse immediately after the rising edge. Several FPGA-friendly algorithms for these filters exist, although some were designed exclusively for offline analysis [42, 59–61]. Figure 6.1 shows the filter output for four different cusp-like filters.

(a) Trapezoid   (b) Cusp-like

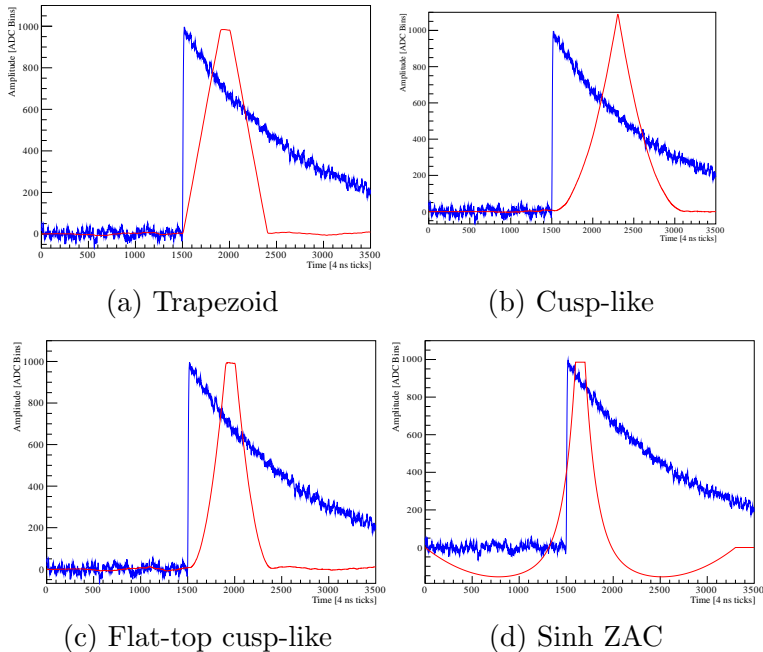(c) Flat-top cusp-like   (d) Sinh ZAC

Figure 6.1: Commonly used filters for energy extraction from solid-state detector signals. The cusp-like shape optimally reduces noise contributions, while the flat-top prevents ballistic deficit. The trapezoid, cusp-like, and flat-top cusp-like filters utilize FPGA friendly algorithms, but the hyperbolic ZAC filter as implemented is intended purely for offline analysis.

Other energy extraction methods use smooth, slowly changing filters such as Butterworth, Bessel, or Gaussian low-pass filters. These filters simply remove high frequency noise from the signal to reduce the variation in extracted time and energy. Recursive implementations of filters based on analog concepts are often developed using the bilinear transform $s = (1 - z^{-1})/(1 + z^{-1})$ such as the Butterworth and Bessel filters developed by the `mkfilter` utility [62, 63], while the Gaussian filter algorithm is developed using the backward difference technique [64]. Although none of these filters are optimal for FPGA implementation, modifications have been made to produce several successful approximations [65]

The last type of filter is meant to determine the particle arrival time, often to sub-sample precision. These filters are designed to determine the particle arrival time using a weighted differentiation. While a simpler finite difference algorithm could be used, such a technique would be susceptible to SNR issues present in most

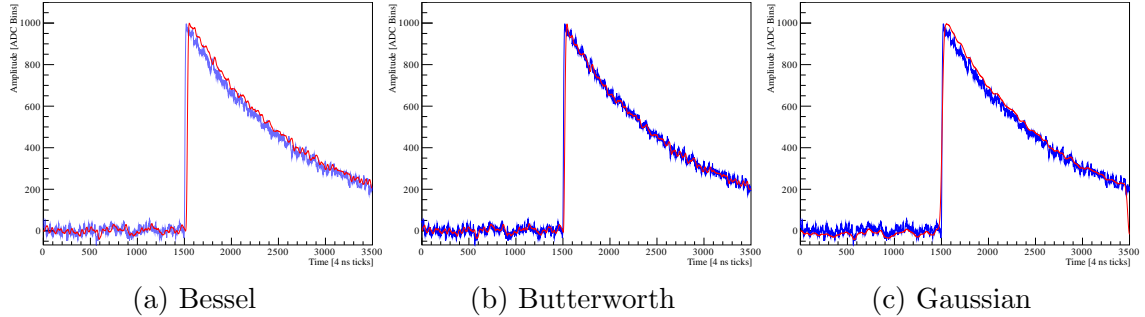(a) Bessel   (b) Butterworth   (c) Gaussian

Figure 6.2: Bessel, Butterworth, and Gaussian smoothing filters for improved noise performance. Each of these low-pass filters attenuates the contributions of high-frequency noise while mostly preserving the desired pulse-shape. Temporal distortion of frequency components near the cutoff frequency is unfortunately an unavoidable consequence, so these filters are usually used for slowly varying signals where timing information is not as critical.

experiments. As such, algorithms such as Gaussian weighted derivatives or RC-$(CR)^2$ differentiation are used to determine the point of maximal slope or the inflection point within the waveform [64,66]. Alternatively, another method called constant fraction discrimination is used to differentiate with a delay and a scaling factor [67]. This method is favorable, as it is able to consistently determine the time at which the pulse amplitude reaches a specific fraction. Unfortunately, a major drawback to these techniques for timing in the Nab experiment is that they rely on a relatively consistent pulse shape for optimal results. As this is not the case due to the detector physics, some other approach will likely be necessary. Nonetheless, timing filters such as these are fairly common place for nuclear spectroscopy; the filter responses are shown in Figure 6.3.

One of the earliest challenges in choosing the filter and trigger logic for the Nab DAQ was how to maintain the noise suppression of the trapezoid filter but easily determine the arrival time of the particle; furthermore, each channel was offset by a DC voltage, which varied from run to run, so baseline subtraction was also required. After some careful thought, the decision was made to approach the problem using the wide-delay derivative of the trapezoid filter, resulting in a double trapezoid shape. Two unit boxcar filters of opposite polarity and separated by some delay were convolved with the output of the trapezoid filter to produce a filter with the noise suppression characteristics of the trapezoid but with a sharp zero-crossing for timing extraction. Figure 6.4 shows the output of the double trapezoid filter after convolution with a tail pulse.

This double trapezoid served as the online filter for the entire UCNB and ^{45}Ca data-set. The rise time and flat top were optimized just like the single trapezoid, while the delay was set to the length of one full trapezoid. The cost of this filter was primarily in fast pileup detection, as it was about 6 μs in length, approximately 50% longer than the single trapezoid. However, this was not a concern in UCNB where the event rate was extremely low, and in ^{45}Ca, the acquired traces were 14 μs in length,

(a) Gaussian 1st derivative (b) Gaussian 2nd derivative
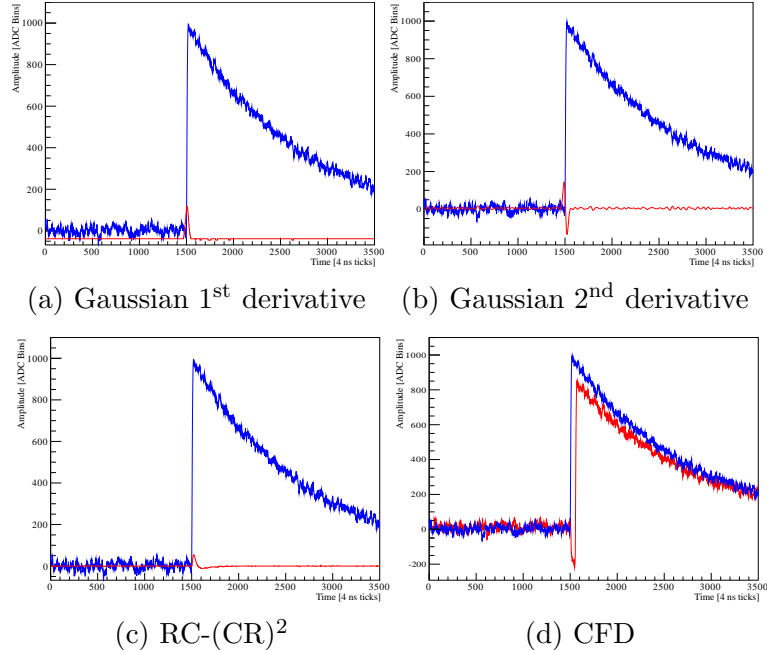
(c) RC-(CR)^2 (d) CFD

Figure 6.3: Commonly used filters for particle arrival time determination. Due to the step-like response of the signal, they all utilize some form of weighted numerical differentiation to detect the maximal slope or inflection point within the waveform. However, the filter output amplitude is substantially lower than that of the integrating filters, as the steps between points, and thus the derivatives, are only a fraction of the original pulse height.

so any pileup events could be found in an offline analysis. For both experiments, the trigger logic discriminated on a rising edge through a negative threshold to arm the trigger, followed by the actual trigger at the zero-crossing. The time of the trigger served as the T_0 for the event, and started a countdown to the middle of the top of the positive trapezoid. The output amplitude at this location was recorded as the energy of the event. This double trapezoid energy and timing filter provided an efficient low-threshold trigger for both of these experiments.

A collection of these filters written in C++ is available online [68]. Both a recursive and an FFT implementation are available, as architectures such as CPU and GPU run more optimally on one versus the other. The code was developed and tested on Ubuntu 14.08 using the `g++` 4.7 compiler. Version 3.3.7 of the Fastest Fourier Transform in the West (FFTW) library was used to calculate the Fourier transforms within the code [69]. It has since been verified to be compatible with Ubuntu 16.04 and `g++` 5.4. A makefile is used to compile and link the code, creating the main executable in the `bin/` directory. The default example shows how to call each of the routines.

The algorithm for each recursive implementation comes from the references cited in this section. Due to the user configurable range of parameters for each of filter, the FFT implementation of each was determined starting from the recursive form. The impulse response function was determined by convolving the appropriately configured
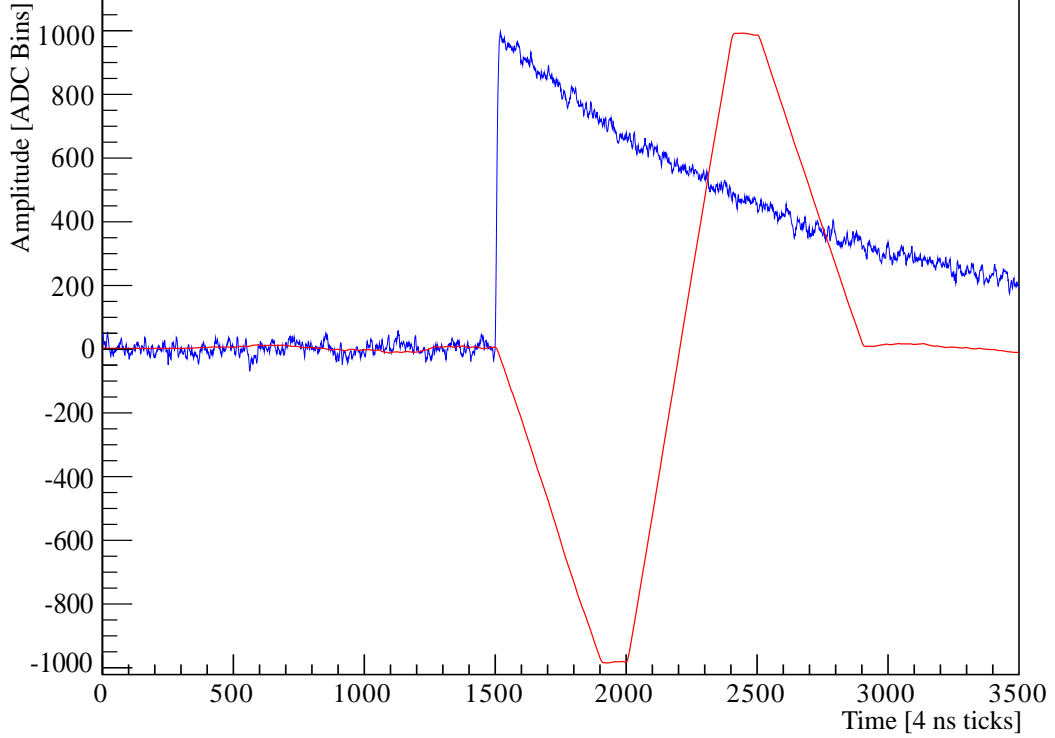
Figure 6.4: Double trapezoid filter output.

filter via the recursion relation with a Dirac delta function. The FFT of this output was then calculated, and point-wise multiplied with the FFT of each the raw waveform trace to be analyzed. Finally, the inverse FFT was calculated to determine the filter output, per the Convolution Theorem:

$$\mathcal{F}\{f * g\} = \mathcal{F}\{f\} \cdot \mathcal{F}\{g\}. \tag{6.1}$$

FFT filter equilibration bands at the beginning of each waveform trace were prevented by padding the beginning of each waveform with a blocks of 0's the same size as the filter. After convolution of the filter with the signal and the subsequent inverse FFT, the padding was removed to return the waveform trace as expected. This was done to maintain the original position of the signal within the waveform trace and prevent the logic from falsely flagging events in the filter equilibration region. The parallel architecture of GPUs naturally favors the FFT implementation of convolution; thus, if this library were to be used for such an effort, the FFT of the filter response only needs to be calculated once at the beginning of the program.

6.2 Waveform Analysis in a Low Signal-to-Noise Regime

One challenge of any precision spectroscopy measurement is the presence of noise features superimposed over the desired detector response. While the aforementioned filters can attenuate problematic noise components relative to the desired signal, the extracted energy resolution, low-threshold trigger efficiency, and timing accuracy and

78

precision inevitably suffer. Much of the detector data taken as part of the UCNB, ^{45}Ca, and Nab-prototyping efforts was plagued by high frequency baseline oscillations. These oscillations were found to be the result of microphonics: whistling within the cryogenic cooling lines generated mechanical vibrations, which coupled to the electronic detector signals. The oscillations were transient in nature, exhibiting variation in both amplitude and frequency, but generally their presence was characterized by sinusoidal features of amplitude 50-100 ADC LSB, and frequency 25 kHz with LN$_2$ cooling or 68 kHz with He gas cooling.

Fortunately, these problems are fairly commonplace, so many methods exist for signal processing in the presence of problematic noise features. The most successful solution if resources and time permit is to manually fit the noise features, subtract them from the trace, and then perform an analysis on the residual waveform. This technique can be slow and cumbersome, as the fits can include non-linear parameters such as the frequency of oscillation or time-varying amplitude; moreover, the algorithms are resource intensive, and do not lend themselves naturally to a high-throughput, real-time FPGA implementation.

Given the prevalence of the problem, however, several approaches for removal of features such as baseline oscillations and real-time fitting have been implemented. One technique used the digitized input from an adjacent mechanical sensor to provide feedback to an adaptive filter for active microphonic cancellation [70]. Another method decomposes the least-squares fitting process into two steps, one to estimate the signal amplitude and the other to precisely determine the arrival time of the event [71]. The most robust implementation utilizes iterative least squares fitting to 100s or 1000s of parameters on a heavily down-sampled data stream to subtract the noise components before using a more traditional filter for energy and time determination [72].

Ultimately, the pitfall of the aforementioned techniques is their rigidity of response. If any features of the expected pulse shape or background noise contributions change sufficiently, the firmware may need to be recompiled to update the filter. Moreover, the recursive least squares techniques can be mathematically and computationally rigorous to implement, and potentially unstable due to the continuous feedback from previous iterations. With these issues in mind, a new idea for real-time fitting was developed.

6.3   Generic Filter Design for FPGA

Given the finite availability of resources on FPGAs, a method was developed for arbitrary filter development. From [42], it can be shown that the n^{th} element of the convolution of a polynomial of order m and length L with an input signal $v[i]$ can be

written as $m+1$ serial truncated-boxcar convolutions:

$$r_L^m[n] = \sum_{i_m=0}^{n} \left(\sum_{i_{m-1}=0}^{i_m} \left(\ldots \sum_{i_0=0}^{i_1} (v[i_0] - v[i_0 - L]) - \ldots - \Lambda_{m-1} v[i_{m-1} - L] \right) - \Lambda_m \right.$$
$$\left. \times \, v[i_m - L] \right). \tag{6.2}$$

where the length scaling coefficient is defined to be

$$\Lambda_m \equiv \frac{1}{m!} \prod_{j=0}^{m-1} (L+j), m \geq 1.$$

Figure 6.5 shows a block diagram of the generic piecewise polynomial kernel. In this scheme, a recursive polynomial kernel of order m requires 1 delay of length L, $2m+1$ multiplies, $m+2$ adds, and $m+1$ accumulators.
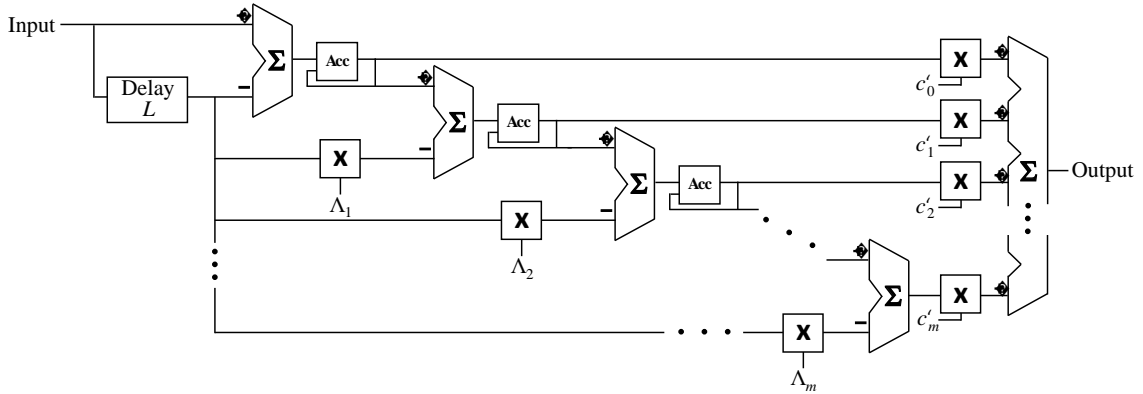
Figure 6.5: Logic diagram for the recursive implementation of a polynomial convolution of order m, requiring m accumulators and $2m+1$ multiplications.

The continuous time equivalent for each polynomial recursion is given by the Pochammer polynomials scaled by $1/m!$,

$$\frac{x^{(m)}}{m!} = \begin{cases} 1, & m = 0, \\ \frac{1}{m!} \prod_{j=0}^{m-1} (x+j), & m \geq 1, \end{cases} \tag{6.3}$$

where the n^{th}-order coefficient of the m^{th}-order Pochammer polynomial is determined as an unsigned Stirling number of the first kind, $|S_n^{(m)}|$. Figure 6.6 shows the relationship between the recursive polynomial convolutions and their continuous time counterparts.

The recursive implementation of each polynomial convolution kernel is defined for $x \geq 1$; as such, transformation from an analytically derived continuous expression to the recursion relation requires a leftward shift of the analytic expression to $x=1$,

Figure 6.6: The equivalent kernel weights for recursive polynomial convolutions of length 5 and orders 0, 1, and 2, as determined from a convolution with a delta function located at $t=1$.

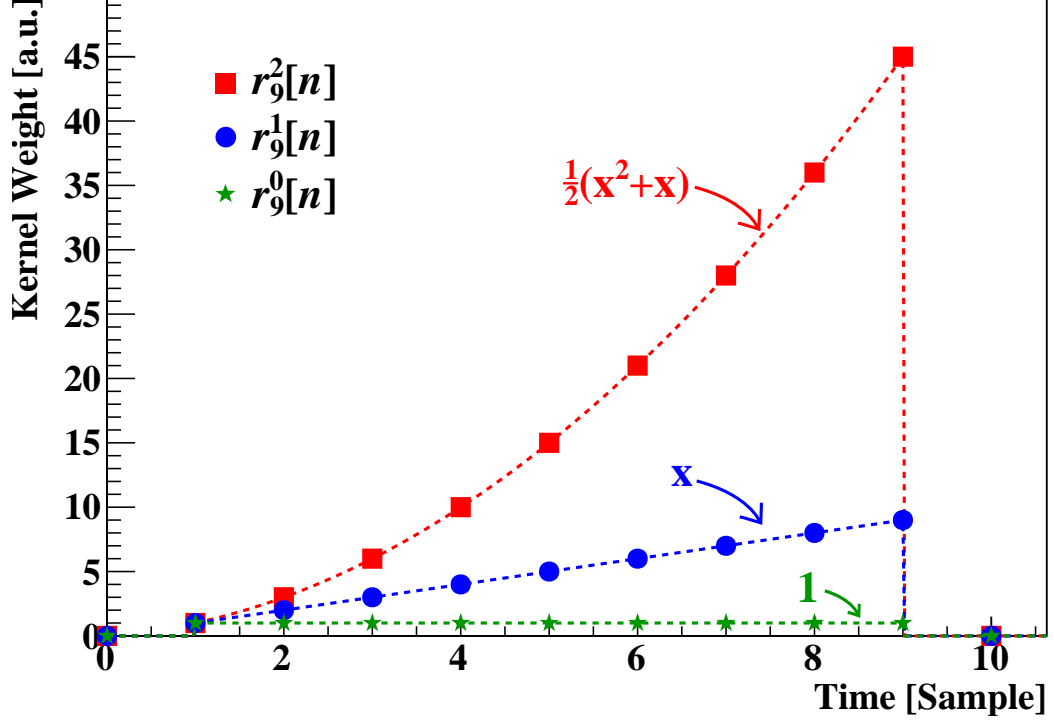followed by a delay in the recursion relation to shift the kernel to the right. Also, the polynomial coefficients must be transformed to the scaled Pochammer basis:

$$
\begin{pmatrix} c_0' \\ c_1' \\ c_2' \\ c_3' \\ \vdots \\ c_m' \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 \\ 0 & 1 & 1/2 & 1/3 & \dots & |S_1^{(m)}|/m! \\ 0 & 0 & 1/2 & 1/2 & \dots & |S_2^{(m)}|/m! \\ 0 & 0 & 0 & 1/6 & \dots & |S_3^{(m)}|/m! \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & |S_m^{(m)}|/m! \end{pmatrix}^{-1} \begin{pmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \\ \vdots \\ c_m \end{pmatrix}. \tag{6.4}
$$

In the simplest implementation of these polynomial convolution kernels, the maximum number of multiplications required for implementation goes as $2m+1$, regardless of the length L of the kernel. Therefore, this recursive implementation requires at least $L-(2m+1)$ fewer multiplications than a traditional scalar product, allowing for development of much longer filters with substantially fewer multiplies. More complex logic can be achieved by concatenating the outputs of multiple polynomial components. For a piecewise polynomial of k segments, multiplications can be combined to reduce the total number to $(k+1)(m+1)$.

A zero-area cusp-like filter using piecewise polynomial-based recursions has been developed as an example of this technique. To keep the filter as generic as possible,
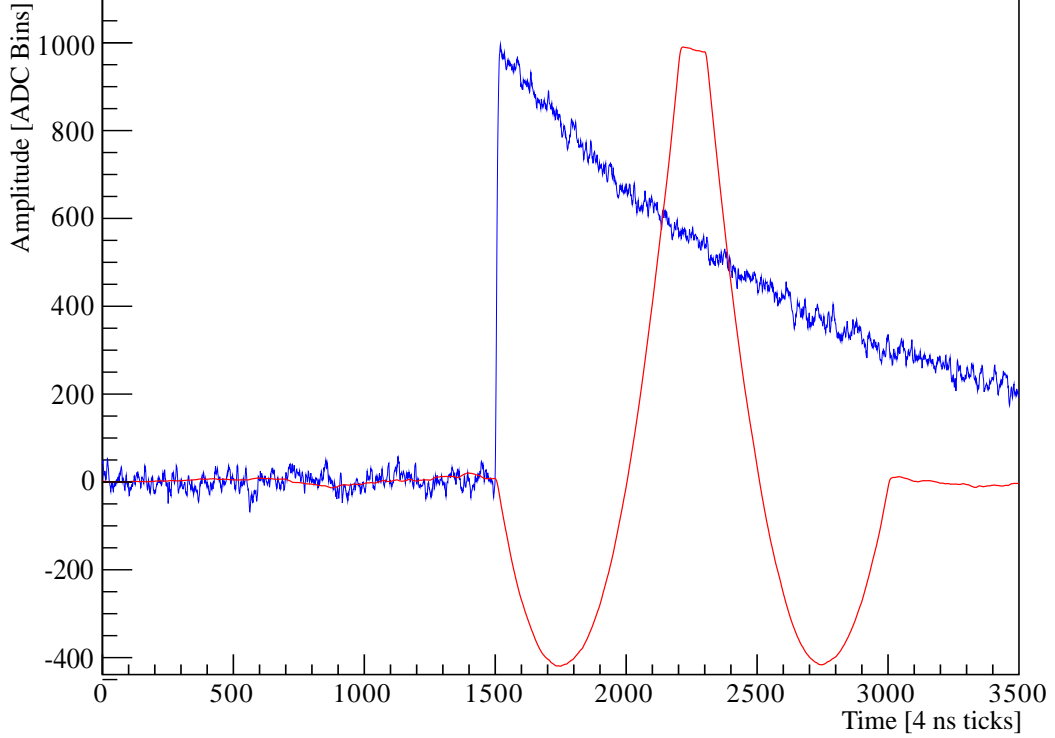
Figure 6.7: Piecewise polynomial implementation of a zero-area, cusp-like filter for tail pulses.

a superposition of the trapezoid and a double quadratic lobe was used to achieve the desired output. Initially, the integral area of the trapezoid and the lobed filters, S_{trap} and S_{lobe}, was calculated, and used to determine a scaling coefficient $a = S_{trap}/S_{lobe}$ to correctly implement the zero-area requirement. Figure 6.7 shows the filter output after convolution with a tail pulse.

Another technique for development of filters using this method is via the use of Equation 6.1. If the input signal, $f(t)$, is known, and an output, $g(t)$, is desired, then the appropriate impulse response function is found as

$$h(t) = \mathcal{F}^{-1}\left\{ \frac{\mathcal{F}(g(t))}{\mathcal{F}(f(t))} \right\}. \tag{6.5}$$

The recursive algorithm can now be developed by fitting $h(t)$ to a set of piecewise polynomials, shifting each in time to $t = 1$, and transforming the coefficients according to Equation 6.4, and adding back the appropriate time delay. Using this method, any arbitrary filter can be developed.

One important consideration for filter design using this piecewise-polynomial technique is the conversion from floating point to fixed point arithmetic. Ultimately, this leads to rounding error, depending on the precision of allocated resources. Such accumulated rounding errors can result in deviations from the true convolution, such as non-zero intergral area. A simple, yet effective, solution is to modify the 0^{th}-order coefficients of the polynomials to correct for the integral offset, and restore the zero area condition.

6.4 FPGA-Based Sliding Linear Least Squares Fitting

Using the generic polynomial filter construction method outlined in the previous section, we developed a real-time algorithm for a sliding linear least squares fit filter for FPGA. Any waveform $\mathbf{Y}$ can be decomposed as the linear sum of some number of basis functions. Mathematically, we can represent this via the product of an $n \times m$ design matrix $\mathbf{A}$, with m columns comprised of the m basis functions of sample length n, multiplied by a $m \times 1$ fit parameter column vector $\mathbf{a}$, which describes the component of, or relative contribution of each basis. The solution for the fit parameters, then, is simply

$$\mathbf{a} = \mathbf{A}^{-1}\mathbf{Y}. \tag{6.6}$$

The number of basis elements m representing template waveforms and standard noise shapes is less than the length n of the template waveforms. These basis functions thus span a subspace of waveforms representing the model space. The remaining $(m-n)$-dimensional subspace spans variations due to random noise. Thus, the matrix $\mathbf{A}$ is not invertible because the matrix is not square. Consequently, the fit parameters a_m must be determined using a generalized left-sided inverse of the design matrix, called the pseudo-inverse $\mathbf{A}^+ = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$, where $\mathbf{a} = \mathbf{A}^+\mathbf{Y}$ minimizes the residual error $\chi^2 = |\mathbf{Y} - \mathbf{A}\mathbf{a}|^2$.

One technique for calculating this pseudo-inverse is by singular value decomposition (SVD). The $n \times m$ design matrix $\mathbf{A}$ can be decomposed as the product of three matrices:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger, \tag{6.7}$$

where $\mathbf{U}$ and $\mathbf{V}$ are $n \times n$ and $m \times m$ unitary matrices, respectively ($U^\dagger U = I$ and $V^\dagger V = I$), and $\mathbf{\Sigma}$ is a diagonal $n \times m$ matrix, with diagonal elements σ_i, called the singular values of $\mathbf{A}$. $\mathbf{V}^\dagger$ is the Hermitian transpose of $\mathbf{V}$. The matrix $\mathbf{\Sigma}$ is usually also not invertible, unless $n = m$ and all singular values are nonzero, but its pseudo-inverse $\mathbf{\Sigma}^+$ can readily be obtained by calculating the reciprocal of the diagonal values and the transpose of the resultant matrix. Any degenerate singular values are ignored and kept as 0. Combining equations 6.6 and 6.7, the solution to the system of equations can thus be found as

$$\mathbf{a} = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^\dagger\mathbf{Y}. \tag{6.8}$$

Numerical Recipes in C provides an algorithm for the singular value decomposition of any arbitrary matrix, as well as an algorithm to calculate the covariance matrix $(A^\dagger A)^{-1}$ for estimation the uncertainty of fit parameters [73]. As the pseudo-code provided was written to reflect the coding style of FORTRAN, a slightly modified version of the code adapted to the index conventions of C was used instead [74].

In addition to $\mathbf{A}^+$ not being square, the length n of the template waveforms is much shorter than that of the raw data $\mathbf{Y}$. In fact, in hardware the raw data is continuous (infinitely long). Therefore the result $\mathbf{a} = \mathbf{A}^+ * \mathbf{Y}$ must be evaluated as a sliding least-squares (SLS) linear fit: one convolution per element of $\mathbf{a}$, of each row of $\mathbf{A}^+$ with $\mathbf{Y}$. In this case $\mathbf{a}(t)$ is actually a time series of vectors, one for each possible

Table 6.1: A summary of the parameters used to for the synthetic data analysis. The filter parameters were chosen to maximize the SNR, while the threshold was chosen to produce approximately an additional 0.5% of triggers strictly on noise. The average number of false triggers for the 15 amplitude settings is presented in the last column.

Filter	Parameters	Effective Threshold [ADC Bins]	False Triggers [Per 10k Traces]
SLS Filter	2 µs fit basis, uniform weighting	44	52.8
Short Trap.	80 ns rise time, 0 ns flat top	87	49.6
RC-(CR)^2	512 ns RC stage, 64 ns CR stages	95	51.9
Long Trap	1.6 µs rise time, 0.4 µs flat top	27	42.3

$t = T_0$. It is possible to calculate χ^2 as a function of T_0 using convolutions. Here we take the simpler approach of maximizing the amplitude of the template waveform $a_1(t)$ as a function of t to determine the final non-linear fit parameter T_0 and the pulse amplitude $a_1(T_0)$. Like the other filters, this methods only requires one convolution and is lightweight to implement in FPGA.

6.5 Comparison of Filter Performance with Large Amplitude Baseline Oscillations

In order to compare the performance of the different filtering algorithms, a synthetic data set was constructed based on data observed in the prototype detection system for the Nab experiment [75]. Each simulated waveform was generated as an integrated semi-Gaussian current pulse with variable rise time to mimic charge collection effects in the silicon, followed by the convolution with a CR-(RC)^2 shaper to emulate the front-end electronics. Random noise, generated from an empirically determined average power spectrum with uncorrelated phases, was then superimposed on top of the idealized pulse shape. Additionally, random-phase oscillations with frequencies ranging from 25-80 kHz and amplitudes of 150 ADC bins were added to mimic the transient microphonics noise observed in our tests. The initial rise of each pulse starts at the same time bin, and the amplitude spans from 20 to 5000 ADC bins, the anticipated range for events in the Nab experiment. Figure 6.8 shows example synthetic pulses used in this analysis.

The test data were analyzed using the least squares fit convolution described in Section 6.4, as well as a long trapezoid (ideal for energy reconstruction), a short trapezoid (used for timing), and an RC-(CR)^2 filter (also used for timing). The output of each filter on an ideal pulse is shown in figure 6.9.
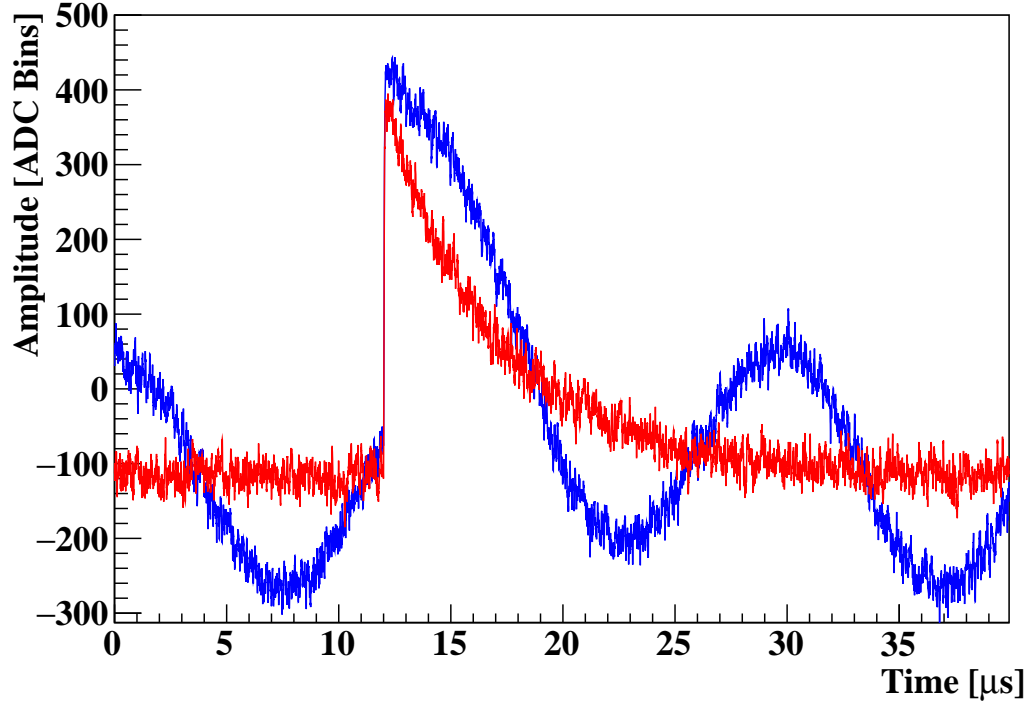
84

Figure 6.8: Sample synthetic waveform traces. Each has baseline noise with an RMS width of approximately 18 ADC bins and an offset of -125 ADC bins. The slow oscillations have an amplitude of 150 ADC bins, and vary in frequency and phase for each event. The pulse shaping characteristics reflect those of the Nab silicon detection system.

The threshold for each filter was set such that approximately 0.5% of the triggers generated were due to noise fluctuations. In the event of a trigger, the energy and time for the short trapezoid, least squares fitter, and RC-(CR)^2 were determined based on the value of the local maximum within a short window of the trigger. The long trapezoid and energy were determined as the midpoint of the flat top region, as determined by a rising and falling edge trigger through a fixed threshold. Each trigger was followed by a fixed dead-time of (5 μs to prevent redundant triggering on the same event.

Figures 6.10 and 6.11 show the extracted pulse amplitude and relative resolution for each of the different filters. As expected, the long trapezoid achieves the best results in the absence of baseline oscillations. The long shaping time effectively filters the high frequency noise, while the flat top properly integrates the charge while compensating for the exponential baseline. Furthermore, determining the energy from a fixed position on top of the trapezoid prevents a systematic bias that occurs when a local maximum is used instead. Nonetheless, the other three filters maintain linearity within ±15 ADC bins for most of the range of interest.

With the addition of slow baseline oscillations, the short trapezoid, RC-(CR)^2, and least squares fitter perform comparably to the cleaner baseline. The two timing filters are unaffected due to their short length, while the least squares filter successfully fit
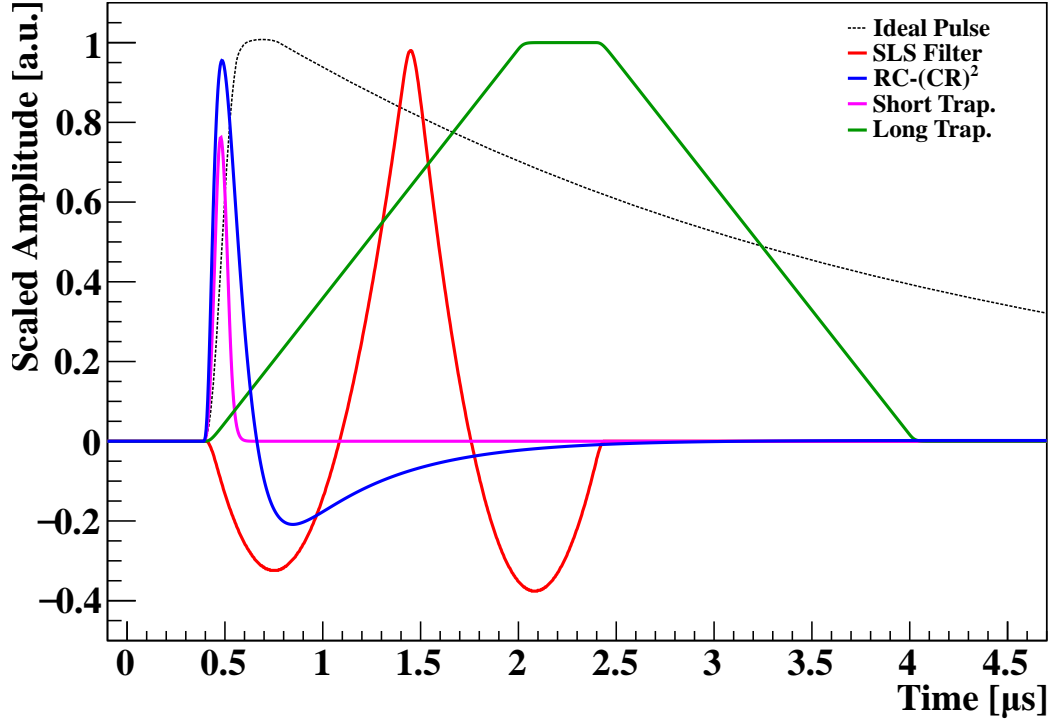
85

Figure 6.9: The ideal responses for each of the four filters used in the synthetic data comparison.

the ideal pulse shape while subtracting the oscillations. The long trapezoid performance was degraded, particularly in the low amplitude regime, as the energy pick-off was determined by the rising and falling edge triggers, which were more difficult to determine among the additional noise features.

Each filter has an inherent offset from the start of the rising edge of a pulse due to the shaping characteristics of that filter. Figure 6.12 shows the extracted T_0 for each filter, after correcting for this offset based on an ideal pulse response. A deviation in the corrected T_0 from 0 as a function of amplitude is indicative of a systematic bias in the results; as such, the ideal filter should minimize this drift. Both timing filters and the least squares fitter had small residuals, while the long trapezoid was much worse

In addition to the accurate reconstruction of the event start time, another important consideration is the jitter in the determined T_0, which can be determined from the width of the trigger time distribution. While the start time of a pulse can be greatly refined in offline analysis, jitter in the online trigger can complicate this process, since the variation of the pulse relative to the pre-trigger length necessitates additional parameterization. Figure 6.13 shows the standard deviation for each of the filters. The long trapezoid exhibits the greatest variation in trigger timing, both with and without the baseline oscillations.

Figure 6.14 shows the fraction of observed triggers as a function of amplitude for each of the filters. With a stable baseline, the long trapezoid, with its heavy averaging, is capable of the lowest threshold for high efficiency triggering. The two
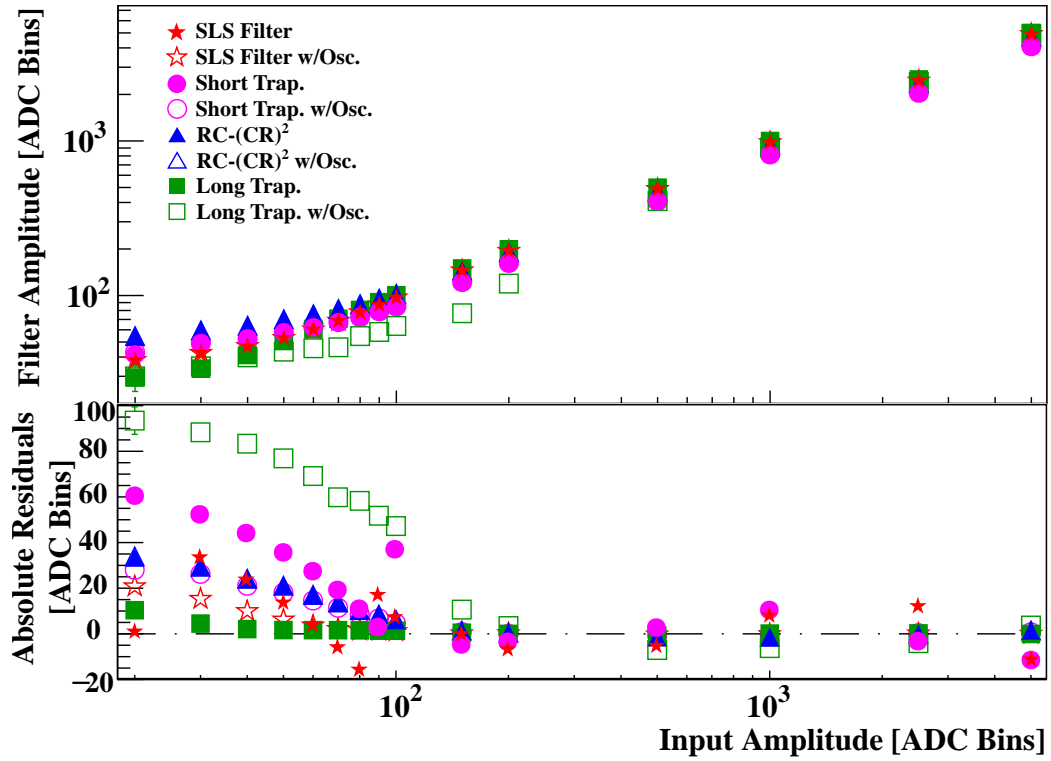
86

Figure 6.10: The measured amplitude for each of the filters with and without noise (top), and the residuals to a linear fit for input amplitudes >150 ADC bins.
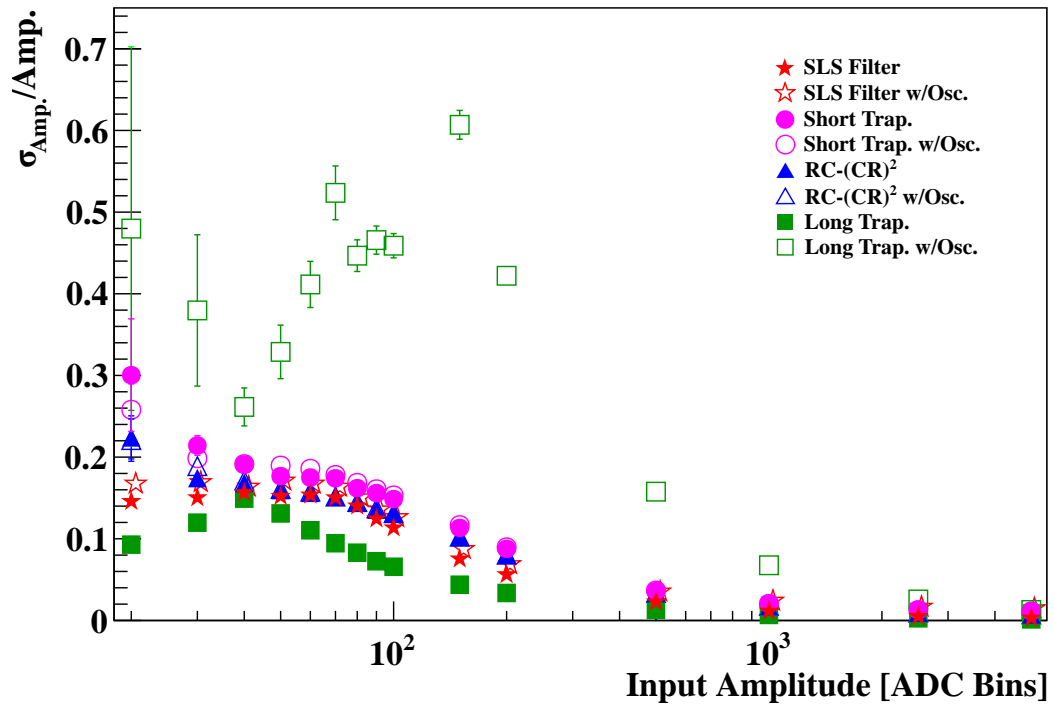

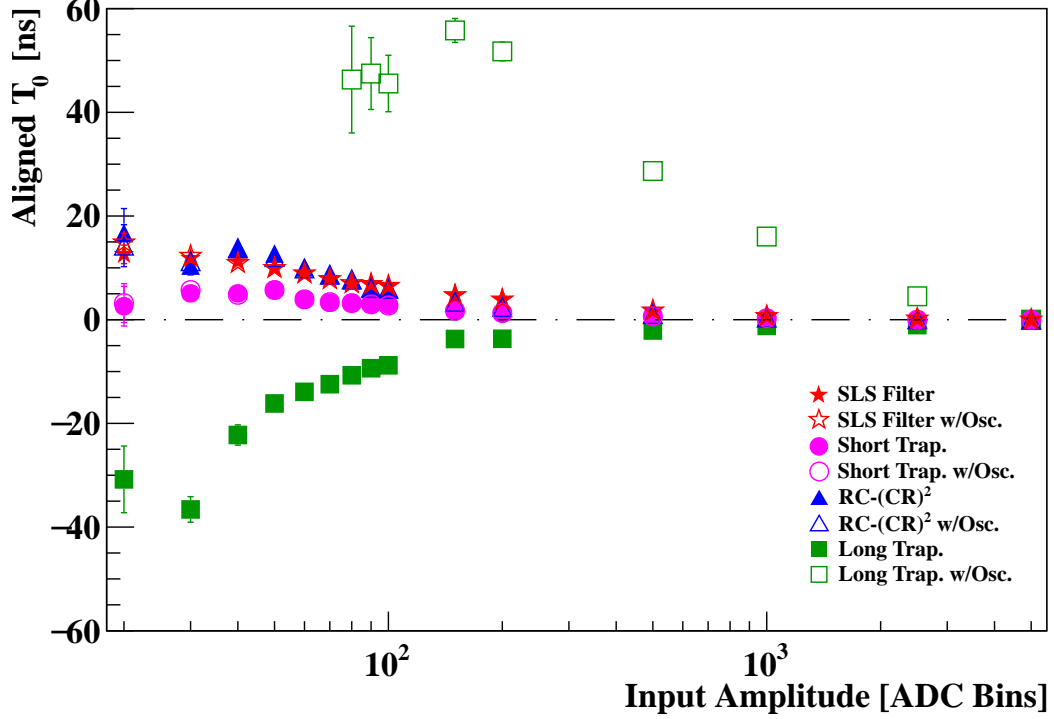Figure 6.11: The relative peak widths as determined by a Gaussian fit.

Figure 6.12: The corrected trigger time, T_0, for each of the filters. The low amplitude data for the long trapezoid with the oscillating baseline is off scale at low amplitudes.

timing filters, however, require larger amplitude pulses for efficient detection. The least squares fitter, while not as efficient as the long trapezoid, still maintains better than 99.9% efficiency for events with energies comparable to than that of a proton in the Nab experiment.

The RC-$(CR)^2$ and short trapezoid perform comparably in the presence of slow baseline oscillations, as their short shaping times integrate little of the oscillation. The long trapezoid is much more susceptible to the baseline fluctuations, and the trigger efficiency suffers dramatically, even at large amplitudes. The least squares fitter fits for the ideal pulse response while subtracting the slow oscillation, and as such preserves its low-threshold trigger efficiency even with the noise.

Ultimately, the SLS fit filter as implemented here was more balanced in timing and energy performance when compared to standard filters, especially when the system exhibited noise. However, in the ideal noise environment, traditional filters still performed better overall, which suggests that the focus moving forward should be to minimize noise in the front-end electronics. Fortunately, the piecewise polynomial filter implementation lends itself to running any filter independently on each channel without firmware modification. Therefore, if noise is present during the Nab run cycle, an SLS fit filter or similar should be implemented to actively subtract background noise for optimal energy and arrival time determination. If, on the other hand, the noise is limited to the standard pink or white spectrum as expected, then the coefficients should be such that the online filter is a more traditional cusp-like filter. In reality, each channel will entail its own unique challenges, so the filters should be
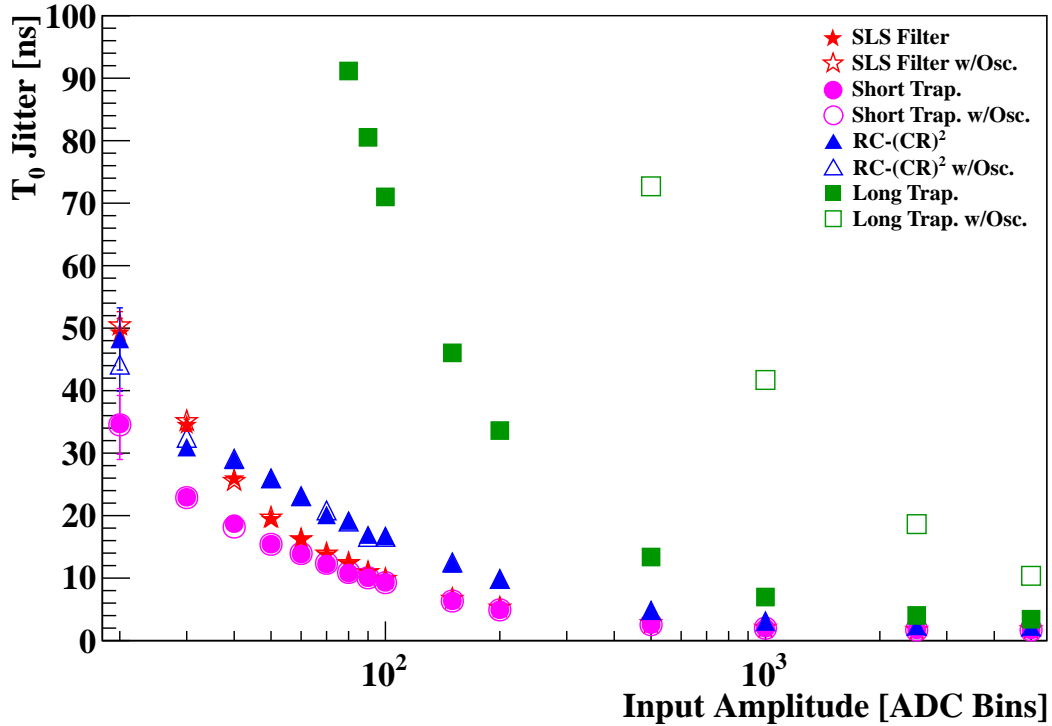
Figure 6.13: The trigger time distribution width for each of the filters. Again, the long trrapezoid data runs off-scale at lower amplitudes.

customized accordingly.

6.6   Filtering on Experimental Data

While analysis of synthetic data yields well-understood results, a true test of these filtering capabilities requires validation against empirical data. Microphonics of varying frequency were seen in both the UCNB and ^{45}Ca efforts, complicating the data analysis problem. The west detector mount was particularly problematic, where a neck-down in the cryogenic cooling line led to high frequency oscillations in the electronic signals. During the UCNB experiment, LN_2 was used for cooling, and the observed frequency was approximately 25 kHz. This prompted a switch to He gas for ^{45}, with the hopes that the single phase system would be less prone to oscillations. While the frequency of microphonics did decrease with the new cooling system, when it appeared, the frequency was about three times higher.

Another change that occurred between UCNB and ^{45}Ca was the switch to a 5 µs shaping time for the front-end electronics. Combined with the higher frequency oscillations, this actually exacerbated the situation for ^{45}Ca, as the period of an oscillation was closer in scale to the fall time of the desired signals, which caused larger distortions to the true signal. Although UCNB came first, the ^{45}Ca data set was chosen first to validate the concept of the SLS fit filter due to data accessibility and the urgency of analyzing this dataset. With the success of the filter for cleaning up the spectrum, the idea was modified to re-assess some of the UCNB data to determine
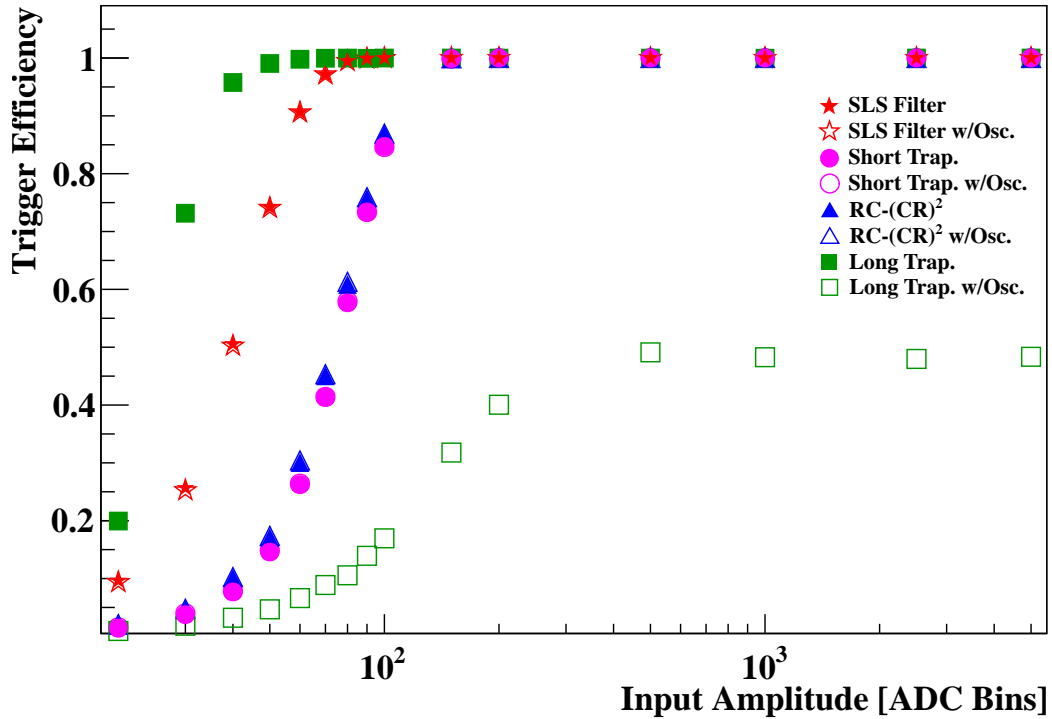
Figure 6.14: The trigger efficiency as a function of pulse amplitude for each of the filters, with and without noise. The fast timing filters and least squares fitter are immune to the baseline oscillations, but the long trapezoid trigger efficiency suffers significantly with the additional noise.

its efficacy for proton detection. What follows is a discussion and the results of these two efforts.

Peak Resolution

The first validation of the SLS fit filter was conducted on calibration source data acquired during the June 2017 ^{45}Ca run. Run 131 was taken with a ^{113}Sn source inserted through the load lock, and the ^{45}Ca source in its stationary holder both illuminating the detector. Due to the source distance from the east and west detectors, only electron events were accumulated with appreciable statistics.

The two central pixels of each detector, 64E and 64W, were compared for energy resolution, as they recorded the majority of the events. Only the 360 keV line from the decay of ^{113}Sn is used for comparison, since it is a monochromatic conversion electron emission and is isolated from all other prominent lines. Signals from the east detector showed no microphonics, as the detector mount cooling lines were larger in diameter than those on the west. As such, pixel 64E offered a much cleaner baseline in comparison to pixel 64W.

Due to its standard use for nuclear spectroscopy, a trapezoid filter with a rise time of 2 μs and a flat top of 500 ns was used as the baseline resolution for comparison. Before applying the filter, each baseline was corrected for DC offset by subtracting

the average of the samples between 3.2 and 3.6 μs from each data point. The trigger threshold was dynamically chosen as 80% of the maximum filter output within the middle 7 μs of the trace; the timestamps of the rising and falling edges through this threshold were used to determine the middle of the trapezoid. The energy of the event was determined as the value of the filter output 125 ns (25% of the flat top length) after the midpoint of the trapezoid output.

The SLS fit filter was implemented via FFT convolution. Four basis functions, comprised of 1) a template waveform formed from empirical averages, 2) a constant, 3) a linear function, and 4) a quadratic function, were used to develop the SLS filter. Each basis was 2 μs in length, and used a uniform weighting. The singular value decomposition was used to determine the pseudoinverse, after which only the vector corresponding to the template waveform function was used for analysis. The FFTW library was used for convolution of the data, and the energy was determined as the maximum value of the output. Figure 6.15 shows the spectra for pixels 64E and 64W from an offline waveform analysis, with the trapezoid output in red and the SLS fit filter in blue.

(a) Pixel 64E   (b) Pixel 64W
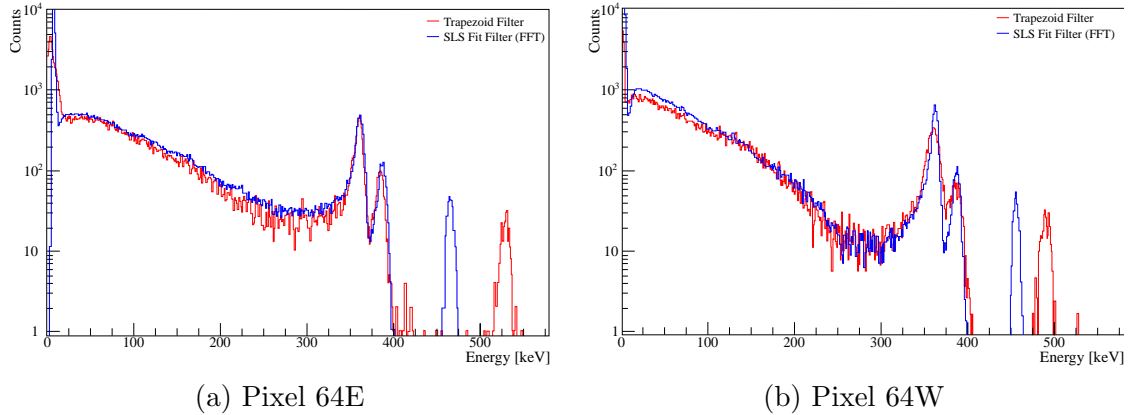
Figure 6.15: ^{113}Sn (peaks at 360 and 390 keV) and ^{45}Ca (continuous spectrum to 256 keV) spectra on the two central pixels, comparing the performance of a standard trapezoid filter against the SLS fit filter. With its relatively clean baseline, the both filters produce comparable spectra on pixel 64E, but due to microphonics, the trapezoid resolution degrades significantly on pixel 64W. The two peaks beyond 400 keV are from the bias pulser.

The peaks situated above 400 keV are amplitudes extracted from the bias pulser. The two factors that determined the amplified pulse amplitude were the pulser input and the detector capacitance; as the east and west bias pulsers were set to roughly the same value, only the detector capacitance would contribute to the observed pulse heights. While the initial design of the bias pulser circuit was assumed to inject a voltage pulse to the junction side of the detector, which would cause a larger output pulse with increased detector capacitance, the opposite was found to be true. As such, the lower-capacitance, 2.0 mm east detector produced a larger amplitude output pulse than the 1.5 mm west detector. The discrepancy in amplitude height for the two filter methods is due to the elongated rise time for bias pulser signals;

while the trapezoid fully integrated the charge to determine the pulse height, the SLS fit filter basis functions did not compensate for the slower peaking time, and thus, the apparent amplitude was an underestimate.

To assess the energy resolution capabilities of each of the two filters, the 360 keV peak was fit to a Gaussian function and the standard deviation was compared. As baseline oscillations were absent on pixel 64E, the width of each peak was comparable; using a trapezoid, the standard deviation of the peak was found as 3.65±0.04 keV, while the standard deviation with the SLS fit filter was found to be 3.72±0.05 keV. The west detector, however, exhibited baseline oscillations on multiple pixels, which led to degraded resolution; the trapezoid filtered spectrum yielded a standard deviation of 8.24±0.17 keV. Use of the SLS fit filter improved the peak resolution substantially, with a standard deviation of 3.77±0.04 keV.

While at first glance these peak widths fell short of the Nab specification of a peak resolution better than 3.0 keV, the spectra were comprised of singles triggers. Peak broadening in the spectrometer was expected for two reasons: multi-pixel depositions and energy degradation in the source foil. Roughly 10-15% of the incident electrons were expected to backscatter, at which point they would be reflected back into the detector via the magnetic mirror from the field expansion or they would deposit energy into the opposite detector array. Without logic in place to properly sum multi-hit events (or reject them to produce a single site deposition spectrum), the peak would exhibit a low energy tail and would appear to be poorer in resolution than it actually was. Additionally, because the magnetic spectrometer was capable of 4π coverage, electrons emitted in all directions were able to reach the detector. Compared to longitudinally ejected electrons that lose less than 1 keV, those with a large initial transverse momentum would necessarily have a longer average path through the source foil, resulting in a greater loss of energy before reaching the detector. This, too, broadens the peak, reducing the apparent line resolution, and thus the observed energy resolution for both pixels is likely within specification. Therefore, we recognize the efficacy of the SLS filter as a means for improving energy resolution in the presence of noise features.

Low-Threshold Efficiency

Approximately 36 hours of UCNB data from January 2016 were analyzed using the piecewise polynomial implementation of the SLS fit filter. During this time, the electronic shaping time of the front-end electronics was nominally 1 μs, the depletion bias voltage was −60 V, and liquid nitrogen was used to cool the detectors. While the data set is likely a poor candidate for a determination of the neutrino asymmetry, the underlying physics does provide an opportunity for assessing the capabilities of our new digital signal processing techniques. The data were taken with combinations of different gate-valve states: open to allow neutrons in the spectrometer, closed to measure backgrounds, high-voltage (−30 keV required for proton observation, and both AFP states to observe the electron-proton coincidence spin asymmetry. Only the west detector was capable of floating at high-voltage for these runs, so same-side coincidences were restricted to this detector.
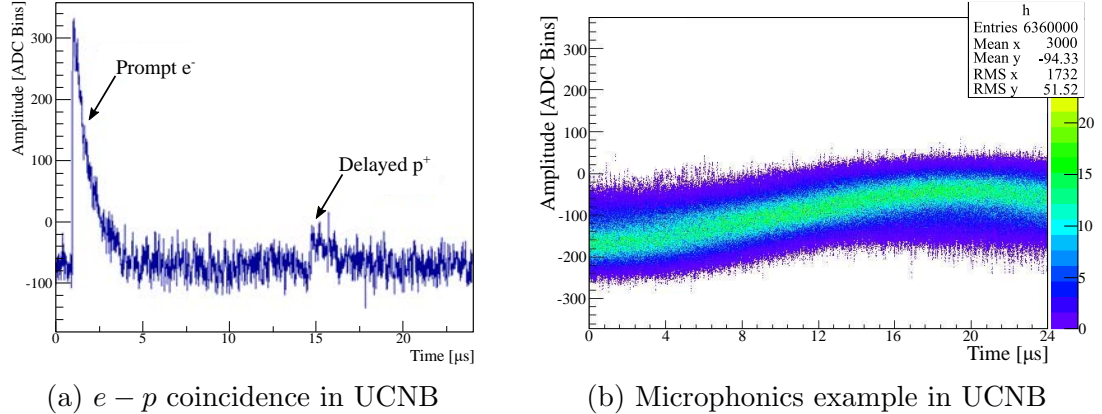
(a) $e - p$ coincidence in UCNB

(b) Microphonics example in UCNB

Figure 6.16: The challenge with analysis in UCNBwas the detection of the low-amplitude proton signals with comparable amplitude oscillations superimposed on top of them. Maintaining a sufficiently low threshold for proton detection resulted in nearly 75% of the data being comprised of noise.

The first evidence of electron-proton coincidences was observed through analysis of the online trigger files from the DAQ. For these measurements, a double trapezoid filter with a risetime of 1 μs, a flat-top of 400 ns, and an overlap of 0 was used; the fall-time for each channel was fixed at 1 μs in accordance with measured data. The threshold for each channel was independently set to provide a background trigger rate of about 0.1 Hz, although periods of large baseline oscillation would cause spikes in the trigger rates as high as 25 kHz.

The data were analyzed with the goal of demonstrating the observation of a clean proton signal, as a proof of concept of high efficiency, low-threshold triggering of the detector and data acquisition system. Four run configurations were considered to unambiguously demonstrate successful detection of the proton-electron coincidence: 1) the gate valve was opened to allow neutrons into the decay trap, and the west detector was left at ground. No same-side coincidences were expected in this configuration, as the protons did not carry sufficient energy to be detected. 2) the gate valve was closed, and the west detector was biased to -30 kV. In these runs, no same-side coincidences should be observed, as there were no neutrons decaying in the trap.

The last two types utilized 3) the gate-valve open and the accelerating potential energized, and 4) the same plus the AFP turned on. In these two configurations, same-side coincidences should be detected, although with different rates due to the neutrino asymmetry B. With the AFP off, the neutron spin was directed towards the west detector, so neutrinos would preferentially be emitted towards the east detector. Additionally, the neutron density was expected to be $\sim 25\%$ lower with the AFP on, as the spin flip provided an energy boost that increased the loss probability in the transport guides between the AFP and SCS [76]. Taking these two effects into account, more same-side $e - p$ coincident events were expected for the unflipped runs versus those with the AFP on.

For each run type, a two dimensional histogram was prepared, where the x-axis showed the current event energy, from $0 - 40$ keV, and the y-axis showed the time

difference from the preceding event, with a range of $0 - 100$ μs. Due to the time-of-flight distribution between the electron and proton, and due to the energy loss in the dead layer of the silicon, same-side coincidences would be apparent as a Gaussian distribution in energy centered around $15 - 20$ keV that decreased in intensity as the elapsed time between the two events extends beyond zero. Figure 6.17 shows the four histograms generated from the online trigger file data.

(a) Neutrons, HV on, and AFP off.

(b) Neutrons, HV on, and AFP on.

(c) Neutrons, HV off, and AFP on/off.
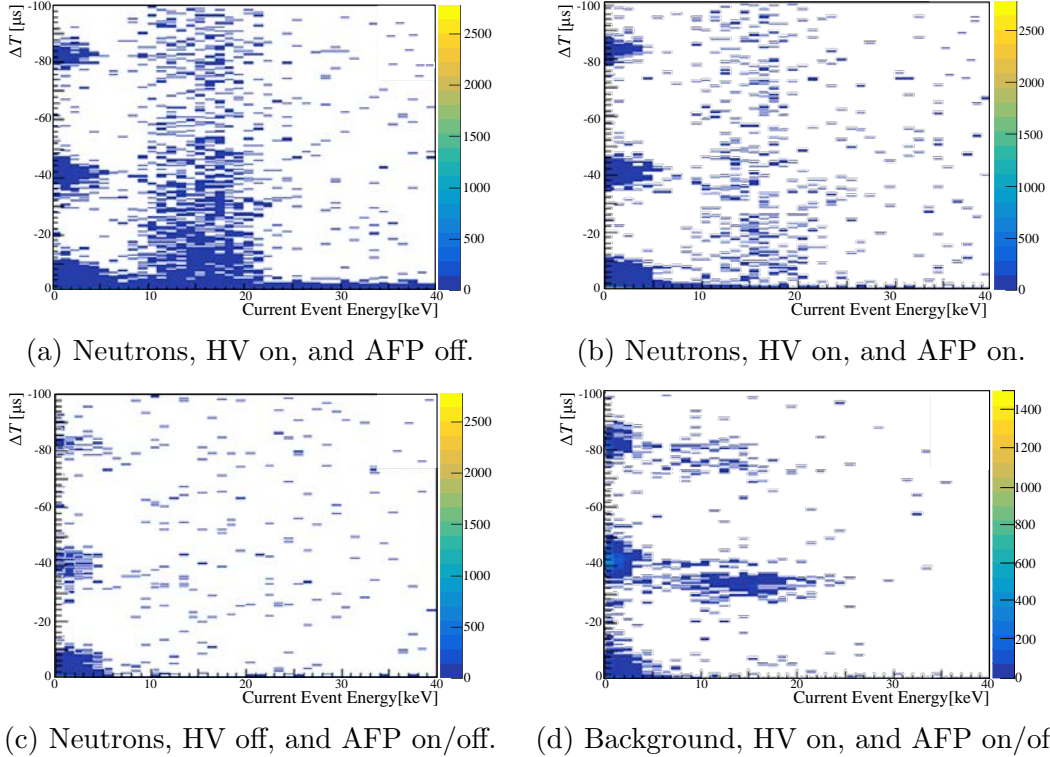
(d) Background, HV on, and AFP on/off.

Figure 6.17: Online trigger data using the double trapezoid filter from the UCNB run in January 2016. The four plots show the different run conditions meant to highlight the

As expected, a distribution showing $e-p$ coincidences was apparent in the two run configurations where they would be possible, and absent in the two run configurations where they would not appear. Moreover, more same-side coincidences were observed in the unflipped spin-state, consistent with the combination of the AFP transport efficiency and the effect of the neutrino asymmetry. Unfortunately, a significant number of noise events were recorded in clusters every $\sim$40 μs. These noise events were recognized as the microphonics due to the LN_2 cooling system; while the events with energies consistent with the pedestal ($<$7 keV) proved to be little more than a nuisance, the abundance of events overlapping the proton peak region, specifically for the background runs, was a point of concern for the integrity of the data.

An offline analysis of the waveform data was conducted to further validate the efficacy of the SLS fit filter for waveform processing for low-threshold triggering in high noise environments. The traces acquired during this time period were only 8 μs long, with approximately 4 μs of pre-trigger, so the fit window was restricted to 2 μs

to provide a full convolution without abrupt truncation at the end of the waveform trace. As with the synthetic data analysis, the basis functions were comprised of a template formed from average waveforms, and the three basis functions for a second order polynomial, as shown in Figure 6.18.
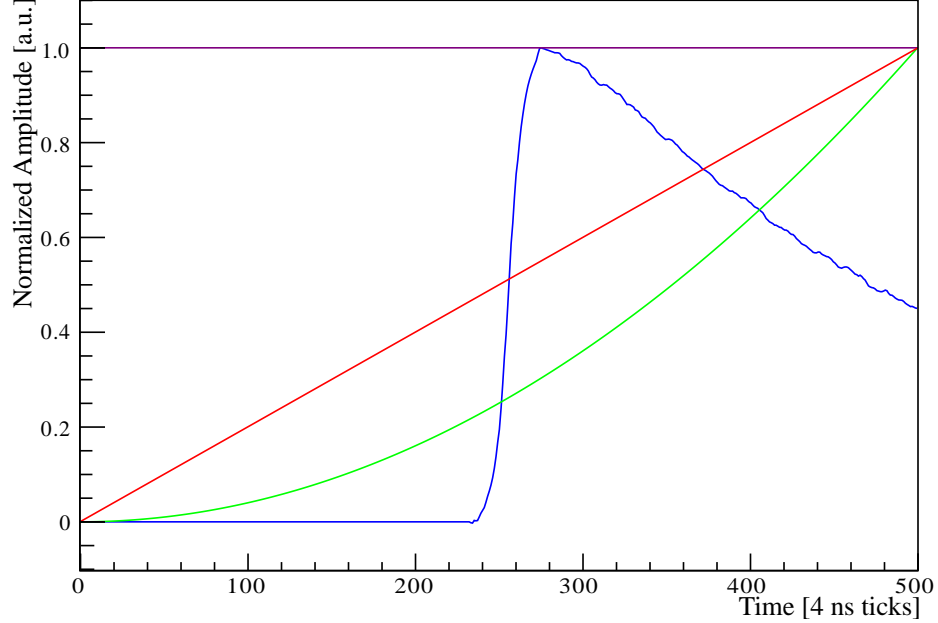

Figure 6.18: The basis functions for analysis of the UCNB data.

After calculating the pseudo-inverse of the design matrix, the vector corresponding to the template waveform basis function was fit to a set of piecewise polynomials. The vector was segmented in three locations: the local minimum at $t = 234$, the inflection point at $t = 257$, and the local maximum at $t = 276$. Each of these five segments was then fit to a quadratic equation to build up the piecewise polynomial approximation. A quadratic fit was chosen for simplicity of the algorithm, although higher order polynomials could have been used to reduce the deviation from the exact fit vector.

$$p(t) = \begin{cases} 1.33082 \cdot 10^{-7}t^2 - 1.28812 \cdot 10^{-4}t + 9.40295 \cdot 10^{-3}, & 1 \leq t < 234 \\ 4.96880 \cdot 10^{-5}t^2 - 2.37043 \cdot 10^{-2}t + 2.81339, & 234 \leq t < 257 \\ -4.42091 \cdot 10^{-5}t^2 - 2.41609 \cdot 10^{-2}t - 3.28474, & 257 \leq t < 276 \\ 2.09523 \cdot 10^{-7}t^2 - 2.72802 \cdot 10^{-4}t + 9.40295 \cdot 10^{-2}, & 276 \leq t < 501. \end{cases}$$
(6.9)

The piecewise polynomial filter was then converted to a recursive implementation, as per Section 6.3. Each segment was shifted in time to $t = 1$, and the coefficients c_i' were calculated according to Equation 6.4. The recursive polynomial form of the
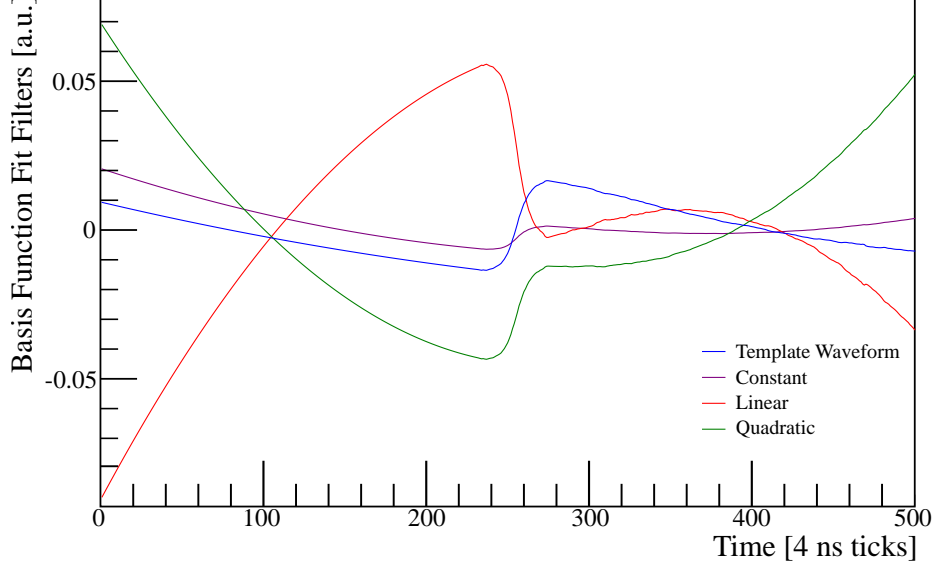
Figure 6.19: The four rows of the pseudo-inverse of the design matrix. Convolved along a waveform trace, each row determines the time-dependent amplitude of that particular basis function while actively excluding the contributions of the other basis functions.

template waveform fitter was

$$
\begin{aligned}
\text{Filter} = \quad & 9.403 \cdot 10^{-3} r_{233}^0[n] + 1.289 \cdot 10^{-4} r_{233}^1[n] + 2.662 \cdot 10^{-7} r_{233}^2[n] \\
& -1.22 \cdot 10^{-2} \cdot r_{22}^0[n - 234] - 5.99 \cdot 10^{-4} \cdot r_{22}^1[n - 234] \\
& +9.94 \cdot 10^{-5} \cdot r_{22}^2[n - 234] + 3.16 \cdot 10^{-3} \cdot r_{18}^0[n - 257] \\
& +1.57 \cdot 10^{-3} \cdot r_{18}^1[n - 257] - 8.84 \cdot 10^{-5} \cdot r_{18}^2[n - 257] \\
& +1.75 \cdot 10^{-2} \cdot r_{266}^0[n - 276] - 1.58 \cdot 10^{-4} \cdot r_{266}^1[n - 276] \\
& +4.19 \cdot 10^{-7} \cdot r_{266}^2[n - 276].
\end{aligned}
\tag{6.10}
$$

Figure 6.21 shows the result of the analysis of the UCNB data with the piecewise polynomial implementation of the SLS fit filter. The same four run configurations are shown as in Figure 6.17, although the x-axis is extended to 100 μs and the y-axis shows the timing distribution of events immediately following (positive) or immediately preceding (negative) the current event.

As expected, the SLS filter dramatically reduced the amount of triggers generated on microphonics noise, and resulted in a more clearly defined proton peak, since it actively rejected that type of background. Similarly, the features that overlapped the region of interest in the background run were completely removed. Most importantly, the observed $e - p$ coincidence rate remained unchanged between the two analyses. Based on these facts, the fit filter was shown to be effective for low-level triggering, even in the presence of large-amplitude noise.
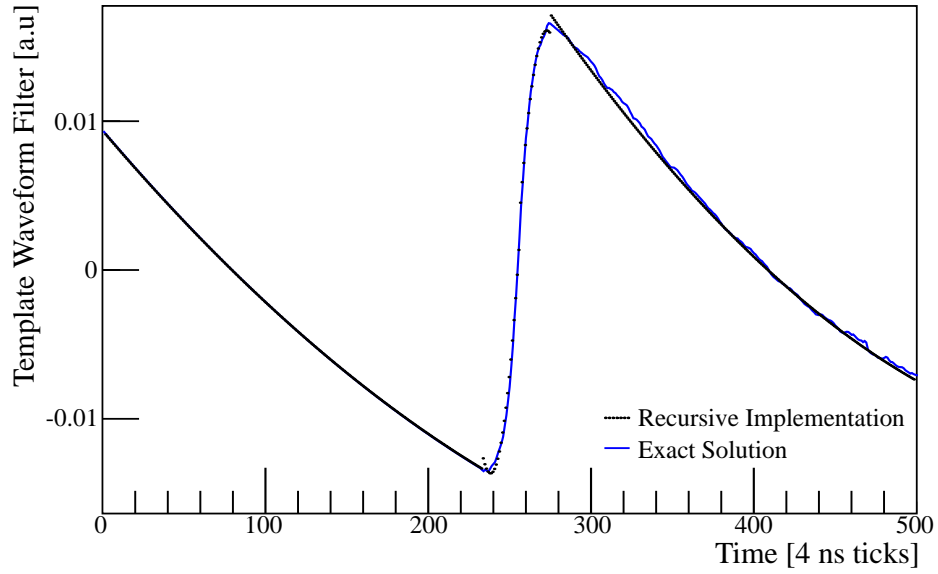
Figure 6.20: The four rows of the pseudo-inverse of the design matrix. Each row determines

(a) Neutrons, HV on, and AFP off.

(b) Neutrons, HV on, and AFP on.

(c) Neutrons, HV off, and AFP on/off.

(d) Background, HV on, and AFP on/off.

Figure 6.21: SLS fit filter implemented with piecewise polynomials on UCNB data. The number of microphonics based triggers is substantially reduced, but the overall count rate for the $e - p$ coincidences remains unchanged.

6.7 Implementation in the DAQ

The piecewise polynomial filter was implemented in the Nab DAQ as a series of four stages of second order polynomials, each of maximum length of 500 samples (4 μs). The polynomial coefficients, c'_m, and length parameters, Λ_m, for each channel and segment were passed to the FPGA at run-time, which allowed for filter adaptation without firmware recompilation. Each coefficient was defined as a signed 32-bit fixed point number, with 27 fractional bits. These values were chosen to balance throughput, so that the filter could process data at 125 MHz, and precision, to maintain the integrity of the filter and minimize rounding errors. After summing the output from each of the polynomial segments, the filter was coerced to a 14-bit number to match the input of the ADCs.

The logic for this filter was implemented in look-up tables and shift registers, and was found to consume an additional 4.7% and 4.3% of these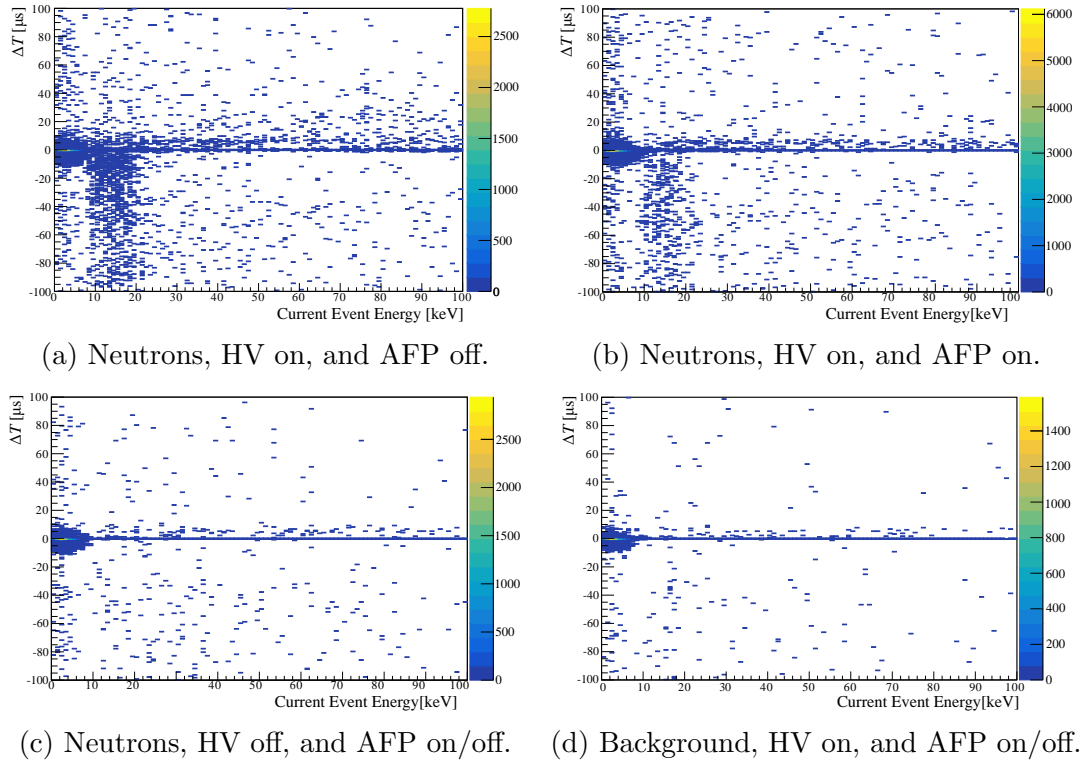 two resources, respectively. Moving forward, it will be worthwhile to investigate migrating the multiplies to the DSP slices, given their relative abundance on the FPGA, although the coefficients may need to be changed, as each multiplier is limited in width to 25 x 18 bits. Even without such a change, however, the online performance of the Nab DAQ will be improved compared to runs for the UCNB and ^{45}Ca experiments.

Chapter 7: Results and Discussion

In the effort to further comprehend our understanding of physics, neutron and nuclear beta decay experiments serve as a rich source of information. The program of correlation and spectroscopy measurements laid out by Jackson, Treiman, and Wyld in 1957 establishes an overconstrained test of the Standard Model and a means to explore the weak sector for BSM physics. These efforts are complementary to the high-energy efforts conducted at the LHC, and are vital as independent studies at the frontier of our knowledge.

UCNB, ^{45}Ca, and Nab are three experiments that are designed to improve the precision of the ratio of the axial vector coupling constant to the vector coupling constant, λ, and to further constrain exotic scalar and tensor couplings in the weak interaction. The UCNB experiment will measure the neutrino asymmetry, B, from polarized neutron decay, which touches on both Standard and BSM physics. ^{45}Ca will measure Fierz interference in the pure Gamow Teller decay via precision spectroscopy, and place a new limit on possible tensor currents. Nab will measure both the electron-neutrino correlation term, a, and Fierz interference in the mixed decay of the neutron, providing competitive results for both λ and BSM contributions.

A thick, highly pixelated silicon detector and the corresponding readout electronics were developed for the Nab experiment. Tests for efficient proton detection and precision spectroscopy were conducted in UCNB and ^{45}Ca, respectively. Additionally, the system was assessed for performance capabilities before the design was finalized for Nab fabrication. Related hardware, such as the high voltage assemblies, were also prototyped and characterized during these initial efforts, so that as Nab comes online, the spectrometer will be thoroughly understood.

A universal data acquisition architecture has been developed for precision nuclear spectroscopy. By taking advantage of the advances in technology, we were able to produce a system that maintains the last 400 ms of data for every channel, and migrates all global trigger logic to software in the host CPU. With these two features, the DAQ can be reconfigured on the fly to retrieve waveform data from any and all channels. For Nab, this means that if initial data analysis dictates a change, the turnaround time will be negligible.

Along with the DAQ architecture, a new way of performing digital signal processing was developed and tested. Rather than relying on a single filter for real-time event reconstruction, a generic piecewise polynomial filter has been developed to serve all applications. By concatenating several stages of weight polynomial recursions, any finite filter can be implemented without the need to recompile firmware. Because the filter coefficients are determined at run-time, and are unique to each channel in the DAQ, the optimal filter can be implemented depending on the noise environment for that channel.

Building upon this, we were able to develop a technique for a real-time linear least squares convolution for event reconstruction with simultaneous background rejection. Any signal can be decomposed into a set of linearly independent basis functions, for

which the coefficients can be determined via the product of the pseudo-inverse with the signal to be fit. Each row of the pseudo-inverse fits that particular basis function while also considering the contributions of the other basis functions. If the number of elements in the basis function is fewer than the number of elements in the signal, and the pseudo-inverse in convolved with the unknown waveform, the coefficients as a function of time can be determined.

While the convolution of this sliding least squares fit filter can readily be implemented via FFT convolution in CPU or GPU, a different approach is necessary for FPGA. By approximating the fit vector as a series of piecewise polynomials, a real-time fitting method can be designed and implemented uniquely for each channel. This technique allows for each channel to be customized according to its optimum response.

As a proof of concept of the power of this technique, spectrum data from ^{45}Ca was analyzed offline. In ^{45}Ca, the lines from calibration source data such as ^{113}Sn were heavily degraded by the presence of microphonics in the west detector. As such, the resolution of the resultant spectra for the east detector was significantly better. To regain the peak resolution in noisy pixels, the waveforms were analyzed offline using the sliding fit filter, where a template waveform was constructed by averaging actual digitized traces and the background was modeled with a second order polynomial. Convolution was performed using the FFT. From this analysis, the 360 keV peak resolution in pixel 64W was improved from 8.24±0.17 keV with a standard trapezoid filter to 3.77±0.04 keV, which was consistent with the 3.65±0.04 keV peak width observed in pixel 64E.

Given the success of the filter in improving energy resolution, the next step involved revisiting the UCNB data set for proton detection. While protons had been observed before in the data, the peak was obfuscated by the presence of microphonics as well. Especially problematic was the overlap of the proton signal by noise elements. Again, a design matrix was constructed from an average template waveform and a quadratic approximation to the background, and the waveforms were fit offline. Additionally, the filter was implemented via the piecewise polynomial approximation to validate its efficacy for real data analysis. As with the ^{113}Sn data, the background rejection significantly improved the observed spectra, resulting in an isolated proton peak and minimal noise contamination.

Finally, with regards to the allowed timing systematic for the Nab experiment, several efforts were developed to understand contributions from the detector and data acquisition systems. The most straightforward was verification that the DAQ was capable of providing the stability in synchronization needed to meet the budgeted 300 ps uncertainty. A minor modification of NI's `Synch` API via the addition of one more synchronization process reduced the channel-to-channel jitter between runs to 30 ps. The next step towards understanding the timing uncertainties within the experiment was the development and implementation of a bias pulser circuit. By capacitively coupling a high frequency signal to the detector bias voltage, a current pulse could be injected into every pixel simultaneously, allowing for an understanding of contributions at both the detector and amplifier circuit level. Last, an effort was made to assess the effects of charge collection on particle arrival time determination

via Monte Carlo simulation and a tagged electron timing test. The simulation showed that work must be done to prevent an unexpected bias as a function of energy and pixel location, and that the electron capture source studies with a fast scintillator can aid in validating these results.

The work presented in this dissertation is just the start of further efforts on each of these fronts, but it shows that these ideas can readily be implemented to improve system performance within these experiments. The ideas for the DAQ and filtering techniques especially show promise for future development, as system modification in software allows experimentalists greater control over the incoming data stream. These concepts will help Nab achieve its experimental goals and improve future nuclear physics efforts.

References

[1] M. Tanabashi et al. Review of particle physics. *Physical Review D*, 98:030001, 2018.

[2] G. Gamow and E. Teller. Selection rules for the β-disintegration. *Physical Review*, 49:895, 1936.

[3] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Physical Review*, 104(1):254–258, October 1956.

[4] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Physical Review*, 105(4):1413–1415, 1957.

[5] R. Garwin, L. Lederman, and M. Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Physical Review*, 105:1415, 1957.

[6] J. D. Jackson, S. B. Treiman, and H. W. Wyld, Jr. Possible tests of time reversal invariance in beta decay. *Physical Review*, 106(3):517–521, May 1957.

[7] William J. Marciano and Alberto Sirlin. Improved calculation of electroweak radiative corrections and the value of V_{ud}. *Phys. Rev. Lett.*, 96:032002, Jan 2006.

[8] N. Cabibbo. Unitary symmetry and leptonic decays. *Phys. Rev. Lett.*, 10:531–533, Jun 1963.

[9] M. Kobayashi and T. Maskawa. CP-violation in the renomalizable thoery of weak interaction. *Progress of Theoretical Physics*, 49(2):652–657, February 1973.

[10] J. C. Hardy and I.S. Towner. Superallowed $0^+ \rightarrow 0^+$ nuclear β decays 2014 critical survey, with precise results for v_{ud} and ckm unitarity. *Phys. Rev. C*, 91:025501, 2015.

[11] J. C. Hardy and I. S. Towner. v_{ud} from nuclear beta decays. *Proceedings of Science*, 291-CKM2016:028, 2016.

[12] Pieter Mumm. Resolving the neutron lifetime puzzle. *Science*, 360(6389):605–606, 2018.

[13] D. Mund, B. Märkisch, M. Deissenroth, J. Krempel, M. Schumann, H. Abele, A. Petoukhov, and T. Soldner. Determination of the weak axial vector coupling $\lambda = g_A/g_V$ from a measurement of the β-asymmetry parameter a in neutron beta decay. *Phys. Rev. Lett.*, 110:172502, Apr 2013.

[14] M. P. Mendenhall, R. W. Pattie, Y. Bagdasarova, D. B. Berguno, L. J. Broussard, R. Carr, S. Currie, X. Ding, B. W. Filippone, A. García, P. Geltenbort, K. P. Hickerson, J. Hoagland, A. T. Holley, R. Hong, T. M. Ito, A. Knecht, C.-Y. Liu, J. L. Liu, M. Makela, R. R. Mammei, J. W. Martin, D. Melconian, S. D. Moore, C. L. Morris, A. Pérez Galván, R. Picker, M. L. Pitt, B. Plaster, J. C. Ramsey, R. Rios, A. Saunders, S. J. Seestrom, E. I. Sharapov, W. E. Sondheim, E. Tatar, R. B. Vogelaar, B. VornDick, C. Wrede, A. R. Young, and B. A. Zeck. Precision measurement of the neutron β-decay asymmetry. *Phys. Rev. C*, 87:032501, Mar 2013.

[15] M. A.-P. Brown, E. B. Dees, E. Adamek, B. Allgeier, M. Blatnik, T. J. Bowles, L. J. Broussard, R. Carr, S. Clayton, C. Cude-Woods, S. Currie, X. Ding, B. W. Filippone, A. García, P. Geltenbort, S. Hasan, K. P. Hickerson, J. Hoagland, R. Hong, G. E. Hogan, A. T. Holley, T. M. Ito, A. Knecht, C.-Y. Liu, J. Liu, M. Makela, J. W. Martin, D. Melconian, M. P. Mendenhall, S. D. Moore, C. L. Morris, S. Nepal, N. Nouri, R. W. Pattie, A. Pérez Galván, D. G. Phillips, R. Picker, M. L. Pitt, B. Plaster, J. C. Ramsey, R. Rios, D. J. Salvat, A. Saunders, W. Sondheim, S. J. Seestrom, S. Sjue, S. Slutsky, X. Sun, C. Swank, G. Swift, E. Tatar, R. B. Vogelaar, B. VornDick, Z. Wang, J. Wexler, T. Womack, C. Wrede, A. R. Young, and B. A. Zeck. New result for the neutron β-asymmetry parameter A_0 from UCNA. *Phys. Rev. C*, 97:035505, Mar 2018.

[16] M. Simson, F. Ayala Guardia, S. Baeßler, M. Borg, F. Glück, W. Heil, I. Konorov, G. Konrad, R. Muñoz Horta, K.K.H. Leung, Yu. Sobolev, T. Soldner, H.-F. Wirth, and O. Zimmer. Measuring the proton spectrum in neutron decay—latest results with aspect. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2):203 – 206, 2009. Particle Physics with Slow Neutrons.

[17] F.E. Wietfeldt, J. Byrne, B. Collett, M.S. Dewey, G.L. Jones, A. Komives, A. Laptev, J.S. Nico, G. Noid, E.J. Stephenson, I. Stern, C. Trull, and B.G. Yerozolimsky. acorn: An experiment to measure the electron–antineutrino correlation in neutron decay. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 611(2):207 – 210, 2009. Particle Physics with Slow Neutrons.

[18] A. P. Serebrov, I. A. Kuznetsov, I. V. Stepanenko, A. V. Aldushchenkov, M. S. Lasakov, Yu. A. Mostovoi, B. G. Erozolimskii, M. S. Dewey, F. E. Wietfeldt, O. Zimmer, and H. Börner. Measurement of the antineutrino escape asymmetry with respect to the spin of the decaying neutron. *Journal of Experimental and Theoretical Physics*, 86(6):1074–1082, Jun 1998.

[19] M. Schumann, T. Soldner, M. Deissenroth, F. Glück, J. Krempel, M. Kreuz, B. Märkisch, D. Mund, A. Petoukhov, and H. Abele. Measurement of the neutrino asymmetry parameter b in neutron decay. *Phys. Rev. Lett.*, 99:191803, Nov 2007.

[20] G. Darius, W. A. Byron, C. R. DeAngelis, M. T. Hassan, F. E. Wietfeldt, B. Collett, G. L. Jones, M. S. Dewey, M. P. Mendenhall, J. S. Nico, H. Park, A. Komives, and E. J. Stephenson. Measurement of the electron-antineutrino angular correlation in neutron β decay. *Phys. Rev. Lett.*, 119:042502, Jul 2017.

[21] Vincenzo Cirigliano, Susan Gardner, and Barry R. Holstein. Beta decays and non-standard interactions in the lhc era. *Progress in Particle and Nuclear Physics*, 71:93 – 118, 2013. Fundamental Symmetries in the Era of the LHC.

[22] G. Soti, F. Wauters, M. Breitenfeldt, P. Finlay, P. Herzog, A. Knecht, U. Köster, I. S. Kraev, T. Porobic, P. N. Prashanth, I. S. Towner, C. Tramm, D. Zákoucký, and N. Severijns. Measurement of the β-asymmetry parameter of ^{67}Cu in search for tensor-type currents in the weak interaction. *Phys. Rev. C*, 90:035502, Sep 2014.

[23] M. G. Sternberg, R. Segel, N. D. Scielzo, G. Savard, J. A. Clark, P. F. Bertone, F. Buchinger, M. Burkey, S. Caldwell, A. Chaudhuri, J. E. Crawford, C. M. Deibel, J. Greene, S. Gulick, D. Lascar, A. F. Levand, G. Li, A. Pérez Galván, K. S. Sharma, J. Van Schelt, R. M. Yee, and B. J. Zabransky. Limit on tensor currents from ^{8}Li β decay. *Phys. Rev. Lett.*, 115:182501, Oct 2015.

[24] K. P. Hickerson et al. First direct constraints on fierz interference in free neutron β decay. *Phys. Rev. C*, 96:042501, 2017.

[25] W. S. Wilburn, V. Cirigliano, A. Klein, M. F. Makela, P. L. McGaughey, C. L. Morris, J. Ramsey, A. Salas-Bacci, A. Saunders, and A. R. Young. Measurement of the neutrino-spin correlation paramater b in neutron decay using ultracold neutrons. *Revista Mexicana De Física*, 55(2):119–122, 2009.

[26] A. Saunders, M. Makela, Y. Bagdasarova, H. O. Back, J. Boissevain, L. J. Broussard, T. J. Bowles, R. Carr, S. A. Currie, B. Filippone, A. García, P. Geltenbort, K. P. Hickerson, R. E. Hill, J. Hoagland, S. Hoedl, A. T. Holley, G. Hogan, T. M. Ito, Steve Lamoreaux, Chen-Yu Liu, J. Liu, R. R. Mammei, J. Martin, D. Melconian, M. P. Mendenhall, C. L. Morris, R. N. Mortensen, R. W. Pattie, M. Pitt, B. Plaster, J. Ramsey, R. Rios, A. Sallaska, S. J. Seestrom, E. I. Sharapov, S. Sjue, W. E. Sondheim, W. Teasdale, A. R. Young, B. VornDick, R. B. Vogelaar, Z. Wang, and Yanping Xu. Performance of the los alamos national laboratory spallation-driven solid-deuterium ultra-cold neutron source. *Review of Scientific Instruments*, 84(1):013304, 2013.

[27] T. M. Ito, E. R. Adamek, N. B. Callahan, J. H. Choi, S. M. Clayton, C. Cude-Woods, S. Currie, X. Ding, D. E. Fellers, P. Geltenbort, S. K. Lamoreaux, C.-Y. Liu, S. MacDonald, M. Makela, C. L. Morris, R. W. Pattie, J. C. Ramsey, D. J. Salvat, A. Saunders, E. I. Sharapov, S. Sjue, A. P. Sprow, Z. Tang, H. L. Weaver, W. Wei, and A. R. Young. Performance of the upgraded ultracold neutron source at los alamos national laboratory and its implication for a possible neutron electric dipole moment experiment. *Phys. Rev. C*, 97:012501, Jan 2018.

[28] T. W. Burrows. Nuclear data sheets for a=45. *Nuclear Data Sheets*, 109:171, 2008.

[29] S. Raman, T. A. Walkiewicz, and H. Behrens. Superallowed $0^+ \to 0^+$ and isospin-forbidden $j^\pi \to j^\pi$ fermi transitions. *Atomic Data and Nuclear Data Tables*, 16(5):451–494, November 1975.

[30] L. Hayen, N. Severijns, K. Bodek, D. Rozpedzik, and X. Mougeot. High precision analytical description of the allowed β spectrum shape. *Rev. Mod. Phys.*, 90:015008, Mar 2018.

[31] Leendert Hayen and Nathal Severijns. Beta Spectrum Generator: High precision allowed β spectrum shapes. 2018.

[32] A. Salas-Bacci, P.L. McGaughey, S. Baeßler, L. Broussard, M.F. Makela, J. Mirabal, R.W. Pattie, D. Počanić, S.K.L. Sjue, S.I. Penttila, W.S. Wilburn, A.R. Young, B.A. Zeck, and Z. Wang. Characterization of large area, thick, and segmented silicon detectors for neutron β-decay experiments. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 735:408 – 415, 2014.

[33] G. Bertuccio and A. Pullia. A method for the determination of the noise parameters in preamplifying systems for semiconductor radiation detectors. *Review of Scientific Instruments*, 64(11):3294–3298, 1993.

[34] E. Miyashita. Reversible polarity & high stability hv power supply: K12-r/k3-r series. Specification Datasheet, 2018.

[35] J. R. Nagel. Solving the generalized poisson equation using the finite difference method (fdm) [lecture notes], February 2012.

[36] A. Hong. Dielectric strength of air. https://hypertextbook.com/facts/2000/AliceHong.shtml, 2000.

[37] G.F. Knoll. *Radiation Detection and Measurement*. Wiley, 2000.

[38] G. Ottaviani, C. Canali, and A. Alberigi Quaranta. Charge carrier transport properties of semiconductor materials suitable for nuclear radiation. *IEEE Transactions on Nuclear Science*, NS-22:192–204, February 1975.

[39] P. Hovington, D. Drouin, and R. Gauvin. Casino: A new monte carlo code in c language for electron beam interaction - part i: Description of the program. *Scanning*, 19:1–14, 1997.

[40] Laurence W. Nagel and D.O. Pederson. Spice (simulation program with integrated circuit emphasis). Technical Report UCB/ERL M382, EECS Department, University of California, Berkeley, Apr 1973.

[41] M. Engelhardt. SPICE Differentation. *LT Journal of Analog Innovation*, 24(4):10–16, 2015.

[42] Valentin T. Jordanov and Glenn F. Knoll. Digital synthesis of pulse shapes in real time for high resolution radiation spectroscopy. *Nucl. Instrum. Methods A*, 345(2):337–345, 1994.

[43] J. L. Campbell and J. A. Maxwell. A cautionary note on the use of the Hypermet tailing function in X-ray spectrometry with Si(Li) detectors. *Nucl. Instrum. Methods B*, 129(2):297–299, 1997.

[44] B. Joly, G. Montarou, J. Lecoq, G. Bohner, M. Crouau, M. Brossard, and P.-E. Vert. An optimal filter based algorithm for PET detectors with digital sampling front-end. *IEEE Trans. Nucl. Sci.*, 57(1):63–70, 2010.

[45] A.P. Jezghani. EC Event Generation. https://gitlab.com/apjezghani/EC_event_generation, 2018.

[46] M.-M. Bé, V. Chisté, C. Dulieu, M.A. Kellett, X. Mougeot, A. Arinc, V.P. Chechev, N.K. Kuzmenko, T. Kibédi, A. Luca, and A.L. Nichols. *Table of Radionuclides*, volume 8 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, Pavillon de Breteuil, F-92310 Sèvres, France, 2016.

[47] M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 8(1):3–30, January 1998.

[48] F.G.A. Quarati, P. Dorenbos, J. van der Biezen, A. Owens, M. Selle, L. Parthier, and P. Schotanus. Scintillation and detection characteristics of high-sensitivity CeBr$_3$ gamma-ray spectrometers. *Nucl. Instrum. Methods*, A729:596–604, 2013.

[49] SensL. Biasing and readout of on semiconductor sipm sensors. White paper, 2017.

[50] J. Kodosky and E. Pérez. Introduction to G programming. Tutorial, 2008.

[51] National Instruments. How many threads does LabVIEW allocate? Knowledge Article, May 2018.

[52] A. Asenjo (National Instruments). NI LabVIEW high-performance FPGA developer's guide. Tutorial, February 2014.

[53] National Instruments. Synchronization explained. White Paper, June 2017.

[54] National Instruments. Three Steps to Using DRAM Effectively. Tutorial, February 2018.

[55] R. Ward and T. Molteno. Csce680 course notes: Table of linear feedback shift registers, October 2007.

[56] Maxim Integrated. Pseudo random number generation using linear feedback shift registers. Application Note, 2010.

[57] J. Kamleitner, S. Code, S. Gnesin, and Ph. Marmillod. Comparative analysis of digitial pulse processing methods at high count rates. *Nucl. Instrum. Methods A*, 736:88–98, 2014.

[58] M. Nakhostin. *Signal Processing for Radiation Detectors*. Wiley, Hoboken, NJ, 2018.

[59] H. Tan, M. Momayezi, A. Fallu-Labruyere, Y. X. Chu, and W. K. Warburton. A fast digital filter algorithm for gamma-ray spectroscopy with double-exponential decaying scintillators. *IEEE Trans. Nucl. Sci.*, 51(4):1541–1545, 2004.

[60] R. Abbiati, A. Geraci, and G. Ripamonti. A new filter concept yielding improved resolution and throughput in radiation detection systems. *IEEE Trans. Nucl. Sci*, 52(4):950–953, 2005.

[61] M. Agostini et al. Improvement of the energy resolution via an optimized digital signal processing in GERDA phase I. *Eur. Phys. J. C*, 75:255, 2015.

[62] A. V. Oppenheim, A. S. Willsky, and I. T. Young. *Systems and Signals*. Prentice-Hall, Englewood Cliffs, NJ, 1983.

[63] T. Fisher. mkfilter. http://www-users.cs.york.ac.uk/ fisher/mkfilter/, 1999.

[64] I.T. Young and L.J. van Vliet. Recursive implementation of the gaussian filter. *Signal Proc.*, 44(2):139–151, 1995.

[65] D. Demigny, L. Kessal, and J. Pons. *Fast Recursive Implementation of the Gaussian Filter*, pages 39–49. Springer US, Boston, MA, 2002.

[66] C. Tintori. Digital pulse processing in nuclear physics. White paper, 2011.

[67] A. Fallu-Labruyere, H. Tan, W. Hennig, and W. K. Warburton. Time resolution studies using digital constant fraction discrimination. *Nucl. Instrum. Methods A*, 579(1):247–251, 2007.

[68] A.P. Jezghani. Digital Signal Processing. https://gitlab.com/apjezghani/DigitalSignalProcessing, 2018.

[69] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proc. IEEE*, 93(2):216–231, 2005.

[70] S. Zimmermann. Active microphnoic noise cancellation in radiation detectors. *Nucl. Instrum. Methods A*, 729:404–409, 2013.

[71] G. Ripamonti and A. Geraci. Towards real-time digital pulse processing based on least-mean-squares algorithms. *Nucl. Instrum. Methods A*, 400:447=455, 1997.

[72] V. Moeller-Chan, T. Hasenohr, T. Stezelberger, M. Turqueti, and S. Zimmermann. Microphonic noise cancellation in radiation detectors using real-time adaptive modeling. In *2014 19th IEEE-NPSS Real Time Conference*, pages 1–4, May 2014.

[73] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C*. Cambridge University Press, New York, NY, 2002.

[74] M. S. Schenewerk. svdfit.c. https://www.ngs.noaa.gov/gps-toolbox/sp3intrp/svdfit.c, 2003.

[75] L. J. Broussard et al. Detection system for neutron β decay correlations in the UCNB and Nab experiments. *Nucl. Instrum. Methods A*, 849:83 – 93, 2017.

[76] B. Plaster, R. Rios, H. O. Back, T. J. Bowles, L. J. Broussard, R. Carr, S. Clayton, S. Currie, B. W. Filippone, A. García, P. Geltenbort, K. P. Hickerson, J. Hoagland, G. E. Hogan, B. Hona, A. T. Holley, T. M. Ito, C.-Y. Liu, J. Liu, M. Makela, R. R. Mammei, J. W. Martin, D. Melconian, M. P. Mendenhall, C. L. Morris, R. Mortensen, R. W. Pattie, A. Pérez Galván, M. L. Pitt, J. C. Ramsey, R. Russell, A. Saunders, R. Schmid, S. J. Seestrom, S. Sjue, W. E. Sondheim, E. Tatar, B. Tipton, R. B. Vogelaar, B. VornDick, C. Wrede, Y. P. Xu, H. Yan, A. R. Young, and J. Yuan. Measurement of the neutron β-asymmetry parameter A_0 with ultracold neutrons. *Phys. Rev. C*, 86:055501, Nov 2012.

Vita

NAME	Aaron Paul (Sprow) Jezghani	
BIRTH PLACE	Frederick, MD	
EDUCATION	University of Kentucky, Lexington, KY *Masters of Science, Physics*	May 2018
	SUNY College at Buffalo, Buffalo, NY *Bachelor of Science, Physics*	May 2012
	Grove City College, Grove City, PA *Bachelor of Science, Mathematics*	May 2012

PUBLICATIONS C. C. Haddock et al. A Search for Possible Long Range Spin Dependent Interactions of the Neutron from Exotic Vector Boson Exchange. *Phys. Lett. B 783* (2018) 227-233.

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. *Phys. Rev. C 97* (2018) 012501(R).

L. J. Broussard et al. Detection System for Neutron β Decay Correlations in the UCNB and Nab Experiments. *Nucl. Instrum. Meth. A 849* (2017) 83-93.

Z. Wang et al. A multilayer surface detector for ultracold neutrons. *Nucl. Instrum. Meth. A 798* (2015) 30-35.

R. C. Rai, A. Delmont, A. Sprow, B. Cai, and M. L. Nakarmi. Spin-charge-orbital coupling in multiferroic LuFe2O4 thin films. *Appl. Phys. Lett. 100* (2012) 212904.